

Timed Relational Abstractions For Sampled Data Control Systems

Aditya Zutshi¹, Sriram Sankaranarayanan¹ and Ashish Tiwari². *

1. University of Colorado, Boulder, CO. {aditya.zutshi, srirams}@colorado.edu

2. SRI International, Menlo Park, CA. ashish.tiwari@sri.com

Abstract. In this paper, we define timed relational abstractions for verifying sampled data control systems. Sampled data control systems consist of a plant, modeled as a hybrid system and a synchronous controller, modeled as a discrete transition system. The controller computes control inputs and/or sends control events to the plant based on the periodically sampled state of the plant. The correctness of the system depends on the controller design as well as an appropriate choice of the controller sampling period.

Our approach constructs a timed relational abstraction of the hybrid plant by replacing the continuous plant dynamics by relations. These relations map a state of the plant to states reachable within the sampling time period. We present techniques for building timed relational abstractions, while taking care of discrete transitions that can be taken by the plant between samples. The resulting abstractions are better suited for the verification of sampled data control systems. The abstractions focus on the states that can be observed by the controller at the sample times, while abstracting away behaviors between sample times conservatively. The resulting abstractions are discrete, infinite-state transition systems. Thus conventional verification tools can be used to verify safety properties of these abstractions. We use k-induction to prove safety properties and bounded model checking (BMC) to find potential falsifications. We present our idea, its implementation and results on many benchmark examples.

1 Introduction

We present techniques for verifying safety properties of sampled data control systems using timed relational abstractions. Sampled data control systems consist of a discrete controller that periodically senses the state of a continuous physical plant, and actuates by setting inputs or sending control commands to the plant. Sampled data control systems are quite common in practice. Complex (network) control systems often involve many control tasks that are scheduled periodically, with each task controlling a different aspect of the plant. The cadencing of these tasks to enforce the safety and stability of the system is an important problem. The choice of sampling period is crucial: a small

* Zutshi and Sankaranarayanan's work supported by NSF Career grant CNS-0953941. Tiwari's work supported in part by DARPA under subcontract No. VA-DSR 21806-S4 under prime contract No. FA8650-10-C-7075, and NSF grants CSR-0917398 and SHF:CSR-1017483. All opinions expressed are those of the authors and not necessarily of the US NSF or DARPA.

sampling time can place infeasible constraints on the scheduling policy, whereas large sampling times can cause instabilities or safety violations.

In this paper, we consider a simple and natural model of a sampled data control system. The controller is modeled by an infinite state (linear) transition system. It communicates synchronously with a plant modeled by an affine hybrid automaton. The controller runs periodically with a fixed sampling period $T_s > 0$. At each time period, the controller senses the state of the plant and performs controller actions that may include (a) setting control input signals for the plant, and (b) “commanding” the plant to execute a controlled discrete transition, resulting in an instantaneous jump and a mode change in the plant.

Our verification approach uses the idea of timed relationalization, extending the idea of untimed relational abstractions [35]. A timed relational abstraction considers the set of states of the plant that are potentially observable by the controller at the sample times, while safely abstracting away all the intermediate states. To this end, we build relations that map a state of the plant at the beginning of a sampling period to states that can be reached at the end. Using these relations, the entire plant can be safely abstracted away by a discrete, infinite-state transition system. This system is composed together with the controller to yield an overall discrete system that can be analyzed by existing tools such as k -induction [36], bounded model-checking [5] and abstract interpretation [10, 22], while exploiting advances in abstract domains, SAT and SMT solvers.

There are two key challenges in constructing the timed abstraction: (a) dealing with the continuous dynamics inside a mode, and (b) handling autonomous transitions that can be taken by the plant between two sampling periods. For systems with affine dynamics the former problem is solved by computing a matrix exponential for the matrices defining the dynamics in each mode [27]. However, the numerical exponential computation is potentially unsound due to round-off and truncation errors. Likewise, solving for autonomous transitions involves computing a symbolic matrix exponential to deal with the unknown switching times for each transition. To solve both problems, we exploit advances in interval arithmetic to compute guaranteed enclosures to the matrix exponential [28, 29, 18, 6]. This yields interval linear relations. We then use a *template*-based mechanism using SMT solvers to abstract the resulting interval linear relations in terms of relations expressible in linear arithmetic.

We have implemented our approach to relationalization and present an extensive evaluation over a set of benchmark systems. Our evaluation performs a relational abstraction of the plant using the techniques described in this paper. The resulting abstraction is analyzed using the SAL tool-set from SRI [32, 40]. The results of our evaluation are quite promising: we show that our techniques can successfully handle complex sampled data control systems efficiently and soundly. We compare our approach with the SpaceEx tool that implements symbolic model checking for affine hybrid automata using support functions [16]. Our approach compares quite favorably with the various options available in SpaceEx: providing safety proofs in many cases where SpaceEx fails to prove safety over finite time horizons. On the other hand, we note that the finite time horizon bounds on the reachable state-space established by SpaceEx can sometimes be used as strengthenings for k -induction in our approach to obtain significant

speedups. Our implementation, the data from our experiments along with an extended version of the paper will be available online ¹.

Motivating Examples: We discuss two simple motivating examples that clearly illustrate the need for verification of sampled data control systems.

Consider a proportional-integral (PI) controller defined by $u' := -30x - y$ composed with a plant defined by $\dot{x} = 5x + u$, $\dot{y} = x$. With a period $T_s = 0.1s$, the controller is able to stabilize the plant but fails to do so with a period $T_s = 0.5s$.

Consider an inverted pendulum controller:

$$u' := \begin{cases} -16, & y \geq 2 \vee 16x - y \leq -10 \\ 16, & y \leq -2 \vee 16x - y \geq 10 \\ u, & \text{otherwise} \end{cases} .$$

The linearized plant has the dynamics $\dot{x} = y$ $\dot{y} = 20x + 16y + 4u$. If the controller is implemented in the continuous domain, it results in a stable system. However, a digital implementation, regardless of the sampling period, is unable to stabilize the pendulum.

We now discuss the related work.

Relational Abstractions: Relational abstractions have been used primarily for checking liveness properties [4, 30]. There are many subtle distinctions between the various forms of relational abstractions used. Transition invariants [30], used in termination proofs, relate the current state to *any* previous state at a given program location. Likewise, progress invariants relate the current state and the *immediately* previous state at a given location [20]. Podelski and Wagner provide a verification procedure for (region) stability properties of hybrid systems [31], wherein they derive *binary reachability relations* over trajectories of a hybrid system, similar in spirit to a relational abstraction. Note that Podelski and Wagner use a hybrid system reachability tool to compute their abstractions in the first place. The techniques in this paper and our previous work [35] are meant to solve the reachability problem using these relations.

Our previous work explored the idea of abstracting the dynamics inside each discrete mode of a hybrid automaton by an *untimed relational abstraction* [35]. The relational abstractions presented here capture the relationship between the current state at time $t = t_0$ and any state reachable at time $t = t_0 + T_s$ units. The consideration of the sampling time T_s is essential for verifying sampled data control systems. Furthermore, it is also important to note that unlike untimed relational abstraction, it is essential for the abstraction presented here to account for the plant's discrete transitions that can be taken in the time interval $t \in [t_0, t_0 + T_s]$.

Abstractions of Hybrid Systems: *Discrete abstractions* have been widely studied and applied for verifying safety properties of hybrid systems [1, 39]. The use of counterexample-guided abstraction refinement has also been investigated in the past [1, 8]. In this paper, the proposed abstraction yields a discrete but infinite state system. Hybridization is a technique for converting nonlinear systems into affine hybrid systems by subdividing the invariant region into numerous subregions and approximating

¹ <http://systems.cs.colorado.edu/research/cyberphysical/relational>

the dynamics as a hybrid system by means of a linear differential inclusion in each region [23, 3, 11]. However, such a subdivision is expensive as the number of dimensions increases and often infeasible if the invariant region is unbounded.

Flowpipe Construction: Reasoning about the reachable set of states for flows of affine hybrid systems through flowpipe approximation has long been dominant approach for checking safety properties [37], using various representations for sets of states including ellipsoids [25], zonotopes [17], template polyhedra [33], and support functions [19]. The tool SpaceEx implements numerous improvements to these techniques with impressive performance on some benchmarks with a large number of system variables [16].

Synchronous Systems: Techniques for verifying synchronous system models, with piecewise constant continuous dynamics, have been studied in the past, notably by Halbwachs et al. [22] and as part of the NBAC tool by Jeannet et al. [24]. Our work considers a synchronous controller with affine hybrid plants. Furthermore, we consider the idea of an up front relationalization of the plant dynamics, enabling a verification procedure to focus purely on discrete systems.

2 Sampled Data Control Systems

Let \mathbb{R} denote the set of real numbers. We use $\mathbf{a}, \dots, \mathbf{z}$ with subscripts to denote (column) vectors and A, \dots, Z to denote matrices. For a $m \times n$ matrix A , the row vector A_i , for $1 \leq i \leq m$, denotes the i^{th} row.

We discuss models for sampled data control systems. Figure 1 shows the schematic diagram for such a control system consisting of a discrete controller communicating with a hybrid plant. The controller has a time period $T_s > 0$. Every T_s time units, the controller “senses” the state of the hybrid plant and synchronizes to change the mode of the plant. The commands can take the form of (a) events enabling a discrete transition of the plant, or (b) values for control inputs that are assumed to be held constant throughout the sample time period. We model the controller as a discrete transition system [26].




Fig. 1. Schematic for a Sampled Data Control System.

Definition 1 (Discrete Transition System). A discrete transition system Π is a tuple $\langle L, \mathbf{x}, \mathcal{T}, \ell_0, \Theta \rangle$ wherein, L is a finite set of discrete locations; $\mathbf{x} : (x_1, \dots, x_n)$ is a set of variables with variable x_i of type $\text{type}(x_i)$; \mathcal{T} is a set of discrete transitions; $\ell_0 \in L$ is the initial location; and $\Theta[\mathbf{x}]$ is an assertion capturing the initial values for \mathbf{x} .

Each transition $\tau \in \mathcal{T}$ is of the form $\langle \ell, m, \rho_\tau \rangle$, wherein ℓ is the pre-location of the transition and m is the post-location. The relation $\rho_\tau[\mathbf{x}, \mathbf{x}']$ represents the transition relation over current state variables \mathbf{x} and next state variables \mathbf{x}' .

We now discuss the overall model for the plant as a hybrid automaton with controlled and uncontrolled transitions.

Definition 2 (Plant Model). A plant \mathcal{P} is an extended hybrid automaton described by a tuple $\langle \mathbf{x}, \mathbf{u}, Q, \mathcal{F}, \mathcal{X}, \mathcal{T}, q_0, X_0 \rangle$, wherein,

- $\mathbf{x} : (x_1, \dots, x_n)$ denotes the continuous state variables, and $\mathbf{u} : (u_1, \dots, u_m)$ the control inputs,
- Q is a finite set of discrete modes. $q_0 \in Q$ is the initial mode and X_0 the initial set of states.
- \mathcal{F} maps each discrete mode $q \in Q$ to an ODE $\frac{d\mathbf{x}}{dt} = F_q(\mathbf{x}, \mathbf{u}, t)$.
- \mathcal{X} maps each discrete mode $q \in Q$ to a mode invariant $\mathcal{X}(q) \subseteq \mathbb{R}^n$.
- \mathcal{T} represents a set of discrete transitions. Each transition $\tau \in \mathcal{T}$ is a tuple $\langle s, t, \gamma, U \rangle$ wherein s, t represent the pre- and post- mode respectively. $\gamma[\mathbf{x}]$ is the transition guard assertion, and U maps each variable $x_i \in \mathbf{x}$ to an update function $U_i(\mathbf{x})$. The transition relation for τ is defined as $\rho_\tau(\mathbf{x}, \mathbf{x}') : \gamma(\mathbf{x}) \wedge \mathbf{x}' = U(\mathbf{x})$.
- We partition the transitions in \mathcal{T} as autonomous transitions \mathcal{T}_{aut} and controlled transitions \mathcal{T}_{ctrl} . Autonomous transitions can be taken by the plant non-deterministically, whenever enabled. On the other hand, controlled transitions are taken upon an explicit command by the controller.

The state of the plant is a tuple $(q, \mathbf{x}, \mathbf{u})$ consisting of the current mode q , state values \mathbf{x} and controller input \mathbf{u} . Note that the control input is set at the beginning of a time period, and is assumed to remain constant throughout the period.

The overall sampled data control system is a tuple $\langle \mathcal{C}, \mathcal{P}, \mu, T_s \rangle$ of a discrete controller transition system \mathcal{C} , a hybrid plant model \mathcal{P} and a mapping μ from variables in \mathcal{C} to control inputs \mathbf{u} of \mathcal{P} . A given sampling time $T_s > 0$ specifies the periodicity of the controller execution. The state of the system is represented by the joint state of the plant σ_P and σ_C of the controller. We assume that the computations of the controller take zero (or negligible time) compared to the sampling period. Furthermore, we assume that the commands issued by the controller are in the form of an input \mathbf{u} for the next time period, and/or a command to execute a discrete transition. Finally, to avoid considering improbable “race conditions”, we assume that the plant itself may not execute autonomous discrete transitions at sample time instances when the controller executes². The overall system evolves in one of two ways:

1. At sample times $t = nT_s$ for $n \in \mathbb{Z}$, a controlled transition is taken based on the current state of the plant and the controller. The transition updates the controller state, the values of the plant inputs and can also command the plant to execute a discrete transition out of its current mode.
2. Between two sample times $t \in (nT_s, (n+1)T_s)$, the state of the plant evolves according to its current mode q , continuous variables \mathbf{x} and input \mathbf{u} . If an autonomous transition $\tau \in \mathcal{T}_{aut}$ is enabled, then it may be non-deterministically executed by the plant, possibly changing the plant’s state instantaneously.

² This assumption can be relaxed to allow such simultaneous executions, provided the plant and the controller do not attempt to update the same state variable.

A plant is *affine* iff (a) for each discrete mode q , the dynamics are of the form $\frac{dx}{dt} = A_q x + B_q u + b_q$, (b) the initial condition Θ and guards γ_τ for each transition τ , are linear arithmetic formulae, and (c) the update functions U_τ are affine.

Why Autonomous Transitions? The ability to model autonomous transitions is quite important in practice. Real life plants are often multi modal with mode changes that can be effected by exogenous user inputs, disturbance inputs, system failures and other exceptional situations. Examples include pump failures or occlusion events observed in models of drug administration using infusion pumps [2, 34], failsafe models of spacecraft control systems, wherein exogenous disturbance inputs can cause mode changes [7]. Another reason for autonomous transitions includes the modeling of actuation delays. Autonomous transitions can be used to model delays between controller commands and their actuation in networked control systems.

3 Relationalization

In this section, we discuss the notion of timed relationalizations for plants in a sampled data control system. The basic idea behind relationalization is to build a relation $R_P(q', x', q, x, u)$ of all possible pairs of states (q', x') and (q, x) such that (a) the plant is in the state (q, x) at the start time $t = t_0$, (b) it reaches the state (q', x') at time $t = t_0 + T_s$, and (c) u is the constant controller input for $t \in (t_0, t_0 + T_s]$. Note that the discrete modes q, q' may be different, depending on whether an autonomous transition is taken by the plant between two samplings.

Let us suppose a relation R_P can be built that can characterize all pairs (q', x') that a controller can observe at the next time step, given that (q, x) was observed at the current time step and u was the control input. As a result, we may construct a purely discrete abstraction of the sampled data control system wherein the behavior of the plant between two samplings is entirely captured by R_P . Therefore, the resulting discrete transition system can be verified using a host of approaches for verification of discrete programs. Furthermore, since our goal is to perform safety verification, we do not need to compute the exact relation R_P , but only an over-approximation of it.

We will now describe techniques for constructing timed relational abstractions.

Example 1. Consider the hybrid plant model shown in Figure 2 with two state variables x, y and no control inputs. The matrices defining the dynamics are

$$A_1 : \begin{pmatrix} -1.5 & 1.2 \\ 1.3 & 0.2 \end{pmatrix} \quad b_1 : \begin{pmatrix} 1.0 \\ -0.5 \end{pmatrix} \quad A_2 : \begin{pmatrix} 2 & 1.2 \\ 0.1 & -3.6 \end{pmatrix} \quad b_2 : \begin{pmatrix} -0.6 \\ -0.6 \end{pmatrix}.$$

There are two modes n_0 and n_1 with an autonomous transition τ_1 from n_0 to n_1 and a controlled transition τ_2 from n_1 back to n_0 . Relationalization of this automaton will need to consider 3 cases: (a) the automaton remains entirely inside the mode n_0 during a sample interval, (b) the automaton remains entirely inside mode n_1 and (c) the automaton switches from mode n_0 to n_1 sometime during a sampling interval.

We first discuss relational abstraction for the case when the plant remains in some mode q during a sampling period without any autonomous transitions taken in between.




Fig. 2. A simple affine hybrid automaton with an autonomous transition τ_1 and controlled transition τ_2 . Some sample trajectories of the automaton are shown with the autonomous transition being taken. The red colored trajectories belong to mode n_0 and the blue colored trajectories to mode n_1 . The guard set is shown in thick lines.

This situation is abstracted by a relation $R_q(\mathbf{x}, \mathbf{u}, \mathbf{x}')$ that relates all plant states (q, \mathbf{x}) at some time t and state (q, \mathbf{x}') at time $t + T_s$ with control input \mathbf{u} . The resulting R_q for each $q \in Q$ will form a disjunct in the overall relation R_P for the plant.

Definition 3 (Timed Relational Abstraction). Consider a continuous system specified by a time invariant ODE $\frac{d\mathbf{x}}{dt} = f(\mathbf{x}, \mathbf{u})$ for $\mathbf{x} \in X$ and control inputs $\mathbf{u} \in U$.

A relation $R(\mathbf{x}, \mathbf{u}, \mathbf{x}')$ is a timed relational abstraction with sample time T_s of the continuous system iff for all time trajectories $\mathbf{x}(t)$ of the ODE with constant control input $\mathbf{u}(t) = \mathbf{u}$, we have $(\mathbf{x}(0), \mathbf{u}, \mathbf{x}(T_s)) \in R$.

Since we assume that the dynamics are time invariant, the starting time of the observation can be arbitrarily set to $t_0 = 0$. Time varying dynamics can be treated by lifting this assumption and specifying the value of the time t as part of the state \mathbf{x} of the system.

We now consider the timed relational abstraction for a system with affine dynamics given by $\frac{d\mathbf{x}}{dt} = A\mathbf{x} + B\mathbf{u} + \mathbf{b}$. We note that the solution of the ODE can be written as $\mathbf{x}(t) = e^{tA}\mathbf{x}(0) + \int_{s=0}^t e^{(t-s)A}(B\mathbf{u}(s) + \mathbf{b}) ds$. If the matrix A is invertible and $\mathbf{u}(s) = \mathbf{u}$ for $s \in [0, T_s]$, we may write the resulting relation as

$$\mathbf{x}(T_s) = e^{T_s A} \mathbf{x}(0) + A^{-1}(e^{T_s A} - I)(B\mathbf{u} + \mathbf{b}).$$

For general A , we write the result as

$$\mathbf{x}(T_s) = e^{T_s A} \mathbf{x}(0) + P(A, T_s)(B\mathbf{u} + \mathbf{b}), \text{ wherein } P(A, t) = \sum_{j=0}^{\infty} \frac{A^j t^{j+1}}{(j+1)!}.$$

In theory, given T_s and A , we may compute the matrices $e^{T_s A}$ and $P(A, T_s)$ to arbitrary precision. This yields an affine expression for $\mathbf{x}(T_s)$ in terms of $\mathbf{x}(0), \mathbf{u}$.

In practice, however, arbitrary precision computation of the exponential map is often impractical, unless the matrix A is known to be diagonalizable with restrictions on its eigenvalues, or nilpotent. Therefore, for a general matrix A , we resort to error-prone numerical computations of $e^{T_s A}$ and $P(A, T_s)$.

The loss of soundness can be alleviated by using sophisticated numerical approximation schemes [27]. In particular, we can estimate both matrices using interval arithmetic calculations as $e^{T_s A} \in [\underline{E}_s, \overline{E}_s]$ and $P(A, T_s) \in [\underline{P}_s, \overline{P}_s]$ by taking into account the arithmetic and truncation errors of the resulting power series expansions [6, 29, 18]. Therefore, the resulting relationalization obtained is *interval linear* of the form $\mathbf{x}' \in [\underline{E}_s, \overline{E}_s]\mathbf{x} - [\underline{P}_s, \overline{P}_s](B\mathbf{u} + \mathbf{b})$ which stands for the logical formula

$$R(\mathbf{x}, \mathbf{u}, \mathbf{x}') : (\exists E \in [\underline{E}_s, \overline{E}_s], P \in [\underline{P}_s, \overline{P}_s]) \mathbf{x}' = E\mathbf{x} - P(B\mathbf{u} + \mathbf{b}).$$

As such, the relation above cannot be expressed in linear arithmetic. We will expand upon the treatment of interval linear relations later in this section.

Example 2. Going back to the system in Ex. 1, we find relational abstractions for mode n_0 when the system does not take an autonomous transition within the sampling period of 0.2 time units. Using a numerically computed matrix exponential, we obtain the relation $\mathbf{x}' = \begin{pmatrix} 0.7669 & 0.214 \\ 0.232 & 1.07 \end{pmatrix} \mathbf{x} + \begin{pmatrix} 0.1635 \\ -0.079 \end{pmatrix}$. On the other hand, using the interval arithmetic based method described by Goldsztejn [18], we obtain the relation

$$\mathbf{x}' \in \begin{pmatrix} [0.7669282852020186, 0.7669282852020187] & [0.2139643726426075, 0.2139643726426076] \\ [0.2317947337083272, 0.2317947337083273] & [1.0700444963848672, 1.0700444963848673] \end{pmatrix} \mathbf{x} + \begin{pmatrix} [0.1635149326785402, 0.1635149326785403] \\ [-0.0789845829507958, -0.0789845829507957] \end{pmatrix}.$$

While pathological cases for matrix exponential computation are known (Cf. Goldsztejn [18]), the rather tight interval bounds for the exponential seem to be quite common in our benchmarks, and therefore, the use of numerically computed matrix exponentials may be quite satisfactory for many applications, wherein the dynamics are obtained as an approximation of the physical reality in the first place.

Applying the same for mode n_2 , we obtain the relation $\mathbf{x}' = \begin{pmatrix} 1.5245 & 0.2181 \\ 0.0182 & 0.4885 \end{pmatrix} \mathbf{x} + \begin{pmatrix} -0.1626 \\ -0.0867 \end{pmatrix}$. Once again, an interval computation yields intervals of width 10^{-16} or less centered around the numerically computed value.

3.1 Dealing with Autonomous Transitions

Thus far, we have described a simple relationalization scheme under the assumption that no autonomous transitions were taken by the plant during a sampling time period. We will now describe the treatment of autonomous transitions that can be taken by the plant between two successive samplings. In general, there is no *a priori* bound on the number of such transitions that a plant can take in any given period ($nT_s, (n+1)T_s$). Even if the plant is assumed to be non-Zeno, any relationalization has to capture the effects of the plant executing a finite but unbounded number of transitions. We remedy this situation by making two assumptions regarding the plant: (a) There is a minimum dwell time $T_D > 0$ for each mode q of the plant. In other words, whenever a run of the plant enters some mode q , it remains there for at least T_D time units before an

autonomous transition is enabled. (b) No autonomous transitions can be taken precisely at the time instant $t = jT_s$ for $j \in \mathbb{Z}$.

The first assumption provides a bound $N = \left\lceil \frac{T_s}{T_D} \right\rceil$ on the maximum number of autonomous transitions taken inside a sampling interval. For this paper, we will assume that $N = 1$ to simplify the presentation, i.e., the controller is assumed to sample the plant fast enough to restrict the number of autonomous transitions in any sample period to at most 1. The second assumption allows us to use the standard *interleaving semantics* when the relationalization of the plant and the system are composed. This assumption fails if the execution time of the controller is not negligible compared to the time scale of the plant dynamics, as is sometimes the case. However, if bounds are known on the execution times, we may compute relationalizations of the plant for two time steps, one for the controller step and the other for the sampling period. Likewise, the basic ideas presented here extend to more sophisticated task execution schedules for control tasks.

Let us assume that a single autonomous transition $\tau : \langle q_1, q_2, \rho \rangle$ is taken during the time $t \in (0, T_s)$. Our goal is to derive a relation $R_\tau((q_1, \mathbf{x}), \mathbf{u}, (q_2, \mathbf{x}'))$ characterizing all possible pairs of states (q_1, \mathbf{x}) and (q_2, \mathbf{x}') so that the plant may evolve from continuous state \mathbf{x} in mode q_1 at time $t = 0$ to the state (q_2, \mathbf{x}') at time $t = T_s$ with the transition τ taken at some time instant $0 < t < T_s$. The resulting R_τ for each $\tau \in \mathcal{T}_{aut}$ will form a disjunct of the overall relation R_P for the plant.

Figure 3 summarizes the situation diagrammatically. We let \mathbf{y} be the valuation to continuous variables just prior to τ being taken and \mathbf{y}' be the valuation just after τ is taken. Let t be the time instant at which τ is taken. Let the dynamics in mode q_i for $i = 1, 2$ be given by $\frac{d\mathbf{x}}{dt} = A_i\mathbf{x} + B_i\mathbf{u} + \mathbf{b}_i$. Therefore, $\mathbf{x} = e^{-tA_1}\mathbf{y} - P(A_1, -t)(B_1\mathbf{u} + \mathbf{b}_1) \wedge \mathbf{x}' = e^{(T_s-t)A_2}\mathbf{y}' - P(A_2, T_s - t)(B_2\mathbf{u} + \mathbf{b}_2)$. The overall relation is given by

$$R_\tau(\mathbf{x}, \mathbf{u}, \mathbf{x}') : (\exists t, \mathbf{y}, \mathbf{y}') \begin{pmatrix} \mathbf{x} = e^{-tA_1}\mathbf{y} - P(A_1, -t)(B_1\mathbf{u} + \mathbf{b}_1) \\ \mathbf{x}' = e^{(T_s-t)A_2}\mathbf{y}' - P(A_2, T_s - t)(B_2\mathbf{u} + \mathbf{b}_2) \\ 0 < t < T_s \wedge \rho_\tau(\mathbf{y}, \mathbf{y}') \end{pmatrix}. \quad (1)$$

Note that we have chosen to encode $\mathbf{x} = e^{-tA_1}\mathbf{y}$ instead of encoding the dynamics in the forward direction $\mathbf{y} = e^{tA_1}\mathbf{x}$. This seemingly arbitrary choice will be seen to make the subsequent quantifier elimination problem easier.

Eliminating Quantifiers: The main problem with the relation R_τ derived in Eqn. (1) is that the matrices e^{tA_i} and $P(A_i, t)$ are, in general, *transcendental functions* of time. It is computationally intractable to manipulate these relations inside decision procedures. To further complicate matters, the variable t is existentially quantified. Removing this




Fig. 3. Schematic for relational abstraction of a single autonomous transition.

quantifier poses yet another challenge. However, our goal will be to derive an over-approximation of R_τ expressible in linear arithmetic.

To this end, the main challenge is to construct a good quality and linear over-approximation $R_\tau^a(\mathbf{x}, \mathbf{u}, \mathbf{x}')$ of the relation R_τ . We address this challenge using interval arithmetic techniques.

Interval Over-approximation We subdivide the interval $[0, T_s]$ into $M > 0$ subintervals each of width $\delta = \frac{T_s}{M}$. Next, we consider each subinterval of the form $t \in [i\delta, (i+1)\delta]$ and use interval arithmetic evaluation for the functions e^{tA_i} and $P(A_i, t)$ to obtain a conservative approximation valid for the subinterval. In effect, we over-approximate R_τ as a disjunction

$$R_\tau^I : \bigvee_{0 \leq i < M} (\exists \mathbf{y}, \mathbf{y}') \left(\begin{array}{l} \mathbf{x} \in [\underline{E}_{i,1}, \overline{E}_{i,1}] \mathbf{y} - [\underline{P}_{i,1}, \overline{P}_{i,1}] (B_1 \mathbf{u} + \mathbf{b}_1) \wedge \\ \mathbf{x}' \in [\underline{E}_{i,2}, \overline{E}_{i,2}] \mathbf{y}' - [\underline{P}_{i,2}, \overline{P}_{i,2}] (B_2 \mathbf{u} + \mathbf{b}_2) \wedge \\ \rho_\tau(\mathbf{y}, \mathbf{y}') \end{array} \right), \quad (2)$$

wherein $[\underline{E}_{i,1}, \overline{E}_{i,1}]$ is a safe interval enclosure of $e^{-(i+1)\delta, -i\delta} A_1$ while $[\underline{E}_{i,2}, \overline{E}_{i,2}]$ is an enclosure of $e^{(T_s - [i\delta, (i+1)\delta]) A_2}$. Likewise, $[\underline{P}_{i,1}, \overline{P}_{i,1}]$ and $[\underline{P}_{i,2}, \overline{P}_{i,2}]$ are safe enclosures of $P(A_1, [-(i+1)\delta, -i\delta])$ and $P(A_2, (T_s - [i\delta, (i+1)\delta]))$, respectively.

The resulting over-approximation is a disjunction of M interval linear relations. In effect, the transcendental relation R_τ in Eq. (1) is over-approximated by an algebraic (bilinear) relation R_τ^I . The over-approximation error can be made as small as necessary by increasing the number of subdivisions M , and by using a more expensive procedure for deriving a better approximation of the exponentials by intervals. The problem of evaluating safe interval enclosures to the matrices $e^{[t_1, t_2]A}$ and $P(A, [t_1, t_2])$ uses the idea of scaling and squaring with Horner's rule for evaluating the truncated power series, precisely as described by Goldsztejn [18]. A convenient trick used in our implementation folds the computation of $e^{A, [t_1, t_2]}$ and $P(A, [t_1, t_2])(B\mathbf{u} + \mathbf{b})$ into a single matrix exponential computation for a block matrix of the form $\begin{pmatrix} A & B & \mathbf{b} \\ 0 & 0 & 0 \end{pmatrix}$.

Example 3. Consider the hybrid automaton described in Ex. 1. We wish to consider the relational abstraction when τ_1 is taken sometime during the sampling period of 0.2 seconds. To this end, we will choose $M = 2$ and consider two possible intervals for the switching time t when the transition τ_1 is taken $J_1 : [0, 0.1]$ and $J_2 : [0.1, 0.2]$. Considering interval J_1 , we obtain the following relation (intervals are rounded to 2 significant digits for presentation):

$$R_{\tau, J_1} : (\exists \mathbf{y}) \left[\begin{array}{l} \mathbf{x} \in \left(\begin{bmatrix} 0.99 & 1.17 \\ -0.14 & 0.0 \end{bmatrix} \begin{bmatrix} -0.13 & 0.01 \end{bmatrix} \right) \mathbf{y} + \left(\begin{bmatrix} -0.11 & 0 \\ -0.01 & 0.05 \end{bmatrix} \right) \wedge \\ \mathbf{x}' \in \left(\begin{bmatrix} 1.23 & 1.53 \\ 0.0 & 0.02 \end{bmatrix} \begin{bmatrix} 0.09 & 0.25 \end{bmatrix} \right) \mathbf{y} + \left(\begin{bmatrix} -0.16 & -0.07 \\ -0.1 & -0.04 \end{bmatrix} \right) \wedge \mathbf{y} \in G_{\tau_1} \end{array} \right].$$

Templatization The next step is to use a *templatization* technique to effectively eliminate the quantifiers \mathbf{y}, \mathbf{y}' from the relation R_τ^I in Equation (2) while, at the same time,

over-approximating the result by means of a linear arithmetic over-approximation. Recall that each disjunct in Equation 2 is an *interval linear assertion* of the form

$$R_j^I : (\exists \mathbf{y}, \mathbf{y}') \left(\begin{array}{l} \mathbf{x} \in [\underline{E}_{j,1}, \overline{E}_{j,1}] \mathbf{y} - [\underline{P}_{j,1}, \overline{P}_{j,1}] (B_1 \mathbf{u} + \mathbf{b}_1) \wedge \\ \mathbf{x}' \in [\underline{E}_{j,2}, \overline{E}_{j,2}] \mathbf{y}' - [\underline{P}_{j,2}, \overline{P}_{j,2}] (B_2 \mathbf{u} + \mathbf{b}_2) \wedge \rho_\tau(\mathbf{y}, \mathbf{y}') \end{array} \right),$$

An interval linear constraint of the form $\sum_{j=1}^n I_j x_j + I_0 \leq 0$ is a place holder for a *bi-linear* constraint $\sum_{j=1}^n w_j x_j + w_0 \leq 0$, wherein, w_0, \dots, w_n are freshly introduced variables and each w_j is constrained by requiring that $w_j \in I_j$.

In order to eliminate \mathbf{y}, \mathbf{y}' from this relation, a technique for eliminating quantifiers for real arithmetic such as Cylindrical Algebraic Decomposition (CAD) [9], or a more efficient version for quadratic polynomials is called for [41, 13, 38]. However, the downside of using such complex techniques include (a) it is well known that QE over non-linear constraints is a hard problem with limited scalability, and (b) the result after elimination will, in general, be a set of polynomial inequalities (semi-algebraic constraint). Therefore, the resulting relationalization may not be easy to reason with for existing tools.

We present a more efficient alternative that side steps the elimination altogether, relying instead on the use of templates and optimization:

1. We choose a set of template expressions $e_k(\mathbf{x}, \mathbf{x}', \mathbf{u})$ involving the variables \mathbf{x} , \mathbf{u} and \mathbf{x}' . We discuss a natural choice for these templates subsequently.
2. For each e_j , we carry out the optimization: $\min e_k$ s.t. $R_j^I(\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{y}', \mathbf{x}')$. If the problem is feasible and bounded, the a_k allows us to conclude that

$$R_j^I(\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{y}', \mathbf{x}') \Rightarrow e_k(\mathbf{x}, \mathbf{u}, \mathbf{x}') \geq a_k.$$

As a result by choosing some $K > 0$ templates e_1, \dots, e_K , we obtain an assertion

$$e_1(\mathbf{x}, \mathbf{u}, \mathbf{x}') \geq a_1 \wedge e_2 \geq a_2 \wedge \dots \wedge e_K(\mathbf{x}, \mathbf{u}, \mathbf{x}') \geq a_K.$$

This assertion serves as an over-approximation to R_j^I with the quantified variables \mathbf{y}, \mathbf{y}' eliminated through optimization. We now provide a natural scheme for the choice of templates, and then discuss how the optimization problem for each template expression can be solved.

The overall relationalization of the plant R_P is the disjunction of the relations R_q , for each mode, and R_τ^I , for each autonomous transition τ .

Theorem 1. *For any pairs of states $\sigma : (q, \mathbf{x})$ and $\sigma' : (q', \mathbf{x}')$ such that \mathbf{x}' is reachable from \mathbf{x} in T_s seconds for constant control input \mathbf{u} , the computed relational abstraction R_P satisfies $R_P(\sigma, \sigma', \mathbf{u})$.*

Choosing Template Expressions A natural choice for template expression presents itself in our setup by considering the midpoints of the intervals used in the matrix exponential computations. We note that \mathbf{y} is the state obtained starting from \mathbf{x} and evolving in mode 1 for time $[i\delta, (i+1)\delta)$. Likewise, \mathbf{x}' is obtained by evolving according to the state \mathbf{y}' for time $[T_s - (i+1)\delta, T_s - i\delta)$. Finally, $\mathbf{y}' = U(\mathbf{y})$, wherein U is

the affine update map for transition τ . In practice, δ is chosen to be small enough to yield tight enclosures to e^{tA_i} and $P(A_i, t)$ matrices. Therefore, a natural choice of template expression is obtained by considering the midpoints of the time intervals involved. Specifically, we consider the affine expressions defined by

$$\mathbf{x}' - e^{(T_s - t_m)A_2} U(e^{t_m A_1} \mathbf{x}), \text{ where } t_m = (i + \frac{1}{2})\delta.$$

Example 4. In Ex. 3, we showed the interval linear relation obtained by considering switching times in the interval $t \in [0.0, 0.1]$. The midpoint of this interval is $t_m = 0.05$. Therefore, we consider the mode n_0 taken for time 0.05 units followed by 0.15 units of mode 1 for generating a suitable template. These template expressions are given by $e_1 : \mathbf{x}' - 1.31\mathbf{x} - 0.25\mathbf{y}$ and $e_2 : \mathbf{y}' - 0.05\mathbf{x} - 0.6\mathbf{y}$. We seek to bound these expressions to obtain a linear arithmetic over-approximation.

Encoding Optimization Next, we turn our attention to setting up the optimization problem for a given template expression $\mathbf{c}\mathbf{x} + \mathbf{d}\mathbf{x}'$. The intermediate states \mathbf{y}, \mathbf{y}' are related by interval linear expressions of the form

$$\mathbf{x}' \in [\underline{E}_2, \overline{E}_2]\mathbf{y}' + [\underline{P}_2, \overline{P}_2], \quad \mathbf{x} \in [\underline{E}_1, \overline{E}_1]\mathbf{y} + [\underline{P}_1, \overline{P}_1].$$

To set up the optimization problem, we substitute these expressions for \mathbf{x}, \mathbf{x}' in the template to obtain $\mathbf{c}([\underline{E}_1, \overline{E}_1]\mathbf{y} + [\underline{P}_1, \overline{P}_1]) + \mathbf{d}([\underline{E}_2, \overline{E}_2]\mathbf{y}' + [\underline{P}_2, \overline{P}_2])$. This is, in fact, an interval linear expression involving \mathbf{y}, \mathbf{y}' . The overall optimization problem reduces to: $\min [\underline{\mathbf{c}}, \overline{\mathbf{c}}]\mathbf{y} + [\underline{\mathbf{d}}, \overline{\mathbf{d}}]\mathbf{y}' + [\underline{c}_0, \overline{c}_0]$ s.t. $\rho_\tau(\mathbf{y}, \mathbf{y}')$. Here $[\underline{\mathbf{c}}, \overline{\mathbf{c}}] = \mathbf{c}[\underline{E}_1, \overline{E}_1]$, $[\underline{\mathbf{d}}, \overline{\mathbf{d}}] = \mathbf{d}[\underline{E}_2, \overline{E}_2]$ and $[\underline{c}_0, \overline{c}_0] = \mathbf{c}[\underline{P}_1, \overline{P}_1] + \mathbf{d}[\underline{P}_2, \overline{P}_2]$. The problem has an interval linear objective and linear constraints. We now show that the constraints can be encoded into a disjunctive linear program.

Theorem 2. *The optimization of an interval linear objective w.r.t linear constraints*

$$\min [\underline{\mathbf{c}}, \overline{\mathbf{c}}] \times \mathbf{z} \text{ s.t. } \mathbf{A}\mathbf{z} \leq \mathbf{b},$$

can be equivalently expressed as a linear program with disjunctive constraints:

$$\min \underline{\mathbf{c}}\mathbf{z}^+ - \overline{\mathbf{c}}\mathbf{z}^- \text{ s.t. } \mathbf{A}\mathbf{z}^+ - \mathbf{A}\mathbf{z}^- \leq \mathbf{b}, \mathbf{z}^+, \mathbf{z}^- \geq 0, z_i^+ = 0 \vee z_i^- = 0$$

where $\mathbf{z} = \mathbf{z}^+ - \mathbf{z}^-$.

Proof. We may decompose any vector \mathbf{z} as $\mathbf{z} = \mathbf{z}^+ - \mathbf{z}^-$, where $\mathbf{z}^+, \mathbf{z}^- \geq 0$, and enforce $z_i^+ z_i^- = 0$. Next, consider the objective $[\underline{\mathbf{c}}, \overline{\mathbf{c}}] \times (\mathbf{z}^+ - \mathbf{z}^-)$. Since correspondent entries in $\mathbf{z}^+, \mathbf{z}^-$ cannot be positive at the same time, we may write the objective as a linear expression $\underline{\mathbf{c}}\mathbf{z}^+ - \overline{\mathbf{c}}\mathbf{z}^-$. Finally, the complementarity condition $z_i^+ z_i^- = 0$ is rewritten as $z_i^+ = 0 \vee z_i^- = 0$.

A simple approach to solve the optimization problem for disjunctive constraints is to use a linear arithmetic SMT solver to repeatedly obtain feasible solutions $\mathbf{z}^+, \mathbf{z}^-$. For a given feasible solution output by the SMT solver, we fix a minimal set of the values for $\mathbf{z}^+, \mathbf{z}^-$ to zero to enforce the complementarity constraints $z_i^+ z_i^- = 0$, leaving the

remaining variables as unknowns. An LP solver is then used to compute an optimal value f^* for the objective function f , based on the remaining constraints. This yields a potential optimum. Next, we add a *blocking constraint* $f > f^*$ to the SMT solver and search for a different solution. The process is carried out until the SMT solver returns UNSAT. At this point, we output the last optimal solution as the final value.

Example 5. Continuing with the examples worked out in Ex. 3, we perform the optimization of the template expressions chosen in Ex. 4 to obtain the relational abstraction:

$$R_{\tau_1, J_1} : -1.0 \leq 1.31x + 0.25y - x' \leq 1.24 \wedge -0.32 \leq 0.05x + 0.6y - y' \leq 0.51 .$$

Likewise, considering the time interval $J_2 : [0.1, 0.2]$ for the switching time, we obtain the abstraction:

$$R_{\tau_1, J_2} : -1 \leq 0.94x + 0.25y - x' \leq 0.88 \wedge -0.54 \leq 0.17x + 0.9y - y' \leq 0.7708 .$$

The overall timed relational abstraction for the sampling period where τ_1 can be taken sometime in between is $R_{\tau_1} : R_{\tau_1, J_1} \vee R_{\tau_1, J_2}$.

4 Experimental Evaluation

We first briefly describe our implementation of the relational abstractor using the techniques presented here.

Implementation: The relational abstractor takes in a plant description including the sample time T_s , and outputs the relation as a SAL transition system [32]. The relationalization is performed for the continuous dynamics in each mode by computing a matrix exponential. A numerical approximation of the matrix exponential function is obtained using Pade's approximation [27]. We have also implemented a procedure that provides a sound interval enclosure of the exponential function over interval matrices using the ideas described by Goldsztejn [18]. However, this procedure is used solely for dealing with autonomous transitions.

Autonomous transition between modes are handled using the algorithms presented so far. We implicitly assume minimum dwell time greater than or equal to the sampling time for the controller. The optimization problems encountered for autonomous transitions are solved using the SMT solver Z3 [12]. SAL provides a k -induction and BMC engine using the solver Yices [14]. This was used for analyzing the resulting composed transition system for our evaluations.

Benchmarks: Table 1 shows the benchmarks used in our evaluation along with their sources. The benchmarks vary in dimensionality and number of transitions. Note that many benchmarks do not contain autonomous transitions. For each benchmark, we performed the relational abstraction for different sampling times T_s , and used SAL to analyze safety properties.

The NAV benchmarks, due to Ivancic and Fehnker [15], model a particle traveling through many $2D$ cells that each have a different dynamics. We consider two versions of this benchmark (a) the transitions in the benchmark are all interpreted as controlled, commanded by a controller, or (b) transitions are autonomous in nature. Starting with

Model	Description	# Var	# Mode	# Trs	Prop.	Description
InvPen	Inverted Pendulum Control	5	1	1	$\theta_b(0.05)$	Angle $\theta \in [-0.05, 0.05]$
SNCS	Network Control System [42]	2	1	2	P_1 P_2	$(x, y) \in [-100, 100]^2$ $(x, y) \in [-10^4, 10^4]^2$
ACC	Adaptive Cruise Control [21]	4	1	2	SAFE	No collision between cars.
ACC-T	ACC + transmission [21]	3	20	24	SAFE	No collision between cars.
Heat-x	Room heater [15]	9	8	20	LB	Lower bounds on temp. Cf. [15]
Nav-y	NAV benchmarks [15]	4	[7,16]	[9,16]	RA RB	Cell A is unreachable Cell B is unreachable
Toy	Example 1	2	2	5	$band(k)$	$n_1 \Rightarrow \mathbf{x} \leq k$
Ring(n,m)	Cf. Section 4	n	m+1	m+1	$band(k)$	$n_4 \Rightarrow \mathbf{x} \leq k$

Table 1. The benchmarks used in our experiments at a glance.

all controlled transitions, we introduce uncontrolled transitions incrementally into these benchmarks.

Ring Benchmarks: We created a set of sampled data control systems with autonomous transitions. We consider a plant with $k + 1$ modes, wherein modes m_1, \dots, m_k are governed by stable dynamics, while mode m_{k+1} is an unstable mode. The controller seeks to stabilize this mode by periodically sensing the plant’s state and applies a control that reverts it back to mode m_1 .

The benchmark instance $\text{Ring}(n, k)$ consists of n state variables and $k + 1$ modes. The autonomous transitions are added from mode i to mode $i + 1$ for $i \leq k$, while the controlled transition leads from mode $k + 1$ to mode 1. The dynamics in each mode is of the form $\frac{d\mathbf{x}}{dt} = A_i(\mathbf{x} - \mathbf{b}_i)$, wherein, for the stable modes A_i is a Hurwitz matrix and \mathbf{b}_i is a designated equilibrium for m_i . For the unstable mode, we ensure that A_i has a positive eigenvalue. The switch from m_i to m_{i+1} takes place inside a box $[\mathbf{b}_i - \epsilon, \mathbf{b}_i + \epsilon]$. The controller periodically senses the plant and whenever $|\mathbf{x}| > c$ for some fixed c , it brings the dynamics back into the box $|\mathbf{x}| < c$ while transitioning to mode m_1 . We wish to check whether all trajectories lie inside a box $|\mathbf{x}| \leq c + d$, for varying tolerances d .

Results: Table 2 shows the experimental results on benchmarks that do not have autonomous transitions. Our experiments are attempted using numerous values of sample times for each property until either a proof is obtained for the controller or the SAL tool fails due to a timeout. In the absence of autonomous transitions, the relationalization time for all these benchmarks was well under 1 second. We also note that the counterexamples generated by SAL can be concretized, since the timed abstractions involving matrix exponentials are seen to be quite precise.

Table 3 shows the results for systems with controlled and autonomous transitions. These include the system from Ex. 1, the $\text{Ring}(n, k)$ systems for varying n and the NAV benchmark instances as we increase the number of autonomous transitions. We observe that making all the transitions autonomous leads to a counterexample. This counterexample may potentially be an artifact of the precision loss due to relationalization of autonomous transitions. Future work will consider the refinement of these

Model	Prop	T_s	Result	Depth	Time
InvPen	$\theta_b(0.05)$	0.1	CE	1	0.1
			P	12	0.9
SNCS	P_1	1.7	P	2	0.1
	P_2	1.8	CE	93	5.6
ACC	SAFE	0.1	P	7	0.1
	gap(100)	0.1	P	4	0.1
ACC-T	SAFE	1	P	14	2.2
Nav 1-7	1	RB	P	<11	<5
	1	RA	CE	<13	<2
Nav 8	1	RB	CE	7	0.27
	0.2	RB	F	25	> 1h
Nav 9	1	RB	P	19	213.05
	1	RA	CE	9	0.37
Nav 10	1	RB	CE	19	28.37
	0.5	RB	F	25	> 1h

Model	T_s	Prop	result	depth	time
Heat1	LB	1	CE	4	0.1
		.2	CE	8	0.1
		0.1	P	37	1967
Heat2	LB	1	CE	4	0.1
		0.2	P	17	160
Heat3	LB	1	CE	2	0.1
		0.2	CE	17	27
		0.1	F	30	> 1h
Heat4	LB	1	CE	2	0.1
		0.1	CE	10	1.22
		0.02	F	25	> 1h

Table 2. Results on benchmarks without autonomous transitions. All timings were measured in seconds on a laptop running Intel Core i7-2820Q 2.30GHz processor (x86_64 arch) with 8GB RAM running Ubuntu 11.04 Linux 2.6.38-13. Legend: **CE** indicates true counter-example, **P** indicates proofs, **F** indicates failure due to timeout.

counterexamples by subdividing the transition switching time intervals further based on spurious counterexamples. We note that the time for relationalization remains a small fraction of the time needed to check the system. The relationalization scheme can be improved further if SMT solvers such as Z3 can be modified to support the optimization of objective functions.

Comparison With SpaceEx: We now compare the results obtained for our approach with the SpaceEx tool over the same set of benchmarks [16]. While performing the comparison with SpaceEx, we reiterate two key points of difference: (a) SpaceEx handles general hybrid systems with support for synchronous time-triggered semantics as well as the standard event-triggered semantics given by guards and resets. Our technique is specialized to sampled data control systems. (b) SpaceEx attempts to characterize the reachable sets for all time instances, whereas our approach focuses on proving properties at the periodic sampling times.

Typically, running the benchmarks in SpaceEx required choosing from a range of parameters including template domains, underlying implementation, flowpipe tolerances, error tolerances, local and global time horizons and limits on the number of iteration. We ran SpaceEx for each benchmark using multiple option sets, choosing the option that provided the “best answer” with as few warnings as possible. However, it may be possible to obtain qualitatively different results using choices for the parameters that were unexplored in our experiments. A detailed table summarizing our experiences is available upon request.

Table 4 presents a summary of the results obtained by running SpaceEx on our benchmarks. We note that in many cases, SpaceEx did not reach a fixed point. Therefore, whenever a property proof was obtained, we report if the proof was obtained over

Model	Prop	result	depth	T_{sal}	T_{rel}
Toy	$bnd(8)$	P	2	0.1	< .1
	$bnd(6)$	P	2	0.1	
	$bnd(5.5)$	P	2	0.3	
	$bnd(5)$	CE	70	204	
Ring(3,4)	$bnd(20)$	P	10	5.2	0.8
	$bnd(15)$	P	30	451	
Ring(5,4)	$bnd(25)$	P	10	34.7	2.8
	$bnd(20)$	P	10	56.2	
	$bnd(15)$	F	20	> 1h	
Ring(7,4)	$bnd(25)$	P	10	157	11.9
	$bnd(20)$	P	10	357	
	$bnd(15)$	F	20	> 1h	
Ring(9,4)	$bnd(25)$	P	10	515	19.1
	$bnd(20)$	P	10	2929	
Ring(11,4)	$bnd(25)$	F	10	> 1h	150

Model	Prop	T_s	# Aut.	result	depth	time
Nav1	RB	0.2	6	P	9	199
			14	P	9	72
			21	P	9	96
			24	F	9	169
Nav2	RB	0.2	All	CE	9	161
			20	P	9	92
			21	F	9	160
			All	CE	7	151
Nav3	RB	0.2	22	P	9	83
			All	CE	6	13
Nav4	RB	0.2	9	P	18	1305
			20	F	18	> 1h
			All	CE	6	7

Table 3. Results on systems with autonomous transitions. For the NAV benchmarks, autonomous transitions between cells were incrementally enabled over the controlled transitions until all transitions were autonomous. T_{sal} refers to running time for SAL and T_{rel} the running time for the relationalization.

a finite time horizon. Likewise, for cases where the property was not proved, we ran SpaceEx for the minimum number of iterations until a potential violation is observed.

The comparison between our approaches clearly showcases some of the relative merits and demerits of our approach vis-a-vis SpaceEx. There are many benchmarks wherein our approach is able to establish the property over an infinite time horizon using k -induction, whereas SpaceEx either proves the property over a finite time horizon or fails. On the other hand, the NAV benchmarks are an interesting case where SpaceEx’s performance is at par or clearly superior to that of our approach.

For some of the Ring examples, we observed that using the bounds obtained by SpaceEx as inductive strengthenings enabled the k -induction technique to prove the property for a smaller value of k , leading to improved running times.

Our future work will focus on an integration of the approaches considered here in combination with tools such as SpaceEx to achieve infinite horizon safety property proofs. Another important area of future research will be to extend our approach to analyze non linear hybrid systems, which are much more challenging.

References

1. R. Alur, T. Dang, and F. Ivančić. Counter-example guided predicate abstraction of hybrid systems. In *TACAS*, volume 2619 of *LNCS*, pages 208–223. Springer, 2003.
2. D. Arney, R. Jetley, P. Jones, I. Lee, and O. Sokolsky. Formal methods based development of a PCA infusion pump reference model: Generic infusion pump (GIP) project. In *Proc. High Confidence Medical Devices, Software Systems and Medical Device Plug and Play Interoperability*, 2007.
3. E. Asarin, T. Dang, and A. Girard. Hybridization methods for the analysis of nonlinear systems. *Acta Informatica*, 43:451–476, 2007.

Model	Prop	T_s	Result	Time
InvPen	$\theta_b(0.05)$	0.05	F	6
SNCS	P_1	1.7	F	371
Heat1	LB	0.1	F	557
Ring(3,4)	$bnd(20)$	0.2	P (FT)	3475
Ring(5,4)	$bnd(25)$	0.2	P (FT)	2051
Ring(7,4)	$bnd(25)$	0.2	F	6323
Ring(9,4)	$bnd(25)$	0.2	F	709

Model	Prop	T_s	Result	Time
Nav 4(A)	RB	—	P (FT)	1501
Nav 6(A)	RB	—	F	1223
Nav 7(A)	RB	—	P (FT)	615
Nav 9(A)	RB	—	F	619
Nav 4(C)	RB	1	P (FT)	109
Nav 6(C)	RB	1	P (FT)	167
Nav 7(C)	RB	1	P	7
Nav 9(C)	RB	1	F	6

Table 4. Results of SpaceX tool on the benchmark. Legend: **A:** All transitions were Autonomous, **C:** All transitions were controlled Autonomous, **Prop:** property, **F:** Found potentially spurious counter-example (our approach proves model+property), **P:** Proved, **FT:** proof valid over a finite time horizon, **Time:** Running time in seconds.

4. J. Berdine, A. Chawdhary, B. Cook, D. Distefano, and P. W. O’Hearn. Variance analyses from invariance analyses. In *POPL*, pages 211–224. ACM, 2007.
5. A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS*, volume 1579 of *LNCS*, pages 193–207, 1999.
6. P. Bochev and S. Markov. A self-validating numerical method for the matrix exponential. *Computing*, 43(1):59–72.
7. M. Bozzano, A. Cimatti, and F. Tapparo. Symbolic fault tree analysis for reactive systems. In *ATVA*, volume 4762 of *LNCS*, pages 162–176. 2007.
8. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.
9. G. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In H.Brakhage, editor, *Automata Theory and Formal Languages*, volume 33 of *LNCS*, pages 134–183. Springer, 1975.
10. P. Cousot and R. Cousot. Abstract Interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *ACM Principles of Programming Languages*, pages 238–252, 1977.
11. T. Dang, O. Maler, and R. Testylier. Accurate hybridization of nonlinear systems. In *HSCC ’10*, pages 11–20. ACM, 2010.
12. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
13. A. Dolzmann and T. Sturm. REDLOG: Computer algebra meets computer logic. *ACM SIGSAM Bulletin*, 31(2):2–9, June 1997.
14. B. Dutertre and L. de Moura. The YICES SMT solver. Cf. <http://yices.csl.sri.com/tool-paper.pdf>, last viewed Jan. 2009.
15. A. Fehnker and F. Ivančić. Benchmarks for hybrid systems verification. In *HSCC*, volume 2993 of *LNCS*, pages 326–341. Springer, 2004.
16. G. Frehse, C. L. Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler. Spaceex: Scalable verification of hybrid systems. In *CAV*, volume 6806 of *Lecture Notes in Computer Science*, pages 379–395. Springer, 2011.
17. A. Girard. Reachability of uncertain linear systems using zonotopes. In *HSCC*, volume 3414 of *LNCS*, pages 291–305. Springer, 2005.
18. A. Goldsztejn. On the exponentiation of interval matrices, 2009. Preprint (Working Paper) # hal-00411330, version 1. Cf. <http://hal.archives-ouvertes.fr/hal-00411330/fr/>.

19. C. L. Guernic and A. Girard. Reachability analysis of linear systems using support functions. *Nonlinear Analysis: Hybrid Systems*, 4(2):250–262, 2010.
20. S. Gulwani, S. Jain, and E. Koskinen. Control-flow refinement and progress invariants for bound analysis. In *PLDI*, 2009.
21. S. Gulwani and A. Tiwari. Constraint-based approach for hybrid systems. In *CAV*, volume 5123 of *LNCS*, pages 190–203, 2008.
22. N. Halbwachs, Y.-É. Proy, and P. Roumanoff. Verification of real-time systems using linear relation analysis. *Formal Methods in System Design*, 11(2):157–185, 1997.
23. T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. Algorithmic analysis of nonlinear hybrid systems. *IEEE Transactions on Automatic Control*, 43:540–554, 1998.
24. B. Jeannot, N. Halbwachs, and P. Raymond. Dynamic partitioning in analyses of numerical properties. In *SAS*, volume 1694 of *LNCS*, pages 39–50, 1999.
25. A. B. Kurzhanski and P. Varaiya. Ellipsoidal techniques for reachability analysis. In *HSCC*, volume 1790 of *LNCS*, pages 202–214. Springer, 2000.
26. Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer, New York, 1995.
27. C. Moler and C. V. Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1):161–208, 2003.
28. R. Moore, R. B. Kearfott, and M. Cloud. *Introduction to Interval Analysis*. SIAM, 2009.
29. E. P. Oppenheimer and A. N. Michel. Application of interval analysis techniques to linear systems. II. the interval matrix exponential function. *IEEE Trans. on Circuits and Systems*, 35(10):1230–1242, 1988.
30. A. Podelski and A. Rybalchenko. Transition invariants. In *LICS*, pages 32–41. IEEE Computer Society, 2004.
31. A. Podelski and S. Wagner. Model checking of hybrid systems: From reachability towards stability. In *HSCC*, volume 3927 of *LNCS*, pages 507–521. Springer, 2006.
32. J. Rushby, P. Lincoln, S. Owre, N. Shankar, and A. Tiwari. Symbolic analysis laboratory (sal). Cf. <http://www.csl.sri.com/projects/sal/>.
33. S. Sankaranarayanan, T. Dang, and F. Ivančić. Symbolic model checking of hybrid systems using template polyhedra. In *TACAS*, volume 4963 of *LNCS*, pages 188–202. Springer, 2008.
34. S. Sankaranarayanan, H. Homaei, and C. Lewis. Model-based dependability analysis of programmable drug infusion pumps. In *FORMATS*, volume 6919 of *Lecture Notes in Computer Science*, pages 317–334, 2011.
35. S. Sankaranarayanan and A. Tiwari. Relational abstractions for continuous and hybrid systems. In *CAV*, volume 6806 of *LNCS*, pages 686–702. Springer, 2011.
36. M. Sheeran, S. Singh, and G. Stålmarck. Checking safety properties using induction and a sat-solver. In *FMCAD*, volume 1954 of *LNCS*, pages 108–125. Springer, 2000.
37. B. I. Silva, K. Richeson, B. H. Krogh, and A. Chutinan. Modeling and verification of hybrid dynamical system using checkmate. In *ADPM 2000*, 2000. available online from <http://www.ece.cmu.edu/~webk/checkmate>.
38. T. Sturm and A. Tiwari. Verification and synthesis using real quantifier elimination. In *ISSAC*, pages 329–336. ACM, 2011.
39. A. Tiwari. Abstractions for hybrid systems. *Formal Methods in Systems Design*, 32:57–83, 2008.
40. A. Tiwari and S. International. HybridSAL: A tool for abstracting HybridSAL specifications to SAL specifications, 2007. Cf. <http://sal.csl.sri.com/hybridsal/>.
41. V. Weispfenning. Quantifier elimination for real algebra—the quadratic case and beyond. In *Applied Algebra and Error-Correcting Codes (AAECC) 8*, pages 85–101, 1997.
42. W. Zhang, M. S. Branicky, and S. M. Phillips. Stability of networked control systems. *IEEE Control Systems Magazine*, 21:84–99, 2001.