

Anonymized Reachability of Hybrid Automata Networks

Taylor T. Johnson¹ and Sayan Mitra²

¹ University of Texas at Arlington, Arlington, TX 76019, USA
taylor.johnson@uta.edu

² University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
mitras@illinois.edu

Abstract. In this paper, we present a method for computing the set of reachable states for networks consisting of the parallel composition of a finite number of the same hybrid automaton template with rectangular dynamics. The method utilizes a symmetric representation of the set of reachable states (modulo the automata indices) that we call anonymized states, which makes it scalable. Rather than explicitly enumerating each automaton index in formulas representing sets of states, the anonymized representation encodes only: (a) the classes of automata, which are the states of automata represented with formulas over symbolic indices, and (b) the number of automata in each of the classes. We present an algorithm for overapproximating the reachable states by computing state transitions in this anonymized representation. Unlike symmetry reduction techniques used in finite state models, the timed transition of a network composed of hybrid automata causes the continuous variables of all the automata to evolve simultaneously. The anonymized representation is amenable to both reducing the discrete and continuous complexity. We evaluate a prototype implementation of the representation and reachability algorithm in our satisfiability modulo theories (SMT)-based tool, Passel. Our experimental results are promising, and generally allow for scaling to networks composed of tens of automata, and in some instances, hundreds (or more) of automata.

Keywords: hybrid automata network, reachability, verification, symmetry.

1 Introduction

Networks consisting of automata that communicate via shared variables are useful for modeling distributed algorithms such as mutual exclusion algorithms, media access control (MAC) such as time-division multiple access (TDMA) protocols, and distributed cyber-physical systems (CPS) such as air-traffic control systems [17]. However, as the state-space of the network consisting of parallel compositions of these automata grows exponentially in the number of automata N , automated analysis is challenging. It is particularly challenging for timed and hybrid systems, where the number of continuous variables (dimensions) also grows. Such networks are often specified in a symmetric manner—such as being composed of instantiations of an automaton template $\mathcal{A}(N, i)$ —and are often amenable to methods that exploit symmetries. Formal analysis and state-space construction methods that exploit symmetries have been thoroughly

investigated for many classes of system models, because such methods ameliorate the state-space explosion problem [1, 2, 5, 6, 9–11, 14, 15, 20–22].

For example, several methods exploiting symmetry have been developed and implemented for the Mur φ verification system [8] for discrete systems, such as the *scalarset* data structure [14], and the *repetitive id* data structure [15]. Advances in tools like UP-AAAL [3] and PAT [23] that exploit state-space symmetries have enabled scaling to larger models. For instance, the scalarset data structure from Mur φ was extended for timed systems and implemented in UPAAAL [11, 12], and a clock-symmetry reduction method has been implemented in the PAT model checker [21]. Quasi-equal clocks and variables for timed [13] and hybrid (multi-rate) [4] automata networks also allow reductions in state-space explosion, but do not require automata in the network to be identical (modulo identifiers), as we do. We focus on safety properties, and to the best of our knowledge, before this paper, such symmetry techniques have not yet been applied to systems with general continuous dynamics like the rectangular differential inclusions we consider (e.g., [4] analyzes multi-rate automata and does not allow differential inclusions). The method described in this paper and implemented in our *Passel* verification tool [16, 18, 19] uses the SMT solver Z3 [7]. The method is used as a subroutine in methods for performing uniform verification of parameterized networks of hybrid automata (e.g., verification for all network sizes, $\forall N \in \mathbb{N}, \mathcal{A}(N, 1) \parallel \dots \parallel \mathcal{A}(N, N) \models \zeta(N)$), although we highlight that this paper addresses fixed, constant choices of N only.

2 Hybrid Automata Network Syntax and Semantics

We specify the behavior of each participant in the network using a syntactic structure called a hybrid automaton template, denoted by $\mathcal{A}(N, i)$.¹ The special symbols N and i are natural numbers that respectively refer to the number of automata, and the i^{th} automaton. For a natural number n , the set $[n]$ is $\{1, \dots, n\}$. For a set S , the set S_{\perp} is $S \cup \{\perp\}$. Fixing a particular value of N gives concrete instances of $[N]$ and $[N]_{\perp}$.

Terms and Formulas. We use a class of formulas to: (a) specify the syntactic components of a hybrid automaton template $\mathcal{A}(N, i)$, and (b) represent sets of states symbolically in the reachability computation. Formulas are built-up from constants, variables, and terms of several types. The grammar for formulas is:

$$\text{ITerm} ::= \perp \mid 1 \mid N \mid i \mid p[i]$$

$$\text{DTerm} ::= l_c \mid q \mid q[\text{ITerm}]$$

$$\text{RTerm} ::= 0 \mid 1 \mid r_c \mid x \mid x[\text{ITerm}]$$

$$\text{RPoly} ::= \text{RTerm} \mid \text{RPoly}_1 + \text{RPoly}_2 \mid \text{RPoly}_1 - \text{RPoly}_2 \mid (\text{RPoly}_1 * \text{RPoly}_2)$$

$$\text{Atom} ::= \text{ITerm}_1 = \text{ITerm}_2 \mid \text{DTerm}_1 = \text{DTerm}_2 \mid \text{RPoly} < 0$$

$$\text{Formula} ::= \text{Atom} \mid \neg \text{Formula} \mid \text{Formula}_1 \wedge \text{Formula}_2 \mid \exists x \text{ Formula}$$

The grammar is composed of *index terms* (ITerm) with type $[N]_{\perp}$, *discrete terms* (DTerm) with type L , and *real terms* (RTerm) with type \mathbb{R} . For a discrete term, l_c is constant from L and q is a discrete variable. For a real term, r_c is a real numerical

¹ Readers interested in additional technical details are referred to [16, Chapters 2 and 4].

constant and x is a real variable. Index ($p[i]$), discrete ($q[\text{ITerm}]$), and real ($x[\text{ITerm}]$) *pointer variables* are names for arrays composed of N elements of the corresponding type, respectively referenced at an index variable i , or an evaluation of an index term ITerm . Atoms (Atom) are composed of ordered relations between real polynomials (RPoly), as well as equality relations between index terms and discrete terms. Formulas are composed of Boolean combinations of atoms and shorter formulas. Comparison operators are expressed using negation (\neg) and conjunction (\wedge) in formulas. Combining the Boolean operators \wedge and \neg with the $<$ operator, other comparison operators like $=$, \neq , \leq , $>$, and \geq , can be expressed. We assume the language contains the standard quantifiers and Boolean operators, even if not explicitly specified in the grammar (e.g., universal quantification \forall , implication \Rightarrow , disjunction \vee , less-than-or-equal \leq , etc.).

Variables. A hybrid automaton $\mathcal{A}(N, i)$ has a set of *variables*, each of which is a name used for referring to state and is a term in the grammar just defined. As specified in the grammar, each variable v is associated with a *type*—denoted $\text{type}(v)$ —that defines a set of values the variable may take. The type of a variable may be: (a) L : a finite set of locations names, (b) $[N]_{\perp}$: a set of automaton indices—called *pointers*—with the special element \perp that is not equal to any automaton’s index, or (c) \mathbb{R} : the set of real numbers. A variable may be a *local* variable with a name of the form $\text{variable_name}[i]$, or *global*, in which case its name does not have a symbolic index $[i]$. For example, $q[i] : L$, $p[i] : [N]_{\perp}$, and $x[i] : \mathbb{R}$ respectively define location, pointer, and real typed local variables, while $g : [N]_{\perp}$ is a global variable of pointer type. The sets of local and global variables are denoted by $V_L(i)$ and $V_G(i)$, respectively. The *valuation* of a variable v is a function that associates the variable name v to a value in its type $\text{type}(v)$. For a set of variables V , $\text{val}(V)$ is the set of valuations of each $v \in V$. For a set of variables V , $V' \triangleq \{v' | v \in V\}$ and $\dot{V} \triangleq \{\dot{v} | v \in V \wedge \text{type}(v) = \mathbb{R}\}$. V' is used for specifying resets of discrete transitions and \dot{V} is used for specifying continuous dynamics. For a formula ϕ , let: (a) $\text{vars}(\phi)$ be the set of variables appearing in ϕ , (b) $\text{ivars}(\phi)$ be the set of distinct index variables appearing in ϕ .

Let N be a symbol representing an arbitrary natural number and i be a symbol representing an arbitrary element of $[N]$. For the remainder of the paper, we fix N and refer to it implicitly in the remaining definitions. When clear from context, we drop the parameter N , for instance, a hybrid automaton template $\mathcal{A}(N, i)$ is written $\mathcal{A}(i)$, etc.

Definition 1. A hybrid automaton template $\mathcal{A}(N, i)$ is specified by the syntactic components: (a) $V(i)$: a finite set of variable names with associated types. (b) L : a finite set of location names. (c) $\text{Init}(i)$: an initial condition formula over $V(i)$. (d) $\text{Trans}(i)$: a finite set of discrete transition statements, each of which is a tuple $\langle \text{from}, \text{to}, \text{grd}, \text{rst} \rangle$, where $\text{from}, \text{to} \in L$, grd is a formula over $V(i)$ called a guard and rst is a formula over $V(i) \cup V'(i)$ called a reset. The guard is an enabling condition that must be satisfied so that a transition may be taken, and the reset models the update of state. (e) $\text{Traj}(i)$: for each element in L , there is a trajectory statement, each of which is a tuple $\langle \text{loc}, \text{inv}, \text{frate} \rangle$, where $\text{loc} \in L$, inv is a formula over the real variables $X(i)$ called an invariant, and frate is a formula over $X(i) \cup \dot{X}(i)$ called a flowrate that specifies how the real variables evolve over time. The invariant is an assertion that must be satisfied while $\mathcal{A}(i)$ is in loc , and the flow rate associates each real-valued variable of $\mathcal{A}(i)$ with a rectangular differential inclusion.




Fig. 1. Hybrid automaton template $\mathcal{A}(i)$ for MUX-INDEX-RECT mutual exclusion algorithm

Let $X(i) \triangleq \{v \in V(\mathbb{N}, i) \mid \text{type}(v) = \mathbb{R}\}$ be the set of variables of $\mathcal{A}(i)$ with real type.

MUX-INDEX-RECT (Figure 1) is a timed mutual exclusion algorithm with an imprecise real clock $x[i]$ that evolves between rates $lb \leq ub$. There is a global variable g with type $[\mathbb{N}]_{\perp}$. Each automaton i starts in `rem` with $x[i] = 0$ and $g = \perp$, then after waiting B time i may enter `try` which also sets the global variable g to the identifier i . After waiting in at least $2B$ time, it may enter the critical section `cs` and stay there for at least $3B$ and at most $6B$ time before returning to `rem` and setting $g = \perp$.

2.1 Semantics of Hybrid Automata Networks

For a hybrid automaton template $\mathcal{A}(i)$, we define a transition system to formalize the semantics of the network where \mathbb{N} instantiations of $\mathcal{A}(i)$ operate concurrently.

Definition 2. Let \mathbb{N} be a symbol representing an arbitrary natural number. A hybrid automata network is a tuple $\mathcal{A}^{\mathbb{N}} \triangleq \langle V^{\mathbb{N}}, Q^{\mathbb{N}}, \Theta^{\mathbb{N}}, T^{\mathbb{N}} \rangle$, where: (a) $V^{\mathbb{N}}$ are the variables of the network, $V^{\mathbb{N}} \triangleq V_G \cup \bigcup_{i=1}^{\mathbb{N}} V_L(i)$, (b) $Q^{\mathbb{N}} \subseteq \text{val}(V^{\mathbb{N}})$ is the state-space, (c) $\Theta^{\mathbb{N}} \subseteq Q^{\mathbb{N}}$ is the set of initial states, and (d) $T^{\mathbb{N}} \subseteq Q^{\mathbb{N}} \times Q^{\mathbb{N}}$ is the transition relation, which is partitioned into sets of discrete transitions $\mathcal{D}^{\mathbb{N}} \subseteq Q^{\mathbb{N}} \times Q^{\mathbb{N}}$ and continuous trajectories $\mathcal{T}^{\mathbb{N}} \subseteq Q^{\mathbb{N}} \times Q^{\mathbb{N}}$.

A state \mathbf{x} of $\mathcal{A}^{\mathbb{N}}$ is a valuation of *all* the variables in $V^{\mathbb{N}}$ and is denoted by boldface \mathbf{v} , \mathbf{v}' , etc. The set of all states is called the state-space and is denoted $Q^{\mathbb{N}}$. If a state $\mathbf{v} \in Q^{\mathbb{N}}$ satisfies a formula ϕ —that is, the corresponding variable valuations result in ϕ evaluating to *true*—we write $\mathbf{v} \models \phi$. For a formula ϕ with $\text{vars}(\phi) \subseteq V(i)$, the corresponding states $\mathbf{x} \in Q^{\mathbb{N}}$ satisfying ϕ are $\llbracket \phi \rrbracket \triangleq \{\mathbf{x} \in Q^{\mathbb{N}} \mid \mathbf{v} \models \phi\}$. For instance, the initial states $\Theta^{\mathbb{N}} \triangleq \llbracket \text{Init}(i) \rrbracket$ are the states satisfying `Init`(i). For some state \mathbf{v} , the valuation of a particular local variable $x[i] \in V_L(i)$ for automaton $\mathcal{A}(i)$ is denoted by $\mathbf{v}.x[i]$, and $\mathbf{v}.g$ for some global variable $g \in V_G(i)$. For a set of variables V , the valuations of each $v \in V$ at state \mathbf{v} is denoted by $\mathbf{v}.V$. For a formula ϕ and a set of variables $V \subseteq \text{vars}(\phi)$, let $\phi \downarrow V$ be the projection of ϕ onto the variables V , such that $\text{vars}(\phi \downarrow V) = V$ and $\llbracket \phi \rrbracket \subseteq \llbracket \phi \downarrow V \rrbracket$, which can be computed by eliminating the existential quantifiers from the formula $\exists \text{vars}(\phi) \setminus V : \phi$. The evolution of the states of $\mathcal{A}^{\mathbb{N}}$ are describing by a transition relation $T^{\mathbb{N}} \subseteq Q^{\mathbb{N}} \times Q^{\mathbb{N}}$. For a pair $(\mathbf{v}, \mathbf{v}') \in T^{\mathbb{N}}$, we use the notation $\mathbf{v} \rightarrow \mathbf{v}'$, where \mathbf{v} is called the *pre-state* and \mathbf{v}' is called the *post-state*. There are two ways variables may be updated by $T^{\mathbb{N}}$. Discrete transitions $\mathcal{D}^{\mathbb{N}}$ model instantaneous changes and continuous trajectories $\mathcal{T}^{\mathbb{N}}$ model evolution over a real time interval.

When necessary to disambiguate state updates owing to either discrete transitions and continuous trajectories, we write $\mathbf{v} \rightarrow_{\mathcal{D}^N} \mathbf{v}'$ or $\mathbf{v} \rightarrow_{\mathcal{T}^N} \mathbf{v}'$, respectively.

Discrete Transitions. Discrete transitions model atomic, instantaneous updates of state due to *one* automaton in the network \mathcal{A}^N . Informally, a discrete transition from pre-state \mathbf{v} to post-state \mathbf{v}' models the discrete transition of one particular hybrid automaton $\mathcal{A}(i)$ by some transition $t \in \text{Trans}(i)$. There is a discrete transition $\mathbf{v} \rightarrow \mathbf{v}' \in \mathcal{D}^N$ iff: $\exists i \in [N] \exists t \in \text{Trans}(i) : \mathbf{v}.V(i) \models \mathbf{grd}(t, i) \wedge \mathbf{v}'.V(i) \models \mathbf{rst}(t, i) \wedge (\forall j \in [N] : j \neq i \Rightarrow \mathbf{v}'.V(j) = \mathbf{v}.V(j))$. From the pre-state \mathbf{v} , any automaton $\mathcal{A}(i)$ in the network \mathcal{A}^N that has some transition where \mathbf{v} satisfies its guard *may* update its post-state according to the transition's reset, while the variable valuations of all the other automata in \mathcal{A}^N remain unchanged.²

Continuous Trajectories. Continuous trajectories model update of state over intervals of real time. Informally, there is a trajectory $\mathbf{v} \rightarrow \mathbf{v}' \in \mathcal{T}^N$ iff some amount of time— t_e —can elapse from \mathbf{v} , such that, (a) the states of *all* automata in the network \mathcal{A}^N are updated to \mathbf{v}' according to their individual trajectory statements, and (b) while ensuring the invariants of *all* automata along the entire trajectory. Formally, trajectories are defined as solutions of differential equations or inclusions specified in the trajectory statements of $\mathcal{A}(i)$. For a state \mathbf{v} , a location m , a real time t , and a real variable $v \in X(i)$, let $\mathbf{flow}(\mathbf{v}, m, v, t) = \mathbf{v}.v + \int_{\tau=0}^t \mathbf{frate}(m, v) d\tau$. Since \mathbf{frate} may specify a differential inclusion, \mathbf{flow} is a set-valued function. There is a trajectory $\mathbf{v} \rightarrow \mathbf{v}' \in \mathcal{T}^N$ iff: $\exists t_e \in \mathbb{R}_{\geq 0} \forall t_p \in \mathbb{R}_{\geq 0} \forall i \in [N] \exists m \in L : t_p \leq t_e \wedge \mathbf{flow}(\mathbf{v}, m, X(i), t_p) \models \mathbf{inv}(m, i) \wedge \mathbf{v}'.X(i) \in \mathbf{flow}(\mathbf{v}, m, X(i), t_e)$. For each $i \in [N]$ and each real variable $x[i]$, $\mathbf{v}.x[i]$ must evolve to the valuations $\mathbf{v}'.x[i]$, in exactly t_e time in some location $m \in L$ according to the flow rates allowed for $x[i]$ in location m . In addition, all intermediate states along the trajectory must satisfy the invariant $\mathbf{inv}(m, i)$.

Executions and Invariants. An execution of the network \mathcal{A}^N models a particular behavior of all the automata in the network. An *execution* of \mathcal{A}^N is a sequence of states $\alpha = \mathbf{v}_0, \mathbf{v}_1, \dots$ such that $\mathbf{v}_0 \in \Theta^N$, and for each index k appearing in the sequence, $(\mathbf{v}_k, \mathbf{v}_{k+1}) \in T^N$. A state \mathbf{x} is *reachable* if there is a finite execution ending with \mathbf{x} . The set of reachable states for \mathcal{A}^N is $\text{Reach}(\mathcal{A}^N)$. The set of reachable states for \mathcal{A}^N starting from an arbitrary subset $\mathbf{V}_0 \subseteq Q^N$ is $\text{Reach}(\mathcal{A}^N, \mathbf{V}_0)$. An *invariant* for \mathcal{A}^N is any set of states that contains $\text{Reach}(\mathcal{A}^N)$.

3 Anonymized State-Space Representation

For any fixed $N \in \mathbb{N}$, let i be a symbol representing an arbitrary element of $[N]$, and for the hybrid automaton template $\mathcal{A}(N, i)$, the composed automaton modeling a network of size N is \mathcal{A}^N (Definition 2). We present an algorithm for computing $\text{Reach}(\mathcal{A}^N)$ that takes advantage of the symmetries in the template $\mathcal{A}(i)$ instantiated in \mathcal{A}^N . The representation of $\text{Reach}(\mathcal{A}^N)$ is *anonymized*, so numerical automaton indices— $1, 2, \dots, N$ —are not explicitly enumerated and are instead modeled using symbolic indices— $i_1,$

² The guard may depend on the variables of some automaton $j \neq i$, so automata may communicate via global variables and local variables, for details, see [16, Chapter 2].

i_2, \dots, i_N . Frequently, the number of symbolic indices needed to represent equivalent states is significantly smaller than the number N of numerical indices. For example, in MUX-INDEX-RECT (Figure 1), a single symbolic index is sufficient independent of N . For a given state $\mathbf{x} \in Q^N$, the set of corresponding states $\mathbf{X} \subseteq Q^N$ that are equivalent modulo indices is obtained by substituting any numerical index i of all local variables $v[i] \in \mathcal{V}_L(i)$ with a symbolic index j with type $[N]$.

Definition 3. *Two states $\mathbf{x}, \mathbf{x}' \in Q^N$ of \mathcal{A}^N , are equivalent modulo indices if there exists a bijection $\pi : [N] \rightarrow [N]$ such that for each $v[i] \in \mathcal{V}(i)$, $\mathbf{x}.v[i] = \mathbf{x}'.v[\pi(i)]$. For a state $\mathbf{x} \in Q^N$ of \mathcal{A}^N , the set of states $\varepsilon(\mathbf{x})$ that is equivalent modulo indices to \mathbf{x} is: $\varepsilon(\mathbf{x}) \triangleq \{\mathbf{x}' \in Q^N \mid \mathbf{x} \text{ and } \mathbf{x}' \text{ are equivalent modulo indices}\}$.*

We note this is the same type of definition as the existence of an automorphism used in [6, 9, 14]. A state is equivalent modulo indices to itself by picking the bijection π to be the identity mapping. For a formula ϕ , we will overload π and write $\pi(\phi)$, which modifies ϕ by applying π to each index variable $i \in \text{ivars}(\phi)$. The anonymized representation takes this idea a step further by utilizing symbolic names for process indices along with counters, and a formula representing the valuations of any global variables. We use the (\cdot) notation to refer to particular elements of tuples. For example, $C.\text{Count}$ refers to the count of anonymized class C , $C.\text{Form}$ refers to C 's formula, etc. If C is clear from context, we refer to $C.\text{Count}$ as Count , etc.

Definition 4. *An anonymized state S of the network \mathcal{A}^N is a tuple $\langle \text{Classes}, G \rangle$, where:*

- (a) *Each anonymized class $C \in \text{Classes}$ is a tuple $C \triangleq \langle \text{Count}, I, \text{Form} \rangle$, where:*
- (i) *Form is a quantifier-free formula over the variables $\mathcal{V}_L(i_1) \cup \dots \cup \mathcal{V}_L(i_I)$, where i_1, \dots, i_I are I distinct symbolic index variables.*
 - (ii) *$I \geq 1$ is a natural number called the class's rank, which is equal to the number of distinct symbolic index variables appearing in Form : $I \triangleq |\text{ivars}(\text{Form})|$.*
 - (iii) *Count is a natural number called the class's count, and satisfies $N \geq \text{Count} \geq |I|$. The count is the number of automata of class C . Additionally, the sum of all the class counts in S equals N : $N = \sum_{C \in S.\text{Classes}} C.\text{Count}$, where $C.\text{Count}$ is the count of class C .*
- (b) *G is a quantifier-free formula over the global variables \mathcal{V}_G .*

For an anonymized class C , requirement (iii) of Definition 4 that $\text{Count} \geq |I|$ means the number of automata satisfying Form is at least the rank (the number of distinct index variables appearing in Form). When the rank I is clear from context, we drop it from the C tuple and write $\langle \text{Count}, \text{Form} \rangle$. We say two anonymized classes C_1 and C_2 over the same symbolic indices ($\text{ivars}(C_1.\text{Form}) = \text{ivars}(C_2.\text{Form})$) are equivalent and write $C_1 \equiv C_2$ iff they have equivalent class formulas and equal class counts:³

³ It is possible for classes with different ranks to represent the same states. For example, consider states arising from MUX-INDEX-RECT, $S_1 = \langle \{\langle 2, 2, q[i_1] = \text{rem} \wedge q[i_2] = \text{rem} \rangle\}, g = \perp \rangle$ and $S_2 = \langle \{\langle 2, 1, q[i_1] = \text{rem} \rangle\}, g = \perp \rangle$, which both represent there are two automata with location rem and g is \perp , i.e., $\llbracket S_1 \rrbracket = \llbracket S_2 \rrbracket$. However, a class of a particular rank may not be expressible as a different rank. For example, there is no way to express the following using rank 1 classes: $\langle \{\langle 2, 2, q[i_1] = \text{rem} \wedge q[i_2] = \text{rem} \wedge x[i_1] \geq x[i_2] \rangle\}, g = \perp \rangle$, which expresses that there are two automata in rem with one's clock at least as large as the other's.

Definition 5. Two classes C_1 and C_2 are equivalent, written $C_1 \equiv C_2$ iff $C_1.\text{Count} = C_2.\text{Count} \wedge C_1.\text{Form} \equiv C_2.\text{Form}$.

Here, equivalence between the class formulas is a semantic and not syntactic notion, and means the formula $C_1.\text{Form} \equiv C_2.\text{Form}$ is valid. We say two anonymized states S_1 and S_2 are equivalent and write $S_1 \equiv S_2$ iff they have the same state counts, the classes in their sets of classes are equivalent, and their global formulas are equivalent. See Footnote 3 for an example from MUX-INDEX-RECT.

Definition 6. Two anonymized states S_1 and S_2 are equivalent, written $S_1 \equiv S_2$, iff $\forall C_1 \in S_1.\text{Classes} \exists C_2 \in S_2.\text{Classes} C_1 \equiv C_2 \wedge G_1 \equiv G_2$.

We make the following assumption about the format of class formulas.

Assumption 1. For an anonymized state S , for each class $C \in \text{Classes}$, the class formula $C.\text{Form}$ is in conjunctive normal form (CNF). For each index $i \in \{i_1, \dots, i_{c.I}\}$, $C.\text{Form}$ contains an equality $q[i] = \text{loc}$ for some location $\text{loc} \in L$.

For example, Equation 1 (arsing from MUX-INDEX-RECT) satisfies this assumption. This assumption ensures that each class corresponds to a concrete state, and has a control location specified to determine the transitions and trajectories that may be possible (recall Definition 1). Under Assumption 1, the interpretation of an anonymized state S corresponds to a set of states of Q^N , which we write as $\llbracket S \rrbracket$ and define formally next. Since the class formulas of S are over the variables of automata with symbolic indices, the interpretation instantiates the symbolic indices with specific elements of $[N]$, which yields the set of states that are equivalent modulo indices.

Definition 7. For an anonymized state

$$S = \langle \underbrace{\{\langle \text{Count}_1, I_1, \text{Form}_1 \rangle\}}_{C_1}, \dots, \underbrace{\{\langle \text{Count}_k, I_k, \text{Form}_k \rangle\}}_{C_k}, G \rangle,$$

we instantiate the set of symbolic indices $\{i_1, \dots, i_{I_k}\}$ with all possible values in $[N]$ as follows. A consistent partition of $[N]$,

$$P = \left\{ \underbrace{\{P_1^1, \dots, P_1^{I_1}\}}_{P_1}, \dots, \underbrace{\{P_k^1, \dots, P_k^{I_k}\}}_{P_k} \right\},$$

is a partition of $[N]$, such that, for any $P_j \in P$, (a) $|P_j| = \text{Count}_j$ and (b) P_j is partitioned into I_j sets $P_j^1, \dots, P_j^{I_j}$ (and we recall that I_j is the rank of C_j).

For a consistent partition P , we note that (a) $\sum_{P_j \in P} |P_j| = N$, since P partitions $[N]$, and (b) $\text{Count}_j \geq I_j$ (by Definition 4, (iii)). For example, consider the anonymized state (arsing from MUX-INDEX-RECT, Figure 1) with count three and rank two:

$$\langle \{ \langle 3, 2, q[i_1] = \text{try} \wedge q[i_2] = \text{rem} \wedge x[i_1] \geq x[i_2] + B \rangle \}, g = i_1 \rangle. \quad (1)$$

One consistent partition is: $P = \{P_1, P_2\}$ where $P_1 = \{1\}$ and $P_2 = \{2, 3\}$. The set $\{\{1, 2, 3\}\}$ is *not* a consistent partition since it is partitioned into one set, but the rank $I = 2$, and Definition 7 requires each $P_j \in P$ be partitioned into I_j partitions.

For an anonymized state S , the set of consistent partitions $ConsPart(S)$ are all consistent partitions of $[N]$. Continuing MUX-INDEX-RECT for Equation 1, $ConsPart(S)$ is $\{\{\{1\}, \{2, 3\}\}, \{\{2\}, \{1, 3\}\}, \{\{3\}, \{1, 2\}\}, \{\{1, 2\}, \{3\}\}, \{\{1, 3\}, \{2\}\}, \{\{2, 3\}, \{1\}\}\}$. All these partitions define the set of states $\llbracket S \rrbracket$ the anonymized state S represents. This is the same as all the states equivalent modulo indices to the states $\llbracket S_P \rrbracket$ for a particular consistent partition P .

Definition 8. For an anonymized state S and a consistent partition $P \in ConsPart(S)$, the set of states of network \mathcal{A}^N represented by S corresponding to P are:

$$\llbracket S_P \rrbracket \triangleq \{ \mathbf{x} \in Q^N \mid \mathbf{x} \models G \wedge Form_1(P_1) \wedge \dots \wedge Form_k(P_k) \}, \quad (2)$$

where each $Form_j(P_j) \triangleq \forall i_j^1 \in P_j^1, \dots, i_j^{I_j} \in P_j^{I_j} : Form_j(i_j^1, \dots, i_j^{I_j})$. The set of states of network \mathcal{A}^N represented by S with all consistent partitions is:

$$\llbracket S \rrbracket \triangleq \bigcup_{P \in ConsPart(S)} \llbracket S_P \rrbracket. \quad (3)$$

We have written $Form_j(i_j^1, \dots, i_j^{I_j})$ to highlight that $Form_j$ is over I_j symbolic index variables. Note that $Form_j(P_j)$ is equivalent to a finite-length conjunction since each $P_j^{I_j}$ is a finite set. The next lemma states that this definition of interpretations of anonymized states yields the same set of states as equivalence modulo identifiers.

Lemma 1. For an anonymized state S , for any $\mathbf{x} \in \llbracket S \rrbracket$, for any $\mathbf{x}' \in \varepsilon(\mathbf{x})$, $\mathbf{x}' \in \llbracket S \rrbracket$.

Continuing the MUX-INDEX-RECT Equation 1 example with the consistent partition $P = \{\{1\}, \{2, 3\}\}$, the states represented by S_P are:

$$\begin{aligned} \llbracket S_P \rrbracket &= \{ \mathbf{x} \in Q^3 \mid \mathbf{x} \models \forall i_1^1 \in P_1^1, i_1^2 \in P_1^2 : q[i_1] = \text{try} \wedge q[i_2] = \text{rem} \wedge \\ &\quad x[i_1] \geq x[i_2] + B \wedge g = i_1 \} \\ &= \{ \mathbf{x} \in Q^3 \mid \mathbf{x} \models (q[1] = \text{try} \wedge q[2] = \text{rem} \wedge q[3] = \text{rem} \wedge \\ &\quad x[1] \geq x[2] + B \wedge x[1] \geq x[3] + B \wedge g = 1) \}. \end{aligned}$$

Applying Lemma 1, $\llbracket S \rrbracket = \varepsilon(\llbracket S_P \rrbracket)$.

4 Anonymized Reachability of Hybrid Automata Networks

Next we describe an on-the-fly algorithm for overapproximating the reachable states of a network \mathcal{A}^N using anonymized states. We note that the CNF requirement (Assumption 1) is not restrictive: if a new class is created during the execution of the algorithm that contains disjunctions, it is split into multiple classes with CNF formulas. Recall from Section 2.1, that $\phi \downarrow V$ is the projection of ϕ onto the variables V .

Pseudocode for the reachability algorithm, areach appears in Figure 2. The algorithm operates on frontiers of reachable states represented by Frontier, which is initialized (line 3) to a singleton set with one class with count N and formula $\text{Init}(i) \downarrow V_L(i)$, which is $\text{Init}(i)$ projected onto the local variables. The global formula is initialized with


```

1  function areach( $\mathcal{A}(i)$ , Init( $i$ ), N)
   AnonReach  $\leftarrow$   $\emptyset$ 
3   Frontier  $\leftarrow$   $\{\{\langle N, \text{Init}(i) \downarrow V_L(i) \rangle, \text{Init}(i) \downarrow V_G(i) \rangle\}$  // initial anonymized state
   while Frontier  $\neq$   $\emptyset$  // repeat until no new states are added to the frontier
5     FrontierNew  $\leftarrow$   $\emptyset$  // initialize next frontier
     AnonReach  $\leftarrow$  AnonReach  $\cup$  Frontier // add frontier to reachable states
     // compute successors of each anonymized state in the frontier
     foreach anonymized state S in Frontier
9       FrontierNew  $\leftarrow$  FrontierNew  $\cup$  discPost(S) // Figure 4
       FrontierNew  $\leftarrow$  FrontierNew  $\cup$  contPost(S) // Figure 5
11      FrontierNew  $\leftarrow$  mergeAndDrop(FrontierNew, AnonReach) // Figure 3
       Frontier  $\leftarrow$  FrontierNew
13  return AnonReach

```

Fig. 2. On-the-fly anonymized reachability algorithm. The inputs are an automaton template $\mathcal{A}(i)$, an initial condition $\text{Init}(i)$, and a constant natural number N . The anonymized reachable states AnonReach are computed as a fixed-point starting from the anonymized initial states.

$\text{Init}(i) \downarrow V_G(i)$, which is $\text{Init}(i)$ projected onto the global variables. The set of reachable anonymized states computed so far is the set AnonReach . Next (line 4), we remove an anonymized state S from Frontier , compute anonymized post-states from S , and continue until no new anonymized states are added to Frontier . Anonymized post-states are added to the frontier using the set $\text{Frontier}_{\text{New}}$ (line 5). Computing successors (post-states)—the states reachable from S in one step—is composed of two parts: (a) computing the discrete successors corresponding to transitions (line 9), and (b) computing the continuous successors corresponding to trajectories (line 10).

Equivalent Class Merging Subroutine. We first describe the `mergeAndDrop` subroutine (Figure 3). It takes a set of anonymized states $\text{Frontier}_{\text{New}}$ and returns a set of anonymized states guaranteed to (a) not have any equivalent classes (lines 7 through 8) and (b) be new (not already represented in AnonReach) (line 3). Invariant 1 states no two class formulas in any reachable anonymized state are equivalent, and Invariant 2 states no two anonymized states in AnonReach are equivalent (Definition 6).

Invariant 1. For any $S \in \text{AnonReach}$, $C_1, C_2 \in S.\text{Classes}$, $C_1.\text{Form} \not\equiv C_2.\text{Form}$.

Invariant 2. For any distinct $S_1, S_2 \in \text{AnonReach}$, $S_1 \not\equiv S_2$.

Discrete Successors. The function `discPost` (Figure 4) computes the discrete successors from an anonymized state S in the frontier (Figure 2, line 9). The post-states $\text{States}_{\text{New}}$ are added to the frontier $\text{Frontier}_{\text{New}}$. First, we iterate over each class C in $S.\text{Classes}$ (line 3), and then we iterate over each index variable i in the set of index variables in the class formula, $\{i_1, \dots, i_{C.I}\}$ (line 5). Next, we iterate over the (syntactic) transitions $\text{Trans}(i)$ of $\mathcal{A}(i)$ (line 6). For a transition $t \in \text{Trans}(i)$ and an anonymized class C , line 7

```

1  function mergeAndDrop(FrontierNew, AnonReach)
   foreach S in FrontierNew
3     if S  $\in$  AnonReach then FrontierNew = FrontierNew  $\setminus$  {S}
   else
5     foreach distinct pair of anonymized classes (C1, C2) in S.Classes
       if  $\neg(C_1.\text{Form} \equiv C_2.\text{Form})$  is UNSAT then
7         C1.Count  $\leftarrow$  C1.Count + C2.Count // if equivalent, sum counts
         S.Classes  $\leftarrow$  S.Classes  $\setminus$  {C2} // if equivalent, drop equivalent class
9  return FrontierNew

```

Fig. 3. `mergeAndDrop` combines classes with equivalent class formulas and sums their counts

```

1  function discPost(S)
   StatesNew ← ∅
3  foreach anonymized class C in S.Classes
   Vs ← V'(i1) ∪ ... ∪ V'(iC,I)
5  foreach symbolic index i in iwars(Vs)
   foreach transition t in Trans(i)
7     CNew.Form ← (C.Form ∧ S.G ∧ grd(t, i) ∧ rst(t, i)) ↓ V'(i) // make post-state class
   // substitute primed variables with unprimed variables
9     CNew.Form ← Substitute(CNew.Form, V'(i), V(i))
   // project onto global variables for global constraint
11    SNew.G ← CNew.Form ↓ VG(i)
   // project onto local variables for local constraint
13    (CNew.Count, CNew.I, CNew.Form) ← ⟨1, 1, CNew.Form ↓ VL(i)⟩
   SNew.Classes ← S.Classes \ {C} // remove pre-state from post-state classes
15    // add pre-state class to post-state classes if count at least rank
   if C.Count > C.I then SNew.Classes ← S.Classes ∪ {(C.Count - 1, C.I, C.Form)}
17    // otherwise, pre-state class no longer exists (count less than rank)
   else SNew.Classes ← S.Classes ∪ {(C.Count - 1, C.I - 1, C.Form ↓ Vs \ V(i))}
19    SNew.Classes ← SNew.Classes ∪ {CNew} // add class to post-state
   StatesNew ← StatesNew ∪ {SNew}
21  return StatesNew

```

Fig. 4. `discPost` computes the post-states of an anonymized state S due to discrete transitions for an automaton with index i and states satisfying C 's formulas.

computes the subsequent class from C by transition t , made by the automaton with index i . This computation can be carried out using quantifier elimination procedures over the types of the variables appearing in the guard and reset of the transition t , and then syntactically unpriming all primed variables (representing successors) following quantifier elimination using `Substitute` (line 9). This step is an overapproximation, since it computes the successors of each class regardless of the number of automata with states satisfying the anonymized class formula `Form`, and just presumes *there is some* automaton with variable valuations satisfying `Form`.

The anonymized post-state S_{New} is constructed using the classes of the anonymized pre-state S along with the new anonymized class, C_{New} (lines 14 through 19). First, the classes for S_{New} are set to be the anonymized classes of S , without the anonymized class of the current iteration, C (line 14). Next, if the class count of C is larger than its rank, then it is added to the classes of the post-state, with its count reduced by one to indicate some automaton has left the set of states satisfying the corresponding class formula (line 16). On the other hand, if the class count is equal or less than its rank, then the pre-state's anonymized class C would no longer satisfy the requirements of Definition 4, (iii), so its class formula is projected onto the variables of all automata except those of automaton i , the one making a transition (line 18). If a class has count or rank equal to 0, then it is removed. This process may result in two classes with equivalent formulas, since the algorithm has not yet detected if any other classes had the same formula and presumed the post-state class C_{New} had a count of one, which is why we use `mergeAndDrop` (Figure 2, line 11).

Lemma 2. (*Discrete Successor Soundness*) *For an anonymized state S , for any corresponding concretized state $\mathbf{x} \in \llbracket S \rrbracket$, if $\mathbf{x} \rightarrow_{\mathcal{D}^N} \mathbf{x}'$, then $\mathbf{x}' \in \llbracket \text{discPost}(S) \rrbracket$.*

Continuous Successors. An overapproximation of continuous successors are computed using `contPost`—shown in Figure 5—called from `symreach` (Figure 2, line 10). For an anonymized state S in the frontier, `contPost` computes an overapproximation of the post-states from S owing to the individual trajectories of all automata in the network for up to the most amount of time that can elapse before any invariant is

```

1  function contPost(S)
    $V_S \leftarrow \bigvee_G$ 
3  // formula to encode trajectories for all automata in the network
    $pf \leftarrow (t_e > 0 \wedge S.G)$ 
5  foreach anonymized class C in S.Classes // iterate over each class in pre-state
    $V_S \leftarrow V_S \cup \bigvee'(i_1) \cup \dots \cup \bigvee'(i_{C.I})$ 
7    $pf \leftarrow pf \wedge C.Form$  // encode pre-state class formula
   // determine locations any automaton may be in (recall Assumption 1)
9   foreach location loc in L
      foreach  $i$  in  $\{i_1, \dots, i_{C.I}\}$  // iterate over all indices (ranks)
11      if  $C.Form \not\models (q[i] = loc)$  is UNSAT then // use loc if automaton  $i$  is in loc
          // add the trajectory semantics overapproximating the post-states
13           $pf \leftarrow pf \wedge \text{inv}(loc, i) \wedge X(i) \in \text{flow}(pf, loc, X(i), t_e)$ 
15   $pf \leftarrow pf \downarrow V_S$ 
    $pf \leftarrow \text{Substitute}(pf, \bigvee'(i), V(i))$ 
    $S_{New} \leftarrow \text{RemapClasses}(S, pf)$  // Figure 6
17  return  $S_{New}$ 

```

Fig. 5. contPost function that computes the continuous successors from an anonymized state S

```

1  function RemapClasses(S, pf)
    $S_{New}.Classes = \emptyset$ 
3  foreach anonymized class C in S.Classes
   // project pf onto variables of indices in each pre-state class
5    $V_S \leftarrow \bigvee(i_1) \cup \dots \cup \bigvee(i_{C.I})$ 
   // create new class with post-state formula and copy pre-state count
7    $(C_{New}.Count, C_{New}.I, C_{New}.Form) \leftarrow (C.Count, C.I, pf \downarrow V_S)$ 
    $S_{New}.Classes \leftarrow S_{New}.Classes \cup C_{New}$  // add post-state class to classes
9   $S_{New}.N \leftarrow S.N$ 
   return  $S_{New}$ 

```

Fig. 6. RemapClasses uses pf and the pre-state indices, class counts, and ranks to create the anonymized post-state S_{New} . It first projects onto variables with indices of each class in the pre-state and then uses the pre-state to ensure class counts and ranks remain constant over trajectories.

violated. The anonymized state specifies a location $loc \in L$ for each automaton in the network (recall Assumption 1). Each location loc specifies a trajectory statement, so trajectories are defined for each automaton in the network. Each new anonymized state $S_{New} \in \text{States}_{New}$ computed corresponds to the trajectory semantics updating the real variables of *all* automata in the network \mathcal{A}^N . The variable pf encodes the trajectory semantics of all automata in the network \mathcal{A}^N (line 4), which is initially the constraint $t_e > 0$, indicating that some positive real amount of time t_e will elapse. However, for an anonymized state S , for distinct anonymized classes C_1, C_2 in $S.Classes$, the symbolic indices appearing in the formulas may be equal, i.e., $\exists i \in \text{ivars}(C_1)$ and $\exists j \in \text{ivars}(C_2)$ such that $i = j$. Since pf encodes the states of all automata in the network, the symbolic index variables appearing in any class formula of any anonymized class must be distinct. Rather than performing these tedious syntactic manipulations, we assume that for an anonymized state S , for distinct classes C_1, C_2 in $S.Classes$, $\forall i \in \text{ivars}(C_1), \forall j \in \text{ivars}(C_2)$, we have $i \neq j$.⁴

Each anonymized class formula $C.Form$ of an anonymized state S specifies the location(s) of the automata, so the first step is to determine the dynamics that will modify each class formula. This is accomplished by first determining the appropriate flow-rate conditions to use for each class in $S.Classes$, which can be detected by finding which

⁴ This is a tedious, but trivial invariant that we maintain in our implementation in *Passel*, so we make this assumption for clarity of presentation only.




Fig. 7. Memory usage comparison of PHAVer and *Passel*'s anonymized reachability. Vertical axis scale is logarithmic and has units of megabytes, and horizontal axis is number of automata, N .

Form implies the location variable $q[i]$ is in some location $\mathbf{loc} \in L$. If the control location of automaton i is found to be equal to location \mathbf{loc} , then the trajectory statement of location \mathbf{loc} is used to define the semantics of the time-evolution of i 's real variables (line 13). The semantics of trajectories result in *all* the automata's real variables evolving over time t_e , so the formula encoding the trajectory statements of all automata is conjuncted (line 13). The post-states are computed by projecting onto the primed variables of all classes, and then renaming primed variables with their unprimed counterparts (line 15).⁵ We call `RemapClasses` with the pre-state S and `pf`, which encodes the post-state constraints, to recreate classes from sub-formulas of `pf` (Figure 6 called at line 16). This is done to ensure the class counts are constant when computing post-states due to trajectories.

Lemma 3. (*Continuous Successor Soundness*) *For an anonymized state S , for any corresponding concretized state $\mathbf{x} \in \llbracket S \rrbracket$, if $\mathbf{x} \rightarrow_{\mathcal{T}^N} \mathbf{x}'$, then $\mathbf{x}' \in \llbracket \text{contPost}(S) \rrbracket$.*

The next invariant states the sum of all class counts equals N . It follows from the definitions of `discPost` and `contPost`, since `discPost` always decreases class counts by the same amount it increases them—so the sum remains invariant—and `contPost` does not change class counts (only formulas). Additionally, `mergeAndDrop` changes class counts, but their sum remains the same since it removes any duplicate classes after adding their counts (Figure 3, lines 7 through 8).

Invariant 3. *For any $S \in \text{AnonReach}$, $N = \sum_{C \in S.\text{Classes}} C.\text{Count}$.*

Theorem 1 states soundness of the algorithm: the concretization of the anonymized reachable states `AnonReach` contains the reachable states for network \mathcal{A}^N . It follows from Lemmas 2 and 3. The approximation comes from: (a) transitions are allowed as long as *some* automaton satisfies a guard, (b) index-typed variables are abstracted to be equal or not equal only, and (c) rectangular dynamics are overapproximated.

⁵ This may result in a DNF formula, and if so, each conjunctive clause is added as a new anonymized state by iterating over the conjunctive clauses so all class formulas are CNF.

Theorem 1. (Soundness) For a fixed $N \in \mathbb{N}$, for the network \mathcal{A}^N composed of N instantiations of the template $\mathcal{A}(N, i)$, the anonymized reachable states AnonReach computed by areach overapproximate the reachable states of \mathcal{A}^N : $\text{Reach}(\mathcal{A}^N) \subseteq \llbracket \text{AnonReach} \rrbracket$.

5 Experimental Results

The anonymized reachability algorithm is implemented in *Passel* [16, 18, 19]. The current implementation of *Passel* uses the SMT solver Z3 [7] for proving validity, checking satisfiability, and performing quantifier elimination. *Passel* is written in C# and uses the managed .NET API to Z3, with experimental results reported using version 4.1. *Passel* proves validity of a formula ϕ by checking unsatisfiability of $\neg\phi$. The variables $V(i)$ used in defining $\mathcal{A}(i)$ are specified to the SMT solver. Each local variable $v[i] \in V_L(i)$ is modeled as an uninterpreted function $v : [N] \rightarrow \text{type}(v)$. *Passel* automatically generates and asserts trivial data-type lemmas that the SMT solver requires. The experiments were conducted in an Ubuntu 12.04 VMWare virtual machine with 4 GB RAM allocated running *Passel* through Mono, executed on a modern laptop with a quad-core Intel i7 processor running Windows 8 with 16 GB RAM physically available. For comparison purposes, we evaluated *Passel*, PHAVer (version 0.38), and SpaceEx (version 0.9.8b). We do not present results for SpaceEx, as the only scenario—out of the PHAVer, LeGuernic-Girard (LGG), and STC scenarios—that can compute the reachable states of systems with rectangular differential inclusion dynamics ($\dot{x} \in [a, b]$ for real constants $a \leq b$) adequately is the PHAVer scenario, so the results are equivalent.

Figures 7 and 8 show, respectively, a runtime and memory usage comparison between PHAVer and *Passel* for several examples as a function of N , the number of




Fig. 8. Runtime comparison of PHAVer and *Passel*'s anonymized reachability. Vertical axis is logarithmic and has units of seconds, and horizontal axis is number of automata, N .

automata.⁶ The examples include several timed mutual exclusion algorithms (such as MUX-INDEX-RECT from Figure 1), a simplified SATS model [16, 18, 19], and several purely discrete examples. All properties were safety properties (invariants), such as mutual exclusion, separation (collision avoidance) in SATS, etc. Comparing all the examples, the anonymized reachability method implemented in *Passel* allows us to compute the reachable states of networks composed of many more automata than PHAVer, which runs out of memory on all examples for $N \geq 11$. The experimental results indicate the primary advantage is reduced memory growth. Even for networks of tens of automata, in all examples, *Passel* never uses more than a few hundred megabytes of memory as shown in Figure 7.⁷ For protocols that are highly asymmetric, the worst-case asymptotic memory growth may be exponential. The runtime required by *Passel* could be reduced by performing some operations more efficiently in the implementation—particularly the checks to determine if a new anonymized state representation is actually new or not—which we plan to implement for future work.

For MUX-INDEX-RECT, PHAVer runs out of memory for $N \geq 8$. As shown in Figures 7 and 8, for $N = 7$, PHAVer uses over 1.3 GB memory and completes in over 3 hours, while *Passel* uses over an order of magnitude less memory at about 70 MB and nearly four orders of magnitude less runtime at about three seconds. Because of the anonymized representation, *Passel* is able to compute the reachable states of $N = 30$ in a few seconds using about 70 MB memory, and we have experimented successfully up to hundreds and even thousands of automata for this example.

6 Summary

In this paper, we present an on-the-fly forward reachability algorithm that computes an anonymized representation of the reachable states for hybrid automata networks consisting of N instantiations of a template $\mathcal{A}(N, i)$. The anonymized representation uses symbolic automata indices instead of explicit ones to avoid generating all permutations of automata indices and states. We showed it to be effective at computing the reachable states of networks with tens of automata for several examples, with significantly lower memory usage than PHAVer. The restriction to rectangular inclusion dynamics is due in part to *Passel*'s implementation, but a future direction is to evaluate the anonymized reachability method on examples with linear and nonlinear dynamics.

Acknowledgments. The authors are grateful for the anonymous reviewers' feedback. This material is based upon work supported by the National Science Foundation under Grant No. NSF CNS 10-54247 CAR. This work was supported by the Air Force Office of Scientific Research Young Investigator Program Award FA9550-12-1-0336.

⁶ *Passel* and the examples may be downloaded from:

<https://publish.illinois.edu/passel-tool/>.

⁷ For small N , PHAVer uses less memory than *Passel* because *Passel* must load runtime components (e.g., the .NET framework via Mono) and libraries (e.g., Z3).

References

1. Basler, G., Mazzucchi, M., Wahl, T., Kroening, D.: Symbolic counter abstraction for concurrent software. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 64–78. Springer, Heidelberg (2009)
2. Behrmann, G., Bouyer, P., Fleury, E., Larsen, K.G.: Static guard analysis in timed automata verification. In: Garavel, H., Hatcliff, J. (eds.) TACAS 2003. LNCS, vol. 2619, pp. 254–270. Springer, Heidelberg (2003)
3. Bengtsson, J., Larsen, K., Larsson, F., Pettersson, P., Yi, W.: UPPAAL: A tool suite for automatic verification of real-time systems. In: Alur, R., Sontag, E.D., Henzinger, T.A. (eds.) HS 1995. LNCS, vol. 1066, pp. 232–243. Springer, Heidelberg (1996)
4. Bogomolov, S., Herrera, C., Muñoz, M., Westphal, B., Podelski, A.: Quasi-dependent variables in hybrid automata. In: 17th International Conference on Hybrid Systems: Computation and Control (2014)
5. Braberman, V., Garbervetsky, D., Olivero, A.: Improving the verification of timed systems using influence information. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 21–36. Springer, Heidelberg (2002)
6. Clarke, E.M., Enders, R., Filkorn, T., Jha, S.: Exploiting symmetry in temporal logic model checking. *Formal Methods in System Design* 9, 77–104 (1996)
7. de Moura, L., Bjørner, N.S.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
8. Dill, D.L.: The $\text{mur}\varphi$ verification system. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 390–393. Springer, Heidelberg (1996)
9. Emerson, E.A., Sistla, A.P.: Symmetry and model checking. *Formal Methods in System Design* 9(1-2), 105–131 (1996)
10. Emerson, E., Wahl, T.: Dynamic symmetry reduction. In: Halbwachs, N., Zuck, L.D. (eds.) TACAS 2005. LNCS, vol. 3440, pp. 382–396. Springer, Heidelberg (2005)
11. Hendriks, M., Behrmann, G., Larsen, K.G., Niebert, P., Vaandrager, F.W.: Adding symmetry reduction to UPPAAL. In: Larsen, K.G., Niebert, P. (eds.) FORMATS 2003. LNCS, vol. 2791, pp. 46–59. Springer, Heidelberg (2004)
12. Hendriks, M.: Model checking timed automata: Techniques and applications. Ph.D. thesis, University of Nijmegen, The Netherlands (2006)
13. Herrera, C., Westphal, B., Feo-Arenis, S., Muñoz, M., Podelski, A.: Reducing Quasi-Equal Clocks in Networks of Timed Automata. In: Jurdiński, M., Ničković, D. (eds.) FORMATS 2012. LNCS, vol. 7595, pp. 155–170. Springer, Heidelberg (2012)
14. Ip, C.N., Dill, D.L.: Better verification through symmetry. *Formal Methods in System Design* 9, 41–75 (1996)
15. Ip, C.N., Dill, D.L.: Verifying systems with replicated components in $\text{Mur}\varphi$. *Formal Methods in System Design* 14(3) (1999)
16. Johnson, T.T.: Uniform Verification of Safety for Parameterized Networks of Hybrid Automata. Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana, IL 61801 (2013)
17. Johnson, T.T., Mitra, S.: Parameterized verification of distributed cyber-physical systems: An aircraft landing protocol case study. In: ACM/IEEE 3rd International Conference on Cyber-Physical Systems (April 2012)
18. Johnson, T.T., Mitra, S.: A small model theorem for rectangular hybrid automata networks. In: Giese, H., Rosu, G. (eds.) FORTE 2012 and FMOODS 2012. LNCS, vol. 7273, pp. 18–34. Springer, Heidelberg (2012)
19. Johnson, T.T., Mitra, S.: Invariant synthesis for verification of parameterized cyber-physical systems with applications to aerospace systems. In: Proceedings of the AIAA Infotech at Aerospace Conference (AIAA Infotech 2013), Boston, MA (August 2013)

20. Obal, W.D., McQuinn, M., Sanders, W.: Detecting and exploiting symmetry in discrete-state Markov models. *IEEE Transactions on Reliability* 56(4), 643–654 (2007)
21. Si, Y., Sun, J., Liu, Y., Wang, T.: Improving model checking stateful timed csp with non-zenoness through clock-symmetry reduction. In: Groves, L., Sun, J. (eds.) *ICFEM 2013*. LNCS, vol. 8144, pp. 182–198. Springer, Heidelberg (2013)
22. Sun, J., Liu, Y., Dong, J.S., Liu, Y., Shi, L., André, E.: Modeling and verifying hierarchical real-time systems using stateful timed csp. *ACM Trans. Softw. Eng. Methodol.* 22(1), 1–29 (2013)
23. Sun, J., Liu, Y., Dong, J.S., Pang, J.: PAT: Towards flexible verification under fairness. In: Bouajjani, A., Maler, O. (eds.) *CAV 2009*. LNCS, vol. 5643, pp. 709–714. Springer, Heidelberg (2009)