

# Parameterized Verification of Infinite-State Processes with Global Conditions

Parosh Aziz Abdulla<sup>1</sup>, Giorgio Delzanno<sup>2</sup>, and Ahmed Rezine<sup>1</sup>

<sup>1</sup> Uppsala University, Sweden  
parosh@it.uu.se, Rezine.Ahmed@it.uu.se

<sup>2</sup> Università di Genova, Italy  
giorgio@disi.unige.it

**Abstract.** We present a simple and effective approximated backward reachability algorithm for parameterized systems with existentially and universally quantified global conditions. The individual processes operate on unbounded local variables ranging over the natural numbers. In addition, processes may communicate via broadcast, rendez-vous and shared variables. We apply the algorithm to verify mutual exclusion for complex protocols such as Lamport’s bakery algorithm both with and without atomicity conditions, a distributed version of the bakery algorithm, and Ricart-Agrawala’s distributed mutual exclusion algorithm.

## 1 Introduction

We consider the analysis of safety properties for *parameterized systems*. A parameterized system consists of an arbitrary number of processes. The task is to verify correctness regardless of the number of processes. This amounts to the verification of an infinite family; namely one for each size of the system. Most existing approaches to automatic verification of parameterized systems make the restriction that each process is finite-state. However, there are many applications where the behaviour relies on unbounded data structures such as counters, priorities, local clocks, time-stamps, and process identifiers.

In this paper, we consider parameterized systems where the individual processes operate on Boolean variables, and on numerical variables which range over the natural numbers. The transitions are conditioned by the local state of the process, values of the local variables; and by *global conditions* which check the local states and variables of the other processes. These conditions are stated as propositional constraints on the Boolean variables, and as *gap-order constraints* on the numerical variables. Gap-order constraints [20] are a logical formalism in which we can express simple relations on variables such as lower and upper bounds on the values of individual variables; and equality, and gaps (minimal differences) between values of pairs of variables. A global condition is either *universally* or *existentially* quantified. An example of a universal condition is “variable  $x$  of a given process  $i$  has a value which is greater than the value of variable  $y$  in all other processes inside the system”. Process  $i$  is then allowed to

perform the transition only if this condition is satisfied. In an existential condition we require that *some* (rather than *all*) processes satisfy the condition. In addition to these classes of transitions, processes may communicate through broadcast, rendez-vous, and shared variables.

There are at least two advantages with using gap-order constraints as a language for expressing the enabling conditions of transitions. First, they allow to handle a large class of protocols where the behaviour depends on the relative ordering of values among variables, rather than the actual values of these variables. The second reason is that they define a natural ordering on the set of system configurations. In fact, it can be shown, using standard techniques (such as the ones in [23]), that checking safety properties (expressed as regular languages) can be translated into the reachability of sets of configurations which are upward closed with respect to this ordering.

To check safety properties, we perform backward reachability analysis using gap-order constraints as a symbolic representation of upward closed sets of configurations. In the analysis, we consider a transition relation which is an over-approximation of the one induced by the parameterized system. To do that, we modify the semantics of universal quantifiers by eliminating the processes which violate the given condition. For instance in the above example, process  $i$  is always allowed to take the transition. However, when performing the transition, we eliminate each process  $j$  where the value of  $y$  is smaller or equal to the value of  $x$  in  $i$ . The approximate transition system obtained in this manner is *monotonic* with respect to the above mentioned ordering, in the sense that larger configurations can simulate smaller ones. In fact, it turns out that universal quantification is the only operation which does not preserve monotonicity and hence it is the only source of approximation in the model. The fact that the approximate transition relation is monotonic, means that upward closedness is maintained under the operation of computing predecessors. A significant aspect of the reachability procedure is that the number of copies of variables (both Boolean and numerical) which appear in constraints whose denotations are upward closed sets is not bounded a priori. The reason is that there is an arbitrary number of processes each with its own local copy of the variables. The whole verification process is fully automatic since both the approximation and the reachability analysis are carried out without user intervention. Observe that if the approximate transition system satisfies a safety property then we can conclude that the original system satisfies the property, too.

Termination of the approximated backward reachability analysis is not guaranteed in general. However, the procedure terminates on all the examples we report in this paper. Furthermore, termination is guaranteed in some restricted cases such as for systems with existential or universal global conditions but with at most one local integer variable.

In order to test our method we have selected a collection of challenging protocols in which integer variables are used either as identifiers, priorities, local clocks, or time-stamps. Almost all of the examples are outside the class for

which termination is guaranteed. In particular, we automatically verify safety properties for parameterized versions of the following algorithms:

- Lamport’s bakery algorithm [19] with atomicity conditions;
- A version of Lamport’s bakery algorithm with non-atomic computation of tickets;
- A distributed version of Lamport’s bakery in which tickets and entry conditions are computed and tested non-atomically by means of a handshake protocol run by each process;
- The Ticket mutual exclusion algorithm with a central monitor for distributing tickets [6];
- The Ricart-Agrawala distributed mutual exclusion algorithm based on the use of logical clocks and time-stamps [21].

We also consider a bogus version of the Lamport’s bakery without atomicity conditions in the computation of tickets. In this version, the *choosing* flag is simply ignored in the entry section. For this example, our procedure returns symbolic traces (from initial to bad states) that explain the subtle race conditions that may arise when the flag is not tested.

Each one of these examples present challenging problems for parameterized verification methods in the following sense:

- Their underlying logic is already hard for manual or finite-state verification.
- They are all instances of multidimensional infinite-state systems in which processes have unbounded local variables and (apart from Ticket) an order over identifiers is used to break the tie in the entry section. For instance, they cannot be modelled without the use of abstractions in the framework of Regular Model Checking [17,5,8,4].
- In all examples, global conditions are needed to model the communication mechanisms used in the protocols (e.g. broadcasts, update, and entry conditions that depend on the local integer variables of other processes).

*Related Work.* The multi-dimensional parameterized models studied in the present paper cannot be analyzed without use of additional abstractions by methods designed for networks of finite-state processes, e.g., Regular Model Checking [17,5,8] and counter abstraction methods [12,16,13,14]. The approximation scheme we apply in our backward reachability procedure works well for a very large class of one-dimensional parameterized systems. In fact, the verification procedure used in [4] is a special case of the current one, where the processes are restricted to be finite-state systems.

Parameterized versions of Lamport’s bakery algorithm have been tested using a semi-automated verification method based on *invisible invariants* in [7], with the help of *environmental abstractions* for a formulation with atomicity conditions in [11], and using heuristics to discover *indexed predicates* in [18]. A parameterized formulation of the Ricart-Agrawala algorithm has been verified semi-automatically in [22], where the STeP prover is used to discharge some of the verification conditions needed in the proof. We are not aware of other

attempts of fully automatic verification of parameterized versions of the Ricart-Agrawala algorithm or of the distributed version (with no atomicity assumptions) of Lamport’s bakery algorithm.

In contrast to the above mentioned methods, our verification procedure is fully automated and it is based on a generic approximation scheme. Furthermore, our method is applicable to versions of Lamport’s bakery both with and without atomicity conditions and may return symbolic traces useful for debugging.

A parameterized formulation of an abstraction of the Ticket algorithm has been analyzed in [9]. The verification procedure in [9] does not handle parameterized universally quantified global conditions. Furthermore, in the abstraction of the Ticket algorithm studied in [9] the central monitor may forget tickets (the update of *turn* is defined by a jump to any larger value). Thus, the model does not keep the FIFO order of requests. With the help of universally quantified guards and of our approximation, we verify a more precise model in which the FIFO order of requests is always preserved.

In contrast to symbolic methods for finite, a priori fixed collections of processes with local integer variables, e.g., those in [10,15], our gap-order constraints are defined over an *infinite* collections of variables. The number of copies of variables needed during the backward search cannot be bounded a priori. This feature allows us to reason about systems with global conditions over any number of processes. Furthermore, the present method covers that of [2] which also uses gap-order constraints to reason about systems with unbounded numbers of processes. However, [2] cannot deal with global conditions which is the main feature of the examples considered here.

*Outline.* In the next two Sections we give some preliminaries and define a basic model for parameterized systems. Section 4 and 5 describe the induced transition system and the coverability (safety) problem. In Section 6 we define the approximated transition system. Section 7 defines the gap-order constraints and presents the backward reachability algorithm, while Section 8 describes the operations on constraints used in the algorithm. Section 9 explains how we extend the basic model to cover features such as shared variables, broadcast and binary communication. In Section 10 we report the results of running our prototypes on a number of examples. Finally, in Section 11, we give conclusions and directions for future work. In the appendix, we give some proofs and detailed descriptions of the case studies.

## 2 Preliminaries

In this section, we give some preliminary notations and definitions. We use  $\mathcal{B}$  to denote the set  $\{true, false\}$  of Boolean values; and use  $\mathcal{N}$  to denote the set of natural numbers. For a natural number  $n$ , let  $\bar{n}$  denote the set  $\{1, \dots, n\}$ .

For a finite set  $A$ , we write a multiset over  $A$  as a list  $[a_1, a_2, \dots, a_n]$ , where  $a_i \in A$  for each  $i : 1 \leq i \leq n$ . We use  $a \in A$  to denote that  $a = a_i$  for some  $i : 1 \leq i \leq n$ . For multisets  $M_1 = [a_1, \dots, a_m]$  and  $M_2 = [b_1, \dots, b_n]$ , we use  $M_1 \bullet M_2$  to denote the union (sum) of  $M_1$  and  $M_2$  (i.e.,  $M_1 \bullet M_2 = [a_1, \dots, a_m, b_1, \dots, b_n]$ ).

We will work with sets of variables. Such a set  $A$  is often partitioned into two subsets: *Boolean* variables  $A_{\mathcal{B}}$  which range over  $\mathcal{B}$ , and *numerical* variables  $A_{\mathcal{N}}$  which range over  $\mathcal{N}$ . We denote by  $\mathbb{B}(A_{\mathcal{B}})$  the set of Boolean formulas over  $A_{\mathcal{B}}$ . We will also use a simple set of formulas, called *gap formulas*, to constrain the numerical variables. More precisely, we let  $\mathbb{G}(A_{\mathcal{N}})$  be the set of formulas which are either of the form  $x = y$  or of the form  $x \sim_k y$  where  $\sim \in \{<, \leq\}$ ,  $x, y \in A_{\mathcal{N}} \cup \mathcal{N}$ , and  $k \in \mathcal{N}$ . Here  $x <_k y$  stands for  $x + k < y$ . We use  $\mathbb{F}(A)$  to denote the set of formulas which has members of  $\mathbb{B}(A)$  and of  $\mathbb{G}(\mathcal{N})$  as atomic formulas, and which is closed under the Boolean connectives  $\wedge, \vee$ . For instance, if  $A_{\mathcal{B}} = \{a, b\}$  and  $A_{\mathcal{N}} = \{x, y\}$  then  $\theta = (a \supset b) \wedge (x + 3 < y)$  is in  $\mathbb{F}(A)$ . Sometimes, we write a formula as  $\theta(y_1, \dots, y_k)$  where  $y_1, \dots, y_k$  are the variables which may occur in  $\theta$ ; so we can write the above formula as  $\theta(x, y, a, b)$ .

Sometimes, we perform substitutions on logical formulas. A *substitution* is a set  $\{x_1 \leftarrow e_1, \dots, x_n \leftarrow e_n\}$  of pairs, where  $x_i$  is a variable and  $e_i$  is either a constant or a variable of the same type as  $x_i$ , for each  $i : 1 \leq i \leq n$ . Here, we assume that all the variables are distinct, i.e.,  $x_i \neq x_j$  if  $i \neq j$ . For a formula  $\theta$  and a substitution  $S$ , we use  $\theta[S]$  to denote the formula we get from  $\theta$  by simultaneously replacing all occurrences of the variables  $x_1, \dots, x_n$  by  $e_1, e_2, \dots, e_n$  respectively. Sometimes, we may write  $\theta[S_1][S_2] \cdots [S_m]$  instead of  $\theta[S_1 \cup S_2 \cup \cdots \cup S_m]$ . As an example, if  $\theta = (x_1 < x_2) \wedge (x_3 < x_4)$  then  $\theta[x_1 \leftarrow 3, x_4 \leftarrow 2][x_2 \leftarrow y] = (3 < y) \wedge (x_3 < 2)$ .

### 3 Parameterized Systems

In this section, we introduce a basic model for parameterized systems. The basic model will be enriched by additional features in Section 9.

A parameterized system consists of an arbitrary (but finite) number of identical processes. Each process is modelled as an extended finite-state automaton operating on local variables which range over the Booleans and the natural numbers. The transitions of the automaton are conditioned by the values of the local variables and by *global* conditions in which the process checks, for instance, the local states and variables of all the other processes inside the system. A transition may change the value of any local variable inside the process (possibly deriving the new values from those of the other processes). A parameterized system induces an infinite family of (potentially infinite-state) systems, namely one for each size  $n$ . The aim is to verify correctness of the systems for the whole family (regardless of the number  $n$  of processes inside the system).

Formally, a *parameterized system*  $\mathcal{P}$  is a triple  $(Q, X, T)$ , where  $Q$  is a finite set of *local states*,  $X$  is a finite set of *local variables* partitioned into  $X_{\mathcal{B}}$  (which range over  $\mathcal{B}$ ) and  $X_{\mathcal{N}}$  (which range over  $\mathcal{N}$ ), and  $T$  is a finite set of *transition rules*. A transition rule  $t$  is of the form

$$t : [q \rightarrow q' \triangleright \theta] \tag{1}$$

where  $q, q' \in Q$  and  $\theta$  is either a *local* or a *global condition*. Intuitively, the process which makes the transition changes its local state from  $q$  to  $q'$ . In the

meantime, the values of the local variables of the process are updated according to  $\theta$ . Below, we describe how we define local and global conditions.

To simplify the definitions, we sometimes regard members of the set  $Q$  as Boolean variables. Intuitively, the value of the Boolean variable  $q \in Q$  is *true* for a particular process iff the process is in local state  $q$ . We define the set  $Y = X \cup Q$ .

To define local conditions, we introduce the set  $X^{next}$  which contains the *next-value* versions of the variables in  $X$ . A variable  $x^{next} \in X^{next}$  represents the next value of  $x \in X$ . A *local condition* is a formula in  $\mathbb{F}(X \cup X^{next})$ . The formula specifies how local variables of the current process are updated with respect to their current values.

Global conditions check the values of local variables of the current process, together with the local states and the values of local variables of the other processes. We need to distinguish between a local variable, say  $x$ , of the process which is about to perform a transition, and the same local variable  $x$  of the other processes inside the system. We do that by introducing, for each  $x \in Y$ , two new variables **self**· $x$  and **other**· $x$ . We define the sets **self**· $Y = \{\mathbf{self} \cdot x \mid x \in Y\}$  and **other**· $Y = \{\mathbf{other} \cdot x \mid x \in Y\}$ . The sets **self**· $X$ , **other**· $X^{next}$ , etc, are defined in the obvious manner. A *global condition*  $\theta$  is of one of the following two forms:

$$\forall \mathbf{other} \neq \mathbf{self} \cdot \theta_1 \quad \exists \mathbf{other} \neq \mathbf{self} \cdot \theta_1 \quad (2)$$

where  $\theta_1 \in \mathbb{F}(\mathbf{self} \cdot X \cup \mathbf{other} \cdot Y \cup \mathbf{self} \cdot X^{next})$ . In other words, the formula checks the local variables of the process which is about to make the transition (through **self**· $X$ ), and the local states and variables of the other processes (through **other**· $Y$ ). It also specifies how the local variables of the process in transition are updated (through **self**· $X^{next}$ ). A global condition is said to be *universal* or *existential* depending on the type of the quantifier appearing in it. As an example, the following formula

$$\forall \mathbf{other} \neq \mathbf{self} \cdot (\mathbf{self} \cdot a) \wedge (\mathbf{self} \cdot x^{next} > \mathbf{other} \cdot x) \wedge \mathbf{other} \cdot q_1$$

states that the transition may be performed only if variable  $a$  of the current process has the value *true*, and all the other processes are in local state  $q_1$ . When the transition is performed, variable  $x$  of the current process is assigned a value which is greater than the value of  $x$  in all the other processes.

## 4 Transition System

We describe the transition system induced by a parameterized system.

A *transition system*  $\mathcal{T}$  is a pair  $(D, \Longrightarrow)$ , where  $D$  is an (infinite) set of *configurations* and  $\Longrightarrow$  is a binary relation on  $D$ . We use  $\Longrightarrow^*$  to denote the reflexive transitive closure of  $\Longrightarrow$ . Let  $\preceq$  be an ordering on  $D$ . We say that  $\mathcal{T}$  is *monotonic* with respect to  $\preceq$  if the following property is satisfied: for all  $c_1, c_2, c_3 \in D$  with  $c_1 \Longrightarrow c_2$  and  $c_1 \preceq c_3$ , there is a  $c_4 \in D$  such that  $c_3 \Longrightarrow c_4$  and  $c_2 \preceq c_4$ . We will consider several transition systems in this paper.

First, a parameterized system  $\mathcal{P} = (Q, X, T)$  induces a transition system  $\mathcal{T}(\mathcal{P}) = (C, \longrightarrow)$  as follows. A configuration is defined by the local states and the values of the local variables in the processes. Formally, a *local variable state*  $v$  is a mapping from  $X$  to  $\mathcal{B} \cup \mathcal{N}$  which respects variables' types. A *process state*  $u$  is a pair  $(q, v)$  where  $q \in Q$  and  $v$  is a local variable state. As mentioned in Section 3, we may regard members of  $Q$  as Boolean variables. Thus, we can view a process state  $(q, v)$  as a mapping  $u : Y \mapsto \mathcal{B} \cup \mathcal{N}$ , where  $u(x) = v(x)$  for each  $x \in X$ ,  $u(q) = \text{true}$ , and  $u(q') = \text{false}$  for each  $q' \in Q - \{q\}$ . A *configuration* is a multiset  $[u_1, u_2, \dots, u_n]$  of process states. Intuitively, the above configuration corresponds to an instance of the system with  $n$  processes. Notice that if  $c_1$  and  $c_2$  are configurations then so is their union  $c_1 \bullet c_2$ .

We define the transition relation  $\longrightarrow$  on the set of configurations as follows. We start by describing the semantics of local conditions. Recall that a local condition corresponds to one process changing state without checking states of the other processes. Therefore, the semantics is defined in terms of two local variable states  $v, v'$  corresponding to the current resp. next values of the local variables of the process; and a formula  $\theta \in \mathbb{F}(X \cup X^{\text{next}})$  (representing the local condition). We write  $(v, v') \models \theta$  to denote the validity of the formula  $\theta[\rho][\rho']$  where the substitutions are defined by  $\rho := \{x \leftarrow v(x) \mid x \in X\}$  and  $\rho' := \{x^{\text{next}} \leftarrow v'(x) \mid x \in X\}$ . In other words, we check the formula we get by replacing the current- resp. next-value variables in  $\theta$  by their values as defined by  $v$  resp.  $v'$ . The formula is evaluated using the standard interpretations of the Boolean connectives, and the arithmetical relations  $<, \leq, =$ . For process states  $u = (q, v)$  and  $u' = (q', v')$ , we use  $(u, u') \models \theta$  to denote that  $(v, v') \models \theta$ .

Next, we describe the semantics of global conditions. The definition is given in terms of two local variable states  $v, v'$ , a process state  $u_1$ , and a formula  $\theta \in \mathbb{F}(\text{self} \cdot X \cup \text{other} \cdot Y \cup \text{self} \cdot X^{\text{next}})$  (representing a global condition). The roles of  $v$  and  $v'$  are the same as for local conditions. We recall that a global condition also checks states of all (or some) of the other processes. Here,  $u_1$  represents the local state and variables of one such a process. We write  $(v, v', u_1) \models \theta$  to denote the validity of the formula  $\theta[\rho][\rho'][\rho_1]$  where the substitutions are defined by  $\rho := \{\text{self} \cdot x \leftarrow v(x) \mid x \in X\}$ ,  $\rho' := \{\text{self} \cdot x^{\text{next}} \leftarrow v'(x) \mid x \in X\}$ , and  $\rho_1 := \{\text{other} \cdot x \leftarrow u_1(x) \mid x \in Y\}$ . The relation  $(u, u', u_1) \models \theta$  is interpreted in a similar manner to the case of local conditions.

Now, we are ready to define the transition relation  $\longrightarrow$ . Let  $t$  be a transition rule of the form of (1). Consider two configurations  $c = c_1 \bullet [u] \bullet c_2$  and  $c' = c_1 \bullet [u'] \bullet c_2$  where  $u = (q, v)$  and  $u' = (q', v')$ . We denote by  $c \xrightarrow{t} c'$  that one of the following conditions is satisfied:

1.  $\theta$  is a local condition and  $(u, u') \models \theta$ .
2.  $\theta$  is a universal global condition of the form of (2), and  $(u, u', u_1) \models \theta_1$  for each  $u_1 \in c_1 \bullet c_2$ .
3.  $\theta$  is an existential global condition of the form of (2), and  $(u, u', u_1) \models \theta_1$  for some  $u_1 \in c_1 \bullet c_2$ .

We use  $c \longrightarrow c'$  to denote that  $c \xrightarrow{t} c'$  for some  $t \in T$ .

## 5 Safety Properties

In this section, we introduce an ordering on configurations, and use it to define the safety problem. Given a parameterized system  $\mathcal{P} = (Q, X, T)$ , we assume that, prior to starting the execution of the system, each process is in an (identical) *initial* process state  $u_{init} = (q_{init}, v_{init})$ . In the induced transition system  $\mathcal{T}(\mathcal{P}) = (C, \longrightarrow)$ , we use  $Init$  to denote the set of *initial* configurations, i.e., configurations of the form  $[u_{init}, \dots, u_{init}]$ . Notice that this set is infinite.

We define an ordering on configurations as follows. Consider two configurations,  $c = [u_1 \cdot \dots \cdot u_m]$  and  $c' = [u'_1 \cdot \dots \cdot u'_n]$ , where  $u_i = (q_i, v_i)$  for each  $i : 1 \leq i \leq m$ , and  $u'_i = (q'_i, v'_i)$  for each  $i : 1 \leq i \leq n$ . We write  $c \preceq c'$  to denote that there is an injection  $h : \overline{m} \rightarrow \overline{n}$  such that the following four conditions are satisfied for each  $i, j : 1 \leq i, j \leq m$ :

1.  $q_i = q'_{h(i)}$ .
2.  $v_i(x) = true$  iff  $v'_{h(i)}(x) = true$  for each  $x \in X_{\mathcal{B}}$ .
3.  $v_i(x) = v_j(y)$  iff  $v'_{h(i)}(x) = v'_{h(j)}(y)$ , for each  $x, y \in X_{\mathcal{N}}$ .
4.  $v_i(x) <_k v_j(y)$  implies that there is a  $l \geq k$  with  $v'_{h(i)}(x) <_l v'_{h(j)}(y)$ , for each  $x, y \in X_{\mathcal{N}}$ .

In other words, for each process in  $c$  there is a corresponding process in  $c'$ . The local states and the values of the Boolean variables coincide in the corresponding processes (Conditions 1 and 2). Regarding the numerical variables, the ordering preserves equality of variables (Condition 3), while gaps between variables in  $c'$  are at least as large as the gaps between the corresponding variables in  $c$  (Condition 4).

A set of configurations  $D \subseteq C$  is *upward closed* (with respect to the ordering  $\preceq$ ) if  $c \in D$  and  $c \preceq c'$  implies  $c' \in D$ . For sets of configurations  $D, D' \subseteq C$  we use  $D \longrightarrow D'$  to denote that there are  $c \in D$  and  $c' \in D'$  with  $c \longrightarrow c'$ .

The *coverability problem* for parameterized systems is defined as follows:

### PAR-COV

#### **Instance**

- A parameterized system  $\mathcal{P} = (Q, X, T)$ .
- An upward closed set  $C_F$  of configurations.

**Question**  $Init \xrightarrow{*} C_F ?$

It can be shown, using standard techniques (see e.g. [23]), that checking safety properties (expressed as regular languages) can be translated into instances of the coverability problem. Therefore, checking safety properties amounts to solving PAR-COV (i.e., to the reachability of upward closed sets).

## 6 Approximation

In this section, we introduce an over-approximation of the transition relation of a parameterized system. The aim of the over-approximations is to derive a

new transition system which is *monotonic* with respect to the ordering  $\preceq$  defined on configurations in Section 5. The only transitions which do not preserve monotonicity are those involving universal global conditions. Therefore, the approximate transition system modifies the behavior of universal quantifiers in such a manner that monotonicity is maintained. Roughly speaking, in the new semantics, we remove all processes in the configuration which violate the condition of the universal quantifier. Below we describe how this is done.

In Section 4, we mentioned that each parameterized system  $\mathcal{P} = (Q, X, T)$  induces a transition system  $\mathcal{T}(\mathcal{P}) = (C, \longrightarrow)$ . A parameterized system  $\mathcal{P}$  also induces an *approximate* transition system  $\mathcal{A}(\mathcal{P}) = (C, \rightsquigarrow)$ ; the set  $C$  of configurations is identical to the one in  $\mathcal{T}(\mathcal{P})$ . We define  $\rightsquigarrow = (\longrightarrow \cup \rightsquigarrow_1)$ , where  $\longrightarrow$  is defined in Section 4, and  $\rightsquigarrow_1$  (which reflects the approximation of universal quantifiers) is defined as follows. For a configuration  $c$ , a formula  $\theta \in \mathbb{F}(\mathbf{self} \cdot X \cup \mathbf{other} \cdot Y \cup \mathbf{self} \cdot X^{next})$ , and process states  $u, u'$ , we use  $c \ominus (\theta, u, u')$  to denote the configuration derived from  $c$  by deleting all process states  $u_1$  such that  $(u, u', u_1) \not\models \theta$ . To explain this operation intuitively, we recall that a universal global condition requires that the current and next states of the current process (described by  $u$  resp.  $u'$ ) together with the state of each other process (described by  $u_1$ ) should satisfy the formula  $\theta$ . The operation then removes from  $c$  each process whose state  $u_1$  does not comply with this condition.

Consider two configurations  $c = c_1 \bullet u \bullet c_2$  and  $c' = c'_1 \bullet u' \bullet c'_2$ , where  $u = (q, v)$  and  $u' = (q', v')$ . Let  $t$  be a transition rule of the form of (1), such that  $\theta$  is a universal global condition of the form of (2). We write  $c \overset{t}{\rightsquigarrow}_1 c'$  to denote that  $c'_1 = c_1 \ominus (\theta_1, u, u')$  and  $c'_2 = c_2 \ominus (\theta_1, u, u')$ . We use  $c \overset{t}{\rightsquigarrow} c'$  for some  $t \in T$ ; and use  $\rightsquigarrow^*$  to denote the reflexive transitive closure of  $\rightsquigarrow$ .

**Lemma 1.** *The approximate transition system  $(C, \rightsquigarrow)$  is monotonic with respect to  $\preceq$ .*

We define the coverability problem for the approximate system as follows.

APRX-PAR-COV  
**Instance**

- A parameterized system  $\mathcal{P} = (Q, X, T)$ .
- An upward closed set  $C_F$  of configurations.

**Question**  $Init \rightsquigarrow^* C_F ?$

Since  $\longrightarrow \subseteq \rightsquigarrow$ , a negative answer to APRX-PAR-COV implies a negative answer to PAR-COV.

## 7 Backward Reachability Analysis

In this section, we present a scheme based on backward reachability analysis for solving APRX-PAR-COV. For the rest of this section, we fix an approximate transition system  $\mathcal{A}(\mathcal{P}) = (C, \rightsquigarrow)$ .

**Constraints.** The scheme operates on *constraints* which we use as a symbolic representation for sets of configurations. For each natural number  $i \in \mathcal{N}$  we make a copy  $Y^i$  such that  $x^i \in Y^i$  if  $x \in Y$ . A *constraint*  $\phi$  is a pair  $(m, \psi)$ , where  $m \in \mathcal{N}$  is a natural number, and  $\psi \in \mathbb{F}(Y^1 \cup Y^2 \cup \dots \cup Y^m)$ . Intuitively, a configuration satisfying  $\phi$  should contain at least  $m$  processes (indexed by  $1, \dots, m$ ). The constraint  $\phi$  uses the elements of the set  $Y^i$  to refer to the local states and variables of process  $i$ . The values of these states and variables are constrained by the formula  $\psi$ . Formally, consider a configuration  $c = [u_1, u_2, \dots, u_n]$  and a constraint  $\phi = (m, \psi)$ . Let  $h : \overline{m} \mapsto \overline{n}$  be an injection. We write  $c \models_h \phi$  to denote the validity of the formula  $\psi[\rho]$  where  $\rho := \{x^i \leftarrow u_{h(i)}(x) \mid x \in Y \text{ and } 1 \leq i \leq m\}$ . In other words, there should be at least  $m$  processes inside  $c$  whose local states and variables have values which satisfy  $\psi$ . We write  $c \models \phi$  to denote that  $c \models_h \phi$  for some  $h$ ; and define  $\llbracket \phi \rrbracket = \{c \mid \phi \models c\}$ . For a (finite) set of constraints  $\Phi$ , we define  $\llbracket \Phi \rrbracket = \bigcup_{\phi \in \Phi} \llbracket \phi \rrbracket$ . The following lemma follows from the definitions.

**Lemma 2.** *For each constraint  $\phi$ , the set  $\llbracket \phi \rrbracket$  is upward closed.*

In all the examples we consider, the set  $C_F$  in the definition of APRX-PARCOV can be represented by a finite set  $\Phi_F$  of constraints. The coverability question can then be answered by checking whether  $Init \xrightarrow{*} \llbracket \Phi_F \rrbracket$ .

**Entailment and Predecessors.** To define our scheme we will use two operations on constraints; namely *entailment*, and *computing predecessors*, defined below. We define an *entailment relation*  $\sqsubseteq$  on constraints, where  $\phi_1 \sqsubseteq \phi_2$  iff  $\llbracket \phi_2 \rrbracket \subseteq \llbracket \phi_1 \rrbracket$ . For sets  $\Phi_1, \Phi_2$  of constraints, abusing notation, we let  $\Phi_1 \sqsubseteq \Phi_2$  denote that for each  $\phi_2 \in \Phi_2$  there is a  $\phi_1 \in \Phi_1$  with  $\phi_1 \sqsubseteq \phi_2$ . Observe that  $\Phi_1 \sqsubseteq \Phi_2$  implies that  $\llbracket \Phi_2 \rrbracket \subseteq \llbracket \Phi_1 \rrbracket$ .

For a constraint  $\phi$ , we let  $Pre(\phi)$  be a set of constraints, such that  $\llbracket Pre(\phi) \rrbracket = \{c' \mid \exists c \in \llbracket \phi \rrbracket. c \rightsquigarrow c'\}$ . In other words  $Pre(\phi)$  characterizes the set of configurations from which we can reach a configuration in  $\phi$  through the application of a single rule in the approximate transition relation. In the definition of  $Pre$  we rely on the fact that, in any monotonic transition system, upward-closedness is preserved under the computation of the set of predecessors (see e.g. [1]). From Lemma 2 we know that  $\llbracket \phi \rrbracket$  is upward closed; by Lemma 1,  $(C, \rightsquigarrow)$  is monotonic, we therefore know that  $\llbracket Pre(\phi) \rrbracket$  is upward closed. In fact, we show in Section 8 that this set is finite and computable. For a set  $\Phi$  of constraints, we let  $Pre(\Phi) = \bigcup_{\phi \in \Phi} Pre(\phi)$ .

**Scheme.** Given a finite set  $\Phi_F$  of constraints, the scheme checks whether  $Init \xrightarrow{*} \llbracket \Phi_F \rrbracket$ . We perform a backward reachability analysis, generating a sequence  $\Phi_0 \supseteq \Phi_1 \supseteq \Phi_2 \supseteq \dots$  of finite sets of constraints such that  $\Phi_0 = \Phi_F$ , and  $\Phi_{j+1} = \Phi_j \cup Pre(\Phi_j)$ . Since  $\llbracket \Phi_0 \rrbracket \subseteq \llbracket \Phi_1 \rrbracket \subseteq \llbracket \Phi_2 \rrbracket \subseteq \dots$ , the procedure terminates when we reach a point  $j$  where  $\Phi_j \sqsubseteq \Phi_{j+1}$ . Notice that the termination condition implies that  $\llbracket \Phi_j \rrbracket = (\bigcup_{0 \leq i \leq j} \llbracket \Phi_i \rrbracket)$ . Consequently,  $\Phi_j$  characterizes the set of all predecessors of  $\llbracket \Phi_F \rrbracket$ . This means that  $Init \xrightarrow{*} \llbracket \Phi_F \rrbracket$  iff  $(Init \cap \llbracket \Phi_j \rrbracket) \neq \emptyset$ .

Observe that, in order to implement the scheme (i.e., transform it into an algorithm), we need to be able to (i) compute  $Pre$ ; (ii) check for entailment between constraints; and (iii) check for emptiness of  $(Init \cap \llbracket \phi \rrbracket) \neq \emptyset$  for a constraint  $\phi$ .

## 8 Constraint Operations

In this section, we show how to perform the three operations on constraints which are used in the scheme presented in Section 7. In the rest of the section, we fix a parameterized systems  $\mathcal{P} = (Q, X, T)$ . Recall that  $Y = X \cup Q$ .

**Entailment.** Consider two constraints  $\phi = (m, \psi)$ , and  $\phi' = (m', \psi')$ . Let  $\mathcal{H}(\phi, \phi')$  be the set of injections  $h : \overline{m} \mapsto \overline{m'}$ . We use  $\psi^h$  to denote the formula  $\psi[\rho]$ , where  $\rho := \{x^i \leftarrow x^{h(i)} \mid x \in Y \text{ and } i \in \overline{m}\}$ . The following lemma gives a logical characterization which allows the computation of the entailment relation.

**Lemma 3.** *Given constraints  $\phi = (m, \psi)$ , and  $\phi' = (m', \psi')$ , we have  $\phi \sqsubseteq \phi'$  iff*

$$\forall y_1 \cdots y_k. \psi'(y_1, \dots, y_k) \supset \bigvee_{h \in \mathcal{H}(\phi, \phi')} \psi^h(y_1, \dots, y_k)$$

**Pre** The following lemma describes the computation of the function *Pre*.

**Lemma 4.** *For a constraint  $\phi$ , we can compute  $Pre(\phi)$  as a finite set of constraints.*

The complete proof of the lemma can be found in [3]. As an example, for a transition  $t$  with a universally quantified global condition  $\theta$ , *Pre* applied to constraint  $(m, \psi)$  returns the constraint  $(m, \psi')$ , where

$$\psi' = \exists Y^\bullet. \bigwedge_{1 \leq j \neq i \leq m} \psi[\rho_1^i] \wedge \theta'_1[\rho_2^{i,j}] \quad (3)$$

Here, for each variable  $x$ , the substitution  $\rho_1^i$  maps variable  $x^i$  to a new variable  $x^\bullet$  in  $Y^\bullet$ , while  $\rho_2^{i,j}$  maps  $\mathbf{self} \cdot x^{next}$  to  $x^\bullet$ ,  $\mathbf{self} \cdot x$  to  $x^i$ , and  $\mathbf{other} \cdot x$  to  $x^j$ . The condition  $\theta'_1$  is obtained from  $\theta$  by adding information on  $q$  and  $q'$ . All fresh variables in  $Y^\bullet$  are then projected away.

**Intersection with Initial States.** For a constraint  $\phi = (m, \psi)$ , we have  $(Init \cap \llbracket \phi \rrbracket) \neq \emptyset$  iff  $[u_{init}, \dots, u_{init}] \models \phi$ , where the multiset  $[u_{init}, \dots, u_{init}]$  is of size  $m$ .

## 9 Additional Features

We add a number of features to the model of Section 2, namely, *Binary Communication*, *Broadcast*, and *Shared Variables*. In Binary Communication, two processes satisfying some global conditions change their process states simultaneously. In Broadcast, a process satisfying a global condition, called the initiator, changes its process state, together with all other processes, called receptors, which satisfy another global condition. Shared Variables may have unbounded domain, and can be read and written by all processes in the system. For all the new features, we can use the same constraint system as in Section 7; consequently checking entailment and intersection with initial states need not be modified. Also, as shown in [3], the definition of the *Pre* operator can be extended to cope with the new classes of guards.

## 10 Experimental Results

We have built two different prototypes that implement our approximated backward reachability procedure, based on an integer resp. a real solver for handling the constraints over the process variables. The results are summarized in Figure 1. For each protocol we give the number of iterations and the time needed for performing the verification. The experiments are performed using a Pentium M 1.6 Ghz with 1G of memory (see [3] for more details).

Model	Iterations		Time		Safe		Trace	
	R	I	R	I	R	I	R	I
Simplified Bakery Alg.	6	6	0.8s	0.3s	✓	✓		
Lamport's Bakery Alg.	9	9	2.1s	2s	✓	✓		
Bogus Bakery	10	6	0.8s	11s			✓	✓
Ticket Mutex Alg.	9	8	0.3s	1.6s	✓	✓		
Ricart-Agrawala Distr. Mutex Alg.	9	11	3.4s	2mn40s	✓	✓		
Lamport's Distr. Mutex Alg.	21	27	9mn19s	146mn	✓	✓		

**Fig. 1.** Experimental results. R and I stand for the real resp. integer solver. Safe and Trace stand for checking safety properties resp. generating a counter-example.

## 11 Conclusion and Future Work

We have presented a method for approximate reachability analysis of systems which consist of an arbitrary number of processes each of which is infinite-state. Based on the method, we have implemented a prototype and automatically verified several non-trivial mutual exclusion protocols. The Bakery example describes a distributed protocol without atomicity assumptions on the transitions. One direction for future research is to develop a methodology for automatic verification of general classes of parameterized systems with non-atomic global conditions. Furthermore, our algorithm relies on an abstract ordering which can be naturally extended to several different types of data structures. We are currently developing similar algorithms for systems with more complicated topologies such as trees and general forms of graphs.

## References

1. Abdulla, P.A., Čerāns, K., Jonsson, B., Yih-Kuen, T.: Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation* 160, 109–127 (2000)
2. Abdulla, P.A., Delzanno, G.: On the coverability problem for constrained multiset rewriting. In: *Proc. AVIS'06*
3. Abdulla, P.A., Delzanno, G., Rezine, A.: Parameterized Verification of Infinite-state Processes with Global Conditions, Technical Report 2007-014, Uppsala University (April 2007)

4. Abdulla, P.A., Henda, N.B., Delzanno, G., Rezine, A.: Regular model checking without transducers. In: Proc. TACAS '07, To appear (2007)
5. Abdulla, P.A., Jonsson, B., Nilsson, M., d'Orso, J.: Regular model checking made simple and efficient. In: Brim, L., Jančar, P., Křetínský, M., Kucera, A. (eds.) CONCUR 2002. LNCS, vol. 2421, Springer, Heidelberg (2002)
6. Andrews, G.: Foundations of Multithreaded, Parallel, and Distributed Programming. Addison-Wesley, Reading (2000)
7. Arons, T., Pnueli, A., Ruah, S., Xu, J., Zuck, L.: Parameterized verification with automatically computed inductive assertions. In *Proc. CAV 2001*. In: Agha, G.A., De Cindio, F., Rozenberg, G. (eds.) Concurrent Object-Oriented Programming and Petri Nets. LNCS, vol. 2102, Springer, Heidelberg (2001)
8. Boigelot, B., Legay, A., Wolper, P.: Iterating transducers in the large. In: Hunt Jr., W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, Springer, Heidelberg (2003)
9. Bozzano, M., Delzanno, G.: Beyond parameterized verification. In: Katoen, J.-P., Stevens, P. (eds.) ETAPS 2002 and TACAS 2002. LNCS, vol. 2280, Springer, Heidelberg (2002)
10. Bultan, T., Gerber, R., Pugh, W.: Model-checking concurrent systems with unbounded integer variables. ACM TOPLAS 21(4), 747–789 (1999)
11. Clarke, E., Talupur, M., Veith, H.: Environment abstraction for parameterized verification. In: Emerson, E.A., Namjoshi, K.S. (eds.) VMCAI 2006. LNCS, vol. 3855, pp. 126–141. Springer, Heidelberg (2005)
12. Delzanno, G.: Automatic verification of cache coherence protocols. In: Wilhelm, R. (ed.) Informatics. LNCS, vol. 2000, Springer, Heidelberg (2001)
13. Emerson, E., Namjoshi, K.: On model checking for non-deterministic infinite-state systems. In: Proc. LICS' 98 (1998)
14. Esparza, J., Finkel, A., Mayr, R.: On the verification of broadcast protocols. In: Proc. LICS' 99, 14<sup>th</sup> IEEE Int. Symp., IEEE Computer Society Press, Los Alamitos (1999)
15. Fribourg, L., Richardson, J.: Symbolic verification with gap-order constraints. In: Gallagher, J.P. (ed.) LOPSTR 1996. LNCS, vol. 1207, Springer, Heidelberg (1997)
16. German, S.M., Sistla, A.P.: Reasoning about systems with many identical processes. Journal of the ACM 39(3), 675–735 (1992)
17. Kesten, Y., Maler, O., Marcus, M., Pnueli, A., Shahar, E.: Symbolic model checking with rich assertional languages. TCS 256, 93–112 (2001)
18. Lahiri, S.K., Bryant, R.E.: Indexed predicate discovery for unbounded system verification. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 135–147. Springer, Heidelberg (2004)
19. Lamport, L.: A new solution of dijkstra's concurrent programming problem. Commun. ACM 17(8), 453–455 (1974)
20. Revesz, P.: A closed form evaluation for datalog queries with integer (gap)-order constraints. Theoretical Computer Science 116(1), 117–149 (1993)
21. Ricart, G., Agrawal, A.K.: An optimal algorithm for mutual exclusion in computer networks. Communications of the ACM 24(1), 9–17 (1981)
22. Sedletsy, E., Pnueli, A., Ben-Ari, M.: Formal verification of the ricart-agrawala algorithm. In: Proc. FSTTCS'00 (2000)
23. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: Proc. LICS '86, June (1986)