A Complete Method for Symmetry Reduction in Safety Verification

Duc-Hiep Chu and Joxan Jaffar

National University of Singapore
hiepcd,joxan@comp.nus.edu.sg

Abstract. Symmetry reduction is a well-investigated technique to counter the state space explosion problem for reasoning about a concurrent system of similar processes. Here we present a general method for its application, restricted to verification of safety properties, but *without* any prior knowledge about global symmetry. We start by using a notion of *weak symmetry* which allows for more reduction than in previous notions of symmetry. This notion is relative to the target safety property. The key idea is to perform symmetric transformations on *state interpolation*, a concept which has been used widely for pruning in SMT and CEGAR. Our method naturally favors "quite symmetric" systems: more similarity among the processes leads to greater pruning of the tree. The main result is that the method is *complete* wrt. weak symmetry: it only considers states which are not symmetric to an already encountered state.

1 Introduction

Symmetry reduction is a well-investigated technique to counter the state space explosion problem when dealing with concurrent systems whose processes are similar. In fact, traditional symmetry reduction techniques rely on an idealistic assumption that processes are *indistinguishable*. Because this assumption excludes many realistic systems, there is a recent trend to consider systems of non-identical processes, but where the processes are *sufficiently similar* that the original gains of symmetry reduction can be still be accomplished, even though this necessitates an intricate step of detecting symmetry in the state exploration.

We start by considering an intuitive notion of symmetry which is based on a standard adaptation of the notion of bisimilarity. We say two states s_1 and s_2 are symmetric if there is a "permutation" π such that $s_2 = \pi(s_1)$, and if each successor state s_1' of s_1 is matched (via π) with a unique successor state s_2' of s_2 such that s_1' is symmetric with s_2' wrt. π. When we consider a safety property ϕ, we further require that computation from s_1 is a safe iff that from s_2 is.

We will call this notion *strong symmetry*. We mention that all recent works which deal with heterogeneous systems (where processes are not necessarily identical) have the desire to capture strong symmetry in the sense that they only consider states which are not strongly symmetric to an already encountered state.

In this paper, we present a general approach to symmetry reduction for safety verification of a parameterized system without any prior knowledge about its global symmetry. In particular, we explicitly explore the possible interleavings of the reachability tree, while applying pruning on "symmetric" subtrees. We now introduce a new notion of symmetry: *weak symmetry*. Informally, this notion weakens the notion of permutation between states so that *only the program counter* is used in consideration of symmetry. In contrast, values of program variables are used in consideration of strong symmetry. The main result is that our approach is *complete* wrt. weak symmetry: it only considers states which are not weakly symmetric to an already encountered state.

In more details, we address the state explosion problem by employing *symbolic learning* on the search tree of all possible interleavings. Specifically, our work is based on the concept of interpolation. Here, interpolation is essentially a form of *backward learning* where a completed search of a *safe* subtree is then formulated as a recipe for pruning (every state/node is a root associated to some subtree). There are two key ideas regarding our learning technique: First, each learned recipe for a node not only can be used to prune other nodes having the same future (same program point), but also can be *transfered* to prune nodes that having *symmetric* futures (symmetric program points). Second, each recipe discovered by a node will be conveyed back to its ancestors, which gives rise to pruning of *larger* subtree. Another important distinction is that our method learns *symbolically* with respect to the safety property and the interleavings. In the final section, we will confirm the effectiveness of weak symmetry experimentally on some classic benchmarks.

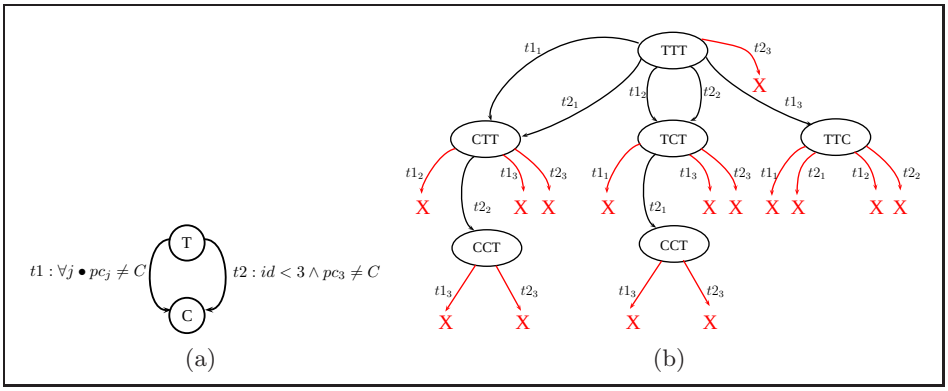

Fig. 1: Modified 3-process reader-writer example and its interleaving tree

We conclude this subsection with two examples in order to demonstrate strong and weak symmetry. First we borrow with modification from [13, 14] wherein are two "reader" processes (indices 1, 2) and one "writer" process (index 3). We denote by C and T the local process states which indicate entering the critical section and in a "trying" state, respectively. See Figure 1(a). Note that pc_j is is the local control location of process j and for each process, id is its *process identifier*. These concepts will be defined more formally in Section 2.

For each process, there are two transitions from T to C. The first, $t1$, is executable by any process provided that no process is currently in its critical section ($\forall j \bullet pc_j \neq C$). The second, $t2$, is however available to only readers ($id < 3$), and the writer must be in a non-critical local state $pc_3 \neq C$. This example shows symmetry between the reader processes, but because of their priority over the writer, we do not have *total* symmetry.

Figure 1(b) shows the full interleaving tree. Transitions are labelled with subscripts to indicate the process to which that transition is associated. *Infeasible* transitions are (red) arrows ending with (red) crosses. Note that nodes CTT and TCT are strongly symmetric, but neither is symmetric with TTC.

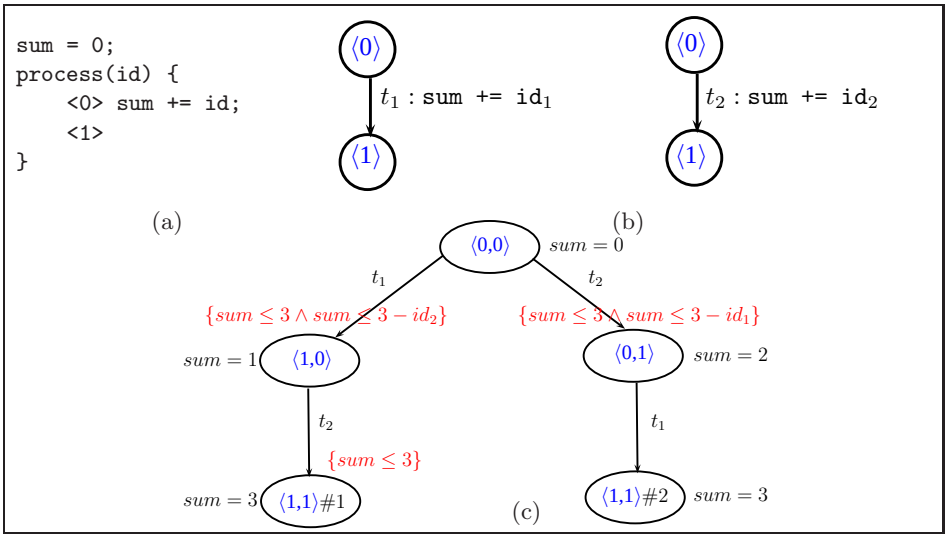

Fig. 2: (a) Sum-of-ids system (b) Its 2-process concretization (c) Full interleaving tree

Our second example is the system in Figure 2(a). Initially, the shared variable `sum` is set to 0. Each process increments `sum` by the amount of its process identifier, namely `id`. The local transition systems for process 1 and process 2 are shown in Figure 2(b). The full interleaving tree is shown in Figure 2(c).

Let π be the function swapping the indices of the two processes. We can see that the subtree rooted at states $\langle\langle 1,0\rangle; sum = 1\rangle$ and $\langle\langle 0,1\rangle; sum = 2\rangle$ share the same shape. However, due to the difference in the value of shared variable `sum`, strong symmetry does not apply (in fact, any top-down technique, such as [13, 14, 11], cannot avoid the subtree rooted at $\langle\langle 0,1\rangle; sum = 2\rangle$, even if the subtree rooted at $\langle\langle 1,0\rangle; sum = 1\rangle$ has been traversed and proved safe).

There is however a *weaker* notion of symmetry that does apply. We explain this by outlining our own approach, whose key feature is the computation of an *interpolant* for a node, by a process of backward learning. Informally, this interpolant represents a *generalization* of the values of the variables such that the traversed tree has the same transition structure, and also remains safe. In the example, we require the safety property $\psi \equiv sum \leq 3$ at every state, and interpolants are shown as formulas inside curly brackets.

The interpolant for state $\langle\langle 1,1\rangle; sum = 3\rangle$ is computed as $sum \leq 3$, and the interpolant for state $\langle\langle 1,0\rangle; sum = 1\rangle$ is computed as $\phi_{\langle 1,0\rangle} \equiv sum \leq 3 \wedge sum \leq 3 - id_2$. Using this, we then infer that $\phi_{\langle 0,1\rangle} \equiv sum \leq 3 \wedge sum \leq 3 - id_1$ (obtained by applying π on $\phi_{\langle 1,0\rangle}$) is a sound interpolant for program point $\langle 0,1\rangle$. As $\langle\langle 0,1\rangle; sum = 2\rangle \models \phi_{\langle 0,1\rangle}$, that subtree can be avoided.

1.1 Related Work

Symmetry reduction has been extensively studied, eg. [4, 2, 8, 5]. Symmetry is traditionally defined as a transition-preserving equivalence, where an automorphism π, other than being a bijection on the reachable states, also satisfies that (s, s') is a transition iff $(\pi(s), \pi(s'))$ is. There, this type of symmetry reduction is enforced by *unrealistic* assumptions about indistinguishable processes. As a result, it does not apply to many systems in practice.

One of the first to apply symmetry reduction strategies to "approximately symmetric" systems is [7], defining notions of *near* and *rough* symmetry. Near and rough symmetry is then generalized in [6] to *virtual symmetry*, which is considered as the most general condition that allows a *bisimilar symmetry quotient*. Though bisimilarity makes them suitable for full μ-calculus model checking, the main limitation of these approaches is that they exclude many systems, where bisimilarity to the quotient is simply not attainable. Also, these approaches work only for the verification of *symmetric properties*. No implementation is provided.

The work [11] allows arbitrary divergence from symmetry, and accounts for this divergence initially by conservative optimism, namely in the form of symmetric "super-structure". Specifically, transitions are added to the structure to achieve symmetry. A *guarded annotated quotient* (GAQ) is then obtained from the super-structure, where added transitions are marked. Loss in precision during exploration is prevented by means of frequent symmetry checks during runtime. This approach works well for programs with syntacticly specified static transition priority. However, in general, the GAQ needs to be *unwound* frequently to compensate for the loss in precision. This might affect the running time significantly as it might need to consider many combinations of transitions which do not belong to the original structure.

In comparison with our technique, this method has a clear advantage that it can handle arbitrary CTL* property. Nevertheless, our technique is more efficient both in space and time. Our technique is required to store an interpolant for each non-subsumed state, whereas in [11], a quotient edge might require multiple annotations. Furthermore, ours does not require a costly preprocessing the program text, such as in order to determine a symmetric super-structure. Also, extending [11] to symbolic model checking does not seem possible.

The most *recent* state-of-the-art regarding symmetry reduction, and also closest to our spirit, is the *lazy approach* proposed by [13, 14]. Here only safety verification is considered. This approach does not assume any prior knowledge about (global) symmetry. Indeed, they initially and lazily ignore the potential lack of symmetry. During the exploration, each encountered state is annotated with information about how symmetry is violated along the path leading to it. The idea

is that more similarity between component processes entails more compression is achieved.

In summary, the two main related works which are not restricted *a priori* on global symmetry are [11] and [13]. That is, these works allow the system to use process identifiers and therefore do not restrict the behaviors of individual processes. This is not the case with the previously mentioned works.

These works, [11] and [13], can be categorized as *top-down* techniques. Fundamentally, they look at the syntactic similarities between processes, and then come up with a reduced structure where symmetric states/nodes are merged into one abstract node. When model checking is performed, an abstract node might be concretized into a number of concrete nodes and each is checked one by one ([11] handles that by unwinding). For them, two symmetric parental nodes are not guaranteed to have correspondingly symmetric children. For us, by backward learning, we *ensure* that is the case. Consequently, and most importantly, they do not exponentially improve the runtime, only compress the state space.

Consider again the first example above (Figure 1). A top-down approach will consider TTC as a symmetric state of CTT, and all three states CTT, TCT, and TTC are merged as one abstract state. While having compaction, it is not the case that the search space traversed is of this compact size. As a non-symmetric state (TTC) is merged with other mutually symmetric states (CTT and TCT), in generating the successor abstract state, the parent abstract state is required to be concretized and both transitions $t2_2$ (emanating from CTT) and transition $t2_1$ (emanating from TCT) are considered (in fact, infeasible transition $t2_3$ is also considered). In general, compaction may not lead to any reduction in the search space.

We finally mention that we consider only safety properties because we wish to employ abstraction in the search process. And it is precisely a judicious use of abstraction that enables us to obtain more pruning in comparison with prior techniques. We prove this in principle by showing that we are *complete* wrt. weak symmetry, and we demonstrate this experimentally on some classic benchmarks.

2 Preliminaries

We consider a parameterized system composed of a fixed number n of processes. In accordance with standard practice in works on symmetry, we assume that the domain of discourse of the program variables is *finite* so as to guarantee termination of the search process of the underlying transition system. (Infinite domains may be accomodated by some use of abstraction, as we show in one benchmark example below.)

We employ the usual syntax of a deterministic imperative language, and communication occurs via shared variables. Each process has a unique and predetermined *process identifier*, and this is denoted parametrically in the system by the special variable `id`. Thus the concrete value of `id` for each individual process ranges from 1 to n. We note that the variable `id` cannot be changed. Even though the processes are defined by one parameterized system, their dy-

namic behaviors can be arbitrarily different. This would depend on how `id` is expressed in the parameterized system. Finally, we also allow a blocking primitive `await(b) s;` where `b` is a boolean expression and `s` is an *optional* program statement.

Consider the 2-process parameterized system in Figure 3(a). Note the (local) program points in angle brackets. Figure 3(b) "concretizes" the processes explicitly. Note the use in the first process of a local variable id_1 which is not writable in the process, and whose value is 1. Similarly for id_2 in the other process.

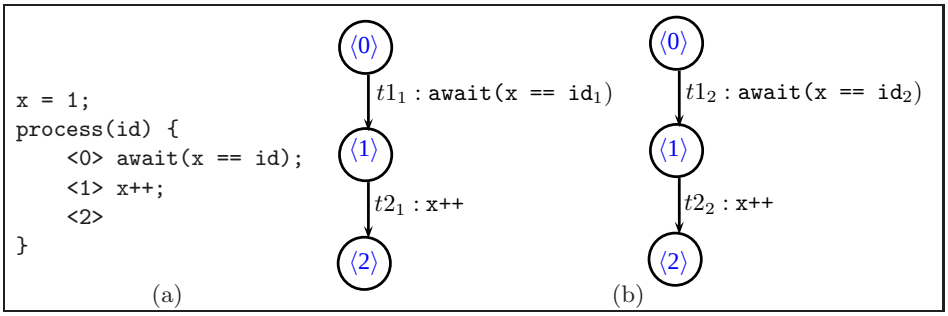

Fig. 3: (a) A parameterized system (b) Its 2-process concretization

In general, where P_i $(1 \leq i \leq n)$ is a process, let V_i be its local variables and V_{shared} be the shared variables of entire system. We note here that V_i does not include the special local variables which represent the process identifiers. Let $pc_i \in V_i$ be a special variable represent the local program counter, and the tuple $\langle pc_1, pc_2 \cdots, pc_n \rangle$ represent the global program point. Let $State$ be the set of all global states of the given program where $s_0 \in State$ is the initial state. A state $s \in State$ comprises of three parts: its *program point* $\texttt{pc}(s)$, which is a tuple of local program counters, its *valuation* over the program variables $\texttt{val}(s)$, and its valuation over the process identifiers `pids`. In other words, we denote a state s by $\langle \texttt{pc}(s); \texttt{val}(s); \texttt{pids} \rangle$. Note that all states from the same parameterized system share the same valuation of the individual process identifiers. Therefore, when the context is clear, we omit the valuation `pids` of a state.

We consider the *transitions* of states induced by the program. A transition t pertains to some process P_i. It transfers process P_i from control location $\texttt{l}_1$ to $\texttt{l}_2$. In general, the application of t is guarded by some condition `cond` (`cond` might be just `true`). At some state $s \in State$, when the i^{th} component of $\texttt{pc}(s)$, namely $\texttt{pc}(s)[i]$, equals $\texttt{l}_1$, we say that t can be *scheduled* at s. And when the valuation $\texttt{val}(s); \texttt{pids}$ satisifies the guard `cond`, denoted by $\texttt{val}(s); \texttt{pids} \models \texttt{cond}$, we say that t is *enabled* at s. Furthermore, we call the enabling condition of t the formula: $(\texttt{pc}(s)[i] \texttt{ == } \texttt{l}_1) \wedge \texttt{cond}$. For each state s, let $Scheduled(s)$ and $Enabled(s)$ denote the set of transitions which respectively can be scheduled at s and are enabled at s. Without further ado, we assume that the effect of applying an enabled transition t on a state s to arrive at state s' is well-understood. This is denoted as $s \xrightarrow{t} s'$.

Again consider in Figure 3 with two processes P_1 and P_2. with variables $\texttt{id}_1 = 1$ and $\texttt{id}_2 = 2$ respectively. In the system, it is specified parametrically that each

process awaits for `x == id`. In P_1, this is interpreted as $\texttt{await(x == id}_1\texttt{)}$ while P_2, this is interpreted as $\texttt{await(x == id}_2\texttt{)}$. Each process has 2 transitions: the first transfers it from control location $\langle 0 \rangle$ to $\langle 1 \rangle$, whereas the second transfers it from control location $\langle 1 \rangle$ to $\langle 2 \rangle$. Initially we have $x = 1$, i.e. the initial state s_0 is $\langle \langle 0, 0 \rangle; x = 1; id_1 = 1, id_2 = 2 \rangle$. We note that at s_0, both $t1_1$ and $t1_2$ can be scheduled. However, among them, only $t1_1$ is enabled. By taking transition $t1_1$, P_1 moves from control location $\langle 0 \rangle$ to $\langle 1 \rangle$, and the whole system moves from state $\langle \langle 0, 0 \rangle; x = 1; id_1 = 1, id_2 = 2 \rangle$ to state $\langle \langle 1, 0 \rangle; x = 1; id_1 = 1, id_2 = 2 \rangle$. We note that here the transition $t1_2$ is still disabled. From now on, we will omit the valuation of process identifiers. The whole system then takes the transition $t2_1$ and moves from state $\langle \langle 1, 0 \rangle; x = 1 \rangle$ to state $\langle \langle 2, 0 \rangle; x = 2 \rangle$. Now, $t1_2$ becomes enabled. Subsequently, the system takes $t1_2$ and $t2_2$ to move to state $\langle \langle 2, 1 \rangle; x = 1 \rangle$ and finally to state $\langle \langle 2, 2 \rangle; x = 3 \rangle$.

Definition 1 (Safety). *We say the given concurrent system is* safe *wrt. a safety property ψ if $\forall s \in State \bullet s$ is reachable from $s_0 \longrightarrow s \models \psi$.*

2.1 Symmetry

Given an n-process system, let $\mathcal{I} = [1 \cdots n]$ denote its *indices*, to be thought of as process identifiers. We write $Sym\ \mathcal{I}$ to denote the set of all permutations π on index set $\mathcal{I}$. Let Id be the identity permutation and π^{-1} the inverse of π.

For an indexed object b, such as a program point, a variable, a transition, valuation of program variables, or a formula, whose definition depends on $\mathcal{I}$, we can define the notion of permutation π acting on b, by simultaneously replacing each occurrence of index $i \in \mathcal{I}$ by $\pi(i)$ in b to get the result of $\pi(b)$.

Example 1. Consider the parameterized system in Figure 3. Let the permutation π swap the two indices $(1 \mapsto 2, 2 \mapsto 1)$. Applying π to the valuation $x = 1$ gives us $\pi(x = 1) \equiv x = 1$, as x is a shared variable. Applying π to the formula $x = id_1 \wedge id_1 = 1$ gives us $\pi(x = id_1 \wedge id_1 = 1) \equiv (x = id_2 \wedge id_2 = 1)$. On the other hand, applying π to the transition $t1_1 \equiv \texttt{await(x = id}_1\texttt{)}$ will result in $\pi(t1_1) \equiv t1_2 \equiv \texttt{await(x = id}_2\texttt{)}$.

Definition 2. *For $\pi \in Sym\ \mathcal{I}$ and state $s \in State$, $s \equiv \langle \mathtt{pc}(s); \mathtt{val}(s); \mathtt{pids} \rangle$, the application of π on s is defined as $\langle \pi(\mathtt{pc}(s)); \pi(\mathtt{val}(s)); \mathtt{pids} \rangle$,*

In other words, permutations do not affect the valuation of process identifiers.

Example 2. Consider again the parameterized system in Figure 3. Assume the π is the permutation swapping the 2 indices $(1 \mapsto 2, 2 \mapsto 1)$. We then can have $\pi(\langle \langle 1, 0 \rangle; x = 1; id_1 = 1, id_2 = 2 \rangle) \equiv \langle \langle 0, 1 \rangle; x = 1; id_1 = 1, id_2 = 2 \rangle$. Please note that while π has no effect on shared variable x and valuation of process identifiers id_1, id_2, it does permute the local program points.

Definition 3. *For $\pi \in Sym\ \mathcal{I}$, a safety property ψ is said to be* symmetric *wrt. π if $\psi \equiv \pi(\psi)$.*

We next present a traditional notion of symmetry.

Definition 4 (Strong Symmetry). *For $\pi \in Sym\ \mathcal{I}$, and a safety property ψ, for $s, s' \in State$, we say that s is π-similar to s' wrt. ψ, denoted by $s \overset{\pi,\psi}{\sim} s'$ if ψ is symmetric wrt. π and the following conditions hold:*
- $\pi(s) = s'$
- *for each transition t such that $s \xrightarrow{t} d$ we have $s' \xrightarrow{\pi(t)} d'$ and $d \overset{\pi,\psi}{\sim} d'$*
- *for each transition t' such that $s' \xrightarrow{t'} d'$ we have $s \xrightarrow{\pi^{-1}(t')} d$ and $d \overset{\pi,\psi}{\sim} d'$.*

One of the strengths of this paper is to allow symmetry by *disregarding* the values of the program variables.

Definition 5 (Weak Symmetry). *For $\pi \in Sym\ \mathcal{I}$, and a safety property ψ, for $s, s' \in State$, we say that s is π-similar to s' wrt. ψ, denoted by $s \overset{\pi,\psi}{\sim} s'$ if ψ is symmetric wrt. π and the following conditions hold:*
- $\pi(\mathit{pc}(s)) = \mathit{pc}(s')$
- $s \models \psi$ *iff* $s' \models \pi(\psi)$
- *for each transition t such that $s \xrightarrow{t} d$ we have $s' \xrightarrow{\pi(t)} d'$ and $d \overset{\pi,\psi}{\sim} d'$*
- *for each transition t' such that $s' \xrightarrow{t'} d'$ we have $s \xrightarrow{\pi^{-1}(t')} d$ and $d \overset{\pi,\psi}{\sim} d'$.*

We note here that, if s is π-similar to s' then s' is π^{-1}-similar to s. Therefore, if s is symmetric with s', then s' is symmetric with s also.

3 Motivating Examples

Figure 4 shows a parameterized system and its 2-process concretization. The shared array `x` contains 2 elements, initially 0. For convenience, we assume that array index starts from 1. Process 1 assigns $\mathtt{id}_1$ (whose value is 1) to `x`[1] while process 2 assigns $\mathtt{id}_2$ (2) to `x`[2].

Consider the safety property $\psi \equiv x[1] + x[2] \leq 3$, interpreted as $\psi \equiv x[id_1] + x[id_2] \leq 3$. The reachability tree explored is in Figure 4(c). Circles are used to denote states, while double-boundary circles denote subsumed/pruned states.

From the initial state $s_0 \equiv \langle\langle 0, 0\rangle; x[1] = 0, x[2] = 0; id_1 = 1, id_2 = 2\rangle$ process 1 progresses first and moves the system to the state $s_1 \equiv \langle\langle 1, 0\rangle; x[1] = 1, x[2] = 0; id_1 = 1, id_2 = 2\rangle$. From s_1, process 2 now progresses and moves the system to the state $s_2 \equiv \langle\langle 1, 1\rangle; x[1] = 1, x[2] = 2; id_1 = 1, id_2 = 2\rangle$. Note that s_0, s_1, and s_2 are all safe wrt. ψ. As there is no transition emanating from s_2, the interpolant for s_2 is computed as $\phi_2 \equiv \psi \equiv x[id_1] + x[id_2] \leq 3$. The pair $\langle\langle 1, 1\rangle; \phi_2\rangle$ is memoized. The interpolant for s_1 can be computed as a conjunction of two formulas. One concerns the safety of s_1 itself, and the other concerns the safety of the successor state from t_2. In other words, we can have $\phi_1 \equiv \psi \wedge \mathsf{pre}(x[id_2] = id_2; \psi)$, where $\mathsf{pre}(t; \phi)$ denotes a precondition wrt. to the program transition t and the postcondition ϕ. Consequently, we can have $\phi_1 \equiv \psi \wedge x[id_1] + id_2 \leq 3$. The pair $\langle\langle 1, 0\rangle; \phi_1\rangle$ is memoized.

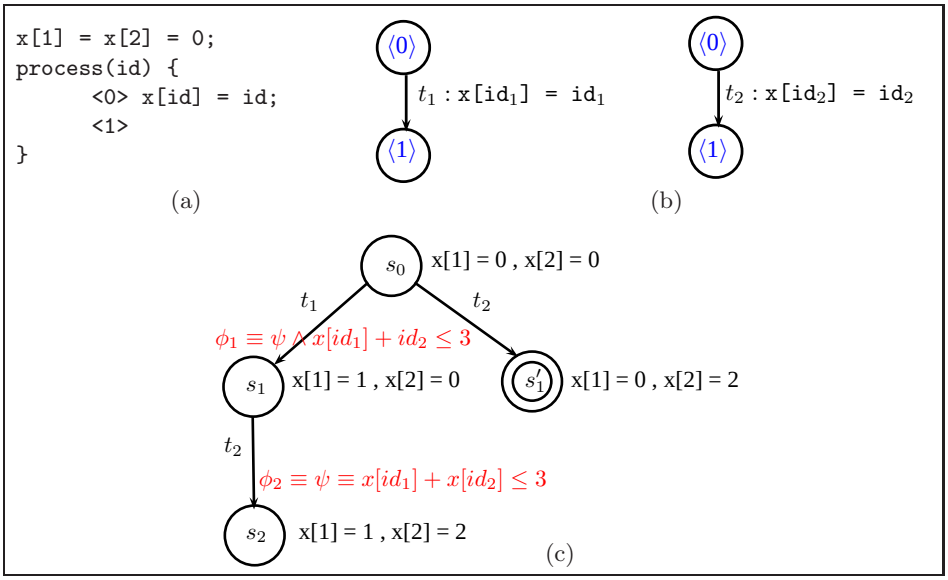

Fig. 4: (a) An example (b) Its 2-process concretization (c) Traversed tree

Now we arrive at state $s'_1 \equiv \langle\langle 0, 1\rangle; x[1] = 0, x[2] = 2; id_1 = 1, id_2 = 2\rangle$. This is indeed a symmetric image of state s_1 which we have explored and proved safe before. Here, we discover the permutation π to transform the program point $\langle 1, 0\rangle$ to program point $\langle 1, 0\rangle$. (Clearly π simply swaps the two indices.) We also observe that the safety property ψ is symmetric wrt. this π, i.e. $\pi(\psi) \equiv \psi$ (in the literature, we may also say that ψ is invariant wrt. π). In the next step, we check whether $\texttt{val}(s'_1)$ conjoined with `pids` implies the *transformed interpolant* $\pi(\phi_1)$. We have $\pi(\phi_1) \equiv \pi(x[id_1] + x[id_2] \leq 3 \wedge x[id_1] + id_2 \leq 3 \equiv x[id_2] + x[id_1] \leq 3 \wedge x[id_2] + id_1 \leq 3$. As $\texttt{val}(s'_1); \texttt{pids} \models x[id_2] + x[id_1] \leq 3 \wedge x[id_2] + id_1 \leq 3$, we do not need to explore s'_1 any further. In other words, the subtree rooted at s'_1 is pruned.

Another example is Figure 5. We are interested in safety property $\psi \equiv x < 2$. As x is a shared variable, ψ is symmetric wrt. all possible permutations.

The reachability tree is depicted in Figure 5(c). From the initials state s_0 we arrive at states s_1, s_2, and s_3, where:

$s_0 \equiv \langle\langle 0, 0, 0\rangle; x = 0; id_1 = 1, id_2 = 2, id_3 = 3\rangle$
$s_1 \equiv \langle\langle 1, 0, 0\rangle; x = 1; id_1 = 1, id_2 = 2, id_3 = 3\rangle$
$s_2 \equiv \langle\langle 1, 1, 0\rangle; x = 1; id_1 = 1, id_2 = 2, id_3 = 3\rangle$
$s_3 \equiv \langle\langle 1, 1, 1\rangle; x = 1; id_1 = 1, id_2 = 2, id_3 = 3\rangle$.

At s_3 we compute its interpolant $\phi_3 \equiv \psi \equiv x < 2$. In a similar manner as before, we compute the interpolant for s_2, which is $\phi_2 \equiv x < 2 \wedge id_3 \neq 1$. When we are at state $s'_2 \equiv \langle\langle 1, 0, 1\rangle; x = 1; id_1 = 1, id_2 = 2, id_3 = 3\rangle$, we look for a permutation π_1 such that $\pi_1(\langle 1, 1, 0\rangle) = \langle 1, 0, 1\rangle$. Clearly we can have π_1 as the permutation which fixes the first index and swaps the last 2 indices. Moreover, $\texttt{val}(s'_2); \texttt{pids} \equiv x = 1; id_1 = 1, id_2 = 2, id_3 = 3 \models \pi_1(\phi_2) \equiv x < 2 \wedge id_2 \neq 1$. Therefore, s'_2 is pruned.

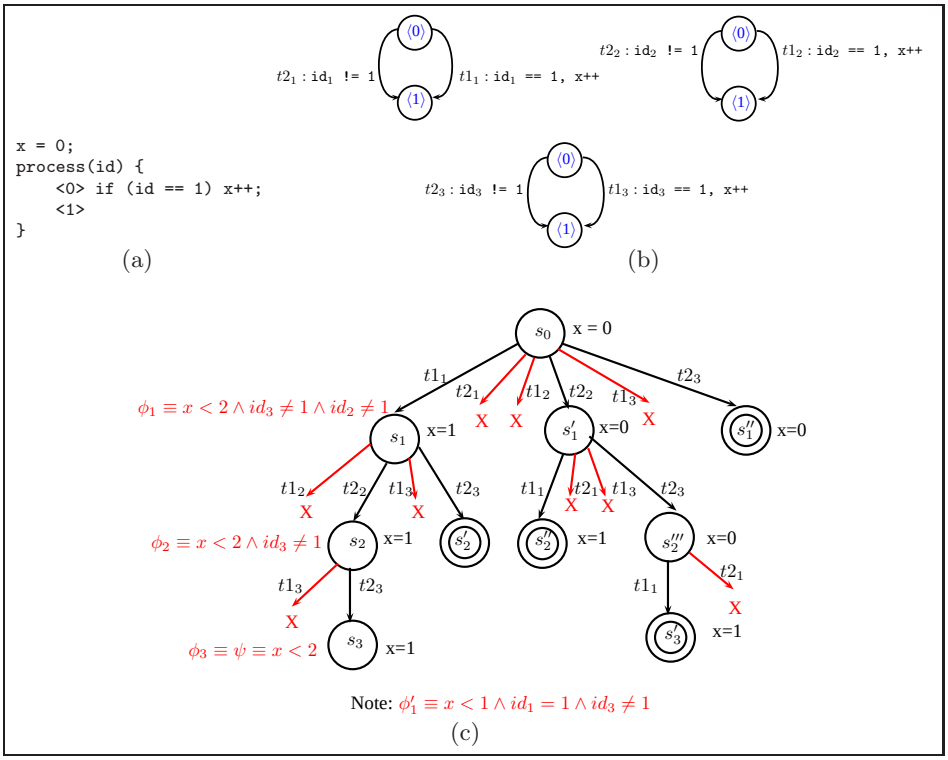

Fig. 5: (a) An example (b) Its 3-process concretization (c) Traversed tree

Similarly, the interpolant for s_1 is computed as $\phi_1 \equiv x < 2 \wedge id_2 \neq 1 \wedge id_3 \neq 1$. When at state $s_1' \equiv \langle\langle 0,1,0\rangle; x = 0; id_1 = 1, id_2 = 2, id_3 = 3\rangle$, we look for a permutation π_2 such that $\pi_2(\langle 1,0,0\rangle) = \langle 0,1,0\rangle$. Clearly we can have π_2 as the permutation which fixes the third index and swaps the first two indices. However, $\texttt{val}(s_1');\texttt{pids} \equiv x = 0; id_1 = 1, id_2 = 2, id_3 = 3 \not\models \pi_2(\phi_1) \equiv x < 2 \wedge id_1 \neq 1 \wedge id_3 \neq 1$. Thus the subtree rooted at s_1' cannot be pruned and it requires further exploration. After having been traversed, the intepolant for s_1' is computed as $\phi_1' \equiv x < 1 \wedge id_1 = 1 \wedge id_3 \neq 1$. Next we arrive at $s_1'' \equiv \langle\langle 0,0,1\rangle; x = 0; id_1 = 1, id_2 = 2, id_3 = 3\rangle$. We can find a permutation π_3 which fixes the first index and swaps the last 2 indices ($\pi_3 \equiv \pi_1$). We have $\pi_3(\langle 0,1,0\rangle) = \langle 0,0,1\rangle$. Also $\texttt{val}(s_1'');\texttt{pids} \equiv x = 0; id_1 = 1, id_2 = 2, id_3 = 3 \models \pi_3(\phi_1') \equiv x < 1 \wedge id_1 = 1 \wedge id_2 \neq 1$. As a result, we can avoid considering the subtree rooted at s_1''.

In the two above examples, we have shown how the concept of backward learning with interpolation can help capture the shape of a subtree. More importantly, computed interpolants can be transformed in order to detect the symmetry as well as the non-symmetry (mainly due to the use of `id`) between candidate subtrees.

4 State Interpolation

State-based interpolation was first described in [9] for finite transition systems. The essential idea was to prune the search space of symbolic execution, informally described as follows. Symbolic execution is usually depicted as a tree rooted at the initial state s_0 and for each state s_i therein, the descendants are just the states obtainable by extending s_i with an enabled transition. Consider one particular feasible path represented in the tree: $s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} s_2 \cdots s_m$. The boundary between each transition can be considered a *program point*, characterizing a point in the reachability tree in terms of all the remaining possible transitions. Now, this particular path is *safe* wrt. to safety property ψ if for all i, $0 \leq i \leq m$, we have $s_i \models \psi$. A (state) interpolant at program point j, $0 \leq j \leq m$ is simply a set of states S_j containing s_j such that for any state $s'_j \in S$, $s'_j \xrightarrow{t_{j+1}} s'_{j+1} \xrightarrow{t_{j+2}} s'_{j+2} \cdots s'_m$, it is also the case that for all i, $j \leq i \leq m$, we have $s'_i \models \psi$. This interpolant was constructed at point j due to the one path. Consider now all paths from s_0 and with prefix $t_1, \cdots, t_{j-1}$. Compute each of their interpolants. Finally, we say that the interpolant for the subtree of paths just considered is simply the intersection of all the individual interpolants. This notion of interpolant for a subtree provides a notion of *subsumption* because we can now prune a subtree in case the state rooted at this subtree are within the interpolant computed for some previously encountered subtree of the same program point.

Definition 6 (Safe Root). *Let s_i be a state which is reachable from the initial state s_0, we say that s_i is a* safe root*, denoted by $\triangle(s_i)$, if all states s'_i reachable from s_i is safe.*

Definition 7 (State Coverage). *Let s_i and s_j be two states which are reachable from the initial state s_0 such that $\mathtt{pc}(s_i) \equiv \mathtt{pc}(s_j)$. We say that s_i covers s_j, denoted by $s_i \succeq s_j$ if $\triangle(s_i) \rightarrow \triangle(s_j)$.*

During the traversal of the reachability tree, if we detect that $s_i \succeq s_j$ while s_i has been proved to be a safe root, the traversal of the subtree rooted at s_j can be avoided. We thus reduce the search space.

In practice, in order to determine state coverage, during the exploration of subtree rooted at s_i we compute a state-interpolant of s_i, denoted as $\mathtt{SI}(s_i, \psi)$. Note that trivially, we should have $s_i \models \mathtt{SI}(s_i, \psi)$. Furthermore, $\mathtt{SI}(s_i, \psi)$ ensures that for all state s_j at program point $\mathtt{pc}(s_i)$, if $s_j \models \mathtt{SI}(s_i, \psi)$ then for all $t \in \mathit{Scheduled}(s_i)$ (note that $\mathit{Scheduled}(s_i) \equiv \mathit{Scheduled}(s_j)$) the two following conditions must be satisfied:

- if t was disabled at s_i, it also must be disabled at s_j
- if t was enabled at s_j (by the above condition, it must be enabled at s_i too) and $s_j \xrightarrow{t} s'_j$ and $s_i \xrightarrow{t} s'_i$, then s'_i must cover s'_j.

This observation enables us to determine the coverage relation as the form of backward learning in a recursive manner. Our symmetry reduction algorithm presented in Section 5 will implement this idea of state interpolation.

5 Symmetry Reduction Algorithm

```
⟨1⟩  Intially : stack = ∅; Explore(s_0)
function Explore(s)
⟨2⟩  if s ⊭ ψ Report Error and TERMINATE
⟨3⟩  if ∃π • π(ψ) ≡ ψ ∧ memoed(PC, φ) ∧ pc(s) ≡ π(PC) ∧ s ⊨ π(φ) return π(φ)
⟨4⟩  if s ∈ stack return true else stack.push(s)
⟨5⟩  φ := ψ
⟨6⟩  foreach t in Scheduled(s) do
⟨7⟩      if t in Enabled(s)
⟨8⟩          s' := succ(s) after t     /* Execute t */
⟨9⟩          φ' := Explore(s')
⟨10⟩         φ := φ ∧ pre(t; φ')
         else
⟨11⟩         φ := φ ∧ pre(t; false)
         endif
⟨12⟩ endfor
⟨13⟩ stack.pop()
⟨14⟩ memo(pc(s), φ) and return φ
end function
```

Fig. 6: Symmetry Reduction Algorithm (DFS)

Our algorithm, presented in Figure 6, naturally performs a depth first search of the interleaving tree. It assumes the safety property to be known as ψ. Initially, *stack* is initialized as empty and we explore the initial state s_0. During the search process, the function Explore will be recursively called.

Base Cases: The first base case is when the current state does not conform to the safety property ψ (line 2). We then immediately report an error and terminate. The second base case applies when the current state (subtree) has a symmetric image (subtree) which has already been traversed and proved safe before (line 3). This case corresponds to a subsumed node. The third base case, we make use of *stack* to handle cycles (line 4). Termination is ensured due to finite setting.

Recursive Traversal and Computing the Interpolants: Our algorithm recursively explores the successors of the current state by the recursive call in line 9. The interpolant ϕ for the current state is computed as from line 5 - 11. The operation $\mathsf{pre}(t; \overline{\phi})$ denotes the precondition computation wrt. the program transition t and the postcondition $\overline{\phi}$. In practice, we implement this as an estimation of the weakest precondition computation [3].

Theorem 1 (Soundness). *Our symmetry reduction algorithm is sound.*

*Proof (**Outline**).* Let the triple $\{\phi\}\ \langle\langle pc_1, pc_2, \cdots, pc_n\rangle; P_1||P_2||\cdots||P_n\rangle\ \{\psi\}$ denote the fact that ϕ is a *sound* interpolant for program point $\langle pc_1, pc_2, \cdots, pc_n\rangle$ wrt. the safety property ψ and the concurrent system $P_1||P_2||\cdots||P_n$. Due to space limit, we will not prove that our interpolant computation (line 5-11) is a sound computation. Instead, we refer interested readers to [9]. Let us assume

that the soundness of that triple is witnessed by the proof $\mathcal{P}$. By consistently renaming $\mathcal{P}$ with the renaming function π, we can derive a new *sound* fact (ie. a proof), which is

$\{\pi(\phi)\}\ \pi(\langle\langle pc_1, pc_2, \cdots, pc_n\rangle; P_1||P_2||\cdots||P_n\rangle)\ \{\pi(\psi)\} \equiv$
$\{\pi(\phi)\}\ \langle\langle pc_{\pi(1)}, pc_{\pi(2)}, \cdots, pc_{\pi(n)}\rangle; P_{\pi(1)}||P_{\pi(2)}||\cdots||P_{\pi(n)}\rangle\ \{\pi(\psi)\}$

Since $P_1, P_2, \cdots P_n$ come from the same parameterized system, we have

$P_{\pi(1)}||P_{\pi(2)}||\cdots||P_{\pi(n)} \equiv P_1||P_2||\cdots||P_n$

Therefore, $\{\pi(\phi)\}\ \langle\langle pc_{\pi(1)}, pc_{\pi(2)}, \cdots, pc_{\pi(n)}\rangle; P_1||P_2||\cdots||P_n\rangle\ \{\pi(\psi)\}$ must hold too. In the case that ψ is symmetric wrt. π, we have $\pi(\phi)$ is a *sound* interpolant for program point $\langle pc_{\pi(1)}, pc_{\pi(2)}, \cdots, pc_{\pi(n)}\rangle$ wrt. the same safety property ψ and the same concurrent system $P_1||P_2||\cdots||P_n$. As a result, the use of interpolant $\pi(\phi)$ at line 3 in our algorithm is *sound*. □

Definition 8 (Completeness). *In proving a parameterized system with global state space State is safe wrt. a property ψ, an algorithm is said to be* complete *wrt. a symmetry relation $\mathcal{R}$ iff for all $s, s' \in State$, $s\ \mathcal{R}\ s'$ implies that the algorithm will avoid traversing either the subtree rooted at s or the subtree rooted at s'.*

Definition 9 (Symmetry Preserving Precondition Computation). *Given a parameterized system and a safety property ψ, the precondition computation* pre *used in our algorithm is said to be* symmetry preserving *if for all $\pi \in Sym\ \mathcal{I}$, for all transition t and all possible interpolants $\phi \bullet \pi(\mathit{pre}(t;\phi)) \equiv \mathit{pre}(\pi(t);\pi(\phi))$.*

This property means that our precondition computation is consistent wrt. to renaming operation. A reasonable implementation of pre can always ensure this.

Definition 10 (Monotonic Precondition Computation). *Given a parameterized system and a safety property ψ, the precondition computation* pre *used in our algorithm is said to be* monotonic *if for all transition t and all possible interpolants $\phi_1, \phi_2 \bullet \phi_1 \rightarrow \phi_2$ implies $\mathit{pre}(t;\phi_1) \rightarrow \mathit{pre}(t;\phi_2)$.*

We emphasize here that the weakest precondition computation [3] does possess the monotonicity property. As is well-known, computing the weakest precondition in all the cases is very expensive. However, in practice (and in particular in the experiments we have performed), we typically observe this property in the implementation of the precondition computation. Incidentally, some possible implementations for this operation are discussed in [1, 9, 10].

Theorem 2 (Completeness). *Our symmetry reduction algorithm is complete wrt. the weak symmetry relation if our operation* pre *is both monotonic and symmetry preserving.*

*Proof (**Outline**).* Assume that $s, s' \in State$ and s is weakly π-similar to s'. Assume we encounter s first. If the subtree rooted at s is avoided (due to subsumption), the theorem trivially holds. W.l.o.g. we assume that the subtree rooted at s is traversed first. The theorem also trivially holds if s is *not* a safe root. Now we

consider that the subtree rooted as s is proved safe and the returned interpolant is ϕ. We will prove by structural induction on that interpolated subtree.

For the base case that ϕ is true (line 4), it trivially holds that $s' \models \pi(\phi)$. For the base case that ϕ is ψ (when there is no schedulable transition from s) due to the definition of weak symmetry relation, there is no schedulable transition from s' and $s' \models \pi(\psi)$. Therefore, the subtree rooted at s' is avoided.

As the induction hypothesis, assume now that the theorem holds for all the descendants of state s. Let assume that $\phi \equiv \psi \wedge \phi_1 \wedge \phi_2 \wedge \cdots \wedge \phi_k \wedge \phi_{k+1} \wedge \cdots \wedge \phi_m$, where $\phi_1 \cdots \phi_k$ are the interpolants contributed by enabled transitions in s and $\phi_{k+1} \cdots \phi_m$ are the interpolants contributed by schedulable but disabled transitions in s (line 10 and 11). Now assume that state s' violates the subsumption test, which means that $s' \not\models \pi(\phi)$. Using the first condition of weak symmetry relation, obviously $s' \models \pi(\psi)$. As such, the must exist some $1 \leq j \leq m$ such that $s' \not\models \pi(\phi_j)$. There are two possible cases: (1) ϕ_j is contributed by an enabled transition; (2) ϕ_j contributed by a schedulable but disabled transition.

Let us consider case (1) first. Assume ϕ_j corresponds to transition $t \in Enabled(s)$ and $s \xrightarrow{t} d$. By definition we have $s' \overset{\pi(t)}{\rightarrow} d'$ and d is π-similar to d'. Let ϕ_d be interpolant for the subtree rooted at d. By induction hypothesis, we have $d' \models \pi(\phi_d)$. Obviously, we have $s' \models \mathsf{pre}(\pi(t); d')$, by monotonicity of pre, we deduce $s' \models \mathsf{pre}(\pi(t); \pi(\phi_d))$. As pre is symmetry preserving, $s' \models \mathsf{pre}(\pi(t); \pi(\phi_d)) \equiv \pi(\mathsf{pre}(t; \phi_d)) \equiv \pi(\phi_j)$. Consequently we arrive at the fact that $s' \models \pi(\phi_j)$ which is a contradiction.

For case (2), by using the symmetry preserving property of pre (and note that $\pi(false) \equiv false$), we also derive a contradiction. □

6 Experimental Evaluation

We used a 3.2 GHz Intel processor and 2GB memory running Linux. Unless otherwise mentioned, timeout is set at 10 minutes, and '-' indicates timeout.

	CSR			RSR			NSR		
# Phil	Visited	Subsumed	T(s)	Visited	Subsumed	T(s)	Visited	Subsumed	T(s)
4	230	134	0.09	328	184	0.13	1246	702	0.81
5	662	446	0.28	1509	981	0.71	7517	4893	4.93
6	1778	1304	0.85	7356	5216	4.18	43580	30908	34.53
7	4584	3552	2.55	35079	26335	28.83	–	–	–
8	11526	9281	7.54	–	–	–	–	–	–
9	28287	23432	22.6	–	–	–	–	–	–
10	67920	57504	58.07	–	–	–	–	–	–
11	159738	137609	226.86	–	–	–	–	–	–

Table 1. Experiments on Dining Philosophers

Our first example is the classic *dining philosophers* problem. It exhibits *rotational* symmetry; importantly, we exploit more symmetry. We verify a *tight* safety property that 'no more than *half* the philosophers can eat simultaneously'.

Table 1 presents three variants: Complete Symmetry Reduction (CSR), Rotational Symmetry Reduction (RSR), and No Symmetry Reduction (NSR). The

number of *stored states* is the difference between the number of visited states (Visisted column) and subsumed states (Subsumed column). Note that although RSR achieves linear reduction compared to NSR, it does not scale well. CSR significantly outperforms RSR and NSR in all the instances.

		Complete Symmetry Reduction			Lazy Symmetry Reduction	
# Readers	# Writers	Visisted	Subsumed	T(s)	Abstract States	T(s)
2	1	35	20	0.01	9	0.01
4	2	226	175	0.19	41	0.10
6	3	779	658	0.93	79	67.80
8	4	1987	1750	3.23	165	81969.00
10	5	4231	3820	9.21	–	–

Table 2. Experiments on Reader-Writer Protocol

Next consider the *Reader-Writer Protocol* from [13, 14]. Here we highlight the aspect of *search space size* as compared to top-down techniques, of which Lazy Symmetry Reduction [14] is chosen as a representative. Table 2 shows that although lazy symmetry reduction has aggressively compressed the state space (which now grows roughly linearly), the running time is still *exponential*. In other words, the abstract states is not representative of the search space. In contrast, our running significantly outperforms [14]. Note that only in the instance of 8 readers and 4 writers, we extended the timeout; even so [14] takes almost 1 day to finish.

	Complete Symmetry Reduction			SPIN-NSR		
# Processes	Visited	Subsumed	T(s)	Visited	Subsumed	T(s)
10	57	45	0.02	6146	4097	0.03
20	212	190	0.04	11534338	9437185	69.70
40	822	780	0.37	–	–	–
60	1832	1770	1.91	–	–	–
80	3242	3160	7.62	–	–	–
100	5052	4950	22.09	–	–	–

Table 3. Experiments on sum-of-ids Example

Next we have the 'sum-of-ids' example mentioned earlier. To the best of our knowledge, there is no symmetry reduction algorithm which can detect and exploit symmetry here. Table 3 shows we have significant symmetry reduction. In term of memory (stored states), we enjoy linear complexity. For reference, we also report the running times of this example, without symmetry reduction, using SPIN 5.1.4 [12].

	Complete Symmetry Reduction			SI		
# Processes	Visited	Subsumed	T(s)	Visited	Subsumed	T(s)
3	65	31	0.10	265	125	0.43
4	182	105	0.46	1925	1089	5.89
5	505	325	2.26	14236	9067	74.92
6	1423	983	11.10	–	–	–

Table 4. Experiments on Bakery Algorithm

In the fourth and last example, we apply our method to handle infinite domain variables and loops. We choose the well-known Bakery Algorithm to perform the experiments, and we use the well-known abstraction of using an inequality to describe each pair of counters. Again, as far as we are aware of, there has been no symmetry reduction algorithm which can detect and exploit symmetry for this example. Table 4 shows the significant improvements due to our symmetry reduction, compared to just symbolic execution with interpolation (`SI`).

7 Conclusion

We presented a method of symmetry reduction for searching the interleaving space of a concurrent system of transitions in pursuit of a safety property. The class of systems considered, by virtue of being defined parametrically, is completely general; the individual processes may be at any level of similarity to each other. We then enhanced a general method of symbolic execution with interpolation for traditional safety verification of transition systems, in order to deal with symmetric states. We then defined a notion of weak symmetry, one that allows for more symmetry than the stronger notion that is used in the literature. Finally, we showed that our method, when employed with an interpolation algorithm which is monotonic, can exploit weak symmetry completely.

References

1. D.H. Chu and J. Jaffar. Symbolic simulation on complicated loops for WCET path analysis. In *EMSOFT*, pages 319–328, 2011.
2. E. M. Clarke, T. Filkorn, and S. Jha. Exploiting symmetry in temporal logic model checking. In *CAV*, pages 450–462, 1993.
3. E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM*, pages 453–457, August 1975.
4. E. A. Emerson and A. P. Sistla. Model checking and symmetry. In *CAV'93*.
5. E. A. Emerson and A. P. Sistla. Utilizing symmetry when model-checking under fairness assumptions. *ACM TOPLAS*, pages 617–638, July 1997.
6. E. A. Emerson, J W. Havlicek, and R. J. Trefler. Virtual symmetry reduction. In *LICS*, pages 121–131,2000.
7. E. A. Emerson and R. J. Trefler. From asymmetry to full symmetry: New techniques for symmetry reduction in model checking. In *CHARME'99*.
8. C. N. Ip and D. L. Dill. Better verification through symmetry. *Formal Methods in System Design*, pages 41–75, 1996.
9. J. Jaffar, J. E. Santosa, and R. Voicu. An interpolation method for CLP traversal. In *CP'09*, pages 454–469.
10. A. Rybalchenko and V. S. Stokkermans. Constraint solving for interpolation. In *VMCAI*, pages 346–362, 2007.
11. A. P. Sistla and P. Godefroid. Symmetry and reduced symmetry in model checking. *ACM TOPLAS*, pages 702–734, July 2004.
12. SPIN Model Checker. *http://spinroot.com*.
13. T. Wahl. Adaptive symmetry reduction. In *CAV*, pages 393–405, 2007.
14. T. Wahl and V. D'Silva. A lazy approach to symmetry reduction. *Form. Asp. Comput.*, pages 713–733, November 2010.