

Programming a Fast Explicit Conflict Checker

Robi Malik

Abstract—This paper describes the implementation of *explicit model checking algorithms* to verify the *nonblocking* or *nonconflicting* property of discrete event systems. Explicit algorithms enumerate and store all reachable states of a synchronous composition. Three alternatives optimised for memory consumption or runtime are described and compared. The algorithms have been implemented in C++ in the discrete event systems library Waters, and experimental results show that they can explore more than 100 million states on standard computers.

I. INTRODUCTION

The nonconflicting property is a weak liveness property commonly used in *supervisory control theory* of discrete event systems to capture the absence of livelocks and deadlocks [1]. A system is nonconflicting if, from any reachable system state, some terminal state is reachable. A lot of effort has been put into the development of algorithms for the automatic verification of this crucial property. Systems of considerable size can be verified using *symbolic model checking* [2] or *compositional verification* [3]–[6].

This paper is concerned with *explicit* verification of the nonconflicting property. Explicit algorithms are the simplest model checking algorithms that construct all reachable states and explore every single transition. As all states are stored in memory, they are only suitable for small to medium-size model checking problems with a few million reachable states, which they may solve faster than more complicated algorithms. Compositional verification [3]–[6], which works by simplifying larger models, relies on an explicit or symbolic algorithm in its final step.

This paper describes the explicit verification algorithms in the Waters library, which is part of Supremica [7], and which can explore state spaces with more than 100 million states. Existing discrete event systems software such as TCT [8] and libFaudes [9] performs conflict check by synchronous composition and nonblocking verification, which is inferior to the direct algorithms shown here. The SPIN model checker [10] has powerful explicit algorithms, but its modelling language neither supports synchronisation with events shared by more than two components nor the nonconflicting property.

In the following, Section II provides the background of finite-state machines and the nonconflicting property. Then Section III describes three design alternatives of explicit conflict check algorithms, and Section IV evaluates them using experiments. Section V adds concluding remarks.

The author is with the Department of Computer Science, The University of Waikato, Hamilton, New Zealand (robi@waikato.ac.nz).

II. PRELIMINARIES

A *finite-state machine (FSM)* is a tuple $G = \langle \Sigma, Q, Q^\circ, Q^\omega, \rightarrow \rangle$, where Σ is a finite set of *events*, Q is a finite set of *states*, $Q^\circ \subseteq Q$ is the set of *initial states*, $Q^\omega \subseteq Q$ is the set of *accepting states*, and $\rightarrow \subseteq Q \times \Sigma \times Q$ is the *transition relation*.

The transition relation is written in infix notation, where $x \xrightarrow{\sigma} y$ means the existence of a transition from state $x \in Q$ to $y \in Q$ with event $\sigma \in \Sigma$. This notation is extended to *traces* $s \in \Sigma^*$ in the standard way. Given state sets $X, Y \subseteq Q$, the notation $X \xrightarrow{s} Y$ means $x \xrightarrow{s} y$ for some states $x \in X$ and $y \in Y$, and $X \xrightarrow{s}$ means $X \xrightarrow{s} Q$, and $X \rightarrow Y$ means $X \xrightarrow{s} Y$ for some $s \in \Sigma^*$. Events not in the event set of an FSM are assumed to be always enabled without state change, so the transition relation is further extended by letting $x \xrightarrow{\sigma} x$ for all $x \in Q$ and $\sigma \notin \Sigma$.

The FSM G is *deterministic* if $|Q^\circ| \leq 1$ and if $x \xrightarrow{\sigma} y_1$ and $x \xrightarrow{\sigma} y_2$ always implies $y_1 = y_2$. The *language* of G is $\mathcal{L}(G) = \{s \in \Sigma^* \mid Q^\circ \xrightarrow{s}\}$, and its *accepting language* is $\mathcal{L}^\omega(G) = \{s \in \Sigma^* \mid Q^\circ \xrightarrow{s} Q^\omega\}$. The *prefix-closure* of a language $L \subseteq \Sigma^*$ is $\text{pre}(L) = \{s \in \Sigma^* \mid st \in L \text{ for some } t \in \Sigma^*\}$.

FSMs are synchronised in lock-step [11]. The *synchronous composition* of two FSMs $G_1 = \langle \Sigma_1, Q_1, Q_1^\circ, Q_1^\omega, \rightarrow_1 \rangle$ and $G_2 = \langle \Sigma_2, Q_2, Q_2^\circ, Q_2^\omega, \rightarrow_2 \rangle$ is

$$G_1 \parallel G_2 = \langle \Sigma_1 \cup \Sigma_2, Q_1 \times Q_2, Q_1^\circ \times Q_2^\circ, Q_1^\omega \times Q_2^\omega, \rightarrow \rangle$$

where $(x_1, x_2) \xrightarrow{\sigma} (y_1, y_2)$ if $x_1 \xrightarrow{\sigma_1} y_1$ and $x_2 \xrightarrow{\sigma_2} y_2$.

Example 1: Fig. 1 shows a faulty version of the “small factory” system [1], modelled with four FSMs M_1 , B , R , and M_2 , and its synchronous composition. The synchronised states represent combinations of the states each FSM is in. For example, 1100 is the state tuple $(1, 1, 0, 0)$, where M_1 and B are in their states 1 and R and M_2 are in their states 0.

This paper is concerned with the nonblocking property of the synchronous composition of several FSMs, which is also referred to as nonconflicting. An FSM is nonblocking, if accepting states are reachable from all reachable states. More precisely, given $G = \langle \Sigma, Q, Q^\circ, Q^\omega, \rightarrow \rangle$, a state $x \in Q$ is *reachable* in G if $Q^\circ \rightarrow x$, and *coreachable* in G if $x \rightarrow Q^\omega$.

Definition 1: An FSM G is nonblocking if and only if every reachable state in G is also coreachable in G . FSMs G_1, \dots, G_n are *nonconflicting* if their synchronous composition $G_1 \parallel \dots \parallel G_n$ is nonblocking.

Example 2: State 1100 in Fig. 1 is coreachable because $1100 \xrightarrow{f_1 s_2 f_2} 0000 \in Q^\omega$, but states 0100 and 2100 are not coreachable. By Def. 1, this FSM is *blocking*, and thus M_1 , B , R , and M_2 are *conflicting*.

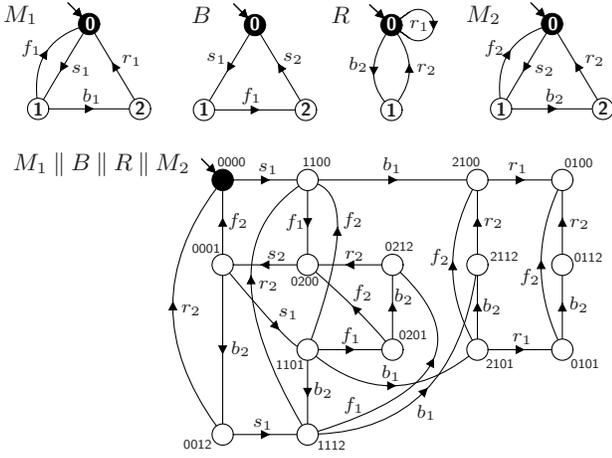


Fig. 1. Blocking small factory.

If G is a deterministic FSM, then its nonblocking property can be stated equivalently in language-based form as $\text{pre}(\mathcal{L}^\omega(G)) = \mathcal{L}(G)$ [1]. Then a system G is nonblocking, or nonconflicting, if every trace in its language $\mathcal{L}(G)$ can be continued to some trace in its accepting language $\mathcal{L}^\omega(G)$.

This paper distinguishes two kinds of blocking or conflict. A state $x \in Q$ is a *deadlock* state if $x \notin Q^\omega$ and for all $y \in Q$ such that $x \rightarrow y$ it holds that $x = y$, i.e., x is a non-accepting state with only selfloop transitions outgoing. For example, 0100 in Fig. 1 is a deadlock state. This is different from the usual definition [11], because here a deadlock state can have selfloops. A state $x \in Q$ is a *livelock* state if x is not coreachable and not a deadlock state.

III. NONBLOCKING VERIFICATION ALGORITHMS

Based on Def. 1, a conflict check can be done by exploring the reachable and coreachable states of a synchronous composition. Section III-A below shows how to compute the set of reachable states, and Section III-B shows how to find the coreachable states. Afterwards, Section III-C describes an alternative that avoids computing coreachable states, and Section III-D discusses counterexample computation.

A. Exploring Reachable States

A basic algorithm to explore the synchronous composition of n deterministic FSMs with exactly one initial state is shown in Algorithm 1. It performs a standard search, using two state sets Q and $Open$. The state set Q collects all state tuples of the synchronous composition, while $Open$ contains states that are yet to be expanded.

First, lines 1–2 add the initial state tuple to both Q and $Open$. Then the loop in line 3 expands each state tuple x in $Open$. For each event σ , line 6 determines whether it is enabled in x , and if so, line 7 computes the successor state y . If y is not already contained in the state space Q , then it is a new state and line 9 adds it to Q and to $Open$ for subsequent expansion. What information about transitions, if any, is recorded in line 10 depends on the second pass of the algorithm, as discussed in Section III-B below. When Algorithm 1 terminates, Q contains all reachable state tuples.

Algorithm 1: Reachability

Input: Deterministic FSMs

$G_i = \langle \Sigma_i, Q_i, \{x_i^\circ\}, Q_i^\omega, \rightarrow_i \rangle$ for $1 \leq i \leq n$

Output: Reachable states Q

```

1  $x^\circ \leftarrow (x_0^\circ, \dots, x_n^\circ)$ ;
2  $Q \leftarrow \{x^\circ\}$ ;  $Open \leftarrow \{x^\circ\}$ ;
3 while  $Open \neq \emptyset$  do
4   remove  $x = (x_0, \dots, x_n)$  from  $Open$ ;
5   foreach  $\sigma \in \Sigma_1 \cup \dots \cup \Sigma_n$  do
6     if  $x_i \xrightarrow{\sigma}_i y_i$  for each  $1 \leq i \leq n$  then
7        $y \leftarrow (y_1, \dots, y_n)$ ;
8       if  $y \notin Q$  then
9          $Q \leftarrow Q \cup \{y\}$ ;  $Open \leftarrow Open \cup \{y\}$ ;
10      record transition  $x \xrightarrow{\sigma} y$ ;
```

Algorithm 1 works for deterministic FSMs. If the FSMs to be composed are nondeterministic, then all initial state combinations of $Q_1^\circ \times \dots \times Q_n^\circ$ must be generated and enqueued in lines 1–2, and likewise several successor state combinations y must be considered in lines 7–10.

Data structures. The memory requirements of Algorithm 1 are determined by the state set Q and the queue $Open$. They contain state tuples, which can be bit-packed to save memory. The states in Fig. 1 can be encoded using one or two bits for each FSM, e.g., the tuple 2100 can be stored as the binary string 1001000. This paper’s implementation packs each tuple into the smallest possible number of 32-bit words.

All state tuples in Q are stored consecutively in a growing array list, so they can be uniquely identified by their *index*. The first initial state tuple has index 0, the next tuple encountered has index 1, etc. A hash table [12] facilitates the membership test in line 8. The hash table contains only state indices, with the hash function computed using the actual state tuples from the list. The set $Open$ is better implemented as a queue. For example, a first-in-first-out queue ensures that states are processed in the order in which they are discovered, resulting in breadth-first search [12]. In fact, $Open$ is represented as a single integer, namely the index of the first state tuple that has not yet been expanded.

Thus, the amount of memory grows linearly with the number $|Q|$ of reachable states in the synchronous composition and with the number n of FSMs. The space complexity of Algorithm 1 is $O(n|Q|)$. In practice, each state typically requires 1–4 words for the bit-packed tuple plus two words in the hash table with 50% load, i.e., 12–24 bytes per state.

State expansion. The main factor influencing the runtime of Algorithm 1 is the loop in line 5 and the if-statement in line 6, which is executed for each state and event in the system. It is important that the test for $x_i \xrightarrow{\sigma}_i y_i$ is evaluated in constant time, which is achieved with an array structure set up in advance. The array contains, for each source state and event, either a null-value to indicate that the event is disabled, or the successor state, or a reference to a list of successor states in the case of nondeterminism.

```

1 case  $M_1$  goto 2,5,9      11 case  $M_2$  goto 12,15,19
2 if-disabled  $B, s_1$  goto 11 12 if-disabled  $B, s_2$  goto 21
3 execute  $s_1$               13 execute  $s_2$ 
4 goto 11                  14 goto 21
5 execute  $b_1$               15 execute  $f_2$ 
6 if-disabled  $B, f_1$  goto 11 16 if-disabled  $R, b_2$  goto 21
7 execute  $f_1$               17 execute  $b_2$ 
8 goto 11                  18 goto 21
9 if-disabled  $R, r_1$  goto 11 19 if-disabled  $R, r_2$  goto 21
10 execute  $r_1$             20 execute  $r_2$ 
                          21 end

```

Fig. 2. Branching program for Fig. 1.

Even more important is the observation that usually less than 10% of events tested for eligibility produce a transition. Performance is improved substantially by making the failing tests fail fast. This is achieved with pre-calculated lists of the FSMs that can disable a given event σ . For example, event r_1 in Fig. 1 only appears in M_1 and R , so only the states of these FSMs are tested to determine whether r_1 is enabled. Further, r_1 is enabled in only 1 out of 3 states of M_1 as opposed to 1 out of 2 states of R . The average number of state tests is reduced by first testing M_1 and then R , i.e., the FSMs should be sorted by event probability. The successor state computation in line 7 only starts after the event has been found to be enabled. In case of nondeterminism, a lot of repetition is avoided by first computing the state components of FSMs that do not synchronise or have only one successor state with the current event.

While the above approach with per-event FSM lists sorted by enablement probability usually works well, its performance degrades for models with hundreds or thousands of events. Further improvement is possible because, in many models, several events can be ruled out as disabled after testing the state of a single FSM. For example, all enablement decisions for Fig. 1 can be made using the *branching program* in Fig. 2. Here, line 1 branches to line 2, 5, or 9 depending on whether M_1 is in state 0, 1, or 2. If M_1 is in state 0, e.g., then M_1 disables f_1 , b_1 , and r_1 , and the only event from M_1 's event set that could be enabled is s_1 . Therefore line 2 checks whether s_1 is disabled by B , using array lookup, and if it is, the program continues in line 11 to check the remaining events s_2 , f_2 , b_2 , and r_2 . Otherwise s_1 is enabled, and line 3 performs the successor state computation as per lines 7–10 of Algorithm 1, before continuing in line 11. Overall, this program inspects the states of only four FSMs to make enablement decisions for eight events. This paper's implementation generates a branching program similar to Fig. 2 for its input FSMs and encodes it as primitive bytecode. Then the execution of the bytecode replaces the loop in line 5 of Algorithm 1.

Complexity. In the worst case, line 6 checks enablement of all events in all component FSMs, for each state in the synchronous composition, which gives up to $n|\Sigma||Q|$ operations. Lines 7–10 are executed once per transition of the synchronous composition, which may exceed $n|\Sigma||Q|$ in the nondeterministic case. Therefore, the worst-case time complexity of Algorithm 1 is $O(n|\Sigma||Q| + |\rightarrow|)$. The branching

Algorithm 2: Coreachability

Input: FSMs $G_i = \langle \Sigma_i, Q_i, Q_i^o, Q_i^\omega, \rightarrow_i \rangle$,
reachable states Q
Output: States marked as coreachable or not

```

1  $Open \leftarrow \emptyset$ ;
2 foreach  $x \in Q$  do
3   if  $x \in Q_1^\omega \times \dots \times Q_n^\omega$  then
4      $Open \leftarrow Open \cup \{x\}$ ;
5     mark  $x$  as coreachable;
6 while  $Open \neq \emptyset$  do
7   remove  $y$  from  $Open$ ;
8   foreach transition  $x \rightarrow y$  do
9     if  $x$  is not marked as coreachable then
10     $Open \leftarrow Open \cup \{x\}$ ;
11    mark  $x$  as coreachable;

```

program often significantly reduces the $n|\Sigma||Q|$ part while maintaining the same worst-case complexity.

Early termination. Algorithm 1 can terminate early in the case of deadlock. After the loop in line 5, if no successors of the expanded state x have been found, or all successors are equal to x , then it can be checked whether x is accepting. If it is not, then x is a reachable deadlock state and the system is conflicting. Only if Algorithm 1 reaches the end without terminating early, a second pass is needed to determine whether the system has a livelock or is nonconflicting.

B. Finding Coreachable States

Algorithm 2 continues after Algorithm 1 and performs a backwards search to find all coreachable states. This is a standard model checking algorithm to verify the CTL property **EF** *accepting* [2]. The loop in line 2 marks as coreachable all reachable states that are accepting, and the loop in line 6 adds to this all predecessors of states already marked as coreachable.

This approach uses one bit per state to record whether or not it is coreachable, which is stored conveniently as part of the bit-packed state tuples. Unlike Algorithm 1, the set $Open$ of unvisited states must be stored explicitly as a queue or stack of state numbers, and its size is bounded by the number of states. Thus, the memory requirements from Algorithm 1 increase by 33 bits or 4 bytes and 1 bit per state. This increase remains linear in the number of states, $O(|Q|)$.

Upon termination of Algorithm 2, exactly the coreachable states are marked, so a final loop can determine whether all reachable states in Q are coreachable. If so, the system is nonconflicting, otherwise it is conflicting.

Line 8 requires the ability to iterate over the predecessor states of a given state y in the synchronous composition. This requires additional data structures or computational effort. The following two alternatives have been implemented.

Stored backwards transitions. As Algorithm 1 explores the transitions, it can set up a data structure for the backwards search. The loop in line 8 of Algorithm 2 requires transitions indexed by target states. As the number of transitions per

target state is not known a-priori, they are stored in linked lists [12]. Each target state is associated with the first node of its list of incoming transitions, and each node contains the source state of a transition and a reference to the next node. Events are not stored because Algorithm 2 does not need them. Thus, each transition occupies two words or eight bytes. Each time Algorithm 1 executes line 10, the current state x is prepended to the predecessors list of the target state y . Some memory is saved by suppressing selfloops and duplicate transitions, which do not affect the result of the conflict check. Duplicates are recognised in constant time, because all transitions $x \rightarrow y$ are encountered while expanding x , when x , if it is already listed as predecessor of y , must be the first entry of the list.

The time to construct the predecessor lists is bounded by the number $|\rightarrow|$ of transitions in the synchronous composition, and so is the runtime of Algorithm 2 using these lists. The complexity of this conflict check is dominated by Algorithm 1, $O(n|\Sigma||Q| + |\rightarrow|)$. However, the predecessor lists require additional memory proportional to the number of transitions, so the space complexity increases to $O(n|Q| + |\rightarrow|)$. This is a serious problem for large systems, as the number of transitions typically is at least ten times the number of states.

Computed backwards transitions. It is also possible to find predecessor states by expanding the transition relation backwards. After reversing the direction of all transitions of the component FSMs, a branching program like Fig. 2 produces the predecessors of any given state tuple. Care needs be taken as the reverse transition relation may be nondeterministic even if all FSMs are deterministic. A more serious problem arises because backwards exploration may lead to unreachable states. Every predecessor state must be checked for reachability using the hash table from Algorithm 1, so only reachable states are enqueued and marked in lines 10–11 of Algorithm 2.

No transitions are stored with this approach, so the space complexity of the conflict check remains $O(n|Q|)$. The time complexity of Algorithm 2 with computed backwards transitions is $O(n|\Sigma||Q| + |\rightarrow|)$ for the same reasons as explained above for Algorithm 1, which also remains the worst-case time complexity of the conflict check. However, now $|\rightarrow|$ includes some unreachable transitions. This rarely is a problem in practice, except for a few models with interrupt or reset events where the extra transitions outweigh the reachable transitions by several orders of magnitude.

C. Strongly Connected Components

This section proposes an alternative conflict check algorithm based on strongly connected components. Given an FSM $G = \langle \Sigma, Q, Q^\circ, Q^\omega, \rightarrow \rangle$, a *strongly connected component* in G is a maximal set $C \subseteq Q$ of states such that, for all states $x, y \in C$ it holds that $x \rightarrow y$. A strongly connected component $C \subseteq Q$ is a *leaf component*, if for all $y \in Q$ such that $C \rightarrow y$ it holds that $y \in C$.

A strongly connected component is a set of states all reachable from each other, and a leaf component is a

Algorithm 3: Conflict Check with Tarjan’s Algorithm

```

1 procedure explore(state index  $i$ )
2    $lowlink[i] \leftarrow i$ ;  $stack.push(i)$ ;  $x \leftarrow tuple[i]$ ;
3   foreach transition  $x \rightarrow y$  do
4      $j \leftarrow$  index of  $y$  in  $Q$ ;
5     if  $j$  undefined then
6        $Q \leftarrow Q \cup \{y\}$ ;  $j \leftarrow$  index of  $y$  in  $Q$ ;
7       explore( $j$ );
8        $lowlink[i] \leftarrow \min(lowlink[i], lowlink[j])$ ;
9     else if  $stack$  contains  $j$  then
10       $lowlink[i] \leftarrow \min(lowlink[i], j)$ ;
11  if  $lowlink[i] = i$  then
12     $comp \leftarrow \emptyset$ ;
13    repeat
14       $j \leftarrow stack.pop()$ ;  $comp \leftarrow comp \cup \{j\}$ ;
15    until  $i = j$ ;
16    foreach  $j \in comp$  do
17      if  $tuple[j] \in Q_1^\omega \times \dots \times Q_n^\omega$  then
18        return
19      else if  $\exists k : tuple[j] \rightarrow tuple[k] \wedge k \notin comp$  then
20        return
21  stop “The system is conflicting”;
```

strongly connected component from which only states in that component can be reached. The FSM $M_1 \parallel B \parallel R \parallel M_2$ in Fig. 1 has seven strongly connected components, e.g., $\{0000, 0001, 0012, 0200, 0201, 0212, 1100, 1101, 1112\}$ and $\{2100\}$. Its only leaf component is $\{0100\}$.

It is known from graph theory that, for any state there exists a leaf component reachable from that state. Therefore, an FSM is blocking if and only if it has a blocking leaf component, i.e., a leaf component without any accepting states. This observation leads to the following result.

Proposition 1: An FSM $G = \langle \Sigma, Q, Q^\circ, Q^\omega, \rightarrow \rangle$ is non-blocking if and only if, for all leaf components $C \subseteq Q$ it holds that $C \cap Q^\omega \neq \emptyset$.

Tarjan’s algorithm [13] is a popular method to find strongly connected components. It performs a single search over all transitions and outputs strongly connected components as it runs. Using it, Algorithm 3 performs a conflict check by testing the conditions of Prop. 1 as new components are detected.

Algorithm 3 is defined by the recursive procedure `explore()`, which must be called exactly once for each initial state. All state tuples are stored in a growing list Q , in the order in which they are detected, as explained in Section III-A above. The argument i to `explore()` is the index of a state in this list, and `tuple[i]` is the state tuple retrieved from the list. In addition to this, Tarjan’s algorithm associates with each state a so-called *lowlink*, which records (roughly) the index of the smallest state known to be in the same component. It also uses a *stack* containing the indices of all states not yet assigned to any component.

Lines 2–10 in Algorithm 3 perform depth-first search as

prescribed by Tarjan’s algorithm. Each state is pushed on the *stack*, before the loop in line 3 expands it as described in Section III-A. New states are added to the list Q and explored recursively, updating the *lowlink*. Tarjan’s algorithm ensures that, if line 11 is reached and the *lowlink* of state number i is still i , then this is the *root state* of a strongly connected component, which includes all states after it on the *stack* [13]. The loop in line 13 adds these states to a temporary set *comp*, and the loop in line 16 checks whether it is a blocking leaf component. If the component is not blocking because it contains an accepting state (line 17), or not a leaf because it has a transition to another component (line 19), the call to *explore()* returns to the previous level of recursion. Otherwise, a blocking leaf component has been found and line 21 terminates the algorithm early.

The appeal of this algorithm is that it explores transitions only in the forward direction and allows early termination even in case of livelock. Each state is expanded once in line 3 and possibly a second time in line 19. The second expansion is often skipped in practice by stopping the loop in line 16 early. Furthermore, Tarjan’s algorithm guarantees that the first component detected is a leaf, so the second expansion can be avoided entirely in the common case of systems with only one component. The worst-case time complexity remains $O(n|\Sigma||Q| + |\rightarrow|)$.

The recursive control structure of Tarjan’s algorithm is problematic for large systems. The depth of recursion is only bounded by the number of states, $|Q|$, and the location of the recursive call in the middle of state expansion requires a lot of context information to be stored for recursive calls, considerably increasing memory consumption or causing stack overflow. Standard iterative algorithms [12] to compute strongly connected components require both forward and backward transitions. This paper is based on an iterative implementation of Tarjan’s algorithm based on [14], which uses a second stack for pending recursive calls. Stack size is bounded by the number of states, and the worst-case memory overhead is 20 bytes per state. The algorithm retains both the linear space complexity in the number of states and the linear time complexity in the number of transitions.

D. Counterexamples

Counterexamples are of great value to users of model checkers, because they explain a detected problem and facilitate the finding of a fix. A *counterexample* to the nonblocking property of an FSM $G = \langle \Sigma, Q, Q^\circ, Q^\omega, \rightarrow \rangle$ is a trace $s \in \Sigma^*$, together with state information in case of nondeterminism, which takes the system to a non-coreachable state, i.e., $Q^\circ \xrightarrow{s} x \not\rightarrow Q^\omega$. As the counterexample is presented to users, it should be as short as possible, and counterexamples that end in a deadlock state or a blocking leaf component, are more specific and thus more helpful.

Counterexample computation is straightforward if breadth-first search is used in Algorithm 1. The end state of the counterexample is either the deadlock state that triggered early termination, or after Algorithm 2, the non-coreachable state with the smallest index in breadth-first order. From

the end state, the trace is constructed backwards until an initial state is reached, either using stored transitions or by backwards expansion. To get a shortest trace, the first predecessor in breadth-first order must be used at each step, which is either the first state added to the list, or the predecessor with the smallest index. When using transition lists without stored events, the event is obtained by expanding the predecessor state a second time forwards. This method expands one state for each step of the trace, whose length can be the total number of states in the worst case. It usually is shorter in practice, so that the overhead for counterexample computation is insignificant. The counterexample is guaranteed to be a shortest trace to a deadlock state, if one exists, and otherwise a shortest trace to a livelock state.

The same method can be used to compute a counterexample after Algorithm 3, but state numbers in depth-first order do not guarantee a shortest trace and the counterexamples are often unusably long. To improve counterexample quality, this paper’s implementation performs a breadth-first search to find a shortest trace to the first blocking leaf component encountered, visiting only states discovered in the depth-first search, and then constructs the counterexample as explained above. This results in overhead of expanding all discovered states up to two more times. While the counterexample is not necessarily the shortest possible, it usually is a fair compromise, and is guaranteed to end in a blocking leaf component.

IV. EXPERIMENTAL RESULTS

The algorithms described above are implemented as part of the *Waikato Analysis Toolkit for Events in Reactive Systems (Waters)*. The Waters library is programmed in C++ and released under the GNU Public License. It is available as part of Supremica [7]. This implementation has been used to check the nonconflicting property of 15 discrete event systems models from industrial applications and case studies.

Table I shows the results of the experiments. It shows for each model, the number of events ($|\Sigma|$), the number of FSMs (n), the number of bits to encode the state tuples (Enc), the number of reachable states in the synchronous composition (State space), and the verification result. Then it shows the runtime (Time), the memory consumption (Mem), and if applicable the length of the counterexample (CE), for four conflict check algorithms.

Stored is the combination of Algorithms 1 and 2 with stored backwards transitions, while **Computed** expands states backwards. **Tarjan** is Algorithm 3. To put the data in perspective, the **BDD** column shows the results of an algorithm based on binary decision diagrams (BDD) [2]. This is a symbolic version of Algorithms 1 and 2, tuned for better performance on discrete event systems models. It uses a variable ordering based on the FORCE heuristics [15] and a disjunctive partitioning and search strategy [16]. It terminates early in case of deadlock but not livelock, and does not ensure shortest counterexamples. BDD-based breadth-first search gives shortest counterexamples but runs up to ten times slower.

TABLE I
EXPERIMENTAL RESULTS

Name	Model					Stored			Computed			Tarjan			BDD		
	$ \Sigma $	n	Enc	State space	Verification result	Time [s]	Mem [MB]	CE	Time [s]	Mem [MB]	CE	Time [s]	Mem [MB]	CE	Time [s]	Mem [MB]	CE
agv	53	16	40	$2.57 \cdot 10^7$	nonconflicting	33.2	1493.5		55.8	590.9		38.2	842.4		0.1	32.0	
agvb	53	17	42	$2.29 \cdot 10^7$	deadlock	4.7	394.1	56	3.9	208.4	56	0.0	80.9	92	0.1	31.6	56
aip0tough_abs1	44	6	64	$1.02 \cdot 10^8$	livelock	84.4	3763.4	143	69.2	1364.2	143	2.6	142.4	153	37.1	63.8	161
aip1efa_2	94	50	112	$3.55 \cdot 10^7$	nonconflicting	89.7	2919.8		259.2	1372.4		75.6	1708.5		23.6	17.3	
aip1efa16_abs1	32	5	53	$3.13 \cdot 10^{10}$	deadlock	39.4	1748.7	91	32.2	665.1	91	2.0	166.4	178	34.2	93.0	101
big_bmw	66	31	57	$3.14 \cdot 10^7$	nonconflicting	105.4	3012.3		692.8	608.0		92.7	854.1		0.1	8.2	
dynamic_prime_sieve_5	4535	25	168	$6.79 \cdot 10^7$	nonconflicting	451.0	5722.0		1000.8	3653.1		457.9	3873.0		53.4	327.2	
fencaiwon09	73	32	79	$1.03 \cdot 10^8$	nonconflicting	212.3	6861.9		485.2	2395.3		168.0	3146.8		1.2	10.0	
fencaiwon09b	73	31	75	$8.93 \cdot 10^7$	deadlock	134.5	6134.7	267	108.1	2388.2	267	0.0	82.5	272	0.4	9.4	267
fencaiwon09s	73	29	67	$3.00 \cdot 10^8$	deadlock	0.5	119.0	41	0.4	99.9	41	0.0	84.5	41	0.1	32.0	41
ftechnik	117	36	120	$1.21 \cdot 10^8$	livelock	Out of memory			448.9	2985.8	0	0.0	115.9	20	0.1	24.4	0
profisafe_i4_abs1	50	4	45	$4.92 \cdot 10^7$	nonconflicting	125.4	3796.7		186.5	1373.6		116.4	1862.0		324.9	197.6	
profisafe_ihost_efa_2_24	498	21	37	$6.69 \cdot 10^7$	nonconflicting	Out of memory			2818.0	1288.0		1820.8	1821.7		4.8	78.2	
tbed_reset1	194	98	254	664128	nonconflicting	0.9	154.6		3.4	141.6		0.9	157.2		613.2	17.8	
verriegel2	88	41	70	$2.18 \cdot 10^7$	nonconflicting	71.9	2146.6		Timeout			48.9	797.0		1.0	12.0	

All experiments were run on a standard PC with a 3.3 GHz microprocessor and 8 GB of RAM. The table shows that this is enough to explore state spaces with 100 million states completely within minutes. **Computed** is the most memory-efficient of the explicit algorithms, but it can be slow and was aborted after failing to solve the **verriegel2** model in an hour. This model has prohibitively many backwards transitions from unreachable states. **Stored** is faster except in case of deadlock, but requires the most memory. The **ftechnik** and **profisafe_i4_abs1** models have too many transitions for this algorithm. **Tarjan** usually runs faster than **Stored** and **Computed**, particularly when it terminates early due to livelock, but uses more memory than **Computed** and sometimes produces longer counterexamples.

The focus of this paper is to compare the explicit algorithms and not necessarily to beat symbolic methods. The **BDD** algorithm usually uses far less memory and runs faster than the explicit algorithms, but there are exceptions. The **tbed_reset1** model has many FSMs and requires many bits to encode, and **aip1efa16_abs1** and **profisafe_i4_abs1** have only few FSMs with thousands of states each. Such models can be difficult to encode as BDDs, causing the **BDD** algorithm to struggle, while their relatively small state numbers in combination with early termination make them amenable for explicit methods. The models **aip1efa16_abs1** and **profisafe_i4_abs1** are results of compositional minimisation of much larger models [6], which is a potential application of explicit algorithms.

V. CONCLUSIONS

Three explicit algorithms to verify the nonconflicting property of discrete event systems have been described and compared. These algorithms maintain linear time and space complexity in the size of the synchronous composition to be explored. They are optimised differently for runtime or memory usage and implemented in C++. Experimental results show that the implementation exhaustively explores state spaces of more than 100 million states on standard computers within minutes.

REFERENCES

- [1] P. J. G. Ramadge and W. M. Wonham, "The control of discrete event systems," *Proc. IEEE*, vol. 77, no. 1, pp. 81–98, Jan. 1989.
- [2] B. Bérard, M. Bidoit, A. Finkel, F. Laroussinie, A. Petit, L. Petrucci, and P. Schnoebelen, *Systems and Software Verification*. Springer, 2001.
- [3] P. N. Pena, J. E. R. Cury, and S. Lafortune, "Verification of nonconflict of supervisors using abstractions," *IEEE Trans. Autom. Control*, vol. 54, no. 12, pp. 2803–2815, 2009.
- [4] H. Flordal and R. Malik, "Compositional verification in supervisory control," *SIAM J. Control and Optimization*, vol. 48, no. 3, pp. 1914–1938, 2009.
- [5] R. Su, J. H. van Schuppen, J. E. Rooda, and A. T. Hofkamp, "Nonconflict check by using sequential automaton abstractions based on weak observation equivalence," *Automatica*, vol. 46, no. 6, pp. 968–978, June 2010.
- [6] C. Pilbrow and R. Malik, "An algorithm for compositional nonblocking verification using special events," *Sci. Comput. Programming*, vol. 113, no. 2, pp. 119–148, Dec. 2015.
- [7] K. Åkesson, M. Fabian, H. Flordal, and R. Malik, "Supremica—an integrated environment for verification, synthesis and simulation of discrete event systems," in *Proc. 8th Int. Workshop on Discrete Event Systems, WODES '06*. Ann Arbor, MI, USA: IEEE, July 2006, pp. 384–385.
- [8] L. Feng and W. M. Wonham, "TCT: A computation tool for supervisory control synthesis," in *Proc. 8th Int. Workshop on Discrete Event Systems, WODES '06*. Ann Arbor, MI, USA: IEEE, July 2006, pp. 388–389.
- [9] T. Moor, K. Schmidt, and S. Perk, "libFaudeS — an open source C++ library for discrete event systems," in *Proc. 9th Int. Workshop on Discrete Event Systems, WODES '08*. Göteborg, Sweden: IEEE, May 2008, pp. 125–130.
- [10] G. J. Holzmann, "The SPIN model checker," *IEEE Trans. Softw. Eng.*, vol. 23, pp. 279–295, 1997.
- [11] C. A. R. Hoare, *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [12] T. H. Cormen, C. E. Leiserson, R. E. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. MIT Press, 2009.
- [13] R. Tarjan, "Depth first search and linear graph algorithms," *SIAM J. Computing*, vol. 1, no. 2, pp. 146–160, June 1972.
- [14] A. M. Shaw, "Partial order reduction with compositional verification," Master's thesis, Dept. of Computer Science, University of Waikato, 2014. [Online]. Available: <http://hdl.handle.net/10289/9001>
- [15] F. A. Aloul, I. L. Markov, and K. A. Sakallah, "FORCE: A fast & easy-to-implement variable-ordering heuristic," in *Proc. 13th ACM Great Lakes Symp. VLSI*, Washington, DC, USA, 2003, pp. 116–119.
- [16] R. Song, "Symbolic synthesis and verification of hierarchical interface-based supervisory control," Master's thesis, Dept. of Computing and Software, McMaster University, Hamilton, ON, Canada, 2006. [Online]. Available: <http://www.cas.mcmaster.ca/~leduc>