

# Saturation: An Efficient Iteration Strategy for Symbolic State–Space Generation\*

Gianfranco Ciardo<sup>1</sup>, Gerald Lüttgen<sup>2</sup>, and Radu Siminiceanu<sup>1</sup>

<sup>1</sup> Department of Computer Science, College of William and Mary  
Williamsburg, VA 23187, USA {[ciardo](mailto:ciardo@cs.wm.edu), [radu](mailto:radu@cs.wm.edu)}@cs.wm.edu

<sup>2</sup> Department of Computer Science, Sheffield University, 211 Portobello Street  
Sheffield S1 4DP, U.K. [g.luetngen@dcs.shef.ac.uk](mailto:g.luetngen@dcs.shef.ac.uk)

**Abstract.** We present a novel algorithm for generating state spaces of asynchronous systems using Multi-valued Decision Diagrams. In contrast to related work, we encode the next-state function of a system not as a single Boolean function, but as cross-products of integer functions. This permits the application of various iteration strategies to build a system's state space. In particular, we introduce a new elegant strategy, called *saturation*, and implement it in the tool SMART. On top of usually performing several orders of magnitude faster than existing BDD-based state-space generators, our algorithm's required peak memory is often close to the final memory needed for storing the overall state space.

## 1 Introduction

*State-space generation* is one of the most fundamental challenges for many formal verification tools, such as model checkers [13]. The high complexity of today's digital systems requires constructing and storing huge state spaces in the relatively small memory of a workstation. One research direction widely pursued in the literature suggests the use of *decision diagrams*, usually Binary Decision Diagrams [7] (BDDs), as a data structure for implicitly representing large sets of states in a compact fashion. This proved to be very successful for the verification of synchronous digital circuits, as it increased the manageable sizes of state spaces from about  $10^6$  states, with traditional explicit state-space generation techniques [14], to about  $10^{20}$  states [9]. Unfortunately, symbolic techniques are known not to work well for *asynchronous systems*, such as communication protocols, which particularly suffer from state-space explosion.

The latter problem was addressed in previous work by the authors in the context of state-space generation using *Multi-valued Decision Diagrams* [18] (MDDs), which exploited the fact that, in event-based asynchronous systems,

---

\* This work was partially supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the authors were in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681, USA. G. Ciardo and R. Siminiceanu were also partially supported by NASA grant No. NAG-1-2168.

each event updates just a few components of a system's state vector [10]. Hence, firing an event only requires the application of *local next-state functions* and the local manipulation of MDDs. This is in contrast to classic BDD-based techniques which construct state spaces by iteratively applying a single, global next-state function which is itself encoded as a BDD [20]. Additionally, in most concurrency frameworks including Petri nets [23] and process algebras [5], next-state functions satisfy a *product form* allowing each component of the state vector to be updated somewhat independently of the others. Experimental results implementing these ideas of locality showed significant improvements in speed and memory consumption when compared to other state-space generators [22].

In this paper, we take our previous approach a significant step further by observing that the reachable state space of a system can be built by firing the system's events in any order, as long as every event is considered often enough [16]. We exploit this freedom by proposing a novel strategy which exhaustively fires all events affecting a given MDD node, thereby bringing it to its final *saturated* shape. Moreover, nodes are considered in a depth-first fashion, i.e., when a node is processed, all its descendants are already saturated. The resulting state-space generation algorithm is not only concise, but also allows for an elegant proof of correctness. Compared to our previous work [10], saturation eliminates a fair amount of administration overhead, reduces the average number of firing events, and enables a simpler and more efficient cache management.

We implemented the new algorithm in the tool SMART [11], and experimental studies indicate that it performs on average about one order of magnitude faster than our old algorithm. Even more important and in contrast to related work, the peak memory requirements of our algorithm are often close to its final memory requirements. In the case of the dining philosophers' problem, we are able to construct the state space of about  $10^{627}$  states, for 1000 philosophers, in under one second on a 800 MHz Pentium III PC using only 390KB of memory.

## 2 MDDs for Encoding Structured State Spaces

**State spaces and next-state functions.** A discrete-state model expressed in a high-level formalism must specify: (i)  $\widehat{\mathcal{S}}$ , the set of *potential states* describing the "type" of states; (ii)  $\mathbf{s} \in \widehat{\mathcal{S}}$ , the *initial state*; and (iii)  $\mathcal{N} : \widehat{\mathcal{S}} \rightarrow 2^{\widehat{\mathcal{S}}}$ , the *next-state function*, describing which states can be reached from a given state in a single step. In many cases, such as Petri nets and process algebras, a model expresses this function as a union  $\mathcal{N} = \bigcup_{e \in \mathcal{E}} \mathcal{N}_e$ , where  $\mathcal{E}$  is a finite set of *events* and  $\mathcal{N}_e$  is the next-state function associated with event  $e$ . We say that  $\mathcal{N}_e(s)$  is the set of states the system can enter when event  $e$  occurs, or *fires*, in state  $s$ . Moreover, event  $e$  is called *disabled* in  $s$  if  $\mathcal{N}_e(s) = \emptyset$ ; otherwise, it is *enabled*.

The *reachable state space*  $\mathcal{S} \subseteq \widehat{\mathcal{S}}$  of the model under consideration is the smallest set containing the initial system state  $\mathbf{s}$  and closed with respect to  $\mathcal{N}$ , i.e.,  $\mathcal{S} = \{\mathbf{s}\} \cup \mathcal{N}(\mathbf{s}) \cup \mathcal{N}(\mathcal{N}(\mathbf{s})) \cup \dots = \mathcal{N}^*(\mathbf{s})$ , where "\*" denotes the reflexive and transitive closure. When  $\mathcal{N}$  is composed of several functions  $\mathcal{N}_e$ , for  $e \in \mathcal{E}$ , we can iterate these functions in any order, as long as we consider each  $\mathcal{N}_e$  often

enough. In other words,  $i \in \mathcal{S}$  if and only if it can be reached from  $\mathbf{s}$  through zero or more event firings. In this paper we assume that  $\mathcal{S}$  is finite; however, for most practical asynchronous systems, the size of  $\mathcal{S}$  is enormous due to the *state-space explosion* problem.

**Multi-valued decision diagrams.** One way to cope with this problem is to use efficient data structures to encode  $\mathcal{S}$  that exploit the system's *structure*. We consider a common case in asynchronous system design, where a system model is composed of  $K$  *submodels*, for some  $K \in \mathbb{N}$ , so that a global system state is a  $K$ -tuple  $(i^K, \dots, i^1)$ , where  $i^k$  is the local state for submodel  $k$ . (We use *superscripts* for submodel indices —not for exponentiation— and *subscripts* for event indices.) Thus,  $\widehat{\mathcal{S}} = \mathcal{S}^K \times \dots \times \mathcal{S}^1$ , with each *local* state space  $\mathcal{S}^k$  having some finite size  $n^k$ . In Petri nets, for example, the set of places can be partitioned into  $K$  subsets, and the marking can be written as the composition of the  $K$  corresponding submarkings. When identifying  $\mathcal{S}^k$  with the initial integer interval  $\{0, \dots, n^k - 1\}$ , for each  $K \geq k \geq 1$ , one can encode  $\mathcal{S} \subseteq \widehat{\mathcal{S}}$  via a (*quasi-reduced ordered*) MDD, i.e., a directed acyclic edge-labelled multi-graph where:

- Nodes are organized into  $K + 1$  *levels*. We write  $\langle k.p \rangle$  to denote a generic node, where  $k$  is the level and  $p$  is a unique index for that level. Level  $K$  contains only a single *non-terminal* node  $\langle K.r \rangle$ , the *root*, whereas levels  $K-1$  through 1 contain one or more non-terminal nodes. Level 0 consists of two *terminal* nodes,  $\langle 0.0 \rangle$  and  $\langle 0.1 \rangle$ . (We use boldface for the node indices 0 or 1 since these have a special meaning, as we will explain later.)
- A non-terminal node  $\langle k.p \rangle$  has  $n^k$  arcs pointing to nodes at level  $k-1$ . If the  $i^{\text{th}}$  arc, for  $i \in \mathcal{S}^k$ , is to node  $\langle k-1.q \rangle$ , we write  $\langle k.p \rangle[i] = q$ . Unlike in the original BDD setting [7, 8], we allow for *redundant* nodes, having all arcs pointing to the same node. This will be convenient for our purposes, as eliminating such nodes would lead to arcs spanning multiple levels.
- A non-terminal node cannot *duplicate* (i.e., have the same pattern of arcs as) another node at the same level.

Given a node  $\langle k.p \rangle$ , we can recursively define the node reached from it through any integer sequence  $\gamma =_{\text{df}} (i^k, i^{k-1}, \dots, i^l) \in \mathcal{S}^k \times \mathcal{S}^{k-1} \times \dots \times \mathcal{S}^l$  of length  $k-l+1$ , for  $K \geq k \geq l \geq 1$ , as

$$\text{node}(\langle k.p \rangle, \gamma) = \begin{cases} \langle k.p \rangle & \text{if } \gamma = (), \text{ the empty sequence} \\ \text{node}(\langle k-1.q \rangle, \delta) & \text{if } \gamma = (i^k, \delta) \text{ and } \langle k.p \rangle[i^k] = q. \end{cases}$$

The substates encoded by  $p$  or reaching  $p$  are then, respectively,

$$\begin{aligned} \mathcal{B}(\langle k.p \rangle) &= \{\beta \in \mathcal{S}^k \times \dots \times \mathcal{S}^1 : \text{node}(\langle k.p \rangle, \beta) = \langle 0.1 \rangle\} && \text{“below” } \langle k.p \rangle; \\ \mathcal{A}(\langle k.p \rangle) &= \{\alpha \in \mathcal{S}^K \times \dots \times \mathcal{S}^{k+1} : \text{node}(\langle K.r \rangle, \alpha) = \langle k.p \rangle\} && \text{“above” } \langle k.p \rangle. \end{aligned}$$

Thus,  $\mathcal{B}(\langle k.p \rangle)$  contains the substates that, prefixed by a substate in  $\mathcal{A}(\langle k.p \rangle)$ , form a (global) state encoded by the MDD. We reserve the indices 0 and 1 at each level  $k$  to encode the sets  $\emptyset$  and  $\mathcal{S}^k \times \dots \times \mathcal{S}^1$ , respectively. In particular,  $\mathcal{B}(\langle 0.0 \rangle) = \emptyset$  and  $\mathcal{B}(\langle 0.1 \rangle) = \{()\}$ . Fig. 1 shows a four-level example MDD and the set  $\mathcal{S}$  encoded by it; only the highlighted nodes are actually stored.

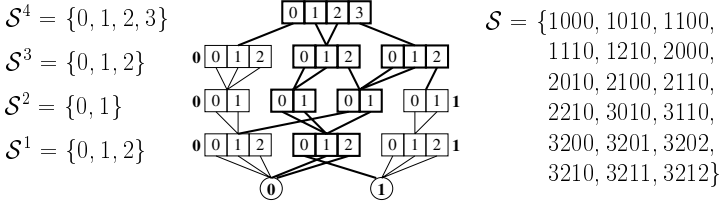


Fig. 1. An example MDD and the state space  $\mathcal{S}$  encoded by it.

Many algorithms for generating state spaces using BDDs exist [20], which can easily be adapted to MDDs. In contrast to those, however, our approach does *not* encode the next-state function as an MDD over  $2K$  variables, describing the  $K$  state components before and after a system step. Instead, we update MDD nodes directly, adding the new states reached through one step of the global next-state function when firing some event. For asynchronous systems, this function is often expressible as the cross-product of local next-state functions.

**Product-form behavior.** An asynchronous system model exhibits such behavior if, for each event  $e$ , its next-state function  $\mathcal{N}_e$  can be written as a cross-product of  $K$  local functions, i.e.,  $\mathcal{N}_e = \mathcal{N}_e^K \times \dots \times \mathcal{N}_e^1$  where  $\mathcal{N}_e^k : \mathcal{S}^k \rightarrow 2^{\mathcal{S}^k}$ , for  $K \geq k \geq 1$ . (Recall that event  $e$  is disabled in some global state exactly if it is disabled in at least one component.) The product-form requirement is quite natural. First, many modeling formalisms satisfy it, e.g., any Petri net model conforms to this behavior for any partition of its places. Second, if a given model does not respect the product-form behavior, we can always coarsen  $K$  or refine  $\mathcal{E}$  so that it does. As an example, consider a model partitioned into four submodels, where  $\mathcal{N}_e = \mathcal{N}_e^4 \times \mathcal{N}_e^{3,2} \times \mathcal{N}_e^1$ , but  $\mathcal{N}_e^{3,2} : \mathcal{S}^3 \times \mathcal{S}^2 \rightarrow 2^{\mathcal{S}^3 \times \mathcal{S}^2}$  cannot be expressed as a product  $\mathcal{N}_e^3 \times \mathcal{N}_e^2$ . We can achieve the product-form requirement by simply partitioning the model into three, not four, submodels. Alternatively, we may substitute event  $e$  with “subevents” satisfying the product form. This is possible since, in the worst case, we can define a subevent  $e_{i,j}$ , for each  $i = (i^3, i^2)$  and  $j = (j^3, j^2) \in \mathcal{N}_e^{3,2}(i)$ , with  $\mathcal{N}_{e_{i,j}}(i^3) = \{j^3\}$  and  $\mathcal{N}_{e_{i,j}}(i^2) = \{j^2\}$ .

Finally, we introduce some notational conventions. We say that event  $e$  *depends* on level  $k$ , if the local state at level  $k$  does affect the enabling of  $e$  or if it is changed by the firing of  $e$ . Let  $First(e)$  and  $Last(e)$  be the first and last levels on which event  $e$  depends. Events  $e$  such that  $First(e) = Last(e) = k$  are said to be *local events*; we merge these into a single *macro-event*  $\lambda^k$  without violating the product-form requirement, since we can write  $\mathcal{N}_{\lambda^k} = \mathcal{N}_{\lambda^k}^K \times \dots \times \mathcal{N}_{\lambda^k}^1$  where  $\mathcal{N}_{\lambda^k}^k = \bigcup_{\{e: First(e)=Last(e)=k\}} \mathcal{N}_e^k$ , while  $\mathcal{N}_{\lambda^k}^l(i^l) = \{i^l\}$  for  $l \neq k$  and  $i^l \in \mathcal{S}^l$ . The set  $\{e \in \mathcal{E} : First(e) = k\}$  of events “starting” at level  $k$  is denoted by  $\mathcal{E}^k$ . We also extend  $\mathcal{N}_e$  to substates instead of full states:  $\mathcal{N}_e((i^k, \dots, i^l)) = \mathcal{N}_e^k(i^k) \times \dots \times \mathcal{N}_e^l(i^l)$ , for  $K \geq k \geq l \geq 1$ ; to sets of states:  $\mathcal{N}_e(\mathcal{X}) = \bigcup_{i \in \mathcal{X}} \mathcal{N}_e(i)$ , for  $\mathcal{X} \subseteq \mathcal{S}^k \times \dots \times \mathcal{S}^l$ ; and to sets of events:  $\mathcal{N}_{\mathcal{F}}(\mathcal{X}) = \bigcup_{e \in \mathcal{F}} \mathcal{N}_e(\mathcal{X})$ , for  $\mathcal{F} \subseteq \mathcal{E}$ . In particular, we write  $\mathcal{N}_{\leq k}$  as a shorthand for  $\mathcal{N}_{\{e: First(e) \leq k\}}$ .

### 3 A Novel Algorithm Employing Node Saturation

In the following we refer to a specific order of iterating the local next-state functions of an synchronous system model as *iteration strategy*. Clearly, the choice of strategy influences the efficiency of state-space generation. In our previous work [10] we employed a naive strategy that cycled through MDDs level-by-level and fired, at each level  $k$ , all events  $e$  with  $First(e) = k$ .

As main contribution of this paper, we present a novel iteration strategy, called *saturation*, which not only simplifies our previous algorithm, but also significantly improves its time *and* space efficiency. The key idea is to fire events node-wise and exhaustively, instead of level-wise and just once per iteration. Formally, we say that an MDD node  $\langle k.p \rangle$  is *saturated* if it encodes a set of states that is a fixed point with respect to the firing of any event at its level or at a lower level, i.e., if  $\mathcal{B}(\langle k.p \rangle) = \mathcal{N}_{\leq k}^*(\mathcal{B}(\langle k.p \rangle))$  holds; it can easily be shown by contradiction that any node below node  $\langle k.p \rangle$  must be saturated, too. It should be noted that the routine for firing some event, in order to reveal and add globally reachable states to the MDD representation of the state space under construction, is similar to [10]. In particular, MDDs are manipulated only locally with respect to the levels on which the fired event depends, and, due to the product-form behavior, these manipulations can be carried out very efficiently. We do not further comment on these issues here, but concentrate solely on the new idea of node saturation and its implications.

Just as in traditional symbolic state-space generation algorithms, we use a *unique table*, to detect duplicate nodes, and *operation caches*, in particular a *union cache* and a *firing cache*, to speed-up computation. However, our approach is distinguished by the fact that only saturated nodes are checked in the unique table or referenced in the caches. Given the MDD encoding of the initial state  $s$ , we saturate its nodes bottom-up. This improves both memory and execution-time efficiency for generating state spaces because of the following reasons. First, our saturation order ensures that the firing of an event affecting only the current and possibly lower levels adds as many new states as possible. Then, since each node in the final encoding of  $\mathcal{S}$  is saturated, any node we insert in the unique table has at least a chance of being part of the final MDD, while any unsaturated node inserted by a traditional symbolic approach is *guaranteed* to be eventually deleted and replaced with another node encoding a larger subset of states. Finally, once we saturate a node at level  $k$ , we never need to fire any event  $e \in \mathcal{E}^k$  in it again, while, in classic symbolic approaches,  $\mathcal{N}$  is applied to the entire MDD at *every iteration*.

In the pseudo-code of our new algorithm implementing node saturation, which is shown in Fig. 2, we use the data types *evt* (model event), *lcl* (local state), *lvl* (level), and *idx* (node index within a level); in practice these are simply integers in appropriate ranges. We also assume the following dynamically-sized global hash tables: (a)  $UT[k]$ , for  $K \geq k \geq 1$ , the *unique table* for nodes at level  $k$ , to retrieve  $p$  given the key  $\langle k.p \rangle[0], \dots, \langle k.p \rangle[n^k - 1]$ ; (b)  $UC[k]$ , for  $K > k \geq 1$ , the *union cache* for nodes at level  $k$ , to retrieve  $s$  given nodes  $p$  and  $q$ , where  $\mathcal{B}(\langle k.s \rangle) = \mathcal{B}(\langle k.p \rangle) \cup \mathcal{B}(\langle k.q \rangle)$ ; and (c)  $FC[k]$ , for  $K > k \geq 1$ , the *firing cache*

<p><i>Generate</i>(in <math>s:\text{array}[1..K]</math> of <math>lcl</math>):<math>idx</math></p> <p>Build an MDD rooted at <math>\langle k.r \rangle</math> encoding <math>\mathcal{N}_e^*(s)</math> and return <math>r</math>, in <math>UT[K]</math>.</p> <pre> declare <math>r.p:idx</math>; declare <math>k:lvl</math>; 1. <math>p \leftarrow 1</math>; 2. for <math>k = 1</math> to <math>K</math> do 3. <math>r \leftarrow \text{NewNode}(k)</math>; <math>\langle k.r \rangle[s[k]] \leftarrow p</math>; 4. <i>Saturate</i>(<math>k, r</math>); <i>Check</i>(<math>k, r</math>); 5. <math>p \leftarrow r</math>; return <math>r</math>;</pre>	<p><i>RecFire</i>(in <math>e:evnt, l:lvl, q:idx</math>):<math>idx</math></p> <p>Build an MDD rooted at <math>\langle l.s \rangle</math>, in <math>UT[l]</math>, encoding <math>\mathcal{N}_{\leq l}^*(\mathcal{N}_e(\mathcal{B}(\langle l.q \rangle)))</math>. Return <math>s</math>.</p> <pre> declare <math>\mathcal{L}</math>:set of <math>lcl</math>; declare <math>f, u, s:idx</math>; declare <math>i, j:lcl</math>; declare <math>sCng:bool</math>; 1. if <math>l &lt; \text{Last}(e)</math> then return <math>q</math>; 2. if <i>Find</i>(<math>FC[l], \{q, e\}, s</math>) then return <math>s</math>; 3. <math>s \leftarrow \text{NewNode}(l)</math>; <math>sCng \leftarrow false</math>; 4. <math>\mathcal{L} \leftarrow \text{Locals}(e, l, q)</math>; 5. while <math>\mathcal{L} \neq \emptyset</math> do 6. <math>i \leftarrow \text{Pick}(\mathcal{L})</math>; 7. <math>f \leftarrow \text{RecFire}(e, l-1, \langle l.q \rangle[i])</math>; 8. if <math>f \neq \mathbf{0}</math> then 9. foreach <math>j \in \mathcal{N}_e^l(i)</math> do 10. <math>u \leftarrow \text{Union}(l-1, f, \langle l.s \rangle[j])</math>; 11. if <math>u \neq \langle l.s \rangle[j]</math> then 12. <math>\langle l.s \rangle[j] \leftarrow u</math>; <math>sCng \leftarrow true</math>; 13. if <math>sCng</math> then <i>Saturate</i>(<math>l, s</math>); 14. <i>Check</i>(<math>l, s</math>); <i>Insert</i>(<math>FC[l], \{q, e\}, s</math>); 15. return <math>s</math>;</pre>
<p><i>Saturate</i>(in <math>k:lvl, p:idx</math>)</p> <p>Update <math>\langle k.p \rangle</math>, not in <math>UT[k]</math>, in-place, to encode <math>\mathcal{N}_{\leq k}^*(\mathcal{B}(\langle k.p \rangle))</math>.</p> <pre> declare <math>e:evnt</math>; declare <math>\mathcal{L}</math>:set of <math>lcl</math>; declare <math>f, u:idx</math>; declare <math>i, j:lcl</math>; declare <math>pCng:bool</math>; 1. repeat 2. <math>pCng \leftarrow false</math>; 3. foreach <math>e \in \mathcal{E}^k</math> do 4. <math>\mathcal{L} \leftarrow \text{Locals}(e, k, p)</math>; 5. while <math>\mathcal{L} \neq \emptyset</math> do 6. <math>i \leftarrow \text{Pick}(\mathcal{L})</math>; 7. <math>f \leftarrow \text{RecFire}(e, k-1, \langle k.p \rangle[i])</math>; 8. if <math>f \neq \mathbf{0}</math> then 9. foreach <math>j \in \mathcal{N}_e^k(i)</math> do 10. <math>u \leftarrow \text{Union}(k-1, f, \langle k.p \rangle[j])</math>; 11. if <math>u \neq \langle k.p \rangle[j]</math> then 12. <math>\langle k.p \rangle[j] \leftarrow u</math>; <math>pCng \leftarrow true</math>; 13. if <math>\mathcal{N}_e^k(j) \neq \emptyset</math> then 14. <math>\mathcal{L} \leftarrow \mathcal{L} \cup \{j\}</math>; 15. until <math>pCng = false</math>;</pre>	<p><i>Find</i>(in <math>tab, key, out v</math>):<math>bool</math></p> <p>If <math>(key, x)</math> is in hash table <math>tab</math>, set <math>v</math> to <math>x</math> and return <i>true</i>. Else, return <i>false</i>.</p>
<p><i>Union</i>(in <math>k:lvl, p:idx, q:idx</math>):<math>idx</math></p> <p>Build an MDD rooted at <math>\langle k.s \rangle</math>, in <math>UT[k]</math>, encoding <math>\mathcal{B}(\langle k.p \rangle) \cup \mathcal{B}(\langle k.q \rangle)</math>. Return <math>s</math>.</p> <pre> declare <math>i:lcl</math>; declare <math>s, u:idx</math>; 1. if <math>p = 1</math> or <math>q = 1</math> then return <math>1</math>; 2. if <math>p = \mathbf{0}</math> or <math>p = q</math> then return <math>q</math>; 3. if <math>q = \mathbf{0}</math> then return <math>p</math>; 4. if <i>Find</i>(<math>UC[k], \{p, q\}, s</math>) then return <math>s</math>; 5. <math>s \leftarrow \text{NewNode}(k)</math>; 6. for <math>i = 0</math> to <math>n^k - 1</math> do 7. <math>u \leftarrow \text{Union}(k-1, \langle k.p \rangle[i], \langle k.q \rangle[i])</math>; 8. <math>\langle k.s \rangle[i] \leftarrow u</math>; 9. <i>Check</i>(<math>k, s</math>); <i>Insert</i>(<math>UC[k], \{p, q\}, s</math>); 10. return <math>s</math>;</pre>	<p><i>Insert</i>(inout <math>tab, in key, v</math>)</p> <p>Insert <math>(key, v)</math> in hash table <math>tab</math>, if it does not contain an entry <math>(key, \cdot)</math>.</p>
<p><i>Locals</i>(in <math>e:evnt, k:lvl, p:idx</math>):set of <math>lcl</math></p> <p>Return <math>\{i \in \mathcal{S}^k : \langle k.p \rangle[i] \neq \mathbf{0}, \mathcal{N}_e^k(i) \neq \emptyset\}</math>, the local states in <math>p</math> locally enabling <math>e</math>. Return <math>\emptyset</math> or <math>\{i \in \mathcal{S}^k : \mathcal{N}_e^k(i) \neq \emptyset\}</math>, respectively, if <math>p</math> is <math>\mathbf{0}</math> or <math>1</math>.</p>	<p><i>Pick</i>(inout <math>\mathcal{L}</math>:set of <math>lcl</math>):<math>lcl</math></p> <p>Remove and return an element from <math>\mathcal{L}</math>.</p>
<p><i>NewNode</i>(in <math>k:lvl</math>):<math>idx</math></p> <p>Create <math>\langle k.p \rangle</math> with arcs set to <math>\mathbf{0}</math>, return <math>p</math>.</p>	<p><i>Check</i>(in <math>k:lvl, inout p:idx</math>)</p> <p>If <math>\langle k.p \rangle</math>, not in <math>UT[k]</math>, duplicates <math>\langle k.q \rangle</math>, in <math>UT[k]</math>, delete <math>\langle k.p \rangle</math> and set <math>p</math> to <math>q</math>. Else, insert <math>\langle k.p \rangle</math> in <math>UT[k]</math>. If <math>\langle k.p \rangle[0] = \dots = \langle k.p \rangle[n^k - 1] = \mathbf{0}</math> or <math>\mathbf{1}</math>, delete <math>\langle k.p \rangle</math> and set <math>p</math> to <math>\mathbf{0}</math> or <math>\mathbf{1}</math>, since <math>\mathcal{B}(\langle k.p \rangle)</math> is <math>\emptyset</math> or <math>\mathcal{S}^k \times \dots \times \mathcal{S}^1</math>, respectively.</p>

Fig. 2. Pseudo-code for the node-saturation algorithm.

for nodes at level  $k$ , to retrieve  $s$  given node  $p$  and event  $e$ , where  $First(e) > k$  and  $\mathcal{B}(\langle k.s \rangle) = \mathcal{N}_{\leq k}^*(\mathcal{N}_e(\mathcal{B}(\langle k.p \rangle)))$ . Furthermore, we use  $K$  dynamically-sized arrays to store nodes, so that  $\langle k.p \rangle$  can be efficiently retrieved as the  $p^{\text{th}}$  entry of the  $k^{\text{th}}$  array. The call  $Generate(s)$  creates the MDD encoding the initial state, saturating each MDD node as soon as it creates it, in a bottom-up fashion. Hence, when it calls  $Saturate(k, r)$ , all children of  $\langle k.r \rangle$  are already saturated.

**Theorem 1 (Correctness).** *Consider a node  $\langle k.p \rangle$  with  $K \geq k \geq 1$  and saturated children. Moreover, (a) let  $\langle l.q \rangle$  be one of its children, satisfying  $q \neq \mathbf{0}$  and  $l = k - 1$ ; (b) let  $\mathcal{U}$  stand for  $\mathcal{B}(\langle l.q \rangle)$  before the call  $RecFire(e, l, q)$ , for some event  $e$  with  $l < First(e)$ , and let  $\mathcal{V}$  represent  $\mathcal{B}(\langle l.f \rangle)$ , where  $f$  is the value returned by this call; and (c) let  $\mathcal{X}$  and  $\mathcal{Y}$  denote  $\mathcal{B}(\langle k.p \rangle)$  before and after calling  $Saturate(k, p)$ , respectively. Then, (i)  $\mathcal{V} = \mathcal{N}_{\leq l}^*(\mathcal{N}_e(\mathcal{U}))$  and (ii)  $\mathcal{Y} = \mathcal{N}_{\leq k}^*(\mathcal{X})$ .*

By choosing, for node  $\langle k.p \rangle$ , the root  $\langle K.r \rangle$  of the MDD representing the initial system state  $s$ , we obtain  $\mathcal{Y} = \mathcal{N}_{\leq K}^*(\mathcal{B}(\langle K.r \rangle)) = \mathcal{N}_{\leq K}^*(\{s\}) = \mathcal{S}$ , as desired.

*Proof.* To prove both statements we employ a simultaneous induction on  $k$ . For the induction base,  $k = 1$ , we have: (i) The only possible call  $RecFire(e, 0, \mathbf{1})$  immediately returns  $\mathbf{1}$  because of the test on  $l$  (cf. line 1). Then,  $\mathcal{U} = \mathcal{V} = \{()\}$  and  $\{()\} = \mathcal{N}_{\leq 0}^*(\mathcal{N}_e(\{()\}))$ . (ii) The call  $Saturate(1, p)$  repeatedly explores  $\lambda^1$ , the only event in  $\mathcal{E}^1$ , in every local state  $i$  for which  $\mathcal{N}_{\lambda^1}^1(i) \neq \emptyset$  and for which  $\langle 1.p \rangle[i]$  is either  $\mathbf{1}$  at the beginning of the “while  $\mathcal{L} \neq \emptyset$ ” loop, or has been modified (cf. line 12) from  $\mathbf{0}$  to  $\mathbf{1}$ , which is the value of  $f$ , hence  $u$ , since the call  $RecFire(e, 0, \mathbf{1})$  returns  $\mathbf{1}$ . The iteration stops when further attempts to fire  $\lambda^1$  do not add any new state to  $\mathcal{B}(\langle 1.p \rangle)$ . At this point,  $\mathcal{Y} = \mathcal{N}_{\lambda^1}^*(\mathcal{X}) = \mathcal{N}_{\leq 1}^*(\mathcal{X})$ .

For the induction step we assume that the calls to  $Saturate(k-1, \cdot)$  as well as to  $RecFire(e, l-1, \cdot)$  work correctly. Recall that  $l = k - 1$ .

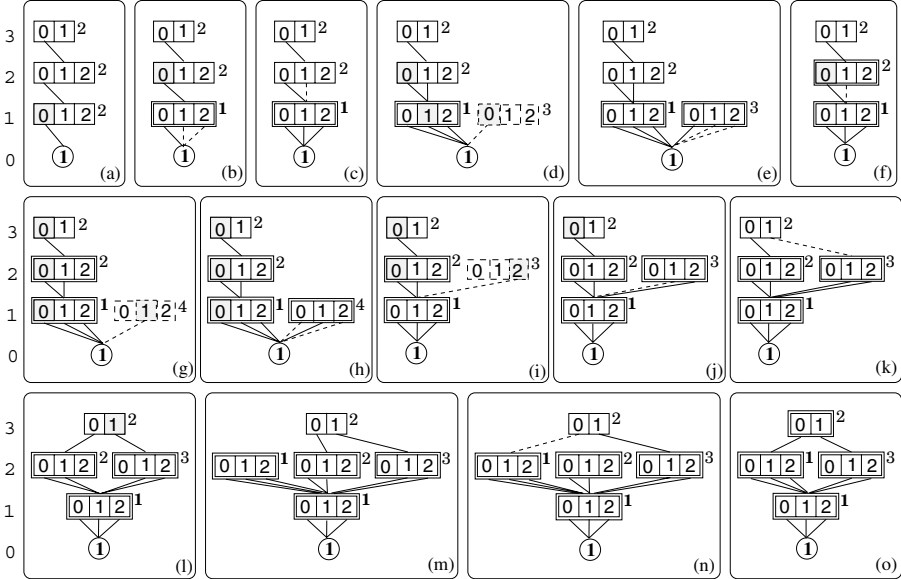
- (i) Unlike  $Saturate$  (cf. line 14),  $RecFire$  does not add further local states to  $\mathcal{L}$ , since it modifies “in-place” the new node  $\langle l.s \rangle$ , and not node  $\langle l.q \rangle$  describing the states from where the firing is explored. The call  $RecFire(e, l, q)$  can be resolved in three ways. If  $l < Last(e)$ , then the returned value is  $f = q$  and  $\mathcal{N}_e^l(\mathcal{U}) = \mathcal{U}$  for any set  $\mathcal{U}$ ; since  $q$  is saturated,  $\mathcal{B}(\langle l.q \rangle) = \mathcal{N}_{\leq l}^*(\mathcal{B}(\langle l.q \rangle)) = \mathcal{N}_{\leq l}^*(\mathcal{N}_e(\mathcal{B}(\langle l.q \rangle)))$ . If  $l \geq Last(e)$  but  $RecFire$  has been called previously with the same parameters, then the call  $Find(FC[l], \{q, e\}, s)$  is successful. Since node  $q$  is saturated and in the unique table, it has not been modified further; note that in-place updates are performed only on nodes not yet in the unique table. Thus, the value  $s$  in the cache is still valid and can be safely used. Finally, we need to consider the case where the call  $RecFire(e, l, q)$  performs “real work.” First, a new node  $\langle l.s \rangle$  is created, having all its arcs initialized to  $\mathbf{0}$ . We explore the firing of  $e$  in each state  $i$  satisfying  $\langle l.q \rangle[i] \neq \mathbf{0}$  and  $\mathcal{N}_e^l(i) \neq \emptyset$ . By induction hypothesis, the recursive call  $RecFire(e, l-1, \langle l.q \rangle[i])$  returns  $\mathcal{N}_{\leq l-1}^*(\mathcal{N}_e(\mathcal{B}(\langle l-1.\langle l.q \rangle[i] \rangle)))$ . Hence, when the “while  $\mathcal{L} \neq \emptyset$ ” loop terminates,  $\mathcal{B}(\langle l.s \rangle) = \bigcup_{i \in \mathcal{S}^l} \mathcal{N}_e^l(i) \times \mathcal{N}_{\leq l-1}^*(\mathcal{N}_e(\mathcal{B}(\langle l-1.\langle l.q \rangle[i] \rangle))) = \mathcal{N}_{\leq l-1}^*(\mathcal{N}_e(\mathcal{B}(\langle l.q \rangle)))$  holds. Thus, all children of node  $\langle l.s \rangle$  are saturated. According to the induction hypothesis, the call

- $Saturate(l, s)$  correctly saturates  $\langle l.s \rangle$ . Consequently, we have  $\mathcal{B}(\langle l.s \rangle) = \mathcal{N}_{\leq l}^*(\mathcal{N}_{\leq l-1}^*(\mathcal{N}_e(\mathcal{B}(\langle l.q \rangle)))) = \mathcal{N}_{\leq l}^*(\mathcal{N}_e(\mathcal{B}(\langle l.q \rangle)))$  after the call.
- (ii) As in the base case,  $Saturate(k, p)$  repeatedly explores the firing of each event  $e$  that is locally enabled in  $i \in \mathcal{S}^k$ ; it calls  $RecFire(e, k-1, \langle k.p \rangle[i])$  that, as shown above and since  $l=k-1$ , returns  $\mathcal{N}_{\leq k-1}^*(\mathcal{N}_e(\mathcal{B}(\langle k-1.\langle k.p \rangle[i] \rangle)))$ . Further,  $Saturate(k, p)$  terminates when firing the events in  $\mathcal{E}^k = \{e_1, e_2, \dots, e_m\}$  does not add any new state to  $\mathcal{B}(\langle k.p \rangle)$ . At this point, the set  $\mathcal{Y}$  encoded by  $\langle k.p \rangle$  is the fixed-point of the iteration

$$\mathcal{Y}^{(m+1)} \Leftarrow \mathcal{Y}^{(m)} \cup \mathcal{N}_{\leq k-1}^*(\mathcal{N}_{e_1}(\mathcal{N}_{\leq k-1}^*(\mathcal{N}_{e_2}(\dots \mathcal{N}_{\leq k-1}^*(\mathcal{N}_{e_m}(\mathcal{Y}^{(m)})) \dots)))),$$

initialized with  $\mathcal{Y}^{(0)} \Leftarrow \mathcal{X}$ . Hence,  $\mathcal{Y} = \mathcal{N}_{\leq k}^*(\mathcal{X})$ , as desired. □

level	event: $l_1$	event: $l_2$	event: $l_3$	event: $e_{21}$	event: $e_{321}$
3	*	*	$1 \rightarrow 0$	*	$0 \rightarrow 1$
2	*	$0 \rightarrow 1, 2 \rightarrow 1$	*	$0 \rightarrow 1$	$0 \rightarrow 2$
1	$0 \rightarrow 1, 1 \rightarrow 2, 2 \rightarrow 0$	*	*	$1 \rightarrow 0$	$0 \rightarrow 1$



**Fig. 3.** Example of the execution of the *Saturate* and *RecFire* routines.

Fig. 3 illustrates our saturation-based state-space generation algorithm on a small example, where  $K = 3$ ,  $|\mathcal{S}_3| = 2$ ,  $|\mathcal{S}_2| = 3$ , and  $|\mathcal{S}_1| = 3$ . The initial state is  $(0, 0, 0)$ , and there are three local events  $l_1$ ,  $l_2$ , and  $l_3$ , plus two further events,  $e_{21}$  (depending on levels 2 and 1) and  $e_{321}$  (depending on all levels). Their effects, i.e., their next-state functions, are summarized in the table at the top of Fig. 3; the symbol “\*” indicates that a level does not affect an event. The



MDD encoding  $\{(0, 0, 0)\}$  is displayed in Snapshot (a). Nodes  $\langle 3.2 \rangle$  and  $\langle 2.2 \rangle$  are actually created in Steps (b) and (g), respectively, but we show them from the beginning for clarity. The level  $lvl$  of a node  $\langle lvl.idx \rangle$  is given at the very left of the MDD figures, whereas the index  $idx$  is shown to the right of each node. We use dashed lines for newly created objects, double boxes for saturated nodes, and shaded local states for substates enabling the event to be fired. We do not show nodes with index  $0$ , nor any arcs to them.

- *Snapshots (a–b)*: The call  $Saturate(1, 2)$  updates node  $\langle 1.2 \rangle$  to represent the effect of firing  $l_1^*$ ; the result is equal to the reserved node  $\langle 1.1 \rangle$ .
- *Snapshots (b–f)*: The call  $Saturate(2, 2)$  fires event  $l_2$ , adding arc  $\langle 2.2 \rangle[1]$  to  $\langle 1.1 \rangle$  (cf. Snapshot (c)). It also fires event  $e_{21}$  which finds the “enabling pattern”  $(*, 0, 1)$ , with arbitrary first component, and starts building the result of the firing, through the sequence of calls  $RecFire(e_{21}, 1, \langle 2.2 \rangle[0])$  and  $RecFire(e_{21}, 0, \langle 1.1 \rangle[1])$ . Once node  $\langle 1.3 \rangle$  is created and its arc  $\langle 1.3 \rangle[0]$  is set to  $\mathbf{1}$  (cf. Snapshot (d)), it is saturated by repeatedly firing event  $l_1$ . Node  $\langle 1.3 \rangle$  then becomes identical to node  $\langle 1.1 \rangle$  (cf. Snapshot (e)). Hence, it is not added to the unique table but deleted. Returning from  $RecFire$  on level 1 with result  $\langle 1.1 \rangle$ , arc  $\langle 2.2 \rangle[1]$  is updated to point to the outcome of the firing (cf. Snapshot (f)). This does not add any new state to the MDD, since  $\{1\} \times \{0\}$  was already encoded in  $\mathcal{B}(\langle 2.2 \rangle)$ .
- *Snapshots (f–o)*: Once  $\langle 2.2 \rangle$  is saturated, we call  $Saturate(3, 2)$ . Local event  $l_3$  is not enabled, but event  $e_{321}$  is, by the pattern  $(0, 0, 0)$ . The calls to  $RecFire$  build a chain of nodes encoding the result of the firing (cf. Snapshots (g–i)). Each of them is in turn saturated (cf. Snapshots (h–j)), causing first the newly created node  $\langle 1.4 \rangle$  to be deleted, since it becomes equal to node  $\langle 1.1 \rangle$ , and second the saturated node  $\langle 2.3 \rangle$  to be added to the MDD. The firing of  $e_{321}$  (cf. Snapshot (k)) not only adds state  $(1, 2, 1)$ , but the entire subspace  $\{1\} \times \{1, 2\} \times \mathcal{S}^1$ , now known to be exhaustively explored, as node  $\langle 2.3 \rangle$  is marked saturated. Event  $l_3$ , which was found disabled in node  $\langle 3.2 \rangle$  at the first attempt, is now enabled, and its firing calls  $Union(2, \langle 3.2 \rangle[1], \langle 3.2 \rangle[0])$ . The result is a new node which is found by  $Check$  to be the reserved node  $\langle 2.1 \rangle$  (cf. Snapshot (m)). This node encoding  $\mathcal{S}_2 \times \mathcal{S}_1$  is added as the descendant of node  $\langle 3.2 \rangle$  in position 0, and the former descendant  $\langle 2.2 \rangle$  in that position is removed (cf. Snapshot (n)), causing it to become disconnected and deleted. Further attempts to fire events  $l_3$  or  $e_{321}$  add no more states to the MDD, whence node  $\langle 3.2 \rangle$  is declared saturated (cf. Snapshot (o)). Thus, our algorithm terminates and returns the MDD encoding of the overall state space  $(\{0\} \times \mathcal{S}^2 \times \mathcal{S}^1) \cup (\{1\} \times \{1, 2\} \times \mathcal{S}^1)$ .

To summarize, since MDD nodes are saturated as soon as they are created, each node will either be present in the final diagram or will eventually become disconnected, but it will never be modified further. This reduces the amount of work needed to explore subspaces. Once all events in  $\mathcal{E}^k$  are exhaustively fired in some node  $\langle k.p \rangle$ , any additional state discovered that uses  $\langle k.p \rangle$  for its encoding benefits in advance from the “knowledge” encapsulated in  $\langle k.p \rangle$  and its descendants.

## 4 Garbage Collection and Optimizations

**Garbage collection.** MDD nodes can become disconnected, i.e., unreachable from the root, and should be “recycled.” Disconnection is detected by associating an *incoming-arc counter* to each node  $\langle k.p \rangle$ . Recycling disconnected nodes is a major issue in traditional symbolic state-space generation algorithms, where usually many nodes become disconnected. In our algorithm, this phenomenon is much less frequent, and the best runtime is achieved by removing these nodes only at the end; we refer to this policy as **LAZY** policy.

We also implemented a **STRICT** policy where, if a node  $\langle k.p \rangle$  becomes disconnected, its “delete-flag” is set and its arcs  $\langle k.p \rangle[i]$  are re-directed to  $\langle k-1.0 \rangle$ , with possible recursive effects on the nodes downstream. When a hit in the union cache  $UC[k]$  or the firing cache  $FC[k]$  returns  $s$ , we consider this entry stale if the delete-flag of node  $\langle k.s \rangle$  is set. By keeping a per-level count of the nodes with delete-flag set, we can decide in routine *NewNode*( $k$ ) whether to (a) allocate new memory for a node at level  $k$  or (b) recycle the indices and the physical memory of all nodes at level  $k$  with delete-flag set, after having removed all the entries in  $UC[k]$  and  $FC[k]$  referring to them. The threshold that triggers recycling can be set in terms of number of nodes or bytes of memory. The policy using a threshold of one node, denoted as **STRICT**(1), is optimal in terms of memory consumption, but has a higher overhead due to more frequent clean-ups.

**Optimizations.** First, observe that the two outermost loops in *Saturate* ensure that firing some event  $e \in \mathcal{E}^k$  does not add any new state. If we always consider these events in the same order, we can stop iterating as soon as  $|\mathcal{E}^k|$  consecutive events have been explored without revealing any new state. This saves  $|\mathcal{E}^k|/2$  firing attempts on average, which translates to speed-ups of up to 25% in our experimental studies. Also, in *Union*, the call  $Insert(UC[k], \{p, q\}, s)$  records that  $\mathcal{B}(\langle k.s \rangle) = \mathcal{B}(\langle k.p \rangle) \cup \mathcal{B}(\langle k.q \rangle)$ . Since this implies  $\mathcal{B}(\langle k.s \rangle) = \mathcal{B}(\langle k.p \rangle) \cup \mathcal{B}(\langle k.s \rangle)$  and  $\mathcal{B}(\langle k.s \rangle) = \mathcal{B}(\langle k.s \rangle) \cup \mathcal{B}(\langle k.q \rangle)$ , we can, optionally, also issue the calls  $Insert(UC[k], \{p, s\}, s)$ , if  $s \neq p$ , and  $Insert(UC[k], \{q, s\}, s)$ , if  $s \neq q$ . This *speculative union heuristic* improves performance by up to 20%.

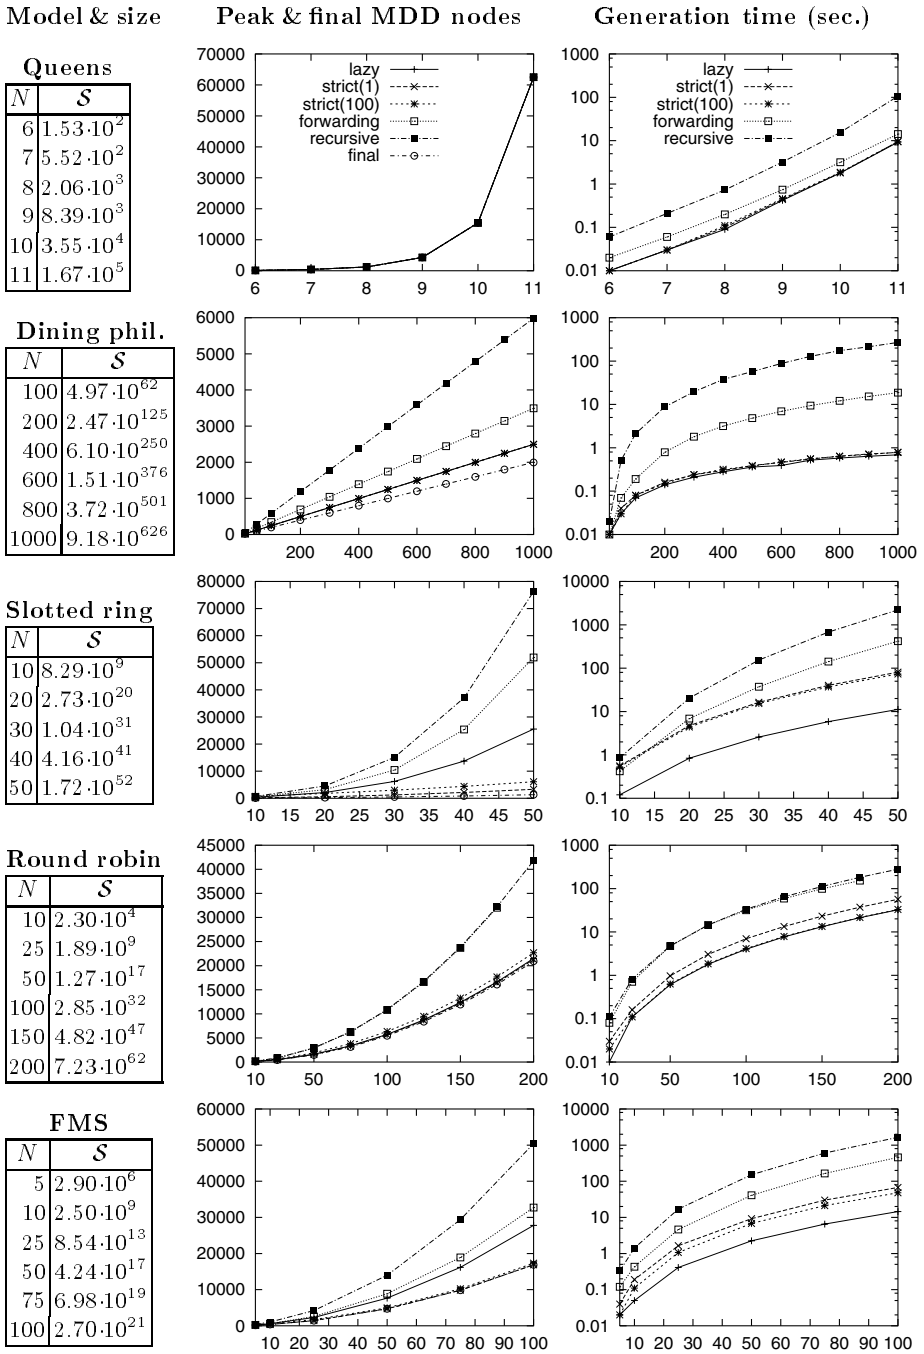
## 5 Experimental Results

In this section we compare the performance of our new algorithm, using both the **STRICT** and **LAZY** policies, with previous MDD-based ones, namely the traditional **RECURSIVE** MDD approach in [22] and the level-by-level **FORWARDING-ARCS** approach in [10]. All three approaches are implemented in **SMART** [11], a tool for the logical and stochastic-timing analysis of discrete-state systems. For asynchronous systems, these approaches greatly outperform the more traditional **BDD**-based approaches [20], where next-state functions are encoded using decision diagrams. To evaluate our saturation algorithm, we have chosen a suite of examples with a wide range of characteristics. In all cases, the state space sizes depend on a parameter  $N \in \mathbb{N}$ .

- The classic *N queens problem* requires to find a way to position  $N$  queens on a  $N \times N$  chess board such that they do not attack each other. Since there will be exactly one queen per row in the final solution, we use a *safe* (i.e., at most one token per place) Petri net model with  $N \times N$  transitions and  $N$  rows, one per MDD level, of  $N + 1$  places. For  $1 \leq i, j \leq N$ , place  $p_{ij}$  is initially empty, and place  $p_{i0}$  contains the token (queen) still to be placed on row  $i$  of the chess board. Transition  $t_{ij}$  moves the queen from place  $p_{i0}$  to place  $p_{ij}$ , in competition with all other transitions  $t_{il}$ , for  $l \neq j$ . To encode the mutual exclusion of queens on the same column or diagonal, we employ *inhibitor arcs*. A correct placement of the  $N$  queens corresponds to a marking where all places  $p_{i0}$  are empty. Note that our state space contains *all* reachable markings, including those where queens  $n$  to  $N$  still need to be placed, for any  $n$ . In this model, locality is poor, since  $t_{ij}$  depends on levels 1 through  $i$ .
- The *dining philosophers* and *slotted ring models* [10, 25] are obtained by connecting  $N$  identical safe subnets in a circular fashion. The MDD has  $N/2$  MDD levels (two subnets per level) for the former model and  $N$  levels (one subnet per level) for the latter. Events are either local or synchronize adjacent subnets, thus they span only two levels, except for those synchronizing subnet  $N$  with subnet 1, which span the entire MDD.
- The *round-robin mutex protocol* model [17] also has  $N$  identical safe subnets placed in a circular fashion, which represent  $N$  processes, each mapped to one MDD level. Another subnet models a resource shared by the  $N$  processes, giving raise to one more level, at the bottom of the MDD. There are no local events and, in addition to events synchronizing adjacent subnets, the model contains events synchronizing levels  $n$  and 1, for  $2 \leq n \leq N + 1$ .
- The *flexible manufacturing system* (FMS) model [22] has a fixed shape, but is parameterized by the initial number  $N$  of tokens in some places. We partition this model into 19 subnets, giving rise to a 19-level MDD with a moderate degree of locality, as events span from two to six levels.

Fig. 4 compares three variants of our new algorithm, using the LAZY policy or the STRICT policy with thresholds of 1 or 100 nodes per level, against the RECURSIVE algorithm in [22] and the FORWARDING algorithm in [10]. We ran SMART on a 800 MHz Intel Pentium III PC under Linux. On the left column, Fig. 4 reports the size of the state space for each model and value of  $N$ . The graphs in the middle and right columns show the peak and final number of MDD nodes and the CPU time in seconds required for the state-space generations, respectively.

For the models introduced above, our new approach is up to two orders of magnitude faster than [22] (a speed-up factor of 384 is obtained for the 1000 dining philosophers' model), and up to one order of magnitude faster than [10] (a speed-up factor of 38 is achieved for the slotted ring model with 50 slots). These results are observed for the LAZY variant of the algorithm, which yields the best runtimes; the STRICT policy also outperforms [22] and [10]. Furthermore, the gap keeps increasing as we scale up the models. Just as important, the saturation algorithm tends to use many fewer MDD nodes, hence less memory. This is most apparent in the FMS model, where the difference between the peak and



**Fig. 4.** State-space sizes, memory consumption, and generation times (a logarithmic scale is used on the y-axis for the latter). Note that the curves in the upper left diagram are almost identical, thus they visually coincide.

the final number of nodes is just a constant, 10, for any STRICT policy. Also notable is the reduced memory consumption for the slotted ring model, where the STRICT(1) policy uses 23 times fewer nodes compared to [22], for  $N = 50$ . In terms of absolute memory requirements, the number of nodes is essentially proportional to bytes of memory. For reference, the largest memory consumption in our experiments using saturation was recorded at 9.7MB for the FMS model with 100 tokens; auxiliary data structures required up to 2.5MB for encoding the next-state functions and 200KB for storing the local state spaces, while the caches used less than 1MB. Other SMART structures account for another 4MB.

In a nutshell, regarding generation time, the best algorithm is LAZY, followed by STRICT(100), STRICT(1), FORWARDING, and RECURSIVE. With respect to memory consumption, the best algorithm is STRICT(1), followed by STRICT(100), LAZY, FORWARDING, and RECURSIVE. Thus, our new algorithm is consistently faster and uses less memory than previously proposed approaches. The worst model for all algorithms is the queens problem, which has a very large number of nodes in the final representation of  $\mathcal{S}$  and little locality. Even here, however, our algorithm uses slightly fewer nodes and is substantially faster.

## 6 Related Work

We already pointed out the significant differences of our approach to symbolic state-space generation when compared to traditional approaches reported in the literature [20], which are usually deployed for model checking [12]. Hence, for a fair comparison, we should extend our algorithmic implementation to that of a full model checker first. Doing this is out of the scope of the present paper and is currently work in progress.

The following paragraphs briefly survey some orthogonal and alternative approaches to improving the scalability of symbolic state-space generation and model-checking techniques. Regarding *synchronous hardware systems*, symbolic techniques using BDDs, which can represent state spaces in sublinear space, have been thoroughly investigated. Several implementations of BDDs are available; we refer the reader to [27] for a survey on BDD packages and their performance. To improve the time efficiency of BDD-based algorithms, breadth-first BDD-manipulation algorithms [4] have been explored and compared against the traditional depth-first ones. However, the results show no significant speed-ups, although breadth-first algorithms lead to more regular access patterns of hash tables and caches. Regarding space efficiency, a fair amount of work has concentrated on choosing appropriate variable orderings and on dynamically re-ordering variables [15].

For *asynchronous software systems*, symbolic techniques have been investigated less, and mostly only in the setting of Petri nets. For safe Petri nets, BDD-based algorithms for the generation of the reachability set have been developed in [25] via encoding each place of a net as a Boolean variable. These algorithms are capable of generating state spaces of large nets within hours. Recently, more efficient encodings of nets have been introduced, which take place invariants into

account [24], although the underlying logic is still based on Boolean variables. In contrast, our work uses a more general version of decision diagrams, namely MDDs [18, 22], where more complex information is carried in each node of a diagram. In particular, MDDs allow for a natural encoding of asynchronous system models, such as distributed embedded systems.

For the sake of completeness, we briefly mention some other BDD-based techniques exploiting the component-based structure of many digital systems. They include partial model checking [3], compositional model checking [19], partial-order reduction [2], and conjunctive decompositions [21]. Finally, also note that approaches to symbolic verification have been developed, which do not rely on decision diagrams but instead on arithmetic or algebra [1, 6, 26].

## 7 Conclusions and Future Work

We presented a novel approach for constructing the state spaces of asynchronous system models using MDDs. By avoiding to encode the global next-state function as an MDD, but splitting it into several local next-state functions instead, we gained the freedom to choose the sequence of event firings, which controls the fixed-point iteration resulting in the desired global state space. Our central contribution is the development of an elegant iteration strategy based on saturating MDD nodes. Its utility is proved by experimental studies which show that our algorithm often performs several orders of magnitude faster than most existing algorithms. Equally important, the peak size of the MDD is usually kept close to its final size.

Regarding future work, we plan to employ our idea of saturation for implementing an MDD-based CTL model checker within SMART [11], to compare that model checker to state-of-the-art BDD-based model checkers, and to test our tool on examples that are extracted from real software.

**Acknowledgments.** We would like to thank the anonymous referees for their valuable comments and suggestions.

## References

- [1] P. Aziz Abdulla, P. Bjesse, and N. Eén. Symbolic reachability analysis based on SAT-solvers. In *TACAS 2000*, vol. 1785 of *LNCS*, pp. 411–425. Springer-Verlag, 2000.
- [2] R. Alur, R.K. Brayton, T.A. Henzinger, S. Qadeer, and S.K. Rajamani. Partial-order reduction in symbolic state-space exploration. In *CAV '97*, vol. 1254 of *LNCS*, pp. 340–351. Springer-Verlag, 1997.
- [3] H.R. Andersen, J. Staunstrup, and N. Maretti. Partial model checking with ROBDDs. In *TACAS '97*, vol. 1217 of *LNCS*, pp. 35–49. Springer-Verlag, 1997.
- [4] P. Ashar and M. Cheong. Efficient breadth-first manipulation of binary decision diagrams. In *ICCAD '94*, pp. 622–627. Computer Society Press, 1994.
- [5] J.A. Bergstra, A. Ponse, and S.A. Smolka. *Handbook of Process Algebra*. Elsevier Science, 2000.

- [6] A. Biere, A. Cimatti, E.M. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS '99*, vol. 1579 of *LNCS*, pp. 193–207. Springer-Verlag, 1999.
- [7] R.E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Trans. on Comp.*, 35(8):677–691, 1986.
- [8] R.E. Bryant. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Comp. Surveys*, 24(3):393–418, 1992.
- [9] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. Symbolic model checking:  $10^{20}$  states and beyond. *Inform. & Comp.*, 98(2):142–170, 1992.
- [10] G. Ciardo, G. Lüttgen, and R. Siminiceanu. Efficient symbolic state-space construction for asynchronous systems. In *ICATPN 2000*, vol. 1639 of *LNCS*, pp. 103–122. Springer-Verlag, 2000.
- [11] G. Ciardo and A.S. Miner. SMART: Simulation and Markovian Analyzer for Reliability and Timing. In *IPDS '96*, p. 60. IEEE Computer Society Press, 1996.
- [12] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: A new symbolic model verifier. In *CAV '99*, vol. 1633 of *LNCS*, pp. 495–499. Springer-Verlag, 1999.
- [13] E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [14] R. Cleaveland, J. Parrow, and B. Steffen. The Concurrency Workbench: A semantics-based tool for the verification of finite-state systems. *TOPLAS*, 15(1):36–72, 1993.
- [15] M. Fujita, H. Fujisawa, and Y. Matsunaga. Variable ordering algorithms for ordered binary decision diagrams and their evaluation. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 12(1):6–12, 1993.
- [16] A. Geser, J. Knoop, G. Lüttgen, B. Steffen, and O. Rütting. Chaotic fixed point iterations. Techn. Rep. MIP-9403, Univ. of Passau, 1994.
- [17] S. Graf, B. Steffen, and G. Lüttgen. Compositional minimisation of finite state systems using interface specifications. *Formal Asp. of Comp.*, 8(5):607–616, 1996.
- [18] T. Kam, T. Villa, R.K. Brayton, and A. Sangiovanni-Vincentelli. Multi-valued decision diagrams: Theory and applications. *Multiple-Valued Logic*, 4(1–2):9–62, 1998.
- [19] K. Larsen, P. Pettersson, and W. Yi. Compositional and symbolic model-checking of real-time systems. In *RTSS '95*, pp. 76–89. Computer Society Press, 1995.
- [20] K.L. McMillan. *Symbolic Model Checking: An Approach to the State-explosion Problem*. PhD thesis, Carnegie-Mellon Univ., 1992.
- [21] K.L. McMillan. A conjunctively decomposed Boolean representation for symbolic model checking. In *CAV '96*, LNCS, pp. 13–24. Springer-Verlag, 1996.
- [22] A.S. Miner and G. Ciardo. Efficient reachability set generation and storage using decision diagrams. In *ICATPN '99*, vol. 1639 of *LNCS*, pp. 6–25. Springer-Verlag, 1999.
- [23] T. Murata. Petri nets: Properties, analysis and applications. *Proc. of the IEEE*, 77(4):541–579, 1989.
- [24] E. Pastor and J. Cortadella. Efficient encoding schemes for symbolic analysis of Petri nets. In *DATE '98*, pp. 790–795. IEEE Computer Society Press, 1998.
- [25] E. Pastor, O. Roig, J. Cortadella, and R.M. Badia. Petri net analysis using Boolean manipulation. In *ICATPN '94*, vol. 815 of *LNCS*, pp. 416–435. Springer-Verlag, 1994.
- [26] M. Sheeran and G. Stålmarck. A tutorial on Stålmarck's proof procedure for propositional logic. *Formal Methods in System Design*, 16(1):23–58, 2000.
- [27] B. Yang, R.E. Bryant, D.R. O'Hallaron, A. Biere, O. Coudert, G. Janssen, R.K. Ranjan, and F. Somenzi. A performance study of BDD-based model checking. In *FMCAD '98*, vol. 1522 of *LNCS*, pp. 255–289. Springer-Verlag, 1998.