

Memory-Efficient Kronecker Algorithms with Applications to the Modelling of Parallel Systems

Anne Benoit, Brigitte Plateau
Laboratoire Informatique et Distribution
ENSIMAG - Antenne de Montbonnot
ZIRST 51, Avenue Jean Kuntzmann
38330 Montbonnot St Martin - France
{benoit,plateau}@imag.fr

William J. Stewart
Department of Computer Science
North Carolina State University
Raleigh, N.C. 27695-8206, USA
billy@csc.ncsu.edu

Abstract

Parallel and distributed systems can be modelled as a set of interacting components. This has an impact on the mathematical structure of the model, namely it induces a product form represented by a tensor product. We present a new algorithm for computing the solution of large Markov chain models whose generators can be represented in the form of a generalized tensor algebra, such as networks of stochastic automata. The tensor structure inherently involves a product state space but inside this product state space, the actual reachable state space can be much smaller. For such cases, we propose an improvement of the standard numerical algorithm, the so-called “shuffle algorithm”, which necessitates only vectors of the size of the actual state space. With this contribution, numerical algorithms based on tensor products can now handle much larger models, even with functional rates and synchronizing events.

1 Introduction

Continuous Time Markov Chains (CTMC) facilitate the performance analysis of dynamic systems in many areas of application [31], and are particularly well adapted to the study of parallel and distributed systems, [1, 16, 29]. They are often used in a high level formalism in which a software package is employed to generate the state space and the infinitesimal generator of the CTMC, as well as to compute stationary and transient solutions. Several high-level formalisms have been proposed to help model very large and complex CTMCs in a compact and structured manner. For example, stochastic automata networks (SANs), [26, 13], queueing networks, [17], generalized stochastic Petri nets, [23], stochastic reward nets, [25] and stochastic activity nets, [28] are, thanks to their extensive mod-

elling capabilities, widely used in diverse application domains, and notably in the areas of parallel and distributed systems. It is therefore possible to generate the CTMC that represents the system to be studied from formalisms such as these which model entire systems from their interconnected subcomponents. Such formalisms often involve a product state space \hat{S} but inside this product state space, the actual reachable state space S can be much smaller ($S \subseteq \hat{S}$).

In this paper, our concern is with the computation of the stationary probability vector $\pi \in R^{|\hat{S}|}$, a row vector whose i^{th} element π_i is the probability of being in state i of the CTMC at a time that is sufficiently long for all influence of the initial starting state to have been erased, and where S is the set of states of the CTMC. The vector π is the solution of the system of linear equations $\pi Q = 0$, subject to $\pi e = 1$, where Q is the generator matrix of the CTMC and e is a vector whose elements are all equal to 1.

The primary difficulty in developing software to handle large-scale Markov chains comes from the explosion in the number of states that usually occurs when certain model parameters are augmented. Indeed, CTMCs which model real parallel and distributed systems are usually huge and sophisticated algorithms are needed to handle them. Both the amount of available memory and the time taken to generate them and to compute their solutions need to be carefully analyzed. For example, direct solution methods, such as Gaussian elimination, are generally not used because the amount of fill-in that occurs necessitates a prohibitive amount of storage space. Iterative methods, which can take advantage of sparse storage techniques to hold the infinitesimal generator, are more appropriate, even though here also, memory requirements can become too large for real life models.

During the generation of the CTMC from the high-level formalism, particular attention must be paid to the way in which the quantities S , Q and π are stored. Indeed, even though the computational cost of a particular numerical so-

lution method may be relatively high, it is essentially storage considerations that determine whether this method may be used or not. Furthermore, iterative methods often need to compute the product of a probability vector and the generator Q many many times. They therefore depend directly upon the manner in which these data are stored.

Different techniques have been proposed for storing the **set of reachable states** S and for performing efficient searches in this set. In certain cases, it is possible to define a reachability function for the model and this may be used to provide access to the reachable states in constant time, [15]. In other cases, multi-level approaches such as decision diagrams [8] may prove to be efficient. In what follows, we suppose that S is stored in an efficient manner and we will no longer be concerned with this particular aspect of the algorithms.

There have been many techniques proposed to represent the **descriptor** Q and the **probability vector** π . A first approach, independent of the high-level formalism, consists of storing the matrix in a row-wise compact sparse format: the nonzero elements of the matrix and their position in the matrix are kept. The probability vectors are the size of the reachable state space $|S|$. Efficient algorithms are available with which to compute a vector-matrix product when the matrix is stored in this fashion [27, 30]. However, for very large models, it is frequently the case that the matrix is too large to be held in memory.

Sanders and his co-workers [9] proposed to generate the elements of Q as they were needed, “on the fly” from a high level model. However, this approach can be computationally expensive, especially when the model contains immediate transitions. A different proposal is to store Q on a fast disk and to seek the elements as and when they are needed [10]. This obviously requires a disk that is sufficiently fast and sufficiently large.

Stochastic Automata Networks (SANs) were introduced by Plateau et al. ([26, 13]) to keep memory requirements manageable. These allow Markov chains models to be described in a memory efficient manner because their storage is based on a tensor formalism. Another formalism based on Stochastic Petri Nets allows us to obtain a similar tensor formalism, as shown by Donatelli in [11, 12]. However, the use of independent components connected via synchronizations and functions may produce a representation with many unreachable states ($|S| \ll |\hat{S}|$). Within this Kronecker framework, a number of algorithms have been proposed. The first and perhaps best-known, is the shuffle algorithm [13, 14, 26, 3], which computes the product but never needs the matrix explicitly. However, as has been shown previously, this algorithm needs to use vectors $\hat{\pi}$ the size of \hat{S} , that we shall call *extended*. We denote this algorithm **E-Sh**, for **extended shuffle**. Some alternative approaches have been proposed in [22, 21, 6, 3]. They consist of first com-

puting the state space S , and then solving the model by using iteration vectors π which contains entries only for these states (size of S), that we shall call *reduced*.

Another approach consists of using decision diagrams in order to represent the Markov chain generator Q . In this tree representation only distinct nonzero elements are stored, and they are kept as the leaves of the tree. On initial construction, the branches of the tree must be pruned to keep it from becoming too large. Multi-terminal binary decision diagrams (MTBDDs) and probabilistic decision graphs (PDGs) can be used in model checking, e.g. [20, 4], or in performance evaluation [19]. This is a promising approach, but its efficiency in memory requirements is directly linked to the number of distinct elements in Q . Moreover, access to the nonzero elements requires following a path from the root of the tree to the appropriate leaf, and this must be done every time an element is needed. As far as we are aware, there are no comparative studies on the effectiveness of these algorithms in actually computing solutions of Markov chains.

In order to improve the computation time, another kind of representation is presented in [7, 24]. This representation of the generator Q as a matrix diagram permits quick access to all of their elements, and the solution time is often satisfactory, even for large models. Furthermore, memory requirement to store Q are barely greater than that needed for a descriptor stored in tensor format and remains negligible compared to that needed to store a probability vector. However, techniques linked to a cache, [7] make theoretical estimations of the complexity of the algorithms difficult, whereas good theoretical results may be obtained with the Kronecker approach, [6].

When $|S| \approx |\hat{S}|$, the gain in memory obtained with the use of the tensor formalism can be enormous compared to the standard approach of an explicit sparse matrix storage. For example, if a model consists of K components of size n_i ($i = 1..K$), the infinitesimal generator is full and the space needed to store it with the standard approach is of the order of $(\prod_{i=1}^K n_i)^2$. The use of a tensor formalism reduces this cost to $\sum_{i=1}^K n_i^2$. The shuffle algorithm E-Sh is very effective in this case, as has been shown in [13, 14, 26, 3].

However, when there are many unreachable states ($|S| \ll |\hat{S}|$), E-Sh is not efficient, because of its use of extended vectors. The probability vector can therefore have many zero elements, since only states corresponding to reachable states have nonzero probability. Moreover, computations are carried out for all the elements of the vector, even those elements corresponding to unreachable states. Therefore, the gain obtained by exploiting the tensor formalism can be lost since many useless computations are performed, and memory is used for states whose probability is always zero. In this case, the approach first described, based on storing all nonzero elements of the generator, per-

forms better since it does not carry out meaningless computations. However, its extensive memory requirements prevent its use on extremely large models.

Due to the large memory requirements of this last approach, it is worthwhile seeking a solution which takes unreachable states into account and at the same time, uses the benefits of the tensor formalism. Thus, we would like to be able to exploit the tensor formalism, even in the presence of an important number of unreachable states. Indeed, the shuffle is an efficient algorithm when $|S| \approx |\hat{S}|$, and we will show how to modify it efficiently for the case in which $|S| \ll |\hat{S}|$. It is interesting to work on this algorithm which conserves its tensor structure and for which we possess certain theoretical results, and which, in addition, is used in diverse domains (signal processing [18], partial differential equations,...).

The use of reduced vectors (of size S) permits a reduction in memory used, and some needless computations are avoided. This leads to significant memory gains when using iterative methods such as Arnoldi or GMRES which can possibly require many probability vectors. A modification to the E-Sh shuffle algorithm permits the use of such vectors. However, to obtain good performance at the computation-time level, some intermediate vectors of size \hat{S} are also used. The algorithm Act-Sh-JCB of [6] transforms the reduced vector π into a vector $\hat{\pi}$ of size \hat{S} before calling E-Sh. An algorithm proposed in [3] permits us to save computations by taking account of the fact that the probabilities corresponding to nonreachable states are always zero in the resulting vector. We refer to this as **partially reduced** and denote the corresponding algorithm **PR-Sh**. However, the saving in memory turns out to be somewhat insignificant for the shuffle algorithm itself.

In summary, the algorithm E-Sh is an efficient algorithm when there are few non reachable states in the system. Modifications to reduce memory needs when $|S| \ll |\hat{S}|$ have been proposed, but the new algorithm PR-Sh uses also intermediate data structures of size \hat{S} [6, 3]. Our goal is to develop algorithms that maintain the desirable complexity of E-Sh, while eliminating the inconvenient data structures of size equal to the size of the product space, which, to our knowledge, has not been done before. We hope that this will minimize memory requirements at the cost of only a small increase in computation time.

In the next section we present the shuffle algorithm E-Sh for the multiplication of a vector and a SAN descriptor. We then present (in Section 3) a new version of the shuffle algorithm, one which concentrates on the amount of memory used, and allows us to handle even more complex models. In this new algorithm, all the intermediate data structures are stored in reduced format. We refer to this as **fully reduced** and denote the corresponding shuffle algorithm **FR-Sh**. A series of tests comparing the two shuffle algorithms is

presented in Section 4. These algorithms were incorporated into the software package PEPS [15] and tested by means of this package.

2 The Shuffle Algorithm

In SANs [13, 14] and in SGSPNs [12], it has been shown that the transition matrix can be expressed as:

$$Q = D + \bigoplus_{i=1}^N Q_i^{(i)} + \sum_{e=1}^E \bigotimes_{i=1}^N Q_e^{(i)} = \sum_{j=1}^{(N+2E)} \left[\bigotimes_{i=1}^N Q_j^{(i)} \right] \quad (1)$$

Here N is the number of automata (resp. sub Petri Nets) in the network and E is the number of synchronizing events (resp. transitions). In the first representation, D is the diagonal of the descriptor; the tensor sum corresponds to the analysis of the local events and is called the local part of the descriptor; the tensor product corresponds to the analysis of the synchronizing events and is called the synchronizing part of the descriptor. The second representation shows that the basic operation of interest in computing the steady state vector of the Markov chain using iterative methods is the product of a row vector $\hat{\pi}$ and a tensor product¹ $\bigotimes_{i=1}^N Q^{(i)}$. This term is composed of a sequence of N matrices denoted $Q^{(i)}$ with $i \in [1 \dots N]$, each associated with an automaton $\mathcal{A}^{(i)}$. We begin by introducing some definitions concerning finite sequences of matrices:

- n_i is the order of the i^{th} matrix in a sequence;
- $nleft_i$ is the product of the order of all the matrices that are to the left of the i^{th} matrix of a sequence, i.e., $\prod_{k=1}^{i-1} n_k$ with the special case: $nleft_1 = 1$;
- $nright_i$ is the product of the order of all the matrices that are to the right of the i^{th} matrix of a sequence, i.e., $\prod_{k=i+1}^N n_k$ with the special case: $nright_N = 1$;

We first assume that there are no functional dependencies. According to the decomposition property of tensor products [13], every tensor product of N matrices is equivalent to the product of N normal factors. To compute the multiplication of a vector by the term $\bigotimes_{i=1}^N Q^{(i)}$ it is therefore sufficient to know how to multiply a row-vector $\hat{\pi}$ and a normal factor i : $\hat{\pi} \times I_{nleft_i} \otimes Q^{(i)} \otimes I_{nright_i}$. Furthermore, the property of commutativity between normal factors allows the multiplication of the normal factors in any desired order. The matrix $I_{nleft_i} \otimes Q^{(i)} \otimes I_{nright_i}$ is a block diagonal matrix in which the blocks are the matrix $Q^{(i)} \otimes I_{nright_i}$. We can treat the different blocks in an independent manner, which suggests the possibility of introducing parallelism into the algorithm, [32]. There are $nleft_i$ blocks of the matrix each of which is to be multiplied by a different piece of the vector, which

¹The indices j have been omitted from the matrices $Q_j^{(i)}$ in order to simplify the notation.

we shall call a *vector slice*. The vector is divided according to I_{nleft_i} and so we shall call these vector slices, l -slices (for left) and denote them by $p_0, \dots, p_{nleft_i-1}$, each of size $nright_i \times n_i$.

Thus, the E-Sh algorithm consists of a loop over the $nleft_i$ l -slices, and at each iteration it computes the resulting l -slice, $r_l = p_l \times Q^{(i)} \otimes I_{nright_i}$. The computation of an element of the resulting l -slice corresponds to the multiplication of p_l by a column of the matrix $Q^{(i)} \otimes I_{nright_i}$. Due to the structure of this matrix, the multiplication therefore boils down to the repeated extraction of components of p_l (at distance $nright_i$ apart), forming a vector called z_{in} from these components and then multiplying z_{in} by a column of the matrix $Q^{(i)}$. Notice carefully that z_{in} is composed of elements of p_l which may not be consecutive. In a certain sense, it is a slice of p_l . The slice z_{in} corresponds to elements of p_l which must be multiplied by the elements of the column of $Q^{(i)}$. The structure of the matrix $Q^{(i)} \otimes I_{nright_i}$ informs us that we must consider $nright_i$ slices z_{in} . Thus, we number the z_{in} from 0 to $nright_i - 1$. Extracting a z_{in} is the same as accessing the vector p_l and choosing the elements at intervals of $nright_i$ positions in the vector.

Once a z_{in} has been obtained, we can compute an element of the result by multiplying the z_{in} by a column of $Q^{(i)}$. The multiplication of a z_{in} by the entire matrix $Q^{(i)}$ therefore provides several elements of the result, a slice of the result, called z_{out} . The positions of the elements of z_{out} in r_l correspond to the positions of the elements of z_{in} in p_l . We number the z_{out} in the same way as we did for z_{in} . Therefore, the multiplication of the z_{in} $n^\circ k$ by $Q^{(i)}$ gives $z_{out} n^\circ k$.

A pseudo-code of the E-Sh can be found in [3].

Complexity: The complexity of the algorithm for multiplying a vector and a classic tensor product may be obtained by observing the number of vector–matrix multiplications that are executed [13, 14]. If we assume that the matrices $Q^{(i)}$ are stored in a sparse compacted form, and letting nz_i denote the number of nonzero elements in the matrix $Q^{(i)}$, then the complexity of E-Sh is given by

$$\prod_{i=1}^N n_i \times \sum_{i=1}^N \frac{nz_i}{n_i} = |\hat{S}| \times \sum_{i=1}^N \frac{nz_i}{n_i}$$

This should be compared to the multiplication of a vector and $Q = \otimes_{i=1}^N Q^{(i)}$ which consists in first computing Q , and then multiplying the result with the vector. This has complexity of the order of $|\hat{S}|^2$ when the matrices $Q^{(i)}$ are full, and $\prod_{i=1}^N nz_i$ for a sparse compacted format. So, E-Sh is better than this last multiplication for $nz_i \geq n_i \cdot N^{\frac{1}{N-1}}$. The complexity of the algorithm handling functional dependencies can be found in [3, 13].

The E-Sh has two drawbacks: it uses vector data struc-

ture of size \hat{S} and computes all vector elements, even the zero entries out of \hat{S} . In what follows, we remove these two drawbacks. The idea is the following: for each normal factor, the nonzero elements of the vector are reordered so that the elements of a z_{in} are consecutive.

3 Reduced Memory Shuffle Algorithm

Data Structures Used: The fact that we shall store only the values of the vector that correspond to reachable states implies that we must keep track of the positions of these elements in the corresponding vector in \hat{S} . We shall assume that the set of reachable states is already known. This may be done either by applying an algorithm that explores all reachable state [22, 8], or by asking the user to provide a function which represents the set of reachable states of the system [15].

We shall let $pos_{\hat{\pi}}$ denote the set of states which have nonzero entries in any vector $\hat{\pi}$. The **probability vectors** are such that $pos_{\hat{\pi}} \subseteq S$, and we represent them with the help of two arrays of size $|S|$: the array $\hat{\pi}.vec$ contains the entry values and the array $\hat{\pi}.positions$ contains the positions of the reachable states in the corresponding vector in \hat{S} . It is however important to notice that processing the multiplication of a vector and a normal factor corresponding to a synchronizing event (with matrix $Q_e^{(i)}$ in Equation 1) may lead to an **intermediate vector** $\hat{\pi}$ for which we do not have $pos_{\hat{\pi}} \subseteq S$. We can equally well store these intermediate vectors in two arrays $\hat{\pi}.vec$ and $\hat{\pi}.positions$ of size $|S|$ if $|pos_{\hat{\pi}}| \leq |S|$. In this case, the positions no longer correspond to the reachable states, but they correspond to $pos_{\hat{\pi}}$. Notice that most often $|pos_{\hat{\pi}}| \leq |S|$ since the intermediate vectors are obtained via multiplication with synchronization matrices, matrices which are generally very sparse. These vectors therefore contain many zero elements. However, it may happen that $|pos_{\hat{\pi}}| > |S|$. In this case, it becomes necessary to dynamically reallocate memory in which to store the vector².

The principal difficulty in applying ideas of algorithm E-Sh with the new reduced vector data structure lies in extracting the slices of the vector z_{in} from an l -slice. Indeed, the previous algorithm used a skipping procedure to extract the vector slices. When the vector is stored in a reduced structure, it is not possible to perform these skips. A somewhat similar method can however be adopted by traversing the vector and extracting all the z_{in} for the l -slice being treated. For this, we use an intermediate structure³ that consists of a set of triplets (number, place, index), each triplet corre-

²We have not yet implemented dynamic reallocation in PEPS. This happens very rarely in our experience and is the reason why we have not had need of dynamic reallocation for the moment.

³To implement this structure, we used the Standard Template Library (STL) *containers*.

sponding to an element of the vector $\hat{\pi}$: **number** represents the number of the z_{in} containing this element; **place** represents the place of the element in the z_{in} ; **index** represents the index (the position) of the element in the reduced vector ($1 \leq index \leq |S|$), it allows us to find the value of the element as $\hat{\pi}.vec[index]$.

The different steps of the multiplication $w = v \times I_{nleft_i} \otimes Q^{(i)} \otimes I_{nright_i}$ for the FR-Sh algorithm are detailed below. We differentiate between the nature of w (probability or intermediate vector) to optimize the performance of the algorithm.

Notice first that an l -slice is a set of consecutive elements and a single sweep over the vector v allows us to determine the limits of l -slices. For $l = 0, \dots, nleft_i - 1$, the l -slice p_l contains elements with index x such that $n_i \times nright_i \times l \leq v.positions(x) < n_i \times nright_i \times (l + 1)$. We handle the l -slices in a sequential manner (but it can be done in parallel): for each l -slice, we begin by collecting the necessary information for an efficient execution of the remainder of the algorithm. The set of triplets *infozin* contains informations on the z_{in} of the l -slice, while *infozout* give informations on the z_{out} when available (the slice to consider in this case is a l -slice of w).

Treatment of an l -slice:

1. Filling *infozin*: We first perform a traversal of the l -slice of $v.positions$, to construct a set of triplets called *infozin*. For each element, we obtain *number* and *place* by performing an integer division of $v.positions[index]$ by $nright_i$: $number = v.positions[index] \bmod nright_i$ and $place = v.positions[index] \div nright_i$.

An integer division by $nright_i$ boils down to fetching the elements that are spaced $nright_i$ apart in the extended vector.

2. Filling *infozout*: When w is a probability vector, the array $w.positions$ is initialized with the positions of the states of S . This means that we can perform a traversal of the l -slice of $w.positions$, to construct a set *infozout* similar to *infozin*, but which takes into account the positions of the vector w (these may be different from those of v if v is an intermediate vector). Thus, the fields *number* and *place* respectively represent for each element of w , the number of the z_{out} and the place of the element in the z_{out} . When w is an intermediate vector, we have no information on w so there is no *infozout*.

3. Sorting *infozin* and *infozout*: We need to perform a sort⁴ on the sets of triplets according to *number* so that all the elements of the same z_{in} (or z_{out}) end up in adjacent positions.

4. Treatment of a z_{in} :

a. Extraction of the z_{in} : The elements of the same z_{in} are now placed consecutively and a single traversal of *in-*

⁴The STL sort used is *introsort*, a variant of *quicksort* which offers a complexity of $O(N \log N)$ in the worst case.

fozin allows us to get the z_{in} one after the other. When w is a probability vector, we extract the z_{in} only if at least one of the states of the corresponding z_{out} is reachable. This is obtained by a single traversal of *infozout*, because the elements of the same z_{out} are now placed consecutively. So we need not necessary to extract all the z_{in} as it is done in E-Sh.

b. Multiplication: When w is a probability vector, *infozout* informs us of which elements need to be multiplied, and we can finally perform the multiplication of the z_{in} by the corresponding columns of the matrix. In this case, we do not have to perform all the products as in E-Sh. On the other hand, when w is an intermediate vector, we do not know a priori the position of the nonzero elements, so we perform the sparse multiplication of z_{in} by the entire matrix, and we store all the nonzero elements as a (position, value) pair.

5. Storing the result: When w is a probability vector, we perform the multiplication for each reachable state and the information contained in *infozout* tells us where to store the value obtained ($w.vec[index]$). The storage therefore takes place as and when the computations are performed. On the other hand, when w is an intermediate vector, we must wait until the end of the computations with all of the z_{in} , and then sort the nonzero elements by increasing position order. A linear traversal of these elements is sufficient to fill the tables $w.positions$ and $w.vec$ ⁵.

When all the z_{in} of a given l -slice have been treated, then the following l -slice is treated.

A pseudo-code of the FR-Sh can be found in [3].

Complexity of the algorithm: Notice first that we no longer use arrays of size $|\hat{S}|$. All the arrays are of size $|S|$ since we assume that there is no intermediate vector $\hat{\pi}$ such that $|pos_{\hat{\pi}}| > |S|$. As far as computation time is concerned, we reduce the number of multiplications to the order of $|S| \times \sum_{i=1}^N \frac{nz_i}{ni}$ when w is a probability vector (where we assume that the number of nonzero elements per column is uniform). However, we introduce some supplementary costs, most notably, the cost of a sort which could reach $O(|S| \log(|S|))$.

When the percentage of nonreachable states is high ($|S| \ll |\hat{S}|$), the improvement is significant. The computation and memory cost are somewhat higher than E-Sh when $|S| \approx |\hat{S}|$. We have therefore fulfilled our objective of not using any structure of size $|\hat{S}|$ while maintaining an efficient algorithm when $|S| \ll |\hat{S}|$.

4 Numerical results

Now that the algorithms have been presented, we shall study their performance and compare them with one another

⁵It is at this moment that a dynamic reallocation may be needed.

on the basis of both memory needs and execution time.

All numerical experiments were performed using the software package, PEPS ([15]), into which we implemented the new algorithm. All execution times were measured to within a tenth of a second on a 531 MHz Pentium III PC with 128 MB of memory. Convergence was verified to an absolute precision of ten decimal places, i.e., the results have a tolerance of the order of 10^{-10} . In all experiments, the initial vectors were chosen to be equiprobable, and we used the *unpreconditioned power method*. Furthermore, since the new algorithm does not alter the speed of convergence of the methods, we only provide the total time needed during execution and not the number of iterations required for convergence.

The memory use is taken from the system during execution. It represents the totality of memory used by PEPS during its execution (i.e., during the solution of a model). This includes the data, memory structures reserved by the procedure, and also the process stack. The only parameters that change from one algorithm to the next are the memory structures reserved by the algorithms (probability vectors, intermediate array structures, and so on).

We present results obtained from two classical parallel systems chosen from the literature ([13, 14, 2]). The first model, called **mutex1**, performs resource sharing with the use of functions; the second, **mutex2**, represents the same model, but with functions replaced by synchronizing transitions. The results obtained from a queuing network model can be found in [3].

In both models, N distinct clients share P identical units of a common resource. A customer requests use of a resource at a rate λ and frees up resource at a rate μ . For the tests, we use the values $\lambda = 6$ and $\mu = 9$. The values of N and P are varied according to the experiment.

Mutex 1

For this first model, a function is used to represent the mechanism by which access to the units of resource is restricted. The semantics of this function is as follows: *access permission is granted if at least one unit of the resource is available*. Freeing up resource, as opposed to acquiring resource, occurs in an independent manner.

The SAN product state space for this model is of size 2^N . Notice that when $P = 1$, the reachable state space is of size $N + 1$, which is considerably smaller than the product state space, while when $P = N$ the reachable state space is the entire product state space. Other values of P give rise to intermediate cases.

The first results we give are for the case in which $N = 16$, and in varying P . In all cases, we have $|\hat{S}| = 65,536$, and, naturally $|S|$ changes with P .

| P | Model | | E-Sh | | FR-Sh | |
|-----|-------|-----------------|-----------|--------|----------|--------|
| | $ S $ | $ S / \hat{S} $ | Time sec. | Mem Kb | Time Sec | Mem Kb |
| 1 | 17 | 0.03% | 27.2 | 4208 | 3.5 | 2124 |
| 4 | 2517 | 3.8% | 112 | 4208 | 22.1 | 2304 |
| 6 | 14893 | 22.7% | 178.2 | 4208 | 93.1 | 3128 |
| 8 | 39203 | 59.8% | 235.4 | 4208 | 267.4 | 4748 |
| 10 | 58651 | 89.5% | 292.7 | 4212 | 478.9 | 6040 |
| 12 | 64839 | 98.9% | 332 | 4212 | 612.9 | 6448 |
| 16 | 65536 | 100% | 27 | 4212 | 377.8 | 6516 |

We notice, as expected, that FR-Sh is better than E-Sh in terms of execution time so long as the percentage of reachable states remains reasonable (less than 50%). During execution, we notice that algorithm FR-Sh permits a reduction in memory needs with respect to algorithm E-Sh as long as the percentage of reachable states is sufficiently high. Indeed, when we are working in reduced vector format, we use structures of the size of $|S|$, but an element of the vector needs additional information which is also stored. Thus, once we exceed 50% of reachable states, we should not hope to produce a gain in memory since a vector is stored with the help of two arrays of size $|S|$, which is greater than $|\hat{S}|$.

We performed additional experiments on larger models to determine the limits and the possibilities of the algorithms developed. When $N = 24$, only algorithm FR-Sh is successful. With $P = 10$, we have $|\hat{S}| = 16,777,216$ and $|S| = 4,540,386$. The solution is obtained in 101,348 seconds and requires 307,768 Kb of memory.

The limits of this algorithm are however reached when the percentage of nonreachable states is high. When $P \geq 12$, the memory needs of algorithm FR-Sh become excessive and the solution cannot be computed.

Mutex 2

Let us now look at how this same system may be modeled without using functional transitions. One possibility is to introduce an additional automaton, a resource pool automaton, which counts the number of units of resource available at any moment. The action of a process in acquiring a resource could then be represented as a synchronizing event requiring the cooperation of the demanding process and the resource pool. A further synchronizing event would be needed for a process to return resource to the resource pool.

The SAN product state space for this model is of size $2^N * (P + 1)$, and the reachable state space size is identical to that of *Mutex1*.

This model allows us to test the efficiency of the new algorithm in the presence of synchronizations, and hence intermediate vectors w which do not satisfy $pos_w \subseteq S$. We observe that FR-Sh remains efficient in this case.

| P | Model | | E-Sh | | FR-Sh | |
|----|-------------|-------------------------|-----------|--------|----------|--------|
| | $ \hat{S} $ | $\frac{ S }{ \hat{S} }$ | Time sec. | Mem Kb | Time Sec | Mem Kb |
| 1 | 131072 | 0.01% | 13.6 | 6480 | 2.2 | 2368 |
| 4 | 327680 | 0.77% | 85.4 | 12672 | 11.1 | 2628 |
| 6 | 458752 | 3.25% | 174.9 | 16800 | 47.9 | 3628 |
| 8 | 589824 | 6.65% | 296.5 | 20928 | 150.3 | 5564 |
| 10 | 720896 | 8.14% | 563.1 | 25060 | 350.9 | 7116 |
| 12 | 851968 | 7.61% | 1002.5 | 29188 | 603.5 | 7628 |
| 16 | 1114112 | 5.88% | 1861.5 | 37452 | 918.2 | 7780 |

Furthermore, the use of synchronizations implies a large number of nonreachable states. Algorithm FR-Sh is therefore the most appropriate when the model contains synchronized transitions, both from the point of view of memory needs as well as from the point of view of execution time.

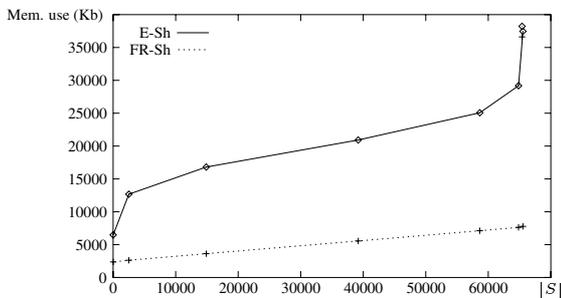


Figure 1. Memory use for the mutex2 model

To verify that the memory required by the new algorithm really depends on $|S|$ and no longer on $|\hat{S}|$, we have drawn some curves that depict memory used by the two algorithms (Figure 1). The memory used by FR-Sh is proportional to $|S|$, whereas this is not the case for algorithm E-Sh. For E-Sh we get a straight line in plotting memory use as a function of $|\hat{S}|$, which goes to show that the memory used is indeed proportional to $|\hat{S}|$. The origin of the straight line corresponds to the minimum memory needed for the execution of algorithm FR-Sh on this model (2400 Kb).

5 Conclusions

In this paper we presented a new algorithm based on the shuffle algorithm. The comparative tests that were performed allow us to reach the following conclusion:

- **Algorithm E-Sh, or the extended shuffle algorithm** is the most efficient when the percentage of reachable states is large. So, whenever more than half the states are reachable, this is the preferred algorithm.

- **Algorithm FR-Sh, or the fully reduced shuffle algorithm** keeps all data structures at the reduced size. This

permits us to handle very large models, models which the other algorithm cannot handle. We notice however a loss in computation time and in memory needs with respect to E-Sh when there are few unreachable states. This is the price we pay in order to be able to handle the largest models. A formal study of the complexity as a function of the percentage of reachable states and of the problem size is part of our future work.

We shall also compare our results with those obtained in generating the global matrix (in Harwell-Boeing (HB) format) and then performing a vector-matrix multiplication using standard sparse matrix multiplication. When it is possible to obtain the global matrix, there is no doubt that time-wise this algorithm performs best. It is from the memory point of view that it is limited. Therefore, the new algorithm allows us to solve models that can not be solved with the HBF algorithm. In the SANs formalism, the use of functions allows a decrease in the size of the product state space. The use of a generalized tensor algebra ([13, 2]) permits tensor operations on matrices to have functional characteristics. However, the cost of matrix evaluation is high and so we try to limit their number. Some techniques have been developed in order to decrease the number of matrix evaluations in the shuffle algorithm **E-Sh** [13, 14]. One possibility that we considered is a reordering of the automata. Some details are provided in [3] on the manner in which the newly presented algorithm need to be modified in order to handle the reordering of automata. Automata grouping ([13]) is another technique that may be used to decrease the number of function evaluations, but this technique is not presented here because the reduced storage of vectors does not cause any change to these procedures.

Finally, the new algorithms was only compared to the extended shuffle algorithm and the algorithm that generates the global matrix in HB format. A comparison with other classical algorithms, and notably algorithms that use Petri nets, [5, 21], were not included. Such a comparison will be performed in future work. We also propose to perform comparisons with our algorithms and matrix diagrams. Indeed, Ciardo in [7] presents a comparison between these two approaches, but he does not use the most recent algorithm developed for the Kronecker approach. His results show that matrix diagrams have a substantial advantage on the Act-Sh-JCB of [6], but it is now necessary to perform new experiments with FR-Sh in order to have comparisons with an algorithm which improves both the memory needs and the computation time when there is a lot of unreachable states.

These improvements to the shuffle algorithm will allow us to model and analyze even larger parallel and distributed systems, something that seems to us to be of primary importance in the evolution of such systems.

References

- [1] S. Alouf, F. Huet, and P. Nain. Forwarders vs. centralized server: an evaluation of two approaches for locating mobile agents. *Performance Evaluation*, 49:299–319, Sept. 2002.
- [2] A. Benoit, P. Fernandes, B. Plateau, and W. Stewart. On the Benefits of Using Functional Transitions in Kronecker Modelling. *Submitted to Performance Evaluation*, 2002.
- [3] A. Benoit, B. Plateau, and W. Stewart. Memory Efficient Iterative Methods for Stochastic Automata Networks. Technical report, Rapport de recherche INRIA n. 4259, France, Sept. 2001.
- [4] M. Bozga and O. Maler. On the Representation of Probabilities over Structured Domains. In N. Halbwachs and D. Peled, editors, *Proc. CAV'99*, volume 1633 of *LNCS*, pages 261–273. Springer, June 1999.
- [5] P. Buchholz. Hierarchical structuring of Superposed GSPNs. In *Proc. 7th Int. Workshop on Petri Nets and Performance Models PNPM'97*, pages 81–90, Saint Malo, France. IEEE CS Press, 1997.
- [6] P. Buchholz, G. Ciardo, S. Donatelli, and P. Kemper. Complexity of memory-efficient Kronecker operations with applications to the solution of Markov models. *INFORMS Journal on Computing*, 12(3):203–222, 2000.
- [7] G. Ciardo and A. Miner. A data structure for the efficient Kronecker solution of GSPNs. In *In proc. 8th Workshop on Petri Nets and Performance models, IEEE CS Press, Zaragoza, 1999*.
- [8] G. Ciardo and A. Miner. Efficient reachability set generation and storage using decision diagram. In *In proc. 20th int. Conf. Application and Theory of Petri Nets, LNCS 1639, Springer, 1999*.
- [9] D. Deavours and W. Sanders. “On-The-Fly” Solution Techniques for Stochastic Petri Nets and Extensions. In *Proceedings of the Seventh International Workshop on Petri Nets and Performance Models*, pages 132–141, Saint Malo, France, June 1997. IEEE Computer Society.
- [10] D. Deavours and W. Sanders. An Efficient Disk-Based Tool for Solving Large Markov Models. *Performance Evaluation*, 33(1):67–84, 1998.
- [11] S. Donatelli. Superposed Stochastic Automata: a class of stochastic Petri nets with parallel solution and distributed state space. *Performance Evaluation*, 18(1):21–36, 1993.
- [12] S. Donatelli. Superposed Generalized Stochastic Petri nets: definition and efficient solution. In Valette, R., editor, *Lecture Notes in Computer Science; Application and Theory of Petri Nets 1994, Proceedings 15th International Conference, Zaragoza, Spain*, volume 815, pages 258–277. Springer-Verlag, 1994.
- [13] P. Fernandes. *Méthodes Numériques pour la solution de systèmes Markoviens à grand espace d'états*. PhD thesis, Institut National Polytechnique de Grenoble, France, 1998.
- [14] P. Fernandes, B. Plateau, and W. Stewart. Efficient Descriptor-Vector Multiplications in Stochastic Automata Networks. *JACM*, 45(3):381–414, 1998.
- [15] J. Fourneau and B. Plateau. PEPS: A package for solving complex Markov models of parallel systems. In *Proceedings of the Fourth International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, Palma, Spain, 1988.
- [16] H. Garavel and H. Hermanns. On Combining Functional Verification and Performance Evaluation using CADP. Technical report, Rapport de recherche INRIA n. 4492, France, July 2002.
- [17] E. Gelenbe and G. Pujolle. *Introduction to Queueing Networks, second edition*. John Wiley & Sons, 1998.
- [18] J. Granata, M. Conner, and R. Tolimeri. Recursive fast algorithm and the role of the tensor product. *IEEE Transactions on Signal Processing*, 40(12):2921–2930, Dec. 1992.
- [19] H. Hermanns, J. Meyer-Kayser, and M. Siegle. Multi Terminal Binary Decision Diagram to Represent and Analyse Continuous Time Markov Chains. In *Proc. 3rd int. workshop on the numerical solution of Markov chains, Prencas Universitarias de Zaragoza, Zaragoza, Spain, 2000*.
- [20] J.-P. Katoen, M. Kwiatkowska, G. Norman, and D. Parker. Faster and Symbolic CTMC Model Checking. *Lecture Notes in Computer Science*, 2165, 2001.
- [21] P. Kemper. Numerical Analysis of superposed GSPNs. *IEEE Trans. on Software Engineering*, 22(9), Sept. 1996.
- [22] P. Kemper. Reachability analysis based on structured representations. In W. R. J. Billington, editor, *Proc. 17th int. conf. Application and Theory of Petri Nets, Springer LNCS 1091*, pages 269–288, 1996.
- [23] M. Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with generalized stochastic Petri nets*. John Wiley & Sons, 1995.
- [24] A. Miner. Efficient solution of GSPNs using Canonical Matrix Diagrams. In *Proc. PNPM'01, 9th International Workshop on Petri Nets and Performance Models*, pages 101–110, Aachen, Germany, Sept. 2001.
- [25] J. K. Muppala, G. Ciardo, and K. S. Trivedi. Modeling Using Stochastic Reward Nets. In *MASCOTS'93 International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, San Diego, CA*, pages 367–372. IEEE Comp. Soc. Press., Jan 1993.
- [26] B. Plateau. *De l'Evaluation du parallélisme et de la synchronisation*. PhD thesis, Université de Paris XII, Orsay (France), 1984.
- [27] Y. Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing Company, Boston, 1996.
- [28] W. H. Sanders and J. F. Meyer. Reduced Base Model Construction Methods for Stochastic Activity Networks. *IEEE Journal on Selected Areas in Communications*, 9(1):25–36, 1991.
- [29] R. Stansifer and D. Marinescu. Petri net models of concurrent Ada programs. *Microelectronics and Reliability*, 31(4):577–594, 1991.
- [30] W. Stewart. MARCA: Markov chain analyzer, a software package for Markov modelling. In W. Stewart, editor, *Numerical Solution of Markov Chains*, Marcel Dekker, 1991.
- [31] W. J. Stewart. *An introduction to numerical solution of Markov chains*. Princeton University Press, New Jersey, 1994.
- [32] C. Tadonki and B. Philippe. Parallel Multiplication of a Vector by a Kronecker product of Matrices. *Journal of Parallel and Distributed Computing and Practices*, 2(4), Dec. 1999.