Appeared in
"Predictably Dependable Computing Systems"
ESPRIT Basic Research Series,
Springer- Verlag, 1995, pp.507-518, ISBN 3-540-59334-9

Computable Dependability Bounds for Large Markov Chains

Pierre-Jacques Courtois[1,2] Pierre Semal[2]

[1]AV Nucléaire, Belgium [2]Université Catholique de Louvain

Abstract. A new method to bound the steady-state solution of large Markov chains is presented. The method integrates the concepts of eigenvector polyhedron and of aggregation. It is specially suited for Markov chains with high locality and very large state spaces.

A model of a repairable fault tolerant system with 16 millions states is used as an example. Bounds on its availability are obtained by considering a small part of its state space only. The method is potentially useful to bound other types of dependability requirements.

1 Introduction

This paper introduces a new technique which can be used to efficiently compute lower and upper bounds on reliability and availability measures in very large Markovian models of computer fault tolerant systems. By large models, we mean models that may have millions of states. Only the principles of the method are presented. The interested reader may find additional information in [Semal 1992].

The technique is based on decomposition and aggregation. The principle of aggregation is simple, but not easy to apply. Subsets of states of the model are replaced by single aggregated states, and the system macroscopic behaviour is modelled by the transitions between these aggregated states only. These transitions between subsets are obtained from the transitions between the individual states of the subsets. The objective is an economy of computation. The difficulty comes from the fact that, except for simple or degenerated cases, the transitions between two subsets are not simply obtained from the superposition or the addition of transitions between their individual states; these transitions depend on, and must be weighted by the relative probabilities of these individual states. These probabilities are

unknown, so that, in general, approximations must be used. It is therefore impossible, except for very restricted cases, to obtain exact results by decomposition and aggregation, and the accuracy of the method is an important issue.

In the 70's and 80's, much research work has been done on the estimation of the approximation error for different techniques of aggregation. The approach taken here is different. In [Courtois and Semal 1984], we found possible, when each subset is considered in isolation, to determine lower and upper bounds on the relative steady-state values of the variables and probabilities associated with each state of that subset.

In this paper we apply this result to the efficient computation of bounds on the (un)availability or reliability of system markovian models that are too large to be generated and analysed in their entirety.

2 An Example

We shall explain the technique by applying it on a real system model taken from [Muntz *et al.*1989]. The system is a fault-tolerant and repairable data base and is shown on Figures 1 and 2. It is made of 24 components, each component (front end, processor switch, processor, memory, bus, disk unit) being supposed to be in one of two states, failed or operational. The system is assumed to remain operational as long as there exists at least one path made of operational components connecting one front-end to one disk unit.

There is one single repair facility, the components being repaired on a fixed priority scheme in case of simultaneous failures. Repair times and inter-failure times are assumed to be exponential random variables, with rates in the ranges of [1.0 - 3.0] and [$1.125\ 10^{-4}$ - $40\ 10^{-4}$] respectively.

The total number of distinct states of this model is 2^{24} states, i.e. more than 16 millions states. In order to compute the availability of the system, one needs to know the steady state probability of all those states in which the system is operational. The vector x which contains these probabilities is the steady state vector of the transition probability matrix $\mathbf{Q}$ of the system, and is solution of the equation

$$x = x\,\mathbf{Q} \qquad (1)$$

Grouping together the states with I failed components in block I leads to a block matrix $\mathbf{Q}$ which is block tridiagonal since, all random variables being

exponential, the probability of more than one component failing or being repaired in a same transition is negligible:

$$\begin{pmatrix} \mathbf{Q}_{00} & \mathbf{Q}_{01} & \mathbf{0} & \mathbf{0} & \dots \\ \mathbf{Q}_{10} & \mathbf{Q}_{11} & \mathbf{Q}_{12} & \mathbf{0} & \dots \\ \mathbf{0} & \mathbf{Q}_{21} & \mathbf{Q}_{22} & \mathbf{Q}_{23} & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{Q}_{32} & \mathbf{Q}_{33} & \dots \\ \mathbf{0} & \mathbf{0} & \dots & \dots & \dots \end{pmatrix} \tag{2}$$

For the same reason, all the elements of each diagonal block $\mathbf{Q}_{II}$ are null except on the diagonal. Transitions between states are non zero between states of different and adjacent blocks only. The vector *x* is too large to be computed exactly; in double precision, its mere storage would already require 128 Mbytes. To obtain upper and lower bounds on some of its elements is however possible. The diagonal structure of **Q** and the differences in magnitude of the failure and repair rates are not indispensable, but will be useful for the computation of these bounds.

3 A Useful Polyhedron

A short interlude in matrix algebra is needed because the method is based on a fundamental result in Linear Algebra [Courtois and Semal 1984, 1985]. Suppose that **B** is a non-negative and in this case stochastic matrix of which we only know a lower bound matrix **L**. More precisely:

$$\mathbf{B} = \mathbf{L} + \mathbf{R} \tag{3}$$

where $\mathbf{L} \geq \mathbf{0}$ is known, $\mathbf{R} \geq \mathbf{0}$ is unknown, and both have same size as **B**. Then, if ν is the steady state vector of B, we can write:

$$\begin{aligned} \nu &= \nu \mathbf{B} = \nu (\mathbf{L} + \mathbf{R}) \\ &= \beta \Sigma^{-1} (\mathbf{I} - \mathbf{L})^{-1} \\ &= \beta \begin{bmatrix} - - r_1 - - \\ - - r_2 - - \\ \dots \\ - - r_i - - \\ \dots \end{bmatrix} \end{aligned} \tag{4}$$

where β is a non negative vector whose elements sum up to one, and Σ^{-1} is the inverse of a diagonal matrix which simply normalizes the rows of $(\mathbf{I} - \mathbf{L})^{-1}$ so that its resulting normalized rows r_i also sum up to one. This result expresses the fact that the eigenvector ν of matrix **B** is a convex combination of the rows of the known matrix $\Sigma^{-1}\ (\mathbf{I} - \mathbf{L})^{-1}$. Therefore every element ν_j is bounded by:

$$\min_{\mathbf{i}} \{\mathbf{r_{ij}}\} \leq \nu_{\mathbf{j}} \leq \max_{\mathbf{i}} \{\mathbf{r_{ij}}\} \qquad (5)$$

These bounds have an intuitive significance. If **L** is a substochastic matrix, and is a block embedded in a larger stochastic matrix, the element r_{ij} is the rate of visit to state *j* when the subsystem **L** is started in state *i* before the occurrence of a transition leading outside **L**. Equation (4) also means that the vector ν belongs to the polyhedron defined by the rows r_i.

In other words, if we are interested in the steady-state vector ν of a matrix **B** for which a lower bound matrix **L** only is known, then ν cannot be exactly determined, but it can be bounded by the space spanned by the convex combinations of the rows r_i, i.e. by the polyhedron that we shall note $P[(\mathbf{I} - \mathbf{L})^{-1}]$, and which is defined by these rows. The inequalities (5) are the most practical consequence of this property. Note also that the larger **L** is, i.e. the closer to **B**, the smaller the polyhedron and the tighter the bounds become. Besides, it is proved in [Courtois and Semal 1984, 1985] that, given **L**, the polyhedron and the bounds are the tightest ones that can be obtained.

4 The Method

How can we apply these results to compute bounds on the availability of a system such as the one described in Section 2 ?

Consider again the matrix **Q,** defined by equation (2), where $\mathbf{Q}_{II}$ is the block of transitions between states with exactly *I* failed components. Apply the same partitioning to its steady state vector *x*, and define $\chi_I \nu_I$ as the steady state vector corresponding to block $\mathbf{Q}_{II}$. χ_I is the *marginal probability* of being in any state of the set *I*. ν_I is the vector of *conditional probabilities* for the individual states of that subset.

First we have to compute the bounds of these conditional probabilities; those of the marginal probabilities will be computed in Section 7.

The vector ν_I can be obtained in the following way. Rewrite the matrix **Q** as:

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_{II} & \mathbf{E}(I) \\ \mathbf{F}(I) & \mathbf{G}(I) \end{pmatrix} \tag{6}$$

Then, solving Equation (1) for ν_I leads to:

$$\nu_I \left(\mathbf{Q}_{II} + \mathbf{E}(I)(\mathbf{I}-\mathbf{G}(I))^{-1}\mathbf{F}(I)\right) = \nu_I \tag{7}$$

where the matrix $\mathbf{E}(I)(\mathbf{I}-\mathbf{G}(I))^{-1}\mathbf{F}(I)$ is called the Schurr complement of $\mathbf{Q}_{II}$, is non-negative, and has the same size as $\mathbf{Q}_{II}$. In practice, the computation of this Schurr complement is prohibitive, since $\mathbf{E}(I)$, $\mathbf{G}(I)$ and $\mathbf{F}(I)$ are almost the same size as the whole state space. However, it is quite possible to bound it from below, and, by application of Equation (4), to obtain the polyhedron which contains ν_I. For instance, if this bound is taken equal to 0, one would obtain that:

$$\nu_I \in P\left[\left(\mathbf{I} - \mathbf{Q}_{II}\right)^{-1}\right] \tag{8}$$

or, if one can easily compute a lower bound matrix $\mathbf{S}(I)$ such that:

$$0 \leq \mathbf{S}(I) \leq \mathbf{E}(I)(\mathbf{I}-\mathbf{G}(I))^{-1}\mathbf{F}(I) \tag{9}$$

one can obtain a tighter polyhedron:

$$\nu_I \in P\left[\left(\mathbf{I}-\mathbf{Q}_{II}-\mathbf{S}(I)\right)^{-1}\right] \tag{10}$$

and tighter bounds (5).

A lower bound matrix $\mathbf{S}_\mathbf{p}(I)$ similar to $\mathbf{S}(I)$ can be obtained at little cost by an iterative procedure which is explained in the following section.

5 Iterative Procedure

Suppose for a moment that we already have a first approximation of the conditional probability vectors ν_J for all the blocks J others than I. Then we can compute the transition probabilities between any pair of blocks not involving I; the block transition probability between block J and K is given

by the scalar $\nu_J \mathbf{Q}_{JK}\mathbf{1}'$, where $\mathbf{1}'$ is a column vector of one. If we keep the bloc $\mathbf{Q}_{II}$ unmodified, then $\mathbf{Q}$ reduces to a matrix $\mathbf{Q_p}$ where each block except I corresponds to a single state:

$$\mathbf{Q_p} = \begin{pmatrix} \nu_0\mathbf{Q}_{00}\mathbf{1}'... & \nu_0\mathbf{Q}_{0I} & ...\nu_0\mathbf{Q}_{0N}\mathbf{1}' \\ ..... & & \\ \mathbf{Q}_{I0}\mathbf{1}'... & \mathbf{Q}_{II} & ...\mathbf{Q}_{IN}\mathbf{1}' \\ ..... & & \\ \nu_N\mathbf{Q}_{N0}\mathbf{1}'... & \nu_N\mathbf{Q}_{NI} & ...\nu_N\mathbf{Q}_{NN}\mathbf{1}' \end{pmatrix}$$

Each block (K,I) in column I is replaced by a row vector $\nu_K\mathbf{Q}_{KI}$ of probabilities out of K to each individual state of I. Each block (I,K) in row I is replaced by a vector $\mathbf{Q}_{IK}\mathbf{1}'$ of transition probabilities out of each individual state of I to block K.

If we rewrite $\mathbf{Q_p}$ as

$$\mathbf{Q}_{\mathbf{p}} = \begin{pmatrix} \mathbf{G}_{\mathbf{p}} & \mathbf{F}_{\mathbf{p}} & \mathbf{G}_{\mathbf{p}} \\ \mathbf{E}_{\mathbf{p}} & \mathbf{Q}_{\mathbf{II}} & \mathbf{E}_{\mathbf{p}} \\ \mathbf{G}_{\mathbf{p}} & \mathbf{F}_{\mathbf{P}} & \mathbf{G}_{\mathbf{P}} \end{pmatrix} \qquad (11)$$

the reduced steady state vector X_p of $\mathbf{Q_p}$ which is equal to

$$X_p = (\chi_0 ... \ \chi_{I-1} \ \ \chi_I \nu_I \ \ \chi_{I+1} ... \ \chi_N) \qquad (12)$$

is solution of the equation

$$X_p = X_p \mathbf{Q_P} \qquad (13)$$

and the conditional distribution ν_I is given by:

$$\nu_{\mathbf{I}} = \nu_{\mathbf{I}} \left(\mathbf{Q_{II}} + \mathbf{E_p}(\mathbf{I}-\mathbf{G_p})^{-1}\mathbf{F_p} \right) \qquad (14)$$

where the Schurr complement involves an inverse $(\mathbf{I} - \mathbf{G_p})^{-1}$ of size (N-1)×(N-1) only.

Now, if we have a first set of lower bounds on the vectors ν_J , $J \neq I$, we can obtain a lower bound matrix $\mathbf{S_p}$ at low cost in terms of these bounds, compute a new polyhedron for ν_I, new bounds, and then a set of new matrices $\mathbf{S_p}$, and so on. This process can be iterated indefinitely, starting with a set of initial values for the vectors ν_J obtained for instance by equation (8), the initial matrices $\mathbf{S_p}$ being taken equal to 0. The sequence of $\mathbf{S_p}$ matrices is non-decreasing, and the sequence of polyhedra non-increasing. The process does not diverge. However, because matrices are at each iteration substituted with lower bounds only, some information is never taken into account, and polyhedra will not eventually reduce to a single point, nor the bounds to the true vector.

An attractive possibility of this process is to restrict the analysis to the conditional distributions of interest, with only the corresponding parts of the Markov chain being needed.

6 Tridiagonal Matrices

The computation of the matrices $\mathbf{S_p}$ is further simplified when the matrix $\mathbf{Q}$, as in our example, is block tridiagonal.

As shown in [Semal 1992], the Schurr complements involve no inverse. What needs to be added to $\mathbf{Q}_{II}$ to obtain the bounds on ν_I , are the transitions to the two immediately adjacent blocks only. The equation (14) which defines ν_I reduces, in this case, to:

$$\nu_{\mathbf{I}} = \nu_{\mathbf{I}}\left(\mathbf{Q}_{\mathbf{II}} + \frac{\left(\mathbf{Q}_{\mathbf{I},\mathbf{I}-1}\mathbf{1}'\right)\left(\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\right)}{\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\mathbf{1}'} + \frac{\left(\mathbf{Q}_{\mathbf{I},\mathbf{I}+1}\mathbf{1}'\right)\left(\nu_{\mathbf{I}+1}\mathbf{Q}_{\mathbf{I}+1,\mathbf{I}}\right)}{\nu_{\mathbf{I}+1}\mathbf{Q}_{\mathbf{I}+1,\mathbf{I}}\mathbf{1}'}\right) \qquad (15)$$

Therefore, bounds on $\nu_{I-1}$and ν_{I+1}only are needed to obtain those on ν_I.

7 Bounding The Marginal Distribution

If bounds on elements of the complete vector are needed, one must also obtain bounds on the marginal distribution vector χ. These bounds can be obtained by a procedure similar to the one described above. The marginal distribution is the steady state solution of the matrix $\mathbf{P}$ of transition probabilities between blocks:

$$\mathbf{P}_{\mathbf{IJ}} = \nu_{\mathbf{I}}\mathbf{Q}_{\mathbf{IJ}}\mathbf{1}' \qquad \mathbf{I},\mathbf{J} = 0,...,\mathbf{N} \qquad (16)$$

With lower bounds on the vectors ν_I , one can construct a lower bound matrix for **P**. If $\mathbf{L_P}$ ($\leq$ **P**) is this matrix, then one has that $\chi \in P[(\mathbf{I} - \mathbf{L_p})^{-1}]$, from which bounds on the elements of χ are readily obtained.

Moreover, if, as it is typically the case (e.g; our example), the matrix **Q** is block tridiagonal, then the matrix **P** is tridiagonal., and tighter bounds can be obtained in the following way. A lower bound on χ_0 is obtained by maximising the elements of the upper diagonal of **P** (i.e. the transitions away from block 0), and minimising those of the lower diagonal (the return transitions to block 0). This tridiagonal matrix can then be solved analytically. The technique can be repeated for each component of the vector χ.

8 Complexity

Thus, when bounds on the complete solution vector *x* are needed, the general method works as follows. For the conditional distributions ν_I of interest, polyhedra are computed iteratively as shown by Box 1 until some criterion of convergence is met.

While (convergence criterion is not met) **do**:

step 1.1 Select *I*

step 1.2 Compute the lower bound matrix $\mathbf{S_p}(I)$

step 1.3 Compute the polyhedron $P[(\mathbf{I} - \mathbf{Q}_{II} - \mathbf{S_p}(I))^{-1}]$

step 1.4 Derive bounds on ν_I from this polyhedron.

BOX 1

step 2.1 Compute the lower bound matrix $\mathbf{L_p}$

step 2.2 Compute the polyhedron $P[(\mathbf{I} - \mathbf{L_p})^{-1}]$

step 2.3 Derive bounds on χ from this polyhedron

BOX 2

Polyhedra for the marginal distribution χ are then computed as shown in Box 2. In [Semal 1992], the computation aspects of each of these steps are discussed, and simplifications and optimisations are given. Some of the most essential ones concern the polyhedra which are all computed from matrix inverses of the form $(\mathbf{I}\text{-}\mathbf{L})^{-1}$ When the matrix **L** tends to a stochastic matrix , the polyhedron tends to a single point, which is the Perron-Frobenius eigenvector of **L**. However, the computation of the inverse becomes ill-conditioned since an eigenvalue tends to zero. The existence of such a small eigenvalue $(1 - \rho(\mathbf{L}))$ is however a good sign. It means that the polyhedron is *almost* uniquely defined by the eigenvector corresponding to this eigenvalue. The other eigenvalues and their eigenvectors introduce perturbations only, while in fact they are responsible for enlarging the polyhedron from a single point to a set.

Attention must therefore be given to this eigenvalue $(1 - \rho(\mathbf{L}))$ during the computation of the inverse. If it reduces to round-off errors, the best alternative is to consider the lower bound matrix **L** as being stochastic, and take its Perron-Frobenius vector as the exact value of the vector to be bounded. In all our numerical experiments, however, this eigenvalue remained away from zero, i.e. between 10^{-6} and 10^{-1}. Note also that the determination of a polyhedron $P[\mathbf{A}]$ requires the computation of the normalized rows of **A** only. Normalisation factors can therefore be introduced at any stage of the polyhedron computation.

Note that the brute force computation of an inverse of the form $(\mathbf{I} - \mathbf{Q}_{II} - \mathbf{S})^{-1}$ requires $O(n_I^3)$ operations, where n_I is the size of $\mathbf{Q}_{II}$ and of **S**. This amount is already prohibitive in many applications. It can be reduced because this inverse can be computed as a rank N update of $(\mathbf{I} - \mathbf{Q}_{II})^{-1}$, where N is the rank of the perturbation matrix **S,** and because $(\mathbf{I} - \mathbf{Q}_{II})^{-1}$ usually can be obtained cheaply. In many models, indeed, and in availability models in particular, the diagonal blocks $\mathbf{Q}_{II}$ are such that the matrix $(\mathbf{I} - \mathbf{Q}_{II})^{-1}$ remains very sparse. In models of repairable fault-tolerant systems for example, the fill ratio of this inverse is at most equal to the inverse of $\binom{N}{I}$, because system states with distinct failed components do not communicate through $\mathbf{Q}_{II}$, and thus $\mathbf{Q}_{II}$ is diagonal. In those cases, one can use the expression (see [Semal 92]):

$$\left(\mathbf{I} - \mathbf{Q}_{II} - \mathbf{S}\right)^{-1} = \left(\mathbf{I} - \mathbf{Q}_{II}\right)^{-1} - \mathbf{UV} \qquad (17)$$

where the matrices $\mathbf{U} \in R_{n_I \times N}$, and $\mathbf{V} \in R_{N \times n_I}$ can be computed in $O(N^3 + N^2 n_I + N d_I)$ operations, where d_I is the number of non-null components in $\left(\mathbf{I} - \mathbf{Q}_{II}\right)^{-1}$. The number of operations remains thus linear in n_I. If $(\mathbf{I} - \mathbf{Q}_{II})^{-1}$ *is* not sparse, $(N n_I^2)$ operations will be necessary, which represents a substantial saving compared to $O(n_I^3)$. Further reductions in complexity are possible, and discussed in [Semal,1992].

9 The Example Revisited

Our 2^{24} state space example of section 2 illustrates very well how the method can yield interesting results at surprising low cost. Because of the size of this state space, our objective is to bound the marginal distribution χ and a few conditional distributions ν_I only.

As said earlier, grouping the states with I failed components in block I leads to a stochastic matrix with diagonal blocks $\mathbf{Q}_{II}$ which are themselves diagonal. The inverses $(\mathbf{I}-\mathbf{Q}_{II})^{-1}$ are thus also diagonal. The bounds on ν_I were obtained using the equation (15). Two procedures were coded. A first procedure is needed for the first iteration and must assume that bounds on ν_{I-1} only are available when computing the bounds on ν_I ; the following lower matrix —derived from (15)— is used for this purpose:

$$\mathbf{S}(\mathbf{I}) \leq \frac{\left(\mathbf{Q}_{\mathbf{I},\mathbf{I}-1}\mathbf{1}'\right)\left(\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\right)}{\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\mathbf{1}'}$$

$\mathbf{S}(I)$ is a rank one matrix; using (17) the rank one matrices $\mathbf{U}$ and $\mathbf{V}$ are determined in $O(n_I)$ operations. This matrix $\mathbf{S}(I)$ is also used to compute the bounds corresponding to the last block which is taken into consideration. Bounds on the elements of ν_I are then obtained from (4) and (5).

In the subsequent iterations, values for both the bounds of ν_{I-1} and ν_{I+1} are available when computing those of ν_I , and the following lower bound matrix is used:

$$\mathbf{S}(\mathbf{I}) \leq \frac{\left(\mathbf{Q}_{\mathbf{I},\mathbf{I}-1}\mathbf{1}'\right)\left(\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\right)}{\nu_{\mathbf{I}-1}\mathbf{Q}_{\mathbf{I}-1,\mathbf{I}}\mathbf{1}'} + \frac{\left(\mathbf{Q}_{\mathbf{I},\mathbf{I}+1}\mathbf{1}'\right)\left(\nu_{\mathbf{I}+1}\mathbf{Q}_{\mathbf{I}+1,\mathbf{I}}\right)}{\nu_{\mathbf{I}+1}\mathbf{Q}_{\mathbf{I}+1,\mathbf{I}}\mathbf{1}'}$$

Using (17), the matrices $\mathbf{U}$ and $\mathbf{V}$, which are this time of rank two, are obtained in $O(n_I^2)$ operations.

The conditional distribution ν_0 corresponding to 0 failed components is degenerated to a single state, and ν_0 =1. We have limited ourselves to the calculation of bounds for the conditional distributions ν_1 to ν_F of the first F blocks. Numerical results for F=5 are given in Table 1 for ν_1 and ν_5 at ten successive iterations. Because of their lengths, the bound vectors ν^{inf} and ν^{sup} are given by their sums only; the closer to one these sums are, the tighter the bounds.

Note that increasing the number of iterations will never reduce the polyhedra to single points since the influence of the blocks *I*, *I* > F, is never taken into account. The theory guarantees however that these bounds are the tightest that one can obtain under those circumstances. One can also see from Table 1 that the neglected blocks have a larger influence on ν_5 than on ν_1, the bounds on the latter being tighter than those on the former.

Iter.	ν_1^{inf} **1**'	ν_1^{sup} **1**'	ν_5^{inf} **1**'	ν_5^{sup} **1**'
0	0.97057296	1.51081352	0.66085877	10747.925
1	0.99590078	1.07123235	0.88154804	3587.743
2	0.99920877	1.01375111	0.94878869	1533.806
3	0.99982092	1.00311225	0.96859991	937.785
4	0.99995537	1.00077566	0.97425200	768.459
5	0.99998807	1.00020725	0.97582994	721.243
6	0.99999653	1.00006017	0.97626554	708.213
7	0.99999879	1.00002086	0.97638519	704.634
8	0.99999941	1.00001020	0.97641799	703.653
9	0.99999958	1.00000729	0.97642698	703.384
10	0.99999962	1.00000650	0.97642944	703.311

Table 1: Quality of ν_I bounds for *I*=1 and *I*=5

The computation of bounds for the marginal distribution χ was done by following the procedure of Section 7. However, this procedure, in principle, assumes that conditional distribution bounds have been computed for *all* blocks, while we have explicitly computed those of the first 5 blocks only. In this example, and in availability models in general, this is feasible because the aggregated matrix is tridiagonal. Non-trivial upper and lower bounds on

the repair rate and the failure rate can be determined for each block I, $I > F$, without information on the conditional distributions of these blocks; minimum/maximum repair and failure rates over the states of each block can

I	χ_I^{inf}	χ_I^{sup}
0	9.7524506e-01	9.7524508e-01
1	2.4196310e-02	2.4196323e-02
2	5.4713023e-04	5.4713505e-04
3	1.1261528e-05	1.1263011e-05
4	2.1278181e-07	2.1323658e-07
5	3.6789961e-09	3.8219805e-09
6	2.6455985e-11	8.3829687e-11
7	1.2670814e-13	2.6554535e-12
8	5.2238309e-16	8.3452321e-14
9	1.8370472e-18	2.6017736e-15
10	5.6948463e-21	7.8513118e-17
11	1.5281171e-23	2.2907590e-18
12	3.4637321e-26	6.4546184e-20
13	6.4079044e-29	1.7541567e-21
14	9.7186549e-32	4.5479612e-23
15	1.1500408e-34	1.1222898e-24
16	9.7753471e-38	2.6291623e-26
17	5.0505960e-41	5.8306321e-28
18	2.1885916e-44	1.1870343e-29
19	7.6600706e-48	2.1792270e-31
20	2.1703533e-51	3.5164800e-33
21	4.7024322e-55	4.7952000e-35
22	7.0536483e-59	5.2601890e-37
23	5.8780403e-63	1.3389572e-39
24	2.4491834e-67	1.7852763e-42

Table 2. Bounds on χ

be used for instance. The bounds obtained for χ using the procedure of Section 7 are given in Table 2. Unavailability bounds are given in Table 3, with SUN 4 CPU computing times, for three different values of F.

The computational complexity of the whole algorithm is dominated by the computation of the bounds of the conditional distributions of the first F blocks, a complexity of $O(n_{F-1}{}^2)$. These blocks have a size n_I which grows approximately with N^I. It is thus imperative to keep F small. The difference

in order of magnitude between repair and failure rates induces a strong locality (or near-decomposability) in the system matrix structure. This is the main factor which allows tight bounds on system availability to be obtained with small F values. For example, the bounds for F=4 were obtained within 6 CPU minutes on a SUN4 workstation. This value of F corresponds to only 12,951 states out of 16,777,216 states. That is, less than 0.1 percent of the total state space is being used by the analysis.

F	n_F	lower bound	upper bound	CPU(sec.)
4	10626	4.63e-08	6.51e-07	325.
5	42504	4.63e-08	8.02e-08	4965.
6	134596	4.63e-08	4.98 e-08	78280.

TABLE 3 (Unavailability bounds)

10 Conclusions

This new iterative computation method to bound conditional and marginal steady-state distributions in Markov chains dispenses from generating the whole state space, and is specially suited - and may even be indispensable - when the state space is too large to use classical procedures. The bounds are proven to be the tightest ones that can be obtained, given the part of the state space being considered. Locality or near-decomposability [Courtois 1977] are important factors contributing to tightness .

The method is limited by the computational requirements of matrix inversions. These requirements are however strongly mitigated when the diagonal blocks $\mathbf{Q}_{II}$ of the Markov chain are such that the inverses $(\mathbf{I}-\mathbf{Q}_{II})^{-1}$ can be economically computed. Fortunately, this is always the case in models of system availability.

Other computational complexity reductions are possible and have been mentioned. Further investigations are needed, however, as it seems that the whole bounding process could remain linear in the part of the state space which is considered.

11 References

[Courtois 1977] P.J. Courtois, *Decomposability: Queueing and Computer System Application*, Academic Press, New York, 1977.

[Courtois and Semal 1984] P.-J. Courtois and P. Semal, "Bounds for the Positive Eigenvectors of Nonnegative Matrices and for their Approximations by Decomposition", *Journal of the Assoc. of Comp. Mach.31,* 1984

[Courtois and Semal 1985] P.J. Courtois and P. Semal,"On Polyhedra of Perron-Frobenius Eigenvectors", *Linear Algebra and its Appl., 65,* 1985.

[Muntz et al. 1989] R.R. Muntz, E. De Souza Silva and A. Goyal, "Bounding Availability of Repairable Systems", *Performance Evaluation Review, 17, 1*, May 1989, 29-38.

[Semal 1992] P. Semal , *Analysis of Large Markov Models: Bounding Techniques and Applications*, Ph. Dissertation, Université Catholique de Louvain, Faculté des Sciences Appliquées Louvain-la-neuve, Belgium, 1992.

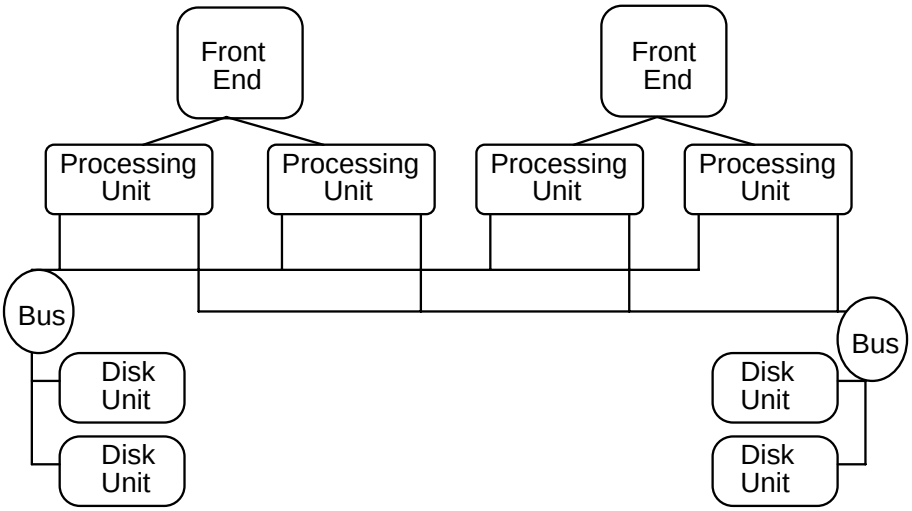

Figure 1. Database Architecture

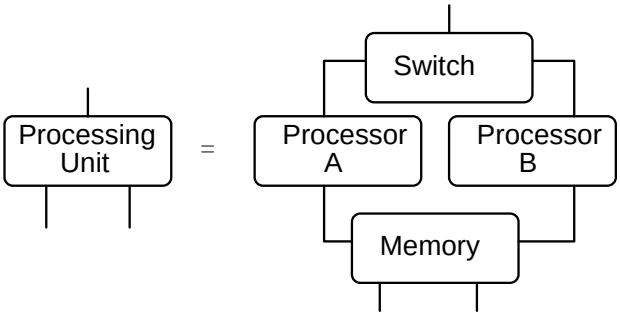

Figure 2. Processing Unit Structure