

The Nsolve File Formats

Peter Buchholz
Informatik IV, University of Dortmund
D-44227 Dortmund, Germany
Email: peter.buchholz@udo.edu

1 A Matrix Format for Structured Markovian Models

The following matrix formats are used to store the information to describe a Markov model in a hierarchical way such that the generator matrix is block structured and each block can be described as a sum of Kronecker products of small component matrices. With this format huge matrices with several billion of entries can be stored in a very compact form. The ideas underlying the format are can be found several papers. We recommend the following sources to obtain some informations about the background. The paper [8] describes two different variants of compact representations for the state spaces and matrices of Markov models, [3] includes an algorithm to generate a hierarchical model from a flat model consisting of interacting components, [4] describes the matrix format in the context of different numerical analysis methods and [2] is an early but comprehensive overview on structured Markovian modeling.

The structured matrix format is the input format for different analysis approaches which are partially public available [6]. In particular, the tool *Nsolve* which includes a wide variety of numerical solution techniques exploiting the compact matrix representation uses the matrix format.

2 File Formats

The section introduces the different file formats which are needed for a structured model. Apart from these files, a solver program usually requires an additional configuration file to set the parameters of the solution algorithms. These files are not presented here.

A hierarchical model consists of two levels. The upper level model *HLM* numbered with 0 describes the interaction between J low level models *LLM* numbered 1 through J . Hierarchical models may result from the specification of a queueing network [7] or a stochastic Petri net [1], or may be generated automatically from a network of interacting automata [2, 3].

The HLM is specified by the transition matrix (Sect. 2.1) and the HLM state space (Sec. 2.2). Each LLM is described by a matrix file (Sec. 2.3). Additionally we introduce the file format for vectors (Sec. 2.4). The matrix formats are derived from the sparse matrix format described in [5].

2.1 HLM Matrix File Format

The HLM matrix contains the visible transitions in the HLM including the weights/rates at this level. For consistency reasons the file contains information about a block structure of the matrix. However, usually the matrix consists of a single block.

```
Number of matrices in the file (integer) (will not be interpreted but should be 1)
Number of partition groups (one partition group is sufficient)
first state of partition group 0 (must be 0!)
...
last state of the last partition group (= Number of states of the HLM)
```

```
For each row of the HLM matrix:
Number of non-zero elements
```

```
For each non-zero element
destination state
transition number
```

weight

(non-zero elements need not be in the right order, but for each transition number and destination state only one element is allowed)

Transition numbers have to correspond to transition numbers used in the LLM matrix files.

2.2 HLM State Space File Format

The file includes the correspondence between a HLM state and the LLM states which form the global state. The number of HLM states and the number of LLMs are not included in the file. The number of HLM states is defined in the file with the HLM matrix and the number of LLMs can be found in the configuration file of the analysis programs. These files are read before the HLM state file.

For each HLM state

Macro state number for LLM $1, \dots, J$

(where J is the number of LLMs)

2.3 LLM Matrix File Format

A LLM is described by a number of matrices. Each matrix includes transitions according to one transition (synchronized) or a set of transitions (local). Transitions are numbered and number 0 is used for local transitions. Synchronized or externally visible transitions are numbered with values > 0 . The numbers of transitions in the LLM matrix files correspond to the transition numbers in the HLM matrix file. If for a synchronized transition in the HLM no matrix is available for a LLM, then it is assumed that the corresponding matrix is an identity matrix. Matrices in a LLM matrix file have to be ordered according to the transition number.

LLM matrices are partitioned into blocks which are numbered consecutively starting with 0. The block numbers correspond to the macro state numbers used in the HLM state file. Thus, if macro state number x appears in the HLM state for LLM j , then the matrix for LLM j has to contain at least $x + 1$ blocks (blocks are numbered from 0).

Number of matrices (one local and one for each non-identity synchronized transition)

Number of blocks

first state in the first block (must be 0!)

...

last state in the last block (= number of states of the LLM)

For each LLM Matrix

transition number

For each row of the matrix

Number of non-zero elements

For each non-zero element

destination state

weight

Observe that diagonal elements are not explicitly printed which is different from the sparse matrix format from [5]. If a LLM matrix contains a diagonal entry, then this value can be described like any other matrix entry by the destination state and the weight.

2.4 Vector Files

Two file formats for vectors exist which differ only by an initial integer. We denote the format with the initial integer as extended format. Most times the normal (non-extended) format is used. Values in vectors can be Integer, Real or Boolean (0/1).

Integer for classification (only for extended format)

Number of elements

Value of the first element

...

Value of the last element

3 Example

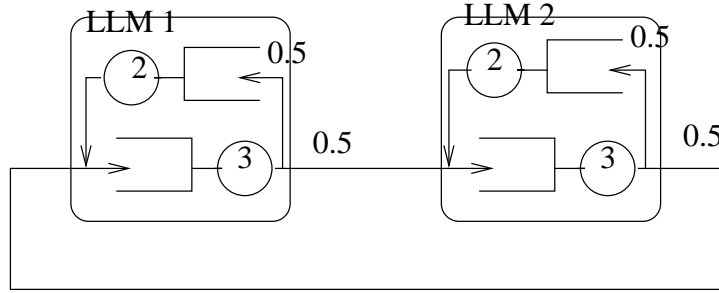


Figure 1: Queueing network example

We consider a simple queueing network example to clarify the concepts. The model is shown in figure 1. It consists of two identical LLMs each containing two exponential stations with service rates 2 and 3, respectively. We denote the lower station in an LLM, from where the LLM can be left, as station 1 and the other station as station 2. Furthermore, we assume that two customers are in the net. Of course, the simple net is a product form net and is lumpable such that a numerical analysis is not necessary.

The global state is defined by the distribution of customers among the LLMs 1 and 2. Thus, the HLM state space includes the three states (2, 0), (1, 1) and (0, 2). The following matrix defines the traveling of customers in the HLM.

$$\begin{pmatrix} - & (1, 1.0) & - \\ (2, 1.0) & - & (1, 1.0) \\ - & (2, 1.0) & - \end{pmatrix}$$

Each non-zero (or better non-empty) entry in the matrix is given by a transition number (number 1 indicates a departure from LLM 1 and an arrival at LLM 2, number 2 indicates a departure from LLM 2 and an arrival at LLM 1). Furthermore, a weight is attached to an entry. In the simple example all weights are 1.0 but arbitrary non-negative values are allowed as well.

Assume that the example is denoted as qn , then file $qn0.mat$ contains the transition matrix and looks as follows (row beginning with # are comments):

```
# number of matrices
1
# number of partition groups
1
# boundaries of the partition group(s)
0
3
# number of non-empty entries in row 0
1
# non-empty entry in the form destination, transition, weight
1 1 1.0
# number of non-empty entries in row 0
2
# non-empty entry in the form destination, transition, weight
0 2 1.0
2 1 1.0
# number of non-empty entries in row 0
1
# non-empty entry in the form destination, transition, weight
1 2 1.0
```

The file $qn.spa$ contains the state space description in the following form:

```
2 0
1 1
0 2
```

The LLM state spaces and matrices for the two submodels are similar. The state space contains 1 state if no customers is present, 2 states, if one customer is present and 3 states if two customers are present. Thus, for two customers in the system we have overall 6 LLM states. We choose the following order of the states $\{(0, 0), (1, 0), (0, 1), (2, 0), (1, 1), (0, 2)\}$ where the first component describes the population at the first station and the second component the population at the second station. LLM 1 is described by three matrices, one for internal transitions, one for departures and one for arrivals. The matrix for internal transitions looks as follows:

$$\left(\begin{array}{c|ccc|ccc} 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ \hline 0.0 & 0.0 & 1.5 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ \hline 0.0 & 0.0 & 0.0 & 0.0 & 1.5 & 0.0 \\ 0.0 & 0.0 & 0.0 & 2.0 & 0.0 & 1.5 \\ 0.0 & 0.0 & 0.0 & 0.0 & 2.0 & 0.0 \end{array} \right)$$

The departure matrix is given by:

$$\left(\begin{array}{c|ccc|ccc} 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ \hline 1.5 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ \hline 0.0 & 1.5 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.5 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \end{array} \right)$$

and the arrival matrix equals

$$\left(\begin{array}{c|ccc|ccc} 0.0 & 1.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ \hline 0.0 & 0.0 & 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 1.0 & 0.0 \\ \hline 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \end{array} \right)$$

Observe that the arrival rates to the LLM which are originated by the environment are set to 1 such that the entries in the matrix that describe arrivals define the routing probabilities upon arrival which are all 1 in our simple example.

The matrices are stored in the file *qn1.mat* which looks as follows:

```
# number of matrices
3
# number of partition groups
3
# boundaries of the partition
0
1
3
6
# Number of the first matrix (internal)
0
# number of non-zeros in row 0
0
# number of non-zeros in row 1
1
# non-zero entries in row 1
2 1.5
# number of non-zeros in row 2
1
# non-zero entries in row 2
1 2.0
# number of non-zeros in row 3
1
# non-zero entries in row 3
```

```

4 1.5
# number of non-zeros in row 4
2
# non-zero entries in row 4
3 2.0
5 1.5
# number of non-zeros in row 5
1
# non-zero entries in row 5
4 2.0
# Number of the second matrix (departure)
1
# number of non-zeros in row 0
0
# number of non-zeros in row 1
1
# non-zero entries in row 1
0 1.5
# number of non-zeros in row 2
0
# number of non-zeros in row 3
1
# non-zero entries in row 3
1 1.5
# number of non-zeros in row 4
1
# non-zero entries in row 4
2 1.5
# number of non-zeros in row 5
0
# Number of the third matrix (arrival)
2
# number of non-zeros in row 0
1
# non-zero entries in row 0
1 1.0
# number of non-zeros in row 1
1
# non-zero entries in row 1
3 1.0
# number of non-zeros in row 2
1
# non-zero entries in row 2
4 1.0
# number of non-zeros in row 3
0
# number of non-zeros in row 4
0
# number of non-zeros in row 5
0

```

The matrices for the second LLM are similar but the departure and arrival matrices are exchanged (i.e., the departure matrix obtains number 2 and the arrival matrix number 1).

References

- [1] P. Buchholz. A hierarchical view of GCSPNs and its impact on qualitative and quantitative analysis. *Journal of Parallel and Distributed Computing*, 15(3):207–224, 1992.

- [2] P. Buchholz. A framework for the hierarchical analysis of discrete event dynamic systems. Habilitationsschrift, Fachbereich Informatik, Universität Dortmund, 1996. available upon request.
- [3] P. Buchholz. Hierarchical structuring of superposed GSPNs. *IEEE Transactions on Software Engineering*, 25(2):166–181, 1999.
- [4] P. Buchholz. Structured analysis approaches for large Markov chains. *Applied Numerical Mathematics*, 31(4):375–404, 1999.
- [5] P. Buchholz. The sparse matrix file format. working paper, 10 2010.
- [6] P. Buchholz. Structured matrix market. http://www4.cs.uni-dortmund.de/Buchholz/struct_matrix_market.html, 2010.
- [7] P. Buchholz. A class of hierarchical queueing networks and their analysis. *Queueing Systems*, 15(1):1994, 59-80.
- [8] P. Buchholz and P. Kemper. Kronecker based matrix representations for large Markov chains. In B. Haverkort, H. Hermanns, and M. Siegle, editors, *Validation of Stochastic Systems*, volume 2925 of *LNCS*, pages 256–295, 2004.