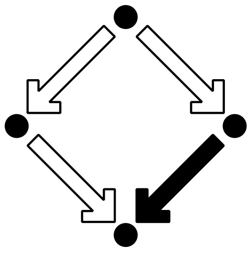

RISC-Linz
Research Institute for Symbolic Computation
Johannes Kepler University
A-4040 Linz, Austria, Europe

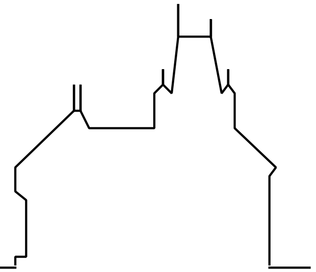

Comparing the Performance Modeling Environment MOSEL and the Probabilistic Model Checker PRISM for Modeling and Analyzing Retrial Queueing Systems

Tamás BÉRCZES, Gábor GUTA, Gábor KUSPER,
Wolfgang SCHREINER, János SZTRIK

(Month 2007)

RISC-Linz Report Series No. 07-17

Editors: RISC-Linz Faculty
B. Buchberger, R. Hemmecke, T. Jebelean, M. Kauers, T. Kutsia, G. Landsmann,
F. Lichtenberger, P. Paule, H. Rolletschek, J. Schicho, C. Schneider, W. Schreiner,
W. Windsteiger, F. Winkler.

Supported by: Austrian-Hungarian Scientific/Technical Cooperation Contract HU 13/2007

Copyright notice: Permission to copy is granted provided the title page is also copied.

Comparing the Performance Modeling Environment MOSEL and the Probabilistic Model Checker PRISM for Modeling and Analyzing Retrial Queueing Systems*

Tamás Bérczes†
tberczes@inf.unideb.hu

Gábor Guta‡
Gabor.Guta@risc.uni-linz.ac.at

Gábor Kusper§
gkusper@aries.ektf.hu

Wolfgang Schreiner‡
Wolfgang.Schreiner@risc.uni-linz.ac.at

János Sztrik†
jsztrik@inf.unideb.hu

December 6, 2007

†Faculty of Informatics, University of Debrecen, Hungary, http://www.inf.unideb.hu

‡Research Institute for Symbolic Computation (RISC), Johannes Kepler University, Linz, Austria, http://www.risc.uni-linz.ac.at

§Esterházy Károly College, Eger, Hungary, http://www.ektf.hu

*Supported by the Austrian-Hungarian Scientific/Technical Cooperation Contract HU 13/2007.

Abstract

We describe the results of analyzing the performance model of a retrial queueing system with the probabilistic model checker PRISM. The system has been previously analyzed with the help of the performance modeling environment MOSEL; we are able to accurately reproduce the results reported in literature. Furthermore, we compare PRISM and MOSEL with respect to their modeling languages and ways of specifying performance queries and benchmark the executions of the tools.

Contents

1 Introduction **4**

2 Problem Description **6**

2.1 Problem Overview 6
2.2 Finite State Model 7
2.3 Mathematical Model 10
2.4 Questions about the System 12
2.5 Different Versions of the System 12

3 Modeling and Analyzing in MOSEL **13**

4 Modeling and Analyzing in PRISM **18**

4.1 Translating the Model to PRISM 18
4.2 Terminals 19
4.3 Server 22
4.4 Rewards 24
4.5 Questions about the System in PRISM 27
4.6 Additional Queries about the PRISM Model 27

5 Experimental Results **28**

5.1 Analysis Results 29
5.2 Tool Benchmarks 29

6 Conclusions **37**

1 Introduction

The *performance analysis* of computing and communicating systems has always been an important subject of computer science. The goal of this analysis is to make predictions about the quantitative behavior of a system under varying conditions, e.g. the expected response time of a server under varying numbers of service requests, the average utilization of a communication channel under varying numbers of communication requests, and so on. Today's multitude of complex but nevertheless quite well functioning Internet-based service and communication infrastructures would not have been possible without an adequate performance analysis accompanying their design and implementation.

To perform such an analysis, however, first an adequate mathematical model of the system has to be developed which comprises the interesting aspects of the system but abstracts away from details that are irrelevant to the questions addressed. Originally, these models were developed purely by manual efforts, typically in formal frameworks based on queuing theory, stochastic Petri networks, and the like, which can be ultimately translated into continuous time Markov chains (CTMCs) as the fundamental mathematical basis [15]. Since the manual creation of complex models was tedious and error prone, then specification languages and corresponding tools were developed that automated the model creation from high-level system descriptions. Since the generated models can be typically not solved analytically, simulation-based techniques were applied in order to predict their quantitative behavior from a large number of sampled system runs. Later on, however, the underlying systems of equations were solved (for fixed parameter values) by iterative numerical calculations, thus deriving (mathematically exact but numerically approximated) solutions for the long-term (steady state) behavior of the system.

One tool of this kind is MOSEL (Modeling, Specification, and Evaluation Language) [11, 3] with its latest incarnation MOSEL-2 [12]. The software has a high-level specification language for modeling interconnected queue networks where transitions execute with certain rates to move entities across queues. The environment supports various backends for simulating the model system or for computing numerical solutions of the derived system of steady-state equations; in particular, it may construct a stochastic petri net model as input to the SPNP solver [8].

While above developments emerged in the *performance modeling and evaluation* community, also the *formal methods* community has recently produced theoretical frameworks and supporting tools that are, while coming from a different direction, nevertheless applicable to performance analysis problems. Originally, the only goal of formal methods was to determine qualitative properties of systems, i.e.

properties that can be expressed by formal specifications (typically in the language of temporal logic) as "yes/no" answers such as "Will the system never run into an invalid state?", "Will every request eventually receive an answer?", and so on. Questions of this kind can be effectively decided for finite state systems (of a certain size) by *model checkers* that investigate the underlying system models (essentially non-deterministic finite-state automata) with respect to the specified properties [5].

In the last couple of years, however, the formal methods community also got more and more interested in systems that exhibit stochastic behavior, i.e. systems whose transitions are executed according to specific rates (respectively probabilities); this gives rise to continuous time (respectively discrete time) Markov chains as those also used by the performance modeling community and to questions about quantitative rather than qualitative system properties. To pursue this new direction of *quantitative verification* [9], model checking techniques were correspondingly extended to *stochastic/probabilistic model checking* [10] and the the accompanying property specification languages were adapted to allow also the formulation of questions like "What is the probability that a certain situation will arrive during the system run?" or "What is the long term the expected value of a certain quantity?", i.e. questions that were originally the realm of the performance evaluation community. Nevertheless, classical qualitative properties can be still queried in this framework, e.g. asking "Is the probability 0 that an invalid state is reached?" amounts to asking "Will the system never run into an invalid state?".

A prominent tool in this category is the probabilistic model checker PRISM [13, 7] which provides a high-level modeling language for describing systems that exhibit probabilistic behavior, with models based on continuous-time Markov chains (CTMCs) as well as discrete-time Markov chains (DTMCs) and Markov decision procedures (MDPs). For specifying system properties, PRISM uses the probabilistic logics CSL (continuous stochastic logic) for CTMCs and PCTL (probabilistic computation tree logic) for DTMCs and MDPs, both logics being extensions of CTL (computation tree logic), a temporal logic that is used in various classical model checkers for specifying properties [5].

The fact that the previously disjoint areas of performance evaluation and formal methods have become overlapping is recognized by both communities. While originally only individual authors hailed this convergence [6], today various conferences and workshops are intended to make both communities more aware of each others' achievements [4, 18]. One attempt towards this goal is to compare techniques and tools from both communities by concrete application studies. The present paper is aimed at exactly this direction.

The starting point of our investigation is the paper [16] which discusses various

performance modeling tools; in particular, it presents the application of MOSEL to the modeling and analysis of a retrial queuing system previously described in [1] and later refined in [14]. The goal of the present paper is

- to construct PRISM models analogous to the MOSEL models presented in [16],
- to construct CSL queries in PRISM that compute the performance measures presented in above paper,
- to compare the results derived by PRISM with those from MOSEL,
- to evaluate the usability and expressiveness of both frameworks with respect to these tasks,
- to benchmark the tools with respect to their efficiency (time and memory consumption),
- and finally draw some overall conclusions about the suitability of PRISM to performance modeling compared with classical tools in this area.

The rest of the paper is structured as follows: Section 2 describes the application to be modelled and the questions to be asked about the model; Section 3 summarizes the previously presented MOSEL solution; Section 4 presents the newly developed PRISM solution; Section 5 gives the experimental results computed by PRISM in comparison to those computed by MOSEL and also gives benchmarks of both tools; Section 6 concludes and gives an outlook on further work.

2 Problem Description

2.1 Problem Overview

In this chapter we give a brief overview the model of the retrial queueing system presented in [16]. The variable names used later in the model are indicated in italics in the textual description. The model is represented by UML [17] class diagram (static view) and state machine diagrams (dynamic behavior).

The system contains a single server and terminals (*NT*) (Figure 1). Their behavior is as follows:

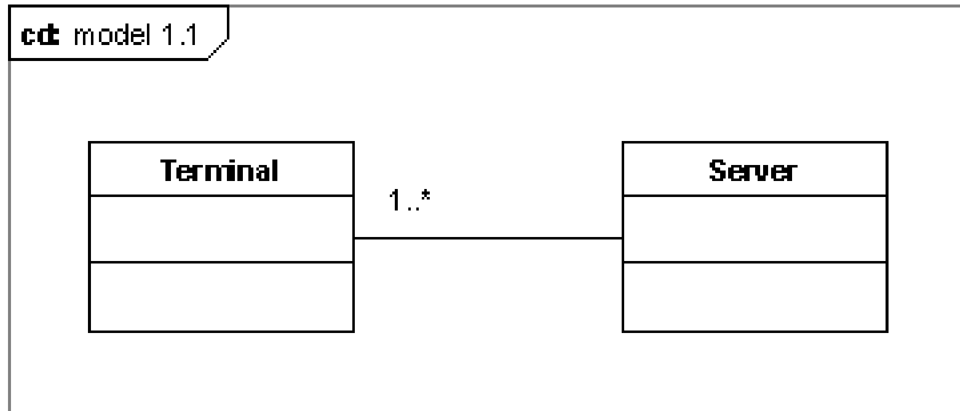

Figure 1: Components of the system

- Intuitively, terminals send requests to the server for processing. If the server is busy, the terminals retry to send the request later. More precisely, the terminals can be in three different states (which are named in parentheses):
  1. ready to generate a primary call (*busy*),
  2. sending repeated calls (*retrying*) and
  3. under service by the server (*waiting*).
- The server according to its CPU state (*cpu*) can be operational (*cpu=cpu_up*) or non-operational (*cpu=cpu_down*): if it is operational we distinguish between two further states (*cpu_state*): idle (*cpu_state=cpu_idle*) and busy (*cpu_state=cpu_busy*).
- In the initial state of the system, the server is operational (*cpu=cpu_up*) waiting for requests (*cpu_stat=cpu_idle*) and all terminals are ready to generate a primary call.

2.2 Finite State Model

The behavior of the system can be described by the state transitions of the terminals and the server, which occur at different rates.

We extend the standard UML [17] state machine diagram notation and semantic to present our model in an easy to read way. According to the standard, the diagram

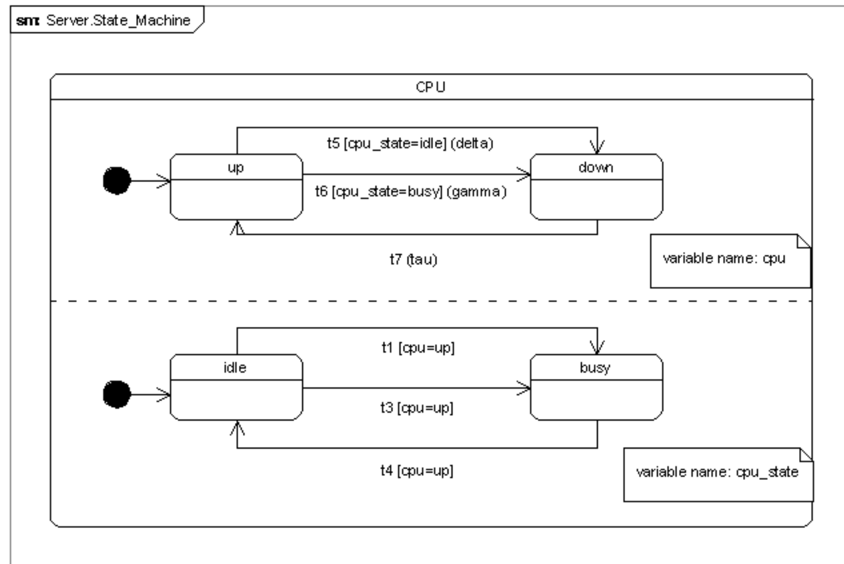

Figure 2: State machine representation of the server

contains states and transitions; the transitions in different swim-lanes can occur independently. Our extensions are the followings:

- Every comments of a swim-lane contains a variable name which is changed by the transition of that lane.
- Each transition is associated with a triple of a label, a guard (in square brackets) and a rate(in parentheses); if there is no rate indicated than the rate equals 1.
- A parallel composition semantics: the set of the states of the composed system is the Cartesian product of the state sets of the two swim-lanes or state machines. The composed state machine can make a transition whenever one of the original state machine can make one, except if multiple transitions in different original state machines have the same label: it that case, they must be taken simultaneously.

In Figure 2 we show the state transitions of the server:

t1 (The server starts to serve a primary call) If the server is in operational state and idle, it can receive a primary call and become busy.

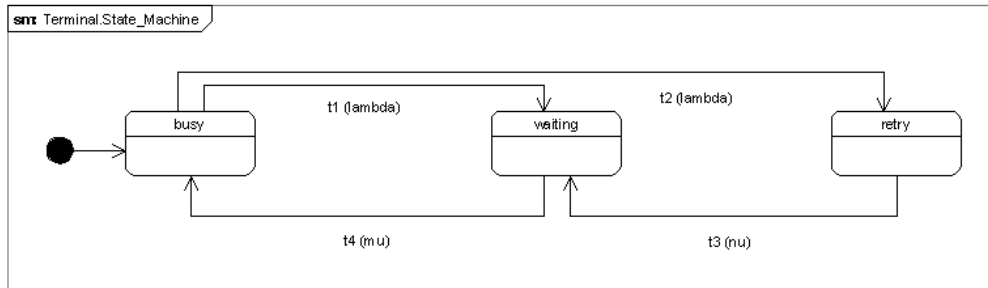

Figure 3: State machine representation of the terminals

t2 (The server rejects to serve a primary call) If the server is operational and busy, it can reject a primary call.

t3 (The server starts to serve a retried call) If the server is in operational state and idle, it can start to serve a repeated call.

t4 (The server finishes a call) If the server is operational and busy, it can finish a processing of the call.

t5 (An idle server becomes inoperable) If the server is in operational state and idle, it can become inoperable with rate δ.

t6 (A busy server becomes inoperable) If the server is in operational state and busy, it can become inoperable with rate γ.

t7 (A server gets repaired) If the server is inoperable, it can become operable again with rate τ.

The state transitions of the client are described in Figure 3:

t1 (The server starts to serve a primary call) The call of a terminal which issues a primary call is accepted and it becomes a waiting terminal with probability λ.

t2 (The server rejects to server a primary call) The call of a terminal which issues a primary call is rejected and it becomes a retrying terminal with probability λ.

t3 (The server starts to server a retried call) The call of a terminal which retries a call is accepted and it becomes a waiting terminal with probability ν.

t4 (The server finishes a call) The call of a terminal is finished and a it becomes ready to generate a new primary call again with rate μ.

The system can be represented alternatively by merging the server and the terminals into a single system as it was modelled in the original MOSEL model [16]: the guard conditions of all transitions with the same label are logically conjoined and their probability are multiplied.

2.3 Mathematical Model

In this section, we describe here the mathematical formulation of the queries. The state of the system at time *t* can be described by the process *X(t)=(cpu(t), cpu_state(t), waiting_terminals(t))*, where *cpu(t)=0* (*cpu_up*) if the server is operable, *cpu(t)=1* (*cpu_down*) if the server is not operable, *cpu_state(t)=0* (*cpu_idle*) if the server is idle and *cpu_state(t)=1* (*cpu_busy*) if the server is busy and *retrying_terminals(t)* describe the number of repeated calls at time *t*. The number of the waiting terminals and busy terminals are not expressed explicitly in the mathematical model. Their values can be calculated according the following:

- *waiting_terminals=0* if *cpu_state=cpu_idle*,
- *waiting_terminals=1* if *cpu_state=cpu_busy*,
- *busy_terminals=NT-(waiting_terminals+retrying_terminals)*,

Because of the exponentiality of the involved random variables and the finite number of sources, this process is a Markov chain with a finite state space. Since the state space of the process $X(t),\ t \geq 0$ is finite, the process is ergodic for all reasonable values of the rates involved in the model construction. From now on, we assume that the system is in the steady-state.

We define the stationary probabilities by:

$$P(q,r,j) = \lim_{t\to\infty} P(cpu(t), cpu_state(t), retrying_terminals(t)),$$
$$q = 0,1, r = 0,1, j = 0,\cdots,NT-1,$$

The main steady-state system performance measures can be derived as follows:

- *Utilization of the servers*

$$cpuutil = \sum_{j=0}^{NT-1} P(0,1,j)$$

- *Availability of the servers*

$$goodcpu = \sum_{r=0}^{1} \sum_{j=0}^{NT-1} P(0,r,j)$$

- *Utilization of the repairman*

$$repairutil = \sum_{r=0}^{1} \sum_{j=0}^{NT-1} P(1,r,j) = 1 - goodcpu$$

- *Mean rate of generation of primary calls*

$$busyterm = E[NT - cpu_state(t) - retrying_terminals(t); cpu(t) = 0]$$
$$= \sum_{r=0}^{1} \sum_{j=0}^{NT-1} (NT - r - j) P(0,r,j)$$

- *Utilization of the sources*

$$termutil = \frac{busyterm}{NT}$$

- *Mean rate of generation of repeated calls*

$$retravg = E[retrying_terminals(t); cpu(t) = 0] = \sum_{r=0}^{1} \sum_{j=0}^{NT-1} jP(0,r,j)$$

- *Mean number of calls staying in the server*

$$waitall = E[cpu_state(t)] = \sum_{q=0}^{1} \sum_{j=0}^{NT-1} P(q,1,j)$$

- *Mean number of calls staying in the orbit*

$$retrall = E[retrying_terminals(t)] = \sum_{q=0}^{1} \sum_{r=0}^{1} \sum_{j=0}^{NT-1} jP(q,r,j)$$

- *Overall utilization*

$$overallutil = cpuutil + repairutil + NT * termutil$$

- *Mean number of calls staying in the orbit or in the server*

$$meanorbit = waitall + retrall$$

- *Mean response times*

$$E[T] = \frac{E[retrying_terminals(t)] + E[cpu_state(t)]}{\lambda * busyterm}$$

2.4 Questions about the System

Our goal is to study various quantitative properties of the presented models to get deeper understanding of the modelled systems. The following properties are analyzed:

cpuutil The ratio of the time the server spends with serving calls compared to the total execution time (0≤cpuutil≤1).

goodcpu The ratio of the time when the server is operable compared to the total execution time (0≤goodcpu≤1).

repairutil The ratio of the time when the server is inoperable compared to the total execution time (0≤repairutil≤1).

busyterm The average number of served terminals while the system is operable (0≤busyterm≤NT).

termutil The ratio of served terminals while the system is operable to the total number of the terminals (0≤termutil≤1).

retravg The average number of retrying terminals while the system is operable (0≤retravg≤NT-1).

waitall The average number of the waiting terminals during the total system execution time (0≤waitall≤1).

retrall The average number of the retrying terminals during the total system execution time (0≤retrall≤NT-1).

overallutil The sum of the system average utilization, i.e. the sum of *cpuutil*, *repairutil* and *NT*termutil* (0≤overallutil≤NT+1).

meanorbit The average number of the retrying terminals and waiting terminals during the total system execution time (0≤retrall≤NT).

resptime The mean response time, i.e: the average waiting time till a call of a terminal is successfully accepted.

2.5 Different Versions of the System

In [16], actually four slightly different systems were described:

continuous The presented model.

non-continuous If the server becomes inoperable, than the call have to be retried (the waiting terminal becomes retrying).

continuous, intelligent It can also reject a call if the server is inoperable (the original model cannot handle a call if the server is inoperable.

non-continuous, intelligent The combination of the non-continuous and intelligent model.

The later three variants are not formally described in the present paper. However, they have been implemented and have been used for the experiments in Chapter 5.

3 Modeling and Analyzing in MOSEL

The MOSEL language (Modeling Specification and Evaluation Language) was developed at the University of Erlangen. The MOSEL system uses a macro-like language to model communication networks and computer systems, like stochastic Petri Nets. The MOSEL tool contains some language features, like variables and functions in the style of the C programming language. The MOSEL calls an external tool after having translated the MOSEL code into the respective tool's format. For example the Petri Net analysis tool SPNP and the state analysis tool MOSES can be used.

In this chapter we demonstrated how we translated the model described in the problem section into a MOSEL model.

Now we present the full source code of the MOSEL model (in verbatim) surrounded by detailed comments.

The MOSEL programs consist of four parts: the declarations, the node definitions, the transition rules and the results.

The MOSEL source code begins with the declarations. In the model we use the same notations and parameters as in the section PRISM.

- `NT` is the number of the terminals,
- λ is the rate of primary call generation,
- μ is the rate of the call servicing,
- ν is the rate of repeated call generation,

- δ is the failure rate in idle state of the CPU,
- γ is the failure rate in busy state of the CPU and
- τ is the repair rate of the CPU.

```
#define NT 6
VAR double lambda;
VAR double nu;
VAR double mu;
VAR double delta;
VAR double gamma;
VAR double tau;
```

The CPU can be in two states: in `idle` or in `busy` state, and it can be in operational (`up`) or non-operational (`failed`) in both states. We define the following enumerations to describe these cases:

- `cpu_states` represents the state of the CPU. It can be in two states: idle and busy.
- `cpu_updown` indicate whether the CPU is operational or not.

```
enum cpu_states {cpu_busy, cpu_idle};
enum cpu_updown {cpu_up, cpu_down};
```

The node part defines the nodes of the system. In our model there are five nodes: one node for the number of busy, retrying, and waiting terminals, respectively, and two nodes for the CPU. The CPU is idle and up and all the terminals are busy at the starting time.

Each terminal can be in three states: generating a primary call, sending repeated calls, and under service. These states are represented by the following variables:

- `busy_terminals` is the number of terminals, which are capable to generate primary calls (they are busy with the local task and may generate calls to the server;
- `retrying_terminals` is the number of retrying terminals. If a generated call finds the CPU in busy state, than the terminal moves to state retrying;

- `waiting_terminals` is the number of the waiting terminals. If a call finds the CPU in idle state, than the terminal moves to state waiting, which issued a successful call to the server and wait for the answer of the call.
- `cpu_state` represent the state of the CPU. It can be `idle` and `busy`.
- `cpu` indicate whether the CPU is operational (`up`) or not (`failed`).

If `cpu_state` is equal to `cpu_busy` it means that that the operational CPU serving of a call and when the `cpu_state` is equal to `cpu_idle` than the CPU waiting a call.

In our model, server contains only one CPU, so the number of the waiting terminals never be more than 1. Initially all terminals are busy terminals.

```
NODE busy_terminals[NT] = NT;
NODE retrying_terminals[NT] = 0;
NODE waiting_terminals[1] = 0;
NODE cpu_state[cpu_states] = cpu_idle;
NODE cpu[cpu_updown] = cpu_up;
```

The transition part describes how the system works.
The first transition rule defines the successful primary call generation: It the CPU is operational then the CPU moves from the `idle` state to `busy` and the terminal from busy to waiting. All busy terminal produce that call with rate λ, so the rate is λ multiplied with the number of busy terminals. After that transition the number of busy terminals decreases by one and the number of the waiting terminals increases by one. Note that the MOSEL generates automatic conditions, so that transition rule will occur when `cpu_state` is `cpu_idle` and number of `busy_terminals` are greater than 0.

```
IF cpu==cpu_up FROM cpu_idle, busy_terminals
TO cpu_busy, waiting_terminals W lambda*busy_terminals;
```

The second rule describes an unsuccessful primary call. It occurs if the CPU is operational and the CPU is in `busy` state. After that transition the busy terminal moves to state retrying. This means, the number of busy terminals decreases by one and the number of the retrying terminals increases by one. All busy terminals produce that call with rate λ, so the rate is λ multiplied by the number of busy terminals.

```
IF cpu==cpu_up AND cpu_state==cpu_busy
FROM busy_terminals
TO retrying_terminals W lambda*busy_terminals;
```

The third rule describes the scenario of a successful repeated call generation. It occurs if the CPU is operational and the CPU is in idle state and there are some retrying terminals. All retrying terminals produce the calls with ν rate, so the rate is ν multiplied by the number of busy terminals. After that transition the CPU moves from the idle state to busy and the terminal from retrying to waiting. That means, the number of retrying terminals decreases by one and the number of the waiting terminals increases by one.

```
IF cpu==cpu_up FROM cpu_idle, retrying_terminals
TO cpu_busy, waiting_terminals W nu*retrying_terminals;
```

The fourth rule describes the request service at the CPU. It occurs if the CPU is operational and the CPU is in busy state. Of course that means, the number of waiting terminals is one. Its rate is determined by the server call serving rate. After that transition the CPU moves from the busy state to idle and the terminal from waiting to busy. So, the number of waiting terminals decreases by one (will be 0) and the number of the busy terminals increases by one.

```
IF cpu==cpu_up FROM cpu_busy, waiting_terminals
TO cpu_idle, busy_terminals W mu;
```

The fifth transition describes the scenario when an idle server becomes failed. It occurs with rate δ if the server is up and idle. If a server becomes failed, it keeps its state.

```
IF cpu_state==cpu_idle FROM cpu_up TO cpu_down W delta;
```

The sixth transition describes the scenario when a busy server becomes failed. It occurs if the server is up and busy with rate δ.

```
IF cpu_state==cpu_busy FROM cpu_up TO cpu_down W gamma;
```

The last transition describes the scenario when a failed server will be repaired. It occurs with rate τ if the server is failed. After it gets repaired, it continues the processing if it was busy.

```
FROM cpu_down TO cpu_up W tau;
```

The last section of a model description is the result part where the calculations of the requested output performance measures are. The mathematical background of the equations used in the result part is described in the previous section.

In the result part of the MOSEL code we used that equations to obtain the performance measures of the system. In MOSEL we can calculate the probability of a certain state or combinations of certain states by the PROB keyword.

The first result is defined the CPU utilization (*cpuutil*).

```
RESULT>> if(cpu==cpu_up AND cpu_state==cpu_busy)
  cpuutil += PROB;
```

The *goodcpu* is the probability of the state where the CPU operational.

```
RESULT>> if(cpu==cpu_up) goodcpu += PROB;
```

The *busyterm* is the mean number of the busy terminals to all states where the server is operational.

```
RESULT if(cpu==cpu_up) busyterm += (PROB*busy_terminals);
```

The *termutil* assigns the ratio of the mean number of busy terminals (*busyterm*) over the total number of terminals (*NT*).

```
RESULT>> termutil = busyterm / NT;
```

The *waitall* is the mean number of the calls staying in the service.

```
RESULT if(waiting_terminals>0)
  waitall += (PROB*waiting_terminals);
```

The *retrall* is the mean number of the calls staying in the orbit.

```
RESULT if(retrying_terminals>0)
  retrall += (PROB*retrying_terminals);
```

The *resptime* is the mean response time.

```
RESULT>> resptime = (retrall + waitall) / NT /
  / (lambda * termutil);
```

Then the last result give us the overall utilization (*overallutil*).

```
RESULT>> overallutil = cpuutil + busyterm;
```

4 Modeling and Analyzing in PRISM

In this chapter we describe how we translate the model described in Section 2 into a PRISM model. Further information about the PRISM system can be found under [13]. In the first subsection we show the source-code of the PRISM model; in the second subsection we formulate questions in the model. This section should give a good overview respect to the capability of the PRISM system for modeling and analysis a queueing system.

4.1 Translating the Model to PRISM

In this subsection and the following ones, we present the full source code of the PRISM model (in `verbatim`) surrounded by detailed comments. The model description has 4 main parts:

- the type of the model,
- the constant declarations,
- the module declarations and
- the reward specifications.

In our case, all models are represented in Continuous-time Markov chains model, which is indicated by the keyword `stochastic`.

```
stochastic
```

Constants can be used in two manners:

- uninitialized constants denote parameters of the model and,
- initialized constants denote fixed values.

```
const int NT;

const double lambda;
const double mu;
const double nu;

const double delta;
const double gamma = delta;
const double tau;
```

The parameters of that model are the following constants:

- `NT` denotes the number of the terminals *NT*,
- `lambda` denotes the rate λ of primary call generation,
- `mu` denotes the rate μ of the call servicing,
- `nu` denotes the rate ν of repeated call generation,
- `delta` denotes the failure rate δ in idle state of the server,
- `gamma` denotes the failure rate γ in busy state of the server and
- `tau` denotes the repair rate τ of the server.

In our simulation we do not distinguish between the failure rate in idle and busy state, so we equal `gamma` with `delta`.

We define constants to represent the state of the server to make the model human-readable. There are two pairs of constants, the two pairs used to assign value for two different variables, namely `cpu` and `cpu_state`. The first pair of the constants `cpu_up` and `cpu_down` indicates whether the server is operable or not. The second pair of constants `cpu_busy` and `cpu_idle` expresses that the server is busy serving a call and that the server is idle waiting for a call, respectively.

```
const int cpu_up = 0;
const int cpu_down = 1;
const int cpu_busy = 0;
const int cpu_idle = 1;
```

The next fragment are the module definitions. A module definition is started with the `module` keyword and is closed with the `endmodule` keyword. All modules contain state variables and state transitions. We have two modules `TERMINALS` and `SERVER` described in the following subsections,

4.2 Terminals

The module `TERMINALS` represents the set of the terminals. We keep track of the number of the terminals in specific states, because in PRISM it is not possible to have multiple instances of a module. Thus all variables range from 0 to the maximal number of the terminals, which is denoted by the range indicator within square brackets in the source code.

```
module TERMINALS

busy_terminals : [0..NT] init NT;
retrying_terminals : [0..NT] init 0;
waiting_terminals : [0..NT] init 0;
```

We have the following variables in the model :

- `busy_terminals` is the number of terminals, which are capable to generate primary calls (they are busy with local tasks and may generate calls to the server);
- `retrying_terminals` is the number of retrying terminals, i.e. terminals which have generated an unsuccessful call and are retrying the same call;
- `waiting_terminals` is the number of the waiting terminals, i.e. terminals which have issued a successful call to the server and wait for the answer of the call.

In the current model, we have only one server, therefore the number of the waiting terminals never be more than 1. Initially all terminals are busy terminals.

The transitions are represented in form `[l] g -> r : u`. The transition with label `l` occurs if the guard `g` evaluates to true; the rate of the transition is `r`, the values of the state variables are updated according to `u`. The labels serve as synchronization identifiers for parallel composition. Transitions with the same label in different modules execute together, i.e. all guards of the transition must be true and the total transition rate is the product of the individual transition rates. We also have to notice that the transitions of the terminals have their counterparts on the server side, which make the transition guards unique.

The transition with label `t1` describes the scenario of a successful primary call:

```
[t1]
  busy_terminals > 0 & waiting_terminals < NT ->
    lambda*busy_terminals :
  (busy_terminals' = busy_terminals-1) &
  (waiting_terminals' = waiting_terminals+1);
```

The transition occurs if there are some busy terminals and the number of the waiting terminals is lower than the number of the terminals. The second part of the guard condition is purely technical to explicitly state that the value of

`waiting_terminals` is not greater than the maximally allowed value. (According to the model semantics we know that it never becomes greater than one, because the server serves only one call at once.) All busy terminals produce that call with rate λ, so the rate is λ multiplied with the number of busy terminals. After that transition, the number of busy terminals decreases by one and the number of the busy terminals increases by one.

The transition with label `t2` describes the scenario of an unsuccessful primary call:

```
[t2]
  busy_terminals > 0 & retrying_terminals < NT ->
    lambda*busy_terminals :
  (busy_terminals' = busy_terminals-1) &
  (retrying_terminals' = retrying_terminals+1);
```

The transition occurs if there are some busy terminals and the number of the retrying terminals is lower than the number of the terminals. The second part of the guard condition is also purely technical to explicitly state that the value of `waiting_terminals` is not greater than the maximally allowed value. (According the model semantics we know that it never becomes grater than maximal number, because the sum of the terminal variables equals the number of the terminals.) All busy terminals produce that call with rate λ, so the rate is λ multiplied by the number of busy terminals. After that transition the number of busy terminals decreases by one and the number of the busy terminal increases by one.

The transition with label `t3` describes the scenario of a successfully repeated call:

```
[t3]
  retrying_terminals > 0 & waiting_terminals < NT ->
    nu*retrying_terminals :
  (retrying_terminals' = retrying_terminals-1) &
  (waiting_terminals' = waiting_terminals+1);
```

The transition occurs if there are some retrying terminals and the number of the waiting terminals is smaller than the number of the terminals. All retrying terminals produce the calls with rate ν, so the rate is ν multiplied with the number of busy terminals. After that transition, the number of retrying terminals decreases by one and the number of the waiting terminals increases by one.

The transition with label `t4` describes the scenario of an answer for a waiting terminal:

```
[t4]
  waiting_terminals > 0 & busy_terminals < NT ->
    1 :
  (waiting_terminals' = waiting_terminals-1) &
  (busy_terminals' = busy_terminals+1);
```

The transition occurs if there are some waiting terminals and the number of busy terminals smaller than the number of the terminals. Its rate is determined by the call serving rate on the server side (see below). After that transition, the number of retrying terminals decreases by one and the number of the waiting terminals increases by one.

```
endmodule
```

4.3 Server

The second module represents the server by two binary state variables. The variable `cpu` expresses the operability of the server by the values 0 and 1, which are denoted by the constants `cpu_up` and `cpu_down`, respectively. The variable `cpu_state` the state of the server by values 0 and 1, which are denoted by the constants `cpu_busy` and `cpu_idle`, respectively.

```
module SERVER

cpu : [cpu_up..cpu_down] init cpu_up;
cpu_state : [cpu_busy..cpu_idle] init cpu_idle;
```

The transition with label `t1` describes the server side scenario of a successful primary call. It occurs, if the server is operable and idle. After the transition, the server becomes busy.

```
[t1]
  cpu = cpu_up & cpu_state = cpu_idle ->
    1 :
  (cpu_state' = cpu_busy);
```

The transition with label `t2` describes the server side scenario of an unsuccessful primary call. It occurs, if the server is operable and busy. After the transition, the state of the server doesn't change.

```
[t2]
  cpu = cpu_up & cpu_state = cpu_busy ->
    1 :
  (cpu' = cpu) &
  (cpu_state' = cpu_state);
```

The transition with label `t3` describes the server side scenario of a successful primary call. It is the same as the transition `t1`, because the server can't distinguish between a primary and a repeated call.

```
[t3]
  cpu = cpu_up & cpu_state = cpu_idle ->
    1 :
  (cpu_state' = cpu_busy);
```

The transition with label `t4` describes the server side scenario of finishing a call (a successful call served). It occurs with rate μ and the server becomes idle after the transition.

```
[t4]
  cpu = cpu_up & cpu_state = cpu_busy &
  mu > 0 ->
    mu :
  (cpu_state' = cpu_idle);
```

The transition with label `t5` describes the scenario when an idle server becomes inoperable. It occurs, if the server is operable and idle with rate γ. If a server becomes inoperable, it keeps its state. After it gets repaired, it continues the processing, if it was busy at the time of the failure.

```
[t5]
  cpu_state = cpu_idle & cpu = cpu_up &
  delta > 0 ->
    delta :
  (cpu' = cpu_down);
```

The transition with label `t6` describes the scenario when a busy server becomes inoperable. It occurs, if the server is operable and busy with rate δ.

```
[t6]
  cpu_state = cpu_busy & cpu = cpu_up &
  gamma > 0 ->
    gamma :
  (cpu' = cpu_down);
```

The transition with label `t7` describes the scenario when a server gets repaired. It occurs, if the server is inoperable with rate τ.

```
[t7]
  cpu = cpu_down &
  tau > 0 ->
    tau :
  (cpu' = cpu_up);

endmodule
```

4.4 Rewards

The last section of a model description is the declaration of rewards. Rewards are numerical values assigned to states or to transitions. Arbitrary many reward structures can be defined over the model and they can referenced by a label. We use rewards to define the various question defined in Section 2.4.

The first reward is the server utilization (*cpuutil*). It assigns a value 1 to all states where the server is operable and busy.

```
rewards "cpuutil"
  cpu = cpu_up & cpu_state = cpu_busy : 1;
endrewards
```

The reward *goodcpu* assigns 1 to all states where the server is operable.

```
rewards "goodcpu"
  cpu = cpu_up : 1;
endrewards
```

The reward *repairutil* assigns 1 to all states where the server is inoperable.

```
rewards "repairutil"
  cpu = cpu_down : 1;
endrewards
```

The reward *busyterm* assigns the number of the busy terminals to all states where the server is operable.

```
rewards "busyterm"
  cpu = cpu_up : busy_terminals;
endrewards
```

The reward *termutil* assigns the ratio of the busy terminals over the total number of terminals to all states where the server is operable.

```
rewards "termutil"
  cpu = cpu_up : busy_terminals/NT;
endrewards
```

The reward *retravg* assigns the number of the retrying terminals to all states where the server is operable.

```
rewards "retravg"
  cpu = cpu_up : retrying_terminals;
endrewards
```

The reward *waitall* assigns the number of the waiting terminals to states with such terminals.

```
rewards "waitall"
  waiting_terminals > 0 : waiting_terminals;
endrewards
```

The reward *waitall* assigns the number of the retrying terminals to states with such terminals.

```
rewards "retrall"
  retrying_terminals > 0 : retrying_terminals;
endrewards
```

The reward *meanoribt* assigns the number of the retrying and waiting terminals to states with such terminals.

```
rewards "meanorbit"
  retrying_terminals > 0 : retrying_terminals;
  waiting_terminals > 0 : waiting_terminals;
endrewards
```

The reward *resptime_comp* assigns a number needed for the mean response time calculation to the states. It doesn't have any intuitive meaning and the reason to calculate is simply technical.

```
rewards "resptime_comp"
  retrying_terminals > 0 : retrying_terminals/(NT*lambda);
  waiting_terminals > 0 : waiting_terminals/(NT*lambda);
endrewards
```

The reward *overallutil* assigns to the all states the total number of all busy elements, i.e. the server, if it is busy or is under repair (a repair unit is busy with its repair), and all busy terminals.

```
rewards "overallutil"
  cpu = cpu_up & cpu_state = cpu_busy : 1;
  cpu = cpu_down : 1 ;
  cpu = cpu_up : busy_terminals;
endrewards
```

4.5 Questions about the System in PRISM

As we mentioned in the introduction, in PRISM the queries about the CTMC models can be formulated in CSL (Continuous Stochastic Logic). CSL is a branching-time logic similar to CTL or PCTL [2]. It is capable to express queries about both transient and steady-state properties. Transient properties refers to the values of the rewards at certain times and the steady-state properties refer to long-run rewards.

The PRISM system support not only evaluating predicates about the rewards, but also queries about the rewards. In our experiments we used only the following one CSL construction: `R{"l"}=? [ S ]`. This query ask for the expected long-run reward of the structure labelled with l. Most questions about the model described in Section 2.4 can be formulated as CSL expressions.

```
R{"cpuutil"}=? [ S ]
R{"goodcpu"}=? [ S ]
R{"repairuti"}=? [ S ]
R{"busyterm"}=? [ S ]
R{"termutil"}=? [ S ]
R{"retravg"}=? [ S ]
R{"waitall"}=? [ S ]
R{"retrall"}=? [ S ]
R{"overallutil"}=? [ S ]
R{"meanorbit"}=? [ S ]
R{"resptime_comp"}=? [ S ]
```

The response time (`resptime`) cannot be expressed as a single query, because the cross-referencing between the reward structures not allowed. As a workaround we calculated `resptime` during the post processing of the results according the relationship `resptime=resptime_comp/termutil`.

4.6 Additional Queries about the PRISM Model

As we mentioned in the previous subsection, the CSL language in PRISM is capable to formulate more sophisticated questions. For example:

> ”What is the probability that the server is busy at time T?”.

The question below can be formulated as:

```
P=? [true U[T,T]
    (cpu=cpu_up & cpu_state=cpu_busy)].
```

We can also ask standard qualitative question about the integrity of the model like:

> ”Is it always true that if the server is busy, then there is one waiting terminal?”

This question can be formulated in PRISM as:

```
(cpu_state=cpu_busy)=>(waiting_terminals=1).
```

The model checker will determine whether this property is true in all states of the system; it is thus able to check classical safety properties. One may also formulate a classical liveness property such as

> “Will the cpu always become busy again?”

by the query

```
P>=1[true U cpu_state=cpu_busy]
```

which asks whether in every state with probability 1 (`true` holds until) eventually the property `cpu_state=cpu_busy` holds.

5 Experimental Results

In this section, we show the result of the experiments carried through with PRISM. The parameters used for the experiments are listed in Figure 5; they are the same as published in [16]. The results of the experiments with PRISM are presented in diagrams (Figure 6, 7, 8, 9, 10, 11) and in tables (Figure 12, 13, 14, 15, 16, 17).

The experiments was performed in two main steps: the execution of the experiments through the GUI of PRISM and the post-processing of the results. We selected the appropriate CSL query according the Figure 4 and set up the parameters according the Figure 5; after the execution of PRISM the results were exported to CSV files for further processing. The post processing happened with a help of Python scripts.

Nr. of the experiment	used reward(s)
1	resptime_comp and termutil
2	overallutil
3	meanorbit
4	resptime_comp and termutil
5	overallutil
6	meanorbit

Figure 4: Rewards calculated in the experiments

Exp. Nr.	NT	λ	μ	ν	γ/δ	τ	X axis
1	6	0.8	4	0.5	X axis	0.1	0. 0.01. ..., 0.12
2	6	0.1	0.5	0.5	X axis	0.1	0. 0.01. ..., 0.12
3	6	0.1	0.5	0.05	X axis	0.1	0. 0.01. ..., 0.12
4	6	0.8	4	0.5	0.05	X axis	0.5. 1.0. ..., 4.0
5	6	0.05	0.3	0.2	0.05	X axis	0.5. 1.0. ..., 4.0
6	6	0.1	0.5	0.05	0.05	X axis	0.5. 1.0. ..., 4.0

Figure 5: Parameters of the experiments

5.1 Analysis Results

The diagrams compared with the ones presented in [16] clearly show that the two models (MOSEL and PRISM) produce identical results for the same parameters. Comparing the raw results of the experiments, it shows that they are differ only after the 5th decimal digit. The quality of the results produced with PRISM is this the same as the ones produced in MOSEL.

5.2 Tool Benchmarks

A benchmark was carried through to compare the efficiency of the two tools. The parameters of the machine that was used for the benchmark: P4 2.6GHz with 512KB L2 Cache and 512MB of main memory. Unfortunately MOSEL is not capable to handle models where the number of terminals (*NT*) is greater than 126, such that the runtime of the benchmarks (which in PRISM especially depend on *NT*) remain rather small.

Both of the tools were tested with the described model using the following parameters: λ =0.05, μ =0.3, ν =0.2, $\gamma = \delta$ =0.05, τ =0.1. The comparison of

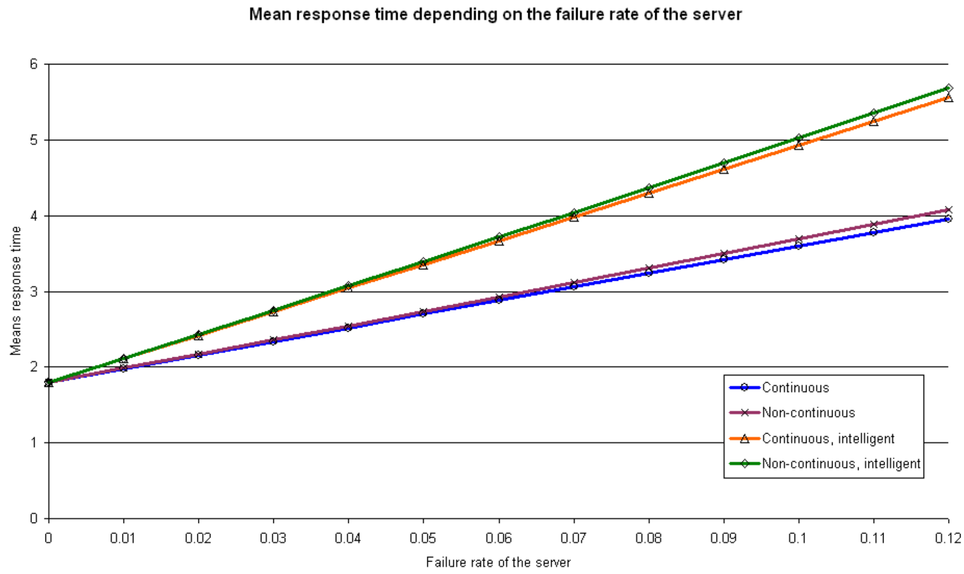

Figure 6: Results of the 1st experiment

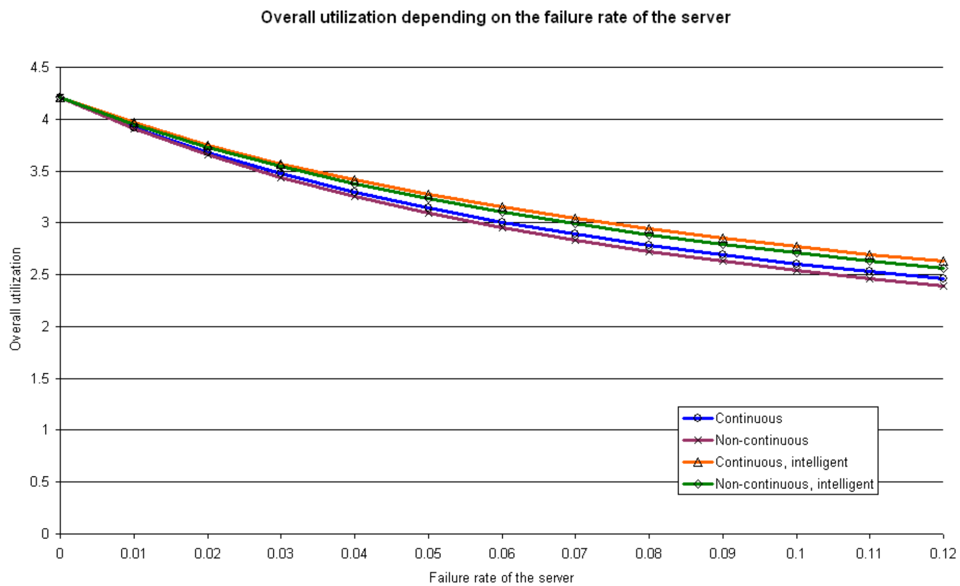

Figure 7: Results of the 2nd experiment

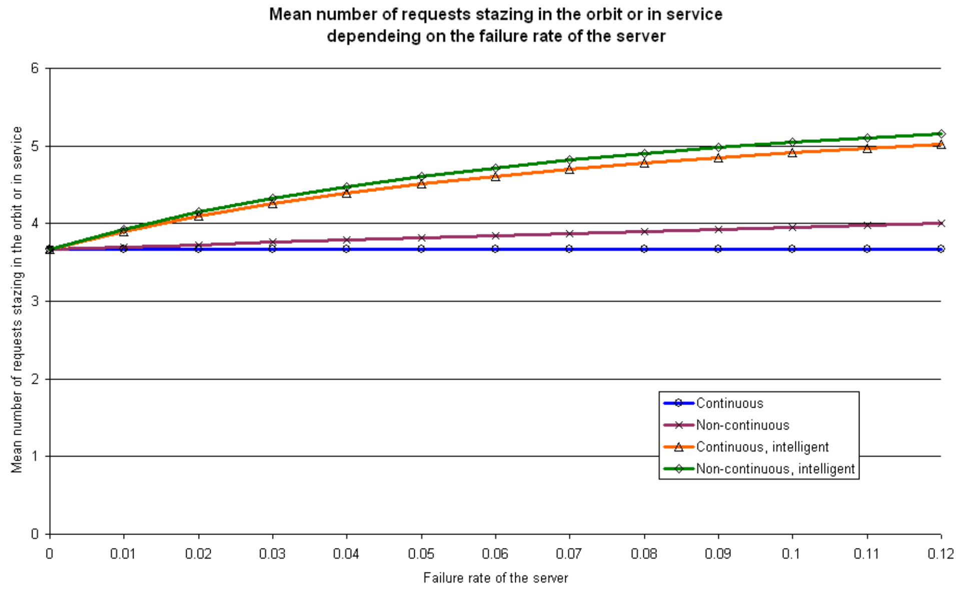

Figure 8: Results of the 3rd experiment

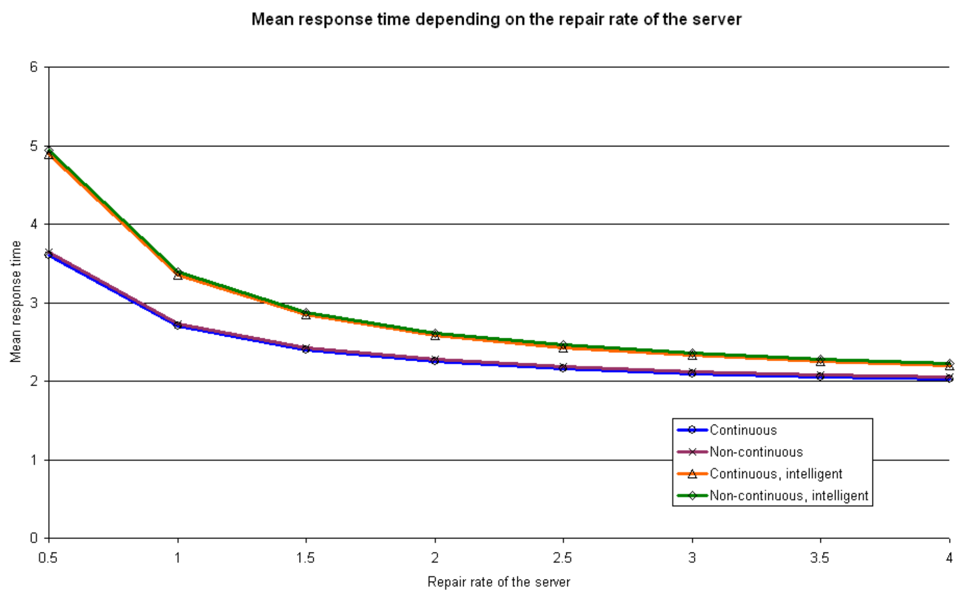

Figure 9: Results of the 4th experiment

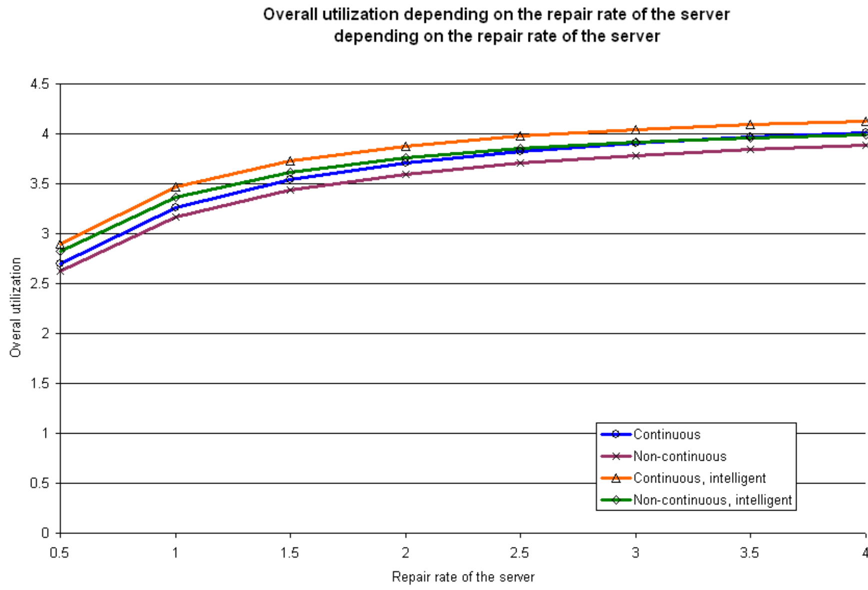

Figure 10: Results of the 5th experiment

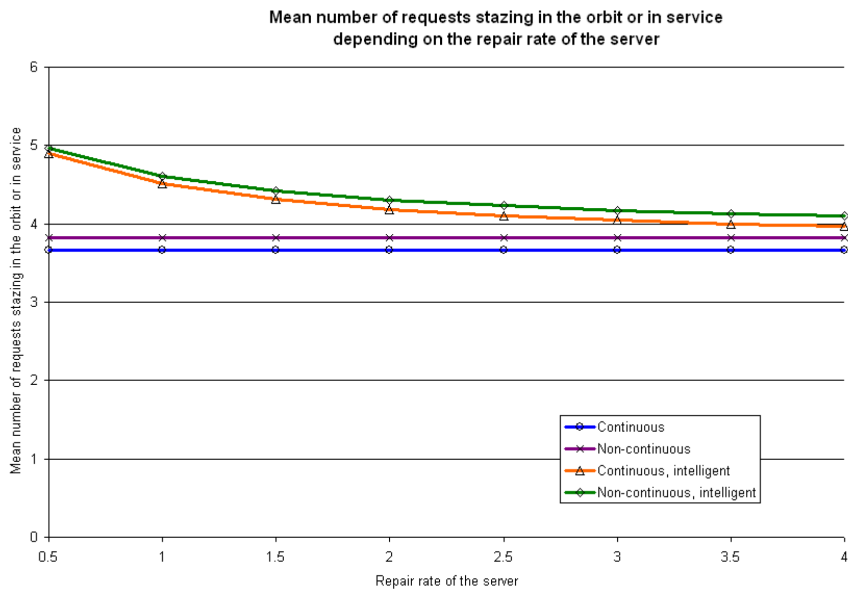

Figure 11: Results of the 6th experiment

γ/δ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0	1.79923998	1.79923998	1.79923998	1.79923998
0.01	1.97916409	1.98417556	2.10889048	2.11433520
0.02	2.15908820	2.17002377	2.41929207	2.43117323
0.03	2.33901212	2.35678481	2.73042432	2.74973388
0.04	2.51893602	2.54445919	3.04226501	3.06999556
0.05	2.69886017	2.73304753	3.35479330	3.39193553
0.06	2.87878407	2.92254971	3.66798924	3.71553851
0.07	3.05870823	3.11296631	3.98183309	4.04078291
0.08	3.23863212	3.30429798	4.29630644	4.36765061
0.09	3.41855599	3.49654449	4.61139161	4.69612416
0.1	3.59848018	3.68970646	4.92707122	5.02618660
0.11	3.77840404	3.88378462	5.24332875	5.35782147
0.12	3.95832788	4.07877860	5.56014834	5.69101349

Figure 12: Results of the 1st experiment

γ/δ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0	4.21273611	4.21273611	4.21273611	4.21273611
0.01	3.92066869	3.90817470	3.96280050	3.95206434
0.02	3.67727900	3.65447477	3.75121400	3.73129996
0.03	3.47133449	3.43989801	3.56946493	3.54161033
0.04	3.29481062	3.25605905	3.41143189	3.37663985
0.05	3.14182326	3.09680955	3.27259315	3.23169038
0.06	3.00795933	2.95753788	3.14953062	3.10320864
0.07	2.88984386	2.83471677	3.03961005	2.98845396
0.08	2.78485256	2.72560416	2.94076485	2.88527499
0.09	2.69091297	2.62803465	2.85134984	2.79195784
0.1	2.60636714	2.54027531	2.77003780	2.70711840
0.11	2.52987349	2.46092472	2.69574429	2.62962623
0.12	2.46033381	2.38883492	2.62757459	2.55854660

Figure 13: Results of the 2nd experiment

γ/δ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0	3.66755364	3.66755364	3.66755364	3.66755364
0.01	3.66755370	3.69889761	3.89985974	3.92983315
0.02	3.66755363	3.72951191	4.09163980	4.14513526
0.03	3.66755356	3.75941838	4.25248179	4.32468088
0.04	3.66755350	3.78863777	4.38919456	4.47642403
0.05	3.66755344	3.81719030	4.50674237	4.60615262
0.06	3.66755338	3.84509556	4.60882671	4.71817419
0.07	3.66755333	3.87237240	4.69826213	4.81575728
0.08	3.66755327	3.89903942	4.77722524	4.90142454
0.09	3.66755322	3.92511430	4.84742498	4.97715239
0.1	3.66755317	3.95061423	4.91022096	5.04450991
0.11	3.66755313	3.97555575	4.96670805	5.10475791
0.12	3.66755308	3.99995523	5.01777728	5.15892033

Figure 14: Results of the 3rd experiment

τ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0.5	3.59848019	3.64406334	4.89208289	4.94183266
1	2.69886017	2.73304753	3.35479330	3.39193553
1.5	2.39898683	2.42937560	2.84173680	2.87461798
2	2.24905016	2.27753963	2.58480692	2.61551916
2.5	2.15908816	2.18643805	2.43037214	2.45975650
3	2.09911350	2.12570367	2.32721545	2.35569575
3.5	2.05627445	2.08232197	2.25338240	2.28120259
4	2.02414516	2.04978550	2.19789281	2.22520681

Figure 15: Results of the 4th experiment

τ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0.5	2.69564779	2.62234687	2.89542691	2.82177740
1	3.26086371	3.16312917	3.46359054	3.36463595
1.5	3.54347167	3.43352032	3.72781517	3.61624283
2	3.71303645	3.59575502	3.87814557	3.75909007
2.5	3.82607964	3.70391149	3.97435645	3.85038104
3	3.90682477	3.78116612	4.04088636	3.91344348
3.5	3.96738362	3.83910677	4.08948305	3.95947336
4	4.01448495	3.88417197	4.12646007	3.99447680

Figure 16: Results of the 5th experiment

τ	Continuous	Non-continuous	Cont., Int.	Non-cont., Int.
0.5	3.66755341	3.81719027	4.89175171	4.96603982
1	3.66755344	3.81719030	4.50674237	4.60615262
1.5	3.66755346	3.81719032	4.30837214	4.42041925
2	3.66755348	3.81719033	4.18667875	4.30632270
2.5	3.66755349	3.81719034	4.10416726	4.22887530
3	3.66755350	3.81719035	4.04444600	4.17276746
3.5	3.66755351	3.81719024	3.99917812	4.13020524
4	3.66755351	3.81719025	3.96366334	4.09679127

Figure 17: Results of the 6th experiment

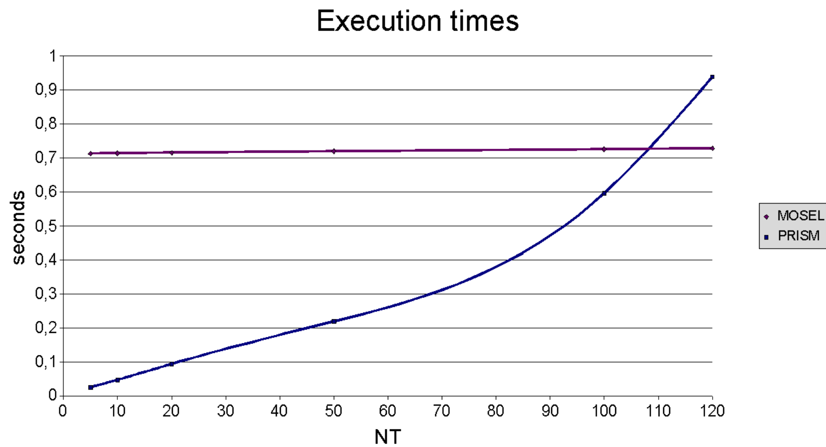

Figure 18: Results of the 2nd experiment

the two tools can be seen in the Figure 19 and Figure 18. In Figure 20, we can see a more detailed description of the PRISM benchmark (the times of the model construction and model checking are indicated separately).

The following preliminary conclusions can be drawn from benchmark:

- The execution times of the MOSEL system almost stay constant independently of *NT*;
- The execution times of the PRISM system increase rapidly with the increase of *NT*.
- The model construction time in PRISM dominates the execution time (rather than the model checking time).

NT	MOSEL	PRISM
5	0.7125	0.025
10	0.7135	0.047
20	0.715	0.094
50	0.719	0.219
100	0.725	0.596
120	0.728	0.938
150	-	1.550
200	-	2.377

Figure 19: Total execution times of the MOSEL and the PRISM in seconds

NT	Model const.	Model checking	Total
5	0.015	0.01	0.025
10	0.031	0.016	0.047
20	0.047	0.047	0.094
50	0.141	0.078	0.219
100	0.391	0.205	0.596
120	0.594	0.344	0.938
150	1.071	0.479	1.550
200	1.609	0.768	2.377

Figure 20: Execution times in seconds

While MOSEL is thus more efficient for smaller models, with PRISM also larger models can be analyzed. Furthermore, once a PRISM model is constructed, it can be arbitrarily often model checked with different parameter values (the PRISM "Experiments" feature). For such scenarios, the model checking time is more relevant than the model construction time.

6 Conclusions

Probabilistic model checkers like PRISM are nowadays able to analyze quantitative behaviors of concurrent systems in a similar way that classical performance analysis tools like MOSEL are. In this paper, we reproduced for the particular example of a retrial queuing system the results of an analysis that were previously generated with the help of MOSEL. The numerical results were virtually identical such that we can put confidence on the quality of the analysis. The construction of the models and the benchmarks of the tools demonstrate the following differences between both tools:

- The PRISM modeling language allows us to decompose a system into multiple components whose execution can be synchronized by combined state transitions; this makes the model more manageable than the monolithic MOSEL model. However, the decomposition can be only based on a fixed number of components such that *NT* terminals must be still represented by a single PRISM module.
- The state transitions in PRISM are described on a lower level than those in MOSEL: all guard conditions have to be made explicit (while the MOSEL `FROM` part of a rule imposes implicit conditions on the applicability of the rule) and all effects have to be exposed (while the MOSEL `TO` part of a rule imposes implicit effects); on the other side, this makes the PRISM rules more transparent than the MOSEL rules. In any case, the difference is syntactic rather than fundamental.
- Several kinds of analysis can be expressed in the property specification language of PRISM (by the definition of "rewards" and CSL queries for the long-term values of rewards) on a higher level than in MOSEL (where explicit calculations have to be written); on the other side, not every kind of analysis can be apparently directly expressed in PRISM (because different reward values can not be combined in a query); it is thus necessary to combine different results by external calculations.

- PRISM is also able to answer questions about qualitative system properties such as safety or lifeness properties and other questions that are beyond the scope of MOSEL.
- The time for an analysis depends in PRISM on the size of the state space of the system while it essentially remains constant in MOSEL, which on the other side puts a rather small limit on the ranges of state variables; the time growth factor in PRISM is super-linear, at least quadratic, and perhaps exponential (this remains to be investigated). While we were thus able to analyze larger systems with PRISM than with MOSEL, it is thus not yet clear whether the analysis will really scale to very large systems.
- As documented by the PRISM web page, the tool is actively used by a large community in various application areas; PRISM is actively supported and further developed (the current release version 3.1.1 is from April 2006, the current development version is from December 2007). On the other hand, the latest version 2.0 of MOSEL-2 is from 2003; the MOSEL web page has not been updated since that time.

The use of PRISM for the performance analysis of systems thus seems a promising direction; we plan to further investigate its applicability by analyzing more systems with respect to various kinds of features.

References

[1] B. Almási, J. Roszik, and J. Sztrik. Homogeneous Finite-Source Retrial Queues with Server Subject to Breakdowns and Repairs. *Mathematical and Computer Modelling*, 42:673–682, 2005.

[2] C. Baier, B. Haverkort, H. Hermanns, and J. Katoen. Model Checking Continuous-time Markov chains by transient analysis. In *12th annual Symposium on Computer Aided Verification (CAV 2000)*, volume 1855 of *Lecture Notes in Computer Science*, pages 358–372. Springer, 2000.

[3] J. Barner, K. Begain, G. Bolch, and H. Herold. MOSEL — MOdeling, Specification and Evaluation Language. In *2001 Aachen International Multiconference on Measurement, Modelling and Evaluation of Computer and Communication Systems*, Aachen, Germany, September 11–14, 2001.

[4] M. Bernardo and J. Hillston, editors. *Formal Methods for Performance Evaluation: 7th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2007*, volume 4486 of

Lecture Notes in Computer Science, Bertinoro, Italy, May 28 – June 2, 2007. Springer.

[5] Jr. Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model checking*. MIT Press, Cambridge, MA, USA, 1999.

[6] Ulrich Herzog. Formal Methods for Performance Evaluation. In Ed Brinksma, Holger Hermanns, and Joost-Pieter Katoen, editors, *European Educational Forum: School on Formal Methods and Performance Analysis*, volume 2090 of *Lecture Notes in Computer Science*, pages 1–37, Lectures on Formal Methods and Performance Analysis, First EEF/Euro Summer School on Trends in Computer Science, Berg en Dal, The Netherlands, July 3-7, 2000, Revised Lectures, 2001. Springer.

[7] Andrew Hinton, Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM: A Tool for Automatic Verification of Probabilistic Systems. In Holger Hermanns and Jens Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 12th International Conference, TACAS 2006, Vienna, Austria, March 27–30*, volume 3920 of *Lecture Notes in Computer Science*, pages 441–444. Springer, 2006.

[8] Christophe Hirel, Bruno Tuffin, and Kishor S. Trivedi. SPNP: Stochastic Petri Nets. Version 6.0. In Boudewijn R. Haverkort, Henrik C. Bohnenkamp, and Connie U. Smith, editors, *Computer Performance Evaluation: Modelling Techniques and Tools, 11th International Conference, TOOLS 2000, Schaumburg, IL, USA, March 27-31, 2000, Proceedings*, volume 1786 of *Lecture Notes in Computer Science*, pages 354–357. Springer, 2000.

[9] M. Kwiatkowska. Quantitative Verification: Models, Techniques and Tools. In *6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*, Cavtat near Dubrovnik, Croatia, September 3–7, 2007. ACM Press.

[10] G. Norman M. Kwiatkowska and D. Parker. Stochastic Model Checking. In M. Bernardo and J. Hillston, editors, *Formal Methods for Performance Evaluation: 7th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2007*, volume 4486 of *Lecture Notes in Computer Science*, pages 220–270, Bertinoro, Italy, May 28 – June 2, 2007. Springer.

[11] MOSEL — Modeling, Specification, and Evaluation Language, June 2003. http://www4.informatik.uni-erlangen.de/Projects/MOSEL.

[12] MOSEL-2, September 2007. http://www.net.fmi.uni-passau.de/hp/projects-overview/mosel-2.html.

[13] PRISM — Probabilistic Symbolic Model Checker, September 2007. http://www.prismmodelchecker.org.

[14] J. Roszik, J. Sztrik, and J. Virtamo. Performance Analysis of Finite-Source Retrial Queues Operating in Random Environments. *International Journal of Operational Research*, 2:254–268, 2007.

[15] William J. Stewart. Performance Modelling and Markov Chains. In *Formal Methods for Performance Evaluation: 7th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2007*, volume 4486 of *Lecture Notes in Computer Science*, pages 1–33, Bertinoro, Italy, May 28 – June 2, 2007. Springer.

[16] János Sztrik and Che Soong Kim. Performance Modeling Tools with Applications. *Annales Mathematicae et Informaticae*, 33:125–140, 2006.

[17] Unified Modeling Language (UML), version 2.1.1, 2007. http://www.omg.org/technology/documents/formal/uml.htm.

[18] Katinka Wolter, editor. *Formal Methods and Stochastic Models for Performance Evaluation*, number 4748 in Lecture Notes in Computer Science, Fourth European Performance Engineering Workshop, EPEW 2007, Berlin, Germany, September 27-28, 2007.