

# Simulating Queueing Networks with OMNeT++

Nicky van Foreest

December 30, 2007

## Abstract

This documents provides a quick introduction to OMNeT++, with a focus on queueing applications. The aim is to get the interested user up and running with a simple M/M/1 fifo queueing demo. After this the user can start exploiting more of OMNeT++'s functionality based on the extensive user manual and the sample directories.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Where to Get the Software? . . . . .	2
1.2	Feedback . . . . .	2
1.3	Structure of This Document . . . . .	2
1.4	Acknowledgments . . . . .	2
1.5	Versions . . . . .	2
<b>2</b>	<b>Simulating the M/M/1 Fifo Queue</b>	<b>2</b>
2.1	Getting the Simulation to Run . . . . .	2
2.2	Running the Demo . . . . .	3
2.3	Analyzing the Simulation Results . . . . .	4
<b>3</b>	<b>Modifying the Basic Example</b>	<b>4</b>
3.1	Changing the Simulation . . . . .	4
3.2	Speeding up the Simulation . . . . .	5
<b>4</b>	<b>The Implementation of the M/M/1 Simulation</b>	<b>6</b>
4.1	A Simple Method to Understanding the Code . . . . .	6
4.2	The .ned Files . . . . .	6
4.3	The .cc and .h Files . . . . .	7
<b>5</b>	<b>Why Use OMNeT++ for Queueing System Simulation?</b>	<b>7</b>

## 1 Introduction

OMNeT++ is a discrete event simulation environment based on C++. András Varga is the principal author and maintainer.

When I started to use it for simulating queueing networks, I was somewhat awed by the amount of functionality that OMNeT++ provides and the size of the user manual. Besides this, I had no experience with C++ programming, which was an extra hurdle in becoming acquainted with the power of the tool. Hence I had to invest some time to find out which parts and functionalities of OMNeT++ are specially useful for queueing systems analysis, and how to get these to work. I wrote this tutorial (this document together with some code examples) with the intention that a user new to OMNeT++ has a simple starting point from where he/she can further explore this simulation tool.

In this tutorial I assume that the reader has experience with programming, at least with C, and has some (very) basic understanding of queueing theory, say the first few sections of the chapter on queueing theory of [Ros93]. Furthermore, I will not explain the details of the code. These are for the reader to study. This may seem like a threat, but the code is not hard to understand.

Perhaps I should have started by mentioning the advantages of using OMNeT++. I refrain from doing this at this point of the document, but instead postpone it to the end. The reason for this is twofold. In the first place the OMNeT++ user manual and the OMNeT++ home page mention plenty of examples. In the second place, I present these advantages as based on my personal experiences, which I consider to be less important than have you start playing with OMNeT++ and this demo.

## 1.1 Where to Get the Software?

You can find OMNeT++ with a search on google. On my homepage <http://www.xs4all.nl/~stimp/nicky/> you can find the latest release of the queueing software, a library of basic queueing systems, such as a multi-server queue, a processor sharing queue, etc. If you want to make advanced plots of the simulation results, make sure to get **gnuplot**.

## 1.2 Feedback

Please do not hesitate to send me your comments, especially if you think the code or this document contains inaccuracies, unclear passages, etc. When you have something to contribute to the software library, please inform me so that I can include it in this package. I grant full acknowledgment with respect to authorship.

My email address is: [vanforeest@gmail.com](mailto:vanforeest@gmail.com).

OMNeT++ related questions can be posted at the OMNeT++ mailing list.

## 1.3 Structure of This Document

In Chapter 2 I discuss a working example of a queueing simulation in OMNeT++. The implemented example is the archetypal M/M/1 fifo queue. Chapter 3 shows some methods to change the working of the simulation, the simulation parameters, the simulation output, and the like. Chapter 4 discusses the implementation of the M/M/1 example.

## 1.4 Acknowledgments

- András: Thanks a lot for OMNeT++.
- Phillip M. Feldman pointed out a few typos and added footnote 1.

## 1.5 Versions

- 0.1 January 2001
- 1.0 May 2001
- 2.0 June 2002
- 2.1 October 2002
- 2.2 January 2003
- 2.3 October 2004 (minor changes)
- 2.4 December 2007 (minor changes)

# 2 Simulating the M/M/1 Fifo Queue

This chapter focusses on how to set up, run, and obtain output of the M/M/1 simulation in OMNeT++. Since I want to show as many features of OMNeT++ as possible, this example is a bit baroque in its output.

## 2.1 Getting the Simulation to Run

Before running the M/M/1 example it is a good idea to check that OMNeT++ and all the software it depends on is installed properly. I usually test this by running one of the standard samples included in the OMNeT++ distribution, such as the fifo queue in `omnetpp-version/samples/fifo/`.

Once you are convinced that OMNeT++ works, run `opp_makemake -f` in the directory that contains the source files of the queueing tutorial, perhaps `tutorial/queuesTut/` or `queues/queuesTut/`,

to make a **Makefile**. The binary **opp\_makemake** should be provided with OMNeT++. In case it cannot be found, try running it directly from the **src/utls/** directory of the OMNeT++ distribution. Once you have a **Makefile**, a simple **make** should do to get the binaries for the fifo simulation.

A general remark is of importance here. In case you change something in one of the files, you should run **make** again. When you decide to copy all files to another directory, or include other files in this directory, run **opp\_makemake -f** again. Run **make clean** to remove all object files and binaries.

The binary to run has the same name as the directory containing the Makefile that was generated by **opp\_makemake**. In my case this is **queuesTut**. Run it, i.e., type **./queuesTut** at the prompt. If everything works the way it should, a few windows should pop up.

I consider it a bit too much work to type in **make** and then the binary again and again. Therefore I included a very simple script **run** which does this for me.

As an aside, the fifo simulation that is part of the tutorial is not the same as **fifo1** and **fifo2** of the **samples** directory. The one provided in the tutorial is a stripped down version that includes just the bare minimum to simulate a fifo queue. The **fifo1** and **fifo2** show some more advanced tricks with OMNeT++.

## 2.2 Running the Demo

Once the simulation has started you should see two windows:

- A graphics window showing a small queueing network: a job generator, a fifo queue, and a job sink.
- The navigation window, called **OMNeT++/TKenv**, containing the main functionality to run the simulation.

Double click on the fifo widget on the canvas of the graphics window. A new box should appear, called **(Fifo) fifonet.fifo[0]**. Click on the tab **Contents**. You will then see a few lines with text. The line with the string **cQueue** contains the word **(empty)**. This will change to the number of customers in queue (= one less the number in the system when a customer is served) once the simulation is running. Then double click on the line with the word **Job Distribution**. You will see a blank blue (on my screen, that is) window. Here a histogram will show, once the simulation started, the probability density function of the number of jobs in the system as seen by an arriving job. Finally, double click the line **Jobs in system**. A yellow canvas will come up. This will contain a simple plot of the number of jobs in the system as a function of time.

Now you can start the simulation by pressing the **RUN**-button in the **OMNeT++/TKenv** window. In the graphics window you should see jobs hopping from the generator to the queue, and from the queue to the sink. Press **STOP** after several seconds. The line **cQueue** of the **(Fifo) fifonet.fifo[0]** window will now, with probability > 0, show that the queue contains some customers. When you double click on this line, another window will tell you which jobs are currently waiting in the queue.

To make a distinction between individual customers may not seem very interesting from a queueing perspective; the M/M/1 queueing model assumes that customers do not have real identity—they only differ by their service requirements and arrival time. However, when studying more complicated models, for instance when jobs have different type, this information becomes relevant. It becomes even more important if you want to study protocol and queue interactions, such as when multiple TCP sources share buffered resources, i.e., routers.

The histogram window should now contain a few black bars. When you put your mouse pointer on one of these, it will change color. (Mine becomes gray.) There is a line at the bottom of window that shows the number of arrivals that observed a certain number of jobs in front of it in the system. In other words, if cell # 1 contains 10, this means that 10 arrivals saw one job in the system<sup>1</sup>.

Clicking the **Job Distribution** line in the **(Fifo) fifonet.fifo[0]**-window with the right mouse button, enables you to choose between a graphical representation of the gathered data, and a textual one, called **Object**. I discussed the graphical one above. The text information will provide you with aggregate statistics such as the mean number of jobs found upon arrival, etc.

The last step will be to press **EXPRESS** in the **OMNeT++/TKenv** window; I leave the other buttons for you to discover. In fact, my best general advice is to press on any buttons you may see, and find out what it does. Do not forget that often both mouse buttons, the left and the right, can be used for different effects.

<sup>1</sup>When the arrival process is Poisson, i.e., the interarrival time is exponentially distributed, the distribution of the number of customers seen by arrivals is the same as the time average of the number of customers in the queue. This phenomenon is due to the PASTA (Poisson Arrivals See Time Averages) property, see for instance [Ros93] or [Tah98]. In general you should be careful to choose a sampling method of the simulation that provides the statistics you want to collect. Beware!

## 2.3 Analyzing the Simulation Results

If you have not interrupted the simulation by pressing the big red **STOP** button, the simulation will finish after having generated 5010 jobs. On my linux machine this takes less than a second.

Notice that the textual and graphical representations of the histogram are updated after the simulation has finished. The **OMNeT++/Tkenv** window contains as well a number of results. You might need to scroll a bit up and down in the window to view all the simulation results.

The directory will now contain a few new files as well.

- **fifo.sca** This file contains the statistics that were gathered in the variable `jobDist` of class `cDoubleHistogram`, see the user manual section 7.9, such as the number of jobs generated, the histogram data, etc. Please have a look at it now, which can be done by pressing **view->Output Scalar File** in the taskbar of the main window.
- **fifo.vec** This file contains info about the dynamics of the number of jobs in the system. You should process this with the OMNeT++ tool **plove**. Chapter 10 of the user manual contains the instructions. Briefly, start up **plove**. The left most button allows you to **File->Load output vector file** the file **fifo.vec**. Send it to the right window with the arrow button in the middle. By double clicking the line **Jobs in System in fifonet.fifo[0]** a plot should show the number in queue as a function of time.

In general you can process these files—I use Python to much avail—to postprocess it. You might instead want to do the post-processing in the function `Fifo::finish`, about which I will talk later. Here comes in the some of the power of OMNeT++’s use of a real programming language: you can incorporate your own post-processing functions in the simulation itself, in case you want to.

Summarizing, OMNeT++ provides a number of ways to output simulation data:

- dynamically, by means of jobs jumping from one node in the network to another;
- graphically, by means of dedicated collection classes such as `cOutVector`;
- verbally, by means of the `ev <<` statement in the code which is output to the canvas of the **OMNeT++/Tkenv** window;
- by means of files.

## 3 Modifying the Basic Example

The previous chapter discussed only one type of queue, the M/M/1 queue, with one specific parameter setting. You probably want to simulate other types of queues, more generic service distribution, etc. Another interesting degree of freedom is to speed up the simulation by leaving out the windows.

### 3.1 Changing the Simulation

I will discuss some ways to change and extend the current simulation environment. The most important way to do this is with the `omnetpp.ini` file, see chapter 9 of the user manual for all details.

**Changing the Interarrival and Service Rate** Change the numbers in `omnetpp.ini` related to the exponential distribution. Mind that the interarrival rate and service rate are the inverse of the parameters you specify, e.g. `fifonet.fifo[1].service_time = exponential(2)` means that the expected service duration  $= 1/\mu = 2$ .

Besides with `omnetpp.ini` file, you can change these values by pressing **Params** in the **(Fifo) fifonet.fifo[0]** window. Click for instance on the line with `service_time`. Now you edit the field in the window that will appear. To try, for instance, another distribution, type in its name and parameters, see the next paragraph for more on this. Do not forget to press the **Enter** key on your keypad to make the edit effective.

**Changing the Interarrival and Service Distribution** OMNeT++ provides quite few standard distributions, such as the uniform and exponential distribution, see sections 4.7.4-4.7.7 and 7.3.2 of the user manual. In case you need other distributions, you have to build them yourself. This is in itself quite interesting, see Ross [Ros93], and, of course, Knuth[Knu97]. Once you know mathematically speaking what to do, OMNeT++ makes the implementation very easy. You define your new distribution in the file `distributions.cc`. I have already provided a few simple

examples. Note that you can use the standard distributions of OMNeT++ in your new functions right away. Once you have implemented your new distribution, you can simply use its name in `omnetpp.ini` and pass appropriate parameters to it. As an example, if you would like to use the distribution `perturbedExponential`, which you can find in the file `distribution.cc`, as service distribution, then you should use the following line in `omnetpp.ini`: `fifonet.fifo[0].service_time = perturbedExponential(2,1)`

**Changing the Network** The `omnetpp.ini` parameter `fifonet.num.buffers` enables you to make a network of tandem queues. Changing it from the current value 1 to  $k$ , say, will put  $k$  queues in tandem. You should as well change the line `fifonet.fifo[0].service_time = exponential(2)` to `fifonet.fifo[*].service_time = exponential(2)`, that is, the 0 has to change to `*` to reference all fifos, instead of just the first one with id 0. I actually chose a slightly different approach in `omnetpp.ini`. First I give `fifo[0]` its value, and then the others by means of the `*`.

You will find included as well a ring network. This is built out of simple fifo queues that are connected sequentially to each other. If you want to run it, you have to change a few lines in `omnetpp.ini`. This file contains a few remarks on where the changes should be made. Since the ring is a closed queueing network, there are no external arrivals of jobs nor sinks. Therefore these two modules are not included in the ring. Furthermore, the buffers should contain some initial number of jobs that will start circling around. You can set these numbers by the parameter `ring.fifo[*].num_init_jobs = 20` to be found in `omnetpp.ini`.<sup>2</sup>

More general networks are for you to build; Chapter 3 of the manual (`usman.pdf`) will tell you how. You should as well consult the following OMNeT++ samples. The sample directory `token` provided with the standard OMNeT++ distribution shows a circular network. The `fddi` sample shows a large network. (Do not forget to click on one of the rings in the `fddi` sample to see how complicated a network you can actually simulate with OMNeT++.)

**Changing the Job Scheduling** I have implemented, at the moment, some other queueing models. In case you are interested, please have a look at the `queues` package, to be found on my homepage or the OMNeT++ server.

**Changing the random number generator** For the more suspicious of you, section 6.4 of the OMNeT++ user manual discusses the random number generator. You can replace this with your own if you like.

## 3.2 Speeding up the Simulation

If you are convinced that everything works the way it should, and you are just interested in numerical output, you can run the simulation straight from the prompt with the `cmdenv` mode. No more windows will appear, only the output files will be produced. Due to this, the simulation will become faster as well.

To achieve this, change these lines in the `Makefile`

```
# User interface (uncomment one) (-u option)
#USERIF_LIBS=$(CMDENV_LIBS)
USERIF_LIBS=$(TKENV_LIBS)

to

# User interface (uncomment one) (-u option)
USERIF_LIBS=$(CMDENV_LIBS)
#USERIF_LIBS=$(TKENV_LIBS)
```

Do a `make clean` and `make` to remove all window related code from the simulation executable. Now it should work. Be aware that a new `opp_makemake -f` reverts the makefile to the old situation, i.e., the simulation with tk windows.

You will want these options in the following section of `omnetpp.ini`:

```
[Cmdenv]
runs-to-execute = 1
module-messages = no
```

---

<sup>2</sup>For the interested. If you analyze the expected number of jobs in the queues, and add them, it will appear as if one job is missing, this is the application of the Arrival Theorem, see e.g. [Ros93].

```
verbose-simulation = no
Display-update = 1h
```

Play with these options to discover that the number of lines of simulation output will be a bit too much to handle. See Chapter 9 of the manual for more info.

## 4 The Implementation of the M/M/1 Simulation

Before starting the main subject of this chapter, I need to define one concept: a *functional entity*. A functional entity is a part of a simulation that carries out a specific action on a job, or a message. For example, a server or a message generator are functional entities. They are, so to say, the essential functional units that take care of one process step in the lifetime of a job, or message. The implementation of such entities will be called *modules*. Now back to the simulator.

The simulation environment is built out of, mainly, two types of files: `.ned` files and `.cc` files. The `.ned` files roughly describe how the functional entities connect to each other and communicate; the `.cc` files contain the C++ code that implement the behavior of the entities. These file types I will discuss in some more detail below. I expect you to have the `.ned` and `.cc` files belonging to this demo at hand.

After having read this tutorial, be sure to check out the demos on the `omnet.html` pages in `src`, and the samples `fifo1` and `fifo2`. They are instructive and show additional functionality of OMNeT++ relevant for queueing systems analysis.

### 4.1 A Simple Method to Understanding the Code

There are various ways to understand the implementation of a new simulation. I found the following approach the most useful. First I run it, of course, to get an understanding of where the various entities reside and how they exchange jobs. Then I work through the files belonging to each entity separately. I start with the `.ned` file to understand what goes in and out of a module, and the parameters it needs. Then I give the header file a brief look to become familiar with the module's specific internal variables and functions. Finally I study the C++ code belonging to the module. Once I somewhat understand what this first module does, I take the next module in the simulation chain, that is, the module that gets its messages from the one I just studied. Working this way, I gain a quicker understanding of what is going on than by first working through all `.ned` files, then all the C++ files, etc.

In this document I will not, however, follow the above suggestion, mainly for brevity. Here I only want to illustrate certain key points of the simulation, and leave the studying to you.

### 4.2 The `.ned` Files

Each entity in a simulation needs to communicate via *messages* with itself and other entities. Messages can be used for various purposes. One is to represent jobs that need service at queues. Another is to convey information of the entity's state to itself in the future, or to other entities.

Specifically, the `fifo` example contains three entities: a source module, a `fifo` module and a sink module. The source module generates messages that represent jobs. These job-messages are sent to the `fifo` module. This in turn delays the messages according to the present queue, then services them, and finally sends the messages to the sink module. The sink module processes these messages to extract some final statistics. The sink is where the messages leave the queueing network. The sink module releases as well the memory allocated to the job messages.

Clearly, all these modules need in- and output gates to receive and send messages. The `.ned` file specifies these gates. In our example, the `fifo` module has an input and an output gate. Furthermore these modules may need parameters such as service rate, queue size, etc. These parameters need be *declared* in the `.ned` files too. The *definition* of the parameters, i.e., giving them a value, takes place when the simulation starts.

The next step is to glue all the modules together to form a network. This should be done by means of another `.ned` file that defines a *compound module*. This new module uses the simple modules and connects them into networks. Besides this, it passes parameter values on to them during the simulation. Finally it gives some directives on where on the canvas the simple modules should appear. To state this in terms of the `fifo` queue, `FifoNet`, defined in `fifonet.ned`, connects the output of the source to the input of the `fifo` queue, and the output gate of the queue to the sink. Consult `fifonet.ned` for all details. Take especially notice of the syntax, and be aware of when to

use a comma ‘,’ and a semicolon ‘;’. I made some time consuming errors by not using them in the appropriate way.

Note that this hierarchy of modules—simple modules forming compound modules, forming in turn other compound modules—provides a very efficient way to build large, complex queueing networks. Start by building the parts, and then connect these parts to larger compound modules up to entire networks.

### 4.3 The .cc and .h Files

Now we have defined all entities that take part in the process cycle of job-messages, and specified the routing between these entities in the compound modules, we should tell how each entity should process job-messages. In concrete terms, the generator should generate jobs with certain interarrival times distributed according to some specified probability distribution. These jobs are represented by a message, so that in fact the generator generates messages. Then it sends these messages out of its output gate. If you will have a look at `gen.cc` you will see that it will produce in total `fifonet.gen.num.messages` messages, each separated some time `ia.time` apart. These two parameters are defined in the file `omnetpp.ini`.

`fifo.cc` contains the heart of the simulation. Once you understand this code, which may be a bit hard at first, you have mastered the most important aspects of OMNeT++.

Once the jobs arrive at the fifo queue they will be queued or not, depending on whether the server is idle or not. These arriving jobs are handled in the function `Fifo::handleMessage`. Lets see what it does. Suppose the server is busy. Then the event stack of the simulator should contain an event that indicates when the job’s service is supposed to be finished. If you think a bit about how to implement such events, taking into account that all events are related to messages, you will understand that the server should send itself an `endService` message at the moment a job’s service starts, to indicate when this job’s service should stop. This `endService` message will be put on the event stack by the event scheduler (not to be confused with the queueing scheduler!). At some point in time, the event scheduler will remove this message from the stack, and will give it to the fifo module. Once this module receives the `endService` message, it knows that the job’s service has ended, and that it should send it to the sink. When the queue is empty after the departure, the server should just wait. On the other hand, if there are jobs in queue, it should take the job in front of the queue, and start another service period. Now have a look at `fifo.cc` and think deeply about it until you understand what is exactly going on. Remember well that a simulation entity has no other means to communicate with itself in the future than via self-messages, and that these messages have to be put on the internal event stack.

`fifo.cc` contains four more functions. The functionality of `Fifo::initialize` should be clear from its implementation. The function `Fifo::finish` processes part of the simulation results when the simulation has finished. The other two functions `Fifo::serviceRequirement` and `endService` enable you to specify what to do when a job starts service—here only a random service duration is generated—and when it leaves. This provides interesting flexibility with respect to queueing simulation. For instance, in multi-class networks, jobs may change class from service station to service station, or a fraction may be sent back for another processing step at the server, etc. etc. Do not be tempted too easily to handle this kind of functionality in the `handleMessage` function. It can become quite complex.

The last module of interest, the sink, should be a breeze once you have mastered the above. It is a good idea to think about why we compute the waiting time of a job in this module, and why this is not possible in `fifo.cc`. This is of course by no means a generic property of OMNeT++, but just a consequence of the current implementation of the queueing entity.

## 5 Why Use OMNeT++ for Queueing System Simulation?

Some time ago I worked for a large telecommunication company. In one of the projects in which I participated, we tried to reduce the convergence time of a distributed network restoration protocol. The interaction between the protocol, the involved timers, and the states of the queues in the equipment became soon completely intractable, analytically speaking. Hence we decided to analyze the behavior of the network and equipment with a simulator. As the developers wrote the protocol in C++, we wanted to use this very code in a simulator. Besides this we needed a graphical environment

to see whether the protocol messages traveled the way we wanted, and ended up in the right queues. OMNeT++ provided this, and more.

The general experiences I gained during this and other projects with respect to simulation complex systems were briefly as follows:

- Using production code in a simulator saves a major amount of time compared to having to redesign the work in some kind of ‘high-level’, seemingly user friendly simulation environment. Besides this, the simulation itself helps to test the real code, that is, the code that will be actually implemented in commercial equipment.
- Having insight into the code of the simulator is helpful to understand details of the simulation. Furthermore, it can be necessary to extend parts of the tool’s functionality. Being dependent on companies to give you support, but keeping you securely away from the details of the implementation of their simulator, is not always what you want. Worse yet, when you have a problem with the simulator you have to wait for their answer, which does not always relate to your question ...:-).

Based on the above experiences as well as on some past work with commercial simulators, I want to point out some advantages, at least in my opinion, of running queueing simulations within OMNeT++.

- The behavior of systems is programmed in a ‘normal’ programming language, in this case C++. This provides the user way more flexibility than a ‘high-end’, tool-specific programming environment. For instance, if you use Bones, your programming skills gained during working with Bones only apply to the Bones environment. Besides this, these high-end languages appear, as long as you stick to the provided tutorials, to be quite generic. But when you try your hands on your own examples, you will feel how restrictive these tool-specific languages sometimes are.
- As a consequence of the use of C++, OMNeT++ is flexible and extensible. As an example, you can easily implement your own job interarrival and service distributions, statistical tests on the simulation results, more general networks, scheduling distributions, etc.
- It is under GPL license.
- OMNeT++ enforces ‘separation of concerns’. You have to specify the separate functional entities, such as queues, schedulers, job generators, separately from routing functionality, that is the definition of the network topology.

Let me say one more thing about the first bullet. Of course, these commercial tools allow you to specify parts of the system in your own code, and connect this code with the simulation platform by means of software hooks. But then you have to figure out how this works plus all debugging, and you still have to program in a real programming language.

With respect to the last item, I hesitated to use the possibly vague words ‘object-oriented’, but that is what it really is. In my experience this is a bonus, as it makes the simulation environment modular. You can easily change only one part, or a module. As long as the interfaces between the modules remain the same, everything will remain working. Changing the topology is easy too. Take for instance this demo. If you want to build Jackson networks—a bunch of M/M/1 queues connected in a network such that jobs can arrive and leave the network—you only have to figure out how to set up such networks. the functionality of job sources, sinks and service stations is already there, ready to use. If you want to change some of the service rates, you can easily do this in the `omnetpp.ini` file.

Let me stop here, and let you convince yourself about the uses of OMNeT++.

## References

- [Knu97] D.E. Knuth. *The art of computer programming*, volume 2, Seminumerical algorithms. 3 edition, 1997.
- [Ros93] S.M. Ross. *Introduction to Probability Models*. Academic Press, 5th edition, 1993.
- [Tah98] M. El-Taha and S. Stidham Jr. *Sample-Path Analysis of Queueing Systems*. Kluwer Academic Publishers, 1998.