Laboratory Lesson 1

1   Instructions

To be able to understand the computer exercise and to complete it within the given time limits, you have to prepare yourself properly by completing all home assignments. If you have any questions regarding the home assignments, you are welcome to ask your tutor. However, when you arrive to the computer exercise, all assignments must have been completed. Further, make sure to be in time for the exercise. To be given access to the computer exercise, you have to complete all home assignments.

In this computer exercise, queuing theory will be used to analyze delay and loss systems both theoretically and through Matlab simulations. To be able to run the simulation programs, some Matlab-files need to be downloaded from the course home page before the start of the computer exercise and copied to your home directory. During the computer exercise, several different tasks have to be solved. For every task, there are a number of questions that have to be answered. Some of the questions are marked (**) and these shall be discussed with the tutor before proceeding with the computer exercise. However, you are supposed to do some thinking yourself before talking to the tutor. If the tutor happens to be occupied elsewhere, it is allowed to continue for a short while until he becomes available.

2   Home Assignments

In the first part of this computer exercise, the behavior of a web server will be studied by making a model based on queuing theory. Out of simplicity reasons, we first assume that the server can be approximated as an M/M/1 system with infinite queue. This assumption is not entirely correct since the queue must be limited in size due to memory restrictions. The arrival and service rates are λ and μ, respectively.

1. Draw the state transition diagram of our model.

2. Derive the steady state probabilities, p_k, meaning the probability for having k jobs in the system.

3. Derive the expression for the mean number of jobs in the system, N, from the steady state probabilities, p_k.

4. Derive the mean time a job spends in the system, T, from the mean number of jobs, N.

In order to make our model more realistic, we now assume that the **queue** is limited to L positions.

5. Draw the corresponding state transition diagram.

6. Derive the steady state probabilities, p_k, of the new model.

7. Express the time and call congestion (tidsspärr respective anropsspärr) using the steady state probabilities, p_k.

8. Study the program code and line comments below. Comment the remaining
   code lines in the table found in the following page.

```
1. clear all; %empty memory
2. rand('seed',0)%set random seed
3. lambda=7; %arrival rate
4. mu=10; %service rate
5. endtime=1000; %simulation length
6. t=0; %current time
7. tstep=1; %average time between consecutive measurement events
8. currcustomers=0; %current number of customers in system
9. event=zeros(1,3); %constructs vector to keep time for next arrival
%(pos 1),next service completion (pos 2) and next
% measurement event (pos 3)
10. event(1)=exprnd(1/lambda); %time for next arrival
11. event(2)=inf; %no next service completion (empty system)
12. event(3)=exprnd(tstep); %time for next measurement event
13. nbrmeasurements=0; %number of measurement events so far
14. nbrdeparted=0; %number of departed customers
15. nbrarrived=0; %number of customers that have arrived throughout the
%simulations
16. while t<endtime
17. [t,nextevent]=min(event);
18. if nextevent==1
19. event(1)=exprnd(1/lambda)+t;
20. currcustomers=currcustomers+1;
21. nbrarrived=nbrarrived+1; %one more customer has arrived to
%the system through the simulations
22. arrivedtime(nbrarrived)=t; %the new customer arrived at time t
23. if currcustomers==1
24. event(2)=exprnd(1/mu)+t;
25. end
26. elseif nextevent==2
27. currcustomers=currcustomers-1;
28. timeinsystem=t-arrivedtime(nbrarrived-currcustomers);
29. nbrdeparted=nbrdeparted+1; %one more customer has departed
%from the system through the
%simulations
30. T(nbrdeparted)=timeinsystem;
31. if currcustomers>0
32. event(2)=exprnd(1/mu)+t;
33. else
34. event(2)=inf;
35. end
36. else
37. nbrmeasurement=nbrmeasurement+1; %one more measurement event
38. N(nbrmeasurements)=currcustomers;
39. event(3)=event(3)+exprnd(tstep);
40. end
41. end
```

5

Fill in your comments below:

Line	Comment
16	
17	
18	
19	
20	
23-25	
26	
27	

Line	Comment
28	
30	
31-32	
33-35	
36	
38	
39	

9. What kind of system is simulated?

10. Which output vectors can be used to obtain the mean number of customers and mean response times from the simulation results?

3   Laboratory assignment: M/M/1 with infinite queue

First, we assume that the web server can be modeled as an M/M/1 system with infinite queue. Obviously, these assumptions are not entirely correct since the queue must be limited in size due to memory restrictions. It is also very doubtful if the time between consecutive arrivals and the service time are exponentially distributed. A Matlab-program simulating an M/M/1 system with $\lambda = 7\,\mathrm{s}^{-1}$ and $\mu = 10\,\mathrm{s}^{-1}$ during 1000 seconds can be found in the Matlab-file MM1.m (download from the course home page).

1. Simulate the system and plot the number of jobs in the system at the instant of a measurement event. Hint: Use the N-vector. Draw the plot in the graph below. report the behavior you observe. what is the largest value for N? why?

2. Calculate the average number of jobs in the system, N, from the simulation results. Hint: Use the Matlabfunction mean.

3. Calculate the theoretical mean value for N (see Home assignment 3) and compare the obtained value with the simulated value.

4. Make a histogram over the obtained measurements for N. Hint: Write histogram, which is described in the appendix, and follow the instructions. What does the histogram show?

5. Normalize the histogram by pressing enter and plot the result in the graph below.

A program called pkMM1, which is described in the appendix, calculates the steady state probabilities for an M/M/1 system with infinite queue. Plot the steady state distribution in the same diagram as the histogram.

6. What does the histogram and the curve show and how have they been obtained? (**)

7. Alter the simulation length to 50 s and then 3000 s. Redo Exercise 3:2 for the new simulation lengths. What impact does the simulation length have on the accuracy of the numerical results?

The number of jobs in a system is interesting for the operator when dimensioning buffer sizes, but for the customers, the total response time is a more interesting measure. Now we will continue to study the behavior of an M/M/1 system using the simulation length of 1000 seconds.

8. Calculate the mean response time, T, directly from output vector (T). Hint: Use the Matlab-function mean.

We will now investigate what will happen to the mean response time when λ is altered.

9. Plot the number of jobs at the instant of a measurement event for $\rho = 1.1$ and compare it with the corresponding plot for $\rho = 0.7$, see Exercise 3:1. Explain the differences. (**)

10. Fill in the mean response time, T, for $\rho = 0.25, 0.5, 0.7$(see Exercise 3:8), $0.9, 1.0, 1.1$ in the table below.

ρ	0.25	0.5	0.7	0.9	1.0	1.1
T						

11. Plot the mean response time, T, for the investigated ρ-values in the graph below. How does the response time change with the system load?

12. For $\rho = 0.25$ and 0.9, plot elements 1-100 of the T-vector as a function of elements 2-101 by writing plot(T(2:101),T(1:100),'*') in the command prompt.

$\rho = 0.25$ $\rho = 0.9$

13. How is the dependency between consecutive response times influenced by the system load?

The accuracy of the measurements can be improved if the measurements are not started until the system has settled. Then, the first measurements, which perhaps are not really representative for the system (very few customers), are discarded. This is especially important for systems where the initial states are very unusual.

14. Which is the initial state and most usual state for an M/M/1-system with $\rho < 1$? Hint: Use the steady state probabilities when determining the most usual state.

15. How long do we have to wait until the system has settled for an M/M/1 system with $\rho \geq 1$? (**)

4 Laboratory assignment: M/G/1- infinite queue

The time between consecutive arrivals and the service time can usually not be approximated as being exponentially distributed. For example, a computer switch handling ATM packets, having a fixed packet length, will have a service time distribution that rather should be approximated as deterministic than exponential. The arrival process to a web server is yet another example of a system where an approximation to an exponential distribution is unsuitable.

Now, we will compare the behavior of a system with one server, infinite queue and exponential interarrival times, but with different service time distributions. Deterministic, uniform, exponential and hyperexponential distributions will be investigated.

Simulate the system for the service time distributions using $\lambda = 7\,\mathrm{s}^{-1}$, a mean service time of 0.1 seconds and a simulation length of 1000 seconds. Hence, the original parameter settings shall be used.

1. Calculate the mean number of jobs and the mean response times for the different simulation scenarios. Hint: Use the files displayed in the table.

Distribution	File	N	T
Deterministic	MD1.m		
Uniform	MU1.m		
Exponential	MM1.m		
Hyperexponential	MH21.m		

2. How does the variance for the service time seem to influence these mean values (N and T). (**)

3. For which λ-value do the systems in Exercise 4:1 become unstable? What is the corresponding λ-value for a G/G/m system? This exercise shall be solved through theoretical reasoning.

In most systems, it is not only the mean response time that is considered to be important but also how many jobs in percent that have a response time less/more than Z seconds. A common criteria used for dimensioning is to not allow more than X percent of the jobs to have a response time longer than Z seconds.

4. The t-values for the response time distribution function (0.1, 0.9) of the different simulation scenarios are displayed in the table below. How does the difference between the distributions alter with the different percent numbers? (**)

Distribution	$P(T \leq t) = 0.1$	$P(T \leq t) = 0.9$
Deterministic	0.1000	0.4299
Uniform	0.0802	0.4506
Exponential	0.0348	0.6764
Hyperexponential	0.0283	0.8256

5   Laboratory assignment: M/M/1*limited queue

Set $\lambda = 7\,\mathrm{s}^{-1}$ and limit the queue to 6 positions by altering the Matlab-file MM1.m. Hint: Insert an admission control feature (if-loop) in the arrival event that investigates if the queue is full.

1. How many jobs can now coexist in the system?

2. Determine the time and call congestion from the simulation results. Hint: Use the histogram function to obtain an estimation of the steady state probabilities.

3. Calculate the theoretical values for the time and call congestion by using the function pkMM1b. Increase ρ to 1.1 by increasing λ.

4. Calculate the time and call congestion. How does the load influence the system stability? Compare with the results obtained for a system with infinite queue, see Exercise 3:9. (**)

Appendix

histogram Plots a histogram over the elements in a vector

pkMM1(lambda,mu) Calculates and plots the steady state probabilities for an M/M/1-system with infinite queue. Input parameters are λ and μ. The numerical value for the steady state probabilities can be obtained by a mouse click in the diagram. To continue, click outside the diagram.

pkMM1b(lambda,mu) Calculates and plots the steady state probabilities for an M/M/1-system with limited queue. Input parameters are λ and μ. Thereafter, the program asks for the queue size. The numerical value for the steady state probabilities can be obtained by a mouse click in the diagram. To continue, click outside the diagram.