Università degli Studi della Campania
Luigi Vanvitelli

Course of
COMPUTER SYSTEMS MODELING AND SEMANTIC WEB

Exercises from the course

Authors:
Prof. Mauro Iacono

November 30, 2021

Abstract

This is a collection of exercises that have been proposed for the Computer Systems Modeling and Semantic Web course in the third year of the Data Analytics international bachelor.

Exercises are inspired to several sources and a possible solution is proposed. Solutions include the use of specific software. We will use the Java Modeling Tools simulator.

1 Exercise: the virtual channels[1]

We have a computer network composed of 4 routers with Poisson distributed service times. Router 1, 2 and 4 can process an average of 10 packets per second. Router 3 can process an average of 20 packets per second.

The network is serving 4 virtual channels. The first channel passes through routers 1, 2 and 3, then delivers the packets to destination. The second channel passes through routers 1, 3 and 4, then delivers the packets to destination. The third channel passes through routers 2, 3 and 4, then delivers the packets to destination. The fourth channel passes through router 3, then delivers the packets to destination. The sources of the channels produce packets according to Poisson distributions as well. The rates are 3, 4, 5 and 6 packets per second respectively.

1. What are the utilizations of routers?
2. What are the average number of packets at the routers?
3. What are the response times of the routers?
4. Which channel has the highest response time at router 2?
5. Which router is the bottleneck of the system?
6. What is the average total time in the system for packets using virtual channel 2?

1.1 Solution

We have a network of four interconnected routers that are visited in different ways by packets that are generated at different sources and travel into separate virtual channels. These channels are using the same 4 routers in different ways.

Since packets of all channels have the same service times in each router, the model can be approached analytically.

As question 6 requests a total time in the system, we can model the system as a queuing network. The routers transparently process packets from any source, making no difference between the channels (e.g., there is no priority or quality of service policy). Anyway, questions 4 and 6 suggest that the four channels have to be dealt with separately, in order to compare what is happening for a specific channel, that is, a specific source. This request suggests the definition of 4 classes.

The system can be modeled as in Fig. 1.

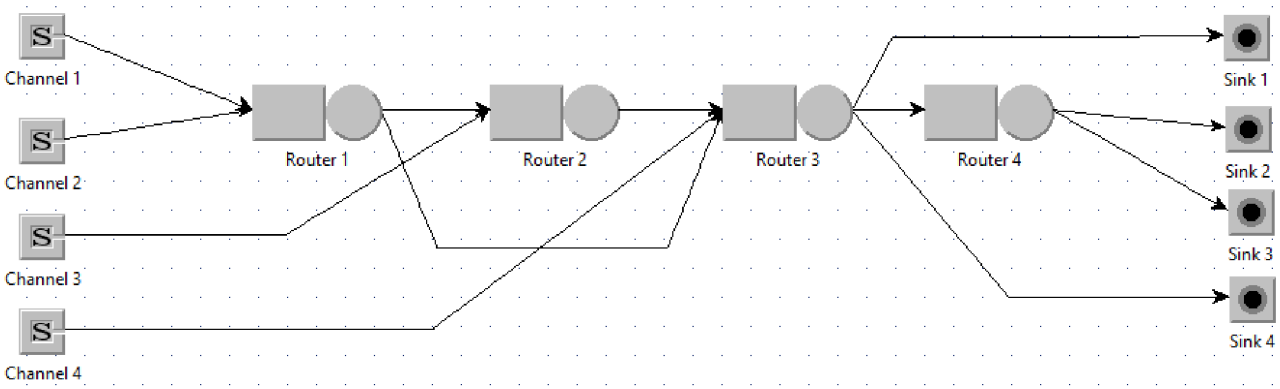

Figure 1: The model

The questions may be answered by applying an analytical approach or by using the JMT simulator.

[1]This exercise is based on the exercise presented in [1] in subsection 18.6.1, page 322.

1.2 Analytical solution

For question 1, $U_k = \lambda_k/\mu_k$, where μ_k are the provided (class-independent) service rates and λ_k should be computed.

As each channel is routed deterministically, $\lambda_{k,c}$ (where k identifies the router or the source if 0, and c identifies the channel) can be immediately obtained for each router:
$\lambda_{1,1} = \lambda_{2,1} = \lambda_{3,1} = \lambda_{0,1} = 3$ job/s,
$\lambda_{1,2} = \lambda_{3,2} = \lambda_{4,2} = \lambda_{0,2} = 4$ job/s,
$\lambda_{2,3} = \lambda_{3,3} = \lambda_{4,3} = \lambda_{0,3} = 5$ job/s,
$\lambda_{3,4} = \lambda_{0,4} = 6$ job/s.

The total $\lambda_k = \sum_c \lambda_{k,c}$ for each router are:
$\lambda_1 = \lambda_{1,1} + \lambda_{1,2} = 3 + 4 = 7$ job/s,
$\lambda_2 = \lambda_{2,1} + \lambda_{2,3} = 3 + 5 = 8$ job/s,
$\lambda_3 = \lambda_{3,1} + \lambda_{3,2} + \lambda_{3,3} + \lambda_{3,4} = 3 + 4 + 5 + 6 = 18$ job/s,
$\lambda_4 = \lambda_{4,2} + \lambda_{4,3} = 4 + 5 = 9$ job/s.

Consequently, average utilizations are:
$U_1 = \lambda_1/\mu_1 = 7/10 = 0.7$,
$U_2 = \lambda_2/\mu_2 = 8/10 = 0.8$,
$U_3 = \lambda_3/\mu_3 = 18/20 = 0.9$,
$U_4 = \lambda_4/\mu_4 = 9/10 = 0.9$.

This answers question 1, and the problem can be solved, as the system is stable; since $U_3 = U_4 = 0.9$ is the highest average utilization value, routers 3 and 4 are the bottleneck, and this answers question 5.

In order to answer question 2, we might consider that

$$N_k = \sum_c N_{k,c} = \sum_c \lambda_{k,c} \cdot R_{k,c} = \sum_c \frac{\lambda_{k,c} \cdot D_{k,c}}{1 - U_k} = \sum_c \frac{U_{k,c}}{1 - U_k} = \frac{U_k}{1 - U_k}.$$

Using the values for average utilizations, the average numbers of packets are:
$N_1 = U_1/(1 - U_1) = 0.7/0.3 = 2.34$ jobs,
$N_2 = U_2/(1 - U_2) = 0.8/0.2 = 4$ jobs,
$N_3 = U_3/(1 - U_3) = 0.9/0.1 = 9$ jobs,
$N_4 = U_4/(1 - U_4) = 0.9/0.1 = 9$ jobs.

For question 3, it holds that

$$R_k = \frac{N_k}{X} = \frac{N_k}{\sum_c X_c} = \frac{N_k}{\sum_c \lambda_c} = \frac{N_k}{\lambda_k},$$

as sums only have to consider the channels that actually use router k, so that
$R_1 = N_1/\lambda_1 = 2.34/7 = 0.34$ s,
$R_2 = N_2/\lambda_2 = 4/8 = 0.5$ s,
$R_3 = N_3/\lambda_3 = 9/18 = 0.5$ s,
$R_4 = N_4/\lambda_4 = 9/9 = 1$ s.

For question 4, we have that

$$R_{k,c} = \frac{D_{k,c}}{1 - U_k} = \frac{v_{k,c} \cdot S_{k,c}}{1 - U_k} = \frac{\frac{\lambda_{k,c}}{\lambda_c} \cdot S_{k,c}}{1 - U_k} = \frac{\frac{1}{\mu_k}}{1 - U_k},$$

because there is no loop in the model and there is no dependency on the channel in the final expression, as all average service rates for the routers are independent from the channel. As this is the only question in which an index is asked that explicitly depends on classes, and dependency on channels disappears, one might consider that classes are not needed, and that the system might be modeled by means of a simple queuing network ans solved by applying the separable models technique: this is anyway not viable, as routing would not be homogeneous, thus violating the third assumption.
In order to answer question 4 we obtain, for router 2, that is used by channels 1 and 3, the same value found while answering question 3 for R_2:

$R_{2,1} = 1/(\mu_2 \cdot (1 - U_2)) = 1/(10 \cdot (1 - 0.8)) = 0.5$ s,
$R_{2,3} = 1/(\mu_2 \cdot (1 - U_2)) = 1/(10 \cdot (1 - 0.8)) = 0.5$ s.

To answer question 6, we have to sum all the average response times for the routers that are used by channel 2, that are routers 1, 3 and 4: the overall average time spent in the systems by packets using channel 2 is $T = R_1 + R_2 + R_4 = 0.34 + 0.5 + 1 = 1.84$ s.

1.3 Using JMT

When creating the 4 classes, each should be assigned one of the sources shown in Fig. 1 as reference station, setting up the correspinding rate as in Fig. 2.

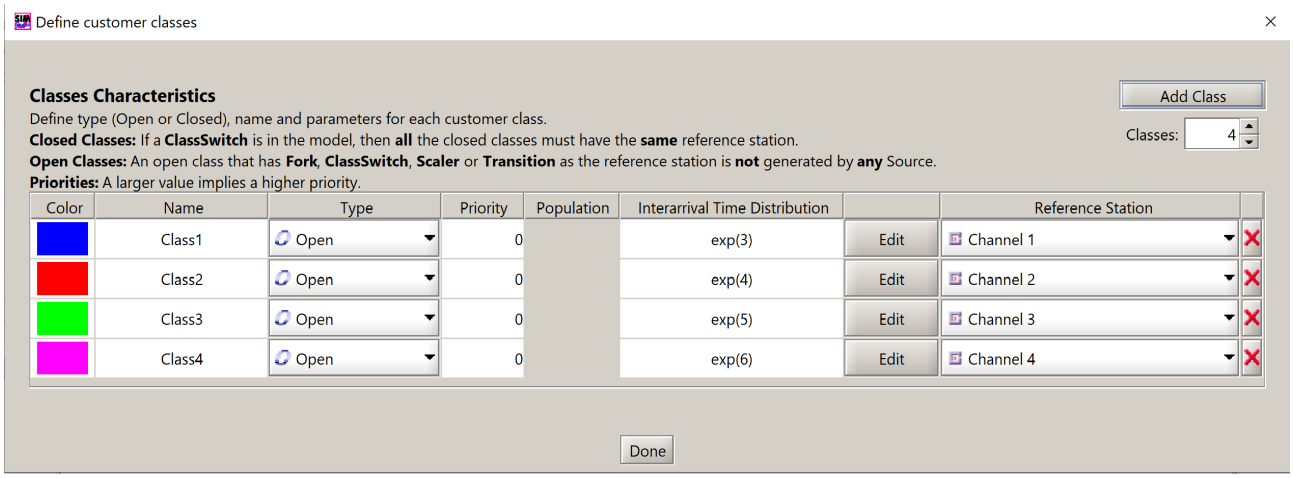

Figure 2: Classes setup

The channels have to be obtained by setting properly the routing section in the queues for each class. As channels have deterministic paths, each router should be set so that jobs from each class are directed to the right next router. This can be done by choosing a probabilistic routing with 100% probability of taking the right path between the physical connections between the routers. Examples for channels 1 and 2 at router 3 are in Fig. 3 and 4.

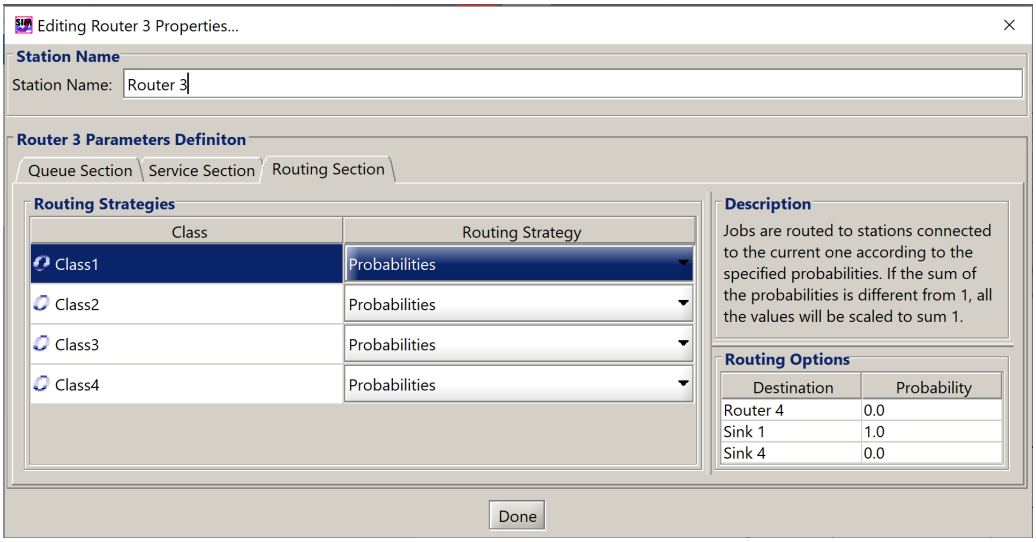

Figure 3: Setup of channel 1 at router 3

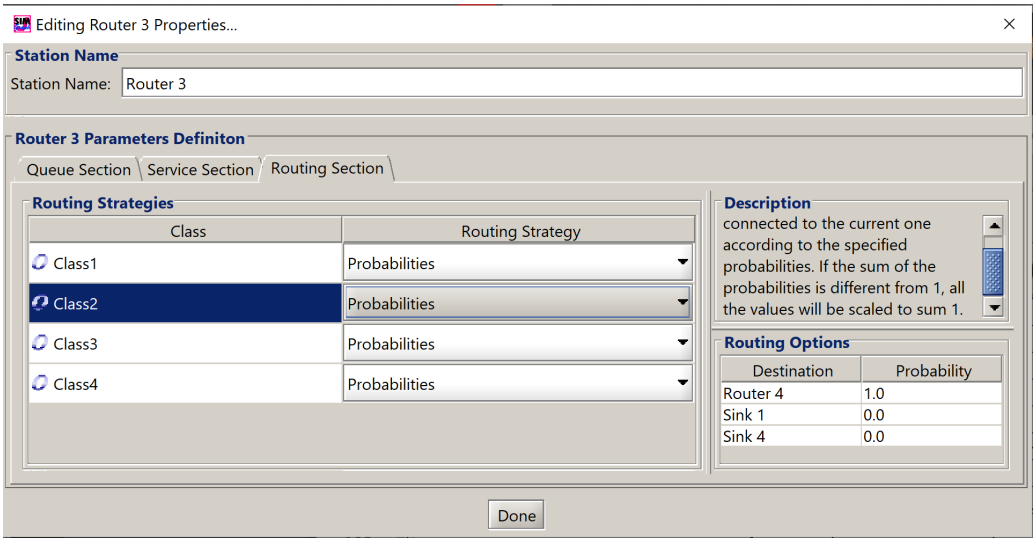

Figure 4: Setup of channel 2 at router 3

As routers process all packets equally for all channels, the service section setup for all classes of each router will be set up so that all classes have the same service rate (10 packets per second for routers 1, 2 and 4, 20 packets per second for router 3, exponentially distributed service times) as in Fig. 5.

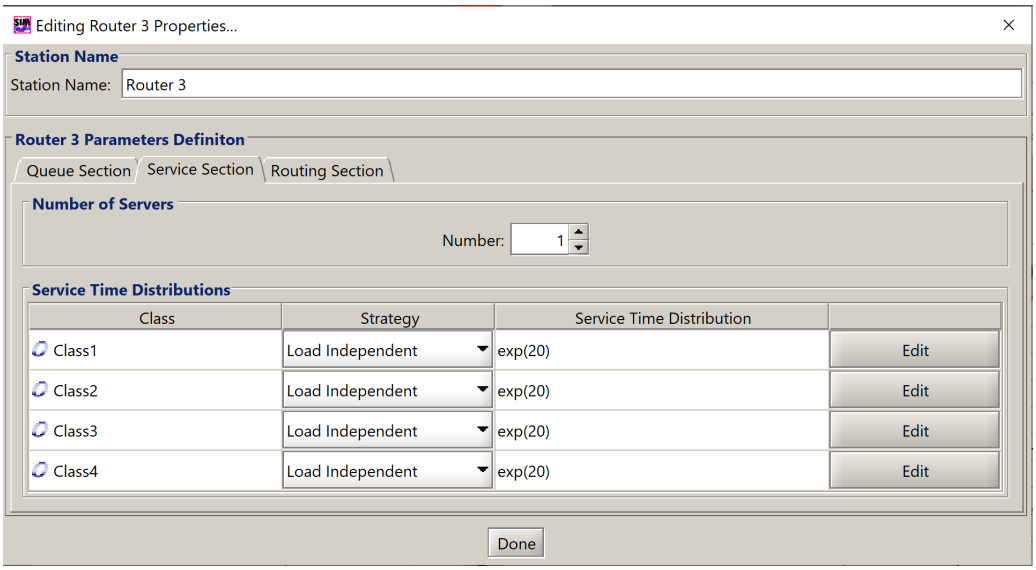

Figure 5: Setup of service times for router 3

The queue section for all routers is set to infinite server, and FCFS service policy for all channels.

In order to answer question 1, the 4 utilizations for the 4 routers with respect to all classes must be added in the performance indices panel. Similarly, to answer question 2, the 4 average number of customers for the 4 routers with respect to all classes must be added. Analogously, in order to answer question 3, the 4 response times for the 4 routers with respect to all classes must be added.

To answer question 4, response times for router 2 with respect to classes 1 and 3 must be added, as channels 2 and 4 do not use router 2.

In order to answer question 5, no additional preformance index is needed, as the bottleneck is the router with the highest utilization and utilizations have been computed to answer question 1.

Finally, to answer question 6, the 3 response times for routers 1, 3 and 4 with respect to class 2 must be added and summed.

The list of needed performance indices is in Fig. 6.

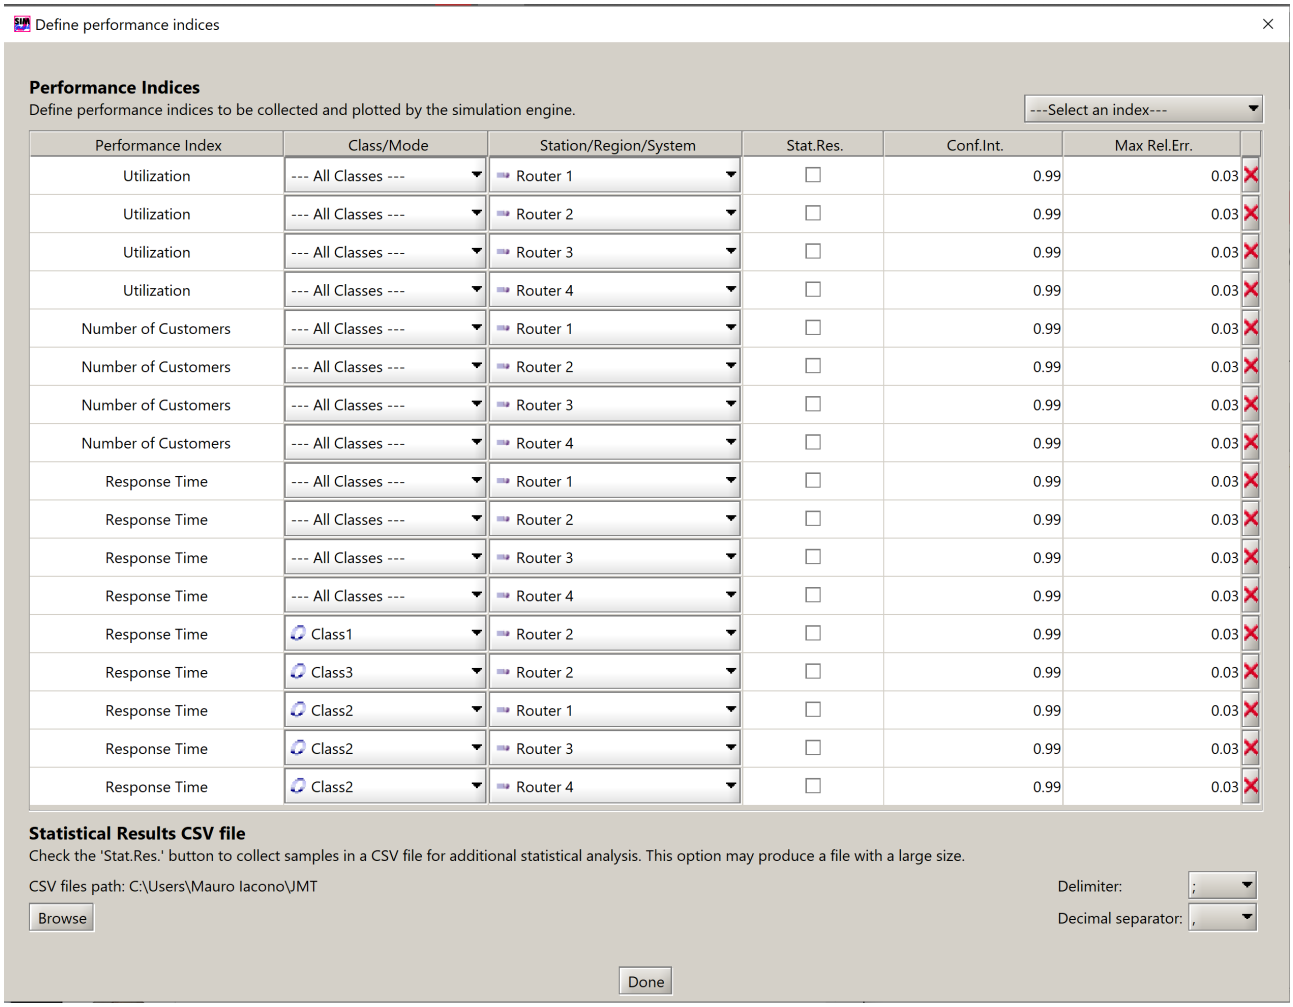

Figure 6: Setup of performance indices

After launching the simulation, some of the indices will not satisfy confidence intervals with standard settings. It is sufficient to augment the maximum number of events from 1000000 to 15000000 in the "Define simulation parameters" panel. If the simulation takes too much time to complete, check if numbers of customers indices all converge to a finite value: otherwise, routing might not be set correctly.

In the following, we obviously expect that the obtained results are coherent with the ones obtained by the analytical solution.

We obtain the values shown in Fig. 7 for the utilizations. All values are in the confidence interval. Utilization of router 1 is 0.70, utilization of router 2 is 0.78 (compatible with 0.8 obtained analytically, as simulation error is 0.03), utilization of router 3 is 0.90, while utilization of router 4 is 0.91 (compatible with 0.90, as simulation error is 0.03, and this answers question 1): consequently, the system bottleneck is router 4, but router 3 is basically in the same conditions of average utilization (this answers question 5, and the analytical results confirm that they both are equally the bottleneck).

We obtain the values shown in Fig. 8 for the average numbers of packets. For router 1 we obtain 2.38, for router 2 we obtain 4.03, for router 3 we obtain 9.08, for router 4 we obtain 9.02 (values for router 1 and 3 are not compatible with the ones obtained by the analytical solution, anyway within a small interval): this answers question 2.

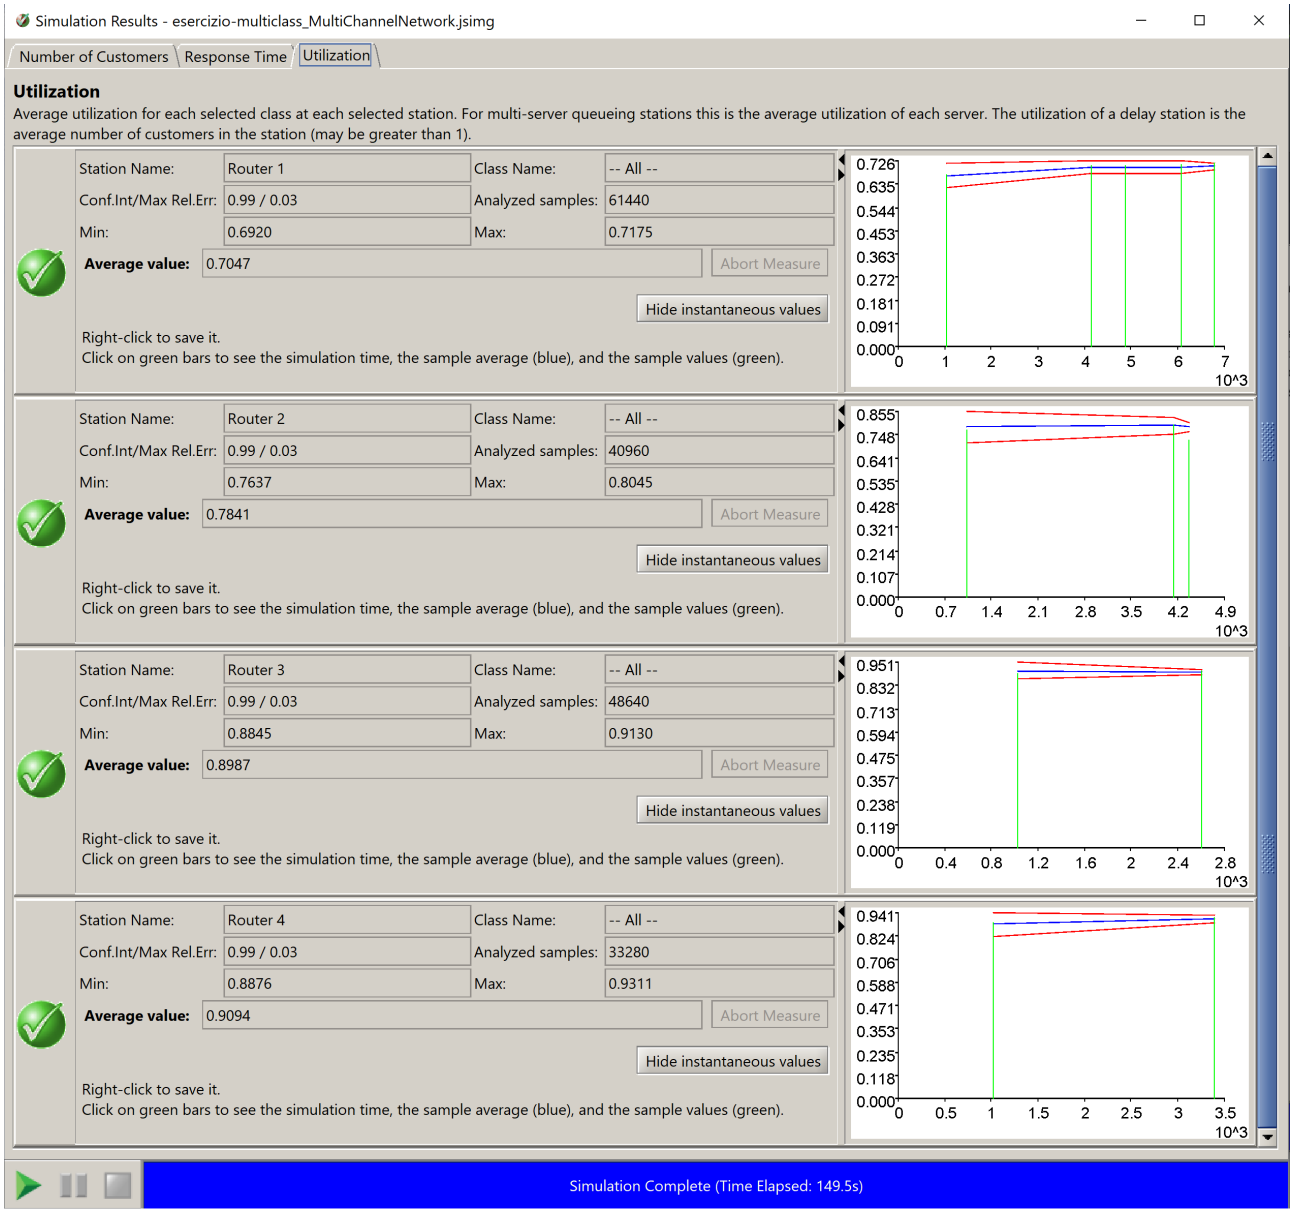

Figure 7: Average utilizations

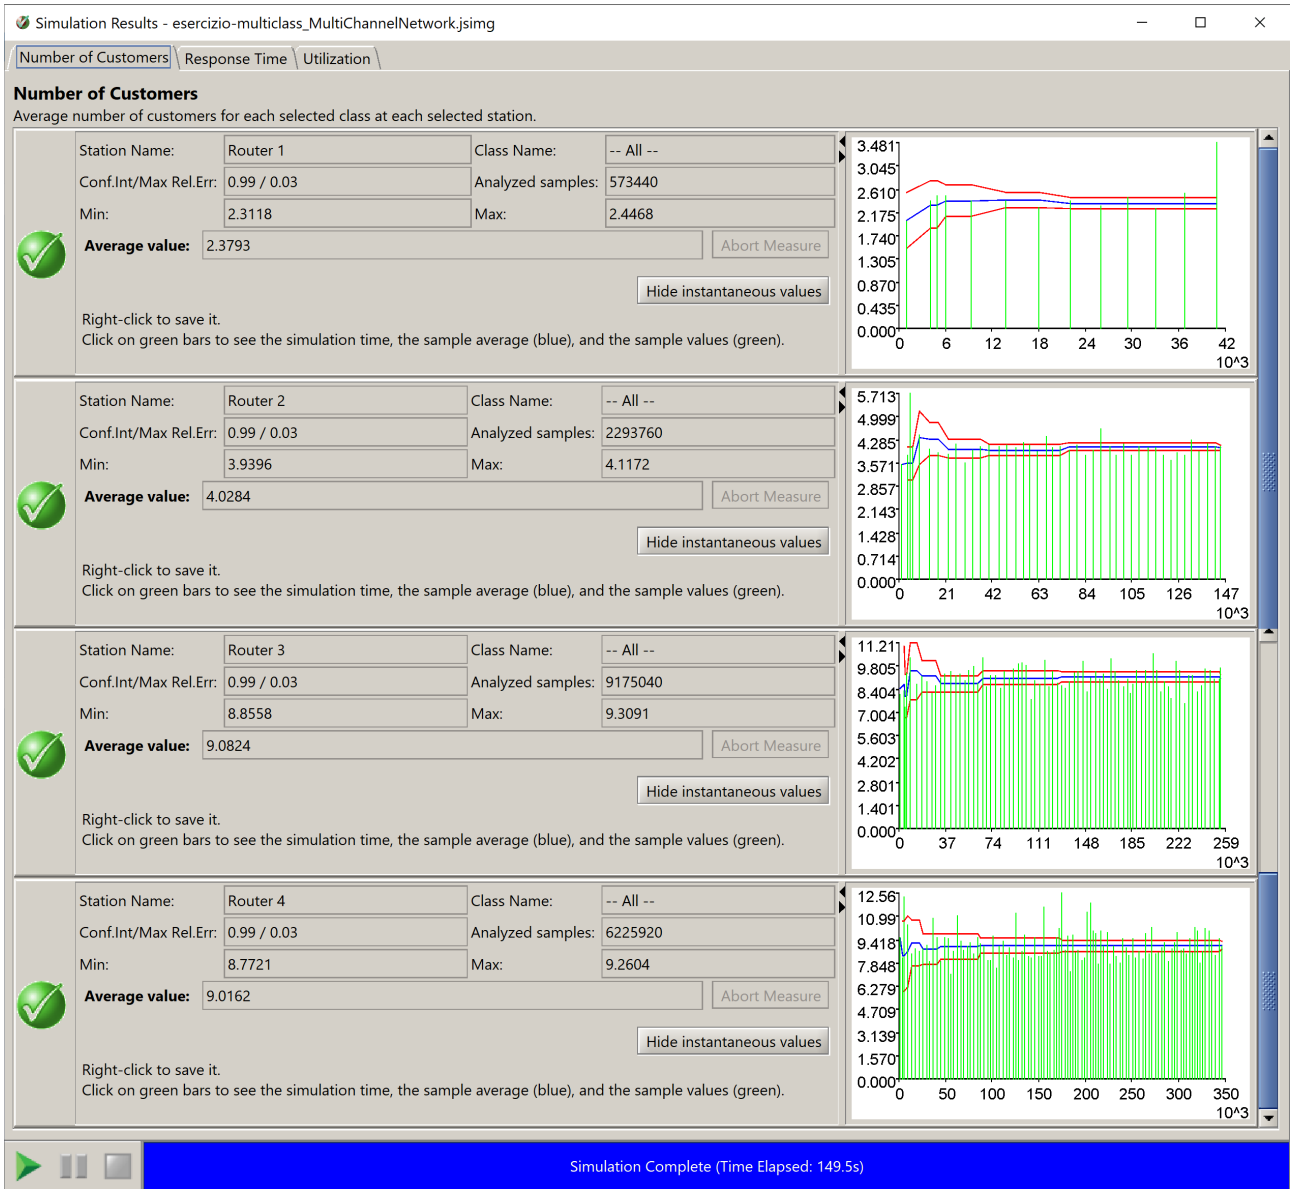

Figure 8: Average numbers of customers

Response times for the routers with respect to all classes are in Fig. 9. Response times for router 1, 2, 3 and 4 are respectively 0.34 s, 0.50 s, 0.50 s and 1.00 s: this answers question 3.

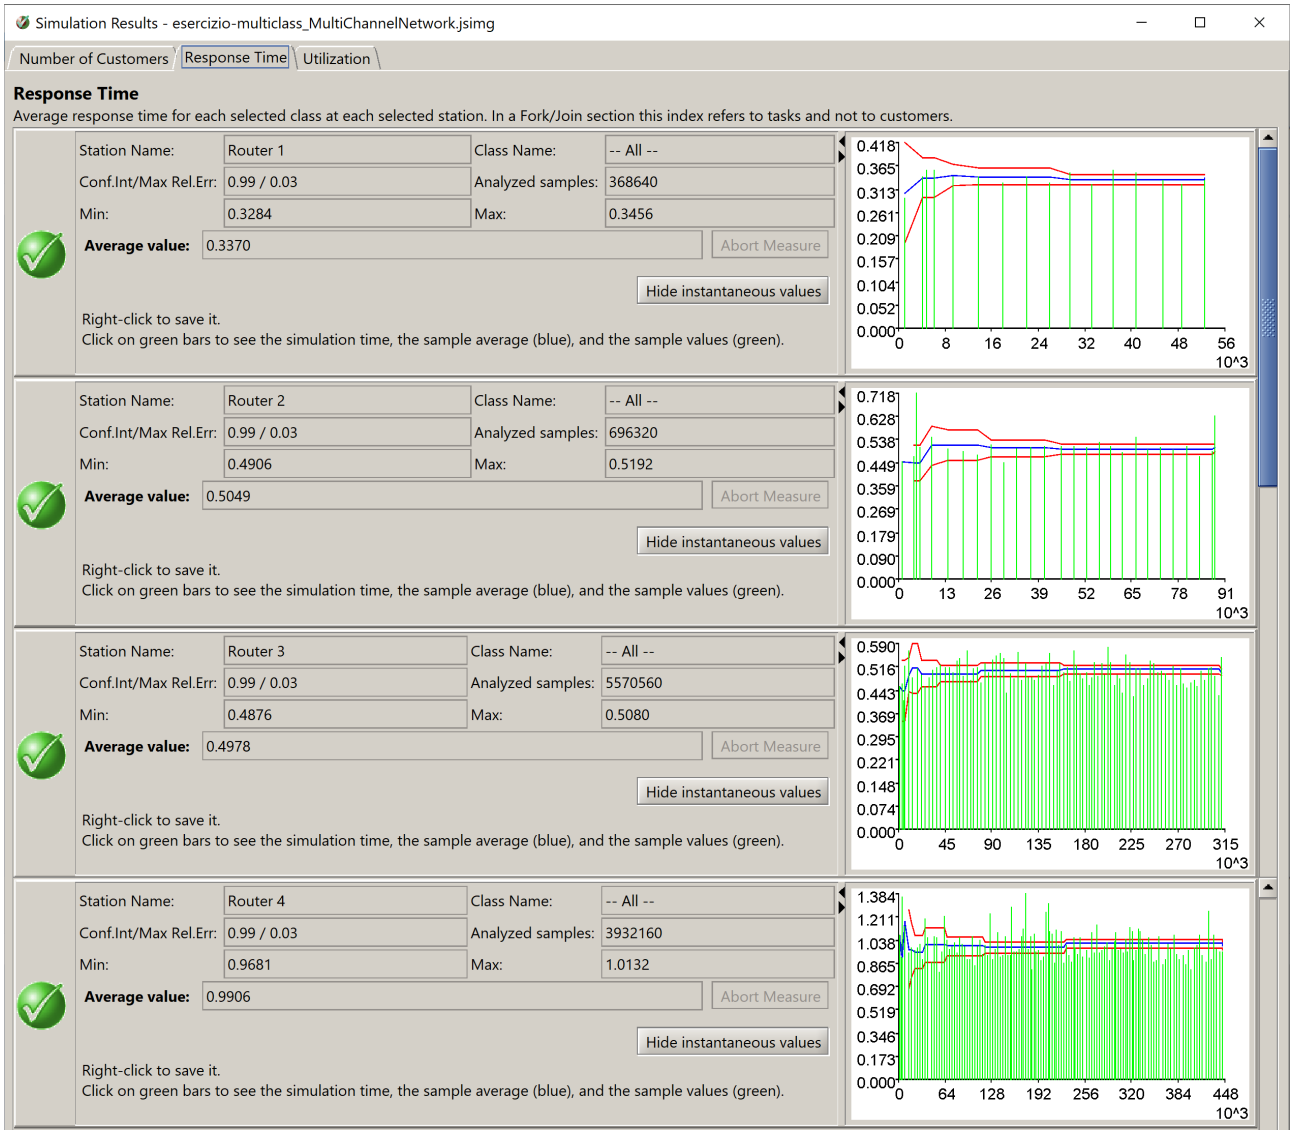

Figure 9: Average response times - question 3

Response times for router 2 with respect to channel 1 and 3 are in the upper part of Fig. 10 (channels 2 and 4 do not use it). Obviously they have the same value of 0.50 s, as we know from previous considerations: this answers question 4.

As for the lower part of Fig. 10, response times for routers 2, 3 and 4 with respect to channel 2 are respectively 0.34 s, 0.51 s and 0.99 s, for an average total time in the system for packets using virtual channel 2 of 1.84 s: this answers question 6.

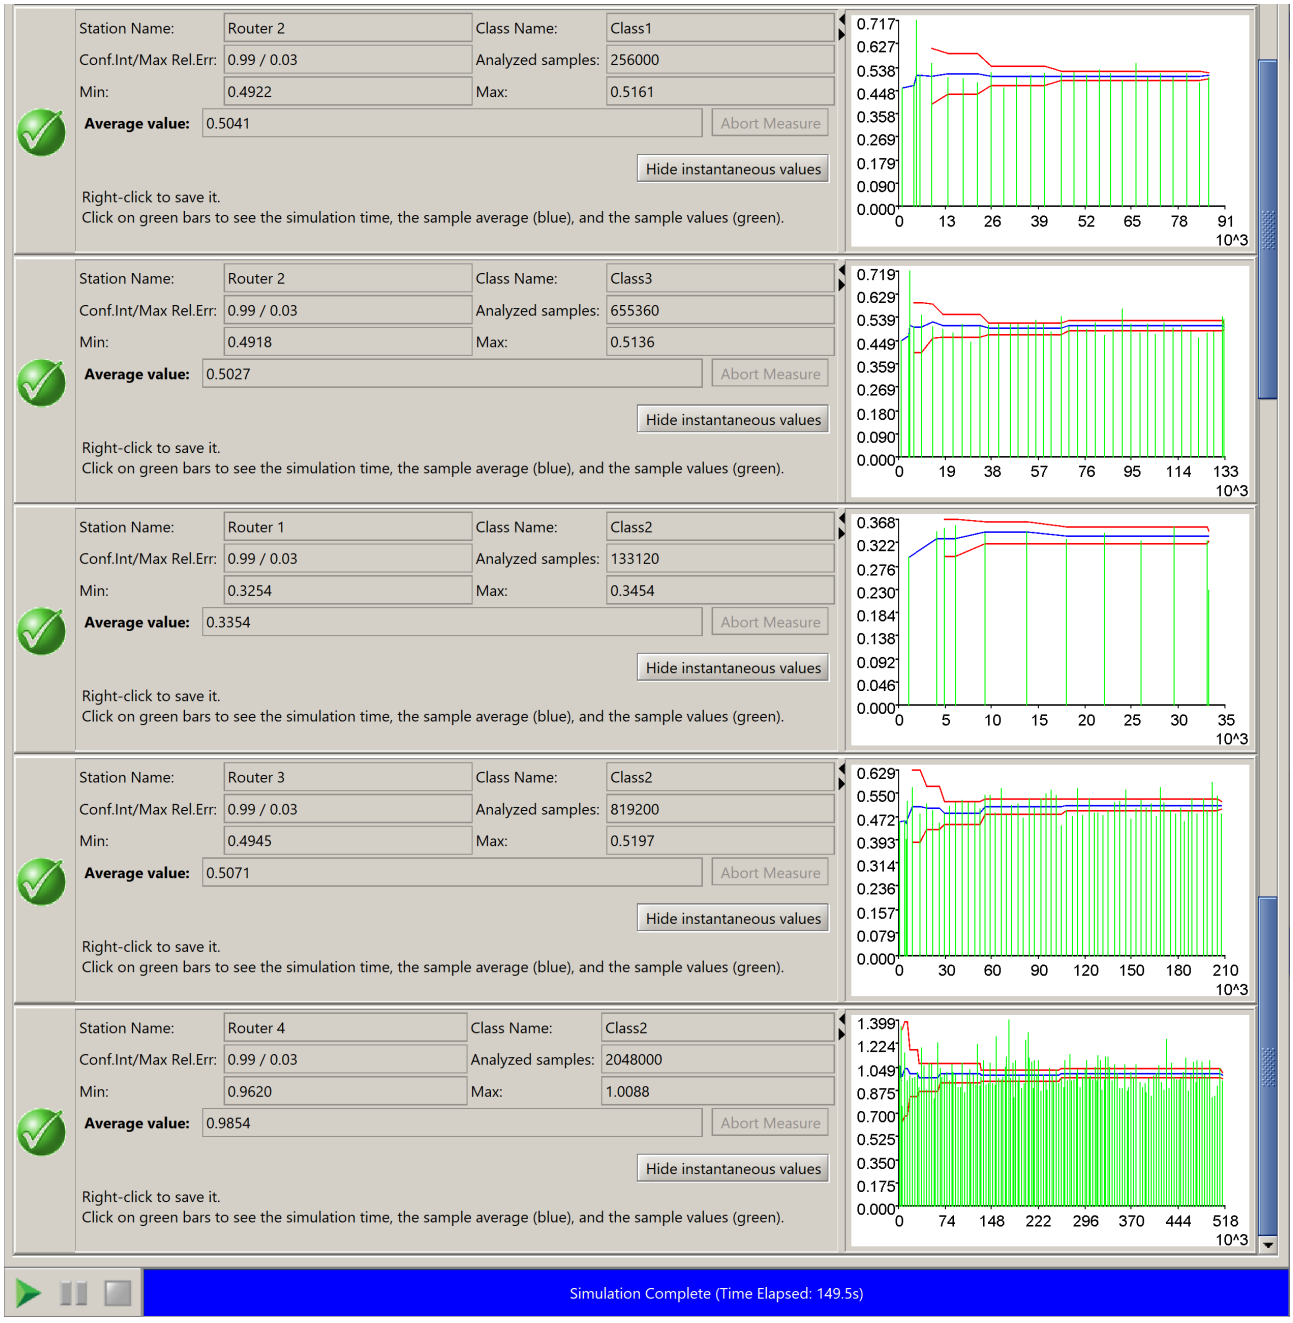

Figure 10: Average response times - questions 4 and 6

2 Exercise: CPU-bound and I/O-bound[2]

A system consists of two devices: a CPU device with exponential service rate 2 jobs/s and an I/O device with exponential service rate 1 job/s. There are two different types of jobs: CPU-bound jobs and I/O-bound jobs.

CPU-bound jobs arrive at the CPU from outside according to a Poisson process of rate 0.2 jobs/s. After serving at the CPU, three things can happen to a CPU-bound job:

1. with probability 0.3, the job leaves the system;
2. with probability 0.65, the job returns to the CPU queue to repeat the process;
3. with probability 0.05, the job goes to the I/O device queue, serves there once, and immediately returns to the CPU queue to repeat the process.

The I/O-bound jobs arrive at the I/O from outside the network according to a Poisson process with rate 0.25 jobs/s. After serving at the I/O, there are three things that can happen to an I/O-bound job:

1. with probability 0.4, the job leaves the system;
2. with probability 0.5, the job returns to the I/O queue to repeat the process;
3. with probability 0.1, the job goes to the CPU device queue; each time the job serves at the CPU device, it has a 0.05 probability of returning to the CPU device and a 0.95 probability of returning to the I/O queue.

Our goal is to answer the following questions:

1. what is the expected time in system of CPU-bound jobs?
2. what is the average number of CPU-bound jobs at the CPU?
3. if the service policy of the CPU is processor sharing and the service policy of the I/O is shortest job first, how do we expect the behavior of the CPU-bound and I/O bound number of customers at the CPU and the I/O and their total time spent in the system will change? Why? Verify if your guess is confirmed by results;
4. in order to use this system for real time applications, the CPU is scheduled with a discriminatory processor sharing (DPS) policy that allocates twice the CPU to CPU-bound jobs with respect to I/O-bound jobs. The deadline for CPU-bound jobs is 3 seconds: is this choice sufficient to satisfy this requirement? Comment results, and compare them to previous cases. If it is not sufficient, what choice may be sufficient?

2.1 Solution

The two devices, namely the CPU and the I/O device, are used by both the types of jobs, and can be represented by two queuing stations. As the routing probabilities are given and depend on the type of job, a classed queuing network with a class per type is a viable model for the system. Both stations are characterized by exponentially distributed service times, that are provided in the form of service rates $\mu_{k,c}$ equal for both classes at each station, so that an analytical solution may be applied if the system is stable.

The system can be modeled as in Fig. 11.

[2] This exercise is based on the exercise presented in [1] in subsection 18.6.3, page 326.

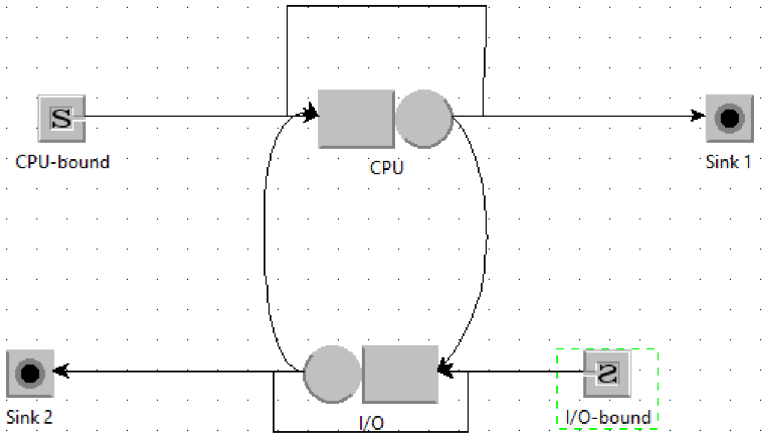

Figure 11: The model

2.2 Analytical solution

The two questions are about CPU-bound jobs, so

$$R_c = \sum_k R_{k,c} = \sum_k \frac{D_{k,c}}{1 - \sum_{c'} U_{k,c'}} = \sum_k \frac{v_{k,c} \cdot s_{k,c}}{1 - \sum_{c'} U_{k,c'},} = \sum_k \frac{\frac{\lambda_{k,c}}{\lambda_c \cdot \mu_{k,c}}}{1 - \sum_{c'} U_{k,c'}}$$

The equations to find $\lambda_{1,CPU-bound}$ and $\lambda_{2,CPU-bound}$ are obtained by inspecting the model and are:

$$\begin{cases} \lambda_{1,CPU-bound} = \lambda_{0,CPU-bound} + \lambda_{1,CPU-bound} \cdot p_{1,1,CPU-bound} + \lambda_{2,CPU-bound} \cdot p_{2,1,CPU-bound} = 0.2 + 0.65\lambda_{1,CPU-bound} \\ \lambda_{2,CPU-bound} = \lambda_{1,CPU-bound} \cdot p_{1,2,CPU-bound} = 0.05\lambda_{1,1,CPU-bound} \end{cases}$$

$$\begin{cases} 0.35\lambda_{1,CPU-bound} = \lambda_{2,CPU-bound} + 0.2 \\ \lambda_{2,CPU-bound} = 0.05\lambda_{1,CPU-bound} \end{cases}$$

$$\begin{cases} 0.3\lambda_{1,CPU-bound} = 0.2 \\ \lambda_{2,CPU-bound} = 0.05\lambda_{1,CPU-bound} \end{cases}$$

$$\begin{cases} \lambda_{1,CPU-bound} = 0.666\ jobs/s \\ \lambda_{2,CPU-bound} = 0.033\ jobs/s \end{cases}$$

Analogously, the equations to find $\lambda_{1,I/O-bound}$ and $\lambda_{2,I/O-bound}$ are obtained by inspecting the model and are:

$$\begin{cases} \lambda_{1,I/O-bound} = \lambda_{1,I/O-bound} \cdot p_{1,1,I/O-bound} + \lambda_{2,I/O-bound} \cdot p_{2,1,I/O-bound} = 0.05\lambda_{1,I/O-bound} + 0.1\lambda_{2,I/O-bound} \\ \lambda_{2,I/O-bound} = \lambda_{0,I/O-bound} + \lambda_{1,I/O-bound} \cdot p_{1,2,I/O-bound} + \lambda_{2,I/O-bound} \cdot p_{2,2,I/O-bound} = 0.25 + 0.95\lambda_{1,I/O-bound} + 0.5 \end{cases}$$

$$\begin{cases} 0.95\lambda_{1,I/O-bound} = 0.1\lambda_{2,I/O-bound} \\ 0.5\lambda_{2,I/O-bound} = 0.25 + 0.95\lambda_{1,I/O-bound} \end{cases}$$

$$\begin{cases} \lambda_{1,I/O-bound} = 0.105\lambda_{2,I/O-bound} \\ 0.5\lambda_{2,I/O-bound} = 0.25 + 0,1\lambda_{2,I/O-bound} \end{cases}$$

$$\begin{cases} \lambda_{1,I/O-bound} = 0.066\ jobs/s \\ \lambda_{2,I/O-bound} = 0.625\ jobs/s \end{cases}$$

We can compute now the average utilizations for the devices by the CPU-bound tasks as $U_{1,CPU-bound} = \frac{\lambda_{1,CPU-bound}}{\mu_{1,CPU-bound}} = \frac{0.666}{2} = 0.3331$ and $U_{2,CPU-bound} = \frac{\lambda_{2,CPU-bound}}{\mu_{2,CPU-bound}} = \frac{0.033}{1} = 0.0331$, and by the I/O-bound tasks as $U_{1,I/O-bound} = \frac{\lambda_{1,I/O-bound}}{\mu_{1,I/O-bound}} = \frac{0.066}{2} = 0.0331$ and $U_{2,I/O-bound} = \frac{\lambda_{2,I/O-bound}}{\mu_{2,I/O-bound}} = \frac{0.625}{1} = 0.6251$, concluding that the system is stable and can be solved analytically.

Applying the formula found for R_c to CPU-bound tasks we obtain

$$R_{CPU-bound} = R_{1,CPU-bound} + R_{2,CPU-bound} = 2.626 + 0.482 = 3.108\ s,$$

since

$$R_{1,CPU-bound} = \frac{\frac{\lambda_{1,CPU-bound}}{\lambda_{0,CPU-bound} \cdot \mu_{1,CPU-bound}}}{1 - U_{1,CPU-bound} - U_{1,I/O-bound}} = \frac{\frac{0.666}{0.2 \cdot 2}}{1 - 0.333 - 0.033} = 2.626\ s$$

and

$$R_{2,CPU-bound} = \frac{\frac{\lambda_{2,CPU-bound}}{\lambda_{0,CPU-bound} \cdot \mu_{2,CPU-bound}}}{1 - U_{2,CPU-bound} - U_{2,I/O-bound}} = \frac{\frac{0.033}{0.2 \cdot 1}}{1 - 0.033 - 0.625} = 0.482\ s,$$

that answers question 1.

In order to answer question 2, we can consider that

$$N_{k,c} = X_c \cdot R_{k,c} = \frac{X_{k,c}}{v_{k,c}} R_{k,c} = \frac{\lambda_{k,c}}{\frac{\lambda_{k,c}}{\lambda_c}} R_{k,c} = \lambda_c \cdot R_{k,c}$$

that, applied to the CPU for CPU-bound tasks, gives an average number of tasks equal to

$$N_{1,CPU-bound} = \lambda_{CPU-bound} \cdot R_{1,CPU-bound} = 0.2 \cdot 2.626 = 0.5252.$$

As the shortest job first policy is used for the I/O device, question 3 cannot be answered by means of the analytical methods presented in the course. Anyway, from a qualitative point of view, we may expect that the effect on the average number of CPU-bound jobs at the CPU will not change significantly, due to the fact that the fraction of I/O-bound jobs that uses the CPU, and that may take advantage of the processor sharing policy with respect to CPU-bound jobs, is small and that they come from the I/O device, that has a lower service rate: we also expect that the average number of I/O-bound jobs at the CPU will be lower, because they will be served faster. For the average number of jobs at the I/O device, we expect a small improvement, because shorter workloads will be promoted to be executed first and will not suffer the operations of longer workloads. For the same reason, we expect shorter average time in the system for both the kind of jobs, with a higher improvement for I/O-bound jobs because they are prevalent in the workload of the I/O device. These hypotheses may be verified by simulation.

Analogously, the analytical approach is not suitable to answer question 4.

2.3 Using JMT

When creating the 2 classes, each should be assigned one of the sources shown in Fig. 11, setting up the correspinding rate as in Fig. 12.

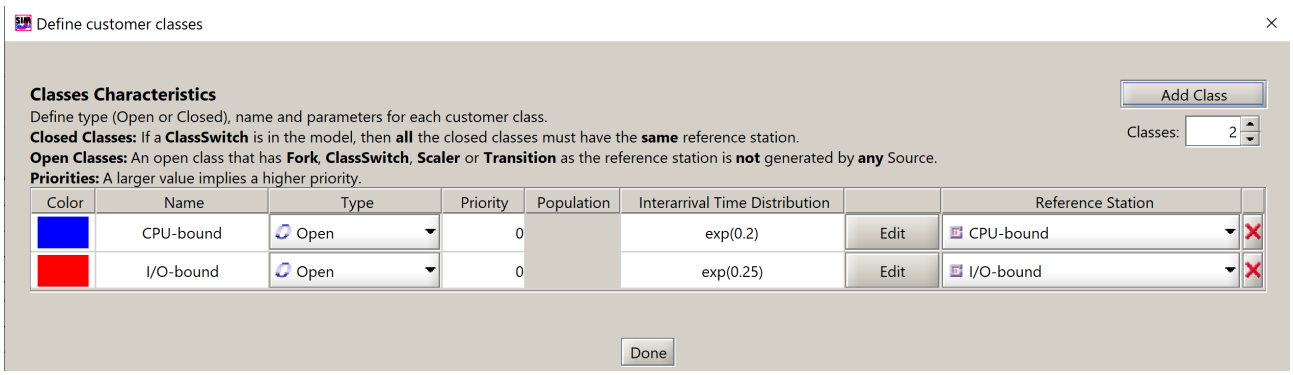

Figure 12: Classes setup

The routing section in the queues can be easily set probabilistically, as described by the specification for each class. Examples for both type of jobs at the CPU are in Fig. 13 and 14.

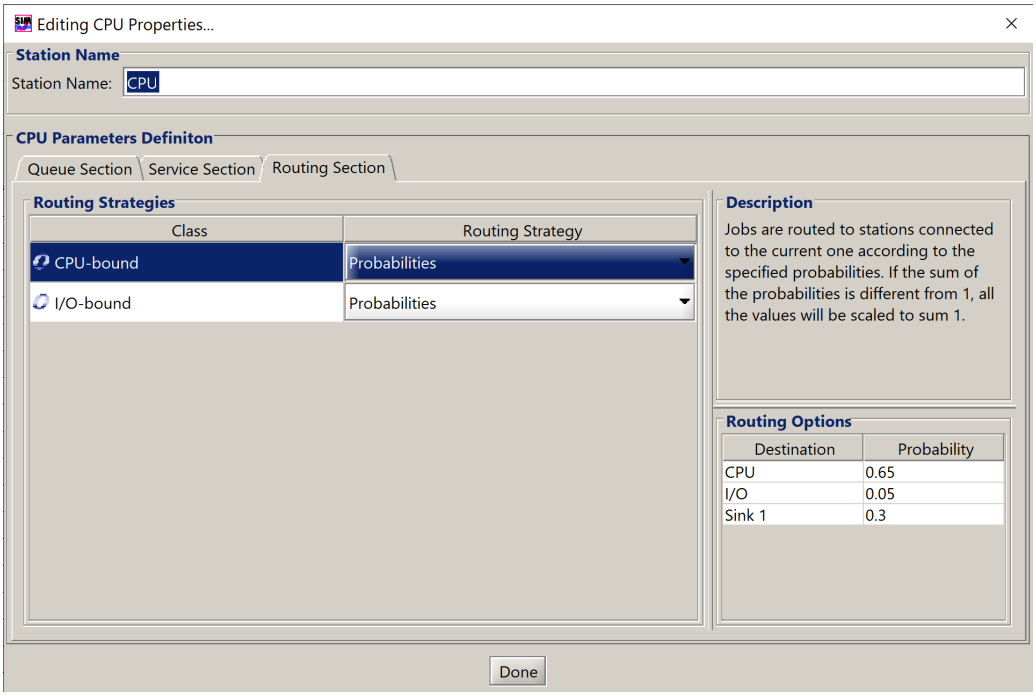

Figure 13: Setup of CPU-bound jobs routing at the CPU

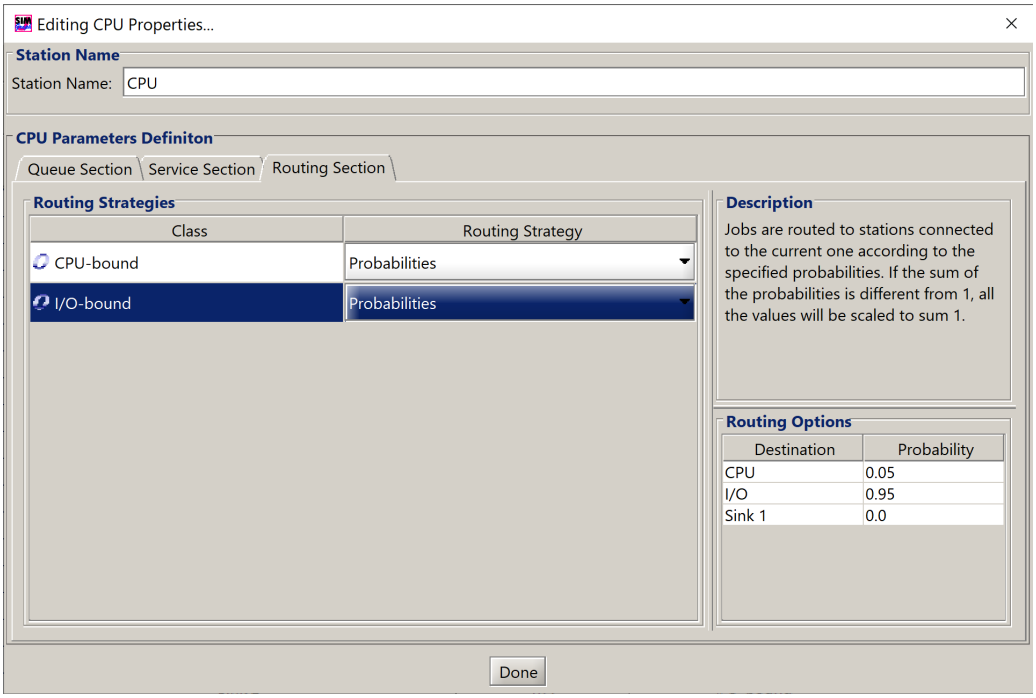

Figure 14: Setup of I/O-bound jobs routing at the CPU

As devices process all jobs with the same rates, the service section setup for all classes of each queuing station will be set up so that all classes have the same service rate (2 jobs per second for the CPU and 1 job per second for the I/O device, exponentially distributed service times) as in Fig. 15.

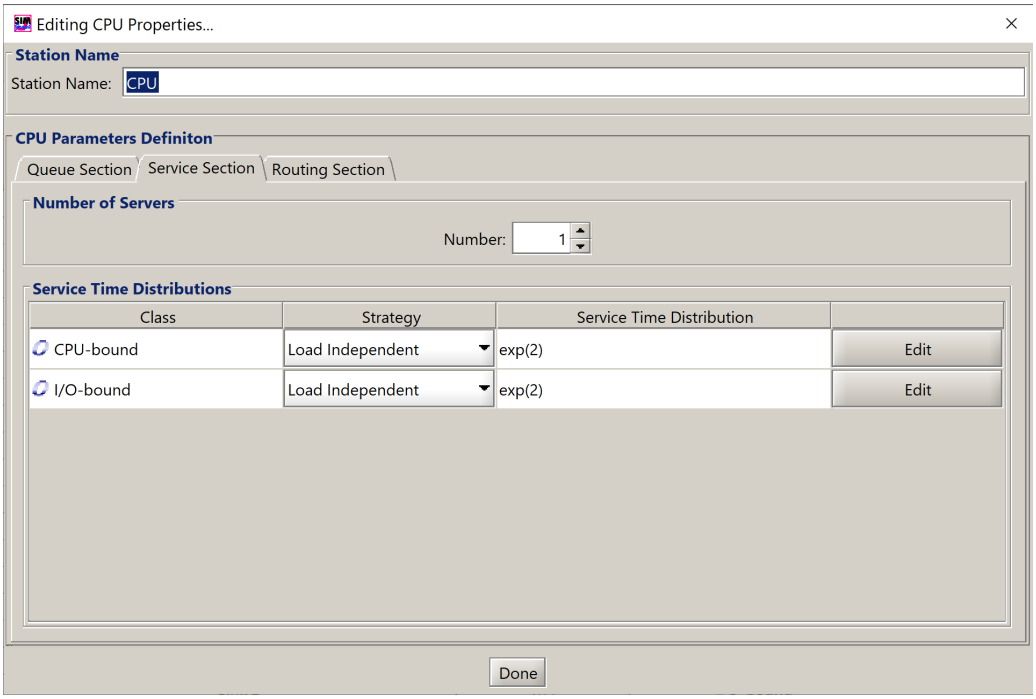

Figure 15: Setup of service times for the CPU

The queue section for all queuing station is set to infinite server, and FCFS service policy for all classes.

In order to answer question 1, the system response time for class CPU-bound must be added in the performance indices panel. Similarly, to answer question 2, the average number of customers for the CPU routers with respect to class CPU-bound must be added.

The list of needed performance indices is in Fig. 16.

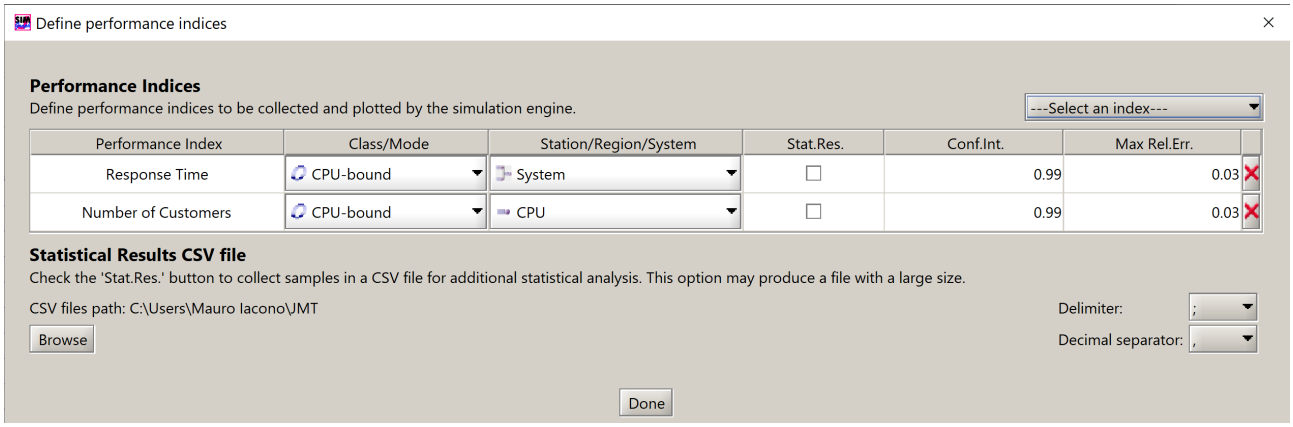

Figure 16: Setup of performance indices

After launching the simulation, we obviously expect that the obtained results are coherent with the ones obtained by the analytical solution.

We obtain the value shown in Fig. 17 for the expected time in system of CPU-bound jobs. The value is in the confidence interval. The result is 3.18, not compatible with 3.10 obtained analytically but sufficiently close, as simulation error is 0.03.

We obtain the value shown in Fig. 18 for the average numbers of CPU-bound jobs at the CPU. The value is in the confidence interval. The result is 0.53, as obtained analytically.

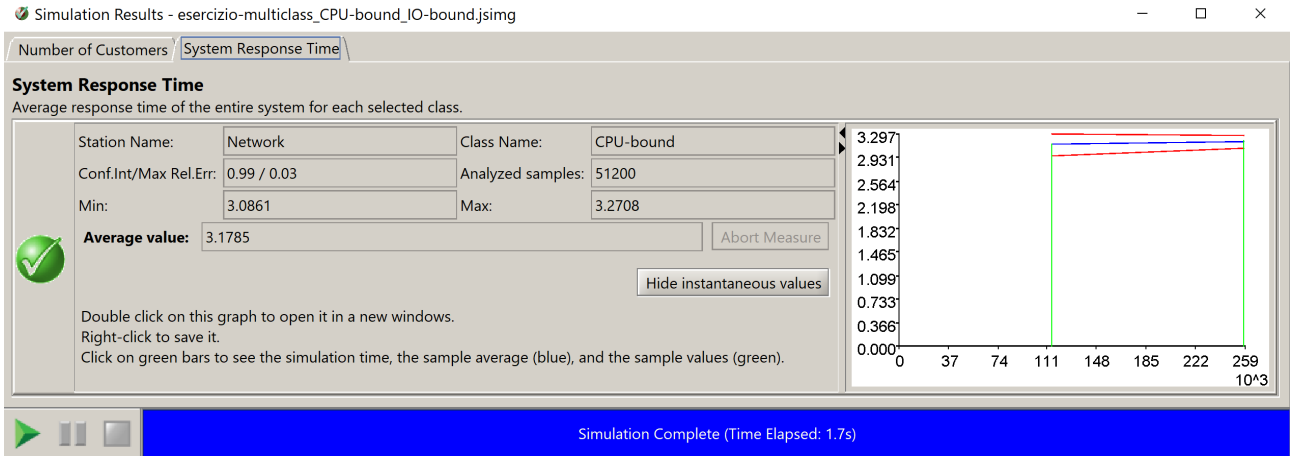

Figure 17: System response time for class CPU-bound

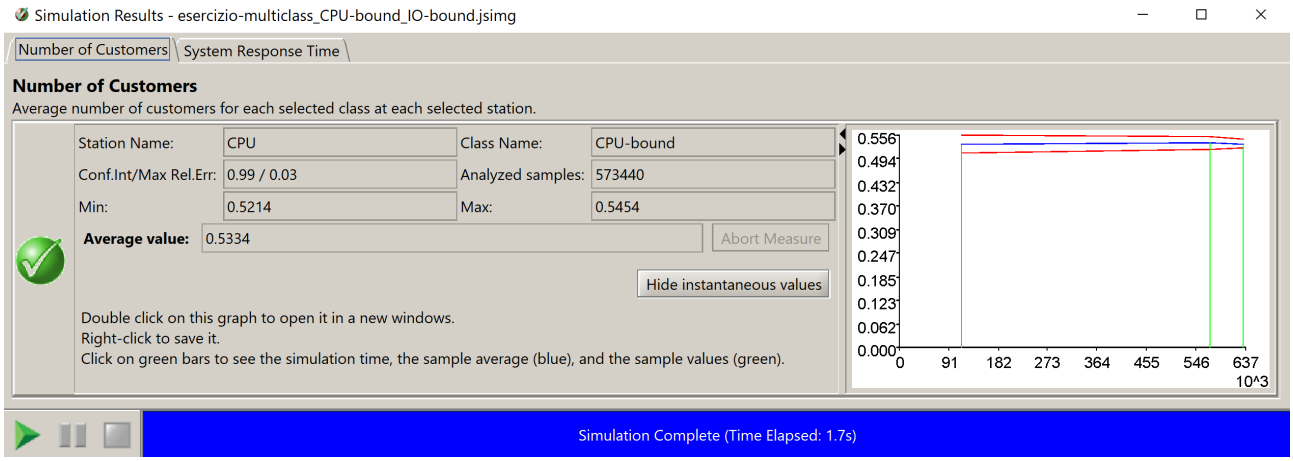

Figure 18: Average number of customers at CPU for CPU-bound jobs

In order to answer question 3, the average number of I/O-bound jobs at the CPU and of I/O-bound and CPU-bound at the I/O, and the system response time for I/O-bound jobs must be computed, as shown in Fig. 19 and Fig. 20.

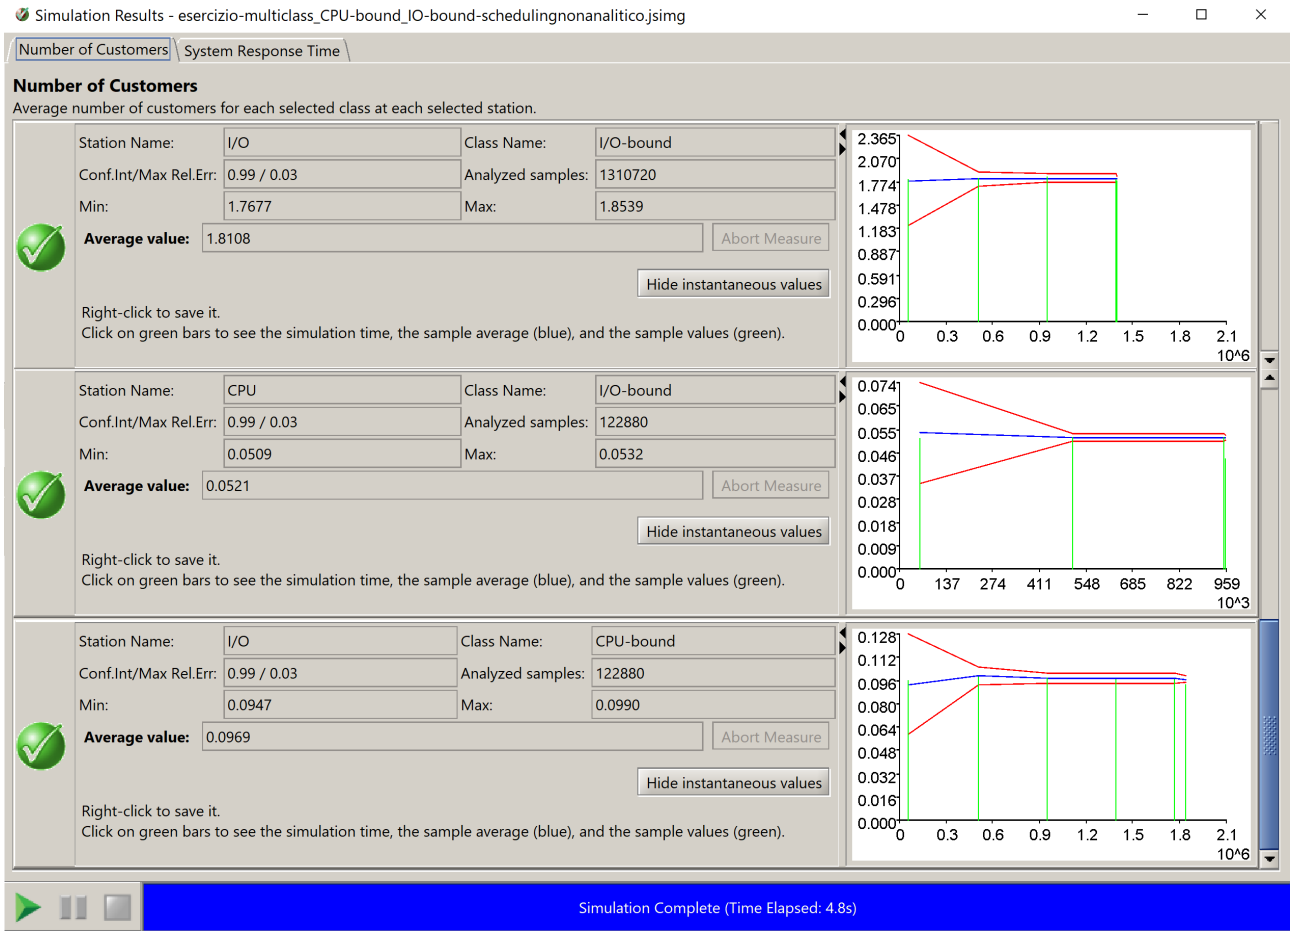

Figure 19: Average number of customers at CPU for I/O-bound jobs and at I/O for CPU-bound and I/O-bound jobs

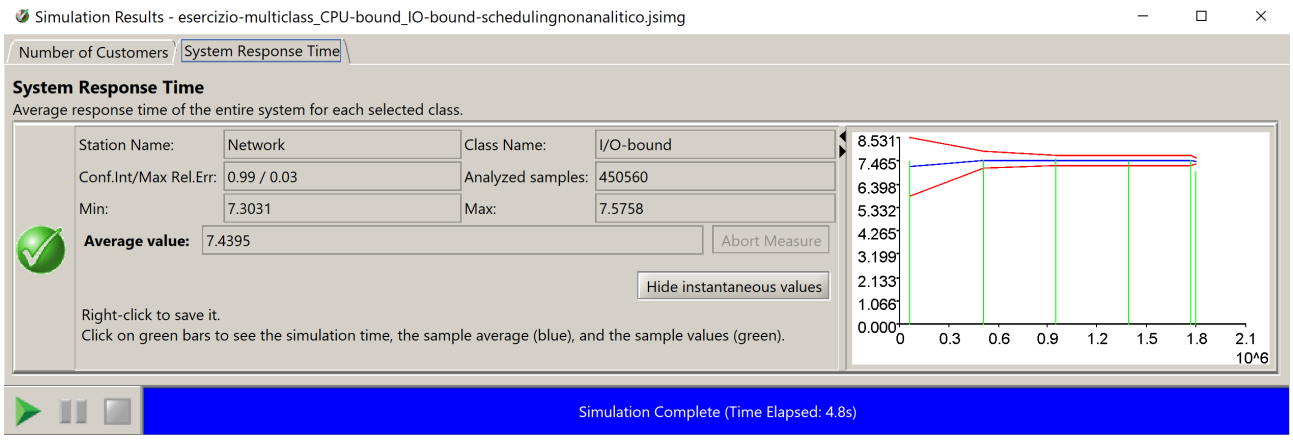

Figure 20: System response time for class I/O-bound

The service discpline of the CPU must be set to PS for both classes, that is accessible in the Preemptive Scheduling set in the Queue Section tab, and the service discipline of the I/O must be set to SJF for both classes. The new results are in Fig. 21 and Fig. 22.

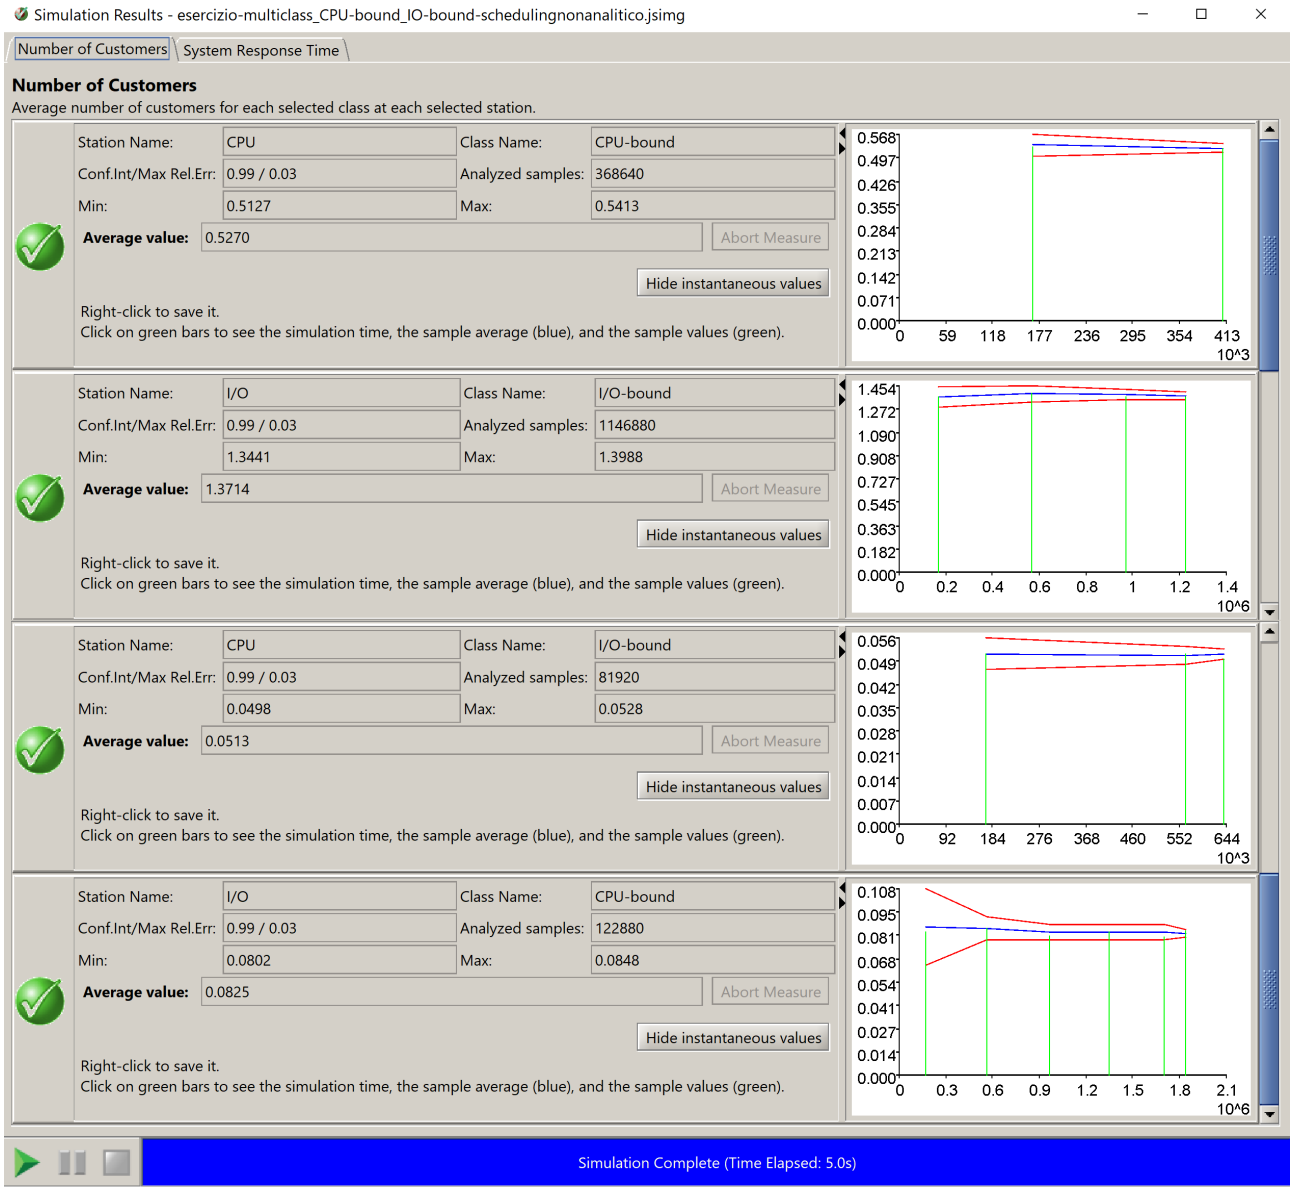

Figure 21: Average number of customers with PS and SJF policies

Results are summarized and compared in the following table:

	FCFS+FCFS	PS+SJF
$N_{CPU,CPU-bound}$	0.527	0.527
$N_{I/O,I/O-bound}$	1.811	1.371
$N_{CPU,I/O-bound}$	0.052	0.051
$N_{I/O,CPU-bound}$	0.097	0.083
$R_{CPU-bound}$	3.126	3.048
$R_{I/O-bound}$	7.440	5.727

and confirm our hypotheses.

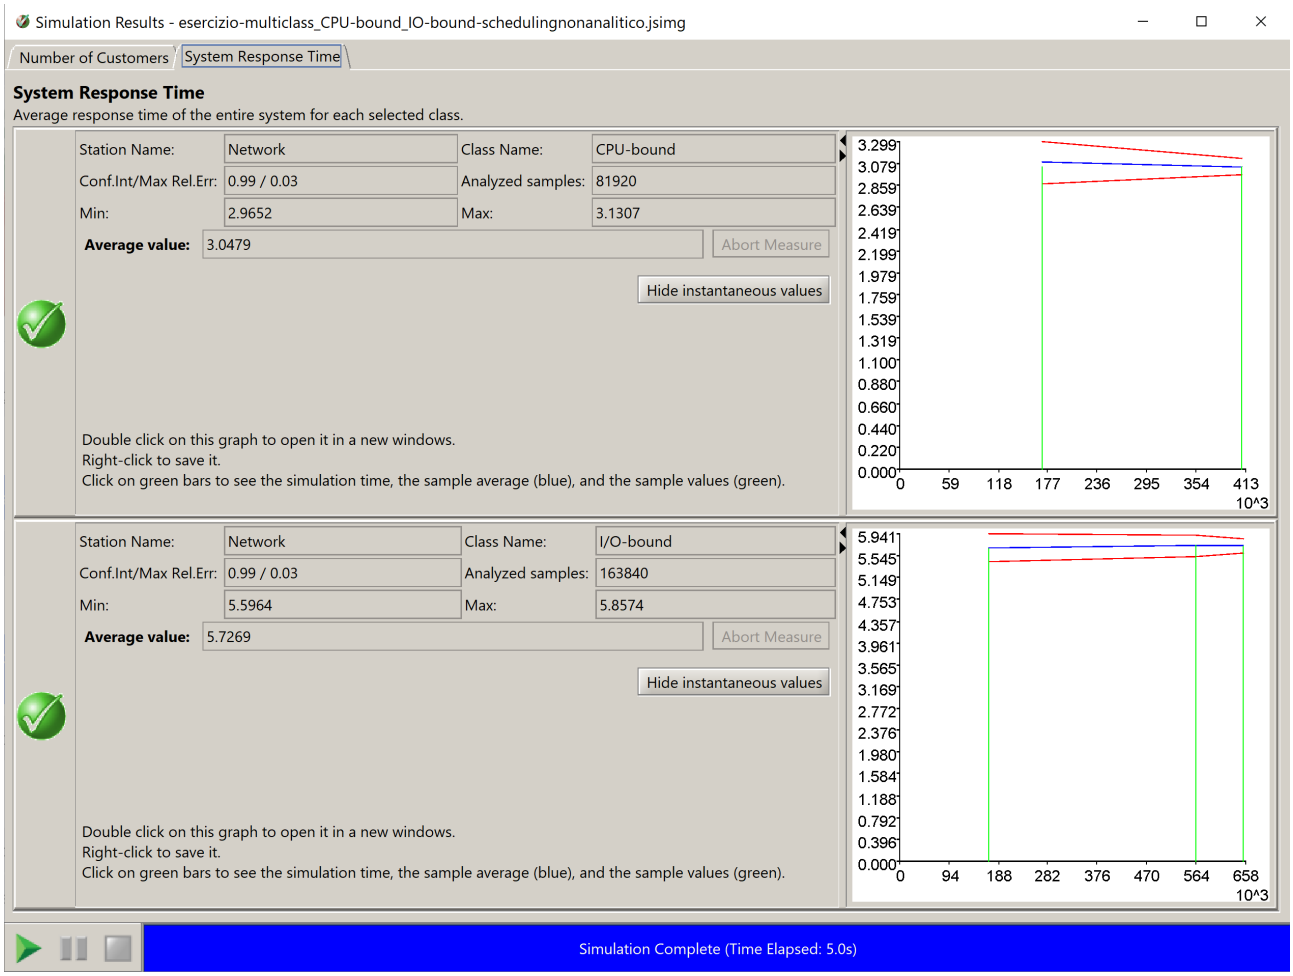

Figure 22: System response times with PS and SJF policies

To answer question 4, it is sufficient to change the queue policy to DPS, with weights 2.0 and 1.0 for CPU-bound and I/O bound jobs, respectively. With this setting, simulation produces the values in column DPS in the following table, that extends the previous one:

	FCFS+FCFS	PS+SJF	DPS+SJF
$N_{CPU,CPU-bound}$	0.527	0.527	0.524
$N_{I/O,I/O-bound}$	1.811	1.371	1.377
$N_{CPU,I/O-bound}$	0.052	0.051	0.063
$N_{I/O,CPU-bound}$	0.097	0.083	0.082
$R_{CPU-bound}$	3.126	3.048	2.980
$R_{I/O-bound}$	7.440	5.727	5.718

As first, checking $R_{CPU-bound}$ we note that the value is 2.980 s, that is less than the 3 s requirement: consequently, this choice is not sufficient, because that value is affected by a 3% error, thus it is 2.980 $\pm$ 0.03, that gives an upper bound of 3.01, and anyway would be too close to the deadline to be viable in practice, considering that it is an average value of a distribution. We also note that, notwithstanding the higher priority (that is assigning, with DPS, twice the CPU time to CPU-bound jobs with respect to I/O jobs), the improvement in the average number of CPU-bound jobs at the CPU is negligible, as well as the average number of I/O-bound ones, with respect to the PS case, as the contribution of the I/O bound ones to the CPU workload is small. The average number of I/O-bound jobs at the I/O is also unaltered, while the average number of CPU-bound ones is significantly higher. The impact of the new setting is also further, even if slightly, improving the value of $R_{I/O-bound}$.

References

Harchol-Balter, M. (2013). *Performance modeling and design of computer systems: Queueing theory in action.* Cambridge University Press. doi: 10.1017/CBO9781139226424