Advanced Systems Lab Report

Autumn Semester 2019

Name: YOUR_NAME
Legi: YOUR_LEGI

October 10, 2019

Notes on writing the report **(remove this page for submission)**

The report for the lab needs not be extensive but it must be concise, complete, and correct. Conciseness is important in terms of content and explanations, focusing on what has been done and explanations of the results. A long report is not necessarily a better report, especially if there are aspects of the design or the experiments that remain unexplained. Completeness implies that the report should give a comprehensive idea of what has been done by mentioning all key aspects of the design, experiments, and analysis. Aspects of the system, be it of its design or of its behavior, that remain unexplained detract from the credibility and grade of the report. Correctness is expected in terms of the explanations being logical and correlated with the numbers in the experiments and the design.

Remember that this is a report about the system you have designed and built, about the experiments you have performed, and about how you interpret the results of the experiments and map them to your design and implementation. Please do not contact us seeking confirmation and assurances about, e.g., whether the report is sufficient, your interpretation of the data, validation of concrete aspects of your design, or whether you have done enough experiments. Making those decisions is your job and part of what the course will evaluate.

The report will be graded together with the code and data submitted. **You will need at least 400 points to pass.** You might be called for a meeting in person to clarify aspects of the report or the system. By submitting the report, the code, and the data, you confirm that you have done the work on your own, the code has been developed by yourself, the data submitted comes from experiments you have done, you have written the report on your own, and you have not copied neither code nor text nor data from other sources.

Each section will have to contain a description of the experimental setup, graphs (where applicable) and tables with most important numbers. More importantly, each experiment and modeling result should be accompanied by an explanation that puts the result in the context of your system. Without proper explanations you risk losing most of the points in the experimental sections. Furthermore, it is expected that the interactive law holds for all experiments. In case throughput and response time do not match, it is imperative that you explain why, otherwise you risk losing most of the points for the experiment in question.

In each section you can find a table summarizing the system configuration that should be used for that particular section. These configurations should be considered as a guideline to help you find interesting configurations where we expect your system to change its behaviour (e.g., going from an under-saturated to a saturated system). If you do not observe such a change inside that range of parameters, you should expand the range and run additional experiments outside the given range. In that case, ask your assistant for advise.

This document provides you with the end-results that have to be included in the report. Please keep in mind that you are encouraged to provide additional graphs and experiments that help you in understanding your system. These additional experiments should be chosen such that they support the claims you make in your explanations and help you illustrate your point. Always backup your claims using results from your experiments and always answer questions in the correct Section.

Formatting guidelines

We expect you to use this template for the report, but you are free to use another text processor. Keep in mind the following:

- We expect you to submit **a single PDF that has the same section structure as this template, and answers all points we outline here**. If you use this file, you should

remove this page with notes, and the short description provided by us at the beginning of sections.

- Keep the same cover page as on this document and **fill out your name and legi number**. Leave the grading table empty.
- Write your name and Legi number on every page.
- The main text should be in **single-column format with 11pt font on A4 paper**. In case you don't start with one of the files provided by us, **for margins use 2.54 cm (1 inch) on all sides**.
- The length of report should be at minimum 20, but should not exceed 35 pages (including figures, tables, etc.).

How to Generate Load and Populate the Data

To generate load on the system, you will use memtier. When using memtier, you can specify two parameters: the number of threads (CT) and the number of virtual clients per thread (VC). The total number of virtual clients per memtier is CPM=CT*VC. When referring to "number of clients" in the system, it is the number of memtier instances times the number of virtual clients per instance, i.e., NumClients=NumMemtier*CPM. For each experiment conducted for this report, the cofiguration setup will be provided in the respective section. In addition, for each section, run all experiments using the same cloud deployment as different deployments might introduce performance discrepancies in your measurements.

The number of keys should always be set to 10000 (`--key-maximum=10000`). Before running experiments, populate the key-value store using a write-only (using the correct value size) workload long enough to write all keys at least once. All sets should have a long expiry time, to ensure that keys are not evicted during the experiments (e.g., `--expiry-range=9999-10000`). The loading of the memcached servers is not part of the measured performance, and the middleware has to be restarted after the initial population, before running experiments.

Plotting Data

The readability of your plots will be part of the grade. For all graphs the main metric being studied (usually the y-axis or z-axix) should start at 0, always use appropriate error metrics/bars, label both axis with units. For graphs that are related try to keep the same range on the corresponding axis. Be careful when using log-scale on any of the axis.

1 System Overview (80 pts)

1.1 System Description (20 pts)

Describe the implementation of your system and highlight design decisions relevant to the experiments. Explain the internal structure of your middleware, which queues and how many are used, minimum and maximum number of threads available for each task, what instrumentation has been included, what you are measuring and where. Use a figure to illustrate all the previous concepts along with the global system architecture (including all the middleware internal components, clients, and servers).

1.2 Middleware Data-Structures (20 pts)

Describe the data-structures used internally in the middleware to manage messages and connections. Explain how you implement/design network threads, queues, request storage, code instrumentation, performance measurement storage, etc.

1.3 Middleware Request Handling (16 pts)

1.3.1 Request Parsing (8 pts)

Describe how the Middleware handles requests, more specifically how do you parse them.

1.3.2 Requests Processing (8 pts)

Describe how the Middleware handles GET requests.

1.4 Work Balancing (8 pts)

Describe how you balance work between servers. Provide experimental results to show it. Note that, for this section, you should use data from the experiments executed in Section 3.

1.5 Data Processing (8 pts)

Describe how you process the data from your experiments to produce the results included in the report.

1.6 Experimental Setup (8 pts)

Describe how you run your experiments. If you scripted your tests, describe how. Otherwise, describe which method you used to run your tests.

1.7 Additional Remarks (optional, remove if not needed)

Describe additional artifacts necessary to understand the behaviour of your system. Problems or bugs that your system presents should be illustrated in this section. Not being able to explain such problems will lead to point deductions.

2 Baseline without Middleware (120 pts)

In these experiments you will study the performance characteristics of the memtier clients, value size for different configurations of memcached servers.

2.1 One Server (50 pts)

Study the throughput and the response time as a function of NumClients and value size for read-only workloads. Each memtier instance is connected to a single memcached instance. Analyse how the behavior of the server changes as more clients and larger message sizes are used.

Number of servers	1
Number of client machines	3
Instances of memtier per machine	1
Threads per memtier instance	3
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value Size (Bytes)	64, 256, 512, 1024
Number of middlewares	N/A
Worker threads per middleware	N/A
Repetitions	3 or more (at least 1 minute each)

2.1.1 Setup (5 pts)

Describe the setup you used for this experiment, how did you launch and run the experiments, if you always used the same deployment or not, etc.

2.1.2 Throughput (10 pts)

Plot the throughput as a function of NumClients and value sizes. All saturation points must be included and described. Plots can be either: i) a 3D plot using bars (x-axis is number o clients, y-axis is value size, z-axis is throughput), or ii) two 2D line plots where the y-axis is always throughput and the x-axis is the number of clients or value size. If you pick the latter, use multiple lines, show the results using the other metric, for example, if the x-axis is the value size, you should have a line for each number of clients.

2.1.3 Response Time (10 pts)

Plot the reponse time as a function of NumClients and value sizes. All saturation points must be included and described. Plots are the same as in the previous section (replace the throughput by the response time in the corresponding axis).

2.1.4 Result Analysis (5 pts)

Explain if results are explained by the interactive law (and why).

2.1.5 Explanation (20 pts)

Describe in which phase the memcached servers are under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients and value sizes have on throughput and response time. Explain what further conclusions can be drawn from the experiment.

2.2 Three Servers (50 pts)

Study the throughput and the response time as a function of NumClients and value size for read-only workloads. Each memtier instance is connected to a single memcached instance. Analyse how the behavior of the server changes as more clients and larger message sizes are used.

Number of servers	3
Number of client machines	3
Instances of memtier per machine	3
Threads per memtier instance	1
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value size (Bytes)	64, 256, 512, 1024
Number of middlewares	N/A
Worker threads per middleware	N/A
Repetitions	3 or more (at least 1 minute each)

2.2.1 Setup (5 pts)

Describe the setup you used for this experiment, how did you launch and run the experiments, if you always used the same deployment or not, etc.

2.2.2 Throughput (10 pts)

Plot the throughput as a function of NumClients and value sizes. All saturation points must be included and described. Use the same plot setup as in Section 2.1.2.

2.2.3 Response Time (10 pts)

Plot the response time as a function of NumClients and value sizes. All saturation points must be included and described. Use the same plot setup as in Section 2.1.3.

2.2.4 Result Analysis (5 pts)

Explain the results are explained by the interactive law (and why) and if they are consistent with previous experiments.

2.2.5 Explanation (20 pts)

Describe in which phase the memcached servers are under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients and value sizes have on throughput and response time. Explain what further conclusions can be drawn from the experiment.

2.3 Summary (20 pts)

Based on the experiments above, fill out the following table:

	Maximum Throughput	Corresponding Response Time	Configuration (number of clients, msg. size)
One memcached server			
Three memcached servers			

2.3.1 Bottleneck Analysis (10 pts)

Describe what are the bottlenecks of this setup is. If the maximum throughput for both experiments is the same, explain why. If it is not the case, explain why not. Explain how the system behaviour when the number of clients and value size increase.

2.3.2 One and Three Servers Configurations (5 pts)

Compare one server and three servers configurations and how it impacts throughput and response time.

3 Baseline with Middleware (240 pts)

In these experiments you will study the performance characteristics of the memtier clients, value size, and number of worker threads for different configurations of middlewares and memcached servers.

3.1 One Middleware, One Server (50 pts)

Study the throughput and the response time as a function of NumClients, value size, and worker threads, for read-only workloads. Each memtier instance is connected to a single middleware instance. Analyse the behavior of the system as more clients, larger message sizes, and more worker threads are used. Measure the response time *both at the client and at the middleware*, and plot the throughput and response time measured in the middleware.

Number of servers	1
Number of client machines	3
Instances of memtier per machine	1
Threads per memtier instance	2
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value size (Bytes)	64, 256, 512, 1024
Number of middlewares	1
Worker threads per middleware	8, 32, 64
Repetitions	3 or more (at least 1 minute each)

3.1.1 Setup (5pts)

Describe the setup you used for this experiment, how did you launch and run the experiments, if you always used the same deployment or not, etc.

3.1.2 Throughput (10pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. You should present six 2D plots configured as follows (the Y axix is always throughput):

1. 8 worker threads, x-axis represents the number of clients, one line for each value size;
2. 8 worker threads, x-axis represents the value sizes, one line for each number of clients;
3. 32 worker threads, x-axis represents the number of clients, one line for each value size;
4. 32 worker threads, x-axis represents the value sizes, one line for each number of clients;
5. 64 worker threads, x-axis represents the number of clients, one line for each value size;
6. 64 worker threads, x-axis represents the value sizes, one line for each number of clients;

3.1.3 Response Time (10pts)

Plot the reponse time as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Plots are the same as in the previous section (replace the throughput by the response time in the corresponding axis).

3.1.4 Result Analysis (5pts

Explain the results are explained by the interactive law (and why) and if they are consistent with previous experiments.

3.1.5 Explanation (20pts)

Describe in which phase the system is under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients, value sizes, and worker threads have on throughput and response time. Explain what further conclusions can be drawn from the experiment. Relate you analysis with system elements such as average queue length.

3.2 One Middleware, Three Server (50 pts)

Study the throughput and the response time as a function of NumClients, value size, and worker threads, for read-only workloads. Each memtier instance is connected to a single middleware instance. Analyse the behavior of the system as more clients, larger message sizes, and more worker threads are used. Measure the response time *both at the client and at the middleware*, and plot the throughput and response time measured in the middleware.

Number of servers	3
Number of client machines	3
Instances of memtier per machine	1
Threads per memtier instance	2
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value size (Bytes)	64, 256, 512, 1024
Number of middlewares	1
Worker threads per middleware	8, 32, 64
Repetitions	3 or more (at least 1 minute each)

3.2.1 Setup (5 pts)

Describe the setup you used for this experiment, how did you launch and run the experiments, if you always used the same deployment or not, etc.

3.2.2 Throughput (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.2.

3.2.3 Response Time (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.3.

3.2.4 Result Analysis (5 pts)

Explain the results are explained by the interactive law (and why) and if they are consistent with previous experiments.

3.2.5 Explanation (pts 20)

Describe in which phase the system is under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients, value sizes, and worker threads have on throughput and response time. Explain what further conclusions can be drawn from the experiment. Relate you analysis with system elements such as average queue length.

3.3 Two Middlewares, One Server (50 pts)

Study the throughput and the response time as a function of NumClients, value size, and worker threads, for read-only workloads. Each memtier instance is connected to a single middleware instance. Analyse the behavior of the system as more clients, larger message sizes, and more worker threads are used. Measure the response time *both at the client and at the middleware*, and plot the throughput and response time measured in the middleware.

Number of servers	1
Number of client machines	3
Instances of memtier per machine	2
Threads per memtier instance	1
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value size (Bytes)	64, 256, 512, 1024
Number of middlewares	2
Worker threads per middleware	8, 32, 64
Repetitions	3 or more (at least 1 minute each)

3.3.1 Setup (5 pts)

Describe the setup you used for this experiment. In particular, describe any changes and the granularity used for virtual client threads, value sizes, and number of worker threads.

3.3.2 Throughput (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.2.

3.3.3 Response Time (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.3.

3.3.4 Result Analysis (5 pts)

Explain the results are explained by the interactive law (and why) and if they are consistent with previous experiments.

3.3.5 Explanation (20 pts)

Describe in which phase the system is under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients, value sizes, and worker threads have on throughput and response time. Explain what further conclusions can be drawn from the experiment. Relate you analysis with system elements such as average queue length.

3.4 Two Middlewares, Three Server (50 pts)

Study the throughput and the response time as a function of NumClients, value size, and worker threads, for read-only workloads. Each memtier instance is connected to a single middleware instance. Analyse the behavior of the system as more clients, larger message sizes, and more worker threads are used. Measure the response time *both at the client and at the middleware*, and plot the throughput and response time measured in the middleware.

Number of servers	3
Number of client machines	3
Instances of memtier per machine	2
Threads per memtier instance	1
Virtual clients per thread	4, 8, 16, 32
Workload	Read-only
Value size (Bytes)	64, 256, 512, 1024
Number of middlewares	2
Worker threads per middleware	8, 32, 64
Repetitions	3 or more (at least 1 minute each)

3.4.1 Setup (5 pts)

Describe the setup you used for this experiment. In particular, describe any changes and the granularity used for virtual client threads, value sizes, and number of worker threads.

3.4.2 Throughput (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.2.

3.4.3 Response Time (10 pts)

Plot the throughput as a function of NumClients, value sizes, and worker threads. All saturation points must be included and described. Use the same plot setup as in Section 3.1.3.

3.4.4 Result Analysis (5 pts)

Explain the results are explained by the interactive law (and why) and if they are consistent with previous experiments.

3.4.5 Explanation (20 pts)

Describe in which phase the system is under-saturated, saturated, or over-saturated. Describe how throughput and response time correlate. Explain the effect that different NumClients, value sizes, and worker threads have on throughput and response time. Explain what further conclusions can be drawn from the experiment. Relate you analysis with system elements such as average queue length.

3.5 Summary (40 pts)

Based on the experiments above, fill out the following table. Miss rate represents the percentage of GET requests that return no data. Time in the queue refers to the time spent in the queue between the net-thread and the worker threads.

Maximum throughput for one and two middlewares.

	Throughput	Response time	Average time in queue	Miss rate
Section 3.1: Measured on middleware				
Section 3.1: Measured on clients			n/a	
Section 3.2: Measured on middleware				
Section 3.2: Measured on clients			n/a	
Section 3.3: Measured on middleware				
Section 3.3: Measured on clients			n/a	
Section 3.4: Measured on middleware				
Section 3.4: Measured on clients			n/a	

3.5.1 Bottleneck Analysis (10 pts)

Describe what are the bottlenecks of this setup is. If the maximum throughput for both experiments is the same, explain why. If it is not the case, explain why not. Explain how the system behaviour when the number of clients, value size, and worker threads increase.

3.5.2 One and Two Middlewaress Configurations (10 pts)

Compare one middleware and two middlewares configurations.

3.5.3 One and Three Server Configurations (10 pts)

Compare one and three servers configurations.

4 2K Analysis (80 pts)

For 3 client machines (with 64 total virtual clients per client VM) measure the throughput and response time of your system in a 2K experiment with repetitions. Value size is fixed to 256Bs. Investigate the following parameters:

- Memcached servers: 1 and 3
- Middlewares: 1 and 2
- Worker threads per MW: 8 and 32

Number of servers	1 and 3
Number of client machines	3
Instances of memtier per machine	1 (1 MW setup), 2 (2 MWs setup)
Threads per memtier instance	2 (1 MW setup), 1 (2 MWs setup)
Virtual clients per thread	32
Workload	Read-only
Value size (Bytes)	256
Number of middlewares	1 and 2
Worker threads per middleware	8 and 32
Repetitions	3 or more (at least 1 minute each)

4.1 2K Model (20 pts)

Present your 2K model and justify your decisions. Present all steps, values, and formulas.

4.2 Throughput and Response time Impact (40 pts)

Discuss the impact of these parameters on throughput and response time.

4.3 Comparing Model Results with Experimental Results (20 pts)

Relate the results from the model with the experimental results of Section 3. Discuss if the model explains your system (or not) and why.

5 Queuing Model (80 pts)

Note that for queuing models it is enough to use the experimental results from the previous sections. It is, however, possible that the numbers you need are not only the ones in the plots we asked for, but also the internal measurements that you have obtained through instrumentation of your middleware. Explain which additional measurements were necessary and how you instrumented the code to obtain them. For each model include an explaination of how you claculate the following parameters: service rate, arrival rate, number of jobs in the queue, time in the queue, and servite time.

5.1 M/M/1 (30 pts)

Build queuing models based on Section 3.2 (1 middleware, 3 servers) for each worker-thread configuration of the middleware. Use one M/M/1 queue to model your entire system. Motivate your choice of input parameters to the model. Explain for which experiments the predictions of the model match and for which they do not.

5.2 M/M/m (30 pts)

Build an M/M/m models based on Section 3.2 (1 middleware, 3 servers), where each middleware worker thread is represented as one service. Motivate your choice of input parameters to the model. Explain for which experiments the predictions of the model match and for which they do not.

5.3 Network of Queues (20 pts)

Based on Section 3.4 (2 middlewares, 3 servers), build a network of queues which simulates your system. Motivate the design of your network of queues and relate it wherever possible to a component of your system. Motivate your choice of input parameters for the different queues inside the network. Perform a detailed analysis of the utilization of each component and clearly state what the bottleneck of your system is. Explain for which experiments the predictions of the model match and for which they do not.