CS 736 Software Performance Engineering
Final Exam – Fall 2010

This is a take home exam. By writing your name and student id in the box below you acknowledge that **you will not seek help nor discuss the exam with anyone until after the exam due date on Thursday, December 16, 2010**. Late exams will not be accepted.

All answers must be either typed or written clearly and neatly. Include your name and page number (e.g., Page 1 of 5) on each page. **E-mail your solutions (including these two pages) as a single file in a PDF or Microsoft Word format** to the instructor at **Katerina.Goseva@mail.wvu.edu** or slide a hard copy under 721 ESB office door at latest by 1:00 pm on December 16 2010.

Student name:	Ekrem Kocaguneli
Student id:	701095814

1	/10
2	/10
3	/30
4	/15
5	/10
6	/10
7	/25
Total	**/110**

1. Describe the difference between event mode and sampling mode performance measurement tools. Specifically address the tradeoffs between the overhead and accuracy. (10 points)

One of the categories we can group performance measurements tools is their measurement mode. According to measurement mode performance measurement tools are either *event mode tools* or *sampling mode tools*.

As the name suggests the ***event mode*** measurement tools are triggered by specific events that were defined earlier and an associated record with that event is generated. The positive side of event mode measurement tools is that they have high accuracy, as they catch the specified events. However, the problem is that when the event rate is too high or too many events are selected to be measured or too much data was chosen to be measured per event; then there will be a considerable overhead due to the measurement tool.

Sampling mode tools measures the events at pre-defined intervals, i.e. the sampling is triggered by timer interrupts rather than events. Positive side of sampling mode tools (when compared to event mode tools) is that we can control the overhead of the tool. The overhead of sampling mode tools is due to number of variables to be measured and the time interval of the measurements. Since we are able to control both of these properties, we can also control the overhead. The negative side of sampling mode in comparison to event mode is that it is less accurate and can miss bursts. For example if a considerable burst has occurred in the performance between two sampling times, then the sampling mode tools would miss it.

Ultimately the comparison between event mode and sampling mode tools is the trade-off between accuracy and overhead. With event mode tools we can attain high accuracy at the cost of high tool overhead. On the other hand, sampling mode tools will let us adjust the overhead, but we will sacrifice from the accuracy.

2. Explain why often we need to use both system level and program level measurement tools. Give an example. (5 + 5 points)

The system level measurement tools provide us *high-level view* by showing the system-wide usage statistics such as global CPU and DISK utilizations as well as total I/O operations and total traffic through a router. They can also provide per-process results. In UNIX systems there are a number of system level tools:

- Ps: reporting information for processes running on the system
- Sar: reports system activity data that is maintained by the kernel
- Vmstat: gives information about virtual memory usage, disk access and CPU utilization

However, system level measurement tools do not provide program specific information. For *low-level view* of the system, i.e. program-related statistics we need program level measurement tools. Program level measurement tools provide statistics such as CPU and memory usage of the program, number of I/O operations per execution of the program and the number of packets per second for a particular program. Furthermore, program-level measurement tools provide profiling information associated with programs (time spent at each function, which functions called each function).

Example: Assume that we are coding up an enterprise resource planning program (ERP) that will run on a local server and we want to measure its performance on a particular server. Then we would need system level measurement tools to give us information about the program as a whole, such as the total CPU and DISK usage of the ERP as well as its total package transfers. However, system level measurement tool would not catch information regarding which particular function used these resources (CPU, DISK, LAN) the most. So as to get a detailed view and to be able to perform profiling work on the ERP system, we would need program level measurement tools.

3. Performance benchmarks are often used for generating workloads and measuring performance of computer systems.

a. Compared to real workload, what are the advantages and disadvantages of generating the workload using benchmarks? (<u>5 points</u>)

As we discussed in the lecture, the best way to observe the behavior of a system is to actually run a real workload and collect the results. However, sometimes this is either not possible or we want an estimate of the future workload for a program. In that case, benchmarks come handy to generate workloads and enable us to collect performance measures. On the other hand the use of benchmarks has both advantages and disadvantages. The advantages are:

- Provides clear definition of performance objectives and workloads
- Enables measurement of performance
- Enables repeatability of the measurements
- Enables comparison of different products
- Enables the testing of new systems
    - Companies, developers etc. can pinpoint performance problems in advance
    - The impact of future changes can be assessed by the companies
- Provides better understanding of system, users, performance criterion to be measured as well as the results

The disadvantages of the benchmarks are:

- Processor rating is usually done via MIPS (million instructions per second), however this is dependent on the instruction set. Therefore, comparison of RISC (reduced instruction set computer) and CISC (complex instruction set computer) does not make much sense.
- If the tested system configurations are not close to actual system configurations, then benchmark does not provide much of a help.

- If benchmark workload is not close to actual workload, then the measured performance may be misguiding.
    - Furthermore, the behavior of the workload cannot always be predicted. For example, unexpected bursts due to particular unexpected events cannot be included in the benchmark.
- Vendors may use industry-standard benchmarks, and may tune their program particularly for that benchmark. However, this is some sort of cheating and a particular program may perform differently at another benchmark.
- Performance benchmarks usually focus on computational speed. However, benchmarks should also be domain specific. For example, in an image rendering program, the quality of the image is as important as the computational speed.

b. Describe the benchmark model of the TPC-C benchmark used for generating workload and measuring performance of transaction processing systems (http://www.tpc.org/tpcc/default.asp). (10 points)

TPC stands for transaction processing performance council and it provides different types of benchmarks for transaction processing and databases. One of the TPC benchmarks is TPC-C, which is an on-line transaction processing (OLTP) benchmark.

TPC-C is capable of simulating a whole environment for a large user population. In this computational environment users execute queries on a database. In the center of the benchmark, there is an order-entry environment, i.e. the fundamental transactions are associated with an order-entry environment, such as entering and delivering orders, recording payments, checking the status of different orders, monitoring stock and warehouse levels etc. Due to that property, TPC-C can simulate the computational environment of any vendor that manages, sells or distributes a particular product or service.

TPC-C is proposed as a more complex OLTP benchmark in comparison to previous OLTP benchmarks. The reason for the complexity is that it enables multiple-transaction types as well as more complex database structures and more complex execution structures. TPC-C provides the users five different types of concurrent transactions and their combinations. The transactions are allowed to be executed on-line or they can be queued for later execution. These transactions are listed as follows:

- New Order: Customer places a new order
- Payment: An update operation, balance is updated for the new payment
- Delivery: A batch operation, order is delivered with this transaction
- Order Status: A retrieval operation, retrieves the last order of a customer
- Stock-level: Monitors the warehouse inventory

The proposed database in TPC-C includes nine different types of tables, where a wide-range of population and record size can be accommodated. The measurement unit of TPC-C is transactions per minute (tpmC). Furthermore, by using this tpmC unit, the TPC-C benchmark has provided rules of thumb for the users, such as 1.2 tpmC per user/terminal to be defined as the maximum amount.

c. TPC-C uses two metrics tpmC and Price/tpmC to compare computer systems. Define both metrics and discuss the situations when you would consider one over the other in selecting a transaction based system. Illustrate your reasoning using the top ten TPC-C price/performance results given at http://www.tpc.org/tpcc/results/tpcc_price_perf_results.asp.
(15 points)

TPC-C is defined as transactions per minute (tpmC). Price/tpmC is a price/performance number, which is derived by taking the price of the entire system (maintenance and all the other aspects of the system) divided by the performance. Depending on the aim of my system, I may prefer using one of these metrics over the other. For example, if my application has a very high demand that requires mission-critical performance, then I would put higher emphasis on tpmC, i.e. throughput metric would be more important for me. On the other hand, if my system has a very light workload and I am constrained with a tight budget, then I would choose to use price/tpmC metric.

Since now we know when to use which performance measure, let us define two scenarios and choose two alternatives from the list of top ten TPC-C price/performance results.

- **Scenario 1:** We are choosing a computer system for a rocket-launch system. In that case the system is a mission-critical system that has to have a very high throughput and the investment in the rocket-launch program is very important. Therefore, I would use tpmC metric and I would choose number 7 (IBM x3850 X5).
- **Scenario 2:** Let's assume that we are the proud owner of a developing market and we want a moderate and not too costly computer system for our store. Then my primary concern would be the budget and I would use price/tpmC metric. In that scenario, I would use either number 1 (HP ProLiant ML350) with a price/tpmC of 0.39 or number 2 (HP ProLiant DL580) with a price/tpmC of 0.39. If both of them are affordable for me, then I would look at the tpmC values. For number 1 the tpmC value is 290,040 whereas for number 2 it is 1,807,347. Since number 2 has a much higher tpmC value, I would eventually choose number 2 (HP ProLiant DL580).

4. The locality principle used for improving performance states that actions, functions, and results should be created close to physical computer resources used to produce them. "Close" may be
- **Special**
- **Temporal (i.e., time)**
- **Effectual (i.e., purpose or intent)**
- **Degree (i.e., intensity or size).**

Explain why. Give examples that illustrate each case. (3 + 12 points)

When designing a system, one can take into consideration some principles so that the designed system will meet the aimed performance objectives. In other words, those principles help the system design activities to occur in accordance with aimed performance goals. These principles provide us guidance when making performance decision; help us identify system requirements (what is expected from the system to do), identify system design (how the tasks are to be done by the system) and identify detailed design (what type of algorithms and data structures are to be used.)

There are 9 such design principles and one of them is the locality principles, which states that system components (actions, functions, results) shall be close to the physical resource they interact

with. In other words, it answers the question of how closely logical tasks match the character of physical resources. The closeness or locality can be defined in different ways. There are 4 locality types. Below is the list of 4 types and brief definitions and examples explaining each one of them.

- **Spatial Locality:** This type of locality deals with the tasks or properties of the system and the physical resources to be close to one another in terms of space. For example in a system that provides menu contents to a user, clustering related actions together is an example of using strong spatial locality. Another example can be from an accounting system, where data of various user types are stored in databases. The use of strong spatial locality principle would be to cluster similar users together and store in same physical disk.
- **Temporal Locality:** This type of locality deals with the locality in time, in other words, frequently used actions are grouped together in their own cluster and non-frequent actions are grouped in another cluster of their own. A good example to the use of temporal locality would be in a menu hierarchy, where frequently used items would be presented first. Another example to the use of temporal locality would be in data organization, where data would be organized depending on how frequent they are accessed, i.e. a highly utilized index table may be preferred to be kept in cache, whereas a less used data table may be preferred to be kept in physical disk and an archive-like bulky data may be kept in remote servers or even on tapes.
- **Effectual Locality:** It depends on the fact whether a particular action is matched to the appropriate computing resource (processor) it is being run on. A good example to the effective use of locality would be the assignment of two different types of actions in a mission-critical system to two different types of processors. For example regular daily actions in this system such as user log-in, routine controls etc. may be assigned to general-purpose processors, whereas failure-recovery actions which need immediate response may be assigned to high-performance processors.
- **Degree Locality:** Degree locality deals with the extent of the action, i.e. the ability of the memory to be able to fit the amount of data required by the system at once. A good use of degree locality is to understand the system requirement and use the appropriate processor with right amount of cache. For example in a monopoly system, knowing the information regarding current state of the market and providing a fast response is very important, so such a system should be capable of storing multiple tables in cache; therefore, a processor with a large cache should be bought. However, a system that is mainly executed for batch actions, i.e. that has enough time to access remote disks, requires a cheaper processor with a smaller cache size.

5. Assume that a system consists of three redundant copies. Explain how redundancy affects performance, availability, and reliability. Give an example (different from those discussed in class). (<u>10 points</u>)

Redundancy is a technique that is heavily used in cloud computing. The idea behind redundancy is the replication of the system or the database, in the sense that the at least one or more exact copies of the current system or database is kept available all the time at a remote site. The benefit of such an approach is that the system performance increases, because depending on the workload, different workloads may be directed to different system copies. Furthermore, it increases the availability of the system, in the sense that whenever one of the systems is down, then there is another system that can be used. Furthermore, it increases the reliability, since the user is not greatly affected by one system failing.

A nice example where redundancy would be useful is an enterprise's database management system. Such an enterprise may choose one of the service providers (Google, Amazon or Microsoft) to have a redundant copy of their DBMS at a remote site and in case of an event that has caused the

actual DBMS to go down, the copy at the remote site can very easily be resumed as the operational DBMS until the primary DBMS is prepared to be back online.

6. **Based on the results presented in the lecture on Web servers it appears that we cannot model the Web servers performance using the types of queueing networks we studied in class. Explain why. Briefly describe the two distinguished characteristics of Web workloads. (10 points)**

The basic modeling principle associated with queuing networks is that they assume that the arrival rates can be modeled via a Poisson distribution. However, for the behavior of web servers cannot be modeled with such distributions. The assumption in Poisson distribution is that it assumes that the request inter-arrival times are independent random variables, which follow exponential distribution. However, it is shown by Popstojanova et. al.[i] that request arrivals in actual web servers do not follow a Poisson distribution. On the contrary, the behavior of the web server appears to be self similar, long-range dependent and with a heavy-tailed distribution.

The two characteristics of web workloads are:

- **Self similarity:** In terms of web workload, we are dealing with time series. Let X be a stationary sequence and let X_k^m be the corresponding aggregated sequence with blocks of size m, then such a stationary sequence is said to second-order-similar if it has the same variance and autocorrelation as another such X for all m.
- **Long-range dependence:** Long-range dependence deals with the correlation structure of a time series. A series is said to have long-range dependence if its correlation structure is the same in coarse or fine time scales, i.e. its correlation structure is irrespective of the time-aggregation.

7. **Read the paper**

Stantchev, V., "Performance Evaluation of Cloud Computing Offerings" Third International Conference on Advanced Engineering Computing and Applications in Sciences, 2009, pp. 187-192.

Then, address the following:

a. Describe the method used for evaluation of performance in this paper. (5 points)

The propose method is different from the traditional methods in the sense that it focuses on the nonfunctional properties (NFPs) of individual services, thereby it provides a better granularity. The proposed methodology is reported to be a continuation of ASTAR, which is described as a method for the evaluation of architectural configurations in the context of service oriented architecture (SOA). Original ASTAR is used to evaluate the replication configurations of operating systems and service-ware, whereas the proposed system applies that method to cloud computing domain.

The proposed methodology is designed to be a generic recommendation tool for different cloud configurations and it is composed of 5 steps:

- Evaluation of suitable benchmarks
- Identify configuration: Evaluation of different service models (IaaS, PaaS, SaaS)
- Run Tests: This step includes the run of actual tests
- Analyze: This is one of the important steps, in the sense that it compares the results of different configurations, evaluates the interdependencies between NFPs

(high transaction rate meaning lower reliability) and it provides a cost-benefit analysis.

- Recommend: The results elicited from the previous step are used here to propose a recommendation.

b. Describe the benchmark used for evaluation. (5 points)

The selected benchmark for the evaluation of their proposed method is WSTest, which provides the operations of EchoVoid(), EchoStruct(), EchoList() and GetOrder(). WSTest is defined as a generic benchmark that is suitable for their methodology. Aside from WSTest, they discuss other benchmarks that they have considered and explain why the others are not suitable for their methods. For example, bench method is rejected because it is designed particularly as a cutomer-to-business method. Furthermore, they reject TPC benchmarks because they are not web-service specific.

c. Describe how the throughput (i.e. Transactions per second) and response time depend on the number of instances used. (10 points)

In the reported experimental evaluation they simulate 200 parallel concurrent clients, where each virtual client generating 100 requests. These requests are then distributed to either one of the instances by a random controller. Therefore, as long as our bandwidth or the upper boundary of the test drivers are not reached, adding a second or a third instance will mean that more than one server is responding to the requests of the clients. Hence, the response times will decrease and more transactions per unit time (higher throughput) will be processed. It is also noted in the paper that the addition of a second copy has provided dramatic increase in throughput and a dramatic decrease in the response times. However, that effect diminishes with the addition of a third copy. The interpretation for the limited effect of the third copy is attributed to the network and driver boundaries, but a thorough discussion on the topic is not provided, as this discussion is considered to be out of the topic of the paper.

d. Describe at least one limitation or threat to validity of this work. (5 points)

The biggest threat to validity that occurs to me is the selection of Amazon EC2. The first threat coming from such a selection is that different cloud computing environments or vendors are using different configurations and the reported results can change from one environment/vendor to another. Therefore, the recommendations and the proposed results may be valid only for Amazon EC2. Furthermore, there is only a single benchmark used in that study. Of course the benchmark is supposed to be able to simulate for the selected application scenario and authors note that benchmarks other than WSTest are not suitable for their experiments. However, particular benchmarks are known to favor particular configurations or particular vendors. Therefore, I see it as a threat to the validity of the results and it should have been noted by authors that their results are limited to the vendor they deployed their tests in as well as to the benchmark that they have used.

[i] Katerina Goseva-Popstojanova, Fengbin Li, Xuan Wang, and Amit Sangle, A Contribution Towards Solving the Web Workload Puzzle, DSN'06.