

CSE 567M

Computer Systems

Analysis

Raj Jain

Washington University in Saint Louis

Saint Louis, MO 63130

Jain@cse.wustl.edu

These slides are available on-line at:

<http://www.cse.wustl.edu/~jain/cse567-15/>




- ❑ Goal of this Course
- ❑ Contents of the course
- ❑ Tentative Schedule
- ❑ Project
- ❑ Grading

Goal of This Course

- ❑ Comprehensive course on performance analysis
- ❑ Includes measurement, statistical modeling, experimental design, simulation, and queuing theory
- ❑ How to avoid common mistakes in performance analysis
- ❑ Graduate course: (Advanced Topics)
 - ⇒ Lot of independent reading and writing
 - ⇒ Project/Survey paper (Research techniques)

Text Book

- ❑ R. Jain, “Art of Computer Systems Performance Analysis,” Wiley, 1991, ISBN:0471503363



Objectives: What You Will Learn

- ❑ Specifying performance requirements
- ❑ Evaluating design alternatives
- ❑ Comparing two or more systems
- ❑ Determining the optimal value of a parameter (system tuning)
- ❑ Finding the performance bottleneck (bottleneck identification)
- ❑ Characterizing the load on the system (workload characterization)
- ❑ Determining the number and sizes of components (capacity planning)
- ❑ Predicting the performance at future loads (forecasting).

Basic Terms

- ❑ **System:** Any collection of hardware, software, and firmware
- ❑ **Metrics:** Criteria used to evaluate the performance of the system. components.
- ❑ **Workloads:** The requests made by the users of the system.

Main Parts of the Course

- ❑ Part I: An Overview of Performance Evaluation
- ❑ Part II: Measurement Techniques and Tools
- ❑ Part III: Probability Theory and Statistics
- ❑ Part IV: Experimental Design and Analysis
- ❑ Part V: Simulation
- ❑ Part VI: Queueing Theory

Part I: An Overview of Performance Evaluation

- ❑ Introduction
- ❑ Common Mistakes and How To Avoid Them
- ❑ Selection of Techniques and Metrics

Example I

- ❑ What performance metrics should be used to compare the performance of the following systems:
 - Two disk drives?
 - Two transaction-processing systems?
 - Two packet-retransmission algorithms?

Part II: Measurement Techniques and Tools

- ❑ Types of Workloads
- ❑ Popular Benchmarks
- ❑ The Art of Workload Selection
- ❑ Workload Characterization Techniques
- ❑ Monitors
- ❑ Accounting Logs
- ❑ Monitoring Distributed Systems
- ❑ Load Drivers
- ❑ Capacity Planning
- ❑ The Art of Data Presentation
- ❑ Ratio Games

Example II

- ❑ Which type of monitor (software or hardware) would be more suitable for measuring each of the following quantities:
 - Number of Instructions executed by a processor?
 - Degree of multiprogramming on a timesharing system?
 - Response time of packets on a network?

Part III: Probability Theory and Statistics

- ❑ Probability and Statistics Concepts
- ❑ Four Important Distributions
- ❑ Summarizing Measured Data By a Single Number
- ❑ Summarizing The Variability Of Measured Data
- ❑ Graphical Methods to Determine Distributions of Measured Data
- ❑ Sample Statistics
- ❑ Confidence Interval
- ❑ Comparing Two Alternatives
- ❑ Measures of Relationship
- ❑ Simple Linear Regression Models
- ❑ Multiple Linear Regression Models
- ❑ Other Regression Models

Example III

- The number of packets lost on two links was measured for four file sizes as shown below:

File Size	Link A	Link B
1000	5	10
1200	7	3
1300	3	0
50	0	1

Which link is better?

Part IV: Experimental Design and Analysis

- ❑ Introduction to Experimental Design
- ❑ 2^k Factorial Designs
- ❑ 2^{kr} Factorial Designs with Replications
- ❑ 2^{k-p} Fractional Factorial Designs
- ❑ One Factor Experiments
- ❑ Two Factors Full Factorial Design without Replications
- ❑ Two Factors Full Factorial Design with Replications
- ❑ General Full Factorial Designs With k Factors

Example IV

- ❑ The performance of a system depends on the following three factors:
 - Garbage collection technique used: G1, G2, or none.
 - Type of workload: editing, computing, or AI.
 - Type of CPU: C1, C2, or C3.

How many experiments are needed? How does one estimate the performance impact of each factor?

Part V: Simulation

- ❑ Introduction to Simulation
- ❑ Types of Simulations
- ❑ Model Verification and Validation
- ❑ Analysis of Simulation Results
- ❑ Random-Number Generation
- ❑ Testing Random-Number Generators
- ❑ Random-Variate Generation
- ❑ Commonly Used Distributions

Example V

- ❑ In order to compare the performance of two cache replacement algorithms:
 - What type of simulation model should be used?
 - How long should the simulation be run?
 - What can be done to get the same accuracy with a shorter run?
 - How can one decide if the random-number generator in the simulation is a good generator?

Part VI: Queueing Theory

- ❑ Introduction to Queueing Theory
- ❑ Analysis of A Single Queue
- ❑ Queueing Networks
- ❑ Operational Laws
- ❑ Mean Value Analysis and Related Techniques
- ❑ Convolution Algorithm
- ❑ Advanced Techniques

Example VI

- ❑ The average response time of a database system is three seconds. During a one-minute observation interval, the idle time on the system was ten seconds.

Using a queueing model for the system, determine the following:

- System utilization
- Average service time per query
- Number of queries completed during the observation interval
- Average number of jobs in the system
- Probability of number of jobs in the system being greater than 10
- 90-percentile response time
- 90-percentile waiting time

The Art of Performance Evaluation

- Given the same data, two analysts may interpret them differently.

Example:

- The throughputs of two systems A and B in transactions per second is as follows:

System	Workload 1	Workload 2
A	20	10
B	10	20

Possible Solutions

- Compare the average:

System	Workload 1	Workload 2	Average
A	20	10	15
B	10	20	15

Conclusion: The two systems are equally good.

- Compare the ratio with system B as the base

System	Workload 1	Workload 2	Average
A	2	0.5	1.25
B	1	1	1

Conclusion: System A is better than B.

Solutions (Cont)

- ❑ Compare the ratio with system A as the base

System	Workload 1	Workload 2	Average
A	1	1	1
B	0.5	2	1.25

Conclusion: System B is better than A.

- ❑ Similar games in: Selection of workload, Measuring the systems, Presenting the results.
- ❑ Common mistakes will also be discussed.

Grading

- ❑ Exams (Best of 2 mid terms + Final) 60%
- ❑ Class participation 5%
- ❑ Homeworks 15%
- ❑ Project 20%

Prerequisites

- ❑ CSE 131: Computer Science I
- ❑ CSE 126: Introduction To Computer Programming
- ❑ CSE 260M: ~~Introduction To Digital Logic And Computer Design~~ (Not required)
- ❑ Basic Probability and Statistics
- ❑ Matrix multiplication and inversion

Prerequisite

❑ Statistics:

- Mean, variance
- Normal distribution
- Density function, Distribution function
- Coefficient of variation
Correlation coefficient
- Median, mode, Quantile

❑ Programming

Tentative Schedule

Class	Day	Date	Topic	Chapter
1	Monday	1/12/15	Course Introduction	
2	Wednesday	1/14/15	Common Mistakes	2
			Selection of Techniques and Metrics	3
	Monday	1/19/15	<i>No Class: MLK Holiday</i>	
3	Wednesday	1/21/15	Types of Workloads	4
			Workload Selection	5
4	Monday	1/26/15	Data Presentation	10
5	Wednesday	1/28/15	Ratio Games	11
6	Monday	2/2/15	Summarizing Measured Data	12
7	Wednesday	2/4/15	Summarizing Measured Data (Cont)	12
8	Monday	2/9/15	Comparing Systems Using Random Data	13
9	Wednesday	2/11/15	Simple Linear Regression Models	14
10	Monday	2/16/15	Mid-Term Exam 1	

Tentative Schedule (Cont)

Class	Day	Date	Topic	Chapter
11	Wednesday	2/18/15	Other Regression Models	15
12	Monday	2/23/15	Project Guidelines	
			Other Regression Models (Cont)	15
13	Wednesday	2/25/15	Experimental Designs	16
			2k Experimental Designs	17
14	Monday	3/2/15	Factorial Designs with Replication	18
15	Wednesday	3/6/15	Project Guidelines (Part 2)	
	Monday	3/9/15	<i>WUSTL Spring Break (No Classes)</i>	
	Wednesday	3/11/15	<i>WUSTL Spring Break (No Classes)</i>	
16	Monday	3/16/15	Fractional Factorial Designs	19
			One Factor Experiments	20
17	Wednesday	3/18/15	Two Factor Full Factorial Design w/o Replications	21
18	Monday	3/23/15	Mid-Term Exam 2	

Tentative Schedule (Cont)

Class	Day	Date	Topic	Chapter
19	Wednesday	3/25/15	Two Factor Full Factorial Designs with Replications	22
			General Full Factorial Designs	23
20	Monday	3/30/15	Introduction to Queueing Theory	30
			Analysis of Single Queue	31
21	Wednesday	4/1/15	Queueing Networks	32
			Operational Laws	33
22	Monday	4/6/15	Mean-Value Analysis	34
23	Wednesday	4/8/15	Random Number Generation	26
24	Monday	4/13/15	Time Series Analysis	37
25	Wednesday	4/15/15	Time Series Analysis(Cont)	37
26	Monday	4/20/15	Heavy Tailed Distributions, Self-Similar Processes, and Long-Range Dependence	38
27	Wednesday	4/22/15	Final Exam	

Projects

- ❑ A survey paper on a performance topic
 - Workloads/Metrics/Analysis: Databases, Networks, Computer Systems, Web Servers, Graphics, Sensors, Distributed Systems
 - Comparison of Measurement, Modeling, Simulation, Analysis Tools: NS2
 - Comprehensive Survey:
Technical Papers, Industry Standards, Products
- ❑ A real case study on performance of a system you are already working on
- ❑ Average 6 Hrs/week/person on project + 9 Hrs/week/person on class
- ❑ Recent Developments: Last 2 to 4 years \Rightarrow Not in books
- ❑ Better ones may be submitted to magazines or journals

Projects (Cont)

- **Goal:** Provide an insight (or information) not obvious before the project.
- **Real Problems:** Thesis work, or job
- **Homeworks:** Apply techniques learnt to your system.

Example of Previous Case Studies

- ❑ Performance of Google App Engine and Amazon Web Service
- ❑ Availability and Sensitivity of Smart Grid Components
- ❑ Modeling and Analysis Issues in x86-based Hypervisors
- ❑ Image Sensor Performance
- ❑ Performance of Solving Laplace's Equation using Auto-Pipe
- ❑ Performance Modeling of Multi-core Processors
- ❑ Performance of Named Data Networking
- ❑ A Measurement Study of Packet Reception using Linux
- ❑ Performance Analysis of Robotics Systems
- ❑ Performance and Measurement Issues of Smart Phones Design
- ❑ Analysis of Online Social Networks
- ❑ Measurement Study on the BitTorrent File Distribution System
- ❑ A Survey of Wireless Sensor Network Simulation Tools

Project Schedule

Mon 2/23	Topic Selection
Mon 3/02	References Due
Mon 3/16	Outline Due
Mon 3/30	First Draft Due -> Peer reviewed
Mon 4/06	Reviews Returned
Mon 4/13	Final Report Due

Office Hours

- ❑ Monday/Wednesday: 11 AM to 12 noon
- ❑ Office: Bryan 523

- ❑ Teaching Assistant:
 - YiJian Li, liyijiankaka@gmail.com
 - Office Hours: Thursday/Sunday 1-2PM

Exams

- ❑ Exams consist of numerical, fill-in-the-blank and multiple-choice (true-false) questions.
- ❑ There is negative grading on incorrect multiple-choice questions. Grade: +1 for correct. $-1/(n-1)$ for incorrect.
⇒ For True-False: +1 for Correct, -1 for Incorrect
This ensures that random marking will produce an average of 0.
- ❑ Everyone including the graduating students are graded the same way.
- ❑ Highest score achieved becomes 100% for that exam.
⇒ Measures relative performance of the student
Effect of all other factors, such as time allotted, hardness of questions are eliminated.

Exams (Cont)

- ❑ All exams are closed book.
One 8.5”X11” cheat sheet with your notes on both sides is allowed.
- ❑ No smart phones allowed.
Only simple TI-30 or equivalent calculator allowed for calculations.
- ❑ Exam dates are fixed and there are no substitute exams
⇒ Plan your travel accordingly.
- ❑ Best of the two mid-terms is used.

Homework Submission

- ❑ All homeworks are due on the following Monday at the beginning of the class unless specified otherwise.
- ❑ Any late submissions, if allowed, will **always** have a penalty.
- ❑ All homeworks should be submitted in hardcopy
- ❑ All homeworks are identified by the class handout number.
- ❑ All homeworks should be on a separate sheet.
Your name should be on every page.
- ❑ Please write CSE567 in the subject field of all emails related to this course.
- ❑ Use word “Homework” in the subject field on emails related homework. Also indicate the homework number.


Homework Grading

- ❑ Grading basis: Method + Correct answer
- ❑ Show how you got your answer
 - Show intermediate calculations.
 - Show equations or formulas used.
 - If you use a spreadsheet, a statistical package, or write a program, print it out and turn it in with the homework.
 - For Excel, set the print area and scale the page accordingly to fit to a page. (See Page Setup)

Quizzes

- There will be a short 5-minute quiz at the beginning of each class to check if you have read the topics covered in the last class.

Summary



- ❑ Goal: To prepare you for correct analysis and modeling of any system
- ❑ There will be a self-reading and writing
- ❑ Get ready to work hard

Quiz 0: Prerequisites

True or False?

T F

- The mean of uniform(0,1) variates is 1.
- The sum of two normal variates with means 4 and 3 has a mean of 7.
- The probability of a fair coin coming up head once and tail once in two throws is 1.
- The density function $f(x)$ approaches 1 as x approaches ∞ .
- Given two variables, the variable with higher median also has a higher mean.
- The probability of a fair coin coming up heads twice in a row is $1/4$.
- The difference of two normal variates with means 4 and 3 has a mean of $4/3$.
- The cumulative distribution function $F(x)$ approaches 1 as x approaches ∞ .
- High coefficient of variation implies a low variance and vice versa.
- If x is 0, then after $x++$, x will be 1.

Marks = Correct Answers _____ - Incorrect Answers _____ = _____

Student Questionnaire

Name: _____

Email: _____

Phone: _____

Degree: _____ Expected Date: _____

Technical Interest Area(s):

Prior probability/statistics related courses/activities:

Prior computer systems related courses (Max 5):

