

COMPUTER NETWORK PERFORMANCE

Gaia Maselli

maselli@di.uniroma1.it

Room: 319



SAPIENZA
UNIVERSITÀ DI ROMA

Overview of first class

- Practical Info (schedule, exam, readings)
- Goal of this course
- Contents of the course
- Beginning of part I



Practical info

Schedule

Monday	12.00 – 14.00
Friday	10.00 – 12.00

Exam

- **Class participation and Homework** (I'll keep track of it)
- Present a paper/project (**oral presentation** during the course)

OR

- **Written exam** at the end of the course (for those who do not present a paper/project)

Prerequisites: Computer networks + basic probability and statistics

Course homepage

twiki.di.uniroma1.it/twiki/view/PDSDR/CNP1617



SAPIENZA
UNIVERSITÀ DI ROMA

Text Book and Suggested readings

BOOK

- R. Jain, “The Art of Computer Systems Performance Analysis,” Wiley
Available in the dept. library

PAPERS

- Research papers available on digital libraries (IEEE, ACM, Elsevier, etc.)



Performance evaluation: why?



Commercially

- What is the reach of the Internet?
- How many individuals are connected in a given area?
- What fraction of users have high speed Internet connection? And how many depends on dial up?
- Where should network access points be placed?
- Will users with wireless connectivity be able to access the Internet?



Socially

- Social implications of Internet use
- Understanding the amount of network activity involving various sites and protocols give considerable insight into social issues
- To characterize web site popularity and content
- Emerging protocols



Technically

- What is the delay (or latency) for a packet to traverse the networks?
- The design of network components (e.g. routers) and protocols is strongly driven by the nature of Internet workloads
 - Router designs strongly depends on the **statistical properties** of **network traffic** and **packet size distribution**
 - The statistical properties **of Web pages** influence the performance and design of Web servers and browsers
- The **popularity** of new applications can drive improvements to associated protocols (as in the case of the explosion of Web traffic, which motivated the improvement of the basic HTTP/1.0 protocol to yield HTTP/1.1)



But also..

1. I have an idea for a new network paradigm and want to test it
 2. I have an idea for a new MAC (or any other) protocol, would it be better than the existing ones?
- ***If you are doing any research, whatever you are doing, you need to***
 - ***analyze that***
 - ***compare different alternatives to find the best one***
 - Performance is a key criterion in the design, procurement, and use of computer networks
 - Get the highest performance at a given (or the lowest) cost



Performance evaluation: goals

Two main goals

1. Assess the performance of an existing network
2. Study a new system before implementing it

Performance evaluation is required at every stage in the life cycle of a system

- Find the best design
- Even if there are no alternatives, evaluating the current system helps in determining how well it is performing and whether any improvements need to be made



How do we evaluate performance?

- Measurements
 - Internet, LAN, WAN
 - Sensor and ad hoc networks, RFID
- Analytical evaluation
 - Queueing networks
 - Probabilistic models (ex. balls and bins, etc.)
- Simulation
 - Network Simulator (NS2)



Goal of this course

- Comprehensive course on performance analysis, including
 - Measurement
 - Modeling
 - Experimental design
 - Simulation
 - Analytical evaluation
- How to avoid common mistakes in performance analysis



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 (input) on the system (workload characterization)
- Determining the number and sizes of components (capacity planning)
- Predicting the performance at future loads (forecasting).



Basic terms

- **System:** anything that you want to study (any collection of hardware, software, and firmware components)
- **Metrics:** criteria used to evaluate the performance of a system (e.g., response time)
- **Workloads:** the requests made by the users of the system (anything that goes in the system).



After the course you should be able to

- **Select appropriate evaluation techniques, performance metrics and workload for a system**
 - Measurement, simulation or analytical modeling?
 - What to measure? (criteria used to evaluate performance of the system)
 - What is the input of the system? (ex. Number of requests per unit of time arriving at a web server)
- After a few lectures you should be able to answer the following question:
 - What performance metrics should be used to compare the performance of the following systems?
 - Two packet retransmission algorithms
 - Two MAC protocols



After the course you should be able to

- **Conduct performance measurements correctly**
 - Which type of tool should be used to measure the results?
 - To measure the performance of a computer network you need at least two tools – a tool to load the system (**load generator**) and a tool to measure the results (**monitor**)
- After this course you should be able to answer to the following problem:
 - Which type of monitor would be more suitable for measuring the response time of packets in a networks?



After the course you should be able to

- **Use proper statistical techniques to compare several alternatives**
 - Most performance evaluation problems consist of *finding the best* among a number of alternatives.
 - If a measurement or simulation is *repeated several times*, generally the *results would be slightly different each time*. How should we compare them?
 - Simply comparing the average result of a number of repeated trials does not lead to correct conclusions, particularly if the variability of the result is high
- After the course you should be able to answer to the following question:
 - The number of packets lost on two links were measured for four file sizes as shown in the table

File size	Link A	LinkB
1000	5	10
1200	7	3
1300	3	0
50	0	1

Which link is better?



After the course you should be able to

- **Design measurement and simulation experiments to provide the most information with the least effort**
 - Given a number of factors that affect the system performance, it is useful to separate out the effects of individual factors
 - Obtain maximum information with a minimum number of experiments (techniques to organize experiments will be given)
- A question you should be able to answer:
 - The performance of the network depends on the following factors:
 - Link rate
 - Number of nodes in the network
 - How many experiments are needed?
 - How does one estimate the impact of each factor on the system performance?



After the course you should be able to

- **Perform simulations correctly**
 - In designing a simulation model, one has to *select*
 - *language*
 - *seeds and algorithms for random number generation*
 - *decide the length of simulation run* and analyze simulation results
- A question you should be able to answer:
 - In order to compare the performance of two MAC protocols:
 - What type of simulation model should be used?
 - How long should the simulation be run?



After the course you should be able to

- **Use simple analytical models to analyze the network performance**
 - Queueing models
 - Probabilistic models
- A question you should be able to answer:
 - The average response time of a web server is 2 seconds. During a 1-minute observational interval, the idle time on the system was 5 seconds. Using a queueing model for the system, determine the following
 - System utilization
 - Average service time per request
 - Number of requests completed during the observation interval
 - 90-percentile response time



Part I: An Overview of Performance Evaluation

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



The art of performance evaluation

- Like a work of art, successful evaluation **cannot be produced mechanically**
- Every evaluation requires an intimate knowledge of the system being modeled and a **careful selection of the methodology, workload, and tools**
- Like an artist, each analyst has a unique style.
 - Given the same problem, two analysts may choose different performance metrics and evaluation methodologies
 - Given the same data, two analysts may interpret them differently.



The art of performance evaluation: example

- 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

An example of the performance games people play to show that their system is better



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.

- In the second and third case the technique known as **ratio game** has been applied
- Intentional game: the proponents of a system want to show the superiority of their proposed alternatives
- Unintentional: simply a result of a lack of knowledge of performance evaluation techniques
- A knowledge of common mistakes and games helps in understanding the importance of proper methodology



Common mistakes in Performance Evaluation and How to Avoid Them



Common mistakes in performance evaluation (1/6)

- **No goals**

- the need for a goal may sound obvious, but many performance efforts are started without any clear goals.
- There is no general-purpose model. Each model must be developed with a particular goal in mind: the metrics, workload and methodology all depend upon the goal
- Describe goals and then design experiments

- **Biased goals**

- Don't show that YOUR system better than HERS
- Analyst = jury (neutral!)



Common mistakes in performance evaluation (2/6)

- **Unsystematic approach**
 - Often analysis adopt an unsystematic approach selecting system parameters, metrics and workloads arbitrarily. This leads to inaccurate conclusions
- **Incorrect performance metrics**
 - Do not choose metrics that can be easily computed or measured but the ones that are relevant (metrics difficult to compute are often ignored)
- **Unrepresentative Workload**
 - should be representative of how the system will work “in the wild” (actual usage of the system)
 - Ex: large and small packets? Don't test with only large or only small packets otherwise you get inaccurate conclusions
- **Wrong Evaluation Technique**
 - Use most appropriate: model, simulation, measurement
 - Do not use the model that you can best solve but that one that best solves the problem



Common mistakes in performance evaluation (3/6)

- **Overlooking important parameters**

- Make a complete list of *system* and *workload characteristics* that affect the performance of the system

- **Ignoring significant factors**

- *Parameters that are varied* in the study are called *factors*.
- The choice of factors should be based on their relevance and not on the analyst's knowledge of the factors.
- *It is important to identify those parameters which if varied will make a significant impact on the performance*
- Example: if packet arrival rate rather than packet size affects the response time of a network gateway, it would be better to use several different arrival rates in studying its performance



Common mistakes in performance evaluation (4/6)

- **Inappropriate experimental design**

- Experimental design relates to the number of measurement or simulation experiments to be conducted and the parameter values used in each experiment.
- Proper selection of these values can lead to more information from the same number of experiment. Improper selection can result in a waste of the analyst's time and resources

- **No analysis**

- If performance analysts lack expertise in *data analysis*, they may collect enormous amounts of data but do not know how to analyze or interpret them (files full of data without any summary)



Common mistakes in performance evaluation (5/6)

- **Erroneous analysis**
 - Taking the average of ratios and too short simulations
- **No sensitivity analysis**
 - Results may be sensitive to the workload and system parameters
 - Without a sensitivity analysis one cannot be sure if the conclusions would change if the analysis was done in a slightly different setting
- **Improper treatment of outliers**
 - *Values that are too high or too low compared to a majority of values in a set are called outliers.*
 - If the outlier is not caused by a real system phenomenon, it should be ignored.
 - If the outlier is a possible occurrence in a real system, it should be appropriately included in the model
 - Deciding which outliers should be ignored and which should be included is part of the art of performance evaluation and requires careful understanding of the system being modeled.



Common mistakes in performance evaluation (6/6)

- **Omitting assumptions and limitations**

- Omitting the assumptions may lead the user to apply the analysis to another context where the assumptions at the will not be valid
- Omitting limitations may lead to make conclusions about the environment to which the analysis does not apply

- **Improper presentation of results**

- The aim of every performance study is to help in decision making
- An analysis that does *not produce any useful results* is a failure
- An analysis with *results that cannot be understood* by decision maker is a failure



A systematic approach to performance evaluation

1. State Goals and Define the System

- State the goals of the study and define what constitutes the system by delineating system boundaries (given the same set of hardware and software, the definition of the system may vary depending upon the goals of the study)

2. List Services and Outcomes

- Each system provides a set of services.
- A computer network allows its users to send packets to specified destination
- When a user requests a service, there are a number of possible outcomes.
 - Desirable (the packet arrives to the destination)
 - Undesirable (the packet gets lost)



A systematic approach to performance evaluation

3. Select Metrics

- Select criteria to compare performance
- Metrics are related to speed, accuracy, and availability of services

4. List Parameters

- Make a list of all the parameters that affect the system (anything that you can change)
 - System parameters (do not vary among various installation of the system)
 - Workload parameters (depend on users' requests and vary from one installation to the next)

5. Select Factors to Study

- Select the parameters that you want to study (they will change during the evaluation)
- If you ignore significant factors you get meaningless results



A systematic approach to performance evaluation

6. Select Evaluation Technique

- Analytical modeling, simulation, or measuring?
- The choice depends upon the time and resources available

7. Select Workload

- list of service requests to the system
 - Probability of various requests
 - Trace of requests measured on a real system

8. Design Experiments

- Decide a sequence of experiments that maximize information with minimal effort

9. Analyze and Interpret Data

- Analysis only produces results (big amount of results) and not conclusions
- Draw conclusions from results (interpretation of data)

10. Present Results

- Make graphs (results must be easy to understood)

Repeat (go back and reconsider some decisions)



Homework

- Choose a network system for performance study. Briefly describe the system and list:
 1. Services
 2. Performance metrics
 3. Workload
 4. Factors
 5. Evaluation technique

