COMPUTER NETWORKS PERFORMANCE

Gaia Maselli
maselli@di.uniroma1.it

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

- I part (From today to mid/end nov.) ➔ professor
- II part (remaining part up to end dec.) ➔ students (exam)

Exam

- Class participation
- You are expected to so something on your own
  - Independent reading
  - Present a paper or a project

Course homepage

http://twiki.di.uniroma1.it/twiki/view/PDSDR/20152016

Suggested readings

BOOKS

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

- M. Crovella, B. Krishnamurthy, “Internet Measurement, infrastructure, traffic & applications”, Wiley

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?
- What is the end-to-end throughput expected when transmitting a large data file across a network?
- 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

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
  - 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
- **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?
  - What is the input of the system?
- **Conduct performance measurements correctly**
  - Which type of tool should be used to measure the results?
- **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?

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
- **Perform simulations correctly**
  - In designing a simulation model, one has to *select*
    - *language*
    - *seeds and algorithms* for *random number generation*
  - And *decide the length of simulation run* and a analyze simulation results
- **Using simple analytical models to analyze the network performance**
  - Queueing models
  - Probabilistic models

Main parts of the course

- Part I: An Overview of Performance Evaluation
- Part II: network measurement (Internet and other)
- Part III: Network simulation
- Part IV: analytical modeling and analysis

Part I: An Overview of Performance Evaluation

- Introduction
- Common Mistakes and How To Avoid Them
- Selection of Techniques and Metrics
- The art of data presentation
  - Any analysis comes out with some results. We need to present results correctly (critical aspect)

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

Part II: network measurement

- Internet measurement
  - Practical issues in Internet measurement (where, role of time, etc.)
  - Infrastructure
  - Traffic
  - Applications (Web and DNS)

Part III: Network simulation

- Introduction to simulation
- Analysis of simulation results
- NS2: a network simulator

Part IV: analytical modeling and analysis

- Introduction to queueing theory and queueing networks
- Analytical performance evaluation: specific network contexts (sensor, RFID)

Part I:
Overview and common mistakes in Performance Evaluation

SAPIENZA
UNIVERSITÀ DI ROMA

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

Common mistakes in performance evaluation (2/6)

- **Unsystematic approach**
  - Identify a complete (not arbitrary) set of goals, parameters, metrics and workloads
- **Incorrect performance metrics**
  - Do not choose metrics that can be easily computed or measured but the ones that are relevant
- **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

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**
   - Compare a new MAC protocol for sensor networks with an existing one
   - Sensor network (wireless sensor nodes with batteries and limited transmission range)
2. **List Services and Outcomes**
   - Send packets to a specific node called sink
   - Success or failure in packet delivery
3. **Select Metrics**
   - Packet delivery ratio, energy consumption, network lifetime
4. **List Parameters**
   - number of nodes, duty cycle, transmission energy cost, etc.
5. **Select Factors to Study**
   - Message inter-arrival period, duty cycle

A systematic approach to performance evaluation

6. **Select Evaluation Technique** (depends upon the time and resources available)
   - simulation
7. **Select Workload**
   - number of data flows in the network
8. **Design Experiments**
   - Decide a sequence of experiments that maximize information with minimal effort
9. **Analyze and Interpret Data**
   - Draw conclusions from results (the duty cycle of 0.1)
10. **Present Results**
    - Make graphs (ex., plot energy consumption by varying message inter arrival time)

Repeat (go back and reconsider some decisions)