EECS 454: Advanced Communication Networks

Spring Quarter 2010

Meeting time: 3:30-4:50 MW

Instructor: Randall Berry
Office: Tech, Rm. M318
Office Hours: by appointment

Course Overview

- Primary goal is to develop *analytical tools* and *conceptual models* that are useful in networking research (and in other fields).
- Focus of the course is on issues identified with data link, network and transport layer - physical layer and application layer issues will not be addressed in detail.
- This graduate level course - you should question what you are learning and think about how what you are learning may be applied in other situations. What are the limitations, key assumptions, etc.

Course Information

Prerequisites:

- **Good understanding of basic probability.** (If you are not comfortable with probability, it may be helpful to take EECS 422 before taking this course.)
- Familiarity with data networks (e.g. EECS 333 or 340) is helpful (this provides motivation/context for the material studied here).

Text: *Data Networks*, 2nd Ed. by D. Betsekas and R. Gallager.

For the first half of the course we will follow parts the text closely; for the second half, the text will be supplemented with other notes and journal papers. Other references are listed on the course information sheet.

Grading

- Problem Sets - can work in groups - write-up your own solutions.
- 1st mid-term - in-class, about 1/2 way through term.
- 2nd mid-term - take home exam - last week of classes
- Final Project - presentation and write-up during last week of classes or final exam week.

EECS 454 vs. EECS 333

- Fewer topics than in 333, but cover each topic in more depth.
- Less emphasis than 333 on describing actual protocols and implementation issues, more emphasis on analytical techniques and performance issues.

Research in Networking

- Multidisciplinary/multifaceted.
- The material in this course focuses on the analytical side.
- But this far from the whole story, also a lot of work on measurements, simulation, implementation, protocol design, etc.

Analytical Modeling

Most of this course consists of describing different analytical techniques that can be used for understanding various aspects of networking.

Uses:

- Performance analysis
  - Parameter setting, network provisioning, comparison of different approaches.
- Improve understanding/intuition
  - qualitative behavior, performance trade-offs, system bottlenecks.
- Establish fundamental limitations
  - What is best performance possible?

Analytical Modeling

- Important point to remember: *All models are wrong.*
- So what makes a good model?

Analytical Modeling

Two uses:

1. Gain deeper insight into performance
   - Often simplified “toy” models - need to abstract away much detail to get something you can really understand.
   - Need a good understanding of what is being modeled and of the modeling techniques used.
   - Often developing such models is an iterative process.
2. Performance analysis
   - More detailed - want to acurately capture actual system performance- often too complicated to analyze by hand.

Communication Networks

- Physically communication networks consist of “communication links” connecting together different “nodes” for the purpose of exchanging information (*bits*).
- Here focus mainly on “wire-line” network models (point-to-point links).
- Each link can be thought of as a (possibly lossy) “bit-pipe” with a certain transmission rate.
- Nodes - sources, sinks, routers/switches.
- Mathematically - represent topology as a *graph*.

Networking motivations

- Networks are used to provide different *information services* to end-users (e.g. e-mail, ftp, voice telephony, streaming video,..).
- Key reason for using networks is *resource sharing* (economies of scale/economies of scope).
- Different applications require different *qualities of service* (QoS).
- A key issue is how to share resources and satisfy QoS.

Packets, Sessions, etc.

- “Users” initiate “sessions” and exchange “messages”.
  - connection-oriented/connectionless
- Messages may be broken into smaller “packets” to be sent over network.
  - Often special names used in particular contexts: frames, cells, ...

Multiplexing

- Multiplexing refers to techniques for sharing a link amoung different sessions.
- Two main techniques
  - circuit-based: e.g. TDM/FDM.
    - Performance: blocking probability.
  - Packet-based: statistical multiplexing/scheduling.
    - performance: delay/packet dropping.
- Various hybrids also possible.

Switching/routing

A related issue to mulitplexing is how packets are switched and routed within the network.

- **Routing:** how to determine the next link for a packet.
- **Switching:** how nodes physically send packets from one link to the next.

Network Control

- Another key issue is how to “control” the traffic in the network.
- A variety of possible techniques:
  - Flow control, congestion control, admission control..

Network Layers and Protocols

- Conceptually, we think of networks as a sequence of vertical *layers*; each layer providing some type of *service* to lower layers.
- Higher layers provide a higher level abstraction of the network.
- Each layer provides its service by implementing one or more*communication protocols*.
- A protocol specifies how different *processes* in a network interact; this includes specifying the format of *messages* that are exchanged and the *algorithm* used to generate these messages.
  - A fundamental characteristic of most of the algorithms used in networks is that they are *distributed*.

Perfomance Analysis

The two most common performance metrics used in networking are

- **Delay** - i.e., how long does it take to send information from its source to its destination.
- **Throughput** - i.e. how much data per second can be sent across the network.

Much of the behaviour in networks is best modeled as random (e.g. user behaviour, failures); thus performance analysis is typically done in a probabalistic framework.

Modleing scales

- Traffic in a network can be modeled at different scales.
- At the finest scale we have the arrivals and departures of individual packets in a session.
    - Such models are refered to as ”packet-level models.”
- At a larger scale, one can look at ”flow level models.”
    - Focus on the arrival and departure of flows as the random quantity of interest.
    - Often model the transmission of packets within a flow as a ”fluid” process.
    - Can think of as looking at system on a longer time-scale.
- In some cases, it is even reasonable to focus on a fixed set of active flows.
    - This leads to deterministic models, e.g. diferential equation models and optimization-based models of TCP.

Course Outline

- Quick Probability review.
- Delay models - queueing theory, Markov chains, networks of queues, long-range dependence.
- Switching - input/output switches/ stability/matching.
- Routing - Bellman-ford, adaptive routing, optimal routing, topological design.
- Flow control - TCP, fairness, token buckets, network calculus, pricing models.