Computer Communication Networks CSI 416/516

Lecture 1: Introduction to Computer Communication Networks

Stephen F. Bush

GE Global Research

September 20, 2009

Lecture Outline

1. Introduction
2. Syllabus and Grading
3. What is Networking?
4. Networks
5. Protocols
6. Network Graphs and Link Capacity

Computer Communication Networks CSI 416/516

Instructor:	Dr. Stephen F. Bush, GE Global Research
Phone:	387-6827
Email:	bushsf@research.ge.com
Office Hours:	Tue 10:05pm HU-132 and by appointment
Course Website:	http://www.cs.albany.edu/~bushsf

Syllabus

Week	Topics	Readings	Notes
1	Intro to Networking	L1 B1	
	Layered Architectures	L2 B2.1-2.2	
2	Layering	L2 B2.3	HW 1 Assigned
	Sockets		
3	Information Theory		HW 1 due
	Compression & Correction	L3.9	HW 2 Assigned
4	Networks	L7	
	The Physical Layer	L3, L12.1-12.3	
5	Multiplexing and Switching	L7.3	
	The Telephone Network	L4	**Proposals Due**
6	Active Networks	B3	
7	Mid Term Exam		

B Active Networks and Active Network Management: A Proactive Management Framework by Stephen F. Bush and Amit Kulkarni, Kluwer Academic/Plenum Publishers, New York, Boston, Dordrecht, London, Moscow, 2001, 196 pp. Hardbound, ISBN 0-306-46560-4.

L Communications Networks: Fundamental Concepts and Key Architectures, Leon-Garcia and Widjaja, McGraw Hill, 2003.

Syllabus

Week	Topics	Readings	Notes
8	Queuing Theory	L App A	HW2 Due
9	Network Simulation		HW3 Assigned
	NS-2		
10	Peer to Peer	L5	
	ARQ and Flow Control		
11	Multiple Access	L6.1-6.2	HW3 Due
	Ethernet	L6.6-6.11	HW4 Assigned
12	Scheduling	L6.3-6.4	
	QoS		
13	Routing	L7	HW4 Due
	Ad Hoc Networking		(optional)
14	Sensor Networks		HW5 Assigned
15	Student Presentation	L App B B4-6	
	Student Presentations		
16	Network Management		HW5 Due
	Final Exam and Vacation!		

B Active Networks and Active Network Management: A Proactive Management Framework by Stephen F. Bush and Amit Kulkarni, Kluwer Academic/Plenum Publishers, New York, Boston, Dordrecht, London, Moscow, 2001, 196 pp. Hardbound, ISBN 0-306-46560-4.

L Communications Networks: Fundamental Concepts and Key Architectures, Leon-Garcia and Widjaja, McGraw Hill, 2003.

Scribe Schedule

Week	Scribe 1	Scribe 2	Scribe 3
1 (Intro)	Caldara,Logan Nathan	Cagan,Ferhat	
2 (Layering)	Crisafulli,Anthony Nicholas	Erbatur,Serdar	
3 (Info Theory)	Krishnaraj Ravindranathan		
4 (Phy Layer)		Kudlack,Edward A	
5 (Switching)	Lin,Xu		
6 (Active Networks)	Krishnaraj Ravindranathan		
7 (Midterm)			
8 (Queuing Theory)		Caldara,Logan Nathan	
9 (ns-2)	Cagan,Ferhat		
10 (ARQ)		Crisafulli,Anthony Nicholas	
11 (Multiple Access)		Erbatur,Serdar	
12 (Sched+Routing)	Lin,Xu		
13 (Proj Pres)		Subramaniam,Sathiyaseelan	
14 (Routing)		Caldara,Logan Nathan	

i>Clickers

- We will use i>Clickers clickers in the lectures for quizzes
- Along with scribe notes, they will be worth 10% of your course grade
- It is your responsibility to attend lecture, bring your clicker, have fresh or spare batteries for your clicker, etc
- To register your clicker, go to `www.iclicker.com/registration`
  - For the Student ID field, give it the first part of your UAlbany Webmail userid (for instance, if your UAlbany webmail was ab123456@albany.edu, you would type in ab123456)
  - Please do not enter your actual UAlbany ID number

Why we are using i>Clickers

- To more accurately determine your preparation and background before learning new material
- To assess how well you understand material that we've covered
- To solicit your anonymous and honest opinion and feedback on how I'm doing
- To help you engage in more active participation in the lecture

Scribe System

- Two different students will be assigned to take notes for each class, starting today
- Your notes will be in LaTeX; notes are due before the next class
- A LaTeXtemplate is available on the class web page: `www.cs.albany.edu/~bushsf`; instructions embedded inside the template
- A shared resource for all class participants; you will be depending on others' taking good notes
- Grade will focus on content of notes; too much improper LaTeXsyntax, then points deducted
- Please email me (`bushsf@research.ge.com`) your contact information

Syllabus

Week	Topics	Readings	Notes
1	Intro to Networking	L1 B1	
	Layered Architectures	L2 B2.1-2.2	
2	Layering	L2 B2.3	HW 1 Assigned
	Sockets		
3	Information Theory		HW 1 due
	Compression & Correction	L3.9	HW 2 Assigned
4	Networks	L7	
	The Physical Layer	L3, L12.1-12.3	
5	Multiplexing and Switching	L7.3	
	The Telephone Network	L4	**Proposals Due**
6	Queuing Theory	L App A	HW2 Due
	Mid Term Exam		

B Active Networks and Active Network Management: A Proactive Management Framework by Stephen F. Bush and Amit Kulkarni, Kluwer Academic/Plenum Publishers, New York, Boston, Dordrecht, London, Moscow, 2001, 196 pp. Hardbound, ISBN 0-306-46560-4.

L Communications Networks: Fundamental Concepts and Key Architectures, Leon-Garcia and Widjaja, McGraw Hill, 2003.

Syllabus

Week	Topics	Readings	Notes
7	Network Simulation		HW3 Assigned
	NS-2		
8	Peer to Peer	L5	
	ARQ and Flow Control		
9	Multiple Access	L6.1-6.2	HW3 Due
	Ethernet	L6.6-6.11	HW4 Assigned
10	Scheduling	L6.3-6.4	
	QoS		
11	Routing	L7	HW4 Due
	Ad Hoc Networking		
12	Sensor Networks		HW5 Assigned
	Active Networks	B3	
13	Network Management	L App B B4-6	
	Student Presentations		
14	Student Presentations		HW5 Due
	Final Exam and Vacation!		

B Active Networks and Active Network Management: A Proactive Management Framework by Stephen F. Bush and Amit Kulkarni, Kluwer Academic/Plenum Publishers, New York, Boston, Dordrecht, London, Moscow, 2001, 196 pp. Hardbound, ISBN 0-306-46560-4.

L Communications Networks: Fundamental Concepts and Key Architectures, Leon-Garcia and Widjaja, McGraw Hill, 2003.

Nanoscale networking

- The topic of nanoscale networking will be highlighted within this course
- This topic will be integrated with the topics of traditional networking
- Special handouts on this topic will be provided
- The material will be included in homework and exams
- Nanoscale networking topics are encouraged for as a final class project

Grading

Component	Weight	Total Grade
Midterm/final	20% each	40%
1 S/W Project	10%	50%
Final Project	30%	80%
Assignments (HW)	10%	90%
Scribe Notes	5%	95%
i>Clicker Responses	5%	100%

Final Project

- Final Project has same weight as Midterm and Final; Start now!
- Final Project is a computer simulation (ns-2)
- Install ns-2 **now** and become familiar with it
- Visit `http://www.isi.edu/nsnam/ns/` to download and install ns-2
- We will discuss ns-2 later, however, if you wait until then to start working with it, it will be too late!
- Try modeling ideas as we progress through the course

Getting Started With NS-2

- Release 2.30 (released Sept 26, 2006) on SourceForge.Net:
- 'install' and 'validate' the code by following instructions
  - Validation can take a long time ($\approx 1/2$ hour)
- Note: Cygwin should have X11 installed in order to run "nam"
- Issue 'startx' – to begin terminal in cygwin

Running ns at CS Dept

- Make sure the following is in your path:
  - `itsunix.albany.edu: /usr/local/depts/cs/ns-2.29/bin/`
- Set the X11 display (if running remotely):
  - set DISPLAY `<hostname>:0`

Documentation

- General ns-2 manual: `http://www.isi.edu/nsnam/ns/`
- Documentation of individual models is poor
  - Some under ns-2.30/doc
  - Look directly at implementation ns-2.30/tcl/lib:
    - ns-queue.tcl
    - ns-mobilenode.tcl
    - ns-default.tcl ***
    - ns-lib.tcl ***

Goals

Goal

This course provides a combined applied / theoretical background in networks and data communications which will focus on the fundamentals of good science and engineering

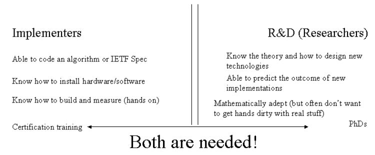

Assumed Background

1. Mathematics
   - Basic Calculus / Numerical Analysis
   - Probability (exponential distribution, poisson processes)
   - Scheduling
   - Linear Algebra and Optimization
2. Computer Science – you must be able to program
   - Computer Architecture
   - Basic Performance Analysis
   - Operating Systems and Systems Programming
   - Software Design
   - Data Structure/Algorithms
3. Good (personal) communications skills – otherwise no one will know what you have done

Teaching Goals

The design of this course and its policies attempts to:

1. Prepare and reward good students because
   - Networks impact quality of life
   - Unskilled practitioners are dangerous!
   - Skilled practitioners are valuable
2. Improve the students skills in:
   - Network design and analysis
   - Network systems software design and implementation
   - Performance analysis
   - Performance tuning
   - Documentation design

   This should help your academic/professional reputation and career

How To Succeed

1. Review prerequisite topics as needed
2. Start projects early
3. If stuck, try:
   - To formulate a clear problem statement
   - Check your resources (including the project's frequently asked questions (FAQ)) on the course website
   - If you are still stuck, ask a classmate (if appropriate) or us but don't resort to cheating!
4. Try the in class exercises and ungraded problems. They might help come exam time!

Hints for Success

- Take the initiative
  - If you don't know something find out or ASK!
- Start homework, projects, and **especially the final project** EARLY
- Become comfortable using a math package (Matlab, Mathematica, Maple, etc...)

Science vs. Engineering a Refresher

Some good approaches to attack problems are traditional methods:

The Scientific Method Used to understand/predict a processes behavior

- **Problem statement** Select which process you want to understand.
- **Hypothesis** Make a measurable guess (prediction) about the behavior.
- **Experiment** Test the hypothesis and statistically confirm or refute it.

Engineering/applied math focuses on controlling a process or deriving a solution:

- **Problem statement** There is a math problem to solve or a phenomena to control.
- **Design** Come up with a solution
- **Analysis** Measure the effectiveness of the solution.

Good Engineering

- Know what you (really your customer) wants
  - Too often engineers have great solutions to the wrong problems
- Measurement is required!
  - Too often there is a lot of verbal hype without any real precision to back it up
- Cost and performance
  - How do these trade-off? A great solution that is too costly will not fly

What is Networking?

Discussion

- Consider:
  - Plain old telephone
  - Web application
- What are the similarities and differences?

What Is Networking?

Common Network Technology:

- Profoundly impacts the quality of life!
- Is digital (encodes data in discrete values)
- Affected by analog phenomena
- Drives and uses theory/experiment
- Is used in ways not foreseen by designers
- Example: www was a surprise to Internet designers

Common Networks

Some well-known networks include:

1. Postal (non-electronic)
2. Telegraph (first digital electronic network)
3. Telephone (DSL)
4. Broadcast (Radio and Television)
5. Cable Television (and modems)
6. Internet
7. Sensor networks
8. Nano networks
9. Biological networks
10. i>Clickers

Question

What are the common elements?

Background Questions

i>Clicker

- What application of networking most interests you?
  - A Telephone (DSL)
  - B Broadcast (Radio and Television)
  - C Cable Television (and modems)
  - D Internet
  - E Sensor networks
  - F Nano networks
  - G Biological networks

Drivers

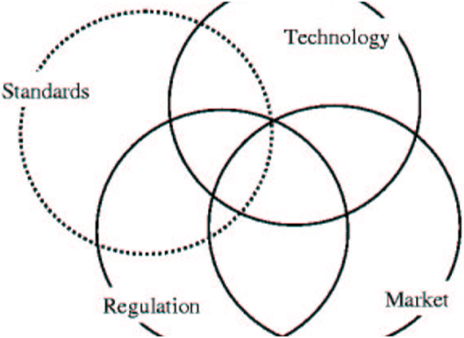

Directions in Network Technology

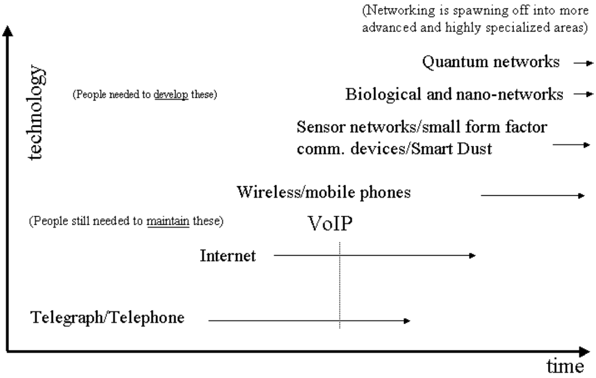

The theory in this course applies regardless of software or technology!

Networks Are Powerful: Phase Change and Power Law

- Each new connection joins more than just two users
  - Everyone connected to those two users is joined
- Consider random connections of bits of string
  - Sudden phase change
- Real Internet does not have random connections
  - Power law degree (k) → some large hubs and many small ones
- $P(k) \approx k^{-\nu}$

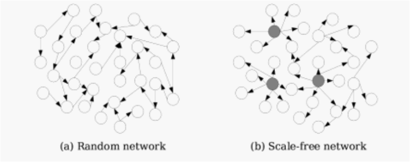

Scale-free Distributions

- A Power Law is a function $f(x)$ where the value y is proportional to some power of the input x, $y = f(x) \propto x^{-a}$
- For binomial, normal, and Poisson distributions the tail probabilities approach zero exponentially fast
- This means that the probability of a large value shrinks quickly and becomes vanishingly small
- Power law distribution have, on the other hand, a fat or heavy tail
    - This means that the probability of encountering a large value is small, but it does happen

Protocol

- $\pi\rho o\tau o$ (first) $\kappa o\lambda\lambda o\nu$ (page)
- Obviously, communication requires that the same language, i.e. network protocol be spoken and understood
- Standards Committees: IETF, IEEE
- Protocols for Communication: TCP/IP
- Protocols for Network Management: SNMP

Deployment Dominates Costs

- Developing and deploying new protocols is tedious, time-consuming, expensive
  - IPv6 disaster (a decade of work and still undergoing testing and deployment)
- Active networking is a brave attempt to solve this problem

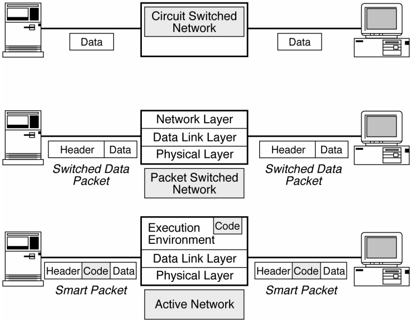

Data

- Data is encoded using *signals*, which are either:

Digital	(discrete) "waves"
Analog	(continuous) waves

- Signals are described by their:

Frequency	Cycles per unit time
Amplitude	Distance between crest and bottom
	Frequency's reciprocal is the *period*
Phase	Horizontal displacement

Software radios – flexibility with interchangeable "waveforms"

Media

- Signals propagate over media, which can be:

  Air/Space Cellular telephone, wireless networks

  Wires Ethernet, traditional telephone, optical fiber, ATM

- Let's consider how to analyze a network...

Bandwidth

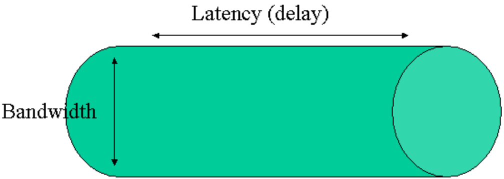

Example: High BW, long delay (CD carried across town)

Example: Low BW, short delay (pager)

Delay-Bandwidth product (bps x s) == bits in pipe

Networks of Links

- Assume we have a directed graph representing network links
- Each link (graph edge) has a maximum capacity K.
- How do we find the maximum flow through an arbitrary network of links?
- Solution: consider a matrix representation of the network and use linear algebra

The Adjacency Matrix

- The adjacency matrix of a finite directed or undirected graph G on n vertices is the $n \times n$ matrix where the non-diagonal entry a_{ij} is the number of edges from vertex i to vertex j
- In the special case of a finite simple graph, the adjacency matrix is a (0,1)-matrix with zeros on its diagonal
- If the graph is undirected, the adjacency matrix is symmetric

- Another matrix representation for a graph is the incidence matrix
- The relationship between a graph and the eigenvalues and eigenvectors of its adjacency matrix is studied in spectral graph theory

Exercise: What is the adjacency matrix for this graph?

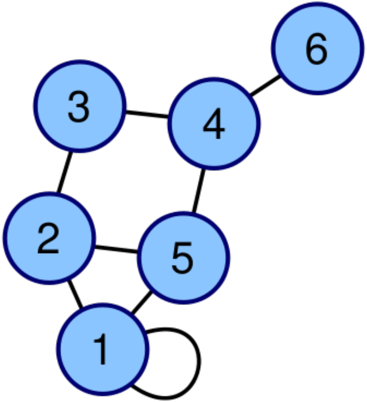

Solution

$$\begin{pmatrix} 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

Properties of the Adjacency Matrix

- The adjacency matrix of an undirected graph is symmetric, and therefore has a complete set of real eigenvalues and an orthogonal eigenvector basis. The set of eigenvalues of a graph is the spectrum of the graph.
- Suppose two directed or undirected graphs G_1 and G_2 with adjacency matrices A_1 and A_2 are given. G_1 and G_2 are isomorphic if and only if there exists a permutation matrix P such that $PA_1P - 1 = A_2$.
- In particular, A_1 and A_2 are similar and therefore have the same minimal polynomial, characteristic polynomial, eigenvalues, determinant and trace. These can therefore serve as isomorphism invariants of graphs. However, two graphs may possess the same set of eigenvalues but not be isomorphic (one cannot 'hear' the shape of a graph).
- If A is the adjacency matrix of the directed or undirected graph G, then the matrix A_n (i.e., the matrix product of n copies of A) has an interesting interpretation: the entry in row i and column j gives the number of (directed or undirected) paths of length n from vertex i to vertex j.
- The matrix $I - A$ (where I denotes the $n \times n$ identity matrix) is invertible if and only if there are no directed cycles in the graph G. In this case, the inverse $(I - A) - 1$ has the following interpretation: the entry in row i and column j gives the number of directed paths from vertex i to vertex j (which is always finite if there are no directed cycles). This can be understood using the geometric series for matrices: $(I - A) - 1 = I + A + A2 + A3 + \ldots$ corresponding to the fact that the number of paths from i to j equals the number of paths of length 0 plus the number of paths of length 1 plus the number of paths of length 2, etc.
- The main diagonal of every adjacency matrix corresponding to a graph without loops has all zero entries.

The Incidence Matrix

- The incidence matrix of a directed graph D is a $p \times q$ matrix $[b_{ij}]$ where p and q are the number of vertices and edges respectively, such that
- $b_{ij} = -1$ if the edge x_j leaves vertex v_i,
- 1 if it enters vertex v_i and
- 0 otherwise
- Many authors may also use the opposite sign convention!

Maximum flow parameters

- We will let f be a scalar value representing the maximum flow
- Let e be a vector with a length being the number of nodes and with +1 for flow input node, -1 for the flow output node, and 0 for all other nodes
- Let x be a vector with a length being the number of edges where each component is the flow though that edge
- 0 here is a vector of length number of nodes where all elements are zero
- We want to find the maximum flow through the network given the topology and the maximum link capacities

Network Link Capacity Limits

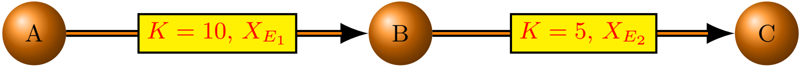

Node by edge incidence matrix:

$$\begin{array}{ccc} & X_{E_1} & X_{E_2} \\ A & -1 & 0 \\ B & 1 & -1 \\ C & 0 & 1 \end{array} \tag{1}$$

Network Link Capacity Limits

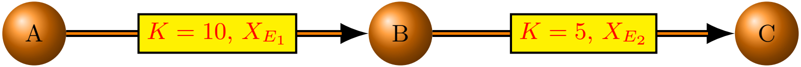

- What is the maximum flow through this simple network?
- I = incidence matrix; II^T = adjacency matrix
- Net input to all nodes 0– expect for source (net inflow f) and drain (net outflow f) otherwise, $Ix = 0$
- No link can exceed its capacity $x \leq K$

$$\begin{aligned} \max f \text{ s.t.} \\ Ix - fe = 0 \\ x \leq k \end{aligned} \tag{2}$$

Representing Network Flow

- Represent the network as an incidence matrix I (node $\times$ edges)
- Values are relative to each node: 0 (no connection to edge), 1 (edge enters the node), -1 (edge leaves the node)
- Let x be a variable representing the actual flow through each edge; then Ix (node $\times$ edge) (edge $\times$ 1) is the flow "stored" in each node
- Assume each node has no storage; whatever goes in must come out
- Then $I \times \bar{x} = \bar{0}$, where I is a matrix, $\bar{x}$ is a vector, and $\bar{0}$ is the zero vector
- Also assume that we know the maximum capacities on each link k; the we have everything we need to determine the maximum flow

Keeping Track of Flows in and Out of Each Node

- Net flow out of A is f : $-X_{E1} + f = 0$
- Net flow into B is zero: $X_{E1} - X_{E2} = 0$
- Net flow into C is f : $X_{E2} - f = 0$
- $\max f$ s.t. $Ix - fe = 0 \; x \leq k$

$$\begin{bmatrix} -1 & 0 \\ 1 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} X_{E1} \\ X_{E2} \end{bmatrix} - f \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \tag{3}$$

- $0 \leq X_{E1} \leq 10$
- $0 \leq X_{E2} \leq 5$
- Solve using any math package: Mathematica, Matlab, etc...

Conclusion

- We've begun to look at the basic elements of networks and communication
- We're starting to see how they are analyzed
- Next lecture will cover:
  - Switching
  - Layering
  - Socket programming
  - More analysis

For next class

- Purchase and register your i>Clicker
- Become familiar with ns2 – install and run it on your personal machine or use the version in the computer lab
  - Begin reading `http://www.isi.edu/nsnam/ns/tutorial/`
  - You will have a significant final project (30% of your grade) that will utilize this
- Find access to Maple/Matlab/Mathematica for future course homework