

My background

- up to last year — CAD for digital ICs
 - functional verification, logic synthesis, physical design
- recently — networking hardware
 - scientific issues

Teaching to learn more

- do not have all answers
- expect input from class

Review

Packet switched network — ensemble of hosts, links, and routers

- Kleinrock: MIT PhD early 1960s, UCLA early 1970s
 - efficient
 - reliability
 - simplicity (end-to-end argument)

EE382n — Interconnection Networks

Adnan Aziz
ECE Department
The University of Texas
adnan@ece.utexas.edu

Introduction

- Internetworking
 - review
 - significance
- Course outline
 - applications
- Administrivia

What do routers do?

A lot

- Compute routing tables
- Forward packets
- Switching
- Manage different classes of service
- Congestion control
- Security (encryption, Denial of Service detection)
- Fragmentation/reassembly
- SNMP, ICMP
- NAT, Load balancing, Firewalling

- Accounting, Peering
- Multicast
- Multiprotocol operation

Networking hierarchy

- Physical layer — lasers, fibers, demodulators
- Datalink layer — Ethernet, SONET
- Network layer — IP
- Transport layer — TCP, UDP

Internet architecture

- LANs (ACES building)
- enterprise (UT system)
- access (MCI connection)
- core (Sprint network)

Design considerations

- Design cost
 - time to design, number of designers
- Manufacturing cost
 - technology, amount of memory, ASIC/FPGA/ μ P, pin count
- Deployment cost
 - administrator time, flexibility, ease of upgrades, downtime

Market – I

	1997	1998	1999	2000	2001	2002
Revenues	149	192	258	330	397	456
Profits	13	16	21	22	27	32

Table 1: US telecoms

	1997	1998	1999	2000	2001	2002
Revenues	83	83	95	123	98	100
Profits	11	8	13	23	14	15

Table 2: Semiconductor industry

Current problems

Formerly — Internet routers general purpose computers connected via bus to Tx/Rx hardware

- link bandwidth bottleneck \Rightarrow conventional processor could implement entire router: check-summing, packet encapsulation, routing, billing, firewalling, security

Today — high-speed optical fiber technology, added functionality \Rightarrow routers are bottleneck

Solution

Many operations not performance-critical \Rightarrow can still be performed with general purpose computers

- Implement router as combination of workstation-class CPU and collection of specialized ICs
 - Processor for “control functions” — compute best routes, analyzing traffic statistics, etc.
 - Chips for line speed operations — packet encapsulation/decapsulation, longest prefix matching, scheduling the switching fabric and output queues, etc.

Hardware: ASICs, network processors, FPGAs

Course outline

Focus on algorithmically challenging problems faced by router designers

- contrast with designing to a specification, e.g., a SONET framer

Overview	1 week
Packet classification	3 weeks
Scheduling	3 weeks
Routing	2 weeks
Miscellaneous	2 weeks
Guest speakers	2 weeks
Student presentations	2 weeks

Miniproject 1 – Classification

How to efficiently compute policies to apply to packet?

Src IP	Dst IP	Prot	Src Port	Dst Port	COS
143.22.*.*	197.32.*.*	6	*	6000	0
143.22.*.*	197.32.*.*	6	*	80	1

...

For IP forwarding, can perform using BDDs

- generalization?
- trade-offs

	1997	1998	1999	2000	2001	2002
Revenues	6	8	12	19	22	17
Profits	1.4	1.9	2.6	3.9	1.6	1.1

Table 3: Cisco

Market – II

- Physical layer: JDS Uniphase, Corning, Sycamore
- Datalink: Analog devices, Broadcom, Lucent
- Routers:
 - Boxes: Cisco, Nortel, Lucent, Juniper, Sonus, Foundry
 - Chips: LSI logic, PMC Sierra, AMC, Lucent
- Carriers: AT&T, SBC, Worldcom, MCI, Qwest, Level 3, Global Crossing
- ISPs: AOL, ATT, MCI

Nontraditional applications: storage, switched backplanes

- Evaluation (subject to change)
 - one midterm (30%), 5 miniprojects (30%), project (40%)
- Get account on ECE machines (ENS 507)
- Add/Drop, Scholastic Dishonesty, Disabilities
 - will follow published guidelines

What will you get from this class?

Techniques for implementing high-performance routers

- fundamental problems and solutions
 - classification, scheduling, routing
 - graph algorithms, stability analysis

Miniproject 2 — Switch scheduling

Fact — doubly substochastic real matrix \mathcal{R} can always be written as $\sum_i \phi_i P_i$, where $0 < \phi \leq 1$, and P_i is a permutation matrix

- doubly substochastic — $0 \leq r_{ij} \leq 1$ and $\sum_i r_{ij} \leq 1, \sum_j r_{ij} \leq 1$

Why is this of any consequence?

- “best” decomposition of \mathcal{R} ?

Administrivia

- Handouts
 - Syllabus, Assignment 0, Lecture 1, Signup sheet
- Prerequisites
 - C/C++ and algorithms (CLRS), digital design (Mano), networks (Peterson)
 - “mathematical maturity”
- Course material
 - Keshav’s book, various papers, my notes