EECS 589:
Advanced Computer Networks

Z. Morley Mao

Fall 2014 (TuTh 10:30-12:00 in 1690 Beyster)

Office hours: TuTh 12:00-1:00PM

Email for appointment

http://www.eecs.umich.edu/courses/eecs589

The Internet: An Exciting Time

- One of the most influential inventions
  - A research experiment that escaped from the lab
  - … to be a global communications infrastructure
- Ever wider reach
  - Today: 1.7 billion users
  - Tomorrow: more users, computers, sensors
- Near-constant innovation
  - Apps: Web, P2P, social networks, virtual worlds
  - Links: optics, WiFi, cellular (4G LTE), WiMax, ...

2

Transforming Everything

- The ways we do business
  - E-commerce, advertising, cloud computing, ...
- The way we have relationships
  - E-mail, IM, Facebook friends, virtual worlds
- How we think about law
  - Interstate commerce? National boundaries?
- The way we govern
  - E-voting and e-government
  - Censorship and wiretapping
- The way we fight
  - Cyber-attacks, including nation-state attacks

3

But, What *is* Networking?

4

A Plethora of Protocol Acronyms?

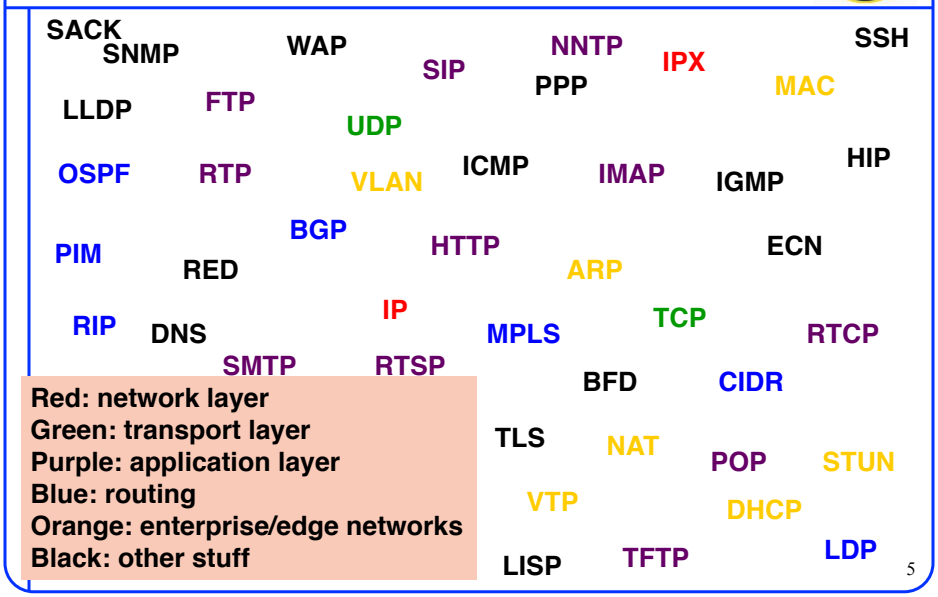

Red: network layer
Green: transport layer
Purple: application layer
Blue: routing
Orange: enterprise/edge networks
Black: other stuff

5

A Heap of Header Formats?

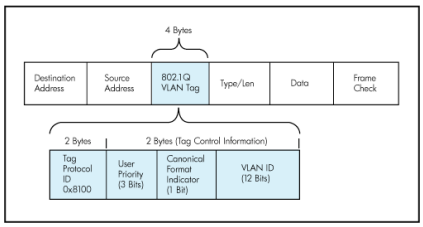

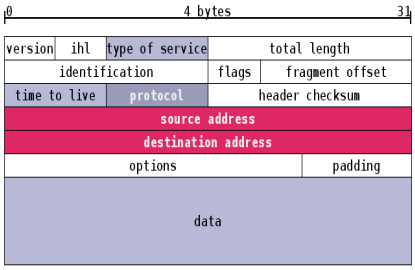

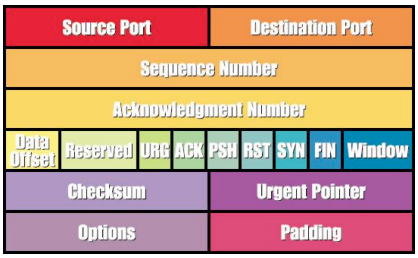

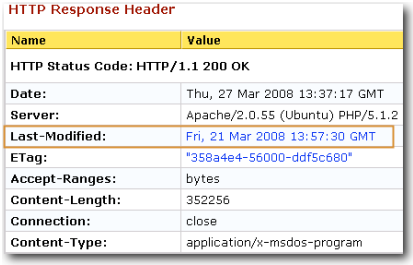

TCP/IP Header Formats in Lego

7

A Big Bunch of Boxes?

Router

Label Switched Router

Load balancer

Switch

Scrubber

Repeater

Gateway

Intrusion Detection System

Bridge

Route Reflector

Deep Packet Inspection

DHCP server

Packet shaper

Firewall

NAT

Hub

Packet sniffer

WAN accelerator

DNS server

Base station

Proxy

8

A Ton of Tools?

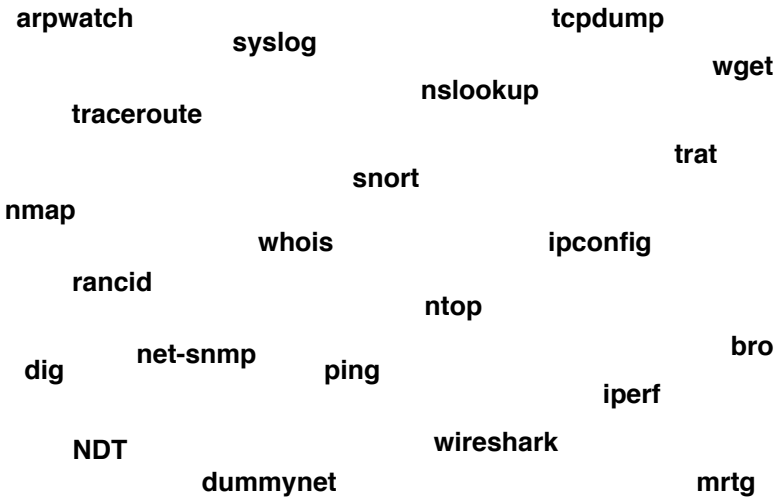

9

What Do Peers in Other Fields Say?

- "You networking people are very curious. You really love your artifacts."
- "In my college networking class I fell asleep at the start of the semester when the IP header was on the screen, and woke up at the end of the semester with the TCP header on the screen."
- "Networking is all details and no principles."

Is networking "just the (arti)facts"?

10

An Application Domain?

11

Application Domain for Theory?

- Algorithms and data structures
- Control theory
- Queuing theory
- Optimization theory
- Game theory and mechanism design
- Formal methods
- Information theory
- Cryptography
- Programming languages
- Graph theory

12

Application Domain for Systems?

- Distributed systems
  - DS assumes underlying connectivity of some sort
- Operating systems
  - End-host networking stack
- Computer architecture
  - Network interface card, router hardware, etc.
- Software engineering
  - Protocol design, understanding requirements, …
- …

13

An Exercise in Entrepreneurship?

- Identify a need or desirable capability
  - Whether previously known or not
- Invent a new feature or system that provides it
- Determine how it fits in the existing network
- Build and/or evaluate your solution
- Pitch or $ell the problem and solution to others
  - Whether to investors or a program committee
- Bask in glory, or lick your wounds

14

What Peers in Other Fields Say?

- "Networking papers are strange. They have a lot of text."
- "What are the top ten classic problems in networking? I would like to solve one of them and submit a paper to SIGCOMM." After hearing that we don't have such a list: "Then how do you consider networking a discipline?"
- "So, these networking research people today aren't doing theory, and yet they aren't the people who brought us the Internet. What exactly are they doing?"
- "Networking is an opportunistic discipline."

Is networking a problem domain or a scholarly discipline?

15

What Do We Teach Networking Students?

16

How Practitioners Learn Networking

- Certification courses
  - On how to configure specific pieces of equipment

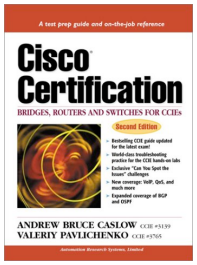

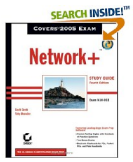

- "On the job" training
  - Aka "trial by fire"

17

How Colleges Teach Networking

- Undergraduates: how the Internet works
- Graduates: read the 20 "best" papers
- Few general principles, little "hands-on" experience

"There is a tendency in our field to believe that everything we currently use is a paragon of engineering, rather than a snapshot of our understanding at the time. We build great myths of spin about how what we have done is the only way to do it to the point that our universities now teach the flaws to students (and professors and textbook authors) who don't know better." **-- John Day**

- (I'm as guilty as anyone)

18

Now That I've Bummed You Out…

Or, Why Should You Stay in
This Class, and This Field?

19

So, Why is Networking Cool?

- Tangible, relates to reality
  - Can measure/build things (we *do* "love our artifacts")
  - Can truly effect far-reaching change in the real world
- Inherently interdisciplinary
  - Well-motivated problems + rigorous solution techniques
  - Interplay with policy, economics, and social science
- Widely-read papers
  - Many of the most cited papers in CS are in networking
  - Congestion control, distributed hash tables, resource reservation, self-similar traffic, multimedia protocols,…
  - Three of top-ten CS authors (Shenker, Jacobson, Floyd)
  - So, *somebody* is interested in reading this stuff… ☺

20

So, Why is Networking Cool? (Cont)

- Young, relatively immature field
  - Great if you like to make order out of chaos
  - Tremendous intellectual progress is still needed
  - *You* can help decide what networking really is
- Defining the problem is a big part of the challenge
  - Recognizing a need, formulating a well-defined problem
  - … is at least as important as solving the problem…
- Lots of platforms for building your ideas
  - Programmability: Click, OpenFlow/NOX, NetFPGA
  - Routing software: Quagga, XORP, and Bird
  - Testbeds: Emulab, PlanetLab, Orbit, GENI, …
  - Measurements: RouteViews, traceroute, Internet2, …

21

But, That Doesn't Say What Networking Really *Is*

Or, What Will This Course Be About?

22

One Take on Defining Networking

- How to
  - Design and operate *components and protocols*
  - That can be used and *combined in many ways*
  - To do *many things*
- Definition and placement of function
  - What to do, and where to do it
- The "division of labor"
  - Between the host, network, and management systems
  - Across multiple concurrent protocols and mechanisms
- But, how to judge a *good* division of labor?
  - We need some sort of "user" in mind

23

Different Users, Different Questions

- End user who runs applications?
  - Too far removed from architectural decisions
- Application developer who writes applications?
  - Somewhat removed, though good to keep in mind
  - E.g., for what API to have, for what properties to ensure
- Protocol designer who creates the protocols?
  - Already presupposes some parts of "dividing the labor"
- Network administrator who runs a network?
  - Responsible for tuning and composing protocols
  - … to achieve a wide variety of specific goals
  - (Internet was not designed with management in mind)

24

What Is This Course About?

- A survey of main topics and current research focus
- An important concept: network management
  - Exploring definition and placement of function
  - With an emphasis on how to manage the network
  - As a *concrete* way to explore architecture questions
- Past and future
  - How to manage *today's* protocols and systems better
  - How to design *future* networks that are easier to manage
- The interplay is important
  - Today's design is one choice for the "division of labor"
  - Exploring the challenges, and existing solutions will help us understand how we can do better

25

Network Management

- Configuration
- Traffic engineering
- Troubleshooting
- Security
- Minimizing disruptions
- Energy efficiency
- Supporting services

Smart distributed protocols with dumb management system vs. dumb network and smart management system?

Enterprise, data-center, home, and backbone networks. (Not much on wireless.)

26

Structure of the Course

- Reading and analyzing research papers
  - Different ways to approach networking problems
  - Summaries, critiques, and comparisons of the papers
- Classroom time
  - Brief overview of background material
  - Discussion and debate about the research papers
- Programming assignments
  - Learning the platforms for evaluating your ideas
  - Click/Emulab, OpenFlow/NOX, and measurement data
- Class project: research project
  - Novel research with a system-building component

27

Course Logistics

- Instructor
  - Z. Morley Mao (zmao@umich.edu, 4629 Beyster)
  - GSI: Sanae Rosen (sanae@umich.edu)
- Lecture time:
  - TuTh, 10:30-12:00 PM,
  - Discussion time: Friday 9:30-10:30
- Location: 1690 Beyster
- Office hour:
  - TuTh 12:00-1:00PM
    - Software reading group: Thu 12:30 in 4901 BBB
    - Security reading group: Tue 12:30 in 4901 BBB
  - email for appointment
  - GSI: Fri: 10:30-11:30AM
- http://www.eecs.umich.edu/courses/eecs589

28

Topics Covered

- Internet design philosophies and principles
- Internet routing, Internet AS relationships
- Software-defined networking (SDN)
- ISP traffic engineering
- Critical network infrastructure services
- Network security: IDS, worms, and honeypots
- CDNs, Peer to peer and overlay networks, DHTs
- Wireless networking, meshed networks, mobile systems
- Sensor networking
- Network measurements
- Network security
- Network models
- Network algorithmics: building fast networks
- Networking in challenging environments
- Data center networking

Interesting new/hot topics

- A research agenda for Web 2.0
- Architectural support for security or availability
- Computing in the cloud: what role for networking research?
- Ensuring correctness for distributed protocols
- Evolution of storage area networks
- Lessons drawn from failed research, and controversial or disruptive topics
- Measurement and management of metro-area WiFi networks
- Network coding: hype or reality?
- Power as a first-class design property; "green" protocols/implementations
- Protocol design for optical switching
- The future role of network processors
- Third-world networking challenges
- Understanding the economics of operational costs
- Unique challenges of massive multi-player game systems
- Validation of measurement-based research: what are our standards?

Administrative Trivia

- Course Web page:
  - http://www.eecs.umich.edu/courses/eecs589
  - Linked from my home page: http://www.eecs.umich.edu/~zmao
  - Check it periodically to get the latest information
- Please come prepared for each class
  - Reading summaries are due before each class
  - Attendance is important
- Assignments are done individually, unless otherwise noted
- Research project are encouraged to be done in groups (at most 3 people)

Goals of this Course

- Critical examination of current topics of computer networks
  - What assumptions are no longer valid
  - What are the new research problems to look at
- Understand solutions in context
  - Goals
  - Assumptions
- Learning how to do research in systems
  - Paper review, writing, and presentation
- Appreciate what is good research
  - Problem selection
  - Solution & research methodology
  - Presentation
- Apply what you learned in a class project

What Do You Need To Do?

- A research-oriented class project
- Paper readings, selective book chapters
- Lead at least one class discussion
- 2-3 assignments

Research Project

- Investigate new ideas and solutions in a class research project
  - Define the problem
  - Execute the research
  - Work with your partner(s)
  - Write up and present your research
- Ideally, best projects will become conference papers (e.g., SIGCOMM, INFOCOM, MOBICOM, Sensys, MobiSys, IMC)

Research Project: Steps

- I'll distribute a list of projects
  - You can either choose one of these projects or come up with your own
- Pick your project, partner, and submit a one page proposal describing:
  - The problem you are solving
  - Your plan of attack with milestones and dates
  - Any special resources you may need
- A midterm presentation of your progress (>10 minutes)
- Final project presentation (>30 minutes)
- Submit project papers

Paper Reviews

- Goal: synthesize main ideas and concepts in the papers
- Number: about two papers per class
- Length: no more than half page per paper
- Content
  - Main points intended by the author
  - Points you particularly liked/disliked
  - Other comments (writing, conclusions…)
- Submission:
  - Submit each review on lecture day in class
  - See class web page for details

Grading

Term project	50%
Assignments	20%
Paper presentation	10%
Reading summaries	10%
Class discussion	10%

- This is a graduate networking class: more important is what you realize/learn than the grade

Self Introduction

- Faculty member in the software lab
- Past and ongoing research:
  - Next-generation Internet design
    - Security, mobility, and management challenges
  - Mobile networks/systems
  - Intrusion detection systems
  - Network troubleshooting, debugging
  - Network configuration management, provisioning
  - Denial of service attacks
  - General interest: Internet networks, mobile systems, network security.

iMAP – Intelligent Mobile App Profiling Tool

- Smartphone app developers are not aware of **resource control** mechanism in cellular network, leading to

  resource-inefficient applications

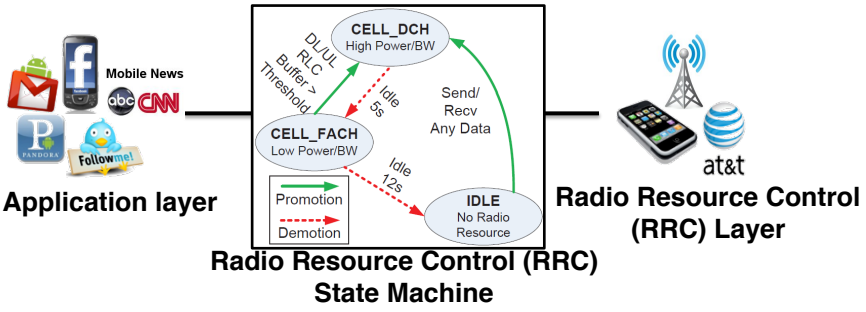

- Our solution: iMAP performs **cross-layer analysis** to pinpoint resource inefficiencies for smartphone applications
  - **Providing visibility** of radio resource / radio energy utilization
  - **Benchmarking** resource efficiencies
  - **Diagnosing** inefficiencies

iMAP: Solution, Findings, and Recommendation

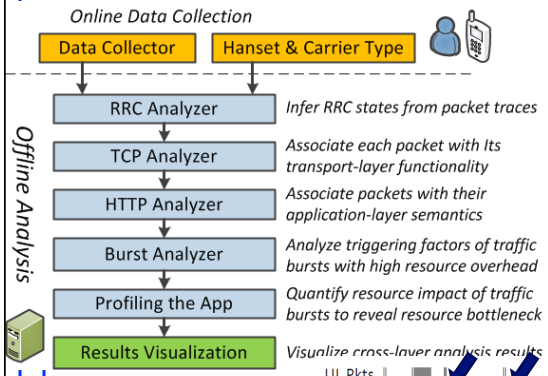

Example: Profiling results for Pandora Streaming

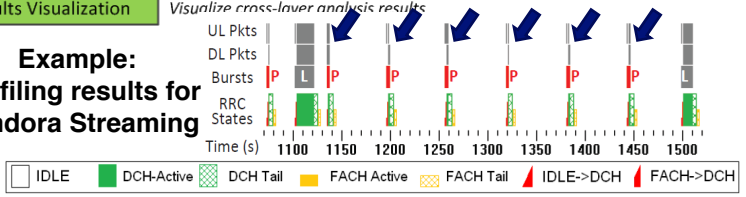

Examples of findings of iMAP for popular Android applications

Pandora Streaming
High radio energy overhead (50%) of periodical measurements

Fox News
High radio energy overhead (15%) due to users' scrolling

Google Search
High radio energy overhead (78%) due to real-time query suggestions

Problem: High resource overhead of periodic audience measurements (every 1 min)
Recommendation: Delay transfers and batch them with delay-sensitive transfers

Student introduction

- Please introduce yourself: name, standing, research area (for grad students)
- Say a few words about what you think you would like to learn about computer networks
- And what you think are "unsolved" problems in computer networks

Getting Started…

"Division of Labor" in Today's Internet:

IP as the Host/Network Interface

42

Host-Network Division of Labor

- Packet switching
  - Divide messages into a sequence of packets
  - Headers with source and destination address
- Best-effort delivery
  - Packets may be lost
  - Packets may be corrupted
  - Packets may be delivered out of order

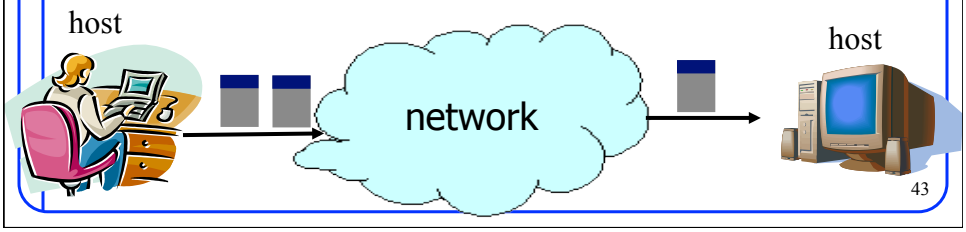

43

Host-Network Interface: Why Packets?

- Data traffic is bursty
  - Logging in to remote machines
  - Exchanging e-mail messages

- Don't want to waste bandwidth
  - No traffic exchanged during idle periods
- Better to allow multiplexing
  - Different transfers share access to same links
- Packets can be delivered by most anything
  - RFC 1149: IP Datagrams over Avian Carriers

44

Host-Network Interface: Why Best-Effort?

- Never having to say you're sorry…
  - Don't reserve bandwidth and memory
  - Don't do error detection & correction
  - Don't remember from one packet to next
- Easier to survive failures
  - Transient disruptions are okay during failover
- Can run on nearly any link technology
  - Greater interoperability and evolution

45

Intermediate Transport Layer

- But, *applications* want efficient, accurate transfer of data in order, in a timely fashion
  - Let the end hosts handle all of that
  - (An example of the "end-to-end argument")
- Transport layer can optionally…
  - Retransmit lost packets
  - Put packets back in order
  - Detect and handle corrupted packets
  - Avoid overloading the receiver
  - <insert your requirement here>

46

IP Suite: End Hosts vs. Routers

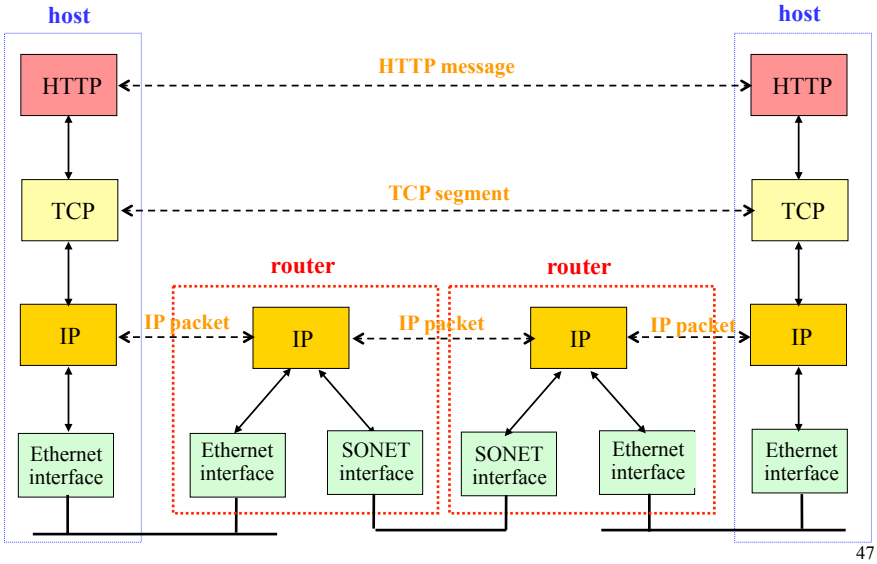

47

The "Narrow Waist" of IP

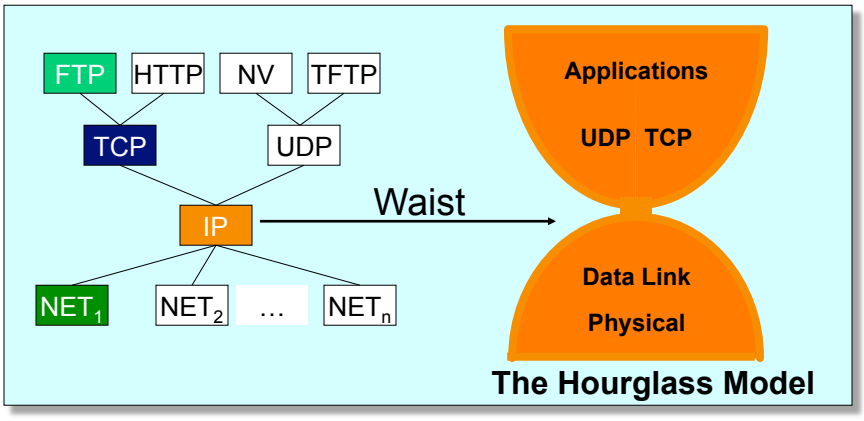

The waist facilitates interoperability

48

Layer Encapsulation

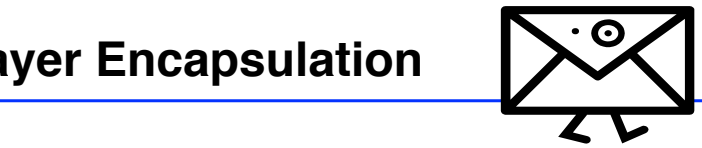

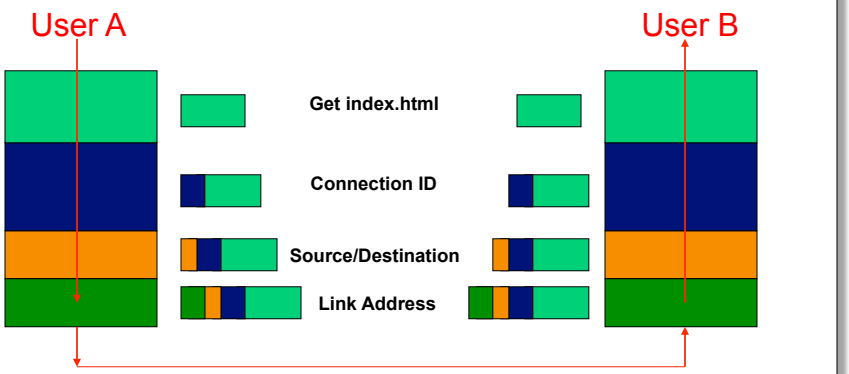

49

But What About the *Inside* of the Network

"Division of Labor" Between Network Elements and the Management System

50

Inside the Network

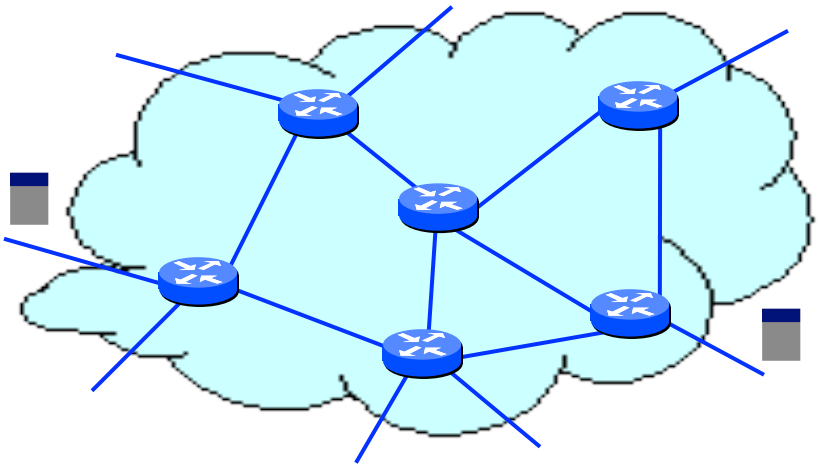

Forward packets from the sender to the receiver

51

Split into Data vs. Control Plane

- Data plane: packets
  - Handle individual packets as they arrive
  - Forward, drop, or buffer
  - Mark, shape, schedule, …
- Control plane: events
  - Track changes in network topology
  - Compute paths through the network
  - Reserve resources along a path

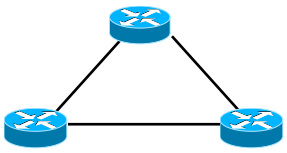

Motivated by need for high-speed packet forwarding

52

Adding the Management Plane

- Making the network run *well*
  - Traffic reaches the right destination
  - Traffic flows over short, uncongested paths
  - Unwanted traffic is discarded
  - Failure recovery happens quickly
  - Routers don't run out of resources
- A control loop with the network
  - Measure (sense): topology, traffic, performance, …
  - Control (actuate): configure control and data planes

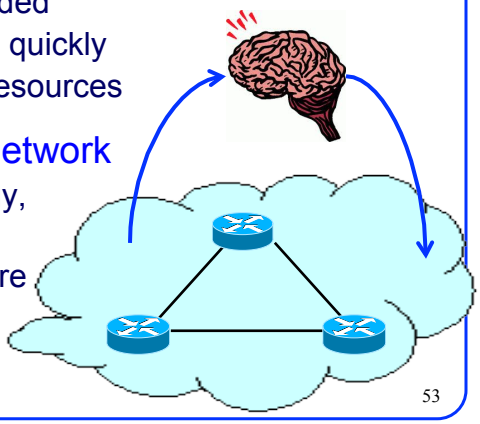

53

Next Three Classes: Review

- Host
  - Network discovery and bootstrapping
  - Resource allocation and interface to applications
- Control plane
  - Distributed algorithms for computing paths
  - Disseminating the addresses of end hosts
- Data plane
  - Streaming algorithms and switch fabric
  - Forward, filter, buffer, schedule, mark, monitor, …
- In addition, we will cover some research papers.

54

How to Read

You May Think You Already Know
How To Read, But…

55

You Spend a Lot of Time Reading

- Reading papers for grad classes (like this one!)
- Reviewing papers for conferences/journals
- Giving colleagues feedback on their papers
- Keeping up with work related to your research
- Staying broadly educated about the field
- Transitioning into a new research area
- Learning how to write better papers ☺

So, it is worthwhile to learn to read *effectively*

56

Keshav's Three-Pass Approach: Step 1

- A ten-minute scan to get the general idea
  - Title, abstract, and introduction
  - Section and subsection titles
  - Conclusion
  - Bibliography
- What to learn: the five C's
  - Category: What type of paper is it?
  - Context: What body of work does it relate to?
  - Correctness: Do the assumptions seem valid?
  - Contributions: What are the main research contributions?
  - Clarity: Is the paper well-written?
- Decide whether to read further…

57

Keshav's Three-Pass Approach: Step 2

- A more careful, one-hour reading
  - Read with greater care, but ignore details like proofs
  - Figures, diagrams, and illustrations
  - Mark relevant references for later reading
- Grasp the content of the paper
  - Be able to summarize the main thrust to others
  - Identify whether you can (or should) fully understand
- Decide whether to
  - Abandon reading the paper in any greater depth
  - Read background material before proceeding further
  - Persevere and continue on to the third pass

58

Keshav's Three-Pass Approach: Step 3

- Several-hour virtual re-implementation of the work
  - Making the same assumptions, recreate the work
  - Identify the paper's innovations and its failings
  - Identify and challenge every assumption
  - Think how you would present the ideas yourself
  - Jot down ideas for future work
- When should you read this carefully?
  - Reviewing for a conference or journal
  - Giving colleagues feedback on a paper
  - Understanding a paper closely related to your research
  - Deeply understanding a classic paper in the field

http://ccr.sigcomm.org/online/?q=node/234

59

Other Tips for Reading Papers

- Read at the right level for what you need
  - "Work smarter, not harder"
- Read at the right time of day
  - When you are fresh, not sleepy
- Read in the right place
  - Where you are not distracted, and have enough time
- Read actively
  - With a purpose (what is your goal?)
  - With a pen or computer to take notes
- Read critically
  - Think, question, challenge, critique, …

60

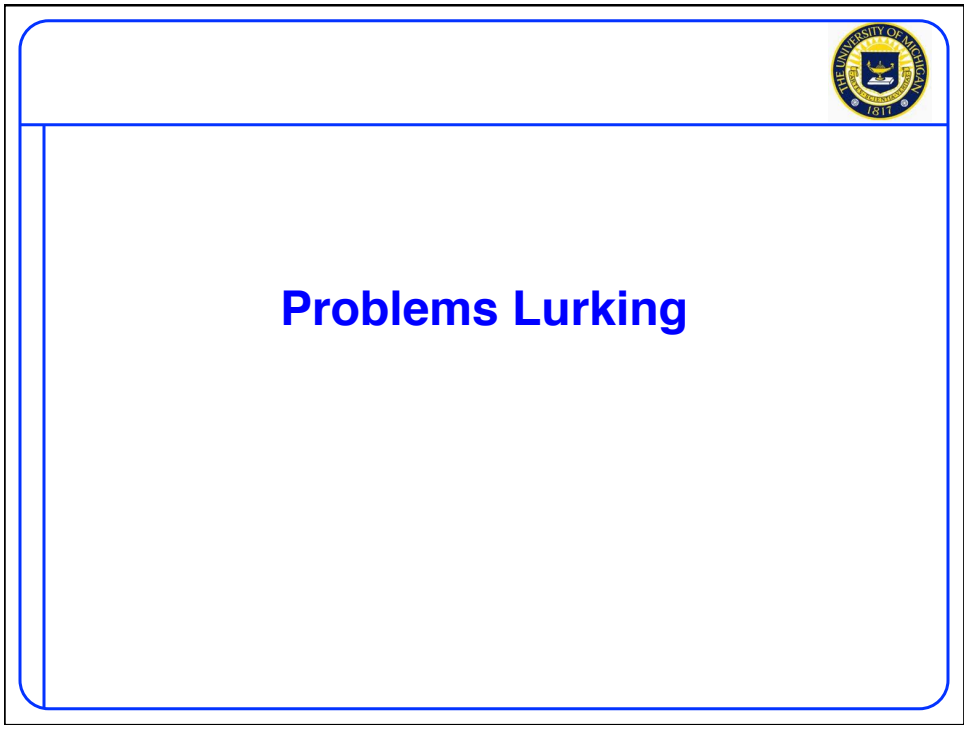

Challenges Tied to Early Design Decisions

- Power of programmable end hosts
  - Easy to spoof IP addresses, e-mail addresses, …
  - Incentives for users to violate congestion control
  - Malicious users launching Denial-of-Service attacks
- Best-effort packet-delivery service
  - Inefficient in high-loss environments (wireless)
  - Poor performance for interactive applications
  - Expensive per-packet handling on high-speed links

Challenges Tied to Early Design Decisions

- Layering and the IP narrow waist
  - Low efficiency due to many layers of headers
  - Poor visibility into underlying shared risks
  - Complex network management due to multiple interconnected protocols and systems
- Decentralized control
  - Hierarchical addressing makes mobility difficult, and requires careful configuration
  - Autonomy makes measurement (and troubleshooting and accountability) hard
  - Autonomy makes protocol changes difficult

Recurring Challenges

- Security
  - Weak notions of identity that are easy to spoof
  - Protocols that rely on good behavior
  - Incomplete or non-existent registries, keys, …
- Mobility and disconnected operation
  - Hierarchical addressing closely tied with routing
  - Presumption that hosts are connected
- Network management
  - Many coupled, decentralized control loops
  - Limited visibility into across layers and networks
- Application performance requirements
  - Real-time, interactive applications
  - Throughput sensitive vs. delay-sensitive

Internet is Not Standing Still

- Partial solutions to these problems
  - Often as “add ons” or “extensions”
  - Hampered by need to be backwards compatible, and work when only partially deployed
  - Rather than complete architectural solutions
- Solutions create problems of their own
  - Violations of architectural assumptions
  - Unexpected interactions with applications
  - Adding complexity to an already complex system

Example: Middleboxes

- Middleboxes are intermediaries
  - Interposed in-between the communicating hosts
  - Often without knowledge of one or both parties
- Examples
  - Network address translators
  - Firewalls
  - Traffic shapers
  - Intrusion detection systems
  - Transparent Web proxy caches
  - Application accelerators

Middleboxes Address Practical Challenges

- Host mobility
  - Relaying traffic to a host in motion
- IP address depletion
  - Allowing multiple hosts to share a single address
- Security concerns
  - Discarding suspicious or unwanted packets
  - Detecting suspicious traffic
- Performance concerns
  - Controlling how link bandwidth is allocated
  - Storing popular content near the clients

Middleboxes Violate Network-Layer Principles

- Globally unique identifiers
  - Each node has a unique, fixed IP address
  - … reachable from everyone and everywhere
- Simple packet forwarding
  - Network nodes simply forward packets
  - … rather than modifying or filtering them

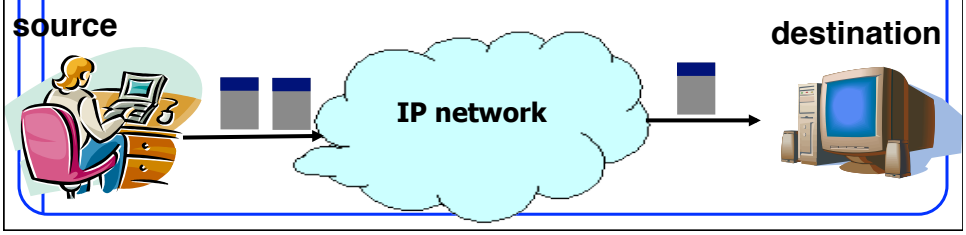

Two Views of Middleboxes

- An abomination
  - Violation of layering
  - Cause confusion in reasoning about the network
  - Responsible for many subtle bugs
- A practical necessity
  - Solving real and pressing problems
  - Needs that are not likely to go away
- Would they arise in any edge-empowered network, even if redesigned from scratch?

Clean-Slate Network Architecture

- Clean-slate architecture
  - Without constraints of today's artifacts
  - To have a stronger intellectual foundation
  - And move beyond the incremental fixes

- Still, some constraints inevitably remain
  - Ignore today's artifacts, but not necessarily all reality
- Such as…
  - Resource limitations (CPU, memory, bandwidth)
  - Time delays between nodes
  - Independent economic entities
  - Malicious parties
  - The need to evolve over time

Conclusions

- Internet architecture is a huge success
  - Functionality at programmable edge nodes
  - Best-effort packet-delivery service
  - Layering and the IP hourglass model
  - Decentralized control of the global system
- These very features are causing problems
  - Security, mobility, manageability, performance, reliability, …
- Rethinking the network architecture
  - For a strong intellectual foundation
  - And long-term improvements to the Internet

See you on Thursday!

72