

# **Course Overview and Internet Architecture**

CS 7260

Nick Feamster

January 8, 2007

# Who Am I?

- Nick Feamster
  - Assistant Professor in CoC
    - Ph.D from MIT, Sept. 2005
    - Thesis: Internet Routing Correctness, Predictability
  - How to reach me
    - feamster at cc.gatech.edu  
( please include “CS7260” in subject line)
    - KACB 3348
  - Office hours: Before class on Mondays  
(and by appointment)

# Primary Goal of This Course

**Provide a survey of the necessary tools, techniques, and concepts to perform research in computer communications**

- This course is **project-based** (not just paper reading)
  - Emphasis on hands-on experience
  - *Realization is key!*
- More in-depth coverage of networking topics
  - Focus on network-layer and above
- Crash-course in available tools for research
  - You may use one or more of these in the project

**Check the course Web page frequently!**

[http://www.cc.gatech.edu/classes/AY2007/cs7260\\_spring/](http://www.cc.gatech.edu/classes/AY2007/cs7260_spring/)

# What This Course is About

- **Lecture:** Learning about cutting-edge research problems in computer networking, and coming up with your own
  - We'll pick up basics along the way as necessary
  - The course topics are (1) breadth-first and (2) not comprehensive
- **Problem Sets:** Developing proficiency with tools and techniques for following through on your research ideas
  - Tons of exciting tools out there!
  - Problem sets will help you with this goal
  - Thinking about network design

**These two components should help you develop great projects.**

# Toolkit for the Networking Researcher

- Exposure to various “hammers”
  - Networking is a domain that draws on many disciplines
- Measurement and deployment experience
  - Realism is key
- Development of design skills
  - How and why the Internet works the way it does
  - Experience thinking about design alternatives

# What This Course is *Not* About

- An introduction to networking
  - Examples of topics that won't be covered:
    - TCP basics
    - Socket programming basics
    - etc.
- An introduction to programming
  - Knowledge of scripting languages will help.
    - If *some* programming language, don't worry if you don't know a scripting language. There's time to learn, since deadlines are spread out.

**Follow the “spirit” of the pre-requisites.**

# How: Course Structure (and Grading)

- One semester-long **project** (50%)
- Two in-class “**quizzes**” (30% total)
  - **February 22 and March 17**
  - Should be relatively easy if you’ve been coming to class and keeping up with readings
- Three **problem sets** (20% total)
  - PS1 will be assigned on 1/17 (Wednesday)
    - A handful of “paper and pencil” problems
    - “Analysis” question, which will require scripting
  - PS2: Experimentation
    - More programming required. Experimentation with traces and scripts
  - PS3: Design
    - Experimentation on Click, Planetlab, etc.

# Project Expectations

- Aim high!
  - A good project can become the basis for:
    - Publication
      - Internet Measurement Conference deadline mid-May, Infocom mid-summer...
    - Ph.D Thesis
- Your project need not be *SIGCOMM*-quality by the end-of-term, but it should be something that could be conference-worthy with a bit more effort
- I am here to help you
- New project ideas posted in a few weeks



# Project Logistics: Five Milestones

- January 22: Project Groups
  - 3 person groups. 2-person groups by rare exception
- February 7: Project Proposal
  - 1-2 page writeup
  - Problem statement, evaluation strategy, metrics for “success”
- March 26: Interim Report and Mini-presentations
- April 23 & 25: Project Presentations (In Class)
- April 27: Writeups Due
  - 8-10 pages. Research paper-style.

Meeting deadlines early is encouraged! I'm happy to look at your progress before these dates.

# Lecture Structure

- ~ 55 minutes lecture, ~ 25 minutes discussion
- ***Read the required paper before class***
- My plan: Thought questions posted at least the day before class.
  - Hopefully will help stimulate discussion.
- I will try to post optional readings, in case you are interested in reading more about some topic.
  - If I don't do so for a topic you're interested in, ask me!

# Differences from Last Year

- New and different themes and topics
  - Multipath routing
  - Network virtualization
  - Strategies for reducing unwanted traffic
- New papers
- Focus on “tool sharpening”
  - Optimization
  - Game theory
  - Network coding
  - Machine learning

# Topic Highlights

- Essentials
  - *Naming and Addressing*: DNS, IPv6, NAT, Flat Names
  - *Routing*: BGP, MPLS, VPNs, etc.
  - Multihoming and reliability
- Measurement and Operations
  - Testbeds: Emulab, PlanetLab, VINI
  - Techniques and tools
  - Network monitoring
  - Troubleshooting

# Topic Highlights

- Abstractions (“Networks on networks”)
  - Overlay routing
  - Network virtualization: techniques and applications
- Security
  - Worms, spam, botnets, etc.
  - Routing security
  - Anomaly detection
- Wireless and “Challenged” Networks
  - Networking in developing regions

# For the Rest of Today

- We will review today's reading and put it in the context of some of the topics we'll be covering through this term.

# Today's Reading

- *Design Philosophy of the DARPA Internet Protocols.* Dave Clark, 1988.
- Conceptual Lessons
  - Design principles/priorities were designed for a certain type of network. As the Internet evolves, we feel the sting of some of these choices.  
*Examples:* Commercialization
  - Engineering/Realization is key to testing an idea.
- Technical Lessons
  - Packet switching
  - Fate Sharing/Soft state

# Fundamental Goal

- “technique for **multiplexed utilization of existing interconnected networks**”
- **Multiplexing** (sharing)
  - Shared use of a single communications channel
- **Existing networks** (interconnection)



# Fundamental Goal: Sharing

## Packet Switching

- No connection setup
- Forwarding based on destination address in packet
- Efficient sharing of resources

***Tradeoff:* Resource management potentially more difficult.**

# Type of Packet Switching: Datagrams

- Information for forwarding traffic is contained in destination address of packet
- No state established ahead of time (helps fate sharing)
- Basic building block
- Minimal assumption about network service

## Alternatives (More on Wednesday)

- **Circuit Switching:** Signaling protocol sets up entire path out-of-band. (cf. the phone network)
- **Virtual Circuits:** Hybrid approach. Packets carry “tags” to indicate path, forwarding over IP
- **Source routing:** Complete route is contained in each data packet

# An Age-Old Debate

## Circuit Switching

- Resource control, accounting, ability to “pin” paths, etc.

## Packet Switching

- Sharing of resources, soft state (good resilience properties), etc.

It is held that packet switching was one of the Internet's greatest design choices.

Of course, there are constant attempts to shoehorn the best aspects of circuits into packet switching.

*Examples: Capabilities (Lecture 21), MPLS (Lecture 15), ATM, IntServ QoS, etc.*

# Stopping Unwanted Traffic is Hard

February 2000


The screenshot shows the BBC News website from February 9, 2000. The main headline is "Yahoo attack exposes web weakness". The article text reads: "Worst outage in Yahoo's history" and "By BBC News Online's Alfred Hermida". A yellow sticky note on the left side of the page contains the text: "The BBC's Rory Cellan Jones 'It was absolutely unprecedented' real 28k" and "The BBC's Alfred Hermida 'A flood of fake e-mails' real 28k". The URL at the bottom is <http://news.bbc.co.uk/1/hi/english/sci/tech/default.stm>.

March 2006

The screenshot shows the CNET News.com website from March 24, 2006. The main headline is "DNS servers do hackers' dirty work". The article is by Joris Evers, Staff Writer at CNET News.com, published on March 24, 2006, at 4:00 AM PST. The article text reads: "In a twist on distributed denial-of-service attacks, cybercriminals are using DNS servers--the phonebooks of the Internet--to amplify their assaults and disrupt online business." and "Earlier this year, VeriSign experienced attacks on its systems that were larger than anything it had ever seen before, it said last week. The Mountain View, Calif.-based company, which helps companies do business on the Web, discovered that the assaults weren't coming from commandeered 'bot' computers, as is common. Instead, its machines were under attack by DNS (domain name system) servers." The URL at the bottom is "Done".

# Research: Stopping Unwanted Traffic

- *Datagram networks*: easy for anyone to send traffic to anyone else...even if they don't want it!



## Possible Defenses

- *Monitoring + Filtering*: Detect DoS attack and install filters to drop traffic.
- *Capabilities*: Only accept traffic that carries a “capability”

**Stay tuned. More detail in Lecture 21.**

# The Design Goals of Internet, v1

- **Interconnection/Multiplexing** (*packet switching*)
- **Resilience/Survivability** (*fate sharing*)
- Heterogeneity
  - Different types of services
  - Different types of networks
- Distributed management
- Cost effectiveness
- Ease of attachment
- Accountability

Decreasing  
Priority

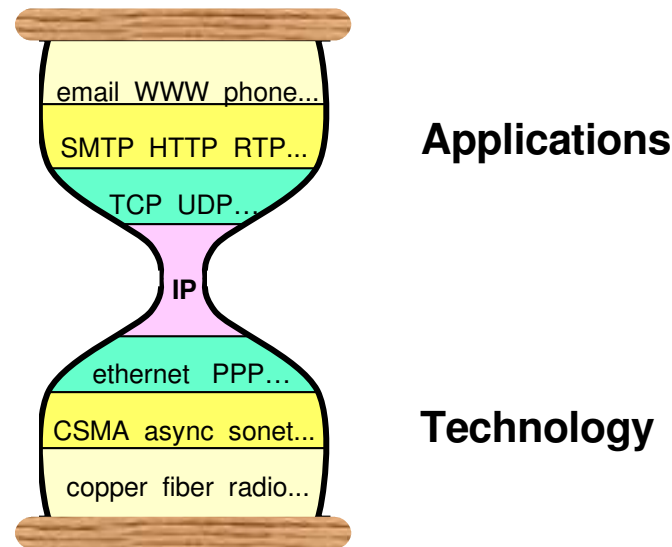


“This set of goals might seem to be nothing more than a checklist of all the desirable network features. It is important to understand that these goals are in order of importance, and **an entirely different network architecture would result if the order were changed.**”

**These goals were prioritized for a military network.  
Should priorities change as the network evolves?**

# Fundamental Goal: Interconnection

- Need to interconnect many existing networks
- Hide underlying technology from applications
- Decisions:
  - Network provides minimal functionality
  - ***“Narrow waist”***



***Tradeoff:*** No assumptions, no guarantees.

# The “Curse of the Narrow Waist”

- IP over anything, anything over IP
  - Has allowed for much innovation both above and below the IP layer of the stack
  - An IP stack gets a device on the Internet
- **Drawback:** very difficult to make changes to IP
  - But...people are trying
  - NSF GENI project: <http://www.geni.net/>



# Interconnection: “Gateways”

- Interconnect heterogeneous networks
- No state about ongoing connections
  - Stateless packet switches
- Generally, *router* == gateway
- But, we can think of your home router/NAT as also performing the function of a gateway




# Network Address Translation

- For *outbound* traffic, the gateway:
  - Creates a table entry for computer's local IP address and port number
  - Replaces the sending computer's non-routable IP address with the gateway IP address.
  - replaces the sending computer's source port
- For *inbound* traffic, the gateway:
  - checks the destination port on the packet
  - rewrites the destination address and destination port those in the table and forwards traffic to local machine

# NAT Traversal

- **Problem:** Machines behind NAT not globally addressable or routable. Can't initiate inbound connections.
- **One solution:** Signalling and Tunneling through UDP-Enabled NAT Devices (STUN)
  - STUN client contacts STUN server
  - STUN server tells client which IP/Port the NAT mapped it to
  - STUN client uses that IP/Port for call establishment/incoming messages



*More on Wednesday.*

# Goal #2: Survivability

- Network should continue to work, even if some devices fail, are compromised, etc.
- Failures on the Abilene (Internet 2) backbone network over the course of 6 months

	instability	unavailability	maintenance	total
node	0	2	22	24
link	0	20	65	85
peer	14	82	77	173
total	14	104	164	282

Thanks to Yiyi Huang

***How well does the current Internet support survivability?***

# Goal #2: Survivability

## Two Options

- Replication
  - Keep state at multiple places in the network, recover when nodes crash
- **Fate-sharing**
  - Acceptable to lose state information for some entity if the entity itself is lost

## Reasons for Fate Sharing

- Can support arbitrarily complex failure scenarios
- Engineering is easier

*Some reversals of this trend:*

NAT (Wednesday), Routing Control Platform (Lecture 4)

# Goal #3: Heterogeneous Services

- TCP/IP designed as a monolithic transport
  - TCP for flow control, reliable delivery
  - IP for forwarding
- Became clear that not every type of application would need reliable, in-order delivery
  - *Example:* Voice and video over networks
  - *Example:* DNS
  - Why don't these applications require reliable, in-order delivery?
  - Narrow waist: allowed proliferation of transport protocols

# Topic: Voice and Video over Networks

- **Deadlines:** Timeliness more important than 100% reliability.
- **Propagation of errors:** Some losses more devastating than others

Loss in “Anchor” Frame (I-Frame)



Propagates to “Dependent” Frames (P and B-Frames)



*More in Lecture 16.*

# Goal #3b: Heterogeneous Networks


- Build minimal functionality into the network
  - No need to re-engineering for each type of network
- “Best effort” service model.
  - Lost packets
  - Out-of-order packets
  - No quality guarantees
  - *No information about failures, performance, etc.*

***Tradeoff:*** Network management more difficult



# Research: Network Anomaly Detection

- Operators want to detect when a traffic flow from ingress to egress generates a “spike”.
- **Problem:** Today’s protocols don’t readily expose this information.
- Management/debuggability not initially a high priority!



***More in Lecture 12.***

# Goal #4: Distributed Management

## Many examples:

- Addressing (ARIN, RIPE, APNIC, etc.)
  - Though this was recently threatened.
- Naming (DNS)
- Routing (BGP)

**No single entity in charge.**

**Allows for organic growth, scalable management.**

***Tradeoff:* No one party has visibility/control.**


# No Owner, No Responsible Party

**“Some of the most significant problems with the Internet today relate to lack of sufficient tools for distributed management, especially in the area of routing.”**

- **Hard to figure out who/what’s causing a problem**
- **Worse yet, local actions have global effects...**

# Local Actions, Global Consequences

“...a glitch at a small ISP... triggered a **major outage in Internet access** across the country. The problem started when MAI Network Services...passed **bad router information** from one of its customers onto Sprint.” --  
*news.com*, April 25, 1997



# Goal #5: Cost Effectiveness

- Packet headers introduce high overhead
- End-to-end retransmission of lost packets
  - Potentially wasteful of bandwidth by placing burden on the edges of the network

**Arguably a good tradeoff. Current trends are to exploit redundancy even more.**

# Goal #6: Ease of Attachment

- IP is “plug and play” Anything with a working IP stack can connect to the Internet (hourglass model)
- A huge success!
  - **Lesson:** Lower the barrier to innovation/entry and people will get creative (e.g., Cerf and Kahn probably did not think about IP stacks on phones, sensors, etc.)
- But....

***Tradeoff:*** Burden on end systems/programmers.

# Goal #7: Accountability

- **Note:** Accountability mentioned in early papers on TCP/IP, but not prioritized
- Datagram networks make accounting tricky.
  - The phone network has had an easier time figuring out billing
  - Payments/billing on the Internet is much less precise
  - (More on this in Lectures 4 and 10)

***Tradeoff:*** Broken payment models and incentives.

# What's Missing?

- Security
- Availability
- Accountability (the other kind)
- Support for disconnected/intermittent operation
- Mobility
- Scaling
- ...



# Today's Reading

- *Design Philosophy of the DARPA Internet Protocols*. Dave Clark, 1988.
- Conceptual Lessons
  - Design principles/priorities were designed for a certain type of network. As the Internet evolves, we feel the sting of some of these choices.  
*Examples:* Commercialization,
  - **Engineering/Realization is key to testing an idea.**
- Technical Lessons
  - Packet switching
  - Fate Sharing/Soft state

# Design Goal Shakeup

- Cost of bandwidth is dropping. IP networks are becoming a commodity.
- Management == Human intervention
  - Costly!!
  - Human error a leading cause of downtime
- More bandwidth: are 40-byte headers still “big”?

# Today's Reading

- *Design Philosophy of the DARPA Internet Protocols.* Dave Clark, 1988.
- Conceptual Lessons
  - Design principles/priorities were designed for a certain type of network. As the Internet evolves, we feel the sting of some of these choices.  
*Examples:* Commercialization,
  - **Engineering/Realization is key to testing an idea.**
- Technical Lessons
  - Packet switching
  - Fate Sharing/Soft state

# Clark's Paper and This Course

- Flexible architectures (Good Thing) leave a lot of "wiggle room".
- **To determine whether something's going to work, it needs to be implemented/engineered.**

# So You've Got an Idea...

- This course will help you figure out how to test it out, measure it, etc..
- Test environments
  - Emulab
  - Planetlab
  - VINI: Virtual Network Infrastructure
- Data Sources
  - Datapository
  - Routeviews
  - Abilene Observatory
- Networking Software
  - Click Modular Router (for forwarding in user-level)
  - XORP Software Router (for routing)

***Details available on course Web site  
Come talk to me.***