CSE 461: Introduction

Arvind Krishnamurthy

Outline

- Administrative trivia
- Goals of the course
- How to study networks?

- Instructor: Arvind Krishnamurthy
  - Email: arvind@cs.washington.edu

- Teaching assistants:
  - John John: runs the sessions, manages projects
  - Alper Sarikaya: in charge of the homeworks
  - Email: jjohn@cs.washington.edu, alpers@cs.washington.edu
  - Office hours: TBA

Administrative Details

- Everything you need is on the course web page
  - http://www.cs.washington.edu/cse461
- Your TODO list:
  - Join the mailing list if not already on it: cse461@cs.washington.edu
  - Gain access to the CSE Labs (form for non-majors)
  - Get Computer Networks by Peterson and Davie
  - Read chapters 1 and 2
  - Go to section
  - Start on project 1
  - Obtain an N800 tablet from the TA (John John in CSE 391) tomorrow

A Network in 461

- "Network" is clearly an overloaded word:
  - Economic networks, regulatory networks, social networks…
  - Telephone, Cable TV, computer clusters
- For 461, a network is what you get anytime you connect two or more computers together by some kind of a link.

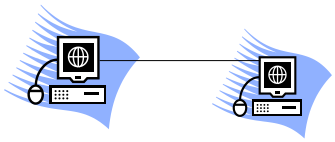

OR

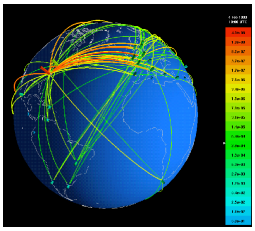

The networks we study

- We are interested in networks that are:
  - Large scale
  - Intrinisically Unreliable
  - Distributed
  - Heterogeneous

The meaning of "Large-scale"

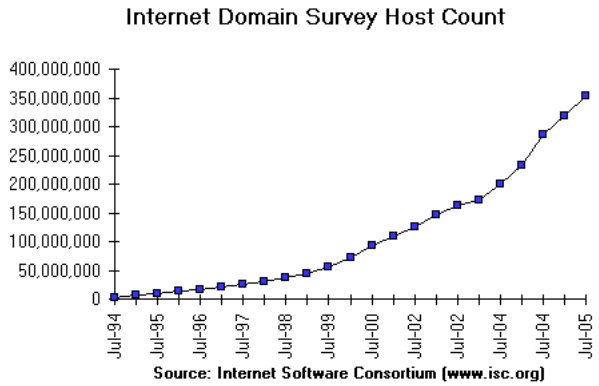

Intrinsic Unreliability

- Information sent from a first place to a second
  - May not arrive
  - May arrive more than once
  - May arrive in garbled fashion
  - May arrive out of order
  - May be read by others
  - May be modified by others
- Why build intrinsically unreliable networks?

Distributed

- (Hopefully) independent failure modes
- Exposed and hidden dependencies

> *"A distributed system is a system in which I can't do my work because some computer has failed that I've never even heard of."* – Lamport

- Independent administrative controls
- Leads to…

Heterogeneous Networks

- Heterogeneous: Made up of different kinds of stuff
- Homogeneous: Made up of the same kind of stuff
- Principles
  - Homogeneous networks are easier to deal with
  - Heterogeneous networks have their own strengths
    - Consider telephone network vs. Internet
  - What are the strengths?

Model of a Network

- Links carry information (bits)
  - Wire, wireless, fiber optic, smoke signals …
  - May be point-to-point or broadcast
- Switches move bits between links
  - Routers, gateways,bridges, CATV headend, PABXs, …
- Hosts are the communication endpoints
  - PC, PDA, cell phone, tank, toaster, …
  - Hosts have names
- Much other terminology: channels, nodes, intermediate systems, end systems, and much more.

Example – Local Area Network

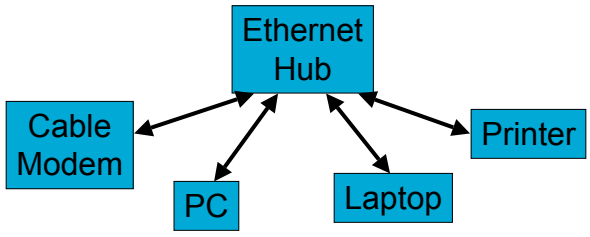

- Your home network
  - Ethernet is a broadcast-capable multi-access LAN

Example – An Internetwork

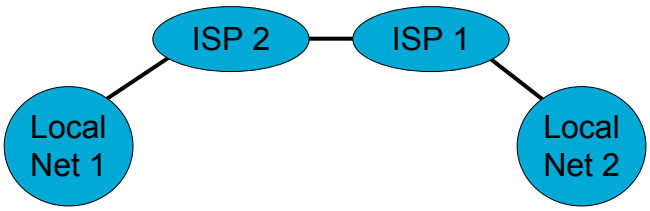

- Internetwork is a network of networks
- The Internet is a global internetwork in which all participants speak a common language
  - IP, the Internet Protocol

Goal of this Course

- You will understand how to design and build *large, distributed computer* networks.
  - Fundamental problems in building networks
  - Design principles of proven value
  - Common implementation technologies
- This is a systems course, not queuing theory, signals, or hardware design.
- We focus on networks, and a bit on applications or services that run on top of them (distributed systems).

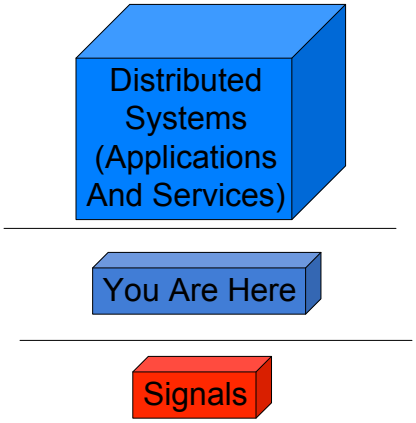

How to study networks?

- Networks in general, and Internet in particular, are complex beasts
- Question: how do we begin to understand Internet's workings?
  - Hands-on programming
  - Tinkering, reverse-engineering the network

Programming Projects

- Develop a peer-to-peer file sharing application for tablet PCs (Nokia N800)

- N800 has a linux variant
- Develop code on desktops, port to N800 (need to cross-compile)
- Programming in C
- Work in groups of two, each person will have a tablet for the course of the quarter

Project Outline

- Phase 1:
  - Simple peer-to-peer file sharing
  - Scan for nearby tablets, allow user to browse for files and download them
- Phase 2:
  - Multi-hop routing: find tablets that are accessible through multiple intermediate hops
  - Share files through intermediate nodes
- Phase 3:
  - Store-and-forward mode of sharing files
  - Register interest in a particular file, percolate information, obtain files and store locally till interested party is again within range

Explore and reverse-engineer networks

Residential access
- Modem
- DSL
- Cable modem
- Wireless

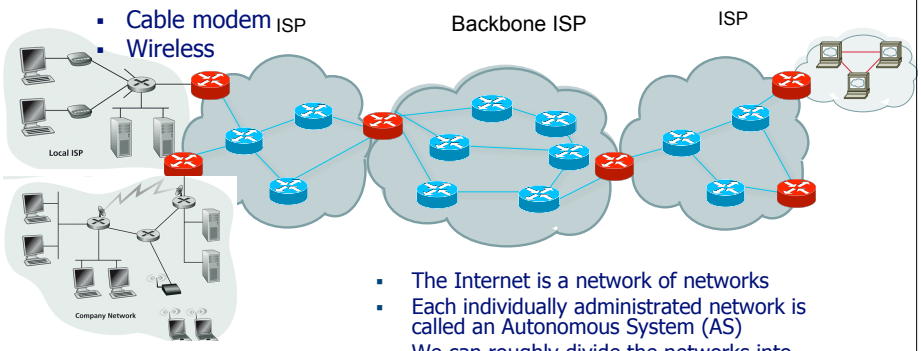

Campus access
- Ethernet
- FDDI
- Wireless

- The Internet is a network of networks
- Each individually administrated network is called an Autonomous System (AS)
- We can roughly divide the networks into access networks and transit networks

A Connectivity Exploration Tool

- Traceroute:
  - Run traceroute host-name on unix machines
  - tracert host-name on windows
- Sends three probes to each intermediate node on the path to the final destination (more details later)
- Reports the IP address, a more readable name, and the round-trip latencies for the probes

Traceroute to an East Coast College

```
-bash-3.1$ traceroute planetx.scs.cs.nyu.edu
traceroute to planetx.scs.cs.nyu.edu (216.165.109.79), 30 hops max, 40 byte packets
 1  acar-hsh-01-vlan75.cac.washington.edu (128.208.2.100)  0.362 ms  0.353 ms  0.396 ms
 2  uwcr-hsh-01-vlan3904.cac.washington.edu (205.175.110.17)  0.407 ms  0.444 ms  0.478
    ms
 3  uwcr-hsh-01-vlan1901.cac.washington.edu (205.175.103.5)  0.592 ms  0.665 ms  0.687 ms
 4  uwbr-ads-01-vlan1902.cac.washington.edu (205.175.103.10)  50.060 ms  50.120 ms
    50.130 ms
 5  hnsp2-wes-ge-0-0-0-0.pnw-gigapop.net (209.124.176.12)  0.703 ms  0.729 ms  0.760 ms
 6  abilene-pnw.pnw-gigapop.net (209.124.179.2)  0.544 ms  0.561 ms  0.588 ms
 7  dnvrng-sttlng.abilene.ucaid.edu (198.32.8.50)  46.984 ms  46.969 ms  47.009 ms
 8  kscyng-dnvrng.abilene.ucaid.edu (198.32.8.14)  63.746 ms  62.699 ms  62.709 ms
 9  iplsng-kscyng.abilene.ucaid.edu (198.32.8.80)  57.320 ms  57.305 ms  57.344 ms
10  chinng-iplsng.abilene.ucaid.edu (198.32.8.76)  70.506 ms  71.011 ms  70.985 ms
11  buf-7600-abilene-chin.nysernet.net (199.109.2.1)  73.003 ms  72.942 ms  72.946 ms
12  nyc-gsr-buf-7600.nysernet.net (199.109.7.14)  81.995 ms  81.966 ms  81.936 ms
13  nyu-nyc-gsr.nysernet.net (199.109.4.22)  82.179 ms  82.249 ms  82.314 ms
14  WWLABGW.NYU.NET (192.76.177.75)  82.350 ms  82.188 ms  82.200 ms
15  delancy.scs.cs.nyu.edu (216.165.108.191)  82.307 ms  82.662 ms  82.558 ms
16  planetx.scs.cs.nyu.edu (216.165.109.79)  82.629 ms  82.493 ms  82.592 ms
```

Abilene I2 Backbone

http://abilene.internet2.edu/maps-lists/

Traceroute to a commercial webserver

```
-bash-3.1$ traceroute www.nytimes.com
traceroute to www.nytimes.com (199.239.136.200), 30 hops max, 40 byte packets
 1  acar-hsh-01-vlan75.cac.washington.edu (128.208.2.100)  0.358 ms  0.357 ms  0.400 ms
 2  uwcr-hsh-01-vlan3904.cac.washington.edu (205.175.110.17)  0.426 ms  0.467 ms  0.502 ms
 3  uwcr-hsh-01-vlan1901.cac.washington.edu (205.175.103.5)  0.609 ms  0.639 ms  0.687 ms
 4  uwbr-ads-01-vlan1902.cac.washington.edu (205.175.103.10)  0.386 ms  0.428 ms  0.445 ms
 5  cnsp1-wes-ge-0-0-0-0.pnw-gigapop.net (209.124.176.8)  0.579 ms  0.643 ms  0.730 ms
 6  129.250.10.194 (129.250.10.194)  70.290 ms  66.878 ms  66.907 ms
 7  xe-1-2-0.r20.sttlwa01.us.bb.gin.ntt.net (129.250.2.206)  1.060 ms  1.063 ms  1.045 ms
 8  ae-0.r21.sttlwa01.us.bb.gin.ntt.net (129.250.2.54)  0.901 ms  0.901 ms  0.883 ms
 9  p64-2-0-0.r20.nycmny01.us.bb.gin.ntt.net (129.250.5.17)  74.106 ms  74.095 ms  74.103 ms
10  xe-4-1.r02.nycmny01.us.bb.gin.ntt.net (129.250.2.187)  141.125 ms  141.209 ms  141.305 ms
11  ge-1-1.a00.nycmny01.us.da.verio.net (129.250.30.113)  73.897 ms  73.997 ms  73.968 ms
12  * * *
13  * * *
14  * * *
```

A Commercial backbone: NTT

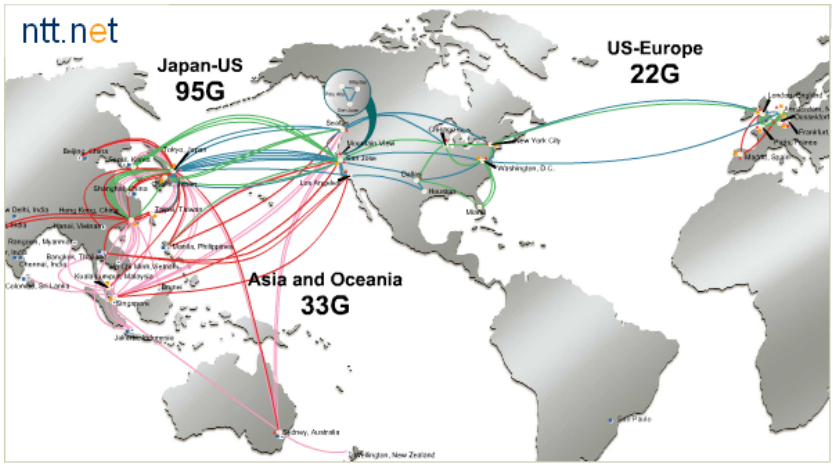

Traceroute to another commercial webserver

```
-bash-3.1$ traceroute www.nyse.com
traceroute to www.nyse.com (209.124.184.150), 30 hops max, 40 byte packets
 1  acar-hsh-01-vlan75.cac.washington.edu (128.208.2.100)  0.327 ms  0.353 ms  0.392 ms
 2  uwcr-hsh-01-vlan3904.cac.washington.edu (205.175.110.17)  0.374 ms  0.412 ms  0.443 ms
 3  uwcr-hsh-01-vlan1901.cac.washington.edu (205.175.103.5)  0.595 ms  0.628 ms  0.659 ms
 4  uwbr-ads-01-vlan1902.cac.washington.edu (205.175.103.10)  0.445 ms  0.472 ms  0.501 ms
 5  ccar1-ads-ge-0-0-0-0.pnw-gigapop.net (209.124.176.32)  0.679 ms  0.747 ms  0.775 ms
 6  a209.124.184.150.deploy.akamaitechnologies.com.184.124.209.in-addr.arpa (209.124.184.150)
      0.621 ms  0.456 ms  0.419 ms
```

What is going on?

```
-bash-3.1$ nslookup www.nyse.com
Name:   a789.g.akamai.net
Address: 209.124.184.137
Name:   a789.g.akamai.net
Address: 209.124.184.150
```

Points to note

- Multi-homed
- Certain routers don't respond
- Variability in response times
- Geography not apparent
  - Geography does not dictate paths
  - Sometimes paths are horribly inflated. Why?
- Content distribution networks operate by returning a nearby cache site
- Reverse engineering is fun!