

6.033 - Networking: Introduction
Lecture 7
Katrina LaCurts, lacurts@mit.edu

Some of the content of this lecture comes from "Why the Internet Only Just Works" by Mark Handley. It's a good paper! It's also from 2006, so much has happened since it was published.

0. Welcome to networking

- How networks work is a big part of systems
- Network = source of failure for many systems. Why?
- Today: intro to network / history of Internet. Look for common themes that you'll see over and over.
- Don't worry about details of specific protocols today; we'll delve into those later on.

1. General networks

- Model networks as graphs. Endpoints on outskirts, switches in the middle. Edge = direct connection between two nodes (perhaps a wire, perhaps not)
- Problems:
 - Addressing
 - Naming
 - Routing
 - Transport
- For small networks:
 - Naming/addressing: Just assign each node a unique name
 - Routing: How do nodes figure out how to get data to other nodes?
 - Transport: Lots of questions here. One problem to think about is reliability: what do we do when a packet gets dropped?
- In 6.033, we care mostly about the Internet. The Internet is not a "normal" network

2. The Internet, Pre-1993

- Launch of Sputnik => creation of ARPA (now DARPA)
- 1970S: ARPAnet. Started small. Combined addressing and transport. Began to grow by connecting existing networks
 - Using hosts.txt, distance-vector, sliding-window, etc.
- 1978: Decide to make the Internet flexible. Encourages a layered model. A common model:

[Application]	- things generating the traffic
[Transport]	- reliable (maybe) delivery, sharing
[Network - IP]	- addressing/routing
[Link]	- point-to-point links

- Ideally, layering lets us swap out protocols.
- ~1983: TCP, provides reliable transport. Now apps don't have to reimplement reliable delivery.

3. Growth => Change

- 1978-79: Link-state routing, EGP. More scalable routing.
- 1982: DNS. More scalable naming. Also enabled growth through distributed management.

4. Growth => Problems

- Mid 80s: Congestion collapse. So many packets in network, but none were useful. TCP added a congestion control mechanism (more on this next week).
 - Why add it to TCP, not as a separate layer? It was already hard to add a new layer, even in the 80s.
- Early 90s: Policy routing. Internet was beginning to be commercialized, parts of it didn't want to provide transit to commercial transit. Policy routing (BGP) is a means to do this.
 - Notice: commercialization is causing problems.
- Addressing. Assign addresses in chunks of different sizes. "Class B" chunks -- 65K -- are typically "just right", and we ran out of those. The protocol CIDR was developed to divide these up.
 - Most interesting thing about CIDR: that it was possible to make this change at all. Changing addressing => changing switches. Happened because all switches were made by Cisco and forwarding was done in software.

5. The Internet, post-1993

- 1993: Commercialization. Changes stopped. "New technologies essentially get deployed for reasons of fear or greed."
(<http://www0.cs.ucl.ac.uk/staff/M.Handley/papers/only-just-works.pdf>)

6. Problems we deal with today

- DDoS: send a lot of traffic to one machine to consume its resources. Hard to prevent. Internet wasn't designed with accountability in mind.
- Security: Internet was not designed for security. DNS, BGP, etc., have no secure infrastructure
- Mobility: No one every imagined this
- Address Space Depletion: IPv4 -> IPv6
- Congestion control: should probably change given the changes we've seen in the Internet (more about this later)

7. So what's new on the Internet?

- Lots. P2P, wireless, mobile, streaming, cloud computing, datacenter networks, security threats/defenses..
- Almost everything happens on the Internet today! Crazy.
- We expect to see continued changes: massive growth, software-defined networks, sensors, robots, embedded devices, privacy, censorship

8. Recurring themes

- Layering and hierarchy (particularly hierarchy as a means of

addressing scale)

- Scalability: If enforcing modularity was the theme of the first part of 6.033, scalability is the theme of the second part.
- Performance/efficiency: As the Internet grew, performance requirements informed its (re-)design.
- Diversity of Applications: As the Internet grew, more and more applications used it. These applications have different demands: some care about throughput, some care about latency, some care about both, some need reliability, some don't, etc. We have to build a network that works for all of them.
- Money: Because people pay to use the Internet, designing protocols can sometimes be trickier than it "should" be.
- End-to-end argument: what part of the network should implement what features (endpoints, switches, etc.)?