Final Review

To do …

- Logistics
- Basic ideas
- Class overview

Northwestern

Exam Friday Dec 13th @ 9-11AM

- Format for the exam
  - Cumulative
  - Open book/notes **but** paper, not digital device (laptops, phones, pads)
  - 1-2 true/false, 4 short, 2-3 long, 1-2 extra points
- What to use to prep
  - Exams
    - Practice midterm | practice final
  - Homework assignments
    - HW answers should be posted before Friday
  - Projects and labs
  - Lecture slides
  - Book Chapters 1-7 (look at the calendar for the specific sections)

Separation of concerns

- Link layer – shares a physical channel among several transmitters/receivers
- Network layer – routes from src to dst, along many hops
- Transport layer – multiplexing, ordering, ack, reliability, congestion
- Applications – HTTP, DNS, eMail, CDN, P2P …

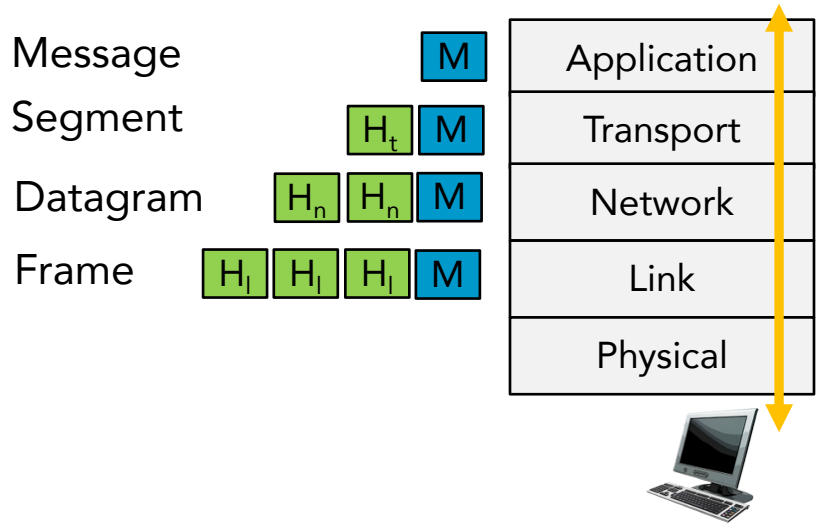

Decentralization

- Little centrally-controlled infrastructure for the Internet
  - ICANN controls the distribution of IP address and domain names
- Internet standards are developed by the IETF through an open process, leading to RFCs
- DNS: 13 root server IP addresses, then top level domains (TLDs)
  - DNS servers at the edge of the network cache results
- BGP uses DV algorithm to compute shortest path
- TCP congestion control mitigates core congestion from observations at the edge

Fault tolerance

- Assume that links and routers are unreliable
  - Bits can be flipped, packets can be dropped
- Bit error checking is included in Ethernet, IPv4, TCP/UDP
- BGP allows routes to change in response to broken links
- TCP provides delivery confirmation, retransmission, and ordering
- HTTP response code can indicate an error
  - (eg., 404 Not Found, 500 Internal Server Error)

Class overview

Introduction

- A networking class focused on the Internet
- From basic building blocks – hosts, links and routers
- Circuit switching (Plain old tel.) and packet switching (Internet)
- Statistical multiplexing, store and forward and best effort service
- Historical context to understand how we got here
  - From the early 1960s ARPA efforts to the Internet explosion of the 90s and the challenges of a new millennium

Edge, core and a bunch of layers

- Network edge, access networks and the network core
- Packets travel along many hops to reach the intended destination
  - Each router has a fixed-size queue; packets are dropped if full or bit error is detected
- Different sources of packet delay – nodal processing, queueing (@ router), transmission, propagation (on link)
- Different access technologies – DSL, Cable, FTTH, satellite
- The network core – the Internet as a "network of networks"
- Protocol, as a set of communication rules, protocol layers and encapsulation

Application-layer protocols

- Client-server, peer-to-peer & the inevitable hybrid architecture
- Application-layer protocols and their transport choices
- Processes and sockets
- HTTP request/response protocol
  - Stateless, persistent and non persistent connections, message formats
  - Cookies and web caching
- SMTP is an earlier application-layer protocol, for sending email
  - Unlike HTTP, stateful, today's mail access with POP3, IMAP or HTTP

Domain Name Service

- DNS is the Internet's directory service
- It's distributed and hierarchical
  - 13 Root servers are run by ICANN
  - Top level domain (TLD) servers manage com, org, edu, cn, au, uk, etc.
  - Each subdomain has a set of authoritative nameservers
- Various types of records for more than just map name → IP
- Domain registrars are accredited by each TLD to sell names
- Resolution process
- Changes in DNS – public DNS, DNS over HTTPS

CDNs and P2P

- New applications with new demands on the underlying network
- Architectural changes are, at best, difficult
- Overlays as a path to deployment and an experimental testbed
- CDNs – ideas and challenges - how to replicate, where, how to choose among replicas, how to redirect clients' requests
- P2P from grassroots efforts and research labs to products
  - Early P2P systems, 2nd generation: unstructured and decentralized, 3rd gen: structured and decentralized (DHTs)
  - Many open hard issues – security, churn, …
- A hybrid example with Akamai's NetSession

Transport layer and UDP

- A transport-layer protocol provides for *logical* communication between app processes running in different hosts
  - Implemented at end hosts, not in routers
- What do they offer – from network layer's host-to-host to process-to-process (port numbers), integrity checking (checksums), reliable data transfer, congestion control
- UDP – not much of a transport
  - Nearly the same as IP
  - Just allowing multiple applications on a host to share one network
  - Plus some error checking
- The value of UDP and some applications

Reliable Transport

- Mechanisms for reliability: checksums, timers, sequence numbers, ACKs and NACKs, and pipelining
- Delivery confirmation & ordering is possible by sending ACKs
  - After a timeout, resend packet that was not ACK'ed
- Pipelining packets allow much better use of link capacity
  - Parallelizes ACK'ed communication
  - Window size determines the number of allowed in-flight packets
- Go Back N is a simple pipelining protocol that uses cumulative ACKs
- Selective Repeat adds buffering to the receiver to avoid unnecessary repetition

TCP

- Cumulative, delayed ACKs
- Receivers buffer out-of-order (early) segments for later reassembly.
- ACK timeout can be set using exponentially-weighted moving average of observed RTT
- # in-flight packets/throughput is determined by window size
- Flow control is implemented with explicit receive window
- Connection setup requires a 3-way handshake
  - Sets initial sequence numbers and receive windows in both directions

TCP Congestion Control

- Congestion control is implemented with a dynamic congestion window, controlled by heuristics that operate in phases:
  - Slow start – exponential growth to find approximate network capacity
  - Congestion avoidance – linear growth, slowly trying to increase throughput
  - Fast recovery – If one packet is lost, then cut window in half
- Adapting to changing network conditions and fairness
- TCP behavior can be controlled with socket options
  - Nagle's algorithm merges small packets to reduce header overhead
  - TCP keepalive message can be sent periodically

Network layer | Inside a router

- Network layer functions – forwarding and routing
- Data and control planes
- Possible services from a network layer and what you get from IP
- Data plane, *what we have seen so far*, and control plane
  - The per-router functions that determine how packets arriving on one of the router input ports is forwarded to one of the output ports
- Routers internal, mechanisms and policies

Internet protocol

- IPv4 fragments packets larger than MTU. Are reassembled at the destination end-host
- IPv4 header is 20B (UDP header is 8B or TCP header is 20B)
- IP subnets define ranges of address that can comm directly
  - CIDR notation (123.100.16.0/28) specifies a range of addresses
- Host's IP config – address, subnet mask, gateway, DNS server
  - Gateway is IP address of the router to route packets outside the subnet
  - DHCP allows newly-arriving machines to request an IP configuration.
- NAT shares one IP address by port mapping (rewriting)
- IPv6 changes, deployment and tunneling

Routing algorithms

- Routing algorithms determine each router's forwarding table
  - Centralized/global or distributed
- Routing is a shortest path problem on this graph
- Dijkstra's Algorithm is a fast centralized link-state algorithm
  - Used by Open Shortest Path First (OSPF) protocol within an AS
- Distance Vector is a distributed shortest path algorithm
  - Used by the BGP to route between AS's
  - Count to infinity leads to slower convergence when links get worse
    - Good news travels quickly, bad news travels slowly

Internet routing

- Routing is hierarchical
  - ASes are groups of routers with a single routing policy
- Interior gateway protocols (eg., OSPF) determines optimal routes within an AS
  - Can use a centralized (Link State) shortest path algorithm, like Dijkstra's
- BGP determines routes between ASes.
  - Uses a distributed shortest-AS-hop path algorithm (Distance Vector).
- BGP advertisement includes a list of routes, each looking like:
  - {PREFIX: 43.5.0.0/16, AS-PATH: [AS4, AS65, AS1], NEXT-HOP: 5.6.7.200)}
  - This tells a neighboring AS that it can forward packets to the prefix

Network management and control

- ICMP – Used by hosts & routers to comm network-level info
  - Mainly error reporting: unreachable host, network, port, protocol
  - Network-layer, but "above" IP
    - ICMP msgs carried in IP datagrams, upper-layer protocol #1
  - Ping and traceroute
- SNMP – application-layer protocol to convey network mgmt control and information between managing server and agents
- From traditional control plane to Software defined networking
  - Logically centralized control
  - Generalized forwarding and OpenFlow
  - Network apps, SDN controller and data-plane switches

Link Layer

- Link-layer handles sharing a link/medium with multiple nodes.
  - Also, handles error detection and correction: Parity, Checksum, and CRC.
- Medium Access Control / Multiple Access Protocol
  - Decide how to share the link.
  - Two nodes sending simultaneously is a collision. Packets are lost.
- Three classes of sharing protocols:
  - Channel Partitioning: Frequency Division Multiplexing, Time Division Mux.
  - Random Access: ALOHA, CSMA/CD (Carrier Sense Multiple Access/Collision Detection)
  - Turn-Taking: Polling, Token-passing

Ethernet Link Layer

- Ethernet adds MAC addresses to identify hosts on a shared link
  - ARP uses Ethernet broadcast to find IP → MAC address mapping
- DHCP requests are sent by Ethernet broadcast
- Old Ethernet hubs broadcasted data to all ports
- Ethernet switches learn which MAC addresses are reachable on each port and relay traffic only to the appropriate ports
  - Reduce broadcast traffic and eliminate collisions
- VLANs create multiple isolated LANs/subnets on one switch
- Virtualizing links with MPLS for faster routing
- Data centers also require carefully designed networks

Wireless and Mobile

- Wireless hosts, links and base stations
- Important differences and specific problems that make wireless networks challenging
  - Decreased signal strength, interference, multipath propagation
- 802.11 LAN, architecture, multiple access, frames
- Component of cellular networks – 2G, 3G and 4G features
- Mobility vocabulary and basic approaches
  - Indirect and direct routing
- Mobility in practice – IP mobility and Cell mobility
- Impact of wireless and mobility on higher layers