Computer Networks
ICS 651

- Layers and Device Names
- NAT and Firewall
- DHCP
- IPv6 autoconfiguration
- home network design
- larger Ethernet network design

Names and Layers

- layer 1, physical layer: Hub, forward to all (physical layer)
- layer 2, data link (MAC) layer:
  - Bridge: something interconnecting LANs
  - Learning Bridge: only forwards on interface that reaches the destination
  - Switch: a bridge that uses hardware to forward packets
- layer 3, network (internet) layer:
  - Gateway: old name for router
  - Router: a box that forwards packets among networks
  - IP switch: a router with hardware forwarding
- layer 4, transport layer:
  - Firewall: a forwarding engine that allows or blocks traffic based on port numbers
  - NAT (Network Address Translation): a forwarding engine that looks at and rewrites network layer and transport layer headers

Ethernet and Wifi home setup

- one host, H, connected to the internet (the “router”)
- internally, 100Mb/s or 1Gb/s switch-based Ethernet and password-protected WiFi
- externally cable modem, DSL, or other technology
- H does Network Address Translation (Layer 4 translation) so:
  - IP packets going out are rewritten to have H as source address, and often a different source port
  - incoming packets are rewritten to have the correct destination host and port number
  - for ICMP Echo, can rewrite identifier
  - for ICMP Error messages, need to look up port number in original header
- H can also:
  - perform firewall functions
  - allocate addresses via DHCP or NDP
  - be the default router for all internal computers

Network Address Translation and Firewall

- for TCP we know when the connection starts and ends, so can de/allocate state in the NAT box
  - and we generally disallow inbound connections
- for UDP/ICMP we don't know when the connection is done, so we must cache: allocate when the first packet between (pair of IPs and ports) is sent from inside to outside, deallocate after a timeout
  - drop inbound UDP/ICMP that do not match a translation
- it generally doesn’t matter what local port we use for outgoing connections
- to run a server, must “poke a hole” in the firewall, i.e. tell the firewall to accept inbound TCP connections to a given port

Dynamic Host Configuration Protocol

- if a host is permanently connected to the Internet, it might as well have a fixed address
  - servers need fixed addresses!
- hosts that connect dynamically need to obtain an address when they connect
- DHCP provides dynamic assignment, on demand, of temporary IP address "leases"
- a lease can be renewed, but no guarantees
- hard to allow TCP connections initiated outside the network, since the IP address may change:
  - bad for servers
  - good for security
  - or, could hard-code some of the DHCP addresses to assign only to specific MAC addresses
- a DHCP block can be configured in the firewall, so no incoming connections would be allowed for DHCP addresses
- harder to automatically snoop on specific host (again, both good and bad)

IPv6 dynamic address assignment

- IPv6 interfaces often have multiple IPv6 addresses
  - at least one of the addresses is typically link-local, fe80::/10, with the low-order 64 bits derived from the MAC address or selected at random
  - when using link-local addresses in a sockets program, must specify the interface
- NDP tells me the network number and the netmask for global IP addresses
  - again, the low-order bits can be selected at random, or derived from the MAC address
  - as in IPv4, servers need one constant IPv6 address
  - mobile clients, which may not wish to be tracked, often create a new random global IPv6 address on a regular basis
- NDP is also used to make sure that the random address I selected is not already in use on the local network

Larger Ethernet Networks

- star topology using hubs
- at most 250 m. (for 100Mb/s) with up to 4 repeaters between any two hosts (one hub to a central hub)
- avoid collisions if possible:
  - split into segments connected by switches, routers
    - if using routers to connect segments, each segment is a different IP network, must assign IP addresses to subnetworks
  - switch to 100Mb/s, Gb, 10Gb Ethernet
  - connect host directly to switch in full-duplex mode

Ethernet Collision Probability

- likelihood that two hosts are waiting to send at end of current packet (assume all packets are the same size)
- load (probability of 1 host wanting to send) is p
- probability of collision is p^2
- cost of collision is: 1 packet lost to collision (plus latency), giving additional load: delta = $\sum_{i = 1 \ldots \infty} p^{2i}$

load p	p^2	delta
0.1	0.01	0.01
0.3	0.09	0.1
0.5	0.25	0.3

Summary

- hub, switch, router, NAT+firewall
  - firewall blocks incoming connections, which prevents port scanning and some attacks
- home network setup is mostly automated
- designing larger networks requires consideration of traffic as well as security
- ethernets should be designed to be either strictly point-to-point (preferably full-duplex), or mostly unloaded