Chapter 5
Link Layer

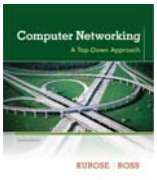

Computer Networking: A Top Down Approach
6th edition
Jim Kurose, Keith Ross
Addison-Wesley
March 2012

A note on the use of these ppt slides:
We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you see the animations; and can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a lot of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) that you mention their source (after all, we'd like people to use our book!)
- If you post any slides on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

The course notes are adapted for Bucknell's CSCI 363
Xiannong Meng
Spring 2014

Link layer, LANs: outline

5.1 introduction, services
5.2 error detection, correction
5.3 multiple access protocols
5.4 LANs
- addressing, ARP
- Ethernet
- switches
- VLANS

5.5 link virtualization: MPLS
5.6 data center networking
5.7 a day in the life of a web request

Multiprotocol label switching (MPLS)

- initial goal: high-speed IP forwarding using fixed length label (instead of IP address)
  - fast lookup using fixed length identifier (rather than shortest prefix matching)
  - borrowing ideas from Virtual Circuit (VC) approach
  - but IP datagram still keeps IP address!

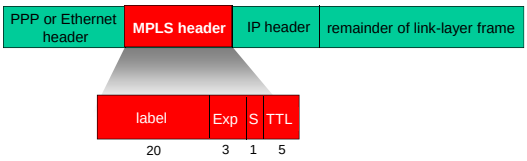

MPLS capable routers

- a.k.a. label-switched router
- forward packets to outgoing interface based only on label value (*don't inspect IP address*)
  - MPLS forwarding table distinct from IP forwarding tables
- *flexibility:* MPLS forwarding decisions can *differ* from those of IP
  - use destination *and* source addresses to route flows to same destination differently (traffic engineering)
  - re-route flows quickly if link fails: pre-computed backup paths (useful for VoIP)

MPLS versus IP paths

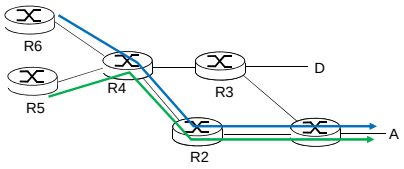

- *IP routing: path to destination determined by destination address alone*

IP router

MPLS versus IP paths

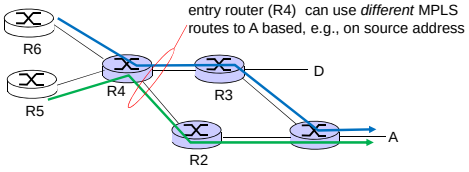

- *IP routing: path to destination determined by destination address alone*
- *MPLS routing:* path to destination can be based on source *and* dest. address
  - *fast reroute:* precompute backup routes in case of link failure

IP-only router

MPLS and IP router

MPLS signaling

- modify OSPF, IS-IS link-state flooding protocols to carry info used by MPLS routing,
  - e.g., link bandwidth, amount of "reserved" link bandwidth
- *entry MPLS router uses RSVP-TE signaling protocol to set up MPLS forwarding at downstream routers*

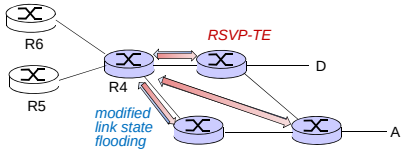

MPLS forwarding tables

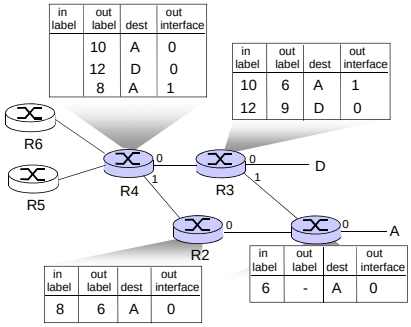

Link layer, LANs: outline

5.1 introduction, services
5.2 error detection, correction
5.3 multiple access protocols
5.4 LANs
- addressing, ARP
- Ethernet
- switches
- VLANS

5.5 link virtualization: MPLS
5.6 data center networking
5.7 a day in the life of a web request

Some Statistics About Google

- June 2006 New York Times reported that Google's Dalles, Oregan Data Center would create 60-120 full-time jobs
- Rate of changes:
  - March 2001, serving 70 million web pages with 8,000 computers
  - In 2003, the number of computers grew to 100,000
  - The June 2006 estimate was 450,000 scattered over at least 25 locations
  - Jeff Dean's presentation anticipates 10 million computers in the next a few years

A Few Other Google Data Centers (1)

- Belgium Data Center:
  - 341 million dollars
  - 120 long-term jobs
  - No chillers, just air-cooling, or called "free-cooling"
  - Open early 2008
- Lenoir, North Carolina Data Center
  - 600 million dollars
  - 150 acres
  - 200 long-term jobs

A Few Other Google Data Centers (2)

- Goose Creek, South Carolina Data Center
  - 600 million dollars
  - 520 acres
  - 200 jobs

Google Data Center Information

- Estimated that Google has about 36 data centers world-wide
  - Google maintains its data center temperature at 80 F
  - Location map of Google data centers http://www.theregister.co.uk/2008/04/11/google_data_center_map/
  - Google data center video: http://www.youtube.com/watch?v=zRwPSFpLX8I
  - Google data center FAQ: http://www.datacenterknowledge.com/archives/2012/05/15/google-data-center-faq/

13

Other Data Centers (Microsoft 1)

- The Chicago Data Center:
  - 700,000 square feet – approximately the size of 16 football fields;
  - Up to 220 shipping containers packed with servers; each container houses 1,800 to 2,500 servers, which makes 396,000 to 550,000 servers at the site! Each container can be wheeled in and out of the center in hours and be connected to the internet.
  - Another source said a total capacity of 112 containers holding 224,000 servers.
  - Go live July 20, 2009.

14

Other Data Centers (Microsoft 2)

- The Dublin Data Center (Ireland):
  - 303,000 square feet, first phase, eventually will reach 550,00 square feet
  - Supported by 5.4 megawatts electricity.
  - Go live on July 1, 2009. (Started in summer 2007.)
  - 500 million dollar project.
  - Expected to create 250 jobs.

15

Other Data Centers (Microsoft 3)

- The Quincy Data Center (Washington):
  - 470,000 square feet room on a 75 acre site
  - A tile floor and a maze of rooms centering around five 12,000-square-foot brain centers that contain tens of thousands of computer servers.
  - Temperature kept in between 60 and 68 degree Fahrenheit
  - Collects rainwater from its roof to use in its cooling system.
  - Row after row of batteries to kick in for 18 seconds if a power failure should occur before the truck-sized backup generators fire up.
  - Go live on March 27, 2007

16

Other Data Centers (Microsoft 4)

- The San Antonio Data Center (Texas):
  - 550 million dollar project
  - 470,000 square foot
  - Two 10 megawatt utility feeds, each expandable to 30 megawatts each.

17

Yahoo Data Center

- Lockport, NY Data Center
  - Using Niagara Fall as its power source
  - 150 million dollars
  - 30 acre site
  - 181,000 square feet rooms
  - First phase : 10 megawatts of hydro-electric power; second phase another 15 megewatts
  - To be completed by May 2010
- Quincy, Washington Yahoo Data Center

18

Data center networks

- 10's to 100's of thousands of hosts, often closely coupled, in close proximity:
  - e-business (e.g. Amazon)
  - content-servers (e.g., YouTube, Akamai, Apple, Microsoft)
  - search engines, data mining
- challenges:
  - multiple applications, each serving massive numbers of clients
  - managing/balancing load, avoiding processing, networking, data bottlenecks

Inside a 40-ft Microsoft container, Chicago data center

Data center networks

load balancer: application-layer routing

- receives external client requests
- directs workload within data center
- returns results to external client (hiding data center internals from client)

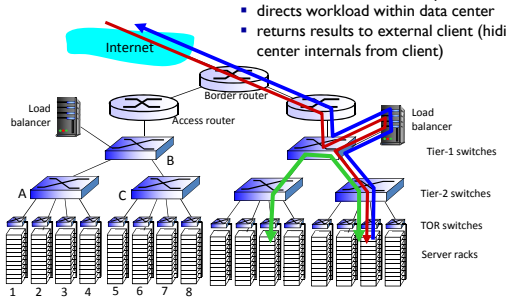

Data center networks

- rich interconnection among switches, racks:
  - increased throughput between racks (multiple routing paths possible)
  - increased reliability via redundancy

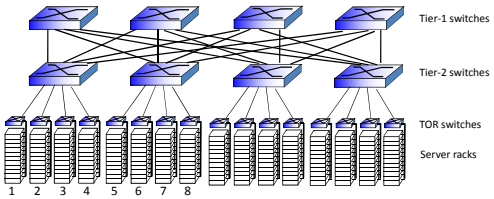

Link layer, LANs: outline

5.1 introduction, services
5.2 error detection, correction
5.3 multiple access protocols
5.4 LANs
- addressing, ARP
- Ethernet
- switches
- VLANS

5.5 link virtualization: MPLS
5.6 data center networking
5.7 a day in the life of a web request

Synthesis: a day in the life of a web request

- journey down protocol stack complete!
  - application, transport, network, link
- putting-it-all-together: synthesis!
  - *goal:* identify, review, understand protocols (at all layers) involved in seemingly simple scenario: requesting www page
  - *scenario:* student attaches laptop to campus network, requests/receives www.google.com

A day in the life: scenario

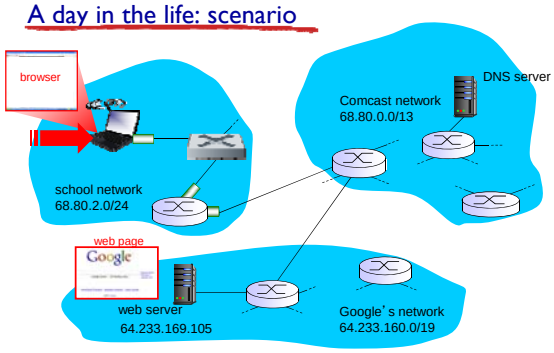

A day in the life… connecting to the Internet

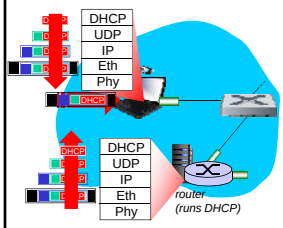

- connecting laptop needs to get its own IP address, addr of first-hop router, addr of DNS server: use *DHCP*
- DHCP request *encapsulated* in *UDP*, encapsulated in *IP*, encapsulated in *802.3* Ethernet
- Ethernet frame *broadcast* (dest: FFFFFFFFFFFF) on LAN, received at router running *DHCP* server
- Ethernet *demuxed* to IP demuxed, UDP demuxed to DHCP

A day in the life… connecting to the Internet

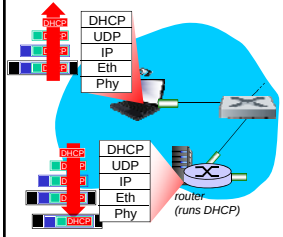

- DHCP server formulates *DHCP ACK* containing client's IP address, IP address of first-hop router for client, name & IP address of DNS server
- encapsulation at DHCP server, frame forwarded (*switch learning*) through LAN, demultiplexing at client
- DHCP client receives DHCP ACK reply

Client now has IP address, knows name & addr of DNS server, IP address of its first-hop router

A day in the life… ARP (before DNS, before HTTP)

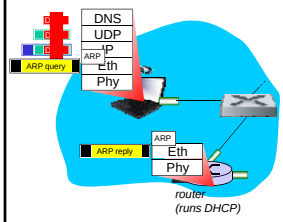

- before sending *HTTP* request, need IP address of www.google.com: *DNS*
- DNS query created, encapsulated in UDP, encapsulated in IP, encapsulated in Eth. To send frame to router, need MAC address of router interface: *ARP*
- *ARP query* broadcast, received by router, which replies with *ARP reply* giving MAC address of router interface
- client now knows MAC address of first hop router, so can now send frame containing DNS query

A day in the life… using DNS

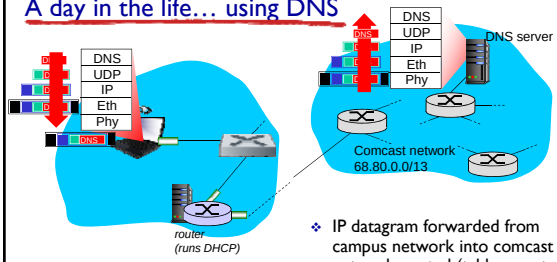

- IP datagram containing DNS query forwarded via LAN switch from client to 1st hop router
- IP datagram forwarded from campus network into comcast network, routed (tables created by *RIP, OSPF, IS-IS* and/or *BGP* routing protocols) to DNS server
- demux'ed to DNS server
- DNS server replies to client with IP address of www.google.com

A day in the life…TCP connection carrying HTTP

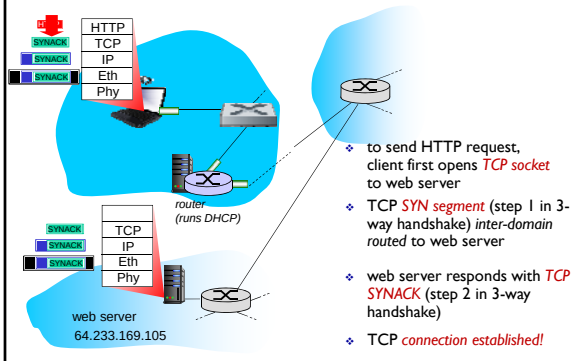

- to send HTTP request, client first opens *TCP socket* to web server
- TCP *SYN segment* (step 1 in 3-way handshake) *inter-domain routed* to web server
- web server responds with *TCP SYNACK* (step 2 in 3-way handshake)
- TCP *connection established!*

A day in the life… HTTP request/reply

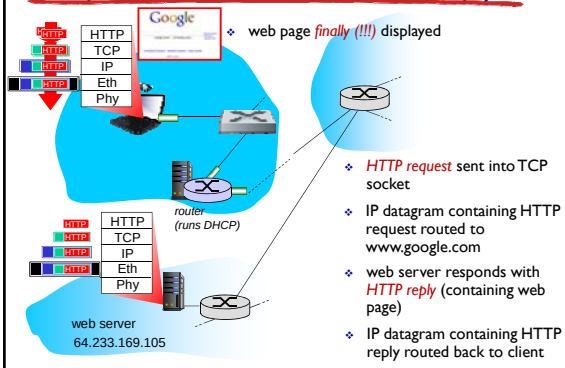

- web page *finally (!!!)* displayed
- *HTTP request* sent into TCP socket
- IP datagram containing HTTP request routed to www.google.com
- web server responds with *HTTP reply* (containing web page)
- IP datagram containing HTTP reply routed back to client

Chapter 5: Summary

- principles behind data link layer services:
  - error detection, correction
  - sharing a broadcast channel: multiple access
  - link layer addressing
- instantiation and implementation of various link layer technologies
  - Ethernet
  - switched LANS, VLANs
  - virtualized networks as a link layer: MPLS
- synthesis: a day in the life of a web request

Chapter 5: let's take a breath

- journey down protocol stack *complete* (except PHY)
- solid understanding of networking principles, practice
- ….. could stop here …. but *lots* of interesting topics!
  - wireless
  - multimedia
  - security
  - network management