Lecture 9:
Enterprise Networking

CS 598: Advanced Internetworking

Matthew Caesar

March 8, 2011

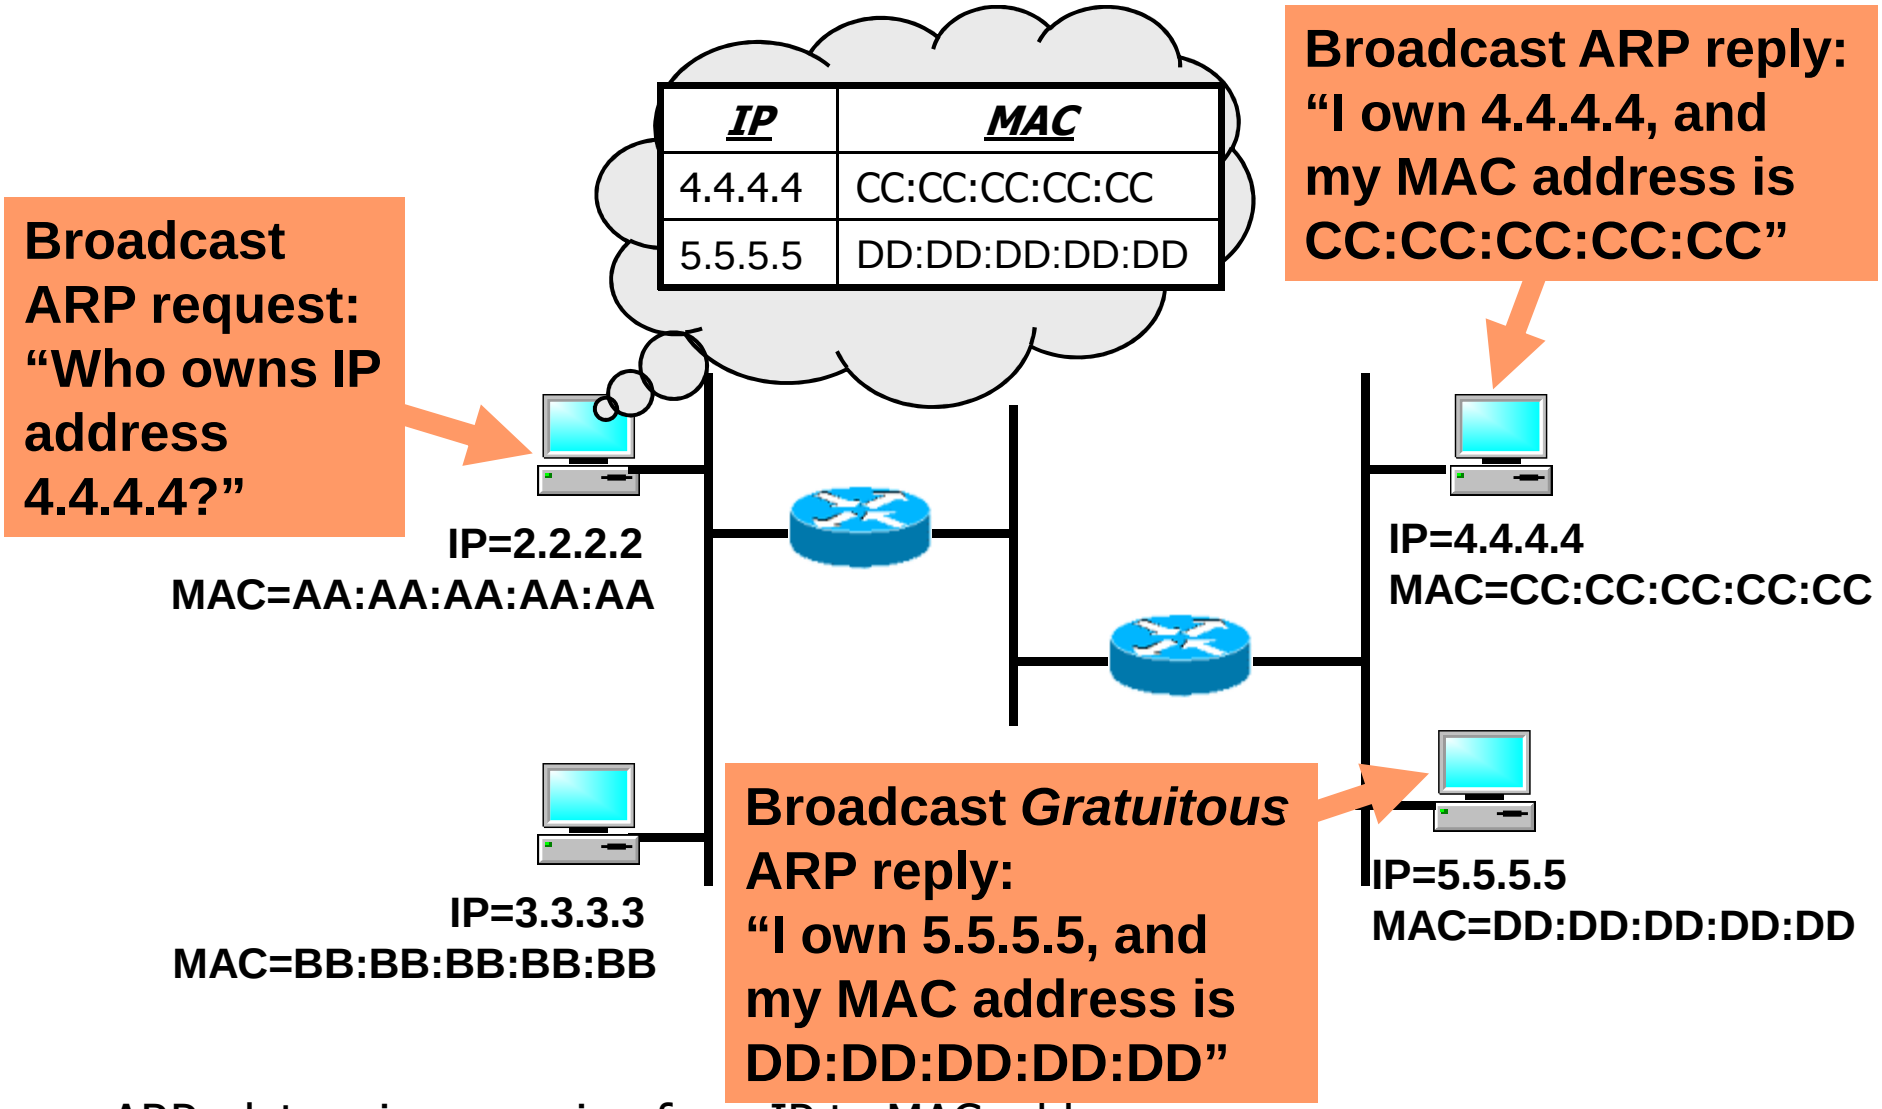

- ARP: determine mapping from IP to MAC address
- What if IP address not on subnet?
  - Each host configured with "default gateway", use ARP to resolve its IP address
- Gratuitous ARP: tell network your IP to MAC mapping
  - Used to detect IP conflicts, IP address changes; update other machines' ARP tables, update bridges' learned information

Physical Layer: Repeaters

- Distance limitation in local-area networks
  - Electrical signal becomes weaker as it travels
  - Imposes a limit on the length of a LAN
- Repeaters join LANs together
  - Analog electronic device
  - Continuously monitors electrical signals on each LAN
  - Transmits an amplified copy

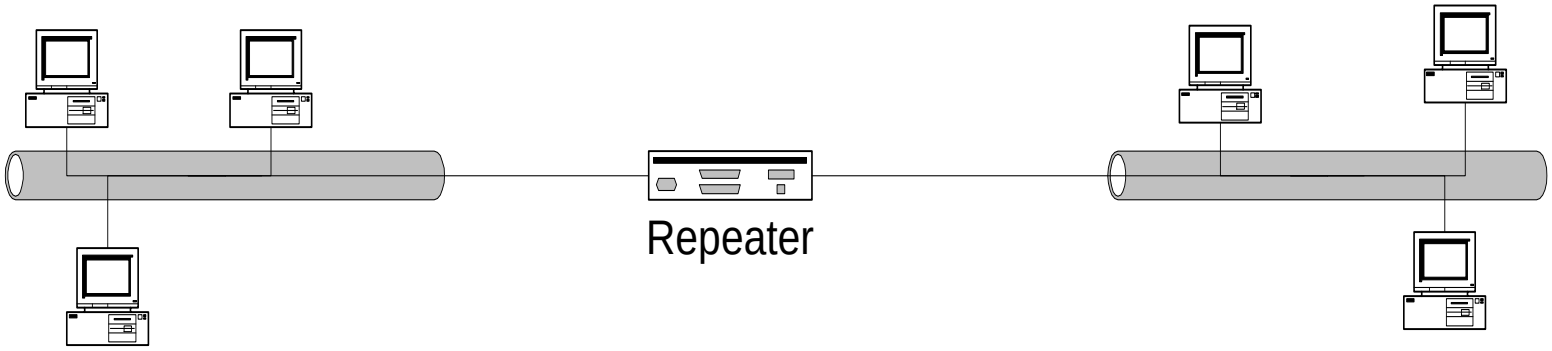

Physical Layer: Hubs

- Joins multiple input lines electrically
  - Do not necessarily amplify the signal
  - Very similar to repeaters
- Disadvantages
  - Limited aggregate throughput due to shared link
  - Cannot support multiple rates or formats (e.g., 10 Mbps vs. 100 Mbps Ethernet)
  - Limitations on maximum # of nodes and physical distance

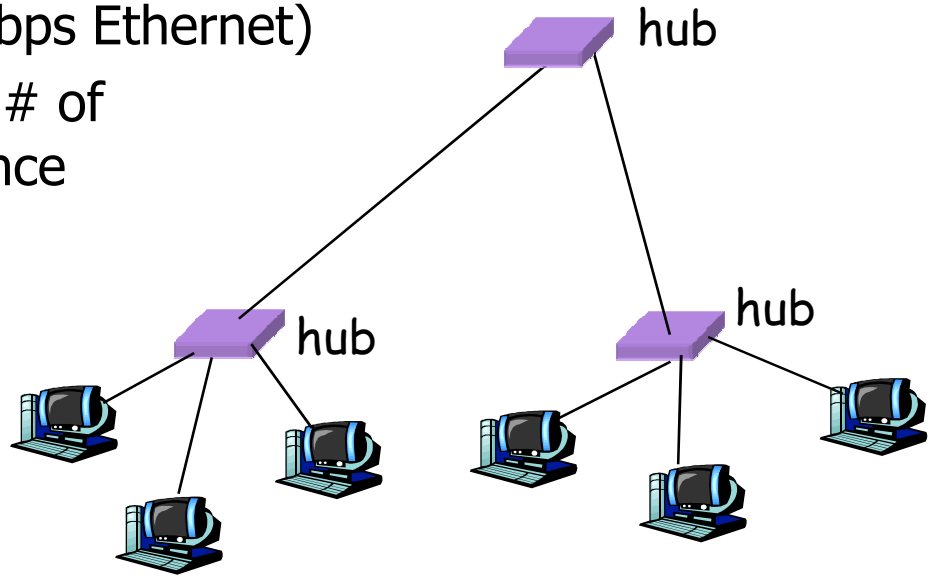

Link Layer: Bridges

- Connects two or more LANs at the link layer
  - Extracts destination address from the frame
  - Looks up the destination in a table
  - Forwards the frame to the appropriate LAN segment
- Each segment can carry its own traffic

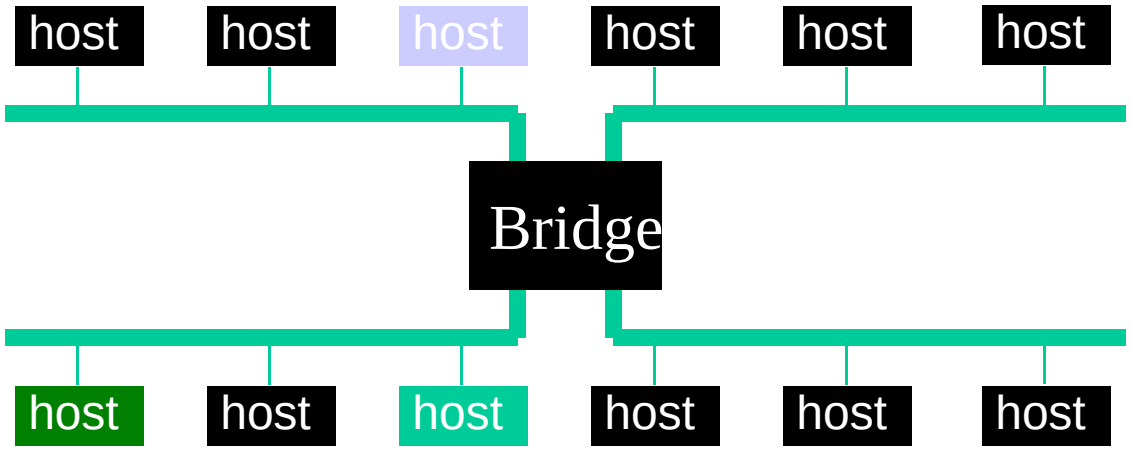

Link Layer: Switches

- Typically connects individual computers
  - A switch is essentially the same as a bridge
  - Supports concurrent communication
- Cut-through switching
  - Start forwarding a frame while it is still arriving

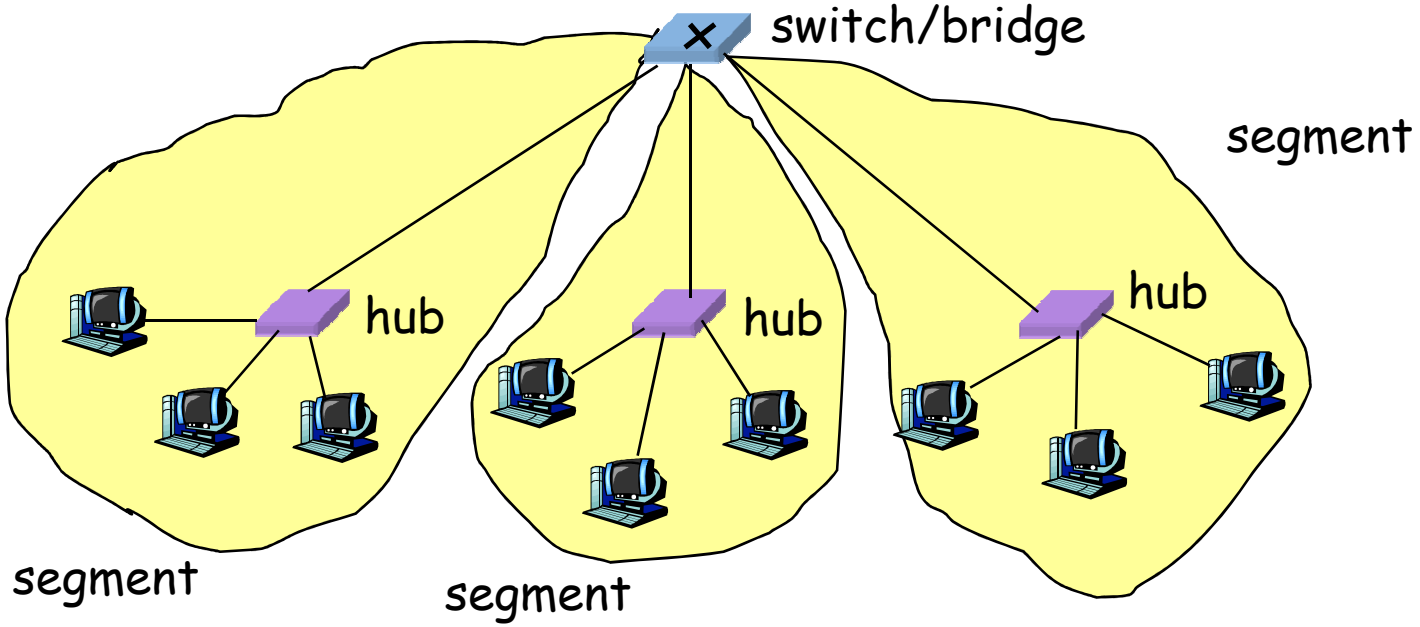

Hubs, Switches, and Routers

	Hub/ Repeater	Bridge/ Switch	Router
Protocol layer	physical	link	network
Traffic isolation	no	yes	yes
Plug and play	yes	yes	no
Efficient routing	no	no	yes
Cut through	yes	yes	no

Enterprise Network Design

Simple Enterprise Design

- A single layer-two subnet
  - Hubs and switches
  - Gateway router connecting to the Internet
  - ISP announces the address block into BGP
- Local services: DHCP and DNS

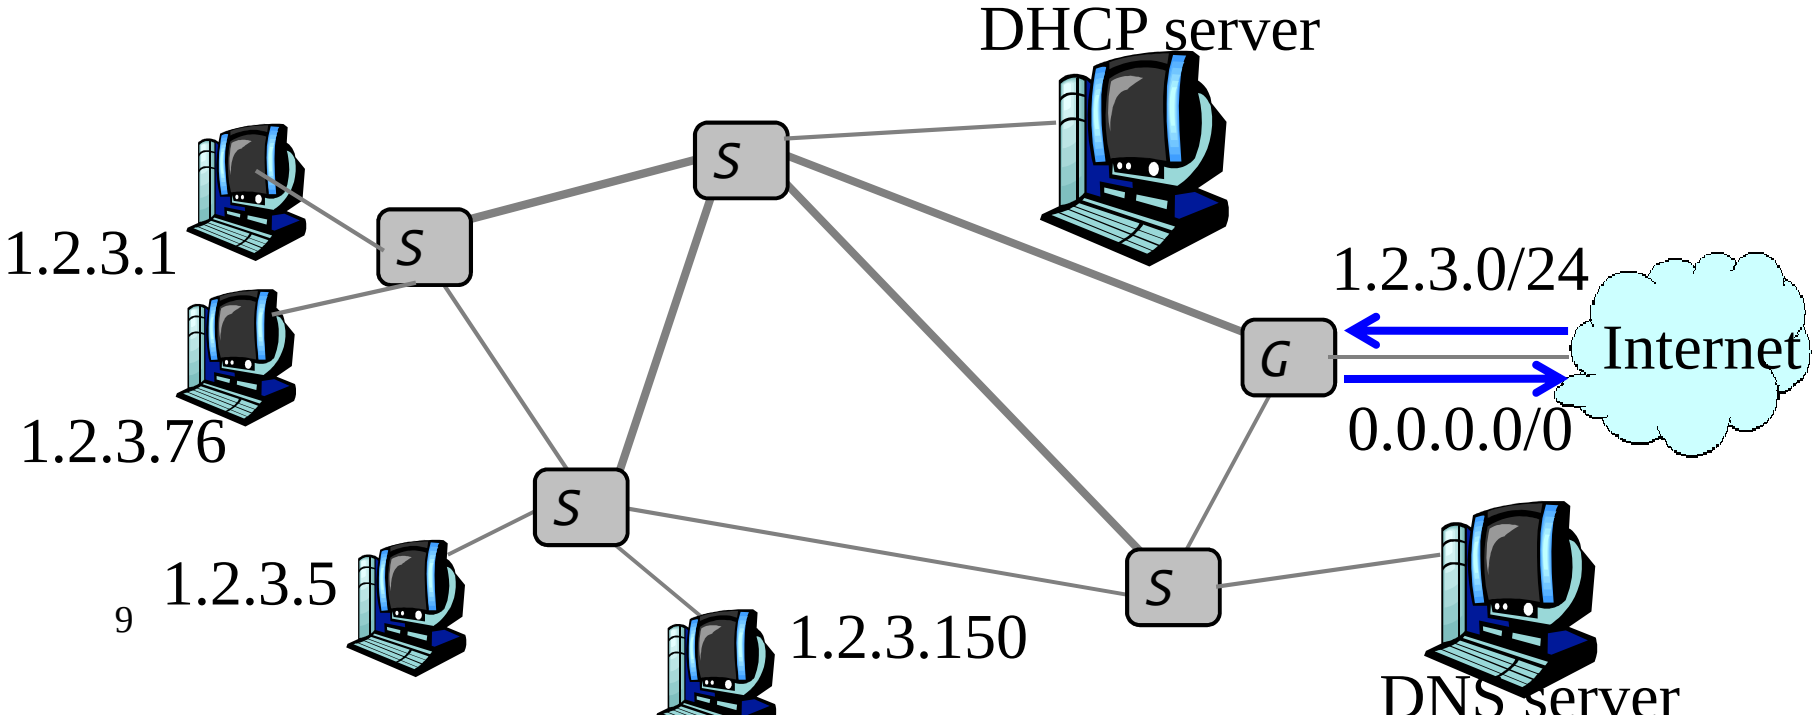

Scalability Limitations

- Spanning tree
  - Paths that are longer than necessary
  - Heavy load on the root bridge
  - Bandwidth wasted for links not in the tree
- Forwarding tables
  - Bridge tables grow with number of hosts
- Broadcast traffic
  - ARP and DHCP
  - Applications that broadcast (e.g., iTunes)
- Flooding
  - Frames sent to unknown destinations

Hybrid of Switches and Routers

- Layer-two subnets interconnected by routers
  - No plug-and-play and mobility between layer-2 subnets
  - Need consistent configuration of IP routing and DHCP

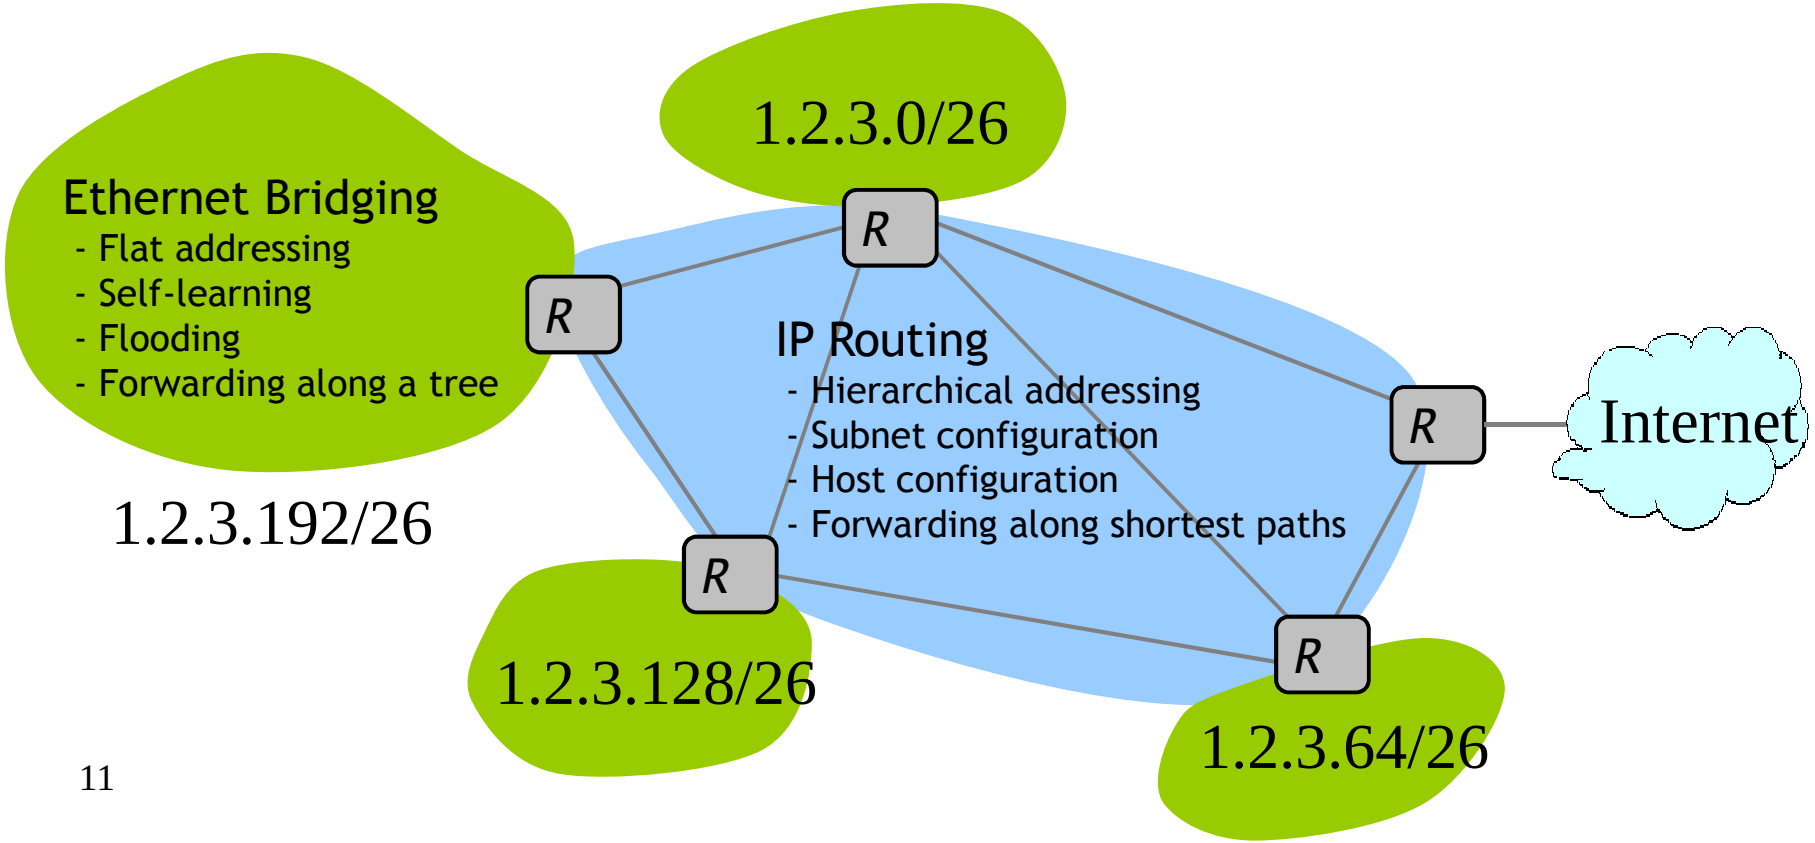

Virtual Local Area Networks (VLANs)

Evolution Toward Virtual LANs

- In the olden days…
  - Thick cables snaked through cable ducts in buildings
  - Every computer they passed was plugged in
  - All people in adjacent offices were put on the same LAN
  - Independent of whether they belonged together or not
- More recently…
  - Hubs and switches changed all that
  - Every office connected to central wiring closets
  - Often multiple LANs (*k* hubs) connected by switches
  - Flexibility in mapping offices to different LANs

Group users based on organizational structure, rather than the physical layout of the building.

Why Group by Organizational Structure?

- Privacy
  - Ethernet is a shared media
  - Any interface card can be put into "promiscuous" mode
  - … and get a copy of any flooded/broadcast traffic
  - So, isolating traffic on separate LANs improves privacy
- Load
  - Some LAN segments are more heavily used than others
  - E.g., researchers running experiments get out of hand
  - … can saturate their own segment and not the others
  - Plus, there may be natural locality of communication
  - E.g., traffic between people in the same research group

People Move, and Roles Change

- Organizational changes are frequent
  - E.g., faculty office becomes a grad-student office
  - E.g., graduate student becomes a faculty member
- Physical rewiring is a major pain
  - Requires unplugging the cable from one port
  - … and plugging it into another
  - … and hoping the cable is long enough to reach
  - … and hoping you don't make a mistake
- Would like to "rewire" the building in software
  - The resulting concept is a Virtual LAN (VLAN)

Example: Two Virtual LANs

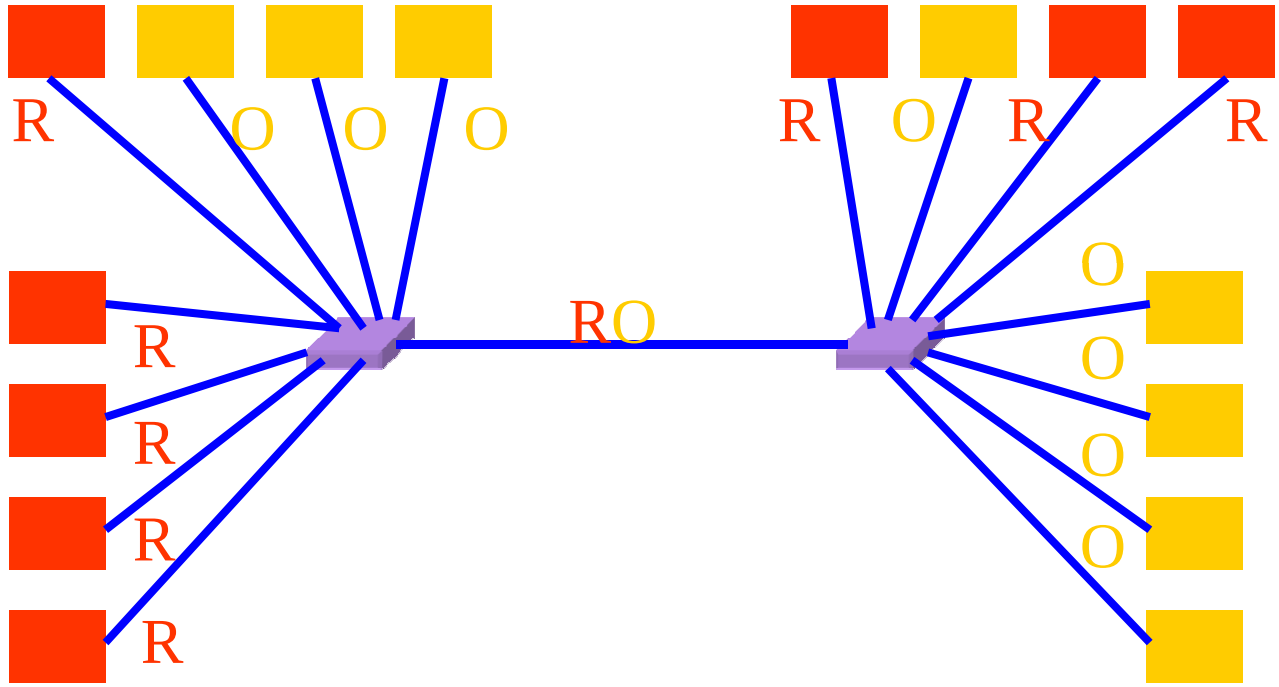

Red VLAN and Orange VLAN

Switches forward traffic as needed

Making VLANs Work

- Changing the Ethernet header
  - Adding a field for a VLAN tag
  - Implemented on the bridges/switches
  - … but can still interoperate with old Ethernet cards
- Bridges/switches trunk links
  - Saying which VLANs are accessible via which interfaces
- Approaches to mapping access links to VLANs
  - Each interface has a VLAN role

Multi-Homing

Motivation for Multi-Homing

- Benefits of multi-homing
  - Extra reliability, e.g., survive single ISP failure
  - Financial leverage through competition
  - Better performance by selecting better path
  - Gaming the 95th-percentile billing model

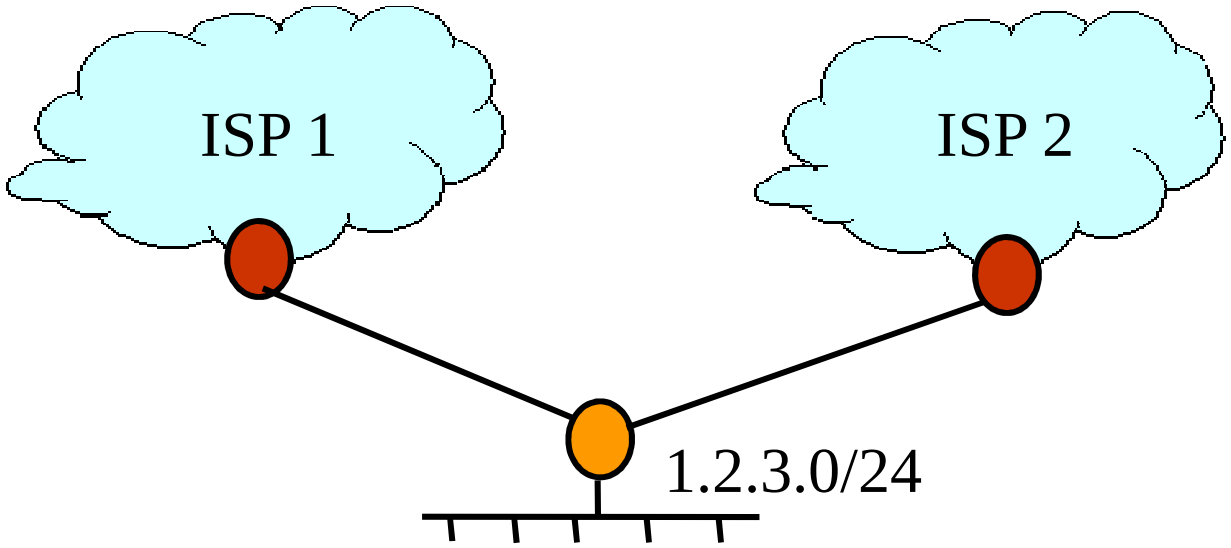

Multi-Homing Without BGP

Inbound Traffic

- Ask each ISP to originate the IP prefix
- … to rest of the Internet

Outbound Traffic

- One ISP as a primary, the other as a backup
- Or simple load balancing of all traffic

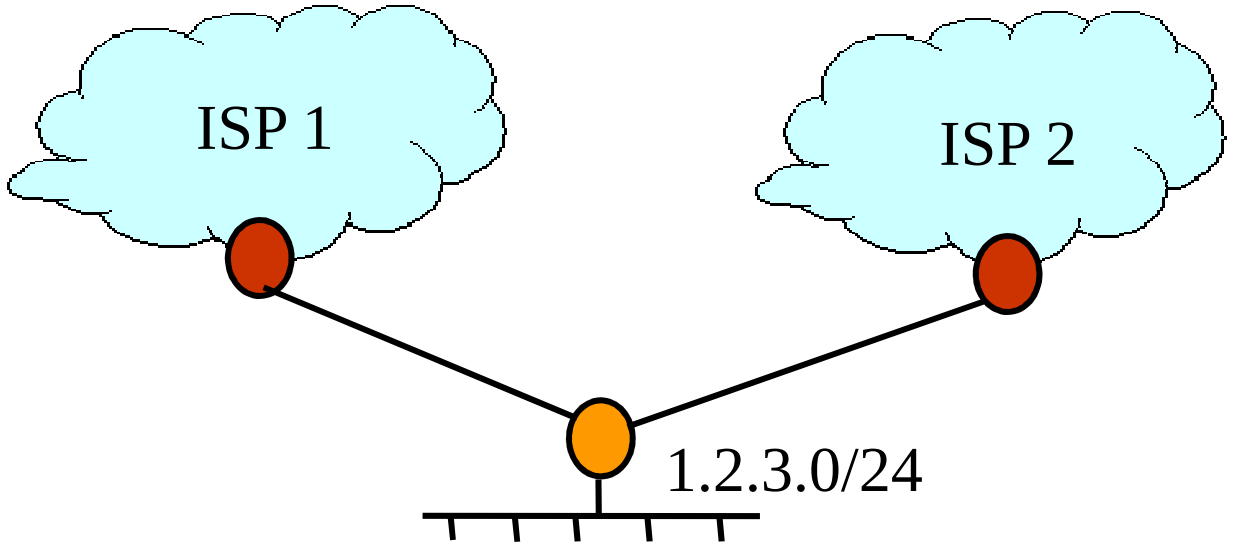

Multi-Homing With BGP

- Inbound traffic
  - Originate the prefix to both providers
  - Do *not* allow traffic from one ISP to another
- Outbound traffic
  - Select the "best" route for each remote prefix
  - Define BGP policies based on load, performance, cost

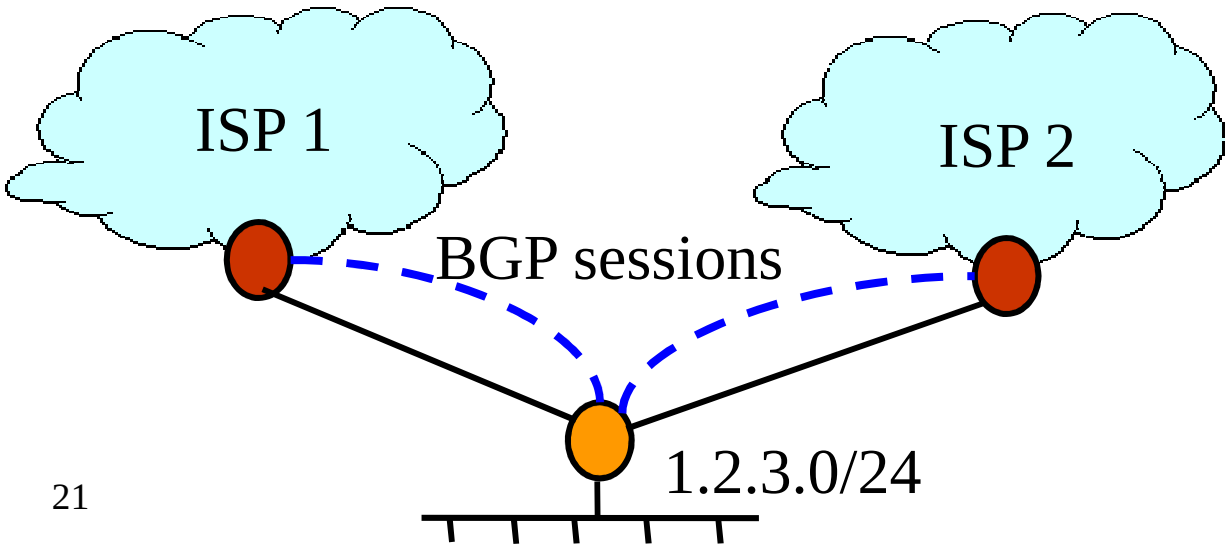

"Intelligent route control" or "multi-homed traffic

Interconnecting Multiple Enterprise Sites

Challenges

- Challenges of interconnecting multiple sites
  - Performance
  - Reliability
  - Security
  - Privacy
- Solutions
  - Connecting via the Internet using secure tunnels
  - Virtual Private Network (VPN) service

Connecting Via the Internet

- Each site connects to the Internet
  - Encrypted tunnel between each pair of sites
  - Packet filtering to block unwanted traffic
  - But, no performance or reliability guarantees

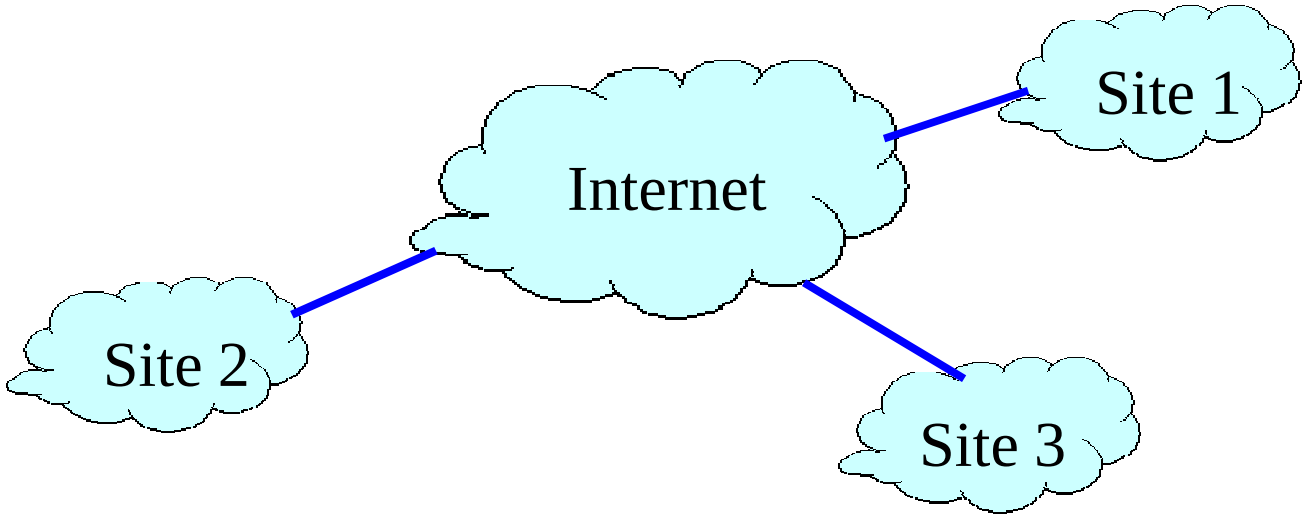

Virtual Private Network (VPN)

- Each site connects to a common VPN provider
  - Provider allows each site to announce IP prefixes
  - Separate routing/forwarding table for each customer
  - Performance guarantees by overprovisioning resources

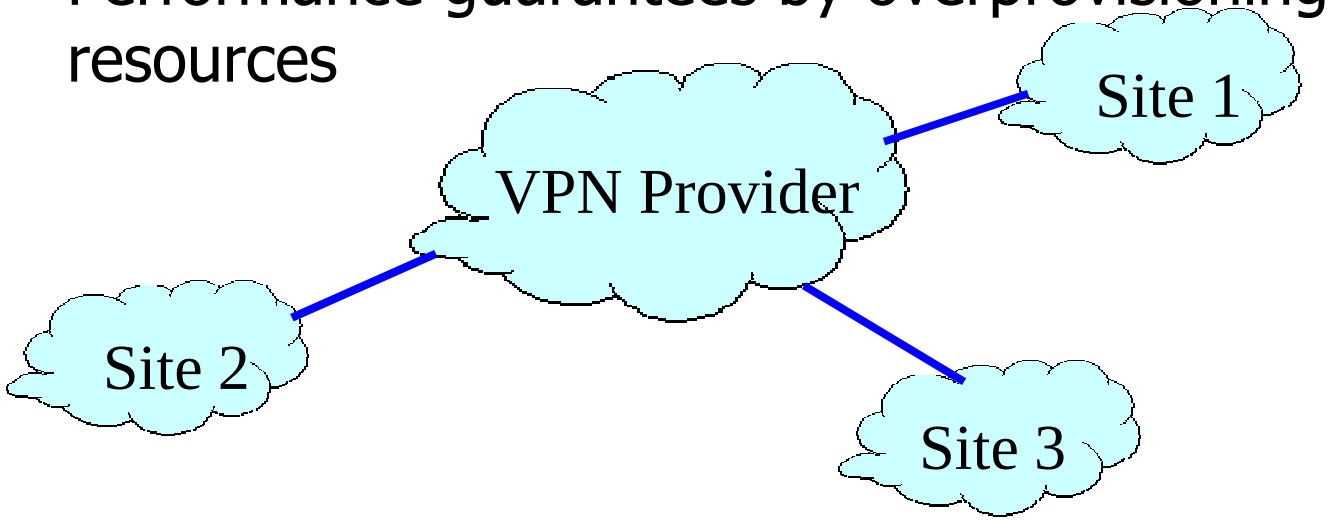

Conclusions

- Simple enterprise network is (mostly) plug and play
  - Ethernet with MAC learning and spanning tree
  - DHCP server to assign IP addresses from single subnet
  - Gateway router with default route to the Internet
- Quickly starts to require configuration
  - Choosing the root bridge in the spanning tree
  - Consistent configuration of DHCP and IP routers
  - VLAN access and trunk link configuration
  - Access control for traffic between VLANs
  - BGP sessions and routing policy
- Discussion of the two papers

Discussion

- Flat vs. hierarchical addressing?
- Roles of the end host vs. the network?
- How to best support flexible policies?
- Alternatives or extensions to VLANs?