ICS 351: Today's plan

- routing in general

- routing protocol comparison

- encapsulation

- network dynamics

- multicasting in general

- IP multicasting

- IGMP

- PIM

Why routing matters

- the routing protocol builds the routing tables, which are essential to correct routing of packets

- incorrect routing tables can lead to packets:

- o being dropped

- o being sent in a loop

- o being sent over slow links

- o being sent over congested links

- all of these cause network "malfunction", even when the hardware is working well!

routing responsibilities

- establish routes to destination networks

- maintain routes in the face of changing configuration: link loss, router loss, new links, new routers

- be trustworthy, do not advertise routes to destinations we can't reach

routing possibilities

- blacklist attacking (DoS) hosts/networks as close to the source as possible

- route depending on packet (flow) type, e.g. low latency, high throughput, constant bit rate (CBR)

- smart routing, based on packet content (may conflict with net neutrality)

what routing is not:

- Ethernet switching does not use IP addresses in any way, and only connects Ethernet segments that use the same IP network number
- Network Address Translation (NAT) allows a single IP address to be used as a "front end" for a number of systems that use TCP/IP or UDP/IP
- A firewall blocks access to most TCP/UDP ports, only allowing selected ports to connect to or from the outside
- each or all of these may be combined in the same box with a routing function, but they are logically separate

Comparison of Routing Protocols

- IGPs vs. EGPs: within an Autonomous System (RIP, OSPF), or among different Autonomous Systems (BGP)
- routing algorithms:
-    o distance vector: RIP (also AODV for wireless networks)
-    o path vector (modified distance vector): BGP
-    o link-state: OSPF (also OLSR for wireless networks)
- link-state distributes the state of the router's links
- distance-vector distributes the routing table distance to each destination -- if that is optimal, the sender is used as the direction (vector) in which to send the data
- OSPF and RIP are optimal, BGP finds acceptable routes

Protocol Encapsulation

- protocols are layered

- packets in one layer are encapsulated within packets in the lower layer

- for example, a UDP header is added to a RIP packet, to support directing to a specific UDP port number

- this UDP packet gets an IP header which includes the destination address (224.0.0.9) and the protocol number 17 (UDP)

- the IP packet receives an Ethernet header (and a CRC "trailer") with the local ethernet address as source, and ff:ff:ff:ff:ff:ff as destination

- in this case, we say that RIP is layered over UDP, UDP over IP, and IP over Ethernet

- what protocol is OSPF layered over? BGP? Ping? Telnet?
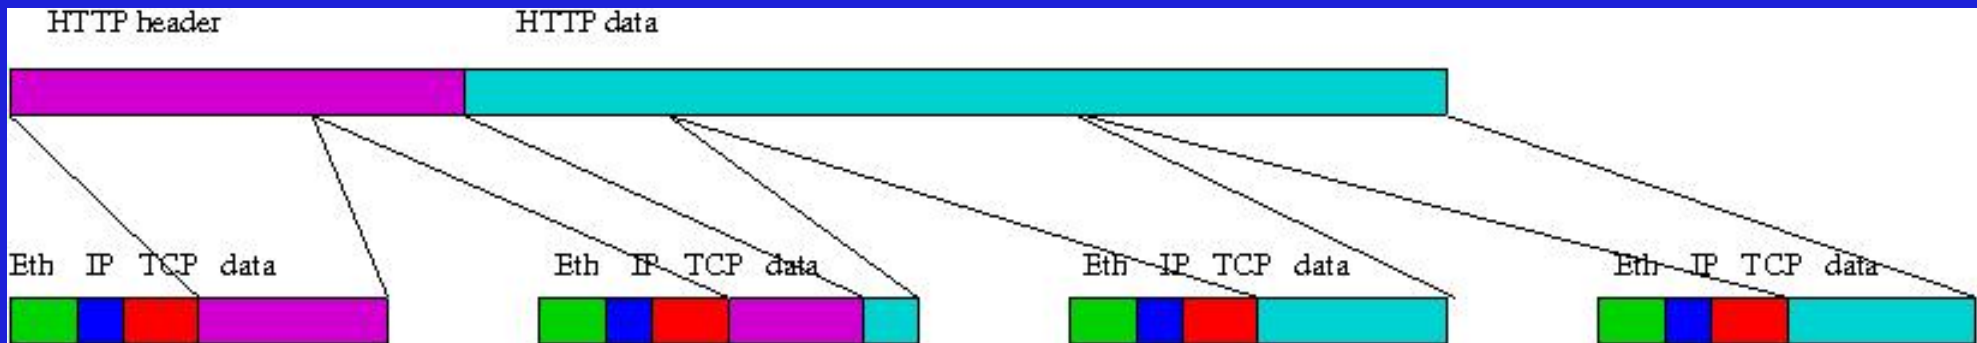
Network dynamics

- a wired network can be modeled of as a graph

- a router or host is a node in the graph

- the links between them are the edges of the graph

- however, this model does not quite capture all the interesting features, for example, a local area network may connect many computers

- when adding a router or a link, all the routing protocols take some time before all the routing nodes recognize the new resource

- what happens during this time? In RIP, in OSPF, in BGP?

- when removing a router or a link, again the routing tables are inconsistent for some time

- what happens during this time?

Multicasting

- if a single sender needs to reach a group of receivers
- or if multiple senders need to reach a group of receivers
- or if multiple computers need to reach each other
- in general, a sender should send each packet (stream) exactly once
- intermediate systems (routers) should duplicate the packet (stream) once for each outgoing link
- overall, this forms a **multicast tree**

Multicast Trees

- The multicast tree can be built starting from the sender towards the receivers (leaves), along the route that unicast packets would take, or

- the tree can be built by each receiver along the reverse path towards the sender, until a router that already carries the multicast data is found. This is **reverse path forwarding**

- for multicast with multiple senders, the tree can be seen as a **spanning tree** with no real root, and a standard **flooding algorithm** can be used to distribute the data throughout this overlay network

- alternately, the tree can have a designated root, a **rendezvous point** or **RP**

Multicast Trees
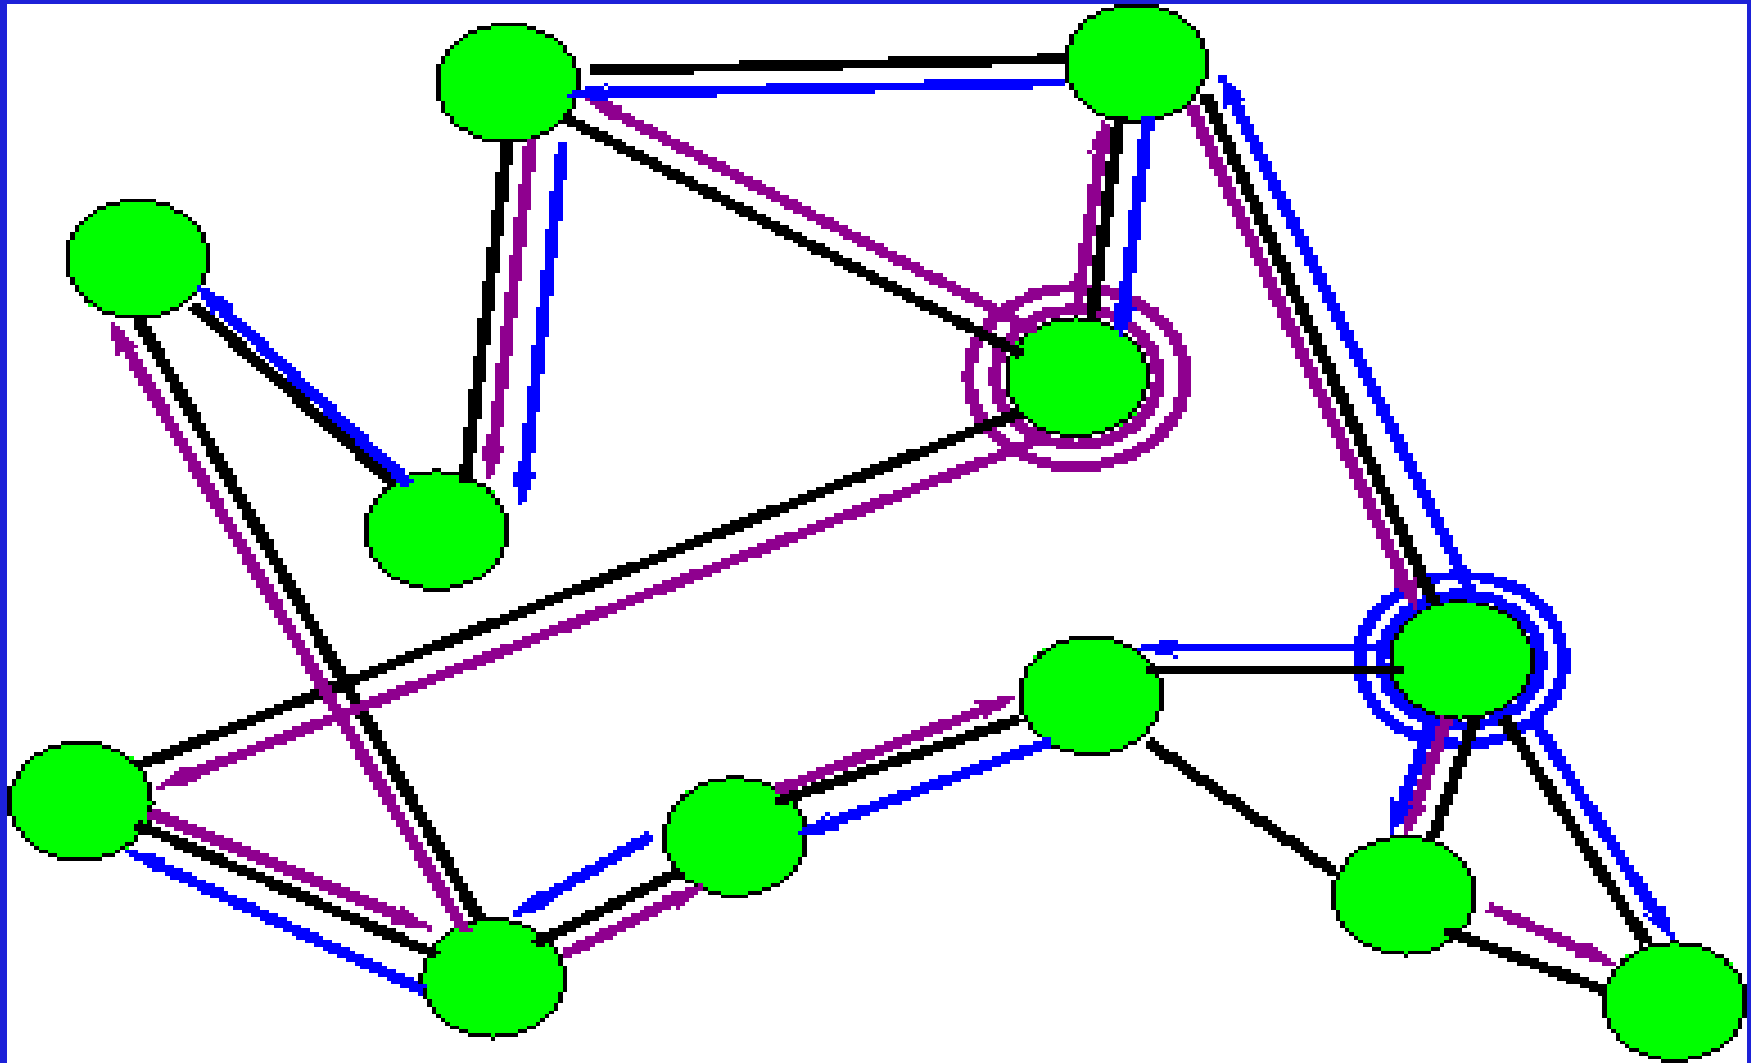
IP Multicast Addresses

- IP multicast addresses are in class D, beginning with 224 through 239
  - 1. the first byte for class A addresses is 0 through 127
  - 2. the first byte for class B addresses is 128 through 191
  - 3. the first byte for class C addresses is 192 through 223
  - 4. the first byte for class D addresses is 224 through 239
  - 5. the first byte for class E addresses is 240 through 255
- for example, 224.0.0.9 for RIP packets, 224.0.0.5 for OSPF packets

IP Multicasting

- IGMP manages group membership in multicast groups within local networks (MLD does the same on IPv6 networks)

- PIM (or MOSPF) are the equivalent of routing protocols for multicast, providing multicast routing when the multicast router is not local

IGMP

- Internet Group Management Protocol version 3
- used to communicate between a multicast router and local multicast hosts
- the router needs to know which hosts require which multicast stream(s), so as to only forward streams that are needed
- a host requesting a stream from its router results in the router recording this information
- this request expires if not refreshed often enough: soft state in the router
- messages sent over IP (protocol number 2) with TTL 1
- IGMP routers send Membership Queries, IGMP hosts send Membership Reports
- RFC 3376

PIM

- Protocol Independent Multicast is protocol-independent in the sense that it relies on a not-otherwise defined routing protocol to find routes

- PIM dense mode (PIM-DM), RFC 3973

- in dense mode, multicast data is sent to all routers except those that send prune messages

- dense mode is only used within an autonomous system (with MSDP used to allow multicast among autonomous systems)

- PIM sparse mode (PIM-SM), RFC 4601

- in sparse mode, multicast data is broadcast over a tree rooted at a designated router called the **Rendezvouz Point (RP)**

- also PIM Source-Specific Multicast (PIM-SSM) and Bidirectional PIM (BIDIR-PIM), a variant of PIM-SM

IGMP

- Internet Group Management Protocol version 3
- used to communicate between a multicast router and local multicast hosts
- the router needs to know which hosts require which multicast stream(s), so as to only forward streams that are needed
- a host requesting a stream from its router results in the router recording this information
- this request expires if not refreshed often enough: **soft state** in the router
- messages sent over IP (protocol number 2) with TTL 1
- IGMP routers send Membership Queries, IGMP hosts send Membership Reports
- RFC 3376

MBone

- for a while, there was a generic multicasting infrastructure called the MBONE (multicast backbone)

- the 6-bone was a similar infrastructure for IPv6 traffic

- the MBONE was a collection of multicast routers willing to carry multicast traffic and to run multicast routing protocols

- a host that was not directly connected to a multicast router could register with a remote MBONE router and exchange packets using unicast IP