CS 204: Multicast

Jiasi Chen

Lectures: MWF 12:10-1pm in WCH 139

http://www.cs.ucr.edu/~jiasi/teaching/cs204_spring16/

Overview

- Basics
- Routing algorithms
  - Flooding
  - Spanning trees
  - RPB
  - TRPB
  - RPM
- Implementations of routing algorithms
  - MOSPF
  - DVMRP
  - PIM
- Paper discussion

What is multicast?

- Group of hosts requesting same content
- What would unicast do?

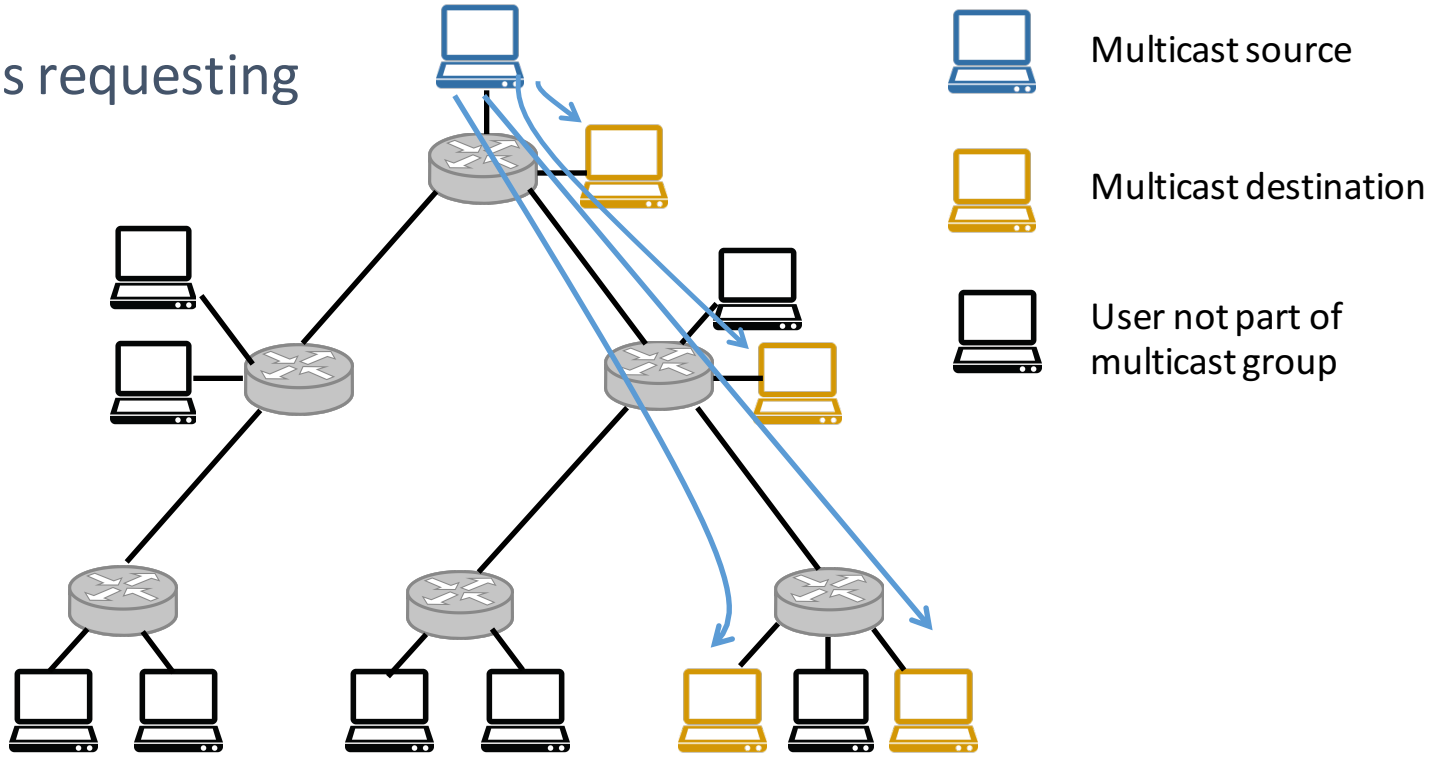

What is Multicast?

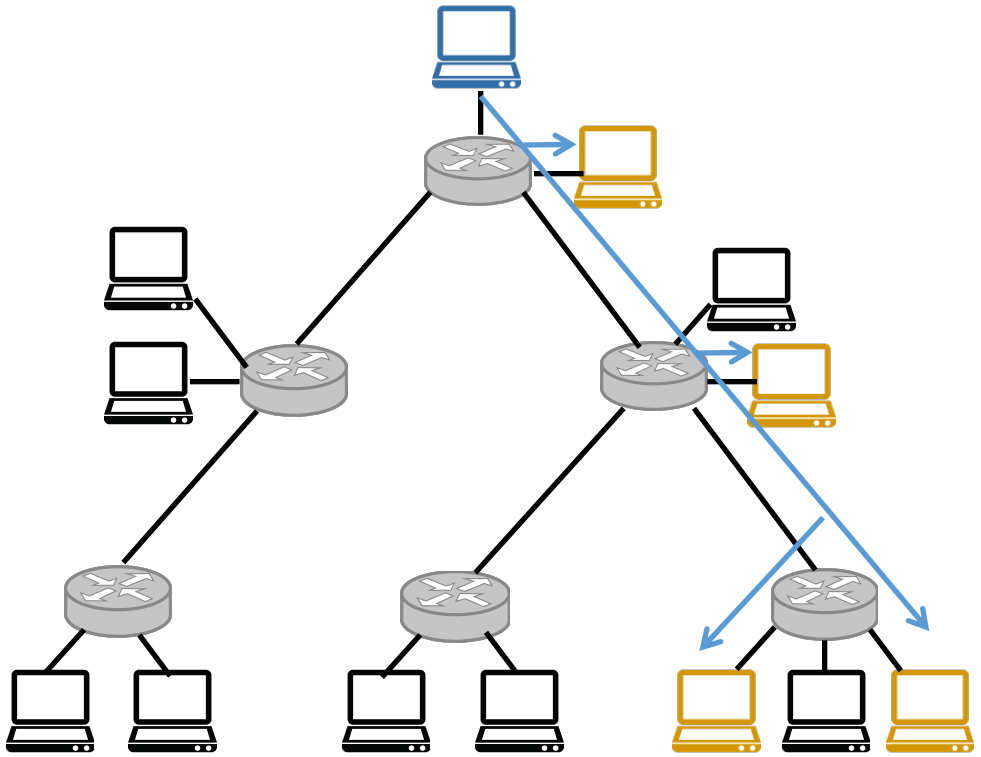

Multicast source

Multicast destination

User not part of multicast group

What is Multicast?

- One-to-many routing
- Main goal: efficiency
- Example applications
  - Audio/video
  - Software distribution
  - Web-cache updates
  - Teleconferencing
  - Games
- Job of the router
  - Know which groups its hosts are subscribed to
  - Forward packets to hosts
  - Forward packets to other routers

Multicast Service Model

- Anyone can join
- Sender need not be part of the multicast group
- Members can join and leave at will
- Group membership is not explicitly known

- Analogy: radio channel

Multicast Addressing

- IP address (32 bits)
  - Class D: 224.0.1.0 to 239.255.255.255

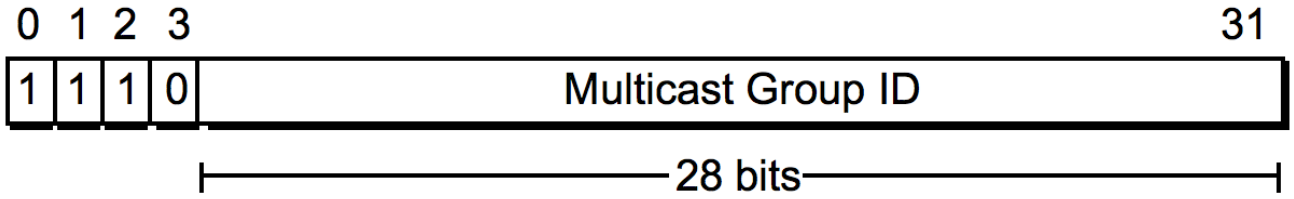

- MAC address (48 bits)
  - 01-00-5E-xx-xx-xx

Converting Multicast IP to MAC Address

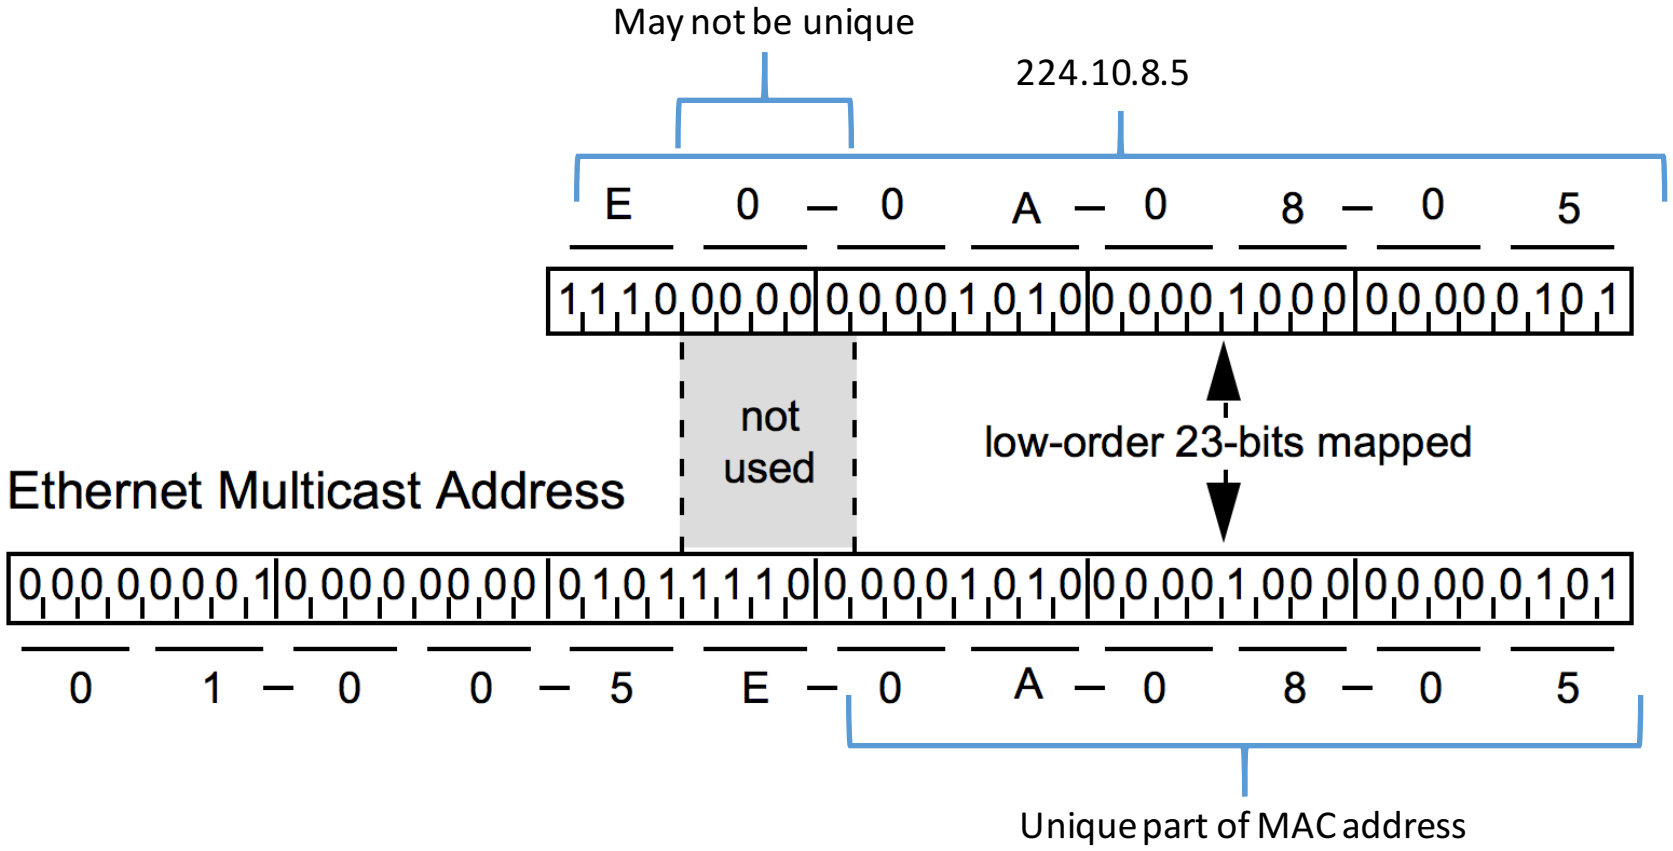

Limiting the Scope of Multicast Packets

- Each interface assigned a TTL
- IP header also contains TTL
- Forward packet iff packet TTL > interface TTL

Initial TTL	Scope
0	Restricted to the same host
1	Restricted to the same subnetwork
32	Restricted to the same site
64	Restricted to the same region
128	Restricted to the same continent
255	Unrestricted in scope

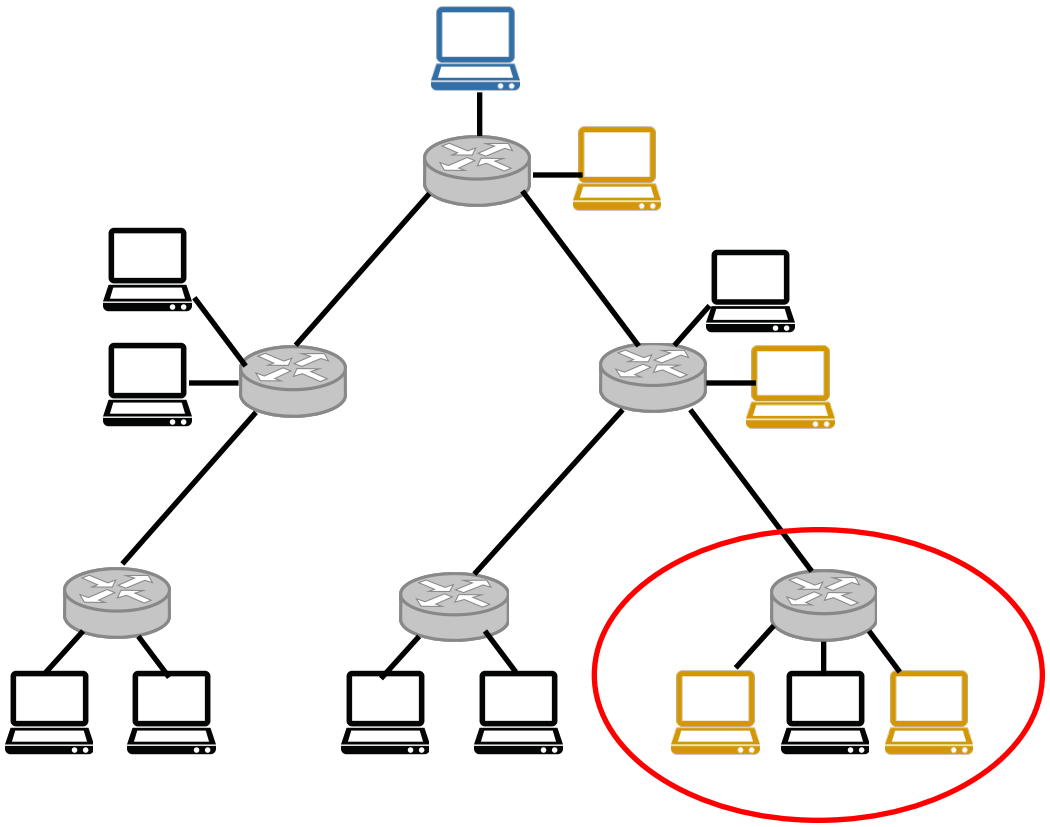

How do hosts talk to routers?

Multicast API

- Sender
  - Same as before: send to multicast IP
- Receiver
  - Need to join the multicast group by sending message to router
  - Join-IP-Multicast-Group
  - Leave-IP-Multicast-Group

IGMP

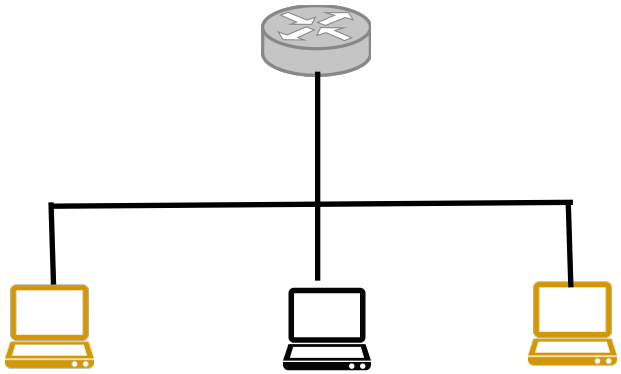

Router

1. QUERY 224.0.0.1, TTL=1

"Which groups are you part of?"

Each client: "I am part of group *g*"

2. For each group *g* I am part of, set a random timer T_g
3. If I hear another report for group *g*, reset T_g
4. When T_g expires, send my report

Overview

- Basics
- Routing algorithms
  - Flooding
  - Spanning trees
  - RPB
  - TRPB
  - RPM
- Implementations of routing algorithms
  - MOSPF
  - DVMRP
  - PIM
- Paper discussion

Classification

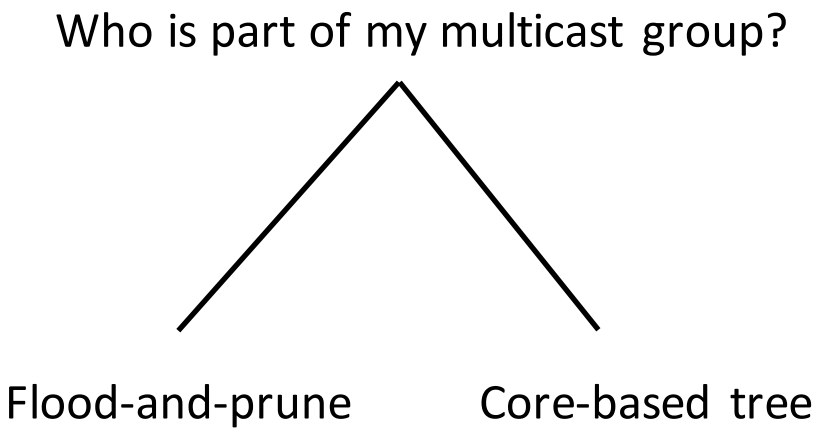

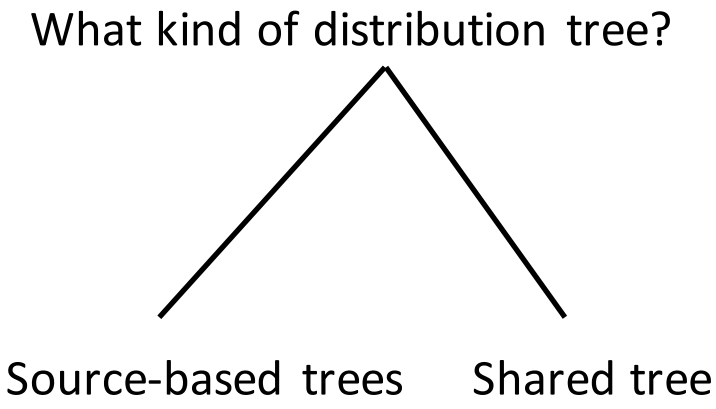

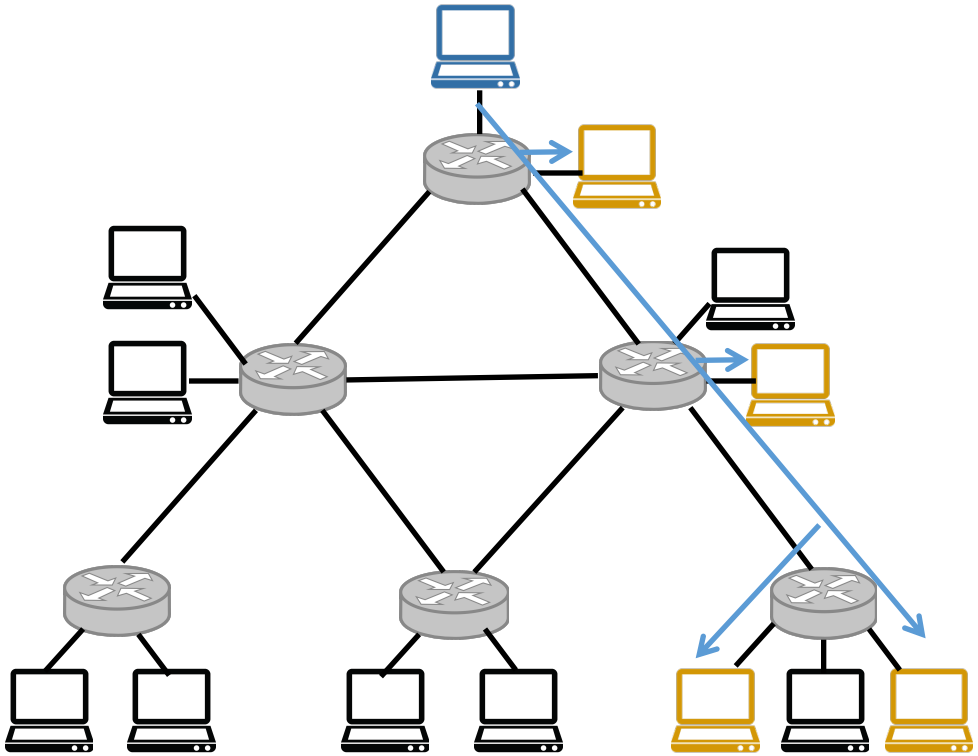

How to route the packets?

Flooding

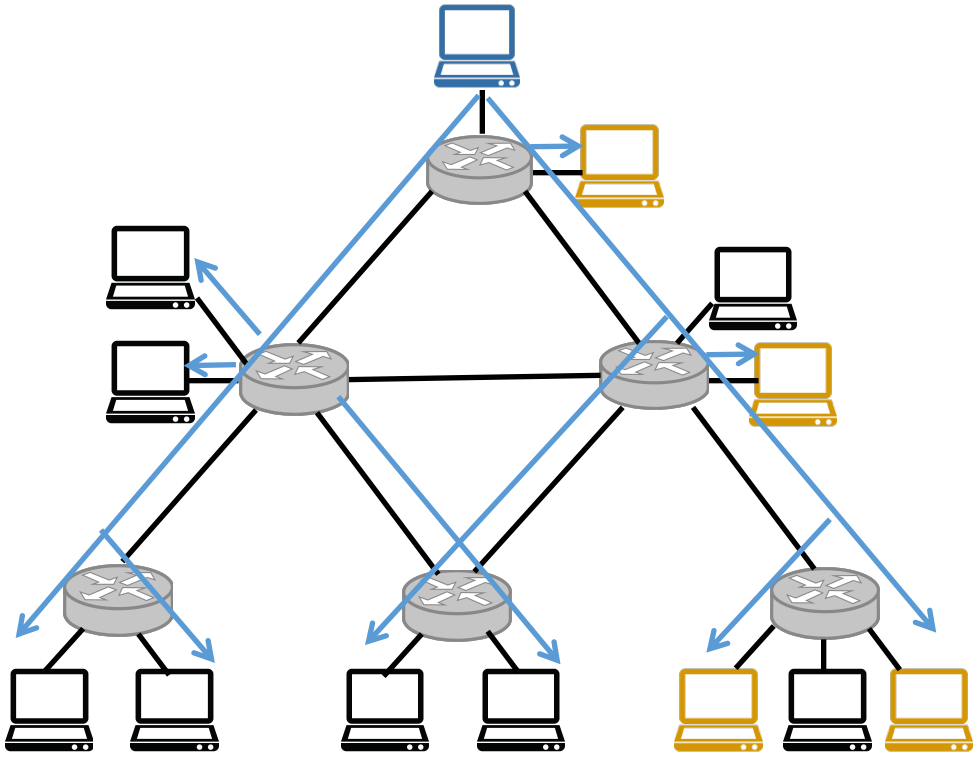

1. Forward the packets
   - On all links except receiving
   - If haven't seen this packet before

Benefits
- Easy to implement

Disadvantages
- Not scalable

Internet-Wide Spanning Trees

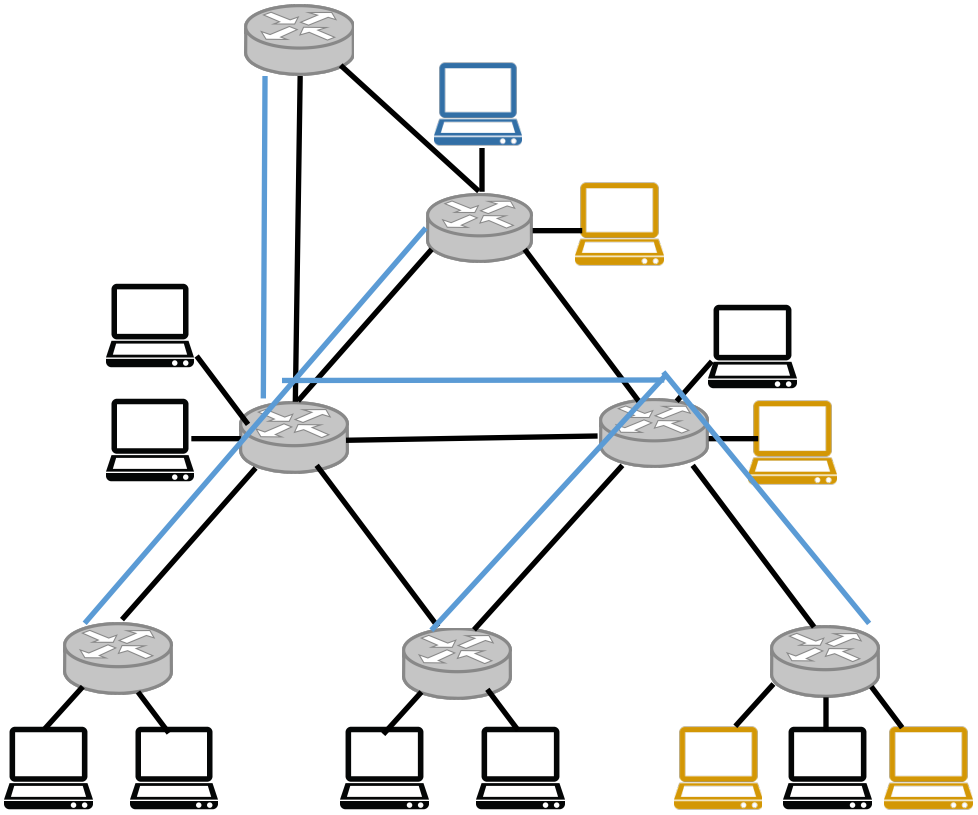

1. Construct a spanning tree
2. Forward the packets
   - on the links of the spanning tree

Advantages
- Spanning tree algorithms are well-known

Disadvantages
- Not the most efficient path
- Concentrates traffic on a few number of links
- Needs the entire Internet topology!

RPB: Source-Specific Spanning Tree

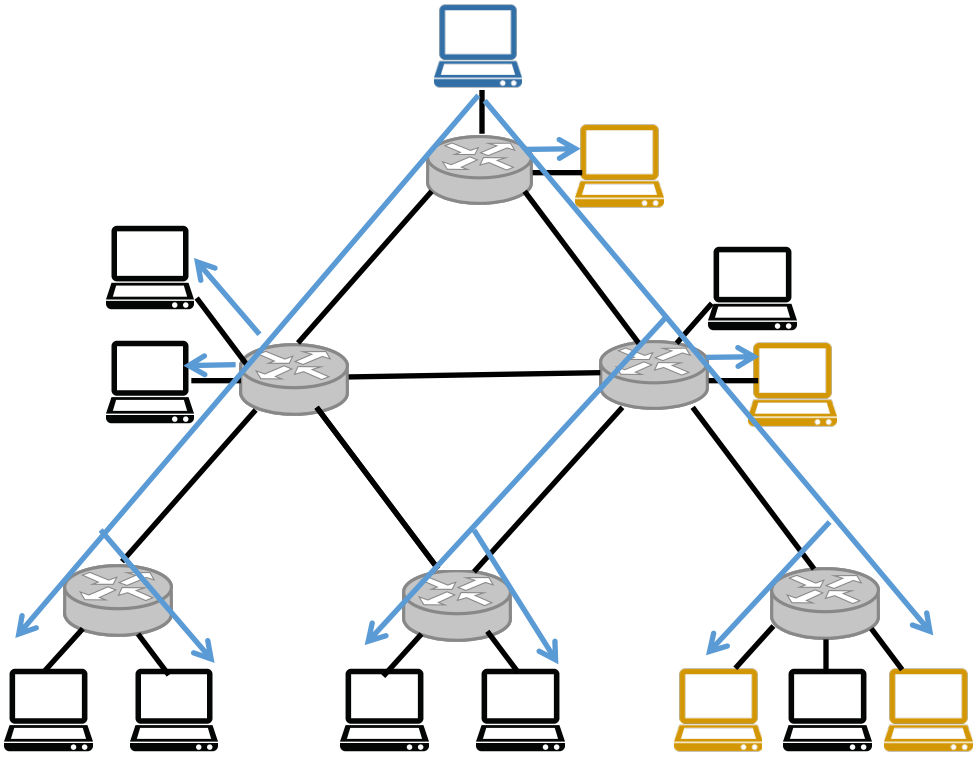

1. Forward packets
   a) If received packet on my shortest path to source
   b) to all "downstream" routers
   c) to hosts on your subnet

Benefits
- Distribute traffic over links because construct a new tree for each source
- Shortest path

Disadvantages
- Forwards packets to routers not connected to multicast groups

How to determine downstream routers?

- Downstream = next router considers me as part of their shortest path to source

- Unicast: route packets **towards destination**
- Multicast: route packets **away from source**
  - Call this reverse path fowarding

- Link-state: already have topology
- Distance vector: need to advertise last hop to neighbors

TRPB: Add in IGMP

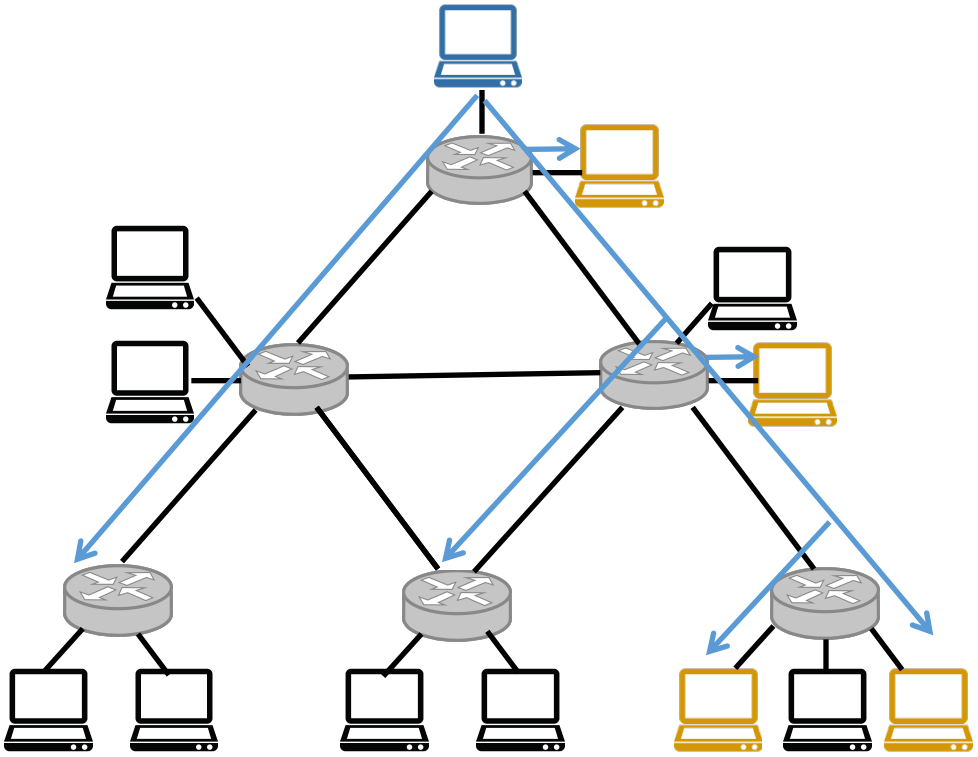

1. Forward packets
   a) If received packet on my shortest path to source
   b) to all "downstream" routers
   c) to hosts on your subnet if they are part of the multicast group

RPM

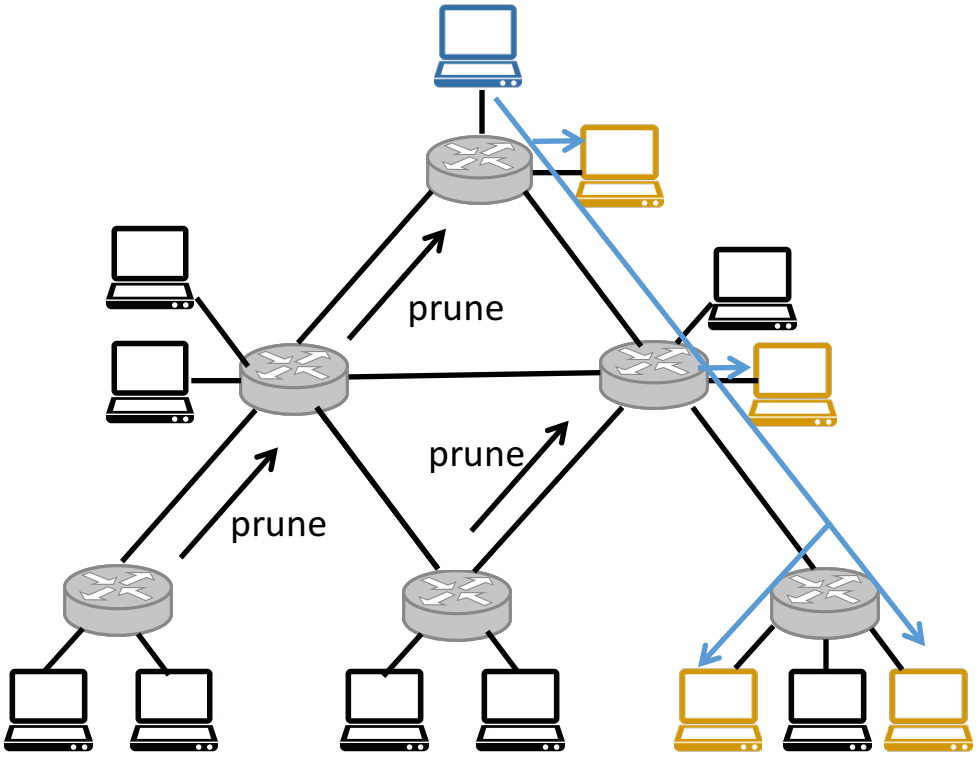

1. Forward packets
   a) If received packet on my shortest path to source
   b) to all "downstream" routers
   c) to hosts on your subnet if they are part of the multicast group
   d) If a "prune" message hasn't been received

Benefits

- Reduce unnecessary traffic in subnets and between routers

Disadvantages

- Periodically, packets sent to all multicast routers

Overview

- Basics
- Routing algorithms
  - Flooding
  - Spanning trees
  - RPB
  - TRPB
  - RPM
- Implementations of routing algorithms
  - MOSPF
  - DVMRP
  - PIM
- Paper discussion

MOSPF

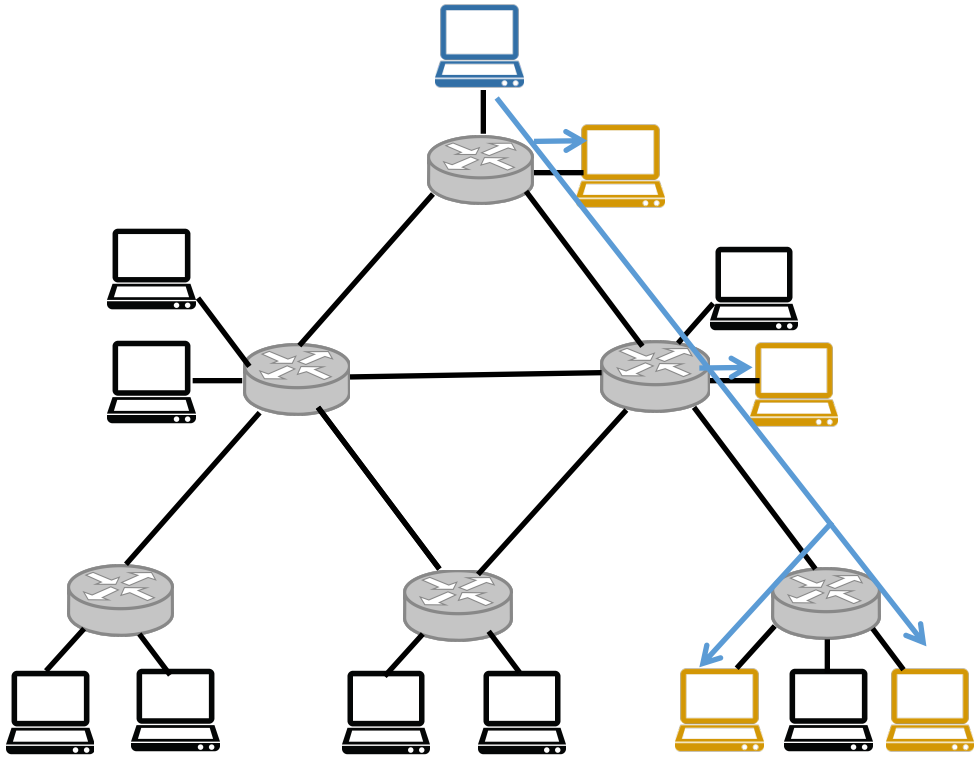

Link state based

- Modify OSPF
- Compute shortest path between source and set of destinations
- Periodically flood with neighbor information

Disadvantage

- Need to re-compute entire shortest path if user joins/leaves

DVMRP

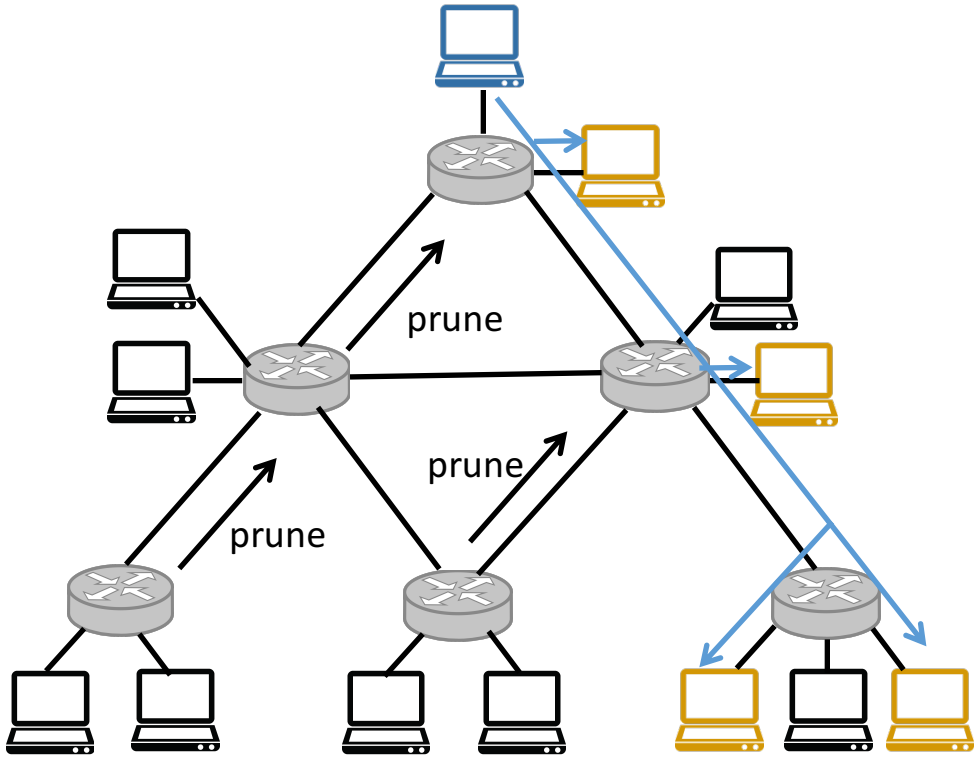

Distance-vector based

- Pass messages with (dest, cost) to neighbors
- If dest = multicast source, pass cost = infinity to upstream router

DVMRP details

- Routing table

Source Subnet	Subnet Mask	From Gateway	Metric	Status	TTL	InPort	OutPorts
128.1.0.0	255.255.0.0	128.7.5.2	3	Up	200	1	2,3
128.2.0.0	255.255.0.0	128.7.5.2	5	Up	150	2	1
128.3.0.0	255.255.0.0	128.6.3.1	2	Up	150	2	1,3
128.4.0.0	255.255.0.0	128.6.3.1	4	Up	200	1	2

- Forwarding table

Source Subnet	Multicast Group	TTL	InPort	OutPorts
128.1.0.0	224.1.1.1	200	1 Pr	2p 3p
	224.2.2.2	100	1	2p 3
	224.3.3.3	250	1	2
128.2.0.0	224.1.1.1	150	2	2p 3

p = prune message sent/received

Protocol Independent Multicast (PIM)

- Why maintain a separate multicast routing table?

- Look at unicast routing table
  - If unicast path (destination = multicast source)
- Agnostic to which unicast routing algorithm is used

- **Dense** mode: assume everyone is part of multicast group, and explicitly remove
- **Sparse** mode: assume nobody is part of multicast group, and explicitly join

PIM-Dense

- Similar to DVMRP
- Reverse-path-forwarding
  - Instead of distance vector, use unicast routing table

PIM-Sparse

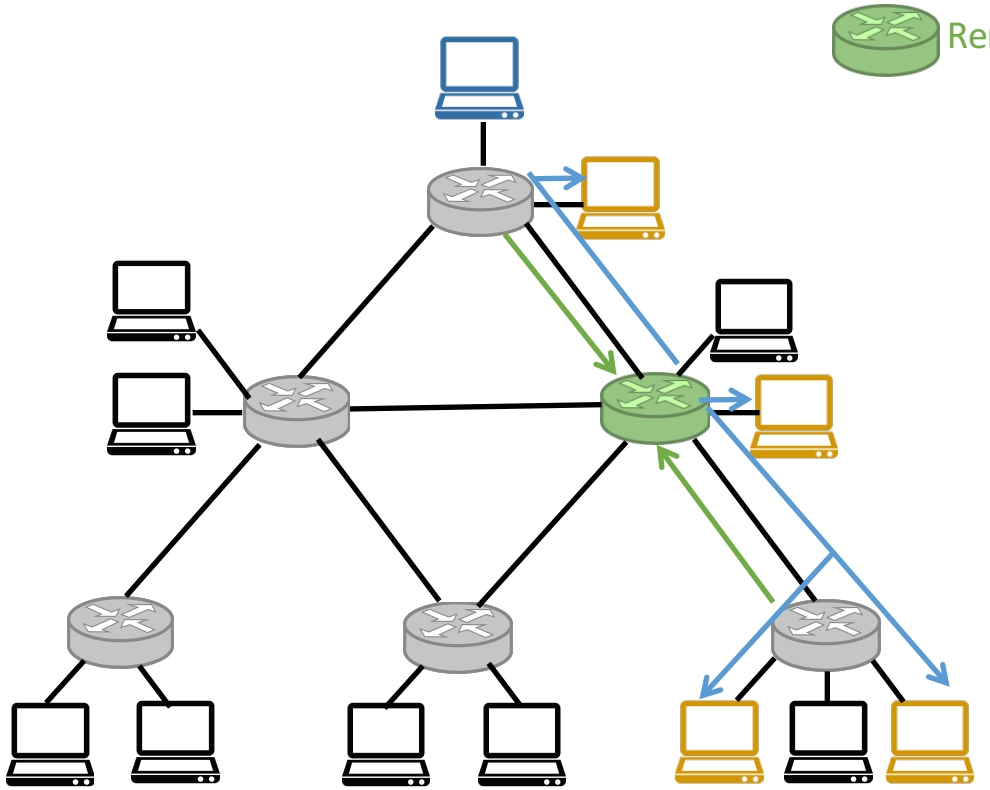

Center-based tree
- RP administratively configured

Source
1. Register with RP
2. Send packets to RP

Receiver
1. Send join message to RP

RP
1. Create the (*,group) tree

Summary

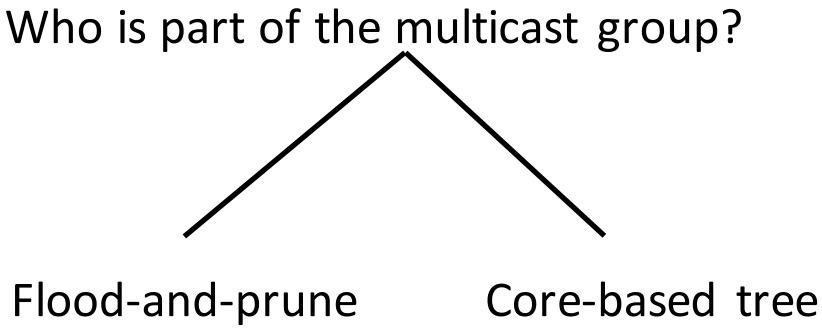

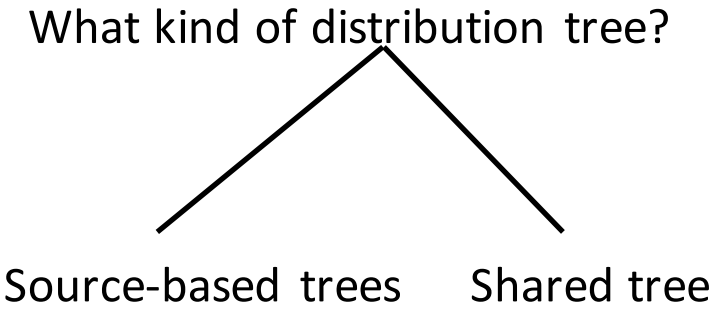

Algorithm	Flood-and-prune or center-based?	Source-based or shared tree?
Flooding/spanning tree	Flood	Shared
RPB	Flood	Source
RPM	Flood-and-prune	Source
DVMRP	Flood-and-prune	Source
PIM-Dense	Flood-and-prune	Source
PIM-Sparse	Center-based	Mostly shared, can be source

Tunneling

Q: how to connect "islands" of multicast routers in a "sea" of unicast routers?

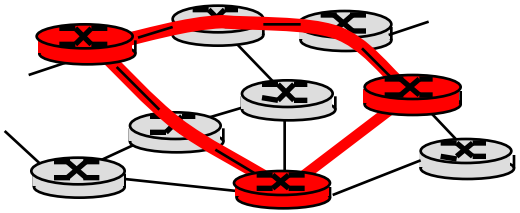

physical topology

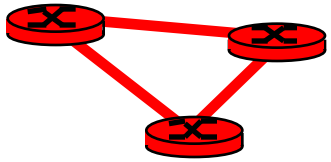

logical topology

- mcast datagram encapsulated inside "normal" (non-multicast-addressed) datagram
- normal IP datagram sent thru "tunnel" via regular IP unicast to receiving mcast router (recall IPv6 inside IPv4 tunneling)
- receiving mcast router unencapsulates to get mcast datagram

Who Uses IP Multicast?

- Testbeds
  - MBONE (DVMRP)
  - Internet2
- Live video CDNs?
- Wireless?

A Case for End System Multicast

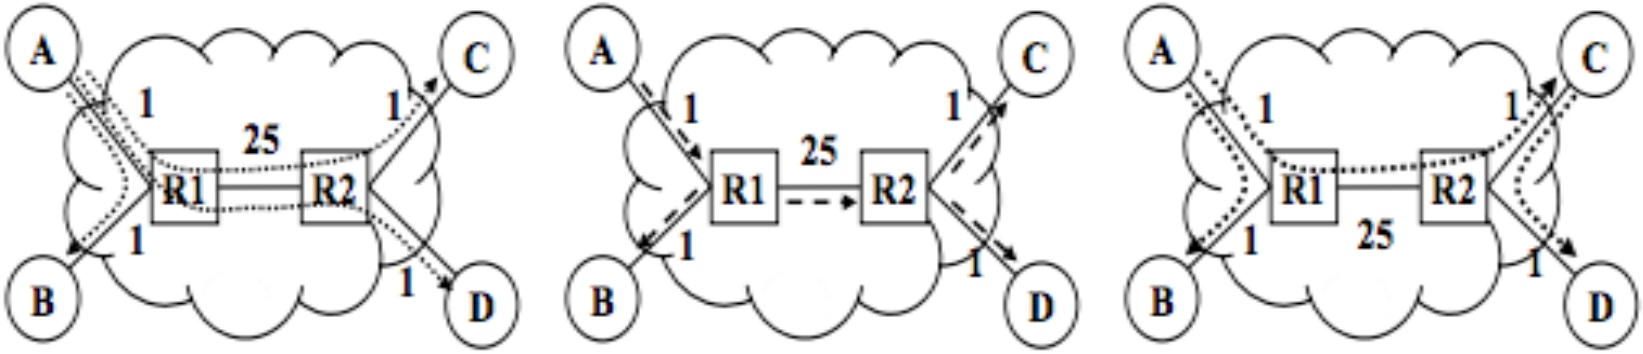

Unicast

Network multicast

End system multicast

Paper Discussion

- Mesh vs source-based tree vs shared tree?
- How did they test?
- What are the drawbacks?

Sources

- *Computer Networking: A Top-Down Approach*, Kurose & Ross
- "A Case for End System Multicast", Yang-hua Chen, Sanjay Rao, Hui Zhang, *SIGMETRICS*, 2000.
- "Introduction to IP Multicast Routing," Chuck Semeria and Tom Maufer