Multicast Communications

Slide Set were original prepared by Dr. Tatsuya Susa

Outline

1. Advantages of multicast
2. Multicast addressing
3. Multicast Routing Protocols
4. Multicast in the Internet
5. IGMP

1. Multicasting

- Unicast
  - A flow from one source to one destination
  - IP packets contain destination IP address
- Broadcast
  - A flow from one source to all destinations
  - IP packets contain broadcast address 255.255
- Multicast
  - A flow from one source to a *Group* of destinations
  - IP packets contain a class D address for destination (if IPv4)

Multicast vs. Multiple Unicast

Multiple Unicast:

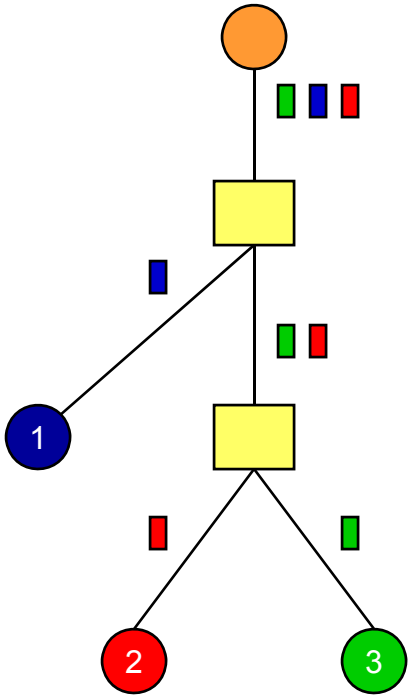

Multicast:

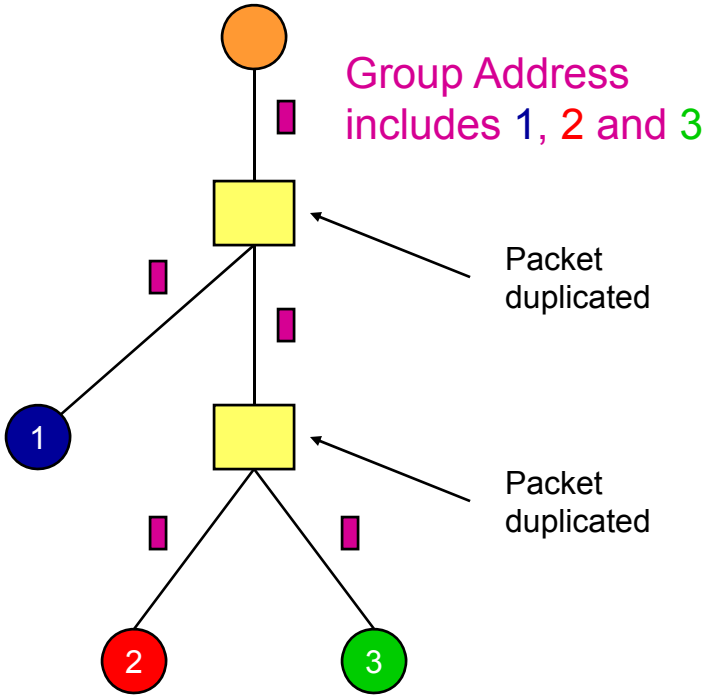

Advantages of Multicasting

- Advantages:
  - Lower overhead at the source
    - Source sends only one packet
  - Bandwidth is conserved on shared links
    - Only one copy of each packet is sent on each link
- Requirements:
  - Group address management
  - Packet duplication at routing nodes

Applications of Multicast

- Push data
  - Video streams, audio streams
- Resource discovery
  - Can send multicast queries (if destination is not known)
- Multimedia collaboration
  - Video conferencing

2. Multicast Addressing

IP Address Classes
(From IPv4 Lecture Notes)

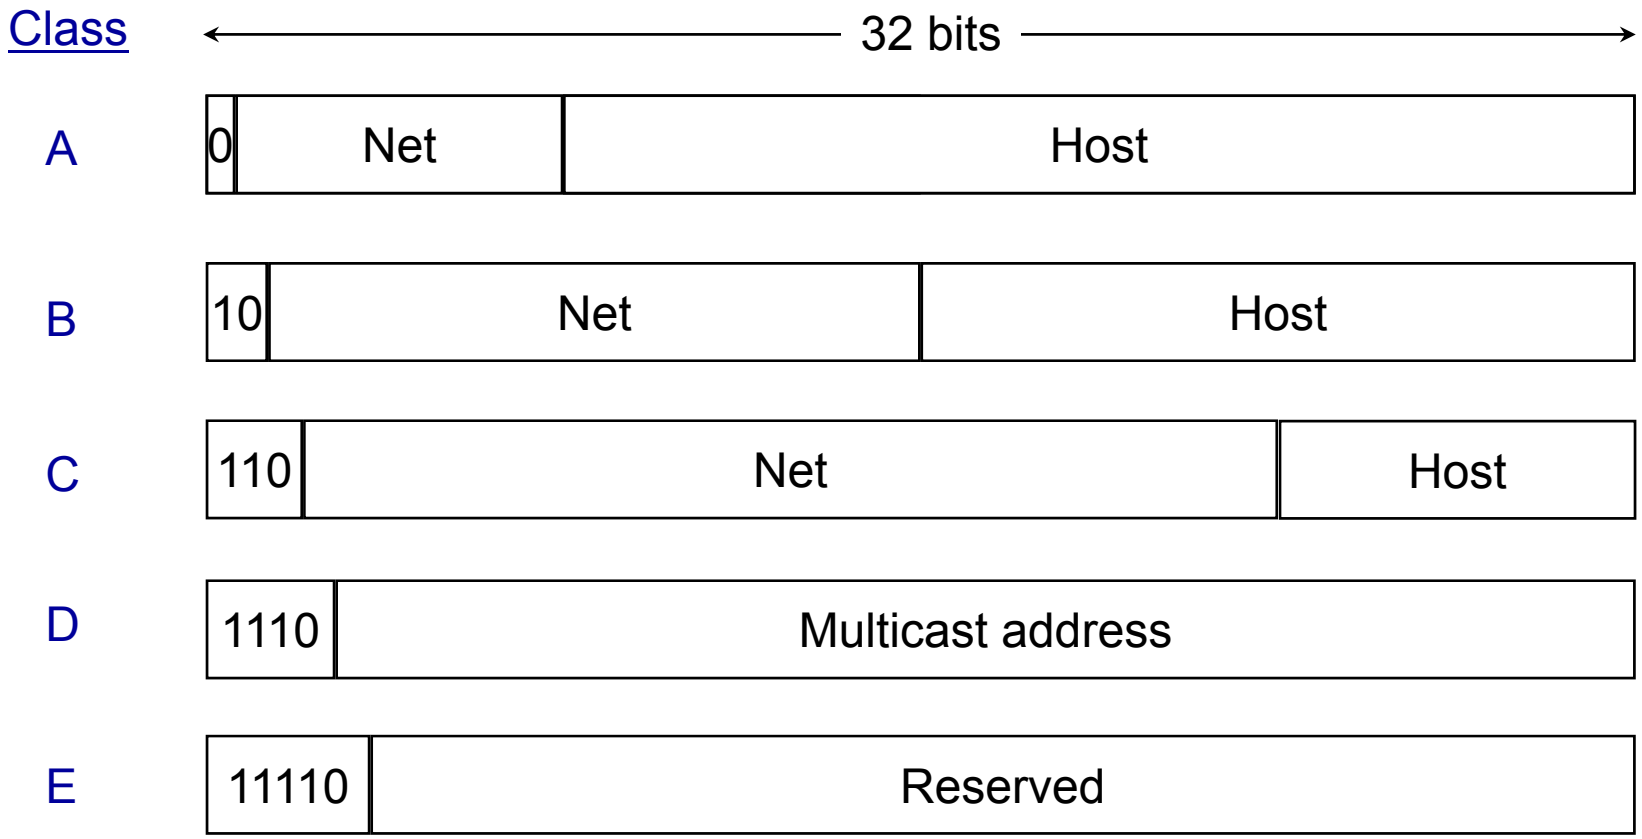

IP Address Classes
(From IPv4 Lecture Notes)

- Class A:
  - For very large organizations
  - 16 million hosts allowed
- Class B:
  - For large organizations
  - 65 thousand hosts allowed
- Class C
  - For small organizations
  - 255 hosts allowed
- Class D
  - Multicast addresses
  - No network/host hierarchy

IP Address Classes
(From IPv4 Lecture Notes)

- Class E
  - reserved
- Loopback
  - 127.xx.yy.zz (127.anything) is reserved for loopback testing
  - packets sent to this address are not put out onto the wire; they are processed locally and treated as incoming packets.
- Broadcast
  - all 1s

MAC Layer Ethernet Frame Format
(From MAC Layer Lecture Notes)

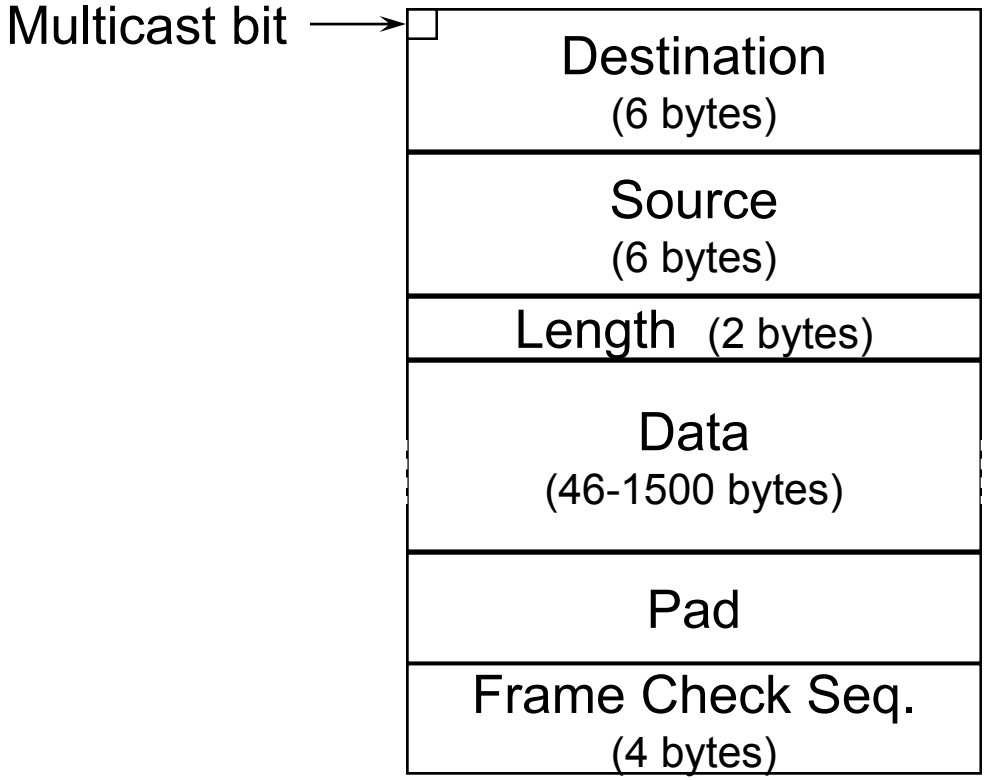

MAC Layer Ethernet Frame Format
(From MAC Layer Lecture Notes)

- Destination and Source Addresses:
  - 6 bytes each
- Two types of destination addresses
  - Physical address: Unique for each user
  - Multicast address: Group of users
  - First bit of address determines which type of address is being used

    0 = physical address

    1 = multicast address

Multicast Addressing

- Class D address

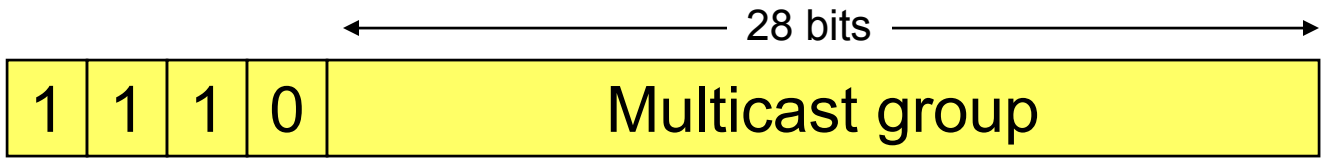

Multicast addresses are in the range 224.0.0.0 through 239.255.255.255

The set of hosts listening to a particular multicast address (i.e., multicast destinations) is called a *host group*

A sender of packets to a multicast group is known as a *multicast source*

Mapping to Ethernet Addresses

- Ethernet has a 48-bit address field
- It has its own multicast address range
  - 01.00.5e.00.00.00 through 01.00.5e.7f.ff.ff
  - Lower order 23 bits can be used for multicast addresses
- IP multicast address has 28 bits for specifying a group address
- Thus, only the lower order 23 bits of IP multicast address are copied into the Ethernet address

3. Multicast Routing Protocols

- Flooding
- Shared Spanning Tree
- Source-Based Spanning Trees
- Reverse Path Forwarding (RPF)
- Truncated Reverse Path Broadcast (TRPB)
- Reverse Path Multicasting (RPM)

3.1 Flooding

- Same algorithm as that for unicast routing
- A router copies a packet and transmits it on all outbound links (except the one the packet came in on)
- Routers keep a list of sequence numbers
  - If a packet with the same sequence number has already been seen, drop the packet

Flooding

A communicates
with group G: {B,C,D}

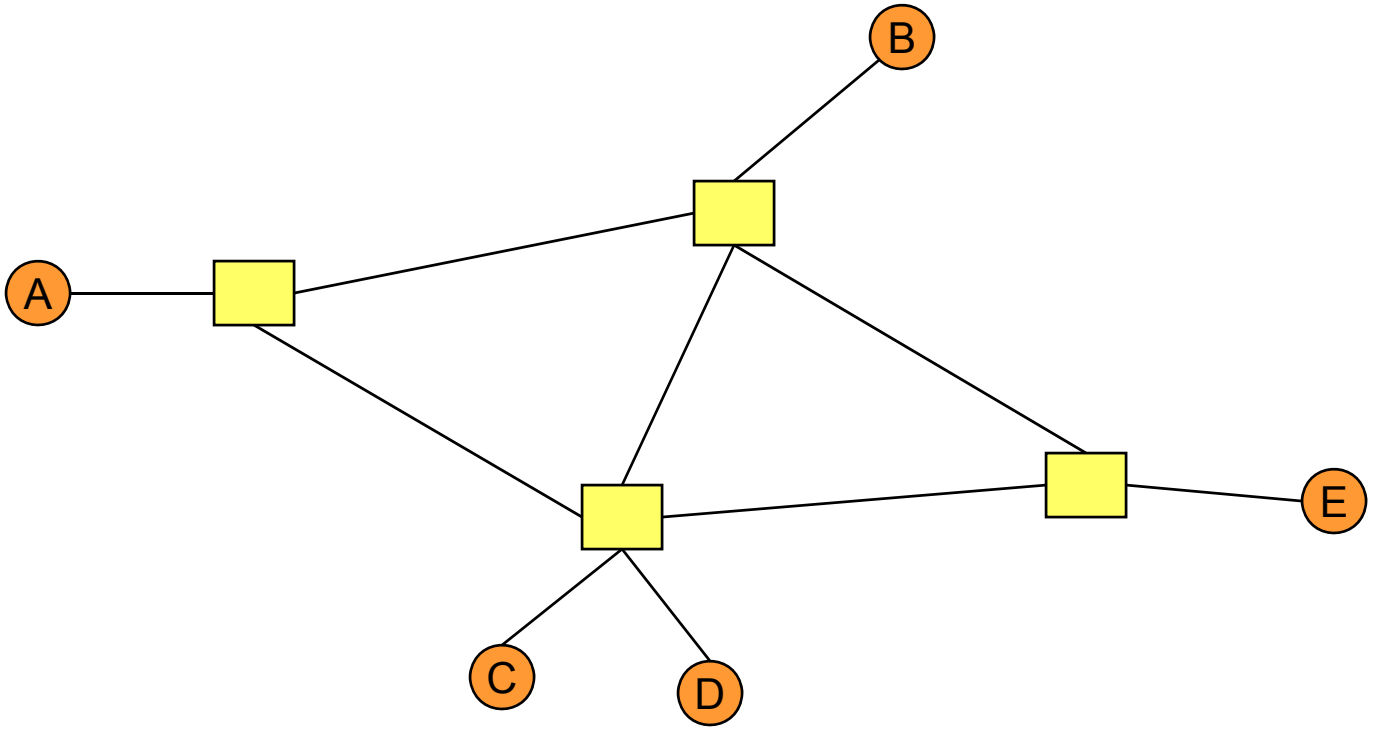

Flooding

A communicates
with group G: {B,C,D}

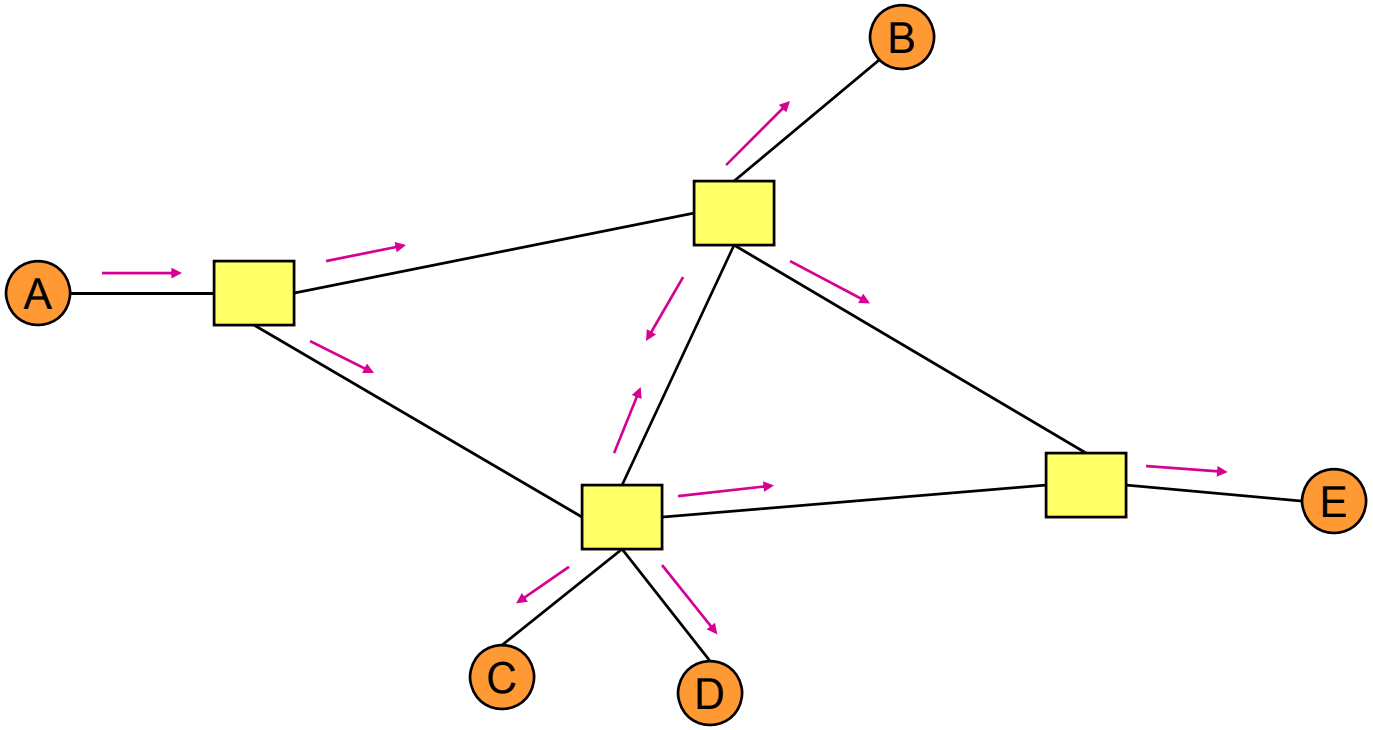

Flooding

Advantages and Disadvantages

- Advantages:
  - Simple to implement
  - No group management protocol required
- Disadvantages:
  - The router sequence number lists can grow very large, so this solution does not scale
  - Routers will still frequently receive the same packet more than once
  - Packets will end up going where they aren't wanted

3.2 Shared Spanning Tree

- First, build a spanning tree for the entire network
  - A tree that spans all the routers in the network
  - May be good within an Autonomous System
- Forwarding multicast packets:
  - A router forwards a packet on all links that are part of the spanning tree except the one on which it received the packet
  - No loops and therefore no duplicate packets

Shared Spanning Tree *(cont'd)*

A communicates
with group G: {B,C,D}

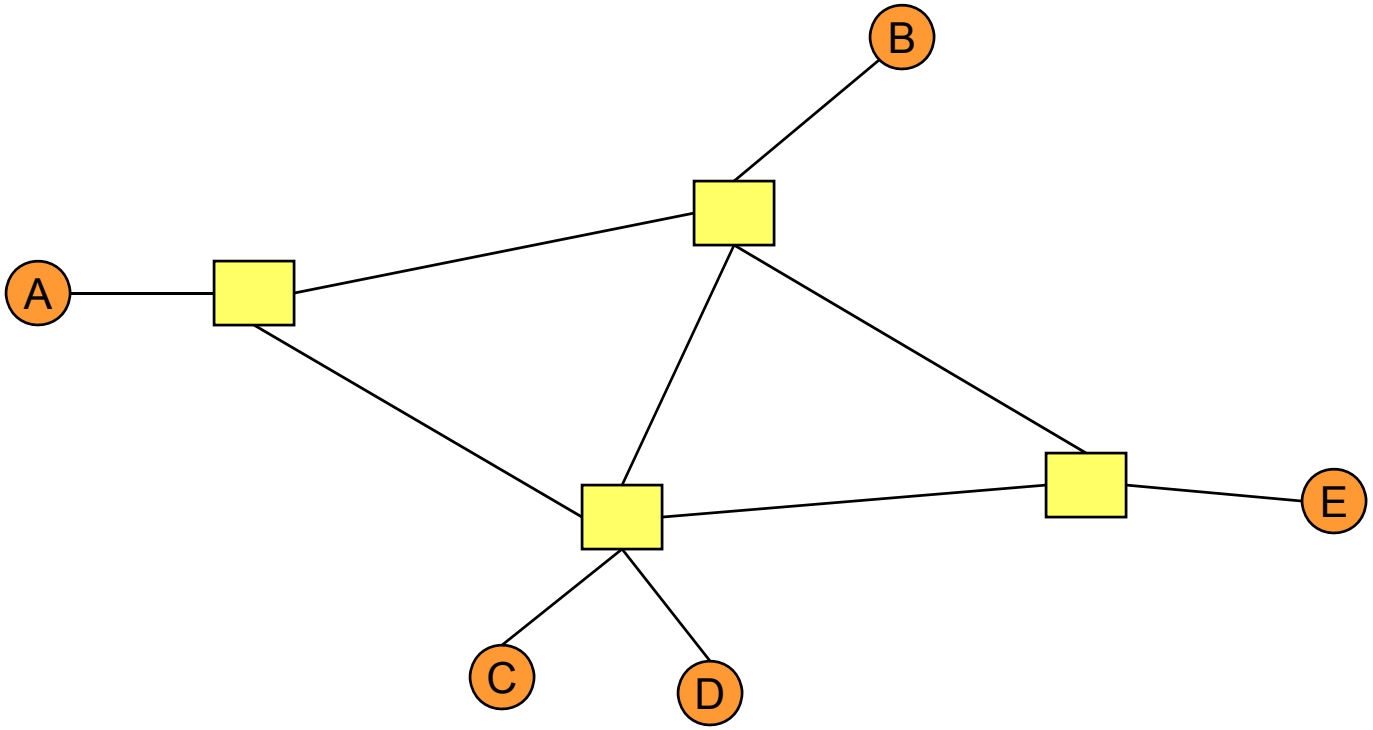

Shared Spanning Tree *(cont'd)*

A communicates
with group G: {B,C,D}

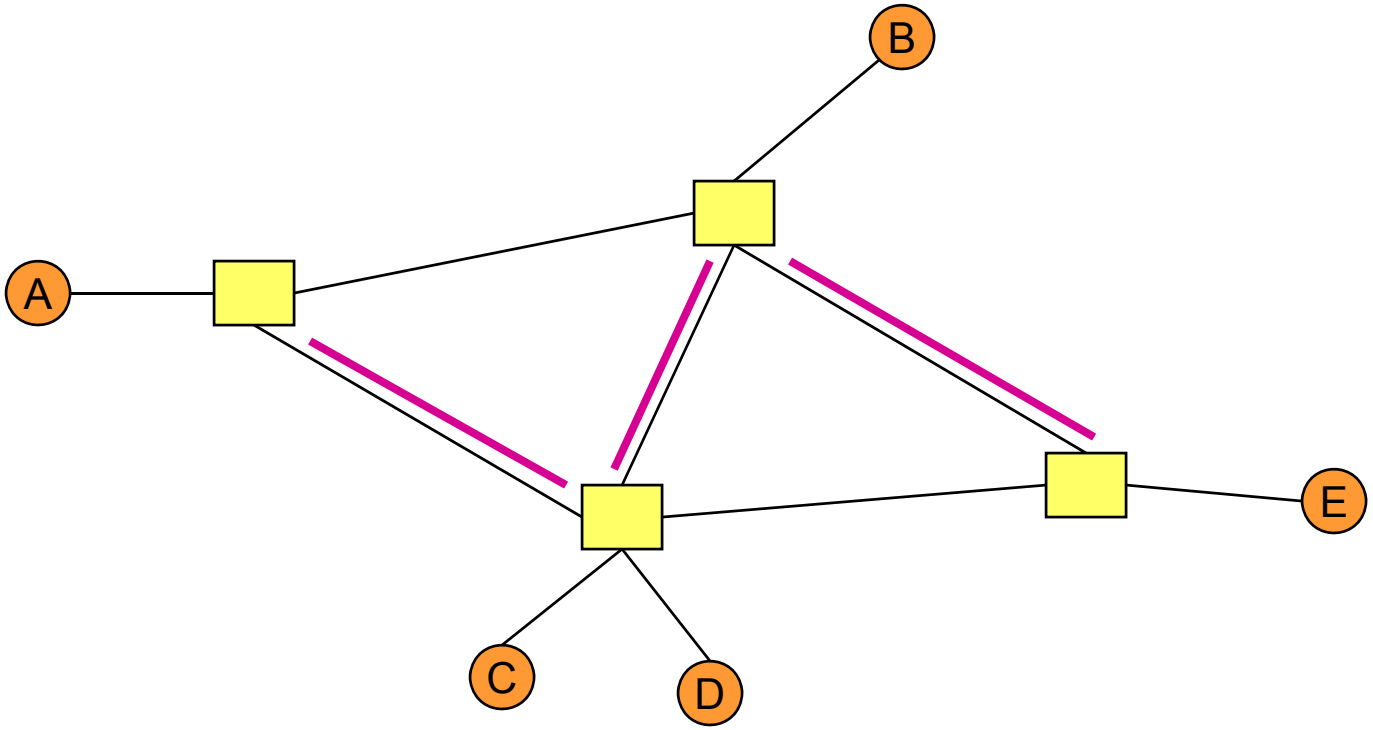

Spanning Tree Path

Shared Spanning Tree *(cont'd)*

A communicates
with group G: {B,C,D}

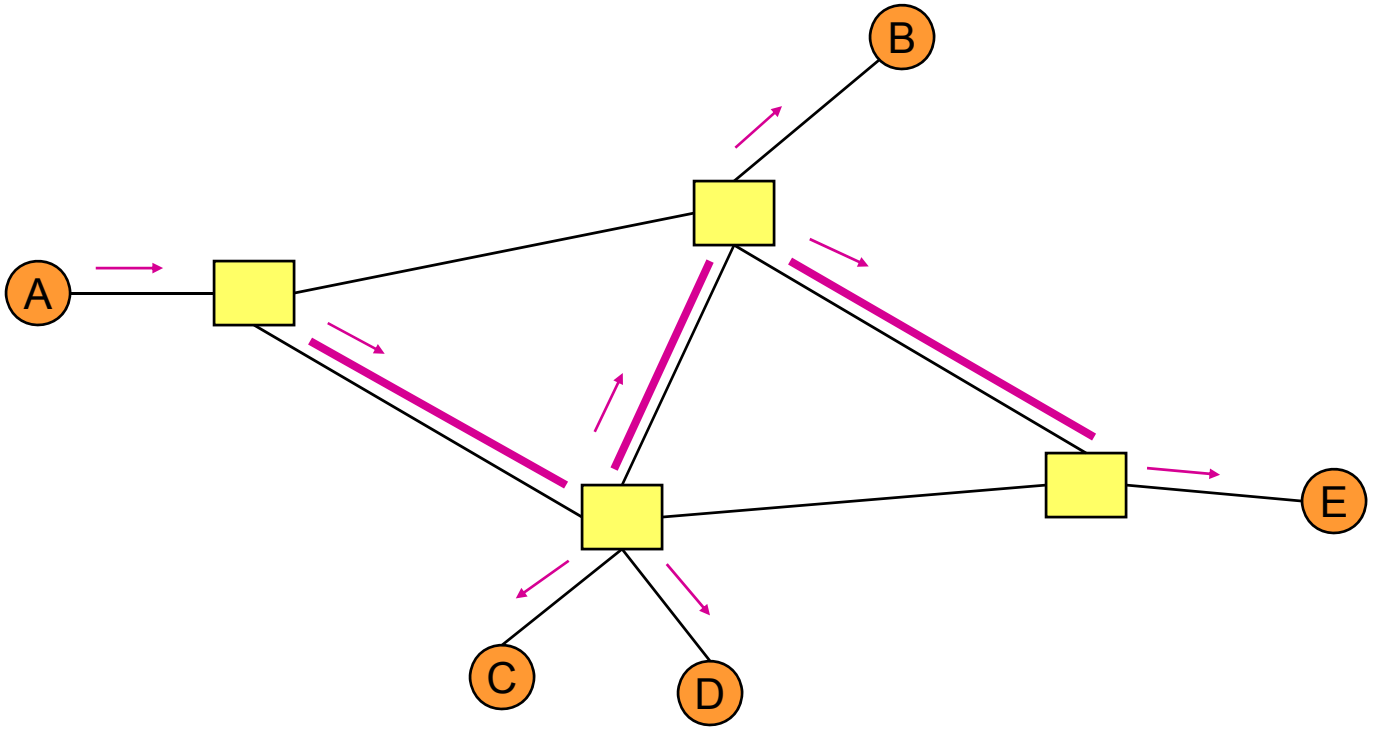

Spanning Tree Path

Shared Spanning Tree

Advantages and Disadvantages

- Advantages:
  - Can centralize traffic on a smaller number of links, so less network bandwidth is used
  - No more duplicate packets at routers
- Disadvantages:
  - The network needs to explicitly construct the shared tree
  - Shared spanning trees do not necessarily create the most efficient paths from the source to all group members
  - Spanning tree paths may become bottlenecks
  - Packets will still end up going where they aren’t wanted

3.3 Source Based Trees

- Instead of building one shared spanning tree for all multicast packets, use a separate spanning tree for each source
- Each source-based spanning tree is *explicitly* constructed using the shortest paths from the source to all other destinations

Source Based Trees *(cont'd)*

A communicates
with group G: {B,C,D}

C communicates
with group G': {A,B,D}

All edges have
equal weights

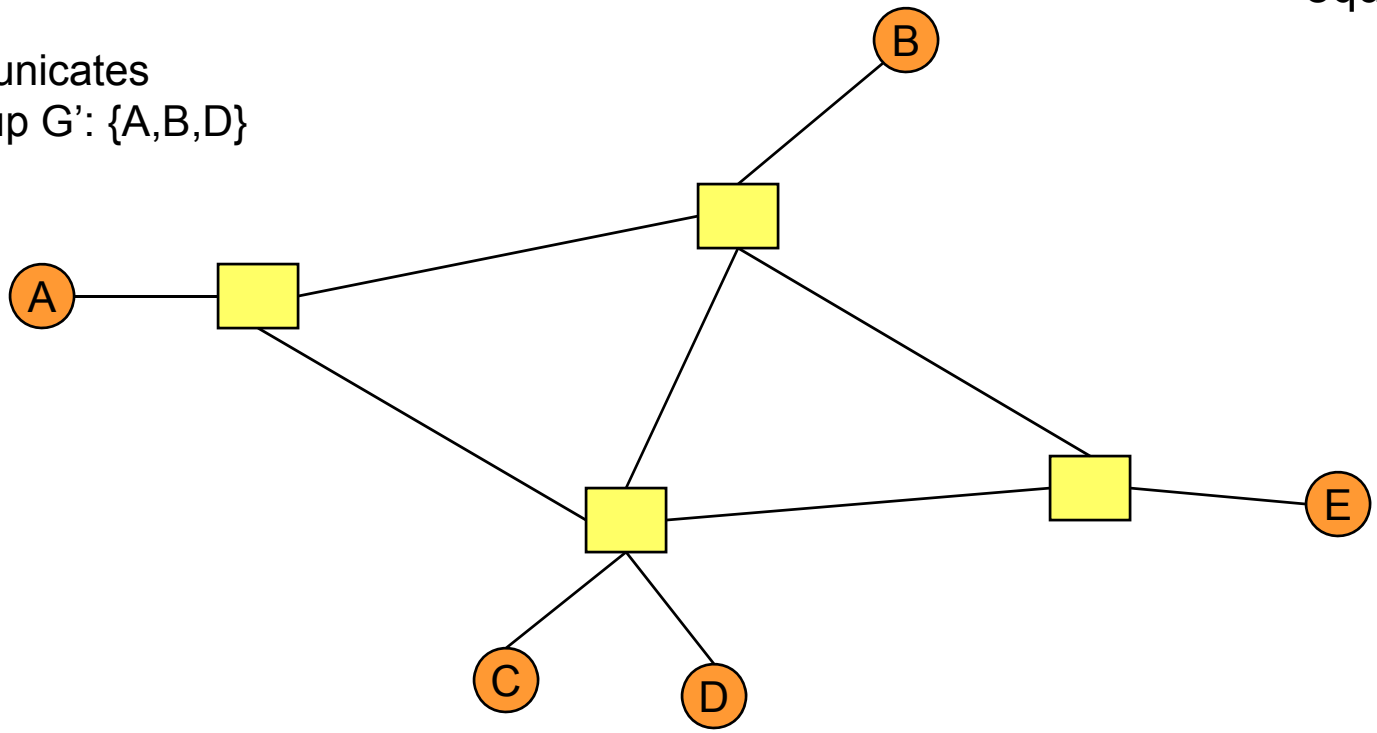

Source Based Trees *(cont'd)*

A communicates
with group G: {B,C,D}

C communicates
with group G': {A,B,D}

All edges have
equal weights

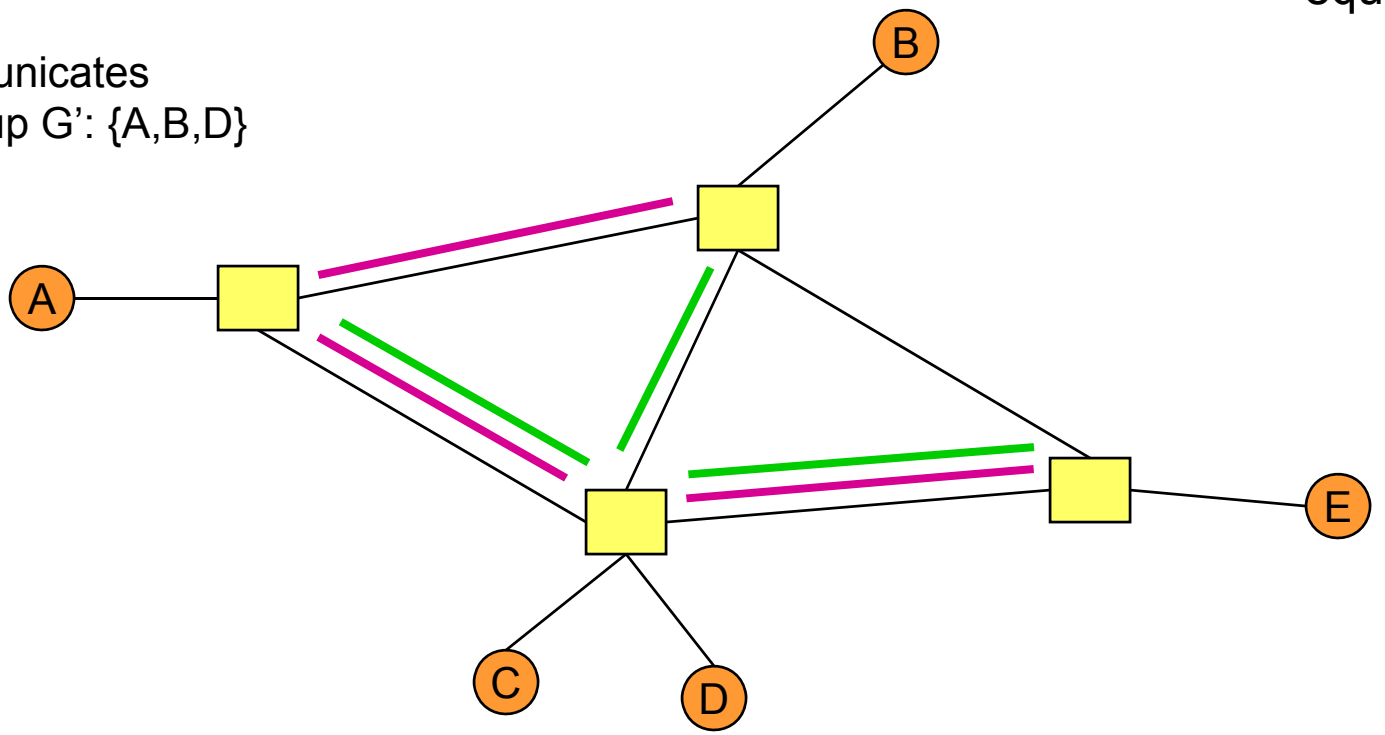

Spanning Tree Path for (A, G)

Spanning Tree Path for (C, G')

etc...

Source Based Trees *(cont'd)*

A communicates
with group G: {B,C,D}

C communicates
with group G': {A,B,D}

All edges have
equal weights

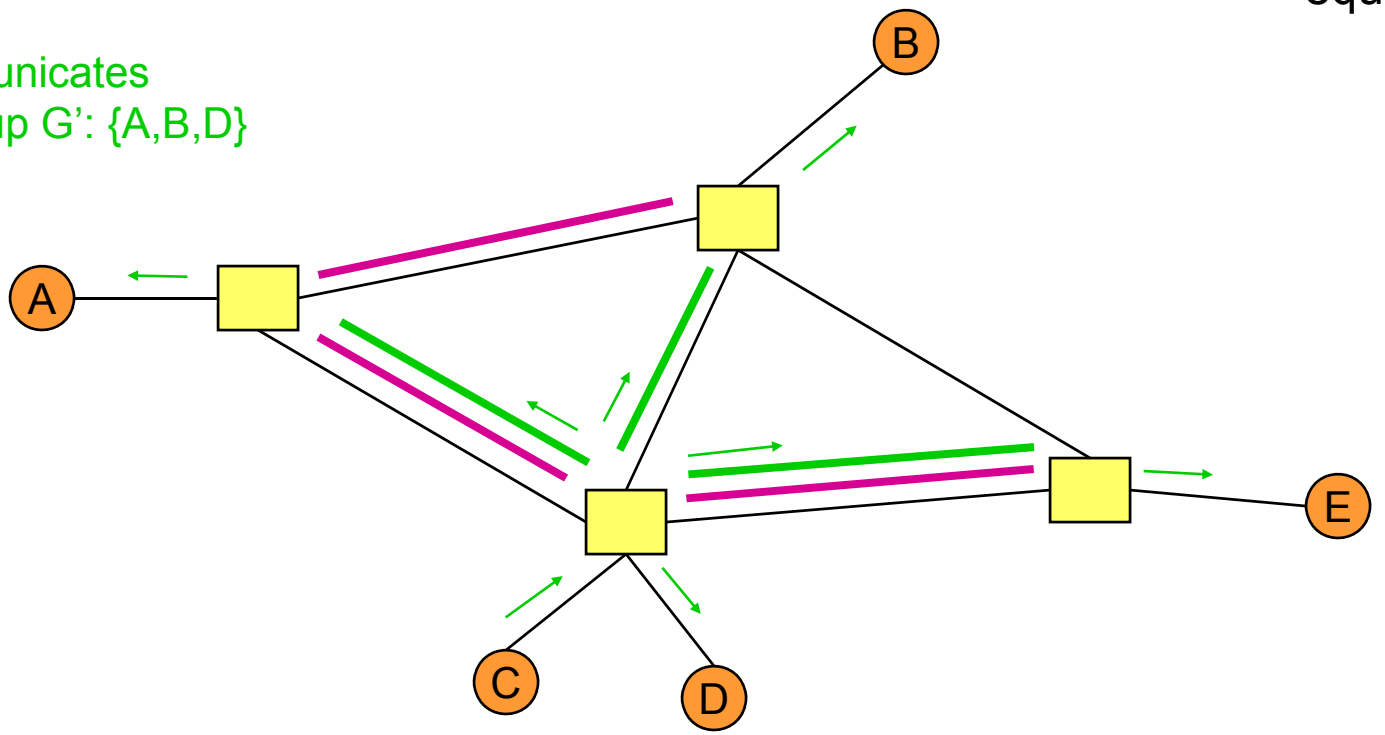

Spanning Tree Path for (A, G)

Spanning Tree Path for (C, G)

etc...

Source Based Trees

Advantages and Disadvantages

- Advantages:
  - Packets follow shortest paths to all destinations
  - No duplicate packets are generated in the network
- Disadvantages:
  - Source Based Trees must be explicitly set up
  - Multicast routing tables can grow very large, since they carry separate entries for each source
  - Packets still arrive where they aren't wanted

3.4 Reverse Path Forwarding

- Also known as Reverse Path Broadcast
- RPF is a simple algorithm used to achieve source-based spanning trees *implicitly*
  - Unicast routing tables are used to make forwarding decisions
- Forwarding algorithm:
  - When packets arrive at a router,
    - If the packet arrives on a link the router would normally use to reach the packet's source, then the router forwards a copy of the packet on all other outgoing links
    - If the packet arrives on *another* link, then it is discarded

RPF *(cont'd)*

All edges have equal weights

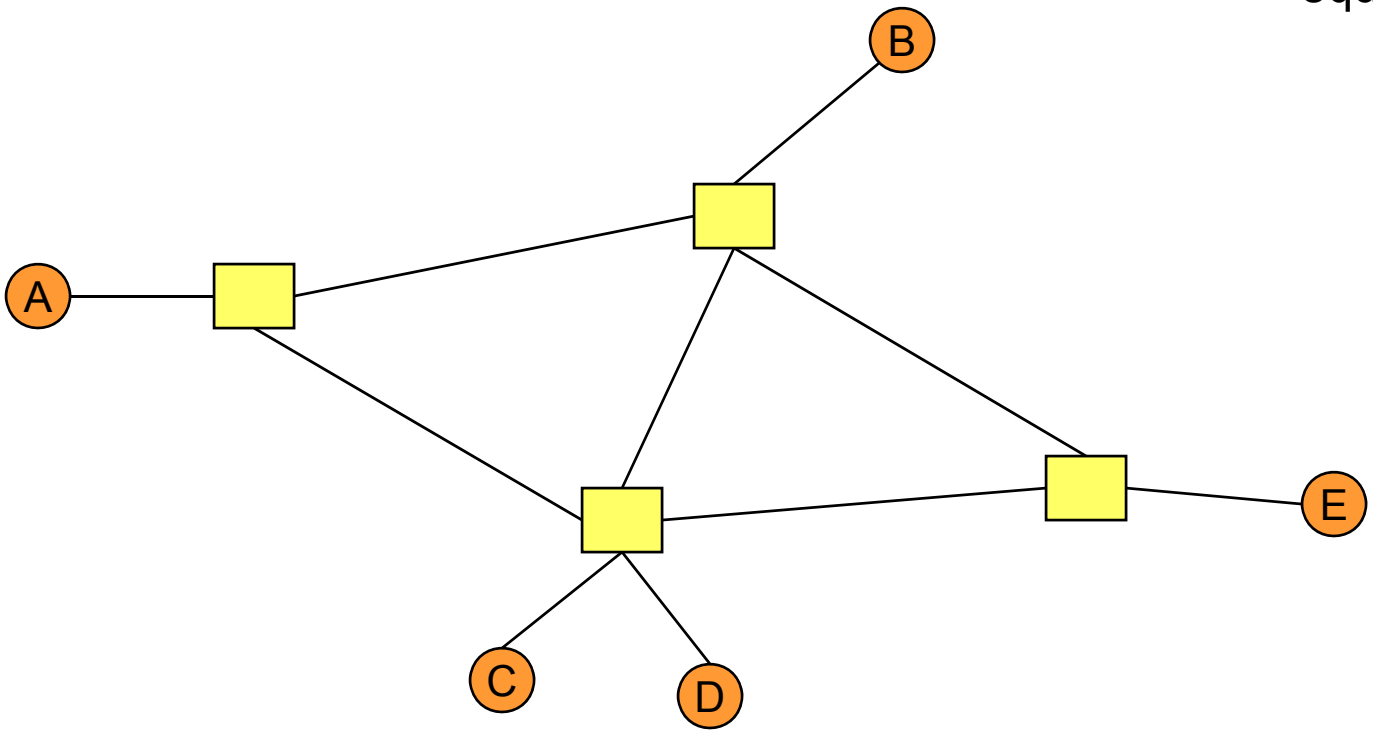

RPF *(cont'd)*

C communicates
with group G: {A,B,D}

All edges have
equal weights

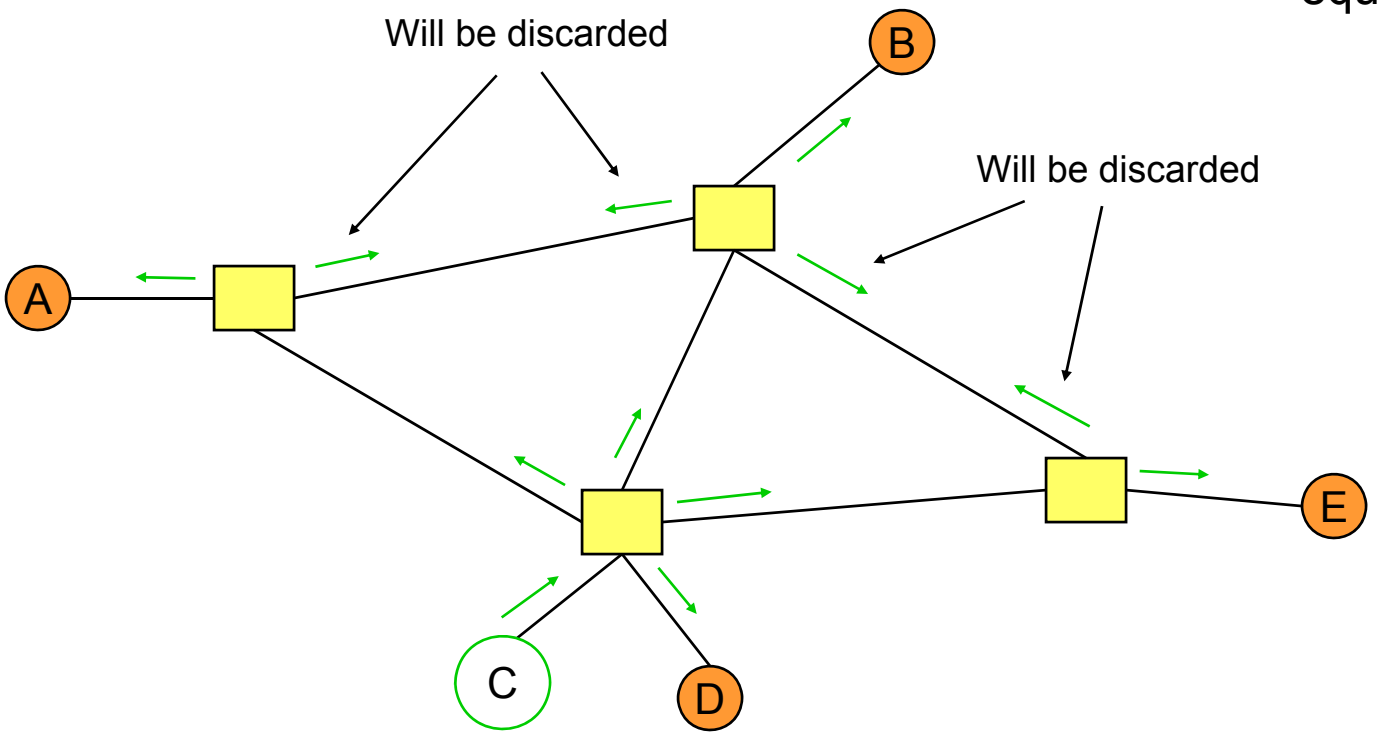

RPF

Advantages and Disadvantages

- Advantages:
  - Efficient and easy to implement
  - No sequence number tracking is required
  - Source-based spanning trees are constructed *implicitly* using unicast routing tables; no explicit tree construction is required
- Disadvantages:
  - Some duplicate packets still arrive at routers
  - Packets still go where they aren't wanted

3.5 TRPB (Truncated RPB)

- Truncated Reverse Path Broadcasting
- Extension of RPF
- Uses special control messages so a router can determine if there are any members of the multicast group on the subnet
  - Control messages are generated by a special protocol called the Internet Group Management Protocol (IGMP)
- If there are no members listening to the multicast group, the router *truncates* the spanning tree and does not forward packets addressed to that group on the subnet

TRPB *(cont'd)*

All edges have equal weights

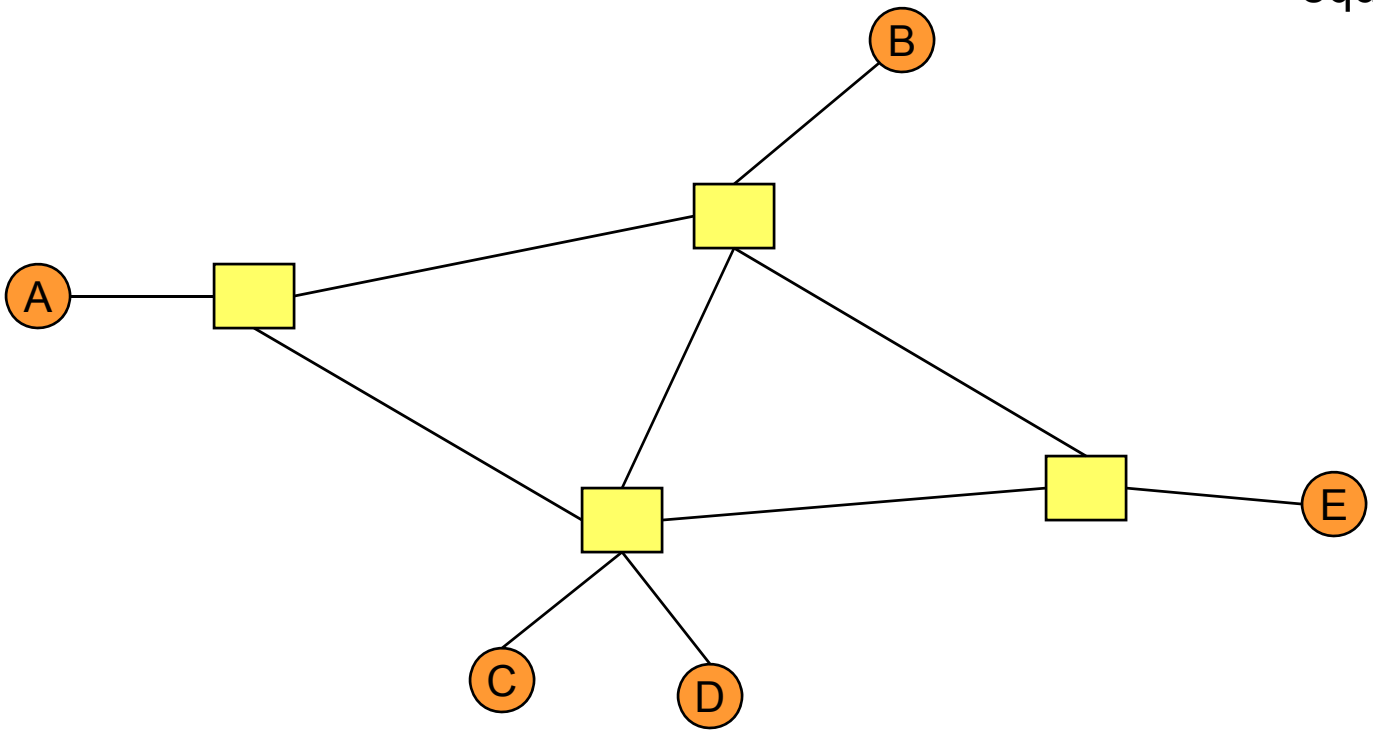

TRPB *(cont'd)*

C communicates
with group G: {A,B,D}

All edges have
equal weights

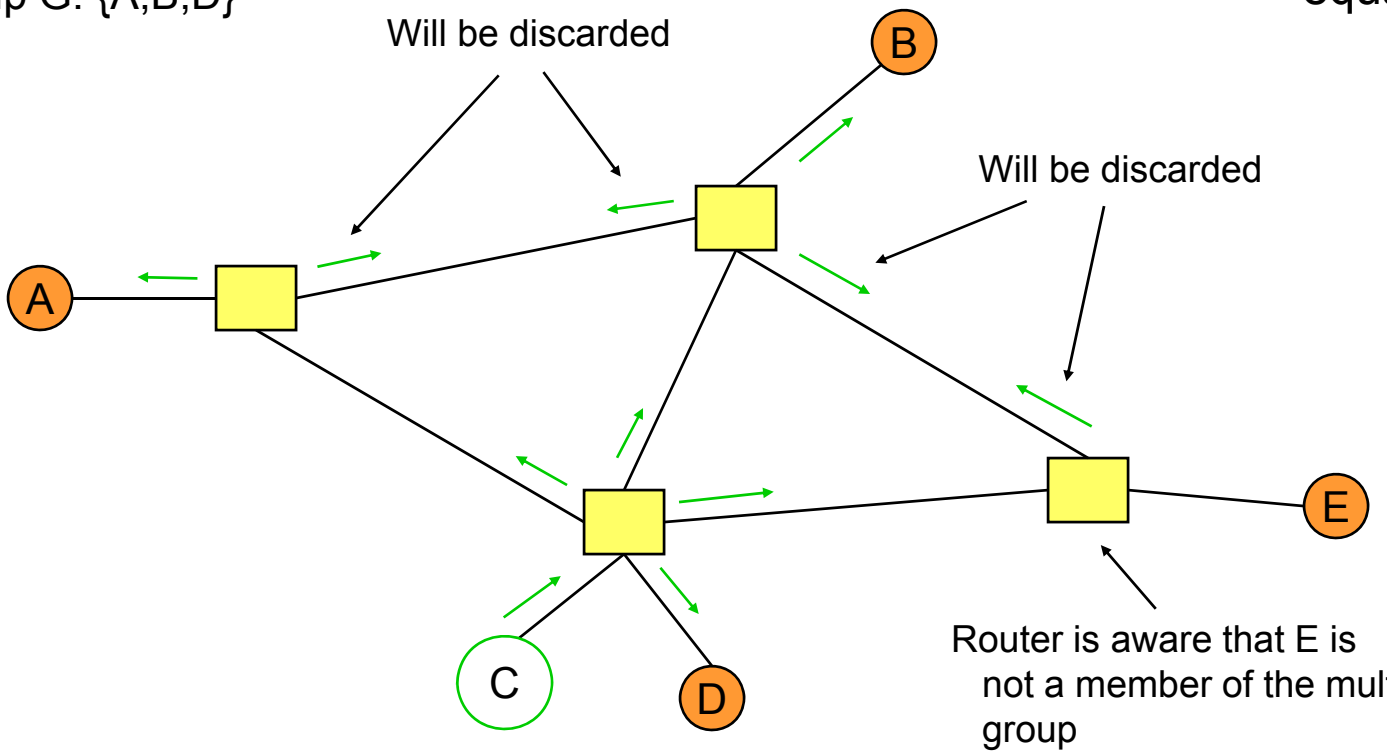

TRPB

Advantages and Disadvantages

- Advantages:
  - Same advantages as RPF but with the improvement that packets no longer go to destination hosts that don’t want them
- Disadvantages:
  - Packets still go to every router in the network, even those that don’t need to receive them

3.6 Reverse Path Multicast (RPM)

- Generalize the concept of truncating all the way back to the source.
- If a packet arrives and there are no group members downstream, a router sends a “prune” message on the link from which the packet arrived.
  - Prune messages allow the parent router stop forwarding the group’s packets down unnecessary branches

RPM *(cont'd)*

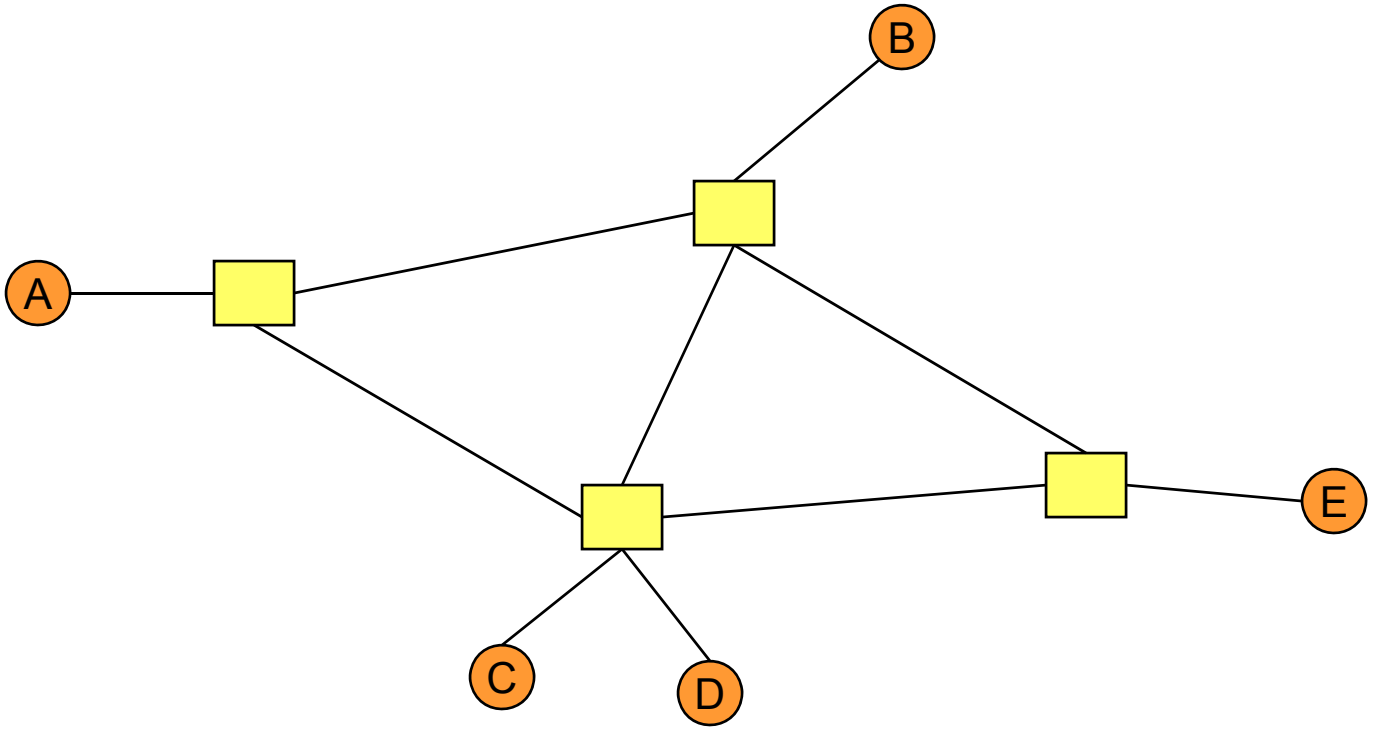

RPM *(cont'd)*

C communicates
with group G: {A,B,D}

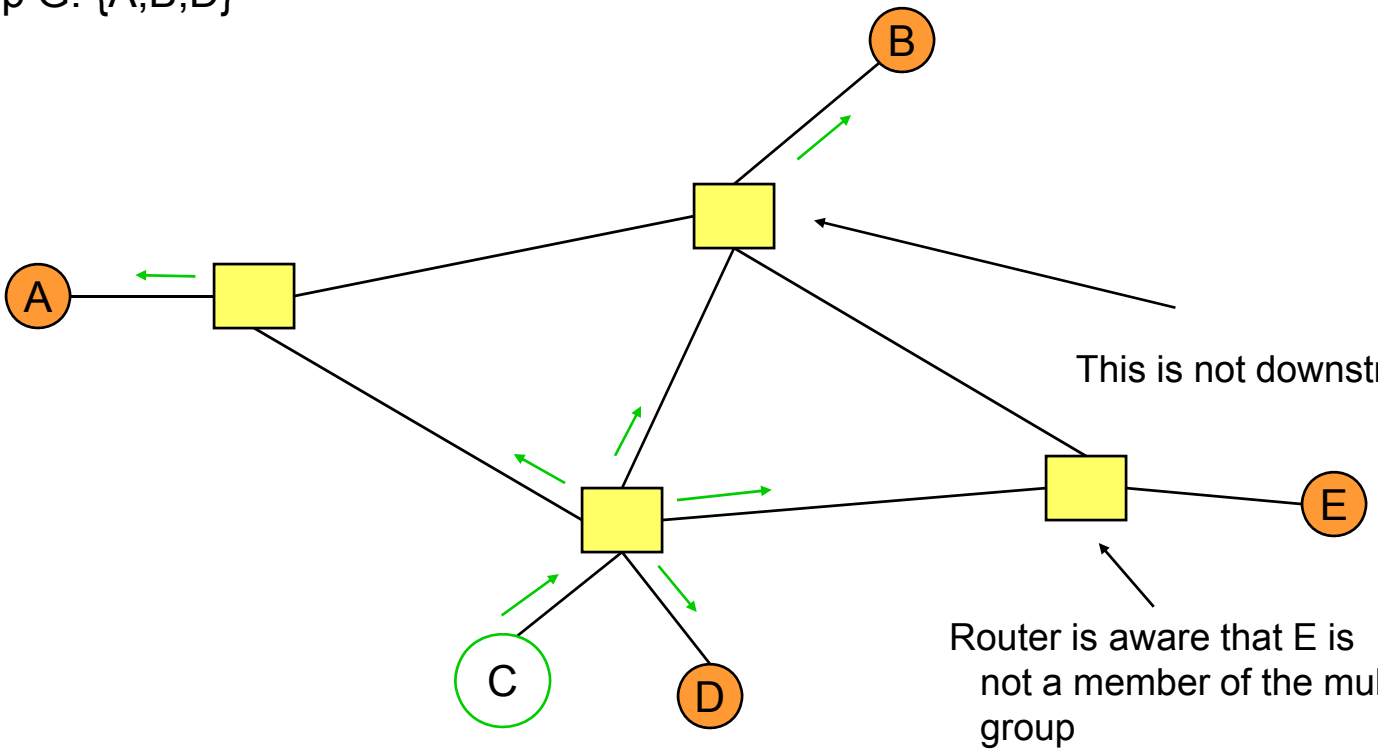

RPM *(cont'd)*

C communicates
with group G: {A,B,D}

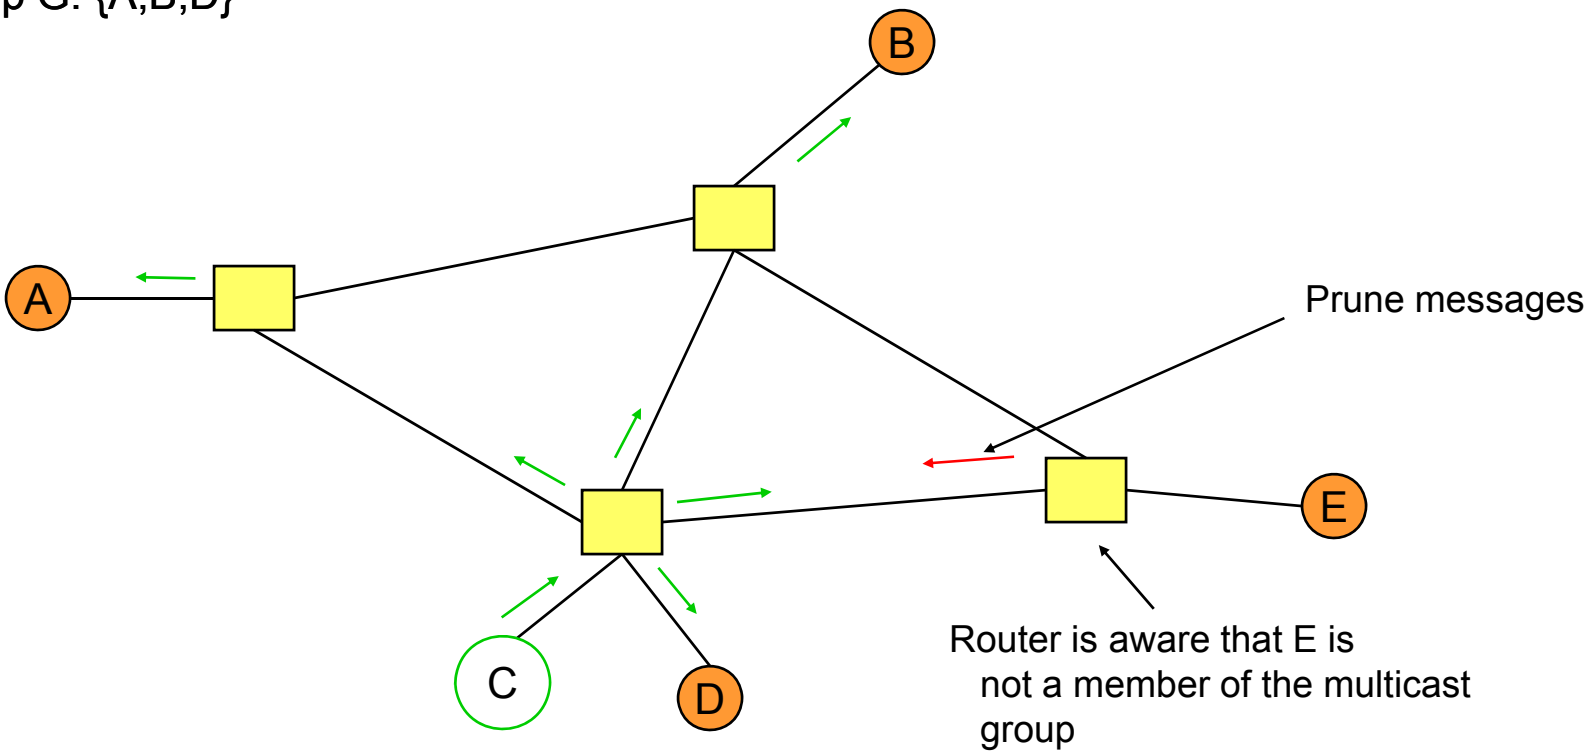

RPM *(cont'd)*

Later...

C communicates
with group G: {A,B,D}

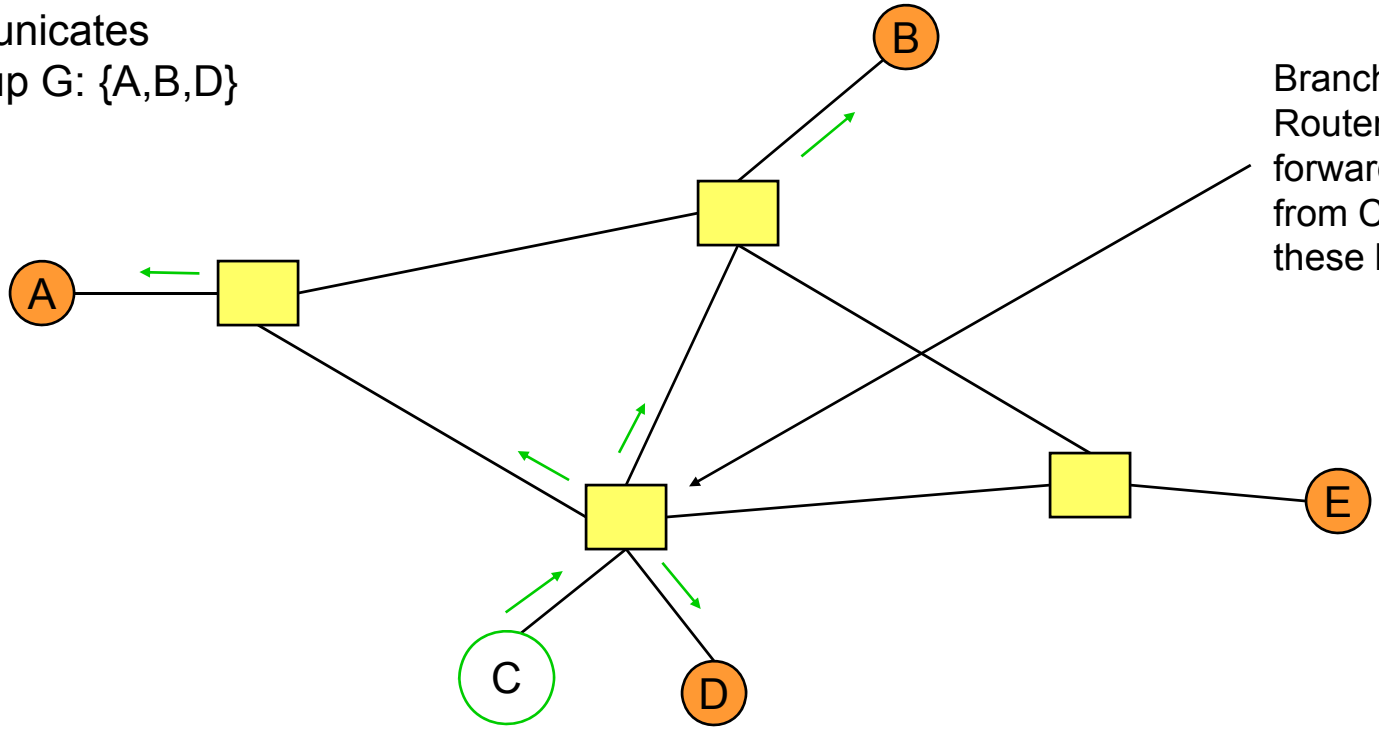

Branches pruned:
Routers no longer
forward packets
from C to G on
these links

RPM *(cont'd)*

- Pruned branches only stay pruned for a limited time. After timing out, the pruned branches "grow back"
  - This allows new (previously pruned) receivers to join the multicast conversation
- Routers also have the option of sending "graft" messages on the parent links when directly connected hosts join a pruned group
  - Graft messages quickly "unprune" a link from a multicast tree

RPM

Advantages & Disadvantages

- Advantages:
  - Same advantages as TRPB but with the additional advantage that multicast packets are not sent to routers that don't need them
- Disadvantages:
  - Greater complexity: requires the use of special "prune" and "graft" messages

4. Multicast in the Internet

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast OSPF (MOSPF)
- Core Based Trees (CBT)
- Protocol Independent Multicast (PIM)

4.1 DVMRP

- DVMRP is the most commonly used multicast routing protocol used by the Internet
- DVMRP = distance vector routing + RPM

4.2 MOSPF

- MOSPF is also widely used in the Internet
- MOSPF = OSPF (link state routing) + RPM

4.3 Core Based Trees (CBT)

- Another disadvantage of all the reverse-path multicast algorithms (RPF, TRPB, RPM) is that they require large multicast routing tables. In other words, they may not be scalable
  - For example, table entries for (source, group) pairs
- Core Based Trees reduce multicast routing table size

CBT *(cont'd)*

- CBT uses a shared tree that connects all receivers in each multicast group
  - One special router in the shared tree is called the "core router"
- When a receiver wishes to join a multicast group, it sends a "join request" message toward the core router
  - As join message passes through non-core routers, branches are added to the shared tree
- When a sender wishes to send packets to a multicast group, it sends the packet toward the core router.
  - The first router (core or non-core) to see the packet will intercept the packet and multicast it to all receivers on the shared tree

CBT *(cont'd)*

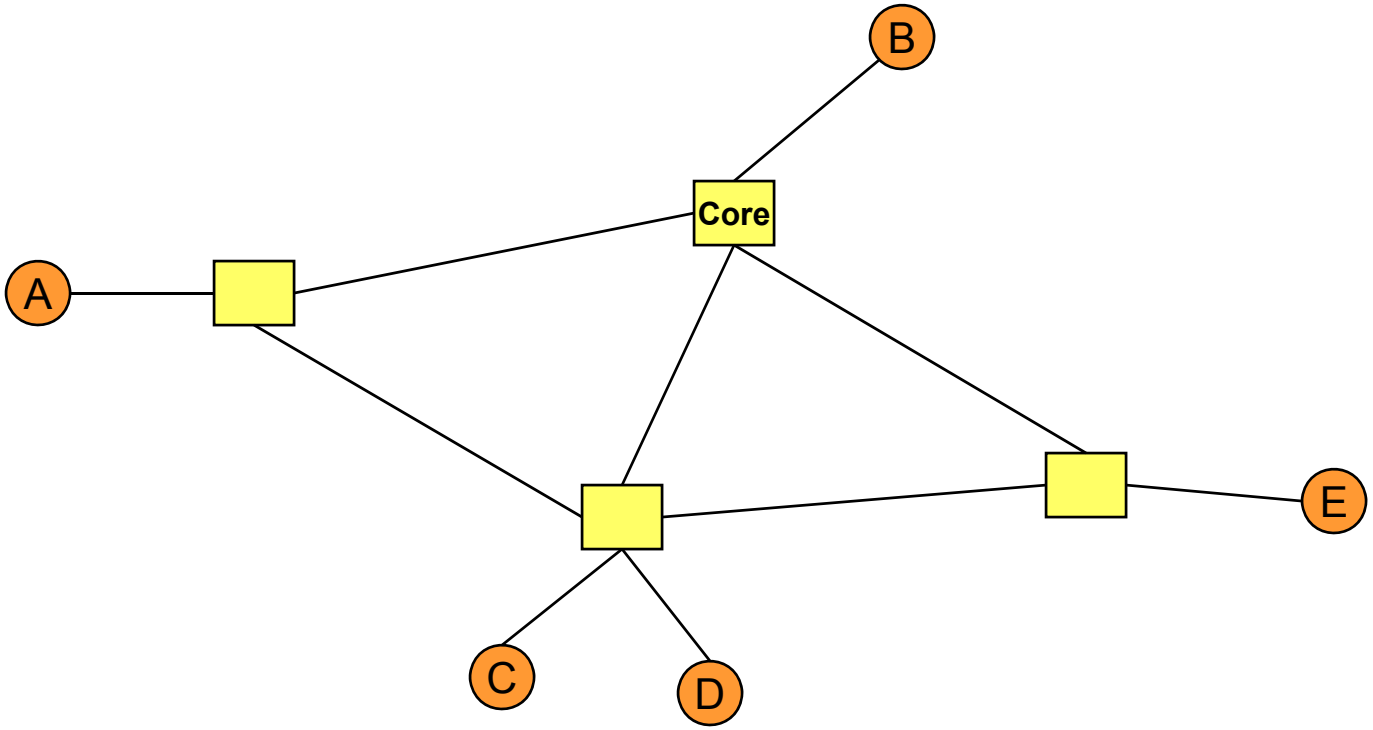

CBT *(cont'd)*

A and C join multicast group G_1

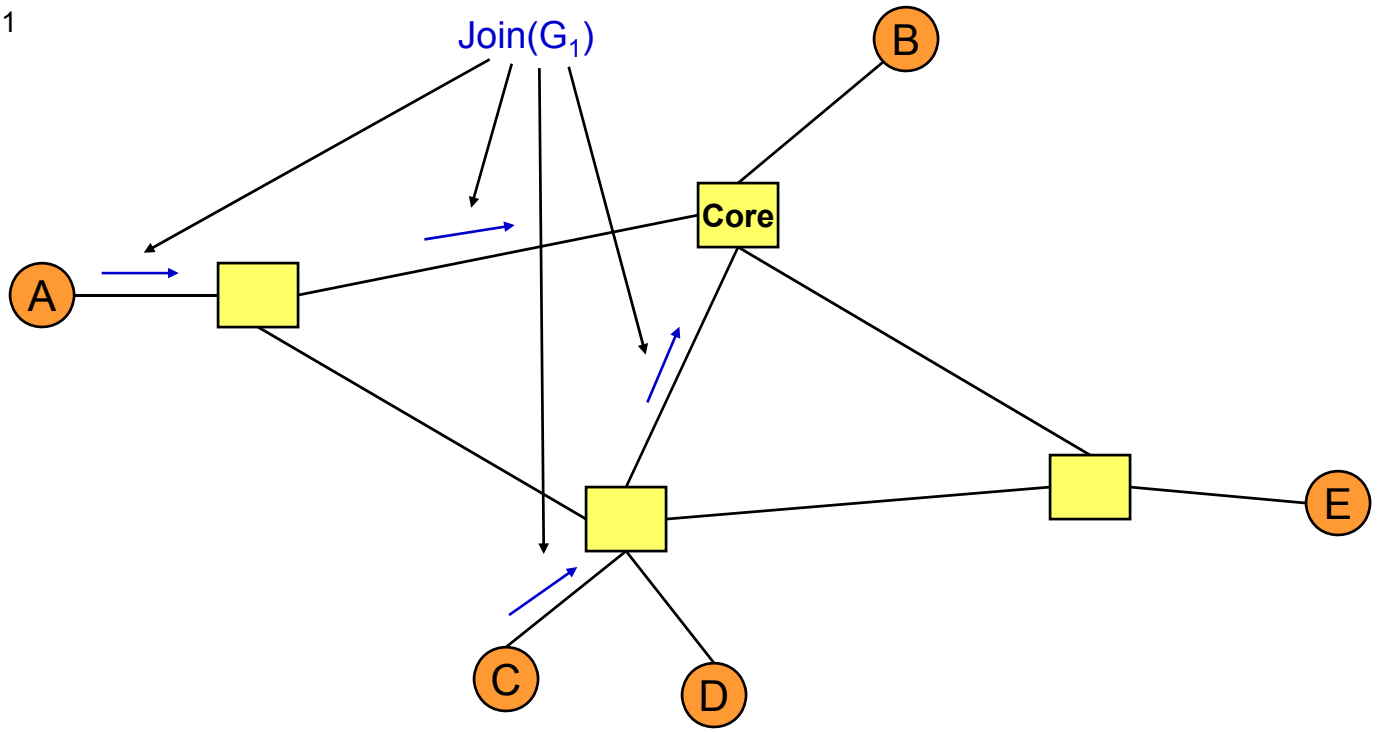

CBT *(cont'd)*

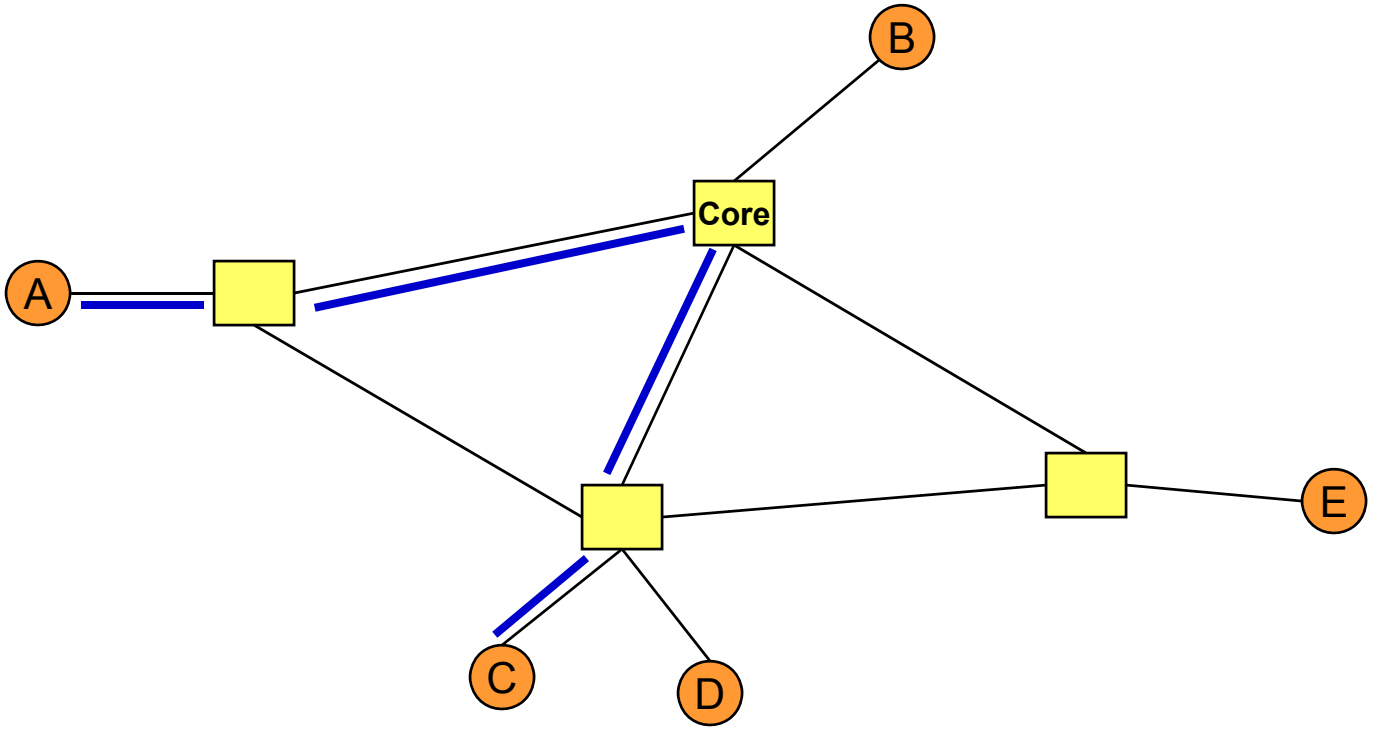

CBT *(cont'd)*

E multicasts a packet to G_1

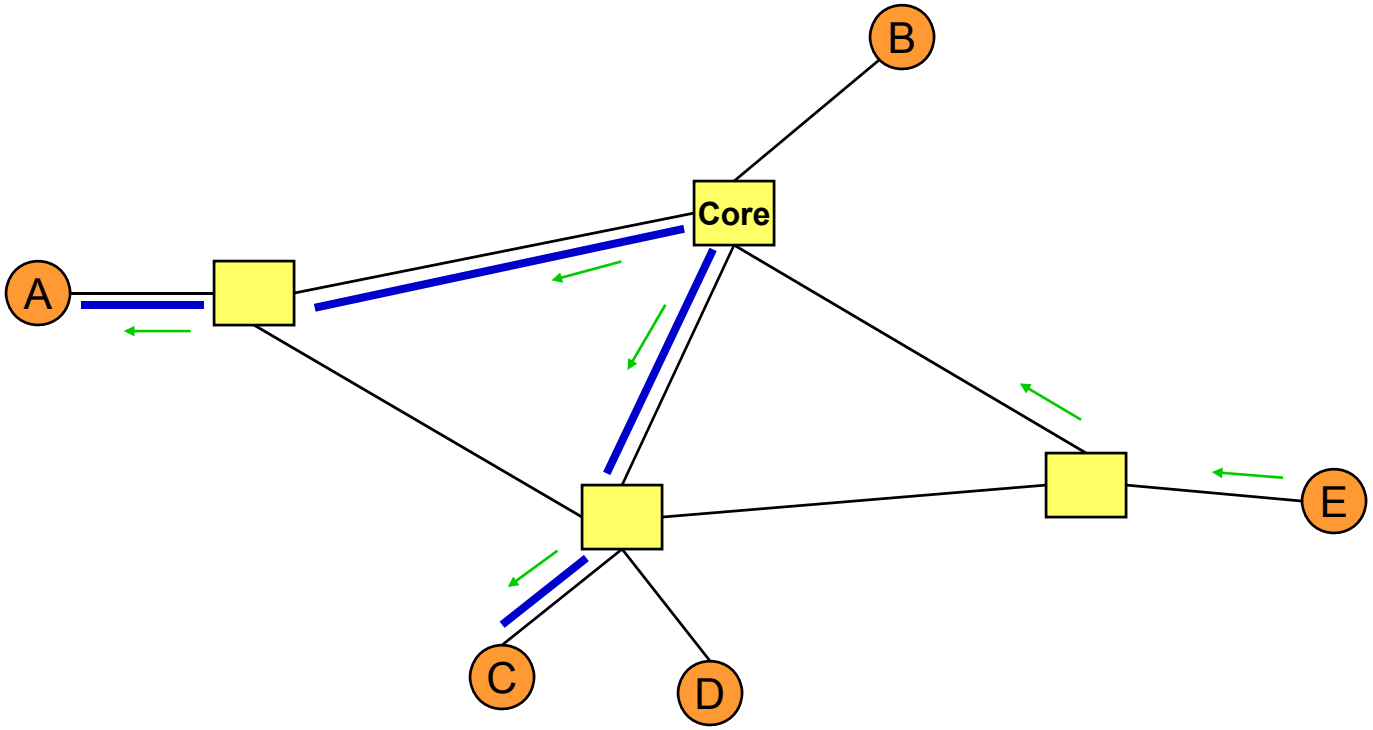

Note that with CBT, the sender does not need to be a member of the multicast group

CBT

Advantages & Disadvantages

- Advantages:
  - Smaller router tables, so more scalable
    - only one entry per multicast group
    - not one entry per (source, group) pair like RPM
  - Senders do not need to join a group to send to it
- Disadvantages:
  - Shared trees are not as optimal as source-based trees
  - Core routers can become bottlenecks

4.4 Protocol Independent Multicast (PIM)

- PIM tries to give you the best of both worlds
- PIM has two modes:
  - Dense mode (PIM-DM):
    - Allows for source-based trees
    - designed for environment where group members are densely packed and bandwidth is plentiful
    - Very similar to DVMRP
  - Sparse mode (PIM-SM):
    - Uses shared trees
    - designed for environment where group members are sparsely distributed and bandwidth is not widely available
    - Very similar to CBT

- Protocol Independence: will work with any underlying unicast routing protocol

5. IGMP

- Internet Group Management Protocol (IGMP)
- runs between a router and its directly connected hosts
- Allows a router to know which of its directly connected hosts belongs to which multicast group
- IGMP is required to support TRPB, RPM, CBT and PIM protocols

IP Address Classes: Review

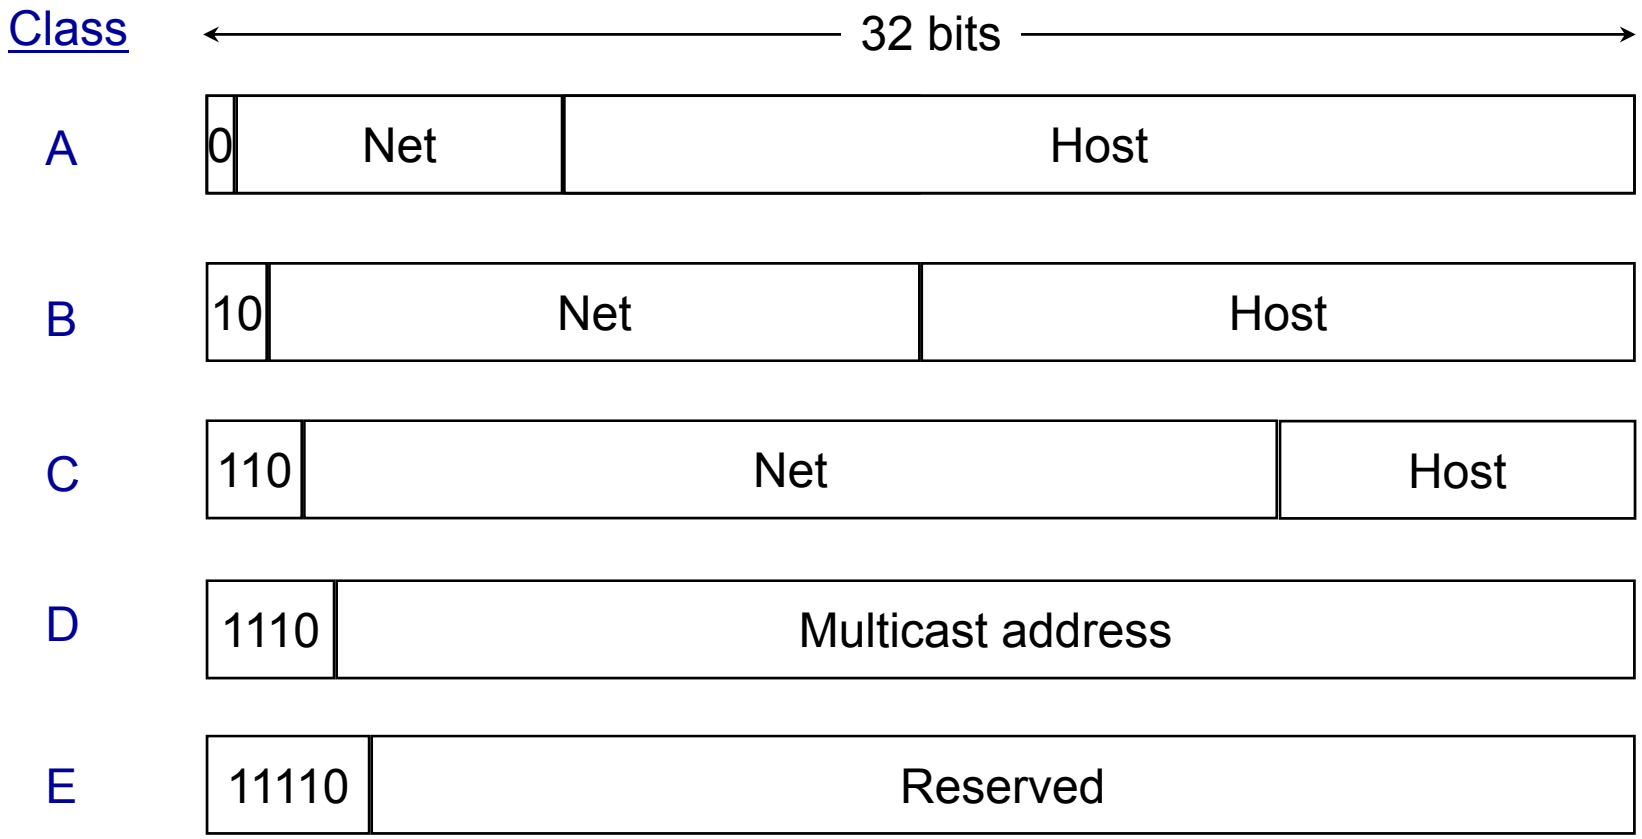

IP Address Classes: Review

- Class A:
  - For very large organizations
  - 16 million hosts allowed
- Class B:
  - For large organizations
  - 65 thousand hosts allowed
- Class C
  - For small organizations
  - 255 hosts allowed
- Class D
  - Multicast addresses
  - No network/host hierarchy

- Class E
  - reserved
- Loopback
  - 127.xx.yy.zz (127.anything) is reserved for loopback testing
  - packets sent to this address are not put out onto the wire; they are processed locally and treated as incoming packets.
- Broadcast
  - all 1s

IGMP *(cont'd)*

- IGMP messages are transmitted in IP packets

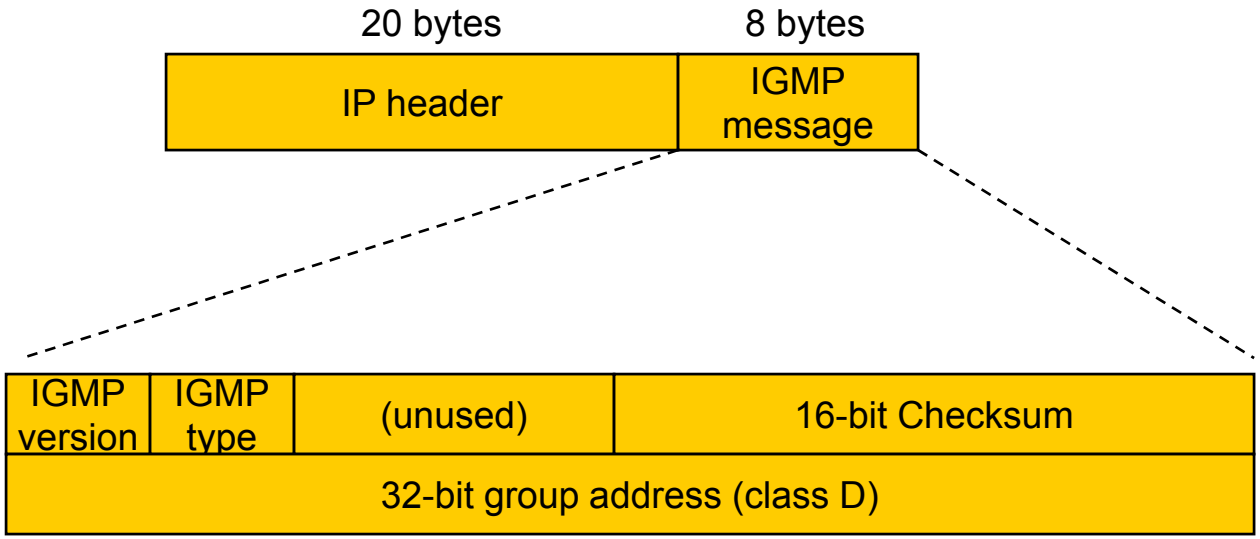

IGMP *(cont'd)*

- IGMP Version 1 [RFC 1112]
  - A router periodically transmits Host Membership Query
  - A query message is addressed to the all-hosts group (224.0.0.1) and have a TTL = 1
  - A host responds with a Host Membership Report for each multicast group to which it belongs.
  - If a router does not receive a Report for a group, that group is removed from the list of multicast groups that the router maintains
  - A host that newly joins a multicast group first transmits a Report for the multicast group rather than waiting for a Query from the router.

IGMP Host Membership Queries

- Routers uses IGMP “query” messages to periodically query hosts on their subnets and learn if they are members of *any* multicast group
  - Hosts who are members of multicast groups respond with one IGMP “report” message for each group they are a member of
  - To improve efficiency, hosts wait a random amount of time before responding
    - During this waiting time, hosts listen to other host responses
    - If another host reports membership in the same group, then the host aborts its report

IGMP Host Queries

An Example

Example: A and B are members of multicast group G_1

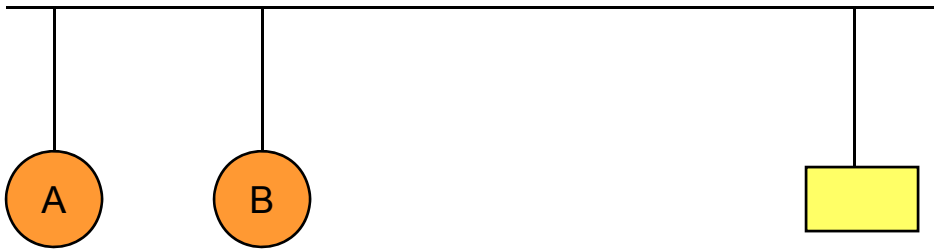

Example: A and B are members of multicast group G_1

IGMP query, TTL=1
IGMP destination group = 0
IP destination address = 224.0.0.1
IP source address = router address

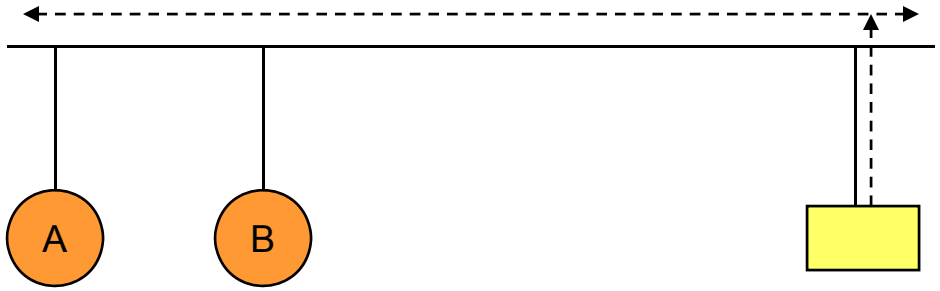

A waits
100 ms

B waits
200 ms

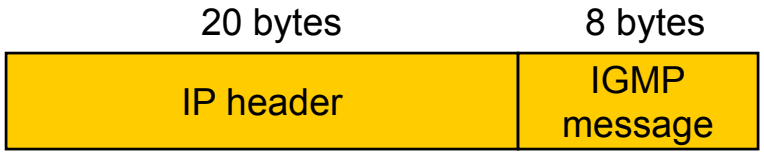

- In the previous figure
  - IGMP type: GMP query
  - IGMP destination group = 0
    - Hosts will ignore this destination group address in the IGMP message
  - IP header TTL value =1
  - IP destination address = 224.0.0.1
    - This address means all host multicast group
  - IP source address = router address

Example: A and B are members of multicast group G_1

IGMP report, TTL=1
IGMP destination group = G_1
IP destination address = G_1
IP source address = A

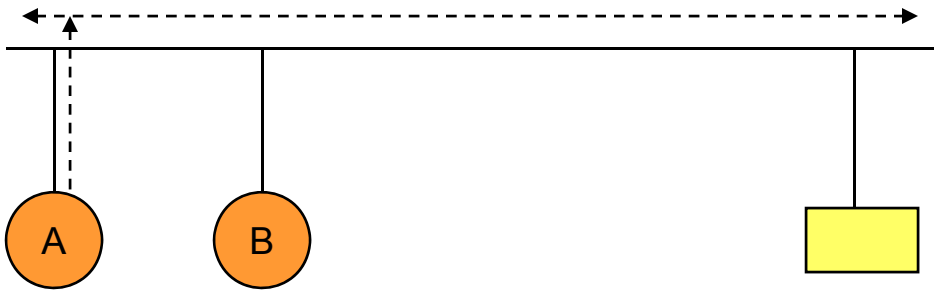

A responds 100 ms later

B aborts response because it hears A's report

Router is now aware that at least one host on its subnet is a member of G_1

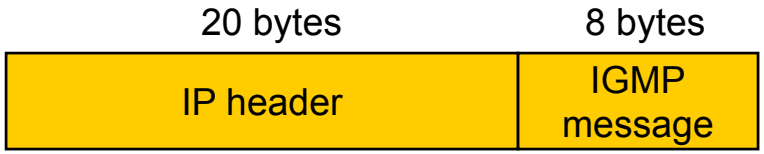

- In the previous figure
  - IGMP type: IGMP report
  - IGMP destination group = G1
    - To indicate that the host is in G1
  - IP header TTL value = 1
  - IP destination address = G1
    - To send this IGMP report to all hosts in G1
  - IP source address = A

IGMP Reports

- Hosts may also send IGMP reports when they first join a multicast group
  - In this case they don't need to wait for an IGMP query first
- When hosts leave a group, they do not need to announce
  - The router will discover if no one is left in the group when it does its next IGMP query