IP Multicast

Basic Concepts

- Unicast:
  - Sent to a specific host.
- Broadcast:
  - Sent to all hosts (normally restricted to the local LAN)
- Multicast:
  - Sent to a group (not everybody)
  - Can go across routers
  - Main reason: efficiency, saves on bandwidth over possible alternatives:
    - Multiple unicast: packet sent many times.
    - Broadcast: hosts that are not interested in the packet get it.

Multipoint Taxonomy

- Two Planes:
  - ***Control Plane***: deals with the establishment of the multipoint session
  - ***Data Plane***: deals with the data transfer amongst the session participants
- Types of Control Planes:
  - ***Rooted***: one special participant, the root, establishes the session, and adds the leaf nodes. Session does not exist without the root.
  - ***Non-rooted***: all participants are leaves. They add themselves to a pre-existing session (this is the same as having an implicit root in the network).

Multipoint Taxonomy (cont.)

- Types of Data Planes:
  - ***Rooted***: there is a special participant, the root. Data flows from the root to all the leaves. Leaves may send data back to to root, but not talk amongst themselves.
  - ***Non-rooted***: all participants are leaves. Any data sent will be delivered to all participants. Non-participants may also send data.

	rooted control plane	**non-rooted control plane**
rooted data plane	ATM ST-II	no known examples
non-rooted data plane	T.120	IP Multicast H.320 (MCU)

IP Multicast

- IP multicasting is the “Internet abstraction” of LAN multicasting.
- Membership of an IP- multicast group is dynamic.
  - A host may join or leave at any time
  - A host may be member of an arbitrary number of multicast groups.
- Membership determines whether the host will receive datagrams sent to the multicast group.
- A host may send datagrams to a multicast group without being a member.
- Basic definition: RFC 1112

IP Multicast Addresses

- A multicast group has a unique class D address

0	1	2	3	4 … 31
1	1	1	0	Group Identification

Addresses range from 224.0.0.0 through 239.255.255.255

224.0.0.0 reserved.

224.0.0.1 all hosts group (all hosts & routers participating IP multicast)

  - Well-known addresses
  - Transient multicast addresses
- Multicast addresses can be used only as destinations

Mapping IP Multicast Address to Ethernet Multicast Address

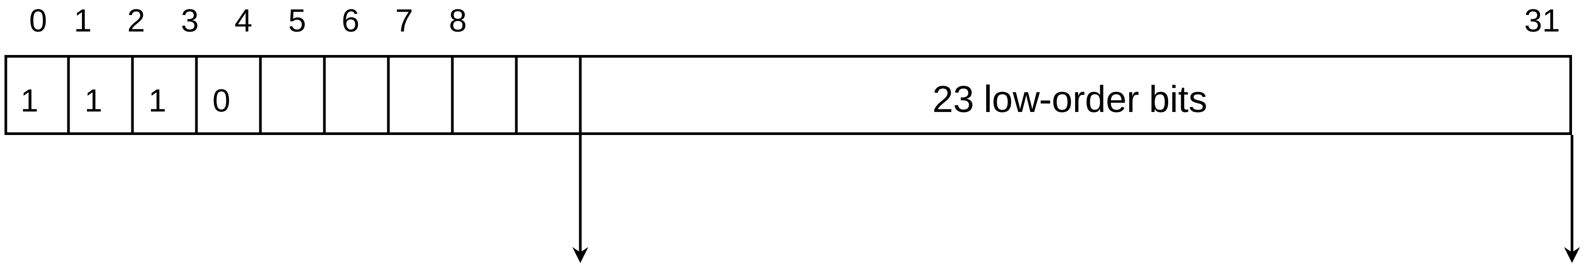

Ethernet special multicast MAC address

$$01.00.5E.00.00.00_{16}$$

with 23 low-order bits replaced by 23 low-order bits of IP multicast address

Note that 32 Class-D addresses map into the same MAC address, so software must be prepared to filter.

Multicast Support In Current Ethernet Devices

- All Ethernet interfaces are capable of “promiscuous mode” - receive all packets. This is the lowest common denominator.
- Some devices can be programmed to receive all multicasts, or be given a list of multicast MAC addresses to listen to.
- Most common today: use of a hash table
  - Run the desired multicast MAC address through the Ethernet CRC generator.
  - Use five or six bits of the CRC to index a flag in memory
  - If the flag is set, the packet is received; otherwise, it is discarded.

Benefits of Multicast (1)

- 1. Resource Discovery
  - In OSPF, sending a hello packet to all neighbors on non-broadcast networks requires the routers to be configured with a list of neighbor addresses.
  - Multicast allows the router to discover this list with little burden on the administrator. (TTL = 1).
  - In BOOTP, servers need not be on the same LAN, use “scope-limited” broadcasts (TTL = 1, the 2, etc.)

Benefits of Multicast (2)

- 2. Efficiency:
  - Fewer transmissions compared to point-to-point
  - Example: Node C sending a file to nodes A, B, D and E.

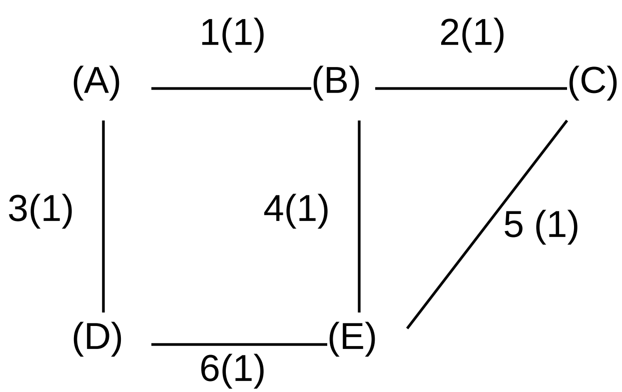

Using For	Point-to-point A	B	E	D	Total	Multicast
Link 1	1				1	1
2	1	1			2	1
5			1	1	2	1
6				1	1	1
	2	1	1	2	6	4

Benefits of Multicast (3)

- 3. Multipoint Communication:
  - Single Source multicasting to a number of destinations (one way communication, one-to-many)
    - Unknown destinations
    - Variable membership
  - Conferencing: many-to-many
    - Multiple sources
    - Multiple destinations

Multicast Routing

- Flooding
- Spanning Trees
- Reverse Path Forwarding
- Steiner Trees
- Core - based Trees

Multicast Routing
Flooding (1)

- Simplest multicast routing algorithm
- Used in:
  – OSPF - flooding link state updates
  – “Usenet news” messages flooded to news servers
- Basically:
  – 1. Node ascertains if received packet is new or duplicate
  – 2. If new, transmit on all interfaces except interface on which it came.

Multicast Routing
Flooding (2)

- Test for 1st reception of a packet
  - OSPF: examine link state database
  - Usenet news:
    - Comparison with a local database (or index) bared on a unique identifier for the news message.
    - Use of trace information listing hosts that have already seen the message
      - (i) If local host is present in list, packet is not processed any further.
      - (ii) If a neighbor is present in list, packet is not forwarded.

$\Rightarrow$ Application-level flooding.

Multicast Routing
Flooding (3)

- At IP Level, similar functions could be implemented.
  - List of recently seen packets (e.g., last 2 minutes)
    - But list could be long
    - Packet may still be received more than once
  - Route recording option - (similar to traces)

Multicast Routing
Spanning Tree (1)

- Similar to what was described in IEEE 802.1
  - Forward packet only on interfaces belonging to spanning tree.

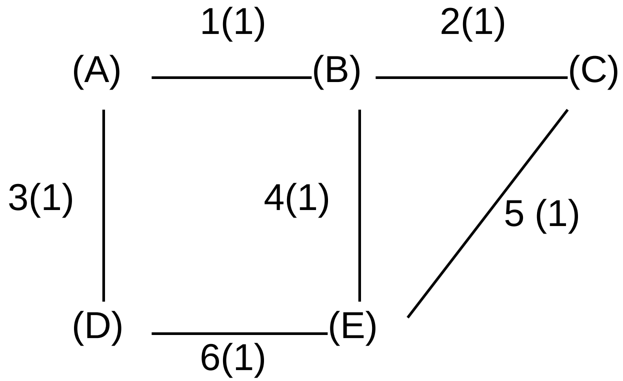

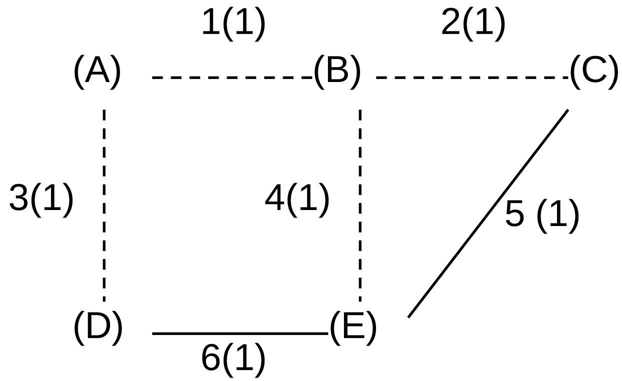

- Drawbacks:
  - 1. It does not take into account group membership
  - 2. May concentrate traffic on a small subset of network links.

Multicast Routing
Spanning Tree (2)

- Example:
  - One group includes only A, B and C
  - Another group includes nodes C, D, E

  ⇒ Packet still propagated to D and E.

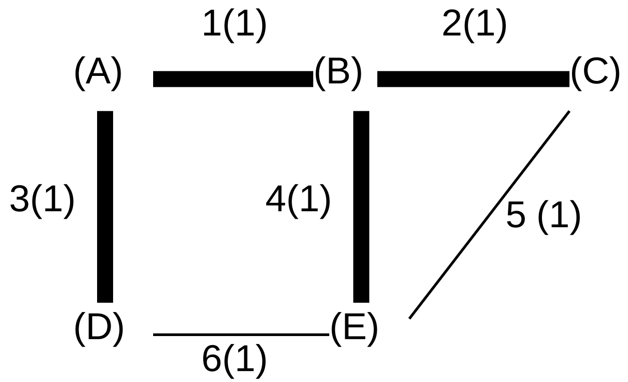

⇒ Packet will propagate on all links of spanning tree. Propagating it on links 5 and 6 is more efficient.
⇒ Group specific graphs to achieve higher efficiency.

Multicast Routing

Reverse-Path Forwarding (1)

- Used in the MBONE
- Takes advantage of unicast routing tables to figure out a spanning tree for each multicast source.
- Algorithm
  - 1. When a multicast packet is received, note source (S), interface (I)
  - 2. If I belongs to the shortest path toward S, forward to all interfaces except I.
  - 3. If the test in step 2 is false, drop the packet.

    Good: Does not require any more resources than the normal unicast routing table.

    Bad: Uses reverse-path metric (should use forward-path metrics!)

Multicast Routing
Reverse-Path Forwarding (2)

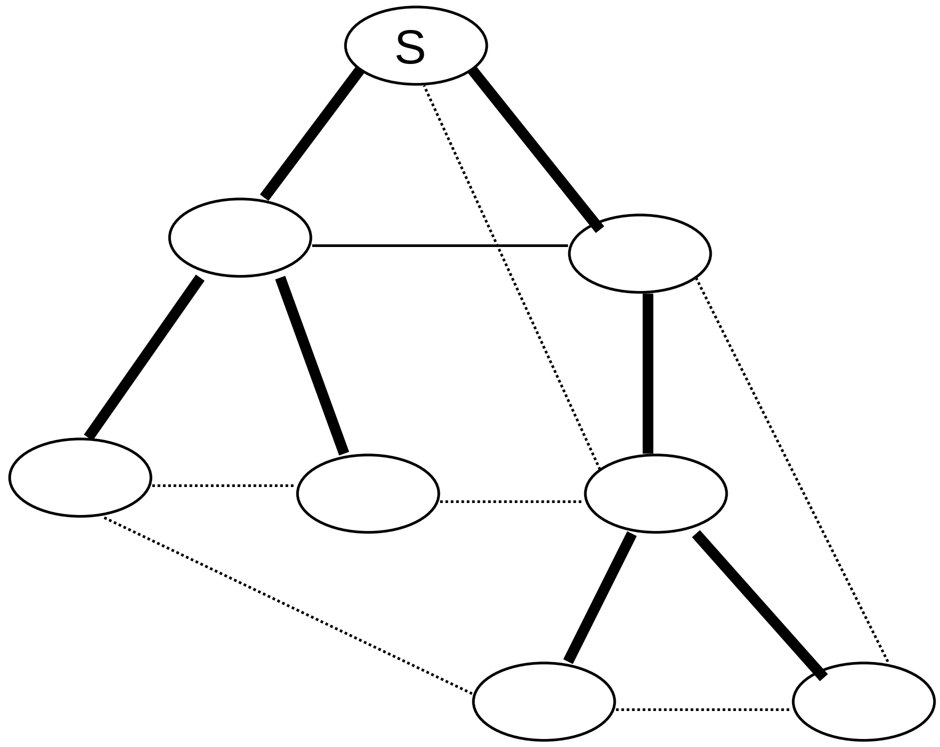

Multicast Routing
Reverse-Path Forwarding (3)

- Simple Improvement:
  - If local router is able to determine whether or not it is on the shortest path between a neighbor and the source, then:
    - If yes, forward to neighbor
    - If no, do not. (neighbor will drop it)

Multicast Routing

Reverse-Path Forwarding (4)

- RPF results in different spanning tree for each source.

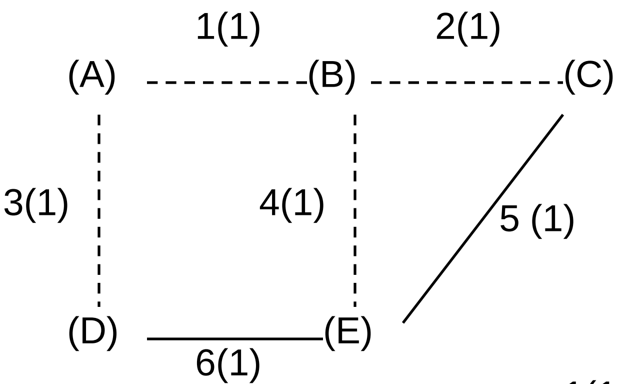

RPF tree from source (A)

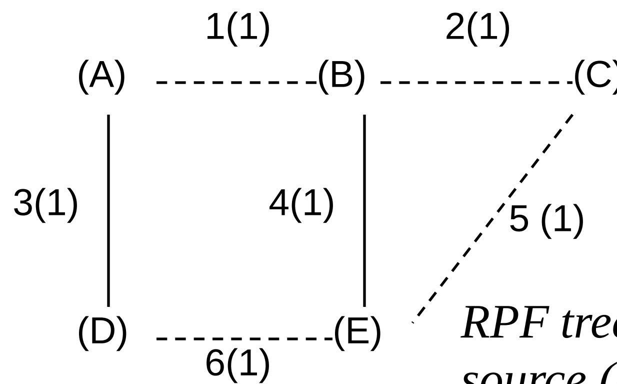

RPF tree from source (C)

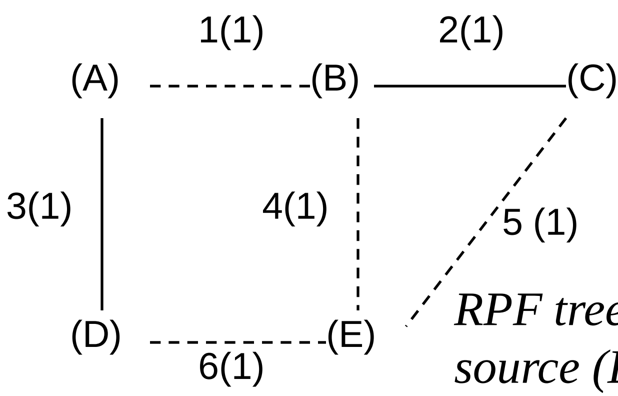

RPF tree from source (E)

Advantages:

- Fastest possible delivery
- Spreading traffic

Multicast Routing

Taking into Account Group Membership

- In RPF and all previous algorithm
  - Multicast packets are broadcast to all nodes. No group membership information is taken into account.
  - Possible improvement: routers connected to leaf subnets may filter multicast traffic according to group membership.

- IGMP (Internet Group Membership Protocol) is introduced to allow hosts to indicate whether or not they are interested in joining a multicast.

Multicast Routing
RPF and Pruning (1)

- Basis:
  - Memorizing group membership within the tree, and forwarding packets down only if group members exist.
  - Note:
    - Group membership information may vary rapidly & require very dynamic updates!
- MBONE RPF Pruning Option: Flood and Prune
  - First packet of multicast transmissions is propagated to all nodes in network;
  - Leaf nodes will send back “Prune Messages” to upstream router if no member exist in leaf network.
  - Intermediate nodes send upwards “Prune Messages” if they receive such messages on all their downstream interfaces.

Multicast Routing
RPF and Pruning (2)

- Drawbacks:
  - 1st packet is flooded.
  - Routers must keep states per group and per source.
    - State varies with membership & network topology.
    - State must be kept for limited lifetime.

    ⇒ periodic flooding & pruning required

Multicast Routing
Steiner Trees (1)

- Minimizing Network Resources.
  - Example: Multicast from C to A and D

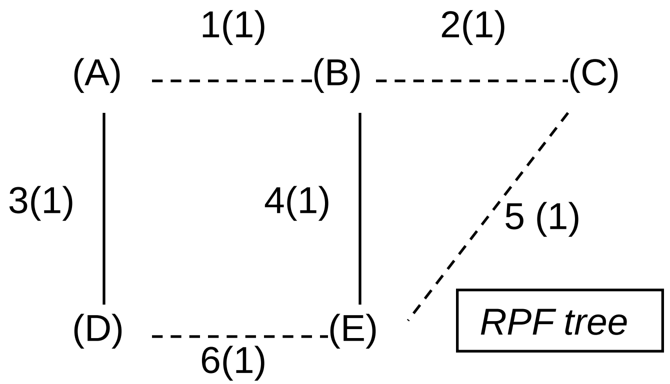

Steiner tree minimizes total # of links to reach all members

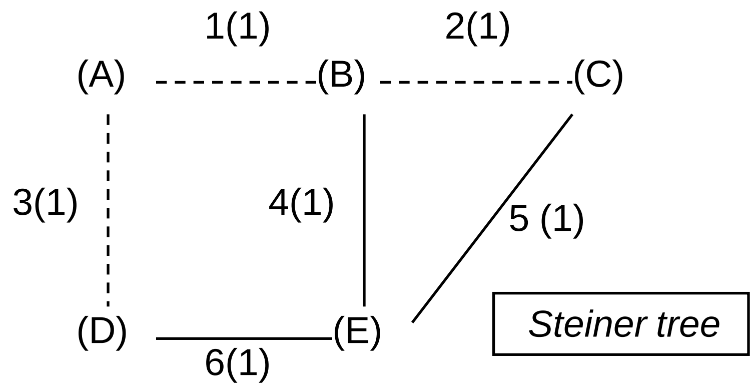

- Drawbacks:
  - Hard to compute (N-P complete)
  - Hard to reconfigure as membership changes.

Multicast Routing
Steiner Trees (2)

- Group A, D, E, and C

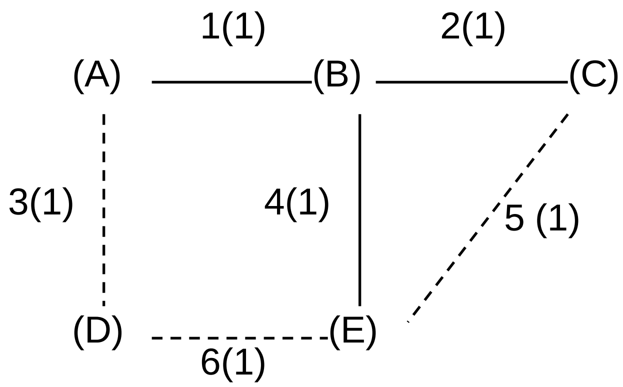

- Adding B to group A, D, E, and C

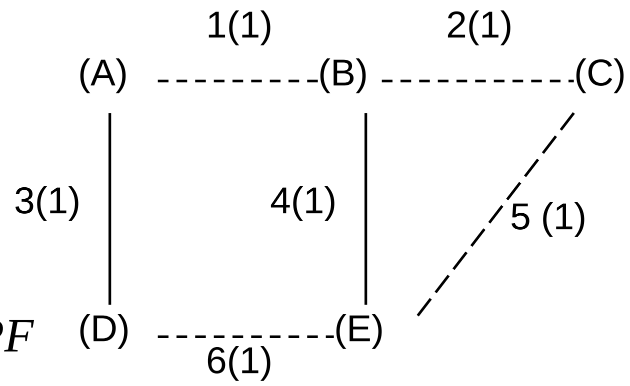

Same as in RPF

Multicast Routing
Core-based Trees (1)

- Introduced to overcome the drawbacks of RPF with Pruning
  - Flooding
  - States for all active groups & sources.
- Choose a *core*: a fixed point in the network serving as the center of the multicast group.
  - Recipients send a “JOIN” towards core.
  - Intermediate routers record interfaces on which JOIN messages are received.
  - If first join, forward it up to the core.

Multicast Routing
Core-Based Trees (2)

- Originators of multicast packets need not belong to group.
  - Packets are forwarded toward the core.
  - As soon as packets reach first router belonging to the tree, multicasting of packets proceed along the CBT tree similar to spanning tree forwarding.
- CBT builds a (spanning) tree for each group, same for all sources.

Multicast Routing
Core-Based Trees (3)

- Advantages of CBT over RPF:
  - More efficient:
    - One state per group instead of state per group and source.
    - No flooding, packets are sent to participants only.
    - Packet forwarding does not rely on routing tables
- Drawbacks:
  - CBT is sub-optimal
  - Traffic concentration (use multiple cores!)