


ELEN 4017


Network Fundamentals


Lecture 28



Purpose of lecture

- Network layer
 - Broadcast and multicast routing






Broadcast /Multicast

- So far we have considered point-to-point routing – **unicast**.
- In **broadcast routing**, network layer provides a service to send a packet from a source to all other nodes on the network.
- In **multicast routing**, a source node sends a packet to a **subset** of other network nodes.

Broadcast routing algorithms



- Most **straight-forward way** to accomplish broadcasting is for sending node **to send a separate copy** of the packet to each destination.
- This N-way unicast approach is simple, since no network layer routing protocol, packet duplication or forwarding functionality is needed.
- Drawbacks:
 - **Inefficiency** – if sending node is connected via a single link, then N separate packets will traverse this link.
- More sensible **if the network node creates the duplicate** (at router)





Figure 4.44 ♦ Source-duplication versus in-network duplication




N-way-unicast

- What other drawback ?
- For N-way-unicast **the sender must have the address of all the destinations.**
- To obtain this there is more overhead (additional protocol mechanisms /lists)
- **Link state protocols use broadcast to distribute link state info** that is used to compute unicast routes.
- Thus it is unwise to use unicast routes to do this broadcast.



Uncontrolled flooding

- A **source node sends a copy of packet** to all its **neighbours**.
- When a node receives a broadcast packet, it duplicates it and forwards it to all its neighbours (except the neighbour from which it received the packet)
- What is the flaw in this scheme?
- **If the graph has cycles**, then one or more copies of each broadcast will cycle indefinitely.
- **If more than two nodes are connected**, the number of packets circulating can increase, leading to a broadcast storm.




Controlled flooding

- To avoid the storm, a node must choose when to flood a packet and when not to.


Sequence number controlled flooding

- Source node puts its address and a broadcast **sequence number** into broadcast packet.
- Each node maintains a list of source address and sequence number of each broadcast packet it has received, duplicated and forwarded.
- Thus if the packet is in the list it is dropped.
- Other techniques exist which don't require router to store the broadcast packet info.






Spanning tree broadcast


- Sequence number flooding avoids the problem of broadcast storms, but we **still have redundant broadcasts in the network.**
- Ideally every node should **receive only one copy of a broadcast packet.**
- Consider the network topology shown
- If we forward packets only along the bold (blue) paths, then each node receives the broadcast packet only once.



Spanning tree




- Thus if we have a network and can arrange it into a **spanning tree**, then we can ensure that the broadcast does not contain redundant packets.
- Spanning tree → no connections between branches




Spanning tree

- A spanning tree of $G = (N, E)$ is a graph $G' = (N, E')$ such that:
 - E' is a subset of E
 - G' is connected
 - G' contains no cycles
 - G' contains all the nodes in G .




Minimum spanning tree

- Each link has an associated cost.
- we define the cost of a tree, as the sum of all the link costs.
- For a given network, multiple spanning tree graphs are possible.
- A spanning tree whose cost is the minimum of all the graphs of spanning trees is called a **minimum spanning tree**.




Benefits of a spanning tree

- When a node receives a broadcast message it forwards it to all its neighbours in the spanning tree.
- Thus the spanning tree eliminates redundant broadcasts.
- Importantly, once it is in place, it can be used by **any node** to initiate a broadcast.
- Note that the node does not need to know the entire tree, only its neighbours.



a. Broadcast initiated at A



b. Broadcast initiated at D



Figure 4.46 ♦ Broadcast along a spanning tree

How to create the spanning tree?




- Many algorithms exist. We consider the centre-based approach.
- A centre-node (rendezvous point/core) is chosen.
- Nodes then unicast **tree-join** messages to centre node.
- A tree-join message is forwarded using uni-cast routing toward the centre until it either arrives at a node belonging to the spanning tree or arrives at the centre.
- In either case, the **path followed by the tree join message defines the branch of the spanning tree.**

Example: centre based approach




a. Stepwise construction of spanning tree




b. Constructed spanning tree

**Node E is chosen as centre node.
Other nodes unicast tree join messages towards E.
A unicast message is forwarded to the centre until arrives at a node already belonging to tree, or the centre node.**

Broadcast algorithms in practice




- A form of sequence number controlled flooding is used in OSPF.
- P2P applications can also employ broadcast mechanisms.
- Gnutella also uses a variant of sequence number controlling.




Multicast routing

- Delivery of packet to a **subset of nodes**.
- Examples of applications requiring multi-cast:
 - Deploying software upgrades to a group of users.
 - Shared data applications (net meeting with whiteboard)
 - Interactive gaming
- 2 immediate problems:
 - How to identify receivers of packet
 - How to address these packets

Addressing



- For unicast there was a single recipient.
- For broadcast all nodes received packet, so there was no need to address any of them
- One approach is to include all destination addresses in each multicast packet → not scalable.
- Furthermore, it may be difficult for the sender to know all the destination addresses.



Address indirection

- A single identifier is used for a group of addresses.
- In the Internet the single identifier is a group D multicast IP address.
- The group of receivers associated with that multicast address is called a **multicast group**.

Multicast group

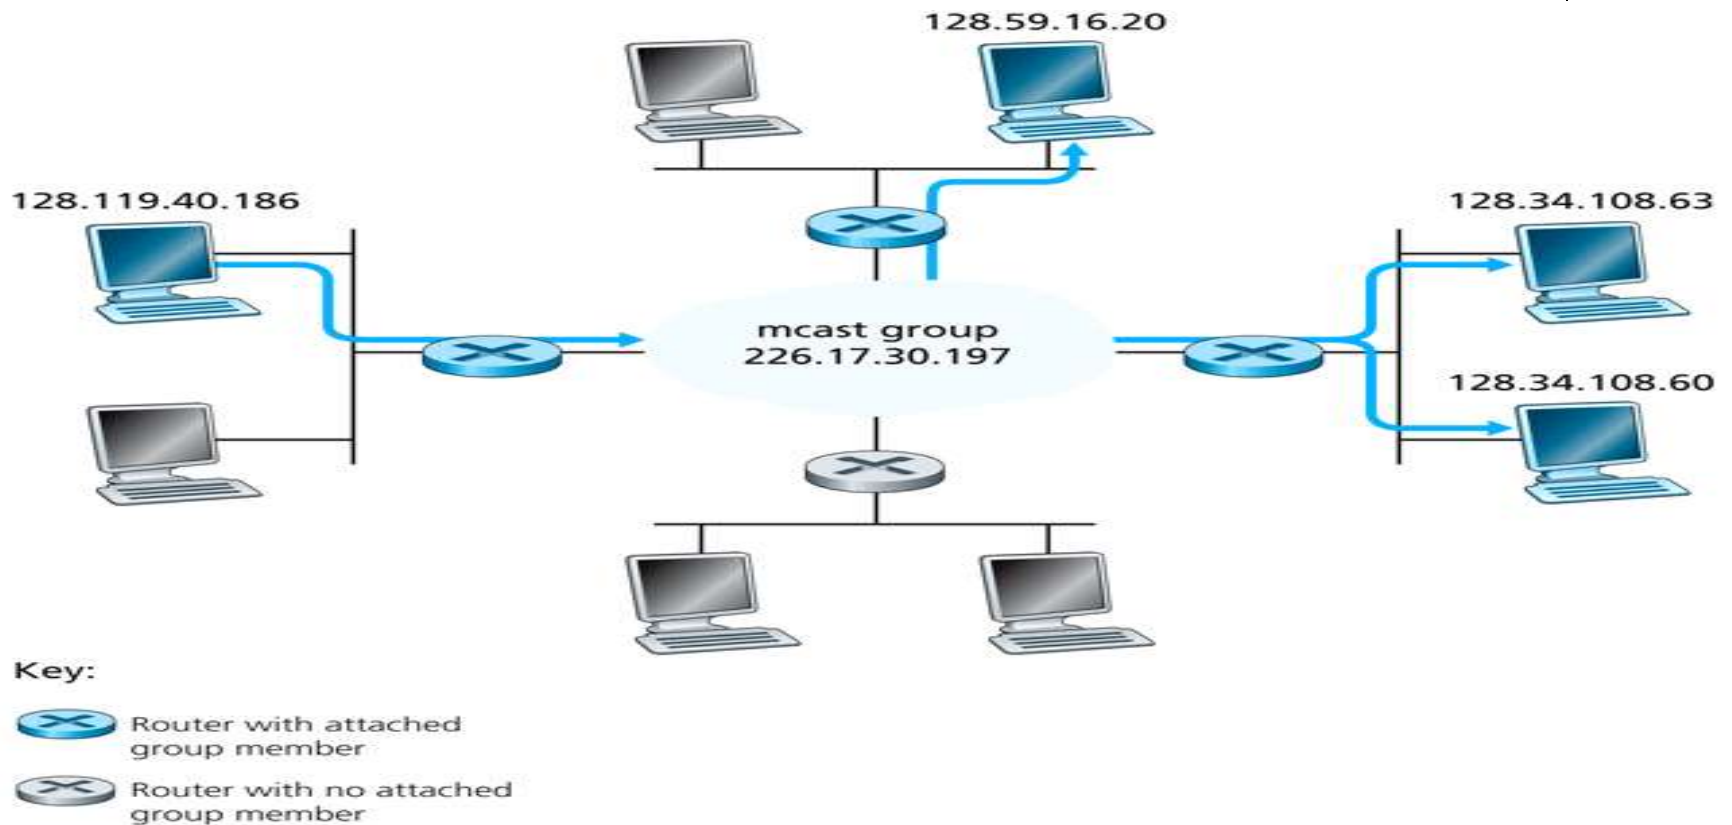




Figure 4.48 ♦ The multicast group: A datagram addressed to the group is delivered to all members of the multicast group.



Issues raised by multicasting

- How does a group start and terminate?
- How is group address chosen?
- How are hosts added /removed?
- Is membership restricted?
- Do members of a group know the addresses of their peers?
- Internet Group Management Protocol (IGMP) provides some answers.
- We wont go into the details.