CpSc 360: Distributed and Network Programming

Broadcasting & Multicasting

James Wang

Chapter 20, 21

Unicast, Broadcast, and Multicast

- **Unicast - "Sending a packet to a single destination host." TCP and UDP.**
- **Broadcast - "A method of delivering a packet to every host on a particular network or internet. May be implemented in hardware (e.g., Ethernet) or software (e.g., IP Broadcast)." UDP only!**
- **Multicast - "A special form of broadcast in which packets are delivered to a specified subgroup of network hosts." (As above can be hardware or software). UDP only!**

2

Protocol Support

- **Multicasting support is optional in IPv4, but mandatory in IPv6.**
- **Broadcasting support is not provided in IPv6. Any IPv4 application that uses broadcasting must be recoded for IPv6 to use Multicasting instead.**
- **Broadcasting and multicasting require datagram transport such as UDP or raw IP; they can not work with TCP.**

3

IP Broadcast Addresses

- **IP Broadcast Addresses**
  - Limited Broadcast: 255.255.255.255, never forwarded by router, used in host configuration if host does not know its IP address or subnet mask.
  - Subnet Directed Broadcast: specific subnet ID but all ones host address, e.g., IP = 165.91.216.50, subnet mask = 255.255.252.0 (22 bits network portion and 10 bits host portion) has a subnet directed broadcast address = 165.91.219.255.
  - All Subnets Directed: obsolete, use multicast instead. Both subnet ID and host address all ones.
  - Net Directed: host ID of all ones, router must forward, but option to disable.
  - On Ethernet all of the above map to all ones 48-bit

4

Unicast of a UDP datagram

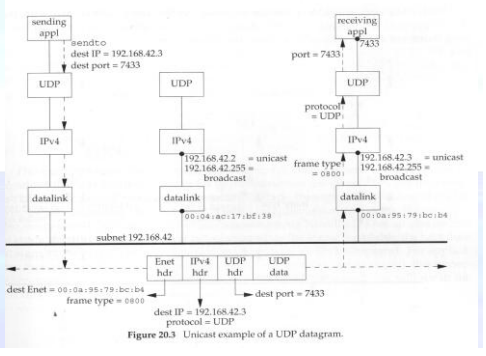

Figure 20.3 Unicast example of a UDP datagram.

5

What happened in unicast?

- **When the IP layer processes the packet, the packet is accepted only if the destination address matches one of its own IP addresses.**
- **Then the IP layer looks at the protocol field in the IPv4 header and the datagram is passed to UDP if the protocol is UDP.**
- **Then the UDP layer looks at the destination port (and possibly the source port, too, if the UDP socket is connected), places the datagram onto the appropriate socket receive queue.**
- **A unicast IP datagram is received by only the one host specified by the destination IP address.**

6

Broadcast of UDP datagram

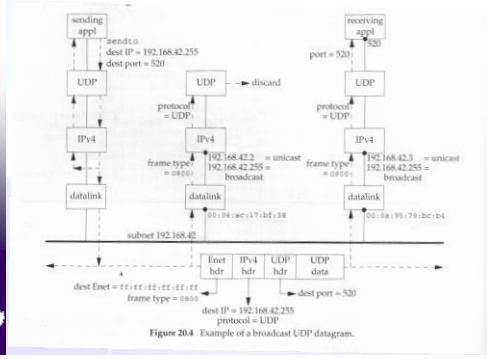

Figure 20.4 Example of a broadcast UDP datagram.

What happened in broadcast?

- When the host broadcasts the datagram, it notices that the destination IP address is the subnet-directed broadcast and maps into proper Ethernet address.
- The hosts on the subnet all have IP addressed matching the destination address, so datagram is passed to UDP layer when the protocol field is 17.
- Nothing special needs to be done by an application to receive a broadcast UDP datagram: it just creates a UDP socket and binds the application's port number to the socket.

IP Multicast Addresses

- Class D IP addresses 224.0.0.0 - 239.255.255.255

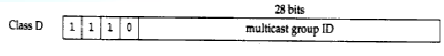

- Mapping Class D IP Address to Ethernet Address

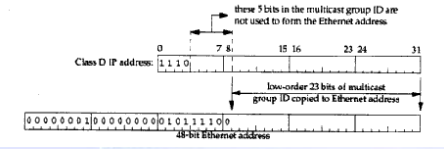

Scope of IP Multicast Addresses

- IPv4
  - Historically, IPv4 TTL field has doubled as multicast scope (i.e., how far do we forward the packet): TTL=0 (node local), TTL=1 (link local), 1<TTL<=32 (site local), 32<TTL<=64 (region local), 64<TTL<=128 (continent), TTL>128 (global).
  - Administrative scoping is proposed Internet draft: address range 239.0.0.0 -- 239.255.255.255 has sections of this address space for link local, site local, organization local, and global.
- IPv6 has a 4-bit scope field.

Multicast of UDP Datagram

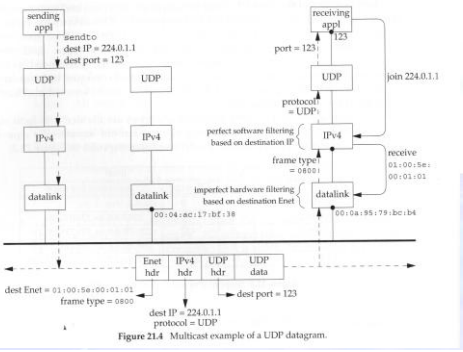

Figure 21.4 Multicast example of a UDP datagram.

IP Multicast

- IP Multicast address can only be used as a destination address. It supports UDP only.
- No ICMP error messages are returned (e.g., destination unreachable, source quench, echo reply, or time exceeded).
- One to many and many to many communications.
- Key technology for multiway conferences.
- Supported in Ethernet, FDDI, and Token Ring
- Experimental MBONE today, vBNS today, full Internet support in IPv6.
- Cisco, Bay Networks, and 3COM support multicasting in current router code.
- Win '95, Win NT, Apple OT support multicast.

Benefits of multicast routing

- **Resource Discovery**
- **Transmitting files to multiple destinations**
- **Audio/Video Conferences**

13

Multicast Distribution of a Hierarchical Layered Stream

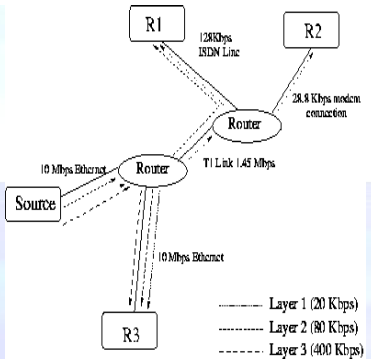

14

Multicast Routing Algorithm (1)

- **Flooding - simplest multicast routing algorithm**
  - Determine if this is the first time you have seen the packet being flooded. If so, transmit a copy on all outgoing links except the one you received it over
  - Not efficient in terms of bandwidth consumption but is robust (no routing tables required)
  - OSPF and earlier Arpanet routing algorithms distribute routing updates using flooding.
  - Usenet news floods to all news servers in a distribution group.

15

Multicast Routing Algorithm (2)

- **Spanning Trees**
  - "Connected subgraph of a network that includes all nodes and has no cycles" (i.e., loops)
  - Multicasting on a spanning tree is more efficient than flooding. Flooding may require less time, depending on the spanning tree selected
  - Used in IEEE 802.1 MAC Ethernet bridges.
  - Easy to add Multicast to Link State Routing (e.g., link OSPF) using spanning tree.

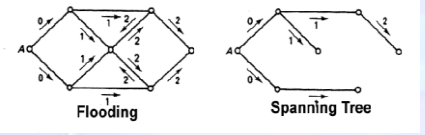

16

Multicast Routing Algorithm (3)

- **Reverse Path Forwarding (RPF)**
  - Multicast for Distance Vector routing (recall that a host does not know the entire topology but only the next hop).
  - Used in MBONE
- **RPF combines the following:**
  - Reverse Path Broadcast
  - Only forward along paths that have group members.

17

Reverse Path Broadcast

- **Since each router knows that the shortest path to a destination goes through the next hop router, the packet is forwarded on all outgoing links except the one it arrived on iff the packet came from the next hop router to S**
- **Reverse Path Broadcast**
  - 1. When multicast packet received, note source (**S**) and interface (**I**).
  - 2. If **I** is on the shortest path to **S**, forward the multicast pack on all interfaces except the one received.
  - 3. If **I** is not on the shortest path to **S**, discard packet.

18

Reverse Path Broadcast (continued)

- RPB results in a different spanning tree for each source (see fig).
- Fastest possible delivery since data follows shortest path from src to dest.
- Because different trees are computed for each source, packets are spread over multiple links.

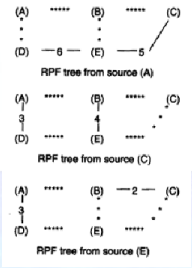

19

Problems with Reverse Path Broadcast

- If the links are not symmetric, we need to modify this procedure. Use shortest path *from* source to node rather than *to* source.
- "Floods" networks that don't have any members of multicast group (at least as far as TTL will carry packet).
- Multicast packet will be forwarded over a LAN by each router connected to that LAN

20

Reverse Path Forwarding

- Need to take group membership into account in Reverse Path Broadcast
  - "Non-pruning" version used initially on MBONE only tested group membership at *leaves* (if the parent router is the only router on the network, it is a *leaf* network). This kept un-subscribed multicast traffic off leaf network, but multicast traffic was flooded over the entire MBONE network only limited by TTL.
  - "Pruning" variant of RPF used in MBONE beginning Sept. 1993. A "prune message" is sent back up tree telling router not to send any more packets for this group. *Tree is pruned*.

21

Reverse Path Forwarding With Pruning

- Keep track of group membership, and only forward multicast packet down a tree if there is a group member.
- Flood and Prune:
  - First multicast packet to a group from a source is flooded to all network nodes.
  - If there are no members of a multicast group on a leaf network, the router for the leaf network sends to its parent a Non Membership Report (NMR). If the parent router finds that there are no members of a group on any of its downstream links, it sends and NMR to its parent router, and so on.
  - Soft-state for group membership with timer.

22

Multicast Tree: Parents, Children, and Leaves

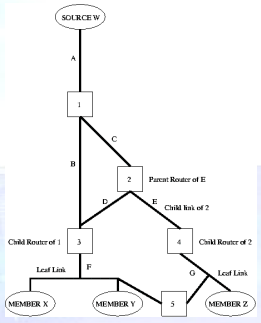

23

Pruning a Multicast Tree

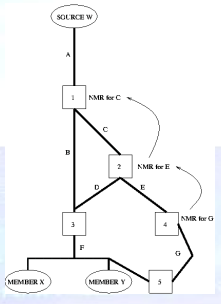

24

Core Based Trees (CBT)

- **Flood and prune method has disadvantages**
  - Requirement to keep state information for each active group and source pair leads to scaling problems (i.e., router memory).
  - Because of soft-state, there will be periodic removal of prunes that will then generate a multicast packet to parts of the network previously pruned. Flooding initial packet doesn't scale well.
- **CBT**
  - Solution to above problems, but no free lunch.
  - Suitable for sparsely distributed receivers.

25

Core Based Trees (continued)

- **CBT Algorithm**
  - Choose a fixed *core node* that will be the center of the group.
  - Potential group members send join messages to the center. Each intermediate node notes the interface the join request was received on and marks this interface as belonging to the group's tree (the router only has to keep one piece of state info per group). If this is the first join the router receives, it forwards the packet one hop towards the *core*.
  - Multicast packets from nonmember senders are forwarded to the tree center until they reach a node that already belongs to the tree.

26

Core Based Trees (continued)

- **Advantages of CBTs**
  - Builds a spanning tree per group, which is the same for all sources and only requires one piece of state information per group.
  - Limits the expansion of multicast transmissions to exactly the set of all recipients.
  - Doesn't depend on routing tables
- **Disadvantages of CBTs**
  - Choosing an optimal *core node* is NP complete.
  - Traffic congestion as traffic converges on core (might have multiple cores).
  - Paths between some source and destination may be suboptimal.

27