Lecture 18

Most high-level network protocols (such as the ISO Transport Protocols or TCP or UDP) only provide a **unicast** transmission service. That is, nodes of the network only have the ability to send to one other node at a time:

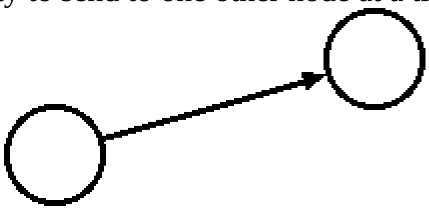

Fig. 1 - Basic Unicast Service

All transmission with a unicast service is inherently point-to-point. If a node wants to send the same information to many destinations using a unicast transport service, it must perform a **replicated unicast**, and send N copies of the data to each destination in turn.

A better way to transmit data from one source to many destinations is to provide a **multicast** transport service. With a **multicast** transport service, a single node can send data to *many* destinations by making just a single call on the transport service:

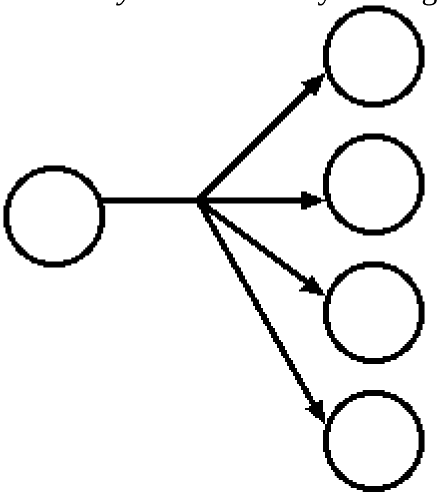

Fig. 2 - Multicast Transport Service

There are number of applications to multicasting

1. Stock Ticker: Stock ticker sends stocks updates to thousands of stations.
2. Teleconferencing: A group of users operating on different or same workstations form a multicast group such that transmission of messages from one user is received by all the users, for ex Yahoo Conference.
3. Database: With multiple copies of database being replicated, any update to any of the copy of the database, results in update of all the copies simultaneously, such that each copy resembles all other copy. For example Lotus Notes.

The notion of group is essential to the concept of multicasting. By definition a multicast message is sent from a source to a group of destination hosts. In IP multicasting, multicast groups have an ID called multicast group ID. Whenever a multicast message is sent out, a multicast group ID specifies the destination group. These group ID's are essentially a set of IP addresses called "Class D". Therefore, if a host (a process in a host) wants to receive a multicast message sent to a particular group, it needs to somehow listen to all messages sent to that particular group.

One half of the Ethernet addresses are reserved for multicast-the low-order bit of the high order octet distinguishes conventional unicast addresses (0) from a multicast addresses (1). In dotted hexadecimal notation, the multicast address is given by

$$(01.00.00.00.00.00)_{16}$$

And rests of the bits specify the different multicast address.

If the source and destination of a multicast message are on the same LAN, then it is quite easy to implement multicasting, since only one copy traverses the network and the station will accept any packet sent to the it's unicast address, multicast address or broadcast address.

Requirements for Multicasting

Since in unicast routing, a router receives the datagram or packet, and depending upon the destination address forwards the packet to the optimum route. But with multicast routing, router may have to generate multiple packets along some of the total routes available.

1. A convention is needed for identifying a multicast address. Class D addresses are for this specific purpose.

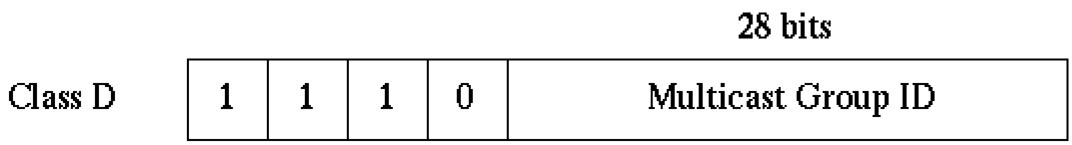

2. Each node must be able to translate between an IP multicast address and list of networks that contain members to the group.
3. Host may leave and join the multicast groups dynamically. Thus a mechanism is needed by which an individual station informs the routers attached to the same network as itself of its inclusion and exclusion from a multicast group.
4. Routers must exchange information. **First**, every router should know, which networks include members of given multicast group? **Second**, routers need sufficient information to calculate the shortest path to each network containing group members.
5. A routing algorithm is needed to calculate the shortest path.
6. Each router must determine multicast routing path on the basis of both source and destination address.

IGMP (Internet Group Management Protocol)

Hosts willing to receive multicast messages (packets) need to inform their immediately-neighboring routers that they are interested in receiving multicast messages sent to certain multicast groups. This way, each node can become a member of one or more multicast groups and receive the multicast packets sent to those groups. The protocol through which hosts communicate this information with their local routers is called Internet Group Management Protocol (IGMP).

The IGMP is also used by the routers to periodically check whether the known group members are still active. In case there is more than one multicast router on a given subnetwork (LAN), one of the routers is elected as the "querier" (The router with the lowest IP address is elected as the querier.) and assumes the responsibility of keeping track of the membership state of the multicast groups which have active members on its subnetwork. Based on the information obtained from the IGMP the router can decide whether to forward multicast messages it receives to its subnetwork(s) or not.

After receiving a multicast packet sent to a certain multicast group, the router will check and see if there is at least one member of that particular group on its subnetwork. If that is the case the router will forward the message to that subnetwork. Otherwise, it will discard the multicast packet.

The fundamental concept related to multicasting is the stations joining and leaving multicast groups. The IGMP provides a method through which a host can join or leave a multicast group. IGMP version 1 was defined in RFC 1112. IGMP which is considered a part of the IP layer has a fixed-size message with no optional data. The format of an IGMP message is shown in Figure 3.

0 3	4 7	8 15	16 31
Version	Type	Code	Checksum
Group Address			

Type=1, specifies a query sent by a multicast router.

Type=0, specifies a query sent by host.

Each host (process) can join a multicast group or leave a multicast group that it previously joined. IGMP messages are used by routers to keep track of group memberships in their immediately connected subnetwork. The following rules apply

1. A host sends an IGMP "report" for joining a group
2. A host will never send a report when it wants to leave a group.
3. Multicast routers send IGMP queries (to the all-hosts group address: 224.0.0.1) periodically to see whether any group members exists on their subnetworks. If no

response is received for a particular group after a number of queries, the router assumes that there is not any group member on the network (physically connected to a particular interface of the router). It should be noted that the TTL field of query messages is set to 1 so that the queries do not get forwarded to other subnetworks.

Based on the reports a router receives from the hosts it can decide whether to forward a multicast packet on a particular interface or not.

Multicasting Routing Algorithms

Flooding

The Flooding algorithm which has been already used in other routing protocols such as OSPF is the simplest technique for delivering the multicast datagrams to the routers of an internetwork. In this algorithm, when a router receives a multicast packet it will first check whether it has seen this particular packet earlier or this is the first time that this packet has reached this router. If this is the first time, the router will forward the packet on all interfaces, except the one from which the packet has been received. Otherwise, the router will simply discard the packet. This way we make sure that all routers in the internetwork will receive at least one copy of the packet.

Although this algorithm is pretty simple, it has some major disadvantages. The flooding algorithm generates a large number of duplicated packets and wastes the network bandwidth. Furthermore, since each router needs to keep track of the packets it has received in order to find out whether this is the first time that a particular packet has been seen or not, it needs to maintain a distinct entry in its table for each recently seen packet. Therefore, the Flooding algorithm makes inefficient use of router memory resources.

Spanning Trees

A better algorithm than Flooding is the Spanning Tree algorithm. This algorithm which has been already used by IEEE-802 MAC bridges is powerful and easy to implement. In this algorithm, subsets of internetwork links are selected to define a tree structure (loop-less graph) such that there is only one active path between any two routers. Since this tree spans to all nodes in the internetwork it is called spanning tree. Whenever a router receives a multicast packet, it forwards the packet on all the links which belong to the spanning tree except the one on which the packet has arrived, guaranteeing that the multicast packet reaches all the routers in the internetwork. Obviously, the only information a router needs to keep is a Boolean variable per network interface indicating whether the link belongs to the spanning tree or not. We use a small network with five nodes and six links to show different trees. For simplicity sake, we do not differentiate between hosts and routers, subnets and links. We also assume that links are symmetric and their costs are shown next to the links. The spanning tree from source node (C) is shown in Figure 4:

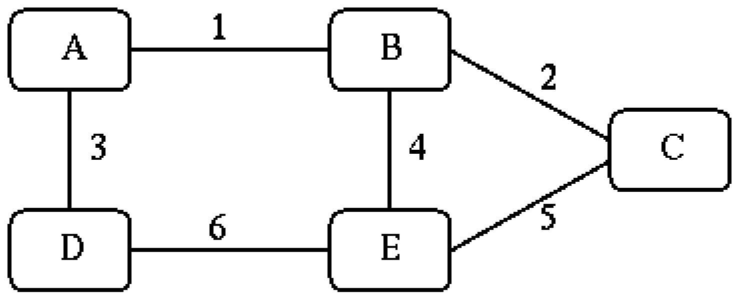

A small test network

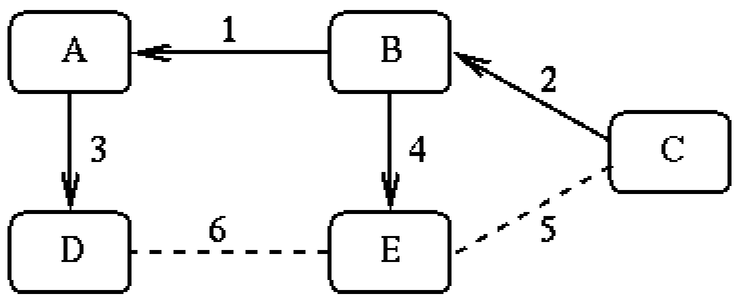

Spanning Tree from source (C)

The spanning tree algorithm has two drawbacks: It centralizes all the traffic on a small set of links and it does not consider the group membership.

Group Membership with IPv6

IGMP was defined for operation with IPv4 and it uses 32-bit addresses. Whereas, IPv6 uses addresses of size 128-bit. So to make IPv6 addresses compatible with IGMP, a new set of protocol is introduced called as ICMPv6, which incorporates all the functionality of ICMP and IGMP.