Broadcasting and Multicasting

Three kinds of IP addresses: unicast, broadcast, and multicast. Broadcasting and multicasting only apply to UDP, where it makes sense for an application to send a single message to multiple recipients. TCP is a connection-oriented protocol that implies a connection between two hosts (specified by IP addresses) and one process on each host (specified by port numbers).

Consider a set of hosts on a shared network such as an Ethernet. Each Ethernet frame contains the source and destination Ethernet address (48-bit values). Normally each Ethernet frame is destined for a single host. The destination address specifies a single interface – called a *unicast*. In this way communication between any two hosts doesn't bother any of the remaining hosts on the cable (except for possible contention for the shared media).

There are times, however, when a host wants to send a frame to every other host on the cable – called a *broadcast* (reference: ARP and RARP). Multicasting fits between incasting and broadcasting: the frame should be delivered to a set of hosts that belong to a multicast group.

To understand broadcasting and multicasting we need to understand that filtering takes place on each host, each time a frame passes by on the cable. Figure shows a picture of this.

First, the interface card sees every frame that passes by on the cable and makes a decision whether to receive the frame and deliver it to the device driver. Normally the interface card receives only those frames whose destination address is either the hardware address of the interface or the broadcast address. Additionally, most interfaces can be placed into a *promiscuous* mode whereby they receive a copy of every frame. This mode is used by *tcpdump,* for example.

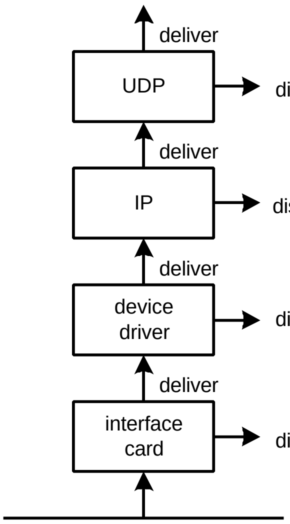

Today most interfaces can also be configured to receive frames whose destination address is a multicast address, or some subset of multicast addresses. On an Ethernet, a multicast address has the low-order bit of the high-order byte turn on. In hexadecimal this bit looks like 01:00:00:00:00:00. (We can consider the Ethernet broadcast address, FF:FF:FF:FF:FF:FF as a special case of the Ethernet multicast address.)

If the interface card receives the frame, it is passed to the device driver. (One reason the interface card might discard the frame is if the Ethernet checksum is incorrect). Additional filtering is performed by the device driver. First, the frame type must specify a protocol that is supported (IP, ARP, etc.). Second, additional multicast filtering may be performed, to check whether the host belongs to the addressed multicast group.

The device driver then passes the datagram up to the next layer, such as IP, if the frame type specifies an IP datagram. IP performs more filtering, based on the source and destination IP addresses, and passes the datagram up to the next layer (such as TCP or UDP) if all is well.

Each time UDP receives a datagram from IP, it performs filtering based on the destination port number, and sometimes the source port number too. If no process is currently using the destination port number, the datagram is discarded and an ICMP port unreachable message is normally generated. (TCP performs similar filtering based on its port numbers.) If the UDP datagram has a checksum error, UDP silently discard it.

The problem with broadcasting is the processing load that it places on hosts that aren't interested in the broadcasts. Consider an application that is designed to use UDP broadcasts. If there are 50 hosts on the cable, but only 20 are participating in the application, every time one of the 20 sends a UDP broadcast, the others 30 hosts have to process broadcast, all way up through the UDP layer, before the UDP datagram is discarded. The UDP datagram is discarded by these 30 hosts because the destination port number is not in use.

The intent of multicasting is to reduce this load on hosts with no interest in the application. With multicasting a host specifically joins one or more multicast groups. If possible, the interface card is told which multicast groups the host belongs to, and only those multicast frames are received.

Broadcasting

Limited Broadcast

The limited broadcast address is 255.255.255.255. This can be used as the destination address of an IP datagram during the host configuration process, when the host might not know its subnet mask or even its IP address.

A datagram destined for the limited broadcast address is never forwarded by a router under any circumstance. It only appears on the local cable.

An unanswered question is: if a host is multihomed and a process sends a datagram to the limited broadcast address, should the datagram be sent out each connected interfaces that support broadcasting? If not, an application that wants to broadcast out all interfaces must determine all the interfaces on the host that support broadcasting, and send a copy out each interface.

Most BSD systems treat 255.255.255.255 as an alias for the broadcast address of the first interface that was configured, and don't provide any way to send a datagram out all attached, broadcast-capable interfaces. The net-directed broadcast address corresponding to that interface is then used as the destination address for datagrams sent out the interface.

Net-directed Broadcast

The net-directed broadcast address has a host ID of all one bits. A class A net-directed broadcast address is netid.255.255.255, where netid is the class A network ID.

A router must forward a net-directed broadcast by default, but it must also have an option to disable this forwarding.

Subnet-directed Broadcast

The subnet-directed broadcast address has a host ID of all one bits but a specific subnet ID. Classification of an IP address as a subnet-directed broadcast address requires knowledge of the subnet mask. For example, if a router receives a datagram destined for 128.1.2.255, this is a subnet-directed broadcast if the class B network 128.1 has a subnet mask of 255.255.255.0, but it is not a broadcast if the subnet mask is 255.255.254.0.

All-subnets-directed Broadcast

An all-subnet-directed broadcast address also requires knowledge of the destination network's subnet mask, to differentiate this broadcast address from a net-directed broadcast address. Both the subnet ID and the host ID are all one bits. For example, if the destination's subnet mask is 255.255.255.0, then the class B IP address 128.1.255.255 is an all-subnets-directed broadcast. But if the network is not subnetted, this is a net-directed broadcast.

Multicasting

IP multicasting provides two services for an application.

1. Delivery to multiple destinations. There are many applications that deliver information to multiple recipients: interaction conferencing and dissemination of mail or news to multiple recipients, for example. Without multicasting these types of services tend to use TCP today (delivering a separate copy each destination). Even with multicasting, some of these applications might continue to use TCP for its reliability.
2. Solication of servers by clients. A diskless workstation, for example, needs to locate a bootstrap server. Today this provided using a broadcast, but a multicast solution would impose less overhead on the hosts that don't provide the service.

Multicast Group Address

Figure shows the format of a class D IP address.

					28 bits
Class D	1	1	1	0	multicast group ID

Unlike the other three classes of IP address (A, B, and C), which we showed in Figure, the 28 bits allocated for the multicast group ID have no further structure.

A multicast group address is the combination of the high-order 4 bits of 1110 and the multicast group ID. These are normally written as dotted-decimal numbers and are in the range 244.0.0.0 through 239.255.255.255.

The set of host listening to a particular IP multicast address is called a host group. A host group can span multiple networks. Membership in a host group is dynamic – host may join and leave host groups at will. There is no restriction on the number of hosts in a group, and a host does not have to belong to a group to send a message to that group.

Some multicast group addresses are assigned a well-known address by the IANA (Internet Assigned Numbers Authority). These are called permanent host groups. This is similar to the well-known TCP and UDP port numbers. Similarly, these well-known multicast addresses are listed Assigned Numbers RFC. Notice that it is the multicast address of the group that is permanent, not the membership of the group.

For example, 224.0.0.1 means 'all systems on this subnet', and 224.0.0.2 means 'all routers on this subnet'. The multicast address 224.0.1.1 is for NTP, the Network Time Protocol, 224.0.0.9 is for RIP-2, and 224.0.1.2 is for SGI's (Silicon Graphics) dogfight application.

Converting Multicast Group Address to Ethernet Address

The IANA owns an Ethernet of address block, which in hexadecimal is 00:00:5E. This is the high-order 24 bits of the Ethernet address, meaning that this block includes address in the range 00:00:5E:00:00:00 through 00:00:5E:FF:FF:FF. The IANA allocates half of this block for multicast addresses. Given that the first byte of any Ethernet address must be 01 to specify a multicast address, this means the Ethernet addresses corresponding to IP multicasting are in the range 01:00:5E:00:00:00 through 01:00:5E:7F:FF:FF.

This allocation allows for 23 bits in the Ethernet address to correspond to the IP multicast group ID. The mapping places the low-order 23 bits of multicast group ID into these 23 bits of the Ethernet address. This is shown in Figure.

Since the upper 5 bits of the multicast group ID are ignored in this mapping, it is not unique. Thirty-two different multicast group IDs map to each Ethernet address. For example, the multicast address 224.128.64.32 (hex E0.80.40.20) both map into the Ethernet address 01:00:5E:00:40:20.

Since the mapping is not unique, it implies that the device driver or the IP module in Figure must perform filtering, since the interface card may receive multicast frames in which the host is really not interested. Also, if the interface card doesn't provide adequate filtering of multicast frames, the device driver may have to receive all multicast frames, and perform the filtering itself.

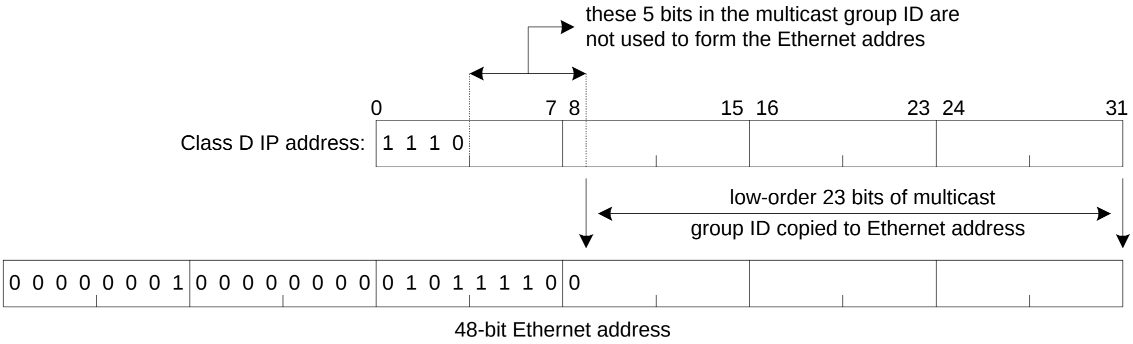

Figure: Mapping of a class D IP address into Ethernet multicast address.

Multicasting on a single physical network is simple. The sending process specifies a destination IP address that is a multicast address, the device driver converts this to the corresponding Ethernet address, and send its. The receiving processes must notify their IP layers that they want to receive datagrams destined for a given multicast address, and the device driver must somehow enable reception of these multicast frames. This is called "joining a multicast group". When a multicast datagram is received by a host, it must deliver a copy to all the processes that belong to that multicast group. This is different from UDP where a single process receives an incoming unicast UDP datagram. With multicasting it is possible for multiple processes on a given host to belong to the same multicast group.

But complications arise when we extend multicasting beyond a single physical network and pass multicast packets through routers. A protocol is needed for multicast routers to know if any hosts on a given physical network belong to a given multicast group. This protocol is called the ***Internet Group Management Protocol (IGMP)***.