CS4254

Computer Network Architecture and Programming

Dr. Ayman A. Abdel-Hamid

Computer Science Department

Virginia Tech

Multicasting – Part I

Outline

- Multicasting (Chapter 21)
  - Multipoint Communications
  - IP Multicast
  - IGMP
  - Multicast Routing
  - IPv4 Multicast addresses
  - Sending and Receiving Multicast Messages

Multipoint Communications

- Multipoint communications support communications between more than two hosts
  - One-to-many
  - Many-to-many
- Unlike broadcast, allows a proper subset of hosts to participate
- Example standards
  - IP Multicast (RFC 1112, standard)

Logical Multipoint Communications

- Two basic *logical* organizations
  - Rooted: hierarchy (perhaps just two levels) that structures communications
  - Non-rooted: peer-to-peer (no distinguished nodes)

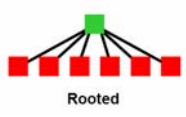

- Different structure could apply to control and data "planes"
  - *Control plane* determines how multipoint session is created
  - *Data plane* determines how data is transferred between hosts in the multipoint session

Logical Multipoint Communications

Control Plane

•The control plane manages creation of a multipoint session

- ➢*Rooted control plane*
  - ✓One member of the session is the root, *c_root*
  - ✓Other members are the leafs, *c_leafs*
  - ✓Normally *c_root* establishes a session
    - ❑Root connects to one or more *c_leafs*
    - ❑*c_leafs* join *c_root* after session established
- ➢*Non-rooted control plane*
  - ✓All members are the same (*c_leafs*)
  - ✓Each leaf adds itself to the session

Logical Multipoint Communications

Data Plane

The data plane is concerned with data transfer

- •*Rooted data plane*
  - ➢Special root member, *d_root*
  - ➢Other members are leafs, *d_leafs*
  - ➢Data transferred between *d_leafs* and *d_roots*
    - ✓*d_leaf* to *d_root*
    - ✓*d_root* to *d_leaf*
  - ➢There is no direct communication between *d_leafs*
- •*Non-rooted data plane*
  - ➢No special members, all are *d_leafs*
  - ➢Every *d_leafs* communicate with all *d_leafs*

Forms of Multipoint Communications

•Server-based -- rooted multipoint communications with server as *d_root*

- ➢Passive or inactive
  - ✓Relay
  - ✓Reflector
- ➢Active
  - ✓Bridge or multipoint control unit (MCU)

•Strictly peer-to-peer multipoint – Non-rooted

- ➢Multicast

Multipoint Servers

Passive Multipoint Server

•a relay or reflector service

•Provides no processing of the data

•Minimum requirement is for transport-level semantics, so can operate at the transport or application level

Active Multipoint Server

•Does application-level processing

- ➢transcoding

•uses application-level semantics

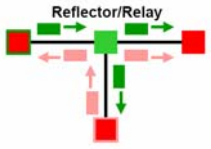

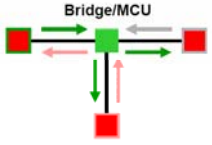

Multicast Communication

•Multicast abstraction is peer-to-peer

➢*Application-level multicast*

➢*Network-level multicast*

✓Requires router support (multicast-enabled routers)

✓Multicast provided at network protocol level → IP multicast

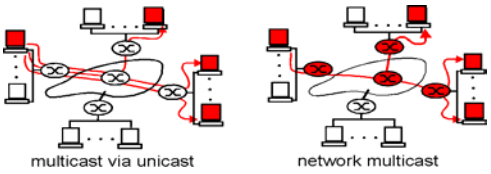

Multicast Communication

•Transport mechanism and network layer must support multicast

•Internet multicast limited to UDP (not TCP)

➢*Unreliable*: No acknowledgements or other error recovery schemes (perhaps at application level)

➢*Connectionless*: No connection setup (although there is routing information provided to multicast-enabled routers)

➢*Datagram*: Message-based multicast

IP Multicast

•IP supports multicasting

➢Uses only UDP, not TCP

➢Special IP addresses (Class D) identify multicast groups

➢*Internet Group Management Protocol* (IGMP) to provide group routing information

➢Multicast-enabled routers selectively forward multicast datagrams

➢IP TTL field limits extent of multicast

•Requires underlying network and adapter to support broadcast or, preferably, multicast

➢Ethernet (IEEE 802.3) supports multicast

IP Multicast: Group Address

•How to identify the receivers of a multicast datagram?

•How to address a datagram sent to these receivers?

➢Each multicast datagram to carry the IP addresses of all recipients? → Not scalable for large number of recipients

➢ *Use address indirection*

✓A single identifier used for a group of receivers

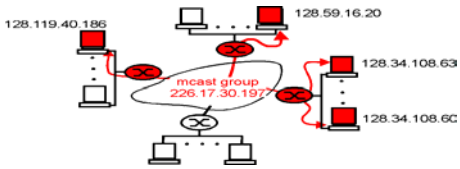

IP Multicast: IGMP Protocol

•RFC 3376 (IGMP v3): *operates between a host and its directly attached router*

•host informs its attached router that an application running on the host wants to join or leave a specific multicast group

•another protocol is required to coordinate multicast routers throughout the Internet → *network layer multicast routing algorithms*

•Network layer multicast → IGMP and multicast routing protocols

•IGMP enables routers to populate multicast routing tables

•Carried within an IP datagram

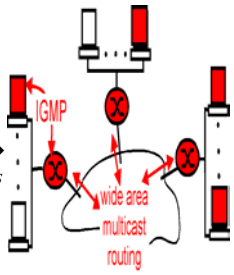

IP Multicast: IGMP Protocol

IGMP v2 Message types

•*membership query: general*

- •Sent by routers → router query multicast groups joined by attached hosts

•*membership query: specific*

- •Sent by routers → query if specific multicast group joined by attached hosts

•*membership report*

- •Sent by host → report host wants to join or is joined to given multicast group

•*leave group (optional)*

- •Sent by host → report leaving given multicast group

IP Multicast: IGMP Protocol

•Joining a group

- ➢Host sends group report when the first process joins a given group
- ➢Application requests join, service provider (end-host) sends report

•Maintaining routing table at the router

- ➢Multicast router periodically queries for group information
- ➢Host (service provider) replies with an IGMP report for each group
- ➢Host does not notify router when the last process leaves a group → this is discovered through the lack of a report for a query

IP Multicast: Multicast Routing

•Multicast routers *do not* maintain a list of individual members of each host group

•Multicast routers *do* associate zero or more host group addresses with each interface

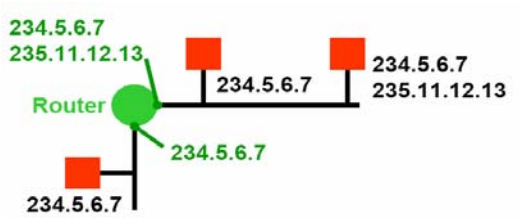

IP Multicast: Multicast Routing

•Multicast router maintains table of multicast groups that are active on its networks
•Datagrams forwarded only to those networks with group members

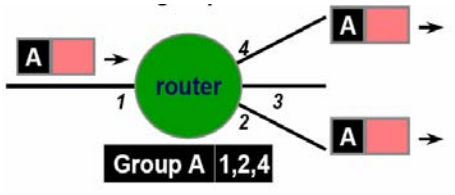

IP Multicast: Multicast Routing

•How multicast routers route traffic amongst themselves to ensure delivery of group traffic?
- ➢*Find a tree of links that connects all of the routers that have attached hosts belonging to the multicast group*
  - ✓Group-shared trees
  - ✓Source-based trees

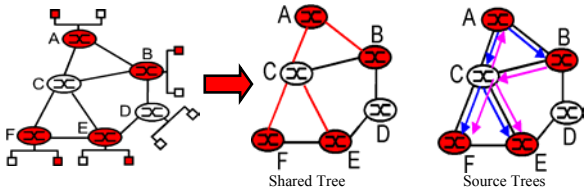

MBONE: Internet Multicast Backbone

•The MBone is a virtual network on top of the Internet (section B.2)
- ➢Routers that support IP multicast
- ➢IP tunnels between such routers and/or subnets

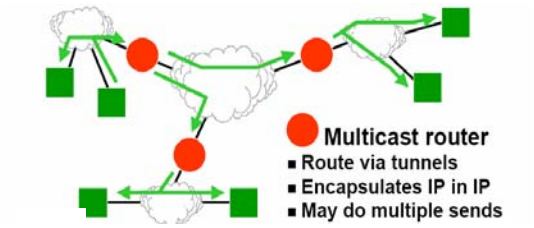

Unicast versus Broadcast versus Multicast

•A unicast address identifies a single IP interface

•A broadcast address identifies all IP interfaces on the subnet

•A multicast address identifies a set of IP interfaces

•A multicast datagram is received only by those interfaces interested in the datagram (applications wishing to participate in the multicast group)

IPv4 Multicast Addresses 1/3

•Class D addresses in the range 224.0.0.0 through 239.255.255.255

•Low order 28 bits of class D address (see appendix A) form the multicast group ID (32-bit address is the group address)

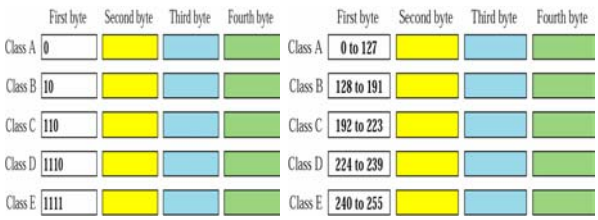

IPv4 Multicast Addresses 1/3

•Mapping of IPv4 multicast address to Ethernet address [RFC 112]

- ➢High-order 24 bits of Ethernet address are always 01:00:5E
- ➢Next bit always 0
- ➢Low-order 23 bits are copied from low-order 23 bits of multicast group address
- ➢High-order 5 bits of group address are ignored in the mapping → 32 multicast addresses map to a single Ethernet address
- ➢Mapping not one-to-one

IPv4 Multicast Addresses 2/3

224.0.1.88 mapped into an Ethernet address?

•Remember an Ethernet address is 48 bits

•The address 224 is E0 in hex, 0 is 00 in hex, 1 is 01 in hex, and 88 is 58 in hex. However, only the low-order 23 bits are used

•Therefore, the IP address of 224.0.1.88 converted to a MAC address is 01-00-5E-00-01-58.

IPv4 Multicast Addresses 3/3

•224.0.0.0 — 224.0.0.255 reserved for routing, topology discovery, maintenance protocols

- ➢Not forwarded by routers (TTL = 1)
- ➢224.0.0.0 reserved
- ➢224.0.0.1 all-host group
- ➢224.0.0.2 all-routers group

•224.0.0.0 — 232.255.255.255 assigned (RFC 1700, http://www.ietf.org/rfc/rfc1700.txt?number=1700)

•239.000.000.000 — 239.255.255.255 are "administratively scoped" (RFC 2365)

- ➢239.192.000.000 — 239.251.255.255 organization-local scope
- ➢239.255.000.000 — 239.255.255.255 site-local scope (TTL < 32)

Sending & Receiving Multicast Messages

Receiving Multicast Messages

- Create a UDP socket
- Bind it to a UDP port, e.g., 1234
  - All processes must bind to the same port in order to receive the multicast messages
- Join a multicast group address
- Use *recv* or *recvfrom* to read the messages

Sending Multicast Messages

- You may use the same socket (you used for receiving) for sending multicast messages or you can use any other UDP socket (it does not have to join any multicast group)