CS551
Multicast Routing: IGMP

Bill Cheng

http://merlot.usc.edu/cs551-f12

Components of the IP Multicast Architecture

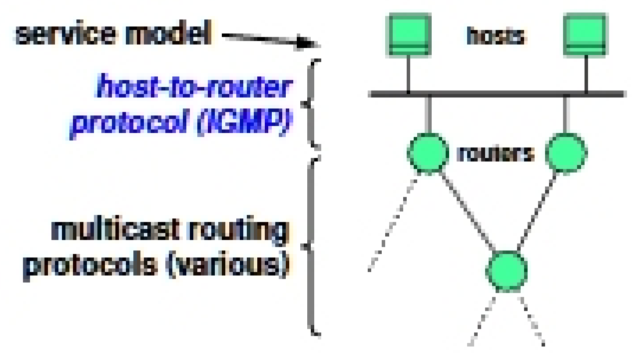

Internet Group Management Protocol (IGMP)

- the protocol by which hosts report their multicast group memberships to neighboring routers
- version 1, the current Internet Standard, is specified in RFC-1112
  - version 2: RFC 2236
- operates over broadcast LANs and point-to-point links
- occupies similar position and role as ICMP in the TCP/IP protocol stack

Link-layer Transmission/reception

- Transmission:
  - an IP multicast packet is transmitted as a link-layer multicast, on those links that support multicast
  - the link-layer destination address is determined by an algorithm specific to the type of link (next slide)
- Reception:
  - the necessary steps are taken to receive desired multicasts on a particular link, such as modifying address reception filters on LAN interfaces
  - multicast routers must be able to receive all IP multicasts on a link, without knowing in advance which groups will be sent to

Mapping to Link-layer Multicast Addresses

- for Ethernet and other LANs using 802 addresses:

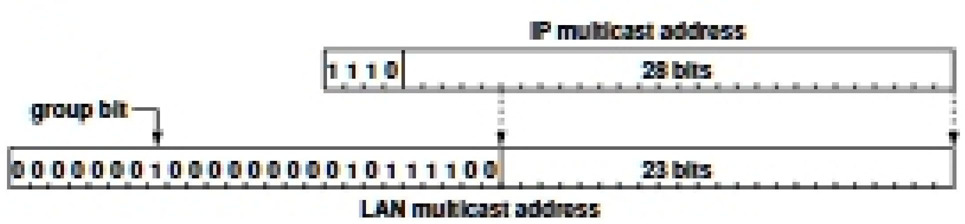

- for point-to-point links: no mapping needed

IGMP Version 1 Message Format

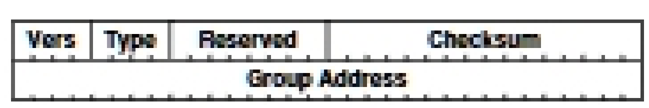

Version : 1

Type : 1 = Membership Query
2 = Membership Report

Checksum : standard IP-style checksum of the IGMP Message

Group Address : group being reported (zero in Queries)

IGMP Goal

- Determine what IP multicast groups have receivers present on the LAN
  - just care about some vs. zero receivers, not how many
- Approach
  - designate one router as IGMP "querier"
  - it asks all hosts
  - get at least one response per active group
  - example of *soft state* (periodically query), so occasional losses are okay

How IGMP Works

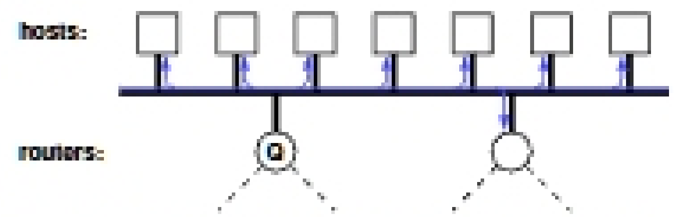

- on each link, one router is elected the "querier"
- querier periodically sends a *Membership Query* message to the all-systems group (224.0.0.1), with TTL = 1
- on receipt, hosts start random timers (between 0 and 10 seconds) for each multicast group to which they belong

How IGMP Works (Cont...)

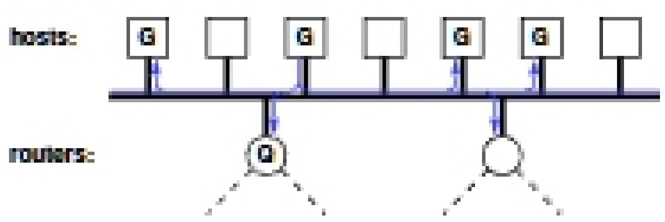

- when a host's timer for group G expires, it sends a *Membership Report to group G*, with TTL = 1
- other members of G hear the report and stop their timers
- routers hear *all* reports, and time out nonresponding groups

IGMP Implications

- In normal case, only one report message per group present is sent in response to a query (routers need not know who all the members are, only that members exist)
- Query interval is typically 60 – 90 seconds
  - IGMPv2 adds explicit leave messages
- To reduce *join latency*, when a host first joins a group, it sends one or two immediate reports (unsolicited responses), instead of waiting for a query

IGMP Version 2

- changes from version 1:
  - new message and procedures to reduce "leave latency"
  - standard querier election method specified
  - version and type fields merged into a single field
- backward-compatible with version 1
- soon to appear as a Proposed Standard RFC
- widely implemented already

CS551
Multicast Routing

Bill Cheng

http://merlot.usc.edu/cs551-f12

Components of the IP Multicast Architecture

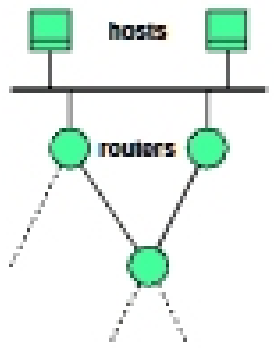

Multicast Routing

- Multicast service model makes it hard to locate receivers
  - anonymity
  - dynamic join/leave
- Options so far (not very efficient)
  - flood data packets to entire network, or
  - tell routers about all possible groups and receivers so they can create routes (trees)

Early Routing Techniques

- *Flood and prune*
  - begin by flooding traffic to entire network
  - prune branches with no receivers
  - *unwanted state where there are no receivers*
  - examples: DVMRP, PIM-DM
- Link-state multicast protocols
  - routers advertise groups for which they have receivers to entire network
  - compute trees on demand
  - *unwanted state where there are no senders*
  - examples: MOSPF

Rendezvous Options

- Specify *rendezvous* (or meeting place) to which sources send initial packets, and receivers join; requires mapping between multicast group address and meeting place
  - examples: CBT, PIM-SM

Multicast Tree Taxonomy

- Multicast routing can build different types of distribution trees
- *Source-based trees*
  - separate shortest path tree (SPT) *for each sender*
    - can have multiple senders per group
  - examples: DVMRP, MOSPF, PIM-DM, PIM-SM
- *Shared trees*
  - single tree shared by all members
  - shared tree rooted at group core/rendezvous point
  - examples: CBT, PIM-SM

Source-based Trees

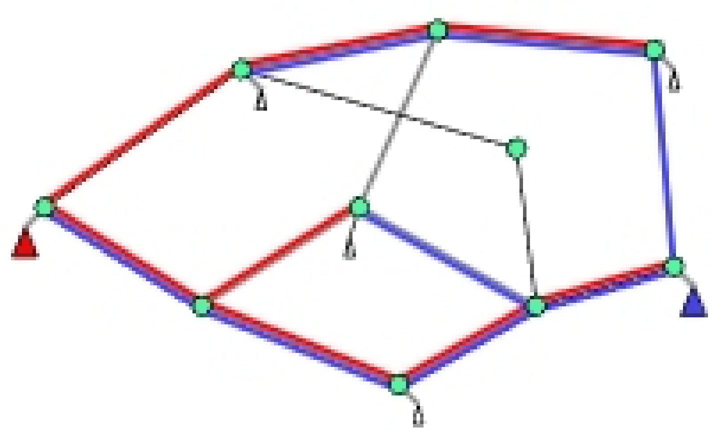

- output link determined from input link, multicast address, and *source* address