


CS 417 – DISTRIBUTED SYSTEMS

Week 4: Part 2
IP Multicast

Paul Krzyzanowski



© 2023 Paul Krzyzanowski. No part of this content may be reproduced or reposted in whole or in part in any manner without the permission of the copyright owner.

IP multicast routing

Deliver messages to a subset of nodes – send to a *multicast address*

- How do we identify the recipients?
 - Enumerate them in the header?
 - What if we don't know?
 - What if we have thousands of recipients?
- Use a **special IP address** to identify a group of receivers
 - A copy of the packet is delivered to all receivers associated with that group
 - IPv4: **Class D multicast IP address**
 - 32-bit address that starts with 1110 (224.0.0.0/4 = 224.0.0.0–239.255.255.255)
 - IPv6:
 - 128-bit address with high-order bits 8 bits all 1 ($\text{ff00:0:0:0:0:0:0:0}/8$)
 - **Host group** = set of machines listening to a particular multicast address
 - A copy of the message is delivered to all receivers associated with that group

IP multicasting

- Can span multiple physical networks
- Dynamic membership
 - Machine can join or leave at any time
- No restriction on number of hosts in a group
- Machine does not need to be a member to send messages
- Efficient: Packets are replicated only when necessary
- Like IP, no delivery guarantees – UDP at the transport layer

IP multicast addresses

Addresses chosen arbitrarily for an application

- Well-known addresses assigned by IANA

Internet Assigned Numbers Authority

IPv4 addresses: <http://www.iana.org/assignments/multicast-addresses/multicast-addresses.xml>

IPv6 addresses: <https://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xhtml>

- Similar to ports – service-based allocation
 - For ports, we have: FTP: port 21, SMTP: port 25, HTTP: port 80
 - For multicast, we have:

224.0.0.1:	all systems on this subnet
224.0.0.2:	all multicast routers on subnet
224.0.23.173:	Philips Health
224.0.23.52:	Amex Market Data
224.0.12.0-63:	Microsoft & MSNBC
FF02:0:0:0:0:0:9:	RIP routers

Internet Group Management Protocol (IGMP)

- Operates between a host and its attached router
- Goal: *allow a router to determine to which of its networks to forward IP multicast traffic*
- IP protocol (IP protocol number 2)

Three message types

1. Membership_query

Sent by a router to all hosts on an interface (i.e., on the LAN) to determine the set of all multicast groups that have been joined by the hosts on that interface

2. Membership_report

Host response to a query or an initial join or a group


3. Leave_group

- Host indicates that it is no longer interested
- Optional: router infers this if the host does not respond to a query

Multicast Forwarding

- IGMP allows a host to subscribe to *receive* a multicast stream
- *What about the sender?*
 - There is no protocol for the source!
 - It just sends one message to a class D address
 - Routers have to do the work

IGMP & Wide-Area Multicast Routing



Multicast Forwarding

- IGMP: Internet Group Management Protocol
 - Designed for routers to talk with hosts on directly connected networks
- PIM: Protocol Independent Multicast
 - Multicast Routing Protocol for delivering packets across routers
 - Topology discovery is handled by other protocols
 - Two forms:
 - Dense Mode (PIM-DM)
 - Sparse Mode (PIM-SM)

PIM-DM: Dense Mode Multicast – *flooding*

Forward multicast packet to all connected routers

- Each router forwards the multicast packet to all of its interfaces
- Feedback & cut off if there are no interested receivers on a link
 - A router sends a *prune* message.
 - Periodically, routers send messages to refresh the prune state
- **Flooding is initiated by the sender's router**
- Use a spanning tree and **reverse path forwarding (RPF)** to avoid loops
 - Packet is duplicated & forwarded ONLY IF it was received via the link that is the shortest path to the sender
 - Shortest path is found by checking the router's forwarding table to the source address

PIM-DM: Dense Mode Multicast – *flooding*

- **Advantage:**
 - Simple
 - Good if the packet is desired in most locations
- **Disadvantage:**
 - wasteful on the network, wasteful extra state & packet duplication on routers

PIM-SM: Sparse Mode Multicast

Initiated by the routers at each receiver

Each router requests a multicast feed with a PIM *Join* message

- Initiated by a router at the destination that gets an IGMP *join*
- Rendezvous Point: meeting place between receivers & source
 - *Join* messages propagate to a defined **rendezvous point**(RP)
 - Sender transmits only to the rendezvous point
 - RP announcement messages inform edge routes of rendezvous points
- A *Prune* message stops a feed

Advantage

- Packets go only where needed
- Creates extra state in routers only where needed
- Requires administration – setting up the RP & ensuring hosts can discover it

IP Multicast in use

- Initially exciting:
 - Internet radio, NASA shuttle missions, collaborative gaming
- But:
 - Few ISPs enabled it
 - For the user, required tapping into existing streams (not good for on-demand content)
 - Industry embraced unicast instead

IP Multicast in use: IPTV

- IPTV has emerged as the biggest user of IP multicast
 - Cable TV networks have migrated (or are migrating) to IP delivery
- Multicast allows one stream of data to be sent to multiple subscribers using a single address
 - Reduces the number of servers needed
 - Reduces the number of duplicate network streams
- IGMP from the set-top box
 - Change channel = join a new multicast group
- Use unicast for video on demand

The End