CSC2209
Computer Networks

Multicast

Stefan Saroiu
Computer Science
University of Toronto

Administrivia

- Thursday class canceled
  - No need to review paper
- No class next week because of OSDI
- Resume on Tuesday, November 14
  - Project mid-term presentations
- Progress report due on November 10th, 11:59 EST
  - Deadline is firm!!!!!

Progress Report

- Draft of your final report
- Start by motivating your problem
  - What is the problem?
  - Why is it interesting?
  - Why it is hard?
- Present what you have accomplished so far
  - This section(s) must have a summary where you list accomplishments bullet-by-bullet
- Brief section of what you expect to accomplish

Project Presentations

- Two options:

1. 5 min. presentations on Tuesday

2. 10 min. presentations on scheduled day
   - Reserve a 2-hour slot for 10 min. presentations

- After presentations, I'll e-mail you my high-level feedback about the project
  - While this will not count towards final grade, they "might" serve as good pointers

Slides Today

- Some of the slides have been adapted from Ramesh Govindan’s slides.

Unicast

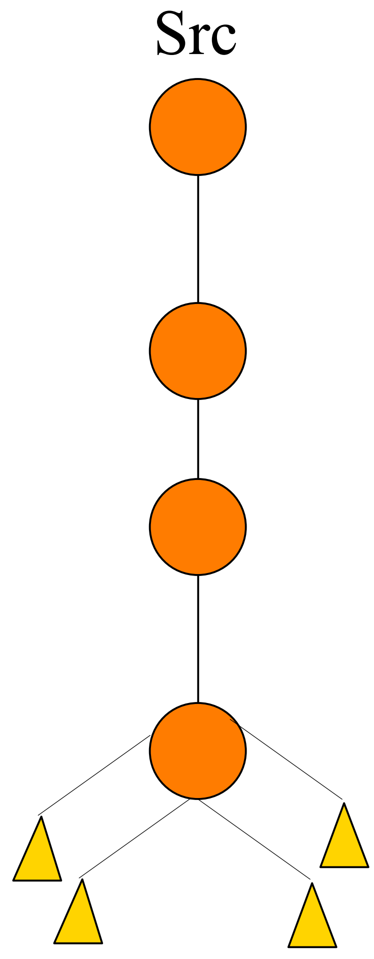

Unicast

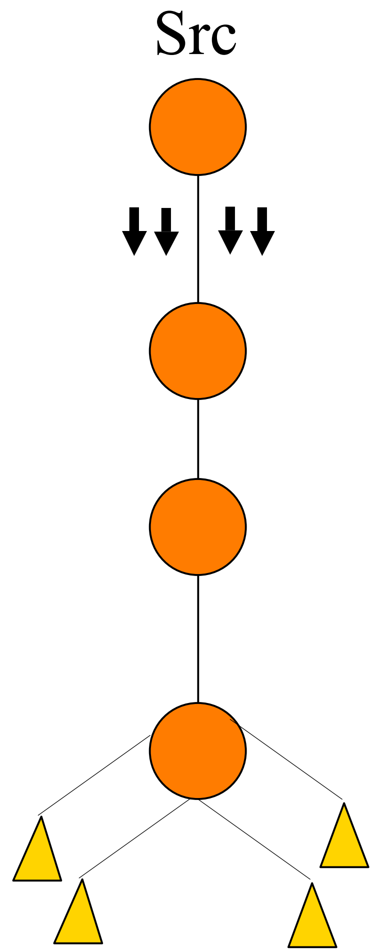

Unicast

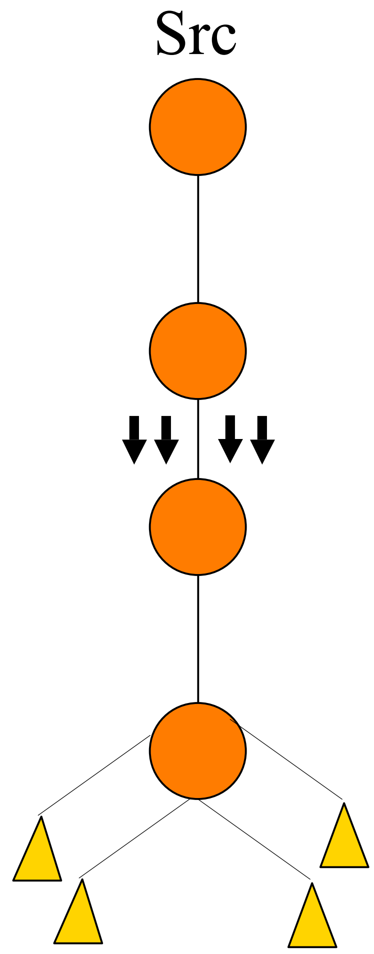

Unicast

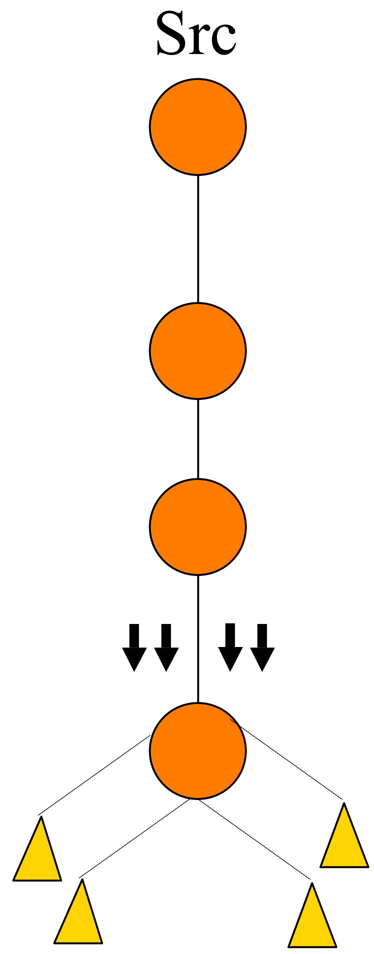

Unicast

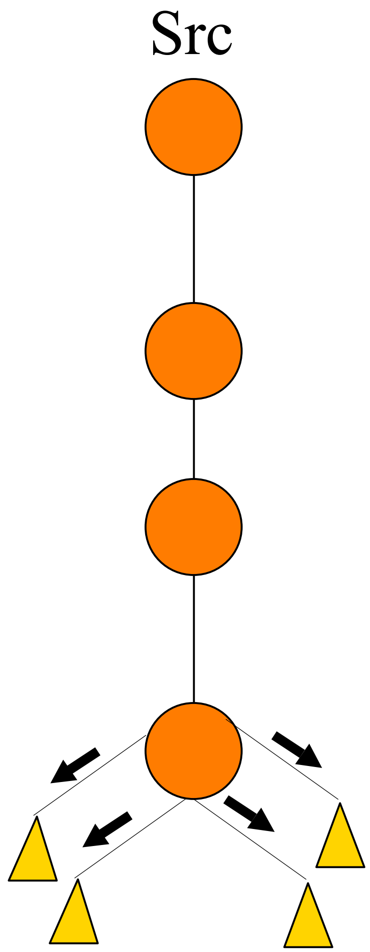

Unicast

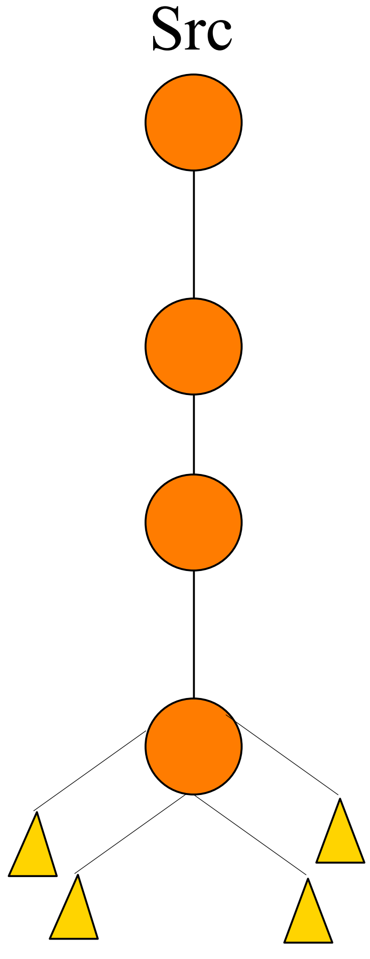

Multicast

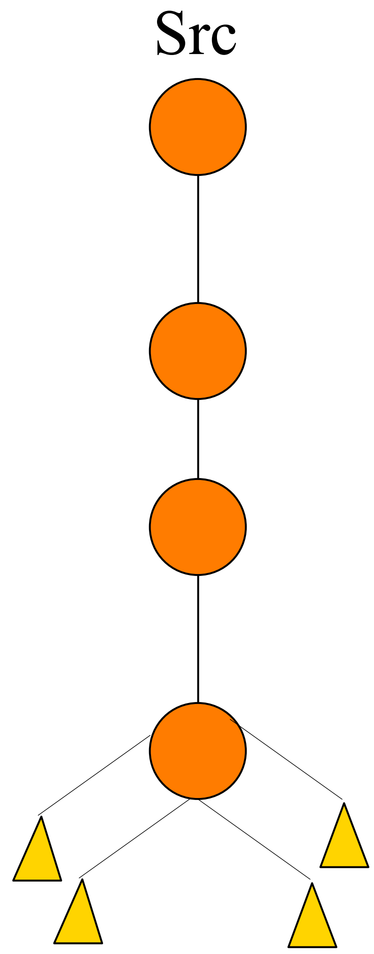

Multicast

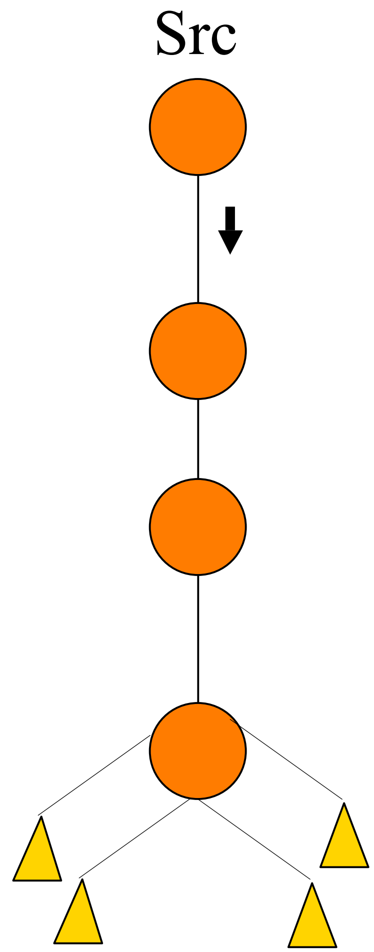

Multicast

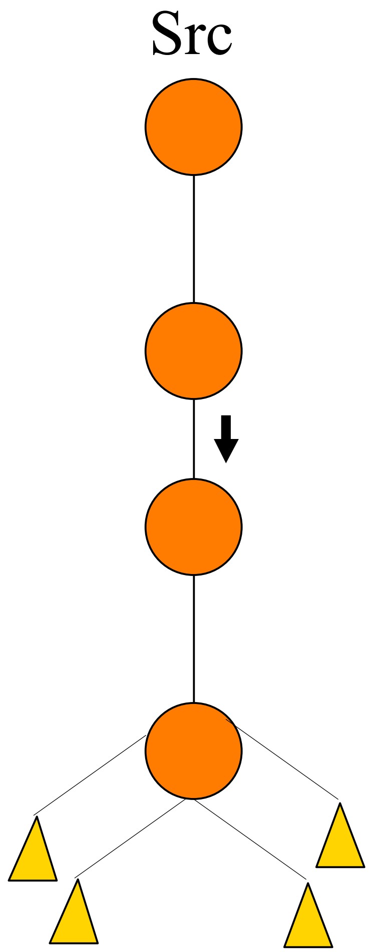

Multicast

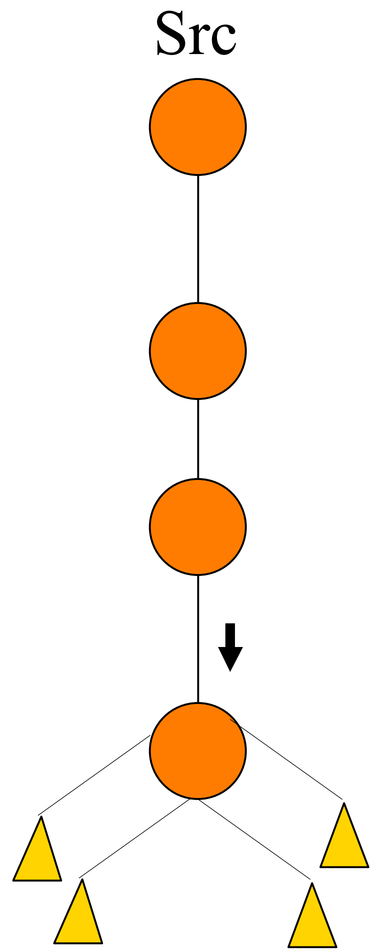

Multicast

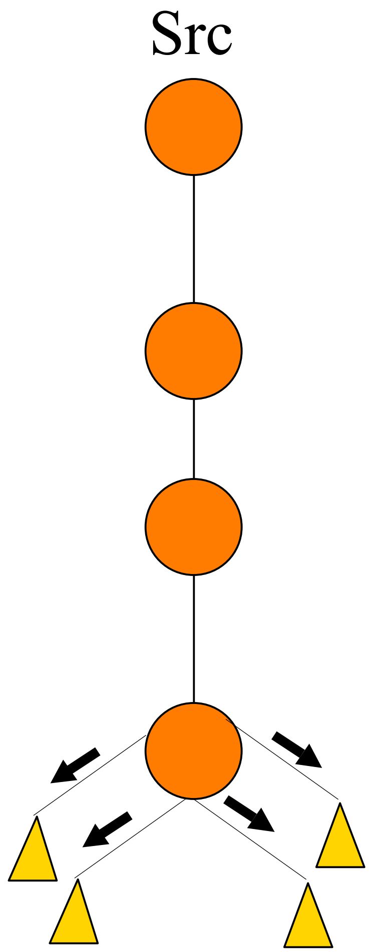

Multicast

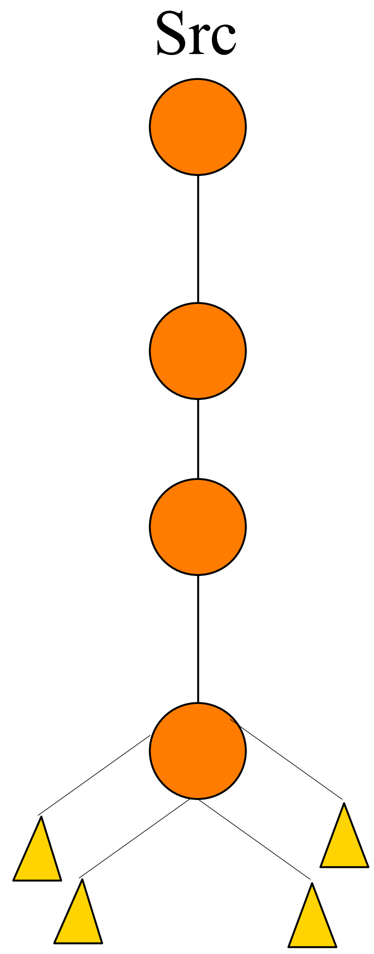

Multicast state

- Router:
  - learn of the existence of multicast groups (advertisement)
  - identify links with group members
  - establish state to route packets
    - replicate packets on appropriate interfaces

- Then, what is the router's state affected by?

Multicast state

- Router:
  - learn of the existence of multicast groups (advertisement)
  - identify links with group members
  - establish state to route packets
    - replicate packets on appropriate interfaces

- Then, what is the router's state affected by?
  - Host joins and leaves affect router state!

Example applications

- Broadcast audio/video
- Push-based systems
- Software distribution
- Web-cache updates
- Teleconferencing (audio, video, shared whiteboard, text editor)
- Multi-player games
- Server/service location
- Distributed apps -- file-sharing P2Ps?

Pondering….

- If so many applications exist, why hasn’t it taken off?

Pondering….

- If so many applications exist, why hasn't it taken off?
- Many reasons have been argued:
  - Don't know how to charge for multicast traffic
  - Require router state/changes (reluctant operators)
  - No "killer app"
  - Hard to debug

- Application-level multicast hasn't taken off either
  - (my take) Internet apps fall into one of the following three categories:
    - Communication (typically one-to-one)
    - Request-based content delivery
    - Content search (e.g., search engines)
  - Multicast model doesn't clearly fit any of these apps.

Components of the IP Multicast Architecture

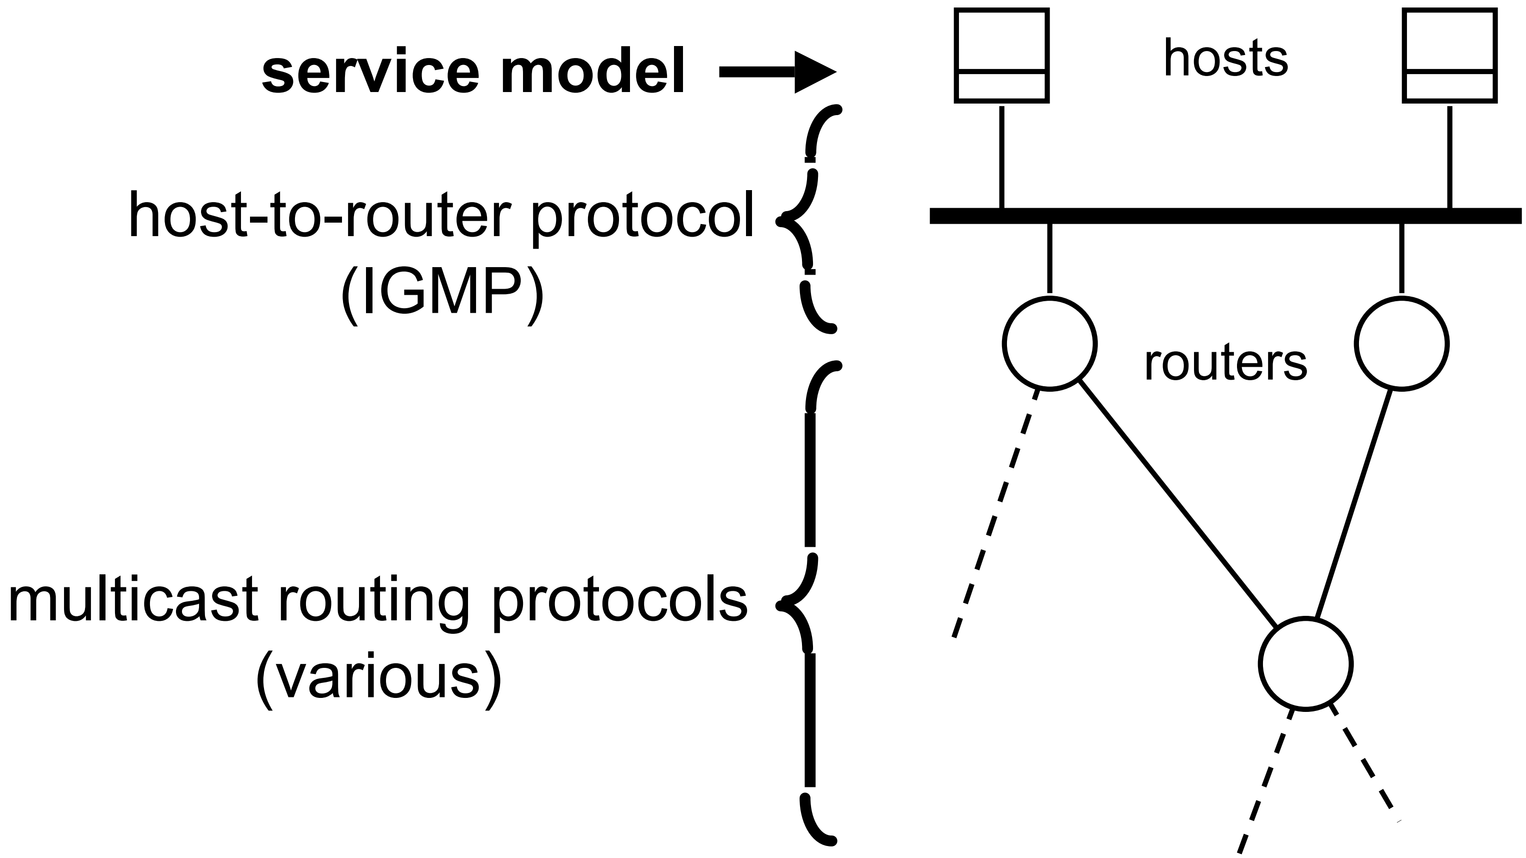

IP Multicast Service Model (RFC-1112)

- each group identified by a single IP address
- groups may be of any size
- members of groups may be located anywhere in the Internet
- members of groups can join and leave at will
- senders need not be members

Service model

- Group membership not known explicitly
- Analogy:
  - each multicast address is like a radio frequency, on which anyone can transmit, and to which anyone can tune-in.

IP Multicast Addresses

Class D IP addresses:

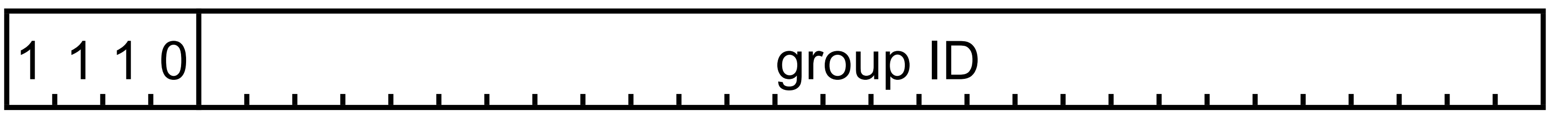

in “dotted decimal” notation: 224.0.0.0 — 239.255.255.255

two administrative categories:

- “well-known” multicast addresses, assigned by IANA
- “transient” multicast addresses, assigned and reclaimed dynamically

IP Multicast Service — Sending

- uses normal IP-Send operation, with an IP multicast address specified as the destination
- must provide sending application a way to:
  - specify outgoing network interface, if >1 available
  - specify IP time-to-live (TTL) on outgoing packet
  - enable/disable loopback if the sending host is a member of the destination group on the outgoing interface

IP Multicast Service — Receiving

- two new operations:

  Join-IP-Multicast-Group (group-address, interface)

  Leave-IP-Multicast-Group(group-address, interface)

- receive multicast packets for joined groups via normal IP-Receive operation

Multicast Scope Control: (1) TTL Expanding-Ring Search

to reach or find a nearby subset of a group

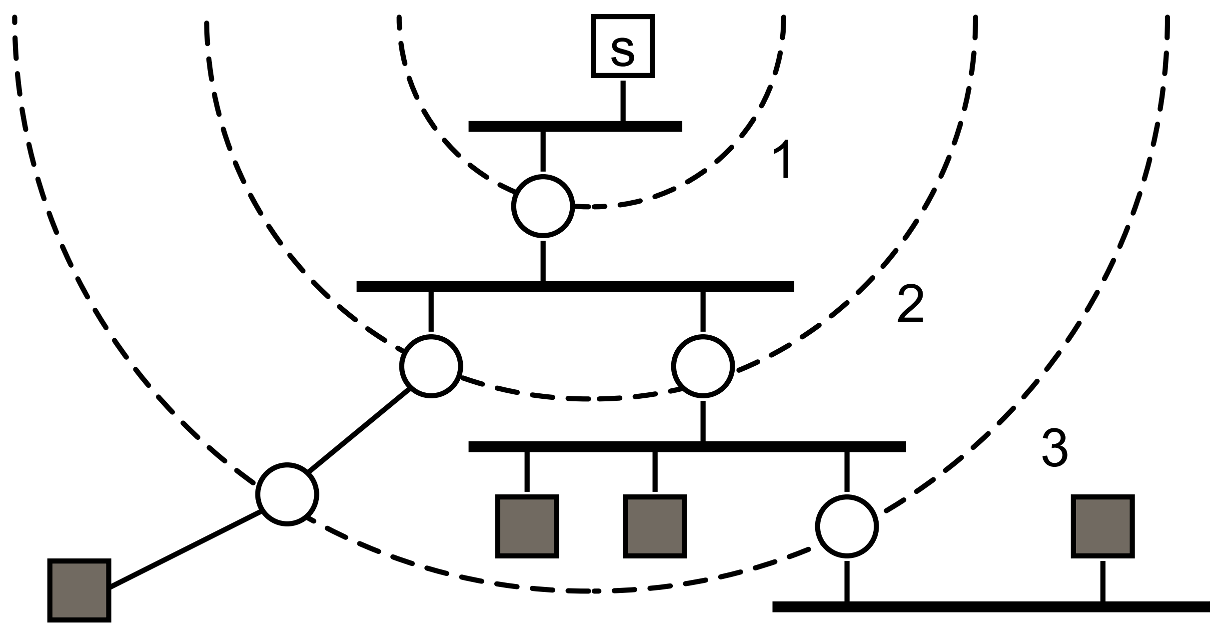

Multicast Scope Control:
(2) Administrative TTL Boundaries

to keep multicast traffic within an administrative domain, e.g., for privacy or resource reasons

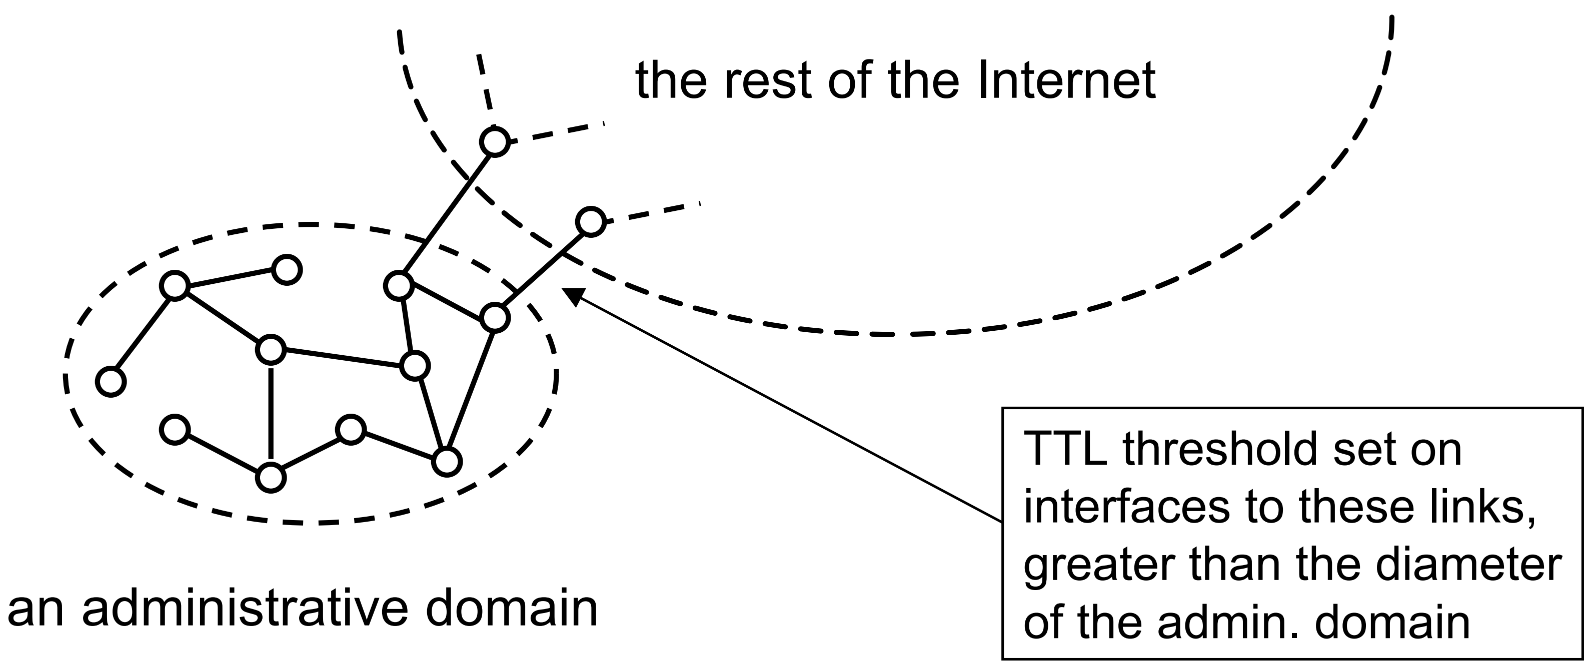

Multicast Scope Control: (3) Administratively-Scoped Addresses

- RFC 1112
- uses address range 239.0.0.0 — 239.255.255.255
- supports overlapping (not just nested) domains

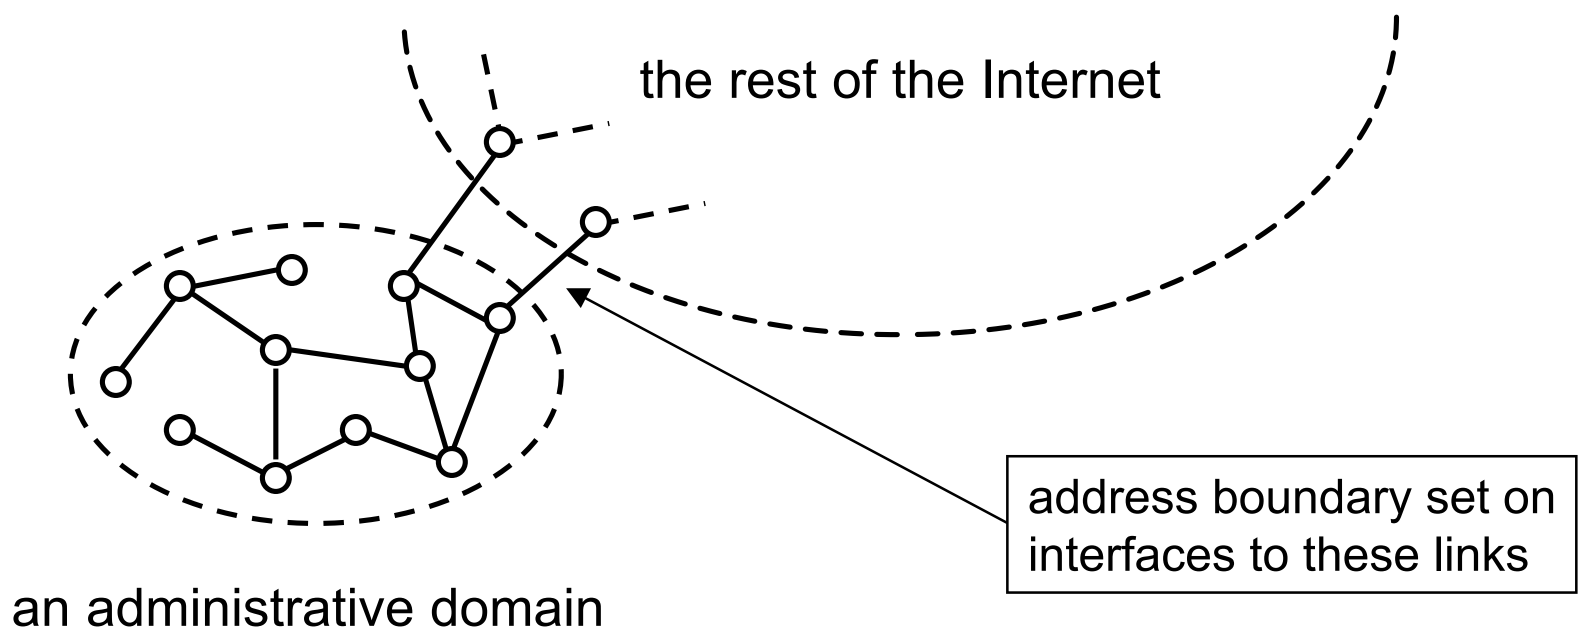

The MBone

- MBone = Multicast Backbone
- an "interconnected" set of multicast-capable routers, providing the IP multicast service in the Internet
- can be thought of as a virtual network, overlaid on the Internet

Components of the MBone

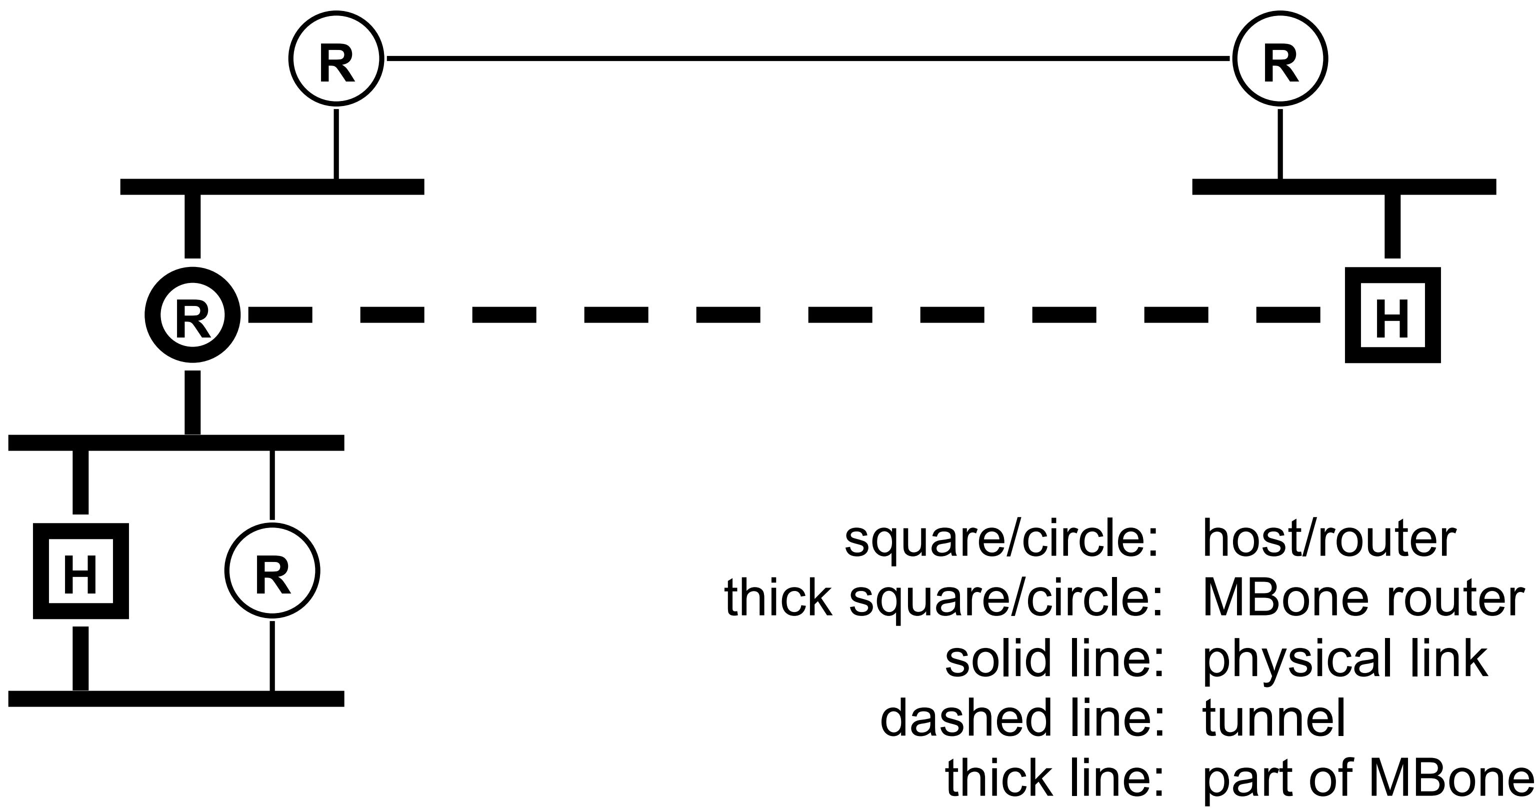

MBone Tunnels

- a method for sending multicast packets through multicast-ignorant routers
- IP multicast packet is encapsulated in a unicast packet addressed to far end of tunnel:

IP header, dest = unicast	IP header, dest = multicast	transport header and data…

- a tunnel acts like a virtual point-to-point link
- each end of tunnel is manually configured with unicast address of the other end

IGMP

- End system to router protocol is IGMP
- Member or process starts an application with mcast address
- IGMP process informed of joined mcast group
- When local router sends IGMP query, host sends IGMP report
- One router on LAN is IGMP querier
- Hosts listen to responses and supress duplicates

Multicast Routing: IGMP

Components of the IP Multicast Architecture

service model

host-to-router protocol (IGMP)

multicast routing protocols (various)

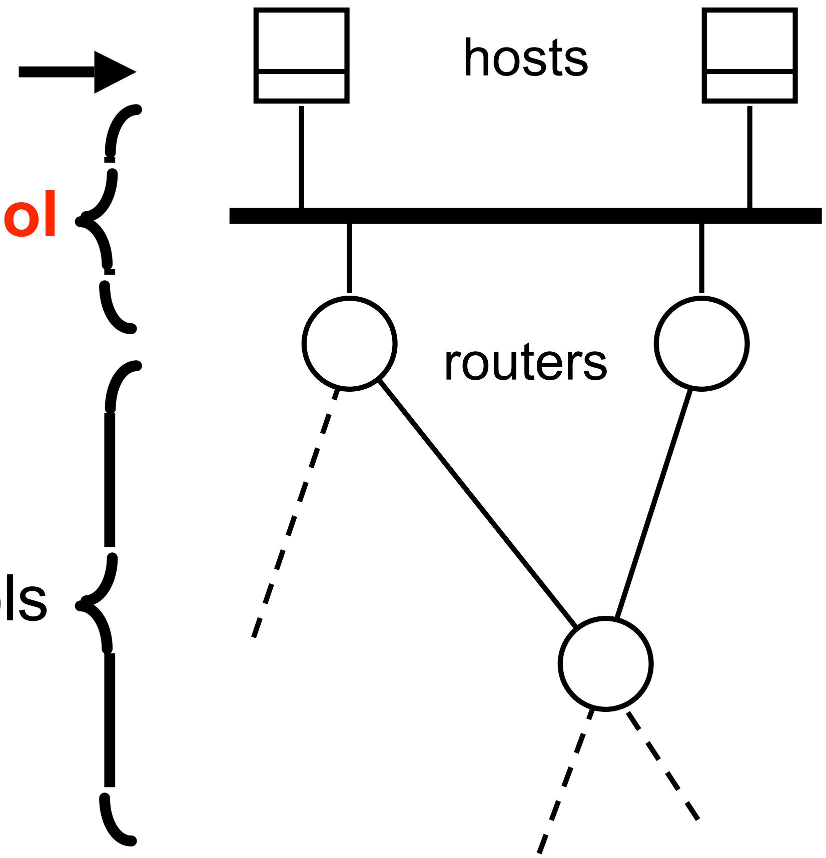

Internet Group Management Protocol(IGMP)

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

How IGMP Works

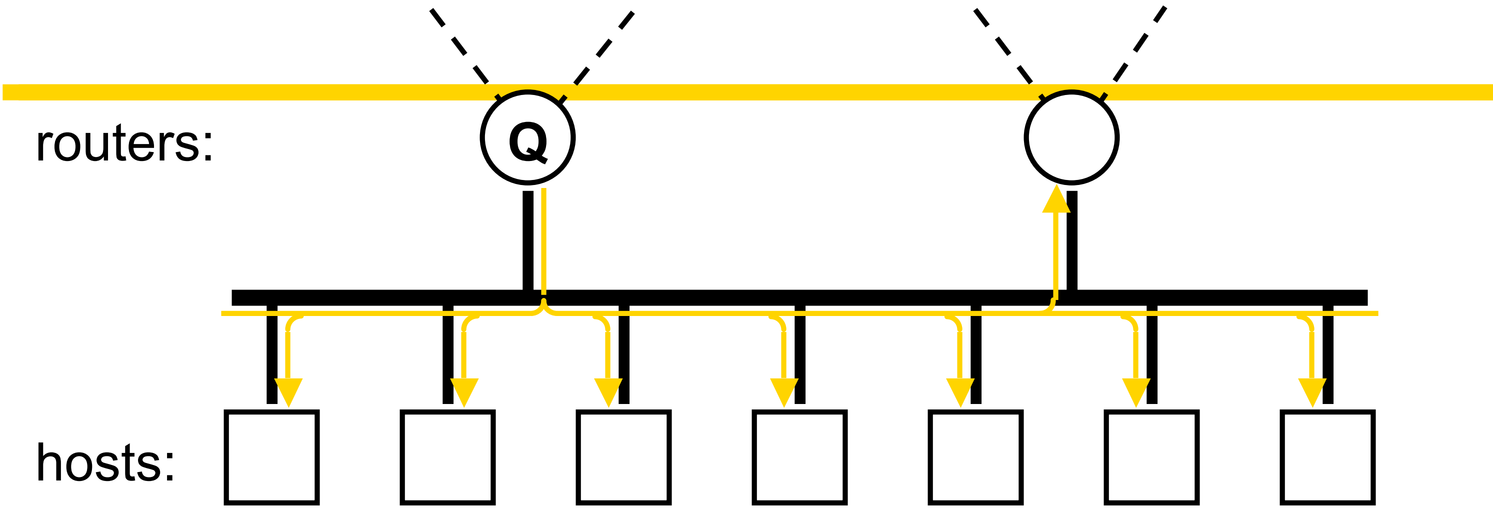

- on each link, one router is elected the “querier”
- querier periodically sends a Membership Query message to the all-systems group (224.0.0.1), with TTL = 1
- on receipt, hosts start random timers (between 0 and 10 seconds) for each multicast group to which they belong

How IGMP Works (cont.)

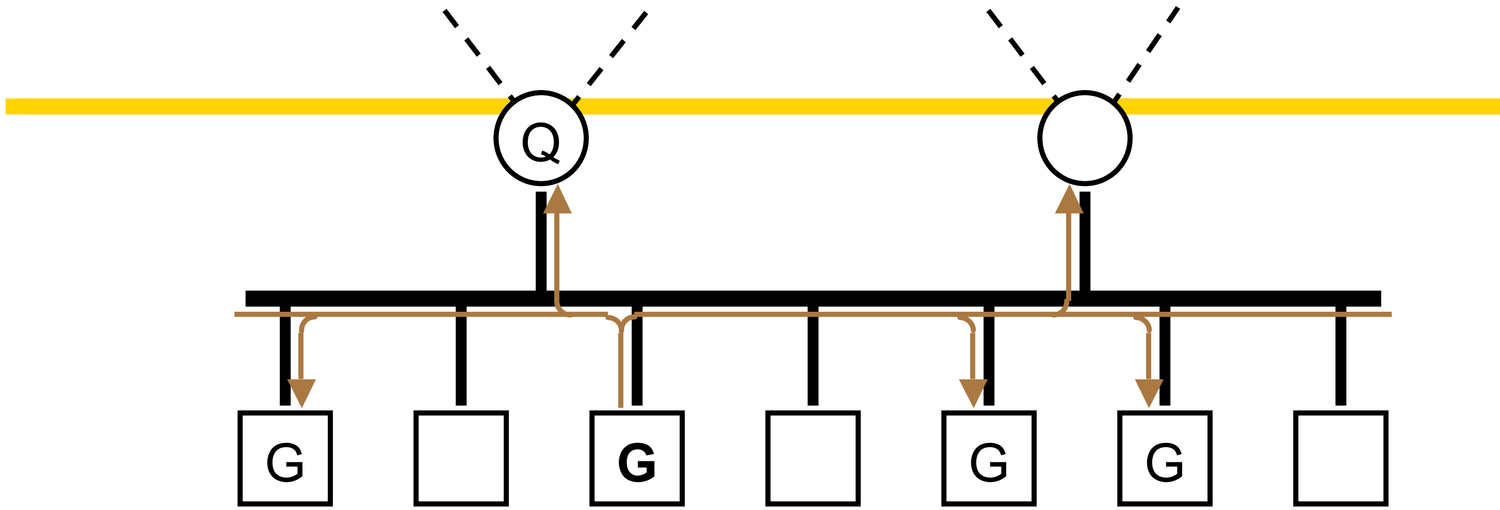

- when a host's timer for group G expires, it sends a Membership Report to group G, with TTL = 1
- other members of G hear the report and stop their timers
- routers hear all reports, and time out non-responding groups

How IGMP Works (cont.)

- note that, 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

- when a host first joins a group, it sends one or two immediate reports, 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
- widely implemented already