Multicasting for Multimedia

N. C. State University

CSC557 ♦ Multimedia Computing and Networking

Fall 2001

Lecture # 27

"Roadmap" for Multimedia Networking

1. Introduction
   - why QoS?
   - what are the problems?
2. Basic operations
   - jitter buffers (at hosts)
   - task scheduling (at hosts)
   - packet shaping (at hosts)
   - packet dropping (at routers)
   - packet scheduling (at routers)
3. Types of service
   - Integrated Services (IntServ) and Resource Reservation Protocol (RSVP)
   - Differentiated Services (DiffServ)
4. Application-level feedback and control
   - Real-time Protocol (RTP), Real-time Control Protocol (RTCP)
5. Application signaling and device control
   - Session Description Protocol (SDP)
   - Real-time Streaming Protocol (RTSP)
   - Session Initiation Protocol (SIP)
   - Media Gateway Control Protocol (MGCP)
6. Routing
   - **Multicasting**
   - Multi-protocol Label Switching (MPLS)

Today's Lecture

Uses

- One-to-many communication between hosts with similar interests
  - intermediate between broadcasting and unicasting

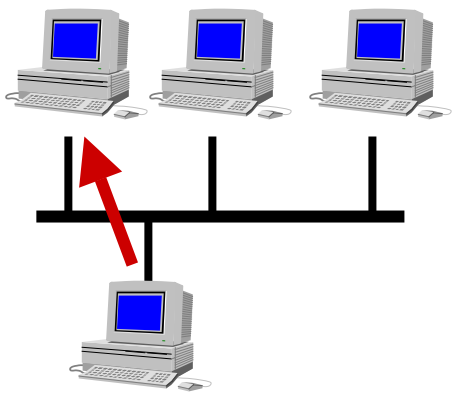

Unicast
1 to 1

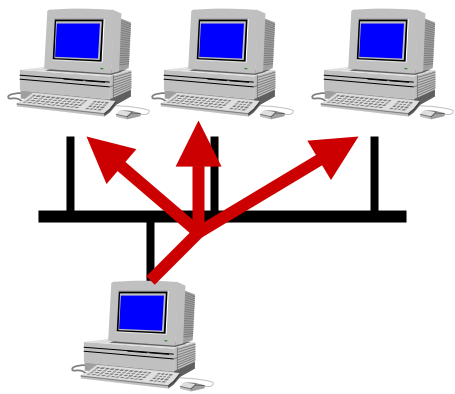

Broadcast
1 to all

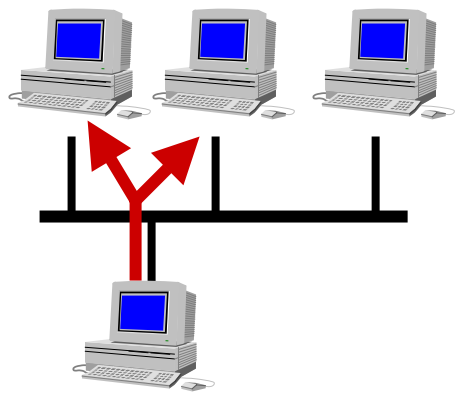

Multicast
1 to many

Uses (cont'd)

- A critical enabling technology for multimedia
  - movies
  - conferencing
  - concerts
  - multi-player games
- "Normal" multicast = UDP
- Can "multicasting" also mean "multiple qualities of service"?

Basic Idea

- Senders send packets to a multicast destination address
- Receivers dynamically subscribe to this address and receive the packets
- Routers make sure traffic is delivered from senders to interested receivers

Multicast Groups

- Groups may be of any size: 1 ... 1,000,000 ...
- Group members may be located anywhere in the Internet
- Hosts can join and leave groups at will
  - membership in a group is receiver-driven
  - there is no centralized "list" of group members
- Each multicast address is like a radio frequency, on which anyone can transmit, and to which anyone can tune-in
- Senders need not be members of the group

Multicast Addressing

- There's a range of addresses available for multicasting
- Senders somehow choose an address that is not already in use
  - multicast address servers within a domain hand out addresses to clients
- Class D addresses are reserved for multicast
  - address range = 224.0.0.0 to 239.255.255.255
  - 256M different multicast groups are possible

Multicast Addressing (cont'd)
(The 4 major classes of IP addresses)

- Class A address
  - 128 networks
  - > 65,536 hosts
- Class B address
  - 16,384 networks
  - 256-65,536 hosts
- Class C address
  - 2^{21} networks
  - less than 256 hosts
- Class D address
  - 28 bit multicast addr.

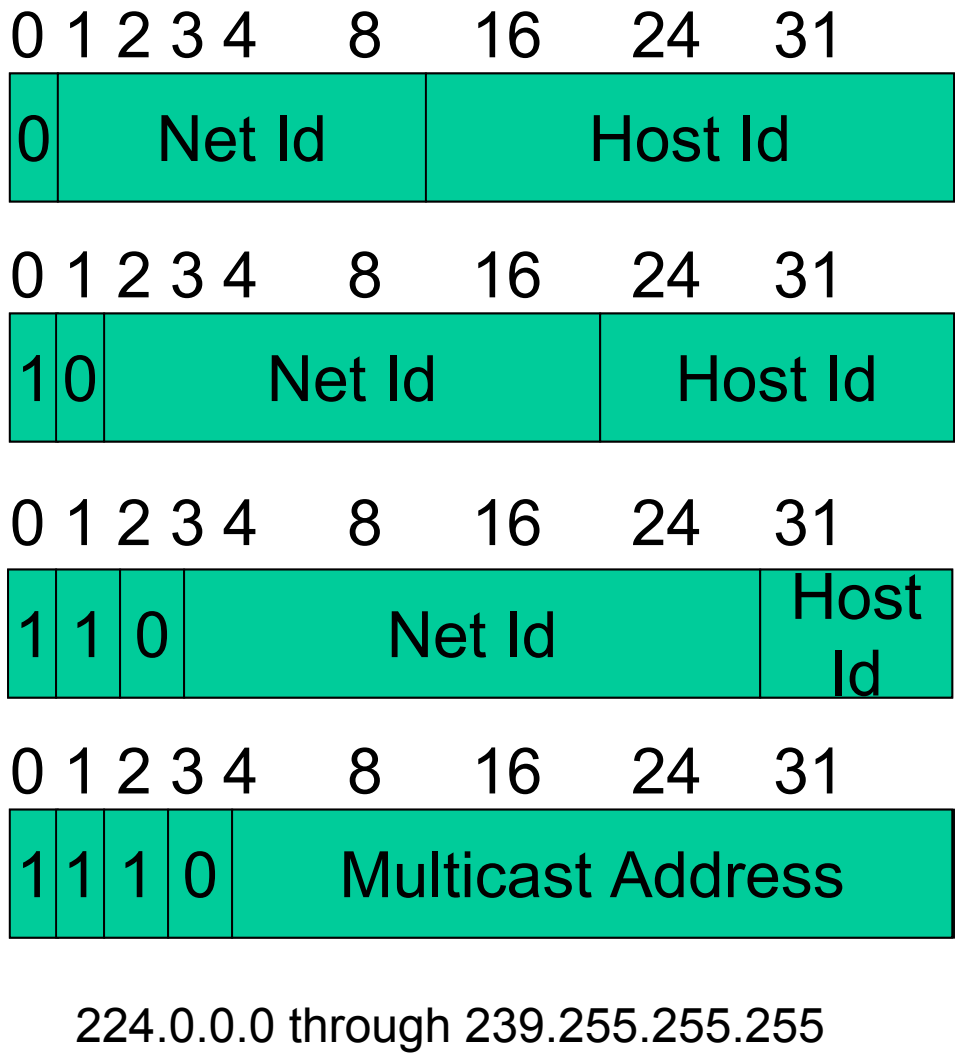

224.0.0.0 through 239.255.255.255

Internet Group Multicast Protocol (IGMP)

- Multicasting over a LAN (i.e., Ethernet)
- Mapping of IP addresses to Ethernet MAC address

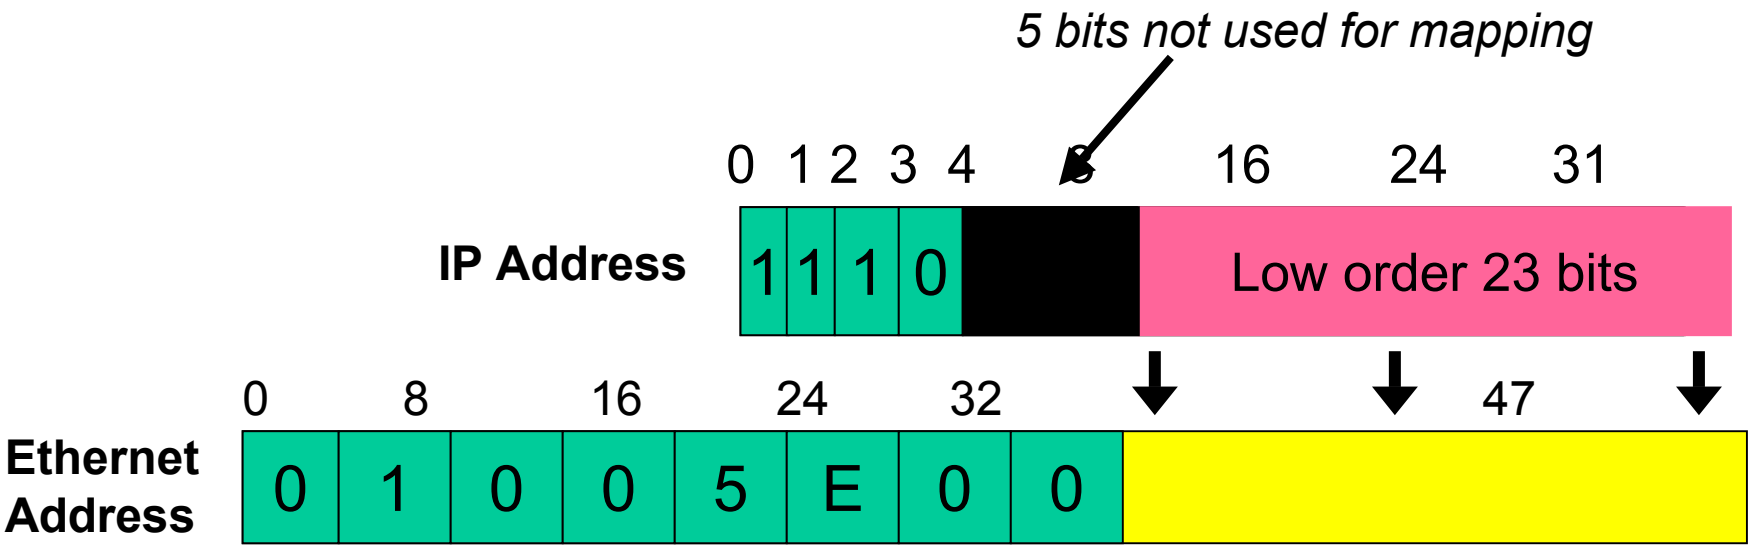

IEEE-802 Ethernet Multicast Address

- As many as 32 groups may be mapped to the same Ethernet address
  - hopefully not a common occurence

IGMP (cont'd)

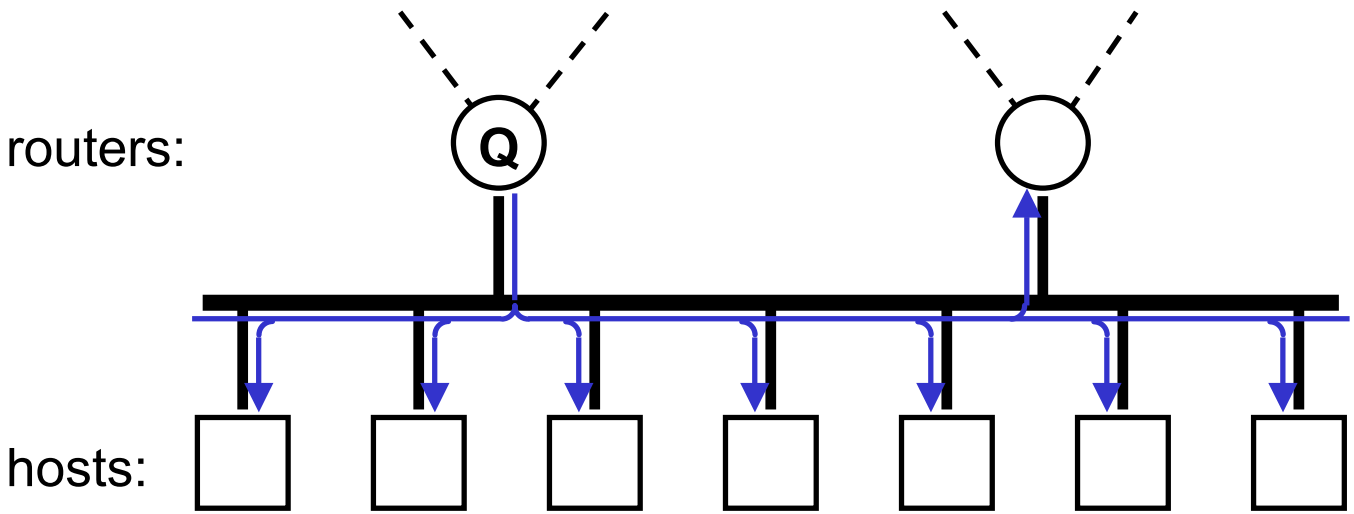

- On each link, one router is elected the "querier"
- Querier periodically sends a Membership Query message to the all-systems group (224.0.0.1), with TTL = 1
- On receipt, hosts start random timers (between 0 and 10 seconds) for each multicast group to which they belong

IGMP (cont'd)

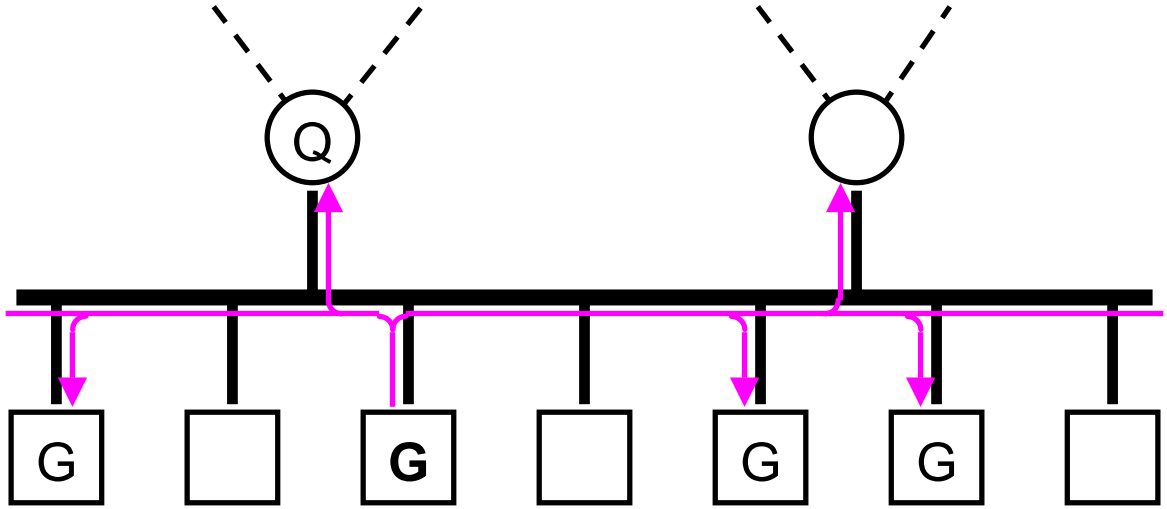

- When a host's timer for group G expires, it sends a Membership Report to group G, with TTL = 1
- Other members of G hear the report and stop their timers
- Routers hear all reports, and time out non-responding groups

Multicast Routing in the Internet

- Multicast service model makes it hard to locate receivers
  - anonymity
  - dynamic joins and leaves
- Options so far...
  - flood data packets to the entire network, or
  - tell routers about all possible groups and receivers so they can create routes (trees)

Multicasting Over the Internet

- DVMRP (Distance Vector Multicast Routing Protocol) & PIM-DM (Dense Mode Protocol Independent Multicast)
- MOSPF (Multicast Open Shortest Path First)
- CBT (Core Based Trees)

DVMRP and PIM-DM

- "Flood and prune" approach
- A source generates traffic to a multicast address
- Basic method
  - a router A receiving a multicast packet checks the interface on which it arrives
  - if *not* on the path from router A to the source, packet is dropped
  - if it *is* on the path from router A to the source, packet is copied and sent out to all other neighbors B (flooded)
- Improvement
  - only send multicast packet out an interface if router A is on shortest path between source and neighbor B

DVMRP / PIM-DM (cont'd)

- When traffic reaches a router that knows (via IGMP) there are no receivers, a "prune" message is sent back towards the source
- Can periodically repeat, so that new receivers can join
- The multicast tree that results is a shortest-path tree that reaches all the receivers
  - good algorithm for smaller networks
  - won't scale up to Internet size, however!

Example Topology

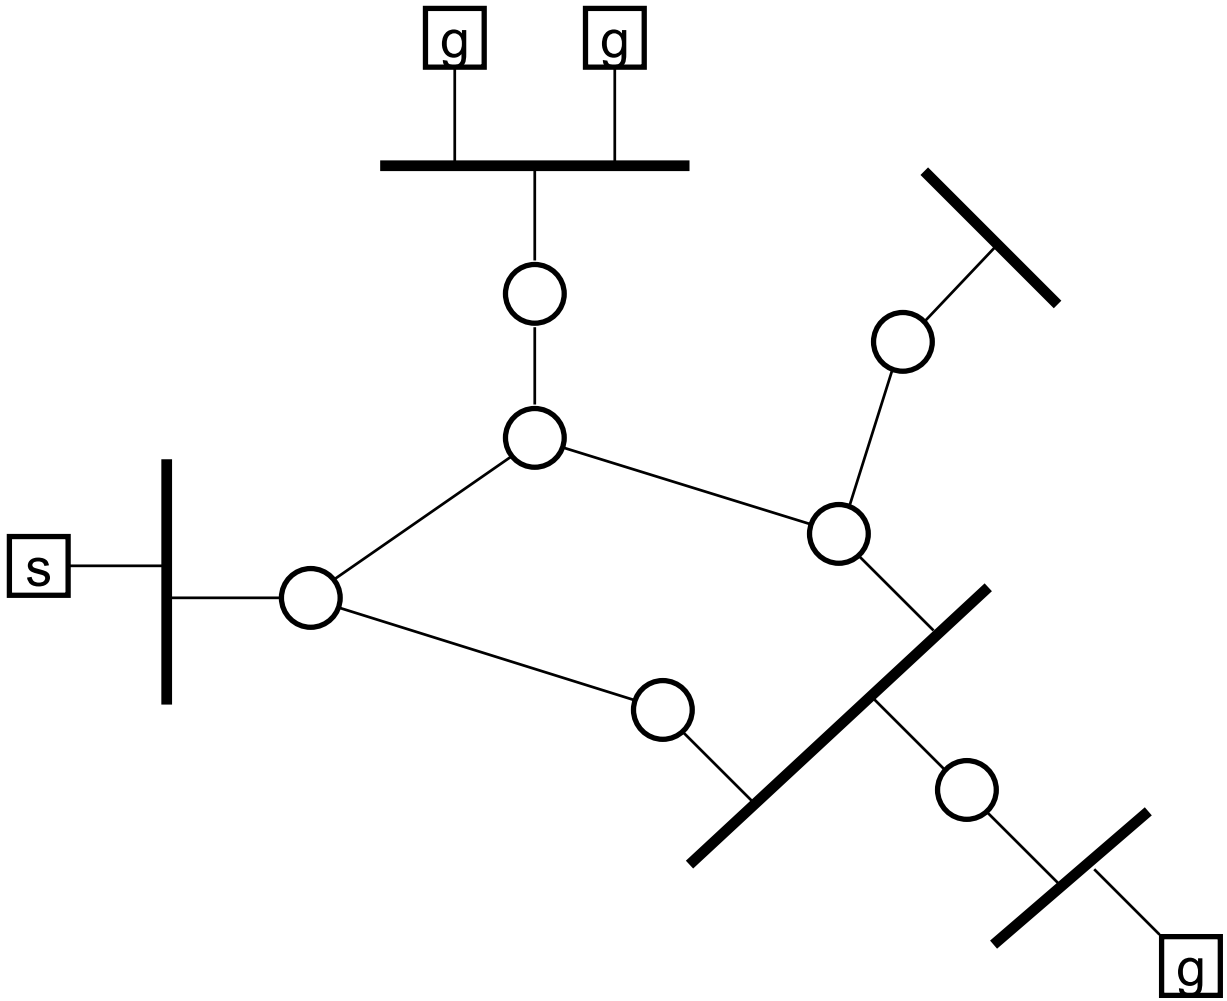

Phase 1: Flood using Truncated Broadcast

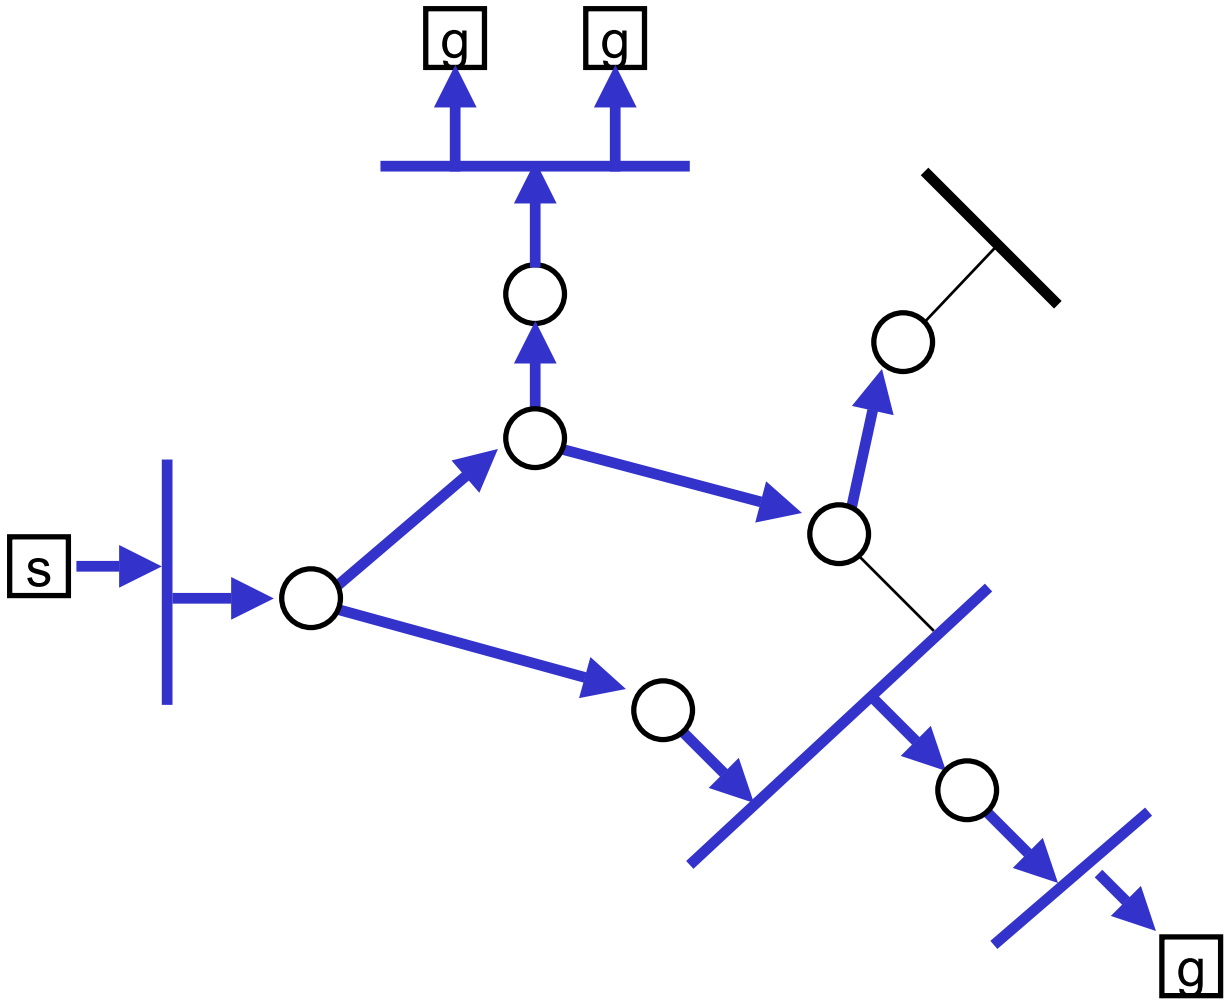

Phase 2: Prune

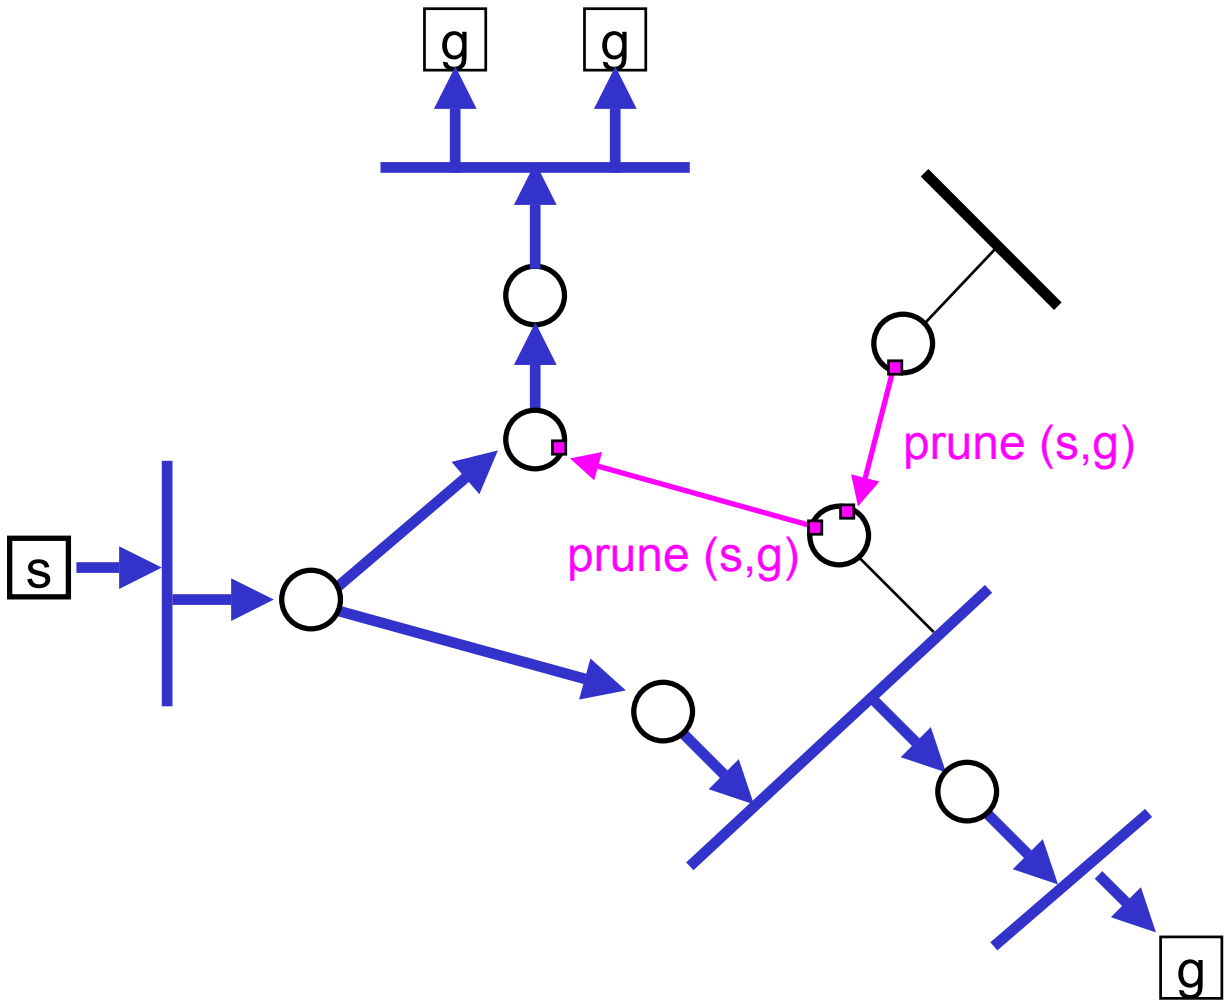

Phase 3: Graft

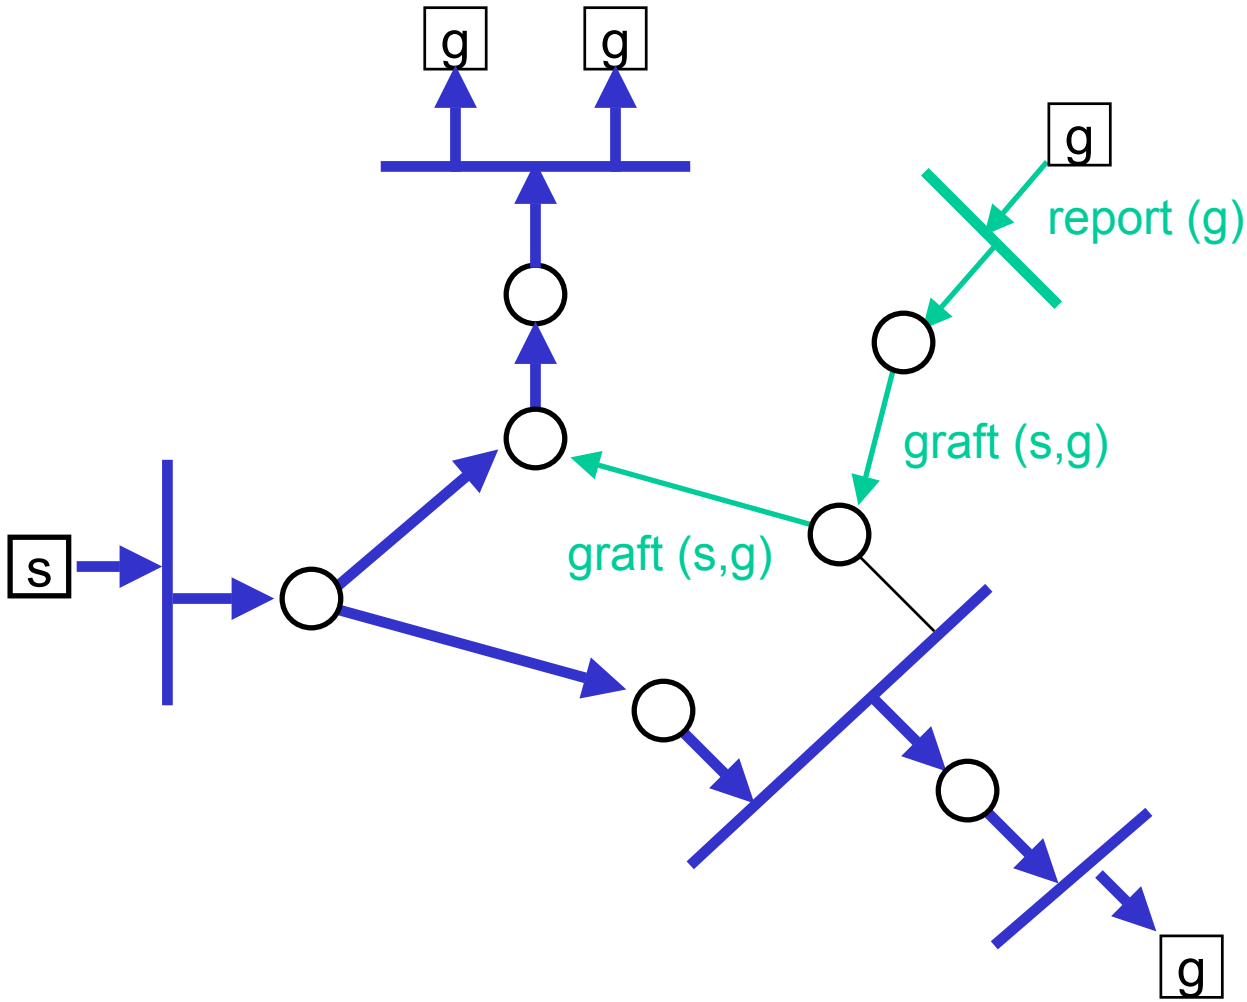

Phase 4: Steady State

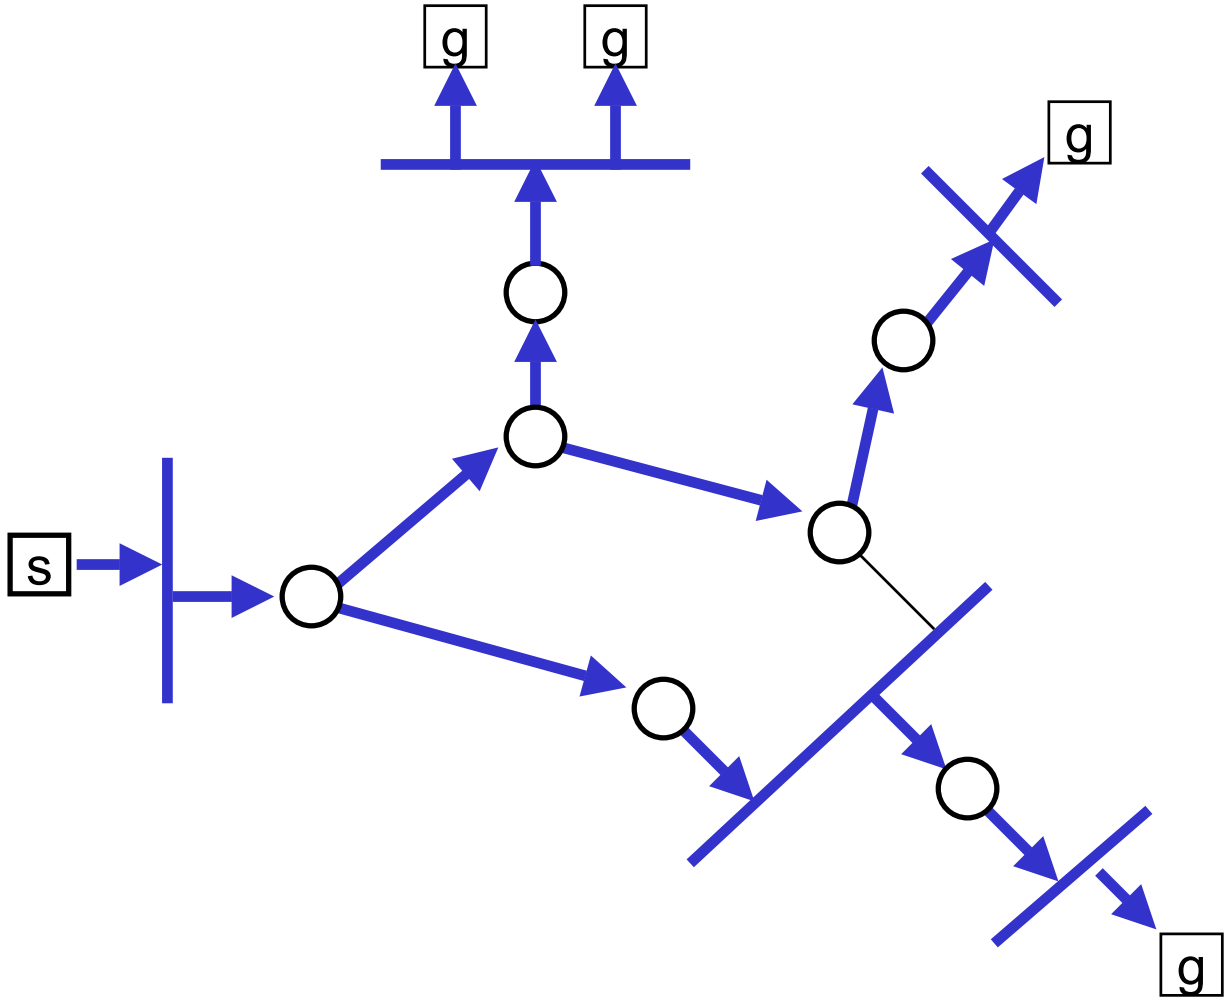

MOSPF (Multicast OSPF)

- OSPF (a unicast routing protocol)
  - all routers exchange information about the links in the network
  - all routers reach a consistent view of the network topology and state
- MOSPF (multicast extension)
  - routers also broadcast info about what multicast groups any of their local receivers are listening to
  - allows all routers to construct an explicit multicast tree
- Scalability problems
  - too much info to distribute, receiver list may change too rapidly

Core-Based Trees

- Cores
  - "meeting places" where sources send initial packets, and receivers join
  - requires mapping from multicast group address to core IP address
  - central distribution points (also called rendezvous points)
  - allows constructing a tree that is somewhat independent of the sources

A shared tree

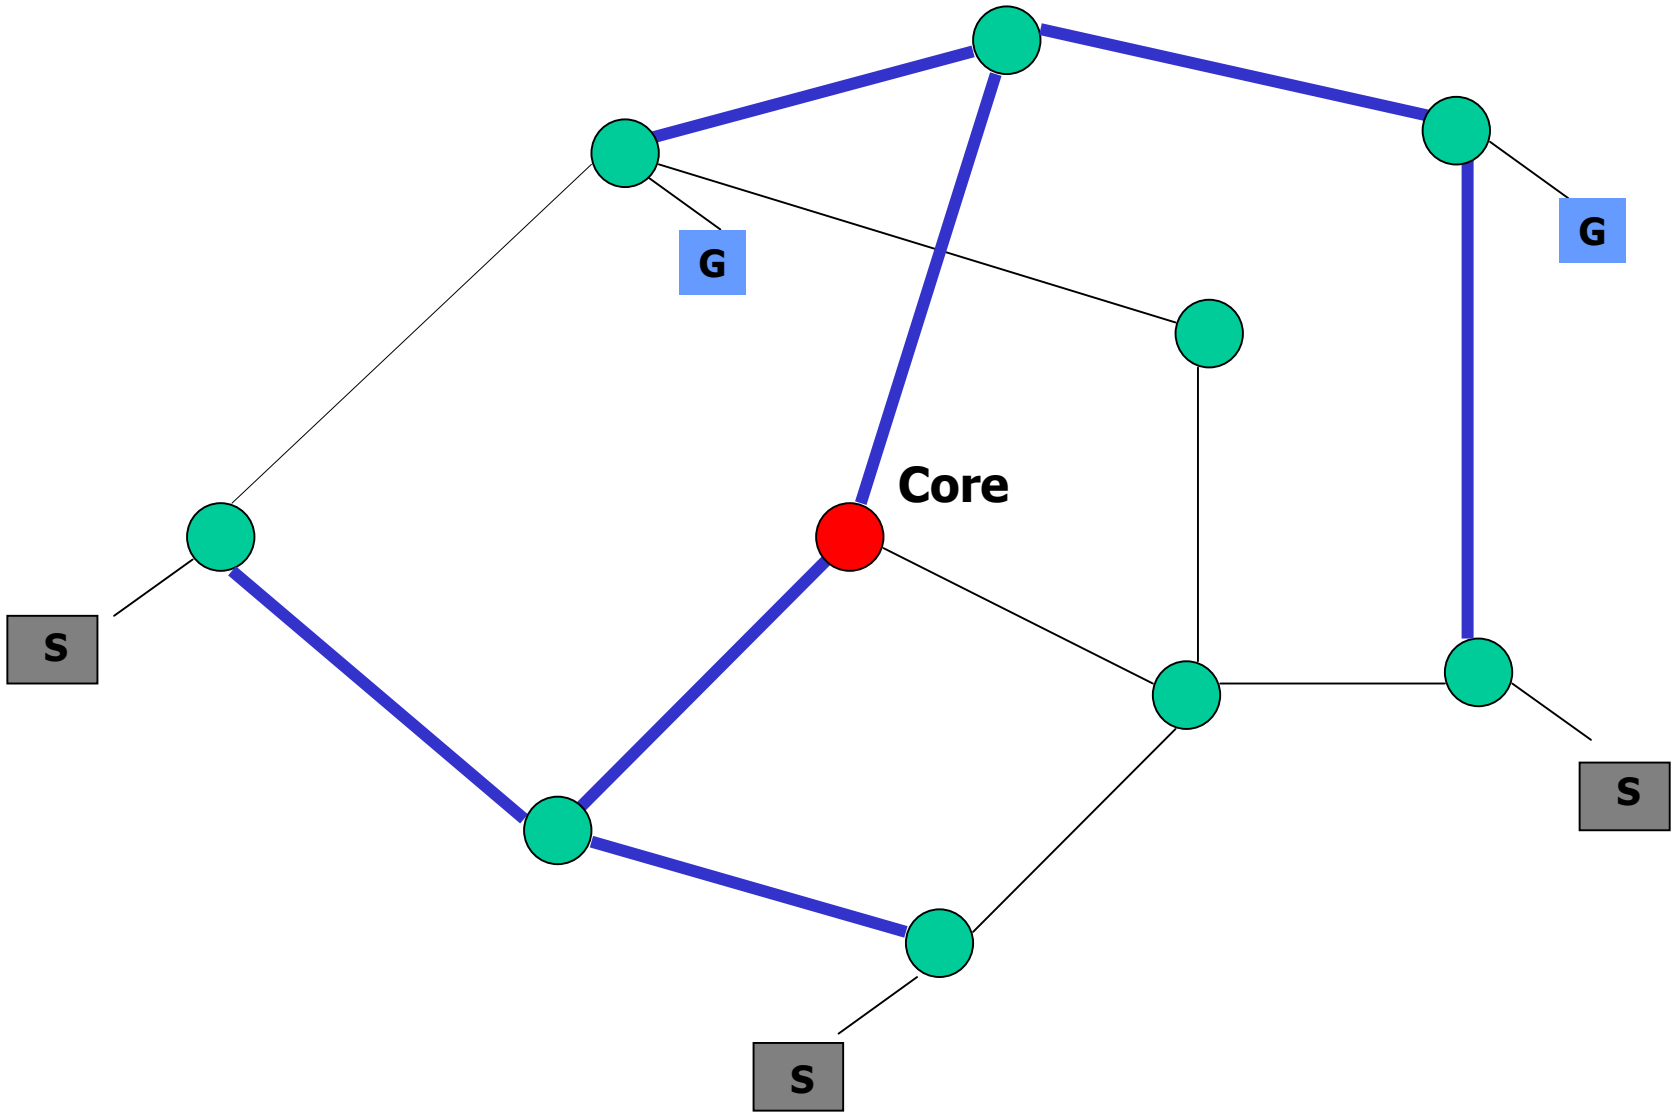

CBT (cont'd)

- Each receiver interested in a multicast group sends an explicit "join" message toward the core
  - each router encountered on the way will *remember* the multicast group and the incoming interface
  - join message stops when it encounters a router already forwarding for that multicast address, or when the message reaches the core
- Multicast trees are bi-directional
  - a sender that is also a receiver sends on the tree
- A sender that is not a receiver sends data towards the core
  - a router on the "tree" will duplicate the message and send out all branches
- CBT scales much better than flood-and-prune algorithms
  - however, there may be reliability and congestion problems at the core

Limiting the Scope of Multicast

- (i.e., determining how widely to multicast the data)
- TTL (time-to-live) field
  - Initialized by sender to some maximum number of hops
  - Decremented by 1 for each router encountered
  - Packet is discarded when TTL = 0
  - Prevents routing loops, etc.
- Multicast router modification: discard packet when TTL < preset "multicast threshold"
- To keep multicast packets within a domain, set multicast threshold at all *border routers* to a fairly large number
  - initialize TTL to less than this threshold

Multicast Scope Control: Administrative TTL Boundaries

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

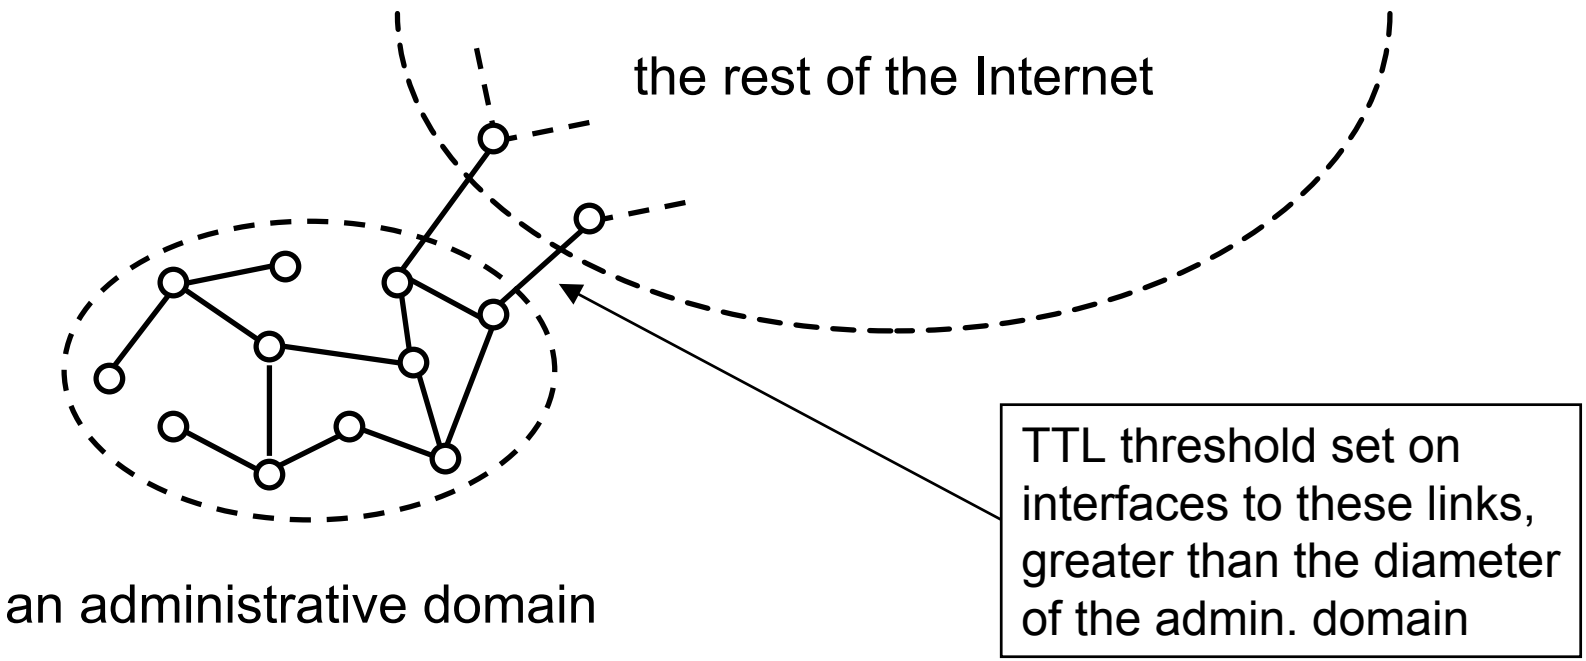

MBONE: The Multicast Backbone

- Many routers do not support or are not configured to provide multicast capabilities
- MBONE is an overlay network in the Internet
  - used for a number of applications: vic, vat, wb, nv, ivs
- Multicast packets are "tunneled" between routers supporting multicast tunneling
  - (i.e.,encapsulate multicast packets between the endpoints inside a "normal"-looking unicast packet)

MBONE (example)

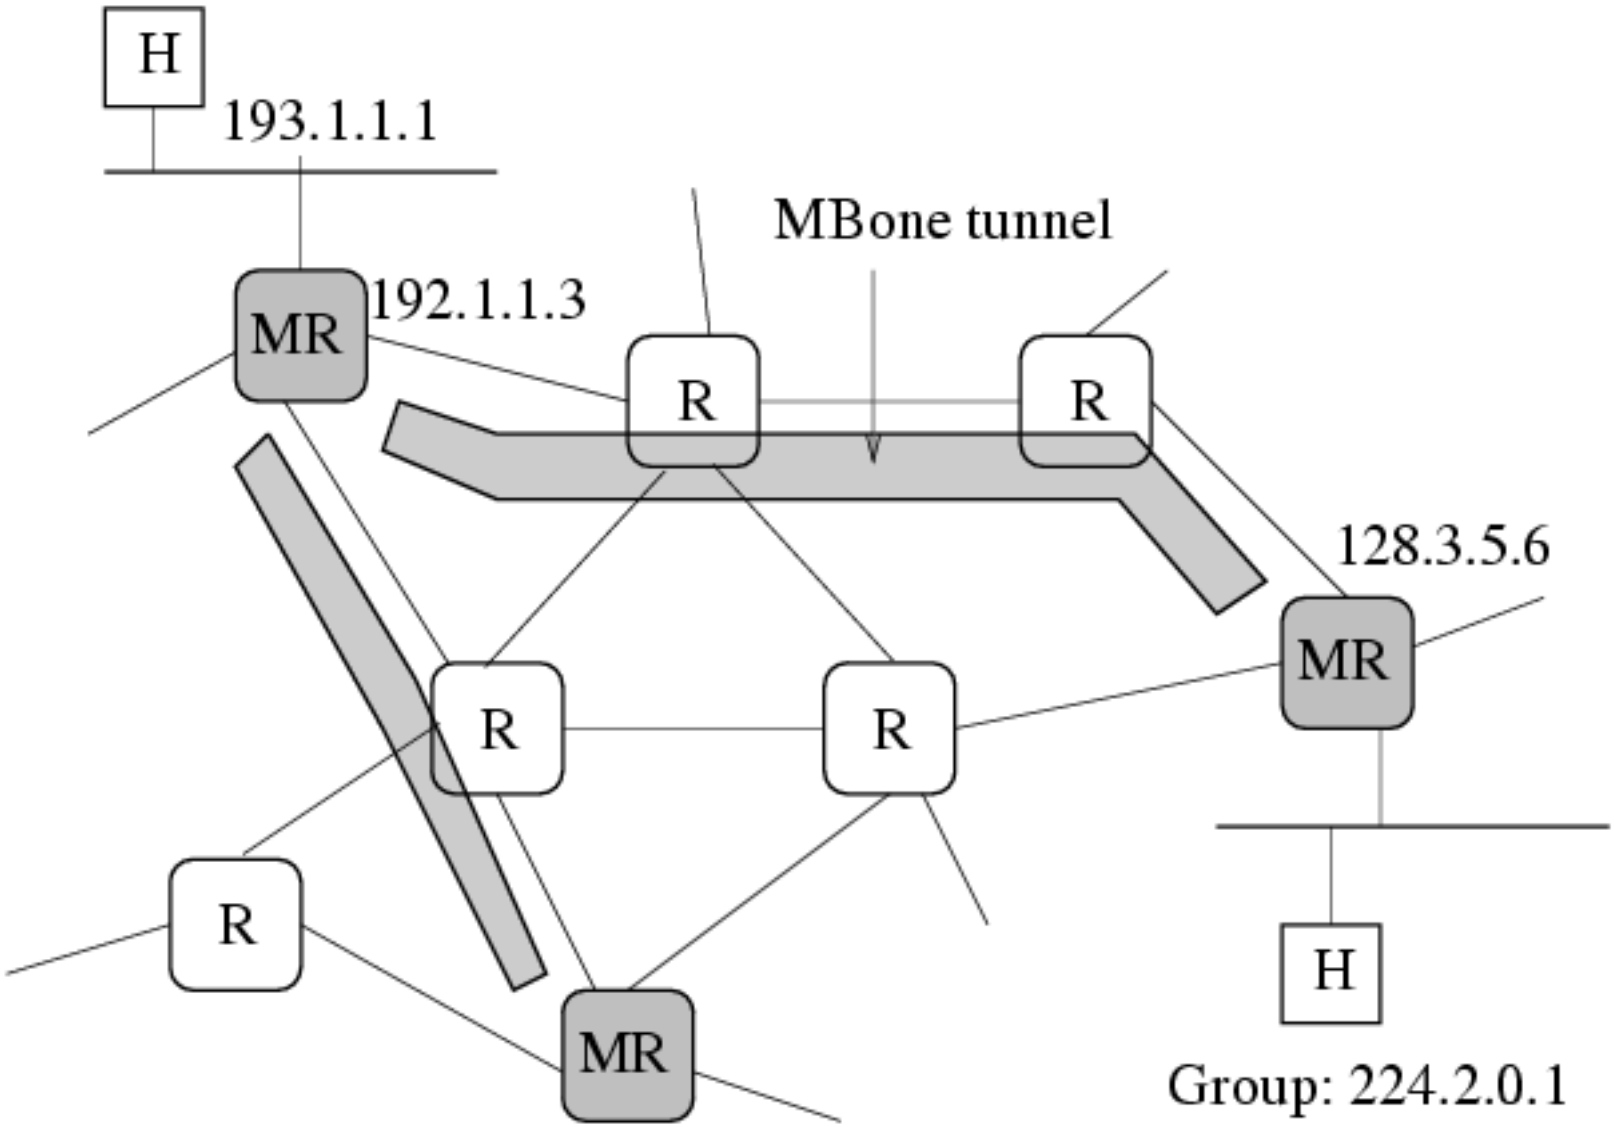

MBONE (cont'd)

- Requires a separate routing protocol for multicasting, since based on the overlay network
- Fixed amount of capacity reserved for use by MBONE
  - use SAP for session announcements and scheduling of MBONE capacity

Sources of Info

- Recommended Books
  - Online!: J. Crowcroft et al., *Internetworking Multimedia*, 1999 (chapter 3)
- Web
  - "Introduction to IP Multicast Routing"
  - "IP Multicast"