Advanced Networking

Multicast

Renato Lo Cigno – Alessandro Russo
Renato.LoCigno@disi.unitn.it - Russo@disi.unitn.it
Homepage:
disi.unitn.it/locigno/index.php/teaching-duties/advanced-networking

Multicasting

- Addresses that refer to group of hosts on one or more networks
- Applications
  - Multimedia “broadcast” and streaming
  - Teleconferencing
  - Distributed Database
  - Distributed computing (GRID??)
  - Real time workgroups
  - File distribution

Example of multicast configuration

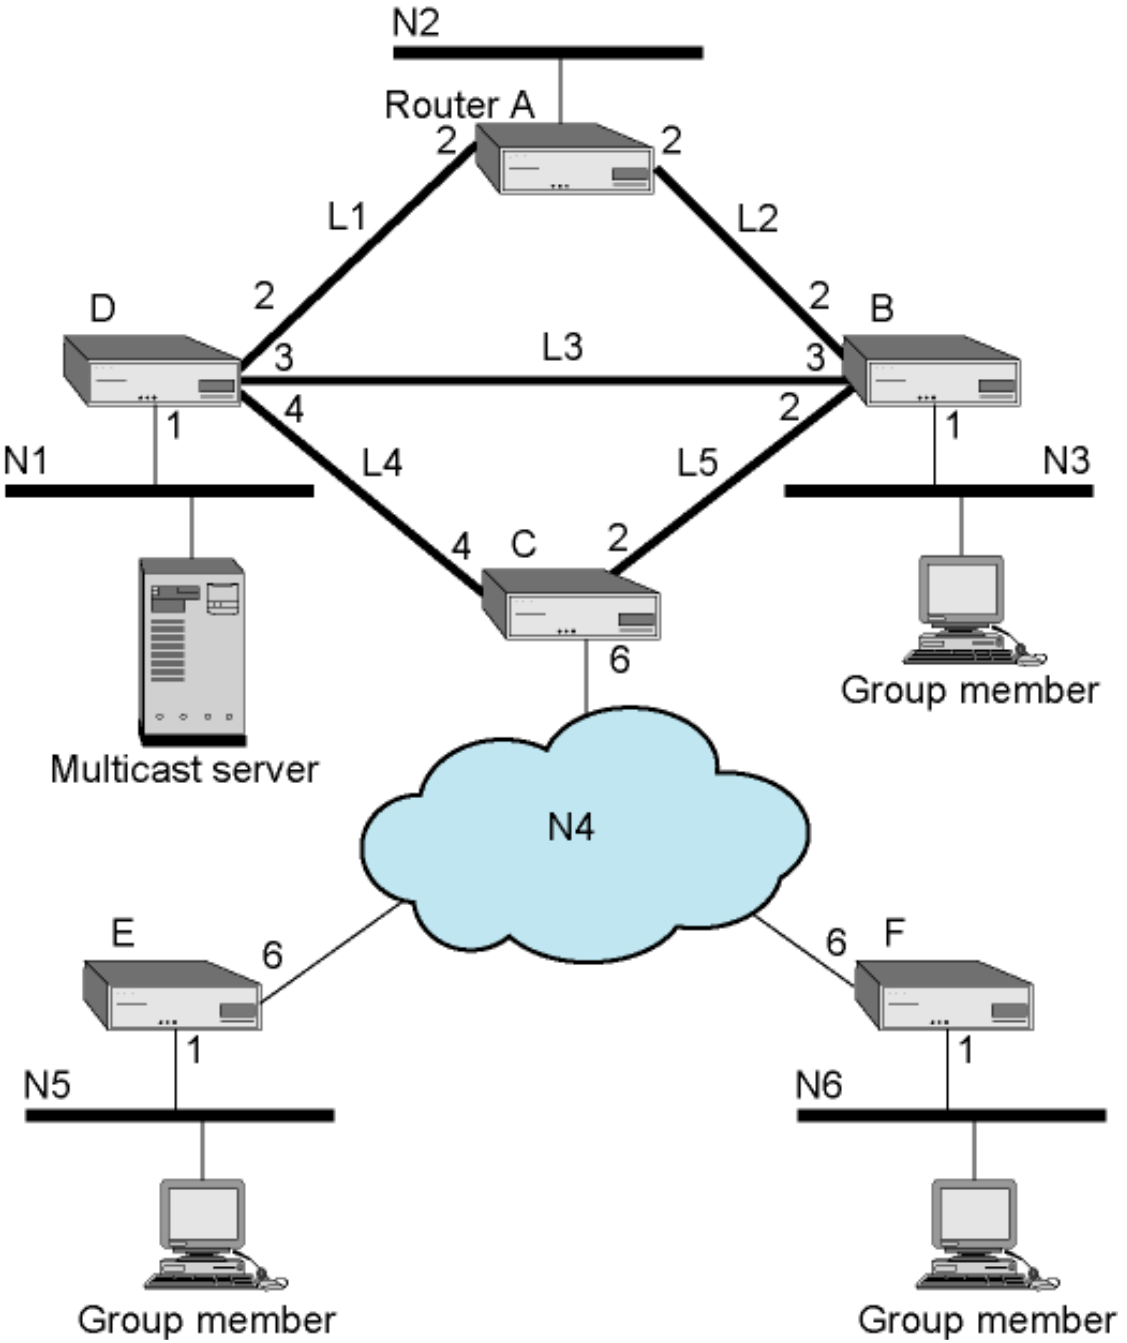

Broadcast and Multiple Unicast

- Broadcast a copy of packet to each network
  - Requires 13 copies of packet
- Multiple Unicast
  - Send packet only to networks that have hosts in group
  - 11 packets

Multicast Routing

- Multicast: delivery of same packet to a group of receivers
- Multicasting is becoming increasingly popular in the Internet (video on demand; whiteboard; interactive games)
- Multiple unicast vs. multicast

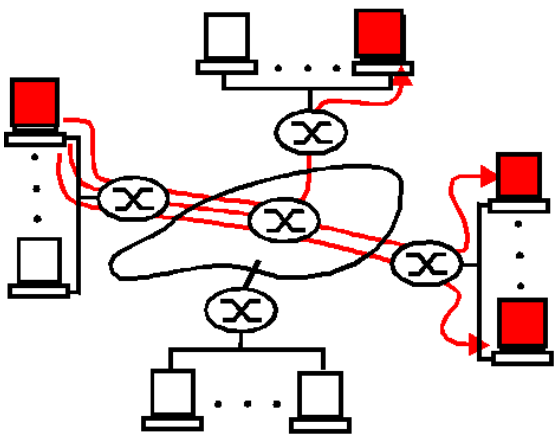

multicast via unicast

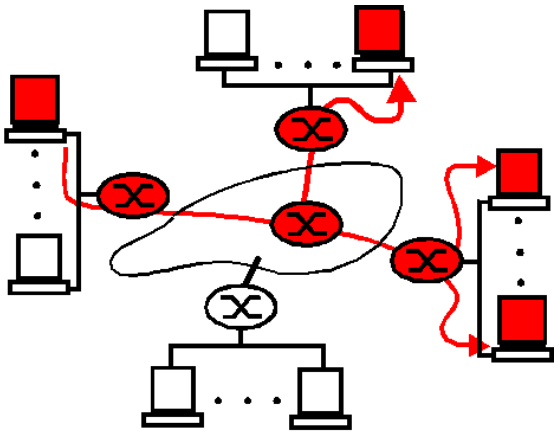

network multicast

True Multicast

- Determine least cost path to each network that has host in group
  - Gives spanning tree configuration containing networks with group members
- Transmit single packet along spanning tree
- Routers replicate packets at branch points of spanning tree
- 8 packets required

Spanning Tree Problem

- Given a graph G=(V,E)
  - nodes are vertices and links are edge
  - connected and undirected
- A Spanning Tree (ST) for G is a subgraph without cycles (i.e., a tree) which covers all vertices
- There are one or more STs for G

Multicast Transmission Example

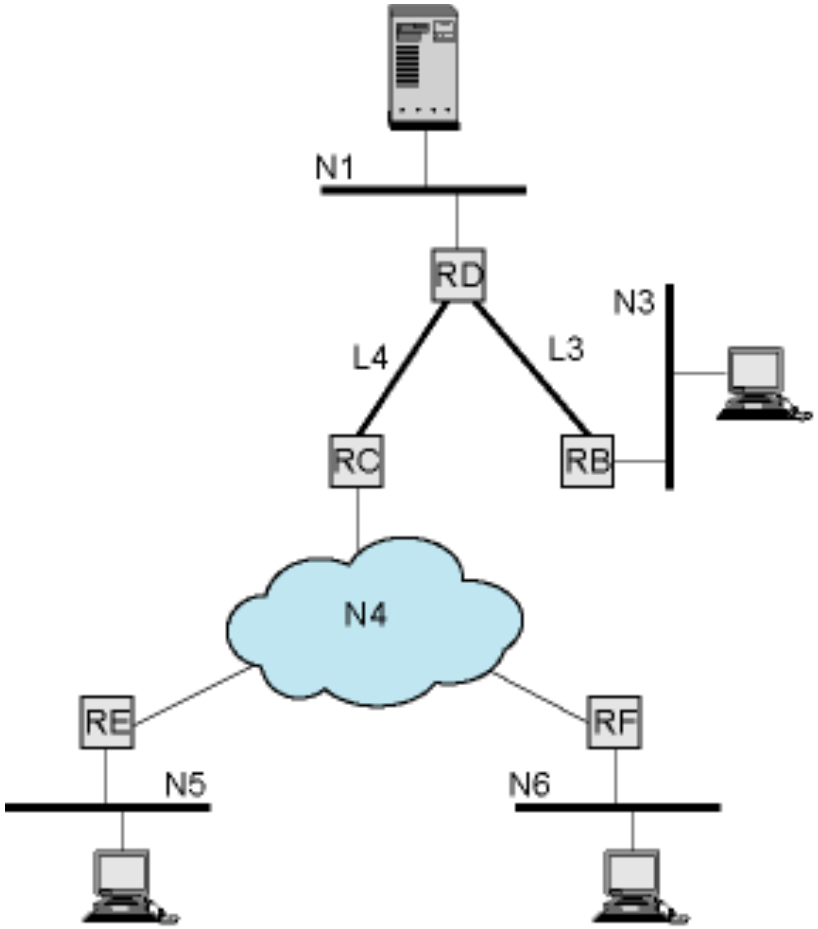

(a) Spanning tree from source to multicast group

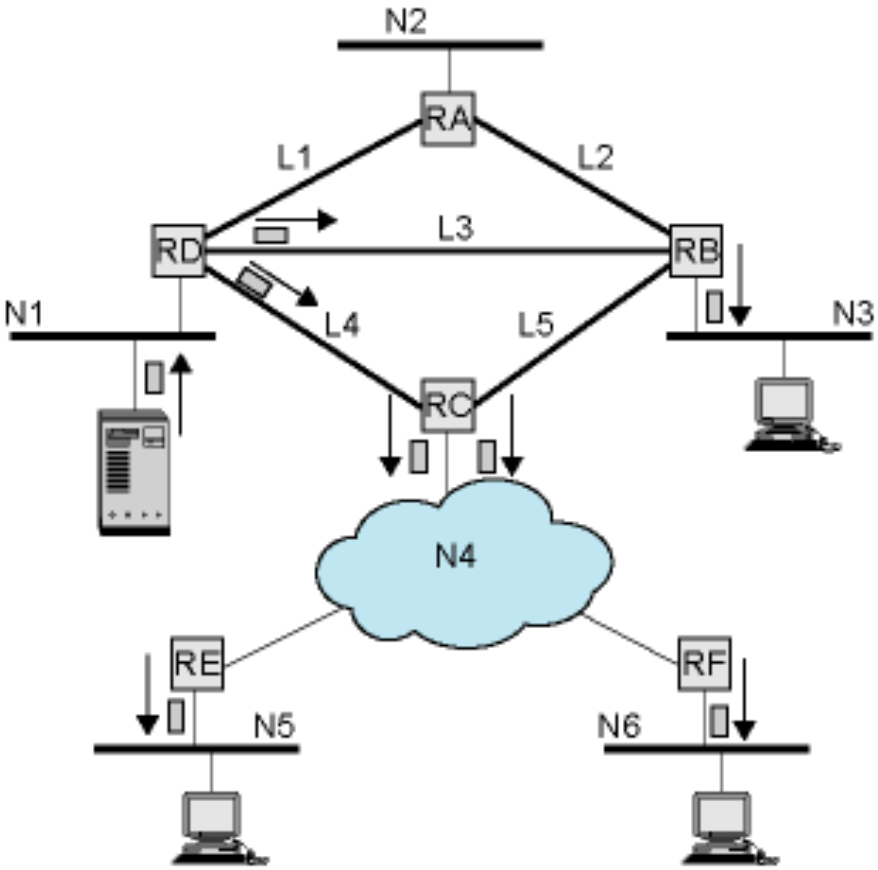

(b) Packets generated for multicast transmission

Multicast Group Address

- M-cast group address “delivered” to all receivers in the group
- Internet uses Class D for m-cast
- M-cast address distribution etc. managed by IGMP Protocol

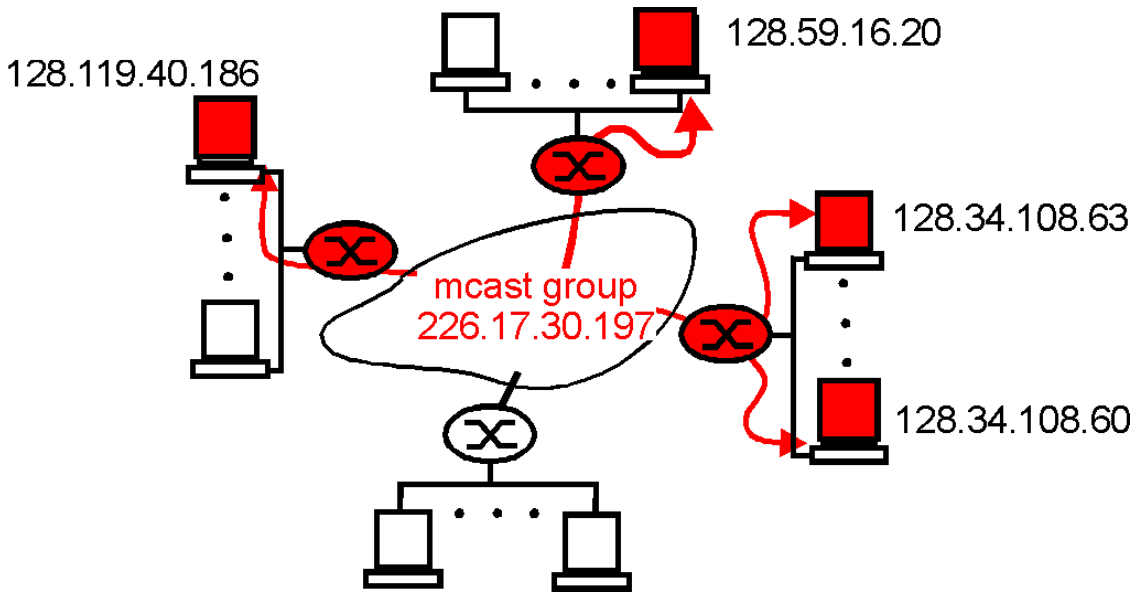

Requirements for Multicasting (1)

- Router may have to forward more than one copy of packet
- Convention needed to identify multicast addresses
  - IPv4 - Class D - start 1110
  - IPv6 - 8 bit prefix, all 1, 4 bit flags field, 4 bit scope field, 112 bit group identifier
- Nodes must translate between IP multicast addresses and list of networks containing group members
- Router must translate between IP multicast address and network multicast address

Requirements for Multicasting (2)

- Mechanism required for hosts to join and leave multicast group
- Routers must exchange info
  - Which networks include members of given group
  - Sufficient info to work out shortest path to each network
  - Routing algorithm to work out shortest path
  - Routers must determine routing paths based on source and destination addresses

Spanning Tree from Router C to Multicast Group

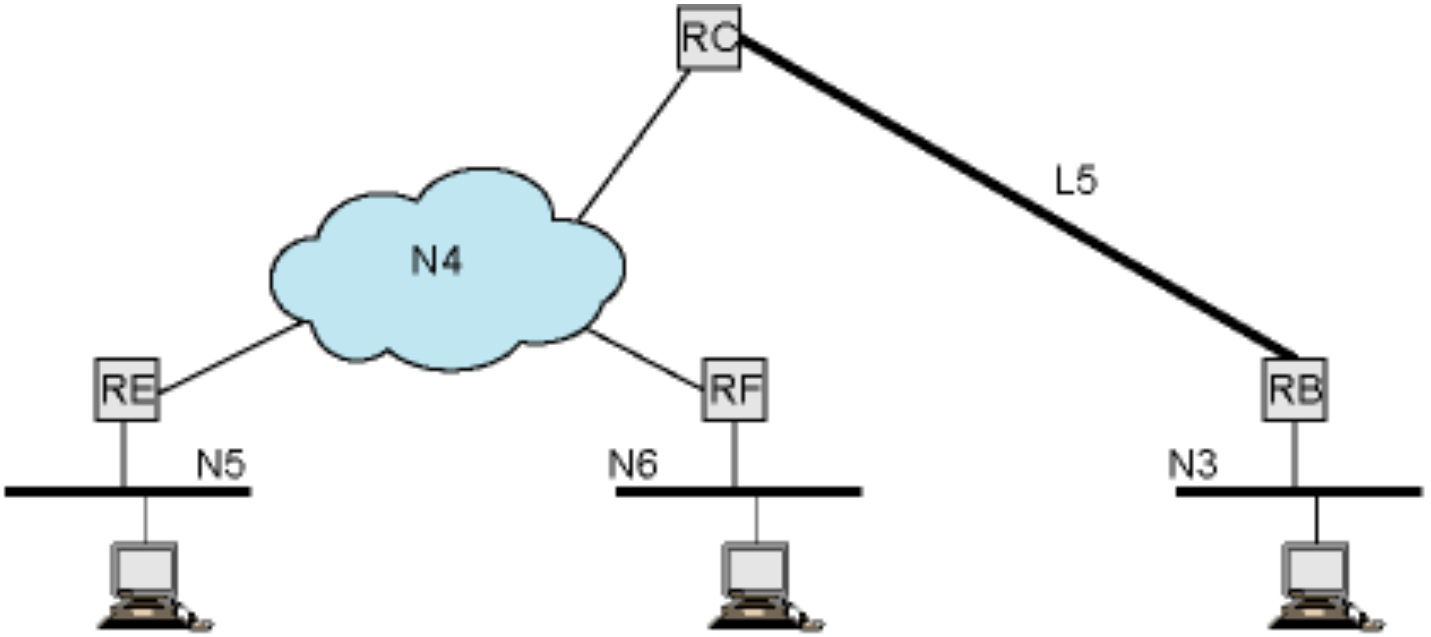

Internet Group Management Protocol (IGMP)

- IGMP v3: RFC 3376 (2002)
- IGMP v2: RFC 2236 (1997)
- IGMP v1 alias Host Extensions for IP Multicasting v3: RFC 1112 (1989)
- Obsoletes: RFCs 988, 1054
  - Host Extensions for IP Multicasting v2: RFC 1054 (1988)
  - Host Extensions for IP Multicasting v1: RFC 988 (1986)
- Host and router exchange of multicast group info
- Use broadcast LAN to transfer info among multiple hosts and routers

Principle of Operations

- Hosts send messages to routers to subscribe to and unsubscribe from multicast group
  - Group defined by multicast address
- Routers check which multicast groups of interest to which hosts
- IGMP currently version 3
- IGMPv1
  - Hosts could join group
  - Routers used timer to unsubscribe members

Operation of IGMP v1 & v2

- Receivers have to subscribe to groups
- Sources do not have to subscribe to groups
- Any host can send traffic to any multicast group
- Problems:
  - Spamming of multicast groups
  - Even if application level filters drop unwanted packets, they consume valuable resources
  - Establishment of distribution trees is problematic
  - Location of sources is not known
  - Finding globally unique multicast addresses difficult

IGMP v3

- Allows hosts to specify list from which they want to receive traffic
  - Traffic from other hosts blocked at routers
- Allows hosts to block packets from sources that send unwanted traffic

IGMP dialogues

- IGMP (Internet Group Management Protocol) operates between Router and local Hosts, typically attached via a LAN (e.g., Ethernet)
- Router queries the local Hosts for m-cast group membership info

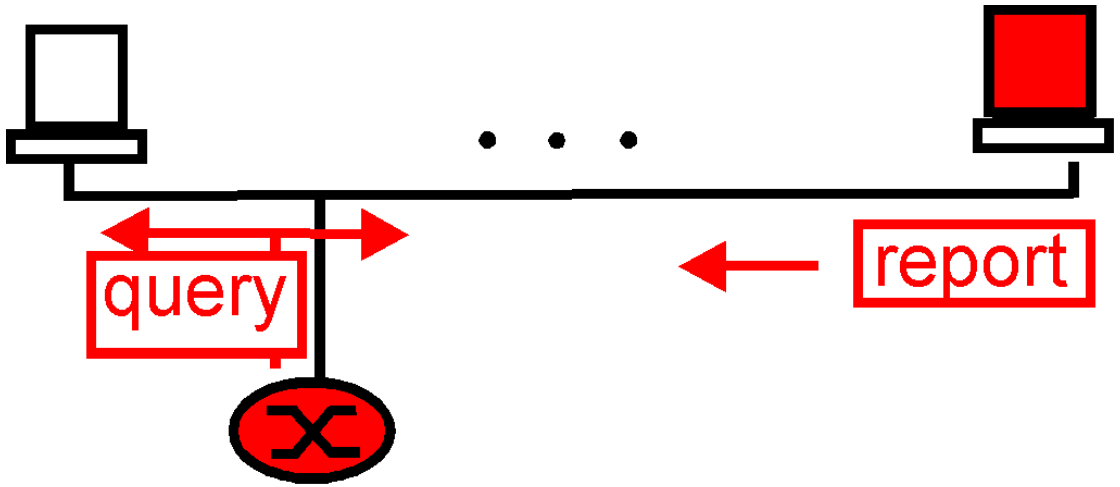

IGMP Protocol

- Router “connects” active Hosts to m-cast tree via m-cast protocol
- Hosts respond with membership reports: actually, the first Host which responds (at random) speaks for all
- Host issues “leave-group” msg to leave; this is optional since router periodically polls anyway (soft state concept)

IGMP message types

IGMP Message type	Sent by	Purpose
membership query: general	router	query for current active multicast groups
membership query: specific	router	query for specific m-cast group
membership report	host	host wants to join group
leave group	host	host leaves the group

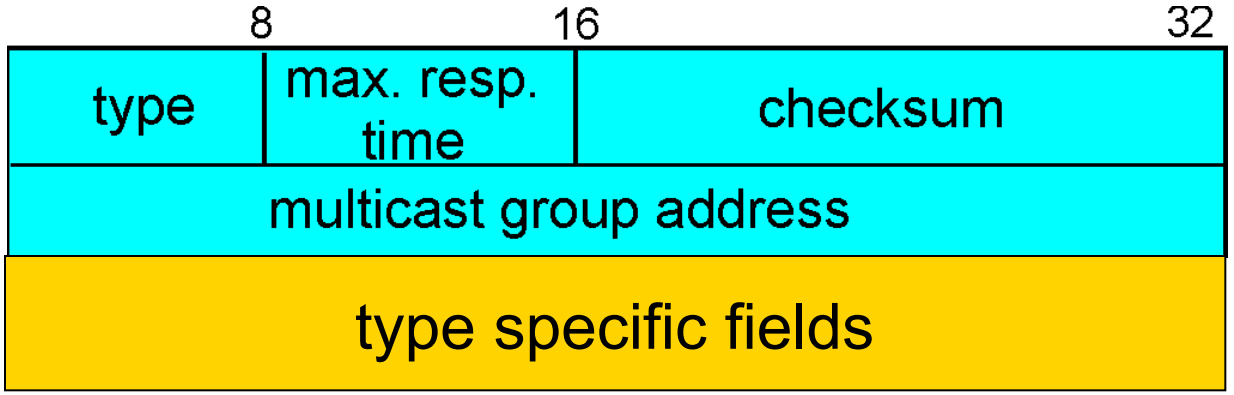

IGMP Message Formats: Membership Query

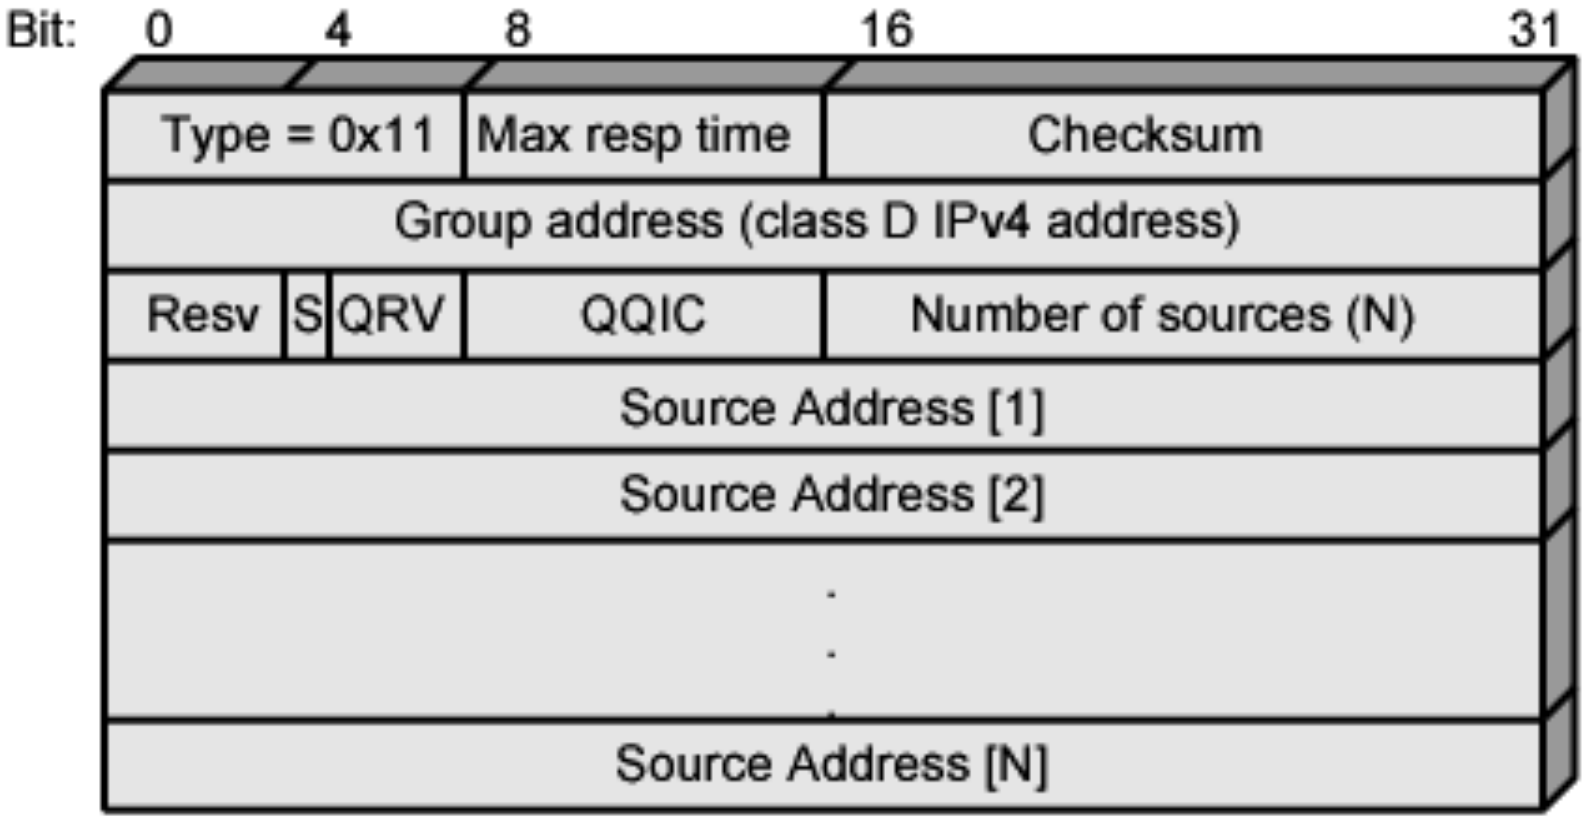

(a) Membership query message

Membership Query

- Sent by multicast router
- General query
  - Which groups have members on attached network
- Group-specific query
  - Does group have members on an attached network
- Group-and-source specific query
  - Do attached device want packets sent to specified multicast address
  - From any of specified list of sources

Membership Query Fields (1)

- Type
- Max Response Time
  - Max time before sending report in units of 1/10 second
- Checksum
  - Same algorithm as IPv4
- Group Address
  - Zero for general query message
  - Multicast group address for group-specific or group-and-source
- S Flag
  - 1 indicates that receiving routers should suppress normal timer updates done on hearing query

Membership Query Fields (2)

- QRV (querier's robustness variable)
  - RV value used by sender of query
  - Routers adopt value from most recently received query
  - Unless RV was zero, when default or statically configured value used
  - RV dictates number of retransmissions to assure report not missed
- QQIC (querier's query interval code)
  - QI value used by querier
  - Timer for sending multiple queries
  - Routers not current querier adopt most recently received QI
  - Unless QI was zero, when default QI value used
- Number of Sources
- Source addresses
  - One 32 bit unicast address for each source

IGMP Message Formats: Membership Report

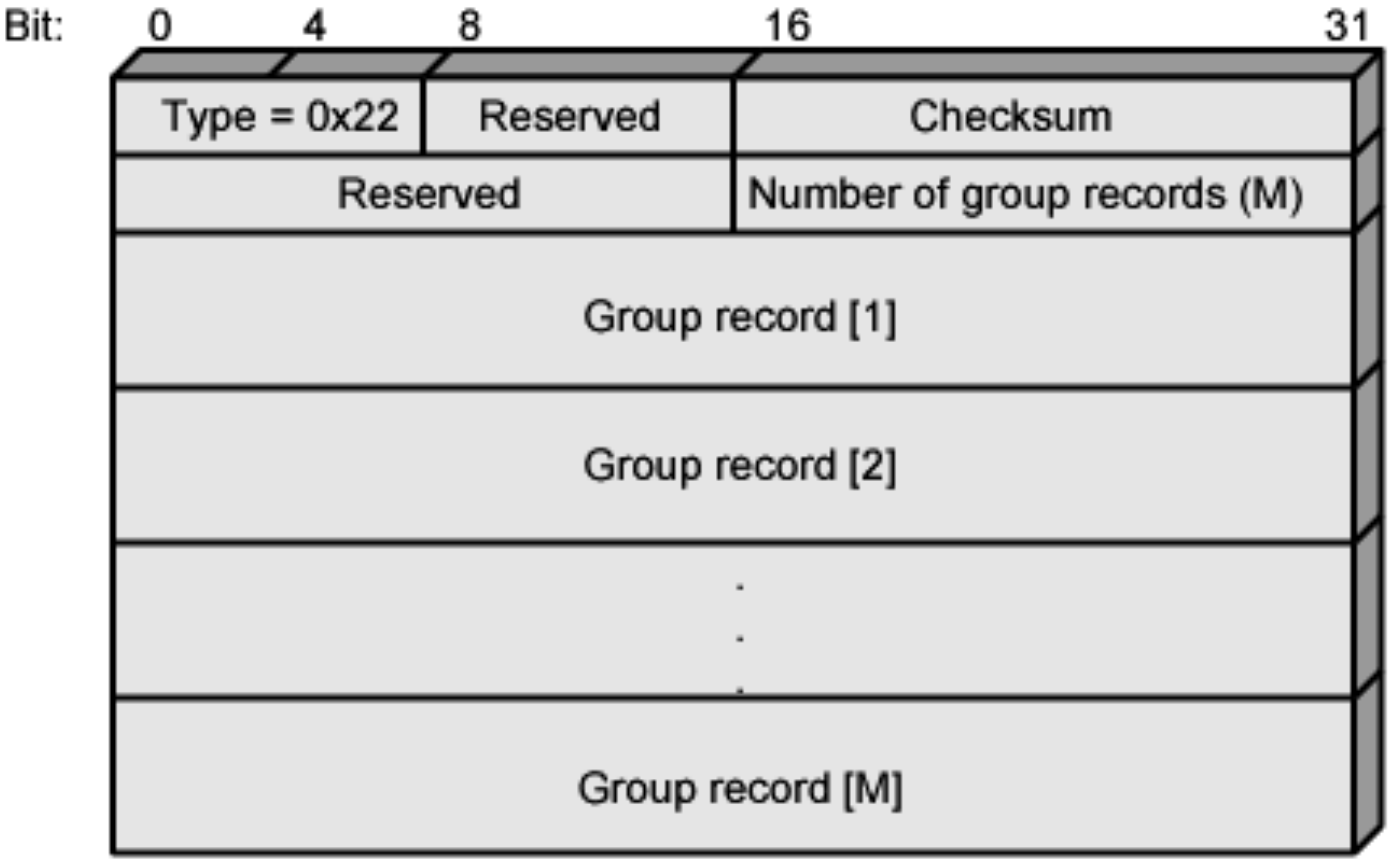

(b) Membership report message

Membership Reports

- Type
- Checksum
- Number of Group Records
- Group Records
  - One 32-bit unicast address per source

IGMP Message Formats: Group Record

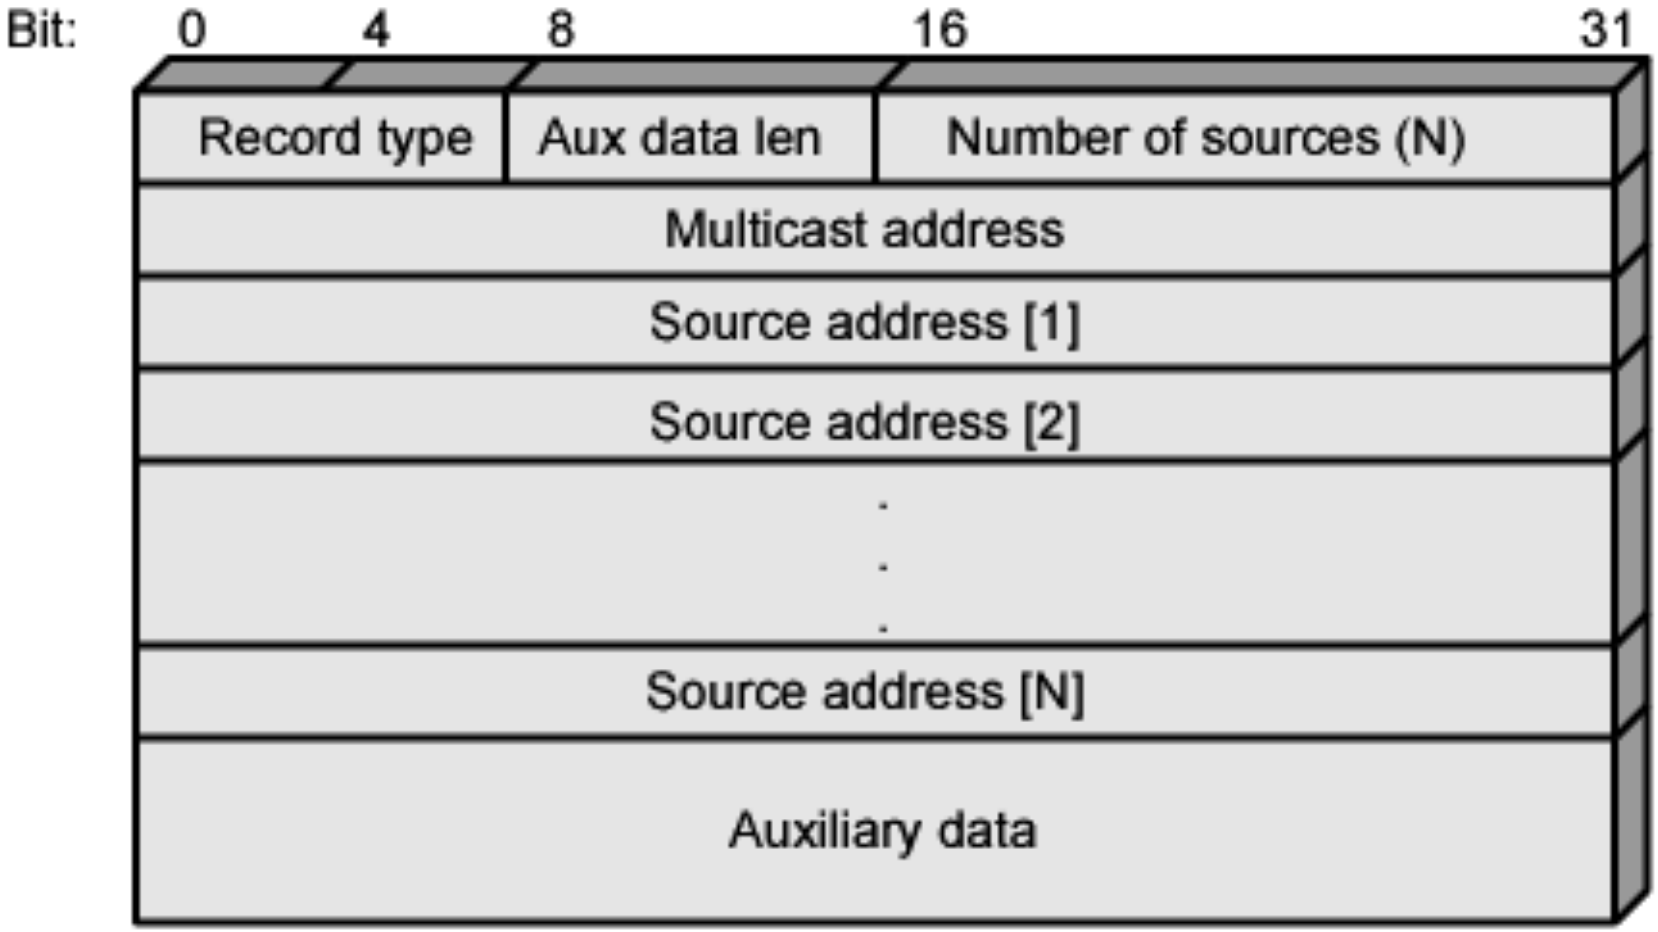

(c) Group record

Group Record

- Record Type
  - **"Current-State Record"** (in response to a Query)
    - MODE_IS_INCLUDE INCLUDE()
    - MODE_IS_EXCLUDE EXCLUDE()
  - **"Filter-Mode-Change Record"** (when the filter mode change)
    - CHANGE_TO_INCLUDE_MODE TO_IN()
    - CHANGE_TO_EXCLUDE_MODE TO_EX()
  - **"Source-List-Change Record"** (when the source list change)
    - ALLOW_NEW_SOURCES ALLOW()
    - BLOCK_OLD_SOURCES BLOCK()
- Aux Data Length
  - In 32-bit words
- Number of Sources
- Multicast Address
- Source Addresses
  - One 32-bit unicast address per source
- Auxiliary Data
  - Currently, no auxiliary data values defined

IGMP Operation - Joining

- Host using IGMP wants to make itself known as group member to other hosts and routers on LAN
- IGMPv3 can signal group membership with filtering capabilities with respect to sources
  - EXCLUDE mode – all group members except those listed
  - INCLUDE mode – Only from group members listed
- To join group, host sends IGMP membership report message
  - Address field contains the group multicast address
  - Sent in IP datagram with Group Address field of IGMP message and Destination Address encapsulating IP header same
  - Current members of group receive and learn about new member
  - Routers listen to all IP multicast addresses to hear all reports

IGMP Operation – Keeping Lists Valid

- Routers periodically issue IGMP general query message
  - In datagram with all-hosts multicast address
  - Hosts that wish to remain in groups must read datagrams with this all-hosts address
  - Hosts respond with report message for each group to which it claims membership
- Router does not need to know every host in a group
  - Needs to know at least one group member still active
  - Each host in group sets timer with random delay
  - Host that hears another claim membership cancels own report
  - If timer expires, host sends report
  - Only one member of each group reports to router

IGMP Operation - Leaving

- Host leaves group, by sending leave group message to all-routers static multicast address
- Send membership report message with EXCLUDE option and null list of source addresses
- Router determine if there are any remaining group members using group-specific query message

Multicast Extension to OSPF (MOSPF)

- Enables routing of IP multicast datagrams within single AS
- Each router uses MOSPF to maintain local group membership information
- Each router periodically floods this to all routers in area
- Routers build shortest path spanning tree from a source network to all networks containing members of group (Dijkstra)
  - Takes time, so on demand only

Forwarding Multicast Packets

- If multicast address not recognised, discard
- If router attaches to a network containing a member of group, transmit copy to that network
- Consult spanning tree for this source-destination pair and forward to other routers if required

Equal Cost Multipath Ambiguities

- Dijkstra's algorithm will include one of multiple equal cost paths
  - Which depends on order of processing nodes
- For multicast, all routers must have same spanning tree for given source node
- MOSPF has tiebreaker rule

Inter-area Multicasting

- Multicast groups may contain members from more than one area
- Routers only know about multicast groups with members in its area
- Subset of area's border routers forward group membership information and multicast datagrams between areas
  - Inter-area multicast forwarders

Inter-AS Multicasting

- Certain boundary routers act as inter-AS multicast forwarders
  - Run and inter-AS multicast routing protocol as well as MOSPF and OSPF
  - MOSPF makes sure they receive all multicast datagrams from within AS
  - Each such router forwards if required
  - Use reverse path routing to determine source
    - Assume datagram from X enters AS at point advertising shortest route back to X
    - Use this to determine path of datagram through MOSPF AS

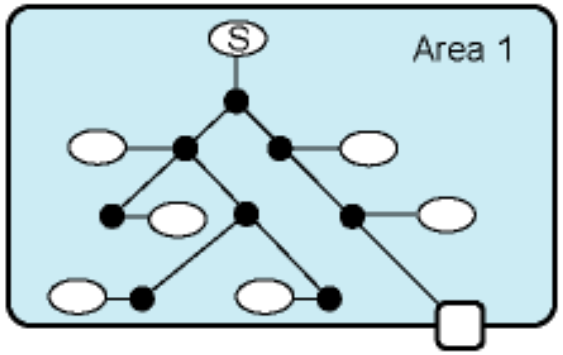

(a) Inter-Area Routing: Source in Same Area

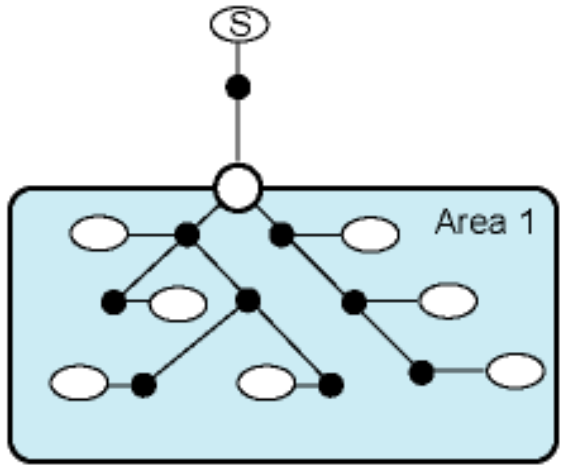

(b) Inter-Area Routing: Source in Remote Area

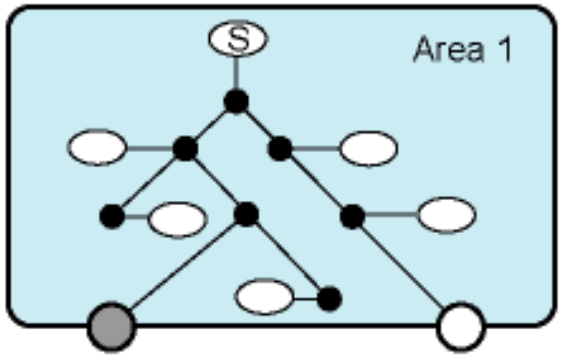

(c) Inter-AS Routing: Source in Same Area

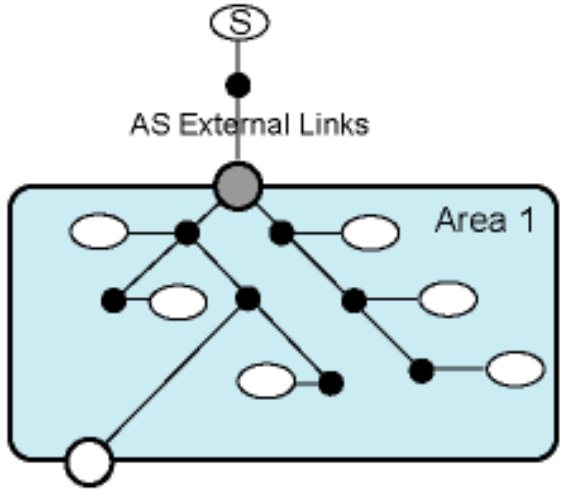

(d) Inter-AS Routing: Source in Different AS

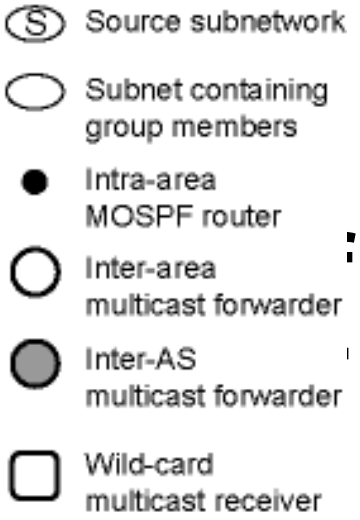

Multicast Routing Protocol Characteristics

- Extension to existing protocol
  - MOSPF v OSPF
- Designed to be efficient for high concentration of group members
- Appropriate with single AS
- Not for large internet