IP Multicast

Hy335 lab

Unicast/Multicast

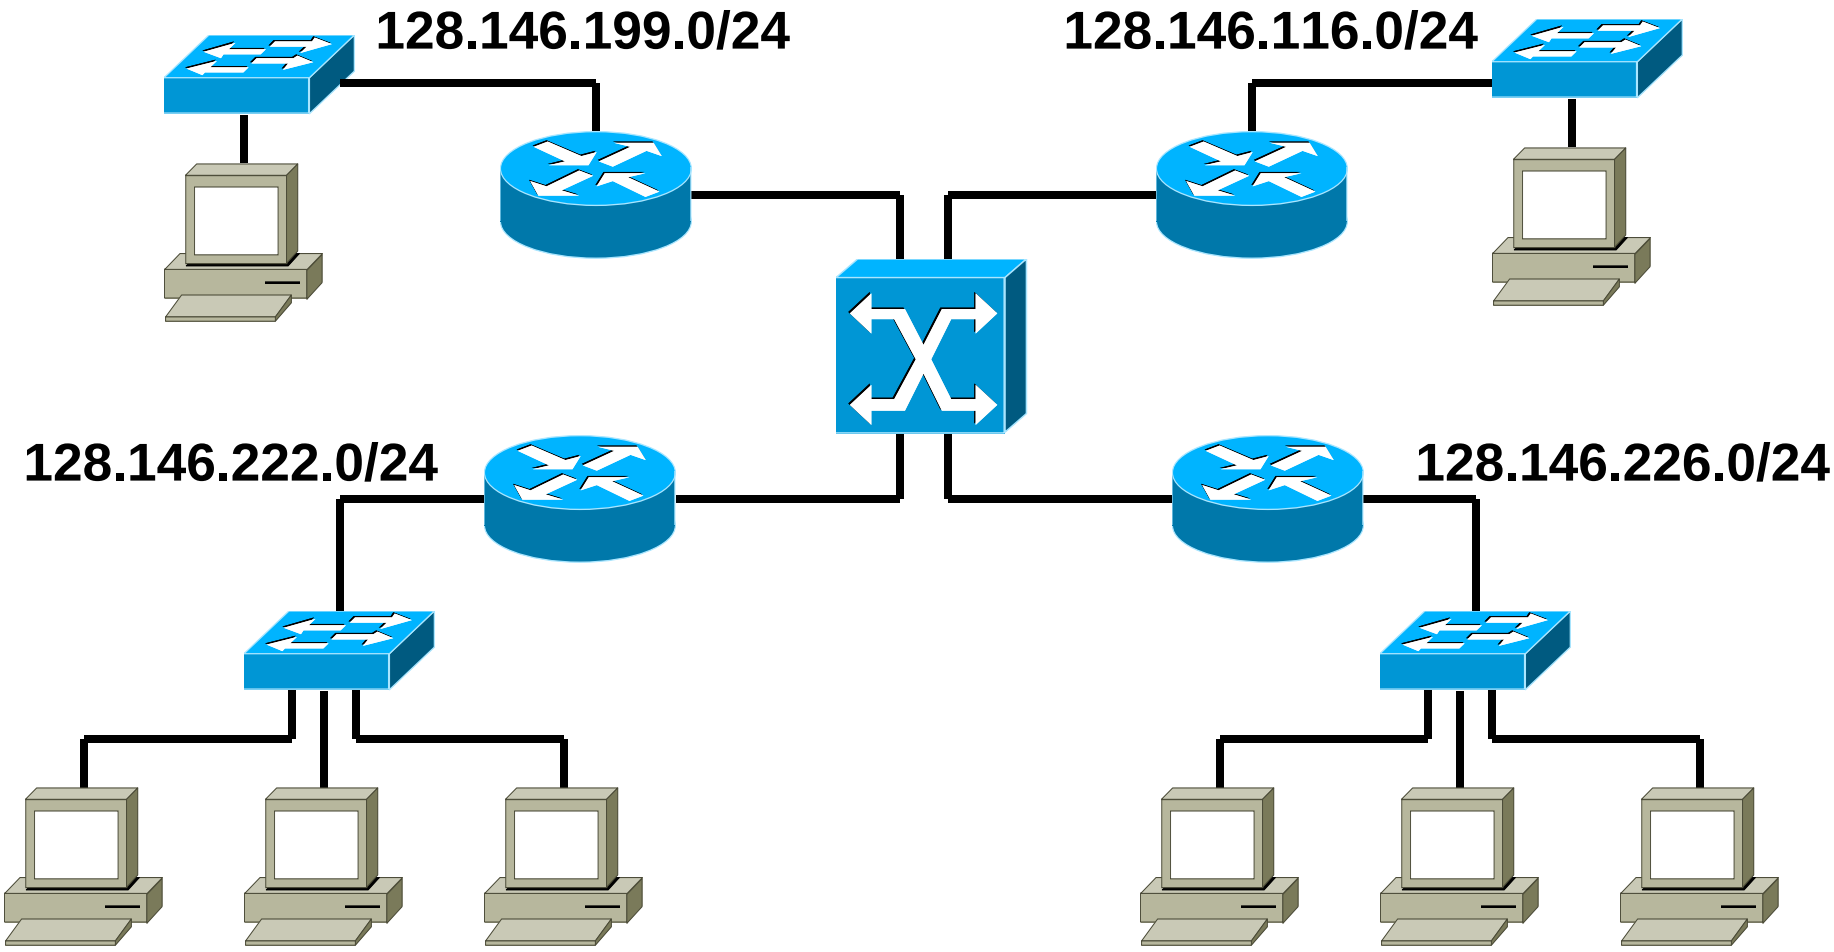

Ref: www.itec.oar.net/oartech/old/oarnet-multicast.ppt

Unicast

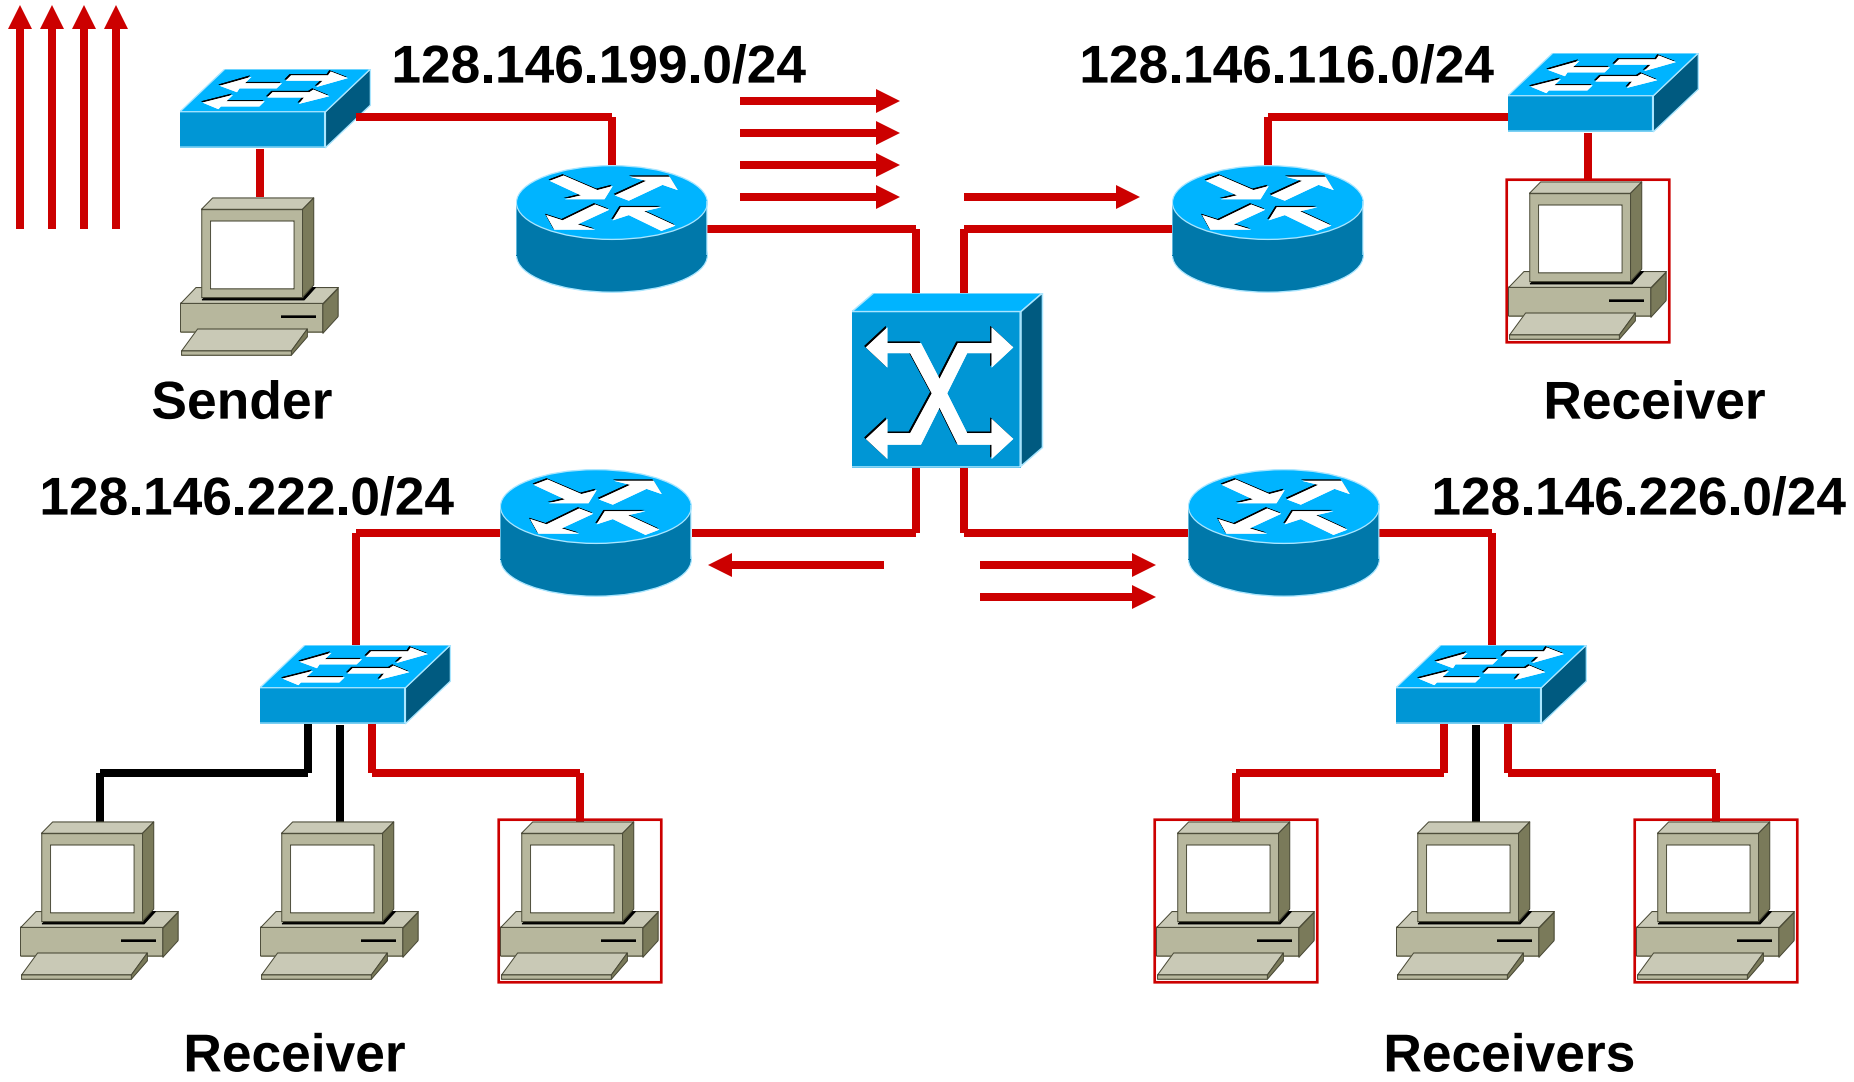

Ref: www.itec.oar.net/oartech/old/oarnet-multicast.ppt

Multicast

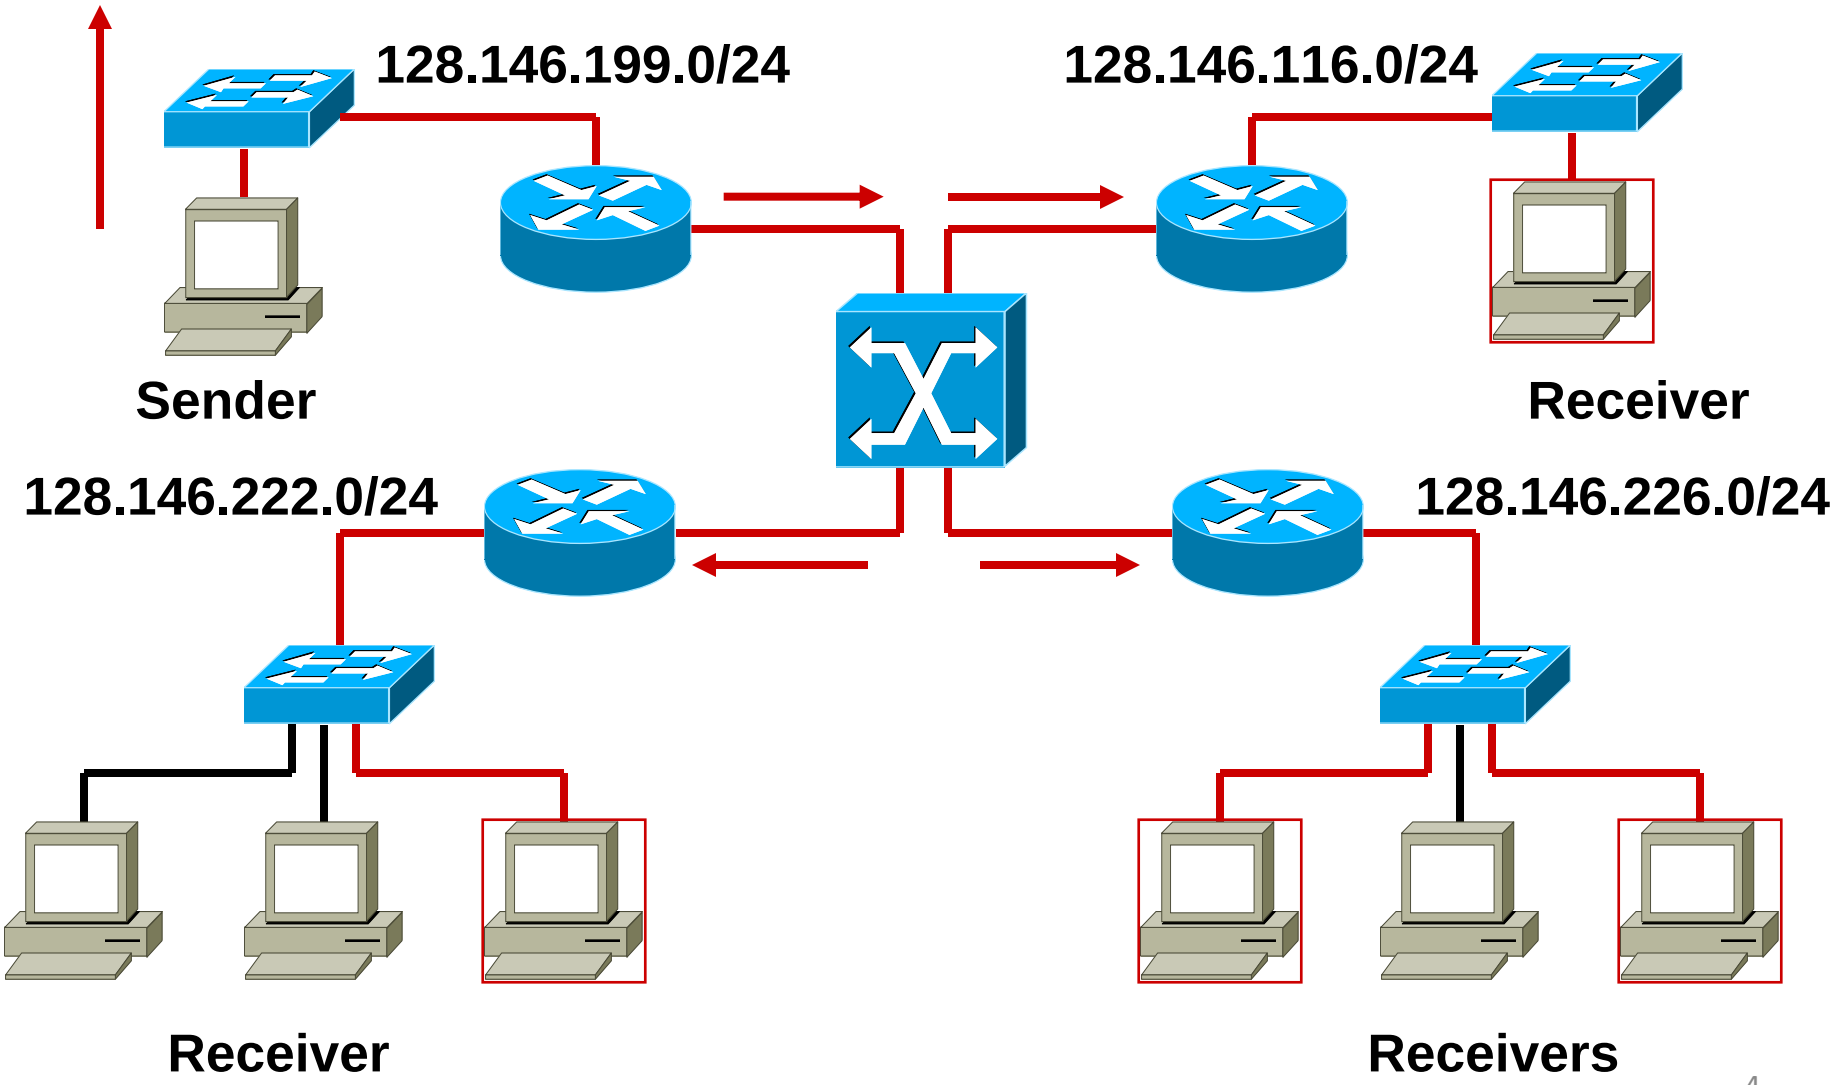

Ref: www.itec.oar.net/oartech/old/oarnet-multicast.ppt

Multicast Basic Concepts

- Previous TCP/UDP examples are all unicast
- Unicast: point to point communication
- Broadcast: packets are sent to all
  – IP supports broadcasting, but the use of broadcasts is strictly limited.
  – Protocols require broadcasts only when there is no alternative
  – Routers limit broadcasts to the local network or subnet, preventing broadcasts form reaching the Internet at large.
- Multicast: send packets to many different hosts, but not to everyone.

Multicast Examples

- Video conferencing
- DNS routers
- News group
- Multiplayer games
- Distributed file systems
- Massively parallel computing
- Database replication
- Name services
- Directory services

Multicast Basic Concepts

- Multicasting has been designed to fit into the Internet as seamlessly as possible.
- Most of the work is done by routers and should be transparent to application programmers.
- An application simply sends datagram packets multicast IP address. The router makes sure that the packets are delivered to all hosts in the multicast group.
- Big problem: multicast routers are not yet ubiquitous ➔ find out if multicasting is supported in your network.

Multicast Basic Concepts

- TTL: time to live in IP header
  - TTL is the maximum number of routers that the datagram is allowed to cross.
  - When a datagram reaches the maximum, it is discarded.
  - Multicasting uses TTL as an ad hoc way to limit how far a packet can travel.
- Multicast: UDP protocol

Multicast Address and Groups

- A multicast address is the address of a group of hosts called multicast group
- Multicast addresses are IP addresses
  - Calss D
  - Range 224.0.0.0 to 239.255.255.255
  - First 4 bits: 1110
- Like any IP address, a multicast address can have a hostname
  - 224.0.1.1 = ntp.mcast.net (network time protocol)

Multicast Address and Groups

- A multicast group is a set of Internet hosts that share a multicast address
- Any data sent to the multicast address is relayed to all the members of the group
- Membership in a multicast group is open; hosts can enter or leave the group at any time
- Groups can be either permanent or transient
  - Permanent groups have assigned address that remain constant
  - Most multicast groups are transient and exist only as long as they have members.

IP Multicast Group Concept

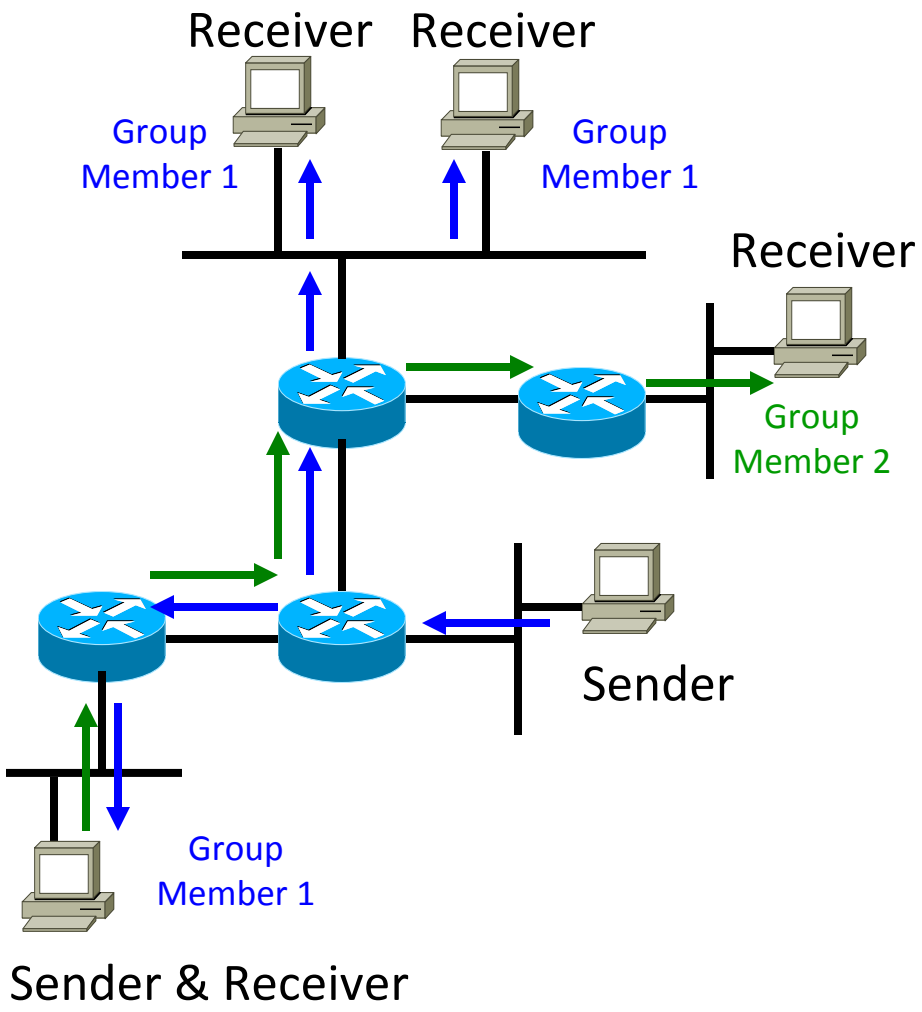

- Receivers must be a member of the group to receive the data stream
- If sender send to group address, all members receive it
- Multicast Router find delivery paths and replicate the packet that enable forwarding across Internet

Router and Routing

- Routing with/without multicasting sockets
- With multicasting:
  - a multicast socket sends one stream of data over the Internet to the clients' router.
  - The router duplicates the stream and sends it to each of the clients.
- Without multicasting:
  - The server sends four separate but indintical stream of data to the router
  - The router each of the stream to a client.

Router and Routing

- Note that real-world routes can be much more complex, involving multiple hierarchies of redundant routers
- Goal of multicast sockets:
  - No matter how complex the network, the same data should never be sent more than once over any given network
  - Programmers don’t need to worry about routing issues.
- To send and receive multicast data beyond the local subnet, you need a multicast router
  - Ping all-routers.mcast.net

Router and Routing

```
>Ping all-routers.mcast.net
Pinging all-routers.mcast.net [224.0.0.2] with 32 bytes of data:

Reply from 224.0.0.2: bytes=32 time<10ms TTL=128
Reply from 224.0.0.2: bytes=32 time<10ms TTL=128
Reply from 224.0.0.2: bytes=32 time<10ms TTL=128
Reply from 224.0.0.2: bytes=32 time<10ms TTL=128

Ping statistics for 224.0.0.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms
```

IP addresses

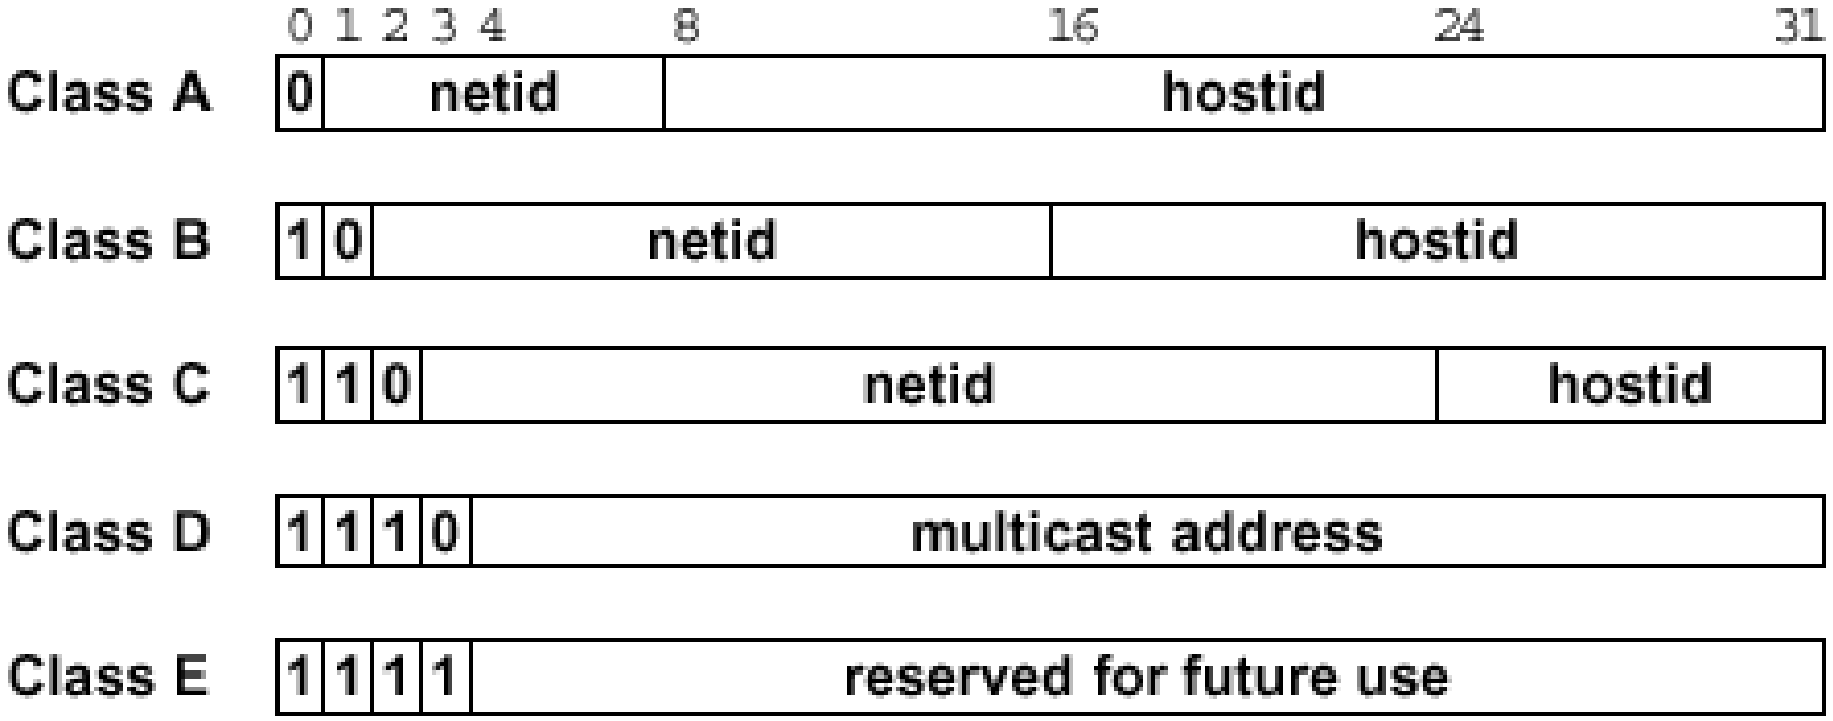

Multicast Addressing

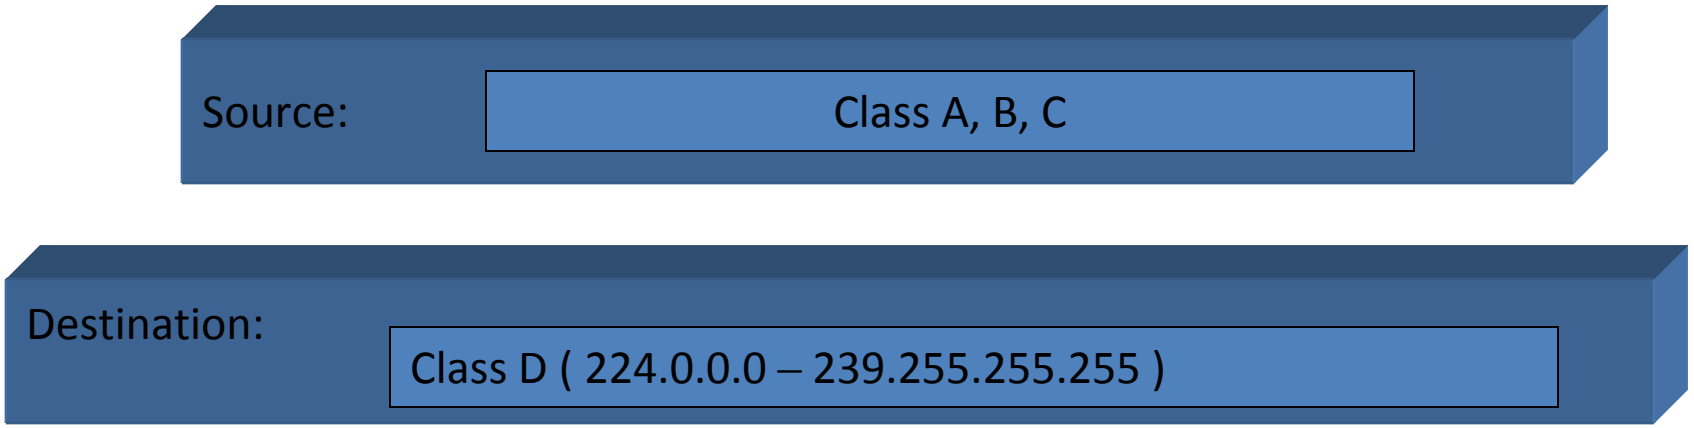

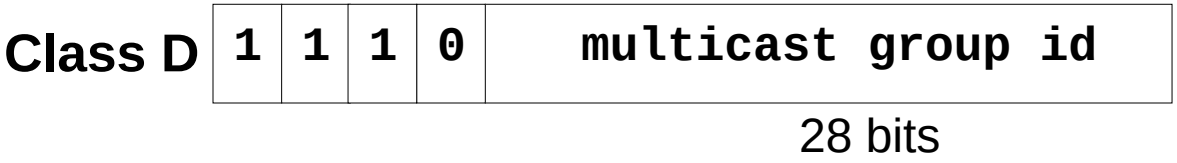

- **Source Address** can never be Class D (Multicast Group Address)

Receiver-Router Signaling: IGMP

- How receivers tell router about group membership
- Router solicit group membership from directly connected receivers

IGMP: Internet Group Management Protocol

- is the protocol which **Receivers** exchange the information about to become a member of multicast group with their **Local Multicast Router**
- provide a method through which a receiver can join or leave a multicast group at any time

What about IGMP rule ?

- Receiver sends an **IGMP report** for joining a group
- Receiver will never send a report when it wants to leave a group
- Multicast routers send **IGMP query** to the all receiver group to see whether any group members exists on their

Receiver-Router Signaling: IGMP

Joining a Group

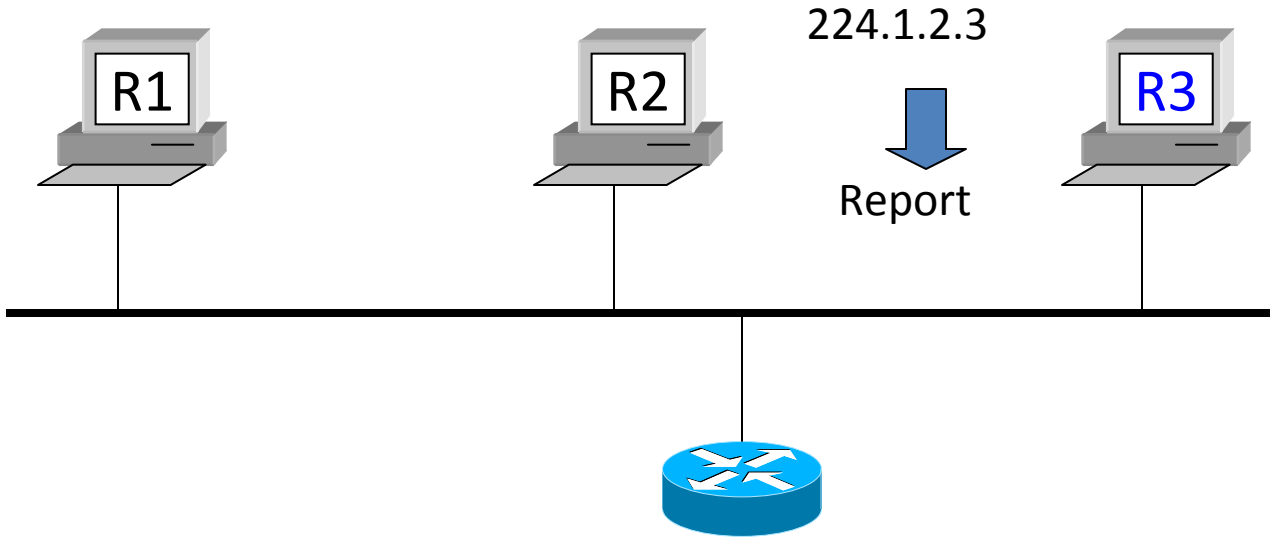

- Receiver sends IGMP Report to join group to Local Multicast Router

Receiver-Router Signaling: IGMP

Maintaining a Group

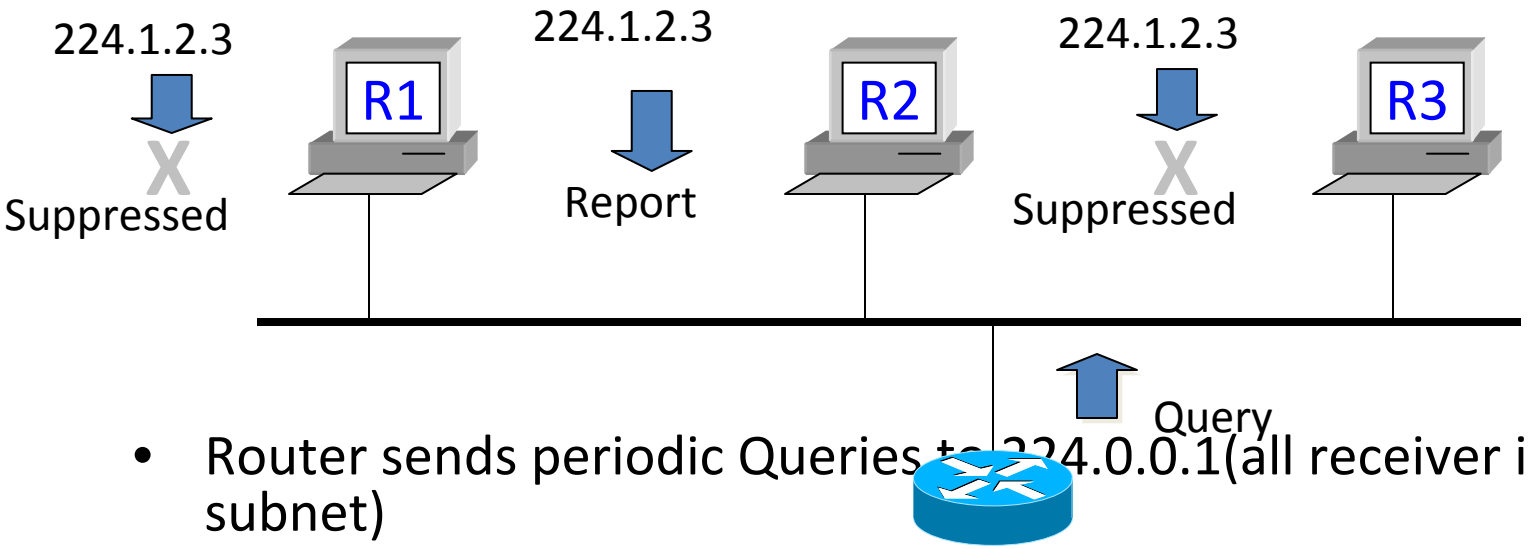

- Router sends periodic Queries to 224.0.0.1(all receiver in the subnet)
- One member per group per subnets reports
- Other members suppress reports

Multicast Address and Groups

- Create a multicast group
  - Pick an random address from 225.0.0.0 to 238.255.255.255
- A number of multicast addresses have been assigned for special purposes.
  - all-systems.mcast.net (224.0.0.1) is a multicast group that includes all systems that support multicasting on local subnet
  - This group is commonly used for local testing
  - Also for local testing experiment.mcast.net (224.0.1.20)

Multicast Address and Groups

- A number of multicast addresses have been assigned for special purposes. (cont.)
  - (224.0.0.0~ 224.0.0.255) are reserved for routing protocols (gateway discovery …)
  - Multicast routers never forward datagrams with destinations in 224.0.0.0~ 224.0.0.255
  - IANA is responsible for handing out permanent multicast addresses
  - About 10,000 have been assigned
  - Still have 248 Million class D addresses can be used.

Clients and Servers

- When a host wants to send data to a multicast group, it puts that data in multicast datagrams (UDP datagrams with an IP address in class D)
- Most multicast data is either audio or video or both.(Small data lost is fine.)
- Multicast data is sent via UDP
- UDP can be as much as three times faster than TCP

Support for Multicast

- The routers must have multicasting ability
  - execute a multicast routing protocol to find delivery paths that enable forwarding of multicast packets
  - replicate the packet and forward it down the appropriate downstream paths if there are multiple downstream paths
- if all router cannot do multicasting, use a virtual network (multicast backbone: MBONE)
  - allow multicast packets to travel through routers that are set up to handle only unicast traffic

Datagram Format

- TTL: time to live
  - One byte

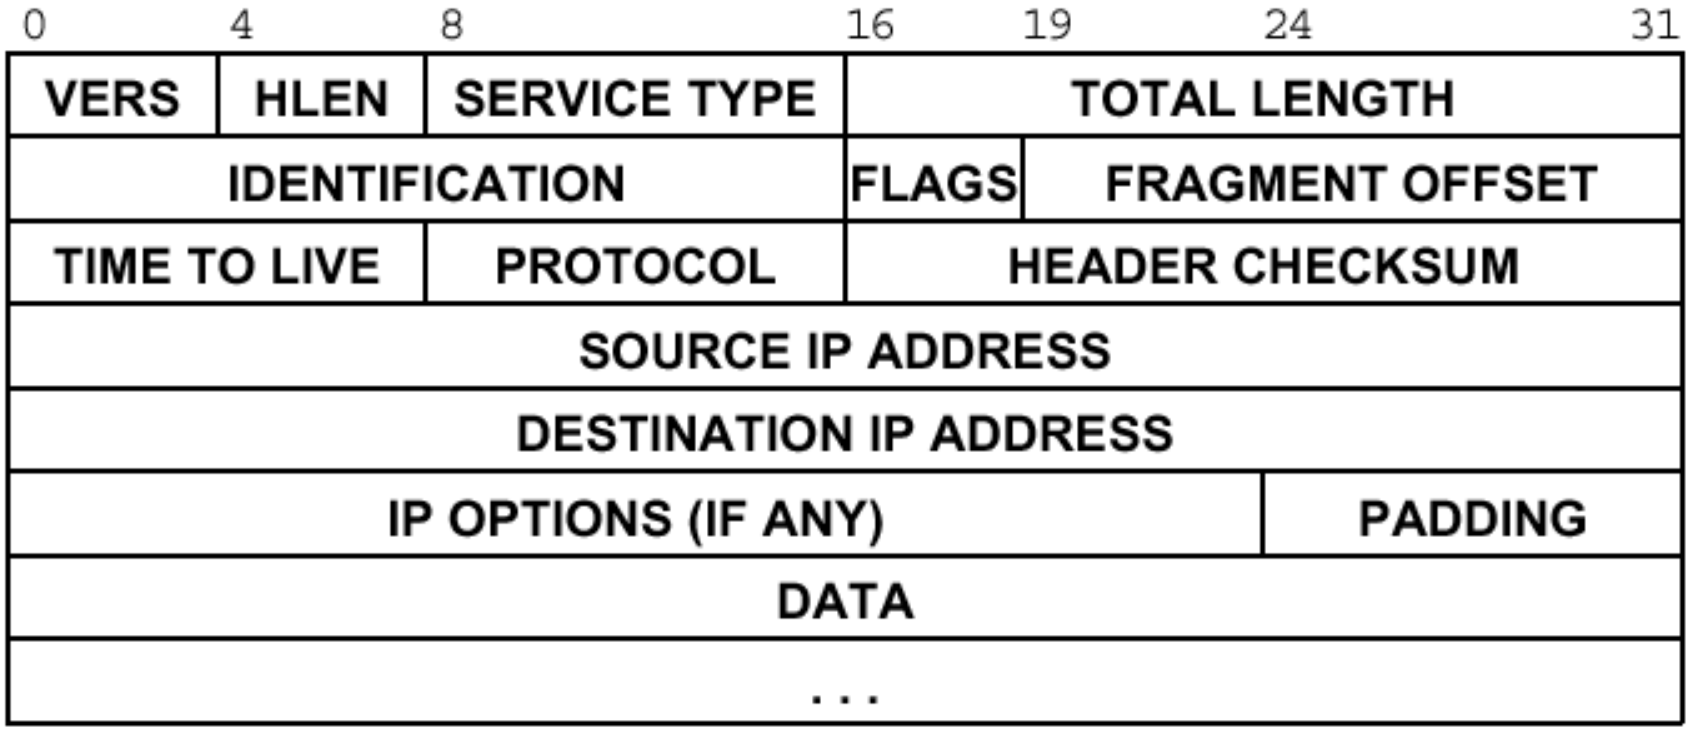

TTL

- Routers and hosts must decrement the *TIME TO LIVE* field by one and remove the datagram from the internet when its time expires.
- In practice, the TTL acts a “hop limit” rather than an estimate of delays.
- Two uses:
  - It guarantees that datagrams cannot travel around an internet forever.
  - Source might want to intentionally limit the journey of the packet.

TTL

- TTL: the number of hops
- Each time a packet passes through a router, its TTL value is decremented by at least one
  - Some routers may decrement the TTL by two or more.
- When the TTL reaches zero, the packet is discarded.
- All packets would eventually be discarded
- TTL may prevent mis-configured routers from sending packets back and forth to each other indefinitely

TTL

- In IP multicasting, TTL is used to limit the multicast geographically.
  - TTL = 0: local host
  - TTL = 1: local subnet
  - TTL = 16: local campus or organization
  - TTL = 32: US backbone
  - TTL = 48: US
  - TTL = 64: North America
  - TTL = 128: high bandwidth sites worldwide
  - TTL = 255: All sites worldwide

Examples

- http://www.csd.uoc.gr/~hy335b/sender.c
- http://www.csd.uoc.gr/~hy335b/listener.c