


Chapter 10

Internet Group Management Protocol (IGMP)

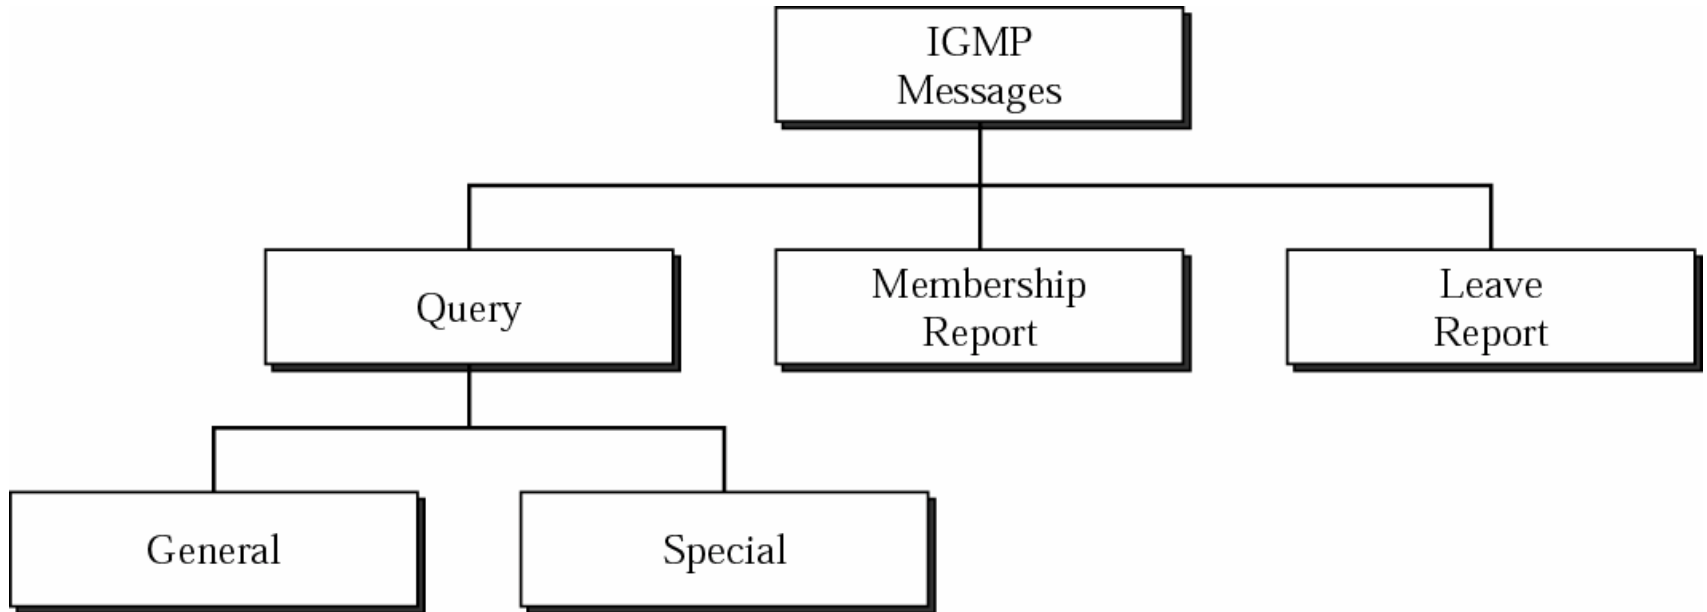
IGMP v2: RFC 2236, Xerox, November 1997

IGMP v1: RFC 1112, Stanford University, August 1989

IGMP is used by IP hosts to register their dynamic multicast group membership. It is also used by connected routers to discover these group members.



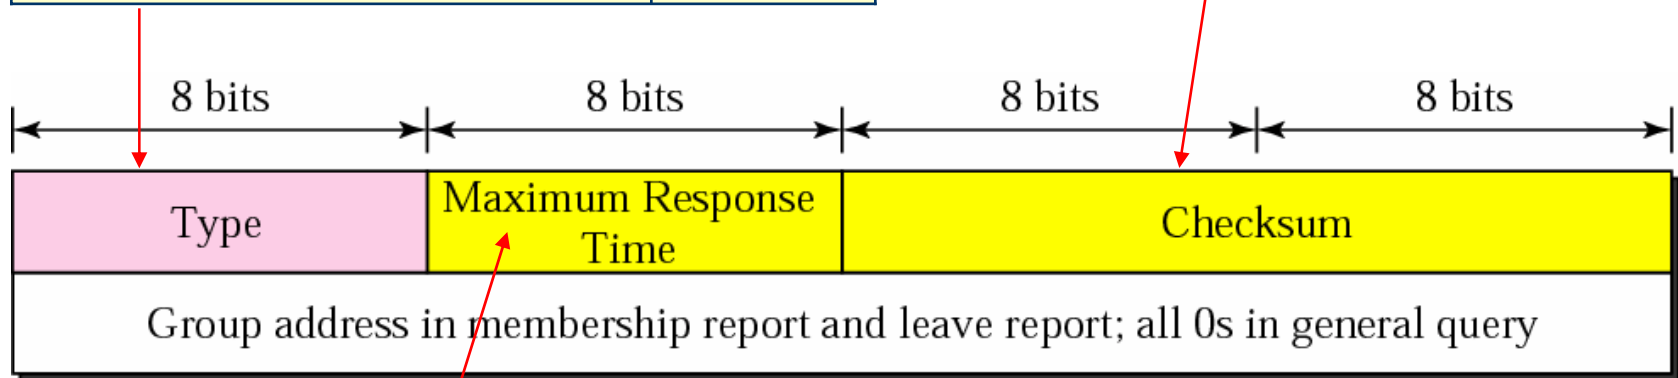
IGMP Message Types



IGMP Message Format

<i>Message Type</i>	<i>Value</i>
General or special query	0x11
Membership report	0x16
Leave report	0x17

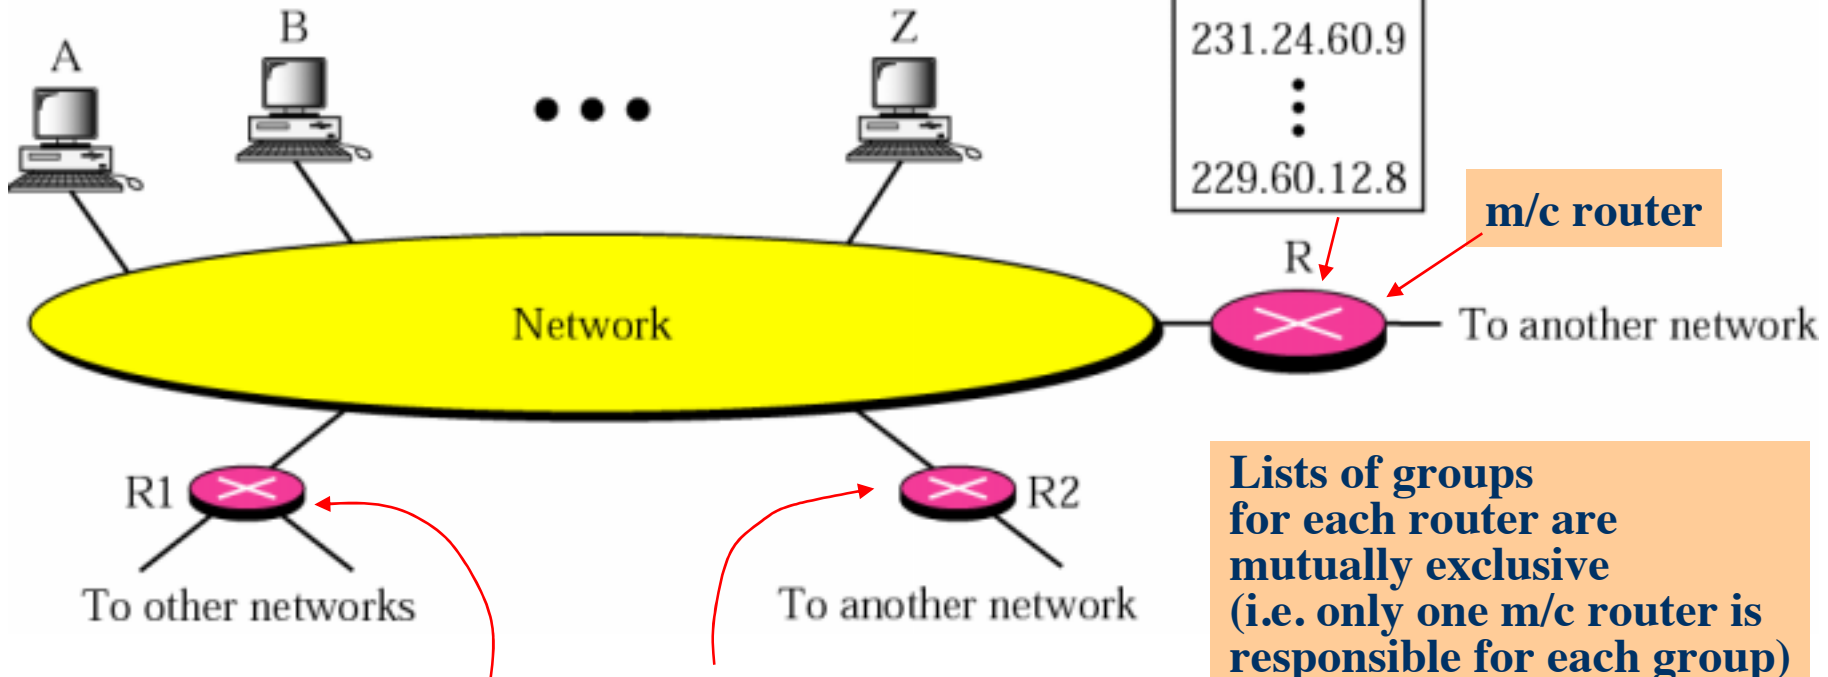
Covers the entire IGMP message



Amount of time (in 1/10 of sec) in which a query must be answered (set to 0 in non-query messages)

IGMP Operation

Each host can have several processes.
The processes can have interest for a group.
Each host maintains a list of groups for which its processes have an interest.



These routers can be distributors for one or several groups from the list of router R, but for other networks, not for the network above.

Lists of groups for each router are mutually exclusive (i.e. only one m/c router is responsible for each group)
A router can also be a member of a group


Membership Report

Host or a router can join a group

A host maintains a list of processes that have membership in a group

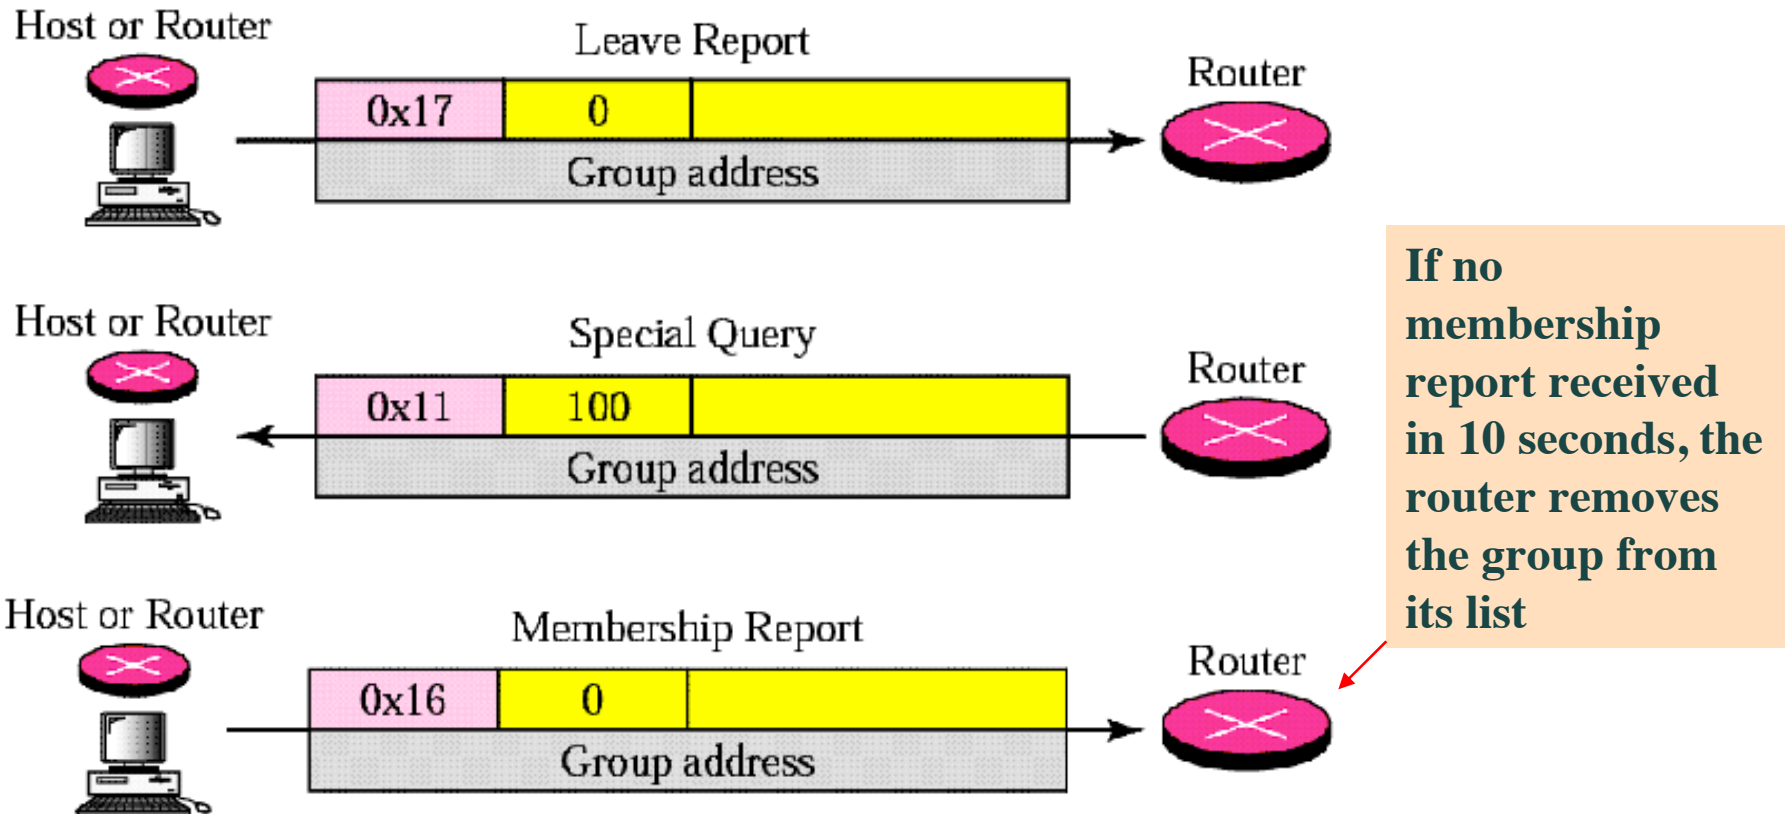
When a process wants to join a new group, it sends its request to the host

The host then adds the name of the process and the name of the requested group to its list, and sends the membership report to the router. The report is send twice (in case the first report get lost or damaged)



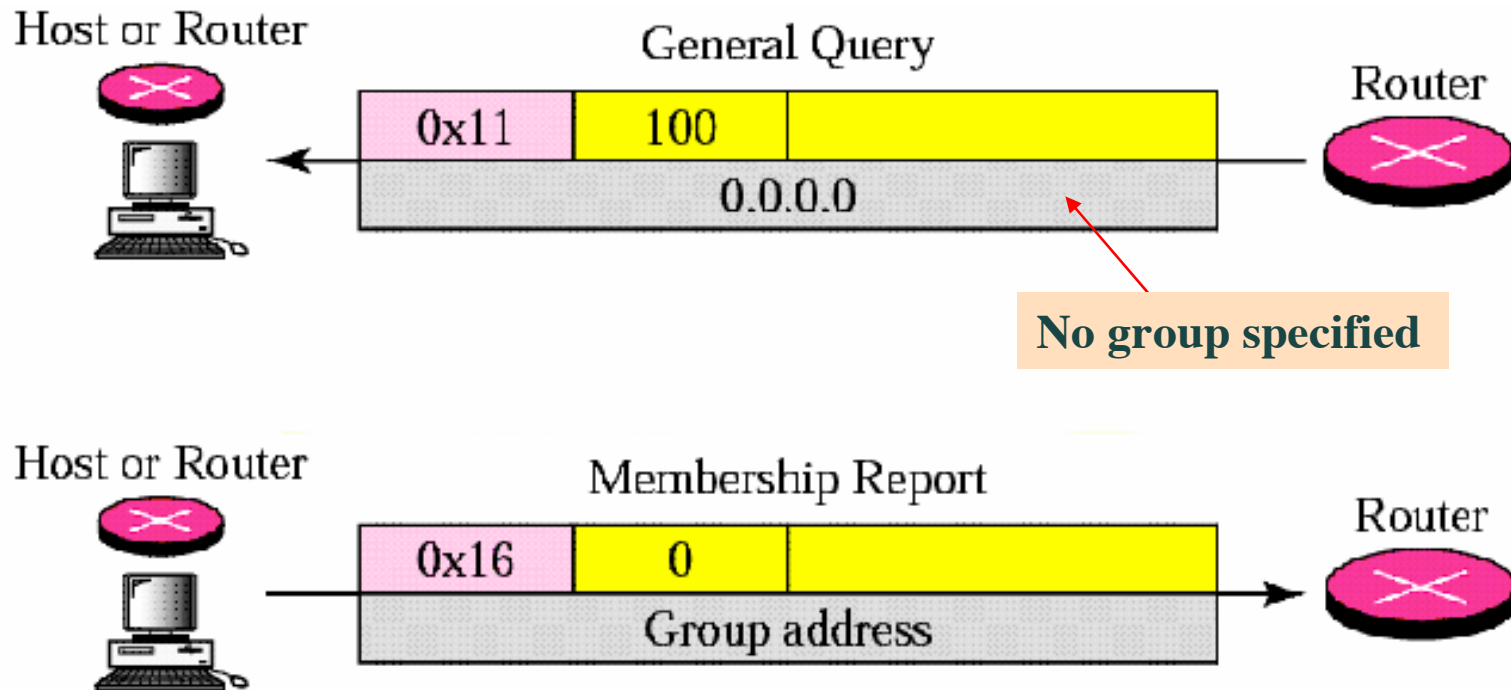
Leave Report

When a host sees that no process is interested in a specific group G, it sends a **leave report**. If router receives a leave report it won't purge the list if there are still other hosts interested in that group. For that purpose the router sends a **special query message** with a specified response time for the group in question to see if there is anyone interested in that group. If there is no response a **membership report**, it purges the list.



General Query Message

Membership report and leave report are not enough to maintain the membership information. Example: a host that is a member of a group can shut down and the m/c router would never receive the leave report. Therefore the m/c router monitors the hosts and routers in LAN by periodically sending (by default every 125 sec) general query message. Hosts/routers respond by membership report if there is still interest in groups



Delayed Response

In order to keep the traffic low the response to general query message must be done by only one host for a given group. How can be made sure that only one host answers the query, while the others which have to report the same group, or groups are not? This is achieved with delayed response:


When a host receives general query message it delays the response: it sets a timer for each group to a different random value between 0 and 10 seconds, then broadcasts the response(s) according to the timers.

If the host receives a response from another host, whose timer for that group has expired earlier, the host cancels the corresponding timer and doesn't send the duplicate response for the group.

Only one router on the LAN is designated for sending the query messages – the query router. This further reduces the traffic.

Example 1

A query message was received at time 0; the random delay time (in tenths of seconds) for each group is shown next to the group address. Show the sequence of report messages.



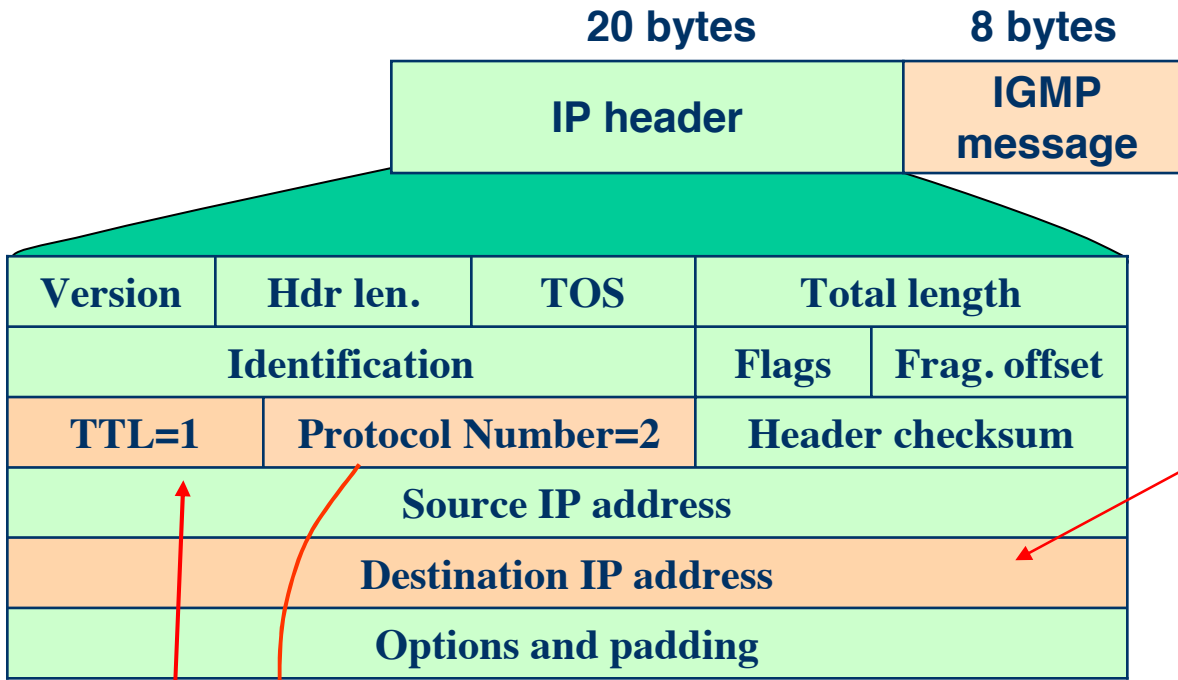
Solution:

<i>Time</i>	<i>Group</i>	<i>From host</i>
12	228.42.0.0	A
30	255.14.0.0	A
50	251.71.0.0	B
70	230.43.0.0	C

Four messages are sent instead of seven messages

Encapsulation of IGMP Messages

IGMP messages are encapsulated in IP datagrams



All systems on this subnet

Query	224.0.0.1
Membership report	m/c addr. of the group
Leave report	224.0.0.2

All routers on this subnet

IGMP message doesn't leave the LAN

<i>Prot. Num.</i>	<i>Protocol</i>	<i>Description</i>
1	ICMP	Internet Control Message
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway
4	IP	IP in IP (encapsulation)
.....

Encapsulation of IGMP Messages (cont.)

Assigned m/c Addresses (Category Addresses)

<i>Address</i>	<i>Group</i>
224.0.0.0	Reserved
224.0.0.1	All systems on this subnet
224.0.0.2	All routers on this subnet
224.0.0.4	DVMRP routers
224.0.0.5	OSPF/IGMP all routers
224.0.0.6	OSPF/IGMP designated routers
224.0.0.7	ST routers
224.0.0.8	ST hosts
224.0.0.9	RIP2 routers
224.0.0.10	IGRP routers
224.0.0.11	Mobile agents

These are only few examples. The full list of assigned multicast addresses is given in <http://www.iana.org/assignments/multicast-addresses>

Encapsulation of IGMP Messages (cont.)

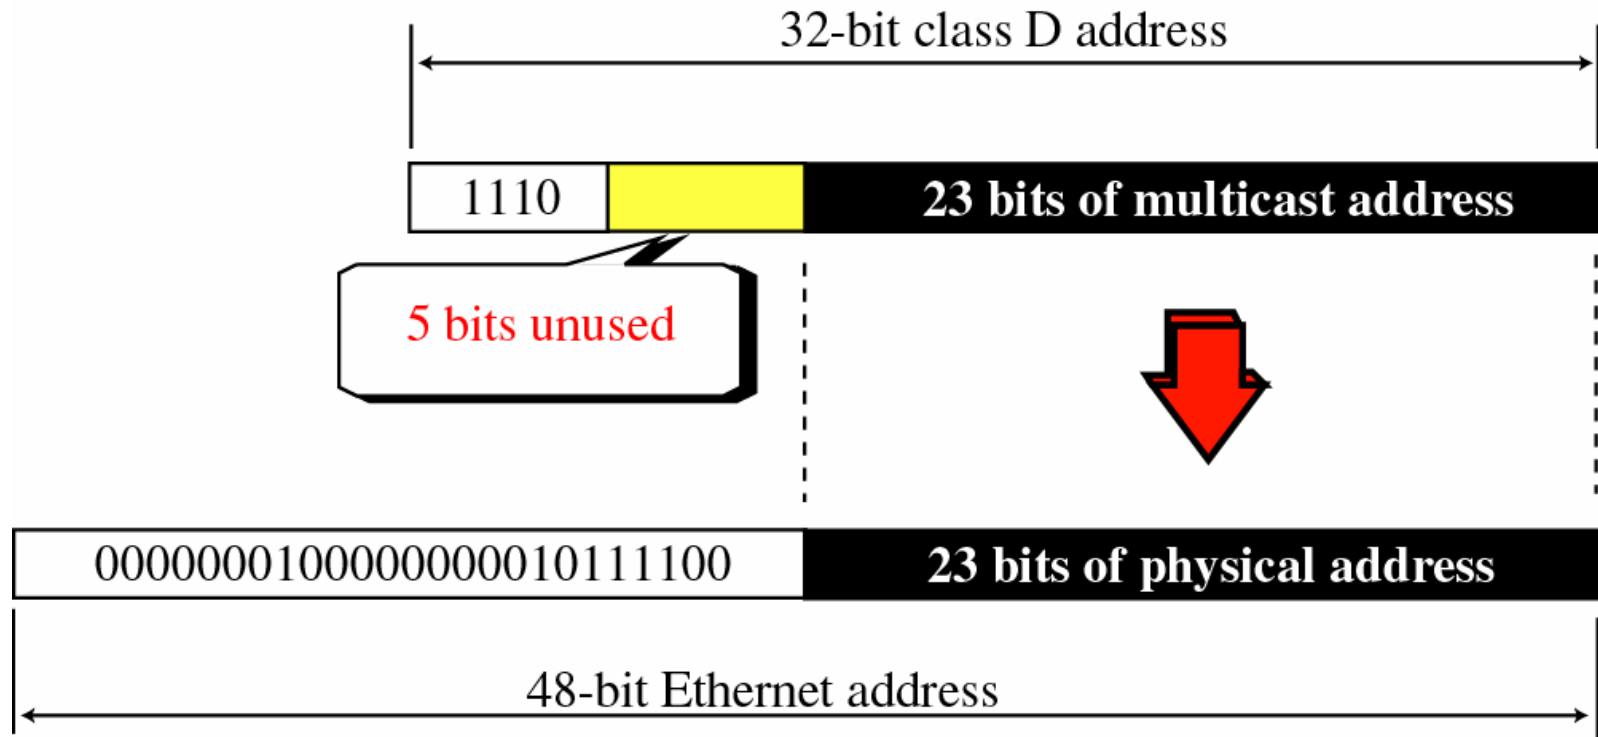
Comment:

Membership report uses the m/c address of reported group as the destination address in the IP header. This message is meant to every host or router that are members of the group, so they can enforce the delayed response principle. Other hosts that are not members of the group need not to receive the membership report. If the group is new to LAN the distributing router would be the only member of the group.

The m/c group address is redundant with the address given in the IGMP message itself. This is a small overhead which however doesn't require the hosts to decapsulate the IGMP message – the delayed request timers can be set based on the IP header.

Mapping class D to Ethernet physical address

In order to encapsulate a m/c IP packet into layer 2 frame the physical address must be resolved. In this case it would be very difficult to apply the usual ARP. Instead, the m/c IP address is mapped into m/c MAC address by a simple mapping as shown in the previous chapter:



Most of LANs support m/c addressing. Ethernet included.

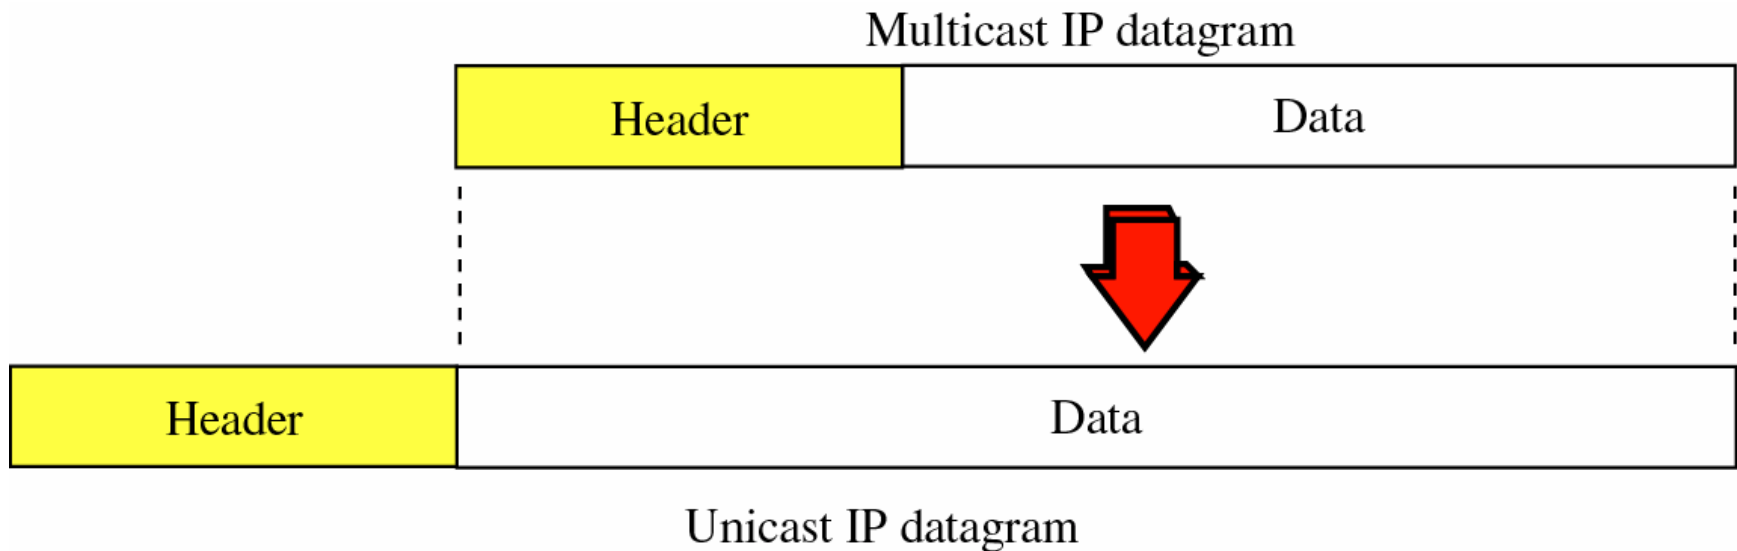
Mapping class D to Ethernet physical address (cont.)

Comment:

The mapping shown on the previous slide is many-to-one mapping because there are $28-23=5$ bits that are supposed to be all zero. In other words 32 different addresses can map into the same MAC address. Therefore the host that receives a m/c packet has to check the m/c IP address and discard the packet if the IP addresses doesn't match.

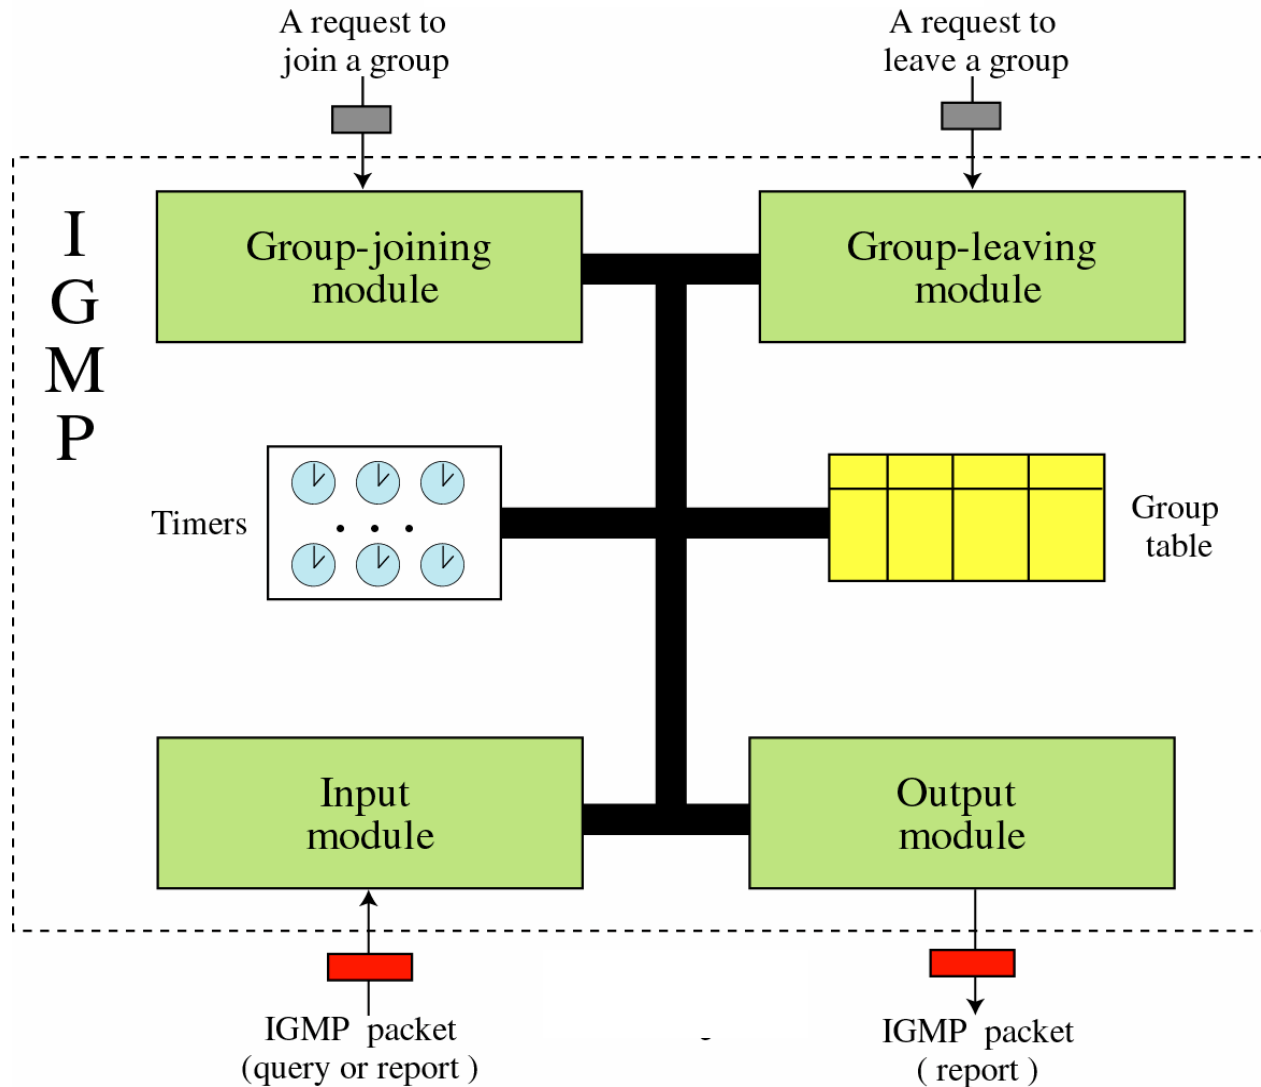
Tunneling

Most WANs do not support m/c MAC addresses (their second and physical layer are not Ethernet for example, but T3, SONET, FR, ATM etc.). In that case the IP tunneling must be applied: the m/c IP datagram is encapsulated in unicast IP datagram:



When the unicast IP packet gets decapsulated at the exit point of the tunnel, the networks there support the m/c addressing on the MAC layer.

IGMP Implementation of the Host



IGMP Implementation of the Host (cont.)

Group Table

The interface through which the m/c packet is sent or received

Number of processes still interested in the group

State	Interface No.	Group Address	Reference Count
.....
.....
.....

FREE – no processes left in the group
DELAYING – a report must be sent for this entry when the timer matures
IDLE – no timer running for this entry

Group-Joining Module

Module invoked by a process that wants to join a group

```
Receive request from a process  $P$  to join a group  $G$ ;  
if (entry  $\langle G, P, RC \rangle$  already exists)  
    Increment the reference count (RC);  
else  
{  
    Create the entry  $\langle G, P, RC=1 \rangle$ ;  
    Request the membership report from the output  
    module;  
    Inform LLC to update its configuration table;  
}  
return;
```

IGMP Implementation of the Host (cont.)

Group-Leaving Module

Module invoked by a process that wants to leave a group

```
Receive request from a process  $P$  to leave a group  $G$ ;  
if (entry  $\langle G, P, RC \rangle$  exists)  
{  
    Decrement  $RC$ ;  
    if ( $RC == 0$ )  
    {  
        Cancel the timer for this entry;  
        Change the state of the list to FREE;  
        Request the leave report from the output  
        module;  
    }  
}  
return;
```

Input Module

Module invoked by an IGMP message

```
Receive an IGMP message;
switch (message type){
  case 'query':
    Start a timer for each entry with state = IDLE;
    Change each IDLE state to DELAYING state1;
    return;
  case 'membership report with <G>':
    if (there is entry <G,P,state=DELAYING>)
    {
      Cancel the timer for this entry;
      Change the state to IDLE;
    }
}
return;
```

¹Membership report will be sent by output module as soon the timer expires (see next slide)

IGMP Implementation of the Host (cont.)

Output Module

Module invoked by a timer or a request joining or leaving module

```
Receive a signal from a timer or a request from
joining or leaving module;
if (signal from timer)
    if (there is an entry with state = DELAYING){
        Create a membership report;
        Reset the state to IDLE;
    }
if (request from group joining module)
    Create a membership report;
if (request from group leaving module)
    Create a leave report;
Send the message;
return;
```