

Networkshop, Nottingham, 2002

G Fairhurst, <http://www.erg.abdn.ac.uk>

(c) April 2002



IP Multicast in LANs

Gorry Fairhurst
Department of Engineering
University of Aberdeen
gorry@erg.abdn.ac.uk

4.v37, (c) Apr 2002

IP Multicast

IGMP

Switched Ethernet

Ten Thorny Issues

Questions & Answers

Internet Multicast Applications

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

One-to-Many

One sender
Audio/Video broadcast
Information push
Multicast file transfer

One-Way Multicast

*Single Source Multicast
(SSM)*

Many-to-Many

Many senders
Audio/video conferences
Distributed computing

Two-Way Multicast

Many-to-One

Many senders
Voting / Auctions
Control protocols

*Any Source Multicast
(ASM)*

IP Multicast Address

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002



Class D IP Multicast Address

32 bits starting with bit pattern "1 1 1 0"
224.0.0.0-239.255.255.255
224.0.0.X are reserved

IPv6

IPv6 has similar address allocations

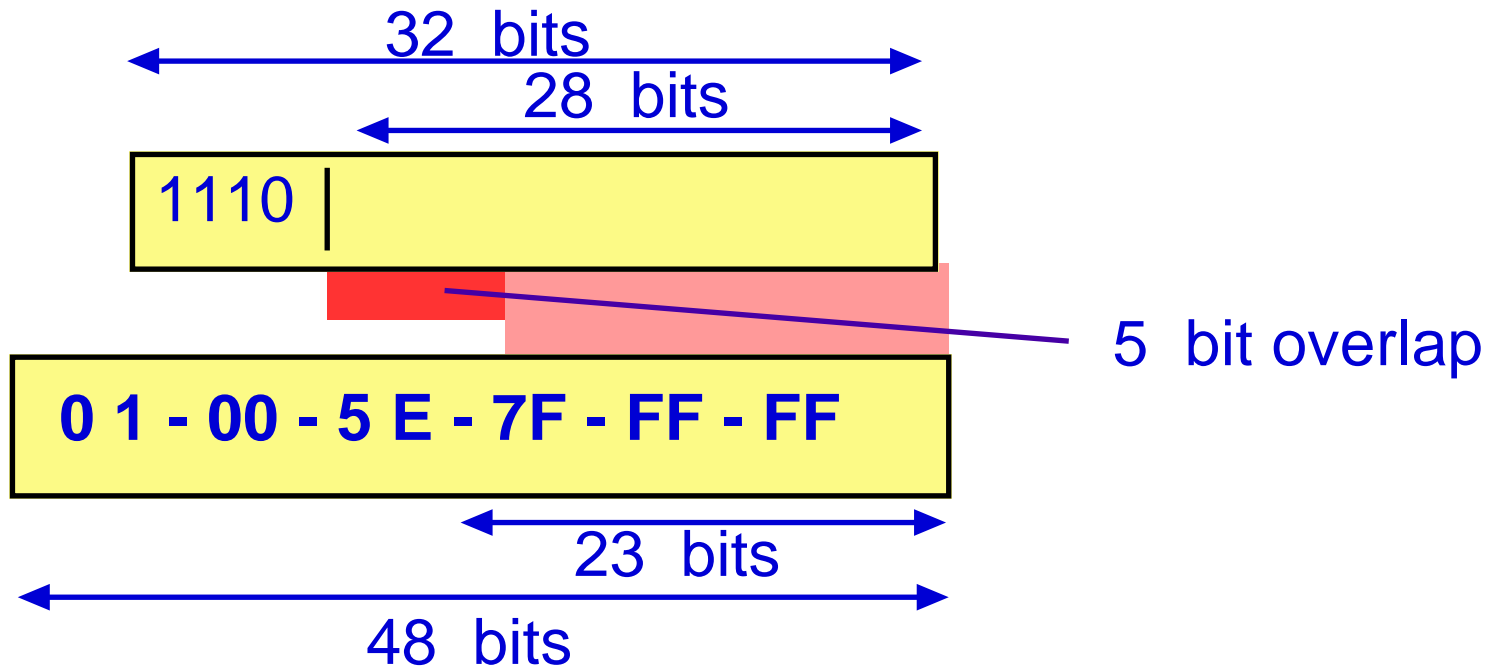
[RFC 3171]

IPv4 Multicast over Ethernet

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Class D IPv4 destination address

224.0.0.0-239.255.255.255



MAC hardware destination address

*One L2 (MAC) address
may carry multiple L3 (IPv4) addresses*

[RFC1122]

224.0.0.X

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

224.0.0.1 All multicast systems

224.0.0.2 All multicast routers

224.0.0.4 DVMRP

224.0.0.9 RIP2 Routers

224.0.0.13 PIM Routers

224.0.0.22 IGMPv3 Reports

**Multicast
Routing**

Packets with these addresses never travel outside LAN

*N.B. avoid using x.0.0.x or x.128.0.x for data
because of Ethernet address overlap*

[RFC1112; RFC3171]

UDP/IP Packet with RTP Payload

G Fairhurst, <http://www.erg.abdn.ac.uk>

ETHER

Packet size = 218 bytes

Destination = 1:0:5e:2:dc:3e, (multicast) (01-00-5e-02-dc-3e)

Source = 0:d0:bb:f7:c6:c0,

Ethertype = 0800 (IPv4)

IP

Version = 4, Header length = 20 bytes

Type of service = 0x00

Total length = 204 bytes (00cc)

ID = 57862, Flags = 0x00, Frags = 0

Time To Live = 113 seconds/hops

Protocol = 17 (UDP)

Header checksum = a1a9

Source address = 132.185.132.118

Destination address = 224.2.220.62

No options

UDP

Source port = 31106 (7982)

Destination port = 31106 (7982)

Length = 184 (00b8)

Checksum = 08a0

RTP

180B of Data

```
0: 0100 5e02 dc3e 00d0 bbf7 c6c0 0800 4500
16: 00cc e206 0000 7111 a1a9 84b9 8476 e002
32: dc3e 7982 7982 00b8 08a0 8005 dbc6 d721
48: 69c0 0752 bb5f fe39 3600 8808 b120 8933
64: 6219 9118 5128 ffc8 1321 bc10 933e aa23
80: 3233 ba00 e892 a00c 1a3c 0a28 37ab 012d
96: aca5 4819 9088 0b39 64ba 43a0 b9a8 04b3
112: 88b8 4bf8 3940 d024 0a98 8b0b 1703 0a3a
128: 8820 a381 a21f 3bc0 9298 e893 90bd 042a
144: 0a88 3287 59ab e980 1211 4002 2208 98b1
160: 7039 0b26 e898 99ab b118 a1aa a702 9ac4
176: 9128 ca21 7822 2971 090a 2194 98d0 27bb
192: 0958 8092 993f b3b0 2922 337a 0f88 8810
208: 8a29 0183 fb15 b888 0d4c
```

IP Multicast

IGMP

Switched Ethernet

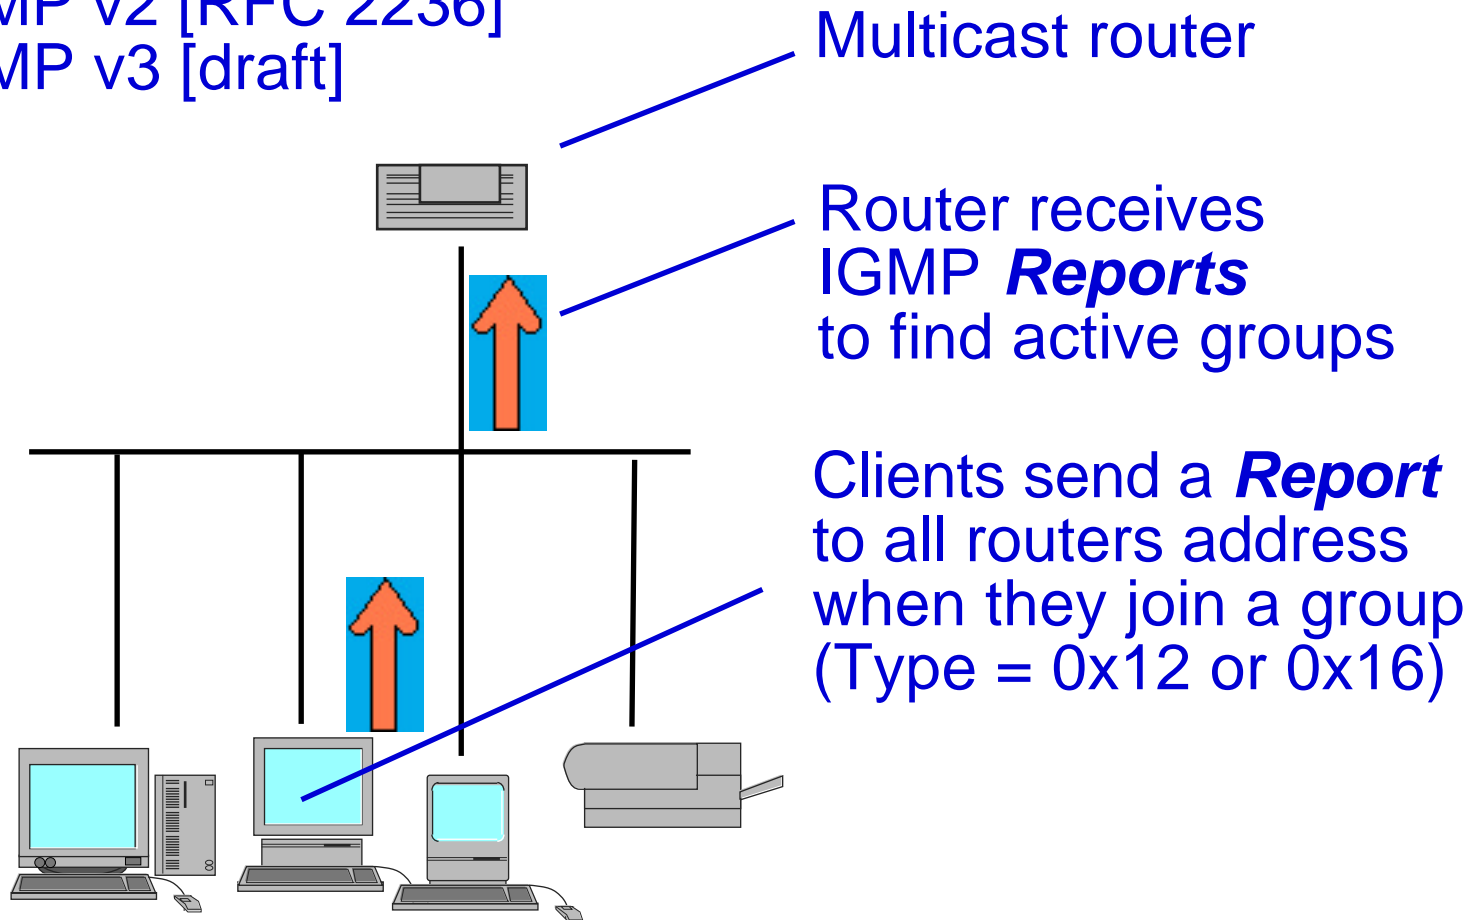
Ten Thorny Issues

Questions & Answers

Internet Group Management

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP v1 [RC 1112]
IGMP v2 [RFC 2236]
IGMP v3 [draft]



Supported on Macs, PCs, UNIX, ...

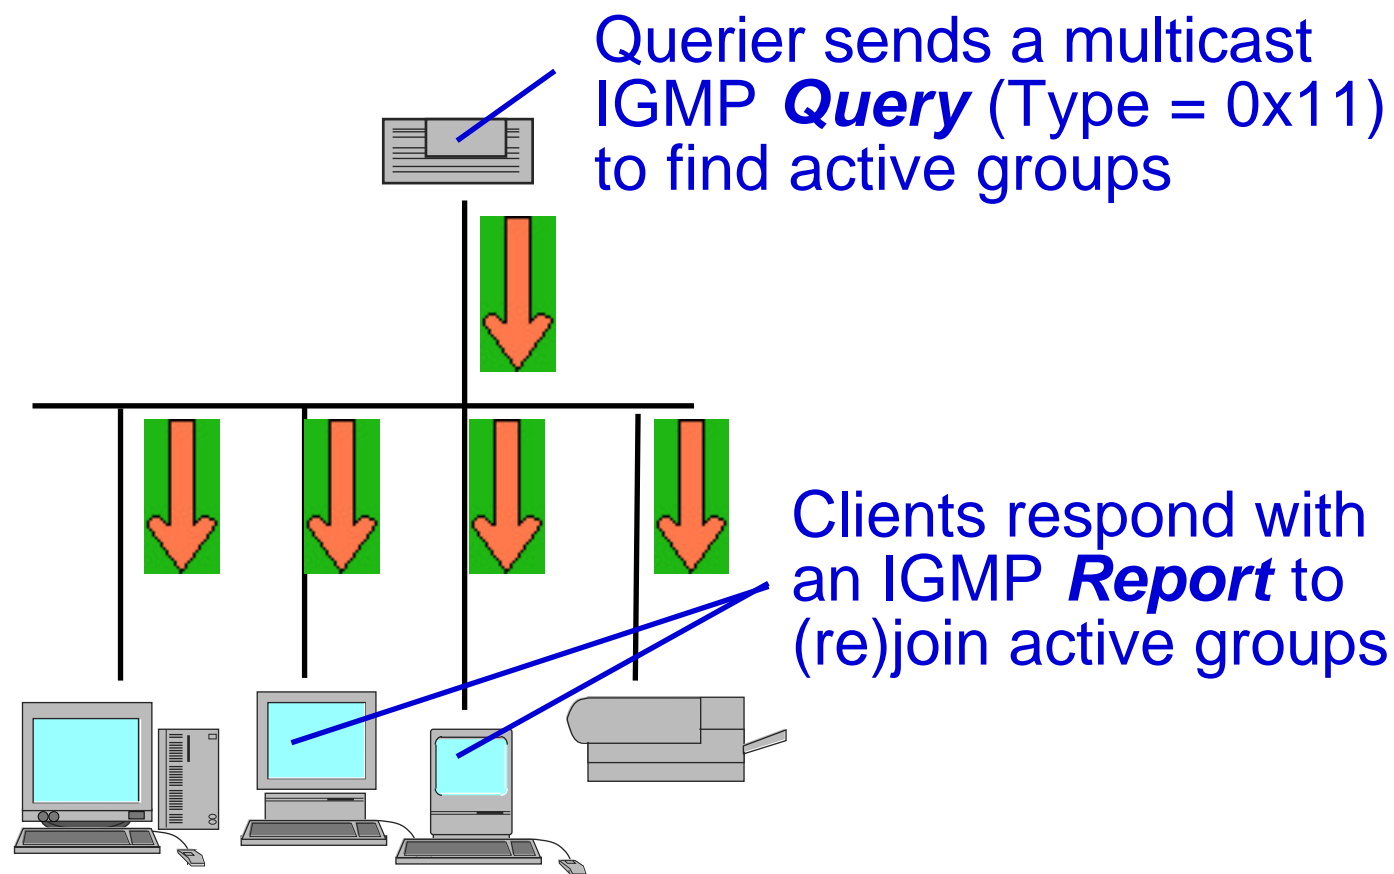
IGMP Query

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP state is “Soft”

Needs to be periodically refreshed

A switch/router acts as an IGMP querier



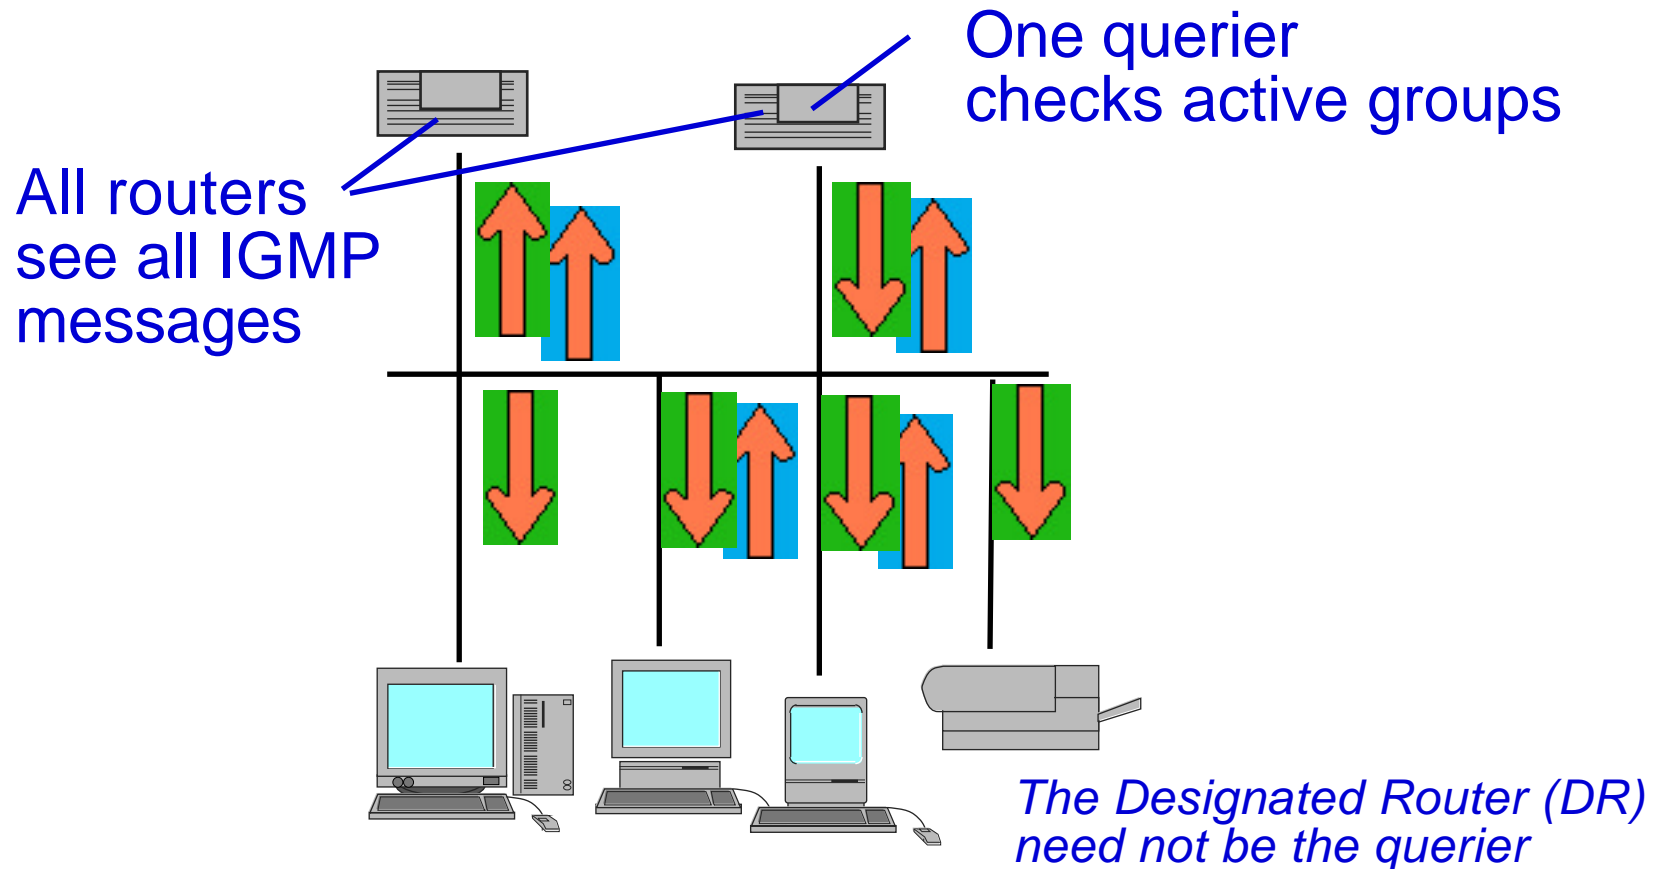
IGMP Querier Election

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Querier Election

Only one IGMP Querier per LAN

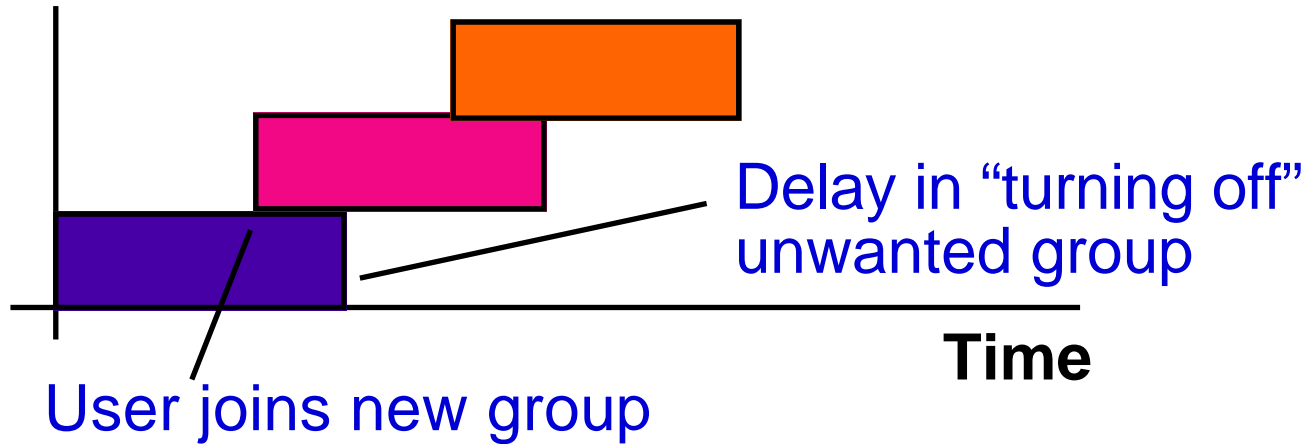
Querier with lowest numbered IP source address (v2)



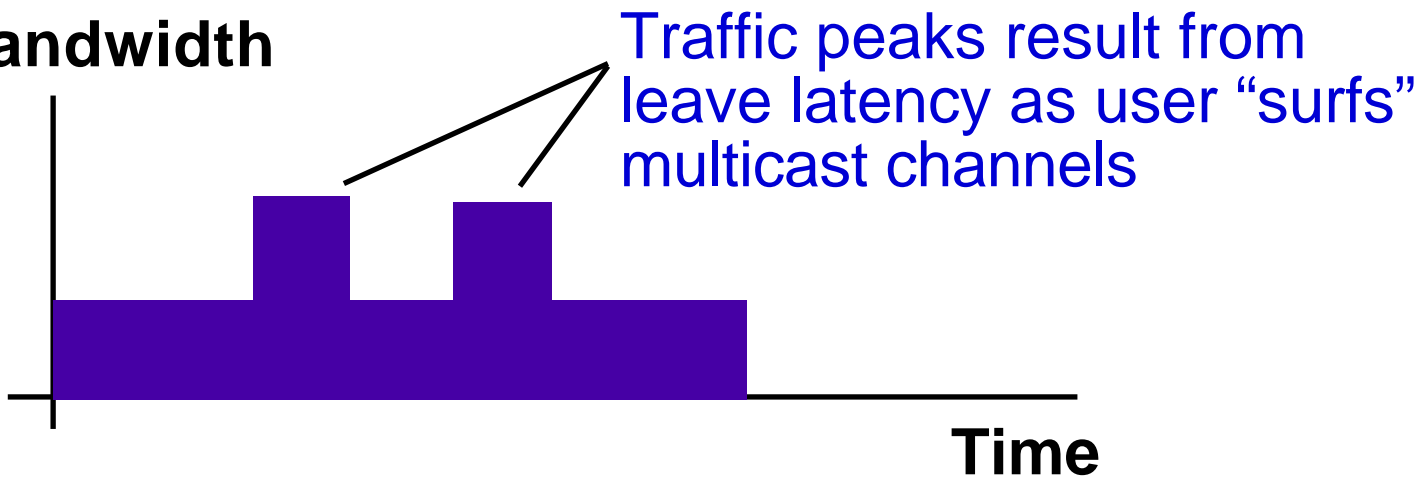
“Channel Surfing” & Leave Latency

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Groups

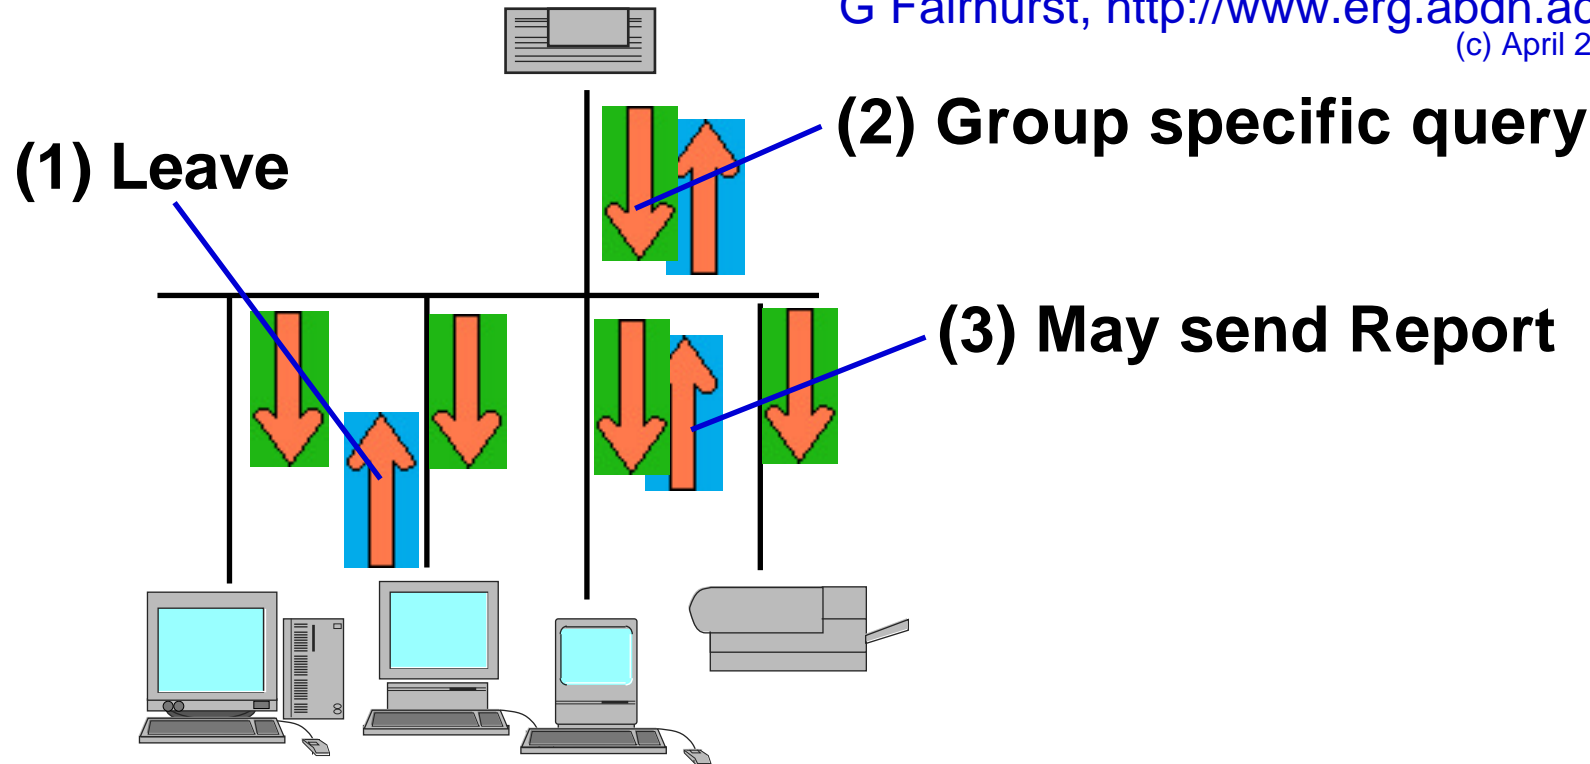


Bandwidth



IGMPv2 Leave / Join Override

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002



Join Override

Querier sends a **group specific query**

(Typically repeated every sec for 3 secs)

If still members, they send a **Report** at random

Router continues to send to the group **[RFC2236]**

Leave Group Message

(Type = 0x17, sent to 224.0.0.2)

Optional if not last member, *SHOULD* if last.

In v2, router sends a **Group Specific Query**

(typically repeated every sec for 3 secs)

Query Interval Response Time

Upper bound for random interval in 1/10 sec

Optimise for group size (*query-interval*)

Larger for bigger groups (increases leave latency)

IGMP Messages carry IP Router Alert Option

4 B IP Option [RFC 2113]

But, not all implementations use Router Alert!

(Strange behaviour when routers expect this)

IGMP MRD (Type = 0x24, 0x25, 0x26)

RFC 2236

Joining

- Join a group (as before)
- Join a group but exclude certain sources
- Join a group only for specific named sources

Queries

- All hosts respond to *queries*
- Routers know state of each member
- Allows instantaneous leave

3 Versions of IGMP

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP v1 - allowed joining multicast groups
(Win 95 - unless patched)

IGMP v2, 1997 - introduced fast leave
(widely implemented - UNIX, Win98)

IGMP v3 - Introduces source specific joins & leaves
Prevents denial of service attacks
Offers fine grain control over sources received
(*IGMP v3 is NOT not yet widely implemented*)

Each version is backwards compatible

Routers and Hosts must use oldest version present

IP Multicast

IGMP

Switched Ethernet

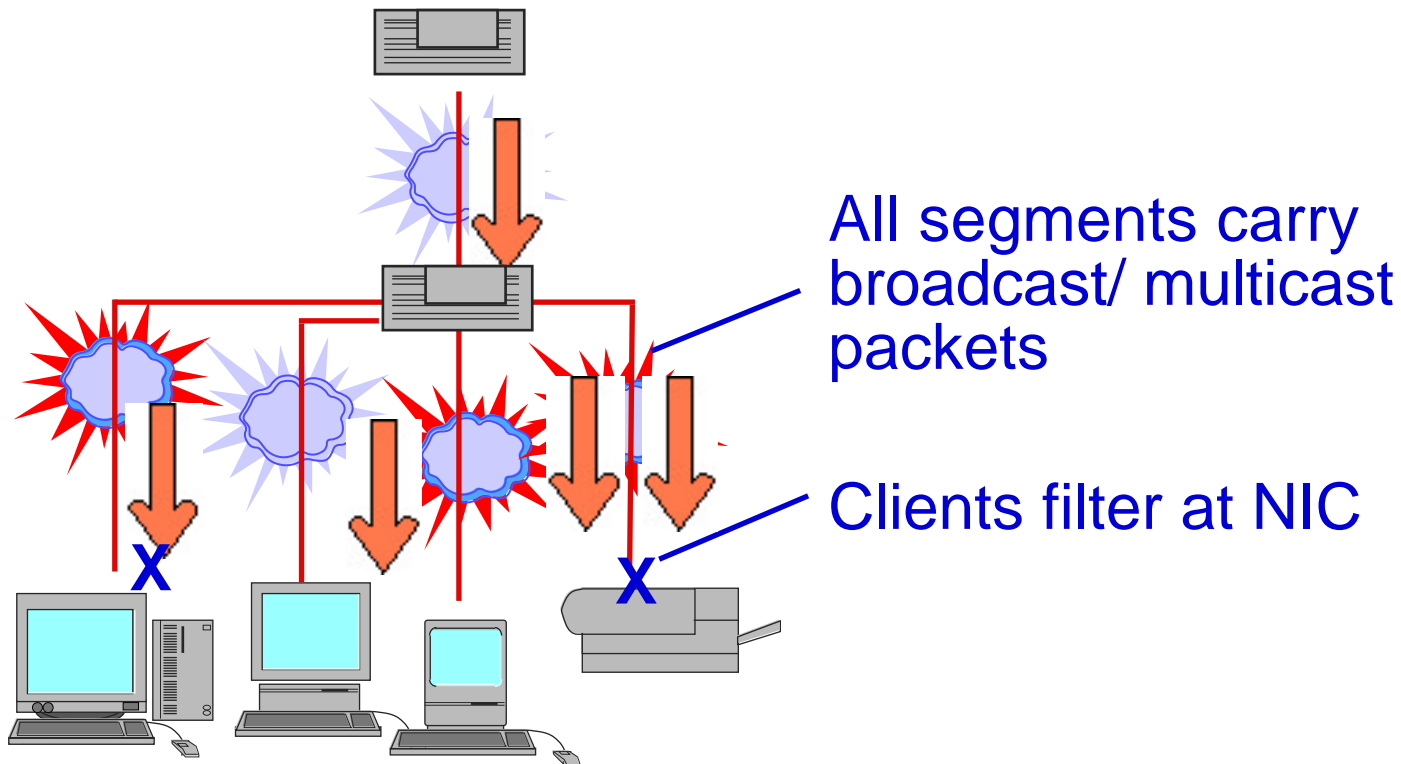
Ten Thorny Issues

Questions & Answers

Switch with Multicast

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Unmanaged layer 2 switches flood **all** multicast

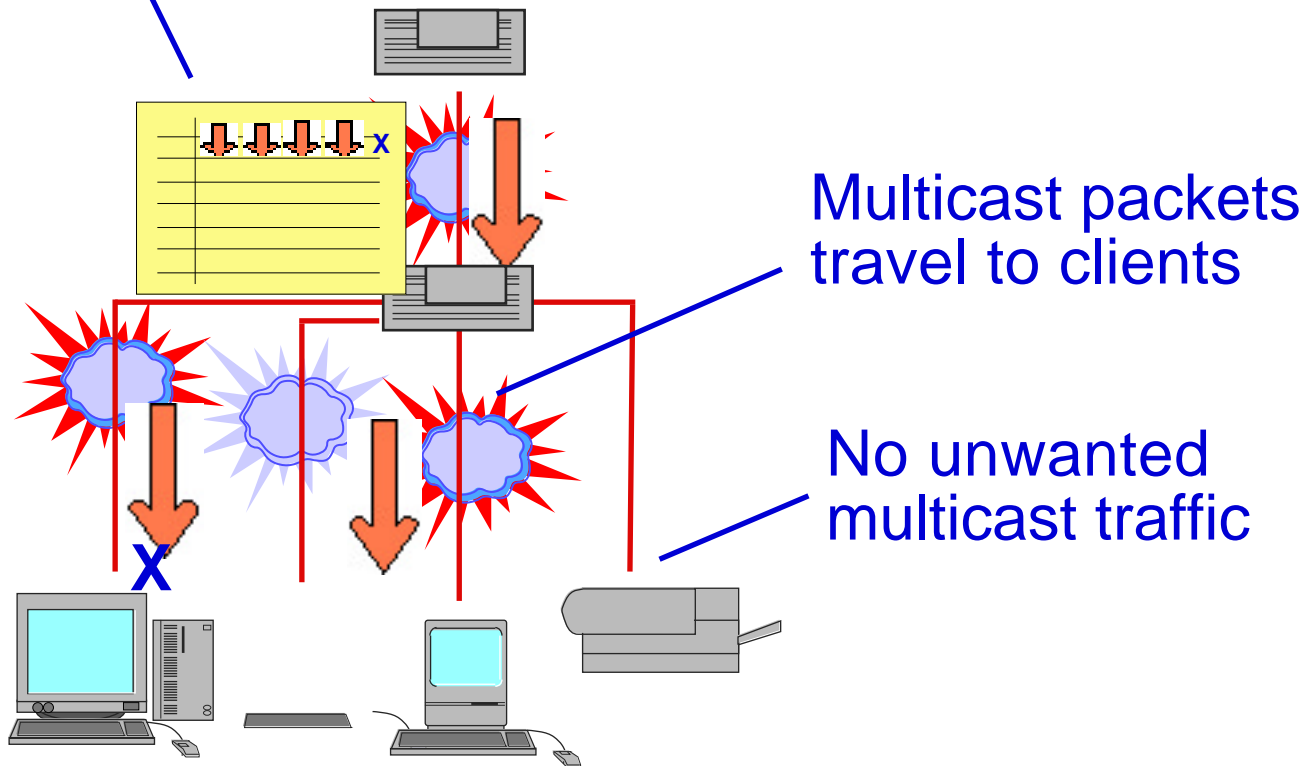


No problem, for low levels of multicast
Big problem, for high levels of multicast

Multicast Filter Table

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Filter (group, port) combinations



Requires manual configuration

IGMP in a Switched LAN

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

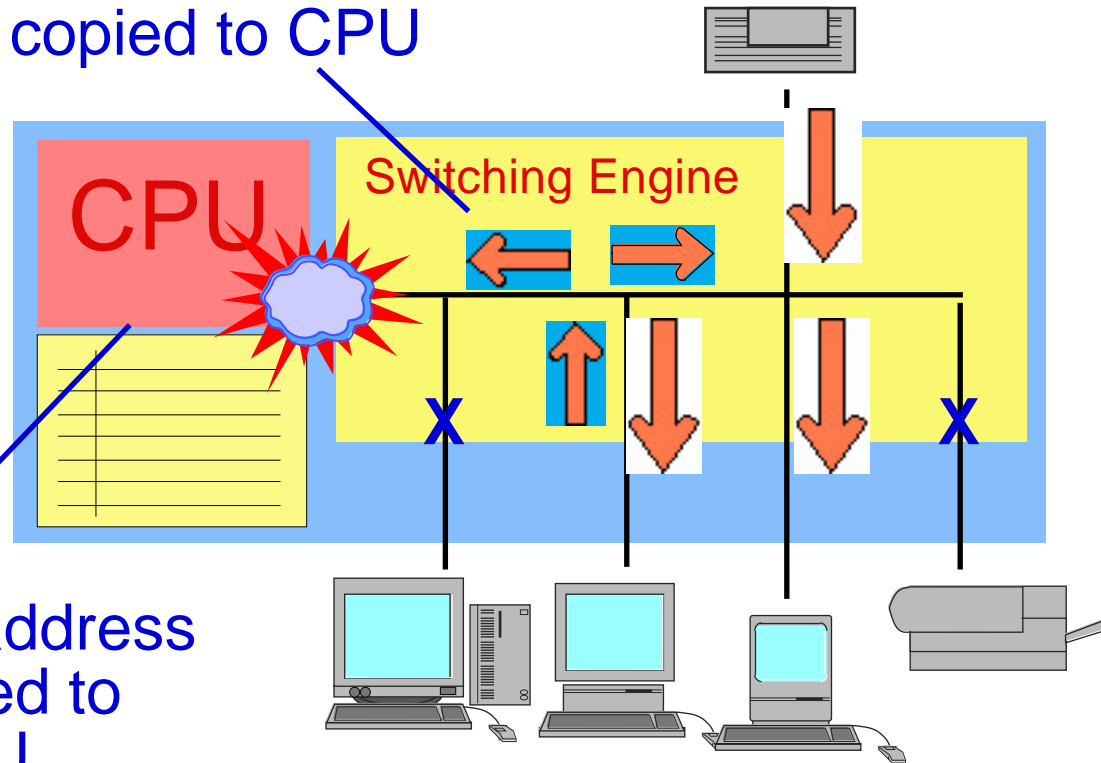
4 solutions:

- 1) IGMP Snooping
 - (i) Switches filter every L2 multicast packet
(heavy processing)
 - (ii) Switch ASICs filter IGMP packets
(expensive ASIC)
- 2) IGMP Proxy
(expensive ASIC, more processing)
- 3) Router signals membership to switch
(simple, e.g., CGMP - CISCO proprietary)
- 4) GMRP/GARP
(simple, routers and hosts use 802.1p)

Multicast in a Switched LAN

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP messages copied to CPU




Multicast (MAC) address
and port lists added to
filter table by CPU


Some groups still need to be flooded:
224.0.0.x; packets from local sources; non-IP frames

Multicast Bridge Filter Table

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002



MAC group address



Port list

Switches normally use MAC group address

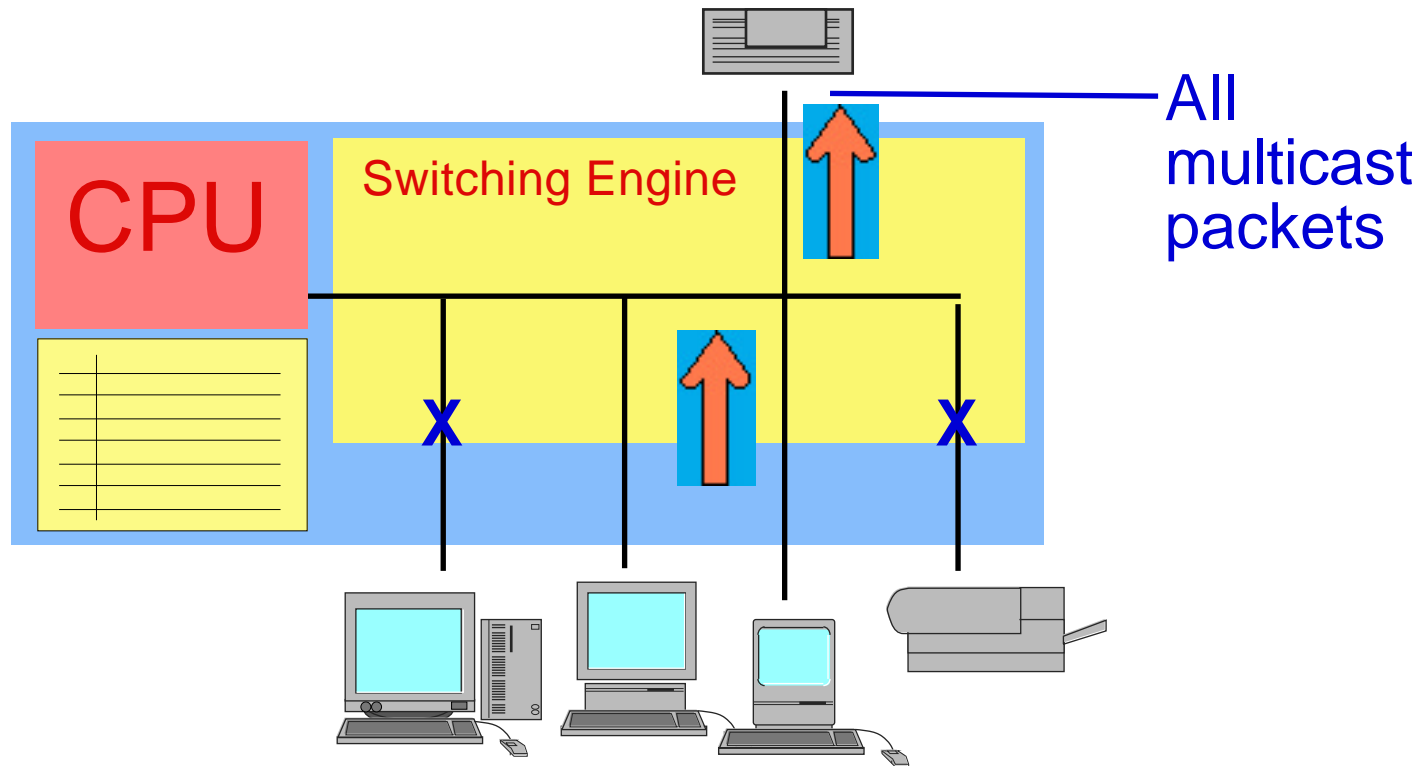
(e.g., 224.0.01 = 01:00:5E:00:00:01
also maps to 32 other groups !!!)

Level 3 switches may forward based on IP address

“Router” Ports

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

All multicast packets sent towards upstream routers

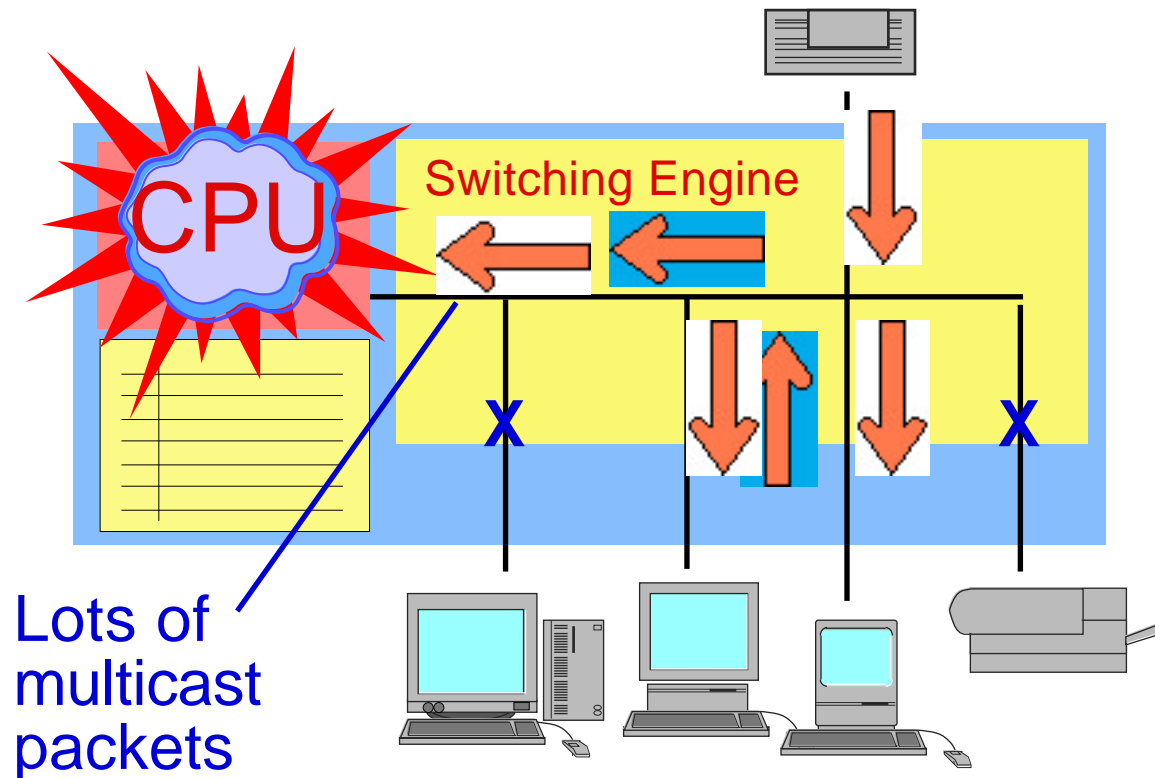


Routers discovered using
PIM HELLO (224.0.0.8); DVMP (224.0.0.13); IGMP MRD

Multicast in a Switched LAN

G Fairhurst, <http://www.erg.abdn.ac.uk>

IGMP v1,v2 messages sent to same group as data

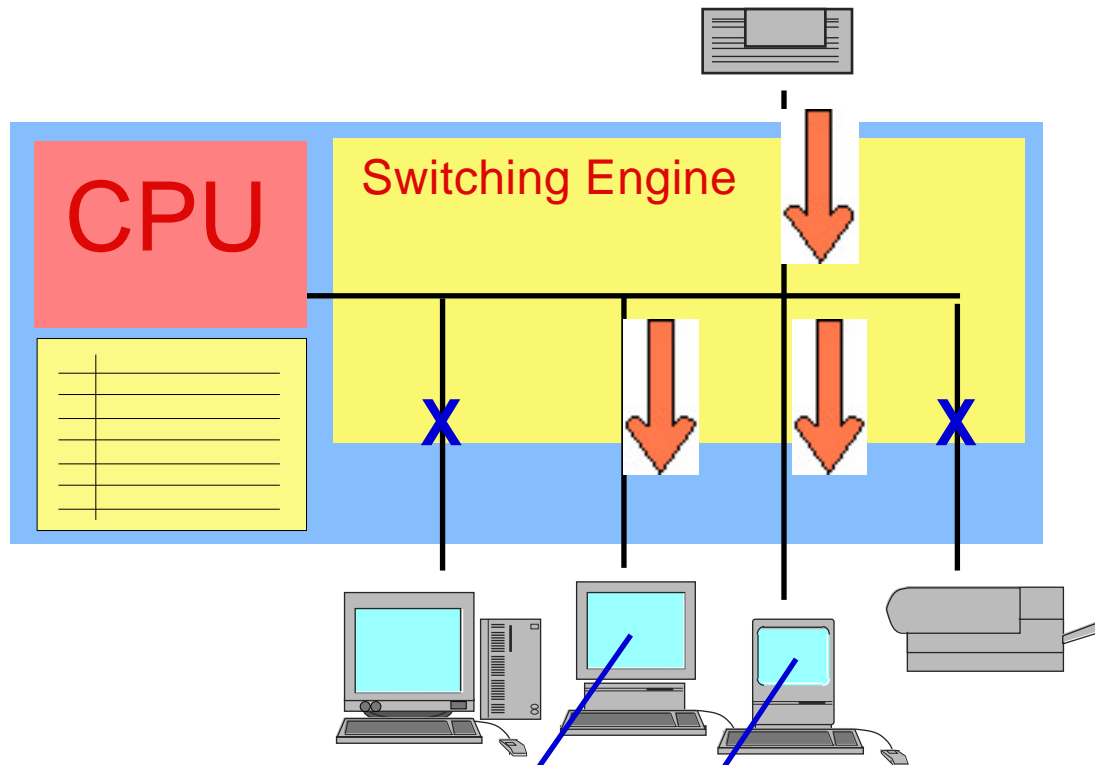


Desirable to use a configurable filter before CPU (ASIC)

Undetected Leaver

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP v1,v2 messages don't identify the end hosts




This end host leaves G

This end host continues to receive G

IGMPv2 Undetected Leaver


G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Groups



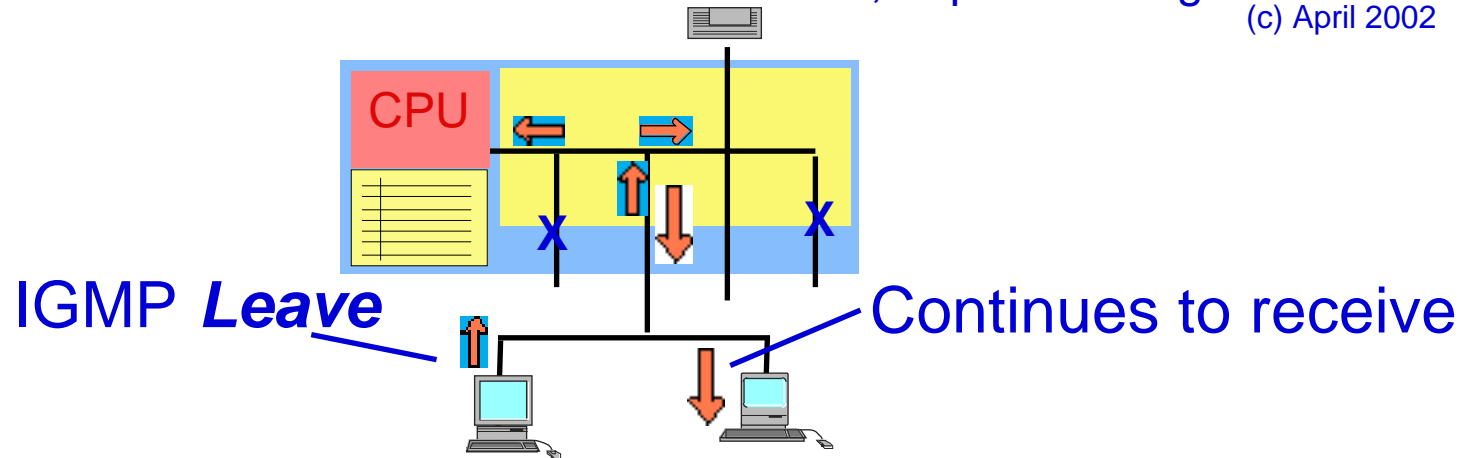
Joins group already sent to LAN

Bandwidth



Leave Processing

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002



Ports with more than one IP address

e.g. a hub or unmanaged downstream switch

Careful processing of received IGMP Leave

Defer sending the *leave* upstream

Need to send a *query* back to the port

May even be safer to send a *general (v1) query*

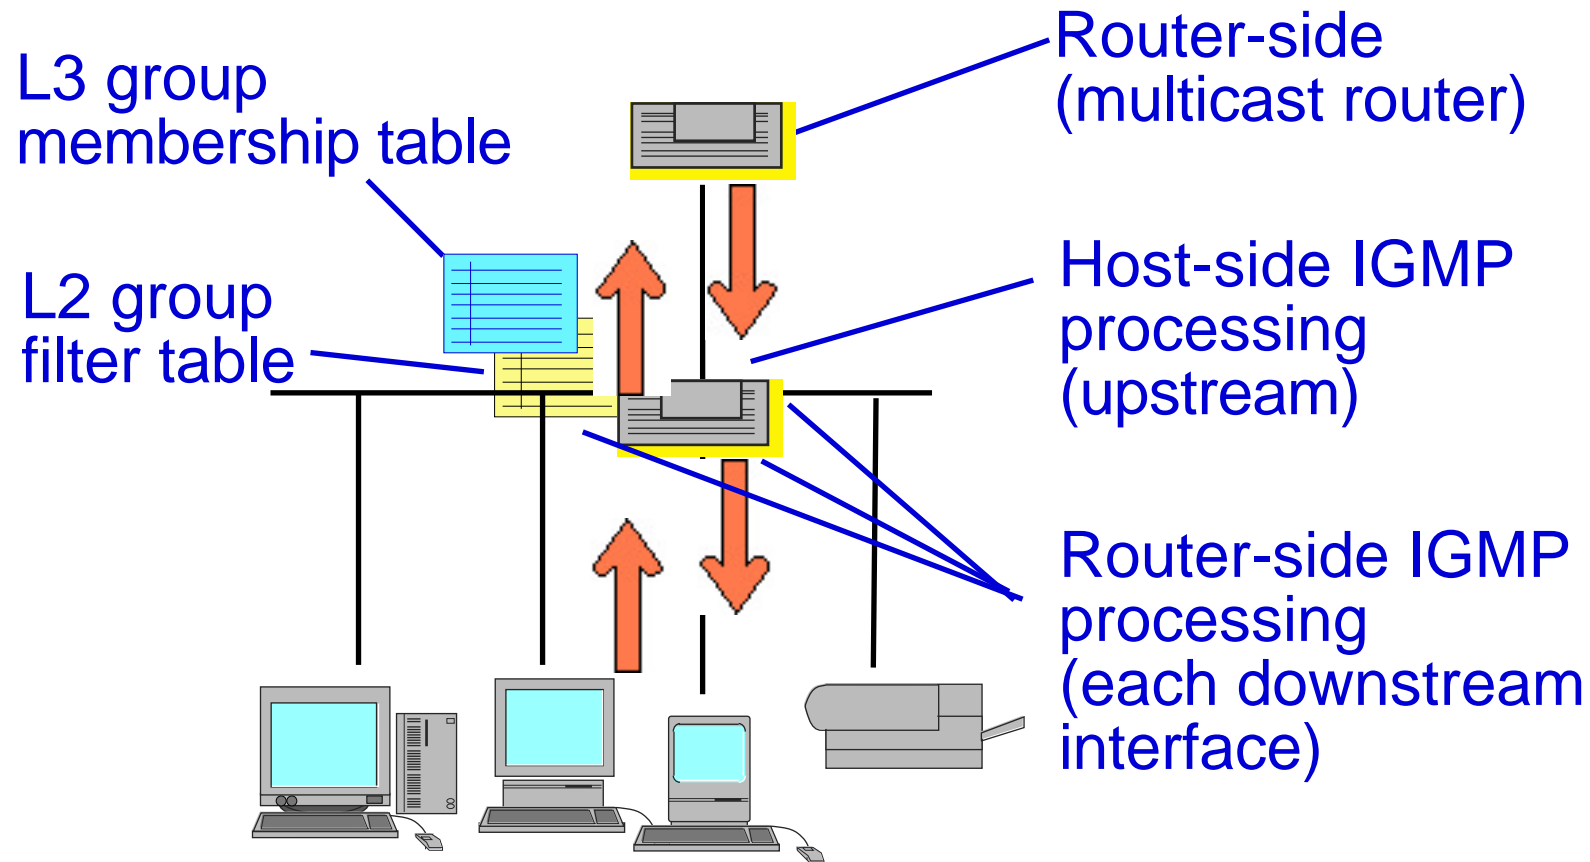
If no response, *and* no other members

send *leave* to all routers

IGMP Proxy

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

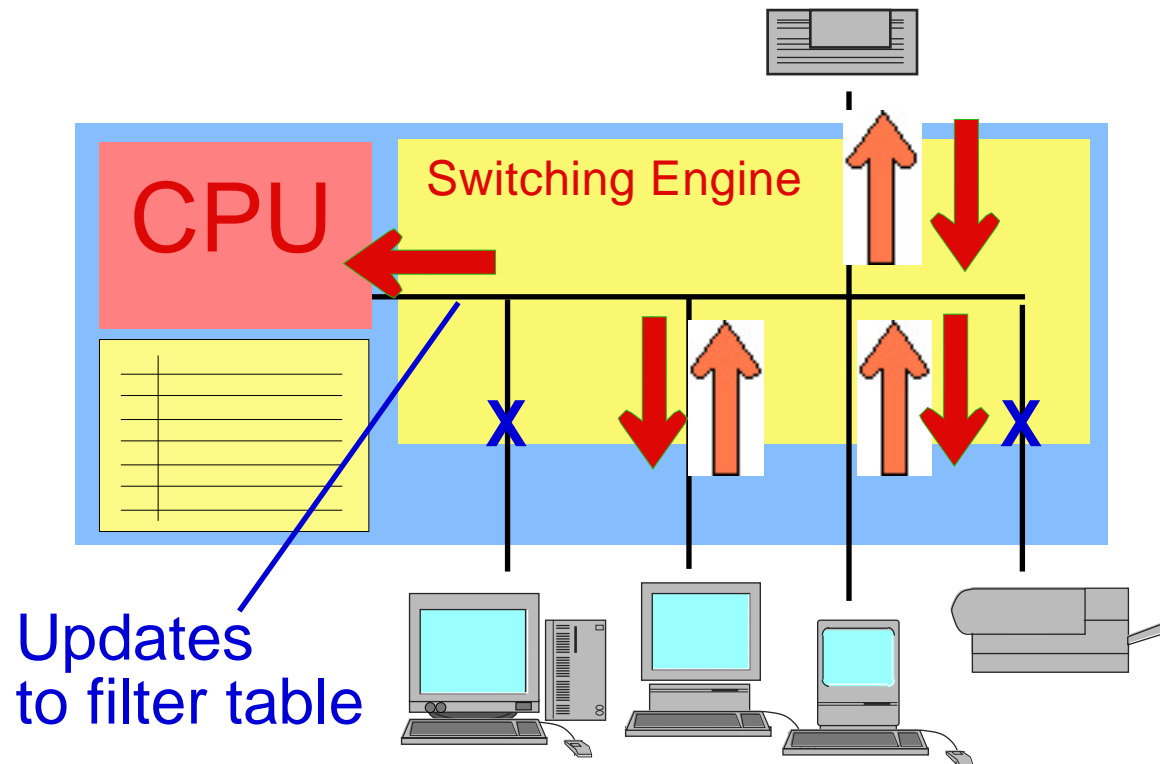
IP multicast forwarding controlled by IGMP at **switch**



CGMP in a Switched LAN

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP messages sent to IGMP querier (router)



Updates
to filter table

Querier sends summary membership to switches
CGMP is multicast (to all capable switches)

CGMP Commands

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Router Port

Every 60 seconds
[GDA = 0; USA = router]

Join a Port

After receiving an IGMP report from user
[GDA = multicast group; USA = switch port]

Delete group from all switches

After receiving last IGMP leave message
[GDA = multicast group; USA = switch port]

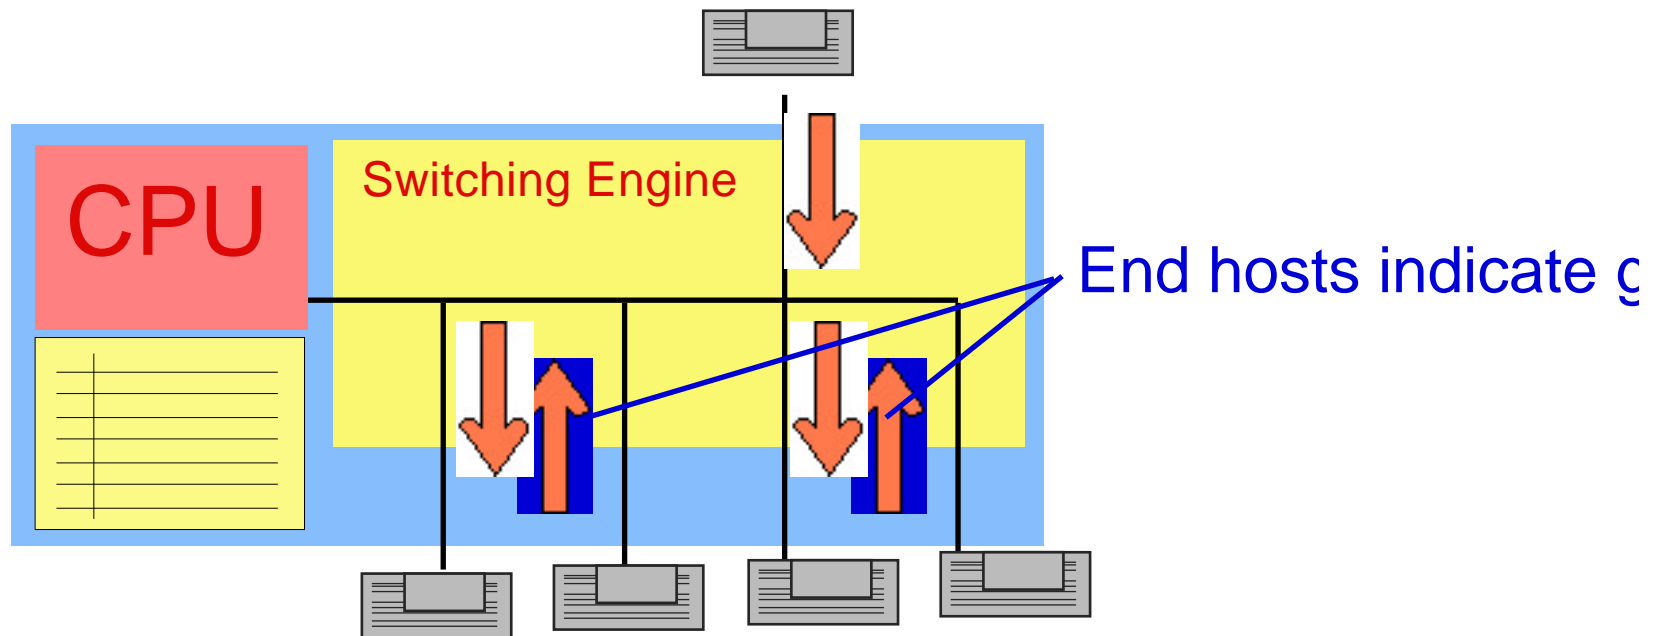
Delete all groups from all switches

After router shuts down CGMP
(Or clear ip cgmp)
[GDA = 0; USA = router/0]

GMRP

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

Uses 802.1p VLAN Tagging (1522 B frames)
Generic Multicast Registration Protocol GARP messages



IP Multicast

IGMP


Switched Ethernet

Ten Thorny Issues

Questions & Answers

Ten Thorny Issues

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002



- (i) PERFORMANCE - software or hardware feature?
- (ii) Forwarding of non-IP multicast traffic
- (iii) SHOULD (MUST?) forward 224.0.0.X
- (iv) MUST support IP Router Alert option
- (v) MUST support Multicast Router Discovery (MRD)
- (vi) Senders on the LAN - flood or forward to routers?
- (vii) How robust? – ST, restart, end host movement?
- (viii) Source address for IGMP Reports (snoop/proxy)
- (ix) Future proof? IGMPv3, MLD & MLDv2
- (x) Debugging & Management Information!!!

Conclusions

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IGMP

v1 - Dead

v2 - Widely supported

v3 - Soon

MLD (MLDv2) used for IPv6

Multicast-enabled LAN Switches

Snooping; Proxy ; CGMP

Widely supported - not all are equal !!!

It works!!!!

Some pitfalls - watch out for them !!!

Ask about v6 ;-)

G Fairhurst, <http://www.erg.abdn.ac.uk>
(c) April 2002

IP Multicast

Ethernet

IGMP

Switched Ethernet

Ten Thorny Issues

Questions & Answers