


EPFL

The MAC Layer




Jean-Yves Le Boudec 2019

Contents

1. MAC as Shared Medium
2. MAC on cabled systems: the world of switches
3. Format and addresses
4. Virtual LANs
5. Architecture Aspects
6. Security aspects

Textbook

Sections 2.6, 3.16 and 4.12




1: Shared Medium Access

MAC = medium access control

Why did engineers invent the MAC layer ?

Share a cable (Ethernet, Token Ring, Powerline)

Use a wireless radio link (GSM, WiFi, Bluetooth)



What is the problem ?

If several systems talk together, how can we decode ?

Solution 1: joint decoding (decode everyone) used in CDMA cellular networks – complex

Solution 2: **mutual exclusion** protocol (only one system can talk at a time) + **distributed** (no central component)

How does MAC layer work?

there are *many* solutions, called: Token Ring, Aloha, CSMA, CMSA/CA, CSMA/CD, RMAC, etc

We describe here the most basic one, called CSMA/CD, used on some Ethernet cables and on CAN buses.

Powerline (“homeplug MAC”) is similar (CSMA/Collision Avoidance) for transmission over home power lines.

Radio links (WiFi) is also similar, with a few complications – see companion textbook for a more detailed explanation.

CSMA/CD derives from Aloha

Aloha is the basis of all non-deterministic access methods. The Aloha protocol was originally developed for communications between islands (University of Hawai) that use radio channels at low bit rates.

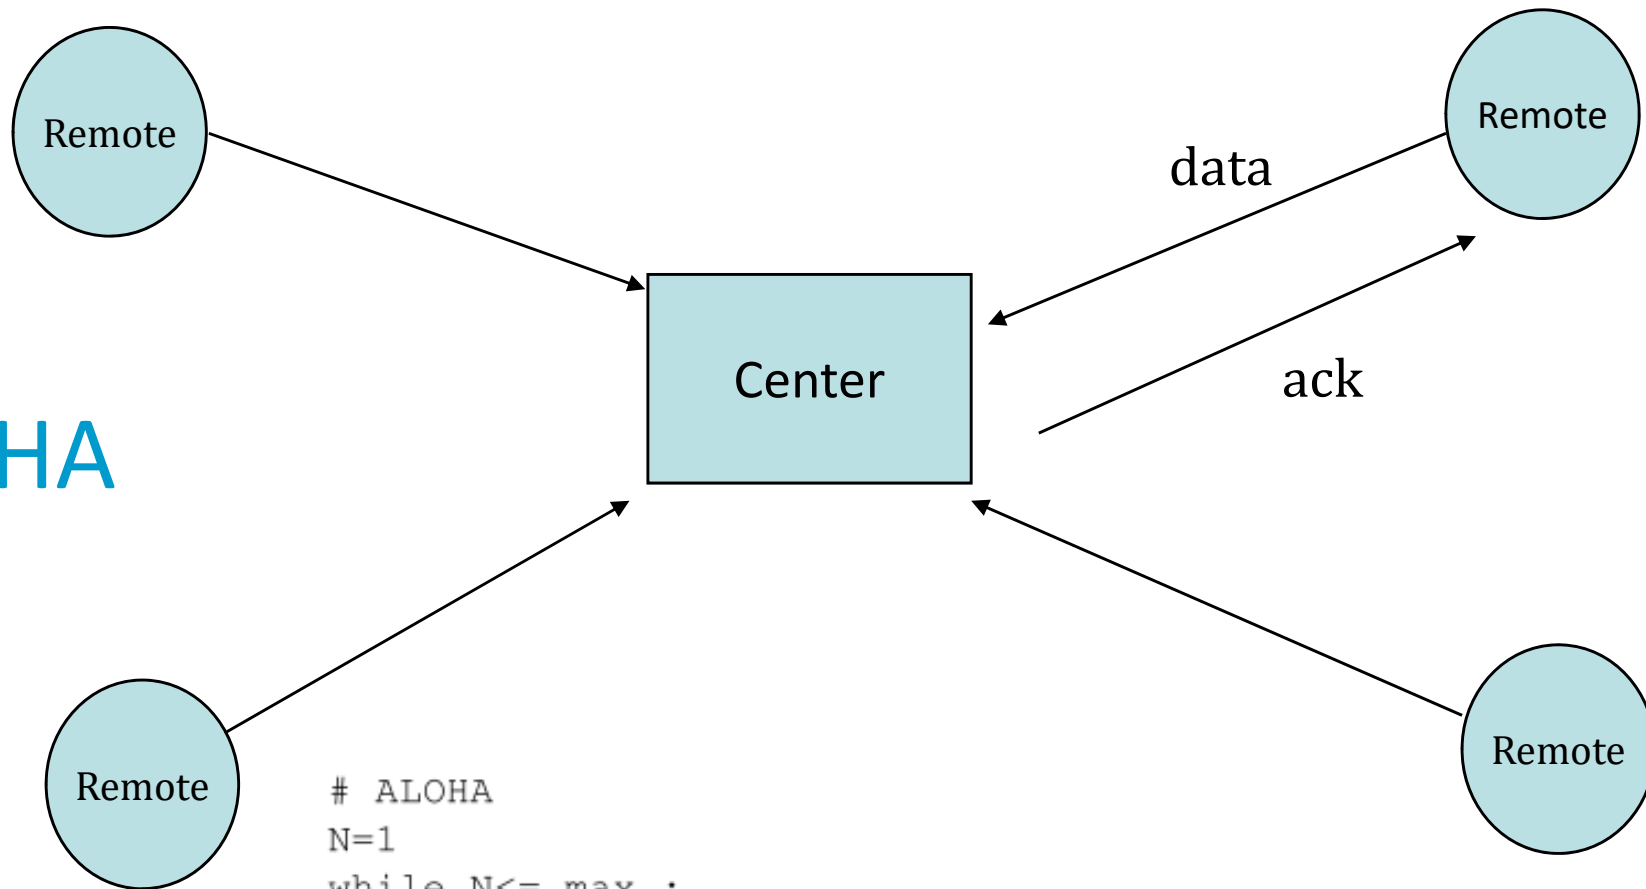
It assumes one shared channel from remote to center, and a separate (non interfering) channel from center to remotes

Collisions occur when two packet transmissions overlap, and if a packet is lost, then source has to retransmit; the retransmission strategy is not specified here; many possibilities exist. We will see the one used for CSMA/CD.

There is no feedback to the source in case of collision (was too complex to implement at that time). The picture shows a radio transmission scenario; Aloha can also be used on a cable (bus). It is used nowadays in cases where simplicity is more important than performance (for example: machine bus)

The maximum utilization can be proven to be 18%, which is small. This is assuming an ideal retransmission policy that avoids unnecessary repetitions of collisions.

ALOHA



```
# ALOHA
N=1
while N<= max :
    send(frame)
    wait(ack_on_return_channel or timeout)
    if (ack_on_return_channel):
        break # transmission was successful
    else:
        # timeout
        wait(random_time)
        N=N+1
else:
    # Too many transmission attempts
```


Transmission procedure at remote

CSMA (carrier sense multiple access) improves on Aloha by requiring that stations listen before transmitting



```
# Non persistent CSMA
N=1
while N<= max :
    listen(channel)
    if free(channel):
        send(frame)
        wait(ack or timeout)
        if received(ack) :
            break # transmission was successful
        else :
            # timeout
            N=N+1
    else:
        wait(random_time)
# end of while loop
# Too many transmission attempts
```

assumes a single *transitive* channel (everyone can hear everyone)



CSMA collisions...

- A. Cannot occur because the medium is transitive
- B. Can occur if the bandwidth-delay product is smaller than 1
- C. Can occur regardless of the bandwidth delay product
- D. I don't know


CSMA / CD (Collision Detection) detects collisions as they occur

```
# CSMA/CD pseudo-code
N=1
while N<= max :
    wait(channel_becomes_free)
    send(frame)
    wait_until (end_of_frame) or (collision)
    if collision detected:
        stop transmitting
        send(jamming)
        k = min (10, N)
        r = random(0, 2k - 1) * slotTime
        wait(r*slotTime)
        N=N+1
    else :
        wait(inter-frame_delay)
        break
# end of while loop
# Too many transmission attempts
```

Ack replaced by CD

CSMA / CD Time Diagram 1


A senses idle channel,
starts transmitting
shortly before T , B
senses idle channel,
starts transmitting



CSMA / CD Time Diagram 2

A senses collision,
continues to transmit
32 bits ("jam")

B senses collision,
continues to transmit
32 bits ("jam")



CSMA / CD Time Diagram 3


A waits random time
 t_1

B waits random time
 t_2

B senses channel idle
and transmits

A senses channel busy
and *defers* to B

A now waits until
channel is idle



CSMA/CD improves on CSMA by requiring that stations detect collisions and stop transmitting (after 32 bits, called *jam* bits, in order to ensure that all circuits properly recognize the presence of collisions).

CSMA/CD has a better performance than Aloha or CSMA

After a collision is detected, stations will re-attempt to transmit after a random time.

Acknowledgements are not necessary because absence of collision means that the frame could be transmitted (see “Minimum Frame Size”).


The interframe delay (“gap”) is 9.6 μ s. It is used to avoid blind times, during which adapters are filtering typical noise at transmission ends.

The random time before retransmission (exponential backoff) is chosen in such a way that if repeated collisions occur, then this time increases exponentially. The effect is that in case of congestion (too many collisions) the access to the channel is slowed down.

A Minimum Frame Size is Necessary to Guarantee Collision Detection


Frame sent by A is too short: collision is *not* visible at A but is visible at C

Frame sent by A is large enough collision *is* visible at A



CSMA/CD imposes a minimum frame size = $64 B$ and a maximum network diameter : 2 km for 10Mb/s, 200m for 100 Mb/s


CSMA/CD = Ethernet on Coax Cables



This is (old style) Ethernet. It uses CSMA/CD.
Also called “**half-duplex** Ethernet”


WiFi = CSMA/CA (Collision Avoidance) = variant of CSMA/CD where Collision Detection is replaced by Acknowledgements

- Does not assume collision can all be detected because some systems may not be within reach; absence of ack interpreted as collision
- CTS heard by C prevents many collisions
- Other optimizations to avoid collisions
- See textbook / mobile network course




In a WiFi network with a single channel and a single base station, how many frames can be transmitted concurrently ?

- A. At most one in total
- B. One from mobile to base station and one from base station to mobile
- C. At most one per mobile to base station and one from base station
- D. At most one per mobile to base station and one per mobile from base station
- E. I don't know



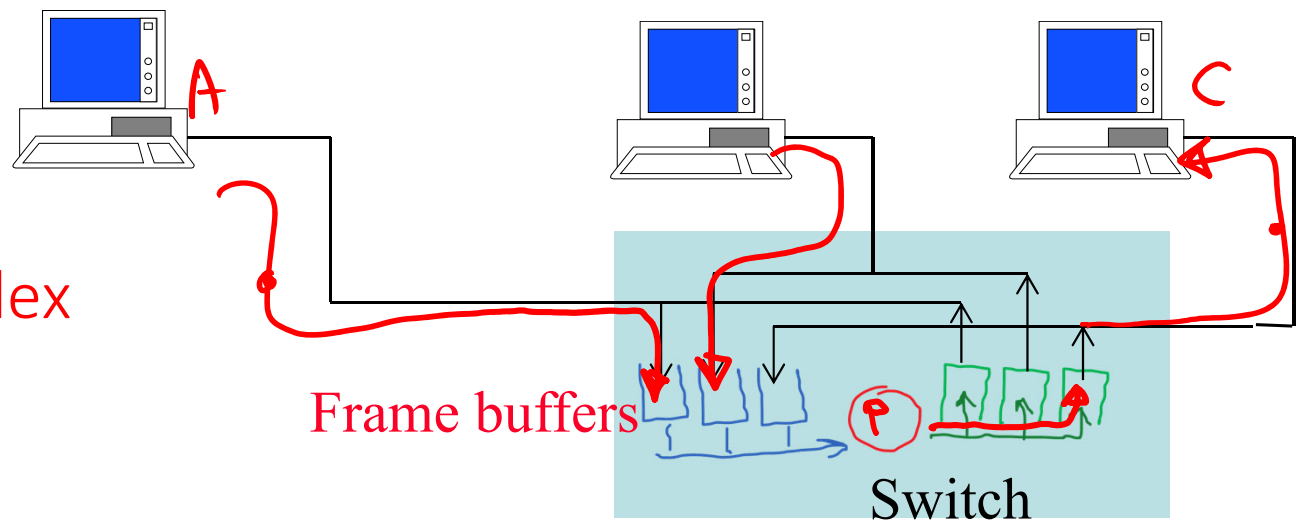
2. Ethernet LANs Use Switches

Ethernet was historically shared medium
But today is based on **switches = bridges**
a bridge is an intermediate system of the
MAC layer = it forwards packet based on
MAC addresses.





A bridge is a queuing
system

Here the links are **full-duplex**
Ethernet: there is no
MAC protocol (no
CSMA/CD) and packets can flow
simultaneously in both directions.




Bridges are Transparent

The design of bridges requires them to be **transparent**, i.e. there should be no difference for the protocol used by end-systems when they are connected to a concentrator (cabling only) or a bridge.



Bridges use a forwarding table with *exact match*



Forwarding table contains list of destination MAC addresses and interfaces (also called “ports”).

Bridge algorithm


- listen to all traffic on all interfaces

- read destination MAC address and obtain destination interface from table if it exists

- if destination interface is same as origin, discard frame; else send to destination interface.

- If destination address not in table, broadcast to all interfaces

Bridges learn addresses by observing traffic



How can a bridge build its forwarding table ?


table built by **learning** from Source Address field in MAC frame

learnt addresses time out if not re-learnt

If destination address not in table, frame is broadcast to all interfaces

Multicast addresses are handled separately – see later.

The method of learning can be extended to networks of bridges



Forwarding Table at B3

Dest MAC Addr	I/f Nb
---------------------	-----------

A	1
B	1
C	3
D	3
E	2

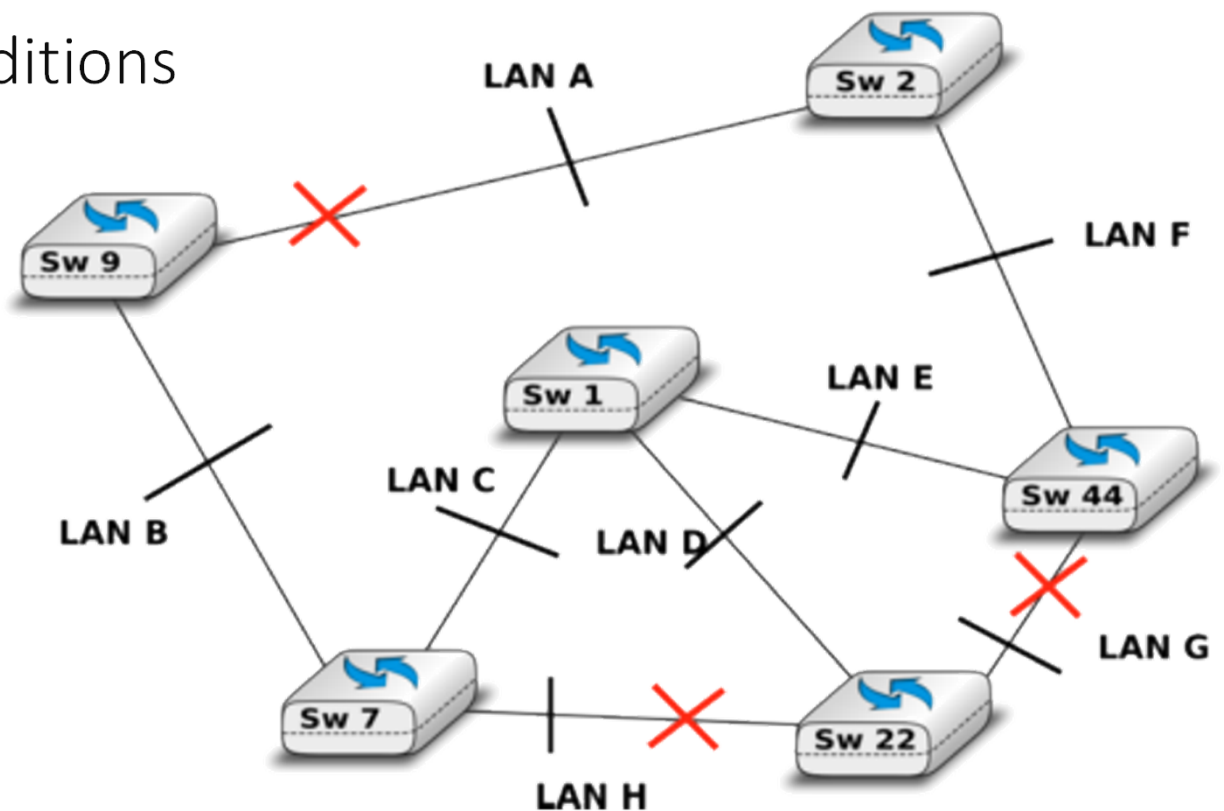
Does this always work ?

- A. Yes because learning propagates
- B. Yes because MAC addresses are unstructured
- C. Only in some special cases
- D. No, never
- E. I don't know

The Spanning Tree Protocol Forces the Active Topology to be Loop-free, i.e. to be a Tree

Bridges run a protocol between neighboring bridges, called the **Spanning Tree Protocol**

What: de-activate some ports such that the remaining topology is a tree that spans all bridges
+ adapt to failures and additions



The Spanning Tree Protocol: How

Algorhyme



I think that I shall never see
a graph more lovely than a tree.
A tree whose crucial property
is loop-free connectivity.
A tree that must be sure to span
so packet can reach every LAN.
First, the root must be selected.

By ID, it is elected.

Least-cost paths from root are traced.

In the tree, these paths are placed.

A mesh is made by folks like me,
then bridges find a spanning tree.

→ using the Bellman-Ford
algorithm, see
“Distance Vector”

Radia Perlman

(inventor of the Spanning Tree Protocol)

https://www.youtube.com/watch?v=iE_AbM8Zyki

Specification of the Spanning Tree Protocol

1. Bridges select one **root** bridge
= bridge with smaller bridge label (configurable)
2. Only links that are on the shortest path to root are active
Spanning Tree = set of **shortest paths** to root

Each LAN between bridges has a **cost**, by default, decreasing function of bit rate.

Distributed (no central control, all of this is done by bridges themselves) using a variant of the Bellman-Ford algorithm (see “distance vector”)

Specification of STP (cont'd)

STP gives a *role* to ports of bridges.

Root ports: One per bridge = port towards root along shortest path; in case of equal costs, lowest port id chosen

Designated ports: Ports that are on the shortest path for other bridges (= downstream)


On every LAN (collision domain), choose one *designated bridge* = bridge that has smallest distance to root

all ports on LAN for which the bridge is designated are designated ports

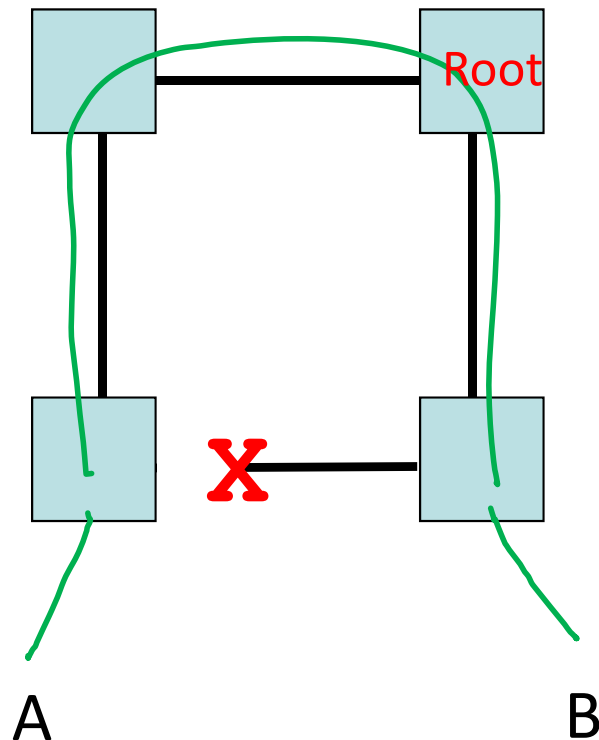
Ports not root nor designated are **blocking**

B1 is the root. What is the role of each port of B3 (b=blocking, r=root, d=designated)?

- A. 1r, 2d, 3b
- B. 1r, 2d, 3d
- C. 1r, 2b, 3b
- D. 1r, 2b, 3d
- E. 1b, 2r, 3b
- F. 1b, 2r, 3b
- G. 1d, 2r, 3b
- H. 1d, 2b, 3d
- I. None of the above
- J. I don't know



What does STP produce on this network of bridges ?



The loop is broken at one port.

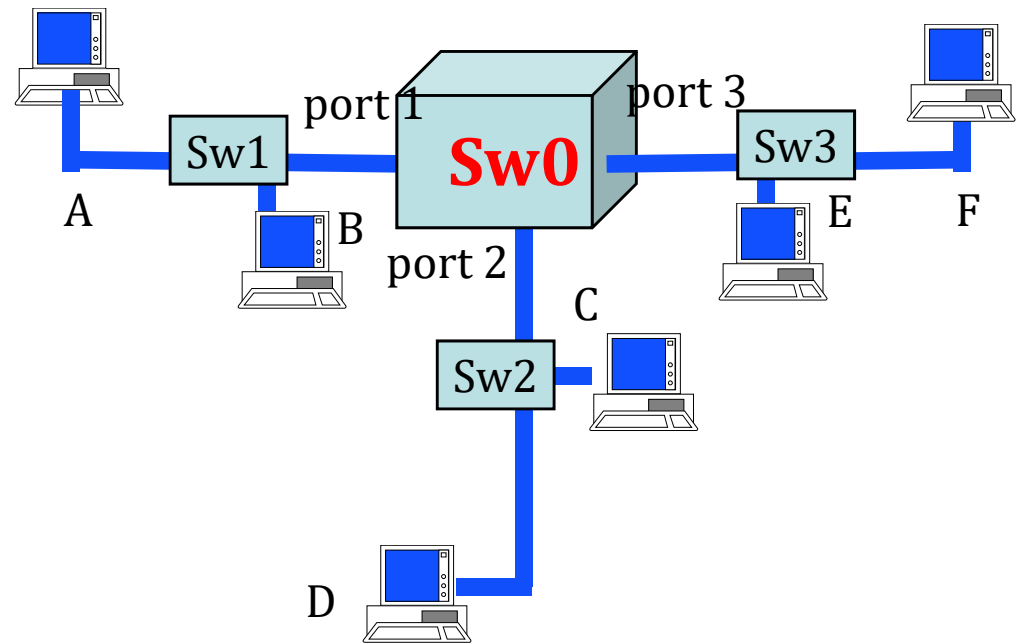
Path from A to B is not optimal:

All frames go through the spanning tree

Less efficient than shortest path
- some more sophisticated bridges implement Shortest Path Bridging (see link state) instead of STP, which is the default

How many frames can be transmitted in parallel (all boxes are switches, i.e. bridges) ?

- A. 1
- B. 3
- C. 6
- D. > 6
- E. I don't know



3. Ethernet Frame format

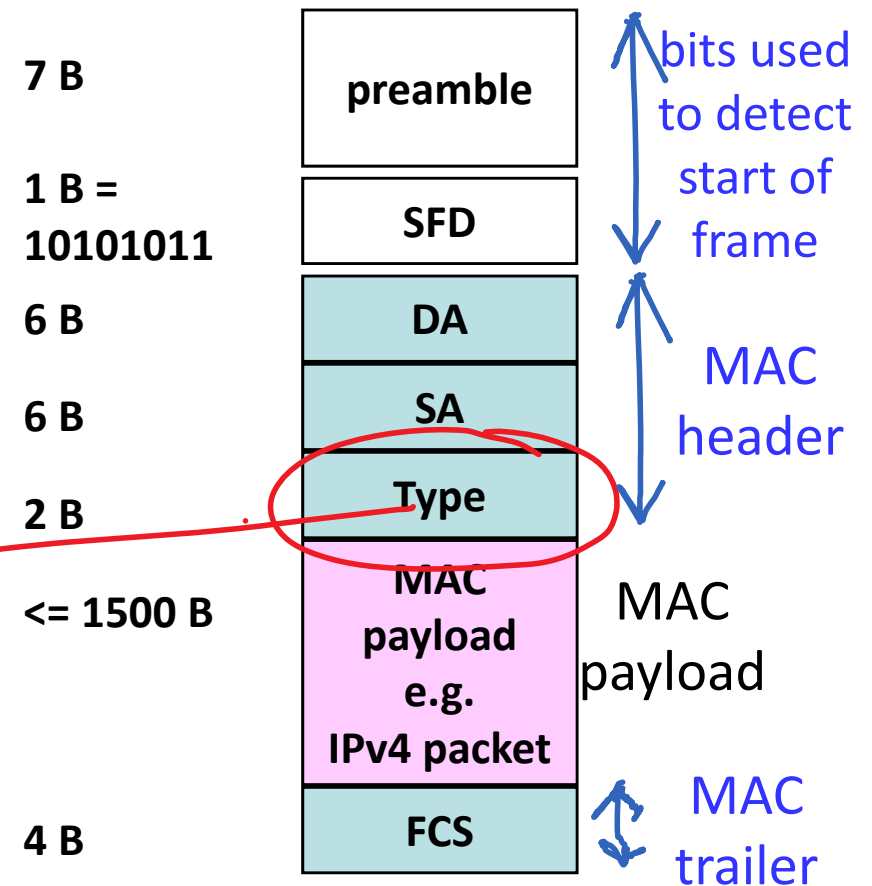
Ethernet frame = Ethernet PDU

An Ethernet frame typically transports an IP packet, sometimes also other

Type of protocol contained in the Ethernet packet (hexa):

- 0800: IPv4
- 0806: ARP (used by IPv4)
- 86DD: IPv6
- 8847: MPLS unicast
- 88F7: Precision Time Protocol

Ethernet V.2 frame



DA = destination address
SA = source address

The preamble is used for the receivers to synchronize (01010101... terminated by 0). With Ethernet, transmission starts asynchronously (stations start independently), and between transmissions, the channel is idle. SFD (start frame delimiter) is used to validate the beginning of a frame.

Destination length is used to indicate the total length before padding. Padding is required if the minimum frame size of 512 bits = 64 bytes is not reached.

There is no frame length field in the header. It is up to the layer using Ethernet to find out the length of the payload. Frames have to be at least 64B and at most 1500B. If needed, padding is inserted to make the frame long enough.

The format shown here is called Ethernet v2. There exists another format on Ethernet called IEEE 802.1. With WiFi the format is similar, but the MAC header contains additional fields.

Error Detection

Frame Check Sequence (FCS) is an error detection mechanism

Why ?

Bit errors due to transmission (Bit error rate $< 10^{-10}$ on cables) can corrupt frame. We want to detect and discard such frames.

How ? A Cyclic Redundancy Checksum (CRC, 32 bits) is computed for every frame and inserted in the FCS field. It is a polynomial code. The receiver recomputes the FCS and checks if equal.

This *detects* frames that have an error. The FCS can detect all single, double, triple errors, all error bursts of length ≤ 32 , most double bursts of length up to 17. The probability that a random collection of bit errors is undetected is $2e-32$.

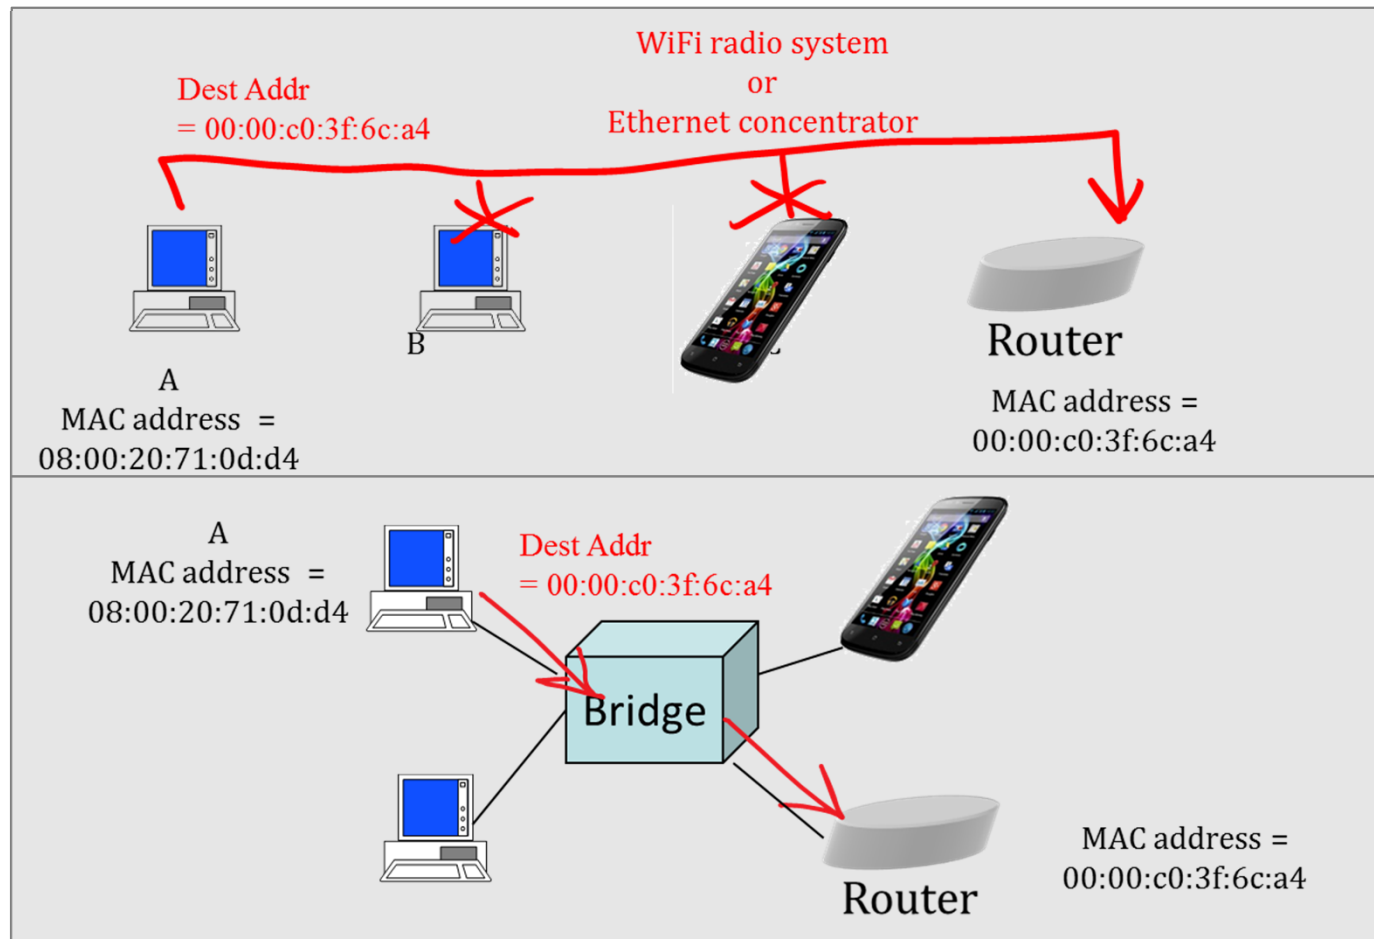
Addressing

MAC address: 48 bits = adapter number (in principle – can be configured). Unique worldwide, in principle. Sender puts destination MAC address in the frame

On shared medium LAN: all stations read all frames; keep only if destination address matches.

Bridges forward frames only to port that needs it (in principle)

Sent in the clear, no encryption



On a shared medium LAN (WiFi or Ethernet concentrator), all systems see the frame and discard it except destination.

With a bridge, only destination sees the frame

Ethernet addresses are known as MAC addresses. Every Ethernet interface has its own MAC address, which is in fact the serial number of the adapter, put by the manufacturer. MAC addresses are 48 bit-long. The 1st address bit is the individual/group bit, used to differentiate normal addresses from group addresses. The second bit indicates whether the address is globally administered (the normal case, burnt-in) or locally administered. Group addresses are always locally administered.

When A sends a data frame to B, A creates a MAC frame with source addr = A, dest addr = B. The frame is sent on the network and recognized by the destination.

Some systems like DEC networks require that MAC addresses be configured by software; those are so-called locally administered MAC addresses. This is avoided whenever possible in order to simplify network management.

Data on Ethernet is transmitted least significant bit of first octet first (a strange behaviour dictated by Intel processors). Canonical representation thus inverts the order of bits inside a byte (the first bit of the address is the least significant bit of the first byte);

The first 24 bits are the Organization Unique Identifier (OUI); the remaining 24 bits are allocated by organization that owns the OUI

examples of addresses:

08:00:20:71:0d:d4 (a SUN machine)

00:00:c0:3f:6c:a4 (a PC)


00:00:0c:02:78:36 (a CISCO router)

ff:ff:ff:ff:ff:ff the broadcast address

The MAC Broadcast Address


ff:ff:ff:ff:ff:ff = broadcast address

On shared medium LAN, all machines receive packet and do not discard it



With Bridges...

Broadcast frames are sent on all nodes and ports on the spanning tree



MAC Multicast Addresses

Multicast = Group ;

MAC Addresses with bit 8 == 1 are multicast addresses


Ethernet adapter discards such packets except if host subscribes to address

Non-smart bridges broadcast such frames; smart bridges send only to nodes that subscribe to multicast address using IPv4 or IPv6 subscription protocols
(see later)

Examples

01-00-5e-XX-XX-XX	IPv4 multicast
33-33-XX-XX-XX-XX	IPv6 multicast
01-80-c2-00-01-81	Precision Time Protocol

B subscribes to 01:00:5e:02:36:cf C subscribes to 01:00:5e:02:36:cf



The destination MAC address is...

- A. A group address derived from the last 23 bits of the IPv6 target address
- B. A group address derived from the last 24 bits of the IPv6 target address
- C. A group address derived from the last 32 bits of the IPv6 target address
- D. A broadcast address
- E. The MAC address of an ARP server
- F. I don't know

```
ETHER: ----- Ether Header -----
ETHER:
ETHER: Packet 1 arrived at 11:55:22.298
ETHER: Packet size = 86 bytes
ETHER: Destination = 33:33:ff:01:00:01
ETHER: Source = 3c:07:54:3e:ab:f2
ETHER: Ethertype = 0x86dd
ETHER:
IP: ----- IP Header -----
IP:
IP: Version = 6
IP: Traffic class = 0x00000000
IP:      .... 0000 00.. ....
IP:      ....      ..0. ....
IP:      ....      ...0 ....
IP: ....      .... 0000 0000 0000 0000 0000 =
IP: Payload length = 32
IP: NextHeader = 58
IP: Hop limit = 255
IP: Source address = 2001:620:618:197:1:80b2:9
IP: Destination address = ff02::1:ff01:1
IP:
```

IPv6

ICMP for IPv6


solicited node mu

The protocol type in an Ethernet frame header indicates whether...

- A. ... the data is a TCP packet or not
- B. ... the data is an IPv4 packet or not
- C. Both
- D. None
- E. I don't know

We have a working environment with IPv4, using Ethernet bridges plus other things. We add IPv6 to the network. Do we need to change the bridge configurations ?

- A. No, since bridges look only at MAC addresses
- B. Yes since the protocol type is different
- C. It depends whether SLAAC is used
- D. I don't know



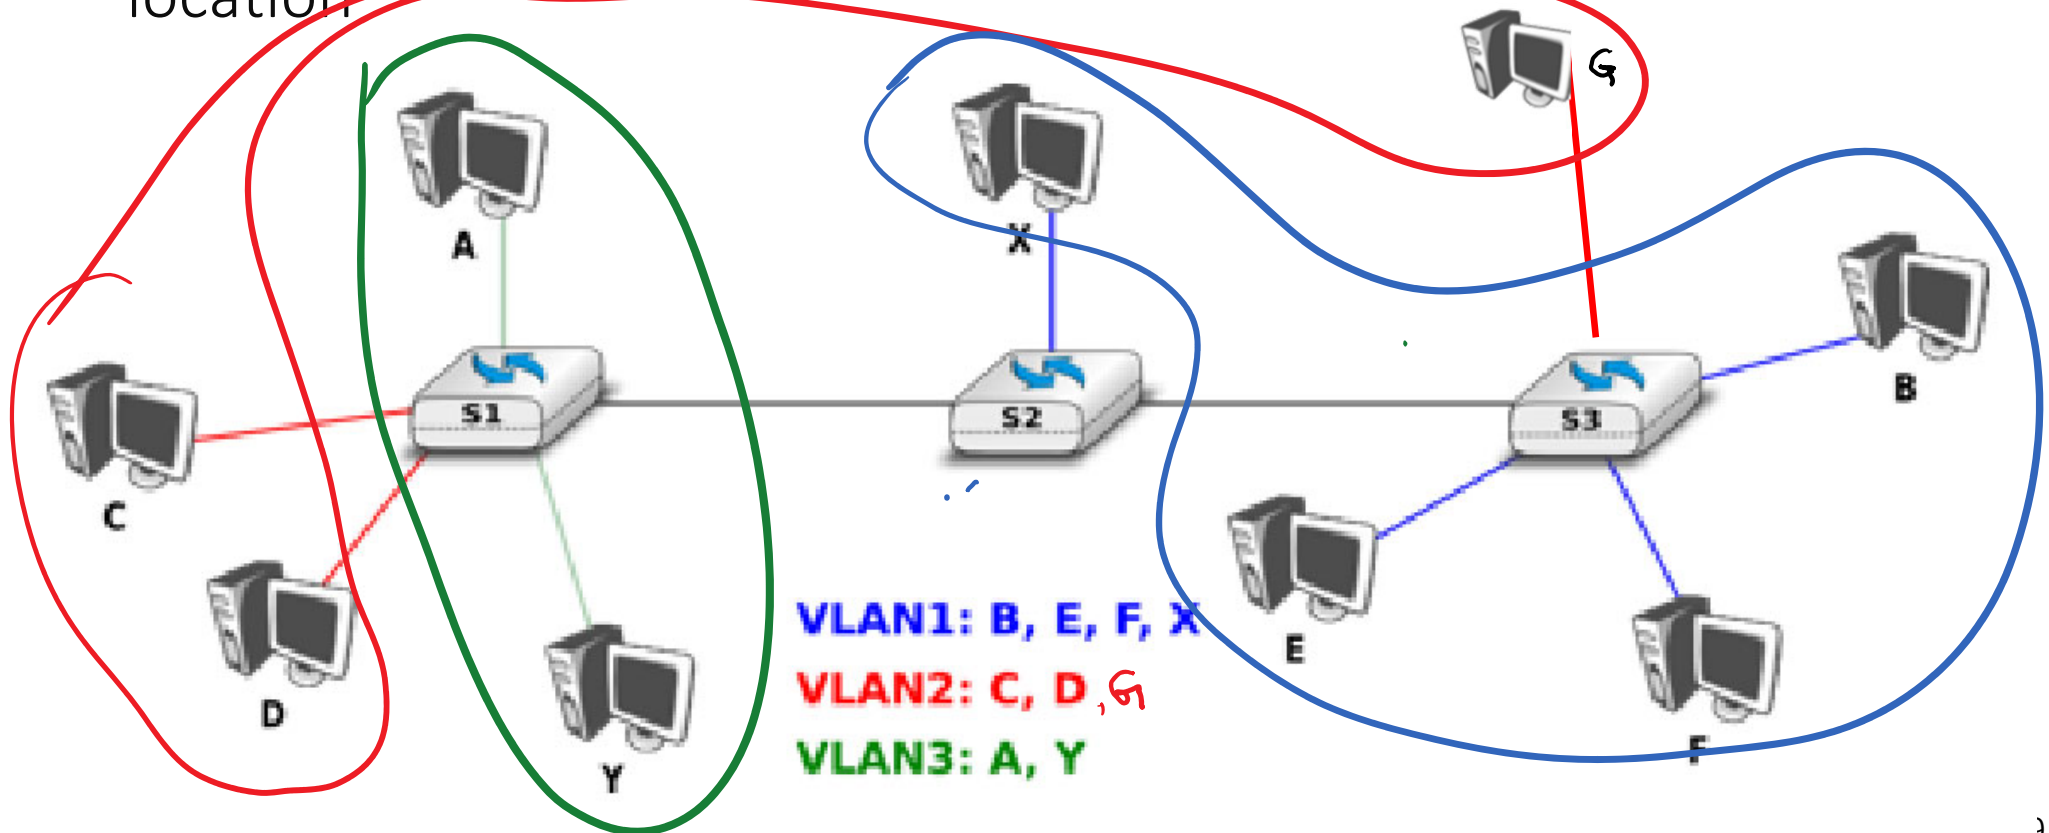
4. Virtual LAN

Why invented ?

Some systems use LAN as access control (e.g. all external web servers machines on one LAN)

Some applications work only on LANs (e.g. iTunes broadcast)

Goal: Decouple who belongs to which LAN from physical (office) location



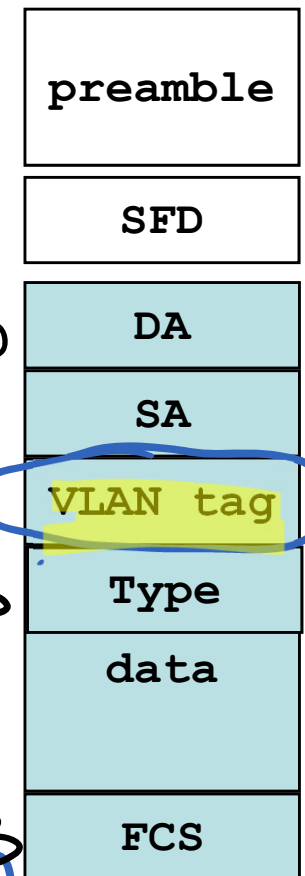
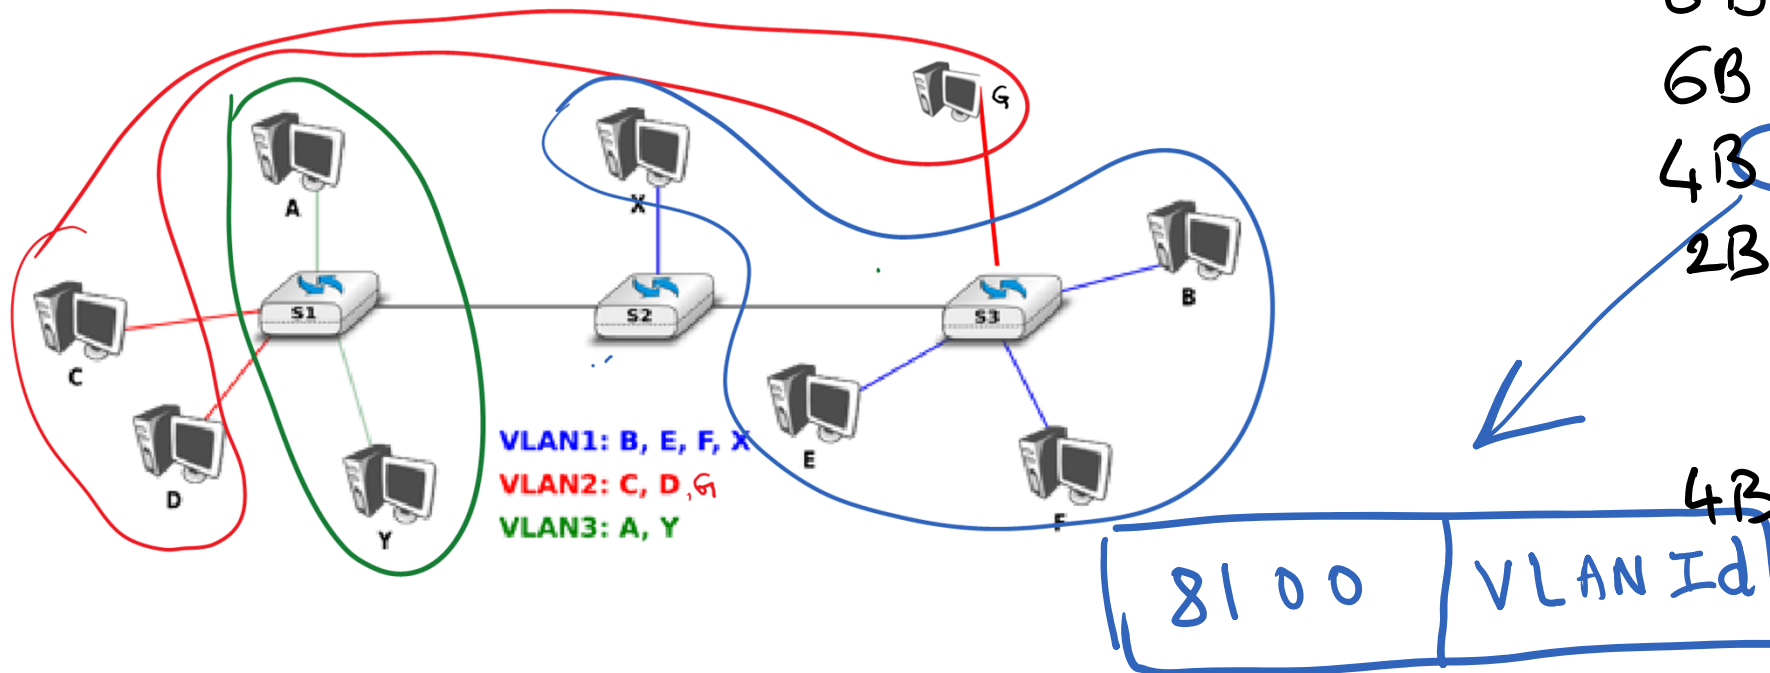
Virtual LAN

VLANs are handled as if they were physically separate: different forwarding tables, different spanning trees. How does it work ?

Configure (by network management) which switch port belongs to which VLAN


Switch handles ports of different VLANs as separate, non communicating worlds

VLAN label in Ethernet header on links between switches

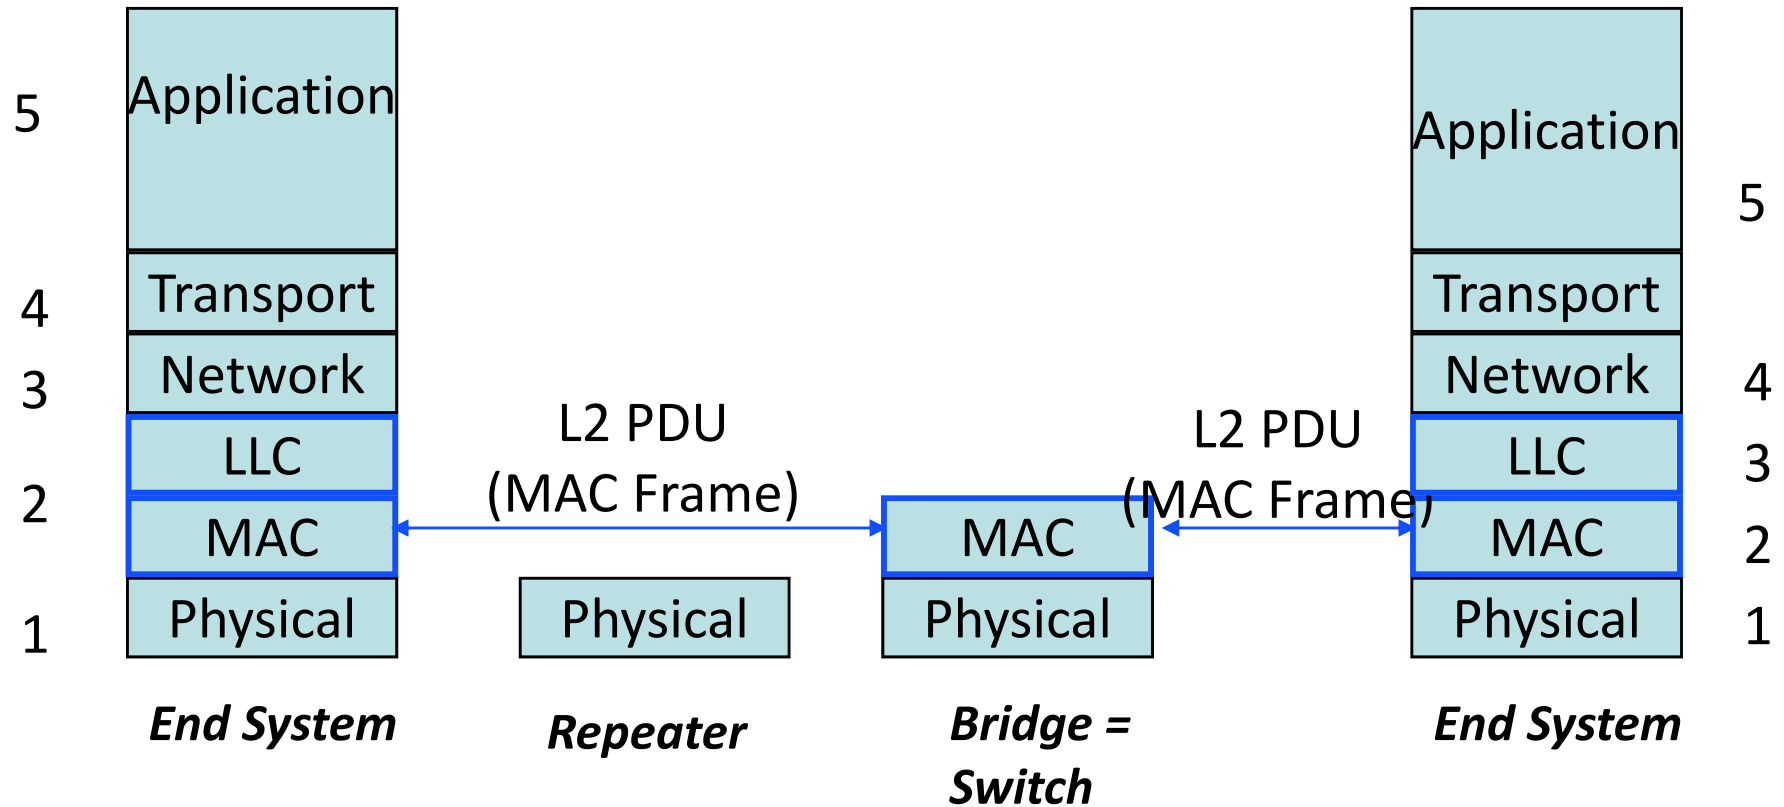


A has an IP packet destined to C; S1 is a switch; what should A do ?

- A. Send an ARP request in order to obtain the MAC address of C (then send packet to this MAC address)
- B. Send an ARP request in order to obtain both the MAC address and the VLAN label of C (then send packet to this MAC address and to this VLAN)
- C. None of the above
- D. I don't know



5. Repeaters and Bridges in TCP/IP Architecture




Bridges are layer 2 intermediate systems

Repeaters are layer 1 intermediate systems – use for example to extend an Ethernet cable through an optical fiber

Layer 3 intermediate systems are IP routers (not shown here)

Bridges can be used between different MAC layer technologies

We have here one single LAN




WiFi Access Points

A WiFi access point is usually the combination of

- A WiFi base station (implements the WiFi MAC protocol)

- A bridge with an Ethernet backend (called the **Distribution System**)

In this case, WiFi frames differ a bit from Ethernet frames in that there can be 3 MAC addresses: source, destination and access-point.

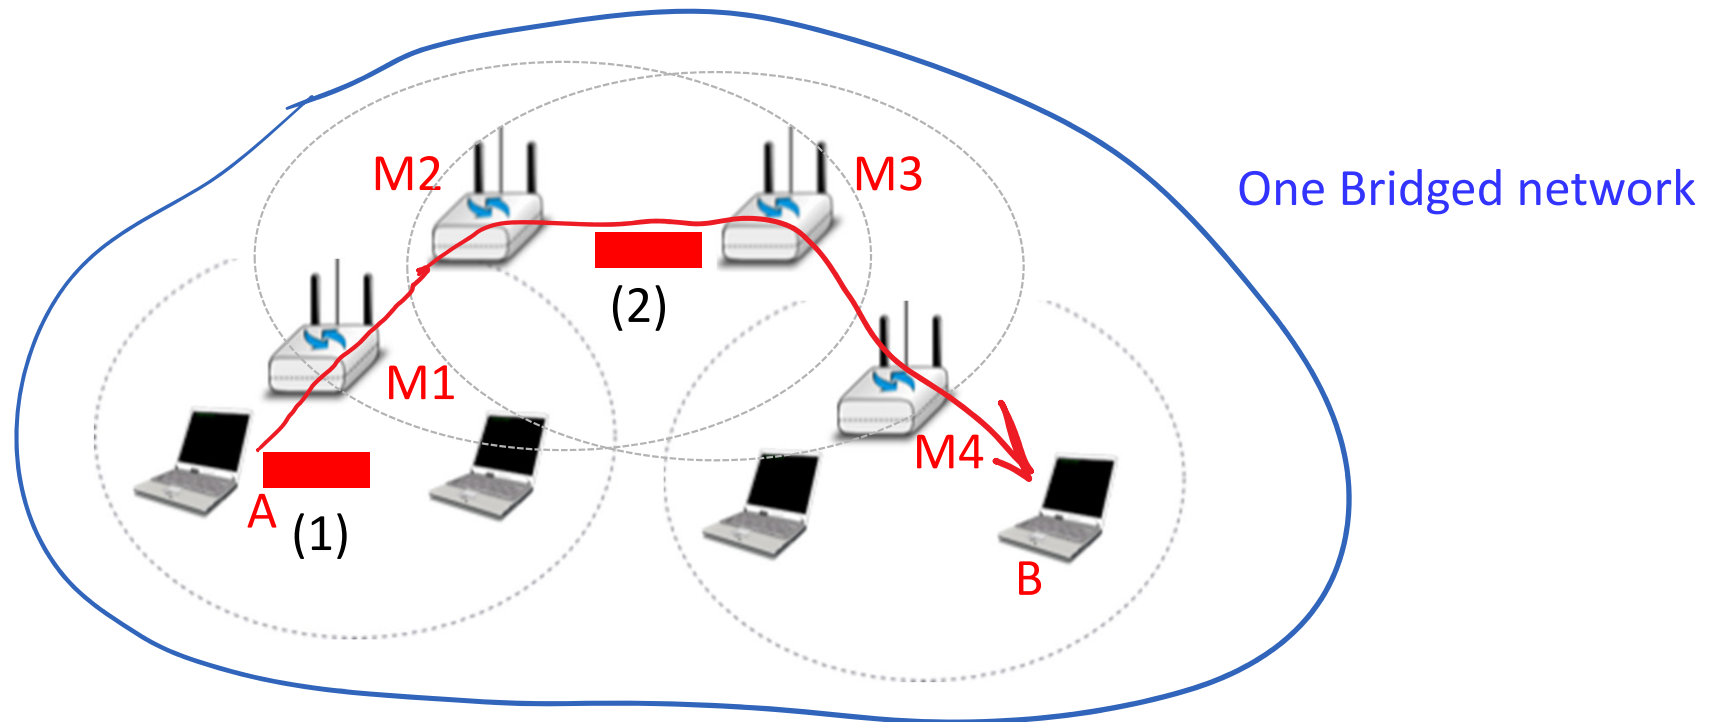


In packet sent by A to G: source MAC addr = A, dest MAC addr = G, access-point address = E

WiFi Access Points with Wireless Distribution System

In some cases, the interconnection between WiFi base stations can be done using WiFi relays, which act as bridges.

In this case, the WiFi frames may carry 4 MAC addresses: source, destination and receiver / transmitter MAC address.



In packet sent by A to B: at (1) source MAC addr = A, dest MAC addr = B, access-point address = M1
At (2) : source MAC addr = A, dest MAC addr = B, tx-access-point address = M2,
rx-access-point address=M3

What is a “Non Routable Protocol” ?

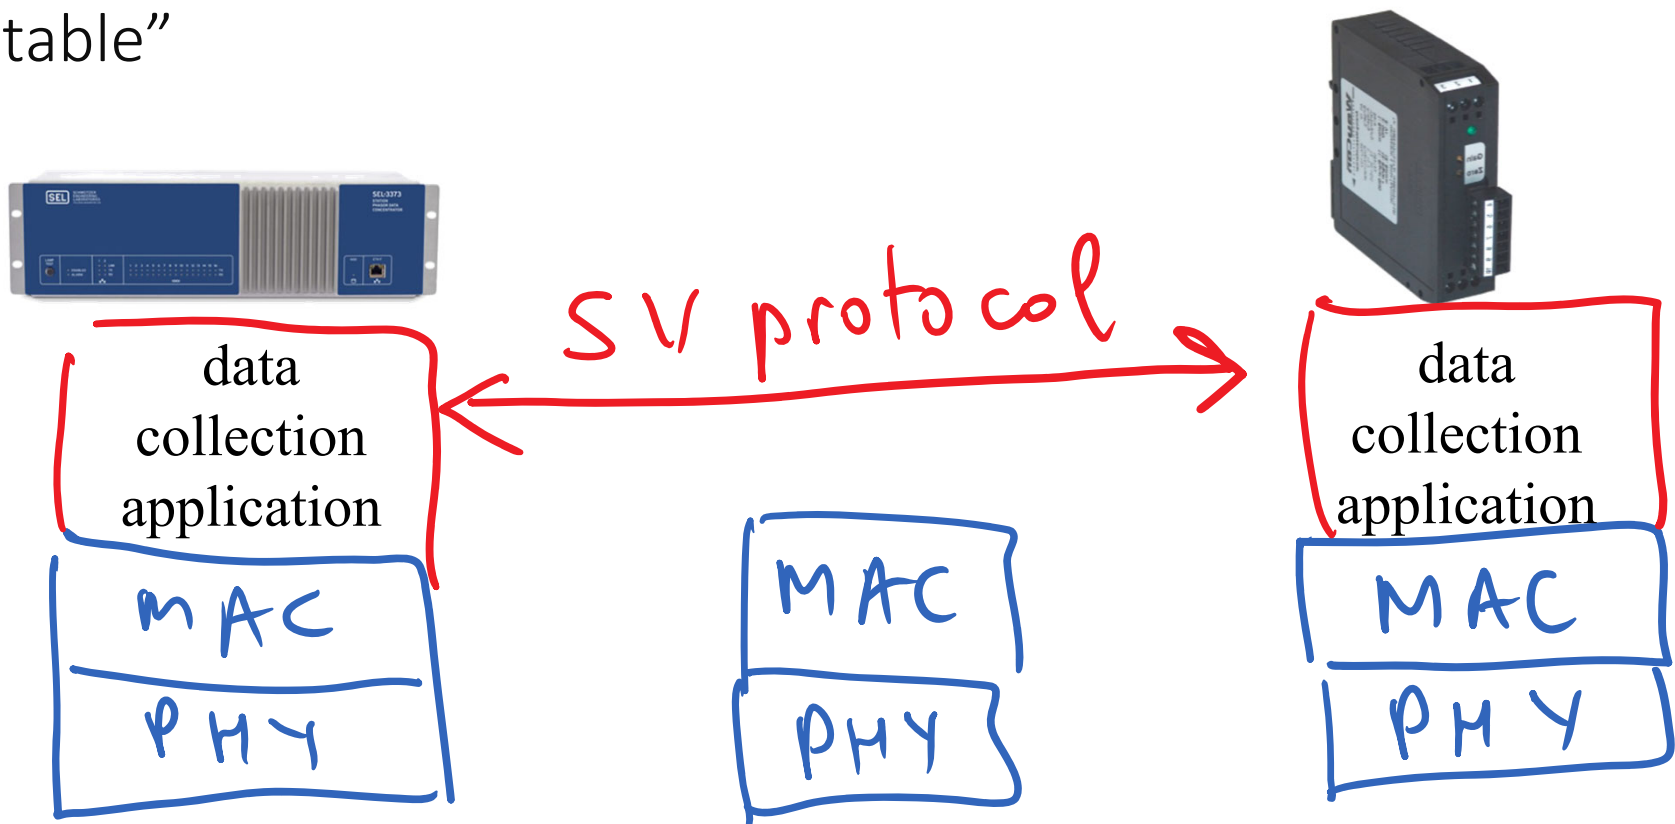
Some (old) application layers are built directly on top of MAC

Ex: some smart grid communication systems

Ex: old versions of NetBIOS (networking for PCs)

Ex: high precision time synchronization (PTP, IEEE 1588)


The corresponding protocol works only in one LAN, and is said “non routable”



SV = “sampled value” protocol


Web server sends one IP packet to Elaine. At (1), the MAC destination address is...

- A. M1
- B. M2
- C. M3
- D. M4
- E. I don't know



Web server sends one IP packet to Elaine. At (1), the MAC destination address is...

- A. M1
- B. M2
- C. M4
- D. I don't know



6. Security Aspects

MAC addresses are sent in the clear. Attacks that result ?

Eavesdropping (hearing someone else's data)

Free riding (inserting a device into a network without authorization)

Impersonation (send data with someone else's MAC address; replay attacks)

Solutions (MACSEC)

Access control require user to show credentials before allowing a given MAC address into the network;

- ▶ shared secret (e.g. WiFi WPA-Enterprise): user knows a password, same for all
- ▶ per-user authentication (e.g. WiFi WPA-Enterprise): user has a personal password

Authentication: every MAC frame is signed using cryptography and numbered

- ▶ prevents free riding
- ▶ prevents impersonation if access control is per user;

Encryption: MAC frame payload is encrypted (not MAC address) -> prevents eavesdropping

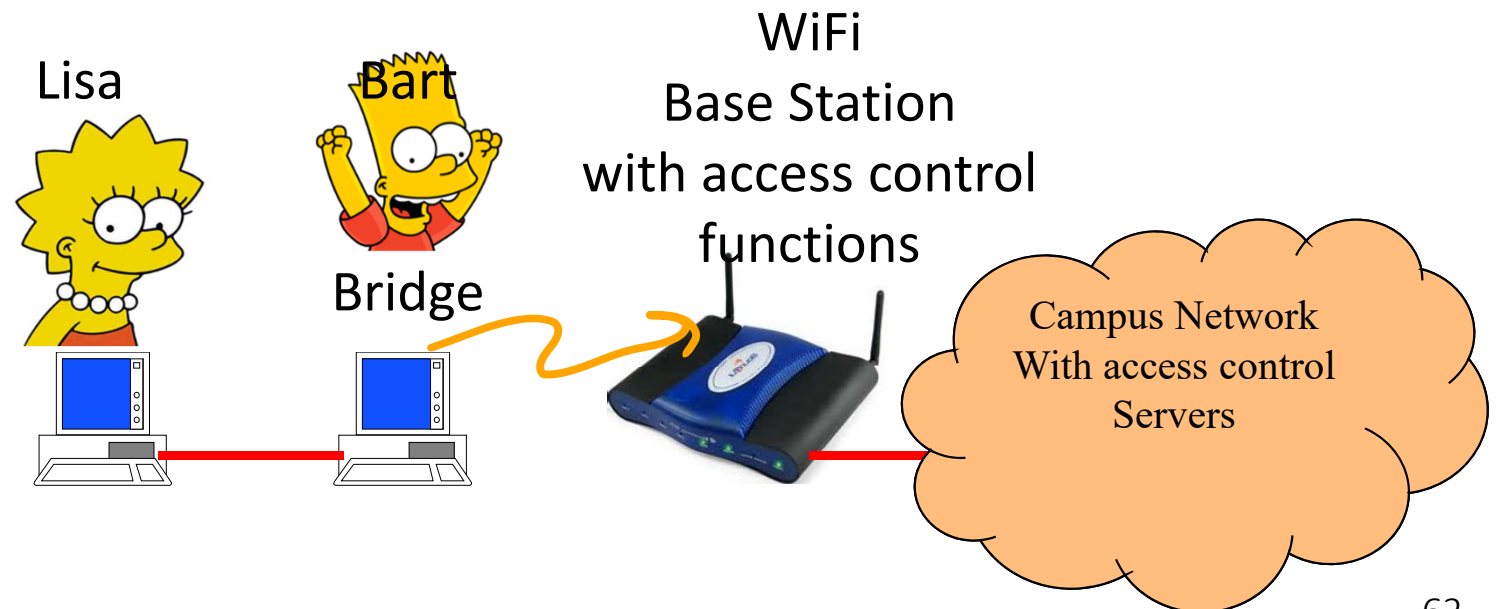
Examples:

At EPFL, WiFi uses access control, authentication and encryption; impersonation is not possible. In contrast, Ethernet is not secured.

At home: WiFi WPA-Personal provides access control but does not authenticate individual users. Impersonation is possible by who knows the shared password.

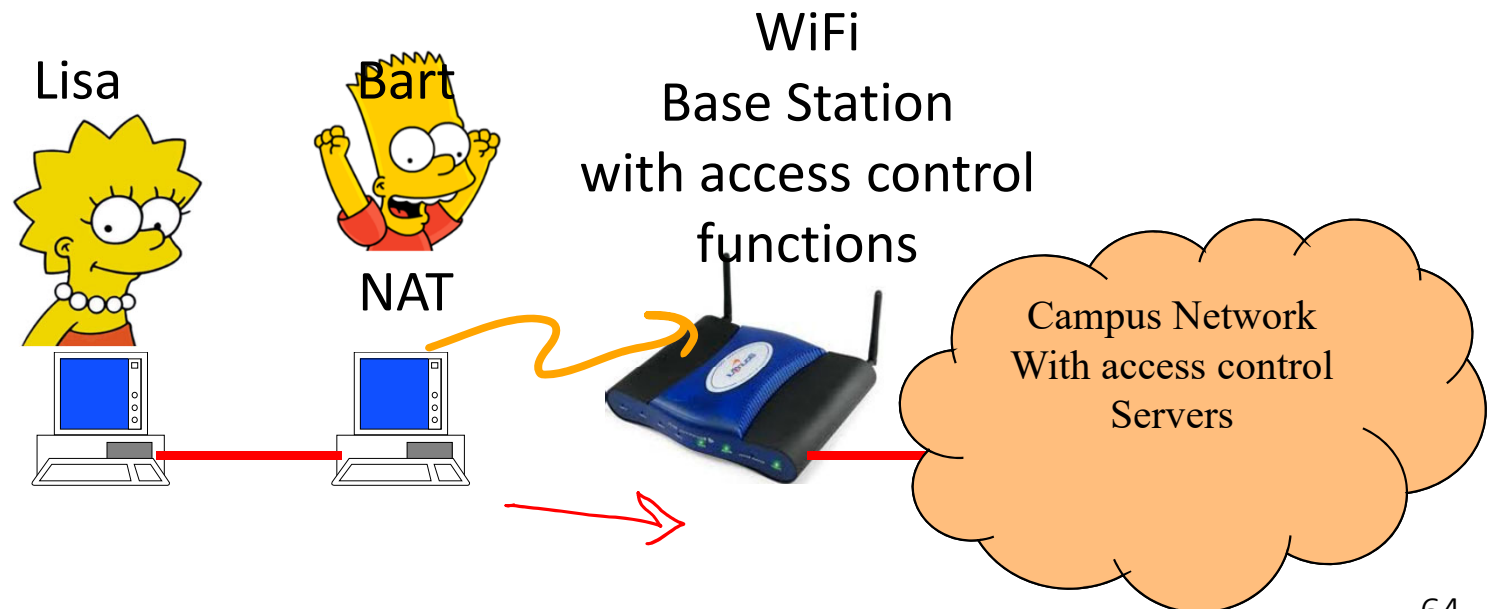
Bart, who is EPFL student, connects his PC to EPFL WiFi, which uses per-user access control, authentication and encryption. Bart configures his PC to be a *bridge* between Ethernet and WiFi and allows Lisa, who is not allowed to access EPFL, to connect to his Ethernet port. Will Lisa be able to connect to the Internet in this way ?

- A. Yes, provided that Lisa can configure an IP address on the same subnet as Bart
- B. Yes, it will work in all cases
- C. No, it will not work in any case
- D. I don't know



Bart, who is EPFL student, connects his PC to EPFL WiFi, which uses per-user access control, authentication and encryption. Bart configures his PC to be a *NAT* between Ethernet and WiFi and allows Lisa, who is not allowed to access EPFL, to connect to his Ethernet port. Will Lisa be able to connect to the Internet in this way ?

- A. Yes if Bart himself is not behind a NAT
- B. Yes
- C. No, it will not work in any case
- D. I don't know



Conclusion 1

The MAC layer for wireless medium (WiFi) takes care of sharing the radio waves in a local environment.

The MAC layer for wired medium (Ethernet) was originally designed for sharing a cable; in this form it uses a protocol, called CSMA/CD very similar to WiFi.

The MAC layer for wired medium (Ethernet) has now got rid of any protocol (“full duplex Ethernet”), using interconnected bridges. It thus forms an interconnection layer of local scope.

Conclusion 2

We have seen two mechanisms to interconnect devices: the network layer (with IP routers) and the MAC layer (with bridges). What are the differences ?

Bridges use routing tables that are not structured. A bridge must search the entire table for *every* packet. The table size and lookup time are prohibitive if network becomes very large. We prefer routers for large networks.

Bridges are independent of whether IPv4 or IPv6 is used. They are plug and play and are independent of IP numbering plans.