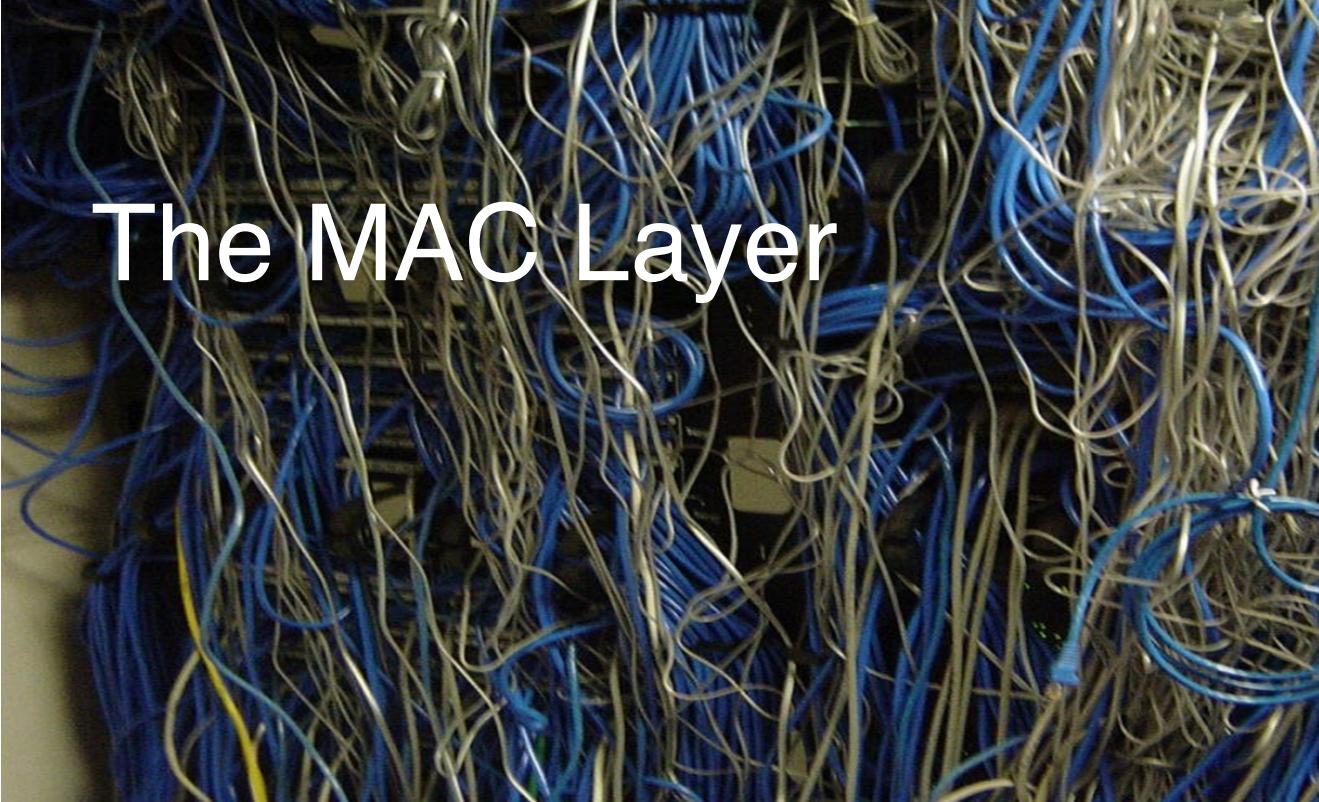




# The MAC Layer

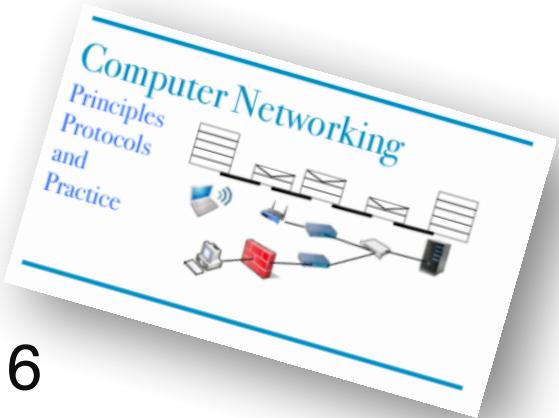


# Contents

1. MAC as Shared Medium
2. Switches
3. Frame format and MAC addresses
4. Virtual LANs
5. WiFi Distribution Systems
6. Security aspects

## Textbooks

Part 6



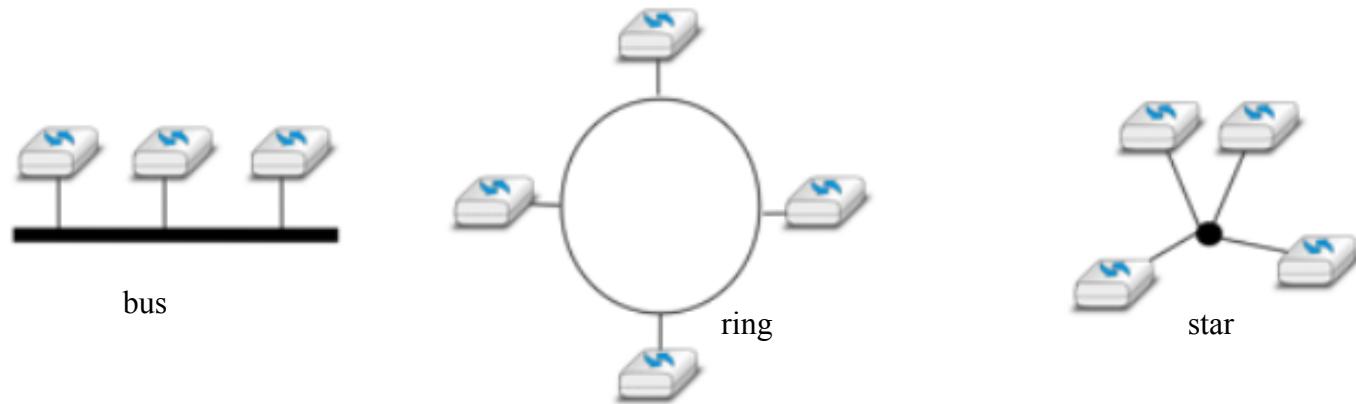
Chapter 6



# 1: MAC: Medium Access Control

Why invented? To *share* a single physical medium to which LAN devices are usually connected:

- cable (old/coaxial Ethernet, Powerline); or
- wireless radio link (GSM, WiFi, Bluetooth)



What is the problem ?

If several systems transmit at the same time, several frames will appear on the transmission medium (channel) and they will *collide/interfere*.

Solution 1: joint decoding (allow everyone to talk and decode) -- used in cellular networks (e.g. CDMA)

Solution 2: *mutual exclusion* (only one system can talk at a time), ideally via a *distributed* protocol

# Ways/protocols to apply mutual exclusion

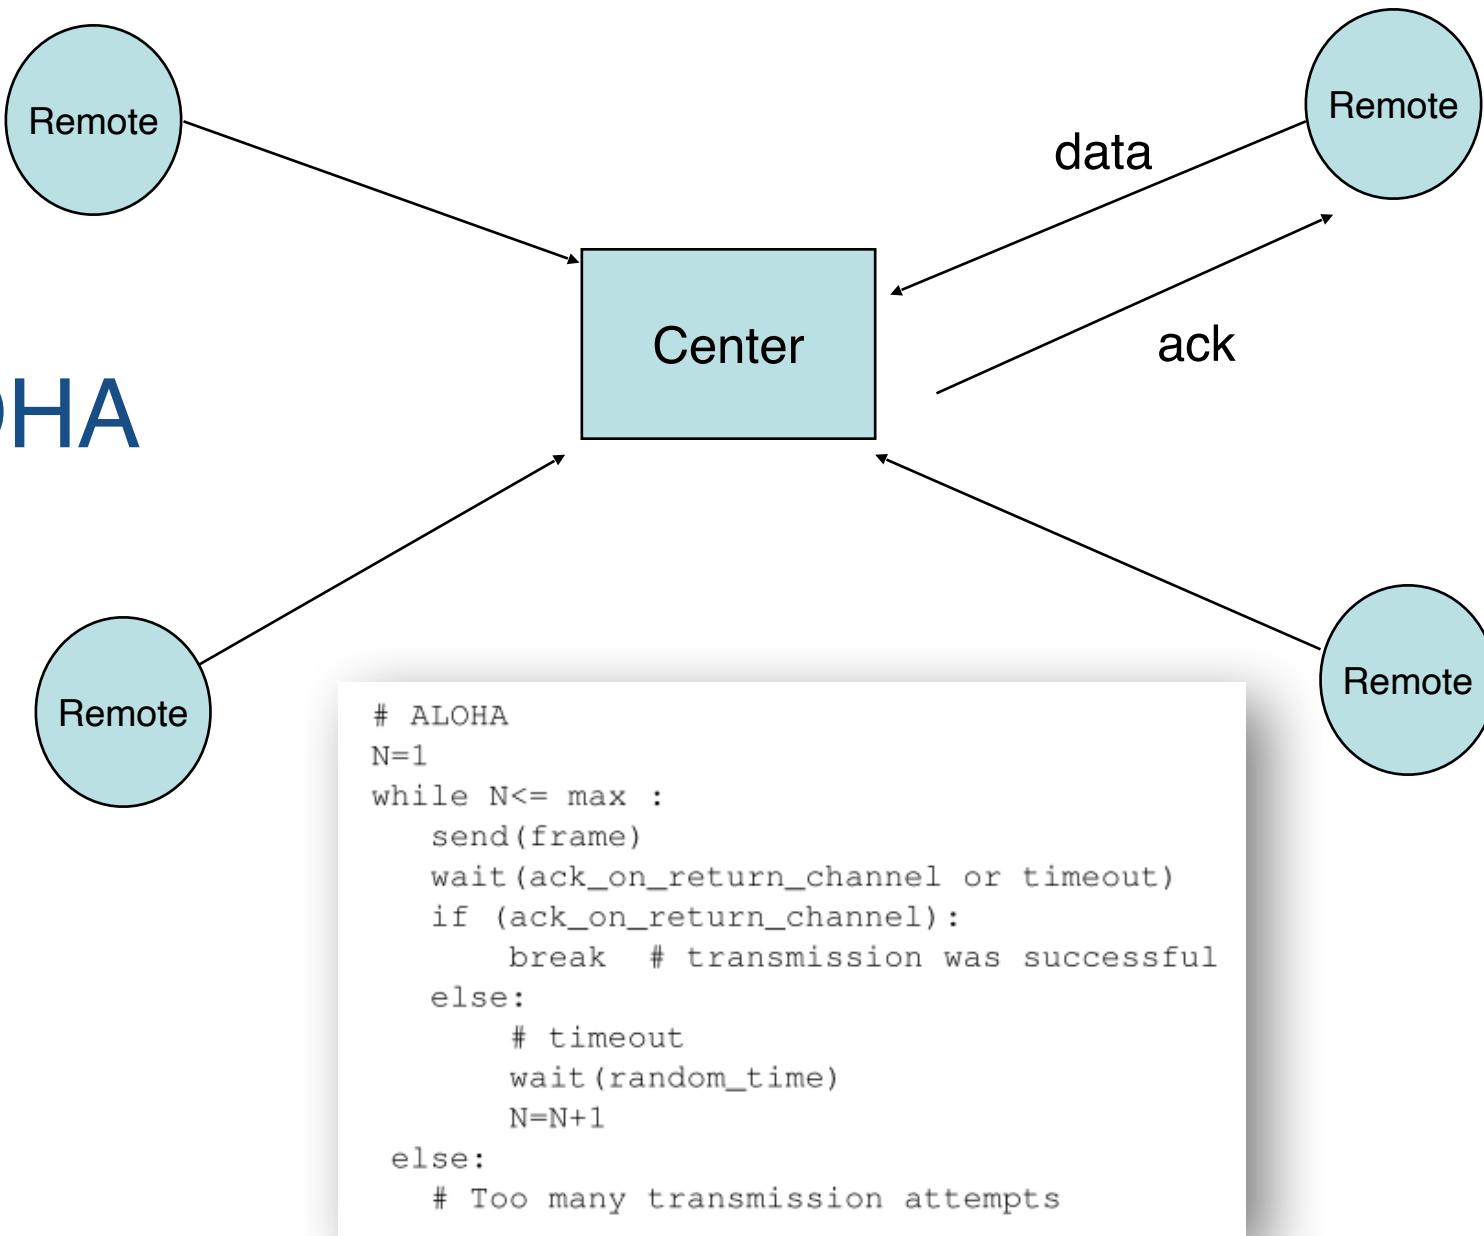
## *Deterministic protocols*

- use an **arbitration** for accessing the transmission medium (channel)
- avoid collisions by construction, **proactively**
  - Time Division Multiple Access (TDMA): each host is given a dedicated time slot to transmit
  - Frequency Division Multiple Access (FDMA): each host is given a dedicated frequency band to use
  - CAN bus: the highest-priority transmitting host dominates other transmissions
  - Token ring, FDDI: every host takes turn and passes a token to next host; whoever has the token, can transmit
    - ♦ more flexible than TDMA and FDMA, may achieve higher throughput
    - ♦ similar concepts appear also in application layer
- but... these work only with few users, because they require synchronization and pre-configuration (if a host goes down or a new one appears, we need to take the network down and reconfigure it)

## *Stochastic (random access) protocols*

- **allow everyone** to transmit at any time, **detect collisions**, and **react**
- hosts transmit at full rate and if a collision is detected, they retransmit after waiting for **random** time (“**back-off**”)
  - WiFi, coaxial/historical Ethernet, Zigbee, signaling channel on cellular networks, Ethernet over Powerline
  - explained in the next slides through ALOHA and CSMA/CD

# ALOHA



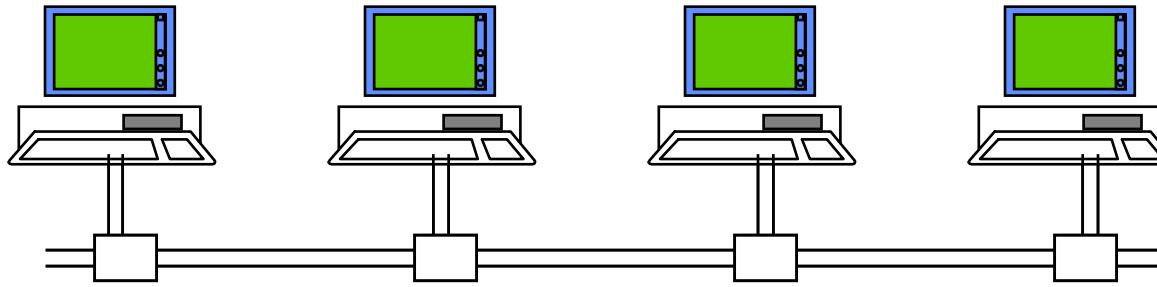
shared channel from  
remotes to center

Transmission procedure at remote

# ALOHA

- Aloha is the basis of all non-deterministic access methods. The Aloha protocol was originally developed for communications between islands (University of Hawaïi) 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** is 18% of the channel capacity, which is very small. This is assuming independent and exponential inter-arrival times of packets and retransmissions.
- **Slotted ALOHA** can improve utilization by a **factor of 2**.

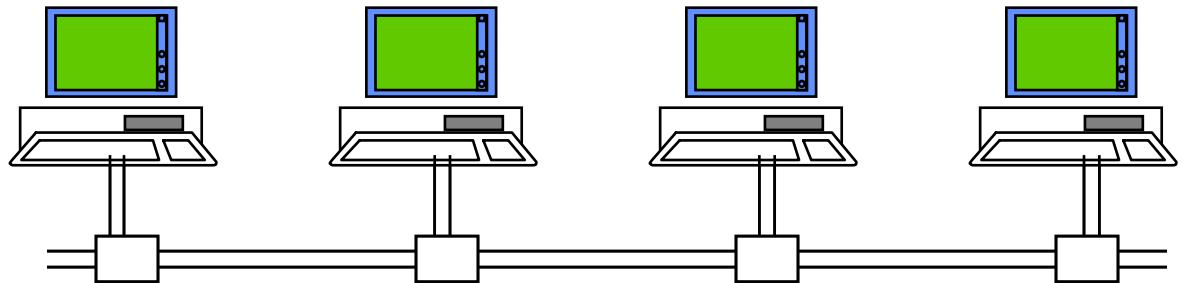
# CSMA (carrier sense multiple access)



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

- drastically *improves efficiency* over Aloha
- requires that hosts “*senselisten* the channel before transmitting
- assumes a single transitive channel (everyone can hear everyone)

# In CSMA, collisions...



- A. cannot occur because the medium is transitive
- B. cannot occur if transmission delays are very small
- C. can occur regardless of the transmission delay
- D. I don't know



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045

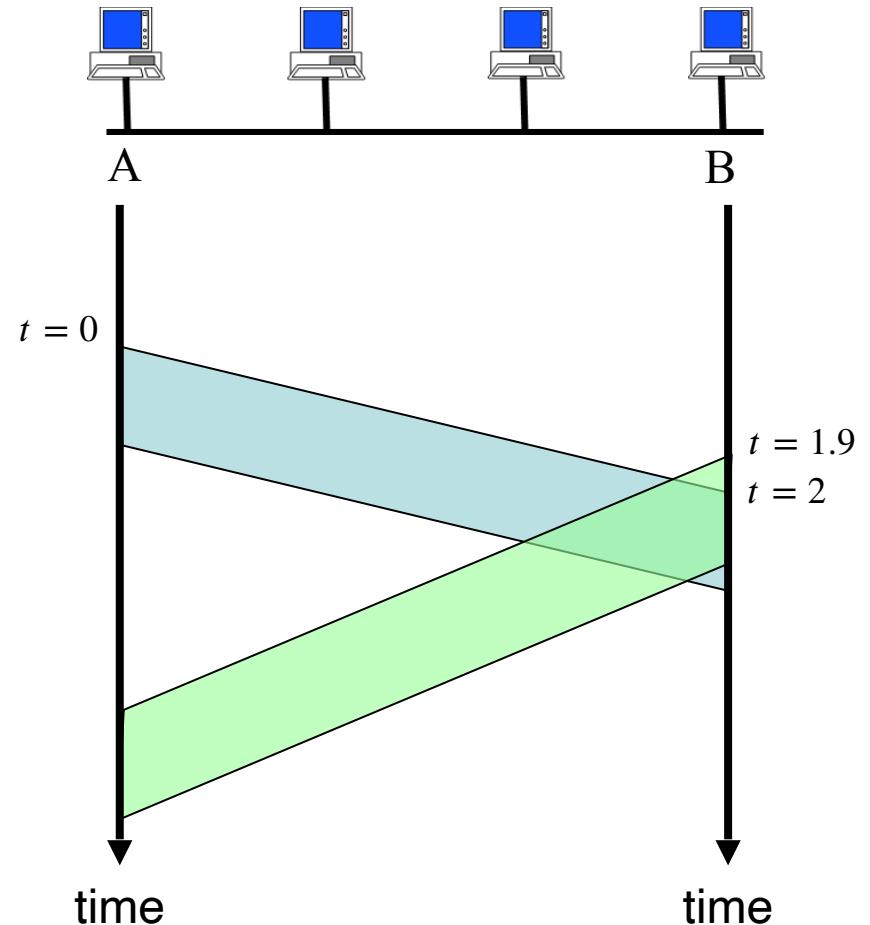
# Solution

Many collisions can be avoided, but not all.

Two or more stations may sense that the medium is idle and start transmitting at time instants that are close enough for a collision to occur.

Example in figure:

- Assume propagation delay between A and B is 2 ms and that all stations are silent until  $t = 0$ .
- At  $t=0$ , station A starts transmitting.
- At  $t=1.9$ ms, station B has not received any signal from A yet, so it also starts transmitting.
- At  $t=2$ ms, station B senses the collision but it is too late —> a collision has occurred.

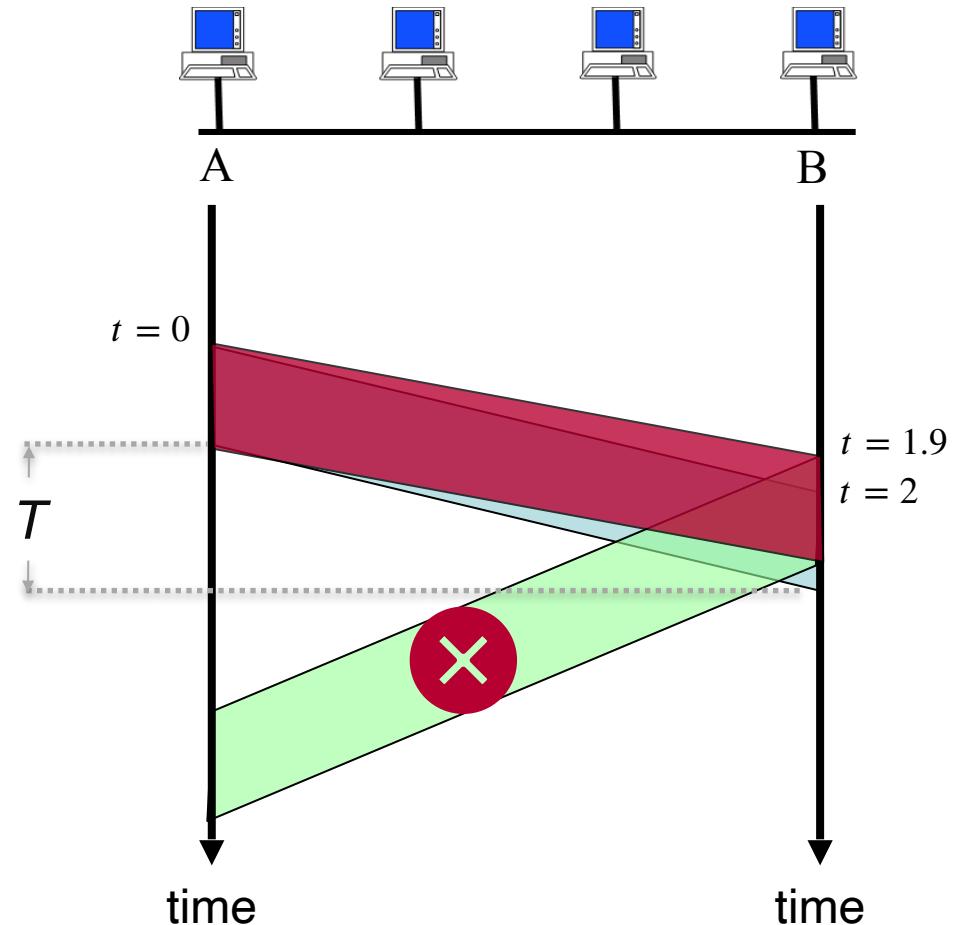


# Importance of propagation delay

The effect of CSMA can be expressed as follows:

- Let  $T$  be maximum propagation delay from station  $A$  to any other stations;
- if no collision occurs during a time duration  $T$  after  $A$  starts transmitting, then  $A$  has *successfully* seized the channel (= no other station can send)

→ CSMA works better if: propagation delay is *small*



# Non-persistent CSMA

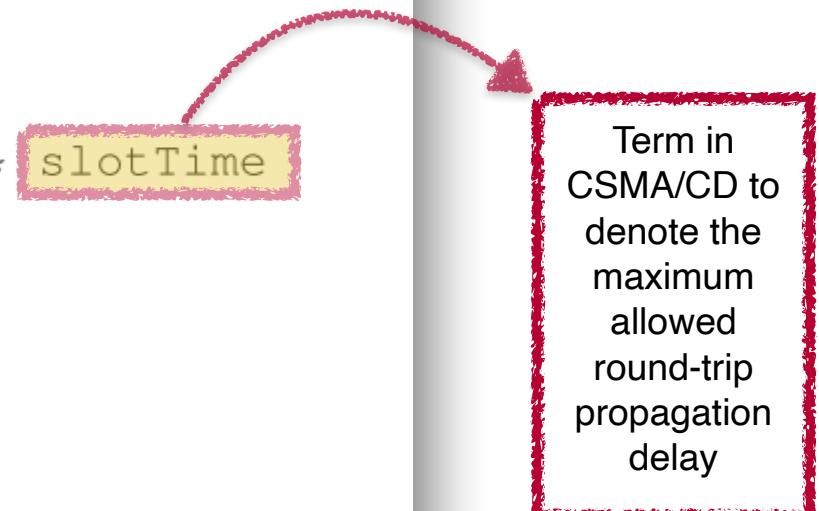
- uses *random* waiting times if channel is busy
- avoids hosts' synchronization
  - if all stations choose the random delays *independently*, and
  - if the waiting times are *larger* than the propagation delay,
  - then, with high probability, only one of the transmitting stations seizes the channel
- improves utilization

```
# 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
```

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

*Ack replaced by CD*

```
# 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
```



# CSMA / CD Time Diagram

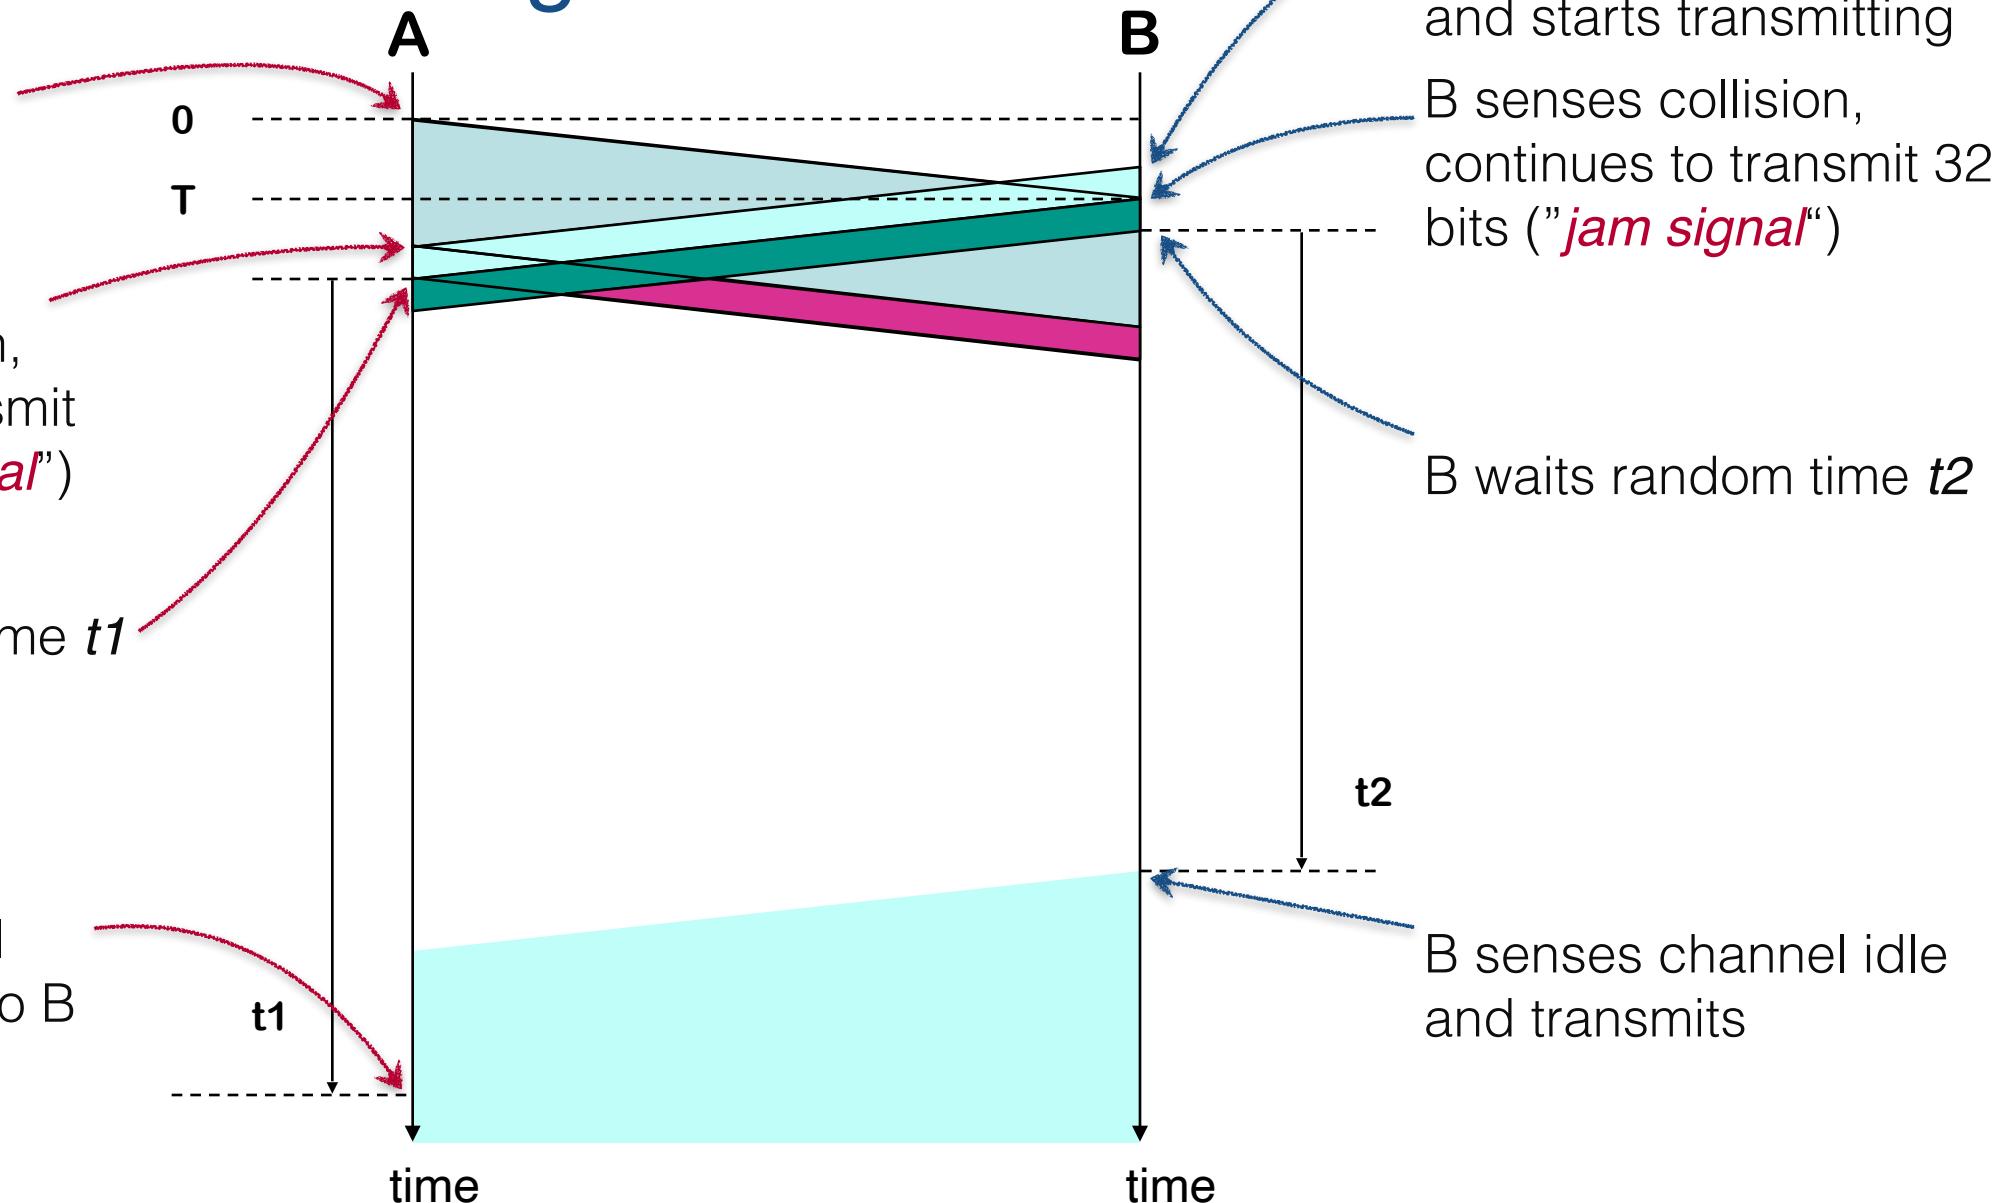
A senses idle channel, starts transmitting

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

A waits random time  $t_1$

A senses channel busy and **defers** to B

A now waits until channel is idle

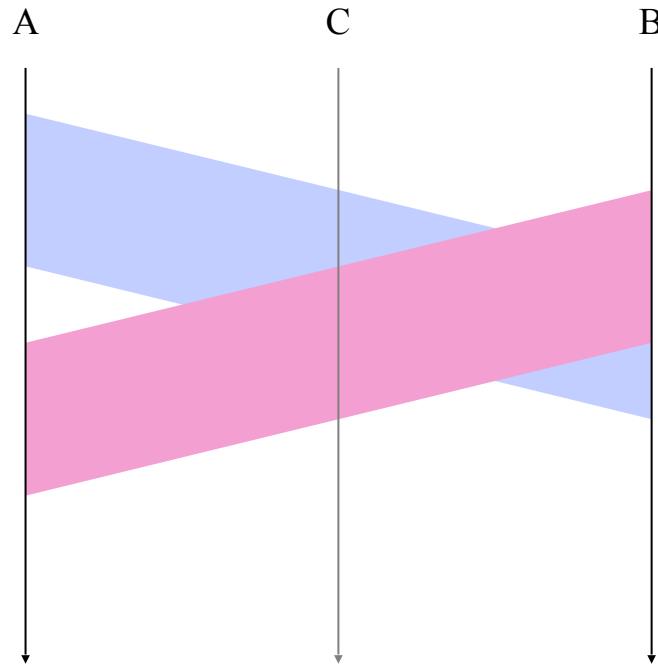


# CSMA / CD

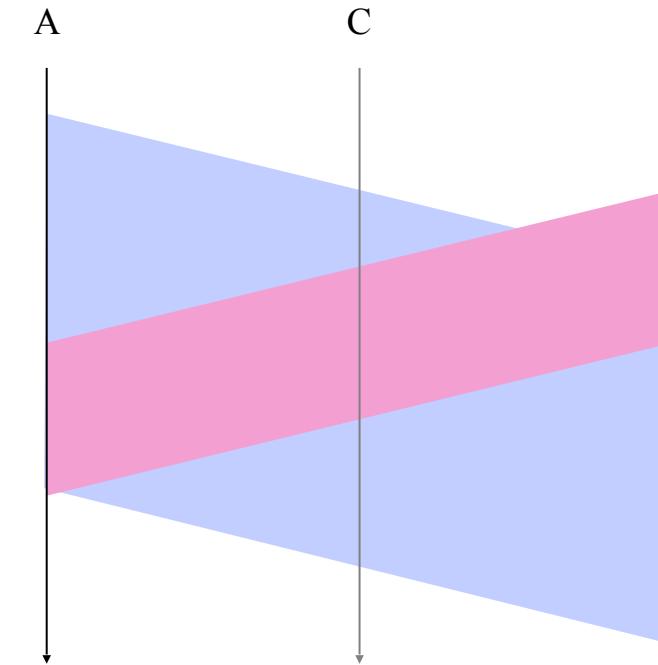
- CSMA/CD improves on CSMA by requiring that stations detect collisions and stop transmitting (after 32 bits, called *jam* bits, in order to ensure that all circuits properly recognize the presence of collisions).
- The jam bits (or jam signal) is a 32-bit binary pattern indicating that a collision has occurred.
- 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.
- The random time before retransmission is chosen in such a way that if repeated collisions occur, then this time increases exponentially (**“exponential backoff”**). To see this, check again the pseudocode from a couple of slides earlier.  
The downside of the exponential backoff is that: in case of congestion (too many collisions), the access to the channel is slowed down.
- Acknowledgements are not necessary because absence of collision means that the frame could be transmitted (see “Minimum Frame Size” in next slide).
- An inter-frame delay (“gap”) of 9.6  $\mu$ s is used to avoid blind times, during which adapters are filtering typical noise at transmission ends.

# A Minimum Frame Size is needed to guarantee collision detection

Problem: Frame sent by A is *too short*;  
collision is not visible at A (but only at C)



Solution: Frame sent by A is *larger*;  
collision is visible at A



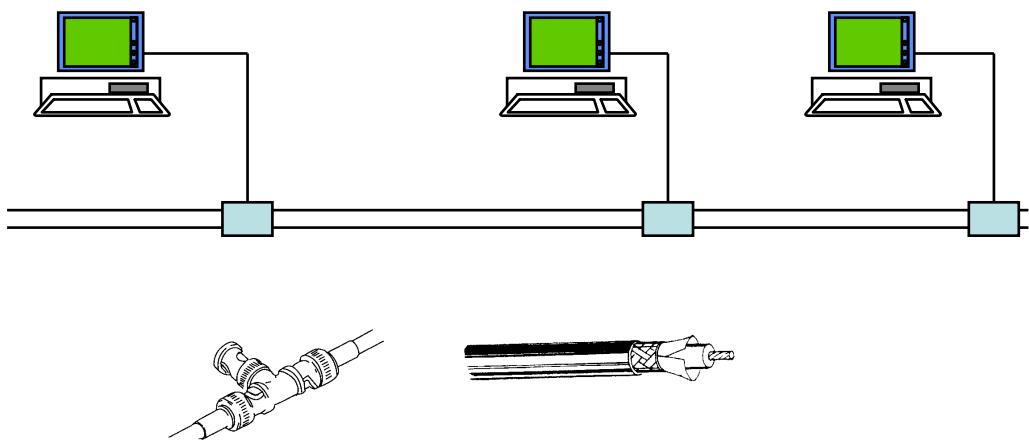
- To avoid undetected collisions, we impose a minimum frame size that depends on the link's *bit rate* and round-trip propagation delay (*slotTime*):

$\text{transmission delay} = \text{frameSize}/\text{bitRate} > \text{slotTime}$

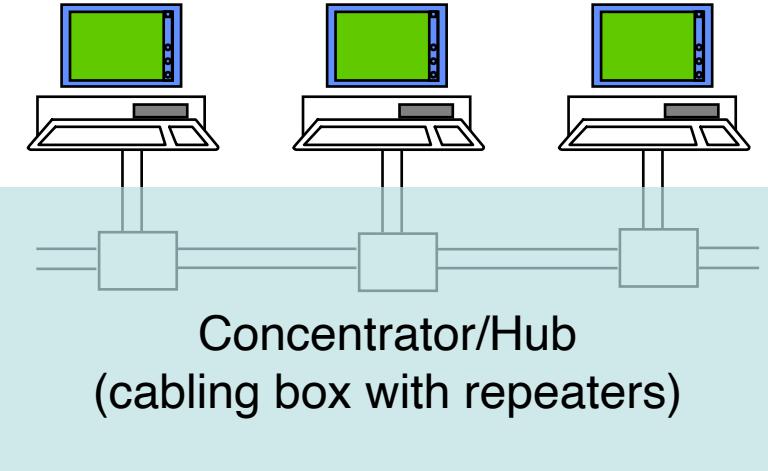
$\Rightarrow \text{frameSize} > \text{slotTime} * \text{bitRate}$  →  $\approx$  bandwidth-delay product

# Historical Ethernet (on coax cables or hubs) = CSMA/CD

Coax cables with T-connectors



Concentrator/Hub-based topology



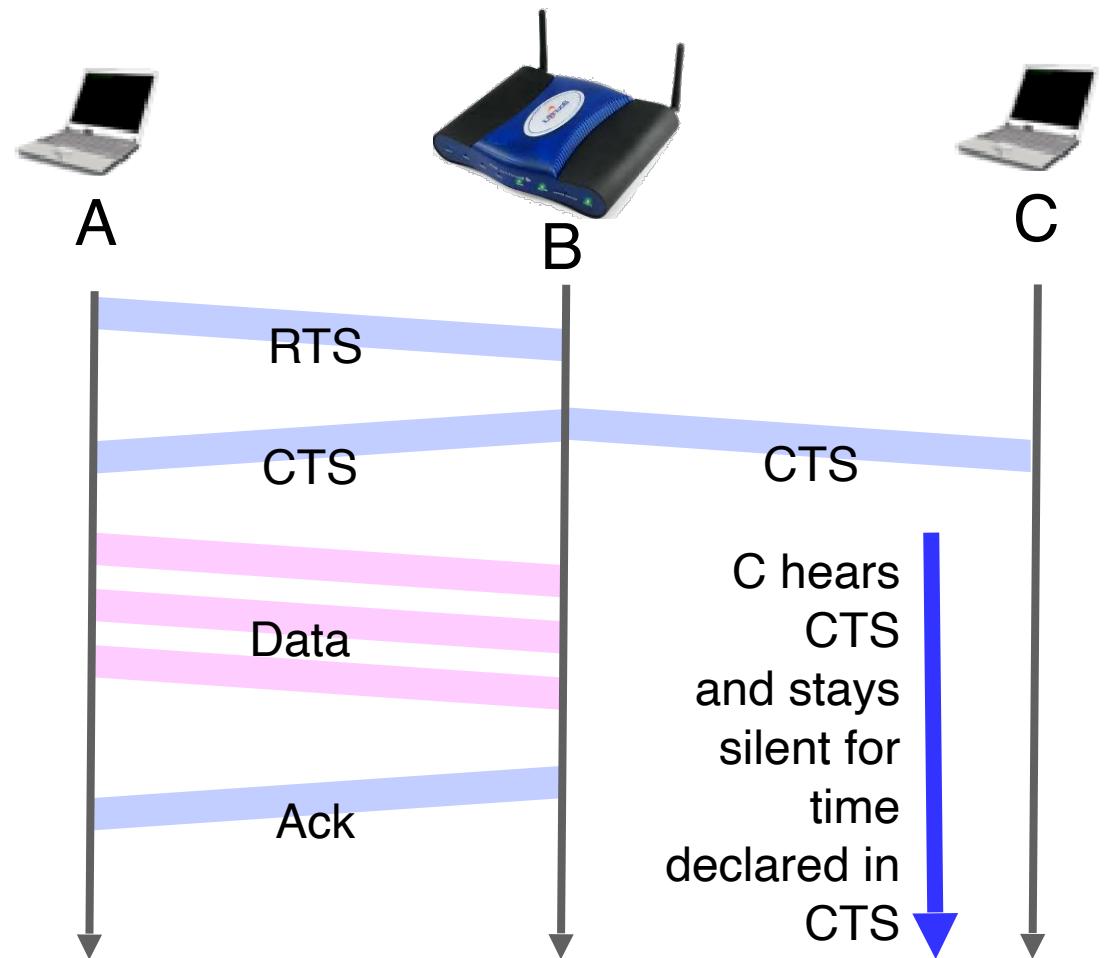
Old Ethernet, a.k.a. **half-duplex** Ethernet, addresses **collisions** by using:

- CSMA/CD
- with restrictions about the slotTime (hence also the min frame size)
  - e.g. 10Mbps half-duplex Ethernet uses  $slotTime = 51.2ms \Rightarrow min\ frame\ size = 64Bytes$

# WiFi = CSMA/CA (Collision Avoidance) = variant of CSMA/CD

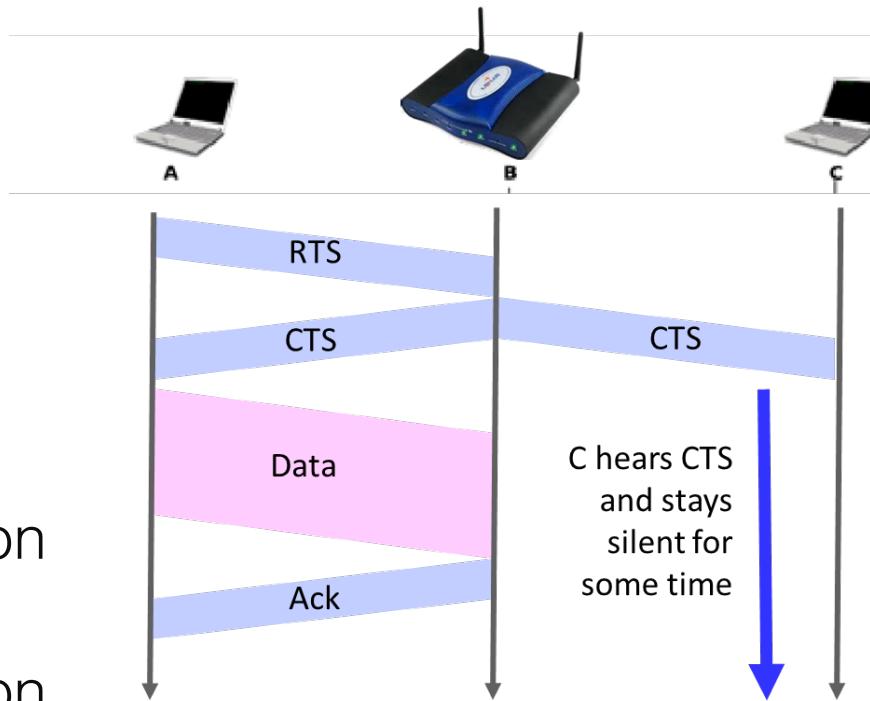
Uses:

- ACKs to detect collisions (i.e. if they are absent)
- Carrier sensing, but does not always work:
  - e.g., C may be “hidden” from from A;
  - and this is *mitigated* through **RTS/CTS**
- Further optimizations to avoid collisions:
  - more **conservative channel backoff**: even if the channel is idle, still wait for random time before transmitting
  - **network allocation vector**: MAC header contains the transmission time needed for the frame; it is a form of **virtual** carrier sensing



# In a WiFi network with a single channel and a single base station, how many data 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



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app  
Join room  
46045

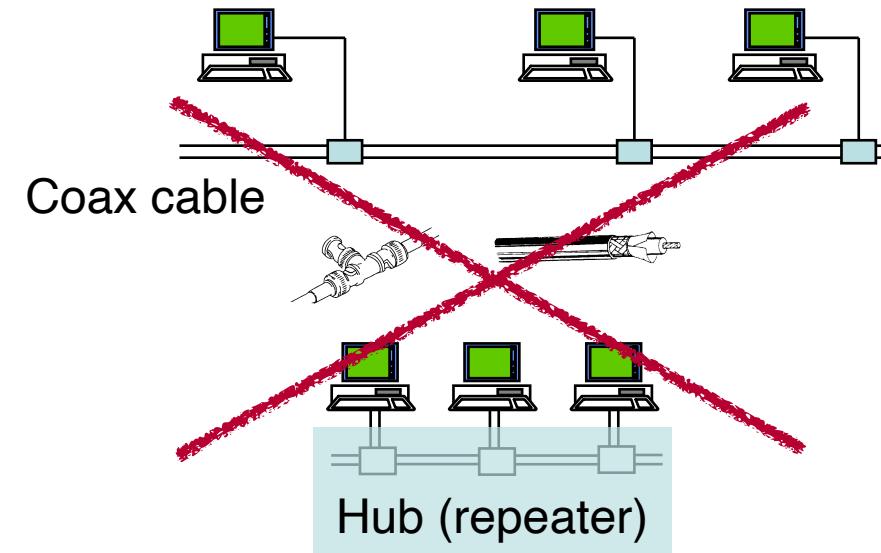
# Solution

Answer A

The WiFi MAC layer controls the access of all senders to the radio channel. So, it makes sure that only one can transmit at a time.

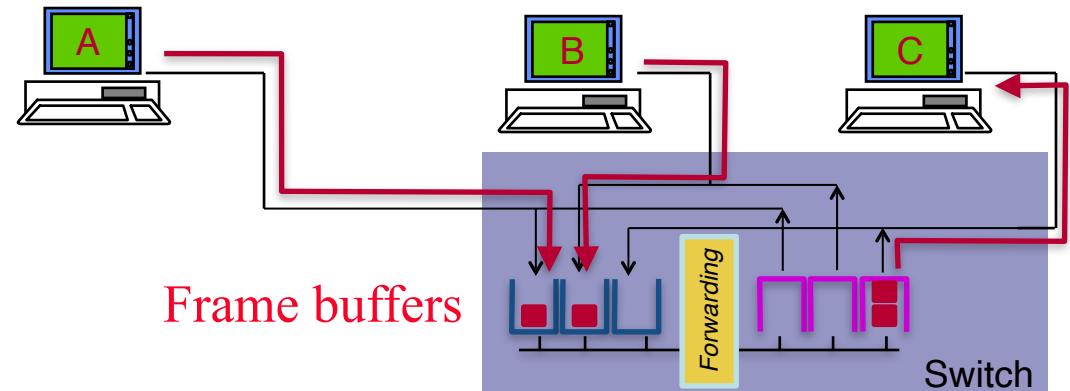
## 2. Today's switched Ethernet

Old-style Ethernet was shared medium (either used coax cables or repeaters/hubs with twisted-pair copper wires)



Today, Ethernet is based on:

- ***switches* (or *bridges*):**
  - = middle boxes of the MAC layer
  - = queuing systems (with frame buffers) that forward frames based on MAC addresses
- ***full-duplex cables*:**
  - frames flow simultaneously in both directions without suffering collisions (***collision-less*** ethernet)
  - e.g. the well-known **UTP** cables with separate downlink/uplink



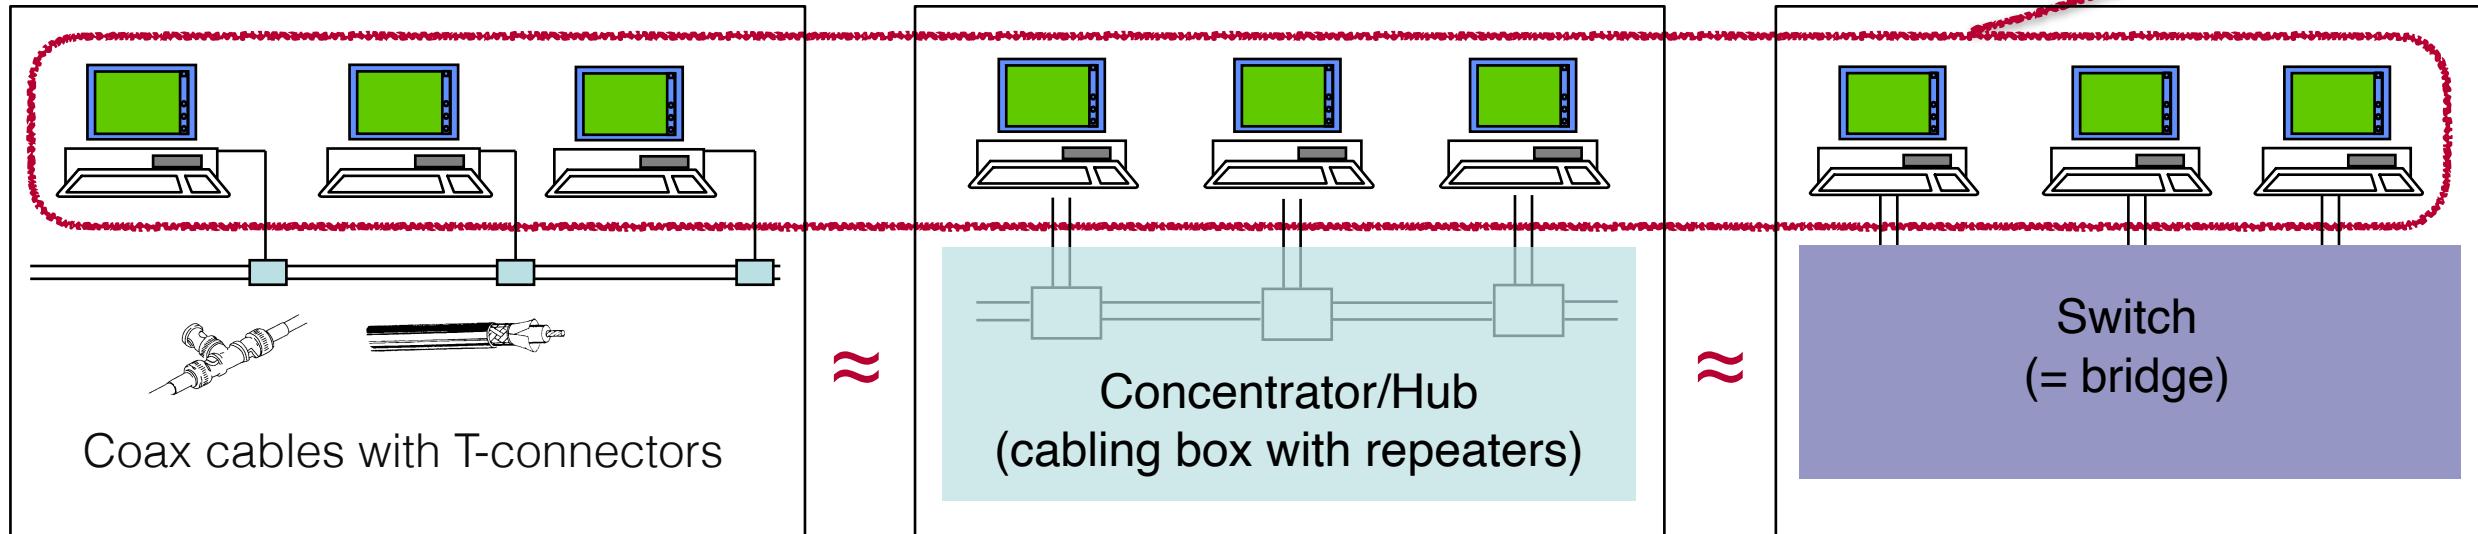
# Is switched Ethernet a different protocol?

No, switches are designed to be *transparent*:

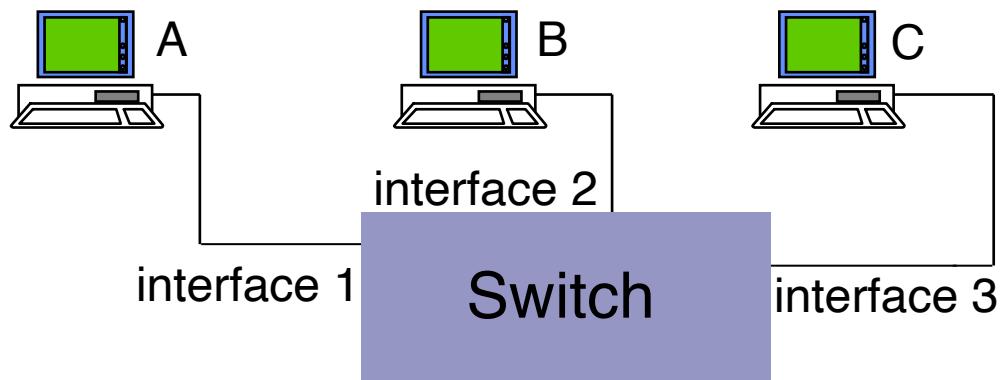
i.e. there should be *no difference* for the protocol used by end-systems when they are connected to a coax cable, concentrator/hub or a switch/bridge

But, since we can't get collisions, *CSMA/CD is disabled* by default.

These see  
no difference  
(same protocol)



# How does a switch forward a frame?



**Switch's Forwarding Table**

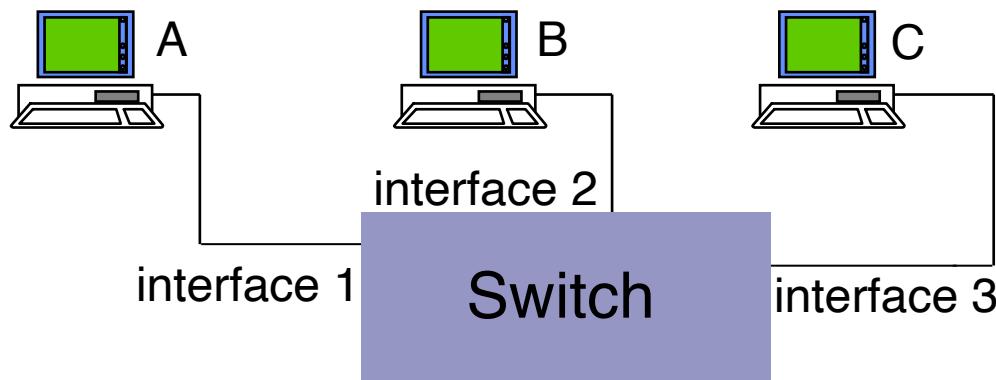
Destination MAC address	Outgoing Interface/Port
MAC of A	1
MAC of B	2
MAC of C	3

- Switches use:
  - a forwarding table with ***exact match*** which maps destination MAC addresses to outgoing interfaces (a.k.a. “ports”)
  - and a ***forwarding algorithm***
- Smart switches handle ***multicast frames differently*** [see later]

**Switch's Forwarding algorithm**

```
Observe all traffic at all ports
For each incoming frame:
  read destination MAC address in MAC header
  if entry for MAC exists in forwarding table:
    obtain matching outgoing port
    if outgoing port = incoming port:
      discard frame (no need to forward)
    else
      forward frame to outgoing port
  else (destination MAC not in the table):
    broadcast to all ports, except for the
    incoming port
```

# How does a switch populate its forwarding table?



*Switch's Forwarding Table*

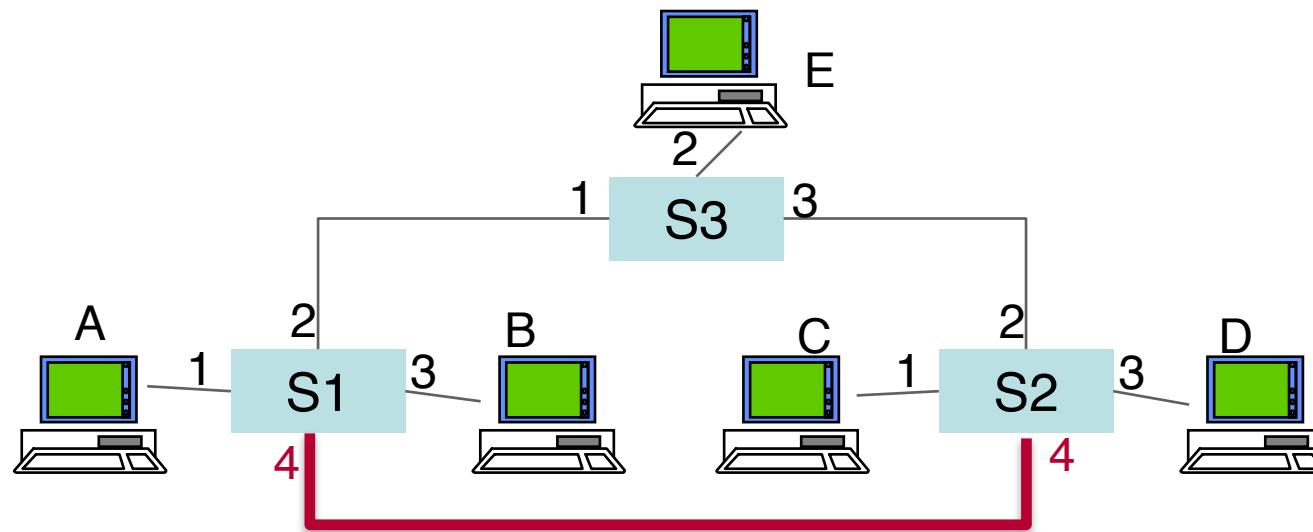
Destination MAC address	Outgoing Interface/Port
MAC of A	1
MAC of B	2
MAC of C	3

By observing traffic (*plug-and-play device*):

- table built by **self-learning** from Source MAC Address field in incoming frames
- learnt addresses time out if not re-learnt after an **ageing time**

→ **different** from IP routers that use routing algorithms

# Can self learning be extended to inter-connected switches?



*Forwarding Table at S3*

Destination MAC address	Outgoing Interface/Port
MAC of A	X 3
MAC of B	1
MAC of C	3
MAC of D	3
MAC of E	2

- 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



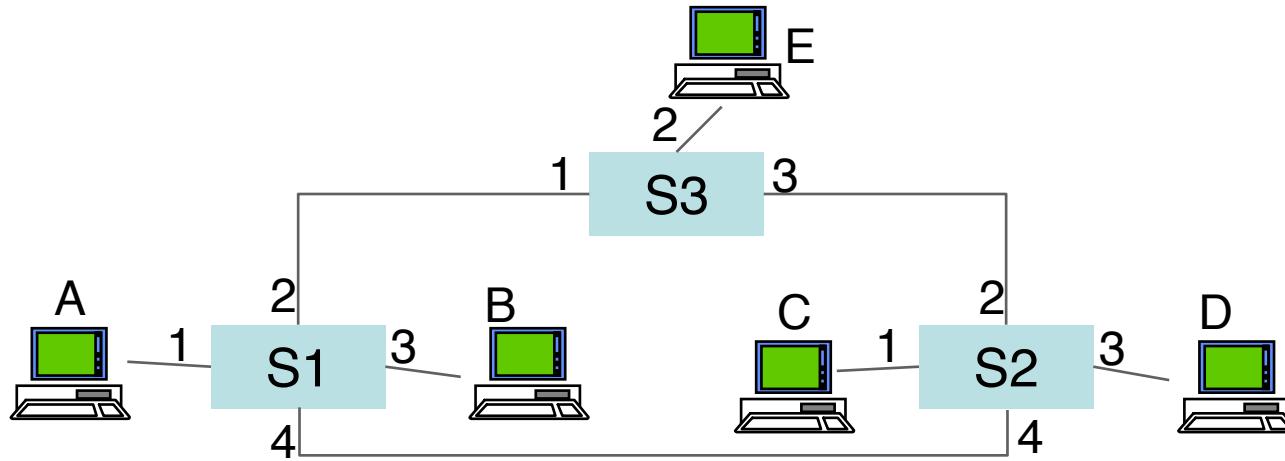
Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045

# Solution

Answer C

It does not work if there are *loops* in the topology



E.g.: here S3 does not know whether A is on port 1 or 3, as flooded/broadcast packets may arrive from either side.

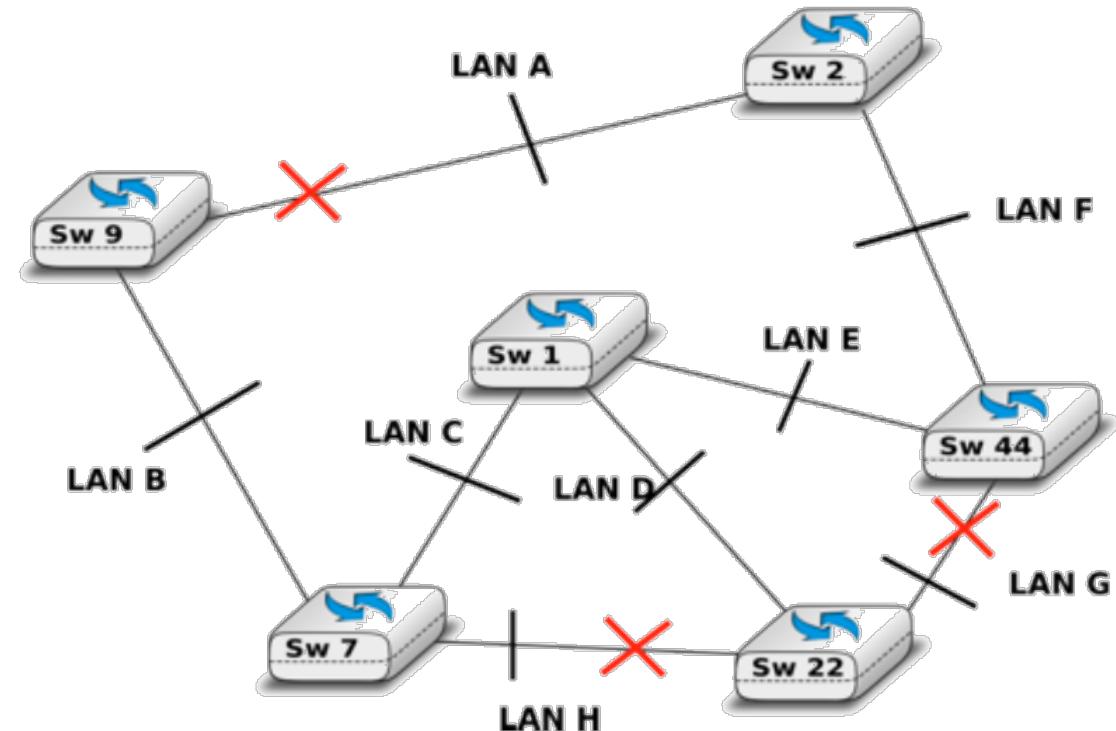
(Recall that switches are queueing systems, hence there exist non-deterministic delays)

- The entries of the forwarding table may change indefinitely and packets may be trapped in a link, as we saw when talking about TTL in IP layer

# Spanning Tree Protocol (STP): the common way to prevent loops

## *What does it do?*

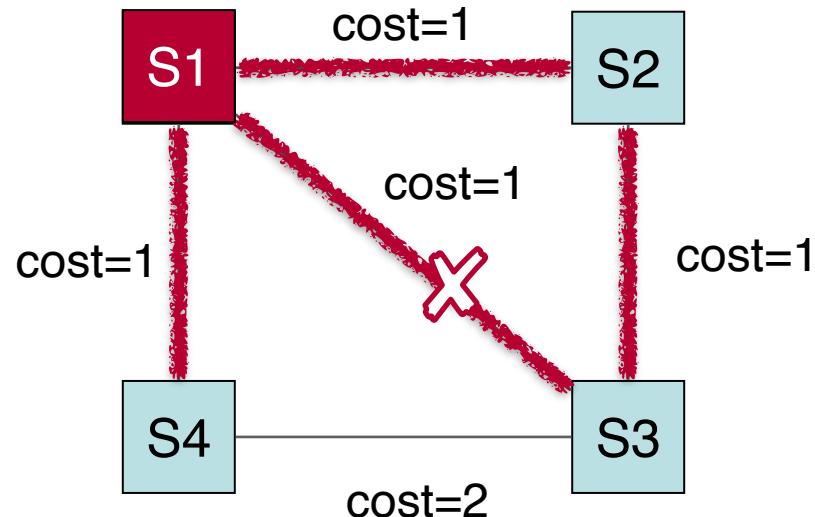
- forces the active topology to be a *tree* (i.e. loop-free) that spans all switches by *deactivating* some links
- adapts to link failures and additions whenever needed



# How does the Spanning Tree Protocol work?

## The switches:

1. elect a **root** switch = switch with smallest (configurable) label ID
2. activate only the links that are along the shortest path to root switch  
spanning tree = set of **shortest paths** to root switch
3. monitor whether the root is reachable, and  
if not, they trigger a re-computation of a new spanning tree



# A few important points

- STP is a *distributed* algorithm
  - no central control,
  - everything is *plug and play*
  - a *gossip* protocol is used to learn the IDs
- Shortest paths = paths incurring the smallest aggregate cost
  - each link between switches has a (configurable) *cost*,
  - typically, the cost is a decreasing function of bit rate
- Switches compute the shortest paths using variant of *Bellman-Ford* algorithm

[Details are a bit more complicated when links between switches are shared medium and not point to point: see [textbook](#).]

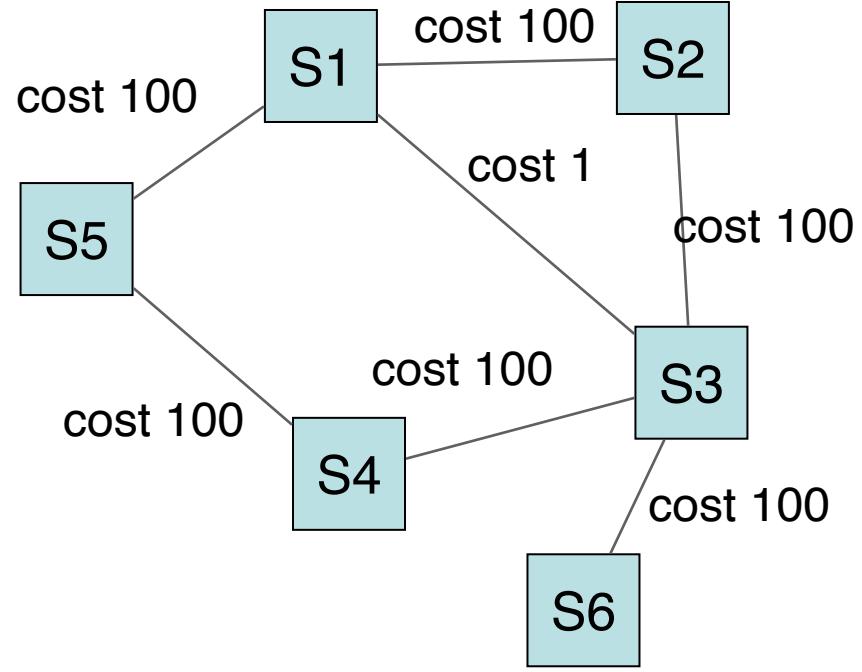
# Which of the following links are on the spanning tree ?

- A. S4—S5
- B. S4—S3
- C. Both
- D. None
- E. I don't know



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045



# Solution

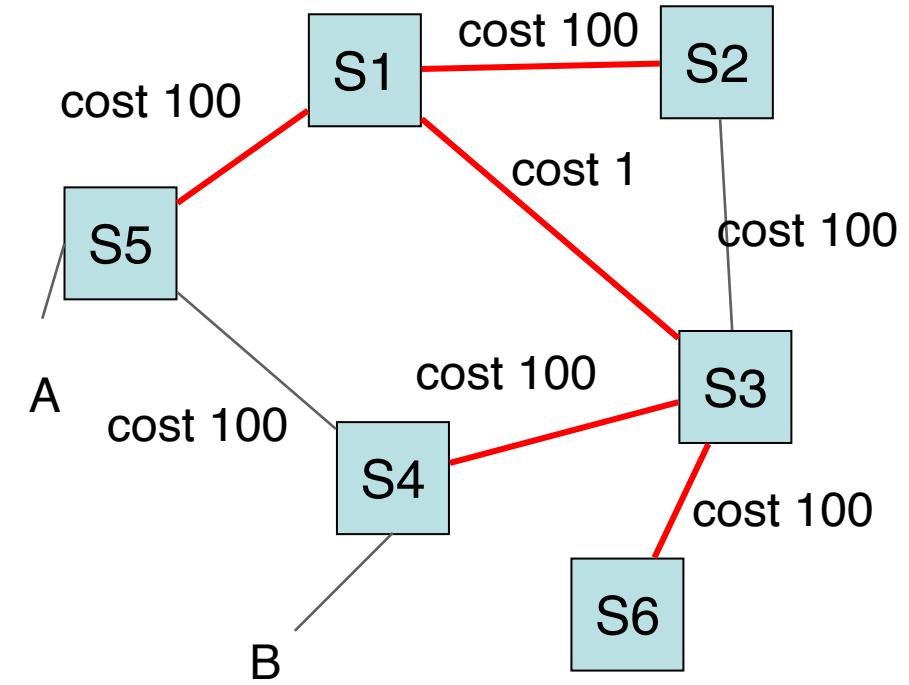
Answer B

The root is S1 (smallest label)

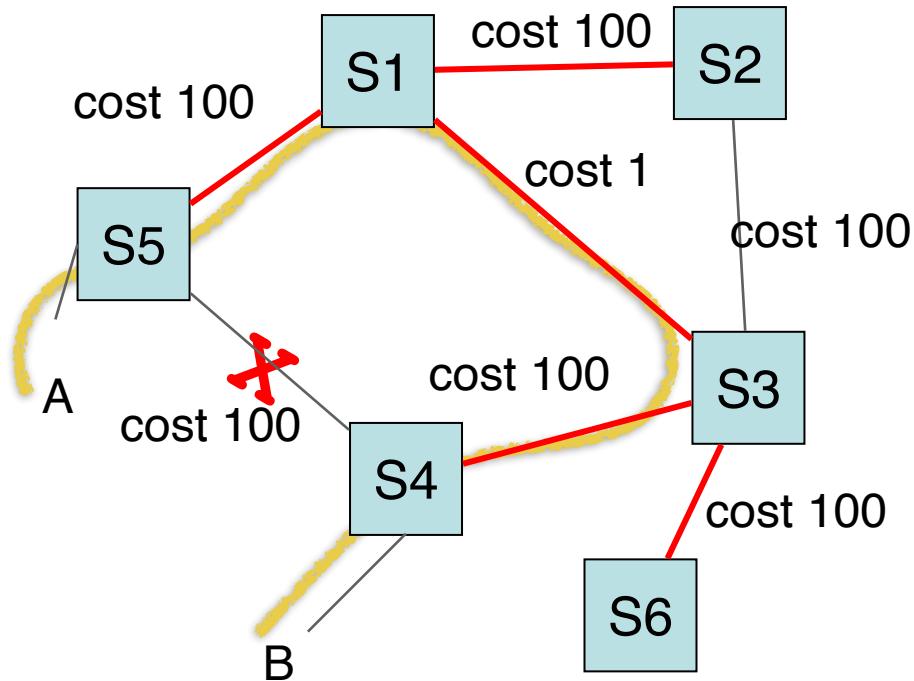
The spanning tree is shown in red →

The path for packets from A to B is:

S5—S1—S3—S4



# A side-effect of the Spanning Tree Protocol (STP)



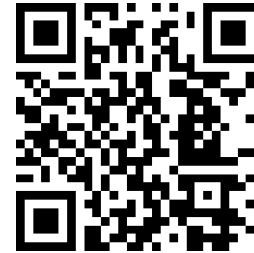
- All frames go through the spanning tree
- A direct link between two (non-root) switches may not be used even if it is optimal
  - e.g., the path from A to B, here, is not optimal

- ▶ **Less efficient** than shortest path
- ▶ some more sophisticated switches implement *Shortest Path Bridging* instead of STP, which is the default
  - [see “link state” lecture]

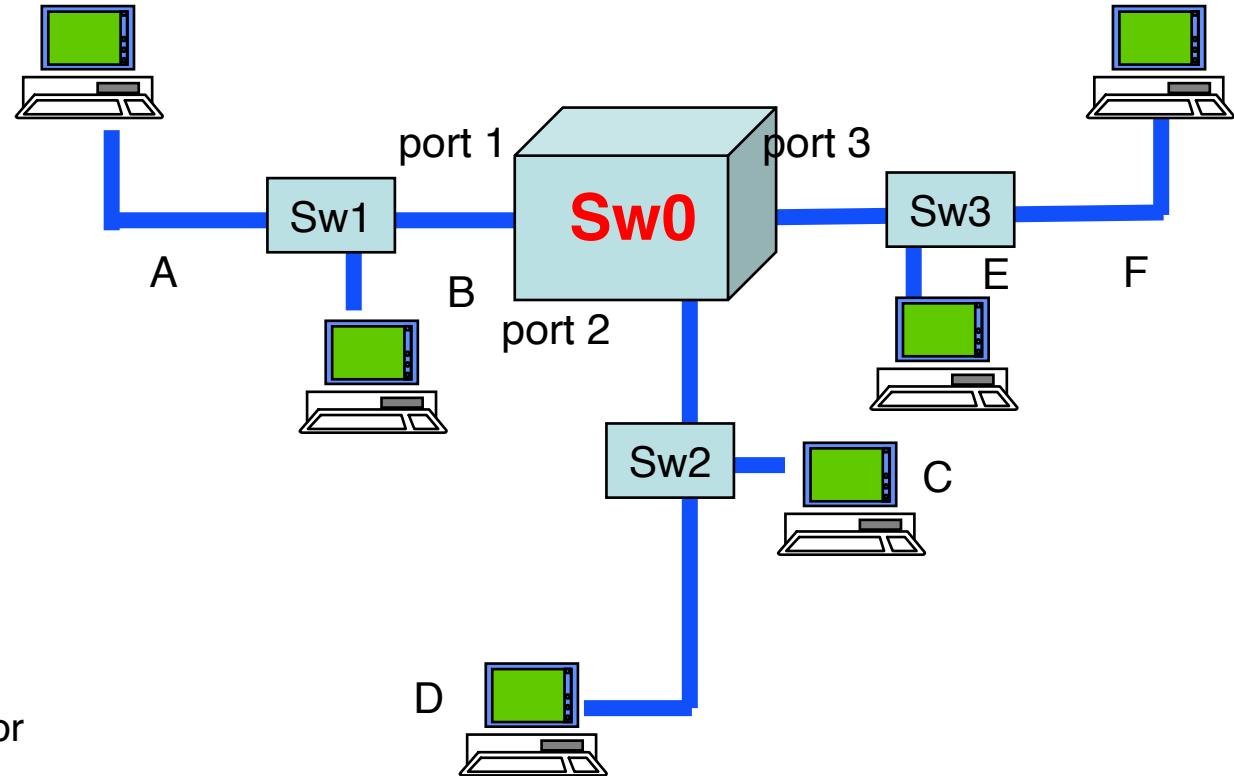
How many frames can be transmitted in parallel in this switched Ethernet LAN (all boxes are switches)?

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



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045

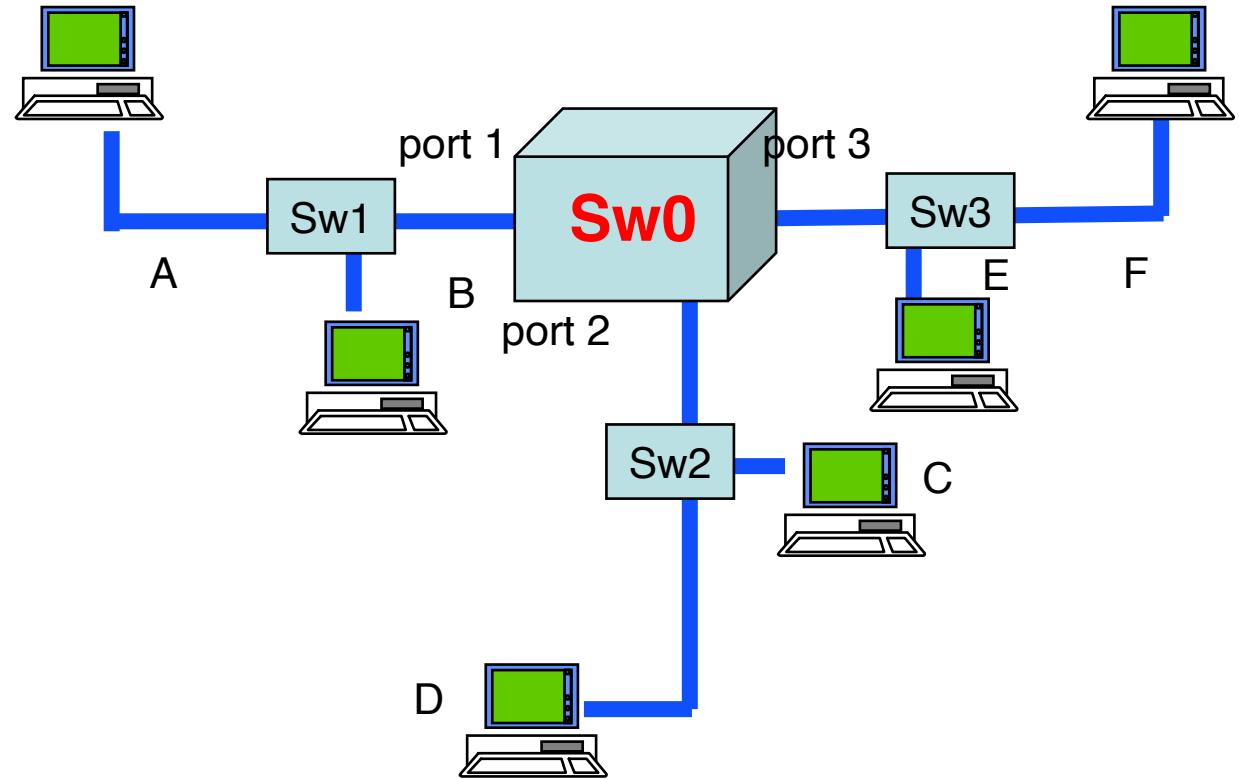


# Solution

Answer D

Assume that switches can transmit on all ports in parallel (one LAN adapter per port).

Since switches are *full duplex*, there can be up to 2 frame transmissions on every link (one in each direction), therefore up to 18 transmissions in parallel.



→ *This is the benefit of today's switches over shared-medium, old-style Ethernet*

# 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 packets 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://youtu.be/ERFohfN\\_H40](https://youtu.be/ERFohfN_H40)

### 3. Ethernet Frame format

- Ethernet frame = Ethernet's Protocol Data Unit (PDU)
- An Ethernet frame typically carries an IP packet in its payload, sometimes also other:

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

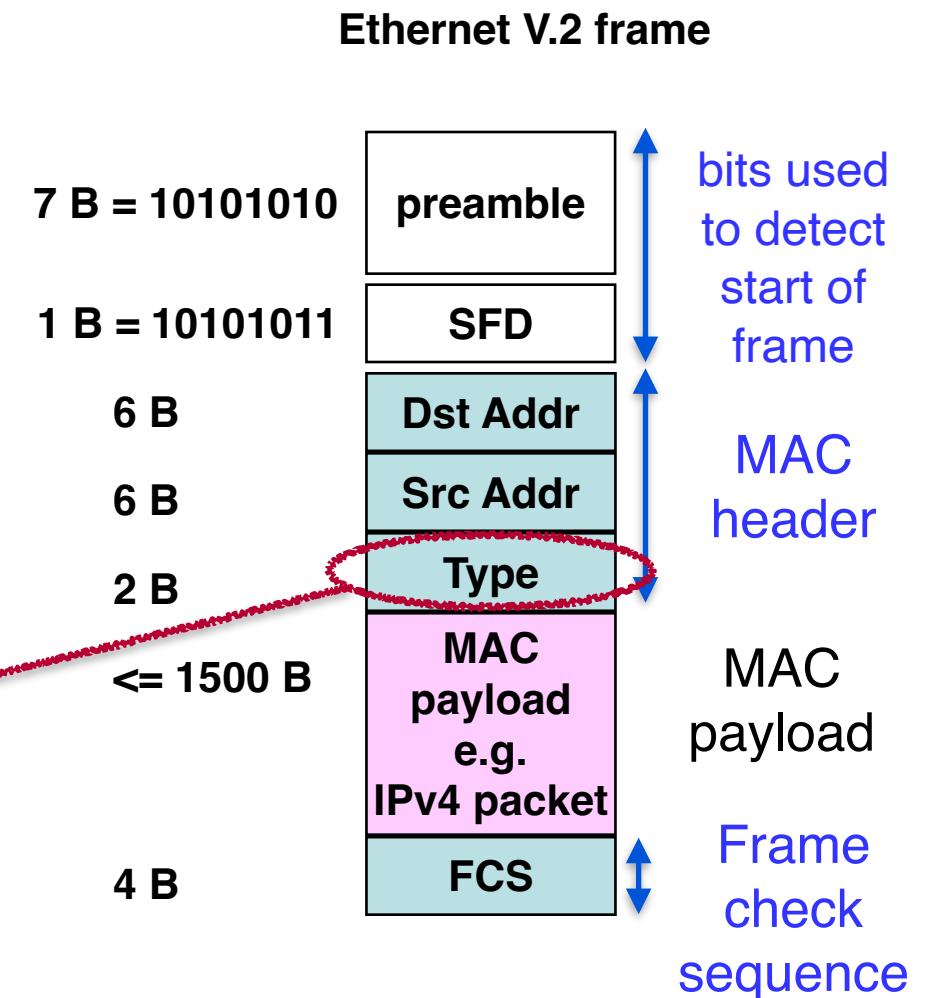
0800: IPv4

0806: ARP (used by IPv4)

86dd: IPv6

88e5: MACSEC (authenticated/encrypted frame)

88f7: Precision Time Protocol



# Ethernet Frame format: more details

The **preamble** is used for the receivers to synchronize their clocks. It is a bit sequence of 7 bytes, where each byte is equal to: 10101010... and it is terminated by 0.

Following the preamble, an additional byte is used to indicate the beginning of the frame. This is called SFD (=start frame delimiter) and it is similar to a preamble byte, but its last two bits are 11 (instead of 10).

In Ethernet, transmission starts **asynchronously** (hosts transmit independently), and between transmissions, the channel is idle. This is why the clocks of the sender and the receiver have to get synchronized at the beginning of each frame, so that the receiver can correctly “read” the received data.

There is no frame length field in the Ethernet header—it is up to the higher layer using Ethernet to find out the length of the payload.

Frames have to be at least 64B (to see why, refer to earlier slide about minimum frame size) and at most 1500B. If needed, padding is inserted to make the frame long enough, and only in that case the Type field is used to indicate the length of the size of the payload before padding.

The format shown of the previous slide 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 via Frame Check Sequence (FCS)

## Why ?

- The channel may cause bit errors during transmission => a frame may be “*corrupted*”.
- We want to detect and discard corrupted frames.

## How ?

- The sender computes a Cyclic Redundancy Checksum (CRC, 32 bits) of the frame and inserts it in the FCS field.
- The receiver recomputes the CRC checksum and checks if it is equal to the FCS.
- CRC can ***detect*** all single, double, triple errors, all error bursts of length  $\leq 32$ , most double bursts of length  $\leq 17$ .
  - The probability that a random collection of bit errors goes undetected is  $\approx 2e-32$ .

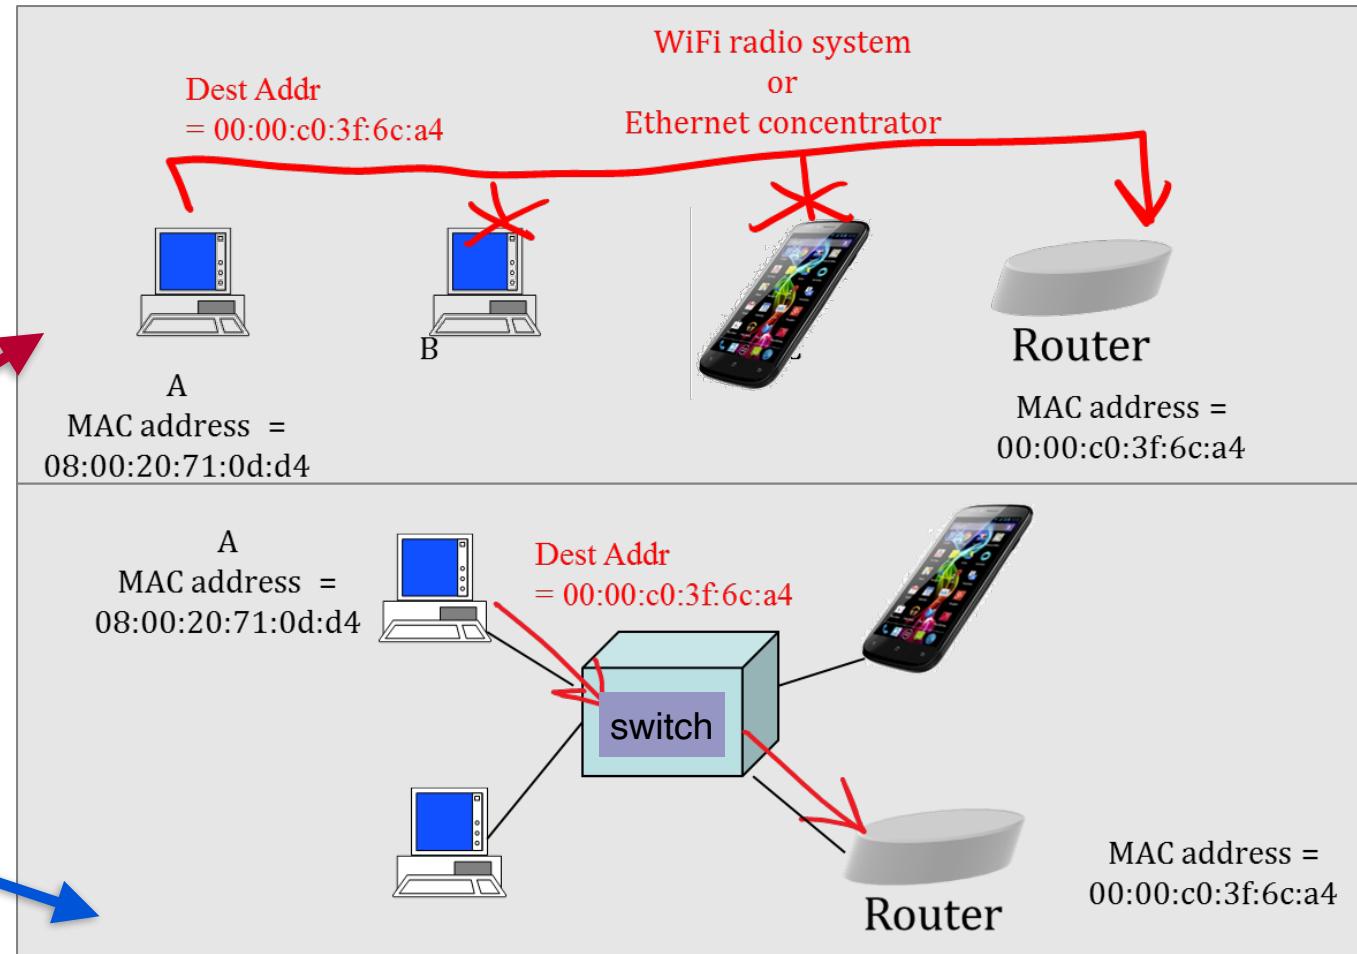
\* *Error recovery* in Ethernet is simple: If a frame fails the check, it is *dropped*

# MAC Unicast

- MAC address: 48 bits = adapter ID  
Unique worldwide (in principle, but can be also configured sometimes)

How?

- Sender puts destination MAC address in the frame
- Address sent in the clear, no encryption
- In *shared-medium LAN*: all stations read *all* frames, but keep only if destination address matches
- In *switched Ethernet*: switches forward frames *only* to ports that need it (in principle)



# Addressing: more details (for the interested students)

Ethernet addresses are known as MAC addresses. Every Ethernet interface has its own MAC address, which is in fact similar to the serial number of the adapter, set 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 (which is the usual case, and also called a “burnt-in MAC address”) 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.

The data is transmitted in Ethernet with the least significant bit of first octet first—this is a strange behavior 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);

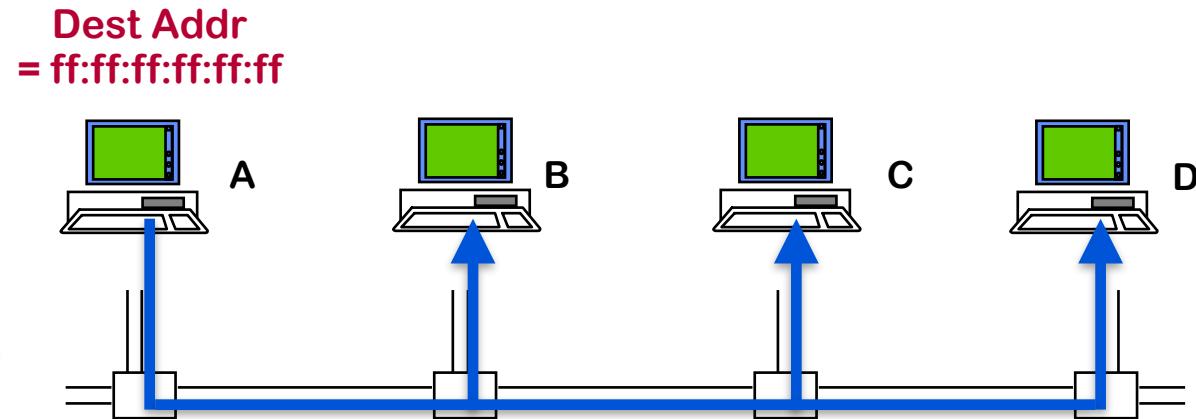
The first 24 bits are the Organization Unique Identifier (OUI); the remaining 24 bits are allocated by organization that owns the OUI.

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

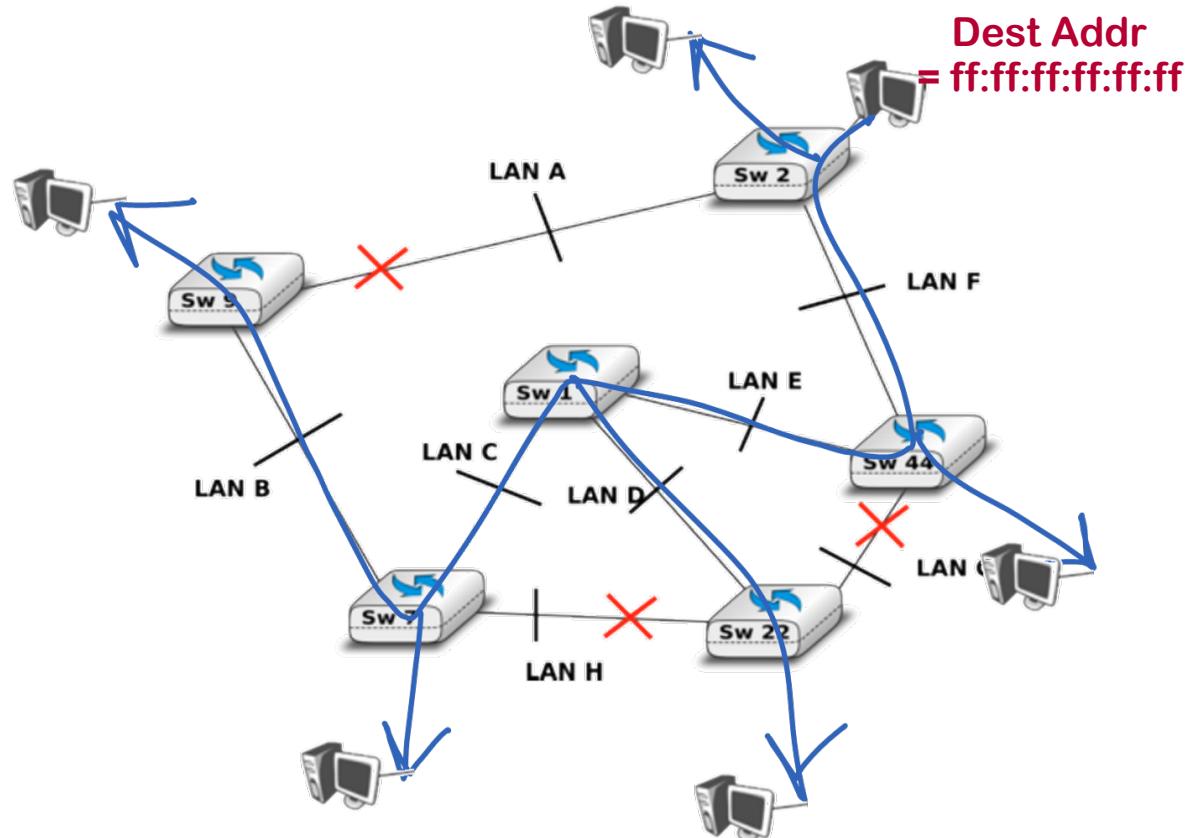
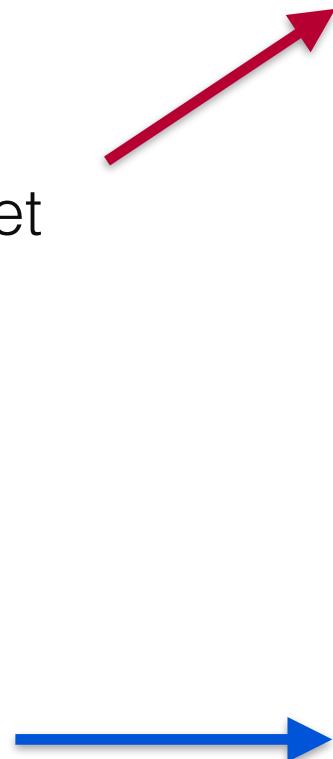
# MAC Broadcast

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



- In *shared-medium* LAN:  
all machines receive packet  
and do *not* discard it

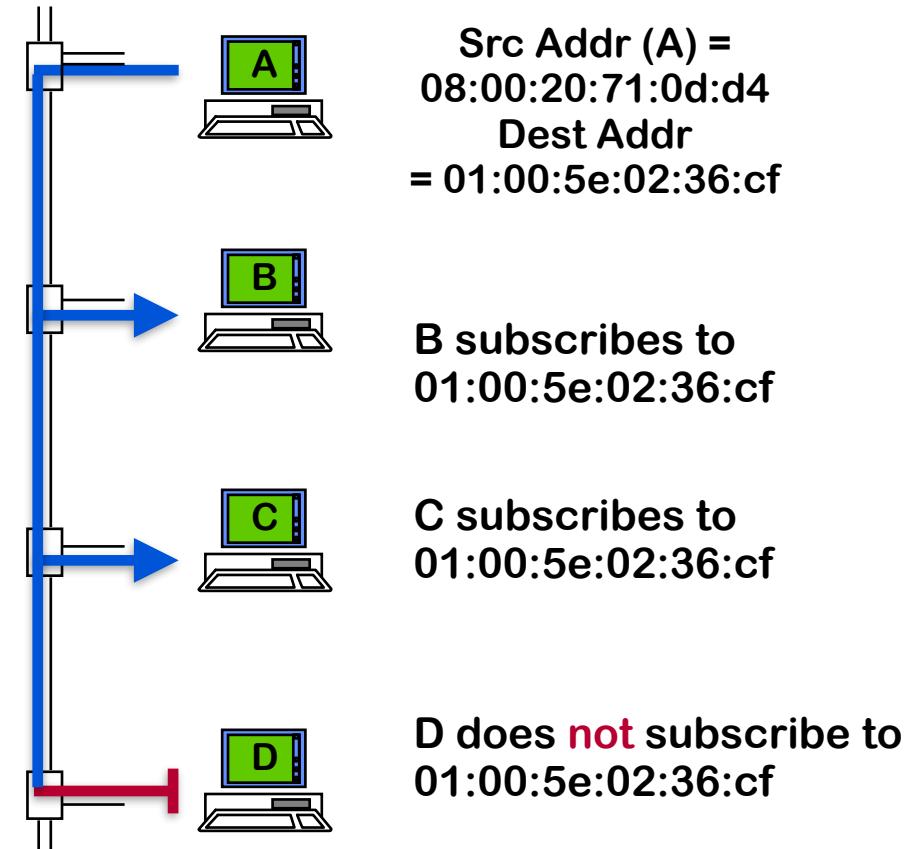
- In *switched Ethernet*:  
broadcast frames are sent  
on all nodes and ports  
*on the spanning tree*



# MAC Multicast (= sending to multiple hosts)

- uses:
  - MAC Addresses with 8th bit == 1 are multicast addresses
  - various multicast addresses to support various protocols (e.g. IPv4, IPv6, PTP)
- makes sense only in *switched Ethernet*:
  - Ethernet adapter discards multicast packets, unless host subscribes to address
  - non-smart switches broadcast such frames
  - smart switches send only to relevant nodes (the ones that have subscribed to the multicast address using IPv4 or IPv6 subscription protocols—see “multicast” lecture)

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



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



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

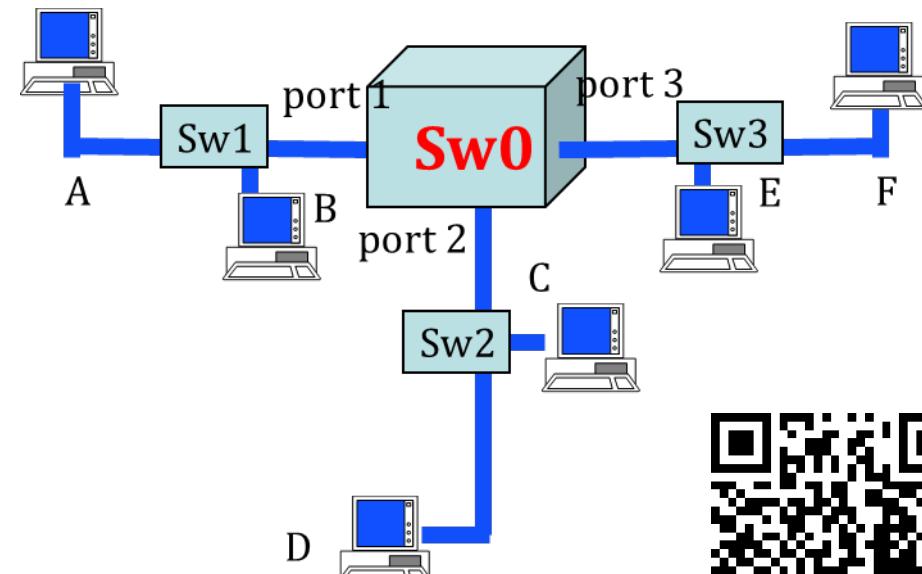
Join room  
46045

# Solution

Answer B: the Ethertype indicates an IPv4 or IPv6 packet

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

- A. No, since switches 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



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045

# Solution

Answer A: the switches work only at MAC layer and continue to work the same whether IPv4, IPv6 (or something else) is used.

One *exception*: “smart” switches are cheating—they build multicast trees by looking at IP layer group membership (see multicast lecture later)

# 4. Virtual LANs

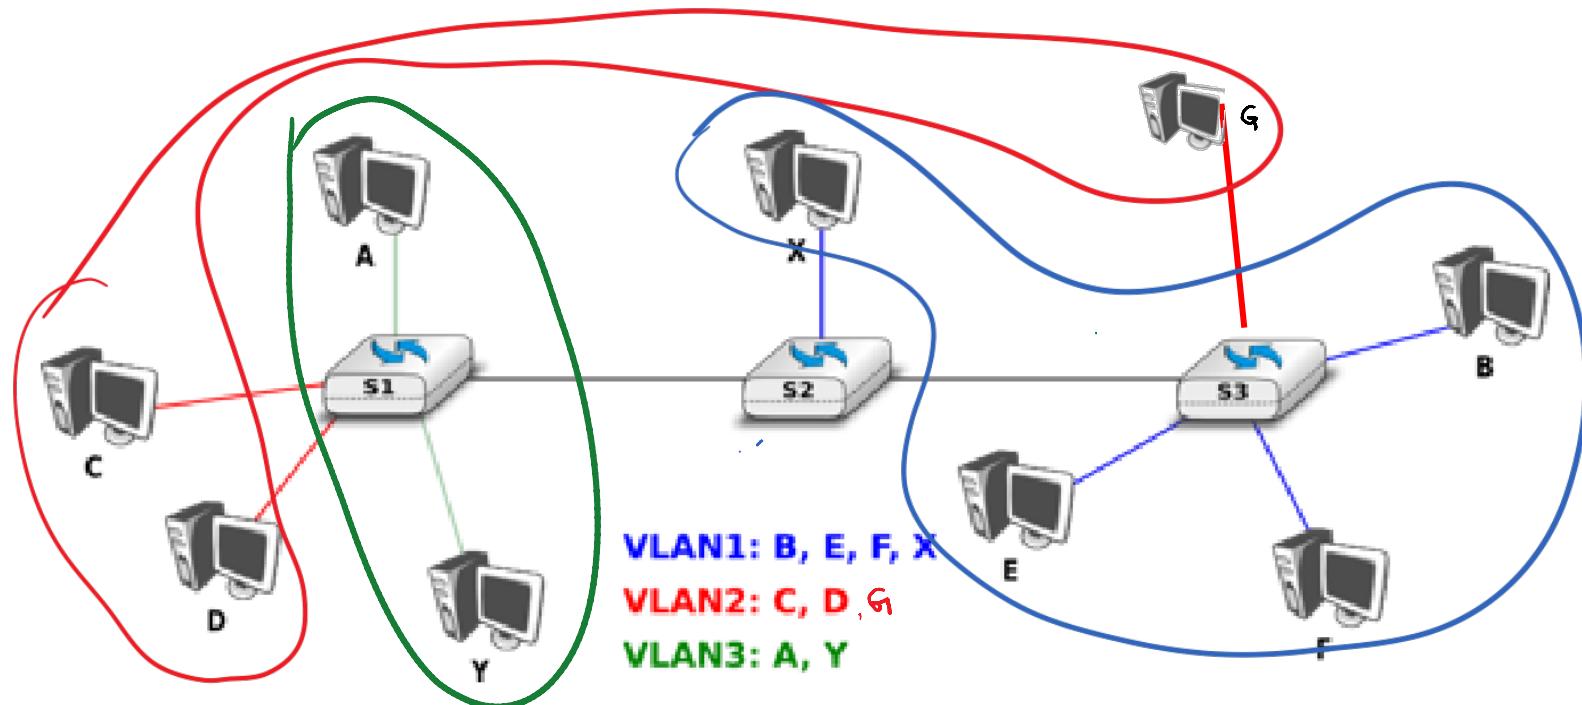
## *Why invented?*

Lack of traffic isolation —> we want to limit the scope of broadcast traffic and have access control

Inefficient use of switches —> we do not want to buy extra switches if others have available ports

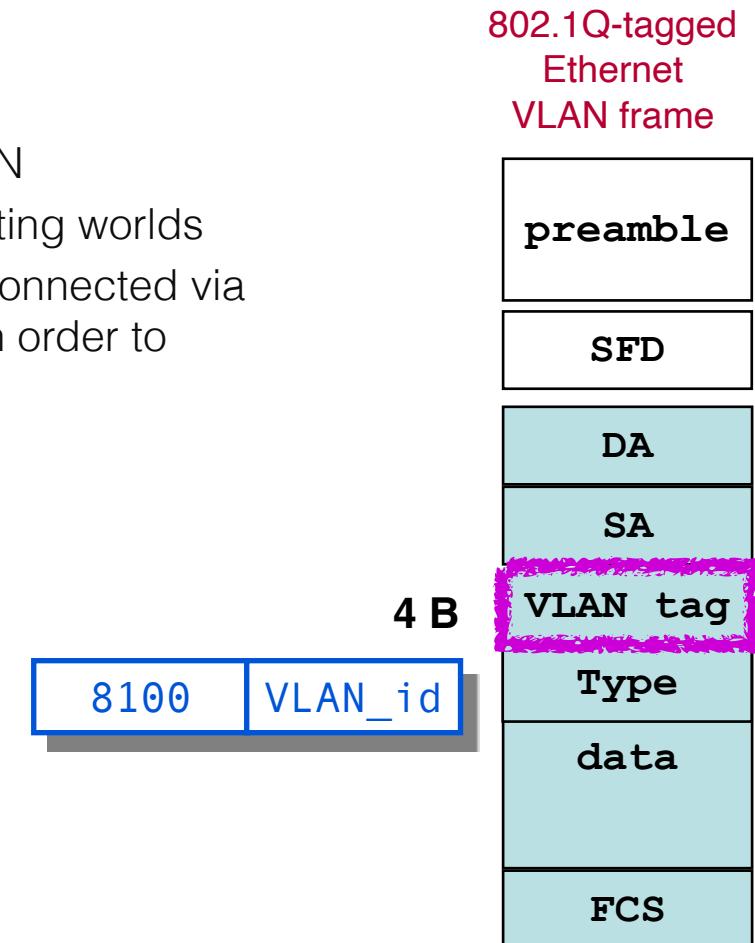
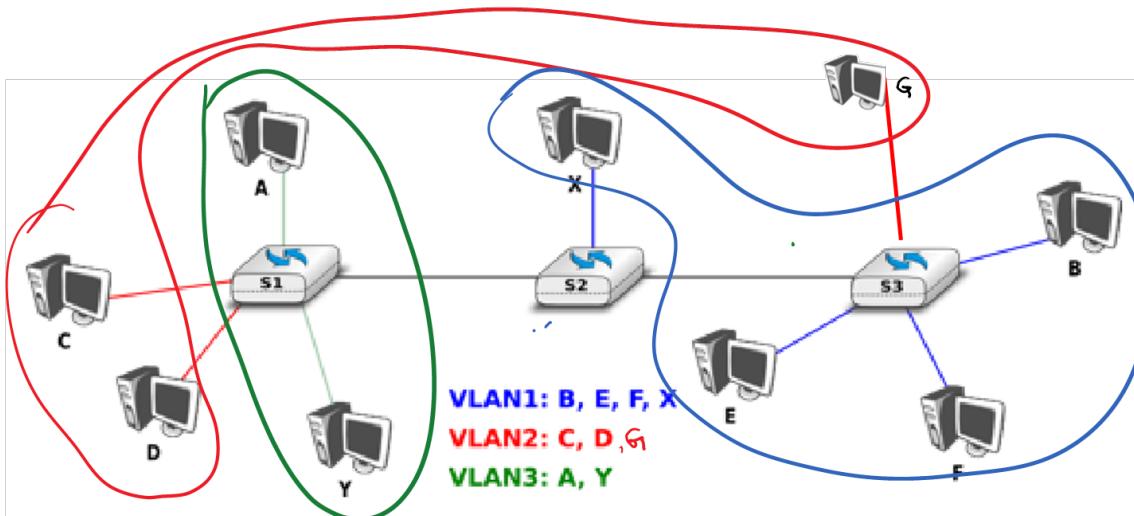
Managing users —> if a user changes LAN, we do not want to also change its physical cabling

**Goal:** Decouple who belongs to which LAN from their physical location

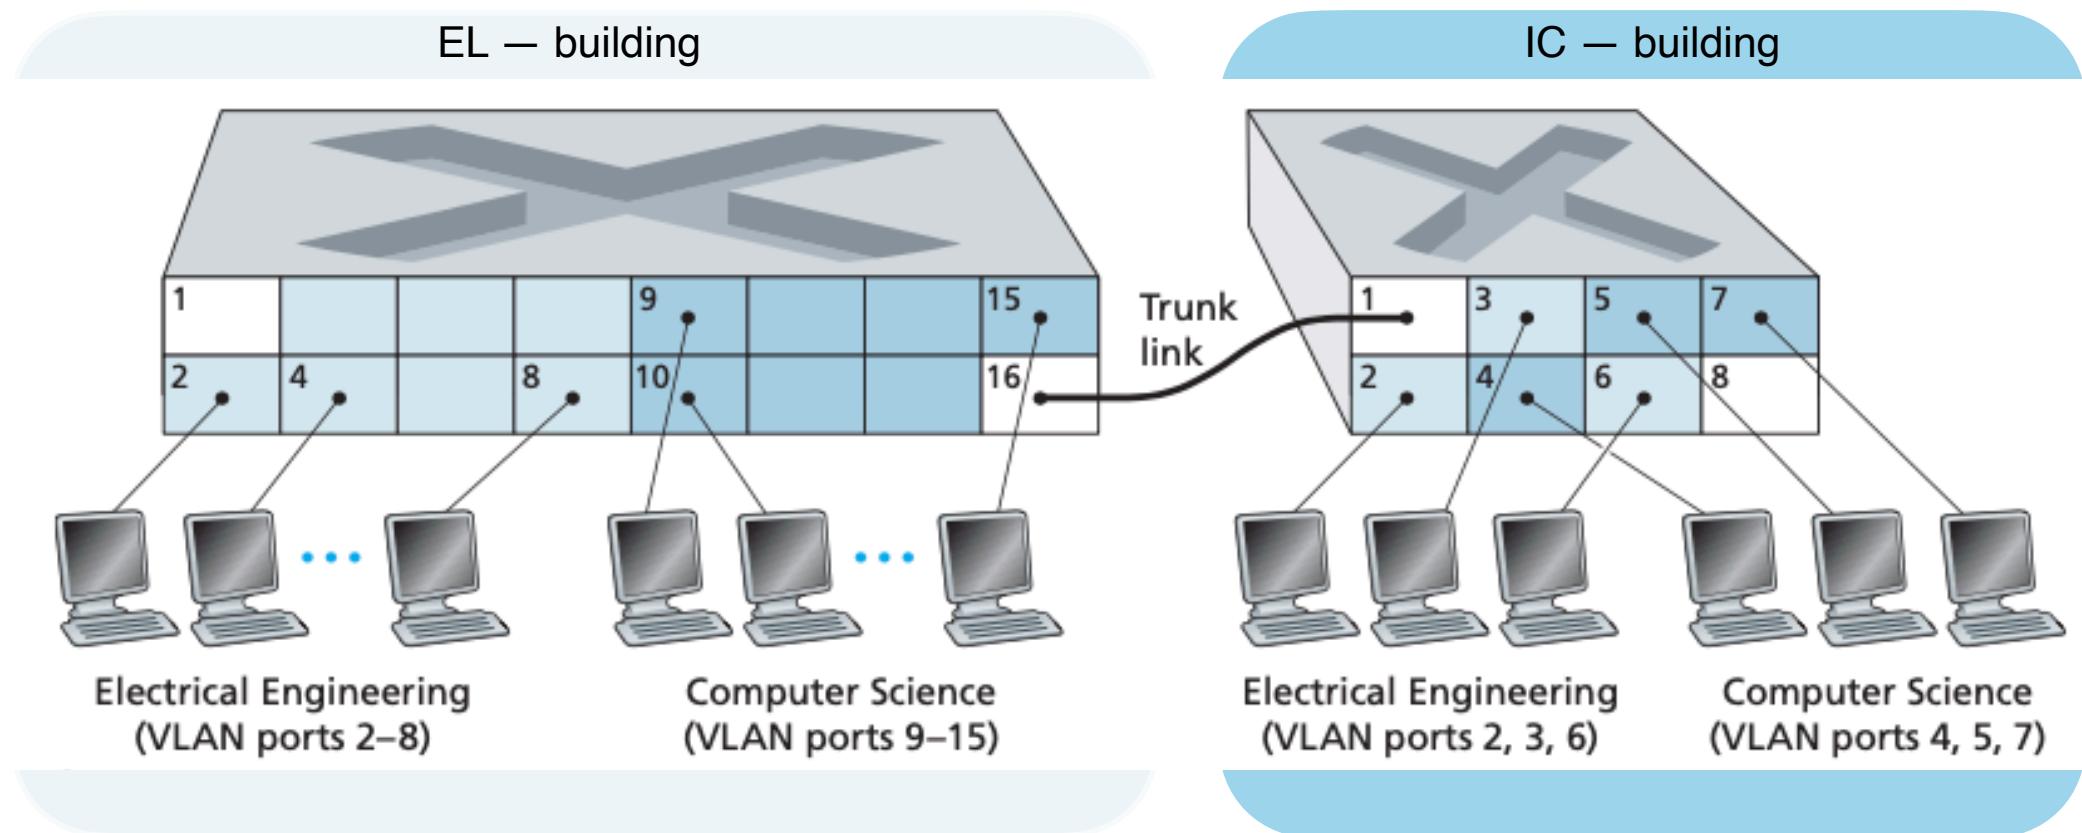


# Virtual LANs

- VLANs are handled as if they were physically separate: different forwarding tables, different spanning trees.
- *How* does it work ?
  - network manager configures which switch port belongs to which VLAN
  - switch handles ports of different VLANs as separate, non communicating worlds
  - if more than one switches implement the same VLANs, they are interconnected via **trunk** ports, which use VLAN tags in an Ethernet's **extended header**, in order to correctly associate frames with VLANs (802.1Q-tagged Ethernet)



# An example of a trunk port



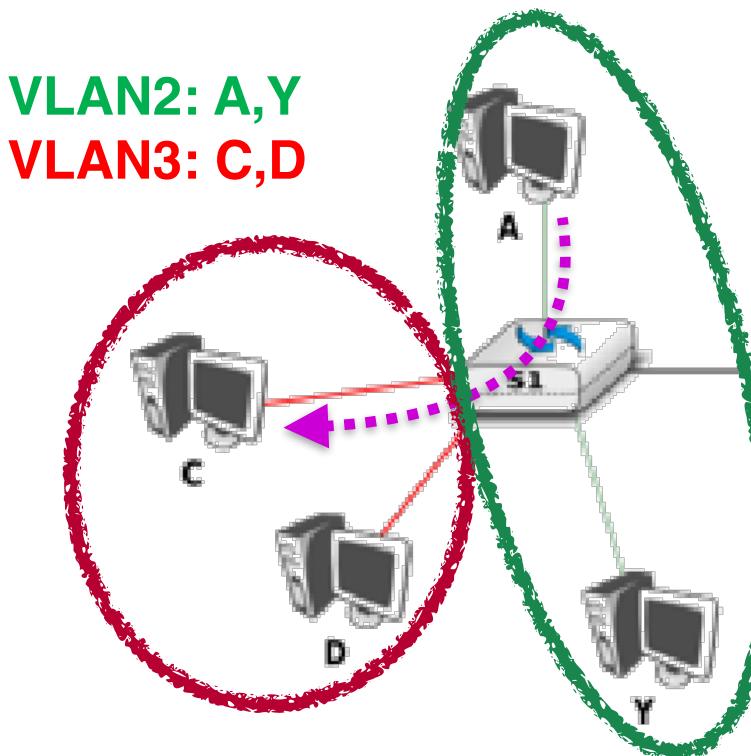
- The switches implement 2 virtual LANs
  - the Electrical Engineering LAN in light blue and
  - the Computer Science LAN in blue color
- The switches are interconnected via trunk ports (16 and 1) that belong to both LANs
- The trunk ports use VLAN tags to notify each other about which arriving frame belongs to which VLAN

# 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 the 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 the packet to this MAC address and to this VLAN
- C. None of the above
- D. I don't know



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app



Join room  
46045

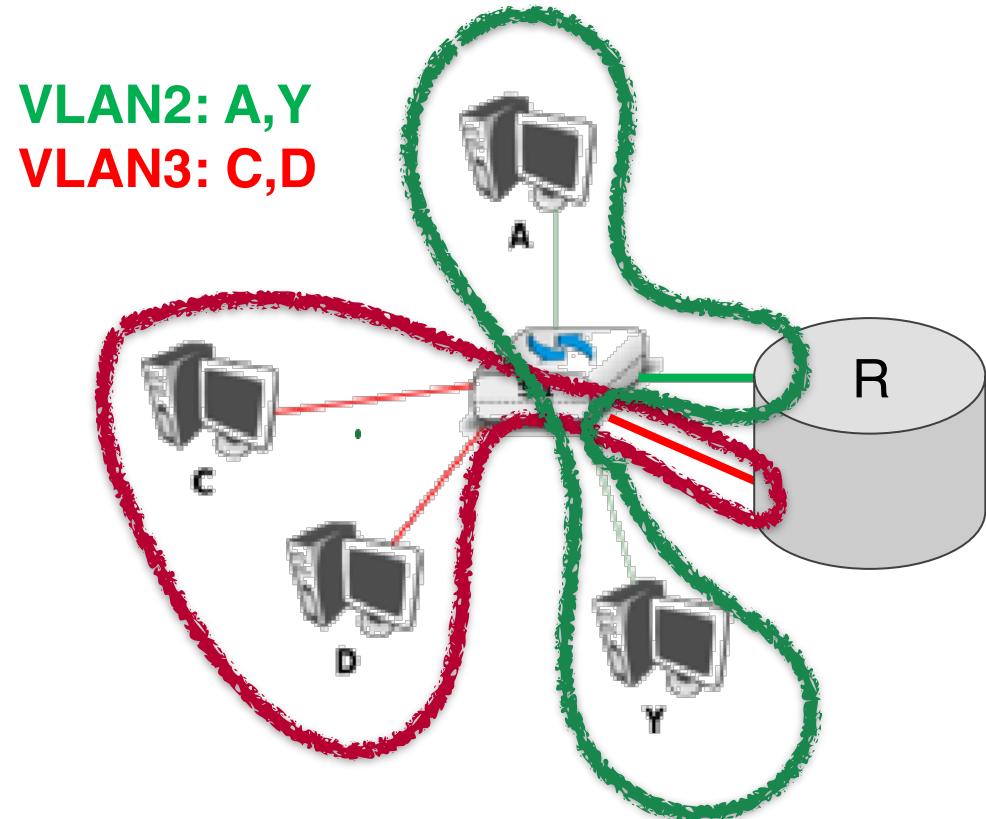
# Solution

Answer C

VLANs 2 and 3 are two *separate* LANs, it is not possible for A to send directly to C—it must go through a router.

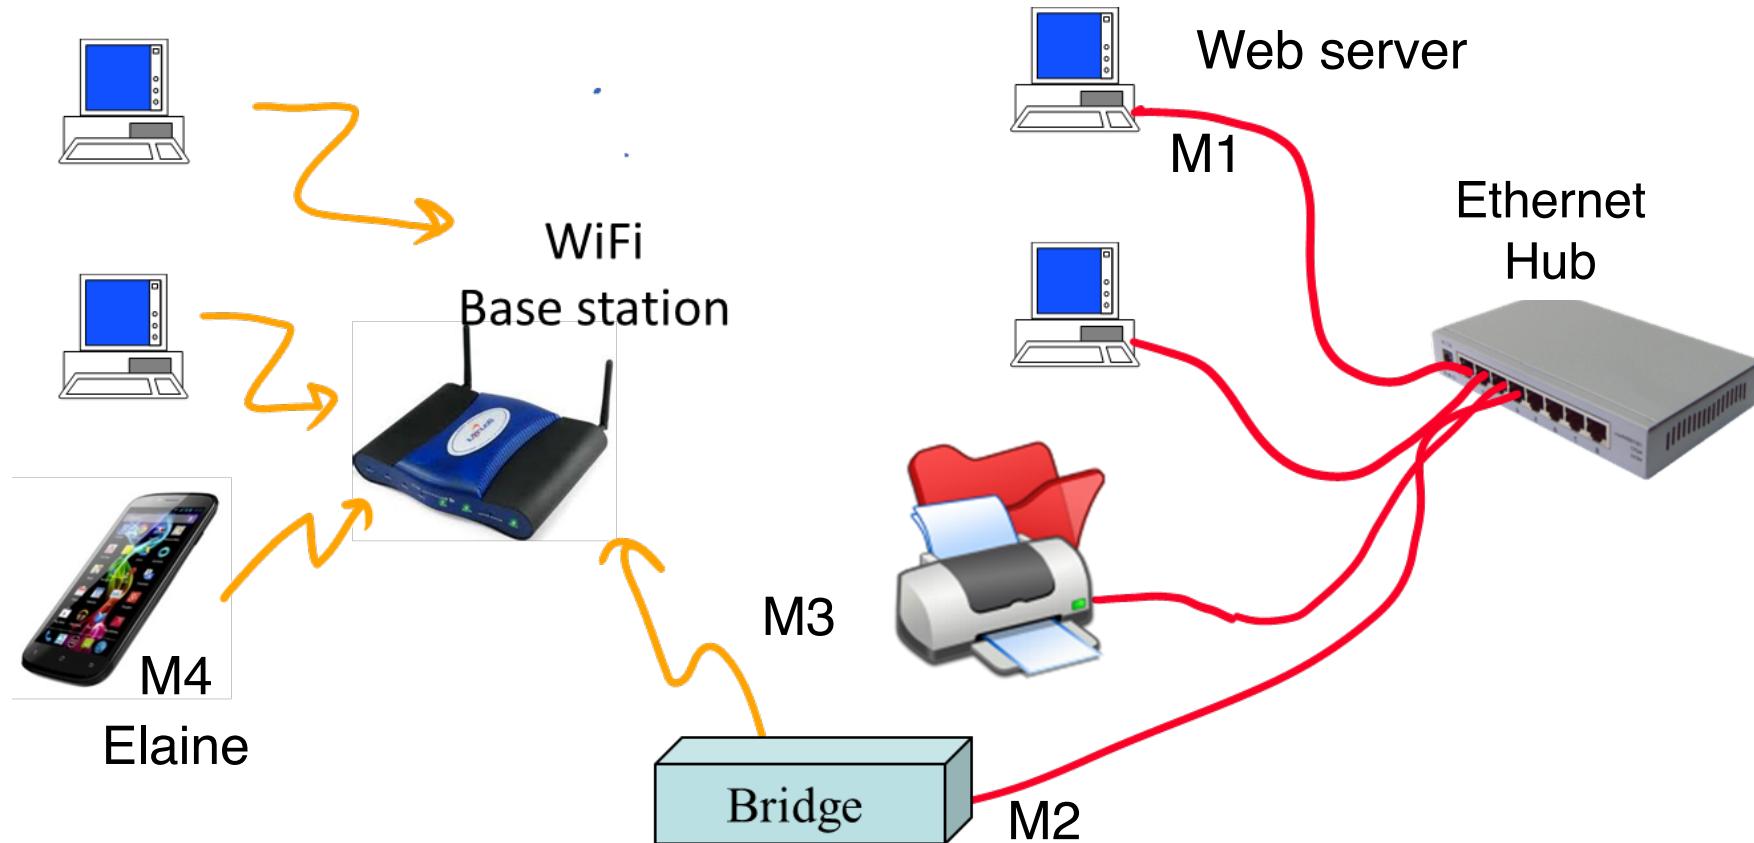
A possible workaround:

- R belongs to both VLAN2 and VLAN 3 (and has several logical interfaces)
- A sends an ARP request to obtain R's MAC address and sends IP packet to R
- R receives the packet on its green VLAN interface, and according to its forwarding table, it has to forward the packet to the red VLAN. To do so:
  - either **R and S1 are in the same box** and R has direct access to all VLAN ports,
  - or **R has a trunk port** (i.e. a port on which VLAN tags are present) and it is connected with S1 via a trunk link.



## 5. Switches can be used between different medium access technologies (e.g. WiFi with Ethernet)

In this case, they are called mostly *bridges*

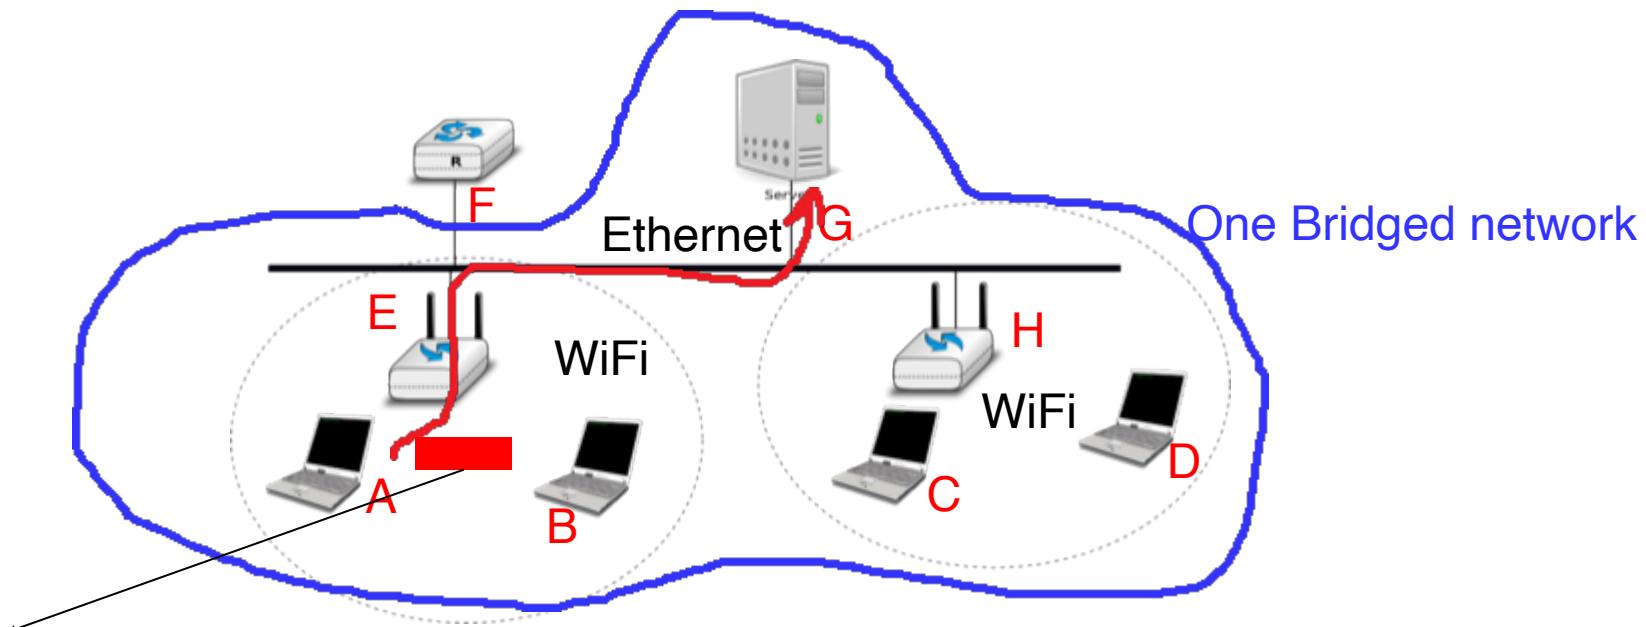


# WiFi Access Point (a typical bridge)

A WiFi access point is a combination of:

- a WiFi base station (implements the WiFi MAC protocol)
- a switch with an Ethernet backend (called the **Distribution System**)

In this case, WiFi frames differ from Ethernet frames, as they may include **3 MAC addresses**: source, destination and access-point.

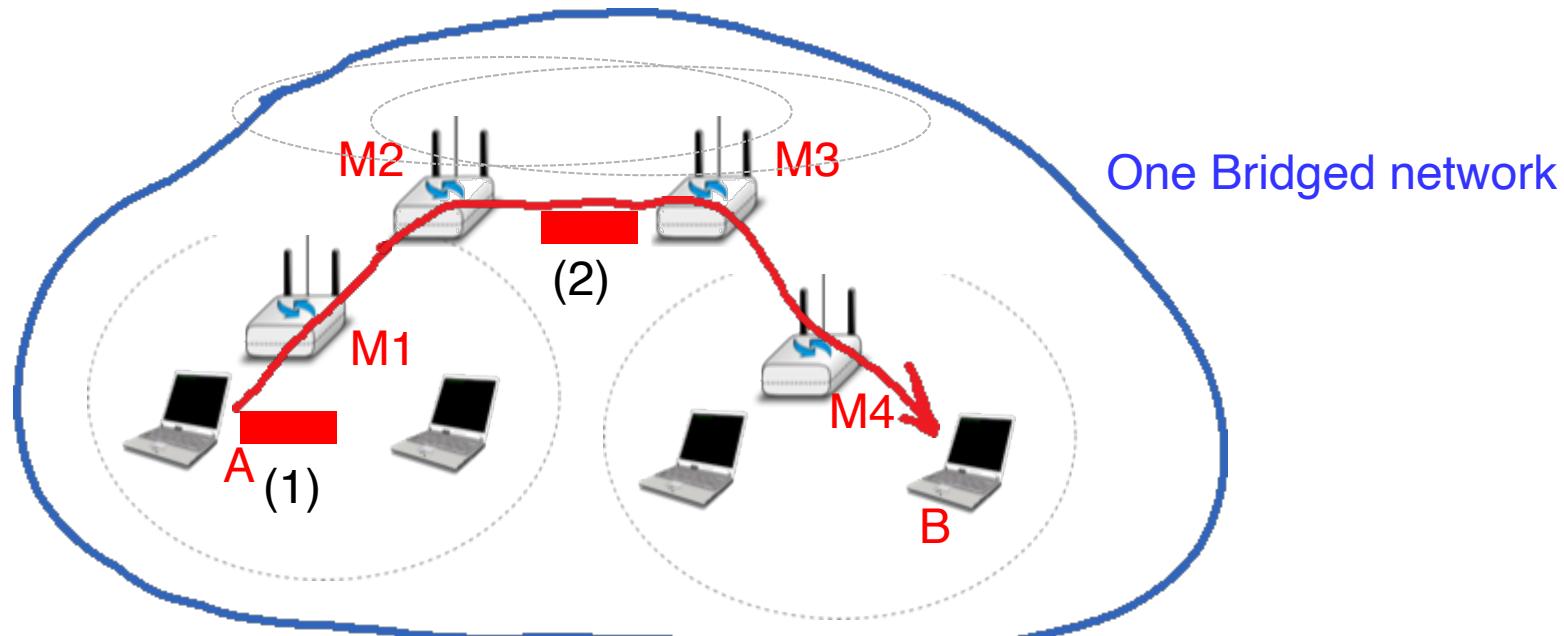


In frame sent by A to G: source MAC addr = A, dest MAC addr = G, access-point address = E

# WiFi Access Point with Wireless Distribution System

WiFi backend instead of Ethernet:

- interconnection between WiFi base stations is done using WiFi relays, which act as switches
- WiFi frames carry **4 MAC addresses**: source, destination, MACs of receiver/transmitter relays



In frame sent by A to B:

At (1) : source MAC addr = A, dest MAC addr = B, access-point address = M1

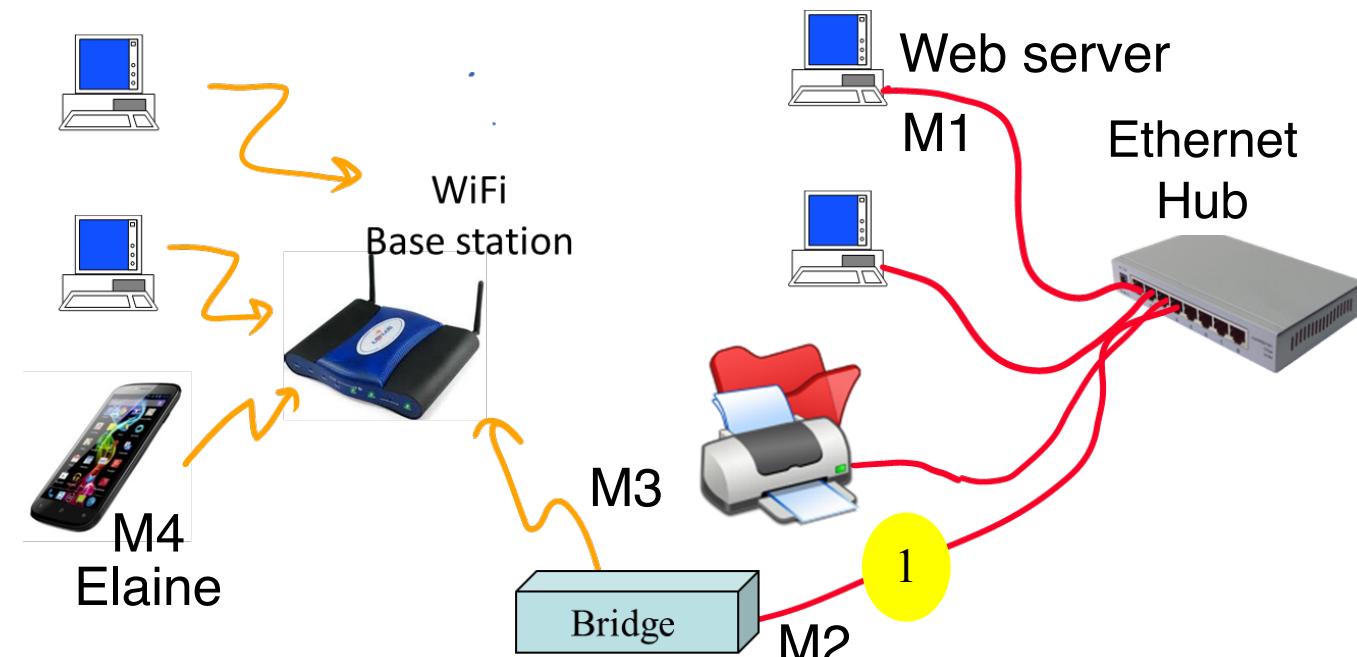
At (2) : source MAC addr = A, dest MAC addr = B, tx-access-point address = M2, rx-access-point address=M3

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



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app



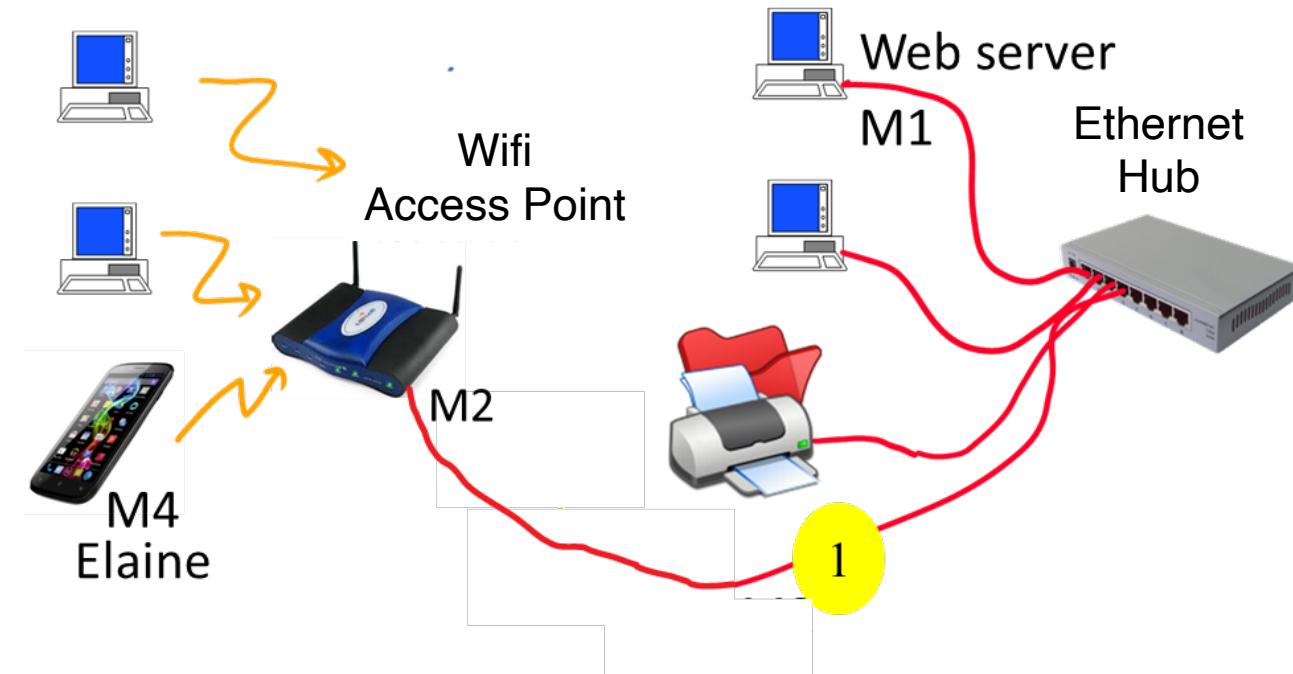
Join room  
46045

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

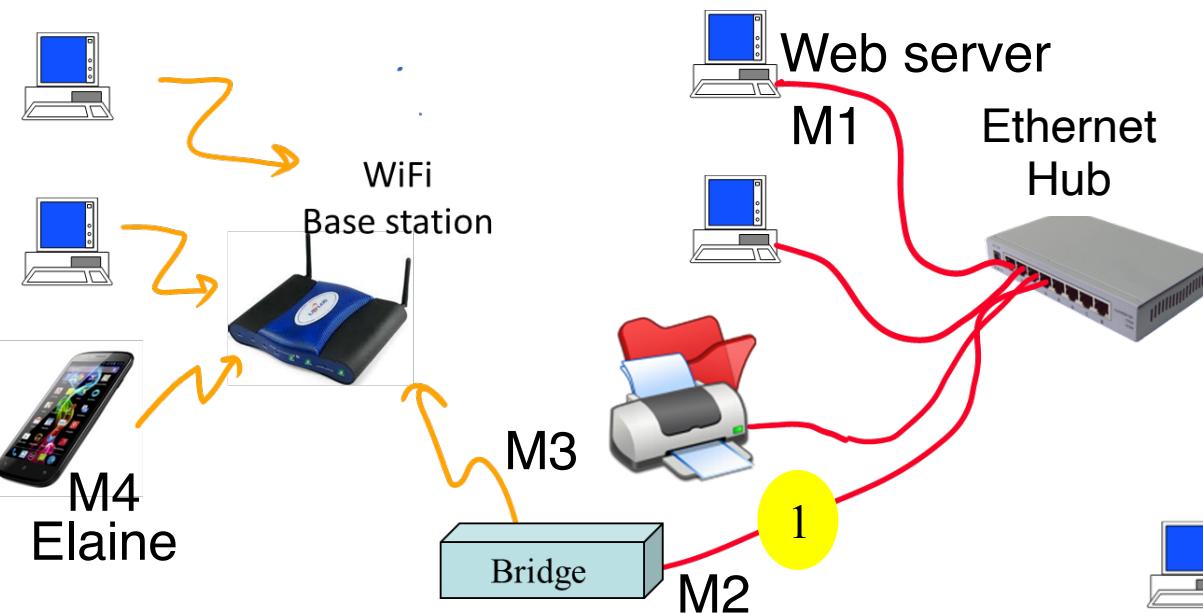


Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

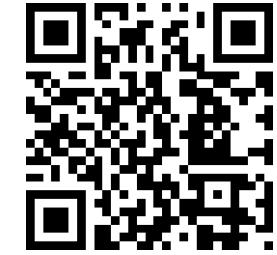
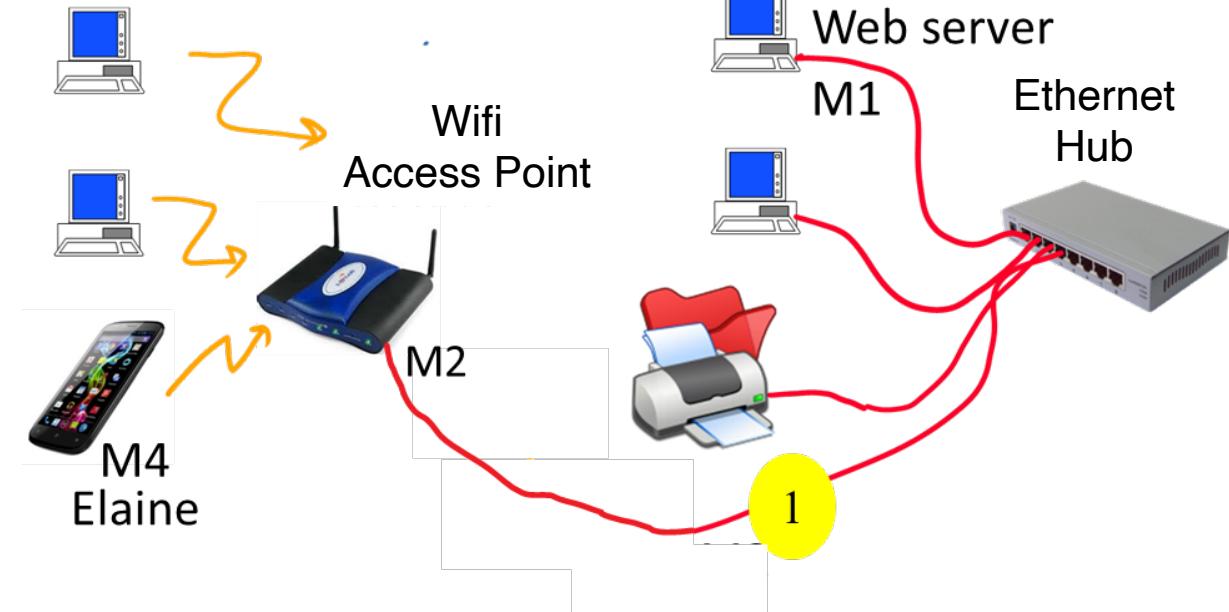


Join room  
46045

# Which topology is more efficient? Which one uses resources better?



- A. Top-left
- B. Bottom-right
- C. I don't know



Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045

# Solution

Answer D

Answer C

Answer B

In both cases the answer is M4.

In both cases there is only one large bridged network.

In the first case, the WiFi base station is not connected to a wired infrastructure. Instead, it is connected wirelessly to a bridge. This is an atypical scenario – it is **not** very efficient because all WiFi traffic that goes to the bridge uses WiFi twice, once to access the base station, once to go from base station to the bridge.

The second case is the *typical use* of an access point. The base station contains a bridge and is itself connected by Ethernet to the rest of the network.

## 6. Security Aspects

- MAC addresses are sent in the clear.

- Possible attacks:

*Eavesdropping*

overhear someone else's data = MAC payload

*Free riding*

insert a device into a network without authorization

*Impersonation*

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

- E.g.: Ethernet at home or EPFL—no authentication of MAC addresses

# Solutions

- **Access control** requires user to show credentials before allowing access to the network;
  - shared = user knows a password, same for all (e.g. WiFi WPA-Personal)
  - per-user = user has a personal password (e.g. WiFi WPA-Enterprise)
- **Authentication**: every MAC frame is digitally signed using cryptography and numbered
  - prevents free riding
  - prevents impersonation if access control is per user;
- **Encryption**: MAC frame payload (but not MAC addresses) is encrypted
  - prevents eavesdropping

Examples:

EPFL:

WiFi uses per-user access control, authentication and encryption

Eavesdropping, Free-riding, Impersonation are not possible

At home:

WiFi WPA-Personal provides access control but does not authenticate individual users

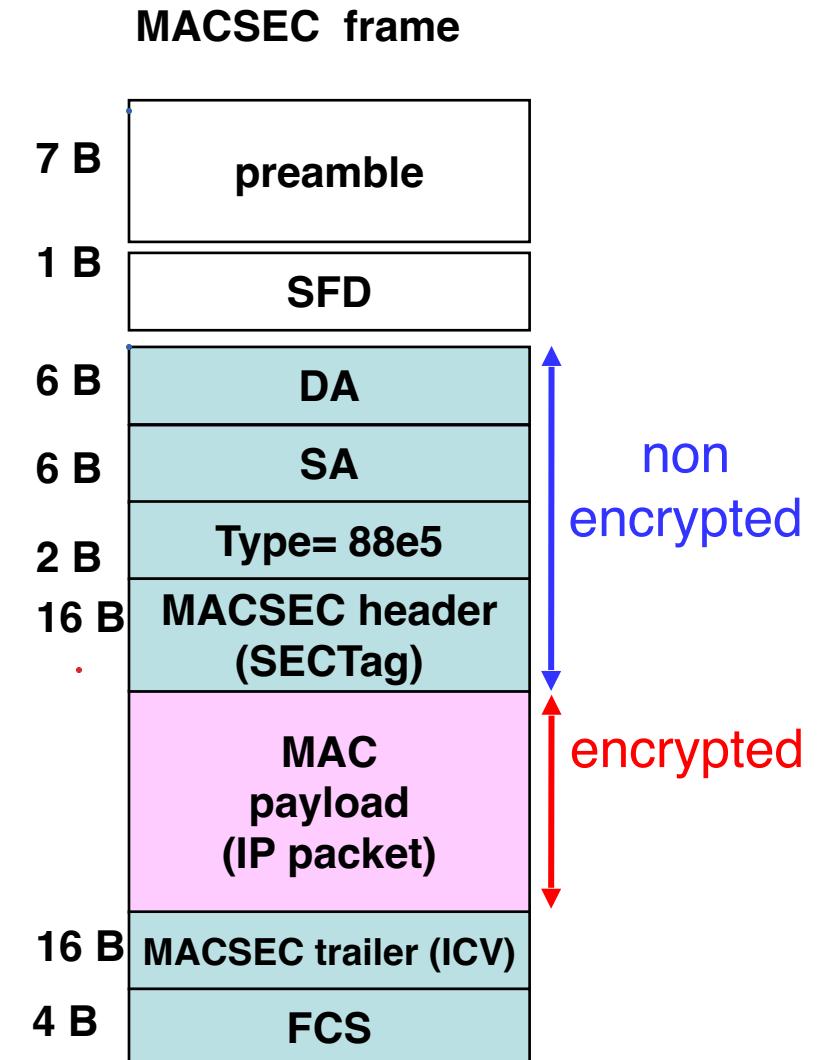
Impersonation is possible by whoever knows the shared WiFi password

# MAC Security (MACSEC)

- i. Host negotiates access with WiFi base station (BS) or Ethernet switch
- ii. WiFi BS/ Ethernet switch contacts an authentication server
- iii. If host is accepted, Ethernet frames between host and WiFi BS/ Ethernet switch are ***authenticated and encrypted***
  - authentication covers MAC payload and MAC source address
  - encryption covers MAC payload

IEEE 802.1X for access control

IEEE 802.1AE for authentication and encryption



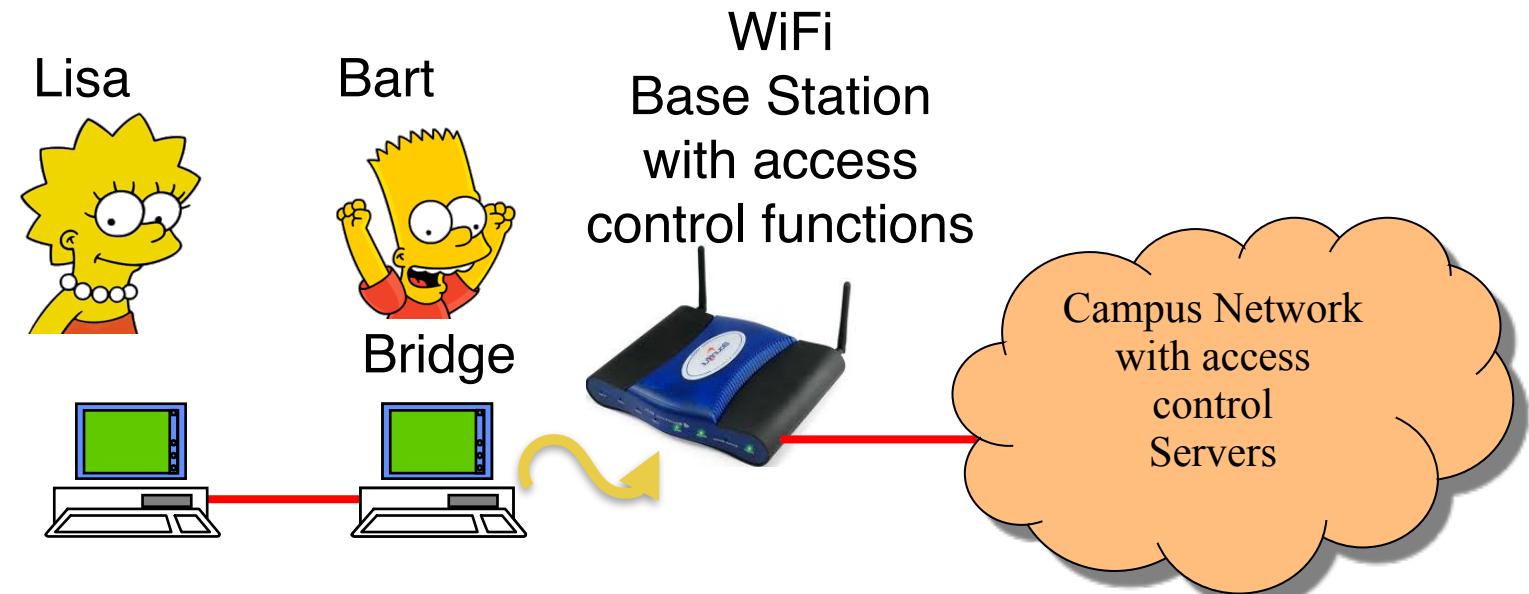
Bart, an EPFL student, connects his PC to EPFL WiFi, which uses per-user access control, authentication and encryption. Bart configures his PC to be a *bridge* between Ethernet and WiFi and allows Lisa, who is not allowed to access EPFL, to connect to his Ethernet port. Will Lisa be able to connect to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

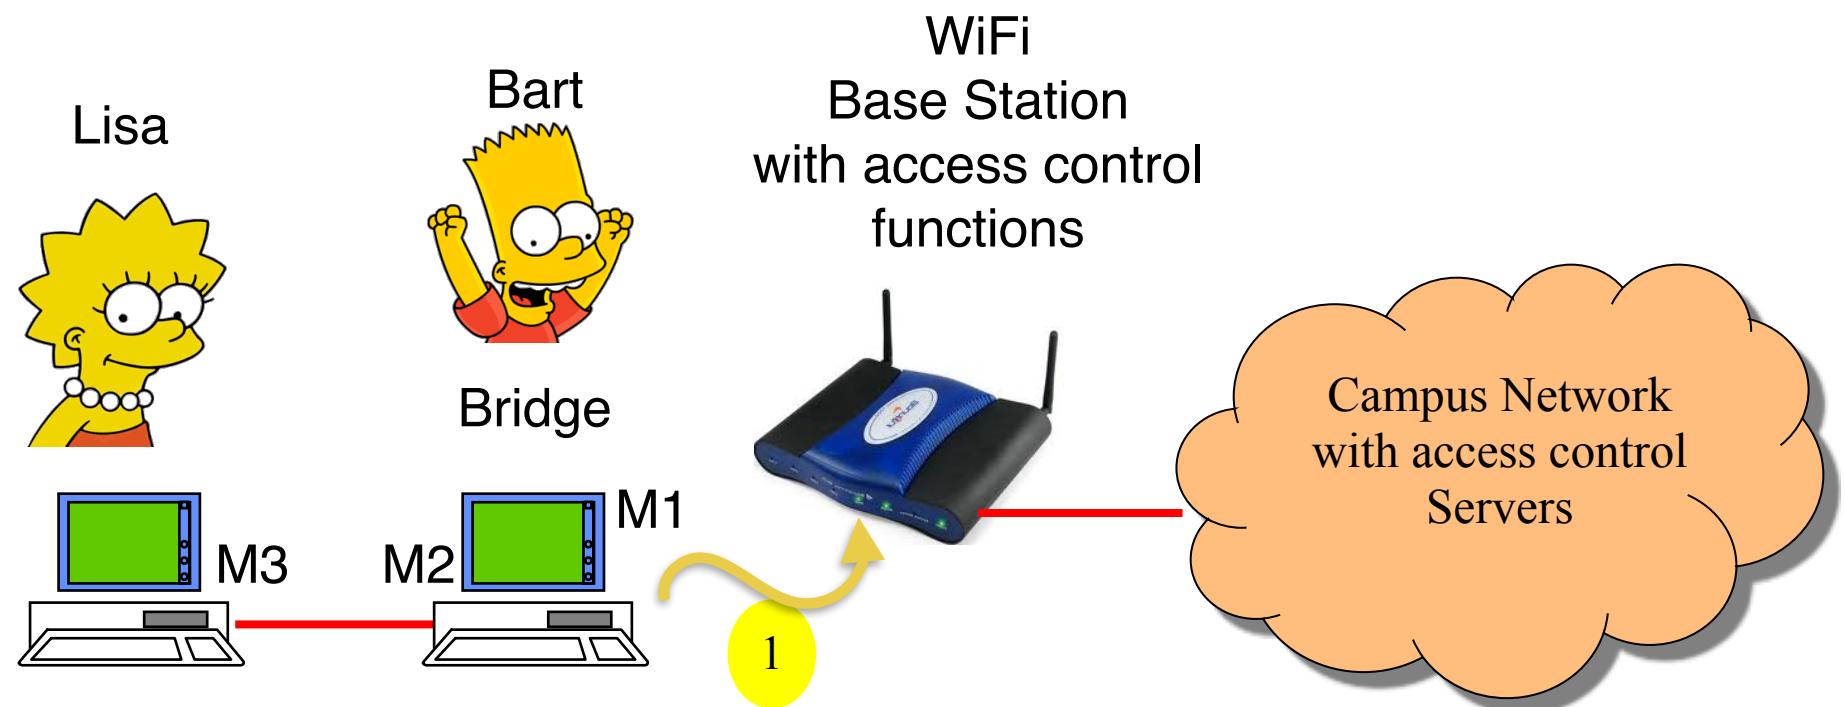


Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045



# Solution



Answer C

Lisa will not be able to connect to campus network in this way.

Access control accepts Bart's MAC address; MAC frames at (1) are authenticated to verify that the source MAC address is indeed Bart's MAC address on WiFi (M1). WiFi base stations accept only authenticated MAC frames.

Lisa's data will appear at (1) with source address = M3( $\neq$  M1).

Even if Bart's PC authenticates Lisa's MAC frames, M3 is not the MAC address that the base station has related to Bart's authentication data/key. So, when trying to verify the authenticity of Lisa's frames, the base station will see this mismatch, and it will discard the frames.

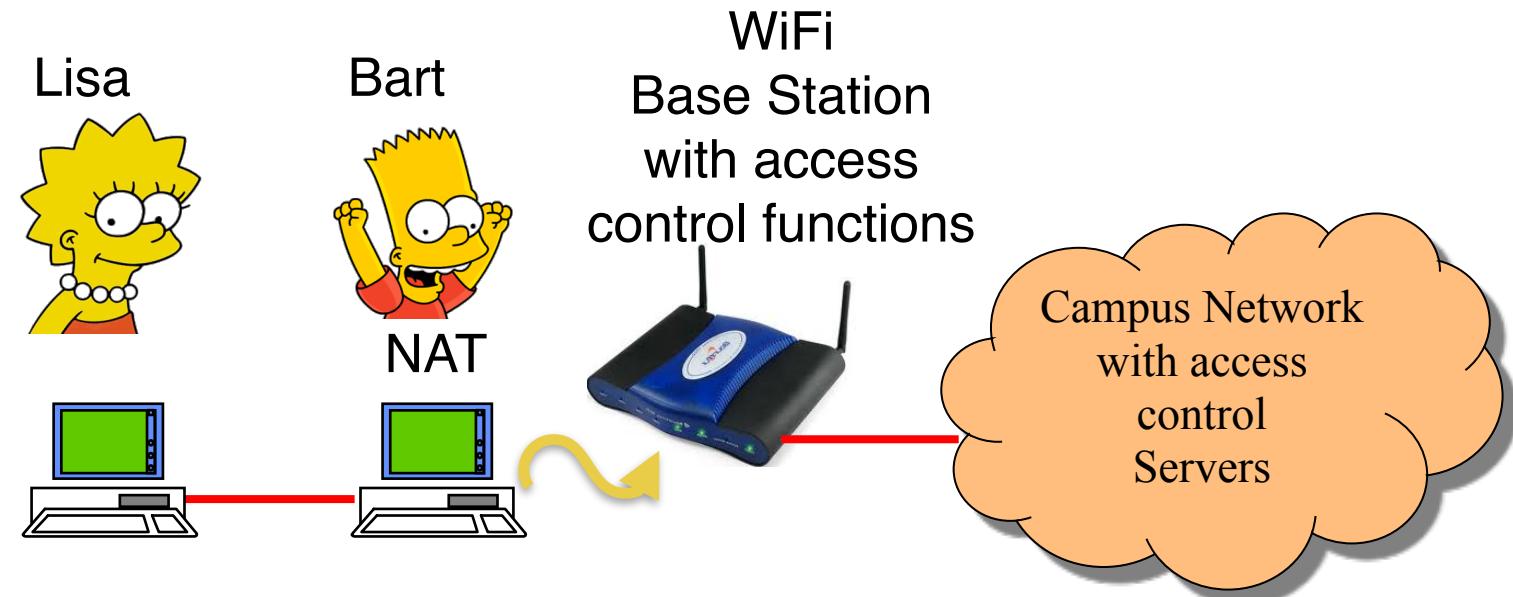
Bart, an EPFL student, connects his PC to EPFL WiFi, which uses per-user access control, authentication and encryption. Bart configures his PC to be a **NAT** between Ethernet and WiFi and allows Lisa, who is not allowed to access EPFL, to connect to his Ethernet port. Will Lisa be able to connect to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

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

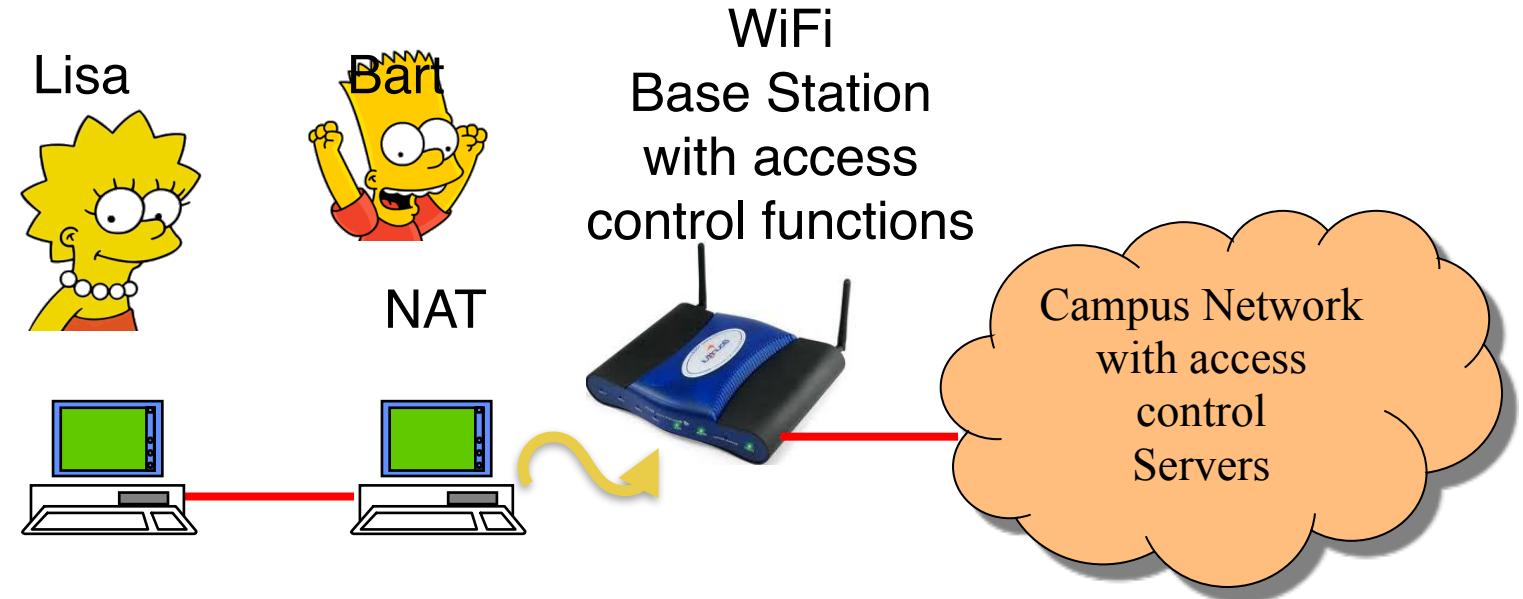


Go to [web.speakup.info](http://web.speakup.info) or  
download speakup app

Join room  
46045



# Solution



Answer B, it will work.

There are 2 layer-2 domains separated by the NAT, which is layer-3 device.

The WiFi frames seen by the WiFi base station belong to the WiFi domain i.e. they are seen as originated by Bart, and thus they are accepted as legitimate.  
Lisa can access the campus network without being a legitimate user !

Answer A is incorrect: it will still work if Bart is behind a NAT. In this case, it is the NAT that must connect to the EPFL WiFi network. Since we assume that Bart could connect to EPFL, it means that the NAT box succeeded to connect to EPFL. In this case, Bart can connect and Lisa also. The EPFL network sees Bart's NAT and thinks that Bart and Lisa are programs on Bart's NAT. An example where this occurs is when Bart connects its laptop to EPFL, so that it is authenticated, and then he uses a guest OS in Virtual box (that has its own NAT).

# 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; so it used a protocol, called CSMA/CD very similar to WiFi.
- Now, the MAC layer for wired medium (Ethernet) does not need any such protocol, because it uses *full-duplex cables* and *interconnected switches*.
  - hence, it forms an *interconnection layer of local scope*.

# Conclusion 2

We saw *2 mechanisms to interconnect* devices:

- the network layer (with IP routers), and
- the MAC layer (with switches).

What are the differences ?

- Switches use forwarding tables that are *not* structured.
  - For every received frame, a switch must search the *entire* table, because of the exact match.
  - The table size and lookup time are *prohibitive*, if network becomes very large.
  - So, we prefer routers for large networks.
- Switches are *plug-and-play* devices, independent from IP numbering plans and whether IPv4 or IPv6 is used (as opposed to routers).
- Historically, routers were in software and switches in hardware; nowadays, every combination exists.

## How do you find the *content* of the lectures so far?

- A. Not interesting at all  
(I know all this stuff from my undergrad)
- B. Somewhat interesting  
(There are few things I missed from my undergrad and I am happy to learn them now)
- C. Interesting  
(I thought I knew much about networks, but I now realize that there are a lot more things to learn)
- D. Other

## How do you find the *way of lecturing* so far?

- A. Very slow pace, I am bored.
- B. Rather slow, many times I feel that time is spent on useless stuff. I want a faster pace.
- C. Engaging, it helps me catch up with all concepts.
- D. Very fast pace, I do not understand critical ideas.
- E. Other



## How do you find the *content* of the labs so far?

- A. Not interesting at all.  
(I know all this stuff from my undergrad)
- B. Somewhat interesting.  
(There are few things I missed from my undergrad and I am happy to have a hands-on experience about them now)
- C. Interesting.  
(I had limited hands-on experience with networks and I am happy to dive into these details)
- D. Other

## How do you find the *labs* so far?

- A. Completely useless, I am losing my time
- B. Somewhat useful, they help me catch up with things I already knew
- C. Useful, they help me understand in detail what is taught in the lectures and all new things I learn
- D. Other