

RCOMP - Redes de Computadores (Computer Networks)

2025/2026

Lecture 03

- Ethernet local area network technologies.
- Virtual local area networks (VLAN).
- Wireless local area networks (WLAN).

ETHERNET Networks – CSMA/CD


Ethernet networks (IEEE 802.3 / ISO 8802-3) were originally developed by Xerox in the 70s. Nowadays, ethernet is undoubtedly the most widely used technology in wired LANs.

Originally, access control to the medium (MAC - Medium Access Control) was a key issue. The CSMA/CD technique used in ethernet is not ideal, it doesn't avoid collisions and, as such, results in low efficiency under heavy traffic.

The early Ethernet networks were based on a coaxial cable to which all nodes were connected (**bus topology**), the most important variants were:

Thick Ethernet - 10base5 - 10 Mbps / Digital Signal⁽¹⁾ / maximum 500 m bus length


Thin Ethernet - 10base2 - 10 Mbps / Digital Signal⁽¹⁾ / maximum 180 m bus length



⁽¹⁾ **base** stands for a baseband transmission medium, therefore, with digital signals being used.

ETHERNET networks – Collision Domain

CSMA/CD (Carrier Sense Multiple Access with **Collision Detection**) requires packet collisions to be detected by all nodes before the emission of the packet ends. This introduces limitations on the relationship between the packet's transmission time and the signal propagation delay.



To ensure collision detection by all nodes, there's a **minimum packet size** of 64 bytes (this sets a minimum transmission time), also there is a **maximum segment size** (sets the maximum propagation delay). These two limits guarantee a collision at any point is detected by the farthest node before the packet is completely transmitted. This is called the **collision domain**.

The collision domain may or may not match the Ethernet network extension, **store & forward** devices do isolate collision domains. Notice that higher transmission rates, result in increasingly smaller collision domains because it takes less time to transmit the minimum size packet (64 bytes).

ETHERNET – packet format and addresses

Ethernet networks evolved significantly over time, but they have always kept the same packet format and addressing schema. This allows full compatibility among various technical developments that have taken place, so even the latest versions of 10 Gbps over optical fibre can work together with old coaxial copper segments using 10base5 and 10base2.

Each node is identified by a 48-bit number node address also known as physical address or MAC address.


Usually, these addresses are represented in hexadecimal as a sequence of six bytes, separated by colons, for example **00:60:B0:3C:93:DB**.

To ensure addresses are unique, to each hardware manufacturer is assigned with a unique fixed sequence for the first 24 bits.

FF:FF:FF:FF:FF:FF address is the broadcast address; a packet sent to this address will be delivered to every node within the same Ethernet network.

Ethernet II packet format

Keeping the same packet format throughout its evolution was a key factor for Ethernet networks success. Within the logical link layer, packets are usually called **frames**. Different Ethernet frame formats have existed, but nowadays **Ethernet II**, also known as DIX (Digital, Intel, Xerox) and is universal:



Among various existing formats, this is the simplest, yet it implements all that's required: source and destination node addresses, a multiplexing identifier (E-TYPE) and an error detection code (FCS - Frame Check Sequence). It can carry up to 1500 bytes of data; this maximum payload size is the **MTU (Maximum Transmission Unit)**. The **synchronisation sequence** depends on the transmission medium; it is used for bit and frame synchronization.

This ethernet packet format is so widespread that new technologies, such as 802.11 access points, support this format to allow direct interconnection between cabled Ethernet networks and wireless local networks.

Ethernet – from the bus to the star topology


The bus topology with coaxial cable, e.g., 10base5 and 10base2, provided extremely low-cost networks, this was a key factor in the initial expansion of Ethernet. Yet, especially with 10base2, reliability became an important issue, any physical problem along the bus makes the entire bus unusable.

In the early 90s, **star topology** Ethernet began to emerge. Based on two twisted copper pairs (10baseT) or two optical fibres (10baseFL and 10baseFB) and a **hub repeating device**. In these variants, each node has two separate connections (TX and RX) to the hub device.

Despite this new topology, initially CSMA/CD was still required and limitations to the collision domain (maximum distance between two nodes) were significant, for instance in 10baseT it's 500 meters.

But now, new potentials arise:

- **Switching.**
- **Full-duplex.**



Ethernet – frame switching

The star topology opens new possibilities, nodes have two separate connections, one for sending (TX) and one for receiving (RX), thus, collisions will not happen here.

If collisions are eliminated, then we could operate in **full-duplex** because when sending (TX) we no longer need the RX connection to listen for collisions detection (CD), so we can use it to receive data at the same time.

If the HUB is modified in such a way it can:

- Simultaneously receive frames on all ports.
- Simultaneously emit frames on all ports.
- Temporarily store frames when needed.
- Record a frame's **source address** when it's received through a port (build the **MAC table**).
- Check a frame's **destination address** and retransmit it only on the required port (by checking the **MAC table**).

This new device can now be renamed as **switch**, and it brings enormous improvements in network performance because it removes the biggest Ethernet issue: **collisions and CSMA/CD**

Switched Ethernet

Far more than any increase in transmission rate, switching was a huge progress since it eliminated the major issue around original Ethernet networks.

- Full-duplex operation, no collisions and no need for medium access control.
- Switching, based on nodes address, thus, a frame is delivered only to the destination node (MAC table management). This massively reduces traffic as frames are propagated only to where they are required, and not to every node.
- Eradication of collision domains, therefore removing restrictions on the maximum size of the Ethernet network.

These factors have increased dramatically the overall efficiency of the network, leading to an increase in the apparent data rate far greater than any increase in the nominal transmission rate.

When mixing repeating hubs and switches, parts of the network covered by hubs will still operate with CSMA/CD and impose local collision domains around them. These collision domains around hubs are, however, contained and confined by switches.

Ethernet – some current LAN technologies

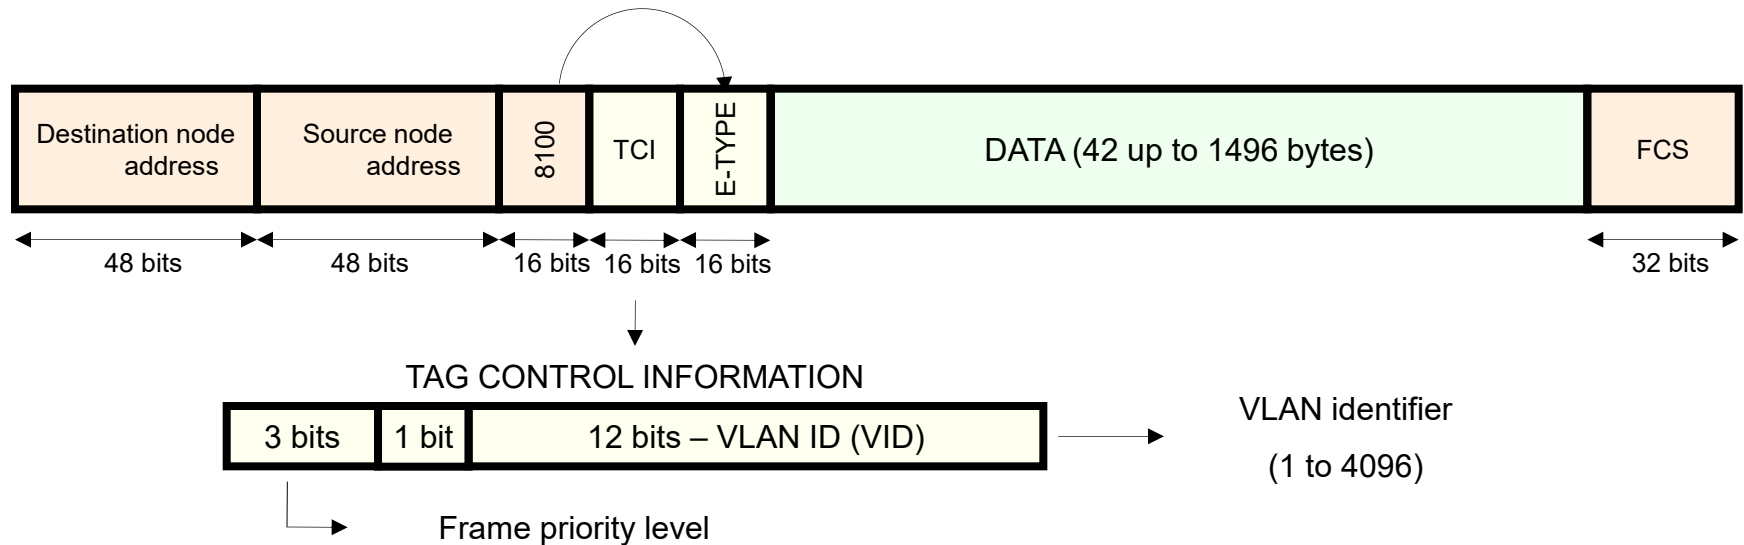
1 Gbps	1000baseT : 4 twisted copper pairs (CAT 5 or upper), maximum cable length is 100 meters. 1000baseSX : 2 multi-mode optical fibres, maximum cable length up to 1 Km. 1000baseLX : 2 single-mode optical fibres, maximum cable length up to 5 Km.
2.5 Gbps	2.5GbaseT : 4 twisted copper pairs (CAT 5E or upper), maximum cable length is 100 meters.
5 Gbps	5GbaseT : 4 twisted copper pairs (CAT 5E or upper), maximum cable length is 100 meters.
10 Gbps	10GbaseT : 4 twisted copper pairs (CAT 6A or upper), maximum cable length is 100 meters. 10GbaseSR : 2 multi-mode optical fibres, maximum cable length up to 500 m. 10GbaseLR : 2 single-mode optical fibres, maximum cable length up to 10 Km.
25 Gbps	25GbaseT : 4 twisted copper pairs (CAT 8 or upper), maximum cable length is 30 meters. 25GbaseSR : 2 multi-mode optical fibres, maximum cable length up to 400 m. 25GbaseLR : 2 single-mode optical fibres, maximum cable length up to 10 Km.
40 Gbps	40GbaseT : 4 twisted copper pairs (CAT 8 or upper), maximum cable length is 30 meters. 40GbaseSR4 : 8 multi-mode optical fibres, maximum cable length up to 125 m. 40GbaseLR4 : 8 single-mode optical fibres, maximum cable length up to 10 Km.
50 Gbps	50GbaseSR : 2 multi-mode optical fibres, maximum cable length up to 100 m. 50GbaseLR : 2 single-mode optical fibres, maximum cable length up to 10 Km.
100 Gbps	100GbaseSR1 : 2 multi-mode optical fibres, maximum cable length up to 125 m. 100GbaseLR1 : 2 single-mode optical fibres, maximum cable length up to 10 Km.

Recent developments for higher data rates: 200 Gbps, 400 Gbps, 800 Gbps, and 1.6 Tbps

Virtual Local Area Networks - IEEE802.1Q

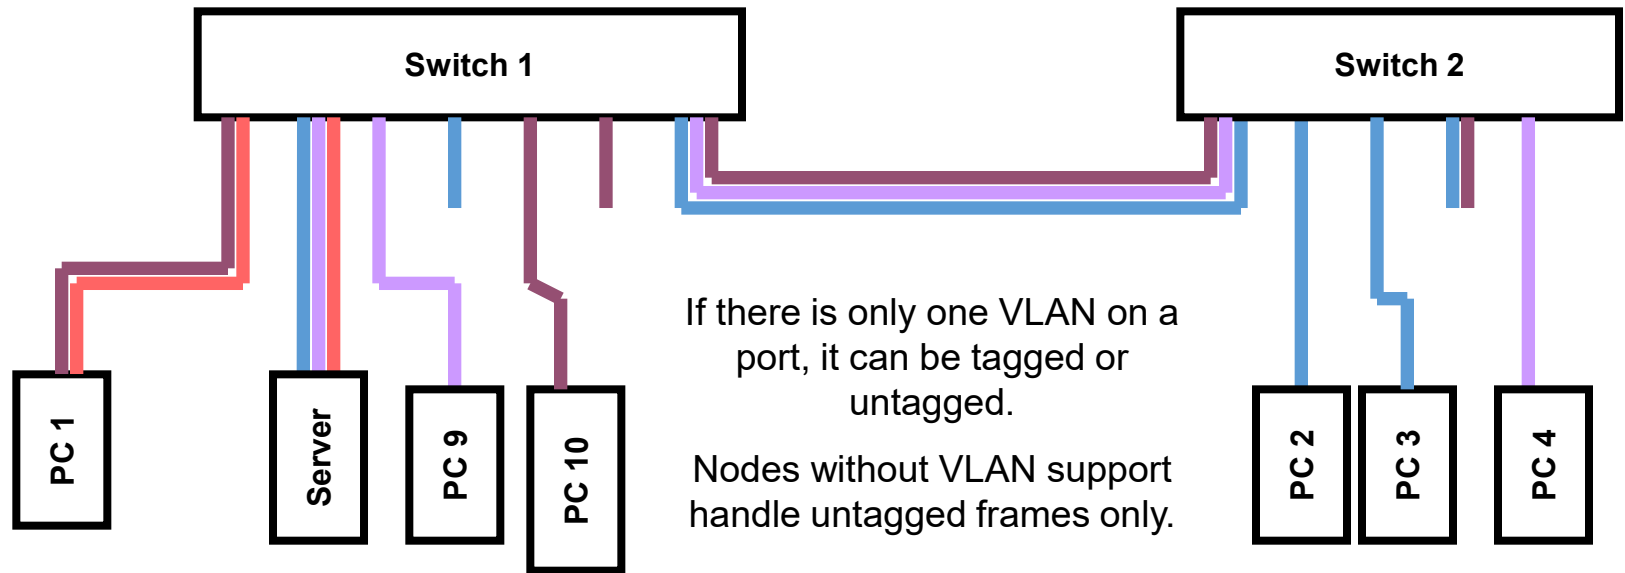
A Virtual Local Area Network (VLAN) is a logical network defined over a physical network. Nevertheless, from all points of view, a VLAN must behave as an independent physical network.

Over the same physical network, frames can be tagged in such a way each tag represents one different VLAN. The IEEE802.1Q standard defines how to place this tags into Ethernet II frames: value 0x8100 is placed on E-TYPE field and its original value is shifted 4 bytes.



Virtual Local Area Networks - switches

Switches may settle VLANs without frame tagging, they can define a VLAN as a subset of their ports. By doing so, the switch will work as several independent virtual switches. One switch port can be associated to more than one VLAN, on that case, only one VLAN can be untagged.



Port-based VLANs are the most often used, but some switches also support MAC address based VLANs, where the VLAN to which a node belongs is determined by the node's MAC address.

Wireless Local Area Networks (WLAN)

Wireless LANs are an important progress towards accessibility and mobility on one hand, and simplification of cable installations, on another hand.

The most important standard is IEEE 802.11 and respective amendments. In early versions, transmission rates were far from wired LANs, i.e., 802.11 at 2 Mbps, but current standards, like 802.11n can operate up to 600 Mbps using multiple simultaneous channels (multiple emitters / receivers).

More challenging than low transmission rates is the return back to shared transmission mediums, and the need for low-efficiency MAC mechanisms like CSMA.

Even if MAC was totally efficient, using it means the maximum available data rate will have to be divided by all nodes that are eager to send frames.

Moreover, on a wireless shared transmission medium, any physical access control is impossible. Being wireless, this brings even more privacy and security issues than on wired shared transmission medium networks.

802.11 - modulation

The original 802.11 standard anticipated an implementation based on infrared light; however, all subsequent developments use radio waves. The used bands are centred in 2.4 GHz and 5.7 GHz frequencies, both in the microwaves zone.

Being analogue signals by nature, data transmission uses digital modulation techniques. The modulation techniques currently used are rather complex, using multiple signals simultaneously with multiple PSK and ASK combinations. Some of these techniques have been developed throughout the evolution of telephone line DSL modems and mobile networks.

The 802.11 standard and its amendments establishes several alternative modulation techniques, leading to various transmission rates. It is up to the nodes trying the different techniques to get the best rate possible with an acceptable error rate. Generally speaking, the lower rate options are more reliable when the signal strength is low.

Frequencies used (microwaves) and legal restrictions on the transmission power (100 mW) result in a very limited reach, especially inside buildings (usually far less than 50 meters).

802.11 – CSMA/CA

The biggest issue on wireless LANs is the use of a shared transmission medium; this implies only one node can emit at a time.

- Even if MAC was 100% effective, and that's not true, the maximum available data rate is always divided by the number of nodes.
- Furthermore, a node can't send and receive at the same time, thus, transmissions are always half duplex. Due to this, CSMA/CD can't be used to detect collisions, for that a node must be **listening while talking**.

As alternative, **CSMA/CA (Collision Avoidance** instead of Collision Detection) is used, a node must check if the transmission medium is free (no signal present). If busy, it must wait for a random period before checking again. If the medium is free, then it can start sending the frame.

Collisions can't be directly detected, instead, the receiving node is compelled to send back an **ACK** signal if it succeeds to receive the frame. This informs the sender everything went ok, and hence, there was no collision. If no ACK is received, the frame is assumed to be lost and will have to be sent again. This is, therefore, a **stop and wait** error control implementation.

802.11 – IEEE 802.11 RTS/CTS


CSMA/CA may be combined with the **RTS/CTS** technique. However, **RTS/CTS** is used only if the frame size is above the **RTSThreshold** value, otherwise, it's disadvantageous.

With **RTS/CTS**, before sending a frame, nodes must send a **Request to Send (RTS)** to the destination node, the destination node may then reply with **Clear to Send (CTS)** meaning it's ready to receive.

When a node hears a third party **RTS** or a **CTS** is obliged to wait for a period of time before trying to send, this avoids collisions when the frame is being sent.

RTS/CTS is especially effective in **infrastructure mode** where there's a central device called **access point (AP)** by which all communications must pass. The access point operates roughly as a wireless switch.


By opposition in a WLAN without an access point, every node can directly send to any other node, and moreover, any node can also retransmit frames. This operation mode is called **ad-hoc mode**.



802.11 – infrastructure mode

The infrastructure mode requires a central device by which all communications pass, truly it may be seen as a wireless star topology.


One major advantage of infrastructure mode comes in combination with the RTS/CTS technique. Because the access point will take part on every RTS/CTS dialog, it can, therefore, avoid most collisions.



802.11 – cells

In infrastructure mode, wireless coverage of a large area can be assured by dividing it into smaller areas called BSS (Basic Service Set), also known as **cells**. Each cell is controlled by a **base station**, also known as AP (Access Point). Each cell has a unique identifier called **BSSID** which is the 48 bits MAC address of the base station.

A set of cells can be part of the same infrastructure, called ESS (Extended Service Set), the ESS is identified by a sequence of 0 to 32 bytes called **SSID** (Service Set Identifier). Often SSID bytes are human readable characters and interpreted as a meaningful string.



By using the same SSID on all cells, all cells will be part of the same ESS. Wireless network nodes can then move freely between cells of the same ESS without losing network access. The transparent transfer of a wireless node from one cell to another cell within the same ESS is known as **roaming**.

802.11 – segmentation

Dividing the area into small cells it's always a good idea; it restrains the negative effects of a shared transmission medium. Each cell is a collision domain, so more cells mean smaller collision domains. Increasing the number of cells (APs), ensures that each cell will contain a smaller number of nodes, and thus, mitigates the negative effects the shared transmission medium.

The number of cells should be the required to ensure full coverage of the desired area, but furthermore, it should also ensure that the number of nodes in each cell is not very high.

It's acceptable to have two or more APs installed in a 20 m² room, it all depends on the number of workstations and the desired efficiency, however, with closely installed APs special care must be taken on avoiding overlapping frequency channels.

Wireless AP interconnection (**wireless distribution system**) is possible, but it should always be avoided as it will simply extend the collision domain to all cells. **Access points should always be connected to a cable infrastructure.**

802.11 – frames

The operation of 802.11 is complex, involving nodes with different roles (i.e., end nodes and access points) and different specific control information. This results in rather complex frame format, for instance a single 802.11 frame can contain up to 4 MAC addresses (two end nodes and two intermediate access point nodes).

Despite these internal complexities direct connection to local wired networks (Ethernet) is simple because the address format is the same and the data and control fields can be transported directly between 802.11 frames and 802.3 frames. This is a mission of the access point that operates as an interface between the WLAN (802.11) and Ethernet (802.3).

In fact, a significant effort was made to keep direct compatibility with 802.3 frames. Often (on high error rates) it's convenient to use very small 802.11 frames, but 802.3 frames can be up to 1518 octets long. To solve the problem 802.11 nodes, have the ability of fragmenting frames 802.3 into segments and later reassemble these segments to rebuild the original 802.3 frame.

802.11 – security

Being a shared medium network, security is always a challenge, in a wireless network, the problem is even worse because unlike a wired network, the medium is publicly accessible without any physical contact.

Access Control

Access points must implement access control mechanisms to avoid public free access to the cell (**called association**). MAC address-based authentication is not safe, better alternatives are user's authentication or the use of a secret Pre-Shared Key (PSK).

Privacy

As transmitted frames are publicly available, to ensure privacy, encryption must be used. Symmetric key encryption requires that both the node and the AP must have a same secret key. The secret key can be manually pre-shared (then it will also work as authentication/access-control) or can be generated during the user's authentication procedure. Another option is the use of public key certificates for authentication/access-control and public key encryption for privacy.