3.1. Introduction

The prevalent high-speed LANs today are Ethernet-based and were developed by the IEEE 802.3 standards committee. These LANs consist of a medium access control layer and a physical layer, which we'll discuss sequentially.

Ethernet employs shared media technology, necessitating a method to manage access to the physical media. This method is known as Carrier Sense Multiple Access Collision Detection (CSMA/CD) circuit.

The concept is straightforward: before sending, each station performs Carrier Sense to check if the media is in use. If available, the station gains Media Access and can transmit data. If two stations access the media simultaneously, a collision occurs. Collision recognition and resolution are handled by the Collision Detect circuit. While sending, each station monitors its own data. Collisions are detected by observing the DC-level on the medium. A jamming signal is sent to ensure all stations acknowledge the collision. Involved stations halt transmission, initiate a randomized timer, and resume access attempts when the timer expires.

The main IEEE 802 Standards

802.1 High-Level Interface, Network Management, Bridging, Glossary
802.2 Logical Link Control
802.3 CSMA/CD Ethernet (LAN)
802.4 Token Bus (LAN)
802.5 Token Ring (IBM LAN)
802.6 Metropolitan Area Network (DQDB: Distributed Queue Dual Bus)
802.7 Broadband LAN Technical Advisory Group
802.8 Fiber Optic Technical Advisory Group
802.9 Integrated Service LAN (IsoEthernet) for isochronous (real-time) communication
802.10 LAN Security (SILS: Standard for Interoperable LAN Security)
802.11 Wireless LAN
802.12 Demand Priority LAN (100VG AnyLAN)
802.14 Cable TV MAN
802.15 Wireless Personal Area Network (WPAN), Bluetooth
802.16 Fixed Broadband Wireless Access (wireless broadband)

3.2. Ethernet

- 802.3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD): Ethernet
- Ethernet 802.3 was initially led by DEC, Intel, and Xerox in the early 1980s. It was one of the first widely adopted Ethernet standards and is still used in many networks today.
- Ethernet 802.3 is considered to be an elementary and easy-to-administer technology. It is relatively simple to implement and troubleshoot.

3.2.1. Ethernet History

- 1973: Bob Metcalfe and David Boggs at Xerox PARC invented Ethernet, originally called Alto Ethernet. It was a 2.94 Mbps coaxial cable network that could connect up to 256 devices.

- 1979: DEC, Intel, and Xerox formed the DIX Consortium to standardize Ethernet and promote its wider adoption. The DIX Consortium published the first Ethernet specification in 1980.

- 1980: The IEEE Project 802 was launched to develop standards for local area networks (LANs). Ethernet was one of the first LAN technologies to be standardized by the IEEE.

- 1985: The IEEE ratified the 802.3 standard, which defined the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol used in Ethernet networks.

- 1989: The IEEE ratified the 802.3u standard, which defined Fast Ethernet, a 100 Mbps version of Ethernet.

- 1995: The IEEE ratified the 802.3z standard, which defined Gigabit Ethernet, a 1 Gbps version of Ethernet.

- 1999: The IEEE ratified the 802.3ab standard, which defined Gigabit Ethernet over twisted-pair copper cables.

- 2002: The IEEE ratified the 802.3ae standard, which defined 10 Gigabit Ethernet (10 Gbps).

- 2007: The IEEE ratified the 802.3an standard, which defined 10GBASE-T, enabling 10 Gigabit Ethernet over twisted-pair copper cables.

- 2010: The IEEE ratified the 802.3ba standard, which defined 40 Gigabit Ethernet and 100 Gigabit Ethernet.

- 2016: The IEEE ratified the 802.3bz standard, which defined 2.5GBASE-T and 5GBASE-T, providing intermediate speeds between 1 and 10 Gigabit Ethernet.

- 2020: The IEEE ratified the 802.3cg standard, which defined 10BASE-T1L, enabling Ethernet connectivity over a single twisted-pair cable for industrial applications.

  10: The data transmission rate is 10 Mbps.
  BASE: The signal is baseband, meaning that only Ethernet signals can be transmitted over the cable.
  T: The cable is twisted pair.
  1: The cable can support link lengths of up to 1 kilometer.
  L: The cable is designed for long-range communication

3.2.2.. Ethernet MAC Addresses

Ethernet MAC addresses are 48-bit unique identifiers assigned to network interface controllers (NICs). They are used to identify devices on a network and to enable communication between them.

MAC addresses are transmitted with the least significant bit first (LSB). The first bit of the MAC address is the individual/group (I/G) bit. It indicates whether the MAC address is an individual address (0) or a group address (1).

The second bit of the MAC address is the universally/locally administered (U/L) bit. It indicates whether the MAC address was assigned by the manufacturer of the NIC (universally administered, 0) or by the network administrator (locally administered, 1).

The first 24 bits of the MAC address form the organizationally unique identifier (OUI). The OUI is assigned to the manufacturer of the NIC by the IEEE. The remaining 24 bits of the MAC address are the serial number of the NIC.

Examples of Ethernet MAC addresses:

- 00:00:0C:00:00:01 (Cisco)
- 00:C0:4F:00:00:01 (DELL)
- 33:33:FF:00:00:01 (multicast address)
- 00:00:00:00:00:01 (broadcast address)

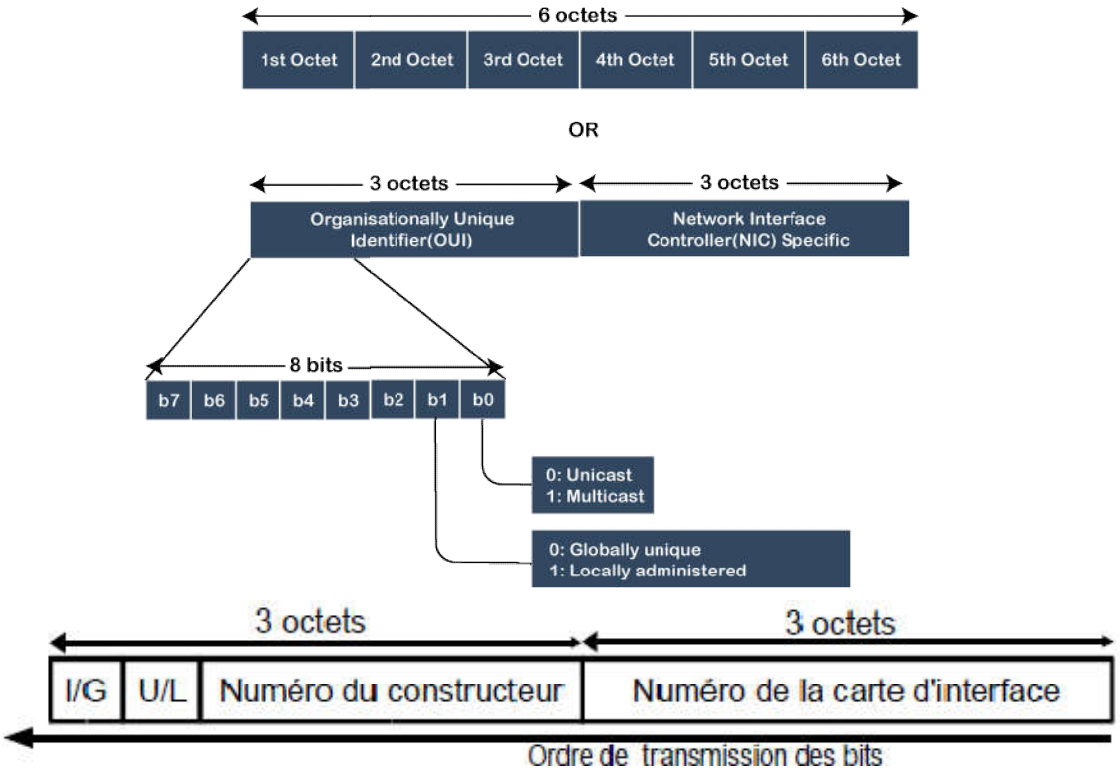

3.3. Frame Encapsulation at the "Network Access" Layer

3.3.1. The 802.3 Frame

The Ethernet frame has a capacity of 64 to 1518 bytes, divided into various fields. It is always preceded by two fields that enable synchronization with the network. These fields are not recorded during a frame capture. The following are the different fields that make up an Ethernet 802.3 frame:

Destination Address	Source Address	Length	Data Field + Padding	FCS
6 octets	6 octets	2 octets	46 to 1500 octets	4 octets

This frame practically no longer exists in modern networks. It has been replaced by the Ethernet II frame, which is designed to carry IP datagrams.

Note: The two fields that are not recorded during a frame capture are the preamble and the start of frame delimiter (SFD). These fields are used to synchronize the receiver with the transmitter, but they are not necessary for the receiver to decode the frame data.

3.3.2. The Ethernet II Frame

The Ethernet II frame was created by a consortium of companies to address specific needs. It can circulate on the same network as the 802.3 frame. The choice between the two depends on the upper-layer protocols used.

Destination Address	Source Address	Layer 3 protocol	Data Field + Padding	FCS
6 octets	6 octets	2 octets	46 to 1500 octets	4 octets

The difference between an Ethernet II frame and an IEEE 802.3 frame lies in the value of the 3rd field. If this value is less than 1500, it is an IEEE 802.3 frame. If this value is greater than 1500, it is then an Ethernet II frame.

Fields of the Ethernet II Frame

• Destination address & source address fields (6 bytes)

These two fields indicate the destination and source addresses. They represent the physical addresses of network interface cards (NICs). These addresses are encoded in 48 bits (6 bytes). The first bit specifies whether it's an individual address (0) or a group address (1). The second bit indicates whether the group address is multicast (0, for a group of stations) or broadcast (1, for general broadcast). The 48-bit addresses are unique. IEEE assigns a specific number to

each manufacturer, which forms the 3 most significant bytes of the address. The manufacturer then manages the remaining bits of the address themselves. Therefore, regardless of the hardware's origin, there's no possibility of a physical address conflict on the network.

• Data field (Information)

The data field is often referred to as the information field; these terms are often used interchangeably.

• Padding field

The padding field fills the data field if it contains less than 46 bytes. Note that the value of 64 bytes for an Ethernet frame is calculated based on the DESTINATION, SOURCE, LENGTH, DATA, INFORMATION, and FCS fields.

• FCS field

The Frame Check Sequence (FCS) field provides reception control for the frame. The sender performs a calculation on the DESTINATION, SOURCE, LENGTH, and INFORMATION fields, and writes the result into the 4 bytes of the FCS. The recipient performs the same calculation and checks for matching results. If there's no match, the frame is blocked by the recipient's MAC sublayer, which then signals the manager. The calculation is called CRC (Cyclic Redundancy Codes) and is based on polynomial division using a predetermined polynomial.

3.4. Encapsulated Frame at the "Internet (IP)" Layer

3.4.1. The IP Packet

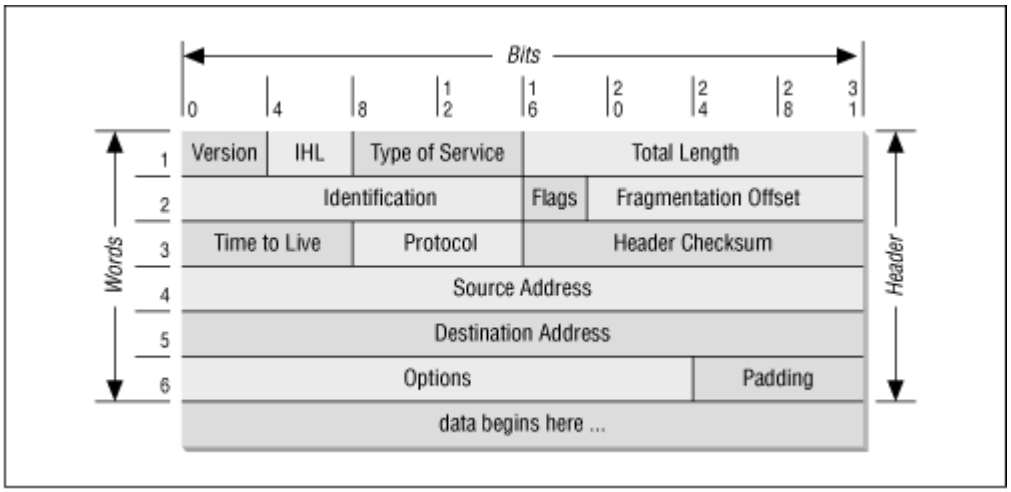

• Version Field: This 4-bit field indicates the version number of the IP protocol being used (typically version 4).

• Header Length Field: Also 4 bits, it indicates the length of the header in terms of 32-bit words (4 bytes).

• Type of Service Field

This 8-bit field comprises:

- 4 bits for "Priority": D for low delay,
- ✓ T for high throughput,
- ✓ R for high reliability,
- ✓ C for low cost.
- 4 bits for "Service": Telnet = 1000,
- ✓ FTP control = 1000,
- ✓ FTP data = 0100,
- ✓ SNMP = 0010.
- Total Length Field:

This 16-bit field is expressed in bytes. If the datagram's length exceeds the maximum size, it's divided into segments. The total length indication helps distinguish padding in an Ethernet frame.

- Identification Field:

With 16 bits, it's used to identify a datagram in case of fragmentation (it's copied into each segment).

- Flags Field:

A 3-bit field:

- ✓ DF (Don't Fragment): Set to 1 if the frame should not be fragmented.
- ✓ MF (More Fragments): Set to 1 if the frame has been fragmented, and if this fragment is not the last one.
- Fragment Offset Field:

With 13 bits, this is used for reconstructing IP frames that had to be fragmented while passing through certain media. This value indicates the relative position, in multiples of 8 bytes, of this frame fragment in the original frame. This counter is also used for reconstructing

fragmented frames on the receiving machine; it's decremented each second until all fragments comprising the original frame have arrived.

- Time to Live (TTL) Field:

This 8-bit field indicates a lifespan, in seconds, for the frame. The frame should be discarded when this field becomes zero. Each node traversal essentially involves decrementing this field.

- Protocol Field:

This 8-bit field indicates the upper-layer protocols being used: ICMP = 1, TCP = 6, UDP = 17.

- Header Checksum Field:

With 16 bits, this is a CRC recalculated by each router before retransmission. It helps detect header inconsistencies and possible transmission errors. The actual data is not included in this calculation.

- Source and Destination Address Fields:

Each 4 bytes in length, they indicate the IP addresses.

- Option Field:

Of variable length, it can be empty, with padding added to make it a multiple of 32 bits.

3.4.2. The ARP Packet

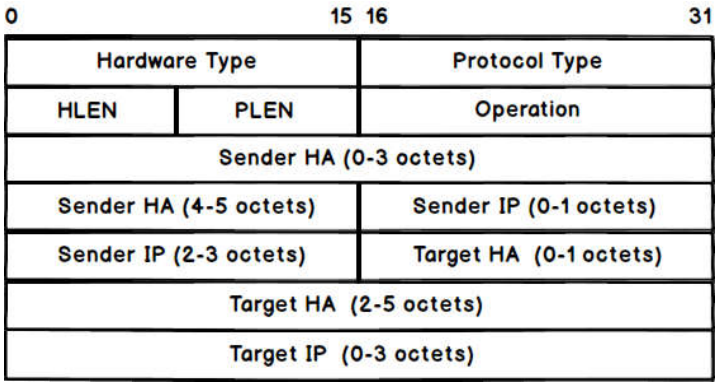

- Field: Physical Network Address Identifier Description: A value of 1 pertains to classic Ethernet networks.
- Field: Protocol Address Identifier Description: It indicates the protocol for which the address is being sought. For ARP, this field is set to 0x0800 (IP).
- Field: Physical Address Length Description: This represents the MAC address in bytes (typically 6).

- Field: Network-Level Protocol Address Length Description: For IP, this field is set to 4.
- Field: Operation Description: It indicates the type of packet:
  - ✓ 1 for an address request,
  - ✓ 2 for a response.
- Field: Sender's Hardware Address Description: It contains the MAC address of the packet sender. In the case of a response, it contains the sought-after information.
- Field: Sender's Protocol Address Description: It contains the IP address of the packet sender.
- Field: Target Hardware Address Description: It contains the MAC address of the packet receiver. In the case of a request, this field is null.
- Field: Target Protocol Address Description: It contains the IP address of the packet recipient.

3.4.3. The ICMP Packet

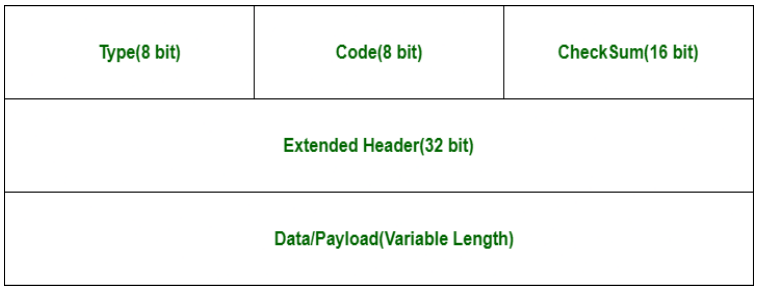

The ICMP header is a variable-length header that is encapsulated within the IPv4 or IPv6 packet. The ICMP header contains the following fields:

- Type (8-bit): The ICMP message type.
- Code (8-bit): The ICMP message code.
- Checksum (16-bit): A checksum used to verify the integrity of the ICMP header.
- Extended Header (optional): A variable-length header that is used to provide additional information about certain types of ICMP messages.
- Data or Payload (variable length): The data portion of the ICMP message.

3.5. Media Access Method

A media access method is a set of rules that define how a computer sends and retrieves data on a communication channel. The name of this media access method, CSMA/CD (Carrier Sense Multiple Access with Collision Detection), is explained by the fact that computers "listen" to the cable (carrier sense). In general, multiple computers on the network attempt to transmit simultaneously (multiple access). They listen to the cable to detect possible collisions (collision detection) and must wait for a random time before retransmitting if a collision occurs.

Collision detection imposes a maximum distance on a CSMA/CD network. Due to signal attenuation, collision detection is not effective beyond 2,500 meters. If multiple computers transmit simultaneously, data collisions occur, resulting in data corruption.

CSMA/CD is also known as a contention method, as computers compete to access the cable. CSMA/CD might seem inefficient, but implementations are fast enough to be invisible to users. The minimum size of a frame is equal to the time it takes for a frame to travel the round-trip distance between the two farthest stations on the network to detect a potential collision, which is the minimum acquisition time of the channel (slot time) multiplied by the data rate. The data rate itself depends on the amount of information a source can send per unit of time (frequency function).

CSMA/CD Principles:

CSMA/CD, an advanced access method compared to ALOHA or slotted ALOHA, minimizes collision inefficiencies by adopting the following principles:

1. Transmit data if the medium is idle; otherwise, proceed to step 2.
2. Continue listening until the medium becomes idle, and then transmit immediately.
3. If a collision is detected during transmission, send a brief jamming signal to signal the collision and stop transmission.
4. After sending the jamming signal, wait for a random backoff period before attempting retransmission, following the sequence from step 1.

CSMA/CD Example : Imagine two stations, A and B, communicating over the same channel. Station A initiates transmission at time t=x (Figure a). Station B, unaware of A's transmission, begins its transmission slightly later at t=x+dt (Figure b). However, due to signal propagation delay (tp), the transmission from B reaches A when A's signal is still propagating in the channel. This results in a collision at time t=x+tp (Figure c), as both signals overlap, causing data corruption.

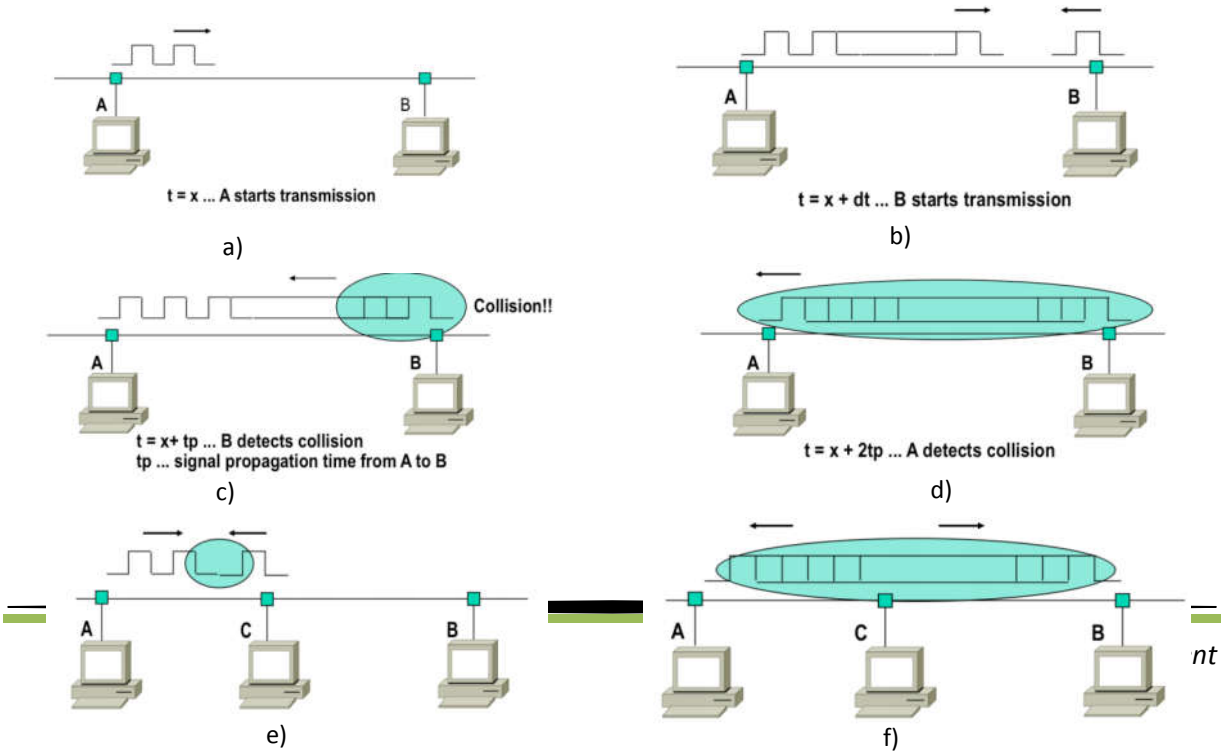

Now, consider a third station, C, located between A and B. When a collision occurs between closely located stations like A and C, a jamming signal is sent out (Figures e and f). This jamming signal extends the collision to a minimum length, ensuring that all network stations, including C, recognize the collision. Without the jamming signal, the collision between A and C might be too short to be accurately detected by other stations' receiving circuits.

In CSMA/CD, before transmitting, stations listen to the channel to detect ongoing transmissions (carrier sense). If the channel is free, they start transmitting. However, while transmitting, stations continuously monitor the channel for collisions (collision detection). When a collision is detected, the colliding stations send jamming signals, signaling other stations to halt their transmissions. After the jamming signal, stations initiate a random backoff timer before reattempting transmission, which helps prevent immediate collisions during retransmission attempts.

Although CSMA/CD may seem inefficient due to possible collisions, its implementation is optimized for speed and remains transparent to users. To prevent distant stations from experiencing excessively delayed collision detection, the protocol mandates a minimum frame size based on the time required for a signal to travel the farthest station distance.