

UNIT 3

The Data Link Layer

(8 Hours) — Computer Networking (CACS303)

The Data Link Layer is the second layer of the OSI model, sitting directly above the Physical Layer. Where the Physical Layer is concerned only with moving raw, unstructured bits across a medium, the Data Link Layer takes those bits and organizes them into meaningful units called frames, while also dealing with the practical realities of a shared or noisy channel: how do multiple devices take turns transmitting, and how do we detect (or even fix) the errors that inevitably creep in? This unit works through both halves of that problem — point-to-point link control, and multi-access channel arbitration — before surveying the real-world standards (Ethernet, Wi-Fi, Token Ring, VLAN) built on top of them.

3.1 Functions of the Data Link Layer

The Data Link Layer provides node-to-node delivery of data across a single link — that is, between two devices directly connected by a physical medium (or sharing one, as on a broadcast LAN). It converts the Physical Layer's unreliable, unstructured bit stream into a link that appears, to the layer above, to be relatively free of undetected transmission errors.

- Framing — dividing the bit stream into manageable, addressable units called frames, marking where each frame begins and ends.
- Physical Addressing — attaching a header to each frame identifying the sender and/or receiver, typically using MAC (hardware) addresses.
- Flow Control — preventing a fast sender from overwhelming a slow receiver's buffer capacity.
- Error Control — detecting, and where possible correcting, bit errors introduced during transmission, typically via checksums and acknowledgement/retransmission schemes.
- Access Control — on a link shared by multiple devices, determining which device has the right to transmit at any given moment.

Sub-layer	Responsibility
Logical Link Control (LLC)	Framing, flow control, error control — independent of the physical medium
Media Access Control (MAC)	Addressing and channel-access arbitration — specific to the medium (e.g. Ethernet, Wi-Fi)

Note: The IEEE 802 project formally splits the Data Link Layer into these two sub-layers, which is why you will see both 'LLC' and 'MAC' referenced throughout the rest of this unit, particularly in the LAN standards.

3.2 Data Link Control: Framing, Flow and Error Control

3.2.1 Framing

Framing solves a deceptively simple problem: given a continuous stream of bits arriving from the Physical Layer, how does the receiver know where one frame ends and the next begins?

Framing Method	How It Works
Character (Byte) Count	A length field at the start of the frame states exactly how many bytes follow
Character/Byte Stuffing	Special flag bytes (e.g. Flag) mark frame boundaries; an escape byte is inserted before any flag byte that appears inside the data
Bit Stuffing	A specific bit pattern (e.g. 01111110) marks boundaries; a 0 is stuffed after every five consecutive 1s in the data to avoid a false boundary
Physical Layer Coding Violations	Certain physical-layer signal encodings reserve invalid symbol patterns purely to mark frame boundaries

Example — Bit Stuffing

Flag pattern: 01111110

Original data: 1101 1111 1111 1101

After stuffing: 1101 11111 0 111 1 0 1101 (a 0 is inserted after every run of five 1s)

The receiver reverses the process: after seeing five consecutive 1s, it removes the next bit if that bit is a 0, restoring the original data exactly.

The byte-count method is simple but fragile — if the length field itself is corrupted, the receiver loses synchronization for the rest of the stream. Stuffing-based methods are self-synchronizing: even after an error, the receiver can resynchronize as soon as it sees the next valid flag pattern, which is why they dominate in practice (HDLC and PPP both use flag-based framing, discussed in Section 3.4).

3.2.2 Flow Control

Flow control regulates the rate at which a sender transmits so that it does not overwhelm a receiver that may process, buffer, or forward data more slowly than data arrives.

Scheme	Mechanism
Stop-and-Wait	Sender transmits one frame, then waits for an acknowledgement (ACK) before sending the next
Sliding Window	Sender may transmit up to W frames before requiring an ACK, keeping the link fuller and improving throughput

Stop-and-Wait is simple to implement but wastes significant time on high-latency links, since the sender is idle while waiting for each acknowledgement. Sliding-window protocols, such as Go-Back-N and Selective Repeat, allow several frames to be "in flight" simultaneously, dramatically improving link utilization — especially over links with a large bandwidth-delay product.

3.2.3 Error Control

Error control ensures that frames are delivered to the destination correctly and, in connection-oriented service, exactly once, in the correct order. It is built from three cooperating mechanisms.

- Error detection — the receiver can tell, with high probability, whether a received frame was corrupted in transit (see Section 3.3).
- Acknowledgement (ACK/NAK) — the receiver explicitly confirms correct receipt, or signals that a frame must be resent.
- Retransmission (ARQ) — if an ACK is not received within a timeout period, or a NAK is received, the sender automatically resends the frame — this overall strategy is called Automatic Repeat reQuest (ARQ).

Key Takeaways — 3.1 & 3.2

- ✓ The Data Link Layer provides framing, addressing, flow control, error control, and access control.
- ✓ It is conceptually split into the LLC (medium-independent) and MAC (medium-specific) sub-layers.
- ✓ Bit/byte stuffing framing is self-synchronizing, unlike simple byte-count framing.
- ✓ Stop-and-Wait is simple but slow; sliding-window flow control improves throughput on long links.
- ✓ Error control combines detection, acknowledgement, and ARQ-based retransmission.

3.3 Error Detection and Correction

No physical medium is perfectly reliable — electrical noise, electromagnetic interference, and signal attenuation can all flip bits in transit. The Data Link Layer must be able to at least detect, and sometimes correct, such errors before passing data up to higher layers.

3.3.1 Types of Errors

- Single-bit error — only one bit in the data unit changes (0 to 1, or 1 to 0). Common on links with low, evenly distributed noise.
- Burst error — two or more consecutive bits are corrupted. Far more common in practice, since noise on real media (e.g. lightning, motor interference) tends to affect a duration of time rather than a single bit.

3.3.2 Error Detection Techniques

Technique	Idea
Parity Check	One extra bit makes the total count of 1-bits even (or odd); catches all single-bit errors but not all burst errors
Checksum	Data is summed in fixed-size blocks; the complement of the sum is sent and re-verified at the receiver (used in IP, TCP, UDP headers)
Cyclic Redundancy Check (CRC)	Data is treated as a binary polynomial and divided by a fixed generator polynomial; the remainder is appended as the CRC field

Example — CRC Calculation (Simplified)

Data: 1101011011, Generator: 10011 (degree 4, so append four 0-bits)

Dividend: 11010110110000

Perform binary (XOR) division of the dividend by the generator.

The 4-bit remainder becomes the CRC checksum, appended to the original data before transmission.

How the Receiver Checks It

The receiver divides the ENTIRE received frame (data + appended CRC) by the same generator polynomial.

If the remainder is exactly zero, the frame is accepted as (very likely) error-free.

If the remainder is non-zero, the frame is discarded, and — depending on the protocol — a retransmission is requested.

Note: CRC is far more powerful than simple parity or even checksums: with a well-chosen generator polynomial, CRC reliably detects all single-bit errors, all double-bit errors, any odd number of errors, and virtually all burst errors up to the length of the generator polynomial.

3.3.3 Error Correction: Hamming Code

While CRC is excellent at detection, it cannot tell the receiver which bit is wrong. Forward Error Correction (FEC) techniques, such as the Hamming Code, add enough redundant bits that the receiver can pinpoint — and flip back — a single-bit error without any retransmission.

Hamming code places parity bits at power-of-two positions (1, 2, 4, 8, ...) within the transmitted word. Each parity bit covers a specific subset of positions, determined by the binary representation of the position index. If an error occurs, the pattern of failed parity checks at the receiver forms a binary number that directly points to the exact position of the corrupted bit.

Approach	Detects	Corrects	Overhead
Parity Bit	Single-bit errors only	No	Very low (1 bit)
Checksum	Most burst and random errors	No	Low
CRC	Virtually all common error patterns	No	Low-moderate
Hamming Code	Single- and double-bit errors	Single-bit errors	Moderate ($\log_2(n)$ bits)

Key Takeaways — 3.3

- ✓ Errors are classified as single-bit or burst; burst errors are more common on real media.
- ✓ Parity, checksum, and CRC are error DETECTION techniques of increasing strength.
- ✓ CRC works by polynomial division; a zero remainder at the receiver indicates (likely) no error.
- ✓ Hamming code is a FORWARD ERROR CORRECTION technique that can locate and fix single-bit errors.
- ✓ Stronger error handling generally costs more redundant bits — there is always a trade-off.

3.4 HDLC and Point-to-Point Protocol (PPP)

3.4.1 High-Level Data Link Control (HDLC)

HDLC is a bit-oriented Data Link Layer protocol, standardized by ISO, that provides reliable data transfer between two nodes over a point-to-point or multipoint link. It uses flag-based bit stuffing (Section 3.2.1) to delimit frames and defines a single, uniform frame format for every kind of frame it sends.

Field	Size (bits)	Purpose
Flag	8	Marks the start and end of the frame (01111110)

Field	Size (bits)	Purpose
Address	8 or more	Identifies the secondary station on multipoint links
Control	8 or 16	Identifies the frame TYPE and carries sequence numbers
Information	Variable	The actual payload (present only in I-frames)
FCS	16 or 32	Frame Check Sequence — a CRC used for error detection
Flag	8	Marks the end of the frame (01111110)

HDLC defines three categories of frame, distinguished by the first one or two bits of the Control field.

- I-frames (Information) — carry actual user data, along with piggybacked sequence numbers for flow and error control.
- S-frames (Supervisory) — carry no user data; used purely for flow and error control, e.g. acknowledging receipt (RR — Receive Ready) or requesting retransmission (REJ — Reject).
- U-frames (Unnumbered) — used for link management: establishing or terminating a connection, and reporting certain error conditions.

HDLC supports three transfer modes: Normal Response Mode (NRM), where a secondary station may transmit only when explicitly permitted by the primary; Asynchronous Response Mode (ARM), where a secondary may transmit without explicit permission; and Asynchronous Balanced Mode (ABM), the most common in practice, where every station is a peer "combined" station able to initiate transmission at will.

3.4.2 Point-to-Point Protocol (PPP)

PPP is the dominant Data Link Layer protocol used to carry multi-protocol traffic across simple point-to-point links — most classically, a dial-up connection to an ISP, but also DSL and many leased-line and VPN links today. PPP's frame format is directly derived from HDLC, but simplified for the point-to-point case (no addressing of multiple stations is required).

Field	Purpose
Flag	01111110, marks frame boundaries (as in HDLC)
Address	Fixed value 11111111 — broadcast, unused for real addressing
Control	Fixed value 00000011 — indicates an unnumbered frame
Protocol	Identifies which network-layer protocol the payload belongs to (e.g. IP)
Payload	The encapsulated packet from the network layer

Field	Purpose
FCS	16-bit CRC for error detection

Because PPP must run over a huge variety of physical links and needs to carry many different network-layer protocols, it is built from three cooperating components.

- Encapsulation — defines how packets from various network-layer protocols are framed for transmission.
- Link Control Protocol (LCP) — establishes, configures, tests, and terminates the data-link connection itself.
- Network Control Protocols (NCPs) — a family of protocols, one per network-layer protocol (e.g. IPCP for IP), that configure the specifics of carrying that particular protocol once LCP has brought the link up.

Note: A useful way to remember the difference: HDLC is a general bit-oriented framing protocol usable point-to-point or multipoint, whereas PPP is a specialised, simplified descendant of HDLC purpose-built for exactly one scenario — a single point-to-point link carrying IP (or other network-layer) traffic.

Key Takeaways — 3.4

- ✓ HDLC is a bit-oriented protocol with I-frames (data), S-frames (control), and U-frames (management).
- ✓ HDLC operates in NRM, ARM, or ABM transfer modes depending on the station relationship.
- ✓ PPP is derived from HDLC but simplified for point-to-point links, e.g. dial-up and DSL.
- ✓ PPP combines an encapsulation format, LCP (link setup), and per-protocol NCPs (e.g. IPCP for IP).

3.5 The Channel Allocation Problem

On a point-to-point link, only two devices share the medium, so deciding who transmits when is trivial. The channel allocation problem arises specifically on a broadcast (multi-access) channel — such as a wireless network or a shared Ethernet segment — where many stations compete for the same finite bandwidth, and an arbitration strategy is needed to decide who gets to transmit, and when.

3.5.1 Static Channel Allocation

The traditional solution divides the available bandwidth into N fixed portions using FDM (Frequency Division Multiplexing) or TDM (Time Division Multiplexing), permanently assigning one portion to each of N stations.

- Advantage — simple, and collisions are impossible by construction.
- Disadvantage — badly wastes capacity when traffic is bursty: a station with nothing to send still "owns" its slice of the channel, which sits idle, while another station with data to send may be forced to wait even though capacity is available elsewhere.

Static allocation performs reasonably only when the number of stations is small and constant, and traffic is smooth and continuous (e.g. telephone circuits) — assumptions that do not hold for typical, bursty computer network traffic.

3.5.2 Dynamic Channel Allocation

Because most computer traffic is bursty, far greater efficiency is achieved by allocating the channel dynamically, on demand, rather than statically pre-assigning fixed slices to fixed stations. Several key assumptions typically underpin dynamic-allocation models.

Assumption	Meaning
Independent traffic	Frame arrivals at each station are independent of other stations
Single channel	All stations share exactly one communication channel
Observable collisions	If two frames overlap in time, all stations can detect the resulting collision
Continuous or slotted time	Transmissions may begin at any instant, or only at fixed slot boundaries
Carrier sense (optional)	Stations may or may not be able to sense whether the channel is currently busy

The multiple-access protocols surveyed in Section 3.6 (ALOHA, CSMA, token-based schemes, FDMA/TDMA/CDMA) are all, in effect, different practical solutions to this same dynamic channel allocation problem, each trading off simplicity, fairness, and efficiency in different ways.

Note: The channel allocation problem is sometimes framed as a special case of resource contention: it is conceptually the same challenge faced whenever many independent actors compete for one shared, finite resource.

3.6 Multiple Access Protocols

Multiple access protocols fall into three broad families, matching the classification used throughout this section.




Fig. 3.1 — The three families of multiple access protocol and their principal members.


3.6.1 Random Access Protocols

In random access, no station has priority over any other, and no station is scheduled to transmit at a particular time. Each station decides for itself, based on some protocol-defined rule, when it may attempt to transmit — which means collisions are always possible and must be handled.

ALOHA, developed at the University of Hawaii, is the earliest and simplest random access scheme.

Variant	Rule	Maximum Throughput
Pure ALOHA	Transmit whenever a frame is ready, with no coordination at all	$\approx 18.4\%$
Slotted ALOHA	Transmission may only begin at the start of a fixed time slot	$\approx 36.8\%$

Pure ALOHA — frames transmitted at arbitrary times
(shaded zones = vulnerable period, width = $2 \times$ frame time)



Slotted ALOHA — transmissions align to fixed time slots
(vulnerable period reduced to $1 \times$ frame time)




Fig. 3.2 — Pure ALOHA's vulnerable period spans two full frame times; aligning transmissions to slots in Slotted ALOHA halves it to one.

Slotted ALOHA doubles pure ALOHA's maximum throughput simply by forcing all stations to align their transmission attempts to slot boundaries, which shrinks the vulnerable period — the window during which a collision can occur — from two frame times down to just one.

3.6.2 Carrier Sense Multiple Access (CSMA)

CSMA improves on ALOHA with one simple but powerful rule: listen before you talk. A station senses the channel before transmitting, and only transmits if the channel is idle — dramatically reducing (but not eliminating) collisions, since two stations can still both sense an idle channel and transmit simultaneously due to propagation delay.




Fig. 3.3 — Three strategies a station may use once it finds the channel busy or has just sensed it become idle.

Persistence Strategy	Behaviour
1-Persistent	Sense continuously; transmit immediately (probability 1) the instant the channel is idle — simple, but leads to near-certain collision if two stations were both waiting
Non-Persistent	If busy, wait a RANDOM period before sensing again, rather than sensing continuously — reduces collisions but adds idle time
p-Persistent (slotted channels)	If idle, transmit with probability p , or defer to the next slot with probability $(1 - p)$ — tunable trade-off between the two extremes

3.6.3 CSMA/CD (Collision Detection)

CSMA/CD extends CSMA with the ability to detect a collision while it is happening, and abort the transmission immediately rather than wasting the rest of the frame time — this is the access method that classic (shared-medium) Ethernet is built on.

- 1. Sense the channel; if busy, defer using the chosen persistence strategy.

- 2. Transmit and simultaneously continue to listen to the channel for a collision (a signal that does not match what was sent).
- 3. If a collision is detected, immediately stop transmitting and send a short jam signal so all stations recognize the collision.
- 4. Back off for a random period (commonly using the binary exponential back-off algorithm) before attempting to retransmit.

Note: CSMA/CD requires that a station be able to transmit and receive simultaneously, and it only works reliably when the frame transmission time is long enough for a collision at the far end of the network to propagate back before transmission finishes — this constraint directly shaped classic Ethernet's minimum frame size of 64 bytes.

3.6.4 CSMA/CA (Collision Avoidance)

On a wireless medium, a transmitting station generally cannot listen for a collision while it is sending (its own strong outgoing signal drowns out any weak incoming one), so collision detection is not practical. CSMA/CA, used by Wi-Fi (Section 3.8), instead focuses on avoiding collisions before they occur.

- Interframe spacing (IFS) — mandatory idle gaps between transmissions give every station a fair, predictable opportunity to sense the channel.
- Random back-off — even on a sensed-idle channel, a station waits a random number of slot times before transmitting, reducing the chance that two stations transmit at the same instant.
- Optional RTS/CTS handshake — Request To Send / Clear To Send frames reserve the channel in advance, addressing the "hidden terminal" problem where two stations cannot hear each other directly.

3.6.5 Controlled Access Protocols

Controlled access eliminates collisions altogether (or nearly so) by coordinating which station is allowed to transmit at any given moment, trading some efficiency and added complexity for predictability.

Method	How It Works
Reservation	Time is divided into intervals; each interval begins with a short reservation sub-interval during which stations announce their intent to transmit
Polling	A designated primary/controller station asks each secondary station, in turn, whether it has data to send

Method	How It Works
Token Passing	A special control frame (the token) circulates among stations; only the station currently holding the token may transmit

Polling requires a controller and works well in a centralized (master/slave) topology, but the controller itself becomes both a bottleneck and a single point of failure. Token passing removes the need for a central controller and guarantees every station a fair, bounded opportunity to transmit — the mechanism underlying Token Ring and Token Bus, described in Section 3.9.

3.6.6 Channelization Protocols

Channelization protocols divide the available bandwidth of a link sharing among multiple stations by dimension — frequency, time, or code — so that each station effectively gets its own private sub-channel, avoiding collisions structurally rather than through arbitration.

Method	Shared Resource Divided By	Typical Use
FDMA (Frequency Division)	Frequency band — each station gets a distinct frequency range	Analogue radio and TV broadcasting, early cellular systems
TDMA (Time Division)	Time — each station gets a fixed, recurring time slot on the same frequency	GSM cellular telephony, satellite links
CDMA (Code Division)	A unique orthogonal code assigned to each station; all stations share the same time and frequency simultaneously	3G cellular telephony, some Wi-Fi/military applications

CDMA is the most elegant of the three, in one sense: rather than physically partitioning the channel, every station transmits across the full bandwidth at the same time, but each station's data is encoded with a unique orthogonal spreading code. A receiver can extract one station's signal from the combined transmission purely by knowing that station's code — even though, to any receiver without the right code, the combined signal looks like noise.

Key Takeaways — 3.6

- ✓ Random access (ALOHA, CSMA, CSMA/CD, CSMA/CA) allows any station to attempt transmission at will, risking collisions.
- ✓ Slotted ALOHA doubles pure ALOHA's throughput by halving the vulnerable period.
- ✓ CSMA/CD detects and aborts on collision (wired Ethernet); CSMA/CA avoids collisions proactively (Wi-Fi).
- ✓ Controlled access (reservation, polling, token passing) coordinates access to avoid collisions structurally.
- ✓ Channelization (FDMA, TDMA, CDMA) divides the physical resource itself so stations never contend at all.

3.7 Wired LAN: Ethernet Standards and FDDI

3.7.1 Ethernet Evolution

Ethernet, standardized as IEEE 802.3, is by far the most widely deployed wired LAN technology. It began as a shared-medium, CSMA/CD-based bus network and has evolved substantially over several generations.

Standard	Speed	Key Characteristic
Standard Ethernet (10BASE5/2)	10 Mbps	Shared coaxial bus; original CSMA/CD access method
Fast Ethernet (100BASE-TX)	100 Mbps	Twisted-pair cabling, star topology via hubs/switches
Gigabit Ethernet (1000BASE-T)	1 Gbps	Full-duplex switched links make CSMA/CD largely unnecessary
10 Gigabit Ethernet and beyond	10 Gbps+	Almost exclusively fibre or high-grade copper; fully switched

A crucial shift occurred as Ethernet moved from hubs (which broadcast every bit to every port, recreating a shared medium electrically) to switches (which forward frames only to the correct destination port over a dedicated, full-duplex link). On a modern switched, full-duplex Ethernet link, both ends can transmit simultaneously with no possibility of collision, which is why CSMA/CD — while still part of the formal 802.3 standard for compatibility — is effectively unused in most modern deployments.

3.7.2 The Ethernet Frame Format

Field	Size	Purpose
Preamble	7 bytes	Alternating 1s and 0s; allows receiver clock synchronization
Start Frame Delimiter	1 byte	Marks the start of the actual frame (10101011)
Destination Address	6 bytes	MAC address of the intended recipient
Source Address	6 bytes	MAC address of the sender
Length/Type	2 bytes	Payload length, or identifies the encapsulated protocol
Data & Pad	46–1500 bytes	Payload; padded if necessary to meet the 64-byte minimum frame size
FCS	4 bytes	CRC-32 checksum for error detection

Note: The 64-byte minimum frame size exists specifically to support CSMA/CD: it guarantees a frame takes long enough to transmit that a collision anywhere on the network will be detected by the sender before it finishes sending, which is essential on a shared (non-switched) segment.

3.7.3 FDDI (Fiber Distributed Data Interface)

FDDI is a high-speed (100 Mbps) LAN standard built on fibre-optic cabling, arranged as a dual counter-rotating ring. It uses token passing (Section 3.9) rather than CSMA/CD for access control, which — unlike Ethernet's probabilistic access — gives every station a deterministic, bounded wait time before it may transmit.

- The dual-ring design provides automatic fault tolerance: if the primary ring breaks, traffic can be wrapped onto the secondary ring to bypass the fault.
- FDDI was historically popular as a campus or building backbone, interconnecting many lower-speed Ethernet segments, before high-speed switched Ethernet made it largely obsolete.

3.8 Wireless LAN: IEEE 802.11x and Bluetooth Standards

3.8.1 IEEE 802.11 (Wi-Fi)

The IEEE 802.11 family of standards defines wireless LAN (Wi-Fi) operation. Because a wireless medium cannot reliably support collision detection (Section 3.6.4), 802.11 relies on CSMA/CA as its fundamental access method.

Amendment	Approx. Max Speed	Frequency Band
802.11b	11 Mbps	2.4 GHz

Amendment	Approx. Max Speed	Frequency Band
802.11a	54 Mbps	5 GHz
802.11g	54 Mbps	2.4 GHz
802.11n (Wi-Fi 4)	600 Mbps	2.4 / 5 GHz
802.11ac (Wi-Fi 5)	~3.5 Gbps	5 GHz
802.11ax (Wi-Fi 6/6E)	~9.6 Gbps	2.4 / 5 / 6 GHz

802.11 networks operate in one of two modes: infrastructure mode, where all stations communicate through a central Access Point (AP) that typically bridges to a wired network, and ad hoc mode, where stations communicate directly with one another with no AP involved.

- Hidden terminal problem — two stations, both in range of the AP but out of range of each other, cannot sense each other's transmissions and may collide at the AP without either realizing it.
- RTS/CTS handshake — the optional Request-to-Send / Clear-to-Send exchange lets a station reserve the channel in advance, which every station within range of the AP will hear and defer to, mitigating the hidden terminal problem.

3.8.2 Bluetooth

Bluetooth (standardized as IEEE 802.15.1, though now primarily governed by the independent Bluetooth SIG) is designed for short-range, low-power Personal Area Networks (PANs) — connecting peripherals such as headphones, keyboards, and fitness trackers rather than building general-purpose LANs.

Characteristic	Bluetooth
Typical range	Up to ~10 m (Class 2), some devices reach ~100 m (Class 1)
Frequency band	2.4 GHz ISM band
Channel access	Frequency Hopping Spread Spectrum (FHSS) — rapidly hops among 79 channels
Network unit	Piconet — one master device plus up to 7 active slave devices
Larger networks	Scatternet — multiple overlapping piconets, linked via shared devices

Note: Frequency hopping gives Bluetooth strong resistance to interference — since a hop only stays on any one frequency for a fraction of a second, even significant narrowband interference on one channel affects only a brief portion of the overall transmission.

3.9 Token Bus, Token Ring, and Virtual LAN

3.9.1 Token Ring (IEEE 802.5)

In a Token Ring network, stations are physically (or logically) connected in a ring, and a special 3-byte control frame — the token — circulates continuously around the ring. A station may transmit data only once it has captured the free token, at which point it converts the token into the start of a data frame, transmits, and then releases a fresh token once its data has circulated the ring.

- Access is completely deterministic — every station is guaranteed a turn within a bounded, predictable time, unlike the probabilistic access of CSMA/CD.
- There are no collisions by design, since only the token holder may transmit at any given moment.
- A single broken link or failed station can, without added fault-tolerance mechanisms, disrupt the entire ring — a key structural weakness compared to switched star topologies.

3.9.2 Token Bus (IEEE 802.4)

Token Bus combines the physical topology of a bus (as in classic Ethernet) with the deterministic, collision-free access method of token passing. Stations are physically attached to a common bus, but the token circulates among them in a logical ring order defined by software, independent of their physical position on the cable.

Aspect	Token Ring	Token Bus
Physical topology	Ring	Bus
Logical topology	Ring	Ring (imposed in software over a bus)
Typical use	Office LANs (largely historical)	Industrial/factory-floor networks (deterministic timing)

Note: Token Bus's deterministic access made it attractive for real-time industrial control applications, where a bounded worst-case delay matters more than raw average throughput — this is why it saw adoption in manufacturing automation (e.g. the MAP standard) even after it lost ground to Ethernet in general-purpose office networking.

3.9.3 Virtual LAN (VLAN)

A VLAN is a logical grouping of devices into the same broadcast domain, regardless of their physical location or which physical switch port they are connected to. VLANs are implemented at the switch level and let network administrators segment a single physical LAN infrastructure

into multiple, isolated logical networks.

- **Broadcast containment** — broadcast traffic within one VLAN does not reach devices in a different VLAN, reducing unnecessary traffic and improving overall network performance.
- **Security isolation** — devices in different VLANs cannot communicate directly without passing through a router or Layer-3 switch, even if they share the same physical switch.
- **Flexible grouping** — devices can be grouped logically (e.g. by department: Sales, Engineering, Finance) rather than by physical wiring, and can be reassigned without any physical recabling.

VLAN membership is most commonly implemented by tagging Ethernet frames according to IEEE 802.1Q, which inserts a 4-byte VLAN tag (including a 12-bit VLAN ID, supporting up to 4094 usable VLANs) into the standard Ethernet frame header, allowing switches to keep each frame's VLAN membership intact as it traverses trunk links between switches.

Key Takeaways — 3.7–3.9

- ✓ Ethernet has evolved from a shared CSMA/CD bus into fully switched, full-duplex, collision-free links.
- ✓ FDDI uses dual counter-rotating fibre rings with token passing for fault-tolerant, deterministic access.
- ✓ 802.11 (Wi-Fi) uses CSMA/CA; Bluetooth uses frequency hopping for short-range personal-area networking.
- ✓ Token Ring and Token Bus both use token passing for deterministic, collision-free access, differing in physical topology.
- ✓ VLANs (IEEE 802.1Q) logically segment a physical LAN into isolated broadcast domains for security and efficiency.

Overall Unit 3 Summary

- ✓ The Data Link Layer frames raw bits, addresses nodes, and provides flow and error control (3.1-3.2).
- ✓ Parity, checksum, CRC (detection) and Hamming code (correction) protect data against transmission errors (3.3).
- ✓ HDLC and PPP are the classic bit-oriented protocols for reliable point-to-point and multipoint links (3.4).
- ✓ The channel allocation problem motivates dynamic, on-demand sharing of a broadcast medium (3.5).
- ✓ Random access (ALOHA/CSMA family), controlled access, and channelization are three distinct solutions to that problem (3.6).
- ✓ Ethernet/FDDI (wired) and 802.11/Bluetooth (wireless) are the dominant real-world LAN standards built on these ideas (3.7-3.8).
- ✓ Token Ring, Token Bus, and VLANs round out the picture with deterministic access and logical network segmentation (3.9).

— End of Unit 3 —