ABHYASMITRA.IN

STUDY NOTES

COMPUTER NETWORKS

Unit II — Data Link Layer

Topics Covered (09 Hours)

Services Provided to the Network Layer
Framing Techniques (Bit/Byte Stuffing & Violations)
Link Addressing Design Issues
Design Issues: Error Control & Flow Control
Error Detection: Parity, Checksum, and CRC Math
Error Correction Concepts (Hamming Codes)
Elementary Data Link Protocols & Stop-and-Wait
Sliding Window Protocols (Pipelining, Go-Back-N, Selective Repeat)
Case Study: Packet over SONET

Compiled in a simple, easy-to-understand, book style format
for students of Computer Science and Engineering

© Copyright — abhyasmitra.in — All Rights Reserved

TABLE OF CONTENTS

2.1 Data Link Layer Services & Framing

The Data Link Layer lies immediately above the Physical Layer and directly below the Network Layer. It is responsible for handling node-to-node (hop-by-hop) communication over a single physical link. Its main goal is to convert the raw bitstream provided by the physical layer into a structured, reliable transmission line. It intercepts packets from the network layer, encapsulates them into structures called frames, and manages access to the physical transmission medium.

Services Provided to the Network Layer

The primary service of the data link layer is transferring data from the network layer on the source machine to the network layer on the destination machine. It provides three primary service models:

- **Unacknowledged Connectionless Service:** The source machine sends independent frames to the destination without establishing a connection. The receiver does not send back acknowledgments. If a frame is lost due to noise, no attempt is made to detect or recover it. This service is ideal for reliable physical links (such as Ethernet) or real-time voice/video systems where speed is preferred over complete reliability.
- **Acknowledged Connectionless Service:** No logical connection is established, but each frame sent is individually acknowledged by the receiver. The sender maintains a timer; if a frame does not arrive within a specific timeout, the sender retransmits it. This is useful for unreliable media, such as wireless Wi-Fi (IEEE 802.11) channels, where local retransmissions are much faster than waiting for higher layers to time out.
- **Acknowledged Connection-Oriented Service:** The source and destination establish a connection before any data is transferred. Every frame is numbered, acknowledged, and guaranteed to arrive exactly once and in the correct order. The connection lifecycle involves setup, data transfer, and teardown.

Detailed Study of Framing Methods

To transmit packets, the data link layer must split the physical layer's continuous bitstream into distinct frames. The receiver must be able to detect the exact start and end of each frame. There are four primary framing techniques used in practice:

1. Character Count

This technique uses a field in the frame header to specify the total number of characters (bytes) in the frame. When the receiver sees the character count, it knows how many bytes follow. However, this method is highly vulnerable to transmission noise. If a single bit in the count field flips during transmission, the count becomes incorrect. The receiver will look for the end of the frame at the wrong position, losing synchronization. The receiver will interpret data bytes as count fields, leading to catastrophic framing errors.

2. Flag Bytes with Byte Stuffing

This method marks the beginning and end of each frame with a special flag byte (typically '0x7E'). If the flag byte occurs inside the actual data (payload), the sender's software inserts a special escape byte ('ESC', e.g., '0x7D') immediately before it. If an ESC byte occurs in the data, it is also prefixed with an ESC byte. This is called Byte Stuffing. When the receiver sees an ESC byte, it discards it and treats the following byte as raw data, successfully detecting frame boundaries.

3. Starting and Ending Flags with Bit Stuffing

To support arbitrary byte sizes, bit-level framing is used. Each frame starts and ends with a specific bit pattern: '01111110' (six consecutive 1s). Whenever the sender's hardware detects five consecutive 1s in the data stream, it automatically stuffs a '0' bit into the outgoing transmission. When the receiver detects five consecutive 1s followed by a '0' bit, it automatically deletes the 0 bit. If the receiver sees six consecutive 1s, it knows it has reached the frame flag boundary. This is called Bit Stuffing.

Example: If the sender has data '011111101', it will transmit '0111110101' (with a stuffed 0 after the five 1s). The receiver will remove the 0, restoring the original data.

4. Physical Layer Coding Violations

This technique is only possible if the physical medium uses redundant coding. In Manchester encoding, for example, every bit has a transition in the middle (e.g., High-to-Low represents 1, Low-to-High represents 0). This means that every bit interval contains a transition. We can use periods *without* transitions (High-High or Low-Low) to mark frame boundaries. Since these violate the code rules, they are easily recognized and do not conflict with the actual data bits.

2.2 Design Issues: Flow Control & Error Control

To maintain reliable communication, the data link layer must implement rules for flow control, error control, and link addressing.

Error Control & Flow Control Principles

Error control is the process of ensuring that all frames are delivered without corruption or loss, in the correct order. Flow control ensures that a fast sender does not overwhelm a slow receiver. If the receiver has limited buffer space or a slow CPU, a fast sender will flood the receiver, causing buffer overflows and packet drops.

Flow Control Strategies

- **Feedback-Based Flow Control:** The sender is allowed to transmit data only after receiving permission or an acknowledgment from the receiver. The receiver controls the transmission rate dynamically.
- **Rate-Based Flow Control:** The protocol restricts the sender to a maximum transmission speed, defined in bytes or packets per second. This is rarely used at the link layer and is more common in transport layer congestion control.

Addressing at the Data Link Layer

When multiple hosts share a single physical medium (like Wi-Fi or Ethernet), each node must have a unique physical address, known as a MAC (Media Access Control) address. The MAC address is a 48-bit hardware address burned into the network interface card (NIC). The frame header contains the source and destination MAC addresses, allowing nodes to filter out frames not addressed to them.

2.3 Error Detection and Correction Math

We add redundant bits to our data to identify or fix errors. To design these codes, we use modular arithmetic and binary polynomials.

Error Detection Math

- **Checksums:** In internet protocols, data is split into 16-bit blocks. The blocks are added using 1s complement arithmetic. The sum is then complemented (inverted) and appended to the header. The receiver performs the same addition; if the result is all 1s (0xFFFF), the data is clean. Checksums are simple but do not protect against swapped byte errors.
- **Cyclic Redundancy Check (CRC):** Also called Polynomial Codes. The data bitstring is represented as a polynomial (e.g., '110001' represents $D(x) = x^5 + x^4 + 1$). The sender and receiver agree on a Generator Polynomial $G(x)$ of degree r. To calculate the CRC, the sender appends r zeros to $D(x)$ and divides it by $G(x)$ using modulo-2 arithmetic (where addition and subtraction are equivalent to XOR). The remainder of this division is the CRC checksum, which is appended to the frame. The receiver divides the entire received frame by $G(x)$; if the remainder is zero, the frame is error-free.

CRC Math Example

Let the data be '1101001110' (10 bits) and the generator polynomial $G(x)$ be '10011' (degree $r = 4$). We append 4 zeros to the data, giving '11010011100000'. We perform modulo-2 division by '10011'. The remainder is '0100'. The sender transmits '11010011100100'. The receiver divides this by '10011', obtaining a remainder of 0, validating the frame.

Error Correction (Hamming Code)

The Hamming Distance is the number of bit positions in which two binary codewords differ. To detect d errors, a code must have a minimum Hamming distance of $d+1$. To correct d errors, the code must have a minimum distance of $2d+1$. This is because we can map corrupted codewords back to the closest valid codeword without ambiguity.

In a Hamming Code, parity bits are placed at positions that are powers of 2 (1, 2, 4, 8, etc.). The remaining positions (3, 5, 6, 7, 9, etc.) contain the data bits. Each parity bit calculates the parity of a subset of data bits. When the receiver checks the parity bits, if some fail, the sum of their positions gives the exact index of the corrupted bit. For example, if parity checks 1 and 4 fail, the error is at bit position 5 ($1 + 4 = 5$).

2.4 Data Link Layer Protocols

We study data link protocols using simple abstract models. We assume peer entities run as independent software processes on host machines.

Elementary Protocols

- **Utopian Simplex Protocol:** An ideal protocol. Data is transmitted in one direction only. The receiver's buffer is infinite, processing speed is infinite, and the channel is completely noise-free. No flow or error control is needed.
- **Simplex Stop-and-Wait Protocol:** Introduced to handle flow control. The sender transmits a single frame and must wait for an acknowledgment (ACK) from the receiver before sending the next frame.
- **Stop-and-Wait ARQ (Noisy Channels):** Handles both flow and error control. The sender numbers each frame (0 or 1 sequence numbers). If a frame is corrupted or lost, the receiver discards it, the sender's timer expires, and the sender retransmits the frame. Sequence numbers prevent duplicate frames if ACKs are lost.

© abhyasmitra.in

2.5 Sliding Window Protocols

Stop-and-wait is highly inefficient on long-distance or high-bandwidth links. The sender spends most of its time waiting for ACKs. Sliding Window Protocols allow the sender to transmit multiple frames before waiting for an acknowledgment. The sender and receiver maintain window boundaries to track outstanding frames.

Pipelining

Pipelining is the technique of keeping the transmission pipe full of frames. If a link has a round-trip propagation delay of 100ms, and we can transmit a frame in 1ms, we can send up to 100 frames before receiving the first ACK, utilizing 100% of the channel capacity.

Two Main Pipelining Recovery Protocols

- **1. Go-Back-N (GBN):** The sender can transmit up to W frames without waiting. The receiver only accepts frames in strict sequential order. If frame n is lost or corrupted, the receiver discards it and all subsequent frames (n+1, n+2, etc.). The sender's timer for frame n will eventually expire, and the sender must retransmit all outstanding frames starting from n.
- **2. Selective Repeat (SR):** The sender can transmit up to W frames. The receiver has a buffer to store out-of-order frames that arrive after a lost frame. The receiver sends a negative acknowledgment (NAK) for the lost frame n, or the sender's timer for n expires. The sender retransmits only the single lost frame n, avoiding redundant retransmissions.

© abhyasmitra.in

2.6 Case Study: Packet over SONET

Packet over SONET (PoS) is a major data link technology used in wide area telecommunications networks. SONET (Synchronous Optical Network) is a physical layer standard for transmitting high-speed digital signals over optical fibers.

PoS acts as the Data Link Layer protocol that encapsulates IP packets directly into SONET frames. It bypasses complex ATM layers, reducing overhead. PoS uses PPP (Point-to-Point Protocol) for framing and link control, and HDLC (High-Level Data Link Control) flag stuffing ('0x7E') to mark frame boundaries. This technology allows routers to send massive amounts of internet traffic directly over high-speed optical trunk lines at speeds of tens of Gigabits per second.

© abhyasmitra.in

Unit Summary

This unit explored the design and protocols of the Data Link Layer.

- The Data Link Layer provides services to the network layer, converting raw physical channels into error-free links.
- Framing methods include character count, byte stuffing (flag bytes), bit stuffing (0-bit insertion after five 1s), and physical violations.
- Design issues focus on flow control (preventing sender flooding) and error control (handling corruption and loss).
- Error detection uses parity, internet checksums, or Cyclic Redundancy Checks (CRC/polynomial codes).
- Error correction (Hamming Code) locates flipped bits using redundant parity bits placed at power-of-2 positions.
- Stop-and-wait requires waiting for an ACK after every frame, which is inefficient for high-bandwidth-delay products.
- Sliding window protocols pipeline frames. Go-Back-N retransmits all frames from the lost frame, while Selective Repeat buffers out-of-order frames and retransmits only the lost one.
- Packet over SONET (PoS) encapsulates IP packets over optical fibers using PPP and HDLC framing.

© Copyright — abhyasmitra.in — All Rights Reserved