

VU LEARNING HUB

vulearninghub.xyz

HANDOUT

CS716

Advanced Computer Networks

CS716 by VU Learning HUB

Prepared for VU students • Free study resources

<https://vulearninghub.xyz>

CS716

Advanced Computer Networks

STUDY HANDOUT

A complete topic-wise study handout for CS716 Advanced Computer Networks, prepared from the course video lectures by Dr. Amir Qayyum (Virtual University, 45 video lectures, also on VU's official YouTube channel). Covers network foundations, multiplexing, framing, error control, LANs, switching, IP internetworking, routing, IPv6, multicast, MPLS, TCP, congestion control, QoS and network security - with definitions, explanations, examples and exam pointers.

VU Learning HUB | vulearninghub.xyz

Table of Contents

1. Introduction and Network Overview
2. Foundations and Building Blocks
3. Multiplexing
4. Performance Metrics
5. Network Architecture and Protocol Machinery
6. Internet Architecture and Network Models
7. Socket Programming and Client-Server Model
8. TCP and UDP
9. Direct Link Networks: Encoding and Modulation
10. Framing
11. Error Detection and Correction
12. Reliable Transmission
13. Shared Access Networks: Ethernet
14. Token Rings and FDDI
15. Wireless LANs and Spread Spectrum
16. Switching and Forwarding
17. Cell Switching: ATM
18. Internetworking: IP, Fragmentation, ARP, ICMP, DHCP
19. Virtual Private Networks
20. Routing Algorithms: Distance Vector and Link State
21. OSPF Routing Protocol
22. Inter-Domain Routing: BGP
23. IPv6
24. Multicast
25. Multiprotocol Label Switching (MPLS)
26. End-to-End Protocols and RPC
27. TCP: Reliable Byte Stream
28. Congestion Control
29. Quality of Service and RSVP
30. Network Security: Concepts and Cryptography
31. Security Protocols: IPsec, SSL/TLS, Firewalls
32. Mobile IP and Wireless Networking Trends
33. Revision: Midterm Practice Set
34. Revision: Final-Term Practice Set

Lecture 1: Introduction and Network Overview

Key terms: computer network, packet switching, circuit switching, Internet

A computer network connects autonomous computers to share resources and information. The two fundamental switching techniques: circuit switching (a dedicated path reserved for the call, e.g. telephone network - predictable but wasteful) and packet switching (data split into packets sharing links statistically, e.g. the Internet - efficient but with variable delay). The Internet is a global packet-switched network of networks speaking TCP/IP.

- Network types by scale: LAN (building/campus), MAN (city), WAN (country/continent).
- Performance concerns: bandwidth, latency, loss and jitter - the vocabulary of the whole course.
- Applications driving networks: web, email, streaming, cloud, IoT - all assume the Internet's best-effort service.

EXAM POINTER: Compare circuit switching and packet switching on resource reservation, efficiency and delay - the course's opening contrast.

Lecture 2: Foundations and Building Blocks

Key terms: nodes, links, protocols, layering, encapsulation

Networks are built from nodes (hosts, routers, switches) connected by links (copper, fibre, wireless). Communication is governed by protocols - agreed rules for format and exchange. Layering organises protocols into a stack where each layer uses the service of the one below and provides service to the one above; data is encapsulated with headers at each layer as it descends the stack.

- OSI 7 layers: Physical, Data Link, Network, Transport, Session, Presentation, Application.
- TCP/IP 4 layers: Network Access, Internet (IP), Transport (TCP/UDP), Application.
- Encapsulation: application data -> TCP segment -> IP packet -> Ethernet frame -> bits.

Lecture 3: Multiplexing

Key terms: multiplexing, FDM, TDM, STDM, WDM

Multiplexing lets multiple signals share one expensive link. Frequency Division Multiplexing (FDM) divides the bandwidth into frequency bands (radio stations); Time Division Multiplexing (TDM) divides time into slots in fixed rotation (classic telephony); Statistical TDM allocates slots on demand for bursty data; Wavelength Division Multiplexing (WDM) is FDM for fibre optics, multiplying fibre capacity enormously.

- FDM needs guard bands; TDM wastes slots on idle users; STDM suits bursty data.

- SONET/SDH uses TDM over fibre; DWDM carries 80+ wavelengths per fibre.
- Demultiplexing at the receiver reverses the process.

Lecture 4: Performance Metrics

Key terms: bandwidth, latency, RTT, bandwidth-delay product, throughput

Bandwidth is capacity in bits per second; latency (delay) is the time for a bit to travel, comprising propagation delay (distance/speed), transmission delay (size/bandwidth) and queueing delay. Round-trip time (RTT) matters for protocols. The bandwidth-delay product = bandwidth x RTT gives the number of bits 'in flight' - the pipe's true capacity, critical for window sizing in TCP.

- Throughput is the achieved rate - always $\leq$ bandwidth, reduced by congestion and protocol overhead.
- High bandwidth-delay product networks (satellite, transcontinental fibre) are 'long fat pipes'.
- Jitter (delay variation) hurts voice/video - the motivation for QoS.

EXAM POINTER: Numericals: compute transmission delay, RTT and bandwidth-delay product - standard midterm questions.

Lecture 5: Network Architecture and Protocol Machinery

Key terms: protocol, service, interface, peer processes, state machines

A protocol defines the format and rules of peer-to-peer exchange; a service is what a layer offers upward; the interface is how it is invoked. Protocol machinery includes: addressing (who), multiplexing/demultiplexing (which conversation), and state machines handling connection setup, data transfer and teardown. Formal specification (e.g. with finite state machines) exposes deadlocks and livelocks before deployment.

- Connection-oriented (TCP: setup, reliable, ordered) vs connectionless (UDP: no setup, best-effort).
- Request/reply, one-way and streaming are basic communication patterns.
- Layering pros: modularity, standardisation; cons: overhead, possible duplicated function.

Lecture 6: Internet Architecture and Network Models

Key terms: Internet architecture, hourglass model, end-to-end principle

The Internet's hourglass architecture converges everything onto IP at the waist: diverse link technologies below, diverse applications above, one universal network layer. The end-to-end principle places intelligence at hosts, keeping the network core simple and fast - new

applications deploy without changing the core. Standard bodies: IETF (RFCs), IEEE (LAN standards), ITU (telecom).

- IP is the 'narrow waist' - every host and router speaks it.
- End-to-end examples: encryption, reliable file transfer implemented at endpoints.
- Where the principle bends: NATs, firewalls and CDNs put function inside the network.

Lecture 7: Socket Programming and Client-Server Model

Key terms: socket, client-server, TCP socket, UDP socket, ports

Sockets are the programmer's interface to the network: an endpoint identified by (IP address, port). In the client-server model, the server passively listens on a well-known port while clients actively connect. TCP sockets give reliable byte streams (socket, bind, listen, accept, connect, send, recv); UDP sockets give message datagrams with no connection setup.

- Ports 0-1023 are well-known (HTTP 80, DNS 53); servers bind them, clients use ephemeral ports.
- Iterative vs concurrent servers: fork/thread per client or event-driven multiplexing.
- Blocking vs non-blocking I/O and select()/poll() for handling many connections.

Lecture 8: TCP and UDP

Key terms: TCP, UDP, ports, segments, datagrams, reliability

UDP is a thin wrapper on IP: ports for demultiplexing, an optional checksum, no reliability - used by DNS, DHCP, streaming and games where speed beats reliability. TCP provides a reliable, ordered, flow-controlled byte stream: connection establishment (three-way handshake), sequence/acknowledgement numbers, retransmission timers, flow control (advertised window) and congestion control.

- TCP header essentials: source/dest ports, sequence and ACK numbers, flags (SYN, ACK, FIN), window size.
- Three-way handshake: SYN -> SYN+ACK -> ACK; closing uses FIN exchanges.
- Choose UDP for low-latency tolerance of loss; TCP when every byte must arrive in order.

Lecture 9: Direct Link Networks: Encoding and Modulation

Key terms: encoding, NRZ, Manchester, modulation, ASK/FSK/PSK

Encoding maps bits to signals on a wire: NRZ (high=1, low=0) is simple but has baseline wander and clock-recovery problems; Manchester encodes clock in every bit (transition mid-bit) at the cost of doubling bandwidth; 4B/5B maps 4-bit groups to 5-bit codes avoiding long runs. Modulation impresses digital data on analogue carriers: ASK (amplitude), FSK

(frequency), PSK (phase) - and QAM combines amplitude and phase for high bit rates.

- Baud rate (symbols/sec) vs bit rate - multilevel signalling sends multiple bits per symbol.
- Nyquist and Shannon limits bound the achievable rate for a given bandwidth and SNR.
- Fibre uses light pulses; wireless uses radio modulation with complex constellations.

Lecture 10: Framing

Key terms: framing, HDLC, PPP, SONET, byte stuffing, bit stuffing

Framing marks where each packet begins and ends on a link. Approaches: byte-oriented (BISYNC sentinel bytes with byte stuffing - escape DLE before sentinel), bit-oriented (HDLC flag 01111110 with bit stuffing - insert 0 after five 1s), and clock-based (SONET's fixed 125-microsecond frames). Point-to-Point Protocol (PPP) frames IP over serial links with byte stuffing; SONET/SDH frames multiplexed tributaries over fibre.

- Bit stuffing example: data 01111110 becomes 011111010 on the wire.
- HDLC frame: flag, address, control, data, CRC, flag.
- SONET STS-1: 810 bytes per 125 us = 51.84 Mbps; STS-n multiplies by n.

EXAM POINTER: Apply bit stuffing to a given bit string - a mechanical question examiners love.

Lecture 11: Error Detection and Correction

Key terms: parity, checksum, CRC, Hamming code, ARQ

Errors are detected with redundancy: parity (one bit catches odd numbers of errors), checksums (Internet checksum sums 16-bit words), and CRC (cyclic redundancy check - treats data as a polynomial divided by a generator; catches all burst errors shorter than the generator). Correction needs more redundancy: Hamming codes add check bits locating the erroneous bit. ARQ (automatic repeat request) retransmits damaged frames: stop-and-wait, Go-Back-N and selective repeat.

- CRC strength depends on the generator polynomial (CRC-32 in Ethernet).
- Hamming distance d detects $d-1$ errors and corrects $\lfloor (d-1)/2 \rfloor$.
- Sliding-window protocols keep the pipe full: window size relates to bandwidth-delay product.

EXAM POINTER: Compute a CRC remainder for small polynomials, or find the Hamming check bits for a 4-bit message.

Lecture 12: Reliable Transmission

Key terms: stop-and-wait, Go-Back-N, selective repeat, sequence numbers

Reliable transmission over an unreliable link needs sequence numbers (to detect loss/duplication), acknowledgements, and timers. Stop-and-wait sends one frame per RTT - simple but utilisation = tiny on long fat pipes. Go-Back-N slides a window of N frames but retransmits everything from a lost frame. Selective repeat buffers out-of-order frames and retransmits only lost ones - efficient but complex; needs window $\leq 2^{(k-1)}$ for k-bit sequence numbers.

- Utilisation = window transmission time / RTT - the key performance formula.
- Piggybacking carries ACKs on reverse data frames.
- These link-layer ideas reappear in TCP (cumulative ACKs, fast retransmit).

Lecture 13: Shared Access Networks: Ethernet

Key terms: Ethernet 802.3, CSMA/CD, MAC address, exponential backoff

Classic Ethernet shares one cable among stations using CSMA/CD: listen before transmitting (carrier sense), transmit, and detect collisions - on collision, jam then retry after a random exponential backoff. Minimum frame size (64 bytes) guarantees a sender still transmitting when the collision signal returns. Modern switched Ethernet gives each port its own collision domain, making CSMA/CD legacy, but the frame format and MAC addressing remain.

- MAC addresses: 48-bit, globally unique, burned into NICs.
- Ethernet frame: preamble, dest/src MAC, type/length, payload (46-1500 B), CRC.
- Exponential backoff: after i collisions wait random(0, $2^i - 1$) slot times.

Lecture 14: Token Rings and FDDI

Key terms: token ring 802.5, FDDI, token passing

Token ring networks circulate a special token frame; only the holder may transmit, giving deterministic, fair access - attractive for real-time traffic. IBM Token Ring (802.5) ran on star-wired rings at 4/16 Mbps; FDDI (Fibre Distributed Data Interface) used dual counter-rotating fibre rings at 100 Mbps with automatic wrap-around on fibre cuts, popular as campus backbones before Gigabit Ethernet.

- Token holding time bounds each station's transmission - fairness by design.
- Ring maintenance: monitor station removes orphan frames and regenerates lost tokens.
- Why Ethernet won: cheaper, simpler, and switched Ethernet removed collisions.

Lecture 15: Wireless LANs and Spread Spectrum

Key terms: 802.11, CSMA/CA, spread spectrum, FHSS, DSSS

Wireless LANs (802.11/Wi-Fi) cannot reliably detect collisions (hidden terminal problem), so they avoid them with CSMA/CA: wait for idle channel, optional RTS/CTS handshake reserving the channel, then transmit and expect an ACK. Spread spectrum spreads signals across frequencies for interference resistance: FHSS hops between frequencies; DSSS multiplies data by a chipping code (used in 802.11b).

- Hidden terminal: A and C both reach B but not each other - RTS/CTS solves it; exposed terminal is the converse inefficiency.
- 802.11 frame has 4 address fields (source, destination, AP, etc.) unlike Ethernet's 2.
- Evolution: 802.11a/b/g/n/ac/ax (Wi-Fi 6) - wider channels, MIMO, OFDMA raise throughput.

Lecture 16: Switching and Forwarding

Key terms: switching, forwarding, bridges, learning bridges, spanning tree

Switches forward frames within a LAN using destination MAC addresses. Learning bridges self-configure: they learn (source MAC -> port) from arriving frames and flood unknown destinations. Loops would cause broadcast storms, so the spanning tree protocol (STP) blocks redundant ports, keeping one loop-free active topology that reconverges after failures.

- Cut-through vs store-and-forward switching: latency vs error checking.
- Extended LANs join segments with bridges/switches - one broadcast domain.
- VLANs partition one physical LAN into logical broadcast domains.

Lecture 17: Cell Switching: ATM

Key terms: ATM, cells, virtual circuits, VPI/VCI

ATM (Asynchronous Transfer Mode) switches tiny fixed 53-byte cells (5 header + 48 payload) - small cells mean predictable queueing delay, suiting voice/video. It is connection-oriented: virtual circuits are set up with VPI/VCI labels swapped at each switch. ATM offered QoS classes (CBR, VBR, ABR, UBR) - ahead of its time, but IP/MPLS simplicity won the market.

- 53-byte cells: compromise between voice delay and data efficiency.
- AAL adaptation layers fit different traffic types into cells.
- MPLS borrowed ATM's label-swapping idea for IP networks.

Lecture 18: Internetworking: IP, Fragmentation, ARP, ICMP, DHCP

Key terms: IP, fragmentation, ARP, ICMP, DHCP, NAT

IP provides best-effort global delivery: each packet carries source/destination addresses and is routed independently. Fragmentation splits packets exceeding a link's MTU, with reassembly at the destination (offset fields). ARP maps IP addresses to MAC addresses on a LAN ('who has this IP?'); ICMP carries error/report messages (ping, traceroute); DHCP dynamically assigns IP addresses, gateways and DNS; NAT lets private addresses share public ones - at the cost of end-to-end transparency.

- IPv4 header: version, TTL, protocol, addresses, fragmentation fields, checksum.
- Path MTU discovery avoids fragmentation using the Don't-Fragment bit.
- ICMP messages: echo, destination unreachable, time exceeded.

EXAM POINTER: Explain ARP and DHCP packet exchanges step by step - both are classic short questions.

Lecture 19: Virtual Private Networks

Key terms: VPN, tunnelling, IPsec, confidentiality, authentication

A VPN creates a private network over public infrastructure by tunnelling: packets are encapsulated and usually encrypted between VPN gateways (site-to-site) or between a remote client and the corporate gateway (remote access). IPsec is the standard suite: AH (authentication/integrity) and ESP (encryption plus authentication), with IKE negotiating keys. VPNs give confidentiality, integrity and authentication across the untrusted Internet.

- SSL/TLS VPNs work through browsers without special client software.
- Tunnelling protocols: PPTP (weak), L2TP/IPsec, OpenVPN, WireGuard.
- Trade-off: encryption overhead and centralised chokepoints vs security.

Lecture 20: Routing Algorithms: Distance Vector and Link State

Key terms: distance vector, link state, Bellman-Ford, Dijkstra

Routers compute forwarding tables with distributed algorithms. Distance vector (Bellman-Ford): each router advertises its distance to each destination to neighbours and adopts the best neighbour's offer + link cost - simple but suffers count-to-infinity; split horizon and poison reverse mitigate loops. Link state: each router floods its local link costs so all routers learn the full topology, then each runs Dijkstra's shortest-path algorithm locally - faster convergence at the cost of more state and flooding.

- Bellman-Ford equation: $D_x(y) = \min_v \{ c(x,v) + D_v(y) \}$ - memorise it.
- Dijkstra: iteratively add the closest unsettled node - needs non-negative costs.
- RIP is distance-vector; OSPF is link-state.

EXAM POINTER: Run Dijkstra on a 5-node graph and show distance-vector convergence over two exchanges - guaranteed exam material.

Lecture 21: OSPF Routing Protocol

Key terms: OSPF, areas, LSAs, designated router

OSPF (Open Shortest Path First) is the standard intra-domain link-state protocol: routers flood Link State Advertisements, build an identical topology map, and run Dijkstra. For scale, an autonomous system is divided into areas - a backbone area 0 interconnects others, limiting flooding scope. Designated routers on broadcast LANs reduce adjacency overhead; OSPF supports authentication, load balancing over equal-cost paths, and fast reconvergence.

- Hierarchy: areas summarise routes to the backbone - less flooding, smaller tables.
- Hello protocol discovers neighbours and elects designated/backup routers.
- Metrics are configurable costs; external routes are redistributed with tags.

Lecture 22: Inter-Domain Routing: BGP

Key terms: BGP, autonomous systems, path vector, policy routing

BGP (Border Gateway Protocol) routes between Autonomous Systems (ASes) - the Internet's inter-domain protocol. It is a path-vector protocol: routers advertise full AS paths, which both prevents loops (reject paths containing own AS) and enables policy routing - ISPs choose routes by business policy (prefer customers over peers over providers), not just shortest path. eBGP runs between ASes, iBGP inside.

- BGP attributes: AS-path, next-hop, local-pref, MED - policy knobs.
- Convergence can be slow (path exploration); hijacks/misconfigurations cause famous outages.
- Route reflectors scale iBGP without full mesh.

Lecture 23: IPv6

Key terms: IPv6, 128-bit addresses, header, transition, dual stack

IPv6 answers IPv4 address exhaustion with 128-bit addresses (340 undecillion), written as eight hex groups with :: compression. Its streamlined 40-byte header drops the checksum and fragmentation (hosts use Path MTU discovery), adds flow labels for QoS, and mandates IPsec support. Address types: unicast, multicast, anycast (no broadcast - uses multicast). Transition mechanisms: dual stack, tunnelling (6to4, Teredo), and translation (NAT64).

- Header fields: version, traffic class, flow label, payload length, next header, hop limit, addresses.

- Stateless autoconfiguration (SLAAC) lets hosts self-configure from router advertisements.
- Design controversies: fixed 128 bits, no NAT need, slow adoption due to IPv4 inertia.

EXAM POINTER: Expand/compress IPv6 addresses and list three header improvements over IPv4.

Lecture 24: Multicast

Key terms: multicast, IGMP, multicast routing, RPF

Multicast delivers one packet to many receivers efficiently - one copy traverses each link, replicated at branch points. Hosts join groups with IGMP (Internet Group Management Protocol); the service model is open (anyone can send to a group). Multicast routing builds distribution trees: source-based trees via Reverse Path Forwarding (forward only packets arriving on the shortest path back to source) or shared core-based trees (CBT, PIM-SM rendezvous points).

- Class D addresses (224.0.0.0/4) identify groups.
- RPF check kills loops and duplicates elegantly.
- Applications: IPTV, conferencing, software distribution; inter-domain multicast deployment remains limited.

Lecture 25: Multiprotocol Label Switching (MPLS)

Key terms: MPLS, labels, LSP, label swapping, traffic engineering

MPLS forwards packets by short fixed-length labels instead of longest-prefix IP lookups: ingress routers push a label stack, core Label Switch Routers swap labels per their tables, egress pops. Label Switched Paths (LSPs) are set up by signalling (LDP, RSVP-TE). Benefits: fast forwarding, traffic engineering (steer flows off congested shortest paths), and VPNs (labels isolate customer routes) - the technology behind carrier backbones.

- Label stack enables hierarchy: transport label outside, VPN label inside.
- MPLS-TE explicitly routes LSPs around congestion - unlike destination-only IP routing.
- Segment routing modernises the idea with source-routed label stacks.

Lecture 26: End-to-End Protocols and RPC

Key terms: end-to-end, UDP vs TCP service model, RPC

End-to-end protocols run on hosts: UDP offers minimal multiplexing; TCP offers the reliable byte stream with flow and congestion control (covered next). Remote Procedure Call (RPC) makes distributed computing look like local calls: client stubs marshal parameters into

messages, the server stub unmarshals and executes - hiding networking behind procedure-call semantics (used by NFS, gRPC).

- RPC issues: parameter marshalling (XDR, Protocol Buffers), binding, failure semantics (at-least/at-most once).
- Presentation formatting (NDR, XDR) handles byte order and data representation between heterogeneous machines.
- Modern RPC: gRPC over HTTP/2 with Protocol Buffers.

Lecture 27: TCP: Reliable Byte Stream

Key terms: TCP reliability, sequence numbers, cumulative ACK, fast retransmit

TCP's reliability machinery: every byte is sequence-numbered; receivers send cumulative ACKs (ACK n = all bytes $< n$ received); senders retransmit on timeout (RTO estimated from RTT variance - Jacobson's algorithm) or on triple duplicate ACKs (fast retransmit without waiting). Flow control via the advertised window prevents overwhelming the receiver; checksums catch corruption.

- Silly window syndrome and Nagle's algorithm handle tiny segments.
- Selective ACK (SACK) reports exactly which blocks arrived - better on lossy links.
- Connection management: simultaneous open/close, TIME_WAIT for late duplicates.

Lecture 28: Congestion Control

Key terms: congestion, AIMD, slow start, congestion avoidance, RED

Congestion control keeps the network stable: TCP's AIMD (additive increase, multiplicative decrease) probes for capacity - slow start exponentially ramps the congestion window, then congestion avoidance grows it linearly; on loss, the window halves (or resets). Queueing disciplines in routers shape fairness: FIFO with tail drop, fair queueing (per-flow queues, round robin), weighted fair queueing (shares by weight), and RED (random early detection) drops probabilistically before queues fill to signal senders early.

- Sawtooth behaviour: window oscillates around capacity - the classic TCP throughput graph.
- TCP-friendly: new algorithms (CUBIC, BBR) compete fairly with classic Reno.
- Bufferbloat: oversized buffers inflate latency - the modern congestion battleground.

EXAM POINTER: Sketch TCP's congestion-window sawtooth and explain slow start vs congestion avoidance - a top exam diagram.

Lecture 29: Quality of Service and RSVP

Key terms: QoS, IntServ, DiffServ, RSVP, traffic classes

Best-effort IP treats all packets equally - inadequate for voice/video needing delay and loss bounds. Integrated Services (IntServ) reserves per-flow resources with RSVP signalling (admission control, guaranteed/controlled-load service) - powerful but unscalable in core routers. Differentiated Services (DiffServ) marks packets into classes (DSCP bits) with per-hop behaviours (expedited forwarding for voice, assured forwarding tiers) - scalable, coarser guarantees.

- RSVP: receiver-oriented reservation, soft state refreshed periodically.
- Policing/shaping at edges: token buckets enforce/mould traffic profiles.
- MPLS-TE plus DiffServ gives carriers engineerable, class-aware backbones.

Lecture 30: Network Security: Concepts and Cryptography

Key terms: confidentiality, integrity, authentication, DES, RSA, message digest

Security goals: confidentiality (no eavesdropping), integrity (no tampering), authentication (right parties), availability. Symmetric ciphers (DES - 56-bit, now broken; AES - the standard) use one shared key and are fast; asymmetric (RSA - security from factoring difficulty) use public/private key pairs for key exchange and signatures. Cipher Block Chaining (CBC) chains blocks so patterns do not leak; message digests (MD5/SHA) fingerprint data for integrity.

- Key distribution is symmetric crypto's hard problem - solved by public-key methods (Diffie-Hellman).
- Digital signatures: sign the digest with the private key; verify with the public key.
- Certificates bind public keys to identities - the basis of HTTPS trust.

EXAM POINTER: Compare symmetric and asymmetric cryptography on speed, key distribution and typical uses.

Lecture 31: Security Protocols: IPsec, SSL/TLS, Firewalls

Key terms: IPsec, SSL/TLS, firewall, IDS

IPsec secures IP itself: AH authenticates, ESP encrypts+authenticates, in transport mode (host-to-host) or tunnel mode (gateway-to-gateway VPNs), with IKE managing keys. SSL/TLS secures above TCP - the padlock in browsers: handshake negotiates ciphers and authenticates the server via certificates, then bulk-encrypts. Firewalls filter by policy (packet filters, stateful inspection, application gateways); IDS/IPS detect and block intrusions.

- TLS 1.3: faster handshake, fewer weak options - know why it matters.
- DMZ architectures isolate public servers from the internal network.

- Defence in depth: no single mechanism suffices - layer them.

Lecture 32: Mobile IP and Wireless Networking Trends

Key terms: Mobile IP, home agent, foreign agent, handoff

Mobile IP lets hosts keep their IP address while roaming: the home agent tunnels packets to the mobile node's care-of address at the foreign agent; return traffic routes directly (route optimisation). Challenges: triangle routing inefficiency, handoff latency, and security of binding updates. Broader trends: 5G network slicing, Wi-Fi 6/7, IoT scale (billions of constrained devices), and edge computing pushing content closer to users.

- Handoff types: hard (break-before-make) vs soft (make-before-break).
- Cellular evolution: 1G voice -> 5G slicing, massive MIMO, mmWave.
- IoT protocols (6LoWPAN, LoRaWAN, NB-IoT) trade rate for range and battery.

Lecture 33: Revision: Midterm Practice Set

1. Compare circuit switching and packet switching. (5 marks)
2. Compute the bandwidth-delay product of a 1 Gbps link with 50 ms RTT and explain its significance. (5 marks)
3. Apply bit stuffing to the data 011111101111110. (5 marks)
4. Explain CSMA/CD with the exponential backoff algorithm. (10 marks)
5. Run Dijkstra's algorithm on a given 6-node graph from source A. (10 marks)

Lecture 34: Revision: Final-Term Practice Set

1. Explain BGP's path-vector approach and how ISP routing policy overrides shortest-path. (10 marks)
2. Describe TCP congestion control: slow start, AIMD and fast retransmit, with the window diagram. (10 marks)
3. Compare IntServ/RSVP with DiffServ for QoS. (10 marks)
4. Write short notes on: MPLS label swapping, IGMP, NAT. (15 marks)
5. Compare DES and RSA; explain how digital signatures provide authentication and integrity. (10 marks)

End of Handout - Good luck with your exams!