THE UNIVERSITY OF AUCKLAND

SECOND SEMESTER, 2012
Campus: City

Computer Science

Modern Data Communications

(Time Allowed: TWO HOURS)

Note:

- The use of calculators is NOT permitted.
- Compare the exam version number on the Teleform sheet supplied with the version number above. If they do not match, ask the supervisor for a new sheet.
- Enter your name and student ID on the Teleform sheet. Your name should be entered left aligned. If your name is longer than the number of boxes provided, truncate it.
- Answer all **Multiple-choice** questions on the Teleform answer sheet provided. Attempt all questions. There are no negative marks.
- Use a dark pencil to mark your answers in the multiple choice answer boxes on the Teleform sheet. If you spoil your sheet, ask the supervisor for a replacement.

Question 1
[1 mark] Which of the following statements is TRUE?

(a) By compression, the size of a file of 9n characters has decreased by 5n characters thus resulting in a 33% reduction.
(b) By compression, the size of a file of 9n characters has decreased by 5n characters thus resulting in a 45% reduction.
(c) By compression, the size of a file of 9n characters has decreased by 5n characters thus resulting in a 44.5% reduction.
(d) By compression, the size of a file of 9n characters has decreased by 5n characters thus resulting in a 50% reduction.
(e) By compression, the size of a file of 9n characters has decreased by 5n characters thus resulting in a 55.5% reduction.

Question 2
[1 mark] Which of the following codewords is a correct Huffman set of codewords for the letters A, B, C, D, E, F having frequencies 45%, 20%, 15%, 10%, 7%, 3%?

(a) A= 0, B=100, C=10, D=110, E=1110, F=1111.
(b) A= 0, B=100, C=101, D=110, E=1110, F=1111.
(c) A= 0, B=100, C=101, D=101, E=1110, F=1111.
(d) A= 0, B=0100, C=101, D=110, E=1110, F=1111.
(e) A= 0, B=100, C=101, D=110, E=111, F=1111.

Question 3
[1 mark] Which of the following statements is TRUE?

(a) There is no frequency-dependent code.
(b) In a frequency-dependent code, less frequent characters have no codewords.
(c) In a frequency-dependent code, more frequent characters have longer codewords.
(d) In a frequency-dependent code, less frequent characters have shorter codewords.
(e) In a frequency-dependent code, more frequent characters have shorter codewords.

Question 4
[1 mark] The numbers 3, 2, 3:

(a) satisfy Kraft's inequality and the prefix code produced by Kraft's theorem is 00, 010, 011.
(b) satisfy Kraft's inequality and the prefix code produced by Kraft's theorem is 00, 0100, 1001.
(c) the numbers do not satisfy Kraft's inequality.
(d) satisfy Kraft's inequality and the prefix code produced by Kraft's theorem is 10, 010, 100.
(e) satisfy Kraft's inequality and the prefix code produced by Kraft's theorem is 00, 010, 1111.

Question 5
[1 mark] All ASCII codewords have the same length, so ASCII

(a) is an infinite code.
(b) is a frequency-dependent code.
(c) is not uniquely decodable.
(d) is a prefix code.
(e) is not a prefix code.

Question 6
[1 mark] Which of the following statements is TRUE?

(a) A parity bit is a redundant bit added to every data unit so that the total number of 1s in the unit (excluding the parity bit) becomes even (or odd).
(b) A parity bit is a redundant bit added to every data unit so that the total number of 0s in the unit (excluding the parity bit) becomes even (or odd).
(c) A parity bit is a redundant bit added to every data unit so that the total number of 1s in the unit (including the parity bit) becomes even (or odd).
(d) A parity bit is a redundant bit added to every data unit so that the total number of 0s and 1s in the unit (including the parity bit) becomes even (or odd).
(e) A parity bit is a redundant bit added to every data unit so that the total number of 0s in the unit (including the parity bit) becomes even (or odd).

Question 7
[1 mark] Which of the following statements is TRUE?

(a) Error correction is as simple as error detection.
(b) Error detection is impossible without error correction.
(c) Error correction is simpler than error detection.
(d) Error detection is as simple as error correction.
(e) Error detection is much simpler than error correction.

Question 8
[1 mark] Which of the following statements is TRUE?

(a) A Hamming code for 7-bit ASCII code has 2 redundancy bits.
(b) A Hamming code for 7-bit ASCII code has 4 redundancy bits.
(c) A Hamming code for 7-bit ASCII code has 3 redundancy bits.
(d) A Hamming code for 7-bit ASCII code has 1 redundancy bits.
(e) A Hamming code for 7-bit ASCII code has 5 redundancy bits.

Question 9
[1 mark] The "Open Systems Interconnection" communications model has seven network layers. Which of the following statements is FALSE?

(a) Network layers allow us to separate a network design into independent functions.
(b) Network layers allow us to implement common functions only once.
(c) Network layers allow us to provide and use alternative implementations of each layer.
(d) Network layers are a theoretical abstraction, they only help us to think about networks.
(e) Network layers make it easier to add or remove functions within each layer.

Question 10
[1 mark] Assume that a fibre from Auckland to Palmerston North is 400 km long, and is used to provide a 10 Mb/s link. Approximately how long does it take for a 1500-byte frame to be sent out and acknowledged? *(Hints: light travels at 2×10^8 m/s in optical fibre; you may ignore the time needed to transmit the short ACK frame).*

(a) 1.2 ms
(b) 4 ms
(c) 3.2 ms
(d) 5.2 ms
(e) 2 ms

Question 11
[1 mark] Time-Division Multiplexing (TDM) and Statistical Multiplexing (SM) are two methods to carry several bit streams through a channel. Which of the following statements is FALSE?

(a) TDM doesn't require data frames to carry a stream identifier.
(b) SM data frames may carry a stream identifier.
(c) SM is harder to implement, which could make Statistical Multiplexers more expensive.
(d) TDM uses fixed time slots, so each bit stream gets a well-defined share of the channel.
(e) SM makes better use of the channel if traffic in the bit streams is bursty.

Question 12
[1 mark] Which of the following statements about the Selective Repeat protocol (SR) is FALSE?

(a) An SR sender's window base cannot have a lower sequence number than the receiver's window base.
(b) SR sends an ACK frame for each correctly-received frame.
(c) An SR sender knows that all ACKed frames have been safely received.
(d) SR resends frames that receive a NAK response.
(e) SR resends timed-out ACK or NAK frames.

CONTINUED

Question 13
[1 mark] How does an Ethernet sender behave when it detects a collision on the medium?

(a) It sends a short 'jamming' signal, so as to minimise the time the medium is occupied, then resends.
(b) It stops sending, waits for the medium to be quiet, then resends.
(c) It sends a long 'jamming' signal, so as to make all stations aware of the collision, then resends.
(d) It discards the frame it was trying to send, and reports the collision to its network layer.
(e) It stops sending, waits for a random time, then resends.

Question 14
[1 mark] Consider Ethernet hubs and switches. Which of the following statements is FALSE?

(a) An Ethernet hub always repeats a signal to all its ports.
(b) An Ethernet switch may receive a packet then re-send it.
(c) An Ethernet switch allows full-duplex Ethernet working.
(d) An Ethernet hub can be used to extend an Ethernet segment.
(e) An Ethernet switch can have more ports than a hub.

Question 15
[1 mark] Category-5 twisted-pair Ethernet wiring has four separate copper wire pairs in an insulating outer covering. How does 1000BaseT (Gigabit Ethernet) use the pairs to achieve its 1 Gb/s data rate?

(a) Each pair carries 250 Mb/s of data, using Manchester encoding; error correction is not needed.
(b) Each pair carries 250 Mb/s of data, using 5-level signaling; error-correction is not needed.
(c) Each pair carries 250 Mb/s of data, using 5-level signaling and error-correcting codes.
(d) Each pair carries 250 Mb/s of data, using Manchester encoding and error-correcting codes.
(e) Two pairs carry 500 Mb/s, using 8-level signaling; the other two are unused.

Question 16
[1 mark] Why is spread spectrum radio transmission used by 802.11?

(a) It can penetrate thin walls, making it suitable for use indoors.
(b) To minimise problems with interference from other wireless transmissions.
(c) Frequency modulation is simpler to implement than amplitude modulation.
(d) Spread spectrum allows 802.11 to have more channels in its radio frequency band.
(e) 802.11 only needed an effective range of a few tens of metres.

Question 17
[1 mark] Which of the following is NOT a reason why 802.11 wireless uses collision avoidance rather than collision detection?

(a) Collision detection performs poorly when wireless channel usage is high.
(b) Resending packets shortens the time between battery charges for portable devices.
(c) A receiving host may not be able to hear signals from all senders in a wireless network.
(d) Wireless channel conditions near a receiver can differ from those near a sender, making collisions hard to detect reliably.
(e) For small devices, their own strong signal drowns any weak incoming signal.

Question 18
[1 mark] How do Ethernet switches learn which port to send packets out through?

(a) Switches build up tables of IP addresses for each port from packets arriving on their ports.
(b) Switches read their forwarding tables from a configuration server on the network.
(c) Switches don't need address tables, they simply forward packets to all their ports.
(d) Switches build up tables of source MAC addresses for each port from packets arriving on their ports.
(e) Switches must be configured 'by hand,' i.e. a new host must be added to a table each time we connect one.

Question 19
[1 mark] Consider an Ethernet network made up of many segments, interconnected by one or more bridges on each segment, configured to use the Spanning Tree protocol. Which of the following statements does NOT describe an effect caused by Spanning Tree?

(a) Some bridges will normally be idle, acting as backup when another bridge fails.
(b) A segment with only one bridge will be left isolated if that bridge fails.
(c) Broadcast packets will reach all the connected hosts in the network.
(d) The network will recover fairly quickly after an active bridge fails.
(e) The bridges will cooperate to form a network in which packets are not forwarded in loops.

Question 20
[1 mark] Which of the following statements describes a DISADVANTAGE of VLANs?

(a) Broadcast packets are only sent to hosts in their source host's VLAN.
(b) VLANs let sets of hosts inter-operate together even if they are not connected to the same network switch.
(c) VLANs may allow network administrators to make more effective use of their inter-building and inter-site cabling.
(d) 802.1q VLAN tags make Ethernet packets four bytes longer.
(e) VLAN tagging can enable network administrators to prioritse network traffic within trunks (links that carry traffic for several VLANs).

Question 21
[1 mark] RIP is an example of a Distance Vector (DV) routing protocol. Why does RIP use a distance value of 16 to mean 'infinity?'

(a) Real networks don't have infinite-length paths.
(b) The DV algorithm requires comparisons between path lengths.
(c) Using a low value for 'infinity' limits the time RIP takes to recompute routes.
(d) Real networks may have paths with more than eight hops.
(e) We cannot use 'infinity' as a value in computer arithmetic.

Question 22
[1 mark] Which of the following is a DISADVANTAGE of Link State routing?

(a) Link State routing relies on distributing network maps rather than routing tables.
(b) Calculating a forwarding table using the Link State information is not a distributed algorithm.
(c) Since Link State routers know the state of all the network's links, their routing tables converge quickly.
(d) Routers need fast processors and large memory to handle the Link State calculation and storage.
(e) A Link State router calculates a new routing table whenever a link state changes.

Question 23
[1 mark] Which of the following best describes BGP4?

(a) BGP4 is a Link State routing algorithm.
(b) BGP4 is a modified Distance Vector routing algorithm.
(c) BGP 4 is a Path Vector routing protocol.
(d) BGP4 is a modified Distance Vector routing protocol.
(e) BGP4 is a Link State routing protocol.

Question 24
[1 mark] BGP4 can aggregate adjacent address prefixes so as to keep the global routing table as small as possible. Consider this set of IPv4 prefixes: 128.66.0.0/22, 128.66.8.0/22, 128.66.16.0/22, 128.66.24.0/22; what could these aggregate to?

(a) 128.66.0.0/19
(b) 128.66.0.0/17
(c) 128.66.0.0/18
(d) 128.66.0.0/20
(e) 128.66.0.0/16

Question 25
[1 mark] TCP/IP is usually implemented as a stack of code modules, one for each layer. What advantage could this give to developers wanting to implement TCP/IP over a new kind of technology platform?

(a) All network devices need the same software.
(b) They only have to produce a new module for one of the layers.
(c) Separate code modules are always more efficient.
(d) Separate code modules need less memory.
(e) Official standards require separate code modules.

Question 26
[1 mark] What does the IP protocol do within the Internet protocol stack?

(a) IP makes sure that packets are delivered correctly.
(b) IP allows packets to be forwarded from one Ethernet to another.
(c) IP allows packets to be routed from their source to their destination.
(d) IP allows data to be sent as a stream of bytes.
(e) IP provides point-to-point connections between hosts.

Question 27
[1 mark] What is the MAIN difference between IPv4 and IPv6?

(a) IPv6 provides more features in its design.
(b) IPv6 is designed for faster forwarding of packets.
(c) IPv6 uses 128-bit addresses instead of IPv4's 32-bit addresses.
(d) IPv6 has simplified headers compared to IPv4.
(e) IPv6 supports connections for smartphones.

Question 28
[1 mark] Why do IPv6 routers not fragment datagrams?

(a) Fragmentation could reduce routers' packet forwarding rate.
(b) Network links these days can carry standard-size Ethernet packets.
(c) IPv6 hosts fragment datagrams that are too big for a network path before sending them.
(d) Modern transport protocols don't rely on datagrams longer than an Ethernet packet.
(e) The IPv6 header has no fields to support fragmentation.

Question 29
[1 mark] How does a host decide whether an IP address is within the same network as itself?

(a) It sends a packet to the host to see whether that host answers.
(b) By comparing that host's network number with its own network number and netmask.
(c) A host can ask its default router.
(d) By requesting that information using DHCP.
(e) Each host is configured with a list of reachable IP addresses.

Question 30
[1 mark] Why should a host refresh its MAC cache obtained using ARP?

(a) To free up the memory consumed by the cache.
(b) So that it does not send packets to the wrong host because of stale information.
(c) The cache represents a single-point of failure and should not be trusted.
(d) To speed up the sending of packets.
(e) Because a MAC address might be dynamically assigned to a different host.

Question 31
[1 mark] What path metric is most commonly used in IP routers?

(a) The path with the lowest latency.
(b) Time for a datagram to travel along the path.
(c) The routers always prefer to use routes with secure links.
(d) Shortest path in number of hops.
(e) The path that has the lowest monetary cost.

Question 32
[1 mark] Why are Traceroute and Tracert different? Choose the most appropriate answer.

(a) Because Treacert is implemented in Windows and it is copyrighted code.
(b) Because Traceroute uses UDP for outgoing packets while Tracert uses ICMP packets.
(c) Because Traceroute records the route in an IP packet using the option in the IP header.
(d) Traceroute is more reliable than Tracert because it uses UDP packets.
(e) Traceroute and Tracert are the same tool, but they are implemented for different OSes.

Question 33
[1 mark] Let [IDENT, D, M, FRAG OFFSET, DATA BYTES] be the fragmentation fields of an IPv4 packet plus the bytes of payload data. The packet with the following information [IDENT=x, D=1, M=0, FRAG OFFSET = 0, DATA BYTES 1800] arrives at a router that has to forward it to a link with MTU = 600. Which of the following statements is TRUE?

(a) The router sends back to the sender an ICMP Type 3 message.
(b) The router requests the sender of the IP packet to split the packet in 3 fragments and send them again.
(c) The router sends the packet to the previous router asking it to select a different route with larger MTU.
(d) The router splits the packet in 3 different fragments [IDENT=x, D=0, M=1, FRAG OFFSET = 0, DATA BYTES 600], [IDENT=x, D=0, M=1, FRAG OFFSET = 75, DATA BYTES 600], and [IDENT=x, D=0, M=0, FRAG OFFSET = 150, DATA BYTES 600].
(e) The router forwards the packet as it is because D is set to 1.

Question 34
[1 mark] When you point your browser to www.google.com which of the following is a correct encapsulation?

(a) An Ethernet frame encapsulates the IP packet that carries the TCP packet that encapsulates the HTTP request.
(b) An ARP packet encapsulates the TCP packet that encapsulates the HTTP request.
(c) The HTTP request encapsulates a TCP packet which in turn encapsulates an Ethernet frame.
(d) The HTTP request encapsulates an IP packet which in turn encapsulates an Ethernet frame.
(e) An IP packet encapsulates the HTTP request that encapsulates the Ethernet frame.

Question 35
[1 mark] Which of the following statements is NOT true?

(a) UDP makes no guarantees about data delivery or the correctness of delivered data.
(b) For UDP traffic, the packet order at transmission may not be the same as the packet order at reception.
(c) Both TCP and UDP use ports to distinguish applications that use them.
(d) TCP guarantees that the packet ordering at reception is the same as the packet order at transmission.
(e) TCP guarantees data delivery as well as correctness of data.

Question 36
[1 mark] UDP is an unreliable transport protocol; which of the following is NOT a reason why some applications use UDP rather than TCP?

(a) UDP is a connectionless protocol, unlike TCP.
(b) Applications might need to specify different congestion strategies than TCP.
(c) Some applications don't need to recover lost data.
(d) Some applications might prefer to cope with lost data rather than wait for the resending of lost packets.
(e) Because creating a UDP packet is simpler than creating a TCP packet.

Question 37
[1 mark] Which of the following is NOT a reason why IP datagrams (carrying TCP segments) might be lost?

(a) Dropped at receiving host because a buffer is full.
(b) Corrupted on a link between two routers.
(c) Dropped at overloaded router.
(d) Caught in a routing loop until its hop count reached zero.
(e) Delayed too long in a low-speed link between routers.

Question 38
[1 mark] If a TCP header has the ACK bit set, what does that tell the host receiving it (choose the most appropriate answer)?

(a) The sender is now sure that the receiver node has processed all the bytes that it sent before.
(b) The receiving host acknowledges the sender to keep sending the segments.
(c) The sender is informed by the receiving host that the route is established.
(d) It is an acknowledgement to inform the sender to send extra bytes up to the sequence number specified in the packet.
(e) The receiving host's next expected sequence number.

Question 39
[1 mark] Consider host A in a TCP connection with host B. A has just received a TCP packet with sequence number 38575 and a content of 200 bytes. What should host A do, and what should it expect next?

(a) Send a packet with acknowledgement number 38774 and expect a packet with sequence number 38775.
(b) Send a packet with acknowledgement number 38775 and expect a packet with sequence number 38775.
(c) Send a packet with acknowledgement number 38576 and expect a packet with sequence number 38576.
(d) Send a packet with acknowledgement number 38575 and expect a packet with sequence number 38576.
(e) Send a packet with acknowledgement number 38775 and expect a packet with sequence number 38776.

Question 40
[1 mark] How does TCP ensure a fair balance between unacknowledged data and the risk of losing a data segment?

(a) TCP uses a sliding window to manage the number of bytes in transit.
(b) TCP uses multiple connections for delivering unacknowledged segments to a destination.
(c) TCP makes sure that acknowledged segments have higher priority than new data segments.
(d) TCP sends further segments only when it receives acknowledges for the segments previously sent.
(e) TCP uses a sliding window to manage the packets in transit.

Rough Working – This page will not be marked