S5 ICT Practice Paper 4 | Subnetting, TCP/IP, databases, complexity, compression and information security

Senior Secondary ICT

S5 ICT practice paper 4: IPv4 subnetting, TCP/IP layers, normalisation and SQL, algorithm complexity, file sizes and information security, with full answers.

1. IPv4 subnetting and usable hosts (18 marks)

Usable hosts = 2^n - 2, where n is the number of host bits (network and broadcast addresses are excluded).

#	Question	Answer and marking notes
1	(a) For the network 192.168.10.0/26, find the number of addresses per subnet, the number of usable hosts, and how many /26 subnets a /24 block yields.	Host bits = 32 - 26 = 6, so each subnet has 2^6 = 64 addresses and 64 - 2 = 62 usable hosts. A /24 block holds 2^(26-24) = 4 /26 subnets.
2	(b) For the network 10.0.0.0/22, find the subnet mask, the number of addresses and the number of usable hosts.	The mask is 255.255.252.0 (22 ones followed by 10 zeros). Host bits = 10, so there are 2^10 = 1024 addresses and 1024 - 2 = 1022 usable hosts.
3	(c) For the host address 172.16.5.130/27, find the network address, the broadcast address and the usable host range.	The block size is 2^(32-27) = 32. Since 130 / 32 = 4 remainder 2, the network address is 172.16.5.128 and the broadcast address is 172.16.5.159; usable hosts run from 172.16.5.129 to 172.16.5.158.
4	(d) A /24 block 192.168.1.0 must be split into at least 5 subnets with at least 25 hosts each. Find the most suitable prefix length and justify it.	Twenty-five hosts need 5 host bits because 2^5 - 2 = 30 is enough, while /28 gives only 2^4 - 2 = 14. The prefix is therefore 32 - 5 = /27 (255.255.255.224), which yields 2^3 = 8 subnets and meets the requirement.

2. TCP/IP layers and protocols (16 marks)

Name the layer and the protocol in each answer.

#	Question	Answer and marking notes
5	(a) List the four layers of the TCP/IP model and give one protocol for each.	Application layer (HTTP, DNS, SMTP), transport layer (TCP, UDP), internet layer (IP, ICMP) and network access layer (Ethernet, IEEE 802.11).
6	(b) Compare TCP and UDP on four points: connection, reliability, overhead and typical use.	TCP is connection-oriented and reliable (acknowledgements, retransmission, in-order delivery) but has a larger header and more overhead, so it suits the web and file transfer. UDP is connectionless, does not guarantee delivery, has a small header and low latency, and suits voice, live streaming and DNS lookups.
7	(c) A link has an MTU of 1500 bytes, with a 20-byte IP header and a 20-byte TCP header. Find the MSS and the payload as a percentage of the frame.	MSS = 1500 - 20 - 20 = 1460 bytes, so the payload is 1460 / 1500 x 100% = 97.33% of the frame.

#	Question	Answer and marking notes
8	(d) Outline the three steps of the TCP three-way handshake and state its purpose.	The client sends SYN with its initial sequence number, the server replies SYN-ACK acknowledging it and sending its own number, and the client returns ACK. This confirms that both sides can send and receive, synchronises the initial sequence numbers and establishes a reliable connection before data flows.

3. Database normalisation and SQL (18 marks)

Use standard SQL syntax and state the dependency when normalising.

#	Question	Answer and marking notes
9	(a) State the requirements of first, second and third normal form.	1NF: every cell holds a single atomic value and there are no repeating groups. 2NF: 1NF plus every non-key column depends on the whole primary key, removing partial dependencies. 3NF: 2NF plus no non-key column depends on another non-key column, removing transitive dependencies.
10	(b) The table Enrol(student_id, student_name, class_code, class_teacher) has two dependency problems. Identify them and decompose the table into 3NF.	student_name depends only on student_id (a partial dependency if the key is student_id plus class_code), and class_teacher depends only on class_code (a transitive dependency). Decompose into Student(student_id, student_name, class_code), Class(class_code, class_teacher) and Enrol(student_id, class_code).
11	(c) Write the SQL that lists each class whose average score is above 70, with its average score.	SELECT class_code, AVG(score) AS avg_score FROM Enrol GROUP BY class_code HAVING AVG(score) > 70; With the sample data only A1 qualifies, averaging 80.33.
12	(d) Write the SQL that lists every student name with the teacher of that student class (a two-table join).	SELECT s.name, c.teacher FROM Student s JOIN Class c ON s.class_code = c.class_code; With the sample data this returns 5 rows.

4. Algorithm complexity and sorting (16 marks)

Show the expression used for each answer.

#	Question	Answer and marking notes
13	(a) For a sorted array of 1,048,576 elements, how many comparisons do linear search and binary search need in the worst case?	Linear search needs n = 1048576 comparisons in the worst case; binary search needs ceil(log2 n) = ceil(log2 1048576) = 20.
14	(b) For 1024 elements, bubble sort makes n(n-1)/2 comparisons in the worst case while merge sort makes about n log2 n. Compute both and compare.	Bubble sort: 1024 x 1023 / 2 = 523776 comparisons. Merge sort: about 1024 x 10 = 10240. That is roughly 51 times fewer, and the gap widens as n grows.

#	Question	Answer and marking notes
15	(c) Take the array [42, 17, 68, 9, 55, 23, 71, 34] and run one Lomuto partition with the last element as pivot. Give the resulting array and the pivot index.	The pivot is 34. Sweep left to right, moving every element not greater than 34 to the front, then swap the pivot into the boundary: the array becomes [17, 9, 23, 34, 55, 68, 71, 42] with the pivot at index 3, everything to its left at most 34 and everything to its right greater than 34.
16	(d) Order O(n^2), O(n log n), O(log n) and O(1) from slowest to fastest growth and give one example of each.	O(1) (array access by index) < O(log n) (binary search) < O(n log n) (merge sort) < O(n^2) (bubble sort). As n grows the highest-order term dominates the running time.

5. Digital file sizes and compression (16 marks)

1 MiB = 1,048,576 bytes and 1 MB = 1,000,000 bytes.

#	Question	Answer and marking notes
17	(a) An uncompressed 1920 x 1080 bitmap at 24-bit colour depth: find its size in MiB and in MB.	Each pixel uses 24 bits = 3 bytes, so the total is 1920 x 1080 x 3 = 6220800 bytes = 5.93 MiB, about 6.22 MB.
18	(b) Find the size of an uncompressed 3-minute WAV file at 44.1 kHz, 16-bit, stereo.	Each second needs 44100 x 2 bytes x 2 channels = 176400 bytes; 3 minutes is 180 s, so the file is 31752000 bytes = 30.28 MiB, about 31.75 MB.
19	(c) A row of indexed-colour pixels reads WWWWBBBBWWWWWWBBB at 1 byte per pixel. Encode it with run-length encoding using 2 bytes per (colour, count) pair, and find the compressed size and the compression ratio.	The runs are (W,4) (B,4) (W,6) (B,3), so 4 pairs x 2 bytes = 8 bytes against 17 bytes originally; the ratio is 17 / 8 = 2.125. RLE works best when colours come in long runs.
20	(d) Compare lossy and lossless compression, giving one example of each and when it is used.	Lossless compression restores the file exactly (PNG, ZIP, FLAC) and suits text, source code and medical images. Lossy compression discards some data for a much higher ratio (JPEG, MP3, H.264) and suits photos, music and video, though repeated re-encoding accumulates visible or audible loss.

6. Information security and legislation (16 marks)

Name the technique and state what it protects.

#	Question	Answer and marking notes
21	(a) Compare symmetric and asymmetric encryption, giving one algorithm and one use for each.	Symmetric encryption uses one shared key for both directions (AES): fast, so it protects bulk data such as disks and HTTPS session traffic. Asymmetric encryption uses a public key to encrypt and a private key to decrypt (RSA): used for key exchange and digital signatures, but computationally slower.

#	Question	Answer and marking notes
22	(b) Explain how a digital signature provides both integrity and non-repudiation.	The sender hashes the message (SHA-256) and signs the hash with a private key. The receiver verifies the signature with the sender public key and recomputes the hash; a match proves the content was not altered (integrity), and only the holder of the private key could have produced the signature (non-repudiation).
23	(c) Name three common network attacks and one defence for each.	Phishing email (check the sender domain and enable multi-factor authentication); password brute force (long passphrases plus account lockout and MFA); man-in-the-middle (enforce HTTPS and validate certificates).
24	(d) Of the six data protection principles in Hong Kong Personal Data (Privacy) Ordinance (Cap. 486), explain the two most relevant to collecting and using customer data, with one example each.	Principle 1 (purpose and manner of collection): data must be collected for a lawful purpose directly related to the function, adequate but not excessive, and the purpose must be notified, for example a trial-class form collecting only a name and a phone number. Principle 3 (use): unless the data subject consents or the law allows, data may be used only for the purpose stated at collection, so a contact list may not be passed to a third party for marketing without consent.