S6 ICT Practice Paper: Data Representation and File Compression, Network Transmission and Subnetting, Algorithms and Program Tracing, Relational Databases and SQL, Multimedia and Streaming, Security and Backup

S6 | ICT | Time: 60 minutes | Total: 100 marks | 24 questions

---

Data Representation and File Compression

Part A: Data representation and file compression (16 marks)

1. A digital camera captures an uncompressed image at 1920 by 1080 pixels with a colour depth of 24 bits per pixel. (a) Calculate the file size in MiB (1 MiB = 1 048 576 bytes). (b) State one benefit and one cost of saving the photo as JPEG instead. [4]

Answer: (a) Total bits = 1920 x 1080 x 24 = 49,766,400; bytes = 49,766,400 / 8 = 6,220,800; MiB = 6,220,800 / 1 048 576 = 5.93 MiB. (b) Benefit: lossy compression greatly reduces file size, saving storage and speeding up web transfer. Cost: fine detail is discarded and repeated editing leaves compression artefacts.

2. A stereo recording uses a sampling rate of 44,100 Hz, a bit depth of 16 bits and lasts 180 seconds. (a) Calculate the uncompressed size in MiB to 1 decimal place. (b) If it is compressed at 128 kbps, calculate the compressed size and the compression ratio. [4]

Answer: (a) Bits = 44,100 x 16 x 2 x 180 = 254,016,000; bytes = 31,752,000; size = 30.3 MiB. (b) 128 bits per second x 180 s / 8 = 2,880,000 bytes = 2.7 MiB, a ratio of about 11.0:1.

3. A plain text file has 2,000 characters: 1,500 ASCII characters and 500 Chinese characters. (a) Calculate the number of bytes needed in UTF-8 (1 byte per ASCII character, 3 bytes per Chinese character) and in UTF-16 (2 bytes per character). (b) For an English-dominant website, state which encoding is preferable and why. [4]

Answer: (a) UTF-8 = 1,500 x 1 + 500 x 3 = 3,000 bytes; UTF-16 = 2,000 x 2 = 4,000 bytes. (b) UTF-8 is preferable: each English character takes only 1 byte, so the file is smaller and transfers faster, and it is the default encoding for web pages with the widest compatibility.

4. A 45 MB raw image is compressed to 9 MB. (a) Calculate the compression ratio and the percentage reduction in file size. (b) Give one example each of lossless and lossy compression, and explain why lossless compression cannot achieve the same reduction. [4]

Answer: (a) Ratio = 45 / 9 = 5:1; reduction = (45 - 9) / 45 x 100% = 80%. (b) Lossless examples: PNG and ZIP; lossy examples: JPEG and MP3. Lossless compression only removes redundancy and must restore the data exactly, so its reduction is limited, while lossy compression can permanently discard detail the eye or ear barely notices.

Network Transmission and Subnetting

Part B: Network transmission and subnetting (17 marks)

5. A 250 MB file is sent over a 100 Mbps link whose actual efficiency is only 80%. (a) Calculate the transfer time in seconds. (b) Explain two reasons why efficiency is below 100%. [4]

Answer: (a) Data = 250 x 8 = 2,000 Mbit; effective rate = 100 x 0.8 = 80 Mbps; time = 2,000 / 80 = 25 s. (b) Every packet carries headers and checksums that consume bandwidth, and the sender must wait for acknowledgements and retransmit lost packets, with congestion and device processing delay adding further overhead.

6. An organisation is given the network 192.168.10.0/24 and divides it into /26 subnets. (a) How many subnets can be created? (b) How many usable host addresses does each subnet have? (c) Write the subnet mask in dotted decimal and explain why two addresses are subtracted. [5]

Answer: (a) 2 bits are borrowed, giving 2^2 = 4 subnets. (b) Host bits = 32 - 26 = 6; subtracting the network and broadcast addresses gives 2^6 - 2 = 62 usable hosts. (c) The mask is 255.255.255.192; the first address identifies the subnet and the last is the broadcast address, so neither can be assigned to a host.

7. Two computers 1,200 km apart are joined by fibre, where signals travel at 2.0 x 10^8 metres per second. The path passes through 4 routers, each adding 0.5 ms of processing delay. (a) Calculate the one-way propagation time. (b) Calculate the round-trip time (RTT). [4]

Answer: (a) Distance = 1,200 x 1 000 = 1,200,000 m; time = 1,200,000 / 2.0 x 10^8 = 6 ms. (b) RTT = 6 x 2 + 4 x 0.5 = 14 ms.

8. A 9,000-byte packet is sent over a 1 Gbps link. (a) Calculate the transmission time for one packet in microseconds. (b) Calculate the total time in milliseconds for 100 consecutive packets, and state one delay this calculation ignores. [4]

Answer: (a) Bits = 9,000 x 8 = 72,000; time = 72,000 / 1 x 10^9 = 72 microseconds. (b) Total = 72 x 100 = 7200 microseconds = 7.2 ms; propagation delay and router processing and queuing delays are ignored.

Algorithms and Program Tracing

Part C: Algorithms and program tracing (16 marks)

9. Consider the code: for i in range(40): for j in range(i): count += 1. (a) Calculate the value of count when the loops finish. (b) State the time complexity and explain how the running time changes when the input doubles. [4]

Answer: (a) The inner loop runs 0 + 1 + 2 + ... + 39 times = 40 x 39 / 2 = 780. (b) The complexity is O(n squared); doubling the input makes the inner total about four times larger, so the running time is about four times longer.

10. A program searches an ordered array of 1,000,000 items using binary search. (a) Calculate the maximum number of comparisons needed. (b) Explain why the array must be sorted and state the advantage over linear search. [4]

Answer: (a) Each comparison halves the range, so at most log base 2 of 1,000,000, rounded up, = 20 comparisons. (b) Binary search decides whether the target lies in the left or right half by comparing with the middle value, which only works if the array is sorted. Its O(log n) cost is far below the O(n) of linear search: one million items need only 20 comparisons.

11. Bubble sort is used on 30 items. (a) Calculate the number of comparisons in the first pass. (b) Calculate the total comparisons in the worst case. (c) In an unordered array of 1,000 items, how many comparisons does linear search need on average? [4]

Answer: (a) The first pass compares 30 adjacent items, needing 30 - 1 = 29 comparisons. (b) Worst-case total = 30 x 29 / 2 = 435. (c) On average (n + 1) / 2 = (1,000 + 1) / 2, about 500 comparisons.

12. A recursive function solves term n by calling itself twice on term n - 1, so the total number of calls is 2 to the power n. (a) Calculate the total calls when n = 20. (b) Explain the problem when n rises to 30 and suggest one improvement. [4]

Answer: (a) 2 to the power 20 = 1,048,576 calls. (b) At n = 30 it needs 2 to the power 30 = 1,073,741,824 calls, about 1 024 times more, so time and memory use explode; add memoisation to cache computed results or rewrite it iteratively, reducing O(2 to the power n) to O(n).

Relational Databases and SQL

Part D: Relational databases and SQL (17 marks)

13. The orders table holds 5 orders with amounts 120, 80, 200, 50, 150. (a) Write one SQL statement that returns the number of orders, the total amount and the average amount. (b) Write a statement listing each customer id whose total spending exceeds 100 with that total, and explain why HAVING is used rather than WHERE. [5]

Answer: (a) SELECT COUNT(*), SUM(amount), AVG(amount) FROM orders; this returns 5 orders, a total of 600 and an average of 120, with the largest single order 200. (b) SELECT customer_id, SUM(amount) FROM orders GROUP BY customer_id HAVING SUM(amount) > 100; WHERE filters individual rows before grouping, when SUM does not yet exist, whereas HAVING filters the aggregated groups after GROUP BY.

14. A table stores all the phone numbers of one customer in a single cell separated by commas. (a) State which normal form this violates and how to fix it. (b) Explain the different roles of a primary key and a foreign key in maintaining data integrity. [4]

Answer: (a) It violates first normal form because the cell is not atomic; create a separate phone table keyed by the customer id with one row per number. (b) A primary key uniquely identifies each row and cannot repeat or be null, while a foreign key references the primary key of another table so that a non-existent customer id cannot be inserted and cascade rules can be applied.

15. Table A has 40 rows and table B has 25 rows, joined on customer id, with 25 matching rows. (a) State the row count of an INNER JOIN. (b) State the row count of a LEFT JOIN with A as the left table. (c) Explain the practical difference between the two. [4]

Answer: (a) INNER JOIN keeps only matched rows: 25. (b) LEFT JOIN keeps all 40 rows of A and fills unmatched B columns with NULL: 40. (c) INNER JOIN is used when only records present in both tables matter, while LEFT JOIN preserves every row of the main table, for example to find customers who never ordered (B columns are NULL).

16. A table has 1,000,000 rows. Without an index a query scans every row and takes 50 ms; with an index the same query takes 2 ms. (a) Calculate the speed-up factor. (b) Explain why an index speeds up search and state one cost when inserting data. [4]

Answer: (a) Speed-up = 50 / 2 = 25 times. (b) An index is like a sorted catalogue, usually a B+ tree, so the database descends a few levels instead of scanning 1,000,000 rows, turning an O(n) scan into O(log n); the cost is that every insert or update must also update the index, using extra storage and slowing writes.

Multimedia and Streaming

Part E: Multimedia and streaming (17 marks)

17. An uncompressed video has frames of 1920 by 1080 pixels, 25 frames per second and 24 bits per pixel. (a) Calculate the uncompressed bit rate in Mbps. (b) If the encoded bit rate is 8 Mbps, calculate the compression ratio to 1 decimal place and explain why such a high ratio is possible. [4]

Answer: (a) Bit rate = 1920 x 1080 x 25 x 24 = 1,244,160,000 bits per second = 1,244.16 Mbps. (b) Ratio = 1,244.16 / 8, about 155.5:1; neighbouring frames and pixels are highly correlated, so the encoder uses motion compensation and a discrete cosine transform to store only differences and discards chroma detail the eye barely sees.

18. A streaming player downloads 30 MB of buffer before playback starts, and the video bit rate is 5 Mbps. (a) Calculate how long playback can continue if the network stops completely. (b) Explain how buffering improves viewing and why an oversized buffer causes problems. [4]

Answer: (a) Buffered data = 30 x 8 = 240 Mbit; time = 240 / 5 = 48 s. (b) A buffer absorbs jitter and short outages so playback does not stall, but an oversized buffer delays the start, consumes more memory and adds noticeable lag in live streams.

19. A 4K video has a resolution of 3840 by 2160 and an HD video is 1280 by 720. (a) Calculate the total pixels of each. (b) Calculate how many times more pixels the 4K video has. (c) State the aspect ratio of both and explain one benefit of them being the same. [4]

Answer: (a) 4K = 3840 x 2160 = 8,294,400 pixels; HD = 1280 x 720 = 921,600 pixels. (b) 8,294,400 / 921,600 = 9 times more. (c) Both are 16:9; sharing the same ratio means the same video can simply be scaled to either resolution without distortion or black bars.

20. The highest frequency humans can hear is about 20 kHz. (a) Using the Nyquist theorem, calculate the minimum sampling rate needed for digital recording. (b) Explain why CD audio uses 44.1 kHz instead of exactly 40 kHz, and what problem an insufficient sampling rate causes. [5]

Answer: (a) The sampling rate must be at least twice the highest frequency: 2 x 20 kHz = 40 kHz. (b) Real filters cannot cut off perfectly, so 44.1 kHz leaves about 4.1 kHz of transition band for the anti-aliasing filter to roll off above 20 kHz without attenuating audible content; too low a rate causes aliasing, where high frequencies are recorded as false low tones.

Security and Backup

Part F: Security and backup (17 marks)

21. A password is built from 62 available characters (26 uppercase, 26 lowercase and 10 digits). (a) Calculate the total number of 8-character passwords. (b) Calculate the total for 6 characters and how many times more combinations the 8-character password has. [4]

Answer: (a) 62 to the power 8 = 218,340,105,584,896 combinations, about 2.18 x 10^14. (b) 62 to the power 6 = 56,800,235,584, about 5.68 x 10^10; the ratio is 62 squared = 3844 times.

22. A system upgrades messages encrypted with a 56-bit key to a 128-bit key. (a) Express how many times larger the 128-bit key space is, as a power of 2. (b) If an attacker tests 10^9 keys per second, calculate the years needed to break the 56-bit key on average (half the key space), to 2 decimal places. [4]

Answer: (a) 2 to the power 128 / 2 to the power 56 = 2 to the power 72, about 4.72 x 10^21 times larger. (b) Half of 2 to the power 56 is 3.60 x 10^16 keys; divided by 10^9 gives 3.60 x 10^7 s, which is about 1.14 years.

23. A website stores user passwords as MD5 hashes, and an MD5 digest is 128 bits long. (a) Using the birthday bound, state approximately how many inputs are needed to find two with the same hash, as a power of 2. (b) Explain why a random per-user salt is added and how it defeats rainbow tables. [4]

Answer: (a) The birthday bound is about 2 to the power 64, roughly 1.84 x 10^19 inputs. (b) A salt is a unique random string hashed together with the password, so two users with the same password get different hashes and precomputed rainbow tables no longer match; the attacker must rebuild a table for every salt.

24. An organisation holds 500 GB of data and applies the 3-2-1 backup rule: 4 weekly full backups plus daily incremental backups on the other 26 days, with about 2% of the data changing each day. (a) Calculate the uncompressed backup volume for the month. (b) If backups compress 2:1, calculate the storage used and explain what the 2 and the 1 mean in the 3-2-1 rule. [5]

Answer: (a) Full backups = 500 x 4 = 2,000 GB; daily change = 500 x 2% = 10 GB, giving 260 GB over 26 days; total = 2,000 + 260 = 2,260 GB. (b) Compressed = 2,260 / 2 = 1,130 GB. The 2 means at least two different media and the 1 means at least one copy off site, so a local disaster cannot destroy every copy at once.