S6 ICT Practice Paper: Networks and Data Transmission, Encryption and Security, Databases and SQL, Algorithms and Programming, System Development and Testing, Ethics and Sustainable Computing

S6 | Information and Communication Technology | Time: 75 minutes | Total: 100 marks | 24 questions

---

Section A Networks and Data Transmission

Part A: 4 marks each, 16 marks in total.

1. A 480 MB file is sent over a link rated at 100 Mbps, but the effective throughput is only 80% of the rated value. Calculate the time required in seconds. [4]

Answer: Effective throughput = 100 x 0.8 = 80 Mbps = 10 MB/s; time = 480 / 10 = 48 s. (In bits: 480 x 8 x 10^6 / 80 x 10^6 = 48 s.)

2. An uncompressed bitmap is 3000 x 2000 pixels with 24 bits per pixel. Find its size in MB. If lossless compression reduces the size by 70%, find the compressed size in MB. [4]

Answer: Pixels = 3000 x 2000 = 6 000 000; total bits = 6 000 000 x 24 = 144 000 000 bits; / 8 = 18 000 000 bytes = 18.0 MB; compressed = 18.0 x 0.3 = 5.4 MB.

3. An IPv4 network uses a /26 prefix. State the number of usable hosts, the subnet mask in dotted decimal, and the broadcast address for 192.168.10.64/26. [4]

Answer: Host bits = 32 - 26 = 6, addresses = 2^6 = 64, usable hosts = 64 - 2 = 62; mask = 255.255.255.192; the block runs from 192.168.10.64 to 127, so the broadcast is 192.168.10.127.

4. A 60 MB file is transferred over a 50 Mbps link with a round-trip time of 40 ms, and two round trips are needed for the handshake. Ignoring other delays, find the total time in seconds. [4]

Answer: Transfer time = 60 x 8 / 50 = 9.6 s; handshake delay = 2 x 0.04 = 0.08 s; total = 9.6 + 0.08 = 9.68 s.

Section B Encryption and Network Security

Part B: 4 marks each, 16 marks in total.

5. Compare symmetric and asymmetric encryption in key distribution: for 10 users communicating pairwise, how many keys does each scheme need? Give one performance advantage and one disadvantage of each. [4]

Answer: Symmetric: one shared key per pair, 10 x (10 - 1) / 2 = 45 keys; asymmetric: one key pair per user, 2 x 10 = 20 keys. Symmetric is fast but needs secure key distribution; asymmetric solves distribution but is far slower, so it is used to exchange a session key that then protects the bulk data.

6. RSA uses primes p = 61, q = 53 and public exponent e = 17. Find n, ϕ(n) and the private key d, then encrypt and decrypt the plaintext m = 65. [4]

Answer: n = 61 x 53 = 3233; ϕ(n) = 60 x 52 = 3120; from e*d = 1 (mod 3120), d = 2753 (17 x 2753 = 46801 = 15 x 3120 + 1). Encryption: c = 65^17 mod 3233 = 2790; decryption: 2790^2753 mod 3233 = 65, so the plaintext is recovered.

7. A system uses passwords of exactly 8 characters drawn from upper case, lower case and digits (62 symbols). Find the total number of combinations, and the average time to crack it in days (one decimal place) if an attacker tries 10^9 combinations per second. [4]

Answer: Combinations = 62^8 = 2.183e+14 (about 2.18 x 10^14); exhaustive search takes 218,340 s; average (half the space) = 1.3 days. Longer passwords, salting and rate limiting are therefore needed.

8. A company has 500 staff and 3% of them clicked a fake link in a phishing test. Calculate the number who clicked and explain what a firewall and staff training each defend against. [4]

Answer: Number who clicked = 500 x 3% = 15. A firewall filters unauthorised ports and traffic, blocking external scanning and some malicious connections; phishing usually rides on normal ports 80/443, so training, mail filtering and multi-factor authentication are also needed, plus a two-person rule for payment requests.

Section C Databases and SQL

Part C: 4 marks each, 16 marks in total.

9. Explain the purpose of a primary key, a foreign key and normalisation, and why an unnormalised table causes update anomalies. [4]

Answer: A primary key uniquely identifies each record and cannot be null; a foreign key references another table key to enforce referential integrity; normalisation (1NF to 3NF) removes repeating groups and partial or transitive dependencies. In an unnormalised table the same fact is stored many times, so an update can leave inconsistent copies and a delete can lose other data.

10. Table Marks holds 5 scores: 72, 58, 91, 45, 84. Write the SQL for the average score and state the result. [4]

Answer: SELECT AVG(score) AS avg_score FROM Marks; sum = 350, average = 350 / 5 = 70.0. ROUND(AVG(score), 0) rounds it, and a WHERE clause can restrict the rows counted.

11. Given Student(student_id, name) and Enrol(student_id, class_code), write SQL that lists the number of students per class, showing only classes with more than 20 students, sorted from the largest group down. [4]

Answer: SELECT class_code, COUNT(*) AS n FROM Enrol GROUP BY class_code HAVING COUNT(*) > 20 ORDER BY n DESC; joining Student is only needed for names; aggregates are filtered with HAVING, not WHERE, and ORDER BY comes last.

12. A table has 1,000,000 records; a full scan costs about 10 microseconds per record, while a B-tree index needs 20 node reads at 5 ms each. Find the average time for each and the speed-up ratio. [4]

Answer: A full scan reads about 500,000 records on average: 500,000 x 10 us = 5.0 s; the index costs 20 x 5 ms = 100 ms = 0.1 s; speed-up = 5.0 / 0.1 = 50 times. Indexes raise write cost and storage use, and help little on low-selectivity columns.

Section D Algorithms and Programming

Part D: 4 marks each, 16 marks in total.

13. Searching 1,000,000 sorted items, how many comparisons does linear search make on average, and how many at most for binary search? Find the ratio. [4]

Answer: Linear search averages 500,000 comparisons; binary search needs at most ceil(log2 1,000,000) = 20; the ratio is 500,000 / 20 = 25,000 times. Binary search requires sorted data, so inserts must keep the order.

14. Bubble sort on 12 items in the worst case (fully reversed) needs how many comparisons? If each comparison plus swap costs 1 microsecond, find the total time. [4]

Answer: Comparisons = n(n - 1)/2 = 12 x 11 / 2 = 66; time is about 66 us = 0.066 ms. The cost is O(n^2): doubling the data roughly quadruples the comparisons, so O(n log n) sorts are used for large inputs.

15. For the naive recursive definition, how many calls does fib(10) make in total? Explain why memoisation reduces this greatly. [4]

Answer: fib(10) makes 177 calls in total because the same fib(k) is recomputed many times; the cost grows about O(1.6^n). With memoisation each k is computed once and stored, giving O(n) work, so fib(10) needs only about ten computations plus lookups.

16. Compare the number of operations for O(1), O(log n), O(n), O(n log n) and O(n^2) at n = 1024, and say which is least suitable for large inputs. [4]

Answer: O(1) = 1; O(log n) = 10; O(n) = 1,024; O(n log n) = 10,240; O(n^2) = 1,048,576. O(n^2) is the least suitable: a tenfold increase in n multiplies the work by 100, so it times out in practice; the growth order matters before constant factors.

Section E System Development and Testing

Part E: 4 marks each, 16 marks in total.

17. List the main stages of the system development life cycle (SDLC) and explain how weak requirements analysis raises later costs. [4]

Answer: Stages: feasibility study, requirements analysis, system design, coding, testing, conversion and rollout, then maintenance and review. Errors in requirements are multiplied through design and coding, and the later they surface the more rework is needed, so requirements must be confirmed with users and validated with prototypes.

18. A decision has 3 independent conditions. Under condition-combination coverage in white-box testing, what is the minimum number of test cases? Distinguish black-box from white-box testing. [4]

Answer: Each condition is true or false, so 2^3 = 8 cases are needed. White-box testing examines internal structure and paths and is designed by developers; black-box testing checks outputs against the specification and is designed from requirements, so the two complement each other.

19. Compare the waterfall model with agile development, and explain why agile suits projects whose requirements change often. [4]

Answer: Waterfall proceeds stage by stage with heavy documentation but little backtracking, suiting stable requirements; agile delivers working increments in short iterations, reviews after each sprint and keeps the customer involved, so changing requirements are corrected early. Risk is spread but the team must be disciplined and communicate continuously.

20. Explain the purpose and practice of user acceptance testing (UAT), and the role of version control such as Git in team development. [4]

Answer: UAT is run by real users on business scenarios and must pass before go-live; it uses realistic data on key workflows, logs defects and ends with sign-off. Version control records every change and author, allows branches and merges, supports rollback, prevents overwriting files and keeps an auditable history.

Section F Ethics and Sustainable Computing

Part F: 5 marks each, 20 marks in total.

21. Compare free and open-source software (FOSS) with proprietary software in licensing, cost and risk, giving one suitable scenario for each. [5]

Answer: FOSS (GPL, MIT) can be used, modified and redistributed freely, has no licence fee and is auditable, but support and maintenance fall on the user; proprietary software charges licence fees and keeps source closed, with vendor support but vendor lock-in. FOSS suits school platforms on tight budgets, proprietary licences suit core banking where contracts matter.

22. State four of the six data protection principles in the Personal Data (Privacy) Ordinance and explain how a school should handle student data. [5]

Answer: The six principles: lawful and adequate collection; accuracy and no excessive retention; use limited to the stated purpose unless consent is given; security safeguards; openly stated policies; and access and correction rights. A school should collect only what teaching needs, encrypt storage, restrict access, set retention periods and tell parents the purpose and how to request access.

23. Explain where bias can enter decisions made with big data or machine learning, and give two measures to reduce the risk. [5]

Answer: Bias enters through historically discriminatory or unrepresentative training data, features that encode protected attributes indirectly, subjective human labelling, and opaque models that hide errors. Measures: audit data distribution and representativeness, evaluate accuracy and error rates per group, keep human review and an appeal route, and publish model cards and decision grounds.

24. A data centre draws 400 kW for IT equipment and 500 kW in total. Find its PUE, then the annual emissions in tonnes if it runs 8 760 hours a year at 0.6 kg CO2 per kWh. Suggest one way to lower the PUE. [5]

Answer: PUE = 500 / 400 = 1.25; annual energy = 500 x 8 760 = 4,380,000 kWh; emissions = 4,380,000 x 0.6 = 2,628,000 kg = 2,628 tonnes. To lower PUE, use free cooling such as cold-aisle containment and higher supply temperatures, heat reuse, or virtualisation to raise utilisation.