

Markscheme

Mock examination — Set 2

Computer science

Standard level

Paper 1

This markscheme is **confidential** and for the exclusive use of the teachers marking this mock examination.

It must **not** be shared with candidates before the mock examination has been sat.

General marking instructions

1. Follow the markscheme provided, award only whole marks and mark only in **RED**.
2. Where a mark is awarded, a tick/check (✓) **must** be placed in the text at the **precise point** where it becomes clear that the candidate deserves the mark. **One tick to be shown for each mark awarded.**
3. Sometimes, careful consideration is required to decide whether or not to award a mark. In these cases write a brief comment in the margin to support your decision, especially for moderation purposes. It should be remembered that the script may be returned to the candidate.
4. Personal codes/notations are unacceptable.
5. Where an answer to a part question is worth no marks but the candidate has attempted the part question, write a zero next to the answer. Where the candidate has not attempted the part question, write “NR”.
6. Ensure that you have viewed **every** page including any additional sheets.
7. A mark should not be awarded where there is contradiction within an answer. Write “CON” next to the contradiction.

Subject details: Computer Science SL Paper 1 markscheme (Mock examination — Set 2)

Mark allocation

Section A: Candidates are required to answer **all** questions. Total 38 marks.

Section B: Candidates are required to answer **all** questions. Total 12 marks.

Maximum total = 50 marks.

General

A markscheme often has more specific points worthy of a mark than the total allows. This is intentional. Do not award more than the maximum marks allowed for that part of a question.

When deciding upon alternative answers by candidates to those given in the markscheme, consider the following points:

- Each statement worth one point has a separate line and the end is signified by means of a semi-colon (;).
- An alternative answer or wording is indicated in the markscheme by a “/”; either wording can be accepted.
- Words in (...) in the markscheme are not necessary to gain the mark.
- If the candidate’s answer has the same meaning or can be clearly interpreted as being the same as that in the markscheme then award the mark.
- Mark positively. Give candidates credit for what they have achieved and for what they have got correct, rather than penalizing them for what they have not achieved or what they have got wrong.
- Remember that many candidates are writing in a second language; be forgiving of minor linguistic slips. In this subject effective communication is more important than grammatical accuracy.
- Occasionally, a part of a question may require a calculation whose answer is required for subsequent parts. If an error is made in the first part then it should be penalized. However, if the incorrect answer is used correctly in subsequent parts then **follow through** marks should be awarded. Indicate this with “FT”.

General guidance

Issue	Guidance
Answering more than the quantity of responses prescribed in the questions	• In the case of an “identify” question, read all answers and mark positively up to the maximum marks. Disregard incorrect answers. • In the case of a “describe” question, which asks for a certain number of facts <i>eg</i> “describe two kinds”, mark the first two correct answers. This could include two descriptions, one description and one identification, or two identifications. • In the case of an “explain” question, which asks for a specified number of explanations <i>eg</i> “explain two reasons ...”, mark the first two correct answers. This could include two full explanations, one explanation, one partial explanation <i>etc.</i>

Section A

1. (a) Award **[1]** max

Decodes each instruction (held in the instruction register/IR);
Sends control signals to coordinate the other components (e.g. the ALU, registers and memory);
Controls/manages the fetch–decode–execute cycle (e.g. using the clock to time each step);
Controls the flow of data between the CPU, memory and input/output devices (along the buses);

Note: Do not accept “carries out calculations” or “performs logical operations” (these are functions of the ALU).

(b) (i) Award **[1]** max

180;

Note: $B4 = 11 \times 16 + 4 = 180$. Do not accept 60 (the red intensity) or 230 (the blue intensity).

(ii) Award **[1]** max

C8B4E6;

*Note: Accept lower case (c8b4e6) and a leading # (#C8B4E6). Award **[0]** if the green or blue digits have been changed.*

(c) Award **[3]** max

Award **[1]** for each correct column (P, R and Q)

A	B	C	P	R	Q
0	0	0	1	0	0
0	0	1	1	1	1
0	1	0	0	1	1
0	1	1	0	0	0
1	0	0	0	1	1
1	0	1	0	0	1
1	1	0	0	1	1
1	1	1	0	0	1

*Note: The eight combinations of A, B and C may be listed in any order, but each must appear exactly once; award **[2]** max if a combination is missing or repeated. Follow through: award the R mark for an R column that is correct for the candidate's P column, and the Q mark for a Q column that is correct for the candidate's R column.*

(d) Award **[2]** max

Award **[1]** per point up to **[2]** max

The image data (e.g. each block of pixels in a frame) is transformed into a different form, such as frequency components (e.g. using the discrete cosine transform);
...which shows which details have the most effect on the picture;
Components that have little visible effect (fine/high-frequency detail) are stored with less precision/rounded (quantized) or discarded;
...so fewer bits are needed to store each frame (many values become zero and can be stored very efficiently);
Some detail is permanently lost (lossy compression), but viewers barely notice the difference;

*Note: Award **[1]** max if the answer only states that data is removed or that the method is lossy, without referring to the data being transformed. Award **[1]** max if the answer describes the transform but does not state that the less visible components are discarded or stored with less precision. Do not credit a description of run-length encoding.*

(e) Award **[3]** max

Award **[1]** per point up to **[3]** max

ASCII uses 7 bits (8 bits in extended ASCII) for each character;
...so it can represent only 128 (256) different characters;
...which is enough only for English letters, digits and punctuation/Arabic, Chinese and Greek characters have no ASCII code;
...so these characters could not be stored, or would be displayed incorrectly;
Unicode gives each character in (almost) every writing system its own unique code (code point);
Unicode uses more bits per character (e.g. UTF-8 uses 1 to 4 bytes, UTF-16 uses 2 or 4 bytes), so over a million different codes are possible;
In UTF-8, English letters keep their ASCII codes/use only one byte, so English text does not take up more space;
The text is displayed correctly on any device or browser, in any country, because they all use the same codes;

*Note: Award **[2]** max if the answer does not explain the difference in terms of the number of bits used for each character or the number of different characters that each code can represent. Award **[2]** max if the answer does not state how Unicode overcomes the limitation (e.g. a unique code for every character/more bits per character).*

2. (a) Award **[1]** max

10.12.200.7;

Note: Do not accept 172.64.10.9; private addresses beginning with 172 run only from 172.16.0.0 to 172.31.255.255.

(b) Award **[4]** max

*Award **[1]** for identifying a benefit and **[1]** for expansion; award **[1]** for identifying a limitation and **[1]** for expansion. Mark as **[2]** + **[2]**.*

Benefits

Adverts start playing more quickly/with less buffering (lower latency);

...because the data travels a shorter distance/through fewer routers, so viewers are less likely to skip the advert;

Higher reliability/availability;

...if one server fails, viewers can receive the advert from another server;

Scalability;

...when an advert becomes very popular, the load is shared between many servers instead of overloading one;

Less traffic on the agency's own servers and internet connection;

...so the agency does not need to buy extra bandwidth or servers;

Limitations

Cost;

...the agency must pay for the service, often according to the amount of data delivered;

Dependence on another company;

...if the service has an outage or changes its prices, the agency's adverts are affected and it cannot fix the problem itself;

Updates take time to reach every server;

...a corrected advert may still be shown from some servers for a while (the copies are inconsistent);

Copies are stored in many countries;

...the agency must comply with different laws (e.g. on data protection or advertising) and has less control over the security of the copies;

(c) Award **[3]** max

*Award **[1]** per point up to **[3]** max*

Each file/message is divided into small packets;

Each packet has a header containing the source and destination (IP) addresses and a sequence number;

The connection is not reserved for one transfer (as it would be with circuit switching);

...packets from different computers/transfers are interleaved and sent one after another along the shared connection;

...so a large video upload does not block the connection, because packets for web pages and emails can be sent between its packets;

Routers forward each packet independently, using the destination address in its header;
Packets arriving for the agency are passed to the correct computer using the addresses (and port numbers) in their headers (e.g. by NAT on the router);
At the destination, the packets of each transfer are put back in order using their sequence numbers;

Note: Award [2] max if the answer only describes how a file is split into packets and reassembled, without explaining how the single connection is shared.

(d) Award [3] max

Award [1] per point up to [3] max

The web server accepts connections from anyone on the internet, so it is the device most likely to be attacked/compromised;
If an attacker takes control of the web server, the firewall blocks/filters traffic from the web server's segment to the staff segment;
...so the attacker cannot easily spread (e.g. malware) to the staff computers or reach the client files on the file server;
The damage from a successful attack is limited to one segment/the rest of the network keeps working;
The traffic in each segment can be monitored separately, so unusual traffic from the web server is easier to detect;

Note: Credit answers that call the separate segment a demilitarized zone (DMZ) or that refer to subnets or VLANs. Award [2] max if the answer does not explain what happens if the web server is compromised.

3. (a) Award [2] max

Award [1] per point up to [2] max

The data type of each field (e.g. VARCHAR(30) for ClientName, INTEGER for Budget);
The size/length of each field;
The indexes (e.g. on ClientID);
How/where the data is stored, e.g. the files and storage structures used;
Details specific to the DBMS being used (e.g. the storage engine or how tables are partitioned);

Note: Do not accept tables, fields/attributes, primary keys, foreign keys or relationships, which are already in the logical schema.

(b) Award [2] max

Award [1] for the UPDATE statement and [1] for the ALTER TABLE statement

The UPDATE statement is DML/it changes the data in an existing record (the budget of P103), so it is allowed and is carried out;
The ALTER TABLE statement is DDL/it changes the structure of the table (it adds a new field), so it is refused;

Note: Do not award a mark for stating which statement is carried out unless the answer classifies it as DML or DDL or explains that it changes the data or the structure.

(c) Award [3] max

Award [1] for selecting Title, Budget and ClientName from CAMPAIGN and CLIENT with a correct join condition (CAMPAIGN.ClientID = CLIENT.ClientID)

Award [1] for WHERE Sector = 'Retail' AND Channel <> 'Social' (or NOT Channel = 'Social')

Award [1] for ORDER BY Budget DESC

```
SELECT CAMPAIGN.Title, CAMPAIGN.Budget, CLIENT.ClientName
FROM CAMPAIGN
INNER JOIN CLIENT ON CAMPAIGN.ClientID = CLIENT.ClientID
WHERE CLIENT.Sector = 'Retail' AND CAMPAIGN.Channel <> 'Social'
ORDER BY CAMPAIGN.Budget DESC;
```

Alternative solution using an implicit join:

```
SELECT Title, Budget, ClientName
FROM CAMPAIGN, CLIENT
WHERE CAMPAIGN.ClientID = CLIENT.ClientID
AND Sector = 'Retail' AND NOT Channel = 'Social'
ORDER BY Budget DESC;
```

Note: Accept either quote style, the string literals in any letter case, the fields in any order, and != in place of <>. Ignore minor syntax errors. Table prefixes are only required where a field name is ambiguous (ClientID). Accept other correct ways of combining the tables, e.g. JOIN ... USING (ClientID), or a subquery on CLIENT. Do not award the first mark if both tables are listed without a join condition (this returns every combination of campaign and client). Do not award the second mark if OR is used instead of AND, if either condition is missing, if Channel = 'Video' is used in place of the Social condition (this misses Ride to

Work), or if hard-coded IDs (e.g. ClientID IN ('C02', 'C04')) are used in place of Sector = 'Retail'.

(d) Award **[5]** max


Award **[1]** for a *FREELANCER* table with *FreelancerID* as its primary key and *FreelancerName*;

Award **[1]** for a *STUDIO* table with *StudioID* as its primary key and *StudioName*;

Award **[1]** for a *SHOOT* table with *ShootID* as its primary key, *ShootDate*, and *StudioID* as a foreign key (and no *StudioName*);

Award **[1]** for a *BOOKING* table that contains *FreelancerID* and *ShootID* as foreign keys, and *Fee*;

Award **[1]** for (*FreelancerID*, *ShootID*) as the composite primary key of the *BOOKING* table;



Note: Accept any clear way of showing keys (e.g. underlining, PK/FK labels or asterisks) and any sensible table names; relationship lines are not required. For the fifth mark, accept a surrogate primary key (e.g. BookingID) only if FreelancerID and ShootID are both included as foreign keys. Do not award the first mark for a FREELANCER table that contains Fee, the third mark for a SHOOT table that contains StudioName, or the fourth mark for a BOOKING table that contains FreelancerName, ShootDate, StudioID or StudioName. Award **[3]** max if StudioName has not been separated from SHOOT.

4. (a) Award [2] max

Award [1] for describing what training requires and [1] for describing what making a prediction requires

Training processes a very large data set many times, adjusting the model's (millions of) parameters as it goes;

...which needs a huge number of calculations (best done in parallel by GPUs) and a lot of memory;

Making a prediction only applies the finished model, whose parameters are fixed, to one visitor's data;

...which needs far fewer calculations, so an ordinary server CPU can do it quickly enough;

Note: Award [1] max if only training or only making a prediction is described.

(b) Award [2] max

Award [1] for identifying a way and [1] for expansion

Financial harm/debt;

...people who are already struggling may be encouraged to borrow money that they cannot repay, so their debt increases;

Exploitation at a vulnerable time;

...the adverts reach people when they are under pressure, so they may make rushed decisions without comparing other options;

Loss of privacy;

...a sensitive fact (their financial difficulty) is inferred from their browsing without their knowledge or consent;

Unfair treatment/bias;

...the predictions may rely on data such as location that act as proxies for income or ethnicity, so some groups are targeted more than others;

Security/fraud;

...data that identifies people in financial difficulty could be stolen and used by criminals to target them with scams;

Distress;

...repeatedly seeing loan adverts may make people feel watched or increase their anxiety about money;

Section B

5. (a) *Award [2] max*

Award [1] for why more steps take longer and [1] for why more steps give finer detail

Each step is a full pass of the whole image through the denoiser (CNN), involving a very large number of calculations, so 50 steps need about five times as much computation as 10 steps; Each step uses the result of the previous step, so the steps must be carried out one after another and cannot be carried out in parallel to save time;

Each step removes only part of the noise, so with more steps the image is refined more gradually and fine features such as edges and textures can be reconstructed; With fewer steps, more noise must be removed at each step, so the denoiser's predictions are less accurate, giving blurred/less detailed images or artefacts; The denoiser is trained to remove noise gradually (as formalized in the DDPM), so more steps follow this gradual process more closely;

Note: Do not award the first mark for only stating that 50 steps is more than 10 steps, or the second mark for only stating that more steps give "better quality".

(b) *Award [4] max*

Award [1] for identifying a way and [1] for expansion

Use an embedding-based approach, e.g. learn a new embedding/special word for the mascot from a few reference images (e.g. textual inversion); ...which is included in every prompt, so the model generates the same learned features each time;

Fine-tune the model on a small set of images of the mascot (transfer learning, e.g. DreamBooth/LoRA); ...so the model learns the mascot's exact appearance and can place it in new scenes;

Give the model an approved image of the mascot as an extra input/reference image (e.g. an image prompt adapter); ...the model is conditioned on an embedding of the reference image, so its features are carried into each new image;

Use image-to-image translation, starting from an approved image or a sketch of the mascot; ...only part of the noise is added, so the mascot's shape and features are kept while other details change;

Use conditional inputs such as an outline, pose or segmentation map of the mascot (e.g. ControlNet); ...so the shape and position of the mascot are fixed in every image;

Keep an approved image of the mascot and generate only the scene around it (inpainting); ...so the pixels of the mascot are not regenerated and cannot change;

Use the same detailed description of the mascot and the same random seed in every prompt; ...so the starting noise and wording stay the same, which reduces variation (award [1] max for this way, because changing the scene still changes the mascot);

Review every image against a reference sheet/style guide for the mascot; ...designers reject or edit images in which the mascot differs, so only consistent images reach the client;

Note: Do not credit two ways that are the same strategy, e.g. two named tools that both learn an embedding for the mascot, or a reference image and image-to-image translation (both regenerate the mascot from an approved image, so credit only one of them; inpainting, which keeps the mascot's pixels unchanged, is a separate way). Do not credit "use the same prompt" on its own, because the stimulus shows that this does not keep the mascot consistent.

Mark as **[2]** + **[2]**

(c) Award **[6]** max

Responses may include some of the following points. The best ones will demonstrate research and give a balanced perspective before presenting a conclusion. Named examples, studies and regulations below are illustrative only; candidates are not required to know them.

Note: Answers that consider only AI disclosure (or only bias mitigation) are unlikely to go beyond the Adequate band.

How bias mitigation can maintain trust

- Dataset curation and auditing: Checking the training dataset for under-represented groups and rebalancing it (e.g. adding licensed images of those groups) reduces exclusionary or inaccurate representations (lines 101–102), so campaigns reflect clients' audiences and are less likely to cause offence.
- Testing outputs: Generating many images from neutral prompts (e.g. "a doctor", "a manager") and checking the range of people shown allows the studio to detect stereotypes before images reach a client.
- Human review: Diverse review teams check campaign images before release, so the studio, not the model, is accountable for what audiences see.
- Reputation: Clients are more likely to trust a studio that can show a documented bias mitigation process, and fewer complaints protect the client's brand.

Limits of bias mitigation

- Bias can be reduced but not removed: the people who select and label images have their own biases, and very large datasets are too big to check image by image.
- Over-correction can also damage trust: in February 2024 Google paused Gemini's image generation of people after it produced historically inaccurate images (e.g. of 1943 German soldiers).
- If the studio uses an external model, it cannot audit the training data and can only partly control bias through prompts and review.
- Bias mitigation does not address other reasons for distrust, such as intellectual property concerns about the training data (lines 99–100) or the effect on human artists.

How AI disclosure can maintain trust

- Clear guidelines (lines 103–105): Telling clients in advance how and where AI is used in the creative process means that they are not surprised later; discovering undisclosed AI use is likely to damage trust more than disclosure.
- Labels and metadata: Visible labels and machine-readable records such as C2PA Content Credentials (added to DALL·E 3 images in February 2024) or invisible watermarks (e.g. Google DeepMind’s SynthID, launched in August 2023) let audiences and platforms check whether an image was AI-generated.
- Legal compliance: Regulations increasingly require disclosure, e.g. the EU AI Act’s transparency obligations, which apply from 2 August 2026, require AI-generated images to be marked in a machine-readable format (with extra time, to December 2026, for systems already on the market) and deepfakes to be disclosed; China’s labelling measures for AI-generated content took effect on 1 September 2025. Compliance protects clients from legal and reputational risk.
- Honesty and respect for audiences support long-term trust in the client’s brand.

Limits of AI disclosure

- As the stimulus suggests, disclosure may reduce trust in the short term: research (e.g. Cicek, Gursoy and Lu, 2024) found that describing products as using AI lowered consumers’ emotional trust and purchase intentions, and Coca-Cola’s AI-generated Christmas advert (November 2024) was criticized on social media as “soulless”.
- Metadata can be removed, e.g. by screenshots or when images are uploaded to social media, and watermarks are not perfectly robust, so the disclosure may not reach the audience.
- Labels applied too broadly can mislead: in 2024 Meta changed its “Made with AI” label to “AI info” after photographers complained that real photos with minor edits were labelled as AI-generated.
- Disclosure shows that AI was used but not whether the images are fair or how they were made, so it does not on its own make the content trustworthy.

Conclusion

A reasoned conclusion that expresses an opinion and is supported by evidence.

Marks	Level descriptor
No marks	• No knowledge or understanding of the relevant issues and concepts. • No use of appropriate terminology.
Basic 1–2 marks	• Minimal knowledge and understanding of the relevant issues. • Minimal use of appropriate terminology. • The answer may be little more than a list. • No reference is made to the case study or independent research.
Adequate 3–4 marks	• A descriptive response with limited knowledge and/or understanding. • A limited use of appropriate terminology. • There is some analysis. • There is evidence of some research.
Competent 5–6 marks	• Knowledge and understanding of the related issues and/or concepts. • Uses terminology appropriately in places. • There is evidence of analysis. • There is evidence of research. • There is a conclusion.
