

Markscheme

May 2019

Computer science

Higher level

Paper 3

No part of this product may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without written permission from the IB.

Additionally, the license tied with this product prohibits commercial use of any selected files or extracts from this product. Use by third parties, including but not limited to publishers, private teachers, tutoring or study services, preparatory schools, vendors operating curriculum mapping services or teacher resource digital platforms and app developers, is not permitted and is subject to the IB's prior written consent via a license. More information on how to request a license can be obtained from <http://www.ibo.org/contact-the-ib/media-inquiries/for-publishers/guidance-for-third-party-publishers-and-providers/how-to-apply-for-a-license>.

Aucune partie de ce produit ne peut être reproduite sous quelque forme ni par quelque moyen que ce soit, électronique ou mécanique, y compris des systèmes de stockage et de récupération d'informations, sans l'autorisation écrite de l'IB.

De plus, la licence associée à ce produit interdit toute utilisation commerciale de tout fichier ou extrait sélectionné dans ce produit. L'utilisation par des tiers, y compris, sans toutefois s'y limiter, des éditeurs, des professeurs particuliers, des services de tutorat ou d'aide aux études, des établissements de préparation à l'enseignement supérieur, des fournisseurs de services de planification des programmes d'études, des gestionnaires de plateformes pédagogiques en ligne, et des développeurs d'applications, n'est pas autorisée et est soumise au consentement écrit préalable de l'IB par l'intermédiaire d'une licence. Pour plus d'informations sur la procédure à suivre pour demander une licence, rendez-vous à l'adresse <http://www.ibo.org/fr/contact-the-ib/media-inquiries/for-publishers/guidance-for-third-party-publishers-and-providers/how-to-apply-for-a-license>.

No se podrá reproducir ninguna parte de este producto de ninguna forma ni por ningún medio electrónico o mecánico, incluidos los sistemas de almacenamiento y recuperación de información, sin que medie la autorización escrita del IB.

Además, la licencia vinculada a este producto prohíbe el uso con fines comerciales de todo archivo o fragmento seleccionado de este producto. El uso por parte de terceros —lo que incluye, a título enunciativo, editoriales, profesores particulares, servicios de apoyo académico o ayuda para el estudio, colegios preparatorios, desarrolladores de aplicaciones y entidades que presten servicios de planificación curricular u ofrezcan recursos para docentes mediante plataformas digitales— no está permitido y estará sujeto al otorgamiento previo de una licencia escrita por parte del IB. En este enlace encontrará más información sobre cómo solicitar una licencia: <http://www.ibo.org/es/contact-the-ib/media-inquiries/for-publishers/guidance-for-third-party-publishers-and-providers/how-to-apply-for-a-license>.

Subject details: **Computer science HL paper 3 markscheme**

Mark allocation

Candidates are required to answer **all** questions. Total 30 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**”.
- Question 4 is marked against markbands. The markbands represent a single holistic criterion applied to the piece of work. Each markband level descriptor corresponds to a number of marks. When assessing with markbands, a “best fit” approach is used, with markers making a judgment about which particular mark to award from the possible range for each level descriptor, according to how well the candidate’s work fits that descriptor.

General guidance

Issue	Guidance
Answering more than the quantity of responses prescribed in the questions	<ul style="list-style-type: none"><li data-bbox="368 333 1481 398">• In the case of an “identify” question read all answers and mark positively up to the maximum marks. Disregard incorrect answers.<li data-bbox="368 405 1481 501">• 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.<li data-bbox="368 508 1481 633">• 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>

1. (a) *Award [2 max].*
Award [1] for any valid function up to [2 max].
Stores/retrieves/searches/edits the data;
Encapsulates/holds the code for the database (accept DML / SQL or any other language);
Abstracts/Simplifies access to the databases;
Personalises the SQL for different databases (e.g. contains optimised queries for Oracle, MySQL etc);
Manages the transactions (anything relating to ACID); [2]
- (b) *Award [2 max].*
Award [1] for any characteristic up to [2 max].
Binding anything relevant (e.g. nodes, port, IP);
Associating a socket with a socket address;
Associated with the IP address and a port number for the local node;
IP address;
Port number;
Sending/receiving data (packets);
Node-to-node communication;
Works at OSI levels 3, 4, and 5 (Accept any from L3 - Network, L4 - Transport, L5 - Session); [2]
2. (a) (i) *Award [2 max].*
Award [1] for any advantage up to [2 max].
Tried and tested;
So, less risk;
- Regularly updated;
Therefore, contains fewer errors;
- More tech support;
e.g. Online chat systems.
- Usually available for various platforms / operating systems;
Therefore, fewer compatibility issues, or can be run on different HW and SW;
- Shorter timeframe for implementation;
Which is important in the context of saving lives/obtain benefits more quickly;
- Books/Training material more accessible;
Easier to train staff on using the system;
- Reviews and feedback exist already for the product;
Therefore, the quality of the final product can be known in advance;
- Off-the shelf is likely to implement common standards;
Whereas, you would need to program compatibility in a bespoke system/to work with a multitude of devices; [2]

(ii) *Award [2 max].*

Award [1] for any disadvantage up to [2 max].

Modifications are not possible without help of developers;

Source code is not usually available to change;

Compatibility/interfaces with existing systems;

It may not be possible to connect / share data with systems already in use;

Standard operating procedures may need to change;

To fit the software capabilities/functions;

Software may contain unused features;

That is potentially confusing for users;

Software is designed for general requirements;

It may not include features needed by Bangbai;

[2]

(b) *Award [4 max].*

Award [1] for explaining what URL-rewriting is in the context of stateful connection and allow [1] for an example or a comparison with session cookies

URL Rewriting changes the URL to include parameters
e.g. www.bangbai.com?userId=Asdf34e3

Mark as [2] and [2].

Works without cookies/cookies disabled/as an alternative to cookies;
URL re-writing will still work as it embeds the session ID inside the request;

URL Rewriting does not result in anything being saved (e.g. cookies) on the client side;

Therefore, the privacy of the client cannot be compromised (e.g. By searching contents of the cookie folder later);

May work with legacy devices/devices without access to cookies;
That are not able to use cookies/do not have access to an alternative to URL re-writing;

Makes the links more descriptive/readable;
Which will help the developer resolve issues (accept example, such as server identified in URL);

(Note to examiners: descriptive cannot be awarded a second mark if not related to the scenario e.g. search engines is not correct)

Allows the details of the identification of the user to be recorded in log files;
Which will help for future analysis/reviews/data mining;

[4]

3. Award **[6 max]**.
Award **[1]** for each point that explains why the use of VPN reduces the effectiveness of the app up to **[6 max]**.

A VPN hides the true IP address;
Instead, it provides the IP address of the VPN's server;
A VPN adds an additional layer of complexity has been added;
Which may increase the likelihood of the connection dropping;
A VPN encrypts/tunnels the data which is transmitted;
Which has a potential overhead (more computational power/increases the size of data packets) and this may reduce emergency responses times;
A VPNs may log traffic that passes through;
This may present a potential security concern if sensitive information is reported;
A VPN may employ a filtering policy/blocks ports;
So certain parts of the app may not work (e.g. VOIP protocols may not work);
If GPS/location service is deactivated, then the system is reliant upon IP Address;
Location information can be obtained from the GPS/GPS is sent in data packets;
However, it may not be as easily verifiable by comparing with a known IP location;
Not having the true IP address would mean that the EMIS server would not have any indication of the location of the source of the report;
VPN usage may give users a feeling of anonymity which may encourage misuse;
Which may divert resources to the hoaxes instead of authentic reports;

[6]

4. *The answer could include the following:*

Definitions:

The source IP hash combines the source and destination IP address to generate a hash key, which is then assigned to a specific server. The benefit of this approach is that a client who experiences a dropped connection can be returned to their session with the same server.

The client-side random load balancing algorithm delivers a list of server IPs to the client, which then selects a server IP at random.

Security

As Client Side Random allows control of the choice of server to be made by the client it would leave the system much more vulnerable to DOS/DDOS attacks.

Persistence

If disconnected, Client Side Random may allocate a new server to the client as it will choose randomly from the list. However, if the IP is the same, then Client-side Hash will pair the client with the previous server. This may mean that the work being done before is not lost (e.g. Client continues from where she left of seamlessly).

Distinguishing between users

Client-Side Hash is affected when a single IP address is used by many users (e.g. proxy servers sharing a connection, VPNs etc). A single user might also disconnect and re-connect with a different IP address due to DHCP so would be seen as a “different” person.

Cost

Client Side Random is a cheaper solution as it doesn't require any expensive hardware on the server side (i.e. the hardware load balancers etc).

Single Point of Failure (SPOF)

Client Side Random eliminates the SPOF as all requests do not depend on one device (i.e. the hardware load balancer provided by the system). Similarly, it avoids **bottlenecks** if the load balancer is overloaded.

Control

When using client side random, Bangbai Government is essentially losing control over the load balancing process. If for some reason they suddenly want to re-direct all queries away from certain servers (e.g. It's been compromised, or the data centre has a fire etc) they can't, as they cannot easily/quickly change the algorithm on each of the client devices.

Intelligence/Adaptiveness/Flexibility

Client Side Random may not take into consideration privileged information about the servers, which it does not have access to. Therefore, it cannot decide which of the servers is **currently** best suited to handle a specific type Such information may include:

- Server Resources are different (Primary Memory, Secondary Memory, Processor)
- Deployment of the services to different servers (Server A and B deal with video reports, Server C with simple text/photo reports etc).
- Current load of each server (Server B has more connections than server A)
- Current Health of each server (Server C just died)

Additional Research

Caching can affect a lot of load balancing algorithms as they use DNS to allocate the next server. This can skew the allocation of servers, but because Client-side Random uses a randomly generated server from a simple list, it avoids this problem.

The HTTP/2 specification, which is now supported by every major browser, has built in support for Client-Side Load Balancing. The citizen reporting app can benefit from this as it uses HTTP for the REST API.

Overall comparison/evaluation

The main issue with the source IP hash load balancing algorithm is that each change can redirect EVERYBODY to a different server, which again lags the time of response for a citizen call in Bangbai.

That is why some good load-balancers have implemented a consistent hashing method to ensure that if a server fails, for example, only the client connected to this server are redirected and not all.

The counterpart of consistent hashing is that it doesn't provide a perfect hash, and so, in a farm of 3 servers, some may receive more clients than others and this can take a toll on the time of response to citizen calls.

When a failed server comes back, its users (determined by the hashing done on the Source IP) will be redirected to it again.

There is no overhead in terms of CPU or memory when using such an algorithm.

Conclusion

A final measured conclusion that links together the various points and recommends one or both of the algorithms as appropriate for the needs of Bangbai.

[12]

Marks	Level descriptor
No marks	<ul style="list-style-type: none"> No knowledge or understanding of the relevant issues and concepts. No use of appropriate terminology.
Basic 1–3 marks	<ul style="list-style-type: none"> Minimal knowledge and understanding of the relevant issues or concepts. Minimal use of appropriate terminology. The answer may be little more than a list. No reference is made to the information in the case study or independent research.
Adequate 4–6 marks	<ul style="list-style-type: none"> A descriptive response with limited knowledge and/or understanding of the relevant issues or concepts. A limited use of appropriate terminology. There is limited evidence of analysis. There is evidence that limited research has been undertaken.
Competent 7–9 marks	<ul style="list-style-type: none"> A response with knowledge and understanding of the related issues and/or concepts. A response that uses terminology appropriately in places. There is some evidence of analysis. There is evidence that research has been undertaken.
Proficient 10–12 marks	<ul style="list-style-type: none"> A response with a detailed knowledge and clear understanding of the computer science. A response that uses terminology appropriately throughout. There is competent and balanced analysis. Conclusions are drawn that are linked to the analysis. There is clear evidence that extensive research has been undertaken.

[12]

Total: [30]