ASSESSMENT and QUALIFICATIONS ALLIANCE

General Certificate in Education

Computing 2510

COMP3 Problem Solving, Programming, Operating Systems, Databases and Networking

Report on the Examination

2010 examination – June series

Further copies of this Report are available to download from the AQA Website: www.aqa.org.uk

Set and published by the Assessment and Qualifications Alliance.

The Assessment and Qualifications Alliance (AQA) is a company limited by guarantee registered in England and Wales (company number 3644723) and a registered charity (registered charity number 1073334). Registered address: AQA, Devas Street, Manchester M15 6EX

General

This is the first time that COMP3 had been examined and it was pleasing to see the quality of many responses. In particular, topics such as Turing machines and syntax diagrams, which are new to this specification, were tackled equally well as topics that appeared on the previous Computing specification. It appears that candidates had enough time to complete all of the questions.

Question 1

Part 1a: The advantages of serial data transmission over parallel data transmission were well understood. The most common correct responses were that data could be transmitted over longer distances and that the cabling cost would be lower as fewer wires would be required. Answers that referred to lower costs without any further explanation did not gain credit. Some candidates also discussed the problem of data skew in relation to parallel transmission. Candidates need to be aware that data going out of synchronisation is not the same as data skew occurring. Similarly, interference/noise is not the same as crosstalk. Lack of synchronisation and interference are problems that can occur with both serial and parallel communications.

Part 1b: Many candidates were able to correctly calculate the parity bit and to identify appropriate values for both the start and stop bits. Credit was awarded for the start and stop bits provided they both had different values.

Part 1c: Candidates demonstrated a reasonable understanding of what asynchronous data transmission is, but descriptions of the purposes of the start and stop bits were very superficial, the most common response being that the start bit indicated where the data started and that the stop bit indicated the end of the data. To gain credit candidates needed to explain that the start bit is used to synchronise the receiver to the data that is being transmitted and that the stop bit is used to give the receiver time to process the data just received or to allow it to recognise subsequently the next start bit. Some candidates lost marks by referring to "it" when it was far from clear what "it" was.

Question 2

Part 2a: Most candidates scored at least one mark for this question part by explaining that the operating system provided a virtual machine or abstracted the user from the complex hardware. Some however only described the provision of a user interface, which was not sufficient to gain the mark for abstraction from the hardware.

Part 2b: Most candidates showed some understanding of the differences between the two types of operating system, with many scoring high marks. The most commonly made mistake was to discuss the same point repeatedly, usually that a desktop operating system would perform a wider range of tasks. Good answers also referred to upgradability, the range of software applications that could be run, the range of hardware that might be supported and the likely minimum hardware requirements of the operating system. Some candidates assumed erroneously that embedded systems and real time systems were the same things. On this occasion we were lenient when marking answers that had made this assumption.
The vast majority of candidates wrote appropriately structured answers with only minor errors in grammar and spelling. A bigger problem was the legibility of some answers, but very few were impossible to read.

Question 3

Part 3a: Most candidates were able to write correctly the largest positive number that could be represented in the given normalised floating point system.

Part 3b: This question part was well answered, with many candidates getting both marks. The most common error was to forget that the most significant bit of the mantissa had a negative value, resulting in an answer of 10.5 rather than -5.5.

Part 3c: The majority of candidates got full marks for this question part. Many of those who failed to do this clearly knew the method to use, but made arithmetic errors.

Part 3d: This question part was well answered, with many candidates correctly calculating both the mantissa and exponent. A small number calculated the mantissa correctly, but failed to normalise this properly, resulting in an incorrect value for the exponent.

Part 3e: This question part was very well answered. Most candidates correctly explained that overflow occurs when the result of a calculation is too large to store in the available number of bits. Many were also able to give an appropriate example of how this might occur.

Question 4

Part 4a: The vast majority of candidates scored full marks for this question part, correctly identifying whether or not the numbers were syntactically valid.

Part 4b: There were some good responses to this question part, although some candidates clearly found it quite difficult. This is the first question that has been asked on this topic so, on this occasion, we were lenient when faced with minor syntactical errors. A small number of candidates confused Backus-Naur Form with regular expressions.

Question 5

Part 5a: Most candidates were able to identify both the time complexity of the algorithm from the graph and also to indentify which of the three algorithms was the most time-efficient.

Part 5b: An intractable problem is one which has a solution, but this solution cannot be found by any algorithm that has polynomial time complexity or better. Most candidates identified the fact that a solution to an intractable problem could not be found in polynomial time so gained one mark but only a small number went on to explain that the problem could be solved to gain the second mark.

Many candidates scored at least one mark for their description of how an intractable problem might be 'solved', with the most common response being that heuristic knowledge would be applied. Some candidates went on to explain what this meant to get a second mark. Other good points included simplifying the problem itself by relaxing some constraints and using an algorithm that produced an acceptable but non-optimal solution. Answers which simply stated that guesswork would be used did not gain credit, as a heuristic is more than simply guessing an answer.

Question 6

Part 6a: Responses to this question part were poor. More candidates were able to state what it meant for a database to be in Third Normal Form (3NF) than were able to explain why it was important. Various definitions of 3NF are possible. The most commonly offered one was that a relation should contain no repeating groups, no partial dependencies and no non-key dependencies. A repeating group is not the same as have repeating attributes. A repeating group refers to the use of non-atomic data whereas repeating attributes occur when multiple fields with similar purposes occur in a relation, e.g. OptionSubject1, OptionSubject2, etc.

Putting a database into 3NF will ensure that there is no redundant data and therefore eliminates the possibility of data inconsistency. It also prevents update, insertion and deletion anomalies such as not being able to add a member to a library database until she borrows a book. It does not ensure the integrity of the database nor does it eliminate duplicated data. Some candidates made very general statements such as "more efficient", "easier to use" and "better organised" which are all insufficient at this level.

Part 6b: The entity-relationship diagram was reasonably well completed, although the number of candidates scoring both marks was slightly disappointing given that this was a straightforward example.

Part 6c: Many candidates gained full marks for this question part, but there were also a lot of candidates who either made no response or clearly had no idea how to use SQL as a DDL. The accuracy of the SQL syntax was poor. The most common errors were to forget to include a data type for the MagazineID field and to use a text data type for the subscription rate. Some candidates simply wrote out the list of fieldnames. A small number of candidates appear to have tried to learn how to answer this type of question from old CPT5 mark schemes and gave both of the two possible ways included on these for defining a primary key.

Part 6d: The quality of responses to this question part varied quite significantly between centres, possibly reflecting which students had used SQL as part of their COMP4 projects. There were many good responses which scored full or almost full marks. The most common mistakes were to use "GET" instead of "SELECT", to miss out "AND" between conditions and to be confused between the symbols for less than and greater than. A small number of candidates errantly wrote tbl before table names or put semicolons at the end of each line. Some candidates added an ORDER BY clause, even though this had not been asked for. They were not penalised for this.

Part 6e: The use of SQL for data manipulation was less well understood than its use for performing queries. Most candidates picked up at least one mark, but the level of understanding was poor.

Question 7

Part 7a: The use of the adjacency matrix was clearly well understood with all but a few candidates achieving full marks.

Part 7b: There were some good responses to this question part, but also quite a lot of confused answers. An adjacency matrix is more appropriate when there are many edges in a graph, or if these edges need to be checked or updated frequently. An adjacency list is appropriate for graphs with few edges (sparse graphs) or where the edges are not checked/updated frequently. Neither the number of vertices in a graph nor the available memory would influence the choice. There was confusion over the use of terminology with some candidates apparently using the term vertex to mean edge.

Part 7c: This question part was poorly answered, with only a third of candidates scoring any marks. A tree is a graph that is connected, undirected and has no cycles. Some candidates gave responses that referred only to specific types of tree, either rooted trees or binary trees. These responses did not gain credit.

Part 7d: The vast majority of candidates knew how to construct a binary search tree. The most common cause of error appeared to be candidates forgetting the order of the letters in the alphabet rather than forgetting the principles that should be used to construct the tree. A small number of candidates mistakenly drew arrows instead of lines between nodes.

Part 7e: There were some good responses to this question part but many were disappointing and a surprising number of candidates did not write a response at all. Many candidates who did answer chose to use a diagram to illustrate their response which was quite acceptable, so long as the diagram included enough detail to make clear that it was a representation of a binary tree.

Question 8

Part 8a: The vast majority of candidates scored both marks for this question. Candidates who only scored one mark usually named the states correctly.

Part 8b: Most candidates gained some of the four available marks and over half gained all four. Some chose to keep the position of the tape head fixed and move the tape, rather than vice-versa which was perfectly acceptable.

Part 8c: The Turing machine outputted 'e' if the tape contained an even number of ones and 'o' if the number of ones was odd. Determining this was a difficult task given the limited trace that candidates were asked to complete. Nevertheless a quarter of candidates were able to do this. Many who did not get the correct answer had managed to understand that the use of the Turing machine related to evenness but that this was whether the number itself was odd or even. A commonly made mistake was to assume that the output of the Turing machine depended only on whether the last digit read was a 0 or 1.

Part 8d: It was pleasing to see that many candidates understood that a problem is computable if and only if it can be computed by a Turing machine. Some went on to make a further point, such as that no computer could be more powerful than a Turing machine, but answers scoring both marks were rare.

Question 9

Part 9a: Many candidates scored full marks by identifying three appropriate IP addresses. A small but not insignificant number of candidates wrote IP addresses consisting of just three decimal numbers, clearly indicating that they did not understand the concept. Candidates also need to be aware that the last octet of an IP address cannot contain certain values. Some lost marks by giving this as either 0 or 255.

Part 9b: The overwhelming majority of candidates correctly identified the topology that was used.

Part 9c: Many candidates were able to state the subnet mask correctly but few were able to explain its purpose. A commonly held misconception, presumably derived from the word mask, was that it would hide a computer's IP address from other computers. The purpose of the subnet mask is to allow a communicating device to determine whether or not another device that it is communicating with is on the same subnet, so that data can be sent to it directly, or a different subnet in which case communication must be via the router.

Part 9d: There were many very good responses to this question. Candidates used the full range of responses that were given in the mark scheme, with the most popular ones being the use of encryption and needing a key to access the network. A small number of candidates lost marks by discussing general network security measures, such as the use of a firewall, when the question had asked specifically about securing the connection to the wireless access point. Some candidates understood that access could be limited based upon the address of the computer, but thought that the IP address would be used rather than the MAC address and so did not gain credit.

Part 9e: Responses to this question part were mixed. The majority of candidates scored some marks, but only a tenth of candidates provided responses that were awarded all six marks. Some candidates lost marks through weak descriptions. A good example of this is that a certain time is not the same thing as a random time.

The Quality of Written Communication was similar to responses to question part 2b.

Question 10

Part 10a: This question part was very well answered with the many candidates getting full marks.

Part 10b: The algorithm inserted an item into an ordered list at the correct position. Some candidates lost marks by failing to refer to the 'correct position'. Others stated that a sort was performed which was not true, as the order of the items in the existing list was not changed.

Part 10c: For part (i), around half of the candidates were able to explain that a static data structure has a fixed size which could not change at compile time, whereas the size of a dynamic data structure can change at runtime.

There were few correct answers to part (ii). The heap is a pool of available memory locations and as the linked list grew, memory would be allocated to it from the heap and deallocated memory would be returned to it. Some candidates erroneously believed that the heap would be used as a temporary store or that the pointers in the linked list would be stored in the heap.

Question 11

Part 11a: Most candidates were able to explain what inheritance meant. Marks were sometimes lost because candidates wrote about a child class sharing 'things' with its parent, rather than methods or properties.

Part 11b: The inheritance diagram was poorly drawn, surprisingly. Many candidates failed to put boxes around the class names or to put arrowheads on the lines and consequently lost marks. Candidates need to remember to do both of these things.

Part 11c: Answers to this question part were quite mixed. Overriding is when a method from a parent class is redefined with the same name by a child class to perform a different function.

Part 11d: Most candidates got some marks for this question part but full mark responses were seen very infrequently. Common mistakes were to incorrectly redeclare the fields and methods from the parent class and to forget to override the PlayFile function.

Mark ranges and award of grades

Grade boundaries and cumulative percentage grades are available on the **Results Statistics** page of the AQA Website.