Version 1.0 0712

AQA

General Certificate of Education (A-level)
June 2012

Computing COMP3

(Specification 2510)

Unit 3: Problem Solving, Programming, Operating Systems, Databases and Networking

Report on the Examination

General

The COMP3 module is the only theory module that is examined at A2 and so the question paper necessarily covers a lot of topics. This year, students' responses suggested that they were better prepared across the full range of topics than was the case in 2010 and 2011. The areas that students tackled least well were hashing and comparing server-based and peer-to-peer, and thick and thin-client network types. The ability of students to work with floating point numbers remained as impressive as in previous years.

Question 1

This question was very well answered, with the majority of students being able to describe three different types of management that would be performed by an operating system and thus scoring full marks.

Question 2

Part 2a: Approximately three quarters of students correctly recognised that Backus-Naur Form had been used. There was a considerable variation in the spelling of the term but, as long as the meaning was clear, spelling mistakes were not penalised in this question part.

Part 2b: The vast majority of students correctly identified which of the two statements was valid.

Part 2c: This question part introduced the idea of using a parse tree to demonstrate that an expression was valid. How the parse tree was built was explained in the question and the tree was partially completed to further aid students. Most students made a reasonable attempt at completing the rest of the tree, with nearly half achieving at least two of the three available marks. The most commonly made error was to treat "21" a single digit instead of decomposing it into two digits.

Question 3

Part 3a: This question part was well answered, with most students correctly identifying that space complexity was the second measure of complexity (in addition to time complexity).

Part 3bi: In this question part, students were required to carry out a trace of a simple algorithm for looking for duplicate values in a file. The primary purpose of this question part was to give students the opportunity to study the algorithm so that they could discuss its complexity in question parts 3bii and 3biii. The majority of students achieved all three marks for fully completing the trace table. However, disappointingly, approximately a third of students scored no marks at all for what was a relatively straightforward trace to complete.

Part 3bii: Pleasingly, approximately two thirds of students recognised that the algorithm had order of time complexity $O(n^2)$.

Part 3biii: Students justified the order of time complexity as being $O(n^2)$ in various ways. Some did it by analysing the code and recognising that there was a for loop nested inside another for loop, with each loop repeating n times. Others did it by considering the number of comparisons made and recognising that each of the n items in the file would be compared to each other item, resulting in n^2 comparisons being made. The most commonly made mistake was to refer to there being two loops in the algorithm, without making clear that these were nested. Some responses were too superficial to be creditworthy, failing to relate the complexity to the steps involved in the algorithm.

Question 4

Part 4a: The relationship between bandwidth and bit rate was well understood. A small number of students failed to achieve the available mark because they simply defined the terms instead of explaining how the bit rate was determined by the bandwidth.

Part 4b: There were many good responses given to this question part. Credit was awarded to responses that were either given in the context of data communication or were generic definitions of latency. In context, latency is the time delay between when a signal is transmitted and when it is received. Some students clearly knew that latency related to time, but gave responses that lacked technical accuracy, such as, “the time delay between data being sent and a reply being received,” or, “the time taken to transmit.”

Part 4c: The correct relationship in the example was that the bit rate was double the baud rate. Somewhat disappointingly, just under half of the students recognised this, with a small number stating the relationship the wrong way around.

Part 4d: The diagram of the Moore machine was very well completed, with the vast majority of students achieving all four marks. A small number labelled the states but forgot to label the transitions.

Question 5

Part 5a: This question part asked for a definition of encryption (using an algorithm and a key to convert message data into a form that is not understandable without the key to decrypt it). Approximately three quarters of students were able to provide a suitable definition.

Part 5bi: This question part was well answered, with most students recognising that B would not have A’s private key so could not decrypt the message. Some students did not understand the asymmetric nature of the process and so wrote responses that assumed that if A’s public key was used to encrypt the data, the same key would need to be used to decrypt it.

Part 5bii: This question part was far less well answered than part 5bi. The correct answer was that this would be insecure as A’s public key, which would be used for decryption, is available to anyone. As in the previous part, some students lost marks because they did not recognise the asymmetric nature of the encryption. The response that, “anyone with A’s public key could decrypt the message,” was not considered to be enough for the mark as it did not make clear that everyone could get this key.

Part 5biii: This question part required students to explain the purpose of a digital signature and how digital signatures are used. Approximately three quarters of students were able to explain the purpose. More disappointingly, less than half were able to describe how digital signatures were used. Nevertheless, students who knew the topic provided excellent, detailed explanations. The most commonly made mistakes were: to be confused between the hash and the digital signature; to believe that the digital signature was attached to the end of the original message after the original message was encrypted rather than before; and to provide an unclear description of how the hash would be regenerated at the receiver and compared to the transmitted hash. Some students confused a digital signature with a digital certificate. The quality of written communication of almost all responses was satisfactory.

Question 6

Part 6a: Students were required to write the most negative number that could be represented, which is the negative number that is furthest away from zero. The vast majority of students correctly identified that the exponent would be 01111. There was more confusion about the mantissa, which may have been caused by students failing to understand how negative values should be normalised. A particularly common but incorrect response for the mantissa was 1.000001

Part 6b: Students were required to perform a conversion from floating point binary to denary. This was well tackled, with the only common error being treating the bit to the left of the binary point as being worth 1 in denary rather than -1.

Part 6c: This question part was well attempted. Students were asked to represent a much bigger number than has been the case on previous papers, but this did not appear to cause them any problems.

Part 6d: Students needed to convert a negative denary number into floating point binary. This conversion was more complex than those in parts 6b and 6c and, as a result, whilst nearly three quarters of students achieved two of the three marks, only around 15% of students scored all three. The most commonly made error was to identify the exponent incorrectly as being -4 instead of 4. Working mistakes were also made when converting a fixed point representation of +12.5 to a two's complement representation of -12.5.

Part 6e: Students were required to identify standard types of error that occur when floating point arithmetic is performed. Overflow and underflow were well known, but very few students correctly identified that the last example would produce a cancellation error. Some students lost marks through imprecise naming such as "overload" or "stack overflow".

Question 7

Part 7a: Students were required to draw an inheritance diagram. Most students scored two of the available three marks which were for identifying correctly the class hierarchy. The third mark was for drawing a correctly styled diagram and many students failed to do this. Students who did achieve the third mark correctly enclosed the class names and also drew arrows that pointed upwards to a class' parent class.

Part 7b: In this question part students had to write a class definition for the Computer class. Most students had a reasonable understanding of how to do this, with almost all achieving some marks, but less than a fifth scored full marks. To achieve all four marks students needed: to make clear that the class inherited from the Device class, to redefine the AddDevice procedure, to declare private variables to store the additional properties, and to declare public functions to provide access to the values in these variables. The most commonly made mistakes were to fail to make the inheritance clear and to forget to redefine the AddDevice procedure. Some students lost marks by unnecessarily redeclaring the functions or variable from the parent class or by giving the functions the same names as the variables.

Part 7c: The purpose of this question was to test if students understood that the Laptop class inherited from the Computer class, rather than the Device class. The vast majority of students who dealt with inheritance in this question part correctly identified this.

Part 7d: Most students were able to identify that Bluetooth is a wireless protocol. Many, but not all of these, then went on to explain that it was designed for use over short distances. Many good examples, such as transferring photos from a mobile phone to a laptop or using a Bluetooth mouse were given. Some students lost marks by giving an example that was not in context.

Question 8

Part 8a: Students were required to explain why a server-based network was more appropriate for a college administration system than a peer-to-peer network. Responses that either explained general advantages of a server-based network or that were written in the context of using a centralised database on a server-based network were both given credit. However, many responses were very weak and amounted to little more than stating that the database could be stored on the server. Good responses considered issues such as security, management, concurrency issues and the creation of backups. Just under half of students achieved any marks for this question part.

Part 8b: For this question part, students needed to explain how a thin-client network worked and the impact of this on the hardware that should be purchased to implement one. Most students were aware of the fundamental concept that in a thin-client network the majority of processing tasks were completed by the server. Disappointingly, far fewer then went on to explain that the clients would function primarily as input/output devices with the network being used to transmit input and output from the clients to the server and back again. The worst tackled aspect of the question was the description of the impact of the method of operation on the choice of hardware. Students often wrote vague statements such as, "a more powerful server would be required," or, "the clients could be quite weak." At this level, students needed to give specific examples of hardware requirements such as that the server might require multiple processors to deal with the workload or that the clients may not need a hard disk drive. The most neglected aspect of the hardware part of the question was the impact on

the network infrastructure of the volume of data that needed to be transmitted. A small but noticeable number of students wrote about the differences between thick and thin-client systems, failing to directly answer the question. Some students also confused thin and thick-client systems with star and bus topologies.
Part 8c: About a quarter of students were able to identify that a gateway was used to link together two networks that used different protocols and would perform protocol conversion. A common mistake was to confuse a gateway with a firewall. Many students just restated what they were told in the question, ie that it would be used so that staff could access the network from outside of the college.

Question 9

Part 9a: Just over half of students correctly identified that the key was a composite key. On this occasion, as the question did not refer directly to the example on the question paper, the answer compound key was also accepted – although it is important to note that these two terms are not equivalent.

Part 9b: This question was not tackled well, but good students recognised that relations in a normalised database would have no repeating groups of attributes, no partial dependencies and no non-key dependencies.

Part 9c: As in previous years, students had difficulty completing the entity-relationship diagram with many incorrectly identifying the degree of the relationships. Only a third of students achieved both marks for this question part.

Part 9d: A pleasing number of students (approximately a quarter) achieved the full five marks on this question part, which required that a query was written in SQL. Common errors made by those who did not score full marks were: to forget to cross-reference the tables using either linking conditions in the WHERE clause or INNER JOIN in the FROM clause, to forget to enclose the string "Lucas Bailey" in quotation marks, and to miss out the AND keywords if they were required in the WHERE clause.

Part 9e: This is the first time that students have been asked to use the INSERT INTO command in SQL. Most students scored at least one mark, for recognising that the first line of the command needed to be INSERT INTO BOOK. More mistakes were made in the second line, which should have been 837023, "Kenyan Safari", "Karen Matu", "African Travel Guides". The most common error was to try to set the values by using assignment commands eg Author = "Karen Matu". Another commonly made mistake was to leave out the quotation marks for the values which were clearly strings. Students were awarded credit regardless of whether the value 837023 was in quotation marks as the data type could have been either a string or numeric type.

Part 9f: Students were required to redesign the database structure so that it would be able to store data about multiple copies of the same book whilst remaining normalised. Good responses recognised that a new relation would need to be created, which would contain the existing BookID and a new primary key Accession ID. The BookID in the Loans relation would then need to be replaced by the AccessionID. Alternative solutions that maintained normalisation were also accepted as were answers in which students rewrote new relations. The most common errors were to create a new primary key and store this in the Book relation, which would have resulted in redundant data (eg author details being repeated for each book), or to produce a solution which stored a quantity in stock field and updated this as books were loaned out. The latter would not have allowed the library to keep track of the individual copies of a book when loans were made.

Part 9gi: This question part asked about hashing, which is used so that searching, adding and deleting records can be performed efficiently. Some students were able to explain the purpose of hashing adequately, but many confused hashing in this context with hashing for security purposes.

Part 9gii: This question part was moderately well answered, but most students did not include enough detail in their responses to achieve both of the available marks. In particular, whilst students recognised that the hash function would calculate a record position, many did not explain that the calculation would be based on the value in the record's key field.

Part 9giii: This question part was better answered than part 9gii. Many students recognised that a collision would occur when two key field values mapped to the same storage location and suggested an appropriate method, eg using the next available location or a linked list, to deal with the problem. A small number confused collisions in this context with collisions on a bus network.

Question 10

Part 10a: Two thirds of students were able to identify one property that a graph must have to be a tree. A small number confused a tree with a rooted tree and made assertions such as that a tree must have a root, which is incorrect.

Part 10b: This question part tested students' understanding of the method being used to represent a maze as a graph. The majority of students correctly identified a feature of the maze that would stop its graph being a tree. The most commonly seen correct response identified that there could be a loop in the maze. Other possibilities included that part of the maze could be inaccessible or that part of the maze might only be traversable in one direction. Some students failed to achieve the mark because they re-answered part 10a, discussing a feature of a graph that would stop it being a tree, rather than a feature of a maze.

Part 10c: Students were asked to represent the graph of the maze as an adjacency matrix. Three quarters of students scored both marks for this question part. Responses where symbols other than 0s and 1s were used in the matrix were accepted, as long as they could be viewed as an accurate representation of the graph.

Part 10di: The vast majority of students were able to identify that a recursive routine would call itself. A small number asserted that a recursive routine would repeat itself, which was not considered to be enough for a mark as this could equally have been a description of iteration.

Part 10dii: Most students scored some marks for this question part, but less than a fifth achieved both. The most widely understood point was that the data would need to be removed from the stack in the reverse of the order that it was put onto it so that the recursion could be unwound. Less well understood was the types of data that would be stored, such as return addresses and local variables.

Part 10diii: Most students achieved some marks on this question part and around a quarter achieved all five for a fully complete trace. The most commonly made mistake was to update, incorrectly, the CompletelyExplored array as the recursive calls were made, as opposed to when the recursion unwound.

Question 11

Part 11a: Most students got at least one mark for this question part, usually for identifying that the first required response was "write a program". Many also went on to achieve the second mark too for identifying that the program being tested would not be run.

Part 11b: This question part was poorly attempted. Students needed either to explain that the Halting problem was an example of a non-computable problem, or to state that there was no solution to the problem, so inspection alone could not always determine if a program would halt on a given set of inputs. Some students stated that the Halting problem could be used to show if a program/algorithm was computable or not. This is not the case, as there is no solution to the Halting problem. Rather, it can be used to demonstrate that some problems are not computable as it is an example of one such problem. A small number of students confused non-computability with intractability.

Question 12

Both question parts were well attempted, with approximately four-fifths of students getting the mark for part 12a and two thirds for part 12b. Almost all students knew how a regular expression should be formed. The most common mistake was to confuse the use of the * and + operators.

Statistical data and information on grade boundary ranges www.aqa.org.uk/over/stat.html

UMS conversion calculator www.aqa.org.uk/umsconversion