Version 0.1 0711

General Certificate of Education (A-level)
June 2011

Computing COMP3

(Specification 2510)

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

Report on the Examination

2011 examination - June series

Further copies of this Report on the Examination are available from: aqa.org.uk

Copyright © 2011 AQA and its licensors. All rights reserved.

Copyright
AQA retains the copyright on all its publications. However, registered centres for AQA are permitted to copy material from this booklet for their own internal use, with the following important exception: AQA cannot give permission to centres to photocopy any material that is acknowledged to a third party even for internal use within the centre.

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 second time that COMP3 had been examined and it was pleasing to see the quality of many responses. Many topic areas were well understood and some candidates achieved very high marks on the paper.

Question 1

Part 1a: The binary search method was well understood and the majority of candidates were able to correctly label the sequence of four items that would be checked to search for the name "Richard". Some candidates missed out on the final mark by not realising that there would be a fourth comparison, i.e. "Richard" compared with "Richard". A minority of candidates applied the linear search method instead, labelling each of the names from Adam down to Richard consecutively from 1 to 11.

Part 1b: This part was not tackled as well as part 1a. Most candidates seemed to realise that the answer involved logarithms or powers of two, but the most common response was seven rather than eight, the correct answer. When calculating the number of comparisons required to search a list of n items, $\log_2 (n)$ should be calculated and the result then rounded up. So $\log_2(137)=7.10$ to 2 decimal places, which rounds up to 8.

Part 1c: This question part was well answered, with over half of the candidates correctly identifying the complexity of the binary search method as $O(\log_2 n)$.

Question 2

Part 2a: This topic was poorly understood. Many candidates did little more than rephrase the terms given on the question paper in their responses. An interactive operating system is one in which the user and computer are in direct two-way communication. A network operating system contains a layer of software that redirects requests to remote resources in a way that is transparent to the user. Stating that a network operating system was used to connect to a network was not enough; various hardware and software components are involved in this. Some candidates confused a network operating system with the use of thin-client computing or assumed that one would only be used on a file server.

Part 2b: The responses to this question part were much better than those to part 2a, with almost half of candidates achieving both available marks. Good candidates identified that an event-driven program would respond to events by calling specific subroutines and were also able to give an example of an event. Some responses also mentioned the use of a system loop or event handlers. The most common mistake was to write an answer that was too vague and could equally be applied to programs that were not event driven, such as, "an event driven program reacts to user inputs or runs certain code based on the inputs made".

Question 3

Part 3a: Most candidates got at least one mark for this question part, usually for correctly identifying that the exponent should be 1000. The mantissa was harder to work out, as the representation used required that this should be normalised, so the correct answer was 0.1000000. Some candidates wrongly stated that the mantissa should be 0.0000001 which would not be valid in a normalised representation.

Part 3b: Most candidates knew how to convert a floating point value into denary, with nearly half achieving both marks. The most common mistake was to forget that the most significant bit in the mantissa had a negative value and so arrived at the answer 5.5 instead of the correct answer of -2.5. Candidates who did this nevertheless usually achieved one mark for their working out. Different methods could be used to arrive at the final answer. More mistakes were made by candidates who used the method of converting the bit pattern from negative to positive by flipping the bits and then adding one. Often, candidates using this method forgot to add the one at the end.

Part 3c: This question part was very well answered with nearly two thirds of candidates getting full marks. A small but surprising number of candidates put a 1 in the most significant bit of the mantissa which would have produced a negative number. This mistake should have been easy to spot if these candidates had checked their answers.

Part 3d: This was the part of question 3 that candidates found hardest. The advantages of normalisation are that it maximises precision within a given number of bits and also that there is a unique representation of each number which makes it simpler to test for equality of numbers. Many candidates gave advantages of floating point over fixed point rather than answering the question that was asked. Some candidates who made responses regarding precision failed to get the associated mark because they did not refer to the important fact that precision is maximised in a fixed number of bits. This is important as a number can always be represented more precisely by using more bits.

Part 3e: Responses to this question part were excellent. Many candidates correctly identified that the new representation would give increased range at the expense of less precision.

Question 4

Part 4a: The overwhelming majority of candidates were able to correctly complete the transition table.

Part 4b: The double circle indicates an accepting state. If the Finite State Automation (FSA) is in an accepting state when it finishes processing the input then the input string is accepted, otherwise it is rejected. Many candidates were confused between an accepting state of an FSA and a halting state of a Turing machine and incorrectly stated that the FSA would halt as soon as it entered the accepting state.

Part 4c: The vast majority of candidates got both marks for this question part, indicating that they were able to follow the processing of a string by an FSA, even though they struggled to answer question part 4b.

Part 4d: This question part was well answered with many candidates correctly responding that the FSA would accept strings which began with a 0 and contained an odd number of 1s. Some responses were too vague such as, “accepts 0s and 1s”. The most common error was to state that, to be accepted, a string would have to begin with 01 when in fact the second digit could be a 0.

Question 5

Part 5a: This question part was very well answered with the majority of candidates getting both marks. The only common mistake was to miss out the brackets in the expression that should be (12+19)*8.

Part 5b: As with part 5a, this question part was also well answered. The most common correct response was that brackets are not required. It would have been nice to see some more detailed explanations of this point, rather than just a brief statement of it. A common incorrect answer was that RPN was easier for a computer to understand. The word “understand” is not appropriate in this context.

Part 5c: Responses to this question part were excellent, with relatively few errors made. The majority of candidates got full marks which is unusual for a question involving a trace table. The only two recurring mistakes were to pop the numbers off the stack in the wrong order, resulting in the transposition of the values in Op1 and Op2 and forgetting to push 50 back onto the stack at the very end.

Part 5d: This question was well answered, with most candidates getting some marks and a significant number more than half marks. The most common mistake was to increment the TopOfStackPointer in the wrong place – either before the `If` construct which tested for the stack full condition or after the value in `ANumber` was stored into the `StackArray`. Some candidates implemented solutions that used a loop to find the first empty position in the array to insert the number into. These were awarded credit if they would have worked, but many failed to test properly for the stack being full. It is important that candidates use the correct variable names when they are given on the question paper.

Question 6

Part 6a: This question part was poorly answered with many candidates giving vague responses or explaining what a simulation is rather than a model. In this context, a model is an abstraction of the real-world problem that leaves out unnecessary details. Some candidates confused a model with a prototype.

Part 6bi: Again, this question part was poorly answered. A significant number of candidates appeared to have no understanding of what was being asked, although more than half got at least one mark. Candidates who made a reasonable attempt at an answer often named two pointers, but then offered inadequate explanations of their purpose. For example, the purpose of the pointer to the end of the list is to enable new items to be added to the list, not simply to know where the end is.

Part 6bii: Some candidates correctly identified that a priority queue was required, but many invented new types of queues.

Part 6c: This question part was well answered with many candidates giving well thought out answers such as determining whether the next person entering the cafeteria was a student or teacher or generating a time taken to serve the person at the front of the queue. The most common incorrect answer was the number of people/students/teachers in a queue. In each case, the number in a queue would be a consequence of other randomly determined occurrences rather than determined randomly itself.

Question 7

Part 7a: Most candidates got at least one of the two marks for this question part. Good responses explained both the theoretical aspects of the problem with the order relation and also the real-world consequences of this. Some candidates mistakenly stated that the database contained partial key or non-key dependencies.

Part 7b: Most candidates got some marks for this question, but full mark responses were quite rare. A small number of candidates ignored the instruction to draw three relationships and thus limited the maximum mark they were allowed to two out of three. The most common error was to show an incorrect degree for the relationship between the Product and Orderline relations.

Part 7c: This question part was very well answered with most candidates getting some marks and many getting all three. The most common mistakes were to define ProductNumber as the primary key, but then to forget to give it a data type, and to declare Price to be an integer data type. When marking this question we allowed the use of data types that were taken from other languages rather than SQL, so long as they were clearly equivalent. In the future it is likely that we will require the use of correct SQL data types and syntax for full marks to be awarded to a response.

Part 7d: A server-side script is a sequence of instructions that is executed on a web server to generate web pages dynamically at the time that a request to view a page is received. The majority of candidates showed a reasonable understanding of this and got at least one of the two available marks. Instead of explaining what a server-side script is, some candidates gave examples of usage or described them in such a way that they might have been templates rather than executable programs. A common statement that was insufficient to be creditworthy was that the scripts were stored on a server or accessed from the server. This was not adequate as the same could be said of static HTML pages.

Part 7ei: This question part was well answered. To achieve both marks, candidates had to make clear that the values were being retrieved from the web server or from the input made in the web browser.

Part 7eii: This question part was well answered with many candidates getting both available marks.

Part 7eiii: Only a minority of candidates achieved the mark for this question part. Most recognised that a calculation would be performed and the result of this would be output, but to achieve the mark a candidate needed to explain that these results would be displayed on the web page or sent back to the client computer and displayed in the web browser.

Part 7f: This question part was well answered. Most candidates clearly understood the structure of an SQL query and many scored high marks. The most common mistakes were to include spurious punctuation such as semicolons or commas in responses and to include the Order relation in the FROM clause, which was not required. The correct command to sort the results was ORDER BY ProductNumber, or alternatively, ORDER BY ProductNumber ASC. Some candidates put brackets around the ASC which is not correct syntax.

Question 8

Part 8ai, ii: Approximately two-thirds of candidates responded with an appropriate IP address for each of these question parts. However, as in 2010, a small but significant minority of candidates gave answers that could not possibly be IP addresses.

Part 8b: The correct subnet mask was 255.255.255.0. Approximately half of the candidates identified that this was the case.

Part 8c: Answers to this question part covered a range of issues: security, reliability and throughput. The most common correct response explained that there would be a reduction in the number of collisions. Candidates needed to ensure that their explanations were sufficiently detailed to achieve both marks.

Part 8di: This topic was reasonably well understood, with candidates explaining that cabling costs would be lower as a single cable would run around the entire network. A small but surprising number of candidates believed that a bus network would be faster because only one cable was involved.

Part 8dii: There were many good responses, covering issues relating to reliability, security and speed. The most common error was to state that a star topology was more reliable

because the failure of a single computer would not affect the others. Rather, it is the limited effect of the failure of a single cable that would improve reliability.
Part 8e: Most candidates managed to write lengthy responses to this question part. The quality of these answers was quite variable, with some candidates demonstrating an extensive understanding of network security when connected to the Internet whilst other responses were quite superficial. When answering this type of question candidates need to make sure that they address the entire question. Some candidates lost marks by focussing only on the security measures that the network manager could put in place, either ignoring or only briefly mentioning the threats to which these would be responses. To achieve marks, candidates needed to describe the threats and measures, not simply name them.

Question 9

For question parts 9ai to 9aiii answers by example were accepted so long as they included a sufficient number of examples to show the pattern clearly and (when appropriate) there was some indication that the pattern continued beyond the examples listed.

Part 9ai: This regular expression would match strings consisting of one or more “a” characters followed by a single “b”. In responses to this type of question, candidates should ensure that their answers make the order of characters explicit. “Strings containing one or more a*s* and a b,” does not make clear that the “a” characters must be before the “b”.

Part 9aii: This regular expression would match the strings “ab” and “b”. Some candidates mistook the ? symbol to be a wildcard that would match any single character between an “a” and a “b”.

Part 9aiii: This regular expression would match a sequence of zero or more occurrences of “ab”. The response, “any number of as and bs,” was not sufficient to be awarded a mark as it did not make clear that the “a” and “b” characters must come in pairs. Candidates who answered by example also have to make clear that the empty string would be accepted.

Part 9bi: This was the best answered part of question 9, with most candidates correctly writing a regular expression that would work.

Part 9bii: Most candidates achieved one mark for this question part – for the correct start and end of the string. Writing a correct expression for the middle part of the string was more difficult. The correct regular expression was 10(0|1)*01. Candidates who wrote an expression that would match the same set of strings, such as 10(0?1?)*01 or 10(0|1|01|10)*01 also gained full credit.

Question 10

Part 10a: The baud rate is the number of signal changes per second that can be supported by a transmission medium. Approximately half of the candidates stated this correctly, but some defined the bit rate instead.

Part 10bi: The bit rate can be higher than the baud rate if more than two different signal levels are supported so that more than one bit can be encoded in each signal change. This was recognised by some, but not many, candidates. A commonly held misconception was that the bit rate could be higher than the baud rate if the data being transmitted contained consecutive bits of the same value - for example, transmitting three 1s followed by a single 0 would only need one signal change.

Part 10bii: This question part was well answered, with the vast majority of candidates getting two of the three marks. Most candidates correctly explained the first two missing steps but got the last one wrong. The final stage in the handshake should have been that the printer indicates it is again ready to receive. Many candidates mistakenly believed that the printer indicated that the data had been received or was being printed instead.

Part 10c: Candidates showed only a limited understanding of baseband and broadband, with just over a third scoring and marks. Many clearly had little idea what the terms meant and were just guessing based on their general knowledge of "broadband" and knowing what the terms LAN and WAN mean. Those that did achieve some marks often made quite superficial points, usually about the number of communication channels, rather than demonstrating a sound technical understanding of how the two systems differed in their method of operation.

Question 11

Part 11a: This question part was very well answered. The majority of candidates knew how to trace the execution of a Turing machine and many got full marks. The two most common mistakes were to change into state S_2 on the fourth transition, i.e. when the head moved right into the first blank cell, and to start to delete 1s on the fifth transition, i.e. when the head moved left for the first time.

Part 11b: The Turing machine deleted the rightmost two 1s from the end of the string on the tape. This was recognised by a third of the candidates. Some however made assertions that were too vague to be creditworthy, such as, "deletes two ones from the tape," or, "erases the string." An alternative valid answer was that the Turing machine subtracted two from a unary number. A small number of candidates referred to the end of the tape. Such responses were rejected as the tape is infinitely long.

Part 11c: A Universal Turing machine (UTM) is a Turing machine that can simulate the behaviour of any other Turing machine. A description of the Turing machine that is being simulated, including the instructions that the machine follows, is written onto the tape of the UTM. The UTM then acts as an interpreter, faithfully executing the operations on the data exactly as the original Turing machine would have. The two most common mistakes that candidates made were to describe an ordinary Turing machine rather than a Universal Turing machine and to state that a UTM would control another Turing machine rather than simulate it.

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

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