AQA

A-LEVEL
COMPUTER SCIENCE

7517/2 Paper 2
Report on the Examination

7517
November 2020

Version: 1.0

General Introduction to the November Series

This has been an unusual exam series in many ways. Entry patterns have been very different from those normally seen in the summer, and students had a very different experience in preparation for these exams. It is therefore more difficult to make meaningful comparisons between the range of student responses seen in this series and those seen in a normal summer series. The smaller entry also means that there is less evidence available for examiners to comment on.

In this report, senior examiners will summarise the performance of students in this series in a way that is as helpful as possible to teachers preparing future cohorts while taking into account the unusual circumstances and limited evidence available.

Overview

Given the different circumstances in which exams were taken this year, it was pleasing to see the range of responses that students made and the high quality of some of them.

Students performed particularly well on questions six and ten. However, they found the final part of question six, which was a Boolean simplification question, harder than normal. This was probably because it required the use of De Morgan’s law. Question two about floating and fixed point was also less well tackled than similar questions have been in the past; it asked about some parts of the topic for the first time, and students found the conversion in part 2.2 harder than usual.

Comments on Individual Questions

Question 1.1

How to calculate the amount of storage space required for a sound sample was well understood, with three quarters of students correctly calculating the answer and almost everyone achieving at least one mark. Common mistakes were using a time of 40 seconds instead of 1 minute 40 seconds, multiplying by 4 instead of 16 (presumably as a result of confusion with calculations about images) or making an arithmetic error during the calculation.

Question 1.2

The topic of analogue to digital conversion was poorly understood, and even markworthy responses were often poorly expressed. Good responses recognised that the amplitude of the waveform, represented as an electrical signal, was measured at regular intervals and each measurement would be coded as the closest binary value. Many students wrote little more than that an ADC would take an analogue signal and convert it into a digital signal. Students who wrote otherwise good responses often stated that samples were taken, without really explaining what this meant, ie measuring the amplitude of the input waveform/signal, and therefore only achieved one of the two available marks.

Question 2.1

This was the first time on this specification that students have been asked to compare fixed and floating point number systems. The majority of students achieved at least one mark.

A mark was awarded for stating than fixed point is more precise, but students should be aware that this is only true for some numbers and in a given number of bits – marks might not be awarded for this point without the additional clarification in the future.

Some students were confused between representing very small numbers (ie close to zero) and representing numbers with a high degree of precision. In a fixed overall number of bits, it is true that a floating point system can represent much smaller numbers (by using a negative exponent) but it is not true that it can represent numbers more precisely, as some of the bits will be used by the exponent so cannot be used in the mantissa.

Students should also note that whilst a floating point system can represent a bigger range of values that a fixed point system, it cannot represent a bigger number of values in the same number of bits than a fixed point system can; in n bits, it is only possible to represent 2^n different values, whatever those values may be.

Question 2.2

Students found this question, which required the conversion of a floating point number into decimal, harder than similar questions in the past. Only a quarter correctly converted the value from binary to decimal. The most common issue was not knowing how to deal with the -1 in the mantissa when the binary point shifted left. Other common mistakes included treating the -16 in the mantissa as +16 and moving the binary point three places right instead of left.

Question 2.3

The responses to this part, which required the conversion of a decimal number into floating point binary, were much better than to part 2.2 with half of the responses achieving full marks. Many of the students who did not achieve full marks were able to correctly represent 1632 as an unsigned binary integer but then went wrong with positioning the binary point and setting the value of the exponent. The most common error was to set this to -11 instead of +11.

Question 2.4

For this question students had to calculate the range of a floating point system. Most students were able to calculate one of the highest and lowest values, but only a small minority correctly calculated both. The most common errors were giving the closest value to zero instead of the most negative value as the lowest value and calculating the value of the mantissa to be 63 and -64 for the highest and lowest values respectively.

It is important that students show their working as some wrong answers might get working marks; students who had incorrect final answers often showed correct values in binary and so achieved some working marks.

Question 2.5

The majority of students achieved some marks for explaining that the number would be rounded in some way as it could not be represented exactly. Responses which suggested redesigning the system used to represent the number were not awarded marks.

Question 3.1

The question was well answered with two thirds of students achieving full marks. With this type of question, it is usually helpful to identify the processor first, if this is possible. Common errors were to mix up the address bus and data bus or processor and main memory.

Question 3.2

This question was not well answered, with only a third of students correctly calculating that 4 gibibytes was equivalent to 4194304 kibibytes. The most common errors were forgetting that there was another unit (mebibytes) between kibibytes and gibibytes and so only multiplying by 1024 once, converting 1 gibibyte to kibibytes and converting 4 gigabytes to kilobytes.

This question was about the von Neumann and Harvard architectures. It was not well answered, with only a third of students achieving any marks. The most common correct answers related to being able to access data and instructions simultaneously and the security that is achieved by it not being possible to execute values in the data memory as program code. The statement that data and instructions are stored in separate memories was not enough for a mark; this was a description of the architecture rather than an advantage of it.

Just over half of the students were able to name all three of the registers used in the Fetch-Execute cycle. The most common mistakes were minor errors in naming the registers (eg “Memory Access Register”, “Control Instruction Register”) and giving the wrong name for register three, which should have been the Current Instruction Register.

Interrupts have not been asked about before, and approximately half of the students achieved some marks for describing their role. Good responses recognised that the role of an interrupt was to cause the processor to suspend the execution of its current task so that it could deal with a task that needed the immediate attention of the processor. Some answers were too vague and did not make the distinction between interrupts and scheduling clear enough to be markworthy. Another common mistake was to state that an interrupt would stop the Fetch-Execute cycle or even the processor.

Question 3.6

The reason for saving the volatile environment was poorly understood, although many responses were just not expressed clearly enough or in enough detail to achieve a mark as opposed to being incorrect. Good responses recognised that the volatile environment was saved because the process executed during the interrupt might overwrite the register values and so it would not be possible to resume execution of the suspended process later. Some students saw the word "volatile" and wrote about the contents of volatile memory being lost when a computer was turned off.

Question 3.7

Good responses recognised that the hardware is the electrical or physical components of the computer system and that software are the programs that execute on the hardware to carry out tasks. The majority of students expressed this in some way,

Question 4.1

The majority of students correctly identified that the REST API is not created and run on a client computer. Students who responded incorrectly chose evenly between options B and C, with just a handful choosing A.

Question 4.2

The vast majority of students correctly identified that the representation used XML rather than JSON.

Question 4.3

The vast majority of students were able to describe at least one advantage of JSON over XML. The most common responses related to the conciseness of the code, the simplicity of parsing it and the fact that humans could more easily understand it. A common error was to refer to the execution of JSON code, which is not executed.

Question 4.4

A quarter of students recognised that the choice of key meant that a buyer could only visit a particular property once per day. Many students who gave incorrect responses recognised that the choice of key would impose a limitation, but gave incorrect limitations such as “A buyer will only choose to look at one property” or “Each house will only be viewed by one person”.

Question 4.5

Two thirds of students achieved some marks for writing the SQL query, but fewer than 10% wrote a fully correct query. Many different mistakes were made, the most common of which was missing the Buyer table out of the FROM clause. This might have been because it only needed to be used in the conditions and none of the query's output fields were taken from it.

Question 5.1

Just under half of students could give an advantage of using DHCP. The most common correct responses related to automatic configuration of a host without the need for expert knowledge and the ability to easily reuse a limited pool of addresses.

Question 5.2

Students struggled to explain how two computers connected to the Internet could have the same IP address and communicate with each other. Good responses recognised that the IP addresses would be private / non-routable and that NAT would be used to convert locally unique private IP addressed to the public IP address of a router as the data passed out of a LAN onto the public Internet. Some students incorrectly believed that this was achieved using MAC addresses or subnet masks.

Question 5.3

Most students were able to achieve at least one mark for their descriptions of how a firewall protected computers but responses often lacked depth and did not achieve many more marks. For example, a response might refer to blocking certain devices or stopping some applications accessing the Internet without referencing how this was achieved using IP addresses or port numbers. A common misconception was that a firewall would scan for viruses. As a firewall typically examines network traffic at a packet level, it is unlikely that the firewall would be able to determine if it contained a virus as it would not have enough data to look at. A further misconception was that stateful inspection simply meant looking at the data in a packet; stateful inspection refers to the analysis of packets in the context of an overall communication, which might for example include rejecting packets that are not replies to requests originated from within a network. It is important to note that firewalls typically only examine a packet's header and do not inspect the data in the payload.

Question 5.4

This was the highest tariff question on the paper, and half of students achieved at least five marks. The question covered three areas and the use of a checksum to determine if the packet contents had been changed was by far the best understood area. Students often wrote about the use of the subnet mask to determine whether a destination could be reached directly on the LAN, but how this was achieved was poorly described; there were references to the use of the AND operation but it often wasn't clear what was ANDed with what or what was done with the result.

The least well understood area was that of routing across the Internet. Good responses recognised that packets would be passed from one router to another, with the route being determined dynamically by each router using routing tables based on the destination IP address. A frequently seen misconception was that the route would be determined at the start of the transmission using Dijkstra's algorithm. The router in the LAN would not have enough information available to it for this type of source routing to take place.

If there is time, it is advisable for a student to read through their answer at the end and check that what they have written makes sense. Marks are often lost because points are made that are unclear or lacking in detail and it is possible that a student could correct this if they read their own answer back.

Question 6.1

The vast majority of candidates correctly identified that table C was not the truth table for one of the listed gates.

Question 6.2

Nearly three quarters of students were able to complete the truth table for the logic circuit correctly. This type of question is usually well tackled and this year was no different.

Question 6.3

This question was well answered, with almost everyone who correctly completed the truth table in question part 6.2 being able to identify which segment of the display the circuit controlled.

Question 6.4

This question part required students to simplify a Boolean expression. Students typically find it harder to use De Morgan's Laws than to use identities for simplifications which was reflected in the quality of the responses. The most common mistake was to apply De Morgan's laws and then not insert brackets into the resulting expression. The effect of this was to produce an incorrect expression because of the order of precedence of the operators.

Students should ensure that they set out their working clearly and also that they only apply one simplification at each stage; this will ensure that if their final answer is incorrect, as was the case for the majority of students this year, they are able to achieve as many working marks as possible. Some students applied multiple simplifications in one step, which was fine if they did it correctly, but if they didn't then they may not have achieved some working marks that they could have done if they had shown the steps individually.

Question 7.1

Just over three quarters of students correctly identified that it was -43 that was rational and real but not irrational and natural.

Question 7.2

Approximately half of students correctly identified that a real number was the most appropriate to use to measure the length of a piece of rope. The most frequently selected incorrect option was rational.

Question 8.1

Most students achieved at least one mark for this question part. The most commonly identified relevant reasons for using RFID tags were that the tags could be scanned more quickly or without having direct sight of them. Fewer students went on to achieve a second mark, for explaining that this was relevant to the scenario as the warehouse would have many products to scan at once when a delivery arrived.

Question 8.2

The most common responses related to the barcodes being cheaper to produce. Many students did not relate their answer directly to the question and so did not achieve the second mark. The best responses did this and recognised, for example, that the lower cost of producing the barcodes would mean the supermarket could sell the product more cheaply or the fact that barcodes were not susceptible to radio interference would mean a lower risk of a nearby product being accidentally scanned.

Question 8.3

For this question part, students were required to explain how RFID worked and how SQL might be used to update the database. A good range of responses were seen but some candidates addressed one aspect of the question in much more detail than the other. The operation of RFID was fairly well understood, but students should be careful to make clear that the transmission of data is achieved using radio waves.

Good responses to the SQL part of the question explained that a SELECT query could be used to identify if the produce type already existed in the database, and if it did an UPDATE query could be used to update the stock level for the product. Otherwise, an INSERT query could be used to create a new record for the product. The SELECT query would not return an error if the product were not in the database, it would simply return an empty relation.

The majority of students correctly worked out that with 10 bits available 1024 memory locations could be addressed. The most common incorrect response was 1023, with students calculating the highest number that could be represented in 10 bits, forgetting that 0 is also a valid address.

Fewer than half of students were able to explain that immediate addressing means that the operand will be used as the datum in the operation. Common incorrect responses confused direct and immediate addressing modes or suggested that immediate addressing meant the instruction would have to be carried out quickly.

Most students achieved some marks for writing the assembly language code, but only a third produced a fully correct solution. Students who made mistakes often did so with the syntax of the language rather than the logic of the solution. There was some flexibility in relation to this, but students who made one mistake often made others as well, for example by using both an invalid register name and incorrect syntax in a memory address like this: LDR Rd, <102>. Whilst the list of commands is published in every exam paper, it is important that students have used these ahead of the exam and so do not make basic syntactical errors.

Most students recognised that the problem with the pseudo-code for the Caesar Cipher was that it would move some letters outside the range of capital letters in the ASCII code. The responses were generally good, with students losing marks for vague responses rather than incorrect ones, for example talking about a character being “outside the range” without being clear what this meant. Students also sometimes wrote about subtracting the difference from 65 off 90 for some ASCII codes without being clear about whether they meant the original ASCII code or the code after a subtraction of 5 had already taken place. Further, in most cases, the subtraction should have been made from 91 rather than from 90. Students could have spotted this error if they had tried to apply the method they were proposing to an example character.

Question 10

This question was about how a social media company could decide which news stories to display to a member and the moral, ethical and legal consequences of doing this. The question was well answered, with a good range of responses being seen. Students should avoid focussing too much on one specific issue, for example how the Data Protection Act might be relevant to this, as to achieve a high mark they need to cover a range of issues.

Question 11.1

Approximately half of the responses to this question part were correct. Almost everyone who failed to get the mark gave the head correctly as “Blackpool” but just gave the tail as “Toronto”, failing to recognise that the tail of the list is all of the items except the head, not just the last item in the list. The head is a single item, whereas the tail is a list (even if it happens to be a single item or empty list) and so should be expressed with square brackets around it as ["Paris", "New Brighton", "Toronto"]. On this occasion, as the question was only worth one mark the mark was still awarded even if the brackets were used incorrectly, but this may not be the case in the future.

Question 11.3

Good responses to this question part recognised that the function was recursive and that it would split the list into its head and tail, passing the tail on to the next call. At each call, the value of the head would be added to the total value of the items in the tail, with the recursion terminating when the list was empty, and a value of 0 being returned for the total of the empty list.

The most commonly achieved mark points were that the function was recursive and that the recursion terminated when the list was empty. Some students failed to achieve marks because they did not make clear that the head was added mathematically to the total of the tail, instead referring to head being added to the tail. Another common error, seen in some otherwise very good responses, was to believe that the recursion terminated when the list contained only one item, when it actually terminated when the list was empty.

Question 11.3

This question part required students to explain what a higher-order function was, and it was not well answered. Good responses recognised that a higher-order function would either take a function as an argument and/or return a function as the result. Students who achieved one of the two marks usually recognised that a higher order function could take another function as an argument but not that that one could return another function as a result – this could be because the specification includes specific examples of higher-order functions that take a function as an argument, such as map.

The most common incorrect explanations were that a higher-order function calls another function, is a built-in function, or is a function that is more complex or more important than other functions.

Question 11.4

Approximately half of the students correctly identified that the result of the fold operation would be the number 12. Common mistakes were giving the list [2, 3, 2] which was the result of multiplying each element by one, applying the map function instead of fold, and the number 7 which is the sum of all of the digits in the list multiplied by one.

Mark Ranges and Award of Grades

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