

BCS THE CHARTERED INSTITUTE FOR IT
BCS HIGHER EDUCATION QUALIFICATIONS
BCS Level 4 Certificate in IT
COMPUTER & NETWORK TECHNOLOGY

Examiner Report

xxxx March 2018 - yyy

General comments on candidates' performance and the process

Overall level of English is pretty poor what sometimes makes it very difficult to judge the knowledge.

Because of some reasons 1 candidate provided answers to Section B of the exam in the book for the Section A.

Also, it has been discovered that some papers in Box #1 were misplaced (i.e. did not followed the sequence as shown in the spreadsheet.

After careful checking and moving some of the papers which were out of the order there were still 2 inconsistencies:

1. There are two papers missing, namely B136-990664372 and B170-990666620 (i.e. they are listed in the spreadsheet but not present in the boxes). I've put the notes in the spreadsheet instead of mark for A1. I've also put a "placeholders" with paper numbers to the place in the pile according to the spreadsheet.
2. There are two papers which are not in the spreadsheet - I've added their marks at the bottom of the list in the "Marks" sheet - B23-990422069 and B23-990668543. I've put these two papers on the top of the box №1.

Thus, the logistics of exam papers handling at the BCS HQ shall be improved...

The most popular question was A1 (73.74%), followed by A2 (50.11%), A4 (37.86%) and A3 (30.20%) – thus, the most popular combination was A1+A2. Please note, that this statistics does not reflect results of the 2 non-listed papers.

Answer pointers

A1 Part (a) 8 marks

A logic circuit has four binary inputs D, C, B, A that represent the sixteen values 0 (0,0,0,0) to 15 (1,1,1,1). The output F is true if the input on D, C, B, A is divisible by 3, 8, 13, or 14. Note that 0 is not divisible by any number. Draw a truth table for this system with inputs D, C, B, A and output F.

Part (a) Truth table

D	C	B	A	Number	Div 3	Div 8	Div 13	Div 14	F
0	0	0	0	0	0	0	0	0	0
0	0	0	1	1	0	0	0	0	0
0	0	1	0	2	0	0	0	0	0
0	0	1	1	3	1	0	0	0	1
0	1	0	0	4	0	0	0	0	0
0	1	0	1	5	0	0	0	0	0
0	1	1	0	6	1	0	0	0	1
0	1	1	1	7	0	0	0	0	0
1	0	0	0	8	0	1	0	0	1
1	0	0	1	9	1	0	0	0	1
1	0	1	0	10	0	0	0	0	0
1	0	1	1	11	0	0	0	0	0
1	1	0	0	12	1	0	0	0	1
1	1	0	1	13	0	0	1	0	1
1	1	1	0	14	0	0	0	1	1
1	1	1	1	15	1	0	0	0	1

To construct the truth table, we put 1 in the column of each of the four specified conditions. A 1 is put in the output column F if and of the four conditions is true.

A1 Part (b) 8 marks

From the truth table, write down an expression for the output F (unsimplified).

F is given by the sum (logical OR) of all condition that set F to 1. That is

$$F = D!C!B!A + D!CBA! + DC!B!A! + DC!B!A + DCB!A! + DCB!A + DCBA! + DCBA$$

A1 Part (c) 8 marks

Simplify this expression

$$F = D!C!B!A + D!CBA! + DC!B!A! + DC!B!A + DC(B!A! + B!A + BA! + BA)$$

$$F = D!C!B!A + D!CBA! + DC!B!(A + A!) + DC \text{ (using } B!A! + B!A + BA! + BA = 1)$$

$$F = D!C!B!A + D!CBA! + DC!B! + DC$$


$$F = D!C!B!A + D!CBA! + D(C!B! + C) = D!C!B!A + D!CBA! + D(B! + C) = D!C!B!A + D!CBA! + DB! + DC$$

$$F = D!C!B!A + D!CBA! + DC + DB! = D!C!B!A + C(D!B!A! + D) + DB! = D!C!B!A + C(BA! + D) + DB!$$

$$F = D!C!B!A + CBA! + DC + DB!$$

The Karnaugh map provides a graphical form of simplification. Locate the 1s on the map, form the smallest number of large groups to get the map below. These groups can be read from the map as

$$F = DB! + DC + CBA! + D!C!B!A$$



A1 Part (d) 6 marks

Design a circuit using NAND logic only to implement the simplified expression for F

$$F = D\bar{B} + DC + C\bar{B}\bar{A} + D\bar{C}BA$$

We need to convert all OR operators to AND operators.


We do this by using DeMorgan's theorem and noting that complementing (negating) twice does not change a value; that is, $F = \overline{\overline{F}}$

$$\text{In this case } D\bar{B} + DC + C\bar{B}\bar{A} + D\bar{C}BA = \overline{\overline{D\bar{B} + DC + C\bar{B}\bar{A} + D\bar{C}BA}}$$

Applying deMorgan's theorem, we get

$$F = \overline{\overline{D\bar{B}} \cdot \overline{DC} \cdot \overline{C\bar{B}\bar{A}} \cdot \overline{D\bar{C}BA}}$$

We can now draw a circuit. Note that we can negate a single term by using $\overline{\overline{F}} = F$



Examiner's comments

As it has been noted this question was the most popular with 73.74% of students who attempted it as well as with 81.01% who passed it.

The structure of the A1 is built in such a way that A1b depends on the answer to A1a, A1c on A1b and A1d on A1c. Thus, it is causing a problem in the case when A1a is answered either partially or mostly incorrectly.

I've marked A1a "proportionally", i.e. giving portion of $X/16$ out of max 8 marks allocated for the A1a. The same approach it was possible to apply to the A1b. Some students have not attempted A1a properly, e.g. either just produced a table from 0000 to 1111 without producing output F or have done some portion of the answer or produced table not in the regular consistent sequence (for what I have not punished the candidate if the answers were correct).

However, A1c was very difficult to mark because using of Karnaugh maps was not explicitly required and the provided sample solution was also showing simplification of the formulae. Thus, I've decided to split marks for A1b into two parts – upto 4 for the Karnaugh map (or 0 if not present) and another 4 for simplification. Not all students who have done A1a and A1b have attempted A1c.

Finally, A1d was attempted by very few students and even less of them have answered it correctly or close enough to the model answer. It seems that the largest confusion was caused by the requirement to use only NAND logic as some students involved NOT components too.

Average mark for the A1 is 18.

Overall, questions such as A1 are much more suitable for the exam than open-ended questions. However, having interdependent sub-questions, e.g. when output from A1a is input to A1b and output from A1b is input to A1c, etc is causing "chain reaction" in case of not fully correct answer in the previous steps. Marking is especially difficult without clear instructions regarding marking.

A2 Part (a) 7 marks

The performance of microprocessors has been improving exponentially since they were first manufactured. However, over the last few years, the advance of microprocessor performance has slowed. Why is this; what is limiting the increase in microprocessor performance?

There are several reasons why the performance of microprocessors has not increased in accordance with Moore's law over the last few years. The main reason is the ability to increase clock speed.

The speed of a processor is dependent on the rate at which it is clocked. In the 1980s clock rates were in the low megahertz region. Eventually, the 1 GHz clock rate barrier was broken. Today, typical desk-top PC have processors operating in the range of 3 to 4 GHz.

A processor's energy consumption increases with clock speed. Consequently, the processor gets hotter at higher speeds. Since the size of the chip is only a few millimetres square, the heat (which can be over 100 W) can only be dissipated by large heat sinks. The practical limit to clock speed is imposed by the need to cool chips.

Another limitation on processor speed is memory access times. The rate at which processors have increased speed has not been matched by a corresponding increase in memory access times. Therefore, memory access is proving to be a bottleneck and limits the performance of a computer system.

Examiner's comments

Just a few candidates followed this question close to the point. Very many are discussing that "microprocessors are getting slower" or just wrote everything they knew about development of computers starting even from vacuum tube based machines. Many students stated that "complex applications are causing slowing down of the microprocessors".

Thus, it is pretty open-ended question which should be avoided in the exam in the future. Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark.

Part (b) 10 marks

How are computer manufacturers attempting to increase the performance of microprocessors?

There are several ways of increasing the performance of a processor. Some of these are:

Pipelining: Instead of completing an instruction before executing the next one, modern processors execute several instructions in a pipeline (like a car assembly line). One instruction is being fetched, while the previous instruction is being decoded, and the one before that is being executed. A computer with N-stages in the pipeline can provide an N-fold increase in speed with no change in clock rate. However, pipelining cannot provide its full capability because of branches that force the pipeline to be flushed.

Out-of-order execution: The von Neuman machine executes instructions sequentially in order. However, if one instruction is waiting for the result of a computation that is still in the pipeline, instructions further in the program can be executed if they don't require data that is not yet available.

Multicore: The multicore processor has 2,3,4 or more CPUs on the same chip. Each CPU can process data in parallel. This allows more computations per second by executing them in parallel (but is limited by the ability to divide a program into parallel operations).

Examiner's comments

Just a few candidates followed this question close to the point. Very many partially repeated what they said in A2a. Pretty many have indicated the introduction of cache memory or increase in its size as a way to increase performance (which was not expected in the model answer but actually shall be considered as a correct one).

Thus, it is pretty open-ended question which should be avoided in the exam in the future. Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark.

Part (c) 7 marks

Cache memory can increase the performance of a microprocessor system. What is a cache memory and, briefly, explain how it improves performance?

Main memory (usually called DRAM) contains programs and data being executed. Unfortunately, its access time is far too low to be used effectively (typically 50ns access time with a CPU clock of less than 1 ns). Cache memory have a very low access time (e.g., 5ns) and is used to overcome the limitations of DRAM. In general, a program tends to executed the same instructions and data frequently (e.g., 5% of a program may be executed for 90 of the time the program is being executed). Cache memory holds only a fraction of the program and data for the DRAM, but for (typically) 95% of the time data is taken from cache rather than DRAM. The secret of cache is in ensuring that the most frequently used data is stored in cache.

Examiner's comments

Not all candidates who attempted A2 have tried A2c and just a few followed this question close to the point. Very many partially repeated what they said in A2b. Very few attempted to draw a diagram showing cache memory, CPU and RAM (as it was not specifically required). Pretty many have indicated use of L1, L2 and L3 cache memories while it was not explicitly required and was not shown in the model answer.

Also, some have referred to the use of cache in the web-browsers as well as in the hard drive controllers (which are valid points ideologically while missing the point of the CPU cache)

Thus, it is pretty open-ended question which should be avoided in the exam in the future. Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark.

It would be better to formulate this question more abstractly about "cache memory idea and its applications in various parts of the computer systems" (and maybe even indicating these parts) which will allow to clearly structurize marking of the answers (e.g. 2+2+2+2 marks in total for the describing the general idea as well as providing example of cache in CPU, cache in browsers and cache in HD controllers – maybe with less of the details which are expected).

Part (d)

The main memory of a computer has an access time of 50 ns. The cache memory has an access time of 5 ns. The hit ratio for the memory system is 0.9 (i.e., 90% of memory accesses are to the cache). What is the speedup ratio of this system?

The hit ratio is 90%. Therefore, for 0.9 of accesses are to cache and 0.1 of accesses are to DRAM.

The average access times is cache access + DRAM access = $50\text{ns} \times 0.1 + 5\text{ns} \times 0.9 = 5 + 4.5 = 9.5\text{ns}$

The average access time without cache is $1.0 \times 50\text{ns} = 50\text{ns}$.

The speedup ratio is access time without cache divide by access time with cache = $50\text{ns}/9.5\text{ns} = 5.26$

Examiner's comments

Very few candidates attempted this question and there were NO CORRECT ANSWERS AT ALL. Thus, question is if such exercises are considered at all in the BCS Learning Centres.

COMMENTS TO THE A2

While this was the 2nd most popular Question, the candidate's performance is much worse than in case of A1 - average mark for the A2 is 4.3 and pass rate is 3.02%. This is the result of not trying A2c and A2d as well as poor answers to A2a and A2b.

A3 Part (a) 5 marks

A computer has a program counter. What is a program counter and what is its function?

The program counter, PC, or instruction counter contains the address of the next instruction to be executed. It is used in a von Neuman computer to step through a program. The instruction at the address in the PC is fetched from memory and executed.

Examiner's comments

Pretty many candidates attempted this question however only few were answering to the point. Thus, question is if such topics are considered properly (or at all) in the BCS Learning Centres.

A3 Part (b) 5 marks

Under what circumstances is the program counter modified by the program?

The program counter is incremented after each instruction is fetched from memory. In this way the program counter steps through the instructions sequentially. However, the program counter can be loaded from memory or a register to execute a jump (or branch) operation.

The program counter can be loaded with a new address if a condition is met. For example, BEQ XYZ may mean "branch on condition 0 to address XYZ. This allows the implementation of loops and operations like IF...THEN...ELSE

Examiner's comments

Pretty many candidates attempted this question however only few were answering to the point. Thus, question is if such topics are considered properly (or at all) in the BCS Learning Centres.

A3 Part (c) 10 marks

In the context of computer architecture, what is a stack pointer register? Describe how it is used to implement subroutines in a computer.

The stack pointer register, SP, contains typically the address of the top of stack. A stack is a last-in-first-out, LIFO, queue because items are retrieved from a stack in the reverse order to which they were entered.

In a typical microprocessor stack, an item X is pushed on the stack by
 $[M[SP]] \leftarrow X$ Push the element pointed on to the stack (in main memory) at the address in the stack pointer
 $[SP] \leftarrow [SP] - 1$ Decrement the stack pointer to point to the next free location above the stack

An item is pulled off the stack by
 $[SP] \leftarrow [SP] + 1$ Increment the stack pointer to point to the item at the top of the stack
Destination $\leftarrow [M[SP]]$ Copy the element at the top of the stack pointer at by SP to its destination

A subroutine is a piece of code that is called, executed and a return is made to the instruction after the calling point. This means that the return address (next instruction after the call) must be preserved.

A typical subroutine call is BSR XYZ which pushed the return address onto the stack and then jumps to location XYZ to execute the subroutine.

At the end of the subroutine an instruction like RTS (return from subroutine) pulls the return address off the stack and puts it in the program counter to execute the instruction after the return. This mechanism allows subroutines to be nested (one subroutine to call another).

Examiner's comments

Pretty many candidates attempted this question (while less than A3a and A3b), however only few were answering to the point and with sufficient level of details (e.g. diagram with stack, etc). Thus, question is if such topics are considered properly (or at all) in the BCS Learning Centres.

A3 Part (d) 10 marks

What is the function of an index register (also called pointer register or an address register) in a CPU. Give a simple example of the use of an index register.

An index register is a pointer register. That is, it contains the address of an operand. If register A0 is an index register then an instruction of the typical form MOVE D0,[A0] means "copy the item pointed at by the contents of index (pointer) register A0 into data register D0".

The index register allows you to use variable addresses because you can change an index register's contents. This allows you to implement data structures such as lists, vectors, tables, and arrays.

A typical example of the use of indexing might be:

```
MOVE A0,Table1  Index register A0 points to Table1 (source)
MOVE A1,Table2  Index register A1 points to Table2 (destination)
Loop MOVE D0,[A0]  Repeat: read item pointed at by A0
  ADD A0,#1      increment A0 to point to next item
  MOVE [A1],D0   copy item to location pointed at by A1
  ADD A1,#1      increment A1 to point to next item
UNTIL Done      EndRepeat ... suitable code to close the loop
```

Note ... any suitable example of indexing is acceptable.

Examiner's comments

Very few candidates attempted this question (even less than A3c), however even less were answering to the point and with sufficient level of details (e.g. ASSEMBLER-like code or diagram with references to the parts of memory, etc). Thus, question is if such topics are considered properly (or at all) in the BCS Learning Centres.

COMMENTS TO THE A3

This was the least most popular Question and the candidate's performance is much worse than in case of A1 and A2 - average mark for the A3 is 3.1 and pass rate is 4.96%. This is the result of not trying A3c and A3d as well as poor answers to A3a and A3b.

A4 Part (a) 10 marks

The operation of the simple von Neumann computer with its fetch/execute cycle has been enhanced by several techniques to improve its performance. Write notes on the following three mechanisms that are found in most computers. In each case, state what the technique is, how it works, and why it increases computer performance.
THE INTERRUPT

A computer executes a program instruction by instruction until the program has been completed. However, there are occasions when a computer needs to interact with its environment. Without an interrupt mechanism, the computer would have to periodically test environment variables (e.g., the time, input and output device status, the readiness of disk drives and printers etc).

An interrupt provides a mechanism whereby any device wanting attention (e.g., a printer, keyboard, or a mouse movement) can inform the computer that attention is required on a hardware line connected to the computer called an interrupt request, IRQ.

Then an interrupt request is received, the computer determines whether the level of the request is high enough to be serviced. If it is, the current operation is completed, the return address saved and a jump made to the appropriate interrupt handler (the software that deals with the source of the interrupt; for example, a printer driver). When the interrupt has been serviced, the return address is loaded into the program counter and execution continues.

Some interrupts (often called exceptions) are triggered by software events such as an illegal operation (invalid instruction) or a memory error (page fault). In this case, the operating system usually deals with the cause of the interrupt.

Interrupts make a computer more efficient and responsive. However, they also make its behaviour harder to predict because interrupts are usually asynchronous and can happen at any time.

Examiner's comments

Very few candidates answered this question and even less followed it to the point. Very many are discussing that "microprocessor and their architecture" or just wrote everything they knew about operating systems.

*It is also unclear to me **why interrupts mechanism shall be considered as a way to improve performance**. It is a way to control input-output operations as well as managing of the control of the programmes (i.e. giving them CPU time) but how it relates to the improved performance?*

Thus, it is pretty controversial question which should be avoided (or more clearly rephrased) in the exam in the future. Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark.

A4 Part (b) 10 marks

VIRTUAL MEMORY

In the simple model of a computer, the CPU generates an address and that is used to access data or an instruction in memory.

In reality, the simple model of CPU and memory is not appropriate for several reasons. First, the actual memory and the memory accessed by the computer may be different in size. This was a major problem in the past because many programs were bigger than the available DRAM storage and could not fit into memory.

A second problem is that memory must be segregated and allocated individual users and the operating system. This is necessary for security. Much of the malware found on computers operates by bypassing security mechanisms and allowing the rogue software to access data that should have been hidden/protected.

A third problem is that the programmer should not have to worry about memory addresses and where data goes in memory.

A virtual memory system uses a memory management unit, MMU, that sits between the address bus from the CPU (the virtual address) and the DRAM memory (the physical address). The MMU translates virtual addresses into physical addresses (location of the actual data).

The MMU is able to check the “rights” or “permissions” of every virtual address generated by the CPU. If the address is accessing a region of memory that has been allocated to it, the address is translated into the appropriate physical address. If not, the access is terminated.

Virtual memory allows programs to reside on the hard disk. When data or program is accessed and it not currently in memory, the operating system intervenes and moves it from disk to DRAM and automatically updates its logical-to-physical address mapping tables.

Examiner’s comments

Very few candidates answered this question and even less followed it to the point. Very many are discussing “virtual memory” while they actually are referring to the “cloud computing”.

*It is also unclear to me **why virtual memory mechanism shall be considered as a way to improve performance**. It is a way to run more programs using smaller RAM, run them maybe slower, but nevertheless be able to run but how it relates to the improved performance?*

Thus, it is pretty controversial question which should be avoided (or more clearly rephrased) in the exam in the future. Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark.

A4 Part (c) 10 marks

DIRECT MEMORY ACCESS (DMA)

Memory is normally accessed by the processor ... instruction codes when executing a program and data when processing information. When data is required from the outside world (e.g., data from a disk drive or from the Internet), it must be read into memory.

In a simple computer, the CPU may be programmed to read data and put it in memory (input) or to read memory and transfer it to a peripheral (output). This process is slow and inefficient and suitable only for simple controller applications.

Direct Memory Access, DMA, is a mechanism whereby the computer’s data and address buses are disconnected from the CPU (electronically) and a DMA controller takes them over. The DMA controller is an autonomous device (a little like a CPU in its own right). The DMA controller is able to control the buses to transfer data between a peripheral and memory rapidly and efficiently.

The DMA controller can be arranged to perform I/O data transfers in periods when the CPU is not using the memory’s bus or it can interleave data operations with the CPU (so called cycle stealing).

In a modern computer like the PC, the DMA function is built into the “bridge” controller chips.

Essentially, any card in a PCI slot can request the bus from the computer, take over the bus (i.e., become the bus master) and then perform the DMI.

Examiner’s comments

Very few candidates answered this question (much less than A4a and A4b) and even less followed it to the point. NONE of the candidates referred to the PCI cards and their operation. Thus, question is if such topics are considered properly (or at all) in the BCS Learning Centres.

Overall, answering such an unstructured question must have been very difficult for the candidates and was a lot of pain to mark. Having PCI mentioned in the question explicitly would be much better indication of the expected direction.

COMMENTS TO THE A4

This was the 3rd most popular Question and the candidate’s performance is much worse than in case of A1 and A2 - average mark for the A4 is 4.9 and pass rate is 11.36%. This is the result of not trying A4c as well as poor answers to A4a and A4b.