

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

Examiner Report

Thursday 26th March 2015

General comments on candidates' performance

The standard of attempts has declined during this sitting. As in previous sittings, many candidates did not write sufficiently in-depth answers for 30 marks awarded for section A. Candidates attempted section B questions well. Centres must provide guidance and support to candidates so that they are prepared to write sufficiently in depth answers in section A. The most popular question was 8 and the least popular was 2.

Below are suggested answers pointers to each question. Full marks were given for alternative valid answers.

SECTION A

(Candidates were required to answer **TWO** out of the four questions set)

A1

Answer Pointers

All information inside a computer is represented in the form of bits (binary digits) that have the values 0 or 1. The basic unit of information accessed in a computer (e.g., in a store or load operation) is the word which varies from 8 bits for simple microcomputers to 64 bits (or more) for high-performance computers.

There is no intrinsic (or inherent) meaning in a string of bits. A string of bits means only what the programmer or computer designer intended. All data in a computer has to be interpreted in the way it was intended (e.g., as a number, an instruction, text, video, image, sound...). However, certain standards have been universally agreed upon to represent information – particularly in the field of number and character representation.

A string of bits; for example, 01001010 can be regarded as an **unsigned binary value** with the 1s having the values (weights) $2^0, 2^1, 2^2, 2^3, \dots$. In this case, 01001010 represents $2^7 + 2^3 + 2^1 = 128 + 8 + 2 = 138$. This form of representing numbers is the default mechanism used by many computers.

An unsigned integer with n bits must fall within the range 0 to 2^{n-1} (inclusive). For 8 bits this is 0 to 255. An error will occur if the possible range of numbers is exceeded. For example, in 8 bits $11111111 + 1$ represents $255 + 1$ but the results in 8 bits is 00000000 because the correct answer 256 cannot be represented in 8 bits. You can create longer integers by chaining numbers together; for example, a 32-bit machine can create 64-bit integers by concatenating two consecutive 32-bit integers (upper 32 bits and lower 32 bits).

Negative numbers can be represented in many different ways. In one system, the most-significant bit is used as a sign bit (0 = positive and 1 = negative); for example in 8 bits 10000011 represents -3 and 000000110 represents +6. This is called **sign plus magnitude** notation but is not used to represent negative integers. Its disadvantage is that there are two values for zero; that is 00000000 for +1 and 10000000 for -0.

Most computers use **two's complement** to represent negative integers. In this case the value of $-N$ is represented by $2^n - N$ where n is the number of bits. For example; in 8 bits -3 is represented by $256 - 3 = 253$ (i.e., 11111101). The advantage of two's complement is that it is easy to calculate (invert the bits of a word and add 1). The most-significant bit is a sign bit. Moreover, a subtractor is not needed because adding a complement automatically performs a subtraction; for example if you add the two's complement of -3 to 7 you would get the binary value corresponding to 4.

A negative two's complement number in n bits is in the range -2^{n-1} to $+2^{n-1} - 1$. Operations that exceed this range result in arithmetic overflow and give an incorrect result.

Fractional numbers can be represented in the same positional notation as integers by imagining a binary point; for example in 8 bits with a binary point in the middle 00111010 would be $0011.101 = 2 + 1 + \frac{1}{2} + \frac{1}{8} = 3.625$

This is called **fixed-point arithmetic** and it is not widely used.

The universal means of representing real numbers is called *floating point*, where a number is represented as

$-1^S \times 1.N \times 2^{E-B}$. This is similar to the everyday decimal notation 1.23×10^6 .

The S is a sign-bit representing negative ($S=1$) or positive ($S=0$). $1.N$ is a mantissa with a leading 1; for example 1.1101. The leading 1, the integer is automatically assumed so that only 1101 would need storing in memory. When the floating point has a mantissa in the range 1.000...0 to 1.111...0 it is said to be normalized (in this form, it has its highest precision).

The mantissa is multiplied by an exponent which is biased by B ; that is, the stored value is B greater than the actual value. If the bias is 128 then the exponent +4 would be stored at $4 + 128 = 132$. The advantage of this is that the most negative exponent is stored as 0.

Floating point arithmetic is called scientific notation and is used in calculations involving very large and very small numbers. Floating point is inexact because not all real numbers can be represented exactly in floating point form. Computer users must appreciate that arithmetic results using floating point may be inexact (although the error may be very small).

Financial calculations in accounting are almost never carried out using floating-point arithmetic because of errors due to fixed-size arithmetic and rounding. Fractions such as $1/3$ can be represented as "1 divided by 3" in some languages to avoid rounding errors. However, most computers use floating point and there is a tiny error in the conversion. The same is true of irrational numbers like the square root of 2 that cannot be represented by a finite string of bits in any base.

Examiner's comments

Most candidates were unable to explain numeric values in binary form. This is a key area of the syllabus and it is expected that centres develop a good understanding of how information is represented in the form of bits. Good answers showed some understanding of the advantages and disadvantages of each system. Overall, very few candidates were able to cover all elements mentioned in the question.

A2

Answer Pointers

a)

This question requires us to evaluate $S = a_0 \times b_0 + a_1 \times b_1 + \dots + a_{n-1} \times b_{n-1}$

```
S = 0 //clear the sum
FOR I = 0 to n-1
  S = S + ai x bi
END_FOR
```

A simple assembly language, explaining each operation.

```
MOV r0,#0           //register r0 is S. Store 0 in it (i.e., clear S). # indicates a literal value
MOV r1,#A           //register r1 is loaded with address of A (i.e., r1 is a pointer to A)
MOV r2,#B           //register r2 is loaded with address of B (i.e., r2 is a pointer to B)
MOV r3,#n           //register r3 is n. Load it with the number of elements to count
Loop MOV r4,[r1]    //copy element pointed at by r1 into r4 (get ai)
MOV r5,[r2]         //copy element pointed at by r2 into r5 (get bi)
MUL r6,r4,r5        //put a1 x b1 in r6
ADD r0,r0,r6        //Add r6 to r0; that is, update the sum by adding the new product
ADD r1,r1,#1        //Point to the next element in A
ADD r2,r2,#1        //Point to the next element in B
SUB r3,r3,#1        //Subtract 1 from the loop counter
BNE Loop           //Branch on not zero to Loop
```

b)

This question asks for **two** advances in computer architecture. Note the question is concerned with architecture and not hardware. Advances due to improvements in manufacturing such as faster clocks and better semiconductor technology will gain few marks. However, if new hardware is related to architecture, then the answer is acceptable (e.g., a faster clock is not what is being asked for but cache memory and multicore processing is).

Changes in computer architecture include wider registers and memory allowing operations on 64 bit words (8-bit processors can also operate on 64 bit words but they have to process 8 bits at a time). Some modern architectures increase the number of registers which reduces the number of accesses required to relatively slow DRAM.

Special instruction sets have been introduced to facilitate operations on data representing video and image processing. These were originally called multimedia instructions (or MMX extensions). Such instruction deal with special-purpose applications such as video or data encoding.

At the structural level (i.e., processor organization) one of the main advances has been pipelining – instructions are executed in a series of stages. As soon as an instruction finishes the first step in its execution, it is passed to the next stage of processing, and the next instruction fetched. At any instant 4 to over 30 instructions may be in the pipeline in different stages of processing. Pipelining increases performance without requiring faster technology.

Superscalar processing is another technique that improves performance without requiring faster hardware. Superscalar processors have multiple processing units. A group of instructions are fetched into buffers and are executed in parallel (however, logic is necessary to prevent one instruction being executed before another if that instruction needs the result of an earlier

instruction). Some computers even implement out-of-order processing in which the strict order of sequential processing is abandoned.

Limitations in memory speed are dealt with in several ways. In particular, cache memory brings instructions and data into very fast memory. The underlying principle is that for 80% of the time only 20% of the instructions are being used.

Cache systems can be speeded up by using prefetching instructions and logic to bring instructions into the processor ahead of their use. This means that an instruction is in place in the processor as soon as its execution begins – cache memory, instruction prefetching, out-of-order execution.

Over the last decade, the clock speed of computers appears to have reached a limit. This is imposed by the need to cool chips. Modern microprocessors dissipate up to 150 W in a small package. Higher clock speeds are not possible without radical changes in packaging technology. In order to overcome the limit on the clock rate, manufacturers have resorted to multicore processing; that is, replicating the CPU so that a processor package may have four or more CPUs operating in parallel (often with a shared cache memory). However, multicore processors are effective only when a task can be shared between processors.

Examiner's comments

Candidates did not show a good understanding of assembly language. Answers were unclear and did not explain the action of each instruction. Candidates did not understand what was required and merely attempted to list down a series of keywords.

A3

Answer Pointers

Modem

Digital information in a computer is transmitted as a string of 1s and 0s; that is, the data is transmitted in its original binary form (one exception to this is the PCI_E bus where data is encoded before being transmitted over the serial lanes).

However, the electrical characteristics of the communications networks that connect your house/business to and across the internet cannot handle simple strings of 1s and 0s (called baseband data). A modem (MODulator DEModulator) is required to match the characteristics of baseband digital signals to data paths that have a restricted bandwidth (band pass channels).

Several decades ago, most computer users used the PSTN (telephone network) for computer communications. The telephone network has a bandwidth of approximately 30 to 3,000 Hz. Consequently, digital data has to be converted into waveforms that occupy the available telephone bandwidth. First generation modems used amplitude modulation where the amplitude of a sine wave was changed to represent 1s and 0s. The frequency of a sine wave can also be changed to transmit data (FM), and the phase of the wave can be changed (PM). Eventually, very efficient modulation systems that change both the phase and amplitude simultaneously were used (QAM – Quadrature Amplitude Modulation). Telephone modems reached speeds of up to 56K bits/s.

Modems can operate in one direction at a time (half duplex) or data can be transmitted in both directions simultaneously (full duplex).

Even the fastest telephone modem speeds are low by today's standards. One solution to high-speed data transmission was to employ part of the existing telephone network – the part that links your home to the telephone switching network. This so-called “last mile” of a telephone path has a bandwidth far higher than the 30-3,000 Hz end-to-end bandwidth. This modem is called an ADSL (asymmetric subscriber line) modem. The term asymmetric indicates that data is transmitted in one direction faster than in the other (because most users download more data than they upload). The high bandwidth downstream channel provided by ADSL is of the region of 8 Mbits/s. ADSL uses frequencies outside the normal telephone band – you can think of it as piggybacking (hitching a ride) on the line between your home and the local exchange. Standards from ADSL range from about 1.5 to 24 Mbit/s (downstream), to 0.5 to 3.5 Mbit/s (upstream).

A cable modem performs exactly the same job as a telephone or ADSL modem: it matches digital data to data on a high-speed coaxial cable (which itself is often connected to a fibre-optic network). Cable modems require the installation of cable networks which are now common in Europe, the Americas, and Asia.

Cable modems are capable of speeds of the order of 30 Mbits/s although the actual speed depends on the amount of traffic on the cable because many users have to share the same data path. Incidentally, in some countries the cable modem belongs to the cable company rather than the user. The output (i.e., interface to the host computer) of both ADSL and cable modems is usually an Ethernet connector. NOTE that most modems also provide router functions (i.e., the modem and router are combined).

Hub

A network of computers in a home may connect several devices (multiple computers, storage devices, printers). A hub is a device that connects multiple Ethernet segments together. It operates at the bit level (layer 1 of the OSI model) and simply routes bits between segments. The hub essentially lengthens the Ethernet and shares its bandwidth between the segments. The Ethernet behaves normally in the sense that only one device can transmit at a time, no matter which segment it belongs to. The hub can be used to extend the number of connections to a router. If a router has 4 connections and you have 7 networked devices, you can use a hub to allow three devices to share one of the router's ports.

WiFi Router

A router is a far more complex device than a hub that simply copies all data between segments of an Ethernet. A router moves data packets between different networks. When a router receives a data package it decides which network to forward the packet to.

Routers perform address translation which allows several computers to be connected to a single cable modem. The router automatically forwards data to and from the appropriate computer to the network. Routers normally include security elements such as a firewall that examines all data packets and decides what to do with them.

A WiFi router often includes the functions described above, but it can also manage a WiFi network. WiFi is a short range communications technology operating in the GHz microwave band. A wireless router can communicate with a group of computers (and storage devices and printers) exactly like a wired network. The router takes care of addressing (i.e., moving information between appropriate pairs of devices). If a user at computer X wants to print something on printer Y, it is necessary only to click the appropriate icon and the router will take care of the actual routing of the wireless packets. To the user, the Wireless router appears exactly as a conventional wired router. Note that the wireless router also takes care of the wireless network's security mechanisms.

Examiner's comments

Candidates made a good attempt to answer each of these key networking devices. Overall, answers were sketchy and lacked depth. Candidates must endeavour to fully explore the operation and use of networking devices. For Section A questions it is not appropriate to just briefly explain what the device does.

A4

Answer Pointers

Computer Performance

A wide range of factors affect a computer's overall performance. Some of these are:

Semiconductor technology. Moore's law states that the number of devices per chip doubles every 18 months or so. Increasing the number of devices tends to make them smaller which increase their speed. In general, increases in the performance of semiconductor technology have been the main contribution to today's faster computers.

Feature size. This is a term used to indicate the size of a transistor on the chip. In 1971 the feature size was 10 micrometre. In 2106 it will be 10 nanometre. This is an increase of 1,000.

Computer Architecture. Register sizes have increased which means that more bits can be processed in a single operation. Moreover, instructions have been developed for special purpose applications such as multimedia processing.

Cache memory. Memory speed is one of the limiting factors of computer performance. The increase in the size of cache memory from a few KB to 8MB or more means that code can run in very fast memory.

Buses. Buses and their protocols have been improved making it possible to transport data from functional unit to functional unit. For example, many PCs now use the PCI express bus which is far faster than previous buses.

GPUs. Graphics processor units were designed to control graphics cards. However, they have taken over some functions from the CPU and can perform parallel processing with multiple CPUs (particularly for scientific calculations and games).

Multiple cores. Multicore processors can divide a task between several processors and increase the overall speed of the processor.

Compilers. Developments have appeared in software; for example, modern compilers are more efficient than earlier versions.

Memory. Although DRAM main store has not increased in speed as much as the CPU, flash memory has been used to speed up operation by implementing secondary storage as SSD (solid state disk) which is far faster than conventional magnetic disks with rotating platters.

Wider data widths. First generation microprocessors had 8-bit data buses. Today's processors use much wider busses (e.g., 64 bits) to move more data at a time. In GPUs, buses may be over 100 bits wide.

Development of chip sets. Today's motherboards contain advanced chip-sets that interface the CPU to memory and to peripherals (especially USB and the SATA disk interface). These chip sets can speed up operation by taking over complex actions that would once have to have been performed by the CPU.

Measuring Performance

Performance can be defined as the speed at which a task is performed. It is difficult to measure the performance of a computer because performance includes so many aspects: the CPU, memory, compiler and software.

The performance of a computer is very much dependent of a range of factors. Some compilers are more efficient than others and produce faster machine code. Even the writing of programs is significant because software can be written to exploit the characteristics of a computer (e.g., exploiting cache memory).

Performance was once measured crudely by comparing clock speeds. This is not a good idea because it gives little idea of how fast processors execute tasks. Some people started to use MIPS (millions of instruction per second) to measure how fast instructions are executed. Again, this sounds good but it does not take into account how much work the instructions perform. At best, clock speed gives only a very crude indication of performance (and even then only when very similar processors are being compared).

Performance is often measured by means of benchmarks which are specific programs that test the processor's capabilities. Because some programs might stress certain operations over others (addressing, data access, arithmetic, floating-point) it is normal to use a suite of test programs and take an average of their results. One of the most popular benchmarks is SPEC which uses a suite of programs to get a time for each individual benchmark. The results are normalized by dividing each result by the figure obtained by a standard base computer and then averaging the results (using geometric averaging). This gives a SPEC mark for the computer.

Power Limits to Performance

For many years, the speed of processors has been increased simply by increasing the clock speed and making it faster.

An important limitation in computing is power. Modern high performance processors use a lot of power (high-performance processors use in the region of 100 W). This power is created in a small slice of silicon and it must be removed rapidly if the temperature is not to rise and destroy the chip itself. For this reason, modern processors have large heat sinks and fans to keep them cool (although the actual chip temperature may be 70 degrees C).

The power dissipated by a transistor is determined by several factors. One factor is the speed at which it switches (i.e., the clock speed). The power dissipation in a gate is given by three components: static power, leakage, and dynamic (switching) power. The dynamic component is proportional to the square of the voltage and the frequency. Reducing chip voltage reduces power – but that cannot be continued to very low voltages because of the nature of silicon semiconductors.

Increasing clock frequency increases power dissipation and we are now close to the limit where clock frequencies cannot be raised above about 4 GHz without requiring new form of cooling. For this reason, manufacturers have resorted to using multiple cores running at a lower frequency.

Examiner's comments

Most candidates who attempted this question scored well. It was good to note that previous guidance given in examiner's reports have been taken on board when covering similar topics. Part b) was poorly attempted by most candidates.

SECTION B

(Candidates were required to answer FIVE out of the eight questions set)

B5

Answer Pointers

TCP/IP model is made of four layers:

Application layer responsible for coding of data packets,

Transport layer monitors the end-to-end path selection of the packets,

Internet layer responsible for sending packets through different networks

Data Link layer provides services to the internet layer.

Examiner's comments

This topic has been examined in previous sittings. Many candidates did not read the question well. Only 4 marks were available for part a). Answers did not require a detailed explanation of the TCP/IP model. Candidates had to cover the 4 layers in details in part b).

B6

Answer Pointers

Anti-virus software – this is needed to stop, detect and remove unwanted computer programs called viruses. Anti-virus software provides a range of facilities both on local and network computers. Anti-virus software must be regularly updated to ensure that the latest viruses being released are detected.

Anti-spyware – this is needed to combat malicious programs called spyware on a computer. The spyware will normally infiltrate a computer without the knowledge of the user. The anti-spyware must be able to regularly perform checks and detects these malicious programs.

Access Control List – this is configured at server or router level. It enables to filtering of data and user IP addresses attempting to access a network system.

Popup blocker – a utility which detects and stops a popup (small program which manifests itself randomly to collect data and information about user). The popup blocker must be up to date and active on a computer.

Examiner's comments

This question was poorly attempted. Candidates were unable to explain key elements of security threats and measures to deal with these. Centres are advised to develop suitable scenarios of these main security topics and show how these are important in day to day use of computers and networks.

B7

Answer Pointers

- a) Virtual LAN, or VLAN, is a networking technology that allows networks to be segmented logically without having to be physically rewired. This is done using Ethernet switches. Each device (host) in the network is assigned a VLAN ID. Broadcasts sent by a host are only received by hosts with the same VLAN ID.
- b) An intranet is a network which is set up and owned by an organization. The intranet is based on the organisation's LAN. It is based on TCP/IP technologies. The intranet allows users to share data and resources among each other. The intranet does not allow access from outside the organization

Examiner's comments

Most candidates were able to clearly distinguish between the two types of networking principles. Good answers included detailed technical information as well as the uses of intranet and VLAN. Weaker candidates merely covered uses of these.

B8

Answer Pointers

- a) The purpose of DRAM (Dynamic Random Access Memory) is to store volatile data inside a computer. Unlike SRAM, DRAM is slower and consumes large of power.
- b) USB – short for Universal Serial Bus is a port available on modern devices including Smartphones. USB provides flexibility on connecting different types of devices including keyboards, printers, video cameras, etc. USB is faster than the serial ports and can be used to connect multiple devices.
- c) The hard disk stores permanent data. With increasing computer capabilities, larger hard disk space is required. The speed at which the hard disk is accessed is also important.

Examiner's comments

The purpose of this question was to test candidates' basic understanding of storage devices. Candidates were able to cover a range of issues on these devices. Good answers included uses of the devices. Weaker candidates did not understand the operation and up to date storage mediums.

B9

Answer Pointers

- a) RJ45 refers to a type of connector with eight wires commonly used to connect computers to a network. RJ45 has been extensively used in Ethernet cabling.
- b) HDMI (High Definition Multimedia Interface) allows information to be transferred from a variety of devices to and from a computer. HDMI enables high quality, high definition images needed for video and graphics to be output.
- c) A technology that uses glass (or plastic) threads (fibres) to transmit data. A fibre optic cable consists of a bundle of glass threads, each of which is capable of transmitting messages modulated onto light waves. Fibre optics has several advantages over traditional metal communications lines. Fibre optic cables have a much greater bandwidth than metal cables. This means that they can carry more data.

- d) Cat5 – Short for Category 5 is a type of cabling which consists of twisted pairs of copper wire terminated by RJ45 connectors. Cat5 is used to connect computers in LANs

Examiner's comments

These topics have been examined in previous sittings. Most candidates who paid attention to guidance given in the past were able to write good answers. Candidates scored well in this question.

B10

Answer Pointers

- a) Bandwidth is the term used to refer to data transmission rate. It is the amount of information that can be transmitted along a channel.
- b) Ipconfig is a command line tool which is used in Windows based system to control network connections. Ipconfig shows all IP addresses of devices connected within the TCP/IP network configuration.
- c) FAT32 (File Allocation Table) used within Microsoft Windows to locate files on a disk. Since files can be fragmented, the FAT keeps track of all the pieces associated with a particular file.
- d) A proxy server is an 'intermediate' server that sits between a client application such as a Web server and a real server. It intercepts all requests to the real server to see if it can fulfil the requests itself. If not, it forwards the requests to the real server.

Examiner's comments

This question was set to test candidates' basic knowledge of practical computing. The percentage of candidates producing satisfactory answers for this question was low. Centres should endeavour to provide guidance on these key networking tools and techniques.

B11

Answer Pointers

A file structure provides a mean of organizing and storing data in a computer system. Different types of file structure (types) that computers handle are:

- a) A structure to store files on a hard disk by having a root directory, sub directory, etc. These allow for easy access of files.
- b) File formats are needed to ensure each application to store data. Typical examples include .EXE, .DOC, .PDF, etc.

Examiner's comments

Candidates were able to provide reasonably structured answers covering a range of reasons why a file structure is needed in a computer. Most candidates were able to identify various file formats. However, answers were often very superficial.

B12

Answer Pointers

With reference to any modern operating system (typically PC based) e.g. MS Windows. The operating system provides some of the facilities:

Kernel – the main part of the operating system which loads first and remains in the main memory. It provides services to the other parts of the operating system and applications which run on the computer. The kernel provides services related to tasks management, disk management and process management.

Multiprocessing – The operating system is able to handle various tasks on several processes at one time. Programs run concurrently. In a multiprocessing environment, the operating system must allocate resources to various processes.

Scheduler - this is the part of the operating system which provides services with handles resource coordination. The scheduler ensures that resources are accessed and serviced when requests are made.

Examiner's comments

Most candidates were able to appreciate the range of functions of operating systems. Some candidates were able to provide detailed answers referring to their own experience using operating systems. Weaker candidates confused operating systems with how the computer hardware operates.