Engineering Tripos Part IIA THIRD YEAR

Module 3F5: Computer and Network Systems

Solutions to 2012 Tripos Paper

Authors: Andrew Gee and Tim Wilkinson

1. Adders

(a) Amdahl's Law says that it is of paramount importance to optimise the speed of those components that are most frequently used in a computer system ("make the common case fast"). For most instruction set architectures, ALUs are used at least once in the execution of just about every instruction. Faster ALUs would therefore have a significant effect on the speed of the overall system. [10%]

(b) Carry lookahead can be used to determine the carry inputs to each full adder without using ripple carry. For each bit i of the adder, we define two signals, generate gi and propagate pi. Bit i generates a carry if the two bits it is adding are both 1, and propagates a carry if either of the two bits it is adding is 1:

$$gi = ai.bi \qquad pi = ai + bi$$

c1, the carry into bit 1, will be 1 if either bit 0 generates a carry or c0 is 1 and bit 0 propagates a carry:

$$c1 = g0 + p0.c0$$

Likewise for c2 and c3:

$$c2 = g1 + p1.g0 + p1.p0.c0$$
$$c3 = g2 + p2.g1 + p2.p1.g0 + p2.p1.p0.c0$$

These expressions show how the carry-in signals can be obtained without waiting for them to ripple through a 4-bit adder. [20%]

(c) (i)

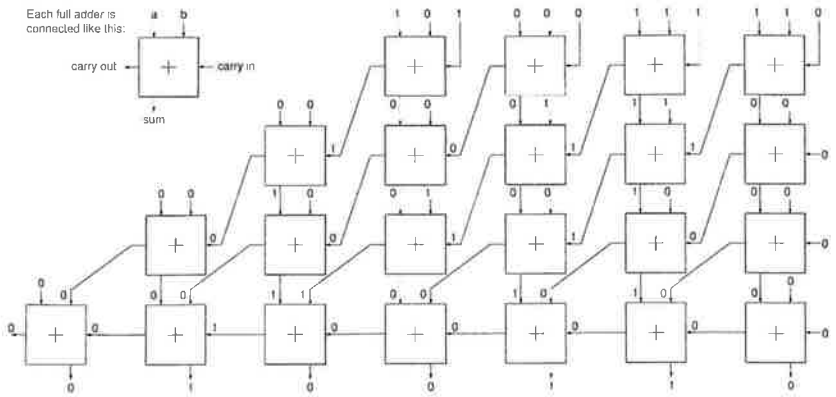

The carry save adder produces the correct result ($0100110_2 = 38_{10}$). [20%]

(ii) There is no ripple carry until the final stage. The adders in the first row can operate immediately. As soon as they are finished (latency T), the adders in the second row can operate, and so on down to the final row which takes longer since it uses ripple carry. So there are $N-2$ fast rows, each with a latency of T, followed by a $m+N-2$ bit ripple-carry row, with latency $(m+N-2)T$. The total latency is therefore $(m+2N-4)T$.

For the more obvious approach, and again assuming a worst case scenario requiring one extra bit each time a number is added in, we have a latency of mT to add the first pair, then $(m+1)T$ to add in the third number, then $(m+2)T$ for the fourth, and so on, giving

$$[m+(m+1)+(m+2)+(m+3)+\ldots+(m+N-2)]\,T = \frac{N-1}{2}(2m+N-2)T$$

since this is the sum of an arithmetic series. [25%]

(iii) One obvious way to improve the latency is to replace the final ripple-carry stage with a carry-lookahead adder. Also, not all the full adders are required, since the adder at the left of the second row can never generate a carry, so the adder at the left of the third row will always be adding three zeros: its outputs might as well be replaced by 0 signals. We can even dispense with the adder at the end of the second row, wiring its carry in signal directly to the input of the adder immediately below. Finally, depending on the particular values of N and m, we will not always need an extra bit to store the sum in each row, further reducing the full adder count. [15%]

(iv) The classic application of carry save adders is to sum partial products in multiplication hardware. [10%]

2. **Pipelined datapaths and hazards**

(a) The term *hazard* is used to describe dependencies between instructions which disrupt the operation of a pipelined datapath. *Data hazards* occur when an instruction requires data before a previous instruction has written it to the register file. *Branch hazards* occur when the address of the next instruction is required (for instruction fetching) before an earlier conditional branch instruction has been evaluated. [20%]

(b) **Pair 1.** Counting from when the first instruction is at the IF stage, `$9` is written on clock cycle 5 and read on clock cycle 3. This is an actual data hazard. It may be resolved by: (i) scheduling three independent instructions between the two adds (`nop` in the worst case); (ii) stalling the pipeline for three cycles; or (iii) forwarding data from the EX/MEM and MEM/WB pipeline registers to the ALU input, in which case no stalls are required.

Pair 2. The memory location is written on clock cycle 4 and read on clock cycle 5. So there is no actual hazard here.

Pair 3. `$9` is written on clock cycle 5 and read on clock cycle 3. This is an actual data hazard. It may be resolved by: (i) scheduling three independent instructions between the `lw` and the `add` (i.e. a long delayed load); (ii) stalling the pipeline for three cycles; or

(iii) forwarding data from the MEM/WB pipeline register (but not the EX/MEM pipeline register, since the data does not exist at the EX stage) to the ALU input. With (iii) we still require one stall, or alternatively we can specify `lw` as a short delayed load, requiring the compiler to schedule an independent instruction in the next slot.

Pair 4. `$9` is written on clock cycle 5 and read on clock cycle 3. This is an actual data hazard. It may be resolved by: (i) scheduling three independent instructions between the `lw` and the `add` (i.e. a long delayed load); (ii) stalling the pipeline for three cycles; (iii) forwarding data from the MEM/WB pipeline register to the ALU input, with a single stall or delayed load; or (iv) forwarding data from the MEM/WB pipeline register to the memory input, in which case no stalls are required. [40%]

(c) (i) We assume a two-way superscalar architecture. When two instructions are issued, one of them must be an arithmetic/logic operation or a branch, while the other must be a load/store.

The pipeline registers need to be widened to hold the intermediate states of two instructions at a time. The path to instruction memory needs be widened to deliver 64 bits (two instructions) at each clock cycle. The decoding logic at the ID stage needs to examine two opcodes at a time. The register file still holds the same number of registers, but needs further inputs and outputs: we might need to read four registers at a time (two for an arithmetic/logic operation and two for a `sw`) and write two registers (one for the arithmetic/logic result and one for a `lw`). We need an extra ALU at the EX stage, so we have one for branch target calculation, one for an arithmetic/logic operation or branch comparison and one for load/store address calculation. The MEM stage can stay as it is, since only one of the concurrent instructions is allowed to be a load/store. We also need extra hazard detection logic.

Superscalar operation will generally enhance throughput, since multiple instructions are issued at each time slot, but this means more dependencies between the instructions in the pipeline and therefore more stalls: it is highly unlikely that the pipeline will be able to run at full capacity.

(ii) For simultaneous multithreading (SMT), we need extra hardware to deal with each process's independent state. This means a separate program counter, a separate register file and a separate page table (implying, at the hardware level, a separate TLB) for each thread.

Since instructions sharing the pipeline can now be from independent threads, with no dependencies, we can expect fewer stalls and higher throughput, though this assumes that there are two or more CPU-hungry processes running concurrently. The only worry comes from the memory system: the cache is now being shared, without time slicing, between multiple threads. If they are totally independent processes, we can expect more cache conflicts and more misses. If they are multiple threads cooperating on the same task in a shared memory framework, things might not be so bad. [40%]

Question 3

a) Layer 1: The physical layer. This defines an interface in terms of the connections, voltage levels and data transmission rate, in order for signals to be transmitted bit by bit.

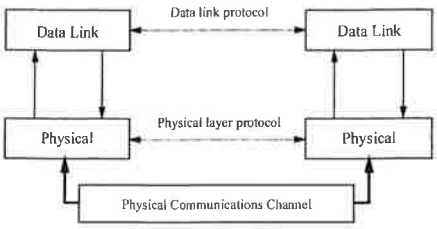

Before designing a transmission system, the designer must answer these questions: How much bandwidth for acceptable quality and service? Over what maximum distance is the transmission? What are the acceptable interference levels? Crosstalk? Can we use existing installed media?

i) Digital subscriber lines: The main limitation on the bandwidth of a Cat 3 UTP telephone line is filter placed at the exchange end before the voice signal is digitised. This limitation means that a normal modem must squeeze all of the data through a 3.7kHz wide band. The Cat 3 UTP is in fact capable of transmitting data over short distances at frequencies in excess of 1.1MHz, which is a property exploited by digital subscriber line modems which bypass the filter in the line card at the exchange and access all of the available line bandwidth.

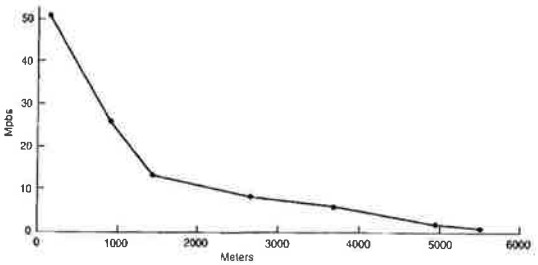

ii) GSM - It was simplest to make mobile networks digital, operating in a new frequency band (1.8GHz). GSM uses both FDM and TDM. The available spectrum is broken up into 200kHz bands; within each band TDM is used to multiplex multiple users. GSM was originally designed for use in the 900MHz band, but later frequencies were allocated at 1.8GHz, and a second system closely patterned on GSM was set up, known as DCS 1800. Full GSM now has up to 200 full duplex channels per cell; each channel has a down-link and an up-link frequency.

iii) Quadrature amplitude modulation (QAM) modem. By combining amplitude and phase modulation it is possible to build up a constellation of data points that are transmitted by the modem at a physical baud rate less that 3.4kHz. The choice of constellation must balance out the number of phase steps and amplitude levels. A strong constellation should have a minimum number of amplitude levels and a quantised even phase step. Much higher transmission rates can be achieved using more complex constellations in conjunction with error correction and compression algorithms.

vi) Pulse code modulation (PCM): Modern digital exchanges offer much higher performance and can accept both digital telephone and computer data through services such as the integrated services digital network (ISDN). Plain ordinary telephony services (POTS) are digitised using PCM. The digitisation of the PCM bit stream is done at the exchange before entering the digital network, unless a specialised ISDN telephone is used.

v) Bulk data transmission: A surprising way of transmitting large quantities of data is via fixed storage media such as hard drives or optical disk. For instance, a compact disk is currently capable of storing 650Mb of data, DVDs have capacities up to 50Gb (HDD is terabytes) and they can transfer data at rates in excess of 2Mb/sec. Imagine the transmission capacity of a van load of compact disks travelling up the M1.

b) i) A coaxial cable consists of a stiff copper wire as an inner conductor, inside a solid insulator, that is itself inside a closely woven braided wire mesh that acts as an outer conductor. This is then covered in an insulating protective cover. It is possible to use Maxwell's equations to calculate the electric and magnetic fields within the coaxial cable to give the attenuation in the cable due to the inner conductor and the solid insulator, along with the frequency cut-off for a given cable dimensions: High frequency cut-off, High immunity to interference and crosstalk.

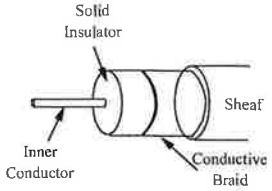

Coaxial cable is a good transmission medium for high data rates over relatively short distances, however it is rather expensive and is not ideally suited to long distance transmission and has been largely replaced by optical fibre. It was extensively used in Ethernet, bus type systems where many nodes can be tapped into a single coaxial backbone. Ethernet was originally a proprietary LAN standard (predating the IEEE 802.3 standard) developed by Xerox. The original design was based on a length of coaxial cable, with 'tee-offs' to individual work stations, with a maximum of around 500 stations. The idea was to simplify the cabling needs of offices with many computers in use. Each time a new device was installed, a new tee-off could be installed from the corridor into the particular office required.

The legacy of coaxial based ethernet is the 1500 byte MTU.

ii) Twisted pair cabling - Almost every home contains twisted pair cabling as normally used with a telephone connection. The bandwidth of this cabling is severely limited by its length and the way in which it is used within the telecommunications networks. Twisted pair consists of two insulated copper wires twisted together. Any pair of wires held close together will have an associated capacitance which is susceptible to outside interference or crosstalk. Theory suggests that random collision of a large number of devices can lead to transmission degradation under heavy traffic, however traffic is rarely random as most LAN transmissions involve a central server system. Traffic problems can be alleviated by subdivision into smaller LANs.

Twisted pair replaced coaxial cable as LAN hubs (or switches) have become necessary as the data rate is too high for shared media. This caused the development of LANs with structured cabling, using hubs and twisted pair (10 and 100baseT), in a star configuration.

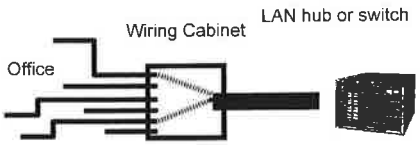

When installed as part of a structured cabling scheme (nowadays the most common realisation of ethernet), twisted pair or coaxial cabling provides for the transmission medium. Multiple twisted pair cables are usually installed in each individual office and near each desk, and are wired back to a wiring cabinet. Next to the wiring cabinet is a LAN hub which replaces the coaxial backbone, so the arrangement is often called a collapsed backbone. The bus topology still exists, but only within the hub itself. Hence the CSMA/CD protocol is no longer required (along with the LLC).

c) When working within the lower levels of the OSI model, it is often useful to add a common communications device such as a modem to utilise the available communications channel. It is useful to split the lower end of the OSI model into two parts: The data terminal equipment (DTE), The data circuit (terminating) equipment (DCE)

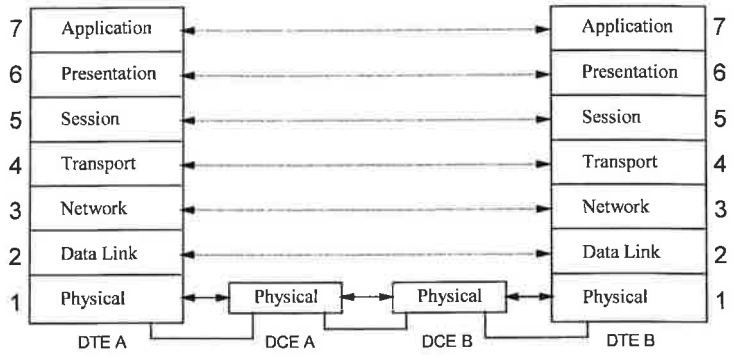

This makes for a common DTE/DCE interface standard. This allows for very important physical layer standards such as connecterisation and voltage levels to be standardised as well as proprietary features to be included in the physical layer. This allows network equipment providers to have a competitive edge as long as it does not encroach on the OSI standard protocols. Physical layers within the DCE connections can have multiple media, different error and flow controls and different metrology and feedback mechanisms for QoS monitoring.

d) Multi-label protocol systems (MPLS) - The process of routing packets cannot easily guarantee delay (latency) across the network. This is a problem for voice services across packet switched networks (such as VoIP). One way to fix this is to group packets through a common network together with a common global header (or label or tag) which gets the group of packets to the other side of the network with minimum delay. This puts a lot of pressure on the routers at the edge of the network to find suitable packet groups.

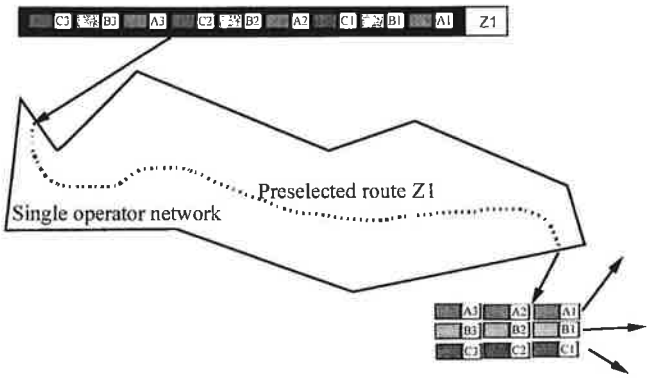

This could be argues as a DTC/DCE scenario as the MPLS information only exists within the MPLS network nodes and the original data packets are encapsulated within the labelling system. It could be argued that this is not a purely physical layer procedure, but in effect it is as the higher layer information is purely used to interpret local routing within the MPLS system. It is of course, highly proprietary and so cannot be seen outside of the MPLS network ingress and egress nodes.

Question 4

a) On a CSMA/CD LAN the terminals do not request permission from a central controller before transmitting data onto the transmission channel; they contend for its use. Before transmitting a packet of data, a sending terminal 'listens' to check whether a path is in use and if so, it waits before transmitting its data. Even when it starts to send data, it needs to continue checking the path to make sure that no other stations have started sending data at the same time. If the sending terminal's output does not match that which it is simultaneously monitoring on the transmission path, it knows there has been a collision. To receive data, the medium access control (MAC) software in each terminal monitors the transmission path, decoding the destination address of each packet passing through to find out whether it is the intended destination. If it is, the data is decoded, if not the data is ignored. Theory suggests that random collision of a large number of devices can lead to transmission degradation under heavy traffic, however traffic is rarely random as most LAN transmissions involve a central server system. Traffic problems can be alleviated by subdivision into smaller LANs.

This standard is the logical link control (LLC) defined by IEEE 802.2 or ISO 8802.2. The LLC provides a standard communication interface equivalent to that provided by OSI layer 2 to OSI layer 3 and uses HDLC calls like those used in the X.25 data-link layer protocols.

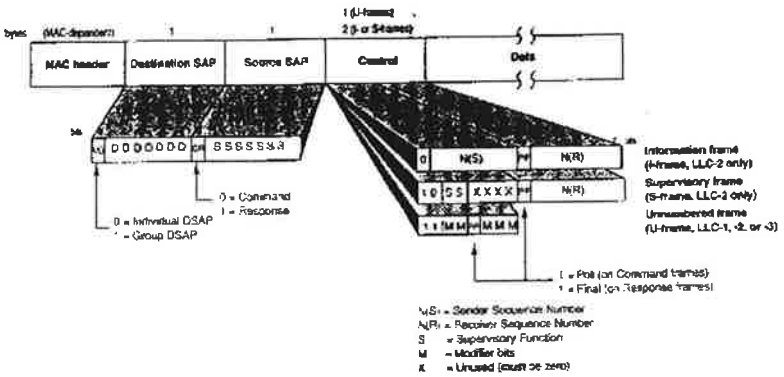

b) Initially, early WLANs were run using CSMA/CD quite efficiently until it was realised that there were occasional miss-transmissions and blocked links dropping frames. This problem was due to the use of CSMA/CD in conjunction with the wireless physical layer.

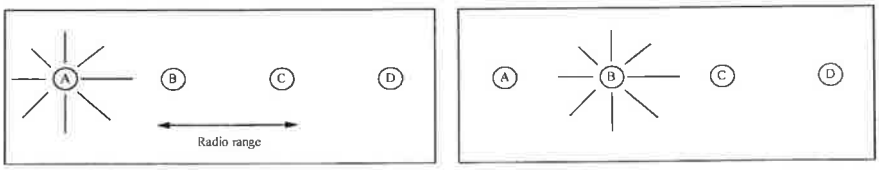

If we consider the above scenario, where four stations are in the local area (they could be base stations or mobile stations).

i) The radio range is such that A and B are within each other range. C can possibly communicate with B and D but not A. When A is transmitting to B and C senses for a carrier, it will not hear A and will start transmitting as well, corrupting the original frame from A. This is called the hidden station problem.

ii) Similarly if B is transmitting to A, then C will sense the carrier and conclude that it cannot transmit, even though it could talk to D if it wanted. CSMA/CD is suited to wires, but not radio waves.

c) The solution is to use multiple access with collision avoidance (MACA), which became the basis for the IEEE standard family 802.11. In MACA, the sender stimulates the receiver to send a short frame which alerts other stations nearby of its impending activity and prevents collisions.

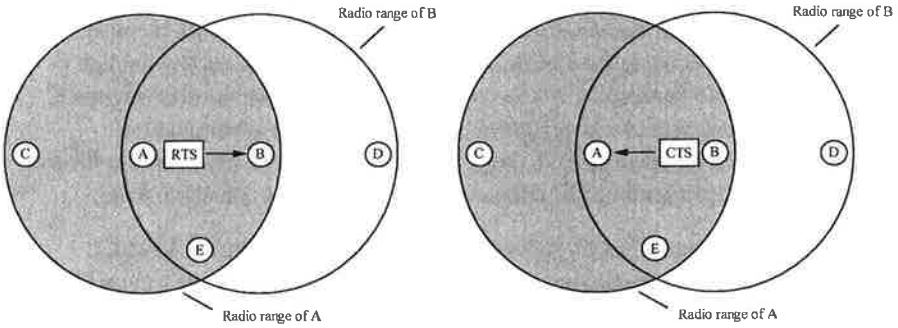

Consider the stations above. A sends a frame to B, known as a request to send (RTS) frame, which is a 30 byte frame which contains the length of the data to follow. B then replies with a clear to send (CTS) frame which also contains the data length. When A gets the CTS it starts transmitting data. Any station hearing the RTS from A must be quiet until the CTS from B has been sent. Any station close to B will hear the CTS and must remain silent for the length of the data frame. C is close to A, but not B so it must wait until the CTS has reached A, then it is free to transmit. D is close to B, but not A so it must wait until the data has been sent to B. E is in range of both A and B so it must be silent for the whole operation. Collisions can still occur if RTS frames are sent at the same time, hence colliding stations will use the same back off procedure as CSMA/CD. Later improvements were added such as acknowledgements and CSMA to improve the protocol (known as MACAW).

d) The IEEE 802.11 standard defines wireless LAN connections in two main frequency bands over relatively short distances, which have become part of the WiFi marketing standard and brand of products. The 802.11b is the international standard for wireless networking that operates in the 2.4 GHz frequency range (2.4 GHz to 2.4835 GHz) and provides a throughput of up to 11 Mbps. This is a very commonly used frequency. Microwave ovens, cordless phones, medical and scientific equipment, as well as Bluetooth devices, all work within the 2.4 GHz frequency band. The 802.11a specification for wireless networking operates in the 5 GHz frequency range (5.725 GHz to 5.850 GHz) with a maximum of 54 Mbps data transfer rate. The 5 GHz frequency band is not as crowded as the 2.4 GHz frequency, hence the 802.11a specification offers more radio channels than the 802.11b. These additional channels can help avoid radio and microwave interference.

Another standard is the 802.11g series which offers up to 54Mbps in the 2.4GHz band (now up to 125Mbps in some cases). Now available is the 802.11n which uses multiple antenna/receiver combinations to increase the reliability and data rate further. There are even WiFi components which use both frequencies and offer a service similar to that of the 10baseT (10Mbits/sec) Ethernet standard. The most challenging aspect of operating a wireless network is avoiding interference, multi-path and external crosstalk (other transmitters). This is done using very sophisticated modulation scheme such as orthogonal frequency division multiplexing (OFDM), frequency-hopping spread spectrum (FHSS), direct-sequence spread spectrum (DSSS).

Part IIA 2012

Module 3F5: Computer and Network Systems

Numerical Answers

1. (c) (ii) Carry save $(m + 2N - 4)T$, sequential ripple carry $\frac{N-1}{2}(2m + N - 2)T$.