ENGINEERING TRIPOS PART IIA

Tuesday 8 May 2007 2.30 to 4

Module 3F5

COMPUTER AND NETWORK SYSTEMS

*Answer not more than **three** questions.*

All questions carry the same number of marks.

*The **approximate** percentage of marks allocated to each part of a question is indicated in the right margin.*

There are no attachments.

STATIONERY REQUIREMENTS
Single-sided script paper

SPECIAL REQUIREMENTS
Engineering Data Book
CUED approved calculator allowed

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

1 (a) What properties of typical computer programs motivate the inclusion of a cache between the CPU and main memory? [10%]

(b) A certain computer has a byte-addressable memory with 4-byte words. Both the virtual and physical addresses are 32 bits wide and the page size is 4 KBytes. The translation lookaside buffer (TLB) is 2-way set associative with 64 entries and a block size of one page table entry. The primary data cache is 16 KBytes large and direct-mapped, with a block size of four words. Sketch a hardware schematic showing the entire process of retrieving a data word, starting from the virtual address. Include details of how the virtual and physical addresses are split into page numbers and page offsets, and cache/TLB tags, indices and offsets. Assume hits at both the TLB and cache levels, so there is no need to show the page table or the main memory system. [50%]

(c) The processor in the above computer has a baseline clock cycles per instruction (CPI) count of 1.0 when all memory references hit in the primary cache, and a clock speed of 1 GHz. The main memory access time is 100 ns, including all the cache miss handling, and the cache miss rate per instruction is 5%. How much faster would the computer be if we were to add a secondary cache (between the primary cache and main memory) that has a 10 ns access time for either a hit or a miss, and a miss rate per access of 40%? State any assumptions you make. [30%]

(d) At first sight, a secondary cache miss rate of 40% per access may seem surprisingly high. Explain why this is in fact perfectly realistic. [10%]

2 (a) In the context of parallel processing, define and very briefly explain the following acronyms: SIMD, MIMD, SMP, UMA, NUMA, SMT. [35%]

(b) Describe a method for enforcing cache coherency in a single bus SMP. [25%]

(c) Why might the optimal cache block size be smaller for a single bus SMP than for the corresponding uniprocessor? Illustrate your answer with two hypothetical memory accesses that would trigger a miss in the SMP with multi-word blocks, but not with single word blocks. [20%]

(d) Briefly compare and contrast network-connected NUMA MIMD machines and clusters. [20%]

(TURN OVER

3 (a) Compare the different requirements of voice services versus data services in the context of a computer or telecommunications network. Give an example of a network which has been specifically designed to implement each type of service and explain how that service is guaranteed. [30%]

(b) Explain how the two network examples given in (a) would react to a service requirement such as live, broadcast-quality video. [20%]

(c) One of the most important mechanisms in providing latency-based quality of service in internet protocol (IP) networks is the concept of the *fast path*. Use a simple diagram to show how the fast path in an IP network is implemented. Explain which fields of the IP datagram in Fig. 1 are vital in defining the fast path. [25%]

(d) An alternative approach to fast path routing is to use tag or label systems to allow a priority mechanism to be implemented. How might labelling be used to implement a video broadcast service? What are the limitations of this mechanism? [25%]

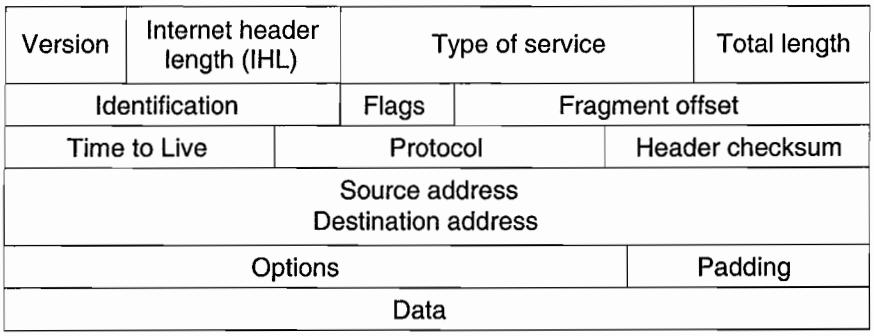

Fig. 1

4 (a) Explain what is meant by the terms *single-mode* and *multi-mode* optical fibre when describing the physical layer of a network. What are the relative advantages and disadvantages of each type of optical fibre? Which one would you use in a long distance telecommunications network and at what wavelength of light? Explain your reasoning. [30%]

(b) Where might you find optical fibre being used in a local area network (LAN)? Why is it becoming a more popular physical layer option in modern LANs? [20%]

(c) The synchronous digital hierarchy (SDH) is based around optical fibre transmission at the physical layer. Sketch a diagram showing how an SDH network would be constructed using synchronous multiplexers. Explain your choice of topology and define the role of the synchronous multiplexer. What is meant by the term *add/drop mode*? [30%]

(d) Give a brief explanation of how computer data, packet-based services can be incorporated into SDH. Why is this difficult when trying to map the internet protocol (IP) packets onto the SDH? [20%]

END OF PAPER

Part IIA 2007

Module 3F5: Computer and Network Systems

Numerical Answers

1. (d) 1.71 times faster.