ENGINEERING TRIPOS PART IIA

Friday 29 April 2011 9 to 10.30

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) Explain why the inclusion of a cache, between the CPU and main memory, generally improves a computer's performance. [20%]

(b) Consider the following three caches.

(i) Direct-mapped, 4 words per block, 65536 blocks.

(ii) 4-way set-associative, 4 words per block, 4096 sets.

(iii) Fully associative, 4 words per block, 32768 blocks.

For each case, state the size of the cache and where (i.e. which byte of which word of which block or set) the byte at address 0xA0973C8F would be stored were it to be fetched into the cache. State also what the value of the tag would be. Assume 32-bit addresses and words. Express cache indices and tags in hexadecimal. [30%]

(c) For a set-associative cache, sketch graphs showing how the miss rate and the miss penalty vary with (i) the block size and (ii) the degree of associativity. [25%]

(d) Consider performing a matrix multiplication on a typical modern PC using the C++ code in Fig. 1. On the same axes, sketch curves of execution time against `matrixSize` for the cases `transpose_c = false` and `transpose_c = true`. Explain briefly the curves' general characteristics. [25%]

(cont.

```
int a[matrixSize][matrixSize];
int b[matrixSize][matrixSize];
int c[matrixSize][matrixSize];
int ct[matrixSize][matrixSize];
bool transpose_c;

....

if (transpose_c) {
  for (i=0; i < matrixSize; i++)
    for (j=0; j < matrixSize; j++)
      ct[i][j] = c[j][i];
}

// Calculate a = b * c
for (i=0; i < matrixSize; i++)
  for (j=0; j < matrixSize; j++) {
    r = 0;
    for (k=0; k < matrixSize; k++)
      if (transpose_c) r += b[i][k] * ct[j][k];
      else r += b[i][k] * c[k][j];
    a[i][j] = r;
  }
```

Fig. 1

2 (a) What are the distinguishing features of reduced instruction set computer (RISC) architectures? [25%]

(b) Figure 2 shows the datapath and control for a single-cycle implementation of the MIPS instructions `lw sw add sub and or slt` and `beq`. The operation times of the main functional units are: 2 ns for the memory (read or write); 2 ns for the ALU and adders; and 1 ns for the register file (read or write). All other units have negligible latencies and there is no pipelining.

(i) Could this datapath be used to execute an instruction set that is not load-store? Justify your answer. [10%]

(ii) A hardware designer proposes a more compact implementation of the ALU and adders, with an operation time of 3 ns. What would be the impact on performance? What if only the adders, and not the ALU, were replaced with the more compact design? [25%]

(c) Consider implementing a new `lw_inc` instruction:

Example:
```
lw_inc $15,100($2) # reg $15 loaded with mem[$2 + 100], also
                   # reg $2 incremented to point to next word
```

Format:

6 bits	5 bits	5 bits	16 bits
op	rs	rt	immediate

(I-format)

Describe any modifications to the datapath, functional units and control required to implement the `lw_inc` instruction. List the state of each of the control signals for `lw_inc`. What impact might this new instruction have on performance? [40%]

(cont.

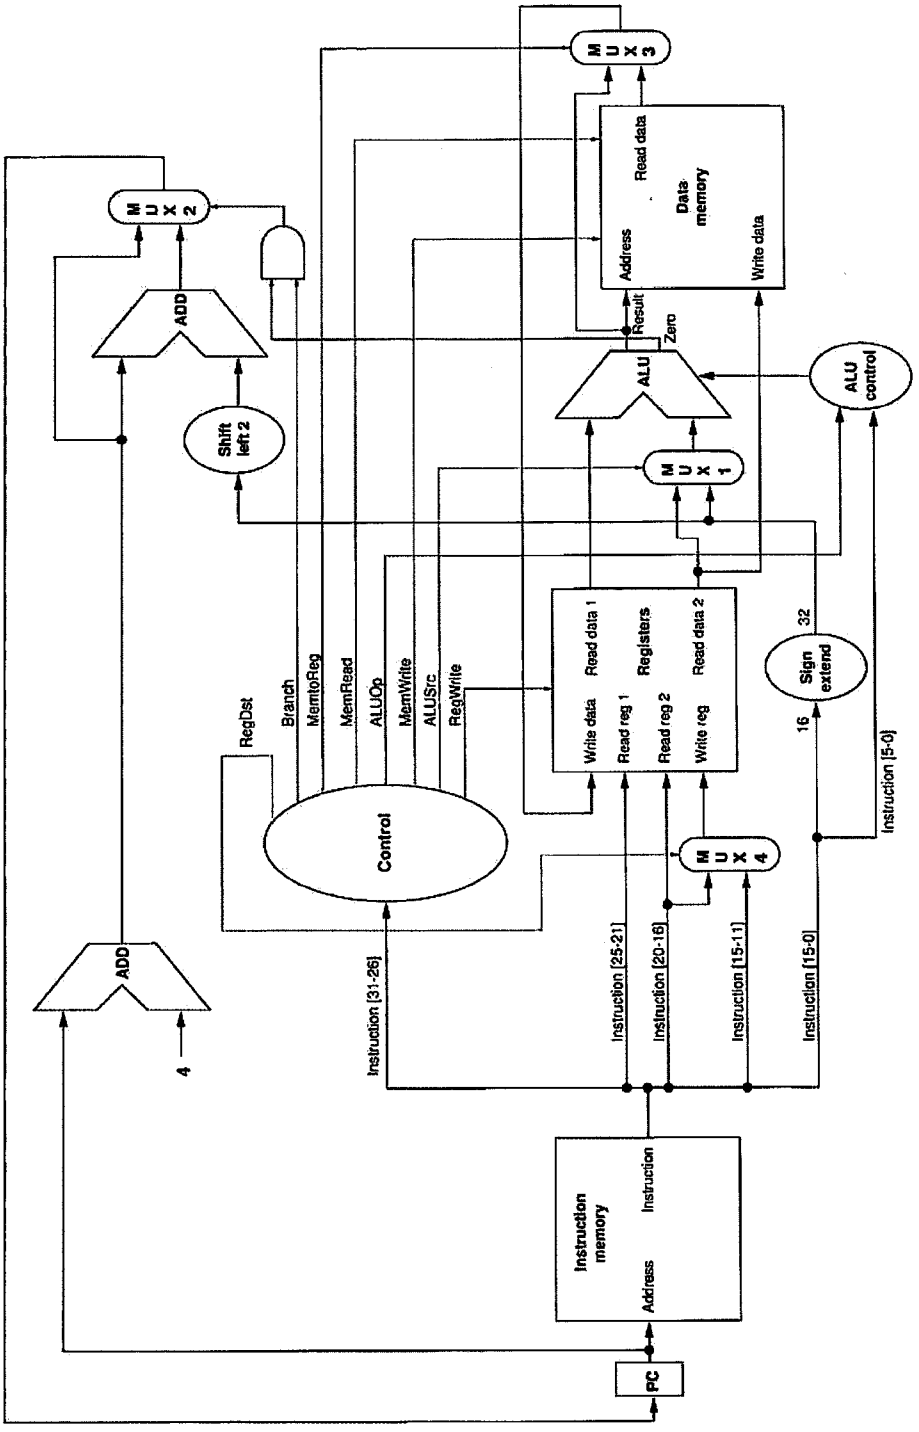

Fig. 2

(TURN OVER

3 (a) Describe the role of each of the network-dependent layers of the open systems interconnect (OSI) seven-layer reference model. Give an example of two services that might be found in each layer. [25%]

(b) Local area networks (LANs) have a unique structure in layer 2 of the OSI model, as defined by the 802 series of standards. Describe this layer 2 structure and explain why it played such an important part in the evolution of LANs in computer communication networks. [25%]

(c) Ethernet has become a dominant LAN protocol at layer 2 over the past five years. Explain what features of ethernet have been removed or bypassed, and discuss why they are no longer required by modern LAN hardware and network architectures. [25%]

(d) If a voice service were to be run over the transmission control protocol/internet protocol (TCP/IP) on top of ethernet, what sort of features could be employed to give a suitable quality of service? [25%]

4 (a) Explain the main differences between a connection-oriented circuit-switched network and a connectionless packet-switched network. Use the examples of data and voice services to highlight the strengths and weaknesses of each network type. [35%]

(b) Eight telephone users are grouped into four pairs of local users, as shown in Fig. 3. The users A, B, C, D, E and F are to be connected to a circuit-switched telephone network with full availability. Sketch a diagram showing how they can be connected using only three trunks/junctions. Describe all the parts of the network and give an example of the limitations of such a connection mechanism. How could its performance be improved? [30%]

(c) The same six users in (b) are now joined by two further users G and H, as indicated in Fig. 3. Sketch a new circuit-switched network showing how all of these users can be connected with full availability. Explain carefully the reasoning behind your switch and junction design. [35%]

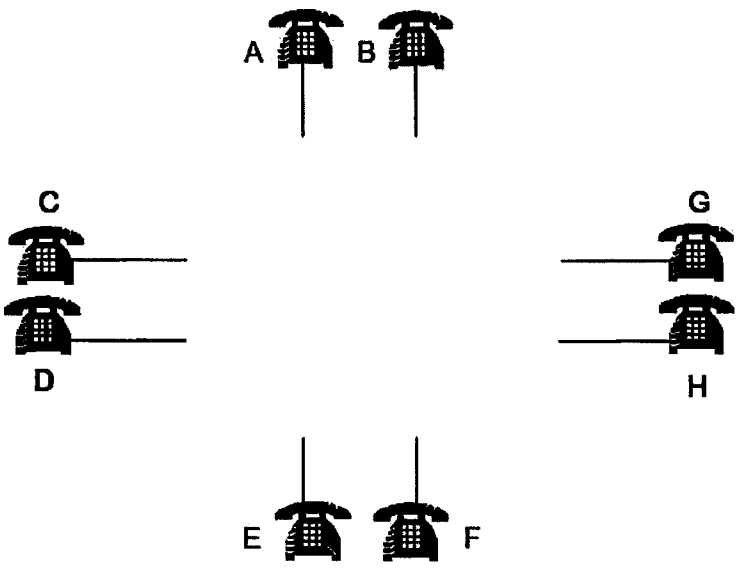

Fig. 3

END OF PAPER