CONFIDENTIAL

Paper 3.32 MEng3 Computing Parallel Architectures

Department of Computing
Imperial College

Examination paper **January 23, 1997**

First examiner: ..... **Second examiner:** .....

Paper 3.32(MEng3 test) Parallel Architectures

This examination is of TWO HOURS' duration.
Answer ALL THREE questions.

1 A major research organisation plans to buy a vector-pipeline supercomputer. They are considering two alternative products, both based on the DLXV described by Hennessy and Patterson. The standard product, $DLXV_1$, has the following characteristics:

$DLXV_1$: 64 64-word vector registers.
One pipelined vector FP add unit (start-up: 6 cycles).
One pipelined vector FP multiply unit (start-up: 7 cycles).
Two Vector load/store units, start-up time 12 cycles.
(assume that a vector load delivers one operand per clock cycle).
Strip mine loop overhead: 15 cycles per iteration.

The "enhanced" product is similar but has two copies of every FP functional unit, with full chaining between them. Because of the extra functional units, and the increased complexity of assigning operations to units, all the start-up overheads are doubled:

$DLXV_2$: 64 64-word vector registers.
Two pipelined vector FP add units (start-up: 12 cycles).
Two pipelined vector FP multiply units (start-up: 14 cycles).
Two Vector load/store units, start-up time 24 cycles.
(assume that a vector load delivers one operand per clock cycle).
Strip mine loop overhead: 30 cycles per iteration.

Consider the following pair of loops:

Loop 1:
```
for i = 1 to 6400 do
    A[i] := x*A[i] + y*A[i]
```
Loop 2:
```
for i = 1 to 128 do
    A[i] := x*A[i]
```

a Write down the DLXV assembly code for Loop 1, including strip mining (explain any assumptions you have to have to make concerning DLXV's instruction set).

b Find the number of clock cycles for Loop 1 on $DLXV_1$.

c Find the number of clock cycles for Loop 1 on $DLXV_2$. How does this compare to $DLXV_1$, and why?

d Find the number of clock cycles for Loop 2 on $DLXV_1$.

e Find the number of clock cycles for Loop 2 on $DLXV_2$. How does this compare to $DLXV_1$, and why?

f Discuss *briefly* the advantages, disadvantages and alternatives to having 32 128-word vector registers, instead of the 64 64-word vector registers assumed at the start of the question.

(The four parts carry, respectively, 20%, 15%, 20%, 10%, 10% and 25% of the marks).

2 Consider the following loop:

```
for i = 1 to N do
  for j = 1 to N do
    A[i+1,j+1] := A[i,j] + A[i+1,j] + A[i,j+1]
```

a Estimate the number of cache misses per iteration when executing the original version of the loop.

b Suppose the CPU has a small, write-back data cache with block size of two words. Estimate the number of cache misses per iteration when executing the original version of the loop. Take care to state any assumptions you need to make.

c Draw a diagram showing the iteration space for this loop, and the dependences between loop iterations.

d Demonstrate by drawing an appropriate diagram that the following blocked implementation of the loop is valid:

```
for ii = 1 to N step B do
  for jj = 1 to N step B do
    for i = ii to ii+B-1 do
      for j = jj to jj+B-1  do
        A[i+1,j+1] := A[i,j] + A[i+1,j] + A[i,j+1]
```

e Suppose the CPU has a 1K-word (8KByte) direct-mapped data cache with block size of *one* word. Estimate the number of cache misses per iteration when executing the blocked implementation of the loop using a suitable value for B. Take care to state any assumptions you need to make.

(The four parts carry, respectively, 10%, 20%, 30%, 20%, and 20% of the marks).

Turn over ...

3 An environmental research consultancy company seeks your advice on its plans top buy a high-performance computer to support a wide range of simulation work. The argument concerns two alternatives, K5V and K5MP, of similar cost. You are to advise on the performance and ease of use of the two.

K5V: 1 vector CPU:

- Clock cycle time: 1ns, average CPI in scalar mode: 2.0 not counting memory access delays
- Vector mode:
  - 64-word vector registers
  - 6-stage vector add/subtract unit, 12-stage vector multiply/divide unit
- Memory:
  - 1MB instruction cache
  - Data cache for scalars only (all vectors/arrays are uncached)
  - 128 memory banks, each with access time 50ns (delivers one 64-bit word every 50ns)
- Interconnection network:
  - Full crossbar connection between CPU and memory banks (communication delay included in memory access time above)

K5MP: 8 CPU shared-memory multiprocessor:

- Clock cycle time: 3ns, average CPI: 1.0 not counting delays due to cache misses; no vector mode
- Fully-pipelined FP unit (2 cycle add, 3 cycle multiply)
- Memory:
  - 1MB instruction cache + 1MB data cache per CPU, all operands cached
  - Cache line size: 4 64-bit words (i.e. 32 bytes)
  - One memory bank per CPU, with access time 50ns (delivers one cache line, ie 4 64-bit words, every 50ns)
  - Cache coherence maintained using invalidation protocol, directories maintained using singly-linked sharing list
- Interconnection network:
  - Hypercube topology (3 dimensions, 8 nodes)
  - Communications latency:
    - For cache lines: 1000ns plus 5ns per per intermediate node
    - For short messages: 200ns plus 5ns per per intermediate node

a What do you consider to be the main advantages of the K5V?

b What do you consider to be the main advantages of the K5MP?

c Describe briefly an example of a benchmark on which K5V would perform better than K5MP.

d Describe briefly an example of a benchmark on which K5MP would outperform K5V.

(The four parts carry, respectively, 30%, 20%, 25% and 25% of the marks).

End of Paper