

Master Copy

Paper Number(s): **E1.9A**

IMPERIAL COLLEGE LONDON

DEPARTMENT OF ELECTRICAL AND ELECTRONIC ENGINEERING
EXAMINATIONS 2005

ISE Part I: MEng, BEng and ACGI

**PRINCIPLES OF COMPUTING AND SOFTWARE ENGINEERING (PART A)
INTRODUCTION TO COMPUTER ARCHITECTURE**

Wednesday 8th June 2005 2:00pm

There are THREE questions on this paper.

Question 1 is compulsory and carries 40% of the marks.

**Answer Question 1 and EITHER Question 2 (carrying 60%)
or Question 3 (carrying 60%).**

This exam is **open book**

Time allowed: 1:30 hours.

Any special instructions for invigilators and information for candidates are on page 1.

Examiners responsible:

First Marker(s): Clarke, T.
Second Marker(s): Demiris, Y.K.

Students may bring any written aids into this examination.

Question 1 is compulsory, and carries 40% of the total mark, this equals a time of 36 minutes.

Answer EITHER Question 2 OR Question 3, each of which carries 60% of the total mark, corresponding to a time of 54 minutes.

Question 1 is compulsory

- a) Convert into a decimal integer or rational number each of the following hexadecimal words, interpreted as specified.
- (i) 0xc4e2, 16 bit unsigned
 - (ii) 0xfde2, 16 bit signed
 - (iii) 0x90600000, IEEE 754 floating point.
- [6]
- b) Write in assembler a length optimal ARM assembler fragment which implements the following pseudo-code, assuming the comparison to be signed:
- ```
if r0 < r9 + 5 then r1 := r2-r3 else r1 := r4+r5
```
- [8]
- c) Noting that  $2159 = (128-1)*(16+1)$ , determine a sequence of ARM data processing instructions that will set r1 equal to  $2159*r0$  in two machine cycles leaving all other registers unchanged.
- [6]
- d) Determine the changes to registers and/or memory locations resulting from each instruction in the ARM interrupt service routine (ISR) shown in Figure 1.1. Explain the use of r13 and hence the overall function of the code. Write an optimised version of this code for use with the ARM Fast Interrupt (FIQ).
- [10]
- e) Trace through execution of code fragment COPYLOOP in Figure 1.2 when the initial value of r2 is 9, giving the sequence of conditionally executing instructions and stating in each case whether the instruction condition is true. You may assume that, when the instruction condition is true, branch instructions take 2 machine cycles, multiple register load and store instructions with  $n$  registers in their masks take  $n + 1$  machine cycles, and all other instructions (including any with instruction condition false) take 1 machine cycle. Determine the number of machine cycles taken by the COPYLOOP loop when it iterates  $n$  times ( $n > 1$ ). Hence or otherwise determine the asymptotic limit for large  $n$  of the ratio of number of memory bytes written to the number of machine cycles used by this code fragment.
- [10]

```

STR r0, [r13], #1
LDR r0, TIME
ADD r0, r0, #1
STR r0, TIME
LDR r0, [r13, #-1]!
SUBS pc, r14, #4
TIME % 4

```

Figure 1.1

```

COPYLOOP SUBS r2, r2, #8
A LDMLIA r0!, {r3-r10}
B STMLIA r1!, {r3-r10}
C BPL COPYLOOP

```

Figure 1.2

2. The subroutine PARITY is shown in Figure 2.1.

a) Which registers are changed as the result of execution of PARITY? Indicate how, by the addition of appropriate instructions, the registers r0 - r13 can be preserved across the subroutine call. Specify precisely where in the PARITY code the additional instructions must be inserted, and what these are. [15]

b) The code between PA and PB calculates a result  $r4 = f(r3)$ . Describe the function  $f$ . [15]

c) State precisely in what way PARITY modifies the contents of memory, using where necessary the function  $f$  from part b. [15]

d) Write an alternative to the code between PA and PB using a lookup from a 256 byte constant table in memory to speed up execution. You need not provide code to initialise the table, or a definition of the table in assembler, but must specify precisely the table's contents. [15]

```

PARITY ADR r0, BUFFER
 MOV r1, #0
 MOV r2, #0
 MOV r7, #0
L1 LDRB r3, [r0,r1]
 EOR r7, r7, r3
PA
 MOV r4, #0
 MOV r5, r3, lsl #24
L2 MOVS r5, r5, lsl #1
 ADDMI r4, r4, #1
 BNE L2
 AND r4, r4, #1
PB
 AND r3, r3, #&7f
 ORR r3, r3, r4, lsl #7
 EOR r7, r7, r3
 STRB r3, [r0,r1]
 ADD r1, r1, #1
 CMP r1, #128
 BNE L1
 STRB r7, [r0,r1]
 MOV pc, r14

BUFFER % 129 ;129 byte memory buffer

```

Figure 2.1

- a) Write paragraphs of no more than 50 words which answer each of the following questions:
- (i) What is an Instruction Set Architecture? [3]
  - (i) How do instructions in the ARM Instruction Set Architecture support the implementation of stacks. [4]
  - (ii) How does the mechanism of shadow registers in the ARM Instruction Set Architecture enable transparent handling of IRQ mode interrupts. [3]
- b) Draw a diagram illustrating how the bits of r0, r1 and the Carry status bit change after the execution of the ARM code fragment in Figure 3.1. Write in ARM assembler a subroutine REVERSE, using an ascending stack in which the register r13 points to the highest memory word used by the stack. On exit register r0 must be set equal to its value on subroutine entry but with bits reversed, so that bit  $n$  becomes bit  $31 - n$  and vice versa. All other registers are unchanged. [20]
- c) A CPU with 8 bit data bus uses a write-back direct access data cache to speed up its access to main memory. The cache contains only 1 line of length 4 bytes. The data memory usage of a program consists of 5 byte pushes, followed by 5 byte pops, on an ascending stack, where the first stack location written has address hexadecimal 0x100. You may assume that instructions are fetched from a separate instruction memory, and that this process does not affect the data cache.
- (i) What is the total size, in bytes, of the data cache? Assuming that the CPU address bus is 20 bits in length, determine the (possibly empty) sets of bits in the CPU address corresponding to the cache tag, index, and select fields. [8]
  - (ii) Determine the sequence of data read and write addresses issued by the CPU to the data cache during the program execution. [4]
  - (iii) Assume initially that valid bits for all data cache lines are false. Trace through sequence of CPU data operations specified in part (ii) indicating the sequence of main memory reads and writes that are required, and the state of data cache valid and dirty bits after each CPU data operation. [18]

```
MOV r0, r0, rrx #1
ADCS r1, r1, r1
```

Figure 3.1

**Question 1**

- a) Determine the numbers represented by the hexadecimal bit-pattern  $0x90003001$  when interpreted as specified.
- (i)  $0xc4e2$ , 16 bit unsigned
  - (ii)  $0xfde2$ , 16 bit signed
  - (iii)  $0x90600000$ , IEEE 754 floating point.
- (i) 50402  
(ii) -542  
(iii)  $-2^{32-127} * 1.75 = -4.42 * 10^{-29}$
- b)
- ```
ADD r10, r9, #5
SUB r0, r10
SUBMI r0, r2, r3
ADDPL r0, r4, r5
```
- c)
- ```
ADD r1, r0, r0, lsl 4
RSB r1, r1, r1, lsl 7
```
- d) The first instruction stores  $r0$  on the stack – notice this is the IRQ stack since the interrupt will have swapped to a shadow  $r13$ . The next three instructions load  $r0$  from a word in memory, increment it, and save to memory. The penultimate instruction restores the old value of  $r0$ . The final instruction returns from the interrupt swapping back to user registers.
- Note that the save & load are needed because  $r0$  is not shadowed. In an FIQ there are extra shadow registers available. One of these ( $r8-r12$ ) could be used instead of  $r0$  in which case the  $r0$  save/restore would not be needed. In fact the value of TIME could be kept permanently in the FIQ register so reducing the FIQ to just 2 instructions.
- e)
- Execution if  $r2=9$ . Condition true unless otherwise specified.
- ```
COPYLOOP
A
B
C
COPYLOOP
A (condition false)
B (condition false)
C (condition false)
```
- Time taken is $1+9+9+2$ cycles for all iterations other than the last, or $1+1+1+1$ cycles for last iteration only. So $t = 21 * n - 17$. 8 words \Rightarrow 32 bytes are transferred each iteration, so asymptotically 32/21 bytes are transferred each cycle.

SOLUTIONS

Question 2

a)

r0,r1,r2,r3,r4,r5,r7 changed.

STMED r13!, {r0,r1,r2,r3,r4,r5,r7} ; insert before first instruction

LDMED r13!, {r0,r1,r2,r3,r4,r5,r7} ; insert before MOV pc, r14,
; or add r15 to register list & replace the MOV pc

b)

f calculates the *xor* of the bottom 7 bits of r3, and sets the bottom 7 bits of r4 = r3, and the top bit equal to the *xor* (the 7 bit parity).

c)

Bytes BUFFER – BUFFER+127 are modified as $mem[x] \leftarrow f(mem[x])$. Byte BUFFER+128 is set to the 8 bit bitwise xor of all the bytes from BUFFER to BUFFER+127.

d)

ADR r1, TABLE
LDR r4, [r1,r3]

Byte i of TABLE initialised to contain $f(i)$

SOLUTIONS

Question 3

- a) An ISA is a precise definition of the operation of a CPU at the level of machine instructions which defines the CPU registers, the function, but not timing, of each instruction and also the function (but not timing) of interrupts and exceptions.

The ARM ISA provides multiple register load & store instructions, which contain a register mask allowing any subset of the 16 ARM registers to be loaded or stored from successive locations in memory as determined by another of the registers. The 4 types of stack: ascending/descending and SP full or empty, each have corresponding versions of these instructions: LDM/STM suffix EA,ED, FA, FD SP points to empty/full location, stack grows Ascending, descending. Furthermore, each LDM/STM instruction can either update the register used as stack pointer, or leave the SP unchanged as when accessing data on the stack.

IRQ mode shadows r14 & r13 & also has a register which saves the user CPSR. The IRQ mode r13 points to a separate IRQ stack allowing interrupt code to execute safely independent of user code. Other registers can be saved & restored using this stack as necessary, and the user instruction stream restored by restoring the CPSR from the IRQ SPSR, and the user pc from the IRQ r14.

b)

```
REVERSE      STMFA   r13!, {r1,r2}
              MOV     r2, #32
REV1         MOV     r0, r0, lsr #1
              MOV     r1, r1, lsl #1
              SUBS   r2, r2, #1
              BPL    REV1
              LDMFA  r13!, {r1,r2,r15}
```


SOLUTIONS

c)

(i) size 4 bytes, select A1:0, tag A19:2, index has no bits.

(ii) (all in hex) W100, W101, W102, W103, W104, R104, R103, R102, R101, R100.
(R=Read, W=Write)

(iii) There is only one cache line, it is always (after the first op) valid.

Data op	Valid	Dirty	Memory operation
W100	1	1	R100-103
W101	1	1	
W102	1	1	
W103	1	1	
W104	1	1	W:100-103 R:104-107
R104	1	1	
R103	1	0	W104-107 R:100-103
R102	1	0	
R101	1	0	
R100	1	0	

Paper Number(s): **E1.9B**

IMPERIAL COLLEGE LONDON

DEPARTMENT OF ELECTRICAL AND ELECTRONIC ENGINEERING
EXAMINATIONS 2005

ISE Part I: MEng, BEng and ACGI

PRINCIPLES OF COMPUTING AND SOFTWARE ENGINEERING (PART B)

OPERATING SYSTEMS

Wednesday 8th June 2005 2:00pm

Corrected Copy

There are TWO questions on this paper.

Answer ONE question.

Q. 2 c).

This exam is CLOSED BOOK

Time allowed: 1 hour

Any special instructions for invigilators and information for candidates are on page 1.

Examiners responsible:

First Marker(s): Demiris, Y. K.

Second Marker(s): Shanahan, M.P.

© University of London 2005

Answer **ONLY ONE** of the following two questions

QUESTION 1:

- (a) Describe the MFQS (Multilevel Feedback Queue Scheduling) scheduling algorithm and list its advantages and disadvantages [3]
- (b) Consider the following set of processes, with their corresponding duration, arrival times, and priority levels [*higher numbers indicate higher priority*]:


Process	Arrival time (ms)	Duration (ms)	Priority level
A	0	3	4
B	3	4	3
C	4	4	2
D	5	2	5
E	8	2	6

Show the order of execution (including timing information) of the processes if the scheduler implements the following scheduling algorithms:

- (i) Priority-scheduling without preemption [3]
- (ii) Priority-scheduling with preemption [3]
- (iii) Round-Robin with a time quantum of 4ms [3]

For each of the algorithms calculate the average waiting time, and the average turnaround time.

- (c) In the “dining philosophers” synchronization problem, five philosophers (P) need to eat, with only 5 chopsticks available to them, arranged as shown in the figure below. Eating requires two chopsticks; a philosopher can only pick one chopstick at a time, and only chopsticks located next to him/her.



Describe how can you ensure that a deadlock will not occur in this situation, and explain why your solution guarantees that. [4]

- (d) In the context of page replacement algorithms, describe the Optimal Page Replacement, and the Least Recently Used (LRU) page replacement algorithms and list their advantages and disadvantages [4]

QUESTION 2:

(a) Using semaphores, describe how can you enforce that the critical section of a process A will always be executed before the critical section of a process B. [3]

(b) In the context of a memory paging system, consider the following scenario:

- You have three available frames
- The reference string is 2-5-1-2-4-5-7-3-5-2

Starting with empty frame contents, show the sequence of frame contents after each request, and count the number of page faults for each of the following page replacement algorithms:

- (i) Optimal page replacement [4]
 (ii) LRU (Least Recently Used) page replacement [4]

4:20 pm
 G.C.

(c) A system has ~~15~~²³ instances of a resource type and there are currently four processes running; their maximum needs and their current allocation are shown in the table below

	Maximum requirements	Current allocation
Process A	12	3
Process B	9	4
Process C	10	3
Process D	12	7
Free: 6		

(i) Determine whether the current state is a safe state, or not, and in either case, demonstrate why. [2]

- (ii) Assume that the system is using the banker's algorithm for dynamic deadlock avoidance. Describe the algorithm's response for the following allocation requests [2]
 a. B requests 4 instances [2]
 b. C requests 3 instances [2]

(d) In the context of memory allocation and dynamic partitioning, describe the "first-fit", "best-fit", and "worst fit" methods of memory allocation, and list their advantages and disadvantages [3]

E1.9 – section B: Operating Systems
Model answers to exam questions 2005

Question 1

(a) [bookwork]


Ready processes are divided into several separate queues, each queue having a different priority associated with it. Queues with higher priority can be given more time-slices for each process. Different scheduling algorithms can be used for scheduling within queues. Processes are allowed to move between queues, depending on their CPU utilisation.

Advantages: Flexible, allows fine control of scheduling, feedback mechanism adds adaptability; this is the most general scheme you can have.

Disadvantages: It's the most complex scheme – you have to decide on several issues, including number of queues, scheduling algorithm for each queue, method for promoting or demoting a process, method for determining which queue a new process should enter; extensive tuning might be required to define the best scheduler.

(b) [new computed example]


Priority Scheduling without preemption



Average waiting time: $(0+0+7+2+1) / 5 = 2$ ms

Average turnaround time: $(3+4+11+4+3) / 5 = 5$ ms


Priority-scheduling with preemption



Average waiting time: $(0+4+7+0+0) / 5 = 2.2$ ms

Average turnaround time: $(3+8+11+2+2) / 5 = 5.2$ ms

Round Robin (4ms)



Average waiting time: $(0+0+3+6+5) / 5 = 2.8$ ms

Average turnaround time: $(3+4+8+8+7) / 5 = 6$ ms

(c) [Bookwork]

Impose ordering on chopsticks, assigning a unique number

(between 1 and 5) to each of them. A dining philosopher must acquire the lower number chopstick first before taking the higher one.

Deadlock now impossible since philosophers 1 and 5 will compete for one of the chopsticks, and whoever acquires it first, will eat, release the chopstick, allowing the other one to complete.

(d) [Bookwork]:

Optimal page replacement: replace page that will not be used for the longest period of time.

- Advantages: Lowest page-fault rate of all algorithms – can be used to evaluate relative performance of other page replacement algorithms.
- Disadvantages: Difficult to impossible to implement since we need to know the stream of requested pages in advance.

Least-recently used: replace the page that has not been used for the longest time

- Advantages: Good performance
- Disadvantages: Not the easiest to implement

Question 2:

(a) [bookwork]

Initialise a semaphore to 0, and require process B to wait on it until A signals that it is complete:

Init(Sem, 0)

Process A:

Critical region
Signal(sem);
End

Process B

....
wait(Sem);
critical region;
signal(Sem);
End

(b) [new computed example]

Optimal page replacement algorithm (6 page faults)

	2	5	1	2	4	5	7	3	5	2
Frame1	2	2	2		2		2	2		
Frame2	-	5	5		5		5	5		
Frame3	-	-	1		4		7	3		

LRU (Least recently used) page replacement algorithm (8 page faults)

	2	5	1	2	4	5	7	3	5	2
Frame1	2	2	2		2	5	5	5		5
Frame2	-	5	5		4	4	4	3		3
Frame3	-	-	1		1	1	7	7		2

(c) [New computed example]

(i) The current state is safe since there exist safe sequences through which resources can be allocated without the possibility of a deadlock: e.g D->C->B->A. (more exist)

(ii) (a) Request granted since, following the assignment of 4 resources to B, there still exist safe sequences (e.g. B->C->D->A)

(ii) (b) Request refused since following the assignment of 3 resources to C, we would have an unsafe state.

(d) [Bookwork]

First-fit: allocate first memory hole that is big enough. Advantages: fast allocation method. Disadvantages: can be very inefficient

Best-fit: allocate the smallest hole that is enough. Advantages: less inefficient in terms of space than first-fit. Disadvantages: tends to produce lots of remaining tiny fragments, and requires search through the entire list of memory holes

Worst fit: allocate the largest hole that is available. Advantages: after allocating the request, the remainder of that hole might still be usable. Disadvantages: requires search through the entire list of holes.