

Internal Assessment Test 2 – MAY
2025

Sub:	MICROCONTROLLERS					Sub Code:	BCS402	Branch:	CSE
Date:	26/05/25	Duration:	90 mins	Max Marks:	50	Sem / Sec:	IV Sem A/B/C		OBE

Answer any FIVE FULL Questions

MARKS

<u>Answer any FIVE FULL Questions</u>		MARKS	CO	RBT												
1. a)	<p>Explain Arm c compiler data types.</p> <p>ANS:</p> <p>Compilers <i>use</i> the datatype mappings which is shown in below table- C compiler datatype mappings.</p> <table><tr><td>C Data Type</td><td>Implementation</td></tr><tr><td>char</td><td>unsigned 8-bit byte</td></tr><tr><td>short</td><td>signed 16-bit halfword</td></tr><tr><td>int</td><td>signed 32-bit word</td></tr><tr><td>long</td><td>signed 32-bit word</td></tr><tr><td>long long</td><td>signed 64-bit double word</td></tr></table>	C Data Type	Implementation	char	unsigned 8-bit byte	short	signed 16-bit halfword	int	signed 32-bit word	long	signed 32-bit word	long long	signed 64-bit double word	[2]	CO2	L2
C Data Type	Implementation															
char	unsigned 8-bit byte															
short	signed 16-bit halfword															
int	signed 32-bit word															
long	signed 32-bit word															
long long	signed 64-bit double word															
1. b)	<p>Consider the following C code to calculate the Checksum of a data packet containing 64 words. Illustrate the compiler output generated for the same code shown below. Summarize the drawbacks of the compiler output.</p> <pre>short checksum_v3(short *data) { unsigned int i; short sum = 0; for (i = 0; i < 64; i++) { sum = (short)(sum + data[i]); } return sum; }</pre> <p>Ans:</p> <p>The compiler output generated for the same code are as follows—</p>	[8]	CO4	L2												

	<pre> MOV r2,r0 ; r2 = data MOV r0,#0 ; sum = 0 MOV r1,#0 ; i = 0 checksum_v3_loop ADD r3,r2,r1,LSL #1 ; r3 = &data[i] LDRH r3,[r3,#0] ; r3 = data[i] ADD r1,r1,#1 ; i++ CMP r1,#0x40 ; compare i, 64 ADD r0,r3,r0 ; r0 = sum + r3 MOV r0,r0,LSL #16 MOV r0,r0,ASR #16 ; sum = (short)r0 BCC checksum_v3_loop ; if (i<64) goto loop MOV pc,r14 ; return sum </pre> <p>The loop is now three instructions longer which are ADD r3,r2,r1,LSL #1 MOV r0,r0,LSL #16 MOV r0,r0,ASR #16</p> <p><u>There are two reasons for the extra instructions:</u> The LDRH instruction does not allow for a shifted address offset as the LDR instruction did in checksum_v2. Therefore the first ADD in the loop calculates the address of item i in the array. The LDRH loads from an address with no offset.</p> <p>The cast reducing total + array[i] to a short requires two MOV instructions. The compiler shifts left by 16 and then right by 16 to implement a 16-bit sign extend. The shift right is a sign-extending shift so it replicates the sign bit to fill the upper 16 bits.</p> <p>How to overcome the drawback We can avoid the second problem by using an int type variable to hold the partial sum. We only reduce the sum to a short type at the function exit.</p> <p>However, the first problem is a new issue. We can solve it by accessing the array by incrementing the pointer <i>data</i> rather than using an index as in data[i]. This is efficient regardless of array type size or element size. All ARM load and store instructions have a post increment addressing mode.</p>			
2 a)	<p>Briefly explain steps to enable IRQ and FIQ mode in ARM processor.</p> <p>Ans: IRQ and FIQ exceptions only occur when a specific interrupt mask is cleared in the <i>cpsr</i>.</p> <p>An IRQ or FIQ exception causes the processor hardware to go through a standard procedure (provided the interrupts are not masked):</p> <ol style="list-style-type: none"> 1. The processor changes to a specific interrupt request mode, which reflects the interrupt being raised. 2. The previous mode's <i>cpsr</i> is saved into the <i>spsr</i> of the new interrupt request mode. 3. The <i>pc</i> is saved in the <i>lr</i> of the new interrupt request mode. 4. <i>Interrupt/s are disabled</i>—either the IRQ or both IRQ and FIQ exceptions are disabled in the <i>cpsr</i>. This immediately stops another interrupt request of the same type being raised. 5. The processor branches to a specific entry in the vector table. <p>Following Table shows how IRQ and FIQ interrupts are enabled. The procedure uses three ARM instructions.</p> <p>The first instruction MRS copies the contents of the <i>cpsr</i> into register <i>r1</i>. The second instruction clears the IRQ or FIQ mask bit. The third instruction then copies the updated contents in register <i>r1</i> back into the <i>cpsr</i>, enabling the interrupt request. The postfix <i>_c</i> identifies that the bit field being updated is the control field bit [7:0] of the <i>cpsr</i>. The interrupt request is either enabled or disabled only once the MSR instruction has completed the execution stage of the pipeline. Interrupts can still be raised or masked prior to the MSR completing this stage.</p>	[6]	CO3	L3

Enabling an interrupt.

<i>cpsr</i> value	IRQ	FIQ
Pre	<i>nzcvqjIFt_SVC</i>	<i>nzcvqjIFt_SVC</i>
Code	enable_irq	enable_fiq
	MRS r1, cpsr	MRS r1, cpsr
	BIC r1, r1, #0x80	BIC r1, r1, #0x40
	MSR cpsr_c, r1	MSR cpsr_c, r1
Post	<i>nzcvqjiFt_SVC</i>	<i>nzcvqjiFt_SVC</i>

Disabling an interrupt.

<i>cpsr</i>	IRQ	FIQ
Pre	<i>nzcvqjift_SVC</i>	<i>nzcvqjift_SVC</i>
Code	disable_irq	disable_fiq
	MRS r1, cpsr	MRS r1, cpsr
	ORR r1, r1, #0x80	ORR r1, r1, #0x40
	MSR cpsr_c, r1	MSR cpsr_c, r1
Post	<i>nzcvqjiFt_SVC</i>	<i>nzcvqjiFt_SVC</i>

2.b) Explain full descending Stack with proper example.

[4]

CO2

L2

Ans:

The ARM architecture uses the load-store multiple instructions to carry out stack operations. The *pop* operation (removing data from a stack) uses a load multiple instruction; similarly, the *push* operation (placing data onto the stack) uses a store multiple instruction. When using a stack we have to decide whether the stack will grow up or down in memory. A stack is either *ascending* (A) or *descending* (D). Ascending stacks grow towards higher memory addresses; in contrast, descending stacks grow towards lower memory addresses. When we use a *full stack* (F), the stack pointer *sp* points to an address that is the last used or full location (i.e., *sp* points to the last item on the stack). In contrast, if we use an empty *stack* (E) the *sp* points to an address that is the first unused or empty location (i.e., it points after the last item on the stack).

ARM has specified an ARM-Thumb Procedure Call Standard (ATPCS) that defines how routines are called and how registers are allocated. In the ATPCS, stacks are defined as being full descending stacks. Thus, the LDMFD and STMFD instructions provide the pop and push functions, respectively.

The STMFD instruction pushes registers onto the stack, updating the *sp*. Figure 3.7 shows a push onto a full descending stack. We can see that when the stack grows the stack pointer points to the last full entry in the stack.

Example--

PRE

r1 = 0x00000002

r4 = 0x00000003

sp = 0x00080014

STMFD sp!, {r1,r4}

STMFD instruction—full stack push operation

PRE	Address	Data	POST	Address	Data
	0x80018	0x00000001		0x80018	0x00000001
<i>sp</i> →	0x80014	0x00000002		0x80014	0x00000002
	0x80010	Empty		0x80010	0x00000003
	0x8000c	Empty	<i>sp</i> →	0x8000c	0x00000002

POST

r1 = 0x00000002

r4 = 0x00000003
sp = 0x0008000c

3 a) With a neat diagram explain ARM processor exceptions and modes.

Ans:

Each exception causes the core to enter a specific mode. In addition, any of the ARM processor modes can be entered manually by changing the *cpsr*. *User* and *system* mode are the only two modes that are not entered by a corresponding exception.

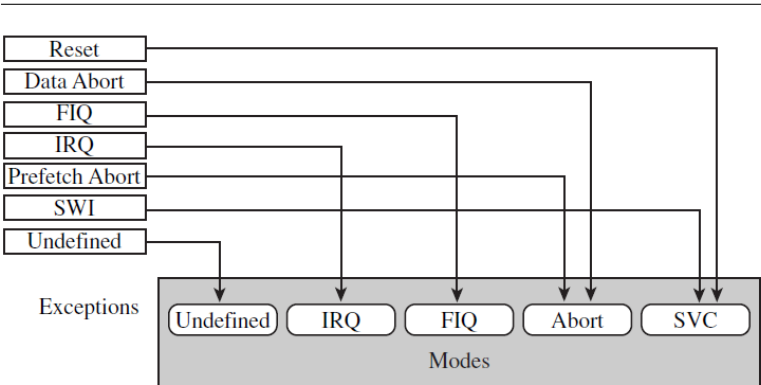
When an exception causes a mode change, the core automatically

- saves the *cpsr* to the *spsr* of the exception mode
- saves the *pc* to the *lr* of the exception mode

ARM processor exceptions and associated modes.

Exception	Mode	Main purpose
Fast Interrupt Request	<i>FIQ</i>	fast interrupt request handling
Interrupt Request	<i>IRQ</i>	interrupt request handling
SWI and Reset	<i>SVC</i>	protected mode for operating systems
Prefetch Abort and Data Abort	<i>abort</i>	virtual memory and/or memory protection handling
Undefined Instruction	<i>undefined</i>	software emulation of hardware coprocessors

Figure shows a simplified view of exceptions and associated modes.



- sets the *cpsr* to the exception mode
- sets *pc* to the address of the exception handler

3. b) What is interrupt latency and how software handler can minimize the interrupt latency.

Ans:

The interval of time from an external interrupt request signal being raised to the first fetch of an instruction of a specific interrupt service routine (ISR).

Interrupt latency depends on a combination of hardware and software. System architects must balance the system design to handle multiple simultaneous interrupt sources and minimize interrupt latency.

The first method is to use a nested interrupt handler

Nested interrupt handler– This allows further interrupts to occur even when currently servicing an existing interrupt. This is achieved by reenabling the interrupts as soon as the interrupt source has been serviced (so it won't generate more interrupts) but before the interrupt handling is complete. Once a nested interrupt has been serviced, then control is relinquished to the original interrupt service routine.

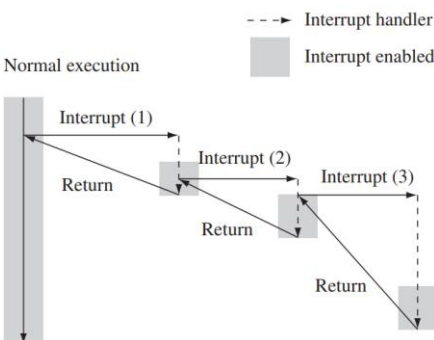


Figure 9.3 A three-level nested interrupt.

The second method involves prioritization-- In this case program the interrupt controller

	<p>to ignore interrupts of the same or lower priority than the interrupt we are handling, so only a higher-priority task can interrupt the handler.</p> <p>The processor spends time in the lower-priority interrupts until a higher-priority interrupt occurs. Therefore higher-priority interrupts have a lower average interrupt latency than the lower-priority interrupts, which reduces latency by speeding up the completion time on the critical time-sensitive interrupts.</p>			
4 a)	<p>Write a C program for ARM micro controller to sort the numbers in ascending order using bubble sort</p> <p>Ans:</p> <pre>#include<lpc21xx.h> int main(void) { unsigned long int temp, arr[4]= {0x00000001, 0x00000002, 0x00000004, 0x00000003}; unsigned char i,j,n=4; for (i=0;i<n-1;i++) { for (j=0;j<n-1;j++) { if (arr[j]<arr[j+1]) { temp = arr[j]; arr[j]=arr[j+1]; arr[j+1]= temp; } } } }</pre>	[05]	CO3	L3
4 b)	<p>Define Pointer aliasing. Analyze the concept of pointer-aliasing by using the code given below.</p> <pre>void timers_v1(int *timer1, int *timer2, int *step) { *timer1 += *step; *timer2 += *step; }</pre> <p>Ans: Two pointers are said to <i>alias</i> when they point to the same address. If we write to one pointer, it will affect the value we read from the other pointer. The following function increments two timer values by a step amount:</p> <pre>void timers_v1(int *timer1, int *timer2, int *step) { *timer1 += *step; *timer2 += *step; }</pre> <p>This compiles to</p> <pre>timers_v1 LDR r3,[r0,#0] ; r3 = *timer1 LDR r12,[r2,#0] ; r12 = *step ADD r3,r3,r12 ; r3 += r12 STR r3,[r0,#0] ; *timer1 = r3 LDR r0,[r1,#0] ; r0 = *timer2 LDR r2,[r2,#0] ; r2 = *step ADD r0,r0,r2 ; r0 += r2 STR r0,[r1,#0] ; *timer2 = r0 MOV pc,r14 ; return</pre> <p>The compiler loads from step twice. Usually a compiler optimization called <i>common subexpression elimination</i> would kick in so that *step was only evaluated once, and the value reused for the second occurrence. However, the compiler can't use this optimization here. The pointers timer1 and step might alias one another. In other words, the compiler cannot be sure that the write to timer1 doesn't affect the read from step. In this case the second value of</p>	[05]	CO3	L4

	*step is different from the first and has the value *timer1.This forces the compiler to insert an extra load instruction.															
5. a)	<p>Describe the features of Red Hat Red Boot firmware tool.</p> <p>Ans:</p> <p>RedBoot is a firmware tool developed by Red Hat. It is provided under an open source license with no royalties or upfront fees. RedBoot is designed to execute on different CPUs (for instance, ARM, MIPS, SH, and so on). It provides both debug capability through GNU Debugger (GDB), as well as a bootloader. The RedBoot software core is based on a HAL.</p> <p><u>RedBoot supports these main features:</u></p> <p><i>Communication</i>—configuration is over serial or Ethernet. For serial, X-Modem protocol is used to communicate with the GNU Debugger (GDB). For Ethernet, TCP is used to communicate with GDB. RedBoot supports a range of network standards, such as <i>bootp</i>, <i>telnet</i>, and <i>ftp</i>.</p> <p><i>Flash ROM memory management</i>—provides a set of filing system routines that can download, update, and erase images in flash ROM. In addition, the images can either be compressed or uncompressed.</p> <p>■ <i>Full operating system support</i>—supports the loading and booting of Embedded Linux, Red Hat eCos, and many other popular operating systems. For Embedded Linux, RedBoot supports the ability to define parameters that are passed directly to the kernel upon booting</p>	[06]	CO4	L2												
5. b)	<p>Illustrate the steps in the execution flow of sandstone code structure.</p> <p>Ans: Sandstone consists of a single assembly file. The file structure is broken down into a number of steps, where each step corresponds to a stage in the execution flow of Sandstone</p> <p>Sandstone execution flow.</p> <table><thead><tr><th>Step</th><th>Description</th></tr></thead><tbody><tr><td>1</td><td>Take the Reset exception</td></tr><tr><td>2</td><td>Start initializing the hardware</td></tr><tr><td>3</td><td>Remap memory</td></tr><tr><td>4</td><td>Initialize communication hardware</td></tr><tr><td>5</td><td>Bootloader—copy payload and relinquish control</td></tr></tbody></table> <p>Step 1: Take the Reset Exception</p> <p>Execution begins with a Reset exception. Only the reset vector entry is required in the default vector table. It is the very first instruction executed. reset vector is used to move the execution flow to the second stage.</p> <p>Step 2: Start Initializing the Hardware</p> <p>The primary phase in initializing hardware is setting up system registers. These registers have to be set up before accessing the hardware. For example, the ARM Evaluator-7T has a seven-segment display, which we have chosen to be used as a feedback tool to indicate that the firmware is active. Before we can set up the segment display, we have to position the base address of the system registers to a known location</p> <p>Step 3: Remap Memory</p> <p>One of the major activities of hardware initialization is to set up the memory environment. Sandstone is designed to initialize SRAM and remap memory. This process occurs fairly early on in the initialization of the system.</p> <p>Step 4: Initialize Communication Hardware</p> <p>Communication initialization involves configuring a serial port and outputting a standard banner. The banner is used to show that the firmware is fully functional and memory has been successfully remapped.</p> <p>N.B—if mark is 10 write the the code for each step.</p>	Step	Description	1	Take the Reset exception	2	Start initializing the hardware	3	Remap memory	4	Initialize communication hardware	5	Bootloader—copy payload and relinquish control	[04]	CO4	L2
Step	Description															
1	Take the Reset exception															
2	Start initializing the hardware															
3	Remap memory															
4	Initialize communication hardware															
5	Bootloader—copy payload and relinquish control															
6.a)	<p>Explain the swap instruction with an example code.</p> <p>Ans: The swap instruction is a special case of a load-store instruction. It swaps the contents of memory with the contents of a register. This instruction is an <i>atomic operation</i>—it reads and writes a location in the same bus operation, preventing any other instruction from reading or writing to that location until it completes.</p>	[04]	CO2	L2												

Syntax: SWP{B}{<cond>} Rd,Rm,[Rn]

SWP	swap a word between memory and a register	$tmp = mem32[Rn]$ $mem32[Rn] = Rm$ $Rd = tmp$
SWPB	swap a byte between memory and a register	$tmp = mem8[Rn]$ $mem8[Rn] = Rm$ $Rd = tmp$

Example—

PRE $mem32[0x9000] = 0x12345678$
 $r0 = 0x00000000$
 $r1 = 0x11112222$
 $r2 = 0x00009000$

SWP $r0, r1, [r2]$

POST $mem32[0x9000] = 0x11112222$
 $r0 = 0x12345678$
 $r1 = 0x11112222$
 $r2 = 0x00009000$

6. b) Given:

PRE:

$r1 = 0x00000002$,

$r4 = 0x00000003$,

Show the stack content and register contents after execution of following instructions.

STMFD $sp!$, { $r1, r4$ } $sp = 0x00080014$.

STMED $sp!$, { $r1, r4$ } $sp = 0x00080010$.

Ans

This instruction:

1. Decrements sp by $4 * \text{number of registers} \rightarrow 2 \text{ registers} \Rightarrow 8 \text{ bytes}$.
2. Stores $r1$ and $r4$ at the new sp (in order: $r1$ first, then $r4$).
3. Updates the stack pointer (due to the !).

i) **New SP = $0x00080014 - 8 = 0x0008000C$**

Memory content after execution:

PRE	Address	Data	POST	Address	Data
	$0x80018$	$0x00000001$		$0x80018$	$0x00000001$
$sp \rightarrow$	$0x80014$	$0x00000002$		$0x80014$	$0x00000002$
	$0x80010$	Empty		$0x80010$	$0x00000003$
	$0x8000c$	Empty	$sp \rightarrow$	$0x8000c$	$0x00000002$

Registers after execution:

$sp = 0x0008000C$

$r1 = 0x00000002$

$r4 = 0x00000003$

ii) Now, **Initial $sp = 0x00080010$** (as given before this instruction).

Again:

- 2 registers $\rightarrow 8 \text{ bytes to be stored}$
- New SP = $0x00080010 - 8 = 0x00080008$

[06]

CO2

L2

- Values stored in order: r1, r4

Memory content after execution:

PRE	Address	Data	POST	Address	Data
	0x80018	0x00000001		0x80018	0x00000001
	0x80014	0x00000002		0x80014	0x00000002
<i>sp</i> →	0x80010	<i>Empty</i>		0x80010	0x00000003
	0x8000c	<i>Empty</i>		0x8000c	0x00000002
	0x80008	<i>Empty</i>	<i>sp</i> →	0x80008	<i>Empty</i>

Registers after execution:

sp = 0x00080008

r1 = 0x00000002

r4 = 0x00000003

Faculty Signature

CCI Signature

HOD Signature