



Australian
National
University

Mid-Semester Examination (First Semester) 2011

Introduction to Computer Systems

(COMP2300/COMP6300)

Writing Period: 1 hour duration

Study Period: 0 minutes duration

Permitted Materials: One A4 page with notes on both sides.

NO calculator permitted.

The questions are followed by labelled, framed blank panels into which your answers are to be written. An additional answer panel is provided (at the end of the paper) should you wish to use more space for an answer than is provided in the associated labelled panels. If you use the additional panel, be sure to indicate clearly the question and part to which it refers to.

The marking scheme will put a high value on clarity so, as a general guide, it is better to give fewer answers in a clear manner than to outline a greater number in a sketchy, half-answered fashion. The Appendix contains a table of powers of 2 in decimal and the rPeANUt spec.

Please write clearly – if we cannot read your writing you may lose marks!

Student Number (please write this one very very clearly):

Official use only:

Q1:

Q2:

Total:

Question 1 [15 marks]

Student number:

(a) Convert the decimal number -23 into an 8 bit 2-complement representation. Give your answer in hexadecimal.

Q1(a) [2 marks]

(b) Convert the octal number 10.34 into decimal.

Q1(b) [1 marks]

(c) The IEEE single-precision floating-point standard is: 1 bit sign, 8 bits exponent with a bias of 127, and the remaining 23 bits are the significand (mantissa). Convert the decimal number -10.4 into IEEE single-precision floating-point standard. Give your answer in binary.

Q1(c) [3 marks]

(d) The IEEE single-precision floating-point numbers can be grouped into three types of values: Normalized Values, Denormalized Values, and Special Values. When are these floating-point numbers in denormalized form? What is the purpose of denormalized values?

Q1(d) [3 marks]

(e) In c write a program that uses a for loop that prints (to the standard output stream) the values 1 to 10 on separate lines.

Q1(e) [4 marks]

(f) What would the below code print if it was compiled and run?

```
#include<stdio.h>
int foo(int *p) {
    int x = 0;
    while (*p != 0) {
        x += *p;
        p += *p;
    }
    return x;
}

int main() {
    int l[] = {1,3,1,-3,1,0};
    printf("%d\n",foo(l));
    printf("%d\n",foo(&l[2]));
    return 0;
}
```

Q1(f) [2 marks]

Question 2 [15 marks]

Student number:

(a) What are three characteristics that make rPeANUt a RISC machine?

Q2(a) [3 marks]

(b) Disassemble the program given in the image below. Exactly what will the program output to the terminal when it is run? (Hint 0x41 is the ASCII code for 'A')

address	data
0x0100	0xc0000000
0x0101	0xc0010001
0x0102	0xc002000a
0x0103	0xc0030043
0x0104	0xa4000107
0x0105	0xd130fff0
0x0106	0x10100000
0x0107	0x20240000
0x0108	0xa4340105
0x0109	0x00000000
0x010a	0x00000000

Q2(b) [5 marks]

(c) The below c code echoes the upper case version of alphabetic characters (the 'upper' method given below does not work with some none alphabetic characters - do not worry about this). Convert this c code into rPeANUt assembler code. Your solution must implement procedures for both 'upper' and 'getchar'. Also, use the conventional rPeANUt stack frame approach for these procedures.

```
#include<stdio.h>
int upper(int c) {
    return 0xDF & c;
}
int main() {
    while(1) {
        putchar(upper(getchar()));
    }
}
```

Q2(c) [7 marks]

Additional answers to Question

Appendix

Powers of 2 in decimal:

$2^{-5}=0.03125$	$2^{-4}=0.0625$	$2^{-3}=0.125$	$2^{-2}=0.25$	$2^{-1}=0.5$	$2^0=1$
$2^1=2$	$2^2=4$	$2^3=8$	$2^4=16$	$2^5=32$	$2^6=64$
$2^7=128$	$2^8=256$	$2^9=512$	$2^{10}=1024$	$2^{11}=2048$	$2^{12}=4096$
$2^{13}=8192$	$2^{14}=16384$	$2^{15}=32768$	$2^{16}=64436$.	.

Specification of the rPeANUt Computer and Assembler (v1.2)

Eric McCreath - 2011- Research School of Computer Science - ANU

Introduction: The rPeANUt Computer is a simple microprocessor which was created for teaching computer systems at the ANU. There is a Java implementation of a simulator along with an assembler. This document aims to precisely describe the microprocessor along with the assembler so students can develop code and in so doing gain an idea of what is involved in developing code for a real microprocessor.


Overview: The rPeANUt is a 32 bit microprocessor with: 16 bit addresses, 32 bit register, and memory that is addressable in words of 32 bits. So the total maximum amount of addressable memory is $2^{16} = 65523$ words or 262144 bytes (256k). Only the addresses 0x0000 to 0x7FFF are connected to actual memory. Address between 0x8000 and 0xFFFF are used for memory mapped IO (although only 3 of these addresses are actually used). When the microprocessor is reset the instruction pointer (IP) is set to 0x0100, so normally a program will be load at this point for execution. Addresses 0x0000 to 0x0FFF are reserved for the interrupt vector and other OS code. Also the last 960 words of actual memory is used for the frame buffer.

The microprocessor contains the following 32 bit registers:

- 8 generally purpose registers these may be used for storing either data or addresses. These are denoted R0, R1, ... R7.
- An instruction register (IR) - which holds the current instruction that is being executed.
- A status register (SR) - contains status information about the CPU. Bit 0 is used for integer overflow (OF), and bit 1 is used for interrupt mask (IM).
- Three constant registers called ONE, ZERO, and MONE. They contain the constants 1, 0, and -1 respectively.

The microprocessor contains the following 16 bit registers:

- A stack pointer (SP) - this points to the top of the stack and is used for method calls, method returns, and interrupts (SP is set to 0x7000 when the microprocessor is reset).
- A program counter (PC) - which contains the address of the next instruction to execute.



Note that the IR registers is not directly accessible via the instruction set. Although clearly the execution of instructions will effect this register.

The microprocessor also contains a control unit which sequences the movement of data around the CPU. The microprocessor goes through the follow execution cycle:

```
do {
    IR = mem[PC];
    PC = PC + 1;
    execute_instruction in IR;
    check for interrupts;
} while(!halt);
```

Instruction set

Instruction are all 1 word long (32 bits). Registers have the labels R0,R1,...R7, SP, SR, PC, ONE, ZERO, MONE and take a nibble (4 bits) in the machine code. The encoding of this nibble is: R0 is 0x0, R1 is 0x1, ... , R7 is 0x7, SP is 0x8, SR is 0x9, PC is 0xA, ONE is 0xB, ZERO is 0xC, and MONE is 0xD. Addresses and values take 16 bits of the 32 bit machine code instruction. Values are sign extended from 16 bits to 32 bits. The description in the table below assumes the word of the instruction has been loaded into the IR and the PC has been moved to point to the next instruction word.

Name	Assembly Instruction	Machine code	Description
addition	add <RS1> <RS2> <RD>	0x1<RS1><RS2><RD>0000	RD <- RS1 + RS2
subtraction	sub <RS1> <RS2> <RD>	0x2<RS1><RS2><RD>0000	RD <- RS1 - RS2
multiply	mult <RS1> <RS2> <RD>	0x3<RS1><RS2><RD>0000	RD <- RS1 * RS2
divide	div <RS1> <RS2> <RD>	0x4<RS1><RS2><RD>0000	RD <- RS1 / RS2
modulo	mod <RS1> <RS2> <RD>	0x5<RS1><RS2><RD>0000	RD <- RS1 % RS2
bit and	and <RS1> <RS2> <RD>	0x6<RS1><RS2><RD>0000	RD <- RS1 & RS2
bit or	or <RS1> <RS2> <RD>	0x7<RS1><RS2><RD>0000	RD <- RS1 RS2
bit xor	xor <RS1> <RS2> <RD>	0x8<RS1><RS2><RD>0000	RD <- RS1 ^ RS2
negate	neg <RS> <RD>	0xA0<RS><RD>0000	RD <- - RS
bit not	not <RS> <RD>	0xA1<RS><RD>0000	RD <- ~ RS
copy register	move <RS> <RD>	0xA2<RS><RD>0000	RD <- RS
call immediate	call <address>	0xA300<address>	SP <- SP +1 mem[SP] <- PC PC <- address
return from call	return	0xA3010000	PC <- mem[SP] SP <- SP-1
trap	trap	0xA3020000	SP <- SP +1 mem[SP] <- PC PC <- 0x0004 SR <- SR (1<<1)
jump	jump <address>	0xA400<address>	PC <- address
jump if zero	jumpz <R> <address>	0xA41<R><address>	if (R == 0x00000000) {

Name	Assembly Instruction	Machine code	Description
			PC <- address }
jump if negative	jumpn <R> <address>	0xA42<R><address>	if ((R&0x80000000) == 0x00000000) { PC <- address }
jump if not zero	jumpnz <R> <address>	0xA43<R><address>	if (R != 0x00000000) { PC <- address }
reset status bit	reset <BIT>	0xA50<BIT>0000	SR <- SR & ~(1<<BIT)
set status bit	set <BIT>	0xA51<BIT>0000	SR <- SR (1<<BIT)
push onto stack	push <RS>	0xA60<RS>0000	SP <- SP + 1 mem[SP] <- RS
pop from stack	pop <RD>	0xA61<RD>0000	RD <- mem[SP] SP <- SP - 1
bit left rotate (using immediate)	rotate #<amount> <RS> <RD>	0xB0<RS><RD>00<amount>	RD <- RS << amount RS >>> (32 - amount)
bit left rotate (using register)	rotate <RA> <RS> <RD>	0xE<RA><RS><RD>00000	RD <- RS << RA RS >>> (32 - RA)
immediate load from memory	load #<label or value> <RD>	0xC00<RD><value>	RD <- ext(value)
absolute load from memory	load <address> <RD>	0xC10<RD><address>	RD <- mem[address]
indirect load from memory	load <RSA> <RD>	0xC2<RSA><RD>0000	RD <- mem[RSA]
base + displacement load from memory	load <RSA> #<value> <RD>	0xC3<RSA><RD><value>	RD <- mem[RSA + ext(value)]
absolute store to memory	store <RS> <address>	0xD1<RS>0<address>	mem[address] <- RS
indirect store to memory	store <RS> <RDA>	0xD2<RS><RDA>0000	mem[RDA] <- RS
base + displacement store to memory	store <RS> #<value> <RDA>	0xD3<RS><RDA><value>	mem[RDA+ext(value)] <- RS
halt	halt	0x00000000	fetch execution stops!

Hardware and Interrupts

A simple terminal is provided via memory mapped IO. Interacting with this device is done via three addresses: dataIO at 0xFFF0, status at 0xFFF1, and control at 0xFFF2. The least most significant bit of the status register (bit 0) is 1 when data is available and 0 otherwise. Bit 1 of the status register is 0 when it is ready to receive data and 1 if not ready. To write to this device simply write to the memory address of the dataIO location, and to read from this

device just read from the dataIO address. Note the status register should be checked prior to reading or writing to this device (although good practice to check the status bit before writing, in this simulator it will always be ready for writing, so you may just write directly to the dataIO location). The control address is used to set interrupts on for this device (bit 0 of the control address is 0 if interrupts are disabled and 1 if enabled). If the interrupt bit is set then when the a key is hit an interrupt is generated. Note interrupts are disabled by default.

The simulated computer has a black and white screen which is 192 pixels wide and 160 pixels high (or 0xC0 wide and 0xA0 high). The contents of this screen is determined by a frame buffer which starts at address 0x7C40 and ends at 0x7FFF. Assuming (0,0) is the top left corner of the screen then pixel (x,y) will be bit $x\%32$ of the word at address $0x7C40 + 6*y + x/32$. If this bit is 1 then the pixel is white and if the bit is 0 then the pixel is black.

When an interrupt occurs the current PC is pushed onto the stack and the PC is set to the address associated with that interrupt. Any registers used by the interrupt service routine must be saved and restored. The interrupt event also sets the interrupt mask high which should be cleared before the interrupt service routine finishes. The standard 'return' instruction is used to return from interrupts. The interrupts and their addresses are given below:

interrupt	address	description
memory fault	0x0000	This happens when memory is accessed that is not addressable.
IO device	0x0001	This happens when interrupts are enabled on the terminal device and a key is hit.
trap	0x0002	This interrupt happens when the trap instruction is executed.

rPeANUt Assembler

The rPeANUt assembler provides a simple way of converting assembly code into the rPeANUt machine code. It works using two phases. The first is a line by line translation into machine code. As this translation takes place both a symbol table is created and a list of addresses that need resolving. The second phase involves resolving all these missing addresses. Note the assembler writes directly into the hardware's simulated main memory (this is just for simplicity).

Each instruction must be placed on a single line. Lines with no instructions or labels are simply ignored. Any characters after the first ";" on a line are considered comments and ignored.

As the code is assembled instructions and data are placed into the next available address. The process starts at address 0x0000 and can only move forward. If you wish to skip forward to a new address location then you can simply place the address before a ":". You can not go backwards!

Address labels may consist of alpha numeric characters but must not start with a numeric character. They also must not be keywords (keywords are instructions, register names, and assembler directives). A single address can have multiple names, however, separate lines must be used to achieve this.

Instructions may be placed on a line by them self or after a “:”. Instructions have the instruction name followed its parameters. These are all space separated. The names of instructions are given in the table in the instruction section of this document. Please note the order of the instruction parameters is important.

The registers are denoted: R0, R1, R2, ... , R7, SP, SR, PC, ONE, ZERO, MONE. Addresses can be given either using an integer (given in base 10 or as hex number, hex numbers are prefixed with 0x) or simply the label. Immediate addresses or values are prefixed with the # symbol. The addressing mode of the load and store operations is determined by the list of parameters.

The “block” keyword is used to reserve a block of memory. If the block keyword then a positive integer k is given then k words are reserved (these are initialised to 0). If block followed by an immediate integer or character then one word is reserved and this word is initialised to the given value. If a string is given using #"<string>" (e.g. #”Hello World”) then a block of words is reserved and initialised to that string. A null terminator is used and the characters are stored in the least most significant byte of each word.
