

Computer Architecture

1 ISA (INSTRUCTION SET ARCHITECTURE) (10%)

ISA	Number of General registers	Number of FP registers
Intel X86 (80386)	8 (eax to esp)	8 (ST0 to ST7)
Intel X86-64	16	32 ZMM registers
Intel Itanium	128	128
ARMv7 (VFPv2)	16	16
ARMv8	31	32

The above table shows the number of registers provided by Intel x86, Intel/HP Itanium, and ARM ISA (Instruction Set Architecture).

- A. (1 point) Please list one or more ISAs that provide **no** registers at all.
- B. (5 points) Registers are usually used to hold temporaries. Modern compilers also use registers to hold frequently used local variables, pointers, and so on. Ideally, the more the registers, the more objects could be allocated to registers and memory references to those objects could be saved. The trend of ISA has clearly shown that new architectures tend to have more registers.
So is it a good idea to take a quantum leap to go for 1024 registers in an ISA. You must argue for your answer, just say yes or no will get no credits.
- C. (2 points) In each ISA's calling convention, the set of general purpose registers is usually divided into caller-save and callee-save two subsets. Which type of procedures could benefit most from such a register partition?
- D. (2 points) Most ISAs have separate register files: general purpose registers, floating point registers, and SIMD (e.g. Intel SSE/ZMM and ARM Neon) registers. Why do most modern ISAs provide different set of registers? Would it be better to use just one set of general purpose registers? For example, we could have instructions such like:

```
Add  R1, R2, R3    /* integer add */
FADD  R1, R2, R3    /* Floating point add */
PADD  R1, R2, R3    /* SIMD Add */
```

Although these are different type of instructions, they could use the same general purpose registers.

見背面

OPCODES, BASE CONVERSION, ASCII SYMBOLS


Table with 8 columns: MIPS opcode (31:26), MIPS func1 (5:0), MIPS func2 (5:0), Binary, Decimal, Hexa-decimal, ASCII Character, Decimal, Hexa-decimal, ASCII Character. It lists various MIPS instructions like sll, srl, sra, etc., and their corresponding binary, decimal, hex, and ASCII values.

(1) opcode(31:26) == 0
(2) opcode(31:26) == 17hex (11hex); if fmt(25:21) == 16hex (10hex), f = s (single); if fmt(25:21) == 17hex (11hex), f = d (double)

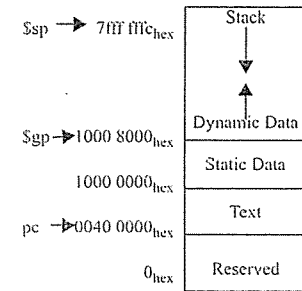
IEEE 754 FLOATING-POINT STANDARD

(-1)^S * (1 + Fraction) * 2^(Exponent - Bias)
where Single Precision Bias = 127, Double Precision Bias = 1023.


IEEE Single Precision and Double Precision Formats:




MEMORY ALLOCATION




STACK FRAME



DATA ALIGNMENT



EXCEPTION CONTROL REGISTERS: CAUSE AND STATUS



BD = Branch Delay, UM = User Mode, EL = Exception Level, IE = Interrupt Enable

EXCEPTION CODES

Table with 4 columns: Number, Name, Cause of Exception, and Cause of Exception. It lists various exception codes such as 0 (Int), 4 (AdEL), 5 (AdES), 6 (IBE), 7 (DBE), and 8 (Sys).

SIZE PREFIXES (10^x for Disk, Communication; 2^x for Memory)

Table with 7 columns: SIZE, PREFIX, SIZE, PREFIX, SIZE, PREFIX, SIZE, PREFIX. It lists size prefixes like Kilo, Mega, Giga, Tera, Peta, Exa, Zetta, Yotta, milli, micro, nano, pico, femto, atto, zepto, yocto.

The symbol for each prefix is just its first letter, except μ is used for micro.

Fig 1. MIPS Reference Data for Computer Architecture question 2.

2 DATA REPRESENTATION (7%)

On a MIPS machine (Fig. 1) running UNIX, we observed the following binary string stored in memory location x.

0011 1111 0111 0000 0000 0000 0000 0000

This binary could mean many different things,

- A. (1 point) If this is an integer number, what value is it?
- B. (1 point) If this is a single precision floating point number, what value is it?
- C. (1 point) If this is an instruction, what instruction is it?
- D. (1 point) If this is a C string, what string is it?
- E. (1 point) If this binary string was observed on your x86 desktop, what would be your answer for (D)?
- F. (2 points) Suppose we have three variables a, b and c. Give a case where $(a+b)+c$ computes a different value than $a + (b+c)$ on a MIPS microprocessor.

3 BRANCH PREDICTION (8%)

For pipelined processors, control hazards could significantly decrease the performance.

Dynamic branch prediction techniques have been successfully adopted in many modern processors to reduce performance penalty caused by control hazards. However, unlike direct branches (e.g. BEQ L1 in ARM or BEQ \$t0,\$t1, L1 in MIPS), indirect branches are usually difficult to predict.

- A. (2 points) Please give at least one static and one dynamic branch prediction scheme used for **direct branches** in microprocessors.
- B. (4 points) List conditions where indirect branches, instead of direct branches, are used? Which one of the listed cases is most frequent?
- C. (2 points) Please explain why indirect branches are hard to predict?

4 (10%)

Cache block size is an important design parameter for cache architecture. Assume a 1-CPI (Cycle-per-Instruction) machine with an average of 1.4 memory references (both instruction and data) per instruction. Assume the CPU stalls for cache misses.

Answer the following questions using the cache miss rates for different block sizes listed in the following table.

見背面

Block Size (Bytes)	8	16	32
Miss rate	8%	4%	3%

- A. (5 points) If the miss penalty is $24 + B$ (block size in bytes) cycles, what is the optimal cache block size? Please show how you derive the answer.
- B. (5 points) If critical-word-first is implemented in the cache, what is the optimal cache block size? Please show how you derive the answer.

5 (10%)

You are given a task to parallelize the following problem in a multi-core architecture:

```

for (i = 0; i < N; i = i+1)
    for (j = 0; j < N; j = j+1) {
        r = 0;
        for (k = 0; k < N; k = k+1)
            r = r + y[i][k]*z[k][j];
        x[i][j] = r;
    };
    
```

- A. (5 points) Is this a weak-scaling or strong-scaling problem? Please explain your answer.
- B. (5 points) Is it possible to partition this problem among cores such that there are no cache coherency misses? Please explain your answer.

6 (5%)

(5 points) In current disk storages, the operating system generally completes all I/Os asynchronously via interrupts. But recent studies show that for future NVM (Non-Volatile Memory) storage system, which has significantly lower access latency than disks, the synchronous approach (i.e., polling) could be more efficient than the interrupt approach. Please provide rationale behind this.

Operating System

7 (15%)

Consider short-term process scheduling. Please answer the following two problems:

- A. (10 points) The Round Robin scheduling algorithm (RR) lets every process run for a given time quantum and then switch to another process so as to have more fair sharing of the CPU time. Please give me your argument why a small time quantum is bad to the average turnaround time, where the turnaround time of a process is the difference between its completion time and its ready time of the process.
- B. (5 points) Why having processor affinity is good to processes in using data or instruction cache.

8 (10%)

(10 points) Consider a process synchronization problem, in which we have 5 people competing for 4 chairs, and each chair can be used by one person at a time. Let the competing go by rounds. That is, 5 people compete for 4 chairs in the beginning of each round, and the 4 winners release their chairs at the end of each round so that the next round starts. Please use semaphores to write programs for the entry and exit sections of the 5 people so that every person will win at least one time in competing the chairs for every 3 consecutive rounds.

9 (15%)

If the number of frames allocated to a low-priority process falls below the minimum number required by the computer architecture, we must suspend the execution of that process. We should then page out its remaining pages, freeing all its allocated frames. A process is thrashing if it is spending more time paging than execution.

- A. (3 points) For a system supporting virtual memory with unlimited swap space, how do we make a process to be thrashing?
- B. (3 points) If we do not allow a thrashing process to steal frames from other processes, will the other processes finish soon and the thrashing process eventually get enough frames to run? Why?
- C. (3 points) Why does working set model prevent thrashing?
- D. (3 points) Would working set model prevent thrashing totally? Why?
- E. (3 points) If processes in a virtual machine are thrashing, how does it affect the performance of the other virtual machines?

見背面

10 (10%)

Vectored I/O or scatter-gather I/O allows one system call to perform multiple I/O operations involving multiple locations.

- A. (5 points) Please describe the advantages using vectored I/O with a real application scenario.
- B. (5 points) When we implement a system call for vectored I/O, such as `readv` or `writen` in POSIX, what need to be considered to avoid the disadvantages using a vectored I/O.

試題隨卷繳回