

Computer Architecture I Final

Chinese Name: _____

Pinyin Name: _____

E-Mail ... @shanghaitech.edu.cn: _____

Question	Points	Score
1	1	
2	44	
3	18	
4	18	
5	15	
6	20	
7	19	
8	15	
Total:	150	

- This test contains 14 numbered pages, including the cover page, printed on both sides of the sheet.
- We will use gradescope for grading, so only answers filled in at the obvious places will be used.
- Use the provided blank paper for calculations and then copy your answer here.
- Please turn **off** all cell phones, smart-watches, and other mobile devices. Remove all hats and headphones. Put everything in your backpack. Place your backpacks, laptops and jackets out of reach.

- You have 120 minutes to complete this exam. The exam is closed book; no computers, phones, or calculators are allowed. You may use three A4 pages (front and back) of handwritten notes in addition to the provided green sheet (one of those can be printed).
- The estimated time needed for each of the 7 topics is given in parenthesis - it is 36 minutes for question 1 (Various Questions) and about 15 minutes for each of the 6 others. The total estimated time is 120 minutes.
- There may be partial credit for incomplete answers; write as much of the solution as you can. We will deduct points if your solution is far more complicated than necessary. When we provide a blank, please fit your answer within the space provided.
- Do **NOT** start reading the questions/ open the exam until we tell you so!
- Unless otherwise stated, always assume a 32 bit machine for this exam.

1. First Task (worth one point): Fill in you name

Fill in your name and email on the front page and your ShanghaiTech email on top of every page (without @shanghaitech.edu.cn) (so write your email in total 14 times).

2. Various Questions (36 minutes)

- 5 (a) This subquestion involves T / F questions. Incorrect answers on T / F questions are penalized with negative credit. Circle the correct answer.

T / F : ECC provides protection from disk failures.

T / F : A single parity bit allows us to detect any bit errors we have, but we need Hamming ECC (or something similar) to correct them.

T / F : All RAID configurations improve reliability.

T / F : All RAID configurations improve performance.

T / F : MapReduce is intended to run on a single, multi-core machine.

T / F : Exceptions in early pipeline stages override exceptions in later stages for a given instruction.

T / F : Exceptions are handled in the pipeline stage where they occur.

T / F : PUE does not measure the power efficiency of Warehouse Scale Computer servers.

T / F : Amdahl's law is restricting the speed increase predicted by Moore's Law.

T / F : We can implement the atomic test-and-set operation in MIPS using `lw` and `sw`.

- 3 (b) Consider attaching a hard disk to a CPU. Should we use polling or interrupts to deal with the hard disk? Should we use a DMA engine? Explain.

- 3 (c) What does AMAT stand for and how is it defined?

- 6 (d) For a single, running process, where in memory are the `heap`, `stack`, `code` and `static data` located? For each memory types, write if they can (typically) change in size and if so, how do they grow.

- 2 (e) Do different processes have (typically) access to each other's memory (Yes/ No). What is the name of the technique that is responsible for this?

(e) _____

- 3 (f) What does TLB stand for and what does it do?

(f) _____

- 3 (g) What are the advantages and disadvantages between using a static library and a dynamically linked library?

(g) _____

- 6 (h) Briefly explain what happens between switching on a computer and it being fully started up (4 keywords and very brief explanations for each).

(h) _____

- 2 (i) Which of the following can increase the availability?
- a. Increasing MTTF
 - b. Decreasing MTTF
 - c. Increasing MTTR
 - d. Decreasing MTTR
 - e. Redundant data copies

(i) _____

- 2 (j) Explain very briefly why RAID1 is the most expensive form of RAID.

(j) _____

- 2 (k) Represent the following decimal number using 2's complement:

$$-68(\text{dec}) = 0x\text{_____} = 0b\text{_____}$$

- 2 (l) Using IEEE 754 representation, what decimal number is encoded?

$$0xC2C00000 = \text{_____}$$

- 1 (m) We represent an unsigned integer. What decimal number is encoded? You may use the size prefixes for memory (MIPS Green Sheet).

$$0xC0000000 = \text{_____}$$

- 4 (n) Consider the truth table below. On the right side, first extract the "Sum of Products" from the table. Afterwards simplify it.

A	B	C	D	X
0	0	0	0	0
0	0	0	1	1
0	0	1	0	0
0	0	1	1	0
0	1	0	0	1
0	1	0	1	1
0	1	1	0	0
0	1	1	1	0
1	0	0	0	0
1	0	0	1	1
1	0	1	0	0
1	0	1	1	0
1	1	0	0	0
1	1	0	1	1
1	1	1	0	0
1	1	1	1	0

3. C Programming (15 minutes)

7

- (a) The following function should allocate space for a new string, copy the string from the passed argument into the new string, and convert every lower-case character in the new string into an upper-case character. Fill in the blanks and the body of the for() loop. You may **not** declare any new variable.

```
char* upcase(char* str){
    char* p;
    char* result;

    result =

    for(
    {

    }
    return result;
}
```

3

- (b) Please explain what a dangling pointer is:

3

- (c) Please correct the following program by writing the correct version on the right side:

```
swap( int* p1, int* p2 )
{
    int *p;
    *p = *p1;
    *p1 = *p2;
    *p2 = *p;
}
```

- 1 (d) The system program that combines separately compiled modules of a program into a form suitable for execution is called:
- A. Assembler
 - B. Loader
 - C. Linker
 - D. None of the Above

(d) _____

- 1 (e) Which flag would you put in a compilation command (e.g. gcc) to include debugging information?
- A. -o
 - B. d
 - C. g
 - D. debug

(e) _____

- 1 (f) At the end of the compiling stage, the symbol table contains the ___ of each symbol.
- A. relative address
 - B. absolute address
 - C. the stack segment beginning address
 - D. the global segment beginning address

(f) _____

- 1 (g) beq and bne instructions produce ____ and they _____.
- A. PC-relative addressing, never relocate
 - B. PC-relative addressing, always relocate
 - C. Absolute addressing, never relocate
 - D. Absolute addressing, always relocate

(g) _____

- 1 (h) j and jal instructions add symbols and ___ to ____.
- A. instruction addresses, the symbol table
 - B. symbol addresses, the symbol table
 - C. instruction addresses, the relocation table
 - D. symbol addresses, the relocation table


(h) _____

4. MIPS & C programming (12 minutes)

Consider a list with nodes defined in C as follows:

```
struct ListNode {
    char name[6];
    int code[3];
    struct ListNode* next;
};
```

The diagram below, not drawn to scale, gives an example of such a list.



- 5 (a) Assume that register \$a0 contains a pointer to the first node of the list. Write a MIPS assembly language function called `getty1` that returns (in \$v0) with the second integer in the second node in the list (with the list pictured above, this will load a 5 into \$v0). You don't need to do any error checking (assume node is existing). Adhere to the standard MIPS programming conventions but be as brief as possible!
- 6 (b) Again assume that register \$a0 contains a pointer to the first node of the list. Write a MIPS assembly language function called `getty2` that returns (in \$v1) the fourth character in the third node in the list (with the list pictured above, this will load 'n' into \$v1). Again, you don't need to do any error checking (assume node is existing and string long enough). Adhere to the standard MIPS programming conventions but be as brief as possible!

7


- (c) Read and understand the following MIPS assembly code. Then translate this function into C code. Your answers should be as concise as possible. The parameter x in register \$a0 is an unsigned int.

```
#####  
## BitCount  
## $a0 = x, $v0 = return value  
#####  
BitCount:  
    addi $sp, $sp, -8  
    sw $ra, 4($sp)  
    sw $s0, 0($sp)  
    add $v0, $0, $0  
    beq $a0, $0, end  
    andi $s0, $a0, 1  
    srl $a0, $a0, 1  
    jal BitCount  
    add $v0, $v0, $s0  
end:  
    lw $ra, 4($sp)  
    lw $s0, 0($sp)  
    addi $sp, $sp, 8  
    jr $ra
```


5. SDS (12 minutes)

5


(a) Consider the following circuit, the delay is 50 ns for the NOT gate and 40 ns for the rest of the logic gates. The register has 30 ns setup time, 30 ns clk-to-q delay and 10 ns hold time. Please calculate the maximum clock frequency for this circuit.



(a)

4



(b) Write the C expression of the following circuit (e.g. $output = (X \& Y)$).



(b)

6

(c) Draw the Timing Diagram for the circuit below. The delay for any logic is 10 ns, the setup time and hold time can be ignored. The clock-to-q delay for a register is also 10 ns. Each clock cycle is 30 ns (updated to 60 ns on the whiteboard), each grid in the following diagram is a unit of 10 ns (output is initially high).



(d) The following questions are based on the CPU we learned in class. First let's recall clocking methodology.

- The input signal to each state element must stabilize before each rising edge.
- Critical path: Longest delay path between state elements in the circuit.
- $t_{clk} \geq t_{clk-to-q} + t_{CL} + t_{setup}$, where t_{CL} is the critical path in the combinational logic.
- If we place registers in the critical path, we can shorten the period by reducing the amount of logic between registers.

The delays of circuit elements are given as follows:

Element	Register clk-to-q	Register Setup	MUX	ALU	Mem Read	Mem Write	RegFile Read	RegFile Setup
Parameter	$t_{clk-to-q}$	t_{setup}	t_{mux}	t_{ALU}	$t_{MEMread}$	$t_{MEMwrite}$	t_{RFread}	$t_{RFsetup}$
Delay(ps)	30	20	25	200	250	200	150	20

1 (e) What instruction exercises the critical path?

(e) _____

3 (f) What are the minimum clock cycle, t_{clk} ? Provide the formula/ details for partial credit.

(f) _____

2 (g) Why is a single cycle CPU inefficient?

(g) _____

2 (h) How can you improve its performance?

(h) _____

7. Parallel Programming (15 minutes)

Read the following C code for matrix addition.

```
void matadd (float * A, float * B, float * C, int n){
    for(int i = 0; i < n; i++){
        for(int j = 0; j < n; j++){
            A[j*n+i] = B[j*n+i] + C[j*n+i];
        }
    }
}
```

Assume no compiler optimizations are applied and big values for n . We can easily see that this piece of code is not optimal. As a computer science student, it is your job to optimize this code to improve efficiency. Answer the following questions:

- 2 (a) There is a very obvious error in one line that is slowing down the program a lot. What is this error? Explain why it is slowing down the program so much.

(a) _____

- 4 (b) Restructure the code using loop unrolling. Also correct the error you identified in the above task. Use 4 statements per iteration. You may assume that n is dividable by 4.

```
void matadd (float * A, float * B, float * C, int n){
```

- 4 (c) Optimize this code further with OpenMP (no explicit loop unrolling in this task!).

```
void matadd (float * A, float * B, float * C, int n){
```

- 5 (d) Use SSE intrinsics to utilize SIMD instructions on modern processors. (Do not use OpenMP here.) Hint: You might need the following intrinsics:

```

__m128 _mm_add_ps (__m128 a, __m128 b)
__m128 _mm_load_ps (float const* mem_addr)
void _mm_store_ps (float* mem_addr, __m128 a)

void matadd (float * A, float * B, float * C, int n){

```

- 2 (e) Amdal's Law: Suppose that we've developed a new method that can run 3 times faster in 60% of the instructions. What is the overall speedup?

- 2 (f) Amdal's Law: Suppose that we can run n times faster in 60% of the instructions. What is the maximum overall speedup that we can achieve?

8. Memory Access (15 minutes)

Consider a 32-bit physical memory space and a 32 KiB 4-way associative cache with LRU replacement. You are told that the hit rate of loop two is 9/16.

```

int ARRAY_SIZE = 32 * 1024;
int arr[ARRAY_SIZE]; // *arr is aligned to a cache block

/* loop one */
for (int i = 0; i < ARRAY_SIZE; i += 4) arr[i] = i;
/* loop two */
for (int i = ARRAY_SIZE - 4; i >= 0; i -= 4) arr[i+1] = arr[i]-1;

```

- 4 (a) Fill the number of bits in the tag, index and offset fields in the figure below.

Tag	Index	Offset

(b) It's Not My Fault

Consider the following OpenMP snippet:

```
int values[size];
#pragma omp parallel
{
    int i = omp_get_thread_num();
    int n = omp_get_num_threads();
    for(int j = i * (size / n); j < (i + 1) * (size / n); j++){
        values[j] = j;
    }
}
```

All cores share the same physical memory and we are running 2 threads. This is the sole process running. Each page is 1 KiB, and you have 2 pages of physical memory. The code snippet above starts at virtual address 0x400, and the values array starts at 0x800. The size, n, i, and j variables are all stored in registers. The functions `omp_get_thread_num` and `omp_get_num_threads` are stored in virtual addresses 0x440 to 0x480. The replacement policy for the page table is Least Recently Used.

At the start of the `pragma omp parallel` call the page table looks as follows:

Virtual Page Number	Valid	Dirty	Physical Page Number
0	0	0	0
1	1	0	0
2	1	1	1
3	0	0	1

3 (c) How many page faults will occur if `size=0x080`?

(c) _____

5 (d) What is the minimum number of page faults that will occur if `size=0x200`? And what is the maximum?

(d) _____

3 (e) How could you reduce the maximum page faults for part (b) (updated to (d))? Choose the valid options amongst the following (- 1 pt for every incorrect answer):

1. increase virtual address space
2. decrease number of threads
3. increase number of threads
4. use SIMD instructions
5. change page table replacement policy to Random
6. increase physical address space
7. add a 4-entry TLB

(e) _____