

CS 3330 Exam 3 Fall 2017

Name: _____ Computing ID: _____

Letters go in the boxes unless otherwise specified (e.g., for **C** 8 write “C” not “8”).

Write Letters clearly: if we are unsure of what you wrote you will get a zero on that problem.

Bubble and Pledge the exam or you will lose points.

Assume unless otherwise specified:

- little-endian 64-bit architecture
- `%rsp` points to the most recently pushed value, not to the next unused stack address.
- questions are single-selection unless identified as select-all

Variable Weight: point values per question are marked in square brackets.

Mark clarifications: If you need to clarify an answer, do so, and also add a ★ to the top right corner of your answer box.

.....

Question 1 [2 pt]: Consider the following C code:

```
int x[3] = {0x1234, 0x5678, 0xABCD};
```

If `ints` are four bytes and stored in little endian order, then the sixth byte of the array `x` is:

- A** there is not enough information to answer
- B** `0x78`
- C** `0x00`
- D** `0xAB`
- E** `0x56`
- F** `0xCD`
- G** none of the above

Answer:

Question 2 [2 pt]: Consider a five-stage pipelined processor like the one described in lecture and the textbook. Including register delays, each of the stages requires the following amount of time to perform its computations:

- fetch — 100 picoseconds
- decode — 50 picoseconds
- execute — 150 picoseconds
- memory — 200 picoseconds
- writeback — 50 picoseconds

How many picoseconds after a program starts fetching its first instruction will it finish the writeback stage of its third instruction? Assume the pipeline does not need to stall. Write your answer as a base-10 number of picoseconds.

Answer:

Question 3 [2 pt]: Which of the following changes to a cache design would **decrease** the number of set index bits it has? Place a \checkmark next to each change that would decrease the number of set index bits. Leave other boxes blank.

- A** doubling the associativity without changing the block size or total cache size
- B** doubling the block size without changing the associativity or total cache size
- C** changing from a random replacement policy to an LRU replacement policy
- D** doubling the total cache size and the associativity without changing the block size

Information for questions 4–5

Suppose a processor has 32-bit addresses and a 32KB 4-way set associative unified L1 cache with 64B blocks; a write-allocate, write-back policy; and a random replacement policy.

Question 4 [2 pt]: (see above) How many block offset bits does this cache use? Write your answer as a base-10 number.

Answer:

Question 5 [2 pt]: (see above) How many set index bits does this cache use? Write your answer as a base-10 number.

Answer:

Question 6 [2 pt]: Suppose a processor has 32-bit virtual addresses, 36-bit physical addresses, and 8KB pages. If the processor uses a single-level page table with 4 byte page table entries, then how large is a page table on this processor?

- A 2^{21} bytes
- B 2^{13} bytes
- C 2^{19} bytes
- D 2^{25} bytes
- E it depends on what memory a program has allocated
- F 2^{23} bytes
- G none of the above

Answer:

Question 7 [2 pt]: Which of the following are ways using **multiple accumulators** is likely improve performance? Place a \checkmark next to each correct answer. Leave other boxes blank.

- A by making it easier for the processor to correctly predict the branches of a loop
- B by decreasing the number of instructions a program executes
- C by decreasing the number of pipeline stalls required due to data dependencies
- D by allowing more instructions to be executed in parallel with each other

Question 8 [2 pt]: Which of the following are true about kernel mode? Place a ✓ next to each true answer. Leave other boxes blank.

- A** kernel mode can be entered via a function call
- B** code running in kernel mode can directly access I/O devices
- C** code running in kernel mode may have its own dedicated regions of the address space
- D** kernel mode can be entered via an interrupt

Information for questions 9–11

For these questions, consider the following Y86-64 assembly snippet:

```
irmovq $0x4000, %rbx
irmovq $0x50, %rdx
mrmovq 0x3(%rbx), %rcx
subq %rcx, %rbx
addq %rdx, %rbx
```

Question 9 [2 pt]: (see above) On the five-stage pipelined processor design we discussed in class and in lecture with stalling and forwarding, how many cycles of stalling will be necessary to execute this assembly snippet? Write your answer as a base-10 number.

Answer:

Question 10 [2 pt]: (see above) Suppose we made a three-stage pipelined processor using the same components we used for our five-stage pipelined processor, such that the stages are:

1. Fetch and Decode
2. Execute
3. Memory and Writeback

Assume this processor implements all forwarding paths that will not require dramatically increasing the cycle time. In this processor, how many cycles of stalling will be required to execute the assembly snippet? Write your answer as a base-10 number.

Answer:

Question 11 [2 pt]: (see above) Suppose the three-stage pipelined processor (from the previous question) did not implement forwarding. How many cycles of stalling would be required to execute the assembly snippet then? Write your answer as a base-10 number.

Answer:

Question 12 [2 pt]: Which of the following are part of a process's *context* that an operating system stores during a context switch? Place a ✓ next to each correct answer. Leave other boxes blank.

- A** the value of the exception table base register
- B** the value of the page table base register
- C** the value of the program counter
- D** the value of the condition code SF

Question 13 [2 pt]: Suppose the variable **a** is declared as an `long*` and stored in the register `%rax` and the variable **b** is declared as an `long*` and stored in the register `%rbx`. Which of the following is a correct translation of the C code `*a += *b;` to x86-64 assembly?

- A** `leaq 8(%rbx,%rax), %rcx`
`movq %rcx, (%rbx)`
- B** `leaq (%rax,8), %rbx`
- C** `movq (%rax), %rcx`
`addq %rcx, %rbx`
`movq (%rbx), %rbx`
- D** `movq (%rax), %rcx`
`addq %rcx, (%rbx)`
- E** none of the above

Answer:

Information for questions 14–15

Consider a system with 16 KB (2^{14} byte) direct-mapped data cache with 64-byte blocks. In each of the below C code snippets, assume that:

- `array` is stored at an address that is a multiple of 2^{30}
- data caches are initially empty when starting the outer `for` loop
- only accesses to `array` use the data cache
- $512 = 2^9$, $128 = 2^7$, $1024 = 2^{10}$, $64 = 2^6$

Question 14 [2 pt]: (see above) Consider the following C code:

```
unsigned char array[1024 * 64];
...
for (int j = 0; j < 10; j++)
    for (int i = 0; i < 128; i++)
        array[i * 512] += 1;
```

How many data cache misses will the above nested loops experience on this system?

A $416 = 10 \times 32 + 96 = 10 \times 128/4 + 3 \times 128/4$

B $81920 = 10 \times 128 \times 64$

C $1280 = 10 \times 128$

D 128

E $992 = 32 + 10 \times 96 = 128/4 + 10 \times 3 \times 128/4$

F none of the above

Answer:

Question 15 [2 pt]: (see above) Consider the following C code:

```
unsigned char array[1024 * 64];
...
for (int j = 0; j < 10; j++)
    for (int i = 0; i < 128; i++)
        array[i * 512] += array[i * 128 + 64];
```

How many data cache misses will the above nested loops experience on this system?

A $2560 = 10 \times 128 \times 2$

B $163840 = 2 \times 10 \times 128 \times 64$

C $2240 = 10 \times 128 \times 2 - 10 \times 128/4$

D $1408 = 10 \times 128 + 128$

E $256 = 2 \times 128$

F none of the above

Answer:

Question 16 [2 pt]: On a system that implements swapping, when a program runs a `mov` instruction that tries to access a memory location that has not been loaded from disk, which of the following may happen? Place a \checkmark next to each correct answer. Leave other boxes blank.

- A the memory management unit of the processor tells the disk to start reading
- B the processor looks up the location of the page fault handler in the exception table
- C the operating system eventually restarts the program starting from the instruction **after** the `mov`
- D the operating system can context switch to another program while the disk read is happening

Question 17 [2 pt]: Which of the following statements about loop unrolling **are true**? Place a \checkmark next to each true statement. Leave other boxes blank.

- A loop unrolling can be combined with cache blocking
- B loop unrolling will not improve performance on a single-cycle processor since there is no way for multiple instructions to execute in parallel
- C compilers have trouble performing loop unrolling themselves because it requires analyzing code from multiple source files
- D too much loop unrolling is likely to cause register spilling, decreasing overall performance

Question 18 [2 pt]: Which of the following statements are true about vector instructions like the SSE extensions the x86-64 instruction set? Place a \checkmark next to each true statement. Leave other boxes blank.

- A code with more data dependencies is more likely to benefit from vector instructions
- B code with more conditionals is more likely to benefit from vector instructions
- C implementing vector instructions complicates the control path of the processor more than implementing out-of-order execution
- D code with more loops over large arrays is more likely to benefit from vector instructions

Information for questions 19–20

Suppose an out-of-order processor has the following execution units (also known as functional units):

- four pipelined multipliers with four cycle latency
- four pipelined adders with two cycle latency

The following questions ask about the amount of time required perform the computations specified by an assembly snippet with these execution units, ignoring any overheads of fetching, decoding, etc. the instructions involved. Assume the overhead of forwarding values between functional units and any other costs to coordinating the computations are negligible.

Question 19 [2 pt]: (see above) Consider the following x86-64 assembly snippet:

```
mulq %r8, %r9
mulq %r9, %rax
mulq %r8, %rbx
```

How long will the computations in this assembly snippet take to perform? Write your answer as a base-10 number of cycles.

Answer:

Question 20 [2 pt]: (see above) Consider the following x86-64 assembly snippet:

```
mulq %r8, %r9
mulq %r10, %r11
addq %r9, %r11
```

How long will the computations in this assembly snippet take to perform? Write your answer as a base-10 number of cycles.

Answer:

Question 21 [2 pt]: Suppose a processor lacked an explicit system call instruction. Which of the following would be an acceptable substitute for this instruction, allowing the operating system running on that processor to provide the same functionality without giving all programs unlimited access to the machine?

- A** have programs perform a deliberate access to an invalid virtual page which is not used for anything else in place of the system call instruction
- B** have programs write the system call arguments to a file instead of using the system call instruction
- C** have programs include the entire operating system code as a library and make a normal function call
- D** have the operating system use a signal handler in place of the system call trap handler

Answer:

Question 22 [2 pt]: When an exception occurs, the *processor* _____. Place a ✓ next to each correct answer. Leave other boxes blank.

- A** performs a context switch to another program
- B** jumps to an address specified by the operating system
- C** starts running in kernel mode
- D** writes the current program counter to the exception table

Question 23 [2 pt]: Which of the following features are **more characteristic of a CISC** (complex instruction set computer) instruction set architecture than a RISC (reduced instruction set computer) architecture? Place a ✓ next to each true answer. Leave other boxes blank.

- A** variable width instructions
- B** more addressing modes
- C** simpler instructions
- D** an instruction that increments a value in memory

Question 24 [2 pt]: In the single-cycle processor design described in our textbook, when the processor executes the `ret` instruction, the data memory output is

- A** used to read the fields of the `ret` instruction itself
- B** used to restore saved registers
- C** used to compute the new stack pointer value
- D** not used
- E** used to set the program counter

Answer:

Information for questions 25–26

Consider a system with 32-bit virtual addresses, 1KB (2^{10} byte) pages, and 34-bit physical addresses. Suppose this system uses 4-byte page table entries, with the following fields, from most significant to least significant bit (when the page table entry is interpreted as a 4-byte integer):

- a 24-bit physical page number
- 1 write-allowed bit
- 1 execute-allowed bit
- 1 kernel-only bit
- 4 unused bits
- 1 valid bit

Question 25 [2 pt]: (see above) If an `unsigned int addr` contains a virtual address on this system, which of the following C expressions would be equal to the page offset of that address? (Recall that `>>` does a logical shift on unsigned values.)

- A** `addr >> 8`
- B** `(addr >> 10) & 0xFFF`
- C** `addr >> 10`
- D** `addr & 0x3FF`
- E** `addr & 0xFF`
- F** `addr >> 34`
- G** none of the above

Answer:

Question 26 [2 pt]: (see above) If we have loaded a page table entry into an `unsigned int pte`, then which of the following C expressions would extract the physical page number from it? (Recall that `>>` does a logical shift on unsigned values.)

- A** `(pte >> 8) ^ 0xFFFFFFFF`
- B** `(pte & 0xFFFFFFFF) >> 8`
- C** `pte >> 24`
- D** `pte >> 8`
- E** `pte & 0xFFFFFFFF`
- F** `pte >> 10`
- G** none of the above

Answer:

Question 27 [2 pt]: Which of the following best describes the locality for this code fragment?

```
for (i = 0; i < 1000; i++)
    for (j = 0; j < 1000; j++)
        b[j * 1000 + i] = a[i * 1000 + j];
```

- A** good spatial locality for **b** and good temporal locality for **b**
- B** good spatial locality for **a** and good temporal locality for **b**
- C** good spatial locality for **b** and poor temporal locality for **a** and **b**
- D** poor spatial locality and temporal locality for both **a** and **b**
- E** good spatial locality for **a** and poor temporal locality for **a** and **b**
- F** poor spatial locality and temporal locality for **a**

Answer:

Question 28 [2 pt]: Consider the following three loops:

```
// Loop A
for (int i = 0; i < 128; ++i) {
    A[i] = (B[i] + 10) * (C[i] - D[i]);
}

// Loop B
for (int i = 1; i < 128; ++i) {
    A[i] = A[i-1] - A[i+1];
}

// Loop C
for (int i = 0; A[i] != 0; ++i) {
    A[i] += B[i * 10];
}
```

Which should be easiest to implement and improve the performance of with vector instructions like the SSE extensions to the x86-64 instruction set?

- A** Loop A
- B** Loop B
- C** Loop C
- D** none of the above; none can benefit from vector instructions

Answer:

Question 29 [2 pt]: Suppose a system has 4KB pages (2^{12} bytes), 39-bit virtual addresses, and three-level page tables, where page tables at each level are one page, containing 512 entries, and each page table entry is 8 bytes. If a program on this system has 3 MB ($2^{21} + 2^{20}$ bytes) of memory that it can access without triggering a fault, what is the minimum **size of its page tables**?

- A 2 pages or $2 \cdot 2^{12}$ bytes or less
- B 3 pages or $3 \cdot 2^{12}$ bytes
- C 4 pages or $4 \cdot 2^{12}$ bytes
- D 5 pages or $5 \cdot 2^{12}$ bytes
- E 6 pages or $6 \cdot 2^{12}$ bytes
- F 7 pages or $7 \cdot 2^{12}$ bytes
- G more than 7 pages or $7 \cdot 2^{12}$ bytes

Answer:

Question 30 [2 pt]: Which of the following are likely to result in a **fault**? (Consider the textbook's terminology for types of exceptions.) Place a \checkmark next to each true answer. Leave other boxes blank.

- A the user pressing control-C
- B accessing a memory location not in the process's address space
- C dividing by zero
- D attempting to execute an invalid instruction

Question 31 [2 pt]: Which of the following optimizations are likely to **increase** the number of L1 instruction cache and/or instruction TLB misses a program experiences? Place a \checkmark next to each correct answer. Leave other boxes blank.

- A function inlining
- B removing redundant calculations from a loop
- C loop unrolling
- D cache blocking

Question 32 [2 pt]: Consider the following C code:

```
int x[8] = {1, 2, 3, 4, 5, 6, 7, 8};
int *ptr;
ptr = x;
ptr += *ptr;
*ptr += *ptr;
ptr += *ptr;
```

What is the value of `*ptr` after this code executes?

- A** 3
- B** 4
- C** 7
- D** 5
- E** 6
- F** unknown; the code invokes undefined behavior
- G** none of the above

Answer:

Question 33 [2 pt]: Both signals and exceptions _____. Place a ✓ next to each true answer. Leave other boxes blank.

- A** must be used in order to access functionality that requires more privileges than normal user programs are allowed
- B** can be triggered by events external to the currently running program
- C** are features of the processor, primarily implemented in hardware
- D** require something outside the signal or exception handler code to save the old program counter

Question 34 [2 pt]: Suppose we wanted to add a new add instruction `addq3` to Y86-64 that took three arguments, two source registers and one destination register. For example,

```
addq3 %rax, %rbx, %rcx
```

would add `%rax` and `%rbx` and store the result in `%rcx`. Which of the following would be true about adding this instruction to our single-cycle Y86-64 processor? Place a \checkmark next to each true answer. Leave other boxes blank.

- A** While this instruction is executing, one of the data inputs to the register file would be equal to the data memory output.
- B** Implementing this instruction would require modifying the register file to support more inputs and/or outputs.
- C** `addq3` would require a register number to be placed in a different part of the machine code than register numbers are placed for any other Y86-64 instruction
- D** While this instruction is executing, one of the register number inputs to the register file would be equal to the ALU output.

Question 35 [2 pt]: Suppose a processor has 16KB (2^{14} byte) pages, 32-bit virtual addresses and 24-bit physical addresses. Which L1 cache design would allow the L1 cache access to start before the TLB accesses finishes, but still allow the L1 cache to ensure that a given physical address always maps to the same cache set? Place a \checkmark next to each correct answer. Leave other boxes blank.

- A** a 128 KB (2^{17} byte), 4-way set associative cache with a write-through policy and 32 byte blocks
- B** a 96 KB ($3 \cdot 2^{15}$ byte), 6-way set associative cache with a write-back policy and 128 byte blocks and an LRU replacement policy
- C** a 4 KB (2^{12} byte) direct-mapped cache with a write-through policy and 16 byte blocks
- D** a 16 KB (2^{14} byte) direct-mapped cache with a write-back policy and 64 byte blocks

Information for questions 36–38

For these questions, consider the following Y86-64 assembly snippet, running on the five-stage pipelined processor with forwarding and branch prediction discussed in lecture and our textbook.

```

    addq %rbx, %rax
    jle after
    rrmovq %rax, %rdx
after:
    subq %rax, %rdx
    xorq %rbx, %rdx

```

Recall that this processor implements branch prediction, in which it predicts all branches as taken, and on a misprediction, fetches the correct instruction when the conditional jump instruction executes its memory stage.

Question 36 [2 pt]: (see above) If **the jle is not taken** and the **addq** instruction performs its fetch stage in cycle 0, during what cycle will the **xorq** instruction perform its writeback stage? Write your answer as a base-10 number.

Answer:

Question 37 [2 pt]: (see above) If **the jle is taken** and the **addq** instruction performs its fetch stage in cycle 0, during what cycle will the **xorq** instruction perform its writeback stage? Write your answer as a base-10 number.

Answer:

Question 38 [2 pt]: (see above) If **the jle is not taken**, which of the following forwarding operations are needed to execute this snippet correctly? Place a \checkmark next to each needed operation. Leave other boxes blank.

- A** %rax from addq to subq
- B** %rax from addq to rrmovq
- C** %rdx from rrmovq to subq
- D** %rdx from rrmovq to xorq

Question 39 [2 pt]: In the single-cycle processor design described in our textbook, when the processor is executing the **pushq** instruction, the value of the output of the ALU is

- A** not used
- B** equal to the value of one the data inputs to the register file
- C** equal to the value of one of the register number inputs to the register file
- D** equal to the register number for **%rsp**
- E** none of the above

Answer:

Information for questions 40–43

Consider a system with:

- 64 KB (2^{16} byte) pages
- 28-bit physical page numbers
- 8 byte page table entries
- 36-bit virtual addresses
- two-level page tables, where first-level page tables are 4KB, and second-level tables are 16KB
- a 3-way, 384-entry TLB

Question 40 [2 pt]: (see above) Suppose the page table base register contains the byte address $0x10000$. What is the physical address of the first-level page table entry for $0x009231234$? Write your answer as a hexadecimal address.

Answer:

Question 41 [2 pt]: (see above) How much information (including metadata) is stored in the TLB?

- A** between 384 bytes and 1KB
- B** between 64KB and 21MB
- C** between 1KB and 6KB
- D** more than 21MB
- E** between 6KB and 64KB

Answer:

Question 42 [2 pt]: (see above) What is the size (in bits) of a physical address on this system? Write your answer as a base-10 number.

Answer:

Question 43 [2 pt]: (see above) When the processor is accessing the virtual address $0x012345678$, it will access the same TLB set as when it accesses _____. Place a \checkmark next to each correct answer. Leave other boxes blank. .

- A** $0x000005678$
- B** $0x5FFB45678$
- C** $0x012340000$
- D** $0x012FF5678$

.....
Pledge:

On my honor as a student, I have neither given nor received aid on this exam.

Your signature here