Name	

UW CSE 351, Summer 2013
Final Exam

9:40am - 10:40am, Friday, 23 August 2013

Instructions:

- Make sure that your exam is not missing any of the 11 pages, then write your full name and UW student ID on the front.
- **Read over the entire exam before starting to work on the problems!** The last page is a reference page that you may tear off to use during the exam; it does not have to be turned in.
- Do not spend too much time on a problem if there are other easy problems that you haven't solved yet.
- Feel free to use the backs of pages for working on problems, but *write your answers in the space provided below each problem*. If you make a mess, clearly indicate your final answer. Be sure to answer all parts of all questions.
- For free-form answers, keep you answers brief and feel free to use short phrases. ***Full sentences NOT required.***
- **No books, notes, or electronic devices may be used during the exam.** You may not communicate with other students during the exam, but please ask the instructor if you need clarification for some problem.
- Thanks for a fun quarter of 351! Cue the final exam music!

Part	Awarded	Possible
Part 1		**20**
Part 2		**32**
Part 3		**18**
Part 4		**30**
Total		**100**

1 Concepts (20 points)

Answer the following questions ***briefly***. ***Full sentences NOT required.***

(a) (5 points) List the two important illusions that the process abstraction provides to programs. For each illusion, list a mechanism involved in its implementation.

(b) (5 points) One purpose of virtual memory is to allow programs to use more memory than is available in the physical memory, by storing some parts on disk transparently. Name one *other* useful things that can be done with the virtual memory system.

(c) (5 points) Does a TLB (Translation Lookaside Buffer) miss always lead to a page fault? Why or why not?

(d) (5 points) Name one difference between Java references and C pointers.

2 Cache Flow (32 points)

You are interviewing for a lucrative position with Accelerated Throughput Memories, Inc. (ATM), a company known for its fast cache hardware. Answer the following questions to get the job.

ATM is evaluating two cache designs for machines with 32-bit physical addresses. The Fast Data and Instruction Cache (FDIC) and the Super-Efficient Cache (SEC) both use the conventional *write-back* and *write-allocate* policies and a true *least-recently-used* replacement policy. The (partial) geometry of these caches is as follows:

Name	Cache Size (bytes)	Block/Line Size (bytes)	Sets	Set Associativity
FDIC	1024	32		2
SEC		16	8	4

(a) (2 points) Fill in the number of sets in the FDIC cache and the size (in bytes) of the SEC cache in the table above. Use decimal notation. (The powers of 2 and exponent rules on page 11 may be useful.)

(b) (6 points) Write the number of tag bits, set index bits, and block/line offset bits for each cache in the table below. Recall the physical address size is 32 bits. (The powers of 2 and exponent rules on page 11 may be useful.)

Cache	Number of Bits		
	Tag	Set Index	Block/Line Offset
FDIC			
SEC			

(c) (4 points) Briefly, why does *write-back* often have better performance than *write-through*? ***Full sentences NOT required.***

Part 2 continues on the next page.

(d) You will calculate the data-cache ***miss rates*** of two code sections on each of the two caches.

- Consider data accesses only. Instructions are handled by a separate cache.
- There is only one level of data cache in the system.
- For each combination of code and cache, the cache starts empty.
- Both code sections use the following declarations:

```
int values[4][128];
int i, j;
```

- All data except arrays are stored in registers; accesses to them never affects the cache.
- The `values` array starts at address 0x0.
- The system's page size is 4096 bytes.
- `sizeof(int) == 4`

Showing calculations or explanations for your answers to the following four questions is ***NOT required****, but could help achieve partial credit if your answers are incorrect.*

(i) (8 points) What is the miss rate of the following code on the FDIC cache? on the SEC cache? Write your answers as fractions.

```
// code section A
int prod = 1;
for (i = 0; i < 8; i++) {
  for (j = 0; j < 128; j++) {
    prod = prod * values[i % 4][j];
  }
}
```

FDIC miss rate = ———— SEC miss rate = ————

(ii) (2 points) The three types of cache misses are cold/compulsory misses, conflict misses, and capacity misses. What type or types of cache misses occur in *both* caches when executing code section A?

Part 2 continues on the next page.

(iii) (8 points) What is the miss rate of the following code on the FDIC cache? on the SEC cache? Write your answers as fractions.

```
// code section B
for (i = 0; i < 128; i++) {
  values[0][i] = values[1][i] + values[2][i] + values[3][i];
}
```

FDIC miss rate = ———— SEC miss rate = ————

(iv) (2 points) The three types of cache misses are cold/compulsory misses, conflict misses, and capacity misses. What single type of cache misses accounts for the majority of misses in the FDIC cache when executing code section B?

3 Virtual Mystery (18 points)

You may detach this page for reference. You do not need to turn it in.

We just dug up a dusty old computing system that uses virtual memory, but we could not find much information about the paging system. We know from the ISA that the machine has 16-bit words and uses 16-bit virtual addresses and we discovered via some tinkering that it uses 14-bit physical addresses. *Note that* `sizeof(short) == 2` *bytes.*

To learn more about the virtual memory system, we ran the following program in a process P. It uses `valloc`, which is just like `malloc`, but allocates the payload to be page-aligned. Assume that all local variables are stored in registers; only the array is stored in memory.

```
short* numbers = (short*)valloc(1024 * sizeof(short));
short i;
for (i = 0; i < 1024; i++) {
  numbers[i] = i;
}
// examine physical memory here...
```

We verified that no disk accesses occurred during process P's execution, so we know all of the virtual pages P used must be mapped to physical pages. Just before the program finished, we used a special hardware tool to dump the contents of *physical* memory. Here is part of what we found. (Note these are all hexadecimal values!)

Partial Contents of Physical Memory							
Address	**Contents**	**Address**	**Contents**	**Address**	**Contents**	**Address**	**Contents**
0x0a00	0x0010	0x0a20	0x0000	0x0a40	0x000e	0x0a60	0x0040
0x0a02	0x0011	0x0a22	0x0001	0x0a42	0x00e0	0x0a62	0x0041
0x0a04	0x0012	0x0a24	0x0002	0x0a44	0x0e00	0x0a64	0x0042
0x0a06	0x0013	0x0a26	0x0003	0x0a46	0xe000	0x0a66	0x0043
0x0a08	0x0014	0x0a28	0x0004	0x0a48	0x0e00	0x0a68	0x0044
0x0a0a	0x0015	0x0a2a	0x0005	0x0a4a	0x00e0	0x0a6a	0x0045
0x0a0c	0x0016	0x0a2c	0x0006	0x0a4c	0x000e	0x0a6c	0x0046
0x0a0e	0x0017	0x0a2e	0x0007	0x0a4e	0x00e0	0x0a6e	0x0047
0x0a10	0x0018	0x0a30	0x0008	0x0a50	0x0e00	0x0a70	0x0048
0x0a12	0x0019	0x0a32	0x0009	0x0a52	0xe000	0x0a72	0x0049
0x0a14	0x001a	0x0a34	0x000a	0x0a54	0x0e00	0x0a74	0x004a
0x0a16	0x001b	0x0a36	0x000b	0x0a56	0x00e0	0x0a76	0x004b
0x0a18	0x001c	0x0a38	0x000c	0x0a58	0x000e	0x0a78	0x004c
0x0a1a	0x001d	0x0a3a	0x000d	0x0a5a	0x00e0	0x0a7a	0x004d
0x0a1c	0x001e	0x0a3c	0x000e	0x0a5c	0x0e00	0x0a7c	0x004e
0x0a1e	0x001f	0x0a3e	0x000f	0x0a5e	0xe000	0x0a7e	0x004f

Assume that, except for physical addresses `0x0a40-0x0a5f`, the physical memory shown above represents part of the `numbers` array in process P.

Part 3 continues on the next page.

(a) (6 points) Given these assumptions and the contents of physical memory, what is the ***largest*** page size the system could be using? *Briefly* explain your reasoning. ***Full sentences NOT required.*** *NOTE: questions (b) and (c) do not depend on this answer.*

(b) (4 points) Assume that the system actually uses 16-byte pages and an 8-way set-associative TLB with 32 total entries. Label the bits of a virtual address below as they are used to determine the following:

VPO Virtual Page Offset

VPN Virtual Page Number

TI TLB Index

TT TLB Tag

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

(c) (8 points) Assume 16-byte pages and assume the `numbers` array begins at *virtual* address 0x00004000. Using the contents of physical memory, reconstruct as many valid entries in process P's page table as you can. Write page table entries in any order, using the notation $VPN \rightarrow PPN$ to mean that the page table entry for virtual page number VPN is valid and maps to physical page number PPN. Write page numbers in hexadecimal.

4 Memory Allocation and Movement (30 points)

You may detach this page for reference. You do not need to turn it in.

Some garbage collectors for languages like Java move allocated objects from one address in memory to another address in memory to defragment the heap. To do this safely, they must update all references to moved objects to refer to the object's new address. Movement is generally unsafe in languages like C, because it is not always possible to determine exactly what values are pointers.

Suppose we are using a special dialect of C that makes it possible to find all pointers and move allocated blocks safely. We would like to extend a memory allocator like the one in Lab 5 with the ability to move allocated blocks in memory to reduce fragmentation.

Movement Algorithm Given an `old` allocated block to move and a `new` allocated block of sufficient size into which to move the `old`, we must (1) copy the payload of the `old` block into the `new` block, (2) replace all pointers that point to the `old` payload with pointers to the `new` payload, and, finally, (3) free the `old` block.

Your Tasks You will complete three small C functions to help find and replace pointers to the `old` payload with pointers to the `new` payload. **Minor coding mistakes will not affect your score since you do not have a compiler at hand. Use pseudocode if short on time.**

Provided Code and Environment You may use the following functions and definitions in your code. Most should be familiar from Lab 5. The machine has a 64-bit word and address size. Blocks are 8-byte aligned. When allocated, blocks contain an 8-byte header, storing their size and allocation tags, and a payload. You will only need to use the headers of blocks when they are free.

We added functions that return the `heap_start()` and `heap_end()` addresses. We also added the magic `holds_pointer` function, which determines whether or not a given memory location holds a pointer to some other memory location.

```
#define WORD_SIZE 8
// Unscaled pointer arithmetic (same as UNSCALED_POINTER_ADD ...)
#define P_ADD(x,y) ((void*)((char*)(x) + (char*)(y)))
#define P_SUB(x,y) ((void*)((char*)(x) - (char*)(y)))
#define TAG_USED 1
#define SIZE(x) ((x) & ~3)
struct BlockInfo {
  size_t sizeAndTags;
  // other fields not needed today
};
typedef struct BlockInfo BlockInfo;

// Returns the address of the first block in the heap.
BlockInfo* heap_start();
// Return the address of the special end word of the heap.
BlockInfo* heap_end();
// Return 1 if memory at addr holds a pointer, 0 otherwise.
int holds_pointer(void* addr);
```

Part 4 continues on the next page.

(a) (12 points) Implement a function `update_heap` that takes pointers to the `old` and `new` versions of the moved block and scans over the heap block by block from `heap_start()` to `heap_end()`, calling `update_block` on every *allocated* block to replace any `old` pointers with `new` pointers. Do not allocate or free any blocks. Do not change headers or footers. Do not use the free list. *Our solution adds 5 lines; 2 lines hold single curly braces.*

```
// Implemented in later question.
void update_block(BlockInfo* b, BlockInfo* old, BlockInfo* new);

void update_heap(BlockInfo* old, BlockInfo* new) {
  BlockInfo* b;   // A pointer for iterating over heap blocks.
  // -- YOUR CODE HERE ---------------------------------------




}
```

(b) (8 points) Implement a function `points_into` that takes a pointer, `ptr`, a starting address, `start`, and a length (in bytes), `nbytes`. Return 1 if `ptr` points to any address in the `nbytes`-long range of addresses starting at `start`. Return 0 otherwise. Assume `start` is at least `nbytes` before the end of the address space. *Our solution adds one (long) line.*

```
int points_into(void* ptr, void* start, size_t nbytes) {
  // -- YOUR CODE HERE -----------------------------




}
```

Part 4 continues on the next page.

(c) (10 points) Implement a function `update_block` that takes a pointer to a block, `b`, and pointers to blocks `old` and `new`, the source and destination of the movement operation. The function should update all pointers stored in the payload of `b` that point to the `old` payload so they point to the `new` payload. The provided code iterates over the payload of `b`, word by word. Your job is to update a word in the payload if it holds a pointer to an address in the payload of the `old` block (see `holds_pointer` and `points_into`), replacing it with a pointer to the corresponding location in the `new` block. Assume the `old` and `new` blocks are the same size. *Our solution adds 3 lines (4 lines when wrapped to fit).*

Example 1 If `old`'s payload starts at 0x8000 and `new`'s payload starts at 0xc000, and a pointer in the payload of `b` points to 0x8008, then this pointer should be replaced with a pointer to 0xc008.

Example 2

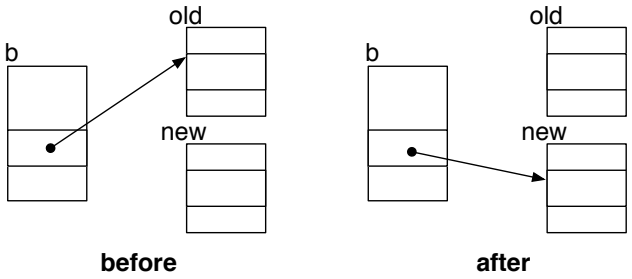


```
void update_block(BlockInfo* b, BlockInfo* old, BlockInfo* new) {
  // Start address and size of the old payload
  void*  oldPayloadStart = P_ADD(old, WORD_SIZE);
  size_t oldPayloadSize  = SIZE(*((size_t*)old)) - WORD_SIZE;
  void*  newPayloadStart = P_ADD(new, WORD_SIZE);
  void** slot;  //  pointer for iterating over payload words
  for (slot = P_ADD(b, WORD_SIZE);
       slot < P_ADD(b, SIZE(b->sizeAndTags));
       slot = P_ADD(slot, WORD_SIZE)) {
  // -- YOUR CODE HERE ------------------------------------
```

```
  }
}
```

5 Just for Fun

Write a pun using 351 material.

6 Reference

You may detach this page for reference. You do not need to turn it in.

Hexadecimal

Hex	Binary	Decimal
0x0	0000	0
0x1	0001	1
0x2	0010	2
0x3	0011	3
0x4	0100	4
0x5	0101	5
0x6	0110	6
0x7	0111	7
0x8	1000	8
0x9	1001	9
0xa	1010	10
0xb	1011	11
0xc	1100	12
0xd	1101	13
0xe	1110	14
0xf	1111	15

Powers of Two

2^0 = 0x0001 = 1	2^7 = 0x0080 = 128
2^1 = 0x0002 = 2	2^8 = 0x0100 = 256
2^2 = 0x0004 = 4	2^9 = 0x0200 = 512
2^3 = 0x0008 = 8	2^{10} = 0x0400 = 1024
2^4 = 0x0010 = 16	2^{11} = 0x0800 = 2048
2^5 = 0x0020 = 32	2^{12} = 0x1000 = 4096
2^6 = 0x0040 = 64	2^{13} = 0x2000 = 8192

$2^a * 2^b = 2^{a+b}$
$2^a / 2^b = 2^{a-b}$