

## CMU 15-418/618: Exam 1 Practice Exercises

### Problem 1: Miscellaneous Short Answer Questions

- A. Today directory-based cache coherence is widely adopted because snooping-based implementations often scale poorly. What is the main reason for this lack of scalability, and how do directory-based approaches avoid this problem?
- B. On your first day of work at Intel, you sit in a design meeting for the company's next quad-core processor, the Intel Core i-15418. Your boss immediately announces that like previous chips, this processor will use directory-based cache coherence using a full bit vector scheme. An engineer slams his notebook down on the table, and yells "What!?! We talked about several better ways to reduce the overhead of directories in 15-418! We should implement one of those!" The room goes silent, and the next day he is transferred to another group. Why was the boss unhappy with this suggestion. (Intel processors have 64-byte cache lines.)

C. Give one reason why a processor architect might decide to adopt a relaxed memory consistency model. (Hint: I'd like to see the term *latency* in your answer.)

D. You and your friend think the OpenMPI library is a bit crufty and decide to design a new message passing API for C++ to replace MPI. You've already designed and implemented new versions of SEND and RECEIVE to your API and now your friend suggests that you also add support for LOCKS in your library. Do you agree with this suggestion or do you argue that it is unnecessary mechanism in the message passing programming model? Why?

E. Recall that in the flat, *cache-based*, sparse directory scheme, the list of sharing processors is maintained in the caches as a doubly-linked list. If this list was maintained as a *singly-linked* (rather than doubly-linked) list, would you expect any significant impact on the latency of the following operations (please explain your answers):

Read misses:

Write misses:

Replacements (the line is evicted to make room for other data):

F. Imagine you are asked to implement ISPC, and your system must run a program that launches 1000 ISPC tasks. Give one reason why it is very likely more efficient to use a fixed-size pool of worker threads rather than create a pthread per task. Also specify how many pthreads you'd use in your worker pool when running on a quad-core, hyper-threaded Intel processor. Why?

G. Your friend suspects that his program is suffering from high communication overhead, so to overlap the sending of multiple messages, he tries to change his code to use asynchronous, non-blocking sends instead of synchronous, blocking sends. The result is this code (assume it's run by thread 1 in two-thread program).

```
float mydata[ARRAY_SIZE];
int dst_thread = 2;

update_data_1(mydata); // updates contents of mydata
async_send(dst_thread, mydata, sizeof(float) * ARRAY_SIZE);

update_data_2(mydata); // updates contents of mydata
async_send(dst_thread, mydata, sizeof(float) * ARRAY_SIZE);
```

Your friend runs to you to say “my program no longer gives the correct results.” What is his bug?

H. Complete the ISPC code below to write an if-then-else statement that causes an 8-wide SIMD processor to run at nearly 1/8th its peak rate. (Assume the ISPC gang size is 8. Pseudocode for an answer is fine.)

```
void my_ispc_func() {  
    int i = programIndex;  
  
  
  
  
  
  
  
  
  
}
```

I. Assume you want to efficiently run a program with very high temporal locality. If you could only choose one, would you add a data cache to your processor or add a significant amount of hardware multi-threading? Why?

- J. Consider the following OpenMP program running on a 4-core processor with infinite bandwidth and 0 memory latency. (Assume memory load and store operations are “free”.)

```
float total = 0.0;

#pragma omp parallel for
for (int i=0; i<N; i++) { // assume N is very, very large
    if (i % 16 < 8)      // assume 0 ops
        out[i] = 1 * in[i]; // 1 op
    else
        out[i] = 2 + in[i]; // 1 op
}

for (int i=0; i<N; i++)
    total += out[i];      // 1 op
```

What speedup will this program realize on a 4 core machine?

- K. Consider a program with a shared counter that is frequently written to by all threads, but rarely read (a stats counter is a good example). The program will run on a parallel system **with a large number of cores, that implemented invalidation-based coherence**. You’ve learned that directories help scaling coherence to high core counts, but your friend suggests that in this case you should design a processor with a snooping-based coherence implementation, claiming that broadcasting coherence messages to all cores is efficient since all cores need to manipulate the counter anyway. Is your friend correct? Why or why not? (e.g., would a directory-based protocol be preferable in this case?)

- L. In class we talked about the `barrier()` synchronization primitive. No thread proceeds past a barrier until all threads in the system have reached the barrier. (In other words, the call to `barrier()` will not return to the caller until its known that all threads have called `barrier()`). Consider implementing a barrier in the context of a message passing program that is only allowed to communicate via **blocking sends and receives**. Using only the helper functions defined below, implement a barrier. Your solution should make no assumptions about the number of threads in the system. **Keep in mind that all threads in a message passing program execute in their own address space—there are no shared variables.**

```
// send msg with id msgId and contents msgValue to thread dstThread
void blockingSend(int dstThread, int msgId, int value);

// recv message from srcThread. Upon return, msgId and msgValue are populated
void blockingRecv(int srcThread, int* msgId, int* msgValue);

// returns the id of the calling thread
int getThreadId();

// returns the number of threads in the program
int getNumThreads();
```

## Problem 2: Buying a New Computer

You write a bit of ISPC code that modifies a  $32 \times \text{height}$  grayscale image based on the contents of a black and white “mask” image of the same size. The code brightens input image pixels by a factor of 1000 if the corresponding pixel of the mask image is white (the mask has value 1.0) and by a factor of 10 otherwise.

The code partitions the image processing work into 64 ISPC tasks, which you can assume balance perfectly onto all available CPU processors.

```
void brighten_image(uniform int height, uniform float image[], uniform float mask_image[])
{
    uniform int NUM_TASKS = 64;
    uniform int rows_per_task = height / NUM_TASKS;
    launch[NUM_TASKS] brighten_chunk(rows_per_task, image, mask_image);
}

void brighten_chunk(uniform int rows_per_task, uniform float image[], uniform float mask_image[])
{
    // 'programCount' is the ISPC gang size.
    // 'programIndex' is a per-instance identifier between 0 and programCount-1.
    // 'taskIndex' is a per-task identifier between 0 and NUM_TASKS-1

    // compute starting image row for this task
    uniform int start_row = rows_per_task * taskIndex;

    // process all pixels in a chunk of rows
    for (uniform int j=start_row; j<start_row+rows_per_task; j++) {
        for (uniform int i=0; i<32; i+=programCount) {

            int idx = j*32 + i + programIndex;
            int iters = (mask_image[idx] == 1.f) ? 1000 : 10;

            float tmp = 0.f;
            for (int j=0; j<iters; j++)
                tmp += image[idx];

            image[idx] = tmp;
        }
    }
}
```

(question continued on next page)



You go to the store to buy a new CPU that runs this computation as fast as possible. On the shelf you see the following three CPUs on sale for the same price:

- (A) 4 GHz *single core* CPU capable of performing one floating point addition per clock (no parallelism)
- (B) 1 GHz *single core* CPU capable of performing one 32-wide SIMD floating point addition per clock
- (C) 1 GHz *dual core* CPU capable of performing one 4-wide SIMD floating point addition per clock




Figure 1: Image masks used to govern image manipulation by `brighten_image`

- A. If your only use of the CPU will be to run the above code as fast as possible, and assuming the code will execute using mask image 1 above, rank all three machines in order of performance. Please explain how you determined your ranking by comparing execution times on the various processors. When considering execution time, you may assume that (1) the only operations you need to account for are the floating-point additions in the innermost loop. (2) The ISPC gang size will be set to the SIMD width of the CPU. (3) There are no stalls during execution due to data access.
- (Hint: it may be easiest to consider the execution time of each row of the image.)

B. Rank all three machines in order of performance for mask image 2? Please justify your answer, but you are not required to perform detailed calculations like in part A.

### Problem 3: Buying a New Computer, Again

You plan to port the following sequential C++ code to ISPC so you can leverage the performance benefits of modern parallel processors.

```
float input[LARGE_NUMBER];
float output[LARGE_NUMBER];
// initialize input and output here ...

for (int i=0; i<LARGE_NUMBER; i++) {
    int iters;
    if (i % 16 == 0)
        iters = 256;
    else
        iters = 8;
    for (int j=0; j<iters; j++)
        output[i] += input[i];
}
```

Before sitting down to hack, you go the store, and see the following CPUs all for the same price:


- 4 GHz single core CPU capable of performing one floating point addition per clock (no parallelism)
- 1 GHz quad-core CPU capable of performing one 4-wide SIMD floating point addition per clock
- 1 Ghz dual-core CPU capable of performing one 16-wide SIMD floating point addition per clock

If your only use of the CPU will be to run your future ISPC port of the above code as fast as possible, which machine will provide the best performance for your money? Which machine will provide the least? Please explain why by comparing expected execution times on the various processors. When considering execution time, you may assume that (1) the only operations you need to account for are the floating-point additions in the innermost loop. (2) the ISPC gang size will be set to the SIMD width of the CPU.

(Hint: consider the execution time of groups of 16 elements of the input and output arrays).

#### Problem 4: Parallelizing a Video Processing Pipeline

Consider a video processing pipeline that is parallelized using two processors that communicate through shared memory. Processor 1 accepts as input a new video frame, and runs the function `boost_contrast` to increase the frame's contrast. It then puts the modified frame in a FIFO buffer large enough to hold 2 frames. Processor 2 then retrieves the modified frame from the buffer and performs face detection using the function `detect_kayvon`. It outputs a bit indicating whether Professor Kayvon is in the frame. This process is repeated for all video frames in a pipelined fashion (that is, after Processor 1 completes contrast detection on frame  $n$  and places its result in the intermediate buffer, it immediately begins working on frame  $n + 1$ ). The buffer is of finite size, so processor 1 blocks if the buffer becomes full. *All frames in the video are independent, so the results of processing a frame do not influence processing of any other frames.*



A. If contrast enhancement takes 1 second and face detection takes 3 seconds, what is the throughput of the pipeline in frames-per-second?

(Assume that costs to store/retrieve data from the buffer, or to synchronize buffer access between the two processors are negligible. Also, assume the input video is very long, containing millions of frames.)

B. You ask your friend to improve the performance of the video processing pipeline by changing the program's implementation. She smiles and says "Oh, that's easy! We just need to allocate a larger buffer to hold the intermediate results between processor 1 and 2." Do you agree with your friend? Why? If yes, what throughput do you expect to see as a result of her optimization?

## Problem 5: Angry Students

Your friend is developing a game that features a horde of angry students chasing after professors for making long exams. Simulating students is expensive, so your friend decides to parallelize the computation using one thread to compute and update the student's positions, and another thread to simulate the student's angriness. The state of the game's  $N$  students is stored in the global array `students` in the code below).

```
struct Student {
    float position; // assume 1D position for simplicity
    float angriness;
};

Student students[N];

////////////////////////////////////

void update_positions() {
    for (int i=0 i<N; i++) {
        students[i].position = compute_new_position(i);
    }
}

void update_angriness() {
    for (int i=0 i<N; i++) {
        students[i].angriness = compute_new_angriness(i);
    }
}

////////////////////////////////////

// ... initialize students here

pthread_t t0, t1;
pthread_create(&t0, NULL, updatePositions, NULL);
pthread_create(&t1, NULL, updateAngriness, NULL);
pthread_join(t0, NULL);
pthread_join(t1, NULL);
```

Questions are on the next page...

A. Since there is no synchronization between thread 0 and thread 1, your friend expects near a perfect  $2\times$  speedup when running on two-core processor that implements invalidation-based cache coherence. She is shocked when she doesn't obtain it. Why is this the case? (For this problem assume that there is sufficient bandwidth to keep two cores busy – “the code is bandwidth bound” is not an answer we are looking for.)

B. Modify the program to correct the performance problem. You are allowed to modify the code and data structures as you wish, **but you are not allowed to change what computations are performed by each thread and your solution should not substantially increase the amount of memory used by the program.** You only need to describe your solution in pseudocode (compilable code is not required).

## Problem 6: Oh, the Students Remain Angry

Due to the great success of the hit iPhone app “Angry Students”, Prof. Kayvon decides to release “Angry Students 2: They are Still Angry”, which uses ISPC to take advantage of the SIMD instructions on the iPhone’s ARM processor. The code is written like this:

```
struct Student {
    float position;
    float angriness;
};

Student students[N];

// ispc function
void updateStudents(int N, Student* students) {
    foreach (i = 0 ... N) {
        students[i].position = compute_new_position(i);
        students[i].angriness = compute_new_angriness(i);
    }
}
```

Performance is lower than expected, so Prof. Kayvon changes the code to this:

```
float positions[N];
float angriness[N];

// ispc function
void updateStudents(int N, float* positions, float* angriness) {
    foreach (i = 0 ... N) {
        position[i] = compute_new_position(i);
        angriness[i] = compute_new_angriness(i);
    }
}
```

The resulting code runs significantly faster. Why?

## Problem 7: Be An ISPC Compiler

- A. Please study the ISPC function `my_ispc_function` given below. The function multiplies all elements of the array 'input' by two.

```
// Recall ISPC's built-in variables:
// 'programCount' is the ISPC gang size
// 'programIndex' is a per-instance identifier between 0 and programCount-1;

void my_ispc_function(uniform int N, uniform float* input, uniform float* output) {

    // do not assume programCount divides N
    uniform int chunkSize = (N+programCount-1) / programCount;

    int start = programIndex * chunkSize;
    int end = start + chunkSize;

    if (end > N)
        end = N;

    for (int i=start; i<end; i++) {
        output[i] = 2 * input[i];
    }
}
```

Imagine you are implementing an ISPC compiler that translates ISPC programs into C programs that use vector intrinsics. To help, we have provided you a library of vector intrinsics similar to the library you used in Assignment 1. The library's methods are given on the next page. NOTE: YOU DO NOT need to study these functions in detail, but note that:

- (a) Although not listed assume that all arithmetic instructions (add, subtract, multiply, divide, etc. are present in the library for your use) and version are present for both vectors of floats and vectors of ints).
- (b) Assume that all binary operations on masks are present: and, or, equal
- (c) Just like Assignment 1, all operations can take an optional mask (1's in the mask mean the lane is ENABLED)
- (d) There are two types of vector load and store methods: **packed loads and stores** (loading consecutive elements) and **scatters and gathers** (loading non-consecutive elements).

On the next page, please translate this ISPC program into its vector equivalent `my_vector_function`. Your implementation can be pseudocode, but it should produce the same mapping of work to vector lanes as the real ISPC compiler implementation.



```

// Assume intVec and floatVect are vectors of ints and floats of size PROGRAM_COUNT
// Assume maskVec is a vector of bools: {111...111} = all lanes enabled

// ARITHMETIC EXAMPLES ////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////

intVec  initVec(int value, maskVec mask);           // set all elements of vector to 'value'
maskVec initMaskVec(bool value, maskVec mask);     // set mask to all ones or zeros
intVec  copyVec(intVec a, maskVec mask);           // result = a;

intVec  addVec(intVec a, intVec b, maskVec mask);   // add vectors (same for: 'mul', 'sub', 'div')

maskVec lessThanVec(intVec a, intVec b, maskVec mask); // result[i] = a[i] < b[i]
maskVec andVec(maskVec a, maskVec b, maskVec mask);  // a && b (same for: 'equal', 'or')
maskVec notVec(maskVec a, maskVec mask);            // !a

int     countOnesVec(maskVec v, maskVec mask);      // returns number of 1's in v

// LOAD/STORE, GATHER/SCATTER EXAMPLES ////////////////////////////////////////////////////////////////////

// load A[0], A[1], A[2], ..., A[PROGRAM_COUNT-1] into vector
intVec loadVec(int* A, maskVec mask);

// store v into A[0], A[1], A[2], ..., A[PROGRAM_COUNT-1]
void storeVec(int* A, intVec v, maskVec mask);

// load A[indices[0]], A[indices[1]], ..., A[indices[PROGRAM_COUNT-1]] into vector
intVec gatherVec(int* A, intVec indices, maskVec mask);

// store elements of v into A[indices[0]], A[indices[1]], ..., A[indices[PROGRAM_COUNT-1]]
void scatterVec(int* A, intVec indices, intVec v, maskVec mask);

// YOUR IMPLEMENTATION GOES HERE (WE'VE STARTED IT FOR YOU) ////////////////////////////////////////////

void my_vector_function(uniform int N, uniform float* input, uniform float* output) {

    intVec programIndex;           /* assume programIndex = {0, 1, 2, 3, ..., PROGRAM_COUNT-1}; */
    intVec programCount = initVec(PROGRAM_COUNT);
    intVec chunkSize      = divVec(vecAdd(initVec(N), vecAdd(programCount, initVec(-1))), programCount);

    intVec start          = mulVec(programIndex, chunkSize);
    intVec end            = addVec(start, chunkSize);
}

```

B. There is a performance problem with the current implementation of `my_ispc_function`. Please explain the problem and then re-write the original ISPC code to remove this problem. (Hint: it has to do with memory access.) For simplicity, your implementation can assume `programCount` divides `N` equally if you wish – though it would be cooler if it did not. (Note: if you've forgotten *exact* ISPC syntax it's okay, just write good pseudocode.)

C. How would the code you produced in part A change as a result of your ISPC program rewrite in part B? You do not need to provide the exact modified code here. A short explanation of the major difference is sufficient.

### Problem 8. Memory Consistency

Consider the following program which has four threads of execution. In the figure below, the assignment to  $x$  and  $y$  should be considered stores to those memory addresses. Assignment to  $r0$  and  $r1$  are loads from memory into local processor registers. (The print statement does not involve a memory operation.)

Processor 0	Processor 1	Processor 2	Processor 3
$x = 1$	$y = 1$	$r0 = y$ $r1 = x$ $\text{print } (r0 \ \& \ \sim r1)$	$r0 = x$ $r1 = y$ $\text{print } (r0 \ \& \ \sim r1)$

- Assume the contents of addresses  $x$  and  $y$  start out as 0.
- Hint: the expression  $a \ \& \ \sim b$  has the value 1 only when  $a$  is 1 and  $b$  is 0.

You run the program on a four-core system and observe that both processor 2 and processor 3 print the value 1. Is the system sequentially consistent? Explain why or why not?

### Problem 9: Parallel Histogram Generation (Yet Again)

Your friend implements the following parallel code for generating a histogram from the values in a large input array `input`. For each element of the input array, the code uses the function `bin_func` to compute a “bin” the element belongs to (`bin_func` always returns an integer between 0 and `NUM_BINS-1`), and increments a count of elements in that bin. His port targets a small parallel machine with only two processors. *This machine features 64-byte cache lines and uses an invalidation-based cache coherence protocol.* Your friend’s implementation is given below.

```
float input[N];           // assume input is initialized and N is a very large
int  histogram_bins[NUM_BINS]; // output bins
int  partial_bins[2][NUM_BINS]; // assume bins are initialized to 0
                                   // assume partial_bins is 64-byte aligned

//////////////////////////////////// Code executed by thread 0 //////////////////////////////////////
for (int i=0; i<N/2; i++)
    partial_bins[0][bin_func(input[i])]++;

barrier(); // wait for both threads to reach this point

for (int i=0; i<NUM_BINS; i++)
    histogram_bins[i] = partial_bins[0][i] + partial_bins[1][i];

//////////////////////////////////// Code executed by thread 1 //////////////////////////////////////
for (int i=N/2; i<N; i++)
    partial_bins[1][bin_func(input[i])]++;

barrier(); // wait for both threads to reach this point
```

- A. Your friend runs this code on an input of 1 million elements ( $N=1,000,000$ ) to create a histogram with eight bins ( $NUM\_BINS=8$ ). He is shocked when his program obtains far less than a linear speedup, and glumly asserts believe he needs to completely restructure the code to eliminate load imbalance. You take a look and recommend that he not do any coding at all, and just create a histogram with 16 bins instead. Who’s approach will lead to better parallel performance? Why?

B. Inspired by his new-found great performance, your friend concludes that more bins is better. He tries to use the provided code from part A to compute a histogram of 10,000 elements with 2,000 bins. He is shocked when the speedup obtained by the code drops. Improve the existing code to scale near linearly with the larger number of bins. (Please provide pseudocode as part of your answer – it need not be compilable C code.)

C. Your friend changes `bin_func` to a function with *extremely high arithmetic intensity*. (The new function requires 100000's of instructions to compute the output bin for each input element). If the histogram code **provided in part A** is used with this new `bin_func` do you expect scaling to be better, worse, or the same as the scaling you observed using the old `bin_func` in part A? Why? (Please ignore any changes you made to the code in part B for this question.)