

COMPUTER ARCHITECTURE (263-2210-00L), FALL 2018  
HW 5: CACHE PARTITIONING, MEMORY COHERENCE, MEMORY CONSISTENCY,  
MULTI-CORE, GPU PROGRAMMING

Instructor: Prof. Onur Mutlu

TAs: Mohammed Alser, Can Firtina, Hasan Hassan, Jeremie Kim, Juan Gómez Luna,  
Geraldo Francisco de Oliveira, Minesh Patel, Giray Yaglikci

Assigned: Thursday, Dec 6, 2018

Due: **Friday, Dec 21, 2018**

- **Handin - Critical Paper Reviews (1).** You need to submit your reviews to <https://safari.ethz.ch/review/architecture18/>. Please, check your inbox, you should have received an email with the password you should use to login. If you didn't receive any email, contact [comparch@lists.ethz.ch](mailto:comparch@lists.ethz.ch). In the first page after login, you should click in "Architecture - Fall 2018 Home", and then go to "any submitted paper" to see the list of papers.
- **Handin - Questions (2-8).** Please upload your solution to the Moodle (<https://moodle-app2.let.ethz.ch/>) as a single PDF file. **Please use a typesetting software (e.g., LaTeX) or a word processor (e.g., MS Word, LibreOfficeWriter) to generate your PDF file. Feel free to draw your diagrams either using an appropriate software or by hand, and include the diagrams into your solutions PDF.**

## 1 Critical Paper Reviews [200 points]

Please read the following handout on how to write critical reviews. We will give out extra credit that is worth 0.5% of your total grade for each good review.

- Lecture slides on guidelines for reviewing papers. Please, follow this format. <https://safari.ethz.ch/architecture/fall2018/lib/exe/fetch.php?media=onur-comparch-f18-how-to-do-the-paper-reviews.pdf>
  - Some sample reviews can be found here: <https://safari.ethz.ch/architecture/fall2018/doku.php?id=readings>
- (a) Write a one-page critical review for the first paper of the following list and at least **one** of the other 4 papers:
- M. Aater Suleman, Onur Mutlu, Moinuddin K. Qureshi, and Yale N. Patt, "Accelerating critical section execution with asymmetric multi-core architectures," ASPLOS 2009. [https://people.inf.ethz.ch/omutlu/pub/acs\\_asplos09.pdf](https://people.inf.ethz.ch/omutlu/pub/acs_asplos09.pdf)
  - Jose A. Joao, M. Aater Suleman, Onur Mutlu, and Yale N. Patt, "Bottleneck Identification and Scheduling in Multithreaded Applications," ASPLOS 2012. [https://people.inf.ethz.ch/omutlu/pub/bottleneck-identification-and-scheduling\\_asplos12.pdf](https://people.inf.ethz.ch/omutlu/pub/bottleneck-identification-and-scheduling_asplos12.pdf)
  - G. Pekhimenko, V. Seshadri, O. Mutlu, P.B. Gibbons, M.A. Kozuch, T.C. Mowry, "Base-Delta-Immediate Compression: Practical Data Compression for On-Chip Caches," PACT 2012. [https://people.inf.ethz.ch/omutlu/pub/bdi-compression\\_pact12.pdf](https://people.inf.ethz.ch/omutlu/pub/bdi-compression_pact12.pdf)
  - V. Seshadri, O. Mutlu, M. A. Kozuch, and T. C. Mowry, "The evicted-address filter: A unified mechanism to address both cache pollution and thrashing," PACT 2012. [https://people.inf.ethz.ch/omutlu/pub/eaf-cache\\_pact12.pdf](https://people.inf.ethz.ch/omutlu/pub/eaf-cache_pact12.pdf)
  - N. Vijaykumar, G. Pekhimenko, A. Jog, A. Bhowmick, R. Ausavarungnirun, C. Das, M. Kandemir, T. C. Mowry, and O. Mutlu, "A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Flexible Data Compression with Assist Warps," ISCA 2015. [https://people.inf.ethz.ch/omutlu/pub/caba-gpu-assist-warps\\_isca15.pdf](https://people.inf.ethz.ch/omutlu/pub/caba-gpu-assist-warps_isca15.pdf)

## 2 Utility-Based Cache Partitioning [100 points]

- (a) Does utility-based cache partitioning guarantee a minimum amount of cache space to each thread/core sharing the cache? Why or why not? Explain.

- (b) If yes, describe (and analyze) the minimum level of guarantee provided by utility based cache partitioning to each thread. If no, describe how the basic utility-based cache partitioning mechanism can be modified to provide a minimum amount of cache space to each thread.

- (c) Describe how you would perform utility based cache partitioning if each core has an identical prefetcher that prefetches into the shared cache. What needs to be modified in the utility based cache partitioning mechanism described by Qureshi and Patt (MICRO 2006) to take into account prefetches? Explain the new hardware design.

### 3 Cache Coherence [100 points]

We have a system with 4 byte-addressable processors. Each processor has a private 256-byte, direct-mapped, write-back L1 cache with a block size of 64 bytes. Coherence is maintained using the Illinois Protocol (MESI), which sends an invalidation to other processors on writes, and the other processors invalidate the block in their caches if *the block is present* (NOTE: On a write hit in one cache, a cache block in Shared state becomes Modified in that cache).

Accessible memory addresses range from 0x50000000 – 0x5FFFFFFF. Assume that the offset within a cache block is 0 for all memory requests. We use a snoopy protocol with a shared bus.

Cosmic rays strike the MESI state storage in your coherence modules, causing the state of a *single* cache line to instantaneously change to another state. This change causes an inconsistent state in the system. We show below the initial tag store state of the four caches, *after* the inconsistent state is induced.

#### Initial State

Cache 0		
	Tag	MESI state
Set 0	0x5FFFFFFF	M
Set 1	0x5FFFFFFF	E
Set 2	0x5FFFFFFF	S
Set 3	0x5FFFFFFF	I

Cache 1		
	Tag	MESI state
Set 0	0x522222	I
Set 1	0x510000	S
Set 2	0x5FFFFFFF	S
Set 3	0x533333	S

Cache 2		
	Tag	MESI state
Set 0	0x5F111F	M
Set 1	0x511100	E
Set 2	0x5FFFFFFF	S
Set 3	0x533333	S

Cache 3		
	Tag	MESI state
Set 0	0x5FF000	E
Set 1	0x511100	S
Set 2	0x5FFFF0	I
Set 3	0x533333	I

- (a) What is the inconsistency in the above initial state? Explain with reasoning.

- (b) Consider that, after the initial state, there are several paths that the program can follow that access different memory instructions. In b.1-b.4, we will examine whether the followed path can potentially lead to incorrect execution, i.e., an incorrect result.

b.1) Could the following path potentially lead to incorrect execution? Explain.

order	Processor 0	Processor 1	Processor 2	Processor 3
1 <sup>st</sup>			ld 0x51110040	
2 <sup>nd</sup>	st 0x5FFFFFF40			
3 <sup>rd</sup>				st 0x51110040
4 <sup>th</sup>		ld 0x5FFFFFF80		
5 <sup>th</sup>		ld 0x51110040		
6 <sup>th</sup>		ld 0x5FFFFFF40		

b.2) Could the following path potentially lead to incorrect execution? Explain.

order	Processor 0	Processor 1	Processor 2	Processor 3
1 <sup>st</sup>				ld 0x51110040
2 <sup>nd</sup>	ld 0x5FFFFFF00			
3 <sup>rd</sup>			ld 0x51234540	
4 <sup>th</sup>	st 0x5FFFFFF40			
5 <sup>th</sup>				ld 0x51234540
6 <sup>th</sup>	ld 0x5FFFFFF00			

After some time executing a particular path (which could be a path *different* from the paths in parts b.1-b.4) and with no further state changes caused by cosmic rays, we find that the final state of the caches is as follows.

**Final State**

<b>Cache 0</b>		
	<i>Tag</i>	<i>MESI state</i>
<i>Set 0</i>	0x5FFFFFFF	M
<i>Set 1</i>	0x5FFFFFFF	E
<i>Set 2</i>	0x5FFFFFFF	S
<i>Set 3</i>	0x5FFFFFFF	E

<b>Cache 1</b>		
	<i>Tag</i>	<i>MESI state</i>
<i>Set 0</i>	0x5FF000	I
<i>Set 1</i>	0x510000	S
<i>Set 2</i>	0x5FFFFFFF	S
<i>Set 3</i>	0x533333	I

<b>Cache 2</b>		
	<i>Tag</i>	<i>MESI state</i>
<i>Set 0</i>	0x5F111F	M
<i>Set 1</i>	0x511100	E
<i>Set 2</i>	0x5FFFFFFF	S
<i>Set 3</i>	0x533333	I

<b>Cache 3</b>		
	<i>Tag</i>	<i>MESI state</i>
<i>Set 0</i>	0x5FF000	M
<i>Set 1</i>	0x511100	S
<i>Set 2</i>	0x5FFFF0	I
<i>Set 3</i>	0x533333	I

- (c) What is the *minimum* set of memory instructions that leads the system from the initial state to the final state? Indicate the set of instructions in order, and clearly specify the access type (ld/st), the address of each memory request, and the processor from which the request is generated.

## 4 Memory Consistency [100 points]

A programmer writes the following two C code segments. She wants to run them concurrently on a multicore processor, called SC, using two different threads, each of which will run on a different core. The processor implements *sequential consistency*, as we discussed in the lecture.

Thread T0		Thread T1	
Instr. T0.0	<code>a = X[0];</code>	Instr. T1.0	<code>Y[0] = 1;</code>
Instr. T0.1	<code>b = a + Y[0];</code>	Instr. T1.1	<code>*flag = 1;</code>
Instr. T0.2	<code>while(*flag == 0);</code>	Instr. T1.2	<code>X[1] *= 2;</code>
Instr. T0.3	<code>Y[0] += 1;</code>	Instr. T1.3	<code>a = 0;</code>

X, Y, and `flag` have been allocated in main memory, while `a` and `b` are contained in processor registers. A read or write to any of these variables generates a single memory request. The initial values of all memory locations and variables are 0. Assume each line of the C code segment of a thread is a *single* instruction.

- (a) What is the final value of `Y[0]` in the SC processor, after both threads finish execution? Explain your answer.

- (b) What is the final value of `b` in the SC processor, after both threads finish execution? Explain your answer.

With the aim of achieving higher performance, the programmer tests her code on a new multicore processor, called RC, that implements *weak consistency*. As discussed in the lecture, the weak consistency model has no need to guarantee a strict order of memory operations. For this question, consider a very weak model where there is *no* guarantee on the ordering of instructions as seen by different cores.

- (c) What is the final value of  $Y[0]$  in the RC processor, after both threads finish execution? Explain your answer.

After several months spent debugging her code, the programmer learns that the new processor includes a `memory_fence()` instruction in its ISA. The semantics of `memory_fence()` is as follows for a given thread that executes it:

1. Wait (stall the processor) until *all* preceding memory operations from the thread complete in the memory system and become visible to other cores.
2. Ensure *no* memory operation from any later instruction in the thread gets executed before the `memory_fence()` is retired.

(d) What *minimal* changes should the programmer make to the program above to ensure that the final value of `Y[0]` on RC is the same as that in part (a) on SC? Explain your answer.

## 5 Building Multicore Processors [200 points]

You are hired by Amdahl's Nano Devices (AND) to design their newest multicore processor.

Ggl, one of AND's largest customers, has found that the following program can predict people's happiness.

```
for (i = 12; i < 2985984; i++) {  
    past = A[i-12]  
    current = A[i]  
    past *= 0.37  
    current *= 0.63  
    A[i] = past + current  
}
```

A is a large array of 4-byte floating point numbers, gathered by Ggl over the years by harvesting people's private messages. Your job is to create a processor that runs this program as fast as possible.

Assume the following:

- You have magically fast DRAM that allows infinitely many cores to access data in parallel. We will relax this strong assumption in parts (d), (e), (f).
  - Each floating point instruction (addition and multiplication) takes 10 cycles.
  - Each memory read and write takes 10 cycles.
  - No caches are used.
  - Integer operations and branches are fast enough that they can be ignored.
- (a) Assuming infinitely many cores, what is the maximum steady state speedup you can achieve for this program? Please show all your computations.

(b) What is the minimum number of cores you need to achieve this speedup?

(c) Briefly describe how you would assign work to each core to achieve this speedup.

It turns out magic DRAM does not exist except in Macondo<sup>1</sup>. As a result, you have to use cheap, slow, low-bandwidth DRAM. To compensate for this, you decide to use a private L1 cache for each processor. The new specifications for the DRAM and the L1 cache are:

- DRAM is shared by all processors. DRAM may only process one request (read or write) at a time.
- DRAM takes 100 cycles to process any request.
- DRAM prioritizes accesses of smaller addresses and write requests. (Assume no virtual memory)
- The cache is direct-mapped. Each cache block is 16 bytes.
- It takes 10 cycles to access the cache. Therefore, a cache hit is processed in 10 cycles and a cache miss is processed in 110 cycles.

All other latencies remain the same as specified earlier. Answer parts (d), (e), (f) assuming this new system.

(d) Can you still achieve the same steady state speedup as before? Circle one: YES      NO

Please explain.

---

<sup>1</sup>An imaginary town featured in *One Hundred Years of Solitude* by the late Colombian author Gabriel García Márquez (1927-2014).

(e) What is the minimum number of cores your processor needs to provide the maximum speedup?

(f) Briefly describe how you would assign work to each core to achieve this speedup.

## 6 Parallel Speedup [200 points]

You are a programmer at a large corporation, and you have been asked to parallelize an old program so that it runs faster on modern multicore processors.

- (a) You parallelize the program and discover that its speedup over the single-threaded version of the same program is significantly less than the number of processors. You find that many cache invalidations are occurring in each core's data cache. What program behavior could be causing these invalidations (in 20 words or less)?

- (b) You modify the program to fix this first performance issue. However, now you find that the program is slowed down by a global state update that must happen in only a single thread after every parallel computation. In particular, your program performs 90% of its work (measured as processor-seconds) in the parallel portion and 10% of its work in this serial portion. The parallel portion is perfectly parallelizable. What is the maximum speedup of the program if the multicore processor had an infinite number of cores?

- (c) How many processors would be required to attain a speedup of 4?

- (d) In order to execute your program with parallel and serial portions more efficiently, your corporation decides to design a custom heterogeneous processor.

- This processor will have one large core (which executes code more quickly but also takes greater die area on-chip) and multiple small cores (which execute code more slowly but also consume less area), all sharing one processor die.
- When your program is in its parallel portion, all of its threads execute **only** on small cores.
- When your program is in its serial portion, the one active thread executes on the large core.
- Performance (execution speed) of a core is proportional to the square root of its area.
- Assume that there are 16 units of die area available. A small core must take 1 unit of die area. The large core may take any number of units of die area  $n^2$ , where  $n$  is a positive integer.
- Assume that any area not used by the large core will be filled with small cores.

- (i) How large would you make the large core for the fastest possible execution of your program?

(ii) What would the same program's speedup be if all 16 units of die area were used to build a homogeneous system with 16 small cores, the serial portion ran on one of the small cores, and the parallel portion ran on all 16 small cores?

(iii) Does it make sense to use a heterogeneous system for this program which has 10% of its work in serial sections?

Why or why not?

(e) Now you optimize the serial portion of your program and it becomes only 4% of total work (the parallel portion is the remaining 96%).

(i) What is the best choice for the size of the large core in this case?

(ii) What is the program's speedup for this choice of large core size?

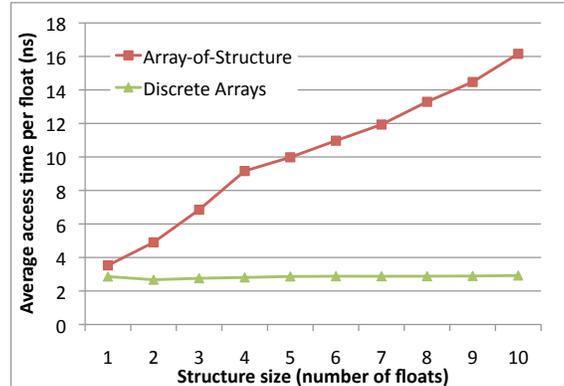
(iii) What would the same program's speedup be for this 4%/96% serial/parallel split if all 16 units of die area were used to build a homogeneous system with 16 small cores, the serial portion ran on one of the small cores, and the parallel portion ran on all 16 small cores?

(iv) Does it make sense to use a heterogeneous system for this program which has 4% of its work in serial sections?

Why or why not?

## 7 AoS vs. SoA on GPU [50 points]

The next figure shows the execution time for processing an array of data structures on a GPU. Abscissas represent the number of members in a data structure. Consecutive GPU threads read consecutive structures, and compute the sum reduction of their members. The result is stored in the first member of the structure.



The green line is the time for a kernel that accesses an array that is stored as discrete sub-arrays, that is, all  $i$ -th members of all array elements are stored in the  $i$ -th sub-array, in consecutive memory locations. The red line is the kernel time with an array that contains members of the same structure in consecutive memory locations.

- Why does the red line increase linearly? Why not the green line?

- How can the effect on the red line be alleviated?

- How would both kernels perform on a single-core CPU with one level of cache? And on a dual-core CPU with individual caches? And on a dual-core CPU with a shared cache?

## 8 Histogram Calculation on GPU [150 points]

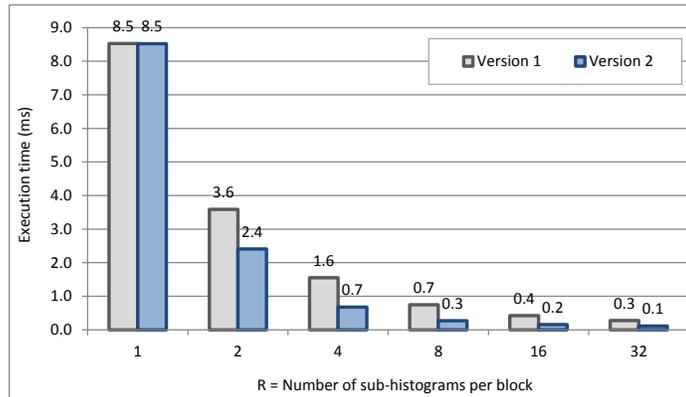
Histograms are a powerful tool in many fields, such as image processing. Their implementation on GPU is challenging because of the need for atomic operations. One way to accelerate their computation is using privatization in the fast shared memory. The following code calculates the histogram of an image "img" using privatization:

```
1 extern __shared__ unsigned int Hs[]; // Dynamic shared memory allocation
2 __global__ void histogram_kernel(
3     unsigned int* histo, unsigned int* img, int size, int BINS){
4     // Block and thread index
5     const int bx = blockIdx.x;
6     const int tx = threadIdx.x;
7     // Constants for read access
8     const int begin = bx * blockDim.x + tx;
9     const int end = size;
10    const int step = blockDim.x * gridDim.x;
11    // Sub-histogram initialization
12    for(int pos = tx; pos < BINS; pos += blockDim.x) Hs[pos]=0;
13    __syncthreads(); // Intra-block synchronization
14    // Main loop
15    for(int i = begin; i < end; i += step){
16        // Global memory read
17        unsigned int d = img[i];
18        // Atomic vote in shared memory
19        atomicAdd(&Hs[d], 1);
20    }
21    __syncthreads(); // Intra-block synchronization
22    // Merge in global memory
23    for(int pos = tx; pos < BINS; pos += blockDim.x){
24        unsigned int sum = 0;
25        sum = Hs[pos];
26        // Atomic addition in global memory
27        atomicAdd(histo + pos, sum);
28    }
29 }
```

- (a) As natural images are smooth (that is, they present spatial correlation), it is very likely that neighboring pixels fall into the same bin. To avoid atomic conflicts, R sub-histograms per block can be used (and later merged). Lets analyze two different ways of accessing the sub-histograms (to replace line 19):

```
atomicAdd(&Hs[BINS * (tx % R) + d], 1); // Version 1
atomicAdd(&Hs[tx % R + d * R], 1); // Version 2
```

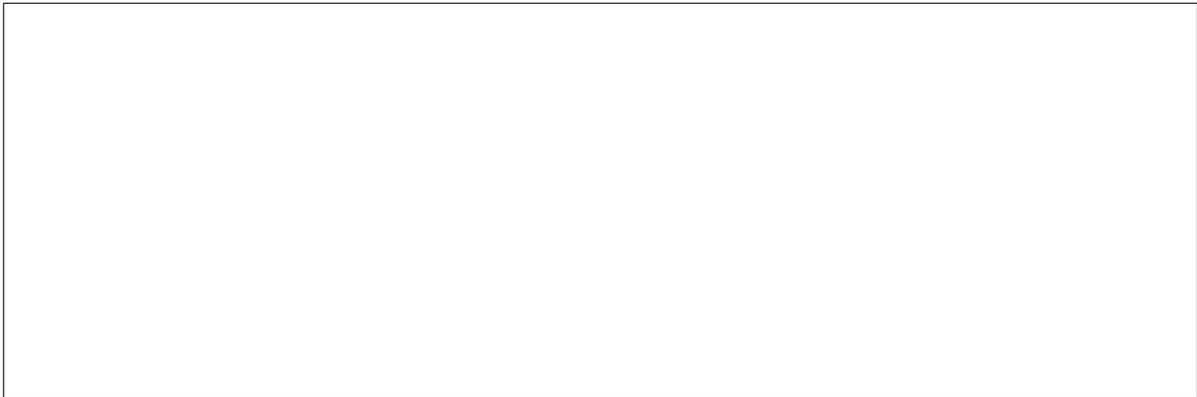
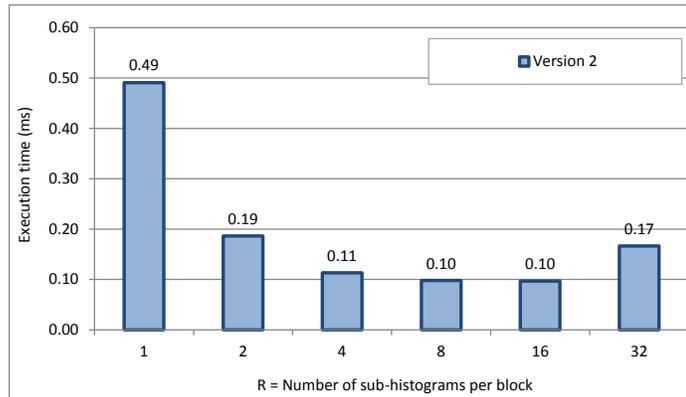
This graph shows the execution time for a 32-bins image histogram:



Why does version 2 obtain better results? What would happen for an odd-number-sized histogram?



- (b) As can be seen in the above graph, increasing the number R of sub-histogram tends to reduce the number of atomic conflicts, and consequently the execution time. Could you then explain the following graph? (Note: Histograms of 256 bins are calculated. Tests have been carried out on a Kepler GPU with a maximum of 64 warps per multiprocessor, and 48 KB of shared memory. Blocks of 256 threads are used).



- (c) For very large histograms, privatization in shared memory is not possible, unless multiple passes are carried out. Assume that, given the limited shared memory availability,  $N$  passes are needed. Atomic operations in shared memory take 2 ns to complete. For each pass, 10% of the input data loads hit the L2 cache. Compare this multi-pass approach to an approach where the histogram resides in global memory. Assume a GPU with global memory atomic operations in L2. Each atomic operation takes 10 ns to complete in L2, and 200 ns to complete in DRAM. 95% of the atomic operations hit the L2 cache. Find the value of  $N$  that makes worthwhile each of the approaches. (Note: the global memory bandwidth is 100 GB/s, and the L2 is 10 times faster).

