

CHALMERS TEKNISKA HÖGSKOLA
 Institutionen för data- och informationsteknik
 Avdelningen för datorteknik

Exam in EDA284 (Chalmers) and DIT361 (GU) Parallel Computer Architecture, Saturday, March 21st, 2020, 8:30h - 12:30h

Teacher/Lärare: During the exam you can send questions to miquelp@chalmers.se (examiner) with cc: to musabdu@chalmers.se (TA), or using the Canvas Inbox. Answers to common questions will be published in the FAQ that you can find in Canvas.

Language/Språk: Answers shall be given in English. The results should be submitted via Canvas in pdf (preferred) or in .docx format. You may include diagrams and answers written on paper by scanning them with your phone, but for readability we prefer that you instead use keyboard and drawing software.

Solutions/Lösningar: Solutions will be posted on Tuesday, March 24th, 2020 on the Canvas page.

Exam review/Granskning: The review date will be posted on the course Canvas page by the time you receive the email from LADOK.

Aids/Hjälpmedel: Since this is a take-home exam, all aids are allowed during examination. Plagiarism is of course strictly prohibited and checks will be conducted on all exams.

Grades:

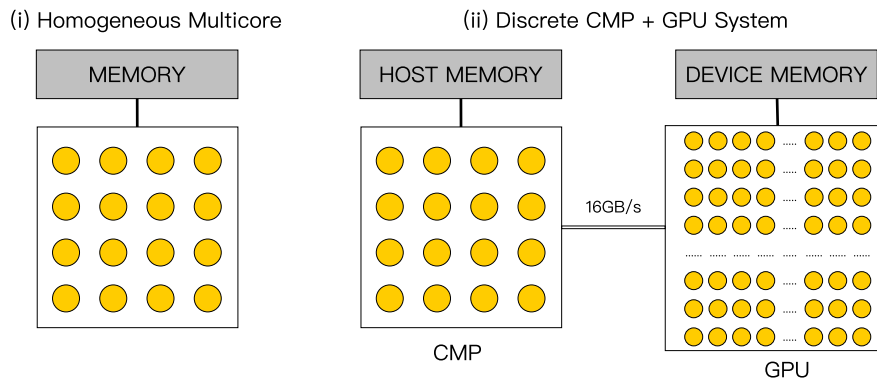
Chalmers				
Points	0-23	24-35	36-47	48-60
Grade	Failed	3	4	5

GU				
Points	0-23	24-41	42-60	
Grade	Failed	G	VG	

Good Luck!

Problem 1 (12 points)

A team of researchers is considering purchasing a new computing system to run a scientific application in the fastest possible time. The team has analyzed the application and observed that it consists of two phases: (a) a serial phase, which runs for 0.5 second on a single core, and (b) a parallel phase, which consists of 16G ($16 \times 1024 \times 1024 \times 1024$) floating point (FP) operations. All FP operations can be executed in parallel. The parallel phase has a streaming behavior and there is no reuse of data, so caches have no impact. The arithmetic intensity of phase (b) is 0.4 Flops/Byte according to the DRAM Roofline model. The team has two architecture choices: (i) a multicore (CMP) architecture with 16 single-threaded cores, and (ii) a two-chip system with the same homogeneous multicore CMP with 16 cores, a discrete GPU with 128 SMs (SIMT Core Clusters), and an interconnect bus between CMP and GPU with a fixed bandwidth of 16GB/s. See the figures below for diagrams of these two architectures.



The peak performance throughput for each core on the CMP is 0.5GFlops/s, and for each SM on the GPU it is 1GFlops/s. The memory bandwidth is 16GB/s for the CMP and 128GB/s for the GPU. The processed dataset must be present in the host memory after the computation has finished. When using the interconnect bus, you may consider that other connection latencies such as link setup or message assembly are negligible. Overlapping of communication and computation is not possible, i.e. these two phases must happen at different times. Finally, note that the serial phase can only be executed on a CMP core, never on the GPU. In order for the team to select the best architecture and the best configuration (i.e. number of cores or SMs), your task is to find the case that leads to the minimum total execution time:

- What will be the execution times of architecture (i) when using 4, 8 or 16 cores? Show the steps to reach the results.
- Next consider architecture (ii): What will be the execution times if the number of used SMs on the GPU is 32, 64 or 128? Given the results from (a) and (b), which configuration should be ultimately chosen? Discuss the impact considering both execution time and energy consumption.
- The memory on the CMP can be upgraded to a memory bandwidth of 32GB/s. If so, what will be the execution times of part (a)? Similarly, if a more advanced memory technology is used on the GPU, delivering 256GB/s, what will be the execution times of part (b)? Does this impact the team's choice of architecture and configuration?
- Considering the above results, what is the main bottleneck that needs to be addressed in order to improve the execution time in each architecture?

Problem 2 (10p)

The invalidate-based MESI cache coherence protocol state diagram was introduced in Lecture 4 - Slide 22. You are given the following machine code

```

1 (1) LW z1, 0(z4) // LW: Load Word
2 (2) LW z2, 0(z5)
3 (3) SW z3, 0(z5) // SW: Store Word
4 (4) SW z2, 0(z4)
5 (5) LW z1, 0(z5)
    
```

Assume that the addresses in z4 and z5 map to different cache lines. The code is executed by two cores, namely C0 and C1. Each has a private L1 cache. Following is the instruction sequence in real time:

C0(1), C0(2), C1(1), C1(2), C0(3), C1(3), C1(4), C0(4), C0(5), C1(5)

Here, C0(1) means: C0 executes instruction 1.

Core(Instruction)	State z4 - C0	State z5 - C0	State z4 - C1	State z5 - C1
C0(1)				
C0(2)				
C1(1)				
C1(2)				
C0(3)				
C1(3)				
C1(4)				
C0(4)				
C0(5)				
C1(5)				

- (a) For the table above, fill in the MESI state of the L1 cache lines for z4 and z5 at C0 and C1 after each call in the sequence.
- (b) For this code, and assuming the system is implemented as a Chip Multiprocessor (CMP), which protocol is preferred: MOESI or MESI? Why?
- (c) Assume a cache line size of 32 bytes. This time, memory locations mapped by z4 and z5 are accessed in offsets of the processor numbers (C0.offset = 0 x 2 = 0. C1.offset = 1 x 2 = 2).

Hence, instruction (3) becomes:

SW z3, 0(z5) on C0.

SW z3, 2(z5) on C1.

Do you see a potential problem? Explain in 1 line, and modify either instruction to fix the problem.

Problem 3 (10p)

In the design exploration of a new chip-multiprocessor system that is designed to have 512 cores, the architecture team decides to use a clustered (hierarchical) organization in which each cluster consists of 64 cores, each core has its own private L1 cache and the 64 cores share a L2 cache. The cache block size is 64 bytes.

- (a) The architecture team first decides to use a presence-flag-based directory cache protocol inside each cluster and across clusters. Calculate the overhead of maintaining cache coherence in each L2 cache block for the intra-cluster protocol, and in each memory block for the inter-cluster protocol.
- (b) Next the architecture team changes to use a limited-pointer directory cache protocol with four pointers inside each cluster, and a coarse-vector directory cache protocol that partitions the clusters into groups of four clusters. Calculate the overhead of maintaining cache coherence in each cache block at the L2 cache level (intra-cluster) and at the memory block level (inter-cluster).

For the intra-cluster interconnect the team is considering four choices of on-chip interconnection networks: a 1D linear array network (LA), a uni-directional ring (UR), a bi-directional ring (BR), and a $N \times N$ mesh network (NM), as exemplified in Figure 1 below for a smaller number of cores. The team has decided to use a directory-based (DB) protocol between the L1 caches as the intra-cluster coherence protocol. The latency to communicate between two cores connected directly by a link is 1 cycle. The router latency is one cycle per router that the packet goes through. The access time to the directory is a fixed 8 cycles. Assume that there is no contention in any link or the directory. In this exercise, latency is considered to be the time it takes to reach the destination cores, not including any acknowledgments.




Fig1.a




Fig1.b




Fig1.c




Fig1.d

- (c) What are the worst-case and average-case latencies for a coherence message to reach all possible destinations (inside the cluster) for each of the four interconnect topologies? Assume that the requesting core is located in the position of core 0 in the above figure.

Problem 4 (10p)

Modern parallel systems are built from multiple CMPs connected over a network, each having a complex cache hierarchy with multiple levels of private and shared caches. In this exercise, we ask you to propose and discuss several CMP-based system architectures tailored for several application scenarios. The only two design parameters to be considered are (1) the cache hierarchy (i.e. levels of private and shared caches), and (2) the hardware support for message passing.

The list of application scenarios is as follows:

- (a) The workload consists of multithreaded SPMD programs (e.g. OpenMP), such that there is a high degree of inter-core data sharing but threads only update their own local data.
- (b) The workload consists of multithreaded SPMD programs (e.g. OpenMP), such that there is high inter-core data sharing and threads update each other's data.
- (c) The workload consists of message passing SPMD programs (e.g. MPI), such that each MPI process consists of a single thread and there is heavy communication between MPI processes.
- (d) The workload consists of message passing SPMD programs (e.g. MPI), such that each MPI process is multithreaded and communication between MPI processes is light and infrequent.
- (e) The workload consists of applications from unrelated users. Assume that each application uses only a single core and that each application has a different working set size.

Your task is to, for each of the five scenarios (a)-(e):

1. Propose a system architecture that can execute the workload targeting high performance and low energy (you may draw a diagram if you want), and
2. Briefly discuss why the selected cache hierarchy and message passing support are appropriate for each scenario

Problem 5 (10p)

The following code approximates an integral using middle Riemann sum.

```

1  typedef float* arr_t;
2  float integrate(arr_t x, arr_t fx, int bound1, int bound2){
3      float integration = 0;
4      for (unsigned int i = bound1; i < bound2 - 1; i++)
5          integration += ((fx[i] + fx[i + 1]) * 0.5) * ( x[i + 1] - x[i]);
6      return integration;
7  }
8  void read_heavy_io_from_file(arr_t& arr, int& N) {
9      // file processing
10 }
11 void display_graph(arr_t x, arr_t fx, int N, float integral) {
12     // image rendering
13 }
14 void main() {
15     arr_t x, fx; int N;           // declare variables
16     read_heavy_io_from_file(x, N); // read x values
17     read_heavy_io_from_file(fx, N); // read fx values
18     float integral = integrate(x, fx, 0, N); // compute integration
19     display_graph(x, fx, N, integral); // display histogram
20 }

```

Below are your profiling results for each single call to the functions for $N = N_H$ on a single CMP:

Function Name	Time
read_heavy_io_from_file	1
display_graph	1
integrate	6

- You decide to parallelize the function `integrate` on $P = 8$ CMPs to speedup the solution for $N = N_H$. 1) What is your expected speedup (write the formula and substitute the terms). 2) What is the maximum speedup that can be achieved for $P = \infty$?
- You decide to run P such problems on P CMPs, i.e., the overall problem size solved on P CMPs is $N_P = P \times N_H$. What is the expected speedup achieved by running the Riemann sum of size N_P on P CMPs (as a function of P)? Note that the execution times shown by the table linearly scale with input size. Describe one advantage and one limitation of such parallelization technique.
- You are asked to enhance the ISA to support vector arithmetic instructions. Name at least 3 vector floating point arithmetic instructions needed to vectorize this code and describe their functionality.
- For the `integrate` function, assume a scalar code time $T_S = 8$. Having `SIMD_LENGTH=128` bits, assume that the code is vectorized by the compiler and is fully scalable to the vector width. What is your expected vector code time T_V given that a `float` is 4 bytes.

Problem 6 (8p)

Consider the following two barrier implementations written in C++:

Barrier #1:

```

1 // global variables
2 int counter = threadnum; // threadnum is the number of threads in barrier
3
4 // per-thread function to call barrier
5 void barrier(){
6     counter--; // decrease pending threads counter
7     while(counter){ }; // wait until all threads have arrived
8 }

```

Barrier #2:

```

1 // global variables
2 struct{
3     int f; // local flag to indicate that a thread has reached the barrier
4 }flags [threadnum] __attribute__((aligned(64)));
5 // initialized to {0, 0, 0, ...}, no false-sharing
6 bool done = false; // indicates if threads can be released from barrier
7
8 // per-thread function to call barrier
9 void barrier(){
10     int l_counter; // counter used by thread 0 to count waiting threads
11     int thread_id = get_thread_id(); // thread_id is the calling thread's
12                                     // identifier = {0, ..., threadnum-1}
13     flags[thread_id].f = 1; // announce #thread_id has reached barrier
14     if(thread_id == 0 // thread 0 checks the state of the barrier
15         while(!done){
16             l_counter = 0;
17             for(int i = 0; i < threadnum; i++)
18                 l_counter += flags[i].f; // increment l_counter if flag is set
19             if(l_counter == threadnum) done = true; // release the barrier
20         }
21     else // wait until all other threads have arrived
22         while(!done){ }; // and the barrier is released
23 }

```

These two barriers are to be used on a bus-based multiprocessor with private L1/L2 caches. Coherence is managed via a snoop-based protocol based on MSI-invalidate. Your task is to discuss the relative merits of these two implementations:

- Assume first that each line is executed atomically and behaves according to sequential consistency. How well does each of these barriers scale when additional cores are added to the system? Discuss the scalability of each barrier in terms of the cache coherence protocol.
- Now relax the assumption of atomic execution, as is the case with real hardware. Will the above codes still execute correctly? Are there any instructions that need special treatment from the hardware? If so, please identify those instructions and describe how the hardware should behave when encountering these special instructions.
- Finally, relax also the assumption of sequential consistency. Will barrier #2 operate correctly under relaxed memory consistency models?