



Re-examination August 2020 EDA283/EDA284/DIT360/DIT361 Parallel Computer Architecture

Contact: Miquel Pericàs, Examiner
Email: miquelp@chalmers.se

Written Component

Examination guidelines:

- The deadline to submit your answers is Friday, August 21st at 8:00h. Please submit your answers via the assignment in Canvas.
- It is preferable for us if you prepare your answers directly in digital format, but scanned, hand-written answers will also be accepted.
- Answers are to be prepared individually. We will conduct checks on all exams to detect violations.
- If you have any questions, please contact the examiner either by email or via the Canvas inbox.

Problem 1 (10p)

Assume the following message passing code, expressed in two versions:

Synchronous version:

CODE FOR THREAD 0:

```
B[1] = 1024;
SEND(&B[1], sizeof(float), T1, SEND_B1);
RECV(&B[0], sizeof(float), T1, SEND_A1);
Unrelated computation;
```

CODE FOR THREAD 1:

```
A[1] = 256;
RECV(&A[0], sizeof(float), T0, SEND_B1);
SEND(&A[1], sizeof(float), T0, SEND_A1);
Unrelated computation;
```

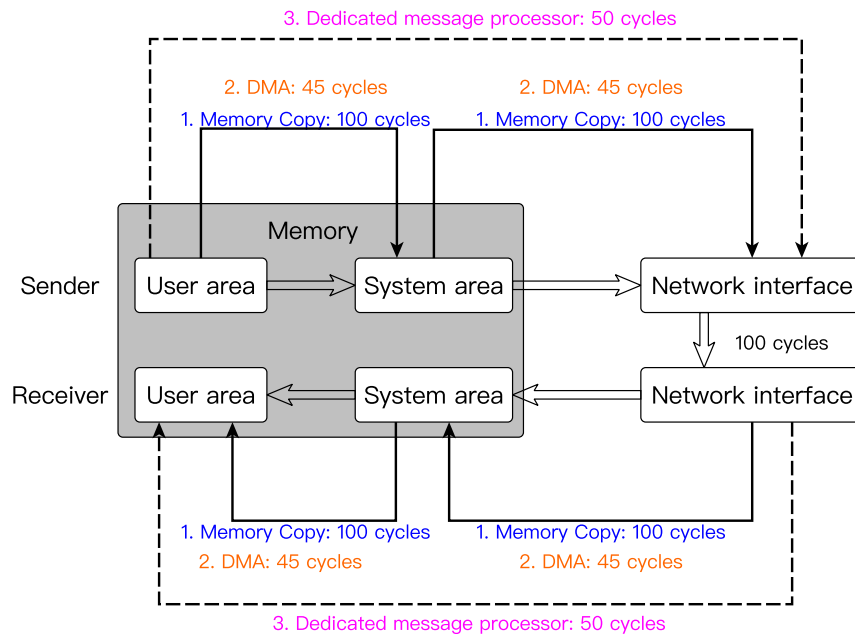
Asynchronous version:

CODE FOR THREAD 0:

```
B[1] = 1024;
ASEND(&B[1], sizeof(float), T1, SEND_B1);
Unrelated computation;
ARECV(&B[0], sizeof(float), T1, SEND_A1);
```

CODE FOR THREAD 1:

```
A[1] = 256;
ASEND(&A[1], sizeof(float), T0, SEND_A1);
Unrelated computation;
ARECV(&A[0], sizeof(float), T0, SEND_B1);
```



Assume that the unrelated computation takes 200 cycles. Consider the following four scenarios:

(a) No special hardware support: It requires five steps of memory copy for a

message from a sender to a receiver. How long does it take to execute the program with synchronous and asynchronous primitives, respectively?

- (b) DMA programmed by O/S without support from user messages: But programming DMA costs 15 cycles overhead per each step. What will be the new execution time of synchronous and asynchronous version respectively?
- (c) User level messages with O/S support and DMA: Programming DMA costs 15 cycles overhead per step. In this scenario, the message could be delivered directly to the user area from the network interface. However, incoming messages are taken care of by users instead of OS, which costs context switches 100 cycles. What will be the new execution time of synchronous and asynchronous version respectively?
- (d) User level messages with a dedicated message processor, shown as dotted line in figure: This dedicated message processor is in the NIC hardware. Hence, it comes for free but still requires a context switch of 100 cycles. What will be the new execution time of synchronous and asynchronous versions, respectively?

- (d) The bus action(s) (BusRd, BusUpgr, BusRdX, flush).
- (e) Design a sequence of 4-5 instructions that result in MESI causing more bus traffic than MSI.

Note: use the tables below to answer parts a - d

Inst	Memory content
-	A(6),B(6)
1	
2	
3	
4	
5	
6	
7	
8	

Inst	P1			
	Cache content	State	PrReq	BusReq
-	-	I	-	-
1				
2				
3				
4				
5				
6				
7				
8				

Inst	P2			
	Cache content	State	PrReq	BusReq
-	-	I	-	-
1				
2				
3				
4				
5				
6				
7				
8				

Inst	P3			
	Cache content	State	PrReq	BusReq
-	-	I	-	-
1				
2				
3				
4				
5				
6				
7				
8				

Problem 3 (10p)

We want to compare a shared memory (SM) machine and distributed memory (DM) system. Each system has a total of 8 cores. In the SM system, the 8 cores share the main memory (DRAM). The bandwidth of this DRAM memory is 25GB/s. Each core has a peak performance of 10 GFLOPs/s. The DM system consists of 4 nodes, each with 2 cores that are identical to the cores of the SM system. Each node has a DRAM with the same characteristics as the SM system. The nodes are interconnected by a network that consists of a shared bus and has a total bandwidth of 1GB/s.

The program to be executed is an 8-point 2D stencil. C-like pseudocode for the sequential version of this program is shown below:

```
1  #define SIZEX 1024
2  #define SIZEY 1024
3
4  // shared space on which the program is operating
5  float area1[SIZEX][SIZEY], area2[SIZEX][SIZEY];
6
7  void 2dstencil(float in[SIZEX][SIZEY], float out[SIZEX][SIZEY])
8  {
9      for(i=2; i<SIZEX-2; i++)
10         for(j=2; j<SIZEY-2; j++)
11             out[i][j] += 0.125*(in[i+1][j] + in[i-1][j] +
12                                 in[i+2][j] + in[i-2][j] +
13                                 in[i][j+1] + in[i][j-1] +
14                                 in[i][j+2] + in[i][j-2]);
15 }
16
17 main()
18 {
19     float *in, *out, *tmp;
20     in = area1;
21     out = area2;
22
23     for(timestep=0; timestep < MAX_TIMESTEPS; timestep++){
24         2dstencil(in, out);           // perform stencil
25         tmp = out; out = in; in = tmp; // exchange pointers
26     }
27 }
```

Assume that the caches are such that data reuse is possible only within a timestep, but not across timesteps. Within a timestep, caches behave perfectly: they retain all input data needed to compute the kernel and provide unlimited bandwidth to

the cores. Your tasks are:

- (a) Calculate the number of memory accesses and floating point operations in `2Dstencil()`. How many arithmetic operations are performed for each byte of data accessed from memory? Determine the performance of a single core when running the `2Dstencil()` function when all other processors are idle. What is the performance when all cores are active for both the SM and for the DM systems?
- (b) To parallelize the application, each core computes an equivalent partition of the output grid in the row-direction (X). For the case of distributed memory, and considering the above sizes of `SIZEX` and `SIZEY`, how much data needs to be transferred between the 4 nodes after each timestep?
- (c) Construct a formula that determines the execution time, for the two above systems, as a function of
 - (a) the number of timesteps (`MAX_TIMESTEPS`), and
 - (b) the number of partitions (`PARTITIONS`).

Assume unlimited cores (case of SM, one partition per core) and unlimited nodes (case of DM, one partition per node). The memory system (DRAM) architecture remains unchanged. How does the execution time change as the number of partitions increases?

- (d) One of your colleagues suggests adding a GPU to accelerate the shared memory system. The GPU has a memory system that has 16 times the bandwidth of the DRAM in the SM system, but has a limited capacity of 8GB. The peak performance of the GPU is 1 TFLOP/s (single-precision). The GPU is connected to the system via an I/O bus such as PCI express, with much lower bandwidth compared to the GPU memory bandwidth. At what speed (FLOPS) can this GPU execute the `2Dstencil()` function? Given the problem size shown above, will it be beneficial to make use of this GPU?

Problem 4 (8p)

The following C++-like pseudocode represents an application segment that is running on a set of threads in a multicore system. The goal of the segment is to accumulate SIZE values stored in the array `acc_accum[] []`. In the proposed parallelization, the array is already pre-partitioned across NTHREADS threads.

```
1  #define SIZE 1000000
2  #define NTHREADS 4
3
4  std::atomic<int> counter = 0;
5  std::atomic<int> acc_array[NTHREADS][SIZE/NTHREADS];
6  // per-thread array initialized to some values
7
8  // NTHREADS threads all execute the following function in parallel
9  void thread()
10 {
11     int tid = get_thread_number(); // obtain an index to current thread
12     std::atomic<int> inc;
13
14     for(inc = 0; inc < SIZE/NTHREADS; inc++)
15     {
16         counter = counter + acc_array[tid][inc]; // accumulate
17     }
18
19     barrier(); // wait until all threads reach this point
20
21     if (tid == 0) printf ("The final value is %d\n", counter);
22 }
```

The tasks to be completed are as follows:

- (a) Discuss the quality of the proposed parallelization. How will it scale if hundreds of threads are executing?
- (b) Propose three optimizations to the code that will make it execute faster.
- (c) What is the function of `std::atomic`? How does it impact the processor's pipeline?
- (d) Your optimized code is likely to require the use of hardware supported atomic instructions. Atomic instructions may imply also sequential consistency, or they may enable more relaxed memory orders. In the case of your code, discuss what memory orders will be required by the atomic instructions.

Problem 5 (8p)

In the design exploration of a new chip-multiprocessor (CMP) system that is designed to have 256 cores, the architecture team decides to use a clustered (hierarchical) organization in which each cluster consists of 16 cores, each core has its own private L1 cache and the 16 cores share a L2 cache. Inside of each cluster, the architecture team decides to use bi-directional ring as the on-chip interconnection networks.

- (a) Assume a directory-based (DB) protocol is used to maintain cache coherence between cores inside of the same cluster. A core generates a write invalidation request. The latency to communicate between two cores connected directly by a link is 1 cycle and the directory lookup time is 10 cycles. What is the worst latency for the coherence message to reach all possible destinations?

In addition, the CMP system allows cluster-based frequency scaling, which means that the user can power off some clusters in order to allow the rest to operate at a higher frequency.

With the other clusters powered off, users can activate different number of clusters to operate at the following latencies (nanosecond per clock cycle):

When only 1 cluster is active, it operates at 0.4 ns/cycle;
When 2 clusters are active, they operate at 0.7 ns/cycle;
When 4 clusters are active, they operate at 1.3 ns/cycle;
When 8 clusters are active, they operate at 2.5 ns/cycle;
When 16 clusters are active, they operate at 4.2 ns/cycle.

Consider a parallel application that features 500 instructions in the serial part and 25600 instructions in the parallel part. Parallelization also adds extra overhead for creating threads, which adds 5 instructions in the serial part for each created thread.

- (b) Assume that running the serial part only needs one cluster to be active. What will be the execution time of the application on this CMP system when using 1, 2, 4, 8 or 16 clusters for the parallel part? Which configuration is the fastest?
- (c) Assume now that the system is redesigned such that 8 clusters (0-7) operate at 1.2 ns/cycle and the other 8 clusters (8-15) operate at 0.4 ns/cycle. All clusters are active at the same time on this new CMP. The serial part can be executed on one of the cores of the high frequency clusters (8-15) or low frequency clusters (0-7). How would you manage the execution of the parallel program in order to minimize the execution time?