

Set 4 - Blocking MPI Routines and Processor Pipelines

Issued: November 8, 2019

Hand in (optional): November 22, 2019 08:00am

Question 1: Diffusion in 2D (50 points)

Mass diffusion in a medium can be described by the following equation

$$\frac{\partial \rho(x, y, t)}{\partial t} = D \nabla^2 \rho(x, y, t) \quad (1)$$

where $\rho(x, y, t)$ represents the concentration of a substance at position \mathbf{r} and time t , and D represents a constant diffusion coefficient. Lets define the domain Ω in two dimensions as $x, y \in [-1, 1]$. We will use the boundary condition:

$$\rho(x, y, t) = 0 \quad \forall t \geq 0 \text{ and } (x, y) \notin \Omega \quad (2)$$

and an initial distribution:

$$\rho(x, y, 0) = \begin{cases} 1 & |x, y| < 1/2 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

For this exercise, use only **blocking** MPI routines.

- In the provided skeleton code, the 2D diffusion equation is computationally solved on a uniform grid, using central finite differences in space and forward Euler in time. Implement the MPI parallelization of the 2D diffusion equation by filling in all parts of the code marked by `TODO`, in functions `advance` and `main`. Decompose the domain using tiling decomposition scheme (described in the lecture notes). (i.e. distribute the rows evenly to the MPI processes).
- Compute an approximation to the integral of $\rho(x, y, t)$ over the entire domain in the function `compute_diagnostics` by implementing the missing MPI parts marked by `TODO`, and plot the result as a function of time using $D = 1$, $L = 2$ and $N = 100$.
- Compute a histogram of $\rho(x, y, t)$ over the entire domain in the function `compute_histogram` by implementing the missing MPI parts marked by `TODO`, and plot the resulting histogram at simulation time $t = 0.5$ using $D = 1$, $L = 2$ and $N = 100$.

Question 2: MPI Bug Hunt (20 points)

Find the bug(s) in the following MPI code snippets and find a way to fix the problem!

a)

```
1  const int N = 10000;
2  double* result = new double[N];
3  // do a very computationally expensive calculation
4  // ...
5
6  // write the result to a file
7  std::ofstream file("result.txt");
8
9  for(int i = 0; i < N; ++i){
10     file << result[i] << std::endl;
11 }
12
13 delete [] result;
```

b)

```
1  // only 2 ranks: 0, 1
2  double important_value;
3
4  // obtain the important value
5  // ...
6
7  // exchange the value
8  if(rank == 0)
9  MPI_Send(&important_value, 1, MPI_DOUBLE, 1, 123, MPI_COMM_WORLD);
10 else
11 MPI_Send(&important_value, 1, MPI_DOUBLE, 0, 123, MPI_COMM_WORLD);
12
13 MPI_Recv(
14 &important_value, 1, MPI_INT, MPI_ANY_SOURCE,
15 MPI_ANY_TAG, MPI_COMM_WORLD, MPI_STATUS_IGNORE
16 );
17
18 // do other work
```

c) What is the output of the following program when run with 1 rank? What if there are 2 ranks? Will the program complete for any number of ranks?

```
1  MPI_Init(&argc, &argv);
2
3  int rank, size;
4  MPI_Comm_size(MPI_COMM_WORLD, &size);
5  MPI_Comm_rank(MPI_COMM_WORLD, &rank);
6
7  int bval;
8  if (0 == rank)
9  {
10     bval = rank;
11     MPI_Bcast(&bval, 1, MPI_INT, 0, MPI_COMM_WORLD);
12 }
13 else
14 {
15     MPI_Status stat;
16     MPI_Recv(&bval, 1, MPI_INT, 0, rank, MPI_COMM_WORLD, &stat);
17 }
18
19 cout << "[" << rank << "]" << bval << endl;
20
21 MPI_Finalize();
22 return 0;
```

Question 3: Processor Pipelining (30 points)

- a) In the lecture we have looked at the pipeline block diagram of the Intel Coffee Lake CPU, which is identical to the pipeline in the Intel Skylake CPUs. You can find this diagram on page 78 in the lecture slides about CPU pipelining or directly at this link

https://en.wikichip.org/w/images/7/7e/skylake_block_diagram.svg

- i) Use this diagram to determine the number of data loads and data stores you can perform in the same cycle.
 - ii) With the help of the lecture slides, explain why the architecture designers have chosen the particular ratio of data load and data store instructions.
- b) Suppose we have the following functional units to process instructions with latencies:

IF	(Instruction Fetch)	2 ns
ID	(Instruction Decode)	2 ns
EX	(Execute Instruction)	3 ns
MEM	(Physical Memory Access)	6 ns
WB	(Write Back Result)	2 ns

- i) If we use these units to build a non-pipelined processor, how long does it take to execute a single instruction?
- ii) If we use these units to build the 5-stage pipelined processor discussed in the lecture, what is the shortest possible cycle time, i.e. the time for executing a single stage?
- iii) How long does it take to execute at least N instructions in a *full* pipeline? Assume the pipeline is empty at the beginning.
- iv) What is the speedup of this pipelined processor over the non-pipelined implementation? Assume a large number of instructions that do not cause pipeline stalls.