ETH zürich

P. Koumoutsakos
S. Martin
ETH Zentrum, ML H44
CH-8092 Zürich

High Performance Computing for Science and Engineering I

Fall semester 2021

Set 6 -MPI

Issued: December 10, 2021
Hand in (optional): December 24, 2021 12:00

Question 1: MPI Bug Hunt (16 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;
```

- There is a segfault hidden in line 9.
  It can be fixed by changing `i <= N` to `i < N`.
- All ranks write simultaneously to the same output file! This is a problem for many reasons:
  On one hand, it leads to an incorrect output file because of many concurrent writes that overwrite each other.
  On the other hand, the same work is done many times. One remedy would be to let only the root rank write the output. This can be implemented as shown in the code below.

```
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  if(rank == 0){
```

```
8       std::ofstream file("result.txt");
9
10      for(int i = 0; i < N; ++i){
11          file << result[i] << std::endl;
12      }
13  }
14
15  delete[] result;
```

Of course, this assumes that every rank has the same data. If that is not the case, one would have to send all the information to the root rank.
In next semester, you will learn to to use MPI to perform I/O (input/output) operations involving many ranks in parallel.

- 5 points in total
  - 1.25 points for each identified bug
  - 1.25 points per fixed bug
  - -2.5 points if another bug is introduced through the "improvement"
  - you cannot have less than zero points in this subquestion

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
```

- The MPI type in the receive call (`MPI_INT`) does not match the MPI type that is sent (`MPI_DOUBLE`). Change it to `MPI_DOUBLE` to ensure defined behaviour.
- This code will deadlock: Both ranks send first and receive later. Neither of the two ranks can return from the send call because the call blocks until the sending is completed. But the sending cannot complete because the corresponding receive is not called yet. Deadlock!
  Possible solution:

```
// only 2 ranks: 0, 1
double important_value;

// obtain the important value
// ...

// exchange the value
if(rank == 0){
    MPI_Send(&important_value, 1, MPI_DOUBLE, 1, 123, MPI_COMM_WORLD);
    MPI_Recv(
        &important_value, 1, MPI_DOUBLE, MPI_ANY_SOURCE,
        MPI_ANY_TAG, MPI_COMM_WORLD, MPI_STATUS_IGNORE
    );
}else{
    MPI_Recv(
        &important_value, 1, MPI_DOUBLE, MPI_ANY_SOURCE,
        MPI_ANY_TAG, MPI_COMM_WORLD, MPI_STATUS_IGNORE
    );
    MPI_Send(&important_value, 1, MPI_DOUBLE, 0, 123, MPI_COMM_WORLD);
}
// do other work
```

- 5 points in total
  - 1.25 points for feach identified bug
  - 1.25 points per fixed bug
  - -2.5 points if another bug is introduced through the "improvement"
  - you cannot have less than zero points in this subquestion

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;
```

For only 1 rank, everything is fine and the output of the program is:

```
[0] 0
```

With 2 ranks, however, there will be a deadlock:
Rank 0 arrives at the broadcast. This is a blocking *collective* operation, that means rank 0 waits until all other ranks in the communicator `MPI_COMM_WORLD` have reached this point and performed the collective operation. But rank 1 never arrives at this point! Instead, it gets stuck in the receive operation.

- 2 points for predicting correct behavior with one rank
- 2 point for predicting correct behavior with two ranks
- 2 points for justifying the behavior of two ranks

Question 2: Implementing a distributed reduction (45 points)

In this question, you will use MPI to calculate the following sum:

$$x_{\text{tot}} = \sum_{n=1}^{N} n = 1 + 2 + 3 + \ldots + (N-1) + N \tag{1}$$

a) Fill in the missing part in the Makefile in order to compile the code with MPI support.
See solution code.
2 points for correct Makefile.

b) Validation of HPC code is an important subject. For example, there is an analytic formula for the above sum. Use this to check if your implementation is correct. To this end, implement the function `exact(N)`. **Hint:** A young C.F. Gauss found the formula in elementary school.
Analytic solution:

$$x_{\text{tot}} = \sum_{n=1}^{N} n = \frac{N(N+1)}{2} \tag{2}$$

This formula can be proofed by induction (not necessary to get points):
$\underline{N = 1:}$

$$x_{\text{tot}} = 1 = \frac{1(1+1)}{2} \tag{3}$$

$\underline{N - 1 \to N:}$

$$\begin{aligned} x_{\text{tot}} &= \sum_{n=1}^{N} n \\ &= N + \sum_{n=1}^{N-1} n \\ &\overset{(*)}{=} N + \frac{(N-1)N}{2} \\ &= N + \frac{1}{2}\left(N^2 - N\right) \\ &= \frac{1}{2}\left(N^2 + N\right) \\ &= \frac{N(N+1)}{2} \end{aligned} \tag{4}$$

In (*), the induction hypothesis for $N - 1$ was used.

- **2 points for correct formula**
- **2 points for correct implementation**

c) Initialize and finalize MPI by filling the corresponding gaps in the skeleton code.
See solution code.

- **1 point for calling `MPI_Reduce`**

- **1 point for correct MPI datatype `MPI_LONG`**
- **1 point for correct reduction operation `MPI_SUM`**
- **1 point for sending all the data to rank 0**
  (other reduction targets are also fine as long as in the end, only one rank has all the data and the receiving rank also prints the result)
- **1 point for using the correct communicator `MPI_COMM_WORLD`**

d) Each rank performs only a part of the sum. Distribute the work load reasonably in order to guarantee load balancing. Each rank should calculate the subsum

$$\text{sum}_{\text{rank}} = N_{\text{start}} + (N_{\text{start}} + 1) + \ldots + N_{\text{end}}, \tag{5}$$

where N_{start} and N_{end} are the corresponding variables in the skeleton file.
See solution code.

- **2 points for correctly determining the start index**
  Note: Starting at zero is also correct, but implies that rank 0 sums one element less than the other ranks if the other code in the solution is not changed accordingly (see solution code). However, in this exercise, this does not represent much of an imbalance.
- **2 points for correctly determining the end index of all ranks**
- **2 points for taking care of the case where N `mod size` $\neq 0$**

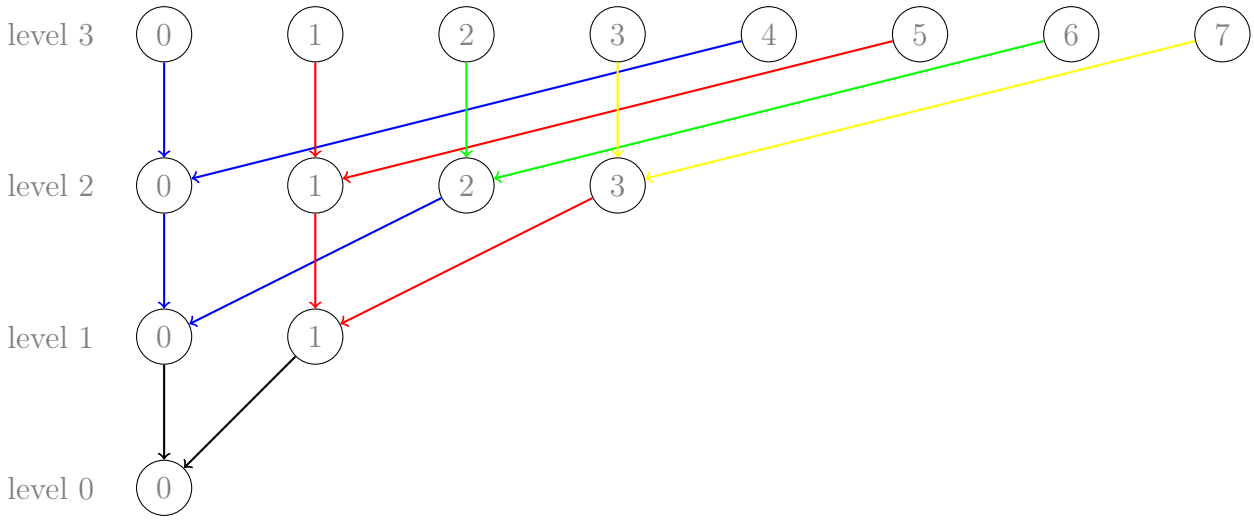

Figure 1: The communication pattern of a tree-like reduction. Each circle represents a rank, the number inside is the rank ID. Communication takes place along the arrows.

e) First Perform the reduction using MPI blocking collectives in the skeleton code. Then, implement your own reduction. This can be done in a tree-like way as depicted in Fig. 1. Your task is to implement this scheme for the special case that the number of ranks is a power of 2, i. e.

$$|\text{ranks}| = 2^l, l \in \mathbb{N}_0 \tag{6}$$

See solution code.

- **20 points for performing the communication "in levels", as shown in Fig. 1.**
  - **7 points for correct separation between sending and receiving ranks in each level**
  - **7 points for correctly determining the communication partner in each level**
  - **4 points for correct usage of `MPI_Send` and `MPI_Recv`**
  - **2 points if only the receiving rank performs the reduction operation (addition)**
- **5 points for testing the implementation at least with 2, 4, 8, 16, 32 ranks.**

f) What is the advantage of this scheme compared to the naive reduction? Name 2 advantages and quickly justify your answer. **Hint:** In the naive approach, every rank sends its elements directly to the master. The master then reduces all obtained elements by repeatedly applying the operation, in our case the sum.

- More ranks communicate with each other. Depending on the network topology, this means that there is potentially more bandwidth available.
- The reduction operation can be performed by many processes in parallel, leading to better load balance and faster execution of the reduction. In the naive version, every process has to wait until the root of the reduction has performed the reduction operation on all items received by another rank. Also, only one process can send at each time, leaving the others idle and wasting an opportunity to do something useful.

1.5 points per each pair of advantage and discussion (3 points in total)
No points if discussion is missing!

Figure 2: Left: Wake of an airplane visualized by condensation behind the engines. The vorticity sheet is generated at the trailing edge of the wings. Right: Vortex sheet at $t = 1$ from the simulation.

Question 3: Roll-up of a vortex line (30 points)

We want to simulate the evolution of a two-dimensional vorticity sheet similar to the wake of an airplane (Figure 2). We use N particles moving with time t, each particle i is located at $(x_i(t), y_i(t))$ and carries a constant value Γ_i (circulation). The velocity field defined as

$$u(x_j, y_j, t) = \sum_{i=0}^{N-1} \frac{\Gamma_i}{2\pi} \frac{-[y_j(t) - y_i(t)]}{\varepsilon^2 + [x_j(t) - x_i(t)]^2 + [y_j(t) - y_i(t)]^2},$$

$$v(x_j, y_j, t) = \sum_{i=0}^{N-1} \frac{\Gamma_i}{2\pi} \frac{[x_j(t) - x_i(t)]}{\varepsilon^2 + [x_j(t) - x_i(t)]^2 + [y_j(t) - y_i(t)]^2}$$

is an approximate solution of Euler equations and corresponds to the vorticity field

$$\omega_j = \frac{\partial v(x_j, y_j, t)}{\partial x_j} - \frac{\partial u(x_j, y_j, t)}{\partial y_j} = \sum_{i=0}^{N-1} \Gamma_i \delta_\epsilon\big(x_j(t) - x_i(t), y_j(t) - y_i(t)\big)$$

where $\delta_\epsilon(x, y) = \frac{1}{\pi} \frac{\epsilon^2}{(\epsilon^2 + x^2 + y^2)^2}$ approximate a Dirac-delta function for a small ϵ. A particle moves with the velocity at its location

$$\frac{dx_i(t)}{dt} = u(x_i(t), y_i(t), t), \quad \frac{dy_i(t)}{dt} = v(x_i(t), y_i(t), t),$$

Initially, particles are placed at $y_i = 0$ and

$$x_i = -\frac{1}{2} + \frac{i + 1/2}{N}, \quad i = 0, ..., N - 1$$

and have circulation $\Gamma_i = \frac{1}{N} \frac{4x_i}{\sqrt{1 - 4x_i^2}}$

a) Implement the interaction function `computeVelocities` in `q3/serial.cpp`. You may check your results by visualizing the csv files with paraview.

- **5 points for correct implementation of the update formula**

b) Parallelize your code using MPI by filling in the TODOs in `q3/mpi.cpp`. Each MPI rank must contain an equal number of particles. The parallelization of the `computeVelocities` can be done using a multi-pass communication, as described in the table:

		process p				
		0	1	$\cdots$	P-2	P-1
	0	D_0	D_1	$\cdots$	D_{P-2}	D_{P-1}
	1	D_{P-1}	D_0	$\cdots$	D_{P-3}	D_{P-2}
pass q	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	P-1	D_1	D_2	$\cdots$	D_{P-1}	D_0

the data necessary to compute the velocities (x_i, y_i and Γ_i) must be communicated to every rank in a cyclic manner until every rank has computed the interactions between its own particles with every particles in the simulation. Parallelize the function `dumpToCsv` by gathering the data on the root.

- **5 points for having a total of one copy of the initial x, y, Γ arrays.**
- **7 points for getting the same result as the serial code.**
- **5 points for correct Gather operations in dumpToCsv**

c) Use non blocking MPI routines to overlap the communication with computation. Write your solution in `q3/mpi_non_blocking.cpp`.

- **2 points for having a total of two copies of the initial x, y, Γ arrays (one for receiving data and one for computing, see solution code).**
- **1 point for swapping pointers for the two copies (see solution code).**
- **2 points for using MPIWait after the computation takes place.**
- **3 points for getting the same result as the serial code.**