

Innovative Computing Laboratory University of Tennessee, Knoxville

HW2 : Heat propagation using MPI

v. 1.0.0

George Bosilca

September 13, 2017




Stencil 2D

Heat transfer



- ▶ Heat transfer is a discipline of thermal engineering that concerns the generation, use, conversion, and exchange of thermal energy (heat) between physical systems.
- ▶ Heat convection occurs when bulk flow of a fluid (gas or liquid) carries heat along with the flow of matter in the fluid.



The square is the 2D surface where the heat propagation is to be observed. The circles are the heat sources, and the heat waves they generate. The red lines are the impact of these heat sources on the boundaries of the 2D surface, and represent the stable boundary condition of the problem (they are stored in an extra column/row and are not supposed to be altered during the execution).



You will take the sequential code and transform it into a fully distributed application using MPI. Assume the data is row-major (contiguous in memory per row). Assume the application will use a cartesian grid of $P \times Q$ processors, and the data will be distributed by 2d blocks. We will use the data distribution described in the lecture (with the ghost region on each node) as indicated in the Figure 1. This is similar to the slicing used in the multi-threaded case, except the ghost region will now exist for each process.

To facilitate the development, here are the steps you should follow:

- ▶ Construct the datatype for the north and south ghost regions
- ▶ Construct the datatype for the east and west ghost regions
- ▶ Build the communication pattern where each process (with the exception of those on the boundary) exchange data with its 4 neighbors: east, west, north and south.
- ▶ Plug the node level Jacobi into the code
- ▶ Make sure you reuse as much code as possible.


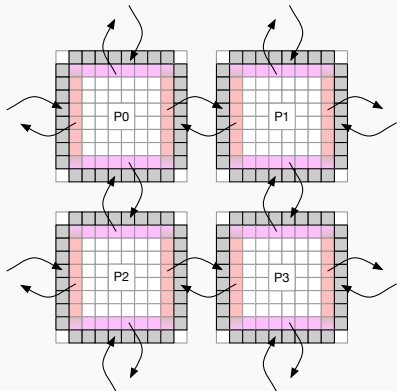


Figure: The 4 regions to be exchanged with the neighbors

Each process needs to exchange 4 regions of data with the neighbors. To facilitate the application writing, these regions should be represented as MPI datatype. Assuming the matrix is stored in row-major format there are only 2 datatype that needs to be created, as 2 ghost regions are identical memory layout with a starting point shifted. Make sure you understand what region of data is sent and what region is received from the Figure 1. Hint: Keep an eye on the color code.



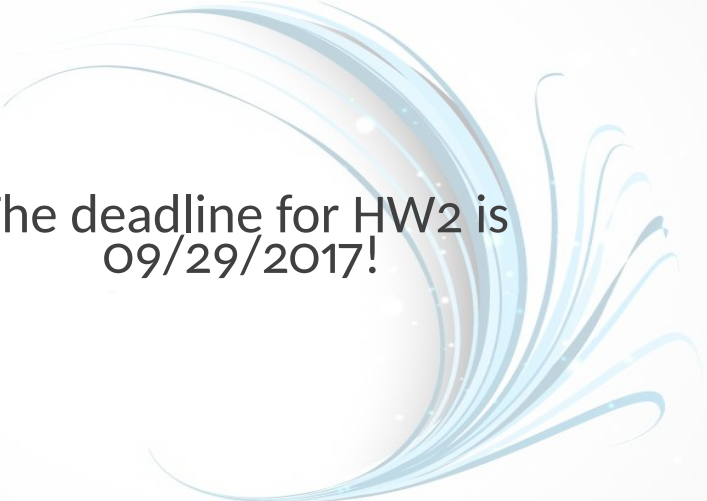
- ▶ Design the communication pattern using blocking communications. Pay attention to ensure no deadlock is possible.
- ▶ Once the version using blocking communications is completed, relax the ordering of message reception and use `MPI_Wait*`.
- ▶ Do you have any overlap between communications and computations? If the answer is no, then you are not yet done with the homework.

Figure: Communication patterns between neighbors



What to check

- ▶ The code provided delivers the correct answer. A correct implementation of the distributed version of the Jacobi should not only remain deterministic, but provide bit-wise reproducibility of the result. Check your result against the provided version (sequential Jacobi).
- ▶ Validate the performance you obtain. The computational part is almost embarrassingly parallel, but its performance is impacted by the cost of the data exchange (the ghost regions) at each iteration. Make sure that, compared with the provided code, your code delivers the expected performance, depending on the number of processes.

A decorative graphic consisting of several overlapping, flowing, wavy lines in shades of light blue and white. The lines curve from the top left towards the bottom right, creating a sense of motion and fluidity. The background is a plain, light gray.

The deadline for HW2 is
09/29/2017!