

MPI
FINAL VERSION
CSI 702- Fall 2010 - Assignment 4
Due: Monday April 5

Overview

For this assignment, you will experiment with several of the code tuning techniques we covered in class. You may write your codes in C, C++, or Fortran. The timing results should be done using the profiling tools for the appropriate compiler. Details for submission will be posted in the final version of this document on-line.

NOTE: Elements of these assignments may be included as part of this exam.

Objectives At the end of this assignment, students will

- Learn how to use MPI to solve practical problems
- Use basic MPI constructs, including sends, recv, broadcasts, etc.
- Compare MPI implementations to serial implementation, and understand the challenges of moving to MPI
- Conduct a performance review to see how well problems are suited to parallel implementations

Requirements:

For this assignment, you will need to

- Complete TWO of the FOUR problems in this assignment.
- You may use C, C++, or Fortran for the serial code.
- You MUST use MPI to parallelize the code.
- You need to create both a serial and parallel version of each code.

continued on next page

- *You need to have a write up for each problem that documents*
 - The parallel approach taken
 - The problems you had during the implementation
 - A quick summary of the MPI commands used
 - A performance analysis of the code, including changing the problem size and the number of nodes that the code runs on
 - Comparisons to the serial code of the performance and ease of implementation
 - Comparisons of the output of the serial and parallel codes. This might include visualization - if appropriate - of the results. Are the results identical?
- Extra credit - complete any of the problems using Sockets and Forks. This code will need to fork to at least four processes, although they will all be on the same machine.

Turning the assignment in For the programs you are turning in, please follow these instructions carefully.

- Create a directory named “youremail-mpihw” where “youremail” is your email account name at gmu.edu
- In this directory, create two subdirectories with the problems you attempt. (Three if you one of the programs using sockets and forks for extra credit.) The directory names should be selected from the following. Only create the directories for the problems you worked on in this assignment.
 - diffusion_mpi
 - manager_mpi
 - nbody_mpi
 - sort_mpi
 - diffusion_fork
 - manager_fork

- nbody_fork
- sort_fork
- Create a clean copy of the source file, the documentation, and Makefile, and any other files into separate directories.
- Create a file using tar and gzip named “youremail-hw5.tgz” that contains the main “youremail-mpihw” directory and its subdirectories.
- Email this file to the instructor.

Grading For this problem, the grades will be based on:

- (45%) The quality of the documentation. How well did you document the speedup, the algorithm, and any problems you used? Did the sample runs show the results clearly?
- (45%) The code - does it compile and run? Does it accomplish the objectives of the assignment? Is it readable, and is there enough internal documentation to make it understandable?
- (10%) Is the assignment turned in following the correct format?

1 Diffusion Equation

The diffusion equation is a simple parabolic partial differential equation that can be solved numerically using a finite difference method.

$$\frac{d\phi}{dt} = \kappa \nabla^2 \phi$$

in one dimension

$$\frac{d\phi}{dt} = \kappa \frac{d^2\phi}{dx^2}$$

The simple finite difference representation of this equation is in one dimension

$$\frac{\phi_{i,j+1} - \phi_{i,j}}{k} = \kappa \frac{\phi_{i+1,j} - 2\phi_{i,j} + \phi_{i-1,j}}{h^2}$$

or

$$\phi_{i,j+1} = \phi_{i,j} + \kappa \frac{k}{h^2} (\phi_{i+1,j} - 2\phi_{i,j} + \phi_{i-1,j})$$

$$\phi_{i,j+1} = \left(1 - 2\kappa \frac{k}{h^2}\right) \phi_{i,j} + \left(\kappa \frac{k}{h^2}\right) (\phi_{i+1,j} + \phi_{i-1,j})$$

where h is the spatial grid size, k is the time-step size, and κ is a diffusion coefficient.

For this problem, let

$$\begin{aligned}x_{min} &= 0 \\x_{max} &= 1 \\ \phi(x_{min}) &= 0 \\ \phi(x_{max}) &= 0 \\ \kappa &= 1 \\ k &= 0.2 \times h^2 \\ t_0 &= 0 \\ t_{final} &= 0.5\end{aligned}$$

Also, let

$$\phi = \begin{cases} 0, & x < 0.4 \\ 1, & 0.4 \leq x \leq 0.6 \\ 0 & x > 0.6 \end{cases}$$

and print the output every 0.05 in time.

For this project

1. Write a simple serial C or Fortran code to solve the problem. Make sure to dump the output every 0.05 time units out. The program should ask the user for the total number of cells to run, then calculate h and k based on the input.
2. Modify the serial code to run on MPI using a one-dimensional spatial decomposition. You will need to communicate with neighbors on every time-step, and gather data for the output.

2 Manager/Worker Mandelbrot Set

The Mandelbrot Set is an iterative map in the complex plane that has chaotic properties. Given an initial location in the complex plane $C = x + iy$ where $i = \sqrt{-1}$, we can go through the following iterations.

$$r_0 = c$$

$$r_i = r_{i-1}^2 + c$$

remember, for complex multiplication

$$(x + iy) \times (x + iy) = x^2 - y^2 + 2ixy$$

After a given number of iterations, many points begin to diverge. Any iteration where $|r_i| > 2$ will leave the complex plane.

To make a Mandelbrot plot, you need to loop through the complex plane from $x = (-2, 2)$ and $y = (-2i, 2i)$. The pixel color at location C is associated with the number of iterations it takes to reach $C > 2$.

It is important to note that some starting points will never reach $C > 2$, so a maximum iteration number is needed in the problem (500 iterations, for example.)

For this problem, you are to make a parallel version of the Mandelbrot problem that uses the manager/worker programming pattern. You cannot use the fine grain scatter decomposition we discussed in the notes for this problem.

1. Write a simple C or Fortran code to calculate the Mandelbrot set. Find the number of iterations needed to exit a given point in the complex plane using the iterative mapping described above. Use a maximum cutoff for the iterations of 500 cycles. You should have the user input the number of points (n) in the x and y direction at the beginning of the code, then divide the sampling points into an $n \times n$ grid.
2. Modify the code to work under MPI using a manager/ work paradigm.
 - (a) To do this, you should designate node 0 to manage the work load.

- (b) Each CPU will be given an initial parcel of work to do.
 - (c) Processor 0 will then listen for communications from the others using asynchronous communication.
 - (d) When any processor finishes its work block, it should send a message to processor 0 asking for another block. At the same time, it should send its results back to node 0 so they can be output to a file. It should go in a receive state to get the message from node 1 about which block to do.
 - (e) Processor 0 will assign the processor a new block, and update the list of things left to do. If there are no blocks left to be calculated, then it should tell the CPU to exit.
3. You will need to experiment with the block sizes and resolutions for a given set of CPU's.

3 N-body Assembly Line

In this problem, you will solve the gravitational potential for a set of particles using an assembly line programming paradigm.

The gravitational potential for a given particle is given by the equations

$$\phi_j = \sum_{i=1, i \neq j}^n \frac{-GM_i M_j}{R_{ij}}$$

So, to calculate the gravitational potential, you need to have a loop over all the particles like

```

for j = 1:n
  potential(j) = 0
  for i = 1:n
    if (i /= j) then
      R_ij = sqrt( (x(i) - x(j))^2 + (y(i)-y(j))^2 + (z(i)-z(j))^2)
      potential(j) = potential(j) -G * M(j) * M(i)/ Rij
    endif
  enddo
enddo

```

To do this in parallel, you will use the “assembly line” programming pattern.

Assume you are running n particles over $nproc$ processors

1. Generate random positions for $n/nproc$ particles for each of $nproc$ files. The file names should be named “particle0”, “particle1”, “particle2”, etc where the number designates the processor ID associated with the file. You can (and perhaps should) generate the particle data with a serial code. If you do the generation in parallel, be aware that you must worry about how the random number generator works on each CPU. Assume the masses for all the particles are 1, and the gravitational constant is 1 as well. For calculational ease, assume the x,y , and z coordinates range from 0 to 1.
2. Write a serial code that reads in all the data files, and calculates the potential for each particle. The code should write out the potentials into a single file, and return the minimum and maximum potential for the entire set of particles.
3. Convert the code to use MPI.
 - (a) The code should read the data files locally off each CPU. This will mean that you need to copy the appropriate data file to each node.
 - (b) After calculating the local contribution of the potential from the particles that are native to the CPU, you need to rotate the particles to the next CPU. After you have received the data from the neighboring processor, update the potential sum by calculating the contribution from each of these particles.
 - (c) Repeat this shifting of particle positions until you have circulated the data through all the CPU’s.
4. Copy the appropriate data file to each CPU you will use in your parallel run.
5. Compare the serial and parallel versions of the code, and be able to present a brief overview of how the code was constructed and how it works in class.

4 Parallel Sort

In this problem, you will sort a list of numbers in parallel.

- Generate a list of n random numbers on processor zero.
- Distribute equal portions of the list to each node. [Note: don't presort the list before distribution - we want to have random samples on each processor to start.]
- Do an internal sort on each list.
- Implement a sample sort or another kind of sort on the list.
- Return the sorted list back to node 1.
- Print the result out and verify it is in ascending order.

Note: the initial list of data and the final list of sorted particles must be located on node zero. The sorting must be done across all of the nodes in parallel.

There are a number of algorithms that work effectively for this. I would encourage you to look at a few algorithms, and then choose one that fits this problem and this architecture. For an initial test, try setting $n = 10^6$ to perhaps 10^8 . You may need to adjust this to get good performance.