

When to Use MPI, OpenMP, and Hybrid MPI+OpenMP

Blue Waters Petascale Institute 2018

Monday, May 28, 4:15pm-4:45pm

Lead Instructor: Colleen Heinemann, NCSA/UIUC

Goals

- Understand what situations it is best to use OpenMP, MPI, or a hybrid MPI+OpenMP model
- Understand why different programming models are better suited for certain scenarios over others

Introduction

Now that we have learned how to program in MPI for distributed memory, OpenMP for shared memory, and a hybrid MPI+OpenMP, we will now take a look at applying such techniques to a variety of problems and decide when it is appropriate to use one technique over another as well as why the chosen programming model is applicable in that scenario.

Exercise #1

For this exercise, we will take a look at several different example problems and determine what programming model (MPI, OpenMP, hybrid MPI+OpenMP) is best for such a situation as well as why.

1. An astronomy model where we want to analyze 16 disjoint galaxies. Calculations will be run on each individual star in each galaxy.

MODEL: _____

WHY: _____

2. We want to calculate the photosynthetic rate of each individual leaf of 10,000 different soybean plants. The photosynthetic rate of one plant does not affect the photosynthetic rate of another, but the photosynthetic rate of one leaf of a plant can affect the photosynthetic rate of another leaf on the same plant

MODEL: _____

WHY: _____

3. We need to analyze the amount of traffic that goes through a neighborhood with 4 intersections to determine whether or not putting traffic lights in would be necessary.

MODEL: _____

WHY: _____

4. We will perform a scalability study on an image processing algorithm to see how well the algorithm scales with really large datasets from material science. The example dataset you will run your strong and weak scalability studies on is 4 TB in size

MODEL: _____

WHY: _____

5. The goal is to build a graph that represents a photograph of a landscape as seen from up above (imagine that it is a picture taken from an airplane). The nodes of the graph will represent different sections of the landscape, but each pixel of the graph requires information about other pixels so that it can determine whether or not the pixels are in the same section of the graph that is being generated. Your image you are building the graph of is 512x512 pixels.

MODEL: _____

WHY: _____

6. The visualization we are attempting to build with VisIt contains 27 billion total elements laid out as a cube. There are 3072 files that VisIt will read in. Each one is a sub-grid of the full cube that contains 192x129x256 cells.

MODEL: _____

WHY: _____

7. We want to run 8 completely different chemical reactions and from the data gathered regarding the chemical reactions, we can then build the output molecular model. Two additional pieces of information are gathered for each chemical reaction: the final amount of the reactant as well as the final amount of the product.

MODEL: _____

WHY: _____

Exercise #2

Now, we want to look at how MPI, OpenMP, or a hybrid MPI+OpenMP model can be used for your projects. To do so, get into groups of 4-5 and provide a brief discussion of your project and gather feedback from your fellow group mates on what they think would be the best model for your project.

After a few minutes, all groups will rejoin and we will ask for some volunteers to give a short description of their project and what model their group decided would be best for them.