

PARALLELISM

- What is it?
- When should you do it?
- Approaches / Type of problems



What is it?

Parallel computing is a type of computation in which many calculations or the execution of processes are carried out simultaneously

Large problems can often be divided into smaller ones, which can then be solved at the same time

In the physical sciences community we typically have two types of parallelism

- 1. Data Parallelism
- 2. Task Parallelism

https://en.wikipedia.org/wiki/Parallel_computing



Types of Parallelization

Job Arrays

- Create a list of independent runs of a program
- Be careful with too much file IO

OpenMP – Open Multi Processing

- Shared Memory System
- Very easy to parallelize a loop
- MPI Message Passing Interface
 - Distributed Memory Systems
 - More difficult to parallelize
 - Requires a change in fundamental programming



Architecture – A pre-requisite for parallel programming

A Single Computer





Serial Programming - i.e. only use one thread







OpenMP Programming – i.e. use all threads on single machine





MPI Programming – Distributed Memory





MPI Programming – Distributed Memory - Serial IO



- Each core gets its own exclusive chunk of RAM
- Communication is needed between threads for them to see each others memory
- All threads can see the hard drive but through one doorway i.e. the file server



MPI Programming – Shared Memory - Serial IO



- All cores can see entire RAM
- All threads can see the hard drive but through one doorway i.e. the file server This can be a major bottleneck in parallel programs



MPI Programming – Shared Memory – Parallel IO



- All cores can see entire RAM
- All threads can see the hard drive and read/write in parallel
 i.e. Lustre file server



MPI Programming – Distributed Memory – Parallel IO



- Each core gets its own exclusive chunk of RAM
- Communication is needed between threads for them to see each others memory
- All threads can see the hard drive and read/write in parallel
 i.e. Lustre file server



MPI + OpenMP – Hybrid Parallelization – Parallel IO



- Each core sees all RAM on its node
- All threads can see the hard drive and read/write in parallel i.e. Lustre file server
- Minimize memory duplication
- Maximize compute



MPI Programming – Distributed Memory – Parallel IO – OffLoading



- Each core gets its own exclusive chunk of RAM
- Off load to GPU or Xeon Phi enabled algorithms for rapid processing
- All threads can see the hard drive and read/write in parallel
 i.e. Lustre file server



Is parallel programming the be-all-and-end-all?





https://wrathematics.github.io/RparallelGuide/

Not so fast





https://wrathematics.github.io/RparallelGuide/

You have to be careful

- Thread Safe?
- Race Conditions?
- Too much file IO?
- Bad domain decomposition?
- Poorly balanced?
- Too much communication?



Key take-home If the library exists – use it!



https://wrathematics.github.io/RparallelGuide/



PARALLELISM

- What is it?
- When should you do it?
- Approaches / Type of problems

When should we do it?

- 1. Too much time!
 - I have a lot of data that takes too much time to read/write
 - Just one of my iterations takes hours to days
 - Multiple iterations take days to weeks!



When should we do it?

- 1. Too much time!
 - I have a lot of data that takes too much time to read/write
 - Just one of my iterations takes hours to days
 - Multiple iterations take days to weeks!
- 1. Not enough memory!
 - I have a lot of data that takes up too much RAM
 - I have to solve large systems of (non)linear equations
 - I have to perform large matrix-vector or matrix-matrix ops



When should we do it?

- 1. Too much time!
 - I have a lot of data that takes too much time to read/write
 - Just one of my iterations takes hours to days
 - Multiple iterations take days to weeks!
- 1. Not enough memory!
 - I have a lot of data that takes up too much RAM
 - I have to solve large systems of (non)linear equations
 - I have to perform large matrix-vector or matrix-matrix ops
- 2. Too much time AND not enough memory!!



When NOT to do it?

- 1. We have poorly written code
 - Trying to navigate Spaghetti Junction in serial is hard enough



http://www.where it is.co.uk/swindon-junctions/magic-roundabout-a-master pi.html



http://urbangreens.tumblr.com/post/80761299/parc-nus-de-la-trinitat-barcelona-this



When NOT to do it?

- 1. We have poorly written code
 - Trying to navigate Spaghetti Junction in serial is hard enough
- 2. We have code that uses slow algorithms that may be replaced

Simple Example: Finding the nearest neighbour in a point cloud

- a. Brute Force is an $O(N^2)$ Operation
 - For each point, compute the distance to all other points
 - Choose the point with minimum distance
- b. Build a KdTree in O(n Log n), and search in O(Log n)

For 1,000,000 points in space, this amounts to

- a. 1x10¹²
- b. $13x10^6 + 13$



When NOT to do it?

- 1. We have poorly written code
 - Trying to navigate Spaghetti Junction in serial is hard enough
- 2. We have code that uses slow algorithms that may be replaced

- 3. Our problem size is too small
 - We can still parallelize the small problem
 - May be slower because of
 - Communication
 - Overhead setting up the parallel region





PARALLELISM

- What is it?
- When should you do it?
- Approaches / Type of problems

Approaches / Types of problems

- 1. Numerical solution of Partial Differential Equations (PDE)
 - Weather
 - Fluid flow
 - Nuclear
- 2. Bayesian Inversion
- 3. Lidar pointcloud processing
- 4. Raster image processing
- 5. Machine Learning
- 6. N-body universe simulations





QUESTIONS?

Leon Foks nfoks@contractor.usgs.gov