

Parallel design patterns

ARCHER course

Practical two: Pipeline for pollution problem



Reusing this material



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

<https://creativecommons.org/licenses/by-nc-sa/4.0/>


This means you are free to copy and redistribute the material and adapt and build on the material under the following terms: You must give appropriate credit, provide a link to the license and indicate if changes were made. If you adapt or build on the material you must distribute your work under the same license as the original.

Acknowledge EPCC as follows: “© EPCC, The University of Edinburgh, www.epcc.ed.ac.uk”

Note that this presentation contains images owned by others. Please seek their permission before reusing these images.


Extending the problem

- The geologists wish to use the pollution calculation code in a more automated, high volume, approach.
- Take some raw measured data and feed this into the calculation code, then generating some final result answer which tells them where about in the pipe the pollution is above a specific threshold and the severity of the pollution.
 - Input, raw, data provided in a directory of files – one for each pipe we are testing



Your task.....


- You are supplied with the logic of each stage, but these are currently unconnected
 - Complete the code so each stage runs on a UE and communicates with neighbouring stages
 - You also need to consider termination via a poisoned pill



- Once you have done this you will calculate the load imbalance
 - And as an advanced exercise look at addressing this


Wash up of practical

- Sample solutions are available
 - MPI P2P messages correspond well to communication between stages
 - For the termination poisoned pill an empty (NULL) message can be sent




- But the stages of the pipeline are heavily imbalanced
 - Not necessarily easy to give lightly loaded stages more work, but can do something to optimise the heavily loaded stage(s)

Let's look at this with Paraver




MPI_Recv

Duplicating the third stage



- All extra UEs make up duplicate stage three.
 - No stage three UEs communicate, but instead work concurrently on different pieces of data
- Fairly simple to do, but termination does require a little more thought

In Paraver



With the LIF what's average?

- $LIF = \text{maximum load} / \text{average load}$
- This tells us how much faster the code could run if the load were perfectly balanced (1.0 being the best.)
- Assume we take the mean (i.e. sum up all values and divide by the number of UEs)
 - But in extreme cases, where we have small amounts of load and one very large value then this can be misleading as the large value pollutes things.
 - Instead the median can sometimes be a better approach

Code	Runtime	Mean LIF	Median LIF
Linear	11.09s	4.95	366
Multiple stage three	0.81s	1.58	13