

Fractals, Task Farms and Load Imbalance: The Outcomes

This is a summary of what should be learned in the Fractals exercise. At the bottom is a set of results which can help with some of the questions posed in the exercise.

1 What We Should Have Accomplished

In the Fractals exercise, the following should have been attempted:

- Accessing a remote HPC resource.
- Obtained and compiled source code.
- Assessed the performance of a task farm across configurations.

2 What We Should Take Away

After attempting the exercise, it is important to reinforce a few key points. **Note:** if anything is unclear or missing, ask the demonstrators who are there to help.

2.1 Key Points

Task Farm, also known as the master/worker pattern, allows a master process to distribute work to a set of workers processes. This is suitable for calculations which can be split up into smaller chunks with little or no inter-worker communications. The pattern can be used for other types of tasks but it complicates the situation and other patterns may be more suitable for implementing. The master process is responsible for creating, distributing and gather the individual jobs. This will often include the collation of the results into an interpretable whole.

Tasks are the units of work in a task farm. These can vary in size and do not have to be of consistent execution time. If execution times are known it can help with load balancing. For a task farm, the tasks should as far as possible be discrete units of work from the overall calculation.

Queues are what the pool of tasks a master generates can be thought of. Each worker will be given a new task when they finish their current task until the pool is empty. If the scheme is to give out tasks as and when requested, then the pool will act like a queue. If there is some filtering of tasks, like known long and short tasks, they may be separated into multiple queues. Some systems will generate the tasks on-the-fly while others will prepare the pool of tasks before starting to distribute tasks and others will take a strategy between these two.

Load Balancing is how a system determines how work or tasks are distributed across processing elements to ensure good performance. Successful load balancing will avoid overloading a single element, maximising the throughput of the system and make best use of resources available. Poor load balancing will lead to elements of the system being under-utilised and reducing the performance as it gets dominated by over-subscribed resources.

The distribution of tasks is generally done in either a pre-planned or responsive mode. In responsive mode, the master will give out tasks as requested by workers. In pre-planned mode, where task lengths are known or estimated, the master will map out a set of tasks for each worker to complete. This allows a clear picture of the overall workload to be developed: most planned systems allocate larger tasks first.

2.2 Other

Cost is increasingly becoming a big factor in how resources are requested and used. Hopefully this exercise will have highlighted that often just throwing more power and resources at something eventually has diminishing returns in performance. This means the monetary cost can be a lot higher for only minor returns over a more efficient and less expensive configuration at run time. Is it necessary to run on 4000 cores, or can the same problem be run on 2000 cores?

Generic patterns and techniques are common to software development and the task farm is a common pattern to recognise. This is commonly use in parallel and high performance computing and while the exercise focus on a fractal problem the underlying framework could be adapted to any problem with similar properties.

3 Sample Results

3.1 Cores, Time, Worker Workloads

Example results for the default image size (768×768 pixels), fixed number of iterations (5000), 64 tasks of the size of 96×96 pixels and varying number of workers :

Workers	Time (s)	Avg. Workload	Max Workload	Min Workload	Load Imbalance Factor
4	1.52	124505763	126124383	122696852	1.01
8	0.96	62252881	77744803	51117022	1.25
16	0.59	31126440	46737752	10968369	1.50
32	0.57	15563220	46114456	67614	2.96
64	0.57	7781610	46089216	9246	5.92

Table 1: Example Run-Times for different number of workers and their Avg/Min/Max Workloads.

3.2 Time, Task Size, Load Imbalance

Example results for the default image size (768×768 pixels), fixed number of iterations (5000), fixed number of workers (16) and varying number of tasks :

Number of Tasks (Task Size)	Time (s)	Load Imbalance Factor
16 (192×192)	1.93	5.034
64 (96×96)	0.59	1.501
256 (48×48)	0.43	1.108
4096 (12×12)	0.4	1.017
36864 (4×4)	0.4	1.003
147456 (2×2)	0.47	1.017
589824 (1×1)	0.80	1.006

Table 2: Example execution Times for 16 workers and varying number of Tasks.