

# F20DP Distributed and Parallel Technology

## Assessed Coursework 1

### Evaluating Low-level Parallel Programming models

#### Purpose

In this coursework you will develop parallel programming, measurement, and technology evaluation skills. Additionally, there needs to be some reflections on the advantages and dis-advantages of the chosen parallel programming technologies. The parallel machines are nodes in the [Robotarium](#) high-performance cluster at Heriot-Watt. You will develop and measure parallel versions of a program using one of the following low-level parallel programming technologies: C+MPI or OpenCL. Finally you will compare the performance and programming models of these technologies.

#### Sequential Program

The program that should be parallelised is the computation of the *sum of Euler totient computations* over a range of integer values. The *Euler totient function* computer, for a given integer  $n$ , the number of integers that are relatively prime to  $n$ , i.e.  $\Phi(n) \equiv |\{m \mid m \in \{1, \dots, n\} \wedge m \perp n\}|$ . Two numbers  $m$  and  $n$  are relatively prime, if the only integer number that divides both is 1. To test this, it is sufficient to establish that their greatest common divisor is 1, i.e.  $m \perp n \equiv \gcd mn = 1$ . Thus, the task for this program is: for a given integer  $n$ , compute  $\sum_{i=1}^n \Phi(i)$ .

The following C code is a direct implementation of the above specification, as starting point for the parallel algorithm: <http://www.macs.hw.ac.uk/~hwloidl/Courses/F21DP/srcs/TotientRange.c>

#### Organisation

The **assessed coursework work is to be carried out in pairs**, and you should choose your own partner. **Each member of the team chooses one of the following technologies**, implements a parallel version of the program in this technology and evaluates the performance of this parallel version:

- C with MPI for explicit message passing;
- C with OpenCL for off-loading computations to a GPGPU;

*Together* you can share ideas about how to parallelise the code and how to tune the parallel performance. However, it needs to be clearly stated who implemented the parallel version using which technology. *Together* you should prepare a comparative report, matching the structure given below. It is relatively easy to produce a simple parallelisation of both programs, however *additional marks are available for thoughtful sequential and parallel performance tuning*.

Deadline: 3:30PM on Friday 24<sup>th</sup> of February, 2017 (Week 7) (**pair project**)

Tools that can help you, and have been discussed [on these slides](#), are `gprof` and `cachegrind` for sequential C. For plotting parallel performance results `gnuplot` is recommended. Further help on these tools is available through the [external links section of the main course information page](#).

## Deliverable

The deliverable is a **report** (including sources) with the structure below. The deliverable is due on Friday 24<sup>th</sup> of February, 2017 (Week 7). Reports not following this structure, or not containing all of the specified results, will be penalised. Submission should be through Vision (sub-menu Assessment) and as hardcopies in the drop-box near the School Office.

### Structure of the report:

#### Section 1: Introduction (4 marks).

This should give a short summary of the task to implement and parallelise, describe the environment it is performed in, the parallel technologies used, and the learning objective of this coursework.

#### Section 2: Sequential Performance Measurements (4 marks).

Run the sequential version of the programs and analyse the performance, possibly using some of the tools mentioned above. Discuss the sequential performance of and possible improvements to these programs. Of interest are in particular hotspots in the program and good cache usage. *max 1 A4 page*

#### Section 3: Comparative Parallel Performance Measurements (12 marks).

You should measure and record the following results in numbered sections of your report. The measurements are based on these inputs:

- DS1: calculating the sum of totients between 1 and 15000.
- DS2: calculating the sum of totients between 1 and 30000.
- DS3: calculating the sum of totients between 1 and 100000.

Runtime measurements should be the middle (median) value of three executions. N.B. For comparison purposes the performance of the different systems must be reported on the *same* graph. You may also plot other graphs to show interesting features, or use larger numbers of cores.

**Section 3a: Runtimes.** Measure the runtime of the sequential C program for DS1 and DS2. Plot **three** runtime graphs

- DS1: execution times for the sequential C program, the MPI parallel program on 1,2,3, ..., 64 cores, and the runtime on a GPU accelerated system.
- DS2: execution times for the sequential C program, the MPI parallel program on 1,2,3, ..., 64 cores, and the runtime on a GPU accelerated system.

Deadline: 3:30PM on Friday 24<sup>th</sup> of February, 2017 (Week 7) (**pair project**)

- DS3: execution times for the sequential C program, and the MPI parallel program using 64,  $2 \times 64$  .., up to the maximal number of cores available and the runtime on a GPU accelerated system.

**Section 3b: Speedups.** Plot **three** relative speedup graphs corresponding to the runtime results for DS1, DS2 and DS3 showing the ideal speedup and the speedups for both parallel programs on 1,2,3, .. 64 cores. Recall that relative speedup is calculated using the runtime of the parallel program on a single core.

**Section 3c** A table summarising the sequential performance and the best parallel runtimes of both parallel programs.

**Section 3d** A discussion of the comparative performance of both parallel implementations. This part should also include a discussion about how the different execution platforms can be reasonably compared. Try to identify a suitable platform-independent measure of performance and derive those figures from your experiments.

*max 1 A4 page*

#### **Section 4: Programming Model Comparison**

*(6 marks).*

An evaluation of the parallel programming models, specifically for the totient application. You should indicate any challenges you encountered in constructing your programs and the situations where each technology may usefully be applied.

The comparison should be based on the TotientRange application, and focus on the technology in general, and be supported by technical arguments.

*max 1 A4 page*

#### **Section 5: Reflection on Programming Models**

*(0 marks).*

Not applicable.

#### **Appendix A and B**

*(10 marks each).*

For each parallel implementation, an appendix should give the listing of your parallel TotientRange program, clearly labelled with the *single author's name*. Also include a paragraph, and possibly diagram(s), identifying the *parallel paradigm* used, and *performance tuning approaches* used.

*Additional marks are available for thoughtful sequential and parallel performance tuning.* *(4 marks)*

#### **Notes**

- Complete the C with MPI and OpenCL Lab exercises before starting the coursework.
- Check [these slides](#) for practical tips on performance measurements.
- [Check the FAQ](#) for more info on C+MPI etc usage and cluster configuration.
- Graphs and tables must have appropriate captions, and the axes must have appropriate labels.
- The Robotarium cluster uses a batch-job system to give you exclusive access to the cluster when you need it for measurements. Follow [this link to familiarise yourself with the batch-job system](#).