**CE816-7-AU (Example)**

**UNIVERSITY OF ESSEX**

Example Examination

---

**High-Performance Computing (Example Examination)**

---

Time allowed: **TWO** hours

Candidates are permitted to bring into the examination room:

Calculators (hand-held, containing no textual information)

The following items are provided:

Graph paper (available on the invigilator's desk)

The paper consists of **FOUR** questions.

Candidates must answer all questions.

All questions are of equal weight.

The percentages shown in brackets provide an indication of the proportion of the total marks for the **PAPER** which will be allocated.

**Please do not leave your seat unless you are given permission by an invigilator.**
**Do not communicate in any way with any other candidate in the examination room.**
**Do not open the question paper until told to do so.**
**All answers must be written in the answer book(s) provided.**
**All rough work must be written in the answer book(s) provided. A line should be drawn through any rough work to indicate to the examiner that it is not part of the work to be marked.**
**At the end of the examination, remain seated until your answer book(s) have been collected and you have been told you may leave.**

## Question 1

HTCondor is a widely-used framework for performing computation on a group of networked computers.

(a) Explain the basic approach that HTCondor uses for identifying suitable machines and distributing computation. **[8%]**

(b) A particular application is to be distributed across a group of identical computers under HTCondor. However, 20% of its computation cannot be distributed and must be carried out centrally. Determine the maximum speed-up that can be achieved if 50 computers are available to HTCondor. **[9%]**

(c) If a speed-up of 5 needs to be obtained, discuss whether this can be achieved. **[8%]**

## Question 2

MPI, the *Message Passing Interface*, is commonly used for high-performance computing on clusters. One of the standard example programs distributed with MPI concerns the computation of $\pi$ from the recurrence relation

$$\int_0^1 \mathrm{d}x \frac{4}{1+x^2} = \pi$$

(a) Describe a strategy for splitting this problem up into one that can be used with MPI. **[10%]**

(b) Write an MPI program in commented pseudo-code that implements your approach. **[15%]**

Marks will be awarded for the algorithm and the correct use of MPI calls.

## Question 3

(a) When describing parallel computing systems, the acronyms SISD, SIMD and MIMD are commonly used. Explain what these terms mean and explain which of these descriptions applies best to computation using OpenCL. **[8%]**

(b) A simulation of the motion of a cloud of charged particles in an electrical field is being developed. The force experienced on each charged particle by all the other charged particles is able to be experienced at any distance. Explain why this might be problematic when programming the simulation in OpenCL. **[8%]**

(c) Despite the difficulties alluded to in (b), OpenCL-based implementations of these type of simulations do exist. Describe a strategy that allows the difficulties to be overcome. **[9%]**

## Question 4

The approach known as *MapReduce* has become popular for processing large amounts of data on clusters and clouds of computers.

(a) Explain the two principal functions of MapReduce. **[6%]**

(b) MapReduce works only when the overall problem has certain properties. Explain what these properties are and discuss what type of parallelism MapReduce can provide. **[10%]**

(c) Apache's `Hadoop` framework is based around an implementation of MapReduce. Discuss what additional functionality a framework such as `Hapdoop` must provide in order to produce an effective and efficient software framework. **[9%]**

**END OF PAPER CE816-7-AU (Example)**