

PGT Exam Performance Feedback

2016/2017 Semester 1

COMP60611 Parallel Programs and their Performance

Graham Riley

Comments Questions 1,2 and 3 were roughly equally popular with 8, 7 and 6 answers respectively. Only one student tackled question 4. The average mark was approximately 68%.

For Question 1a, (overheads), the most popular omissions were not to point out that overhead values can added to give the total overhead, and the fact that overheads are generally positive.

For 1b, very few students produced a table to quantify the overheads (as the question required). No one provided a column of the naive ideal times from which a total overhead column could be computed.

In Question 2, not many students produced clear iteration space diagrams, with the result that explanations were frequently lacking precision. Also, the potential for false sharing was generally ignored in choosing a chunk size for the suggested scheduling solution.

Also in Question 2, in the definition of Amdahl's law, few people pointed out that an Amdahl curve smoothly goes to an asymptotic limit (equal to the performance level of the sequential fraction of code) and that the maximum speedup cannot exceed P (so the curve is always below the naive ideal line).

In Question 3, no one pointed out that with data parallelism there are generally data dependencies that have to be managed (this being the main feature that distinguishes data parallelism from "pleasingly" parallel algorithms). Also, few people pointed out that reduction parallelism explicitly involves some sequential work.

Also in Question 3, part e (multi-stack implementation), often, the analysis did not clearly explain the implications of initialising the separate stacks and work-stealing between them for overheads
