

Real Time and Embedded Systems: Problem Set 1

Dr. Colin Perkins

10th January 2008

The first week of the module has outlined the concepts behind real time and embedded systems, introduced some terminology, and provided an overview of some real time scheduling techniques. This problem set seeks to test your understanding of that material, and your ability to draw new conclusions based on material provided. You should answer all questions.

Question 1: Briefly describe what is a real time system, and explain the difference between hard and soft real time systems. Discuss whether the distinction between hard and soft real time systems is always clear cut. [4 marks]

Question 2: We considered several priority-driven scheduling algorithms. It was noted that these algorithms make *locally optimal* decisions about which job to run, but the resulting schedules are often not globally optimal. Explain the difference between locally and globally optimal, and discuss why priority-driven scheduling algorithms typically do not produce globally optimal schedules. [4 marks]

Question 3: It was noted that priority-driven scheduling is difficult to validate, since *scheduling anomalies* can occur. Explain what is a scheduling anomaly, and outline a situation when one might occur. [2 marks]

Question 4: The lecture and book outline a proof that earliest deadline first (EDF) scheduling is optimal on a single processor, as long as pre-emption is allowed and jobs do not contend for resources (section 4.6 of Liu's book). Using a similar method, prove that least slack time (LST) scheduling is optimal, and explain the constraints under which it is optimal. [5 marks]

This problem set is worth 5% of the mark for this module. Your answers must be submitted by 2:00pm on 21st January 2008 via the locked box outside the Teaching Office. You must include your pink declaration of authorship form with your submission. Any late submission will be awarded zero marks unless accompanied by a valid special circumstances form.