EEE4084F Quiz 1: Lectures 1-4

This quiz is for marks! 6 Mar 2014

PLEASE ANSWER ON A SEPARATE PAGE

MAKE SURE YOUR STUDENT NUMBER IS ON YOUR ANSWER PAGE(s)

TOTAL NUMBER OF QUESTIONS : FIVE

#	Question	/	time (mins)
Q1	UML is an effective means to concisely model systems. This question relates to the following model of a computer system, where the PCBox is the fundermental part of this system. 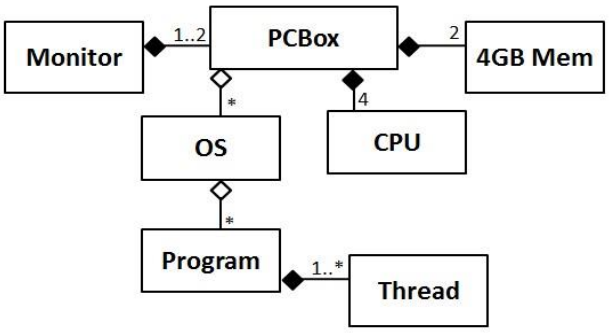 (a) What type of relation is it that connects the CPU element to the PCBox element? Is it aggregation or composition? (b) According to this model, would the following statement be true or false? "A PCBox can work with 4GB of Mem" (c) Provide motivation as to why a connector with an open diamond, instead of one with a coloured in diamond, is used to connect Operating System to Program. (d) In this system, what can be said about the amount of data that a program needs in order to operatate? (e) There's at least one inaccuracy in this diagram considering the reality of how a standard PC system is configured. Identify an inaccuracy and explain what it should rather show in order to more closely represent a standard PC. [5 x 2 marks each = 10]	10	9.0
Q2	(a) Explain the acronmy HPEC, as is used frequently in the textbook [1 mark]. Provide some insight into why, for some applications such as radar signal processing and radio astronomy, there is a move from the standard concept of HPC application development, which uses multiprocessors and programming languages such as Java, towards HPEC and a broader variety of languages. [4 marks] (b) Broadly speaking, a computation can be implemented on any one of the following types of architecture: 1) A totally hardware solution; 2) A solution implemented on a reconfigurable computer, and 3) A software solution running on a general purposes computer. Explain each of these three architecture tyes and provide one advantage and one disadvantage of implementing a computation on each one of these. [3 x 2 = 6 marks]	10	9.0
Q3	**These question relates to Seminar 1.** (a) The Berkeley's Landscape of Parallel Computing suggests processor architectures follow a design of using thousands of processor cores per chip. Provide an explanation of why this approach could be effective for massively parallel system and what sort of savings could be gained from such an approach. [4 marks] (b) The paper discusses the concept of a 'DWARF', and mentions 13 of these. Explain what is meant by a DWARF and why this could protentially be an effective strategy for building parallel applications. Provide an example of a DWARF. [6 marks] (c) What is meant by the terms Old Convension and New Convention. Give an example of one old convention and a corresponding new convention in terms of parallel computer system design. [4 marks]	14	12.6

Q4	The processors used for HPC nowadays are sometimes considered “power hungry”. a) Discuss some of the power concerns for processors nowadays and why in many cases it is so important to 'keep a lid' on the amount of power that processors draw - whether they are used in supercomputing centers or in other types of system. [3 marks] b) In terms of Moore's Law, there is a trend in terms of the number of transistors in a cirucit doubling about every 18 months. Does this trend still hold in terms of recent processors? Provide a short motivation for your answer. [3 marks] c) Discuss the similarities and/or differences between Moore's Law (i.e. in the classic sense of increasing transistor use), compared to trends in terms of power consumption of processors over the past few decades. You are welcome to provide an graph to illustrate your answer. [4 marks]	10	9.0
Q5	Consider the following scenario: There exists a computationally intensive algorithm X that can be partitioned into three parallel operations X1, X2 and X3. X1 and X2 are independent of each other but use the same input data, a few 100 MBytes; but both operations only return results that are under 1Kbyte of data. X3 is dependent on the results of both X1 and X2, but also needs access to the same input data. The result of X3 is only a few bytes which can for example be displayed to a terminal. Respond the the following: If this system were to be used on a cluster comprising two dualcore machines, each with loads of memory and harddrive space, recommend the way you would go about implementing the system, highlighting relvant programming model(s) to use. Provide clear motivation for your suggested strategy? You may provide a diagram. [6 marks]	6	5.4
Q6	*Optional Extra Question for a bonus mark:* What was the MFLOPs/s per watt rating for the Intel Paragon system. Choose the closest option below: (i) <20 (ii) 30-40 (iii) 60-90 (iv) 65–320 (v) 200 - 100,000 (vi) 100,000 - 1,000,000	+1	
	TOTAL :	50	45

You can use this space for rough work if you like: