Halmstad University
School of Information Science, Computer and Electrical Engineering
Bertil Svensson, CC-lab

Mobile: 070 648 0843

3 pages

Written Exam
Embedded Parallel Computing (DO8003)

March 20, 2009

Closed book exam.

Aids allowed: A dictionary, translating between English and your native language, or an English Thesaurus. It must be in printed form, not electronic!

Welcome to the exam!

READ THIS FIRST:

In most of the assignments in this exam you are asked to "describe" or "explain". Please write short and concise. It is not necessary that you cover everything that can possibly be connected to the topic – it is more important that what you write is clear, coherent, correct and relevant.

The facts that you find in the course literature are not "the law". If you have different opinion, don't be afraid to mention it. Of course, you are also free to take up things that are not to be found in the course literature, as long as they are of importance to the topic.

Please, read the assignments carefully, so that you give answer to the correct questions – and to all questions!

Good luck!

-- Bertil

Number of assignments: 5
Maximum points: 30
Bonus points from the seminars (maximum 16) will be added to the points of this written exam.

Required points:	22	30	38
Grade:	3	4	5

Grades may be raised based on excellent seminar achievements.

Assignment 1: Contrast Concepts (3 p)

For each of the following concept-pairs, give a clear and concise explanation of **the significant difference**. (Just a couple of lines, NOT a long description of each concept).

a) Instruction level parallelism vs. data parallelism
b) Superscalar vs. VLIW architecture
c) crossbar vs. multi-stage interconnection network

Assignment 2: SIMD Processing (6 p)

Assume that one field in the memory of a SIMD computer with *N* PEs stores integer values (as shown in the figure). The computer has *N* one-bit PEs, each with a one-bit wide interface to the memory.

Further, assume that we want to read out the numbers as a sorted list, starting with the largest number.

a. Describe (in words or in some kind of "pseudo-code") how this is done in the most efficient way.

b. Does the speed of the read-out (measured in numbers per second) depend on the value of *N*? Why?/Why not?

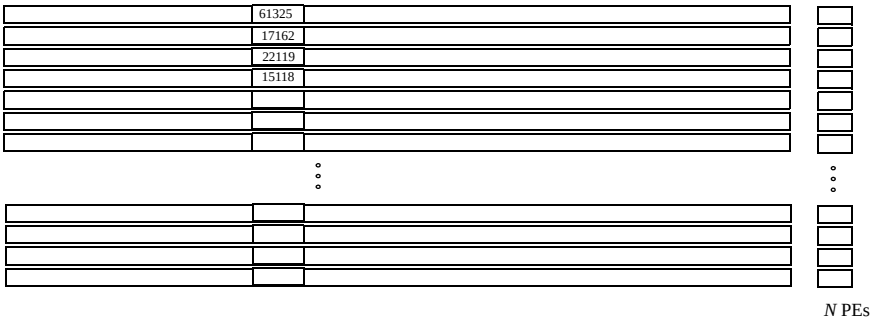

Assignment 3: Cache Coherence in Distributed Shared Memory Architectures (6 p)

In Centralized Shared Memory Architectures using a bus as interconnection network, caches are kept coherent using the snooping-based technique. On Distributed Shared Memory Architectures (also known as NUMA) with other interconnection networks, this is, however, not possible.

a) Explain why these architectures are called NUMA.
b) Explain why the snooping technique cannot be used on these architectures.
c) Describe a technique for cache coherence that can be used on these architectures. In particular, describe where the caching information is stored and what this information consists of.

Assignment 4: Static Interconnection Networks (9 p)

a) Determine the condition under which a binary tree with k levels has a larger diameter than each of the following: (i) an n-dimensional hypercube, and (ii) an $r \times r$ 2D mesh.

b) For the three kinds of networks mentioned above calculate the number of nodes as well as the diameter when $k = 10$, $n = 16$, and $r = 12$.

c) When choosing a static network for a multicomputer, the diameter is one of the criteria to consider. Mention at least two more and discuss the tree, the hypercube and the 2D mesh in the light of these criteria.

Assignment 5: Speedup and Efficiency (6 p)

Amdahl's law states the maximum speedup that can be achieved with n processors when the fraction f of the entire sequential work cannot be parallelized, while the rest can.

a) Show the inequality that relates f and n if a 25% efficiency is to be achieved. Assume that there is no communication overhead.

b) For the following two cases, determine whether or not 25% efficiency can be achieved:
(i) $f = 0.1$, $n = 10$
(ii) $f = 0.1$, $n = 100$

--- END ----------