

ST: Parallel and Distributed Computing

Final Exam

Thursday 14 December 2000
Open Book / Open Notes

Instructions: Answer all 9 questions, each on a separate sheet of paper. Write your name on each page, and staple together the answers to questions (1, 7, 8, & 9), (2 & 3), and (4, 5, & 6). You have roughly 15 minutes for each question.

1. MPI and PVM are two popular packages for message-passing in a distributed environment. Summarize the major differences between these two packages.
2. Analyze the divide-and-conquer method of assigning one processor to each node in a tree for adding numbers (section 4.1.2) in terms of communication, computation, overall parallel execution time, speedup, and efficiency.
3. Asynchronous computing using message-passing (e.g., clusters) and programming with shared memory are similar in some ways, and yet very different. Identify the major difference between these two paradigms for asynchronous parallel computation and indicate the advantages/disadvantages of each.
4. Discuss the following claims:
 - a) Dr. Flynn was wrong
 - b) Complexity analysis is wrong
5. Assume that the basic measure of efficiency is the number of data bits processed per instruction. Is this good or bad for SIMDs or MIMDs? Why or why not?
6. Pertaining to ASC/ACE:
 - a) Why are structure codes used in associative computing? How are they used?
 - b) What is the result of the following code (i.e. show the contents of age, wage, rate and time after execution):

```
bonus = 5.0;
if age[$] .gt. 50 then
    wage[$] = rate[$] * age[$];
else
    wage[$] = (rate[$] + bonus) * time[$] + wage[$];
```

given:

age	wage	rate	time
25	33	5.50	40
37	17.7	6.00	42
53	12.5	7.30	37
61	205	3.90	52

7. Which of the following machines -- the Illiac IV, the CM-2, the CM-5, or the Cray T3D -- could most easily be modified to support associative computing, and why? Describe the architectural changes that would be necessary to support associative computing.
8. Consider the Illiac IV and the CM-2, both SIMD machines. The Illiac IV has a 2D grid interconnection network, while the CM-2 has a grid that can be configured in multiple dimensions as well as a hypercube interconnection network. For what type of applications would each of these interconnection networks be best suited, and why?
9. Most SIMD machines use custom processors while most MIMD machines use commercial CPUs. Explain why this is the common practice, yet may not always be the best choice.