

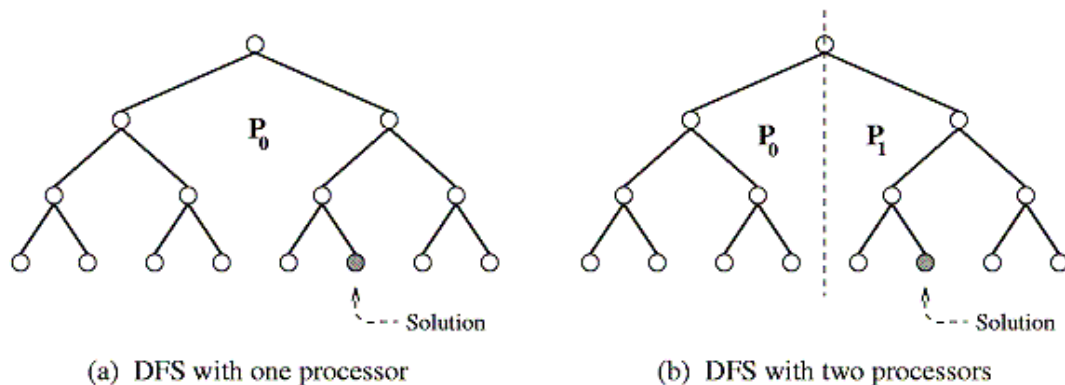
**CS546 “Parallel and Distributed Processing”  
Homework 1 (Due Jan. 31, 2012)**

---

**Submission:**

- **Late penalty: 10% penalty for each day late**
  - **Please upload your assignment to “Digital Drop Box” on the Blackboard with the following name: CS546\_SectionNumber\_LastName\_FirstName\_HW1. Please do NOT email your assignment to the instructor and TA!**
- 

1. What are the two main styles of parallelism? Explain.
2. What are the two main types of locality? Explain.
3. Define the terms of NUMA, Data Intensive computing and Cloud computing.
4. Assume for a given machine, the cache 1 latency is 3 ns and the cache 2 latency is 20 ns, the main memory latency is 70 ns, and remote memory is 5 us (microsecond), and assume in average the cache 1 hit ratio is 50%, cache 2 hit ratio is 70%, and main memory hit ratio is 95%. Compute the average memory access time assuming there is no concurrent memory access, nor latency hiding.
5. Discuss the difference between SIMD and MIMD machines. What is SPMD model? Discuss the advantages and disadvantages of SIMD and SPMD. Which category is MPI belong to?
6. Discuss the difference between shared address space machines and distributed address space machines. And discuss the advantages and disadvantages of both architectures.
7. Consider the search tree shown in the following figure, in which the dark node represents the solution.



- (a) If a sequential search of the tree is performed using the standard depth-first search (DFS) algorithm, how much time does it take to find the solution if traversing each arc of the tree takes one unit of time? Note: DFS begins by expanding the initial node and generating its successors. In each subsequent step, DFS expands one of the most recently generated nodes. If this node has no successors (or cannot lead to any solutions), then DFS backtracks and expands a different node.
  - (b) Assume that the tree is partitioned between two processing elements that are assigned to do the search job, as shown in figure b. If both processing elements perform a DFS on their respective halves of the tree, how much time does it take for the solution to be found? What is the speedup? Is there a speedup anomaly? If so can you explain the anomaly?
8. Current computers may combine different levels of parallelism. Please describe three levels of parallelism.
9. Read the lecture note of lecture 5 and go through all the examples

**Note on cheating: There are penalties for cheating. Don't find out the hard way. Working in groups is fine for discussing approaches and techniques. Copying problem solutions or code is cheating. Both the person copying and the person giving the answers will be equally penalized. Make sure you do your own work.**