

Course ISE 431: Distributed Information Systems

Recitation 1 Exercise

Michael J. May

October 20, 2009

Answer the following questions:

1. Name two advantages and two disadvantages of distributed systems over centralized ones.
2. What is the difference between a multiprocessor and multicomputer?
3. The terms *loosely coupled system* and *tightly coupled system* are often used to describe distributed computer systems. What is the difference between them?
4. What is the difference between a MIMD computer and a SIMD computer?
5. Crossbar switches allow a large number of memory requests to be processed at once, giving excellent performance. Why are they rarely used in practice?
6. A multicomputer with 256 CPUs is organized in a 16×16 grid. What is the worst case delay (in hops) that a message might have to take?
7. Now consider a 256 CPU hypercube. What is the worse-case delay here, again in hops?
8. What is meant by a *single system image*?
9. What is the main difference between a distributed operating system and a network operating system?
10. What are the primary tasks of a microkernel?
11. Name two advantages of a microkernel over a monolithic kernel?
12. Concurrency transparency is a desirable goal for distributed systems. Do centralized systems have this property automatically?
13. An experimental file server is up $\frac{3}{4}$ of the time and down $\frac{1}{4}$ of the time, due to bugs. How many times does this file server have to be replicated to give an availability of at least 90%?
14. Suppose you have a large source program consisting of m files to compile. The compilation is to take place on a system with n processors, where $n \gg m$. The best you can hope for is an m -fold speedup over a single processor. What factors might cause the speedup to be less than this maximum?