

Systems Comprehensive Exam, Spring 2005

January 11, 2005

1 Instructions

This is a closed-book, closed-notes exam with a total of 100 points. You may not use any external source for answering these questions. Please direct any questions about this exam to Professor Bridges. Professor Bridges may be reached either in person in his office in 301B Farris, by phone at 277-3032 or 363-8798, or by email at bridges@cs.unm.edu. Turn your exam in to Professor Bridges or the front office by 5:00 PM MDT on Tuesday, January 11, 2005. Exams *will not* be accepted after this time except by prior arrangement with Professor Bridges.

Type or write your answers to the stated number of questions in each of the following three sections. Make any *reasonable* assumptions necessary to answer the question, but be sure to state any assumptions that you make.

2 Short Answer - Answer 3 of 4 (30 points)

Briefly answer 3 of the following 4 questions. Your answer should be no longer than *one* paragraph.

1. **Caching.** Compare and contrast the use of write-through and write-back caches in multiprocessor systems.
2. **File Systems.** What operations do log-structured file systems attempt to optimize? What is the difference between the locality optimized by log-structured file systems and more traditional UNIX file system implementations (e.g. Berkeley FFS)?
3. **TCP/IP.** Consider the state transition diagram for TCP connections shown in figure 1. Why is the TIME WAIT state needed? Why has the TIME WAIT state become a problem for Web servers?

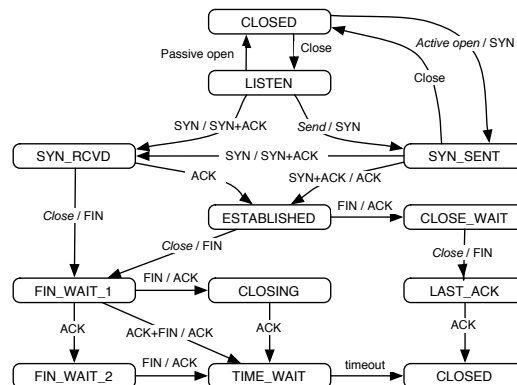


Figure 1: TCP Connection States

4. **Consistency.** What is the fundamental difference between weakened consistency models such as weak consistency, entry consistency, and release consistency and more traditional consistency models such as sequential and PRAM consistency?

3 Medium Answer - Answer 2 of 3 (40 points)

Provide detailed answers to two of the following three questions. Be sure to state any assumptions you make and to fully justify your answers. Limit your answers to approximately one to two pages in length.

1. **Addressing Modes.** Processors designed for embedded applications (e.g., the ARM) do not typically have support for paged address translation, using other techniques to ensure address space isolation. For example, in base and bound addressing, the application uses a base register to denote the base of its current addressing range and offsets from that base as virtual addresses. Discuss the design tradeoff that motivates this decision.
2. **Locking.** Compare and contrast the advantages and disadvantages of spin-locks and blocking locks in multi-threaded programs. A study of lock acquisition times has shown that they are highly bi-modal—locks are generally either acquired quickly or after a very long wait. Describe in detail a solution that takes advantage of this distribution to achieve the advantages of both spin locks and blocking locks.
3. **Operating System Structure.** Compare and contrast the following general approaches to OS design: monolithic systems, monolithic systems with loadable modules, microkernels, and exokernels/library operating systems.

4 Design - Answer 1 of 2 (30 points)

Provide a *full* and *detailed* answer to one of the following two questions. Be sure to state any assumptions you make and to fully justify your answer.

4.1 Multiprocessor Scheduling

Consider the scheduler for an operating system for a shared memory multi-processor. In such systems, the scheduler functionality can be distributed in a variety of ways. Describe in detail the advantages and disadvantages of:

- A single global scheduler
- Individual per-processor schedulers that interact with a global scheduler that assigns jobs to processors.
- A hierarchy of interacting processor schedulers.

In writing your answer, be sure to consider:

- The size of the machine (which can range from two processors up to several hundred processors)
- Application characteristics (e.g. job lengths, job-level parallelism, and synchronization behavior)
- Operating system parameters (e.g. length of scheduling quantum)

4.2 Intelligent Networks

Suppose that you have a small router (10-12 port pairs) that is capable of actively processing every packet that comes through the router. That is, the router can execute on the order of 1000 instructions per packet (in addition to routing the packet). In addition, the router has a small amount of memory (on the order of a couple of megabytes) that can be used to store state information. As an example, the router could examine an incoming packet and, based on a value in the packet, decide that the packet should be dropped, a counter incremented, and several new packets constructed and sent on different output ports.

How might you use such a router in the implementation of a parallel file system where one of the port pairs will be used to connect the parallel file system to the outside world and the remaining port pairs would be used to connect systems that provide the processing and storage space needed for the file system? Pay particular attention to how you would partition functionality between the switch and the systems connected to the switch.