

OS Comprehensive Exam
Spring 2007

1. Both SPIN and the exokernel focused on providing safe system extensibility. One characterization of these two systems is that SPIN provided these features by pushing type-safe extension code into the kernel, while the exokernel focused on making hardware memory protection cheap enough to expose all critical hardware services to the user level.

a) Do you agree with this description? Why or why not?


b) Explain what safety properties each system guaranteed of its extensions?

c) Imagine you are using SPIN and exokernel today. Assuming that each system works as described, would you feel comfortable installing and using a third-party file compression extension off the Internet? Why or why not?

2. Answer the following questions about Rosenblum and Ousterhout's "The Design and Implementation of a Log-Structured File System" paper.

- a) Describe the tradeoff in deciding how many segments to clean at once. I.e., what are the pros and cons of cleaning many segments at the same time? Few segments?
- b) Why does a 'hot-and-cold' file access pattern cause the greedy cleaner to clean segments at a higher average utilization?
- c) What was the purpose of the 'cost-benefit' cleaning policy?
- d) With the cost-benefit policy, which segments are cleaned with low utilization? High?
- e) Now suppose there was a new workload, 'cold-and-hot', which contained 90% hot files and 10% cold files (where 'cold' and 'hot' are defined as before). How would the performance of 'LFS Cost-Benefit' compare to the original two workloads?

3. Consider the three processes p1, p2 and p3 executing the following sequences of code:



The arrow in each column indicates which instruction in the corresponding process is currently executing. For example, p1 is executing P(z) and p2 is executing code before V(y). All semaphores are initially set to 1.

a) Explain why, given the programs above, each semaphore x, y and z can each be thought of as a resource with a single unit, and where P on that semaphore requests the resource and V on the semaphore releases the allocated resource.

b) Draw a resource allocation graph describing the current state of the system consisting of the three processes p1, p2 and p3 and the three resources x, y, and z. Using this resource allocation graph, show that the system is deadlocked.

c) Explain how this program could not have satisfied a hierarchical locking policy.

d) Explain how this deadlocked state would not have occurred had the processes used the Banker's Algorithm.

4.a. Compare RAID levels 1 and 5 storage schemes; discuss advantages and disadvantages for both. Suppose level 5 group contains n disks.

4.b.I. For r small reads per second directed to the group, what is the average workload per disk in both the fault free and one failed disk modes.

4.b.II. For w small writes per second directed to the group, what is the average workload per disk in both the fault free and one failed disk modes.