

Due in class on Monday 21 October 2002

1. Why is a context switch between threads faster than a context switch between processes?
2. As discussed in class, when a signal operation is performed on a semaphore, if there are processes waiting, one is woken up. Generally the waiting processes are kept in a first-in first-out (FIFO) queue, which gives a type of semaphores called “strong semaphores”. “Weak semaphores”, in contrast, are removed from the queue in arbitrary order. Are there any disadvantages in weak semaphores?
3. If a signal operation is performed on a semaphore, and the corresponding wait operation is performed later, the thread performing the wait operation does not actually wait. Explain.
4. (Exercise 15.9 from OSC) What are the advantages of using...
5. In multilevel feedback queue CPU scheduling, why doesn't the mechanism for reducing the priority of a process cancel out the aging mechanism?