COSC 6360 MIDTERM OCTOBER 23, 2002

This exam is closed book. You can have one sheet of notes. Cheaters will be expelled from UH. Answers that are not specific to the question being asked will not be read.

1. Spring treats the thread issuing a cross-domain call and all downstream threads as a single *scheduling entity* called a *shuttle*. Why does it not merge them instead into a single thread within the calling domain? (2×5 points)

Merging the threads into a single thread managed by the calling domain would cause problems in two cases:

a) when the domain being called resides on another machine, and

b) when the calling domain wants to suspend a thread and the thread is inside a critical section of the target domain.

2. In their presentation of the SPIN system, Bershad *et al*. argue that is difficult to build a system that is extensible, remains safe and achieves a good performance.

a) Explain why systems that are safe and extensible tend to have a poor performance. (5 points)

Most systems that are safe and extensible use microkernels and let user-level servers handle most system calls. As a result these system calls require four context switches instead of two.

b) Give one example of such systems (5 points):_Mach_____

3. Why are guarded command essential to CSP (5 points) and to CSP only? (5 points)

CSP use blocking receives and requires the receiver process to specify the name of the sender process in its receive (<*source_process*>?<*target value*>). Other message passing mechanism use indirect naming and require the receiver process to specify a mailbox or a port.

4. What are the advantages and disadvantages of *partial subbblocking* in terms of (a) efficient use of the TLB memory and (b) efficient use of the main memory? 2×5 points)

Partial subblocking uses the TLB memory in a more efficient fashion because it only requires one PPN per TLB entry. On the other hand, it requires all physical pages within the same subblock to occupy contiguous locations in main memory. Hence accessing even a single page within a subblock forces the virtual memory to allocate enough physical memory space for all the pages in the subblock.

5. How can Mach implement *kernel-supported threads* without having a special system call to create such threads? (10 points)

Mach creates kernel-supported threadsthrough the same mechanism it creates regular processes. The sole difference is that the inheritance attribute of the data segment of the parent process is then marked shared.

6. You are given a multithreaded C program that was running on a shared memory multiprocessor and asked to port it to a cluster of workstations managed by Munin.. Needless to say, you want each thread to run on a separate workstation. What should you do with:

a) The variables that are shared by several threads? (5 points)

Declare them to be shared data variables.

b) The variables that are local to each thread?(5 points)

Leave them as they are.

c) The mutex semaphores used by the program to control access to the shared variables? (5 points)

Declare them to be shared synchronization variables.

d) The mutex calls surrounding the program critical sections ? (5 points)

Replace them by calls to the appropriate Munin library procedures.

7. Babaoglu and Joy decided to limit the *speed* at which the hand of the clock sweeps the circular list of active pages to 300 pages/sec.

a) Why did they do that? (5 points)

To guarantee that the software simulation of page referenced bits will never take more than 10% of the total CPU time.

b) Could you give another reason to limit the speed of a clock hand? (5 points) (**Hint**: It is not in the paper.)

A clock policy whose hand is moving too fast will cycle too quickly through the main memory and expel pages that are still in use.

8. You are to write a program to let students submit their assignments on Bayou by copying them into a subdirectory of the TA account. Your solution should guarantee the confidentiality of the programs before and after they are submitted. It cannot use electronic mail nor a client/server pair. Which UNIX security mechanism should you use? (5 points) What would be the ownership and access rights of your program? (5 points)

The best solution is to use the UNIX set user_ID bit. The program should be owned by the TA, writable by her only and executable by the students.

9. How does Totem order messages that are created within the same ring? (10 points)

Totem uses a circulating token that contains the sequence number of the last message that was sent on the ring. Processes that want to send a message must wait for the arrival of the token. They can then assign an new sequence number to their message and update the token.