

Department of Computer Science, National Tsing Hua University
Ph.D. Qualification Examination
Operating Systems, Fall 2017

- (10%) (a) (3%) What is the purpose of interrupt? (b) (2%) How does an interrupt (signal) differ from trap? (c) (5%) Explain the steps of handling an interrupt in details.
- (13%) (a) (3%) What is the purpose of TLB? (b) (5%) Illustrate how TLB works. (c) (5%) Given a two-level paging, let the memory access time and TLB access time be 300ns and 40ns, respectively. If the TLB hit ratio is 90%, what is the effective memory access time?
- (12%) All modern OS systems implement both paging and segmentation for address translation, and support virtual memory. (a) (3%) Explain the difference between paging and segmentation. (b) (3%) Why do segmentation before paging? (c) (3%) Under such memory management design, what is the cause of "segmentation fault" and "page fault"? (d) (3%) Who (programmer/OS) should handle those faults, respectively?
- (10%) Recall the various deadlock detection and prevention algorithms we've discussed in this course, and consider the following snapshot of a system with five processes (P1, P2, P3, P4, P5) and four resources (R1, R2, R3, R4). There are no current outstanding queued unsatisfied requests.

Currently Available Resources

R1	R2	R3	R4
2	1	2	0

Process	Current Allocation				Max Need				Still Needs			
	R1	R2	R3	R4	R1	R2	R3	R4	R1	R2	R3	R4
P1	0	0	1	2	0	0	3	2	0	0	2	0
P2	2	0	0	0	2	7	5	0	0	7	5	0
P3	0	0	3	4	6	6	5	6	6	6	2	2
P4	2	3	5	4	4	3	5	6	2	0	0	2
P5	0	3	3	2	0	6	5	2	0	3	2	0

- (5%) Is this system currently deadlocked, or can any process become deadlocked? Why or why not? If not deadlocked, give an execution order.
- (5%) If a request from a process P2 arrives for (0, 1, 2, 0), can the request be immediately granted? Why or why not? If yes, show an execution order.

5. (10%) Short questions. Brief reasoning is required.
- (a). (3%) What needs to be saved and restored on a context switch between two threads in the same process?
 - (b). (3%) Why would two processes want to use shared memory for communication instead of using message passing?
 - (c). (4%) [True/False] Shortest remaining time first (or called shorting remaining job first) is the best preemptive scheduling algorithm that can be implemented in an operating system.
6. (15%) The “H2O” problem. You have just been hired by Mother Nature to help her out with the chemical reaction to form water, which she does not seem to be able to get right due to synchronization problems. The trick is to get two H atoms and one O atom all together at the same time. The atoms are threads. Each H atom invokes a procedure *hReady* when it is ready to react, and each O atom invokes a procedure *oReady* when it is ready. For this problem, you are to write the code for *hReady* and *oReady*. The procedures must delay until there are at least two H atoms and one O atom present, and then one of the threads must call the procedure *makeWater* (which just prints out a debug message that water was made). After the *makeWater* call, two instances of *hReady* and one instance of *oReady* should return. Your solution should avoid starvation and busy-waiting. You may assume that the semaphore implementation enforces FIFO order for wakeups—the thread waiting longest in *wait()*, which is also known as *P()*, is always the next thread woken up by a call to *signal()*, which is also known as *V()*.
- (a). (5%) Specify the correctness constraints. Be succinct and explicit.
 - (b). (10%) Write down the pseudo implementation of *hReady* and *oReady* using semaphores.
7. (15%) (a) (10%) Please explain how the LRU algorithm works. (b) (5%) Please prove that the LRU algorithm is a stack algorithm.
8. (15%) (a) (7%) Please describe how the memory mapped I/O works. (b) (8%) Please discuss the advantage and disadvantage of memory mapped I/O.