

CS 414 Assignment 5

3 points per answer for 3a-g, 10 points per answer for 4a-c, 5a-b, 12 points for 2 and 14 points for 1 (100 max). Due Wednesday April 05.

1. In TCP slow start, at first the sending host will double the number of packets it sends every time it receives all acknowledgements for the previous packets it sent. In other words, the sending host will send one packet, receive the ACK for that packet, send two packets, receive the ACKs for those packets, send four packets, and so on. This exponential growth continues until there is a packet loss. (Subsequent to this loss, the sending host will increase the number of packets linearly, not exponentially.) Consider the scenario where a host H1 has two connections, one to a host HL on the same LAN, and one to a host HW thousands of miles away. The one-way latency to host HL is 1 micro-second, the one-way latency to host HW is 50 milli-seconds and the round-trip latency is twice the one-way latency. Assuming no packet losses, and assuming near-infinite bandwidth and processing speeds at all hosts, how long does it take to send the first 1023 packets to each host?
2. The operating system uses segmentation with paging. The virtual address is 32-bits long. Consider the scenario where a process has 6 segments with bases and limits given below:

	Base	Limit (base 16)
0	0e457000	0100
1	2a983400	8000
2	7c567000	1040
3	0e468800	4008
4	0c945c00	0810
5	0e756900	1000

Assume a page size of 1KB (1024 bytes). What is the size of the page table (number of entries) and how much space is lost to internal fragmentation?

3. Consider a demand-paging system with the following time-measured utilizations:
 - CPU utilization 20%
 - Paging Disk 97.7%
 - Other I/O defices 5%

For each of the following, say whether it will (or is likely to) improve CPU utilization. Explain your answers.

- (a) Install a faster CPU
 - (b) Install a bigger paging disk
 - (c) Increase the degree of multiprogramming
 - (d) Decrease the degree of multiprogramming
 - (e) Install more main memory
 - (f) Install a faster hard disk or multiple controllers with multiple hard disks
 - (g) Add prepaging to the page-fetch algorithms
 - (h) Increase the page size
4. Consider the page reference string shown below. Assuming that $\Delta = 3$ (for the demand paging algorithms, Δ is the memory size; for working set algorithms Δ is the working set window size), fill in the table for the (a) LRU and the (b) WS algorithms. Time starts at 0 and we've filled in the first three entries to warm start the algorithms.

(a) (LRU)

R(t)	1	2	5	3	2	6	3	1	5	3	2	5	3	1	4	1	5	4	5	4	3
S(t)	1	1	1																		
		2	2																		
			5																		
In(t)	1	2	5																		
Out(t)	$\emptyset$	$\emptyset$	$\emptyset$																		

(b) (WS)

R(t)	1	2	5	3	2	6	3	1	5	3	2	5	3	1	4	1	5	4	5	4	3
S(t)	1	1	1																		
		2	2																		
			5																		
In(t)	1	2	5																		
Out(t)	$\emptyset$	$\emptyset$	$\emptyset$																		

(c) For each of the two algorithms, compute the hit ratio starting from time $t = 3$ (skip the part that we filled in for you).

5. When modifying the Linux kernel late one night, Doug accidentally modifies the context switching software so that sometimes, when context switching from process P0 to process P1, the system forgets to flush the TLB and main-memory (L2) cache.

- (a) What sorts of problems might result from this mistake?
- (b) Doug realizes what he has done, but in fixing it, modifies Linux to flush the TLB and cache not just on every context switch operation, but also on every interrupt (including page faults, system calls, device interrupts, etc). What impact would you expect this to have on the system?