

CS 5565 Midterm

This is a closed-book, closed-internet, closed-cellphone and closed-computer exam. However, you may refer to your sheet of prepared notes. Your exam should have 8 pages with 4 questions totaling 64 points. You have 75 minutes. Please write your answers in the space provided on the exam paper. If you unstaple your exam, please put your initials on all pages. You may use the back of pages if necessary, but please indicate if you do so we know where to look for your solution. You may ask us for additional pages of scratch paper. You must submit all sheets you use with your exam. However, we will not grade what you scribble on your scratch paper unless you indicate you want us to do so. Answers will be graded on correctness and clarity. You will lose points if your solution is more complicated than necessary.

Name (printed) _____

I accept the letter and the spirit of the Virginia Tech graduate honor code – I have not given or received aid on this exam.

(signed) _____

#	Problem	Points	Score
1	Reliable Data Transmission and TCP	20	
2	Congestion Control and DCCP	16	
3	Implementation Issues	20	
4	Know Your Current Events	8	
	Total	64	

1 Reliable Data Transmission and TCP (20 pts)

a) (8 pts) Layering vs. Performance.

Name two *performance-enhancing* techniques TCP implementations must employ that directly result from the fact that these implementations reside in a *layered architecture*. For each technique, describe the relationship to layering!

i. (4 pts) Technique #1

ii. (4 pts) Technique #2

b) (4 pts) Autotuning

Windows Vista added support for TCP autotuning. According to Microsoft, Vista “tunes the TCP receive window size based on the bandwidth delay product (BDP) and the rate at which the application reads data from the connection.” Microsoft claims to have observed throughput improvements of up to 10x.

Under which circumstances would you expect such throughput improvements to occur? Be specific!

c) (4 pts) Download Accelerators

Some so-called download accelerators purport to increase download speeds by opening multiple TCP connections to download files (for instance, by simultaneously retrieving different byte ranges using HTTP's Range: headers) Define why and under which circumstances a user would benefit from these accelerators!

d) (4 pts) USB

The USB specification supports a number of modes for data transfer, including one mode that provides reliable data transfer (this mode is used, for instance, by mass storage devices such as USB hard drives or flash drives). However, rather than using a more general sliding window protocol with adjustable window sizes, a USB host controller uses a single "data toggle synchronization" bit to implement a simple alternating bit protocol in hardware.

Why do you think the designers of USB made this choice, despite the known weaknesses of using a single bit only?

2 Congestion Control and DCCP (16 pts)

a) (8 pts) Modeling congestion control

Bansal and Balakrishnan in [MIT-LCS-TR-806] introduce the following generalization of the additive increase, multiplicative decrease congestion control model used in TCP into a binomial model that uses two parameters k and l to describe a class of congestion control schemes.

$$I: w_{t+R} \leftarrow w_t + \alpha/w_t^k; \alpha > 0$$

$$D: w_{t+\delta t} \leftarrow w_t - \beta w_t^l; 0 < \beta < 1$$

α and β are constants in this model. w_t is the window size at time t . I stands for Increase after the successful receipt of an acknowledgement, D stands for Decrease after a loss event. Consider TCP's slow start and congestion avoidance phases. What values are used for exponents k and l and for α and β ?

	k	l	α	β
TCP in Slow Start				
TCP in Congestion Avoidance				

b) (4 pts) DCCP

What motivation led the developers of DCCP to propose to implement congestion control as a new layered protocol, rather than leaving its implementation to individual applications?

scratch space

4 Know Your Current Events (8 pts)

Conficker C Worm

As of today, April 1, 2009, the Conficker C worm has emerged as a potential threat to stability on the Internet. This worm, which by some estimates has infected over 11.9M machines, is programmed to contact a sample of 50,000 randomly generated domain names every day in order to receive updates for its code. An alliance of researchers and industry, the Conficker Cabal, has formed to combat the worm. This group recently decoded the algorithm the worm uses to compute these domain names based on the current date and time.

One possible approach to fight the worm would be to reprogram the root DNS servers to redirect these domain names away from the servers that would provide these potentially malicious updates. Since there are only 13 distinct root servers (presented by a total of 169 IP-anycast-based replicas), this could be done relatively cheaply.

Would that approach work? Justify your answer!