

THE UNIVERSITY OF CHICAGO

DEPARTMENT OF COMPUTER SCIENCE

WINTER QUARTER 2003

CSPP 54001

A COURSE IN NETWORKS

I. FOSTER, INSTRUCTOR

MIDTERM COMPREHENSIVE EXAMINATION

This is an in-class, closed-book and no-notes test.

Thursday, February 6, 2003

TO ENSURE ACCURATE GRADING PLEASE MIND YOUR HANDWRITING.

Show your reasoning clearly. If your reasoning is correct, but your final answer is wrong, you will receive most of the credit. If you just show the answer without reasoning, and your answer is wrong, you may receive no points at all.

In the box below enter your name and mailslot #:

Full name:	mslot #:
------------	----------

READ THE FOLLOWING STATEMENT ABOUT ACADEMIC HONESTY
AND ACKNOWLEDGE IT BY SIGNING YOUR NAME IN THE SPACE BELOW IT

I am aware of the University of Chicago's academic honesty policy that states:

As students and faculty of the University of Chicago, we belong to an academic community with high scholarly standards of which we are justly proud. Our community also holds certain fundamental ethical principles to which we are deeply committed. We believe it is contrary to justice, to academic integrity, and to the spirit of intellectual inquiry to submit the statements or ideas of work of others as one's own. To do so is plagiarism or cheating, offenses punishable under the University's disciplinary system. Because these offenses undercut the distinctive moral and intellectual character of the University, we take them very seriously and punishments for them may range up to permanent expulsion from the University.

Sign in the space above.

Round-Trip Time (20 points)

(a) If the Round-Trip Time (RTT) from Chicago to Sydney, Australia, is 250ms and all links in the network have a 100M bits/second bandwidth, how much data can fit in the "pipe" in bytes?

$$(250 \text{ ms}) \times (100 \text{ Mbits/second}) = 25 \times 10^6 \text{ bits} = 3,125,000 \text{ bytes.}$$

(b) Imagine that now you want to send data directly to a geostationary satellite. If the satellite is 36,000 km away, what is the RTT in this case? Assume the speed of light in the void/air is 3×10^8 m/s.

$$2 \times (36,000 \text{ km}) / (3 \times 10^8 \text{ m/s}) = 240 \text{ ms}$$

(c) Imagine that now you want to send data to the Oort Cloud at the same data rate. If the distance to this cloud (which is on the border of our Solar System) is 150×10^9 km, what is the RTT?

$$2 \times (150 \times 10^9 \text{ km}) / (3 \times 10^8 \text{ m/s}) = 1 \times 10^6 \text{ seconds} = 277.78 \text{ hours.}$$

General (20 points)

(a) We examine the buffers of a particular router and find that, on average, they contain 1000 packets. We also measure the average time that a packet is delayed by the router, and find it to be 10ms. Assuming that no packets are dropped by the router, what is the average arrival rate?

The answer to this question requires material on queuing theory that we did not cover in class. Thus, I've given everyone full marks.

The answer is as follows:

By Little's Result, we know that $1000 \text{ packet avg. queue occupancy} = (10 \text{ ms avg. waiting time}) \times (\text{avg. arrival rate})$
 $\hat{O} \text{ avg. arrival rate} = 100,000 \text{ packets/second.}$

(b) You are given a communication link that transmits R bits per second. The objective is to transmit a file of length F bits. The bits are sent in packets that are all P bits long. Each packet contains a checksum of length C bits, where $C < P$. (I.e. the length P includes the length of the checksum.) Consecutive packets must be separated by at least G seconds. Write down an expression for the total time to transmit the file as a function of R , F , P , C and G .

Each packet has $(P-C)$ bits of useful information. To transmit F bits of data, we need

$$\lceil F/(P-C) \rceil = N \text{ packets}$$

(we can assume the last packet will be padded out to the required P bits).

If there are N packets, then there are $N-1$ inter-packet delays of at least G seconds. Thus the total transmission time is

$$\geq (N-1)G + N \times (\text{packet transmission time}).$$

The time to transmit a packet is P/R . Thus, total time is

$$\begin{aligned} &\geq (\lceil F/(P-C) \rceil - 1)G + \lceil F/(P-C) \rceil (P/R) \\ &= \lceil F/(P-C) \rceil (P/R + G) - G. \end{aligned}$$

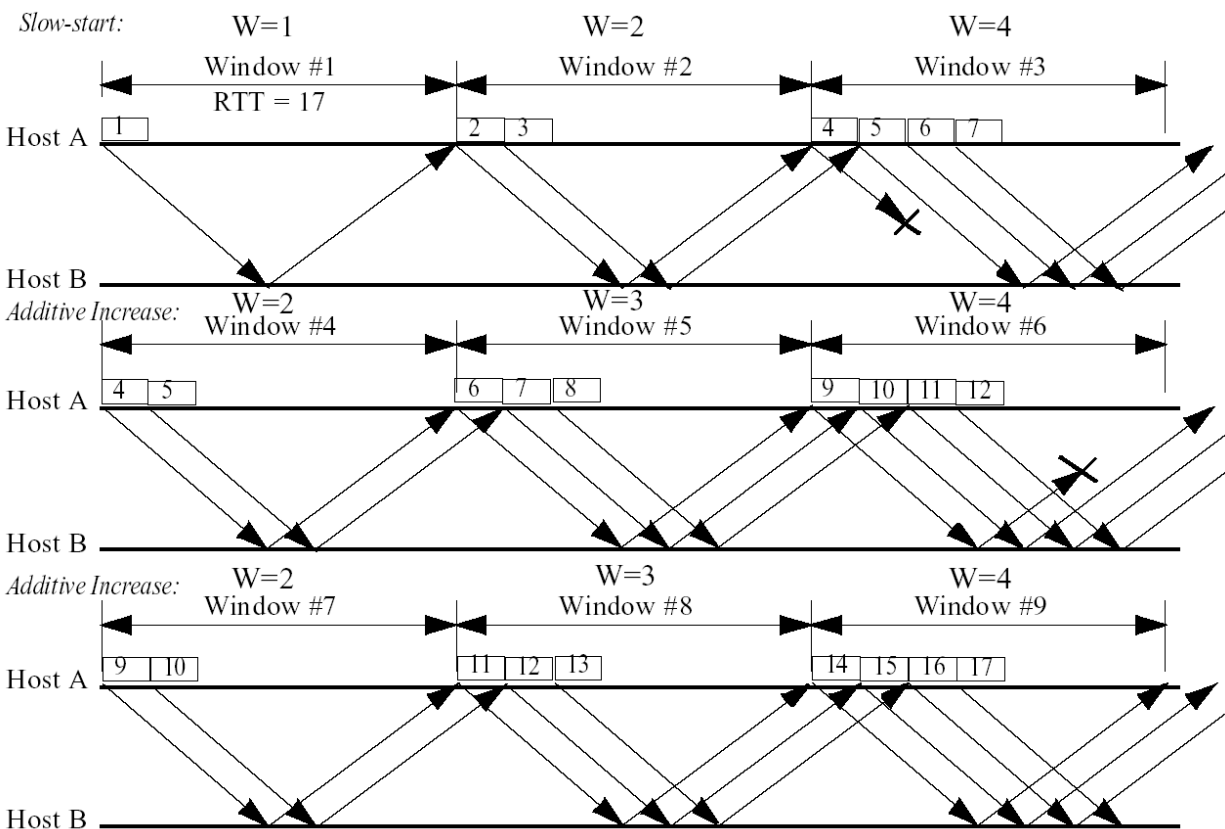
TCP Congestion Control (20 points)

We will use simplified rules for TCP congestion control:

1. **Initial window size.** The window size is measured in packets, not bytes. We'll assume that all packets are the same length. Initially, the window size is 1 packet.
2. **Slow-start.** The connection is in slow-start from the time the connection starts until the first time a packet is dropped. Each time an acknowledgement is received, the window size is increased by one packet.
3. **Multiplicative decrease.** Each time a packet (or acknowledgement) is dropped, the window size is halved.
4. **Additive increase.** When not in slow-start, the window is increased by one packet when acknowledgements have been received for all the packets in the previous window.
5. **Fast Retransmission and Fast Recovery.** This mechanism is not used in our simplified version of TCP.

Host A has 100 packets to send to Host B using TCP. The two hosts have established a TCP connection, and Host A has sent its first data packet, as shown in the figure below. We will assume that the RTT always equals the time taken by Host A to transmit 17 packets, and that the retransmission timeout is set to exactly RTT. In other words, there is no variability in the RTT and Host A retransmits unacknowledged packets after exactly RTT. We will further assume that the receiver has an infinite buffer size and so never invokes the flow-control mechanism, and that the source uses "Go-Back-N" when retransmitting packets. "Packet 4" is lost the first time it is transmitted; and when "Packet 9" is transmitted for the first time, its acknowledgement is lost. No other packets or acknowledgements are lost in the first 10 windows.

Draw the sequence of packets that are sent by Host A up to and including the 9th window. Clearly mark each packet with its sequence number (starting from 1), showing which packets are transmitted (or retransmitted) in each window. Show that Packets 1,...,17 have been sent by the end of Window #9.



Network Characterization (20 points)

(a) What characteristic(s) distinguish(es) a small world network from a random network?

A small world network has a large clustering coefficient; a random network with the same number of nodes and links has a significantly smaller clustering coefficient.

(“Regular networks are highly clustered, i.e., there is a high density of connections between nearby nodes, but have long path lengths, i.e., to go from one distant node to another one must pass through many intermediate nodes.

Random networks are highly un-clustered but have short path lengths. This is because the randomness makes it less likely that nearby nodes will have lots of connections, but introduces more links that connect one part of the network to another.

Small world networks have a "clustering coefficient" that is similar to that of a regular graph, but the "characteristic path length" is less than 20% of what it would be for a regular graph.”)

(b) Give two examples of small world networks.

Many: citation graph, actor graph, power grid, neural network of nematode worm.

(c) Discuss how small world structure could be exploited when designing search algorithms.

Key notion is that we preferentially use the highly connected nodes when forwarding queries.

Public Key Cryptography (20 points)

Consider two participants, A and B, with private and public keys A_{priv} , A_{pub} , B_{priv} , and B_{pub} . Explain what steps would be followed when:

(a) A wants to send a confidential message M to B.

A computes $M' = E(M, B_{\text{pub}})$

A sends M' to B

B receives M' from A

B computes $M = D(M', B_{\text{priv}})$

(b) B wants to send a confidential message M to A.

B computes $M' = E(M, A_{\text{pub}})$

B sends M' to A

A receives M' from B

A computes $M = D(M', A_{\text{priv}})$

(c) A wants to authenticate a message M that it sends to B.

A computes $M' = E(M, A_{\text{priv}})$

A sends M' to B

B receives M' from A

B computes $M = D(M', A_{\text{pub}})$

(d) A wants to authenticate a confidential message M that it sends to B.

A computes $M' = E(E(M, B_{\text{pub}}), A_{\text{priv}})$

A sends M' to B

B receives M' from A

B computes $M = D(D(M', A_{\text{pub}}), B_{\text{priv}})$