

# Data Networks, Summer 2007

## Homework #1

Assigned May 16, 2007  
Due May 21, 2007 in class

Name: \_\_\_\_\_


Email: \_\_\_\_\_

Student ID: \_\_\_\_\_

Problem	Points
1	
2	
3	
4	
Total	

In this homework, for simplicity, **K** means  $10^3$ , **M** means  $10^6$ , and **G** means  $10^9$ . But read the green box on page 44-46 of Peterson & Davie to understand the precise meaning of these notations in different contexts. Show your work when appropriate, do not just write down a single number as your answer if you want full credit.

Problem 1 (25 points)– Consider the following network path with 3 links and two store-and-forward packet switches. The speeds and the delays of the links are as indicated in the picture. Packets flow from left to right. Note that each switch has 1.5MB of memory for doing store-and-forward.



- (a) Assume that the path is completely empty initially, what's the total time required to transmit a 500KB packet across this path (i.e. the time from the first bit is injected until the last bit emerges)? (6 points)

Total time = propagation delay + transmission delay + queuing delay + processing delay  
 Since the path is empty, queuing delay is zero. We assume processing delay is also zero.

$$\text{Total time} = 40\text{ms} + 5\text{ms} + 10\text{ms} + 500\text{KB} \left( \frac{1}{10\text{Gbps}} + \frac{1}{2\text{Gbps}} + \frac{1}{2\text{Gbps}} \right) = 59.4\text{ms}$$

- (b) Assume that someone has just finished injecting a 500KB packet into the first link before you start transmitting your 500KB packet. In this case, what is the total time required to transmit your 500KB packet across the path? (7 points)

Since a 500KB packet is ahead and the link speed becomes slower after the first hop, queuing delay will be incurred at the first switch.

Queuing delay at first switch is  $500\text{KB} / 2\text{Gbps} - 500\text{KB} / 10\text{Gbps} = 1.6\text{ms}$

Queuing delay at second switch is 0 because link speeds and sizes of the 2 packets are equal

So, total time now is  $59.4\text{ms} + 1.6\text{ms} = 61\text{ms}$

- (c) Given that other arbitrary network traffic can exist to compete for transmission, if a 500 KB packet is successfully transmitted across the path, what is the maximum (worst case) time it could have taken? (6 points)

Two answers may be provided based on two different assumptions about the network:

Assumption 1: traffic can only be injected at the ingress of the network (first link). In this case, the maximum delay is incurred when we inject a 1.1MB packet (P1) just before transmitting the 500KB packet (P2). This is because, by the time P2 has completely entered router 1, exactly  $(500\text{KB} / 10\text{Gbps}) * 2\text{Gbps} = 0.1\text{MB}$  of P1 will have been transmitted by router 1, leaving exactly 1MB of data ahead of P2 in the queue, thereby leading to the maximum possible queuing delay for P2 of  $1\text{MB} / 2\text{Gbps} = 4\text{ms}$ . The queuing delay at router 2 will be  $0.6\text{MB} / 2\text{Gbps} = 2.4\text{ms}$ . So the total delay for P2 will be  $(59.4 + 4 + 2.4)\text{ms} = 65.8\text{ms}$ .

Assumption 2: traffic can be generated everywhere (even at routers). Thus we can incur the maximum queuing delay of  $1\text{MB} / 2\text{Gbps} = 4\text{ms}$  at every router.

Thus the total delay will be  $(59.4 + 4 + 4)\text{ms} = 67.4\text{ms}$

- (d) What's the maximum number of bits of data this network path carries at any moment? A MP3 file is roughly 4MB, how many MP3 files can the path carry at any moment? (6 points)

This is just the delay-bandwidth product plus all the memory space.  
So, it is  $10\text{Gbps} \cdot 40\text{ms} + 2\text{Gbps} \cdot 5\text{ms} + 2\text{Gbps} \cdot 10\text{ms} + 3\text{MB} = 56.75\text{MB}$   
That's roughly 14.2 MP3 files just flowing along in the network!

Problem 2 (30 points) – (This is adopted from Peterson & Davie p.59, #25 (a) and (b)) Assume you wish to transfer an n-byte file along a path composed of the source, destination, five point-to-point links, and four switches. Suppose each link has a propagation delay of 1 ms, bandwidth of 3 Mbps, and that the switches support both circuit and packet switching. Thus you can either break the file up into 1 KB packets, or set up a circuit through the switches and send the file as one contiguous bit stream. Suppose that packets have 24 bytes of packet header information and 1000 bytes of payload, that store-and-forward packet processing at each switch incurs a 1 ms delay after the packet has been completely received, that packets may be sent continuously without waiting for acknowledgements, and that circuit setup requires a 1 KB message to make one round-trip on the path incurring a 1 ms delay at each switch after the message has been completely received. Assume switches introduce no delay to data traversing a circuit. You may also assume that file size is a multiple of 1000 bytes.

- (a) For what file size n bytes is the total number of bytes sent across the network less for circuits than for packets? (15 points)

Packet switching: packets pay an ongoing per packet cost of 24 Bytes for a total count of  $1024 \times (\text{number of packets}) = 1024 \cdot n / 1000$  bytes

Circuit switching: circuits pay an up-front penalty of 1KB being sent across the network twice for a total data count of  $2000 + n$  bytes

So the question really asks how many packet headers does it take to exceed 2000 bytes, which is  $\sim 83.34$ .

Thus, for files longer than  $\sim 83,340$  bytes, using packets results in more total data sent.

- (b) For what file size  $n$  bytes is the total latency incurred before the entire file arrives at the destination less for circuits than for packets? (15 points)

Packet switching:

$$\begin{aligned}\text{Total transfer latency} &= (\text{transmit delay for all pkts at the source}) \\ &+ (\text{per-pkt transmit delay introduced by each switch}) \\ &+ (\text{per-pkt processing delay introduced by each switch}) \\ &+ (\text{propagation delay for all links}) \\ &= (c*t) + (x*t) + (x*0.001) + ((x+1) * 0.001) \\ &= 2.731*10^{-6}n + 0.01992\end{aligned}$$

$c$ : number of packets =  $n/1000$

$t$ : transmission time of a packet =  $1024*8\text{bit}/3\text{Mbps}$

$x$ : number of switches = 4


$b$ : bandwidth = 3Mbps

Circuit switching:

$$\begin{aligned}\text{Total transfer latency} &= (\text{transmit delay for the whole file at source}) \\ &+ (\text{propagation delay for all links}) \\ &+ (\text{setup time for the circuit}) \\ &= 8n/b + ((x+1)*0.001) \\ &+ 2*(1000*8/b + (x*1000*8/b) + (x*0.001) + ((x+1)*0.001)) \\ &= 2.667*10^{-6}n + 0.04967\text{seconds}\end{aligned}$$

Solving  $2.731*10^{-6}n + 0.01992 > 2.667*10^{-6}n + 0.04967$   
 $\Rightarrow n > 465\text{KB}$  (rounding to the next 1000 bytes)

Problem 3 (25 points) – Abstractly, a router (or switch) has some input ports and some output ports. When a packet arrives at an input port, it is forwarded to the correct output port for outbound transmission. A router typically has memory buffers for queueing packets that are waiting to be forwarded or transmitted. In one type of router, called input queueing router, the memory buffers exist only at the input side. A queue is called FIFO (first-in-first-out), if the packets leave the queue in the order they arrived. The figure below depicts a simple input queueing router with a FIFO queue at each input port. The small numbered boxes inside the FIFO queues denote packets destined for the corresponding numbered output port. Assume that the transmission paths indicated by the arrows have the same transmission speed. Assume also that the input elements can coordinate to eliminate collisions at the output port



- (a) This simple input queueing router design is known to achieve very poor utilization. That is, the output ports are often not kept busy transmitting as much as possible. Explain why you think this design has such a low utilization problem. (10 points)

The reason is that when packets from different inputs are waiting for the same output (e.g. the two packet “3”s, one has to wait), other packets behind a waiting packet must also wait despite their destination outputs may be idle (e.g. packets “7” behind “3”). It has been shown that under assumptions, a FIFO input queued router will reach a utilization of only 58%. This is called the Head-of-Line Blocking problem.

- (b) How would you change the design of the input side of the router to reach higher utilization? A critical objective here is that the time it takes to insert or remove a packet from memory should not depend on the number of packets in memory. Drawing your design may help with your explanation. (15 points)

The solution to this problem is to use  $N$  FIFO queues at each input port, with each queue corresponding to an output port. An incoming packet is queued in one of the  $N$  queues according to its destination output. In the example above, packet “3” and “7” will be in different queues #3 and #7, so while “3” is waiting, “7” can be transmitted. Thus HOL Blocking is eliminated.

Problem 4 (20 points) – Simple networking tools. Login to a Unix machine in the CIP pool. Do the following:

- (a) Run “man ping”, skim the manual page that is displayed. In a couple of sentences, describe what the “ping” tool does. (2.5 points)

ping sends and receives ICMP ECHO messages to measure the RTT and packet loss to a remote host.

- (b) Now, run “ping www.mpi-sws.org”. Stop it after a few seconds by hitting “Control-C”. What information is reported by the “ping” tool? (2.5 points)

You should see RTT on the order of 1ms, and 0% loss. Real domain name of www.mpi-sws.mpg.de is shown as www.mpi-sws.mpg.de

- (c) Now, try “ping [www.cs.cmu.edu](http://www.cs.cmu.edu)”. Stop it after a few seconds by hitting “Control-C”. What are the most important differences between this result and what you saw in part (b)? (2.5 points)

RTT is now much larger, on the order of 100ms, possibly with some loss. The messages are making round-trips to Pittsburgh, PA!

- (d) What network problems do you think “ping” can help diagnose? (2.5 points)

Remote machine not reachable; poor network performance, etc

- (e) Run “man traceroute”, skim the manual page that is displayed. In a couple of sentences, describe what the “traceroute” tool does. (2.5 points)

traceroute discovers the IP routers on the path from the source to the destination. It also measures RTT to each router along the way.

- (f) Now, run “traceroute loki10.mpi-sws.org”. Let it run until it has finished. What information is reported by the “traceroute” tool? (2.5 points)

The path is: CSgate.cs.uni-sb.de, rz2mpi.net.uni-saarland.de, swsao0910.mpi-sws.mpg.de (the last hop is just an alias for loki10.mpi-sws.org). RTT information is also shown.

- (g) Now, run “traceroute [www.cs.cmu.edu](http://www.cs.cmu.edu)”. Let it run until it has finished. What are the most important differences between this result and what you saw in part (f)? (2.5 points)

There are many more routers on the path to www.cs.cmu.edu. Approx 14 hops. If you look closely at the names, you can get an idea where the path goes through!

- (h) What network problems do you think “traceroute” can help diagnose? (2.5 points)

Mainly performance problems on the path to the remote host. This is similar to ping, but using traceroute you can potentially observe loops in the route, or a particular hop is having high delay, etc.