

CS244

Advanced Topics in Computer Networks

Midterm Exam – Monday, May 2, 2016

OPEN BOOK, OPEN NOTES, INTERNET OFF

Your Name: _____

SUNet ID: _____@stanford.edu

In accordance with both the letter and the spirit of the Stanford Honor Code, I neither received nor provided any assistance on this exam.

Signature: _____

- The exam has 11 questions totaling 90 points.
- You have 75 minutes to complete them.
- Some questions may be much harder than others.
- All questions require you to justify your answer to receive full credit, even multiple choice questions for which you circle the correct answer(s).
- Keep your answers concise. We will deduct points for a correct answer that also includes incorrect or irrelevant information.

1	/5
2	/5
3	/5
4	/5
5	/5
6	/5
7	/12
8	/15
9	/9
10	/15
11	/9

Total	/90
--------------	------------

1. [5 points]:

PIM. If PIM runs to completion in each cell time (i.e. until no more connections can be added to the match), what is the maximum number of iterations required for an $N \times N$ switch?

- A $\log_2(N)$.
- B $N/2$.
- C N .
- D $\log_2(N) + 1$.

Justification:

2. [5 points]:

PIM. The paper authored by Tom Anderson proves a lower bound on the expected fraction of unresolved requests that are resolved in each iteration. What is the bound?

- A $3/4$.
- B $1/4$.
- C $1/2$.
- D None.

Justification:

3. [5 points]:

PIM We would like to find ways to make the PIM algorithm converge faster (i.e. in fewer iterations). The proof of the convergence time for the PIM algorithm starts by assuming that an output grants a requesting input that receives *no* other grants with probability k/n . Imagine that the algorithm is “tuned” to make it twice as likely to pick an input that receives *no* other grants, i.e. $2k/n$. What effect will this have on the convergence time?

- A** No effect.
- B** It will make it converge faster, because it expects to resolve at least $7/8$ of unresolved requests in each iteration.
- C** It will make it converge faster, because it expects to resolve at least $3/8$ of unresolved requests in each iteration.
- D** It will make it converge slower, because it expects to resolve at least $1/8$ of unresolved requests in each iteration.

Justification:

4. [5 points]:

PGPS. Which of the following are true for Packetized Generalized Processor Sharing (PGPS)? Circle the correct option. Note that more than one option may be correct.

- A PGPS serves packets in the increasing order of their packet lengths.
- B PGPS provides the same average throughput guarantees for each flow as GPS.
- C PGPS serves packets in decreasing order of their finishing time, which is computed based on GPS scheduling.
- D PGPS can provide packets with an end-to-end delay guarantee across the network, but only if the source is constrained in some way, to prevent it sending in too much traffic.

Justification:

5. [5 points]:

Accountability. Which of the following assists in achieving the last goal in Clark's "Design Philosophy of the DARPA Internet Protocols" (The resources used in the internet architecture must be accountable)? Circle all that apply.

- A Flow Groups and Tunnel Groups in B4
- B Precise RTT measurement in TIMELY
- C Leaky Bucket tokens in GPS
- D Separation of control and data planes in SDN

Justification:

6. [5 points]:

Throughput. Consider an $N \times N$ input-queued switch carrying packets all of the *same* size. All links in the network operate at rate R . When packets arrive to the switch, it simply forwards the packet to a randomly picked output (i.e. it doesn't consult a lookup table based on the destination address). The output is picked uniformly and at random from all outputs. Which of the following statements are true?

- A If each input maintains a single FIFO queue for all arriving packets, then the throughput can be limited to approximately 58% because of head-of-line blocking.
- B If each input maintains a single FIFO queue, the input queues are always empty.
- C If each input maintains virtual output queues (VOQs), the input queues are always empty.
- D If each input maintains virtual output queues (VOQs), then the throughput is 100% for all arrival patterns.

Justification:

7. [12 points]:

Congestion control. Below is a list of congestion protocols.

- A TCP Reno
- B TIMELY
- C QUIC
- D TCP Reno with ECN—if the ECN congestion bit is marked in the packet ACK, perform multiplicative decrease with TCP Reno’s congestion window.

For each of the following, write down the letter(s) of the congestion protocol(s) that has the corresponding feature. For example, if the feature is “packet modifications and window sizes adjusted at the sender,” list A, B, C, and D, and justify briefly: “congestion signals trigger the size of the next window of packets sent by the sender.”

- a. Single-bit congestion signal
- b. Multi-bit congestion signal
- c. Records some measure of RTT
- d. In the single *lossless* link case, signals congestion only when switch buffer has reached capacity
- e. Packets modified at intermediate hops
- f. Implements loss recovery

8. [15 points]:

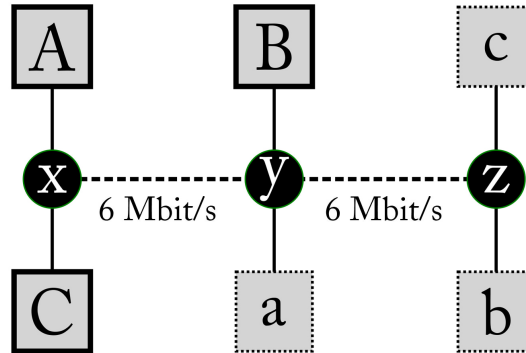
Congestion control.

- a. At the time Van Jacobson and Karels added congestion control to TCP, which of these mechanisms did TCP implementations use to detect lost segments? *Please circle T or F. No justification needed for this part. (1/2 point each)*
- T F** A timeout
 - T F** Four acknowledgments asking for the same next byte
 - T F** Selective acknowledgment of segments delivered while prior segments are missing
 - T F** NACKs for segments still missing after later segments have been delivered
 - T F** Explicit congestion notification from the network
- b. Which of these mechanisms do current TCP implementations use to detect lost segments? *Please circle T or F. No justification needed for this part. (1/2 point each)*
- T F** A timeout
 - T F** Four acknowledgments asking for the same next byte
 - T F** Selective acknowledgment of segments delivered while prior segments are missing
 - T F** NACKs for segments still missing after later segments have been delivered
 - T F** Explicit congestion notification from the network
- c. If the current loss-detection mechanisms had been in TCP from the start, do you think Van Jacobson and Karels would still have wanted to add congestion control to TCP? Give at least one reason why and why not. **(5 points)**

- d. What assumptions about the network stack and application does the VJ+K algorithm seem to make? Which of these assumptions would you question today? **(5 points)**

9. [9 points]:

Flow-rate fairness. In the network below, sender A sends a flow to receiver a , sender B sends a flow to receiver b , and sender C sends a flow to receiver c . The link rate between x and y is 6 megabits per second, as is the link rate between y and z .



a. (3 points) What is the *max-min fair* allocation of throughputs to the three flows?

	$A \rightarrow a$	$B \rightarrow b$	$C \rightarrow c$
throughput (Mbit/s)			

b. (3 points) What is the *max-utilization* allocation of throughputs to the three flows? (This allocation maximizes the total of the throughputs of the three flows.)


	$A \rightarrow a$	$B \rightarrow b$	$C \rightarrow c$
throughput (Mbit/s)			

c. (3 points) What is the *proportionally fair* allocation of throughputs to the three flows? (This allocation maximizes the sum of the log of each flow's throughput.)

	$A \rightarrow a$	$B \rightarrow b$	$C \rightarrow c$
throughput (Mbit/s)			

10. [15 points]:

Leaky-bucket constrained traffic. In this question we'll assume that flows consist of a stream of bits rather than packets. In the PGPS paper we saw how a leaky bucket regulator can constrain a flow so that the number of bits sent in an interval of duration τ is given by: $A(t, t + \tau) \leq \sigma + \rho\tau$. If you think about it, this leaky bucket would allow a burst of up to σ bits to depart from the source in zero time (i.e. at an infinite rate)! This is clearly not possible; so in practice, the source is constrained by the data rate of its outgoing link. We'll assume that the flow is constrained by a leaky bucket at the source, and that a burst can depart no faster than R bits/second (the data rate of the link to which the source is connected). The figure shows how the traffic is constrained. The flow passes through multiple routers, all of which use bit-by-bit generalized processor sharing (GPS) to serve multiple flows.



- a. (5 points) If the first router allocates the flow a service rate greater than or equal to some constant c , write down expressions for the maximum delay of a bit in the router, and the size of the buffer needed to prevent overflows. (Write your answers in terms of σ , ρ , and R , but not c).

b. (5 points) What happens to your answer above when $R = \rho$? Explain your answer.

c. (5 points) Why do think it is customary to assume that a burst can leave a leaky bucket regulator at infinite rate?

11. [9 points]:

TIMELY. TCP achieves max-min fair share rates per connection. That is, no flow could have a higher rate without decreasing the rate of an equal or smaller rate flow. B4 also uses max-min fairness for Traffic Engineering in assigning rates to flow groups.

(a) (4 points) Does TIMELY also achieve max-min fairness? Explain your reasoning.

(b) (3 points) Does bufferbloat occur to a similar degree in networks that use TIMELY compared to networks that use vanilla TCP? Explain your reasoning.

(c) (2 points) What are the benefits of using RTT over a packet drop as a congestion signal?