Exam - Final

Due: 11:00am, May 13th, 2011

Closed Book. Maximum points: 100

NAME:

1. TCP congestion control [30 pts]

The figure below shows the evolution of the congestion window of a TCP Reno connection.

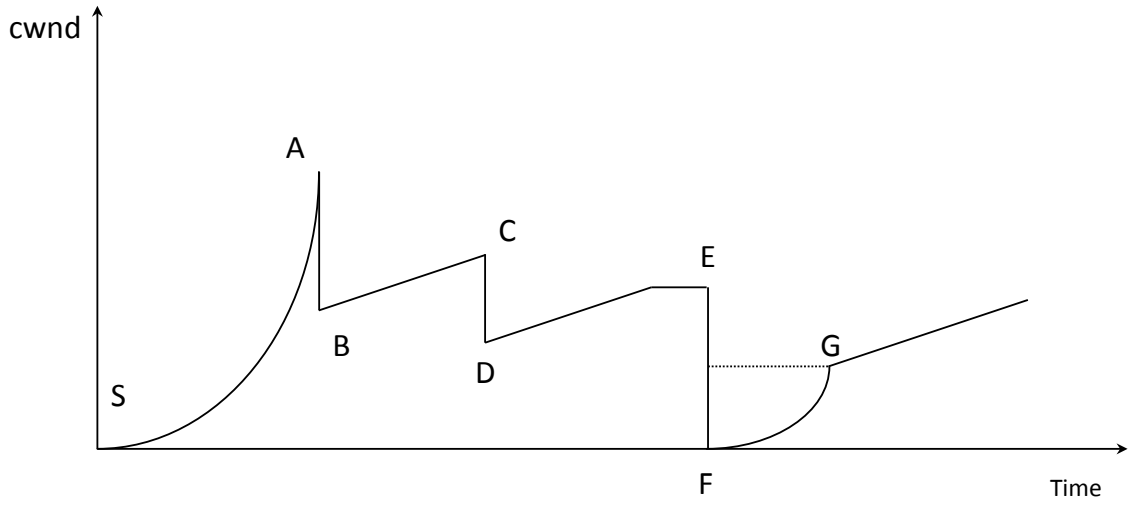

a. To what ranges of the protocol operation do the following sections of the curve correspond: [4 pts]

   i. S-A
   ii. B-C
   iii. F-G

b. What causes the change in window sizes at the following instants:[4 pts]

   i. A
   ii. C
   iii. E
   iv. G

c. A TCP Reno flow in steady state (AIMD) presents a characteristic sawtooth pattern in the congestion window size, even if there is only one flow in the link.

i. Why does this pattern arise? When does the congestion window increase? When does it decrease? [4 pts]

ii. What is the average window size if W is the maximum window size that does not incur in losses due to congestion (assume losses only occur when the window surpasses W)? [6 pts]

iii. Why does TCP never stop changing the rate, even if there is a single flow on a path? [4 pts]

d. The way TCP probes for available bandwidth tends to cause high latency in routers with FIFO drop-tail queues.

   i. Why? [4 pts]

   ii. How does RED alleviate the problem? [4 pts]

2. End-to-end Argument [8 pts] The end-to-end argument serves as a guideline for a designer to place functionality in a system. Take reliability in a BitTorrent transfer over TCP over a wireless link as an example.

a. BitTorrent has its own hashing mechanism to identify incorrect blocks. Why do you need this in spite of TCP's reliability efforts? [4 pts]

b. If BitTorrent implements reliability anyway, why then do you have TCP, and the wireless link underneath it, both implement reliability mechanisms? Isn't that redundant? [4 pts]

3. DNS [12 pts]

a. Suppose you want to lookup `www.sigcomm.org`, through a local resolver that has nothing but the root name servers in cache. Give a list of name servers that your local resolver would have to go through to give you the answer. Assume each label in the DNS name will be resolved by a different nameserver, and use the notation 'ask nameserver of (...) for (...)', since you don't know their names or IP addresses. [4 pts]

b. You performed two successive lookups for `www.sigcomm.org` through the local resolver in the CS department (128.148.31.200), and got the following two responses (these are actual responses from `dig`, edited for compactness). [4 pts]

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15112
;; QUESTION SECTION:
;www.sigcomm.org.       IN  A
;; ANSWER SECTION:
www.sigcomm.org.    14400   IN  CNAME   sigcomm.org.
sigcomm.org.        14400   IN  A   63.118.7.16
;; AUTHORITY SECTION:
org.            40797   IN  NS  a2.org.afilias-nst.info.
org.            40797   IN  NS  b2.org.afilias-nst.org.
;; ADDITIONAL SECTION:
a2.org.afilias-nst.info. 26338   IN  A   199.249.112.1
b2.org.afilias-nst.org. 29856   IN  A   199.249.120.1
;; Query time: 25 msec
;; SERVER: 128.148.31.200#53(128.148.31.200)
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 45215
;; QUESTION SECTION:
;www.sigcomm.org.       IN  A
;; ANSWER SECTION:
www.sigcomm.org.    14263   IN  CNAME   sigcomm.org.
sigcomm.org.        14263   IN  A   63.118.7.16
;; AUTHORITY SECTION:
org.            40660   IN  NS  b2.org.afilias-nst.org.
org.            40660   IN  NS  a2.org.afilias-nst.info.
;; ADDITIONAL SECTION:
a2.org.afilias-nst.info. 26201  IN  A   199.249.112.1
b2.org.afilias-nst.org. 29719   IN  A   199.249.120.1
;; Query time: 0 msec
;; SERVER: 128.148.31.200#53(128.148.31.200)
```

Explain the following differences between the two:

i. `id:` in line 1:

ii. Query time, 25ms on the left, 0ms on the right, in line 13:

iii. The order of the two NS records in the Authority Section:

iv. The numbers in line 5: 14400 (left) and 14263 (right):

c. Assuming no other queries were made to this domain on this server, can you determine which one was made later, and how much time elapsed between the two? Briefly justify your answer. [4 pts]

4. Anycast [8 pts] Anycast is a communication pattern in which one node communicates with one of a group of other nodes. It is implemented on the Internet in different ways.

a. At the IP layer, anycast allows a single IP address to map to more than one node. How is IP anycast implemented? Give an example of where it is used. [4 pts]

b. Anycast can also be effectively implemented using DNS: a given DNS name can map to different IP addresses. How is anycast implemented using DNS? Given an example. [4 pts]

5. Overlay Networks [14 pts]

a. (Chord) A problem in Chord, and in many overlays, is that there is little to no correlation between proximity in the ID space, and proximity in the network. This is particularly bad when the latency between nodes that are part of a Chord path are much higher than the latency between the start and end nodes.

i. Give an example illustrating this problem. [4 pts]

ii. One nice aspect of the Chord algorithm is that there is flexibility in how to fill the finger table. In particular, the ith finger entry for node with id n can be *any* node responsible for id n_i such that $n+2^i(\mod 2^m) \leq n_i < n+2^{i+1}(\mod 2^m)$. (Most implementations choose the finger to be the first node in this range for each i.) Given this flexibility, how can you choose fingers to minimize the problem in (a)? [4 pts]

b. (RON) We saw that RON is a routing overlay that involves N nodes. Each node in RON maintains complete connectivity information (latency, bandwidth, loss probability) to all other nodes. The surprising aspect of RON is that the triangle inequality does not hold on the Internet: the latency between A and B is not necessarily lower than the latency of A to C plus C to B. RON also reacts to failed paths much quicker than BGP. Explain, using your knowledge of BGP updates and route selection, why RON can find better paths and react more quickly to failures. [6 pts]

6. Wireless [12 + 7 pts] One of the functions of the link layer is to control access to the medium when multiple parties want to transmit. We saw that Ethernet uses carrier sense with collision detection (CSMA/CD) to avoid multiple nodes transmitting at the same time.

a. Explain why this doesn't work in wireless. [4 pts]

b. Explain what the hidden terminal problem is (you can draw a diagram if it helps). [4 pts]

c. RTS/CTS (Request-to-send/Clear-to-send) is a mechanism used to alleviate the hidden terminal problem. Explain how it works, and give two reasons why it is not used in practice.[4 pts]

d. (Bonus) In class we saw a new scheme for full duplex wireless, in which a node can transmit and receive at the same time, using the same radio. If you have this capability on all nodes, propose a scheme that solves the hidden terminal problem. [7 pts]

7. Security [16 pts]

a. We saw several desirable properties for secure communication. Briefly (in one or two sentences) describe what the following properties mean:[4 pts]

i. Authentication:

ii. Integrity:

iii. Confidentiality:

b. Suppose I implement a secure server that uses the following protocol: Client A sends a request (in the clear). B replies with its public key. A supplies its public key, encrypted with B's public key. All further traffic is encrypted with the appropriate key. Assume that an attacker can snoop, delete, modify, and inject traffic in the path between A and B. Why is this not secure? Structure your answer around the three properties above. [4 pts]

c. SSL/TLS provide the three properties above. Describe how SSL/TLS guarantee *authentication* and *confidentiality*. In particular, your answer must describe the roles of certificates and of certificate authorities. [8 pts]