EXAM Telematics Networks (192620000)

1 November 2011, 13:45–17:15

- This is an open-book exam: you are allowed to use the book ("Computer Networking" by Kurose & Ross), the reader, copies of the lecture slides, and a dictionary. Use of other written material, such as your own notes, is not allowed, nor is the use of laptops, notebook computers, PDAs, mobile phones, etc. ***Please remove any such material and equipment from your desk, now!***
- You are allowed to answer in either English or Dutch.
- Each problem is worth 10 points, except problem 2 which is worth 8 points.
- You should always explain or motivate your answers, with so much detail that the grader can judge whether you understand the material; so just saying "yes" or giving a formula without explanation is not enough.
- Besides the exam, you are also given a questionnaire about the course. Please do fill out that form, and hand it in when leaving the room. Of course, you may fill out the questionnaire after handing in the exam answers, so the questionnaire doesn't cost you time that would be better spent on the exam itself.
- Note that your exam will not be graded unless/until you have also completed the *Wireshark* assignment, and that 10% of your final grade will be determined by the homework multiple-choice questions.

---

1. Information and communication theory

(a) From an information-theory point of view, how much information is there (approximately) in the sequence number field in the TCP header of a *random* TCP packet? Explain.

(b) Same question as (a), but now assuming that you already know the header and contents of the *previous* TCP packet with the same source and destination addresses and port numbers.

(c) Consider this graph of bit-error-rate (BER) as a function of signal-to-noise ratio (SNR). One curve is a typical curve for some transmission system *without* error-correcting coding, one curve is for the same transmission system but *with* error-correcting coding, and one curve is *incorrect* (impossible).
Which curve is which?
Motivate your answer.

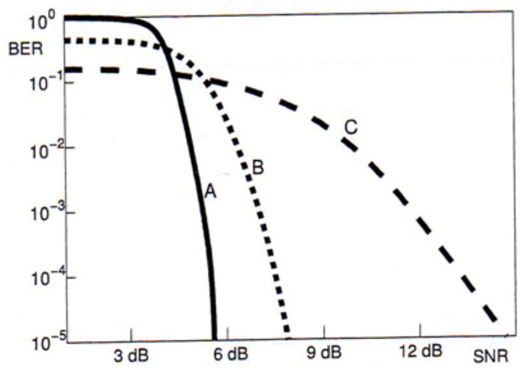

(d) According to a result from Shannon, we can make the bit-error-rate arbitrarily low; so why do the curves in the above graph not show BER=0? In other words, how does the above graph relate to Shannon's results?

2. IPv6 addressing

Consider the following notations:

1) 2000:5678:98ab:cdef:1234:5678:98ab:cdef
2) 2234:0000:98ab:cdef:1234:0:0:cdef
3) 2234::5678:98ab:1234::cdef
4) 2234:0:98ab:cdef:1234::cdef
5) 3234:5678::98ab:cdef
6) 3234:5678:98ab:cdef
7) face:b00c::g00:g1e

(a) Which of these is/are *invalid* IPv6 address notations? Explain.

(b) Do any two of the above notations represent the *same* IPv6 address? Explain.

(c) Which of the above addresses is/are inside the 2000::/3 subnet? Explain.

(d) Which of the above addresses is/are 6to4 addresses? Explain.

---

3. Network layer issues

Consider some home network containing several computers. The home network has a single connection to an internet service provider (ISP), and this ISP only gives a single public IPv4 address, and no public IPv6 address. A "box" connects the home network to the outside world.

(a) Can that "box" be a bridge, and still allow all home computers to talk to computers in the outside world? Explain.

(b) Does the "box" need to speak BGP to the ISP? Explain.

(c) Suppose the "box" is a NAT. When a packet comes in from the ISP, how does the NAT know to which home computer the packet should be sent?

(d) Suppose the "box" is a 6to4 router, so the home computers can all use IPv6. When a 6to4 packet comes in from the ISP, how does the 6to4 router know to which home computer the packet should be sent?

Continued on next page...

4. TCP flow control and the Silly Window Syndrome

Suppose two hosts A and B have just set up a TCP connection, and host A has a lot of data to be sent through this connection to B. Assume that initially, host A has a congestion window of 1 MSS. The MSS (Maximum Segment Size) for this connection is 1000 bytes. Host B has a receive buffer of 4000 bytes. The RTT of this link is negligible; furthermore, for simplicity assume delayed ACK is not used (i.e., TCP sends an ACK for each incoming data packet immediately, instead of waiting a little to combine multiple ACKs). The application software on host B is rather slow: it reads 1 byte out of the receive buffer every 200 ms.

(a) Describe what packets will be sent; you may stop after host A has transmitted at least 4003 bytes of data. A good way of doing this is by completing the below time-sequence diagram, and writing some explanation to go with it:

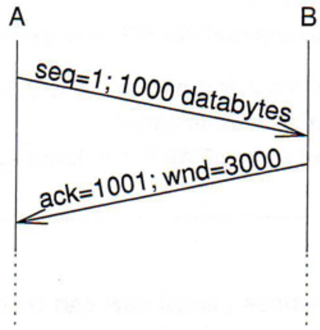

Note 1: although this question may look a lot like the problem about TCP congestion control that we did in the lecture, it is actually rather different (and easier); the emphasis here is on flow control, not congestion control.

Note 2: the description of handling a full receive buffer in the book (p. 290 in the 5th edition) is unfortunately not quite complete:

- the text suggests that probe segments with one byte of data are sent *continuously*, but in reality they are sent at intervals of several seconds (otherwise, they would overload the network).
- when space becomes available in the receive buffer, TCP sends an extra ACK to notify the sending side of this; the probe segments are only needed in case this extra ACK is lost.

As you (hopefully) see, many packets contain far fewer data bytes than the MSS (which is 1000 bytes).

(b) Why is this undesirable?

This phenomenon is called the *Silly Window Syndrome* (SWS).

(c) Propose an extra rule to be implemented in TCP to avoid the silly window syndrome, and explain why your rule works well.
(Your rule should not require any changes to the TCP header format. Furthermore, it should be generic, so it also works in slightly different situations, e.g., when the application at B reads groups of 2 bytes. And of course, it should not affect TCP performance if the application software at host B reads fast enough.)

Note: this is not a hypothetical question just made up for this exam; real TCP implementations indeed contain an algorithm to avoid the SWS.

Continued on next page...

5. Explicit Congestion Notification (ECN)

(a) What is the advantage of using ECN, compared to the "old" mechanism of detecting congestion?

(b) Suppose ECN is used; should the end-hosts then still also reduce their congestion window if a packet is lost? Why?

ECN uses some previously unused bits in the IP and TCP headers; previously, these bits should be zero. Unfortunately, some overzealous firewalls get nervous when they see these bits being non-zero.

(c) Suppose a firewall sets the ECN bits in the IP header to zero in all *incoming* packets. What consequence(s) will this have?
E.g., will congestion not be detected anymore? Will TCP no longer work at all? Etc.

(d) Suppose a firewall drops all incoming packets in which the CWR bit in the TCP header is set to 1. What consequence(s) will this have?
E.g., will congestion not be detected anymore? Will TCP no longer work at all? Etc.

---

6. Real-time traffic

Assume we have a data source whose packet flow can be described by a leaky-bucket model with rate r tokens/second and bucket size B tokens (with 1 token per packet).

(a) What is the largest burst this source is allowed to send?

Suppose this source sends a burst of k packets once every 100 ms, and nothing in between.

(b) Give an expression for the maximal allowed value of k in terms of r and B.

Now assume this source, which sends a burst of k packets every 100 ms, shares a link in the network with another source that sends a burst of ℓ packets every 200 ms. All packets are S-bits long, and the outgoing link speed is R bits/sec.

(c) If Round Robin scheduling is used to fairly share the link among these two sources, what is the maximum possible delay for "our" source (the one sending k packets every 100ms)?

ATM was originally intended (among others) as a telephone network, and thus engineered to provide quality of service guarantees.

(d) Compare ATM to the approaches proposed for QoS in the internet: Integrated Services, Differentiated Services, and Overprovisioning. To which of them do you think ATM is most similar? Explain.

End of this exam.