EXAM Telematics Networks (192620000)

26 October 2010, 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!***
- Although the questions are only written in English, you are allowed to answer in either English or Dutch.
- Each problem is worth 10 points.
- 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!
- Note also that 10% of your final grade will be determined by the homework multiple-choice questions.

1. Communication theory

(a) Assume we have a channel with a (Shannon) capacity of 900 bits/s. Can we send 890 bits/s through this channel with a bit error rate of not more than 0.01% ? Explain.

(b) In the IP header, there is a 4-bit field for the version number. How much information does this field actually contain, from an information-theory point of view? Explain.
(Note: if you can't compute a precise number, indicate how it could be computed and how large you approximately expect it to be.)

(c) Assume we have a channel with a (Shannon) capacity of 900 bits/s. Describe a way to squeeze 1000 bits/s through this channel with (on average) not more than 6% bit errors.

2. Wireless medium access mechanisms

Assume for this problem that the ranges over which wireless signals may propagate are symmetric: e.g., if A is within B's detection range, then B is also within A's detection range.

(a) Is the "hidden terminal" problem caused by the fact that the *interference* range may be larger than the *detection* and *transmission* ranges, or can it still occur if for each station these three ranges are equally large? Explain.

(b) Suppose all stations are within each other's *interference* range but not all are within each other's *transmission* range. Does the Carrier Sense part of CSMA/CA still work properly if all stations are within each other's *detection* range? Explain.

(c) CSMA/CA tries to avoid collisions of the data packets by using RTS/CTS packets to 'reserve' the channel. Does this still work correctly if for each station the *detection* range is not equally large as its *interference* range and/or not equally large as its *transmission* range? Explain.

3. Network Address Translation

Consider a Network Address Translator (NAT) for IPv4, connecting several hosts on a private network to the global internet. Assume *no* fragmentation happens *within* the private network.

(a) Suppose a packet is fragmented somewhere on the global internet on its way *to* the NAT. Can the NAT perform address translation for each of the fragments individually, or does it need to reassemble the packet first? Explain.

Many NATs for IPv4 do not check or change the 'identifier' field in the IPv4 header of packets that go through it.

(b) Explain in your own words what the purpose of the 'identifier' field is.

Such a NAT that does not do anything with the 'identifier' field can cause problems for outgoing packets (i.e., packets going *from* the private network via the NAT *to* the global internet).

(c) Explain this problem, e.g. by giving an example in which it goes wrong.

---

4. 6to4 tunnelling

Some people find the 128-bit IPv6 addresses too long: such long addresses make the headers longer than necessary, thus wasting bandwidth.

(a) Does tunnelling the IPv6 packets inside IPv4 packets (which have shorter headers) help to reduce this waste? Explain.

Tunnels are used for connecting IPv6 "islands" to the rest of the IPv6 network.

(b) Explain (in your own words) how 6to4 makes such tunnelling much easier.

The following statement is made about 6to4 tunnelling in the reader (page 53):

> *Packets from the IPv6 Internet to 6to4 systems must be sent to a 6to4 relay router by normal IPv6 routing methods. The specification states that such relay routers must only advertise 2002::/16 and not subdivisions of it to prevent IPv4 routes polluting the routing tables of IPv6 routers.*

(c) Explain the concern that is expressed here: how could IPv4 routes influence the IPv6 routing tables, and why would that be undesirable ("pollution")?

---

5. BGP

(a) In BGP, "path" information is exchanged instead of just "distance". What does this mean, and why is it useful?

Consider the following network, consisting of 5 autonomous systems and 5 links between them:

Continued on next page...

Normally, AS1 would advertise itself to its neighbours with a path of just one entry, namely itself.
However, suppose that AS1 indeed decides to advertise itself to AS2 like that, but to AS5 it advertises itself with the path 'AS1,AS1,AS1'. (Such tricks are actually used in reality.)
Assume that no (other) BGP policies or local preferences have been installed in any AS.

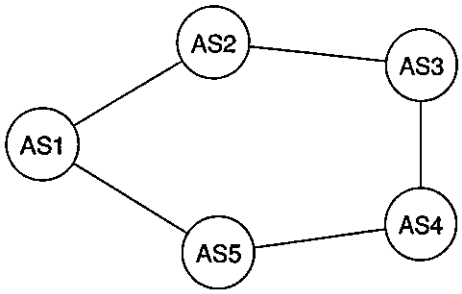

(b) What BGP information about AS1 will then be exchanged by the other ASs?

(c) What routes will AS2, AS3, AS4 and AS5 use to AS1?

(d) What would happen if the link between AS3 and AS4 breaks?

Now suppose that AS1 wants to achieve that AS4 *never* routes packets for AS1 via AS5, not even if the link AS3-AS4 breaks.

(e) How can AS1 achieve this by changing the 'AS1,AS1,AS1' advertisement it sends to AS5?

---

6. Problems of standard TCP in high-speed networks

(a) Consider the "receive window field too small" problem. If the round-trip time (RTT) increases, does this problem get worse (i.e., reduce TCP's speed even further)? Explain.

(b) Consider the "sequence number wrap" problem. If the RTT increases, does this problem get worse? Explain.

(c) Consider the problems related to the congestion control algorithm. If the RTT increases, do these problem get worse? Explain.

---

7. TCP congestion control, and techniques for real-time support

Real-time flows are often specified using a leaky-bucket model.

(a) Could we also use a leaky-bucket model to describe a TCP flow which is used for a file transfer, and if so, does it make sense? Explain.

Suppose in some network there are two (standard) TCP connections active, which share only one link, and this link happens to be the bottleneck for both of them. The capacity of this bottleneck link is 120 Mbps. Straight-forward FIFO scheduling is used on this link; we'll assume that this leads to both flows having the same packet loss rate. The flows have different round-trip-times: flow 1 has an RTT of 0.1 ms, while flow 2 has an RTT of 0.2 ms.

(b) What throughputs do the connections achieve? Why?

(c) Assuming fair queueing would be used instead of FIFO, with each flow treated as a separate class, what would the throughputs be? Why?

(d) Assuming TCP CUBIC would be used instead of standard TCP (but still using FIFO queueing), what would the throughputs be? Show your calculation.
Hint: use one of the formulas on page 100 of the reader.
If you don't have a calculator at hand, give your answer as a formula.

End of this exam.