

EXAM Telematics Networks (192620000)

30 October 2012, 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 a simple (non-graphical) calculator is allowed (but not needed).
Use of other written material, such as your own notes, is not allowed, nor is the use of laptops, notebook computers, graphical calculators, 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.
- 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

Consider an information source which produces 1000 messages per second, each message being either Y or an N (think “yes” and “no”), with 75% probability for Y and 25% for N, and consecutive messages being independent of each other.

- (a) Calculate Shannon’s information rate (i.e., amount of information per unit of time) for this stream.
- (b) Design a way to encode these messages using on average less than 1 bit per message, and show that your method indeed achieves this.
(Note: you’re not asked to actually achieve the Shannon rate.)

Consider Hamming codes for error correction.

- (c) Suppose we have a message of 16 data bits, in which we wish to be able to correct single bit errors. How many parity bits need to be added at least to make this possible? Explain.

There are practical systems in which one either is sure that a bit is received correctly, or that it is “erased”. Erasure means that the receiver simply doesn’t know at all whether it was a 0 or a 1 (but it does know that there was *some* bit, e.g. based on timing). Error correcting codes can be used here too, not to correct bits (because all received bits are correct), but to *fill in* the erased bits.

- (d) Consider the two-dimensional parity matrix. How many erasures can it correct, and why?

Continued on next page...

2. Transmission media and MAC protocols

- (a) Is a MAC protocol needed on a full-duplex twisted-pair copper cable link between a host and a router? If so, which one is used in practice? If no, why not?

One possible architecture for Fiber-To-The-Home networks is the so-called Passive Optical Network, where a single fiber comes from a central node (which is connected to the rest of the internet), and is split using passive splitters to individual fibers to individual homes. Any optical signal from the central node thus reaches all homes ("downlink"); and the optical signals from the homes ("uplink") get combined together onto the fiber going to the central node (but the signal from one home does not reach the others homes).

- (b) Clearly, the uplink direction needs some kind of multiple-access mechanism. Propose two solutions for this, and give an advantage and a disadvantage of each of them.

Consider a step-index glass fiber.

- (c) If we *increase* the difference in the index of refraction of the cladding and the core, does that *increase* or *decrease* the dispersion? And what influence does this have on the signal bandwidth of the fiber? Explain your answers.

3. Faulty middleboxes

A "middlebox" is a device sitting somewhere on the path between internet hosts, which does something more than just forwarding packets as a normal router should; examples of middleboxes are firewalls and NATs. Middleboxes may do undesirable things to your packets, like modify them or drop them inappropriately, either due to bugs in the software of the middlebox, or due to excessive paranoia of the designer.

In this problem, we will consider a middlebox that sits between a single host (let's call it "our host") and the rest of the internet. You are asked, for each of the middlebox behaviours given below, to tell what consequences this has (for example, "outgoing TCP connections will fail", or "the host will catch fire"), and explain how/why that happens. Where applicable, be careful to distinguish between (outgoing) connections set up by our host to another host somewhere on the internet, and (incoming) connections set up by another host to ours.

A few remarks:

- The words 'incoming' and 'outgoing' are from the point of view of "our" host.
 - Unless otherwise specified, IPv4 is assumed.
 - Whenever the middlebox modifies a packet, it also recomputes checksums that cover the modified part of the packet.
 - Apart from the described behaviour, the middlebox simply forwards the packets as it should.
 - It *may* be that the action of the middlebox is innocent, i.e., has no consequences that the endhosts or users notice; in that case, say so, and explain why it is innocent.
- (a) Of every incoming packet, the middlebox makes four copies and sends them all to "our" host.
What are the consequences? Why?
- (b) The middlebox sets the 'more fragments follow' bit in the IPv4 header of every incoming packet to zero.
What are the consequences? Why?
- (c) If the middlebox sees a DNS request for `www.thepiratebay.org`, it drops it. Does this make the website `http://www.thepiratebay.org` unreachable for our host? Explain.

- (d) In incoming IPv6 Packet-Too-Big ICMP messages, the middlebox sets the MTU value to 65535 bytes.
What are the consequences? Why?
- (e) If *both* ECN bits in the IP header of incoming packets are 1, the middlebox drops the packet.
What are the consequences? Why?
-

4. Routing

Three routing protocols used in the Internet are RIP, OSPF and BGP.

- (a) Describe, for each of them, briefly what information is exchanged between two neighbouring nodes about each destination (node or (sub)network).
- (b) Can a transition from IPv4 to IPv6 help to reduce the size of the BGP routing tables? How and why, or why not?
- (c) The mathematical description of the Distance Vector algorithm contains a *minimization* operation. Suppose we would replace this minimization by a *maximization*: would this make the algorithm converge to the *longest* (most expensive) instead of the shortest (cheapest) paths? Why, or why not?
-

5. Transport layer performance

Someone in Europe is downloading a 10 gigabyte file from a website that is hosted on a server in America. All links have a link speed of 1 Gbps and are not congested, and the computers on both sides are fast enough. On both sides, TCP-Reno is used with the window scaling and timestamp options.

- (a) Explain in your own words the purpose of the window scaling option.

Despite the use of these TCP options, the download speed is still much less than 1 Gbps.

- (b) Explain in your own words what causes this.
- (c) Would it help if the file were split into two halves, and this person would use 2 TCP connections, one for the first half and one for the second half of the file, at the same time? Explain.
- (d) Can his download performance be improved by installing TCP CUBIC on his computer? Or should it be installed on the web server? Or on both? Why or why not?
-

Continued on next page...

6. Packet pairs and scheduling

For network measurements “packet pairs” are often used. This means that two packets are sent into the network immediately after each other, and the time between their arrival (the “inter-arrival time”) at the destination is measured.

In this problem, assume the packet pair consists of two packets of 1000 bits each.

Assume the packet pair travels over a single link of 10 Mbit/s, with no other traffic.

(a) What is the inter-arrival time at the destination? Explain.

Next, assume the path from source to destination goes over 3 links; the first has a speed of 100 Mbit/s, the second of 10 Mbit/s, and the last of 200 Mbit/s, again with no other traffic.

(b) What is the inter-arrival time at the destination? Explain.

Still considering the three links as discussed above, and assume that (only) on the 10 Mbit/s link there is also other traffic (called “cross-traffic”), whose packets can be anything between 1000 and 10000 bits long.

(c) Assume FIFO scheduling is used, and the cross-traffic satisfies a leaky-bucket policer with a bucket size of 10 tokens, and a token rate of 10 per second (counting 1 token per packet). What is the minimum and maximum possible inter-arrival time of our packet pair at the destination? Explain.

(d) Assume weighted-fair-queueing is used, with equal weight for our packet pair and the other traffic. What is the minimum and maximum possible inter-arrival time of our packet pair at the destination? Explain.

End of this exam.