Network Systems (201300179/201400431), Test 4

April 2, 2015, 13:45–15:15

- This is an open-book exam: you are allowed to use the book by Peterson & Davie and the reader that belongs to this module, and the handout about peer-to-peer communication (i.e., the part of the Kurose&Ross book distributed via Blackboard). Furthermore, use of a dictionary is allowed. Use of a simple (non-graphical) calculator is allowed.
- Other written materials, and laptops, tablets, graphical calculators, mobile phones, etc., are not allowed. ***Please remove any such material and equipment from your desk, now!***
- Although the questions are stated in English, you may answer in English or Dutch, whichever you are more comfortable with.
- 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.
- Visiting the toilet without explicit permission of the supervisor is not allowed. During the last 30 minutes of the exam, no toilet visits are allowed.

---

1. TCP Congestion Control

Consider a TCP connection between hosts A and B which has been active for a while, with host A sending data to host B. At some point, the `CongestionWindow` = 1 MSS and `SlowstartThreshold` = 4 MSS (see footnote[1]).

2 pt (a) How was the most recent packet loss detected? Explain.

4 pt (b) Assuming no further packet loss occurs, how many RTTs does it take to transmit the next 15 data packets? Explain your answer by drawing a time-sequence diagram, in which you indicate not only all packets, but also the values of host A's `CongestionWindow` and `SlowstartThreshold`.

The book says about the AIMD phase (p. 503):

> The congestion window is incremented as follows each time an ACK arrives:
> CongestionWindow += MSS × (MSS/CongestionWindow).

Unfortunately, taking this literally allows cheating by a receiver which wants to download faster. By sending multiple ACKs, each acknowledging a small part of a received data packet, the receiver can make the sender increase the congestion window extra quickly.

2 pt (c) Give an algorithm which does not suffer from this problem, and explain why it works.

Continued on next page...

[1] By `SlowstartThreshold`, we mean the same thing which is called `CongestionThreshold` in the book. The former name is much more common though.

2. QoS

A network administrator of a company wants to share the company's 10 Mbit/s internet link between VoIP (Voice over IP, i.e., telephony over the internet) traffic and web browsing traffic. He wants to ensure that, no matter how much web browsing traffic there is, half of the link speed is still available for all VoIP calls together.

The administrator wants to use his router's Fair Queueing option for this.

2 pt (a) What's the best solution: configure the router to distinguish two flows, with all VoIP in one and all web traffic in the other flow; or to configure the router to treat every separate VoIP and web connection as a separate flow? Explain.

Suppose there is one VoIP flow and one web flow, sharing this 10 Mbit/s link using Fair Queueing. The web flow uses packets of 1000 bytes each. The VoIP flow uses 100 byte packets, and sends according to a token bucket model with a bucket size of 400 bytes, and a token rate of 1000 bytes/s.

3 pt (b) What is the maximum delay a VoIP packet may encur at this router?

---

3. Security

3 pt (a) A firewall at the border of a network (i.e., at the point where this network is connected to the rest of the internet) is often configured to drop outgoing packets whose source address is not in the range allocated to this network.
Consider the following three goals of security: *confidentiality*, *authentication*, *availability*. For each of these, explain whether the mentioned filtering contributes to it or not.

2 pt (b) When using HTTPS, can the client be sure about the identity of the server, and/or can the server be sure about the identity of the client? Explain.

2 pt (c) In PGP, the session key is encrypted using the recipient's public key. So a man-in-the-middle could replace this session key by his own session key (and replace the message by something else encrypted with that new session key). Isn't this risky? Shouldn't we also *sign* the session key to prevent this? Explain.

---

4. Time synchronization and localization

2 pt (a) Is NTP suitable for sensor networks, explain why (in max 10 sentences)?

3 pt (b) What is the difference between location using time-of-flight (ToF) and localization using time difference of arrival (TDoF)? What are the basic consequences of these two techniques in terms of scalability and communication overhead?

3 pt (c) (i) Suppose we want to to have the Earth and the Moon synchronised in time. What is the main problem that this time synchronization is difficult?
(ii) Suppose we want to have time synchronisation between two wireless acoustic devices, one at the water surface and one at the button of the sea at 1km depth. What is the main problem that this time synchronization is difficult?

End of this exam.