Advanced Networking 2013/14

Examples of Written Exercises
with solution hints and an estimation of the time to solve them

Instructions:

The time for the written part of the exam is always 90min (1h:30). There might be one or more exercises (the three sample exercises here are not "one exam", but examples of what you can find in the exam itself), sometimes they can be broad enough so that there is no single "correct" solution (See Exercise 2), and there are many different ways of solving it, some may be brighter than others, and some may be entirely wrong, in the sense that are solutions where there are serious conceptual bugs.
The time you have is limited, but the correction process keeps this into account: in any case sometimes you have to take decisions and include what is more important leaving out details or less important issues.
Sample solutions to these exercises will be published shortly, also with the indication of the time needed to produce the solution itself (i.e., the time it took to write it multiplied by 1.5 rounded to the 5 minutes).

Exercise 1

Consider a TCP connection at steady state supporting a large file transfer. The receiver window is exactly 40 MSS, and there have been no losses in the initial part of the transfer. The connection RTT is 100 ms, and the network capacity is much larger than the connection throughput. Also queuing times at routers are negligible.

1. What is the connection throughput in this state?

At some point during the connection, say t=0 for convention, due to an error in the operating system of a router exactly 1 packet every 40 is discarded.

2. Compute the connection throughput in this case

Suppose the connection is ECN enabled and the router instead of discarding the packet marks it with the ECN bit.

3. Compute the connection throughput also in this case and discuss the differences (if any) with the answer at point 2.

Exercise 2

Describe SIP Methods and Operations, with specific regard to the role of the different logical entities; consider in particular the Registrar and Proxy Servers.

Exercise 3

Explain the difference between store-and-forward switches and cut-through switches. Estimate the difference in end-to-end delay within a single data center with a three-tier hierarchical architecture (i.e., servers are connected to tier-1 switches, which are interconnected with tier-2 switches, which are finally interconnected with tier-3 switches.
All switches support 10Gbit/s links and jumbo-frames (suppose for simplicity that a jumbo-frame is up to 5000 bytes long), and all interconnections' length is identical and equal to 30m. Finally, the time needed to access the forwarding table is 10ns.

Solution Traces

Exercise 1 (Solution time ~ 50min)

1) The throughput Thr of TCP (as any window-protocol) in the conditions indicated by the text is:

Thr = W/RTT = 40x1460x8/0.1 [bit/s] ≅ 4.67 [Mbit/s]

2) Given one packet every 40 is dropped, and assuming that FR/FR works perfectly in its basic form, i.e., the dropped packet is recovered in exactly 1 RTT, then the windows is halved and congestion avoidance is entered, TCP works in “cycles” as represented in the figure below. The figure represents also the transient phase that TCP spends to reach a new steady-state starting from W=40 MSS until it reaches are “regular pattern”. The regular pattern is highlighted by the shadowed area.
In order to make calculations easy, we suppose that in halving the window TCP always truncates it to the nearest integer in MSS, and that retransmitted packets are not counted in the “loss pattern” (counting them changes little, but makes computations more cumbersome). Thus after the first loss the window is set to 20, after 2 RTT (W equal to 21) another packet is lost, next W is set to 10, and the next loss happen at window 13 (10+11+12+7 = 40). W is set to 6 and the next loss happen with window 10, then the window is set to 5 and the next loss happen again at window 10, thus the regular pattern is a cycle of 40 good packets sent over 7 RTTs. In practice TCP spends 1 RTT to retransmit the lost packet, which does not count in the throughput, as it is a re-transmission, then spends 6 RTTs with W=5,6,7,8,9,10, then another loss occur.

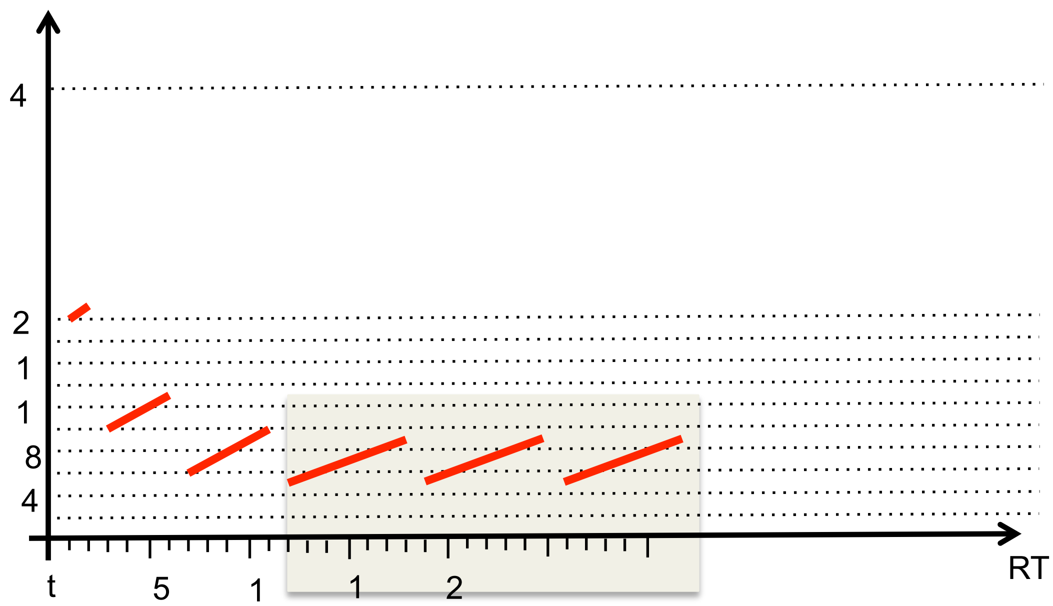

Thus the approximated steady state throughput is:

Thr = 40/(7RTT) = 40x1460x8/0.7 [bit/s] ≅ 0.67 [Mbit/s]

The mere loss of a packet every 40 leads to a loss in performance of a factor of 7, clearly showing the highly non-linear behavior of TCP.
More precise calculations can be done, as this is an approximate behavior, but they do not add much insight in TPC performance in these situations. The key approximation done lies in the fact that TCP keeps sending packets after the one dropped, thus the behavior of the window is more complex, and sometimes grows above 10, leading to a steady-state which is not a single “cycle” but a repetition of multiple similar cycles.

3) In this case there is no much difference in marking or dropping packets, as the performance is dominated by the new, much smaller average transmission window reached in presence of drops. Indeed, the main advantage of ECN would be in not re-transmitting the lost packet, so that the RTT spent (in our approximate computation) in the retransmission and that do not contribute to the throughput is not present, leading to a cycle lasting only 6 RTTs and not 7. We can assume that the throughput with ECN would be 7/6 larger, i.e.:

Thr = 40/(6RTT) = 40x1460x8/0.6 [bit/s] ≅ 0.78 [Mbit/s]

Exercise 2 (Solution time ~ 30min)

There are 6 standard methods (or primitives) in the SIP protocol.

REGISTER: Registers the user (URI) in the To: field to a Registrar server or modify its properties. Can affect also location and Proxy behavior.
INVITE: Attempt (invitation) to join a SIP session, can be a simple telephone call or a multimedia streaming session or anything else described by the SDP payload of the message.
ACK: Final answer to an INVITE establishing the call or connection to the SIP session from a specific terminal and UAC
BYE: A UAC announces the leaving of a SIP session.
CANCEL: Terminates one outstanding INVITE that has not received and ACK.
OPTIONS: Allows the negotiations of capabilities and additional parameters between the terminals (UACs and UASs) while establishing a session.

Other Methods have been defined by successive standards and more can be defined to specialize or customize SIP, but they are not mandatory and may not be implemented in all devices/slients.

The operations of SIP are carried out through the exchange of Methods, for example to establish a call an INVITE is sent by the Caller UAC and (if the call is answered) and ACK is sent back by the Callee (UAS). The figure below represents a typical successful call, where both users exploit Proxies.

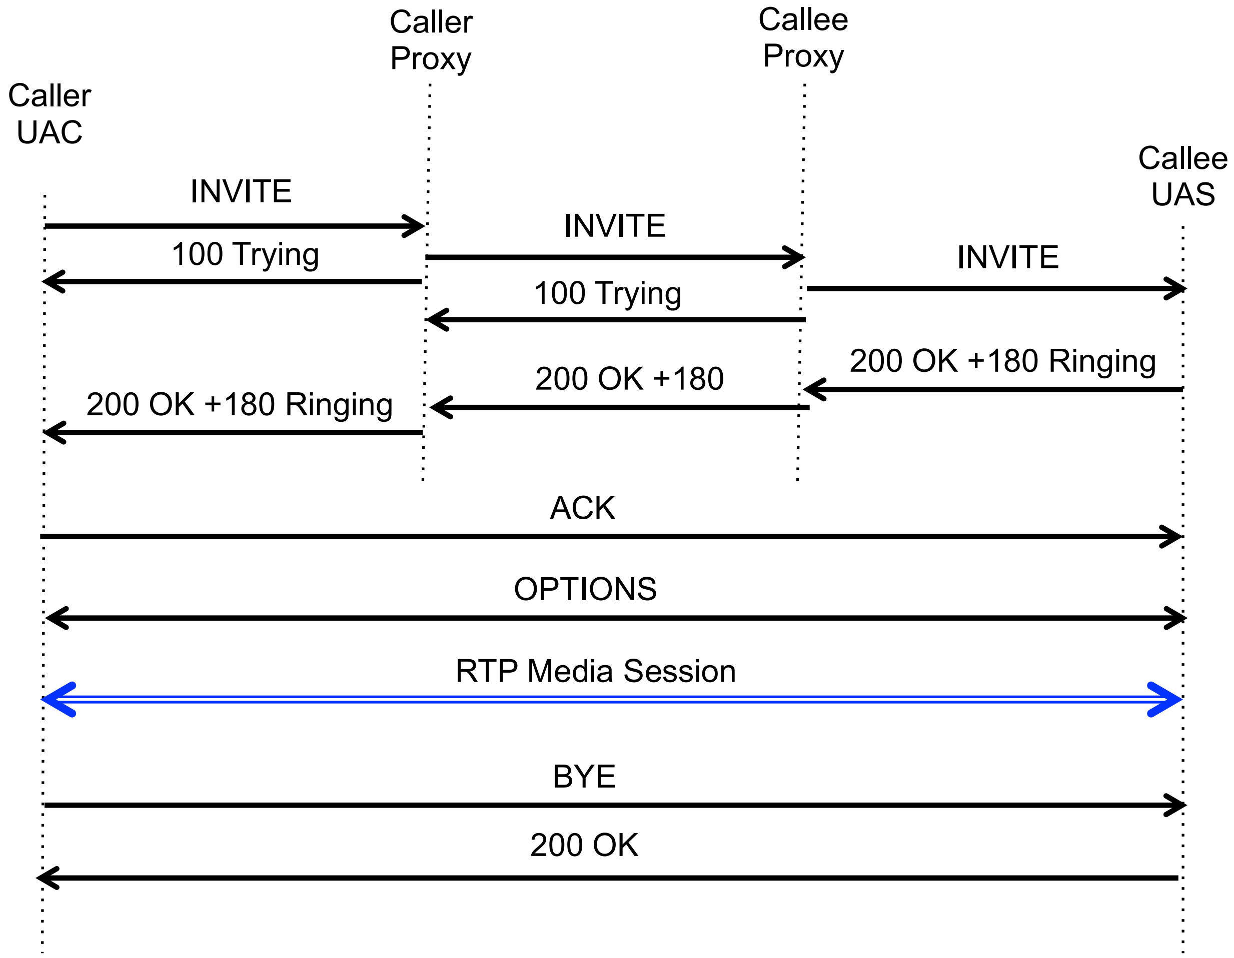

Proxies and Registrars are servers that mediate the SIP services for UAS, they are fundamental:

1) for reachability (Users can be on different terminals, behind NATs, etc.);
2) for accountability and tracking of usage.

Normally a Registrar is associated to a Proxy and it only does the authentication and registration procedures, while call processing is performed by Proxies.

Exercise 3 (Solution time ~ 40min)

A store-and-forward switch operates by receiving the entire frame, checking its integrity, then lookup the forwarding tables and transfer the frame to the desired output. The minimum transfer delay within the switch is 1 frame time plus the forwarding tables lookup time (assuming that integrity check time is negligible, which is valid only of integrity is computed directly in hardware as the frame arrives, otherwise if a CPU is used, then it cannot be negligible compared to the time to lookup tables).

A Cut-through switch instead does not memorize the frame, nor check its integrity, but the bytes are forwarded to the output in pipeline. The minimum transfer time is the time to lookup the forwarding table plus the time to receive the first bytes until the destination MAC is read. Let's, assume for the sake of this exercise solution, that 64 bytes are received before the forwarding tables are read.

Based on these considerations and the different operations of the two categories of switches, the switch transfer times Tsf and Tct of a single store-and-forward or cut-through switch respectively is:

$$Tsf = 5000 \text{x} 8 \text{x} 10^{-1} + 10 \text{ [ns]} = 4010 \text{ ns}$$
$$Tct = 64 \text{x} 8 \text{x} 10^{-1} + 10 \text{ [ns]} = 112.4 \text{ ns}$$

The picture below sketches the architecture of the data center considered, the redundancy of devices always present for reliability is not depicted, as it has no impact on the latency to be computed.

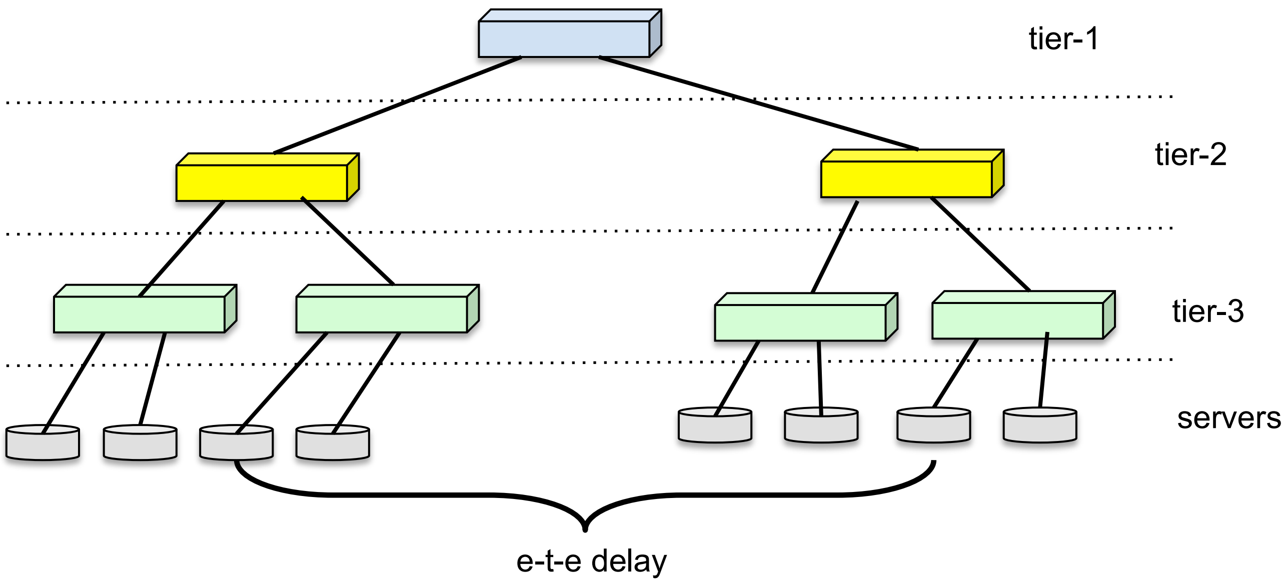

The e-t-e latency Le is given by 5 switch traversals, plus the propagation over 6 cables plus 1 frame time (the transmission from the source server to the first switch):

$$Le = 5Tst + 6 \text{x} 30/(3 \text{x} 10^{-1}) + Tf \text{ [ns]} = 5Tst + 4600 \text{ [ns]}$$

$$Lesf = 5 \text{x} 4010 + 4600 = 20050 + 4600 \text{ [ns]} = 24.650 \ \mu s$$

$$Lect = 5 \text{x} 112{,}4 + 4600 = 562 + 4600 \text{ [ns]} = 5.562 \ \mu s$$