

---

**Computer Communication**  
**Exam: EDA343, DIT423; Re-exam EDA344, LEU062**

---

*Time and Place:* Tuesday 31/05, 2022, 08.30-12.30

*Course Responsible:* Romaric Duvignau (EDA343, DIT423, LEU062), Marina Papatriantafilou (EDA344, DIT423) (Tel: 031 772 6976, 031 772 5413)

**Grading:**

CTH - EDA344, EDA343, LEU062: 30-40, 41-50, 51-60 → 3, 4, 5

GU - DIT423: 30-45, 46-60 → G, VG

**Allowed material:**

- English-X (X can be French, German, Swedish, etc) dictionary
- *No other books, no notes, no calculators, no electronic devices.*


*Instructions*

- **Write clearly your course-code (EDA343/DIT423/LEU062/EDA344)**
- **Start answering each assignment on a new page; use only one side of each sheet of paper; please sort the sheets according to the question-ordering and number them.**
- Write in a **clear manner** and **motivate** (explain, justify) your answers. If it is not clear what is written for some answer, it will be considered wrong. If some answer is not explained/justified, it will get **significantly** lower marking.
- If you make any **assumptions** in answering any item, do not forget to clearly state what you assume.
- A good **rule-of-thumb for the extent of detail to provide**, is to include enough information/explanation so that a person, whose knowledge on computer communication is at the level of our introductory lecture, can understand.
- Please answer in English, if possible. If you have large difficulty with that (with all or some of the questions) and you think that your grade might be affected, feel-free to write in Swedish.
- Inspection of exam: date and time will be announced on the front page of the course in the canvas system.

**Good Luck !!! Lycka till !!!!**

## 1. Computer Networks and the Internet (10pts)

Consider the following network, made of 2 computers (PC1, PC2), 2 servers (S1,S2) and 2 routers (R1,R2):




Above each link of the network is mentioned the values “ $D, B$ ” where  $D$  is the length of the link and  $B$  is the bandwidth of the link. In this exercise, we will assume that:

- All ping packets are exactly **1000 bits** long.
  - Propagation speed is exactly **200km/ms**.
  - Nodal processing time is exactly **1ms**.
  - Queuing delay is **0ms**.
  - Bandwidth is shared equally among concurrent TCP connections.
  - As usual, 1 Mbps =  $10^6$  bits per second = 1000 bits per ms.
- (a) (4pts) Under our assumptions, calculate the latency returned by running from PC1 (i) **ping** PC2 and (ii) **ping** S2. You can round your answer to closest millisecond.
- (b) (3pts) Suppose that using HTTP, PC1 downloads a very large file stored on the server S1. What will be the average throughput of the download assuming (i) this is the only traffic on the network and (ii) PC2 is simultaneously downloading a very large file from S2.
- (c) (3pts) The Internet is using a packet-switching approach for handling communication. There is an alternative approach called “circuit switching” that does not rely on *packets* in order to communicate. What are the pros and cons of using packet switching versus circuit switching? (Feel free to use a table to organize your answer).

## 2. Application Layer (10pts)

In this exercise, we suppose a client PC1 opens the webpage <https://www.cineflix.se/> and retrieves a URL to a video resource pointing to <https://video.cineflix.se/abcd>. Cineflix is a streaming service that uses DASH protocol to stream the videos to the clients and uses the CDN provider amassonCDN to host all its video content. As often, we assume that amassonCDN uses DNS to intercept and redirect the client requests, that is Cineflix will be sending an alias upon receiving a request for `video.cineflix.se`.

- (a) (5pts) Reproduce and complete the following figure illustrating the process on how PC1 gets to know using its Local DNS (LDNS) a CDN server that will provide it with the requested content:



On the figure, `dns.cineflix.se` is an authoritative name server on `cineflix.se` domain, `dns.amassonCDN.se` is an authoritative name server on `amassonCDN.se` domain, `dns.se` is a TLD name server for `.se` TLD. We assume both `www.cineflix.se`, `dns.cineflix.se` and `dns.se`'s IP addresses are already cached on LDNS (and nothing else) and PC1 has already retrieved `www.cineflix.se`'s IP address. Also, the LDNS resolves DNS queries in an iterative fashion.

You should add to the figure the different communication by using **numbered directed edges** (where the number indicates the order of the interaction). The first edge is given and corresponds to the HTTP GET request from PC1.

List all the edges you added and for each give a short explanation for what it corresponds to, eg "HTTP GET request" for (1).

- (b) (2pts) What information can amassonCDN use to decide which CDN server to return to the client?
- (c) (3pts) What are the advantages for DASH to use HTTPS and thus TCP/TLS compared with using UDP for better efficiency?


### 3. Transport Layer (10pts)

Let's consider in this exercise packets of fixed length of exactly  $L = 1100$  bytes, acknowledgments of negligible size ( $0$  bytes to simplify) and that a client  $C$  is communicating with a remote server  $S$  crossing several routers. We further assume that the access link of the client which has a bandwidth of  $R = 1100$  bytes/ms is the bottleneck link, that there are no packet losses and that the RTT between the client and the server is constant and equals to  $10$ ms. For TCP, we assume  $MSS = L$  and  $ssthresh = 8$  MSS.

- (a) (3pts) Suppose a "stop-and-wait" protocol is used (i.e., waiting for an acknowledgment before proceeding to next packet). What is then the maximum sending rate between the client and the server? (Illustrate your answer by a simple time-diagram; your answer can be in bytes/ms or KB/s).
- (b) (4pts) Suppose now we use TCP as the transport layer. After how much time will TCP start sending packets at the maximum sending rate? Illustrate your answer with a small graphic showing TCP's congestion window size in function of the number of elapsed RTT with the different TCP phases.
- (c) (3pts) Explain why TCP's congestion control algorithm is labelled as *Additive-Increase-Multiplicative-Decrease* (AIMD).

### 4. Network Layer (10pts)

Consider the following network topology:



The above network is made of 12 routers belonging to 3 Autonomous Systems (AS). The AS have been allocated with the following IP blocks **129.16.0.0/24** for AS1, **129.16.1.0/24** for AS2 and **129.16.2.0/24** for AS3. BGP is used for inter-AS routing while the 3 AS use an intra-AS routing protocol that minimizes the latency (displayed next to each link). No policy has been set for the route preferences by the administrators of the 3 AS, hence the rest of BGP's usual list of criteria is used here. We assume the routers use both destination-based forwarding and longest prefix matching.

- (3pts) Some links of the network will never get used by any traffic. List all such links.
- (3pts) Provide the routing table of router's 1c assuming all a-routers have one interface configured with an IP address with 1 as host part, b-routers with 2 as host part, c-routers with 3 and d-routers with 4 (to simplify, we assume 1 IP address only per router).
- (4pts) Let's suppose AS1 and AS2 decide to merge. What will be their new IP address block? Will this have consequences on the routers and the routing tables and which ones?


## 5. Link Layer and Wireless (10pts)

Let's consider a hypothetical link-layer protocol ALOHA2 that provides *error detection* and *multiple access* to a shared link. Error detection is obtained using single-bit parity check (even parity) by adding an extra parity bit for every seven bits of the original data. Multiple access is done using the slotted-ALOHA protocol using a probability of 1/2 for sending data during each timeslot. The goal of the exercise is to compare ALOHA2 with Ethernet and 802.11 link-layer protocols.

- (2pts) Review how a single parity-bit is working by giving 2 examples of 8-bits words: 1 byte where no error is detected and 1 byte that contains at least one error.
- (2pts) Now consider a short acknowledgment link-layer packet of **70 bytes** (excluding all error-detection bits), how many bytes will need to be sent with ALOHA2 and with Ethernet/Wifi (both using 32-CRC, a 32-bits Cyclic Redundancy Check)? What about sending a **1400 bytes** long packet?
- (2pts) Actually, **32-CRC** looks very similar to a message's hash. Why would it be though a very poor choice to sign messages?
- (4pts) ALOHA2 will not be very efficient to share the channel as the probability that one out of  $n$  senders successfully sends its transmission during one timeslot would be  $n/2^n$ . What are the mechanisms in place in Ethernet's protocol that allow it to be much more channel-efficient than ALOHA2? (No need to enter in details here, just mention the different mechanisms and compare them with ALOHA2). Are all of these mechanisms also used in Wifi's protocol and if not, why is that so?

## 6. Software-Defined Networking (10pts)

Let's consider the following network:



In the above network, each link connecting 2 routers is labelled  $\ell, B$  where  $\ell$  is a latency measure (in ms) and  $B$  is the bandwidth capacity of the link (in Gbps).

The 2 clients are both communicating with the server **h3** (with IP 10.0.0.3) but each is rather interested in utilizing different routes: **h1** is interested by the route with lowest latency whereas **h2** would like to have the highest possible throughput for the connection. SDN will be used to make this possible and satisfy both clients.

- (a) (4pts) Assume first that destination-based forwarding is in place on the routers. Draw a directed graph showing the paths used by packets with destination **h3** in the following two situations:
1. All the routes have been configured to use the lowest latency path (equivalent to running least-cost Dijkstra algorithm with latency  $\ell$  as cost).
  2. All the routes have been configured to use the maximum bandwidth path (equivalent in this graph to running least-cost Dijkstra algorithm with  $1/B$  as cost).

That is, for each case, determine the next-hop router for packet with destination **h3**.

- (b) (3pts) The administrator of the network decided to configure the routers so that, by default, the low-latency routes are used for routing packets based on their destination. On top of those already configured flows, we will add flows so that **h2** uses the high-bandwidth paths to reach **h3**. Copy and complete the following table with the necessary flows to add for this to happen:

Switch to add a new flow	Match fields	Actions
...	...	...

Use eg `IP,nw_src=...` and `IP,nw_dst=...` to specify source/destination IP addresses and `output:x` to forward packet on the interface that is connected to x.

- (c) (3pts) Suppose that **h1** was sending UDP traffic whereas **h2** was sending TCP traffic. Would it be possible to make the same paths that we configured using SDN but keeping destination-based forwarding and using some priority-based *Packet Scheduler* on each router? (here UDP packets can be one class of traffic and TCP another class).