

Communication Networks

Prof. Laurent Vanbever

Exercise 1 - Introduction

Internet Structure

1.1 Layer Model


Internet protocol stack	OSI reference model
Application	Application
	Presentation
	Session
Transport	Transport
Network	Network
Link	Link
Physical	Physical

Internet communication layers: Internet protocol stack and the OSI reference model.

Communication over the Internet can be decomposed into independent layers. In the lecture, we have discussed the Internet protocol stack which contains 5 layers. Another often used model is the OSI (Open Systems Interconnection) model with 7 layers. Find the best matching layer for the following operations/devices. You can use the already known 5-layer model. **Hint:** Some of the operations could be implemented in different layers.

- a) Bit-to-bit transmission between two devices.
- b) Encryption of a message.
- c) A switch in a network.
- d) Routing path search.
- e) Adding a sequence number to each packet.
- f) A router in a network.
- g) A middlebox in a network performing deep packet inspection (DPI) to find malware in Web traffic.

1.2 Internet Organization



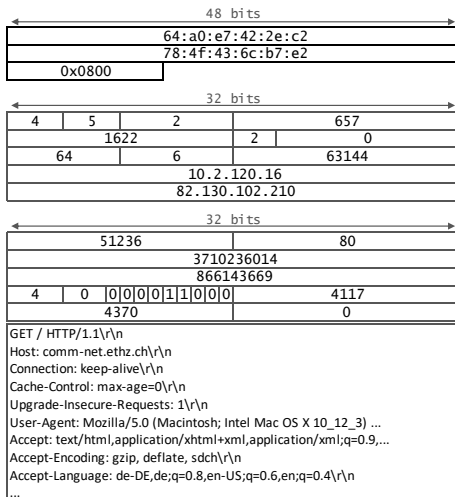
A network of multiple autonomous systems (AS).

The network on the left consists of multiple autonomous systems (AS). Single-headed arrows point from providers to their customers. Double-headed arrows represent peer connections.

- For each AS, identify if it is a Tier-1, Tier-2 or Tier-3 ISP network or an IXP.
- AS7 has two different providers (AS1 and AS2). How is this type of interconnection called? What are the advantages of multiple different providers for AS7? Can you see any disadvantages?

1.3 Internet Communication

The figure on the left shows a (simplified version of a) packet that was recorded at your machine's network interface using Wireshark^a.



Packet consisting of data in four layers (link, network, transport and application layer).

- What kind of activity does this packet belong to?
- From the lecture, you know that a packet is composed of data in different layers. The fields in the left-hand figure are already grouped to four blocks representing the link, network, transport and application layer. Do you know which protocol is used in each of the four layers? *Hint: You may want to have a look at the list of protocols on Wikipedia^b.*
- After being sent out at your machine's network interface, the packet will traverse a switch which will forward it based on the destination address. On which layer does the switch operate and which is the destination address it will look at?
- After leaving your local network, the packet will traverse a router. On which layer does the router operate and which is the destination address it will look at?

^a<https://www.wireshark.org/>

^bhttps://en.wikipedia.org/wiki/Internet_protocol_suite

1.4 Decipher an Internet Path using Traceroute

When two end-systems (hosts or servers) communicate with each other over the Internet, their communication traverses multiple routers.

The traceroute (or `tracert` on Windows) utility lists the routers that are on the path from one end-system to another. E.g., if you type `traceroute <target>` in the command line, where `<target>` is an IP address or a domain name, that will display a list of router IP addresses and/or domain names and the round-trip-times (RTTs) that were measured between your computer and each router on the path to `<target>`.

There also exist traceroute servers (e.g., see www.traceroute.org) that allow to traceroute from them to any other computer.

```
1 wblindex.sdv.fr 212.95.66.126 0.102 ms
2 border-gateway2.sdv.fr 212.95.69.227 0.85 ms
3 ipv4.de-cix.fra.de.as559.switch.ch 80.81.196.147 9.221 ms
4 swiB51-100GE-0-0-0-0.switch.ch 130.59.37.34 9.437 ms
5 swiPS1-100GE-0-0-1-3.switch.ch 130.59.37.190 10.535 ms
6 swiPS2-100GE-0-0-1-4.switch.ch 130.59.37.58 10.476 ms
7 swiZH3-100GE-0-0-0-2.switch.ch 130.59.36.170 11.268 ms
8 swiZH1-B1.switch.ch 130.59.37.65 11.117 ms
9 swiEZ2-B2.switch.ch 130.59.38.109 13.437 ms
10 swiEZ3-B1.switch.ch 130.59.36.126 11.066 ms
11 rou-gw-lee-tengig-to-switch.ethz.ch 192.33.92.1 10.915 ms
12 rou-fw-rz-rz-gw.ethz.ch 192.33.92.169 10.916 ms
13 cms-publish.ethz.ch 129.132.19.216 10.545 ms
```

Traceroute output from a server in France to ethz.ch

We have already used `http://traceroute.sdv.fr` to traceroute from a server in France to `ethz.ch` (cf. left-hand figure):

- How many routers are there between the server in France and `ethz.ch`?
- How many of these routers are inside the ETHZ network? How many are inside ETHZ's Internet Service Provider (ISP)?

```
traceroute to 129.132.19.216 (129.132.19.216), 30 hops max, 40 byte packets
 1 core-ns-router (128.112.128.2) 0.930 ms 0.828 ms 0.847 ms
 2 rtr-core-west-router.princeton.edu (128.112.12.229) 0.910 ms 0.666 ms 0.741 ms
 3 fw-border-hprc-router.princeton.edu (128.112.12.14) 1.103 ms 1.129 ms 1.043 ms
 4 rtr-border-hprc-router.princeton.edu (204.151.48.253) 1.475 ms 1.403 ms 1.251 ms
 5 172-96-130.unassigned.userdns.com (172.96.130.49) 3.524 ms 3.058 ms 3.187 ms
 6 bundle-ether240.200.core1.newy2laaa.net.internet2.edu (163.253.5.38) 6.312 ms 5.722 ms 4.129 ms
 7 198.71.45.237 (198.71.45.237) 77.297 ms 75.866 ms 76.052 ms
 8 ae7.mxl1.gen.ch.geant.net (62.40.98.238) 83.937 ms 83.921 ms 83.851 ms
 9 swicel-100ge-0-3-0-1.switch.ch (62.40.124.22) 85.288 ms 87.816 ms 85.516 ms
10 swiEZ2-400GE-0-0-0-0.switch.ch (130.59.38.82) 90.262 ms 90.852 ms 89.309 ms
11 swiEZ3-B1.switch.ch (130.59.36.126) 88.471 ms 88.290 ms 88.074 ms
12 rou-gw-lee-tengig-to-switch.ethz.ch (192.33.92.1) 88.290 ms 88.272 ms 88.239 ms
13 rou-fw-rz-rz-gw.ethz.ch (192.33.92.169) 88.514 ms 88.316 ms 88.252 ms
14 * * *
15 * * *
16 * * *
17 * * *
18 * * *
19 * * *
20 * * *
21 * * *
22 * * *
23 * * *
24 * * *
25 * * *
26 * * *
27 * * *
28 * * *
29 * * *
30 * * *
```

Traceroute output from a server in Princeton to ethz.ch


Next, we have tracerouted from a server in Princeton University (<https://www.net.princeton.edu/traceroute.html>) to `ethz.ch` (cf. left-hand figure):

- Between which routers does the path cross the Atlantic?
- Why do we see asterisks at certain lines of the traceroute output?
- Does the network path from the server in France to `ethz.ch` overlap with the path from the Princeton server to `ethz.ch`?
- If we were to traceroute again tomorrow from the Princeton server to `ethz.ch`, do you think that the traceroute output would necessarily be the same? Why?

Bandwidth and Delay

1.5 Network Characterization

Calculate and compare the bandwidth and the delay for different communication methods. (For this task, assume that the delay only consists of propagation delay.)




The pneumatic tube network of Paris in 1907.

- a) Pigeon post: Pigeons can be used as messengers. They are trained to transport messages from one location to another. Assuming you want to send a USB flash drive with 16 GB^a from Zürich to a friend in Paris (500 km) (only in one direction). Calculate the bandwidth and the delay for one pigeon carrying the USB drive and traveling at an average speed of 80 km/h.
- b) Pneumatic tube: These systems were introduced in the late 19th century to transport small, urgent items within buildings or even within cities. The capsules travel at an average speed of 8 meters per second. Assuming you send an external hard drive with 2 TB of storage through a tube from ETZ to the main building (distance 400 m), calculate the bandwidth and delay.
- c) AWS Snowmobile: Amazon uses a truck to move data from its customers to their data center. The truck houses a container which can store 100 PB of data.^b Assuming the truck is transporting data from New York to an AWS data center in San Francisco (distance 4700 km) at an average speed of 100 km per hour, calculate the bandwidth and the delay.

^a1 GB = 10⁹ bytes, 1 TB = 10¹² bytes, 1 PB = 10¹⁵ bytes

^b<https://techcrunch.com/2016/11/30/amazon-will-truck-your-massive-piles-of-data-to-the-cloud-with-an-18-wheeler/>

1.6 Packet order, parallel paths, and performance metrics



Two laptops connected by two links and a switch.


Consider the network on the left, consisting of laptops A and B connected by two links and a switch. Each link has bandwidth $R = 10\text{ Mbps}$, length $l = 10\text{ km}$, and propagation speed $c = 2 \cdot 10^8\text{ m/s}$. The switch has zero processing delay, an infinite buffer, and performs packet switching: the switch can start processing and transmitting a packet on an egress link only as soon as it has received the *full* packet on an ingress link.

Laptop A sends two back-to-back packets to B; the first packet has size $L_1 = 2\text{ kbit}$, the second packet has size $L_2 = 1\text{ kbit}$.

- What is the transmission delay experienced by each packet until it reaches B?
- What is the propagation delay experienced by each packet until it reaches B?
- What is the queueing delay experienced by each packet until it reaches B?
- What is the total transfer time, i.e., the time it takes from the moment A starts transmitting the first bit of the first packet until the moment B receives the last bit of the last packet?
- What is the throughput of transferring the two packets from A to B? Is that the same as the link bandwidth? Why (not)?

Assume now that we change the order of packets, i.e., laptop A first sends the 1 kbit packet, followed by the 2 kbit packet.

- Which of the aforementioned performance metrics will change compared to the initial scenario? Calculate the changed metrics.



Two laptops connected by two parallel paths.


Consider now the network on the left, which is the original network with an additional parallel path between the two laptops. The two paths are similar and do not share any links.

Laptop A starts sending the two packets simultaneously, one on each path, using the full bandwidth of each link.


- Which of the aforementioned performance metrics will change compared to the initial scenario? Calculate the changed metrics.

Sharing

1.7 Packet vs. Circuit Switching



Network with a shared link.



Demand distributions for node A and B.

Consider the network on the left. A and B are sending data towards C over a switch S and a shared link. All the links in the network have a bandwidth of 10 Mbps and you can assume that the propagation and switch processing time are negligible. For circuit switching, assume that (end-to-end) circuit establishment and teardown each take 50 ms. For packet switching, you can assume that switch S already knows how to reach C.

- a) How long does it take if node A is sending a 50 Mbit file to C using packet switching? B sends nothing.
- b) How long does it take if node B is sending a 50 Mbit file to C using circuit switching? A sends nothing.

Assume now that A and B are using packet switching and are each sending a 50 Mbit file to C at the same time.

- c) What will happen if the switch has no buffer?
- d) How large has the buffer to be (in Mbit) such that both senders can successfully transmit their files to C when they are simultaneously sending at full speed?
- e) For obvious reasons, switch buffers cannot be of unlimited size. How is it possible to successfully send data over the Internet even if some packets are dropped due to full buffers?

Finally, assume that A and B have to send data with a demand according to the diagram on the left.

- f) How long does it take to send all data if A and B use circuit switching (reserving for the peak demand)?
- g) How long does it take to send all data if A and B use packet switching (you can assume an unlimited buffer size on S)?