Chair of Distributed Systems and Security
School of Computation, Information and Technology
Technical University of Munich

Computer Networking and IT Security (INHN0012)

Tutorial 1

Problem 1 Layer models

In this exercise, a layer model consisting of a total of **3 layers** is to be developed, which describes the writing, sending, receiving and reading of an advertising brochure. Since most recipients do not read ads, we assume it is the menu from the nearest pizza place that the recipient is very interested in.

a)* Is the sending of advertising material a *bidirectional* communication, i. e. will the recipient respond by mail?

No. The recipient may order a pizza, but he will not do that via mail. It is therefore a unidirectional communication.

b)* The figure below serves as a template for the layer model. Think of meaningful names for the missing layers as well as the transmission channel and add them to the figure.

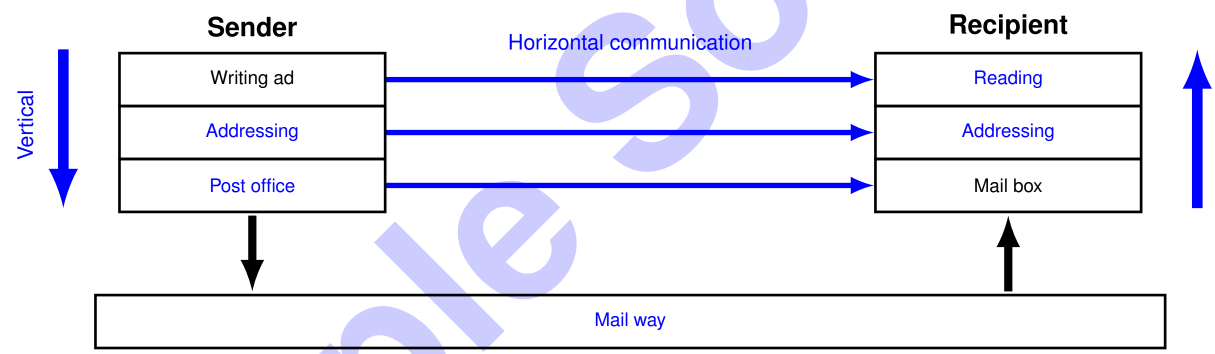

c) Describe which services are offered by the different layers.

Sender:

- Writing: Advertising text is put on paper (presentation of information in written form)
- Addressing: The brochure is packed in an envelope, which is marked with the sender's and recipient's address
- Post office: The letter is taken (along with many others) to the post office and sent out

Receiver:

- Mailbox: The letter is delivered and placed in the recipient's mailbox
- Addressing: The recipient usually checks again whether the delivered letter was really addressed to him, and is then taken out of the envelope
- Reading: The brochure included in the letter is read

d) What is meant by *horizontal* and *vertical communication* in the context of layered models? Draw both communication types in the figure from Subproblem b).

Vertical communication: communication between layer N and $N-1$ the layers on the respective system.
Horizontal communication: communication between N layers on different hosts.

We now have a closer look at Layer 2. From the lecture we know the following illustration:

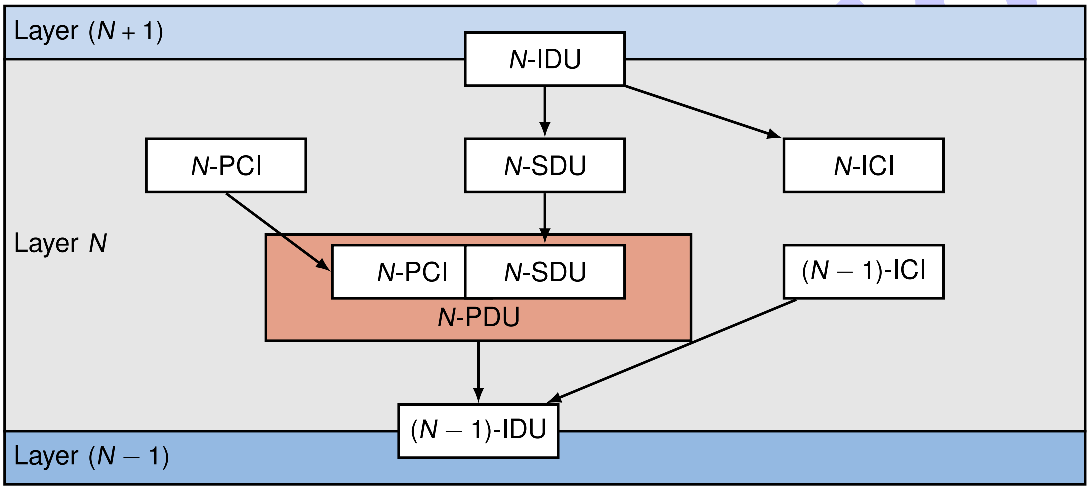

e)* Which parts of the letter correspond to the PCI (Protocol Control Information), SDU (Service Data Unit) and PDU (Protocol Data Unit) from a layer 2 perspective?

- PCI: The address information located on the envelope
- SDU: The advertising brochure itself, that is, the contents of the letter
- PDU: The sealed and labeled letter

Problem 2 Binary prefixes

The difference between binary and SI prefixes always causes confusion. The problem consists in contradictory information, especially on the part of the operating systems: The memory allocation on mass storage devices such as hard disks is commonly counted with binary prefixes while the stated units contain SI prefixes.
An example: you are buying a hard disk with a manufacturer's stated capacity of 3 TB. In the small print on the packaging you will find the note „1 TB = 10^{12} B". Obviuosly, SI prefixes are used. Now assume that your operating system counts the storage using binary prefixes.

SI prefix		Value	Binary prefix		Value
k	(kilo)	10^3	Ki	(Kibi)	2^{10}
M	(Mega)	10^6	Mi	(Mebi)	2^{20}
G	(Giga)	10^9	Gi	(Gibi)	2^{30}
T	(Tera)	10^{12}	Ti	(Tebi)	2^{40}
P	(Peta)	10^{15}	Pi	(Pebi)	2^{50}

Table 2.1: Comparison between SI and binary prefixes

a)* State the capacity of your hard disk in TiB .

$$3\,\text{TB} = 3 \cdot 10^{12}\,\text{B} = \frac{3 \cdot 10^{12}}{2^{40}}\,\text{TiB} \approx 2.73\,\text{TiB}$$

b)* Determine the percentage difference between SI and binary prefixes for all values given in Table 2.1.

$$\frac{\text{k}}{\text{Ki}} = \frac{10^3}{2^{10}} \approx 97.66\,\% \quad \Rightarrow \quad e = 2.34\,\%$$

$$\frac{\text{M}}{\text{Mi}} = \frac{10^6}{2^{20}} \approx 95.37\,\% \quad \Rightarrow \quad e = 4.63\,\%$$

$$\frac{\text{G}}{\text{Gi}} = \frac{10^9}{2^{30}} \approx 93.13\,\% \quad \Rightarrow \quad e = 6.87\,\%$$

$$\frac{\text{T}}{\text{Ti}} = \frac{10^{12}}{2^{40}} \approx 90.95\,\% \quad \Rightarrow \quad e = 9.05\,\%$$

$$\frac{\text{P}}{\text{Pi}} = \frac{10^{15}}{2^{50}} \approx 88.82\,\% \quad \Rightarrow \quad e = 11.18\,\%$$

Note: The specification of binary prefixes is only usual for byte values. Bit values, e. g. kbit or Mbit , are specified with SI prefixes only.

Problem 3 Sneakernet

BigCorp™ is moving its data to the cloud as it has outgrown its data centre in Akron, Ohio. In order to do so, it has to move all of its data to Amazing Web Services (AWS), the cloud provider it chose to go with. Moving the total of 91 PB is no easy task. Therefore, BigCorp™ has chosen to shut down all of its services for as long as necessary, move the data and finally restart all services in the cloud.

a)* How long does it take to transfer the data using BigCorp's high-speed 8 Gbit/s internet uplink? As for now, ignore the time it takes to serialize the data onto the link. Compare the given uplink with your uplink at home.

$$t = \frac{91\,\text{PB}}{8\,\text{Gbit/s}} = \frac{728\,\text{Pbit}}{8\,\text{Gbit/s}} = \frac{728 \cdot 10^6\,\text{Gbit}}{8\,\text{Gbit/s}} = 91 \cdot 10^6\text{s} \approx 25\,277.78\,\text{h} \approx 1053.24\,\text{d} \approx 2.89\,\text{yr}$$

Seeing these numbers, BigCorp™ realizes that it clearly cannot shut down its services for so long without going out of business. AWS offers a product called Snowmobile: a truck, able to carry 100 PB, is sent to the customer's on-premise location. Once arrived, the customer transfers the data onto the storage medium inside the truck. Finally, the Snowmobile is sent to the nearest AWS location and the data is transferred onto AWS' systems.

b) Assume the infrastructure and Snowmobile allow loading and unloading data via eight parallel 400 Gbit/s fibre connections. What is the total data rate? How long does it take to load or unload all of BigCorp's data to and from the Snowmobile?

$$r_{load} = r_{unload} = 400\,\text{Gbit/s} \cdot 8 = 3200\,\text{Gbit/s} = 3.2\,\text{Tbit/s}$$

$$t_{load} = t_{unload} = \frac{91\,\text{PB}}{3.2\,\text{Tbit/s}} = \frac{728 \cdot 10^3\,\text{Tbit}}{3.2\,\text{Tbit/s}} = 227.5 \cdot 10^3\,\text{s} \approx 2.63\,\text{d}$$

Assume that the distance between BigCorp's datacentre and AWS' data centre is 186 km. The truck travels at an average of 93 km/h.

c) For this subtask, assume the entire 100 PB capacity of the truck as utilized. Given the capacity and travel time, what data rate in Tbit/s does the truck achieve? For this subtask, ignore the time it takes to load and unload the truck.

$$t_{travel} = \frac{186\,\text{km}}{93\,\text{km}/\text{h}} = 2\,\text{h} = 7200\,\text{s}$$

$$r_{fulltruck} = \frac{100\,\text{PB}}{7200\,\text{s}} = \frac{800 \cdot 10^3\,\text{Tbit}}{7200\,\text{s}} \approx 111.11\,\text{Tbit}/\text{s}$$

d) Now consider the amount of data BigCorp™ actually has, as well as the time it takes to load and unload the truck. How long does the entire data transfer take? Which data rate in Gbit/s is achieved overall?

$$t_{all} = t_{load} + t_{travel} + t_{unload} = 227.5 \cdot 10^3\,\text{s} + 7200\,\text{s} + 227.5 \cdot 10^3\,\text{s} = 462.2 \cdot 10^3\,\text{s} \approx 5.35\,\text{d}$$

$$r_{overall} = \frac{91\,\text{PB}}{462.2 \cdot 10^3\,\text{s}} = \frac{728 \cdot 10^3\,\text{Tbit}}{462.2 \cdot 10^3\,\text{s}} = 1575.03\,\text{Gbit}/\text{s}$$

e) Assume that BigCorp™ still cannot shut down its business for that long. Which alternative does it have to moving these vast amounts of data at once. Consider how real networks function and try applying it to this scenario. Sketch how BigCorp™ can reduce the downtime of the entire system.

BigCorp™ can split its data into smaller parts, given that it can migrate parts of its system at a time (e. g. one service at a time). Splitting the data into smaller chunk allows cutting down on the time needed to load and unload the truck, which is the majority of time needed. Therefore, smaller chunks can be transferred quicker than all data at once, allowing to migrate one service at a time, causing multiple shorter downtimes. This allows to e.g. move the entire data piece by piece over multiple weekends.

"Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway." — Andrew S. Tanenbaum