

Solution of Exercise Sheet 1

Exercise 1 (Different Client-Server Scenarios)

Company X runs 8,000 computer workplaces.

- Scenario 1: Fat clients (PC)
 - Electrical power rating per desktop: 400 watts
 - Electrical power rating per screen: 100 watts
- Scenario 2: Thin clients
 - Electrical power rating per thin client: 30 watts
 - Electrical power rating per screen: 100 watts
 - Electrical power rating per server blade: 400 watts
 - Each server blade has enough resources to run 40 virtual desktops

What are the electricity costs per year for 24/7 operation when the electricity price is 0,24€/kWh?

Scenario 1:

Electricity costs per year (including the leap year) for 8,000 computer workplaces:

$$0.5 \text{ kW} * 24 \frac{\text{h}}{\text{Day}} * 365.25 \frac{\text{Day}}{\text{Year}} * 0.24 \frac{\text{€}}{\text{kWh}} * 8,000 = 8,415,360 \frac{\text{€}}{\text{Year}}$$

Scenario 2:

Electricity costs per year (including the leap year) for 8,000 computer workplaces:

$$0.13 \text{ kW} * 24 \frac{\text{h}}{\text{Day}} * 365.25 \frac{\text{Day}}{\text{Year}} * 0.24 \frac{\text{€}}{\text{kWh}} * 8,000 = 2,187,993.6 \frac{\text{€}}{\text{Year}}$$

200 server blade are required to run the 8.000 computer workplaces.

Electricity costs per year (including the leap year) for 200 server blades.

$$0.4 \text{ kW} * 24 \frac{\text{h}}{\text{Day}} * 365.25 \frac{\text{Day}}{\text{Year}} * 0.24 \frac{\text{€}}{\text{kWh}} * 200 = 168,307.2 \frac{\text{€}}{\text{Year}}$$

Electricity costs per year for the computer workplaces and the server blades.

$$2,187,883.6 \frac{\text{€}}{\text{Year}} + 168,307.2 \frac{\text{€}}{\text{Year}} = 2,356,190.8 \frac{\text{€}}{\text{Year}}$$

Exercise 2 (Types of Clients)

Four types of clients exist in the client-server model. Describe the four types. Focus the differences.

- **X-Terminal** or **Text-Terminal**
 - Only display the (graphical) user interface and transfer the user interaction to the server
 - Calculation of the (graphical) user interface, data processing and data storage, data management are tasks of the server
- **Thin Clients** or **Zero Clients**
 - Calculate and display the graphical user interface
- **Applet Clients** or **Network Computers**
 - Calculate and display the graphical user interface and do a part of the data processing
 - The clients process the applications (applets) themselves
- **Fat Clients**
 - Only data management and data storage are located on the (file or database) server

Exercise 3 (Storing and transmitting Data)

Common assumptions about data are:

- It is easy to store data today.
- It is easy to transport or transmit data today.

In this exercise, we verify the correctness of these statements.

1. A scientific experiment produces 15 PB of data per year, which need to be stored. What is the height of a stack of storage media, if for storing the data...
 - CDs (capacity: 600 MB = $600 * 10^6$ Byte, thickness: 1.2 mm) are used?
 - DVDs (capacity: 4.3 GB = $4.3 * 10^9$ Byte, thickness: 1.2 mm) are used?
 - Blu-rays (capacity: 25 GB = $25 * 10^9$ Byte, thickness: 1.2 mm) are used?
 - HDDs (capacity: 2 TB = $2 * 10^{12}$ Byte, thickness: 2.5 cm) are used?

Attention: Calculate the solutions for both options:

- 15 PB = $15 * 10^{15}$ Byte \Leftarrow this way, the hardware manufacturer calculate
- 15 PB = $15 * 2^{50}$ Byte \Leftarrow this way, the operating systems calculate

*Solution for CDs with 15 PB = $15 * 10^{15}$ Byte:*

Number of CDs: $\frac{15 * 10^{15} \text{ Byte}}{600 * 10^6 \text{ Byte}} = 25,000,000$

CD stack height: $25,000,000 * 1.2 \text{ mm} = 30,000,000 \text{ mm}$
 $= 3,000,000 \text{ cm}$
 $= 30,000 \text{ m}$
 $= 30 \text{ km}$

*Solution for CDs with 15 PB = $15 * 2^{50}$ Byte:*

Number of CDs: $\frac{15 * 2^{50} \text{ Byte}}{600 * 10^6 \text{ Byte}} = 28,147,498$

CD stack height: $28,147,498 * 1.2 \text{ mm} = 33,776,997.6 \text{ mm}$
 $= 3,377,699.76 \text{ cm}$
 $= \text{approx. } 33,777 \text{ m}$
 $= 33.78 \text{ km}$

*Solution for DVDs with 15 PB = $15 * 10^{15}$ Byte:*

Number of DVDs: $\frac{15 * 10^{15} \text{ Byte}}{4.3 * 10^9 \text{ Byte}} = 3,488,372.093$

An integer number is required $\implies 3,488,373$

DVD stack height: $3,488,373 * 1.2 \text{ mm} = 4,186,047.6 \text{ mm}$
 $= 418,604.76 \text{ cm}$
 $= 4,186.0476 \text{ m}$
 $= 4.1860476 \text{ km}$

*Solution for DVDs with 15 PB = $15 * 2^{50}$ Byte:*

Number of DVDs: $\frac{15 * 2^{50} \text{ Byte}}{4.3 * 10^9 \text{ Byte}} = 3,927,557.814$

An integer number is required $\implies 3,927,558$

DVD stack height: $3,927,558 * 1.2 \text{ mm} = 4,713,069.6 \text{ mm}$
 $= 471,306.96 \text{ cm}$
 $= 4,713.0696 \text{ m}$
 $= 4.7130696 \text{ km}$

*Solution for Blu-rays with 15 PB = 15 * 10¹⁵ Byte:*

Number of Blu-rays: $\frac{15 \cdot 10^{15} \text{ Byte}}{25 \cdot 10^9 \text{ Byte}} = 600,000$

Blu-ray stack height: $600,000 * 1.2 \text{ mm} = 720,000 \text{ mm}$
 $= 72,000 \text{ cm}$
 $= 720 \text{ m}$

*Solution for Blu-rays with 15 PB = 15 * 2⁵⁰ Byte:*

Number of Blu-rays: $\frac{15 \cdot 2^{50} \text{ Byte}}{25 \cdot 10^9 \text{ Byte}} = 675,539.944$
An integer number is required $\implies 675,540$

Blu-ray stack height: $675,540 * 1.2 \text{ mm} = 810,648 \text{ mm}$
 $= 81,064.8 \text{ cm}$
 $= 810.648 \text{ m}$

*Solution for HDDs with 15 PB = 15 * 10¹⁵ Byte:*

Number of HDDs: $\frac{15 \cdot 10^{15} \text{ Byte}}{2 \cdot 10^{12} \text{ Byte}} = 7,500$

HDD stack height: $7,500 * 2.5 \text{ cm} = 18,750 \text{ cm}$
 $= 187.5 \text{ m}$

*Solution for HDDs with 15 PB = 15 * 2⁵⁰ Byte:*

Number of HDDs: $\frac{15 \cdot 2^{50} \text{ Byte}}{2 \cdot 10^{12} \text{ Byte}} = 8,444.2493$
An integer number is required $\implies 8,445$

HDD stack height: $8,445 * 2.5 \text{ cm} = 21,112.5 \text{ cm}$
 $= 211.125 \text{ m}$

2. The data of the scientific experiment is transmitted via networks that use fiber-optic cables and provide a bandwidth of 40 Gbit/s.
- How long does it take to transfer the 15 PB via a 40 Gbit/s network?
 - How long does it take to transfer the 15 PB via a 100 Mbps Ethernet?

Attention: Calculate the solutions for both options:

- 15 PB = 15 * 10¹⁵ Byte
- 15 PB = 15 * 2⁵⁰ Byte

*Solution for the 40 Gbit/s network with 15 PB = 15 * 10¹⁵ Byte:*

40 Gbit/s bandwidth: $40 \text{ Gbit/s} = 40,000,000,000 \text{ Bit/s}$
 $= 5,000,000,000 \text{ Byte/s}$

Duration of transmission: $\frac{15 \cdot 10^{15} \text{ Byte}}{5 \cdot 10^9 \text{ Byte/s}} = 3 \cdot 10^6 \text{ s} = 3,000,000 \text{ s}$
 $= 50,000 \text{ m}$
 $= 833.333333333 \text{ h}$
 $= 34.722222222 \text{ d}$

\implies approx. 34 Days, 17 Hours, 20 Minutes

*Solution for the 40 Gbit/s network with 15 PB = 15 * 2⁵⁰ Byte:*

40 Gbit/s bandwidth: $40 \text{ Gbit/s} = 40,000,000,000 \text{ Bit/s}$
 $= 5,000,000,000 \text{ Byte/s}$

Duration of transmission: $\frac{15 \cdot 2^{50} \text{ Byte}}{5 \cdot 10^9 \text{ Byte/s}} = 3,377,699.72 \text{ s}$
 $= \text{approx. } 56,295 \text{ m}$
 $= \text{approx. } 938.25 \text{ h}$
 $= \text{approx. } 39.09 \text{ d}$

\implies approx. 39 Days, 2 Hours, 15 Minutes

*Solution for the Ethernet network with 15 PB = 15 * 10¹⁵ Byte:*

Ethernet bandwidth: $100 \text{ Mbit/s} = 100,000,000 \text{ Bit/s}$
 $= 12,500,000 \text{ Byte/s}$

Duration of transmission: $\frac{15 \cdot 10^{15} \text{ Byte}}{12,500,000 \text{ Byte/s}} = 1,200,000,000 \text{ s}$
 $= 20,000,000 \text{ m}$
 $= 333,333.3333 \text{ h}$
 $= 13,888.88889 \text{ d}$
 $= 38.02570538 \text{ y}$

\implies approx. 38 Years, 13 Days, 21 Hours, 20 Minutes

*Solution for the Ethernet network with 15 PB = 15 * 2⁵⁰ Byte:*

Ethernet bandwidth: $100 \text{ Mbit/s} = 100,000,000 \text{ Bit/s}$
 $= 12,500,000 \text{ Byte/s}$

Duration of transmission: $\frac{15 \cdot 2^{50} \text{ Byte}}{12,500,000 \text{ Byte/s}} = 1,351,079,888 \text{ s}$
 $= 22,517,998.13 \text{ m}$
 $= 375,299.9688 \text{ h}$
 $= 15,637.4987 \text{ d}$
 $= 42.81313812 \text{ y}$
(each year has 365,25 days!)

\implies approx. 42 Years, 296 Days, ...

Exercise 4 (Laws and Limitations)

1. What is the central statement of Moore's law?

The number of transistors per area unit on an integrated circuit doubles every 24 months.

2. What is the Von Neumann bottleneck?

Each additional CPU decreases the relative performance gain.

3. How can the Von Neumann bottleneck be weakened?

Caches reduce the bottleneck impact.

4. What is the central statement of Amdahl's law?

The performance gain, when executing a programs in parallel on multiple CPUs is limited mainly by the sequential part of the problem.

5. Which important factor is ignored by Amdahl's law?

Amdahl's law does not take into account the cache and the effects, which are caused by the cache in practice. A growing number of CPUs also increases the quantity of fast memory which is available.

6. What is the central statement of Gustafson's law (highlight the difference against Amdahl's law)?

The bigger a parallelizable problem is, the smaller is the portion of the sequential part. A problem, which is sufficiently large, can be parallelized efficiently. The sequential part is not limiting, because it gets more and more unimportant as the number of CPUs rises.

Exercise 5 (Parallel Computers)

1. Describe the shared memory architecture in just a few words.

In systems with shared memory, the entire memory is part of a uniform address space, which is accessed by all CPUs.

2. Name two challenges of shared memory architectures.

Write operations of the CPUs must be coordinated.

Data inside the CPU caches. If a memory cell duplicated in multiple CPU caches, any change in the memory cell must be propagated to all caches.

3. What is the difference between asymmetric and symmetric multiprocessing (SMP)?

Symmetric multiprocessing (SMP) allows to dynamically distribute the running processes to all available CPUs.

In asymmetric multiprocessing systems, each CPU must be assigned to a fixed task. One or more CPUs run the operating system. The other processes are distributed to the remaining CPUs.

4. Give an example for a system in practice which implements the asymmetric multiprocessing architecture.

the IBM Cell processor.

5. Give an example for a system which implements the symmetric multiprocessing (SMP) architecture.

Dual-Mainboards which are used often in Workstations or Server systems.

6. Describe the distributed memory architecture in just a few words.

In systems with distributed memory, each CPU can only access its own local memory.

7. Name a drawback of distributed memory architectures.

Network connections are much slower, compared with the data rate between CPU and memory.