

Last name:

First name:

Student number:

Question 1)

Points:

Maximum points: 8

Imagine you have trained a pigeon to carry a USB flash memory drive with a storage capacity of 64 GB. The pigeon can fly with an average speed of 72 km/h.

For what range of distance has the pigeon a better data rate than a computer network whose data rate (excluding overhead) is 250 Mbps?

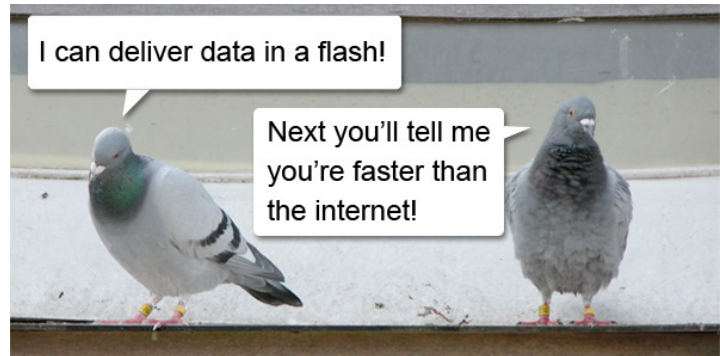


Image source: <http://www.usb-flashdrive.co.uk>

A speed of 72 km/h is equal to $\frac{72 \text{ km}}{3600 \text{ s}} = 0.02 \text{ km/s} \implies 20 \text{ m/s}$

To calculate the data rate, we must calculate the time, the pigeon needs to transport the 512 Gbit for an unknown distance.

The time to travel the unknown distance of x km is: $\frac{x \text{ km}}{0.02 \text{ km/s}} = \frac{x \text{ s}}{0.02} = 50x \text{ s}$

Data rate of the pigeon: $\frac{512 \text{ Gbit}}{50 x \text{ s}} \implies \frac{10.24 \text{ Gbit}}{x \text{ s}} \implies \frac{10,240 \text{ Mbit}}{x \text{ s}}$

If the data rate of the pigeon shall be faster than the computer network, the following inequality must be true:

$$\frac{10,240 \text{ Mbit/s}}{x} > 250 \text{ Mbit/s.}$$

Via two equivalence transformations, the inequality can be solved for x .

$$\frac{10,240 \text{ Mbit/s}}{x} > 250 \text{ Mbit/s} \quad | \quad \cdot x$$

$$10,240 \text{ Mbit/s} > 250 \text{ Mbit/s} \cdot x \quad | \quad : 250 \text{ Mbit/s}$$

For $\frac{10,240 \text{ Mbit/s}}{250 \text{ Mbit/s}} > x \implies x < \frac{10,240 \text{ Mbit/s}}{250 \text{ Mbit/s}} = 40.96 \text{ (km)}$, the pigeon has a faster data rate than the transmission line.

Last name:

First name:

Student number:

Question 2)

Points:

Maximum points: 1+1+1+1+1+1=6

- a) What is the central statement of Moore's law?

The number of transistors on an integrated circuit double every 24 months.

- b) What is the Von Neumann bottleneck?

The data and control bus is increasingly becoming a bottleneck between the CPU and memory.

- c) How can the Von Neumann bottleneck be weakened?

Caches reduce the bottleneck impact.

- d) What is the central statement of Amdahl's law?

A program can never be fully executed in parallel. The performance gain is limited mainly by the sequential part of the problem.

- e) Which important factor is ignored by Amdahl's law?

A growing number of CPUs also increases the quantity of fast memory (cache) which is available.

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

A problem, which is sufficiently large, can be parallelized efficiently. The problem needs to scale with the number of CPUs. If the parallel portion of the problem grows with the number of CPUs, the sequential part is not limiting, because it gets more and more unimportant as the number of CPUs rises.

Last name:

First name:

Student number:

Question 3)

Points:

Maximum points: 1+1+1+1+1=5

- a) Describe the shared memory architecture in just a few words.

The entire memory is part of a uniform address space, which is accessed by all CPUs. The memory is accessed via an interconnect.

- b) Name two challenges of shared memory architectures.

Write operations of the CPUs must be coordinated.

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

- c) What is the difference between asymmetric and symmetric multiprocessing (SMP)?

In multiprocessor systems, which operates according to the asymmetric multiprocessing principle, 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.

SMP allows to dynamically distribute the running processes to all available CPUs. All CPUs can access the memory with the same speed.

- d) Describe the distributed memory architecture in just a few words.

Each CPU can only access its own local memory. The communication between the CPUs takes place via a network connection.

- e) Name a drawback of distributed memory architectures.

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

Last name:

First name:

Student number:

Question 4)

Points:

Maximum points: $1+2+2+1+1=7$

- a) For exercise sheet 4, you implemented a graphical remote desktop solution for a Linux instance. Which protocol did you use to implement the remote desktop solution with Linux?

Remote Framebuffer Protocol (RFB), Remote Desktop Protocol (RDP), X11,...

- b) If you create a cluster of virtual server instances in EC2, you can distribute the instances over multiple regions. Give an advantage and a drawback of this method.

Advantage: Reliability.

Drawback: Performance of network communication, cost for network communication between regions.

- c) If you create a cluster of virtual server instances in EC2, you can distribute the instances over multiple availability zones. Give an advantage and a drawback of this method.

Advantage: Reliability.

Drawback: Performance of network communication.

- d) For exercise sheet 5, you implemented with the infrastructure services of the Amazon Web Services a highly available High Throughput Cluster of virtual web servers. Which web server software did you use?

Apache, nginx,...

- e) For exercise sheet 5, you implemented with the infrastructure services of the Amazon Web Services a highly available High Throughput Cluster of virtual web servers. The web server data was stored in EBS volumes. Which Linux file system did you deploy on the EBS volumes?

ext3, ext4, ReiserFS, XFS, btrfs,...

Last name:

First name:

Student number:

Question 5)

Points:

Maximum points: 10

For exercise sheet 5, you implemented with the infrastructure services of the Amazon Web Services a highly available High Throughput Cluster of virtual web servers. One part of this task was to attach an EBS volumes as persistent storage to each web server instance. All storage volumes contained identical data.

But not for all possible scenarios, it is the optimal approach, when all instances have EBS volumes with identical content attached.

For exercise sheet 6, you investigated other approaches and summarized their advantages and drawbacks. The focus of your investigation were the aspects availability, cost, throughput, scalability and complexity of the different approaches.

Explain one of the approaches you investigated and summarize its advantages and drawbacks. Focus on the aspects availability, cost, throughput, scalability and complexity. Name the services and software solutions you used. Maybe a diagram is helpful too.

One option is using a protocol like the Network File System (NFS). One node hosts the NFS server and exports the file system and each other node execute a NFS client, which accesses the file system of the NFS server.

Advantage: Only the NFS server requires an EBS volume.

Drawbacks: The NFS server is a single point of failure. If multiple nodes import the file system via NFS, the NFS server may become a bottleneck.

Another option is using a distributed file system like GlusterFS, Ceph or PVFS2. Some of these distributed file systems also provide internal replication.

Advantages: Reliability. A single node will not become a bottleneck.

Drawback: All nodes require an EBS volume.

Last name:

First name:

Student number:

Question 5 – Additional Page)

Maximum points: 10

Last name:

First name:

Student number:

Question 6)

Points:

Maximum points: $3+7+2=12$

Your local time in Frankfurt am Main is Monday 09:00 (UTC+1). You need to copy 3 TB of data into the storage service S3. You have two options:

- **Scenario 1:** You immediately start at 09:00 (UTC+1) to upload the 3 TB of data to S3 via the internet. Consider the data rate between your computer and S3 is 100 Mbit/s.

- **Scenario 2:** You use the AWS Import/Export service. Therefore you copy the data to a HDD, which is connected via USB 3.0. The transfer rate (for write) is 125 MB/s.

After you copied the data, you pack the HDD into a parcel and send it via a package delivery company to Amazon. DHL, UPS and FedEx can deliver a parcel from Frankfurt am Main in less than 24 hours to most places in Europe.

You need 15 Minutes to put the HDD into a parcel and another 15 Minutes to bring the parcel to the branch office of a package delivery company.

The parcel must arrive at the branch office of the package delivery company no later than 16:30 (UTC+1) to arrive at Amazon in Ireland at 9:00 (UTC) the next working day.

An Amazon employee needs to copy the data from the HDD to the S3 service. The transfer rate of the HDD (for read) is 150 MB/s.

Consider 3 hours additional overhead for the in-house mail at Amazon to ship the HDD to the correct employee.

Calculate...

- a) for the first scenario, how long it takes until the data is copied to S3.
- b) for the second scenario, how long it takes until the data is copied to S3.
- c) the data rate of the second scenario.

(For all subtasks, the calculation steps must be visible.)

Last name:

First name:

Student number:

Question 6 – Additional Page)

Maximum points: 3+7+2=12

Scenario 1:

$$\frac{100 \text{ Mbps}}{8} = 12.5 \text{ MB/s} = 12.5 * 10^6 \text{ Byte/s}$$

$$\frac{3 * 10^{12} \text{ Byte}}{12.5 * 10^6 \text{ Byte/s}} = 0.24 * 10^6 \text{ s} = 240,000 \text{ s} \implies \frac{240,000 \text{ s}}{60} = 4,000 \text{ m} \implies \frac{4,000 \text{ m}}{60} = 66.\bar{6} \text{ h}$$

\implies the data transmission requires 2 days, 18 hours, 40 minutes

Scenario 2:

Day 1, 9:00 (UTC+1)

$$\text{Write data: } \frac{3,000,000 \text{ MB}}{125 \text{ MB/s}} = 24,000 \text{ s} = 400 \text{ m} = 6 \text{ h}, 40 \text{ m}$$

Day 1, 15:40 (UTC+1)

30 minutes to put the HDD into a parcel and to bring it to the package delivery company.

Day 1, 16:10 (UTC+1)

Because the packet arrived at the package delivery company before 16:30 (UTC+1), it is delivered to Amazon the next working day at 9:00 (UTC).

Day 2, 9:00 (UTC) = 10:00 (UTC+1)

3 hours are required to ship the HDD via in-house mail at Amazon to the correct employee.

Day 2, 12:00 (UTC) = 13:00 (UTC+1)

$$\text{Read data: } \frac{3,000,000 \text{ MB}}{150 \text{ MB/s}} = 20,000 \text{ s} = 333.33\bar{3} \text{ m} = 5 \text{ h}, 34 \text{ m}$$

Day 2, 17:34 (UTC) = 18:34 (UTC+1)

\implies the data transmission requires 1 day, 9 hours, 34 minutes = 120,840 seconds

$$\frac{3 * 10^{12} \text{ Byte}}{120,840 \text{ s}} = \frac{3,000,000 * 10^6 \text{ Byte}}{120,840 \text{ s}} = 24.8262164846 * 10^6 \text{ Byte/s}$$

$$24.8262164846 * 10^6 \text{ Byte/s} * 8 = 198,609,731.877 * 10^6 \text{ Bit/s} = \textit{approx.} 199 \text{ Mbps}$$

The data rate of the second scenario is approx. two times better compared with the first scenario!

Last name:

First name:

Student number:

Question 7)

Points:

Maximum points: 1+1+1+1+1+1=6

- a) Give a short definition of Cluster Computing.

Clustering is parallel computing on systems with distributed memory.

- b) What is a Cluster of Workstations?

The nodes are only available at specific times. During normal working times, the employees use the nodes of such a cluster system as workstations.

- c) How can the availability of a system be calculated?

mean uptime = Mean Time Between Failures

mean downtime = Mean Time To Repair

$$\text{availability} = \frac{\text{mean uptime}}{\text{mean uptime} + \text{mean downtime}}$$

- d) By which approach does High Availability Clustering achieve its objective?

With redundancy of nodes and their components and by avoiding a single point of failure.

- e) Explain the difference between Shared Nothing and Shared Disk Clusters.

In a Shared Nothing cluster, each node has its own storage resource. Even, when a resource is physically connected to multiple nodes, only a single node is allowed to access it.

In a Shared Disk cluster, all nodes have access to a shared storage.

- f) Explain the main difference between a SAN (Storage Area Network) and a NAS (Network Attached Storage).

A SAN provides block-level access to storage devices via the network.

A NAS provides file system-level access to storage devices via the network.

Last name:

First name:

Student number:

Question 8)

Points:

Maximum points: 10

For exercise sheet 10, you implemented a MPI program, which calculates π via Monte Carlo simulation. Several options exist to solve this task. Two of them are:

- The master node creates a specified amount of random numbers and distributes them to the nodes. In this case, additional nodes just reduce the required computing time.
- Each node creates a specified amount of random numbers. In this case, additional nodes will not reduce required computing time but the result will be closer to π .

Implement the solution you created for exercise sheet 10 in pseudo-code. Use comments to explain what operations are carried out for what reason.

These MPI functions may be useful for your implementation.

- `MPI_Bcast()`;
- `MPI_Comm_rank()`;
- `MPI_Comm_size()`;
- `MPI_Finalize()`;
- `MPI_Get_processor_name()`;
- `MPI_Init()`;
- `MPI_Reduce()`;

```
1 main-Function
2 {
3
4 MPI_Init();           // Start/Create the MPI environment
5 MPI_Comm_size();     // Determines the number of processes
6 MPI_Comm_rank();    // Determines the rank (ID) of the calling process
7
8 calculate_pi         // Each process calculates its own estimate of pi
9
10 MPI_Reduce();       // Send all values to the master and sum them
11
12 If id == master {
13     calculate_avarage_value // The master calculates the avarage value
14     print_result           // Print out the result
15 }
16
17 MPI_Finalize();     // Stop the MPI environment
18 }
```

Last name:

First name:

Student number:

Question 9)


Points:

Maximum points: 9

- PR_p = PageRank of a web page p
- $L_{IN}(p)$ = Set of documents, which refer to $p \implies$ incoming links
- $L_{OUT}(p)$ = Set of documents, to which p refers \implies outgoing links
- d = damping factor between 0 and 1

$$PR(p) = (1 - d) + d * \sum_{p_i \in L_{IN}(p)} \frac{PR(p_i)}{\text{amount } L_{OUT}(p_i)}$$

Calculate the missing iterations of the PageRank algorithm for the given example scenario with $d = 0.75$.



- $PR(A) = (1 - d) + d * PR(B)$
- $PR(B) = (1 - d) + d * (\frac{PR(A)}{2} + PR(C))$
- $PR(C) = (1 - d) + d * \frac{PR(A)}{2}$
- Conversion to iteration equations with $d = 0.75$:
- $PR_{n+1}(A) = 0.25 + 0.75 * PR_n(B)$
- $PR_{n+1}(B) = 0.25 + 0.75 * (\frac{PR_n(A)}{2} + PR_n(C))$
- $PR_{n+1}(C) = 0.25 + 0.75 * \frac{PR_n(A)}{2}$

	0	1	2	3	4	5	PR
A	1	1	1,28125	1,0703125	1,1494140625	1,1494140625	1,127166748
B	1	1,375	1,09375	1,19921875	1,19921875	1,1695556641	1,1918029785
C	1	0,625	0,625	0,73046875	0,6513671875	0,6810302734	0,6810302734

Last name:

First name:

Student number:

Question 10)

Points:

Maximum points: $1+1+1+1+1=5$

- a) What is WSDL and for what purpose is it used?

WSDL (Web Services Description Language) is an XML-based interface description language to specify...

- *how a web service can be called*
- *what parameters a web service expects*
- *what data structures a web service returns*

- b) What is UDDI and for what purpose is it used?

UDDI (Universal Description, Discovery and Integration) provides a standardized directory structure for the metadata (e.g. technical characteristics, requirements and provider information) of web services.

- c) Explain the difference between UDDI and WS-Inspection.

UDDI = few, centralized directories, where different providers publish their services.

WS-Inspection = many decentralized, small directories, in which few providers publish their services.

- d) Describe the difference between the theoretical implementation of SOAP web services and the way, SOAP web services operate in practice.

Publicly accessible web services are usually offered without using UDDI. The reason is that individual web services and the required access information are already known to the user or client application. For this reason, the users and client applications no longer need to search for a web services in a registry service.

- e) For exercise sheet 12, you implemented a Private Cloud storage service, which implements the S3 API. Which one of the existing solutions did you use?

OpenStack Swift, Eucalyptus Walrus, Nimbus Cumulus, S3 ninja,...

Last name:

First name:

Student number:

Question 11)

Points:

Maximum points: $1+1+1+1+1+1+1+1=8$

- a) What is the difference between emulation and virtualization?

Virtualization allows to split the resources of a computer system and to execute multiple independent operating system instances.

Emulation allow the execution of an unmodified operating system, which is designed for a different hardware architecture (CPU).

- b) How works application virtualization?

Applications are locally executed inside a virtual environment, which uses local resources and provides all the components, which are required by the application. The VM is located between the executed application and the operating system.

- c) What is the function of the Virtual Machine Monitor (VMM)?

The VMM distributes hardware resources to VMs.

- d) Where runs the Virtual Machine Monitor (VMM)?

The VMM runs *hosted* as an application in the host operating system.

The VMM runs *bare metal* and replaces the host operating system.

- e) Can all physical hardware resources be virtualized when full virtualization is used? If this is not possible, give an example where it does not work.

Some hardware components are emulated, because they are not designed for the concurrent access from multiple operating systems. Example: Network adapters.

- f) Where runs the hypervisor when paravirtualization is used?

The hypervisor runs *hosted* as an application in the host operating system.

The hypervisor runs *bare metal* and replaces the host operating system.

- g) Why is for paravirtualization a host operating system required?

A host operating system is required because of the device drivers.

- h) Name a drawback of operating system-level virtualization (containers/jails).

Only independent instances of the same operating system are started. It is impossible to start different operating systems at the same time because all virtual environments use the same kernel.

Last name:

First name:

Student number:

Question 12)

Points:

Maximum points: 4

How long does it take to transfer 7.5 TB via a 1 Gbps (= 1,000 Mbps) Ethernet?

$$\frac{1,000 \text{ Mbps}}{8} = 125 \text{ MB/s} = 125 * 10^6 \text{ Byte/s}$$

$$\frac{7.5 * 10^{12} \text{ Byte}}{125 * 10^6 \text{ Byte/s}} = 0.06 * 10^6 \text{ s} = 60,000 \text{ s}$$

$$\frac{60,000 \text{ s}}{60} = 1,000 \text{ m}$$

$$\frac{1,000 \text{ m}}{60} = 16.\bar{6} \text{ h}$$

Better (more correct) solution:

$$7.5 \text{ TB} * 1024 * 1024 = 7,864,320 \text{ MB}$$

$$\frac{7,864,320 \text{ MB}}{125 \text{ MB/s}} = 62,914.56 \text{ s}$$

$$\frac{62,914.56 \text{ s}}{60} = 1,048.576 \text{ m}$$

$$\frac{1,048.576 \text{ m}}{60} = \text{approx. } 17.476 \text{ h}$$