Sample solution of the written examination in

Cloud Computing

February 19th 2019

Last name: _____

First name: _____

Student number: _____

I confirm with my signature that I will process the written examination alone and that I feel healthy and capable to participate this examination.
I am aware, that from the moment, when I receive the written examination, I am a participant of this examination and I will be graded.

Signature: _____

- Use the provided sheets. Own paper must *not* be used.
- You are allowed to use a *self prepared*, *single sided DIN-A4 sheet* in the exam. Only *handwritten originals* are allowed, but no copies.
- You are allowed to use a non-programmable calculator.
- Do *not* use a red pen.
- Time limit: *90 minutes*
- Turn off your mobile phones!

Result:

Question:	1	2	3	4	5	6	7	8	**Σ**	Grade
Maximum points:	22	5	10	12	12	9	10	10	90	—
Achieved points:										

1.0: 90.0-85.5, **1.3**: 85.0-81.0, **1.7**: 80.5-76.5, **2.0**: 76.0-72.0, **2.3**: 71.5-67.5,
2.7: 67.0-63.0, **3.0**: 62.5-58.5, **3.3**: 58.0-54.0, **3.7**: 53.5-49.5, **4.0**: 49.0-45.0, **5.0**: <45

Last name: First name: Student number:

Question 1)

Points: .....

Maximum points: 3+12+7=22

Your local time in Frankfurt am Main is Monday 10:00 (UTC+1). You need to copy 4 TB of data into a storage service which is located in Ireland (UTC). You have two options:

- **Scenario 1:** You immediately start at 10:00 (UTC+1) to upload the 4 TB of data to the storage service via the Internet. Consider the data rate between your computer and the storage service is 200 Mbit/s.
- **Scenario 2:** You use an Import/Export service that is offered by many storage service providers today. Therefore you copy the data to a SDD. The transfer rate (for write) is 200 MB/s.

  After you copied the data, you pack the SDD 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 SDD into a parcel and another 30 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 the cloud service provider in Ireland at 8:00 (UTC) the next working day.

  An employee of the cloud service provider needs to copy the data from the SDD into the cloud storage service. The transfer rate of the SDD (for read) is 320 MB/s.

  Consider two hours additional overhead for the in-house mail at the cloud service provider to ship the SDD to the correct employee.

a) Calculate for the first scenario, how long it takes until the data is copied to the storage service:

   Days: _____ 1 _____ **Hours**: _____ 20 _____ **Minutes**: _____ 24 _____

b) Calculate for the second scenario, how long it takes until the data is copied to the storage service:

   Days: _____ 1 _____ **Hours**: _____ 4 _____ **Minutes**: _____ 29 _____

c) Calculate for the second scenario the **Data Rate [Mbps]**: _____ 312 _____

(Fill out the empty fields. The calculation steps of all subtasks must be visible.)

Last name: First name: Student number:

Question 1)

Points: .....

Maximum points: 3+12+7=22

(**This page is free space for your calculation steps of Question 1. Do not forget to write the results into the correct fields.**)

Scenario 1:

$$\frac{200\,\text{Mbps}}{8} = 25\,\text{MB/s} = 25 * 10^6\,\text{Byte/s}$$

$$\implies \frac{4 * 10^{12}\,\text{Byte}}{25 * 10^6\,\text{Byte/s}} = 0,16 * 10^6\,\text{s} = 160,000\,\text{s} \implies \frac{160,000\,\text{s}}{60\,\text{s/m}} = 2,666.\overline{6}\,\text{m} \implies \frac{2,666.\overline{6}\,\text{m}}{60\,\text{m/h}} = 44.\overline{4}\,\text{h}$$

$\implies$ *the data transmission requires 1 day, 20 hours, 24 minutes*

Scenario 2:

Day 1, 10:00 (UTC+1)

Write data: $\frac{4{,}000{,}000\ \text{MB}}{200\ \text{MB/s}} = 20,000\ \text{s} = 333.\overline{3}\ \text{m} \approx\ 5\ \text{h}, 34\ \text{m}$

Day 1, 15:34 (UTC+1)

45 minutes to put the SDD into a parcel and to bring it to the package delivery company.

Day 1, 16:19 (UTC+1)

Because the packet arrived at the package delivery company before 16:30 (UTC+1), it is delivered to the cloud service provider the next working day at 8:00 (UTC).

Day 2, 8:00 (UTC) = 9:00 (UTC+1)

2 hours are required to ship the SDD via in-house mail at the cloud service provider to the correct employee.

Day 2, 10:00 (UTC) = 11:00 (UTC+1)

Read data: $\frac{4{,}000{,}000\ \text{MB}}{320\ \text{MB/s}} = 12,500\ \text{s} = 208.\overline{3}\ \text{m} \approx\ 3\ \text{h}, 29\ \text{m}$

Day 2, 13:29 (UTC) = 14:29 (UTC+1)

$\implies$ *the data transmission requires approx. 1 day, 4 hours, 29 minutes ≈ 102,540 seconds*

$$\frac{4 * 10^{12}\,\text{Byte}}{102,540\,\text{s}} = \frac{4,000,000 * 10^6\,\text{Byte}}{102,540\,\text{s}} \approx\ 39 * 10^6\,\text{Byte/s}$$

$$\implies\ 39 * 10^6\,\text{Byte/s} * 8 \approx\ 312 * 10^6\,\text{Bit/s} \approx\ 312\,\text{Mbps}$$

Last name: First name: Student number:

Question 2)

Points: .....

Maximum points: 5

Explain one possible way of how the multiplication of two matrices can be done in parallel by using a cluster system. (*In other words: Which parts of the multiplication process can be carried out in parallel by the nodes of a cluster and how is it done and what is the task of the master?*)

See the solution of the MPI Special Challenge.

Last name: First name: Student number:

Question 3)

Points: .....

Maximum points: 10

Provide a pseudocode solution that implements a parallel matrix multiplication. Focus on the MPI-Functions that are required for the implementation of a parallel matrix multiplication.

See the solution of the MPI Special Challenge.

Last name: First name: Student number:

Question 4)

Points: .....

Maximum points: 2+10=12

This two diagrams show the total execution time of a Mergesort application from SS2018 for two different problem sizes = number of integer values to be sorted.

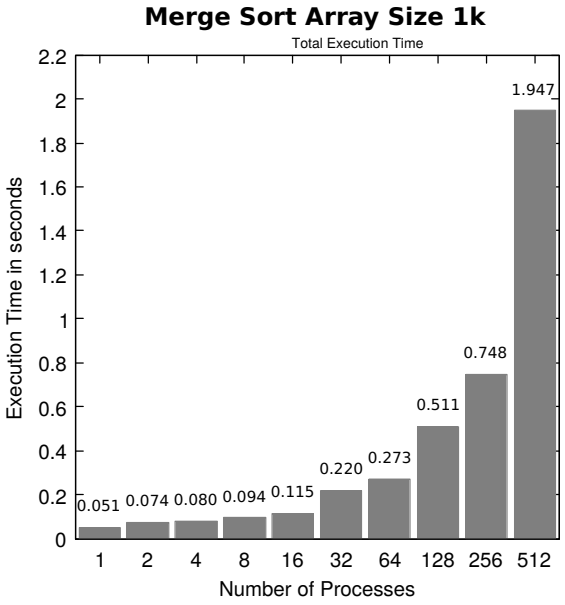

Figure 1: Problem Size = 1,000 values

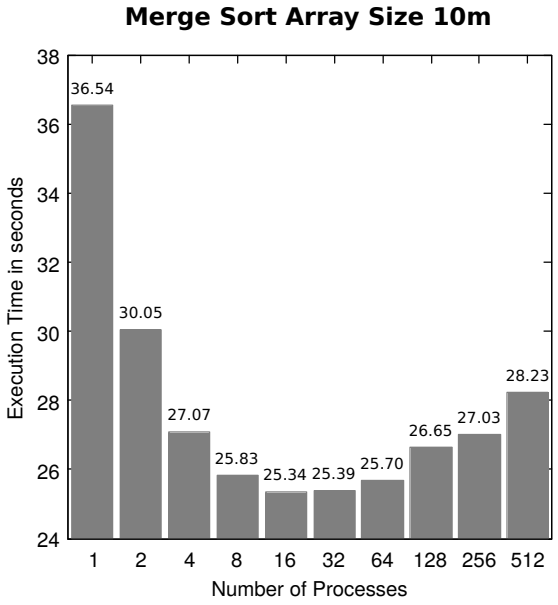

Figure 2: Problem Size = 10,000,000 values

The two diagrams demonstrate two fundamental laws and limitations of parallel computing.

a) Name the two fundamental laws and limitations of parallel computing which are relevant here.

Amdahl's law and Gustafson's law.

b) Explain the two fundamental laws and limitations of parallel computing by using the two diagrams.

Last name: First name: Student number:

Question 5)

Points: .....

Maximum points: 4+8=12

Company X runs 500 computer workplaces.

- Scenario 1: Fat clients (PC)
  - Electrical power rating per desktop: 450 watts
  - Electrical power rating per screen: 80 watts
- Scenario 2: Thin clients
  - Electrical power rating per thin client: 30 watts
  - Electrical power rating per screen: 80 watts
  - Electrical power rating per server blade: 600 watts
  - Each server blade has enough resources to interact with 30 thin clients

What are the electricity costs per year for 24/7 operation when the electricity price is 0,32 €/kWh?

Scenario 1:
Electricity costs per year (including the leap year) for 500 computer workplaces:

$$0.53\,\text{kW} * 24\,\frac{\text{h}}{\text{Day}} * 365.25\,\frac{\text{Day}}{\text{Year}} * 0.32\,\frac{€}{\text{kWh}} * 500 = \mathbf{743{,}356.8}\,\frac{€}{\text{Year}}$$

Scenario 2:
Electricity costs per year (including the leap year) for 500 computer workplaces:

$$0.11\,\text{kW} * 24\,\frac{\text{h}}{\text{Day}} * 365.25\,\frac{\text{Day}}{\text{Year}} * 0.32\,\frac{€}{\text{kWh}} * 500 = 154,281.6\,\frac{€}{\text{Year}}$$

17 server blades are required to run the 500 computer workplaces.

Electricity costs per year (including the leap year) for 17 server blades.

$$0.6\,\text{kW} * 24\,\frac{\text{h}}{\text{Day}} * 365.25\,\frac{\text{Day}}{\text{Year}} * 0.32\,\frac{€}{\text{kWh}} * 17 \approx\ 28,612.23\,\frac{€}{\text{Year}}$$

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

$$154,281.6\,\frac{€}{\text{Year}} + 28,612.23\,\frac{€}{\text{Year}} \approx\ \mathbf{182{,}893.83}\,\frac{€}{\text{Year}}$$

Last name: First name: Student number:

Question 6)

Points: .....

Maximum points: 4+1+2+2=9

a) Explain what the Twelve-Factor-App is.

It is a a Methodology for building SaaS Apps that...

- *Use declarative formats for setup automation*
- *Have a clean contract with the underlying operating system*
- *Are suitable for deployment on modern cloud platforms*
- *Minimize divergence between development and production*
- *Can scale up without significant changes to tooling, architecture or development practices*

It is a triangulation on ideal practices for app development, paying particular attention to the dynamics of the organic growth of an app over time, the dynamics of collaboration between developers working on the app's codebase, and avoiding the cost of software erosion.

b) To which sort of applications and programming languages is the Twelve-Factor-App compatible?

The Twelve-Factor-App can be applied to Apps, written in any programming language, and which use any combination of backing services.

c) Explain the purpose of the MPI function `MPI_Get_processor_name`.

It determines the name of the processor. The name identifies the hardware, where MPI runs. The exact output format is implementation-dependent and may by equal with the output of `gethostname`.

d) Explain what the communicator of an MPI application is.

A communicator contains a group of processes and a communication context.

Last name: First name: Student number:

Question 7)

Points: .....

Maximum points: 10

a) Name <u>five</u> of the main components (services) of OpenStack and explain what the purpose of each mentioned component (service) is.

See the solution of the OpenStack Special Challenge.

Last name: First name: Student number:

Question 8)

Points: .....

Maximum points: 2+2+2+2+2=10

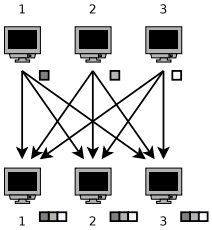

a) Name an MPI function that implements this sort of communication.

`MPI_Allgather`

b) Name an MPI function that implements this sort of communication.

`MPI_Send` and `MPI_Recv`

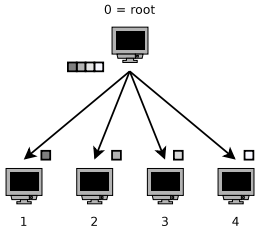

c) Name an MPI function that implements this sort of communication.

`MPI_Scatter`

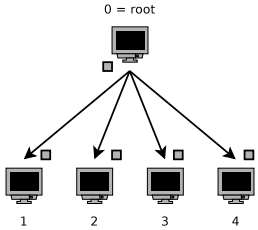

d) Name an MPI function that implements this sort of communication.

`MPI_Bcast`

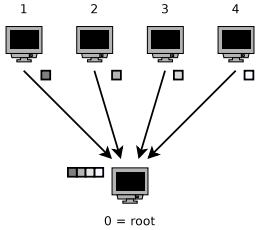

e) Name an MPI function that implements this sort of communication.

`MPI_Gather`