

# Sample solution of the written examination

## Operating Systems

*February 25th 2020*

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: \_\_\_\_\_

- You are allowed to use a *self prepared, single sided DIN-A4 sheet* in the exam. Only *handwritten originals* are allowed, but no copies.
- Use the provided sheets.
- 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 | 9 | 10 | 11 | $\Sigma$ | Grade |
|------------------|---|---|---|----|---|---|----|---|---|----|----|----------|-------|
| Maximum points:  | 4 | 8 | 7 | 11 | 9 | 8 | 10 | 9 | 8 | 7  | 9  | 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:  $0.5+0.5+0.5+0.5+0.5+0.5+1=4$

- a) Explain what scheduling is.

*Automatic creation of an execution plan (schedule), which is used to allocate time limited resources to users or their processes.*

- b) Explain what swapping is.

*Process of storing and removing data to/from main memory from/into background memory (HDDs/SSDs).*

- c) Explain how memory protection works.

*The memory is split and running programs are separated from each other.*

- d) Explain why memory protection is useful.

*This way, a bug or crash of a single program does not affect the stability of other programs and the total system.*

- e) Name the scheduling concept that modern operating systems implement.

*Multilevel feedback scheduling.*

- f) Operate modern operating systems in singletasking oder multitasking mode?

*Multitasking.*

- g) Explain the difference between preemptive and non-preemptive scheduling.

*When using non-preemptive scheduling, a process, which gets the CPU assigned by the scheduler, remains control over the CPU until its execution is finished or it gives the control back on a voluntary basis.*

*When using preemptive scheduling, the CPU may be removed from a process before its execution is completed.*

Last name:

First name:

Student number:

---

## Question 2)

Points: .....

Maximum points: 1+1+1+1+2+2=8

- a) Explain the difference between 8-bit, 16-bit, 32-bit and 64-bit operating systems.

*The bit number indicates the memory address length, with which the operating system works internally.*

- b) Name the two essential criteria of real-time operating systems.

*Response time (short latency) and deadlines must be met.*

- c) Name the two types of real-time operating systems.

*Hard real-time operating systems and soft real-time operating systems.*

- d) Name and explain two advantages of hybrid kernels.

- *Better performance as with microkernels (because fewer context switching)*
- *The stability is (theoretically) better as with monolithic kernels*

- e) Your colleague recommends you to relocate frequently used server daemons, such as web server, email server, SSH server and FTP server, from the user mode to the kernel mode. What is your opinion about this idea? Give reasons for your answer. Explain at least one benefit and one drawback on this idea.

*One advantage would be that the operating system and the server daemons overall perform better, because in the scenario described, fewer context switches between user mode and kernel mode are required.*

*More serious, however, is the resulting drawback. There is a security risk. Complex software such as server daemons should not run in kernel mode. Bugs in the server daemons may cause system crashes or allow attackers to gain complete control of the system.*

- f) Explain why the Linux kernel does not belong to the microkernel category.

*The Linux kernel modules run in kernel mode and not in user mode. Therefore, the Linux kernel is a monolithic kernel.*

Last name:

First name:

Student number:

---

## Question 3)

Points: .....

Maximum points:  $1+1+1.5+0.5+1+0.5+0.5+0.5+0.5=7$

- a) Name two rotating magnetic digital data storages.

*Hard Disk Drive, Drum memory, Floppy Disk.*

- b) Name two non-rotating magnetic digital data storages.

*Magnetic-core memory, Magnetic tape, Magnetic stripe card, Compact cassette (Datasette), Bubble memory.*

- c) Name three benefits of data storage without moving parts compared to data storage with moving parts.

*Lower power consumption, lesser waste heat, mechanical robustness, no noise generation.*

- d) Name one non-persistent data storage.

*Main memory (DRAM).*

- e) Explain what random access is.

*The media does not need to be searched sequentially from the start – such as with magnetic tapes – to locate a specific location (file).*

- f) Is the GNU Bash a compiler or an interpreter?

*It is an interpreter.*

- g) Is the GCC a compiler or an interpreter?

*It is a compiler.*

- h) Is the Javac a compiler or an interpreter?

*It is a compiler.*

- i) Is Python a compiler or an interpreter?

*It is an interpreter.*

Last name:

First name:

Student number:

---

## Question 4)

Points: .....

Maximum points:  $1.5+1.5+3+1+1+2+1=11$

- a) Name the three components that a CPU contains.

*Arithmetic logic unit, control unit, memory.*

- b) Name the three digital bus systems that each computer system includes according to the Von Neumann architecture.

*Control bus, address bus, data bus.*

- c) Explain the tasks that are carried out by the three digital bus systems of subtask b).

*Control bus: Transmits commands (e.g. read and write requests) from the CPU and returns status signals from the I/O devices*

*Address bus: Transmits memory addresses.*

*Data bus: Transmits data between CPU, main memory and I/O devices.*

- d) Explain what the Front Side Bus is.

*It is the bus between CPU and chipset. It contains the address bus, data bus and control bus.*

- e) The chipset contains two components. Name them.

*Northbridge und Southbridge.*

- f) Name the tasks of the components of subtask e).

*The Northbridge is used for the connection of main memory and graphics card(s) with the CPU. It is located close to the CPU for rapid data transfer. The Southbridge is used for „slow“ connections like Ethernet, SATA and USB.*

- g) Explain why it is impossible to implement the optimal replacement strategy OPT.

*Because it is not possible to predict the future and therefore the future request sequence is unknown.*

Last name:

First name:

Student number:

---

## Question 5)

Points: .....

Maximum points:  $1+0.5+0.5+0.5+0.5+1+1+1+3=9$

- a) Name the two basic cache write policies.

*Write-through and write-back.*

- b) Name the cache write policy of subtask a) in which inconsistencies may occur.

*Write-back.*

- c) Name the cache write policy of subtask a) which causes a lower system performance.

*Write-through.*

- d) Name the cache write policy of subtask a) which uses so called dirty bits.

*Write-back.*

- e) Explain why dirty bits are used.

*For each page inside the cache, a dirty bit specifies whether the page was modified.*

- f) Name the factors that influence the access time of HDDs.

*Average seek time, average rotational latency time*

- g) Describe the factors of subtask f).

- *Average seek time: The time it takes for the arm to reach a desired track.*
- *Average rotational latency time: Delay of the rotational speed, until the required disk sector is located under the head. Depends entirely on the rotational speed of the disks.*

- h) Explain what causes a page fault exception to occur.

*A process tries to access a page, which is not located in the physical main memory.*

- i) Explain the reaction of the operating system, when a page fault exception occurs.

*The operating system handles the page fault exception by executing these steps:*

- *Allocate the page by using the controller and the device driver on the swap memory (SSD/HDD).*
- *Copy the page into a free page of the main memory.*
- *Update the page table.*
- *Return control to the process. The process next tries to execute again the instruction that caused the page fault.*

Last name:

First name:

Student number:

---

## Question 6)

Points: .....

Maximum points:  $0.5+0.5+0.5+0.5+0.5+1+1+2+1.5=8$

- a) Name one RAID level, that improves the data transfer rate for write.

*RAID-0 or RAID-5.*

- b) Name one RAID level, that improves the reliability.

*RAID-1 or RAID-5.*

- c) How many drives are allowed to fail in a RAID 0 array without data loss?

*None.*

- d) How many drives are allowed to fail in a RAID 1 array without data loss?

*At least a single drive must work properly.*

- e) How many drives are allowed to fail in a RAID 5 array without data loss?

*A single drive maximum is allowed to fail.*

- f) Name one advantage and one drawback of software RAID compared to hardware RAID.

*Benefit: No cost for additional hardware.*

*Drawback(s): Operating system dependent, additional CPU load.*

- g) Explain the key message of Laszlo Belady's anomaly.

*FIFO result in worse results for certain access sequences with a bigger memory.*

- h) Some scheduling algorithms were not discussed in detail during class because they cannot be used under realistic conditions. Explain why they cannot be used under realistic conditions.

*These algorithms require to know how long it takes for each process until its termination. In other words, these algorithms need to know, how long is the execution time of each process. In practice this is almost never the case  $\implies$  unrealistic.*

- i) Name three unrealistic scheduling algorithms from subtask h).

*SJF, SRTF, LJF, LRFB, HRRB*

Last name:

First name:

Student number:

---

## Question 7)

Points: .....

Maximum points:  $2+0.5+0.5+2+1+1+1+1+1+1=10$

- a) Name one advantage and one drawback of small clusters in the file system compared to large clusters.

*Advantage: Decreasing capacity loss due to internal fragmentation.*

*Drawback: Rising overhead for large files.*

- b) UNIX file systems ignore uppercase and lowercase characters.

true       false

- c) Most operating systems operate according to the principle...

write-back       write-through

- d) Name one advantage and one disadvantage of a cache in the main memory, which is used by the operating system to accelerate the requests to stored data.

*Benefit: Better system performance.*

*Drawback: System crashes may cause inconsistencies.*

- e) `opsys_course/week_05/opsys_slides_05_en.tex` is...

an absolute path name       a relative path name

- f) Name two Linux file systems.

*ext2/3/4, btrfs, XFS, ReiserFS, JFS...*

- g) Name two Windows file systems.

*FAT12/16/32, NTFS, ReFS,...*

- h) Name one file system that implements a journal.

*ext3/4, ReiserFS, NTFS, XFS,...*

- i) Name one file system that implements extents.

*JFS, XFS, btrfs, NTFS, ext4,...*



Last name:

First name:

Student number:

---

## Question 8)

Points: .....

Maximum points:  $0.5+0.5+1+2+2+2+1=9$


x86-CPU's contain 4 privilege levels („rings“) for processes.

- a) Name the ring in which the kernel of the operating system runs.

*In ring 0 (= kernel mode) runs the kernel.*

- b) Name the ring in which the applications of the users run.

*In ring 3 (= user mode) run the applications.*



- c) Explain which processes (in which ring!) have full access to the hardware.

*Processes in kernel mode (ring 0) have full access to the hardware.*

- d) Explain what a system call is.

*If a user-mode process must carry out a higher privileged task (e.g. access hardware), it can tell this the kernel via a system call. A system call is a function call in the operating system, which triggers a switch from user mode to kernel mode ( $\implies$  context switch).*

- e) Explain what a context switch is.

*A process passes the control over the CPU to the kernel and is suspended until the request is completely processed. After the system call, passes the kernel, the control over the CPU back to the user-mode process. The process continues its execution at the location, where the context switch was previously requested.*

- f) Name two reasons why user mode processes should not call system calls directly.

*Working directly with system calls is unsafe and the portability is poor.*

- g) Explain the alternative concept to system calls, when user mode processes shall not call system calls directly.

*Modern operating systems provide a library, which is logically located between the user mode processes and the kernel.*

Last name:

First name:

Student number:

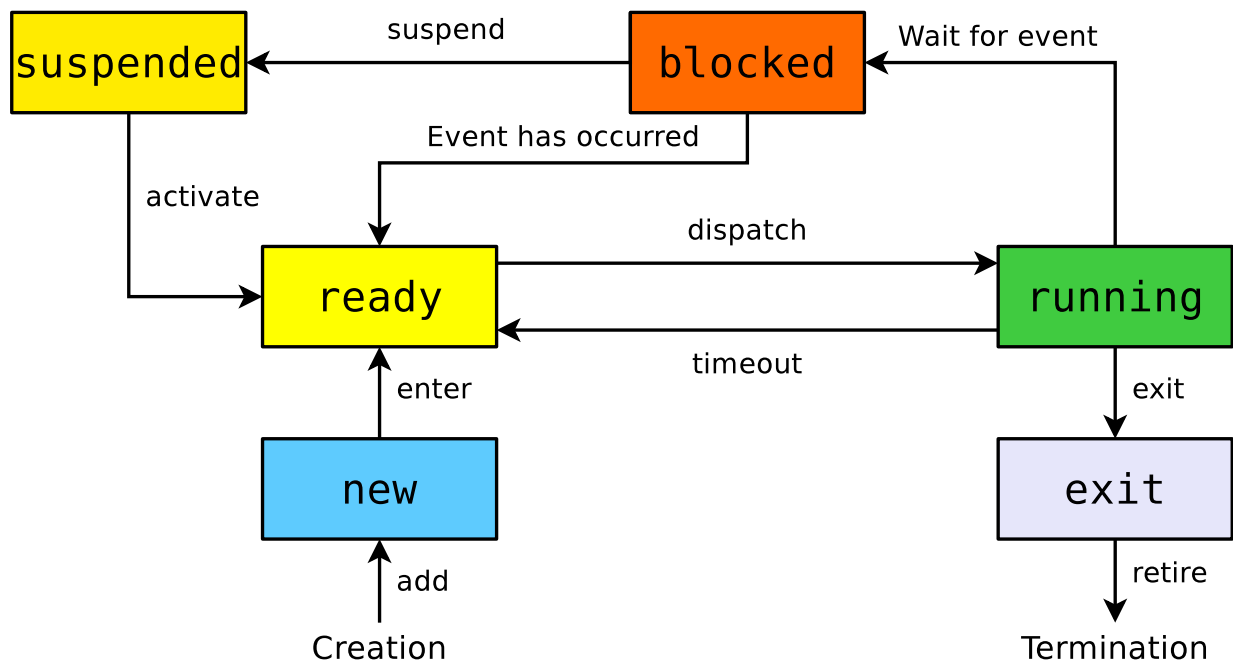
---

## Question 9)

Points: .....

Maximum points: 6+2=8

a) Enter the names of the states in the diagram of the process state model with 6 states.



b) Explain what a zombie process is.

*A zombie process has completed execution (via the system call `exit`) but its entry in the process table exists until the parent process has fetched (via the system call `wait`) the exit status (return code). Its `PID` still cannot be assigned to another process.*

Last name:

First name:

Student number:

---

## Question 10)

Points: .....

Maximum points:  $1+0.5+0.5+1+1+1+1+1=7$

- a) Explain what must be considered, when inter-process communication via shared memory segments is used.

*The processes need to coordinate the accesses themselves and to ensure that their memory accesses are mutually exclusive. A receiver process, cannot read data from the shared memory, before the sender process has finished its current write operation. If access operations are not coordinated carefully, inconsistencies occur.*

- b) According to which principle operate message queues?

Round Robin       LIFO       FIFO       SJF       LJF

- c) How many processes can communicate with each other via a pipe?

2

- d) Explain the effect, when a process tries to write data into a pipe without free capacity.

*The writing process is blocked.*

- e) Name the two different types of pipes.

*Anonymous pipes and named pipes.*

- f) Name the two different types of sockets.

*Connection-less sockets (also called: datagram sockets) and connection-oriented sockets (also called: stream sockets).*

- g) Explain what a critical section is.

*Processes carry out read and write operations on common data. Critical sections may not be processed by multiple processes at the same time.*

- h) Explain what a race condition is.

*It is an unintended race condition of two processes, which want to modify the value of the same record.*

Last name:

First name:

Student number:

---

## Question 11)

Points: .....

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

- a) Explain why external fragmentation of main memory is not a problem in the main memory management concept, implemented by modern operating systems.

*Because of the virtual memory concept paging, external fragmentation is irrelevant. The virtual memory is mapped to the physical memory. Processes do not need to be located in one piece inside the main memory.*

- b) Explain where internal fragmentation happens in the main memory management of modern operating systems.

*Modern operating systems implement virtual memory with paging. In this concept, internal fragmentation can only occur in the last page of each process.*

- c) Explain the advantage of signal and wait compared to busy waiting.

*When using busy waiting, computing time of the CPU is wasted because it is again and again occupied by the waiting process. Using signal and wait causes lesser CPU workload because the waiting process is blocked and later deblocked.*

- d) Explain what a barrier is.

*A barrier synchronizes the involved processes at one point.*

- e) Name the two problems that can arise from blocking.

*Starvation and deadlock.*

- f) Explain the difference between signaling and blocking.

*Signaling specifies the execution order of the critical sections of processes.*

*Blocking secures critical sections. The execution order of the critical sections of the processes is not specified. It is just ensured that the execution of critical sections does not overlap.*