

Name:	Neptun code:
-------	--------------

I. Test questions**Σ / 30 points****Description:**

Mark the proper answers! Mark Y, if you think that the statement is correct. Mark N, if you think that the statement is false. If you do not know the answer, please mark the "?" (question mark) sign. If you would like to correct your answer, write your answer in line with the question on the right side of the paper, your written answer should be correct, false, or "I don't know".

Scoring: A good answer is 1 point, a bad answer is -1 point. Minimum zero points can be earned from a question group (negative points are counted only for a question group, not for the whole test).

1st question group. Hardware architecture.

1.	The race condition between the DMA and the CPU can be solved by disabling the cache for memory areas accessed through DMA.	Y	N	?
2.	The support of multiple execution units on SMP machines are automatic, if the system supports the HW architecture with a single execution unit.	Y	N	?
3.	NUMA architecture can be cache-coherent or non-cache-coherent.	Y	N	?
4.	NUMA architecture is used only in High Performance Computing (HPC) applications, regular PCs (including servers) are implemented using the SMP architecture.	Y	N	?

2nd question group. Operating system architectures.

5.	Monolithic kernels are not used today, because they are not flexible enough.	Y	N	?
6.	Modern operating systems are implemented using the microkernel architecture.	Y	N	?
7.	Modular kernel allows the loading of kernel modules (drivers, OS components) run-time.	Y	N	?
8.	Microkernel based operating systems are essentially more stable, for example, they can survive faults in drivers.	Y	N	?

3rd question group. Time scale of scheduling.

9.	The task of short term scheduling is to select tasks for RUNNING from the tasks in the READY TO RUN state.	Y	N	?
10.	Medium term scheduling makes decisions about which of the tasks in the system are swapped to disk or swapped back to memory.	Y	N	?
11.	Long term scheduling selects which of the tasks are allowed to enter the system, in effect; it delays the running of the task issued by the user as long as they can be started then run efficiently.	Y	N	?
12.	Modern operating systems implements long, medium and short term scheduling.	Y	N	?

4th question group. Multiprocessor scheduling.

13.	Most modern general purpose operating systems use master-slave scheduling, i.e., one execution unit schedule tasks for all execution units.	Y	N	?
14.	Processor affinity of tasks influences system performance in SMP and NUMA systems.	Y	N	?
15.	In PULL type load balancing a special system task pulls tasks from highly utilized execution units and reschedules them on execution units with lower utilization.	Y	N	?
16.	Gang scheduling can increase the performance of highly interrelated tasks if properly implemented.	Y	N	?

5th question group. Processes and threads.

17.	Processes can be organized into a process tree because all processes have parents except the first process started in the system	Y	N	?
18.	If the parent process of a process is terminated, the process also terminates.	Y	N	?
19.	Processes share the heap, code, data memory, resources, but they have their own stack and virtual CPU.	Y	N	?
20.	User space threads (green threads) can be scheduled by the OS, because the users space thread library maps them to OS threads.	Y	N	?

6th question group. Windows HW and architecture specific questions.

21.	Windows 7 is only available on the x86 platform.	Y	N	?
22.	Windows has hardware specific, low level code in the HAL.	Y	N	?
23.	The Windows NT API is open, and documentation is available from Microsoft (downloadable from www.microsoft.com).	Y	N	?
24.	The Windows kernel and the Executive layers are in the same binary image (ntoskrnl.exe).	Y	N	?

7th question group. Windows scheduler related questions.

25.	Windows schedules processes.	Y	N	?
26.	Client version of windows uses lower time slices.	Y	N	?
27.	Windows threads have scheduling and I/O priorities.	Y	N	?

8th question group. Shared resources.

28.	A shared resource may be accessed by more than one parallel task.	Y	N	?
29.	In a single CPU system any machine instructions of the HW platform is atomic that cannot be interrupted.	Y	N	?
30.	A function called by multiple parallel tasks must be reentrant.	Y	N	?

II/1. Task**Σ / 10 points**

A system uses demand paging. A process may use 3 or 4 physical memory pages. The following pages are referenced during the run time of the process:

0, 1, 3, 2, 3, 0, 1, 4, 3, 2, 1, 2, 3, 4

How many page faults occur in the system, if the initially the physical memory pages are empty?

- First evaluate the FIFO algorithm for 3 and 4 physical memory pages (1.5-1.5 points).
- Then evaluate the Least Recently Used (LRU) algorithm for 3 and 4 physical memory pages. (2.5-2.5 points).

Explain how the results fit the theoretical background! (2 points)

The operation of algorithms must be also shown; we cannot accept only the final results!

II/2. Task

Σ / 10 points

Explain how paging hardware operates if TLB (Translation look-aside buffer) is used. Draw the block diagram of the paging hardware. (6 points)

Give the sequence of events if a TLB miss (2 points) or a TLB hit (2 points) occurs!