

**Operating Systems**  
**DIT 400, EDA092**

*Exam 2012-04-11*

---

*Date, Time, Place:* Wednesday 11/4 2012, 14:00-18:00, V building

*Course Responsible:* Arne Dahlberg Tel: 772 1705

*Auxiliary material:* You may have with you

- An English-Swedish, Swedish-English dictionary.
- No other books, notes, calculators, PDA's etc.

*Grade-scale ("Betygsgränser"):*

CTH:3:a 30-38 p, 4:a 39-47 p, 5:a 48-60 p

GU: Godkänd 30-47p, Väl godkänd 48-60 p

*Instructions*

- Do not forget to write your personal number, if you are a GU or CTH student and at which program ("linje").
- Start answering each assignment on a new page; number the pages and use only one side of each sheet of paper.
- Write in a **clear manner** and **motivate** (explain, justify) your answers. If it is not clear what is written, your answer will be considered wrong. If it is not explained/justified, even a correct answer will get **significantly** lower (possibly zero) marking.
- If you make any assumptions in answering any item, do not forget to clearly state what you assume.
- The exam is organized in groups of questions. The credit for each group of questions is mentioned in the beginning of the respective group. Unless otherwise stated, all questions in a group have equal weight.
- Please answer in English, if possible. If you have large difficulty with that and you think that your grade can be affected, feel free to write in Swedish.

**Good luck !!!!**

1. (10 p)
  - (a) Describe the SCAN and SSTF scheduling methods for optimizing head movement in disk memories. (4p)
  - (b) In what respect is C-SCAN more fair than SCAN? (1p)
  - (c) There is a number of different RAID levels for using multiple disks as one unit. Describe how RAID 5 and RAID 6 works. (4p)
  - (d) What is the most important difference between RAID 5 and RAID 6? (1p)
2. (10 p)
  - (a) File system operations in UNIX use a VFS (Virtual File System) layer for calling the actual file system code. Describe how the VFS interface works. (2p)
  - (b) A method to keep track of used blocks in a filesystem is FAT (File Allocation Table). Describe how the FAT method works. (3p)
  - (c) Describe two problems with allowing a filesystem to have a general graph structure (cyclic graph). (2p)
  - (d) Why do UNIX hard links not generate cycles in the file system graph? (1p)
  - (e) Protection mechanisms are sometimes described by an access matrix. What is an access matrix? (1p)
  - (f) What is the main difference between access lists and capability lists? (1p)
3. (10 p)
  - (a) Describe how a server initiated method for cache validation in a distributed filesystem works. (2p)
  - (b) The Andrew filesystem uses a consistency semantics called *Session Semantics*. Explain how *Session Semantics* works. 1(p)
  - (c) What is the main advantage with using *Session Semantics* compared to UNIX semantics? (1p)
  - (d) What does it mean that the naming method in a distributed file system is *location transparent*? (1p)
  - (e) VMware workstation can use a technique called *binary translation* to run an unmodified guest OS atop a hypervisor on X86 processors.
    - i. Explain how *binary translation* works. (3p)
    - ii. Why was *binary translation* used in VMware. (1p)
    - iii. How would it today be possible to run an unmodified guest OS atop a hypervisor on an X86 processor without using *binary translation* (1p)
4. (10 p)
  - (a) Explain how code can be shared among different processes in a system that uses virtual memory based on paging. (2p)
  - (b) Explain what a page table entry is and the typical information fields it contains. (2p)
  - (c) In a paging virtual memory system, a page fault will happen if the referenced page is invalid. Explain what the operating system have to do when a pagefault happens to make it possible for the process that caused the page fault to be restarted. (4p)
  - (d) In most cases dynamic linking is used in current operating systems. There are however cases when static linking may be preferred. Give to reasons why static linking may be preferred. (2p)

5. (10 p) The following processes with specified "burst" time, and priority have arrived:

Process	"Burst" time	Prioritet
P1	10	3
P2	1	1
P3	2	3
P4	1	4
P5	5	2

The processes have arrived in the order P1, P2, P3, P4 and P5 all at time 0 (zero).

- Draw a Gantt chart that illustrate the order in which the processes execute for the scheduling methods FCFS, SJF, Round Robin (quantum time = 1) and non preemptive priority scheduling (Low numeric value indicates high priority). (4p)
- What is the wait time for processes P1 and P2 for each of the scheduling methods in question a. (2p)
- A tricky issue in multiprocessor scheduling is how to design the ready queue. Describe this problem and its common solutions. (2p)
- In scheduling, the concept of **aging** is sometimes used. What does it mean and what problem does it try to solve? (2p)

6. (10 p)

- Many operating systems support semaphores. Explain what different types of semaphores that exist and how they work. (2p)
- Give an example of how a resource allocation graph can be used for deadlock detection. (2p)
- Why do Solaris, Linux and Windows XP use spin-locks as a synchronization mechanism in multiprocessor systems and not on single-processor systems? (2p)
- Design a solution to the mutual exclusion problem for arbitrary number of processes/threads in a multiprocessor system where the following *atomic instruction*, called swap is available by the hardware. Discuss carefully the correctness and the other properties of your solution. (4p)

```
void swap(int *a, int *b)
{
    int temp;
    temp = *a;
    *a = *b;
    *b = temp;
}
```