Chalmers Un. of Technology and Gothenburg Un. Comp. Science and Engineering Department

Operating Systems
DIT 400, EDA092

Exam 2011-12-14

Date, Time, Place: Tuesday 14/12 2011, 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 scheduling method for optimizing head movement in disk memories. (2p)

   (b) Describe how starvation could occur when using a SCAN algorithm to optimize the head movement in a disk memory. (1p)

   (c) Describe an improved version of SCAN that is free from starvation. (2p)

   (d) There is a number of different RAID levels for using multiple disks as one unit. Describe how RAID 0+1 and RAID 1+0 works. (2p)

   (e) Mention two types of errors that RAID do not protect against. (1p)

   (f) A file consists of 4 disk blocks. Which disk operations are needed to modify one block in the file if the filesystem uses RAID 5? (2p)

2. (10 p)

   (a) In almost all operating systems the first block in a disk partition is reserved for a special purpose. Which purpose? And is it necessary to reserve the block for this purpose? (2p)

   (b) What is the function of the open system call from the users point of view? (1p)

   (c) What can you expect to find in a UNIX file system directory for a normal file? (1p)

   (d) UNIX have a *link* system call for creating a hard link to a specified file. Describe the internal function (implementation) of the *link* system call. (2p)

   (e) Why do UNIX hard links not generate cycles in the file system graph? (1p)

   (f) Unix systems use inodes to keep track of used blocks in a filesystem. Describe this inode based method for block allocation. (3p)

3. (10 p)

   (a) What does it mean that a distributed filesystem has UNIX semantics? (1p)

   (b) Give one advantage and one disadvantage with using a stateless file server. (2p)

   (c) Describe two fundamental differences between AFS and NFS version 3. (2p)

   (d) Explain the difference between a type 1 hypervisor and a type 2 hypervisor. (1p)

   (e) What did Popek and Goldberg mean with
   i. A sensitive instruction (1p)
   ii. A privileged instruction (1p)

   (f) What was the Popek and Goldberg criteria for an architecture to be *classically virtualizable*? (1p)

   (g) Explain how paravirtualization works. (1p)

4. (10 p)

   (a) Describe the two main methods for load balancing on an SMP (Symmetric Multi Processor). (2p)

   (b) Explain the term *processor affinity*. (1p)

   (c) Why can load balancing sometimes be in conflict with processor affinity? (1p)

   (d) Describe two strategies for *deadlock prevention*. (4p)

   (e) A resource-allocation graph can be use to test for deadlocks. What does a loop in a resource-allocation graph indicate? Note that there are two different cases here. (2p)

5. (10 p)

   (a) One use of a MMU (Memory Management Unit) is to implement a demand paging virtual memory. Give two other reasons why UNIX systems require a MMU. For each alternative, motivate why a MMU is needed. (4p)

   (b) The following reference string is given:
   1, 2, 3, 4, 5, 3, 1, 4, 5, 1
   Assume that the memory have 3 page frames (empty from the beginning). Explain which pagefaults that happens and which pages are replaced when the reference string is executed with the following replacement algorithms (Replacement is done only if the memory is full).
   i. LRU (2p)
   ii. Optimal (2p)

   (c) Dynamic linking requires some support from the operating system. What support is needed and which system call is usually used? (2p)

6. (10 p)

   (a) Consider the following program:

```
const int n = 50;
int tally;

void total()
{int count;
for (count = 1; count <= n; count++)
{tally++;}
}

void main()
{tally = 0;
parbegin (total(), total());
write (tally);
}
```

   Determine the proper lower bound and upper bound on the final value of the shared variable tally output by this concurrent program. Assume processes can execute at any relative speed and that a value can only be incremented after it has been loaded into a register by a separate machine instruction. Justify your answer. (3p)

   (b) Suppose that an arbitrary number of these processes are permitted to execute in parallel under the assumptions of part (a). What effect will this modification have on the range of final values of tally? Justify your answer. (2p)

   (c) Describe how it is possible to devise a solution to the mutual exclusion problem using the TestAndSet hardware primitive. Argue about the correctness and other properties of the solution. (3p)

   (d) What is the risk of using busy-waiting for solving the mutual exclusion problem in combination with non-preemptive priority scheduling at a single CPU machine? Explain and justify your answer. (2p)