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

Operating Systems
DIT 400, EDA092

Exam 2008-03-12

Date, Time, Place: Wednesday 12/3 2008, 8:30-12:30, V building

Course Responsible: Arne Dahlberg, Marina Papatriantafilou (Tel: 772 1705, 772 5413)

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

Exam review ("Granskningstid"):
Monday 31st March, 13.00-14.00, Room 5128, EDIT building (west wing).

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.
- For questions 4, 5 and 6, 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) 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)

   (c) How is RAID 0+1 and RAID 1+0 affected by two erroneous disks? (2p)

   (d) What are the two most important reasons to use RAID systems? (2p)

   (e) 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) A method to keep track of used blocks in a filesystem is FAT (File Allocation Table). Describe how the FAT method works. (3p)

   (b) Describe two problems with allowing a filesystem to have a general graph structure (cyclic graph). (2p)

   (c) Which strategy is used in the fast Berkeley filesystem to decide in which cylinder group new data blocks for files are allocated. Also explain why this strategy is used. (2p)

   (d) The BSD filesystem used rotationally optimal placement of data blocks. Newer filesystems like the Linux ext3 filesystem use sequential placement.
   i. Which change in hardware have lead to sequential access being preferred today? Explain why. (2p)
   ii. Why are optimizations based on rotational position not useful any more? (1p)

3. (10 p)

   (a) The NFS filesystem implementations use a method called RPC (Remote Procedure Call) for communication between client and server. Describe the different steps in an RPC from the call until the result arrives. (3p)

   (b) Why is RPC often used for process communication in distributed file systems? (1p)

   (c) When doing the original implementation of NFS, an extra layer called VFS (Virtual File System) was added to the system interface. Why was the VFS layer added? (1p)

   (d) In NFS, an encoding called XDR (External Data Representation) is used for data sent between client and server. What is the reason for using this encoding? (1p)

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

   (f) What does it mean that the naming in a distributed file system is *location-independent*? (1p)

   (g) Explain why NFS cannot use a server-initiated method for cache validation. (1p)

   (h) What is the purpose with the NFS mount protocol? (1p)

4. (10 p)

   (a) The following data is given for a system: Demand paging, with the page table in internal registers. The memory access time is 100 ns. A page fault and the following read operation requires 8 ms plus another 12 ms if a modified page is replaced. Assume that 70% of the replaced pages are modified. What is the maximum page fault frequency that can be accepted if an effective access time of maximum 200 ns is wanted? (4 p)

   (b) Describe the clock algorithm for page replacement. Assume that you are monitoring the rate at which the pointer in the clock algorithm (the one that indicates the candidate page for replacement) moves. What can you say about the system if you notice that it is moving fast? What if it is moving slow? (4 p)

(c) Consider a logical address space of eight pages of 1024 bytes each, mapped onto a physical memory of 32 frames. i. How many bits are there in the logical address? ii. How many bits are there in the physical address? (2 p)

5. (10 p)

(a) Consider 3 processes, Alice, Bob and Charlie, that need exclusive access to a resource (e.g. a study room). Does the following solution proposal satisfy the desired requirements for mutual exclusion, progress, no starvation? Justify your response by discussing how each of the requirements is satisfied, if the solution is correct; or by giving a counter-example for the requirements that are not satisfied, if the solution is not correct.

The proposed solution uses 3 shared boolean variables `FA, FB, FC` (flag of Alice, Bob and Charlie, respectively) and a shared variable `SIGN`, whose value can be the initial of the "name" of any of the processes i.e. A, B, or C. Initially all flags are false and `SIGN` has value A.

```
Process A
repeat
  [do other things];
  FA := true ;
  SIGN := A;
  while (((FB == true) OR (FC == true)) AND (SIGN == A)) do [nothing];
  [use the shared resource];
  FA := false;
forever;

Process B
repeat
  [do other things];
  FB := true ;
  SIGN := B;
  while (((FA == true) OR (FC == true)) AND (SIGN == B)) do [nothing];
  [use the shared resource];
  FB := false;
forever;

Process C
repeat
  [do other things];
  FC := true ;
  SIGN := C;
  while (((FA == true) OR (FB == true)) AND (SIGN == C)) do [nothing];
  [use the shared resource];
  FC := false;
forever;
```

(b) i. Why is interrupt-disabling not appropriate for implementing synchronization primitives in multiprocessor systems?
ii. What is the meaning of the term "busy-waiting"? What other kinds of waiting can there be in an operating system?
ii. Why can busy-waiting solutions for synchronization be considered not appropriate for single-processor systems, while this is not exactly so in multiprocessor systems?

6. (10 p)

(a) Can a multithreaded program using user-level threads achieve better performance on a multiprocessor system than on a single-processor one?

(b) Describe how communication and synchronization among threads can be a factor that can affect scheduling decisions in multiprocessor systems.

(c) Why is it important for the scheduler to distinguish between I/O-bound and CPU-bound programs?

(d) What would happen if you executed the following piece of code:

```
main()
{ for( ; ; )
  fork();
}
```