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

Operating Systems
DIT 400, EDA092

Exam 2009-03-10

Date, Time, Place: Tuesday 10/3 2009, 8:30-12:30, V building

Course Responsible: Arne Dahlberg, Thomas Lundqvist (Tel: 772 1705, 160649)

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

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

   (c) Give two reasons why it is difficult for an operating system to improve the rotational delay when reading from a disk. (2p)

   (d) Assume that one disk in a RAID array becomes useless. Which measures need to be taken to recover full access to the data if using
   i. RAID 1 (1p)
   ii. RAID 5 (1p)

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

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

2. (10 p)

   (a) Explain the meaning of memory mapped files. (1p)

   (b) Explain the main steps needed for copying a file using memory mapped files. (2p)

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

   (d) What is the reason that symbolic links can reference across mount points but that this is impossible with hard links. (1p)

   (e) Which cache update strategy is used on the client caches in most NFS implementations? (1p)

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

   (g) Under which circumstances is a *location independent* naming method needed. (1p)

   (h) Describe two fundamental differences between NFS version 3 and NFS version 4. (2p)

3. (10 p)

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

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

   (c) What was the Popek and Goldberg criteria for an architecture to be *classically virtualizable*? (2p)

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

   (e) Give one advantage and one disadvantage with paravirtualization compared to traditional virtualization. (2p)

   (f) A common way for a hypervisor to handle I/O devices is by emulation. Explain how this works. (2p)

4. (10 p)

   (a) When a program is executed it is represented as a process in the operating system. A process can be in different states. Describe different process states and explain when transitions occur between the states. Include basic states as well as states related to virtual memory management. (3p)

   (b) In a system using preemptive scheduling, processes can be preempted. Describe the context (data) and the operations involved when a process gets preempted and a context switch is done. (3p)

   (c) Explain why switching between threads is faster than switching between processes. (2p)

   (d) Consider an operating system kernel that doesn't implement threads but only processes. Can threads still be used? If yes, explain how and also describe the advantages and disadvantages compared to threads handled by the kernel? (2p)

5. (10 p)

   (a) Describe the typical hardware present in modern processors to support virtual memory based on paging. Include at least these concepts (and explain them): Page table, page table entry, page number, frame number, TLB, TLB miss, page fault, page table pointer register. Use figures! Make sure to explain how addresses get translated. (6p)

   (b) Virtual memory systems always include the capability of memory protection. Explain how and why memory protection is used. (2p)

   (c) When using demand paging, trashing might occur. Explain (i) what trashing is, (ii) why it can happen and (iii) also mention a way to avoid it. (2p)

6. (10 p)

   (a) A software-based solution to the critical-section problem is Peterson's solution:

```
flag[i] = TRUE;
turn = j;
while (flag[j] && turn == j)
  ;
(critical section)
flag[i] = FALSE;
```

   There are also hardware-based instructions that can be used like "test-and-set" or "swap".

   A critical sections solution should satisfy three important requirements. List these requirements and explain which of these requirements that are satisfied when using (i) Peterson's solution and (ii) a hardware "test-and-set" instruction. (3p)

   (b) Many operating systems support semaphores. Explain what different types of semaphores that exist and how they work. (2p)

   (c) Show how to solve a basic readers-writers problem using semaphores. (3p)

   (d) Give an example of how a resource allocation graph can be used for deadlock detection. (2p)