Name: email: ID:

Final Exam
CS402
14 Dec 00

You have 2 hours for this exam. The exam has 8 pages. There are 100 possible points. Show all your work for partial credit.

Definitions

Each question includes the number of points. Answer all questions in this section. You may use unambiguous abbreviations.

1. Give a term that describes the following:

   a) The four requirements for deadlock (4 pts)

      Answer: Mutual exclusion, non-preemption, hold and wait, circular wait

   b) The observation that FIFO page replacement may suffer an increase in page faults if memory is added (1 pt.)

      Answer: Belady's anomaly

   c) Operating system code executed when an asynchronous device signals the CPU (1 pt)

      Answer: interrupts service route (ISR)

   d) A UNIX® file that contains filename to i-node mappings. (1 pt.)

      Answer: directory

   e) The condition when the work done by the operating system is greater than the work done by the user processes. For example, an OS that is spending more time servicing page faults than running computations is doing this. (1 pt.)

      Answer: thrashing

   f) A disk arm scheduling policy (1 pt.)

      Answer: FIFO, Shortest seek first (SSF), elevator algorithm

   g) The layer of the ISO stack concerned with carrying data end-to-end (1 pt.)

      Answer: transport

Name: email: ID:

Short Answer

Each question is worth 5 points. Do all questions in this section.

2. What is the main difference between a process and a thread? (3 pts.) Explain this difference in terms of nachos data structures. (2 pts.)

   Answer: A thread is a running computation that shares an address space with other running computations. A process is a thread that is alone in its address space. In terms of nachos, a process would be an address space with a single thread running in it. The Thread data structure encapsulates threads.

3. Show the results of the following memory transactions on the buddy system illustrated below. The system has 1MB of memory. Each blank should be filled with the state of the system after one of the given transactions. An allocation for a process is a box with the process identifier (A), an unallocated amount of memory is a box with the unallocated size (in KB) in it (1024). This is the same convention used in the examples in class and in Tannenbaum. (5 pts).

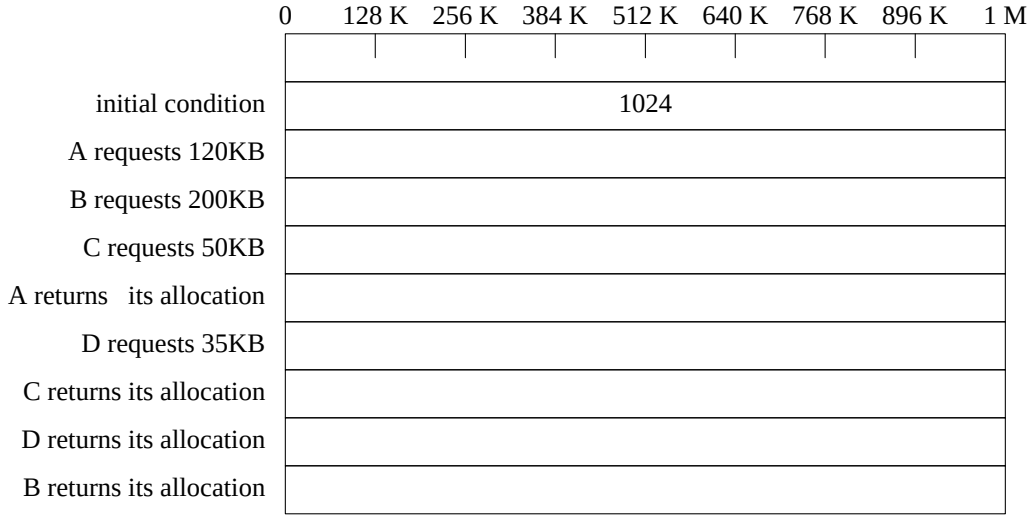

4. Consider a filesystem that keeps file attributes in the directory entries (as MS-DOS™ does) and one that keeps attributes in a separate data structure in the file system (for example UNIX i-nodes or nachos file headers). In which system is it easier to implement hard links? (1 pt.) In which system is it easier to implement soft links (also called symbolic links)? (1 pt.) Justify both answers. (3 pts.)

   Answer:

5. UNIX treats devices as files, by putting them into the file system name space and supporting file operations on them. Give an advantage of this approach and a disadvantage of this approach. (2 pts.) Support both your advantage and disadvantage by an example including a specific device. You should give an example for your advantage and your disadvantage, and the example devices may be different. (4 pts.) [This question is worth 6 points]

   Answer: Putting devices into the file system extends the file protection system to cover devices (advantage). This allows the administrators to restrict access of a device to only certain users of the system - for example allowing only administrators to use a CD-ROM burner or tape drive for backups. Using the file system as an interface is less useful when the device supports more interesting primitives than read/write/execute (disadvantage). Examples of devices that have interesting capabilities that are not well handled by the file system are CD-ROM drives that can play music or graphics coprocessors that draw lists of polygons or shadings.

Name: email: ID:

6. The following jobs are given in order of their order of arrival in the run queue. Give their completion order under the given scheduling policies.

Job	RunTime	Priority
A	5	1
B	10	4
C	7	3
D	11	2

Discipline	Order
Shortest Job First, run-to completion (1 pt.)	
Priority, run-to completion (1 pt.)	
FIFO, preemptive multitasking with a small quantum (1 pt.)	

Which discipline gives the shortest average response time? (1 pt.) [This question is worth 4 points]

Answer:

Discipline	Order
Shortest Job First, run-to completion (1 pt.)	A, C, B, D
Priority, run-to completion (1 pt.)	A, D, B, C
FIFO, Preemptive multitasking with a small quantum (1 pt.)	A, C, B, D

SJF gives the shortest average response time.

7. What problem does RPC (remote procedure call) argument marshalling address? (1 pt.) Describe the marshalling that happens on both the caller and the callee on an RPC call of a function that takes one integer parameter. (4 pts.)

Answer: RPC marshalling allows processes with one internal data representation to request services on a machine with a different internal representation. A good example is a big-endian machine requesting services from a little endian machine. When a single integer call is made in RPC, the integer is translated into a common data format and transmitted to the callee who translates it from the common format to its internal format to do the computation. Any results undergo the reverse transformation.

8. For caching to be effective, the cached elements must share a property. Name or describe that property. (3 pts.) Describe an application or data access pattern for which caching is not particularly effective. (2 pts.)

   Answer: Caching requires locality to be effective. Data elements asked for in the recent past must be good candidates for future requests. Data without locality are poor candidates for caching. Examples include random accesses into a large file, or applications like the world wide web which encourage non-locality. The web encourages people to seek new sites.

9. Sectors from the following tracks have been requested in the given order. What tracks will the head traverse if under the following algorithms:
   Queued tracks: 11, 35, 10, 9, 8, 21, 3, 33

Discipline	Order
FIFO (1 pt.)	
Shortest Seek First (2 pts.)	
Elevator Algorithm (2 pts.) (Head starts on 11 going to higher tracks)	

Answer:

Discipline	Order
FIFO (1 pt.)	11, 35, 10, 9, 8, 21, 3, 33
Shortest Seek First (2 pts.)	11, 10, 9, 8, 3, 21, 33, 35
Elevator Algorithm (2 pts.) (Head starts on 11 going to higher tracks)	11, 21, 33, 35, 10, 9, 8, 3

10. Although they are both hierarchical name spaces, the DNS (Internet domain name system) and the UNIX filesystem differ. Name a goal that the DNS has that the UNIX filesystem does not share. (2 pts.) Describe how the DNS meets the goal you described. (3 pts.)

    Answer: The domain name system must support autonomous administration of parts of the name space and queries from the world. The autonomy is accomplished by dividing the name space between multiple servers. The large numbers of queries are accomplished by splitting the load among servers, so each has only a fraction of the load. Widespread caching also assists.

11. The computer's main memory is used for both physical pages for the VM system and for the filesystem block cache. Some systems allocate some memory for each of these when the system starts, but other systems use the same memory area for both, dividing it dynamically. Describe 2 advantages to using a combined VM/filesystem cache. (5 pts.)

    Answer: Allocating each of the caches separately means that the allocation must be tuned correctly at startup. If the allocation is incorrect for the workload, or if the workload changes, some of the memory will be unused while the rest is overused. Dynamically allocating to each allowed the system to adapt. Also when an executable is loaded, data from the file is first loaded into the filesystem cache when it's read and then into the VM. Using a combined cache allows the data to migrate from one cache to the other, avoiding the duplication. With a little work, such data can be shared among multiple running processes.

Name: email: ID:

Long answer

Each question gives its value. Do all questions in this section.

12. Caching is present in practically all subsystems of an operating system. This question will explore some caching related issues.

    a) Modern CPUs have instruction and data caches that keep the most recent memory accesses. In VM systems these caches are indexed by virtual address. If these caches are not properly managed by the OS, process separation may be compromised. Describe the potential problem (3 pts.) and how to correctly manage them to avoid the problem (2 pts.).

    Answer: If the caches are not cleared when a context switch occurs, the newly scheduled process will be able to see data from the old process, and may even erroneously execute instructions from that process. The way to avoid this is to flush the CPU caches during a context switch.

    b) The file system caches store disk blocks that may be required again. To a first approximation this cache can be managed as a write back (i.e., delayed write) least recently used cache. A block is kept in the cache until the cache is full, a new block must be cached, and the outgoing block is the oldest in the cache. If modified, it is written to disk then. Describe 2 modifications to basic LRU that are needed in managing the file system cache and why they are needed. (5 pts.)

    Answer: File system metadata is important for the correct function of the operating system, and must be stored write through or the filesystem may be corrupted by a badly timed system failure (3). Dirty blocks must be occasionally flushed to disk, or often accessed dirty blocks will never be written to disk. This isn't a problem for dirty but popular VM blocks, because they don't have to survive power outages.

    c) Some modern disk drives pre-fetch and cache blocks that reside on disk near the one read on the disk controller. What aspects of disk access speed does this alter? (2 pts.) How can an operating system take advantage of this caching? (3 pts.)

    Answer: The caching masks the effects of seek and rotation, because the next blocks are likely to be read from cache instead of disk. A system can spend less time optimizing file layout on the disk which may speed access and simplify code.

    d) Describe the use of hardware caching in the virtual memory system. (5 pts.)

    Answer: The hardware caching in the VM system is the translation lookaside buffer that caches virtual to physical mappings.

Name: email: ID:

13. Consider an ATM (automated teller machine) system. The system consists of multiple machines, each holding a set of numbered accounts, and a set of ATM machines. To economize, the bank is using the global Internet to connect these machines and ATMs. Each ATM knows which machine holds each account.

    a) What pairs of elements need to authenticate themselves to each other? (4 pts.) For each pair, explain the possible consequences of not authenticating. (4 pts)

    Answer: user → ATM an unauthenticated user could alter another user's account
    ATM → user a user should not give private information to an ATM it does not trust.
    ATM → account machine an untrusted ATM cannot be allowed to access bank resources
    account machine → ATM information could be lost and transactions not completed.

    b) For each money transfer between accounts the ATM locks the higher numbered account and then the lower numbered account and transfers the data. Is there a danger of deadlock here? (1 pt.) If so, give an example, if not explain why not. (3 pts)

    Answer: No danger. The ordering of resources and acquiring them in order makes circular wait impossible. Because one of the four conditions for deadlock is impossible, deadlock is impossible.

    c) As the number of ATM machines increases, you should expect that the system performance will be reduced even if the network performance keeps pace. Explain why. (4 pts.) Suggest and justify a strategy to improve performance. (4pts.)

    Answer: The performance will be reduced because the number of requests per account machine will rise. To reduce it add more account machines and make each responsible for fewer accounts. It's possible to do more complex things, but this is the simplest solution.