

2010 Autumn Answers [OS II]

Question 1

Scheduling is an important function of the operating system kernel.

a) Present the scheduling process in general, pointing out the sequence of operations that are carried out. [10 marks]

- The purpose of process scheduling is to order the processes that are ready to execute in such a manner that allows them to be completed with the highest level of efficiency. The processes are organized in a queue from where the scheduler selects the next one to take control of the CPU.

b) What is priority scheduling? [5 marks]

- This is one implementation of dynamic priorities.
- Initially, a process gets a priority that puts it on a certain level. After each time slice, the priority is lowered to the next level, until it reaches the lowest acceptable priority. At that level, the strategy is round robin.
- However, after being blocked, the process gets a higher priority (priority boost).

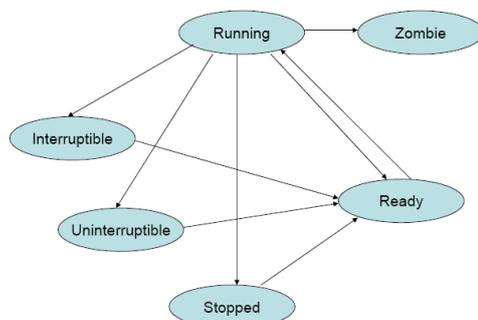
c) Explain the concept of process group scheduling used by multi core systems. [5 marks]

- Process group scheduling is grouping processes together to be scheduled in order to minimise shared resource contention.
- Threads and processes sharing memory segments can be co-scheduled on processors sharing the same cache. Since they share data and other context, this minimises resource contention and resource duplication.

Question 2

While managed by the operating system, a process switches between different states.

a) Consider Linux and show graphically the states of a process. Define each state and state transition from one state to another. [10 marks]



- Running refers to both running and ready

- Interruptible: process blocked, can be awakened by signals
- Uninterruptible: process blocked cannot be awoken by signals
- Stopped: process halted
- Zombie: child process terminated early, still in memory

b) How does Linux assign the process ID? [5 marks]

- The function `alloc_pid` maintains a pointer to the last process id allocated.
- It will try to allocate `last_pid + 1`, but since process ids can be reused, it must check for availability.
- A bitmap of process id values being used is checked, starting at `last_pid + 1`. If the desired id is not available, the next available in the bitmap will be located using linear search.
- If there are no available ids, the error code `EAGAIN` will be issued. `alloc_pid()` can be attempted again when a process terminates and its process id becomes free.
- When a free id is found, it is returned and used as the process id for the newly created process.

c) What is load balancing? [5 marks]

- ‘Refer to balancing of process load between cores in a scheduling domain’

Question 3

Memory allocation to processes is an important function of the operating system.

a) Explain the Buddy system of memory allocation. [10 marks]

- The buddy memory allocation algorithm’s main feature is the splitting of a free block of memory into two equally sized “buddy” blocks, where one is returned for allocation and the other is kept free. These blocks can be combined again at a later stage to form the original sized free block.
- All free memory blocks have a size of some power of 2. If the smallest block size equal or greater than a memory request cannot be found, a block of twice that size is split into two buddies. The first buddy is offered for allocation.

b) What is swapping? [5 marks]

c) In the context of swapping, present two page replacement policies, pointing out their features. [10 marks]

- First In First Out starts from the idea that pages are used for a finite amount of times after which they become “old”. The page selected here is the one that has been in memory the longest. Implementation is done by a queue – all new pages are added to the tail of the queue.
- Second chance is an extension of FIFO: when a page is pulled off the head of the queue, the accessed (A) bit is examined. If its 0, the page is swapped out, else the bit is cleared and the page is reinserted at the tail of the queue.

Question 4

The file system is part of the data management in any computing system.

a) Explain the notion of file system metadata. [5 marks]

- File system metadata is information about a file system and its contents.
- The file system metadata includes the total size of the system, the amount of free space, the date of the last mount etc.

b) How is the free bitmap used to manage the free space? [5 marks]

c) Present the way a directory is implemented. [5 marks]