

Lifecycle

- ▶ Suppose we have two processes that require the CPU. The first one had the CPU and you would like to let the second process run, ie context switch. Should you do it at this time?
 - Cost of context switch
 - Opportunity cost of flushing TLB/cache
 - Cost of losing IO locality for file system
 - Cost of flushing buffers to disks and bringing in new pages
 - Pages might be wired during transfer preventing new process from running (by making them wait for memory to be freed by previous process which was context switched and hence is not running anyways)
- ▶ A good scheduler would optimize across all these parameters: quickly




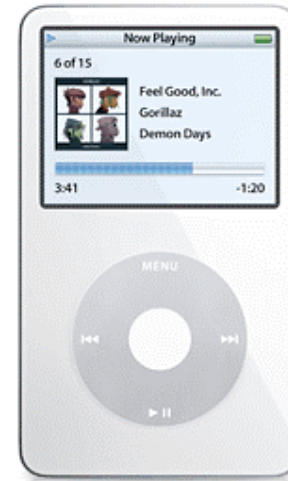
Designing an Operating Systems

- ▶ Goal is to understand how the technologies that we studied so far apply to typical machines
- ▶ First we focus on PDAs and Laptops
 - Both are mobile, inexpensive
 - Battery is a big concern
 - Quick startup
 - Quick shutdown
 - Frequent suspends



PDA

- ▶ Small mobile devices
- ▶ Important design elements:
 - Inexpensive
 - Mobile (small, rugged, good battery life)
 - Constrained CPU, memory, storage, screen
 - CPU: 200 MHz
 - Memory: 64 MB
 - Storage: Flash or Microdrive
- ▶ OS: Symbian, PalmOS, MS Windows Mobile, QNX, Linux?



PDA and Process

- ▶ Usually: only one user, process at any one time
 - Palm context switches by “freezing” process state and unthawing old process
- ▶ Process Synchronization: Little system support.
- ▶ Many multimedia applications (video, audio, cellular calls)



PDA and memory/storage

- ▶ Usually no MMU
- ▶ Storage: Flash or Microdrive
 - Flash has no moving components, however can only be rewritten a finite number of times
 - Mobile device and so storage should be consistent



PDA and security

- ▶ Heavily uses physical security feature
- ▶ Overall: What is the roll of PDA and whatever we learnt?
 - Why do we even discuss PDA class machines?



Laptop class

- ▶ Important design factors:
 - Cost, weight
 - CPU: as fast as your lap can tolerate
 - Memory: up to 4 GB
 - Disk: up to 200 GB
 - Sandisk 32 GB flash disk
 - Energy consumed depends on amount of resource
- ▶ OS: MS Windows, Mac OSX, Linux, FreeBSD, ...



Laptops and Processes

- ▶ Modern laptops are multi-core
 - Mostly interactive tasks and hence prefer interactive applications
 - Frequent suspend - does that affect scheduling?
- ▶ Process synchronization
- ▶ Users use productivity apps, multimedia apps and solitaire



Laptops - memory and storage

- ▶ What do you do with 4 GB on a laptop?
 - Leave memory of exited programs to quicken startup?
 - Energy cost
 - Use massive buffered IO?
 - Reliability when memory runs out
- ▶ Disks and Flash
 - Disks support fully operational, spin-down, park modes



Laptops and protection

- ▶ Physical security still possible
- ▶ Rarely multiuser



Desktop

- ▶ Dual processor/quad core
 - 3+ GHz dual core x2 and 64 bit processor
- ▶ GBs of memory
- ▶ Multiple hard disks
 - Hard disk can be up to 750 GB per disk!!



Desktop and Process scheduling...

- ▶ What do you do with these beasts?
 - Web browse
 - Emails
 - Word
 - Multimedia encoding/creation
- ▶ Scheduling a balance of interactive and batch processing




Memory and File system

- ▶ RAID becomes increasingly necessary for most machines, given that 250 GB hard drive is ~\$60
- ▶ Desktops, if they knew that they would be on UPS, can afford to really use a lots of caching and buffering
- ▶ Security wise, desktops are similar to workstations in that they are single user at a time



Data center server

- ▶ One of the specification is the size that server will take in a rack. 1U is the smallest size and blade servers, which fit one unit are all the rage
- ▶ Dual (Quad Core Xeon, 2x4MB Cache, 2.66 GHz, 1333 MHz FSB), 16 GB memory, 2x73GB 15k rpm hard disk - \$10000
- ▶ 1 rack - 60 racks
 - (\$ 0.6 m)



Servers

▶ Mission critical systems

- Three tier systems - production, backup and test
- Virtual hosting to protect against interference with other processes
- Data center support service level agreements (SLA) - OS should be aware of these
- On demand computing
- Autonomic management

▶ Each rack can consume 10 Kw

- Additional 10 Kw in cooling
- Data center can be powered exclusively by a 300 MW power station.



Hot topics

► Hot research areas:

- Energy management for servers/laptops
- Virtual machine support for isolation (Java, Xen, VMWare, Parallels, Wine etc.)
- Grid/cluster computing to harness lots of machines
- Autonomic OS/storage etc.

