

# MODERN OPERATING SYSTEMS

Third Edition

ANDREW S. TANENBAUM

## Chapter 1 Introduction

# What Is An Operating System (1)

A modern computer consists of:

- One or more processors
- Main memory
- Disks
- Printers
- Various input/output devices

Managing all these components requires a layer of software – the **operating system**

# What Is An Operating System (2)




Figure 1-1. Where the operating system fits in.

# The Operating System as an Extended Machine

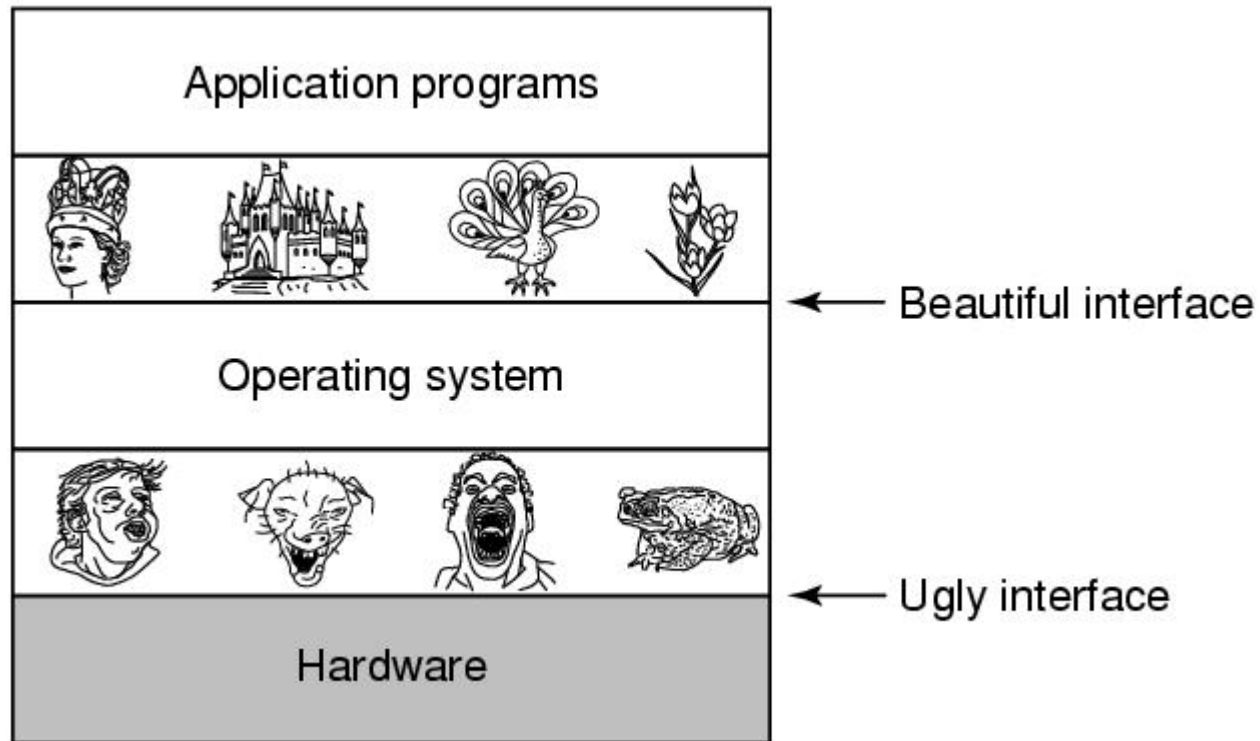


Figure 1-2. Operating systems turn ugly hardware into beautiful abstractions.

# The Operating System as a Resource Manager

- Allow multiple programs to run at the same time
- Manage and protect memory, I/O devices, and other resources
- Includes multiplexing (sharing) resources in two different ways:
  - In time
  - In space

# ICs and Multiprogramming




Figure 1-5. A multiprogramming system with three jobs in memory.

# Computer Hardware Review

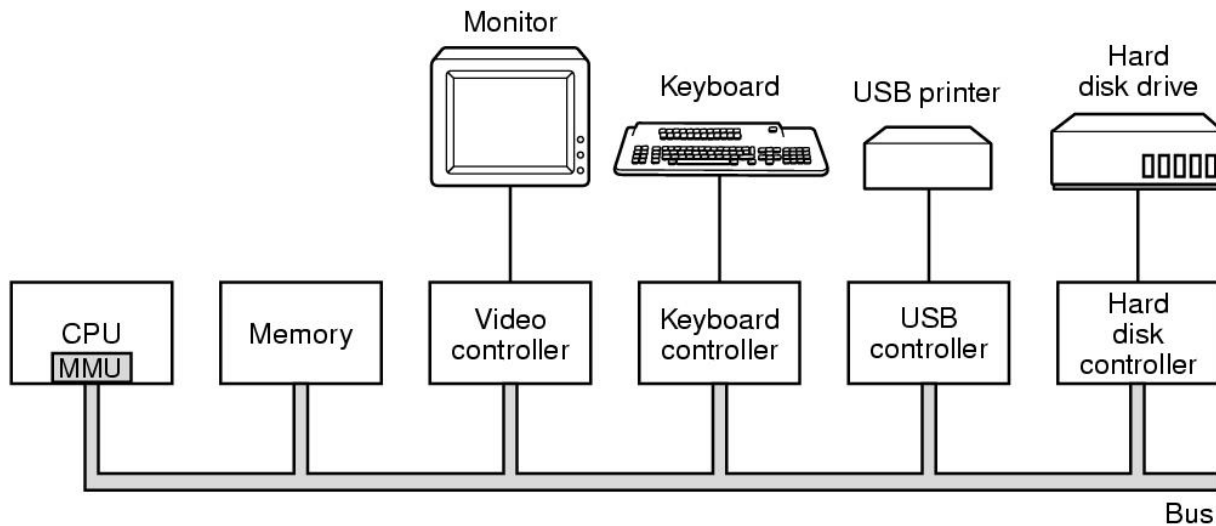


Figure 1-6. Some of the components of a simple personal computer.

# Multithreaded and Multicore Chips

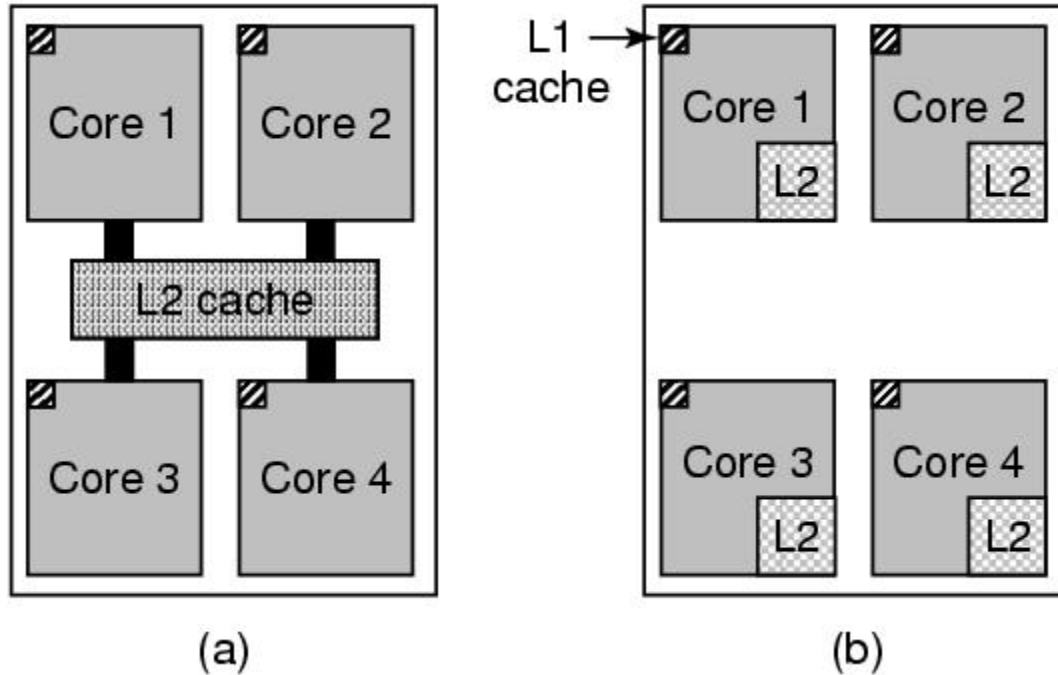


Figure 1-8. (a) A quad-core chip with a shared L2 cache.  
(b) A quad-core chip with separate L2 caches.



# Memory (1)

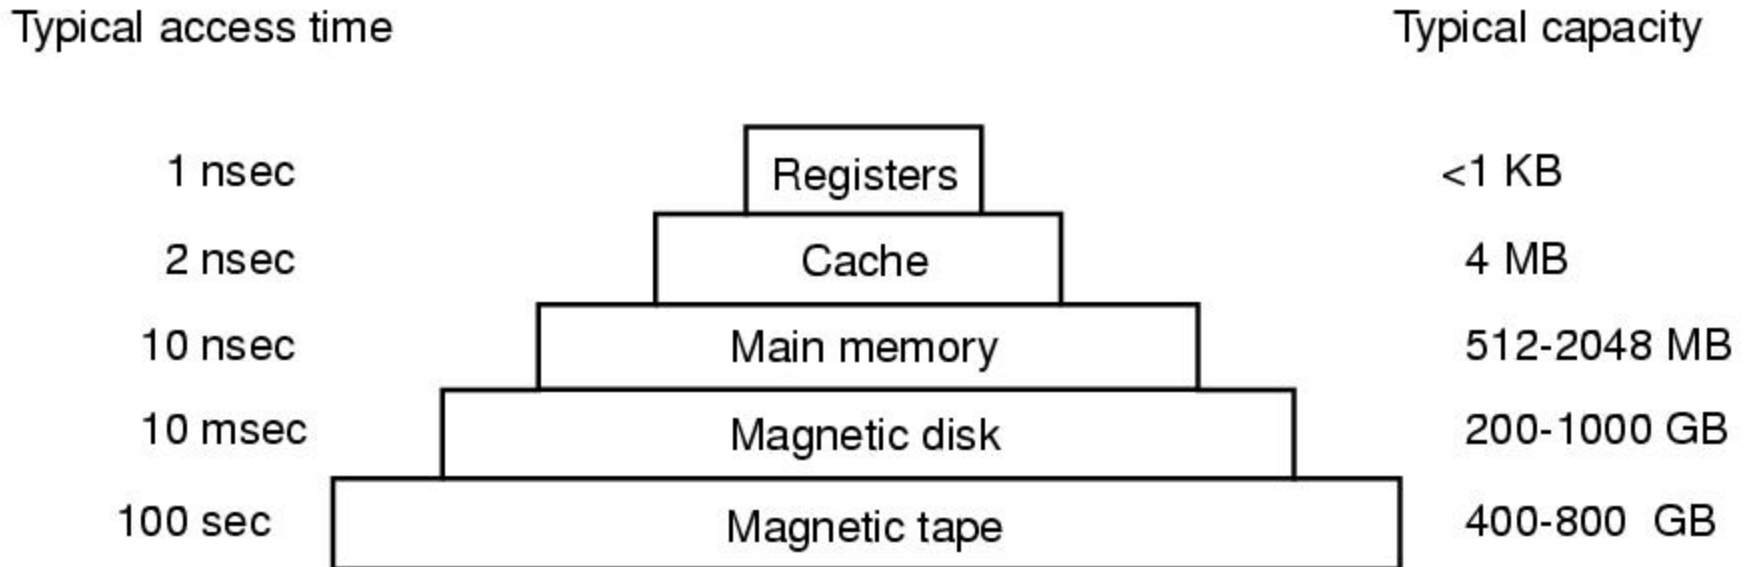


Figure 1-9. A typical memory hierarchy.  
The numbers are very rough approximations.

# Disks




Figure 1-10. Structure of a disk drive.

# I/O Devices

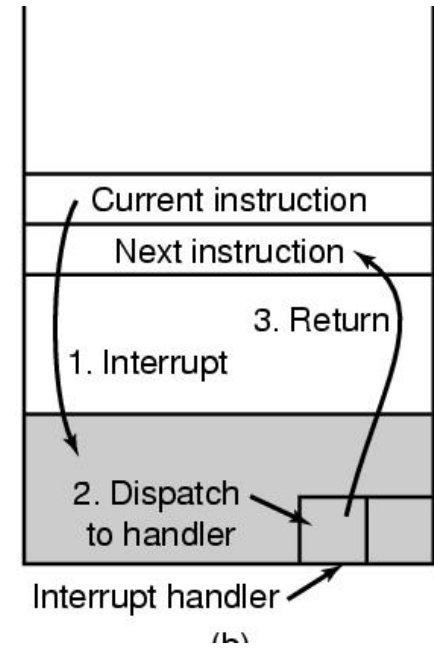
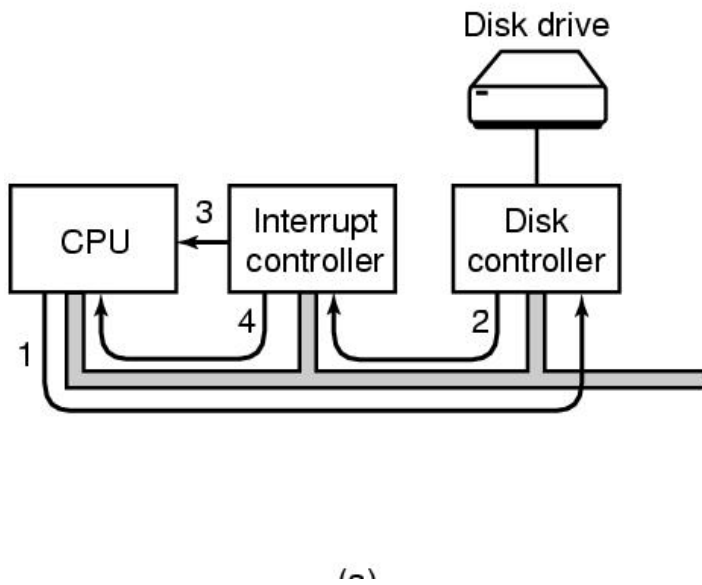
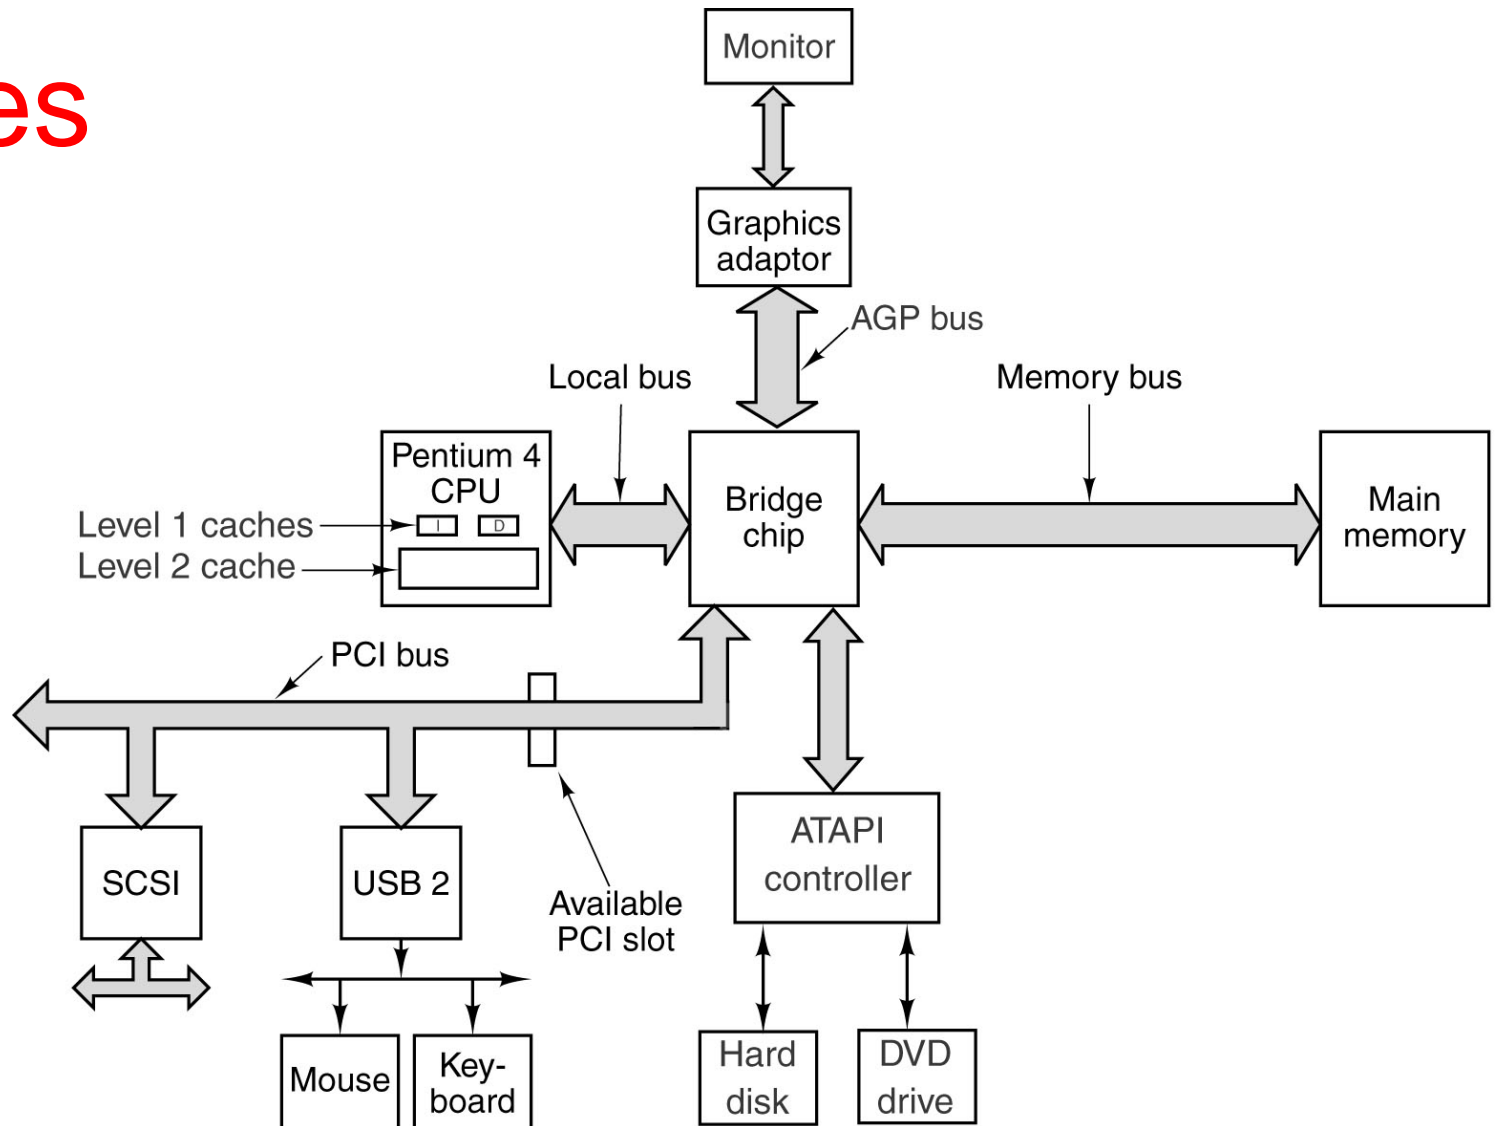



Figure 1-11. (a) The steps in starting an I/O device and getting an interrupt.

# Buses



The bus structure of a modern Pentium 4.

# Enterprise: 2008 Nehalem Based Two Socket System Architecture



## Nehalem-EP Platform:

Two sockets each with Integrated Memory Controller

Turbo mode operation

Intel QuickPath Architecture

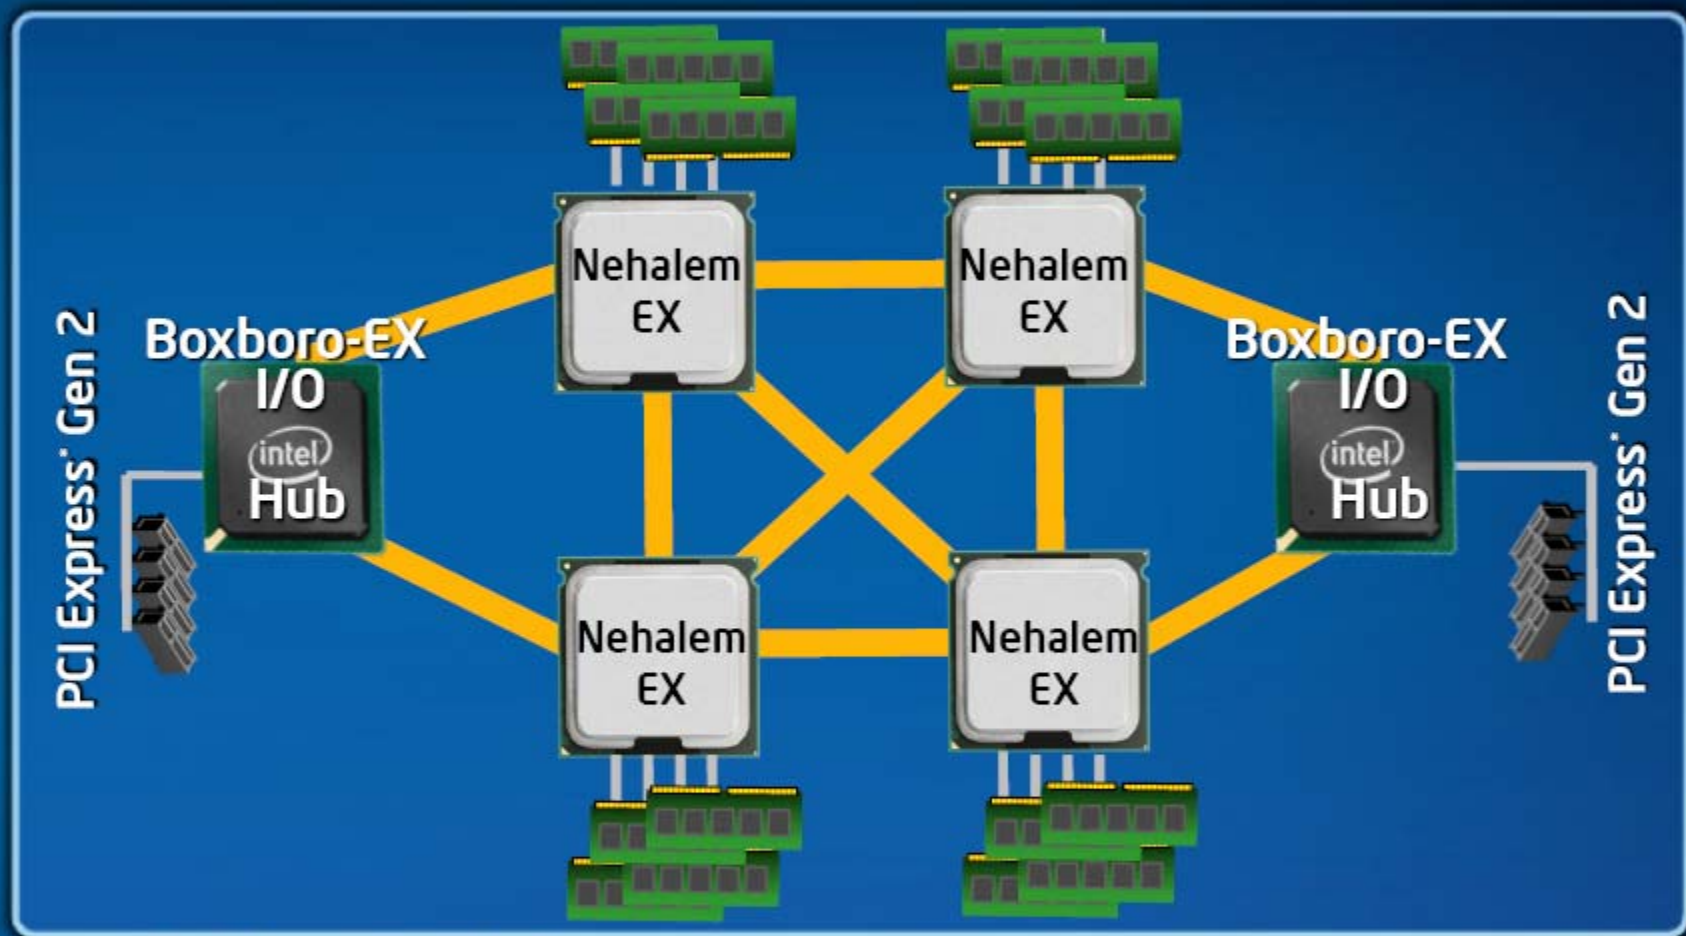
DDR3 Memory: 3 Channel, 3 DIMMs per channel

Intel Virtualization Technology

PCI Express\* Gen 2

Intel® QuickPath Interconnect

# Enterprise: 2009 Nehalem Based Four Socket System Architecture



## Boxboro-EX Platform:

Intel® QuickPath Interconnect

Four processors with Intel® QuickPath Interconnects  
PCI Express® Gen 2, Integrated Memory Controller



# The Operating System Zoo

- Mainframe operating systems
- Server operating systems
- Multiprocessor operating systems
- Personal computer operating systems
- Handheld operating systems
- Embedded operating systems
- Sensor node operating systems
- Real-time operating systems
- Smart card operating systems

# Operating System Concepts

- Processes
- Address spaces
- Files
- Input/Output
- Protection
- Ontogeny recapitulates phylogeny
  - Large memories
  - Protection hardware
  - Disks
  - Virtual memory



# Processes




Figure 1-13. A process tree. Process A created two child processes, B and C. Process B created three child processes, D, E, and F.

# Files (1)

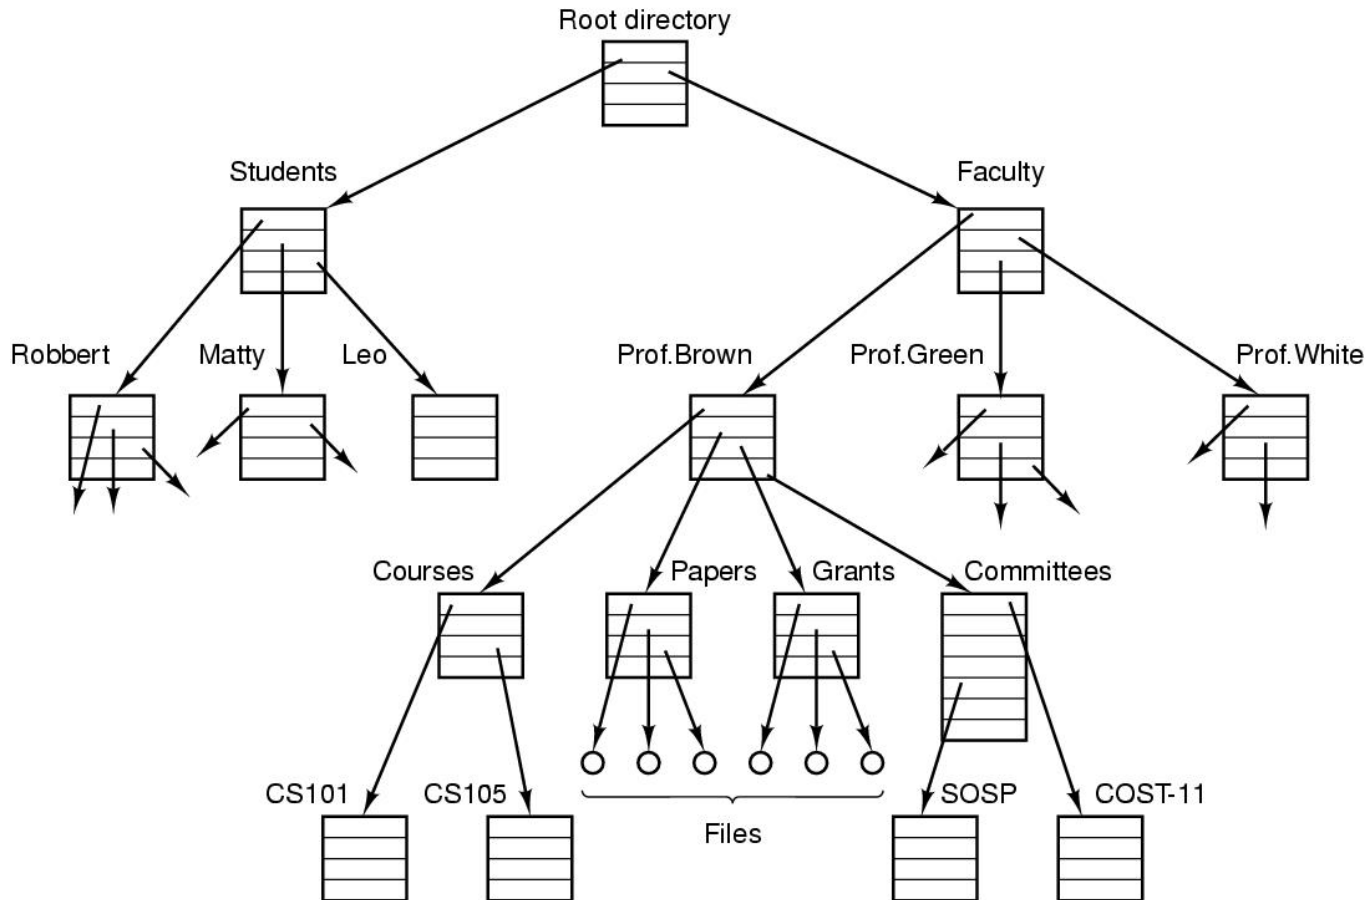
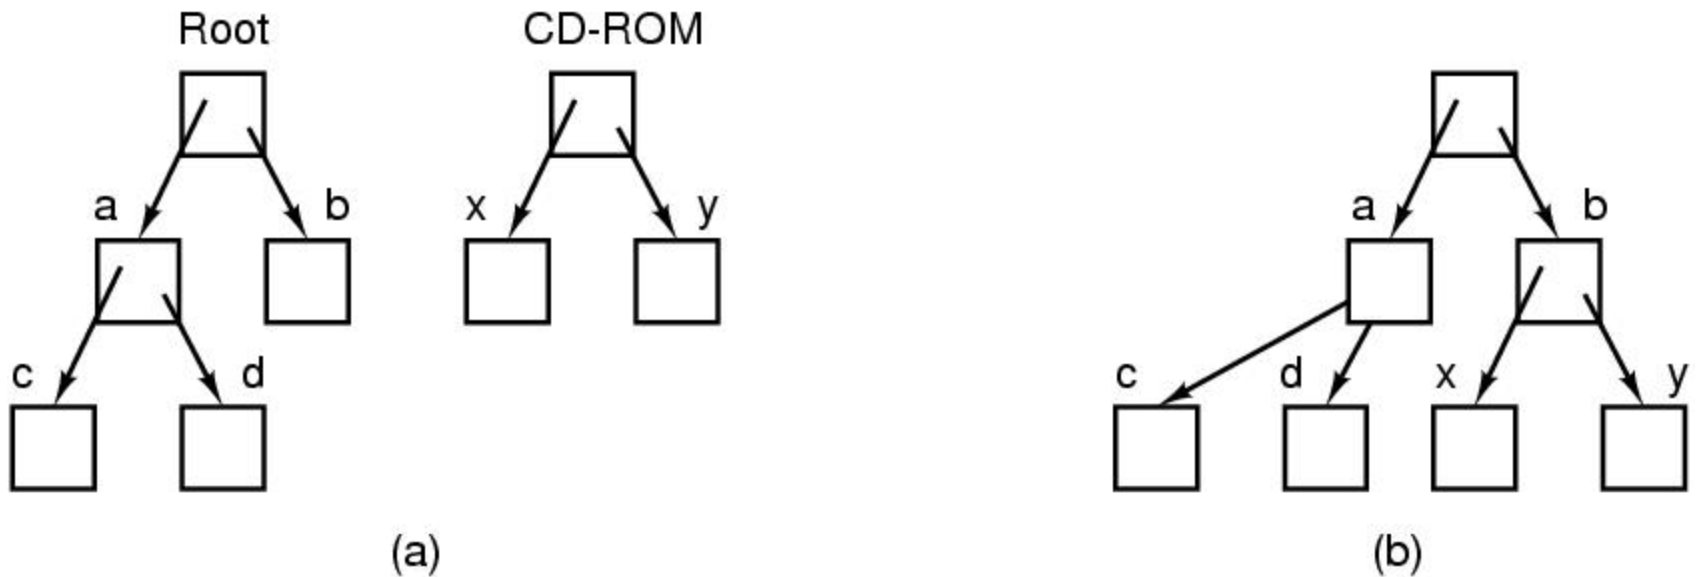


Figure 1-14. A file system for a university department.

# Files (2)



(a)

(b)

Figure 1-15. (a) Before mounting, the files on the CD-ROM are not accessible. (b) After mounting, they are part of the file hierarchy.

# Files (3)




Figure 1-16. Two processes connected by a pipe.

# System Calls




Figure 1-17. The 11 steps in making the system call `read(fd, buffer, nbytes)`.

# System Calls for Process Management

## Process management

Call	Description
<code>pid = fork( )</code>	Create a child process identical to the parent
<code>pid = waitpid(pid, &amp;statloc, options)</code>	Wait for a child to terminate
<code>s = execve(name, argv, environp)</code>	Replace a process' core image
<code>exit(status)</code>	Terminate process execution and return status

Figure 1-18. Some of the major POSIX system calls.

# System Calls for File Management (1)

File management	
Call	Description
fd = open(file, how, ...)	Open a file for reading, writing, or both
s = close(fd)	Close an open file
n = read(fd, buffer, nbytes)	Read data from a file into a buffer
n = write(fd, buffer, nbytes)	Write data from a buffer into a file
position = lseek(fd, offset, whence)	Move the file pointer
s = stat(name, &buf)	Get a file's status information

Figure 1-18. Some of the major POSIX system calls.

# System Calls for File Management (2)

<b>Call</b>	<b>Description</b>
s = mkdir(name, mode)	Create a new directory
s = rmdir(name)	Remove an empty directory
s = link(name1, name2)	Create a new entry, name2, pointing to name1
s = unlink(name)	Remove a directory entry
s = mount(special, name, flag)	Mount a file system
s = umount(special)	Unmount a file system

Figure 1-18. Some of the major POSIX system calls.



# Miscellaneous System Calls

<b>Call</b>	<b>Description</b>
<code>s = chdir(dirname)</code>	Change the working directory
<code>s = chmod(name, mode)</code>	Change a file's protection bits
<code>s = kill(pid, signal)</code>	Send a signal to a process
<code>seconds = time(&amp;seconds)</code>	Get the elapsed time since Jan. 1, 1970

Figure 1-18. Some of the major POSIX system calls.

# Memory Layout




Figure 1-20. Processes have three segments: text, data, and stack.

# Windows Win32 API

UNIX	Win32	Description
fork	CreateProcess	Create a new process
waitpid	WaitForSingleObject	Can wait for a process to exit
execve	(none)	CreateProcess = fork + execve
exit	ExitProcess	Terminate execution
open	CreateFile	Create a file or open an existing file
close	CloseHandle	Close a file
read	ReadFile	Read data from a file
write	WriteFile	Write data to a file
lseek	SetFilePointer	Move the file pointer
stat	GetFileAttributesEx	Get various file attributes
mkdir	CreateDirectory	Create a new directory
rmdir	RemoveDirectory	Remove an empty directory
link	(none)	Win32 does not support links
unlink	DeleteFile	Destroy an existing file
mount	(none)	Win32 does not support mount
umount	(none)	Win32 does not support mount
chdir	SetCurrentDirectory	Change the current working directory
chmod	(none)	Win32 does not support security (although NT does)
kill	(none)	Win32 does not support signals
time	GetLocalTime	Get the current time

Figure 1-23. The Win32 API calls that roughly correspond to the UNIX calls of Fig. 1-18.