

*Operating  
Systems:  
Internals  
and Design  
Principles*

# Chapter 2 Operating System Overview

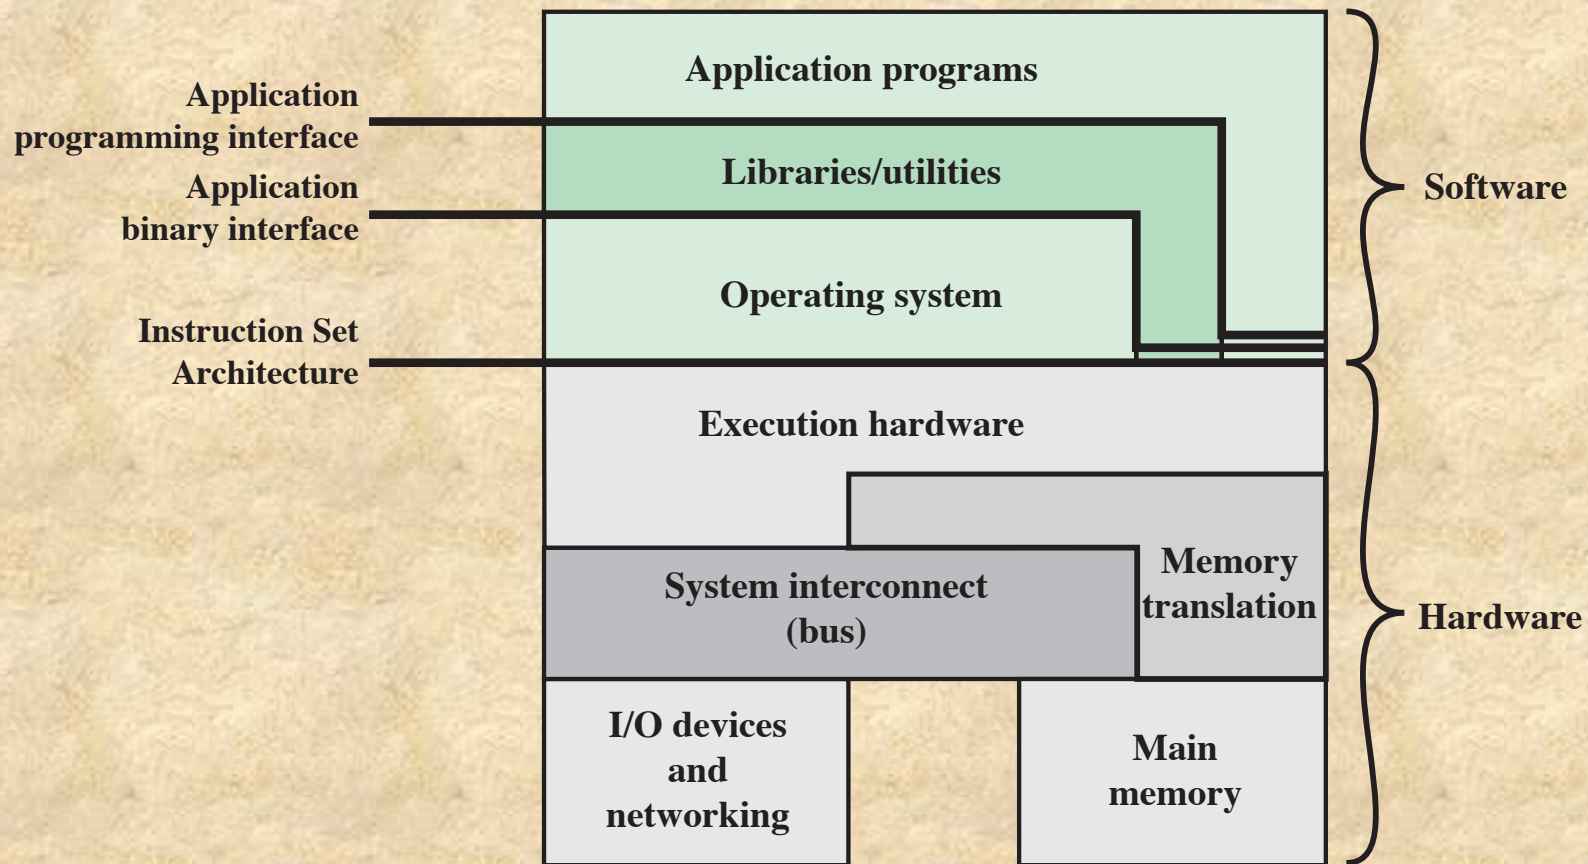
Ninth Edition  
By William Stallings

# Operating System

- A program that controls the execution of application programs
- An interface between applications and hardware

## Main objectives of an OS:

- Convenience
- Efficiency
- Ability to evolve



**Figure 2.1 Computer Hardware and Software Structure**



# Operating System Services

- Program development
- Program execution
- Access I/O devices
- Controlled access to files
- System access
- Error detection and response
- Accounting



# Key Interfaces

- Instruction set architecture (ISA)
- Application binary interface (ABI)
- Application programming interface (API)

# The Operating System as Resource Manager

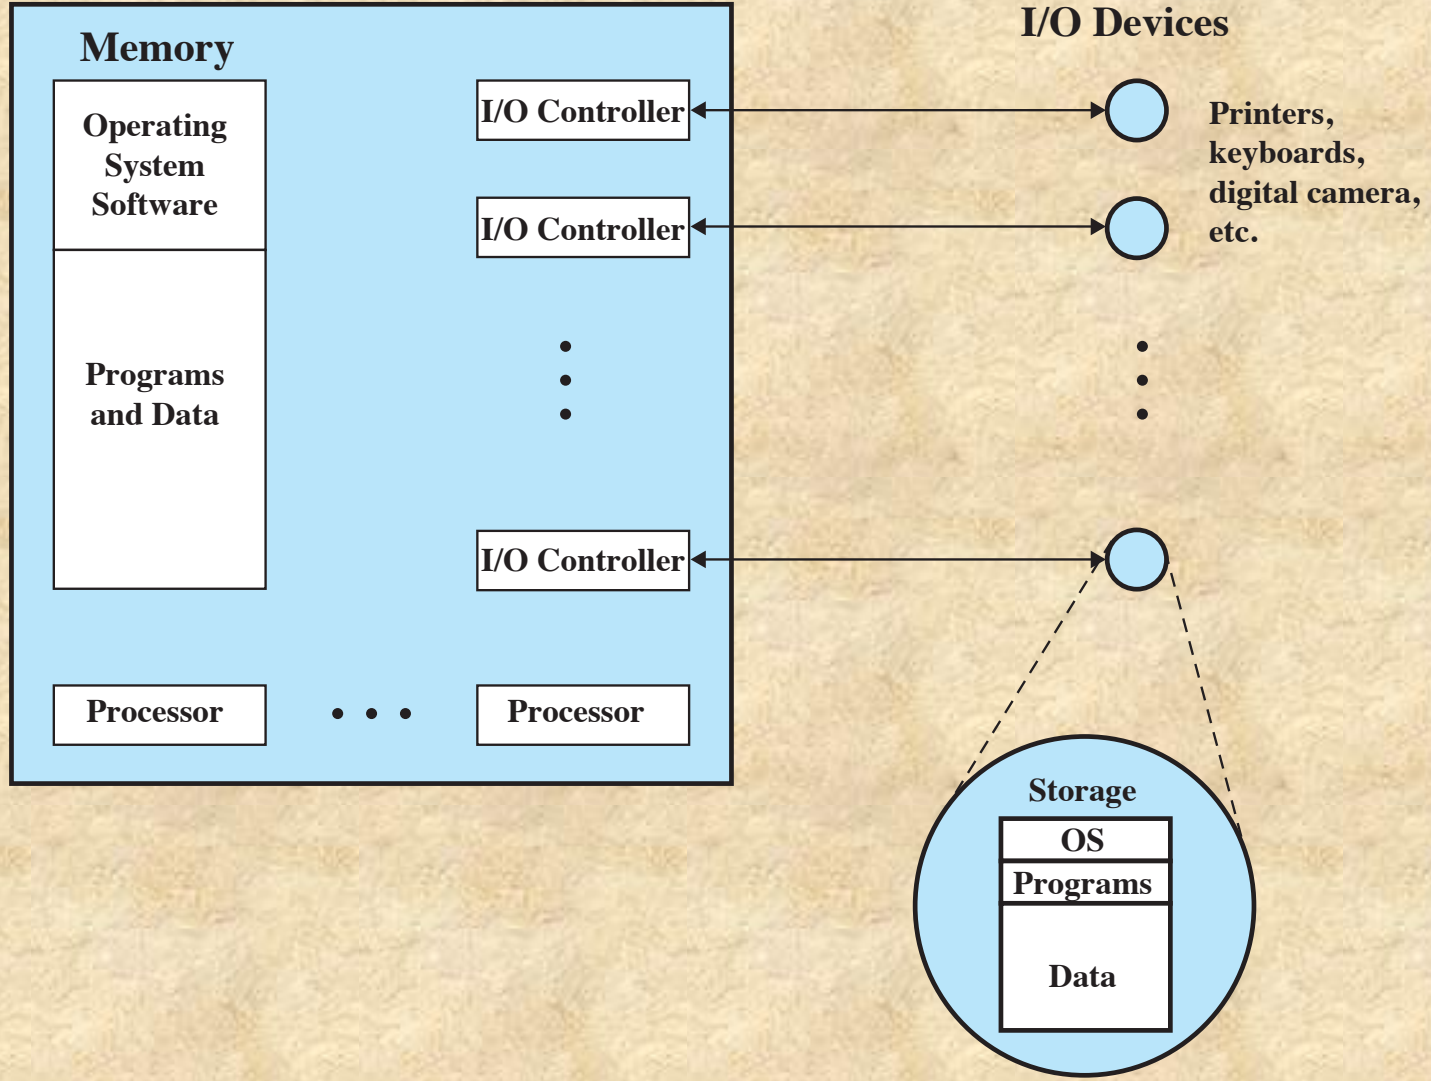
- The OS is responsible for controlling the use of a computer's resources, such as I/O, main and secondary memory, and processor execution time

# Operating System as Resource Manager

- Functions in the same way as ordinary computer software
- Program, or suite of programs, executed by the processor
- Frequently relinquishes control and must depend on the processor to allow it to regain control



# Computer System



**Figure 2.2 The Operating System as Resource Manager**

# Evolution of Operating Systems

- A major OS will evolve over time for a number of reasons:

Hardware upgrades

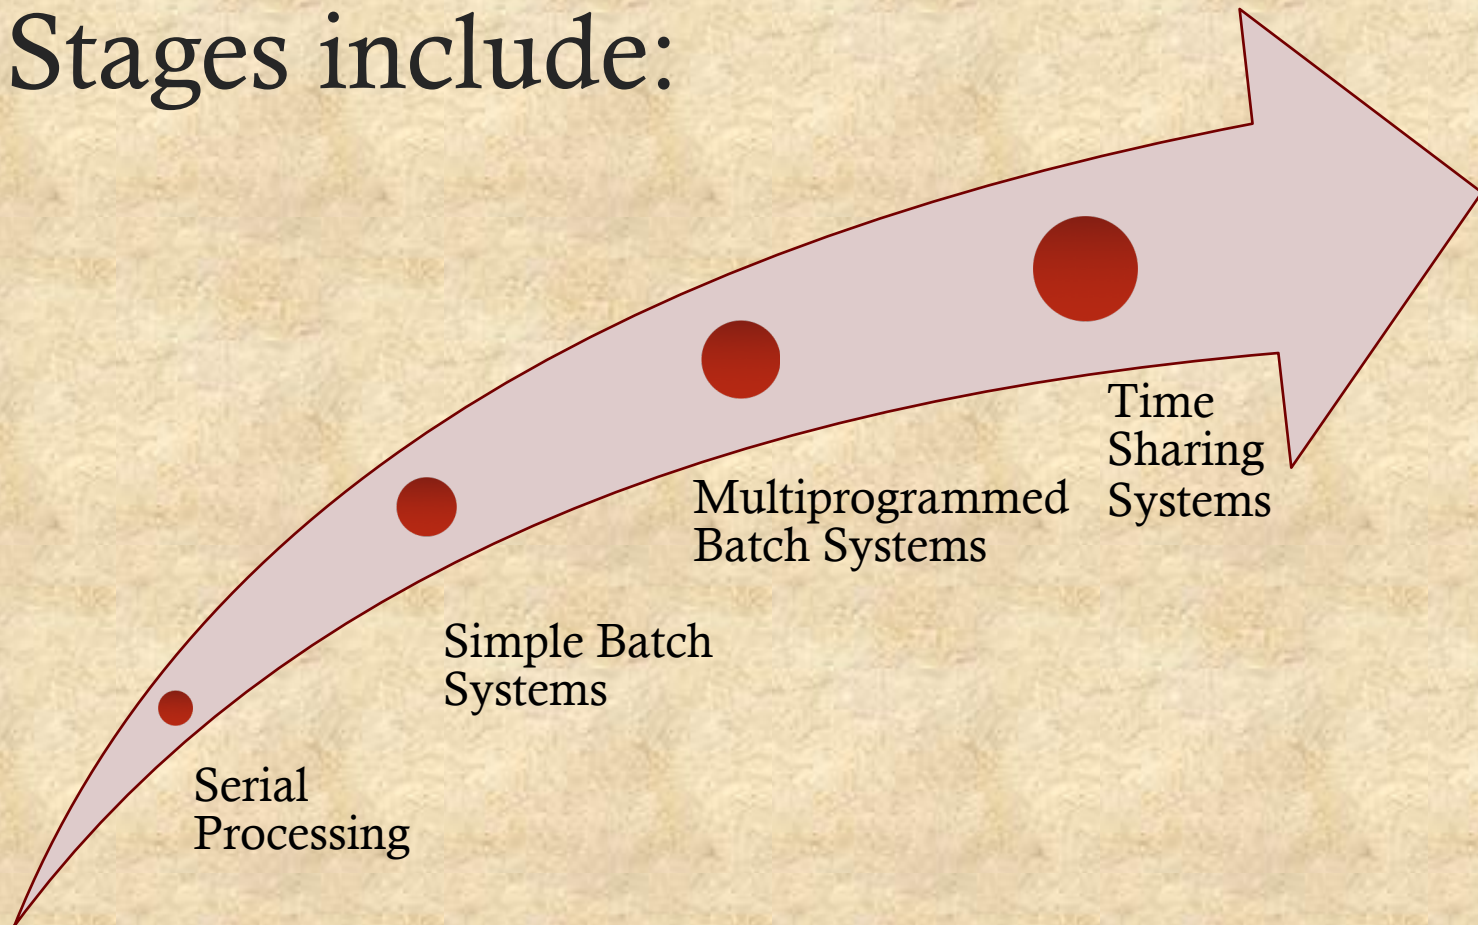
New types of hardware

New services


Fixes

# Evolution of Operating Systems

- Stages include:







# Serial Processing

## Earliest Computers:

- No operating system
  - Programmers interacted directly with the computer hardware
- Computers ran from a console with display lights, toggle switches, some form of input device, and a printer
- Users have access to the computer in “series”

## Problems:

- Scheduling:
  - Most installations used a hardcopy sign-up sheet to reserve computer time
    - Time allocations could run short or long, resulting in wasted computer time
- Setup time
  - A considerable amount of time was spent on setting up the program to run

# Big Data Processing





# Simple Batch Systems

- Early computers were very expensive
  - Important to maximize processor utilization
- Monitor
  - User no longer has direct access to processor
  - Job is submitted to computer operator who batches them together and places them on an input device
  - Program branches back to the monitor when finished



# Monitor Point of View

- Monitor controls the sequence of events
- *Resident Monitor* is software always in memory
- Monitor reads in job and gives control
- Job returns control to monitor

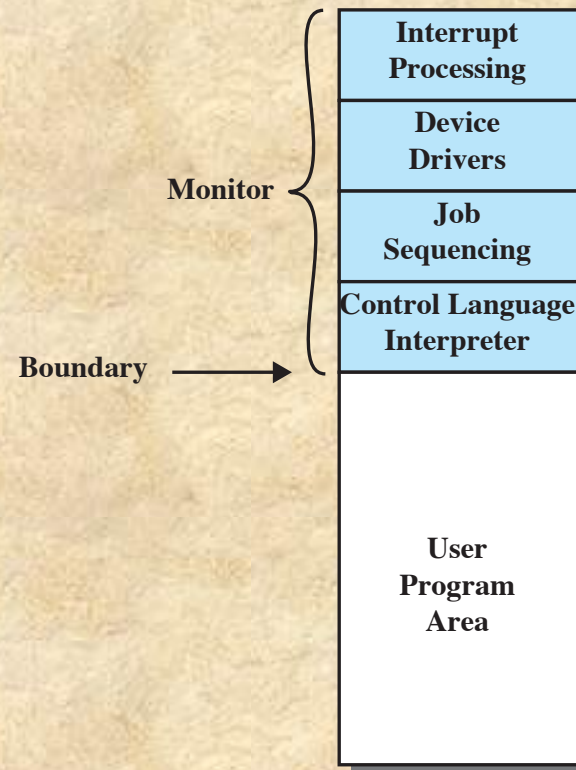


Figure 2.3 Memory Layout for a Resident Monitor

# Processor Point of View

- Processor executes instruction from the memory containing the monitor
- Executes the instructions in the user program until it encounters an ending or error condition
- “*Control is passed to a job*” means processor is fetching and executing instructions in a user program
- “*Control is returned to the monitor*” means that the processor is fetching and executing instructions from the monitor program



# Job Control Language (JCL)

Special type of programming language used to provide instructions to the monitor



What compiler to use



What data to use



# Desirable Hardware Features

## Memory protection

- While the user program is executing, it must not alter the memory area containing the monitor

## Timer

- Prevents a job from monopolizing the system

## Privileged instructions

- Can only be executed by the monitor

## Interrupts

- Gives OS more flexibility in controlling user programs

# Modes of Operation

## User Mode

- User program executes in user mode
- Certain areas of memory are protected from user access
- Certain instructions may not be executed

## Kernel Mode

- Monitor executes in kernel mode
- Privileged instructions may be executed
- Protected areas of memory may be accessed

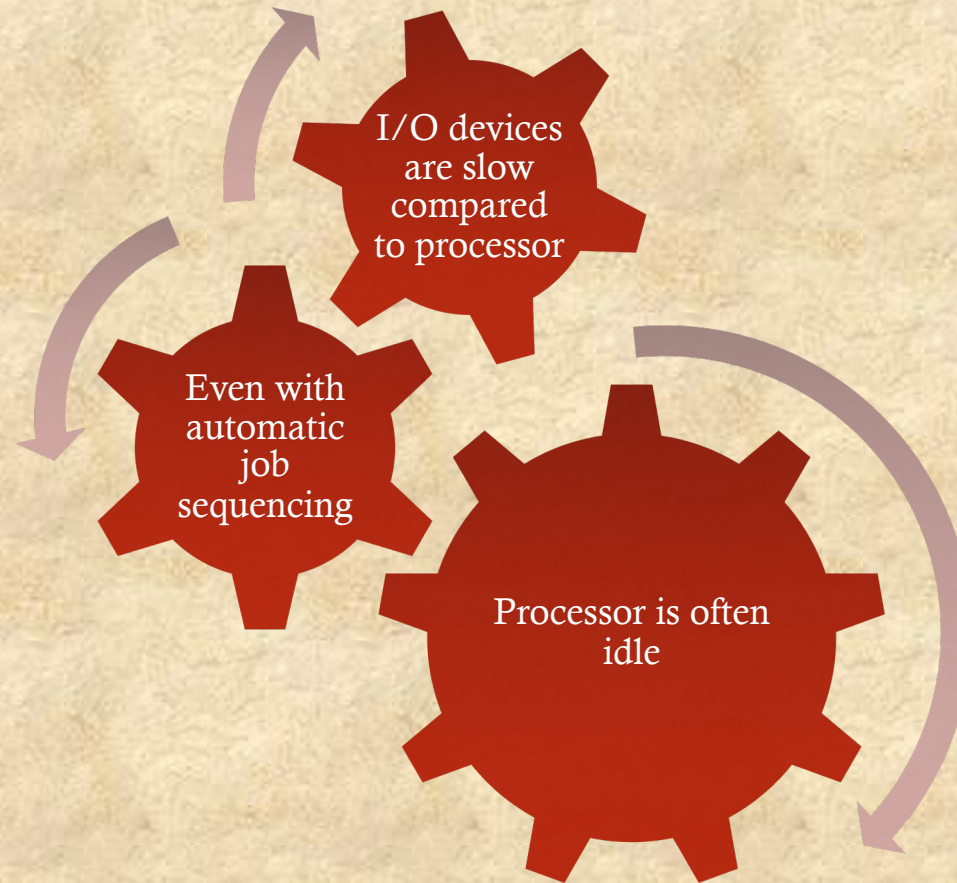


# Simple Batch System Overhead

- Processor time alternates between execution of user programs and execution of the monitor
- Sacrifices:
  - Some main memory is now given over to the monitor
  - Some processor time is consumed by the monitor
- Despite overhead, the simple batch system improves utilization of the computer



# Multiprogrammed Batch Systems



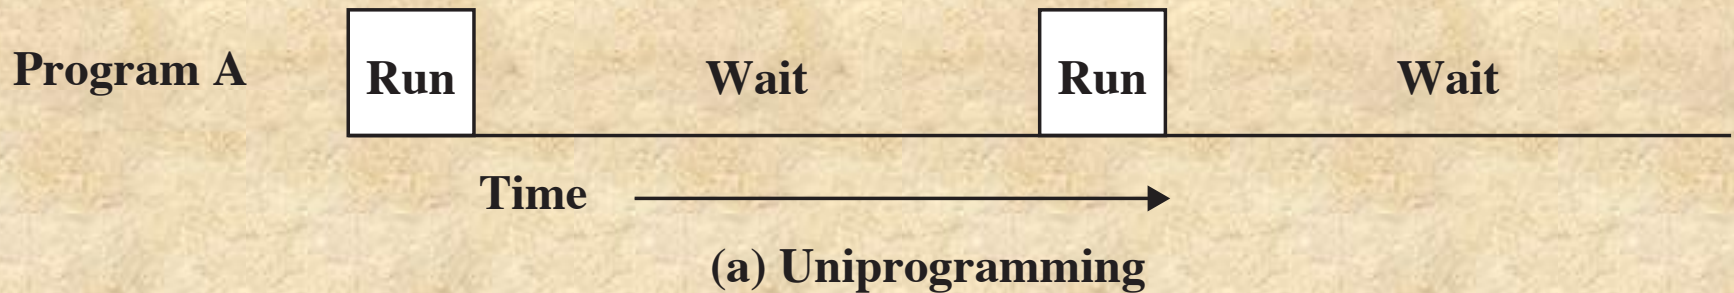
---

Read one record from file	15 $\mu$ s
Execute 100 instructions	1 $\mu$ s
Write one record to file	<u>15 <math>\mu</math>s</u>
TOTAL	31 $\mu$ s

$$\text{Percent CPU Utilization} = \frac{1}{31} = 0.032 = 3.2\%$$

**Figure 2.4 System Utilization Example**

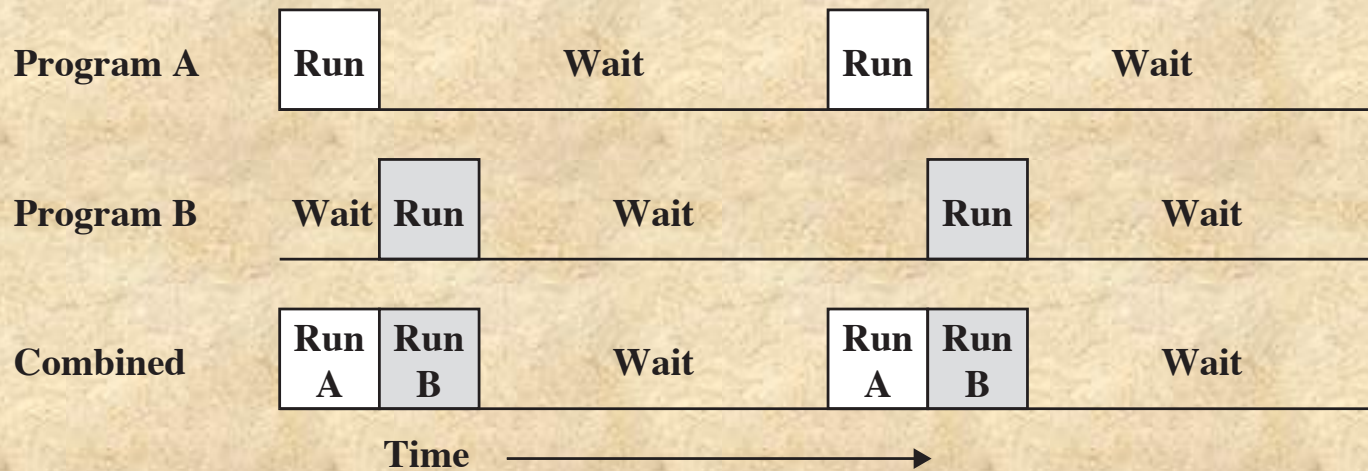
# Uniprogramming



The processor spends a certain amount of time executing, until it reaches an I/O instruction; it must then wait until that I/O instruction concludes before proceeding



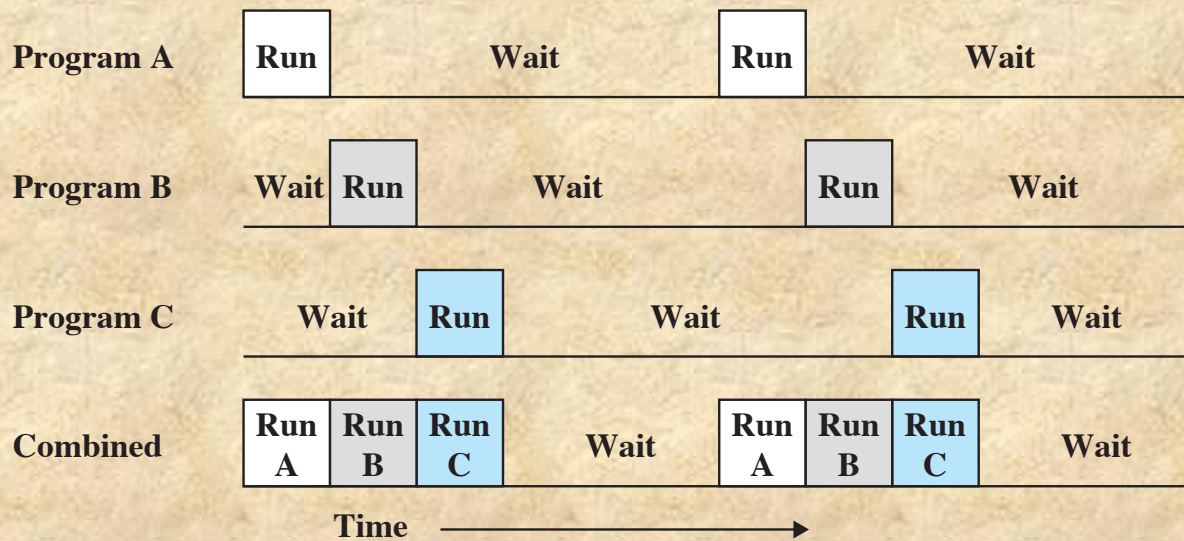
# Multiprogramming



(b) Multiprogramming with two programs

- There must be enough memory to hold the OS (resident monitor) and one user program
- When one job needs to wait for I/O, the processor can switch to the other job, which is likely not waiting for I/O

# Multiprogramming



(c) Multiprogramming with three programs

- Also known as multitasking
- Memory is expanded to hold three, four, or more programs and switch among all of them



# Time-Sharing Systems

- Can be used to handle multiple interactive jobs
- Processor time is shared among multiple users
- Multiple users simultaneously access the system through terminals, with the OS interleaving the execution of each user program in a short burst or quantum of computation



	<b>Batch Multiprogramming</b>	<b>Time Sharing</b>
Principal objective	Maximize processor use	Minimize response time
Source of directives to operating system	Job control language commands provided with the job	Commands entered at the terminal

**Table 2.3 Batch Multiprogramming versus Time Sharing**

# Compatible Time-Sharing System (CTSS)

- One of the first time-sharing operating systems
- Developed at MIT by a group known as Project MAC
- The system was first developed for the IBM 709 in 1961
- Ran on a computer with 32,000 36-bit words of main memory, with the resident monitor consuming 5000 of that
- Utilized a technique known as *time slicing*
  - System clock generated interrupts at a rate of approximately one every 0.2 seconds
  - At each clock interrupt the OS regained control and could assign the processor to another user
  - Thus, at regular time intervals the current user would be preempted and another user loaded in
  - To preserve the old user program status for later resumption, the old user programs and data were written out to disk before the new user programs and data were read in
  - Old user program code and data were restored in main memory when that program was next given a turn

# Major Achievements

- Operating Systems are among the most complex pieces of software ever developed
- Major advances in development include:
  - Processes
  - Memory management
  - Information protection and security
  - Scheduling and resource management
  - System structure



# Process

- Fundamental to the structure of operating systems

A *process* can be defined as:

A program in execution

An instance of a running program

The entity that can be assigned to, and executed on, a processor

A unit of activity characterized by a single sequential thread of execution, a current state, and an associated set of system resources

# Causes of Errors

## ■ Improper synchronization

- It is often the case that a routine must be suspended awaiting an event elsewhere in the system
- Improper design of the signaling mechanism can result in loss or duplication

## ■ Failed mutual exclusion

- More than one user or program attempts to make use of a shared resource at the same time
- There must be some sort of mutual exclusion mechanism that permits only one routine at a time to perform an update against the file

## ■ Nondeterminate program operation

- When programs share memory, and their execution is interleaved by the processor, they may interfere with each other by overwriting common memory areas in unpredictable ways
- The order in which programs are scheduled may affect the outcome of any particular program

## ■ Deadlocks

- It is possible for two or more programs to be hung up waiting for each other



# Components of a Process

- A process contains three components:
  - An executable program
  - The associated data needed by the program (variables, work space, buffers, etc.)
  - The execution context (or “process state”) of the program
- The execution context is essential:
  - It is the internal data by which the OS is able to supervise and control the process
  - Includes the contents of the various process registers
  - Includes information such as the priority of the process and whether the process is waiting for the completion of a particular I/O event



# Process Management

- The entire state of the process at any instant is contained in its context
- New features can be designed and incorporated into the OS by expanding the context to include any new information needed to support the feature

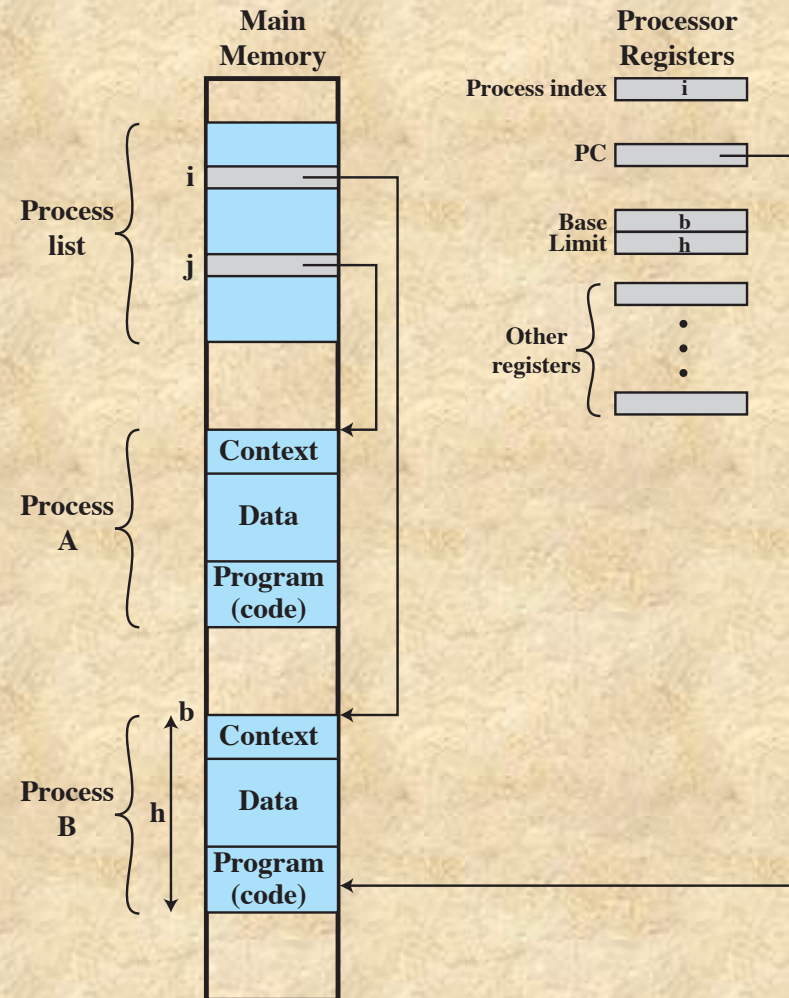


Figure 2.8 Typical Process Implementation

# Memory Management

- The OS has five principal storage management responsibilities:

Process  
isolation

Automatic  
allocation  
and  
management

Support of  
modular  
programming

Protection  
and access  
control

Long-term  
storage

# Virtual Memory

- A facility that allows programs to address memory from a logical point of view, without regard to the amount of main memory physically available
- Conceived to meet the requirement of having multiple user jobs reside in main memory concurrently



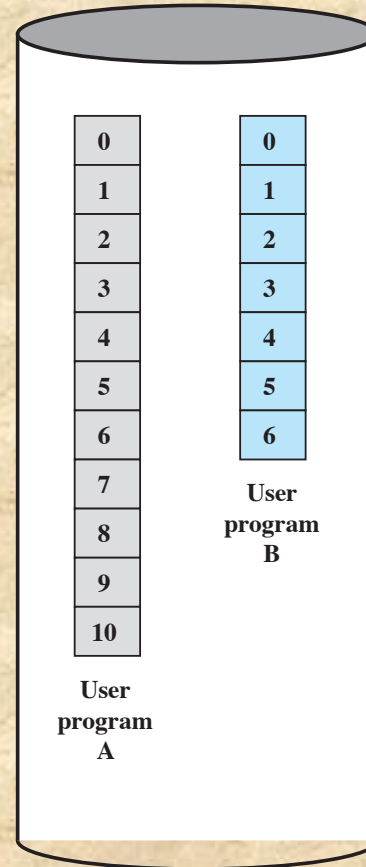
# Paging

- Allows processes to be comprised of a number of fixed-size blocks, called pages
- Program references a word by means of a *virtual address*, consisting of a page number and an offset within the page
- Each page of a process may be located anywhere in main memory
- The paging system provides for a dynamic mapping between the virtual address used in the program and a *real address* (or physical address) in main memory

A.1			
	A.0	A.2	
	A.5		
B.0	B.1	B.2	B.3
		A.7	
	A.9		
		A.8	
	B.5	B.6	

**Main Memory**


Main memory consists of a number of fixed-length frames, each equal to the size of a page. For a program to execute, some or all of its pages must be in main memory.



**Disk**

Secondary memory (disk) can hold many fixed-length pages. A user program consists of some number of pages. Pages for all programs plus the operating system are on disk, as are files.

**Figure 2.9 Virtual Memory Concepts**




**Figure 2.10 Virtual Memory Addressing**



# Information Protection and Security

- The nature of the threat that concerns an organization will vary greatly depending on the circumstances
- The problem involves controlling access to computer systems and the information stored in them



# Different Architectural Approaches

- Demands on operating systems require new ways of organizing the OS

Different approaches and design elements have been tried:

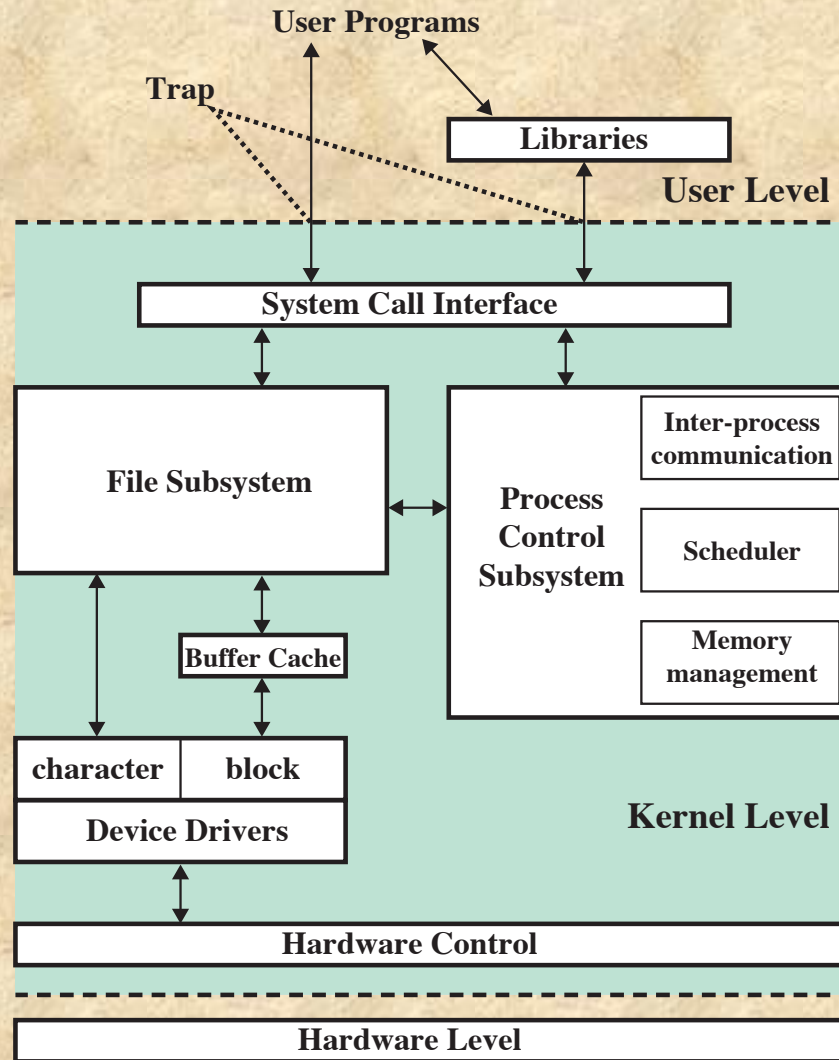
- Microkernel architecture
- Multithreading
- Symmetric multiprocessing
- Distributed operating systems
- Object-oriented design



# Traditional UNIX Systems

- Developed at Bell Labs and became operational on a PDP-7 in 1970
- The first notable milestone was porting the UNIX system from the PDP-7 to the PDP-11
  - First showed that UNIX would be an OS for all computers
- Next milestone was rewriting UNIX in the programming language C
  - Demonstrated the advantages of using a high-level language for system code
- Was described in a technical journal for the first time in 1974
- First widely available version outside Bell Labs was Version 6 in 1976
- Version 7, released in 1978, is the ancestor of most modern UNIX systems
- Most important of the non-AT&T systems was UNIX BSD (Berkeley Software Distribution), running first on PDP and then on VAX computers





**Figure 2.15 Traditional UNIX Kernel**

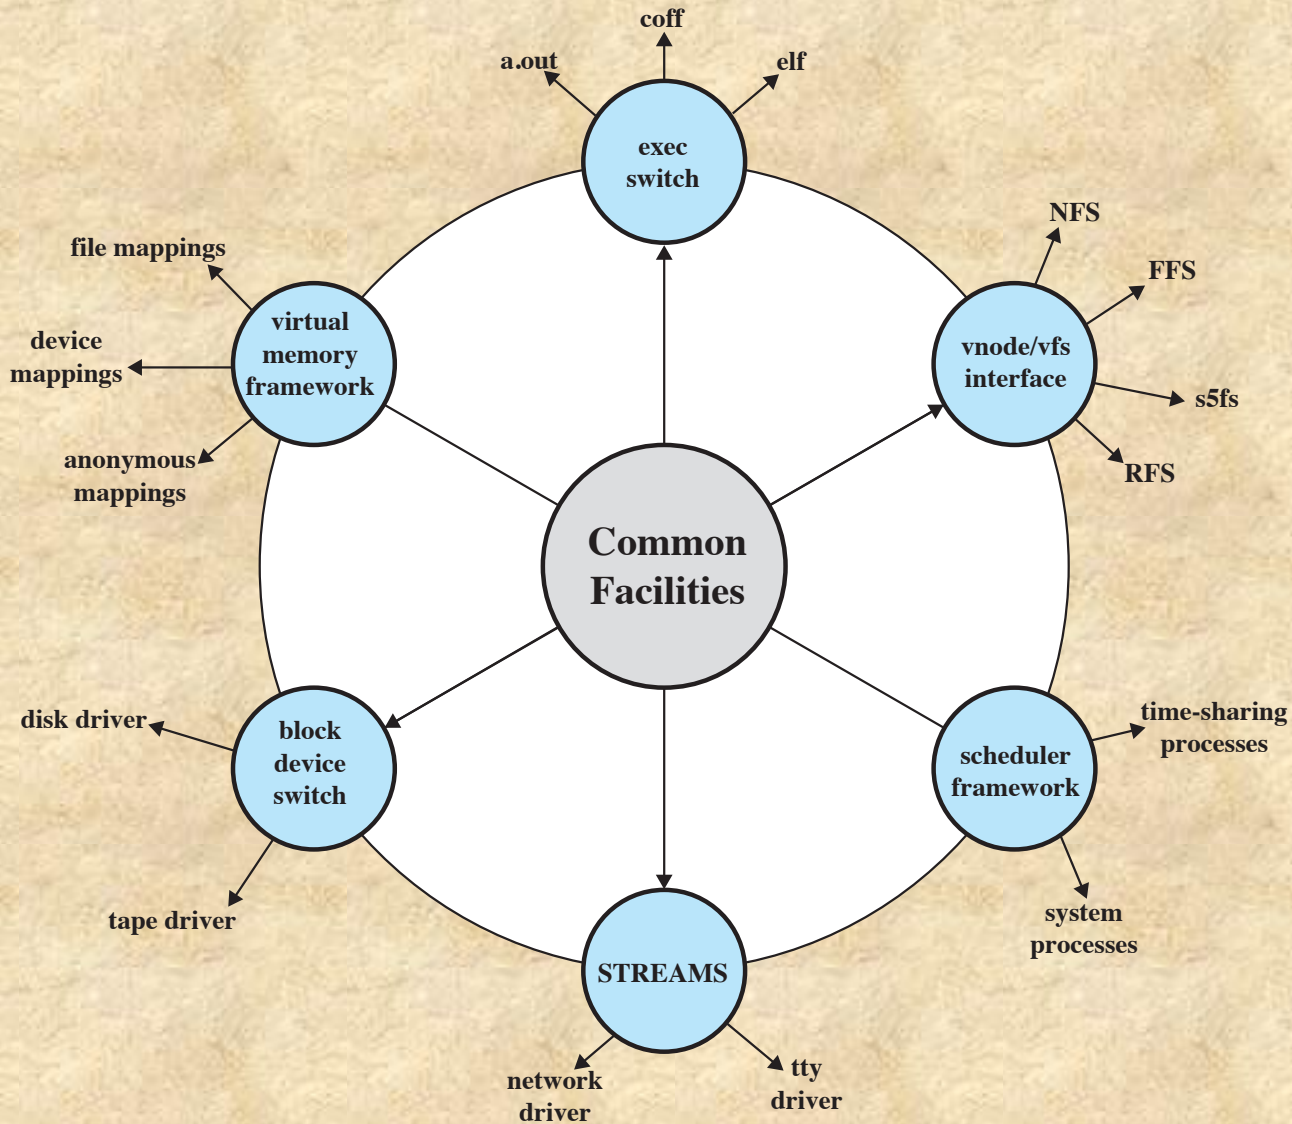
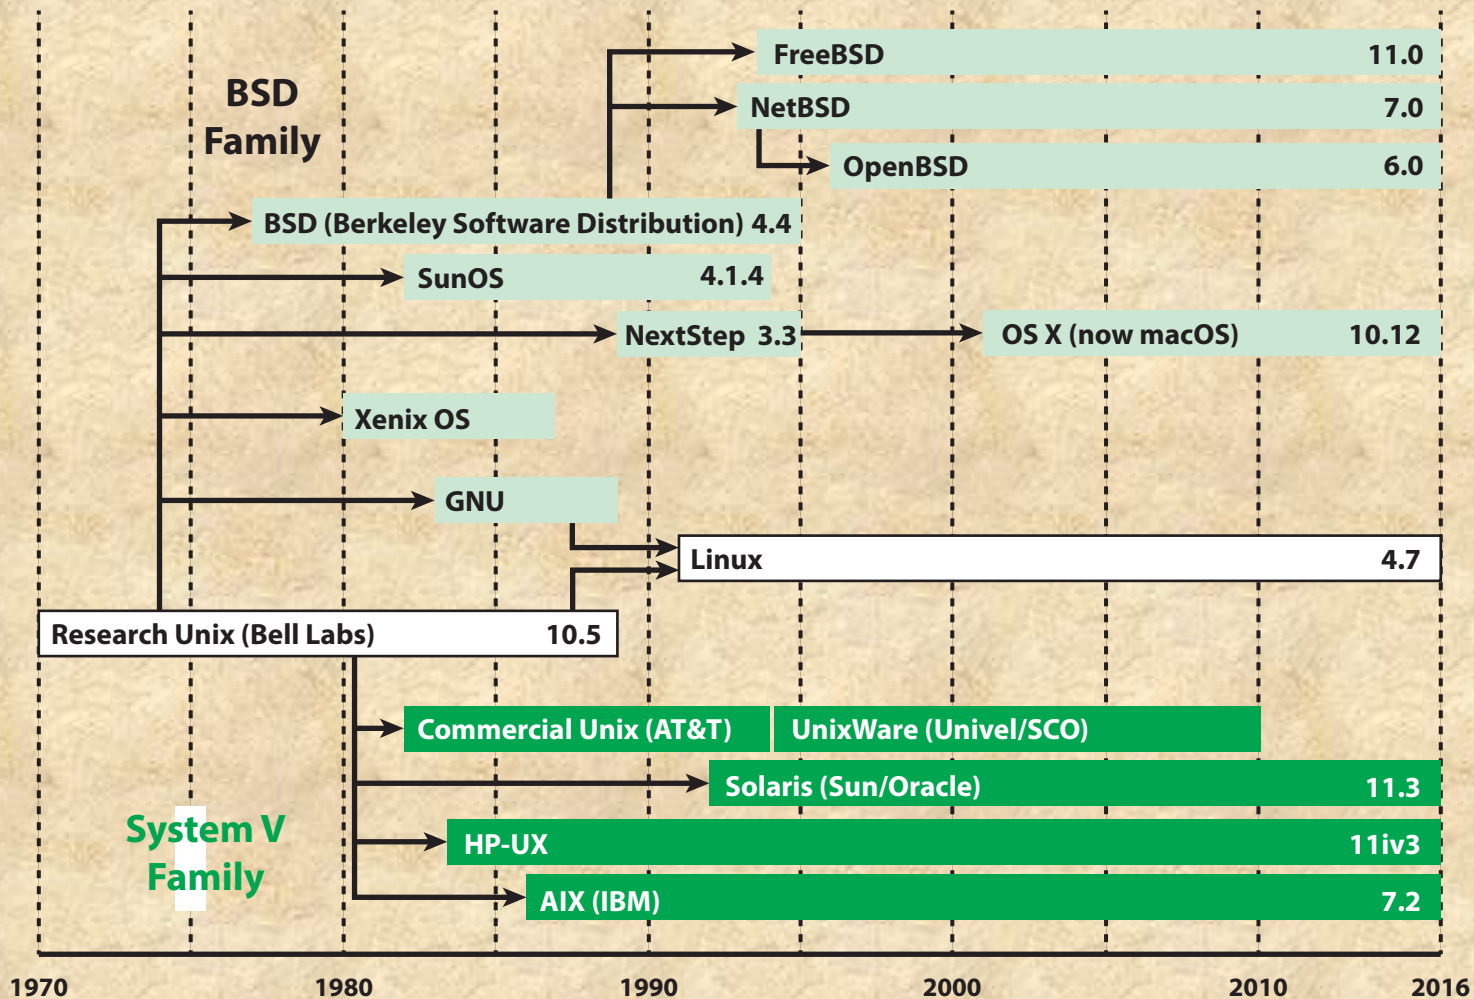


Figure 2.16 Modern UNIX Kernel



**Figure 2.17 Unix Family Tree**



# System V Release 4 (SVR4)

- Developed jointly by AT&T and Sun Microsystems
- Combines features from SVR3, 4.3BSD, Microsoft Xenix System V, and SunOS
- New features in the release include:
  - Real-time processing support
  - Process scheduling classes
  - Dynamically allocated data structures
  - Virtual memory management
  - Virtual file system
  - Preemptive kernel

# BSD

- Berkeley Software Distribution
- 4.xBSD is widely used in academic installations and has served as the basis of a number of commercial UNIX products
- 4.4BSD was the final version of BSD to be released by Berkeley
- There are several widely used, open-source versions of BSD
  - FreeBSD
    - Popular for Internet-based servers and firewalls
    - Used in a number of embedded systems
  - NetBSD
    - Available for many platforms
    - Often used in embedded systems
  - OpenBSD
    - An open-source OS that places special emphasis on security



# Solaris 11

- Oracle's SVR4-based UNIX release
- Provides all of the features of SVR4 plus a number of more advanced features such as:
  - A fully preemptable, multithreaded kernel
  - Full support for SMP
  - An object-oriented interface to file systems



# LINUX Overview

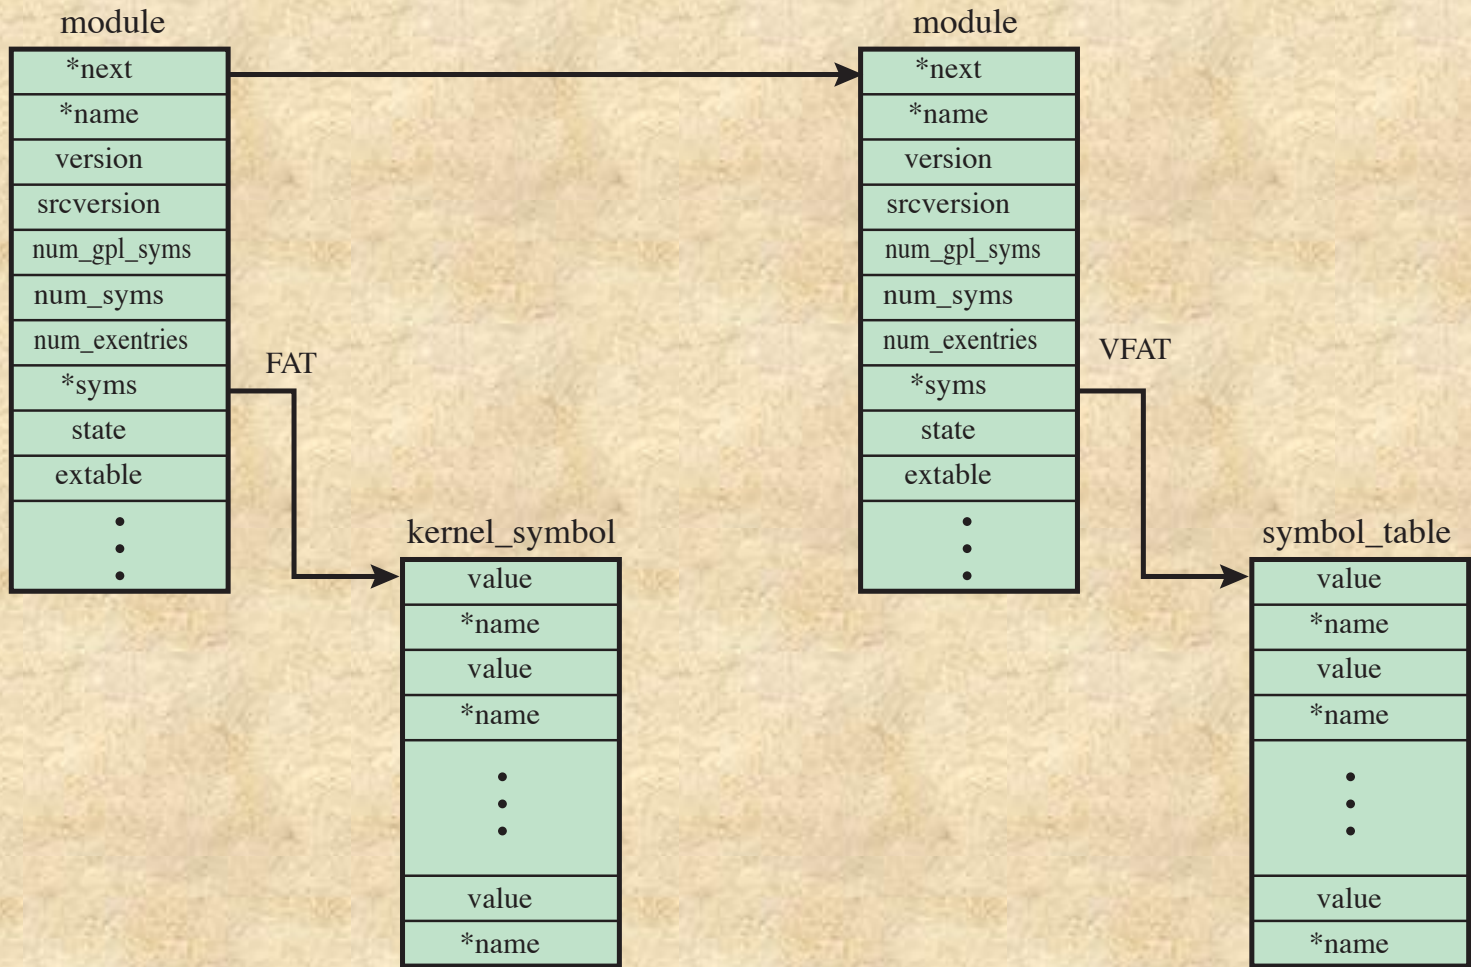
- Started out as a UNIX variant for the IBM PC
- Linus Torvalds, a Finnish student of computer science, wrote the initial version
- Linux was first posted on the Internet in 1991
- Today it is a full-featured UNIX system that runs on virtually all platforms
- Is free and the source code is available
- Key to the success of Linux has been the availability of free software packages under the auspices of the Free Software Foundation (FSF)
- Highly modular and easily configured

# Modular Structure

- Linux development is global and done by a loosely associated group of independent developers
- Although Linux does not use a microkernel approach, it achieves many of the potential advantages of the approach by means of its particular modular architecture
- Linux is structured as a collection of modules, a number of which can be automatically loaded and unloaded on demand

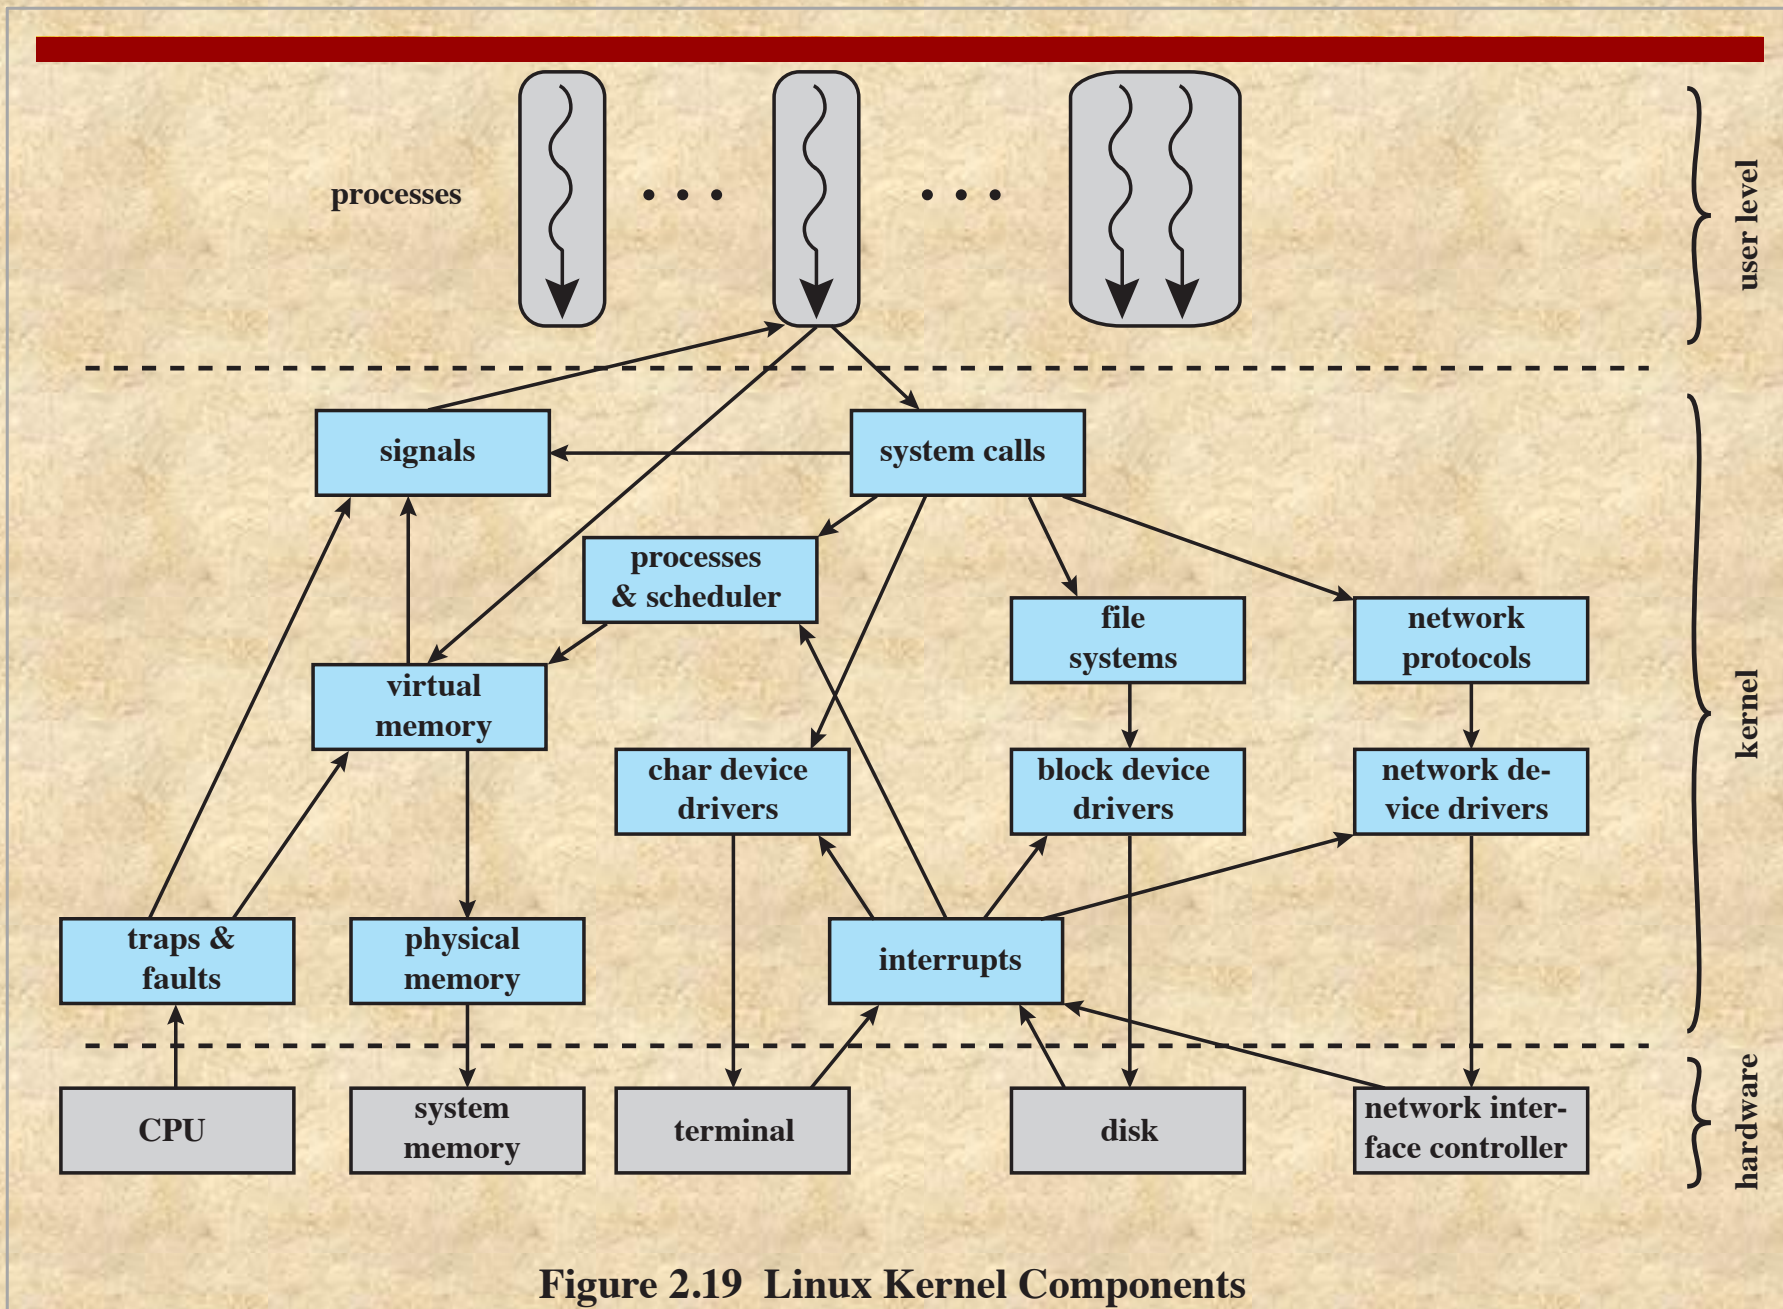
## Loadable Modules

- Relatively independent blocks
- A module is an object file whose code can be linked to and unlinked from the kernel at runtime
- A module is executed in kernel mode on behalf of the current process
- Have two important characteristics:
  - Dynamic linking
  - Stackable modules



**Figure 2.18 Example List of Linux Kernel Modules**





**Figure 2.19 Linux Kernel Components**

# Linux Signals

SIGHUP	Terminal hangup	SIGCONT	Continue
SIGQUIT	Keyboard quit	SIGTSTP	Keyboard stop
SIGTRAP	Trace trap	SIGTTOU	Terminal write
SIGBUS	Bus error	SIGXCPU	CPU limit exceeded
SIGKILL	Kill signal	SIGVTALRM	Virtual alarm clock
SIGSEGV	Segmentation violation	SIGWINCH	Window size unchanged
SIGPIPT	Broken pipe	SIGPWR	Power failure
SIGTERM	Termination	SIGRTMIN	First real-time signal
SIGCHLD	Child status unchanged	SIGRTMAX	Last real-time signal

**Table 2.6 Some Linux Signals**

Filesystem related	
<b>close</b>	Close a file descriptor.
<b>link</b>	Make a new name for a file.
<b>open</b>	Open and possibly create a file or device.
<b>read</b>	Read from file descriptor.
<b>write</b>	Write to file descriptor
Process related	
<b>execve</b>	Execute program.
<b>exit</b>	Terminate the calling process.
<b>getpid</b>	Get process identification.
<b>setuid</b>	Set user identity of the current process.
<b>ptrace</b>	Provides a means by which a parent process may observe and control the execution of another process, and examine and change its core image and registers.
Scheduling related	
<b>sched_getparam</b>	Sets the scheduling parameters associated with the scheduling policy for the process identified by <code>pid</code> .
<b>sched_get_priority_max</b>	Returns the maximum priority value that can be used with the scheduling algorithm identified by <code>policy</code> .
<b>sched_setscheduler</b>	Sets both the scheduling policy (e.g., FIFO) and the associated parameters for the process <code>pid</code> .
<b>sched_rr_get_interval</b>	Writes into the <code>timespec</code> structure pointed to by the parameter <code>tp</code> the round robin time quantum for the process <code>pid</code> .
<b>sched_yield</b>	A process can relinquish the processor voluntarily without blocking via this system call. The process will then be moved to the end of the queue for its static priority and a new process gets to run.

**Table 2.7 Some Linux System Calls (page 1 of 2)**



<b>Interprocess Communication (IPC) related</b>	
<b>msgrcv</b>	A message buffer structure is allocated to receive a message. The system call then reads a message from the message queue specified by msqid into the newly created message buffer.
<b>semctl</b>	Performs the control operation specified by cmd on the semaphore set semid.
<b>semop</b>	Performs operations on selected members of the semaphore set semid.
<b>shmat</b>	Attaches the shared memory segment identified by shmid to the data segment of the calling process.
<b>shmctl</b>	Allows the user to receive information on a shared memory segment, set the owner, group, and permissions of a shared memory segment, or destroy a segment.
<b>Socket (networking) related</b>	
<b>bind</b>	Assigns the local IP address and port for a socket. Returns 0 for success and -1 for error.
<b>connect</b>	Establishes a connection between the given socket and the remote socket associated with sockaddr.
<b>gethostname</b>	Returns local host name.
<b>send</b>	Send the bytes contained in buffer pointed to by *msg over the given socket.
<b>setsockopt</b>	Sets the options on a socket
<b>Miscellaneous</b>	
<b>fsync</b>	Copies all in-core parts of a file to disk, and waits until the device reports that all parts are on stable storage.
<b>time</b>	Returns the time in seconds since January 1, 1970.
<b>vhangup</b>	Simulates a hangup on the current terminal. This call arranges for other users to have a "clean" tty at login time.

**Table 2.7 Some Linux System Calls (page 2 of 2)**

# Summary

- Operating system objectives and functions
  - User/computer interface
  - Resource manager
- Evolution of operating systems
  - Serial processing
  - Simple/multiprogrammed/time-sharing batch systems
- Major achievements
- Developments leading to modern operating systems
- Fault tolerance
  - Fundamental concepts
  - Faults
  - OS mechanisms
- OS design considerations for multiprocessor and multicore
- Microsoft Windows overview
- Traditional Unix systems
  - History/description
- Modern Unix systems
  - System V Release 4 (SVR4)
  - BSD
  - Solaris 10
- Linux
  - History
  - Modular structure
  - Kernel components
- Android
  - Software/system architecture
  - Activities
  - Power management