

cse

# Operating System Overview

Chapter 1.5 – 1.9

THE UNIVERSITY OF NEW SOUTH WALES

1

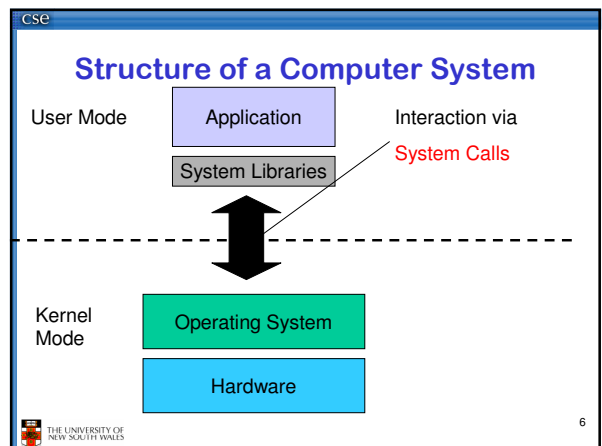
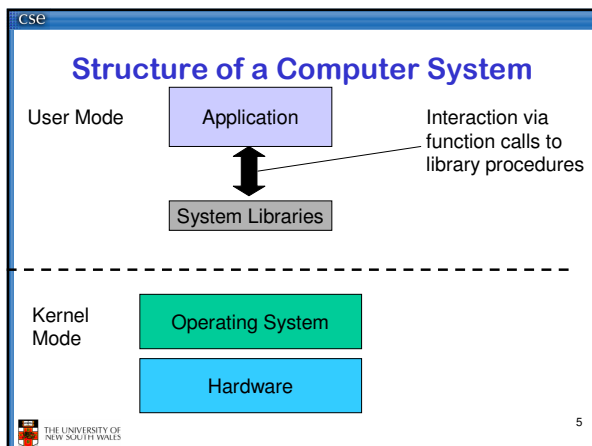
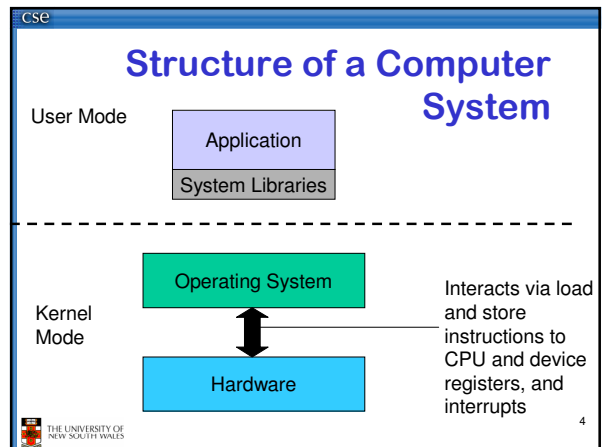
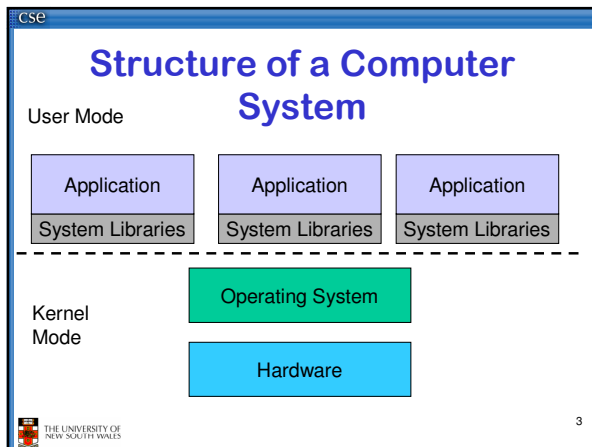
cse

# Operating System

- A program that controls execution of applications
  - The resource manager
- An interface between applications and hardware
  - The extended machine

THE UNIVERSITY OF NEW SOUTH WALES

2



cse

## A note on System Libraries

- System libraries are just that, libraries of support functions (procedures, subroutines)
  - Only a subset of library functions are actually systems calls
    - strcmp(), memcpy(), are pure library functions
    - open(), close(), read(), write() are system calls
  - System call functions are in the library for convenience

THE UNIVERSITY OF NEW SOUTH WALES 7

cse

## Operating System Objectives

- Convenience
  - Make the computer more convenient to use
- Abstraction
  - Hardware-independent programming model
- Efficiency
  - Allows the computer system to be used in an efficient manner
- Ability to evolve
  - Permit effective development, testing, and introduction of new system functions without interfering with existing services
- Protection

THE UNIVERSITY OF NEW SOUTH WALES 8

cse

## Services Provided by the Operating System

- Program development
  - Editors, compilers, debuggers
    - Not so much these days
- Program execution
  - Load a program and its data
- Access to I/O devices
- Controlled access to files
  - Access protection
- System access
  - User authentication

THE UNIVERSITY OF NEW SOUTH WALES 9

cse

## Services Provided by the Operating System

- Error detection and response
  - internal and external hardware errors
    - memory error
    - device failure
  - software errors
    - arithmetic overflow
    - access forbidden memory locations
  - operating system cannot grant request of application

THE UNIVERSITY OF NEW SOUTH WALES 10

cse

## Services Provided by the Operating System

- Accounting
  - collect statistics
  - monitor performance
  - used to anticipate future enhancements
  - used for billing users

THE UNIVERSITY OF NEW SOUTH WALES 11

cse

## Operating System Software

- Fundamentally, OS functions the same way as ordinary computer software
  - It is a program that is executed (just like apps)
  - It has more privileges
- Operating system relinquishes control of the processor to execute other programs
  - Reestablishes control after
    - System calls
    - Interrupts (especially timer interrupts)

THE UNIVERSITY OF NEW SOUTH WALES 12

cse

## Kernel

- Portion of the operating system that is running in *privileged mode*
- Usually resident in main memory
- Contains fundamental functionality
  - Whatever is required to implement other services
  - Whatever is required to provide security
- Contains most-frequently used functions
- Also called the nucleus or supervisor

13

THE UNIVERSITY OF NEW SOUTH WALES

cse

## Major OS Concepts

- Processes
- Concurrency and deadlocks
- Memory management
- Files
- Information Security and Protection
- Scheduling and resource management

14

THE UNIVERSITY OF NEW SOUTH WALES

cse

## Processes

- A program in execution
- An instance of a program running on a computer
- The entity that can be assigned to and executed on a processor
- A unit of resource ownership
- A unit of activity characterized by a single sequential thread of execution, a current state, and an associated set of system resources
  - Nowadays the execution abstraction is separated out: *Thread*
  - Single process can contain many threads

15

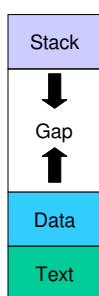
THE UNIVERSITY OF NEW SOUTH WALES

cse

## Process

Memory

- Consist of three segments
  - Text
    - contains the code (instructions)
  - Data
    - Global variables
  - Stack
    - Activation records of procedure
    - Local variables
- Note:
  - data can dynamically grow up
  - The stack can dynamically grow down



16

THE UNIVERSITY OF NEW SOUTH WALES

cse

## Process

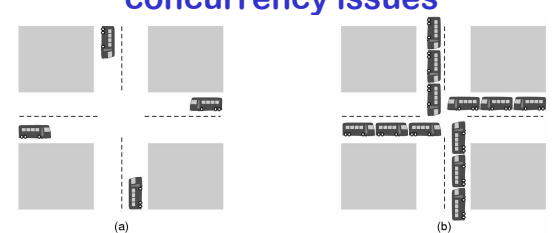
- Consists of three components
  - An executable program
    - text
  - Associated data needed by the program
    - Data and stack
  - Execution context of the program
    - All information the operating system needs to manage the process
      - Registers, program counter, stack pointer, etc...
    - A multithread program has a stack and execution context for each thread

17

THE UNIVERSITY OF NEW SOUTH WALES

cse

## Multiple processes creates concurrency issues



(a) (b)

(a) A potential deadlock. (b) an actual deadlock.

18

THE UNIVERSITY OF NEW SOUTH WALES

cse

## Memory Management

- The view from thirty thousand feet
  - Process isolation
    - Prevent processes from accessing each others data
  - Automatic allocation and management
    - Don't want users to deal with physical memory directly
  - Support for modular programming
  - Protection and access control
    - Still want controlled sharing
  - Long-term storage
  - OS services
    - Virtual memory
    - File system

THE UNIVERSITY OF NEW SOUTH WALES 19

cse

## Virtual Memory

- Allows programmers to address memory from a logical point of view
  - Gives apps the illusion of having RAM to themselves
  - Logical addresses are independent of other processes
  - Provides isolation of processes from each other
- Can overlap execution of one process while swapping in/out others.

THE UNIVERSITY OF NEW SOUTH WALES 20

cse

## Virtual Memory Addressing

The diagram shows a Processor sending a Virtual Address to a Memory Management Unit. The Memory Management Unit then sends a Real Address to Main Memory and a Disk Address to Secondary Memory.

Figure 2.10 Virtual Memory Addressing

THE UNIVERSITY OF NEW SOUTH WALES 21

cse

## Paging

- Allows process to be comprised of a number of fixed-size blocks, called pages
- Virtual address is a page number and an offset within the page
- Each page may be located any where in main memory
- A page may actually exist only on disk

THE UNIVERSITY OF NEW SOUTH WALES 22

The diagram illustrates Main Memory as a grid of frames (A.1 to B.6) and Disk as a cylinder with pages for User program A (0-10) and User program B (0-6).

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

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

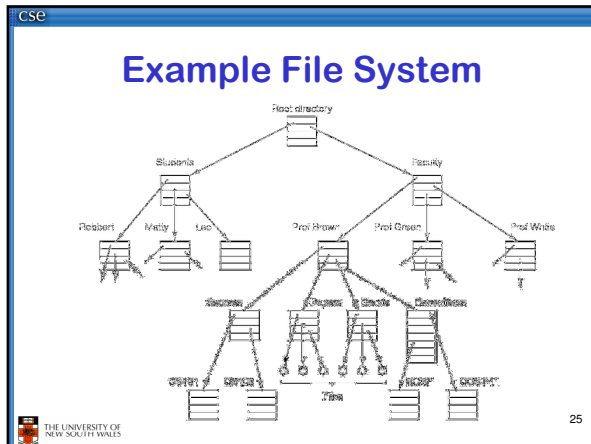
THE UNIVERSITY OF NEW SOUTH WALES 24

cse

## File System

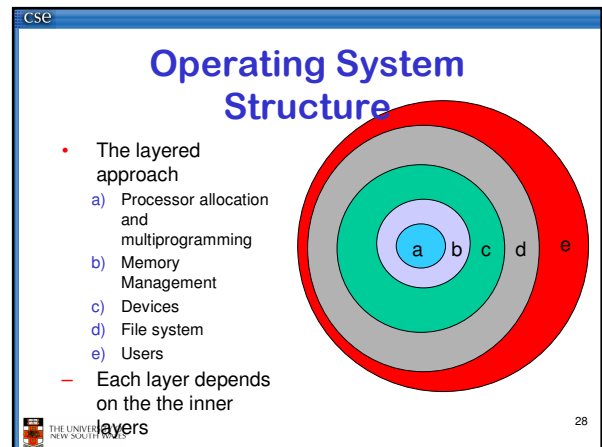
- Implements long-term store
- Information stored in named objects called files

THE UNIVERSITY OF NEW SOUTH WALES 24

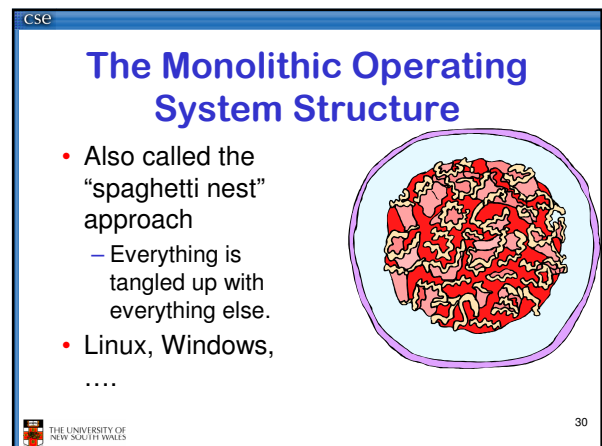


- cse
- ## Information Protection and Security
- Access control
    - regulate user access to the system
    - Involves authentication
  - Information flow control
    - regulate flow of data within the system and its delivery to users
- 26
- THE UNIVERSITY OF NEW SOUTH WALES

- cse
- ## Scheduling and Resource Management
- Fairness
    - give equal and fair access to all processes
  - Differential responsiveness
    - discriminate between different classes of jobs
  - Efficiency
    - maximize throughput, minimize response time, and accommodate as many uses as possible
- 27
- THE UNIVERSITY OF NEW SOUTH WALES



- cse
- ## Operating System Structure
- In practice, layering is only a guide
    - Operating Systems have many interdependencies
      - Scheduling on virtual memory
      - Virtual memory on I/O to disk
      - VM on files (page to file)
      - Files on VM (memory mapped files)
      - And many more...
- 29
- THE UNIVERSITY OF NEW SOUTH WALES



CSE

## The Monolithic Operating System Structure

- However, some reasonable structure usually prevails

31

CSE

## UNIX

- Provides a good hardware abstraction
  - Everything is a file (mostly)
- Runs on most hardware
- Comes with a number of user services and interfaces
  - shell
  - C compiler

32

CSE

## Traditional UNIX Structure

Figure 2.15 General UNIX Architecture

33

CSE

## Traditional UNIX Kernel

34

CSE

## Microkernel-based Systems

- Assigns only a few essential functions to the kernel
  - Address space
  - Interprocess Communication (IPC)
  - Basic scheduling
  - Minimal hardware abstraction
- Other services implemented by user-level servers
- Traditional “system calls” become IPC requests to servers
- Extreme view of a microkernel
  - A feature is only allowed in the kernel if required for security

35

CSE

36

