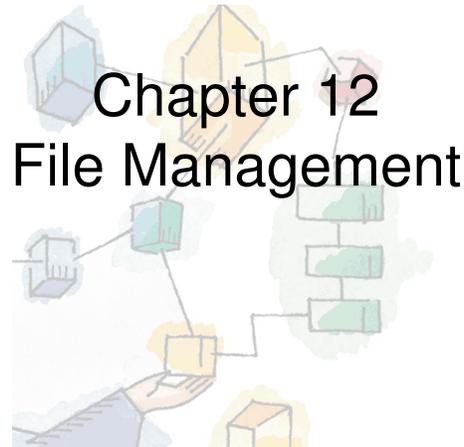
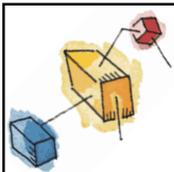


*Operating Systems:  
Internals and Design Principles, 6/E*  
William Stallings



Dave Bremer  
Otago Polytechnic, N.Z.  
©2008, Prentice Hall

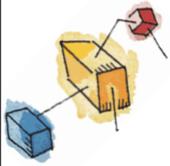


## Roadmap

### → Overview

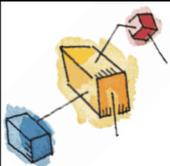
- File organisation and Access
- File Directories
- File Sharing
- Record Blocking
- Secondary Storage Management
- File System Security
- Unix File Management
- Linux Virtual File System
- Windows File System





# Files

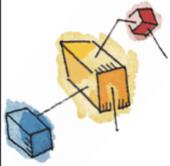
- Files are the central element to most applications
- The File System is one of the most important part of the OS to a user
- Desirable properties of files:
  - Long-term existence
  - Sharable between processes
  - Structure



# File Management

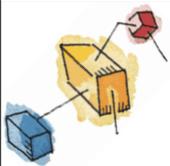
- File management system consists of system utility programs that run as privileged applications
- Concerned with secondary storage





## Typical Operations

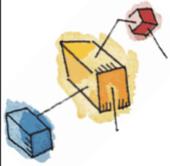
- File systems also provide functions which can be performed on files, typically:
  - Create
  - Delete
  - Open
  - Close
  - Read
  - Write



## Terms

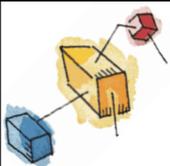
- Four terms are in common use when discussing files:
  - Field
  - Record
  - File
  - Database





## Fields and Records

- Fields
  - Basic element of data
  - Contains a single value
  - Characterized by its length and data type
- Records
  - Collection of related fields
  - Treated as a unit



## File and Database

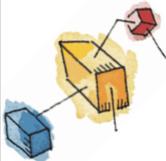
- File
  - Have file names
  - Is a collection of similar records
  - Treated as a single entity
  - May implement access control mechanisms
- Database
  - Collection of related data
  - Relationships exist among elements
  - Consists of one or more files





## File Management Systems

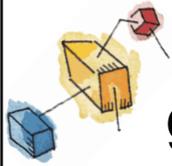
- Provides services to users and applications in the use of files
  - The way a user or application accesses files
- Programmer does not need to develop file management software



## Objectives for a File Management System

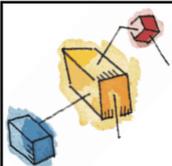
- Meet the data management needs of the user
- Guarantee that the data in the file are valid
- Optimize performance
- Provide I/O support for a variety of storage device types
- Minimize lost or destroyed data
- Provide a standardized set of I/O interface routines to user processes
- Provide I/O support for multiple users (if needed)





## Requirements for a general purpose system

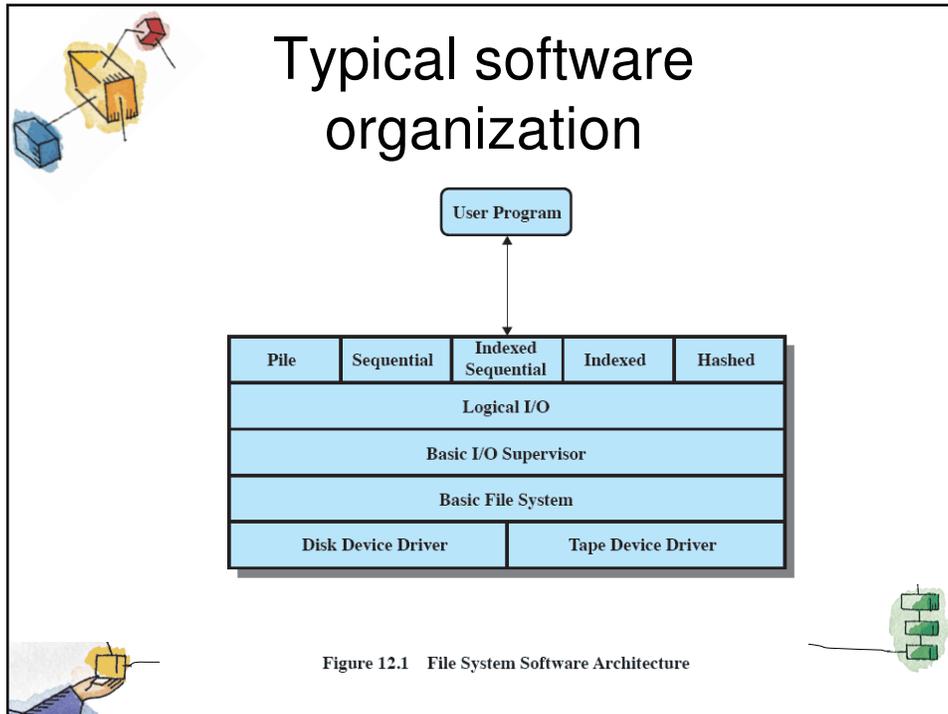
1. Each user should be able to create, delete, read, write and modify files
2. Each user may have controlled access to other users' files
3. Each user may control what type of accesses are allowed to the users' files
4. Each user should be able to restructure the user's files in a form appropriate to the problem



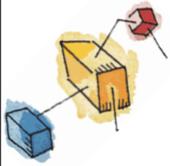
## Requirements cont.

5. Each user should be able to move data between files
6. Each user should be able to back up and recover the user's files in case of damage
7. Each user should be able to access the user's files by using symbolic names



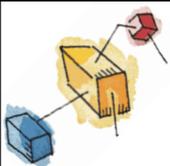


- ## Device Drivers
- Lowest level
  - Communicates directly with peripheral devices
  - Responsible for starting I/O operations on a device
  - Processes the completion of an I/O request



## Basic File System

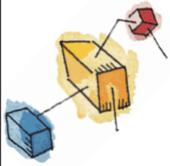
- Physical I/O
- Primary interface with the environment outside the computer system
- Deals with exchanging blocks of data
- Concerned with the placement of blocks
- Concerned with buffering blocks in main memory



## Basic I/O Supervisor

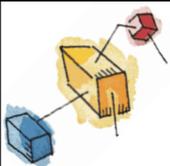
- Responsible for all file I/O initiation and termination.
- Control structures deal with
  - Device I/O,
  - Scheduling,
  - File status.
- Selects and schedules I/O with the device





## Logical I/O

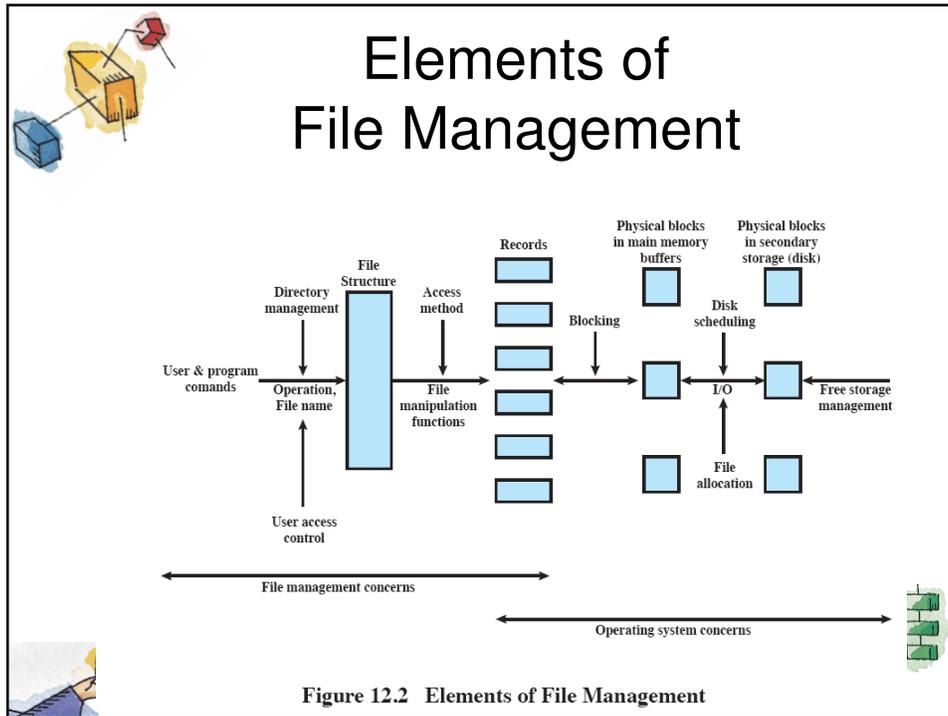
- Enables users and applications to access records
- Provides general-purpose record I/O capability
- Maintains basic data about file



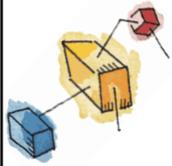
## Access Method

- Closest to the user
- Reflect different file structures
- Provides a standard interface between applications and the file systems and devices that hold the data
- Access method varies depending on the ways to access and process data for the device.



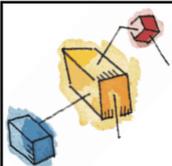


- # Roadmap
- Overview
  - ➔ **File organisation and Access**
  - File Directories
  - File Sharing
  - Record Blocking
  - Secondary Storage Management
  - File System Security
  - Unix File Management
  - Linux Virtual File System
  - Windows File System



## File Organization

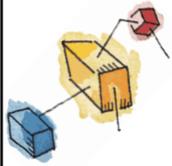
- File Management Referring to the logical structure of records
  - Physical organization discussed later
- Determined by the **way** in which files are accessed



## Criteria for File Organization

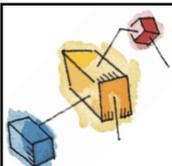
- Important criteria include:
  - Short access time
  - Ease of update
  - Economy of storage
  - Simple maintenance
  - Reliability
- Priority will differ depending on the use (e.g. read-only CD vs Hard Drive)
  - Some may even conflict





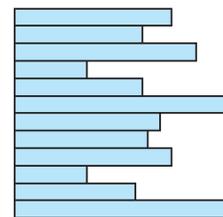
## File Organisation Types

- Many exist, but usually variations of:
  - Pile
  - Sequential file
  - Indexed sequential file
  - Indexed file
  - Direct, or hashed, file



## The Pile

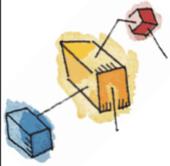
- Data are collected in the order they arrive
  - No structure
- Purpose is to accumulate a mass of data and save it
- Records may have different fields
- Record access is by exhaustive search



Variable-length records  
Variable set of fields  
Chronological order

(a) Pile File



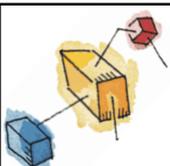


## The Sequential File

- Fixed format used for records
- Records are the same length
- All fields the same (order and length)
- Field names and lengths are attributes of the file
- Key field
  - Uniquely identifies the record
  - Records are stored in key sequence


Fixed-length records  
Fixed set of fields in fixed order  
Sequential order based on key field

(b) Sequential File

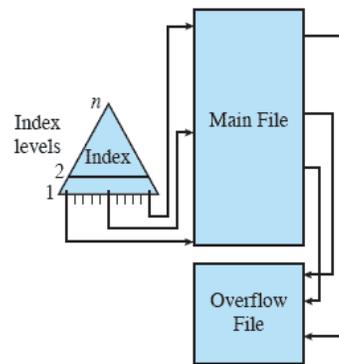


## Indexed Sequential File

- Maintains the key characteristic of the sequential file:
  - records are organized in sequence based on a key field.

Two features are added:

- an index to the file to support random access,
- and an overflow file.



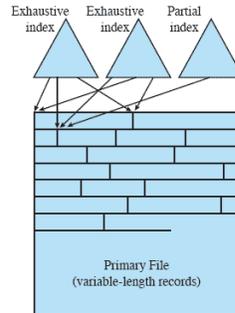
(c) Indexed Sequential File



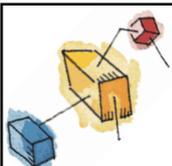


## Indexed File

- Uses multiple indexes for different key fields
  - May contain an exhaustive index that contains one entry for every record in the main file
  - May contain a partial index
- When a new record is added to the main file, all of the index files must be updated.



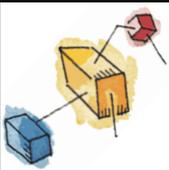
(d) Indexed File



## File Organization

- Access directly any block of a known address.
- The Direct or Hashed File
  - Directly access a block at a known address
  - Key field required for each record





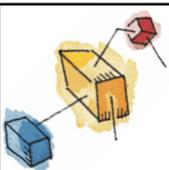
# Performance

Table 12.1 Grades of Performance for Five Basic File Organizations [WIED87]

File Method	Space Attributes		Update Record Size		Retrieval		
	Variable	Fixed	Equal	Greater	Single record	Subset	Exhaustive
Pile	A	B	A	E	E	D	B
Sequential	F	A	D	F	F	D	A
Indexed sequential	F	B	B	D	B	D	B
Indexed	B	C	C	C	A	B	D
Hashed	F	B	B	F	B	F	E

A = Excellent, well suited to this purpose  $\approx O(r)$   
 B = Good  $\approx O(o \times r)$   
 C = Adequate  $\approx O(r \log n)$   
 D = Requires some extra effort  $\approx O(n)$   
 E = Possible with extreme effort  $\approx O(r \times n)$   
 F = Not reasonable for this purpose  $\approx O(n^2)$

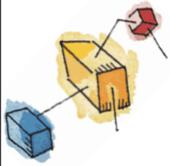
where  
 $r$  = size of the result  
 $o$  = number of records that overflow  
 $n$  = number of records in file



# Roadmap

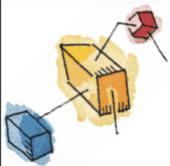
- Overview
- File organisation and Access
- ➔ **File Directories**
  - File Sharing
  - Record Blocking
  - Secondary Storage Management
  - File System Security
  - Unix File Management
  - Linux Virtual File System
  - Windows File System





## Contents

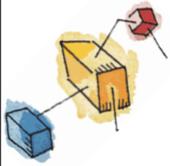
- Contains information about files
  - Attributes
  - Location
  - Ownership
- Directory itself is a file owned by the operating system
- Provides mapping between file names and the files themselves



## Directory Elements: Basic Information

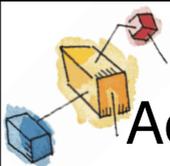
- File Name
  - Name as chosen by creator (user or program).
  - Must be unique within a specific directory.
- File type
- File Organisation
  - For systems that support different organizations





## Directory Elements: Address Information

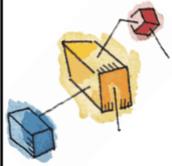
- Volume
  - Indicates device on which file is stored
- Starting Address
- Size Used
  - Current size of the file in bytes, words, or blocks
- Size Allocated
  - The maximum size of the file



## Directory Elements: Access Control Information

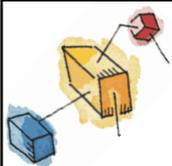
- Owner
  - The owner may be able to grant/deny access to other users and to change these privileges.
- Access Information
  - May include the user's name and password for each authorized user.
- Permitted Actions
  - Controls reading, writing, executing, transmitting over a network





## Directory Elements: Usage Information

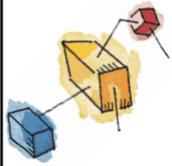
- Date Created
- Identity of Creator
- Date Last Read Access
- Identity of Last Reader
- Date Last Modified
- Identity of Last Modifier
- Date of Last Backup
- Current Usage
  - Current activity, locks, etc



## Simple Structure for a Directory

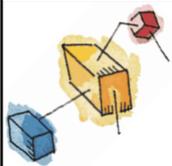
- The method for storing the previous information varies widely between systems
- Simplest is a list of entries, one for each file
  - Sequential file with the name of the file serving as the key
  - Provides no help in organizing the files
  - Forces user to be careful not to use the same name for two different files





## Operations Performed on a Directory

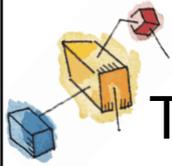
- A directory system should support a number of operations including:
  - Search
  - Create files
  - Deleting files
  - Listing directory
  - Updating directory



## Two-Level Scheme for a Directory

- One directory for each user and a master directory
  - Master directory contains entry for each user
  - Provides address and access control information
- Each user directory is a simple list of files for that user
  - Does not provide structure for collections of files





## Hierarchical, or Tree-Structured Directory

- Master directory with user directories underneath it
- Each user directory may have subdirectories and files as entries

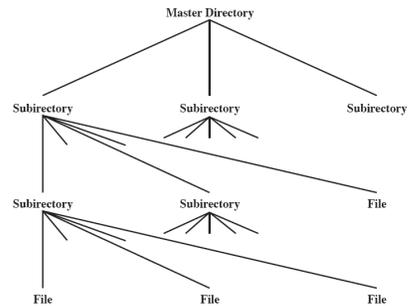
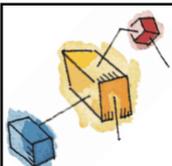


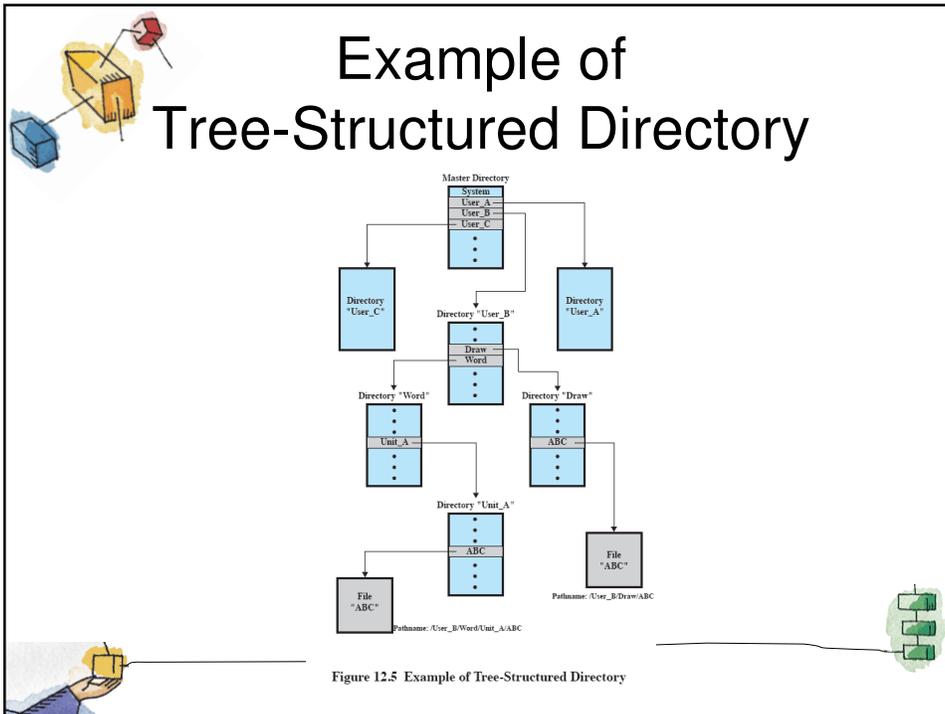
Figure 12.4 Tree-Structured Directory



## Naming

- Users need to be able to refer to a file by name
  - Files need to be named uniquely, but users may not be aware of all filenames on a system
- The tree structure allows users to find a file by following the directory path
  - Duplicate filenames are possible if they have different pathnames





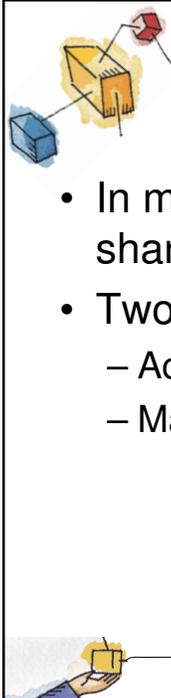
## Working Directory

- Stating the full pathname and filename is awkward and tedious
- Usually an interactive user or process is associated with a **current** or **working directory**
  - All file names are referenced as being relative to the working directory unless an explicit full pathname is used



# Roadmap

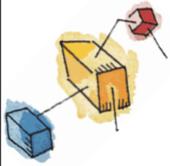
- Overview
- File organisation and Access
- File Directories
- ➔ **File Sharing**
- Record Blocking
- Secondary Storage Management
- File System Security
- Unix File Management
- Linux Virtual File System
- Windows File System



# File Sharing

- In multiuser system, allow files to be shared among users
- Two issues
  - Access rights
  - Management of simultaneous access





## Access Rights

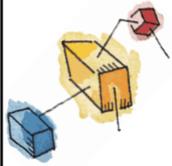
- A wide variety of access rights have been used by various systems
  - often as a hierarchy where one right implies previous
- None
  - User may not even know of the files existence
- Knowledge
  - User can only determine that the file exists and who its owner is



## Access Rights cont...

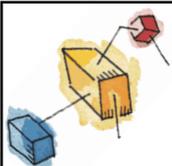
- Execution
  - The user can load and execute a program but cannot copy it
- Reading
  - The user can read the file for any purpose, including copying and execution
- Appending
  - The user can add data to the file but cannot modify or delete any of the file's contents





## Access Rights cont...

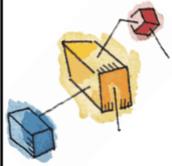
- Updating
  - The user can modify, delete, and add to the file's data.
- Changing protection
  - User can change access rights granted to other users
- Deletion
  - User can delete the file



## User Classes

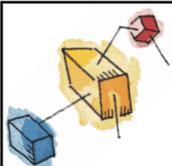
- Owner
  - Usually the files creator, usually has full rights
- Specific Users
  - Rights may be explicitly granted to specific users
- User Groups
  - A set of users identified as a group
- All
  - everyone





## Simultaneous Access

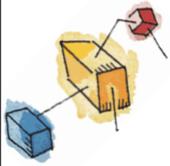
- User may lock entire file when it is to be updated
- User may lock the individual records during the update
- Mutual exclusion and deadlock are issues for shared access



## Roadmap

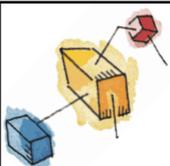
- Overview
- File organisation and Access
- File Directories
- File Sharing
- ➔ **Record Blocking**
- Secondary Storage Management
- File System Security
- Unix File Management
- Linux Virtual File System
- Windows File System





## Blocks and records

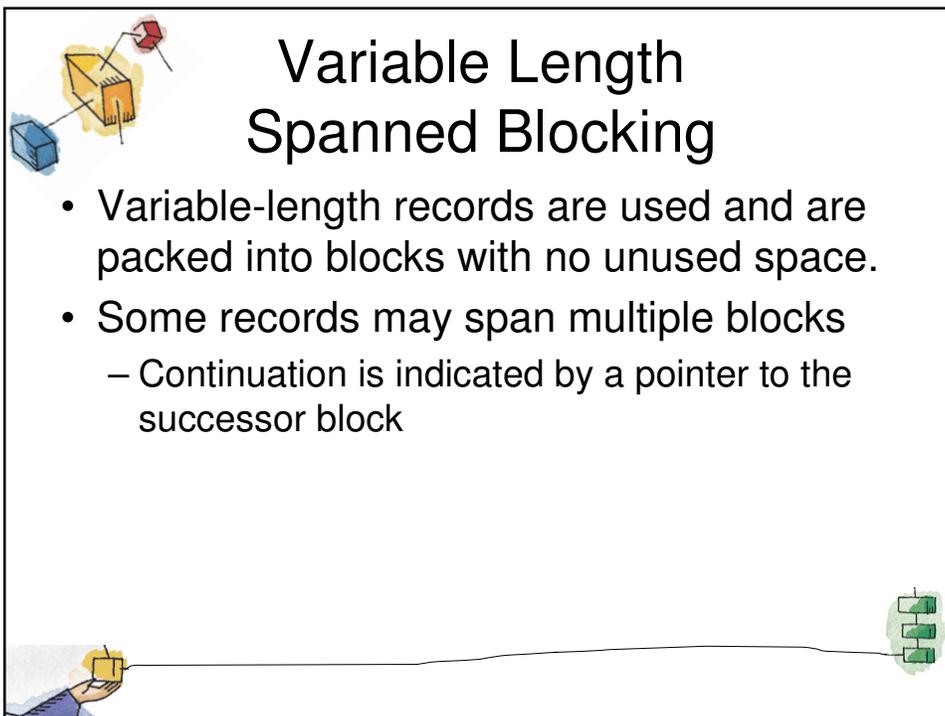
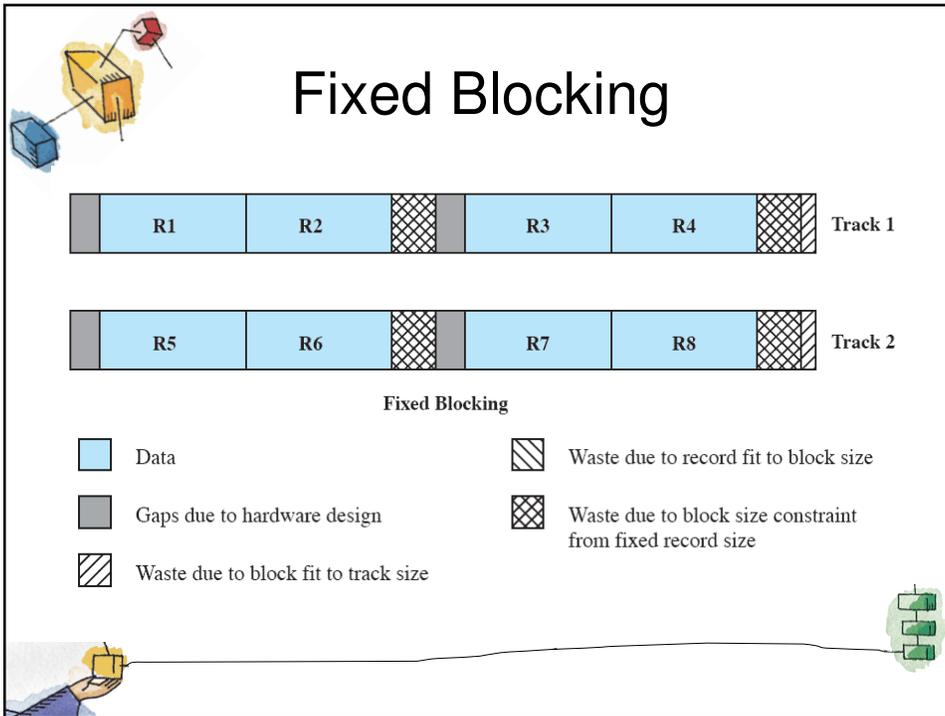
- Records are the logical unit of access of a structured file
  - But blocks are the unit for I/O with secondary storage
- Three approaches are common
  - Fixed length blocking
  - Variable length spanned blocking
  - Variable-length unspanned blocking

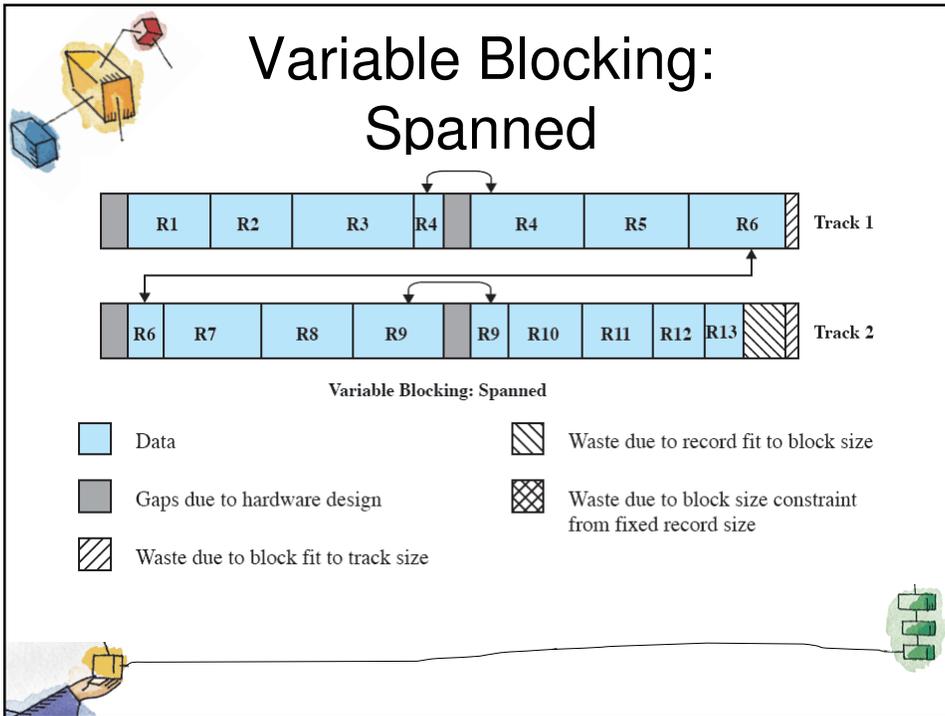


## Fixed Blocking

- Fixed-length records are used, and an integral number of records are stored in a block.
- Unused space at the end of a block is ***internal fragmentation***

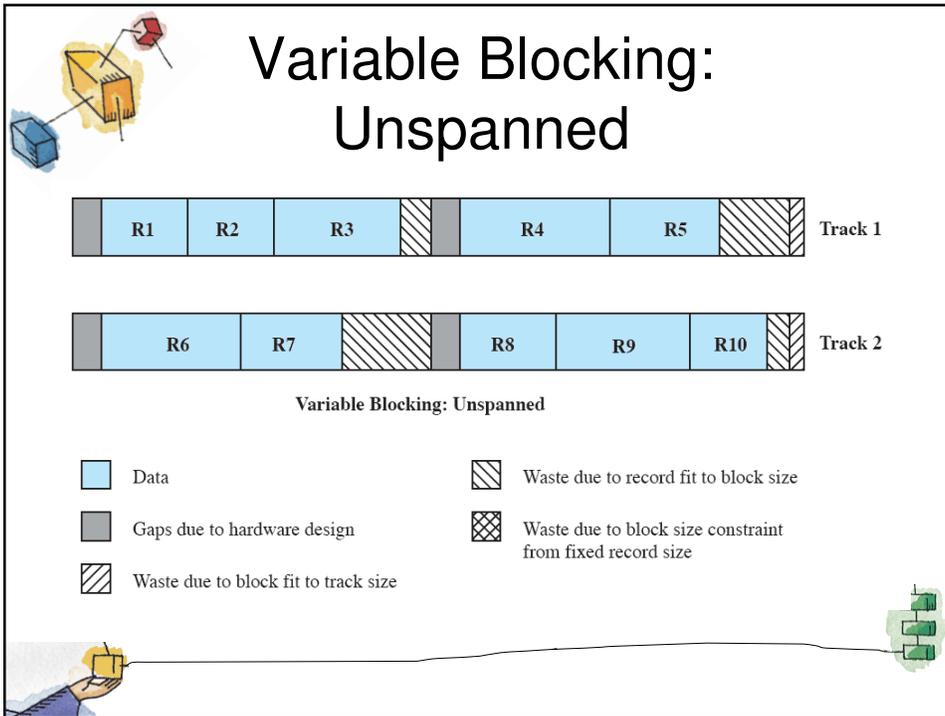




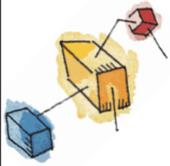


## Variable-length unspanned blocking

- Uses variable length records without spanning
- Wasted space in most blocks because of the inability to use the remainder of a block if the next record is larger than the remaining unused space.

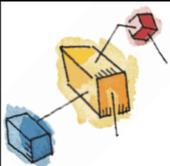


- ## Roadmap
- Overview
  - File organisation and Access
  - File Directories
  - File Sharing
  - Record Blocking
  - ➔ **Secondary Storage Management**
    - File System Security
    - Unix File Management
    - Linux Virtual File System
    - Windows File System



## Secondary Storage Management

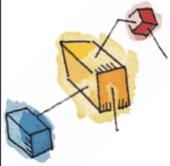
- The Operating System is responsible for allocating blocks to files
- Two related issues
  - Space must be allocated to files
  - Must keep track of the space available for allocation



## File allocation issues

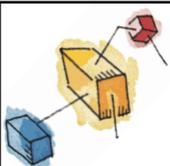
1. When a file is created – is the maximum space allocated at once?
2. Space is added to a file in contiguous 'portions'
  - What size should be the 'portion'?
3. What data structure should be used to keep track of the file portions?





## Preallocation vs Dynamic Allocation

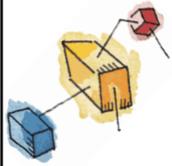
- Need the maximum size for the file at the time of creation
- Difficult to reliably estimate the maximum potential size of the file
- Tend to overestimated file size so as not to run out of space



## Portion size

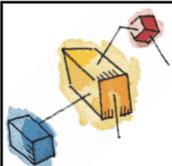
- Two extremes:
  - Portion large enough to hold entire file is allocated
  - Allocate space one block at a time
- Trade-off between efficiency from the point of view of a single file, or the overall system efficiency





## File Allocation Method

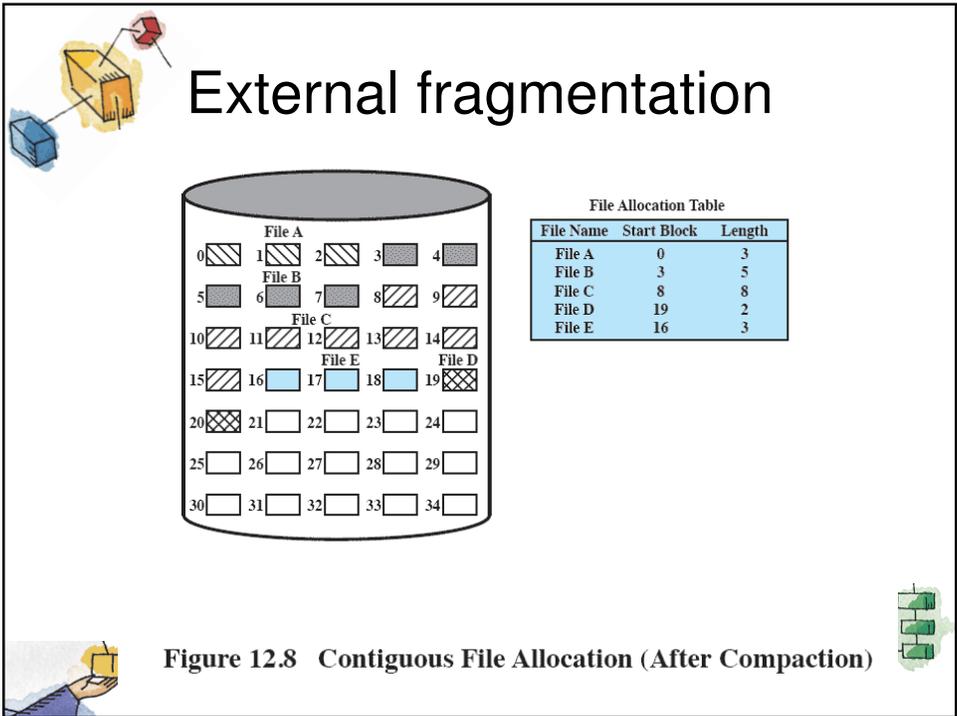
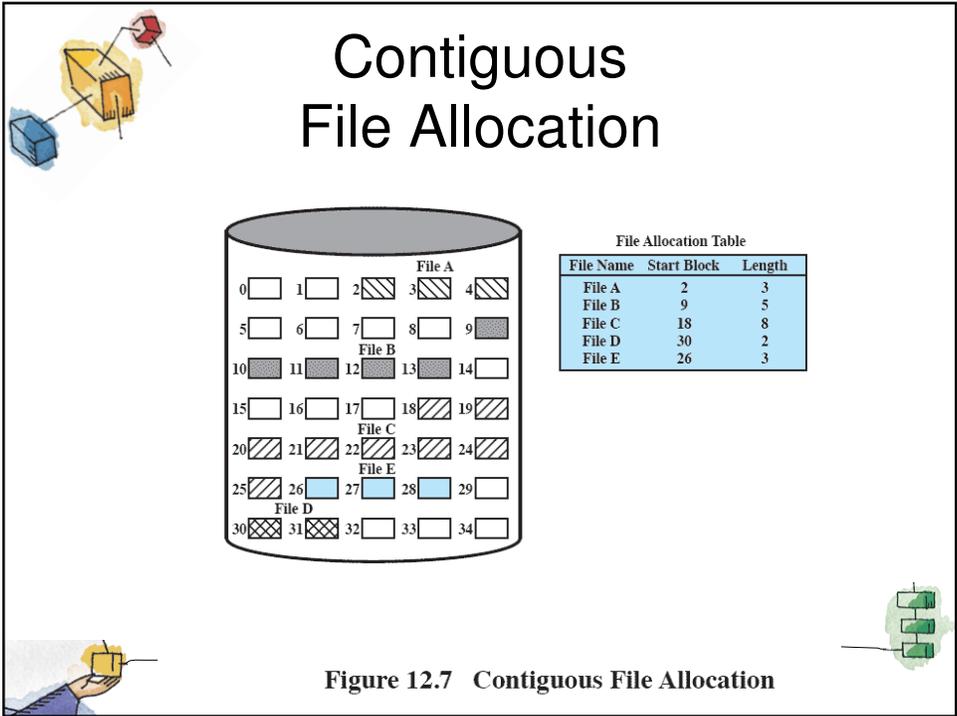
- Three methods are in common use:
  - contiguous,
  - chained, and
  - indexed.

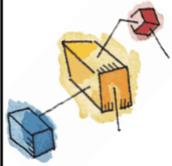


## Contiguous Allocation

- Single set of blocks is allocated to a file at the time of creation
- Only a single entry in the file allocation table
  - Starting block and length of the file
- External fragmentation will occur
  - Need to perform compaction

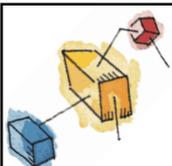




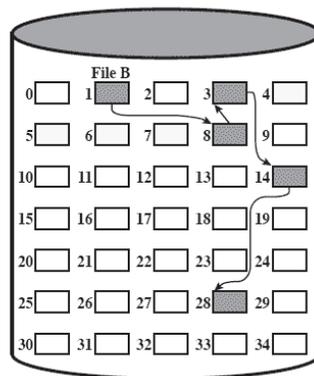


## Chained Allocation

- Allocation on basis of individual block
- Each block contains a pointer to the next block in the chain
- Only single entry in the file allocation table
  - Starting block and length of file
- No external fragmentation
- Best for sequential files

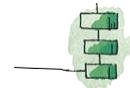


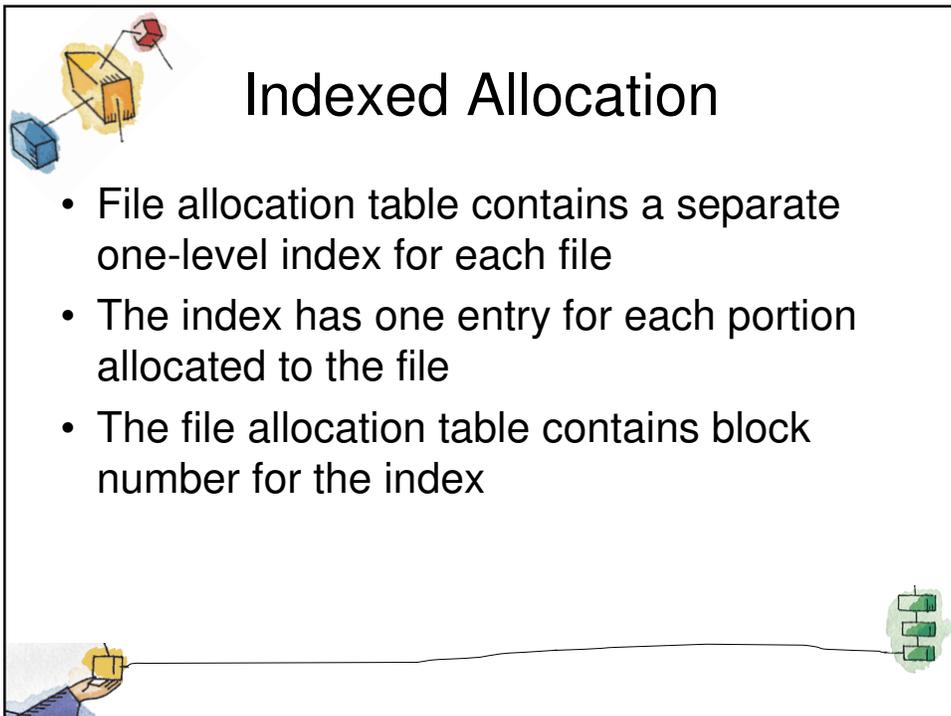
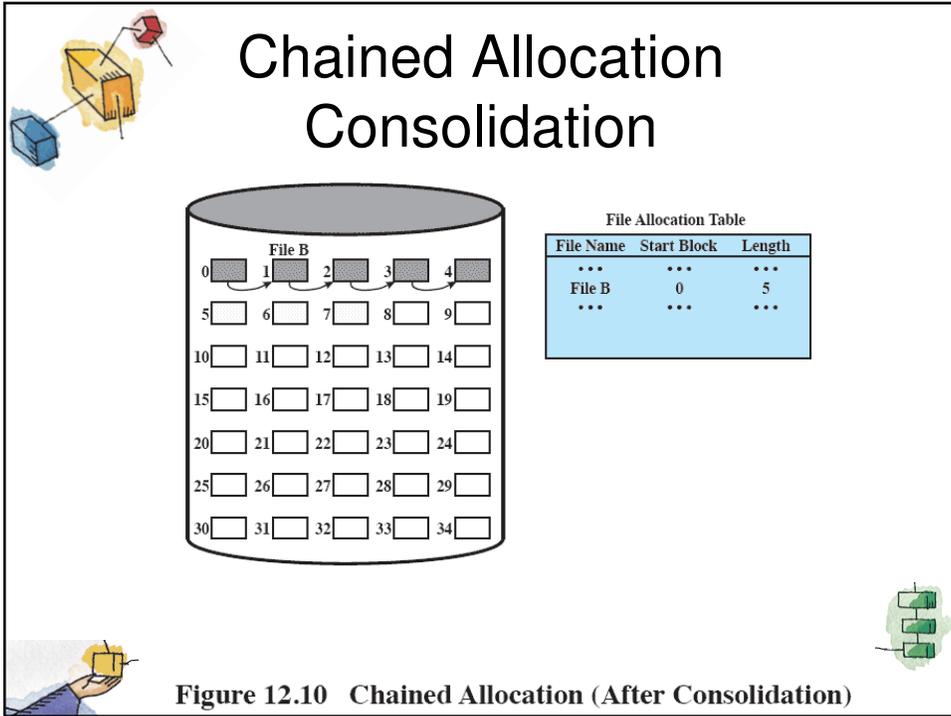
## Chained Allocation

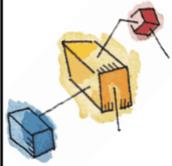


File Allocation Table		
File Name	Start Block	Length
...	...	...
File B	1	5
...	...	...

Figure 12.9 Chained Allocation

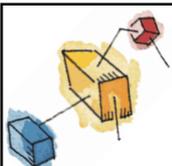






## Indexed Allocation Method

- Allocation may be either
  - Fixed size blocks or
  - Variable sized blocks
- Allocating by blocks eliminates external fragmentation
- Variable sized blocks improves locality
- Both cases require occasional consolidation



## Indexed allocation with Block Portions

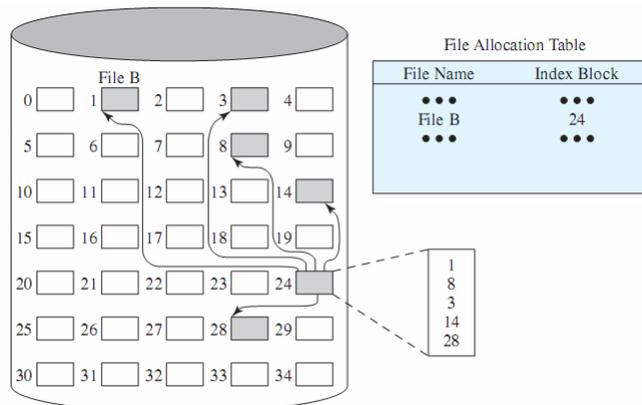
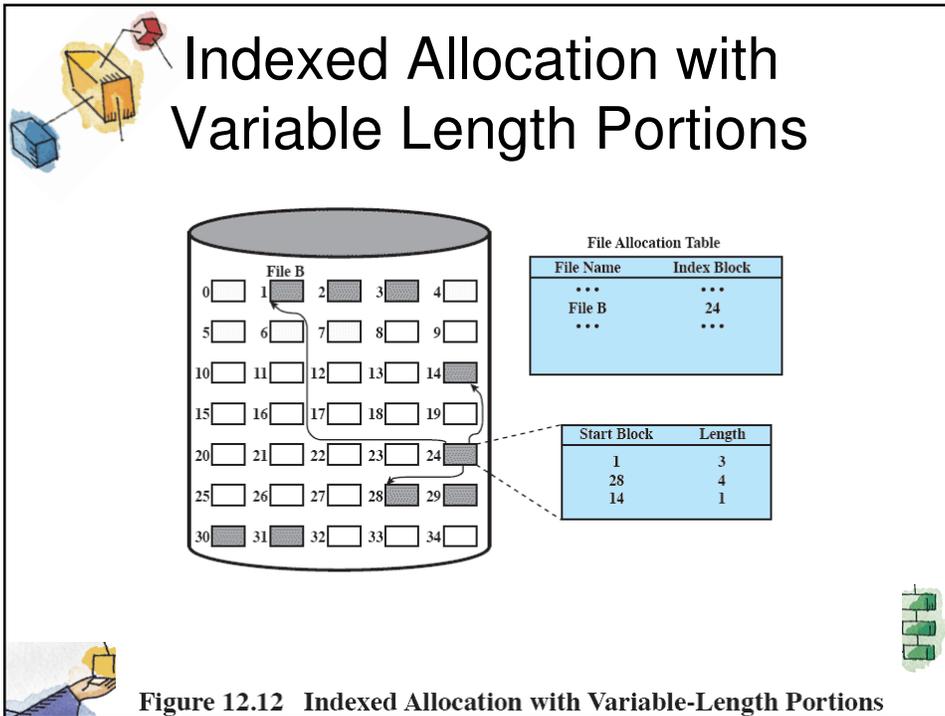


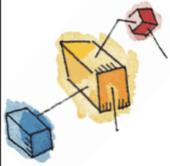
Figure 12.11 Indexed Allocation with Block Portions





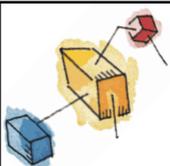
**Free Space Management**

- Just as allocated space must be managed, so must the unallocated space
- To perform file allocation, we need to know which blocks are available.
- We need a disk allocation table in addition to a file allocation table



## Bit Tables

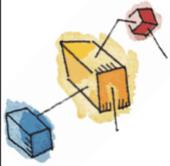
- This method uses a vector containing one bit for each block on the disk.
- Each entry of a 0 corresponds to a free block,
  - and each 1 corresponds to a block in use.
- Advantages:
  - Works well with any file allocation method
  - Small as possible



## Chained Free Portions

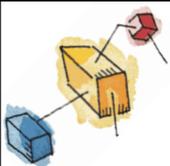
- The free portions may be chained together by using a pointer and length value in each free portion.
- Negligible space overhead
- Suited to all file allocation methods
- Leads to fragmentation





## Indexing

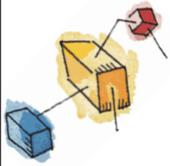
- treats free space as a file and uses an index table as it would for file allocation
- For efficiency, the index should be on the basis of variable-size portions rather than blocks.
  - Thus, there is one entry in the table for every free portion on the disk.
- This approach provides efficient support for all of the file allocation methods.



## Free Block List

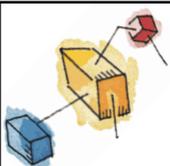
- Each block is assigned a number sequentially
  - the list of the numbers of all free blocks is maintained in a reserved portion of the disk.





# Volumes

- A collection of addressable sectors in secondary memory that an OS or application can use for data storage.
- The sectors in a volume need not be consecutive on a physical storage device;
  - instead they need only appear that way to the OS or application.
- A volume may be the result of assembling and merging smaller volumes.



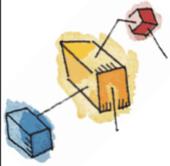
# Roadmap

- Overview
- File organisation and Access
- File Directories
- File Sharing
- Record Blocking
- Secondary Storage Management

## → File System Security

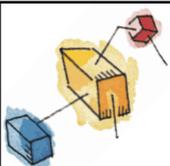
- Unix File Management
- Linux Virtual File System
- Windows File System





# Access Control

- By successfully logging on to a system, the user is identified
- The OS can then enforce rules
  - Granting access to files and applications (or denying)
- The OS needs a rule-set to enforce



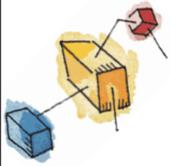
# Access Matrix

- One such rule set is an Access Matrix

	File 1	File 2	File 3	File 4	Account 1	Account 2
User A	Own R W		Own R W		Inquiry Credit	
User B	R	Own R W	W	R	Inquiry Debit	Inquiry Credit
User C	R W	R		Own R W		Inquiry Debit

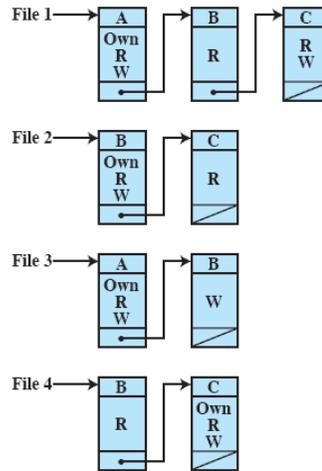
(a) Access matrix



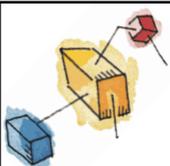


## Access Control Lists

- A matrix may be decomposed by columns
- Giving an Access Control List (ACL) for each file.

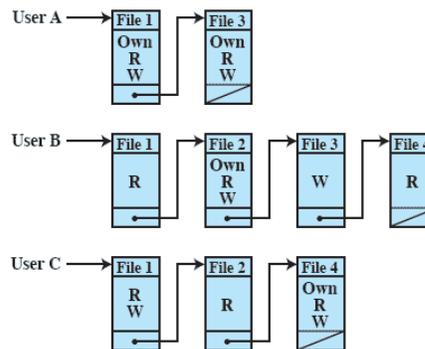


(b) Access control lists for files of part (a)



## Capability Lists

- Decomposition by rows yields capability lists (or ticket)
  - specifies authorized objects and operations for a user.



(c) Capability lists for files of part (a)





# Roadmap

- Overview
- File organisation and Access
- File Directories
- File Sharing
- Record Blocking
- Secondary Storage Management
- File System Security

→ **Unix File Management**

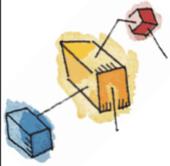
- Linux Virtual File System
- Windows File System



# UNIX File Management

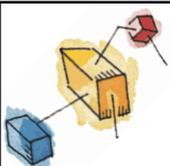
- Six types of files
  - Regular, or ordinary
  - Directory
  - Special
  - Named pipes
  - Links
  - Symbolic links





# Inodes

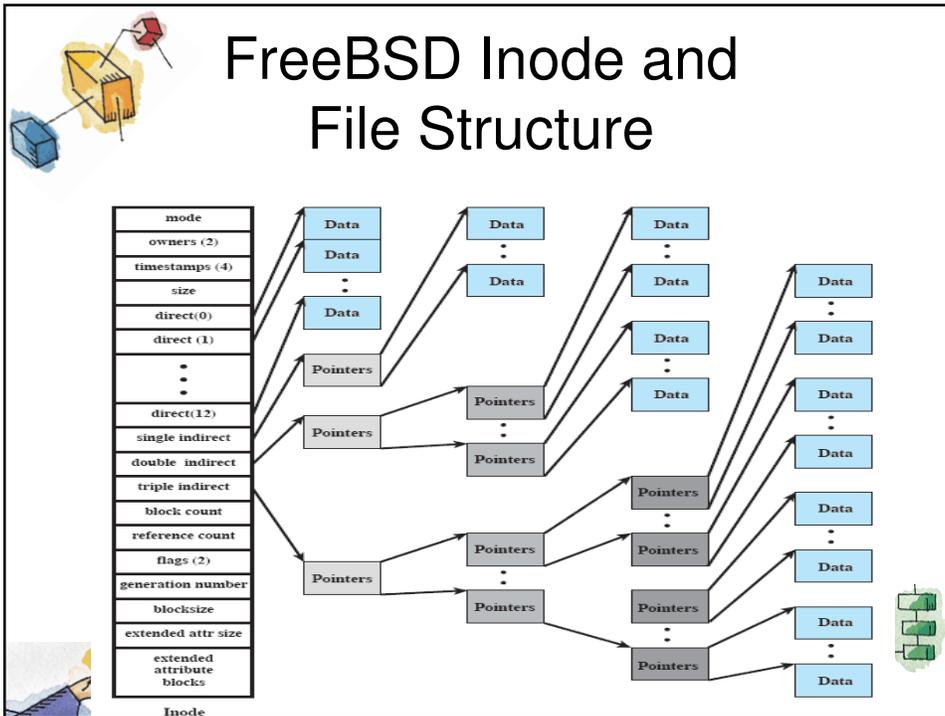
- Index node
- Control structure that contains key information for a particular file.
- Several filenames may be associated with a single inode
  - But an active inode is associated with only one file, and
  - Each file is controlled by only one inode



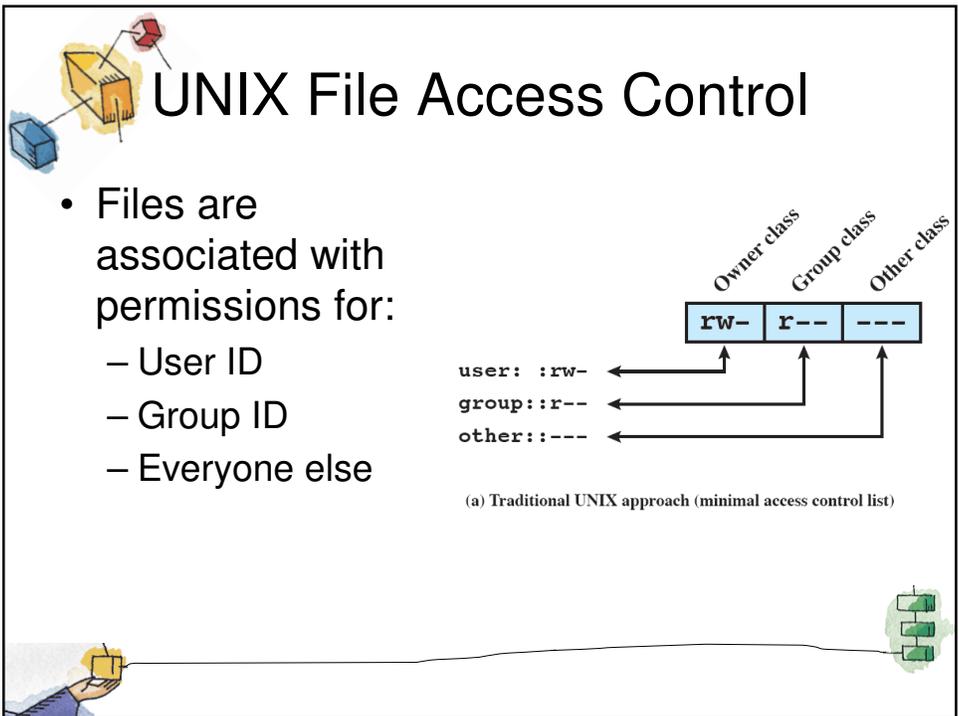
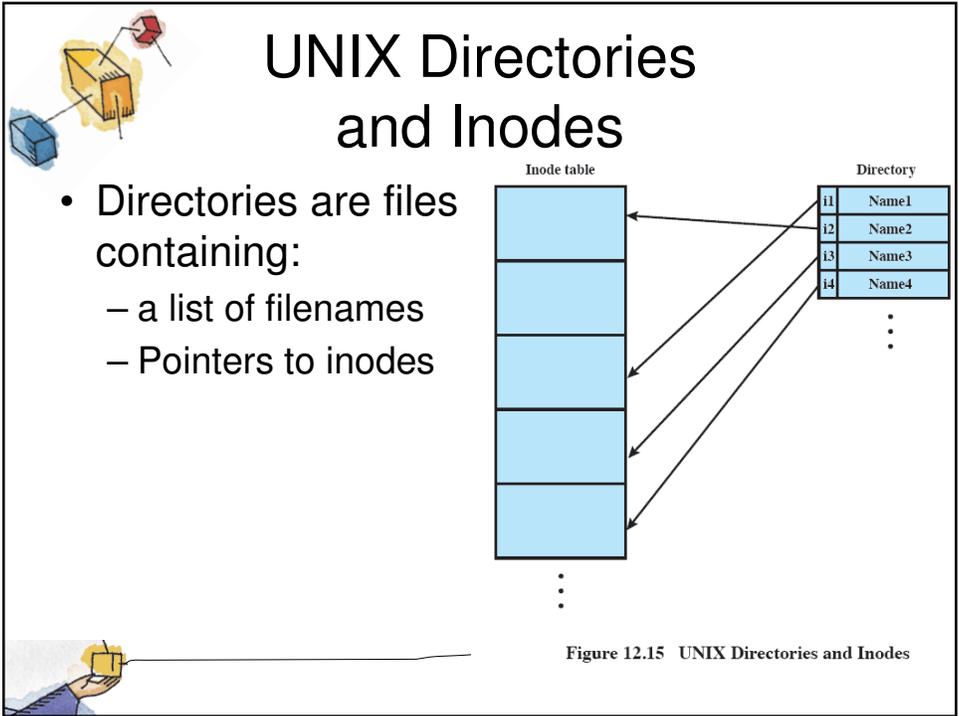
# Free BSD Inodes include:

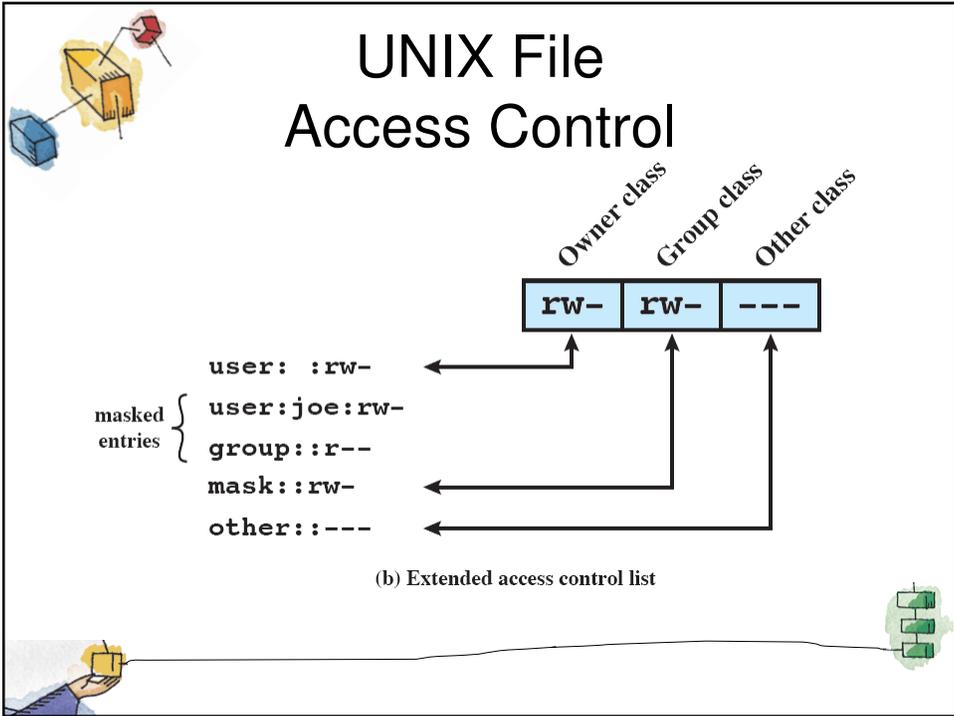
- The type and access mode of the file
- The file's owner and group-access identifiers
- Creation time, last read/write time
- File size
- Sequence of block pointers
- Number of blocks and Number of directory entries
- Blocksize of the data blocks
- Kernel and user settable flags
- Generation number for the file
- Size of Extended attribute information
- Zero or more extended attribute entries



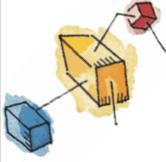


- ## File Allocation
- File allocation is done on a block basis.
  - Allocation is dynamic
    - Blocks may not be contiguous
  - Index method keeps track of files
    - Part of index stored in the file inode.
  - Inode includes a number of direct pointers
    - And three indirect pointers





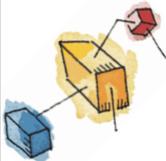
- # Roadmap
- Overview
  - File organisation and Access
  - File Directories
  - File Sharing
  - Record Blocking
  - Secondary Storage Management
  - File System Security
  - Unix File Management
  - ➔ Linux Virtual File System
  - Windows File System



# Linux Virtual File System

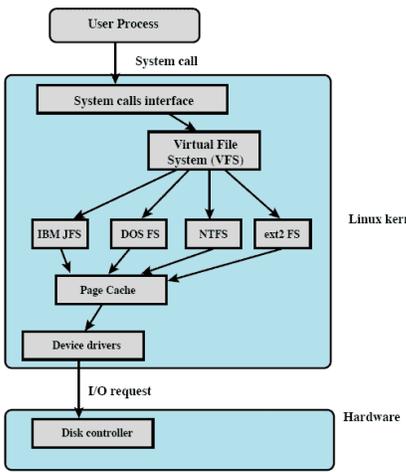
- Uniform file system interface to user processes
- Represents any conceivable file system's general feature and behavior
- Assumes files are objects that share basic properties regardless of the target file system





## Key ingredients of VFS Strategy

- A user process issues a file system call (e.g., read) using the VFS file scheme.
  - The VFS converts this into an internal file system call
  - The new call is passed to a mapping function for a specific file system



```

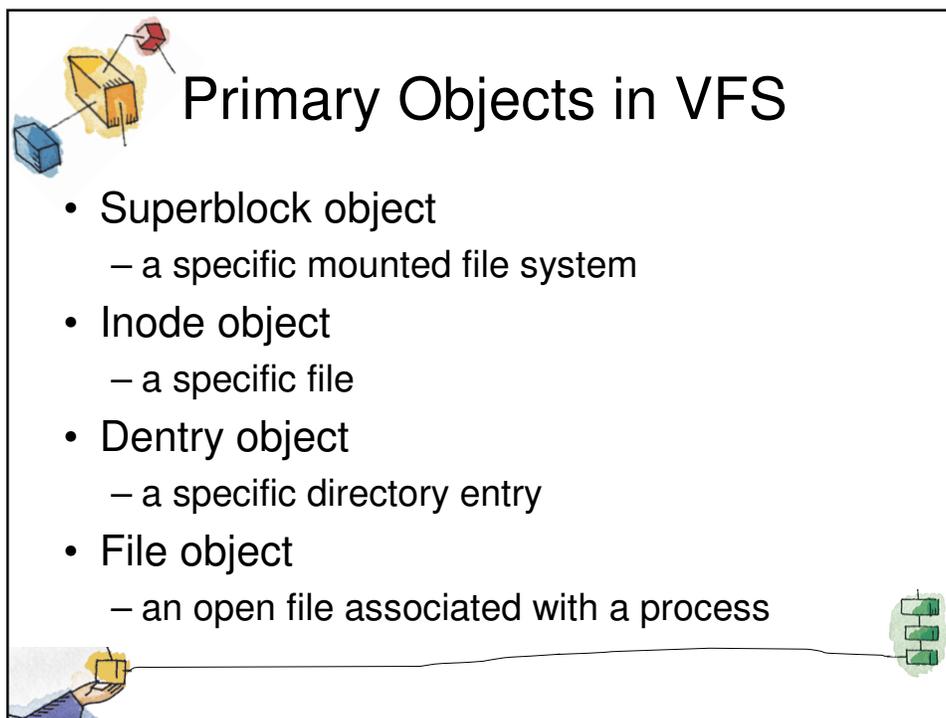
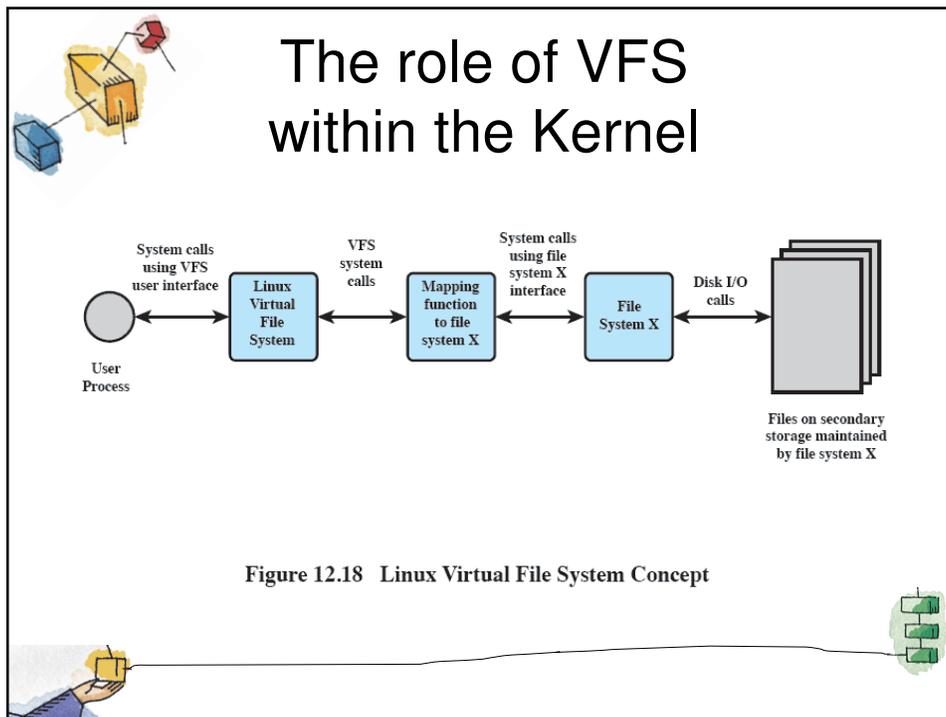
graph TD
    UP[User Process] -- System call --> SSI[System calls interface]
    SSI --> VFS[Virtual File System VFS]
    VFS --> IBM[IBM JFS]
    VFS --> DOS[DOS FS]
    VFS --> NTFS[NTFS]
    VFS --> ext2[ext2 FS]
    IBM --> PC[Page Cache]
    DOS --> PC
    NTFS --> PC
    ext2 --> PC
    PC --> DD[Device drivers]
    DD -- I/O request --> DC[Disk controller]
  
```

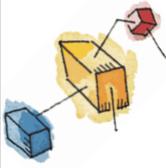
Linux kernel

Hardware

Figure 12.17 Linux Virtual File System Context



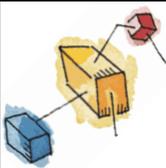




# Roadmap

- Overview
- File organisation and Access
- File Directories
- File Sharing
- Record Blocking
- Secondary Storage Management
- File System Security
- Unix File Management
- Linux Virtual File System

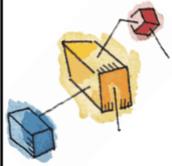
→ Windows File System



# Windows File System

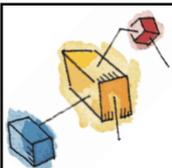
- Key features of NTFS
  - Recoverability
  - Security
  - Large disks and large files
  - Multiple data streams
  - Journaling
  - Compression and Encryption





## NTFS Volume and File Structure

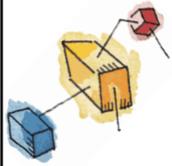
- Sector
  - The smallest physical storage unit on the disk
  - Almost always 512 bytes
- Cluster
  - One or more contiguous sectors
- Volume
  - Logical partition on a disk



## Efficient with Large Files

Volume Size	Sectors per Cluster	Cluster Size
512Mbyte	1	512bytes
512Mbyte – 1 Gbyte	2	1K
1–2 Gbyte	4	2K
2–4 Gbyte	8	4K
4–8 Gbyte	16	8K
8–16 Gbyte	32	16K
16–32 Gbyte	64	32K
>32Gbyte	128	64K





## NTFS Volume Layout

- Every element on a volume is a file, and every file consists of a collection of attributes.
  - Even the data contents of a file is treated as an attribute.

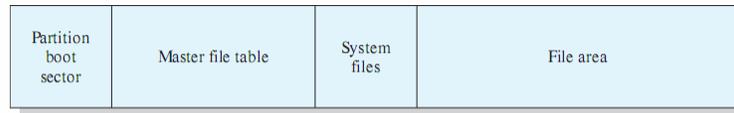
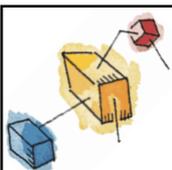


Figure 12.19 NTFS Volume Layout



## Windows NTFS Components

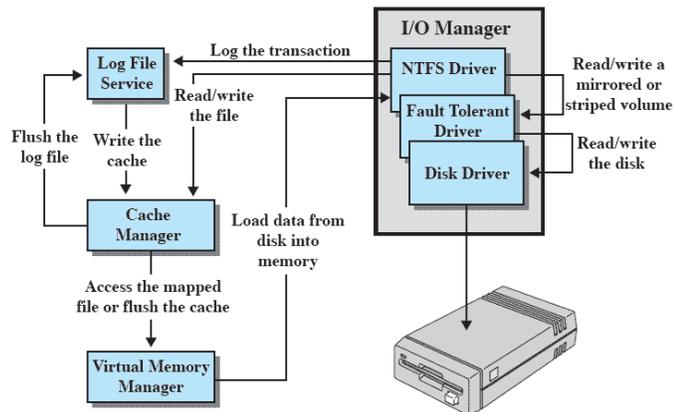


Figure 12.20 Windows NTFS Components

