

Lecture Note 6. File System Basic

April 27, 2026

Jongmoo Choi

Dept. of Software
Dankook University

<http://embedded.dankook.ac.kr/~choijm>

(본 교재는 2026년도 과학기술정보통신부 및 정보통신기획평가원의 'SW중심대학사업' 지원을 받아 제작 되었습니다.)

Contents

- Chap 35. A Dialogue on Persistence
- Chap 36. I/O Devices
- Chap 37. Hard Disk Drives
- Chap 38. RAID
- Chap 39. Interlude: Files and Directories
 - ✓ APIs for file, directory and file system
- Chap 40. File System Implementation
 - ✓ **Layout**: superblock, bitmap, inode, data blocks, ...
 - ✓ **Interface (access method)**: open, read, write, close, lseek, fsync, mount, ...

Chap. 35 A Dialogue on Persistence

Professor: *And thus we reach the third of our four ... err... three pillars of operating systems: **persistence**.*

Student: *Did you say there were three pillars, or four? What is the fourth?*

Professor: *No. Just three, young student, just three. Trying to keep it simple here.*

Student: *OK, fine. But what is persistence, oh fine and noble professor?*

Professor: *Actually, you probably know what it means in the traditional sense, right? As the dictionary would say: "a firm or obstinate continuance in a course of action in spite of difficulty or opposition."*

Student: *It's kind of like taking your class: some obstinance required.*

Professor: *Ha! Yes. But persistence here means something else. Let me explain. Imagine you are outside, in a field, and you pick a —*

Student: *(interrupting) I know! A peach! From a peach tree!*

Professor: *I was going to say apple, from an apple tree. Oh well; we'll do it your way, I guess.*

Student: *(stares blankly)*

Professor: *Anyhow, you pick a peach; in fact, you pick many many peaches, but you want to make them last for a long time. Winter is hard and cruel in Wisconsin, after all. What do you do?*

Student: *Well, I think there are some different things you can do. You can pickle it! Or bake a pie. Or make a jam of some kind. Lots of fun!*

Professor: *Fun? Well, maybe. Certainly, you have to do a lot more work to make the peach **persist**. And so it is with information as well; making information persist, despite computer crashes, disk failures, or power outages is a tough and interesting challenge.*



Student: *Nice segue; you're getting quite good at that.*

Professor: *Thanks! A professor can always use a few kind words, you know.*

• Persistence : Making information durable despite of computer crashes, disk failures and so on

Chap. 36 I/O Devices

- 36.1 System Architecture
- 36.2 A Canonical Device
- 36.3 The Canonical Protocol
- 36.4 Lowering CPU overhead with **Interrupt**
- 36.5 More Efficient Data Movement with **DMA**
- 36.6 Methods of Device Interaction
- 36.7 Fitting into the OS: The **Device Driver**
- 36.8 Case Study: A simple IDE Disk Driver
- 36.9 Historical Notes



(Source: <https://gcallah.github.io/OperatingSystems/IOHardware.html>)

36.1 System Architecture

■ Computer system focusing on Bus

✓ Hierarchical structure

- Memory bus (System bus): CPU and Memory

- Fast, Expensive, Short

- I/O bus: SCSI, SATA, USB (and/or separated bus for Graphic Cards)

- Slow, Less expensive, long, pluggable

✓ Modern system

- Special interconnect: Memory interconnect (e.g. QPI, Hyperport), Graphic interconnect

- Make use of specialized chipsets: I/O chips with different interfaces

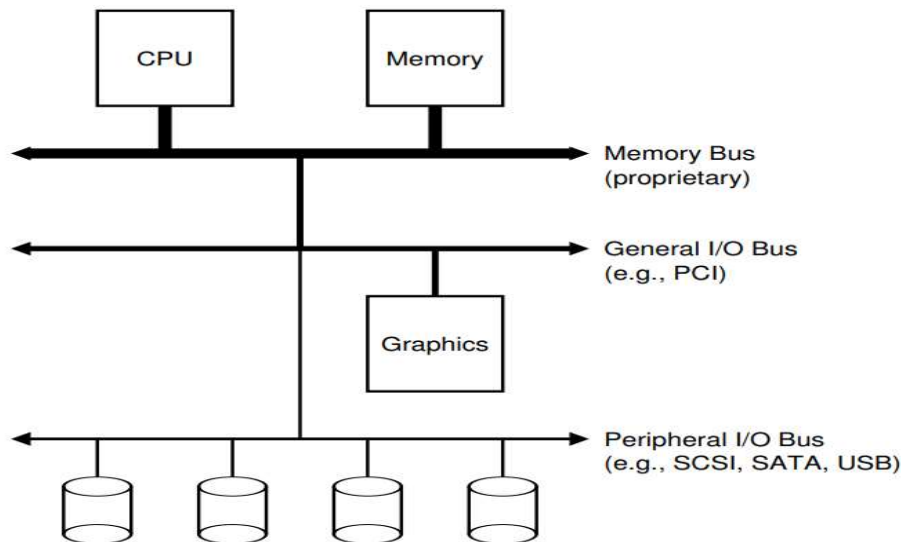


Figure 36.1: Prototypical System Architecture

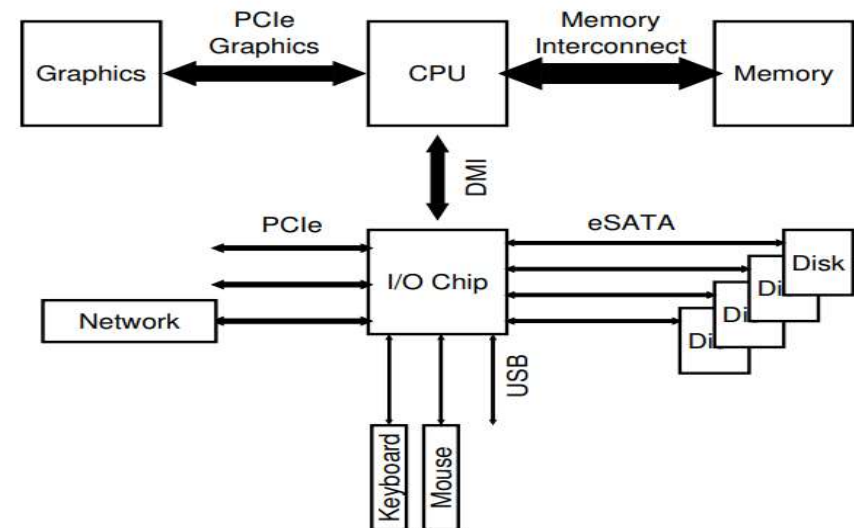


Figure 36.2: Modern System Architecture

36.2 A Canonical Device / 36.3 The Canonical Protocol

■ Devices

✓ Interface parts

- Registers: 1) command, 2) status, 3) data, ...

✓ Internals

- Logic: controller and special chips (device specific) + SW (called **firmware**)
- Memory: I/O Buffer (e.g. store receiving packet, delayed write, ...)

■ Protocol

✓ How to interact with devices?

- Example: Four steps 1) idle check, 2) data, 3) command, 4) finish check

✓ 3 mechanisms: 1) PIO (Programmed I/O), 2) Interrupt, 3) DMA

- **PIO**: CPU performs all steps including idle/finish checking (**polling**)

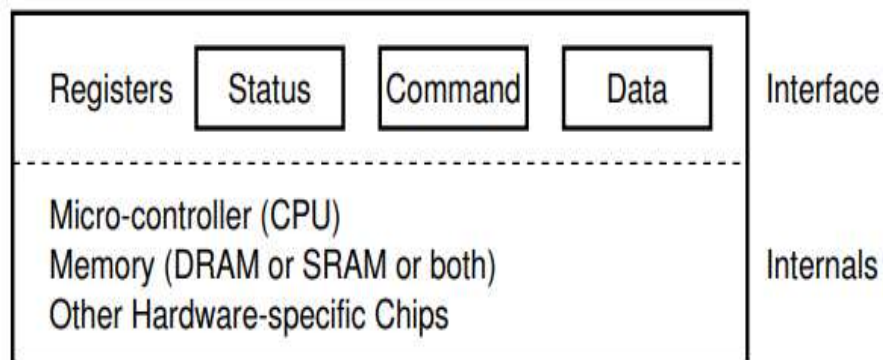


Figure 36.3: A Canonical Device

```
While (STATUS == BUSY)
    ; // wait until device is not busy
Write data to DATA register
Write command to COMMAND register
    (starts the device and executes the command)
While (STATUS == BUSY)
    ; // wait until device is done with your request
```

36.4 Lowering CPU overhead with Interrupt

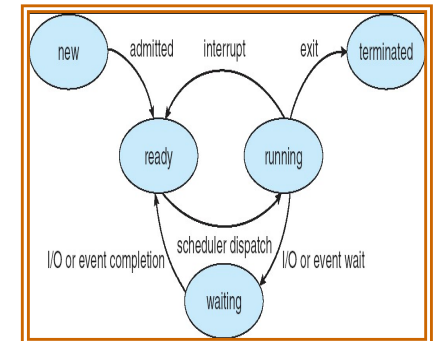

■ Interrupt vs Polling (PIO)

✓ Comparison

- Polling: An I/O thread checks status (busy or idle, like spin) → thread state: running (still hold CPU while its usage is only checking device status)
- **Interrupt**: An I/O thread is blocked, Notified when device is idle (or work is done) via interrupt → thread state: sleeping (CPU: used by other threads)
 - Note) Interrupt definition: a mechanism that informs an event to OS

✓ Example

- Thread 1 requests disk access (read or write)



✓ Tradeoffs

- Benefit of Interrupt: **overlapping**
 - Interrupt: CPU can do other useful job (for thread 2) while doing I/Os (for thread 1)
 - Polling: CPU just polling (actually waiting) while doing I/Os
- New requirement for Interrupt
 - Handling mechanism: call **interrupt handler** via **interrupt table** (see LN 2)
 - Sleep queue management (Context switch overhead)
- Usage suggestion (depend on devices)
 - Slow device: Interrupt, Fast device: Polling (like spin and sleep lock)
 - Optimization: Hybrid, Interrupt coalescing

36.5 More Efficient Data Movement with DMA



■ DMA (Direct Memory Access)

✓ Comparison

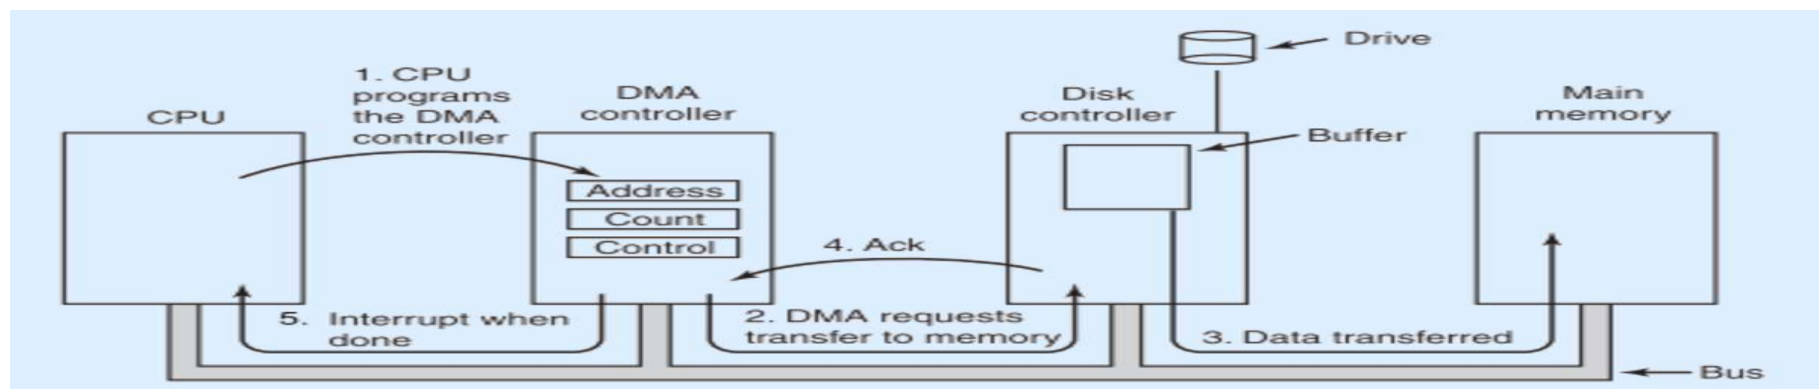
- PIO (Programmed I/O): CPU manages data copy between memory and devices
 - Concern: Devices are too slow for CPU (note CPU: ns, Disk: ms)
- DMA controller performs data copy between memory and devices
 - CPU can do other useful job (better overlapping)

✓ Example

- Thread 1 requests disk write without/with DMA using Interrupt
- Data copy (denoted as “c” in the figure is done by CPU vs. DMA)




✓ DMA mechanism




36.6. Methods of Device Interaction

- How to address registers in devices?
 - ✓ Two approaches
 - Direct I/O
 - Separated address space
 - Explicit I/O instruction (e.g. in/out + port)
 - Memory-mapped I/O
 - Single address space: DRAM + I/Os
 - Memory access instruction (e.g. load/store + I/O address space)
 - ✓ Privileged instruction for I/O space
 - Kernel mode: okay vs. User mode: protection fault (limited direct exec.)
 - Usually accessed in a kernel component called **device driver**



Direct I/O

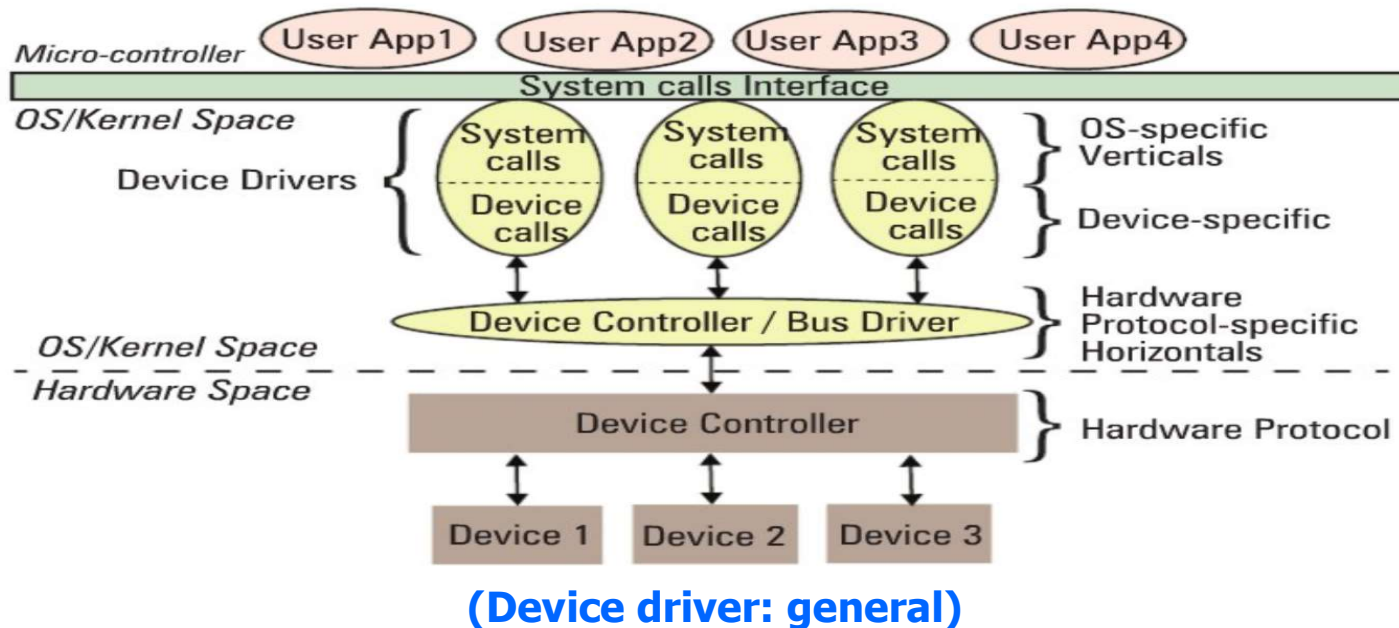


Memory-mapped I/O

36.7 Fitting into the OS: The Device Driver

■ Device driver

- ✓ A set of **software in kernel** that abstracts devices
- ✓ Two layers
 - Device-specific: Manage 1) device registers (command, status, data), 2) interrupt, and 3) DMA
 - OS-specific: Support generic interface such as open, read, write, close, ... (like **file**)
- ✓ 70% of codes in Linux is device drivers (mostly kernel module)



36.7 Fitting into the OS: The Device Driver

■ Two types

✓ Character device driver (or raw mode)

- Device accessed by user directly
 - System call → Driver → Devices
- Byte granularity (sequential access)

✓ Block device driver

- Device accessed by user through file system (FS)
 - System call → FS → Block layer (buffer cache, scheduler) → Driver → Devices
 - Also allow to be accessed in a raw mode
- Block granularity (random access)

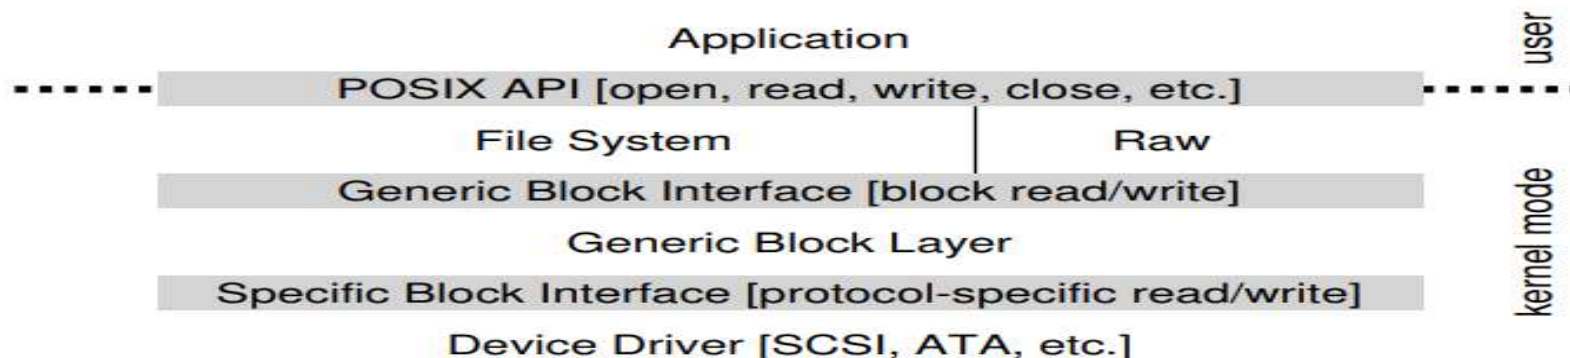
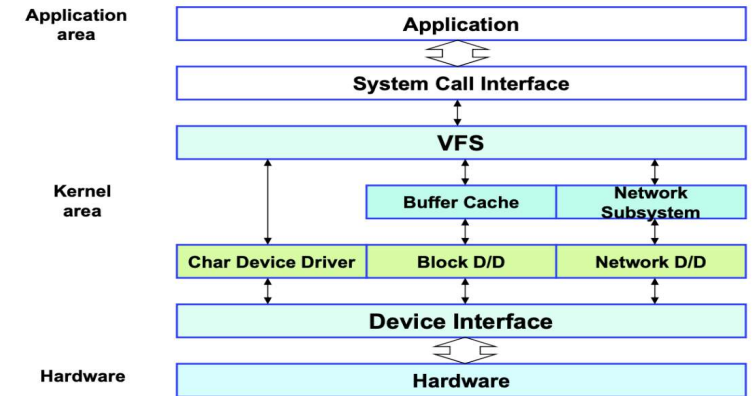


Figure 36.4: The File System Stack

36.8 Case Study: A Simple IDE Disk Driver (Optional)

- Device-specific layer (Device control)
 - ✓ Direct I/O (separated I/O address), I/O instruction: in/out
 - ✓ Registers
 - Control (0x3f6), Command block (0x1f2~1f6), Command or Status (0x1f7), Data port (0x1f0), Error (0x1f1)
 - Note) 1) **LBA**: Logical Block Address, 2) Status: Busy/Ready, 3) Error: bad block,...
 - Example of low-level interface (follow the device protocol)
 - Wait for drive ready: read 0x1f7 until the READY bit is on
 - Give info. about where and how much: write sector count and LBA in 0x1f2~1f6
 - Start I/O: write the WRITE command in 0x1f7
 - Data transfer: wait until READY and DRQ (Drive Request for Data), write data into the Data port
 - Handling errors: read 0x1f7 and 0x1f1

```
Control Register:
Address 0x3F6 = 0x08 (0000 1RE0): R=reset,
                E=0 means "enable interrupt"

Command Block Registers:
Address 0x1F0 = Data Port
Address 0x1F1 = Error
Address 0x1F2 = Sector Count
Address 0x1F3 = LBA low byte
Address 0x1F4 = LBA mid byte
Address 0x1F5 = LBA hi byte
Address 0x1F6 = 1B1D TOP4LBA: B=LBA, D=drive
Address 0x1F7 = Command/status

Status Register (Address 0x1F7):
 7 6 5 4 3 2 1 0
BUSY READY FAULT SEEK DRQ CORR IDDEX ERROR

Error Register (Address 0x1F1): (check when ERROR==1)
 7 6 5 4 3 2 1 0
BBK UNC MC IDNF MCR ABRT TONF AMNF

BBK = Bad Block
UNC = Uncorrectable data error
MC = Media Changed
IDNF = ID mark Not Found
MCR = Media Change Requested
ABRT = Command aborted
TONF = Track 0 Not Found
AMNF = Address Mark Not Found
```

Figure 36.5: The IDE Interface

36.8 Case Study: A Simple IDE Disk Driver (Optional)

■ OS specific layer (OS-level interface)

✓ Generic

- Character driver: open, read, write, close, intr, ...
- Block driver: open, close, intr, rw (or request, strategy), ...
 - Note: [dynamic loadable kernel module](#) interface for Linux (insmod, rmmod)

✓ IDE disk driver example: 4 main functions

- From FS: `ide_rw()` → `ide_start_request()` → `ide_wait_ready()`
- From disk: `ide_intr()`

```
static int ide_wait_ready() {
    while (((int r = inb(0x1f7)) & IDE_BSY) || !(r & IDE_DRDY))
        ; // loop until drive isn't busy
}

static void ide_start_request(struct buf *b) {
    ide_wait_ready();
    outb(0x3f6, 0); // generate interrupt
    outb(0x1f2, 1); // how many sectors?
    outb(0x1f3, b->sector & 0xff); // LBA goes here ...
    outb(0x1f4, (b->sector >> 8) & 0xff); // ... and here
    outb(0x1f5, (b->sector >> 16) & 0xff); // ... and here!
    outb(0x1f6, 0xe0 | ((b->dev&1)<<4) | ((b->sector>>24)&0x0f));
    if (b->flags & B_DIRTY) {
        outb(0x1f7, IDE_CMD_WRITE); // this is a WRITE
        outsl(0x1f0, b->data, 512/4); // transfer data too!
    } else {
        outb(0x1f7, IDE_CMD_READ); // this is a READ (no data)
    }
}


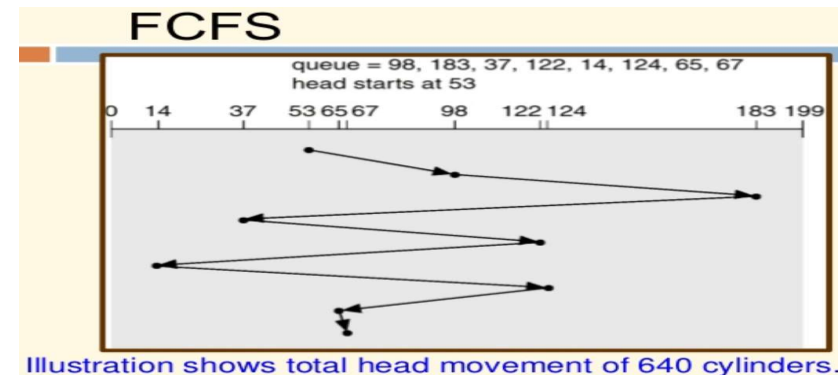

void ide_rw(struct buf *b) {
    acquire(&side_lock);
    for (struct buf **pp = &side_queue; *pp; pp=&(*pp)->qnext)
        ; // walk queue
    *pp = b; // add request to end
    if (ide_queue == b) // if q is empty
        ide_start_request(b); // send req to disk
    while ((b->flags & (B_VALID|B_DIRTY)) != B_VALID)
        sleep(b, &side_lock); // wait for completion
    release(&side_lock);
}

void ide_intr() {
    struct buf *b;
    acquire(&side_lock);
    if (!(b->flags & B_DIRTY) && ide_wait_ready() >= 0)
        insl(0x1f0, b->data, 512/4); // if READ: get data
    b->flags |= B_VALID;
    b->flags &= ~B_DIRTY;
    wakeup(b); // wake waiting process
    if ((ide_queue = b->qnext) != 0) // start next request
        ide_start_request(ide_queue); // (if one exists)
    release(&side_lock);
}
```

Figure 36.6: The xv6 IDE Disk Driver (Simplified)

37 Hard Disk Drives

- 37.1 The interface
- 37.2 Basic Geometry
- 37.3 A Simple Disk Drive
- 37.4 I/O Time: Doing the Math
- 37.5 Disk Scheduling



(Source: <https://www.slideshare.net/PareshParmar6/disk-scheduling-algorithms-71247712>)



37.1 The interface / 37.2 Basic Geometry

■ Basic Geometry

- ✓ Platter (two surface) → Track (thousands tracks per surface) → Sectors (512-byte)
- ✓ Head: sensing data
 - Multiple heads (one per each surface), connected into an arm
- ✓ Data access: seek time + rotation latency (time) + transfer time
 - Cylinder: a set of same tracks in each surface (no seek time required)

■ Interface: Addressing

- ✓ Sector addressing
 - Disk consists of a large number of sectors (0 ~ N-1 sectors or address space)
- ✓ LBA (Logical Block Address) addressing
 - Multi-sector addressing (usually called as a disk block): 4KB or 8KB
 - Kernel developer's viewpoint: **disk is a set of disk blocks** whose size is 4KB



37.3 A Simple Disk Drive

■ Scenarios

- ✓ Access in a same track: Figure 37.2
 - Assumptions: 1) Original head position is 6, target is 10, 2) 1) 12 sectors in a track, 3) 10,000 RPM (rotation per minute),
 - Total rotational latency: 10,000 RPM → 1/6 rotation per ms → a rotation takes 6ms
 - Rotational latency in this case: 2ms (3ms on average)
- ✓ Access in multiple tracks: Figure 37.3
 - Assumptions: 1) Original head position is 30, target is 11
 - Need not only rotational latency but also seek time (ms)
 - Note that seek and rotational latency perform in parallel
- ✓ Track skew: Figure 37.4
 - To optimize sequential access (e.g. read sector 10, 11, 12, 13)
 - Other optimizations: multi-zones, disk cache (track buffer)




Figure 37.2: A Single Track Plus A Head

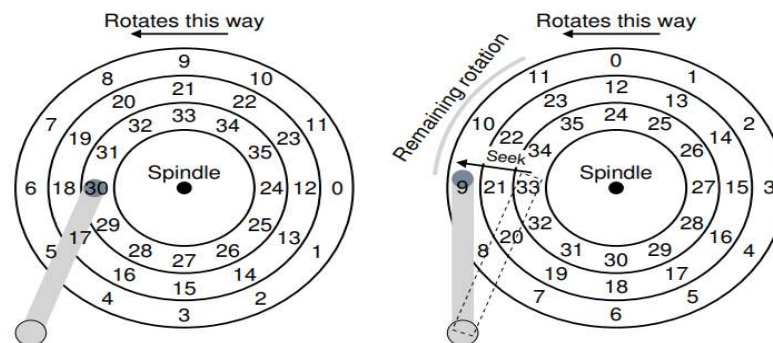


Figure 37.3: Three Tracks Plus A Head (Right: With Seek)




Figure 37.4: Three Tracks: Track Skew Of 2

37.4 I/O Time: Doing the Math

■ Metrics

✓ I/O time (latency) $T_{I/O} = T_{seek} + T_{rotation} + T_{transfer}$

✓ I/O rate (bandwidth, MB/s) $R_{I/O} = \frac{Size_{Transfer}}{T_{I/O}}$

■ Workload

- ✓ Random: issues small (e.g., 4KB) reads to random locations on disk
- ✓ Sequential: reads a large number of sectors consecutively (100 MB)

■ Disk considered: Figure 37.5

- ✓ Cheetah: a high-performance SCSI drive
- ✓ Barracuda: a drive built for capacity

| | Cheetah 15K.5 | Barracuda |
|--------------|---------------|-----------|
| Capacity | 300 GB | 1 TB |
| RPM | 15,000 | 7,200 |
| Average Seek | 4 ms | 9 ms |
| Max Transfer | 125 MB/s | 105 MB/s |
| Platters | 4 | 4 |
| Cache | 16 MB | 16/32 MB |
| Connects via | SCSI | SATA |

Figure 37.5: **Disk Drive Specs: SCSI Versus SATA**

37.4 I/O Time: Doing the Math

■ Metrics

✓ I/O time (latency)

$$T_{I/O} = T_{seek} + T_{rotation} + T_{transfer}$$

✓ I/O rate (bandwidth, MB/s)

$$R_{I/O} = \frac{Size_{Transfer}}{T_{I/O}}$$

■ Calculation (cheetah case)

✓ I/O rate calculation: 1) I/O component time, 2) I/O time, 3) I/O rate

✓ Random: Seek + Rotation + Transfer per 4KB

- I/O time: 4ms + 2ms (total rotational latency = $60 \cdot 1000 / 15000 \text{RPM} = 4\text{ms}$, hence 2ms on average) + 0.032ms ($4\text{KB} / 125\text{MB/s} = 4\text{KB} \cdot 1000 / 125 \cdot 1000\text{KB/ms}$)

- I/O rate: $4\text{KB} / 6\text{ms} = 0.66 \text{ MB/s}$

✓ Sequential: One seek/rotation per large data (e.g. 100MB)

- I/O time = $4\text{ms} + 2\text{ms} + 800\text{ms}(100\text{MB}/125\text{MB/s})$, I/O rate = $100\text{MB}/0.8\text{s}$

■ Implication

✓ Sequential is much faster than random in disk

✓ SW engineers need to make programs that **access disks in sequential**

| | Cheetah 15K.5 | Barracuda |
|--------------|---------------|-----------|
| Capacity | 300 GB | 1 TB |
| RPM | 15,000 | 7,200 |
| Average Seek | 4 ms | 9 ms |
| Max Transfer | 125 MB/s | 105 MB/s |
| Platters | 4 | 4 |
| Cache | 16 MB | 16/32 MB |
| Connects via | SCSI | SATA |

Figure 37.5: Disk Drive Specs: SCSI Versus SATA

| | Cheetah | Barracuda |
|----------------------|-----------|-----------|
| $R_{I/O}$ Random | 0.66 MB/s | 0.31 MB/s |
| $R_{I/O}$ Sequential | 125 MB/s | 105 MB/s |

Figure 37.6: Disk Drive Performance: SCSI Versus SATA

37.5 Disk Scheduling

- Disk scheduler
 - ✓ Role: Examines I/O requests and decides which one to schedule next
- Representative policies
 - ✓ FCFS (First Come First Serve)
 - Pros) simple, Cons) may cause long seek distance
 - ✓ SSTF (Shortest Seek Time First)
 - Pros) reduce seek distance, Cons) unfair (especially boundary tracks)
 - ✓ SCAN (a.k.a. Elevator) and C-SCAN
 - Moves back and forth across all tracks
 - C-SCAN: handle requests from inner-to-outer, then go back inner tracks directly and handling requests again from inner-to-outer (or reverse)





Figure 37.7: SSTF: Scheduling Requests 21 And 2

37.5 Disk Scheduling

■ Representative policies (cont')

✓ SPTF (Shortest Positioning Time First)

- Consider seek and rotation latency
- Why? Issues that consider seek only → not optimal (Figure 37.8)
 - Head position: 30 (sector), Next requests: 16 and 8
 - SSTF: 16 and then 8 → 1 seek + 5/6 rotation + 1 seek + 2/6 rotation
 - How about 8 and then 16 → 1 seek (relatively further) + 1/6 rotation + 1 seek + 4/6 rotation
 - Performance depends on disk characteristics (seek vs. rotation)
- SPTF select a request who has the smallest position time (seek + rotation time)

■ Other scheduling issues

- ✓ Merge: requests 33, 4, 34, ...
- ✓ Anticipatory disk scheduling

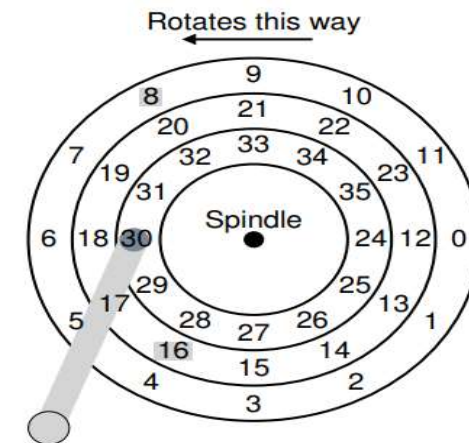


Figure 37.8: SSTF: Sometimes Not Good Enough


Chap. 39 Interlude: Files and Directories

- 39.1 Files and Directories
- 39.2 File System Interface
- 39.3 Creating Files
- 39.4 Reading and Writing Files
- 39.5 Reading and Writing, But **Not Sequentially**
- 39.6 Shared file table entries: `fork()` and `dup()`
- 39.7 Writing immediately with `fsync()`
- 39.8 Renaming files
- 39.9 Getting information about files
- 39.10 Removing files
- 39.11 Making **Directories**
- 39.12 Reading Directories
- 39.13 Deleting Directories
- 39.14 Hard **Links**
- 39.15 Symbolic Links
- 39.16 Permission Bits and Access Control Lists
- 39.17 Making and **Mounting a file system**

Chap. 39 Interlude: Files and Directories

■ Computer system

- ✓ Four key abstractions: process (thread), virtual memory, lock, and file
- ✓ Files are in Storage (Hard disk, Solid State Drive)
 - Storage vs. Memory
 - Non-volatility
 - Advantages: Support persistence (store information permanently)
 - Issues: 1) Integrity, 2) Space-efficiency, 3) Consistency, 4) Crash consideration (fault-tolerance), 5) Access control, 6) Security, ...
 - These issues are managed by a file system
- ✓ How to explore file system?
 - **Interface**: open, read, write, close, mkdir, link, mount, ... (Chapter 39)
 - **Layout**: file, directory, inode, FAT, superblock, ... (Chapter 40)



39.1 Files and Directories

■ File

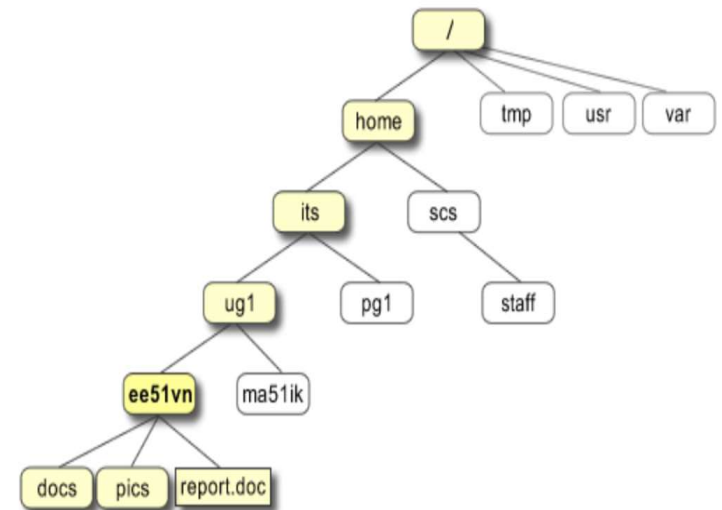
- ✓ Definition: A linear array of characters (bytes), stored persistently
 - Each file has various data structure (text, c code, record, multimedia, ...)
 - But, OS don't care its content, just treating it as a stream of bytes
- ✓ Each file has its name (absolute path, relative path)
- ✓ It also has some kind of low-level name in OS (e.g. [inode](#))
 - Like each process has a unique PCB

■ Directory

- ✓ A special file that constructs a [hierarchy](#) (file hierarchy)
 - Root directory
 - Home directory
 - Working directory
- ✓ Contain <file name, inode>
 - or low-level name or first disk block

■ Others are also treated as a file

- ✓ Device, pipe, socket, and even process



39.2 File System Interfaces

■ APIs

- ✓ System call: 1) open (return a **file descriptor**), 2) name resolution (directory hierarchy traverse), 3) special file, 4) attribute, 5) I/O, 6) file system management, 7) directory management, ...
- ✓ Internals: 1) allocate/free block, 2) allocate/free inode, 3) namei (name-to-inode), 4) buffer related

Filesystem system calls

| Filesystem system calls | | | | | | | |
|---------------------------------------|--|--|----------------------------------|------------------------|------------------------|----------------------------------|-----------------|
| Return a descriptor | Use namei | | Allocate inode | Attributes | I/O | File System Structure Management | |
| open creat dup pipe close | open creat chdir chroot chown chmod | stat link unlink mknod mount umount | creat mknod link unlink | chown chmod stat | read write lseek | mount umount | chdir chroot |

| Filesystem low level functions | | | | | | |
|--------------------------------|------|------|-------|------|--------|-------|
| namei | | | alloc | free | ialloc | ifree |
| iget | iput | bmap | | | | |

| buffer allocation algorithms | | | | |
|------------------------------|--------|-------|--------|--------|
| getblk | brelse | bread | breada | bwrite |

(Source: <http://slideplayer.com/slide/9118590/>)

39.3 Creating Files / 39.4 Reading and Writing Files

■ Create API

- ✓ `open()` with create flag (refer to LN1 or Figure 2.6 io.c in OSTEP)


```
int fd = open("foo", O_CREAT|O_WRONLY|O_TRUNC, S_IRUSR|S_IWUSR);
```

- Arguments: 1) name, 2) flags, 3) permissions
 - Return: fd (file descriptor)
- ✓ `creat()`: less used (but famous by Ken Thompson's answer about redesigning UNIX)

```
int fd = creat("foo"); // option: add second flag to set permissions
```

■ Read/Write API

- ✓ `read_size = read(fd, buf, request_size);`
- ✓ `written_size = write(fd, buf, request_size);`
 - Arguments: 1) fd, 2) buffer that points memory space for data, 3) request size
 - Return: read or written size



39.4 Reading and Writing Files

■ Read and write example

- ✓ Command line viewpoint

```
prompt> echo hello > foo
prompt> cat foo
hello
prompt>
```

- ✓ System call viewpoint (using strace)

```
prompt> strace cat foo
...
open("foo", O_RDONLY|O_LARGEFILE)      = 3
read(3, "hello\n", 4096)                = 6
write(1, "hello\n", 6)                  = 6
hello
read(3, "", 4096)                       = 0
close(3)                                 = 0
...
prompt>
```


TIP: USE `strace` (AND SIMILAR TOOLS)

The `strace` tool provides an awesome way to see what programs are up to. By running it, you can trace which system calls a program makes, see the arguments and return codes, and generally get a very good idea of what is going on.

The tool also takes some arguments which can be quite useful. For example, `-f` follows any fork'd children too; `-t` reports the time of day at each call; `-e trace=open,close,read,write` only traces calls to those system calls and ignores all others. There are many more powerful flags — read the man pages and find out how to harness this wonderful tool.

39.5 Reading and Writing, But Not Sequentially

- Conventional accessing mechanism for a file
 - ✓ Sequential
 - ✓ From the begin, increasing the **offset** while reading or writing



- How to access random position? (not sequentially)
 - ✓ lseek()
 - Arguments: 1) fd, 2) relative offset from whence, 3) reference point

```
off_t lseek(int fildes, off_t offset, int whence);
```

- Whence: SEEK_SET, SEEK_CUR, SEEK_END
- Explicit update the current offset (c.f. read/write: implicit update)
- Do not confuse lseek() with disk seek :-)

☞ **Also do not confuse process and processor**

39.6 Shared File Table Entries: fork() and dup()

■ file table sharing

- ✓ fork case: shared by different PCBs
- ✓ dup case: shared by different fd entries in the same PCB

```
int main(int argc, char *argv[]) {
    int fd = open("file.txt", O_RDONLY);
    assert(fd >= 0);
    int rc = fork();
    if (rc == 0) {
        rc = lseek(fd, 10, SEEK_SET);
        printf("child: offset %d\n", rc);
    } else if (rc > 0) {
        (void) wait(NULL);
        printf("parent: offset %d\n",
              (int) lseek(fd, 0, SEEK_CUR));
    }
    return 0;
}
```

Figure 39.2: Shared Parent/Child File Table Entries (**fork-seek.c**)

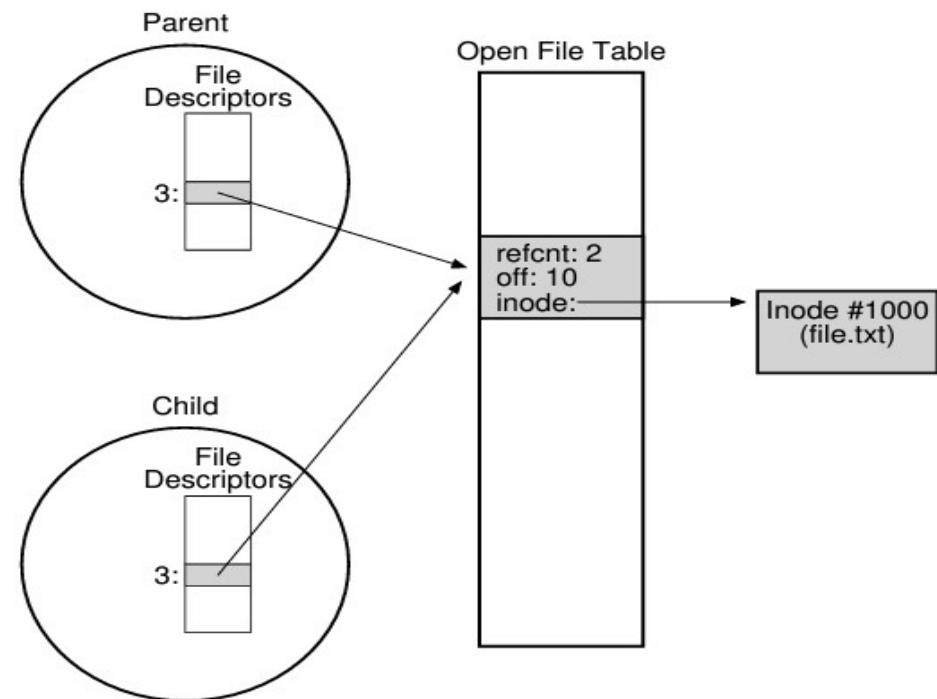


Figure 39.3: Processes Sharing An Open File Table Entry

```
int main(int argc, char *argv[]) {
    int fd = open("README", O_RDONLY);
    assert(fd >= 0);
    int fd2 = dup(fd);
    // now fd and fd2 can be used interchangeably
    return 0;
}
```


Figure 39.4: Shared File Table Entry With dup () (**dup.c**)

```
prompt> ./fork-seek
child: offset 10
parent: offset 10
prompt>
```

39.7 Writing Immediately with fsync()

■ Performance consideration for write

- ✓ Write to DRAM vs Disk: 100ns vs 10,000,000ns (10ms)
- ✓ Delayed write
 - Write data into DRAM (called **buffer** or **page cache**) and set them **dirty**
 - Later write all dirty data into disk in a batch manner (5 or 30 seconds periodically)
 - Write grouping and write reordering indeed enhance performance
 - **Synchronous** vs. **Asynchronous I/O**



■ Concern of delayed write

- ✓ Durability
 - User think his/her data is permanent but not in actuality
- ✓ How to guarantee durability
 - `fsync()` system call

```
int fd = open("foo", O_CREAT|O_WRONLY|O_TRUNC, S_IRUSR|S_IWUSR);
assert(fd > -1);
int rc = write(fd, buffer, size);
assert(rc == size);
rc = fsync(fd);
assert(rc == 0);
```

👉 **Different failure semantic before/after fsync()**

39.8 Renaming Files / 39.10 Removing Files

■ Change a file name

- ✓ Command line viewpoint

```
prompt> mv foo bar
```

- ✓ API (system call) viewpoint: editor example

```
int fd = open("foo.txt.tmp", O_WRONLY|O_CREAT|O_TRUNC,  
             S_IRUSR|S_IWUSR);  
write(fd, buffer, size); // write out new version of file  
fsync(fd);  
close(fd);  
rename("foo.txt.tmp", "foo.txt");
```

- rename(old name, new name)
- conducted atomically

■ Remove a file

- ✓ API

- unlink(file name)

```
prompt> strace rm foo  
...  
unlink("foo")           = 0  
...
```

☞ Why not remove() or delete() instead of unlink()? Then, what is link()?

39.9 Getting Information about Files

■ Contents in a file system

- ✓ Two types of data in file system: User data vs. Metadata
 - User data (or just data): data written by users
 - **Metadata**: data written by a file system for managing files (in inode) and file system (in superblock)
- ✓ API to see the metadata for a certain file
 - `stat(file_name, struct stat)`
 - `fstat(fd, struct stat)`

```
struct stat {
    dev_t      st_dev;      /* ID of device containing file */
    ino_t      st_ino;     /* inode number */
    mode_t     st_mode;    /* protection */
    nlink_t    st_nlink;   /* number of hard links */
    uid_t      st_uid;     /* user ID of owner */
    gid_t      st_gid;     /* group ID of owner */
    dev_t      st_rdev;    /* device ID (if special file) */
    off_t      st_size;    /* total size, in bytes */
    blksize_t  st_blksize; /* blocksize for filesystem I/O */
    blkcnt_t   st_blocks;  /* number of blocks allocated */
    time_t     st_atime;   /* time of last access */
    time_t     st_mtime;   /* time of last modification */
    time_t     st_ctime;   /* time of last status change */
};
```

```
prompt> echo hello > file
prompt> stat file
  File: `file'
  Size: 6 Blocks: 8          IO Block: 4096   regular file
Device: 811h/2065d Inode: 67158084    Links: 1
Access: (0640/-rw-r-----) Uid: (30686/  remzi) Gid: (30686/  remzi)
Access: 2011-05-03 15:50:20.157594748 -0500
Modify: 2011-05-03 15:50:20.157594748 -0500
Change: 2011-05-03 15:50:20.157594748 -0500
```

39.11 Making Directories / 39.13 Deleting Directories

■ API for making directory

- ✓ mkdir(name, permission)

```
prompt> strace mkdir foo
...
mkdir("foo", 0777) = 0
...
prompt>
```

- ✓ After making

- Two entries: parent directory and itself

```
prompt> ls -a
./ ../
prompt> ls -al
total 8
drwxr-x---  2 remzi remzi   6 Apr 30 16:17 ./
drwxr-x--- 26 remzi remzi 4096 Apr 30 16:17 ../
```

■ API for deleting directory

- ✓ rmdir(file_name)
- ✓ We need to use it carefully

39.12 Reading Directories

■ APIs for reading directory

- ✓ opendir(dp), readdir(dp), closedir(dp)
- ✓ “ls”: like the below example (c.f. “ls -l”: readdir() + stat())

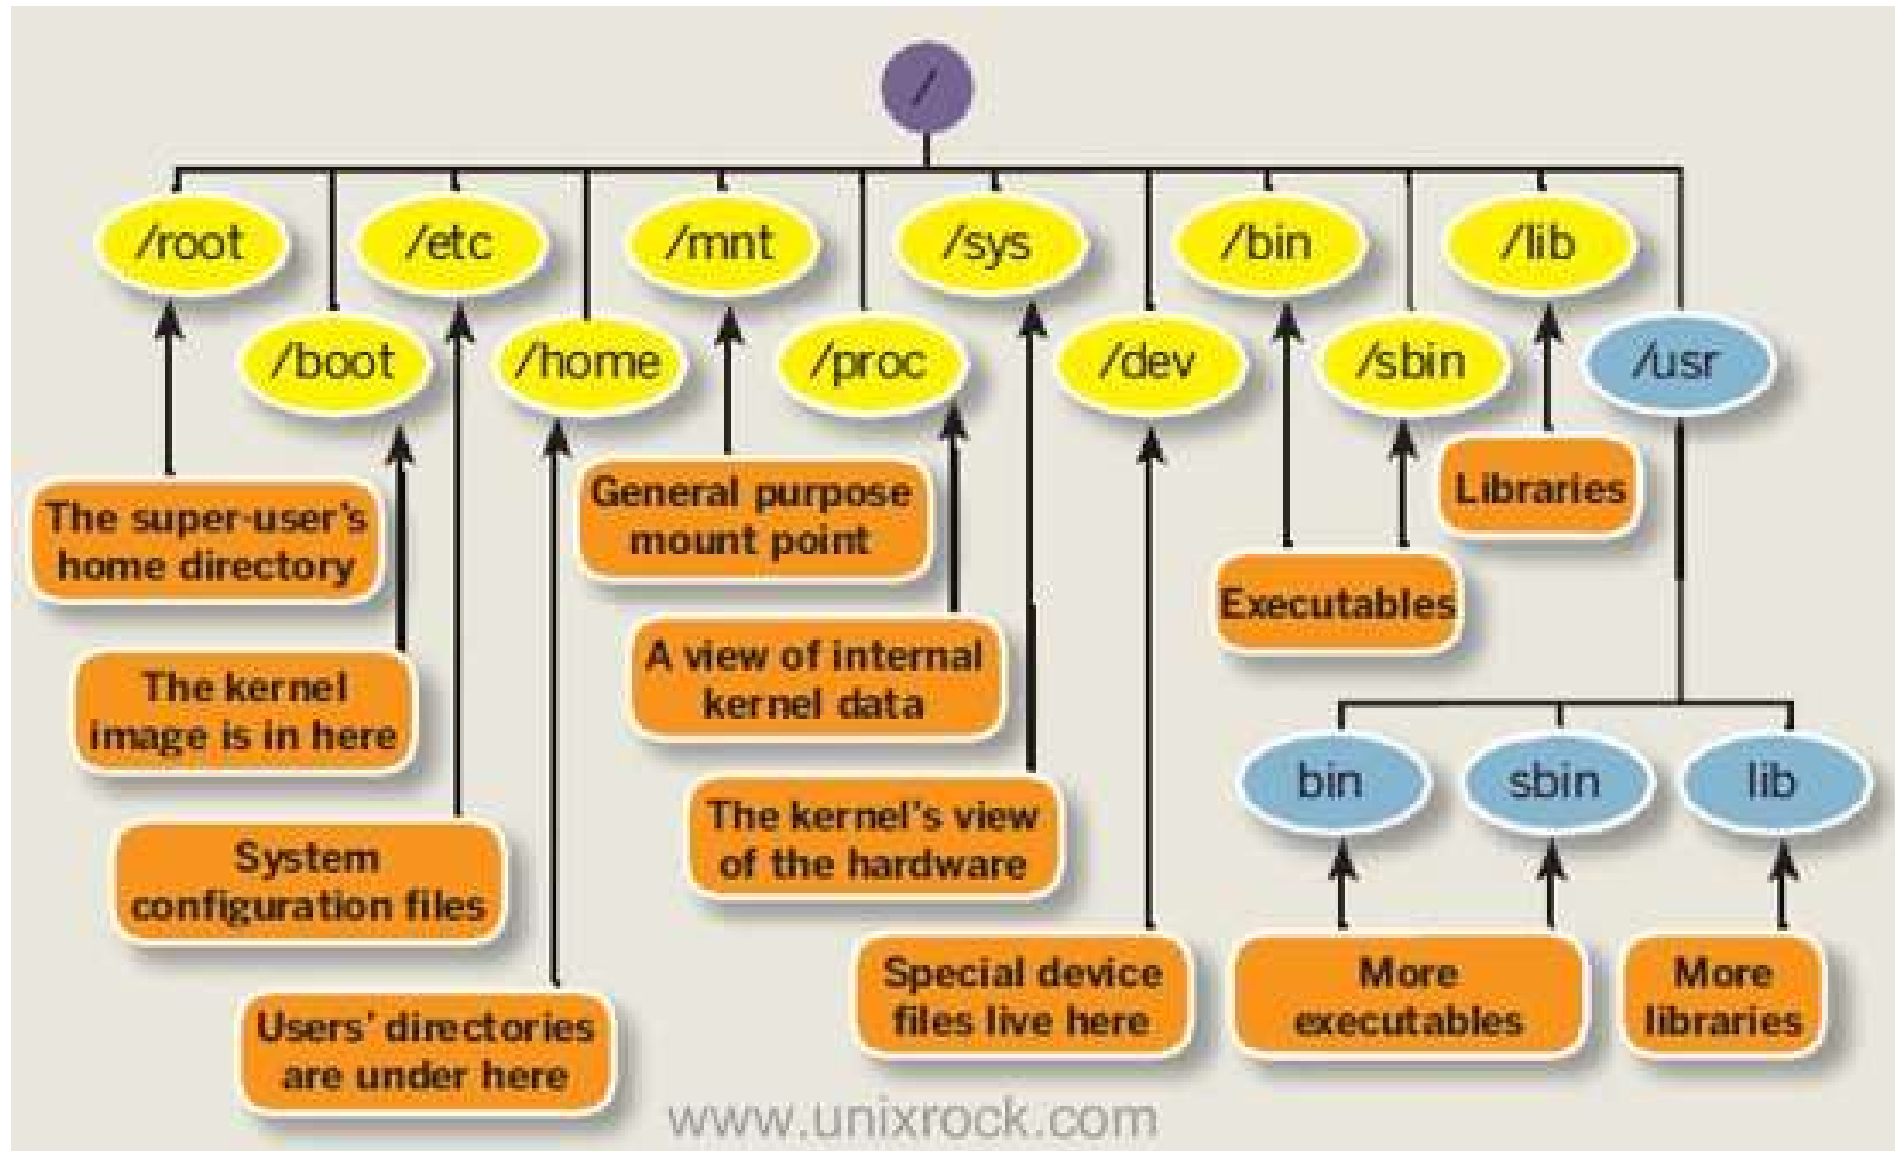
```
int main(int argc, char *argv[]) {
    DIR *dp = opendir(".");
    assert(dp != NULL);
    struct dirent *d;
    while ((d = readdir(dp)) != NULL) {
        printf("%d %s\n", (int) d->d_ino, d->d_name);
    }
    closedir(dp);
    return 0;
}

struct dirent {
    char          d_name[256]; /* filename */
    ino_t         d_ino;       /* inode number */
    off_t         d_off;       /* offset to the next dirent */
    unsigned short d_reclen;   /* length of this record */
    unsigned char d_type;      /* type of file */
};
```

👉 Why there is no writedir()?

39.12 Reading Directories

- Directory name convention



(Source: <http://www.unixrock.com/2013/04/solaris-directory-hierarchy.html>)

39.14 Hard Links

■ Link

- ✓ Make another file name to access an existing file

- Connect [a file name with an inode](#)

- ✓ Command line viewpoint

- Either file or file2

```
prompt> echo hello > file
prompt> cat file
hello
prompt> ln file file2
prompt> cat file2
hello
```

```
prompt> ls -i file file2
67158084 file
67158084 file2
prompt>
```

- ✓ API

- `link(old_name, new_name)`

- ✓ After remove one of them

- Use `unlink()`
- Still remain data

```
prompt> rm file
removed 'file'
prompt> cat file2
hello
```

```
prompt> echo hello > file
prompt> stat file
... Inode: 67158084    Links: 1 ...
prompt> ln file file2
prompt> stat file
... Inode: 67158084    Links: 2 ...
prompt> stat file2
... Inode: 67158084    Links: 2 ...
prompt> ln file2 file3
prompt> stat file
... Inode: 67158084    Links: 3 ...
prompt> rm file
prompt> stat file2
... Inode: 67158084    Links: 2 ...
prompt> rm file2
prompt> stat file3
... Inode: 67158084    Links: 1 ...
prompt> rm file3
```

- ✓ Link count

- Delete data when link count is 0

39.15 Symbolic Links

■ Link

- ✓ Hard link: share inode number
 - Create a new file name and share the existing inode
- ✓ Symbolic link (Soft link): different inode number, but its data is the linked file name
 - Create not only a new file name but also a new inode (set it as a symbolic link)
 - Can link between different file systems, Can link to a directory
- ✓ Dangling reference in symbolic link

```
oslab@osLab: ~/os_test
oslab@osLab:~/os_test$ ls -al
합계 8
drwxrwxr-x  2 oslab oslab 4096  4월 23 12:09 .
drwxr-xr-x 22 oslab oslab 4096  4월 23 12:03 ..
oslab@osLab:~/os_test$
oslab@osLab:~/os_test$ echo "hello world" > file1
oslab@osLab:~/os_test$ ln file1 file2
oslab@osLab:~/os_test$ ln -s file1 file3
oslab@osLab:~/os_test$
oslab@osLab:~/os_test$ ls -ail
합계 16
8297 drwxrwxr-x  2 oslab oslab 4096  4월 23 12:09 .
8196 drwxr-xr-x 22 oslab oslab 4096  4월 23 12:03 ..
 380 -rw-rw-r--  2 oslab oslab   12  4월 23 12:09 file1
 380 -rw-rw-r--  2 oslab oslab   12  4월 23 12:09 file2
 386 lrwxrwxrwx  1 oslab oslab    5  4월 23 12:09 file3 -> file1
oslab@osLab:~/os_test$
oslab@osLab:~/os_test$ rm file1
oslab@osLab:~/os_test$
oslab@osLab:~/os_test$ cat file2
hello world
oslab@osLab:~/os_test$ cat file3
cat: file3: 그런 파일이나 디렉터리가 없습니다
oslab@osLab:~/os_test$
oslab@osLab:~/os_test$ ls -li
380 file2 386 file3
oslab@osLab:~/os_test$
```

Permission bits: 3 categories (u/g/o), 3 bits (r/w/x)

39.17 Making and Mounting a File System

■ File system

✓ Make a file system


- Assemble directories and files
- Related metadata: [superblock](#), [bitmap](#), ... (main topic in chapter 40)
- Command: [mkfs](#)
 - Make an empty file system (only root directory) in a disk partition

```
howtogeek@ubuntu: ~  
howtogeek@ubuntu:~$ sudo mkfs  
mkfs      mkfs.ext2      mkfs.ext4dev  mkfs.ntfs  
mkfs.bfs  mkfs.ext3      mkfs.minix    mkfs.vfat  
mkfs.cramfs  mkfs.ext4      mkfs.msdos  
howtogeek@ubuntu:~$ sudo mkfs.ext4 /dev/sda5
```

- How to make partitions?: [fdisk](#)

✓ Example

- Partitioning and mkfs
- Ext2/3/4, NFS, LFS, proc, sysfs, ... per a partition



```
pritam@PPH:~$ sudo fdisk -l  
[sudo] password for pritam:  
Disk /dev/sda: 500.1 GB, 500107862016 bytes  
255 heads, 63 sectors/track, 60801 cylinders, total 976773168 sectors  
Units = sectors of 1 * 512 = 512 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk identifier: 0x1707a8a5  


| Device    | Boot | Start     | End       | Blocks    | Id | System               |
|-----------|------|-----------|-----------|-----------|----|----------------------|
| /dev/sda1 | *    | 2048      | 97656831  | 48827392  | 7  | HPFS/NTFS/exFAT      |
| /dev/sda2 |      | 97656832  | 293177343 | 97760256  | 7  | HPFS/NTFS/exFAT      |
| /dev/sda3 |      | 293177344 | 879116287 | 292969472 | 7  | HPFS/NTFS/exFAT      |
| /dev/sda4 |      | 879118334 | 976773119 | 48827393  | 5  | Extended             |
| /dev/sda5 |      | 879118336 | 888881151 | 4881408   | 82 | Linux swap / Solaris |
| /dev/sda6 |      | 888883200 | 976773119 | 43944960  | 83 | Linux                |


  
pritam@PPH:~$
```

39.17 Making and Mounting a File System

■ File system

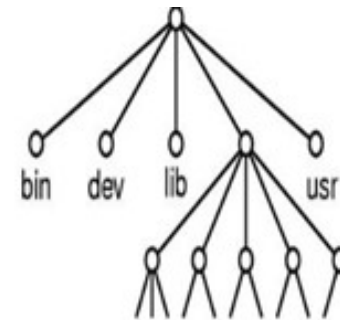
✓ Mount

- Make a file system visible to users
- Connect multiple file systems within the uniform directory tree
 - mount arguments: 1) FS type, 2) partition, 3) mount point



Before mount

```
$mount -t ext3 /dev/sda4 /mnt
```



After mount

- mount point: mnt in the previous example → point the root of the mounted FS

☞ **Why multiple partitions?**

Chap. 40 File System Implementation

■ Objective of this chapter

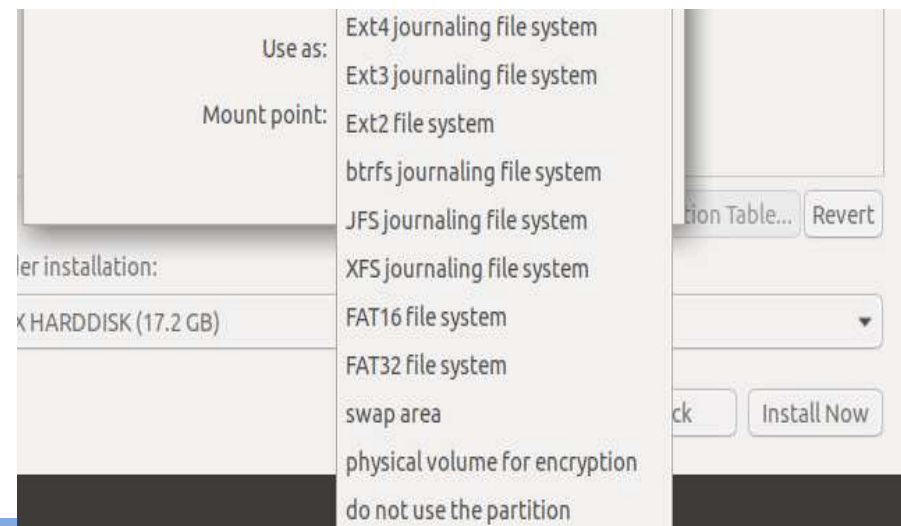

✓ A variety of file systems

- UFS, FFS, EXT2/3/4, JFS, LFS, NTFS, F2FS, FUSE, RAMFS, NFS, AFS, ZFS, GFS, FATFS, Btrfs,

✓ Make a new file system: called VSFS(Very Simple File System)

- Simplified version of UFS (Unix File System)
- 1) On-disk structures: inode, bitmap, directory, ...
- 2) Access method: read, write, ...
- 3) Various policies: cache, delayed write, ...


✓ More complex file systems → next chapters



40.1 The Way to Think / 40.2 Overall Organization

■ Disk

- ✓ Consist of partitions
- ✓ A file system is created in each partition




```
Disk /dev/sda: 320.1 GB, 32007293376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0xaa692010
```

| Device | Boot | Start | End | Blocks | Id | System |
|-----------|------|-------|-------|-----------|----|----------------------|
| /dev/sda1 | * | 1 | 7681 | 61690880 | 7 | HPFS/NTFS |
| /dev/sda2 | | 7681 | 14182 | 52219904 | 7 | HPFS/NTFS |
| /dev/sda3 | | 14182 | 20556 | 51200000 | 7 | HPFS/NTFS |
| /dev/sda4 | | 20556 | 38913 | 147453953 | f | W95 Ext'd (LBA) |
| /dev/sda5 | | 20556 | 32030 | 92160000 | 7 | HPFS/NTFS |
| /dev/sda6 | | 34324 | 34770 | 3583999+ | 82 | Linux swap / Solaris |
| /dev/sda7 | | 34771 | 38913 | 33276928 | 7 | HPFS/NTFS |
| /dev/sda8 | | 32030 | 34323 | 18423808 | 83 | Linux |

Partition table entries are not in disk order

■ Partition

- ✓ Consist of disk blocks
- ✓ User data is stored in a disk block (usually **same size with the page**)
- ✓ Assume a partition having 64 disk blocks (or simply blocks)

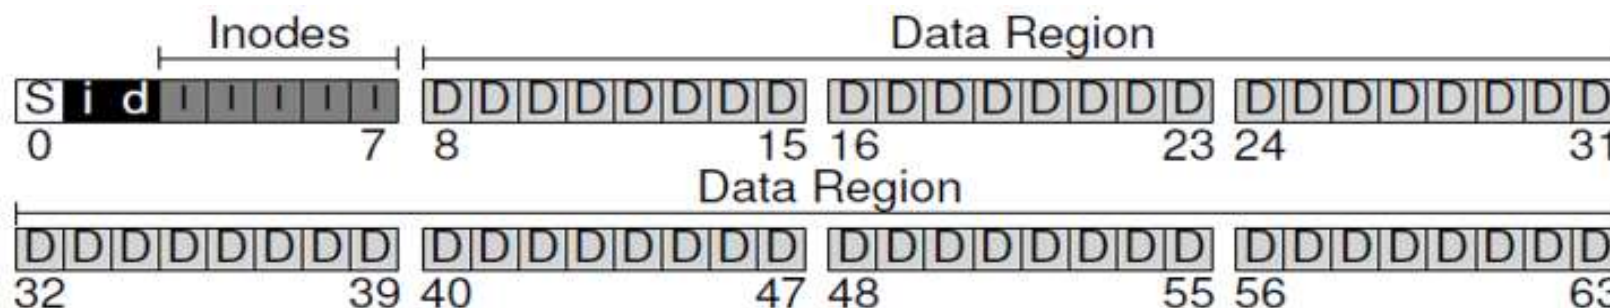


👉 **Now consider what data structures are required for making a FS?**

40.2 Overall Organization

■ Layout of a file system (VSFS)

- ✓ Superblock: 0 blocks
 - Metadata for managing a file system (**one per a file system**)
 - Information: how many data blocks, inodes, where they begin, ...
 - Used during a mount function
- ✓ Bitmap: 1~2 blocks
 - Metadata for managing free space (allocation structure)
 - Two bitmaps: one for data blocks and the other for inodes
- ✓ Inode: 3~7 blocks
 - Metadata for managing files (**one per a file**)
 - Inode size = 256B → 16 inodes per a block → 5 blocks for inode → total 80 files can be created
- ✓ User data: 8 ~ 63 blocks (can be dynamically adjusted)
 - Data written by users

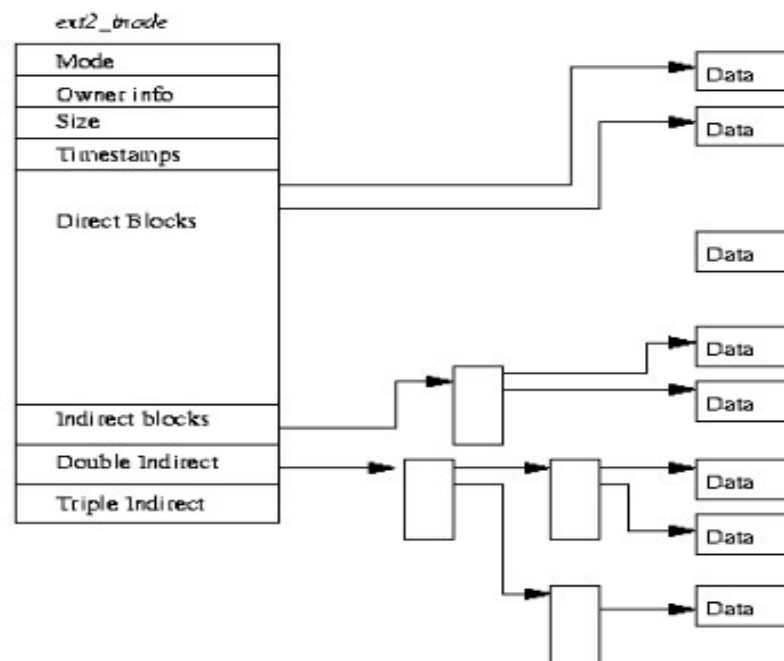


40.3 File Organization: The inode

- How to manage metadata for a file
 - ✓ inode (index node)
 - File information such as mode, uid, size, time, link count, blocks, ...
 - Can be accessed using stat()
 - Locations of User data blocks: **Multi-Level index** (a.k.a **Imbalanced tree**)
 - Direct block pointers (10 or 12 or 15), Single/Double/Triple indirect block pointers(1/1/1)

| Size | Name | What is this inode field for? |
|------|-------------|---|
| 2 | mode | can this file be read/written/executed? |
| 2 | uid | who owns this file? |
| 4 | size | how many bytes are in this file? |
| 4 | time | what time was this file last accessed? |
| 4 | ctime | what time was this file created? |
| 4 | mtime | what time was this file last modified? |
| 4 | dtime | what time was this inode deleted? |
| 2 | gid | which group does this file belong to? |
| 2 | links_count | how many hard links are there to this file? |
| 4 | blocks | how many blocks have been allocated to this file? |
| 4 | flags | how should ext2 use this inode? |
| 4 | osd1 | an OS-dependent field |
| 60 | block | a set of disk pointers (15 total) |
| 4 | generation | file version (used by NFS) |
| 4 | file_acl | a new permissions model beyond mode bits |
| 4 | dir_acl | called access control lists |

Figure 40.1: Simplified Ext2 Inode

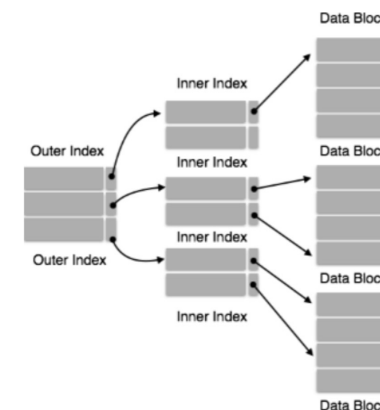


👉 **How large size can be supported by direct block pointers? How about an indirect pointer?**

40.3 File Organization: The inode

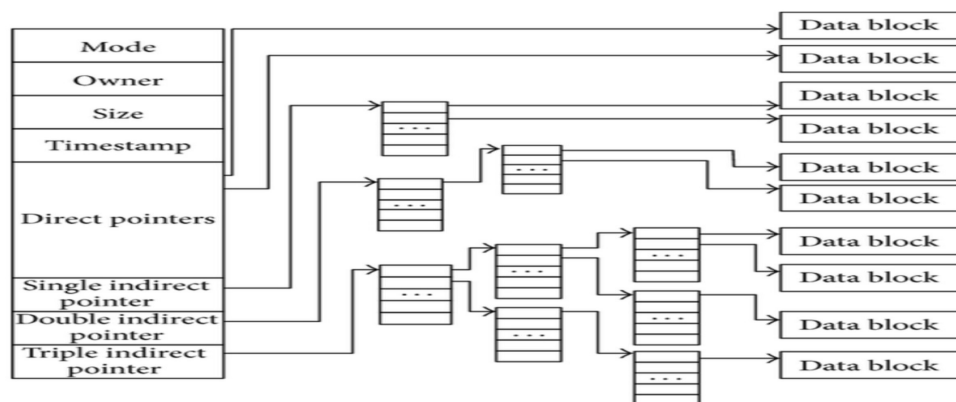
■ Maximum file size supported by an inode

- ✓ Sum up: 48KB + 4MB + 4GB + 4TB
 - Direct block point: 12 x 4KB
 - Single indirect block pointers: 1 x 1024 x 4KB
 - Why 1024: $4\text{KB} / \text{pointer size} = 4\text{KB}/4\text{B} = 1024$
 - Double indirect block pointers: 1 x 1024 x 1024 x 4KB
 - Triple indirect block pointers: 1 x 1024 x 1024 x 1024 x 4KB
- ✓ Benefits of imbalance tree (compare to balanced tree)
 - Small file: **fast** (direct access via direct pointer)
 - Large file: **support large size with the simple structure** of inode



■ Other approach

- ✓ FAT (linked based), Extent-based, Log-based, ..




Most files are small
 Average file size is growing
 Most bytes are stored in large files
 File systems contains lots of files
 File systems are roughly half full
 Directories are typically small

Roughly 2K is the most common size
 Almost 200K is the average
 A few big files use most of the space
 Almost 100K on average
 Even as disks grow, file systems remain ~50% full
 Many have few entries; most have 20 or fewer

Figure 40.2: File System Measurement Summary


40.3 File Organization: The inode

- inode manipulation example (assume 12 direct blocks)
 - ✓ When we create a new file (named hello.c whose size is 7KB) in a root directory?
 - ✓ Then, we compile it? (a.out whose size is 70KB)



| iblock 0 | | | | iblock 1 | | | | iblock 2 | | | | iblock 3 | | | | iblock 4 | | | |
|----------|----|----|----|----------|----|----|----|----------|----|----|----|----------|----|----|----|----------|----|----|----|
| 0 | 1 | 2 | 3 | 16 | 17 | 18 | 19 | 32 | 33 | 34 | 35 | 48 | 49 | 50 | 51 | 64 | 65 | 66 | 67 |
| 4 | 5 | 6 | 7 | 20 | 21 | 22 | 23 | 36 | 37 | 38 | 39 | 52 | 53 | 54 | 55 | 68 | 69 | 70 | 71 |
| 8 | 9 | 10 | 11 | 24 | 25 | 26 | 27 | 40 | 41 | 42 | 43 | 56 | 57 | 58 | 59 | 72 | 73 | 74 | 75 |
| 12 | 13 | 14 | 15 | 28 | 29 | 30 | 31 | 44 | 45 | 46 | 47 | 60 | 61 | 62 | 63 | 76 | 77 | 78 | 79 |

Note: 23 is the index block while others (22, 24, ...) are data blocks



- ./: 0
- ../: 0
- hello.c: 1
- a.out: 2

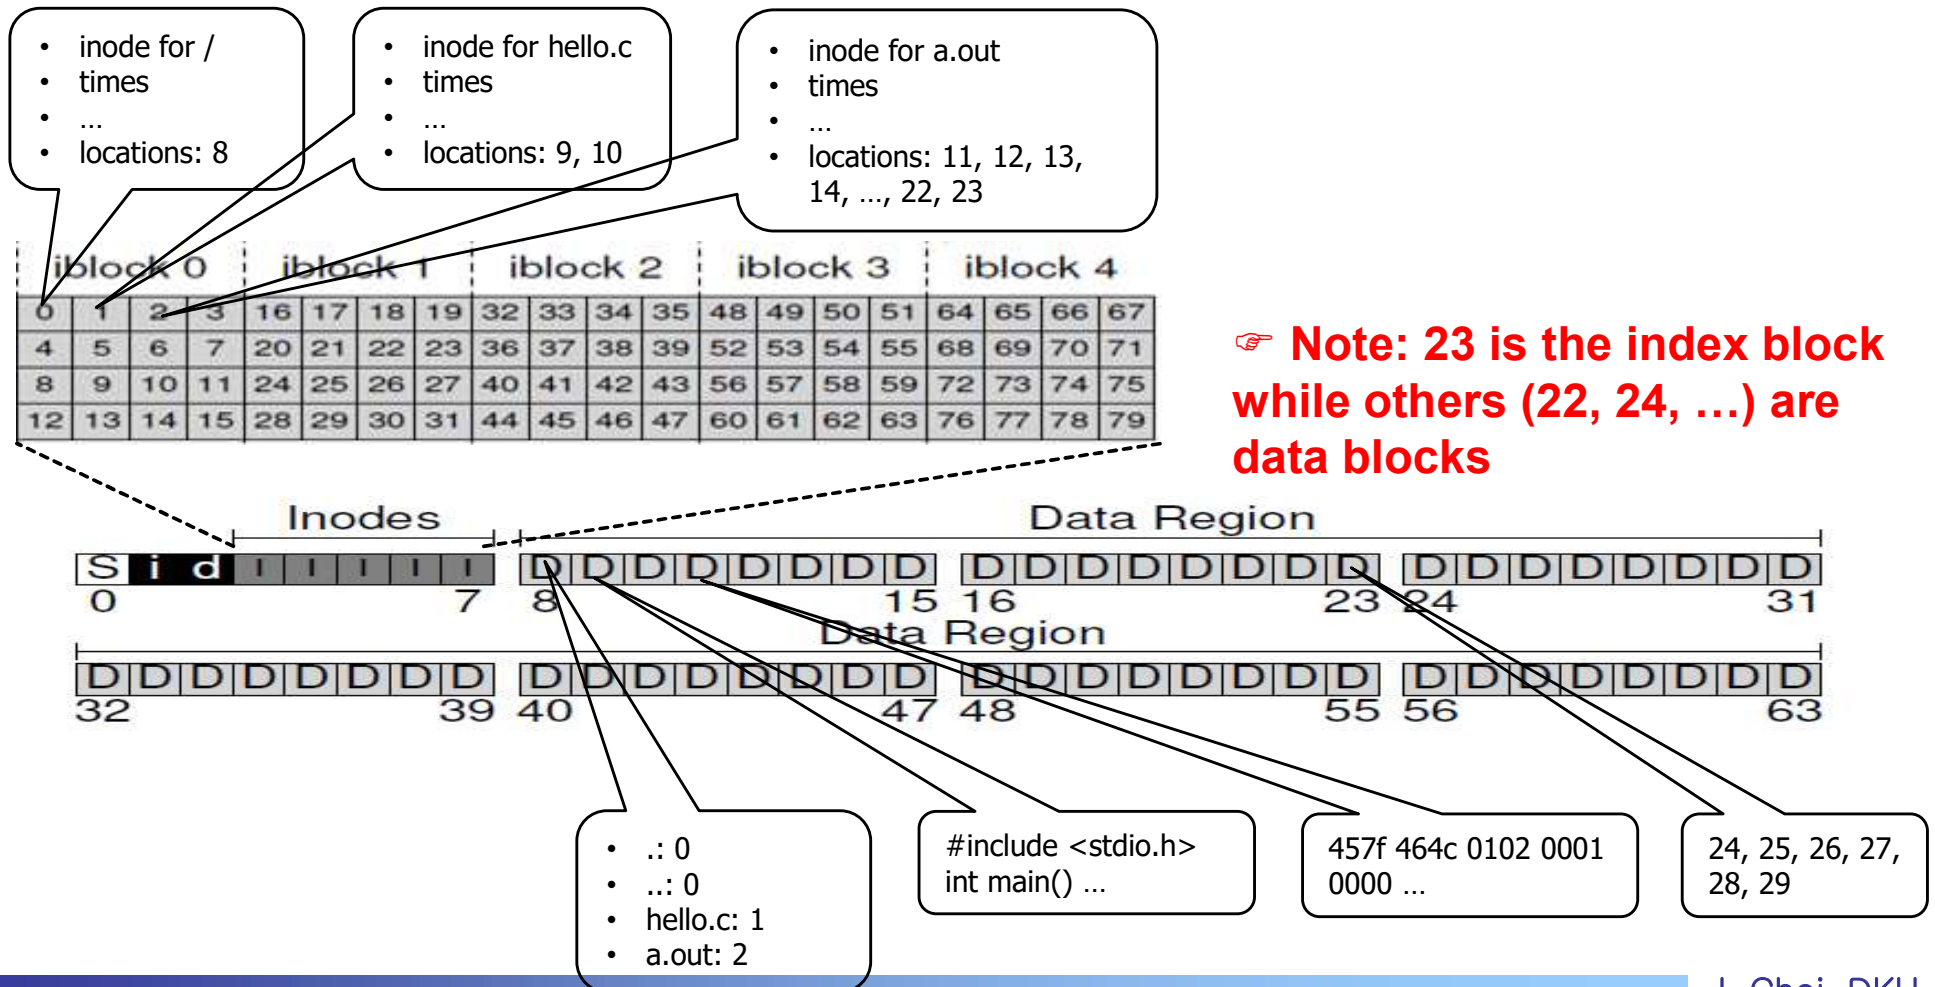
```
#include <stdio.h>
int main() ...
```

```
457f 464c 0102 0001
0000 ...
```

24, 25, 26, 27, 28, 29

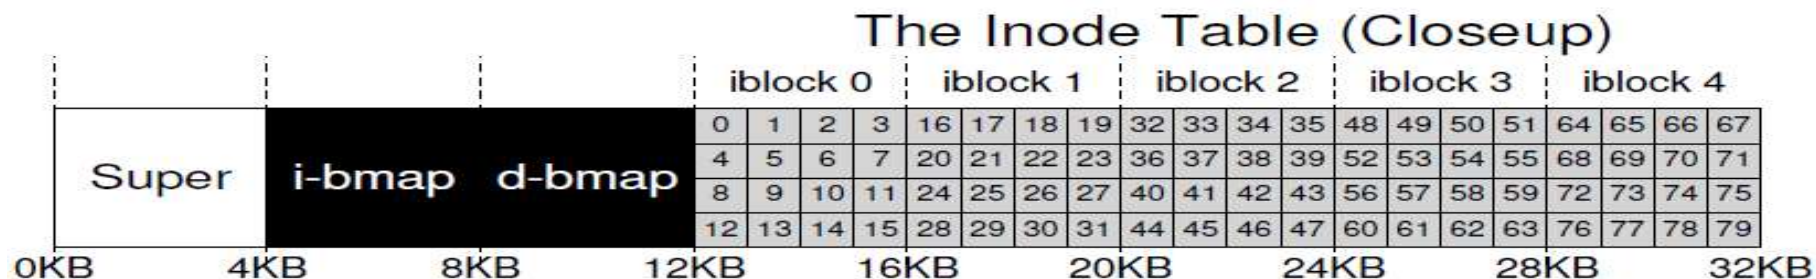
40.3 File Organization: The inode

- inode manipulation example (assume 12 direct blocks)
 - ✓ How to read the a.out?
 - e.g. `fd = open("/a.out", O_RDONLY);`



40.3 File Organization: The inode

- Find a location: inode and data in a real
 - ✓ How to find the location of an inode?
 - Directory entry: <file name, i_number>
 - i_number is used as **the index in inode table** (quotient and remainder)
 - e.g.) i_number = 33 → $33 / (\text{inodes per block}) = 33/16 = 2 \dots 1$ → inode table start + 4KB x 2 = 12KB + 8KB = 20KB → read a block starting 20KB → go to the offset of inode_size x 1 = 256B



- ✓ How to find the location of User data?
 - 1) Find inode, 2) file's current_offset / disk block size = quotient ... remainder, 3) quotient is used to find a **pointer** in the inode (multi-level index), 4) remainder is used as the offset in the disk block
 - e.g.) file's current_offset=5000 → $5000/(\text{block size}) = 5000/4096 = 1 \dots 904$
 - index 1 in inode (e.g. block 12 in the previous slide when the file is a.out)
 - read block 12 → go to the 904 in the block

40.4 Directory Organization / 40.5 Free Space Mgmt.

■ Directory

- ✓ User viewpoint: containing files at a same location
- ✓ System viewpoint: a list of pairs <file name, inode number>
- ✓ For fast search, add the file name length and record length (total bytes including left over space)


| inum | reclen | strlen | name |
|------|--------|--------|--------|
| 5 | 4 | 2 | . |
| 2 | 4 | 3 | .. |
| 12 | 4 | 4 | foo |
| 13 | 4 | 4 | bar |
| 24 | 8 | 7 | foobar |

- ✓ Can use more complex structure for directory (e.g. B-tree in XFS)

■ Free space

- ✓ Bitmap: one bit per block (or inode), indicating whether it is free or used
- ✓ Alternative approach: free-list, tree, ...
- ✓ **Pre-allocation**: allocate free disk blocks in a batch manner → less overhead, contiguous allocation, ...

Bit Map:



Memory:



40.6 Access Paths: Reading and Writing

■ Reading a file from disk

- ✓ open a file “/foo/bar” whose size is 12KB, read data and close it
- ✓ Timeline

| | data bitmap | inode bitmap | root inode | foo inode | bar inode | root data | foo data | bar data[0] | bar data[1] | bar data[2] |
|-----------|----------------|-----------------|---------------|--------------|--------------|--------------|-------------|----------------|----------------|----------------|
| open(bar) | | | read | read | | read | read | | | |
| read() | | | | | read | | | read | | |
| read() | | | | | write | | | | read | |
| read() | | | | | write | | | | | read |

Figure 40.3: File Read Timeline (Time Increasing Downward)

- Open: directory tree traverse → connect fd to inode
- Read: current_offset → find disk block location using the inode and read it → update the last access time in the inode
- Close: deallocate fd and related data structure in OS, No actions in disk
- Note: **repeated reads** for the bar’s inode → How about caching it!

40.6 Access Paths: Reading and Writing

■ Writing a file into disk

- ✓ Create a file “/foo/bar”, write data (also 12KB) and close it
- ✓ Timeline

| | data bitmap | inode bitmap | root inode | foo inode | bar inode | root data | foo data | bar data[0] | bar data[1] | bar data[2] |
|----------------------|----------------|-----------------|---------------|--------------|---------------|--------------|---------------|----------------|----------------|----------------|
| create (/foo/bar) | | read write | read | read | read write | read | read write | | | |
| write() | read write | | | write | read | | | write | | |
| write() | read write | | | | write read | | | | write | |
| write() | read write | | | | write read | | | | | write |

Figure 40.4: File Creation Timeline (Time Increasing Downward)

- Open: 1) create a new inode for bar and update i-bitmap, 2) insert a new entry into foo’s data block (10 I/Os for just creating a file)
- Write: 5 I/Os per a write (d-bitmap read/update, inode read/update, actual user data write)

40.7 Caching and Buffering

■ Issues

- ✓ Disk is too slow.


■ Solutions

✓ 1. Caching

- Caching directories (e.g. / inode, / data, current directory, ...) in DRAM
- Caching recently used file's inodes and data in DRAM
- Management: **LRU (Least Recently Used)** replacement policy, dynamic cache size management

THE CRUX: HOW TO REDUCE FILE SYSTEM I/O COSTS

Even the simplest of operations like opening, reading, or writing a file incurs a huge number of I/O operations, scattered over the disk. What can a file system do to reduce the high costs of doing so many I/Os?



(Source: http://www.atmarkit.co.jp/ait/articles/0810/01/news134_2.html)

40.7 Caching and Buffering

■ Solutions

✓ 2. Write buffering (Delayed write)

- Consolidate several writes into a single one: e.g.) d-bitmap
- Schedule multiple writes so that they have less seek overhead: e.g.) bar data
- Avoid writes: e.g.) temporary file (create and delete immediately)
- Concern: Data loss due to power fault or crash → `fsync()` or direct I/O

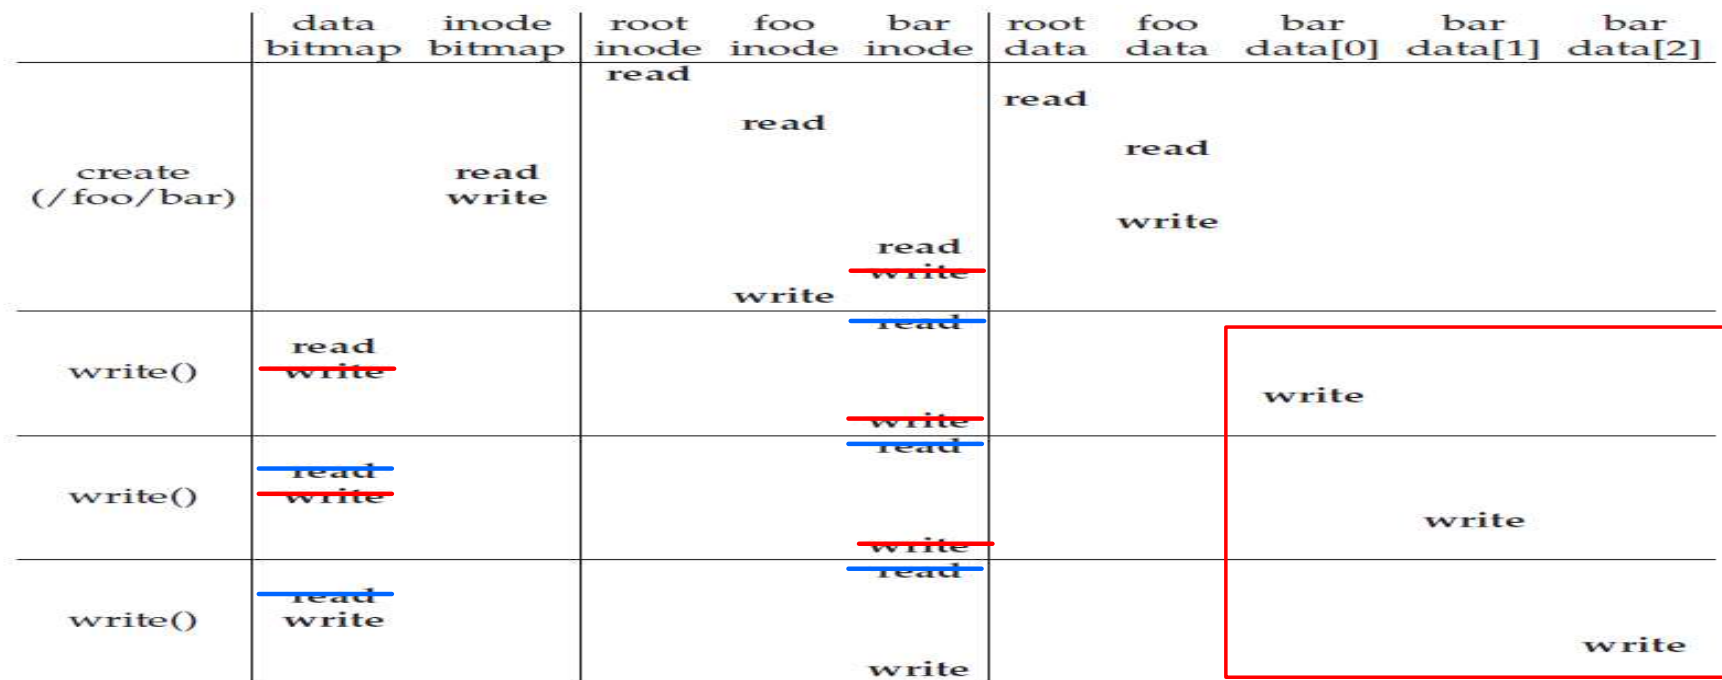


Figure 40.4: File Creation Timeline (Time Increasing Downward)

40.8 Summary

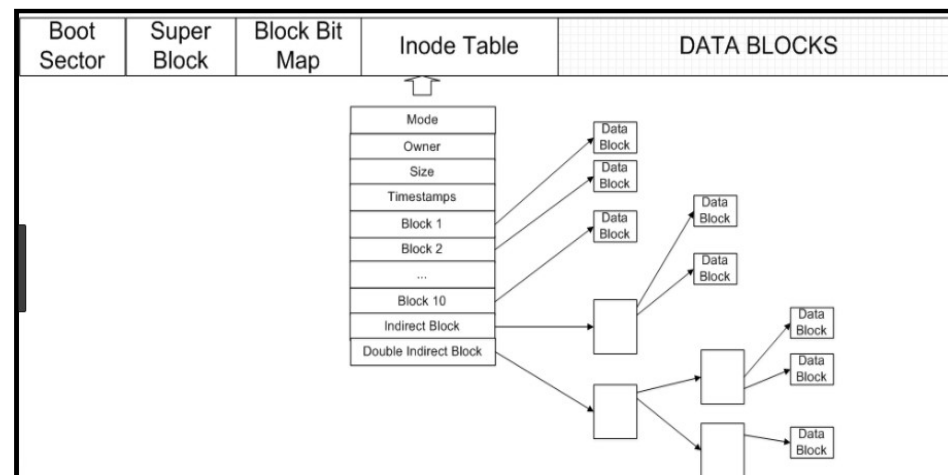
- Device and Driver
- Disk: I/O rate and Scheduling
- File system

- ✓ Interface

- open(), read(), write(), ...
- mkdir(), readdir(), ...
- mount(), mknod(), ...

- ✓ Layout

- Data blocks
- Inode, Bitmap, Superblock
- Boot block



- Importance of **mental model** for OS study (also system study)

ASIDE: MENTAL MODELS OF FILE SYSTEMS

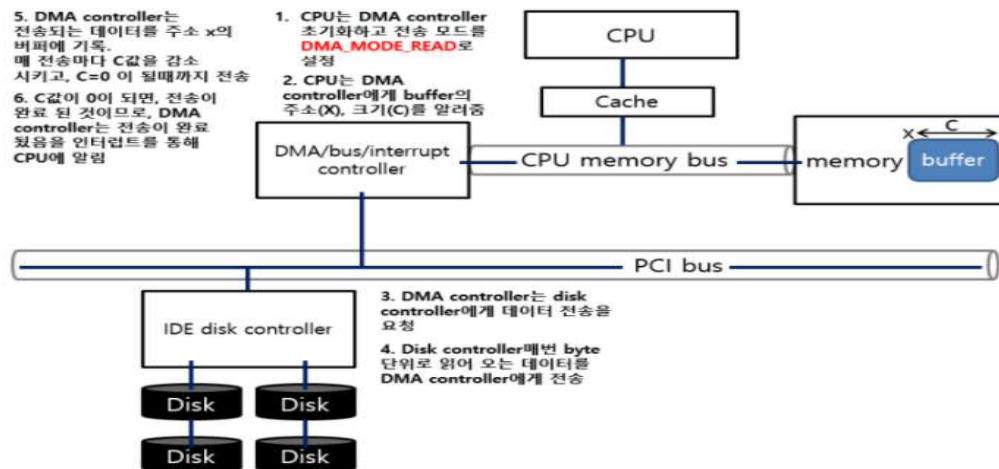
As we've discussed before, mental models are what you are really trying to develop when learning about systems. For file systems, your mental model should eventually include answers to questions like: what on-disk structures store the file system's data and metadata? What happens when a process opens a file? Which on-disk structures are accessed during a read or write? By working on and improving your mental model, you develop an abstract understanding of what is going on, instead of just trying to understand the specifics of some file-system code (though that is also useful, of course!).



Quiz for this Lecture

■ Quiz

- ✓ 1. Discuss the merit/demerit of interrupt, compared with polling. What is the additional merit when we use DMA with interrupt?
- ✓ 2. Discuss the differences between a character device driver and block device driver (at least two differences)
- ✓ 3. Calculate the T_{seek} , $T_{rotation}$, $T_{transfer}$, $T_{I/O}$ and $R_{I/O}$ for the random and sequential workload using Barracuda (hint: refer to 6~8 pages of the Chapter 37 in OSTEP).
- ✓ 4. Discuss why we need fsync(). (hint: using the term of buffer cache and asynchronous write).



| | Cheetah 15K.5 | Barracuda |
|--------------|---------------|-----------|
| Capacity | 300 GB | 1 TB |
| RPM | 15,000 | 7,200 |
| Average Seek | 4 ms | 9 ms |
| Max Transfer | 125 MB/s | 105 MB/s |
| Platters | 4 | 4 |
| Cache | 16 MB | 16/32 MB |
| Connects via | SCSI | SATA |

Figure 37.5: Disk Drive Specs: SCSI Versus SATA

| | Cheetah | Barracuda |
|----------------------|-----------|-----------|
| $R_{I/O}$ Random | 0.66 MB/s | 0.31 MB/s |
| $R_{I/O}$ Sequential | 125 MB/s | 105 MB/s |

Figure 37.6: Disk Drive Performance: SCSI Versus SATA

(Source: <https://m.blog.naver.com/PostView.nhn?blogId=bycho211&logNo=220975324334&proxyReferer=https:%2F%2Fwww.google.com%2F>)

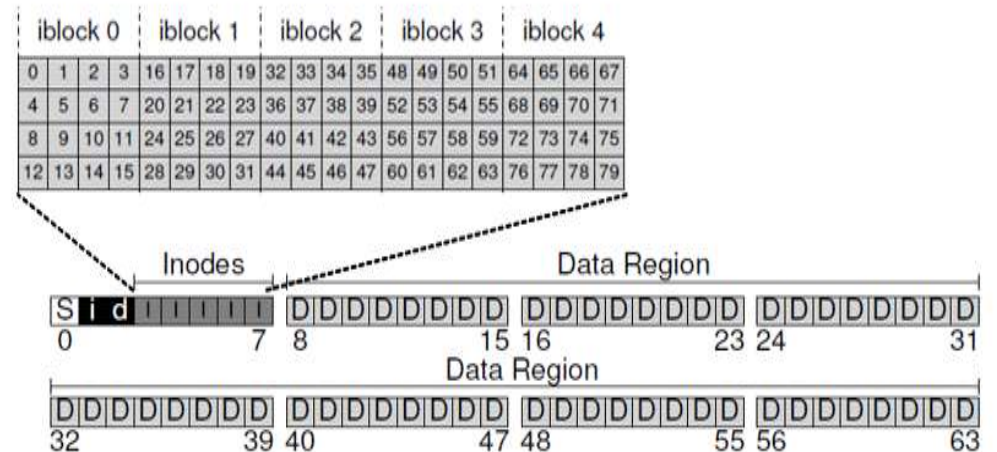


Quiz for this Lecture

■ Quiz

- ✓ 5. The below figure is the snapshot that I conduct with “ls -l” for “current directory”, “/dev/tty”, “/dev/sda”, “/dev/sda1” in our Lab. environment. What are the meaning of “-”, “l”, “b”, “c”, “rw” “5”, “8”, “1” in the figure?
- ✓ 6. Discuss 4 components and their role when we create the VSFS using “mkfs” command.
- ✓ 7. How many disk blocks are allocated from the data region when we create a file “Lab2_report.doc” whose size is 75KB (disk block: 4KB, also explain which block is used for an index block) .
- ✓ 8. When we read (or write) a file we need to access an inode and data alternately, which may cause a long seek distance. Propose **your own idea** for reducing this long seek overhead.

```
choijm@choijm-VirtualBox: ~/OS/chap39
choijm@choijm-VirtualBox:~/OS/chap39$ ls -l
합계 0
choijm@choijm-VirtualBox:~/OS/chap39$ echo "hello DKU" > file1
choijm@choijm-VirtualBox:~/OS/chap39$ ln -s file1 file2
choijm@choijm-VirtualBox:~/OS/chap39$ ls -l
합계 4
-rw-rw-r-- 1 choijm choijm 10  3월 31 13:52 file1
lrwxrwxrwx 1 choijm choijm  5  3월 31 13:53 file2 -> file1
choijm@choijm-VirtualBox:~/OS/chap39$ ls -l /dev/tty
crw-rw-rw- 1 root  tty  5, 0  3월 31 12:22 /dev/tty
choijm@choijm-VirtualBox:~/OS/chap39$ ls -l /dev/sda
brw-rw---- 1 root  disk 8, 0  3월 31 12:22 /dev/sda
choijm@choijm-VirtualBox:~/OS/chap39$ ls -l /dev/sda1
brw-rw---- 1 root  disk 8, 1  3월 31 12:22 /dev/sda1
```



Appendix 1

■ mmap()

ASIDE: mmap () AND PERSISTENT MEMORY (Guest Aside by Terence Kelly)

Memory mapping is an alternative way to access persistent data in files. The `mmap ()` system call creates a correspondence between byte offsets in a file and virtual addresses in the calling process; the former is called the **backing file** and the latter its **in-memory image**. The process can then access the backing file using CPU instructions (i.e., loads and stores) to the in-memory image.

By combining the persistence of files with the access semantics of memory, file-backed memory mappings support a software abstraction called **persistent memory**. The persistent memory style of programming can streamline applications by eliminating translation between different data formats for memory and storage [K19].

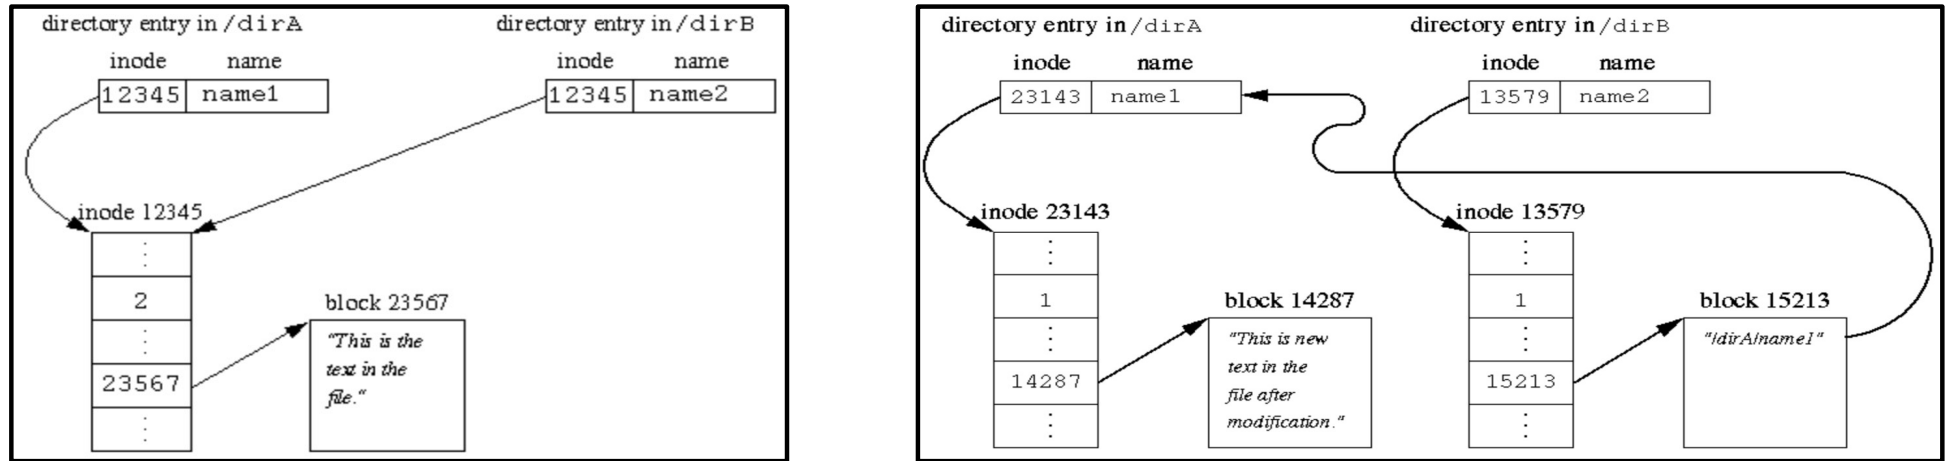
```
1 p = mmap(NULL, file_size, PROT_READ|PROT_WRITE,
2         MAP_SHARED, fd, 0);
3 assert(p != MAP_FAILED);
4 for (int i = 1; i < argc; i++)
5     if (strcmp(argv[i], "pop") == 0) // pop
6         if (p->n > 0) // stack not empty
7             printf("%d\n", p->stack[--p->n]);
8     } else { // push
9         if (sizeof(pstack_t) + (1 + p->n) * sizeof(int)
10            <= file_size) // stack not full
11             p->stack[p->n++] = atoi(argv[i]);
12     }
```

The program `pstack.c` (included on the OSTEP code github repo, with a snippet shown above) stores a persistent stack in file `ps.img`, which begins life as a bag of zeros, e.g., created on the command line via the `truncate` or `dd` utility. The file contains a count of the size of the stack and an array of integers holding stack contents. After `mmap ()`-ing the backing file we can access the stack using C pointers to the in-memory image, e.g., `p->n` accesses the number of items on the stack, and `p->stack` the array of integers. Because the stack is persistent, data push'd by one invocation of `pstack` can be pop'd by the next.


A crash, e.g., between the increment and the assignment of the `push`, could leave our persistent stack in an inconsistent state. Applications prevent such damage by using mechanisms that update persistent memory atomically with respect to failure [K20].

Appendix 2

■ Hard link vs. Symbolic link(Soft link)



■ fd (file descriptor), file table and inode



(Source: <http://classque.cs.utsa.edu/classes/cs3733/notes/USP-05.html>)

사사

- 본 교재는 2026년도 과학기술정보통신부 및 정보통신기획평가원의 ‘SW중심대학사업’ 지원을 받아 제작 되었습니다.
- 본 결과물의 내용을 전재할 수 없으며, 인용(재사용)할 때에는 반드시 과학기술정보통신부와 정보통신기획평가원이 지원한 ‘SW중심대학’의 결과물이라는 출처를 밝혀야 합니다.

IITP 정보통신기획평가원
디지털인재양성단 SW인재팀

SW중심대학이 무엇인가요?

SW중심대학은
대학교육을 SW중심으로 혁신함으로써,
SW전문인력을 양성하고
학생·기업·사회의 SW경쟁력을 강화해
진정한 SW가치 확산을 실현하는 대학을 말합니다.

