

CS 344: OPERATING SYSTEMS I

01.30: FILES

M/W 12:00 – 1:50 PM (LINC #200)

Sanghyun Hong

sanghyun.hong@oregonstate.edu



Oregon State
University

SAIL
Secure AI Systems Lab

NOTICE

- Announcements
 - PA I submissions we cannot compile

NOTICE

- Deadlines (~2 weeks)
 - (1/30 11:59 PM) Midterm quiz 1 – Today
 - (2/06 11:59 PM) Programming assignment 2
 - (2/13 11:59 PM) Midterm quiz 2

PRELIMINARIES: UNIX AND LINUX

- *NIX: Operating Systems
 - UNIX ([Paper](#))
 - 1969: The OS was developed by AT&T Bell Lab, written in assembly languages
 - 1972: C was developed by the same lab
 - “UNIX can run on hardware costing as little as \$40,000” [1]

PRELIMINARIES: UNIX AND LINUX

- *NIX: Operating Systems

- UNIX ([Paper](#))

- 1969: The OS was developed by AT&T Bell Lab, written in assembly languages
 - 1972: C was developed by the same lab
 - “UNIX can run on hardware costing as little as \$40,000” [1]

- Linux

- 1991: Open-source OS, developed by Linus Torvalds
 - Studied at the University of Helsinki
 - His master’s thesis: “Linux: a Portable Operating System” ([Thesis](#)) [2]
 - » “while the Linux project has been closely associated with me personally, partly due to the name, I’d like to make it very clear that the Linux OS is a huge project done co-operatively by lots of people all over the world ... Thanks to all of you.”

[1] The UNIX Time-Sharing System, *ACM SOSP 1973*

[2] Linux: a Portable Operating System, Linus Torvalds

PRELIMINARIES: UNIX AND LINUX

- *NIX: Operating Systems

- UNIX ([Paper](#))

- 1969: The OS was developed by AT&T Bell Lab, written in C
 - 1972: C was developed by the same lab
 - “UNIX can run on hardware costing as little as \$8

- Linux

- 1991: Open-source OS, developed by Linus Torvalds
 - Studied at the University of Helsinki
 - His master’s thesis: “Linux: a Portable Operating System”
 - » “while the Linux project has been closely associated with the name, I’d like to make it very clear that it was developed co-operatively by lots of people all over the world”
 - **Linus Torvalds lives in Oregon**



PRELIMINARIES: *NIX VS. POSIX

- *NIX: Operating Systems
 - UNIX ([Paper](#))
 - Linux ([Thesis](#))
- POSIX: **P**ortable **O**perating **S**ystem **I**nterface (for Uni**X**)
 - OS standard specified by IEEE
 - Defines standard interfaces for system- and user-level APIs
 - Made the applications are portable between Unix-like OSes

TOPICS FOR TODAY

- Part II: Files and File System Basics
 - Provide abstractions
 - What is a file (and a directory)?
 - What is the access control/permission?
 - Offer standard interface
 - How do we create/read/write a file?
 - How do we modify access/permission?
 - Manage resources
 - How OS manages files (and directories)?

PROVIDE ABSTRACTION: WHAT IS A FILE?

- File
 - **Definition:** a named collection of data (*e.g.*, movie.csv containing movie data)
 - **POSIX** : a sequence of data bytes
 - ***NIX OS** : **everything** is a file
 - Files on secondary storages, *e.g.*, disks
 - Devices (mouse, keyboard, monitor, ...)
 - Network devices (network card, sockets in OS, ...)
 - Inter-process communications (pipes, sockets, ...)

PROVIDE ABSTRACTION: WHAT IS A FILE?

- File
 - **Definition:** a named collection of data (*e.g.*, movie.csv containing movie data)
 - **POSIX** : a sequence of data bytes
 - ***NIX OS** : **everything** is a file

- Directories
 - **Definition** : a folder containing files and directories
 - **Motivation:**
 - Scenario: one day you create 100k+ files and the next day, you want to use them

PROVIDE ABSTRACTION: WHAT IS A FILE?

- Directories
 - **Definition** : a folder containing files and directories
 - **Motivation:**
 - Scenario: one day you create 100k+ files and the next day, you want to use them
 - **Solution** :
 - **S0:** You are Von Neumann; remember all the files
 - **S1:** Your system creates a folder containing all the files for each user
 - **S2:** Your system creates multiple folders containing the same kinds

PROVIDE ABSTRACTION: FILES AND DIRECTORIES (IN LINUX)

```
os1 ~/lecture/CS344-OS1$ ls -lh
```

```
total 312K
```

drwxrwx---	6	sahong	upg1xxxx	186	Apr 10 22:14	.
drwxrwx---	3	sahong	upg1xxxx	73	Apr 5 19:58	..
drwxrwx---	2	sahong	upg1xxxx	95	Apr 5 19:58	bufferoverflow
drwxrwx---	2	sahong	upg1xxxx	52	Apr 4 09:02	bufferoverrun
lrwxrwxrwx.	1	sahong	upg1xxxx	22	Apr 10 22:14	home -> /nfs/stak/users/hongsa
-rw-rw----	1	sahong	upg1xxxx	44	Apr 4 08:15	README.md
drwxrwx---	2	sahong	upg1xxxx	79	Apr 5 20:07	thread
Permission	# hard-link	owner	owner-group	size (b)	last modified	name

PROVIDE ABSTRACTION: FILES AND DIRECTORIES (IN LINUX)

```
os1 ~/lecture/CS344-OS1$ ls -alh
```

total 312K

drwxrwx---	6	sahong	upg1xxxx	186	Apr 10 22:14	.
drwxrwx---	3	sahong	upg1xxxx	73	Apr 5 19:58	..
drwxrwx---	2	sahong	upg1xxxx	95	Apr 5 19:58	bufferoverflow
drwxrwx---	2	sahong	upg1xxxx	52	Apr 4 09:02	bufferoverrun
drwxrwx---	8	sahong	upg1xxxx	299	Apr 10 21:56	.git
-rw-rw----	1	sahong	upg1xxxx	430	Apr 5 19:56	.gitignore
lrwxrwxrwx.	1	sahong	upg1xxxx	22	Apr 10 22:14	home -> /nfs/stak/users/hongsa
-rw-rw----	1	sahong	upg1xxxx	44	Apr 4 08:15	README.md
drwxrwx---	2	sahong	upg1xxxx	79	Apr 5 20:07	thread
Permission	# hard-link	owner	owner-group	size (b)	last modified	name

Hidden files!

PROVIDE ABSTRACTION: ACCESS CONTROL (IN LINUX)

- Users and groups
 - **Users** : an account, tied to actual users or that exists for specific applications
 - Physical users: Alice, Bob, ...
 - Accounts for applications: root (sudo), httpd (Apache), ec2-user (AWS), ...
 - **Groups**: a logical expr of an organization, tying users together for a common purpose
 - Linux services: daemon, ...

PROVIDE ABSTRACTION: USERS AND GROUPS

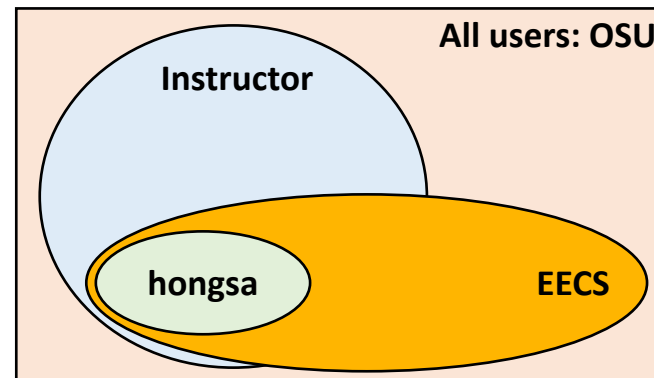
```
os1 ~/lecture/CS344-OS1$ ls -alh
```

```
total 312K
drwxrwx---.    6   sahong upg1xxxx    186   Apr 10 22:14   .
drwxrwx---.    3   sahong upg1xxxx    73    Apr  5 19:58   ..
drwxrwx---.    2   sahong upg1xxxx    95    Apr  5 19:58   bufferoverflow
```

... <omit the entries>

Permission	# hard-link	owner	owner-group	size (b)	last modified	name
------------	-------------	-------	-------------	----------	---------------	------

- Linux controls the access to files or directories based on three categories:
 - **u**ser : owner of a file or a directory
 - **g**roup : the group where users are
 - **o**thers: all the other users



PROVIDE ABSTRACTION: PERMISSION

- Permission
 - **Read** : one can read files and directories with 'r' permission
 - **Write** : one can write files and dirs. with 'w' permission
 - **Execute**: one can execute files and dirs. with 'x' permission

PROVIDE ABSTRACTION: PERMISSION (IN LINUX)

```
os1 ~/lecture/CS344-OS1$ ls -alh
```

```
total 312K
```

Permission	# hard-link	owner	owner-group	size (b)	last modified	name
-rw-rw----	1	sahong	upg1xxxx	44	Apr 4 08:15	README.md
drwxrwx---	2	sahong	upg1xxxx	79	Apr 5 20:07	thread

... <omit the entries>

- Permission representation

- drwxrwx---

[Type] d: directory, -: file

[User] the first three letters (rwx)

[Group] the second three letters (rwx)

[Others] the last three letters (---)

PROVIDE ABSTRACTION: PERMISSION (IN LINUX)

```
os1 ~/lecture/CS344-OS1$ ls -alh
```

```
total 312K
```

Permission	# hard-link	owner	owner-group	size (b)	last modified	name
-rw-rw----	1	sahong	upg1xxxx	44	Apr 4 08:15	README.md
drwxrwx---	2	sahong	upg1xxxx	79	Apr 5 20:07	thread
... <omit the entries>						

- Permission representation

- drwxrwx---

- 770
 - 111111000

Interpretation

- Decimal #: 1st (user), 2nd (group), 3rd (others)
 - + ex. 770 : 7 (user), 7 (group), 0 (others)
- Each #: Binary number
 - 1st (read), 2nd (write), 3rd (execute)
 - + ex. 7 : 111 (rwx)
 - + ex. 6 : 110 (rw)
 - + ex. 600 : 110 000 000 (your ssh key)

TOPICS FOR TODAY

- Part II: Files and File System Basics
 - Provide abstractions
 - What is a file (and a directory)?
 - What is the access control/permission?
 - Offer standard interface
 - How do we create/read/write a file?
 - How do we modify access/permission?
 - Manage resources
 - How OS manages files (and directories)?

OFFER STANDARD INTERFACE: SYSTEM CALLS

- System call
 - **Definition:** a user-level function call to request a service from the OS
 - **Example:** when we run a program “`exec(<a program file>)`”

- Two ways to use a system call
 - **Terminal:** run a command (that is a system call)
 - **C** : call a system call function
 - **Example:** run “`exec`” in Terminal or use “`exec(<a program file>)`” function in C

OFFER STANDARD INTERFACE: USERS AND GROUPS

- System calls (in Terminal)
 - Print the user and group IDs : “exec”
 - Create/modify/delete users : “useradd” / “usermod” / “userdel”
 - Create/modify/delete groups: “groupadd” / “groupmod” / “groupdel”

- System calls (in C)
 - Print the user and group IDs : “getuid()” / “getgid()”
 - Create/modify/delete users : No C APIs; we can use “system('useradd ...')”
 - Create/modify/delete groups: No C APIs; we can use “system('groupadd ...')”

OFFER STANDARD INTERFACE: USERS AND GROUPS

```
os1 ~/lecture/CS344-OS1$ ls -alh
```

```
total 312K
drwxrwx---.    6   sahong  upg1xxxx    186   Apr 10 22:14  .
drwxrwx---.    3   sahong  upg1xxxx    73     Apr  5 19:58  ..
drwxrwx---.    2   sahong  upg1xxxx    95     Apr  5 19:58  bufferoverflow
```

... <omit the entries>

Permission	# hard-link	owner	owner-group	size (b)	last modified	name
------------	-------------	-------	-------------	----------	---------------	------

- An example of 'id' system call

```
os1 ~/lecture/CS344-OS1$ id
```

uid=1xxxxx (sahong)

My user ID

gid=4xxxx (upg1xxxx)

My group ID

groups=4xxxx (upg1xxxx), 3xxx (cs-faculty)

Groups that I am associated with

OFFER STANDARD INTERFACE: PERMISSION

- System calls (in Terminal)
 - Change the ownership : “`chown -R <user>:<group>`”
 - Change the permission: “`chmod -R <mode to set>`”

- System calls (in C)
 - Change the ownership : “`chown(const char *path, uid_t owner, gid_t group)`”
 - Change the permission: “`chmod(const char *pathname, mode_t mode)`”

REVISIT: PATH (IN LINUX)

- Two types of paths
 - **Absolute path**: a complete file/dir path from the root `'/'`
 - **Relative path** : a file/dir path relative from my current working dir `'cwd'`
- Examples:
 - **Absolute path**:
 - `'/nfs/stak/users/sahong/example'` (that you can get from `'pwd'` command)
 - **Relative path** :
 - `'./example_program'`
 - Absolute path for this file: `'/nfs/stak/users/sahong/example/example_program'`
 - **Q1**: If we move to `'/nfs/stak'` what's its relative path?
 - **Q2**: If we move to **HOME** (`'~/`) what's its relative path?

REVISIT: PATH (IN LINUX)

- Useful programming practices
 - 1) Suppose that a program will be used by multiple users.
 - 2) Suppose that the program needs to read a common configuration file.
 - 3) Suppose that a user who runs the program asks to read their file.

Scenario 1) Absolute path.

A developer should put the program binary file (e.g., git) to a common location and users should use the absolute path to run this program in their shell.

Scenario 2) Both will be fine.

As a developer will use the same path for everyone, one can use an absolute path or a relative path from the program binary file.

Scenario 3) Relative path.

As the file shouldn't be read by any other users, the file will be located under a user's home dir. So, we can use a relative path from our home '~/' to the file.

OFFER STANDARD INTERFACE: PERMISSION

```
os1 ~/lecture/CS344-OS1$ ls -alh
```

```
total 312K
```

```
... <omit the entries>
```

-rw-rw----	1	sahong	upg1xxxx	44	Apr 4 08:15	README.md
drwxrwx---	2	sahong	upg1xxxx	79	Apr 5 20:07	thread
Permission	# hard-link	owner	owner-group	size (b)	last modified	name

- Examples of “chown” and “chmod” commands

- \$ chown -R <someone>:<upg1xxxx> thread
- \$ chmod 644 README.md
- \$ chmod o+rw README.md
- \$ chmod 700 thread
- \$ chmod g-rwx thread

Rules

- Use a number, e.g., 644
- Use a string: user/group/others +/- perm.
 - + ex. u+x (user can execute the file/dir)
 - + ex. g-wx (group cannot write or execute it)

OFFER STANDARD INTERFACE: SOME USEFUL SYSTEM CALLS


- System calls (frequently used)
 - Get the details about a file : `$ stat <file/dir name>`
 - Create an empty file : `$ touch <file/dir name>`
 - Total size of a directory : `$ du -alh <file/dir name>`
 - Total filesystem size and info : `$ df -h`
 - ...

TOPICS FOR TODAY

- Part II: Files and File System Basics
 - Provide abstractions
 - What is a file (and a directory)?
 - What is the access control/permission?
 - Offer standard interface
 - How do we create/read/write a file?
 - How do we modify access/permission?
 - Manage resources
 - How OS manages files (and directories)?


PROBLEM: HOW TO STORE FILES TO STORAGE

- Scenario 1: store a file to a disk drive
 - File: a sequence of data




MANAGE RESOURCES: INODE STRUCTURE

- Scenario 1: store a file to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.



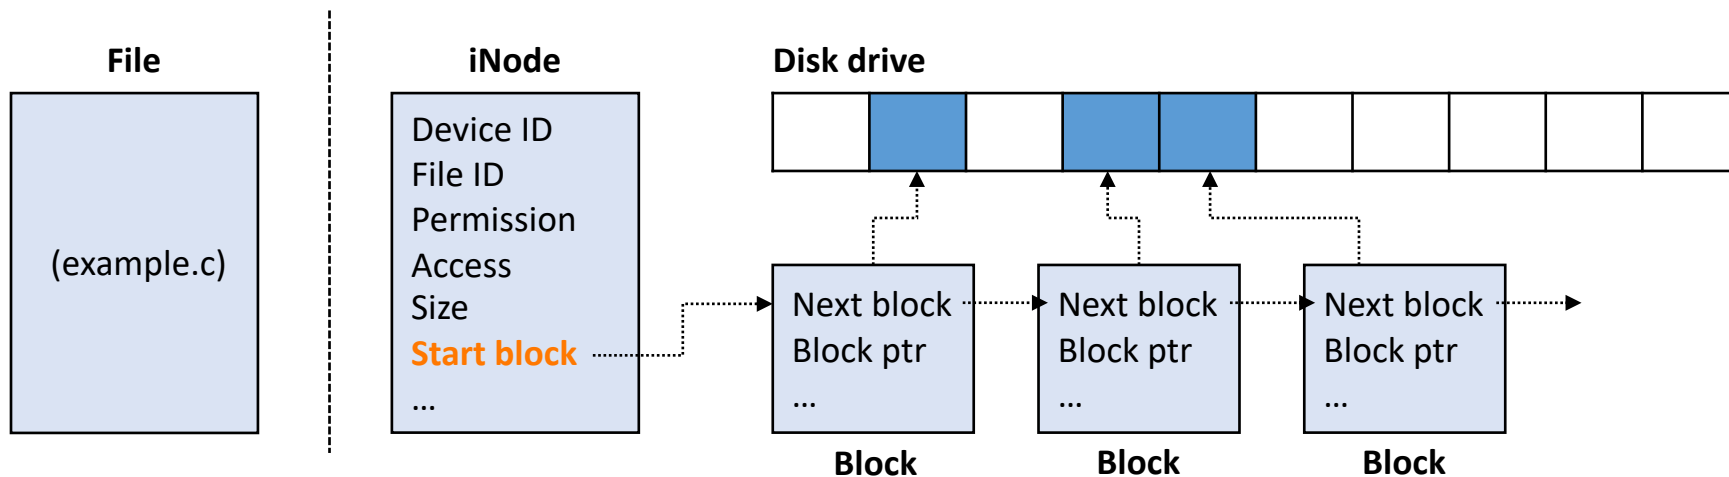
MANAGE RESOURCES: INODE STRUCTURE – CONT'D

- Scenario 2: store **multiple (> 2+)** files to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.




MANAGE RESOURCES: BLOCK STRUCTURE

- Scenario 2: store **multiple (> 2+)** files to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS




MANAGE RESOURCES: BLOCK STRUCTURE

- Scenario 2: store **multiple (> 2+)** files to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS




MANAGE RESOURCES: BLOCK STRUCTURE

- Scenario 2: store **multiple (> 2+)** files to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS




MANAGE RESOURCES: BLOCK STRUCTURE

- Scenario 2: store **multiple (> 2+)** files to a disk drive
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS



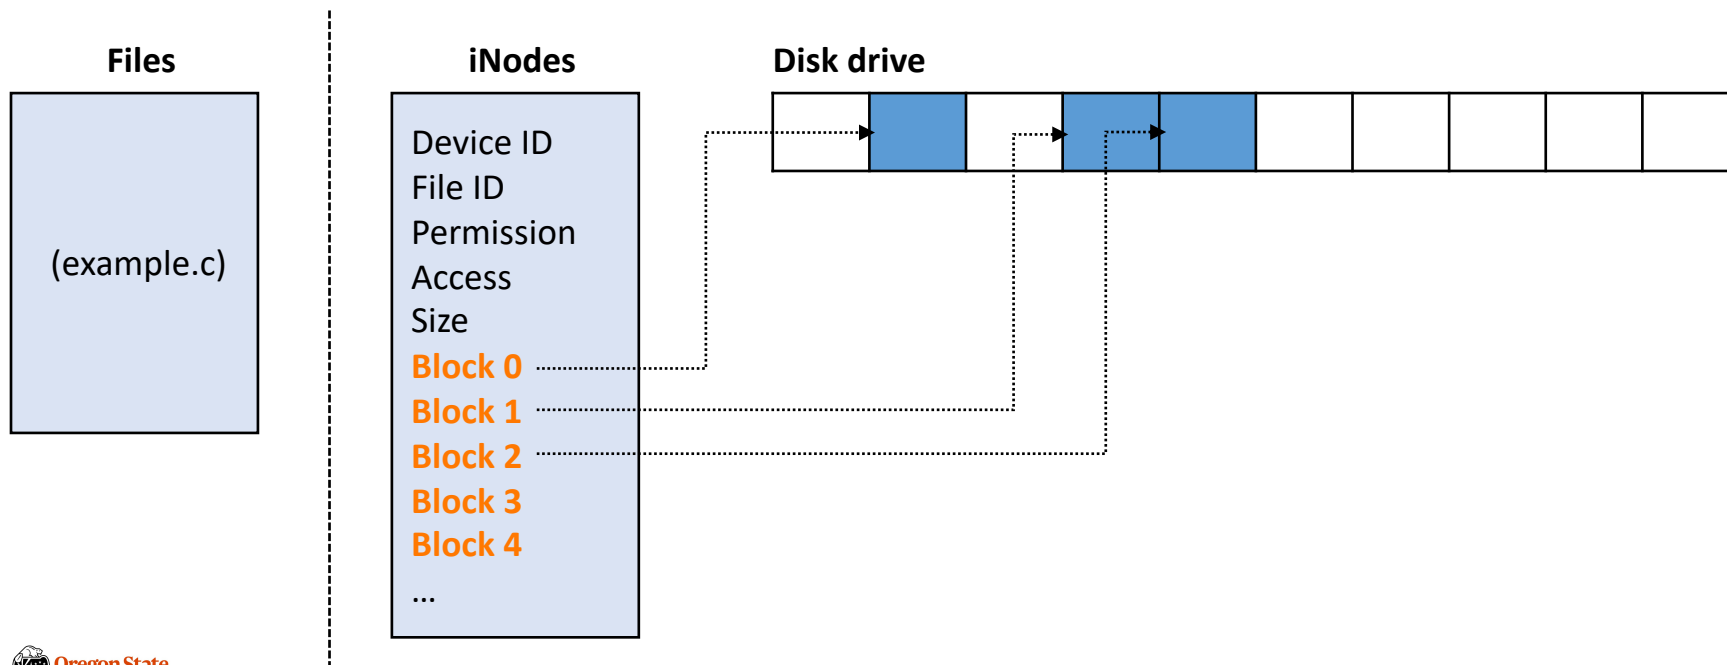
MANAGE RESOURCES: BLOCK STRUCTURE FOR *EFFICIENCY*

- Scenario 3: Users access a certain block(s)
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS

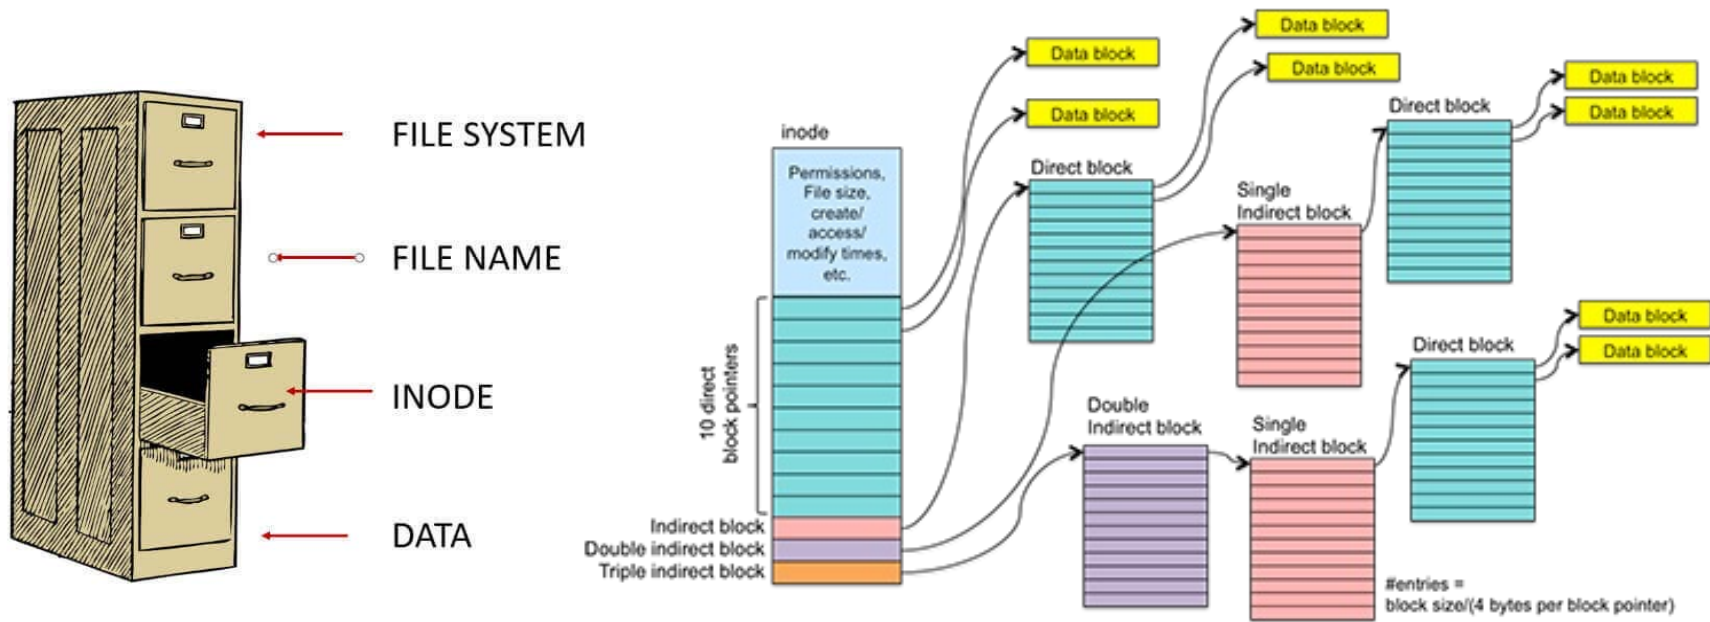


MANAGE RESOURCES: BLOCK STRUCTURE FOR *EFFICIENCY*

- Scenario 3: Users access a certain block(s)
 - **i(ndex)Node**: a data-structure that describes a file-system object, e.g., a file/dir.
 - **Block** : a small(est) unit of data storage, defined by OS



MANAGE RESOURCES: FILESYSTEM STRUCTURE OVERVIEW



[1] What Are inodes in Linux and How Are They Used? <https://helpdeskgeek.com/linux-tips/what-are-inodes-in-linux-and-how-are-they-used/>

[2] File System Design Case Studies, <https://people.cs.rutgers.edu/~pxk/416/notes/13-fs-studies.html>

TOPICS FOR TODAY

- Part II: Files and File System Basics
 - Provide abstractions
 - What is a file (and a directory)?
 - What is the access control/permission?
 - Offer standard interface
 - How do we create/read/write a file?
 - How do we modify access/permission?
 - Manage resources
 - How OS manages files (and directories)?

Thank You!

M/W 12:00 – 1:50 PM (LINC #200)

Sanghyun Hong

sanghyun.hong@oregonstate.edu



Oregon State
University

SAIL
Secure AI Systems Lab