


```

1 CS5600, Handout Week12.b
2
3 // 1. Read/write disk in Lab4 (CS5600 file system)
4 // borrowed from Lab4, cs5600.c
5
6 /* disk access.
7  * All access is in terms of 4KB blocks; read and
8  * write functions return 0 (success) or -EIO.
9  *
10 * read/write "nblks" blocks of data
11 * starting from block id "lba"
12 * to/from memory "buf".
13 * (see implementations in misc.c)
14 */
15 → extern int block_read(void *buf, int lba, int nblks);
16 → extern int block_write(void *buf, int lba, int nblks);
17
18
19 /*
20 * below are two toy examples of
21 * reading from and writing to a disk block
22 */
23
24 // Reading one block
25
26 char buf[FS_BLOCK_SIZE]; // FS_BLOCK_SIZE=4096; see panel 2
27 int inum = 100; // block number to read from
28 int ret = block_read(&buf, inum, 1);
29 if (ret < 0) { // error; ret should be -EIO
30     return ret;
31 }
32
33 // Writing one block
34
35 char buf[FS_BLOCK_SIZE]; // again, 4KB buffer
36 ... // update buf
37
38 int ret = block_write(&buf, 99, 1); // update the 99th block
39 if (ret < 0) { // error; ret should be -EIO
40     return ret;
41 }
42
43
44

```



4KB

```

45
46
47 // 2. CS5600 file system data structures
48 // borrowed from fs5600.h with minor changes
49
50 /*
51 * file: fs5600.h
52 * description: Data structures for CS5600 file system.
53 *
54 * CS 5600, Computer Systems, Northeastern CCIS
55 * Peter Desnoyers, November 2016
56 *
57 * Modified by CS5600 staff in fall 2021.
58 */
59
60 #define FS_BLOCK_SIZE 4096
61 #define FS_MAGIC 0x30303635
62
63 #define INODE_NUM_PTRS (FS_BLOCK_SIZE/4 - 5)
64
65 /* Superblock - holds file system parameters.
66 */
67 struct fs_super {
68     uint32_t magic;
69     uint32_t disk_size; // in blocks *
70
71     /* pad out to an entire block */
72     char pad[FS_BLOCK_SIZE - 2 * sizeof(uint32_t)];
73 };
74
75
76 /* inode = 4096 bytes */
77 struct fs_inode {
78     uint16_t uid; // user id
79     uint16_t gid; // group id
80     uint32_t mode; // change the
81     uint32_t ctime; // modified time
82     uint32_t mtime; // modified time
83     uint32_t size; // in bytes
84     uint32_t ptrs[INODE_NUM_PTRS];
85 };
86
87
88 /* Entry in a directory
89 */
90 struct fs_dirent {
91     uint32_t valid : 1; // 1 bit
92     uint32_t inode : 31; // 31 bits
93     char name[28]; // with trailing NUL
94 };

```

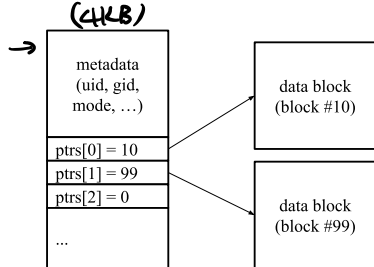
number = inode
block #id

16MB

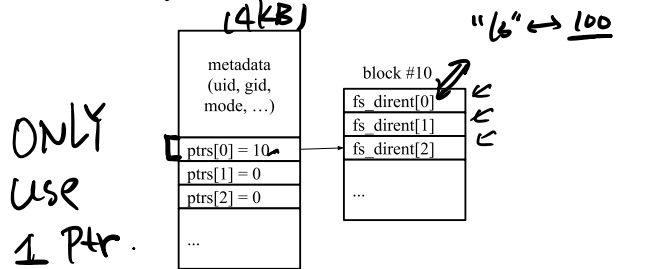
string → inode number
(27 + '\0')

1. fs5600 visualization

fs5600 file inode:




fs5600 directory (also an inode):



Q: how large is fs600?


2. fs5600 block layout



$$4KB \cdot 8 \text{ bit/bytes} = 32 \cdot 2^{10} \text{ bits}$$

$$= 4KB \cdot 8$$

3. fs5600 inode mode



4. interfaces

- fs_init - constructor (i.e. put your init code here)
- fs_statfs - report file system statistics
- fs_getattr - get attributes of a file/directory
- fs_readdir - enumerate entries in a directory
- fs_read - read data from a file
- fs_rename - rename a file
- fs_chmod - change file permissions
- fs_create - create a new (empty) file
- fs_mkdir - create new (empty) directory
- fs_unlink - remove a file
- fs_rmdir - remove a directory
- fs_write - write to a file
- fs_truncate - delete the contents of a file
- fs_utime - change access and modification times

ls -l

--- rwx...

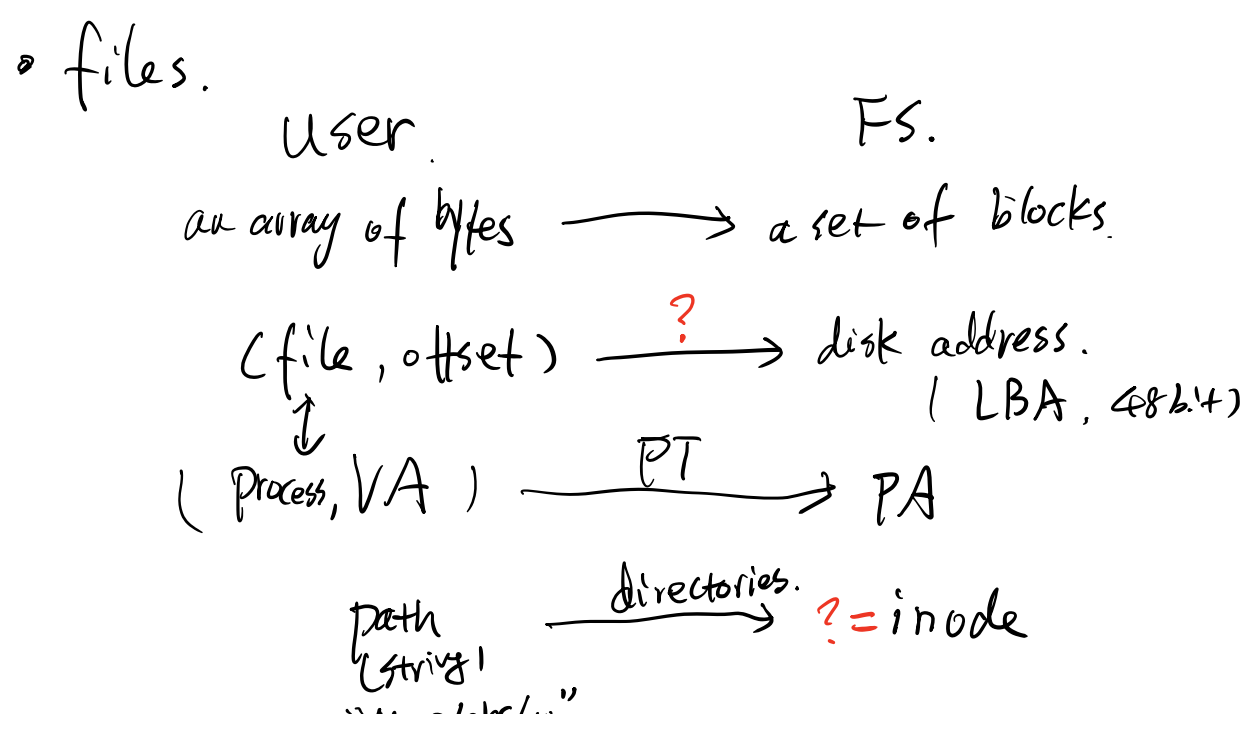
0666

110 bins

1. Last time ←
 2. Files
 3. Directories
 4. fs5600 (Lab4) ←
 - A. fs5600 disk
 - B. data structures
 - C. interfaces
- bottom
 ↓
 up
-

- SSD
 - read/erase/program.
 - Cannot program twice w/o erasing.
 - Wear-out T-TL


- fs.
 - ① persistence
 - ② named bytes ⇒ files
 - ⇒ ③ locate bytes ⇒ directories



~/tmp/1904/...


- FS. design parameters.
 - small files vs. large files.
 - access patterns:
 - sequential access.
 - random access.
 - keyed access.
 - disk utilization
- Candidates for design? (files)

A. Contiguous allocation




- extend file
- fragmentation

B. linked files



C. indexed files.

file2 file1





where is this array?
How large is this?

Unix inode

Q: Why?

Q: Why not
PT be
imbalanced?



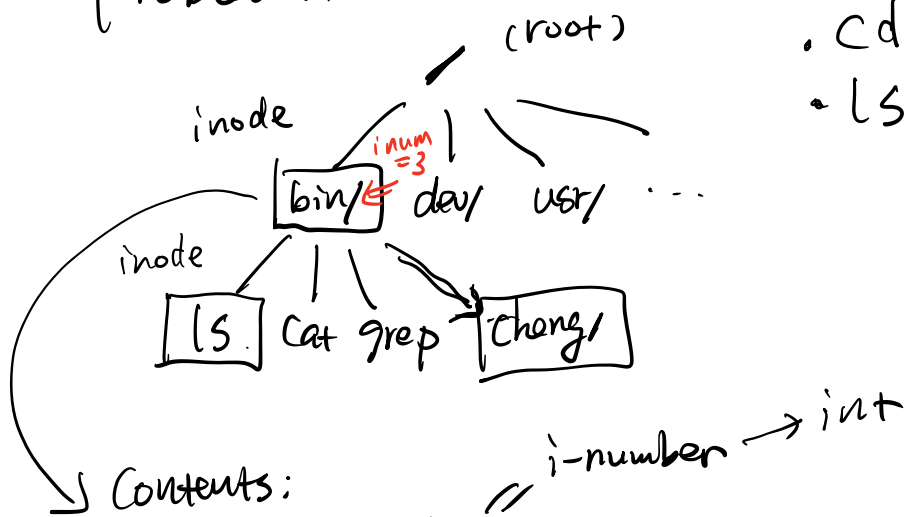
① TLB ←

② HW.


③ 1 level vs. 4-level does not
make diff

- directories.

Problem.



Contents:
 <filename, inode number>
 "ls", 1001
 "cat", 2001
 "cheng", 3001



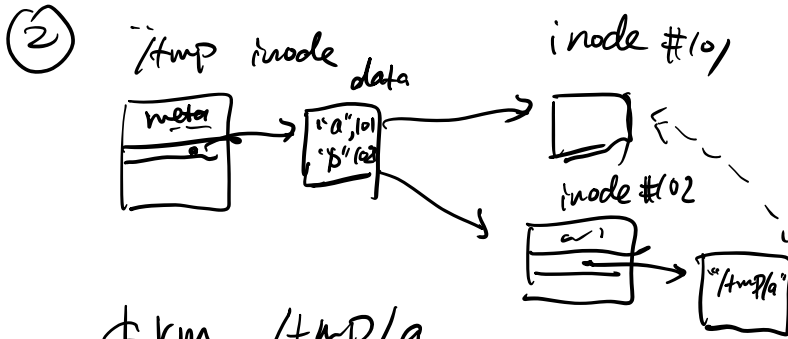

* links.

- ① • hard link
 & ln /tmp/a /tmp/b (alias)

② • soft link.

`$ ln -s /tmp/a /tmp/sb`

(create an inode, link to /tmp/a)



`$ rm /tmp/a`

① `$ cat /tmp/b` ? → okay.

② `$ cat /tmp/sb` ? → error