

Week 11.a
CS3650
03/18 2024
<https://naizhengtan.github.io/24spring/>

1. FS interface ↙
2. Lab4: fs3650

- Lab3.
- Lab4
- Midterm

• strace → syscalls

\$ strace cat file ⇒ ?

- fstat
- read
- write

FZFL*
fprintf


some fs syscalls:

- * statfs - report file system statistics
- { * fstat - get attributes of a file/directory
- * readdir - enumerate entries in a directory
- * read - read data from a file

- * rename - rename a file
- * chmod - change file permissions
- * creat - create a new (empty) file
- * mkdir - create new (empty) directory
- * unlink - remove a file
- * rmdir - remove a directory
- * write - write to a file
- * truncate - delete the contents of a file
- * utime - change access and modification times

• Lab4: CS3650 fs (fs3650)

(a) disk of lab4




```
int block_read ( void * buf, int lba, int nblks );
```

Annotations for the code above:

- An arrow points from `buf` to `char buf [4096];`
- An arrow points from `lba` to `block # (lba)`
- An arrow points from `nblks` to `#blocks`

(b) fs3650 data structures

- superblock
- inode
- inode (dir)



cd etc; pwd
etc

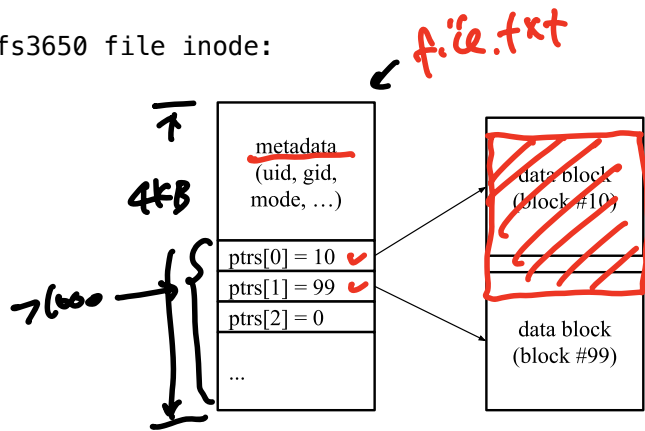
C: \xyz \abc.txt
/ | \

D:
/ | \

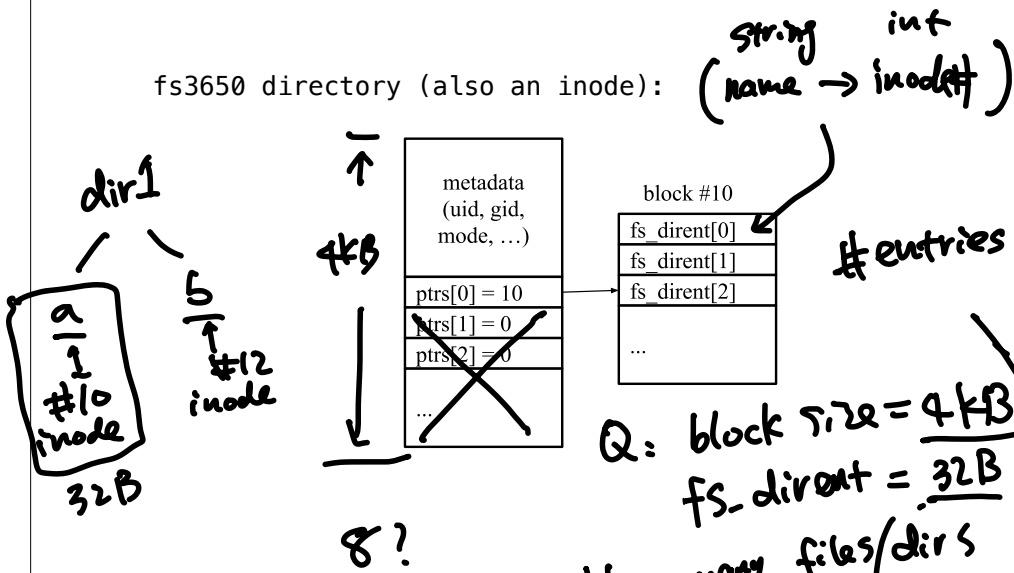
C: \etc

1. fs3650 visualization


fs3650 file inode:



fs3650 directory (also an inode):




Q: $\text{block size} = 4\text{KB}$
 $\text{fs_dirent} = 32\text{B}$
 How many files/dirs under dir?



$\text{inode\#} = \text{block\#}$

3. fs3650 inode mode



file
 ← F: 1
 ← D: 1
 dir

4. interfaces

(a) fs_getattr - get attributes of a file/directory

```
int fs_getattr(const char *path, struct stat *sb,
               struct fuse_file_info *fi)
```

(b) fs_readdir - enumerate entries in a directory

```
int fs_readdir(const char *path, void *ptr, fuse_fill_dir_t filler, off_t offset,
               struct fuse_file_info *fi, enum fuse_readdir_flags flags)
```

(c) fs_read - read data from a file

```
int fs_read(const char *path, char *buf, size_t len, off_t offset,
            struct fuse_file_info *fi)
```

```

1 CS3650 file system
2
3 // 1. Read/write disk in Lab4 (CS3650 file system)
4 // borrowed from Lab4, homework.c
5
6 /* disk access. All access is in terms of 4KB blocks; read and
7 * write functions return 0 (success) or -EIO.
8 */
9 extern int block_read(void *buf, int lba, int nblks);
10 extern int block_write(void *buf, int lba, int nblks);
11
12 /* below is a toy example of reading from a disk block
13 */
14
15 char buf[FS_BLOCK_SIZE]; // FS_BLOCK_SIZE=4096
16 int bnum = 100; // block number to read from
17 int ret = block_read(&buf, bnum, 1);
18 if (ret < 0) { // error; ret should be -EIO
19     return ret;
20 }
21
22
23 // 2. CS3650 file system data structures
24 // borrowed from fs3650.h with minor changes
25
26
27 #define FS_BLOCK_SIZE 4096
28 #define FS_MAGIC 0x33363530
29
30 #define MAX_PATH_NAMES 20 /* max depth of a path */
31 #define MAX_PATH_BYTES 1024 /* max length of a path */
32
33 /* how many buckets of size M do you need to hold N items?
34 */
35 #define DIV_ROUND_UP(N, M) ((N) + (M) - 1) / (M)
36
37 /* Entry in a directory
38 */
39 struct fs_dirent {
40     uint32_t valid : 1;
41     uint32_t inode : 31;
42     char name[28]; /* with trailing NUL */
43 };
44
45 /* Superblock - holds file system parameters.
46 */
47 struct fs_super {
48     uint32_t magic;
49     uint32_t disk_size; /* in blocks */
50     uint32_t root_inode; == 1 /* block number */
51
52     /* pad out to an entire block */
53     char pad[FS_BLOCK_SIZE - 2 * sizeof(uint32_t)];
54 };
55

```

```


56 4KB struct fs_inode {
57     uint16_t uid;
58     uint16_t gid;
59     → uint32_t mode;
60     uint32_t ctime;
61     uint32_t mtime;
62     uint32_t size;
63     uint32_t ptrs[FS_BLOCK_SIZE/4 - 5]; /* inode = 4096 bytes */
64 };
65
66
67 // 3. "struct stat" from Linux
68 // borrowed from "man fstat"
69
70 struct stat {
71     dev_t st_dev; /* ID of device containing file */
72     ino_t st_ino; /* Inode number */
73     mode_t st_mode; /* File type and mode */
74     nlink_t st_nlink; /* Number of hard links */
75     uid_t st_uid; /* User ID of owner */
76     gid_t st_gid; /* Group ID of owner */
77     dev_t st_rdev; /* Device ID (if special file) */
78     off_t st_size; /* Total size, in bytes */
79     blksize_t st_blksize; /* Block size for filesystem I/O */
80     blkcnt_t st_blocks; /* Number of 512B blocks allocated */
81
82     /* Since Linux 2.6, the kernel supports nanosecond
83     precision for the following timestamp fields.
84     For the details before Linux 2.6, see NOTES. */
85
86     struct timespec st_atim; /* Time of last access */
87     struct timespec st_mtim; /* Time of last modification */
88     struct timespec st_ctim; /* Time of last status change */
89
90     #define st_atime st_atim.tv_sec /* Backward compatibility */
91     #define st_mtime st_mtim.tv_sec
92     #define st_ctime st_ctim.tv_sec
93 };

```

metadata

- fs3650 critical functions

- walking path



- ① inode#1
- ② find "a"'s inode#
- ③ load "a"'s inode
- ④ find b's inode#
- ⑤ load b's inode
- ⑥ return c.txt's inode#

read("/a/file.txt", buf, 10, 4096);

offset len

↖ ↗

① walk "a/file.txt" ⇒ inode#

