

CSS503A

Lecture # 13

2019-05-13

Administration

- Lab 3 posted
- will try to get assignments out by end of week
- no progress grading Ass 1
- still accepting Ass 2 until Wednesday

## Our Story So Far

- filesystem as data structure
- file control block aka inode
  - file metadata (creation time/owner/size)
- multiple filesystem types
- directory: "special file"
  - unix land: just mapping between string (filename) & inode
- hard link vs soft (symbolic) link
- Superblock
  - where are the inodes (free/allocated blocks)
  - = root inode
- disk allocation
  - a) linked list (serial data only)
  - b) contiguous block (random access, but fragmentation)

## FAT Filesystem

FAT: File Allocation Table

- ms DOS


- table at beginning of filesystem, one entry per block

→ block # is array index

- contents of table entry:

pointer to table entry (index) of next block in file

Block # / index



- linked list containing only "next" field but pointer  $\equiv$  array index  $\equiv$  block #  
→ block # is node index

- + no overhead in data blocks
- + faster random access (than linked list)
- multiple disk seeks to read links
  - unless table is cached
  - thumb drive: no seeking

## FAT Filesystem (cont.)

FAT16 FAT entry: 16 bits (2 bytes)

$$- 2^{16} = 64k$$

block size: 32k

$$- 128kb \text{ for table} \Rightarrow 4 \text{ blocks}$$

total disk size:

$$64k \times 32k = 2^6 \times 2^{10} \times 2^5 \times 2^{10}$$
$$= 2^{31} \text{ bytes}$$

$$\Rightarrow 2Gb$$

## FAT32

- 28 bits of 32 entry size 256M entries
- 32kb block size  $2^8 \times 2^{12}$
- limit 2T (overhead)

~~$$2^{28} \approx (10^3)^6 \approx 10^{18}$$~~

$$2^{28} \times 2^{15} = 2^{43}$$

## Linux Filesystem

ext 2

Block size:	1k	2k	4k	8k
max filesize	16G	256G	2T <sub>2</sub>	2T <sub>16</sub>
max filesystem	4T	8	16	32


- many critical programs cannot handle files larger than 2 GB

ext 4

up to 1 exabyte filesystem  
up to 16 TiB file - 4k block size  
 $\sim 16 \times 10^{12} = 16 \text{ trillion}$

## Indexed File Allocation

index block: pointers to data blocks



- index block(s): per file  
(variable FAT)
- index block size:
  - too small: cannot support large files
  - too large: waste disk space
- multi-level indexed allocation
  - e.g. Unix inode

## Unix Inode Encoding

mode (permissions)

owner/group

timestamps (3)

size, block count

12 direct block pointers

1 pointer to indirect block (single indirection)

1 pointer to double indirection

1 pointer to triple indirection

⇒ fast for small files

Support for large files


4k bytes per block

128 pointers per block

→ 32 bytes per

block address

→ 128 bits



## Journaled File Systems

- journal : transaction log
  - if file system crashes, recover incomplete operations from journal



## Free Space Management

- bit vector

⊕ easy to find contiguous space

⊖ expensive use of memory

- linked list

⊕ Space efficient

- use free space

⊖ Hard to find contiguous space

⇒ essentially, a memory management problem

RAID → redundant array of inexpensive disks

RAID 0

- even split "strips"
- no parity info → no fault tolerance
  - failure of one drive loses all data
- benefit: speed
- large "virtual" volume

RAID 1

- mirroring
- no parity info
- duplicate writes
  - slow
- faster reads

RAID 2/3

- not used much in practice

Hamming code

Hamming code parity

RAID 4

- block-level striping
- dedicated parity disk
  - writing bottleneck

(Reed Solomon?)

## Read 5

- block-level sharding
- distributed parity
- tolerate failure of single drive
  - lost data reconstructed from parity
- faster read performance (than parity)

## Read 6

- Read 5 + additional parity block
- tolerate 2 concurrent drive failures
- no performance penalty for read
- slower writes (additional parity block)

## Hamming Code

- error correction / detection
  - single-bit correction
  - double-bit detection
- parity bit: single-bit detection
- Hamming (7,4) code: 4 data bits + 3 parity bits
  - correct single-bit error
  - add even parity bits to get double-bit detection

Parity : even/odd  
→ detect 1-bit flip

Hamming Code : (7, 4)

7 bits to make 4

→ 3 "check bits" - analogous to Parity

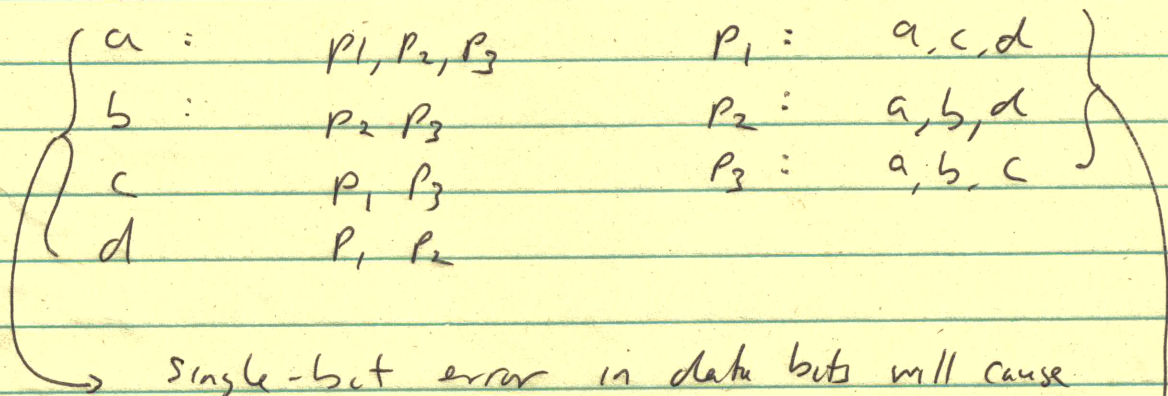
- correct any single bit flip

detect any 2 bit change

} Parity or  
Overlapping Subsets

- number the bits from 1
- place parity bits when bit # is a power of 2
- each parity bit is responsible for bits that have 1-bit in corresponding position (of bit #)

111	110	101	100	11	10	1
a	b	c	p <sub>3</sub>	d	p <sub>2</sub>	p <sub>1</sub>



Single-bit error in data bits will cause unique parity error from which we can determine the error bit

Single-bit parity error will be detected because other parity bits will be correct