

Section 11: File System Details and Reliability

November 9, 2015

Contents

1	Warmup	2
1.1	Conceptual Questions	2
2	Vocabulary	3
3	Problems	4
3.1	Unix-like File System Details	4
3.1.1	inodes	4
3.1.2	Free Block Bitmaps	5
3.2	Reliability	6
3.3	mmap	7

1 Warmup

1.1 Conceptual Questions

For each of the ACID properties, discuss the implications of a transaction system lacking that property.

Atomicity - data left in intermediate state
Consistency - data left in invalid state
Isolation - concurrent transactions may interfere with each other
Durability - crashes may destroy committed transactions

Suppose that we wish to create a reliable storage system. How can we ensure reliability even if part of the system temporarily crashes? How about if part of the system permanently crashes?

Use transactions and redundant systems.

Discuss briefly the different kinds of conflicts that may occur in concurrently executed transactions, in terms of reads and writes.

write-read – reading uncommitted data.
read-write – unrepeatable reads.
write-write – overwriting data.

2 Vocabulary

- **Unix-like File System** -

Unix-like File System is a file system used by many Unix and Unix-like operating systems, although there may be some minor changes over different operating systems.

Basically, disk space is divided into five parts

Boot Sector

Super Block

Free Block Bitmap: This part indicates which block is reserved and which is not. Each bit represents one block. If a bit in the Free Block Bitmap is set to 1, it indicates the corresponding block is free; otherwise the corresponding block has been allocated. Free Block Bitmap is much more efficient than File Allocated Table since one bit is used to tell whether a block is allocated or not. This feature allows Unix-like file system to scale up to large disk space without wasting too much memory. Assume each block has size of 8 KiB. A 1 TiB disk contains 2^{27} blocks. So the Free Block Bitmap only takes 16 MiB space, which is 1/64000 of total disk space. Meanwhile, with FAT, it would take 1 GiB to store File Allocation Table (1/1000 of disk space).

Inode Table: Inode Table stores all the inodes.

Data Blocks

- **Inodes** -

Inode is the data structure that describes the meta data of files (regular files and directories). One inode represents one file or one directory. Inode is composed of several fields such as Ownership, Size, Last Modify Time, Mode, Link Count and Data Block pointers (which point to the data blocks that store the content of the file).

Each inode has 12 direct block pointers (different file system may have different number of direct block pointers). Every direct block pointer directly points to a data block. The maximum size of a file using only direct block pointers can be as big as 80 KiB.

To extend file size, inode introduces the concept of indirect block pointer. Indirect block pointer does not point to file content; instead, it points to the data blocks which store direct block pointers. An inode has one singly, one doubly indirect, and one triply indirect block pointer.

- **Transaction** - A transaction is an indivisible set of data operations that must either succeed or fail as a single unit, that takes data from one valid state to another valid state.

Commonly, reliable transactions are said to satisfy the ACID properties.

- **ACID** - An acronym standing for the four key properties of a reliable transaction.

Atomicity - the transaction must either occur in its entirety, or not at all.

Consistency - transactions must take data from one consistent state to another, and cannot compromise data integrity or leave data in an intermediate state.

Isolation - concurrent transactions should not interfere with each other; it should appear as if all transactions are serialized.

Durability - the effect of a committed transaction should persist despite crashes.

- **Idempotent** - Idempotency is the property on operations that defines the ability of an operation to be repeated without effect beyond the first occurrence. If an operation can be repeated multiple times, but only the first such operation changes the outcome, the operation is said to be idempotent.

- **Logging file system** - A logging file system is a file system in which all modifications are done via transactions. Instead of modifying the disk directly, intended changes are written to an append-only log. Once the transaction is committed, it becomes safe to copy those changes onto disk, since in the event of a crash, transactions committed to the log can safely be re-copied onto disk.
- **2PL** - Two phase locking is a transaction protocol which is used to ensure atomic and isolated commitment of concurrent transactions. Each transaction must acquire a shared lock for any data it wishes to read, and an exclusive lock for any data it wishes to write. Each transaction is broken into two distinct phases: the growing phase, in which it acquires the lock as it needs them, and the shrinking phase, in which it begins releasing its lock. A transaction may not acquire any locks after releasing a single lock.

One key variant of 2PL is strict 2PL, in which all locks must be released at the very end of the transaction.

2PL guarantees that the dependency graph of transactions is acyclic, and that conflicting transactions will be serialized in some order.

3 Problems

3.1 Unix-like File System Details

3.1.1 inodes

For the following questions, assume a FFS inode with 12 direct block pointers, one singly indirect block pointer, one doubly indirect block pointer, and one triply indirect block pointer. Additionally, assume a 4 KiB block size, and 32-bit block pointers.

Consider the `inode_disk` struct in Pintos, which only allows a file to be allocated contiguously on the disk. Create one or more structs to represent an on-disk FFS inode, which allows for direct block pointers as well as indirect block pointers, and does not require a file to have its data allocated contiguously.

```
typedef uint32_t block_sector_t;

/* On-disk inode.
   Must be exactly BLOCK_SECTOR_SIZE bytes long. */
struct inode_disk
{
    block_sector_t start;           /* First data sector. */
    off_t length;                  /* File size in bytes. */
    unsigned magic;                /* Magic number. */
    uint32_t unused[125];         /* Not used. */
};
```

```
/* On-disk inode.
   Must be exactly BLOCK_SECTOR_SIZE bytes long. */
struct inode_disk
{
```

```

    block_sector_t start;           /* First data sector. */
    off_t length;                  /* File size in bytes. */
    block_sector_t direct[12];
    block_sector_t indirect[3];
    unsigned magic;                /* Magic number. */
    uint32_t unused[110];          /* Not used. */
};

```

Consider a random access to the 1200th data block of a file. Supposing that the inode is already in memory, how many disk accesses does it take to read the contents of that data block using Pintos's inode structure? What about using FFS's inode structure?

Each indirect block can hold up to 1024 block pointers, so the 1200th block is doubly indirect, and requires three accesses.

Under Pintos' inode system, file allocation is contiguous, so only one access is required.

Fill out the following struct to represent an in-memory indirect block, then fill out the function in order to return the `block_sector_t` corresponding to the `nth` byte inside that indirect block. Suppose that you have a function `read_indirect_block(block_pointer_t block)` which reads an indirect block from the disk into an `indirect_block_t`. Assume that the file is large enough to contain the desired byte of the file, and ignore triply indirect pointers for this problem. (Bonus: Figure out why triply indirect pointers need to be ignored in this problem.)

```

typedef struct indirect_block {
    block_sector_t block; // block_sector_t of this indirect block
    uint_32t indirection_level;

    _____ block_pointers[_____];
} indirect_block_t

block_sector_t sector_of_byte(indirect_block_t* iblock, uint_32t byte_number) {

    if (iblock->indirection_level > 1) {
        int index = byte_number >> (_____);
        indirect_block_t next_iblock = read_indirect_block(iblock->block_pointers[index]);
        return sector_of_byte(_____);
    }
    return iblock->block_pointers[byte_number >> 12];
}

```

```

typedef struct indirect_block {
    block_sector_t block; // block_sector_t of this indirect block
    uint_32t indirection_level;

    block_sector_t block_pointers[1024];
} indirect_block_t

block_sector_t sector_of_byte(indirect_block_t* iblock, uint_32t byte_number) {

```

```

    if (iblock->indirection_level > 1) {
        index = byte_number >> (10 * iblock->indirection_level + 2);
        next_iblock = read_indirect_block(iblock->block_pointers[index]);
        return sector_of_byte(&next_iblock, byte_number % (1 << (10*iblock->indirection_level + 2)));
    }
    return iblock->block_pointers[byte_number >> 12];
}

```

3.1.2 Free Block Bitmaps

Consider a free block bitmap given as an integer array. Fill in the blanks of `is_set` to check if the bit corresponding to the `i`th block is set, assuming that the input is valid and in-bounds.

```

int bitmap[1024];

int is_set(int index) {
    int offset = index % 32;
    return (bitmap[index / 32] _____);
}

```

```

int bitmap[1024];

int is_set(int index) {
    int offset = index \% 32;
    return (bitmap[index / 32] << (31 - offset)) >> 31;
}

```

3.2 Reliability

Consider the following approach to a transactional key-value store which only stores integers, and each transaction consists of a single arithmetic operation which updates the value associated with a single key. Each transaction consists of the following steps.

```

Read T[i];
Write "T[i] = T[i] + 5" to log
Write "commit T[i]" to log
Commit
Copy changes to disk
Clear log

```

To recover from any potential crashes, the system iterates through the entire log and replicates any committed transactions that are still in the log. Identify a key problem with this system.

```

This system is not idempotent, since if the system crashes after committing changes to disk, but before clearing the log, it will repeat some arithmetic operation, with the final result being the

```

output of the composition of two identical arithmetic operations.

Consider the following approach to file creation in FFS. For each of the ACID properties, discuss whether the given approach satisfies that property. Assuming that we cannot afford to scan the whole disk to check for inconsistencies, which of the ACID properties does using a log provide?

```
Find free block
Find free inode
Write used bit to bitmask (in log)
Write inode to inode table in log
Write data to log
Write directory entry for inode to log
Commit to disk
Copy log changes to disk
```

This is vulnerable to race conditions without the use of a global lock for finding free resources. Without the logging system, we lose atomicity and consistency – it is now possible for a file creation to partially complete, leaving data in an inconsistent state.

For the following schedule of transactions operating on data objects A and B, identify any pairwise conflicts between operations from different transactions.

This information could be used to construct a graph where each transaction is represented by a node, with edges representing conflicting operations between transactions on a particular data object. The significance of this graph will be explained in a later lecture.

T1	R(A)		W(A)					
T2		R(A)						
T3					W(A)		W(B)	
T4				W(A)		W(A)		W(B)

Any write-write, read-write, or write-read conflict on the same data object.

3.3 mmap

Recall the concept of memory mapped files, which allows a block of virtual memory to have a direct correlation with some portion of a file on disk.

Suppose that it takes twice as long to read a 4000 bytes than 2000 bytes from a memory mapped file, but it takes ten times longer to read 8000 bytes. What are some likely causes for the non-linear read time?

Page fault

Discuss the advantages and drawbacks of memory mapped file accesses compared to traditional disk accesses for small random file reads and writes to many files of varying size.

Memory mapped files allow for in-place updates without the need to seek into the file to the appropriate location.

However, since this access pattern accesses data in many different files, it will incur overhead for paging the files in and out. This may or may not be worse than the overhead of traditional I/O,

which involves overhead in the syscall framework and copying memory.