

## 28 File Locking

When files are shared among processes, concurrent accesses can lead to race conditions, similar to those that can occur in shared memory. To allow cooperating processes to coordinate access to shared files, filesystem provides lock primitives. Dimensions in file locking:

- granularity
  - file lock: can lock entire file only
  - record lock: can lock arbitrary region of file (locks held on non-overlapping regions of a file do not conflict)
- semantics
  - shared lock: multiple processes can acquire shared locks for the same (region of) a file
  - exclusive lock: if a process hold an exclusive lock on a (region of) a file, no other process can hold a shared or exclusive lock on the same (region of a) file.
- enforcement
  - mandatory lock: OS prohibits access to files by processes that do not hold the appropriate lock.
  - advisory locks: OS does not restrict access to files based on lock ownership. Processes can chose to ignore locking.

## 29 Filesystem Reliability

Loss of data in a filesystem can have catastrophic effects (much worse than failure of hardware.) Need to ensure reasonable level of safety against data loss in the event of system failures. Threats:

1. media (disk) failure
2. accidental or malicious deletion of data by users
3. system crash during filesystem modifications, leaving data on disk in an inconsistent state
4. during a security compromise, an attacker can delete or tamper with data

Backup: copy entire filesystem onto low-cost media (tape) at regular intervals (e.g., once a day). In the event of a failure of the primary media (disk), can replace media and restore data from backup media. Amount of data loss is limited to modifications that occurred since the last backup. (1, 2, 3, 4)

Mirrored disks: multiple copies of the filesystem are maintained on independent disks. Disk writes update all the redundant disks in parallel. Used in applications that cannot tolerate any data loss (e.g., banking). (1)

RAID (Redundant Array of Inexpensive Disks): use multiple parallel disk drives for higher throughput and increased reliability. For example: each bit of a data byte is stored on one of eight disks. A ninth disk stores a *parity bit* for each data byte. Parity bit is the result of an XOR among the eight data bits of the corresponding data byte. As a result, can still recover the data if one of the nine disks fails. (1) (Also get eighth times the sequential r/w bandwidth; but: small writes are expensive.)

Versioning filesystem: filesystem creates a new version every time a file is modified. Old versions are kept until explicitly deleted. Can use copy-on-write to reduce storage requirements. (1, 2)

Snapshotting filesystem: filesystem takes a snapshot of the entire filesystem at fixed intervals (e.g., daily). Snapshots are kept until explicitly deleted. Can use copy-on-write to reduce storage requirements. (1, partly 2)

After a system crash in the middle of a filesystem operation, filesystem metadata may be in an inconsistent state. (Invariants of the on-disk metadata may not hold.) Example: a file was deleted, but its disk blocks have not yet been added to the freelist. (3)

Solution 1: run a program during system startup that examines the entire filesystem, detects inconsistencies, and restores the invariants. In Unix, that program is called `fsck`. Some of the conditions `fsck` checks and repairs:

- correct i-node reference counts
- missing blocks in the freelist
- blocks that are both in the freelist and part of a file
- incorrect information in the superblock
- out-of-range block and i-node numbers
- disconnected (unreachable) files and directories

This used to be the preferred solution, but today's filesystems have become so big that running `fsck` could take many hours, during which the filesystem would remain unavailable.

Solution 1a: Soft updates Allows delayed writes to disks, but tracks dependencies among metadata updates and ensures they are written in a particular order that ensures consistency in the case of a failure.

Solution 2: keep a separate log (journaling filesystem)

- write log entry describing operations about to be performed
- perform the operation on the filesystem
- add a log entry indicating the operation was completed
- after a crash:
- check the log
- if there are incomplete operations, perform operations described in the log

Requires additional (log) writes; fortunately, log writes are efficient because they are sequential. Recovery time is proportional to the number of operations not yet committed to disk, thus quite short.

Solution 3: maintain only a log (log-structured filesystem)

- write updates only into the log
- writes are very efficient (sequential)
- must read log in reverse to find most recent version of each data structure
- → requires a large cache for efficient reads
- when the log wraps, need to check if existing entries are still live
- if so, need to compact live entries (log cleaning)

## 30 Transactions

Application programs have their own invariants that must hold about the state (contents) of their files. After a system crash, the system cannot restore these invariants, since it has no knowledge of them. Need tools that allow applications to maintain their invariants on persistent data despite crashes.

Transactions are a high-level mechanism that allows applications to specify transitions from one consistent state to another. They can be used by the system to both ensure synchronization (in the event of concurrent access), and all-or-nothing semantics (in the event of system crashes.)

Transaction Primitives:

```
tid = startTransaction();
{sequence of read/write operations to one or more files}
if (anythingGoesWrong)
    abortTransaction(tid);
else {
    {more read/write operations}
    commitTransaction(tid);
}
```

Properties of transactions (ACID):

- Atomicity: either all or no part of a transaction completes
- Consistency: transactions take files from one consistent state to another
- Isolation: concurrent transactions do not interfere with each other
- Durability: once a transaction commits, the changes are permanent

Implementation of transactions:

- atomicity, durability:

- Sol 1: Shadowing. Keep separate, private copies of each modified file blocks, leaving the original file unchanged. When the transaction commits, atomically copy the shadow pages to the original pages. (Difficult to make this work for transactions that write multiple files/tables.)
- Sol 2: Transaction log (write ahead logging). Modify original files directly, but write first into the log all changes that were made (with the original values) . If the transaction aborts or the system fails before the transaction commits (i.e., no commit record in the log upon reboot), use that log to undo all the changes that were made (“rollback”).
- Consistency, Isolation: Two-phase locking (2PL)

In phase 1, locks are acquired and none released. In phase 2, locks are released and none acquired.

Locks can be shared (read lock) or exclusive (write lock).

Guarantees *Serializability*: The effect of a set of transactions is as if the transaction had been executed serially in some order.

There is much more to transactions. Transactions play a central role in databases.