CS 555: Distributed Systems [Google File System]

Shrideep Pallickara
Computer Science
Colorado State University

Frequently asked questions from the previous class survey

Topics covered in this lecture

- Google File System

Chunk size

- This is fixed at **64 MB**
  - Much larger than typical filesystem block sizes (512 bytes)
- **Lazy space allocation**
  - Stored as plain Linux file
  - Extended only as needed

But why this big?

- **Reduces client interaction** with the master
  - Can cache info for a multi-TB working set
- Reduce network **overhead**
  - With a large chunk, client performs more operations
  - Persistent connections
- Reduce **size of metadata** stored in the master
  - 64 bytes of metadata per 64 MB chunk

Why keep the entire metadata in memory?

- **Speed**
- Master can **scan** its **state** in the background
  - Implement chunk garbage collection
  - Re-replicate if there are failures
  - Chunk migration to balance load and space
- **Add** extra memory to increase file system size

Size of the file system with 1 TB of RAM: Assume file sizes are exact multiples of chunk sizes

- Number of entries = $2^{40}/2^{6}$
- MAXIMUM SIZE of the file system
  = Number of entries x Chunk size
  $= \frac{2^{40}}{2^{6}} \times 2^{6} \times 2^{20}$
  $= 2^{60} = 1$ EB

Tracking the chunk servers

- Master **does not** keep a persistent copy of the location of chunk servers
- List maintained via **heart-beats**
  - Allows list to be in *sync* with reality despite failures
  - Chunk server has final word on chunks it holds

Caching at the client/chunk servers

- Clients **do not cache** file data
  - At client the working set may be *too large*
  - Simplify client; eliminate *cache-coherence* problems
- Chunk servers **do not cache** file data either
  - Chunks are stored as local files
  - Linux's buffer cache *already keeps* frequently accessed data in memory

MANAGING MUTATIONS

Mutations

- **Mutation** changes the content or metadata of a chunk
  - Write
  - Append
- Each mutation is performed at **all** chunk replicas

GFS uses leases to maintain consistent mutation order across replicas

- Master grants **lease** to one of the replicas
  - **PRIMARY**
- Primary picks **serial-order**
  - For all mutations to the chunk
  - Other replicas *follow* this order
    - When applying mutations

Lease mechanism designed to minimize communications with the master

- Lease has initial *timeout* of 60 seconds
- As long as chunk is being mutated
  - Primary can request and receive **extensions**
- Extension requests/grants piggybacked over heart-beat messages

Revocation and transfer of leases

- Master may **revoke** a lease before it expires
- If communications lost with primary
  - Master can safely give lease to another replica
    - **Only After** the lease period for old primary *elapses*

How a write is actually performed

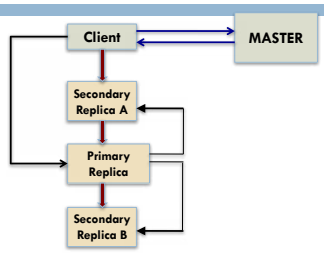

Client pushes data to all the replicas (I)

- Each chunk server stores data in an **LRU buffer** until
  - Data is used
  - Aged out

Client pushes data to all the replicas (II)

- When chunk servers acknowledge receipt of data
  - Client sends a *write request to primary*
- Primary assigns *consecutive serial numbers* to mutations
  - Forwards to replicas

Data flow is decoupled from the control flow to utilize network efficiently

- Utilize each machine's network bandwidth
- Avoid network bottlenecks
- Avoid high-latency links
- **Leverage** network topology
  - Estimate distances from IP addresses

What if the secondary replicas could not finish the write operation?

- Client request is considered **failed**
- Modified region is **inconsistent**
  - *No attempt* to delete this from the chunk
  - Client must handle this inconsistency
- Client **retries** the failed mutation

GFS client code implements the file system API

- Communications with master and chunk servers done *transparently*
  - On behalf of apps that read or write data
- Interact with master for **metadata**
- **Data-bearing** communications directly to chunk servers

Traditional writes

- Client specifies **offset** at which data needs to be written
- Concurrent writes to same region
  - Not serializable
  - Region ends up containing **data fragments from multiple clients**

Atomic record appends

- Client specifies **only** the data **not** the offset
- GFS **appends** it to the file
  - **At least once atomically**
  - At an offset of GFS' choosing
- No need for a distributed lock manger

The control flow for record appends is similar to that of writes

- Client pushes data to replicas of the ***last*** chunk of file
- Primary replica checks if the record **fits** in this chunk

Primary replica checks if the record append will breach the size (64MB) threshold

- If chunk size would be breached
  - **Pad** the chunk to maximum size
  - Tell client, that operation should be retried on *next chunk*
- If the record fits, the primary
  - Appends data to its replica
  - Notifies secondaries to **write at the exact offset**

Record sizes and fragmentation

- Size is restricted to ¼ the chunk size
- Minimizes worst-case fragmentation
  - Internal fragmentation in each chunk …

What if record append fails at one of the replicas

- Client **must retry** the operation
- Replicas of same chunk may contain
  - Different data
  - **Duplicates** of the same record
    - In whole or in part
- Replicas of chunks are **not bit-wise identical**!
  - In most systems, replicas are identical

GFS only guarantees that the data will be written at least once as an atomic unit

- For an operation to return *success*
  - Data must be **written at the same offset** on **all** the replicas
- After the write, all replicas are *as long as* the end of the record
  - Any future record will be assigned a higher offset or a different chunk

Creating snapshots

Snapshots allow you to make a copy of a file very fast

- Master **revoke**s outstanding leases for any chunks of the file (source) to be snapshot
- **Log** the operation to disk
- Update in-memory state
  - Duplicate metadata of the source file
- Newly created file points to the same chunks as the source

When a client wants to write to a chunk C after the snapshot operation

- Master sees the *reference count* to $C > 1$
- Pick **new chunk-handle C'**
- Ask chunk-server with current replica of C
  - Create new chunk C'
  - Data is *copied locally, not over the network*
- From this point chunk handling of C' is no different

GFS does not have a per-directory structure that lists files in the directory

- Name spaces represented as a **lookup** table
  - Maps **full pathnames** to metadata
- File creation does not require a lock on the directory structure
  - No inode needs to be protected from modification

Each master operation acquires a set of locks before it runs

- If operation involves `/d1/d2/…/dn/leaf`
  - Acquire read locks on directory names
    - `/d1, /d1/d2, …, /d1/d2/…/dn`
  - Read or write lock on full pathname
    - `/d1/d2/…/dn/leaf`
- Used to **prevent** operations during snapshots
  - For e.g. cannot create `/home/user/foo`
  - While `/home/user` is being snapshotted to `/save/user`

Locks are used to prevent operations during snapshots

- For e.g. cannot create `/home/user/foo`
  - While `/home/user` is being snapshotted to `/save/user`
- Read locks on `/home` and `/save`
- Write lock on **`/home/user`** and `/save/user`
- File creation does not require write lock on parent directory … there is no "directory"
  - Read locks on `/home` and **`/home/user`**
  - Write lock on `/home/user/foo`
- Read lock prevents a directory from being deleted

Replication

Reasons why chunk replicas are created

- Chunk creation
- Re-replication
- Rebalancing

Chunk replica creation

- Place replicas on chunk servers with below average **disk space** utilization
- **Limit** the number of **recent creations** on a chunk server
  - Predictor of imminent heavy traffic
- Spread replicas **across racks**

Re-replicate chunks when replication level drops

- How **far** is it from replication goal
- Preference for chunks of **live** files
- Boost priority of chunks **blocking client progress**

Rebalancing replicas

- **Examine** current replica distribution and **move** replicas
  - Better disk space
  - Load balancing
- **Removal** of existing replicas
  - Chunk servers with below-average disk space

Incorporating a new chunk server

- **Do not swamp** new server with lots of chunks
  - Concomitant traffic will bog down the machine
- **Gradually** fill up new server with chunks

Consistency in GFS

In GFS the state of file region after mutation depends on …

- **Type** of the mutation
- **Success/Failure** of the mutation
- Whether there were **concurrent** mutations

GFS has a relaxed consistency model

- **Consistent**: See the same data
  - On all replicas
- **Defined**: If it is consistent AND
  - Clients see mutation writes in its **entirety**

File state region after a mutation

	Write	Record Append
Serial success	**defined**	**defined** interspersed with *inconsistent*
Concurrent success	Consistent but *undefined*	
Failure	Inconsistent	

Implications for applications

- Rely on **appends** instead of overwrites
- Checkpoint
- Write **records** that are
  - Self-validating
  - Self-identifying

Deletion of Files & Garbage Collection

Garbage collection in GFS

- After a file is deleted, GFS does not reclaim space immediately
- Done **lazily** during garbage collection at
  - File and chunk levels

Master logs a file's deletion immediately

- File is **renamed to a hidden name**
  - Includes deletion timestamp
- Master scans the file system namespace
  - Delete if hidden file existed for more than 3 days
- When file removed from namespace
  - *In memory metadata* is also removed
  - **Severs links** to all its chunks!

Garbage collection: When Master scans its chunk namespace

- Identifies **orphaned chunks**
  - Not reachable from any file
- Erase metadata for these chunks

The role of heart-beats in garbage collection

- Chunk server reports subset of chunks it currently has
- Master replies with identity of chunks no longer present
  - Chunk server free to delete its replica of such chunks

The contents of this slide-set are based on the following references

- Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung: The Google file system. Proceedings of SOSP 2003: 29-43.