

Weak Consistency and Disconnected Operation in git

Raymond Cheng
ryscheng@cs.washington.edu

Motivation

How can we support disconnected or weakly connected operation?

Applications

- File synchronization across users / devices (e.g. Dropbox)
- Source code control (e.g. git)
- Disconnected / intermittent connectivity (e.g. laptops, mobile, 3rd world)

Consistency Recap

Sequential consistency:

everyone sees same read/write order (cache coherence, Paxos)

Release consistency:

everyone sees writes in unlock order (x86/ARM)

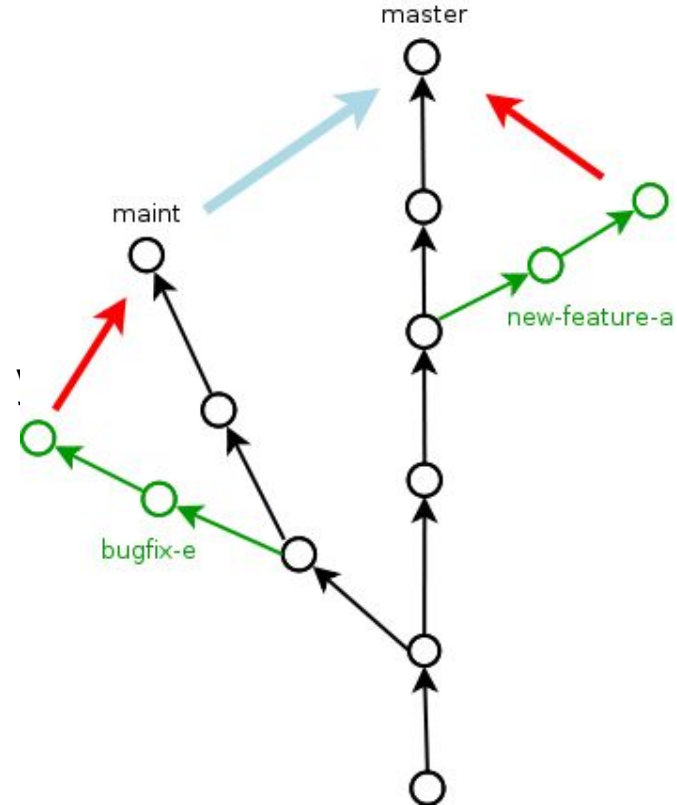
(more generally, reads/writes forced to complete at memory barriers)

Sequential/release consistency is slow:

- wait at memory barrier
- communication needed if recently modified by another node or if write and the local copy is not exclusive

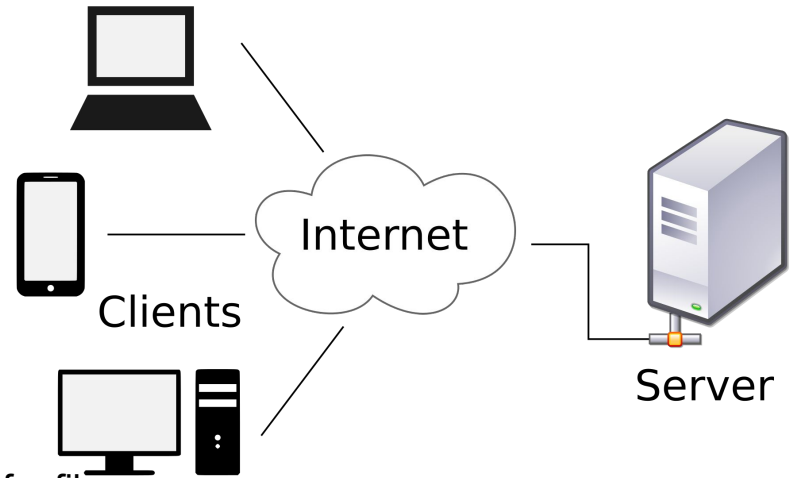
Source Code Control

- Eventual Consistency
 - Okay to read/write cached copy (different versions)
 - Check if it's okay later, recover if necessary
- Track history (with metadata)
- Concurrent editing / Many contributors
- Working copy: Don't want files to change beneath:
 - Push / Pull to server/peers
 - Contributors may be offline / disconnected
- Cheap Branches / Merging
- Centralized / Distributed

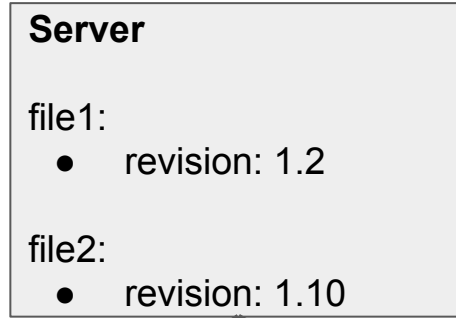


CVS (1990)

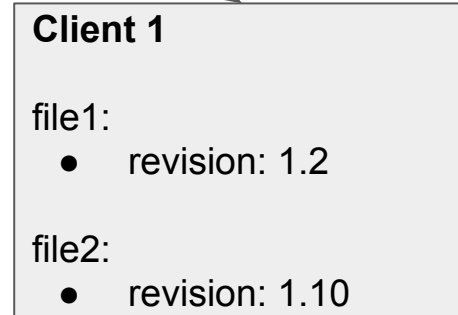
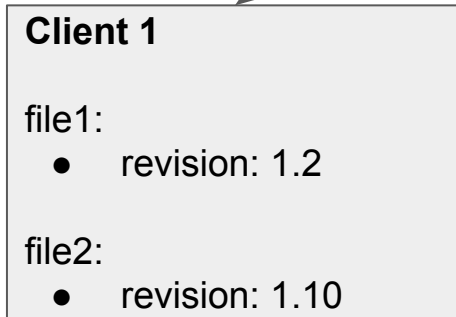
- Client-server architecture
 - Check out a working copy
 - Check in your changes
- Server arbitrates order of changes
 - Only accept changes to the most recent version of a file
 - Developers must always keep their files up to date



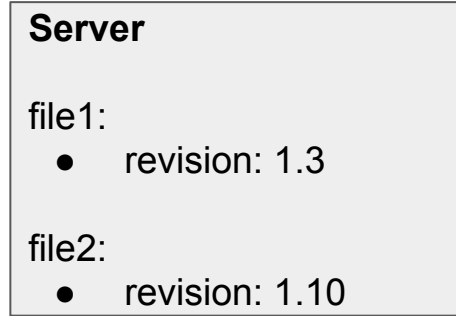
CVS Workflow



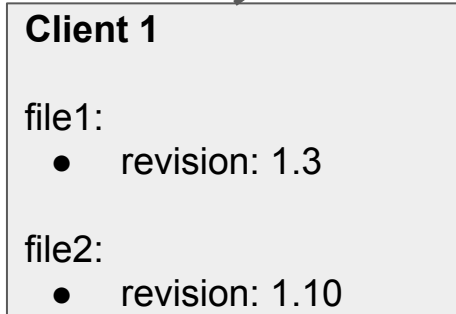
checkout



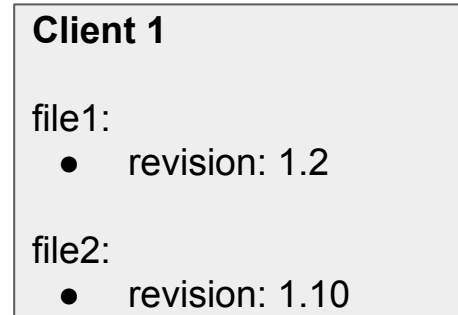
CVS Workflow



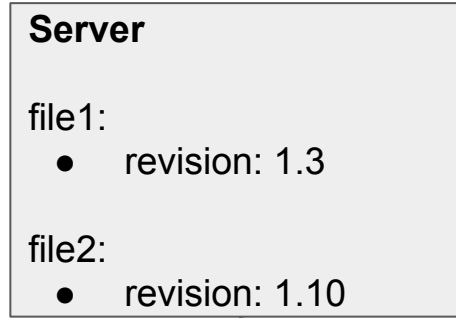
**commit file1
r1.3**



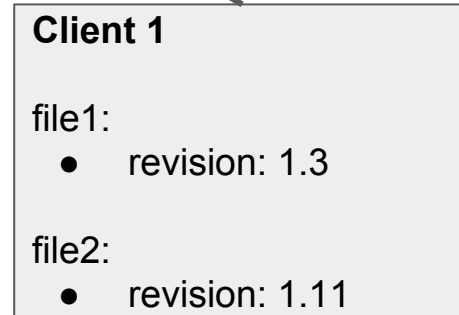
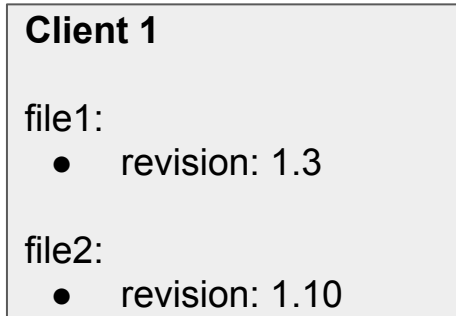
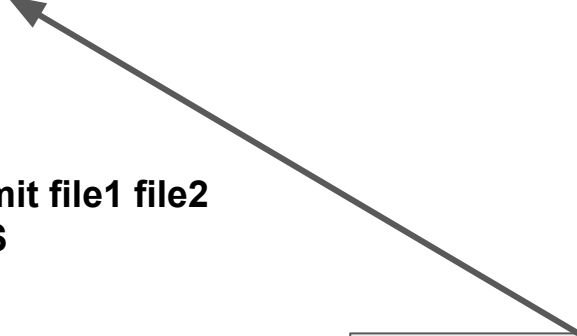
Edit file1



CVS Workflow

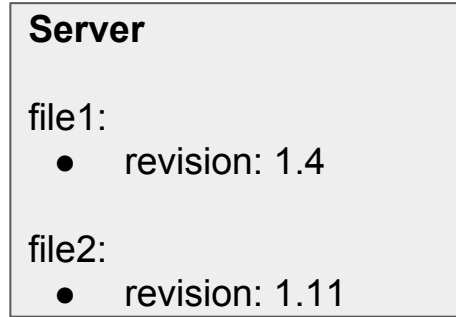


**Commit file1 file2
FAILS**



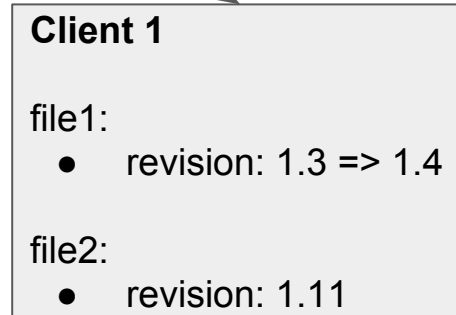
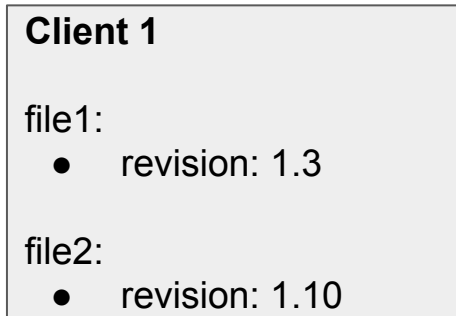
Edit file1, file2

CVS Workflow



**update file1
r1.3**

**commit file1 file2
SUCCEEDS**



**Fix file1
conflicts**

CVS Limitations

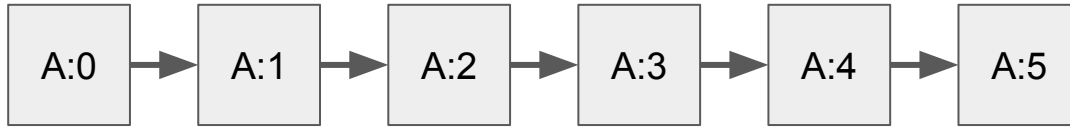
- Everyone is editing the same repository
 - How do you implement a complex feature without constantly conflicting?
- No local version control
 - `cvs commit ~ git commit && git push`
- No log
 - How do I time travel?
- No versioning of moving / renaming files
- Depends on live server to operate
 - Needs to be scaled / backed up / always up / reachable
- Branches were expensive
 - Manual locks are common
- No atomicity (e.g. network failure lead to inconsistency)

Apache SVN (2000)

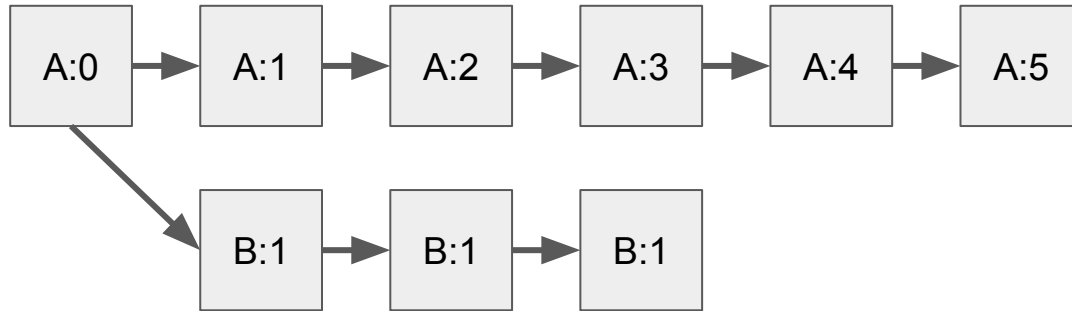
- “CVS done right”
- Improvements
 - Atomic commits
 - Renamed / moved / copied files retained version history
 - Versioning of directories and metadata
 - Cheap branches / tagging
- Centralized - server/client architecture
- Still active
 - All of Facebook’s source code was in a single SVN repository until 2014



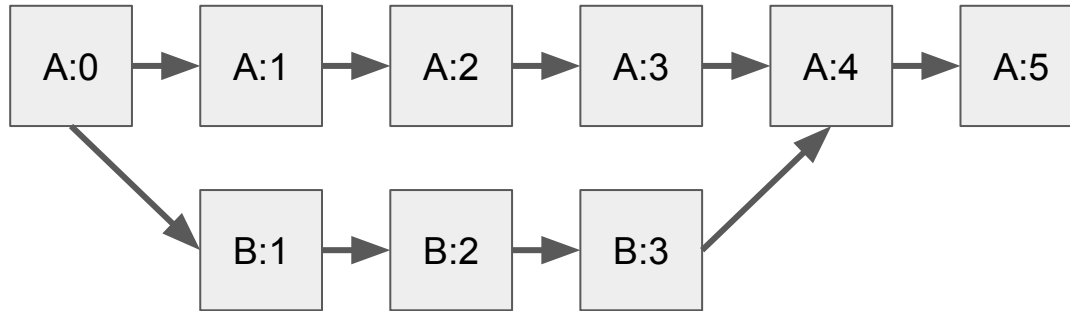
Commit Log



Branching



Merging



A:5 Ancestry Set
{A: 0-4, B:1-3}

Merging and Causal Ordering

Operations that potentially are causally related are seen by every node of the system in the same order

Example:

C1: f=1 -> C2

C2: f=2 -> C3

C3: f=3 -> C1

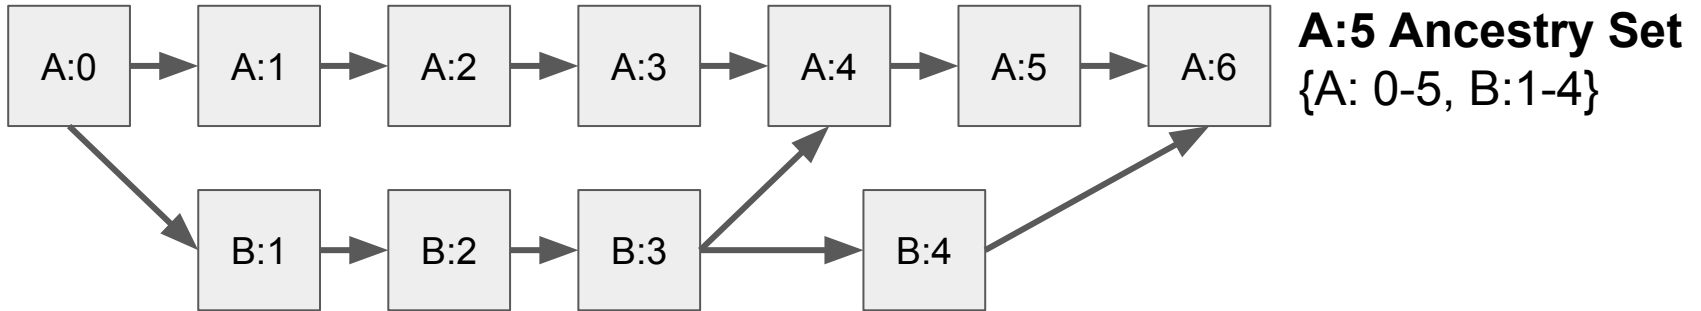
Example:

C1: a=1 -> C2

C2: b=2 -> C3

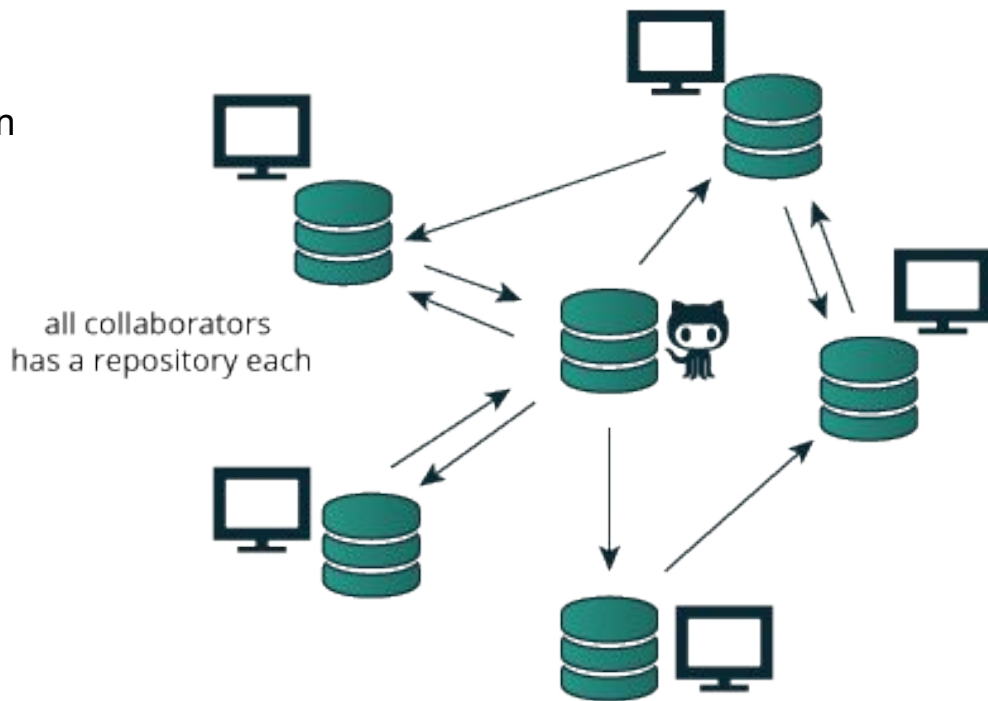
C3: c=3 -> C1

Merging



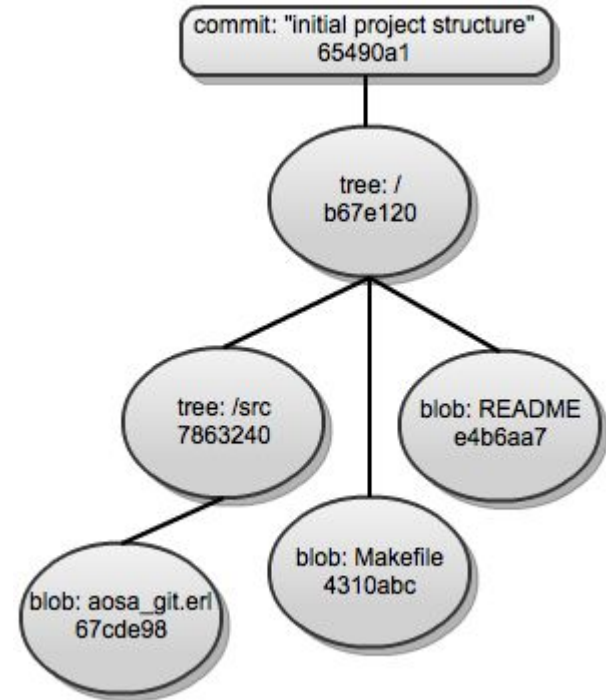
git (2005)

- Distributed!
 - Everyone is a replica
- Consistency and performance
 - Protects from memory, disk corruption
- Cheap branches / merges
- .git/
 - Config
 - Content-addressable filesystem
 - Log of changes (commit history)



Logs (Commit Histories)

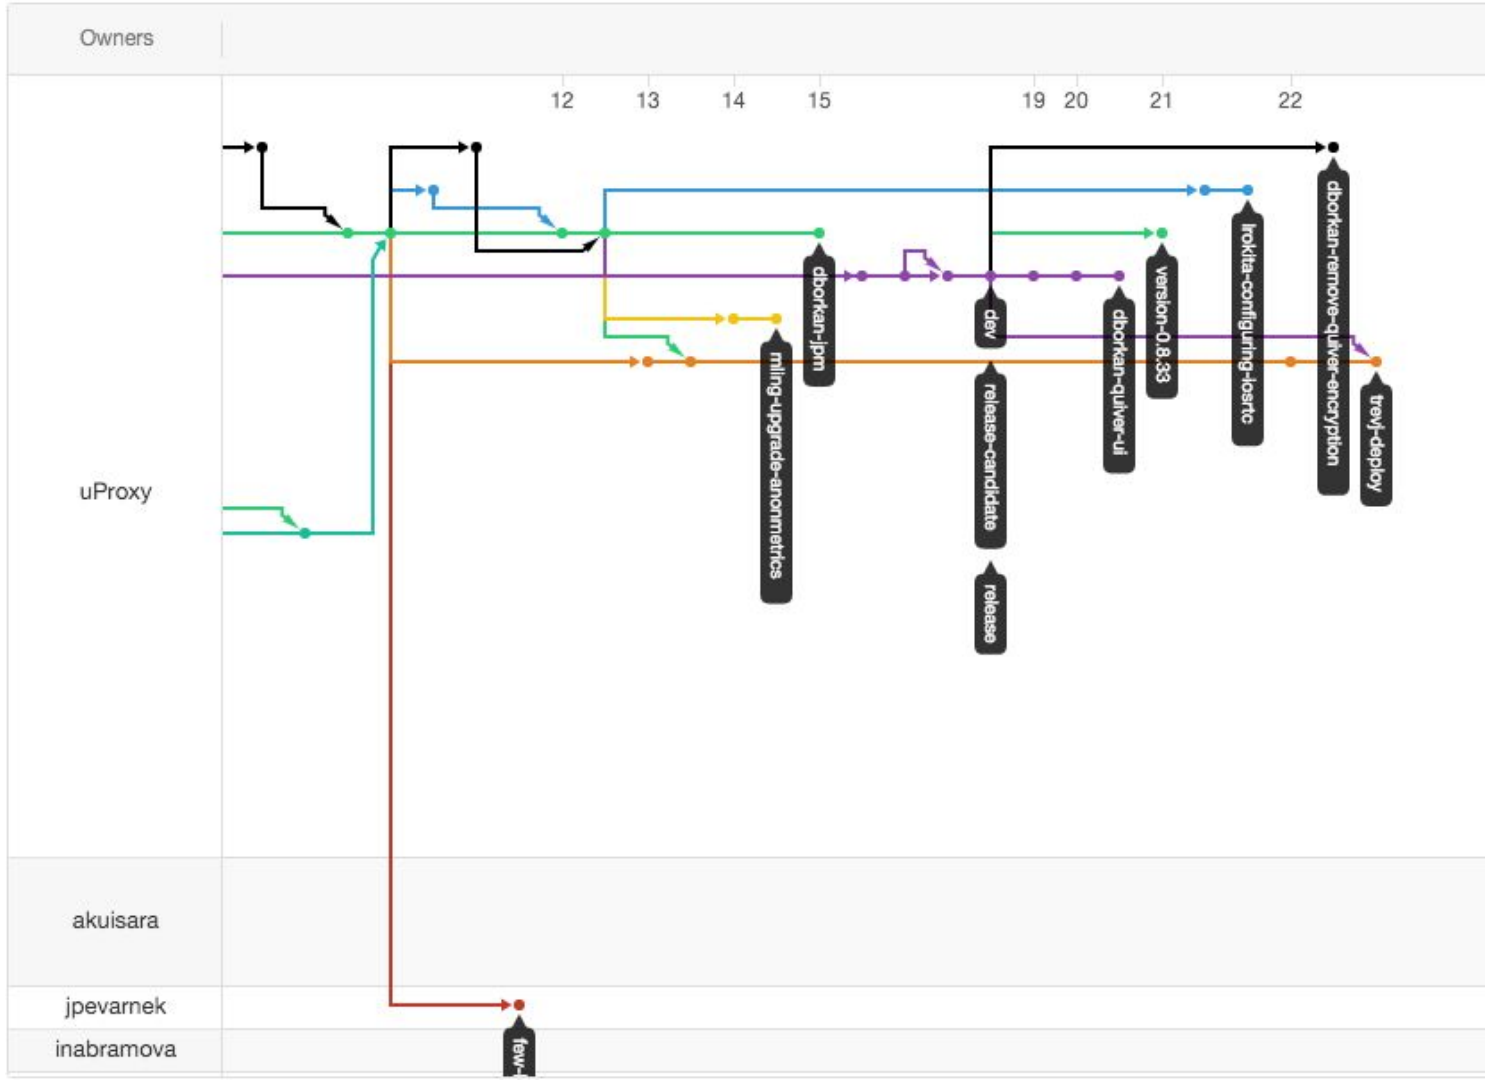
- Complete log of changes (needed for time travel with source code control)
 - Directed acyclic graphs (DAG)
- commit
 - parents
 - deltas (changes to content)
 - hash - for consistency
 - metadata



```

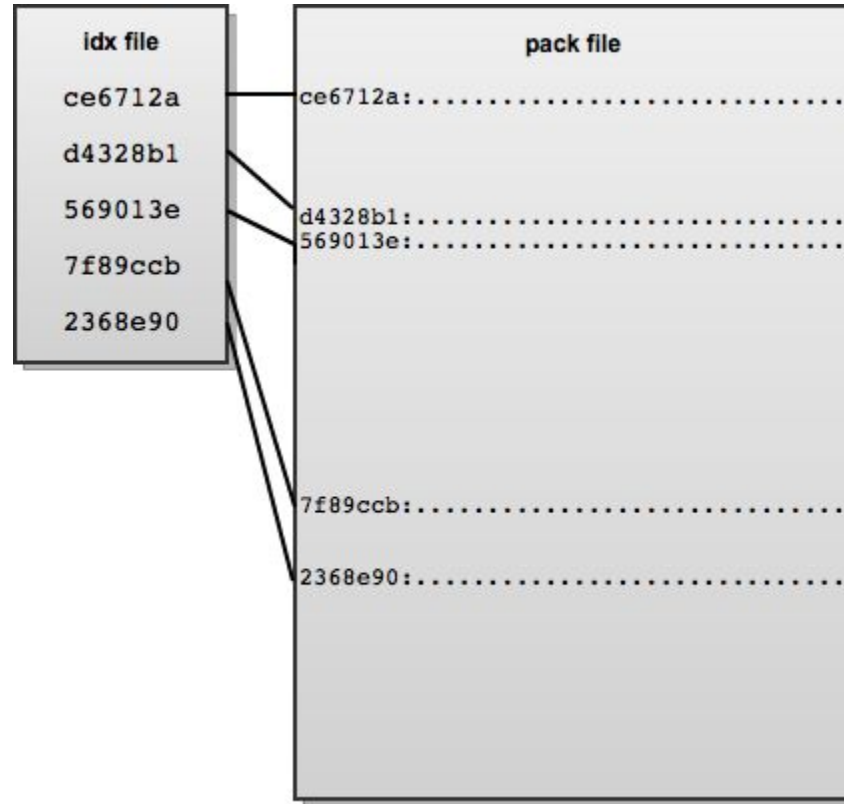
* 686995a jab (9 days ago): improve cloud invite confirm message (HEAD -> dev, tag: v0.8.32, origin/release-candidate, origin/release, ori
* 764f8e5 dborkan (9 days ago): Merge pull request #2180 from uProxy/dborkan-quiver-reconnect
/ \
| * dfd26a2 Daniel Borkan (9 days ago): replace == with ===
| * 9218ea9 Daniel Borkan (9 days ago): Fix Quiver disconnect and automatic reconnect
| * f64ec8c Jonathan Pevarnek (12 days ago): Merge pull request #2172 from jpevarnek/switch-to-lib-rule
/ \
| * a66e18d Jonathan Pevarnek (5 weeks ago): Switch to using buildAndRunTest from uproxy-lib
| * 1987beb Jonathan Pevarnek (12 days ago): Merge pull request #2171 from jpevarnek/separate-model
/ \
| / \
| / \
| * | 0db725e Jonathan Pevarnek (5 weeks ago): Separate the model from the user_interface file
| / \
| * | d50d2fa Jonathan Pevarnek (13 days ago): Merge pull request #2147 from jpevarnek/update-ts
| * | d7d411d Jonathan Pevarnek (13 days ago): Merge branch 'dev' into update-ts
| * | a26fe6a Jonathan Pevarnek (5 weeks ago): Update grunt-ts to the latest version
| * | a991855 dborkan (13 days ago): Merge pull request #2170 from uProxy/version-0.8.32
/ \
| / \
| * | b777228 Daniel Borkan (13 days ago): Bump versions
| / \
| * | 0266021 dborkan (13 days ago): Merge pull request #2169 from uProxy/trevj-uproxy-lib-update (tag: v0.8.31)
/ \
| / \
| * | babbd50 Trevor Johnston (13 days ago): upgrade uproxy-lib for new cloud invite fixes and additions (origin/trevj-uproxy-lib-update)
| / \
| * | e9b5c6c lucyhe (2 weeks ago): Merge pull request #2159 from uProxy/lucyhe-cloudinvite
| * | 17f6057 Lucy He (2 weeks ago): Update getInviteUrl API to take optional parameter too (origin/lucyhe-cloudinvite)
| * | 90b988e Lucy He (2 weeks ago): Merge branch 'dev' of https://github.com/uProxy/uproxy into lucyhe-cloudinvite
/ \
| / \
| * | 6456e95 Laura Rokita (2 weeks ago): Merge pull request #2165 from uProxy/gitlaura-build-for-ios
| * | fa6ca2b lrokita (2 weeks ago): Update gruntfile and index.html to build uProxy on ios (origin/gitlaura-build-for-ios)
/ \
| * | d2cb231 Lucy He (2 weeks ago): Add Cloud descriptions to Share tab.
| * | eee3f4d Lucy He (2 weeks ago): Use showDialog instead of fireSignal to show invite URL.
| * | b4515ba Lucy He (2 weeks ago): Make userId optional for getInviteUrl
| * | ffa5391 Lucy He (3 weeks ago): Fix enum equality bug
| * | c7c30d9 Lucy He (3 weeks ago): Merge with dev
/ \
| / \
| / \

```



Content Addressable Filesystem

.git/objects



Reconciling Conflicts

- Last writer wins (AFS, NFS)
- User-specified conflict handler
- Manual reconciliation (git, svn)
- Conflict-free replicated data types (CRDT)
 - Data types such that conflicts are impossible
 - Operation-based CRDT (e.g. commutative replicated data types)

Operational Transforms

