

Outline

- *Exploiting Weak Connectivity for Mobile File Access*
- Lily B. Mummert, Maria R. Ebling, M. Satyanarayan. In ACM Symposium on Operating Systems Principles (SOSP '95)



Lineage

- Andrew File System (@ CMU) – Mid '80s
 - Wide area distributed file system
 - Cells (ucns.uga.edu), volumes (/students/cs/)
 - /afs/ucns.uga.edu, /afs/acpub.duke.edu/, /afs/ibm.com/
 - Transarc (bought by IBM) commercialized AFS
 - Transarc created DFS (which is used by arches)
- Followup to AFS is Coda (1987-)
 - Provide disconnected service at system level (File system)
 - www.coda.cs.cmu.edu
- Followup to Coda is Odyssey



AFS vs Coda





図 1 Coda のディスクコネクテッド・オペレーション Coda ではネットワークに障害が発生しても処理を続けることができる。ユーザは障害が発生したことに気づかない。

Hoarding

- The process of “storing” files on your mobile device for access on the road
 - E.g.
 - You hoard your calendar in Palm
 - You hoard UGA today in Palm
 - You hoard web pages for disconnected access in IE5
 - You hoard files in laptop that you think will be useful
- Hoard profile says what should be hoarded
 - E.g. (/home/profs/surendar, /bin, /var/mail/surendar)
- “Hoard walk” goes through the hoard profile to download relevant file



Hoarding



図5 Stickyファイルは最後までキャッシュにとどまる Codaではキャッシュするファイルやディレクトリに優先度をつけることができる。高い優先度を備えたファイルやディレクトリほどローカル・ディスクにとどまっている確率が高い。一番高い優先度を Sticky という。

Disconnected Access

- Venus manages file access
 - Venus operates in three modes
 - Hoarding:
 - Normal mode of operation when connected
 - Preserves cache coherence using call backs. When you access a file, you make a local copy and tell the server that if someone else modifies the file, you will get a call back to invalidate your copy
 - Emulating:
 - On disconnection, logs client accesses in Client Modify Log
 - Performs log optimizations to reduce log size
 - Reintegrating state
 - On reconnection, reconciles CML with servers
 - Application specific conflict resolution



Design Rationale

- Don't punish strongly connected clients
 - No write-backs
- Don't make life worse than when disconnected
- Do it in the background if you can
- When in doubt, seek user advice



Transport protocol refinements

- End-to-end arguments
 - Sharing keepalive between Venus and lower layers
 - Change retransmission based on Round Trip Times



Rapid Cache Validation

- Coda - Cache coherency at file level
- Add volume level checks
(/home/students/janedoe/csci4770) before file level
- Volume callbacks



Trickle Reintegration

- Propagate updates asynchronously, in the background
- We don't want reintegrations to take up all the foreground bandwidth when users connect using a slow link
- Used Reintegration barrier to preserve log optimization
- Reintegration chunk based on bandwidth
- User advice



Handling misses

- Augmenting the hoard database
- Modelling user patience

Discussion

