

Eidetic storage: Cheap versions in the cloud

Xianzheng Dou, Peter M. Chen, Jason Flinn
University of Michigan

Cloud-based storage

Ease-of-management
Cost-effectiveness
Reliability

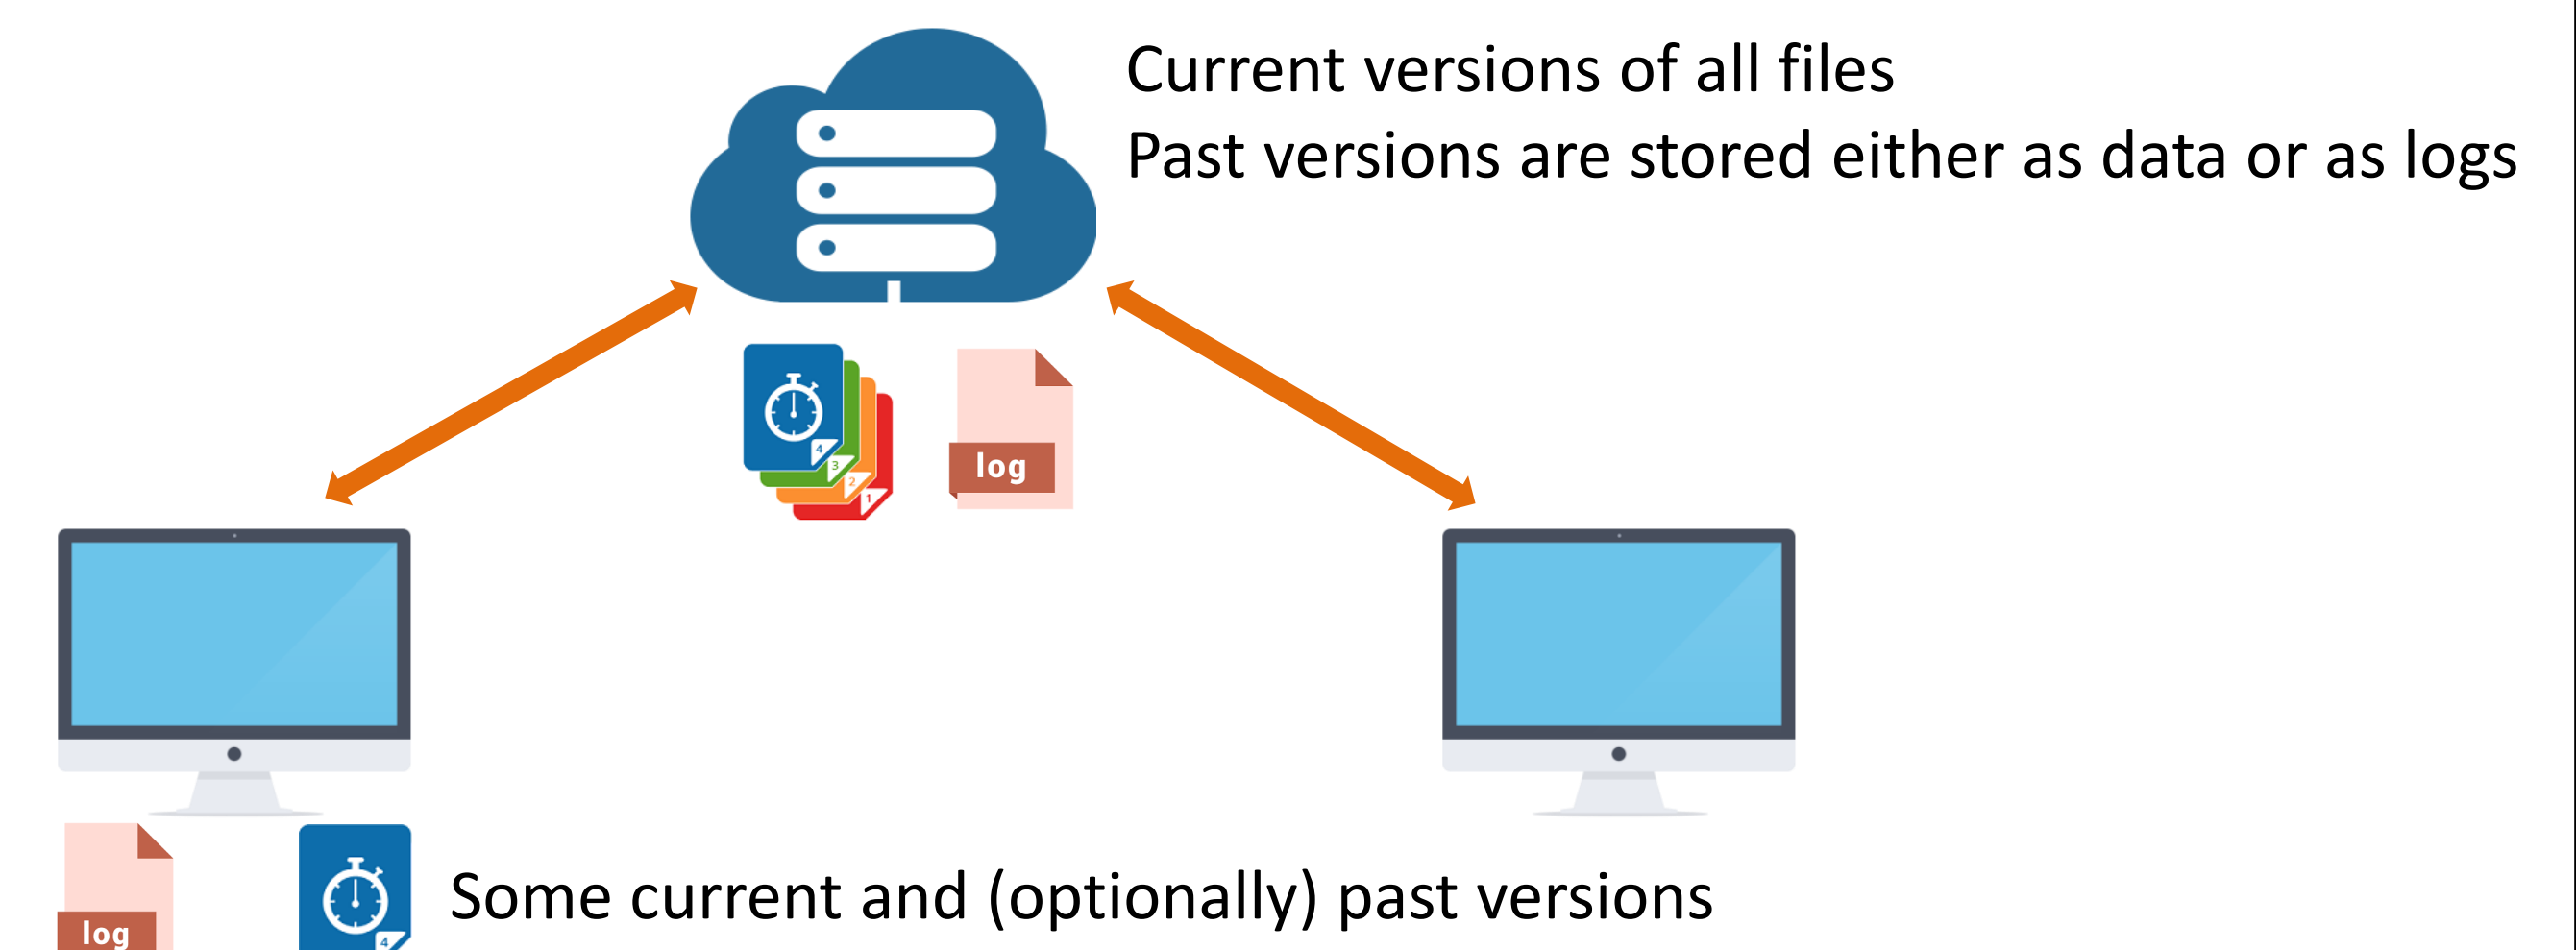
Versioning
Recovery of lost data
Auditing
Troubleshooting

Customers are storing more of their data at cloud-based storage providers
They are also increasingly expecting to have access to past versions of their data

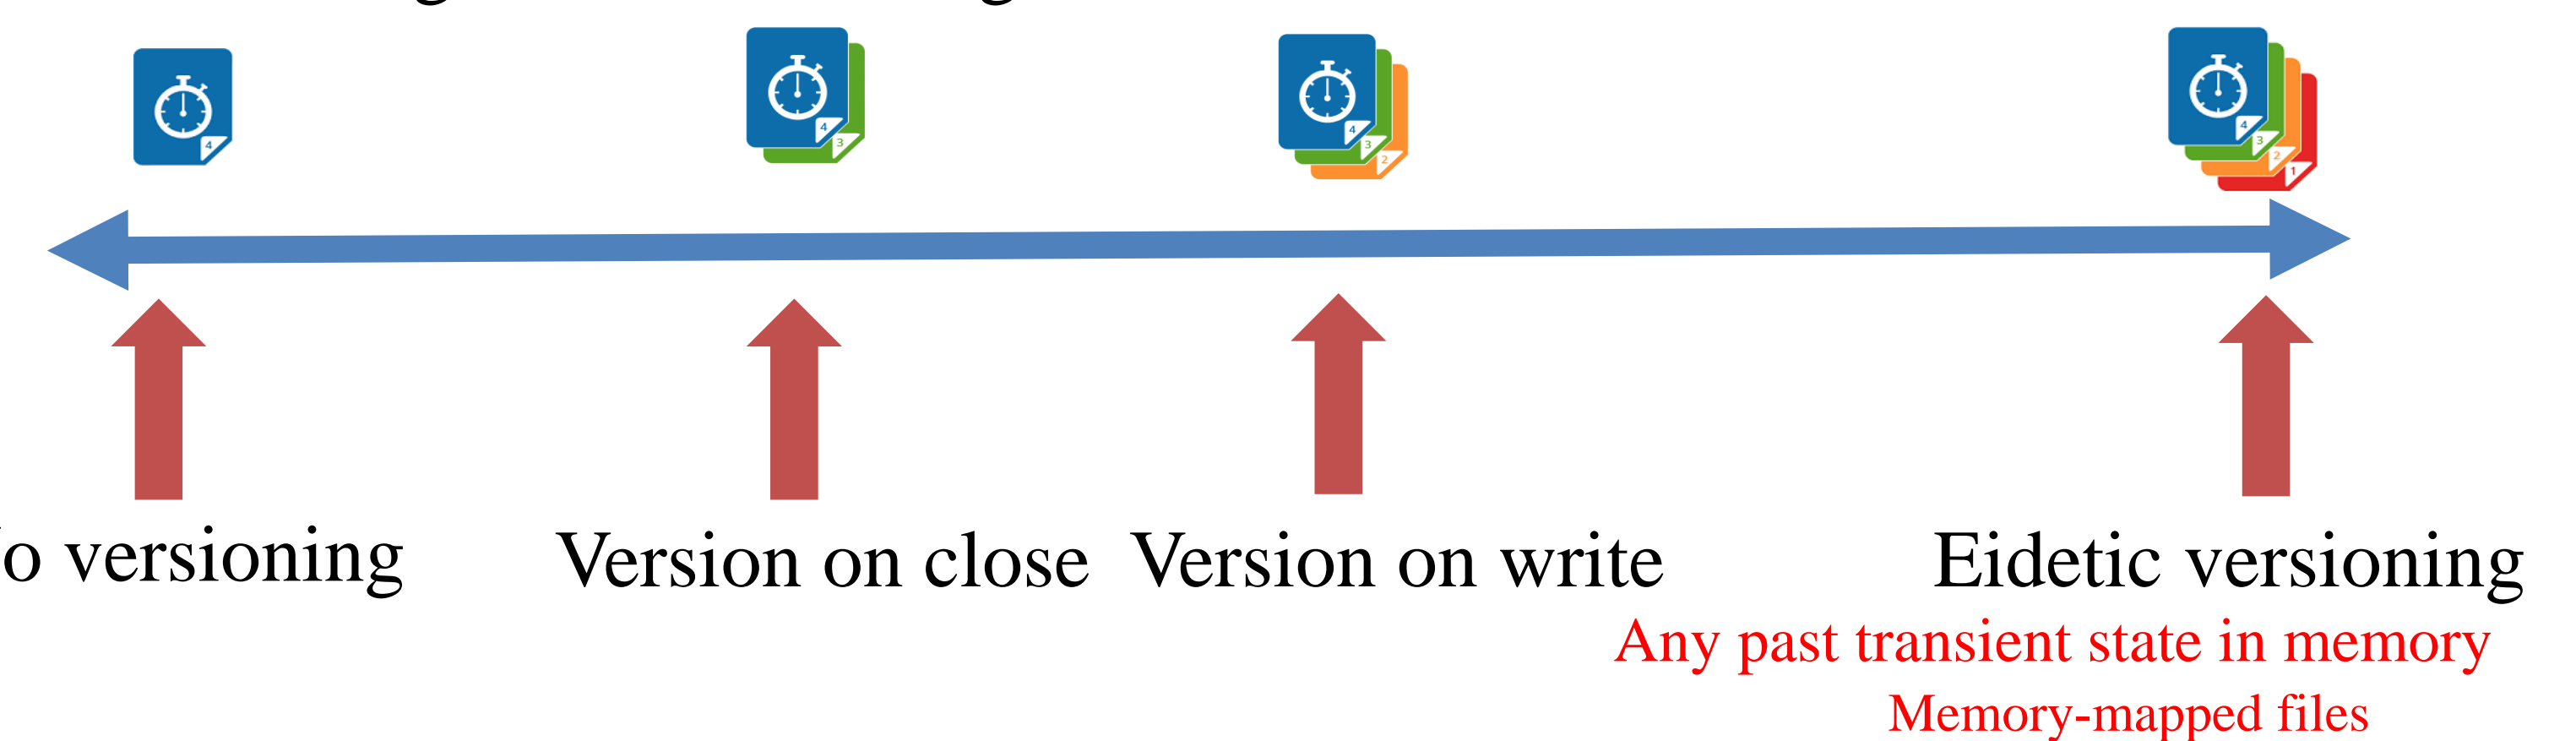


Challenges:
Storage costs
Communication costs


Design



- A general-purpose distributed file system with deterministic recomputation
 - Substantially reduce communication costs
 - Reduce storage costs for existing versioning policies
 - Enables finer-grained versioning

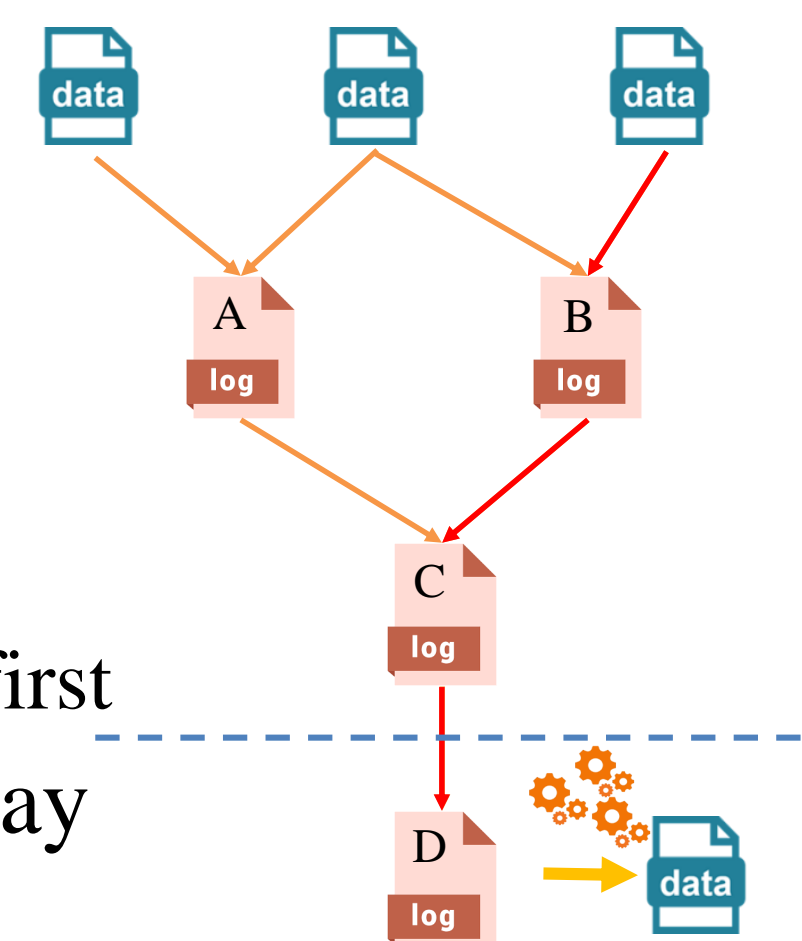


Writing data by values or by operations



Storing data


- Current versions: by values
- Past versions: by values or by operations
 - Recursive regeneration
 - In order to regenerate D, we need to regenerate A, B and C first
 - Regeneration time bound: maximum materialization delay



File: data or computation?


Substitute computation for data

any data can be represented by the computation that generates that data



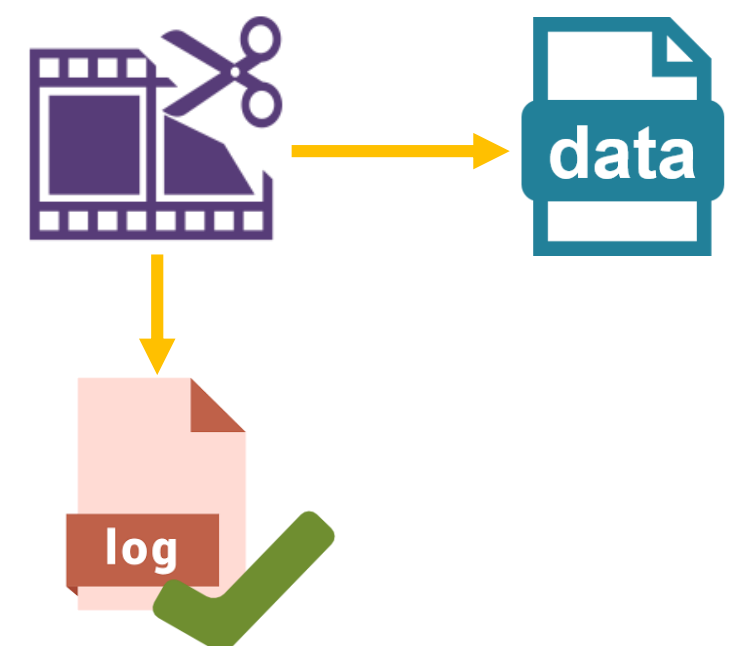
Limitations of previous systems: non-deterministic outputs

Deterministic record and replay




- We can selectively represent a file as:
 - Normal file data (by value)
 - Logs of the nondeterminism needed to recompute the file (by operation)

photo editing



cryptographic key generation



Optimization: log compression


- Chunk-based deduplication is effective for file data
 - Can it also be applied to computation (logs of nondeterminism)?

Log entry	Values
1 open	rc=3
2 pthread_lock	rc=0
3 mmap	rc=<addr>,file=<id,version> rc=<size>,file=<id,version>
4 read	rc=<size>,file=<id,version>
5 gettimeofday	rc=0,time=9.90
6 write	rc=<size>
7 pthread_unlock	rc=0
...	...

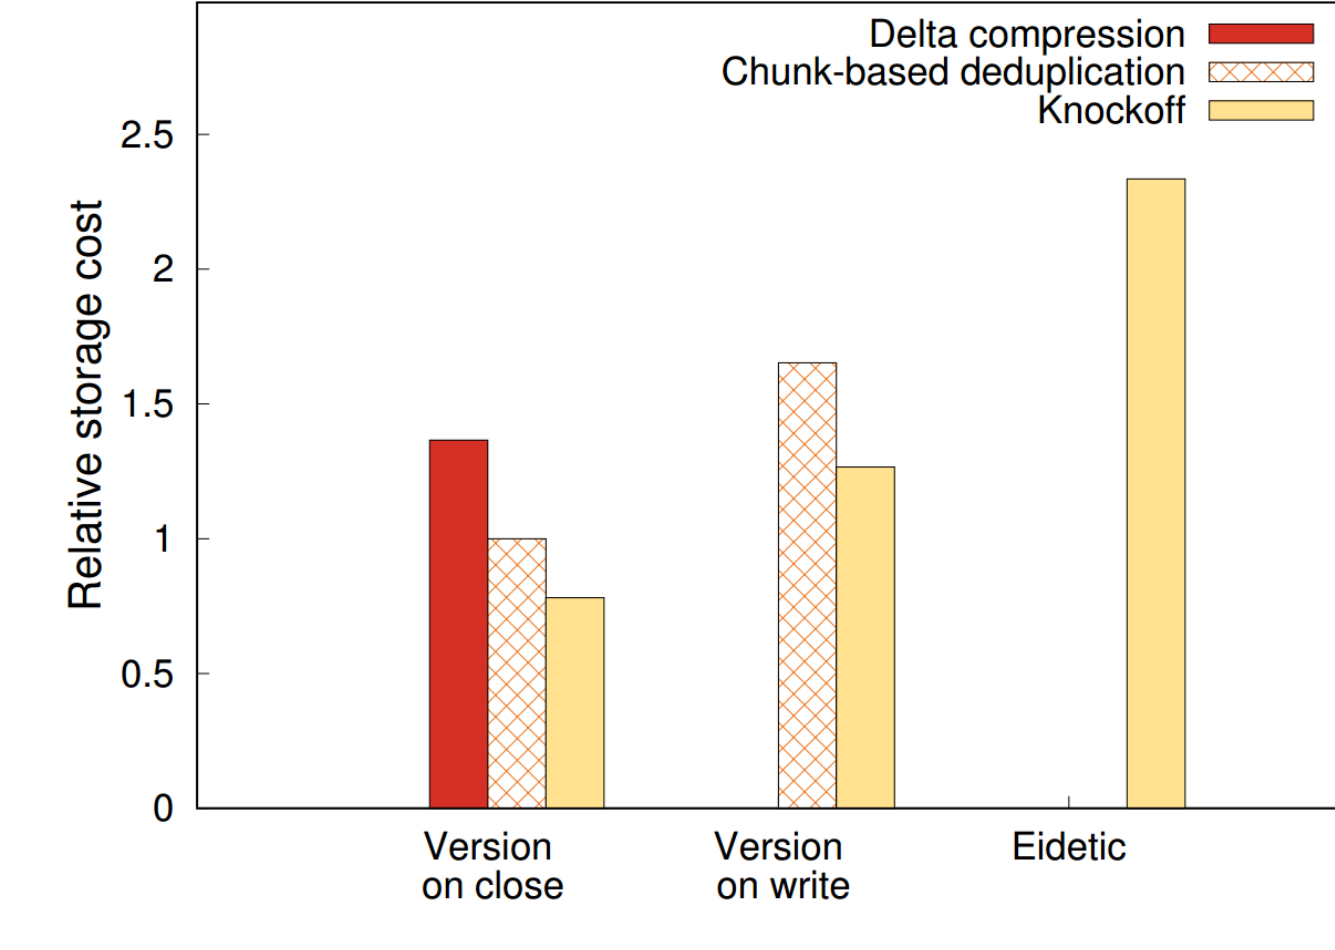
Small differences:
Use delta compression instead

Results

Total bytes sent to the server



Relative storage costs



Variances across users and applications

