


8th Workshop on Principles and Practice of
Consistency for Distributed Data (PaPoC'21)

Access Control Conflict Resolution in Distributed File Systems using CRDTs

**Elena Yanakieva, Michael Youssef, Ahmad Hussein Rezae,
Annette Bieniusa**

Distributed File Systems

- Organize data in files and directories
- Files and directories form a hierarchical structure
- Each entity has a unique identifier, eg. Its path
- Allow multiple users to work concurrently
- High availability
- Security




Operations on distributed file systems

- Creation of files and directories
- Move operations on files and directories
- Content modification of files
- Setting permissions for users and files based on an access control policy
- POSIX access control policy
 - Types of authorization: owner, group, others
 - Types of access: read, write, execute

Challenges

- High availability
 - Replication
- Multiple users work concurrently on the file system content and structure
 - Avoid conflicts without losing data
- **Multiple users work concurrently on the access control**
 - **Avoid data leakage and corruption**
- Security properties
 - Confidentiality – prohibits unauthorized read access
 - Integrity – prohibits unauthorized write access
 - Accessibility – guarantees that authorized access will succeed

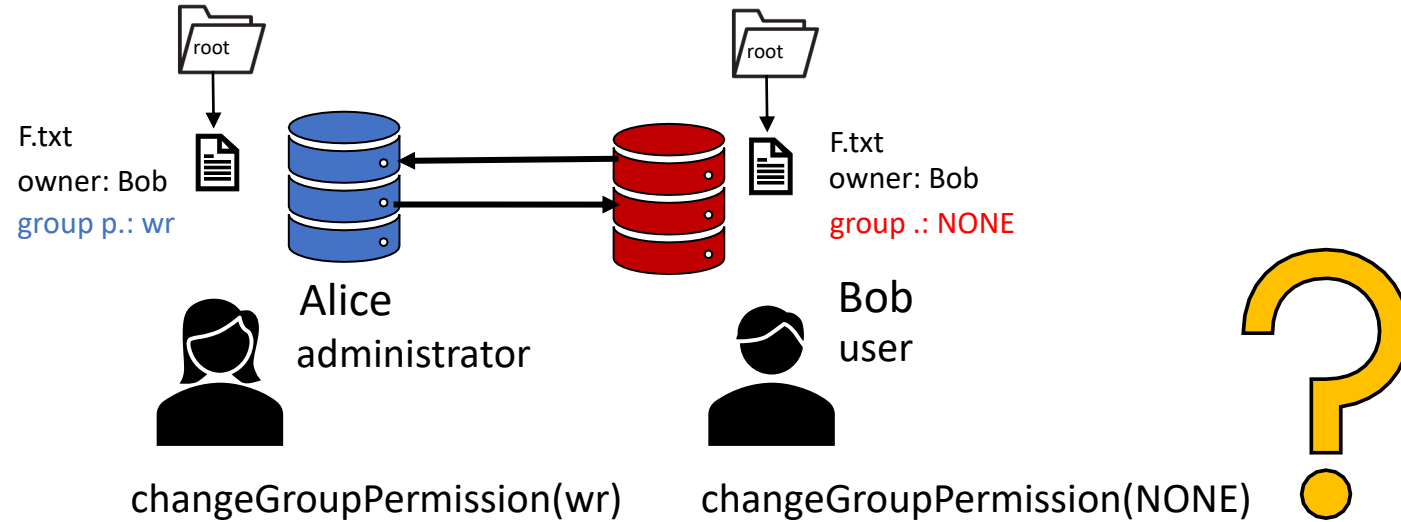
Distributed File Systems using CRDTs



References:

- Mehdi Ahmed-Nacer, Stéphane Martin, and Pascal Urso. 2012. File system on CRDT. CoRR abs/1207.5990 (2012). arXiv:1207.5990 <http://arxiv.org/abs/1207.5990>
- Vinh Tao, Marc Shapiro, and Vianney Rancurel. 2015. Merging semantics for conflict updates in geo-distributed file systems. In Proceedings of the 8th ACM International Systems and Storage Conference, SYSTOR 2015, Haifa, Israel, May 26-28, 2015, Dalit Naor, Gernot Heiser, and Idit Keidar (Eds.). ACM, 10:1–10:12. <https://doi.org/10.1145/2757667.2757683>

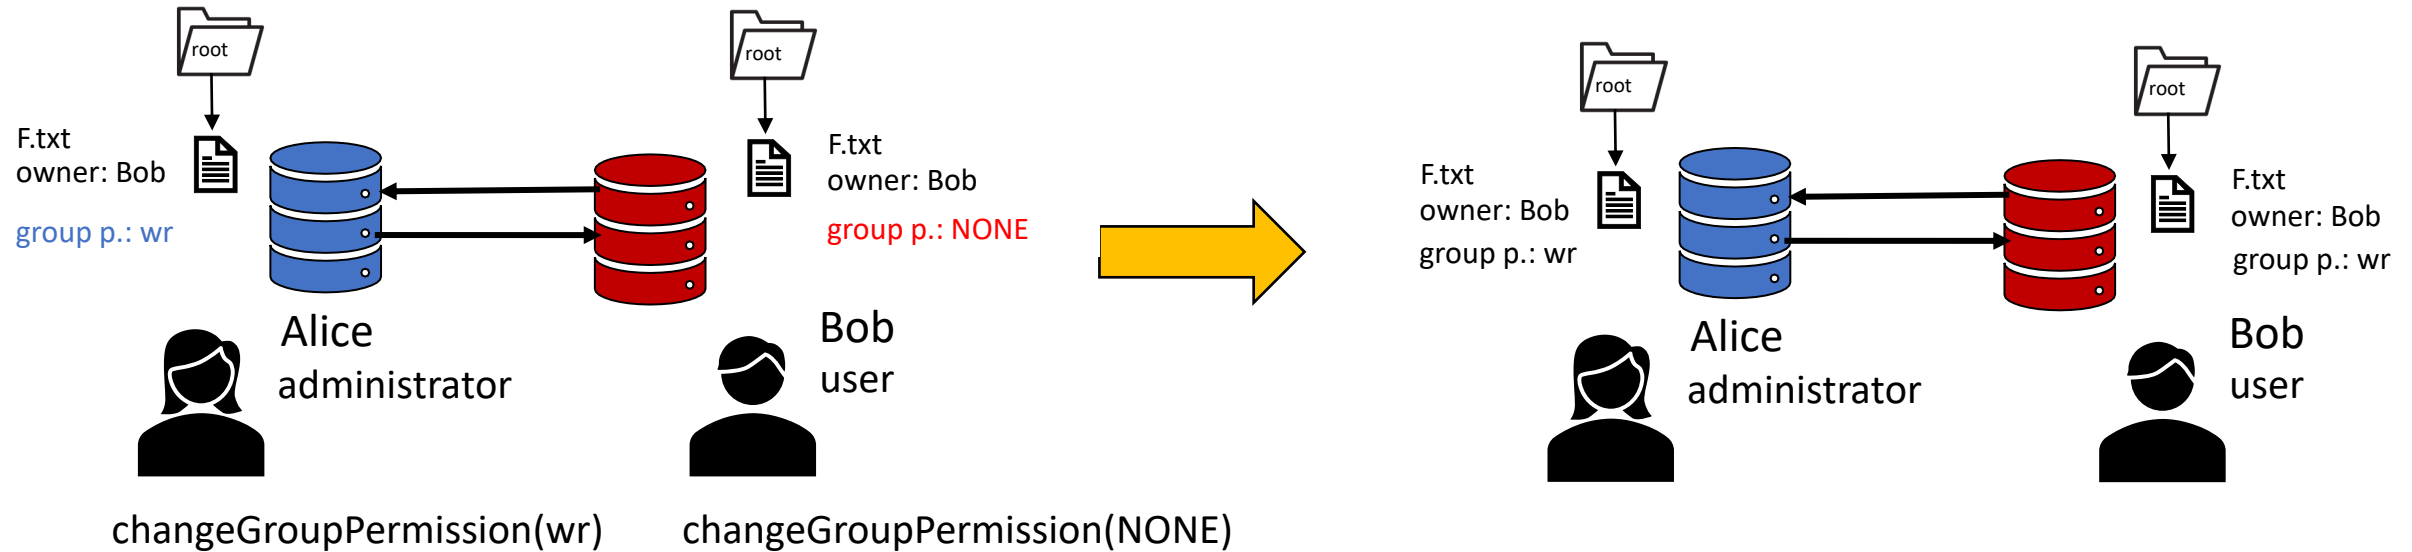
So what happens when...?



Our Contribution

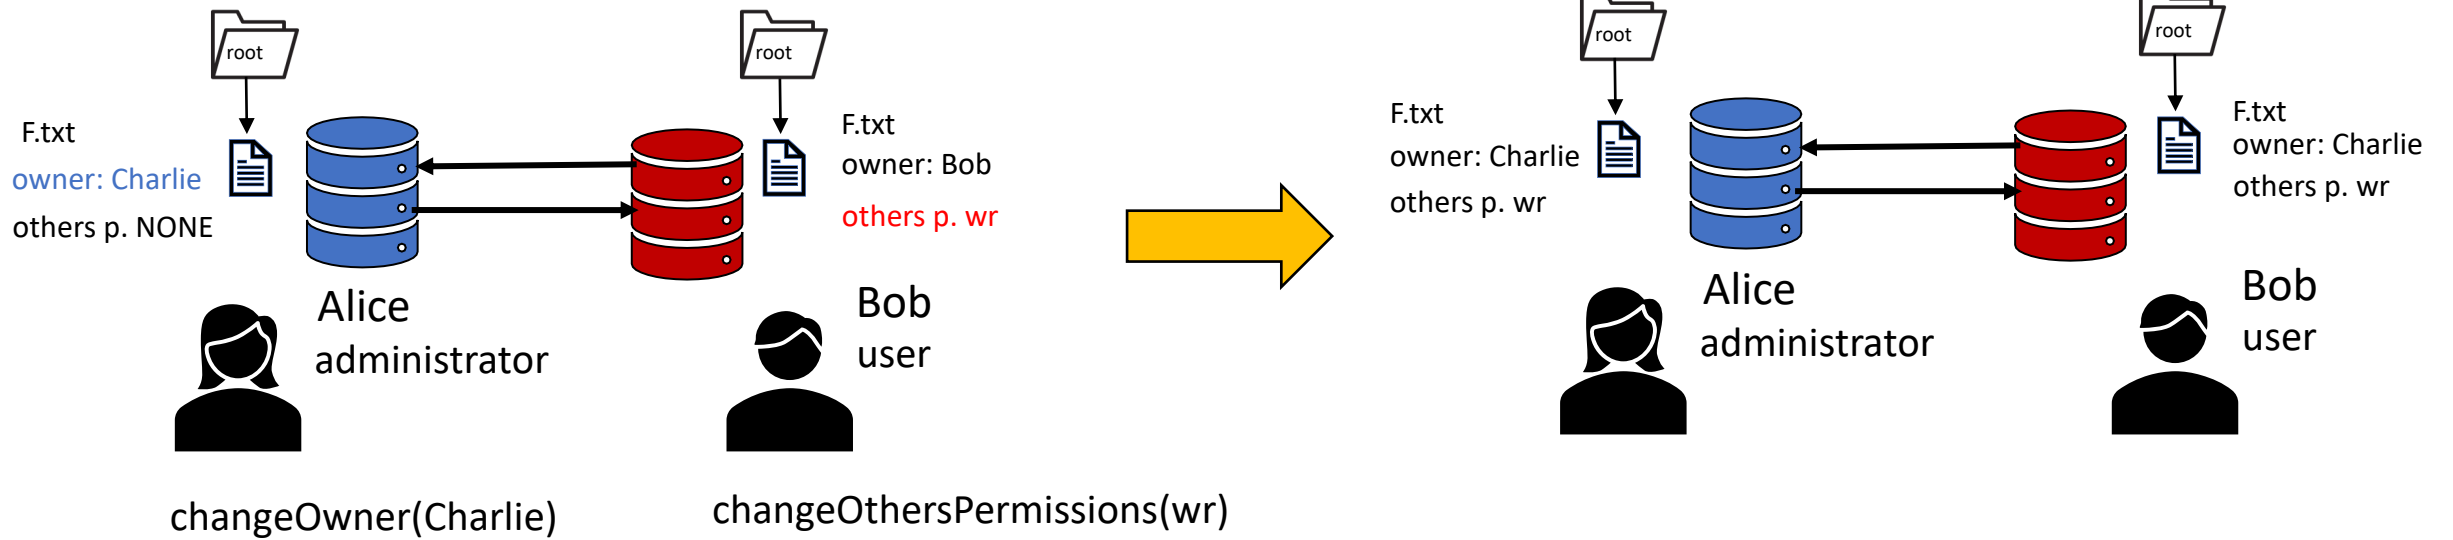
- Formal model of a simplified distributed file system based on CRDTs
 - Focuses on handling access control conflicts
 - Formalized in Repliss, a verification and testing tool with embedded CRDT support
 - Model tested against two invariants for confidentiality and accessibility
-
- <https://github.com/AntidoteDB/crdt-filesystem>

How does our model handle the situation



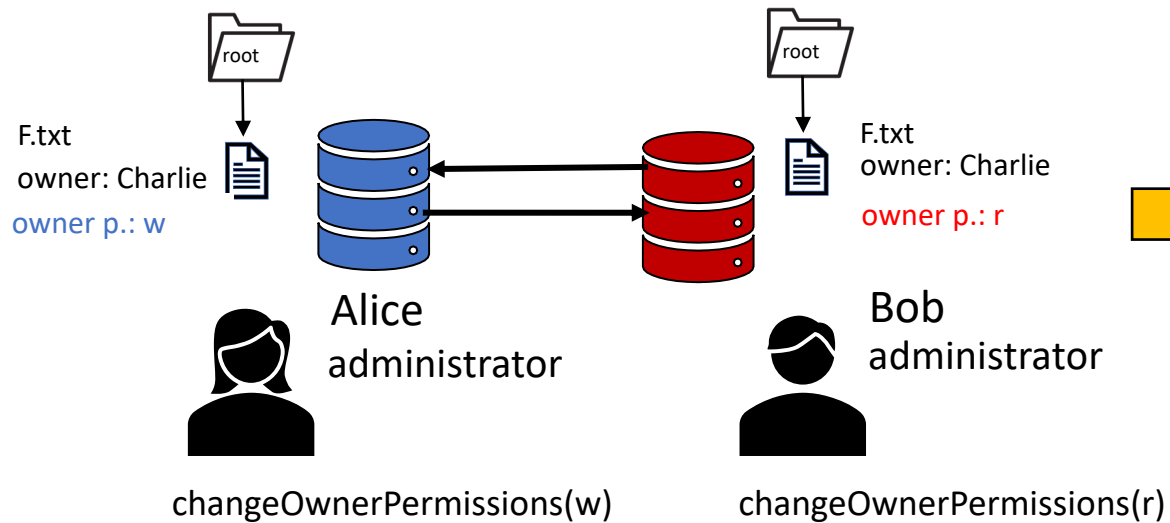
Results.

Atomic CRDT modifications




Results.


Impossibility of confidentiality, integrity and accessibility in a distributed setting



Strategy 1:



Strategy 2:



Summary

- Previous research on CRDT-based distributed file systems
 - Focus on merging content-related conflicts
 - Missing research on resolving conflicts on metadata level
- Our contributions
 - A CRDT-based file system model
 - focuses on access control conflicts
 - CRDTs representing the access control
 - Important insights gathering from testing the model
 - Atomicity of actions is important
 - Impossibility of confidentiality, integrity and accessibility

Thanks for watching! Looking forward to your questions! :)