# PhD thesis proposal – PILOT project PEPR eNSEMBLE

## Collaboration over a distributed file system

**Director of the thesis**: Claudia-Lavinia Ignat, DR Inria, Inria center of Lorraine University
**Co-supervisor of the thesis**: Gérald Oster, MCF, Lorraine University
**Hosting lab**: Inria Center of the University of Lorraine, Nancy
**Lab description**: The Inria Center of the University of Lorraine is one of Inria's nine centers and has twenty project teams, located in Nancy, Strasbourg and Saarbrücken. Its activities occupy over 400 people, scientists and research and innovation support staff, including 45 different nationalities. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institutes, etc.
**Research team**: COAST (Web Scale Trustworthy Collaborative Service Systems)

## Proposal description

File system services are essential for data sharing and collaboration among users. Most of the collaborative file system services such as GoogleDrive and Dropbox rely on a central authority and place personal information in the hands of a single large corporation which is a perceived privacy threat. Users must provide their data to the vendors of these services and trust them to preserve the privacy of their data, but they have little control over the usage of their data after sharing it with other users. Moreover, the centralisation of the platforms hosting these services makes their scalability and reliability very costly. They often limit the number of persons that can simultaneously modify shared data, they generally rely on costly infrastructures and do not allow sharing of infrastructure and administration costs, and centralisation is not suitable for collaboration among a federation of organizations that want to keep control over their data and do not want to store their data at a third party.

A collaborative file system has to support hybrid collaboration including several collaboration modes:
- connected where user modifications are immediately shared and visible to the other users
- disconnected where users are not connected to the network. User modifications will be transmitted to the other users at the reconnection
- ad-hoc collaboration where subgroups of users can work together and synchronise at a later time with other members of the group

Additionally, the collaboration over the file system has to be secure and offer an adapted access control. It should be possible that multiple dynamic administrators can modify users access rights to the shared file system.

We want to build a distributed collaborative file system where control over the data is given to users who can share it directly only with the users they trust and without having to store it at a central authority. The distributed collaborative file system has to support the mentioned collaboration modes and seamless switch from one mode to the others. Additionally, it has to offer a suitable dynamic access control.

Data replication algorithms have to be reliable (i.e. after the reception of all modifications the replicas have to converge) and explainable (i.e., the decisions taken by these algorithms have to be understood by users and their intentions have to be respected). These algorithms have to be suitable for a large community of users that produces a large number of modifications with a high frequency.

As data replication mechanism we propose to use CRDTs (Conflict-free Replication Data Types) [1] that respect Strong Eventual Consistency, a property that ensures convergence as soon as every replica has integrated the same modifications without further message exchange among replicas. CRDTs are suitable for end-to-end encryption in a peer-to-peer environment where data will be decrypted only at the receiver side and conflicts can be resolved locally. There is therefore no need to decrypt data during data transmission as it is the case for centralised architectures where servers require un-encrypted data in order to perform merging. There are two main families of CRDTs: state-based and operation-based [1]. They differ in the way payloads are defined, i.e., how the updates are shared. A payload under state-based CRDT contains the whole data, while the payload under operation-based CRDT carries only a single update. In the context of a hybrid collaboration including connected, disconnected and ad-hoc modes, sending the entire document state after each modification would be inefficient. Operation-based CRDTs are more suitable for our targeted use case.

Several works proposed CRDTs for file systems [2,3] or for trees [4,5]. Most of them rely on state-based CRDTs. For the solutions relying on operation-based CRDTs it rests to be investigated whether the proposed merging semantics satisfy user intentions.

None of the proposed collaborative file systems offers security mechanisms including access control. In order to avoid the use of a central server that stores access rights, we propose that in addition to the replication of data, access rights are also replicated. We want to propose CRDTs for managing replicated file systems and replicated access control by integrating the solution proposed in [6].

We propose to integrate our proposed solution into MUTE [7], our peer-to-peer collaborative editor.

An evaluation with users will be done to test suitable solutions for resolving conflictual changes over the file system and between file system changes and access control rights. The implementation of the proposed solution in MUTE will be also tested with users to evaluate the acceptability of the solution.

Steps:
- Study of literature on CRDTs
- Study CRDTs for file systems
- Design of an operation-based file system CRDT with merging semantics that satisfy user intentions including user studies for validating the resolution strategies in the case of conflicts
- Implement the proposed CRDT into MUTE
- Enhance the proposed CRDT to deal with access rights as described in [6]
- Implement the proposed CRDT for managing replicated file systems and replicated access control into MUTE
- Design user studies for testing the proposed system

## Collaboration aspects

The collaborative aspect of this PhD thesis focuses on remote data sharing, supporting both synchronous and asynchronous collaboration modes, as well as the ability to switch between them. Collaboration can take place in real-time or over extended periods. The aim is to enable collaboration among hundreds of users, leveraging CRDTs, which are inherently scalable and independent of the number of participants. However, challenges related to user awareness and interaction in large-scale collaboration scenarios are beyond the scope of this PhD work.

## Expected results and impact

The contribution of this thesis is threefold.
The primary contribution is theoretical, aiming to design CRDTs that ensure both convergence and security properties. Additionally, the work has a technical dimension, as the proposed approach will be integrated into the MUTE collaborative system, and an empirical component, as it will be evaluated through user studies.

This research lies at the intersection of three domains: distributed systems, security, and human-computer interaction. The results of this work will be disseminated through publications in conferences and journals relevant to these fields.

IPFS is an open-source, peer-to-peer file system based on content addressing, designed as a storage layer for the decentralized web. It benefits from a large and active community of developers and users, which is expected to continue growing. However, IPFS does not natively support mutable data. Integrating a collaborative file system solution with IPFS represents a contribution to peer-to-peer cloud storage, creating a link with the PEPR Cloud initiative and the startup Hivenet (https://www.hivenet.com/), whose infrastructure is based on IPFS.

## Integration into PILOT project:

This subject integrates into the PILOT axis on open technical frameworks for long-term collaboration as it contributes to building a sustainable and safe infrastructure for future forms of collaboration.

# Bibliography:

[1] Marc Shapiro, Nuno M. Preguiça, Carlos Baquero, and Marek Zawirski. Conflict-Free Replicated Data Types. In Xavier Défago, Franck Petit, and Vincent Villain, editors, Stabilization, Safety, and Security of Distributed Systems - 13th International Symposium, SSS 2011, Grenoble, France, October 10-12, 2011. Proceedings, volume 6976 of Lecture Notes in Computer Science, pages 386–400. Springer, 2011. doi:10.1007/978-3-642-24550-3_29.

[2] Mehdi Ahmed-Nacer, Stéphane Martin, and Pascal Urso. 2012. File system on CRDT. https://arxiv.org/abs/1207.5990

[3] Vinh Tao, Marc Shapiro, Vianney Rancurel. Merging semantics for conflict updates in geo-distributed file systems. SYSTOR 2015: 10:1-10:12

[4] Gérald Oster, Hala Skaf-Molli, Pascal Molli, Hala Naja-Jazzar:
Supporting Collaborative Writing of XML Documents. ICEIS (4) 2007: 335-341

[5] Martin Kleppmann, Dominic P. Mulligan, Victor B. F. Gomes, Alastair R. Beresford: A Highly-Available Move Operation for Replicated Trees. IEEE Trans. Parallel Distributed Syst. 33(7): 1711-1724 (2022)

[6] Pierre-Antoine Rault, Claudia-Lavinia Ignat, Olivier Perrin. Access Control based on CRDTs for Collaborative Distributed Applications. TrustCom 2023: 1369-1376

[7] Matthieu Nicolas, Victorien Elvinger, Gérald Oster, Claudia-Lavinia Ignat, François Charoy: MUTE: A Peer-to-Peer Web-based Real-time Collaborative Editor. ECSCW Panels, Demos and Posters 2017