




Region Based File Sharing

The sharing of mutable resources or, more specifically, file sharing has become ever more prevalent in today's interconnected world. As organizations grow, so do their needs for a fast and scalable file sharing infrastructure.

Imagine a network of nodes represented as a weighted undirected graph with each node having specific demands for each file. It is easy to see that having a single node serve as the host for all the files is not a scalable solution. The introduction of replica nodes tries to alleviate this problem but introduces the additional challenge of mitigating synchronization overhead between the different replica nodes whenever a file is updated.

In this thesis, we try to exploit the underlying geographical and organizational structure and then use this information to divide the original network into several connected sub-networks that we call

regions. Each region is then responsible for hosting a distinct subset of the files.



We are interested in investigating some of the arising questions and difficulties in this scenario:

- What is the best file access strategy?
- How best to measure and compare read and write costs of files?
- What is a sensible approach to define a replication network?
- How to assign nodes and files to regions such as to minimize overall read and write costs?
- What additional flexibilities does this approach offer?

Contact

- Robin Fritsch: rfritsch@ethz.ch, ETZ G63
- Béni Egressy: begressy@ethz.ch, ETZ G94