DIAS: Data-Intensive Applications and Systems Laboratory
School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne
Building BC, Station 14
CH-1015 Lausanne
URL: http://dias.epfl.ch/

Efficient Load Balancing of Erasure Encoded Data

Keywords: DBMS, RAID, Storage, Erasure Coding

Problem: Erasure Encoding (EC) is frequently used to prevent data loss when a physical drive or server fails. Data is segmented into blocks and striped across drives/servers. Each stripe of blocks is encoded which results in 1 or more additional parity blocks. For example, in RAID-5, N blocks are encoded to generate one parity block. These N+1 blocks are striped across N+1 drives. If any single drive fails, the block on the failed drive can be recovered.

However, re-encoding data can be expensive. This can be a problem when a system wants to move blocks across drives as it requires costly re-encoding. Systems may wish to move blocks in order to load balance the requests made to each drive but often choose not to due to this cost.

This project aims to devise techniques to mitigate the cost of encoding to enable data management systems to load balance data across multiple drives. Ultimately increasing the end-to-end performance of such systems.

Project: In this project the student will first implement an erasure coding storage interface, potentially leveraging existing open-source implementations of erase coding. The initial implementation will support erasure encoding blocks stored on multiple drives on a single server. The student will analyze the cost of moving blocks across drives, breaking down the cost into the time the data transfer takes and the re-encoding time.

After the initial analysis, the student will work with collaborators to examine possible techniques to reduce the re-encoding cost when blocks are moved across devices.

Plan:

1. Implement an erasure coding storage system in one of C/C++/Rust
2. Develop a benchmarking suite for the system
3. Analysis of initial results
4. Work with a collaborator to leverage the analysis to determine techniques to improve the efficiency of load balancing operations

Supervisor: Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch
Responsible collaborator(s): Hamish Nicholson, hamish.nicholson@epfl.ch

Duration: 6 months