Optimizing Data Locality with Swift/T and FusionFS

Mermer Dupree, Mike Wilde, Justin Wozniak, Ioan Raicu

Abstract

Many of the high-performance computing (HPC) systems use a centralized storage system that is separate from the the compute system. This system is not going to be scalable as we approach exa-scale performance. Distributed file systems can provide the scalability needed for exa-scale computing. FusionFS is a file system designed for HPC systems that achieves large scalability. Swift/T is a high level implicitly parallel scripting language for HPC systems. Swift/T provides automated parallelism and load balancing on a massive scale. Additional optimizations can be achieved by utilizing the features in FusionFS and Swift/T to take advantage of locality.

Swift/T

- Implicit parallel scripting language
- High level dataflow language
- Extremely scalable

FusionFS

- FusionFS is a distributed file system designed with HPC systems in mind
- Meta data is decoupled from data to remove bottlenecks found in most distributed file systems

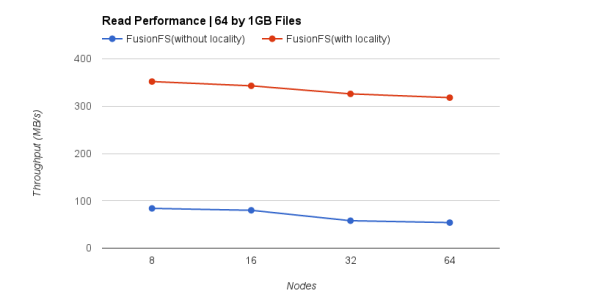

Figure 1: Comparing read performance of FusionFS with locality to FusionFS without locality when scaling the number of nodes. Experiments done with 64 1GB files.

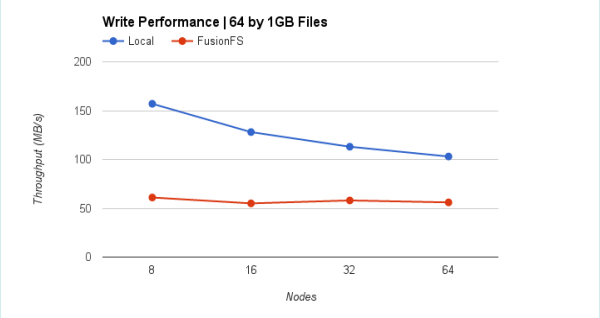

Figure 2: Comparing write performance in FusionFS to writes on local disk when scaling then number of nodes. Experiments done 64 1GB files.

Optimizing Locality

- Use Swift language features to direct tasks close to its data
- The goal is to reduce data movement in the network and improve read performance

Results

- FusionFS gets write performance close to writes on local disk
- Optimizing locality shows noticeable improvements in read throughput

Conclusions

- Our results suggest integrating Swift-T with distributed file systems is a promising area of research
- Future work includes performance tests on larger clusters and comparisons with other file systems

THE UNIVERSITY OF CHICAGO

ILLINOIS INSTITUTE OF TECHNOLOGY

Argonne NATIONAL LABORATORY

UIC UNIVERSITY OF ILLINOIS AT CHICAGO

Mermer Dupree: mdupre2@uic.edu, University of Illinois as Chicago
Mike Wilde: wilde@mcs.anl.gov, University of Chicago; Argonne National Laboratory
Justin Wozniac: wilde@mcs.anl.gov, University of Chicago; Argonne National Laboratory
Ioan Raicu: iraicu@cs.iit.edu, Illinois Institute of Technology; Argonne National Laboratory