

# CS554 Project Ideas

---

## FusionFS:Cache - Cost-Effective Caching for Distributed File Systems

### Overview

FusionFS [1] is a new distributed file system designed for exascale systems. One potential drawback with FusionFS is the limited space and high cost of non-volatile memory on each compute node. We have designed and implemented a middleware, namely HyCache [2], comprised of a hierarchy of heterogonous storage devices (e.g. SSD, HDD, etc.) to achieve an SSD-level performance while retaining the HDD-level capacity and cost. HyCache has been evaluated on the Hadoop File System, and showed its effectiveness for some distributed applications. In this project, you will implement a HyCache-like middleware for FusionFS, and evaluate it with both micro benchmarks and real applications. The implementation would be merged into the next release of FusionFS.

### Relevant Systems and Reading Material

Please read the following papers and their references before submitting your proposal:

[1] Dongfang Zhao, Chen Shou, Zhao Zhang, Iman Sadooghi, Xiaobing Zhou, Tonglin Li and Ioan Raicu. FusionFS: a distributed file system for large scale data-intensive computing, *2<sup>nd</sup> Greater Chicago Area System Research Workshop*, 2013. Available online: [http://datasys.cs.iit.edu/~dongfang/download/poster\\_fusionfs\\_revised.pdf](http://datasys.cs.iit.edu/~dongfang/download/poster_fusionfs_revised.pdf)

[2] Dongfang Zhao and Ioan Raicu. HyCache: a User-Level Caching Middleware for Distributed File Systems, *International Workshop on High Performance Data Intensive Computing*, 2013. Available online: [http://datasys.cs.iit.edu/~dongfang/download/hycache\\_HPDI\\_CRC\\_v1.pdf](http://datasys.cs.iit.edu/~dongfang/download/hycache_HPDI_CRC_v1.pdf)

### Preferred/Required Skills

Principles: operating system, distributed systems, computer network

Programming: Shell Script, Perl/Python, C, C++, PThread, sockets, FUSE

Operating systems: Linux

### Project Mentor

Dongfang Zhao

Email: [dzhao8@hawk.iit.edu](mailto:dzhao8@hawk.iit.edu)