Big Data Systems (Ioan Raicu)

Extracting knowledge from increasingly large data sets produced, both experimentally and computationally, continues to be a significant challenge for scientific discovery. Experimental fields, such as high-energy physics report that experimental data sets are expected to grow by six orders of magnitude or so in coming years. For computational fields (e.g. fusion science) in which energy codes run on a million cores, data is generated in bursts of 2 PB/sec with checkpoints every 10 minutes, producing an average of 3.5 TB/sec over the entire run of an experiment. This could be orders of magnitude larger in future systems.

We will develop methods to index vast amounts of scientific data irrespective of its location and the particular storage system in which it is stored. We aim to provide a scalable search index that allows researchers to find, browse, and discover disparate scientific data based upon file system metadata and metadata buried within scientific data formats. In this project, students will explore methods for indexing large amounts of distributed data. The resulting index is intended to be applicable both to tightly coupled distributed file systems and also loosely coupled distributed storage systems such as those managed by Globus. Students will develop models for extracting an array of file system and file-based metadata, investigate methods for efficiently indexing large amounts of data to support common query models (e.g., free-text search), and analyze the performance of their approaches under a range of scenarios and using real data. Students may also explore scalable methods for identifying and extracting scientific metadata as well as algorithms for synchronizing the index with frequently changing, distributed data.

This work will evaluate the practicality of several information retrieval libraries (e.g. Lucene, CLucene, Xapian and Lemur) as well as more fundamental classical data structures (e.g. inverted indexes, multi-level skip lists, B-Trees, and trie). Evaluation will focus on metrics such as indexing throughput, query throughput, and query latency. The analysis would continue with a thorough look over the inherent parallel capabilities of building indexes and the eventual distributed nature of the index placement, while also providing metrics for fairly comparing network usage for inter component communication (in the distributed case). In exploring solutions for the distributed case, students would study topology-aware strategies for placing indexes, scalable distributed data structures and minimalistic communication interferences.

This work will extend the PI's NSF funded project FusionFS [**1**, 2, 3] project by adding distributed indexing and search capabilities in the data management layer. The work will also aim to deliver such indexing and search capabilities to other mainstream storage systems, such as GPFS, Lustre, and NFS, as well as those accessible via Globus. Initial evaluation will be over file system dumps from production petascale storage systems from large supercomputing centers. Some early research results have been explored in [4]

These projects are suitable for undergraduate students as they typically require knowledge from students on data structures, algorithm analysis, and systems programming, all typical courses in most undergraduate programs around the country. With one on one guidance from graduate students in the lab, undergraduate students will thrive on projects related to distributed indexing and search, giving them a deeper appreciation and understanding of tools they use every day in their lives (e.g. google.com, spotlight search, etc.).

[1]   Dongfang Zhao, Ning Liu, Dries Kimpe, Robert Ross, Xian-He Sun, and Ioan Raicu. "Towards Exploring Data-Intensive Scientific Applications at Extreme Scales through Systems and Simulations", IEEE Transaction on Parallel and Distributed Systems (TPDS) Journal 2015

[2]   Dongfang Zhao, Zhao Zhang, Xiaobing Zhou, Tonglin Li, Ke Wang, Dries Kimpe, Philip Carns, Robert Ross, and Ioan Raicu. "FusionFS: Towards Supporting Data-Intensive Scientific Applications on Extreme-Scale High-Performance Computing Systems", IEEE International Conference on Big Data 2014

[3]   Ioan Raicu. "NSF CAREER: Avoiding Achilles' Heel in Exascale Computing with Distributed File Systems", NSF Award Number 1054974, 2011 - 2018. http://datasys.cs.iit.edu/grants/NSF-CAREER/index.html

[4]   Alexandru Iulian Orhean, Itua Ijagbone, Dongfang Zhao, Kyle Chard, Ioan Raicu. "Toward Scalable Indexing and Search on Distributed and Unstructured Data", IEEE Big Data Congress 2017