

Semantic-Aware Metadata Organization

Yifeng Zhu

zhu@eece.maine.edu

Electrical and Computer Engineering



Hong Jiang

jiang@cse.unl.edu

Computer Science and Engineering



Need New Namespace Scheme

Needle in a haystack

- Approaching Exa-byte
- Billions of files
- Need better search and index capability

“Hierarchical File Systems are Dead”

Margo Seltzer, et al, Harvard, HotOS 2009

“Consideration for very revolutionary ideas such as new approaches to name spaces ... needs to be explored. ... There has been basically no progress in this alternative tree based metadata area.”

- *HEC FSIO 2008 Metadata Gap Report*

Alternative to Hierarchical Directory Tree

- **Limitation of conventional file systems: scalability and functionality**
 - Only provide pathname-based access
 - Linearly brute-force searching for complex metadata queries
 - Physical metadata layout cannot adaptively match workload patterns for better performance
 - Particularly serious for I/O intensive workloads in HEC

What is the alternative to hierarchical directory tree?

Alternative to Hierarchical Directory Tree

- Metadata indexing and query-based model have great potentials.
- Database management systems (DBMS) do not satisfy the scalability requirements of metadata search in large-scale file systems
 - Designed for generic purposed applications, ill-suited for file systems
 - We need light-weight indexing that takes full advantage of important characteristics of file systems

Approach: Semantic-aware Metadata Organization

- A new metadata organization foundation with semantic-awareness
 - Provide fast and flexible metadata query services and abstract file systems in a highly scalable and reliable way
 - Enhance functionality for users and user applications
 - effectively limit search scope of a complex metadata query
 - avoid or alleviate brute-force search in the entire system
 - Improve I/O performance by leveraging semantic correlations
 - De-duplication, disk layout, prefetching, buffer caching

Semantic-aware Metadata Organization

- Files are often semantically correlated
 - Attribute-based semantics
 - Content semantics
 - Context semantics
 - Application semantics
 - Workload semantics (e.g, file access sequence or dependence)

Is metadata semantics really useful?

Preliminary work: metadata prefetching

- Evaluate Attribute-based Semantic Distance

$$Sim(A, B) = \frac{|A \cap B|}{|max(A, B)|}$$

A, B : selected attribute vectors of file A and B

$A \cap B$: count of equal or similar attributes

$max(A, B)$: maximum number of attributes

- Estimate Access Frequency

$$F(A, B) = N_{AB}/N$$

$F(A, B)$: probability that file B is the successor of file A

N_{AB} : number of occurrences that file B is the successor of file A

N : total number of accesses of file A

- File Correlation Degree

$$R(A, B) = Sim(A, B) \cdot p + F(A, B) \cdot (1 - p)$$

Query Abstraction

- **Point Query: conventional pathname-based search**
- **Range Query: locate files whose attribute values are within a given range**

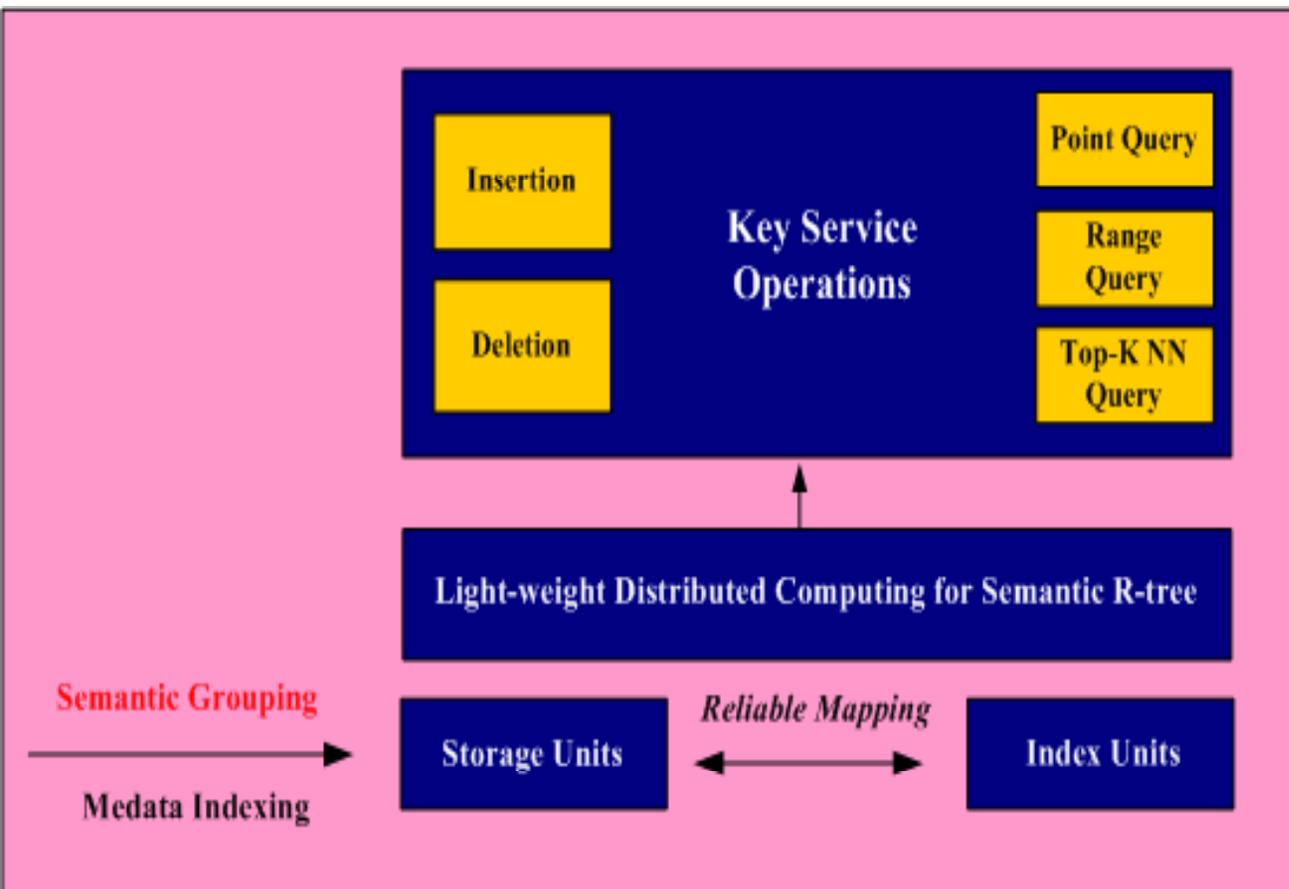
Which experiments did I run yesterday that took less than 30 minutes and generated files larger than 2.6GB?

- **Top-K Query: locate k files which are closest to a given description**

I cannot accurately remember a previously created file but I know that its file size is around 300MB and it was last visited around Jan.1, 2008. Can the system show 10 files that are closest to this description?

System Architecture

- Key functional components:
 - The **grouping component** that classifies metadata into storage and index units based on indexing and retrieval methods, such as the Latent Semantic Indexing (LSI)



- The **construction component** that iteratively builds semantic R-trees in a distributed environment
- The **service component** that supports insertion, deletion in R-trees and multi-query service.

Education Focus

- Graduate Training
- Undergraduate Research
 - NSF REU site Award in Supercomputing (PI: Zhu)
 - Support ten students each summer for ten-week research on campus
- K-12 Outreach
 - NSF ITEST Award (co-PI: Zhu)
 - Maine Laptop programs in middle schools and high schools
 - Understand computing and performance in an age-appropriate way

Acknowledgement



National Labs