Describing HPC Filesystem Trees with the Grand Unified File-Index

Jenna M. Kline, The Ohio State University | Jason Lee (advisor), Rusty Davis (advisor), Los Alamos National Lab HPC Design Division

Introduction

PROBLEM STATEMENT

Analyze the shape of HPC filesystem trees to describe both their immediate characteristics and how these characteristics change over time.

MOTIVATION

- HPC file systems are extremely complex and difficult to predict how they grow and change over time
- Previous studies of file system characteristics largely pre-date modern HPC file systems of the last decade

With GUFI, it is possible to quickly query HPC file system metadata for large-scale system analysis.

GRAND UNIFIED FILE-INDEX (GUFI)

- Makes querying file systems that would be impossible to accomplish with POSIX tools doable
- Single index for all filesystems
- Allows for novel analysis of file system tree metadata as this volume of data was previously very difficult to obtain due to time and resource constraints.

SCORECARD FOR SYSTEM ADMINISTRATORS

- Add-on to GUFI to allow system admins to quickly generate high-level view of filesystem directory
- Command line tool returns queries & analysis for depth, density, weight, width, sparsity, and popularity for a file system by directory and subtree
- Use case: identifying a subtree with many children taking up a disproportionate amount of metadata

Methodology

HPC FILE SYSTEMS ANALYZED

LANL maintains numerous massive filesystems in order to store data from scientific simulations.

File System	Total Size	Directories Count	Files Count	Depth
anony	411 TB	7.4 M	157.6 M	37
scr3	975 TB	2.2 M	59.2 M	59
scr4	1 PB	5.1 M	118.9 M	37
ttscratch	423 TB	5.5 M	117.5 M	37
yellprojs	28 TB	14.1 M	133.3 M	112
yellusers	2 TB	1.6 M	12.9 M	48
Totals	**2.8 PB**	**36 M**	**600 M**	**-**

METRICS TO DESCRIBE HPC FILE SYSTEMS

- Used GUFI to analyze 6 HPC file systems
- Identified 8 metrics that uniquely distinguish file system trees

Metric	Short Name	Definition
Directory Avg File Count	**Density$_d$**	Average number of files in a directory by level
Subtree Avg File Count	**Density$_s$**	Average number of files in a subtree by level
Directory Avg Size by Level	**Weight$_d$**	Size (bytes) of a directories averaged by level
Subtree Avg Size by Level	**Weight$_s$**	Size (bytes) of subtrees averaged by level
Directory Width	**Width**	Total number of directories at each level
Tree Depth	**Depth**	Maximum level of a tree
Directory Sparsity	**Sparsity**	Count of empty directories and files in a tree
User popularity	**Popularity**	Count of unique user ids

DATA ANALYSIS

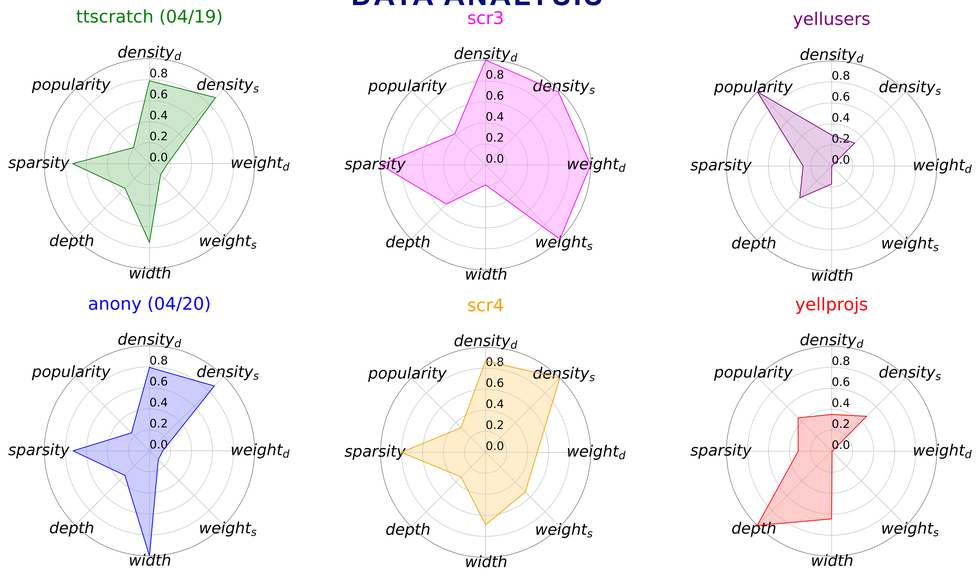

Results

KEY TAKEAWAYS

- **Top ~20 levels are most unique to each file system**
- **Branching factor not a meaningful distinguishing measure**
- **Directory shape stays consistent over time**
- **Directory width may widen as it ages**
- **Scratch directories tend to be larger and denser**

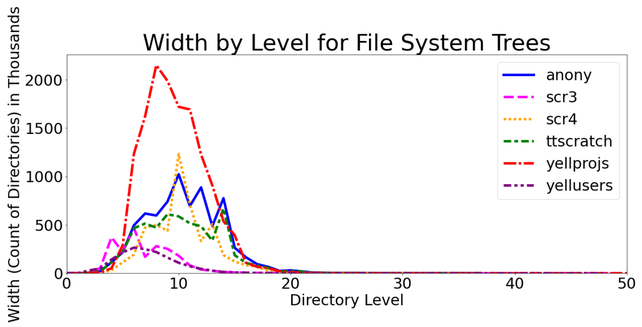

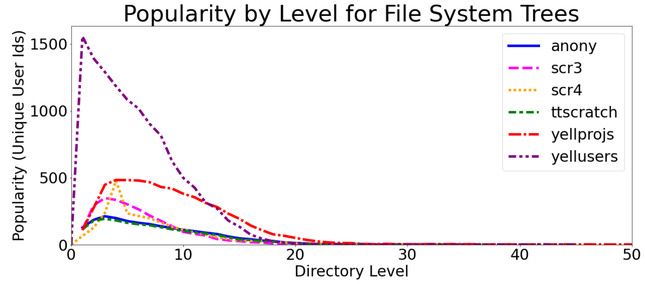

Charts. Width and Popularity by Level for File System Trees

Future Work

- Analyze additional HPC file system directories and how they change shape over time
- Use GUFI to build more complex metrics to describe directories
- Predict requirements of future HPC file systems

Managed by Triad National Security, LLC, for the U.S. Department of Energy's NNSA.

LA-UR-23-28917

References

[1] A. Leung, S. Pasupathy, G. Goodson, and E. L. Miller, "Measurement and Analysis of Large-Scale Network File System Workloads," presented at the USENIX Annual Technical Conference, Jun. 2008.
[2] D. Manno et al., "GUFI: Fast, Secure File System Metadata Search for Both Privileged and Unprivileged Users," in SC22: International Conference for High Performance Computing, Networking, Storage and Analysis, Nov. 2022, pp. 1–14. doi: 10.1109/SC41404.2022.00062.
[3] S. Dayal, "Characterizing HEC Storage Systems at Rest," 2008.
[4] N. Agrawal, W. J. Bolosky, J. R. Douceur, and J. R. Lorch, "A five-year study of file-system metadata," ACM Trans. Storage, vol. 3, no. 3, pp. 9-es, Oct. 2007, doi: 10.1145/1288783.1288788.