Snapshot Judgements: Obtaining Data Insights Without Tracing

Wenxuan Wang
Emory University

Ian F. Adams
Intel Labs

Avani Wildani
Emory University

Metadata snapshots are a favored method for gaining filesystem insights due to their small size and relative ease of acquisition compared to access traces [2]. Since snapshots do not include an access history; typically they are used for relatively simple analyses such as file lifetime and size distributions, and researchers still gather and store full block or file access traces for any higher level analysis such as cache prediction or scheduling variable replication [1, 3]. We claim that one can gain rich insights into file system and user behavior by clustering metadata snapshots and comparing the entropy within clusters to the entropy within natural partitions such as directory hierarchies or single attributes. We have preliminary results indicating that agglomerative clustering methods produce groups of data with high information purity, which may be a sign of functional correlation.

While many studies have analyzed metadata snapshots, most focus on simple statistics, such as file size, age, or extension, or they attempt to reconstruct dynamic trace information from a series of snapshots by interpolating inter-snapshot accesses. We focus instead on what can be learned about a system by looking at metadata correlations within a small set of widely spaced snapshots. For example, timestamps can give insight into the dynamic activity of the system from a purely static viewpoint. UIDs can be used in conjunction with file paths to figure out if there is a "typical" namespace structure users create. Entropy between members of a namespace can help us relate different segments of a trace [4].

Full I/O traces are always superior, but keeping complete logs of accesses is prohibitive in many systems because of the computational overhead to collect the logs and the storage overhead to keep them. For a modern storage system with hundreds of thousands of I/Os per second, storing even minimal representations of the I/O without any metadata is very costly. For example, an enterprise storage system may create over 16 GB of block-level I/O logs per day [5]. Moreover, storing complete traces with metadata is even harder than storing raw accesses because there is more overhead both in terms of size and performance, thus this information is usually lost.

We examined a series of clusterings using HPC and archival snapshots from Los Alamos National Laboratory (LANL) and a set of workstation snapshots collected at Emory. Figure 1 shows a representative LANL snapshot under different clusterings. To analyze storage snapshots, we need an unsupervised learning algorithm that can support n-dimensional, non-linear heterogeneous data with low inter-cluster separation while ideally encoding hierarchical relationships between both data and labels without overfitting. Additionally, our algorithm should handle arbitrary numbers of sparse, binary dimensions to encode Boolean queries about the files in the snapshot , such as queries over the "permissions" and "path" fields. Agglomerative clustering is a well established, interpretable unsupervised machine learning algorithm that fulfills all of the requirements for snapshot analysis as well as providing a natural hierarchy for correlating user locality with path. We also tested k-means and a simple grouping based on a single attribute at a time. To measure cluster validity we calculate the joint Shannon entropy across a set of features of interest within each clustering.

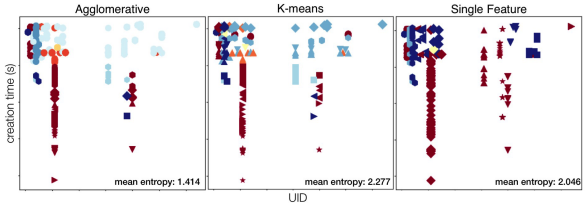

Figure 1: Sample clusterings view for a single snapshot. Clusters are indicated by shape and modification time is indicated by color.

Overall, we found that, regardless of the features considered, the entropy was lower (and thus, the cluster purity was higher) in the agglomerative clusterings than either k-means or single-feature. The takeaway of this is that while picking a single attribute, such as UID, does not lead to highly informative groupings, with agglomeration we can begin a meaningful discussion – based on high information clusters – of a storage system based on a single snapshot. With a series of snapshots, we can calculate shifts in the clusterings and track how these correlate to changes in the directory structure and usage.

References

[1] N. Agrawal, W. J. Bolosky, J. R. Douceur, and J. R. Lorch. A five-year study of file-system metadata. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST)*, pages 31–45, Feb. 2007.

[2] T. Gibson, E. L. Miller, and D. D. E. Long. Long-term file activity and inter-reference patterns. In *Proceedings of the 24th International Conference for the Resource Management and Performance and Performance Evaluation of Enterprise Computing Systems (CMG98)*, pages 976–987, Anaheim, CA, Dec. 1998. CMG.

[3] N. Rowe and S. Garfinkel. Finding anomalous and suspicious files from directory metadata on a large corpus. In *3rd Intl. ICST conference on digital forensics and cyber crime*, 2011.

[4] A. Wildani, I. F. Adams, and E. L. Miller. Single-snapshot file system analysis. In *Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS), 2013 IEEE 21st International Symposium on*, pages 338–341. IEEE, 2013.

[5] A. Wildani, E. L. Miller, and O. Rodeh. Hands: A heuristically arranged non-backup in-line deduplication system. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 446–457. IEEE, 2013.