Data Access Pattern Analysis for dCache Storage System

Amanda Pereira[1], Alex Sim[2], John Wu[2], Shinjae Yoo[3], Hiro Ito[3]
Drake University[1], Lawrence Berkeley National Laboratory[2], Brookhaven National Laboratory[3]

ABSTRACT

dCache is a storage management system with disk space for temporary and permanent file storage and connectivity to a larger high performance storage system (HPSS) for additional file storage. Ideally, all files will be stored locally on the dCache disks when they need to be accessed. Analysis of past dataset access patterns for the files stored on dCache is necessary to determine how dataset popularity varies over time and what factors might affect the fluctuation. Findings suggest that datasets are generally accessed only a few days at a time with an inconsistent day to day popularity measurement, and therefore could only be stored locally for the few days that they are popular.

BACKGROUND

- Particle collision events being studied produce a very high volume of data
  - A majority of the data is stored on HPSS tapes, and the rest is stored on dCache disks.

Data is retrieved through a request to the dCache System

- dCache searches local cache for the file

Local copy found in cache: File can be read from or written to with no HPSS interaction.

Local copy not found in cache: a copy of the file has to be moved for access.

Motivation:
Caching data shortens access time and reduces latency. Forecasting dataset popularity for a future time frame allows the dataset to be cached prior to when it needs to be accessed.

RESEARCH QUESTION

How can future dataset popularity be predicted given information on past dataset popularity?

ANALYSIS METHODS

The analysis of data transfer logs and metadata information are used to address this research question by grouping files into datasets and studying the popularity fluctuation.

Datasets are made up of files that all contain information about the same instance

$$popularity = \frac{number\ of\ file\ accesses}{number\ of\ files}$$

- Energy measurement of particles being studied is described using Teraelectronvolts (TeV).
- 83% of data files in the logs are files with information about particle collisions at an energy level of 13TeV

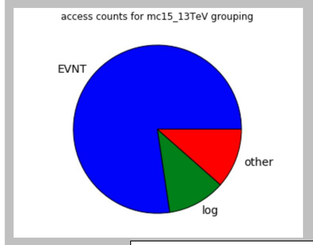

- 69.68 % of the 13TeV files include "mc15_13TeV" or "mc16_13TeV in the file path
- 88.51% of mc15_13TeV files contain "EVNT" or "log" + 8-digit number
- A majority of accesses in mc16_13TeV grouping are for files that contain "DOAD" or "AOD" in their name

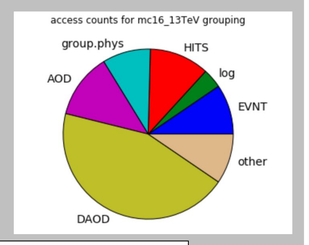

/pnfs/usatlas.bnl.gov/BNLT0D1/rucio/mc15_13TeV/3a/25/EVNT.15927037._062531.pool.root.1

All filepaths with the same 8 digit number are grouped together.

RESULTS

Mc15 group plots:

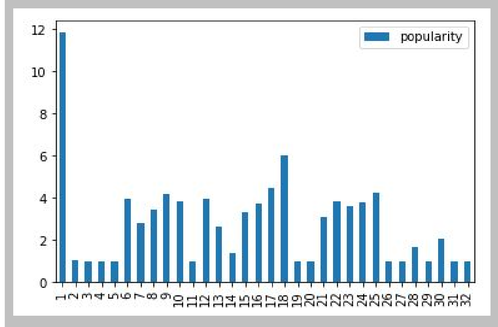

April 3 2021 event popularity

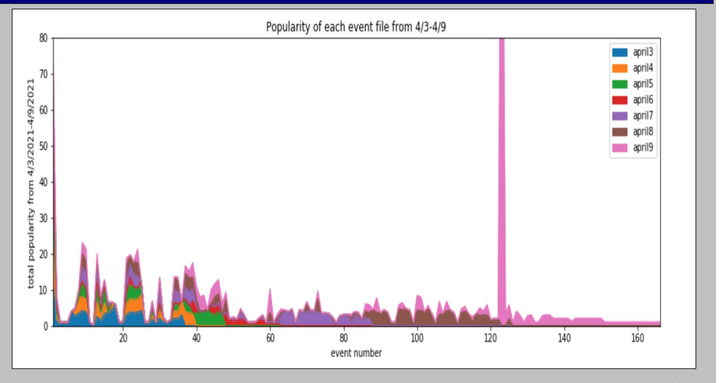

April 3-9 event popularity

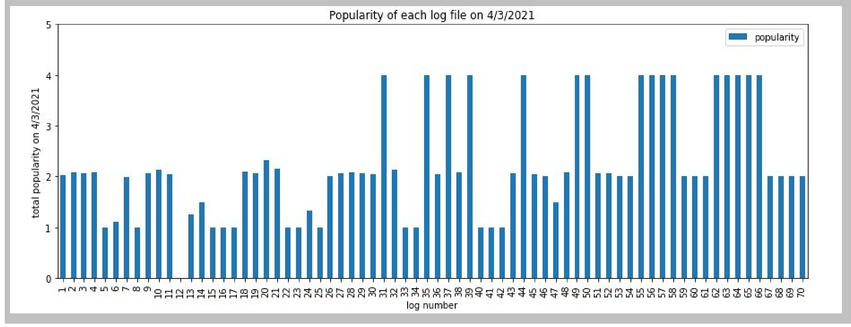

- spike in graph - popular event
- More colors - accessed on more days
  - Greater fluctuation in popularity of event datasets than log datasets.

SUMMARY

- **Search for patterns in popular and unpopular files may help what factors may lead to a high popularity**
- **For maximum storage efficiency, datasets could be stored in the local cache during the time period in which they are popular, and then moved to the HPSS for storage during the time they are not popular.**

FUTURE WORK:

- **Develop a machine learning model to predict the popularity of a given file in the next time frame**
- **Expand study to longer time periods**

FURTHER READING

ACKNOWLEDGMENTS

This work was supported in part by the Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and also used resources of the National Energy Research Scientific Computing Center (NERSC). This work was also supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program.