Understanding Data Access Patterns for dCache System

Julian Bellavita[1], Caitlin Sim[1], John Wu[2], Alex Sim[2], Shinjae Yoo[3], Hiro Ito[3], Vincent Garonne[3], Eric Lancon[3]

[1] Univ. of California, Berkeley
[2] Lawrence berkeley National Lab
[3] Brookhaven National Lab

Introduction

- Sophisticated scientific instruments produce many terabytes of data daily
  - LHC, SPS
- Hierarchical distributed data sharing
- Disk caches (e.g., dCache) work well, but have limited space
- Popularity based caching policy could make disk caches more effective

Introduction (2)

- **dCache is a storage management system for HEP data produced by ATLAS experiment**
  - HEP data is primarily stored on HPSS, but access from HPSS has long latency
  - Disk cache such as dCache could reduce access latency
  - But, disk cache is not able to store all files, so need to focus on popular data
  - Data popularity varies over time, so we need some way to <u>forecast popularity</u>

- **<u>Overarching Goal</u>: Forecast how dataset popularity one day in the future**

- **<u>Main benefit</u>: Improve caching policy to reduce access latency**

Data Description

- Data was extracted from dCache server logs
  - Door, billing, DSN
- Final DataFrame contains data for <u>20 months</u> worth of transactions
  - in two different date ranges: **10/2020 : 04/2021, 01/2022 : 01/2023**
- Files are grouped into <u>datasets</u> based on their <u>task ID</u> (TID)
  - Each dataset has a unique TID

/pnfs/usatlas.bnl.gov/BNLT0D1/rucio/valid1/63/d6/EVNT.08549528._000011.pool.root.1

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

Key:
Storage class
End of project tag
Subdirectories
Event number (included in filename)

Summary Statistics

	10/2020 – 4/2021	1/2022 – 1/2023	TOTAL
TIDs Accessed	576577	677215	1166312
Total Accesses	247598874	302615187	550214061

Table 1. Total numbers of unique TIDs accessed and total number of accesses across all TIDs for both date ranges.

Monthly Access Totals and Means

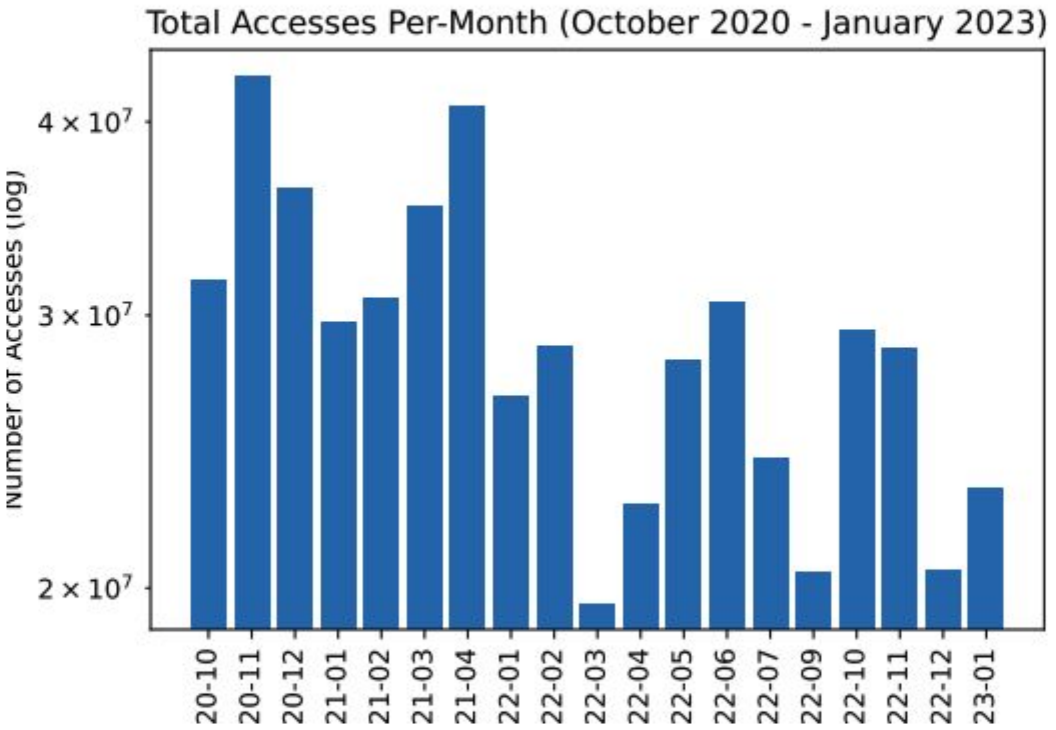

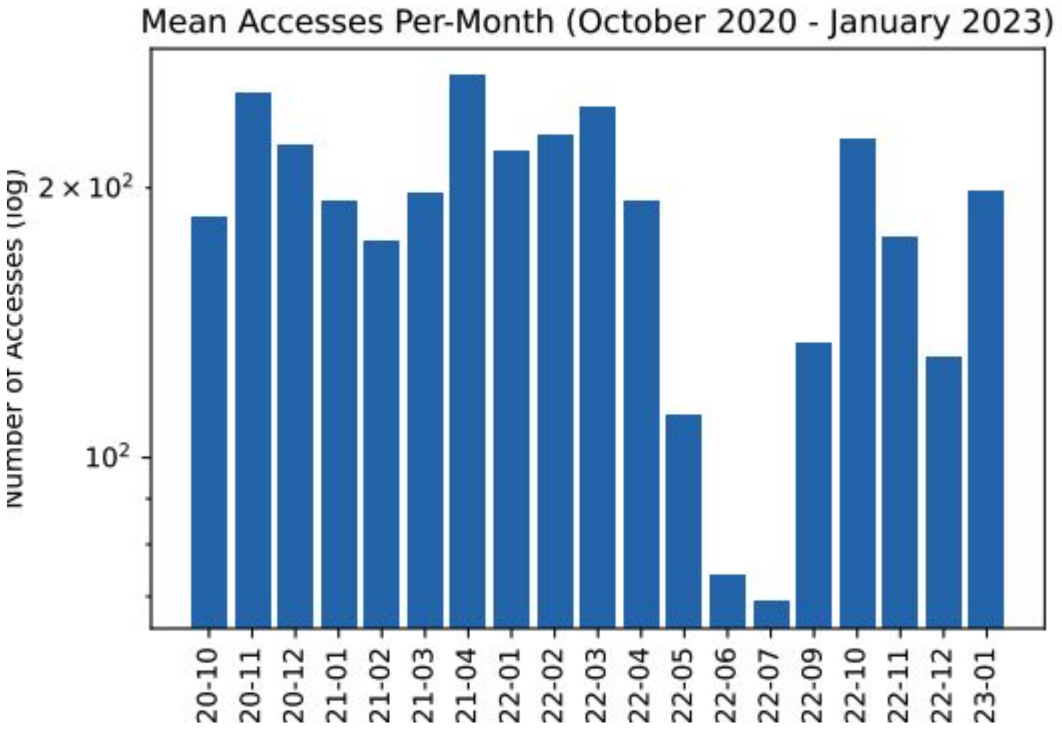

Number of Accesses per Dataset is Highly Skewed

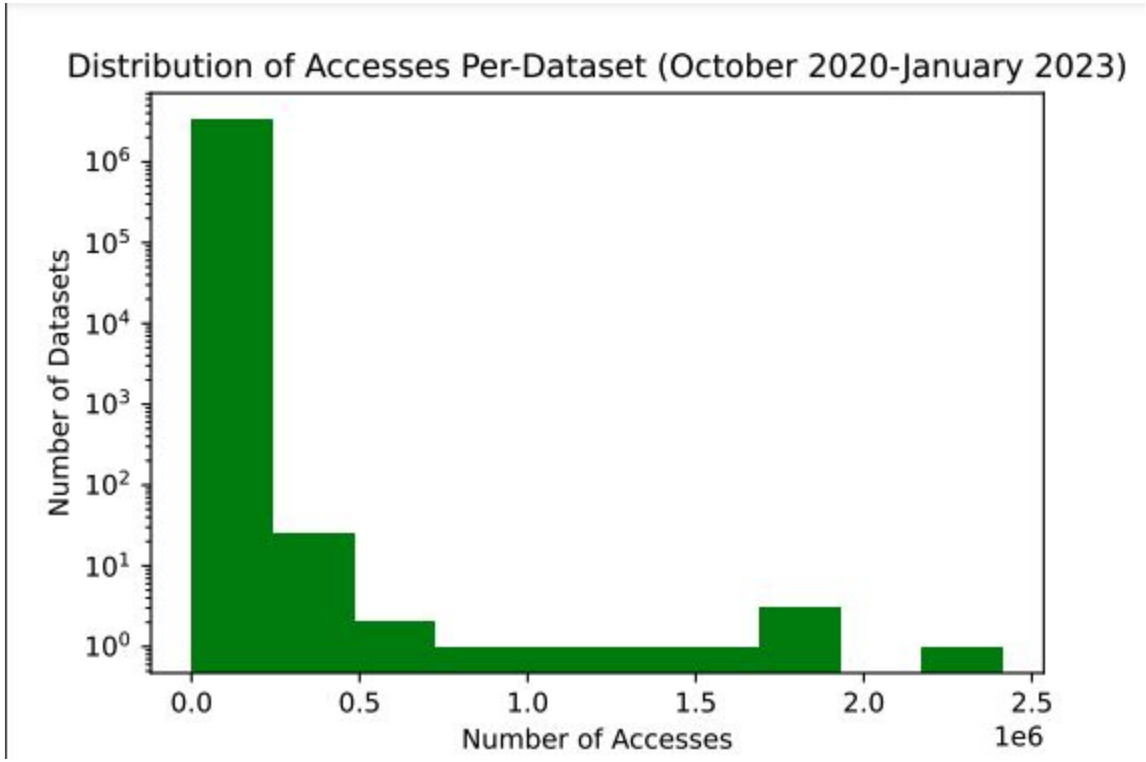

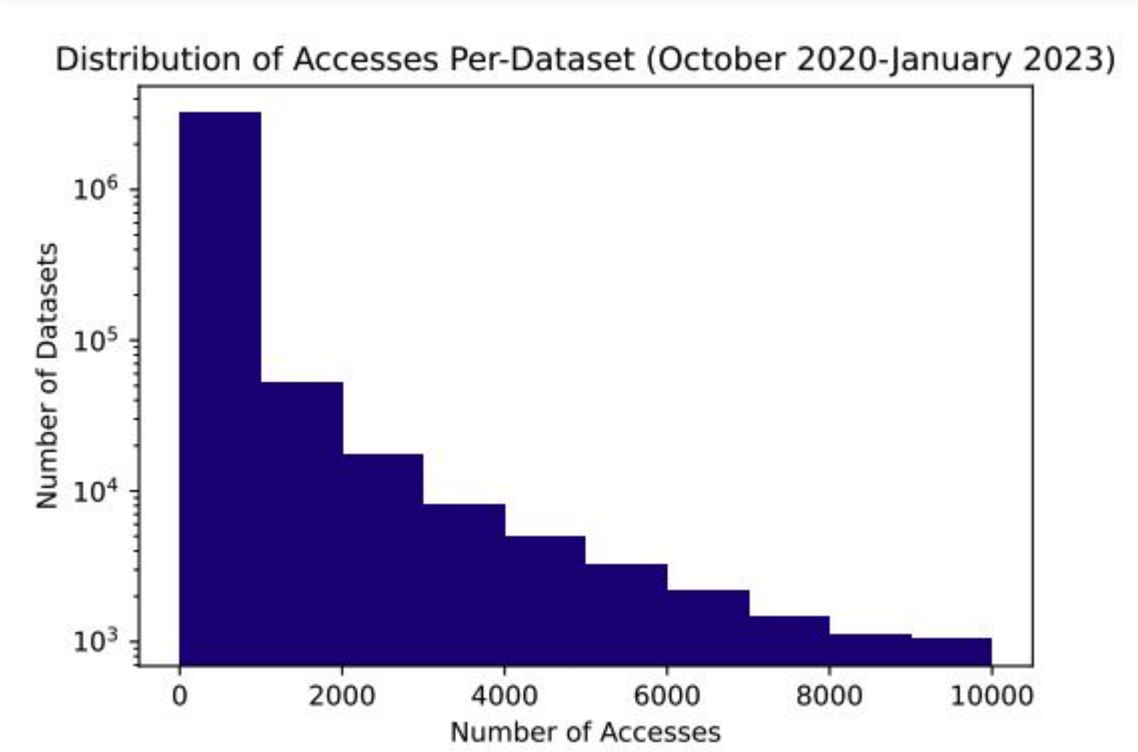

K-Means Clustering

10/2020 – 4/2021 (k=4)

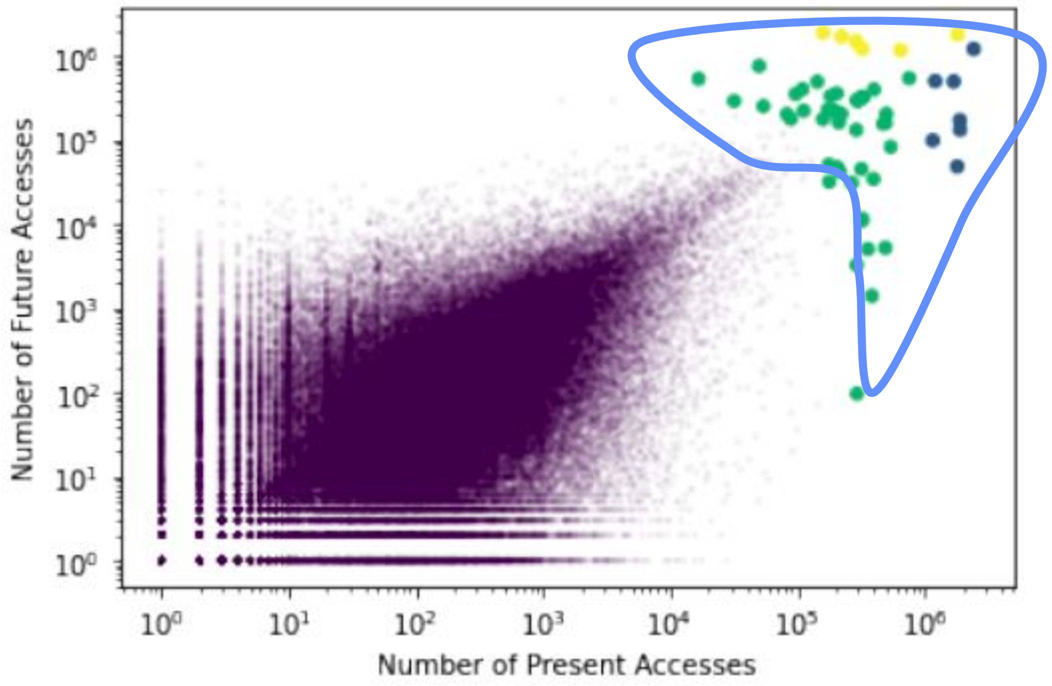

1/2022 – 1/2023 (k=4)

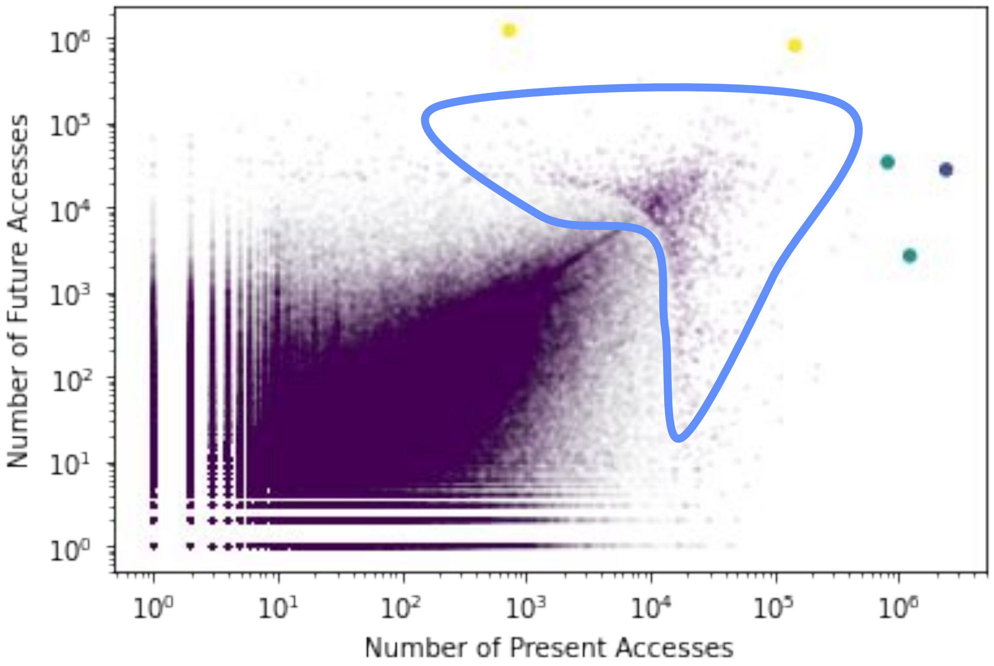

Datasets with more than 10,000 (10^4) accesses per day form a separate cluster from majority of datasets → **popular datasets remains popular**

Forecasting Popularity

- **Model details**
  - Built using PyTorch
  - 2 dense layers, Tanh activation function
  - 70/30 train/test split

- **Trained using data from the first date range (10/2020 : 04/2021), but verified on both date ranges**
  - One model applied to both date ranges

Access Forecasts

10/2020 – 4/2021

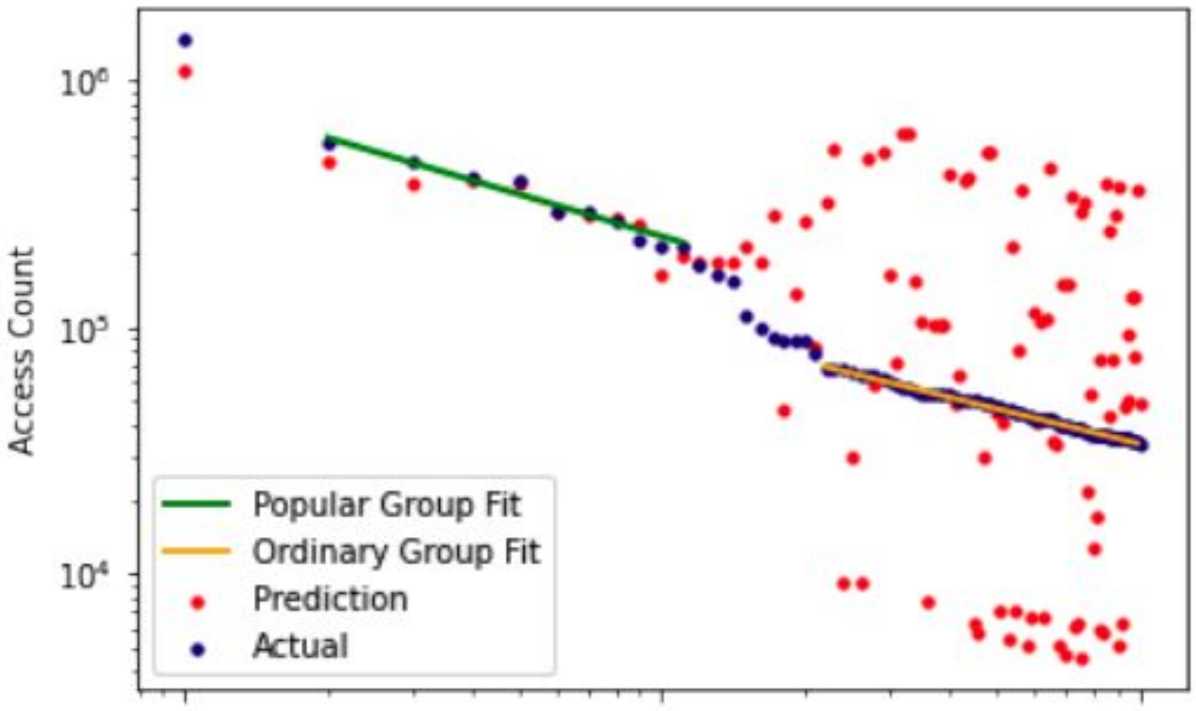

Popular fit exponent := -0.57
Ordinary fit exponent := -0.47

1/2022 – 1/2023

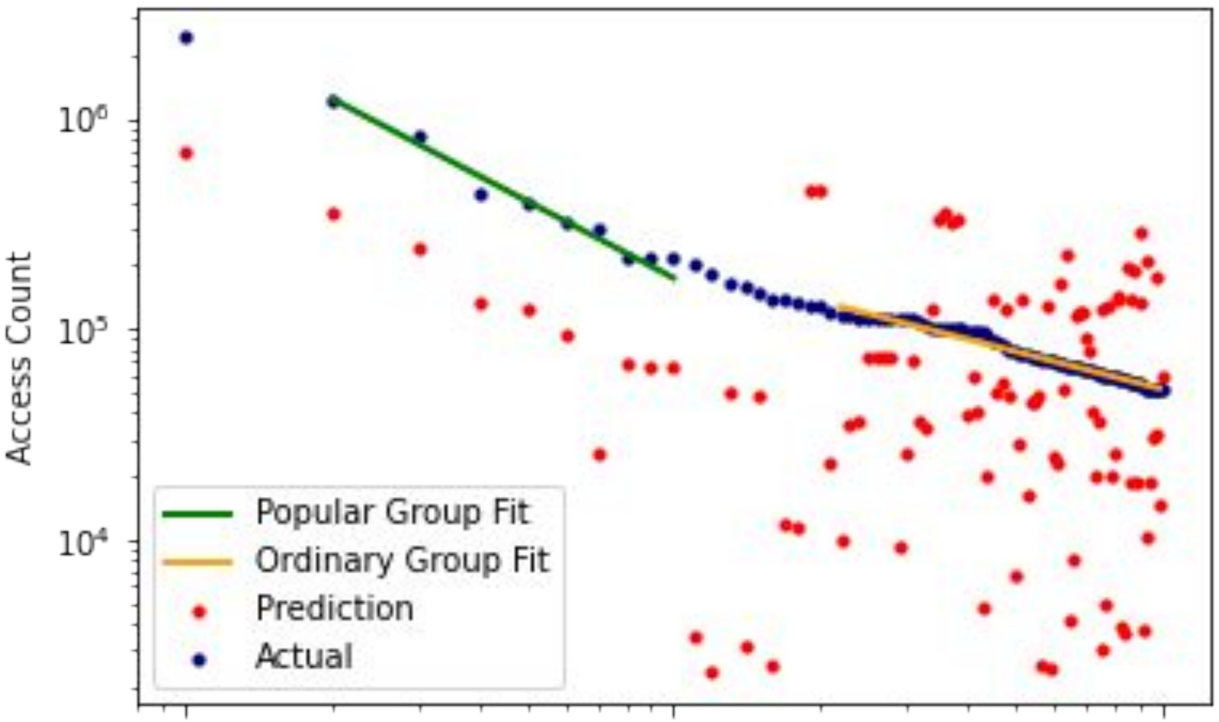

Popular fit exponent := -1.27
Ordinary fit exponent := -0.59

Conclusion

- There exists a small group of highly popular datasets
- These are the datasets we want to pin in dCache
- The neural network <u>cannot</u> predict popularity for <u>unpopular datasets</u>, but it <u>can</u> predict which datasets will be <u>popular</u>
  - Verified across <u>two separate date ranges</u>
- Therefore, we can use the neural network to identify which datasets should be pinned in dCache
- Future work
  - Gradual model update for better accuracy
  - Develop and simulate hypothetical cache policies
  - Try to develop another model for different date ranges