Comparative Study of the Cache Utilization Trends for Regional Scientific Data Caches

Ronak Monga[1], Alex Sim (advisor)[3], K. John Wu (advisor)[3]

[1]Indiana University, [2]Lawrence Berkeley National Laboratory

ABSTRACT

Large scientific collaborations often have many users accessing the same data files, creating repeated file transfers over long distances. Data accesses to the distant data sources cause long latency to the applications and can be further delayed due to limited network bandwidth. XCache-based in-network regional data caching system stores scientific data and can reduce the network traffic and access latency. We examine the established Southern California Petabyte Scale Cache (So Cal Cache) and the newly deployed Chicago Regional Cache (Chicago Cache) for a high-energy physics experiment to analyze cache utilization trends and compare regional data access patterns. The results of the cache utilization trends show that the cache contributed to sharing a majority of data, and regional differences can be explained by the comparative study. Additionally, predictions of cache behavior show low error values in both regions, providing a useful tool for future resource planning.

BACKGROUND

- High-Luminosity Large Hadron Collider expects to have data volume growing by more than 10 times by 2029.
- Study uses data from Southern California Petabyte Scale Cache and from Chicago Regional Cache
- SoCal Cache consists of 23 XCache nodes from Caltech, ESnet, and UCSD
- Chicago Cache consists of 6 XCache nodes from ESnet and Univ. of Wisconsin, Madison
- Regional data caches are expected to handle 30x more data in 2029 than 2023
- Regional scientific data caches
  - Allow data sharing among users in the same region
  - Reduce the repeated data transfers over WAN
  - Decrease data access latency
  - Increase data access throughput
  - Improve overall application performance
- Study period (6 months): Nov 2022 - May 2023

Table 1: Summary of the regional cache storage

	SoCal Cache	Chicago Cache
Number of Nodes	23	6
Storage Capacity	2 PB	345 TB

DATA

Table 2: Summary data access from SoCal Cache (11/22-5/23)

	Number of accesses	Shared data size (TB)	Data transfer size (TB)	Shared data Percentage
Total	3,615,642	5,649.6	415.0	93.1%
Daily Avg	16,975	26.5	1.9	

Table 3: Summary data access from Chicago Cache (11/22-5/23)

	Number of accesses	Shared data size (TB)	Data transfer size (TB)	Shared data Percentage
Total	5,392,415	7,246.5	8,102.9	47.2%
Daily Avg	26,050	35.0	39.7	

RESEARCH QUESTION

How can we use existing regional cache usage data to predict and plan for future regional cache deployments?

CACHE UTILIZATION TRENDS

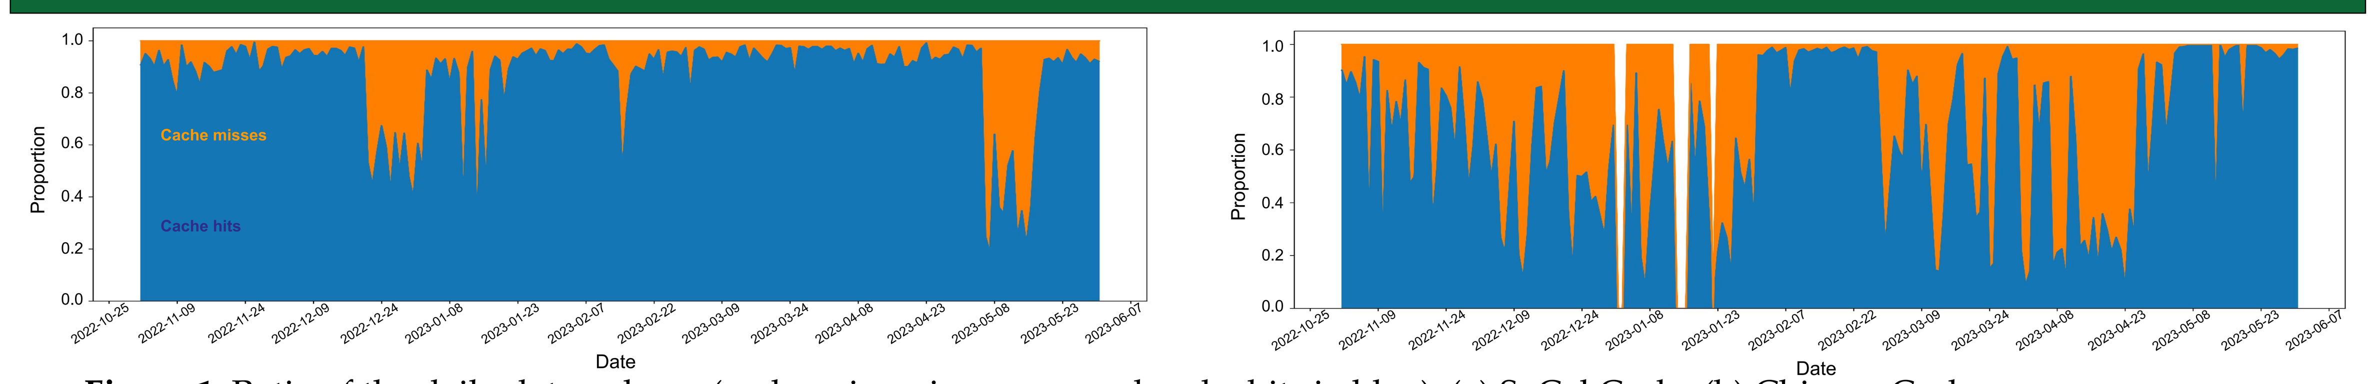

Figure 1: Ratio of the daily data volume (cache misses in orange and cache hits in blue): (a) SoCal Cache (b) Chicago Cache

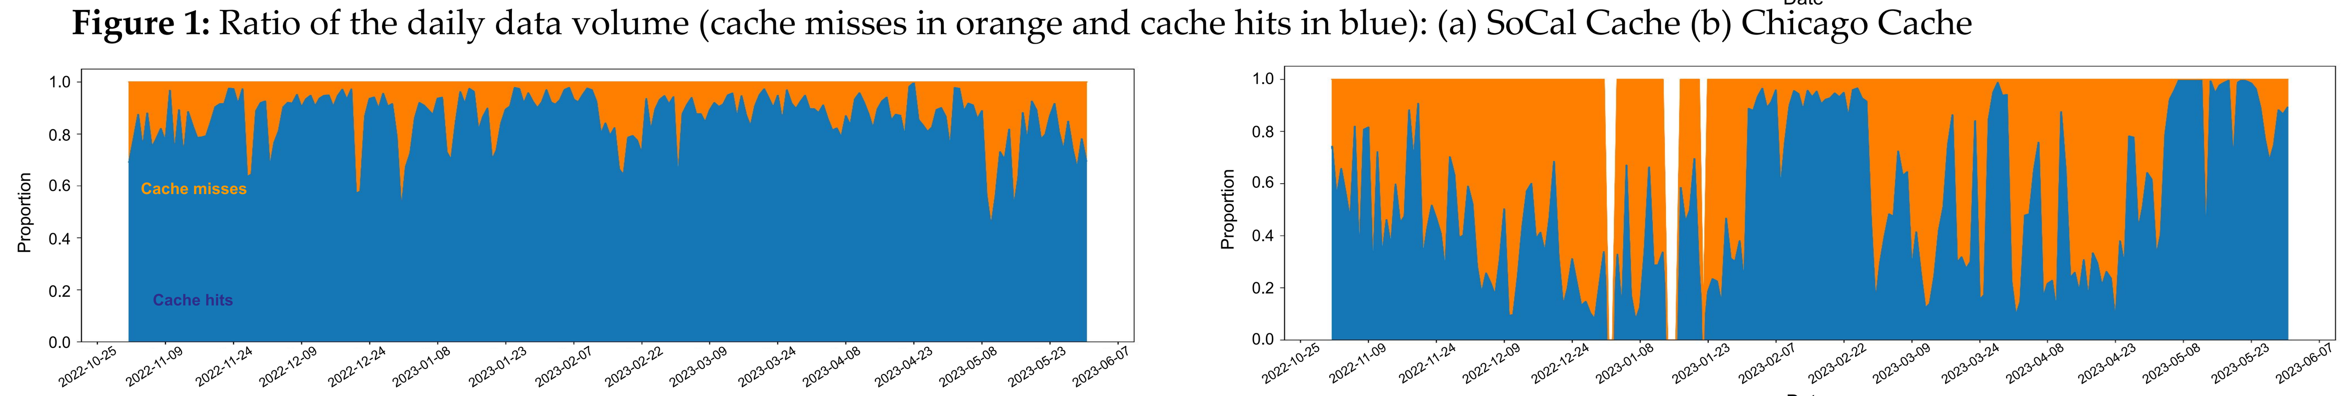

Figure 2: Ratio of the daily number of requests (cache misses in orange and cache hits in blue): (a) SoCal Cache, (b) Chicago Cache

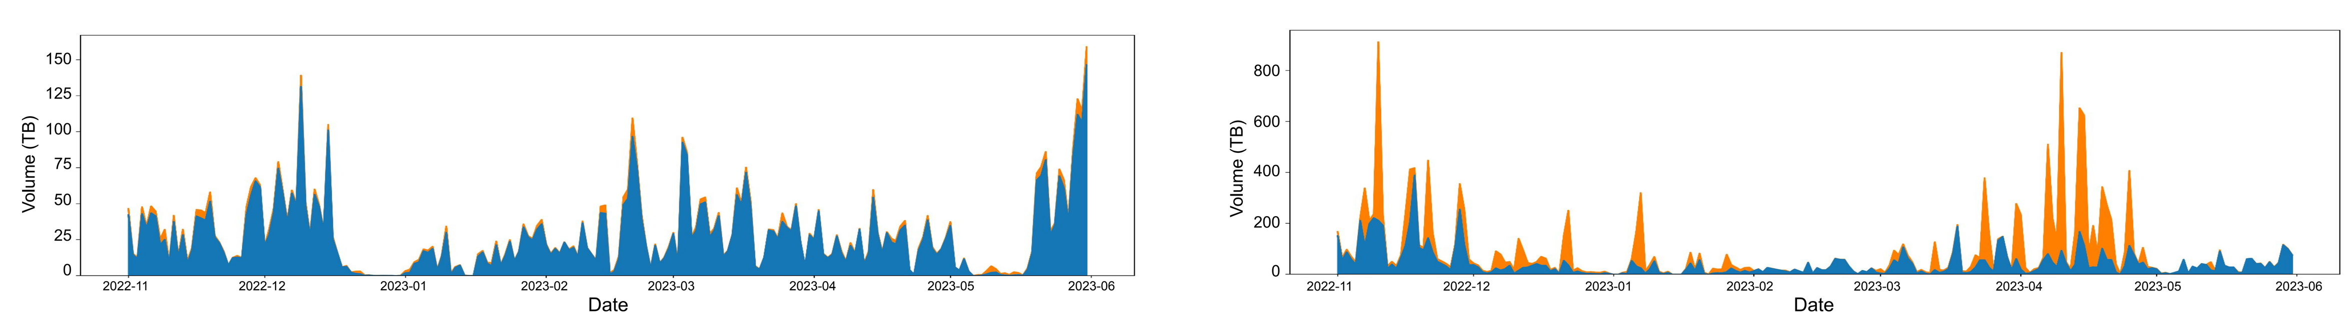

Figure 3: Daily volume of cache hits (in blue) and cache misses (in orange): (a) SoCal Cache (b) Chicago Cache

Figure 4: 12-day moving sum of total access sizes for the Chicago regional cache. The Chicago Cache's maximum capacity (345 TB) is only able to accommodate about 50% of requested bytes. If the cache stores files for 12 days, then initial investigation suggests that increasing capacity to 700 TB would allow the cache to serve upwards of 80% of requested bytes. This might further improve the Chicago Cache's performance and make its impact comparable to that of the SoCal Cache. Note that 12-day is the average lifespan for a file in cache based on the previous study of SoCal repository.

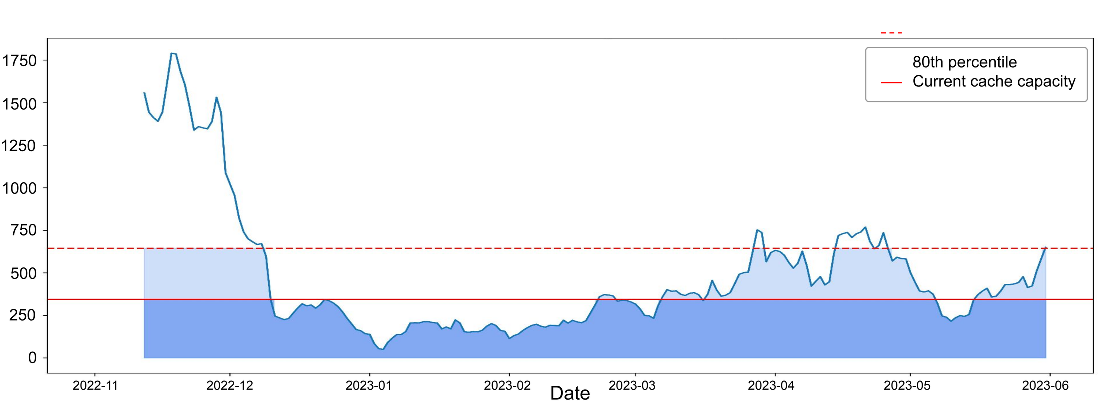

PREDICTIONS

LSTMs: 128 units, dropout rate 0.04, tanh activation function

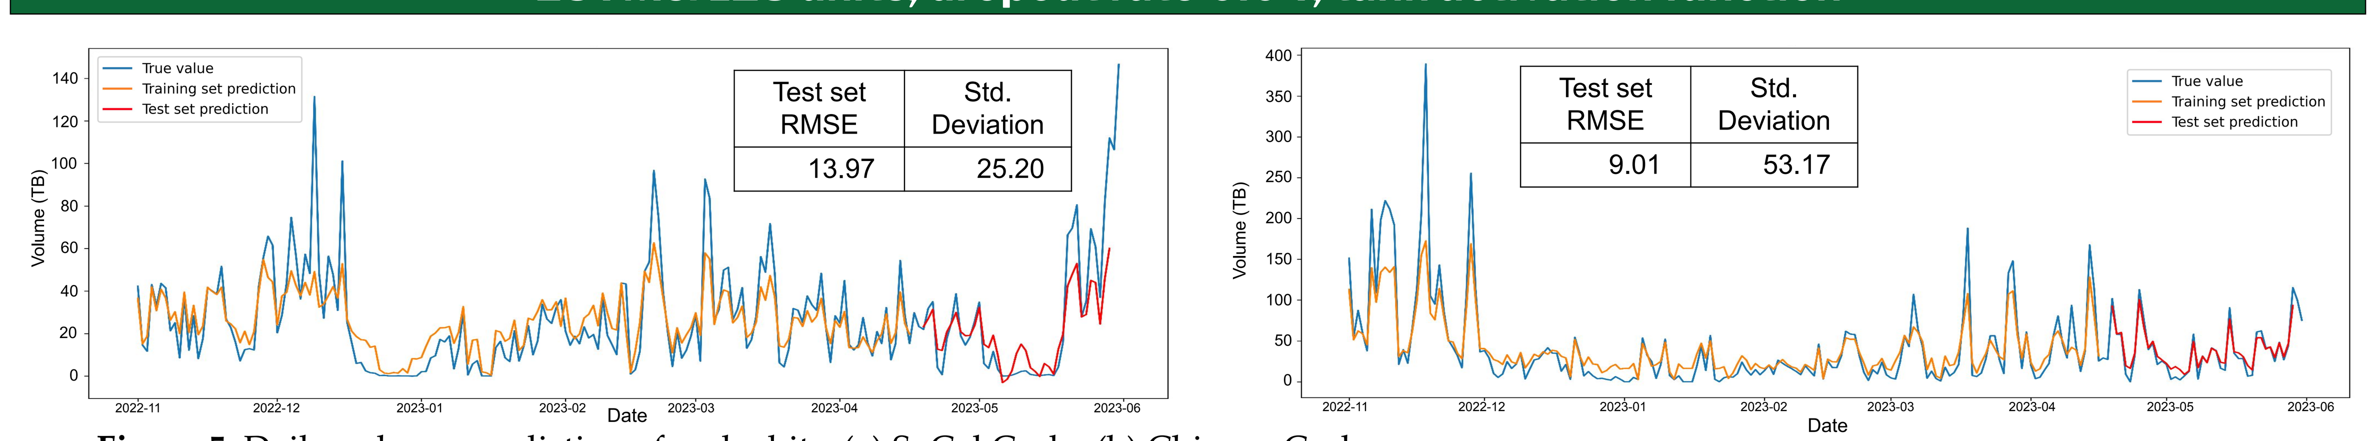

Figure 5: Daily volume prediction of cache hits: (a) SoCal Cache (b) Chicago Cache

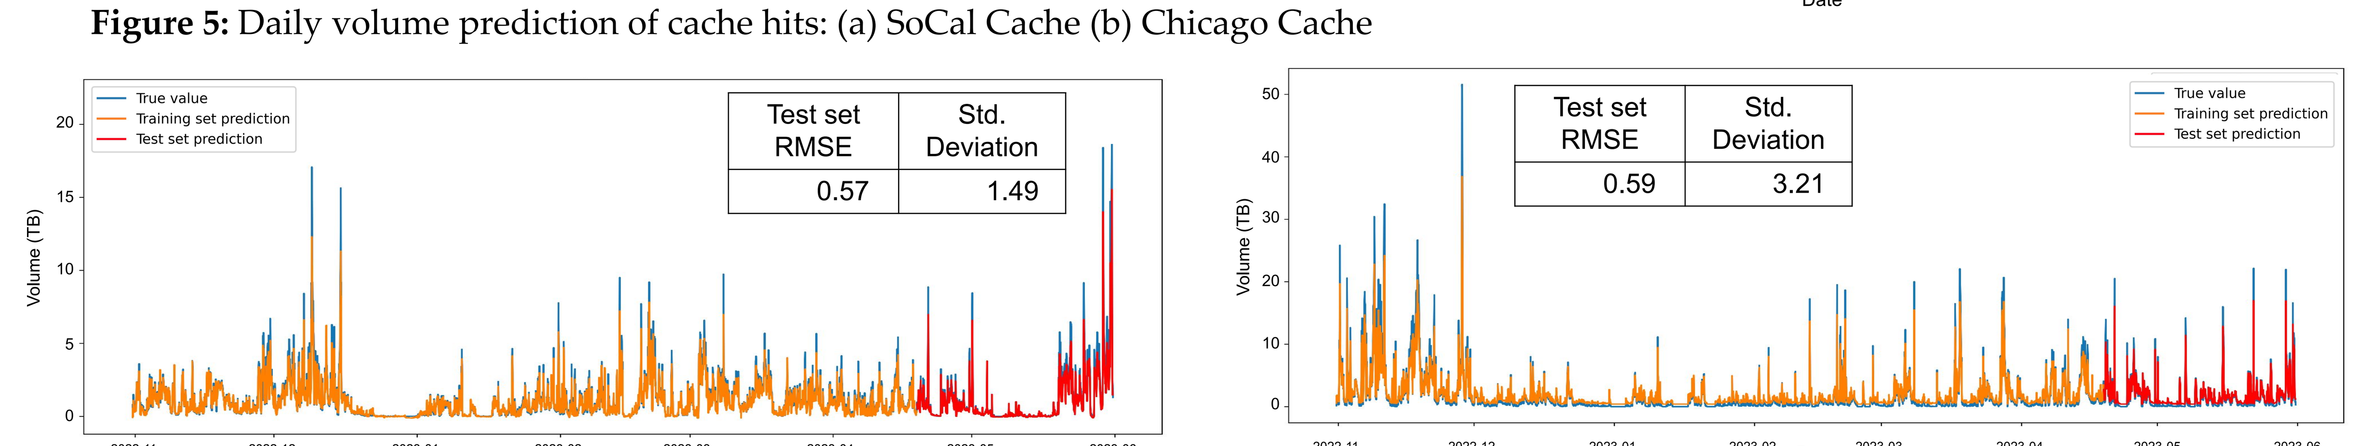

Figure 6: Hourly volume prediction of cache hits: (a) SoCal Cache, (b) Chicago Cache

DISCUSSION

- Chicago Cache experienced higher amounts of peak traffic than SoCal Cache (Fig. 3) while having a smaller cache capacity (Table 1).
- SoCal Cache is better able to serve data as cache hits during peak traffic than the Chicago Cache (Fig. 3)
- Chicago Cache was deployed in Oct 2022, potentially explaining the higher percentage of cache misses in early months. (Fig. 1b, 2b)
- Chicago Cache experienced another period of high cache miss percent in Mar-Apr 2023 that requires investigation (Fig. 1b, 2b).
- Chicago Cache's total access size 12-day moving sum suggests that increasing its capacity could greatly improve its cache performance (Fig. 4).
- LSTM models with 128 units, dropout rate 0.04, and tanh activation functions produce low-error predictions for both regions (Fig. 5, 6).
- LSTM models' testing RMSE are consistently considerably lower than one standard deviation (Fig. 5, 6).
- LSTM models trained on hourly data tend to have less error than models for daily data, likely due to the fact that there are more data points (Fig. 5,.6).

CONCLUSIONS

- Regional data caches supplement existing local storages and benefit the wider user community, allow data sharing, decrease data access latency, and improve overall application performance.
- Chicago regional cache could see improved performance by increasing its capacity.
- Further studies:
- Cache utilization trends are predictable by LSTM models.
- Predictions become more accurate with more granular data.
  - Can a model trained on one region create accurate predictions about behavior in another region?
  - Use cache-modeling to determine the optimal capacity for the Chicago regional cache.

ACKNOWLEDGMENTS

Thanks to my mentors and collaborators, Alex Sim, John Wu, Chin Cuok, Damian Hazen, Inder Monga, Diego Davila, Frank Würthwein, Justas Balcas and Harvey Newman. This work was supported in part by the Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Dept. of Energy under Contract No. DE-AC02-05CH11231, and also used resources of the National Energy Research Scientific Computing Center (NERSC).