Data Throughput Performance Trends of Regional Scientific Data Cache

Caitlin Sim
University of California, Berkeley
caitlinsim@berkeley.edu

I. Introduction

There has been a significant increase in data volume from various large scientific projects. These projects, including the Large Hadron Collider (LHC) experiment, are typically distributed among a wide geographic area [1]. The High Energy Physics (HEP) community requires increased data volume on the network, as the community expects to produce almost thirty times annual data volume between 2018 and 2028 [2]. The same dataset can be transferred multiple times in the same region to the same institution, meaning that much of the data transfer is repetitive on the network and therefore inefficient for the client jobs. This issue can be mitigated by regional data caching mechanism [3], or in-network cache [4], reducing network traffic for regional users [5]. Overall application performance can also be improved, with the number of redundant data transfers over the wide-area network decreasing.

In this study, we examined the trends in data volume and data throughput rates from the Southern California Petabyte Scale Cache (SoCal Repo) [6], where the network traffic is primarily carried out by the Energy Sciences Network (ESnet). From the trends, we also determined how much a machine learning model can predict the network access patterns for the regional data cache. The fluctuation in the daily cache utilization, as shown in Figure 1, is high, and it is challenging to build a learning model to follow the trends. The study results can be used to determine the network utilization for the data transfers and predict the transfer throughput performance with the Long Short-Term Memory (LSTM) models. Cache utilization, network resources, and application workflow performance will be optimized with these results.

II. Methods

SoCal Repo has approximately 2.5PB of total storage with 24 federated caching nodes for LHC CMS experiment. The measurement data has been collected from July 2021 and June 2022 with 8.02 million data access records for 8.2PB of traffic volume for cache misses and 4.5PB of traffic volume for cache hits. The NERSC Cori and Perlmutter supercomputers are used to determine the data access trends on 1-year of measurement data. The information from the cache misses, or data transfers from remote sources to local cache, and the cache hits- data from the local cache to client jobs- is analyzed for access trends. These patterns are further analyzed with moving averages in days and hours, and explored with LSTM models [7] with TensorFlow.

III. Results

Figure 1 shows the daily proportion of cache hits volume to cache misses volume. Observations show a drop in the proportion between Oct. 2021 and Jan. 2022, which was attributed to the unusual streaming of data for the analysis jobs. Average network traffic volume reduction rate, as shown in Figure 2, is 1.55 for the whole study period, and under the normal use such as between July 2021 and Sep. 2021, the average network traffic volume reduction rate is 2.35.

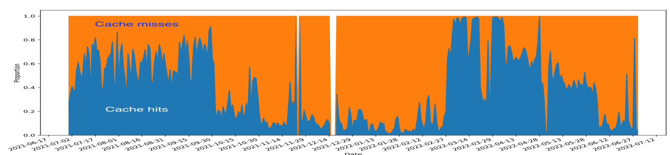

Fig. 1. Daily proportion of cache hits volume and cache misses volume

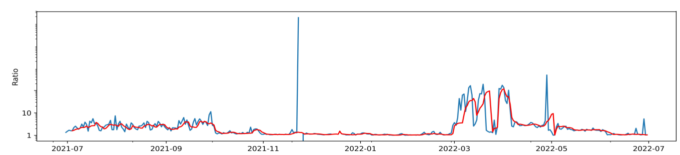

Fig. 2. Daily network traffic volume reduction rates

We studied the hourly and daily average data throughput performance for cache hits and misses. The data throughput was defined as the data transfer size over the transfer time that would determine network performance and anomalous behaviors.

The daily LSTM model results are shown in Figure 3 for the average throughput performance and 4 the 7-day moving average performance trend. In each figure, (a) represents cache hits and (b) represents cache misses. Although the model deviates slightly in the extremes, the RMSE values are fairly low, indicating that the predicted throughput values were close to the actual values.

The results from the hourly LSTM models in Figures 5 and 6 show that the RMSE values are generally lower than ones from the daily model. The predicted hourly values are overall closer to the actual values than the daily model. This could be attributed to the fact that there are more data points for training in the hourly model, leading to better predictions. Results in Figures 4 and 6 show better prediction than Figures

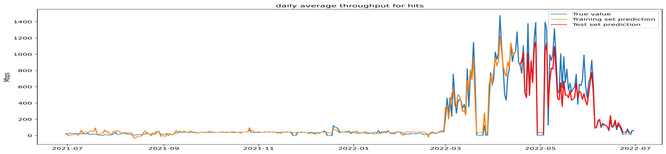

(a) Cache hits; training set RMSE=52.29, testing set RMSE=149.06

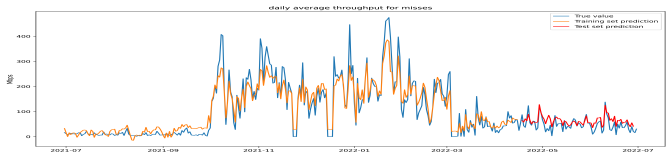

(b) Cache misses; training set RMSE=34.57, testing set RMSE=18.27

Fig. 3. Daily average throughput performance per access.

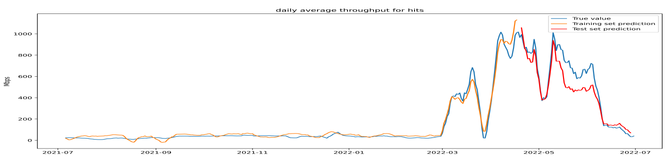

(a) Cache hits; training set RMSE=34.94, testing set RMSE=106.86

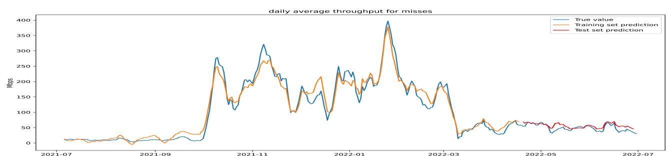

(b) Cache misses; training set RMSE=19.72, testing set RMSE=6.93

Fig. 4. Daily average throughput performance with 7-day moving average.

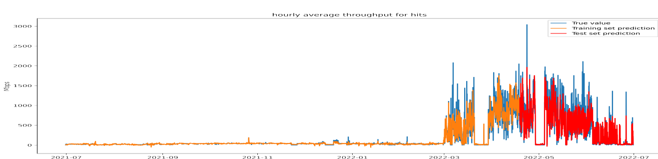

(a) Cache hits; training set RMSE=62.34, testing set RMSE=136.85

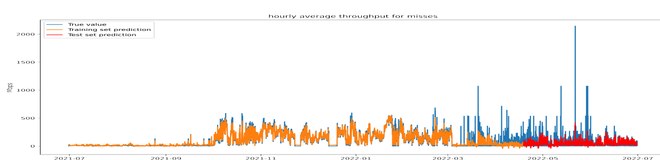

(b) Cache misses; training set RMSE=33.26, testing set RMSE=78.87

Fig. 5. Hourly average throughput performance per access.

3 and 5 respectively, as the moving average reduces extreme values resulting in smaller RMSE values for both cache hits and cache misses.

IV. Conclusion

In this study, we analyzed the data access trends for data volume and the average data throughput rates of the SoCal Repo. Network traffic volume was reduced by an average factor of 2.35 in normal use, indicating that the data caching system does aid in reducing network traffic. Cache utilization and the average data throughput performances are predictable

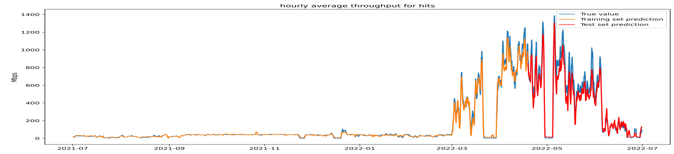

(a) Cache hits; training set RMSE=18.86, testing set RMSE=81.68

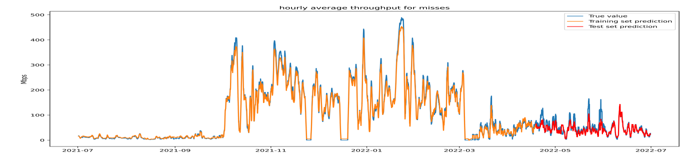

(b) Cache misses; training set RMSE=9.76, testing set RMSE=14.10

Fig. 6. Hourly average throughput performance with 24-hour moving average.

by LSTM models, with the RMSE values being fairly small. The prediction results can be better with more data records and less extreme values.

Further studies are considered by exploring characteristics of other data repositories, as well as examining longer term network requirements for the data repositories.

Acknowledgment

This work was supported 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.

References

[1] X. Espinal, S. Jezequel, M. Schulz, A. Sciabà, I. Vukotic, and F. Wuerthwein, "The quest to solve the hl-lhc data access puzzle," *EPJ Web of Conferences*, vol. 245, p. 04027, 2020. [Online]. Available: https://doi.org/10.1051/epjconf/202024504027

[2] B. Brown, E. Dart, G. Rai, L. Rotman, and J. Zurawski, "Nuclear physics network requirements review report," Energy Sciences Network, University of California, Publication Management System Report LBNL-2001281, 2020. [Online]. Available: https://www.es.net/assets/Uploads/20200505-NP.pdf

[3] E. Fajardo, D. Weitzel, M. Rynge, M. Zvada, J. Hicks, M. Selmeci, B. Lin, P. Paschos, B. Bockelman, A. Hanushevsky, F. Würthwein, and I. Sfiligoi, "Creating a content delivery network for general science on the internet backbone using XCaches," *EPJ Web of Conferences*, vol. 245, p. 04041, 2020. [Online]. Available: https://doi.org/10.1051/epjconf/202024504041

[4] E. Copps, H. Zhang, A. Sim, K. Wu, I. Monga, C. Guok, F. Wurthwein, D. Davila, and E. Fajardo, "Analyzing scientific data sharing patterns with in-network data caching," in *4th ACM International Workshop on System and Network Telemetry and Analysis (SNTA 2021)*, ACM. ACM, 2021.

[5] R. Han, A. Sim, K. Wu, I. Monga, C. Guok, F. Wurthwein, D. Davila, J. Balcas, and H. Newman, "Access trends of in-network cache for scientific data," in *5th ACM International Workshop on System and Network Telemetry and Analysis (SNTA 2022)*, ACM. ACM, 2022.

[6] E. Fajardo, A. Tadel, M. Tadel, B. Steer, T. Martin, and F. Würthwein, "A federated xrootd cache," *Journal of Physics: Conference Series*, vol. 1085, p. 032025, 2018.

[7] A. Sherstinsky, "Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network," *Physica D: Nonlinear Phenomena*, vol. 404, p. 132306, 2020.