Predictive Modeling of Effective Peak Throughput for Regional In-Network Caching

Sunmin (Annie) Eom[1], Alex Sim (advisor)[2], K. John Wu (advisor)[2]

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

U.S. DEPARTMENT OF ENERGY — Office of Science

ABSTRACT

Large-scale scientific experiments, particularly in high-energy physics, generate massive volumes of data that are repeatedly accessed by researchers worldwide. To reduce redundant long-distance transfers, institutions employ regional in-network cache nodes that enable efficient data sharing within the same geographic region. This work characterizes three years of access records from the Southern California Petabyte-Scale (SoCal) Cache, computing daily peak throughputs for cache misses and hits across three metrics: the bursty maximum and the more stable 99th and 95th percentiles. Long Short-Term Memory (LSTM) networks are then trained to forecast peak throughput. Though accuracy degrades at longer horizons, 1-day-ahead forecasts capture the general trend well and can inform future bandwidth allocation and network efficiency improvements.

BACKGROUND

- Regional data caches
  - Reduce redundant wide-area network (WAN) transfers
  - Enable efficient data sharing within the same region
  - Lower data access latency
  - Improve application performance
- Study uses data from Southern California Petabyte Scale (SoCal) Cache, which
  - consists of 28 XCache nodes from Caltech, ESnet, and UCSD, and
  - supports high-energy physics experiments (e.g., Large Hadron Collider, LHC)
- High-Luminosity LHC expects 10x increase in network use by 2029, making regional caching increasingly critical

DATA

- Study period: July 2022 - September 2025 (1,188 days)
- Preprocessing removes 16.8% of 15.6 M raw access records with:
  - Operation time ≤ 0 or > 2 days and/or
  - Data size < 1 MB and operation time > 1 hour
- Of the remaining 13.0 M accesses (9.7 PB), 8% are misses and 92% hits, both by count and volume
  - Each access is classified as a miss (fetched remotely, reflecting WAN traffic) or a hit (served from cache, reflecting local traffic)

Table 1: Summary of data access counts from SoCal Cache

	Total Accesses	Cache Misses	Cache Hits
Total	12,953,244	1,016,241 (7.8%)	11,937,003 (92.2%)
Daily Avg	10,903	855 (7.8%)	10,048 (92.2%)

Table 2: Summary of data access sizes (TB) from SoCal Cache

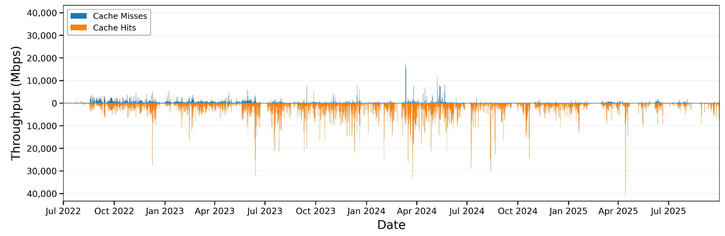
	Total Accesses	Cache Misses	Cache Hits
Total	9,731.0	799.2 (8.2%)	8,931.8 (91.8%)
Daily Avg	8.2	0.7 (8.2%)	7.5 (91.8%)

RESEARCH QUESTION

How do cache hit and miss throughputs on the SoCal Cache behave over a three-year period, and how accurately can LSTMs forecast daily peak throughput metrics to guide bandwidth allocation?

PEAK THROUGHPUT ANALYSIS (MAX, P99, P95)

Figure 1: Per-minute throughput of cache misses (blue) and hits (orange) in Mbps. We compute per-minute throughput separately for misses and hits by proportionally allocating each transfer's data size across the minutes it spans, then summing and dividing by 60 seconds.



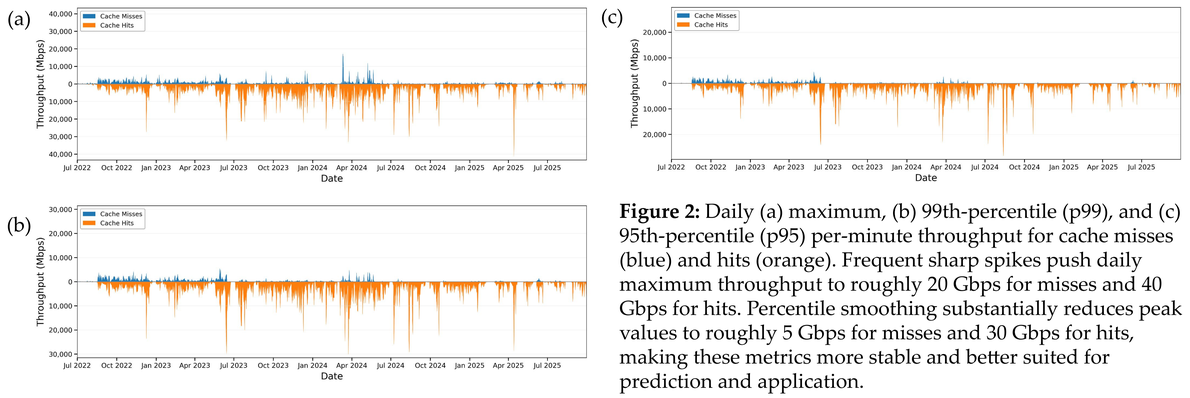


Figure 2: Daily (a) maximum, (b) 99th-percentile (p99), and (c) 95th-percentile (p95) per-minute throughput for cache misses (blue) and hits (orange). Frequent sharp spikes push daily maximum throughput to roughly 20 Gbps for misses and 40 Gbps for hits. Percentile smoothing substantially reduces peak values to roughly 5 Gbps for misses and 30 Gbps for hits, making these metrics more stable and better suited for prediction and application.

FORECASTING

- Goal: Predict daily max, p99, and p95 cache miss and hit throughputs to inform bandwidth allocation.
- Challenges: High variance and user-driven unpredictability make throughput forecasting difficult.
- Approach: Use Long Short-Term Memory (LSTM) networks, well-suited to learning long-term dependencies in sequential data, for prediction.
- Architecture: Models are trained on the first 80% of the data and tested on the rest. Longer prediction horizons favor two-layer architectures, higher learning rates, and longer input windows, though no single architecture is clearly best. The best 1-day-ahead and 7-day-ahead models use input windows ranging from 7 to 30 days.

Table 3: Test-set RMSE and RE for the best 1- and 7-day-ahead models, evaluated on day-1 and day-7 predictions. The standard deviation over the full study period, which is used to compute RE, is also reported. Each value is the average of ten runs; the best model is the one with lowest overall RE (averaged across miss and hit).

Metric	Target	1-Day Model, Day 1 Pred.	7-Day Model, Day 1 Pred.	7-Day Model, Day 7 Pred.	Std. Dev.
Max	Miss	297.1 (0.246)	312.3 (0.259)	368.1 (0.305)	1207.4
	Hit	3192.0 (0.705)	3383.8 (0.747)	4004.0 (0.884)	4527.3
p99	Miss	146.5 (0.227)	155.0 (0.241)	173.6 (0.270)	643.9
	Hit	2530.6 (0.672)	2554.7 (0.678)	3070.6 (0.815)	3768.0
p95	Miss	89.7 (0.200)	103.6 (0.231)	109.7 (0.245)	448.7
	Hit	1795.9 (0.615)	1733.5 (0.594)	2070.0 (0.709)	2917.8

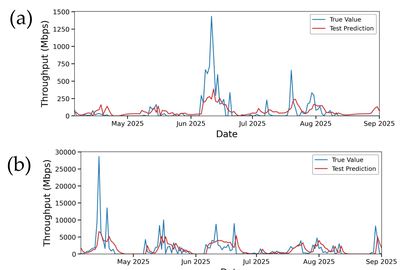


Figure 3: Daily p99 throughput day 1 test-set prediction of (a) cache misses, (b) cache hits, from the 1-day-ahead model

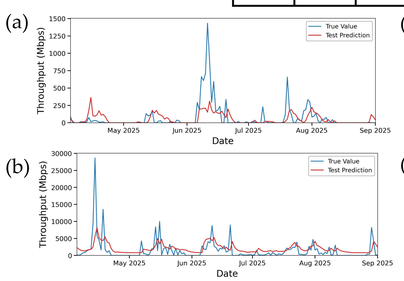


Figure 4: Daily p99 throughput day 1 test-set prediction of (a) cache misses, (b) cache hits, from the 7-day-ahead model

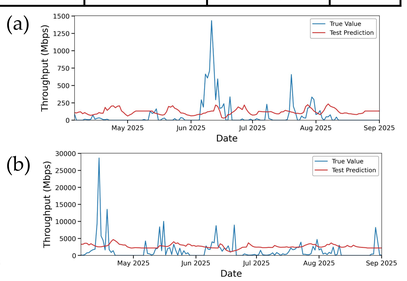


Figure 5: Daily p99 throughput day 7 test-set prediction of (a) cache misses, (b) cache hits, from the 7-day-ahead model

DISCUSSION

- For 1-day-ahead predictions, cache misses achieve low REs of 0.20-0.25 and cache hits 0.62-0.71 (Table 3); predictions follow the general trend despite missing the extreme peaks (Fig. 3).
- RE rises by 6-20% from 1st- to 7th-day predictions of the best 7-day-ahead model (Table 3), with predictions tracking less of the day-to-day variation (Fig. 4, 5).
- At 14 and 30 days, RE reaches 0.28-0.31 for misses and 0.73-0.99 for hits.
- For all horizons, the best model's RE decreases from maximum to 99th to 95th percentile (Table 3), confirming the smoother metrics are more predictable.
- Forecasts beat the persistence (previous day's value) and weekly (previous week's values, repeated) baselines at nearly every horizon; against the harder input-window average, models win at 1 and 7 days ahead for all metrics, but at longer horizons only for some misses.
- At the same prediction distance, models trained for 1-, 7-, 14-, and 30-day horizons perform comparably, so training for a longer horizon doesn't itself degrade performance significantly at a given prediction distance (Table 3; Fig. 3, 4).

CONCLUSION

- Cache hits dominate traffic (92.2% of accesses, 91.8% of data), confirming SoCal Cache absorbs most traffic locally rather than over the WAN.
- Per-minute throughput is bursty, peaking at ~20 Gbps (misses) and ~40 Gbps (hits); p99/p95 smoothing drops these to ~5 Gbps and ~30 Gbps.
- 1-day-ahead forecasts are robust, following the general shape of demand at a low relative error.
- At longer horizons, predictions become progressively less reliable, and from 14 days onward baselines outperform the models.
- Forecasts of each peak metric offer distinct guidance: the predicted maximum covers the day's peak but leaves most capacity unused, while 99th or 95th percentile covers nearly all minutes at far less bandwidth.

FUTURE WORK

- Enhanced models for better long-horizon performance
- Hourly models for max/p99/p95 throughputs for more granular data points

ACKNOWLEDGMENTS

This work was supported by the Office of Advanced Scientific Computing Research and Office of Workforce Development for Teachers and Scientists, Office of Science, of the U.S. Dept. of Energy under Contract No. DE-AC02-05CH11231.