

Comparing Cache Replacement Policies with Paired Synthetic Replay

SciPepper • Controlled synthetic demonstration • Seed 42 / 2,000 requests

Research question

How do LRU and FIFO behave under the same capacity constraint? Compare replacement rules using a shared request trace, identical capacities and an empty starting state.


The aim is a transparent, reproducible comparison. All values below come from synthetic replay; they are not measurements from a production system.

Source p. 1

Method: paired trace replay

Generate 2,000 requests over 128 keys with seed 42 and weights proportional to $1 / (\text{key} + 1)$. Replay the identical sequence at capacities of 16, 32 and 64 entries.

LRU updates recency after a hit; FIFO preserves insertion order. Both evict one entry before inserting into a full cache. Record hits; do not measure execution time.



Controlled synthetic demonstration: no latency or generalization claim

Source p. 2

Controlled comparison

A hit means the requested key is present before the current access. Rate equals the recorded hit count divided by 2,000; cold-start accesses remain in the denominator.


Hold request order, key universe, capacity and initial state constant. The replacement rule is the only changed condition. One generated workload is evaluated; repeated seeds and confidence intervals are outside this demonstration.

Source p. 3

Results on the synthetic trace

LRU: 913, 1248, 1606 hits at capacities 16, 32, 64. FIFO: 811, 1144, 1524 hits under the same conditions. The plot uses these exact counts and a fixed denominator of 2,000.

The rates describe one synthetic trace. No latency, throughput or statistical significance measurement was performed.



Source p. 4

Interpretation & limitations

Pairing the workload and capacity makes the comparison inspectable. Preserve conditions beside the result figure and avoid converting a source-specific observation into a general scientific claim.

One synthetic seed does not cover real access patterns, concurrency or retrieval cost. A stronger study would need multiple workloads, repeated trials and uncertainty analysis. This example verifies layout and provenance only.

Source p. 5