An Analysis of Facebook Photo Caching


Qi Huang

Ken Birman

Robbert van Renesse

Wyatt Lloyd

Sanjeev Kumar

Harry C. Li

Instrumented Stack

Deep and Distributed

1. 4 layers of cache and storage.
2. ~12M user IPs, ~20 Point-of-Presence, 4 Datacenters.

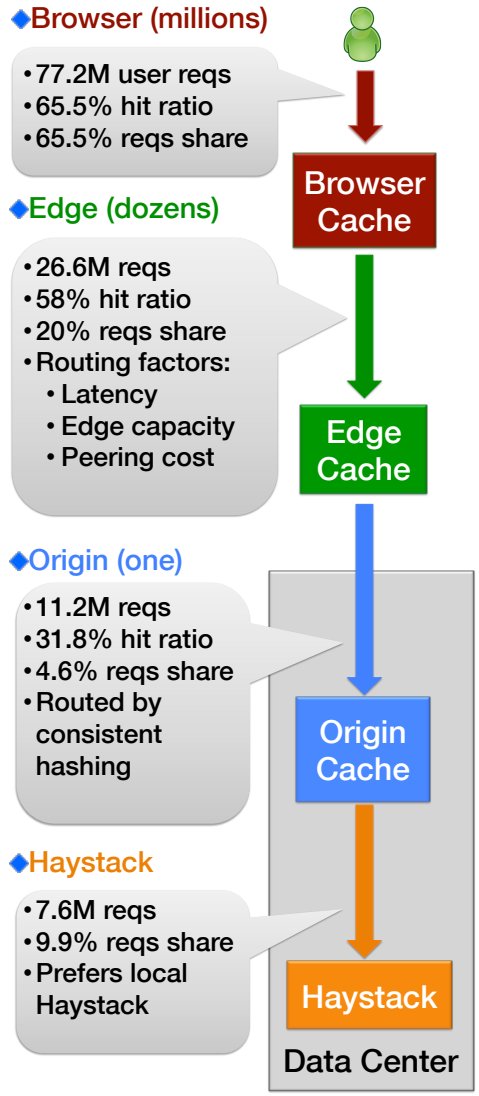

Workload

At top layers, req popularity follows a power-law dist., but curve flattens as reqs tunnels deeper.

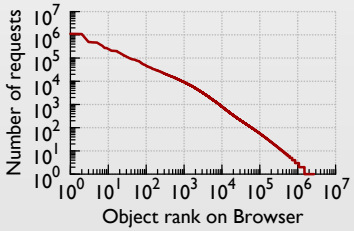

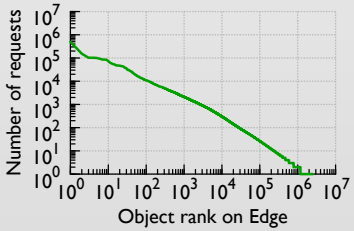

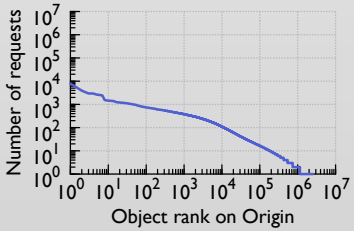

Haystack sees a Stretched Exponential dist.

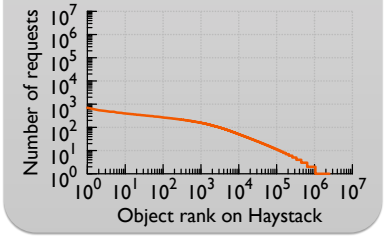

Cache Performance

Traffic Share by Photo Popularity

Cache traffic share drop for less popular items.

1. Top 1K photos attract 25% traffic.
2. Cache serves 99.93% reqs for them.
3. Haystack handles the tail.

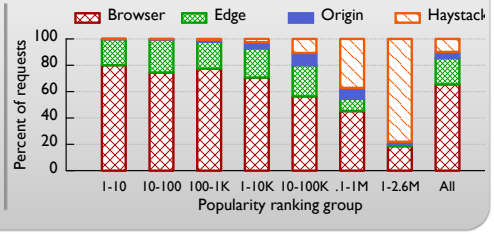

Browser Caching

1. Clients with <10 reqs send 37% traffic.
2. Active clients have higher hit ratio.
3. Increasing cache size helps.

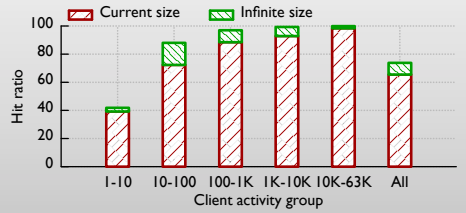

Edge Caching & Origin Caching

1. Request from clients are often routed to remote Edges.
2. Collaborative Edges (collab bar) increases hit ratio by 18%.

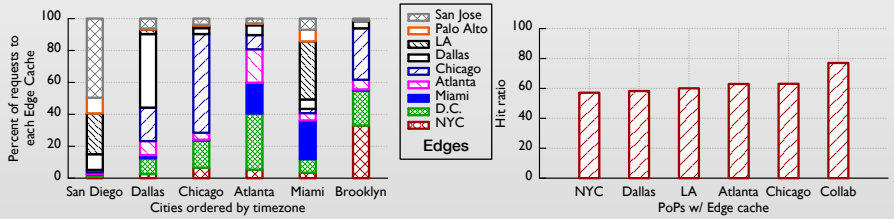

3. S4LRU increases hit ratio significantly both at Edge and Origin.

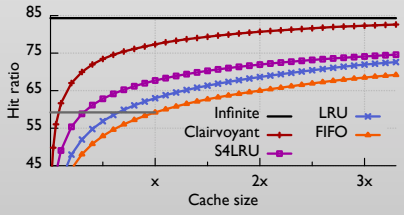

San Jose Edge

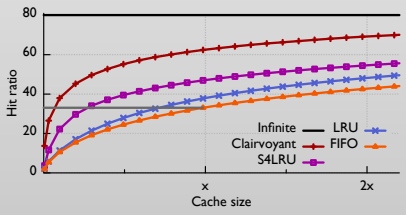

Origin