

On Caching Search Engine Query Results

Josh Miller

09.26.02

Goals of the Paper

- Evaluate the effectiveness of several variations of search engine query result caching
- Determine if temporal locality exists between searches, and if so, how much
- Determine if static or dynamic caching is most efficient

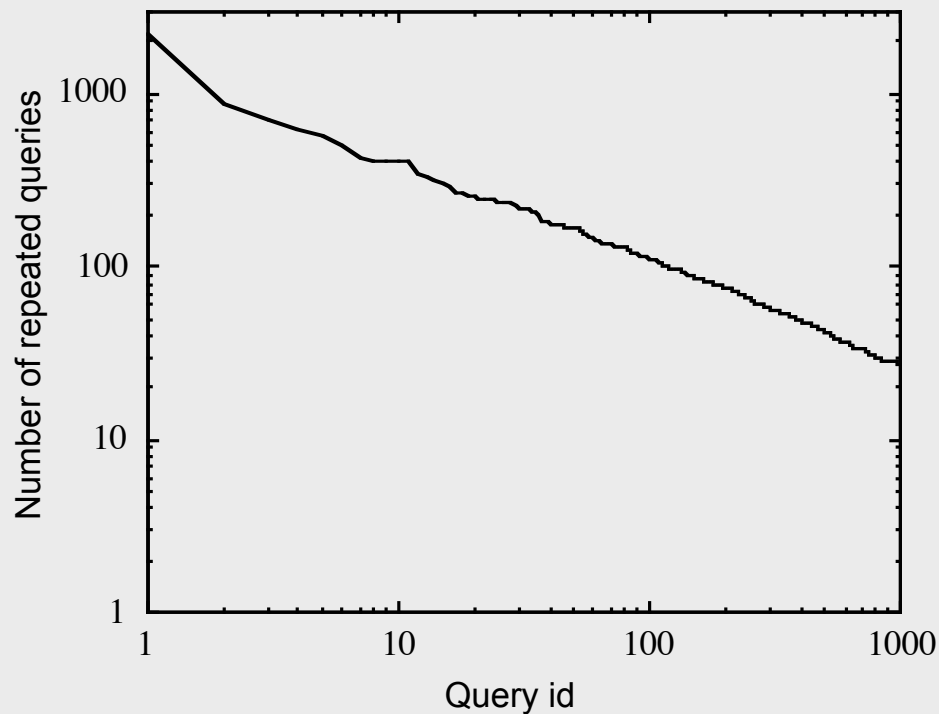
Background

- Caching improves responsiveness on the web by storing frequently accessed pages (query results).
- Caches are used as “accelerators” on search engines to reduce latency.
- Depending on the implementation, caching can affect the freshness of the data returned

Terms

- Temporal Locality = Query Interval
- Hit Rate = Performance = Fraction of all accesses satisfied by the cache

Repeated Queries




**Number of accesses for 1,000
most popular queries**

- Excite: average query is submitted 1.33 times
So, with caching, the hit rate could reach up to 25%
- AltaVista: average query submitted 3.97 times - Hit rate ~75%

(note: AltaVista ~1,000 times larger trace than Excite data)

Temporal Locality





- There are ~1,500 repeated queries which are submitted in an interval of ~100 other queries— So, temporal locality in queries exist

Distances between submissions of the same query

Replacement Algorithms

- **LRU**- Replaces the least recently used query results in the cache
- **FBR**- Like LRU, but accounts for how recently the query was performed, and how frequently it was performed
- **LRU/2**- Like LRU, but keeps track of the second to last most recent access
- **SLRU**- A more intelligent FBR algorithm


Replacement Algorithms in Action



Static vs. Dynamic Caching

- Static caching — Small — Generally used to hold most popular results.
May include human interaction, and may be out-of-date
- Dynamic caching — Large (as discussed in previous slides), uses some replacement algorithm based on the most recent accesses

Static vs. Dynamic Caching



Conclusions

- Locality in search engines exists — Nearly every term has been searched for before
- Medium-sized caches have a hit rate of about 20 to 25%, which eliminates 1/4 of the load on search engines
- Frequency and recency are important for replacement algorithms (because LRU performed the worst)
- Static caching is good for small caches, dynamic is good for large

Questions

- How could static and dynamic caching be used together?
- Your questions...