HyperDex

A Distributed, Searchable Key-Value Store for the Cloud

Robert Escriva[†] Bernard Wong[‡] Emin Gün Sirer[†]

[†] Department of Computer Science
Cornell University

[‡] School of Computer Science
University of Waterloo

EaGL-SysNet, August 20, 2011

Motivation

- Key-value stores are fast with a limited interface
- SQL databases are slow with an expressive interface

Can we expand the key-value store's interface to support equality and range searches?

Consistent Hashing

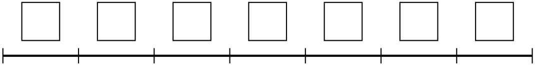

Consistent Hashing

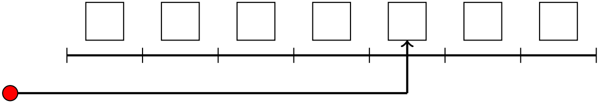

FIRST="Robert" LAST="Escriva", PHONE=(555) 123-4567

Consistent Hashing

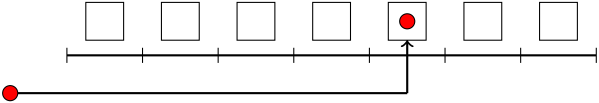

FIRST="Robert" LAST="Escriva", PHONE=(555) 123-4567

Consistent Hashing

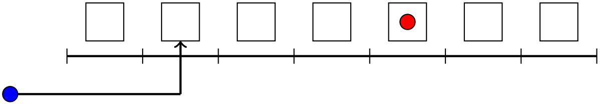

FIRST="Robert" LAST="De Niro" PHONE=(555) 890-1928

Consistent Hashing

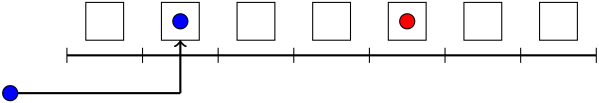

FIRST="Robert" LAST="De Niro" PHONE=(555) 890-1928

Consistent Hashing

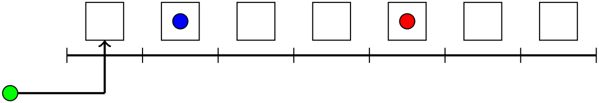

FIRST="Gün" LAST="Sirer" PHONE=(555) 828-3156

Consistent Hashing

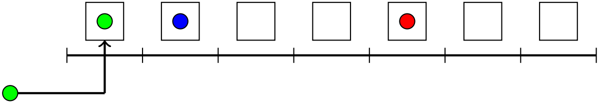

FIRST="Gün" LAST="Sirer" PHONE=(555) 828-3156

Consistent Hashing

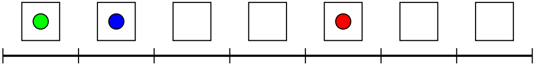

HyperSpace Hashing

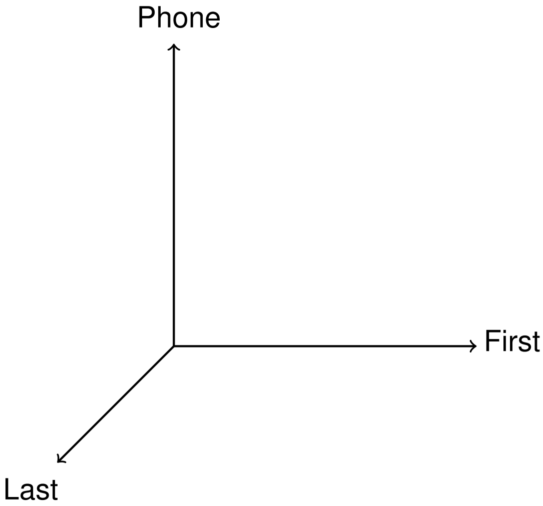

HyperSpace Hashing

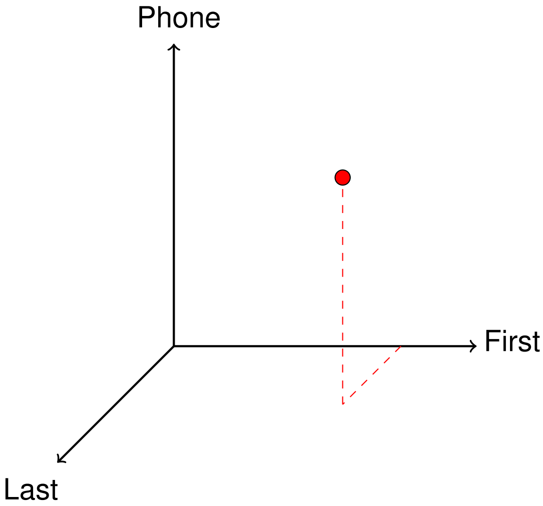

HyperSpace Hashing

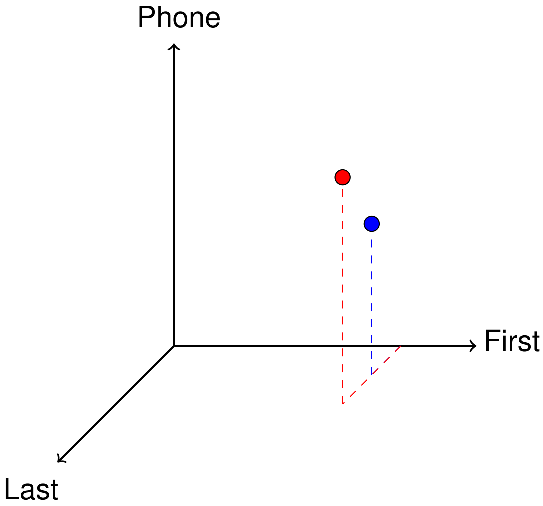

HyperSpace Hashing

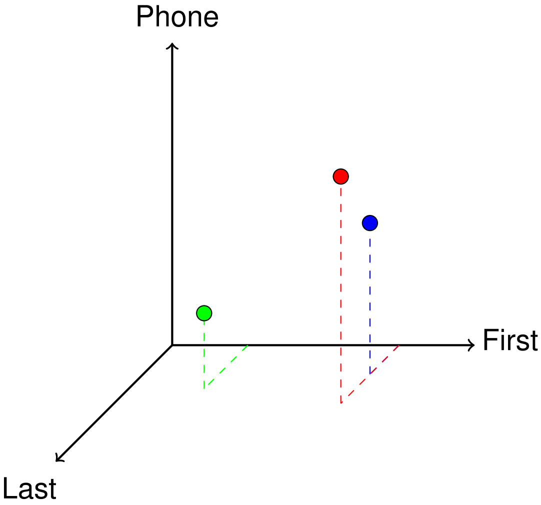

HyperSpace Hashing

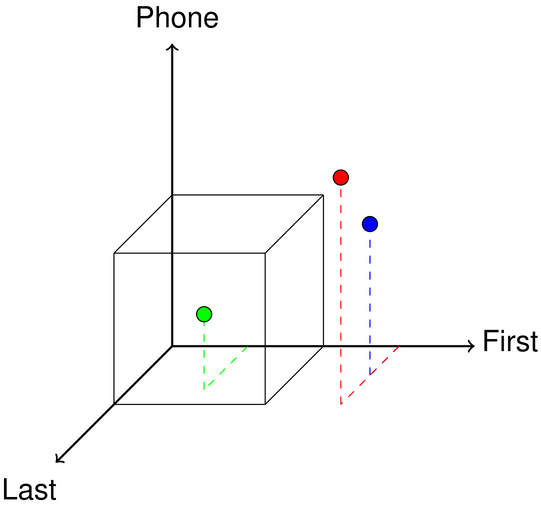

HyperSpace Hashing

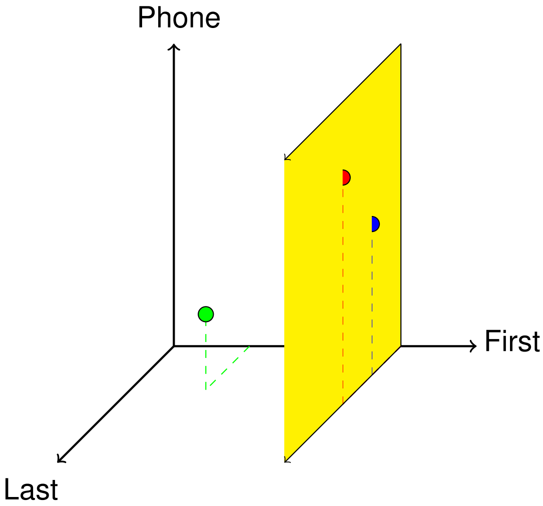

HyperSpace Hashing

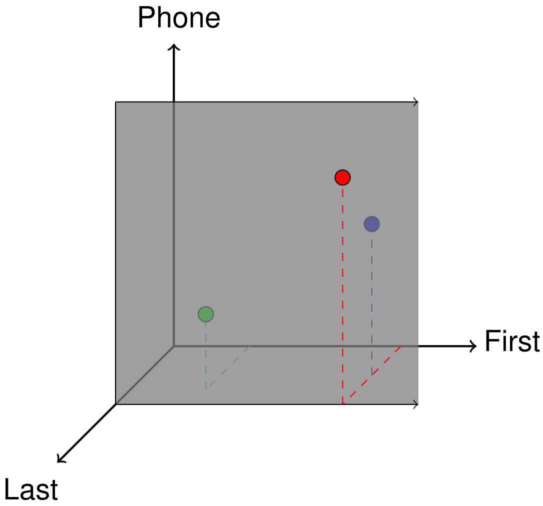

HyperSpace Hashing

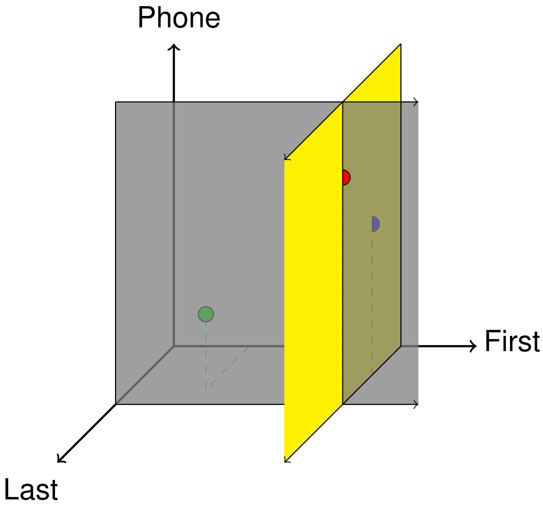

HyperSpace Hashing

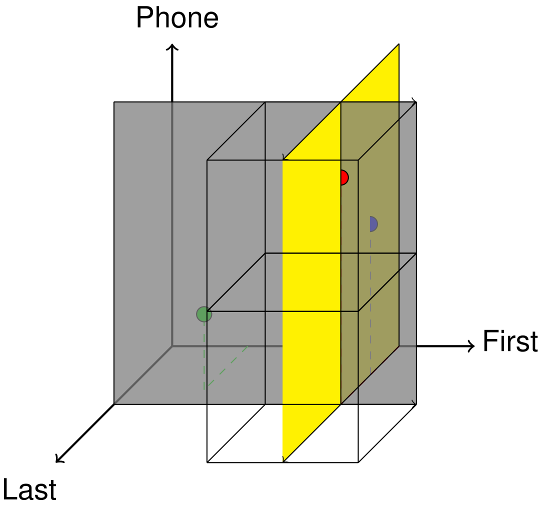

Issues

- Curse of hyperspace
    - Subspace folding
- Fault tolerance
    - Chain-replication-inspired design
- Consistency
    - Linearizability
- Performance with 10% Searches
    - Cassandra with indexing 555 ops/s
    - MySQL 4,925 ops/s
    - HyperDex 36,034 ops/s

Summary

- HyperDex uses hyperspace hashing to efficiently search values
- Contact: `escriva@cs.cornell.edu`