




Histogram-based P2P Main Memory Database for Locality-Aware Data

Tobias Scholl, Bernhard Bauer, Benjamin Gufler, Richard Kuntschke
Daniel Weber, Angelika Reiser, and Alfons Kemper



Lehrstuhl Informatik III: Datenbanksysteme - Technische Universität München - Fakultät für Informatik


The Challenge



The Multi-Wavelength-Milkyway:
<http://adc.gsfc.nasa.gov/mw/>

- Globally distributed archives
- Correlation of data from TeraByte-scale data sets
- Skewed data
- Region-based queries


Distributed P2P Architecture



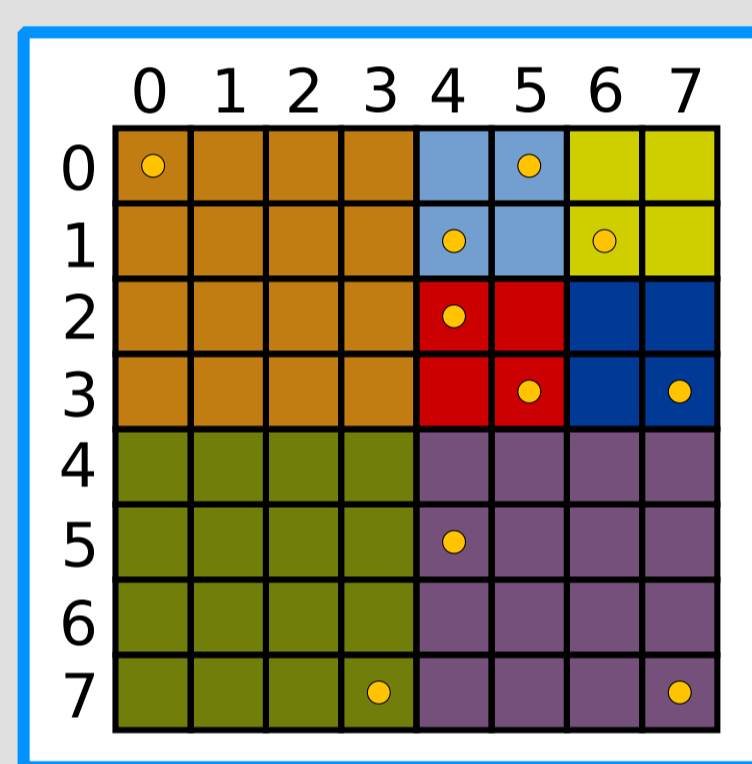
- Data feeding
- High throughput query processing
- Resource sharing (CPU and memory)

"Distribute by Region — not by Archive!"

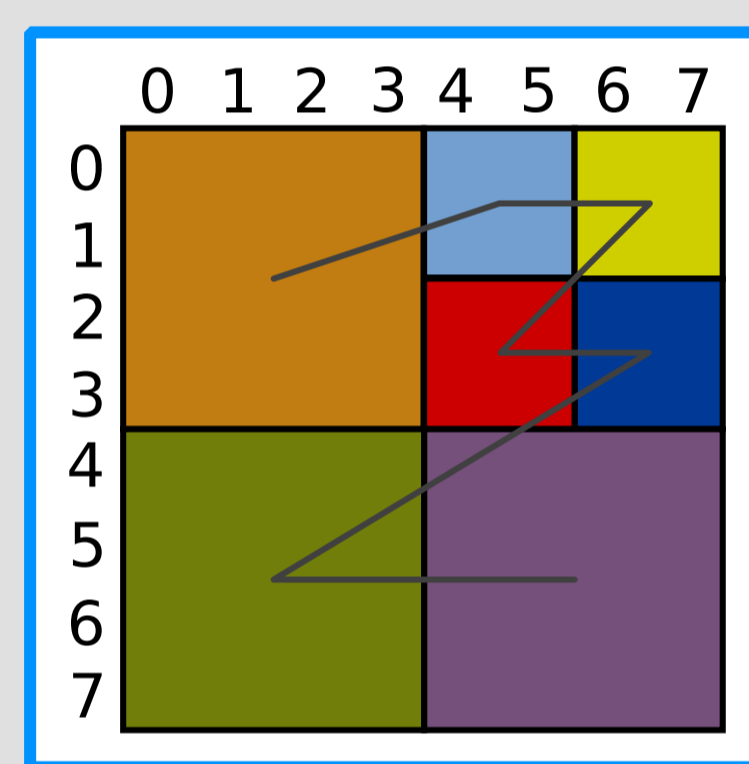
Z-Quadtrees




Training set



Histogram regions

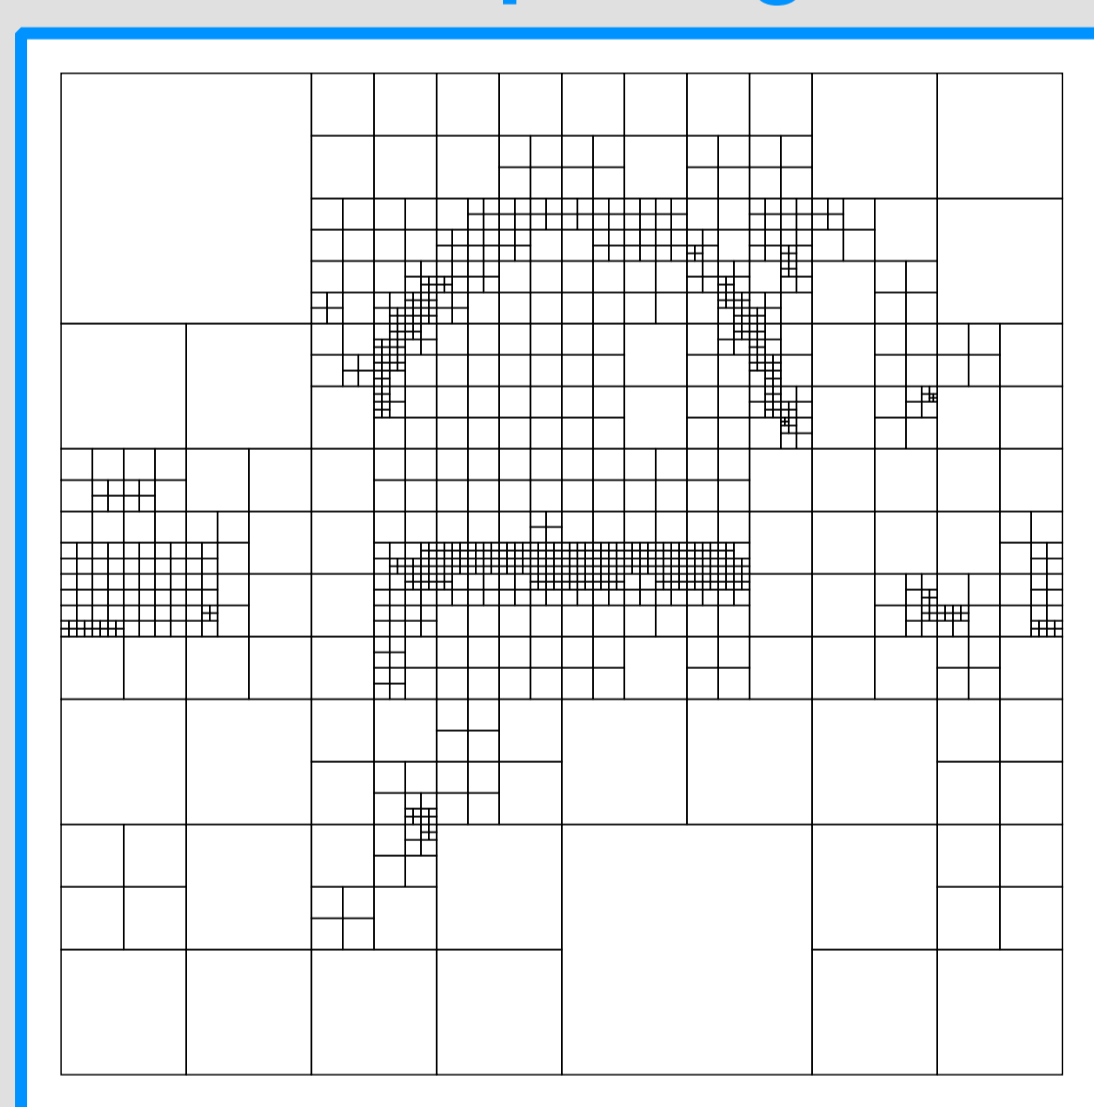


Z-Linearization




Quadtree

Center splitting rule




Median splitting rule



- Regular decomposition
- D-dimensional region: 2^d subregions
- Quadtree leaf \leftrightarrow histogram region
- Adaptive resolution


Example Training Set



Data set fragments:


- ROSAT: 25.617.873 objects
- SDSS: 84.323.326 objects
- TWOMASS: 28.445.694 objects

Data & Peer Mapping




- Uniform region mapping
- Random peer mapping
- Routing towards region-ids (from Peer a to Region 3)

Query Processing



- 1 Client (I) issues query to peer (P).
- 2 P determines relevant regions (D1 - D4) and coordinator region.
- 3 Delegate query to coordinator (C).
- 4 Coordinator starts query processing and routes to relevant regions.
- 5 Coordinator collects intermediate results and sends final result to client.

Local Network Benchmarks



- Standalone HiSbase (2GB Memory) vs. 16-peers (0.5 GB)
- 730 (1430) Cross-matching queries
- Throughput measurements
- Standalone: Highest throughput at 10-20 queries
- 16-peers: Highest throughput at 600 queries [9600 total]