Structured and Unstructured Overlays for Distributed Data

Arvind Krishnamurthy
Yale University

Joint work: Chi Zhang, Anthony Young, Randy Wang

Evolution of Overlays in P2P Systems

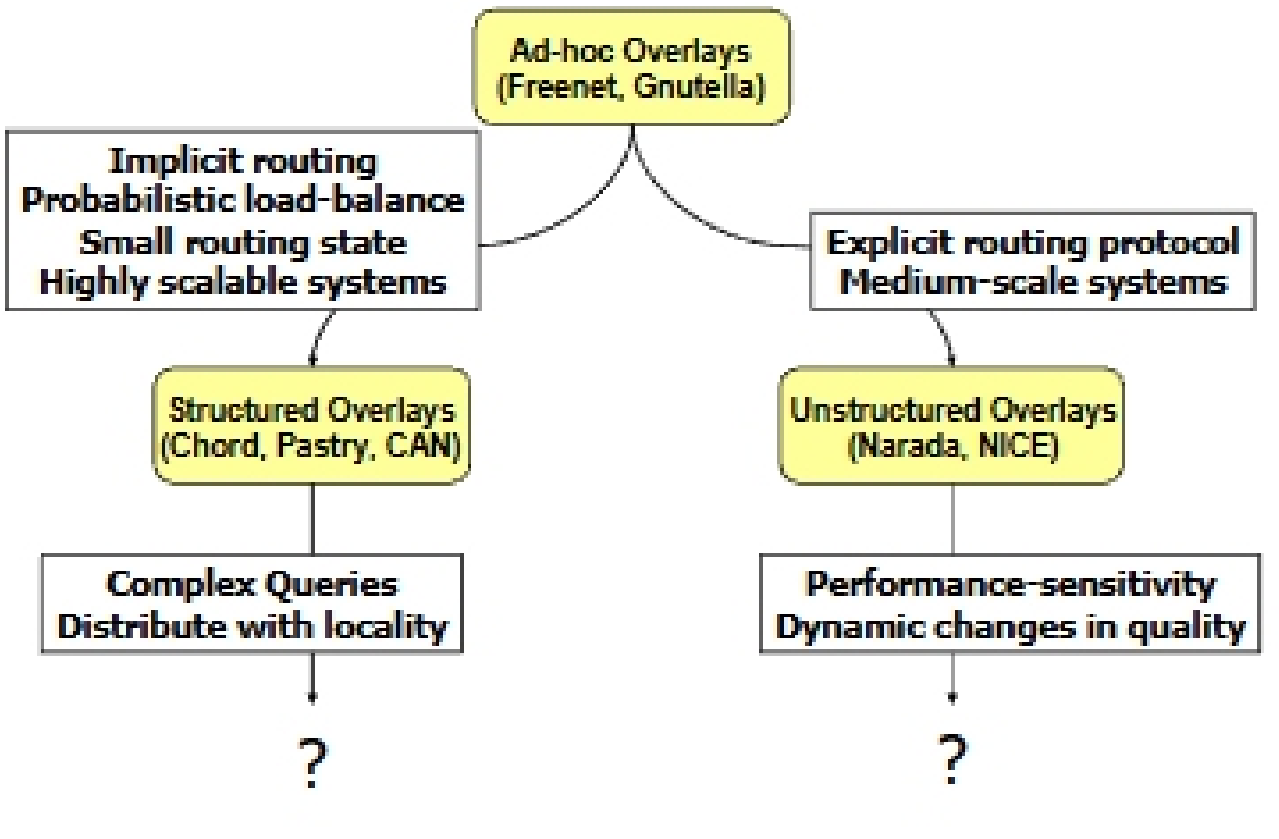

Motivating Factors

1) Complex data
   - Massive, distributed data collections (web documents, multimedia files)
   - Dynamic data (sensor data, network monitoring, etc.)
   - High-dimensional (or at least multi-dimensional)

2) Complex queries
   - Similarity searches, nearest-neighbor queries
   - Range queries

3) Locality-preserving mapping for better performance

Approach #1: Use DHTs

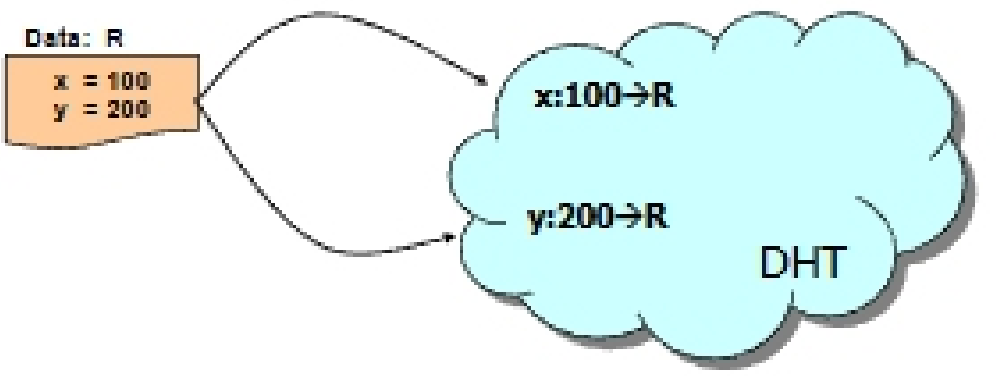

- Store attributes of a multi-dimensional object in a DHT
- Separate entry made for each attribute to allow for queries on individual attributes
- Can support only exact searches
- Related objects are hashed to arbitrary locations
- Exemplified by PIER; PIER supports arbitrary SQL commands

Approach #2: Hashless DHTs

- Exemplified by Mercury
- Maintain a multi-dimensional index as a set of one-dimensional indexes
- Each index stored in a Chord-like ring *without hashing*

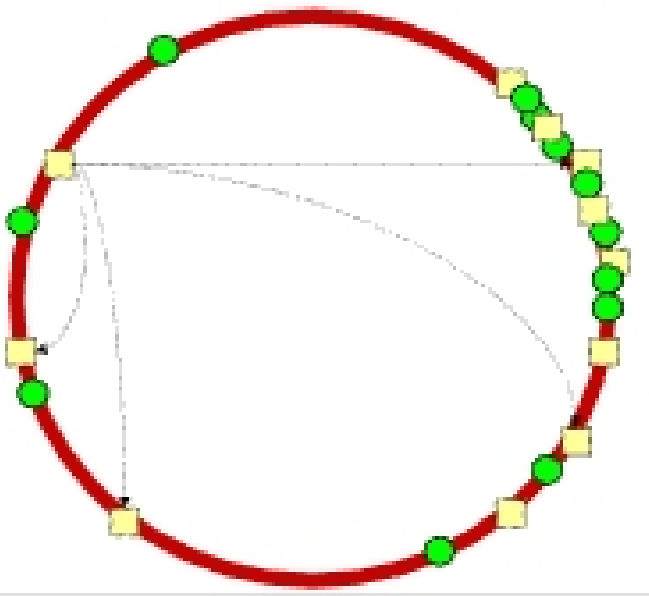

Data skews require:

1) processor redistibution
2) adjustments to "fingers"

Support for Multi-dimensional Objects

- Maintain a collection of rings for one-dimensional indexes
  - Objects are published in all rings
  - Query sent to only one ring

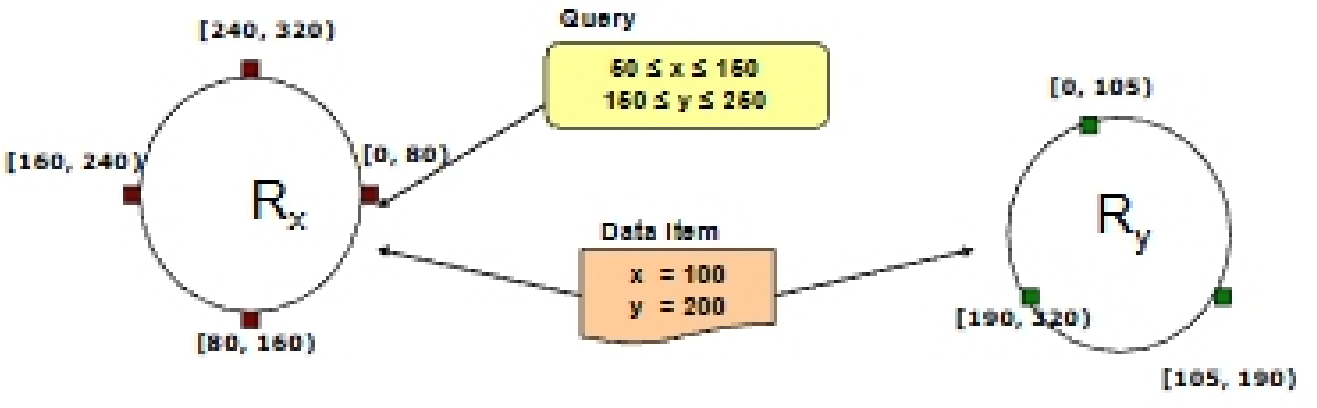

- Can support complex queries, but incurs storage overhead; not suitable for high-dimensional data