# What Can Databases Do for Peer-to-Peer

Steven Gribble, Alon Halevy,
Zachary Ives, Maya Rodrig, Dan Suciu

Presented by: Ryan Huebsch
CS294-4 P2P Systems – 11/03/03

## Outline

- Disclaimer: This is a position paper, not a technical/system paper (no graphs)
- Author's Mindset
- Data Placement
- Complexity
- Piazza

## Why P2P?

- Desirable properties of P2P system amplified with new peers
  - Robustness
  - Availability
  - Performance
- Decentralization for trust reasons & administration
  - No proprietary interests
  - Trust is diffused over all participants

## What is the problem?

- Gnutella failed to attract people because of
  - Weak application semantics (search for filename, what does the filename mean?)
  - Technical flaws limit scaling (short term problem?)
- Ad-hoc membership
  - Difficult to predict resources and load
  - Thus, data placement is demand driven (for lack of better mechanism)
  - May cause fundamental limits on consistency and availability

## Why Databases?

- The problem is placement and retrieval of ***data***… that would be a data management (or DB) problem
- P2P world is lacking
  - Semantics
  - Data transformation
  - Data relationships
- All of which are core strengths of the DB community
- P2P brings a new environment for DB query processing systems
  - increased scalability, reliability, and performance
- This paper focuses on the data placement problem

## Data Placement Problem

- Setup
  - Set of cooperating nodes (no adversaries)
  - Bottlenecks: network, CPU, or memory
  - Nodes serve four roles
    - Data Origin – producers
    - Storage Provider
    - Query Evaluator
    - Query Initiator – consumers
  - Cost of query = Origin or Storage ➔ Evaluator + Evaluator ➔ Initiator

## Design Choices

- Score of decision making
  - Global (hard, optimal) or local (easy, short-sided)
  - Similar to multi-query optimization
- Extent of knowledge sharing
  - Knowledge of materialized views on other nodes (a catalog)
  - Centralized or distributed? Hierarchical (like DNS)?
- Heterogeneity of information sources
  - Few authoritative sources, lots of data producers
  - Heterogeneous data ➔ different schemas

## Design Choices II

- Dynamicity of participants
  - Node churn
  - Some nodes act like servers, some like workstations
  - Could place all data on servers ➔ reduced flexibility and performance
- Data granularity
  - Atomic granularity ➔ indivisible objects (complete file)
  - Hierarchical granularity ➔ groups (albums, directories)
  - Value based granularity ➔ Objects composed of atomic value (tuples composed of values)

## Design Choices III

- Degrees of replication
  - One copy all the way to fully replicated
  - More replicas make updates harder
  - Also makes retrieval harder (more choices)
  - Consistency is harder, typical solution is to have a master replica
- Freshness and update consistency
  - Invalidation messages, pushed by server on update or pulled by client on request
  - Timeout based, lower overhead, looser guarantees about freshness and consistency

## Complexity of Problem

- The papers goes to some trouble to formally define the problem
- Defines a small sub-problem of data placement,
  - Static P2P network
  - Queries are zero-cost
  - Problem: Which nodes an item go on?
- Problem is NP complete, proof comes from vertex-cover, not in this paper

## Piazza

- Peers form small groups called spheres of cooperation.
  - May follow administrative boundaries
  - Spheres of cooperation are nested
- Query Optimization problems:
  - Exploit commonalities between queries
  - Decide where to place data
  - What queries to materialize (store answers)
- To make the problem tractable, optimization occurs within a sphere of cooperation.

## Piazza II

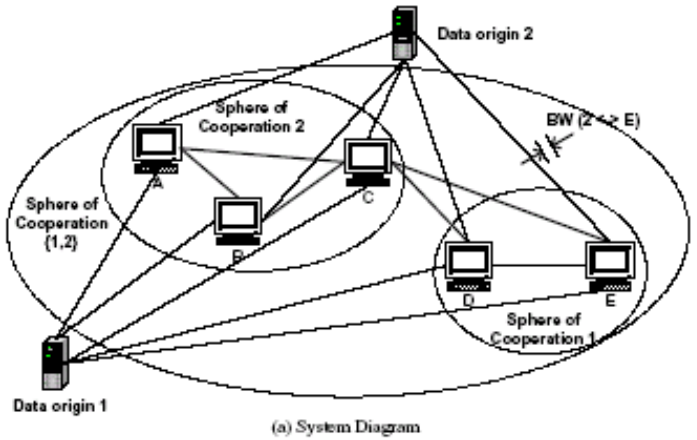

(a) System Diagram

## Piazza III

- Propagating Information
  - Node advertises its materialized views to its neighbors
  - Nodes consolidate info they receive and propagate
  - Type of gossiping protocol
- Consolidating Queries
  - Some queries can not be evaluated if data is not locally available
  - Broadcast all un-evaluatable queries to local sphere of cooperation, and try to answer them collectively

## Where is Piazza now?

- Focusing more on data semantics and information integration
- Every nodes has its view of what the data schema is
- ***Very Difficult*** problem that most people in the database community have ignored.