

Supporting Relational Queries in Peer-to-Peer Systems

Joseph Spadavecchia and Stratis D. Viglas



Introduction



- *Our aim is to support relational queries in a P2P environment – to harness distributed resources.*
- *The problem is that data is not located on local disk, but instead we must go out into the network to **locate** it.*
- *The system should **scale well** while nodes join and leave; have **guaranteed bounds** on the number of routing hops used for lookup; and be **fault tolerant** when nodes fail.*

Problem

Currently, P2P systems supporting multiple attribute indexing are inefficient and have **no guarantees** on information retrieval.

Problem

Currently, P2P systems supporting multiple attribute indexing are inefficient and have **no guarantees** on information retrieval. → **unstructured systems**

Problem

Currently, P2P systems supporting multiple attribute indexing are inefficient and have **no guarantees** on information retrieval. → **unstructured systems**

- *Gnutella/Kazaa/*Mule*
- *Research projects: SONS [1], Schema-Based P2P System [2]*

Problem

Currently, P2P systems supporting multiple attribute indexing are inefficient and have **no guarantees** on information retrieval. → **unstructured systems**

- *Gnutella/Kazaa/*Mule*
- *Research projects: SONs [1], Schema-Based P2P System [2]*

Systems with guaranteed bounds on retrieval, known as **structured systems**, have **no support for multiple attribute indexing**.

Problem

Currently, P2P systems supporting multiple attribute indexing are inefficient and have **no guarantees** on information retrieval. → **unstructured systems**

- *Gnutella/Kazaa/*Mule*
- *Research projects: SONs [1], Schema-Based P2P System [2]*

Systems with guaranteed bounds on retrieval, known as **structured systems**, have **no support for multiple attribute indexing**.

- *Chord/Koorde/Kademlia/Viceroy and DBTs (Distributed Balanced Tables)*
- *Query Processing [3], Data Indexing [4]*

Our Approach

Create a d -dimensional DBT.

Our Approach

Create a d -dimensional DBT.

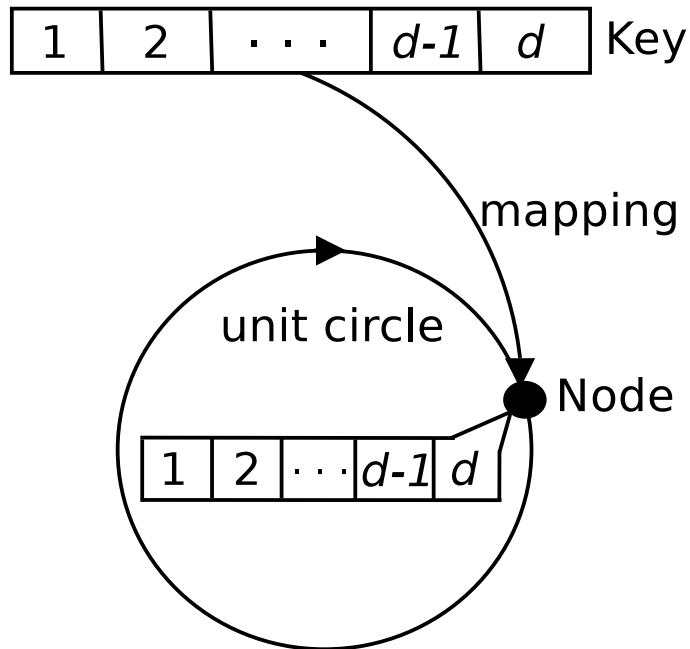
1	2	...	$d-1$	d
---	---	-----	-------	-----

 Key

- *Index data with d -dimensional keys.*

Our Approach

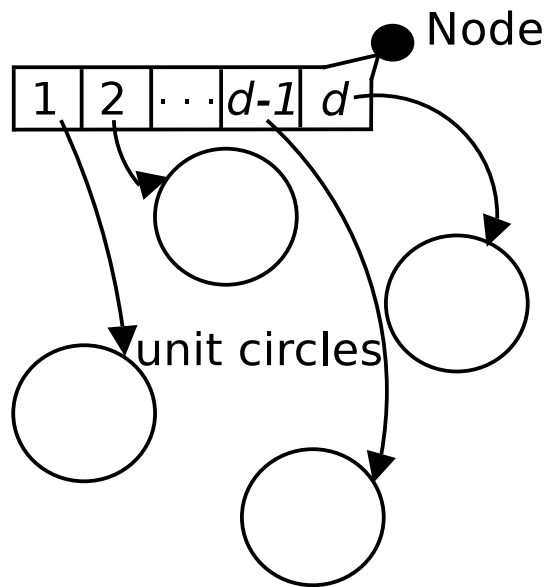
Create a d -dimensional DBT.



- *Index data with d -dimensional keys.*
- *Map the keys to nodes, which also have d -dimensional identifiers.*

Our Approach

Create a d -dimensional DBT.



- *Index data with d -dimensional keys.*
- *Map the keys to nodes, which also have d -dimensional identifiers.*
- *Each node belongs to more than one overlay, one for each dimension.*

Our Approach

Each node must maintain a different set of routing table entries for each dimension.
This can be costly.

Our Approach

Each node must maintain a different set of routing table entries for each dimension. This can be costly.

- *Space-filling Curves to reduce the number of dimensions. (The Hilbert Curve).*

Our Approach

Each node must maintain a different set of routing table entries for each dimension. This can be costly.

- *Space-filling Curves to reduce the number of dimensions. (The Hilbert Curve).*
- *Making the system adapt according to the attributes (dimensions) that are most frequently searched.*

Work in Progress

- *Finishing up simulator*
- *Have many experiments to conduct*
- *Application benchmarks*

References

References

- [1] Arturo Crespo And Hector Garcia-Molina "Semantic Overlay Networks for P2P Systems," Technical Report, Computer Science Department, Stanford University, 2002.
- [2] Ingo Brunkhorst, Hadami Dhraief, Alfons Kemper, Wolfgang Nejdl and Christian Wiesner "Distributed Queries and Query Optimization in Schema-Based P2P Systems," in *International Workshop On Databases, Information Systems and Peer-to-Peer Computing*, September 7-8, 2003 Humboldt University, Berlin, Germany.
- [3] O. Sahin and A. Gupta and D. Agrawal and A. Abbadi "Query processing over peer-to-peer data sharing systems," Technical Report UCSB/CSD-2002-28, University of California at Santa Barbara, 2002.
- [4] P.A. Felber and E.W. Biersack and L. Garcs-Erice and K.W. Ross and G. Urvoy-Keller "Data Indexing and Querying in DHT Peer-to-Peer Networks," Institut EURECOM, 06904 Sophia Antipolis, France.