

# Supporting Multi-dimensional Range Queries in Peer-to-Peer Systems

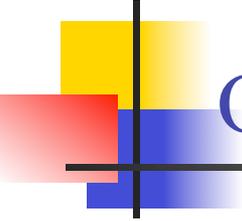
---

**Yanfeng Shu, Beng Chin Ooi, Kian-Lee Tan**

School of Computing, National University of Singapore

**Aoying Zhou**

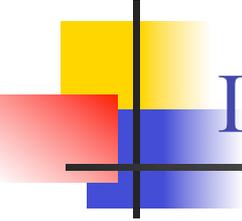
Department of Computer Science, Fudan University, China



# Outline

---

- n Introduction
- n ZNet
- n Experimental Results
- n Conclusion



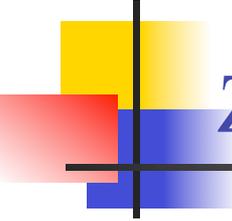
# Introduction

---

- ⌘ Today's P2P systems are unable to cope well with range queries on multi-dimensional data
  - ⌘ *Unstructured systems (e.g., Gnutella)*
    - ⌘ mainly depend on flooding
    - ⌘ no performance guarantee
    - ⌘ cannot ensure data availability
  - ⌘ *Structured systems (e.g., Chord)*
    - ⌘ Have guarantee on search efficiency and data availability
    - ⌘ mainly designed for *exact* key lookup

To extend existing P2P systems and thus support range queries on multi-dimensional data *efficiently* and *effectively*:

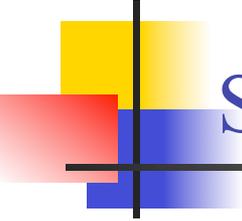
- n *Space Partitioning and Mapping*
  - how is a data space partitioned and mapped onto nodes?
- n *Query Processing*
  - what overlay is based on and what strategy is employed for query processing?
- n *Load Balancing*



# ZNet

---

- n The whole data space is dynamically partitioned and mapped
  - *space is partitioned into different granularities*
  - *Space Filling Curves(SFCs) at different orders are used to preserve data locality*
- n Range queries are efficiently supported
  - n *Skip graphs are extended for query routing, with each node maintaining only  $O(\log N)$  states;*
  - n *An efficient range query resolution strategy is proposed, which evaluates queries in a specific way to avoid unnecessary node visits.*
- n Both Static and Dynamic load balancing are addressed



# Space Partitioning and Mapping

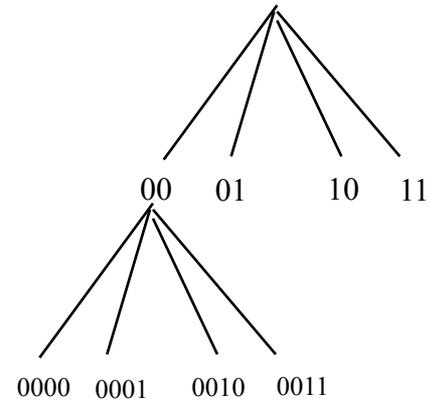
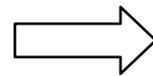
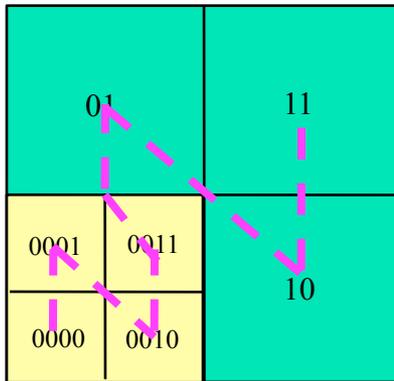
---

- n When?
  - n *Node join, leave/failure*
  - n *Load balancing*
- n How?
  - n *halved in all dimensions*
  - n *each subspace is called a zone, which is numbered according to its position w.r.t. the split dimensions*
- n The whole partitioning process can be modelled as a tree, *partition tree*

# Space Partitioning and Mapping

cont

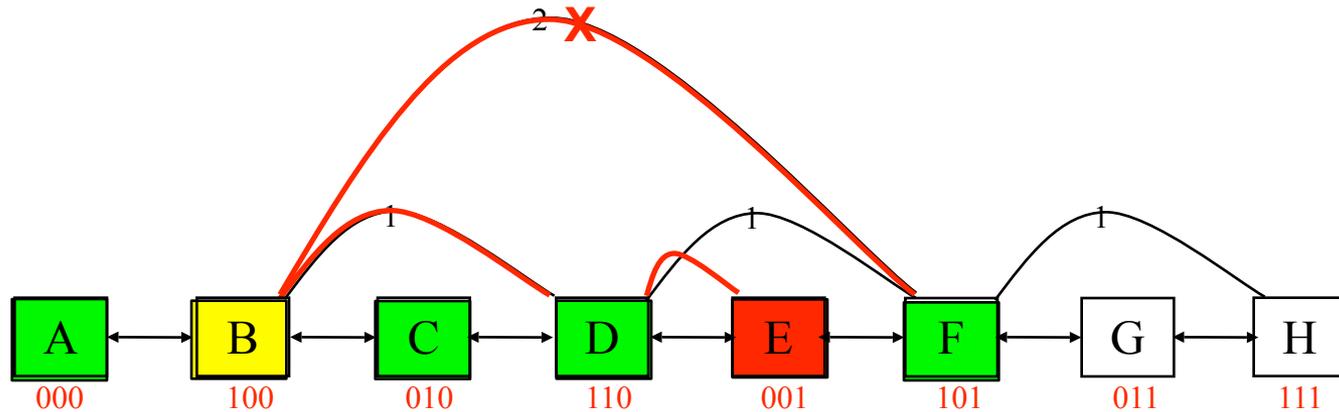
...

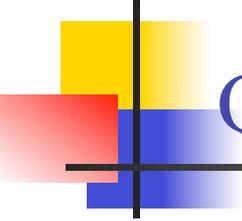


- n This equals to filling zones (from one partitioning) with a first order z-curve. The whole space is filled with z-curves at different orders, and a zone' code is just its Z-address
- n Continuous zones whose Z-addresses are continuous are mapped to the same node or nearby nodes

# Query Processing

Skip Graphs: multiple-level linked lists with each node's links (neighbors) are selected at random





# Query processing

---

*cont*

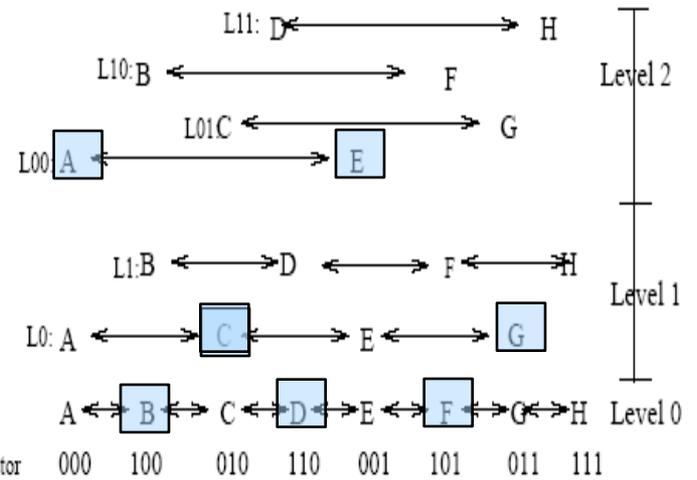
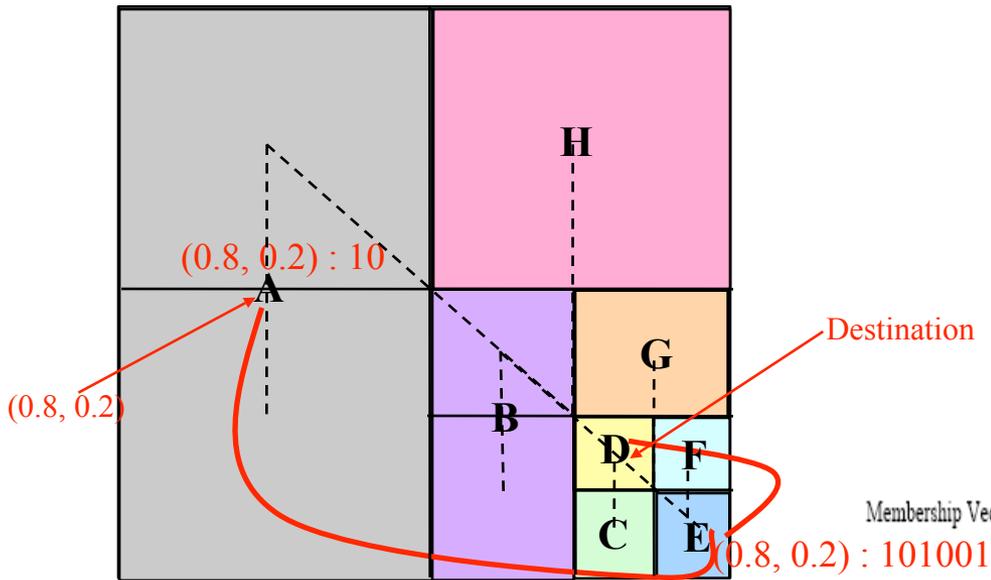
...

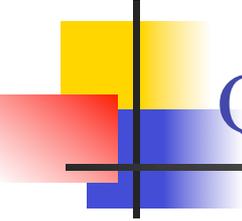
- n ZNet extends Skip Graphs by assigning continuous zones to nodes
- n Each node only needs to maintain  $O(\log N)$  states, where  $N$  is the number of nodes
- n When given a search key ( a point), a node will first transform the point to a z-address, then the search request is routed according to the z-address by following the routing process of Skip Graphs
- n Complexities rise in that each node has only incomplete knowledge about space partitioning

# Query processing

cont

...





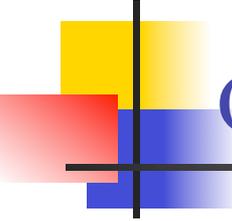
# Query Processing

---

*cont*

...

- n Point /Range queries
- n One method: convert the query range QR to a set of continuous Z-address ranges covered by QR, then send each request to each node which contains the minimum Z-address of each range
  - n *Inefficient*
  - n *Cannot deal with the dynamic case*
- n Our method: convert QR to a superset of zones covered by QR. As the query is routed, the superset is refined
  - n *The query is routed along two opposite directions to avoid unnecessary node revisits*



# Query Processing

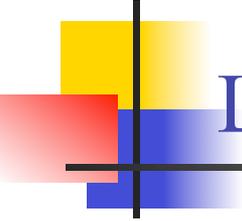
---

cont

...

An Example:

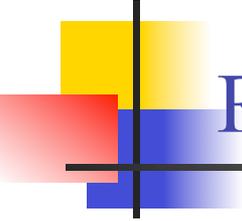
Given QR  $[(0.8,0.1),(0.9,0.2)]$ , which overlaps with C, D, E, and F, A first computes a superset of zones  $[10,10]$ , then routes it to E, where the superset is refined to  $[101000, 101011]$ . Since E's space overlaps with QR, the query will be partitioned into two parts:  $[101000,101001]$  and  $[101011,101011]$ , and so on....



# Load Balancing

---

- n For new nodes, appropriate joining destinations are chosen
  - *try several possible joining destinations, and the one which has the heaviest load will be chosen as the joining destination.*
- n At run-time, heavily loaded nodes can migrate some of their load to lightly loaded nodes
  - n *A node is heavily loaded if its load is larger than  $\bar{L} + \delta\bar{L}$ ; and lightly loaded if its load is less than  $\bar{L} - \delta\bar{L}$ .*
  - n  *$\bar{L}$  is average load, which is approximated by sampling loads of a node's neighbors*
  - n *A node periodically changes load information with its neighbors, and heuristics are adopted to find lightly loaded nodes.*



# Recent Proposals for comparison

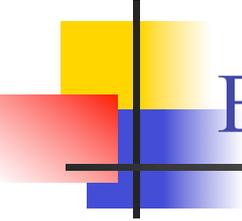
---

## n Squid and SCRAP

- n *Both use space-filling curves to map multi-dimensional space to one dimensional space as ZNet*
- n *They partition the space **statically** – the partitioning level needs to be decided beforehand*
- n *Squid's performance may be affected by data skewness, as it is based on Chord. SCRAP has no efficient query processing scheme*

## n MURK and SkipIndex

- n *Both use K-d trees for space partitioning and mapping*
- n *They cannot cope well with dynamic load balancing*



# Experimental Setup

---

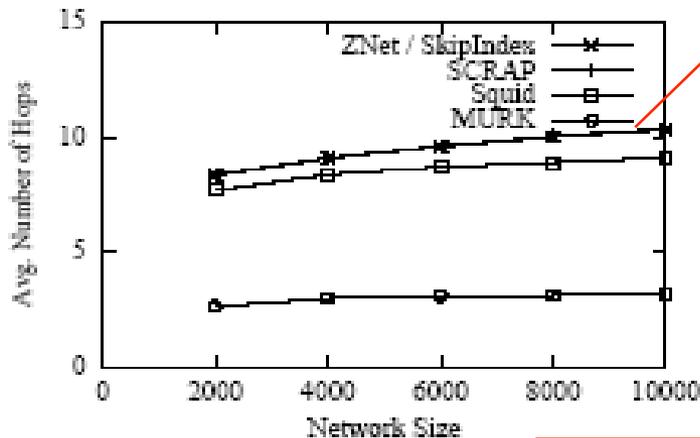
- n Experiments are done via simulation on two kinds of synthetic datasets: one is skewed datasets based on normal distribution; the other is uniform datasets
- n By default, we use data sets with skewed 8-dimensional 300,000 data, and 6,000 nodes in the network
- n The dimensionality is varied from 2 to 20, and the number of nodes is varied from 2,000 to 10,000

# Routing Cost

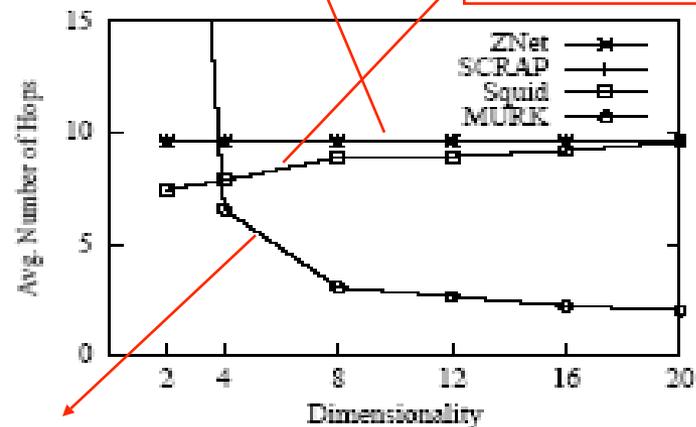
- Assume the maximum partition level of ZNet is globally known in SCRAP and Squid

The add-on dynamicity of Znet does not affect its routing performance

Squid is affected by data skewness



(a) Dimensionality = 8



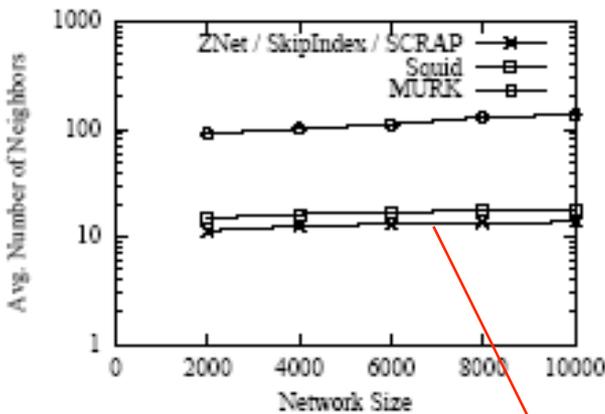
(b) Network size = 6000

MURK behaves badly when dimensionality is low

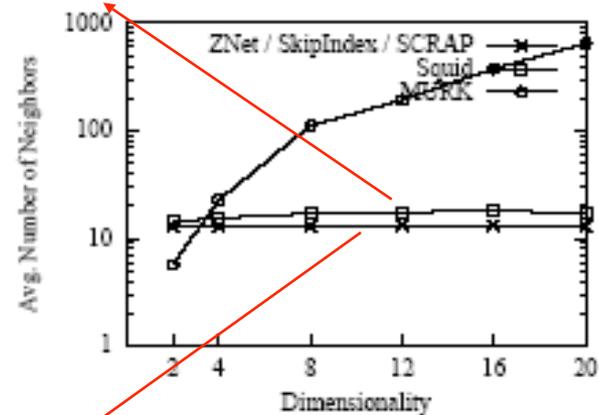
# Maintenance Cost

- n We measure maintenance cost in terms of the number of neighbors maintained by each node

Affected by data skewness



(a) Dimensionality = 8

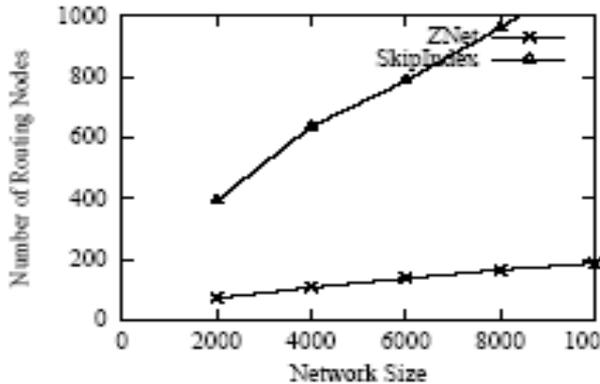


(b) Network size = 6000

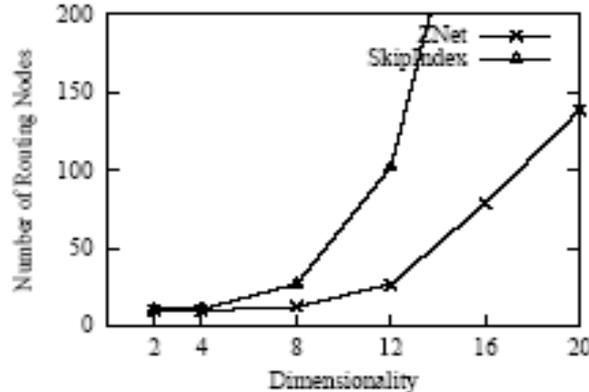
Increase logarithmically with the network size; independent of dimensionality

# Range Search Cost

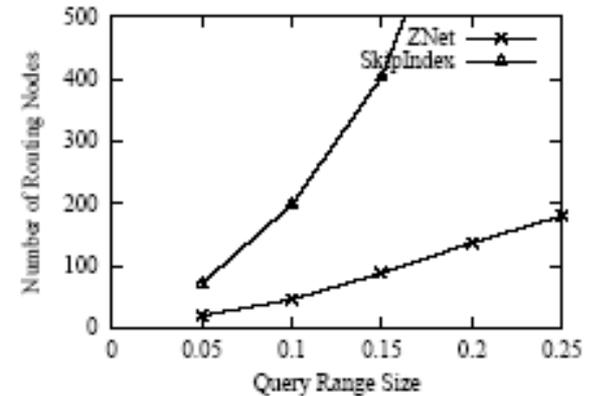
- n Range search cost is measured by the number of routing nodes
- n Four factors are involved: network size, dimensionality, query range size, and data distribution



(a)



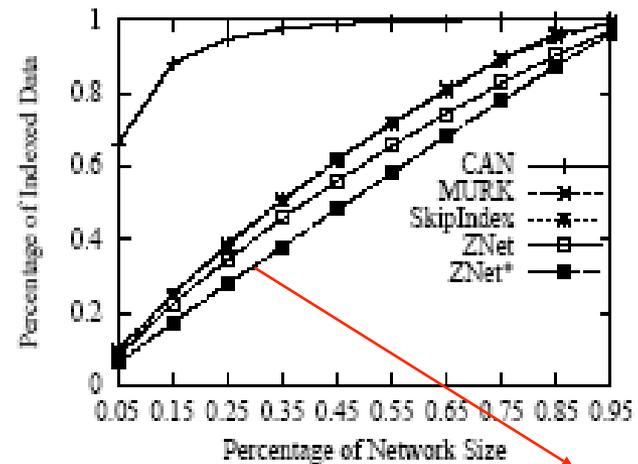
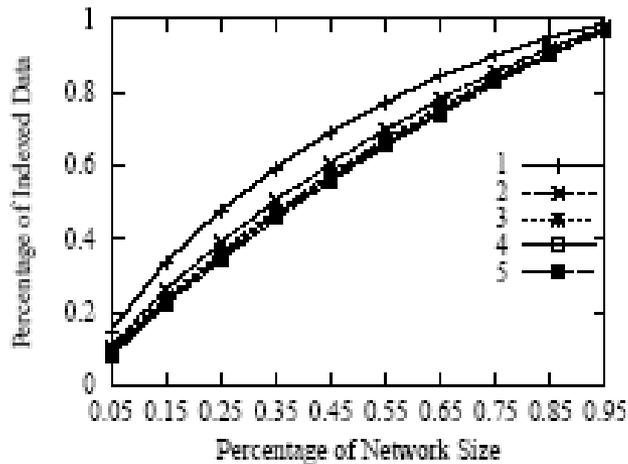
(b)



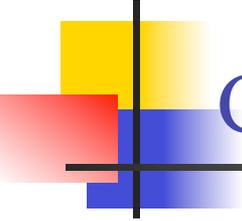
(c)

# Load Balancing

- n Load Balancing is measured by the amount of data indexed by each node



Znet\*: dynamic load balancing in ZNet



# Conclusion

---

- n Znet possesses nearly all desirable properties, while others typically fail in one or another:
  - n *has low routing cost and low maintenance cost*
  - n *Increase logarithmically with the network size*
  - n *independent of data dimensionality and distribution*
  - n *Low range search cost*
  - n *Can achieve better load balancing*