

LDHT: Locality-aware Distributed Hash Tables

Sagar Gaikwad
Big Data Lab

Agenda

- Distributed Hash Tables
- Related Work
- Design of LDHT
- Workflow of LDHT
- Performance Analysis
- Conclusion and Future Scope
- References

Distributed Hash Tables

- Substrate of P2P systems
- Scalable storage and retrieval of {key,value} pairs
- Data object is stored at the node whose node identifier is closely related to data object's key
- Latency issues

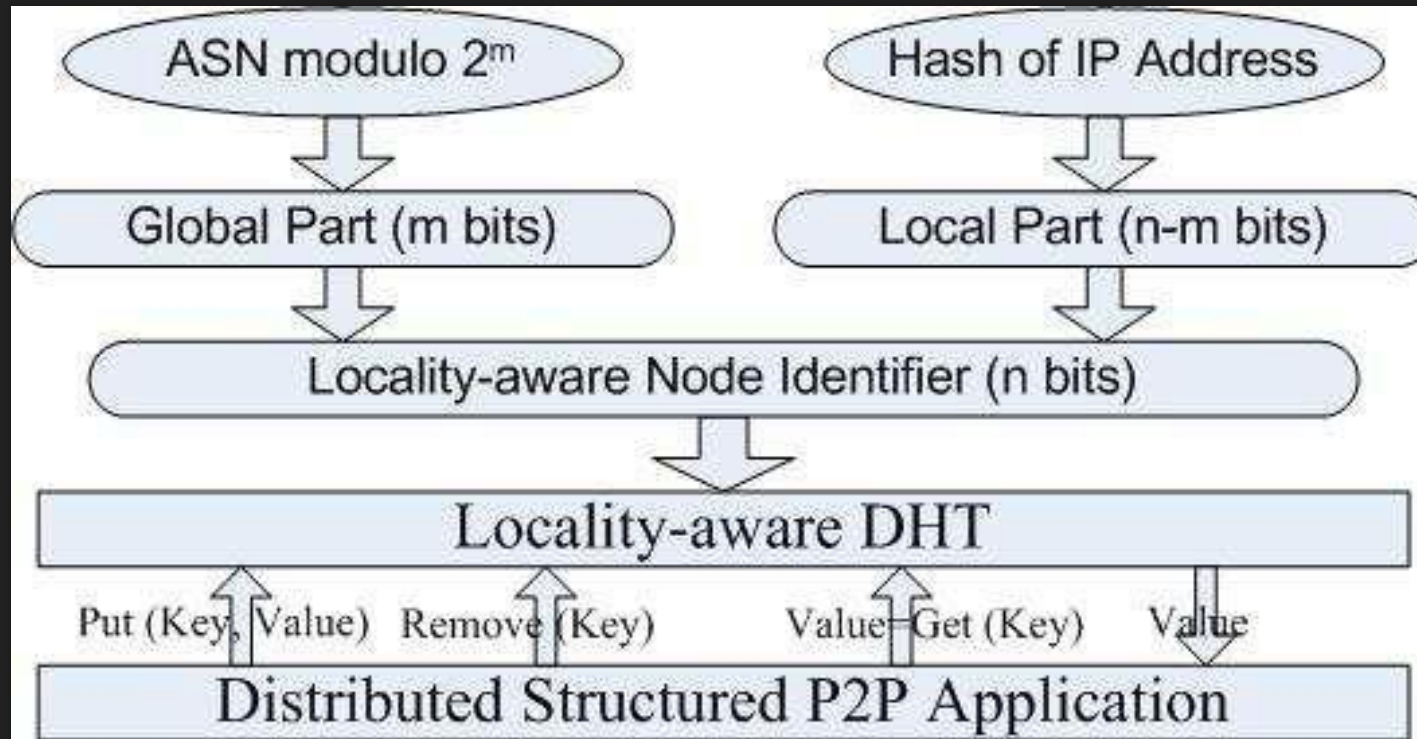
Related Work

- Proximity routing
 - Balance between making progress towards destination and selection of closest routing table entry according to locality
- Proximity neighbor selection
 - Variant of the above idea
 - Proximity criterion is used for not only for next hop selection but also choosing neighbors
- Geographic layout
 - Nodes closer in network topology are also closer in node identifier space.

Design of LDHT

- Basic idea
- Identifier assignments
 - ASN-based identifier assignment
 - Node identifiers divided into two subparts –
 - Global part
 - Local part
 - Global part + Local part = Locality-aware node identifier




Workflow of LDHT



Performance Analysis

- Basic DHT protocols –
 - Chord, Symphony, Kademia
- Evaluation metrics –
 - Path length
 - Relative Delay Penalty (RDP)
 - Hop Count
- Tested against two topologies
 - Topology 1 – generated by GT-ITM[2] , consists of 4000 nodes
 - Topology 2 – real-world internet data set with 4520 nodes

Path length per query




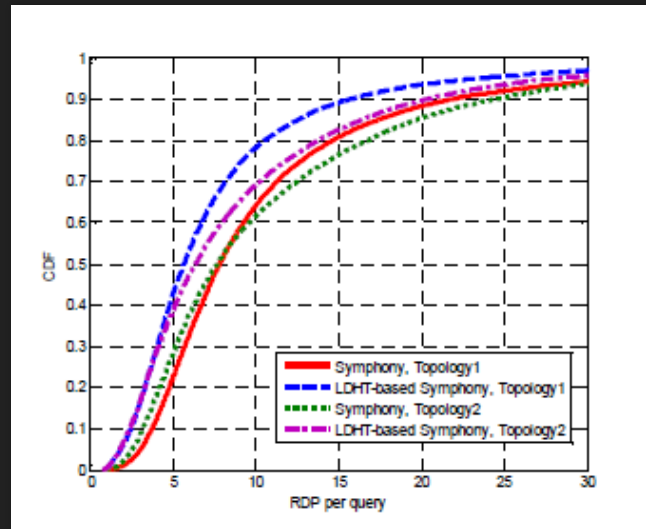
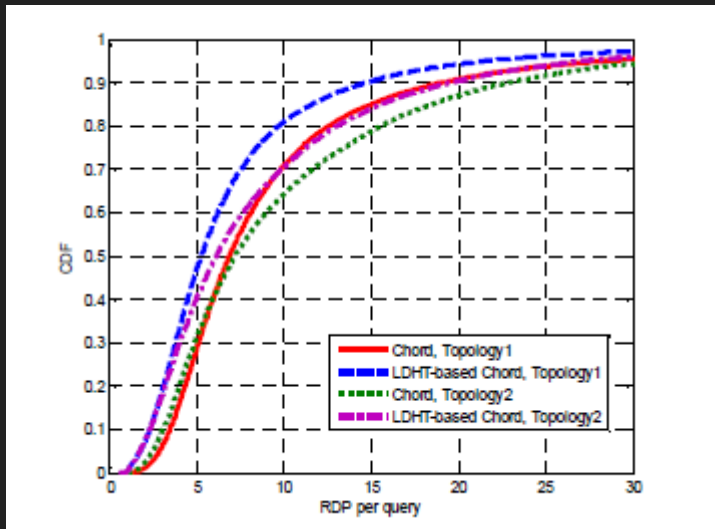
Plots for path length per query

Path length per query(contd.)

TABLE I AVERAGE PATH LENGTH (MS)

	Chord		Symphony		Kademlia	
	Orig	LDHT	Orig	LDHT	Orig	LDHT
TP1	525	407	595	443	449	383
TP2	1024	869	1083	897	884	853

Relative Delay Penalty



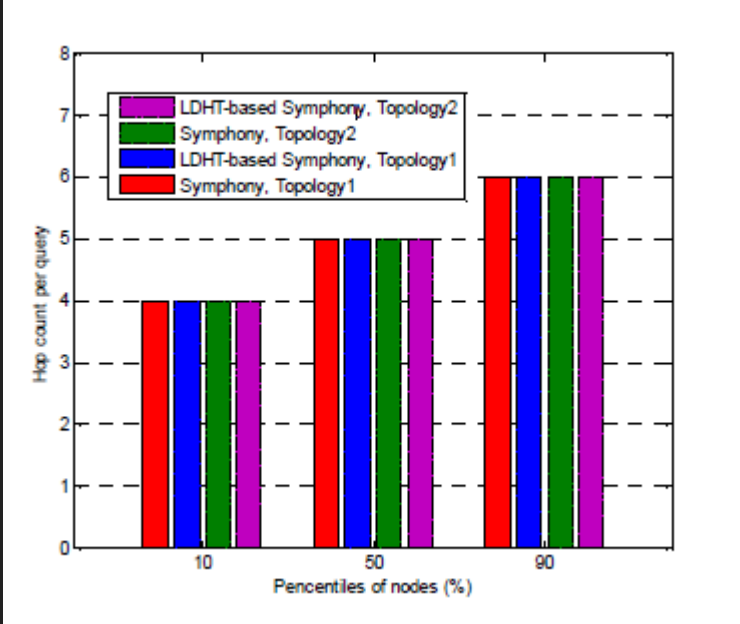
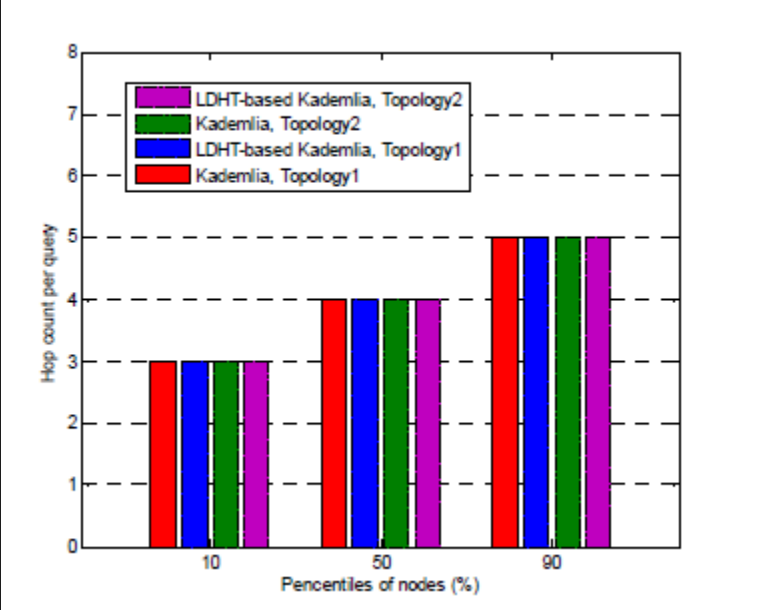
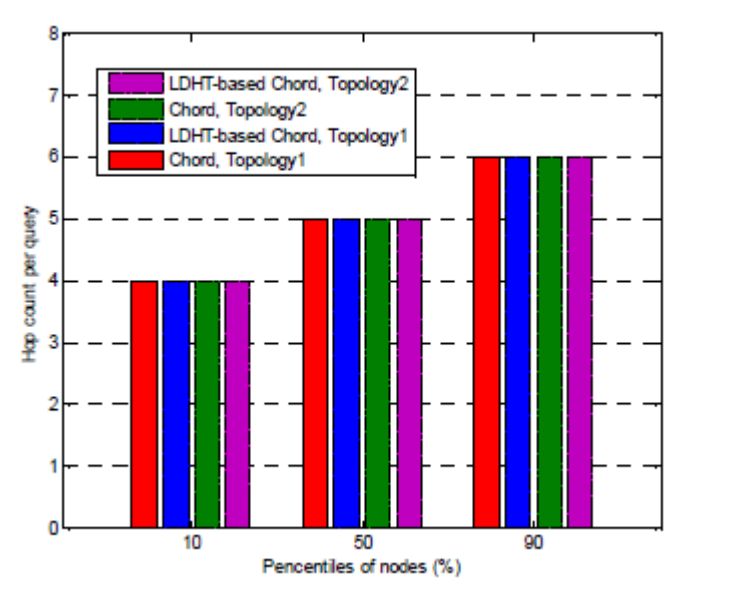
Relative Delay Penalty per query

Relative Delay Penalty(contd.)

TABLE II AVERAGE RDP

	Chord		Symphony		Kademlia	
	Orig	LDHT	Orig	LDHT	Orig	LDHT
TP1	10.71	8.22	12.19	8.64	9.11	7.50
TP2	14.24	13.16	15.48	12.80	13.54	10.82

Hop Count per query



Hop Count per query

Conclusion and Future Scope

- LDHT achieves better performance in terms of end-to-end latency
- Do not introduce additional hops to achieve the goal
- Proximity among ASes to improve performance of LDHT

References

- Weiyu Wu ; Yang Chen ; Xinyi Zhang ; Xiaohm Shi ; Lin Cong ; Beixing Deng ; Xing Li," LDHT: Locality-aware Distributed Hash Tables", International Conference on Information Networking, 2008.
- (2007) The GT-ITM homepage. [Online]. Available: <http://www.cc.gatech.edu/projects/gitim/>

Questions?