



An Adaptable Service Overlay for Wide Area Network Service Discovery

**A. Brown, M. Kolberg – U. Stirling
J. Buford – Avaya Labs Research**

**1st IEEE Workshop on Enabling the Future Service-Oriented Internet
at IEEE Globecom 2007**



Topics

- “Service Overlay”
- Problem Statement
- Related Work
- Chameleon Design
- Simulation Analysis
- Open Issues

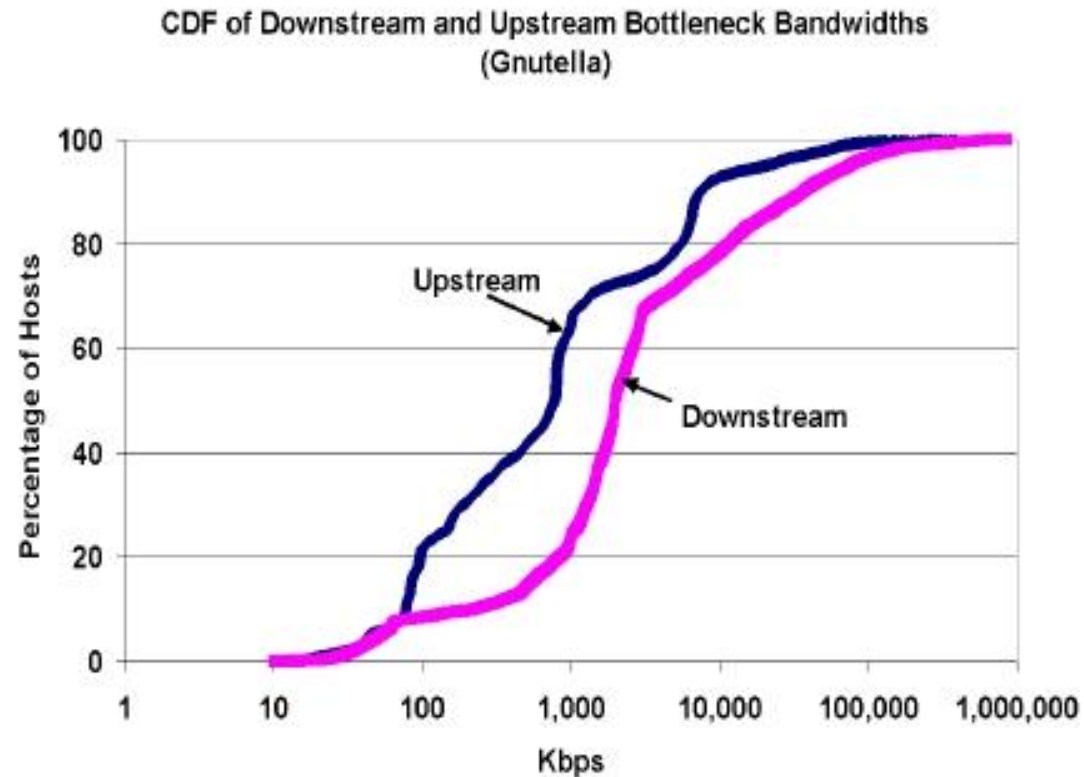
Service Overlay

- “P2P Service Overlay”
 - X. Gu, K. Nahrstedt, and B.Yu. SpiderNet: An Integrated Peer-to-Peer Service Composition Framework. In *Proc. of the 13th IEEE Intl Symposium on High Performance Distributed Computing (HPDC'04)* 110-119.
- Peers provide not only media content but also application service components
 - Requires service description and service discovery mechanism
- Example types of services
 - Streaming infrastructure: relays, mixers, transcoders
 - Communication services: ALM, call features, IMP
 - End user: search, group membership, publish/subscribe, ...

Problem Statement

- Service discovery in structured peer-to-peer overlays
 - Using a P2P overlay offers the potential for highly-scalable, wide-area service discovery
 - Known techniques for indexing service descriptions (e.g., INS/Twine)
- Issues
 - How to improve latency performance (versus multi-hop)
 - (see various studies of using multi-hop Chord as DNS)
 - How to support heterogeneous peer population
 - Chameleon variable-hop overlay
 - How to deal with hot-spots
 - Compound keys which produce uniform distribution

Heterogeneity in P2P Overlays

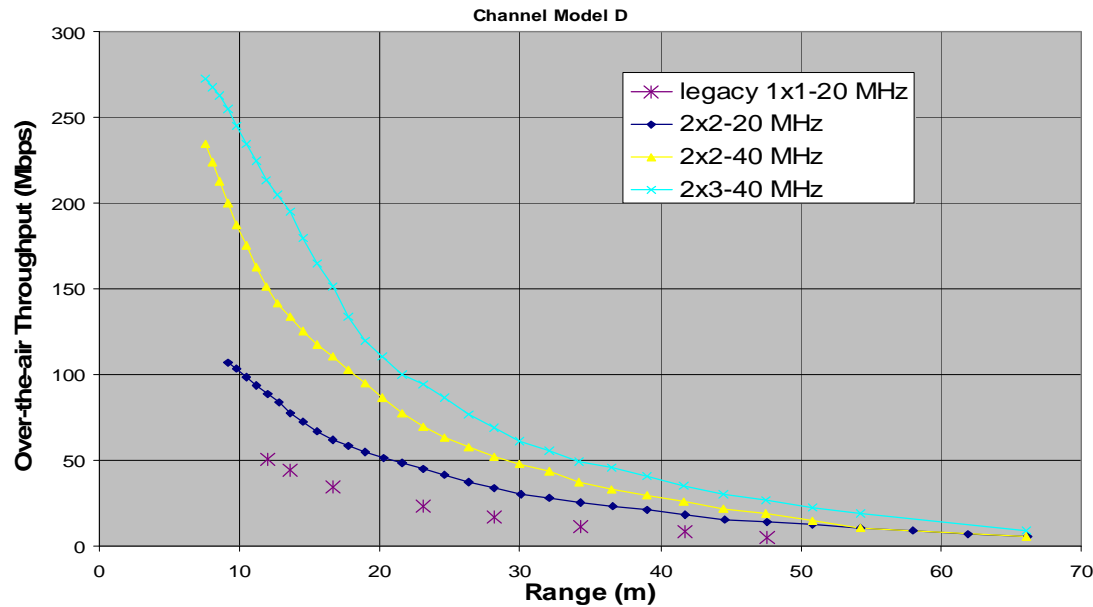


Source:


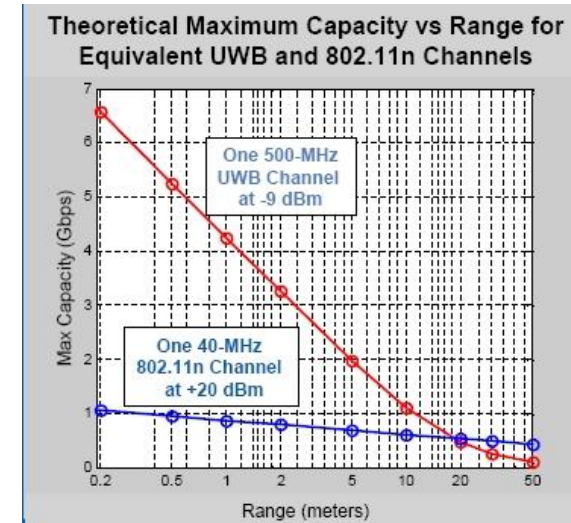
S. Saroiu, P. K. Gummadi, S. D. Gribble. A Measurement Study of Peer-to-Peer File Sharing Systems. In Proceedings of Multimedia Computing and Networking 2002 (MMCN'02) (San Jose, CA, Jan. 2002)

Future Overlays with Wireless Peer Populations

802.11n (Source: R. Stacey, Intel)



UWB vs 802.11n (Source: D. Leeper, Intel)



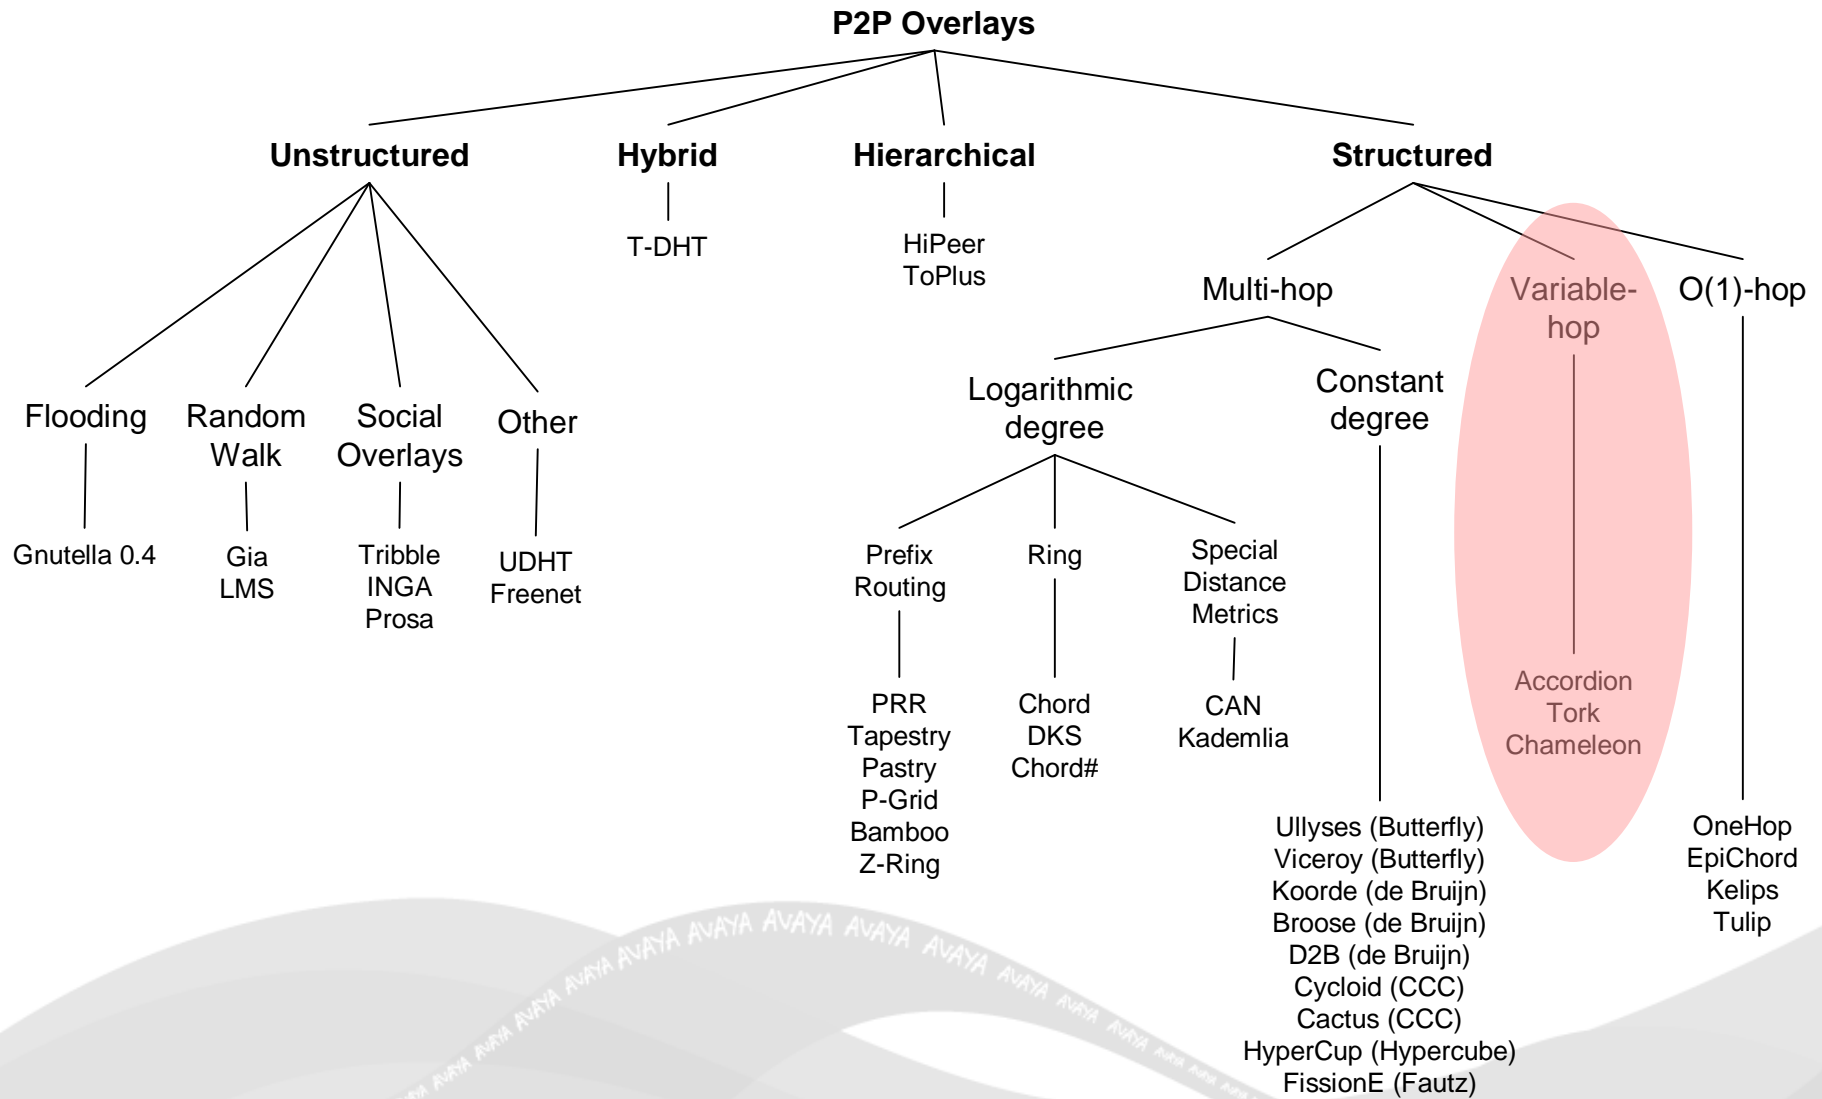
Heterogeneity vs Adaptivity

- Heterogeneity
 - Different devices in the overlay have different capacities for CPU, storage, access BW
 - These capacities for each device are relatively stable over time
- => There are few transitions between H-, M-, L- states.
- Adaptivity
 - In a given access network, access BW varies
 - Distance from access point / base station
 - Interference
 - => could be frequent transitions between H-, M-, and L- states
 - Devices support multiple network interfaces
 - Devices roam and encounter different access BW
 - => If multi-homed, transitions could be masked by high bandwidth interface

The Variable Hop Overlay Approach

- What?
 - Each peer in the overlay has bandwidth budget that is allocated to routing table maintenance
 - Higher budget means more routing table updates are exchanged, leading to higher routing table accuracy
 - Each peer manages its budget independently
- Why?
 - Devices have heterogeneous resources and access network capacity
 - Latency matters
 - Many nodes have the capacity for more routing table accuracy
 - Doesn't penalize the low bandwidth nodes

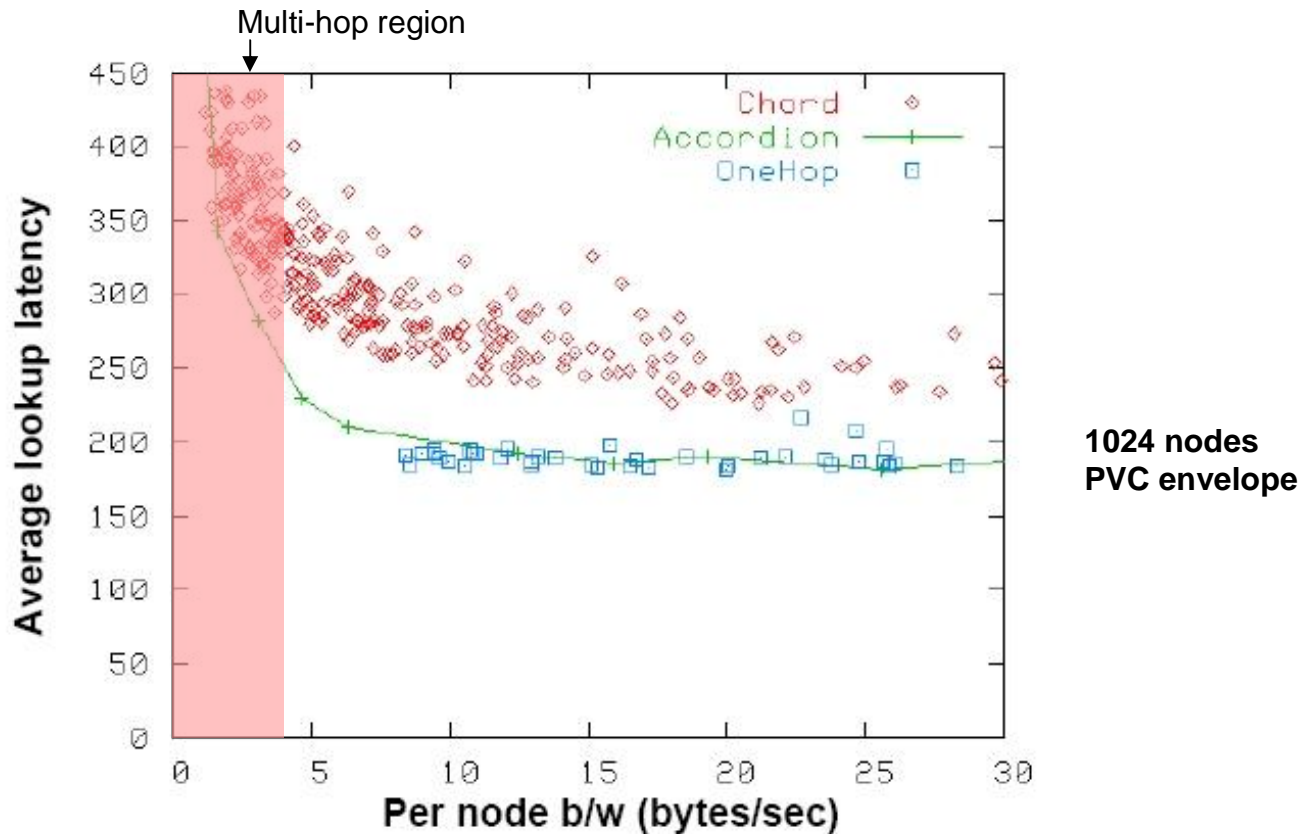
P2P Overlay Taxonomy



Related Work: Accordion

- Routing table density decreases with distance from the node
- Lookups
 - Recursive, since each node has more knowledge about local region of address space
 - Each node forwards lookups with parallelism determined by its available bandwidth
- Routing table maintenance
 - Learns routing table entries from incoming lookups and from responses to lookups
 - Estimates node liveness based on the node's lifetime and when it last heard from the node
 - Excess budget used to explore regions of routing table
 - Exploration lookups biased towards nodes with higher bandwidth budgets
- Bandwidth budget
 - Has an average budget and a burst budget
 - Each node counts (in time window) only outgoing requests and responses for its budget


Related Work: Accordion



Jinyang Li, Jeremy Stribling, Robert Morris and M. Frans Kaashoek. Bandwidth-efficient management of DHT routing tables. In the Proceedings of the 2nd USENIX Symposium on Networked Systems Design and Implementation (NSDI '05), Boston, MA, 2005.

Chameleon Design


- Variable hop overlay aimed at high performance range
- Heterogeneous peers
 - Low bandwidth (L-peers) and high bandwidth (H-peers) peers
 - Bandwidth capacity can change over time
- Combine two $O(1)$ -hop overlay design techniques
 - L-peers: Opportunistic maintenance, achieves routing table accuracy ~ 50%
 - H-peers: Active maintenance, achieves routing table accuracy ~90%



J. Xu, A. Kumar, and X. Yu. On the Fundamental Tradeoffs between Routing Table Size and Network Diameter in Peer-to-Peer Networks, IEEE JSAC, Nov. 2003.

Lowest latency region.
E.g., OneHop, EpiChord, D1HT, Kelips, Tulip

L-Peers




- Based on EpiChord design
- RT organized in slices where density increases close to peer
- RT updates carried in lookup responses
 - Can operate ~ 50% RT accuracy
- If lookup traffic is too low, does active probing

H-Peers

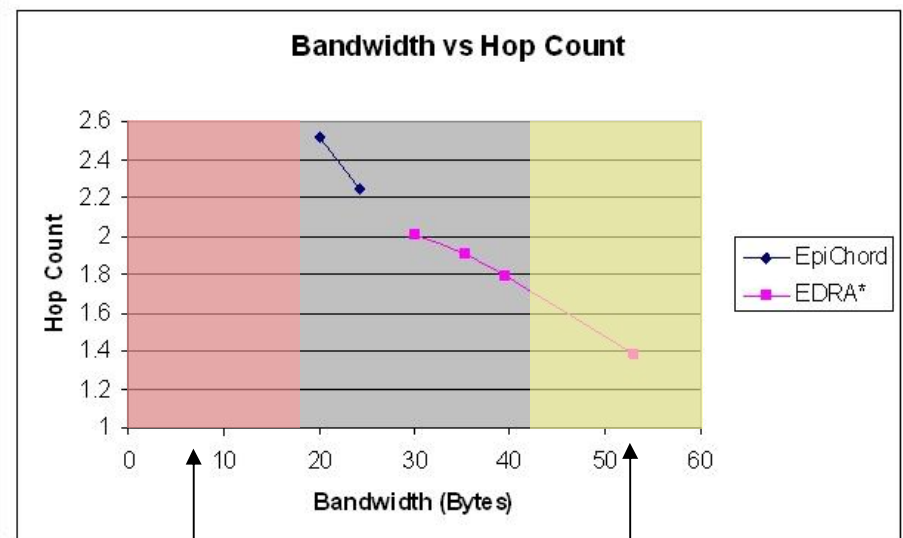
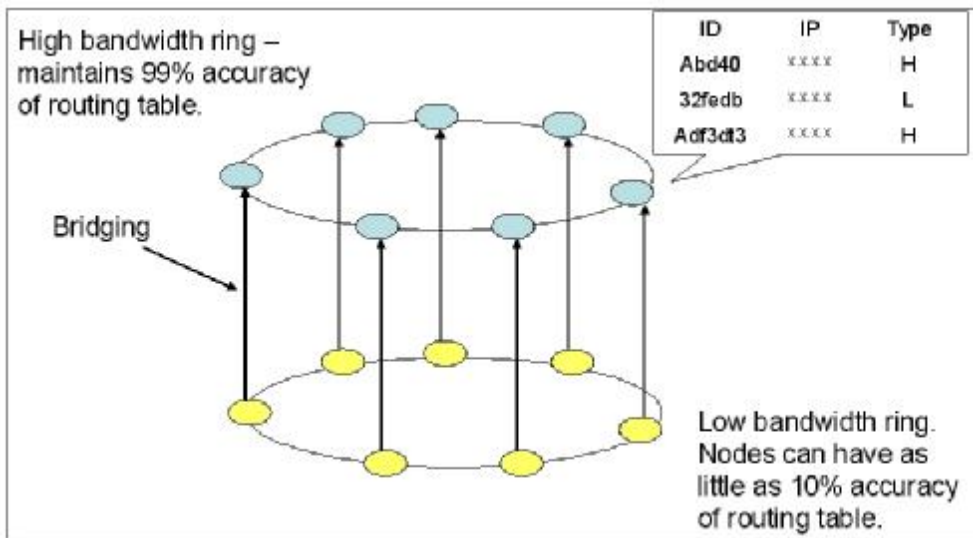
- Modeled after modified D1HT/EDRA algorithm (EDRA*)
- Peers maintain logarithmic-degree connections
 - EDRA protocol to propagate join/leave events throughout overlay
- Can achieve up 99% RT accuracy

Logarithmic Degree Active Maintenance Event Propagation in D1HT/EDRA



Chameleon

- H-Peers maintain and propagate overlay membership state about all L- and H-peers
 - Propagation is within the H-peer set
- L-Peers bridge to another H-peer



Bridging for maintenance traffic only
Lookups use normal routing table

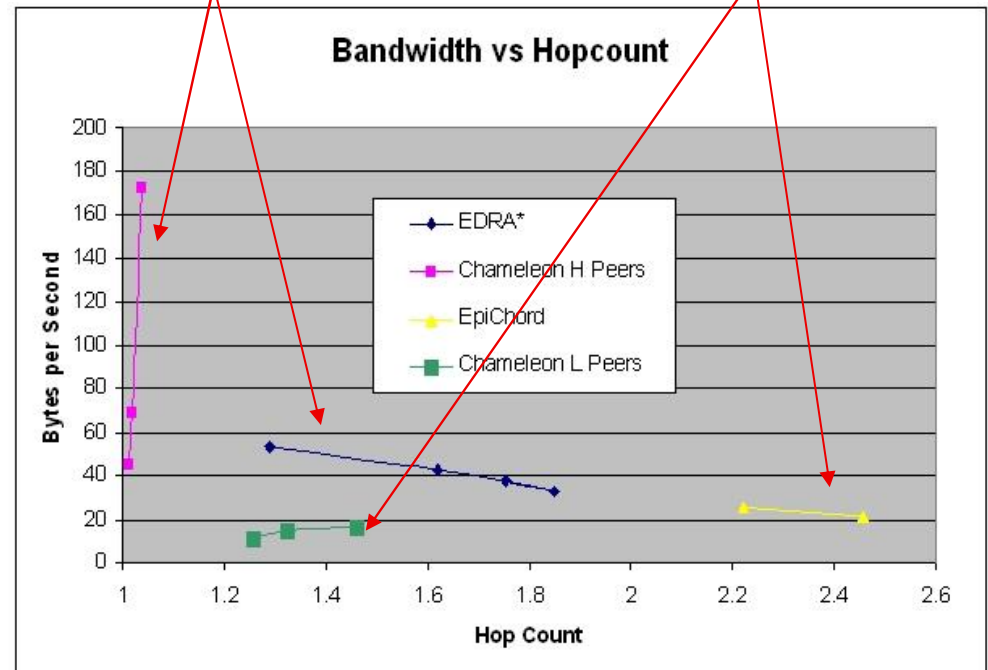
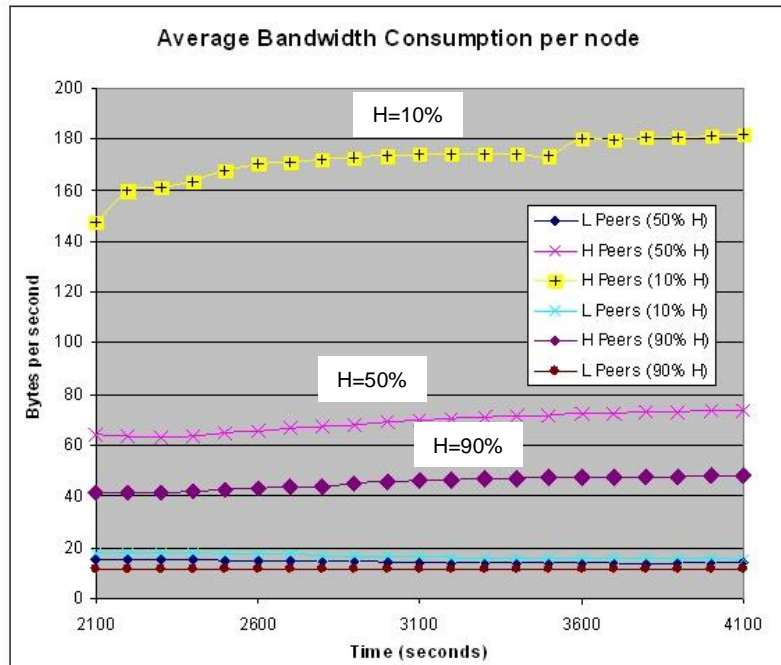
Multi-hop region

O(1)-hop region

Simulation Analysis

- Evaluate lookup performance over range of L:H populations
 - Stability
 - Performance compared to EpiChord and EDRA*
- Service overlay simulation
 - Use synthetic service discovery lookup workload based on Yellow Page services
- Simulation parameters
 - SSFNet simulator
 - 1200 nodes on 10K host network
 - Median node lifetime 60 mins, heavy-tailed distribution

Chameleon

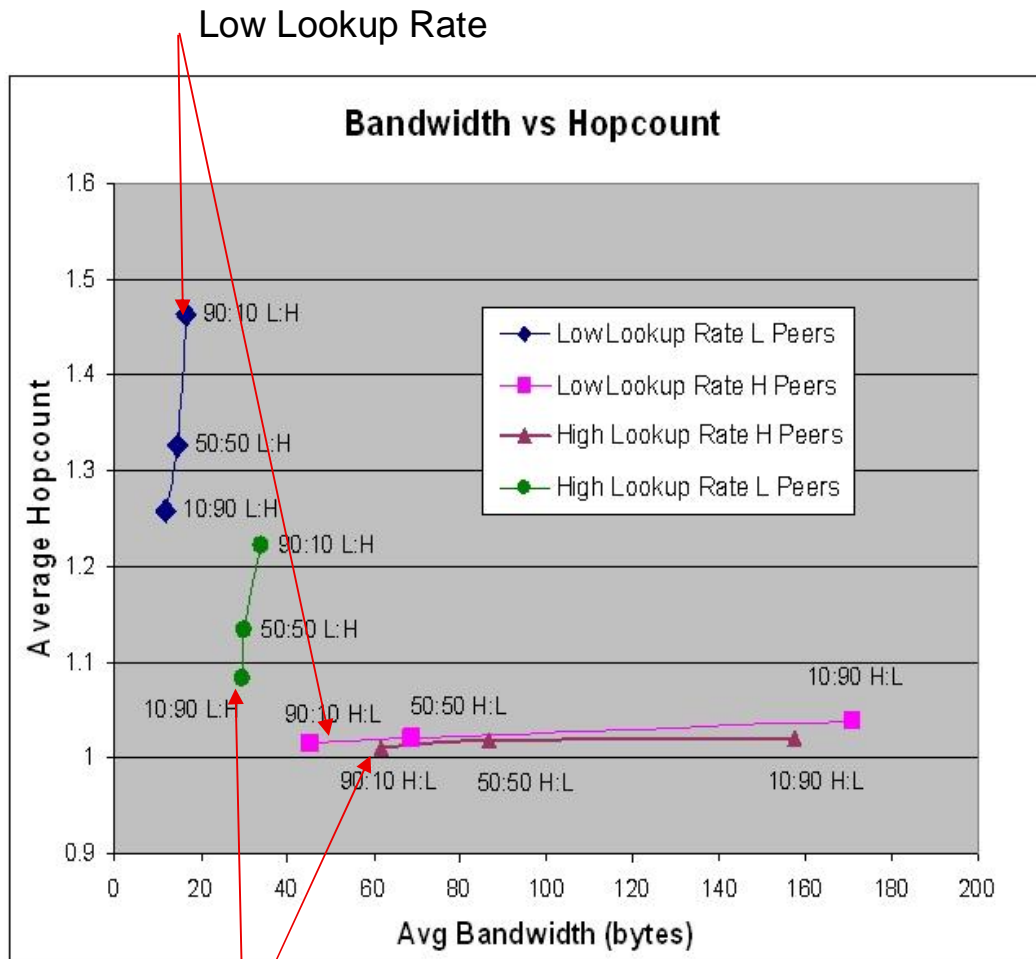


- L-peers benefit from larger proportion of H-peers
 - Increased RT accuracy for all peers
- As H-peer level > 50% then load on H-peers grows
- H < 10% may not be feasible

Simulation Analysis: Service Discovery Workload

- Service descriptions drawn from Yellow Pages
 - 2500 service categories
 - Compound keys formed from service meta-data
 - Domain, Service Discovery Protocol, Latt-Long Location, Operator
 - Indexing shown in prior work to produce uniform key distribution in overlay
- Two lookup rates tested:
 - Low: lookups per node 1 every 10 min
 - High: lookups per node 1 every 20 sec

Bandwidth vs Hopcount



- In the range of L:H from 10:90 to 90:10,
- Higher lookup load increases performance for L peers because routing table updates are carried in lookups, increasing RT accuracy

High Lookup Rate


Open Issues

- Lookup load per H-Peer vs L-Peer
- Scalability for overlays with $N > 10^6$
- Evaluation of Chameleon vs Tork
 - Chameleon creates two separate populations
 - Tork uses both techniques in different ratios as bandwidth changes


AVAYA *labs*



Distribution and Density



Certain regions of the network may be heavily M or H



What is the performance over the range?

