
Ekta: An Efficient DHT Substrate for Distributed Applications in Mobile Ad Hoc Networks

Himabindu Pucha, Saumitra Das, Y. Charlie Hu

Distributed Systems and Networking Lab,
School of ECE, Purdue University

Mobile ad hoc networks (MANETs)

- Wireless networks in which wireless hosts act as forwarding nodes as well as end systems
 - No base station or routing infrastructure
 - Network topology changes frequently and unpredictably
 - Challenge lies in routing packets with changing topology while minimizing overhead
 - Specialized routing protocols: DSR, AODV
-

Structured peer-to-peer routing protocols in the Internet

- Operate in an overlay p2p network in which nodes act as clients as well as servers
- Rely on underlying Internet infrastructure to route packets between overlay hops
- Implement a DHT in a scalable, robust manner
- Challenge lies in routing packets in a network with changing membership while limiting state at each node
 - Specialized protocols: Pastry, Tapestry, Chord, CAN

Motivation

- DHTs provide a useful platform for building scalable and robust distributed applications in the Internet
- DHTs can potentially provide an efficient way to construct distributed applications and services in MANETs
 - Applications such as file sharing, resource discovery could benefit from the insert/lookup convergence
- **Main Challenge**
 - **Provide an efficient DHT abstraction in a MANET and demonstrate its usability**

Outline

- How to support an efficient DHT abstraction in an ad hoc environment?
 - Can off-the-shelf protocols be used?
 - If not, what is an efficient architecture to provide the DHT abstraction?
- How should the DHT abstraction be used in an ad hoc environment ?
 - Can a MANET application benefit from the DHT?

Implementing DHT: Layered Approach

- Layer **Pastry** (structured p2p protocol) on top of **DSR** (MANET routing protocol)
- Pastry operates in the application layer similar to the Internet
- DSR used as underlying routing protocol

Pastry : Features

- Nodes have unique Id, messages have keys
 - Typically 128 bits long
- Primitive: **Route(msg, key)**
 - Delivers **msg** to the currently alive node whose Id is numerically closest to **key**
- Scalable, efficient
 - Per node routing table contains $O(\log(N))$ entries
 - Routes in $O(\log(N))$ steps
- Fault tolerant
 - Self-fixes routing tables when nodes are added, deleted or fail

Pastry : Routing table (# 65a1fcx)

Row 0

Row 1

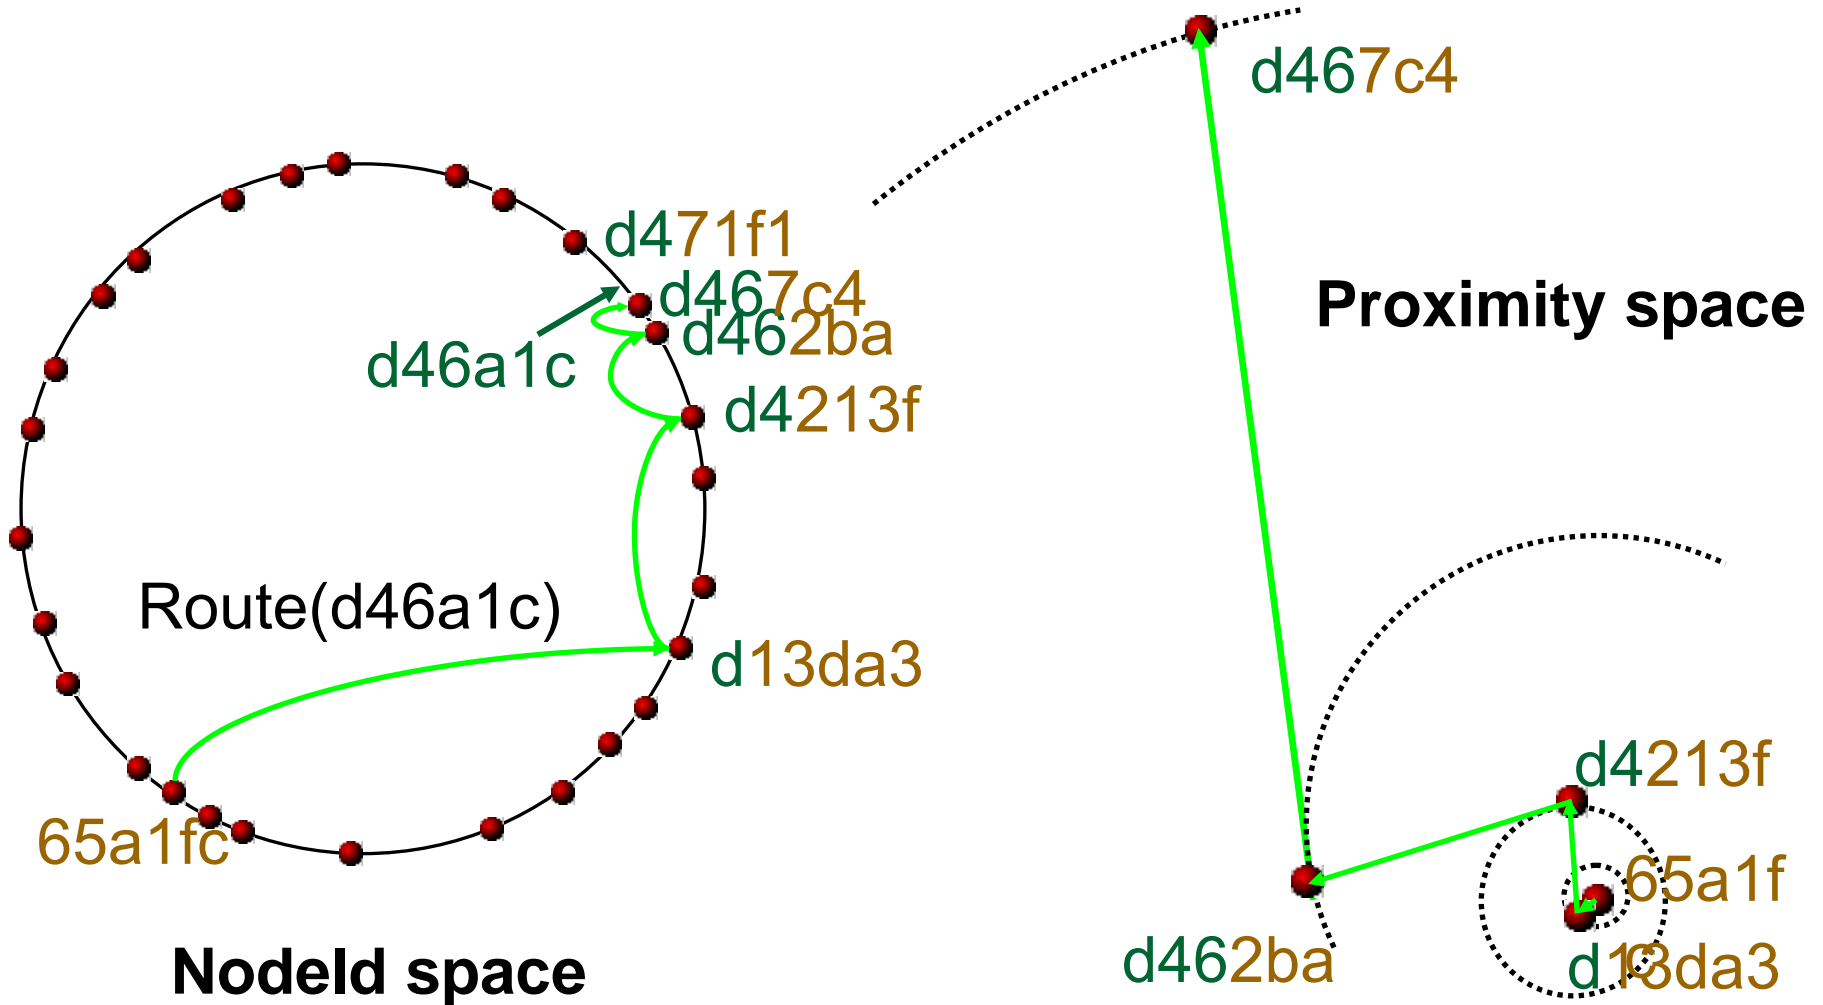
Row 2

Row 3

<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>		<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>
<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>		<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>
<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>		<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>
<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>		<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>
<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>5</i>		<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>
<i>a</i>		<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>0</i>		<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>


$\log_{16} N$
rows

Pastry: Routing Example



DSR

- Reactive routing protocol based on source routing
- Operation: route discovery and maintenance
- Route caching: path and link cache



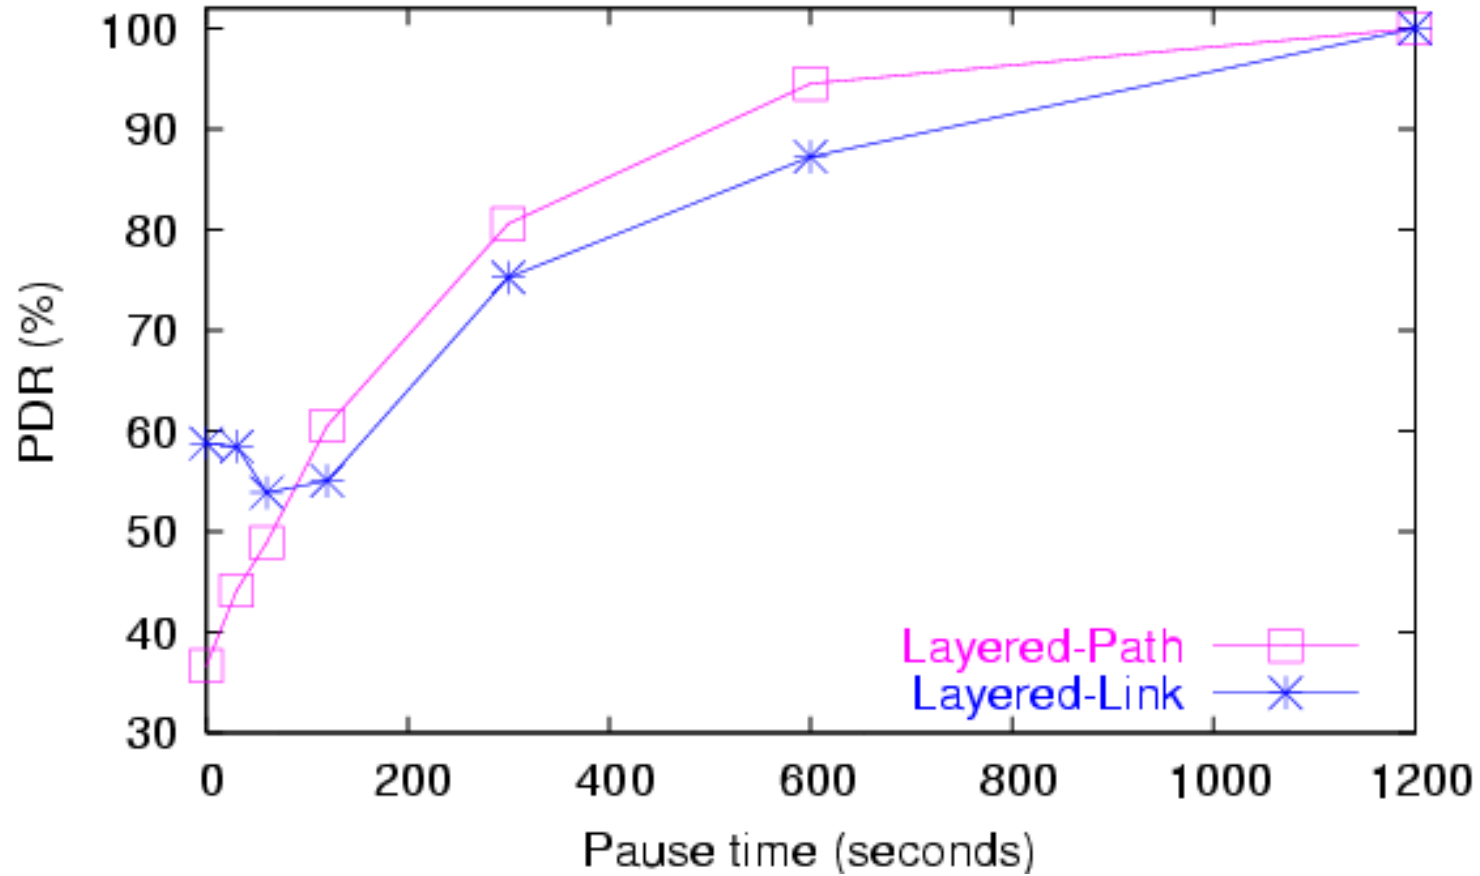
Layered approach

- Direct layering is not practical
- Pastry uses 'ping' to measure delay for proximity
 - Proximity of Pastry routing tables uses hop count instead of delay from 'ping'
 - DSR exports API to answer proximity probes from Pastry using its route cache
- Expanding ring search for bootstrap node

Simulation setup

- Pastry implemented in ns-2 on top of DSR
- Parameters: 50 nodes, 1500mx300m, 2Mbps, 250m, 1-19m/s
- Traffic: 40 sources, 3 pkts/sec, random keys generated
- Metrics
 - Packet delivery ratio (PDR): ratio of successfully and correctly delivered packets to packets sent
 - Overhead
 - Delay

Layered approach performance



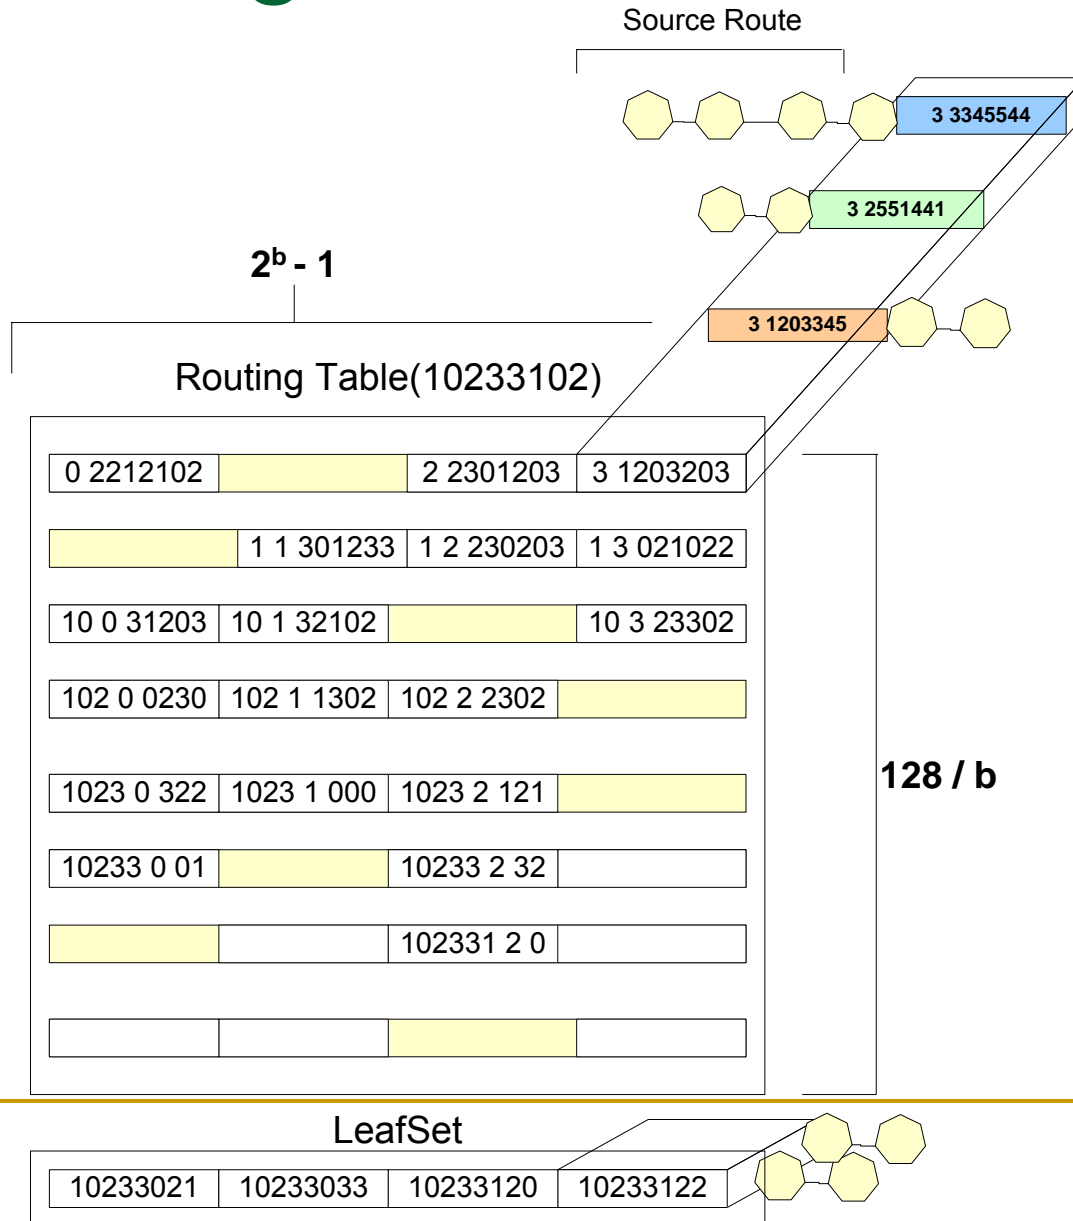
Problems with layered approach

- High overhead
 - Periodic maintenance, proximity probing
- Choice of next logical hop independent of DSR
- Stale proximity information with Pastry
- Mismatch between routing state of Pastry and DSR

Implementing DHT: Integrated Approach

- **Integrates** Pastry and DSR
- A unified DHT substrate at the network layer
- Referred to as *Ekta* (unity)

Ekta: Routing




Ekta: Join and leave

- Join
 - Flood “JOIN REQUEST”
 - Potential leafset members send “JOIN ACK”
 - Node closest sends “JOIN COMPLETE”
- Graceful leave
 - Flood “LEAVE”
 - Leafset members send “LEAVE ACK”
 - Exchange leafsets
- Node failure
 - Reactive failure handling
 - Node sends “proxy leave”

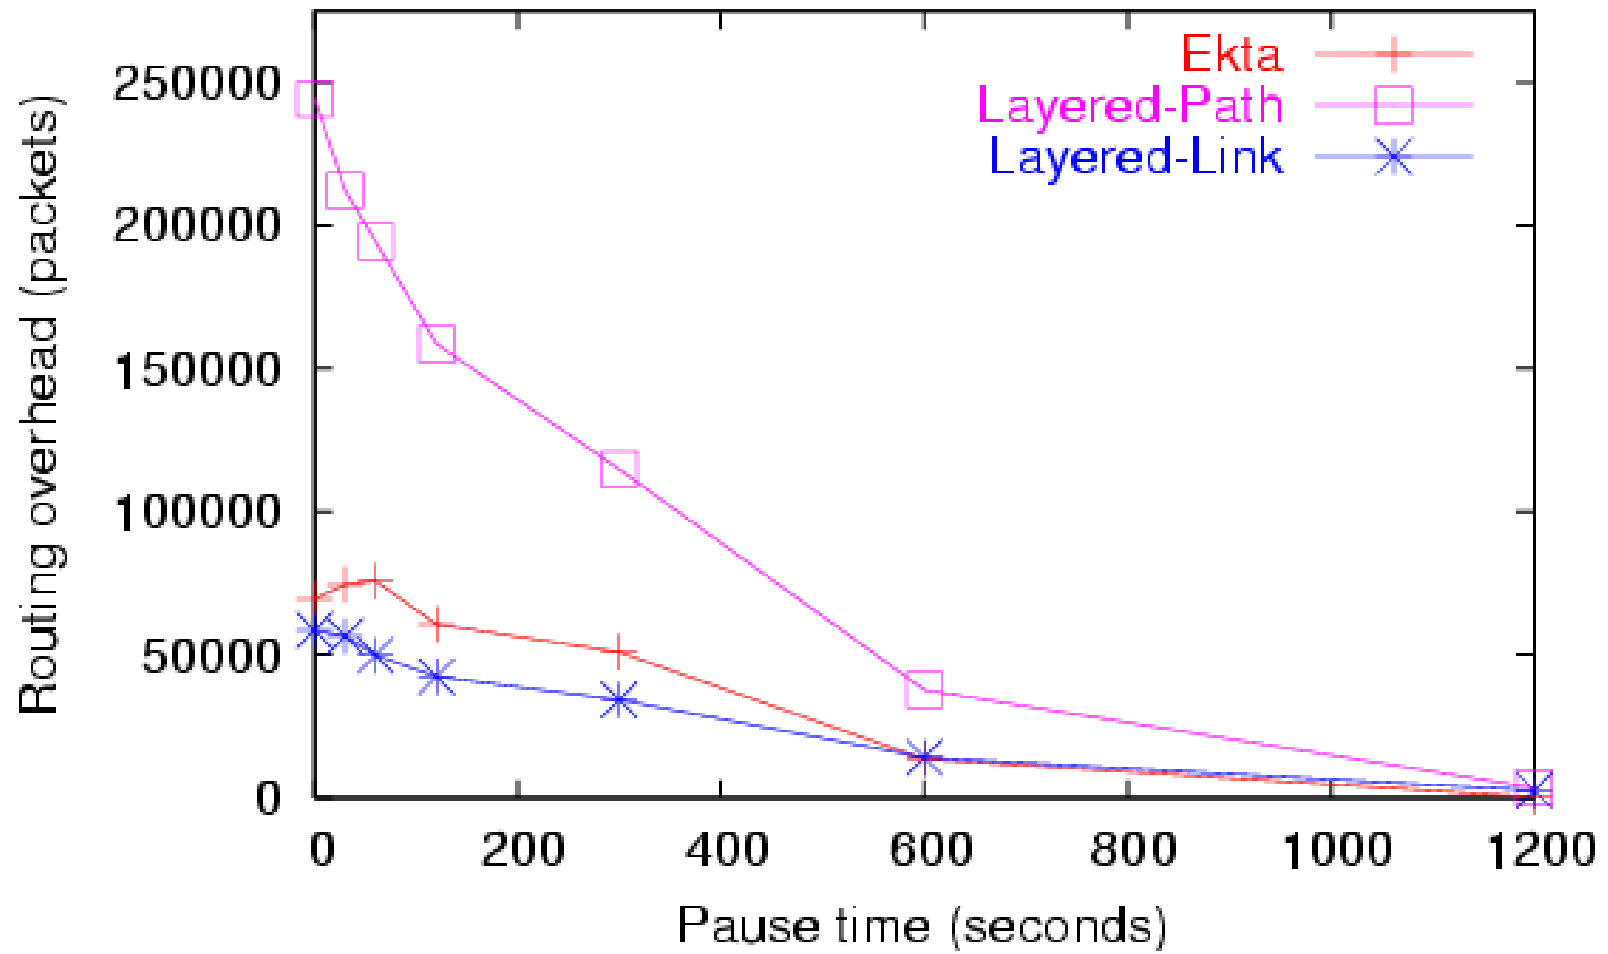
Ekta: Optimizations

- Prefix based route requests
- Routes updated using snooping and overhearing


PDR



Overhead



Delay



Summary

- Integrated approach is an efficient architecture for implementing a DHT in MANETs
- Ekta is superior
 - No proximity probing and periodic maintenance
 - Better coordination between routes available in routing table and choice of logical hops
 - Prefix route requests
 - Fresher proximity information from snooping and overhearing

Can a MANET application benefit from the DHT?

Resource discovery using Ekta

Problem

- Nodes in a MANET possess heterogeneous capabilities and resources
- Cooperative resource sharing is useful in MANETs
 - Requires **resource discovery**
- Two schemes
 - Ekta-RD
 - DSR-RD

Ekta-RD vs. DSR-RD

- Simulations in ns-2 of both protocols
- Number of unique resources = number of nodes
- Each resource replicated on average on 10% of nodes
- Traffic: Poisson arrival of resource requests at each node, each request chooses random resource, varying λ
- Metrics
 - Success ratio: resource requests successfully satisfied
 - Overhead: control overhead for routing

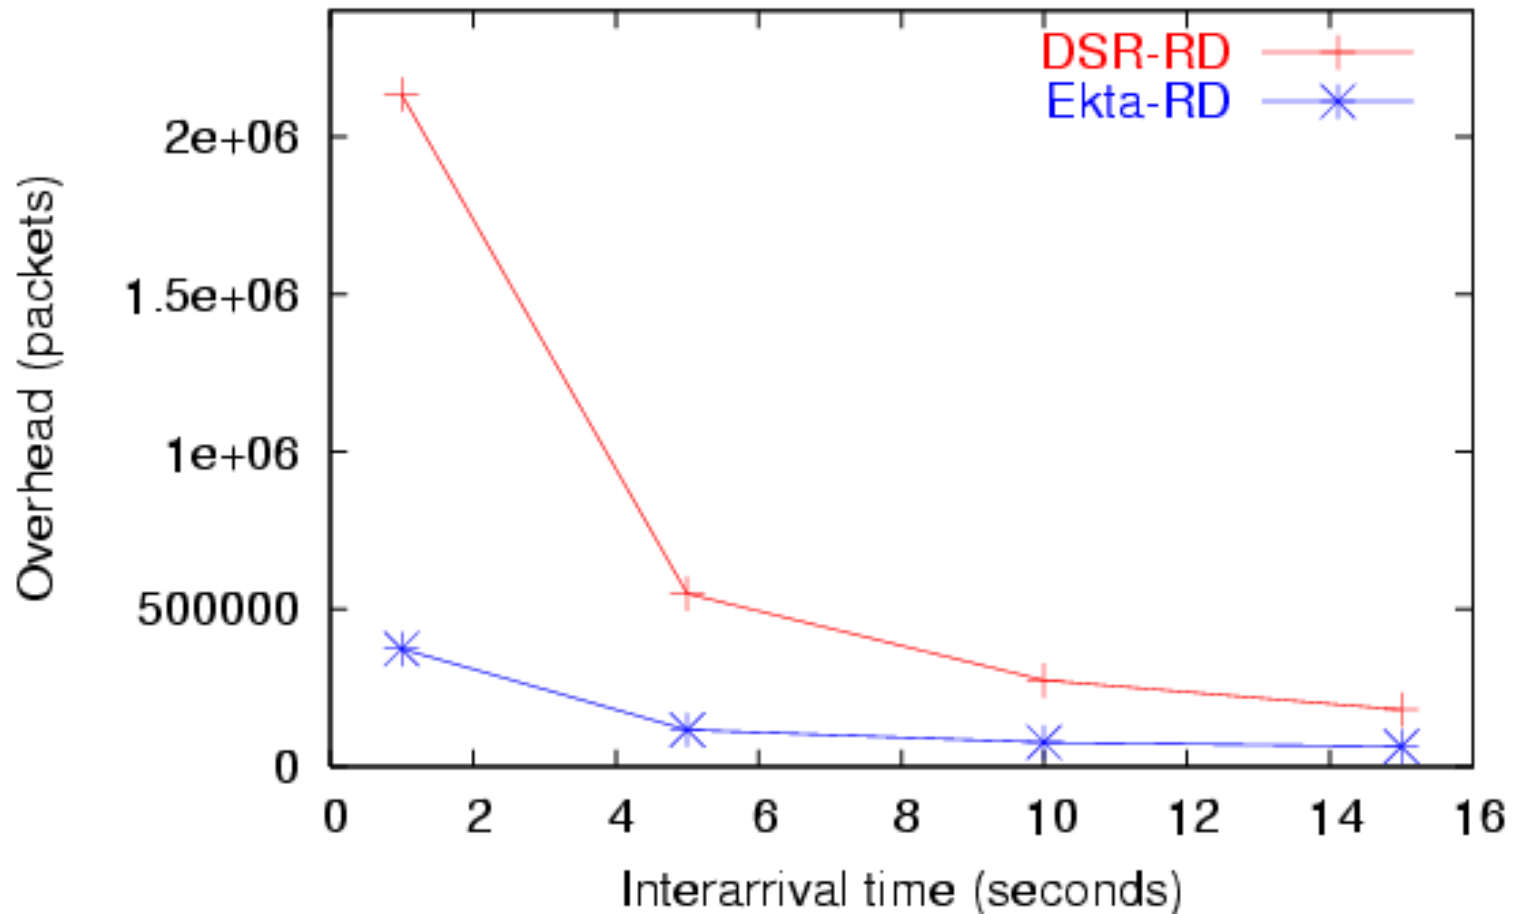
Overhead analysis

- N = network size
- λ = average num of resource requests/node
- q = average degree of replication of any resource
- P = average hops between 2 nodes
- P_b = probability a route is not cached or cached but stale

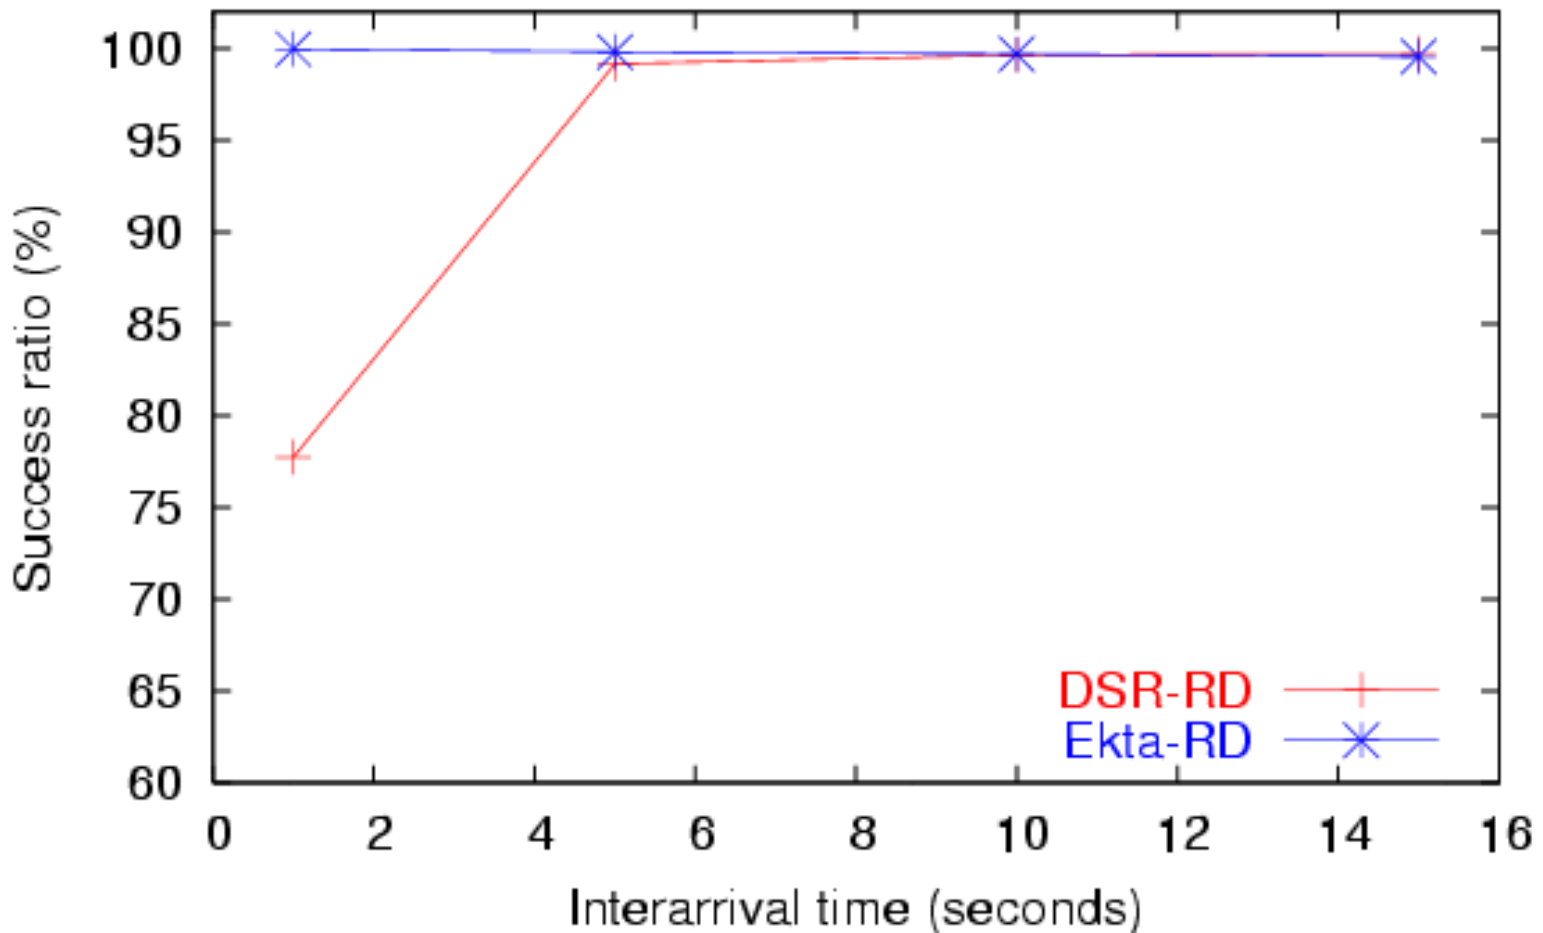
- Overhead (DSR-RD) = $\lambda \cdot N^2 + \lambda \cdot q \cdot N^2 \cdot P$
 - Independent of mobility
 - Grows as $O(N^2)$
 - Decreases with increasing λ

- Overhead (Ekta-RD) = $N (\lambda \cdot \log_2^b N \cdot P + \lambda \cdot P)$
 $+ N^2 \cdot P_b \cdot \lambda (\log_2^b N + 1)$
 - Increases with mobility
 - Grows as $O(N \cdot \log_2^b N)$
 - Decreases with increasing λ

Overhead: Varying λ

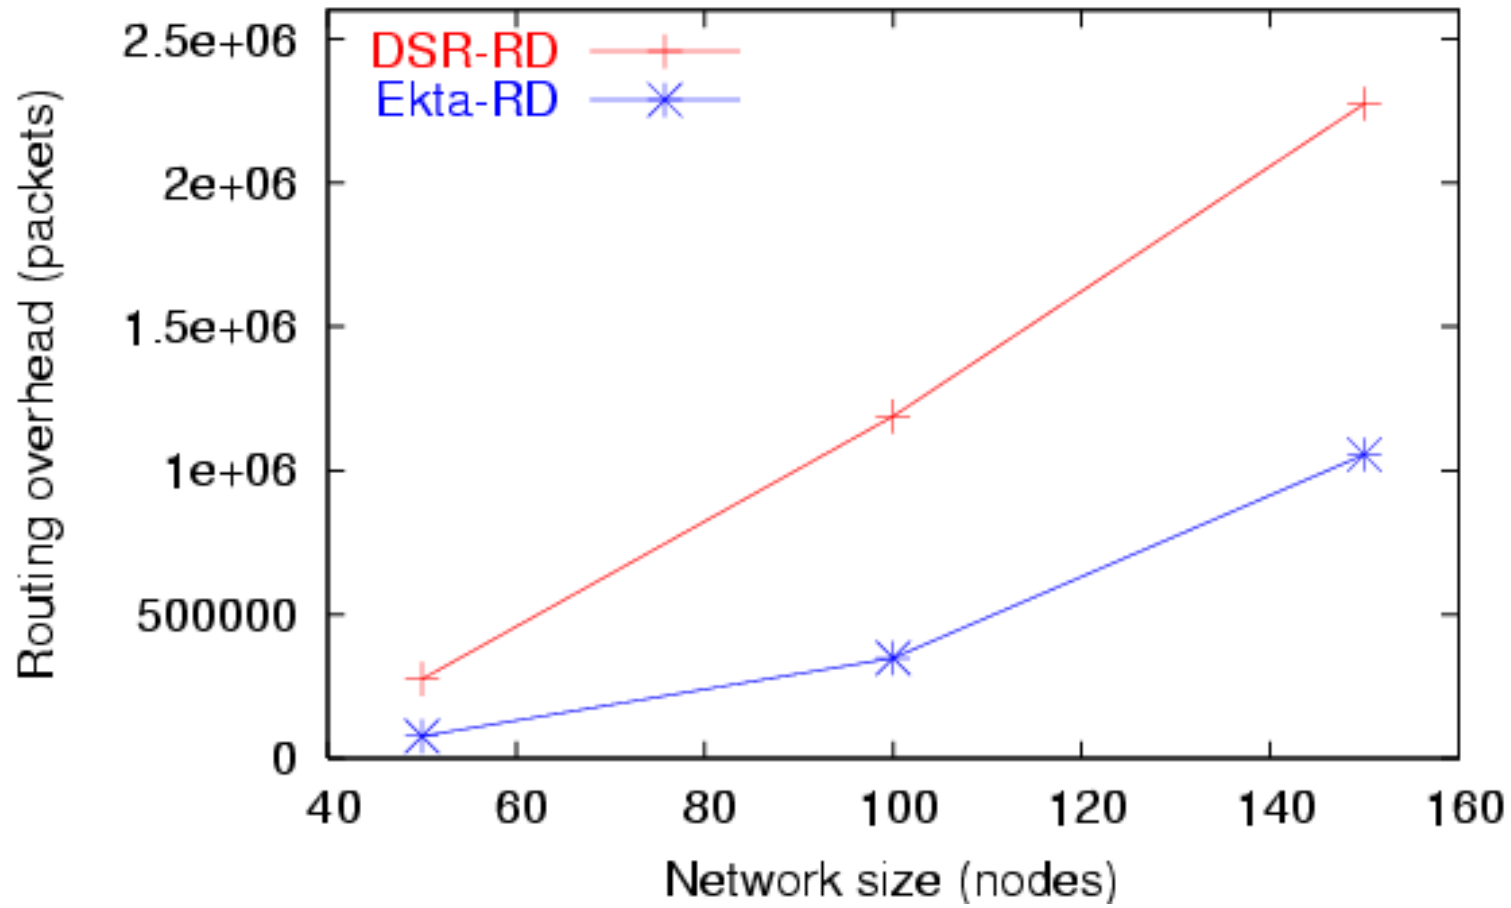


Success ratio: Varying λ



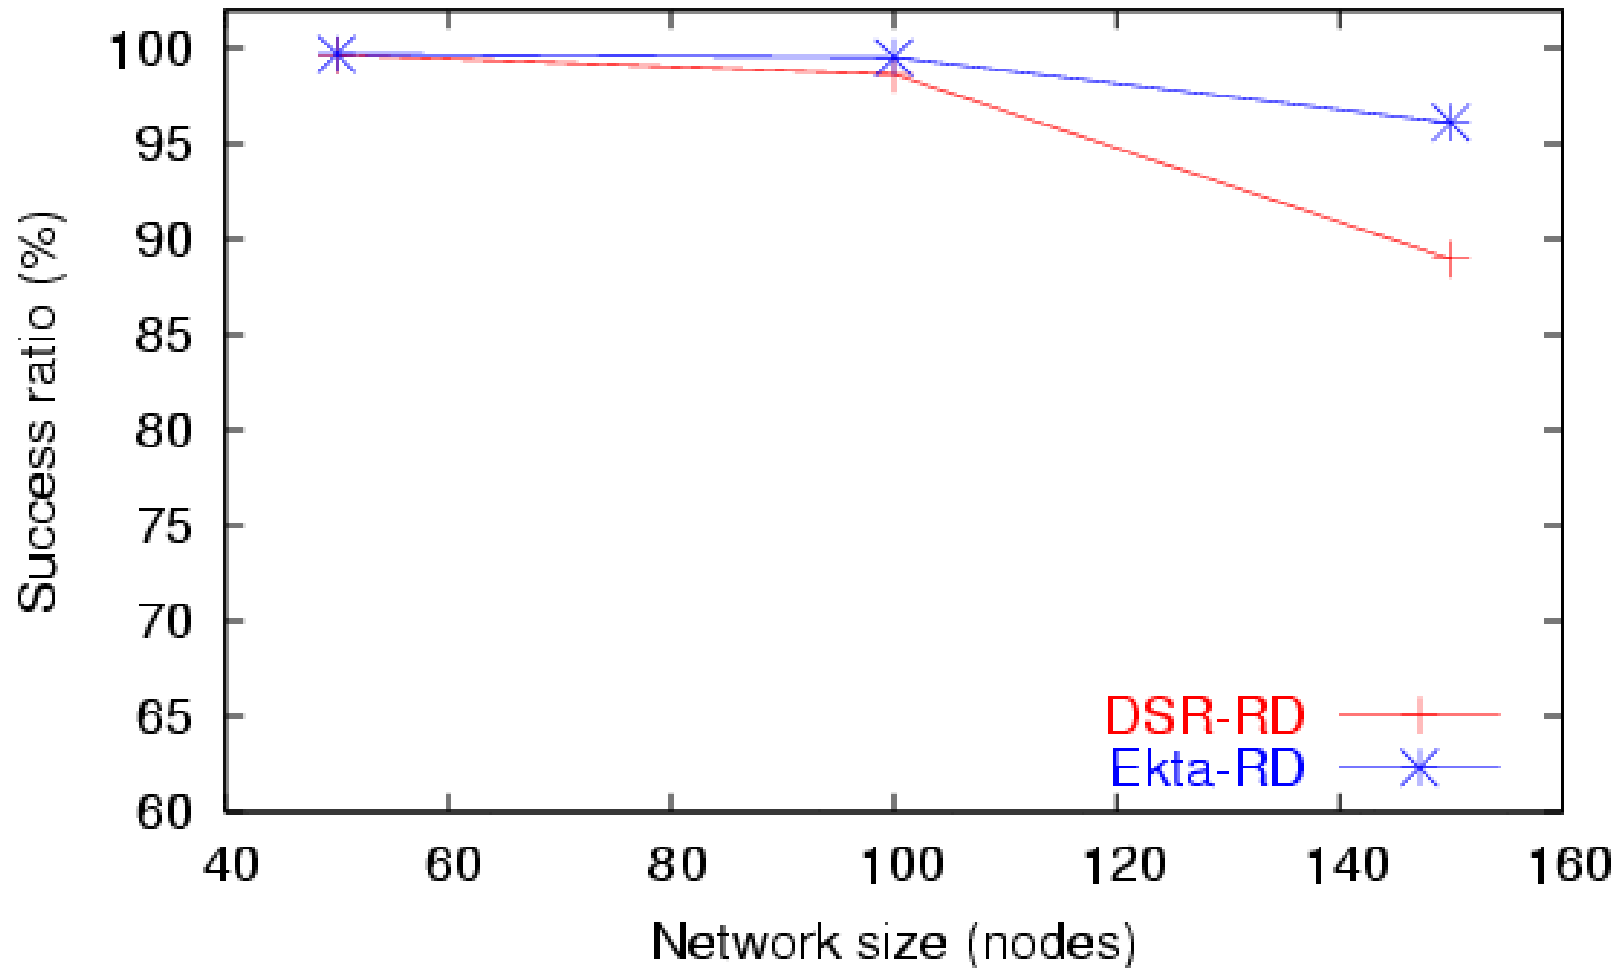
Overhead: Varying N

$1/\lambda = 5$ seconds



Success ratio: Varying N

$1/\lambda = 5$ seconds



Conclusions

- Integrated approach is an efficient architecture for implementing a DHT in MANETs
- MANET applications can benefit from DHTs as demonstrated by the resource discovery application
- Ekta can potentially be used as an efficient substrate for other applications

Q & A
