

Embracing the Peer Next Door: Proximity in Kademlia



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Sebastian Kaune, Tobias Lauinger,
Aleksandra Kovacevic, and Konstantin Pussep
P2P '08 (Aachen)



KOM - Multimedia Communications Lab
Prof. Dr.-Ing. Ralf Steinmetz (Director)
Dept. of Electrical Engineering and Information Technology
Dept. of Computer Science (adjunct Professor)
TUD – Technische Universität Darmstadt
Merckstr. 25, D-64283 Darmstadt, Germany
Tel.+49 6151 166150, Fax. +49 6151 166152
www.KOM.tu-darmstadt.de

Dipl.-Inf. **Sebastian Kaune**

Sebastian.Kaune@KOM.tu-darmstadt.de
Tel.+49 6151 166138

Outline



-
- 1 Motivation**
 - 2 Related Work in Brief**
 - 3 Challenges**
 - 4 Approach**
 - 5 Methodology**
 - 6 Results**
 - 7 Summary and Future Work**

Motivation (1)

Peer-to-peer (P2P) accounts for more than 50% of today's Internet traffic

- High transit and peering costs for ISPs

Many broadband users subscribe to flat rate tariffs

- No incentive to limit bandwidth consumption:
 - Generates disproportional high inter-domain traffic for Internet Service Providers (ISP)

Locality mechanisms attempt to alleviate this problem

- Takes into account a peer's position in the underlay network

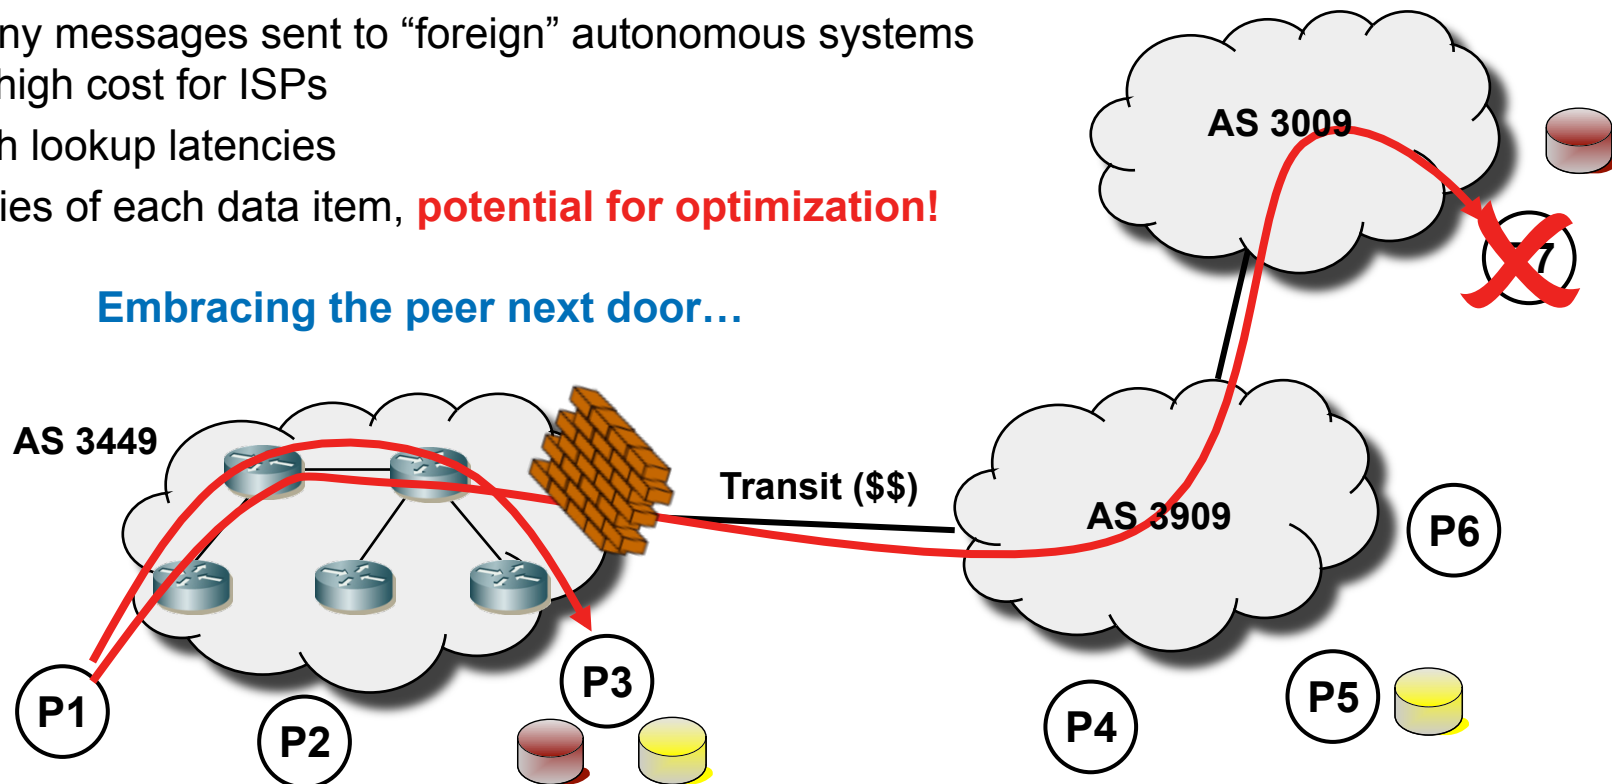
Our Goal: reduce inter-domain traffic at the DHT level

Motivation (2)

Context: Kademlia DHT

- (Geographic) distribution of data is “random”
- Lookup messages routed needlessly long paths
 - many messages sent to “foreign” autonomous systems
→ high cost for ISPs
 - high lookup latencies
- k copies of each data item, **potential for optimization!**

Embracing the peer next door...



Related Work in Brief

Proximity algorithms in DHTs focus on recursive routing

- Goal: **Optimize routing for latency reduction**

Classification into two important categories

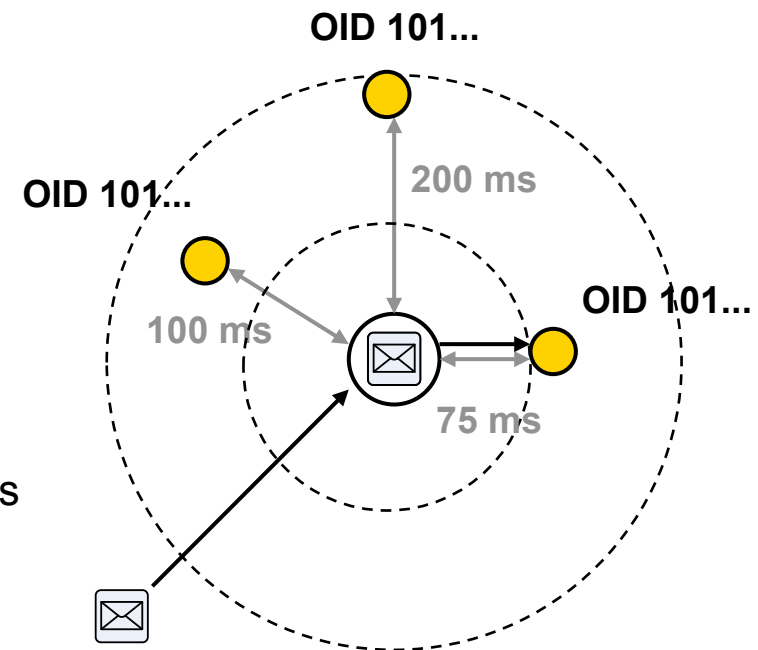
- Proximity Routing Selection (**PRS**)
 - Query **closest** node in the underlying network
- Proximity Neighbor Selection (**PNS**)
 - Routing table entries: select the **closest** one

Proximity is often measured in round-trip time

- Mostly requires active probing of routing table entries
- **Introduces additional communication costs**

Concluding challenges

- Proximity algorithms without active probing?
- Latency reduction = locality of traffic?



Challenges

Widely deployed P2P Applications implement the Kademlia DHT

- Azureus (BitTorrent)
- KAD: Overnet, eMule, aMule

Kademlia uses iterative routing!

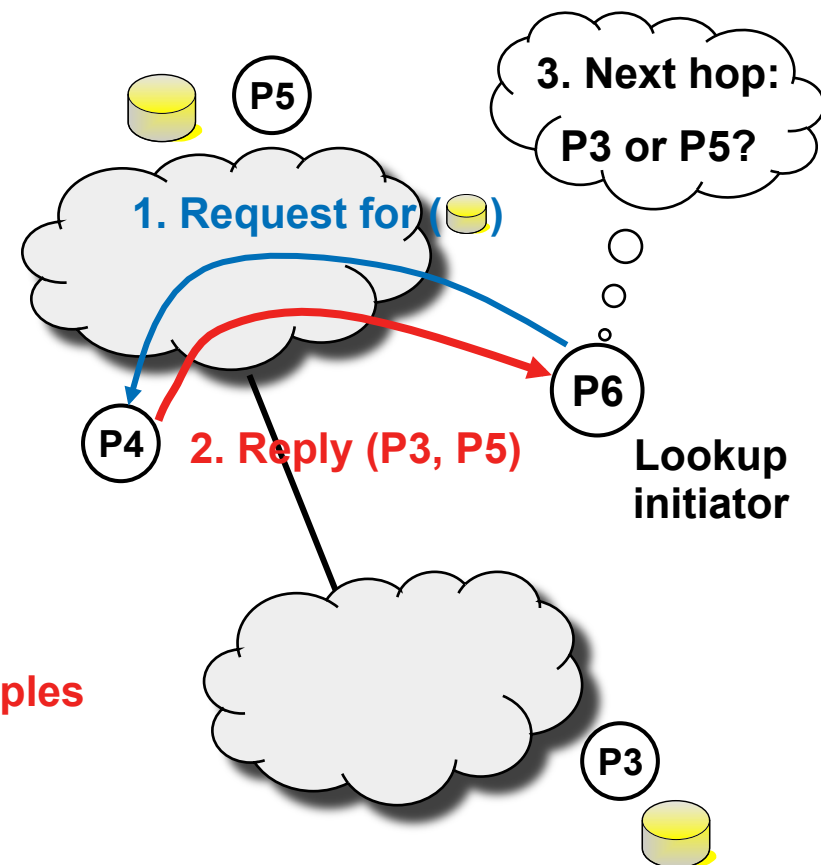
- Lookup initiator is in continuous control of the lookup process

Implications on PRS algorithm

- Next hop candidates: **proximity attributes must be available without time delay**

→ **Iterative routing prohibits taking RTT samples**

- Does not allow to apply related work



Approach (1)

Vast array of conceivable definitions of underlay proximity

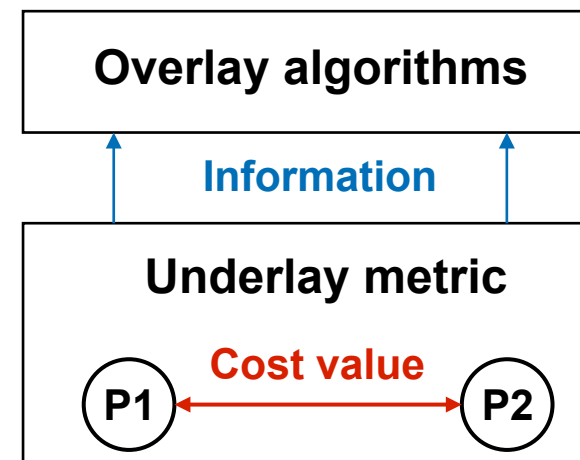
→ Introduction of an abstraction called **underlay metric (UM)**

Underlay metric

- **Definition:** calculates the **underlay-communication costs** between two peers

Idea

- UM provides information (cost value) about the underlay to the overlay network
- Cost value calculation **without time delay**
 - Locally performed by each peer
 - No external entity



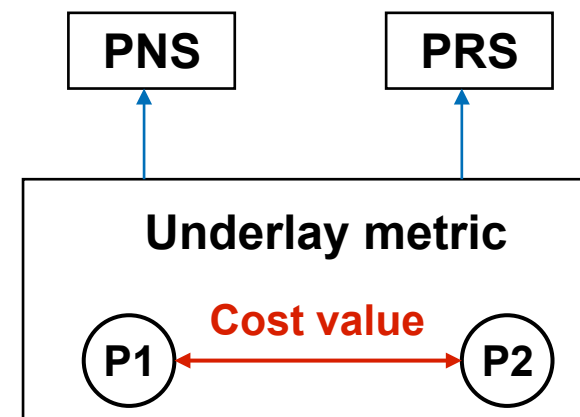
Approach (2)

Extending Kademlia for PRS and PNS

- **Goal:** Improve iterative routing with respect to the underlay

Proposed proximity algorithms

- Allow to take into account information provided by an **underlay metric**
- Preserves all key characteristics of Kademlia
- Passive PNS algorithm: Does not actively search for better peers
 - Introduces no additional overhead



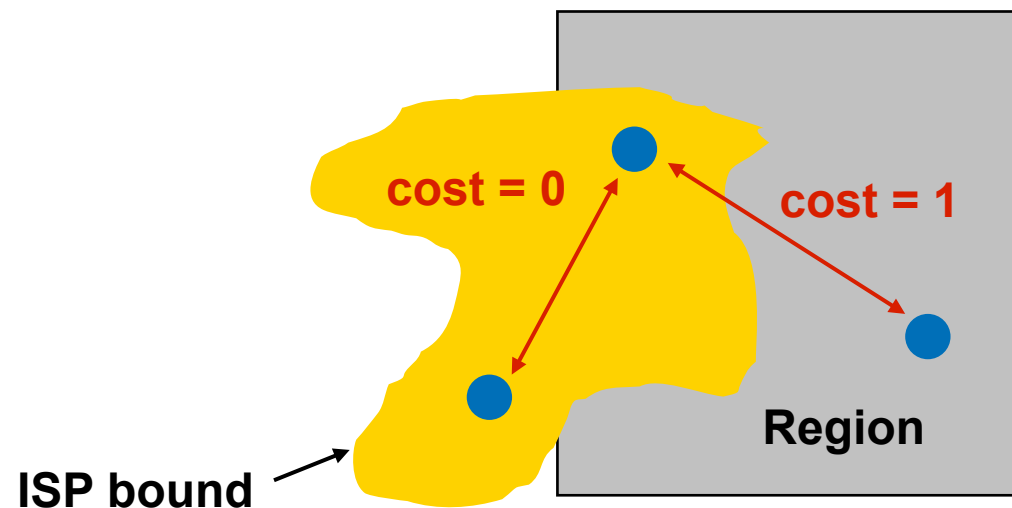
Underlay Metrics: Cluster

Cluster Underlay Metric

- **Objective:** keep largest possible fraction of messages **inside an ISP subnet**

Cost calculation

- Based on geographic information about distant peers
- Lookup of distant peer's IP address in a local copy of MaxMind [1] database
 - Returns ISP, region, country, continent of distant peer
- Local peer ↔ distant peer
 - Same ISP: **cost 0**
 - Same region: **cost 1**
 - Same country: **cost 2**
 - ...



Underlay Metrics: Vivaldi

Vivaldi Underlay Metric [2]

- Objective: **Optimize routing for latency reduction**

Each peer maintains a synthetic network coordinate

- Network coordinate locates peer in multidimensional round-trip time space
- Exchange and update of network coordinates during lookups

Cost calculation

- Euclidean distance between two peers in multidimensional space
 - ➔ **Predicts round-trip time between local and distant peer**

Evaluation Methodology



Goal of evaluation

- Comparison between Plain Old Kademia (POK) and underlay metric-enabled variants

Objectives

- Can locality of traffic be improved by making use of underlay metrics?
- Do low latencies imply locality of traffic?
- Does PNS+PRS perform better than PNS alone?

Evaluation tool

- Using PeerfactSim.KOM, a large-scale simulator for P2P systems [3]
 - Developed at Technische Universität Darmstadt (KOM)

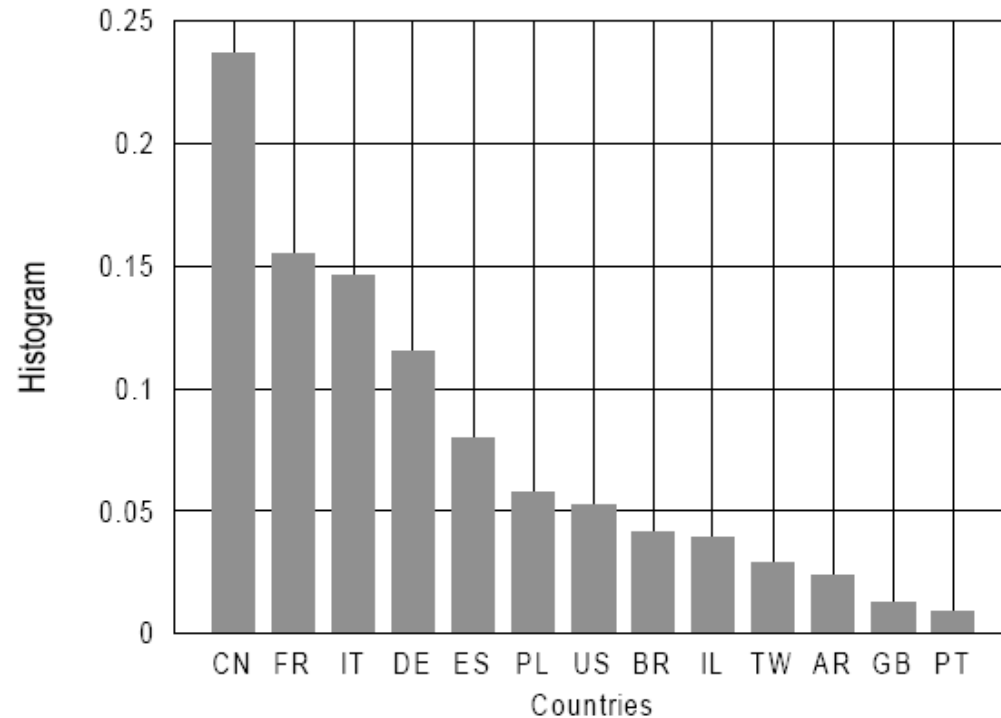
Evaluation Methodology: Experimental Design

Geographical distribution of peers

- Using *M. Steiner et. al*, “A global view of KAD”, in *IMC’07*
 - Session/Inter-session times
 - Geographical distribution of peers

Simulation setup

- 10.000 Peers
- Simulated real-time: 6 hours



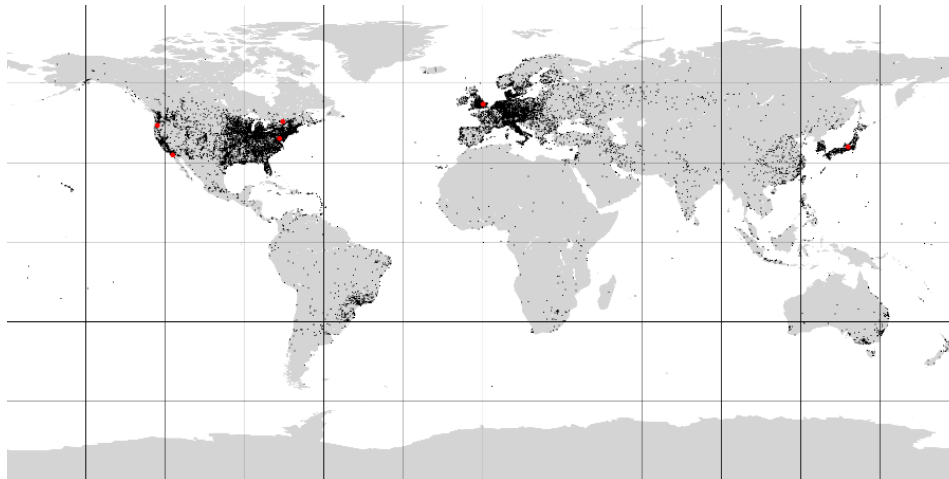
Evaluation Methodology: Internet Delay Space Model

Objective

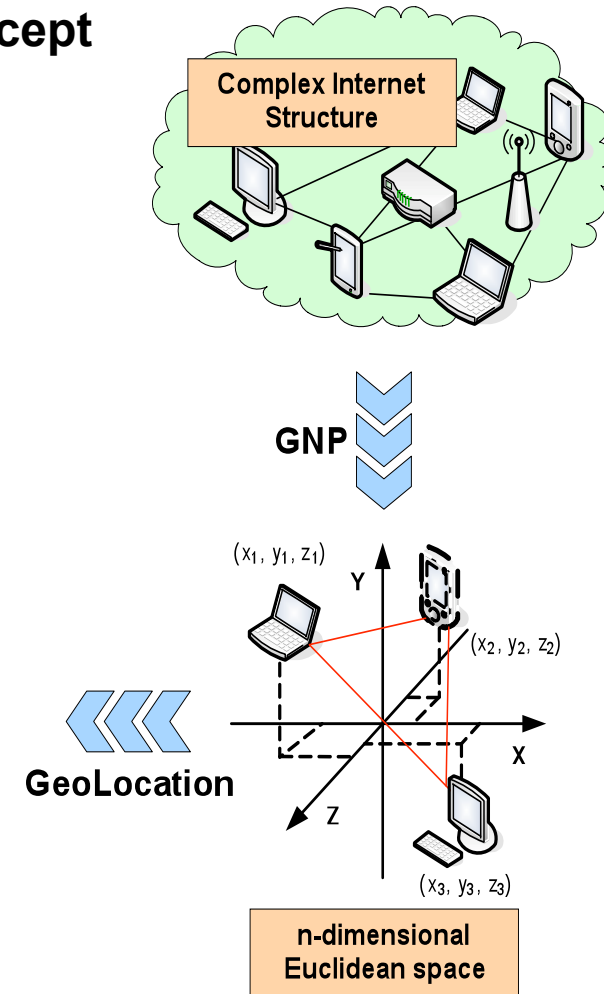
- Lifelike delays (including jitter) between simulated hosts

Solution

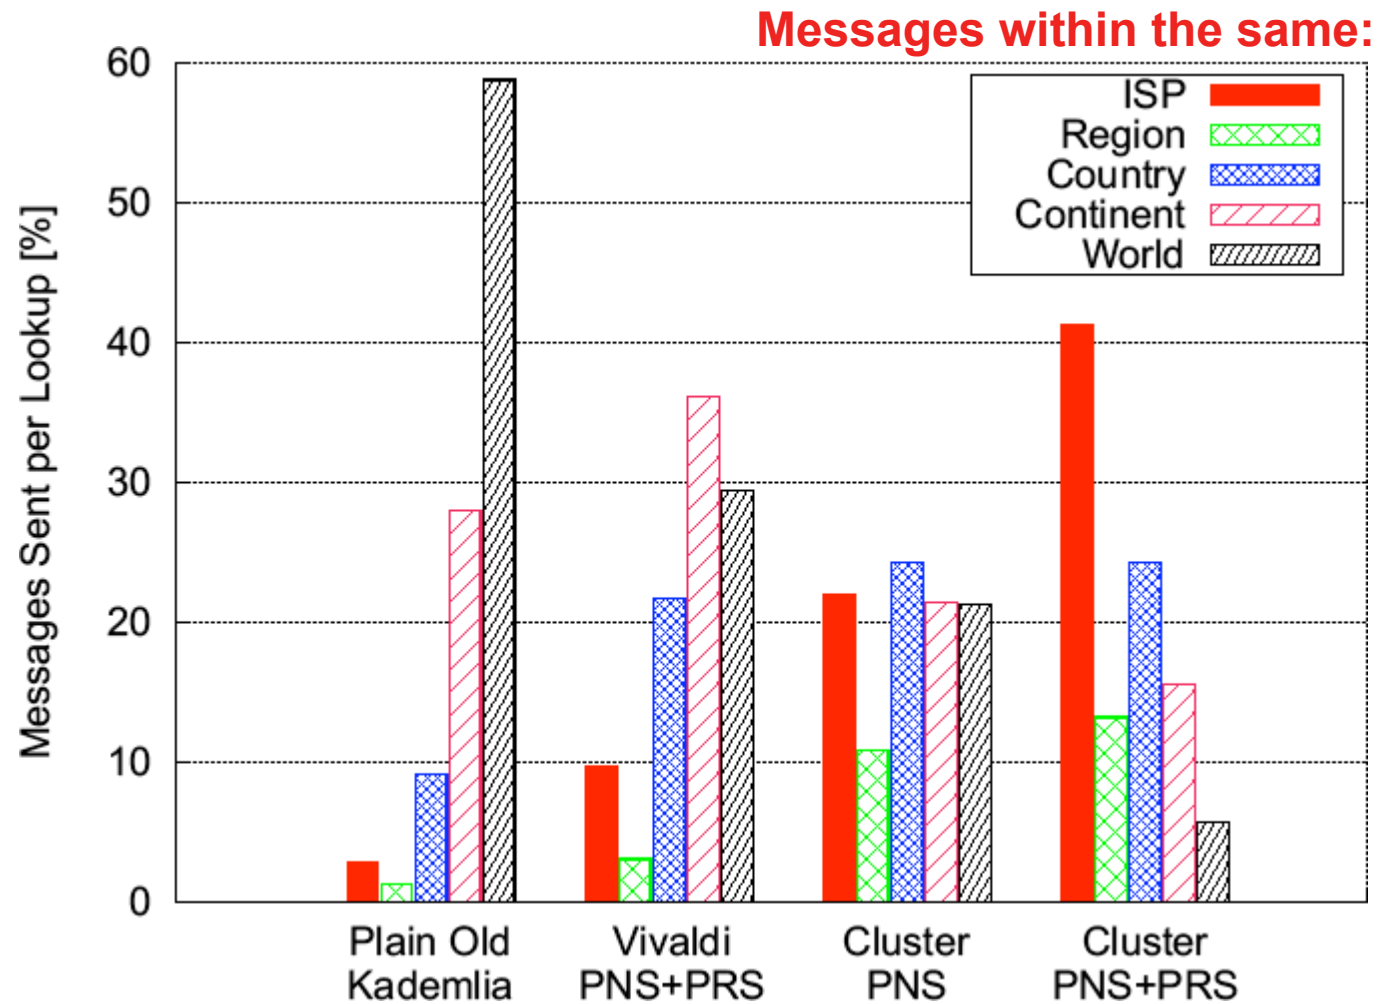
- Use data from two Internet measurement projects (CAIDA and PingER)
- Embed hosts into multidimensional space
 - Technique: Global Network Positioning (GNP)



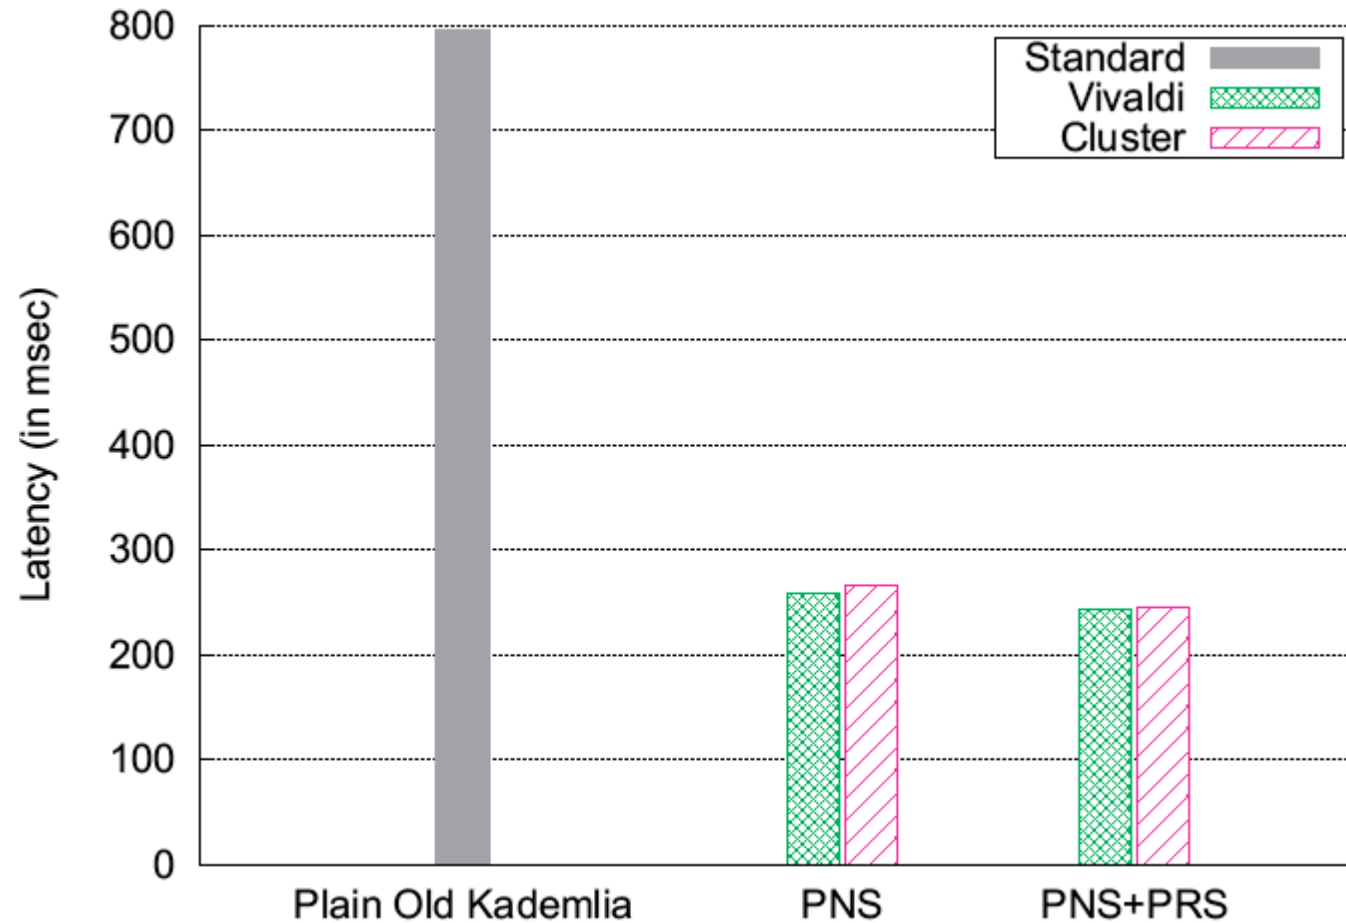
Concept



Results: Message Distribution



Results: Lookup Latencies



Conclusion


Summary

- Reduction of inter-domain traffic at the DHT level
- Comparison between POK and devised underlay metrics
 - **Cost reduction for ISP**
 - 40% of all messages stay inside peer's own ISP
 - **Increased lookup performance in Kademlia**
 - UM reduces lookup latencies of POK by 66 %
 - **Low latencies do not imply locality of traffic**

Future work

- Further underlay metrics: trusted routing, bandwidth, ...
- Trade-off between overlay metric (XOR) and underlay metric

Questions?



Department of Electrical Engineering
and Information Technology
Multimedia Communications Lab - KOM



TECHNISCHE
UNIVERSITÄT
DARMSTADT

Dipl.-Inform. Sebastian Kaune

Sebastian.Kaune@KOM.tu-darmstadt.de

Merckstr. 25
64283 Darmstadt
Germany

Phone +49 (0) 6151/166138
Fax +49 (0) 6151/166152
www.kom.tu-darmstadt.de