

Copyright © 2008
DOCOMO Communications Laboratories
Europe GmbH

The Chordella Overlay Network - Testbed Experiences and Requirements

Dr. Wolfgang Kellerer*

KiVS 2009 workshop on Setup and Operation of Open Testbed Infrastructures in the Context of NGN and Future Internet - Status Quo and Quo Vadis, Kassel, March 6, 2009

*joint work with Z. Despotovic, M. Michel, S. Zöls, Q. Hofstätter, J. Eberspächer

- The Chordella Overlay Network
- Evaluation
 - Simulator
 - Planet Lab
- Requirements for large-scale testbeds

- Peer-to-Peer overlays present a valuable solution space for the Future Internet
 - Self organized
 - Robust and scalable
 - Infrastructure and operation cost saving
- Research usually focuses on the core algorithms for overlay graph maintenance
- Operational aspects are neglected, but are the basis for successfully deployed systems (e.g., Skype, Joost)
 - control of additional (non-functional) features related to system management and control
 - reacting on various environmental (= real-world) conditions, e.g., wireless networks
- Requires sufficient real-world testing

Real world constraints

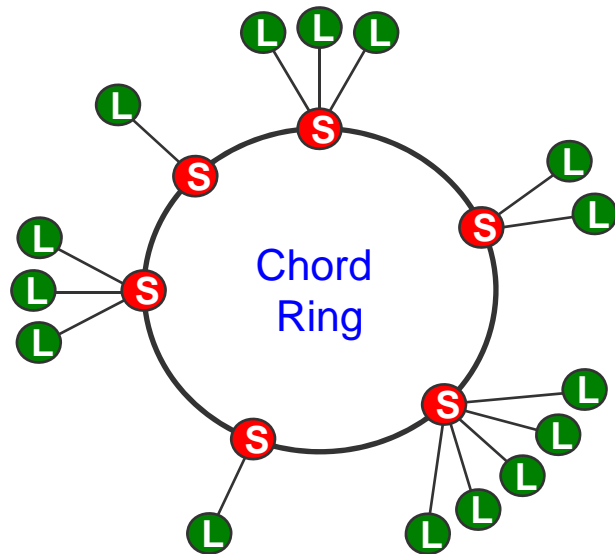
- Heterogeneity of peers ...



- ... and access connections
- DSL 15 Mbit/s
- DSL 1 Mbit/s
- WLAN
- HSDPA
- UMTS
- ISDN
- GPRS

Current P2P design space:
homogeneous peers

Proposed overlay system: "Chordella" Moving Peers out of the main DHT network



S: Superpeer
L: Leafnode

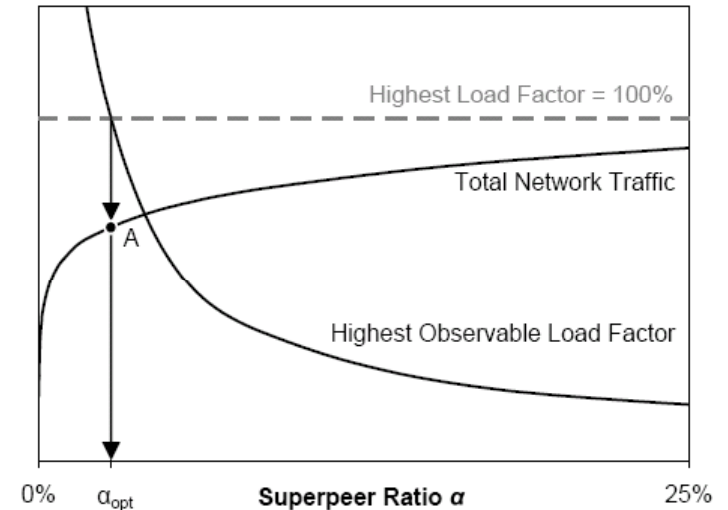
- Hierarchical P2P overlay network
- Superpeer/leafnode concept as in Gnutella 0.6
- Chord as DHT between superpeers
- Superpeers:
 - Proxies for their leafnodes
 - Routing of all search traffic
- Leafnodes:
 - Low capability terminals are connected to a superpeer.
 - No core DHT routing
- Design objectives:
 - (Self-)operability
 - Minimization of operation cost

Zöls et al.: IEEE P2P 2006

Dynamic Cost Optimization: Determine Optimal Superpeer Ratio α_{opt}

- Minimize total cost without overloading a superpeer
 - Total network traffic increases monotonically
 - Load per node decreases

$$\alpha = \frac{\text{\# of superpeers}}{\text{\# of peers}}$$



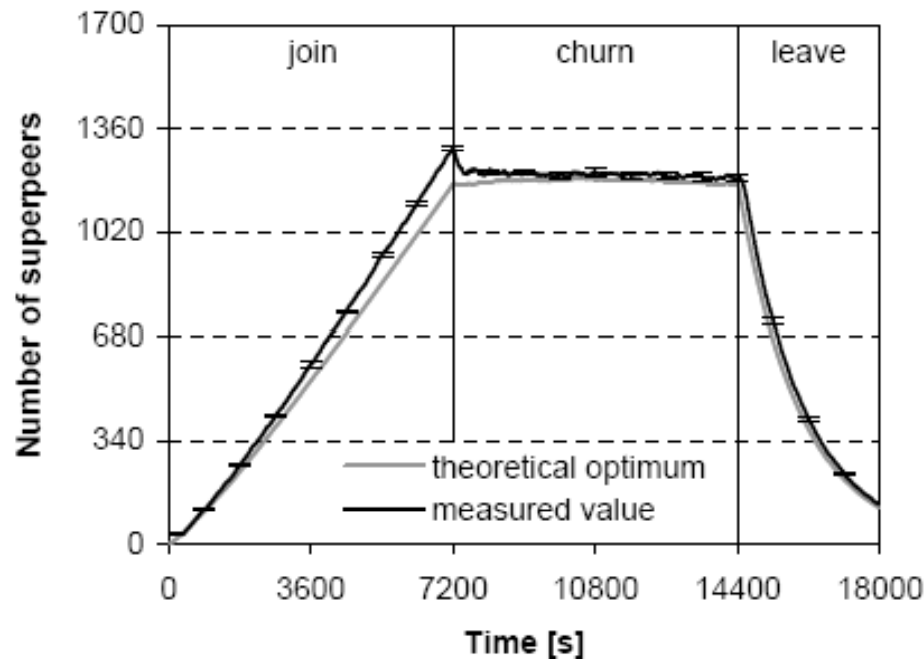
- *Tune alpha*: distributed algorithms for achieving and maintaining optimal superpeer ratio
 1. Load balacing
 2. Distributed parameter estimation
 3. Load calculation

$$s^* = \frac{\frac{1}{N^*} \left(R^* + \frac{F^*}{T_{REP}} \right) (\log N^* + 2) + \frac{3}{T_{STAB}} + \frac{\log N^* \cdot 2}{T_{FIX}} + \frac{X^* - N^*}{N^* \cdot T_{PING}}}{C^*}$$

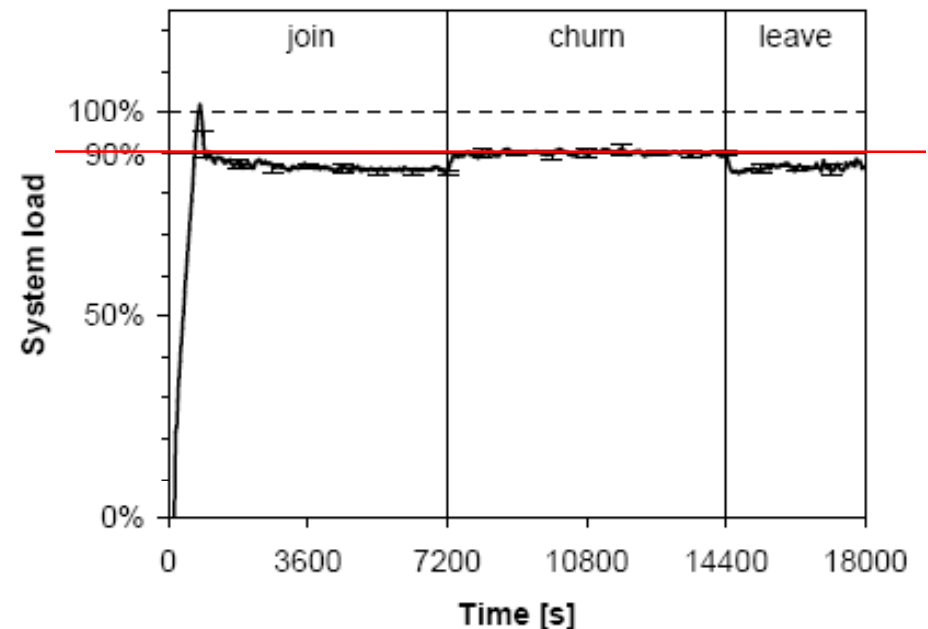
4. Determination of change necessary
5. Local decision (probability based)

Evaluation by simulation

- Simulation based on custom-made layer 7 event-based simulator to achieve scenario with over 10 000 peers



Number of superpeers
 Black: tune alpha
 Grey: optimum with global knowledge



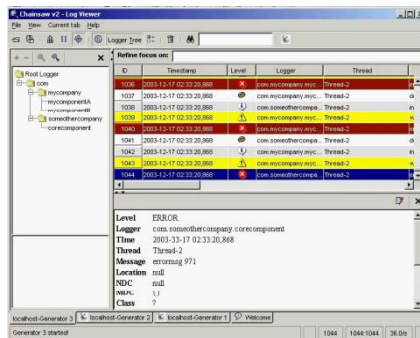


System load
 Target line: 90%

Evaluation in Planet Lab (1)

Chainsaw log viewer
Based on
log4j framework

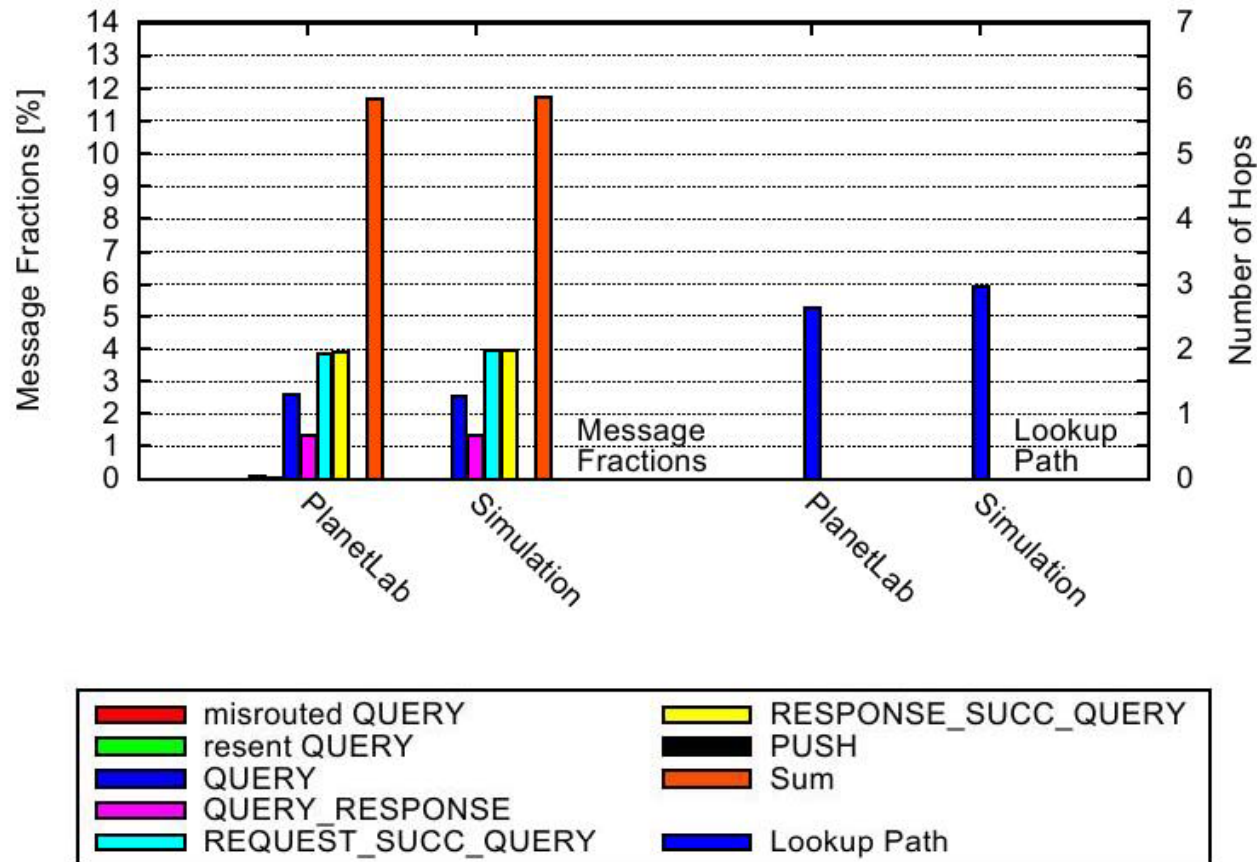
- Verification of simulation results in Planet Lab
- Same codebase used → 500 Planet Lab nodes distributed worldwide
- From 900 nodes only < 600 useable
- Central monitor



Problems:

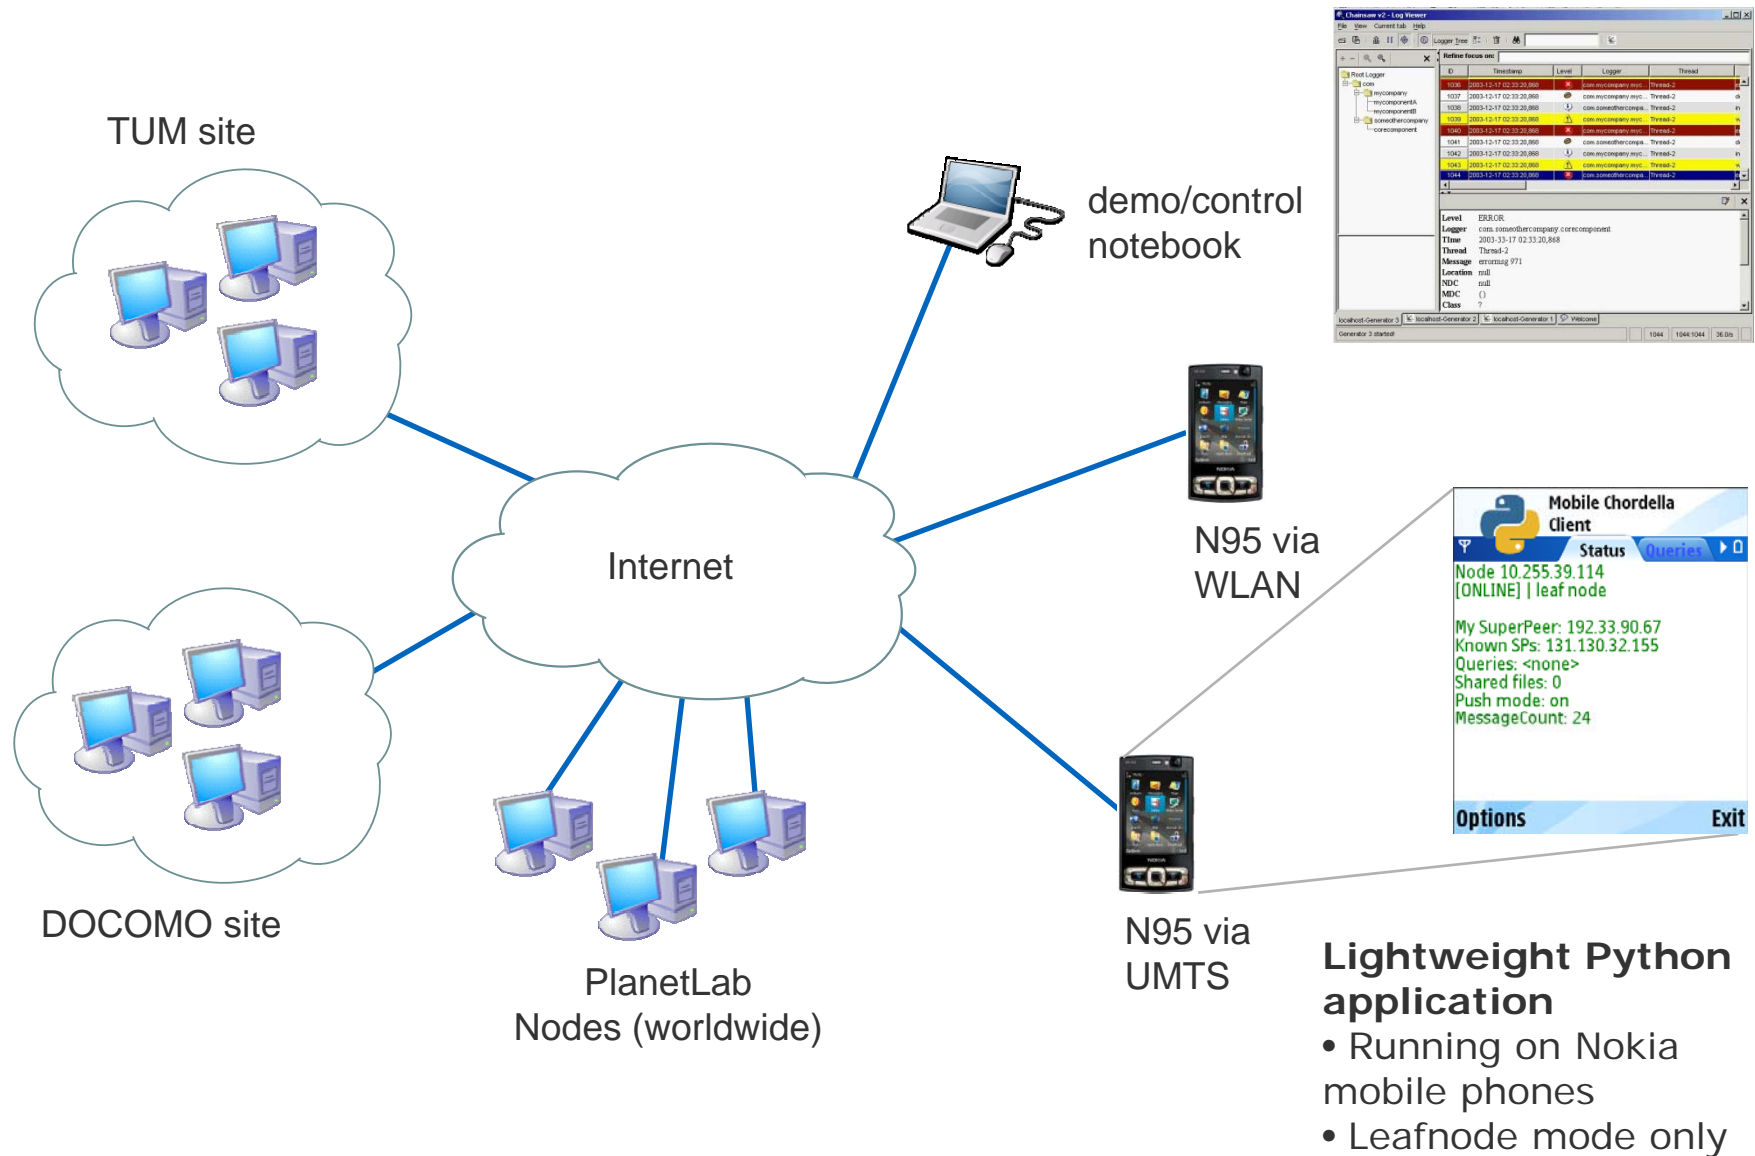
- too many open TCP connections
- non reproduceable results
- no no-churn mode

Evaluation in Planet Lab (2)



- message-fractions match almost perfectly the simulation results
- small deviation in lookup path length due to unpredictability of Planet Lab nodes even in a no Churn scenario

Live deployment in Planet Lab (1)



Live deployment in Planet Lab (2)

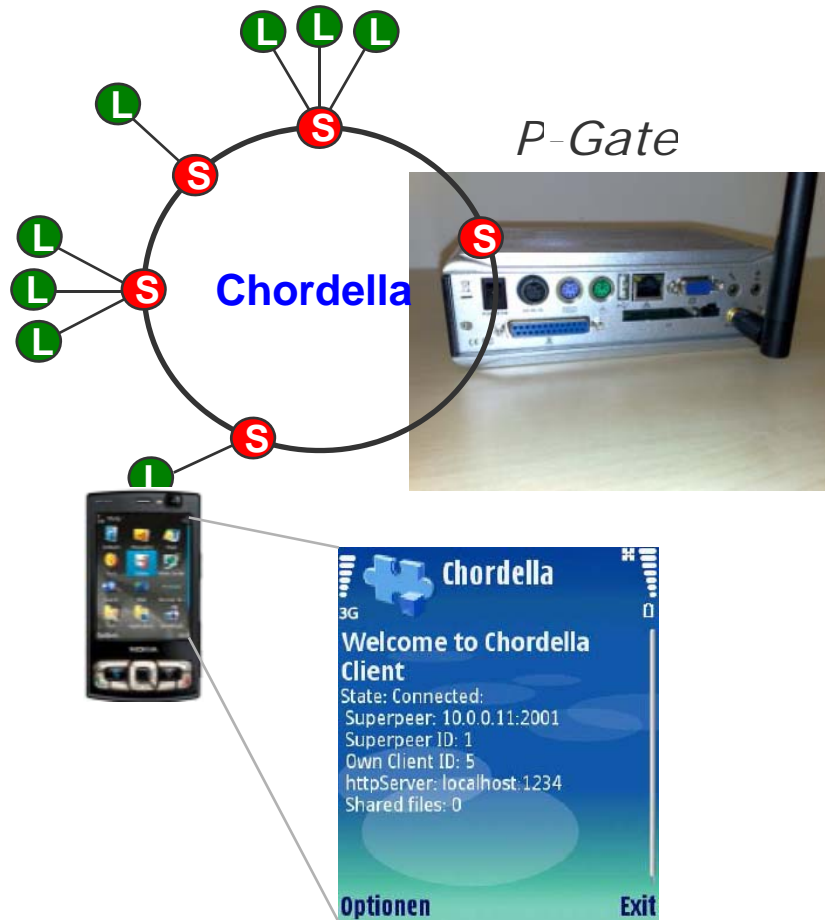
	hostname	Message
:25,816	planetlab1.cs.uchicago.edu	LN9143<- QUERY_RESPONSE[146] for ID {0} from 31331
:25,895	planetlab1.3s.uni-hannover.de	SP31331<- CONNECT[1] from 1639
:25,154	freedom.ri.uni-tuebingen.de	SP37661<- CONNECT_FORWARDING from 31331
:25,929	swsat1501.mpi-sws.mpg.de	SP10099<- CONNECT_FORWARDING from 37661
:25,950	nizza.dvs.informatik.tu-darm...	SP7837<- CONNECT_FORWARDING from 10099
:27,098	planetlab3.cs.uchicago.edu	LN12721<- CONNECT_ACK from 37964
:26,905	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT[1] from 12721
:25,993	planetlab2.georgetown.edu	LN1639<- CONNECT_ACK from 7837
:26,745	freedom.ri.uni-tuebingen.de	SP37661<- QUERY[12] for ID {0} from 52841
:27,573	planetlab1.cs.cornell.edu	LN52841<- QUERY_RESPONSE[99] for ID {0} from 37661
:28,267	swsat1501.mpi-sws.mpg.de	SP10099<- CONNECT[1] from 34735
:28,294	planetlab1.unineuchatel.ch	LN34735<- CONNECT_ACK from 7837
:28,288	nizza.dvs.informatik.tu-darm...	SP7837<- CONNECT_FORWARDING from 10099
:28,257	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT[1] from 56961
:28,519	planetlab1.iii.u-tokyo.ac.jp	LN56961<- CONNECT_ACK from 37964
:29,058	planetlab1.informatik.uni-kl.de	SP62724<- QUERY[12] for ID {3} from 41881
:28,319	freedom.ri.uni-tuebingen.de	SP37661<- QUERY[174] for ID {3} from 62724
:29,119	planetlab1.informatik.uni-kl.de	SP62724<- QUERY_RESPONSE[110] for ID {3} from 37661
:29,160	planetlab-1.iscte.pt	LN41881<- QUERY_RESPONSE[176] for ID {3} from 62724
:29,701	planetlab1.informatik.uni-kl.de	SP62724<- QUERY[12] for ID {3} from 52534
:28,961	freedom.ri.uni-tuebingen.de	SP37661<- QUERY[183] for ID {3} from 62724
:29,738	planetlab1.informatik.uni-kl.de	SP62724<- QUERY_RESPONSE[119] for ID {3} from 37661
:29,753	planetlab1.informatik.uni-kl.de	SP62724<- CONNECT[1] from 6334
:29,641	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT_FORWARDING from 62724
:29,795	planet2.att.nodes.planet-lab...	LN52534<- QUERY_RESPONSE[184] for ID {3} from 62724
:29,828	planetlab-9.CS.Princeton.EDU	LN6334<- CONNECT_ACK from 37964
:31,291	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT[1] from 15862
:31,863	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT[1] from 51393
:32,115	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT_FORWARDING from 62724
:32,287	iraplab1.iralab.uni-karlsruhe.de	SP37964<- QUERY[12] for ID {302} from 45089
:32,338	iraplab1.iralab.uni-karlsruhe.de	SP37964<- QUERY_RESPONSE[141] for ID {302} from 37661
:32,371	iraplab1.iralab.uni-karlsruhe.de	SP37964<- CONNECT_FORWARDING from 7837

Load balancing

xterm

Requirements for large-scale testbeds

- Simple to deploy
 - One script fits all
- Simple to evaluate
 - Central monitoring
- Setting of conditions
- Reproduceable conditions
 - Not unpredictable such as in Planet Lab



Chordella

- hierarchical P2P overlay for mobile networks
- Distributed overlay
→ deployment cost and maintenance cost reduction
→ self optimizing
- Performed large-scale evaluation
- Current work:
live-testbed deployment on WLAN routers

"P-Gate"

Mobile Adventure

Thank you

Dr. Wolfgang Kellerer
kellerer@docomolab-euro.com

DOCOMO Communications Laboratories Europe GmbH
Landsberger Strasse 312 – 80687 Munich, Germany
Phone: +49 (89) 56824-0 | www.docomolab-euro.com

^{NTT}**docomo**
DOCOMO Euro-Labs

References

- S. Zoels, Q. Hofstätter, Z. Despotovic, W. Kellerer. Achieving and Maintaining Cost-Optimal Operation of a Hierarchical DHT System. IEEE Intl. Conference on Communications (IEEE ICC 2009), Dresden, Germany, June 14-18, 2009
- Quirin Hofstätter, Zoran Despotovic, Wolfgang Kellerer, Maximilian Michel. Chordella – A Hierarchical Peer-to-Peer Overlay Implementation for Heterogeneous, Mobile Environments. 8th IEEE Intl. Conference on P2P Computing, Sept. 7-11, 2008, Aachen, Germany.
- W. Kellerer, Z. Despotovic, G. Kunzmann, S. Zoels. Peer-to-Peer Communication – Operational Aspects in Heterogeneous Systems (Tutorial). IEEE ICC 2008, Beijing, China, May 19-23, 2008.
- S. Zoels, Z. Despotovic, W. Kellerer. On Hierarchical DHT Systems – An Analytical Approach for Optimal Designs. Computer Communications (COMCOM), Elsevier, Vol. 31, Issue 3 (Feb. 2008), pp. 576-590.
- W. Kellerer. The Bright Future of P2P: a Telecom Operator's Perspective (Invited keynote). IEEE Intl. Conference on Peer-to-Peer Computing, (P2P 2007), Galway, Ireland, September 2-5, 2007.
- S. Zoels, Z. Despotovic, W. Kellerer. Load-Balancing in Hierarchical DHT-based P2P Systems. CollaborateCom 2007, White Plains, New York, USA, November 12-15, 2007.
- S. Zoels, Z. Despotovic, W. Kellerer. Cost-Based Analysis of Hierarchical DHT Design. In proc. 6th IEEE Intl. Conference on P2P Computing, Sept. 6-8, Cambridge, UK.
- W. Kellerer, G. Kunzmann, R. Schollmeier, S. Zöls. Structured Peer-to-Peer Systems for Telecommunications and Mobile Environments. AEÜ, Elsevier, Vol. 60, Issue 1, January 2006, pp. 25-29.