DOCTORAL DISSERTATION

Structured Peer-to-Peer Services for Mobile Ad Hoc Networks

Dissertation
zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften
im Fachbereich Mathematik und Informatik
der Freien Universität Berlin
vorgelegt von
Dipl.–Inform. Thomas Zahn
aus Berlin

Datum der Disputation: 26. Juli 2006

Gutachter:

Prof. Dr.–Ing. Jochen H. Schiller

Prof. Dr. Peter Druschel

Acknowledgements

The work presented herein was carried out under a research fellowship in the Berlin Berlin-Brandenburg Graduate School on Distributed Information Systems mainly at the Freie Universität Berlin, Germany and partly at Rice University, USA. This document is the accumulation of my work as a PhD student at the Computer Systems and Telematics group headed by Professor Jochen Schiller. I would like to take this opportunity to thank all the people that supported me during this time personally and professionally.

First of all, I would like thank Prof. Jochen Schiller who has been my primary adviser over the past three years. Without his support and inspiration as well as the confidence and freedom that he has given me to pursue my ideas, this work would never have been possible. He is a most extraordinary and inspiring person with a passion for his work and joie de vivre that I have never encountered before. I am immensely grateful for the opportunity that he has given me to be part of his most remarkable research group. This has been one of the best times of my life. I truly owe him my deepest gratitude.

Furthermore, I would like to thank Prof. Peter Druschel for being my second adviser. The research visit that I spent in his amazing research group at Rice University substantially shaped the key ideas of this work. The personal advice that he has provided over the last two years has been priceless. Knowing how busy his schedule is makes his accepting to be my adviser a true honor.

Additionally, I would like to thank Prof. Agnès Voisard who has been my second adviser within the Berlin-Brandenburg Graduate School on Distributed Information Systems. I am very grateful for her support and encouragement over the past three years. Coming from a slightly different background, our discussions have always been immensely helpful in staying on track and not becoming lost in technical details. I also owe my deep gratitude to her for her personal support concerning my next career steps.

Moreover, I am very grateful to the members of the Berlin-Brandenburg Graduate School on Distributed Information Systems for their continuous feedback and advise. The degree of scientific freedom provided by the school has been most remarkable and inspiring.

Rolf Winter deserves my special thanks. I have been incredibly fortunate to share an office with him during the entire time of my PhD studies. Our joint work at the beginning of our studies has significantly helped in jump-starting my research. The literally daily discussions with Rolf have greatly inspired this work and always served as a "quality measure" for my ideas. In addition to working on scientific challenges, we challenged each other on the Squash court and on Thursday evenings. We have become very good friends.

I am also very grateful to my colleagues in the Computer Systems and Telematics group. The work atmosphere in this group has been the most inspiring and enjoyable I could ever imagine.

I would also like to express my deep gratitude to Tonia-Alexa Richter. Her love and support during all phases of this work and during all times over the past years have given me the strength to carry out this work. I owe her my eternal gratitude.

Finally, I would like to thank my parents and family not only for their support during my PhD studies, but also for the unconditional love, support, encouragement and freedom that they have given me throughout my entire life. Thank you for always being there for me all these years.

Abstract

As large mobile ad hoc networks (MANETs) become realistic, it also becomes more and more interesting to build the distributed network applications that one is accustomed to from the domain of the Internet on top of those MANETs. In the Internet, Distributed Hash Tables (DHTs) have recently proven themselves an efficient building block for such distributed applications. However, conventional DHTs are not well-suited for direct deployment in MANETs as they are largely oblivious to the physical routing.

This document presents MADPastry, a DHT explicitly designed for the use in MANETs. MADPastry considers physical locality and integrates the functionality of a conventional DHT and a reactive ad hoc routing protocol at the network layer to provide an efficient indirect (i.e. key-based) routing primitive in MANETs. Furthermore, it is also demonstrated how to build scalable distributed applications for MANETs on top of the presented DHT.

Zusammenfassung

Große mobile Ad-hoc-Netze (MANETs) werden immer realistischer. Daher wird es ebenfalls immer interessanter, die verteilten Netz-Applikationen, die man aus dem Bereich des Internets kennt, für solche MANETs zu erstellen. Im Internet haben sich kürzlich Verteilte Hash-Tabellen (DHTs) als effiziente Bausteine für solche verteilten Applikation herausgestellt. Allerdings sind konventionelle DHTs nicht gut geeignet für einen direkten Einsatz in MANETs, da diese DHTs weitestgehend losgelöst sind vom physischen Routing.

In dieser Arbeit wird MADPastry, eine explizit für MANETs entworfene DHT, vorgestellt. MADPastry berücksichtigt physische Lokalität und integriert die Funktionalität einer konventionellen DHT und eines reaktiven Ad-hoc-Routingprotokolls auf der Netzwerkschicht, um effizientes indirektes (d.h. schlüsselbasiertes) Routing in MANETs zu ermöglichen. Zusätzlich wird demonstriert, wie skalierbare verteilte Applikationen für MANETs auf der präsentierten DHT erstellt werden können.

Table of Contents

1 Introduction 12
1.1 Problem Statement 14
1.2 Contributions 16
1.3 Thesis Overview 17
2 Background 18
2.1 Peer-to-Peer Overlay Networks 18
2.1.1 Unstructured Overlays 19
2.1.2 Structured Overlays 20
2.2 Mobile Ad Hoc Networks 24
2.2.1 Flat Routing Protocols 24
2.2.2 Hierarchical Routing Protocols 26
2.3 Topology-Aware Structured Overlay Networks 27
2.4 Summary 34
3 Related Work 35
3.1 Conventional Distributed Hash Tables 35
3.2 Unstructured Peer-to-Peer Networks for MANETs 36
3.2.1 Flooding-Based Protocols 36
3.2.2 Unstructured Key Lookup 37
3.2.3 Proactive Search Routing for Mobile Peer-to-Peer Networks 38
3.3 Structured P2P Network for MANETs 39
3.3.1 Chord-Based 39
3.3.2 Pastry-Based 42
3.4 Service Discovery in MANETs 44
3.4.1 Broadcast-Based Service Discovery 44
3.4.2 Multicast-Based Service Discovery 45
3.4.3 Cluster-Based Service Discovery 46

3.4.4 Hierarchical Service Discovery 49

3.4.5 Geographic Service Discovery 50

3.4.6 Geographic Service Discovery Without Location Information 51

3.5 Summary 52

4 The MADPastry Architecture 55

4.1 Motivation and Architectural Overview 55

4.2 Clusters 57

4.3 Routing Tables 59

4.4 Routing 62

4.5 Routing Table Maintenance 65

4.6 Bootstrapping 67

4.7 Summary 68

5 Experimental Results 69

5.1 Basic Results 71

5.2 Load Distribution 74

5.3 Node Velocity 76

5.4 Lookup Rates 79

5.5 Churn 81

5.6 Handovers 84

5.7 Summary 89

6 Application I: Peer-to-Peer Based Name Service for MANETs 90

6.1 Introduction 90

6.2 Resource Advertisement 90

6.3 Resource Discovery 92

6.4 Local Replications 93

6.5 Handovers and Caching 95

6.6 Experimental Results 96

6.6.1 Basic Results 97

6.6.2 Local Replications 99

6.7 Related Work 101

6.8 Summary 102

7 Application II: A DHT-based Unicast for MANETs 103

7.1 Address Servers 103

7.2 MADPastry's Unicast 104

7.3 Shortest Routing Paths vs. MADPastry's Routing Paths 105

7.4 Experimental Results 107

7.4.1 Network Size 108

7.4.2 Node Velocity 111

7.5 Related Work 113

7.6 Summary 114

8 Conclusion and Future Work 115

8.1 Conclusion 115

8.2 Future Work 118

9 Appendix 120

9.1 Abbreviations 120

9.2 MADPastry Parameters and Default Values 121

10 References 123

List of Figures

Figure 1.1 Overlay vs. physical routing. 15
Figure 2.1 P2P overlay network vs. physical network.. 18
Figure 2.2 Example of a Pastry routing table and leaf set. Overlay ID digits have base 8 (b=3) and the leaf set contains 2^b entries.. 21
Figure 2.3 Example of the Pastry overlay routing process. 23
Figure 2.4 Example of an AODV routing table.. 25
Figure 2.5 Overlay stretch.. 28
Figure 2.6 Spatial distribution of overlay ID prefixes in a Pastry network. Equal symbols and shades of grey represent equal overlay ID prefixes. 30
Figure 2.7 Spatial distribution of overlay ID prefixes in a RLM network. Equal symbols and shades of grey represent equal overlay ID prefixes. 30
Figure 2.8 Overlay stretch for Pastry and RLM mobile networks. 31
Figure 2.9 Total number of overlay messages exchanged during an average 24h simulation......... 32
Figure 2.10 Success rate of Pastry and RLM in a 100-node AODV network.................................. 32
Figure 2.11 Average number of physical packets per overlay hop. 33
Figure 2.12 Physical traffic generated by Pastry and RLM.. 33
Figure 3.1 Node address tree with the l level-k siblings of node address 100. [14]........................ 49
Figure 4.1 MADPastry in the network stack... 56
Figure 4.2 Spatial distribution of overlay prefixes in a 250-node MADPastry network. 58
Figure 4.3 Beacon cluster broadcast. 59
Figure 4.4 MADPastry routing tables. 61
Figure 4.5 MADPastry routing - local view. 63
Figure 4.6 MADPastry routing - network view.. 65
Figure 5.1 Success rates of the respective routing agents – 1.4m/s... 71
Figure 5.2 Total number of messages - 1.4m/s. 72
Figure 5.3 Overall generated traffic in Kbytes - 1.4m/s. 72
Figure 5.4 Overlay stretch.. 74
Figure 5.5 Percentage deviation of the individual accumulated node loads from the median........ 75
Figure 5.6 Success rates vs. node velocity. 77
Figure 5.7 Total number of messages vs. node velocity.. 78
Figure 5.8 Overall traffic vs. node velocity. 78
Figure 5.9 Success rates vs. lookup intervals... 79
Figure 5.10 Number of packets vs. lookup interval. 80
Figure 5.11 Overall traffic vs. lookup intervals.. 80
Figure 5.12 Success rates vs. churn rates. 82
Figure 5.13 Overall traffic vs. churn rates... 83
Figure 5.14 Success rate vs. number of objects – 1 lookup per 10 seconds per node........................ 86
Figure 5.15 Number of packets vs. number of objects – 1 lookup per 10 seconds per node. 86
Figure 5.16 Overall traffic vs. total number of objects – 1 lookup per 10 seconds per node............ 87
Figure 5.17 Success rate vs. number of objects – 1 lookup per second per node............................... 87
Figure 5.18 Number of packets vs. number of objects – 1 lookup per 1 second per node.................. 88
Figure 5.19 Overall traffic vs. total number of objects – 1 lookup per second per node.................... 88
Figure 6.1 The MAPNaS architecture.. 91
Figure 6.2 Indirect routing of a MAPNaS resource advertisement using MADPastry. 92
Figure 6.3 MAPNaS resource discovery using MADPastry routing.. 93
Figure 6.4 Resource advertisement and discovery with local replications.. 94
Figure 6.5 Success rates - basic results. 97
Figure 6.6 Total number of packets sent. 98
Figure 6.7 Overall traffic. 98
Figure 6.8 Success rate vs. traffic pattern... 100
Figure 6.9 Overall traffic vs. traffic pattern. 101
Figure 7.1 Address publication with MADPastry's unicast scheme.. 104
Figure 7.2 Address resolution... 105
Figure 7.3 MADPastry's unicast. 106

Figure 7.4 Success rate vs. network size.. 109
Figure 7.5 Total number of packets sent vs. network size. 110
Figure 7.6 Total amount of generated network traffic vs. network size. 111
Figure 7.7 Success rate vs. node velocity.. 112
Figure 7.8 Total amount of generated network traffic vs. node velocity... 113

List of Tables

Table 3.1 Assessment of related approaches. 54
Table 5.1 Simulation parameters and values.. 71
Table 5.2 Individual accumulated node traffic load.. 75
Table 5.3 Simulation parameters and values – varying node velocities. 76
Table 5.4 Simulation parameters and values – varying lookup rates.. 79
Table 5.5 Simulation parameters and values – varying churn rates. 81
Table 5.6 Simulation parameters and values – handovers.. 85
Table 6.1 Simulation parameters and values – basic results. 97
Table 6.2 Simulation parameters and values – local replications. 99
Table 7.1 Simulation parameters and values – different network sizes.. 108
Table 7.2 Simulation parameters and values – different node velocities.. 112
Table 8.1 Comparison of MADPastry and related approaches.. 117