

Lecture Notes in Computer Science

2429

Edited by G. Goos, J. Hartmanis, and J. van Leeuwen

**Springer**

*Berlin*

*Heidelberg*

*New York*

*Barcelona*

*Hong Kong*

*London*

*Milan*

*Paris*

*Tokyo*

Peter Druschel Frans Kaashoek  
Antony Rowstron (Eds.)

# Peer-to-Peer Systems

First International Workshop, IPTPS 2002  
Cambridge, MA, USA, March 7-8, 2002  
Revised Papers



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany  
Juris Hartmanis, Cornell University, NY, USA  
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Peter Druschel  
Rice University  
MS 132, Houston, TX 88005, USA  
E-mail: druschel@cs.rice.edu

Frans Kaashoek  
MIT Laboratory of Computer Science  
200 Technology Square, Cambridge MA, USA  
E-mail: kaashoek@lcs.mit.edu

Antony Rowstron  
Microsoft Research Ltd.  
7 J J Thomson Avenue, Cambridge, CB3 0FB, UK  
E-mail: antr@microsoft.com

Cataloging-in-Publication Data applied for

Bibliographic information published by Die Deutsche Bibliothek  
Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie;  
detailed bibliographic data is available in the Internet at <http://dnb.ddb.de>

CR Subject Classification (1998): C.2, H.3, H.4, D.4, F.2.2, E.1

ISSN 0302-9743

ISBN 3-540-44179-4 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York  
a member of BertelsmannSpringer Science+Business Media GmbH  
<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2002  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Boller Mediendesign  
Printed on acid-free paper SPIN: 10873895 06/3142 5 4 3 2 1 0

## Preface

Peer-to-peer has emerged as a promising new paradigm for large-scale distributed computing. The International Workshop on Peer-to-Peer Systems (IPTPS) aimed to provide a forum for researchers active in peer-to-peer computing to discuss the state of the art and to identify key research challenges.

The goal of the workshop was to examine peer-to-peer technologies, applications, and systems, and also to identify key research issues and challenges that lie ahead. In the context of this workshop, peer-to-peer systems were characterized as being decentralized, self-organizing distributed systems, in which all or most communication is symmetric.

The program of the workshop was a combination of invited talks, presentations of position papers, and discussions covering novel peer-to-peer applications and systems, peer-to-peer infrastructure, security in peer-to-peer systems, anonymity and anti-censorship, performance of peer-to-peer systems, and workload characterization for peer-to-peer systems. To ensure a productive workshop environment, attendance was limited to 55 participants.

Each potential participant was asked to submit a position paper of 5 pages that exposed a new problem, advocated a specific solution, or reported on actual experience. We received 99 submissions and were able to accept 31. Participants were invited based on the originality, technical merit, and topical relevance of their submissions, as well as the likelihood that the ideas expressed in their submissions would lead to insightful technical discussions at the workshop.

A digest of the discussions that took place at the workshop is provided in the first chapter. Thanks to Richard Clayton for editing this digest based on notes taken at the workshop by himself and George Danezis. We would like to thank the steering committee for their role in initiating the workshop, MIT for hosting the workshop, Neena Lyall for local arrangements and for taking care of countless logistical details, Kevin Fu for setting up the secure Web registration site, and last but not least, the program committee for selecting a superb technical program.

Finally, we wish to thank all participants for making IPTPS 2002 a great success. Plans are currently underway for IPTPS 2003, which is sure to build on the success of this first workshop.

July 2002

Peter Druschel, Frans Kaashoek, and Antony Rowstron

## Organization

### Steering Committee

Peter Druschel	Rice University, USA
Frans Kaashoek	MIT, USA
Antony Rowstron	Microsoft Research, UK
Scott Shenker	ICIR, Berkeley, USA
Ion Stoica	UC Berkeley, USA

### Program Committee

Ross Anderson	Cambridge University, UK
Roger Dingledine	Reputation Technologies, Inc., USA
Peter Druschel (co-chair)	Rice University, USA
Steve Gribble	University of Washington, USA
David Karger	MIT, USA
John Kubiawicz	UC Berkeley, USA
Robert Morris	MIT, USA
Antony Rowstron (co-chair)	Microsoft Research, UK
Avi Rubin	AT&T Labs Research, USA
Scott Shenker	ICIR, Berkeley, USA
Ion Stoica	UC Berkeley, USA

### Organizing Chairs

Frans Kaashoek	MIT, USA
Antony Rowstron	Microsoft Research, UK

### Sponsors

Microsoft Research	<a href="http://www.research.microsoft.com/">http://www.research.microsoft.com/</a>
--------------------	---

## Table of Contents

Workshop Report for IPTPS'02: 1st International Workshop on Peer-to-Peer Systems . . . . .	1
<i>Richard Clayton (Cambridge University)</i>	
 <b>Structure Overlay Routing Protocols: State of the Art and Future Directions</b>	
Observations on the Dynamic Evolution of Peer-to-Peer Networks . . . . .	22
<i>David Liben-Nowell, Hari Balakrishnan, and David Karger (MIT)</i>	
Brocade: Landmark Routing on Overlay Networks . . . . .	34
<i>Ben Y. Zhao, Yitao Duan, Ling Huang, Anthony D. Joseph, and John D. Kubiatowicz (UC Berkeley)</i>	
Routing Algorithms for DHTs: Some Open Questions . . . . .	45
<i>Sylvia Ratnasamy, Ion Stoica (UC Berkeley), and Scott Shenker (ICSI, Berkeley)</i>	
Kademlia: A Peer-to-Peer Information System Based on the XOR Metric . . . . .	53
<i>Petar Maymounkov and David Mazières (New York University)</i>	
Efficient Peer-to-Peer Lookup Based on a Distributed Trie . . . . .	66
<i>Michael J. Freedman (MIT) and Radek Vingralek (Oracle Corporation)</i>	
Self-Organizing Subsets: From Each According to His Abilities, to Each According to His Needs . . . . .	76
<i>Amin Vahdat, Jeffrey Chase, Rebecca Braynard, Dejan Kostić, Patrick Reynolds, and Adolfo Rodriguez (Duke University)</i>	
 <b>Deployed Peer-to-Peer Systems</b>	
Mapping the Gnutella Network: Macroscopic Properties of Large-Scale Peer-to-Peer Systems . . . . .	85
<i>Matei Ripeanu and Ian Foster (The University of Chicago)</i>	
Can Heterogeneity Make Gnutella Scalable? . . . . .	94
<i>Qin Lv (Princeton University), Sylvia Ratnasamy (UC Berkeley), and Scott Shenker (ICSI, Berkeley)</i>	
Experiences Deploying a Large-Scale Emergent Network . . . . .	104
<i>Bryce Wilcox-O'Hearn (zooko.com)</i>	

## Anonymous Overlays

Anonymizing Censorship Resistant Systems .....	111
<i>Andrei Serjantov (Cambridge University)</i>	

Introducing Tarzan, a Peer-to-Peer Anonymizing Network Layer .....	121
<i>Michael J. Freedman, Emil Sit, Josh Cates, and Robert Morris (MIT)</i>	

## Applications

Mnemosyne: Peer-to-Peer Steganographic Storage .....	130
<i>Steven Hand (Cambridge University) and Timothy Roscoe (Sprint Advanced Technology Lab)</i>	

ConChord: Cooperative SDSI Certificate Storage and Name Resolution ..	141
<i>Sameer Ajmani, Dwaine E. Clarke, Chuang-Hue Moh, and Steven Richman (MIT)</i>	

Serving DNS Using a Peer-to-Peer Lookup Service .....	155
<i>Russ Cox, Athicha Muthitacharoen, and Robert T. Morris (MIT)</i>	

Network Measurement as a Cooperative Enterprise .....	166
<i>Sridhar Srinivasan and Ellen Zegura (Georgia Institute of Technology)</i>	

The Case for Cooperative Networking .....	178
<i>Venkata N. Padmanabhan (Microsoft Research) and Kunwadee Sripanidkulchai (Carnegie Mellon University)</i>	

Internet Indirection Infrastructure.....	191
<i>Ion Stoica, Dan Adkins, Sylvia Ratnasamy (UC Berkeley), Scott Shenker (ICSI, Berkeley), Sonesh Surana, and Shelley Zhuang (UC Berkeley)</i>	

Peer-to-Peer Caching Schemes to Address Flash Crowds .....	203
<i>Tyron Stading, Petros Maniatis, and Mary Baker (Stanford University)</i>	

## Evaluation

Exploring the Design Space of Distributed and Peer-to-Peer Systems: Comparing the Web, TRIAD, and Chord/CFS.....	214
<i>Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble (University of Washington)</i>	

Are Virtualized Overlay Networks Too Much of a Good Thing? .....	225
<i>Pete Keleher, Samrat Bhattacharjee, and Bujor Silaghi (University of Maryland)</i>	



## Searching and Indexing

Locating Data in (Small-World?) Peer-to-Peer Scientific Collaborations . .	232
<i>Adriana Iamnitchi, Matei Ripeanu, and Ian Foster (University of Chicago)</i>	
Complex Queries in DHT-based Peer-to-Peer Networks . . . . .	242
<i>Matthew Harren, Joseph M. Hellerstein, Ryan Huebsch, Boon Thau Loo (UC Berkeley), Scott Shenker (ICSI, Berkeley), and Ion Stoica (UC Berkeley)</i>	
The Sybil Attack . . . . .	251
<i>John R. Douceur (Microsoft Research)</i>	
Security Considerations for Peer-to-Peer Distributed Hash Tables . . . . .	261
<i>Emil Sit and Robert Morris (MIT)</i>	
Dynamically Fault-Tolerant Content Addressable Networks . . . . .	270
<i>Jared Saia (University of Washington), Amos Fiat (Tel Aviv University), Steven Gribble, Anna R. Karlin, and Stefan Saroiu (University of Washington)</i>	

## Data Management

Scalable Management and Data Mining Using Astrolabe . . . . .	280
<i>Robbert van Renesse, Kenneth Birman, Dan Dumitriu, and Werner Vogels (Cornell University)</i>	
Atomic Data Access in Distributed Hash Tables . . . . .	295
<i>Nancy Lynch (MIT), Dahlia Malkhi (Hebrew University), and David Ratajczak (UC Berkeley)</i>	
Dynamic Replica Placement for Scalable Content Delivery . . . . .	306
<i>Yan Chen, Randy H. Katz, and John D. Kubiatowicz (UC Berkeley)</i>	
Peer-to-Peer Resource Trading in a Reliable Distributed System . . . . .	319
<i>Brian F. Cooper and Hector Garcia-Molina (Stanford University)</i>	
Erasur Coding Vs. Replication: A Quantitative Comparison . . . . .	328
<i>Hakim Weatherspoon and John D. Kubiatowicz (UC Berkeley)</i>	
Author Index . . . . .	339