

Tutorial at ITC18: Network Efficient P2P-Services and their Management

Authors: Dr. K. Tutschku

Universität Würzburg, Lehrstuhl für Verteilte Systeme,
Am Hubland, 97074 Würzburg.
Tel.: +49-931-8886641, Fax: +49-931-8886632,
e-Mail: tutschku@informatik.uni-wuerzburg.de,

Dr. H. de Meer

University College London, Department of Electronic & Electrical Engineering,
Torrington Place, London WC1E 7JE, United Kingdom.
Tel.: +44-20-7679-7308. Fax: +44-20-7388-9325,
e-Mail: h.demeer@ee.ucl.ac.uk

Description:

Over the last three years, peer-to-peer (p2p) services [Oram01] [Barkai02] have evolved to one of the most important sources of Internet traffic. P2P file sharing protocols caused more than 25% of the traffic volume and more than 20% of the flows at some core routers in the US (as of November 21st 2002) [Sprint02]. P2P file sharing is challenging web traffic for the top spot of traffic sources in the Internet [Internet2]. Rarely before has a new class of network services evolved in such a short time as p2p services have done.

The impact of the high traffic load of today's p2p applications is already evident. Network operators complain about high traffic volume and users about decreasing reliability and attractiveness. In order to cope with this, it is necessary to implement an efficient performance and traffic management. From the user perspective, p2p performance management should provide a consistent quality of locating and accessing resources as well as minimize the adverse effects of p2p on other network services. From the operators' perspective, traffic management should facilitate an optimal traffic distribution and efficient network operation.

A main characteristic of p2p services is their highly distributed, server-less architecture using autonomous *peers*. P2P services trade the client/server paradigm for decentralization by extensive use of communication. P2P services typically form application-specific virtual network structures, so-called *overlays*. The peers may join or leave the overlay arbitrarily. The attractiveness of p2p services increases exponentially with the number of peers contributing to the service, but also the amount of signaling traffic increases almost exponentially with the number of participants.

The possibilities for performance and traffic management of p2p services appear to be limited at first. However, new approaches for p2p management on Application Layer Active Networking and dynamic overlay operation appear to be promising. This tutorial will provide an introduction into p2p computing and discusses the possibilities of performance and traffic management of p2p services on network layer and on application layer.

The tutorial is divided into seven sections. Section 1 provides an overview of the landscape of p2p services and outlines their aims, the basic terminology, and architectural classes. This is followed by Section 2, where the main components and communication mechanisms of p2p services are introduced. The third part of the tutorial deals with the technical features of p2p services and Section 4 presents selected P2P applications. In Section 5 we analyse p2p services and in Section 6 we focus in depth on new results in performance and traffic management for p2p services. Section 7 concludes the tutorial with a short summary.

Literature:

Recent publications of the authors on p2p services:

H. de Meer, K. Tutschku: *A Performance Management Architecture for Peer-to-Peer Services on Application-level Active Networks*. Proceedings of the Network Operation and management Symposium 2002 NOMS 2002, Florence, April 2002.

K. Tutschku, H. de Meer: *A Measurement-based Investigation on Signaling in Gnutella Overlay Networks – Providing Insights to Third Generation Internet Applications?* Proceedings of the 1th Workshop Design and Performance of 3G Internet Technologies collocated with MASCOTS 2002, Forth Worth, October 2002.

H. de Meer, K. Tutschku: *Dynamic Operation of Peer-to-Peer Overlays*. Proceedings Fourth Annual International Working Conference on Active Networks, ETH Zürich, Switzerland, Dec. 4-6 2002.

K. Tutschku, H. de Meer: *A Measurement Study on Signaling on Gnutella Overlay Networks*. Proceedings of “Fachtagung – Kommunikation in Verteilten Systemen (KiVS) 2003”, Leipzig, February 2003.

H. de Meer, K. Tutschku, P. Tran-Gia: *Dynamic Operation in Peer-to-Peer Overlay Networks*. Praxis der Informationsverarbeitung und Kommunikation - Special Issue on Peer-to-Peer Systems (PIK Journal), June 2003.

General publications on p2p services:

[Oram01] A. Oram (Ed.): *Peer-to-Peer – Harnessing the Power of Disruptive Technologies*, O’Reilly, Sebastopol, CA, 2002.

[Barkai02] D. Barkei: *Peer-to-Peer Computing*, Intel Press, Hillsboro, OR, 2002.

Measurements of p2p services:

[Internet2] information available at <http://netflow.internet2.edu/>

[Spint] information available at <http://ipmon.sprintlabs.com/>

Biography of the Authors:

Kurt Tutschku is an Assistant Professor at the Chair of Distributed Systems, University of Würzburg. He has received a doctoral degree from University of Würzburg in 1999. Kurt Tutschku is leading the department’s group on network management and future network services. His research interest include event- and traffic-oriented management of IP-based communication networks, planning and optimization of telecommunication systems, demand-oriented design of future generation, large scale IP networks, methods for the estimation and characterization of the spatial tele-traffic distribution in mobile communication networks, design of algorithms and tools for demand based computer aided mobile network planning, and the modeling of future network services in third generation internet, particular of peer-to-peer systems.

Additional information is available at:

<http://www3.informatik.uni-wuerzburg.de/staff/tutschku>

Hermann de Meer has led several national and international projects in modeling and performance evaluation, data communications, and quality of service. He has been an Assistant Professor at Hamburg University, Germany, a Visiting Professor at Columbia University in New York, USA, and a Research Fellow of the Deutsche Forschungsgemeinschaft (DFG). He is currently appointed as Reader for communication networks and distributed systems at University College London and as Visiting Professor at Karlstad University, Sweden. Dr. Hermann de Meer has received his doctoral degree with distinction from the University of Erlangen-Nuremberg, Germany, has been received an MMB best-paper award, and is co-authoring a well-cited book on Queuing Networks and Markov Chains – Modeling and Performance Evaluation with Computer Science Applications published by John Wiley in 1998.

Additional information is available at:
<http://www.ee.ucl.ac.uk/netdist/>

Preliminary Contents: Network Efficient P2P-Services and their Management

1. Introduction: Overview of p2p services
 - Aims, terminology, architecture classes
2. Components, structures and algorithms of p2p services
 - a. P2P and Distributed Systems
 - b. P2P and Overlay Networks
 - c. P2P and Self-Organization
 - d. P2P and Ad-Hoc Networks
3. Technical Features of p2p services
 - a. Topology and topology models
 - Group Forming Networks
 - Small World and Scale Free networks
 - b. Signaling:
 - Routing and forwarding strategies in p2p overlays
 - Comparison with Ad-Hoc networks
 - c. Efficient resource management
 - Multiple source downloads: *Swarming*
 - b. Distributed Hash Tables (DHT) and Content Addressable Networks (CAN)
4. Case Studies:
Presentation of selected p2p applications and p2p frameworks such as Seti@Home, Gnutella, Kazaa, EDonkey2000, Freenet, JXTA and .Net
5. Analyses of p2p services
 - a. Decentralization and scalability
 - b. Security, anonymity and trust
 - c. Fault tolerance and resilience
 - d. Performance and self-organization
 - e. P2P traffic measurements
 - f. P2P traffic models
6. Performance and Traffic Management:
 - a. Overview on *Network-layer* traffic management:
 - Port throttling, shapers and firewalls
 - Network layer based traffic engineering
 - b. A case study in *Application-layer* performance management:
 - Application Layer Active Networking (ALAN), dynamic overlay networks
 - Application-layer based load balancing
7. Conclusions