

06131 Abstracts Collection
Peer-to-Peer-Systems and -Applications
— Dagstuhl Seminar —

Anthony D. Joseph¹, Ralf Steinmetz², Ion Stoica³ and Klaus Wehrle⁴

¹ UC Berkeley, US

adj@eecs.berkeley.edu

² TU Darmstadt, DE

ralf.steinmetz@kom.tu-darmstadt.de

³ UC Berkeley, US

⁴ Univ. Tübingen, DE

Klaus.Wehrle@uni-tuebingen.de

Abstract. From 26.03.06 to 29.03.06, the Dagstuhl Seminar 06131 “Peer-to-Peer-Systems and -Applications” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Peer-to-Peer, self-organisation, massively distributed systems

06131 Executive Summary – Peer-to-Peer-Systems and -Applications

This is the executive summary of Dagstuhl Seminar 06131, "Peer-to-Peer-Systems and -Applications". The seminar was held from March 26th to March 29th, 2006, at the International Conference and Research Center for Computer Science at Castle Dagstuhl, Germany.

Keywords: Peer-to-Peer, self-organisation, massively distributed systems

Joint work of: Joseph, Anthony D.; Steinmetz, Ralf; Stoica, Ion; Wehrle, Klaus

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/852>

06131 Summary of the Workgroup "Information Retrieval in Peer-to-Peer Systems"

Today's information provisioning in peer-to-peer systems has to deal with several types of information. First there are data files or textual content for sharing that can be more or less described by adequate metadata. In contrast there is process data, which is continuously created, like for instance log information, intermediate results or state information. This data is essential in supporting long-running applications, where users interact with the infrastructure over a significant time-span, e.g., media streaming for IPTV or online experiments. Such data can easily fill 5-10 MB of sequential, non-indexed data files and has to be monitored or searched at run time to detect events that enable the intelligent adaptation of parameters for quality of service (QoS) control or graceful recovery from failures.

The type information in such process data, however, is quite diverse ranging from simple structured information (like CPU usage of a certain node, delays within the network or concise information about a video frame, etc.) to largely unstructured data (like verbose failure descriptions or newly discovered service functionalities). Keeping this data readily accessible is a crucial challenge. Major problems are on one hand the creation of light-weight structures for logging or keeping intermediate states, which allow supporting IR-style retrieval especially with respect to online-algorithms and real-time IR techniques. Moreover, the support of functions for aggregation of data, finding correlations in data and conditional expressions, as well as effectively indexing (highly discriminating) terms or phrases and ranges of values. On the other hand, peers need to offer this information (possibly in different resolutions) and all information needs to be persistent even in the presence of churn (i.e. a certain degree of replication is necessary) to allow for graceful recovery or checkpointing.

For these ends lightweight Ontologies and dictionaries, as well as indexed aggregated information in terms of ranges or sets (for instance kept in a DHT structure) can be expected to be useful. This is primarily because for different applications they enable the explicit formulation of a trade-off between the performance and necessary precision of queries.

Keywords: Information retrieval peer-to-peer P2P challenges

Joint work of: Balke, Wolf-Tilo; Nahrstedt, Klara; Niedermeyer, Heiko

Information Retrieval in Peer-to-Peer Systems

Wolf-Tilo Balke (L3S Research Center and Hanover University, D)

This talk focuses on the challenges of applying information retrieval techniques in peer-to-peer infrastructures. Peer-to-peer systems are already being used for a vast number of applications in content exchange, but mostly searching is done by simple keyword lookups.

In contrast, information retrieval means that not only some more or less matching objects have to be retrieved, but a list of the best matching objects over the entire network given a user's information needs. Since the 1960ies the information retrieval community considers ways to efficiently and effectively query document collections and designs dedicated retrieval systems like e.g. SMART.

Usually a query is seen as a (possibly weighted) set of keywords a user specifies to express his/her information need. Documents that contain those (or sufficiently similar) keywords are considered to be relevant to the user's information need as expressed by the query. Thus, for information retrieval in peer-to-peer infrastructures the challenge is not only to retrieve documents efficiently, but also to effectively find a set of best matching objects. Generally speaking in this scenario retrieval effectiveness can only be traded for improved efficiency to a very limited degree. Moreover, in IR scenarios collection-wide statistical information is heavily used to improve retrieval effectiveness clashing with the very distributed nature of peer-to-peer scenarios. During the talk we will discuss the main issues and point to possible solutions and some ongoing research.

Keywords: Information retrieval peer-to-peer P2P challenges

IR in P2P Workgroup

Wolf-Tilo Balke (L3S Research Center and Hanover University, D)

Information Retrieval techniques vastly differ from retrieval in pure data management systems. But until very recently P2P techniques focussed on supporting simple database-style retrieval functionality only. Since the content offered in P2P networks currently evolves from simple files that can be adequately described by some meta-data tags, to collections of complex documents that need somewhat more detailed descriptions, the development of refined retrieval techniques has to be fostered.

The aim of this working group was to

- identify basic differences and characteristics of metadata-driven vs. content-driven retrieval in P2P networks
- discuss the practical usefulness (e.g. in terms of effectiveness, scalability, etc.) of current P2P techniques for indexing, routing, etc. in the light of sophisticated content-based searches
- spotlight today's most important research questions and novel techniques that might provide viable solutions

Keywords: Information retrieval peer-to-peer P2P challenges

Joint work of: Balke, Wolf-Tilo; Nahrstedt, Klara; Niedermeyer, Heiko

Fast Stochastic Exploration of Tree-based Content Distribution Architectures

Ernst Biersack (Eurecom - Sophia Antipolis, F)

We consider the problem of distributing a content of finite size to a group of users connected through an overlay network that is built by a peer-to-peer application. We examine and compare different distribution architectures based on linear and tree topologies built on top of the P2P overlay, including in the analysis the presence of heterogeneous bandwidths, both symmetric and asymmetric access links.

We propose an analytical solution of the distribution process that not only yields the mean download time but also the distribution of the download times. We validate the analytical model against a Monte Carlo based numerical solution, which can also be used to analyze scenarios where correlation and dynamic behavior make the theoretical analysis too approximate. The work presented is joint work with D. Carra and R. Lo Cigno from University of Trento, Italy

Keywords: Peer to peer, performance evaluation

Reputation based computational models of trust in P2P networks

Zoran Despotovic (DoCoMo Euro-Labs - München, D)

Reports about a project targeting on reputation management by aggregating opinions about any particular peer and make trust scores to be used when interacting with the concerned peer to provide a convergence of the system toward the trustworthy behavior of the P2P system.

Keywords: P2P, reputation, trust, feedback, agents

Joint work of: Despotovic, Zoran; Kellerer, Wolfgang

Fairness in Peer-to-Peer Networks

Kolja Eger (TU Hamburg-Harburg, D)

The first Peer-to-Peer (P2P) networks were based mainly on the altruistic behaviour of its peers. Although newer implementations incorporate some kind of incentive mechanism to award sharing peers, no P2P network assures some quality of service.

Our work is meant as a first step towards the development of P2P networks with quality of service. We propose a distributed resource allocation algorithm where peers control the service rate to its neighbours. This algorithm is based

on the congestion pricing principle known from IP networks and ensures some form of fairness. Hence a peer gets a fair share of the resources available in the P2P network weighted by its contribution to the network.

We present the first simulation results about the convergence of our algorithm and its functionality in large and varying networks.

Keywords: Pricing, Rate Control, Resource Allocation, Distributed Optimisation, Fairness, P2P Networks

Joint work of: Eger, Kolja; Killat, Ulrich

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/645>

Trust Matters: Formation of Peer-to-Peer Networks

Kai Fischbach (Universität Köln, D)

We investigate the role of trust in the formation of peer-to-peer (p2p) networks. Our approach is based on the work of Bala and Goyal [1], Jackson and Wolinsky [2] and Watts [3]. Unlike and in addition to these models, we take into consideration that agents may establish links with other agents not only based on cost-benefit assessments but also based on trust. We present an agent-based simulation model to analyze the dynamic process of network formation and to evaluate which structures emerge and how efficient these structures are. Our results show that (i) even p2p communities with a subset of untrustworthy nodes can learn to coordinate, (ii) larger communities have a higher ability to deal with untrustworthy peers than smaller communities, and (iii) a periphery-sponsored star is a stable state.

Literature:

- [1] Bala, Venkatesh und Sanjeev Goyal: A Noncooperative Model of Network Formation. *Econometrica*, 68(5): 1181-1229, 2000.
- [2] Jackson, Matthew O. und Asher Wolinsky: A Strategic Model of Social and Economic Networks. *Journal of Economic Theory*, 71(1): 44-47, 1996.
- [3] Watts, Alison: A Dynamic Model of Network Formation. *Games and Economic Behavior*, 34(2): 331-341, 2001.

Keywords: Trust, coordination, learning dynamics, network formation

Scalable Routing for Embedded Networks and Mobile P2P Applications

Thomas Fuhrmann (Universität Karlsruhe, D)

Scalable source routing (SSR) is a full-fledged routing protocol that directly provides the semantics of a structured peer-to-peer overlay.

It combines source routing in the physical network with Chord-like routing in the virtual ring formed by the address space. Message forwarding greedily decreases the distance in the virtual ring while preferring physically short paths. Unlike previous approaches, scalability is achieved without imposing artificial hierarchies or assigning location-dependent addresses. Extensive simulation studies demonstrate that SSR outperforms mobile ad hoc protocols like, for example, AODV. By virtue of its structured peer-to-peer overlay semantics, SSR can hence directly serve as an efficient basis for fully decentralized peer-to-peer applications on mobile devices.

Keywords: Cross-layer, Mobile P2P, Embedded Networks, Sensor-Actuator Networks, Mesh Networks

P2P Communication in Massive Multiplayer Games

Carsten Griwodz (University of Oslo, N)

Online multiplayer games are a multi-million industry that is anticipated to grow even more. The games can be classified roughly in three genres, one of them being role-playing games (RPGs). With more than 300.000 concurrently active participants in a single, latency-critical, interactive application (Lineage II), these distributed applications are quite resource-demanding.

Currently, all commercial RPGs are based on a centralized infrastructure. However, player's demand for even higher interactivity and for in-game streaming media communication makes this infeasible. We have therefore begun an investigation into the maintenance of dynamic peer groups that implement overlay multicast among each other and avoid communication with the central servers. Although the results are preliminary, we can see the need to trade QoS parameters such as overall bandwidth consumption and pair-wise maximum latency against each other. Due to the necessary trade-offs, several overlay networks with different optimization conditions are required for the various traffic types that comprise the needs of an RPG. The talk will present some early results.

Keywords: Multiplayer games, QoS, overlay multicast

Quality in P2P Systems

Oliver Heckmann (TU Darmstadt, D)

In this talk, we present the work of the newly established DFG Forschergruppe QuaP2P (see www.quap2p.de). The goal of this research group is to increase the quality of peer-to-peer systems and to compare the achievable quality is centralised solution like client-server applications.

Keywords: Quality of service, p2p

Joint work of: Steinmetz, Ralf; Heckmann, Oliver

Full Paper:

<http://www.quap2p.de>

Mobile Peer-to-Peer teamwork

Wolfgang Kellerer (DoCoMo Euro-Labs - München, D)

Mobile and wireless systems can especially benefit from the self-organizing characteristics of P2P systems, e.g., with respect to robustness, however the heterogeneity and dynamicity of those systems also poses severe challenges to P2P concepts. Moreover exploiting mobile systems features, new applications for P2P systems are opened such as location based services. The degree of controllability as an operator requirement versa self organization has to be considered as another important challenge for mobile P2P approaches. This focus group is going to elaborate on wether the current state of the art in P2P solutions is sufficient for mobile systems and discusses open research issues as well as emerging application classes.

Core - A Peer-To-Peer Based Connectionless Onion Router

Olaf Landsiedel (Universität Tübingen, D)

The talk presents our work on anonymous communication: It introduces "CORE - Connectionless Onion Router". Today's low-delay onion routers such as Tor or Tarzan build a static tunnel through a peer-to-peer mix network for each connection using layered encryption. As all the connection's traffic uses the same tunnel it is susceptible to attacks based on pattern analysis as recent publications show. Compared to today's systems, CORE routes each packet through a different communication path and so is not susceptible to this class of attacks. In this talk we describe the design of our connectionless onion router, evaluate its performance and address the communication overhead. Furthermore, we present address virtualization to abstract from the user's and system's identity and to provide transparent application support. Thus, no application level gateways or proxies are required to strip away the user's identity from the communication.

Keywords: Anonymity, Onion Routing, Connectionless Onion Routing

Assessing Peer Reliability in Cluster based Overlay Networks

Andreas Mauthe (Lancaster University, GB)

In order to create a stable structured overlay using clusters it is crucial that the peers in the cluster are well distributed and take a role according to their capabilities. To do this their reliability and potential up-time has to be assessed while considering the costs and benefits of assigning a specific role at a specific time. This is done using a burn-in scheme and calculating the potential costs and benefits of this process

Keywords: Uptime distribution, role specialisation, burn-in, conditional reliability

Utilizing P2P Overlay Multicast using standard IP-Multicast Applications

Dragan Milic (Universität Bern, CH)

IP Multicast is designed to decrease the network load by eliminating redundancy of data transfers from one server to a group of receivers. This redundancy elimination makes IP Multicast ideal for disseminating multimedia contents such as audio and video broadcasts.

However, at the time, IP Multicast is not supported by most of the commercial Internet service providers (ISPs). There are numerous solutions (such as MBONE) proposed to bridge the gaps in IP Multicast routing in the Internet. The administrative overhead of the existing solutions makes them unavailable for the typical end-user accessing the Internet via xDSL or TV cable.

Application Layer Multicast (ALM) using peer-to-peer overlay networks could solve the problem of sparse IP Multicast support in the Internet. A limitation of this approach is the lack of standardized interfaces for existing IP Multicast applications. We propose a solution, which bridges ALM and IP Multicast and uses a peer-to-peer network to transport Multicast data. Our solution – including a 'proof-of-concept' prototype implementation – demonstrates the feasibility of the proposed solution.

We present an example application of our solution that enables video broadcasting over the Internet using existing IP Multicast applications without requiring additional service deployment or updating existing software.

We also demonstrate integration of our solution with an existing ALM protocol NEMO.

Keywords: Overlay-Networks, Multicast, Middleware

Joint work of: Milic, Dragan; Brogle, Marc; Braun, Torsten

Non-Remuneration in P2P Incentive Mechanisms

Jens Oberender (Universität Passau, D)

Incentive mechanisms protect P2P networks from free-riders. Pricing controls the cost and benefit of cooperation, e.g., the granted upload-download ratio. Some file-sharing networks enforce cooperation by credit-based mechanisms. The locally-stored credit value rates the willingness to grant service to a certain peer. As an incentive, future requests of cooperative peers are accelerated. We examine the correlation of available bandwidth, received data rate, credits, and pricing. In our results we identify the limitations of this incentive mechanism, especially when selfish nodes deny cooperation to maximize their benefit.

Keywords: Incentive, Cooperation, Filesharing, Peer-to-Peer

Peer-to-Peer-based Infrastructure Support for Massively Multiplayer Online

Simon Rieche (Universität Tübingen, D)

Massively multiplayer online games (MMOGs) are becoming increasingly popular today. However, even high-budget titles like World of Warcraft that have gone through extensive beta-testing suffer from downtimes because of hardware and software problems and lags on overloaded servers. Our approach is to use structured P2P technology for the server infrastructure of massively multiplayer online games to improve reliability and scalability of these applications.

Structured P2P networks are able to adapt to the current state of the game world and handle uneven distributions of the players. For load balancing, we propose algorithms based on the virtual server concept.

Taming Dynamic and Selfish Peers

Stefan Schmid (ETH Zürich, CH)

Peer-to-peer systems are often faced with the problem of frequent membership changes. However, many systems are only proven efficient or correct in static environments. In my talk, I will present techniques to maintain desirable properties of a distributed hash table (low peer degree, low network diameter) in spite of ongoing and concurrent dynamics. I will then go on and study the effect of peers not acting according to our protocols. Concretely, I assume that peers are selfish and choose the behavior which maximizes their utility. I will report on our results concerning the impact of selfishness on the peer-to-peer topology.

Keywords: Churn, Selfishness, P2P Topologies

Joint work of: Schmid, Stefan; Moscibroda, Thomas; Wattenhofer, Roger; Kuhn, Fabian

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/647>

On the Topologies Formed by Selfish Peers

Stefan Schmid (ETH Zürich, CH)

Many P2P systems are only proven efficient for static environments. However, in practice, P2P systems are often very dynamic in the sense that peers can join and leave a system at any time and concurrently. In the first part of my talk, I will present a DHT we have developed recently in our group which maintains desirable properties under worst-case churn. In the second part of my talk, we will briefly look at another challenge of prime importance in P2P computing, namely selfishness. Concretely, some results are presented concerning the impact of selfish behavior on the performance of P2P topologies.

Keywords: Churn, Selfishness, P2P Topologies

Joint work of: Schmid, Stefan; Moscibroda, Thomas; Wattenhofer, Roger

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2006/643>

Siemens Peer-to-Peer Technologies and their Industrial Application

Alan Southall (Siemens - München, D)

This presentation gave a brief overview of the Siemens Resource Management Framework (RMF), the RMF Service Platform (RSP) and various industrial applications which have been based on these systems, e.g. VoIP clients and fault tolerance. The RMF aims to provide a generalized API and set of semantics for peer-to-peer protocols, e.g. DHT, unstructured or central indexes. Building on the RMF, the RSP provides an OSGi environment for deploying web-services which can be used by other peers in the network; the RMF provides NAT traversal for all peer-to-peer communication. In addition to the technical issues, disruptive business models which are being employed by the Internet community were discussed.

Full Paper:

http://p2p.ingce.unibo.it/2002/NonRevisedPapers/11_rusitschka.pdf

Hybrid DHT Design for Mobile Communications Environments

Stefan Zöls (TU München, D)

We propose a hybrid design concept for Distributed Hash Tables (DHTs) in order to increase the performance of DHTs in scenarios with mobile participants. By defining two classes of nodes (static and temporary) and assigning critical overlay networking tasks to reliable static nodes our concept allows the disburdening of resource-constraint temporary nodes such as PDAs or mobile phones. Further we show simulation results that prove the significant advantages of our extension in comparison to conventional DHTs.

Keywords: Peer-to-Peer, DHT, Mobile, Hybrid

See also: Zöls S., Schubert S., Kellerer W., Despotovic Z.; Hybrid DHT Design for Mobile Environments; AP2PC 2006.