Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences

AW1 Ausarbeitung

Christian Vogt

Peer-2-Peer Networking using Open Web Technologies

Fakultät Technik und Informatik
Studiendepartment Informatik

Faculty of Engineering and Computer Science
Department of Computer Science

Christian Vogt

Peer-2-Peer Networking using Open Web Technologies

AW1 Ausarbeitung eingereicht im Rahmen der Veranstaltung MINF-AW1

im Studiengang Master of Science Angewandte Informatik
am Department Informatik
der Fakultät Technik und Informatik
der Hochschule für Angewandte Wissenschaften Hamburg

Eingereicht am: 15. Februar 2013

Contents

1 Introduction

In this paper, the basic principles of Peer-To-Peer systems and modern Web Technologies will be introduced, proceeding as follows. First of all, each topic will be discussed separately. Afterwards, possible overlapping issues of both topics will be pointed out, focussing on the opportunities that can result from converging them.

What is *Peer-To-Peer* about?

A definition of a *Peer-To-Peer* (also referred to as *P2P*) system is given by [Oram 2001]:

> "A P2P system is a self-organizing system of equal, autonomous entities (peers) [which] aims for the shared usage of distributed resources in a networked environment avoiding central services."

This implies that, instead of relying on a central entity, every node acts as a client **and** server (*servent*[1]) to add its computing resources to the system.

The Peer-To-Peer paradigm for Internet-based applications was originally designed for file sharing applications like Napster or Gnutella but is applicable to any kind of system that depends on distributed resources like Content Distribution Networks or Internet telephony. [Steinmetz und Wehrle 2005] identified three main requirements of future Internet-based applications.

- *Scalability* allows a system to scale by several orders of magnitude without the loss of efficiency by eliminating bottlenecks caused by the systems design.
- *Security* and *Reliability* are of high importance for Internet-based applications where they are facing DDoS[2] attacks and consumer frustration when not constantly available.
- *Flexibility* and *Quality of Service* requirements form core criteria for modern Internet-based applications to accommodate features they were not conceived for.

[1]Servent is a portmanteau for **serv**er and cli**ent**
[2]Distributed Denial of Service attacks aim at rendering a remote system unusable by flooding it with requests

In comparison to the classic client-server based approach which can no longer fully cope with these main requirements due to its centralized structure, Peer-To-Peer systems try to provide new solutions to the stated problems, especially focusing on scalability issues.

What is Web about?

Tim Berners-Lee et al. introduced the *HTTP* protocol in 1996 [Berners-Lee u. a. 1996], which still forms the basis of the current *World Wide Web* (*WWW*, *Web*). The ongoing evolution of the Web manifests itself in the development of the *HTML5* specification by the W3C consortium which is scheduled for a final release in 2014[3]. The Web is often confused with the Internet as a whole but according to the W3C consortium ([Jacobs und Walsh 2004]) it is defined as follows:

> "The World Wide Web (WWW, or simply Web) is an information space in which the items of interest, referred to as resources, are identified by global identifiers called Uniform Resource Identifiers (URI)."

While the Internet is a global system of interconnected computer networks (thus, a network of networks), the Web is an information space on top of the Internet.

Coherence

The classic request/response approach which is being pursued by HTTP does not always fit the modern Web's use cases. New technologies are being developed within the HTML5 specification in order to allow a direct peer to peer connection between browsers without the need for plugins. These technologies (e.g. *WebRTC*), when fully implemented in the near future, reveal new opportunities to the ecosystem of the Web and a rewarding coherence between the Web and the Peer-To-Peer paradigm.

A more detailed view on the specifications and research efforts of both Peer-To-Peer and Web will be presented in chapter 2.

[3] http://dev.w3.org/html5/decision-policy/html5-2014-plan.html#introduction

2 Technologies And Research Areas

In this chapter the technologies and research areas of Peer-To-Peer systems (chapter 2.1) and Web technologies (chapter 2.2) will be presented.

2.1 Peer-To-Peer Systems

The idea and core requirements of the Peer-To-Peer paradigm which have been introduced in chapter 1 will now be explained from a technical point of view. An investigation into current areas of research follows.

Evolution of Peer-To-Peer systems

Peer-To-Peer systems can be classified in two main generations as shown in figure 2.1.

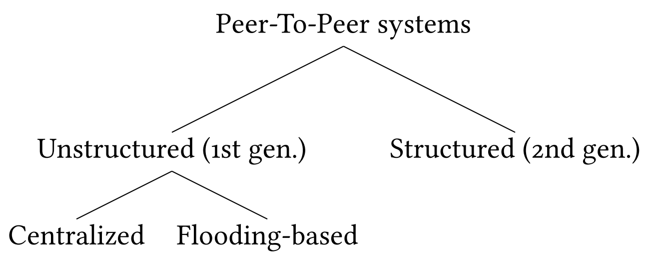

Figure 2.1: Overview of first/second generation Peer-To-Peer systems

First generation Peer-To-Peer systems suffer from the same scalability problems that occur in the classic client/server approaches as they require a central server for lookup and management purposes (e.g. Napster[1]) or use flooding-based approaches (e.g. early versions of Gnutella[2]) with tremendous bandwidth requirements between the peers. It is apparent that the straightforward unstructured approaches, mostly used for file-sharing and instant messaging purposes, suffer from the lack of providing a solution for the requirements introduced in chapter 1.

[1] http://www.napster.com/
[2] http://rfc-gnutella.sourceforge.net/developer/stable/index.html

However, the idea of completely decentralized, self-organizing applications sparked interest in the research community and the second generation Peer-To-Peer systems were developed in response to that. These *Structured Peer-To-Peer systems* are typically based on the idea of the *Distributed Hash Tables* (*DHT*) which will now be further elaborated on.

Structured Peer-To-Peer systems

A typical application that incorporates a structured Peer-To-Peer system (also known as *Structured Overlay Application*) depends on a Distributed Hash Table which, from the application's point of view, can be used like a traditional hash table. Figure 2.2 shows the architecture of such a system and illustrates the DHT's interface:

To store a value, the `put()` operation is used in conjunction with a hash function that calculates the corresponding key for the value. This value will be stored at a corresponding peer that is in charge for a specific key space. The `get()` operation is used to look up a value. While this is transparent to the application, the DHT-layer is responsible for efficiently routing the request to the corresponding peer. In such a system, every peer is assigned a distinct ID out of a key space by using hash functions like *SHA-1*. Content, that is to be stored in the DHT can now be allocated to the peer that is in charge of the specific key. Thereby, logarithmic behavior can be assured for routing complexity and state (Steinmetz und Wehrle [2005]).

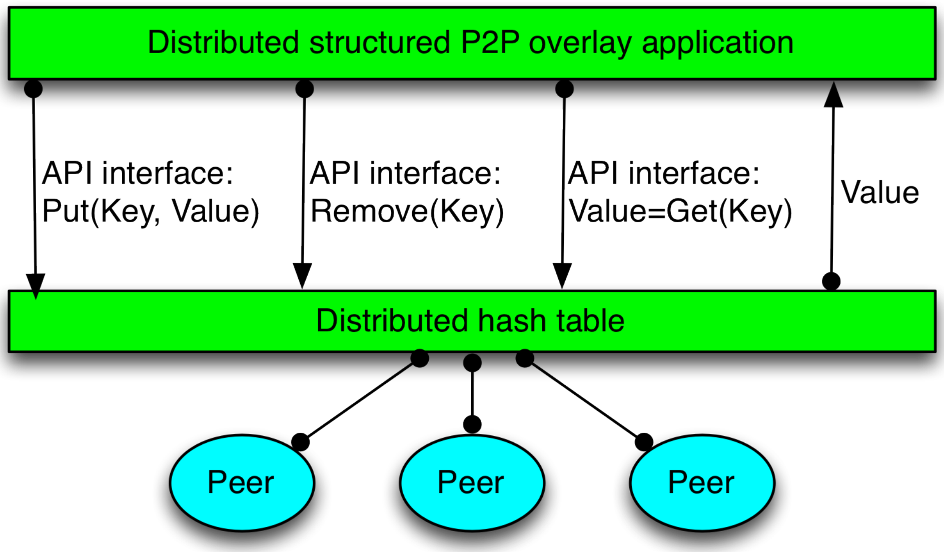

Figure 2.2: Application interface for structured DHT-based Peer-To-Peer overlay systems [Lua u. a. 2005]

DHTs are typically based on similar designs like ring- (e.g. *Chord* [Stoica u. a. 2001], or *Pastry* [Rowstron und Druschel 2001]) or geometry-based approaches (e.g. *Content Adressable Networks* [Ratnasamy u. a. 2001]). The difference between the approaches are the search and management strategies, as well as use case specific enhancements like geographical aspects for routing decisions or data redundancy. A more detailed view on the different DHTs can be found in [Lua u. a. 2005].

Application Layer Multicast

Apart from the obvious usage for distributed content storage, DHTs provide the ability to support *Application Layer Multicast* (*ALM*). ALM provides multicast features to applications without any needed support from the network except for the elementary unicast forwarding service.

The publish/subscribe system *Scribe* ([Rowstron u. a. 2001]) is built on top of Pastry and illustrates such an ALM system. Other approaches like *Splitstream* ([Castro u. a. 2003]) or *BIDIR-SAM* ([Wahlisch u. a. 2009]) focus on high-bandwidth distribution and, in this respect, highly flexible network adaptivity.

Research Interests

While the fundamentals of Peer-To-Peer systems have been known for years, in this case it is important to focus on very recently released publications in order to identify the current research areas of Peer-To-Peer systems. These can still be classified based on the identified requirements of Peer-To-Peer applications stated in chapter 1 which have been introduced by [Steinmetz und Wehrle 2005] back in 2005. Afterwards, an overview of conferences and research groups will be presented.

Scalability

[Vijendran und Thavamani 2012] present a survey of the current research efforts on replication and caching algorithms for content distribution in Peer-To-Peer overlay networks. In a DHT, a single peer is responsible for a specific key space. If that peer fails, other peers have to incur the requests, thus implementing a replication algorithm like [Zaman und Grosu 2011].

Security and reliability

The publication *Impact of Feedback on Trust in P2P Networks* by [Huang u. a. 2012] introduces a feedback-based system that is able to implicitly determine the peer's characteristics to find out if it can be trusted. This is especially important for a Peer-To-Peer system because of its decentralized nature where malicious participators are more prevalent than in centralized approaches.

Quality of Service

When operating a distributed system relying on clients while having no direct control over these, giving guarantees in terms of Quality of Service requirements is a big challenge. [Yang und Lou 2012] propose a contract-ruled model for these guarantees that uses an economic model where each peer is contracted to receive a penalty if it fails to meet the systems requirements.

Flexibility

[Bakht u. a. 2012] try to enhance mobility aware Peer-To-Peer Overlay Networks to mitigate the problem of efficient neighborhood discovery using geographical proximity. As long-range radios such as Wi-Fi can be too expensive in terms of power demand, a cluster using low-power radios (e.g. Bluetooth) is established and distributes the load of high-power, long-range scanning.

Conferences and Researchers

Even though the topic of Peer-To-Peer computing is subject to numerous research approaches, two conferences that solely focus on Peer-To-Peer are mentioned here.

The **IEEE P2P** is considered the oldest and largest conference dedicated to Peer-To-Peer computing and not only focuses on the theoretical basics but also brings in application developers and Internet Service Provider (ISP) executives to the table. *P. Druschel* from the *Max Planck Institute for Software Systems* is last year's keynote speaker and recognised for the widely used structured Peer-To-Peer system *Pastry*.

The **HotPOST**'s current "Hot Topic" clearly are *Online Social Networks* (*OSN*) and related subjects like Peer-To-Peer microblogging. *M. Wählisch* from the *Institute of Computer Science* at the *Freie Universität Berlin*, who is one of the authors of the Application Layer Multicast solution *BIDIR-SAM*, holds this years publicity chair.

2.2 Open Web Technologies

Unlike the Peer-To-Peer topic that is mostly driven by non commercial stakeholders like universities, a lot of Web research is completed behind the closed doors of the browser vendors. However, the W3C standards organization, headed by Tim Berners-Lee, exists and tries to enforce compatibility among the mostly industry-related members. This subchapter first gives an insight into the status quo and follows up with current research topics. Instead of elucidating Web's history (Raman [2009] gives an detailed overview), this paper focuses on the new possibilities introduced by HTML5.

New HTML5 standards

The document-oriented nature of the Web is evolving more and more into a rich software platform. [Anttonen u. a. 2011] summarizes the new trends and technologies and concludes, that the vast majority of future end user software will be developed using web technologies. The keychain of new technologies resides under the HTML5 umbrella and consists of several key components of which a handful will now be further elaborated upon.

One of the more prominent new features of the upcoming technologies is the **WebRTC** specification ([Fette und Melnikov 2011]) which allows for direct peer to peer connections between browsers and enables *Real-Time Communications* (*RTC*) capabilities for e.g. video conferencing or IP Telephony.

The **Geolocation API** ([Popescu 2009]) allows a user to provide their geographical location to the web application. By automatically determining the current position of the user, applications can establish location-aware mash-ups or use the geographical information for efficient neighbor discovery.

A client-side storage that is persistent across restarting a web application has long been a desired feature. The **IndexedDB** ([Mehta 2010]) is able to store significant amounts of structured data and allows for efficient searches on the data.

There are more new features in HTML5[3], but the ones presented show the ongoing evolution of Web technologies and subsidize the assumptions made by [Anttonen u. a. 2011].

[3] https://github.com/SirPepe/SpecGraph gives an overview

Research Interests

The benefits from the new possibilities manifest themselves in new topics in the research community. Several publications released very recently are presented below. These are based on the technologies introduced.

[Wang u. a. 2012] tries to improve **mobile** browser speeds by speculatively loading websites on WebKit-based browsers.

The trend towards user-generated web content has an impact on the content distribution **performance** when replicated to Points of Presence around the world (CDN approach). To save bandwidth, [Traverso u. a. 2012] propose the *TailGate* approach to use geographical closeness for enhanced efficiency.

[Erbad u. a. 2012] demonstrates how HTML5 capable browsers can be used for latency sensitive and processing intensive **Real Time** Applications on multi-core CPU setups.

Currently, HTTP traffic can be secured using the HTTPS protocol. Unfortunately, this is only possible if no intermediate servers (e.g. caching proxies) are used. Therefore [Singh u. a. 2012] proposes *HTTPi* for end-to-end **security** while still leveraging the benefits of intermediate caching.

[Miranda u. a. 2012] explains the need to avoid the vendor lock-in effect in the deployment of **cloud-based** applications. The risk to hinder interoperability is high because no standardization has been introduced to the use and development of cloud-based services.

[Meij u. a. 2012] uses the principles of semantic Web approaches to connect Twitter-like microblog postings and use them for marketing, searches, information dissemination, and online reputation management. An algorithm based on machine learning is used to filter out non-informative postings from **Social Networks**.

Conferences

The broad range of web related subjects leads to a vast number of conferences. Two conferences that are not focused on special subjects and therefore provide a decent overview, are the *International World Wide Web Conferences*[4] and the *International Conference on Web Services*[5]. The former debates about the evolution of the Web (e.g. performance, security) and was conceived by Tim Berners-Lee, whereas the later is directed more to enterprise use cases (e.g. Software As A Service).

[4]http://wwwconference.org/
[5]http://www.icws.org/

3 Opportunities And Approaches

Having provided a detailed insight on both the Peer-To-Peer paradigm and modern Web technologies, an outline on how to combine these two paradigms to leverage benefits for a new type of application will be presented in this chapter.

Where are Peer-To-Peer Applications today?

As of today, applications based on the Peer-To-Peer paradigm are not very widespread except for file sharing (e.g. BitTorrent) or proprietary communication purposes (e.g. Microsoft Skype). Web technologies have several promising benefits which can be of use for Peer-To-Peer systems, namely the vast dissemination of Web-enabled devices (e.g. Computer, Smartphone, SmartTV) and the ease of use with already built-in technologies. These form a great opportunity for new approaches to common multi-user use cases like *Video-On-Demand* (*VOD*), communication and general content distribution over the Web. Potential users no longer have to install a new piece of software but use the built-in HTML5-capable browser as an application platform.

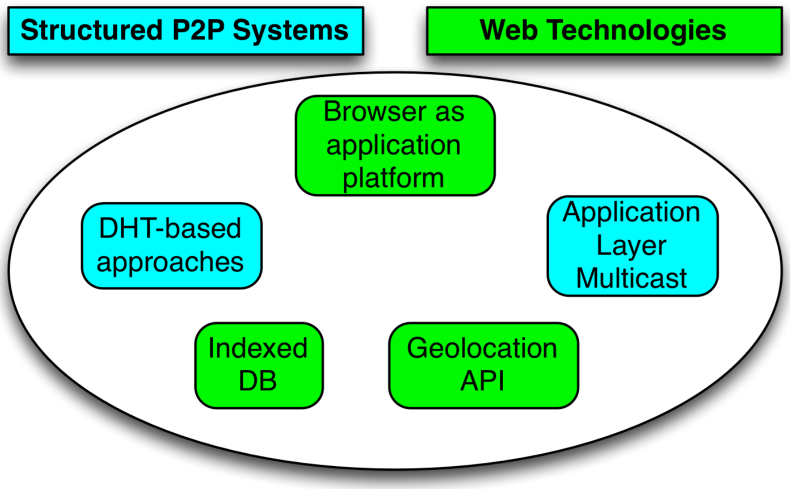

Figure 3.1: Possible overlapping issues of Peer-To-Peer and Web

As shown in figure 3.1, possible benefits of the coherence between structured Peer-To-Peer and Web can be leveraged by combining the particular subtopics discussed in chapter 2. The Geolocation API would allow Peer-To-Peer systems to use geographical information for efficient routing and IndexedDB could be used for peer-side storage and efficient searching, all running on any type of device that is capable of running a browser. On top of such a web based Structured Peer-To-Peer system, Application Layer Multicast could be used to bring multicast support to web applications.

Examples for early implementations of Peer-To-Peer with Web technologies are presented by [Meyn 2012], who focuses on browser to browser media streaming. [Davids u. a. 2011] aims at providing an API for voice communication implemented in the browser. [Nicolas u. a. 2011] focuses on reducing the gap between the legacy telco world and web services, proposing a framework that uses the new web technologies. Proprietary solutions also exist, namely the Akamai NetSession Interface[1] which uses home computers as caching peers for the Akamai CDN.

Current approaches are mostly limited to server based solutions and do not actually span an overlay network over the peers but focus on one-to-one communication. Open questions remain, including but not limited to device discovery, the use of browser capabilities like Geolocation, Indexed DB or the WebRTC Data Channel Protocol as discussed in chapter 2.

Risks

Relying on Web technologies implies dependency on the browser. In the past, browser vendors like *Google*, *Microsoft* or *Mozilla* have often been either slow in adopting new technologies or have not entirely implemented them according to the specification. An example is Microsoft's recently released proposal for an alternative to the WebRTC specification called CU-RTC-Web[2]. These alternative proposals are often a result of economic interests and therefore carry a risk when depending on a technology that later cannot be used in some of the browsers.

Another risk is in the abstraction that is introduced when programming web applications which run sandboxed in the browser. As these applications do not have access to all of the operating systems features, trustworthy third party implementations need to be available (e.g. for Cryptography). Nevertheless, well-known attacks like *Cross-Site Request Forgery* (*XSRF)* or *Cross-Site-Scripting* (*XSS*) can lever out the browser's security.

[1]http://www.akamai.com/html/solutions/client_faq.html

[2]http://blogs.msdn.com/b/interoperability/archive/2013/01/17/ms-open-tech-publishes-html5-labs-prototype-of-a-customizable-ubiquitous-real-time-communication-over-the-web-api-proposal.aspx

Bibliography

[Anttonen u. a. 2011] Anttonen, Matti ; Salminen, Arto ; Mikkonen, Tommi ; Taivalsaari, Antero: Transforming the web into a real application platform: new technologies, emerging trends and missing pieces. In: *Proceedings of the 2011 ACM Symposium on Applied Computing*. New York, NY, USA : ACM, 2011 (SAC '11), S. 800–807. – URL http://doi.acm.org/10.1145/1982185.1982357. – ISBN 978-1-4503-0113-8

[Bakht u. a. 2012] Bakht, Mehedi ; Carlson, John ; Loeb, Alexander ; Kravets, Robin: United we find: enabling mobile devices to cooperate for efficient neighbor discovery. In: *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*. New York, NY, USA : ACM, 2012 (HotMobile '12), S. 11:1–11:6. – URL http://doi.acm.org/10.1145/2162081.2162097. – ISBN 978-1-4503-1207-3

[Berners-Lee u. a. 1996] Berners-Lee, T. ; Fielding, R. ; Frystyk, H.: *Hypertext Transfer Protocol – HTTP/1.0*. RFC 1945 (Informational). Mai 1996. – URL http://www.ietf.org/rfc/rfc1945.txt

[Castro u. a. 2003] Castro, Miguel ; Druschel, Peter ; Kermarrec, Anne-Marie ; Nandi, Animesh ; Rowstron, Antony ; Singh, Atul: SplitStream: high-bandwidth multicast in cooperative environments. In: *SIGOPS Oper. Syst. Rev.* 37 (2003), Oktober, Nr. 5, S. 298–313. – URL http://doi.acm.org/10.1145/1165389.945474. – ISSN 0163-5980

[Davids u. a. 2011] Davids, Carol ; Johnston, Alan ; Singh, Kundan ; Sinnreich, Henry ; Wimmreuter, Wilhelm: SIP APIs for voice and video communications on the web. In: *Proceedings of the 5th International Conference on Principles, Systems and Applications of IP Telecommunications*. New York, NY, USA : ACM, 2011 (IPTcomm '11), S. 2:1–2:7. – URL http://doi.acm.org/10.1145/2124436.2124439. – ISBN 978-1-4503-0975-2

[Erbad u. a. 2012] Erbad, Aiman ; Hutchinson, Norman C. ; Krasic, Charles: DOHA: scalable real-time web applications through adaptive concurrent execution. In: *Proceedings of the 21st international conference on World Wide Web*. New York, NY, USA : ACM, 2012

(WWW '12), S. 161–170. – URL http://doi.acm.org/10.1145/2187836.2187859. – ISBN 978-1-4503-1229-5

[Fette und Melnikov 2011] FETTE, I. ; MELNIKOV, A.: *The WebSocket Protocol.* RFC 6455 (Proposed Standard). 2011. – URL http://www.ietf.org/rfc/rfc6455.txt

[Huang u. a. 2012] HUANG, Zhihua ; LU, Songnian ; ZHANG, Aixin ; GU, Jie: Impact of Feedback on Trust in P2P Networks. In: *Journal of Networks* 7 (2012), Nr. 8. – URL http://ojs.academypublisher.com/index.php/jnw/article/view/jnw070811821188

[Jacobs und Walsh 2004] JACOBS, Ian ; WALSH, Norman: *Architecture of the World Wide Web, Volume One.* The World Wide Web Consortium (W3C), Dezember 2004. – URL http://www.w3.org/TR/webarch/

[Lua u. a. 2005] LUA, Eng K. ; CROWCROFT, J. ; PIAS, M. ; SHARMA, R. ; LIM, S.: A survey and comparison of peer-to-peer overlay network schemes. In: *Commun. Surveys Tuts.* 7 (2005), April, Nr. 2, S. 72–93. – URL http://dx.doi.org/10.1109/COMST.2005.1610546. – ISSN 1553-877X

[Mehta 2010] MEHTA, Nikunj: Indexed Database API / W3C. URL http://www.w3.org/TR/2010/WD-IndexedDB-20100105/, Januar 2010. – W3C Working Draft

[Meij u. a. 2012] MEIJ, Edgar ; WEERKAMP, Wouter ; RIJKE, Maarten de: Adding semantics to microblog posts. In: *Proceedings of the fifth ACM international conference on Web search and data mining.* New York, NY, USA : ACM, 2012 (WSDM '12), S. 563–572. – URL http://doi.acm.org/10.1145/2124295.2124364. – ISBN 978-1-4503-0747-5

[Meyn 2012] MEYN, A. J.: *Browser to Browser Media Streaming with HTML5.* Asmussens Alle, Building 305, DK-2800 Kgs. Lyngby, Denmark, Technical University of Denmark, DTU Informatics, E-mail: reception@imm.dtu.dk, Diplomarbeit, 2012. – URL http://www.imm.dtu.dk/English.aspx. – Supervised by Associate Professor Christian Probst, probst@imm.dtu.dk, DTU Informatics

[Miranda u. a. 2012] MIRANDA, Javier ; MURILLO, Juan M. ; GUILL'EN, Joaqu'in ; CANAL, Carlos: Identifying adaptation needs to avoid the vendor lock-in effect in the deployment of cloud SBAs. In: *Proceedings of the 2nd International Workshop on Adaptive Services for the Future Internet and 6th International Workshop on Web APIs and Service Mashups.* New York, NY, USA : ACM, 2012 (WAS4FI-Mashups '12), S. 12–19. – URL http://doi.acm.org/10.1145/2377836.2377841. – ISBN 978-1-4503-1566-1

[Nicolas u. a. 2011] Nicolas, Gerard ; Sbata, Karim ; Najm, Elie: Websocket enabler: achieving IMS and web services end-to-end convergence. In: *Proceedings of the 5th International Conference on Principles, Systems and Applications of IP Telecommunications*. New York, NY, USA : ACM, 2011 (IPTcomm '11), S. 3:1–3:3. – URL http://doi.acm.org/10.1145/2124436.2124441. – ISBN 978-1-4503-0975-2

[Oram 2001] Oram, Andy (Hrsg.): *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. Sebastopol, CA, USA : O'Reilly & Associates, Inc., 2001. – ISBN 059600110X

[Popescu 2009] Popescu, Andrei: Geolocation API Specification / W3C. URL http://www.w3.org/TR/2009/WD-geolocation-API-20090707/, Juli 2009. – Last Call WD

[Raman 2009] Raman, T. V.: Toward 2W, beyond web 2.0. In: *Commun. ACM* 52 (2009), Februar, Nr. 2, S. 52–59. – URL http://doi.acm.org/10.1145/1461928.1461945. – ISSN 0001-0782

[Ratnasamy u. a. 2001] Ratnasamy, Sylvia ; Francis, Paul ; Handley, Mark ; Karp, Richard ; Shenker, Scott: A scalable content-addressable network. In: *SIGCOMM Comput. Commun. Rev.* 31 (2001), August, Nr. 4, S. 161–172. – URL http://doi.acm.org/10.1145/964723.383072. – ISSN 0146-4833

[Rowstron u. a. 2001] Rowstron, Antony ; Kermarrec, Anne-Marie ; Castro, Miguel ; Druschel, Peter: Scribe: The Design of a Large-Scale Event Notification Infrastructure. In: *Networked Group Communication* (2001), S. 30–43. – URL http://dx.doi.org/10.1007/3-540-45546-9_3

[Rowstron und Druschel 2001] Rowstron, Antony I. T. ; Druschel, Peter: Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems. In: *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg*. London, UK, UK : Springer-Verlag, 2001 (Middleware '01), S. 329–350. – URL http://dl.acm.org/citation.cfm?id=646591.697650. – ISBN 3-540-42800-3

[Singh u. a. 2012] Singh, Kapil ; Wang, Helen J. ; Moshchuk, Alexander ; Jackson, Collin ; Lee, Wenke: Practical end-to-end web content integrity. In: *Proceedings of the 21st international conference on World Wide Web*. New York, NY, USA : ACM, 2012 (WWW '12), S. 659–668. – URL http://doi.acm.org/10.1145/2187836.2187926. – ISBN 978-1-4503-1229-5

[Steinmetz und Wehrle 2005] Steinmetz, Ralf ; Wehrle, Klaus: *Peer-to-Peer Systems and Applications (Lecture Notes in Computer Science)*. Secaucus, NJ, USA : Springer-Verlag New York, Inc., 2005. – ISBN 354029192X

[Stoica u. a. 2001] Stoica, Ion ; Morris, Robert ; Karger, David ; Kaashoek, M. F. ; Balakrishnan, Hari: Chord: A scalable peer-to-peer lookup service for internet applications. In: *SIGCOMM Comput. Commun. Rev.* 31 (2001), August, Nr. 4, S. 149–160. – URL http://doi.acm.org/10.1145/964723.383071. – ISSN 0146-4833

[Traverso u. a. 2012] Traverso, Stefano ; Huguenin, K'evin ; Trestian, Ionut ; Erramilli, Vijay ; Laoutaris, Nikolaos ; Papagiannaki, Kostantina: TailGate: handling long-tail content with a little help from friends. In: *Proceedings of the 21st international conference on World Wide Web*. New York, NY, USA : ACM, 2012 (WWW '12), S. 151–160. – URL http://doi.acm.org/10.1145/2187836.2187858. – ISBN 978-1-4503-1229-5

[Vijendran und Thavamani 2012] Vijendran, Anna S. ; Thavamani, S.: Survey of caching and replica placement algorithm for content distribution in Peer to Peer overlay networks. In: *Proceedings of the Second International Conference on Computational Science, Engineering and Information Technology*. New York, NY, USA : ACM, 2012 (CCSEIT '12), S. 248–252. – URL http://doi.acm.org/10.1145/2393216.2393259. – ISBN 978-1-4503-1310-0

[Wahlisch u. a. 2009] Wahlisch, M. ; Schmidt, T.C. ; Wittenburg, G.: BIDIR-SAM: Large-scale content distribution in structured overlay networks. In: *Local Computer Networks, 2009. LCN 2009. IEEE 34th Conference on*, oct. 2009, S. 372 –375

[Wang u. a. 2012] Wang, Zhen ; Lin, Felix X. ; Zhong, Lin ; Chishtie, Mansoor: How far can client-only solutions go for mobile browser speed? In: *Proceedings of the 21st international conference on World Wide Web*. New York, NY, USA : ACM, 2012 (WWW '12), S. 31–40. – URL http://doi.acm.org/10.1145/2187836.2187842. – ISBN 978-1-4503-1229-5

[Yang und Lou 2012] Yang, Libin ; Lou, Wei: A contract-ruled economic model for QoS guarantee in mobile peer-to-peer streaming services. In: *Proceedings of the 2012 IEEE 20th International Workshop on Quality of Service*. Piscataway, NJ, USA : IEEE Press, 2012 (IWQoS '12), S. 34:1–34:4. – URL http://dl.acm.org/citation.cfm?id=2330748.2330782. – ISBN 978-1-4673-1298-1

[Zaman und Grosu 2011] Zaman, Sharrukh ; Grosu, Daniel: A Distributed Algorithm for the Replica Placement Problem. In: *IEEE Trans. Parallel Distrib. Syst.* 22 (2011), September, Nr. 9, S. 1455–1468. – URL http://dx.doi.org/10.1109/TPDS.2011.27. – ISSN 1045-9219