GKLR08] *Kalman Graffi, Aleksandra Kovacevic, Nicolas Liebau, and Ralf Steinmetz;* **From Cells to Organisms: Long-Term Guarantees on Service Provisioning in Peer-to-Peer Networks.** In: ACM: 8th ACM SIGAPP International Conference on New Technologies of Distributed Systems (NOTERE '08), June 2008.

From Cells to Organisms: Long-Term Guarantees on Service Provisioning in Peer-to-Peer Networks

Kalman Graffi, Aleksandra Kovacevic, Nicolas Liebau, and Ralf Steinmetz
Multimedia Communications Lab (KOM), Technische Universität Darmstadt
Merckstr. 25, 64283 Darmstadt, Germany
{graffi[1],sandra[1],liebau,steinmetz}@kom.tu-darmstadt.de

ABSTRACT

The peer-to-peer paradigm gained more and more impact in the last years. The reason for P2P arising now is related to the continuous development of device capabilities in the last years, like CPU power, storage space and bandwidth. However, the demand for services and resources is permanently increasing, although the peers have a variety of other resources themselves. In this paper we present the idea of a P2P system acting as a service provider using the resources of participating peers and stating guarantees on the quality of the service it provides. In order to fulfill these service level agreements, the peers confederate to a *distributed supervisor of peer resources* (DISPRO), monitoring the network, predicting trends on resource availabilities and deciding on resource allocation strategies. This paper discusses the challenges and a solution draft of the concept of DISPRO.

Categories and Subject Descriptors

D.3.3 [**Information Storage and retrieval**]: Systems and Software—*Distributed Supervision of Peer Resources*

General Terms

Design, Algorithms, Reliability

Keywords

Peer-to-Peer, System Management, Information Architecture, Distributed Service-Oriented Architecture

1. INTRODUCTION

The peer-to-peer (P2P) paradigm gained in the last years more and more impact on communication networks. Approximately 80% of todays internet traffic in the backbone networks is caused by P2P traffic [2]. The reason is that millions of peers build networks and cooperate self-organized in order to provide services and functionalities, for which previously expensive server clusters were needed. Nowadays, two main visible application areas exist for P2P networks: content distribution and communication/collaboration. While P2P file sharing systems gained a negative reputation due to copyright violations by the users, the principles of the P2P paradigm enabled very efficient solutions for global content distribution. On the other hand, Skype [30] provides a P2P based telephony system made up completely of user clients. Skype uses a server only for authentication and login purposes.

What are the lessons learned from the tremendous success of P2P based applications? Due to the technical capabilities of today's computer systems, it is possible for average consumers' computers to build mature complex systems scaling to millions of participants, and providing new functionality. Still great potentials are wasted, due to unused available resources in P2P networks. Computer systems of nowadays typically do not use their resources fully, the capacity utilization is even very low. In the following we present a selection of resources of interest:

- Online time enables the usage of the other resources.
- Computational power for programs running in a sandbox environment, e.g. for simulations/calculations.
- Encrypted storage space, e.g. for distributed data backup.
- Available bandwidth, e.g. for fast data dissemination
- Human presence, e.g. for tasks needing human interaction

This motivates the idea of a P2P system, which monitors the available resources in the P2P network and acts as a service provider for participating peers, harnessing the unused resources in the P2P system. We name this system *DIstributed Supervisor of Peer ResOurces* (DISPRO). Furthermore, DISPRO provides guarantees for the services it offers in form of service level agreements (SLA) and manages the available resources in order to ensure that these SLAs are fulfilled.

In this paper we discuss the vision and challenges of peers cooperating to create DISPRO. The paper is structured as follows. In Section 2 we describe in detail the goals of DISPRO and in Section 3 attractive application areas. In Section 4 we present the challenges to build a DISPRO and present preliminary solutions. In Section 5 we give a brief overview on related work on service guarantees and resource management in P2P networks. Finally we give a conclusion in Section 6.

[1]Authors supported by the DFG Research Group 733, "QuaP2P: Improvement of the Quality of Peer-to-Peer Systems" [6].

2. FROM CELLS TO ORGANISMS: VISION OF A DYNAMIC PEER-TO-PEER GRID

Current P2P applications are quite limited in the service they provide to individual users. In typical applications (content distribution, communication/collaboration) the P2P system enables peers to find other peers that provide either a service of interest or find the peer corresponding to a specific user. To be more general: P2P systems provide currently only information about which other peers in the network provide the specific service of interest. Any further cooperation between the peers is negotiated from single peer to single peer. Consider the case that one peer demands 1 GB storage space from the P2P network for 10 days. It is not feasible to find another peer in the network offering exactly this service, that wants to stay online for this long period of time, and additionally requests resources the first peer offers. Long-term guarantees on service provisioning cannot be made by single peers, as churn and the autonomy of peers is characteristic for P2P networks.

In our vision of a P2P Grid system the P2P network forms a virtual service provider, the DISPRO, which can be contacted by each peer in order to request any type of resource for a time-span defined by the requesting peer. To pick up the example from above: A peer requests 1 GB storage space that should be available for 10 days. Furthermore the requesting peer can describe requirements on the quality of service it requested, e.g. an availability of 99,9%. The request is expressed in form of a SLA (request-SLA).

In [10] we presented a dispatcher of single service requests considering QoS in content distribution. DISPRO considers service requests in a scale that no single peer can fulfill. The DISPRO checks whether the requested resource reservation can be granted based on resource monitoring and predictions and determines which kind of resources will be mostly needed in the future. Based on this prediction the DISPRO makes an offer consisting of two parts: a success probability and a "price" (or "reward"). The reward is a set of services provided to the DISPRO in return for the requested service. The reward is expressed in form of a SLA (reward-SLA). The requesting peer can negotiate with the DISPRO an acceptable price by modifying the requirements for the quality of service. For example the customer peer can decrease the probability for fulfilling of the request-SLA e.g. from 99.9% to 99%. The harder the requirements stated by the customer peer are, the more peers will be needed to provide the service supervised by the DISPRO. When both sides agree on a contract, the requesting peer receives the service it requested for the amount of time and with the quality of service that is fixed in the contract. For the requesting peer the inner structure of the DISPRO is irrelevant.

After having negotiated a SLA, the DISPRO picks peers from the network and commands them to provide specific service in order to meet the SLA requirements. Based on the DISPRO's knowledge on the amount of resources available on each peer, the most appropriate peers are chosen. The peer that requested the service is in return contributing the specific amount of resources determined by the reward-SLA. With this organized giving and taking of resources, both the DISPRO exists (as it is made up of peers) and the individual peers receive resources they demand but do not have themselves.

Before discussing the research challenges arising from this vision, we briefly present a selection of application areas.

3. APPLICATION AREAS

The DISPRO provides a variety of resources that can be used for several applications. In this section we present a selection of three applications from which we assume to gain considerate impact.

3.1 Distributed Computing

In 1988 Litzkow et al. introduced a system called *Condor* [22], that is designed to schedule computational jobs on remote machines. Machines that participate in the Condor cluster provide information about their status to a central cluster head, which determines which machine out of the cluster should be used when starting a new computational job or porting preliminary results from another machine.

In a P2P system where the peers cooperate and provide a DISPRO, peers could negotiate with the DISPRO SLAs guaranteeing that several computational jobs are calculated.

Typical jobs are simulation runs in research communities, rendering in computer graphics communities and analysis of various data like in Seti@Home [29] or in the GIMPS project [9] looking for Mersenne primes. Machines are manually added to the Condor cluster. Their computational power and much more the amount of available memory determine the speed of the calculations. The most important resource they have to provide is their online time, as computational jobs should not be interrupted.

Machines that suit best to these requirements are chosen by the DISPRO to provide service for the network. The processing time of the jobs can be shortened by dynamically choosing the best suited machines and migrating jobs from leaving machines. In return machines requesting computational power may provide e.g. storage space.

3.2 Backup Service

A reliable backup service is desirable in many areas. One requirement in this context is, that the backuped data is not located physically close to the original data. In a P2P system with DISPRO, peers can demand backup space with a specific guaranteed availability and provide in return storage space as well (for other peers). As replicas of the backuped data are needed and peers may leave the network or fail, the "price" for reliable backup space is high, e.g. storage space 10 times larger than the requested backup space may be demanded in the reward-SLA. Still, local storage space is very cheap and we assume that providing reliable backup space "somewhere else" is an attractive feature of a P2P system.

3.3 Globally Distributed Testbed

PlanetLab [25] is a network consisting of currently 759 nodes distributed over the world. Only 200 are online on average [31]. This network can be used to research novel network applications and protocols. PlanetLab is not large enough for the evaluation of large scale distributed applications. Although it is one of the largest testing networks in the world, it is not large enough. Furthermore experiences from PlanetLab users show that the nodes are constantly overloaded [31].

A P2P network consists typically of much more peers than PlanetLab provides. Using a DISPRO based P2P system which chooses appropriate peers to participate at network-testing applications leads to more detailed testbed results. A sandbox environment and limitation that peers can only communicate with other peers out of the sandbox could provide a minimal level of security.

4. CHALLENGES AND APPROACHES

In this section we discuss the challenges arising from the idea of a DISPRO consisting of cooperating individual peers.

4.1 The DISPRO organism: cooperative peers

Long-term SLAs can only be negotiated between single peers and a cooperative of other peers, as only a cooperative built of peers can cope with churn and failing or malicious peers. One of the most challenging issues in the design of a DISPRO is, how the peers should be organized in order to provide a higher intelligence deciding which peer should provide which kind of resource. The peers have to monitor the network, predict trends, decide on SLA negotiations and determine pricing strategies that make sure that relevant resources will be available in future as well.

We propose that peers shall organize themselves in quorums. It is important to distribute the responsibilities of the DISPRO in order to avoid arbitrariness of single peers and to solve the byzantine generals problem [19]. By the size of the quorum one can adjust up to how many malicious peers the quorum may contain before results yield by the quorum are to be assumed as faulty.

Being part of the quorum is a service every peer can provide. Making decisions and analyzing data does not require any resource in great amount, as it consists only of various short-time calculations. However, membership and reputation models for distributed systems are necessary to adopt. We assume that we have a dynamical set of peers participating in quorums. Solutions for this challenge exist, e.g. a secure quorum-based membership mechanism for P2P systems has been proposed in [12] and [21]. These quorums define the DISPRO and decide actions needed to keep the DISPRO functional and to fulfill the guarantees asserted by the DISPRO.

4.2 Resource Monitoring

Having achieved a cooperative of peers we proceed to the next step needed to create a DISPRO. The participating peers provide different kinds of resources. We assume that peers will contribute resources even without special incentives as most of the resources, like computational power, online time and bandwidth cannot be saved[1] and are wasted if not used. The amount of these resources has to be measured locally and aggregated in the network in order to provide coarse-grained information on the available resources in the network to the DISPRO. Besides the local measurements of each peer, user-given information or measurements by other peers should be used as well in order to verify the results. The information about the available resources should be acknowledged by the quorum, reliably stored in the P2P system and updated periodically.

However, resource monitoring and presentation in a network with up to millions of participants is a challenging task.

Several approaches for peer-monitoring have been proposed in literature. Most of them rely on the Key-based Routing functionality [5] which enables the key-based routing of messages in a DHT. DHTs like Chord [32], Pastry [28], CAN [27] and Kademlia [24] enable this functionality, which the information management architecture use to build a second overlay on the DHT. A classification of the information management architectures can be made according how the routing functionality is used. T-Man [16] and P2P-Diet [15] use the existing overlay topology in order to flood the peer information either purely unstructured [16] or hybrid unstructured [15] in the network. The other class of solutions build a structured over-overlay on top of the KBR-covered DHT. This over-overlay establishes a tree-based structure in which the information flows are managed. Astrolabe [34], SDIMS [36], Somo [37] and CONE [3] are to name in this category. Besides these architecture, which are loosely bound to the underlying overlay (e.g. [11]), some solutions like DASIS [1] and Willow [35] exist which are tightly bound to an overlay [1] or provide the overlay functionality themselves [35].

For DISPRO it is irrelevant which specific information management architecture is chosen. The gathered information have to be analyzed afterward. Figure 1 shows the main architecture of DISPRO for managing the required peer resources. A selection of resources of interest has been presented in Section 1. Research questions that arise are: In which detail resource information should be aggregated? Should the resource status of a peer be measured more frequently or should it be predicted? How to store the status information in a reliable distributed manner? Is a "global view" desirable or is it sufficient to restrict to local views?

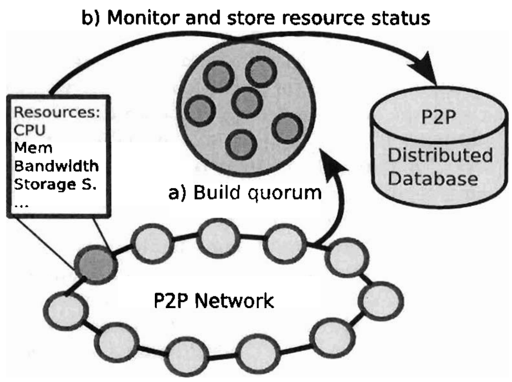

Figure 1: The main architecture of DISPRO: a cooperative of peers

4.3 Predictions on Resource Trends

For predictions we assume that the DISPRO knows which resources are available in the network and also for which amount of time. We plan the modeling of the whole resources available in the P2P network. Each local resource status can be related to a peer specific value. Based on the peer's online behavior, probabilities on the availability of its resources can be calculated for several points of time in the future. Furthermore, that average resource utilization can be predicted based on usage histories.

Calculating valid predictions for several points of time in the future is a severe challenge.

In order to evaluate the quality of DISPRO we propose to compare the amount of service provided by the DISPRO ($Serv_{DISPRO}$) with the amount of requested service ($Serv_{Req}$) or the total amount of resources ($Serv_{Tot}$) respectively. The following equation gives us a metric on the

[1] Several applications exit where people contribute without special incentives, like computational power for Condor cluster [22] or storage space for the Freenet project [7].

amount of unused resources in the system:

$$Serv_{unused} = min\{Serv_{Req}, Serv_{Tot}\} - Serv_{DISPRO}$$

One design goal of the DISPRO is to minimize $Serv_{unused}$.

4.4 SLA Negotiation

Assuming that the DISPRO knows trends on the resource availability it can negotiate SLAs with peers with more certainty. When a peer requests a specific type of resource for a given time and with a certain quality of service, the DISPRO first calculates its costs to fulfill the requested quality of service.

For example when a peer requests backup space of 1 GB for 10 days and guaranteed availability of 99,9%, the DISPRO checks which kind of peers are expected to be online in the next 10 days. Based on their average availability and the storage space they provide, the DISPRO selects e.g. 9 peers providing storage space up to 1 GB. The availability rate of the backuped data is calculated based on the availability probabilities of the chosen peers.

For modeling, let us consider 10 peers p_i, each with an expected online availability of $A_i(t)$ depending on the time and each peer providing 1 GB storage space. The available storage space has to be modeled, in order to guarantee the availability of 1 GB storage space for 10 days. In this simple model the probability $P_{av}(t)$ that at least 1 GB is available is

$$P_{av}(t) = 1 - \prod(1 - A_i(t)).$$

In order to make the model more realistic, it needs to be extended with information on the available storage space on each peer and the time needed to replicate data from one peer to another. This depends on the available bandwidth of each peer. Modeling the current resource status and the requirements to fulfill a requested service is a challenging endeavor we aim to solve in future.

However, based on the predictions made in the step before, the DISPRO can determine its costs to provide the requested service. In return the DISPRO states a reward-SLA consisting of requested resources in order to gain a benefit from the service contract. The DISPRO may have predicted that a special kind of resource will be or is already scarce in the system. Based on the resources the requesting peer has available, a peer specific reward-SLA can be offered that satisfies the needs of the system. Contract details are calculated redundantly by peers chosen by a quorum. Decisions are made by the whole quorum.

In Figure 2 we summarize the steps that are taken in the process of service registration in a P2P system with DISRPO.

4.5 SLA Enforcement

One main issue in payment based models is the enforcement of service delivery and payment. A question that arises in this context is, how to ensure that peers do not receive service without paying for it? Although payment in our DISPRO-based system consists of service that a peer has to provide, we cannot assume that peers voluntarily provide the requested service. The service each peer provides in return to service it receives from the DISPRO has to be monitored. Typically the peer provides the service fixed in the reward-SLA it negotiated with the DISPRO. In order to detect misbehavior, the reputation of peers has to be

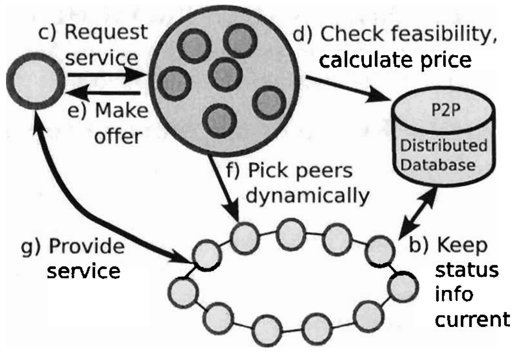

Figure 2: SLA negotiation with DISPRO: a protocol walk-through

maintained. Liebau et al. present in [20] a token-based accounting scheme for P2P networks. Based on their work, mechanisms to monitor the contribution of each peer can be applied.

Nevertheless, if a peer is unwilling to contribute, actions has to be taken. In BitTorrent [4] the tit-for-tat strategy enforces that each peer that downloads a chunk must upload as well. This defines a clear incentive to contribute. The challenging aspect why tit-for-tat cannot be used easily is that the SLA is negotiated between a single peer and the DISPRO, but service is provided from that single peer to another peer in the network. Thus no direct tit-for-tat mechanism can be used. Still, other incentives schemes exist. Appropriate incentive and reputation mechanisms [17] have to be investigated and applied.

4.6 Security

Security is one of the main challenges in the idea of a *distributed supervisor of peer resources* as it has to be applied in any monitoring, decision making, and execution process. Several questions arise: How to authenticate and identify peers? How to verify correct resource measurement? How to agree on which actions are needed to be done for the DISPRO? How to agree on their priority? Furthermore, how to enforce the payment of requested services? How should the peers' reputations be stored, maintained and determined? Due to lack of space, we do not go into detail which further challenges on security exist and which solutions are discussed in literature. In the context of functional challenges the security requirements has to be kept in mind. With these, the research questions are more challenging.

4.7 Summary

In summary, peers that acted previously as cells of single interests, now cooperate in order to create a higher intelligence, the so called DISPRO. DISPRO decides what peers have to do in order to preserve the functionality of the system and to provide best possible service quality to the peers. The DISPRO monitors the available resources in the network, predicts trends and offers services with guarantees on the quality of service using the resources of its peers. In order to fulfill these SLAs, the DISPRO calculates rewards. These rewards absorb resource leaks that are predicted by the system. In conclusion we state that the predictions on resource trends and the SLA enforcement have very promising benefits but introduce severe challenges to P2P systems.

5. RELATED WORK

The idea of management of distributed resources is well known in literature. Common approaches discussed in current literature can be classified in two categories: advancements in Grid computing or in P2P systems.

In Grid systems globally distributed service providers, like companies or institutes, participate in a virtual organization. Each participant is bound by contracts to provide a specific kind of resource, e.g. storage space. Stockinger et al. present in [13] a snapshot on the current view on Grid research. The main difference of Grid systems to our P2P-based approach lies in the uncontrollable behavior of individual peers. Peers join and leave voluntarily, they also control all of their resources and with this do not provide any guarantees.

In Grid systems resource providers are well known and can be directly or indirectly controlled. Research on Grid computing aims at the development of standards and tool kits which can be used to organize resource providers. Approaches to loose the strict requirements for the participating resource sites try to keep the Grid layer structure [33], the CORBA middleware [18] or Virtual Organizations [14] and use P2P based communication between the sites. The main challenges are to provide a virtualization of the underlying resources in order to have a commonly usable resource platform.

Several approaches for resource sharing for P2P systems are discussed in literature. Most P2P based approaches adopt the principles of publish/subscribe systems to P2P networks in order to manage available resources in the network, like in [15], [8] and [23]. These publish/subscribe systems only act as a broker between single peers.

However, the proposed solutions focus on the discovery of resources of single peers and not on the challenge, how to provide long-term guaranteed service in P2P networks in context of unreliable peers. For reliable service provisioning several peers are needed, so that failures and churn can be balanced. For this a higher abstraction level is needed, a cooperative of peers, deciding on the required actions to take and picking suitable peers. This is the focus of DISPRO in our solution.

6. CONCLUSION

This paper discusses the challenges that arise designing a P2P system that is able to provide long-term service level agreements. We introduce the idea of a *distributed supervisor of peer resources* (DISPRO) which acts as a service provider in the P2P system using the resources of the peers. The DISPRO is a cooperative of peers that monitors the resource status in the P2P network, predicts resource availability trends and enters into SLA negotiations with individual peers.

In the P2P network we consider a set of resources, like storage space, computational power or online time. Peers in the network can request specific amounts of these resources for a time span. The DISPRO negotiates with the requesting peer a guaranteed QoS based on the resource availability predictions and states a reward, which the requesting peer has to contribute in the future in return. This way it can be prevented that specific resources are missing in the future.

We discuss how the DISPRO can be designed to cope with P2P characteristic constraints, like churn or the unreliability of the nodes. These constraints define the difference to solutions discussed in literature for P2P and Grid systems.

In Grid computing, participating sites are bound by contracts to provide service. This lack of freedom leads to the case that resource availabilities are easy to predict and with this, service provision can be guaranteed. In contrast to this, we find in P2P systems voluntarily participating peers that have totally freedom on provision of their resources. With this freedom comes the lack of guarantees in P2P systems.

DISPRO describes a balanced trade-off between freedom of the peers on the one hand and guaranteed service provisioning by the system on the other hand: a system that is based on voluntary participation of the peers, but which is also able to provide guaranteed QoS. By the cooperation of individual cells an organism of higher quality is created, that provides a service each cell benefits from.

For future work, the authors aim at solving systematically the open challenges in order to create a DISPRO and with this enable P2P systems to provide long-term guarantees on service provisioning.

Acknowledgment

The authors would like to thank Konstantin Pussep, Sebastian Kaune and Osama Abboud for proof reading and giving helpful comments. Further the authors would like to thank Abdulmotaleb El Saddik for the fruitful discussions. The work in the paper was realized within the DFG Research Group 733 *QuaP2P* [6].

7. REFERENCES

[1] K. Albrecht, R. Arnold, M. Gähwiler, and R. Wattenhofer. Aggregating Information in Peer-to-Peer Systems for Improved Join and Leave. In *Proc. of IEEE P2P '04*, pages 227–234. IEEE Computer Society, 2004.

[2] N. Azzouna and F. Guillemin. Experimental analysis of the impact of peer-to-peer applications on traffic in commercial ip networks. *European transactions on Telecommunications: Special issue on P2P networking and P2P services, ETT*, 15/6:511 – 522,, 2004.

[3] R. Bhagwan, G. Varghese, and G. Voelker. CONE: Augmenting DHTs to Support Distributed Resource Discovery. Technical Report CS2003-0755, University of California, San Diego, 2003.

[4] B. Cohen. Incentives Build Robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[5] F. Dabek, B. Zhao, P. Druschel, J. Kubiatowicz, and I. Stoica. Towards a Common API for Structured Peer-to-Peer Overlays. In *Proc. of IPTPS '03*, 2003.

[6] DFG Research Group 733. QuaP2P: Improvement of the Quality of Peer-to-Peer Systems. http://www.quap2p.de.

[7] Freenet. Freenet homepage. http://freenetproject.org/cgi-bin/twiki/view/Main/WebHome, 2001,.

[8] P. Garbacki, D. H. J. Epema, and M. van Steen. An Algorithm for Broker Placement in Peer-to-Peer Overlay Networks. Technical Report PDS-2004-011, Delft University of Technology, October 2004.

[9] GIMPS project. Great Internet Mersenne Prime Search. 1996. http://www.mersenne.org.

[10] K. Graffi, S. Kaune, K. Pussep, A. Kovacevic, and R. Steinmetz. Load Balancing for Multimedia Streaming in Heterogeneous Peer-to-Peer Systems. In *ACM NOSSDAV '08: 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, May 2008.

[11] K. Graffi, A. Kovacevic, and R. Steinmetz. Towards an Information and Efficiency Management Architecture for Peer-to-Peer Systems based on Structured Overlays. Technical Report KOM-TR-2008-2, Multimedia Communications Lab KOM, Technische Universitaet Darmstadt, Merckstr. 25, 64283 Darmstadt, Germany, Mar 2008.

[12] D. Hausheer and B. Stiller. PeerMint: Decentralized and Secure Accounting for Peer-to-Peer Applications. volume 3462 of *Lecture Notes in Computer Science*, pages 40–52. Springer, 2005.

[13] Heinz Stockinger et al. Defining the Grid: A Snapshot on the Current View. under submission, 2007.

[14] W. Hong, M. Lim, E. Kim, J. Lee, and H. Park. GAIS: Grid Advanced Information Service based on P2P Mechanism. In *HPDC '04: Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC'04)*, pages 276–277. IEEE Computer Society, 2004.

[15] S. Idreos, M. Koubarakis, and C. Tryfonopoulos. P2P-DIET: An Extensible P2P Service that Unifies Ad-Hoc and Continuous Querying in Super-Peer Networks. In *Proc. of ACM SIGMOD '04*, pages 933–934. ACM Press, 2004.

[16] M. Jelasity and O. Babaoglu. T-Man: Gossip-based Overlay Topology Management. In *Proc. of ESOA'05*, 2005.

[17] A. Josang, R. Ismail, and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. *Decis. Support Syst.*, 43(2):618–644, 2007.

[18] U. G. Kang, H. J. Hwang, S. H. Lee, and C. J. Wang. SORBA: Shared Object Request Broker Architecture for Peer Computing. In *Proceedings of the Fifth IASTED International Conference Internet and Multimedia Systems and Applications (IMSA 2001)*, pages 169–174. IASTED/ACTA Press, 13 Aug. 2001.

[19] L. Lamport, R. Shostak, and M. Pease. The Byzantine Generals Problem. *ACM Trans. Prog. Lang. Syst.*, 4(3):382–401, 1982.

[20] N. Liebau, V. Darlagiannis, A. Mauthe, and R. Steinmetz. A token-based accounting scheme for p2p-systems. Technical Report TR-2004-05, Technische Universität Darmstadt, Jan 2004.

[21] N. Liebau, O. Heckmann, A. Kovacevic, O. Tafreschi, M. Fidler, A. Mauthe, and R. Steinmetz. A Secure Quorum Based Membership Mechanism for P2P Systems. In *Proceedings of 12th Americas Conference on Information Systems*, Aug 2006.

[22] M. Litzkow, M. Livny, and M. Mutka. Condor - a hunter of idle workstations. In *Proceedings of the 8th International Conference of Distributed Computing Systems*, June 1988.

[23] Matthieu Weber. Broker-Based Architecture for Advertisement Propagation in P2P Networks. In *Proceedings of the 9-th International Conference on Theory and Tools for Transmitting, Receiving and Processing of Information*, 2003.

[24] P. Maymounkov and D. Mazieres. Kademlia: A Peer-to-Peer Information System Based on the XOR Metric. In *IPTPS*, 2002.

[25] PlanetLab. An Open Platform for Developing, Deploying, and Accessing Planetary-Scale Services. www.planetlab.com, 2006.

[26] K. Pussep, P. Knezevic, N. Liebau, and R. Steinmetz. Improving Xpath Query Execution in P2P XML Storage by Using a Decentralized Index. In *VLDB Int. Worksh. on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P '07)*. Springer, Sep 2007.

[27] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A Scalable Content-Addressable Network. In *SIGCOMM '01: Proc. of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, volume 31, pages 161–172. ACM Press, October 2001.

[28] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems. In *Proc. of IFIP/ACM Middleware '01*, pages 329–350. Springer-Verlag, 2001.

[29] SETI. Seti@home: Search for extraterrestrial intelligence at home. http://setiathome.ssl.berkeley.edu/.

[30] Skype. Peer-to-peer Internet Telephony. http://www.skype.com/skype_p2pexplained.html, http://www.skype.com, 2004.

[31] N. Spring, L. L. Peterson, A. C. Bavier, and V. Pai. Using PlanetLab for network research: myths, realities, and best practices. *Operating Systems Review*, 40(1):17–24, 2006.

[32] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Trans. Netw.*, 11(1):17–32, 2003.

[33] P. Uppuluri, N. Jabisetti, U. Joshi, and Y. Lee. P2P Grid: Service Oriented Framework for Distributed Resource Management. In *SCC '05: Proc. of the 2005 IEEE International Conference on Services Computing*, pages 347–350, USA, 2005. IEEE Computer Society.

[34] R. van Renesse, K. P. Birman, and W. Vogels. Astrolabe: A Robust and Scalable Technology for Distributed System Monitoring, Management, and Data Mining. *ACM Trans. Comput. Syst*, 21(2):164–206, 2003.

[35] R. van Renesse and A. Bozdog. Willow: DHT, Aggregation, and Publish/Subscribe in one Protocol. In *Proc. of IPTPS '04*, pages 173–183. Springer, 2004.

[36] P. Yalagandula and M. Dahlin. Research Challenges for a Scalable Distributed Information Management System. Technical Report CS-TR-04-48, The University of Texas at Austin, Department of Computer Sciences, 2004.

[37] Z. Zhang, S. Shi, and J. Zhu. SOMO: Self-Organized Metadata Overlay for Resource Management in P2P DHT. In *Proc. of IPTPS '03*, volume 2735, pages 170–182. Springer, 2003.