The state of peer-to-peer network simulators

Article (Accepted Version)

Basu, Anirban, Fleming, Simon, Stanier, James, Naicken, Stephen, Wakeman, Ian and Gurbani, Vijay (2013) The state of peer-to-peer network simulators. ACM Computing Surveys, 45 (4). 46:1-46:25. ISSN 0360-0300

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/id/eprint/42240/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

The state of peer-to-peer network simulators

ANIRBAN BASU, Tokai University, Japan and the University of Sussex, UK

SIMON FLEMING, University of Sussex, UK

JAMES STANIER, University of Sussex, UK

STEPHEN NAICKEN, University of Sussex, UK

IAN WAKEMAN, University of Sussex, UK

VIJAY K. GURBANI, Bell Laboratories, Alcatel-Lucent and Illinois Institute of Technology, USA

Networking research often relies on simulation in order to test and evaluate new ideas. An important requirement of this process is that results must be reproducible so that other researchers can replicate, validate and extend existing work. We look at the landscape of simulators for research in peer-to-peer (P2P) networks by conducting a survey of a combined total of over 280 papers from before and after 2007 (the year of the last survey in this area), and comment on the large quantity of research using bespoke, closed-source simulators. We propose a set of criteria that P2P simulators should meet, and poll the P2P research community for their agreement. We aim to drive the community towards performing their experiments on simulators that allow for others to validate their results.

Categories and Subject Descriptors: I.6.7 [**Simulation and Modelling**]: Simulation Support Systems

General Terms: Design, Experimentation, Measurement, Performance, Standardisation

Additional Key Words and Phrases: Peer-to-Peer, P2P, Simulator Evaluation, Simulator Usage

ACM Reference Format:

Basu, A., Fleming, S., Stanier, J., Naicken, S., Wakeman, I., and Gurbani, V. K. 2012. The state of peer-to-peer network simulators. *ACM Comput. Surv.* V, N, Article A (January YYYY), 26 pages.

DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

Peer-to-peer (P2P) is a distributed computer architecture which facilitates the direct exchange of information and services between individual nodes (called *peers*) rather than relying on a centralised server. In P2P networks, users (i.e. human beings) participate only through their representative agents, i.e. software agents running on networked devices.

P2P forms the basis of many distributed computer systems, permitting each peer node to act as both a client and a server, consuming services from other available peers, whilst providing its own service to the rest of the network. Peers within a P2P

Authors' addresses:

A. BASU. (abasu@cs.dm.u-tokai.ac.jp, a.basu@sussex.ac.uk) Rm. 1603, Tokai University, 2-3-23 Takanawa, Minato-ku, Tokyo 108-8619, Japan

S. FLEMING, J. STANIER, S. NAICKEN and I. WAKEMAN. (s.fleming@sussex.ac.uk, j.stanier@sussex.ac.uk, stephenn.ianw@sussex.ac.uk) School of Informatics, University of Sussex, Chichester Building, Falmer, Brighton BN1 9QJ, UK

V. K. GURBANI (vkg@bell-labs.com) Bell Laboratories, Alcatel-Lucent, 1960 Lucent Lane, Rm. 9C-533 Naperville, IL 60566, USA

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0360-0300/YYYY/01-ARTA \$15.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

network engage in direct exchanges with their known neighbours, in order to submit requests and serve responses.

The definition of what specifically constitutes a P2P system is broad. For example, in theory, a P2P system is thought of as having no centralised authority, when in reality many existing P2P applications rely on one. For example, early versions of the BitTorrent¹ protocol required a “tracker” in order to rendezvous the peers and perform membership management of the swarm (a *swarm* is a collection of peers that are interested in distributing the same content). We find that the following definition by Risson and Moors [2006] is well suited to classifying P2P systems:

Peer-to-peer systems are distributed systems consisting of interconnected nodes able to self-organise into network topologies with the purpose of sharing resources such as content, CPU cycles, storage and bandwidth, capable of adapting to failures and accommodating transient populations of nodes while maintaining acceptable connectivity and performance, without requiring the intermediation or support of a global centralised server or authority.

P2P offers many advantages. These include scalability, high resource availability, no need for a centralised authority (eliminating a single point of failure) and robustness. However, the consequence of using a P2P architecture is that the quality and usefulness of the services on offer rely entirely on the participating members of the group.

The power of P2P is apparent when considering Metcalfe’s Law [Hendler and Golbeck 2008], which states that the *value* of a network is proportional to the square of the number of connected users. The number of possible connections within a P2P network can be exponential in relation to the number of network nodes, n . All nodes can potentially connect to all other nodes, giving a theoretical maximum number of connections of $n(n - 1)/2$; the same number as in a fully connected mesh network. Peers may end up communicating directly, however multi-hop routes should exist between any pair of nodes connected to a P2P network through emergent pathways via other peers.

With the enormity of the modern Internet and the efforts to expand it further with the introduction of the Internet Protocol Version 6 (IPv6) [Deering and Hinden 1998], the number of devices that can be explicitly addressed is always increasing. Therefore, P2P will continue to offer the potential for highly scalable and efficient networking technologies in the future on many more devices than just traditional computers, switches and routers.

Researchers and developers continue to develop new technology and protocols for networks. In order to propose, design, implement and evaluate new P2P technologies, there is a key problem: how can experiments be performed and monitored in a “real” environment? Due to the complexity and scale of real P2P networks, a popular option is to use computer simulations. A P2P network can be simulated on one or more computers, allowing for the behaviour of peers to be monitored precisely, giving the possibility for iterative design, and perhaps most importantly, detailed evaluation. However, as no single simulation platform or set of standards has been agreed upon, many different simulators have been created and used for these purposes in the literature.

In this survey we are interested in looking at the landscape of current P2P simulators. We present the following:

- We give an overview of P2P networking, highlighting key features of such systems. We also look at some of the most well-known P2P applications that are being used currently, and describe the benefits and drawbacks of P2P simulation.

¹<http://www.bittorrent.com>

- We present our evaluation criteria for categorising and evaluating the features of P2P simulators.
 - We poll the P2P networking community for their agreement on our proposed criteria, and discuss our findings.
 - We perform a survey of over 100 post-2007 papers in order to discover which P2P simulators are being used in the recent literature, and combine these findings with 180 papers from our previous pre-2007 survey.
 - Using our findings, we motivate the need for unity and repeatability in the P2P simulation community. We propose that a set of standards should be adhered to which could be based upon our evaluation criteria.
- We are also socialising our findings with standards communities, in particular the Internet Engineering Task Force (IETF) and the Internet Research Task Force (IRTF), for a broader dissemination of this information in the form of IETF Request for Comments (RFC) standard documents [Gurbani et al. 2011].

1.1. P2P networking

In this section we will define our terminology, identifying the key features of P2P systems and the terms with which we shall work.

1.1.1. Neighbours and connectivity. Each peer typically knows about a small number of other peers in the network. A peer's neighbours are the possible routes for forwarding messages, and they also act as gateways to receive responses from the network. The neighbours of a peer are directly reachable. Routes beyond these neighbours provide connectivity to the entire P2P network as a whole. Peers need to keep track of their neighbours to ensure that they still remain connected to the P2P network. Without neighbours, no routes to or from any single peer exist. As such, peers may need to drop existing neighbours, or request new neighbours, throughout the duration of their P2P session to ensure that they remain connected to the network.

1.1.2. Network Churn. Peers arrive and leave the network at various rates. Current peers are the active population within the network. The action of peers arriving and leaving is called network churn. Churn is induced by “willing” actions e.g. node departures and mobility, and also “unwilling” actions e.g. inadvertent failure of a node. Peers will arrive at some point in time and begin to participate within a P2P network, remain for some period of time (called the session duration) while possibly making requests and serving other users, and then eventually leave. They may return at some point in the future. High levels of churn indicate many peers coming and going frequently, while low levels exhibit longer peer session durations with less frequent arrivals and departures.

By nature, different P2P applications will have different levels of churn. “Seeders”, or peers that have the entire content may stay connected to the swarm for a longer period of time, while “leechers”, or peers that have portions of the content may connect with the swarm until the content is downloaded and then leave the swarm. For a swarm to continue serving content, the leechers may have to act altruistically by becoming seeders after acquiring the complete content.

The quality of service in P2P networks may vary under different levels of churn. Therefore, protocol designers need to take into account the anticipated levels of churn to provide an appropriate service for users. For example, in a file sharing service with very high churn, replication of critical files across peers would give a greater chance of them being available.

1.1.3. Bootstrapping. Peers need a means of gaining access, and becoming a participant of, a given P2P network. This process is known as bootstrapping. Often, peers will

bootstrap using some kind of centralised resource, commonly called a bootstrapping node or rendezvous host, which gives the unconnected peer an entrance point into the network in the form of a set of active network member addresses, or a list of other centralised repositories to query. Once these addresses have been obtained, the peer can start communicating with the network.

1.1.4. Routing Tables. Peers are typically aware of some or all of the other members of the network. To keep track of these nodes and possible routes, a local collection of peer addresses is managed via a routing table. A routing table may indicate a neighbour address directly, or alternatively via a next hop in a multi-hop route, and also associates a cost with each potential path.

1.1.5. Bandwidth and Resources. Bandwidth is a measure of the capacity of a peer has for uploading and downloading data and is often asymmetrical. The amount of bandwidth that each peer has can vary greatly depending on the type of device. A peer could be a dedicated server with a fast broadband connection, or conversely it could be a mobile phone with only a GPRS connection. Resources such as memory, CPU speed and number of cores, number of sockets or connection blocks that can be maintained are also important contributors. Peers hosted on a well-provisioned host will have more resources at their disposal than peers running on resource constrained devices such as personal digital assistants (PDA) or mobile phones.

1.1.6. Network Topologies. The connections between nodes in a network form a graph $G = (V, E)$, such that G consists of the set of vertices V (peers) and the collection of edges in E (links to neighbours). The topology of a P2P network is the structure created by the active connections from peers to other peers. The topology makes use of an underlying physical network topology consisting of the networking devices and physical communication channels and links.

1.2. Example P2P applications and algorithms

There are a number of P2P systems and algorithms that have become popular, with many receiving mainstream attention for both good and bad reasons. Academic systems tend to focus on routing algorithms, while real world systems facilitate a specific need such as file sharing.

1.2.1. P2P Algorithms. There are two types of P2P networks: structured and unstructured. The primary aim of these networks is to host resources at one or more peers in a network (a storage function) and to allow other peers to find these resources (a routing function) in a distributed space. Generally speaking, the resources can be viewed as “values” and the search characteristics used to locate the resources can be viewed as “keys”. Looked in this light, the function of a P2P network can be specified as a pair of methods: *put(key, value)* and *value = get(key)*, i.e., store a value with an associated key (or set of keys) in the network and retrieve the value given a key at some later time. The value stored in the network represents the resource of interest. The resource could be an image, a media file (movies, songs), books, or the latest Linux kernel. P2P networks typically do not care what is stored in the network as long as there is some search criteria by which the peers in the network can find the resource.

The difference between structured and unstructured P2P networks is how they approach the task of finding resources. This is where most of the algorithmic research in P2P networks has been focused on [Risson and Moors 2006]. Unstructured P2P networks did not impose a strict formation on the peers themselves; a joining peer could simply connect to any existing peer to become part of the P2P network. In structured P2P networks (or overlays), the joining peer is admitted based on an identifier that places the joining peer in relation to a set of existing peers. Thus structured

P2P overlays are characterized as forming a distributed hash table (DHT). DHTs are distributed data structures that map a key to a value; each peer is responsible for a certain portion of the namespace where the keys are hashed. The result of this deterministic hashing is that there exists a strict upper-bound on the search time for a resource. Comparatively, unstructured P2P networks used blind flooding and random walks to search for the resource. Because the flooding was unbounded, routing costs accrued rapidly. The flooding was costly when the resource was not present in the P2P overlay, but determining the non-existence of the resource expended considerable computational and network complexity. To remedy this, some structure in the form of super-peers or peer clustering was developed with a centralized index to rapidly find a resource.

Gnutella [Klingberg and Manfredi 2002] was an early protocol that used flooding while Napster evolved to support decentralized content and a centralized index. Early structured algorithms included Pastry [Rowstron and Druschel 2001], Tapestry [Zhao et al. 2001], Chord [Stoica et al. 2001], Content Addressable Network (CAN) [Ratnasamy et al. 2001] and Kademlia [Maymounkov and Mazières 2002]. While the topological arrangement of these overlays varied — Pastry and Tapestry could be visualized as spanning trees, Chord arranged its peers in a ring, CAN created zones in a torus and assigned peers to these zones and Kademlia effectively treats nodes as leaves in a binary tree — their biggest benefit was that resources could be located using a fixed number of hops. Furthermore, the state information maintained by peers in order to traverse the network was minimal as well. Each peer maintained a small fixed-length routing table where the number of entries in the routing table was in logarithmic proportion to the number of peers in the P2P network.

1.2.2. P2P Applications. Some of the most well-known P2P systems are given below.

Distributed event management. Astrolabe [Van Renesse et al. 2003] is a DNS-like distributed information management system that monitors the dynamically changing state of a collection of distributed resources, reporting summaries of this information to its users. The updates disseminate on the order of seconds or tens of seconds. Astrolabe achieves robustness through the use of a P2P protocol to communicate between the distributed hosts.

P2P Multimedia. With the nexus of communication moving steadily towards the Internet, decentralized, and sometimes free, voice and video communications have proliferated. In many cases, these systems are composed of P2P networks. Skype² is the most visible of such P2P communication networks. RELOAD [Jennings et al. 2012] is another P2P protocol for multimedia communications based on the Session Initiation Protocol (SIP [Rosenberg et al. 2002]). More recently P2P networks have been used to deliver streaming content, both static recorded content as well as live content. Early versions of Joost³ (2007-2008) used P2P techniques to deliver static content. PPS.tv (PPStream)⁴ is P2P video streaming service popular in China. QQlive⁵ is another Chinese video P2P streaming technology.

Distributed data stores. The database functions related to large scalable systems such as Facebook and Twitter place more emphasis on search performance, real-time nature and eventual consistency when compared to enterprise back office functions that use traditional databases where the ACID (atomicity, consistency, isolation and durability) properties dominate. Dynamo [DeCandia et al. 2007] and Cas-

²URL: <http://www.skype.com>

³URL: <http://www.joost.com>

⁴URL: <http://www.pps.tv>

⁵URL: <http://live.qq.com>

sandra [Lakshman and Malik 2010] are the prime examples of such distributed storage systems. Such systems manage very large amounts of loosely structured data spread out across commodity server farms while providing robust service with no single point of failure. P2P techniques such as consistent hashing [Karger et al. 1997] is used to partition data across a cluster and replication is used to achieve high availability. The Tahoe Least-Authority Filesystem (Tahoe-LAFS [Wilcox-O’Hearn and Warner 2008]) is an open source, decentralized, fault-tolerant and secure P2P distributed data store.

Anonymity through P2P. While we make no judgment on the use of anonymity to access illegal content, we do note that the act of being anonymous has high value in societies. Anonymous sources have a rich tradition in journalism while anonymous dissidents can hide from repercussions while providing insight to a society. There are a number of P2P systems built for anonymity. They typically hide the identity of the participants by special routing techniques. Freenet [Clarke et al. 2000] is a censorship resistant distributed file system for anonymous publishing. I2P⁶ [Zantout and Haraty 2011] is a decentralized overlay network for strong anonymity and end-to-end encryption.

1.3. Advantages and Disadvantages

P2P networks allow for a distributed system to exist which benefits from the combined resources of the networked group of clients. Connecting together a collection of nodes/peers provides robust services that operate without significant set-up costs to one particular organisation or individual. As each node in the network operates as both a client and a server it promotes incentives to participate rather than spectate within the system.

Unlike centralised systems, as the number of P2P clients grow the total resources available to the peer-to-peer network (such as CPU, memory, files and storage capabilities) available also increase. The perceived growth of total resources along with the growth of nodes echoes the ideas of the grid and cloud computing communities. This flexible growth of resources has often been called “elasticity”.

P2P does not suffer from a single point of failure as seen in centralised services, as the core functionalities are distributed across the entire network. There can be points of failure, such as BitTorrent network trackers going down, but there is no single failure that disrupts the entire network. There is no centralised authority on P2P networks which gives greater freedom of operation for the network participants, but as ever, such freedom comes at a cost: without administration, malicious code and files may exploit vulnerabilities of nodes in the network.

Since there is no way to authenticate the data being transferred around P2P networks, illegal material may be shared. Also, a resource may not actually be what it claims to be: for example, a fake file server sharing viruses instead of the advertised files.

1.4. P2P Simulation

When developing new networking technology, such as protocols or routing algorithms, researchers and developers need to evaluate the technology with regards to several properties such as its quality, overhead and robustness. However, real P2P networks may consist of thousands of unique entities and experience unpredictable levels of churn. This makes evaluation of new systems with real users and devices a complicated task. How can this new technology be deployed cheaply for testing, and on a large enough scale to show useful results? What if something harmful were to happen

⁶URL: <http://www.i2p2.de>

to one of these devices? Generally, the expense of real user testing is too high to be considered. Furthermore, the difficulty of attracting real users to what is an engineering improvement without any visible user benefits makes *P2P simulation* an attractive option.

A P2P protocol could be implemented on top of a real world test bed environment such as the global research network PlanetLab [Chun et al. 2003]. Such a test bed consists of over 1000 computers situated at participating sites across the globe⁷. While such a setup provides valuable real-world data in terms of network latency, the disadvantage is that a researcher has to depend on the machines on which an experiment was conducted being constantly available and appropriately configured on PlanetLab in order to reproduce results. This itself is difficult task since the machines are owned by different administrative domains, each with its own support staff and hardware/software upgrade schedule. Therefore, tracking issues and problems across such a large number of real computers is a challenging and daunting task.

Simulation is therefore a valuable tool to develop and validate technology before deployment and exposure to the vagaries of the real world. At their essence, P2P simulators allow for the representation of nodes (which simulate peers) with incoming and outgoing message queues. Over time steps these message queues are processed, passing messages between peer outgoing and incoming queues, simulating the communication in a P2P network. All nodes are aware of their neighbours, and therefore send and receive messages to their neighbours via their message queues.

General purpose network simulators could and have been used to evaluate P2P systems, even though they are generally tailored to the emulation of low level networking mechanics. The Network Simulator (ns2 and its successor ns3) [McCanne and Floyd. 1999; Henderson 2008], for example, provides a vast range of networking devices and channels to allow the accurate and realistic representation of the underlying physical network.

If P2P simulation is able to accurately simulate real P2P networks, it can give the ability to design, implement and evaluate new technologies in a controlled environment, while allowing for the collection of statistics with greater ease, compared to deployment with real users. The importance of P2P simulation is apparent in the literature, with a majority of new P2P developments being validated by implementation in a simulator.

There exist many different P2P simulation environments that all aim to simulate a P2P network by providing an abstracted networking model on which users can build their specific simulation models and behaviours. A timeline of major simulator releases is shown in Figure 1. We will look at these later in the survey. However, as shown previously by Naicken et al. [2007], many P2P developments are evaluated by the use of bespoke P2P simulator technologies, where the simulator is written by the authors of the paper. This decision is interesting considering that many open source P2P simulators already exist. Three reasons are identified for this: the steep learning curve of P2P simulators; the excess simulator functionality that is not required for many experiments, which may hinder the implementation of experiments and the analysis of results; and where required functionality is absent, modification of the simulator code may prove to be complex and time consuming.

However, this raises a concern. Since a range of open source P2P simulators exist, and additionally, since many bespoke simulators are used, verifying and repeating results in the literature is difficult. While we are positive that no researchers would wish to present false or misleading results, the community would benefit from being able to repeat, verify, and build upon new findings [Zhang et al. 2009]. It is perhaps tempt-

⁷<http://www.planet-lab.org/>

Fig. 1: A timeline of network simulator releases.

ing to conjecture whether simulator certification should be a requirement for publication. Certainly, one could run a simulator in parallel with deployment and observe the behaviour for small- and medium-scale experiments. Simulation results that match empirical results will provide higher confidence in larger-scale simulations. Unfortunately, as attractive as certification sounds, there are drawbacks. The biggest question would be who will be willing to act as an impartial certifying authority? Certifying a simulator's results will require detailed knowledge of the P2P protocol being simulated and the expected behaviour of the P2P system. Furthermore, simulations are typically run in order to confirm the behaviour of a P2P system if certain changes to the protocol were suggested. Ergo, there will exist a chicken-and-egg problem where certifying the results of a simulation will require the empirical system to support the change that is being simulated. In the end, perhaps the best that can be done is to ensure that the simulation results are repeatable such that future generation of researchers can rapidly repeat past experiments to have a greater confidence in the simulation that will, in turn, allow them to continue the research by building on the previous work.

1.4.1. Alternatives to simulation. Even though the main focus of this paper is P2P simulators, it is worth noting that there are alternative approaches, such as deployment and development frameworks, as well as networking libraries that allow running P2P protocols both on simulation and on a real network. These include PlanetLab [Chun et al. 2003], SPLAY [Leonini et al. 2009], NEST [Dupuy et al. 1990] and NEKO [Urban et al. 2001]. The system proposed in [Wang et al. 2008] allows for trace replay, which can replay a deployed application from trace gathered from simulations. The WIDS toolkit [Lin et al. 2005] achieves higher scalability at the cost of accuracy, by relaxing the synchronization model between nodes.

1.5. Impact of previous survey of P2P simulators

In our previous survey of P2P simulators [Naicken et al. 2007], conducted in 2007, we postulated that the significant use of bespoke simulators in academic research was motivated by poor or missing functionality in existing P2P simulations. It was hoped that the findings of the survey would lead to closer collaboration between the designers of P2P simulators and academics conducting P2P research. A survey of papers that cite this work shows that the opposite has occurred and that there has been further fragmentation in the domain of P2P simulators.

Much of the research that cited the 2007 paper did so in order to justify the use of a P2P simulator, be it custom or otherwise, however there were also contributions to address some of the issues that we raised. Efforts were concentrated on both the development of new simulators and the addressing of issues, such as the transition of simulation code to deployment and APIs.

Rather than collaborate to work on a single simulators, a number of general purpose simulators were created, such as DEUS [Amoretti et al. 2009], OverSim [Baumgart et al. 2007], PeerfactSim.KOM [Stingl et al. 2011]. The ProtoPeer framework allows for the transition of simulator code to the network, an issue we previously highlighted, while the authors of the P2PAM framework [Agosti et al. 2008] adopt an approach that we expected to be more common, which is the modification of existing simulators to improve functionality.

Juxtaposed to the development of general purpose simulators are application-specific frameworks that have been developed post-2007, in particular for BitTorrent [Barcellos et al. 2008] [De Vogelee et al. 2008]. This shows that there may be two contradictory schools of thought, either P2P simulators can address different use cases or that they can not. If the latter is true it would go some way to explaining the popularity of custom simulators.

The main concern of previous and the current work is the repeatability of experiments and mechanisms to facilitate their reproducibility. A common API for simulators to allow for the repeatability of experiments across simulators was proposed by Gross et al. [Gross et al. 2010].

In hindsight, it was somewhat naive of us to assume that the P2P academic community would pool their efforts into the design and use of a single P2P simulator, as we did not work alongside a body that would undertake the responsibility for leading such efforts. To ensure that efforts continue beyond the publication of this survey, we have worked closely alongside the IRTF P2P group. This will ensure that the research community will be able to co-ordinate a collaborative effort to ensure the principal goal, which is the repeatability of academic P2P experiments, is met.

2. EVALUATION CRITERIA

The abundance of P2P algorithms and related applications requires a wide range of features from a simulator in order to perform realistic simulations. The purpose of a peer-to-peer simulator is to provide the building blocks for P2P researchers to implement and test new protocols. In order to evaluate the existing P2P simulators, we construct a set of criteria which we would expect to be met by all simulators in order for them to be accepted and used widely. Further to that, we also present a set of proposals, most of which we envision should be considered when designing and developing new P2P simulation tools. The criteria and the set of proposals are developed and extended from the evaluation criteria presented in our previous works [Naicken et al. 2007; Naicken et al. 2006b; Naicken et al. 2006a] as well as opinions gathered by conducting user surveys.

2.1. Evaluation criteria

Simulation architecture and features. (1) There are two main classes of simulators: query-cycle and discrete event. A query-cycle simulation loops through each node in the overlay network, and carries out queries for networks as and when required. In contrast, a discrete event simulator maintains a thread scheduler which synchronises a queue of messages to be transferred between the simulated nodes of the overlay network. A P2P simulator should be either discrete event or query-cycle, or enable both modes of operation. The classical trade-off in simulation involves sacrificing the level of detail of a simulation in order to gain speed and scalability (in terms of number of hosts being simulated). Session-level simulators also exist, however they do not simulate packet-level events, but rather consider a flow, or a session as the important attribute of a simulation. This allows for simulation of larger networks, albeit at a lower level of detail. (2) The simulator should enable specification of event input files, either as a complete list of pre-determined events, or events that are dynamically generated may generate other causal events; (3) random number generation seeding should allow for repeatability in the simulator; (4) it should support modelling of realistic transport protocols, e.g. message queueing; (5) the simulator should not present any bootstrapping problems⁸ to simulation experiments; (6) it should be written in a well-known programming language.

Underlying network simulation. While P2P networks are overlay networks, many a time P2P researchers are interested in testing their protocols with different implementations of the underlying network, for example: (1) topology generation, e.g. random, circular; (2) network latency and bandwidth including any variation between any pair of nodes and congestion between nodes; (3) different types of access links, e.g. Ethernet, WiFi, 3G, LTE; (4) the ability to model a nodes computational capabilities in terms of processing power and storage specified in compute units, as is done in Amazon EC2⁹.

Behaviour. P2P overlay networks often need to model the behaviour of individual nodes in terms of features, for example: (1) misbehaviour (intentional or unintentional); (2) message loss; (3) mobility models, such as the mobile waypoint model; (4) churn with parametrised distribution models; and (5) node failure.

Statistics. The ability to generate statistical data about a simulator run is an essential feature for a P2P simulator. Certain properties of such statistics generation are to be expected, for example: (1) support for output formats, e.g. SQL, XML, text; (2) mechanisms for logging experiments; (3) network graph properties, e.g. HITTS, PageRank, power law, centrality, betweenness; and (4) support for GraphML (a versatile graph file format) or similar.

Usability. One of the most important criterion that P2P simulators should meet is usability. Despite having a host of features, many simulators do not get used simply because of their learning curves or unavailability of community support. The simulators should: (1) adhere to standard programming interfaces, such as the Common API [Dabek et al. 2003] when simulating DHTs; (2) provide clear documentation; (3) have user community support through forums, mailing list, amongst others; (4) have a reference implementation of well known P2P protocols, e.g. Chord, Kad network (which is a Kademlia-based DHT and implemented by eMule [Kulbak and Bickson 2005], among other applications), CAN, Tapestry, and so on; and (5) have a GUI for visualisation of various aspects of the simulation.

⁸To simulate a dynamic network requires the ability to acquire bootstrap nodes, such as uniformly at random (random graph), from a distinct set of nodes (super nodes) or possibly a single node (central control). Without such facilities, these will need to be manually designed and implemented.

⁹<http://aws.amazon.com/ec2/>

2.2. Recommendations

We make the following generic recommendations in view of improving existing and new simulation platforms.

Efficiency and scalability. In order to evaluate how P2P algorithms will perform in large scale networks, it is essential to be able to simulate large networks. Often, this is restricted by simulator bottlenecks arising from the computational power of single machines. Ideally, simulators should make use of multi-core, multi-processor machines through the use of multi-threading. To grow out of the boundary of a group of CPU cores connected through a high-speed memory bus, further concurrency can be achieved if simulators scale to distributed setups, such as a cluster. Network latencies affecting delays in such distributed simulation will not affect discrete event simulations because of the discrete time scale. However, on a networked setup, it is also harder to maintain a unified clock. We envision that distributed P2P simulation is a research area in itself with a number of unsolved challenges. For example, there is research on concurrency using distributed Java Virtual Machines [Zhu et al. 2002] that spawn over multiple nodes in a network allowing Java programs to see a globally addressable memory space and computing cores larger than one node. On the other hand, [Dean and Ghemawat 2008] proposed the generalised MapReduce concurrent computational paradigm, which we envision can also be used to run discrete event simulations now that robust MapReduce libraries such as Apache Hadoop are freely available. In addition, cloud computing also provides substantial non-volatile storage space and reasonable computational power, which may be useful in some simulation scenarios. Massive parallelism can also be obtained by delegating parallelisable tasks to a graphics processing unit (GPU). For example, the nVidia Tesla “Fermi” series “desktop supercomputer”¹⁰ platform allows virtual addressable memory space spreading across more than one GPU chipset in a distributed fashion.

Inter-operability with an emulator. While complicated scenarios can be simulated, many researchers also opt for emulation on testbeds such as PlanetLab and EmuLab. Emulators allow capturing unforeseen network behaviour that cannot always be modelled in a simulator. The ability to port simulation scenario code from a simulator to an emulator will help researchers obtain faster results from their evaluation prototypes. New simulators should consider inter-operability of simulation code with emulators, such that researchers can deploy emulation scenarios from their existing simulation scenarios with minimal or no additional coding.

Realistic simulation. One of the fundamental objectives of simulation is to model realistic scenarios as closely as possible. This means that simulators should consider realistic network conditions. They should also simulate different capabilities of network nodes in terms of their computational power, device types, and connectivity.

Flexibility. One aspect of flexibility in a simulator platform is node (or peer) scalability as discussed above. Beyond scalability, a simulator framework should also exhibit flexibility in allowing transient state of a simulation to be saved and subsequently resurrected so that long running simulations can be stopped and restarted at a later point in time, or stopped and migrated to a completely different host. A simulation framework should also allow trace-driven simulation for easy replay of events captured in a log file. This allows real-world dynamics to be replayed in the simulator.

¹⁰http://www.nvidia.com/object/tesla_computing_solutions.html

3. PAPER REVIEW

Continuing in the same vein as our previous survey [Naicken et al. 2007] (Figure 2), we aimed to evaluate 200 P2P-related published papers. We took the first 100 search results from Google Scholar¹¹ for the search terms “peer-to-peer” and “P2P”, sorted by number of citations, in order to locate research using P2P simulation. Our previous paper used the same methodology to survey papers at the time of writing (2007) and before. To continue from that work, and to provide a snapshot of the current state of P2P simulators, we only considered papers that were published since 2007. From the search results, we only considered papers we were able to access in an electronic format through our institution’s subscriptions on the day of evaluation. Out of the 200 papers we found, 125 were available to download. Papers which used no implementations (e.g. theoretical papers), emulators and real test platforms (such as PlanetLab [Chun et al. 2003]), or where no simulator was mentioned, are omitted from our results (Figure 3).

Simulator	Papers
Custom	43
ns-2	8
Chord (SFS)	7
Javasim	2
PeerSim	2
Aurora	1
CSIM 19	1
ModelNet	1
Nab	1
Narses	1
NeuroGrid	1
P2PSim	1
SOSS	1

Fig. 2: Simulator usage pre-2007.

The most notable result from both time periods is that custom simulators, are predominantly used; ns-2 is the second most used simulator in both sets of results. It is worth mentioning that ns-2 is not a P2P simulator; it is a general network simulator. Therefore, a P2P simulation overlay needs to be written for it, and in most, if not all, cases this is a custom implementation (we found no specific mentions of openly available ns-2 P2P overlays). The usage of other open source simulators is more fragmented. Clearly, it can be seen that the research community prefers to write their own simulator software rather than using existing simulators. The consequence of this fact is that the repeatability of results, and the ability to build upon existing findings, is poor.

4. FEATURES QUESTIONNAIRE

We polled the P2P community to see how useful our criteria would be in a potential P2P simulator. We gathered information from two groups: the authors of the papers that

¹¹<http://scholar.google.co.uk>

Simulator	Papers
Custom	47 (Fig 4)
ns-2	7 (Fig 5)
PeerSim	4 (Fig 6)
SALSA	2 (Fig 7)
P2PSim	1 [Ghinita et al. 2007]
OMNeT++	1 [Awad et al. 2008]
NGS	1 [Webb and Soh 2007]
NEW	1 [Chen et al. 2008]
GnuSim	1 [Karakaya et al. 2008]

Fig. 3: Simulator usage post-2007.

[Annapureddy et al. 2007b; Zhou et al. 2007; Yunhao et al. 2007; Xie et al. 2008; Mol et al. 2008; Silva et al. 2008; Iliofotou et al. 2009; Pianese et al. 2007; Annapureddy et al. 2007a; Leonardi et al. 2008; Doukeridis et al. 2007a; Cheng et al. 2007; Chi et al. 2007; Rowaihy et al. 2007; Ahmed and Boutaba 2007; Belenkiy et al. 2007; Kantere et al. 2009; Huang et al. 2007b; Garbacki et al. 2007; Buchegger et al. 2009; Crainiceanu et al. 2007; Zhou and Hwang 2007; Bharambe et al. 2008; Terpstra et al. 2007; Zaharia and Keshav 2008; Wang et al. 2007; Guo et al. 2007; Hefeeda and Saleh 2008; Ren et al. 2008; Zhou et al. 2008; Zhuge and Li 2007; Haridasan and van Renesse 2008; Feng and Li 2008; Magharei and Rejaie 2009a; Liu 2007; Yang and Chen 2008; Do et al. 2008; Do et al. 2009; Boufkhad et al. 2008; Duminuco and Biersack 2008; Yiu et al. 2007; Doukeridis et al. 2007b; Skobeltsyn et al. 2009; Johansson et al. 2007; Huebsch 2008; Liu et al. 2008; Mol et al. 2007]

Fig. 4: Papers using custom P2P simulators post-2007

[Mavlankar et al. 2008; Magharei and Rejaie 2009b; Tanin et al. 2007; Eger et al. 2007; Huang et al. 2007a; Setton et al. 2008]

Fig. 5: Papers using ns-2 P2P simulators post-2007

[Laredo et al. 2008; Jelasity et al. 2007; Baldoni et al. 2007; Silverston et al. 2008]

Fig. 6: Papers using PeerSim P2P simulators post-2007

[Mittal and Borisov 2008; 2009]

Fig. 7: Papers using SALSA P2P simulators post-2007

we surveyed earlier, and also the users and developers of existing simulators. We prepared a questionnaire using the Likert scale [Likert 1932] to judge the importance of our criteria, and also gave free text areas to gain insight into individual users' requirements and concerns. The Likert scale is used to measure positiveness or negativeness of response to a question. The scale's response categories are as follows: (1) strongly disagree; (2) disagree; (3) neither agree nor disagree; (4) agree; and (5) strongly agree. The questionnaire was sent by email to the authors of our surveyed papers, and also to several active P2P simulator-related mailing lists¹². Although the inclusion of the P2P simulator-related mailing lists could have introduced bias in the sample, we argue that this is countered by the need to ensure that end users of P2P simulators were included in the sample and to maximise response to the questionnaire. We received a total of 81 responses, 45 of which were from the paper authors, and 36 from mailing list users. We now present the results of the questionnaire, organised into the categories they were presented in. In each section we provide a list of the statements we provided, along with a reference to our results. In each histogram that we present, the x -axis is labelled according to the response given. The y -axis represents the number of responses we received. Typically users were asked to rate questions according to how important they believe the feature to be. On the scale, 1 means not at all important, 3 is neutral and 5 means essential.

4.1. Simulation architecture and features

This section encompasses the design of the simulator and the features it should simulate. These questions represent what the high-level features of a P2P simulator should be. The questions were:

- How important is the ability to simulate in both query-cycle and discrete event manners? (Figure 8a)
- What language should the simulator be written in? (Figure 8b)
- How many P2P nodes would you expect to simulate? (Figure 8c)
- How important is the option of distributing the load of the simulator over multiple computers? (Figure 8d)
- How important is it that the simulator models the low-level details of the underlying network? (Figure 8e)

There was a trend towards having the option of simulating using either query-cycle or discrete event simulators (Figure 8a). The most popular languages for the implementation of the simulator were C/C++ and Java (Figure 8b), possibly due to the ubiquity

¹²These lists were: p2p-hackers, OverSim, PeerSim, ns, PlanetSim and PlanetLab.

uity of the languages in the community and also in teaching. We also allowed for free text responses, and we received two additional responses: one for Erlang and one for Lua. Generally speaking, our respondents preferred high-level imperative programming languages for implementation. Our respondents strongly preferred the smallest range of nodes for simulations (Figure 8c). This suggests that 0–100,000 nodes is sufficient for most P2P simulations. It is worth noting that 15 respondents are interested in P2P networks consisting of nodes in excess of 1M. Simulation is a computationally expensive activity. Thus, the need to distribute computation was also agreed upon (Figure 8d). This is especially important to those respondents interested in networks in excess of 1M nodes. There is also a strong trend towards wanting to realistically simulate the underlying network layer (Figure 8e). We hypothesise that this would give more credibility to any simulator results, as they would be similar to those that would be experienced on real networks.

4.2. Statistics and simulation run

This section considers the generation and collection of statistics from simulation runs. Current P2P simulators offer varying degrees of data logging in a variety of formats. We wanted to find out more detail about this element of P2P simulators. The questions were as follows:

- How important is it to generate statistics for a given run? (Figure 9a)
- Should different output formats be supported, such as xgraph, MATLAB, etc.? (Figure 9b)
- Should detailed log information be generated for every simulator run? (Figure 9c)
- Should it be possible to pause the simulation and inspect the current state at any time? (Figure 9d)
- How important is a specific simulator debugger that allows breakpoints and step-based simulation? (Figure 9e)

Generally speaking, any simulation aims to generate results in some form. Therefore, it is not surprising that Figure 9 shows strong positive responses for all questions. The questions related to pausing the simulation (Figure 9d) and debugging (Figure 9e) show a clear positive response, albeit not as strongly correlated as the others. Clearly, any tools that can assist the developer in locating and solving problems are welcome additions to any simulator. It is worth noting that a vast range of statistical functions may be desirable in order to analyse results, and a simulator may facilitate this by incorporating a statistical package such as SSPS¹³ or open source libraries that implement the same functionality.

4.3. Usability

The easier it is to implement an algorithm or scenario on a P2P simulator, the better. This leads to faster implementation and dissemination of results, and crucially, the minimisation of errors. Therefore we were interested in what features determine good usability. The questions were:

- How important is community support in the form of mailing lists, newsgroups and forums, for example? (Figure 10a)
- How important is a well-documented simulator API? (Figure 10b)
- How important is a specific declarative scripting language to generate overlays and events? (Figure 10c)

¹³<http://www.spss.com>

Fig. 8: Simulation architecture and feature responses

— How important are graphical visualisations during or after simulator runs? (Figure 10d)

Generally speaking, the more features available to increase usability, the better the reception of the simulator. Access to a well-documented API (Figure 10b) and also the use of a specific language (Figure 10c) are paramount. Community support is also favoured (Figure 10a) as an active community can help answer questions and also fix bugs in the simulator. Interestingly, the ability to visualise simulations shows uncer-

Fig. 9: Statistics and simulation run responses

tainty (Figure 10d). This is possibly because visualising very large networks is unfeasible due to the vast quantity of nodes and edges. In summary, visualisation may be a good debugging and demonstration tool and not necessarily a tool to view the entire simulated network.

4.4. Miscellaneous

In addition to the questions above, we asked each respondent their profession. As seen in Figure 11, the overwhelming majority of respondents were research students. This

Fig. 10: Usability Responses

Fig. 11: Profession of questionnaire respondents

may suggest that research students are the main workforce behind P2P simulator development. No industry researchers responded to the poll, however this could be because they are not authoring papers or subscribed to the mailing lists we chose.

5. SIMULATOR FEATURE COMPARISON

Using the simulator feature criteria discussed in Section 2.1 and compiled in Figures 8, 9, 10 and augmented with the data from our paper reviews, we present a comparison table of six simulators — P2PSim, PeerSim, PlanetSim, OverSim, ns (ns-2 and

ns-3), and SimGrid. We note that the list of these simulators is not authoritative and neither is it driven by popularity counts as measured by the number of papers referring to the said simulator. Rather, this list is driven by information gathered from our questionnaire and the aggregate personal experiences from using some of the simulators in a P2P context. Thus, for instance, while Figure 8-a indicates PlanetLab as a simulator, we do not include it in Table I because it is not a simulator as much as it is a reasonable substitute for observing program and code behaviour on a real network. Similarly, while Figure 8-a did not include any mention of SimGrid [Casanova et al. 2008] (see also <http://simgrid.gforge.inria.fr/>), we include it in our list since it is actively supported and in use.

The six simulators along with an evaluation criteria discussed in Section 2.1 are shown in Table I. Some inferences can easily be drawn from the data in Table I, although we stress that these are only inferences and not authoritative results. Most simulators appear to favour discrete event simulation; Java and C/C++ appear to be widely supported. SimGrid supports both of these languages in addition to Ruby and Lua bindings. All simulators support logging as well as some manner of textual output of statistics. None of the simulators listed in the table allow creating topologies using drag- and-drop palettes, though some of the simulators allow for a visualization component once the topology is set up and the simulation is running (although this visualization is rather basic).

Furthermore, while the simulation frameworks specifically designed for P2P networks (such as PeerSim and OverSim) inherently contain behaviour specific to overlay networks — churn, node misbehaviour, node failure, etc. — general purpose simulators such as NS-3 and SimGrid also allow such behaviour to be specified by the application programmer in the specific P2P behaviour model. And finally, although no P2P simulator contains all of the features in our proposed criteria, it can be seen that OverSim contains a majority of them. Older simulators such as P2PSim do not match the increased requirements that the community expects presently. It is also instructive to note that general purpose simulation frameworks such as NS-3 and SimGrid satisfy many of our evaluation criteria features. However, the implicit understanding here is that programmers using these simulation frameworks may have to spend extra time fine-tuning the application model to allow it to behave in a certain fashion instead of having a parametrized application model whose behaviour can be changed through the parameters.

6. CONCLUSION

We have presented an in-depth review of P2P papers investigating the use of simulators and simulations in this research.

By examining our new set of papers, since 2007 we have found a greater proportion of research being tested with real systems and real trace data. The P2P community appears to be focusing less on algorithms and more on network software optimisations such as streaming video (perhaps motivated by the inevitability of Internet Television).

An overwhelming body of work relies heavily on simulations to produce results and comparisons. The most alarming fact about our findings is that a large majority of highly cited P2P research papers produce their results on custom written simulators. These simulators are not publicly available. Most importantly this does not allow the research community to verify the results of these papers by repeating experiments. It also hinders the overall development of simulators and tools because others are not able to build upon existing work. This has a profound effect on the time taken for new research to emerge, as many academics and developers alike have to reproduce existing code that is not available, despite the fact that a wide variety of open source P2P simulators already exist. Prior to 2007, this issue was identified. The problem

Criteria	Simulator					
	P2PSim	PeerSim	PlanetSim	OverSim	ns	SimGrid
Simulation architecture and features.						
1. Query-cycle and discrete event	DE	DE, QC	DE	DE	DE	DE
2. Scripted Events	–	–	–	•	–	–
3. Random number seeding	•	•	–	•	•	•
4. Realistic transport layer	–	•	–	•	•	–
5. Bootstrapping features	–	–	•	•	–	–
6. Implementation language	C++	Java	Java	C++	C++, Python	C, C++ Ruby, Lua
Underlying network simulation.						
7. Topology generator	•	•	•	•	–	–
8. Latency and bandwidth	•	•	–	•	•	•
9. Various access links	–	–	–	•	•	•
11. Computation ability modelling	–	–	–	–	–	–
Behaviour.						
12. Misbehaviour	–	–	•	•	•	•
13. Message loss	•	•	•	•	•	•
14. Mobility models	–	–	–	•	•	•
15. Churn	•	•	–	•	•	•
16. Node failure	•	•	•	•	•	•
Statistics.						
17. Output formats	T	T, G, O	T, G, O	T	T, P	T
18. Logging	•	•	•	•	•	•
19. Graph properties	–	•	–	–	–	–
20. GraphML	–	•	•	–	–	–
Usability.						
21. Programming interfaces	–	–	•	•	•	•
22. Documentation	–	•	•	•	•	•
23. Community	•	•	•	•	•	•
24. Reference implementations	•	•	•	•	•	•
25. GUI	–	–	–	•	–	–
– information not known or not readily available DE = discrete-event, QC = query-cycle T = text, G = gml, O = other, X = XML, P = pcap						

Table I: Criteria present in simulators.

still persists today in similar proportions. Additionally, even though new simulators continue to be released, the research community tends towards only a handful of open source simulators. The demand for features in simulators, as shown by our criteria and survey, is high. Therefore, the community should work together to get these features in open source software. This would reduce the need for custom simulators, and hence increase repeatability and reputability of experiments.

One outlet to socialize the need for common simulation frameworks is the Internet Research Task Force (IRTF), a sister organization to the Internet Engineering Task

Force (IETF). IRTF meetings are co-located with IETF meetings, which are held three times a year. The locations of IETF meetings vary, but to facilitate the reach of Internet as widely as possible, IETF typically arranges meetings at Europe, North America and Asia during the calendar year. The attendees to IETF meetings include both academia and industry participants. Given the composition of the IETF attendees and the reach of the organization to all countries, we believe that socializing the need for a common simulation framework in IETF will have the most impact. To that extent, we have started working with the IETF community to distribute the results from this survey to the IRTF Peer-to-peer Research Group (P2PRG). We will continue working with the IRTF P2PRG to socialize, disseminate and publish the results.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their time, feedback, and a detailed reading of this manuscript. Their reviews have contributed greatly to the quality of the manuscript. The authors would like to thank Barney Livingston, Sethalat Rodhetbhai and Dan Chalmers for their inputs on an earlier related work.

REFERENCES

- AGOSTI, M., ZANICHELLI, F., AMORETTI, M., AND CONTE, G. 2008. P2pam: a framework for peer-to-peer architectural modeling based on peersim. In *Proceedings of the 1st International Conference on Simulation tools and techniques for communications, networks and systems & Workshops*. Simutools '08. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Marseille, France, 22:1–22:10.
- AHMED, R. AND BOUTABA, R. 2007. Distributed pattern matching: a key to flexible and efficient P2P search. *IEEE Journal on Selected Areas in Communications* 25, 1, 73–83.
- AMORETTI, M., AGOSTI, M., AND ZANICHELLI, F. 2009. Deus: a discrete event universal simulator. In *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*. Simutools '09. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Rome, Italy, 58:1–58:9.
- ANNAPUREDDY, S., GUHA, S., GKANTSIDIS, C., GUNAWARDENA, D., AND RODRIGUEZ, P. 2007a. Exploring VoD in P2P swarming systems. In *Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM)*. IEEE, Anchorage, Alaska, USA, 2571–2575.
- ANNAPUREDDY, S., GUHA, S., GKANTSIDIS, C., GUNAWARDENA, D., AND RODRIGUEZ, P. R. 2007b. Is high-quality VoD feasible using P2P swarming? In *Proceedings of the 16th International Conference on World Wide Web*. WWW '07. ACM, Banff, Alberta, Canada, 903–912.
- AWAD, A., GERMAN, R., AND DRESSLER, F. 2008. P2P-based routing and data management using the virtual cord protocol (VCP). In *Proceedings of the 9th ACM International Symposium on Mobile Ad-Hoc Networking and Computing*. MobiHoc '08. ACM, Hong Kong SAR, China, 443–444.
- BALDONI, R., BERALDI, R., QUEMA, V., QUERZONI, L., AND TUCCI-PIERGIOVANNI, S. 2007. TERA: topic-based event routing for peer-to-peer architectures. In *Proceedings of the 2007 Inaugural International Conference on Distributed Event-Based Systems*. DEBS '07. ACM, Toronto, Ontario, Canada, 2–13.
- BARCELLOS, M., MANSILHA, R., AND BRASILEIRO, F. 2008. Torrentlab: investigating bittorrent through simulation and live experiments. In *IEEE Symposium on Computers and Communications (ISCC)*. IEEE, Marrakech, Morocco, 507–512.
- BAUMGART, I., HEEP, B., AND KRAUSE, S. 2007. OverSim: A Flexible Overlay Network Simulation Framework. In *IEEE Global Internet Symposium*. IEEE, Anchorage, Alaska, USA, 79–84.
- BELENKIY, M., CHASE, M., ERWAY, C. C., JANNOTTI, J., KÜPÇÜ, A., LYSYANSKAYA, A., AND RACHLIN, E. 2007. Making P2P accountable without losing privacy. In *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*. WPES '07. ACM, Alexandria, Virginia, USA, 31–40.
- BHARAMBE, A., DOUCEUR, J. R., LORCH, J. R., MOSCIBRODA, T., PANG, J., SESHAN, S., AND ZHUANG, X. 2008. Donnybrook: enabling large-scale, high-speed, peer-to-peer games. In *Proceedings of the ACM SIGCOMM Conference on Data Communication*. SIGCOMM '08. ACM, Seattle, WA, USA, 389–400.
- BOUFKHAH, Y., MATHIEU, F., DE MONTGOLFIER, F., PERINO, D., AND VIENNOT, L. 2008. Achievable catalog size in peer-to-peer video-on-demand systems. In *Proceedings of the 7th International Conference on Peer-to-Peer Systems*. IPTPS'08. USENIX Association, Tampa Bay, Florida, 4–10.

- BUCHEGGER, S., SCHIÖBERG, D., VU, L.-H., AND DATTA, A. 2009. PeerSoN: P2P social networking: early experiences and insights. In *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*. SNS '09. ACM, Nuremberg, Germany, 46–52.
- CASANOVA, H., LEGRAND, A., AND QUINSON, M. 2008. SimGrid: a Generic Framework for Large-Scale Distributed Experiments. In *10th IEEE International Conference on Computer Modeling and Simulation*. IEEE, Cambridge, UK.
- CHEN, M., PONEC, M., SENGUPTA, S., LI, J., AND CHOU, P. A. 2008. Utility maximization in peer-to-peer systems. In *Proceedings of the 2008 ACM International Conference on Measurement and Modeling of Computer Systems*. SIGMETRICS '08. ACM, Annapolis, MD, USA, 169–180.
- CHENG, B., JIN, H., AND LIAO, X. 2007. Supporting VCR functions in P2P VoD services using ring-assisted overlays. In *IEEE International Conference on Communications*. ICC '07. IEEE, Glasgow, Scotland, 1698–1703.
- CHI, H., ZHANG, Q., JIA, J., AND SHEN, X. 2007. Efficient search and scheduling in P2P-based media-on-demand streaming service. *IEEE Journal on Selected Areas in Communications* 25, 1, 119–130.
- CHUN, B., CULLER, D., ROSCOE, T., BAVIER, A., PETERSON, L., WAWRZONIAK, M., AND BOWMAN, M. 2003. PlanetLab: an overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communication Review* 33, 3, 3–12.
- CLARKE, I., SANDBERG, O., WILEY, B., AND HONG, T. W. 2000. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *Workshop on Design Issues in Anonymity and Unobservability*. Vol. 320. Springer, Berkeley, California, USA.
- CRAINICEANU, A., LINGA, P., MACHANAVAJJHALA, A., GEHRKE, J., AND SHANMUGASUNDARAM, J. 2007. P-ring: an efficient and robust P2P range index structure. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. SIGMOD '07. ACM, Beijing, China, 223–234.
- DABEK, F., ZHAO, B., DRUSCHEL, P., KUBIATOWICZ, J., AND STOICA, I. 2003. Towards a common API for structured peer-to-peer overlays. *Peer-to-Peer Systems II* 58, 33–44.
- DE VOGEELEER, K., ERMAN, D., AND POPESCU, A. 2008. Simulating bittorrent. In *Proceedings of the 1st International Conference on Simulation tools and techniques for communications, networks and systems & Workshops*. Simutools '08. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Marseille, France, 2:1–2:7.
- DEAN, J. AND GHEMAWAT, S. 2008. MapReduce: simplified data processing on large clusters. *Commun. ACM* 51, 107–113.
- DECANDIA, G., HASTORUN, D., JAMPANI, M., KAKULAPATI, G., LAKSHMAN, A., PILCHIN, A., SIVASUBRAMANIAN, S., VOSSHALL, P., AND VOGELS, W. 2007. Dynamo: Amazon's highly available key-value store. *ACM SIGOPS Operating Systems Review* 41, 6, 205–220.
- DEERING, S. AND HINDEN, R. 1998. Internet Protocol, Version 6 (IPv6) specification.
- DO, T., HUA, K., JIANG, N., AND LIU, F. 2009. PatchPeer: A scalable video-on-demand streaming system in hybrid wireless mobile peer-to-peer networks. *Peer-to-Peer Networking and Applications* 2, 182–201.
- DO, T., HUA, K., AND TANTAOU, M. 2008. Robust video-on-demand streaming in peer-to-peer environments. *Computer Communications* 31, 3, 506–519.
- DOULKERIDIS, C., NORVAG, K., AND VAZIRGIANNIS, M. 2007a. DESENT: decentralized and distributed semantic overlay generation in P2P networks. *IEEE Journal on Selected Areas in Communications* 25, 1, 25–34.
- DOULKERIDIS, C., VLACHOU, A., KOTIDIS, Y., AND VAZIRGIANNIS, M. 2007b. Peer-to-peer similarity search in metric spaces. In *Proceedings of the 33rd International Conference on Very Large Databases*. VLDB '07. VLDB Endowment, Vienna, Austria, 986–997.
- DUMINUCO, A. AND BIRSACK, E. 2008. Hierarchical codes: how to make erasure codes attractive for peer-to-peer storage systems. In *Proceedings of the 2008 Eighth International Conference on Peer-to-Peer Computing*. IEEE Computer Society, Washington, DC, USA, 89–98.
- DUPUY, A., SCHWARTZ, J., YEMINI, Y., AND BACON, D. 1990. Nest: a network simulation and prototyping testbed. *Communications of the ACM* 33, 10, 63–74.
- EGER, K., HOSSFELD, T., BINZENHOFER, A., AND KUNZMANN, G. 2007. Efficient simulation of large-scale P2P networks: packet-level vs. flow-level simulations. In *Proceedings of the Second Workshop on Use of P2P, GRID and Agents for the Development of Content Networks*. UPGRADE '07. ACM, Monterey, California, USA, 9–16.
- FENG, C. AND LI, B. 2008. On large-scale peer-to-peer streaming systems with network coding. In *Proceedings of the 16th ACM International Conference on Multimedia*. MM '08. ACM, Vancouver, British Columbia, Canada, 269–278.

- GARBACKI, P., EPEMA, D. H. J., AND VAN STEEN, M. 2007. An amortized tit-for-tat protocol for exchanging bandwidth instead of content in P2P networks. In *Proceedings of the First International Conference on Self-Adaptive and Self-Organizing Systems*. SASO '07. IEEE Computer Society, Washington, DC, USA, 119–128.
- GHINITA, G., KALNIS, P., AND SKIADOPOULOS, S. 2007. MOBIHIDE: a mobile peer-to-peer system for anonymous location-based queries. In *Proceedings of the 10th International Conference on Advances in Spatial and Temporal Databases*. SSTD'07. Springer-Verlag, Boston, MA, USA, 221–238.
- GROSS, C., LEHN, M., STINGL, D., KOVACEVIC, A., BUCHMANN, A., AND STEINMETZ, R. 2010. Towards a common interface for overlay network simulators. In *16th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, Shanghai, China, 59–66.
- GUO, L., CHEN, S., XIAO, Z., TAN, E., DING, X., AND ZHANG, X. 2007. A performance study of BitTorrent-like peer-to-peer systems. *IEEE Journal on Selected Areas in Communications* 25, 1, 155–169.
- GURBANI, V. K., BASU, A., SCHMIDT, T., FLEMING, S., KOLBERG, M., AND WAEHLISCH, M. 2011. Peer-to-peer simulation frameworks: a survey, IETF Internet-Draft, Work in Progress, draft-gurbani-p2prg-simulation-survey-00.
- HARIDASAN, M. AND VAN RENESSE, R. 2008. Gossip-based distribution estimation in peer-to-peer networks. In *Proceedings of the 7th International Conference on Peer-to-Peer Systems*. IPTPS'08. USENIX Association, Tampa Bay, Florida, 13–19.
- HEFEEDA, M. AND SALEH, O. 2008. Traffic modeling and proportional partial caching for peer-to-peer systems. *IEEE/ACM Trans. Netw.* 16, 1447–1460.
- HENDERSON, T. 2008. The NS3 discrete event simulator.
- HENDLER, J. AND GOLBECK, J. 2008. Metcalfe's law, Web 2.0, and the semantic web. *Web Semantics Science Services and Agents on the World Wide Web* 6, 1, 14–20.
- HUANG, C.-M., HSU, T.-H., AND HSU, M.-F. 2007a. Network-aware P2P file sharing over the wireless mobile networks. *IEEE Journal on Selected Areas in Communications* 25, 1, 204–210.
- HUANG, Y., FARN CHEN, Y., JANA, R., JIANG, H., RABINOVICH, M., REIBMAN, A., WEI, B., AND XIAO, Z. 2007b. Capacity analysis of MediaGrid: a P2P IPTV platform for fiber to the node networks. *IEEE Journal on Selected Areas in Communications* 25, 1, 131–139.
- HUEBSCH, R. J. 2008. PIER: Internet scale P2P query processing with distributed hash tables. Ph.D. thesis, EECS Department, University of California, Berkeley.
- ILIOFOTOU, M., KIM, H.-C., FALOUTSOS, M., MITZENMACHER, M., PAPPU, P., AND VARGHESE, G. 2009. Graph-based P2P traffic classification at the internet backbone. In *Proceedings of the 28th IEEE International Conference on Computer Communications Workshops*. INFOCOM '09. IEEE Press, Rio de Janeiro, Brazil, 37–42.
- JELASITY, M., VOULGARIS, S., GUERRAOUT, R., KERMARREC, A.-M., AND VAN STEEN, M. 2007. Gossip-based peer sampling. *ACM Transactions on Computers Systems* 25, 8:1–8:36.
- JENNINGS, C., LOWEKAMP, B., RESCORLA, E., BASET, S. A., AND SCHULZRINNE, H. 2012. REsource LOcation And Discovery (RELOAD) Base Protocol. IETF Internet-Draft (work-in-progress), draft-ietf-p2psip-base-20.
- JOHANSSON, B., RABI, M., AND JOHANSSON, M. 2007. A simple peer-to-peer algorithm for distributed optimization in sensor networks. In *46th IEEE Conference on Decision and Control*. IEEE, New Orleans, Louisiana, USA, 4705–4710.
- KANTERE, V., TSOUMAKOS, D., SELLIS, T., AND ROUSSOPOULOS, N. 2009. GrouPeer: dynamic clustering of P2P databases. *Inf. Syst.* 34, 62–86.
- KARAKAYA, M., KÖRPEOĞLU, I., AND ULUSOY, O. 2008. Counteracting free riding in peer-to-peer networks. *Comput. Netw.* 52, 675–694.
- KARGER, D., LEHMAN, E., LEIGHTON, T., PANIGRAHY, R., LEVINE, M., AND LEWIN, D. 1997. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*. ACM, El Paso, Texas, USA, 654–663.
- KLINGBERG, T. AND MANFREDI, R. 2002. Gnutella 0.6.
- KULBAK, Y. AND BICKSON, D. 2005. The eMule protocol specification.
- LAKSHMAN, A. AND MALIK, P. 2010. Cassandra: a decentralized structured storage system. *SIGOPS Operating Systems Review* 44, 35–40.
- LAREDO, J. L., EIBEN, A. E., STEEN, M., CASTILLO, P. A., MORA, A. M., AND MERELO, J. J. 2008. P2P evolutionary algorithms: A suitable approach for tackling large instances in hard optimization problems. In *Proceedings of the 14th International Euro-Par Conference on Parallel Processing*. Euro-Par '08. Springer-Verlag, Las Palmas de Gran Canaria, Spain, 622–631.

- LEONARDI, E., MELLIA, M., HORVATH, A., MUSCARIELLO, L., NICCOLINI, S., AND ROSSI, D. 2008. Building a cooperative P2P-TV application over a wise network: the approach of the European FP-7 strep NAPA-WINE. *IEEE Communications Magazine* 46, 4, 20–22.
- LEONINI, L., RIVIÈRE, É., AND FELBER, P. 2009. Splay: distributed systems evaluation made simple (or how to turn ideas into live systems in a breeze). In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*. USENIX Association, Boston, MA, USA, 185–198.
- LIKERT, R. 1932. A technique for the measurement of attitudes. *Archives of Psychology* 22, 40, 1–55.
- LIN, S., PAN, A., GUO, R., AND ZHANG, Z. 2005. Simulating large-scale p2p systems with the wids toolkit. In *Proceedings of the 13th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*. IEEE, Atlanta, GA, USA, 415–424.
- LIU, Y. 2007. On the minimum delay peer-to-peer video streaming: how realtime can it be? In *Proceedings of the 15th International Conference on Multimedia*. MULTIMEDIA '07. ACM, Augsburg, Germany, 127–136.
- LIU, Z., SHEN, Y., ROSS, K., PANWAR, S., AND WANG, Y. 2008. Substream trading: towards an open P2P live streaming system. In *IEEE International Conference on Network Protocols*. INCP '08. IEEE, Orlando, FL, USA, 94–103.
- MAGHAREI, N. AND REJAIE, R. 2009a. Overlay monitoring and repair in swarm-based peer-to-peer streaming. In *Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video*. NOSSDAV '09. ACM, Williamsburg, VA, USA, 25–30.
- MAGHAREI, N. AND REJAIE, R. 2009b. PRIME: Peer-to-peer receiver-driven mesh-based streaming. *IEEE/ACM Transactions on Networking* 17, 4, 1052–1065.
- MAVLANKAR, A., NOH, J., BACCICHET, P., AND GIROD, B. 2008. Peer-to-peer multicast live video streaming with interactive virtual pan/tilt/zoom functionality. In *15th IEEE International Conference on Image Processing*. ICIP '08. IEEE, San Diego, CA, USA, 2296–2299.
- MAYMOUNKOV, P. AND MAZIERES, D. 2002. Kademlia: A Peer-to-Peer Information System Based on the XOR Metric. *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)* 258, 263.
- MCCANNE, S. AND FLOYD, S. 1999. The Network Simulator - ns2.
- MITTAL, P. AND BORISOV, N. 2008. Information leaks in structured peer-to-peer anonymous communication systems. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*. CCS '08. ACM, Alexandria, Virginia, USA, 267–278.
- MITTAL, P. AND BORISOV, N. 2009. ShadowWalker: peer-to-peer anonymous communication using redundant structured topologies. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*. CCS '09. ACM, Chicago, Illinois, USA, 161–172.
- MOL, J.-D., EPEMA, D. H. J., AND SIPS, H. J. 2007. The orchard algorithm: building multicast trees for P2P video multicasting without free-riding. *IEEE Transactions on Multimedia* 9, 1593–1604.
- MOL, J.-D., POUWELSE, J. A., EPEMA, D. H. J., AND SIPS, H. J. 2008. Free-riding, fairness, and firewalls in P2P file-sharing. In *Proceedings of the Eighth International Conference on Peer-to-Peer Computing*. IEEE Computer Society, Washington, DC, USA, 301–310.
- NAICKEN, S., BASU, A., LIVINGSTON, B., AND RODHETBHAI, S. 2006a. A Survey of Peer-to-Peer Network Simulators. In *Proceedings of The Seventh Annual Postgraduate Symposium (PGNet)*. Liverpool John Moore University, Liverpool, UK.
- NAICKEN, S., BASU, A., LIVINGSTON, B., RODHETBHAI, S., AND WAKEMAN, I. 2006b. Towards Yet Another Peer-to-Peer Simulator. In *Proceedings of The Fourth International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks (HET-NETs)*. University of Bradford, Ilkley, UK.
- NAICKEN, S., LIVINGSTON, B., BASU, A., RODHETBHAI, S., WAKEMAN, I., AND CHALMERS, D. 2007. The State of Peer-to-Peer Simulators and Simulations. *ACM Computer Communications Review* 37(2), 95–98.
- PIANESE, F., PERINO, D., KELLER, J., AND BIERSECK, E. 2007. PULSE: An adaptive, incentive-based, unstructured P2P live streaming system. *IEEE Transactions on Multimedia* 9, 8, 1645–1660.
- RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., AND SHENKER, S. 2001. A scalable content-addressable network. In *Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. SIGCOMM '01. ACM, San Diego, California, United States, 161–172.
- REN, D., LI, Y.-T., AND CHAN, S.-H. 2008. On reducing mesh delay for peer-to-peer live streaming. In *The 27th IEEE Conference on Computer Communications*. INFOCOM '08. IEEE, Phoenix, AZ, USA, 1058–1066.

- RISSON, J. AND MOORS, T. 2006. Survey of research towards robust peer-to-peer networks: search methods. *Comput. Netw.* 50, 3485–3521.
- ROSENBERG, J., SCHULZRINNE, H., CAMARILLO, G., JOHNSTON, A., PETERSON, J., SPARKS, R., HANDLEY, M., AND SCHOOLER, E. 2002. SIP: Session Initiation Protocol. RFC 3261 (Proposed Standard).
- ROWAIHY, H., ENCK, W., MCDANIEL, P., AND LA PORTA, T. 2007. Limiting Sybil attacks in structured P2P networks. In *26th IEEE International Conference on Computer Communications*. INFOCOM '07. IEEE, Anchorage, Alaska, USA, 2596–2600.
- ROWSTRON, A. I. T. AND DRUSCHEL, P. 2001. 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*. Middleware '01. Springer-Verlag, London, UK, 329–350.
- SETTON, E., BACCICHET, P., AND GIROD, B. 2008. Peer-to-peer live multicast: A video perspective. *Proceedings of the IEEE* 96, 1, 25–38.
- SILVA, A. P. C. D., LEONARDI, E., MELLIA, M., AND MEO, M. 2008. A bandwidth-aware scheduling strategy for P2P-TV systems. In *Proceedings of the 2008 Eighth International Conference on Peer-to-Peer Computing*. IEEE Computer Society, Washington, DC, USA, 279–288.
- SILVERSTON, T., FOURMAUX, O., AND CROWCROFT, J. 2008. Towards an incentive mechanism for peer-to-peer multimedia live streaming systems. In *Proceedings of the 2008 Eighth International Conference on Peer-to-Peer Computing*. IEEE Computer Society, Washington, DC, USA, 125–128.
- SKOBELOTSYN, G., LUU, T., PODNAR ARKO, I., RAJMAN, M., AND ABERER, K. 2009. Query-driven indexing for scalable peer-to-peer text retrieval. *Future Gener. Comput. Syst.* 25, 89–99.
- STINGL, D., GROSS, C., RUCKERT, J., NOBACH, L., KOVACEVIC, A., AND STEINMETZ, R. 2011. Peerfactsim.kom: A simulation framework for peer-to-peer systems. In *2011 International Conference on High Performance Computing and Simulation (HPCS)*. IEEE, Istanbul, Turkey, 577–584.
- STOICA, I., MORRIS, R., KARGER, D., KAASHOEK, M. F., AND BALAKRISHNAN, H. 2001. Chord: A scalable peer-to-peer lookup service for internet applications. *SIGCOMM Comput. Commun. Rev.* 31, 149–160.
- TANIN, E., HARWOOD, A., AND SAMET, H. 2007. Using a distributed quadtree index in peer-to-peer networks. *The VLDB Journal* 16, 165–178.
- TERPSTRA, W. W., KANGASHARJU, J., LENG, C., AND BUCHMANN, A. P. 2007. Bubblestorm: resilient, probabilistic, and exhaustive peer-to-peer search. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. SIGCOMM '07. ACM, Kyoto, Japan, 49–60.
- URBAN, P., DÉFAGO, X., AND SCHIPER, A. 2001. Neko: A single environment to simulate and prototype distributed algorithms. In *Proceedings of the 15th International Conference on Information Networking*. IEEE, Beppu City, Oita, Japan, 503–511.
- VAN RENESSE, R., BIRMAN, K. P., AND VOGELS, W. 2003. Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining. *ACM Trans. Comput. Syst.* 21, 164–206.
- WANG, S., OOI, B. C., TUNG, A., AND XU, L. 2007. Efficient skyline query processing on peer-to-peer networks. In *IEEE 23rd International Conference on Data Engineering*. ICDE '07. IEEE, Istanbul, Turkey, 1126–1135.
- WANG, Y., CARZANIGA, A., AND WOLF, A. 2008. Four enhancements to automated distributed system experimentation methods. In *Proceedings of the 30th International Conference on Software Engineering*. ACM, Leipzig, Germany, 491–500.
- WEBB, S. AND SOH, S. 2007. Round length optimisation for p2p network gaming. In *8th Postgraduate Electrical Engineering and Computing Symposium*. Curtin University, Perth, Australia, 23–28.
- WILCOX-O'HEARN, Z. AND WARNER, B. 2008. Tahoe: the least-authority filesystem. In *Proceedings of the 4th ACM International Workshop on Storage Security and Survivability*. StorageSS '08. ACM, Alexandria, Virginia, USA, 21–26.
- XIE, H., YANG, Y. R., KRISHNAMURTHY, A., LIU, Y. G., AND SILBERSCHATZ, A. 2008. P4P: provider portal for applications. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*. SIGCOMM '08. ACM, Seattle, WA, USA, 351–362.
- YANG, S. J. AND CHEN, I. Y. 2008. A social network-based system for supporting interactive collaboration in knowledge sharing over peer-to-peer network. *International Journal of Human-Computer Studies* 66, 1, 36–50.
- YIU, W.-P., JIN, X., AND CHAN, S.-H. 2007. VMesh: Distributed segment storage for peer-to-peer interactive video streaming. *IEEE Journal on Selected Areas in Communications* 25, 9, 1717–1731.
- YUNHAO, L., XIAO, L., AND NI, L. 2007. Building a scalable bipartite P2P overlay network. *IEEE Transactions on Parallel and Distributed Systems* 18, 9, 1296–1306.

- ZAHARIA, M. AND KESHAV, S. 2008. Gossip-based search selection in hybrid peer-to-peer networks. *Concurr. Comput.: Pract. Exper.* 20, 139–153.
- ZANTOUT, B. AND HARATY, R. 2011. I2p data communication system. In *Proceedings of the 10th International Conference on Networks*. IARIA, St. Maarten, The Netherlands Antilles, 401–409.
- ZHANG, D., CHEN, X., AND YANG, H. 2009. State of the art and challenges on peer-to-peer simulators. In *5th International Conference on Wireless Communications, Networking and Mobile Computing*. IEEE, Beijing, China, 1–4.
- ZHAO, B., KUBIATOWICZ, J., AND JOSEPH, A. 2001. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. *Computer* 74, 11–20.
- ZHOU, R. AND HWANG, K. 2007. Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing. *IEEE Transactions on Parallel and Distributed Systems* 18, 460–473.
- ZHOU, R., HWANG, K., AND CAI, M. 2008. GossipTrust for fast reputation aggregation in peer-to-peer networks. *IEEE Trans. on Knowl. and Data Eng.* 20, 1282–1295.
- ZHOU, Y., CHIU, D. M., AND LUI, J. 2007. A simple model for analyzing P2P streaming protocols. In *IEEE International Conference on Network Protocols*. ICNP '07. IEEE, Beijing, China, 226–235.
- ZHU, W., WANG, C.-L., AND LAU, F. 2002. JESSICA2: a distributed Java Virtual Machine with transparent thread migration support. In *IEEE International Conference on Cluster Computing*. IEEE, Chicago, USA, 381–388.
- ZHUGE, H. AND LI, X. 2007. Peer-to-peer in metric space and semantic space. *IEEE Transactions on Knowledge and Data Engineering* 19, 6, 759–771.

Received August 2011; revised Date; accepted July 2012