2 P2P or Not 2 P2P?

Mema Roussopoulos[1], Mary Baker[2], David S. H. Rosenthal[3], TJ Giuli[4], Petros Maniatis[5], and Jeff Mogul[2]

[1] Harvard University, Cambridge, MA
[2] HP Labs, Palo Alto, CA
[3] Stanford University Libraries, Stanford CA
[4] Stanford University, Stanford, CA
[5] Intel Research, Berkeley, CA

Abstract. In the hope of stimulating discussion, we present a heuristic decision tree that designers can use to judge how suitable a P2P solution might be for a particular problem. It is based on characteristics of a wide range of P2P systems from the literature, both proposed and deployed. These include budget, resource relevance, trust, rate of system change, and criticality.

1 Introduction

Academic research in peer-to-peer (P2P) systems has concentrated largely on algorithms to improve the efficiency [31], scalability [22], robustness [12], and security [33] of query routing in P2P systems, services such as indexing and search [20], dissemination [17], and rendezvous [27] [30] for applications running on top of these systems, or even many of the above [18]. While these improvements may be essential to enhancing the performance of some P2P applications, there has been little focus on what makes a problem "P2P-worthy," or on which other, previously ignored problems may benefit from the application of P2P techniques. What questions should a system designer ask to judge whether a P2P solution is appropriate for his particular problem?

In this position paper, we hope to stimulate discussion by distilling the experience of a broad range of proposed and deployed P2P systems into a methodology for judging how suitable a P2P architecture might be for a particular problem. In Section 2, we identify some salient characteristics axes in typical distributed problems. In Section 3, we describe a spectrum of specific problems for which P2P solutions have been proposed. In Section 4, we propose an arrangement of problem characteristics into a heuristic decision tree. We walk through the tree explaining its choices and why we believe certain paths may lead to successful P2P solutions to important problems, while other paths may encounter difficulties. While any particular set of characteristics axes or fixed decision graph may be inadequate for all purposes, we present the arrangement that has proved most useful in our work so far.

2 Problem Characteristics Axes

In this section, we describe the characteristics we believe are important in assessing the P2P-worthiness of distributed problems. Paraphrasing the call for papers of this workshop, we identify as peer-to-peer those environments that satisfy the following three criteria:

- *Self-organizing*: Nodes organize themselves into a network through a discovery process. There is no global directory of peers or resources.
- *Symmetric communication*: Peers are considered equals; they both request and offer services, rather than being confined to either client or server roles.
- *Decentralized control*: Peers determine their level of participation and their course of action autonomously. There is no central controller that dictates behavior to individual nodes.

Milojičić et al. [26] identify similar criteria.

Our axes are the problem's budget, the relevance of resources to individual peers, the rate of system change, the need for mutual trust, and the criticality of the problem. In more detail:

Budget: If the budget for a centrally controlled solution is ample, a designer is unlikely to consider worthwhile the inefficiencies, latencies and testing problems of a P2P solution. If the budget is limited, a key motivator in the choice of P2P architectures is the lowest possible cost of entry for individual peers, despite increased total system cost. Assembling a system from local, often surplus, components can be justified as a small part of many budgets and may be the only economically feasible approach.

Resource relevance to participants: Relevance is the likelihood that a "unit of service" within a problem (e.g., a single file in a file sharing problem) is interesting to many participants. When resource relevance is high, cooperation in a P2P solution evolves naturally. If relevance is low, cooperation may require artificial or extrinsic incentives to make a possible P2P solution viable.

Trust: The cost to a P2P system of handling mutually distrusting peers is high. Distrust may be a necessary evil of the problem, or it may be desirable as a means of imposing fault isolation throughout a peer community to reduce the risks posed by misbehaving peers.

Rate of system change: Different problems have different requirements for timeliness and consistency. Problems or solutions with high rates of change in the participants, the data or the system parameters make it difficult to meet high requirements for timeliness and consistency.

Criticality: If the problem being solved is critical to the users, they may demand centralized control and accountability irrespective of technical criteria. Even if a P2P solution is not ruled out, the need for expensive fault-tolerance or massive over-provisioning may make it uneconomic.

One question that arises is why we do not consider the physical constraints of a problem along with the budget; some problems have physical constraints such as scale or geographic size that require a distributed solution, regardless of budget. One example is latency due to the speed of light in an interplanetary

internet [3]. However, we have found no examples of problems that also require decentralized control or self-organizing peers.

We have excluded other characteristics which, while potentially important, did not enter into this decision tree as far as we have elaborated it. First, it may be important whether resources are public or private; private resources requiring confidentiality may be more difficult to protect and manage in P2P systems, and they may have less relevance to participants. Second, it may be important whether resources are naturally distributed; resources that exist naturally in many places, such as the usage statistics of many individual networks, may be more amenable to a distributed solution, and even a P2P solution.

3 Candidate Problems

We analyze a variety of problems with proposed P2P solutions to determine which of our characteristics they exhibit. These problems come from routing, backup, monitoring, data sharing, data dissemination, and auditing.

3.1 Routing Problems

All distributed systems need a routing layer to get messages to their intended recipients. Routing takes on P2P characteristics when the scale is large enough (e.g., the Internet) or when centralization is ruled out (e.g., wireless ad hoc networks).

Internet Routing Internet routers must communicate to cope with a dynamically changing network topology to determine how to route outbound packets to their destination. They are arranged into "autonomous systems" which "peer" with each other across organizational boundaries, frequently between competitors.

Routing protocols have historically assumed that economic incentives and legal contracts are sufficient to discourage misbehavior. At the application layer (e.g., Resilient Overlay Networks (RON) [1]) or at the network layer (e.g., BGP [21]), routers trust information from known peers. They cooperate because the information being exchanged is relevant to all peers and important to their function. This cooperation tends to fail if error, misbehavior or usage patterns cause the data to change too fast. To scale to the size of the Internet, BGP tries to limit the rate of change by aggregating routes instead of having ISPs propagate internal routing updates. Aggregation reduces the ability to detect path outages quickly [19]. RON instead gives up scaling to large numbers of nodes in favor of more fine-grained route information exchanges.

Ad hoc Routing in Disaster Recovery The ad hoc routing problem is to use transient resources, such as the wireless communication devices of a disaster recovery crew, to deploy temporary network infrastructure for a specific

purpose. Because each individual node's wireless range does not reach all other nodes, peers in the network forward packets on behalf of each other. The costly alternative is to provide more permanent infrastructure for all possible eventualities in all possible locations. The network is of relevance and critical to all participants, and pre-configured security can give a high level of mutual trust. Once established, the participants (humans in the crew) typically change and move slowly, and do not exchange huge volumes of data.

Metropolitan-area Cell Phone Forwarding Ad hoc routing has also been proposed in less critical settings, such as that of public, ad hoc cellular telephony in dense metropolitan areas. The motivation is to reduce the need for base stations, to use the radio spectrum more efficiently, and to avoid payment for air time where traffic does not pass through base stations. Unlike the disaster recovery problem, the participants do not trust each other and they change and move rapidly. In its current state, this problem suffers from the "Tragedy of the Commons" [14]. We doubt that a practical P2P solution to this problem exists, unless either on-going research [2, 4] devises strong, "strategy-proof" mechanisms to combat selfishness, or the scope of the problem is limited to close-knit communities with inherent incentives for participation.

3.2 Backup

Backup, the process in which a user replicates his files in different media at different locations to increase data survivability, can benefit greatly from the pooling of otherwise underutilized resources. Unfortunately, the fact that each peer is interested only in its own data opens the way to selfish peer behavior.

Internet Backup The cost of backup could be reduced if Internet-wide cooperation [9, 10] could be fostered and enforced. For example in Samsara [9] peers must hold real or simulated data equivalent to the space other peers hold for them. But there is no guarantee an untrusted node will provide backup data when requested, even if it has passed periodic checks to ensure it still has those data. Such a misbehaving or faulty node may in turn have its backup data elsewhere dropped in retaliation. If misbehaving, it may already have anticipated this reaction and, if faulty, this is exactly why it would participate in a backup scheme in the first place. We believe that data backup is poorly suited for a P2P environment running across trust boundaries.

Corporate Backup In contrast, when participants enjoy high mutual trust, e.g., within the confines of an enterprise intranet, P2P backup makes sense (HiveCache [15] is one such commercial offering). This is because selfish behavior is unlikely when a sense of trusting community or a top-down corporate mandate obviate the need for enforceable compliance incentives.

3.3 Distributed Monitoring

Monitoring is an important task in any large distributed system. It may have simple needs such as "subscribing" to first-order events and expecting notification when those events are "published" (e.g., Scribe [27]); it may involve more complicated, on-line manipulation, for instance via SQL queries, of complex distributed data streams such as network packet traces, CPU loads, virus signatures (as in the on-line network monitoring problem motivating PIER [16]); it may be the basis for an off-line, post mortem longitudinal study of many, high-volume data streams, such as the longitudinal network studies performed by Fomenkov et al. [11].

Although the abstract monitoring problem is characterized by natural distribution of the data sources monitored, specific instances of the problem vary vastly. A longitudinal off-line network study, though important, is not necessarily critical to its recipients, and has low timeliness constraints. In contrast, an ISP may consider the on-line, on-time monitoring of its resources and those of its neighbors extremely critical for its survival. Similarly, the mechanisms for complex network monitoring described by Huebsch et al. [16] may be appropriate for administratively closed, high-trust environments such as PlanetLab [6], and they may be quite inappropriate in environments lacking mutual trust and rife with fraud or subversion. In contrast, an off-line long-term network study affords its investigators more time for validating data against tampering.

3.4 Data Sharing

File sharing In file sharing systems, participants offer their local files to other peers and search collections to find interesting files. The cost of deployment is very low since most peers store only items that they are interested in anyway. Resource relevance is high; a great deal of content appeals to a large population of peers. In typical file sharing networks, peer turnover and file addition is high, leading to a high rate of system change. Peers trust each other to deliver the advertised content and most popular file sharing networks do not have the capacity to resist malicious peers. File sharing is mainly used to trade media content, which is not a critical application.

Censorship Resistance The goal of the FreeNet project [7] is to create an anonymous, censorship-resistant data store. Both publishing and document requests are routed through a mix-net [5] and all content is encrypted by the content's creator. These steps are necessary because peers are mutually suspicious and some peers may be malicious. Peers share their bandwidth as well as disk space, which means that the cost of entry is low, promoting incremental rapid growth; this growth is unstructured, which strengthens the system against legal attacks. FreeNet is intended to provide a medium for material that some group wishes to suppress, thus data are relevant to publishers, readers and attackers alike. Fortunately, censorship-resistance does not require immediate availability, making this a low rate-of-change problem.

Tangler [32] has similar goals. A peer stores a document by encoding it using erasure codes and distributing the resulting fragments throughout the community. To prevent an adversary from biasing where those fragments are distributed, a peer combines its document with pseudo-randomness derived from other peers' documents before erasure coding. To retrieve its own document, a peer must store this randomness (i.e., other peers' documents) locally. Although the problem lacks inherent incentives for participation, this solution ingeniously supplies them.

3.5 Data Dissemination

Data dissemination is akin to data sharing, with the distinction that the problem is not to *store* data indefinitely but merely to *spread* the data for a relatively short amount of time. Often storing is combined with spreading.

Usenet Usenet, perhaps the oldest and most successful P2P application, is a massively distributed discussion system in which users post messages to "newsgroups." These articles are then disseminated to other hosts subscribing to the particular newsgroup, and made available to local users. Usenet has been a staple of the Internet for decades, arguably because of the low cost barrier to peer entry and the high relevance of the content to participating peers. Unfortunately, although the system flourished at a time when mutual trust was assumed, it remains vulnerable to many forms of attack, perhaps jeopardizing its future in less innocent times.

Non-critical Content Distribution Dissemination of programs, program updates, streaming media [8, 17], and even cooperative web caching [34] are all non-critical content distribution problems.

One successful application is BitTorrent [8], which mitigates the congestion at a download server when relevant (i.e., popular) but non-critical new resources such as programs or updates are posted. Its tit-for-tat policy is effective despite low peer trust.

Cooperative Web caching, although superficially attractive, has not succeeded, for complex and subtle reasons [34]. Although it offered some benefits for large organizations in very low latency environments, and for low-relevance (i.e., unpopular) documents, those benefits were only marginal.

Critical Flash Crowds Other specific instances of dissemination have been proposed to address flash crowds [28, 29], which could be used to distribute critical data, such as news updates during a major disaster.

3.6 Auditing

Digital Preservation The *LOCKSS* system preserves academic e-journals in a network of autonomous web caches. Peers each obtain their own complete replicas of the content by crawling the publisher's web site. If the content becomes

unavailable from the publisher, the local copy is supplied to local readers. Slowly, in the background, a P2P "opinion poll" protocol [24] provides mutual audit and repairs any damage it detects. Peers trust the consensus of other peers but not any individual peer. Mutual distrust is essential to prevent cascade failures which could destroy every copy of the preserved content. The automatic audit and repair process allows peers to be built from cheap, unreliable hardware with very little need for administration, an important factor given library budgets. Publication of new content and damage to preserved content causes system change; the rates of both are limited. The content preserved is highly relevant to many peers.

Distributed Time Stamping A secure time stamping service [13] acts as the digital equivalent of a notary public: it maintains a history of the creation and contents of digital documents, allowing clients who trust the service to determine which document was "notarized" first. Correlating the histories of multiple, mutually distrustful secure time stamping services [23] is important, because not everyone doing business in the world can be convinced to trust the same centralized service; being able to map time stamps issued elsewhere to a local trust domain is essential for critical documents (such as contracts) from disparate jurisdictions. Luckily, sensitive documents such as contracts tend to change little or not at all, and high latencies for obtaining or verifying secure time stamps are acceptable, facilitating the development of an *efficient enough* P2P solution to the problem.

4 2 P2P or Not 2 P2P?

Figure 1 is a decision tree organizing our characteristics to determine whether the application of P2P techniques to a particular problem is justified. There are other metrics, other decision trees and other decision graphs than those presented here, but we have found this arrangement particularly useful. We examine our example problems and suggested solutions by traversing the tree in a breadth-first manner.

At the top of the tree we have the "budget" axis. We believe that limited budget is the most important motivator for a P2P solution. With limited budget, the low cost for a peer to join a P2P solution is very appealing. Otherwise, a centralized or centrally controlled distributed solution can provide lower complexity and higher performance for the extra money. Our tree thus continues only along the "limited" budget end of the axis.

Our next most important characteristic is the "relevance" of the resource in question. The more relevant (important to many) the resource, the more motivated peers in a P2P architecture are to participate. Good P2P solutions for problems with low relevance exist but have other mitigating characteristics, as we explain below.

The next axis in the tree is "mutual trust." Successful P2P solutions with trusting peers exist, as do those whose other characteristics justify the perfor-

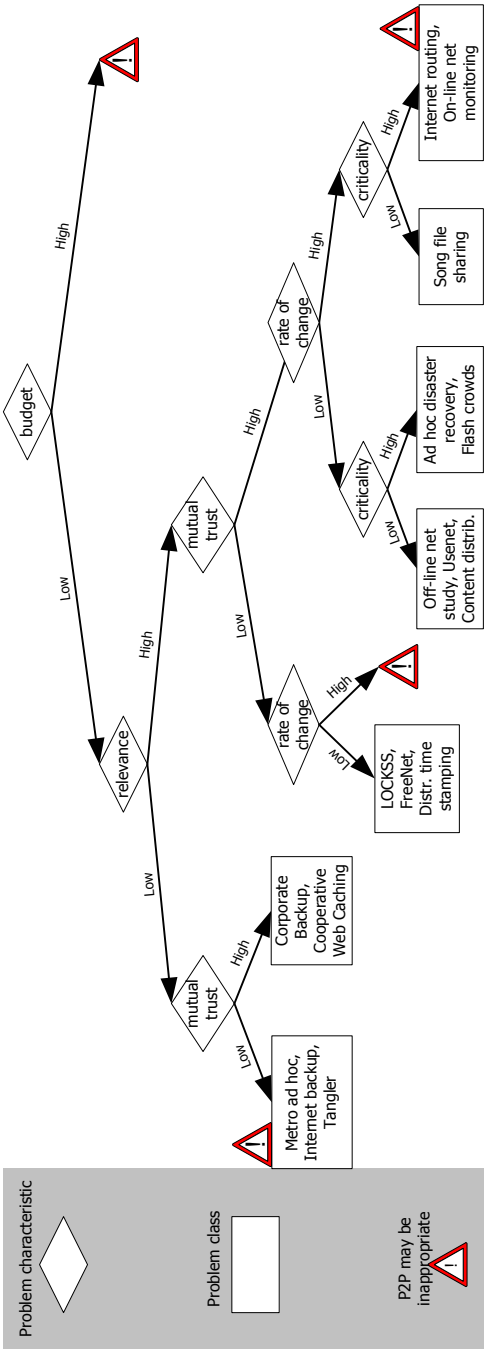

Fig. 1. A decision tree for analyzing the suitability of a P2P solution to a problem. Diamonds indicate decision points. Boxes contain problems or specific P2P solutions to problems. A warning sign over a particular box indicates that the box is a "trouble spot"; a P2P solution for the problems in that box may be inappropriate. In some cases, we include particular P2P solutions (e.g., Tangler) and explain in the text how those solutions overcome the difficulties of their box.

mance and complexity cost of measures to cope with mutual distrust. Those problems with low relevance and low trust have the burden of fostering cooperation. While Tangler is a good example, we believe that metropolitan ad hoc wireless networks and Internet backup have not yet succeeded. Motivation for these problems seems inadequate to overcome the low relevance of the resources and the overheads of protecting against uncooperative or malicious peers. Where peers are assumed to cooperate, problems such as corporate backup may succeed with P2P solutions, since corporate mandate compensates for low relevance. Similarly with cooperative web caching, proposed solutions [34] indicate that some benefits may be obtainable with P2P techniques; note, however, that actual benefits from cooperative web caching have thus far been only marginal compared to centralized solutions.

Where relevance is high, the required level of trust between peers still has an impact on the suitability of a P2P solution for the problem. Creating artificial economies or "trading" schemes to provide extrinsic incentives for cooperation (as in MojoNation) is generally unsuccessful [25]. The overhead in terms of complexity and performance for managing mutually distrustful peers suggests that solutions will be difficult to implement successfully in a P2P manner, unless other characteristics intercede to simplify the problem.

Such a characteristic is the rate of change in the system. Problems with a low rate of change, such as digital preservation, censorship resistant repositories, and distributed time stamping, may succeed despite mutually distrustful peers. For these problems, mutual distrust among peers is an inherent part of the problem, and thus its cost must be born by any proposed solution. The cost, however, is reduced by the low rate of change, which makes it possible to detect anomalies in the system in time to address them, and reduces the performance impact of the measures to protect against malicious peers. Problems with a high rate of change in untrustworthy environments are unlikely to find successful P2P solutions.

The rate of change in the system remains important even for problems in which peers may trust each other to cooperate. If the rate of peers entering and leaving the system is kept low, then both non-critical problems (such as off-line network studies, Usenet, and content distribution) and critical problems (such as ad hoc wireless network deployment for disaster recovery and flash crowd mitigation) may succeed. If the system moves quickly, we believe that it is easier to deploy non-critical applications such as file sharing that can tolerate inconsistent views among peers. When the problem involves critical information that also changes quickly (as in the case of Internet routing and on-line network monitoring), the designer should consider whether the application benefits sufficiently from other features. To the degree that Internet routing is successful, it is because it is amenable to trading accuracy for scalability through techniques such as aggregation of data. If P2P network monitoring succeeds, it will be because the natural distribution and high volume of the data allow few other architectures.

5 Conclusions

To summarize, the characteristics that motivate a P2P solution are limited budget, high relevance of the resource, high trust between nodes, a low rate of system change, and a low criticality of the solution. We believe that the limited budget requirement is the most important motivator. Relevance is also very important but can be compensated for by "saving graces" such as assumed trust between nodes or strong imposed incentives. Lacking these, we believe that problems with low relevance are not appropriate for P2P solutions. Trust between nodes greatly eases P2P deployment, however there are some applications, such as LOCKSS, FreeNet and distributed time stamping, where deployment across trust domains is a requirement. These applications must pay the overhead of distrust between nodes, but are feasible in a P2P context because a low rate of change makes these costs manageable.

While P2P solutions offer many advantages, they are inherently complex to get right and should not be applied blindly to all problems. In providing a framework in which to analyze the characteristics of a problem, we hope to offer designers some guidance on whether their problem warrants a P2P solution.

6 Acknowledgments

We would like to thank the following people for their very helpful feedback and suggestions: John Apostolopoulos, Sujata Banerjee, Kevin Lai, Dejan S. Milojičić, Mitch Trott, Susie Wee, and Zhichen Xu.

References

1. D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *SOSP*, 2001.
2. S. Bansal and M. Baker. Observation-based Cooperation Enforcement in Ad Hoc Networks. Technical report, Stanford University, 2003.
3. S. Burleigh, K. Fall, V. Cerf, R. Durst, K. Scott, H. Weiss, L. Torgerson, and A. Hooke. Delay-Tolernat Networking: An Approach to Interplanetary Internet. *IEEE Communications Magazine*, June 2003.
4. L. Buttyán and J.-P. Hubaux. Stimulating Cooperation in Self-organizing Mobile Ad hoc Networks. *Mobile Networks and Applications*, 2003.
5. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2), 1981.
6. B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. PlanetLab: An Overlay Testbed for Broad-Coverage Services. *ACM Computer Communication Review*, 33(3):3–12, July 2003.
7. I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *Workshop on Design Issues in Anonymity and Unobservability*, 2000.
8. B. Cohen. Incentives Build Robustness in BitTorrent. In *P2P Econ Workshop*, 2003.

9. L. P. Cox and B. D. Noble. Samsara: Honor Among Thieves in Peer-to-Peer Storage. In *SOSP*, 2003.
10. F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area Cooperative Storage with CFS. In *SOSP*, 2001.
11. M. Fomenkov, K. Keys, D. Moore, and kc claffy. Longitudinal study of Internet traffic from 1998–2003. `http://www.caida.org/outreach/papers/2003/nlanr/`.
12. K. Gummadi, R. Gummadi, S. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica. The Impact of DHT Routing Geometry on Resilience and Proximity. In *SIGCOMM*, 2003.
13. S. Haber and W. S. Stornetta. How to Time-stamp a Digital Document. *Journal of Cryptology: the Journal of the Intl. Association for Cryptologic Research*, 3(2):99–111, 1991.
14. G. Hardin. The Tragedy of the Commons. *Science*, 162, 1968.
15. HiveCache, Inc. Distributed disk-based backups. Available at `http://www.hivecache.com/`.
16. R. Huebsch, J. M. Hellerstein, N. Lanham, B. T. Loo, S. Shenker, and I. Stoica. Querying the Internet with PIER. In *VLDB*, 2003.
17. D. Kostić, A. Rodriguez, J. Albrecht, and A. Vahdat. Bullet: High Bandwidth Data Dissemination Using an Overlay Mesh. In *SOSP*, 2003.
18. J. Kubiatowicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. OceanStore: An Architecture for Global-Scale Persistent Storage. In *ASPLOS*, 2000.
19. C. Labovitz, A. Ahuja, A. Abose, and F. Jahanian. Delayed Internet Routing Convergence. In *SIGCOMM*, 2000.
20. J. Li, B. T. Loo, J. Hellerstein, F. Kaashoek, D. R. Karger, and R. Morris. On the Feasibility of Peer-to-Peer Web Indexing and Search. In *IPTPS*, 2003.
21. K. Lougheed and Y. Rekhter. RFC 1267: Border Gateway Protocol 3, October 1991.
22. D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: A Scalable and Dynamic Emulation of the Butterfly. In *CHI*, 1989.
23. P. Maniatis and M. Baker. Secure History Preservation Through Timeline Entanglement. In *USENIX Security*, 2002.
24. P. Maniatis, M. Roussopoulos, TJ Giuli, D. S. H. Rosenthal, M. Baker, and Y. Muliadi. Preserving Peer Replicas By Rate-Limited Sampled Voting. In *SOSP*, 2003.
25. J. McCoy. Lessons Learned from MojoNation. Personal Communication, April 2002.
26. D. S. Milojičić, V. Kalogeraki, R. Lukose, K. Nagaraja, J. Pruyne, B. Richard, S. Rollins, and Z. Xu. Peer-to-Peer Computing. Technical Report HPL-2002-57, HP Labs, 2002.
27. A. I. T. Rowstron, A.-M. Kermarrec, M. Castro, and P. Druschel. SCRIBE: The design of a large-scale event notification infrastructure. In *Networked Group Communication*, 2001.
28. T. Stading, P. Maniatis, and M. Baker. Peer-to-Peer Caching Schemes to Address Flash Crowds. In *IPTPS*, 2002.
29. A. Stavrou, D. Rubenstein, and S. Sahu. A Lightweight, Robust P2P System to Handle Flash Crowds. In *ICNP*, 2002.
30. I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *SIGCOMM*, 2002.
31. I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. In *SIGCOMM*, 2001.

32. M. Waldman and D. Mazières. Tangler: A Censorship-Resistant Publishing System Based On Document Entanglements. In *ACM Conf. on Computer and Communications Security*, 2001.
33. D. Wallach. A Survey of Peer-to-Peer Security Issues. In *Intl. Symposium on Software Security*, 2002.
34. A. Wolman, G. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. Levy. On the Scale and Performance of Cooperative Web Proxy Caching. In *SOSP*, 1999.