A Survey of Solutions to the Sybil Attack

Brian Neil Levine[1] Clay Shields[2] N. Boris Margolin[1]

[1] Dept. of Computer Science, Univ. of Massachusetts, Amherst
[2] Dept. of Computer Science, Georgetown University
{brian, margolin}@cs.umass.edu, clay@cs.georgetown.edu

1 Introduction

Many security mechanisms are based on specific assumptions of *identity* and are vulnerable to attacks when these assumptions are violated. For example, *impersonation* is the well-known consequence when authenticating credentials are stolen by a third party. Another attack on identity occurs when credentials for one identity are purposely shared by multiple individuals, for example to avoid paying twice for a service. Such shared accounts are common in practice: friends exchange iTunes passwords to share purchased music; BugMeNot.com is a community that shares website registration passwords; and network address translation [29] devices allow multiple users to pay for a single IP address which is then shared among them.

In this paper, we survey the impact of the Sybil attack [26], an attack against identity in which an individual entity masquerades as multiple simultaneous identities. The Sybil attack is a fundamental problem in many systems, and it has so far resisted a universally applicable solution.

Many distributed applications and everyday services assume each participating entity controls exactly one identity. When this assumption is unverifiable or unmet, the sevice is subject to attack and the results of the application are questionable if not incorrect. A concrete example of this would be an online voting system where one person can vote using many online identities. Notably, this problem is currently only solved if a central authority, such as the administrator of a certificate authority, can guarantee that each person has a single identity represented by one key; in practice, this is very difficult to ensure on a large scale and would require costly manual attention.

2 The Sybil Attack

While it has only been recently named and described, the Sybil attack has appeared in many forms in both academic work and in the real world. It is a severe and pervasive problem in many areas. For example, it is possible to rig Internet polls by using multiple IP addresses to submit votes [47], to gain advantage in any results of a chain letter [12], and is a well-known and potentially major problem in real-world elections [86]. A Sybil attack is also used by companies that increase the Google PageRank rating of the pages of their customers [9], and

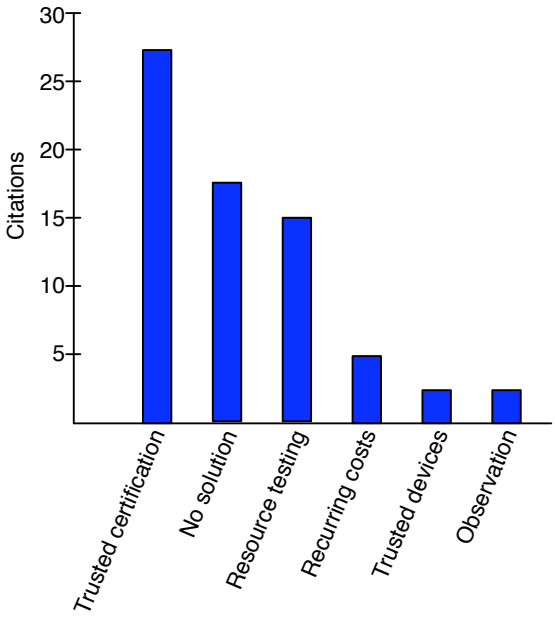

Fig. 1. Sybil attack approaches in the literature, summarized

has been used to link particular search terms to unexpected results for political commentary [66]. Reputation systems are a common target for Sybil attacks [18] including real-world systems like eBay [8].

Spammers can use this attack to gain access to multiple accounts on free-email systems. Peer-to-peer computing systems which use voting to verify correct answers, such as SETI@home, are also susceptible to accepting false solutions from a Sybil attacker [94]. Ad hoc mobile network routing can be manipulated when a Sybil attacker appears to be many different mobile nodes at once [44]. In systems that provide anonymity between peers, such as Tor, the Sybil attack is generally capable of revealing the initiator of a connection [89] and there is no defense against this attack [24]. It also allows free riding in services in cooperative file storage systems such as Pastiche [20].

Formal analyses of the attack have been done in the context of peer-to-peer applications [18,26]. Despite this work, there is no general solution to the attack. Proposed solutions most commonly use resource testing, though Douceur has shown this cannot prevent the attack in practical situations [26]. A wide variety of applications have considered the effects of the attack [22, 49, 67, 71, 83].

Below we first summarize approaches that have been cited in the literature to protecting against or detecting the attack. We then review results that are specific to different applications vulnerable to the attack.

2.1 General Approaches

Since the first analysis of the Sybil attack, some eleven different approaches have been proposed to prevent or mitigate the attack. In this paper, we categorize 90 papers that mention either the Sybil attack or pseudospoofing [23] (an earlier term for the use of multiple false identities) into these eleven categories and describe each approach.

Approximately half of the published papers either suggest certification as a solution to the Sybil attack, following Douceur's approach, or simply state the problem without giving a solution. The remaining papers use one of nine distinct strategies. In Figure 1 we show the number of citations for different approaches to the Sybil problem; some papers are counted in multiple categories.

• **Trusted certification** [2,15,16,26–28,35,40,42,46,51,61–63,65,67,68,75–78,84,85,90,95,96,98]

Douceur [26] has proven that trusted certification is the only approach that has the potential to completely eliminate Sybil attacks. Accordingly, it is cited as the most common solution. However, trusted certification relies on a centralized authority that must ensure each entity is assigned exactly one identity, as indicated by possession of a certificate. In fact, Douceur offers no method of ensuring such uniqueness, and in practice it must be performed by a manual or in-person process. This may be costly or a create a performance bottleneck in large-scale systems. Moreover, to be effective, the certifying authority must ensure that lost or stolen identities are discovered and revoked. If the performance and security implications can be solved, then this approach can eliminate the Sybil attack.

• **No solution** [1,6,7,21,25,32,43,45,53,54,58,79–82,88,97]

Though many researchers are aware that the Sybil attack is a potential problem, they present no solution to it for in their work. We cite these publications to point out that the Sybil attack remains an unsolved problem that is correctly acknowledged where applicable, and not to disparage the works.

• **Resource testing** [4,15,19,33,38,41,51,55,56,59,61,67,87,93,95]

The goal of resource testing is to attempt to determine if a number of identities possess fewer resources than would be expected if they were independent. These tests include checks for computing ability, storage ability, and network bandwidth, as well as limited IP addresses. Cornelli et al. [19] and Freedman and Morris [33] specifically propose testing for IP addresses in different domains or autonomous systems. Requiring heterogeneous IP addresses prevents some attacks but does not discourage others (such as zombie networks) and limits the usability of an application.

Douceur has proven the ineffectiveness resource tests, but a number of researchers suggest them as a minimal Sybil attack defense. In these cases the stated goal is to discourage rather than prevent Sybil attacks, and the number of identities an attacker can have is, in theory, limited. For many applications this is insufficient if an attacker can obtain enough identities for a successful attack, even if it is expensive. In the Tor communication system, for example, only two identities are required for an attack on anonymity [89].

In a type of resource test, Yu et al.'s SybilGuard technique [93] relies on limited availability of real-world friendship edges between nodes. However, the p2p application in use may have little intersection with the real-world friends represented in the graph. These friendship relationships may also be expensive to construct since the proposal requires out-of-band key sharing and a stronger trust relationship than is typical in social networks. It instead requires procedures at least as onerous as the GPG key signing tree. These costs, however, are one-time only and can be amortized over time by honest or malicious users alike.

• **Recurring costs and fees** [5, 27, 36, 55, 56]

In a variation on resource testing, in several papers [5, 55, 56] identities are periodically re-validated using resource tests. The approach limits the number of Sybil attacker with constrained resources can introduce in a period of time. However, as we noted above, in many applications very few Sybil identities are required for an effective attack. Additionally, in these papers, computational power is tested. Computational power mostly involves a one-time cost (for example, the purchase of computing hardware), so an attacker could recover over time even a high initial cost of claiming a large number of identities.

Awerbuch and Scheidler [5] suggest the use of Turing tests, for example CAPTCHAs, to impose recurring fees. Dragovic et al. [27] require certification of identities, but this certification is not trusted; rather, it is seen as a way of imposing identity creation costs. Gatti et al.'s *Sufficiently Secure Peer-to-Peer Networks* [36] uses an economic, game-theoretical approach to examine when attacks on censorship resistant networks are cost-effective. In recent work [57], we showed that charging a recurring fee for each participating identity is quantitatively more effective as a disincentive against successful Sybil attacks than charging one-time fees. For many applications, recurring fees can incur a cost to the Sybil attack that increases linearly with the total number of identities participating; one-time fees incur only a constant cost.

• **Trusted devices** [67, 76]

In a defense related to trusted certification authorities, entities in an application can be linked in some secure fashion to a specific hardware device. Analogous to any central authority handing out cryptographic certificates, there are no special methods of preventing an attacker from obtaining multiple devices other than manual intervention. The cost of acquiring multiple devices may be high, however.

2.2 Application Domains

In the remainder of this section, we summarize results regarding the Sybil attack that are specific to a broad set of application domains.

• **Mobile Networks** [14, 72] Wireless, mobile networks provide a unique avenue for detecting Sybil attackers. Observation of location can distinguish different devices, and limits of realistic mobility can constrain attacker movement. For an attacker with a single device, all Sybil identities will always appear to move together. The defense is not applicable beyond mobile networks, and it

does not protect against a single entity controlling multiple devices, each having a non-recurring cost.

• **Auditing** [3, 83, 94] In some cases, the correctness of identity behavior can be determined through audit. If the audit is cheap, the Sybil attack has little benefit: for instance, a large number of apparently independent identities cannot successfully convince another entity that they have factored a large number unless they have actually done so.

In some cases, audits are as costly as performing a requested computation. Here the probability of a successful misrepresentation must be factored into the cost of a Sybil attack. Yurkewych et al. [94] study the effect of the Sybil attack on p2p computing schemes (e.g., SETI@HOME). They determine the most effective strategy is to offer a large reward for for a correct calculation of a result with limited auditing. Because of the Sybil attack, it is less effective to redundantly give the same computation to several participants and reward a majority that returns with the same result with limited auditing.

• **Cash economies** [52, 91, 92] In these applications, identities explicitly exchange currency for desired goods or services. In most cases, such applications are not susceptible to the Sybil attack, since they do not rely on redundancy. Yokoo et al. [91, 92] describe a Sybil attack in combinatorial auction protocols, where the independence of different valuations of goods is attacked. The benefit of the Sybil attack can be eliminated by pricing goods appropriately in that it must not be more expensive to buy a bundle of goods than to buy each good separately.

• **Reputation Systems** For many p2p systems, including ad hoc networks and online markets, reputation systems have received a significant amount of attention as a solution for mitigating the affects of malicious peers. In an important work, Cheng and Friedman [17] evaluated the vulnerability of reputation systems to the Sybil attack, classifying them as *symmetric* or *asymmetric* approaches.

A symmetric reputation system is one in which an identity's reputation depends solely on the topology of the trust graph, and not the naming or identity of nodes. An attacker that wishes to increase its reputation simply uses Sybil identities to create a copy of the existing graph representing trust relationships. A symmetric reputation system cannot distinguish original nodes from the copies, and thus some Sybil node has reputations equal or better to any original node. Cheng and Friedman [17] prove formally that such reputation systems are susceptible to Sybil attacks. Examples of symmetric reputation systems include Google's PageRank algorithm [9, 50], EigenTrust [48], and others [27, 61, 69, 70]

In asymmetric reputation systems, there are specifically trusted nodes from which all reputation values propagate. Alternatively, each entity separately computes a trust value along their unique paths to every other identity in the system. Since the trusted nodes cannot be impersonated, no Sybil attacker can create a duplicate graph as explained above in the symmetric case. This trust value can change over time as the entity interacts with and observes the behavior of different identities. This is typical of social networks.

Asymmetric reputation systems can be effective at raising the cost of Sybil attacks because attackers are forced to build up trust before effectively launching attacks. Unfortunately, these systems inevitably penalize newcomers who must prove themselves by offering benefits before getting anything in return. Examples of asymmetric systems include Feldman et al. [30], Guha et al. [39], Domingos et al. [74] among others [3, 10, 11, 13, 14, 17, 27, 31, 34, 37, 45, 52, 60, 64, 68, 73]

3 Conclusion

There are a variety of attacks that hinge on the issue of identity. In this paper, we have presented an overview of work related to analyzing or solving the Sybil attack, in which one entity appears as or controls many different identities. We have demonstrated the breadth of applications that are subject to the attack, including the widely used systems Google, eBay, SETI@HOME, and Tor. The attack also presents a problem for peer-to-peer networks, mobile networks, and reputation systems. While we lack an efficient, general solution that scales well to large systems, there are a variety of solutions that can limit or prevent the attack in several individual application domains.

References

1. A. Acquisti, R. Dingledine, and P. Syverson. On the Economics of Anonymity. In *Proc. Financial Cryptography (FC)*. Springer-Verlag, LNCS 2742, January 2003.
2. A. Adya, W. J. Bolosky, M. Castro, R. Chaiken, G. Cermak, J. R. Douceur, J. Howell, J. Lorch, M. Theimer, and R. P. Wattenhofer. Farsite: Federated, available, and reliable storage for an incompletely trusted environment. In *Proc. OSDI*, pages 1–14, Dec. 2002.
3. K. Anagnostakis and M. Greenwald. Exchange-based incentive mechanisms for peer-to-peer file sharing. In *Proc. Intl Conf on Distributed Computing Systems (ICDCS)*, Mar. 2004.
4. J. Aspnes, C. Jackson, and A. Krishnamurthy. Exposing computationally-challenged Byzantine impostors. Technical Report YALEU/DCS/TR-1332, Yale University Department of Computer Science, July 2005.
5. B. Awerbuch and C. Scheideler. Group Spreading: A Protocol for Provably Secure Distributed Name Service. In *Proc. Automata, Languages and Programming (ICALP)*, pages 183–195, 2004.
6. B. Awerbuch and C. Scheideler. Robust Distributed Name Service. In *Proc. Intl Wkshp on Peer-to-Peer Systems*, pages 237–249, 2004.
7. A. Bhargava, K. Kothapalli, C. Riley, C. Scheideler, and M. Thober. Pagoda: a dynamic overlay network for routing, data management, and multicasting. In *Proc. ACM Symp on Parallel Algorithms*, pages 170–179, 2004.
8. R. Bhattacharjee and A. Goel. Avoiding Ballot Stuffing in eBay-like Reputation Systems. In *ACM SIGCOMM Workshop on the Economics of Peer-to-Peer Systems*, August 2005.
9. M. Bianchini, M. Gori, and F. Scarselli. Inside pagerank. *ACM Trans. Inter. Tech.*, 5(1):92–128, 2005.

10. R. Böhme, G. Danezis, C. Díaz, S. Köpsell, and A. Pfitzmann. Mix cascades vs. peer-to-peer: Is one concept superior? In *Proc. Wkshp on Privacy Enhancing Technologies*, pages 243–255, 2004.
11. S. Buchegger and J.-Y. L. Boudec. A Robust Reputation System for P2P and Mobile Ad hoc Networks. In *Proc. Wkshp on the Economics of Peer-to-Peer Systems*, 2004.
12. J. Bulgatz. *More Extraordinary Popular Delusions and the Madness of Crowds*. Three Rivers Press, 1992.
13. S. Čapkun and J.-P. Hubaux. BISS: building secure routing out of an incomplete set of secure associations. In *Proc. ACM Wireless Security Conference*, pages 21–29, 2003.
14. S. Capkun, J.-P. Hubaux, and L. Buttyan. Mobility helps peer-to-peer security. *IEEE Transactions on Mobile Computing*, 5(1):43–51, Jan 2006.
15. M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proc. OSDI*, pages 299–314, Dec 2002.
16. M. Castro, P. Druschel, A. Kermarrec, and A. Rowstron. One ring to rule them all: Service discovery and binding in structured peer-to-peer overlay networks. In *Proc. SIGOPS European Wkshp*, 2002.
17. A. Cheng and E. Friedman. Sybilproof Reputation Mechanisms. In *Proc. ACM Wkshp on Economics of Peer-to-Peer Systems*, pages 128–132, 2005.
18. A. Cheng and E. Friedman. Sybilproof Reputation Mechanisms . In *ACM Wkshp on the Economics of Peer-to-Peer Systems*, August 2005.
19. F. Cornelli, E. Damiani, and S. Samarati. Implementing a reputation-aware gnutella servent. In *Proc. Intl Wkshp on Peer-to-Peer Computing*, 2002.
20. L. Cox and B. Noble. Pastiche: Making backup cheap and easy. In *Proc. USENIX OSDI*, Dec. 2002.
21. D. Cvrcek and V. Matyas. On the role of contextual information for privacy attacks and classification. In *Proc. Wkshp on Privacy and Security Aspects of Data Mining*, pages 31–39, Nov. 2004.
22. G. Danezis, C. Lesniewski-Laas, M. F. Kaashoek, and R. Anderson. Sybil-resistant DHT routing. In *Proc. ESORICS*, pages 305–318, 2005.
23. L. Detweiler. Snakes of medusa and cyberspace: Internet identity subversion. The Risks Digest: Forum on Risks to the Public in Computers and Related Systems, ACM Committee on Computers and Public Policy, Peter G. Neumann, moderator, Nov. 1993. `http://catless.ncl.ac.uk/Risks/15.25.html`.
24. R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proc. USENIX Security Symp*, Aug. 2004.
25. R. Dingledine, V. Shmatikov, and P. Syverson. Synchronous batching: From cascades to free routes. In *Proc. Privacy Enhancing Technologies workshop (PET)*, volume 3424 of *LNCS*, May 2004.
26. J. Douceur. The Sybil Attack. In *Proc. Intl Wkshp on Peer-to-Peer Systems (IPTPS)*, Mar. 2002.
27. B. Dragovic, E. Kotsovinos, S. Hand, and P. R. Pietzuch. Xenotrust: Event-based distributed trust management. In *Proc. Intl Wkshp on Database and Expert Systems Applications*, 2003.
28. B. Dutertre, S. Cheung, and J. Levy. Lightweight key management in wireless sensor networks by leveraging initial trust. Technical Report SRI-SDL-04-02, SRI International, 2002.
29. K. Egevang and P. Francis. The IP network address translator (NAT). RFC 1631, Internet Engineering Task Force, May 1994.

30. M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust Incentive Techniques for Peer-to-Peer Networks. In *Proc. ACM E-Commerce Conference (EC)*, May 2004.
31. M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *ACM Conf on Electronic Commerce*, 2004.
32. D. R. Figueiredo, P. Nain, and D. Towsley. On the analysis of the predecessor attack on anonymity systems. Computer Science Technical Report 04-65, University of Massachusetts, July 2004.
33. M. J. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proc. ACM CCS*, Nov. 2002.
34. Y. Fu, J. Chase, B. Chun, S. Schwab, and A. Vahdat. SHARP: An architecture for secure resource peering. In *Proc. ACM SOSP*, Oct. 2003.
35. A. C. Fuqua, T.-W. J. Ngan, and D. S. Wallach. Economic behavior of peer-to-peer storage networks. In *Wkshp on Economics of Peer-to-Peer Systems*, 2003.
36. R. Gatti, S. Lewis, A. Ozment, T. Rayna, , and A. Serjantov. Sufficiently secure peer-to-peer networks. In *Wkshp on the Economics of Information Security*, May 2004.
37. P. Gauthier, B. Bershad, and S. Gribble. Dealing with cheaters in anonymous peer-to-peer networks. Technical Report 04-01-03, University of Washington, January 2004.
38. S. Goel, M. Robson, M. Polte, and E. G. Sirer. Herbivore: A Scalable and Efficient Protocol for Anonymous Communication. Technical Report 2003-1890, Cornell University, February 2003.
39. R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proc. Intl Conf on World Wide Web*, pages 403–412. ACM Press, 2004.
40. N. Harvey, M. B. Jones, S. Saroiu, M. Theimer, and A. Wolman. Skipnet: A scalable overlay network with practical locality properties. In *Proc. USENIX Symp on Internet Technologies and Systems (USITS)*, March 2003.
41. K. Hildrum. *Finding Nearby Objects in Peer-to-Peer Networks*. PhD thesis, University of California, Berkeley, 2004.
42. K. Hildrum and J. Kubiatowicz. Asymptotically efficient approaches to fault-tolerance in peer-to-peer networks. In *Proc. Intl Symp on Distributed Computing*, pages 321–336, 2003.
43. K. Hildrum, J. D. Kubiatowicz, S. Rao, and B. Y. Zhao. Distributed object location in a dynamic network. In *Proc. ACM Symp on Parallel Algorithms and Architectures*, pages 41–52, Aug. 2002.
44. Y.-C. Hu, A. Perrig, and D. Johnson. Ariadne: A secure on-demand routing protocol for ad hoc networks. In *Proc. ACM MOBICOM*, Sept. 2002.
45. R. Huebsch, B. N. Chun, J. M. Hellerstein, B. T. Loo, P. Maniatis, T. Roscoe, S. Shenker, I. Stoica, and A. R. Yumerefendi. The architecture of pier: an internet-scale query processor. In *Proc. CIDR*, pages 28–43, 2005.
46. M. Jelasity, A. Montresor, and O. Babaoglu. Towards secure epidemics: Detection and removal of malicious peers in epidemic-style protocols. Technical Report UBLCS-2003-14, University of Bologna, Department of Computer Science, Nov. 2003. presented at FuDiCo II: S.O.S, Bertinoro, Italy, June, 2004.
47. P. Judge. ZDNet: .net vote rigging illustrates importance of web services. http://news.zdnet.co.uk/software/ 0,39020381,2102244,00.htm, 2002.
48. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proc. Intl Conf on World Wide Web*, pages 640–651. ACM Press, 2003.

49. C. Karlof and D. Wagner. Secure routing in wireless sensor networks: Attacks and countermeasures. *Ad hoc Networks Journal (Elsevier)*, 1(2–3):293–315, Sept. 2003.
50. P. Lawrence, B. Sergey, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford University, 1998.
51. J. Ledlie and M. Seltzer. Distributed, secure load balancing with skew, heterogeneity, and churn. In *Proc. IEEE INFOCOM*, Mar. 2005.
52. R. L. Levien. Attack resistant trust metrics. Master's thesis, University of California at Berkely, 1995.
53. B. N. Levine, M. K. Reiter, C. Wang, and M. K. Wright. Timing attacks in low-latency mix-based systems. In *Proc. Financial Cryptography*. Springer-Verlag, LNCS 3110, February 2004.
54. P. Maniatis, T. Giuli, M. Roussopoulos, D. Rosenthal, and M. Baker. Impeding attrition attacks on p2p systems. In *Proc. 11th ACM SIGOPS European Wkshp*, Sept. 2004.
55. P. Maniatis, D. S. H. Rosenthal, M. Roussopoulos, M. Baker, T. Giuli, and Y. Muliadi. Preserving peer replicas by rate-limited sampled voting. In *Proc. ACM SOSP*, pages 44–59, 2003.
56. P. Maniatis, M. Roussopoulos, T. J. Giuli, D. S. H. Rosenthal, and M. Baker. The lockss peer-to-peer digital preservation system. *ACM Trans. Comput. Syst.*, 23(1):2–50, 2005.
57. N. B. Margolin and B. N. Levine. Quantifying and discouraging sybil attacks. Computer Science Technical Report 2005-67, University of Massachusetts Amherst, Dec. 2005.
58. S. Marti, P. Ganesan, and H. Garcia-Molina. SPROUT: P2P Routing with Social Networks. In *Proc. EDBT Wkshps*, pages 425–435, 2004.
59. S. Marti and H. Garcia-Molina. Examining metrics for peer-to-peer reputation systems. Technical Report 2003-39, Stanford University, July 2003.
60. S. Marti and H. Garcia-Molina. Identity crisis: Anonymity vs. reputation in p2p systems. In *Proc. 3rd Intl Conference on Peer-to-Peer Computing*, pages 134–141, 2003.
61. S. Marti and H. Garcia-Molina. Limited reputation sharing in p2p systems. In *Proc. 5th ACM conference on Electronic commerce*, 2004.
62. J. Martin and L. Alvisi. A framework for dynamic byzantine storage. Technical Report TR04-08, University of Texas at Austin, 2004.
63. G. Mathur, V. N. Padmanabhan, and D. R. Simon. Securing routing in open networks using secure traceroute. Technical Report MSR-TR-2004-66, Microsoft Research, July 2004.
64. R. Morselli, J. Katz, and B. Bhattacharjee. A game-theoretic framework for analyzing trust-inference protocols. In *2nd Wkshp on Economics of Peer-to-Peer Systems, Cambridge, MA, USA*, June 2004.
65. M. Narasimha, G. Tsudik, and J. H. Yi. On the utility of distributed cryptography in p2p and manets: the case of membership control. In *Proc. IEEE Intl Conference on Network Protocols (ICNP)*, 2003.
66. B. News. 'miserable failure' links to bush, 'December 7' 2003. http://news.bbc.co.uk/2/hi/americas/3298443.stm.
67. J. Newsome, E. Shi, D. Song, and A. Perrig. The Sybil attack in sensor networks: analysis & defenses. In *Proc. Intl Symp on Information Processing in Sensor Networks (IPSN)*, pages 259–268, 2004.
68. T.-W. J. Ngan. Incentives and fair sharing in peer-to-peer systems. Master's thesis, Rice University, 2004.

69. T.-W. J. Ngan, D. S. Wallach, and P. Druschel. Incentives-compatible peer-to-peer multicast. In *2nd Wkshp on the Economics of Peer-to-Peer Systems*, June 2004.
70. N. Ntarmos and P. Triantafillou. Seal: Managing accesses and data in peer-to-peer sharing networks. In *Peer-to-Peer Computing*, pages 116–123, 2004.
71. A. Perrig, J. Stankovic, and D. Wagner. Security in wireless sensor networks. *Commun. ACM*, 47(6):53–57, 2004.
72. C. Piro, C. Shields, and B. N. Levine. Detecting the Sybil Attack in Ad hoc Networks. In *Proc. IEEE/ACM Intl Conf on Security and Privacy in Communication Networks (SecureComm)*, August 2006.
73. B. C. Popescu, B. Crispo, and A. S. Tanenbaum. Safe and Private Data Sharing with Turtle: Friends Team-Up and Beat the System. In *Proc. 12th Cambridge Intl Wkshp on Security Protocols*. Springer-Verlag, April 2004.
74. M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *Intl Semantic Web Conference*, pages 351–368, 2003.
75. R. Rodrigues. An agenda for robust peer-to-peer storage. In *First IRIS Student Wkshp*, Aug. 2003.
76. R. Rodrigues, B. Liskov, and L. Shrira. The design of a robust peer-to-peer system. In *Proc. ACM SIGOPS European Wkshp*, Sept. 2002.
77. T. Roscoe and S. Hand. Transaction-based charging in Mnemosyne: a peer-to-pee steganographic storage system. In *Proc. Intl Wkshp on Peer-to-Peer Computing at Networking 2002*, May 2002.
78. N. Saxena, G. Tsudik, and J. H. Yi. Admission control in peer-to-peer: design and performance evaluation. In *Proc. ACM Wkshp on Security of ad hoc and sensor networks*, pages 104–113. ACM Press, 2003.
79. A. Serjantov. Anonymizing censorship resistant systems. In *Proc. 1st Intl Wkshp on Peer-to-Peer Systems*, Mar. 2002.
80. A. Serjantov and R. Anderson. On dealing with adversaries fairly. In *Wkshp on Economics and Information Security*, May 2004.
81. J. Shneidman and D. C. Parkes. Overcoming rational manipulation in mechanism implementations. Technical Report TR-12-03, Harvard, 2003.
82. J. Shneidman, D. C. Parkes, and L. Massoulie. Faithfulness in internet algorithms. In *Proc. Wkshp on Practice and Theory of Incentives and Game Theory in Networked Systems (PINS)*, 2004.
83. M. Srivatsa and L. Liu. Vulnerabilities and security threats in structured overlay networks: A quantitative analysis. In *Proc. ACSAC*, pages 252–261, 2004.
84. M. Srivatsa, L. Xiong, and L. Liu. XChange: A distributed protocol for electronic fair-exchange. In *Proc. IEEE Intl Parallel and Distributed Processing Symp*, 2005.
85. M. Čagalj, I. A. Saurabh Ganeriwal, and J.-P. Hubaux. On Cheating in CSMA/CA Ad Hoc Networks. Technical Report IC/2004/27, EPFL-DI-ICA, Mar. 2004.
86. A. Viglucci, J. Tanfani, and L. Getter. Herald special report: Dubious tactics tilted mayoral votes. Miami Herald, February 8, 1998.
87. V. Vishnumurthy, S. Chandrakumar, and E. G. Sirer. KARMA: A Secure Economic Framework for P2P Resource Sharing. In *Proc. Wkshp on the Economics of Peer-to-Peer Systems*, 2003.
88. H. J. Wang, Y.-C. Hu, C. Yuan, Z. Zhang, and Y.-M. Wang. Friends troubleshooting network: Towards privacy-preserving, automatic troubleshooting. In *IPTPS*, pages 184–194, 2004.
89. M. Wright, M. Adler, B. N. Levine, and C. Shields. The Predecessor Attack: An Analysis of a Threat to Anonymous Communications Systems. *ACM Transactions on Information and System Security (TISSEC)*, 4(7), 2004.

90. H. Yang, H. Luo, Y. Yang, S. Lu, and L. Zhang. HOURS: Achieving DoS Resilience in an Open Service Hierarchy. Technical report, UCLA, 2003.
91. M. Yokoo. Characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proc. Intl Joint Conference on Artificial Intelligence (IJCAI)*, Aug. 2003.
92. M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: new fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, January 2004.
93. H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman. Sybilguard: Defending against sybil attacks via social networks. In *Proc. ACM SIGCOMM*, 2006.
94. M. Yurkewych, B. N. Levine, and A. L. Rosenberg. On the cost-ineffectiveness of redundancy in commercial p2p computing. In *Proc. ACM CCS*, pages 280–288, 2005.
95. B. Zhao, A. D. Joseph, and J. Kubiatowicz. Supporting rapid mobility via locality in an overlay network. Technical Report UCB/CSD-02-1216, UC Berkeley, Nov. 2002.
96. B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. D. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, Jan. 2004.
97. Y. Zhao. *Decentralized Object Location and Routing: A New Network Paradigm*. PhD thesis, University of California, Berkeley, 2004.
98. S. Zhu, S. Setia, and S. Jajodia. LEAP: Efficient security mechanisms for large-scale distributed sensor networks. In *ACM Conference on Computer and Communications Security (CCS)*, Oct. 2003.