An Offline Foundation for Online Accountable Pseudonyms

Bryan Ford Jacob Strauss
Massachusetts Institute of Technology

ABSTRACT

Online anonymity often appears to undermine accountability, offering little incentive for civil behavior, but accountability failures usually result not from *anonymity* itself but from the *disposability* of virtual identities. A user banned for misbehavior—e.g., spamming from a free E-mail account or stuffing an online ballot box—can simply open other accounts or connect from other IP addresses. Instead of curtailing freedom of expression by giving up anonymity, online services and communities should support *accountable pseudonyms*: virtual personas that can provide both anonymity and accountability. We propose *Pseudonym parties*, a scheme for creating accountable pseudonyms, which combine in-person social occasions (parties) with technical infrastructure (a pseudonymous sign-on service) to enforce the rule that *one real person* gets *one virtual persona* on any participating online service. Pseudonym parties enable the user to adopt different personas in different online spaces without revealing the connection between them, while ensuring that each user has only one accountable pseudonym in *each* space. Pseudonym parties can be started incrementally in a fully decentralized fashion, can run on volunteer labor with minimal funds, and may even be fun.

1. INTRODUCTION

The right to *anonymity*, often seen as a necessary component of free expression, has long seemed at odds with the principle of *accountability*, an equally basic foundation of social justice and the rule of law [37]. The ability to participate anonymously in online communities is a widely cherished feature of the Internet [30, 34], particularly in that it enables people and groups with controversial or unpopular views to communicate and interact without fear of personal reprisal [27]. Opponents on the other hand contend that this anonymity often harbors and encourages antisocial or criminal behavior [7].

Indeed, many of the Internet's current maladies reduce to failures of accountability: given the ability to create online identities at will, there is little incentive for the user controlling any given identity to behave. Because open-access messaging systems cannot reliably identify a message's source for the purpose of suppressing abuse, spam has already relegated USENET to historical obscurity [31], threatens the usability of E-mail [36], and is advancing on voice-over-IP [6]. The automated "Turing tests" many web sites now employ to prevent automated abuses [33] also lock out visually impaired users [5, 22] and are vulnerable to attack using artificial intelligence [4] or social engineering [9]. Wikipedia progressively tightens its editing rules to combat the rising tide of anonymous vandalism [13, 18, 32]. Online voting and peer review systems like Slashdot operate reliably only to the extent that nobody cares about the results enough to bother opening multiple accounts and stuffing the ballot boxes [15]. Banning detected abusers by IP address frequently prevents access by other legitimate users on the same ISP [17], and many attacks come from compromised zombie machines not under the control of their owners [11].

This tension between anonymity and accountability may not be fundamental, but merely an indication that our current *mechanisms* to provide them are too primitive. Consider a masquerade ball in which everyone dresses unrecognizably in costume, and no latecomers are admitted once the event has started. If some attendee, Bob, breaks the rules of the event, the ball's organizers have at least two avenues of punishment. First, they could strip off Bob's mask in front of everyone as in the film *Eyes Wide Shut*, destroying his anonymity and potentially exposing him to reprisals in the real world. Alternatively, they could merely eject him from the ball, holding him accountable for his actions and preventing him from further disrupting the event while respecting his anonymity. The rule against latecomers is a crucial mechanism enabling the latter, anonymity-preserving form of punishment: without it, Bob could merely change costumes and re-enter after ejection.

SocialNets'08, April 1, 2008, Glasgow, Scotland, UK

The Internet's vulnerability to spam, ballot stuffing, and many similar attacks results not directly from the *anonymity* of users, but rather from the *disposability* of online identities. As if changing costumes, an attacker can create a new online identity and evade accountability simply by signing up for a new web account, connecting from another IP address, or sending spam from another compromised host. If an online community could reliably enforce the rule that *one real person* may obtain only *one virtual persona* over some significant period of time, then these online personas could still be anonymous but would no longer be disposable, providing a degree of user accountability. Online communities could revoke the access rights of abusers, for example, such as E-mail spammers or Wikipedia vandals, without affecting innocent users or permitting an abuser to reappear immediately under a different name. Voting systems for peer review or online deliberation could protect voter anonymity while preventing ballot box stuffing.

We will refer to online identities that combine anonymity with accountability in this way as *accountable pseudonyms*. We might create accountable pseudonyms in many ways: here we explore one mechanism, *pseudonym parties*, which takes advantage of the fact that real humans can be in only one place at a time. On a specific day every year, participating organizations or ad hoc groups of people host parties in their local areas, at which they pass out certificates to anyone who shows up in person. The physical presence requirement, combined with suitable procedures, ensures that each user may obtain only one such certificate per year. Given this certificate, a user can create any number of pseudonymous identities at a variety of participating online services—but only one such identity per service.

Accountable pseudonyms need not be deployed pervasively before they benefit users. Online services might still permit access by unauthenticated users, but offer privileges to holders of accountable pseudonyms, such as the right to participate in votes protected from ballot stuffing, and automatic exemption from IP blacklists, waiting periods, and other invasive protections against anonymous abuse. Though pseudonym parties require some "real-world" infrastructure, the costs of this infrastructure should be small enough initially to be borne by voluntary donations of time and resources, and should scale in proportion to the number of participants.

The rest of this paper is organized as follows. Section 2 outlines previous proposals for user accountability in more detail. Section 3 then presents and discusses pseudonym parties. Section 4 briefly outlines deployment issues, focusing more on the scheme's social aspects than on its technical details in order to promote discussion. Finally, Section 5 concludes.

2. BACKGROUND AND RELATED WORK

The abuse of an online system by creating many virtual personas has become known in the peer-to-peer community as a *sybil attack* [10]. Although peer-to-peer systems appear especially vulnerable to sybil attacks due to their decentralized nature, variants of sybil attacks currently plague many online services such as E-mail, web-based discussion forums, and other online public spaces.

Existing proposals addressing the sybil attack generally fall under four categories: associating users with network endpoints, authenticating users' real-world identities, limiting the rate or extent of attacks, and removing incentives to engage in sybil attacks. A related work from the economics field is a discussion of *unreplaceble pseudonyms* [12].

"Authenticating" Users by IP Address:.

Many online services attempt to protect themselves from abuse by associating users with the IP addresses from which they connect. Despite such protections, MIT students successfully rigged a Slashdot poll via ballot stuffing in November 1999 [33], and later a Doonesbury poll in 2006 [15]. Such feats are facilitated by MIT's inheritance of a vast, still mostly unused block of 2^{24} IP addresses from the early days of the Internet. Legitimate users behind large NATs or web proxies [3], on the other hand, may be unable to cast even one vote if another user already voted from the same shared IP address. In effect, your voting power depends on how early your organization joined the Internet.

E-mail spam is a form of sybil attack that has largely defeated attempts to control abuse through IP address blacklists [24]. As a result of modern botnets that send spam from a constantly-changing set of compromised hosts, the volume of spam continues to rise [8, 26] even as legitimate E-mail disappears into increasingly sensitive spam filters [19, 25].

Authenticating User Identities:.

Many proponents of accountability see anonymity as the root of the problem, and propose disclosure of the user's true identity as the solution [7]. Public certificate authorities such as Verisign offer personal, authenticated "digital IDs" for use in secure E-mail, but few users are even aware of these services, let alone willing to bother buying and using one. A few web sites such as PayPal ask the user to enter a credit card or bank account number for identification purposes, even if the user is not (immediately) making a purchase. Most web site operators however are reluctant to impose any unnecessary barriers to attracting new users. PGP key signing parties [2] authenticate user identities in a decentralized manner, but these identities cannot be pseudonymous, and are typically used only for signing E-mail.

Single sign-on initiatives such as Windows Live ID [35], the Liberty Alliance [20], and OpenID [23] address the inconvenience to the user of entering personal information everywhere by centralizing the user's information at a single "identity provider," which various online services contact to authenticate the user. In addition to the practical and security challenges to widespread deployment, however, these schemes create new privacy concerns [14].

Limiting Attack Rate or Extent:.

Another defense against sybil attacks is to increase the cost of creating new identities. On-line "Turing tests" such as CAPTCHAs [33] can prevent fully-automated attacks when they are effective [4], but cannot protect against determined (or paid) users who simply solve puzzles repeatedly [28]. Computational puzzles [1] similarly slow down attacks only by some factor, and are easily countered by an abuser with a large botnet. Out-of-band approaches such as sending an invitation to a cell phone limit accessibility to users who have cell phones, don't exclude attackers who could simply purchase multiple cheap phone accounts, and break the anonymity goal. Heuristics based on social network graph properties can similarly limit the power of large clusters of sybil identities [39]. These rate- and extent-limiting defenses may counter large-scale automated abuse, but do not prevent widespread small-scale attacks on systems in which *everyone* has an incentive to cheat "just a bit," such as with online ballot stuffing or sock puppetry [29].

Removing Attack Incentives:.

It may be possible to design certain applications so that users have no *incentive* to engage in sybil attacks by creating multiple identities. This idea has been studied formally for combinatorial auctions, yielding some positive results together with negative results suggesting that there is no general substitute for accountability [12, 38].

3. PSEUDONYM PARTIES

Pseudonym parties are a scheme for creating accountable pseudonyms that preserve the user's ability to be anonymous while keeping him accountable for his actions. Pseudonym parties ensure that a given real-life person can only operate under one accountable pseudonym at a time within the context of a given online service. An "online service" for our purposes could be a web-based community like Wikipedia or Slashdot, a traditional application such as digitally signed E-mail, or a fully decentralized peer-to-peer system. We first introduce pseudonym parties in the context of a geographically localized community, then explore how it can be decentralized and scaled over larger geographical regions.

3.1 One Body, One Pseudonym

We can imagine many ways of enforcing a *one person, one persona* rule by leveraging various secondary "labels" associated with people: e.g., require the user to provide a unique and verifiable credit card number, cell phone number, home address, social security number, etc. Privacy issues aside, this approach suffers from the fact that people often have more (or fewer) than one such label. The number of such labels one can acquire is usually limited only by effort, financial resources, and—in the case of labels that are legally required to be one-to-one—risk of getting caught. People can acquire several credit cards, several phone numbers, several home addresses, several government identity cards under different names, several national citizenships with a separate passport for each. Using secondary labels is also unfair to the disadvantaged: not everyone has *any* credit card, phone, home address, or national citizenship [21]. Barring certain sci-fi scenarios, however, everyone still has one and only one body. *Pseudonym parties* leverage the "offline foundation" of a user's physical presence at an event to guarantee a one-to-one relationship with online pseudonyms. On a particular day every year—let's call it *Pseudonym Day*—people who desire accountable pseudonyms gather locally and throw a party. Everyone who shows up at the party receives a *pseudonym certificate* and a hand stamp that takes a few days to wear off, ensuring that they can obtain only one such certificate until next year's party.

Each pseudonym certificate confers upon its holder the right to create one and only one accountable pseudonym on each of any number of online services that support such pseudonyms. A user might for example use his certificate to create one accountable pseudonym on Wikipedia, a second one on Slashdot, and a third on a peer-to-peer storage cloud. The user cannot create two separate Wikipedia accounts with the same certificate, however, or two identities on the same P2P storage cloud.

One reason why online communities can grow so quickly is that signing up for new services is quick and easy–a few clicks online is a much lower barrier to entry than any physical transaction. Since a single certificate can be used on multiple services, including ones which did yet exist when the certificate was issued, pseudonym parties do not inhibit this ease of growth, so long as users already have a certificate to re-use from some other service.

3.2 Anonymous Single Sign-On

A pseudonym certificate itself is simply a paper with a login name and password usable on a designated *pseudonym server* run by the pseudonym party's organizers. Suppose the user wishes to create an accountable pseudonym on a web-based online service such as Wikipedia. Wikipedia's web server temporarily redirects the user to his pseudonym server, where he logs in using his pseudonym certificate. The pseudonym server then returns the user to Wikipedia, where he finds himself in his new pseudonymous account. When the user later logs into Wikipedia again with the same certificate, the pseudonym server sends him back to the same account. The user's pseudonym server might in similar fashion provide the user with sybil-proof identities for peer-to-peer systems the user may join.

The pseudonym server acts like an "identity provider" in a single sign-on service [20, 23, 35], except that it does not *identify* the user but merely enforces the *one person, one persona* rule. The pseudonym server *cannot* directly reveal the user's identity even if compromised, since the user never provided any identification or personal information when obtaining his certificate.

The pseudonym server hides not only the true identity of the user, but also the *association* between the user's various pseudonyms for different online services, from both the users and operators of those services. If Bob uses his certificate to create a professional profile on `LinkedIn.com` and a steamy personal profile on `AdultFriendFinder`, for example, no one can tell that the two profiles represent the same person even if the two web sites collude or are hacked—unless, of course, Bob gives away the connection.

Online services could still allow traditional unauthenticated access, but offer special privileges to users of accountable pseudonyms, permitting incremental transition toward stronger accountability. A web-based forum like Slashdot may subject unauthenticated users to CAPTCHA puzzles [33] to deter automated attacks, impose initial waiting periods [13] and posting rate limits to discourage uncivil behavior, and disallow voting by unauthenticated users to prevent ballot stuffing. Users with accountable pseudonyms would be exempt from these restrictions, since misbehavior using an accountable pseudonym can be halted for the year merely by disabling that pseudonym. E-mail from users of accountable pseudonyms could be exempt from heuristic spam filters, avoiding loss due to false positives [19, 25]. Peer-to-peer protocols might prioritize information obtained from neighbors with accountable pseudonyms over information from unauthenticated neighbors, since only the former can be trusted to represent a real person and not a sybil identity.

3.3 Security and Trust Model

Although a user need not trust her pseudonym party's organizers or servers not to divulge her personal information directly, she must trust them to protect the *relationship* between the different accountable pseudonyms she obtains from the same certificate. If there are multiple pseudonym parties in her area on Pseudonym Day, she may freely choose which party to attend and thus which pseudonym provider to trust.

Operators of online services must similarly trust pseudonym providers to enforce the one person, one persona principle. A web service's administrators might simply configure their site with an explicit list of pseudonym providers it considers trustworthy, in the same way a web browser vendor chooses the default set of SSL root certificates for its browser.

Bringing the nodes of a decentralized peer-to-peer system into agreement on a set of pseudonym providers to trust might be more of a challenge. One approach is to "slice" the peer-to-peer cloud by pseudonym provider, so that instead of one large DHT for example, each node joins one DHT for *each* pseudonym provider it trusts, each DHT containing all nodes that trust a particular provider. Each node then performs a given lookup in the DHTs for each of its trusted providers and combines the results. More efficient solutions are obviously desirable, however.

3.4 Federated Pseudonym Parties

Not everyone can show up at one location on the same day, of course, so to scale geographically, many pseudonym parties must occur in different locations. If Pseudonym Day occurs at approximately the same time everywhere and all pseudonym parties follow adequately standardized procedures, a user should be able to drop into any nearby pseudonym party wherever he happens to be on Pseudonym Day to obtain his yearly certificate.

In theory any group could independently organize a pseudonym party anywhere in the world, generating its own certificates and running its own pseudonym servers. In practice, however, groups will need to federate into larger organizations with procedural controls and peer review, in order to persuade operators of online services that the federation's certificate handout procedures and pseudonym servers are trustworthy. With inadequate security or organizational transparency, for example, malicious organizers could generate more certificates than the number of people who showed up to a party and use the extras themselves, or insert a "back door" in a pseudonym server allowing themselves to generate certificates on demand. The accountability problem thus shifts from keeping *users* accountable to keeping *groups* of organizers accountable.

The organizational and security challenges of administering a federation of pseudonym parties resemble in some ways those of administering a democratic election, suggesting similar considerations and structures. All pseudonym parties need not fit under a single administration, however: several federations might evolve independently, covering distinct or overlapping geographic regions, each with its own policies and pseudonym server infrastructure.

3.5 Operating Costs

If the organizational shape of a federation of pseudonym parties vaguely resembles an election administration, we might likewise expect the costs of running pseudonym parties to bear some similarity to the costs of administering an election. Pseudonym parties present three notable differences in cost model, however:

- Election costs are "over" once the election is decided, but a pseudonym party federation must operate pseudonym servers throughout the subsequent year. These costs should be predictable and not very labor-intensive, however, since the servers merely need to be kept running and provisioned to meet the demand of the fixed number of users who obtained certificates that year.

- On the other hand, governmental elections involve registering voters and verifying citizenship and voting eligibility; the costs of these procedures do not apply to pseudonym parties since by definition anyone who shows up is eligible.

- Election commissions must be provisioned to handle *all* eligible voters who might show up to vote. If a pseudonym party reaches capacity, however, people can go to other parties in the area or, in the worst case, wait and organize their own party next year. Starting a new party should be relatively easy and inexpensive, and costs should scale together with available volunteer time and funding, in proportion to the number of active local participants.

A recent UN study of several countries found election costs to be typically \$1–3 per voter in developed countries and \$4–8 per voter in stable countries with less electoral experience [16]. If costs can be kept to similar levels *per capita*, pseudonym parties should be able to cover their costs through voluntary donations or a nominal cover charge.

4. DEPLOYMENT

Unlike identity-based single sign-on services or traditional public-key infrastructure (PKI), pseudonym account services do not need to be widely deployed "all at once" before they become useful at all.

Non-profit organizations and special-interest groups that operate primarily within a local geographic region, for example, might initially both run online services of interest to the local public and organize pseudonym parties to provide pseudonymous credentials for accessing their own online services, protecting their own online community forums from abusers both geographically local and remote. Ad hoc groups and organizations might in this way start with a purely local focus and gradually expand the useful geographical scope of the pseudonymous credentials they hand out by federating with other similarly developing groups and organizations in other geographic areas. Ideally a pseudonym account obtained on Pseudonym Day anywhere in the world should eventually be usable to create accountable pseudonymous identities on online services anywhere else in the world, but this long-term ideal need not be achieved all at once.

Popular web sites that represent global participatory communities operating using deliberative democratic procedures, such as Wikipedia and Slashdot, are particularly sensitive to sybil attacks in the form of ballot stuffing or sock puppetry, but these same communities also tend to have many users who are concerned with preserving privacy and the ability to participate anonymously. Since pseudonym parties currently appear to be the only proposed solution that can address both strong accountability and privacy at the same time, these online services could benefit greatly from such a scheme, and might therefore represent a likely context for initial experimentation with and deployment of pseudonym parties.

There are of course many additional issues and details to work out in the implementation of such a scheme, though we wish to avoid specifying too many technical details at this point in the interest of focusing the discussion for now on higher-level social and usability issues.

Here are a few such areas for discussion:

- Is there a safe way to give new users "first-time" pseudonym accounts immediately when they learn about the system, without forcing them to wait up to nearly a year until the next Pseudonym Day?

- Should there be "backup" mechanisms to obtain pseudonym accounts in case a person is sick or otherwise immobile on Pseudonym Day, or is at a location where there is not yet any organized pseudonym party?

- Might pseudonym parties be allowed to give users the choice of showing ID and attaching personal information to their pseudonym account, so that they could use the same account for both anonymous and identity-based single sign-on if they wish to?

- Can we (and should we try to) prevent a rich person or organization from paying people to attend pseudonym parties and collecting the resulting certificates?

- What specific software do pseudonym account servers and participating online services need, and what is the protocol by which they interact? Could existing identity-based single sign-on infrastructure be reused and adapted to this purpose?

- Can we avoid requiring that pseudonym accounts be accessible at all times in order to log in to services, thus risking reduced availability?

- To what extent, if any, should pseudonym parties and affiliated supporting organizations be allowed or encouraged to build ties or accept the support of governments or for-profit corporations?

5. CONCLUSION

Combating the sybil attacks at the heart of many online problems such as spam, wiki vandalism, and online ballot box stuffing, need not and should not force us to give up our privacy. Pseudonym parties would protect users' ability to maintain multiple disconnected, potentially anonymous online personas, while ensuring accountability and allowing online services to enforce the democratic "one person, one vote" principle when appropriate.

Acknowledgements

We would like to thank Frans Kaashoek, Robert Morris, and the anonymous reviewers for their many helpful comments. This research is sponsored by the T-Party Project, a joint research program between MIT and Quanta Computer Inc., Taiwan, by the National Science Foundation under Cooperative Agreement ANI-0225660 (Project IRIS) and the NSF FIND program, with additional support from Nokia Research Center Cambridge.

6. REFERENCES

[1] Adam Back. Hashcash — a denial of service counter-measure, August 2002. http://www.cypherspace.org/adam/hashcash/.

[2] V. Alex Brennen. The keysigning party howto. http://cryptnet.net/fdp/crypto/keysigning_party/en/keysigning_party.html.

[3] Martin Casado and Michael J. Freedman. Peering through the shroud: The effect of edge opacity on IP-based client identification. In *4th NSDI*, Cambridge, MA, April 2007.

[4] Kumar Chellapilla, Kevin Larson, Patrice Simard, and Mary Czerwinski. Computers beat humans at single character recognition in reading based human interaction proofs (HIPs). In *2nd Conference on E-mail and Anti-Spam*, July 2005.

[5] Curtis Chong. Graphical verification: Another accessibility challenge. *The Braille Monitor*, November 2003.

[6] Ram Dantu and Prakash Kolan. Detecting spam in voip networks. In *Proc. SRUTI*, 2005.

[7] David Davenport. Anonymity on the Internet: why the price may be too high. *Communications of the ACM*, 45(4):33–35, April 2002.

[8] Distributed Checksum Clearinghouse. http://www.dcc-servers.net/dcc/graphs/.

[9] Cory Doctorow. Solving and creating CAPTCHAs with free porn. *Boing Boing*, January 2004.

[10] John R. Douceur. The sybil attack. In *1st International Workshop on Peer-to-Peer Systems*, March 2002.

[11] Joris Evers. ISPs versus the zombies. *CNET News*, July 2005.

[12] E. Friedman and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199.

[13] Katie Hafner. Growing Wikipedia refines its 'anyone can edit' policy. *New York Times*, June 2006.

[14] Rosa R. Heckle and Wayne G. Lutters. Privacy implications for single sign-on authentication in a hospital environment (poster). In *Symposium on Usable Privacy and Security*, July 2007.

[15] Hannah Hsieh. Doonesbury online poll hacked in favor of mit. *MIT Tech*, 126(27), June 2006.

[16] IFES and UNDP. Getting to the CORE: A global survey on the cost of registration and elections, June 2006.

[17] Adam Kalsey. Why IP banning is useless, February 2004. http://kalsey.com/2004/02/why_ip_banning_is_useless.

[18] Will Knight. Wikipedia tightens editorial rules after complaint. *New Scientist*, December 2005.

[19] Gene J. Koprowski. Spam filtering and the plague of false positives. *TechNewsWorld*, September 2003.

[20] Liberty Alliance Project. http://www.projectliberty.org/.

[21] M. Lynch. Lives on hold: the human cost of statelessness, February 2005.

[22] Matt May. Inaccessibility of CAPTCHA: alternatives to visual turing tests on the web, November 2005. W3C Working Group Note 23.

[23] OpenID. http://openid.net/.

[24] Anirudh Ramachandran, David Dagon, and Nick Feamster. Can DNS-based blacklists keep up with bots? In *3rd Conference on Email and Anti-Spam*, July 2006.

[25] Russell Shaw. Avoid the spam filter. *iMedia Connection*, June 2004.

[26] Spamnation. Spam statistics. http://spamnation.info/stats/.

[27] Edward Stein. Queers anonymous: Lesbians, gay men, free speech, and cyberspace. *Harvard Civil Rights-Civil Liberties Law Review*, 38(1), 2003.

[28] Brad Stone. Breaking Google Captchas for some extra cash. *New York Times*, March 2008.

[29] Brad Stone and Matt Richtel. The hand that controls the sock puppet could get slapped. *New York Times*, July 2007.

[30] Al Teich, Mark S. Frankel, Rob Kling, and Ya ching Lee. Anonymous communication policies for the Internet: Results and recommendations of the aaas conference. *Information Society*, May 1999.

[31] Brad Templeton. I remember USENET. *O'Reilly Network*, December 2001.

[32] Bill Thompson. Not as wiki as it used to be. *BBC News*, August 2006.

[33] Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford. CAPTCHA: using hard AI problems for security. In *Eurocrypt*, 2003.

[34] Jonathan D. Wallace. Nameless in cyberspace: Anonymity on the internet, December 1999. Cato Institute briefing paper No. 54.

[35] Windows Live ID. http://www.passport.net/.

[36] Paul Wouters. Personal spam statistics 1997-2004, January 2005. http://www.xtdnet.nl/paul/spam/.

[37] The constitutional right to anonymity: Free speech, disclosure and the devil. *Yale Law Journal*, 70(7):1084–1128, June 1961.

[38] Makoto Yokoo, Yuko Sakurai, and Shigeo 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.

[39] Haifeng Yu, Michael Kaminsky, Phillip B. Gibbons, and Abraham Flaxman. SybilGuard: defending against sybil attacks via social networks. *SIGCOMM Computer Communications Review*, 36(4):267–278, 2006.