

FaceTrust: Assessing the Credibility of Online Personas via Social Networks

Michael Sirivianos
Duke University
msirivia@cs.duke.edu

Kyungbaek Kim
University of California, Irvine
kyungbak@uci.edu

Xiaowei Yang
Duke University
xwy@cs.duke.edu

1. Introduction

The success of the Internet has significantly changed how people interact with each other. Rich social interactions nowadays take place online. Users read, shop, chat, work, and play on the Internet. However, unlike many social interactions in the physical world, the Internet has largely hidden the identity attributes of online users. “On the Internet, nobody knows you are a dog,” says the famous Peter Steiner cartoon.

While anonymity has brought much benefit, including protecting user privacy and free speech, it also poses considerable security threats to online activities. What to believe and whom to believe on the Internet remains extremely challenging. Naive users may easily become victims of online scams by individuals that hide their real identity. There have been numerous incidents where scammers defrauded users [1,6] through email or online social networks. Users with vested interest in a company have been caught creating fake positive reviews for the company’s products or services [5]. Pedophiles and underage users may lie about their ages during online chat sessions and to gain access to age-restricted sites.

This problem largely stems from the fact that there is currently no lightweight and effective way to verify the identity attributes (such as age and location) of online personas. The typical approaches for establishing online identities involve offline manual verification of users. For instance, a bank may require a user to bring a government-issued ID before opening an online account. Additionally, users may purchase digital certificates that are verified and issued by trustworthy authorities such as VeriSign.

These approaches, albeit effective, are heavyweight and often an overkill for many online interactions. Manual verification is slow and costly. It can become the bottleneck that prevents an online service from scaling to hundreds of millions of users. In addition, strict user authentication typically controls access to sensitive or critical resources such as bank accounts or internal networks. Instead in many realistic Internet settings, services or users do not require strong authentication, yet they may benefit greatly from partial or likely-to-be-true user identity information. For example, it suffices for an age-restricted site to know whether a user belongs to an age group, not who the user is or his exact age. Simi-

larly, an online dating service user may desire to know whether another user’s location or profession information is likely to be credible before initiating contact.

These examples motivate the design of FaceTrust, a system that enables online personas to cost-effectively obtain credentials that verify the credibility of their identity statements. In this paper, we refer to credibility as a measure of the likelihood that a user’s assertion is correct or true. FaceTrust achieves this goal by mining and enriching information embedded in Online Social Networks (OSNs), and extends an OSN to provide lightweight, extensible, and relaxed digital credentials.

We observe that OSNs already allow users to express a limited form of trust relationships using friend links. We propose to extend this ability by allowing users to tag how credible they consider their friends’ assertions, such as the identity information they post on their profiles. This process is similar to real-world background check employed by government agencies but is greatly automated by using online social networks.

As an example, a user that wishes to obtain an age certificate from his OSN provider may post on his profile that he is above 18. The user would request from his friends to tag his assertion as true or false. The OSN provider would analyze the annotated social graph to obtain the credibility of the user’s friends and subsequently compute the credibility of the user’s assertion. It would then issue to the user a credential in the form of {assertion, credibility}. Online services or users could use this OSN-issued relaxed credential to inform their interactions with the user.

We face several main challenges in realizing the above vision. First, how can the OSN reliably verify assertions made by users (§ 3.1, § 3.2, § 3.3)? Second, how can an OSN provider export the credibility information of assertions to verifiers without violating a user’s privacy (§ 3.4)? Lastly, how can we evaluate the effectiveness of our design (§4)? The main body of this paper describes our initial approaches towards addressing these challenges. We begin with a high-level overview of FaceTrust and several of its motivating examples.

2. FaceTrust Overview

System Components: Figure 1 presents an overview




Figure 1: FaceTrust overview and an age verification example. We use d and w to denote direct and tagger credibility, respectively.

of FaceTrust. The FaceTrust architecture consists of three main components: a) an OSN provider that maintains the social graph and its users' profiles; b) online users that maintain accounts with the OSN and attempt to access online services by presenting OSN-issued credentials; and c) verifying online services that regulate access to their resources or characterize user inputs based on the user's credentials.

Assumptions and Threat Model: We assume that the OSN provider is fully trusted and can issue credentials to the best of its knowledge based on the input of its users. We also assume that the OSN provider protects the privacy of its users by not revealing their tagging information. Yet, some users may choose not to reveal many of their identity attributes to the OSN. We further assume that verifying services may wish to track a user against its will. They may collude in order to link user accounts and derive a more accurate profile of a user's activities. On the other hand, we assume that users wish to remain anonymous and unlinkable to verifying services.

We make the assumption that, typically, benign users do not lie on behalf of others. Therefore, the collective information gathered from a user's acquaintances is likely to correlate positively with the truth. Of course, this assumption does not hold if a user is motivated to lie about his friends, *e.g.*, when a group of users collude to misrepresent their identities or when a single user creates multiple fake OSN accounts (Sybil attack [12]) aiming at authenticating fake identity attributes. We discuss FaceTrust's defenses against the colluder and Sybil attack in § 3.3, and evaluate the effectiveness of these defenses in § 4.

A Usage Example: Before describing FaceTrust in detail, we first use an age-verification example to shed light on how its components interact. As shown in Figure 1, User u attempts to access an age-restricted movie at the Netflix website. At the same time, u is concerned

with his anonymity and does not wish to reveal neither his real identity nor a linkable pseudonym to Netflix.

With FaceTrust, Netflix may demand an OSN-issued age credential from the user to allow access to its content. To obtain this credential, the user u must have posted an age assertion on his OSN profile, and requested his friends to tag the credibility of his age assertion before he attempts to access the age-restricted content. In this example, user u has asserted that his age is 21, and three of his friends, users x , y , and z , have tagged the assertion with boolean values *TRUE*, *TRUE*, and *FALSE* respectively. Since not all users are equally credible, the OSN provider has computed a credibility score (w) for each user x , y , and z by analyzing the social graph and their tagging history as we soon describe in § 3.3. The OSN provider computes an overall credibility score for user u 's age assertion (0.6 in this example) by aggregating u 's friends' tagged values weighted by their credibility scores (§ 3.2).

As shown in Figure 1, the OSN issues an age credential with an overall credibility that certifies that the user belongs to the restricted age group, and the user presents this credential to a Netflix software process to gain access to Netflix content. FaceTrust implements identity attribute credentials using *idemix* [9], an anonymous, unlinkable and non-transferable credential system to preserve user privacy.

More Motivating Examples: In addition to age verification, we envision that FaceTrust credentials may benefit Internet users and online services in many other ways. A few more examples include but are not limited to:

Assessing the authority or relevance of online reviews or ratings with profession credentials. Many Internet users read online reviews before making purchase decisions. Intuitively, expert opinions of an online product may appear more authoritative to others. For instance, a review on a networking textbook from a computer science professor may carry more weight than that from an average user. With FaceTrust, if an expert user desires to appear more authoritative, he may request a profession credential from his OSN provider and present this credential to an online review site when submitting his reviews.

Verifying participant eligibility. A citizen journalism site [4] may wish to verify that a user actually resides in a specified area before it accepts its report on an event that took place in that area. Similar defenses can be employed by online fora, online auction sites, and in general by any online service that wishes to restrict participants to certain groups of people such as women groups, residents in a certain geographic area, or people of certain age groups. FaceTrust can assist legitimate participants to obtain credentials that certify their eligibility.

Preventing online frauds. Scammers commonly respond to online postings alleging to be prospective participants in legitimate transactions (*e.g.*, a potential tenant of an apartment) but in reality aiming to commit “advance-fee” fraud [2]. Such attacks could possibly be averted if scammers were unable to lie about their location, affiliation, or age. To this end, a classifieds service such as Craigslist could employ FaceTrust to verify identity attributes of users that post or respond to ads. The classifieds service can then attach to each ad post or reply the corresponding verified assertions, enabling users to make more informed decisions.

3. FaceTrust Design

We now describe our design in more detail. A key challenge of the design is to accurately assess the credibility of user assertions, as malicious users may attempt to lie or collude to obtain credentials to their favor. As an initial step, we have developed a social graph analysis algorithm by leveraging prior work on attack-resistant trust metrics [17]. Our preliminary evaluation in § 4 suggests that this initial design is promising in mitigating various attacks.

3.1 Social Tagging

FaceTrust uses “social tagging” to obtain the credibility of online personas. By social tagging, we refer to OSN users posting identity assertions on their profile and their friends assigning a binary *direct credibility* value TRUE or FALSE to them.

FaceTrust categorizes user assertions into various types such as age, address, profession, expertise etc. A user posts his assertions of assorted types in his OSN profile. For instance, for the type age, an assertion has the format ‘[<, =, >], *number*’, *e.g.*, [> 18] means that the user claims to be older than 18. For the type location, the assertion has the format [{*country, state, city...*}, *string*].

For each assertion A_j^t of type t posted by a user j , j ’s friend i may tag a direct credibility score d_{ij}^A . d_{ij}^A takes two values: a) TRUE, indicating that i believes j ’s assertion; and b) FALSE, vice versa. TRUE is mapped to the integer value 1, and FALSE to -1 in the current design. A posted assertion and the associated tags are valid for a specified period of time, which is set by the OSN provider depending on the assertion type. An assertion is uniquely identified by its {type, assertion} pair, thus a user cannot repost the same assertion and reset unfavorable tags before it expires.

We note that this design assumes that users are willing to tag their friends. There is abundant evidence that suggests social tagging may be adopted by users. For example, the “How well do you know me?” Facebook application (quiz) enables users to answer questions that their friends post about themselves and has

amassed ~ 1.2 million monthly active users. It is our future work to conduct a usability study to validate the adoptability of social tagging. In addition, our evaluation (§ 4) on the effectiveness of the trust metric we employ is sensitive to the frequency with which users tag each other, and we present results for varying degrees of tagging adoptability.

To further motivate tagging among users we employ a rudimentary incentive mechanism under which users reciprocate tags to each other in a tit-for-tat fashion. Users that wish to be able to authenticate to online services need their friends to tag them. The tit-for-tat scheme dictates that a user that wishes to be tagged by a friend has to tag his friend in return. In particular, when a user i posts an assertion and wishes to be issued credentials for it, i may explicitly request from a friend f to tag the assertion. f may choose to demand that i tags one of his assertions in return. If i does not reciprocate his friend f ’s tagging, the system does not consider f ’s tag on i in computing the credibility of i ’s assertion.

The direct credibility tags are stored by the OSN provider and are known only to the OSN and the taggers. They are never made available to other users, as they represent sensitive information.

3.2 Assertion Credibility

In the FaceTrust design, an OSN provider plays the role of an authority that issues inexpensive and relaxed credentials. By relaxed, we mean that unlike a conventional certificate authority, the OSN does not guarantee that an assertion is absolutely correct. Instead, each credential is associated with an *assertion credibility* measure in $[0, 1]$ that reflects the probability of the assertion being true as estimated by the OSN. This metric resembles a “wisdom of crowds” approach.

Let F_j denote the set of friends a user j has. To compute an *assertion credibility* score a_A on an assertion A_j^t of type t and posted by user j , the OSN provider aggregates the direct credibility tags by j ’s friends as follows:

$$a_A = \max\left(\sum_{i \in F_j} w_i^t \cdot d_{ij}^A / \sum_{i \in F_j} w_i^t, 0\right) \quad (1)$$

Tags are weighted by weight w_i^t because users are not equally credible, *e.g.*, a teenager’s tag on another user’s age assertion should carry less weight than those from more trustworthy users. We employ the additional condition that if the sum of the weights of j ’s friends is below a specified threshold, a_A is 0. We discuss how we obtain these weights in the next section.

3.3 Tagger Credibility

How can FaceTrust reliably determine the weights w_i^t ? We refer to $w_i^t \in [0, 1]$ as *tagger credibility* for the assertion type t . We observe that the problem is similar to determining the trustworthiness of a user i

and thus resorting to trust metric computation. Trust metric computation refers to the set of mechanisms that compute the trustworthiness of a node in a trust graph. There are two types of trust metrics: global, where the trustworthiness of a node is the same to all other nodes; and pairwise, where a node’s trustworthiness is relative to another node. Since FaceTrust issues credentials on “ground truth” facts, such as age and profession, and not on perceptions that are relative to the querier, such as recommendations or taste, we consider global trust metrics, *e.g.* [8, 16, 17, 24] more appropriate than pairwise and subjective ones, *e.g.* [7, 14, 19].

We face two challenges in incorporating trust metric computation to determine the weights w_i^t . The first is defining how the edges in the trust graph are formed. A trust metric is computed using a trust graph, where an edge between two nodes i and j is explicitly labeled with the degree of trust that i places on j . However, this explicit trust information is not available in a social network graph. One design choice is to require users to explicitly tag other their friends with a trust estimate. However, unlike ground truths such as a friend’s age group or profession, we consider it difficult for a user to gauge the abstract trustworthiness of a friend. Instead, FaceTrust automatically extracts trust by computing the similarity between the tags of two friends using a formula that resembles the Jaccard [15] index as follows.

Let N be the total number of tags by friends i and j that involve assertions of type t tagged by both i and j . Let C be the number of tags on common assertions for which i and j are in agreement. The tagging similarity between i and j for type t is equal to C/N . If $N = 0$, the similarity is equal to 0. After this computation, we translate a social graph with tagging history into a trust graph where each edge between two friends i and j is labeled with an explicit tagging similarity. We refer to this transformed graph as the tagging similarity graph.

The second issue is to compute the tagging credibility w_i^t of each user i from the tagging similarity graph. To this end, we adopt Levien’s Advogato trust metric [17], a graph analysis algorithm based on maximum flow. We choose this max-flow-based trust metric because it has been shown to be resistant to various attacks [18, 21, 25]. Next, we briefly summarize how we compute tagger credibility w_i^t using the Advogato algorithm.

We apply the Advogato trust metric on the tagging similarity graph by treating the tagging similarity as the level of trust between two nodes. The Advogato algorithm determines the set of nodes that can be trusted at a certain trust level x , *i.e.*, whose tagger credibility w_i^t is no less than x . In the first step, the algorithm picks a highly trusted user, *e.g.*, a trusted employee of the OSN provider that is also burdened with verifying and tagging assertions of many of his acquaintances. This user

acts as the source node in the max-flow computation. Next, the algorithm prunes all edges with tagging similarity less than x .

Subsequently, the tagger credibility of users in the social graph is computed as follows. An integer capacity is assigned to each node as a function of the user’s shortest path distance from the source. Users at the same distance from the source are said to be at the same level. To obtain the taggers that have at least x credibility, the capacity of the source node is set approximately to the expected number of users that are at least x credible. The sum of the capacity of users at each subsequent level from the source should be approximately equal to the capacity of the source. Thus, as we move away from the source and the network fans out, the capacity of the users at each subsequent level diminishes.

The tagging similarity graph is then transformed into a new graph with additional edges from the users to an additional artificial supersink user. In the new graph, capacities are assigned to edges instead of users. A user i with capacity c_i is split into two nodes i^- and i^+ and one edge of capacity $c_i - 1$ is added from i^- to i^+ . The incoming and outgoing edges of i become incoming edges and outgoing edges of i^- and i^+ , respectively. In addition, one edge of capacity 1 is added between i^- and the supersink.

We compute the maximum flow from the source to the supersink using the Ford-Fulkerson algorithm in $O(E * c_{source})$ time, where c_{source} is the capacity of the source. For a graph in which trust edges correspond to tagging similarity greater than or equal to x and for assertion type t , if the edge $i^- \rightarrow i^+$ has flow greater than 0, i is accepted as being a user that is at least x credible with respect to assertions of type t . We run this algorithm multiple times for edges that correspond to tagger credibility at increasing values: $x \in \{0.5, 0.6, \dots, 1\}$. For each user i we assign tagger credibility w_i^t (Equation 1) equal to the highest credibility x among the trust graphs in which i was accepted. If i is not accepted for any tagger credibility x , $w_i^t = 0$.

Based on the analysis by Levien [17], the number of Sybils or otherwise malicious users that can be accepted as being at least x credible is bound by $\sum_{i \in S} (c_i - 1)$, where S is the set of honest users that have greater than or equal to x tagging similarity with dishonest users. Under the assumption that it is more difficult for malicious users to have high tagging similarity with high capacity nodes closer to the source than it is with lower capacity nodes further from the source, the number of accepted malicious users should be low.

3.4 OSN-Issued Credentials

After the OSN provider obtains the assertion credibility score for a user j ’s assertion A_j^t , it can issue a credential for this assertion. As shown in Figure 1, a credential

issued by an OSN will include the assertion type t , the assertion A_j^t , and the assertion credibility score.

A credential must be authenticated by cryptographic primitives such as an OSN’s public key signature. In the FaceTrust design, we use the *idemix* [9] anonymous unlinkable credential system because users may desire to preserve their anonymity and untraceability of online activities. The *idemix* system is based on an efficient non-transferable anonymous and unlinkable credential scheme introduced by Camenisch et al. [10]. An *idemix* credential does not reveal any identifying information of a user that possesses the credential, which is ideal for online verifications such as age checking. It also prevents one user from transferring his credentials to other users. More details on how we integrate *idemix* with FaceTrust can be found in [22].

4. Evaluation

To gain a better understanding on how our initial design works, we would like to evaluate all of the following aspects of FaceTrust:

- Effectiveness: How well do credibility scores correlate with the truth, and can the design withstand incorrect user tagging and colluder or Sybil attacks?
- Computational feasibility: A social network may consist of several hundreds of millions of users. Will an OSN provider have sufficient computational resources to process the social graph?
- Usability: How often and how accurately will a user tag his friends to help them obtain credentials?

It will take a full system implementation and experimentation on a real-world OSN to answer these questions. The question regarding effectiveness is particularly difficult, because trust is inherently subjective, and it might not even be feasible to obtain the ground truth. As a first step, we describe our preliminary approaches to evaluate the effectiveness of the design. We defer the computational feasibility and the usability study for future experimentations on computer clusters and with a Facebook application, respectively.

The goal of our evaluation is to demonstrate that truthful assertions get high credibility, while dishonest assertions get low credibility even in the presence of Sybil attacks. To this end, we evaluate the effectiveness of the Advogato-based trust metric and Equation 1 used by FaceTrust using a 200K sample of a crawled Facebook social graph, obtained from a previous study [13]. The average number of friends of each user in the graph is about 12 and the maximum number of friends is 313.

In our simulation, each user in the social graph posts a single assertion of the same type on his profile. We have two types of users: honest and dishonest. Honest users always post truthful assertions and dishonest users

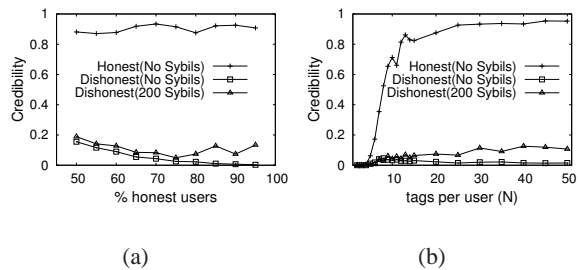


Figure 2: a) Credibility of true and false assertions as a function of the fraction of honest nodes when the maximum number of friends N a user tags is equal to 20; b) Credibility of true and false assertions as a function of N when 80% of users are honest.

always post false assertions. Both honest and dishonest users are randomly distributed in the social graph. In addition, each user tags the assertions of at most N of his friends. We vary N to reflect various degrees of adoptability of social tagging.

The honest users tag their friends truthfully, i.e., they tag as true the assertions made by their honest friends and as false the assertions made by their dishonest friends. The remaining dishonest users tag all assertions as true, regardless of whether the users that post them are honest or not. In this way, dishonest users collude to increase the credibility of each other’s assertions. By truthfully tagging assertions of honest users, dishonest users increase their tagging similarity with trustworthy users.

To evaluate the scheme’s resilience to Sybil attacks, several dishonest users each create 200 Sybil nodes, which are only connected to their creators. Sybils tag the assertions of their creator as true in order to increase the credibility of his dishonest assertions. The creator arranges to have 1.0 tagging similarity with all its Sybils. Since we assume that the users near the source are more reliable, only all the dishonest users whose distance from the source is more than five hops create Sybils.

We obtain the tagger credibility according to § 3.3. We set the capacity of the source c_{source} equal to 80% of the number of users in the graph. We randomly sample 3000 honest and 3000 dishonest assertions and compute the average credibility of each. Figure 2 plots the credibility of honest and dishonest assertions for the case in which dishonest users do not employ Sybils and the case in which they do. The credibility of honest assertions with Sybils is not included because it is almost equal to their credibility without Sybils.

In Figure 2(a), we vary the fraction of users that are honest from 50% to 95% in order to assess the trust metric’s resilience to attacks. The fraction of honest users is computed excluding the Sybil users. Each user tags at most 20 of its friends. We observe that the average credibility of honest assertions is approximately 0.9, regardless of the fraction of honest nodes. On the other hand, the credibility of dishonest assertions is very

small, *i.e.*, ~ 0.2 even when 50% of users are dishonest and Sybils are deployed. When dishonest users do not employ Sybils and the fraction of dishonest users is 5%, their opportunities for colluding by tagging each other are substantially reduced, thus the credibility of their assertions drops to almost 0.

Figure 2(b) shows the credibility of honest and dishonest assertions as a function of the maximum number of friends N that each user tags, for the fraction of honest nodes equal to 80%. As N increases, users obtain more accurate tagging similarity with their friends, increasing the credibility of true assertions and decreasing the credibility of false ones (for $N > 6$). When tagging is infrequent, *i.e.*, $N < 6$, a large portion of edges between honest users do not have high tagging similarity, as it becomes less likely for honest users to tag the same assertions. This lack of tagging information results in honest assertions getting relatively low credibility. In order to achieve reasonable assertion credibility values, N should be greater than 10.

5. Related Work

The goal of FaceTrust is mostly related to the PGP Web of Trust [3, 24, 27]. Like PGP, FaceTrust aims to circumvent the expensive and often monopolized Certificate Authorities such as VeriSign to provide lightweight credentials. Unlike PGP, FaceTrust uses the intuitive OSN interface, and employs social tagging rather than key signing to obtain trust metrics. Furthermore, FaceTrust is easily extensible, and is not limited to certifying only public keys. Users can tag each other regarding multiple types of identity attributes, and this set can be extended by adding fields into a user's profile.

FaceTrust adapts a trust metric proposed in previous work [17]. However, our contribution is not the trust metric per se. Instead, our contributions lie in the novelty of using OSNs to provide lightweight, extensible, and relaxed credentials, and the overall design and preliminary evaluation of FaceTrust. Several systems have employed trust in social networks to improve system security [11, 20, 23, 25, 26]. To the best of our knowledge, this is the first work that proposes to use OSNs to provide relaxed credentials for online personas. We provide a more extensive comparison with related work in [22].

6. Conclusion

Despite the large volume of social interactions taking place on the Internet, it is still hard to assess the credibility of statements made by online users. We present FaceTrust, a system that leverages online social networks to provide lightweight, flexible, and relaxed credentials that enable users to assess the credibility of others and their assertions. In our design, OSN users tag as true

or false their friends' identity assertions on their social network profiles. An OSN provider analyzes the social graph and the tags to assess the credibility of the assertions, and issues credentials annotated by credibility scores. Our preliminary evaluation suggests that FaceTrust is effective in obtaining credible and otherwise unavailable identity information for online personas.

7. Acknowledgements

This work was supported by NSF awards CNS-0627787 and CNS-0627166. We are obliged to the anonymous reviewers, and to Jeff Chase and Landon Cox for the long and fruitful discussions on this work.

8. References

- [1] Craigslist scams. www.craigslist.org/about/scams.
- [2] Nigerian Advance Fee Fraud. www.state.gov/www/regions/africa/naffpub.pdf.
- [3] Thawte web of trust. www.thawte.com/secure-email/web-of-trust-wot/.
- [4] Unedited, Unfiltered. News. iReport.com. www.ireport.com.
- [5] Belkin's Amazon Rep Paying For Fake Online Reviews. hardware.slashdot.org/article.pl?sid=09%2F01%2F17%2F166226&from=rss, 2009.
- [6] Teen Accused of Sex assaults in Facebook Scam. www.msnbc.msn.com/id/29032437/, 2009.
- [7] R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based Recommendation Systems: An Axiomatic Approach. In *WWW*, 2008.
- [8] S. Brin and L. Page. The Anatomy of a Large-scale Hypertextual Web Search Engine. In *Computer Networks and ISDN Systems*, 1998.
- [9] J. Camenisch and E. V. Herreweghen. Design and Implementation of the idemix Anonymous Credential System. In *ACM CCS*, 2002.
- [10] J. Camenisch and A. Lysyanskay. An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In *EUROCRYPT*, 2001.
- [11] G. Danezis and P. Mittal. SybilInfer: Detecting Sybil Nodes using Social Networks. In *NDSS*, 2009.
- [12] J. R. Douceur. The Sybil Attack. In *IPTPS*, March 2002.
- [13] M. Gjoka, M. Sirivianos, A. Markopoulou, and X. Yang. Poking Facebook: Characterization of OSN Applications. In *WOSN*, 2008.
- [14] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of Trust and Distrust. In *WWW*, 2004.
- [15] P. Jaccard. Etude Comparative de la Distribution Florale dans une Portion des Alpes et des Jura. In *Bulletin del la Soci Vaudoise des Sciences Naturelles* 37, 547-579, 1901.
- [16] S. D. Kamvar, M. Schlosser, and H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. In *WWW*, 2003.
- [17] R. Levien. Attack-resistant Trust Metrics. www.levien.com/thesis/compact.pdf, 2003.
- [18] R. Levien and A. Aiken. Attack-resistant trust metrics for public key certification. In *Usenix Security*, 1997.
- [19] P. Massa and P. Avesani. Controversial Users Demand Local Trust Metrics: An Experimental Study on opinions. com Community. In *AAAI*, 2005.
- [20] A. Mislove, A. Post, P. Druschel, and K. P. Gummadi. Ostra: Leveraging Social Networks to Thwart Unwanted Traffic. In *NSDI*, 2008.
- [21] M. Reiter and S. Stubblebine. Authentication Metric Analysis and Design. In *ACM TISSEC*, 1999.
- [22] M. Sirivianos, X. Yang, and K. Kim. FaceTrust: Assessing the Credibility of Online Personas via Social Networks. www.cs.duke.edu/~msirivia/publications/facetrust-tech-report.pdf, 2009.
- [23] Y. Sovran, A. Libonati, and J. L. Pass it on: Social Networks Stymie Censors. In *IPTPS*, 2008.
- [24] W. Stallings. Protect Your Privacy: A Guide for PGP Users. In *Prentice-Hall*, 1995.
- [25] D. N. Tran, B. Min, J. Li, and L. Subramanian. Sybil-Resilient Online Content Rating. In *NSDI*, 2009.
- [26] H. Yu, C. Shi, M. Kaminsky, P. B. Gibbons, and F. Xiao. DSybil: Optimal Sybil-Resistance for Recommendation Systems. In *IEEE S&P*, 2009.
- [27] P. Zimmerman. The Official PGP Users Guide. In *MIT Press*, 1995.