Mistaking friends for foes: An analysis of a social network-based Sybil defense in mobile networks

Abedelaziz Mohaisen
University of Minnesota
Minneapolis, MN 55455, USA
mohaisen@cs.umn.edu

Tamer AbuHmed
Inha University
Incheon, South Korea
tamer@isrl.kr

Hun Jeong Kang
University of Minnesota
Minneapolis, MN 55455, USA
hkang@cs.umn.edu

Yongdae Kim
University of Minnesota
Minneapolis, MN 55455, USA
kyd@cs.umn.edu

DaeHun Nyang
Inha University
Incheon, South Korea
nyang@isrl.kr

ABSTRACT

The Sybil attack is very challenging in the context of distributed systems; Sybil nodes with multiple identities try to deviate the behavior of the overall system from normal behavior. Recently, there have been a lot of interests in social-network based Sybil defenses weighing the trust in social networks to detect Sybil nodes. Such defenses use some algorithmic properties relating to the topological structure of the social networks. However, the use of those properties without validating them in realistic settings makes their applicability impossible in the real-world applications.

In this paper, we discuss such inapplicability by analyzing MobID, a recently proposed defense for mobile environments which claims that existing defenses have largely been designed for peer-to-peer networks. MobID uses the *betweenness*, a graph-theoretic property in the social graph, as a metric of the goodness of nodes in order to defend against the Sybil attacks. By using this betweeness, MobID operates on two fundamental assumptions: i) highly enmeshed nodes in the social graphs have a nonzero betweenness, and ii) verifiers and suspects in an honest social graph have common friends. However, extensive experiments and detailed analysis with real-world social network traces show that these assumptions do not hold well. Accordingly, MobID does not work for a great portion of the network, which is in some cases greater than 50% of the network size, even when not using a threshold on the betweenness. By setting a very low, highly-precise threshold of the betweenness (e.g., less than 10^{-4}), we observe a dramatic loss in the performance of MobID, which corresponds to $8\% - 30\%$ overall acceptance rates of honest nodes (and the remaining nodes are rejected). On the other hand, we observe that existing work, as well as other recently proposed work that is based on the community structure, can be used as an alternative for Sybil defenses in the same context.

ICUIMC'11, February 21–23, 2011, Seoul, Korea.

Categories and Subject Descriptors

C.2.0 [**Computer Communication Networks**]: General – *Security and Protection*; C.4 [**Performance of Systems**]: Design studies.

General Terms

Security, Design, Experimentation.

Keywords

Social networks, Sybil defenses, Betweenness, Mobile networks, Measurement.

1. INTRODUCTION

The Sybil attack is very challenging in the context of distributed systems [5]. In the simple form of this attack, a single node in the distributed system pretends and acts as if she is multiple nodes—by generating multiple identities of the same physical entity, and trying to deviate the behavior of the overall system from normal behavior. The attack is known to be very effective in many settings, including peer-to-peer systems [5, 2], sensor networks [20, 33, 22], file sharing systems [10], and mobile networks [34, 8, 21], among others.

The main challenge facing defenses against the Sybil attack is the lack of a centralized authority which can bind users identity to other credentials. Indeed, if one is to assume the existence of such centralized entity, defending against the Sybil attack is straightforward [20, 2, 26]. However, for that no such an authority exists in a decentralized and fully distributed system—in most deployment scenarios, the centralized solution for the problem is costly and unrealistic.

A more appealing solution to the problem in the context of distributed systems is a distributed defense that meets the structure and assumptions of the distributed systems on top of which the defense operates. Recently, there have been a lot of interests in the research community about the potential of using social networks for defending against the Sybil attack in a fully decentralized manner and using local information at each node [31, 32, 30, 3]. The basic idea of

this class of defenses is to weigh the trust exhibited in social graphs and a topological (or algorithmic) property that is known to be in the honest region—the region in which all honest nodes are. The existence of some Sybil nodes with many Sybil identities associated with these nodes, in what is called the dishonest region, would disrupt the algorithmic property of the overall graph (the honest and dishonest regions put together). This disruption can be used to isolate and detection the Sybil nodes, using property of the graph and random walk theory. Furthermore, one can bound the number of the Sybil identities introduced by the Sybil nodes, also in a distributed manner. Variations of the basic defense idea are used in many state-of-the-art social network-based defenses, as reported in [31, 32, 29, 30, 13, 3, 23].

Most recently, Quercia et al. in [21] introduced MobID, yet another design of the Sybil defenses that use social networks. This design has recently attracted attention because it newly provides a robust defense for mobile network environments while existing defenses have largely been designed for peer-to-peer networks. While it is similar to other designs in literature in the sense that it uses social networks as bootstrapping graphs, MobID differs to many of these designs in the sense that it uses the betweenness, a graph theoretic measure, as the algorithmic property for ranking the "goodness" of nodes in social graphs. The authors of MobID believe that other designs in literature, which assume a fast mixing social graph [19] and use the random walk theory, are not suitable for mobile networks. In short, [21] shows a circumstantial evidence that MobID works and detects Sybil nodes in mobile network environment. However, to reach this conclusion, the authors contradict the "trust assumption" in similar designs for defending against Sybil attack proposed in literature, and use traces of mobile encounters, which have less value of trust. Put it another way, while the design might be of a great value to defend against the Sybil attack, the use of mobile encounter traces is misleading. Whether MobID can operate in a general context and when using trust-possessing social graphs is untested.

This paper analyzes MobID and shows several fundamental findings. The contributions of the paper are as follows

- We first analyze MobID, a recently proposed social network-based Sybil defense, and show several fundamental shortcomings of it. Our analysis emphasize that the shortcomings come from the fact that MobID puts constraints on the nodes in the honest (and dishonest) region that are hard to meet, even with honest nodes. These shortcomings make the operation of MobID in realistic social network settings almost impossible.
- By experimenting with several real-world social network traces, we quantitatively show the extent to which MobID is inapplicable.

Complementary to the contributions above, we suggest that a large body of previous work can be used as a solution for the Sybil attack in the context of mobile networks. These works are not proposed by us in the first place, and their sufficiency (and shortcomings as well) can be found in the paper where they are proposed.

The rest of this paper is organized as follows. In section 2 we discuss the preliminaries required for understanding the rest of this paper. In section 3 we review the details of MobID and its operation. In section 4, we analyze MobID and show the fundamental shortcomings that make its applicability in realistic settings impossible. In section 5 we provide empirical evidence of our analysis of MobID, by experimenting with several real-world social graphs. In section 6 we discuss the implication of the findings and look for alternatives to MobID by recalling other solutions from the literature. In section 7 we review some of the related work. Finally, we draw concluding remarks in section 8.

2. PRELIMINARIES

In the section we review the preliminaries needed to understand the rest of this paper. We formalize the settings of the problem by formally describing the social network as a graph, and define the betweenness and related definitions.

2.1 Network model

The social network can be viewed as an unweighted undirected graph $G = (V, E)$ where V is the set of nodes in the graph, representing the set of people in the social network, and E is the set of edges in the graph, representing the set of interdependencies between people in the social network. For convenience, we set $V = \{v_1 \ldots, v_n\}$ (meaning that $|V| = n$) and $E = \{e_{ij}\}$ where $e_{ij} \in E$ if and only if v_i and v_j in V are adjacent in G—we also assume $|E| = m$. We define A as the adjacency matrix, where $A = [a_{ij}]^{n \times n}$ and

$$a_{ij} = \begin{cases} 1 & v_i \sim v_j \\ 0 & \text{otherwise} \end{cases}. \tag{1}$$

A random walk on G is defined using the transition matrix P. The transition matrix P captures the transition probability from any node to any other node in G. In particular, $P = [p_{ij}]^{n \times n}$ where p_{ij} is the transition probability for movingfrom node i to node j in one-hop defined as:

$$p_{ij} = \begin{cases} \frac{1}{\deg(v_i)} & v_i \sim v_j \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Notice that the movement here is conceptual; imagine a random walk capturing the flow of data over a network, then the event of moving is the propagation from node v_i to v_j using the link between them. This formulation of the transition matrix can be viewed as the norm of the adjacency matrix A by the row norm of A itself (row representing the row sum of A). In other words, let A be the adjacency matrix defined above and define D as a diagonal matrix with the diagonal element d_{ii} in D equal to $\deg(v_i) = \sum_j A_{ji}$, then P is defined as

$$P = D^{-1}A \tag{3}$$

2.2 The Betweenness Centrality

Centrality measures in general capture the significance of nodes in graphs. The betweenness centrality, as defined by Freeman [7], captures the significance of nodes to the flow between other nodes. Two types of betweenness centrality are known and defined in literature: shortest-path betweenness and random walk-based betweenness.

The shortest path based betweenness, as the name implies, weighs the significance of nodes as a result of their occurrence at the shortest path between other nodes. Let $w : E \to \mathbb{R}$ be a weight function that assigns real-valued weights

to edges in G[1]. The weight of path $p = \langle v_1, v_2, \ldots, v_\ell \rangle$ is

$$w(p) = \sum_{r=1}^{\ell} w(v_{r-1}, v_r). \quad (4)$$

The shortest path between nodes v_i and v_j is then defined as:

$$\delta(v_i, v_j) = \begin{cases} \min\{w(p) : v_i \overset{p}{\rightsquigarrow} v_j\} & \text{path from } v_i \text{ to } v_j \\ \infty & \text{otherwise} \end{cases}.$$

Since it may not be unique, a shortest path between v_i and v_j is any path with weight $w(p) = \delta(v_i, v_j)$. Now, suppose that $\sigma_{st} = \sigma_{ts}$[2] is the number of the shortest paths from the node v_s to node v_t and let $\sigma_{st}(i)$ be the number of shortest paths from v_s to v_t that pass through the node v_i. Accordingly, we define the shortest path betweenness of the node v_i as

$$b_i = \frac{2}{(n-1)(n-2)} \sum_{s \neq t \neq i} \frac{\sigma_{st}(i)}{\sigma_{st}} \quad (5)$$

While the shortest path betweenness requires a global knowledge of the whole topology—an issue that raises a lot of concerns in social network-based applications, the random walk based betweenness can be implemented at a fully decentralized settings. Recall the random walk defined earlier and recall that the random walk connecting nodes v_i and v_j, denoted as $v_i \overset{p_r}{\rightsquigarrow} v_j$, is the set of nodes where each node on the random walk is selected uniformly at random by its prior node. Similarly, we define the betweenness for a node v_i as the normalized expected number of times it is hit by the random walk. In short, the same model of the shortest path-based betweenness model can be used to compute the random walk-based betweenness by replacing the short path component in the first one by the random walk in the latter one as a measure of connecting nodes.

Note that the betweenness is one of the centrality measures. Other measures include closeness (which also uses the shortest path), degree centrality (which ranks nodes based on their degrees), and eigenvalue centrality (which ranks nodes based on their corresponding eigenvalues of P).

3. MOBID SCHEME

The scheme around which is rest of the paper is proposed in [21] for defending against Sybil attack using social networks. The main motivation of this scheme is that existing social network-based Sybil defenses (such as SybilGuard, SybilLimit, SybilInfer and others with the same flavor) are not suitable for mobile networks. In this section, we review the details of the scheme in question (MobID).

The basic idea of MobID is to use a decentralized defense mechanism that weighs the value of friends and tries to *blacklist* some others as foes. This would not be possible without the use of a social "bootstrapping" graph. The bootstrapping graph represents a graph of offline relationships that govern what people initially know about the whole network—or the list of people that they initially know in advance. Once a user - marked as a suspect - in the network wants to interact with another user - marked as the verifier, the verifier asks the the suspect to prove his honesty. The suspect exchanges some attestation information by sending his list of contacts to the verifier. The verifier then uses this information to compute the suspect's betweenness when plugging the suspects sub-network within the verifier's network of friends, where shared friends between both the suspect and verifier are only considered. If the suspect's betweenness, computed as the shortest path based betweenness, is greater than some threshold then the verifier accepts him as an honest user. Otherwise the verifier marks the suspect as a foe. Without any further due, we now review the detailed description of MobID.

1. ***Establishing trust relation:*** For those nodes where the verifier is a friend in reality, trust is established using the bootstrapping graph and maintained using an authenticated public key. For other unknown friends, measuring the extent to which the node is trusted is done indirectly. In particular, the suspect sends to the verifier a set of "links" that represent the suspect's friends signed by those friends. Used signatures are typical RSA-like signatures. Given that the verifier already has the set of authenticated public keys for his own friends, he can easily verify what is being sent to him by the suspect if it is reporting a relationship with a common friend. By doing so, the verifier verifies and incorporates edges between the suspect and the verifier's friends.

2. ***Reasoning about the friends and foes networks:*** after the first step, the verifier incorporates the list of valid links of the suspect into his network of friends and foes. This excludes two sets of the suspect friends: friends unknown to the verifier and friends who are known to the verifier but their signatures did not verify against their public keys, which is authenticated and known to the verifier. Then, the verifier ranks the suspect into both networks; the network of friends, and network of foes. The ranking in the network of friends is referred to as GoodRank and the ranking in the network of foes is referred to as the BadRank.

3. ***Deciding whether to accept or reject a suspect:*** Depending on both of suspect ranks in the good network and bad network, the verifier decides whether to accept or reject the suspect node: if the GoodRank gives a betweenness greater than a particular value, referred to as the betweenness threshold b_t, the suspect is accepts the suspect, and rejected otherwise.

To this end, MobID makes several assumptions about the underlying bootstrapping graph and the key distribution model. In the following we summarize the major used assumptions.

First, people have offline relationships that govern their "friendship". Initially each node in the graph has a list of nodes adjacent to it which the node trusts without interaction and has some attestation information about each of them.

Second, the attestation information of friendship between nodes is public keys. Public keys are shared between nodes who are adjacent to each other in an offline phase, hence

[1]In case of undirected unweighted graphs, the weight will be equal to 1 across all edges. This happens to be the case in [21] and many other social network based designs.

[2]The equality results from the optimality of the shortest path; if a shortest path between two nodes is the shortest one then it is optimal and does not matter what order is used for computing it.

not assuming or paying the cost of a centralized public key authentication entity. MobID does not assume the existence of a public key infrastructure but rather a manual public key authentication mechanism that uses a face-to-face interaction.

Third, people do not meet at random but have regular encounters, where such encounters are small in number. Also people move in groups, where members of the same community move all together as a single group. That means the encounter of a suspect implies that such suspect is more likely to be in the list of people that the verifier had an earlier encounter with, or with whom the verifier already has a relationship through the bootstrapping graph.

Last but not least, honest nodes are well-connected in social networks while Sybil nodes are sitting at the outlier part of the graph, with less connectivity to the rest of the graph.

While three of these assumptions are regularly assumed previous related work, some of these assumptions are simply contradicting the conventional literature. In particular, the first assumption is widely used as in [30, 31, 32, 3, 29, 13, 18, 25], the second assumption is being used in [9] and the last assumption is being the assumption in almost all other work that uses social networks for security and communication designs, including the work in [31, 32, 29, 30, 13, 3, 18, 25, 23]. However, the third assumption contradicts other work where mobility is assumed to be a uniform random walk on the graph in relation with prefixed locations as shown in [6]. It is worth noting that the assumption is not conventional in the context of social network-based designs, which makes our analysis even more conservative, by being less critical to the assumption of MobID. Although, challenging this assumption is not the concern of this study.

It is worth noting that the first assumption makes it necessary to use a bootstrap graph where the relationships, according to which trust is realized, exist in an offline and prior deployment phase. Though the social network application is used in mobile network context, the bootstrapping graph, with the first assumption in mind, can be any arbitrary social structure. We use this fact, in section 5 to experiment with several real-world social graphs and show the shortcomings of MobID as pointed out earlier.

4. ANALYSIS

Now we are in a position to analyze MobID, according to the description we have provided earlier. MobID has two major shortcomings that make its applicability in the context of real-world social networks almost impossible. In the following, we elaborate on these problems in more details.

The first problem in MobID is what we refer to as the "common friend(s) problem". In order for a verifier to verify a suspect—this chance by itself is not a guarantee for accepting the suspect as we will see by examples—the suspect has to have a friend with the verifier himself. Not only that, but also the number of friends common to both of the suspect and verifier has to be large enough so that some of the shortest paths between the verifier and the list of his friends flow through the suspect. For example, it is always easy to show that when the verifier and the suspect have one friend in common that it is impossible for verifier's flows to go through the suspect. This impossibility follows from the fact that the shortest path is optimal and incorporating a new node in the verifier's social graph would make the path between the verifier to any other nodes in his graph 1-hop longer than previously known shortest path.

To understand how subtle is this problem for the deployment of MobID in real social network-based design context, consider the following example. In a network with n nodes, each node can have the role of the suspect and the verifier, thus the number of suspect-verifier pairs is $n(n-1)$. At anytime, *and without considering an attacker in the honest bootstrapping graph*, the only condition needed for telling whether the performance of MobID is acceptable or not is that (*almost*) all honest nodes are accepted with high probability by any and every verifier. From the point of view of the designer, any node must accept any other node requiring any node to have at least "some" friends with every other node in the bootstrap social graph. The constraint requires the social graph to be of a special form where the radius and diameter are too small (e.g., a diameter of at most 2, where most nodes are reachable in almost one hop)—the definition of diameter and radius are in section 5. This conservative condition required for the operation of MobiID might not be possible in reality, as the case with all of the social graphs we considered in our experiments—in Table 1.

The second problem is very much related to the first problem though more critical in theoretical and application contexts. The problem is due that the ranking used by MobID assigns scores to nodes based on their betweenness, which is computed based on the frequencies of a shortest path crossing that node. Given that the shortest path is optimal, there is no guarantee that a particular node—even those with high degrees—will be located at the path of the shortest path between other nodes. Accordingly, there is no guarantee for such nodes to have a valued (*non-zero*) betweenness.

In the following we give two simple examples to show where the design of MobID does not work for normal-looking nodes. In the first example we show the case of two nodes with common friends but MobID does not work; for that the betweenness being assigned to the suspect node is still zero. In the second example, we show another case where a suspect is not an edge node, yet it is identified as a Sybil node for that it shares no common friends with the verifier.

Example 1: Consider the topology in Figure 1 where node v_{17} is the verifier and node v_{14} is the suspect, with the corresponding neighbors of each node as shown. The common neighbor among both nodes is the node v_{18}. However, although node v_{14} is not an edge node—it has a lot of friends and located at the path between other nodes; in fact it has more friends than the verifier himself—the verifier will always assign zero-valued betweenness to v_{14}, because it uses the shortest path betweenness and none of the shortest paths from the verifier to its neighbors passes through the suspect.

Example 2: Consider the same topology but the different order of verifier-suspect as shown in Figure 2. In this topology, let the verifier be v_{17} and the suspect be v_6. Even if v_6 is not an terminal node (a node that is connected to the rest of the graph through a single link) and is significant to the shortest path flows from the subgraph behind v_5 to the rest of the network, the node v_6 is simply identified as a Sybil node for that it does not share a common friend with the verifier. Put it another way, whatever the value of the betweenness of v_6 that it gains from its location at the path between other nodes, this value will not count when v_6 is introduced to v_{17} because v_6 has no common friend with v_{17} through which it should be introduced.

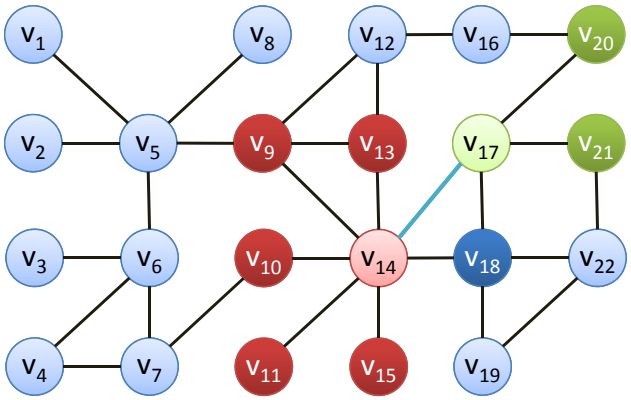

Figure 1: A social network topology scenario where MobID is applied; in order for a suspect to be introduced to the verifier in the social graph, both nodes must have common neighbors. In this topology, despite having common neighbors, the use of shortest path betweenness makes verifiers identify honest nodes as Sybil because none of verifier's shortest paths to other nodes in its network pass through the suspect.

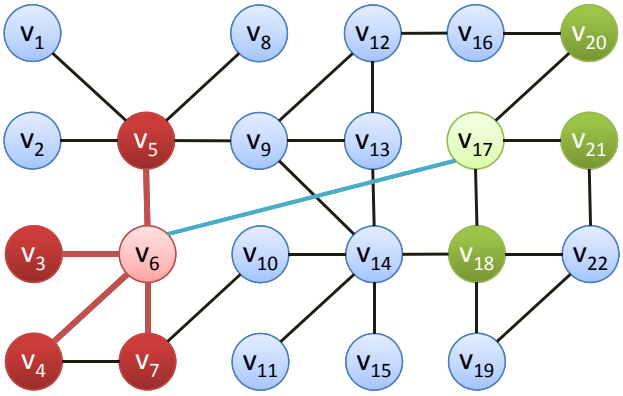

Figure 2: A social network topology scenario where MobID is applied. The case of a verifier identifying a suspect as a Sybil node for that they do not share common friends, despite that the suspect is being with high betweenness when considering it in its own community.

5. EMPIRICAL RESULTS

By experimenting with real-world social graphs, we show that MobID has several shortcomings that prevent its applicability in reality for defending against the Sybil attack. In this section we introduce the results of our study. To gain insight on the behavior of MobID, we relax the operation settings of MobID by assuming that each node has a global knowledge of the honest graph and tries to rank the different nodes according to their betweenness.

Table 1: Social graphs with their size, diameter, and radius. Physics 1, 2, 3 are relativity, high-energy, and high-energy theory co-authorship respectively. D stands for the dimeter and R stands for the radius of the graph.

Social network	Nodes	Edges	D	R
Physics 1 [12]	4,158	13,428	17	9
Physics 2 [12]	11,204	117,649	13	7
Physics 3 [12]	8,638	24,827	18	10
Wiki-vote [11]	7,066	100,736	7	4
DBLP [15]	10,000	20,684	8	4
Enron [12]	10,000	108,373	4	2
Facebook [27]	10,000	81,460	4	2
Youtube [16]	10,000	58,362	4	2

5.1 Data sets

The data sets (social graphs) used in this study are shown in Table 1 and a complementary degree distribution of each of the graphs is shown in Figure 3. The different social networks are selected "carefully" so that to reflect several tendencies in the topological structures in the underlying social graphs. First let's define the eccentricity as the longest shortest path from a source to every possible destination in the social graph. The diameter is the longest eccentricity and radius is the shortest, for a particular source. Then, we list the tendencies of the different social graphs used in this study into three types, as follows:

- ***Type 1:*** social graphs with large radius and diameter, which exhibit strong community structure. This type is represented by Physics 1, Physics 2, and Physic 3.
- ***Type 2:*** social graphs with medium radius and diameter, which exhibit medium community structure. This type is represented by Wiki-vote and DBLP.
- ***Type 3:*** social graphs with small radius and diameter, which are flat and exhibit less community structures. This type is represented by Enron, Facebook, and Youtube.

For more details, see Table 1. It is also worth noting that these graphs are sampled from larger graphs—using the breadth-first-search algorithm beginning from an arbitrary node in the original large graph. For more details on the original graphs and their topological characteristics, see [18] and [19].

5.2 Measuring the betweenness of social graphs

Measuring the betweenness of a social graph is not any harder than computing all pairs of shortest path, which has a cubic complexity in terms of the input size. Using the definition of the shortest path based betweenness explained in section 2, we compute the betweenness of the different social graphs in Table 1. The results are shown in Figure 4. Each of the sub-figures in Figure 4 plots the CDF of the betweenness for the given graph. We draw the following observations on these experiments:

- In some of the graphs, particularly which have large diameter and radius, the betweenness of almost half of the nodes in the network is *zero*. This is understandable by understanding the definition of the shortest path betweenness and imagining the existence of several small communities within these networks, which are connected to the rest of the network through a few nodes. Nodes between communities have high betweenness and nodes within communities have smaller one.

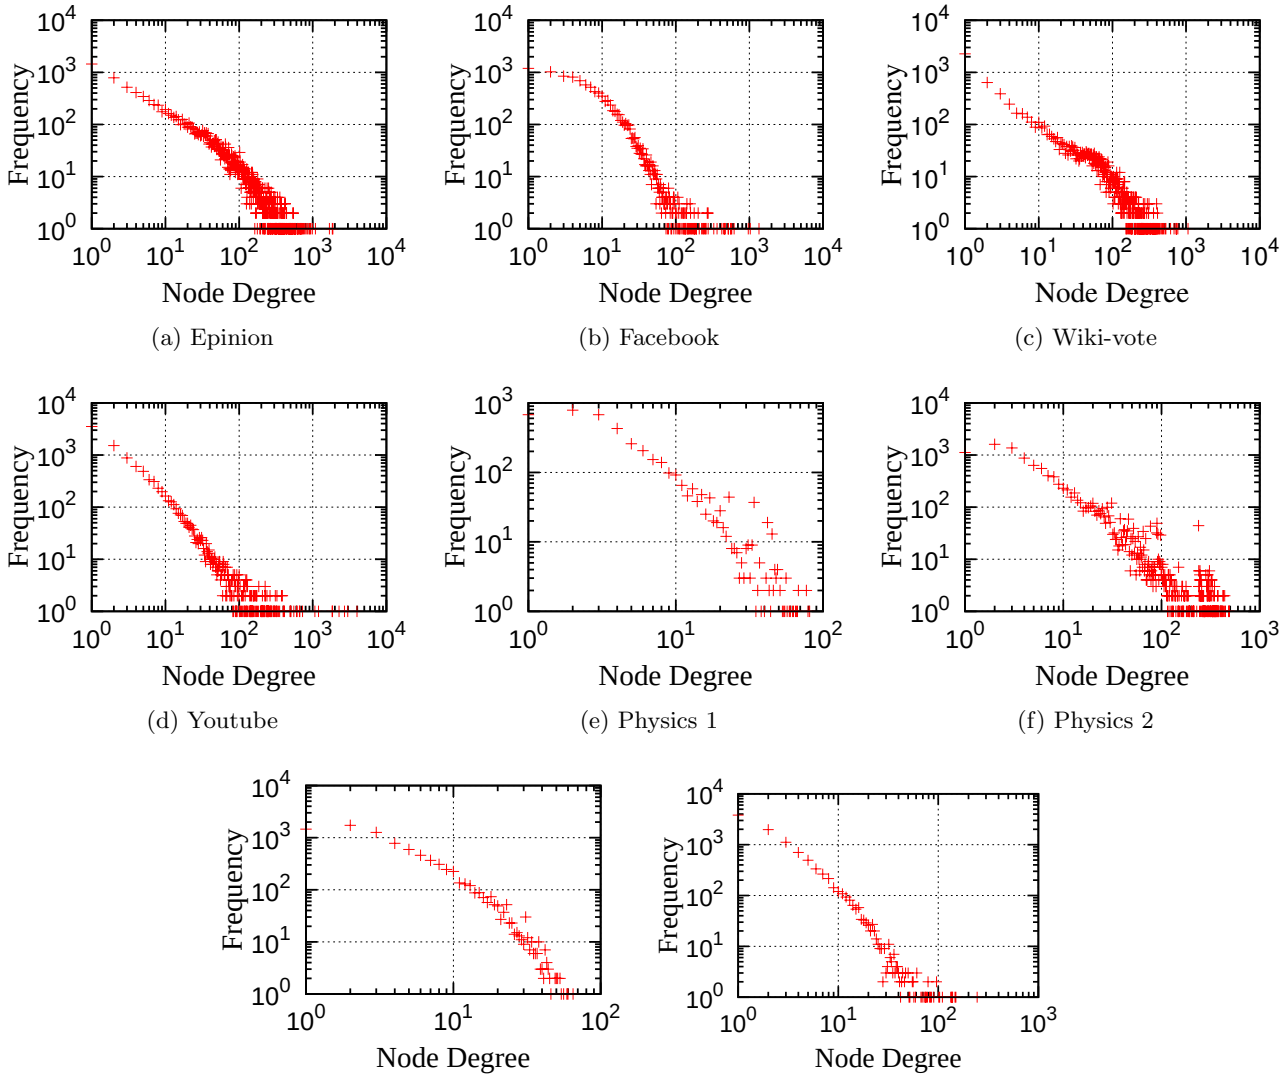

(a) Epinion (b) Facebook (c) Wiki-vote

(d) Youtube (e) Physics 1 (f) Physics 2

(g) Physics 3 (h) DBLP

Figure 3: Degree distribution of the different social graphs used in our experiments follows the power-law distribution (preferential attachment), which is the typical distribution in social networks.

- Even in those graphs with small to medium diameter and radius which have less community structure, 20% to 40% of the nodes have *zero* betweenness. One explanation for this behavior is that the shortest path is the optimal shortest distance between two nodes in the graph and this implies that some nodes, even though they are connected to other nodes in the graph, they may not be used in the shortest path (see example 1 and example 2 in section 4).

6. IMPLICATIONS AND ALTERNATIVES

In this section, we discuss the implications of the raw measurements and the implication when using a threshold in ranking nodes as Sybil and non-Sybil according to MobID. We also discuss some alternatives to the betweenness.

6.1 Implications

The implications of these measurements are striking. Considering even when a node has a global knowledge about the list of good nodes in the social graph (represented by the bootstrapping graph) there is 1/2 chance in some cases that a good node is going to be labeled as a Sybil node, even when considering a very low threshold for the betweenness. This probability is even higher in some cases. For example, even at the highest precision of the betweenness computed for DBLP, Youtube, Physics 1, 2 and 3, as shown in Figure 4, with probability exactly 1/2, a good node is going to be labeled as a bad node. What does that mean? In other words, from the point of view of each node in the social graph, half of the social graph is Sybil, despite being legitimate and honest nodes in the social graph.

6.2 Considering threshold on the betweenness

Is that the worst that MobID can achieve? In fact, MobID assumes a *threshold* on the betweenness for accepting any node in the social graph. In the previous subsection, we considered the threshold to be any value greater than zero. If we use a threshold, then how worse the result could be?

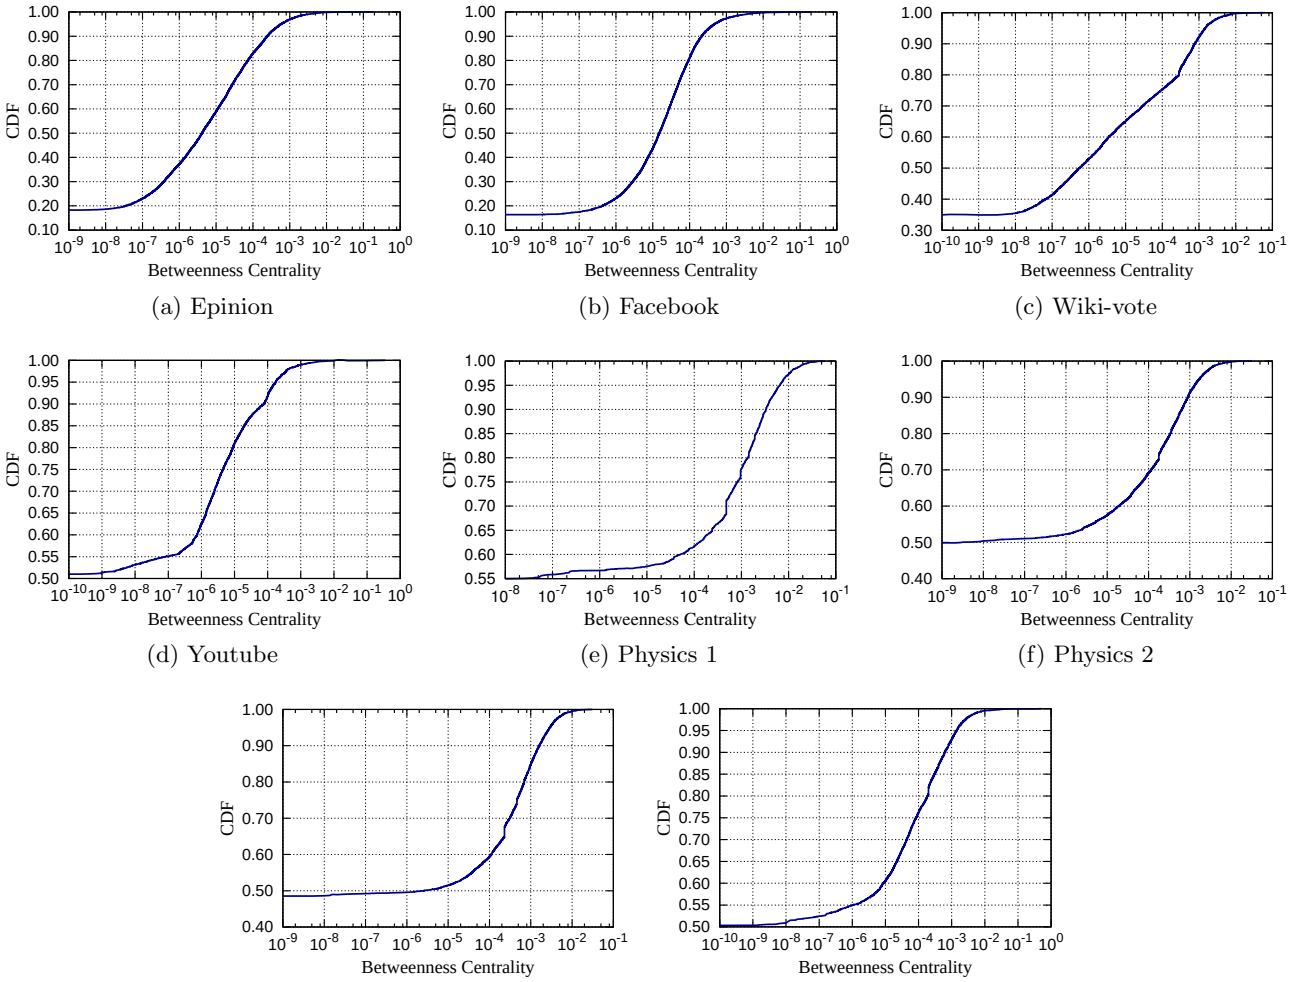

(a) Epinion (b) Facebook (c) Wiki-vote

(d) Youtube (e) Physics 1 (f) Physics 2

(g) Physics 3 (h) DBLP

Figure 4: The CDF of betweenness in the different social graphs where high percents of nodes have zero betweenness. Notice the log (base 10) scale of the x-axis.

Figure 4 answers this question. Recall the best that MobID can achieve in the settings of the social graphs used in our experiments, as shown in Figure 4(a) to 4(h), which are as follows: Epinion (80% at $b_t = 10^{-7}$), Facebook (85% at $b_t = 10^{-7}$), Wiki-vote (65% at $b_t = 10^{-8}$), Youtube (48% at $b_t = 10^{-9}$), Physics 1 (45% at $b_t = 10^{-8}$), Physics 2 (50% at $b_t = 10^{-8}$), Physics 3 (50% at $b_t = 10^{-6}$), and DBLP (50% at $b_t = 10^{-9}$).

What if one is to use a fixed threshold across the different social graphs? How is the rate of acceptance of MobID? By referring to Figure 4, we get the following: Epinion (20% at $b_t = 10^{-4}$), Facebook (20% at $b_t = 10^{-4}$), Wiki-vote (25% at $b_t = 10^{-4}$), Youtube (8% at $b_t = 10^{-4}$), Physics 1 (38% at $b_t = 10^{-4}$), Physics 2 (30% at $b_t = 10^{-4}$), Physics 3 (40% at $b_t = 10^{-4}$), and DBLP (25% at $b_t = 10^{-4}$). While none of these is an acceptable result for a Sybil defense in realistic settings, we observe that originally bad-performing social graphs are more resilient to the assigned threshold.

7. RELATED WORK

Up to now, we have shown that the operation of MobID is almost impossible in the context of mobile networks, as well as any other context—by showing that it mislabels nodes in the social graph into foes when they are legitimate honest nodes. Now, we discuss potential alternatives that one may use instead of MobID. These alternatives are basically conventional alternatives or social network-based alternatives.

Newsome et al. [20] discussed defenses for Sybil attacks on sensor networks, a closely related networking context to that of MobID's environment. In short, such defenses may include resources testing, which includes radio resource testing, the use of pre-shared keys, position verification, and code attestation. All of these can be used to replace MobID. The native environment of such defenses are sensor networks, which share a lot of common characteristics with general mobile networks; which makes the applicability of these defenses in mobile networks possible.

Social network-based Sybil defenses, such as those proposed in [31, 32, 29, 30, 13, 3, 23] can also be used. Such deployment could be with some modifications to meet the mobile environment. Also, other recent design based on community detection in [25] might be a potential alternative.

There is a large body of work on social networks, their analysis, and designs based on them. Many applications –

not only the Sybil defenses – capitalize on the trust exhibited in these social networks and benefit from these networks in trust-demanding settings. These applications include applications for routing, recommendation systems, access control, and admission control, among others. We limit ourselves in this position to the related work on Sybil defenses and analysis of Sybil defenses design. For a relatively recent survey on related work, see [18].

Sybil defenses based on social networks are ever-increasing and include SybilGuard [31], SybilLimit [30], SybilInfer [3], SumUp [24], and Whānau [14]. Each of these designs uses two different properties that are assumed to be in social networks: trust and an algorithmic property, namely the fast mixing of the social graph. The fast mixing of the social graph implies that there is no sparse cut in social graphs on top of which the Sybil defense is built, or that the honest region of the social graph is a single large community and that random walks beginning from an arbitrary good node in that region would reach every other node in the graph with probability driven according to the "stationary distribution" of the random walk. In [19] and [4], the mixing time is measured, where it is shown in [19] that some of the social graphs are not as fast mixing as thought, and so social network-based defenses that use the mixing time as an algorithmic property may have weaker guarantees. This result was supported by the work in [25] where such Sybil defenses are shown to be sensitive to community structures in social graphs, which correspond to slower mixing time.

Recent state-of-the-art online social network analysis can be seen in [16] and [1], and other that measures the community structure of the social graphs can be seen in [12]. While these works measure different characteristics of the social graphs, none of them consider the betweenness of nodes despite its importance. This lack of measurements of the property is probably due to the complexity of computations required for the betweenness, which as we have shown earlier required computing all shortest paths between nodes in the graph. Especially, this is a computationally non-trivial task for very large social graphs (at the scale of millions of nodes) which happens to be the case in most social graphs. While our work is sufficient to the extent it's proposed, we in principle use breadth-first-search based sampling to perform the measurements on a representiative graph. In further investigations, we would like to study the possibility of measuring the betweenness on larger graphs.

Last but also equally important, interaction in social networks and interaction graphs are studied in [27, 17] and strength of links is modeled in [28], which all share the flavor of measurements, and sometimes reasoning about systems built on top of social networks, with our work.

8. CONCLUSION AND FUTURE WORK

In this paper we revised MobID, a recently proposed design for defending against Sybil attack in mobile networks using social networks. In particular, we show and empirically demonstrate that the applicability of MobID in realistic contexts is almost impossible. Furthermore, we discuss alternatives from literature that can be used instead of MobID in mobile environments. While there has been some other work in the literature that tries to solve the problem of Sybil defenses without using social networks, the solutions are limited in many aspects, and further investigation of elegant solutions that are practical and use social networks are worthy of investigation. While the mean contribution of this paper is refuting MobID, in the near future we will look for alternatives that use social networks in mobile networks environments. Such alternative in part would consider properties that are robust against malicious behavior in the underlying social networks (such as infiltration attacks), hold for the good nodes, and can be used to distinguish such nodes from other dishonest nodes. We will consider the specific nature of mobile networks dynamics and how to accommodate for such dynamics when defending against the Sybil attack.

Acknowledgement

The second and last author are supported by Basic Science Research Program through the National Research Foundation of Korea (KRF) funded by the Ministry of Education, Science and Technology (2010-0013254). The first, third, and fourth authors were supported by a research grant from Korea Advanced Institute of Science and Technology (KAIST). We would like to thank Daniele Quercia for discussions on the ideas in his work (MobID) discussed in this paper, Rob Jansen for his help proofreading this work, and Alan Mislove and Ben Y. Zhao for sharing their social graph datasets used in this work..

9. REFERENCES

[1] Y.-Y. Ahn, S. Han, H. Kwak, S. B. Moon, and H. Jeong. Analysis of topological characteristics of huge online social networking services. In *WWW*, pages 835–844, 2007.

[2] N. Borisov. Computational puzzles as sybil defenses. In *P2P '06: Proceedings of the Sixth IEEE International Conference on Peer-to-Peer Computing*, pages 171–176, Washington, DC, USA, 2006. IEEE Computer Society.

[3] G. Danezis and P. Mittal. Sybilinfer: Detecting sybil nodes using social networks. In *NDSS*, 2009.

[4] M. Dell'Amico and Y. Roudier. A measurement of mixing time in social networks. In *STM 2009, 5th International Workshop on Security and Trust Management, September 24-25, 2009, Saint Malo, France*, 09 2009.

[5] J. R. Douceur. The sybil attack. In *IPTPS '01: Revised Papers from the First International Workshop on Peer-to-Peer Systems*, pages 251–260, London, UK, 2002. Springer-Verlag.

[6] Y. Fernandess and D. Malkhi. On spreading recommendations via social gossip. In *SPAA '08: Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures*, pages 91–97, New York, NY, USA, 2008. ACM.

[7] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, March 1977.

[8] S. Hashmi and J. Brooke. Authentication mechanisms for mobile ad-hoc networks and resistance to sybil attack. In *SECURWARE '08: Proceedings of the 2008 Second International Conference on Emerging Security Information, Systems and Technologies*, pages 120–126, Washington, DC, USA, 2008. IEEE Computer Society.

[9] T. Isdal, M. Piatek, A. Krishnamurthy, and T. E. Anderson. Privacy-preserving p2p data sharing with oneswarm. In *SIGCOMM*, pages 111–122, 2010.

[10] H. S. Kim, E. E. Jung, and H. Y. Yeom. Load-balanced and sybil-resilient file search in p2p networks. In *OPODIS '08: Proceedings of the 12th International Conference on Principles of Distributed Systems*, pages 542–545, Berlin, Heidelberg, 2008. Springer-Verlag.

[11] J. Leskovec, D. P. Huttenlocher, and J. M. Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, pages 641–650, 2010.

[12] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Statistical properties of community structure in large social and information networks. In J. Huai, R. Chen, H.-W. Hon, Y. Liu, W.-Y. Ma, A. Tomkins, and X. Zhang, editors, *WWW*, pages 695–704. ACM, 2008.

[13] C. Lesniewski-Laas. A sybil-proof one-hop dht. In *SocialNets '08: Proceedings of the 1st Workshop on Social Network Systems*, pages 19–24, New York, NY, USA, 2008. ACM.

[14] C. Lesniewski-Lass and M. F. Kaashoek. Whānau: A sybil-proof distributed hash table. In *7th USENIX Symposium on Network Design and Implementation*, pages 3–17, 2010.

[15] M. Ley. The DBLP computer science bibliography: Evolution, research issues, perspectives. In *String Processing and Information Retrieval*, pages 481–486. Springer, 2009.

[16] A. Mislove, M. Marcon, P. K. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Internet Measurement Comference*, pages 29–42, 2007.

[17] A. Mislove, B. Viswanath, K. P. Gummadi, and P. Druschel. You are who you know: Inferring user profiles in Online Social Networks. In *Proceedings of the 3rd ACM International Conference of Web Search and Data Mining (WSDM'10)*, New York, NY, February 2010.

[18] A. Mohaisen, N. J. Hopper, and Y. Kim. Keep your friends close: Incorporating trust into social network-based sybil defenses. In *INFOCOM'11: Proceedings of the 30th conference on Information communications*, 2011.

[19] A. Mohaisen, A. Yun, and Y. Kim. Measuring the mixing time of social graphs. In *ACM Internet Measurements Conference*. ACM, 2010.

[20] J. Newsome, E. Shi, D. Song, and A. Perrig. The sybil attack in sensor networks: analysis & defenses. In *IPSN '04: Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 259–268, New York, NY, USA, 2004. ACM.

[21] D. Quercia and S. Hailes. Sybil attacks against mobile users: friends and foes to the rescue. In *INFOCOM'10: Proceedings of the 29th conference on Information communications*, pages 336–340, Piscataway, NJ, USA, 2010. IEEE Press.

[22] K.-F. Ssu, W.-T. Wang, and W.-C. Chang. Detecting sybil attacks in wireless sensor networks using neighboring information. *Comput. Netw.*, 53(18):3042–3056, 2009.

[23] N. Tran, J. Li, L. Subramanian, and S. S. Chow. Brief announcement: improving social-network-based sybil-resilient node admission control. In *PODC '10: Proceeding of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 241–242, New York, NY, USA, 2010. ACM.

[24] N. Tran, B. Min, J. Li, and L. Subramanian. Sybil-resilient online content voting. In *NSDI'09: Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, pages 15–28, Berkeley, CA, USA, 2009. USENIX Association.

[25] B. Viswanath, A. Post, K. P. Gummadi, , and A. Mislove. An analysis of social network-based sybil defenses. In *SIGCOMM*, 2010.

[26] A. Vora, M. Nesterenko, S. Tixeuil, and S. Delaët. Universe detectors for sybil defense in ad hoc wireless networks. In *SSS '08: Proceedings of the 10th International Symposium on Stabilization, Safety, and Security of Distributed Systems*, pages 63–78, Berlin, Heidelberg, 2008. Springer-Verlag.

[27] C. Wilson, B. Boe, A. Sala, K. P. Puttaswamy, and B. Y. Zhao. User interactions in social networks and their implications. In *EuroSys '09: Proceedings of the 4th ACM European conference on Computer systems*, pages 205–218, New York, NY, USA, 2009. ACM.

[28] R. Xiang, J. Neville, and M. Rogati. Modeling relationship strength in online social networks. In *WWW '10: Proceedings of the 19th international conference on World wide web*, pages 981–990, New York, NY, USA, 2010. ACM.

[29] H. Yu, P. B. Gibbons, and M. Kaminsky. Toward an optimal social network defense against sybil attacks. In *PODC '07: Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, pages 376–377, New York, NY, USA, 2007. ACM.

[30] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. Sybillimit: A near-optimal social network defense against sybil attacks. In *SP '08: Proceedings of the 2008 IEEE Symposium on Security and Privacy*, pages 3–17, Washington, DC, USA, 2008. IEEE Computer Society.

[31] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman. Sybilguard: defending against sybil attacks via social networks. In *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 267–278, New York, NY, USA, 2006. ACM.

[32] H. Yu, M. Kaminsky, P. B. Gibbons, and A. D. Flaxman. Sybilguard: defending against sybil attacks via social networks. *IEEE/ACM Trans. Netw.*, 16(3):576–589, 2008.

[33] Q. Zhang, P. Wang, D. S. Reeves, and P. Ning. Defending against sybil attacks in sensor networks. In *ICDCSW '05: Proceedings of the Second International Workshop on Security in Distributed Computing Systems (SDCS) (ICDCSW'05)*, pages 185–191, Washington, DC, USA, 2005. IEEE Computer Society.

[34] T. Zhou, R. R. Choudhury, P. Ning, and K. Chakrabarty. Privacy-preserving detection of sybil attacks in vehicular ad hoc networks. In *MobiQuitous'07*, pages 1–8, Washington, DC, USA, 2007. IEEE Computer Society.