Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 10-019

Keep your friends close: Incorporating trust into social network-based Sybil defenses

Abedelaziz Mohaisen, Nicholas J. Hopper, and Yongdae Kim

August 24, 2010

Keep your friends close: Incorporating trust into social network-based Sybil defenses

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

Nicholas Hopper
University of Minnesota
Minneapolis, MN 55455, USA
hopper@cs.umn.edu

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

ABSTRACT

Social network-based Sybil defenses exploit the algorithmic properties of social graphs to infer the extent to which an arbitrary node in such a graph should be trusted. However, these systems do not consider the different amounts of trust represented by different graphs, and different levels of trust between nodes, though trust is a crucial requirement in these systems. For instance, co-authors in an academic collaboration graph are trusted in a different manner than social friends. Furthermore, some social friends are more trusted than others. However, previous designs for social network-based Sybil defenses have not considered the inherent trust properties of the graphs they use. In this paper we introduce several designs to tune the performance of Sybil defenses by accounting for differential trust in social graphs and modeling these trust values by biasing random walks performed on these graphs. Surprisingly, we find that the cost function, the required length of random walks to accept all honest nodes with overwhelming probability, is much greater in graphs with high trust values, such as co-author graphs, than in graphs with low trust values such as online social networks. We show that this behavior is due to the greater number of close-knit communities in high-trust graphs, requiring longer walk to traverse multiple communities. Furthermore, we show that our proposed designs to account for trust increase the cost function of graphs with low trust value.

1. INTRODUCTION

The Sybil attack is a well-known and powerful attack in distributed systems, such as sensor networks and peer-to-peer systems. In the basic form of this attack, a peer representing the attacker generates as many identities as she can and acts as if she is multiple peers in the system, which are then utilized to influence the behavior of the system [1]. The number of identities that an attacker can generate depends on the attacker's resources such as bandwidth, memory, and computational power. With the sharp hardware growth—in terms of storage and processing capacities—and the popularity of broadband Internet, even attackers who use "commodity" hardware can cause a substantial harm to large systems.

Despite being known for long time, this attack lacked technical defenses and many papers have reported its existence without suggesting any defense while many proposed defenses are limited in many aspects [2]. The majority of defenses proposed in literature to defend against, limit, or mitigate the Sybil attack can be classified into centralized defenses and decentralized defenses. In the centralized defenses (e.g., [1, 3, 4, 5]), a centralized authority is responsible for verifying the identity of every user in the systems. Because they depend on a centralized authority, these defenses are ruled out in many distributed settings. On the other hand, the decentralized defenses (e.g., [6, 7, 8, 9, 10]) utilize distributed approaches to bind credentials to the identities of peers, and verify the peers authenticity.

A recent class of the decentralized defenses uses social networks, where peers in the network are not merely computational entities—the human users behind them are tied to each other to construct a social network. The social network is then used for bootstrapping the security and detecting Sybils under two assumptions: algorithmic and sociological. The algorithmic assumption is the existence of a "sparse cut between the Sybil and non-Sybil subgraphs" in the social network, which implies a limited number of attacker edges; edges between Sybil and non-Sybil nodes. The sociological assumption is a constraint on the trust in the underlying social graph: the social graph used in these defenses needs to exhibit strong trust as evidenced, for example, by face-to-face interaction demonstrating social actors' knowledge of each other [9, 11]. While the first assumption has been recently questioned in [12], where it is shown that even honest subgraphs may have cuts that disrupt the algorithmic property, the trust—though being a crucial requirement for these designs to perform well—was not considered carefully. Even worse, many of these defenses [9, 11, 13, 14]—when verified against real-world social networks—have considered samples of online social graphs, which are known to possess weaker value of social trust.

We have recently measured the mixing time, a concrete measure of the algorithmic property required in social networks, in [15], and demonstrated that it is greater than the values used in literature. Also, we pointed out that social graphs with same size have different mixing times implying

that social networks, even algorithmically, cannot be taken equally for the purpose of these designs (see sec. 5). However, the different mixing times are not arbitrary: social graphs that exhibit knowledge (e.g., co-authorship) or intensive interaction (e.g., social blogs) are slower mixing than social graphs that require less interaction or where edges are less meaningful (e.g., wiki-vote and online social networks such as Orkut and Facebook), which suggest that the algorithmic and trust properties in social graphs are at odds. To this end, we explore designs to model trust in social graphs in order to base the performance of the Sybil defenses more accurately on both assumptions: algorithmic and sociological.

We model the trust exhibited in the social graph as parameters of modified and biased random walks, as opposed to the uniform random walks used in Sybil defenses—where social graphs are presumed to have similar trust value. The proposed designs use two observations: nodes in the social graph trust themselves more than they trust others, and they trust other nodes unequally. We use the first observation to incur gravitational probability in the random walk – at either the current or originator node of the walk – and use the second observation to incur weights on edges between the different nodes. In the first direction we introduce the lazy and originator-biased random walks. In the second direction we introduce the similarity and interaction-biased random walks to model trust. We investigate their power in modeling trust and influencing the Sybil defenses.

Perceiving that even online social networks are potentially going to be used for Sybil defenses—as well as other applications based on social networks such as routing [16, 17, 18, 19, 20, 21], access control [22], among other applications, including those in [23, 24, 25, 26, 27, 28, 29, 30]—the end result of this paper is to prepare for these networks for such applications by accounting for weaker trust in these networks. While both strong and weak trust exhibiting networks exist and can be options for these designs, the tools we provide in this paper can be used to quantitatively compare different data sets, and different options, based on their true values and merits.

Through this paper, we only show the insight and the power of the different designs to control the behavior of the social network-based Sybil defenses. We do not insist on any particular values for parameters used in any design for a simple reason: the different designs are meant to provide the Sybil defense designer with further parameters that she can use to control the network behavior and how the Sybil defense proceed to detect Sybil nodes or admit non-Sybil nodes. The adjustment of the parameter is based on the designers perception for the trust exhibited in the social network. Put it another way, the different parameters associated with the different Sybil defenses can be further assigned by the honest users within the network based on their estimation and understanding of the trust in their own network, the ultimate number of Sybil nodes they are willing to accept, or the ultimate number of non-Sybil nodes they are willing to reject. The length of the random walk used in the Sybil defense, which is a parameter that we can control using our designs, can control both types of these issues.

Also, while we experiment the different designs with different parameters on different social graphs, even including those hypothesized to have good trust values as in section 5, we do this to understand the behavior and impact of these designs on different networks with different properties. It is worth noting that some of the networks, e.g., those hypothesized to have good trust, may not require any adjustment for their behavior. This is, it is may be possible that small (close to zero) parameter value such network would guarantee the adjustment. On the other hand, networks that are known to provide weaker trust may require larger parameters to compensate for the weaker trust in these networks. Ultimately, the parameter value is determined by each node independently or assigned by an operator who oversees the value of the trust associated with overall network. This paper provides the tools for understanding such decisions and their impact on the performance of the defense.

Contributions: The novel and original contributions of this paper are as follows. First, motivated by the observed relationship between the quality of the algorithmic property and hypothesized trust in social graphs, we propose several designs, each in the form of modified random walk, to model trust in social networks. Second, we learn the impact of the different designs on the performance of the Sybil defenses by comparing them to each other when operated on top of SyilLimit, a known work in literature on defending against the Sybil attack using social networks. For this part, we use several real-world social graphs that exhibit different levels of knowledge and trust. We provide several insights through discussions that relate to observations on the measurements.

Organization: The organization of this paper is as follows. We review some of the related work in section 2. In section 3 we introduce the preliminaries of our work. In section 4 we introduce several designs to model trust in social networks, which are used for Sybil defenses. In section 5 we discuss the main results, which include experiments on real-world social networks. In section 6 provide concluding remarks and present implications of the findings. In section 7 we discuss some of the open problems and future work.

2. RELATED WORK

There is a large body of work on social networks, their analysis, and designs based on them. Many applications – and 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. To this end, we present a select of the related work to this paper. In particular, we classify the papers of the related work into papers on Sybil defenses—where this paper is mainly aimed

at, papers on generic applications that use the trust in social network, papers on analyzing social networks, papers on understanding trust in social networks, and finally, papers on analyzing the security and requirements of the Sybil defenses and their assumptions.

Sybil defenses based on social networks are reported in SybilGuard [11], SybilLimit [9], SybilInfer [13], SumUp [31], and Whānau [14]. In principle, the performance of these defenses depends on the quality of the algorithmic property of the underlying graph by assuming good trust in the underlying social network. These studies can benefit from our findings in quantifying their performance by making up for the trust exhibited in the social graphs they operate on. A study on analyzing these designs can be seen in [12]. We further summarize the operation of one of these defenses, namely SybilLimit, in section 3.

Aside from these systems built to defend against the Sybil defense, several other systems are introduced in literature on using the trust in social graphs for building systems. For instance, Daly et al. [19, 32] studied the use of social networks for routing in disconnected delay tolerant networks, where nodes with higher similarity and betweenness (two graph theoretic properties) are favored for forwarding messages from source to destination. The end result of this work has shown that the new construction is better than the gossiping alone protocol. The protocol also explicitly assumed good trust in the underlying social graph, so that nodes are honest when reporting their locally computed graph theoretic measures. Similar results in similar settings are also reported in [17]. In [21], Marti et al. studied constructing DHT over social networks, where nodes in the social graph act as "routing entries" for their social contacts. In [27], Pai et al used the trust in social graphs for bootstrapping trust in ad-hoc networks, that can be then used for building reputation, routing, etc. In [23], the trust in social graphs is used for worm detection (similar results are in [33] and [18]). Selfishness of social node and its impact in social network-based routing applications is studied in [16]

Understanding, predicting, and analyzing interactions in social networks are studied by Viswanathet al. in [34] and by Wilson et al. in [35]. The later model is what we use for the interaction-biased random walk in this study. Understanding negative and positive links in social networks – which can be utilized in settings like our designs – in social graphs are studied in by Leskovec et al. in [36]. The similarity and centralities in social graphs are studied and evaluated on social networks in [37, 38, 39]. In [40], Quercia even used betweenness of nodes to defend against the Sybil attack. The social capital exhibited in social networks is used in [41] to replace the tit-for-tat concept in peer-to-peer systems. Classical studies on the analysis of topological features in online social and information networks can be seen in [42] and in [43].

While all of the applications above consider trust to be the main feature used from the social graphs and utilized for their performance, some papers have considered quantifying and modeling trust in social networks to understand the behavior of some online social networks (e.g., recommender systems). Such studies can be found in [44] and [45].

Finally, many applications can be built on top of social graphs to benefit from the well connectivity of these networks. This well-connectivity features a good tool for anonymity – as evidenced by being able to sample from the whole social graph after short random walk on the graph beginning from any arbitrary initial source – which is studied in [46]. In [46], Nagaraja has shown that social graphs can be good mixers where a random walk of length of about 10 hops is good enough accumulate entropy of 18 in an anonymity set of 4 million nodes – which ideally has an entropy of 22. Though with a different motivations, the mixing time of different social graphs, with the assumptions used in Sybil defenses in mind, has been recently measured in [15].

3. PRELIMINARIES

3.1 Network Model

We view the social network as an undirected unweighted graph $G = (V, E)$ where $|V| = n$, $V = \{v_1, v_2, \ldots, v_n\}$, $|E| = m$, $e_{ij} \in E = v_i \rightarrow v_j$ if $v_i \in V$ is adjacent to $v_j \in V$ for $1 \leq i \leq n$ and $1 \leq j \leq n$. We refer to $\mathbf{A} = [a_{ij}]^{n \times n}$ as the adjacency matrix where $a_{ij} = 1$ if e_{ij} is in E and $a_{ij} = 0$ otherwise. We refer to $\mathbf{P} = [p_{ij}]^{n \times n}$ as the transition matrix

$$p_{ij} = \begin{cases} \frac{1}{\deg(v_i)} & e_{ij} \in E \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where $\deg(v_i)$ is the degree v_i, or the row-norm of $\mathbf{A}$:

$$\deg(v_i) = \sum_{k=1}^{n} \mathbf{A}_{ik}. \quad (2)$$

The set of neighbors of v_i is $N(v_i)$ and $\deg(v_i) = |N(v_i)|$.

3.2 Simple Random Walks and Mixing Time

The "event" of moving from a node to another is captured by a Markov Chain (MC) which represents a random walk over G. A random walk of length w over G is a sequence of vertices in G beginning from an initial node v_i and ending at v_t, the terminal node, using the transition matrix (1). The MC is said to be ergodic if it is irreducible and aperiodic, meaning that it has a unique stationary distribution π and the distribution after random walk of length w converges to π as $w \rightarrow \infty$. The stationary distribution of the MC is a probability distribution that is invariant to the transition matrix $\mathbf{P}$ (i.e., $\pi\mathbf{P} = \pi$). The mixing time of the MC, T is defined as the minimal length of the random walk in order to reach the stationary distribution. More precisely, Definition 1 states the mixing time of MC on G parameterized by a variation distance parameter ϵ.

DEFINITION 1 (MIXING TIME). *The mixing time (parameterized by ϵ) of a Markov chain is defined as*

$$T(\epsilon) = \max_i \min\{t : |\pi - \pi^{(i)}\mathbf{P}^t|_1 < \epsilon\}, \quad (3)$$

where π is the stationary distribution, $\pi^{(i)}$ is the initial distribution concentrated at vertex v_i, $\mathbf{P}^t$ is the transition matrix after t steps, and $|\cdot|_1$ is the total variation distance. The MC is rapidly mixing if $T(\epsilon) = poly(\log n, \log \frac{1}{\epsilon})$.

Papers such as [13, 14, 9, 11] refer to this as "fast mixing" and strengthen the definition by considering only the case of $\epsilon = \Theta(\frac{1}{n})$, and requiring $T(\epsilon) = O(\log n)$.

THEOREM 1 (STATIONARY DISTRIBUTION). *For undirected unweighted graph G, the stationary distribution of the MC over G is the probability vector $\pi = [\pi_{v_i}]$ where $\pi_{v_i} = \frac{\deg v_i}{2m}$. This is, $\pi = [\frac{\deg(v_1)}{2m} \ \frac{\deg(v_2)}{2m} \ldots \frac{\deg(v_n)}{2m}]$.*

THEOREM 2 (SECOND LARGEST EIGENVALUE [47]). *Let $\mathbf{P}$ be the transition matrix of G with ergodic random walk, and λ_i for $1 \leq i \leq n$ be the eigenvalues of $\mathbf{P}$. Then all of λ_i are real numbers. If we label them in decreasing order, $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_{n-1} \geq \lambda_n > -1$ holds. We define the second largest eigenvalue modulus (SLEM) as $\mu = \max(|\lambda_2|, |\lambda_{n-1}|)$. Then, $T(\epsilon)$ is bounded by $\frac{\mu}{2(1-\mu)} \log(\frac{1}{2\epsilon}) \leq T(\epsilon) \leq \frac{\log(n)+\log(\frac{1}{\epsilon})}{1-\mu}$.*

We observe that the mixing time captures the connectivity of the graph. Well-connected graphs have small mixing time while weakly connected graphs have large mixing time [47]. Also, the second largest eigenvalue used for measuring the mixing time bounds the graph conductance, a measure for the community structure [12]. In short, the conductance $\Phi \geq 1 - \mu$.

3.3 Social Network based Sybil Defenses

As mentioned in section 2, there are several defenses to the Sybil attack using social networks. Here we limit ourselves to SybilLimit, which we use to measure our designs.

Unlike SybilGuard which uses one long random route for verification, SybilLimit [9] uses several shorter instances of random routes. A verifier as well as the suspect perform $O(\sqrt{m})$ random routes each of length $w = O(\log n)$ to obtain samples of the honest region – since $O(\sqrt{m}) = r_0\sqrt{m}$, SybilLimit fixes $r_0 = 4$ to ensure high intersection probability. The verifier determines to accept a suspect if he is registered at one of the tails in his sample. SybilLimit accepts a suspect if intersection with the verifier happens on a tail, which is the last edge of the random routes. In SybilLimit, if a tail ends up in the Sybil region, it will always end-up in it due to the random routes one-to-one pre-computed permutation structure. Also, if a tail ends up in the Sybil region, it may advertise many non-existent intersections with routes initiated by Sybil nodes. To avoid that, SybilLimit limits the number of intersections into $g \times w \times m$ intersections on honest tails – where g is the number of attack edges and w is the random walk length. This means that SybilLimit accept at most $w = O(\log n)$ Sybil identities per attack edge. SybilLimit greatly depends on w for its security and uses benchmarking techniques for estimating it. However, since these techniques are not provable, underestimating or overestimating the parameters is problematic. SybilLimit works as long as $g \leq o(\frac{n}{\log n})$.

4. DESIGNS TO ACCOUNT FOR TRUST

In most of the literature that considered social networks for building Sybil defenses, the simple uniform random walk highlighted in section 3 has been used. In this section, we introduce several designs of modified random walks that consider a "trust" parameter between nodes. In all of the proposed modified random walks, the purpose is to assign "trust-driven" weights and thus deviate from uniform. We do this by either capturing the random walk in the originator or current node, as the case of originator-biased and lazy random walks, or by biasing the random walk probability at each node, as the case of interaction and similarity-biased random walks, or a combination of them. The intuition of the lazy and originator-biased random walk is that nodes trust "their own selves" and other nodes within their community more than others. On the other hand, interaction and similarity-biased trust assignments try to weigh the natural social aspect of trust levels. Given the motivation for these designs, we now formalize them by deriving $\mathbf{P}$ and π required for characterizing them. We omit the details for lack of space (see the Appendix, theorem 3 through 7, for the complete proofs).

4.1 Lazy Random Walks

To accommodate for the trust exhibited in the social graph, we assume a global single parameter α in the network which is used to characterize this trust level and used in the different schemes to enforce and apply the trust along with other parameters used (e.g., driven from the algorithmic property in the graph). The transition matrix

$$\mathbf{P}' = \alpha\mathbf{I} + (1-\alpha)\mathbf{P} \quad (4)$$

which yields a transition according to p_{ij} defined as follows:

$$p_{ij} = \begin{cases} \frac{1-\alpha}{\deg(v_i)} & v_j \in N(v_i) \\ \alpha & v_j = v_i \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

We note that for the transition probability defined in (4), by adding self loops it does not alter the final stationary distribution from that in Theorem 1. This is, since $\mathbf{P}' = \alpha\mathbf{I}+(1-\alpha)\mathbf{P}$, by multiplying both sides by π, we get $\pi\mathbf{P}' = \pi(\alpha\mathbf{I}+(1-\alpha)\mathbf{P}) = \alpha\pi\mathbf{I}+(1-\alpha)\pi\mathbf{P} = \alpha\pi+\pi-\alpha\pi = \pi$.

4.2 Originator-biased Random Walk

We incorporate the concept of biased random on the social graph walks to characterize the bias introduced by the trust

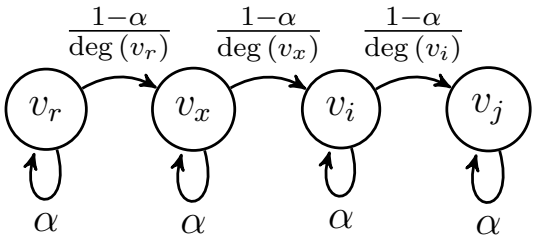

Figure 1: An illustration of the lazy random walk. For simplicity, α is equal for each node though it can be determined by each node locally to reflect what that node perceive as the trust of the network. The random walk in this example is (v_r, v_x, v_i, v_j)—adjacent nodes other than the next in the walk to each node are omitted to simplify the illustration.

among different social actors (nodes). At each time step, each node decides to direct the random walk back towards the node that initiates the random walk, i.e., node v_r, with a fixed probability α or follow the original simple random walk by *uniformly* selecting among its neighbors with the total remaining probability $1-\alpha$. The transition probability that captures the movement of the random walk, initiated by a random node v_r, and moving from node v_i to node v_j is defined according to p_{ij} as follows

$$p_{ij} = \begin{cases} \alpha & j = r, v_r \notin N(v_i) \\ \alpha + \frac{1-\alpha}{\deg(v_i)} & j = r, v_r \in N(v_i) \\ \frac{1-\alpha}{\deg(v_i)} & j \neq r, v_j \in N(v_i) \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

We note that, unlike the lazy random walks, the transition probability here considers moving the state back to the originator of the random walk, a state that may not be connected to the current state in the social graph. This requires a virtual connection between each node through the walk – every node in the graph – and each originator of a random walk. To mathematically model this transition loop, for each node $v_r (1 \leq r \leq n)$, we define $\mathbf{A}_r$ as an all-zero matrix with the exception of the r^{th} row which is 1's. Using $\mathbf{A}_r$, we further define the originator-biased transition matrix, for the walk originated from v_r, as

$$\mathbf{P}' = \alpha \mathbf{A}_r + (1-\alpha)\mathbf{P}. \tag{7}$$

We can show that $\mathbf{P}'$ is stochastic since each row in it sums to 1. Furthermore, since $\mathbf{P}'$ depends on the initial state v_r, we observe that the "stationary" distribution is not unique among all initial states, and so we refer to it as the "bounding distribution" for the walk initiated from v_r. The bounding distribution in that case is $\pi^{(v_r)} = [\pi_i]^{1\times n}$ where π_i is

$$\pi_i = \begin{cases} (1-\alpha)\frac{\deg(v_i)}{2m} & v_i \in V \setminus \{v_r\} \\ \alpha + \frac{\deg(v_i)}{2m} & v_i = v_r \end{cases} \tag{8}$$

We note also that the bounding distribution in (8) is a valid probability distribution since $\alpha + \frac{\deg(v_r)}{2m} + \sum_{v_i \in V\setminus\{v_r\}}(1-\alpha)\frac{\deg(v_i)}{2m} = \alpha + \sum_{i=1}^{n}(1-\alpha)\frac{\deg(v_i)}{2m} = \alpha + (1-\alpha)\sum_{i=1}^{n}\frac{\deg(v_i)}{2m} = \alpha + (1-\alpha) = 1$. It is also easy to show that given distribution bounds the random walk since $\pi\mathbf{P}' = \pi$.

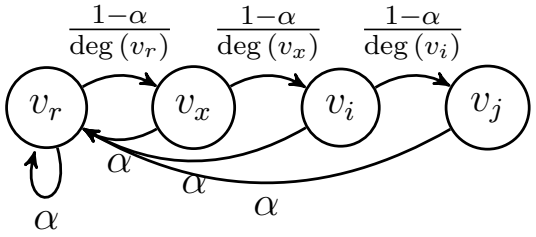

Figure 2: An illustration of the originator biased random walk. For simplicity, α is equal for each node though it can be determined by each node locally to reflect what that node perceive as the trust of the social network. The random walk in this example is (v_r, v_x, v_i, v_j)—adjacent nodes other than the next in the walk to each node are omitted to simplify the illustration.

4.3 Interaction-biased Random Walk

The interaction between nodes can be used to measure the strength of the social links between nodes in the social network [35]. In this model, high weights are assigned to edges between nodes with high interaction and low weights are assigned to edges between nodes with low interaction. Formally, let $\mathbf{B}$ be the raw interaction measurements between nodes in G and $\mathbf{D}$ be a diagonal matrix representing the row norm of $\mathbf{B}$, computed as in (2). The transition matrix $\mathbf{P}$ of the random walk based on interaction is then computed as $\mathbf{P}' = \mathbf{D}^{-1}\mathbf{B}$. The stationary distribution of the random walk on G following to the probability in $\mathbf{P}'$ is $\pi = [\pi_i]^{1\times n}$ where

$$\pi_i = (\sum_{j=1}^{n}\sum_{k=1}^{n} b_{jk})^{-1}(\sum_{z=1}^{n} b_{zi}). \tag{9}$$

We observe that this distribution makes a valid probability distribution since $\sum_{i=1}^{n}\pi_i = 1$ and is a stationary distribution since $\pi\mathbf{P} = \pi$.

Wilson et al. [35] introduced a slightly different model to capture interaction between nodes in the social graph. The interaction graph $G' = (V, E')$ is defined for a social graph $G = (V, E)$ where $E' \subseteq E$ and $e_{ij} \in E'$ if $I(v_i, v_j) \geq \delta$, where I is an interaction measure to assign weights on edges between v_i and v_j for all i, j, and δ is a threshold parameter. The interaction measure used in [35] is the number of interactions over a period of time. This model further simplifies the random walk where the $\mathbf{P}'$ is defined over G', as well as the stationary distribution.

4.4 Similarity-biased random walk

The similarity between social nodes in social networks is used for measuring the strength of social links and predicting future interactions [37, 39]. For two nodes v_i and v_j with sets of neighbors $N(v_i)$ and $N(v_j)$, respectively, the

similarity is $\frac{N(v_i)\cap N(v_j)}{N(v_i)\cup N(v_j)}$. For $\mathbf{a}_i$ and $\mathbf{a}_j$, two rows in $\mathbf{A}$ corresponding to the entries of v_i and v_j, we use the cosine similarity measure given as $S(v_i, v_j) = \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{|\mathbf{v}_i|_2 |\mathbf{v}_j|_2}$, where $|\cdot|_2$ is the L2-Norm. To avoid disconnected graphs resulting from edge cases, we augment the similarity by adding 1 to the denominator to account for the edge between the nodes. Also, we compute the similarity for adjacent nodes only, so that $\mathbf{S} = [s_{ij}]$ where $s_{ij} = S(v_i, v_j)$ if $v_j \in N(v_i)$ or 0 otherwise. The transition matrix $\mathbf{P}$ of a random walk defined using the similarity is given as $\mathbf{P} = \mathbf{D}^{-1}\mathbf{S}$ where $\mathbf{D}$ is a diagonal matrix with diagonal elements being the row norm of $\mathbf{S}$. Accordingly, the stationary distribution of random walks on G according to $\mathbf{P}$ is $\pi = [\pi_i]^{1\times n}$ where

$$\pi_i = (\sum_{z=1}^{n} s_{zi})(\sum_{j=1}^{n}\sum_{k=1}^{n} s_{jk})^{-1}$$

.

4.5 Mixed random walks

It is intuitive and natural to consider a hybrid design that constitutes more than one of the aforementioned random walks. In particular, the interaction and similarity-biased models "rank" different nodes differently and "locally" assign weights to them. Though this limits the mixing time of social graphs, as we will see later, it does not provide nodes any authority on the random walk once they are a "past state". On the other hand, benefits of these models are shortcomings in other models. It's hence technically promising and intuitively sound to consider combinations of these designs. Another potential of a mixed design is to use both the lazy and originator-biased random walk in a single walk. As we will see later, in some rapidly mixing social graphs where the underlying social trust is hypothesized to be weak, the lazy random walk poorly captures the behavior of the Sybil defense.

5. RESULTS AND DISCUSSION

In this section we outline the results of this study. We first measure the mixing time of the social graphs used in this study (in Table 1) and highlight its variable nature among networks with similar size. We follow this by examining the impact of using proposed models on the mixing time and the performance of SybilLimit, a well-known Sybil defense. We limit ourselves to this defense scheme though our conclusions apply to all other schemes that using the mixing time as the underlying property for their performance.

5.1 Social graphs: the data sets

The social graphs used in our experiments are in Table 1. These graphs are carefully selected to feature different models of knowledge between nodes in the social networks. These networks are categorized as follows. (1) Social networks that exhibit knowledge between nodes and are good for the trust assumptions of the Sybil defenses—e.g., physics coauthorships and DBLP. These are slow mixing (see Figure 3).

Table 1: Datasets, their size and their second largest eigenvalues of the transition matrix. Physics 1, 2, 3 are relativity, high energy and high-energy theory coauthorship respectively.

Social network	Nodes	Edges	SLEM
Physics 1 [48]	4,158	13,428	0.998133
Slashdot [49]	82,168	582,533	0.987531
Physics 2 [48]	11,204	117,649	0.998221
Physics 3 [48]	8,638	24,827	0.996879
Wiki-vote [36]	7,066	100,736	0.899418
Enron [48]	33,696	180,811	0.996473
Epinion [50]	75,879	13,428	0.998133
DBLP [51]	614,981	1,155,148	0.997494
Facebook A [35]	1,000,000	20,353,734	0.982477
Facebook B [35]	1,000,000	15,807,563	0.992020
Livejournal A [42]	1,000,000	26,151,771	0.999387
Livejournal B [42]	1,000,000	27,562,349	0.999695
Youtube [42]	1,134,890	2,987,624	0.997972

(2) Graphs of networks that may not require face-to-face knowledge but require interaction—e.g., Youtube and Livejournal, which are slow mixing, but faster than the first category. (3) Datasets that may not require prior knowledge between nodes or where the social links between nodes are less meaningful to the context of the Sybil defenses—e.g., Facebook and wiki-vote, which are shown to be very fast mixing.

While these graphs are used for demonstrating the first part of the results, measuring the performance of SybilLimit and the impact of our designs on the mixing time is done over samples of these graphs. For feasibility reasons, we sample only 10K nodes, using the breadth-first search algorithm, from each graph larger than 10K in Table 1. The resulting sub-graphs are in Table 2. The diameter is the maximal eccentricity (set of maximal shortest paths from each source in the graph) and the radius is the minimal eccentricity. We compute them to show some insight on the structure of the graphs. For Facebook and Livejoural datasets, the sub-graphs are from dataset A of each.

5.2 Measuring the mixing time

While measuring the mixing time using SLEM as explained in section 3 requires computing μ, the computed mixing time might be an overestimation for quality which is necessary in the Sybil defenses. In principle, the overestimation occurs because the computed mixing time using SLEM is the maximal, where a few outlier nodes may capture the mixing time of the entire graph, while the majority of nodes may have relatively smaller mixing time than these outliers [15]. For that, we limit ourselves to measuring the mixing time using Definition 1, and considering a few initial distributions. We classify graphs, shown in Table 1, based on their size into large ($> 600,000$ nodes) and small ($< 100,000$ nodes) graphs.

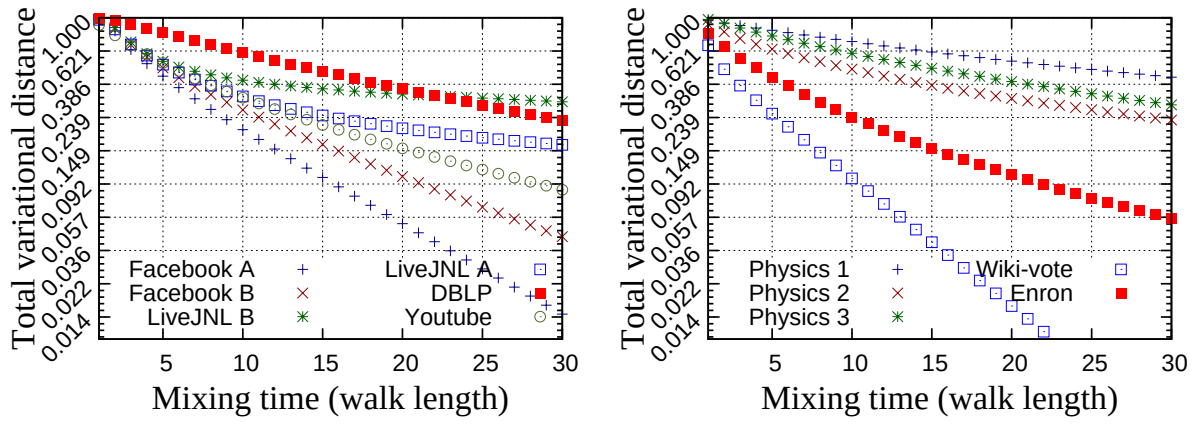

(a) Large datasets

(b) Small to meduim datasets

Figure 3: The average mixing time of a sample of 1000 initial distributions in each graph in Table 1 using the sampling method for computing the mixing time by its definition over P.

Table 2: Social graphs with their size, diameter, and radius. Physics 1, 2, 3 are relativity, high energy and high energy theory co-authorship respectively. Dim stands for Diameter and Rid stands for Radius

Social network	Nodes	Edges	Dim	Rad
Physics 1 [48]	4,158	13,428	17	9
Sdot [49]	10,000	14,6469	6	3
Physics 2 [48]	11,204	117,649	13	7
Physics 3 [48]	8,638	24,827	18	10
Wiki-vote [36]	7,066	100,736	7	4
Enron [48]	10,000	108,373	4	2
Epinion [50]	10,000	210,173	4	2
DBLP [51]	10,000	20,684	8	4
Facebook [35]	10,000	81,460	4	2
Livejournal [42]	10,000	135,633	6	3
Youtube [42]	10,000	58,362	4	2
Rice-cs-grad [52]	501	3255	9	5
Rice-cs-ugrad [52]	1221	43153	6	3

For each social graph, we compute the mixing time according to Definition 1 for a sample of $1,000$ initial distributions (nodes). We then compute the total variation distance for a given walk length w as the *average* distance among the $1,000$ nodes. The results are shown in Figure 3. In short, two things to observe from these measurements [15]. First, the mixing time is larger than used in literature (e.g., 10 to 15 in [8, 9] for 10^6-node graphs). For example, for $\epsilon \approx 1/4$, which is required for $\approx 99\%$ admission rate in SybilLimit, $w = 30$ is required in Physics 1. Second, we observe that the mixing time is variable among social graphs with similar size where graphs with meaningful edges are slower mixing than others with less meaningful links.

5.3 Implication on the mixing time

Along with the simple random walk-based design, we implement three of the proposed designs: lazy, originator, and similarity biased random walks. We use the simple random walk-based implementation over the interaction graph of Wilson et al.'s [35] to learn the performance of the interaction-based model. We examine the impact of each design on the mixing time on some graphs from Table 2. The results are shown in Figure 4 and Figure 5. We observe that, while they bound the mixing time of the different social graphs, the originator-biased random walk is too sensitive even to a small α. For example, as in 5(a) for Facebook social graph in Table 2, $\epsilon \approx 1/4$ is realizable at $w = 6$ with the simple random walk, $w > 10$ for both lazy and originator-biased random walk. However, this happens with $\alpha = 0.5$ in the lazy against $\alpha \approx 0.1$ in the originator-biased walk. This observation is made clearer on Figure 5 which compares the mixing time of four different social graphs with different characteristics when using the simple and modified random walks.

We also observe in Figure 4 and Figure 5 that the linear increments in the parameters do not necessarily have linear effect on the measured mixing time. Furthermore, this behavior is made clearer in the experiments performed on SybilLimit and shown in Figure 6 and Figure 7. This however is not surprising, at least with the originator-biased random walk, since the probability of intersection when sampling from the stationary distribution is $\leq 1 - e^{-8(1-\alpha)^4}$ (the proof is in Theorem 8 in the appendix) from which one can see the exponential effect of α on the admission rate.

While this explains the general tendency in the admission rates of SybilLimit, it does not answer some inconsistency shown in Fig. 7(b) for the transition between $\alpha = 0.12, 0.16$, and 0.20. One additional explanation for that is the community structure in this graph, which is shown in [12] to be clear in Physics 1 and problematic for Sybil defenses (results for the same graph are in Fig. 6(b) and Fig. 7(b)). We believe that the big jumps in these measurements happen at some values of these parameters, which make an entire community unreachable from another community. On the other hand, some graphs are less sensitive to the same value of

these parameters, e.g., Facebook with the results shown in figures 4(a), 5(a), 6(d), and 7(d). One possible explanation for this behavior is that this graph has less community structure. This claim is further supported by observing the small diameter and radius (shown in Table 2) and the high density of the graph, along with supporting clustering coefficient of 0.188. Reasoning about this behavior and its quantification is to be our future work.

5.4 Performance over simple random walks

To understand the necessary mixing time quality required for the operation of SybilLimit, we measure the performance of SybilLimit using simple random walks, where the evaluation metric is the percent of honest nodes accepted by other honest nodes. We do not consider the accepted Sybil nodes because they are bounded per attack edge. For each walk with length $w(0 \leq w \leq 30)$, we compute the number of accepted nodes as a percent out of $n(n-1)$—total verifier/suspect pairs. Since SybilLimit accepts nodes on edges only, it does not work for $w < 2$. The results are shown in Figure 8 and the variable mixing time shown earlier is further highlighted by observing the percent of accepted nodes when varying w. We observe that, unlike claims in [53] where one should expect 95% admission rate at $w = 4$, some graphs require $w = 30$ to achieve that. It is also worth noting that graphs which admit high percent of nodes for small w are those with poor trust.

5.5 Performance over modified random walks

Now we study the impact of the modified random walks on the performance of SybilLimit. We select four datasets with different characteristics from Table 2: DBLP, Facebook, Facebook (Rice grad), and Physics 1 (relativity theory). We implement modified SybilLimit versions that consider changes introduced by the modified random walks and test the admission rate of honest nodes under different values of α and w.

5.5.1 Performance over lazy random walk

we measure the performance of SybilLimit operating with the lazy random walks – results are shown in Figure 6. We vary w from 0 to 30 with steps of 2. We further vary α associated with the lazy random walk from 0 to 0.80 with steps of 0.16—$\alpha = 0$ means simple random walk. While the performance of SybilLimit is generally degraded when increasing α, we observe that the amount of degradation varies and depends on the initial quality of the graph. For example, by comparing DBLP (6(c)) to Facebook (6(d)) we observe that for $w = 10$, DBLP and Facebook admit about 97% and 100% of the honest nodes respectively for $\alpha = 0$. For the same w and $\alpha = 0.64$, the accepted nodes in Facebook are still close to 100% while the accepted nodes in DBLP are only 50% suggesting variable sensitivity of different graphs to the same α. Once we increase α to 0.80, the number of accepted nodes in Facebook decreases to 80% while giving only 25% in DBLP. One explanation of this behavior is what we have discussed in section 5.3. Also, since the ultimiate goal of this model is to characterize trust, which already differs in these graphs, we know that α should not necessarily be equal in both cases. For instance, if one is concerned about achieving same admission rate for the same w in both cases, one may choose $\alpha = 0.48$ in DBLP and $\alpha = 0.80$ in Facebook where $w = 10$ in both cases which yields 80% admission rate in both cases.

5.5.2 Performance over originator-biased random walk

The same settings in section 5.5.1 are used in this experiment but here we vary α from 0 to 0.2 with 0.02 steps since the originator-biased walk is more sensitive to smaller α than the lazy-random walk. Similar to the lazy walk, the originator-biased walk, as shown in Figure 7, influences the performance of SybilLimit on different graphs differently, and depending on the underlying graph. However, two differences are specific to the originator-biased walk over the lazy random walk.

First, the insensitivity shown earlier is even clearer in the originator-biased model. Second, while the end result of SybilLimit operating with lazy random walk is identical to the simple random walk if one allows long enough walk to compensate for the laziness, the behavior of the originator-biased walk is different. The indirect implication of the originator assigned probability to herself is "discontinuity" in the graph (with respect to each node), where each node gives up some of the network by not trusting nodes in it. To cover the whole graph with that the same α, w needs to be exponentially large. To challenge the insensitivity of the fast mixing social graphs, we extend α beyond the values used in Figure 7 with Facebook from Table 2 and use $\alpha(0 \leq \alpha \leq 0.5)$ with 0.1 steps and compute the admission rate. The result is shown in Figure 9. We observe that the originator-biased walk limits the number of accepted nodes, even in fast mixing graphs.

5.5.3 Similarity and interaction-biased random walks

The similarity and interaction-biased random walks as used in this paper are unparameterized. We compute the similarity for Facebook in Table 2, as explained in 4.4. The similarity is then used to assign weights to edges between nodes, and bias the transition matrix. We run SybilLimit with similarity-biased random walks on Facebook in Table 2, where the result is shown in Figure 10. In short, the similarity – while expected to capture some truth about the underlying graph – has less influence on the behavior of SybilLimit. It is worth noting that the impact of the similarity-biased random walk is clearer on other social graphs, such as DBLP and Physics, which have strong community structures.

For the interaction-biased design, we borrow the interaction graph of Wilson et al. [35] on Facebook (same dataset in Table 2). The interaction model introduces a richer model than the mere connections between nodes: it shows how

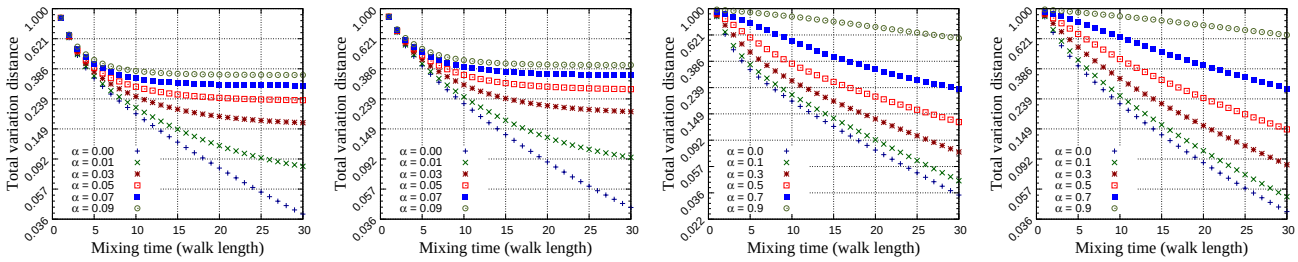

(a) Facebook A (b) Livejournal A (c) Facebook A (d) Livejournal A

Figure 4: The impact of the originator and lazy walks on the mixing time—(a) and (b) are for originator-biased while (c) and (d) are for lazy random walks.

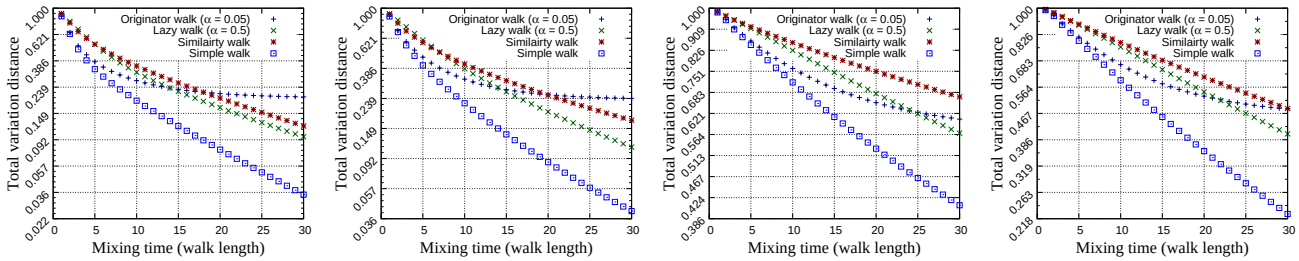

(a) Facebook A (b) Livejournal A (c) Physics 1 (d) DBLP

Figure 5: The mixing time of four different social graphs when using simple vs. lazy, originator, and similarity-biased random walks, for each graph. While they are similar in size, a mixing time (parameterized by the same ϵ) is variable.

strong are the links between nodes in the graph. With the same settings as earlier, we run SybilLimit – as a simple random walks – over the interaction graph. The results are shown in Figure 10.

5.6 All designs: comparative study

Finally, we consider all designs at the same time. Because we only have interaction measurements for the Facebook dataset, we limit ourselves to that dataset. The result is shown in Figure 10. While the performance of the similarity-biased random walk produces *almost* same results as the simple random-walk, the interaction-biased walk affects the number of the accepted nodes. Furthermore, the lazy random walk captures the behavior of model when deviated from the simple random-walk. As shown for this dataset, the interaction model behavior is characterized by the behavior of the lazy random walk for two given parameters ($\alpha = 0.48$ and $\alpha = 0.64$) suggesting that the interaction model can be further modeled as a lazy random walk where the problem is to find the proper parameters to match its behavior. Note that the value of α works for this dataset in particular. However, other datasets may be characterized by other values. This observation is made clearer on Figure 11 where the same results in Figure 10 are smoothened using the cubic spline interpolation.

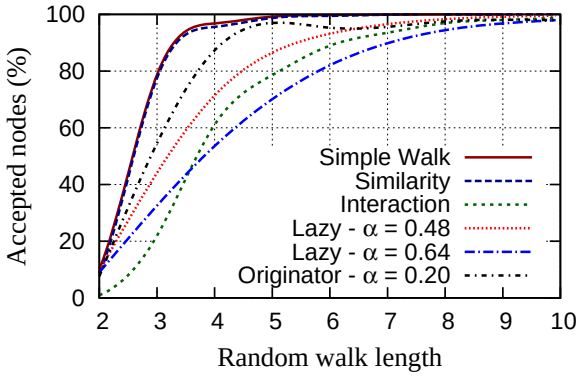

Figure 11: The measurements in Figure 10 smoothened using the cubic spline interpolation to show that the different designs can be made equivalent for some set of parameters. The interaction model can be further considered as an optimization problem in terms of other designs, such as the lazy and originator-biased designs

6. IMPLICATIONS OF FINDINGS

To sum up, we find in this study that one can control the behavior of the social network-based Sybil defenses by incorporating parameters for trust. For this purpose, we introduced and experimented the behavior of four designs. In graphs that are empirically-proven to be fast mixing and well-performing for the utility of the Sybil defense – though having poor value of trust – we have shown that one can select the necessary parameters to account for trust and make the performance of the defense on that graph equivalent to stronger and richer version of the same graph – e.g., the case of the interaction-based model versus the mere connections on the Facebook dataset. With these designs being intuitive in characterizing trust, the results being in agreement one another, and with this paper being the first of its own type in this direction, we believe that this study is a first step in the

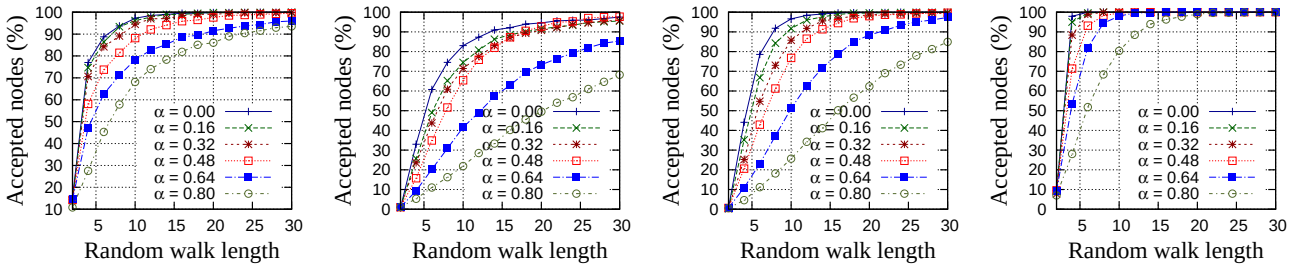

(a) Rice CS Grad (b) Physics 1 (c) DBLP (d) Facebook

Figure 6: The performance of SybilLimit measured for accepted honest nodes when using different lengths of lazy random walk for different social graphs.

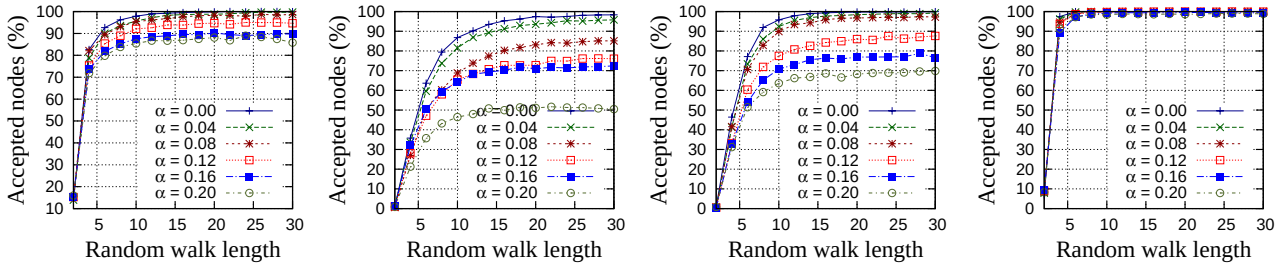

(a) Rice CS Grad (b) Physics 1 (c) DBLP (d) Facebook

Figure 7: The performance of SybilLimit depends on the underlying social graph, where different graphs require different walk lengths to ensure the same number of accepted nodes. The originator-biased random walk can further influence the number of nodes accepted in each graph.

direction of bringing well-received theoretical results into practice. The implications of our findings can be summarized as follows.

First, the mixing time and utility of the Sybil defense depends on the underlying graph. Through measurements, we supported our hypothesis that the quality of the social graph depends on the characteristic of the social links between the nodes. On one hand, social links that are easier to make result in well enmeshed graphs but are bad in principle for the Sybil defense since they already tolerate bad edges. However, these are shown to provide good honest nodes acceptance rate even with shorter random walks. On the other hand, social links that are harder to make result in graphs with more community structure, which are bad for the detection (as shown in [12]) and require longer walks to operate for the honest nodes.

Second, it is now possible for the Sybil defense operator, when given multiple options of social graphs, to further derive the utility of the Sybil defense using several criteria. Our study empowers the operators by an additional dimension that influences the behavior of the Sybil defense: trust.

Third, our findings answer a recently called for question in [12] of studying the behavior of Sbyil defenses when operated on the interaction-based model rather than the mere social connections, which are sometimes less meaningful. In short, our study shows that the interaction model can influence the behavior of the Sybil defense, by requiring longer random walk for the defense to work for honest nodes. However, this finding also suggests that a more community-structure is in the interaction model than in the mere social graph. This implies that, while the original social graph does not possess clear community structure, the use of the interaction model would add sensitivity for the detection part of the defense and result in weaker detection. However, the underlying graphs in both cases are different and the interpretation of the results should also consider the trust value in the interaction model, which is a better fit to the trust required in the Sybil defense.

Finally, online social graphs are known to possess weaker value of trust [45]. However, their potential for being used for Sybil defenses is very high since alternatives are limited, too expensive, and may not fit into the Sybil defense settings. For example, co-authorship social graphs, which are known for their trust value, may not necessarily include most users of a particular online system that tries to deploy the Sybil defense. On the other hand, given the popularity of online social networks, Sybil defenses may benefit from them, across systems and networks. To this end, the main finding of the paper is to open the door wide open for investigating trust, its modeling, and quantification for these systems.

7. OPEN QUESTIONS AND FUTURE WORK

In all of the designs that accept parameters, which are proposed in this paper, we have considered some assumptions to

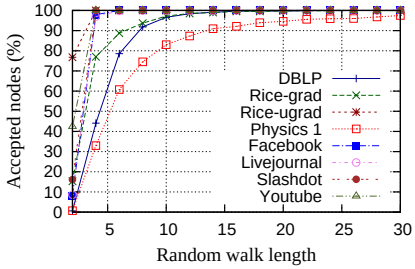

Figure 8: Accepted honest nodes in SybilLimit versus walk length, with simple random walk. Graphs have different quality of algorithmic property though having same size.

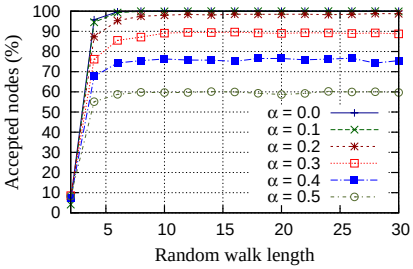

Figure 9: The originator-biased random walk limits the number of accepted nodes in very well-enmeshed social graphs with poor trust. The case of large α on Facebook.

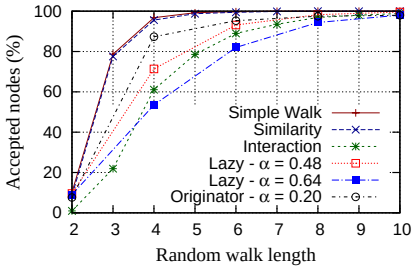

Figure 10: Accepted honest nodes in SybilLimit versus walk length, when using different designs to model trust in the social graph. The social graph of Facebook in Table 2.

simplify the operation of the designs. For example, we considered that each design uses a globally fixed parameter (α) that is used by each node in the graph. While this simplifies the analysis and make it possible to derive a clean formula for the transition probability and the bounding (or stationary) distribution of the random walk of the graph, it is natural to consider different random walks with mixed values for the same parameter – assigned locally by each node depending on its perception of the overall graph and trust in it. Providing clean formulation for this node-wise (as opposed to graph-wise) parameters stay an open question.

Through the experiments with the interaction and similarity biased random walks, we concerned ourselves by the percent of nodes accepted rather than "which nodes" are accepted. The later part of knowing which nodes are accepted is particularly interesting and can be used further to explore what additional attack strategies an attacker can use to be ranked higher than other nodes (see [12] for insights on such attacks). In the near future we will investigate the characteristics of the similarity and interaction graphs and explore how different are they from the original social graphs.

The main idea being exploited in the Sybil defenses, a long with the algorithmic property, is trust. It is assumed that the attack edges are limited. However, recent studies [54] have shown that it is easy to penetrate into the social network by introducing attack edges, in several settings. While the attack in [54] is based on automated identity theft, which has its own shortcomings in the context of Sybil defenses, it would be interesting to study that attack using entirely fake identities, on different social networks, and study the properties of the attacker network, which is our future work.

In the original SybilLimit, the escape probability is defined as the probability of ending up from the honest region into the Sybil region. As we assign probabilities to more familiar nodes (from the point of view of a particular node), it is intuitive that the escaping probability is going to be reduced. However, quantifying the difference remains an open question that we would like to investigate in the future. In fact, some these, e.g., SybilLimit, provide mechanisms to bound the number of accepted nodes by escaping walks.

Several other directions are worth investigation in the near future. First, we would like to investigate generalized node-wise parameterized designs that consider different parameters for different users, or categories of them. Second, we would like to theoretically formulate the behavior of the different designs considering other features of the underlying graph, e.g., its eigenvalues, mixing time, etc. Finally, we would like to investigate the applicability of these designs in other contexts where the trust of social networks is used. Also, we will investigate other intuitive approaches to characterize trust, use other extra information along with these designs to reduce the number of accepted Sybil nodes per attack edge. For example, one of the promising directions that we will investigate is to use landmarks (e.g., list of good nodes or bad nodes) that one can use to infer the goodness of a set of nodes that are close from these landmarks. Another direction is to use these list of previously known good nodes for teleportation and see how this is going to affect the "mixing time" of the social graph and the overall performance of the defense built on top of it.

Acknowledgement

We are grateful to Aaram Yun and Eugene Vasserman for their feedback on earlier versions of this work, and Alan Mislove, Ben Y. Zhao, and Bimal Viswanath for providing us with the social graphs used in this work. This work was supported by the NSF under grant CNS-0917154 and a research grant from Korea Advanced Institute of Science and Technology (KAIST).

8. REFERENCES

[1] J. Douceur and J. S. Donath, "The sybil attack," in *IPTPS*, 2002, pp. 251–260.

[2] B. Levine, C. Shields, and N. Margolin, "A survey of solutions to the sybil attack," University of Massachusetts Amherst, Tech. Rep., 2006.

[3] M. Castro, P. Druschel, A. J. Ganesh, A. I. T. Rowstron, and D. S. Wallach, "Secure routing for structured peer-to-peer overlay networks," in *OSDI*, 2002.

[4] A. Adya, W. J. Bolosky, M. Castro, G. Cermak, R. Chaiken, J. R. Douceur, J. Howell, J. R. Lorch, M. Theimer, and R. Wattenhofer, "Farsite: Federated, available, and reliable storage for an incompletely trusted environment," in *OSDI*, 2002.

[5] J. Ledlie and M. I. Seltzer, "Distributed, secure load balancing with skew, heterogeneity and churn," in *INFOCOM*, 2005, pp. 1419–1430.

[6] F. Lesueur, L. Mé, and V. V. T. Tong, "An efficient distributed pki for structured p2p networks," in *Proceeding of P2P*. IEEE, 2009, pp. 1–10.

[7] N. Borisov, "Computational puzzles as sybil defenses," in *Peer-to-Peer Computing*, A. Montresor, A. Wierzbicki, and N. Shahmehri, Eds. IEEE Computer Society, 2006, pp. 171–176.

[8] H. Yu, M. Kaminsky, P. B. Gibbons, and A. D. Flaxman, "Sybilguard: defending against sybil attacks via social networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 3, pp. 576–589, 2008.

[9] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao, "Sybillimit: A near-optimal social network defense against sybil attacks," in *IEEE Symposium on Security and Privacy*, 2008, pp. 3–17.

[10] N. Tran, J. Li, L. Subramanian, and S. S. M. Chow, "Brief announcement: improving social-network-based sybil-resilient node admission control," in *PODC*, 2010, pp. 241–242.

[11] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman, "Sybilguard: defending against sybil attacks via social networks," in *SIGCOMM*, 2006, pp. 267–278.

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

[13] G. Danezis and P. Mittal, "Sybilinfer: Detecting sybil nodes using social networks," in *NDSS*. The Internet Society, 2009.

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

[15] A. Mohaisen, A. Yun, and Y. Kim, "Measuring the mixing time of social graphs," UMN, Tech. Rep., 2010.

[16] Q. Li, S. Zhu, and G. Cao, "Routing in socially selfish delay tolerant networks," in *INFOCOM'10: Proceedings of the 29th conference on Information communications*. Piscataway, NJ, USA: IEEE Press, 2010, pp. 857–865.

[17] G. Bigwood and T. Henderson, "Social dtn routing," in *CoNEXT '08: Proceedings of the 2008 ACM CoNEXT Conference*. New York, NY, USA: ACM, 2008, pp. 1–2.

[18] J. Davitz, J. Yu, S. Basu, D. Gutelius, and A. Harris, "ilink: search and routing in social networks," in *KDD '07: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2007, pp. 931–940.

[19] E. M. Daly and M. Haahr, "Social network analysis for routing in disconnected delay-tolerant manets," in *MobiHoc '07: Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*. New York, NY, USA: ACM, 2007, pp. 32–40.

[20] G. Danezis, C. Lesniewski-laas, M. F. Kaashoek, and R. Anderson, "Sybil-resistant dht routing," in *In ESORICS*. Springer, 2005, pp. 305–318.

[21] S. Marti, P. Ganesan, and H. Garcia-Molina, "Dht routing using social links," in *IPTPS*, ser. Lecture Notes in Computer Science, G. M. Voelker and S. Shenker, Eds., vol. 3279. Springer, 2004, pp. 100–111.

[22] F. Lesueur, L. Mé, and V. V. T. Tong, "A sybil-resistant admission control coupling sybilguard with distributed certification," in *WETICE*, 2008, pp. 105–110.

[23] Z. Zhu, G. Cao, S. Zhu, S. Ranjan, and A. Nucci, "A social network based patching scheme for worm containment in cellular networks," in *INFOCOM*. IEEE, 2009, pp. 1476–1484.

[24] A. Beach, M. Gartrell, S. Akkala, J. Elston, J. Kelley, K. Nishimoto, B. Ray, S. Razgulin, K. Sundaresan, B. Surendar, M. Terada, and R. Han, "Whozthat? evolving an ecosystem for context-aware mobile social networks," *IEEE Network*, vol. 22, no. 4, pp. 50–55, 2008.

[25] S. Xu, X. Li, and T. P. Parker, "Exploiting social networks for threshold signing: attack-resilience vs. availability," in *ASIACCS*. ACM, 2008, pp. 325–336.

[26] R. J. Clark, E. Zasoski, J. Olson, M. H. Ammar, and E. W. Zegura, "D-book: a mobile social networking application for delay tolerant networks," in *Challenged Networks*, 2008, pp. 113–116.

[27] S. Pai, T. Roosta, S. B. Wicker, and S. Sastry, "Using social network theory towards development of wireless ad hoc network trust," in *AINA Workshops (1)*. IEEE Computer Society, 2007, pp. 443–450.

[28] H. A. Kautz, B. Selman, and M. A. Shah, "Referral web: Combining social networks and collaborative filtering," *Commun. ACM*, vol. 40, no. 3, pp. 63–65, 1997.

[29] E. Vasserman, R. Jansen, J. Tyra, N. Hopper, and Y. Kim, "Membership-concealing overlay networks," in *Proceedings of the 16th ACM conference on Computer and communications security*. ACM, 2009, pp. 390–399.

[30] E. Vasserman, "Towards freedom of speech on the Internet: Censorship-resistant communication and storage," Ph.D. dissertation, UNIVERSITY OF MINNESOTA, 2010.

[31] 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*. Berkeley, CA, USA: USENIX Association, 2009, pp. 15–28.

[32] E. M. Daly and M. Haahr, "Social network analysis for information flow in disconnected delay-tolerant manets," *IEEE Trans. Mob. Comput.*, vol. 8, no. 5, pp. 606–621, 2009.

[33] A. G. Miklas, K. K. Gollu, K. K. W. Chan, S. Saroiu, K. P. Gummadi, and E. De Lara, "Exploiting social interactions in mobile systems," in *UbiComp'07: Proceedings of the 9th international conference on Ubiquitous computing*. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 409–428.

[34] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi, "On the evolution of user interaction in facebook," in *Proceedings of the 2nd ACM SIGCOMM Workshop on Social Networks*, August 2009.

[35] 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*. New York, NY, USA: ACM, 2009, pp. 205–218.

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

[37] D. J. Crandall, D. Cosley, D. P. Huttenlocher, J. M. Kleinberg, and S. Suri, "Feedback effects between similarity and social influence in online communities," in *KDD*, Y. Li, B. Liu, and S. Sarawagi, Eds. ACM, 2008, pp. 160–168.

[38] E. Le Merrer and G. Trédan, "Centralities: capturing the fuzzy notion of importance in social graphs," in *SNS '09: Proceedings of the Second ACM EuroSys Workshop on Social Network Systems*. New York, NY, USA: ACM, 2009, pp. 33–38.

[39] D. Liben-Nowell and J. M. Kleinberg, "The link prediction problem for social networks," in *CIKM*. ACM, 2003, pp. 556–559.

[40] D. Guercia and S. Hailes, "Sybil attacks against mobile users: friends and foes to the rescue," in *INFOCOM 2010*, 2010.

[41] R. Landa, D. Griffin, R. Clegg, E. Mykoniati, and M. Rio, "A sybilproof indirect reciprocity mechanism for peer-to-peer networks," in *Proceedings of IEEE Infocom*, vol. 9, 2009.

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

[43] 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*, C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, Eds. ACM, 2007, pp. 835–844.

[44] J. Golbeck, "Trust and nuanced profile similarity in online social networks," *ACM Trans. Web*, vol. 3, no. 4, pp. 1–33, 2009.

[45] C. Dwyer, S. Hiltz, and K. Passerini, "Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace," in *Proceedings of AMCIS*, 2007.

[46] S. Nagaraja, "Anonymity in the wild: Mixes on unstructured networks," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, N. Borisov and P. Golle, Eds., vol. 4776. Springer, 2007, pp. 254–271.

[47] A. Sinclair, "Improved bounds for mixing rates of marcov chains and multicommodity flow," *Comb., Probability & Computing*, vol. 1, pp. 351–370, 1992.

[48] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graphs over time: densification laws, shrinking diameters and possible explanations," in *KDD '05: Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. New York, NY, USA: ACM, 2005, pp. 177–187.

[49] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, "Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters," *CoRR*, vol. abs/0810.1355, 2008.

[50] M. Richardson, R. Agrawal, and P. Domingos, "Trust management for the semantic web," in *International Semantic Web Conference*, ser. Lecture Notes in Computer Science, D. Fensel, K. P. Sycara, and J. Mylopoulos, Eds., vol. 2870. Springer, 2003, pp. 351–368.

[51] M. Ley, "The DBLP computer science bibliography: Evolution, research issues, perspectives," in *String Processing and Information Retrieval*. Springer, 2009, pp. 481–486.

[52] 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.

[53] H. Yu, C. Shi, M. Kaminsky, P. B. Gibbons, and F. Xiao, "Dsybil: Optimal sybil-resistance for recommendation systems," in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2009, pp. 283–298.

[54] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda, "All your contacts are belong to us: automated identity theft attacks on social networks," in *WWW '09: Proceedings of the 18th international conference on World wide web*. New York, NY, USA: ACM, 2009, pp. 551–560.

APPENDIX

Generalized random walk

Here we introduce a generalized random walk that characterizes the trust in social network as a combination of the different random walks presented in this paper (state diagram shown in Figure 12). In particular, given the transition matrix of the generalized random walk is given as $\mathbf{P} = [p_{ij}]n \times n$ where:

$$p_{ij} = \begin{cases} \alpha & v_j = v_i \\ \beta & v_j = v_r \\ \Pi_{ij} & v_j \in N(v_i) \setminus v_r \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

Π_{ij} in (10) combines the uniform (degree), similarity, and interaction measures presented earlier. Π_{ij} is computed as:

$$\Pi_{ij} = \Pi_{ij}^{\gamma} + \Pi_{ij}^{\eta} + \Pi_{ij}^{\mu}, \tag{11}$$

$$\Pi_{ij}^{\gamma} = \frac{\gamma}{\deg(v_i)}, \tag{12}$$

$$\Pi_{ij}^{\eta} = \frac{\eta s_{ij}}{\sum_y s_{iy}}, \tag{13}$$

$$\Pi_{ij}^{\nu} = \frac{\nu b_{ij}}{\sum_y b_{iy}}, \tag{14}$$

$$1 = \alpha + \beta + \gamma + \eta + \nu. \tag{15}$$

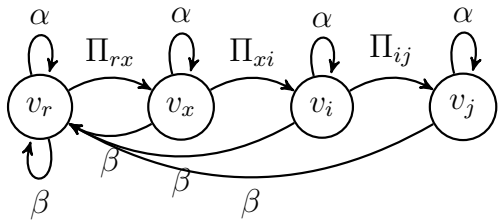

Figure 12: An illustration of the generalized random walk to characterized trust in social graphs. The parameters are defined in (11) through (15) and transition probability is defined in (10).

Proofs

THEOREM 3 (SIMPLE RANDOM WALK). *For the simple random walk with transition matrix* $\mathbf{P}$ *defined uniformly over edges,* $\pi = [\frac{\deg(v_1)}{2m} \dots \frac{\deg(v_n)}{2m}]$ *is a stationary distribution (state diagram shown in Figure 13).*

PROOF. *First, it is easy to show that*

$$\sum_i \pi_i = \sum_{v_i \in V} \frac{\deg(v_1)}{2m} = 1,$$

hence π *is a valid probability distribution. Then, let* $\mathbf{P}$ *be the transition matrix defined as* $\mathbf{P} = [p_{ij}]$ *where* $p_{ij} = \frac{1}{\deg(v_i)}$ *if* $v_j \in N(v_i)$ *or* 0 *otherwise. We know that in order for* π *to be a stationary distribution then* $\pi\mathbf{P} = \pi$ *must hold. Let* $\mathbf{x} = \pi\mathbf{P}$. *By applying that to* $\mathbf{P}$ *and* π *defined above, we compute* $\mathbf{x}$. *In particular, the* i^{th} *element in* $\mathbf{x}$ *is* $\cdots + \frac{deg(v_j)}{2m}\frac{1}{\deg(v_j)} + \cdots + \frac{deg(v_r)}{2m}\frac{1}{\deg(v_r)} + \dots$ *(*$\deg(v_i)$ *linear components)* $= \frac{\deg(v_i)}{2m}$ *for* $1 \leq i \leq n$. *In other words,* $\mathbf{x} = [\frac{\deg(v_i)}{2m}] = \pi = \pi\mathbf{P}$.

□

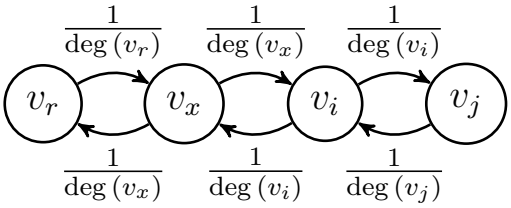

Figure 13: An illustration of the simple random walk. Other states incident to each of the states shown in this illustration are omitted for simplicity.

THEOREM 4 (LAZY RANDOM WALK). *For the lazy random walk with transition matrix* $\mathbf{P}$ *defined uniformly over edges, and self-loops at each node with probability* α *to maintain in the same state,* $\pi = [\frac{\deg(v_1)}{2m} \dots \frac{\deg(v_n)}{2m}]$ *is a stationary distribution. In other words, self loops do not bias the lazy random walk from the simple random walk (state diagram shown in Figure 14).*

PROOF. *First, it is easy to show that* $\sum_i \pi_i = 1$ *hence* π *is a valid probability distribution. Then, let* $\mathbf{P}$ *be the transition matrix defined as* $\mathbf{P} = [p_{ij}]$ *where* $p_{ij} = \frac{1-\alpha}{\deg(v_i)}$ *if* $v_j \in N(v_i)$, $p_{ij} = \alpha$ *if* $v_i = v_j$ *and* $p_{ij} = 0$ *otherwise. We know that in order for* π *to be a stationary distribution then* $\pi\mathbf{P} = \pi$ *must hold. Let* $\mathbf{x} = \pi\mathbf{P}$. *By applying that to* $\mathbf{P}$ *and* π *defined above, we compute* $\mathbf{x}$. *In particular, the* i^{th} *element in* $\mathbf{x}$ *is* $\alpha\frac{\deg(v_i)}{2m} + (\cdots + \frac{deg(v_j)}{2m}\frac{1-\alpha}{\deg(v_j)} + \cdots + \frac{deg(v_r)}{2m}\frac{1-\alpha}{\deg(v_r)} + \dots$ *(*$\deg(v_i)$ *linear components)* $= \alpha\frac{\deg(v_i)}{2m} + \sum_{r=1}^{\deg(v_i)}(\frac{1-\alpha}{2m}) = \alpha\frac{\deg(v_i)}{2m} + \deg(v_i)(\frac{1-\alpha}{2m}) = \frac{\deg(v_i)}{2m}$ *= for* $1 \leq i \leq n$. *In other words,* $\mathbf{x} = [\frac{\deg(v_i)}{2m}] = \pi = \pi\mathbf{P}$. □

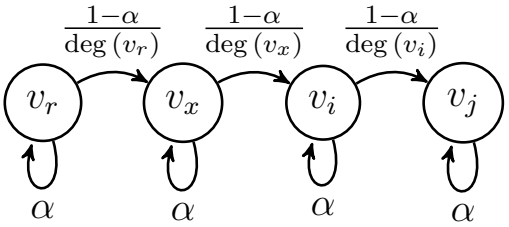

Figure 14: An illustration of the originator biased random walk. Other states incident to each of the states shown in this illustration are omitted for simplicity.

THEOREM 5 (ORIGINATOR-BIASED RANDOM WALK). *Let* α *and* β *be two numbers where* $0 \leq \alpha \leq 1$, $0 \leq \alpha \leq 1$,

and $\alpha+\beta=1$. Define a random walk on $G=(V,E)$ following $\mathbf{P}=[p_{ij}]$, where p_{ij} is defined as (state diagram is in Figure 15):

$$p_{ij}=\begin{cases}\alpha & j=r, v_r\notin N(v_i)\\ \alpha+\frac{1-\alpha}{\deg(v_i)} & j=r, v_r\in N(v_i)\\ \frac{1-\alpha}{\deg(v_i)} & j\neq r, v_j\in N(v_i)\\ 0 & \textit{otherwise}\end{cases} \tag{16}$$

The random walk initiated from v_r has a bounding distribution defined as $\pi=[p_{ij}]$ where

$$\pi_i=\begin{cases}\alpha+\frac{\beta\deg(v_i)}{2m}\approx\alpha & i=r\\ \frac{\beta\deg(v_i)}{2m} & i\neq r\end{cases} \tag{17}$$

PROOF. *First, we show that $\sum_i\pi_i=1$. We sum (17) for all i to get*

$$\alpha+\sum_{v_i\in V}\frac{\beta\deg(v_i)}{2m}=\alpha+\beta=1,$$

which implies that $\sum_i\pi_i=1$ and π is a valid probability distribution. To prove that π is a bounding distribution for the random walk initiated from the node v_r, we show that $\pi\mathbf{P}=\pi$ hold. Let $\mathbf{x}=\pi\mathbf{P}$. By applying that to $\mathbf{P}$ and π defined above, we compute $\mathbf{x}$. In particular, the i^{th} element in $\mathbf{x}$ is given considering two cases:

Case 1— $i=r$: this case includes three sub-cases. Let v_j be any node in V then v_j can be one of the following cases: $v_j\in N(v_r)$, $v_j\notin N(v_r)$, or $v_j=v_r$. Each of these cases sum as part of the probability assigned to v_i — i.e., to give x_i — when $i=r$. Let $x_i=x_{i1}+x_{i2}+x_{i3}$ where x_{i1},x_{i2},x_{i3} are the parts of x_i resulting from each case of the sub-cases independently. Each of these is computed as:

1. *$v_j\notin N(v_r)$*

$$x_{i1}=\sum_{v_j\in V\setminus N(v_r)}\alpha\frac{\beta\deg(v_j)}{2m}=\alpha\beta\sum_{v_j\in V\setminus N(v_r)}\frac{\deg(v_j)}{2m}\approx\alpha\beta \tag{18}$$

2. *$v_j\in N(v_r)$*

$$x_{i2}=\sum_{v_i\in N(v_r)}\frac{\beta\deg(v_j)}{2m}\frac{\beta}{\deg(v_j)}=\beta^2\frac{\deg(v_r)}{2m}\approx 0 \tag{19}$$

3. *$v_j=v_r$*

$$x_{i3}=\alpha(\alpha+\frac{\beta}{2m}\deg(v_r))\approx\alpha^2 \tag{20}$$

By summing up the three above cases, we get $x_i=x_{i1}+x_{i2}+x_{i3}=\alpha\beta+0+\alpha^2=\alpha-\alpha^2+\alpha^2=\alpha$, from which we conclude case 1.

Case 2 — $i\neq r$: similar to above, we compute $x_i=x_{i1}+x_{i2}$ where x_{i1} for π_i when $i=r$ and x_{i2} for the rest of cases. We compute x_{i1} as:

$$x_{i1}=\sum_{v_j\in N(v_i)}(\frac{\beta\deg(v_i)}{2m})(\frac{\beta}{\deg(v_i)})=\frac{\beta^2\deg(v_i)}{2m} \tag{21}$$

and x_{i2} as

$$\begin{aligned}x_{i2}&=\left(\alpha+\frac{\beta}{2m}\deg(v_r)\right)\left(\frac{\beta\deg(v_i)}{2m}\right)\\&=\frac{\alpha\beta\deg(v_i)}{2m}+\frac{\beta^2\deg(v_i)\deg(v_j)}{4m^2}\\&\approx\frac{\alpha\beta\deg(v_i)}{2m}=\frac{\beta\deg(v_i)}{2m}-\frac{\beta^2\deg(v_i)}{2m}\end{aligned} \tag{22}$$

By summing (21) and (22) we get $x_i=\frac{\beta\deg(v_i)}{2m}$. From case 1 and case 2, we conclude that $\mathbf{x}=\pi\mathbf{P}=\pi$, from which we conclude that $\pi=[\pi_i]$ defined in (17) is a bounding distribution for the walk on G following $\mathbf{P}$ in (16). □

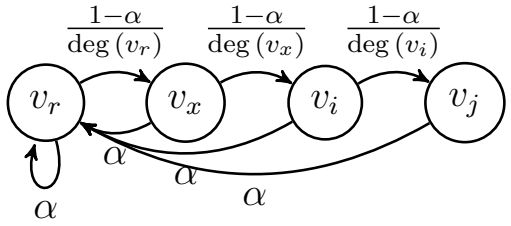

Figure 15: An illustration of the originator biased random walk.

THEOREM 6 (SIMILARITY-BIASED WALK). *For the similarity biased random walk with transition matrix $\mathbf{P}$ defined according the similarity measure over the links between the nodes as $\mathbf{P}=[p_{ij}]^{n\times n}$ where $p_{ij}=\frac{s_{ij}}{\sum_k s_{ik}}$ if $v_j\in N(v_i)$, $\pi=[\pi_i]$ where $\pi_i=\frac{\sum_j s_{ij}}{\sum_i\sum_k s_{ik}}$ is a stationary distribution (state diagram shown in Figure 16).*

PROOF. *First, it is easy to show that $\sum_i\pi_i=1$ hence π is a valid probability distribution. Then, similar to the previous walks above, let $\mathbf{P}$ and π be the transition matrix and the stationary distribution defined as above. We know that in order for π to be a stationary distribution then $\pi\mathbf{P}=\pi$ must hold. Let $\mathbf{x}=\pi\mathbf{P}$. By applying that to $\mathbf{P}$ and π defined above, we compute $\mathbf{x}$. In particular, the i^{th} element in $\mathbf{x}$ is $\frac{s_{ij}}{\sum_k s_{ik}}\frac{\sum_j s_{ij}}{\sum_i\sum_k s_{ik}}+\cdots+\frac{s_{ij}}{\sum_k s_{ik}}\frac{\sum_j s_{ij}}{\sum_i\sum_k s_{ik}}=\pi_i$ for $1\leq i\leq n$. Hence, $\pi=\mathbf{x}=\pi\mathbf{P}$.* □

THEOREM 7 (INTERACTION-BIASED WALK). *For the similarity-biased random walk with transition matrix $\mathbf{P}$ defined according the similarity measure over the links between the nodes as $\mathbf{P}=[p_{ij}]^{n\times n}$ where $p_{ij}=\frac{b_{ij}}{\sum_k b_{ik}}$ if $v_j\in N(v_i)$, $\pi=[\pi_i]$ where $\pi_i=\frac{\sum_j b_{ij}}{\sum_i\sum_k b_{ik}}$ is a stationary distribution (state diagram shown in Figure 17).*

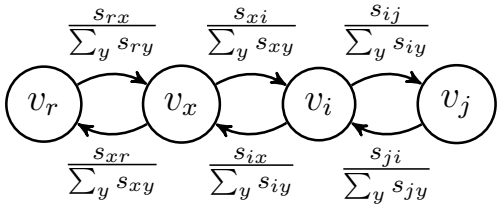

Figure 16: An illustration of the originator-biased random walk. Other states incident to each of the states shown in this illustration are omitted for simplicity.

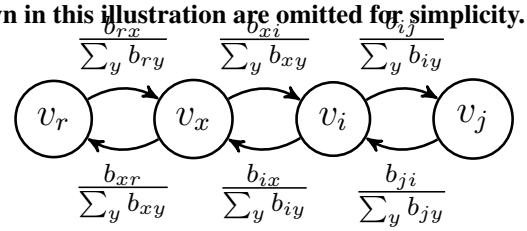

Figure 17: An illustration of the originator biased random walk.

PROOF. *Follows as in the proof of Theorem 6.* □

THEOREM 8. *Let ℓ be the effective length of the random walk, and $\alpha (0 \leq \alpha \leq 1)$ is a graph-wise parameter, then the number of random walks with length greater than 2 is $P_r(\ell \geq 2) = (1-\alpha)^2$. Furthermore, the ideal number of intersections in SybilLimit when using the originator-biased random walk is $\leq 1 - e^{-8(1-\alpha)^4}$*

PROOF. *The proof of this is straightforward by considering the possibilities of a random walk with effective length greater than 2. Let $A(i)$ be the event of having a random walk effective length $\ell = i$. We can write all the possible events of effective lengths greater than 2 as $A(2), A(3), \ldots, A(w)$. We know that $P_r(A(2)) = P_r(\ell = 2) = (1-\alpha)^2$, $P_r(A(3)) = P_r(\ell = 3) = (1-\alpha)^3$, and $P_r(A(w)) = P_r(\ell = w) = (1-\alpha)^w$. But since $A(w) \subseteq A(w-1) \cdots \subseteq A(3) \subseteq A(2)$, we know that $P_r(A(2))$ already includes the cases of the events $A(3) \ldots A(w)$. That is, $P_r(\ell \geq 2) \leq P_r(\ell = 2) = (1-\alpha)^2 = \beta^2$ where $\beta = 1-\alpha$. We know that random walk length needs to at least 2 in order for SybilLimit to work by intersecting on the tails. By plugging this into the number of collected samples (i.e., $4\sqrt{m}$), we get $4\beta^2\sqrt{m}$. By plugging this into the birthday paradox bound, we get the probability of intersection, P_r, from the effective expected number of random walks with length greater than 2 as follows:*

$$P_r = 1 - e^{-[(4\beta^2\sqrt{m})(4\beta^2\sqrt{m}-1)/2]/m} \quad (23)$$

$$= 1 - e^{(\frac{2\beta^2}{\sqrt{m}} - 8\beta^4)} \approx 1 - e^{-8\beta^4} = 1 - e^{-8(1-\alpha)^4} \quad (24)$$

which concludes the proof □