

# SybilLimit: A Near-Optimal Social Network Defense Against Sybil Attacks

Haifeng Yu, Phillip B. Gibbons, *Member, IEEE*, Michael Kaminsky, and Feng Xiao

**Abstract**—Open-access distributed systems such as peer-to-peer systems are particularly vulnerable to *sybil attacks*, where a malicious user creates multiple fake identities (called *sybil nodes*). Without a trusted central authority that can tie identities to real human beings, defending against sybil attacks is quite challenging. Among the small number of decentralized approaches, our recent SybilGuard protocol leverages a key insight on social networks to bound the number of sybil nodes accepted. Despite its promising direction, SybilGuard can allow a large number of sybil nodes to be accepted. Furthermore, SybilGuard assumes that social networks are fast-mixing, which has never been confirmed in the real world. This paper presents the novel SybilLimit protocol that leverages the same insight as SybilGuard, but offers dramatically improved and near-optimal guarantees. The number of sybil nodes accepted is reduced by a factor of  $\Theta(\sqrt{n})$ , or around 200 times in our experiments for a million-node system. We further prove that SybilLimit’s guarantee is at most a  $\log n$  factor away from optimal when considering approaches based on fast-mixing social networks. Finally, based on three large-scale real-world social networks, we provide the first evidence that real-world social networks are indeed fast-mixing. This validates the fundamental assumption behind SybilLimit’s and SybilGuard’s approach.

**Index Terms**—Social networks, sybil attack, sybil identities, SybilGuard, SybilLimit.

## I. INTRODUCTION

**S**YBIL attacks [1] refer to individual malicious users creating multiple fake identities (called *sybil identities* or *sybil nodes*) in *open-access* distributed systems (such as peer-to-peer systems). These open-access systems aim to provide service to any user who wants to use the service (instead of, for example, only to a predetermined group of 10 users). Sybil attacks have already been observed in the real world [2] in the Maze peer-to-peer system. Researchers have also demonstrated [3] that it is surprisingly easy to launch sybil attacks in the widely used eMule system [4].

When a malicious user’s sybil nodes comprise a large fraction of the nodes in the system, that one user is able to “outvote” the honest users in a wide variety of collaborative tasks. Examples

Manuscript received January 13, 2009; revised July 14, 2009; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor P. Rodriguez. First published November 06, 2009; current version published June 16, 2010. This work was supported in part by NUS Grants R-252-050-284-101 and R-252-050-284-133. A preliminary version of this paper appeared in the IEEE Symposium on Security and Privacy 2008.

H. Yu and F. Xiao are with School of Computing, National University of Singapore, Singapore 117543, Singapore (e-mail: haifeng@comp.nus.edu.sg; xiaof@comp.nus.edu.sg).

P. B. Gibbons and M. Kaminsky are with Intel Labs Pittsburgh, Pittsburgh, PA 15213 USA (e-mail: phillip.b.gibbons@intel.com; michael.e.kaminsky@intel.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2009.2034047

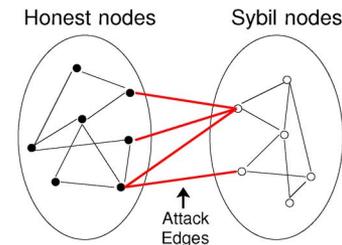


Fig. 1. The social network.

of such collaborative tasks range from Byzantine consensus [5] and voting schemes for e-mail spam [6] to implicit collaboration in redundant routing and data replication in distributed hash tables (DHTs) [7]–[9]. The exact form of such collaboration and the exact fraction of sybil nodes these collaborative tasks can tolerate may differ from case to case. However, a generic requirement for thwarting such attacks is that the number of sybil nodes (compared to the number of honest users) needs to be properly bounded.

Sybil attacks can be thwarted by a trusted central authority if the authority can tie identities to actual human beings, but implementing such a capability can be difficult or impossible, especially given the privacy concern of the users. Another approach is for the central authority to impose a monetary charge on each identity, which is, however, undesirable in many applications. Without these trusted central authorities, defending against sybil attacks is much harder. Among the small number of approaches, the simplest one perhaps is to bind identities to IP addresses or IP prefixes [10]. Another approach is to require every identity to solve puzzles that require human effort, such as CAPTCHAs [11]. Both approaches can provide only limited protection—the adversary can readily steal IP addresses with different prefixes in today’s Internet [12], while CAPTCHAs can be reposted on an adversary’s Web site to be solved by users seeking access to that site. In fact, Douceur’s initial paper on sybil attacks [1] already proved a negative result showing that sybil attacks cannot be prevented unless special assumptions are made.

### A. The SybilGuard Approach

Recently, we proposed SybilGuard [13], a new protocol for defending against sybil attacks without relying on a trusted central authority. SybilGuard leverages a key insight regarding *social networks* (Fig. 1). In a social network, the vertices (nodes) are identities in the distributed system and the (undirected) edges correspond to human-established trust relations in the real world. The edges connecting the honest region (i.e., the region containing all the honest nodes) and the sybil region (i.e., the region containing all the sybil identities created by

malicious users) are called *attack edges*. SybilGuard ensures that the number of attack edges is independent of the number of sybil identities and is limited by the number of trust relation pairs between malicious users and honest users. SybilGuard observes that if malicious users create too many sybil identities, the graph will have a small *quotient cut*—i.e., a small set of edges (the attack edges) whose removal disconnects a large number of nodes (all the sybil identities). On the other hand, “fast-mixing” [14] social networks do not tend to have such cuts. SybilGuard leverages the small quotient cut to limit the size of sybil attacks.

SybilGuard is a completely decentralized protocol and enables any honest node  $V$  (called the *verifier*) to decide whether or not to *accept* another node  $S$  (called the *suspect*). “Accepting” means that  $V$  is willing to do collaborative tasks with  $S$ . SybilGuard’s provable (probabilistic) guarantees hold for  $(1 - \epsilon)n$  verifiers out of the  $n$  honest nodes, where  $\epsilon$  is some small constant close to 0. (The remaining nodes get degraded, not provable, protection.) Assuming fast-mixing social networks and assuming the number of attack edges is  $o(\sqrt{n}/\log n)$ , SybilGuard guarantees that any such verifier, with probability of at least  $1 - \delta$  ( $\delta$  being a small constant close to 0), will accept at most  $O(\sqrt{n} \log n)$  sybil nodes per attack edge and at least  $(1 - \epsilon)n$  honest nodes.

While its direction is promising, SybilGuard suffers from two major limitations. First, although the end guarantees of SybilGuard are stronger than previous decentralized approaches, they are still rather weak in the absolute sense: Each attack edge allows  $O(\sqrt{n} \log n)$  sybil nodes to be accepted. In a million-node synthetic social network, the number of sybil nodes accepted per attack edge is nearly 2000 [13]. The situation can get worse: When the number of attack edges  $g = \Omega(\sqrt{n}/\log n)$  (or  $g > 15\,000$  in the million-node synthetic social network), SybilGuard can no longer bound the number of accepted sybil nodes *at all*. Second, SybilGuard critically relies on the assumption that social networks are fast-mixing, an assumption that had not been validated in the real world.

### B. SybilLimit: A Near-Optimal Protocol for Real-World Social Networks

In this paper, we present a new protocol that leverages the same insight as SybilGuard but offers dramatically improved and near-optimal guarantees. We call the protocol *SybilLimit* because: 1) it limits the number of sybil nodes accepted; and 2) it is near-optimal and thus pushes the approach to the limit. For any  $g = o(n/\log n)$ , SybilLimit can bound the number of accepted sybil nodes per attack edge within  $O(\log n)$  (see Table I). This is a  $\Theta(\sqrt{n})$  factor reduction from SybilGuard’s  $O(\sqrt{n} \log n)$  guarantee. In our experiments on the million-node synthetic social network used in [13], SybilLimit accepts on average around 10 sybil nodes per attack edge, yielding nearly 200 times improvement over SybilGuard. Putting it another way, with SybilLimit, the adversary needs to establish nearly 100 000 real-world social trust relations with honest users in order for the sybil nodes to outnumber honest nodes, as compared to 500 trust relations in SybilGuard. We further prove that SybilLimit is at most a  $\log n$  factor from optimal in the following sense: For any protocol based on the mixing time of a social network, there is a lower bound of  $\Omega(1)$  on

TABLE I  
NUMBER OF SYBIL NODES ACCEPTED PER ATTACK EDGE (OUT OF AN UNLIMITED NUMBER OF SYBIL NODES), BOTH ASYMPTOTICALLY FOR  $n$  HONEST NODES AND EXPERIMENTALLY FOR A MILLION HONEST NODES. SMALLER IS BETTER

Number of attack edges $g$ (unknown to protocol)	SybilGuard accepts	SybilLimit accepts
$o(\sqrt{n}/\log n)$	$O(\sqrt{n} \log n)$	$O(\log n)$
$\Omega(\sqrt{n}/\log n)$ to $o(n/\log n)$	unlimited	$O(\log n)$
below $\sim 15,000$	$\sim 2000$	$\sim 10$
above $\sim 15,000$ and below $\sim 100,000$	unlimited	$\sim 10$

the number of sybil nodes accepted per attack edge. Finally, SybilLimit continues to provide the same guarantee even when  $g$  grows to  $o(n/\log n)$ , while SybilGuard’s guarantee is voided once  $g = \Omega(\sqrt{n}/\log n)$ . Achieving these near-optimal improvements in SybilLimit is far from trivial and requires the combination of multiple novel techniques. SybilLimit achieves these improvements without compromising on other properties as compared to SybilGuard (e.g., guarantees on the fraction of honest nodes accepted).

Next, we consider whether real-world social networks are sufficiently fast-mixing for protocols like SybilGuard and SybilLimit. Even though some simple synthetic social network models [15] have been shown [16], [17] to be fast-mixing under specific parameters, whether real-world social networks are indeed fast-mixing is controversial [18]. In fact, social networks are well-known [19]–[22] to have groups or communities where intragroup edges are much denser than intergroup edges. Such characteristics, on the surface, could very well prevent fast-mixing. To resolve this question, we experiment with three large-scale (up to nearly a million nodes) real-world social network datasets crawled from [www.friendster.com](http://www.friendster.com), [www.livejournal.com](http://www.livejournal.com), and [dblp.uni-trier.de](http://dblp.uni-trier.de). We find that despite the existence of social communities, even social networks of such large scales tend to mix well within a rather small number of hops (10 to 20 hops), and SybilLimit is quite effective at defending against sybil attacks based on such networks. These results provide the first evidence that real-world social networks are indeed fast-mixing. As such, they validate the fundamental assumption behind the direction of leveraging social networks to limit sybil attacks.

## II. RELATED WORK

### A. Sybil Defenses Leveraging Social Networks

Since SybilGuard [13], [23] pioneered using social networks to defend against sybil attacks, there have been a number of research efforts (e.g., Ostra [24], SybilInfer [25], SumUp [26]) adopting such an approach. Same as SybilLimit, all these efforts leverage the fact that sybil attacks will result in small quotient cuts in the social network. Small quotient cuts, in turn, translate to a poor expansion property and a large mixing time of the graph.

If one is willing to assume global knowledge of the continuously changing social network (i.e., one node maintains an up-to-date copy of the entire social network graph), then simply

running an approximation algorithm [27] for minimal quotient cut will bound the number of sybil identities accepted within  $O(\log n)$  per attack edge, where  $n$  is the number of honest identities. Also assuming global knowledge and further focusing on scenarios where only  $o(n)$  honest identities are seeking to be accepted, SumUp [26] uses adaptive maximum flow on the social network to bound the number of sybil identities (voters) accepted per attack edge within  $1 + o(1)$ .

Similarly, the complete design of Ostra [24] and SybilInfer [25] also assume global knowledge about the social network. Even though both works [24], [25] allude to decentralized designs, none of them provides a complete design that is decentralized. Ostra does not provide guarantees that are provable. SybilInfer only proves that sybil nodes will increase the mixing time of the graph and thus affect the probability that a random walk starting from a region will end within that region. There is no result proven on *how much* the probability is affected. SybilInfer determines the probability via sampling, which by itself has unknown estimation error. As a result, SybilInfer is not able to prove an end-to-end guarantee on the number of sybil nodes accepted.

In contrast to all these above efforts, SybilLimit avoids the need for any global knowledge by using a decentralized secure random route technique. It provably bounds the number of sybil identities accepted per attack edge within  $O(\log n)$  while accepting nearly all honest nodes. The relationship between SybilGuard and SybilLimit is discussed in more detail in Sections IV and V-C.

Finally, orthogonal to SybilLimit’s goal of limiting the number of accepted sybil nodes, Ostra and SumUp further leverage feedback to modify the weight of the edges in the social network dynamically. Neither of these two feedback-based heuristics offers a provable guarantee. Our recent work on DSybil [28] also uses feedback to defend against sybil attacks in the context of recommendation systems and provides strong provable end-to-end guarantees. In scenarios where feedback is available, we expect that combining these feedback-based techniques with SybilLimit can further strengthen the defense.

### B. Other Sybil Defenses

Some researchers [29] proposed exploiting the *bootstrap tree* of DHTs. Here, the insight is that the large number of sybil nodes will all be introduced (directly or indirectly) into the DHT by a small number of malicious users. Bootstrap trees may appear similar to our approach, but they have the drawback that an honest user may also indirectly introduce a large number of other honest users. Such a possibility makes it difficult to distinguish malicious users from honest users. Instead of simply counting the number of nodes introduced directly and indirectly, SybilLimit distinguishes sybil nodes from honest nodes based on graph mixing time. It was shown [29] that the effectiveness of the bootstrap tree approach deteriorates as the adversary creates more and more sybil nodes, whereas SybilLimit’s guarantees hold no matter how many sybil nodes are created. Some researchers [30] assume that the attacker has only one or small number of network positions in the Internet. If such assumption holds, then all sybil nodes created by the attacker will have similar network coordinates [31]. Unfortunately, once the attacker

has more than a handful of network positions, the attacker can fabricate arbitrary network coordinates.

In reputation systems, colluding sybil nodes may artificially increase a (malicious) user’s rating (e.g., in Ebay). Some systems such as Credence [32] rely on a trusted central authority to prevent this. There are existing distributed defenses [33], [34] to prevent such artificial rating increases. These defenses, however, cannot bound the number of sybil nodes accepted, and in fact, all the sybil nodes can obtain the same rating as the malicious user. Sybil attacks and related problems have also been studied in sensor networks [35], [36], but the approaches and solutions usually rely on the unique properties of sensor networks (e.g., key predistribution). Margolin *et al.* [37] proposed using cash rewards to motivate one sybil node to reveal other sybil nodes, which is complimentary to bounding the number of sybil nodes accepted in the first place.

### C. Social Networks and Their Fast-Mixing Properties

Besides sybil defense, social networks have been used elsewhere (e.g., for digital library maintenance in LOCKSS [38]). Social networks are one type of trust networks. Unlike many other works [32]–[34] on trust networks, SybilLimit does not use trust propagation in the social network.

Concurrent with the preliminary version [39] of this paper, Leskovec *et al.* [40] independently validate the fast-mixing property of real-world social networks. Their study investigates over 70 large real-world social and information networks (from a few thousand nodes to over 10 million nodes). They show that in nearly every dataset, at nontrivial size scales ( $>100$  nodes), social communities gradually “blend in” more and more with the rest of the network and thus become less and less “community-like.” This in turn implies that, at a nontrivial scale, social networks are expander-like and will likely not contain small quotient cuts in the absence of sybil attacks.

## III. SYSTEM MODEL AND ATTACK MODEL

SybilLimit adopts a similar system model and attack model as SybilGuard [13]. The system has  $n$  honest human beings as *honest users*, each with one *honest identity/node*. Honest nodes obey the protocol. The system also has one or more malicious human beings as *malicious users*, each with one or more identities/nodes. To unify terminology, we call all identities created by malicious users as *sybil identities/nodes*. Sybil nodes are byzantine and may behave arbitrarily. All sybil nodes are colluding and are controlled by an *adversary*. A compromised honest node is completely controlled by the adversary and hence is considered as a sybil node and not as an honest node.

There is an undirected social network among all the nodes, where each undirected edge corresponds to a human-established trust relation in the real world. The adversary may create arbitrary edges among sybil nodes in the social network. Each honest user knows his/her neighbors in the social network, while the adversary has full knowledge of the entire social network. The honest nodes have  $m$  undirected edges among themselves in the social network. For expository purposes, we sometimes also consider the  $m$  undirected edges as  $2m$  directed edges. The adversary may eavesdrop on any messages sent in the protocol.

Every node is simultaneously a suspect and a verifier. As in SybilGuard, we assume that each suspect  $S$  has a locally generated public/private key pair, which serves to prevent the adversary from “stealing”  $S$ ’s identity after  $S$  is accepted. When a verifier  $V$  accepts a suspect  $S$ ,  $V$  actually accepts  $S$ ’s public key, which can be used later to authenticate  $S$ . We do not assume a public key infrastructure, and the protocol does not need to solve the public key distribution problem because the system is not concerned with binding public keys to human beings or computers. A malicious user may create multiple different key pairs for her different sybil nodes.

#### IV. BACKGROUND: SYBILGUARD

To better understand the improvements of SybilLimit over SybilGuard and the challenges involved, this section provides a concise review of SybilGuard.

##### A. Random Walks and Random Routes

SybilGuard uses a special kind of random walk, called *random routes*, in the social network. In a random walk, at each hop, the current node flips a coin on the fly to select a uniformly random edge to direct the walk (the walk is allowed to turn back). For random routes, each node uses a precomputed random permutation—“ $x_1x_2 \dots x_d$ ,” where  $d$  is the degree of the node—as a *one-to-one mapping* from incoming edges to outgoing edges. A random route entering via edge  $i$  will always exit via edge  $x_i$ . This precomputed permutation, or *routing table*, serves to introduce *external correlation* across multiple random routes. Namely, once two random routes traverse the same directed edge, they will merge and stay merged (i.e., they *converge*). Furthermore, the outgoing edge uniquely determines the incoming edge as well; thus the random routes can be *back-traced*. These two properties are key to SybilGuard’s guarantees. As a side effect, such routing tables also introduce *internal correlation* within a single random route. Namely, if a random route visits the same node more than once, the exiting edges will be correlated. We showed [13] that such correlation tends to be negligible, and moreover, in theory it can be removed entirely using a more complex design. Thus, we ignore internal correlation from now on.

Without internal correlation, the behavior of a single random route is exactly the same as a random walk. In connected and nonbipartite graphs, as the length of a random walk goes toward infinity, the distribution of the last node (or edge) traversed becomes independent of the starting node of the walk. Intuitively, this means when the walk is sufficiently long, it “forgets” where it started. This final distribution of the last node (or edge) traversed is called the node (or edge) *stationary distribution* [14] of the graph. The edge stationary distribution (of any graph) is always a uniform distribution, while the node stationary distribution may not be. *Mixing time* [14] describes how fast we approach the stationary distribution as the length of the walk increases. More precisely, mixing time is the walk length needed to achieve a certain *variation distance* [14],  $\Delta$ , to the stationary distribution. Variation distance is a value in  $[0,1]$  that describes the “distance” between two distributions—see [14] for the precise definition. A small variation distance means that the two distributions are similar. For a graph (family) with  $n$  nodes, we say that it is *fast-mixing* if its mixing time is  $O(\log n +$

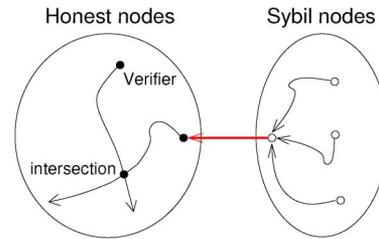


Fig. 2. Routes over the same edge merge.

$\log(1/\Delta)$ ). In this paper, we only care about  $\Delta = \Theta(1/n)$ , and we will simply say that a fast-mixing graph has  $O(\log n)$  mixing time. The following known result follows directly from the definition of mixing time and a useful interpretation of variation distance [41, Theorem 5.2]. This result is all we need in this paper about mixing time.

**Theorem 1:** Consider any fast-mixing graph with  $n$  nodes. A random walk of length  $\Theta(\log n)$  is sufficiently long such that, with probability of at least  $1 - (1/n)$ , the last node/edge traversed is drawn from the node/edge stationary distribution of the graph.

In SybilGuard, a random walk starting from an honest node in the social network is called *escaping* if it ever crosses any attack edge.

**Theorem 2:** (From [13]) In any connected social network with  $n$  nodes and  $g$  attack edges, the probability of a length- $l$  random walk starting from a uniformly random honest node being escaping is at most  $gl/n$ .

##### B. Accepting Honest Nodes

In SybilGuard, each node performs a random route of length  $l = \Theta(\sqrt{n} \log n)$ . A verifier  $V$  only accepts a suspect  $S$  if  $S$ ’s random route intersects with  $V$ ’s. Theorem 2 tells us that  $V$ ’s random route will stay in the honest region with probability of at least  $1 - gl/n = 1 - o(1)$  for  $g = o(\sqrt{n}/\log n)$ . Theorem 1 further implies that with high probability, a random route  $\Theta(\sqrt{n} \log n)$  long will include  $\Theta(\sqrt{n})$  independent random nodes drawn from the node stationary distribution. It then follows from the generalized Birthday Paradox [42] that an honest suspect  $S$  will have a random route that intersects with  $V$ ’s random route with probability  $1 - \delta$  for any given (small) constant  $\delta > 0$ .

##### C. Bounding the Number of Sybil Nodes Accepted

To intersect with  $V$ ’s non-escaping random route, a sybil suspect’s random route must traverse one of the attack edges. Consider Fig. 2, where there is only a single attack edge. Because of the convergence property of random routes, all the random routes from all sybil suspects must merge completely once they traverse the attack edge. All these routes differ only in how many hops of the route remain after crossing the attack edge (between 1 and  $l - 1$  hops for a length- $l$  route). Because the remaining parts of these routes are entirely in the honest region, they are controlled by honest nodes. Thus, there will be fewer than  $l = O(\sqrt{n} \log n)$  random routes that emerge from the sybil region. In general, the number of such routes will be  $O(g\sqrt{n} \log n)$  for  $g$  attack edges. SybilGuard is designed such that only one public key can be *registered* at the nodes on each

<p><b>Executed by each suspect <math>S</math>:</b></p> <ol style="list-style-type: none"> <li>1. <math>S</math> picks a uniformly random neighbor <math>Y</math>;</li> <li>2. <math>S</math> sends to <math>Y</math>: <math>\langle 1, S</math>'s public key <math>K_S, \text{MAC}(1  K_S)\rangle</math> with the MAC generated using the edge key between <math>S</math> and <math>Y</math>;</li> </ol> <p><b>Executed by each node <math>B</math> upon receiving a message <math>\langle i, K_S, \text{MAC}\rangle</math> from some neighbor <math>A</math>:</b></p> <ol style="list-style-type: none"> <li>1. discard the message if the MAC does not verify or <math>i &lt; 1</math> or <math>i &gt; w</math>;</li> <li>2. if <math>(i = w)</math></li> <li>3. record <math>K_S</math> under the edge name "<math>K_A \rightarrow K_B</math>" where <math>K_A</math> and <math>K_B</math> are <math>A</math>'s and <math>B</math>'s public key, respectively;</li> <li>4. else</li> <li>5. look up the routing table and determine to which neighbor (<math>C</math>) the random route should be directed;</li> <li>6. <math>B</math> sends to <math>C</math>: <math>\langle i + 1, K_S, \text{MAC}((i + 1)  K_S)\rangle</math> with the MAC generated using the edge key between <math>B</math> and <math>C</math>;</li> </ol>
---

Fig. 3. Protocol for suspects to do random routes and register their public keys.

random route. This means that the adversary can register only  $O(g\sqrt{n}\log n)$  public keys for all the sybil nodes combined. In order to accept a suspect  $S$ ,  $V$  must find an intersection between its random route and  $S$ 's random route and then confirm that  $S$  is properly registered at the intersecting node. As a result, only  $O(\sqrt{n}\log n)$  sybil nodes will be accepted per attack edge. For  $g = o(\sqrt{n}/\log n)$ , the total number of sybil nodes accepted is  $o(n)$ .

#### D. Estimating the Needed Length of Random Routes

While the length of the random routes is  $\Theta(\sqrt{n}\log n)$ , the value of  $n$  is unknown. In SybilGuard, nodes locally determine the needed length of the random routes via sampling. Each node is assumed to know a rough upper bound  $Z$  on the mixing time. To obtain a sample, a node  $A$  first performs a random walk of length  $Z$ , ending at some node  $B$ . Next,  $A$  and  $B$  each perform random routes to determine how long the routes need to be to intersect. A sample is *bad* (i.e., potentially influenced by the adversary) if any of the three random walks/routes in the process is escaping. Applying Theorem 2 shows that the probability of a sample being bad is at most  $3gl/n = o(1)$  for  $g = o(\sqrt{n}/\log n)$ .

### V. SYBILLIMIT PROTOCOL

As summarized in Table I, SybilGuard accepts  $O(\sqrt{n}\log n)$  sybil nodes per attack edge and further requires  $g$  to be  $o(\sqrt{n}/\log n)$ . SybilLimit, in contrast, aims to reduce the number of sybil nodes accepted per attack edge to  $O(\log n)$  and further to allow for  $g = o(n/\log n)$ . This is challenging because SybilGuard's requirement on  $g = o(\sqrt{n}/\log n)$  is fundamental in its design and is simultaneously needed to ensure the following:

- **Sybil nodes accepted by SybilGuard.** The total number of sybil nodes accepted,  $O(g\sqrt{n}\log n)$ , is  $o(n)$ .
- **Escaping probability in SybilGuard.** The escaping probability of the verifier's random route,  $O(g\sqrt{n}\log n/n)$ , is  $o(1)$ .
- **Bad sample probability in SybilGuard.** When estimating the random route length, the probability of a bad sample,  $O(g\sqrt{n}\log n/n)$ , is  $o(1)$ .

Thus, to allow for larger  $g$ , SybilLimit needs to resolve all three issues above. Being more "robust" in only one aspect will not help.

SybilLimit has two component protocols: a *secure random route protocol* (Section V-A) and a *verification protocol* (Section V-B). The first protocol runs in the background and

TABLE II  
KEY NOTATIONS

$n$	total number of nodes in the honest region
$m$	total number of edges in the honest region
$g$	total number of attack edges
$r$	number of random routes that each verifier and suspect performs
$w$	length of individual random routes (in SybilLimit)
$l$	length of individual random routes in SybilGuard
$V$	verifier node
$S$	suspect node

maintains information used by the second protocol. Some parts of these protocols are adopted from SybilGuard, and we will indicate so when describing those parts. To highlight the major novel ideas in SybilLimit (as compared to SybilGuard), we will summarize these ideas in Section V-C. Later, Section VI will present SybilLimit's end-to-end guarantees. Table II summarizes the key notations used in the following discussion.

#### A. Secure Random Route Protocol

1) *Protocol Description:* We first focus on all the suspects in SybilLimit, i.e., nodes seeking to be accepted. Fig. 3 presents the pseudo-code for how they perform random routes—this protocol is adapted from SybilGuard with little modification. In the protocol, each node has a public/private key pair and communicates *only* with its neighbors in the social network. Every pair of neighbors share a unique symmetric secret key (the *edge key*, established out of band [13]) for authenticating each other. A sybil node  $M_1$  may disclose its edge key with some honest node  $A$  to another sybil node  $M_2$ . However, because all neighbors are authenticated via the edge key, when  $M_2$  sends a message to  $A$ ,  $A$  will still route the message as if it comes from  $M_1$ . In the protocol, every node has a precomputed random permutation  $x_1x_2 \dots x_d$  ( $d$  being the node's degree) as its routing table. The routing table never changes unless the node adds new neighbors or deletes old neighbors. A random route entering via edge  $i$  always exits via edge  $x_i$ .

A suspect  $S$  starts a random route along a uniformly random edge (of  $S$ ) and propagates along the route its public key  $K_S$  together with a counter initialized to 1. Every node along the route increments the counter and forwards the message until the counter reaches  $w$ , the length of a random route. SybilLimit's end guarantees hold even if sybil nodes (on the route) modify the message (see Section VI). In SybilLimit,  $w$  is chosen to be the mixing time of the honest region of the social network; given a fast-mixing social network,  $w = O(\log n)$ . As in SybilGuard [13], SybilLimit assumes that the nodes know a rough

upper bound on the graph's mixing time. Such an assumption is reasonable because the mixing time is logarithmic with  $n$  and, thus, is rather insensitive to  $n$ . Our later experiments show that choosing  $w = 10$  usually suffices for a million-node system, where SybilLimit ends up accepting 10 to 20 sybil nodes per attack edge. For smaller system sizes (e.g., 100 nodes), the only drawback of using  $w = 10$  is that the number of sybil nodes accepted per attack edge will remain at 10 to 20 (while with a more appropriate and smaller  $w$ , this number will drop).

Let " $A \rightarrow B$ " be the last (directed) edge traversed by  $S$ 's random route. We call this edge the *tail* of the random route. Node  $B$  will see the counter having a value of  $w$  and thus record  $K_S$  under the name of that tail (more specifically, under the name of " $K_A \rightarrow K_B$ ," where  $K_A$  and  $K_B$  are  $A$ 's and  $B$ 's public key, respectively). Notice that  $B$  may potentially overwrite any previously recorded key under the name of that tail. When  $B$  records  $K_S$ , we say that  $S$  registers its public key with that tail. Our verification protocol, described later, requires that  $S$  know  $A$ 's and  $B$ 's public keys and IP addresses. To do so, similar to SybilGuard, SybilLimit invokes the protocol in Fig. 3 a second time, where every node uses a "reversed" routing table (i.e., a random route entering via edge  $x_i$  will exit via edge  $i$ ). This enables  $A$  and  $B$  to propagate their public keys and IP addresses backward along the route so that  $S$  can learn about them.

Different from SybilGuard, SybilLimit invokes  $r$  independent instances (called *s-instances*) of the previous protocol for the suspects. The value of  $r$  should be  $\Theta(\sqrt{m})$ , and later we will explain how nodes can automatically pick the appropriate  $r$ . In every  $s$ -instance, each suspect uses the protocol in Fig. 3 to perform one random route and to register its public key with the tail. Across all  $s$ -instances, a suspect will thus register its public key with  $r$  tails, and each edge may be registered with up to  $r$  public keys. Additionally in every  $s$ -instance, SybilLimit invokes the protocol a second time for each suspect using reversed routing tables, so the suspects know their tails. The routing tables used in different  $s$ -instances are completely independent. Note, however, that all suspects share the same  $r$   $s$ -instances—this is critical to preserve the desirable convergence/back-traceability property among their random routes in the same  $s$ -instance.

Similarly, every verifier performs  $r$  random routes. To avoid undesirable correlation between the verifiers' random routes and the suspects' random routes, SybilLimit uses another  $r$  independent instances (called *v-instances*) for all verifiers. Verifiers do not need to register their public keys—they only need to know their tails. Thus, in each  $v$ -instance, SybilLimit invokes the protocol in Fig. 3 once for each verifier with reversed routing tables.

2) *Performance Overheads*: While SybilLimit uses the same technique as SybilGuard to do random routes, the overhead incurred is different because SybilLimit uses multiple instances of the protocol with a shorter route length. Interestingly, using  $\Theta(\sqrt{m})$  instances of the random route protocol does not incur extra storage or communication overhead by itself. First, a node does not need to store  $\Theta(\sqrt{m})$  routing tables since it can keep a single random seed and then generate any routing table on the fly as needed. Second, messages in different instances can be readily combined to reduce the number of messages. Remember that in all  $\Theta(\sqrt{m})$  instances, a node communicates only with its

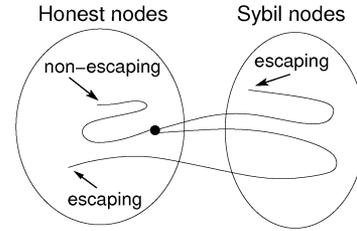


Fig. 4. Escaping and non-escaping tails.

$d$  neighbors. Thus, a node needs to send only  $d$  messages instead of  $\Theta(\sqrt{m})$  messages. SybilLimit inherits the idea from SybilGuard that an honest node should not have an excessive number of neighbors. This restriction helps bound the number of additional attack edges the adversary gets when an honest node is compromised. If there are too many neighbors, SybilLimit will (internally) only use a subset of the node's edges while ignoring all others. This implies that  $d$  will not be too large on average (e.g., 20). Finally, the total number of bits a node needs to send in the protocol is linear with the number of random routes times the length of the routes. Thus, the total number of bits sent in the  $d$  messages in SybilLimit is  $\Theta(\sqrt{m} \log n)$ , as compared to  $\Theta(\sqrt{n} \log n)$  in SybilGuard.

All these random routes need to be performed only one time (until the social network changes) and the relevant information will be recorded. Further aggressive optimizations are possible (e.g., propagating hashes of public keys instead of public keys themselves). We showed [13] that in a million-node system with average node degree being 10, an average node using SybilGuard needs to send 400 KBs of data every few days. Under the same parameters, an average node using SybilLimit would send around  $400 \times \sqrt{10} \approx 1300$  KB of data every few days, which is still quite acceptable. We refer the reader to [13] for further details. For much larger social networks (e.g., with billions of nodes), the overhead may become substantial. Further reducing such overhead (e.g., via indirect multihop validation) is part of our future work.

3) *Basic Security Properties*: The secure random route protocol provides some interesting basic security guarantees. We first formalize some notions. An honest suspect  $S$  has one *tail* in every  $s$ -instance, defined as the tail of its random route in that  $s$ -instance. We similarly define the  $r$  tails of a verifier. A random route starting from an honest node is called *escaping* if it ever traverses any attack edge. The tail of an escaping random route is called an *escaping tail* (Fig. 4), even if the escaping random route eventually comes back to the honest region. By directing the random route in specific ways, the adversary can control/influence to which directed edge an escaping tail corresponds, but the adversary has no influence over non-escaping tails.

In any given  $s$ -instance, for every attack edge connecting honest node  $A$  and sybil node  $M$ , imagine that we perform a random route starting from the edge " $M \rightarrow A$ ," until either a subsequent hop traverses an attack edge or the length of the route reaches  $w$ . Because the adversary can fake a series of routes that each end on one of the edges on this route, these edges are called *tainted tails*. Intuitively, the adversary may register arbitrary public keys with these tails. In a given  $s$ -instance, one can easily see that the set of tainted tails is disjoint

1.  $S$  sends to  $V$  its public key  $K_S$  and  $S$ 's set of tails  $\{(j, K_A, K_B) \mid S\text{'s tail in the } j\text{th s-instance is the edge } "A \rightarrow B"$  and  $K_A$  ( $K_B$ ) is  $A$ 's ( $B$ 's) public key};  
*// Apply the intersection condition (the instance number is ignored when determining intersection)*
2.  $V$  computes the set of intersecting tails  $X = \{(i, K_A, K_B) \mid (i, K_A, K_B) \text{ is } V\text{'s tail and } (j, K_A, K_B) \text{ is } S\text{'s tail}\}$ ;
3. For every  $(i, K_A, K_B) \in X$ ,  $V$  authenticates  $B$  using  $K_B$  and asks  $B$  whether  $S$  is registered under " $K_A \rightarrow K_B$ "  
 If not, remove  $(i, K_A, K_B)$  from  $X$ ;
4. If  $X$  is empty then reject  $S$  and return;
- // Apply the balance condition ( $c_i$  is the counter for  $V$ 's tail in the  $i$ th v-instance)*
5. Let  $a = (1 + \sum_{i=1}^r c_i)/r$  and  $b = h \cdot \max(\log r, a)$ ; *// see text for description of  $h$*
6. Let  $c_{min}$  be the smallest counter among those  $c_i$ 's corresponding to  $(i, K_A, K_B)$  that still remain in  $X$  (with tie-breaking favoring smaller  $i$ );
7. If  $(c_{min} + 1 > b)$  then reject  $S$ ; otherwise, increment  $c_{min}$  and accept  $S$ ;

Fig. 5. Protocol for  $V$  to verify  $S$ .  $V$  has  $r$  counters  $c_1, \dots, c_r$  initialized to zero at startup time.

from the set of non-escaping tails from honest suspects. The reason is that random routes are back-traceable, and starting from a non-escaping tail, one can always trace back to the starting node of the random route, encountering only honest nodes. This means that an honest suspect will never need to compete with the sybil nodes for a tail as long as its random route is non-escaping.

After the secure random route protocol stabilizes (i.e., all propagations have completed), the following properties are guaranteed to hold:

- In every s-instance, each directed edge in the honest region allows only one public key to be registered.
- In every s-instance, an honest suspect  $S$  can always register its public key with its non-escaping tail (if any) in that s-instance.
- In every s-instance, among all the directed edges in the honest region, sybil nodes can register their public keys only with tainted tails. This is because nodes communicate with only their neighbors (together with proper authentication) and also because the counter in the registration message is incremented at each hop.
- In every s-instance (v-instance), if an honest suspect  $S$  (an honest verifier  $V$ ) has a non-escaping tail " $A \rightarrow B$ ," then  $S$  ( $V$ ) knows  $A$ 's and  $B$ 's public keys.

4) *User and Node Dynamics:* Most of our discussion so far assumes that the social network is static and all nodes are online. All techniques in SybilGuard to efficiently deal with user/node dynamics, as well as techniques to properly overwrite stale registration information for preventing certain attacks [13], apply to SybilLimit without modification. We do not elaborate on these due to space limitations.

## B. Verification Protocol

1) *Protocol Description:* After the secure random route protocol stabilizes, a verifier  $V$  can invoke the verification protocol in Fig. 5 to determine whether to accept a suspect  $S$ .  $S$  must satisfy both the *intersection condition* (Steps 2–4 in Fig. 5) and the *balance condition* (Steps 5–7) to be accepted.

The intersection condition requires that  $S$ 's tails and  $V$ 's tails must intersect (instance number is ignored when determining intersection), with  $S$  being registered at the intersecting tail. In contrast, SybilGuard has an intersection condition on nodes (instead of on edges or tails). For the balance condition,  $V$  maintains  $r$  counters corresponding to its  $r$  tails. Every accepted suspect increments the "load" of some tail. The balance condition

requires that accepting  $S$  should not result in a large "load spike" and cause the load on any tail to exceed  $h \cdot \max(\log r, a)$ . Here,  $a$  is the current average load across all  $V$ 's tails, and  $h > 1$  is some universal constant that is not too small (we use  $h = 4$  in our experiments). In comparison, SybilGuard does not have any balance condition.

2) *Performance Overheads:* The verification protocol can be made highly efficient. Except for Steps 1 and 3, all steps in the protocol involve only local computation. Instead of directly sending  $\Theta(r)$  public keys in Step 1,  $S$  can readily use a Bloom Filter [14] to summarize the set of keys. In Step 3, for every intersecting tail in  $X$ ,  $V$  needs to contact one node. On average, the number of intersections between a verifier  $V$  and an honest suspect  $S$  in the honest region is  $O(1)$  with  $r = \Theta(\sqrt{m})$ , resulting in  $O(1)$  messages. The adversary may intentionally introduce additional intersections in the sybil region between  $V$ 's and  $S$ 's escaping tails. However, if those extra intersecting nodes (introduced by the adversary) do not reply,  $V$  can blacklist them. If they do reply and if  $V$  is overwhelmed by the overhead of such replies, then the adversary is effectively launching a DoS attack. Notice that the adversary can launch such a DoS attack against  $V$  even if  $V$  were not running SybilLimit. Thus, such attacks are orthogonal to SybilLimit.

## C. Key Ideas in SybilLimit, vis-à-vis SybilGuard

This section highlights the key novel ideas in SybilLimit that eventually lead to the substantial end-to-end improvements over SybilGuard.

1) *Intersection Condition:* To help convey the intuition, we will assume  $g = 1$  in the following. In SybilLimit, each node uses  $r = \Theta(\sqrt{m})$  random routes of length  $w = \Theta(\log n)$  instead of a single random route of length  $l = \Theta(\sqrt{n} \log n)$  as in SybilGuard.<sup>1</sup> In SybilGuard, each node along a random route corresponds to a "slot" for registering the public key of some node. The adversary can fake  $l$  distinct random routes of length  $l$  that cross the attack edge and enter the honest region. This means that the adversary will have  $1+2+\dots+l = \Theta(l^2) = \Theta(n \log^2 n)$  slots for the sybil nodes in SybilGuard.

In SybilLimit, the tail of each random route corresponds to a "slot" for registration. In any given s-instance, the adversary can fake  $w$  distinct random routes of length  $w$  that cross the

<sup>1</sup>As an engineering optimization, a degree- $d$  node in SybilGuard can perform  $d$  random routes of length  $\Theta(\sqrt{n} \log n)$ , but this does not improve SybilGuard's asymptotic guarantees.

attack edge and enter the honest region. Notice that here SybilLimit reduces the number of such routes by using a  $w$  that is much smaller than  $l$ . Furthermore, because we are concerned only with tails now, in the given  $s$ -instance, the adversary will have only  $w$  slots. With  $r$   $s$ -instances, the adversary will have  $r \cdot w = \Theta(\sqrt{m} \log n)$  such slots total for all the sybil nodes. This reduction from  $\Theta(n \log^2 n)$  slots to  $\Theta(\sqrt{m} \log n)$  slots is the first key step in SybilLimit.

However, doing  $r$  random routes introduces two problems. The first is that it is impossible for a degree- $d$  node to have more than  $d$  distinct random routes if we directly use SybilGuard's approach. SybilLimit observes that one can use many independent instances of the random route protocol while still preserving the desired convergence/back-traceability property. The second problem is more serious. SybilGuard relies on the simple fact that the number of distinct routes from the adversary is  $l$ . All slots on the same route must have the same public key registered. This ensures that the total number of sybil nodes registered is  $l$ . In SybilLimit, there are  $r \cdot w$  distinct routes from the adversary. Thus, a naive design may end up accepting  $r \cdot w = \Theta(\sqrt{m} \log n)$  sybil nodes, which is even worse than SybilGuard. SybilLimit's key idea here is to perform intersections on edges instead of on nodes. Because the stationary distribution on edges is always uniform in any graph, it ensures that the *flip-side* of the Birthday Paradox holds. Namely,  $\Theta(\sqrt{m})$  slots are both sufficient and *necessary* for intersection to happen (with high probability). Together with earlier arguments on the number of slots in SybilLimit, this will eventually allow us to prove that the number of sybil nodes with tails intersecting with  $V$ 's non-escaping tails (more precisely,  $V$ 's *uniform* non-escaping tails—see later) is  $O(\log n)$  per attack edge.

2) *Balance Condition*: In SybilGuard, the verifier's random route is either escaping or non-escaping, resulting in an "all-or-nothing" effect. For SybilGuard to work, this single random route must be non-escaping. Because of the large  $l$  of  $\Theta(\sqrt{n} \log n)$ , the escaping probability will be  $\Omega(1)$  once  $g$  reaches  $\Omega(\sqrt{n}/\log n)$ . Using much shorter random routes of length  $w$  in SybilLimit decreases such escaping probability. On the other hand, because a verifier in SybilLimit needs to do  $r$  such routes, it remains quite likely that *some* of them are escaping. In fact, with  $r = \Theta(\sqrt{m})$  and  $w = \Theta(\log n)$ , the probability of at least one of the  $r$  routes being escaping in SybilLimit is even larger than the probability of the single length- $l$  random route being escaping in SybilGuard. Thus, so far we have only made the "all-or-nothing" effect in SybilGuard fractional.

SybilLimit relies on its (new) balance condition to address this fraction of escaping routes. To obtain some intuition, let us imagine the verifier  $V$ 's tails as bins that can accommodate up to a certain load. When  $V$  accepts a suspect  $S$ , out of all of  $V$ 's tails that intersect with  $S$ 's tails,  $S$  conceptually increments the load of the least loaded tail/bin. Because of the randomness in the system, one would conjecture that all of  $V$ 's tails should have similar load. If this is indeed true, then we can enforce a quota on the load of each tail, which will in turn bound the number of sybil nodes accepted by  $V$ 's escaping tails. Later, we will show that the balance condition bounds the number within  $O(g \log n)$ .

3) *Benchmarking Technique*: The SybilLimit protocol in Figs. 3 and 5 assumes that  $r = \Theta(\sqrt{m})$  is known. Obviously, without global knowledge, every node in SybilLimit needs to estimate  $r$  locally. Recall that SybilGuard also needs to estimate some system parameter (more specifically, the length of the walk). SybilGuard uses the sampling technique to do so, which only works for  $g = o(\sqrt{n}/\log n)$ . To allow any  $g = o(n/\log n)$ , SybilLimit avoids sampling completely. Instead, it uses a novel and perhaps counterintuitive *benchmarking technique* that mixes the real suspects with some random *benchmark suspects* that are already known to be mostly honest. The technique guarantees that a node will never overestimate  $r$  regardless of the adversary's behavior. If the adversary causes an underestimation for  $r$ , somewhat counterintuitively, the technique can ensure that SybilLimit still achieves its end guarantees despite the underestimated  $r$ . We will leave the detailed discussion to Section VII.

## VI. PROVABLE GUARANTEES OF SYBILLIMIT

While the intersection and balance conditions are simple at the protocol/implementation level, it is far from obvious why the designs provide the desired guarantees. We adopt the philosophy that all guarantees of SybilLimit must be proven mathematically because experimental methods can cover only a subset of the adversary's strategies. Our proofs pay special attention to the correlation among various events, which turns out to be a key challenge. We cannot assume independence for simplicity because, after all, SybilLimit exactly leverages external correlation among random routes. The following is the main theorem on SybilLimit's guarantee.

*Theorem 3*: Assume that the social network's honest region is fast-mixing and  $g = o(n/\log n)$ . For any given constants (potentially close to zero)  $\epsilon > 0$  and  $\delta > 0$ , there is a set of  $(1 - \epsilon)n$  honest verifiers and universal constants  $w_0$  and  $r_0$ , such that using  $w = w_0 \log n$  and  $r = r_0 \sqrt{m}$  in SybilLimit will guarantee that for any given verifier  $V$  in the set, with probability of at least  $1 - \delta$ ,  $V$  accepts at most  $O(\log n)$  sybil nodes per attack edge and at least  $(1 - \epsilon)n$  honest nodes.

For the remaining small fraction of  $\epsilon n$  honest verifiers, SybilLimit provides a degraded guarantee that is not provable. Because of space limitations, we will provide mostly intuitions in the following and leave formal/complete proofs to our technical report [43].

### A. Intersection Condition

1) *Preliminaries: Classifying Tails and Nodes*: As preparation, we first carefully classify tails and nodes. Later, we will need to use different techniques to reason about different types of tails and nodes. Table III summarizes the key definitions we will use. Consider a given verifier  $V$  (or suspect  $S$ ) and a given  $v$ -instance (or  $s$ -instance). We classify its tail into three possibilities: 1) the tail is an *escaping tail* (recall Section V-A); 2) the tail is not escaping and is drawn from the (uniform) edge stationary distribution (i.e., a *uniform tail*); or 3) the tail is not escaping and is drawn from some unknown distribution on the edges (i.e., a *non-uniform tail*).<sup>2</sup> In a given  $v$ -instance, the routing tables

<sup>2</sup>A finite-length random walk can only approach but never reach the stationary distribution. Thus a small fraction of tails will be non-uniform (also see Theorem 1).

TABLE III  
TERMINOLOGY USED IN PROOFS (SEE TEXT FOR PRECISE DEFINITIONS)

escaping route	random route from an honest node that traverses an attack edge
escaping tail	tail of an escaping route
tainted tail	any edge in the honest region on a length- $w$ random route starting from an attack edge
uniform tail	non-escaping tail from uniform edge distribution
non-uniform tail	non-escaping tail that is not a uniform tail
non-escaping node	honest node such that a length- $w$ random walk has a uniform tail with $1 - o(1)$ probability
escaping node	honest node that is not a non-escaping node
uniform tail set	the set of all uniform tails of a given honest node
tainted tail set	set of all tainted tails

of all honest nodes will entirely determine whether  $V$ 's tail is escaping and in the case of a non-escaping tail, which edge is the tail. Thus, the adversary has no influence over non-escaping tails.

Since the distribution of the non-uniform tails is unknown, few probabilistic properties can be derived for them. Escaping tails are worse because their distribution is controlled by the adversary. We thus would like to first quantify the (small) fraction of non-uniform tails and escaping tails. Assuming that the honest region of the social network is fast-mixing, our technical report [43] proves that for most honest nodes, most of their tails are uniform tails.

*Lemma 4:* Consider any given constant (potentially close to zero)  $\epsilon > 0$ . We can always find a universal constant  $w_0 > 0$ , such that there exists a set  $H$  of at least  $(1 - \epsilon)n$  honest nodes (called *non-escaping nodes*) satisfying the following property: If we perform a length- $w$  random walk starting from any non-escaping node with  $w = w_0 \log n$ , then the tail is a uniform tail (i.e., a uniformly random directed edge in the honest region) with probability of at least  $1 - O(g \log n/n)$ .

As a reminder, the probability in the above lemma is defined over the domain of all possible routing table states—obviously, if all routing tables are already determined, the tail will be some fixed edge.

It is still possible for the tail of a non-escaping node to be escaping or non-uniform—it is just that such probability is  $O(g \log n/n) = o(1)$  for  $g = o(n/\log n)$ . We will not ignore this  $o(1)$  fraction of tails, but knowing that they are of  $o(1)$  fraction will facilitate our proof later. An honest node that is not non-escaping is called an *escaping node*. By Lemma 4, we have at most  $\epsilon n$  escaping nodes; such nodes are usually near the attack edges. Notice that given the topology of the honest region and the location of the attack edges, we can fully determine the probability of the tail of a length- $w$  random walk starting from a given node  $V$  being a uniform tail. In turn, this means whether a node  $V$  is escaping is not affected by the adversary. In the remainder of this paper, unless specifically mentioned, when we say “honest node/verifier/suspect,” we mean “non-escaping (honest) node/verifier/suspect.” We will not, however, ignore escaping nodes in the arguments since they may potentially disrupt the guarantees for non-escaping nodes.

For each verifier  $V$ , define its *tail set* as

$$\{(i, e) | e \text{ is } V\text{'s tail in the } i\text{th } v\text{-instance}\}.$$

$V$ 's *uniform tail set*  $\mathcal{U}(V)$  is defined as

$$\mathcal{U}(V) = \{(i, e) | e \text{ is } V\text{'s tail in the } i\text{th } v\text{-instance and } e \text{ is a uniform tail}\}.$$

Notice that the distribution of  $\mathcal{U}(V)$  is not affected by the adversary's strategy. We similarly define the tail set and uniform tail set for every suspect  $S$ . We define the *tainted tail set*  $\nabla$  as:  $\nabla = \cup_{i=1}^r \nabla_i$ , where

$$\nabla_i = \{(i, e) | e \text{ is a tainted tail in the } i\text{th } s\text{-instance}\}.$$

Again, the definition of  $\nabla$  is not affected by the behavior of the adversary, as all these tails are in the honest region. Further notice that in a given  $s$ -instance for each attack edge, we can have at most  $w$  tainted tails. Thus,  $|\nabla_i| \leq g \times w$  and  $|\nabla| \leq rgw = O(rg \log n)$ .

With slight abuse of notation, we say that a tail set *intersects* with a tail  $e$  as long as the tail set contains an element  $(i, e)$  for some  $i$ . The *number of intersections* with  $e$  is defined to be the number of elements of the form  $(i, e)$ . We double count  $e$  in different instances because for every element  $(i, e)$ , an arbitrary public key can be registered under the name of  $e$  in the  $i$ th  $s$ -instance. For two tail sets  $T_1$  and  $T_2$ , we define the *number of intersections* between them as:  $\sum_{(j,e) \in T_2} (\# \text{ intersections between } e \text{ and } T_1)$ . For example,  $\{(1, e_1), (2, e_1)\}$  and  $\{(2, e_1), (3, e_1)\}$  have four intersections.  $T_1$  and  $T_2$  *intersect* if and only if the number of intersections between them is larger than 0.

2) *Tail Intersection Between the Verifier and Honest Suspects:* The *intersection condition* requires that for a verifier  $V$  to accept a suspect  $S$ ,  $V$ 's tail set and  $S$ 's tail set must intersect with  $S$  being registered at some intersecting tail. Intuitively, a verifier  $V$  and an honest suspect  $S$  will satisfy the intersection condition because most of their tails are uniform tails. Thus, for proper  $r = \Theta(\sqrt{m})$ , the Birthday Paradox ensures that the two tail sets intersect. This is true despite the fact that we are not able to reason about the distribution of non-uniform tails and escaping tails. More precisely, our technical report [43] proves that for any given constant  $\delta > 0$ ,  $V$  and  $S$  will satisfy the intersection condition with probability  $1 - \delta$  when  $r = r_0 \sqrt{m}$ , with  $r_0$  being an appropriately chosen constant.

3) *Tail Intersection Between the Verifier and Sybil Suspects:* A tail of a verifier  $V$  can be uniform, non-uniform, or escaping. Regardless of the classification, it can potentially intersect with the tail sets of sybil suspects. Our balance condition later will take care of  $V$ 's non-uniform tails and escaping tails. For now, we want to reason about the intersections between  $V$ 's uniform tails and the tails of the sybil suspects.

By definition, all uniform tails of  $V$  are in the honest region. From the secure random route property, the tainted tail set  $\nabla$  contains all tails that the sybil suspects can possibly have in the honest region. Let random variable  $X$  be the number of intersections between  $\nabla$  and  $\mathcal{U}(V)$ . Under a given value of  $X$ , it is obviously impossible for more than  $X$  sybil suspects to have tails intersecting with  $\mathcal{U}(V)$ . Thus, upper-bounding  $X$  is sufficient for our reasoning. Our technical report [43] proves that for any given constant  $\delta > 0$ , with probability of at least  $1 - \delta$ ,  $X = O(g \log n)$ , and thus the number of sybil suspects with tails intersecting with  $\mathcal{U}(V)$  is also  $O(g \log n)$ .

## B. Balance Condition

In this section, for any verifier  $V$ , we treat all of its non-uniform tails as escaping tails. Obviously, this only increases the adversary's power and makes our arguments pessimistic. The goal of the balance condition is to bound the number of sybil nodes accepted by  $V$ 's escaping tails without significantly hurting honest suspects (who are subject to the same balance condition). While the condition is simple, rigorously reasoning about it turns out to be quite tricky due to the external correlation among random routes and also adversarial disruption that may intentionally cause load imbalance. This introduces challenges particularly for proving why most honest suspects will satisfy the balance condition despite all these disruptions.

1) *Effects on Sybil Suspects:* Here, we focus on the intuition. An honest verifier  $V$  has a total of  $r$  tails, out of which roughly  $O(r \cdot (g \log n/n)) = o(r)$  are escaping tails (by Lemma 4). Those  $r - o(r)$  uniform tails together will accept at most  $n$  honest suspects and at most  $O(g \log n)$  sybil suspects (from the intersection condition). Given the various randomness in the system, one would hope that each uniform tail should accept roughly  $(n + O(g \log n))/(r - o(r)) = O(n/r)$  suspects.

The balance condition thus intends to impose a limit of  $O(n/r)$  on the number of suspects that each tail can accept. This is effective because  $V$  has only roughly  $O(r \cdot (g \log n/n))$  escaping tails. Thus, the total number of sybil suspects accepted by the escaping tails will be roughly  $O(n/r) \cdot O(r \cdot (g \log n/n)) = O(g \log n)$ , which is precisely what we are aiming for. Of course,  $n$  is unknown to SybilLimit, thus the protocol uses a floating bar as in Fig. 5. Our technical report [43] explains further and formalizes the above arguments.

2) *Effects on Honest Suspects:* Reasoning about the effects of the balance condition on honest suspects is the most complex part in SybilLimit's proof. For space limitations, we leave the arguments and proofs to our technical report [43], which shows that most of the honest suspects will satisfy the balance condition.

## VII. ESTIMATING THE NUMBER OF ROUTES NEEDED

We have shown that in SybilLimit, a verifier  $V$  will accept  $(1 - \epsilon)n$  honest suspects with probability  $1 - \delta$  if  $r = r_0 \sqrt{m}$ . The constant  $r_0$  can be directly calculated from the Birthday Paradox and the desired end probabilistic guarantees. On the other hand,  $m$  is unknown to individual nodes. Adapting the sampling approach from SybilGuard (as reviewed in Section IV) is not possible because that approach is fundamentally limited to  $g = o(\sqrt{n}/\log n)$ .

### A. Benchmarking Technique

SybilLimit uses a novel and perhaps counterintuitive *benchmarking technique* to address the previous problem by mixing the real suspects with some random *benchmark nodes* that are already known to be mostly honest. Every verifier  $V$  maintains two sets of suspects: the *benchmark set*  $K$  and the *test set*  $T$ . The *benchmark set*  $K$  is constructed by repeatedly performing random routes of length  $w$  and then adding the ending node (called the *benchmark node*) to  $K$ . Let  $K^+$  and  $K^-$  be the set of honest and sybil suspects in  $K$ , respectively. SybilLimit does not know which nodes in  $K$  belong to  $K^+$ . However, a key prop-

erty here is that because the escaping probability of such random routes is  $o(1)$ , even without invoking SybilLimit, we are assured that  $|K^-|/|K| = o(1)$ . The *test set*  $T$  contains the real suspects that  $V$  wants to verify, which may or may not happen to belong to  $K$ . We similarly define  $T^+$  and  $T^-$ . Our technique will hinge upon the adversary not knowing  $K^+$  or  $T^+$  (see later for how to ensure this), even though it may know  $K^+ \cup T^+$  and  $K^- \cup T^-$ .

To estimate  $r$ , a verifier  $V$  starts from  $r = 1$  and then repeatedly doubles  $r$ . For every  $r$  value,  $V$  verifies all suspects in  $K$  and  $T$ . It stops doubling  $r$  when most of the nodes in  $K$  (e.g., 95%) are accepted, and then makes a final determination for each suspect in  $T$ .

### B. No Overestimation

Once  $r$  reaches  $r_0 \sqrt{m}$ , most of the suspects in  $K^+$  will indeed be accepted, regardless of the behavior of the adversary. Furthermore, because  $|K^+|/|K| = 1 - o(1)$ , having an  $r$  of  $r_0 \sqrt{m}$  will enable us to reach the threshold (e.g., 95%) and stop doubling  $r$  further. Thus,  $V$  will never overestimate  $r$  (within a factor of 2).

### C. Underestimation Will Not Compromise SybilLimit's Guarantees

It is possible for the adversary to cause an underestimation of  $r$  by introducing artificial intersections between the escaping tails of  $V$  and the escaping tails of suspects in  $K^+$ . This may cause the threshold to be reached before  $r$  reaches  $r_0 \sqrt{m}$ .

What if SybilLimit operates under an  $r < r_0 \sqrt{m}$ ? Interestingly, SybilLimit can bound the number of sybil nodes accepted within  $O(\log n)$  per attack edge not only when  $r = r_0 \sqrt{m}$ , but also for  $r < r_0 \sqrt{m}$  (see [43] for proofs). To obtain some intuition, first notice that the number of sybil nodes with tails intersecting with  $V$ 's uniform tails (Section VI-A) can only decrease when  $r$  is smaller. Second, the arguments regarding the number of sybil nodes accepted by  $V$ 's escaping tails and non-uniform tails (Section VI-B) hinges only upon the *fraction* of those tails and not the value of  $r$ .

Using  $r < r_0 \sqrt{m}$ , however, will decrease the probability of tail intersection between the verifier and an honest suspect. Here, we leverage a second important property of the benchmark set. Namely, conditioned upon the random routes for picking benchmark nodes being non-escaping, the adversary will not know which nodes are picked as benchmark nodes. (If the adversary may eavesdrop messages, we can readily encrypt messages using edge keys.) As a result, given an honest suspect, the adversary cannot tell whether it belongs to  $K^+$  or  $T^+$ . If most (e.g., 95%) of the suspects in  $K$  are accepted, then most suspects in  $K^+$  must be accepted as well, since  $|K^+|/|K| = 1 - o(1)$ . If most suspects in  $K^+$  are accepted under  $r < r_0 \sqrt{m}$ , the adversary must have intentionally caused intersection between  $V$  and the suspects in  $K^+$ . Because the adversary cannot tell whether an honest suspect belongs to  $K^+$  or  $T^+$ , it cannot introduce intersections *only* for suspects in  $K^+$ ; it must introduce intersections for suspects in  $T^+$  as well. Thus, most suspects in  $T^+$  will be accepted as well under the given  $r$ .

### D. Further Discussions

The benchmarking technique may appear counterintuitive in two aspects. First, if SybilLimit uses an underestimated  $r$ , it

- |   |
|---|
| <ol style="list-style-type: none"> <li>1. <math>V</math> starts with two sets of suspects, <math>K</math> and <math>T</math>;</li> <li>2. Let set <math>A = \emptyset</math> and <math>r = 1</math>;</li> <li>3. While <math>( A \cap K / K  &lt; 95\%)</math></li> <li>4.   For every suspect <math>S \in ((K \cup T) \setminus A)</math></li> <li>5.     Verify <math>S</math> using the protocol from Fig. 5;</li> <li>6.     If <math>S</math> is accepted, <math>A = A \cup \{S\}</math>;</li> <li>7.     Double <math>r</math>;</li> <li>8. <math>V</math> accepts (and only accepts) the suspects in <math>A \cap T</math>;</li> </ol> |
|---|

Fig. 6. Pseudo-code for the benchmarking technique.

will be the adversary that helps it to accept most of the honest nodes. While this is true, SybilLimit is still needed to bound the number of sybil nodes accepted and also to prevent  $r$  from growing beyond  $r_0\sqrt{m}$ . Second, the benchmark set  $K$  is itself a set with  $o(1)$  fraction of sybil nodes. Thus, it may appear that an application can just as well use the nodes in  $K$  directly and avoid the full SybilLimit protocol. However, the set  $K$  is constructed randomly and may not contain some specific suspects that  $V$  wants to verify.

For a more rigorous understanding of the benchmarking technique, we can view the process as a sampling algorithm for estimating the fraction of the suspects accepted in the set  $T^+ \cup K^+$ . We take  $|K^+|$  samples from the set and observe that fraction  $f$  of the samples are accepted. Classic estimation theory [44] tells us that if  $|K^+| = \Theta((1/\epsilon^2) \log(1/\delta))$ , then the fraction of the accepted suspects in  $T^+$  is within  $f \pm \epsilon$  with probability of at least  $1 - \delta$ . It is important to see that the needed size of  $K^+$  (and thus  $K$ ) is independent of the size of  $T$ . Simple simulation experiments show that  $|K| = 30$  gives us an average  $\epsilon$  of 0.0322.

Care must be taken when implementing the benchmarking technique. The technique hinges on the fact that the adversary cannot distinguish suspects in  $K^+$  from suspects in  $T^+$ . A naive implementation would gradually increase  $r$  and invoke the verification protocol from Fig. 5 multiple times (under different  $r$ ) for each suspect. This will leak (probabilistic) information to the adversary. Namely, if the adversary notices that  $V$  still increases  $r$  even after a certain honest suspect  $S$  is accepted, then the conditional probability that  $S$  belongs to  $K^+$  decreases. Under the increased  $r$ , the adversary may then choose to cause other suspects to be accepted while causing  $S$  to be rejected. This will then violate the assumption that  $K^+$  is a set of uniform samples from  $K^+ \cup T^+$ .

To ensure that  $K^+$  is a set of uniform samples, we automatically consider a suspect  $S$  that is accepted under a certain  $r$  to be accepted under larger  $r$  values, without re-verifying this. Fig. 6 presents the pseudo-code, which maintains a set  $A$  including all suspects accepted so far. Now, imagine that the adversary notices that  $V$  still increases  $r$  despite those suspects in  $A$  being accepted. This tells the adversary that the suspects in  $A$  are less likely to belong to  $K^+$  than those suspects not in  $A$ . However, the adversary can no longer reverse the determinations already made for those suspects in  $A$ . The adversary can still influence future determinations on those suspects not in  $A$ . However, all these suspects have the same probability of being in  $K^+$ , so it does not help the adversary to favor some suspects (i.e., cause them to be accepted) over others.

## VIII. LOWER BOUND

SybilLimit bounds the number of sybil nodes accepted within  $O(\log n)$  per attack edge. A natural question is whether we can further improve the guarantees. For example, it may appear that SybilLimit does not currently have any mechanism to limit the routing behavior of sybil nodes. One could imagine requiring nodes to commit (cryptographically) to their routing tables so that sybil nodes could not perform random routes in an inconsistent fashion. We will show, however, that such techniques or similar techniques can provide at most a  $\log n$  factor of improvement because the total number of sybil nodes accepted is lower-bounded by  $\Omega(1)$  per attack edge.

SybilLimit entirely relies on the observation that if the adversary creates too many sybil nodes, then the resulting social network will no longer have  $O(\log n)$  mixing time. Our technical report [43] proves that for any given constant  $c$ , any  $g \in [1, n]$ , and any graph  $G$  with  $n$  honest nodes and  $O(\log n)$  mixing time, it is always possible for the adversary to introduce  $c \cdot g$  sybil nodes via  $g$  attack edges so that the augmented graph's mixing time is  $O(\log n')$ , where  $n' = n + c \cdot g$ . There are actually many ways to create such an augmented graph. One way (as in our proof) is to pick  $g$  nodes arbitrarily from  $G$  and attach a group of  $c$  sybil nodes to each of them (using a single attack edge). It does not matter how the  $c$  sybil nodes in a group are connected with each other, as long as they are connected. Now, because the augmented graph has the same mixing time (i.e.,  $O(\log n')$ ) as a "normal" social network with  $n'$  nodes, as long as the protocol solely relies on mixing time, we cannot distinguish these sybil nodes from honest nodes. In other words, all protocols based on mixing time will end up accepting  $\Omega(1)$  sybil nodes per attack edge.

## IX. EXPERIMENTS WITH ONLINE SOCIAL NETWORKS

### A. Goal of Experiments

We have proved that SybilLimit can bound the number of sybil nodes accepted within  $O(\log n)$  per attack edge, which improved upon SybilGuard's guarantee of  $O(\sqrt{n} \log n)$ . However, these provable guarantees of SybilLimit (and SybilGuard as well) critically rely on the assumption that social networks have small (i.e.,  $O(\log n)$ ) mixing time. Our experiments thus mainly serve to validate such an assumption, based on real-world social networks. Such validation has a more general implication beyond SybilLimit—these results will tell us whether the approach of leveraging social networks to combat sybil attacks is valid. A second goal of our experiments is to gain better understanding of the hidden constant in SybilLimit's  $O(\log n)$  guarantee. Finally, we will also provide some example numerical comparisons between SybilGuard and SybilLimit. However, it is *not* our goal to perform a detailed experimental comparison because SybilLimit's improvement over SybilGuard is already rigorously proved.

### B. Social Network Data Sets

We use three crawled online social network data sets in our experiments: Friendster, LiveJournal, and DBLP (Table IV). They are crawls of <http://www.friendster.com>, <http://www.livejournal.com>, and <http://dblp.uni-trier.de>, respectively. The DBLP data set is publicly available, but the other two are not.

TABLE IV  
SOCIAL NETWORK DATA SETS

Data set	Friendster	LiveJournal	DBLP	Kleinberg
Data set source	[45]	[46]	[47]	[15]
Date crawled	Nov-Dec 2005	May 2005	April 2006	not applicable
# nodes	932,512	900,822	106,002	1,000,000
# undirected edges	7,835,974	8,737,636	625,932	10,935,294
$w$ used in SybilLimit	10	10	15	10
$r$ used in SybilLimit	8,000	12,000	3,000	10,000

We also experiment with Kleinberg’s synthetic social network [15], which we used [13] to evaluate SybilGuard.

Strictly speaking, DBLP is a bibliography database and not a social network. To derive the “social network” from DBLP, we consider two people having an edge between them if they have ever co-authored a paper. Because of the closely clustered co-authoring relationships among researchers, we expect such a social network to be more slowly mixing than standard social networks. Thus, we use DBLP as a bound on the worst-case scenario. Table IV presents the basic statistics of the four social networks after appropriate preprocessing (e.g., converting pairs of directed edges to undirected edges, removing low ( $<5$ ) degree nodes, taking the largest connected component—see [43]). Same as in [13], we then select random nodes as attacker nodes (that launch the sybil attack) until the number of attack edges reaches  $g$ .<sup>3</sup>

### C. Results: Mixing Time of Real-World Social Networks

In SybilLimit, the only parameter affected by mixing time is the length of the random routes ( $w$ ). Namely,  $w$  should be at least as large as the mixing time. It is not possible to directly show that our data sets have  $O(\log n)$  mixing time since  $O(\log n)$  is asymptotic behavior. It is not necessary to do so either since all we need to confirm is that rather small  $w$  values are already sufficient for SybilLimit to work well.

For Friendster and LiveJournal, we use  $w = 10$  (see Table IV). Random routes do not seem to reach good enough mixing for SybilLimit with  $w$  values much smaller than 10 (e.g., 5) in these two social networks. We use  $w = 15$  for DBLP. As expected, DBLP has a worse mixing property than the other social networks. Our results will show that these small  $w$  values are already sufficient to enable good enough mixing in our large-scale social networks (with  $10^5$  to around  $10^6$  nodes) for SybilLimit to work well. More precisely, under these settings, an average verifier in SybilLimit is able to accept over 95% of the honest nodes.

It is worth noting that social networks are well known to have groups or communities where intragroup edges are much denser than intergroup edges. In fact, there are explicitly defined communities in LiveJournal for users to join, while people in DBLP by definition form research communities. Our results thus show that somewhat counterintuitively and despite such groups, the sparse intergroup edges in these real-world social networks are sufficient to provide good mixing properties.

<sup>3</sup>We also consider the “cluster” placement of attack edges from [13]; the results are qualitatively the same.

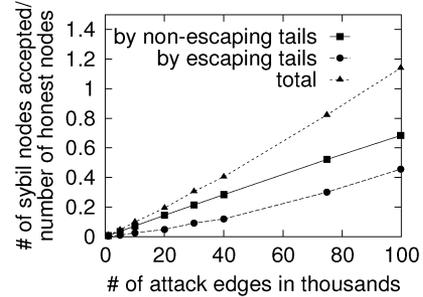


Fig. 7. Friendster.

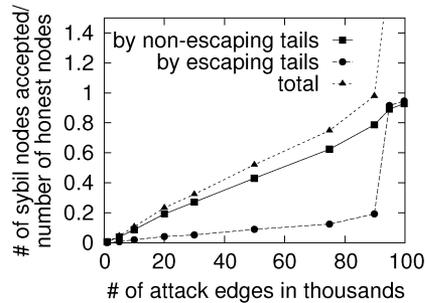


Fig. 8. LiveJournal.

### D. Results: SybilLimit’s End Guarantees

We use the  $w$  values from Table IV to simulate SybilLimit and determine the number of sybil nodes accepted. Our simulator does not implement the estimation process for  $r$ . Rather, we directly use the  $r$  values from Table IV, which are obtained based on the value of  $m$  and the Birthday Paradox. We use 4 for the universal constant  $h$  in all our experiments. We have observed (results not included) that  $h = 2.5$  is already sufficient in most cases, while excessively large  $h$  (e.g., 10) can unnecessarily weaken the guarantees (though not asymptotically). We always simulate the adversary’s optimal strategy (i.e., worst-case for SybilLimit).

Figs. 7–10 present the number of sybil nodes accepted by a randomly chosen verifier  $V$  (as a fraction of the number of honest nodes  $n$ ) in each social network. We present a fraction to allow comparison across social networks with different  $n$ . We have repeated the experiments from a number of verifiers, yielding similar results. For all cases, we experiment with  $g$  up to the point where the number of sybil nodes accepted reaches  $n$ . The figures further break down the sybil nodes accepted into those accepted by  $V$ ’s non-escaping tails versus those accepted by  $V$ ’s escaping tails. The first component is bounded by the

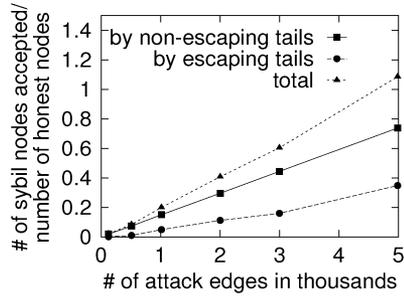


Fig. 9. DBLP.

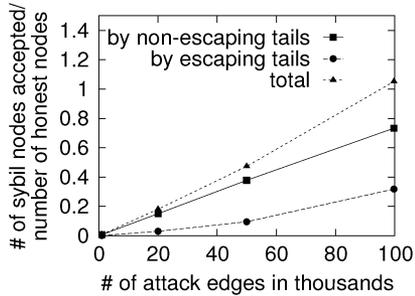


Fig. 10. Kleinberg.

intersection condition, while the second is bounded by the balance condition. In all figures, the number of sybil nodes accepted grows roughly linearly with  $g$ . The asymptotic guarantee of SybilLimit is  $O(\log n)$  sybil nodes accepted per attack edge. Figs. 7–10 show that this  $O(\log n)$  asymptotic term translates to around between 10 (in Friendster, LiveJournal, and Kleinberg) to 20 (in DBLP). As a concrete numerical comparison with SybilGuard, SybilGuard [13] uses random routes of length  $l = 1906$  in the million-node Kleinberg graph. Because SybilGuard accepts  $l$  sybil nodes per attack edge, this translates to 1906 sybil nodes accepted per attack edge for Kleinberg. Thus, numerically in Kleinberg, SybilLimit reduces the number of sybil nodes accepted by nearly 200-fold over SybilGuard.

One can also view Figs. 7–10 from another perspective. The three data sets Friendster, LiveJournal, and Kleinberg all have roughly one million nodes. Therefore, in order for the number of sybil nodes accepted to reach  $n$ , the number of attack edges needs to be around 100 000. To put it another way, the adversary needs to establish 100 000 social trust relations with honest users in the system. As a quick comparison under Kleinberg, SybilGuard will accept  $n$  sybil nodes once  $g$  reaches around 500 (since  $l = 1906$ ). Some simple experiments further show that with  $g \geq 15\,000$ , the escaping probability of the random routes in SybilGuard will be above 0.5, and SybilGuard can no longer provide any guarantees at all. Finally, DBLP is much smaller (with 100 000 nodes), and because of the slightly larger  $w$  needed for DBLP, the number of sybil nodes accepted will reach  $n$  roughly when  $g$  is 5000.

Finally, we have also performed experiments to investigate SybilLimit’s guarantees on much smaller social networks with only 100 nodes. To do so, we extract 100-node subgraphs from our social network data sets. As a concise summary, we observe that the number of sybil nodes accepted per attack edge is still around 10 to 20.

## X. CONCLUSION

This paper presented SybilLimit, a near-optimal defense against sybil attacks using social networks. Compared to our previous SybilGuard protocol [13] that accepted  $O(\sqrt{n} \log n)$  sybil nodes per attack edge, SybilLimit accepts only  $O(\log n)$  sybil nodes per attack edge. Furthermore, SybilLimit provides this guarantee even when the number of attack edges grows to  $o(n/\log n)$ . SybilLimit’s improvement derives from the combination of multiple novel techniques: 1) leveraging multiple independent instances of the random route protocol to perform many short random routes; 2) exploiting intersections on edges instead of nodes; 3) using the novel balance condition to deal with escaping tails of the verifier; and 4) using the novel benchmarking technique to safely estimate  $r$ . Finally, our results on real-world social networks confirmed their fast-mixing property and, thus, validated the fundamental assumption behind SybilLimit’s (and SybilGuard’s) approach. As future work, we intend to implement SybilLimit within the context of some real-world applications and demonstrate its utility.

## ACKNOWLEDGMENT

The authors would like to thank J. Roozenburg for allowing them to use his Friendster data set and R. Wouhaybi for allowing them to use her LiveJournal data set. The authors also would like to thank C. Lesniewski-Laas and the anonymous reviewers of Oakland’08 and IEEE/ACM TRANSACTIONS ON NETWORKING for many helpful comments on the paper.

## REFERENCES

- [1] J. Douceur, “The Sybil attack,” in *Proc. IPTPS*, 2002, pp. 251–260.
- [2] Q. Lian, Z. Zhang, M. Yang, B. Y. Zhao, Y. Dai, and X. Li, “An empirical study of collusion behavior in the Maze P2P file-sharing system,” in *Proc. IEEE ICDCS*, 2007, p. 56.
- [3] M. Steiner, T. En-Najjary, and E. W. Biersack, “Exploiting KAD: Possible uses and misuses,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 5, pp. 65–70, Oct. 2007.
- [4] “E-Mule.” [Online]. Available: <http://www.emule-project.net>
- [5] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” *ACM Trans. Prog. Languages Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [6] V. Prakash, “Razor.” [Online]. Available: <http://razor.sourceforge.net>
- [7] B. Awerbuch and C. Scheideler, “Towards a scalable and robust DHT,” in *Proc. ACM SPAA*, 2006, pp. 318–327.
- [8] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach, “Secure routing for structured peer-to-peer overlay networks,” in *Proc. USENIX OSDI*, 2002, pp. 299–314.
- [9] A. Fiat, J. Saia, and M. Young, “Making Chord robust to byzantine attacks,” in *Proc. ESA*, 2005, pp. 803–814.
- [10] E. Damiani, D. C. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante, “A reputation-based approach for choosing reliable resources in peer-to-peer networks,” in *Proc. ACM CCS*, 2002, pp. 207–216.
- [11] L. von Ahn, M. Blum, N. J. Hopper, and J. Langford, “CAPTCHA: Using hard AI problems for security,” in *Proc. IACR Eurocrypt*, 2003, pp. 294–311.
- [12] A. Ramachandran and N. Feamster, “Understanding the network-level behavior of spammers,” in *Proc. ACM SIGCOMM*, 2006, pp. 291–302.
- [13] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman, “SybilGuard: Defending against sybil attacks via social networks,” *IEEE/ACM Trans. Netw.*, vol. 16, no. 3, pp. 576–589, Jun. 2008.
- [14] M. Mitzenmacher and E. Upfal, *Probability and Computing*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [15] J. Kleinberg, “The small-world phenomenon: An algorithm perspective,” in *Proc. ACM STOC*, 2000, pp. 163–170.
- [16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, “Gossip algorithms: Design, analysis and applications,” in *Proc. IEEE INFOCOM*, 2005, vol. 3, pp. 1653–1664.

- [17] A. Flaxman, "Expansion and lack thereof in randomly perturbed graphs," Microsoft Research, Tech. Rep., 2006 [Online]. Available: <ftp://ftp.research.microsoft.com/pub/tr/TR-2006-118.pdf>
- [18] T. Anderson, "SIGCOMM'06 public review on 'SybilGuard: Defending against sybil attacks via social networks,'" 2006 [Online]. Available: <http://www.sigcomm.org/sigcomm2006/discussion/>
- [19] L. Backstrom, D. Huttenlocher, J. Kleinberg, and X. Lan, "Group formation in large social networks: Membership, growth, and evolution," in *Proc. ACM KDD*, 2006, pp. 44–54.
- [20] M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proc. Nat. Acad. Sci.*, vol. 99, no. 12, pp. 7821–7826, 2002.
- [21] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in *Proc. ACM/USENIX IMC*, 2007, pp. 29–42.
- [22] S. Wasserman and K. Faust, *Social Network Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1994.
- [23] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman, "SybilGuard: Defending against sybil attacks via social networks," in *Proc. ACM SIGCOMM*, 2006, pp. 267–278.
- [24] A. Mislove, A. Post, K. Gummadi, and P. Druschel, "Ostra: Leveraging trust to thwart unwanted communication," in *Proc. USENIX NSDI*, 2008, pp. 15–30.
- [25] G. Danezis and P. Mittal, "SybilInfer: Detecting sybil nodes using social networks," presented at the NDSS, 2009.
- [26] N. Tran, B. Min, J. Li, and L. Subramanian, "Sybil-resilient online content voting," in *Proc. USENIX NSDI*, 2009, pp. 15–28.
- [27] T. Leighton and S. Rao, "An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms," in *Proc. FOCS*, 1988, pp. 422–431.
- [28] H. Yu, C. Shi, M. Kaminsky, P. B. Gibbons, and F. Xiao, "DSybil: Optimal sybil-resistance for recommendation systems," in *Proc. IEEE Symp. Security Privacy*, 2009, pp. 283–298.
- [29] G. Danezis, C. Lesniewski-Laas, M. F. Kaashoek, and R. Anderson, "Sybil-resistant DHT routing," in *Proc. ESORICS*, 2005, Springer-Verlag LNCS 3679, pp. 305–318.
- [30] R. Bazzi and G. Konjevod, "On the establishment of distinct identities in overlay networks," in *Proc. ACM PODC*, 2005, pp. 312–320.
- [31] T. S. E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches," in *Proc. IEEE INFOCOM*, 2002, vol. 1, pp. 170–179.
- [32] K. Walsh and E. G. Sirer, "Experience with an object reputation system for peer-to-peer filesharing," in *Proc. USENIX NSDI*, 2006, p. 1.
- [33] A. Cheng and E. Friedman, "Sybilproof reputation mechanisms," in *Proc. ACM P2PEcon*, 2005, pp. 128–132.
- [34] M. Feldman, K. Lai, I. Stoica, and J. Chuang, "Robust incentive techniques for peer-to-peer networks," in *Proc. ACM Electron. Commerce*, 2004, pp. 102–111.
- [35] J. Newsome, E. Shi, D. Song, and A. Perrig, "The Sybil attack in sensor networks: Analysis & defenses," in *Proc. ACM/IEEE IPSN*, 2004, pp. 259–268.
- [36] B. Parno, A. Perrig, and V. Gligor, "Distributed detection of node replication attacks in sensor networks," in *Proc. IEEE Symp. Security Privacy*, 2005, pp. 49–63.
- [37] N. B. Margolin and B. N. Levine, "Informant: Detecting sybils using incentives," in *Proc. Financial Cryptography*, 2007, pp. 192–207.
- [38] P. Maniatis, M. Roussopoulos, T. J. Giuli, D. S. H. Rosenthal, and M. Baker, "The LOCKSS peer-to-peer digital preservation system," *ACM Trans. Comput. Sys.*, vol. 23, no. 1, pp. 2–50, Feb. 2005.
- [39] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao, "SybilLimit: A near-optimal social network defense against sybil attacks," in *Proc. IEEE Symp. Security Privacy*, 2008, pp. 3–17.
- [40] J. Leskovec, K. Lang, A. Dasgupta, and M. Mahoney, "Statistical properties of community structure in large social and information networks," in *Proc. Int. Conf. WWW*, 2008, pp. 695–704.
- [41] T. Lindvall, *Lectures on the Coupling Method*. New York: Dover, 2002.
- [42] I. Abraham and D. Malkhi, "Probabilistic quorums for dynamic systems," in *Proc. DISC*, 2003, pp. 60–74.
- [43] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao, "SybilLimit: A near-optimal social network defense against sybil attacks," School of Computing, Nat. Univ. Singapore, Tech. Rep. TRA2/08, Mar. 2008 [Online]. Available: <http://www.comp.nus.edu.sg/~yuhf/sybillimit-tr.pdf>
- [44] Z. Bar-Yossef, R. Kumar, and D. Sivakumar, "Sampling algorithms: Lower bounds and applications," in *Proc. ACM STOC*, 2001, pp. 266–275.

- [45] J. Roozenburg, "A literature survey on Bloom filters," Delft Univ. Technol., 2005, unpublished.
- [46] R. H. Wouhaybi, "Trends and behavior in online social communities," Intel Corp., Hillsboro, OR, 2007, unpublished.
- [47] "Dataset: DBLP." [Online]. Available: <http://kdl.cs.umass.edu/data/dblp/dblp-info.html>



**Haifeng Yu** received the B.E. degree from Shanghai Jiaotong University, Shanghai, China, in 1997, and the M.S. and Ph.D. degrees from Duke University, Durham, NC, in 1999 and 2002, respectively.

He is currently an Assistant Professor with the Department of Computer Science, National University of Singapore. Prior to joining National University of Singapore, he was a Researcher with Intel Research Pittsburgh, Pittsburgh, PA, and an Adjunct Assistant Professor with the Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA. His current research interests cover distributed systems, distributed systems security, distributed algorithms, and distributed systems availability. More information about his research is available at <http://www.comp.nus.edu.sg/~yuhf>.



**Phillip B. Gibbons** (M'89) received the B.A. degree in mathematics from Dartmouth College, Hanover, NH, in 1983, and the Ph.D. degree in computer science from the University of California at Berkeley in 1989.

He is currently a Principal Research Scientist at Intel Labs Pittsburgh, Pittsburgh, PA. He joined Intel after 11 years at (AT&T and Lucent) Bell Laboratories, Murray Hill, NJ. His research interests include parallel/distributed computing, databases, and sensor systems, with over 120 publications. More information about his research is available at <http://www.pittsburgh.intel-research.net/people/gibbons/>.

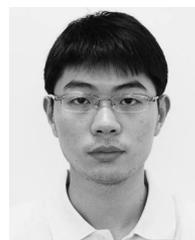
Dr. Gibbons is a Fellow of the Association for Computing Machinery (ACM). He has served as an Associate Editor for the *Journal of the ACM*, the IEEE TRANSACTIONS ON COMPUTERS, and the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS. He has served on over 40 conference program committees, including serving as Program Chair/Co-Chair/Vice-Chair for the SPAA, SenSys, IPSN, DCOSS, and ICDE conferences.



**Michael Kaminsky** received the B.S. degree from the University of California at Berkeley in 1998, and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 2000 and 2004, respectively.

He is currently a Research Scientist at Intel Labs Pittsburgh, Pittsburgh, PA, and an Adjunct Research Scientist at Carnegie Mellon University, Pittsburgh, PA. He is generally interested in computer science systems research, including distributed systems, networking, operating systems, and network/systems security. More information about his research is available at <http://www.pittsburgh.intel-research.net/people/kaminsky/>.

Dr. Kaminsky is a Member of the Association for Computing Machinery (ACM).



**Feng Xiao** received the B.E. degree from the University of Science and Technology of China, Hefei, China, in 2007.

He is currently a graduate program student with the Department of Computer Science, National University of Singapore. His research interests cover distributed system security, distributed computing, and peer-to-peer systems. More information about his research is available at <http://www.comp.nus.edu.sg/~xiaof>.