Trust in Social Network-based Sybil Defenses

Abedelaziz Mohaisen, Nicholas Hopper, and Yongdae Kim
Computer Science and Engineering Department, University of Minnesota
{mohaisen, hopper, kyd}@cs.umn.edu—17th ACM CCS, Chicago, IL October 4-8, 2010.

UNIVERSITY OF MINNESOTA

Contribution

We propose designs to account for trust in social graphs used for Sybil defenses. We model trust as modified random walks. Our designs are motivated by the observed relationship between the algorithmic property required for the defenses to perform well and a hypothesized trust value in the underlying social graphs.

Designs to Account for Trust

Two designs weigh locality of trust, by weighting the current or originator nodes high, while two weigh differential trust among neighbors.

Lazy Random walks: each node captures the random walk with probability α or a neighbor uniformly with probability $\frac{1-\alpha}{\deg(v_i)}$.

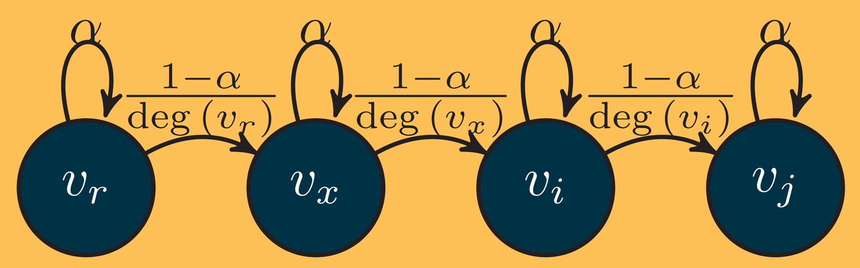

$$\mathbf{P}' = \alpha\mathbf{I} + (1-\alpha)\mathbf{P},\ \pi = \left[\frac{\deg(v_i)}{2m}\right] \quad (1)$$

Originator-biased Random Walk: each node over any path returns the random walk to its originator with probability α or follow the normal protocol with probability $1-\alpha$.

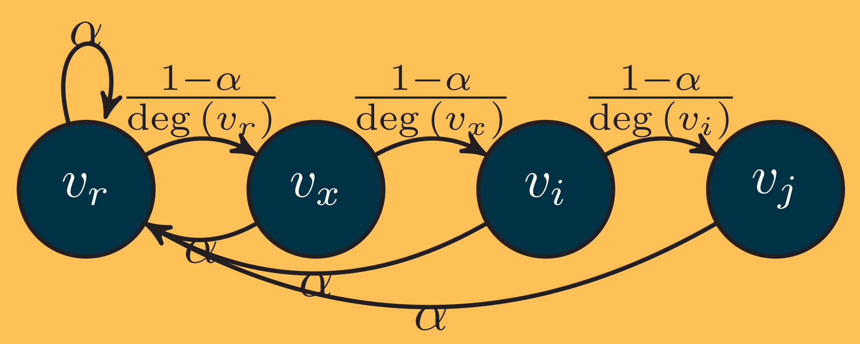

$$\mathbf{P}' = \alpha\mathbf{A}_r + (1-\alpha)\mathbf{P}, \pi = [\pi_i]^{n\times 1}. \quad (2)$$

$$\pi_i = \begin{cases} (1-\alpha)\frac{\deg(v_i)}{2m} & v_i \in V \setminus \{v_r\} \\ \alpha + \frac{\deg(v_i)}{2m} & v_i = v_r \end{cases} \quad (3)$$

Similarity-biased Walk: uses the cosine measure, a graph-theoritic similarity measure to determine how close are nodes to each other

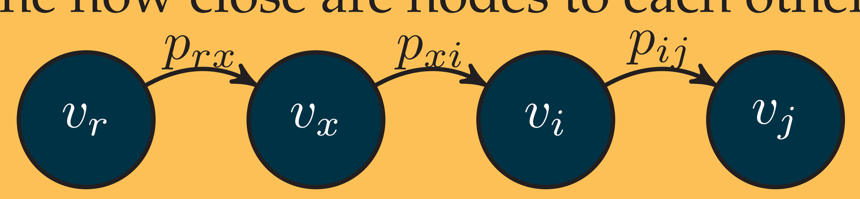

$$\mathbf{P}' = [p_{ij}]^{n\times n} = \mathbf{D}^{-1}\mathbf{S},\ \mathbf{D} = \text{diag}(\sum_{k=1}^{n} \mathbf{S}_{ik}) \quad (4)$$

$$\mathbf{S} = [s_{ij}]^{n\times n}, s_{ij} = \frac{\mathbf{a}_i \cdot \mathbf{a}_j}{|\mathbf{a}_i|_2 |\mathbf{a}_j|_2} \text{ iff } v_i \sim v_j \quad (5)$$

$$\pi_i = (\sum_{z=1}^{n} s_{zi})(\sum_{j=1}^{n}\sum_{k=1}^{n} s_{jk})^{-1} \quad (6)$$

Interaction-biased Walk: similar to the similarity-biased random walk, but weighing the frequency of interactions between nodes.

$$\mathbf{P}' = \mathbf{D}^{-1}\mathbf{B},\ \mathbf{D} = \text{diag}(\sum_{k=1}^{n} \mathbf{B}_{ik}) \quad (7)$$

B's entries are observed locally by each node (interactions) and π is computed same as in (6).

Mixed Random Walks: use a combination of the four different walks.

Defenses Model

- $G = (V, E)$ is undirected and unweighted social graph, where $|V| = n, |E| = m$.
- **A** is an adjacency matrix and **P** is a transition matrix defined as row norm of **A**.

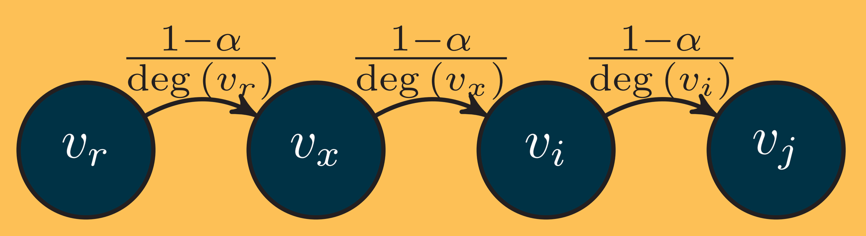

- Probability of walks landing on v_i after walk length of the mixing time of G is proportional to $\deg(v_i)$, i.e., $\pi = [\frac{\deg(v_i)}{2m}]$.

Defenses Assumptions

- Fast mixing social graph: $w = O(\log(n))$.
- Strong trust in social graphs.
- Hard to establish edges with Non-Sybil.
- Number of attack edges is limited.
- Sparse-cut between Sybil and Non-Sybil.

SN-based Sybil Defenses

Include SybilLimit, SybilGuard, SybilInfer, etc.

- Each node performs random walks and sample/register last node in the walk from "close to" stationary distribution.
- Suspects send their "authentic" last nodes address to the verifier, which also has a verification list.
- The verifier checks for collision in the verifier/suspect sampled nodes (probability guaranteed by the birthday paradox).
- The attacker has a limited number of nodes (edges) to register on.
- Probability of honest node's random walk ending in the dishonest region is bounded by the walk length w.
- Guarantees: SybilLimit and SybilGuard allow w Sybil identities per attack edge.
- The parameter w depends on the mixing time of the social graph.

Results

• We first measure the mixing time of different social graphs, and observe a relationship between the mixing time and the level of knowledge (trust) in the underlying graph.
• Then, we learn the impact of the different proposed designs on the mixing time. We show that parameters associated with the different mixing models for characterizing trust control the mixing time.
• Finally, we learn the cost of SybilLimit to accept all non-Sybil nodes in some social graphs, under varying parameters for the different designs. While these proposed designs characterize trust, we show that trust—once it is incorporated into the Sybil defense—comes at some cost. The datasets (n/deg/μ) are below.

Physics 1(4.2K /3.23/ 0.998133) • Youtube (1.1M/2.63/ 0.997972) • Facebook (63.4K/12.87/ 0.998133) • Facebook A (1M/20.35/0.982477) • Wiki-vote (7.1K /14.256/0.899418) • Livejournal B (1M/27.56/0.999695 • Physics 2 (11.2K /10.50/0.998221) • DBLP (615K/1.88/ 0.997494) • Livejournal A (1M/26.15/0.999387) • Physics 3 (8.6K/2.87/0.996879) • Enron (33.7K/5.37/ 0.996473).

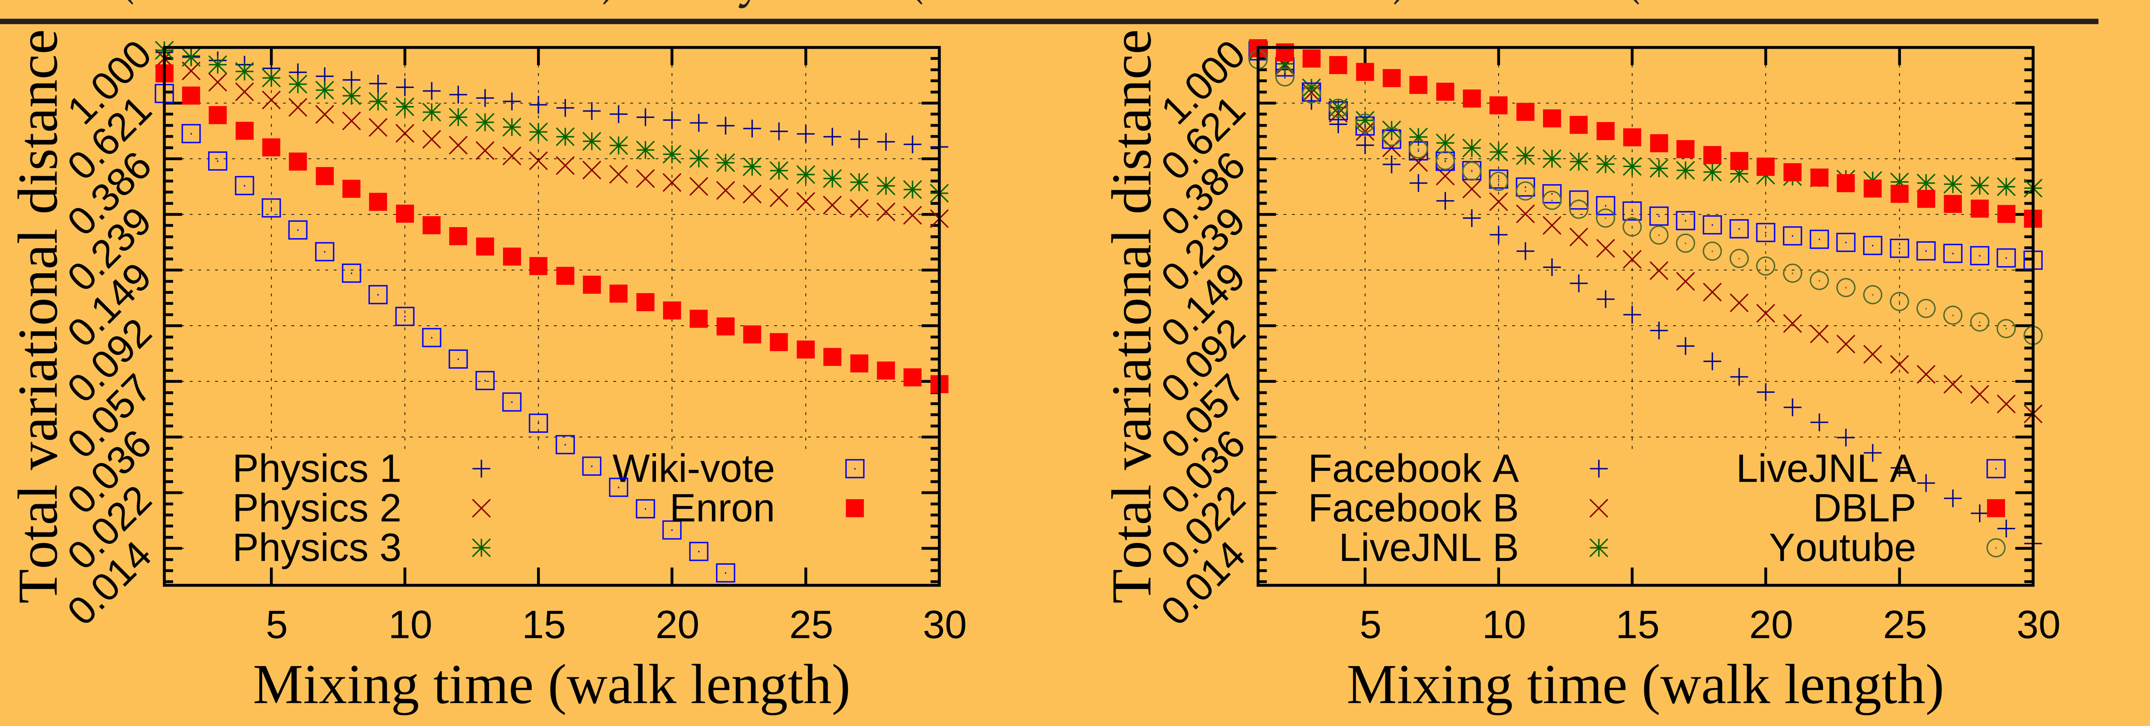

Measuring the mixing time of social graphs (small networks on left, large networks on right)

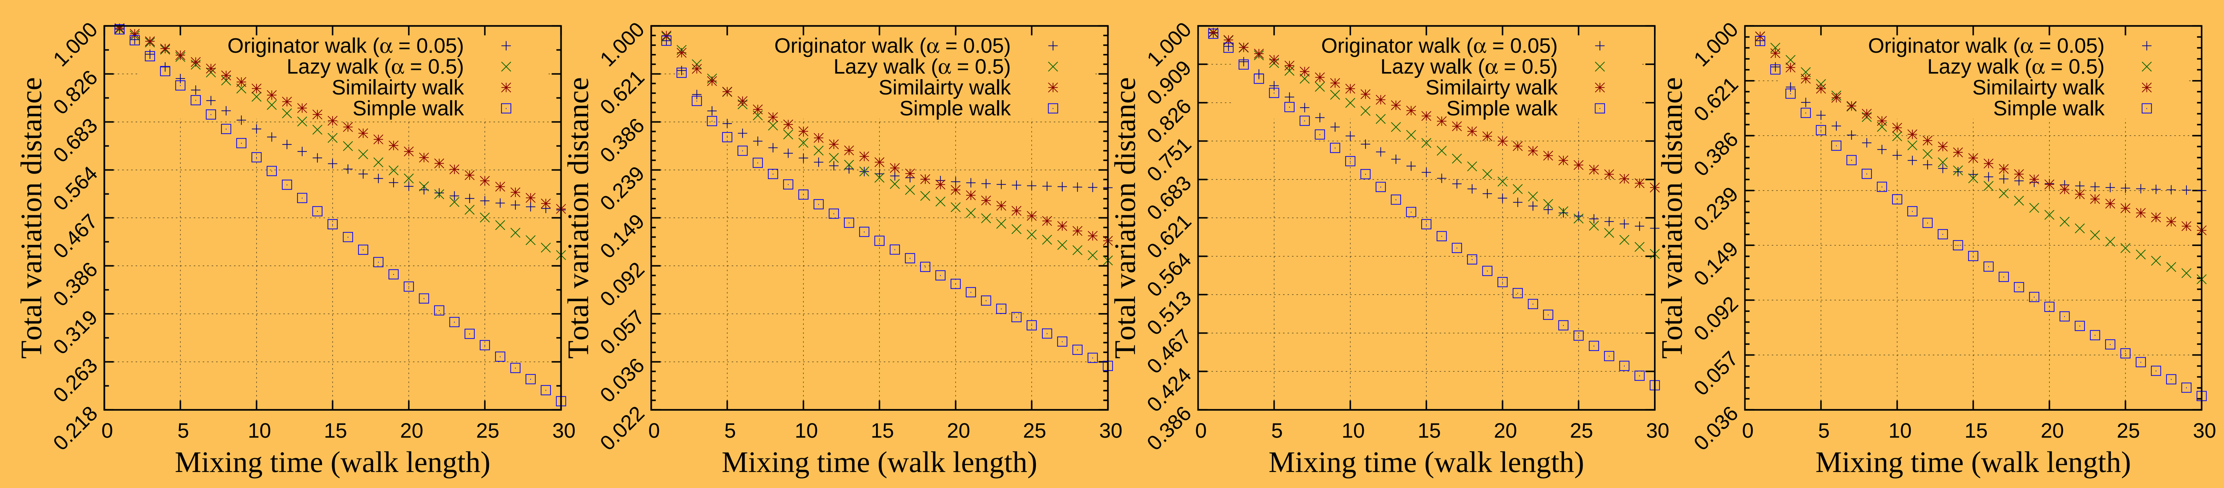

Accounting for trust impacts the mixing time (DBLP, Facebook, Physics, and Livejournal)

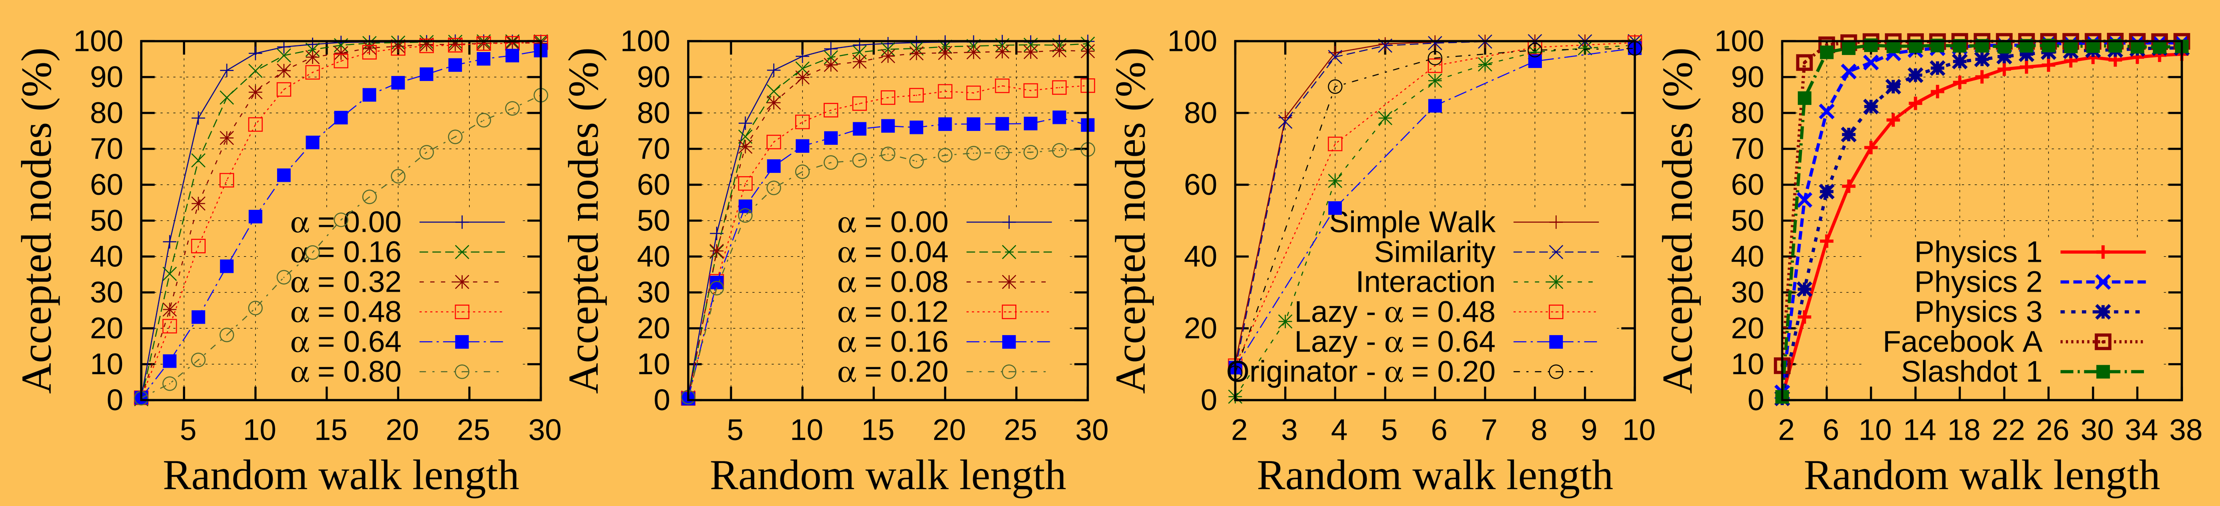

Incorporating trust into social graphs impacts Sybil defenses (Lazy, Originator—DBLP)

References

[1] A. Mohaisen, A. Yun, Y. Kim. Measuring the mixing time of social graphs In *USENIX/ACM SIGCOMM Internet Measurements Conference*, November 2010.

[2] A. Mohaisen, N. Hopper, Y. Kim. Keep your friends close: Incorporating trust in social network-based Sybil defenses, In *Technical Report*, UMN, August 2010.

Funding

This research was supported by the NSF grant CNS-0917154 and a grant from KAIST.