

Solving Internet Problems via Social Networks

Motivation

- Internet faces many security problems
 - DoS, spams, phishing, malware, sybil attacks
- Difficult to solve within today's Internet
 - Recall: grade "F" for accountability
- What do we do?
 - Approach 1: Patching today's Internet ("arms race")
 - Approach 2: Clean-slate redesign of the Internet
 - Approach 3: Looking for help outside the Internet

This lecture: How social networks can help us

Social Networks

- Online social networks
 - MySpace, Facebook, LinkedIn, Twitter, Flickr, Orkut, LiveJournal ...
- SN defined by personal interaction
 - E.g., social network for E-mails
 - Can be defined either explicitly through address book, or implicitly through past (non-spam) email exchanges

KEY: Each edge is associated with certain trust

3

SybilGuard: Defending Against Sybil Attacks via Social Networks

Haifeng Yu, Michael Kaminsky,
Phillip B. Gibbons, Abraham Flaxman

Based on SIGCOMM'06 talk slides by Haifeng Yu

Etymology of Sybil

Sybil is a well known character of the 70s, a woman possessed with multiple personality disorder, of 16 characters



5

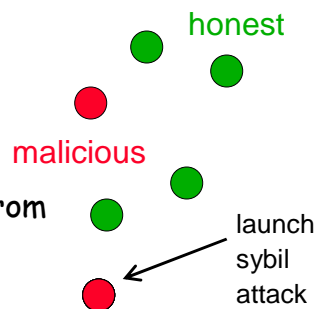
Background: Sybil Attack

- **Sybil attack:** Single user pretends many fake/sybil identities

- Creating multiple accounts from different IP addresses

- Sybil identities can become a large fraction of all identities

- Out-vote honest users in collaborative tasks



6


Background: Defending Against Sybil Attack

- Using a trusted central authority
 - Tie identities to actual human beings
- Not always desirable
 - Can be hard to find such authority
 - Sensitive info may scare away users
 - Potential bottleneck and target of attack
- Without a trusted central authority
 - Impossible unless using special assumptions [Douceur'02]
 - Resource challenges not sufficient -- adversary can have much more resources than typical user

7

SybilGuard Basic Insight: Leveraging Social Networks

Our Social Network Definition




- Undirected graph
- Nodes = identities
- Edges = **strong** trust
 - E.g., colleagues, relatives

8

SybilGuard Basic Insight

- n honest users: One identity/node each
- Malicious users: Multiple identities each (sybil nodes)

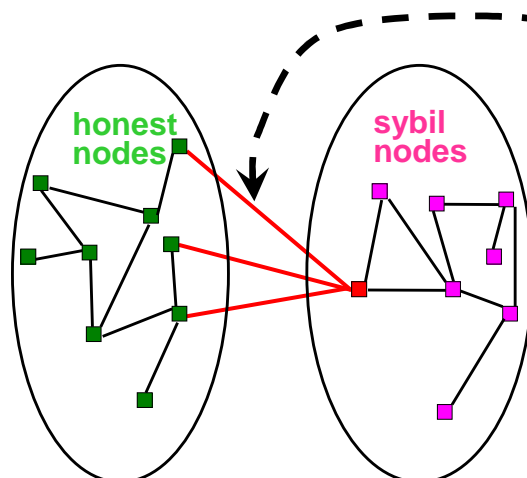


Sybil nodes may collude – the adversary

Observation: Adversary cannot create extra edges between honest nodes and sybil nodes

9

SybilGuard Basic Insight



Dis-proportionally small cut disconnecting a large number of identities

But cannot search for such cut brute-force...

10

Outline

- √ Motivation and SybilGuard basic insight
- Overview of SybilGuard: Random routes
- Properties of SybilGuard protocol
- Evaluation results
- Conclusions


11

Goal of Sybil Defense

- Goal: Enable a *verifier* node to decide whether to **accept** another *suspect* node
 - **Accept**: Provide service to / receive service from
 - Idealized guarantee: An honest node accepts and only accepts other honest nodes
- SybilGuard:
 - Bounds the number of sybil nodes accepted
 - Guarantees are with high probability
 - Approach: Acceptance based on **random route intersection** between verifier and suspect

12

Random Walk Review



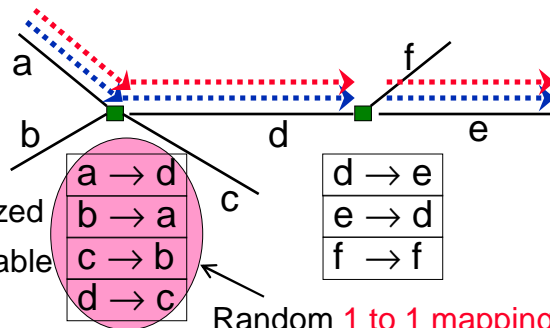
pick random edge d

pick random edge e

pick random edge c

13

Random Route: Convergence



randomized
routing table

a → d
b → a
c → b
d → c

d → e
e → d
f → f

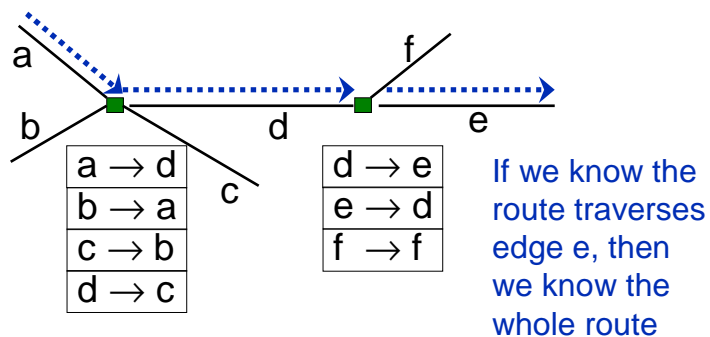
Random **1 to 1 mapping** between incoming edge and outgoing edge

Using routing table gives Convergence Property:

Routes merge if crossing the same edge

14


Random Route: Back-traceable



Using 1-1 mapping gives Back-traceable Property:
Routes may be back-traced

15

Random Route Intersection: Honest Nodes



- Verifier accepts a suspect if the two routes intersect
 - Route length w : $\sim \sqrt{n} \log n$
 - W.h.p., verifier's route stays within honest region
 - W.h.p., routes from two honest nodes intersect

16

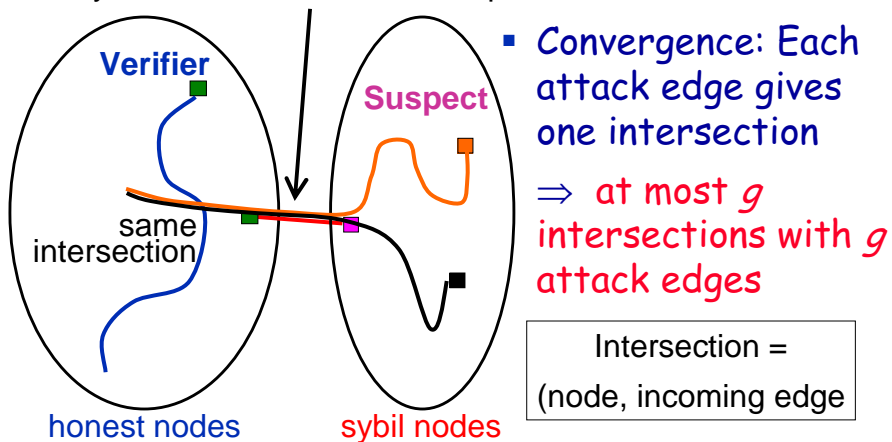
Random Route Intersection: Sybil Nodes

- SybilGuard bounds the number of accepted sybil nodes within $g^* w$
 - g : Number of attack edges
 - w : Length of random routes
- Next ...
 - Convergence property to bound the **number of intersections** within g
 - Back-traceable property to bound the **number of accepted sybil nodes per intersection** within w

17


Bound # Intersections Within g

must cross attack edge to intersect even if sybil nodes do not follow the protocol



18

Bound # Sybil Nodes Accepted per Intersection within w



- Back-traceable: Each intersection should correspond to routes from **at most w honest nodes**
- Verifier accepts at most w nodes per intersection
 - Will not hurt honest nodes

19

Summary of SybilGuard Guarantees

- Power of the adversary:
 - *Unlimited* number of *colluding* sybil nodes
 - Sybil nodes may not follow SybilGuard protocol
- W.h.p., honest node accepts $\leq g^* w$ sybil nodes
 - g : # of attack edges
 - w : Length of random route

If SybilGuard bounds # accepted sybil nodes within	Then apps can do
$n/2$	byzantine consensus
n	majority voting
not much larger than n	effective replication

20

Outline

- √ Motivation and SybilGuard basic insight
- √ Overview of SybilGuard
 - Properties of SybilGuard protocol
 - Evaluation results
 - Conclusions

21

SybilGuard Protocol

- **Security:**
 - Protocol ensures that nodes cannot lie about their random routes in the honest region
- **Decentralized:**
 - No one has global view
 - Nodes only communicate with direct neighbors in the social network when doing random routes

22

SybilGuard Protocol (continued)

- Efficiency: Random routes are performed only once and then "remembered"
 - No more message exchanges needed unless the social network changes
 - Verifier incurs $O(1)$ messages to verify a suspect
- User and node dynamics:
 - Different from DHTs, node churn is a non-problem in SybilGuard ...
- See paper for all the details ...

23

Evaluation Results

- Simulation based on synthetic social network model [Kleinberg'00] for 10^6 , 10^4 , 10^2 nodes
- With 2500 attack edges (i.e., adversary has acquired 2500 social trust relationships):
 - Honest node accepts honest node with 99.8% prob
 - 99.8% honest node properly bounds the number of accepted sybil nodes
 - See paper for full results ...

24

SybilGuard Summary

- Sybil attack: Serious threat to collaborative tasks in decentralized systems
- SybilGuard: Fully decentralized defense protocol
 - Based on random routes on social networks
 - Effectiveness shown via simulation and analysis
- Follow-up work: SybilLimit [IEEE S&P 08]
 - Much better guarantee
honest node accepts $\leq g * \log(n)$ sybil nodes
 - g : number of attack edges
 - Close to optimal -
any algorithm would accept at least $O(g)$ sybil nodes

25

Defending Against SPAM via Social Networks

Applied in Two Ways

- Reduce false positives of spam filtering by white-listing friends and friends of friends
 - RE: Reliable Email [NSDI'06]
 - Infer friendship from non-spam email exchanges
 - Protocol preserves privacy
 - Can reduce FP by 87% (71% direct, 16% FoF)
- Social spam filtering
 - E.g. TrustMyMail.com
 - Infer social network from past email exchanges
 - Compute trust metric based on distance in social network
 - Filter SPAM if trust metric is too low

27

Any Other Applications of SN?

The Journey of CS386M

- **Networking review (3 weeks)**
 - A background equalizer
- **Network design principles (4.5 weeks)**
 - Signaling, data vs. control, hard vs. soft state, randomization, indirection, multiplexing, virtualization, design for scale
 - Internet design: end-to-end principle, rethinking the design
- **Current practice & research (4.5 weeks)**
 - Living with today's Internet
 - Network security (overview + worm fingerprinting)
 - Network operations & management (overview + network anomaly detection)
 - Details of technologies: MPLS + VPN + traffic engineering
 - Addressing the limitations of the Internet
 - Clean-slate redesign
 - External help: social networks
- **Project ideas + presentations (2.5 weeks)**

29

Thank you!

- **A great learning experience for me!**
 - Learned a lot of new materials
 - Thank you all for your thought-provoking questions
- **Feedback highly appreciated!**
 - How to better organize the class?
 - How to improve my teaching in general?
 - Please give me feedback via either CIS or email
- **Next 2 weeks: show time!**
 - Look forward to your presentations & reports

30