Paper presentation, by Hong Liu

SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks

Haifeng Yu, Phillip B. Gibbons
Michael Kaminsky, Feng Xiao

In 2008. IEEE Symposium on Security and Privacy

Contributions

to give O(log n) bounds on Sybil attacks using graph properties

- no centralized trusted users
- not trust propagation
- simple protocol

Assumptions

- Established social network G(n,m)
  - unstructured
  - undirected
  - fast-mixing: mixing time O(log N)
  - Sybils penetrate the network through ”attack edges”
- Nodes identified by their public keys

warming up

- random walk on graphs (see the whiteboard)
  - equilibrium distribution
  - mixing time
  - fast mixing graphs
    - theorem 1
- reference
  - A Sinclair, Improved bounds for mixing rates of markov chains and multicommodity flow, Combinatorics, Probability & Computing, 1:351–370, 1992.

Get ready for Sybil attacks

- Each honest node only knows its neighbors
- Sybils know the entire graph
- Sybils try to slip into the honest zone by fooling the verifiers
- Sybils are byzantine

Basic operation in SybilLimit

- random route on SybilLimit
  - one-to-one mapping from incoming edges to outcoming edges
  - fixed length w=O(log n)
  - head is registered by the tail node
  - new head rewrites old head for a single tail
- and … no more!

SybilLimit protocol

only execute once until the graph changes !!

- for suspects
  - do random routes and register keys at the end nodes – invoke r independent instances (s-instances)
  - all suspects share the same r s-instances
- for verifiers
  - do random route and record the tails – also invoke r independent instances (v-instances)
  - all verifiers share the same r v-instances
  - do the verification on
    - the intersection condition
    - the balance condition

Why does this work?

- For a given tail, the route led to it is determined. There are no more routes that end up with the same tail.
  - to guarantee that honest suspects satisfy the intersection condition: Birthday Paradox
  - to bound the number of Sybils per attack edge: the balance condition

Why balance condition?

- Adversary can register up to g*r*w public keys (r*w per attack edge)
- For each attack edge, w Sybils enter the honest zone as honest suspects do.
- Remaining Sybils only reply on escaping routes to get more intersections with the verifier's tails set.
- Adversary modifies random route info to have more Sybils accepted, which increases the loads of the tails nearby the attack edges.
- bound the Sybil num within O(g*w)

Estimating r

- Benchmarking technique
  - set up a benchmark set by repeating random route and adding the end nodes.
  - starts from r = 1, doubles r until most nodes in the benchmark set are accepted.
- does not overestimate r
- underestimation r does not degrade defence performance

O(log n) is the lower bound

- fast mixing time O(log(n))
- for any g in [1,n], G(n,m) fast mixing, it is always possible to introduce c*g Sybils into the graph s.t. the augmented graph is still fast mixing

Limits of SybilLimit

- undirected unweighted graph
- a honest network exists already – no bootstrap stage
  - newcomers obtain links out-of-band
- favorable to newcomers with many links; unfavorable to those with few links
- the network must be fast mixing
- cannot handle user churns well