

Does Overlay Routing Security Require Admission Control?

Chris Lesniewski–Laas and M. Frans Kaashoek

`{ctl,kaashoek}@mit.edu`

November 7, 2004

Routing Overlays

- Function:
 - starting from node A ...
 - contact node B responsible for ID $target$.
 - In an Internet overlay, generally boils down to finding B 's IP address.
- Key layer of:
 - DHTs (thus filesystems, caches, I3, etc.)
 - P2P routing protocols, multicast overlays
 - censorship-resistant or anonymity systems
 - etc.

Attacks

- Impersonation / forgery
 - Self-certifying IDs
- Misdirection: prevent A from finding B .
 - M : "sorry, nobody is closer to $target$ than me."
 - Attack can be targeted if bad guy can choose ID.

Previous Approaches

- Byzantine Fault Tolerance
 - Not efficient, but can guarantee correctness if less than 1/3 of node IDs are bad.
- Castro et al, OSDI2002
 - central ID-assignment authority
 - constrained routing tables
 - density-based routing failure test
 - improves security when <25% of node IDs are bad.

The Counting Model

- Goal of previous approaches:

Any two nodes A and B should be able to communicate via the overlay if the adversary controls fewer than M out of the N total nodes.

- Can always be thwarted by an adversary who can manufacture IDs!

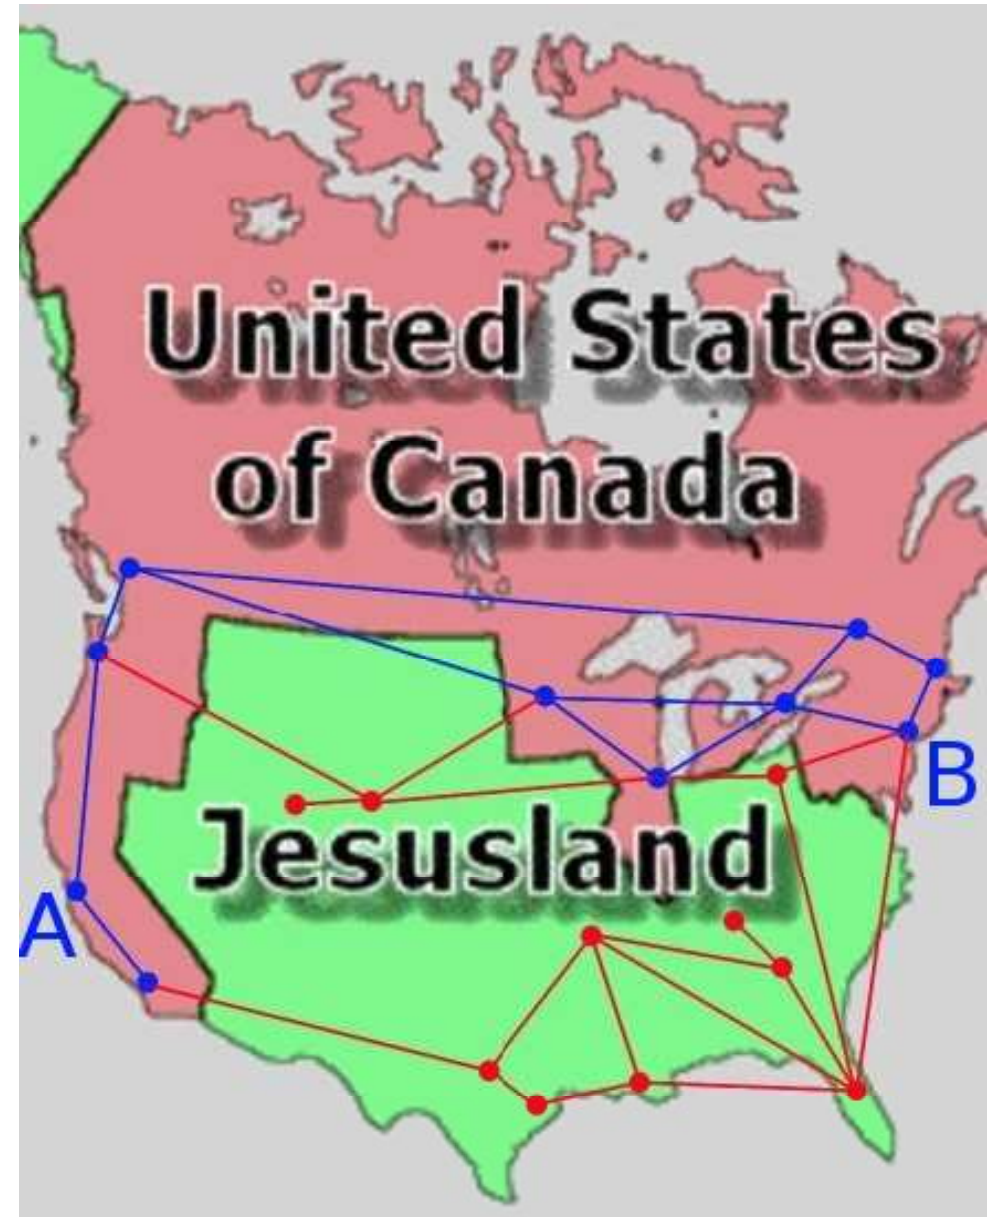
The Sybil Tarpit

- Requires centralized gatekeeper to keep out the unverified masses.
- How does it decide who to admit? (too restrictive vs. too permissive)
 - passports, letterhead (like Verisign?)
 - money (oops: RIAA, Diebold. Hello, moral hazard.)
 - crypto puzzles (oops: zombie networks)
 - IP addresses (oops: MIT, zombies, IPv6)
- Doesn't help small overlays.
- Centralized gatekeeper goes against the grain of what we're trying to accomplish with *decentralized* systems.

The Bootstrap Graph Model

- Set of out-of-band initial links
- If entire initial set colludes, you're doomed anyway.
- Security Criterion:

Any two nodes A and B should be able to communicate via the overlay if there exists a path between A and B, in the bootstrap graph, consisting entirely of good nodes.



Scalability Criterion

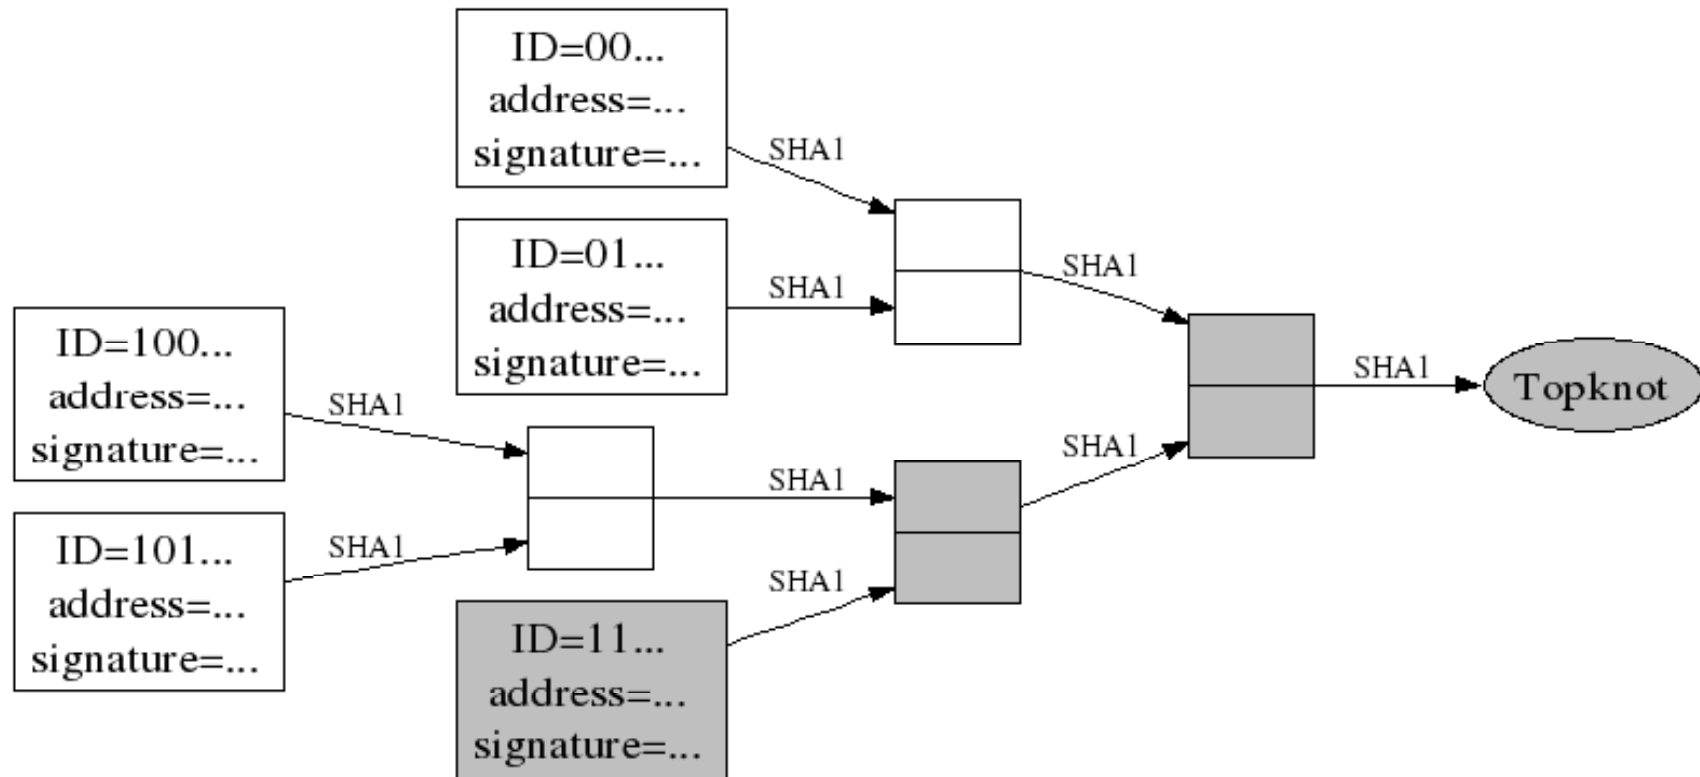
- Normal-case behavior (when not under attack) should be scalable.
- System shouldn't amplify DoS attacks.
- Scalability Criterion:

The average bandwidth and storage cost to a node with degree D in the bootstrap graph should be $O(D \log^k N)$.

- Per: search, maintenance cycle
- **Local** costs for bad decisions
 - incentives in the right places!

Topknot

- Secure routing protocol for a **treelike** bootstrap graph.
- Augments structured routing table with a certified hash tree
 - Hash tree is used to verify the correctness of every hop.



Summary

Bootstrap graph:

Out-of-band links via which nodes joined the overlay.

Security criterion:

A can talk to B **iff** there is a good path in the bootstrap graph.

Scalability criterion:

Damage inflicted by bad nodes is localized.

The upshot:

The bootstrap graph model enables protocols to defeat the Sybil attack without centralized admission control.