

CERIAS

The Center for Education and Research in Information Assurance and Security

Attacking and Improving the Tor Directory Protocol

Zhongtang Luo¹, Adithya Bhat¹, Kartik Nayak², Aniket Kate^{1,3}
Purdue University¹, Duke University², Supra Oracle³

Appearing at IEEE S&P 2024



What is the Tor Directory Network?

- The Tor network enhances clients' privacy by routing traffic through an overlay network of volunteered intermediate relays.
- Tor employs a distributed protocol among nine hard-coded **Directory Authority (DA) servers** to securely disseminate information about these relays to produce a new consensus document every hour.

Cool. Why is it vulnerable?

- The Tor network itself does not defend against attacks on the relay list (e.g. Sybil relays, relays with irregular information). Therefore, **all defense relies on external audits**.
- Tor uses an **outdated consensus system** that uses two rounds of broadcast...

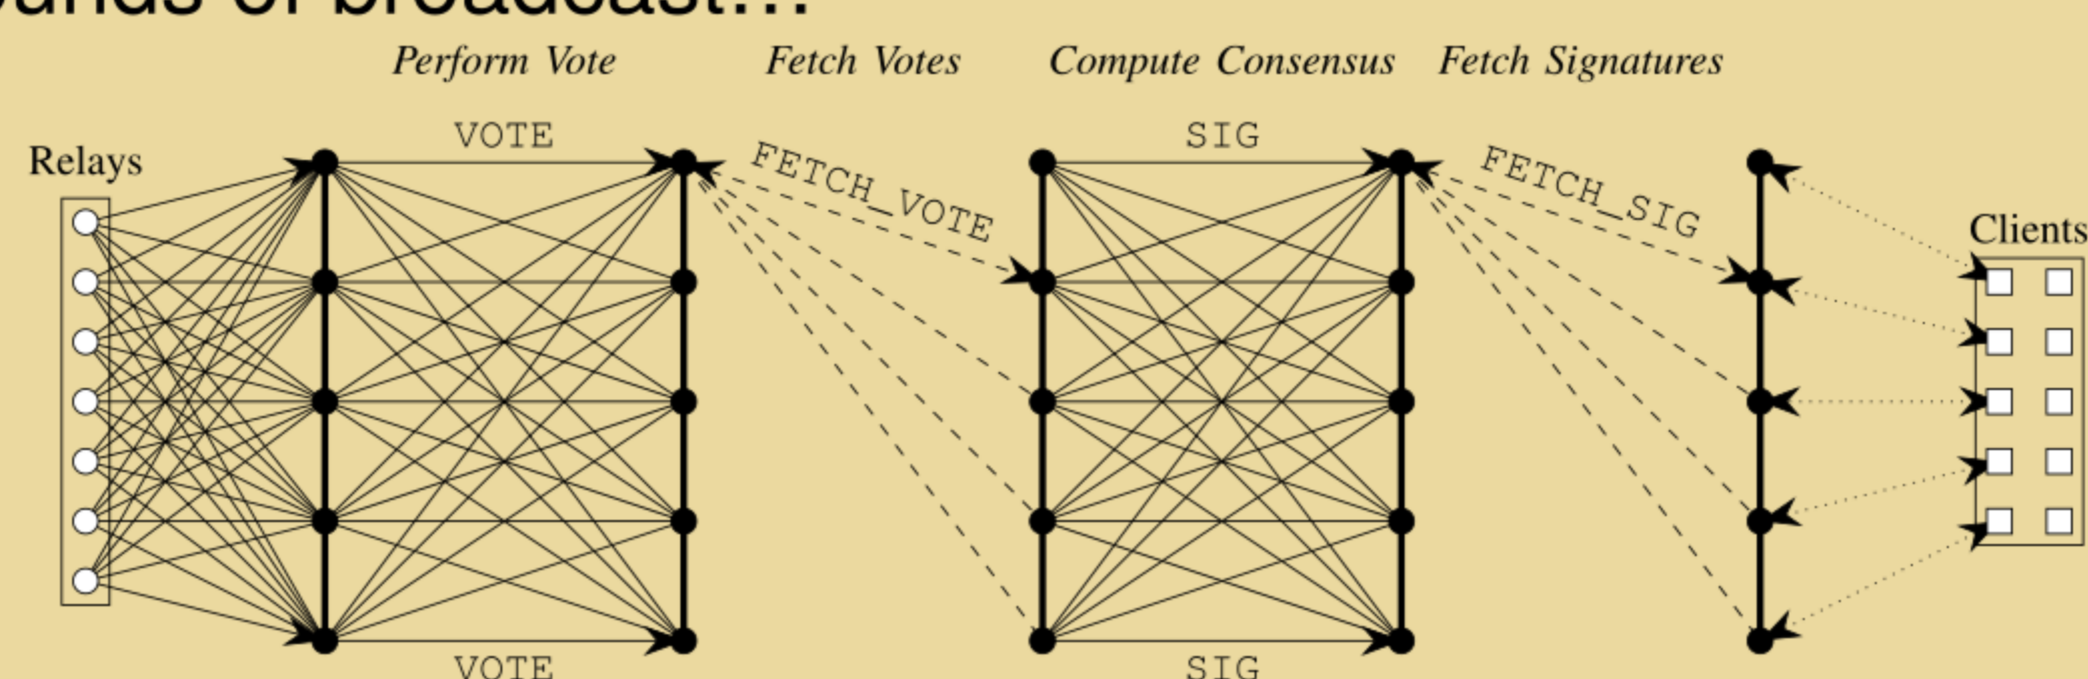


Figure: Two rounds of all-to-all broadcast (and very little else) happen within the procedure.

This is vulnerable to an equivocation attack!

How can we attack the protocol?

An attacker needs to...

- Play nice with half of the authorities.
- Lie to the other half of the authorities and **inject some incorrect information on the relay**.

He can then run away with an **incorrect relay list signed by a majority of the authorities** without being found!

Paper



Poster



That sounds very convoluted. What is so bad about an incorrect relay list?

```
r test010r kNeiqbQsrPh/JPuJiTrcz1bNDTY Nf2VyvkI...
2022-04-05 17:27:05 127.0.0.1 5010 0
.....
w Bandwidth=14597871
.....
-----BEGIN SIGNATURE-----
KtR7wLvXNtat1Kly71bjJVyWp9gwuPbggnQYBdZI8dWLM7M...
.....
-----END SIGNATURE-----
```

```
Apr 05 13:27:20.657 [warn] A consensus needs 5 good
signatures from recognized authorities for us to
accept it. This ns one has 2 (test003a test004a).
7 (test005a test000a test006a test002a test007a
test008a test001a) of the authorities we know
didn't sign it.
```

Figure: A demonstration of the attack from an experiment. Note the very large bandwidth 14597871 (although in a very small font).

The attacker can use incorrect parameters (e.g. very large bandwidth) to attract users to use only his relays, which **totally breaks the anonymity** without anyone finding out about it.

How should we fix it?

We provide two fixes:

- Patch the consensus health monitor so that it includes an equivocation detection mechanism

Sender	Receiver
moria1	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
tor26	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
dizum	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
gabelmoo	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
dannenberg	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
maatuska	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
longclaw	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet
bastet	moria1 tor26 dizum gabelmoo dannenberg maatuska longclaw bastet

Already online and working!

- Patch the protocol so that it is a robust consensus protocol

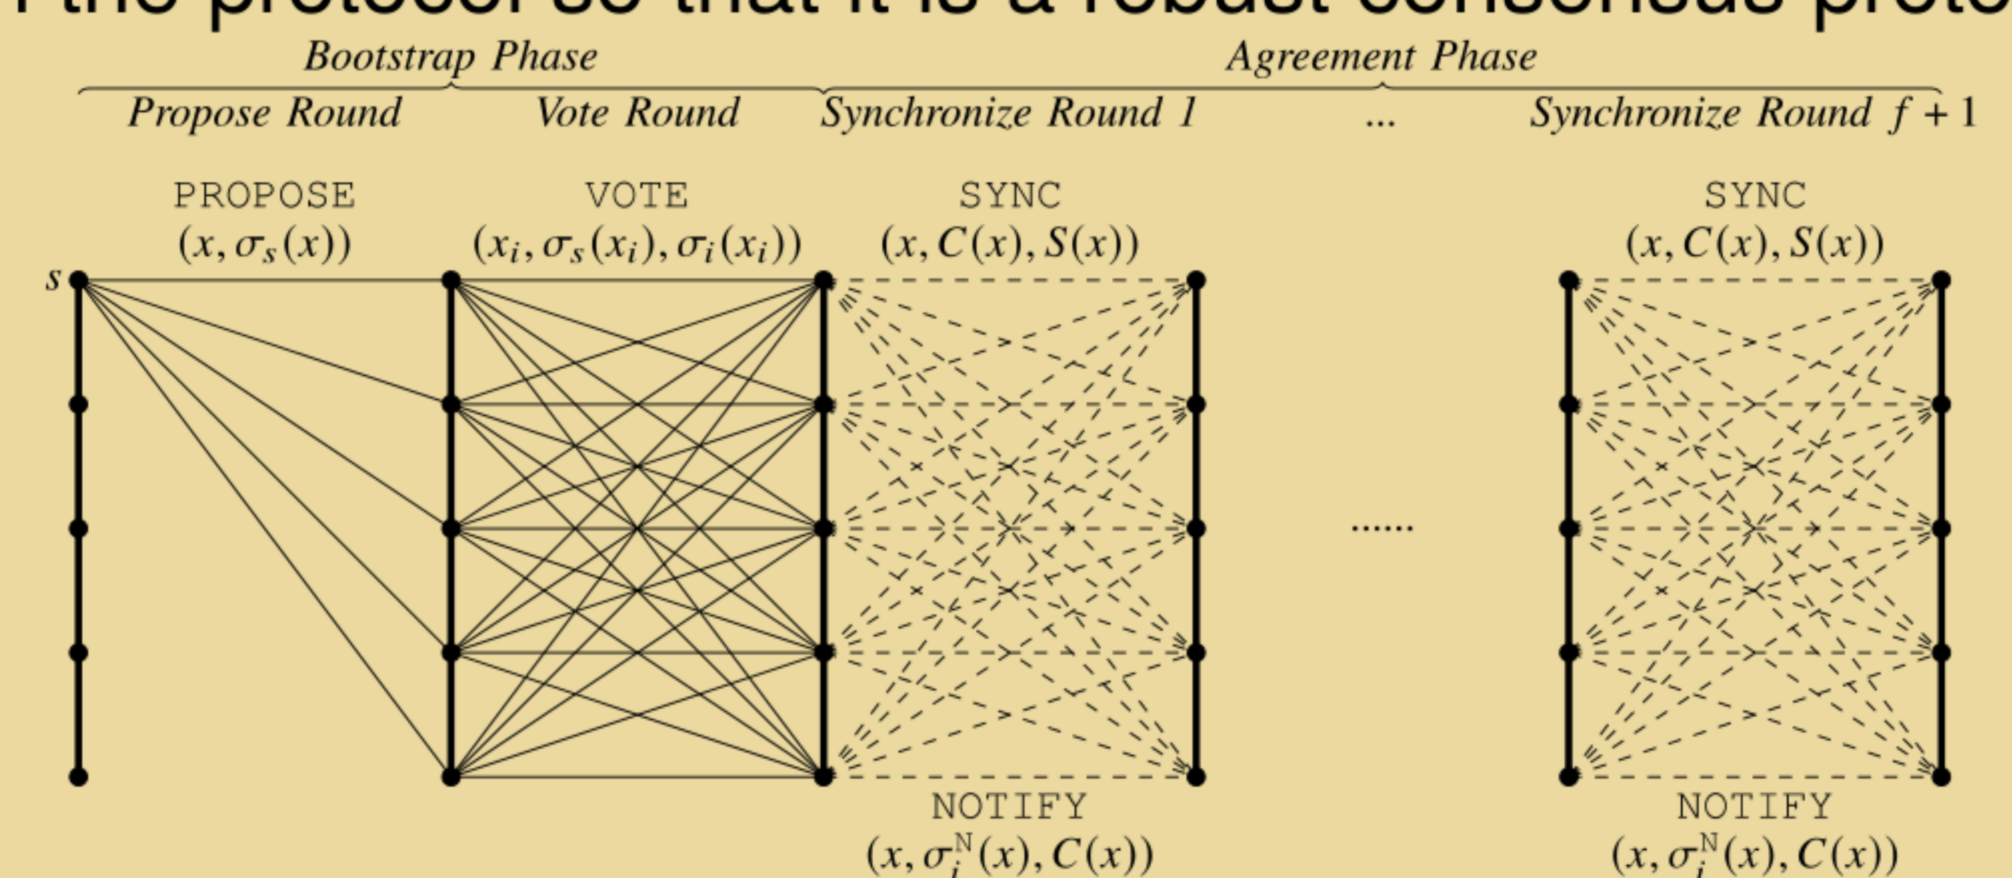


Figure: Inspired by the famous Dolev-Strong protocol, we design a protocol that secures the directory protocol.




Figure 8: Average network and CPU time for different phases of our protocol. PP - Bootstrap Phase. CP - Agreement Phase. Figure 9: Network throughput in three test scenarios, measured in consensus per hour. Figure 10: Prediction of throughput up to 10,000 relays. The protocol can generate up to 500 consensus documents per hour.

Comparable performance with the original protocol!