Assessing or incentivising correct mixing without authorities

Jeff Burdges

We describe a probabilistic assessment technique for mix networks that employs cover traffic roughly similar to Loopix [2], which similarly provides a scheme to incentivise correct packet routing.

Anonymity systems like Tor rely upon the altruism of relay operators and sometimes their financial doners. There is ocasional discussion around providing incentives to relay operators, but asking payments from users destroys anonymity by reducing the user base, aka anonymity set, even assuming an anonymous payment system with sufficient performance and market penetration. At the same time, anonymity systems do assess routers for some combination of performance, capacity, and reliability, but do so using centralised infrastructure.

We propose a decentralised probabilistic sampling technique for assessment of, or incentives for, mix relay availability and reliability. Incentives are funded by inflation of the rewards token, so users need not pay.

Any mix network has mix clients send covert traffic, but we require a mix network in which at least some mix relays, or mixes, send cover traffic too. In Loopix [2], mix nodes send cover traffic messages that traverse each strata, eventually looping back to the first strata, and continuing back to their sending mix. These mix relay generated loops significantly reduces the advantage an adversary gains by monitoring mix nodes, see Theorem 1 vs 2 in §4.1.3 of [2]. They also enable defenses against active attacks like n-1 attacks, see §4.2.1 of [2].

At some level, our sampling procedure resembles Ouroboros Praos [1] except, instead of producing blocks only when winning a verifiable random function (VRF) contest, all our VRF outputs produce cover packets, and the winners are selected in some future epoch.

In epoch i, we have a limited number of "sampler" or "staked" mixes produce cover traffic using random seeds provided by a verifiable random function (VRF), with the VRF key $V = vG$ registered in some prior epoch.

$$p_{V,i,j} := VRF_v(r_i || j)$$

These VRFs are themselves seeded by $r_i || j$ where i denotes the current epoch, r_i denotes the value of some collaborative random beacon in the ith epoch, and $j < j_{\max}$ is a counter that ranges from zero to number of eligible cover traffic packets per epoch.

We expect all messages are either routed through the mix network or dropped within k epochs. In epoch $i+k$, we hold a lottery using the collaborative randomness r_{i+k}, in which a few cover traffic packets by several of these sampler mixes' win, and then reveal the seeds for their winning packets.

$$H(p_{V,i,j} || r_{i+k}) < \texttt{difficulty}$$

At this point, anyone may recompute the private key and route of the winning packets from the seed $p_{V,i,j}$ and global network consensus. We have ensured the routes taken by these packets are hard to bias because the winning pool of mixes is limited, and they each produce only $j_{\max}$ their packets using VRFs. We argue these winning packets provide an sufficiently unbiased random sample of mixes' routing behavior.

We also require that mixes in a winning packet's route provide a proof, from a commitment to their in coming and out going packets, that they forwarded the winning cover traffic paket message. If all do so, then all receive credit for correct routing, but none receive credit if any drop the packet.

We assess mixes primarily to provide useful information to clients, especially each mixes' chances for dropping packets. We learn about dropped packets only if V publish their winning $p_{V,i,j}$ though, even when the cover packet itself got dropped.

We therefore always credit the sampler V for their winning packets, even dropped packets. In this way, samplers have reason to publish $p_{V,i,j}$, even when doing so degrades some mix node.

In terms of anonymity, we impact the defenses provided by the loop cover traffic, as well as protections afforded by total traffic. We think active defenses are fully utilised before the lottery, but we require tha only a few packets win so that our sampling procedure only minimally increases an adversary's knowledge about mix fullness.

We need the shared random values r_i to have minimal adversarial bias, ala malicious mining. We want to adapt ideas from the Ouroboros Praos [1] analysis to determine how many winners we require to limit adversarial bias acceptably, although unbiased collaborative randomness schemes like RandHerd or VDFs work too.

1. REFERENCES

[1] B. David, P. Gaži, A. Kiayias, and A. Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake protocol, 2017. https://eprint.iacr.org/2017/573.

[2] A. M. Piotrowska, J. Hayes, T. Elahi, S. Meiser, and G. Danezis. The loopix anonymity system. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1199–1216, 2017.