Reward Sharing for Mixnets

Claudia Diaz, Harry Halpin, Aggelos Kiayias
Nym Technologies SA

https://nymtech.net/nym-cryptoecon-paper.pdf

What is a mixnet?

- Type of overlay anonymous communication network

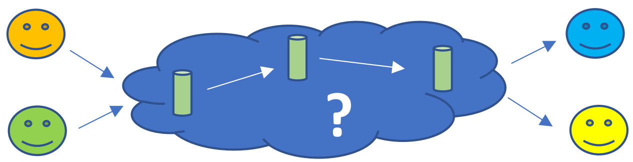

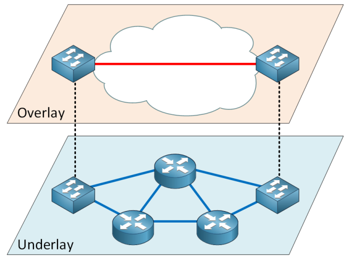

- Multi-hop, layered encryptions, source-routed

- Packet-based, per-mix reordering of packet flows (different from OR)

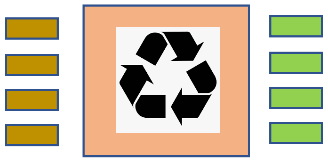

- Nym mixnet: layered structure, uniform routing

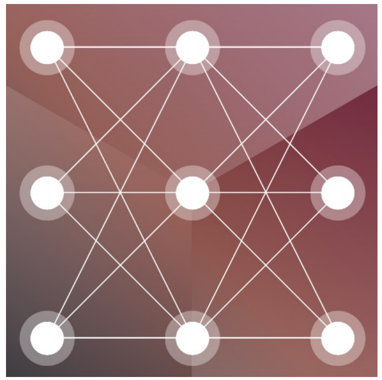

Why incentivized?

- **Scalability** : mixnet can add nodes to meet arbitrarily large user demand
  - Volunteer-operated networks : inelastic pool of volunteers to bear operational costs
  - Incentivized : extra income can fund growth needed to serve increased demand
    - **Market** for consuming/providing private bandwidth

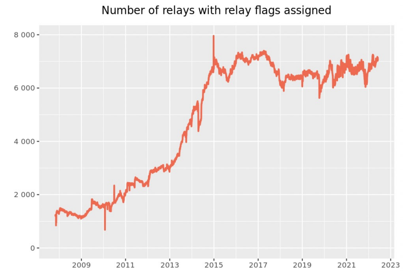

- Scale with good **quality of service** (low packet loss)
  - P2P architectures where all users are also providers for others do not work
  - Distinguish profesionalized providers (paid for the work) and consumers (pay for the service)
    - Privacy for consumers; verifiability and transparency for providers (intermediaries)

- **Goal** of incentives : populate sufficiently big mixnet with reliable mix nodes
  - The number of mix nodes that is *sufficient* depends on service demand (traffic load)
  - Nodes compete on quality : select well-performing mix nodes and weed out weak nodes

Mitigate Sybil attacks

- At least one intermediary node must be honest to provide privacy to a communication
  - If adversary controls all the intermediaries: can reconstruct path and link sender to receiver

- How to prevent the adversary from fully capturing a significant amount of routes?
  - Volunteer networks + variable node capacity : adversary setting up high-bandwidth nodes can route (and deanonymize) a large fraction of paths
  - Uniform routing (same resources required from all nodes) removes the high-bandwidth advantage (forcing adversary to set up more nodes)
  - Longer routes (more mixnet layers) : impact on latency and resources
  - … how to raise the cost of Sybil attacks and select nodes for the mixnet in a *decentralized* manner

- Given an excess of mix node candidates competing to provide the service:
  - Allow all stakeholders to signal which mix node they want to endorse for active service provision
  - Select mix nodes for service provision proportionally to their stakeholder support
  - In addition to setting up nodes, the Sybil adversary now needs to either become itself a major stakeholder (expensive) or gather support from many stakeholders for each of its Sybil nodes (effortful)

Stake as reputation

- “Stakeholder support” for mix nodes must be meaningful
  - Limited supply: nodes compete for stakeholder support
  - Incentivize stakeholders to support “best nodes” for the network:
    - **Reliability and performance**: high uptime, no packet loss
    - **Cost effectiveness**
    - Trust in the operator : node lifetime, **operator stake**, history of engagement and contributions to the ecosystem, geolocation, donation to a good cause, endorsements
- “Reputation” is represented by the total stake associated to a node
  - Includes stake bonded by the operator to register the node and stake delegated from other stakeholders to support the node
  - Reputation maxes out when a *stake saturation point* is reached
    - Prevent stake from over-concentrating on too few nodes, ensure stakeholders spread their support over sufficient nodes

Reputation-based selection of nodes

- The mixnet is periodically (hourly) reconstituted : sample fresh set of nodes to route packets for the next time period
  - Nodes are selected with probability proportional to their reputation
  - Additional selection of **standby set** to incentivize spare capacity and allow for fast mixnet growth

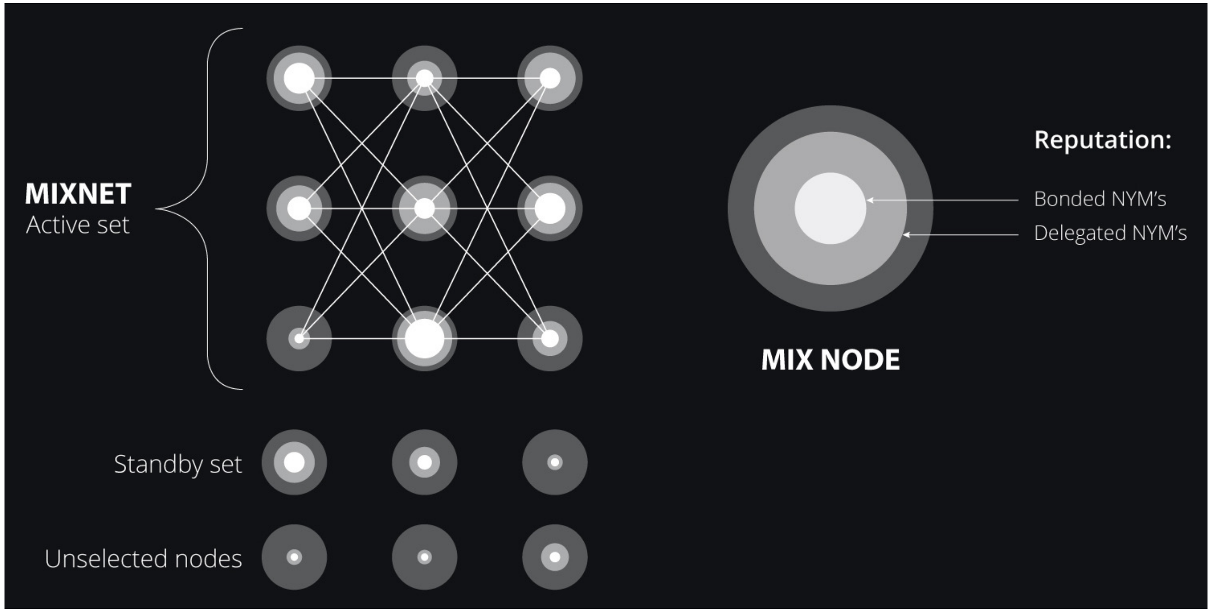

Mix nodes are rewarded based on performance and reputation

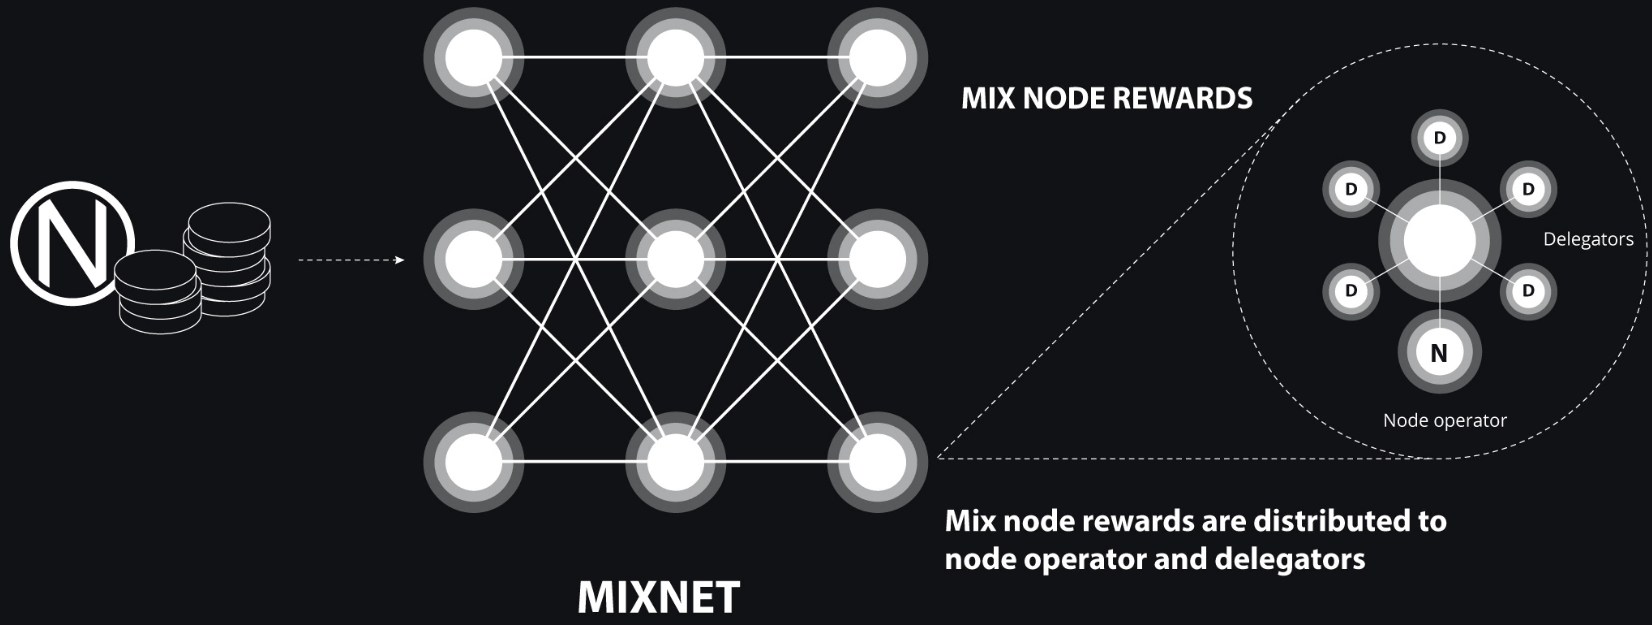

Enables decentralized decision-making

- No centralized entity making or executing decisions
  - Which nodes should be part of the network
  - How much they are rewarded for their work

- Collective decision-making by stakeholders requires:
  - All participants have access to all the relevant network information
  - Ability to verify the authenticity and integrity of data and operations

- Blockchain
  - Public record of: node registrations (keys, addresses), network parameters, staking state, node performance measurements, etc.
  - Smart contracts for network management, reward algorithms
  - Integrity, availablity, governance mechanisms for updating software / parameters

Bootstrapping reserve

- Chicken and egg problem:
  - Anonymity grows with the user base
    - Little incentive to pay at the start and thus no initial income to fund operations
  - Low quality of service at the start (due to poor funding) precludes usage growth

- Initial funding needed to support infrastructure while usage picks up
  - Part of the token supply is locked in a reserve that provides initial rewards
    - Released gradually over time
  - After some years: income from user fees needed to sustain network operations
  - Somewhat similar to Bitcoin mining/fees (though with important differences)

Nym economic model

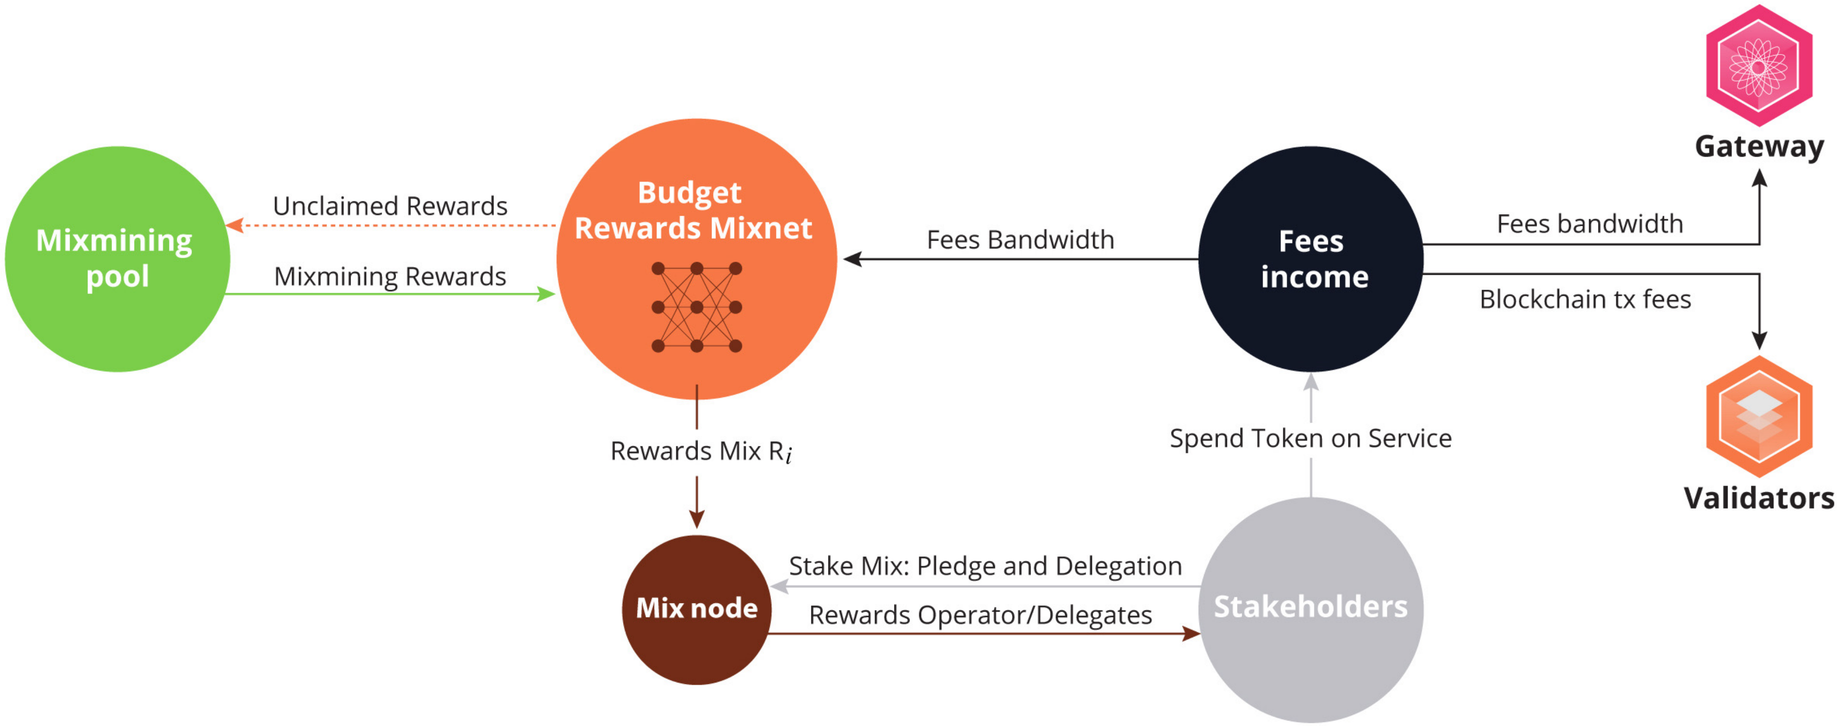

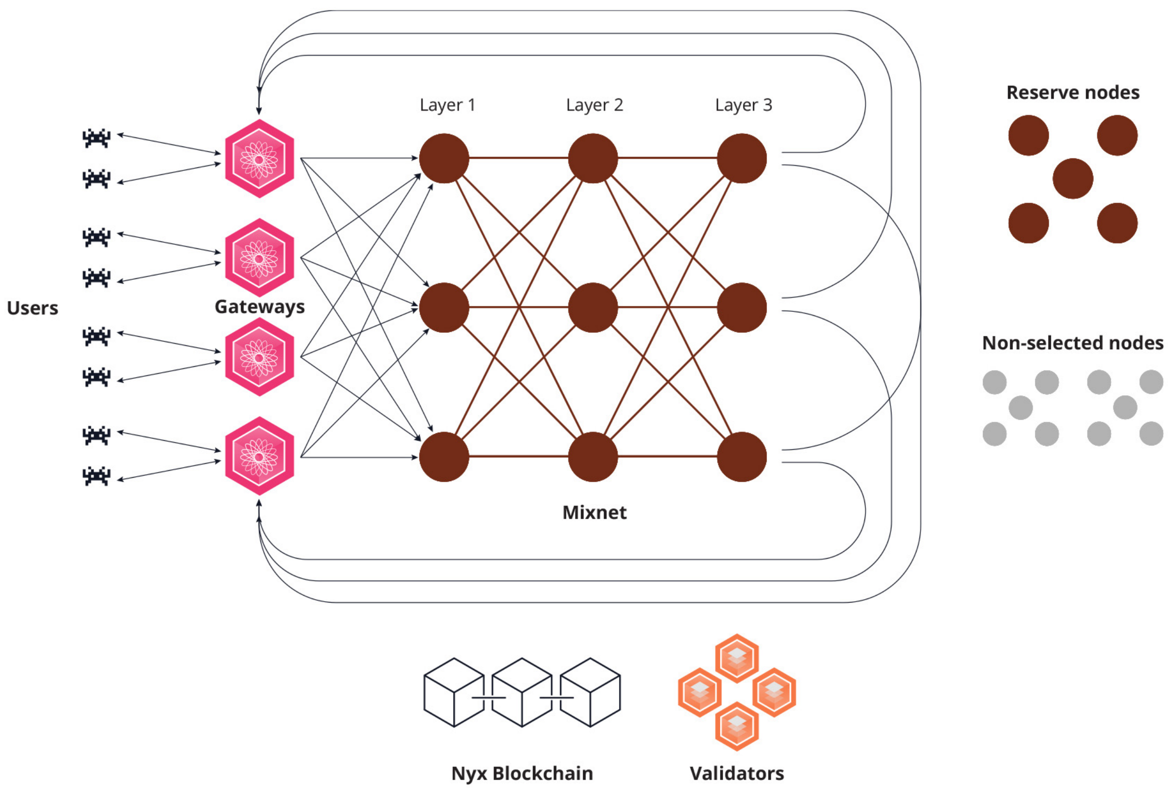

- Validators:
  - Function: maintain the blockchain, network state, execute smart contracts
  - Third-party service paid by blockchain transaction fees from all participants
  - Nyx chain: anyone can write general-purpose Web Assembly smart contracts
    - Can support (and be paid for) any other services (not exclusive to Nym mixnet)

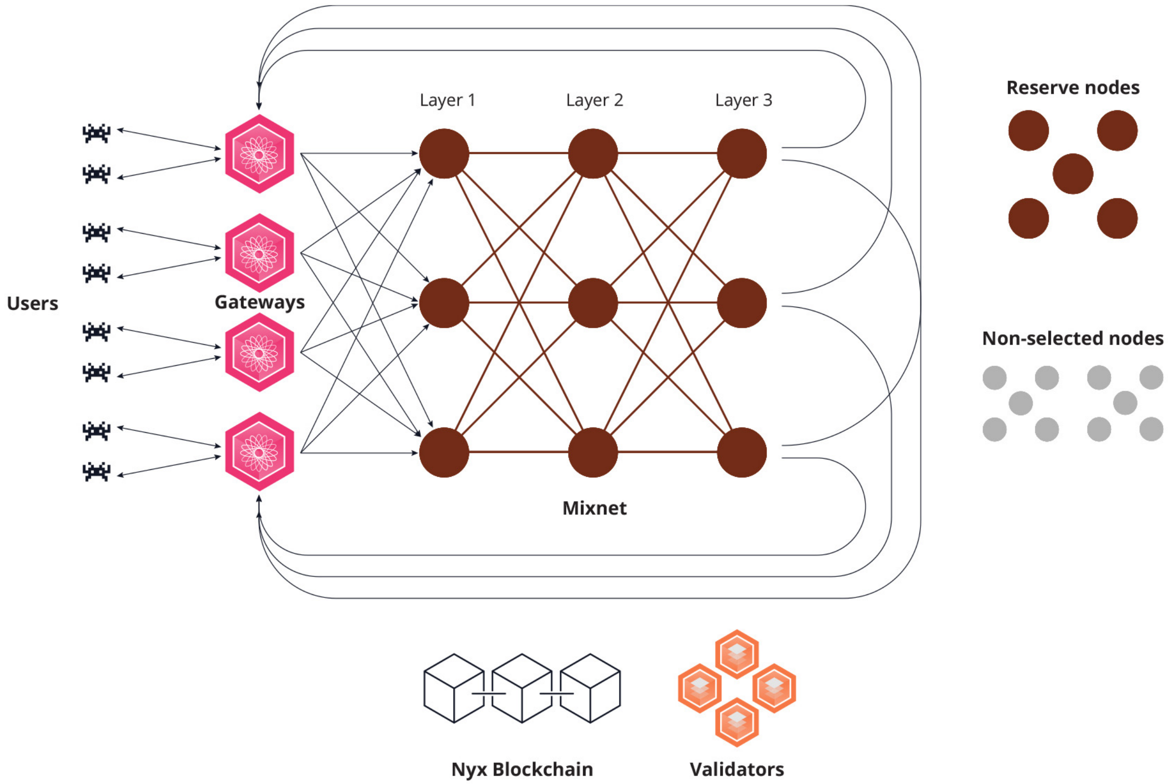

- Gateways:
  - Function: interface between users and mixnet, collecting user payments, forwarding packets, caching received packets, censorhip circumvention access
  - Chosen by the user rather than automatically assigned (unlike nodes in route)
  - Paid by a fraction of the bandwidth fees
    - Compete for users, may offer additional services

Components of the reward scheme (1)

1. Node registration by any stakeholder
   - operator bond (pledge), node cost, profit margin
2. Delegation of stake to a registered node to increase its reputation
   - maxes out at the "stake saturation point" (disincentives to stake more)
   - stake saturation point = available staking supply / target number of nodes (K)
3. Selection of nodes for the mixnet
   - sampling K nodes without replacement, weighed by (capped) reputation
   - **active set**: populate L layers of width W, sufficient to serve demand (first LW)
   - **standby set**: spare capacity to allow for fast mixnet growth (next K-LW)
     - rewarded at a lower rate than active nodes
   - unselected nodes: not rewarded for the epoch

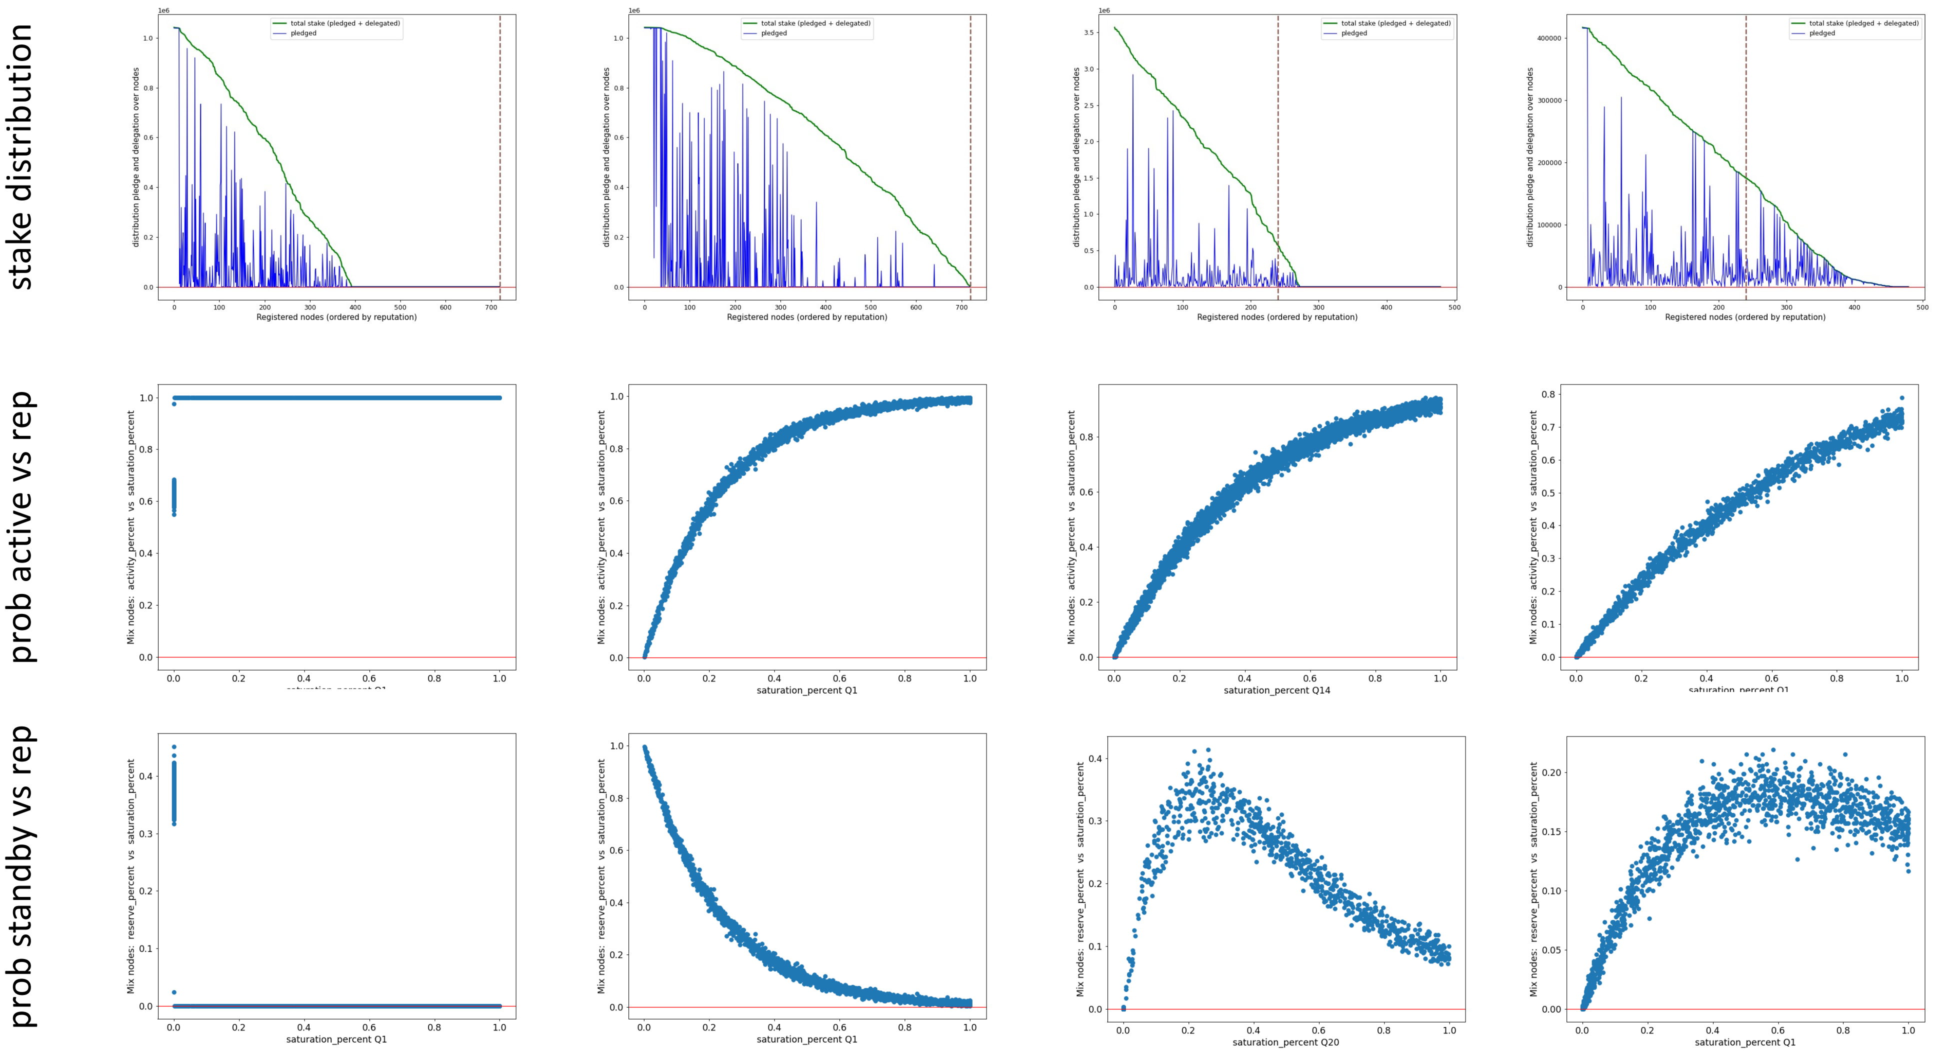
stake distribution
prob active vs rep
prob standby vs rep
total stake (pledged + delegated)
pledged
distribution pledge and delegation over nodes
Registered nodes (ordered by reputation)
Mix nodes: activity_percent vs saturation_percent
saturation_percent Q1
saturation_percent Q14
Mix nodes: reserve_percent vs saturation_percent
saturation_percent Q20

Components of the reward scheme (2)

4. Node performance measurements (a whole topic by itself)
    - Decentralized solution: "secret shoppers" to sample node performance
    - Placeholder solution: validators send test packets through all nodes
    - Result: performance score $\boldsymbol{\rho_i}$ for each node (value between zero and one, representing estimated fraction of correctly routed packets)
5. Reward budget
    - Mixmining emission schedule:
        - 25% of token (250m) locked in the "mixmining reward pool"
        - each month 2% of reserve is made available for rewards (5m in the first month)
        - unallocated rewards are fed back to the reserve (softens exponential decay)
    - Bandwith fees:
        - dynamic posted price approach considering node operational costs
        - computed to cover operational costs plus a system-wide profit fee $\boldsymbol{\tau}$
6. Distribution of rewards:
    - Algorithm to distribute rewards to nodes: performance, reputation, active/standby, operator bond
    - Algorithm to distribute node rewards among the node operator and delegates: cost, profit margin

Distribution of rewards to nodes

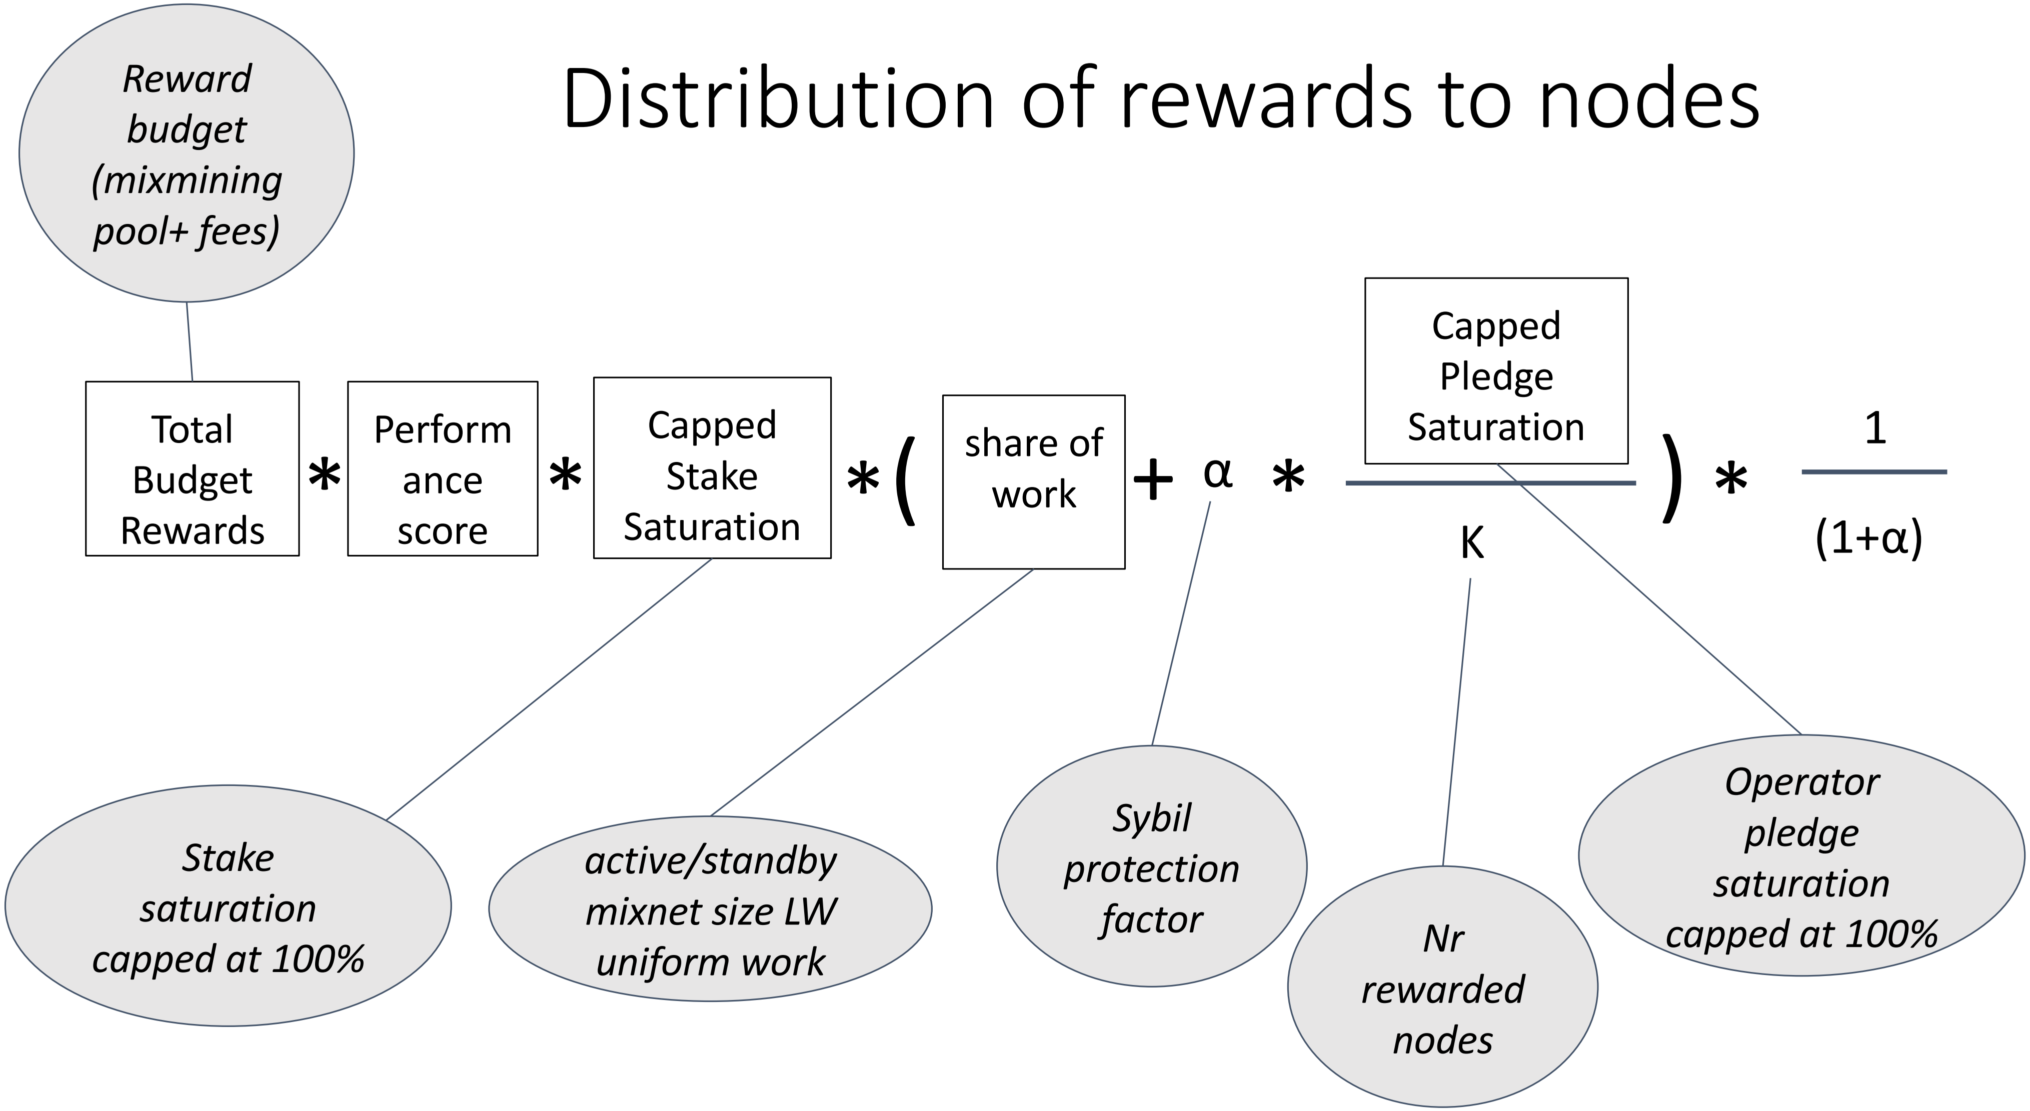

Properties of node reward algorithm

- Rewards proportional to performance, reputation, and partly operator pledge
- Some rewards may not be allocated due to eg, low performance or low reputation (rewards maximally distributed at equilibrium)
  - Equilibrium: exactly K nodes with saturated reputation and perfect performance
  - Unallocated rewards are fed back to the mixmining pool
- Size of network (K)
  - Capped reputation incentivizes spread of reputation over K nodes
- Sybil protection (α)
  - Financial penalty for operators splitting their own stake over multiple nodes

Distribution of node rewards

$$\text{Profit} = \text{Total Node Rewards} - \text{Declared Node Cost}$$

Given to operator

$$\text{Profit Operator} = \text{Profit} * \left(\text{profit margin} + (1 - \text{profit margin}) * \frac{\text{Pledge}}{\text{Total node stake}} \right)$$

$$\text{Profit Delegate} = \text{Profit} * (1 - \text{profit margin}) * \frac{\text{stake delegate}}{\text{Total node stake}}$$

Given to delegates

Properties of node reward sharing algorithm

- Prioritize covering operational costs before distributing profits
- Nodes compete on cost-effectiveness and profit margin
  - Untruthful cost declarations are not advantageous (proof in the paper)
  - Profit margins are discovered through market competition between nodes
- Diminished returns for all node delegates when a node becomes oversaturated

Simulator

- Study reward distributions when the system is not in equilibrium
- Scenarios with various staking distributions, service demand, and network parameters
- Useful for testing impact of network parameters and staking behaviours

- Available: https://github.com/nymtech/rewardsharing-simulator

Examples empirical results

stake distribution over nodes

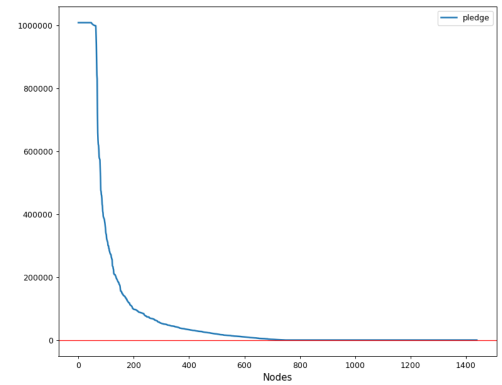

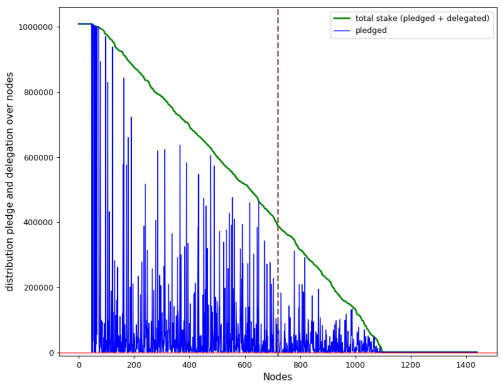

Low demand

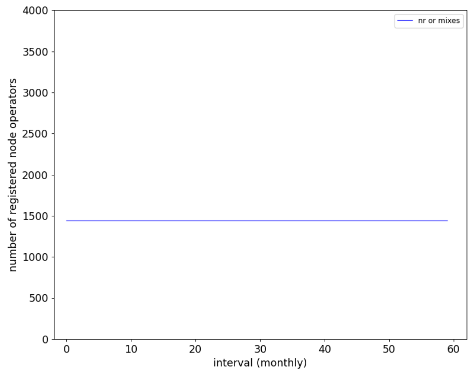

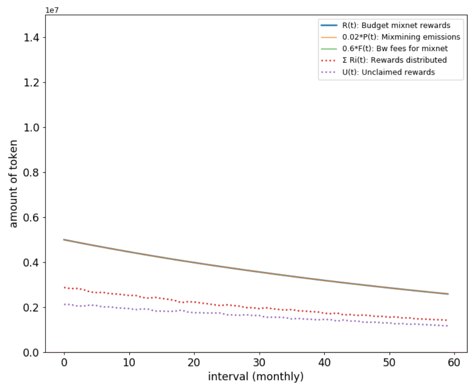

fast-growing demand

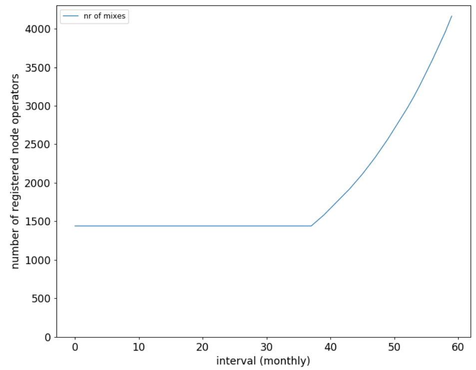

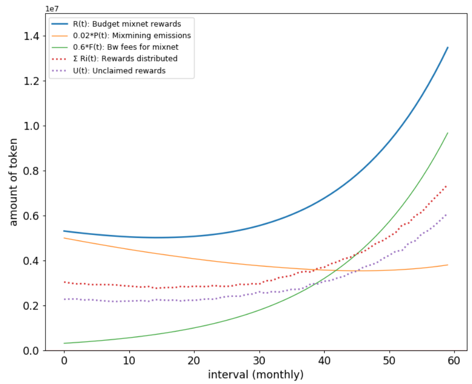

Node rewards vs reputation

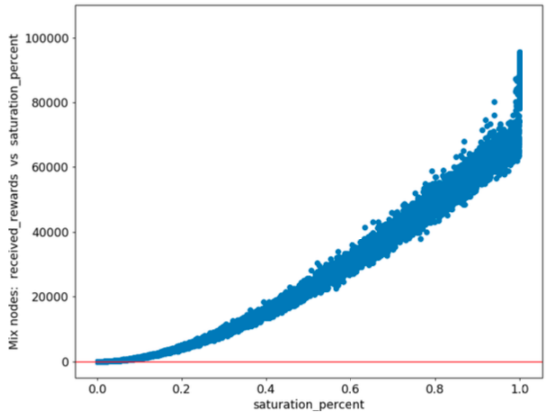

(a) Annualized node rewards in S_0 (first year).

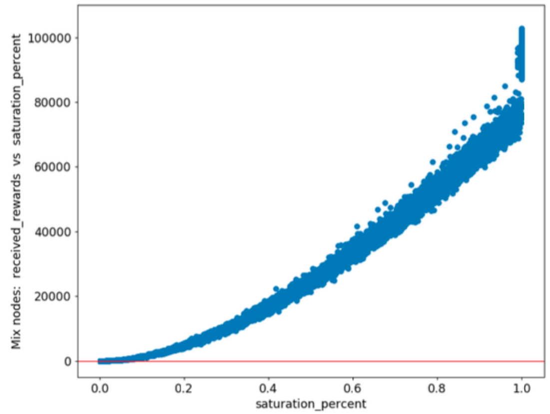

(b) Annualized node rewards in S_1 (first year).

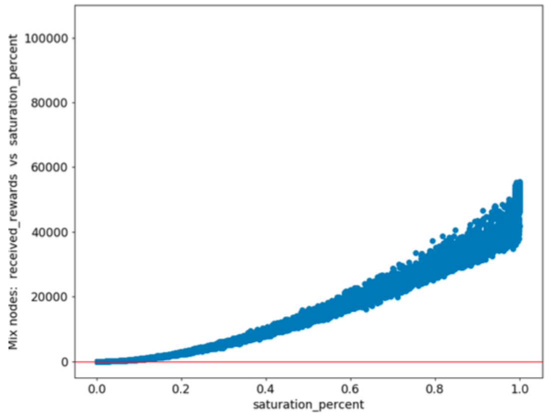

(c) Annualized node rewards in S_0 (fifth year).

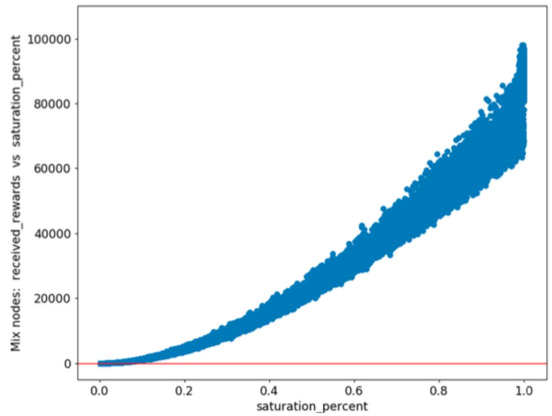

(d) Annualized node rewards in S_1 (fifth year).

Operator rewards vs pledge saturation level

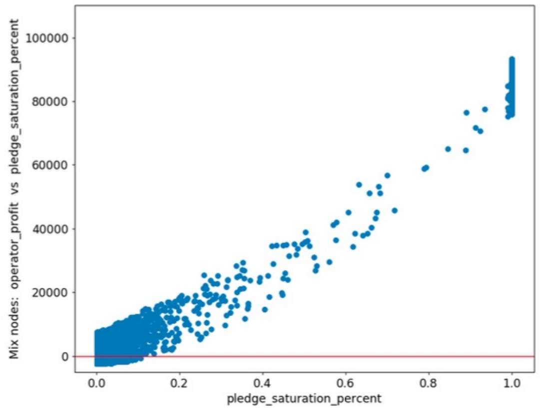

(a) Annualized operator rewards in S_0 (first year).

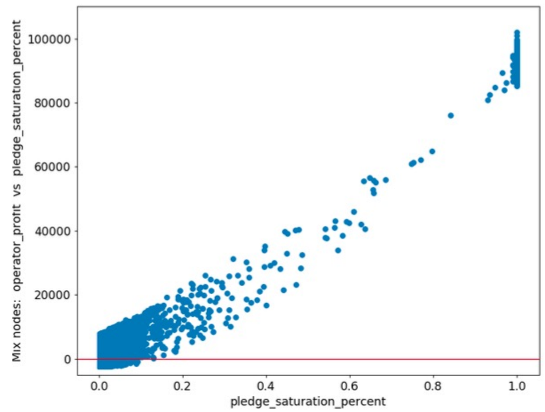

(b) Annualized operator rewards in S_1 (first year).

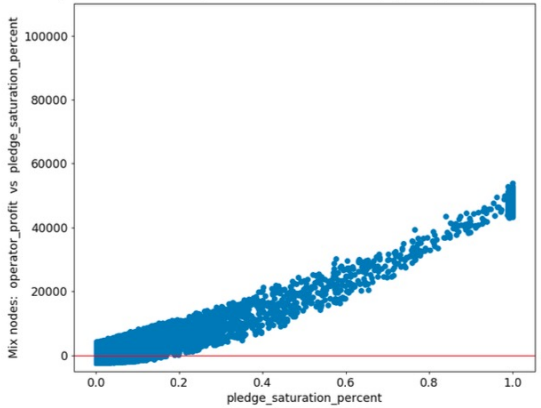

(c) Annualized operator rewards in S_0 (fifth year).

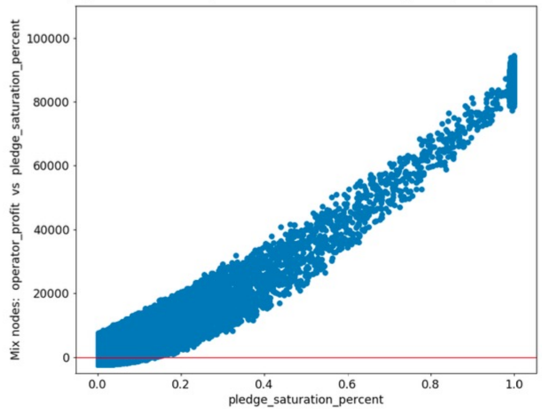

(d) Annualized operator rewards in S_1 (fifth year).

Annualized Return on Stake (RoS) for delegates

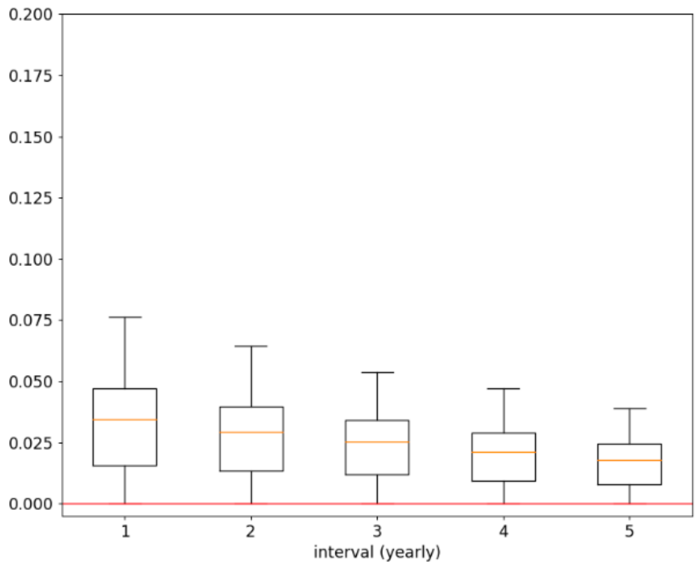

(a) Annualized ROS for delegates in S_0.

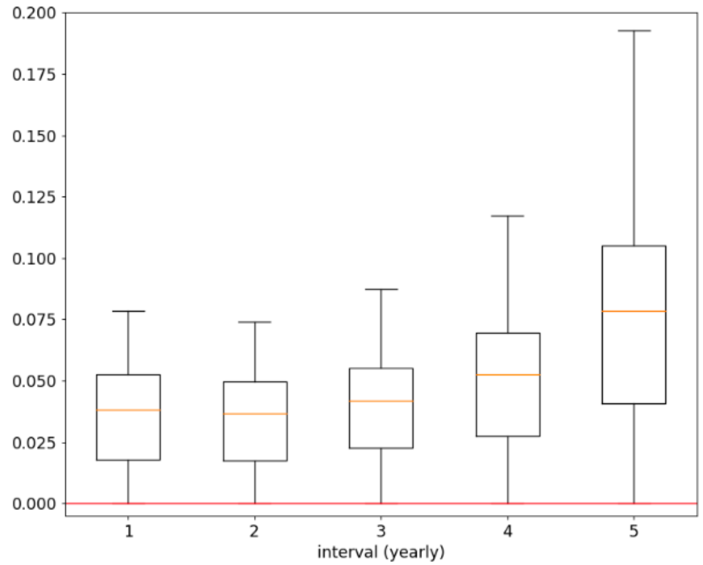

(b) Annualized ROS for delegates in S_1.

RoS vs node reputation

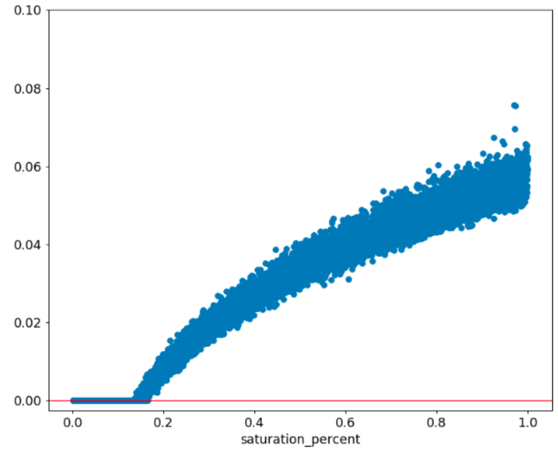

(a) Delegate ROS vs node saturation in S_0 (first year).

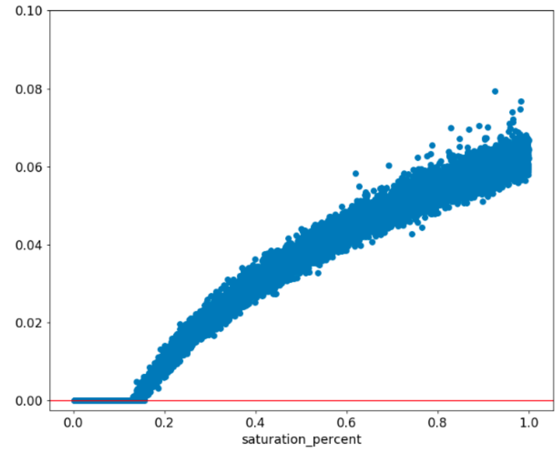

(b) Delegate ROS vs node saturation in S_1 (first year).

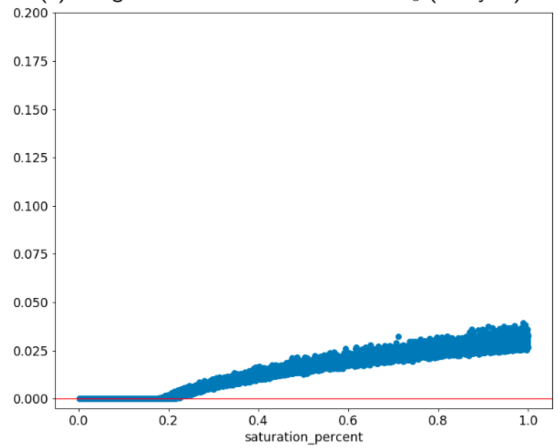

(c) Delegate ROS vs node saturation in S_0 (fifth year).

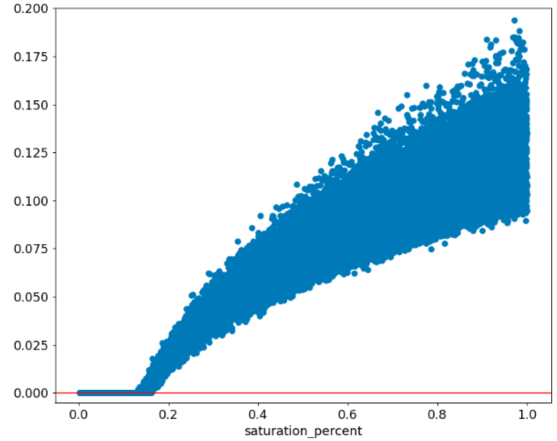

(d) Delegate ROS vs node saturation in S_1 (fifth year).

Parameters example simulations

Name	Value	Notation	Notes
Reference Mix Node			
Minimum node pledge	1000 NYM		Constant
Number CPUs per node	16		Constant
Peak packets/second per CPU	3125 p/s		Grows 1% monthly (12.7% yearly)
Monthly costs per node	$200	$C_i(\cdot)$	Constant
Node performance	1.0 (100%)	ρ_i	Constant
Node profit margin	0.1 (10%)	μ_i	Constant
Mixnet parameters			
Layers of mixnet	3	L	Constant
Width of mixnet	≥ 120	W	Proportional to demand
Active nodes	≥ 360	A	$A = L \cdot W$
Idle (reserve) nodes		B	$B = A$
Rewarded nodes	≥ 720	K	$K = A + B = 6 \cdot W$
Total node candidates	≥ 1440	N	$N = 2 \cdot K$
Average mixnet load	20%		Network absorbs 5x peaks
Simulation parameters			
Epoch	1 hour		
Reward interval	1 month	t	720 hours (epochs)
Simulated period	60 months (5 years)		
Data routed per interval		$M(t)$	Dependent on Scenario S_0, S_1
Scenario S_0 "low demand"	$M_0(0) = 0$	S_0	$M_0(t) = 0$ p/month
Scenario S_1 "growing demand"	$M_1(0) = 500 \cdot 10^9$	S_1	$M_1(t+1) = 1.06 \cdot M_1(t)$ p/month
Exchange rate NYM	1 NYM = $1		Constant
Price for users	$1 for 10^6 packets		Constant
Income from fees in S_0	$F_0(0) = 0$	S_0	$F_0(t) = 0$ NYM/month
Income from fees in S_1	$F_1(0) = 500 \cdot 10^3$	S_1	$F_1(t+1) = 1.06 \cdot F_1(t)$ NYM/month
Token distribution and staking parameters			
Mixmining pool reserve	$P(0) = 250$m NYM	$P(t)$	$P(t+1) = P(t) - 0.02 \cdot P(t) + U(t)$
Monthly pool emissions	2%		$0.02 \cdot P(t)$
Budget rewards entire mixnet		$R(t)$	$R(t) = 0.02 \cdot P(t) + 0.6 \cdot F(t)$
Rewards for node i (out of K)		$R_i(t)$	Eq. (4)
Unclaimed rewards		$U(t)$	$U(t) = R(t) - \sum_i R_i(t)$
Available staking supply	initial: 750m NYM		1 billion minus $P(t)$
Per-node stake saturation point	initial: 1.04m NYM		Available supply divided by K
Pledged stake	0.15		Constant at 15% of available stake
Delegated stake	0.6		Constant at 60% of available stake
Unallocated stake	0.25		Constant at 25% of available stake
Sybil resilience parameter	0.3	α	Constant

Summary

- Economic model for incentivized mixnets
- Market for private bandwidth that can scale to serve demand
- Promotes quality of service and cost effectiveness
- Leverages staking and stake delegation as *node reputation*
- Participation in service provisioning is proportional to reputation
- Rewards are proportional to performance and reputation
  - Need for accurate performance estimations
- Algorithmic rewards and decentralized network management with input from all stakeholders
- Gory details: https://nymtech.net/nym-cryptoecon-paper.pdf