This study investigates the process of validator committee selection in permissionless blockchain networks operating on the Proof-of-Stake algorithm. The task addressed relates to the vulnerability of conventional static selection schemes to identity-forging (Sybil) attacks.

A fixed baseline weight facilitates stake splitting among numerous fictitious entities, allowing attackers to gain control over the network. In response to these challenges, a method for the dynamic stabilization of consensus based on an adaptive control law has been devised. This method automatically regulates the weight mixing intensity using the smoothed Gini coefficient.

The concept of Proof-of-Persistence has been proposed, which replaces the uniform baseline distribution with a time-weighted reputation of the participants. The analytical and experimental analyses of data from 10 real-world networks were conducted, demonstrating that the proposed mechanism reliably reduces the aggregate weight of a potential attacker. The result is attributed to the fact that when new entities are created, their prior participation experience is not considered, and the loss of reputational weight outweighs the benefits of acquiring new baseline shares. This makes the stake-splitting strategy economically unviable.

An important distinct feature is that the system's adaptation is carried out exclusively on the basis of deterministic on-chain data, without the need for external identification. The proposed system functions autonomously: under a normal mode, intervention is minimized, while under the risk of an oligopoly, protection is strengthened. The results could be practically applied to the architecture of permissionless blockchain networks as the method might be integrated both at the network protocol core level and in the form of smart contracts to enhance the security of distributed ledgers without additional manual adjustments

Keywords: Proof-of-Stake, Sybil resistance, adaptive consensus, committee selection, Gini coefficient, Proof-of-Persistence

UDC 004.75:519.8:519.2
DOI: 10.15587/1729-4061.2026.355853

DEVISING A METHOD FOR DYNAMIC STABILIZATION OF POS CONSENSUS USING ADAPTIVE WEIGHT MIXING AND A TIME FACTOR

Ihor Solomka
Corresponding author
PhD Student*
E-mail: ihor.r.solomka@lpnu.ua
ORCID: https://orcid.org/0009-0004-3705-8801
Bohdan Liubinskyi
PhD, Associate Professor*
ORCID: https://orcid.org/0000-0002-0715-8068
*Department of Applied Mathematics
Institute of Applied Mathematics and Fundamental Sciences
Lviv Polytechnic National University
Bandery str., 12, Lviv, Ukraine, 79013

Received 09.01.2026
Received in revised form 13.03.2026
Accepted 20.03.2026
Published 30.04.2026

How to Cite: *Solomka, I., Liubinskyi, B. (2026). Devising a method for dynamic stabilization of PoS consensus using adaptive weight mixing and a time factor. Eastern-European Journal of Enterprise Technologies, 2 (9 (140)), 19–28. https://doi.org/10.15587/1729-4061.2026.355853*

1. Introduction

Blockchain networks based on the Proof-of-Stake (PoS) algorithm have formed a developed infrastructure to support decentralized financial services, digital assets, e-governance, and other distributed ledger applications. As the scale of networks and the volume of assets circulating in them grow, the issues of efficiency, reliability, and security of consensus mechanisms become particularly relevant. This is important not only from a technical but also from an economic and organizational point of view, since stability determines the level of trust of participants and the sustainability of the ecosystem.

Analysis of the functioning of real PoS networks revealed a significant uneven distribution of stakes between validators. In mature networks, a significant share of the stake is concentrated in the hands of a limited number of participants, which creates risks of centralization and oligopolistic influence on the consensus process. In the long term, this reduces the network's resistance to coordinated attacks and limits the possibilities of fair and transparent transaction confirmation.

An additional threat is Sybil attacks, in which an entity artificially distributes assets across multiple pseudo-independent addresses in order to increase its representation among validators or in committees. Open blockchain systems are particularly vulnerable to such manipulations since the creation of new addresses does not require identification. In the absence of effective countermeasures, this could lead to disproportionate influence of individual players, which leads to transaction censorship and violates the integrity of the consensus.

Most modern PoS protocols are based on rigid, fixed rules for validator selection and immutable parameters, which provide predictability but limit the flexibility of the network. Such architecture does not take into account stake concentration, changes in validator activity, or anomalous behavior, as a result of which the protocols limit the potential of decentralized networks or prove ineffective in critical situations.

In this regard, special attention is paid to scientific research aimed at dynamic stabilization of consensus, which takes into account not only the size of the stake but also the duration and stability of the participation of validators. Such approaches make it possible to design adaptive mechanisms for responding to threats in real time, which strengthens the principles of decentralization. In addition, the implementation of dynamic models could increase the trust of institutional participants and open new prospects for the large-scale application of PoS technologies in critical industries.

Therefore, it is a relevant task to carry out studies aimed at devising methods for dynamic stabilization of consensus and mechanisms for adaptive selection of committees in open blockchain networks.

2. Literature review and problem statement

Most current publications on PoS consensus mechanisms summarize that the security and decentralization of open blockchain networks are determined not only by cryptographic procedures but also by the rules for assigning weights to participants and the mechanism for forming validator committees [1]. It is pointed out that attempts to enhance fairness by changing weighting rules have often come into conflict with the requirements of resistance to identity forgery and strategic manipulation of share allocation. One option for overcoming such a conflict may be the transition from static settings to adaptive, context-sensitive consensus management mechanisms.

Study [1] systematizes the classes of blockchain consensus protocols and outlines typical adversary models and trade-offs between throughput, latency, robustness, and decentralization. The value of such a generalization was that the selection of validators and the method of weighing their influence were chosen as the central factor that linked security guarantees with the economic structure of the network. At the same time, the analysis was mainly taxonomic in nature. It fixed the classes of solutions and general constraints but did not derive a protocol control law that would change the selection parameters depending on the measured state of the network and at the same time limit the motivation to fake the plurality of participants. Therefore, the very fact of systematization left open the question of how exactly in a PoS network the requirements of security and decentralization should be combined under conditions of a changing share distribution [1]. The quantitative side of the problem was considered in work [2], in which decentralization metrics for weighted consensus were formulated, an empirical analysis of a number of PoS networks was conducted. and a significant concentration of weight among a limited number of validators was revealed. To reduce this effect, sublinear weighting of the share was proposed, which improved the value of the decentralization metrics. At the same time, a key dilemma was demonstrated: any deviation from the linear proportion between the share and the chance of participation changed the incentive structure and could turn the weight function into a tool for manipulation. Since sublinear transformations enhanced the "return" from smaller shares, they potentially made it profitable to strategically split a large share into several identities to increase its total weight. This is especially important if the protocol does not introduce internal counterincentives or does not require expensive identity confirmation. Thus, in [2], an important drawback of static weight transformations was identified: they could improve the metrics on average, but did not provide resistance to the rational behavior of an attacker who optimized the identity structure for a specific rule [2].

The discussion of the economic foundation is deepened in [3], in which PoS consensus is interpreted as an incentive mechanism capable of ensuring agreement without the energy costs characteristic of PoW. It is shown that the rules of rewards and participation form the equilibria necessary for the stability and viability of the protocol [3]. However, in [3], attention is focused on the conditions for reaching consensus in an abstract form; the issues of concentration of influence in committee mechanisms and the motivation for splitting a stake into many identities were not considered in detail. In [4], the evolution of stakes in a PoS cryptocurrency was studied, and it was shown that under certain conditions stakes can have stabilizing properties, and "automatic" increase is not universally inevitable [4]. However, these conclusions depend on market assumptions, and in [4] a mechanism was not designed to control the distribution of influence in the selection of committees that would respond to changes in concentration in real time. This leaves an open question: even if the distribution of shares does not always degrade, the protocol must remain resistant to highly concentrated regimes and strategic attempts to reinterpret "participation" through multiple identities [3, 4].

A modern approach to quantifying and enhancing decentralization in PoS networks is proposed in [5]. It presents the mathematical metrics devised for accurately measuring influence among validators and suggests ways to algorithmically improve consensus to combat oligopoly. However, while the study provides a powerful tool for determining the level of centralization, the proposed improvement mechanisms leave it vulnerable to identity manipulation. Using their metrics to redistribute influence without taking into account time experience (continuity of node operation) inevitably stimulates Sybil attacks in open environments.

In [6], the results of research on the nominated proof-of-stake protocol are reported. It is shown that the mechanism mitigates concentration by complex optimization of the election and delegation process. However, issues related to the dynamic adaptation of the system to changing conditions remain unresolved. The reason is the lack of mathematical apparatus of the theory of automatic control in the basic protocol. An option to overcome the difficulties may be algorithmic economic self-stabilization of the network. This is the approach considered in [7], in which a simulator was designed for testing PoS mechanisms. It was shown that the combination of protocol parameters can lead either to a close "stable" distribution of wealth or to a significant degradation of fairness depending on the settings. An important contribution of [7] is that instability is associated with sensitivity to parameters, and it was emphasized that constant parameter values are not universal for different networks. However, the approach presented in [7] remained mainly offline: stability was determined by selecting and testing parameters in a simulation, rather than through intra-protocol control that would automatically adapt to concentration in the network. In addition, paper [7] did not consider the attack of forging the multiplicity of participants through shard splitting, although this strategy can significantly change the results of any "leveling" rules [2]. Thus, although [7] showed ways to achieve stability through tuning, it remained unknown how to achieve such stability autonomously and without incentives to manipulate identities.

The issue of committeeization as an element of scaling was investigated in [8], in which an approach to the formation of various committees for sharded architectures was proposed. In the work, attention was focused on the structure of committees and their selection criteria for local network segments, which reduced the risk of concentration in individual subnetworks and facilitated scaling in accordance with certain security requirements. However, the emphasis was on engineering and architectural aspects: the optimization criterion was the suitability for sharding and the functioning

of subcommittees, while the single global structure of the committee and its resistance to participant forgery remained unresolved protocol issues. Therefore, it remains relevant to study how to take into account not only structural diversity in PoS-committeeization but also the ability to adapt depending on changes in the concentration of the stake in the network, which directly affects the risk of committee capture.

In [9], attempts are described to replace or supplement the weight with reputation indicators. A mechanism for selecting nodes based on reputation is proposed using a machine learning model to predict its values. The authors argue that this approach enhances fairness and reduces the probability of attracting unwanted nodes to the network and also demonstrate its applicability in reputation-based consensus. However, the work also encounters typical limitations of such solutions: the complexity of evaluation and training, problems of reproducibility and determinism of the calculated indicators, as well as high vulnerability to attacks if the reputation criteria are observable. In the context of an open PoS network, this means that without an intra-protocol, transparent and simple weight update law, reputation can turn into an off-chain trust system, which undermines autonomy and complicates the formal justification of incentives.

In [10], formal steps are provided to combine fairness and identity fraud resistance, using a selection mechanism based on node clustering and comparing aggregated metrics to improve fairness. The study recognizes that long-term PoS implementation could lead to concentration and proposes a design to mitigate this effect. However, the approach presented in [10] has two significant drawbacks: first, it relies on complex calculations and difficult-to-interpret intermediate quantities, which makes it difficult to estimate costs and implement in high-traffic networks; second, the mechanism remains largely static and does not take into account changes in validator concentration over time, which could lead to incorrect behavior under low and high concentration conditions, with the risk of splitting or capture.

The concept of self-stabilization in the context of blockchain was developed in [11], in which a self-stabilizing version of Byzantine consensus was proposed, capable of restoring correctness after temporary failures. It was shown that the agreement invariants can be restored even after arbitrary transient state violations, which is important for the long-term operation of distributed ledgers. However, in [11], self-stabilization was considered as a property of the protocol state, and not as a characteristic of the economic incentives of participants, and no mechanism was proposed that purposefully regulates the selection weights depending on the concentration and behavior of the adversary. Therefore, work [11] does not eliminate the need for a special mechanism that turns "self-stabilization" into a combination of correctness restoration and incentive management in PoS committee selection [11, 12].

The threat perspective was systematized in [12], which analyzed the risk of identity spoofing attacks, known as Sybil attacks, and double-spending in blockchain systems. It was indicated that attack scenarios should be described in terms of assets, vulnerabilities, adversary capabilities, and expected benefits. The advantage of [12] was that the problem of identity spoofing is related to the economic motivation of the attacker, which emphasized the importance of considering incentives in the threat model. At the same time, [12] did not address the issue of a specific protocol mechanism for PoS-selection of the committee, which would reduce concentration and complicate short-term splitting of the share without external verification of the person.

That problem was studied in [13], in which a committee selection rule with a mixing parameter between the partial component and the limited basic component was applied. It was proven that such mixing reduces the risk of committee capture under certain conditions of adversary concentration. It is found that even moderate mixing can reduce the expected share of adversary seats in the committee and reduce the probability of catastrophic deviations. At the same time, the study notes that static mixing with a fixed parameter is not universal for networks with different levels of concentration, and the concept of a basic component requires the assumption of a "high identity cost" or attribution at the entity level to avoid incentives for strategic splitting. Thus, the study justifies the usefulness of mixed weights for reducing the risk of capture, but at the same time the task of making mixing context-sensitive and resistant to spoofing of multiplicity of participants without external identification remains unsolved.

Thus, our review of the literature has demonstrated that key consensus approaches fall into several incomplete categories. Survey and economic studies have explained the nature of PoS trade-offs and the role of incentives but have not provided protocol rules for managing weights depending on concentration [1, 3, 4]. Work on weight alignment and stability has shown that changing rules or selecting parameters can improve metrics and long-term behavior but remains vulnerable to strategic share splitting or does not offer real-time intra-protocol adaptation [2, 7]. Solutions based on reputation scores or complex grouping procedures have improved fairness in some cases, but have relied on heavy computation, complex verification, and possible manipulation, making them difficult to implement in open networks [9, 10]. Self-stabilizing approaches helped restore correctness after failures but did not eliminate the need for a mechanism to regulate incentives for identity fraud [11, 12]. Mixed committee weighting is theoretically sound, but in a static form it required assumptions incompatible with open participation and did not take into account that the optimal mixing intensity should depend on the level of concentration [13].

The problem identified in our literature review prompted research to devise a method for an adaptive committee selection mechanism in PoS networks that would use intra-protocol concentration indicators to control the intensity of mixing weights and at the same time take into account the time factor of participation as a means of counteracting instantaneous splitting of the share, without relying on external verification of the person and positional control.

3. The aim and objectives of the study

The aim of our research is to devise a method for dynamic stabilization of the Proof-of-Stake consensus based on adaptive weight mixing and time factor. This will make it possible to automatically adapt the level of protection of open blockchain networks to current threats to capital concentration. At the same time, the benefit from stake splitting attacks should be significantly reduced solely due to deterministic on-chain data, without the need to implement external verification of the person.

To achieve the goal, the following tasks were set and subsequently solved:

– to build a mathematical model of the law of adaptive control over the weight mixing parameter based on a smoothed global inequality signal (Gini coefficient);

– to formulate the concept of a time-weighted basic level of reputation (Proof-of-Persistence) and analytically prove its

anti-Sybil properties for one-time deviations of the replacement type;

– to perform comparative simulation of the effectiveness of the proposed adaptive method and classical static approaches based on a set of historical data from ten real blockchain networks (in particular, Solana, Avalanche, Cosmos Hub, etc.).

4. The study materials and methods

The object of our study is the process of forming a validator committee in open (permissionless) blockchain networks based on the Proof-of-Stake algorithm.

The principal hypothesis assumes that the implementation of the law of dynamic control over the weight mixing parameter could make the Sybil-splitting strategy mathematically strictly dominated and economically unprofitable. Such a law would be based on a smoothed Gini coefficient with the replacement of a uniform basic weight with a time-weighted reputation (Proof-of-Persistence). In this case, the consensus should automatically stabilize without involving external identification mechanisms.

The following definitions were used in the work. Proof-of-Stake (PoS) is an algorithm in which the probability of selecting a validator increases in proportion to its stake. Proof-of-Persistence (PoP – "proof of participation") is the time component of the selection weight, which depends on the duration of the validator's continuous activity. The Sybil attack in this paper was considered in the form of "Sybil splitting" – a single economic entity divides a stake among many identities to gain additional advantage from the identity base component.

Since its formalization, the Sybil attack has remained a key threat to systems where the reward depends on the number of identities. In standard linear PoS, stake splitting is selection-neutral: for a fixed total sum S distributed among k identities, the cumulative probability of selection does not change. In contrast, mixed-weight schemes can unintentionally subsidize splitting by providing an additional base component to each identity.

In [13], a λ-mixing rule was proposed and analyzed

$$q_i(\lambda) = (1-\lambda) w_{stake,i} + \lambda u_{base,i}, \quad (1)$$

where $w_{stake,i}$ is the normalized stake weight and $u_{base,i}$ is the base component. If $u_{base,i}$ depends only on the number of validators (e.g., $u_{base,i} = 1 / N$), then an adversary who divides a stake by k identities receives k copies of the base stake instead of one.

Assumptions adopted in the study:

1. The time index t is defined as an "epoch (or day)" is a discrete observation step at which the set of active validator identifiers is formed.

2. Validators were considered as rational economic agents whose main goal is to maximize the cumulative probability of their selection to the consensus committee (reward maximization). In the epoch index t, let $V_t = 1,...,N_t$ denote the active validator identifiers. Let $s_i(t) \geq 0$ be the share delegated to the identifier $i \in V_t$ then the normalized stake weights will be

$$w_i(t) := \frac{s_i(t)}{\sum_{j \in V_t} s_j(t)}, \quad \sum_{i \in V_t} w_i(t) = 1.$$

3. When simulating a substitution attack, the total amount of blocked capital $s_i(t)$ of the attacker remained unchanged, and the stake was only redistributed among k new pseudo-independent identities.

4. The blockchain network has the technical ability to deterministically track the continuous activity experience (age) of each validator $\tau_i(t)$ at the on-chain data level.

5. The committee selection mechanism assigns each identification a sample weight $q_i(t) \in [0,1]$ with $\sum_i q_i(t) = 1$. Then $q_i(t) \in [0,1]$ determines the expected share of committee seats obtained by i from an independent sample. The total sample weight is used as an evaluation metric.

6. The value of the Gini coefficient G is used as a control signal and can be calculated by the system automatically at each epoch without involving external oracles.

Simplifications accepted in the study:

1. For analytical proof of anti-Sybil properties, only the one-shot deviation model was considered, in which a mature validator was instantly replaced by k fresh identities. Strategies of long-term "farming" or gradual accumulation of experience by many sleeping nodes with a minimal stake were deliberately not modeled in this work.

2. As the main metric of the efficiency of the entity E, the cumulative selection weight $Q_E(t)$ was used, which is interpreted as the expected share of seats in committees. This avoids the complex modeling of explicit cryptographic draw procedures in the empirical part.

3. Considering the honest state (single identity) and the attacked state (same stakes replaced by k identifiers) in the same epoch, the relative change is defined as

$$\Delta_{atk} := 100 \cdot \frac{Q_E^{atk}(t) - Q_E^{hon}(t)}{Q_E^{hon}(t)}. \quad (2)$$

4. Negative Δ_{atk} indicates that the partitioning reduces the cumulative weight of the attacker's selection in this epoch.

5. To ensure the reproducibility of the numerical experiment on historical data of 10 networks, the parameters of the concave age transformation (in particular, the coefficient $\beta = 0.7$ and the maximum age $\tau_{\max} = 730$ of epochs are fixed as invariant constants.

Historical snapshots of ten real PoS networks were used for validation. At step $t_{atk} = 50$, the identity with the largest stake is replaced by $k = 50$ new identities while preserving the total stake.

The experiment compared a static model with $\lambda = 0.3$ and uniform $u_i = 1 / N_t$ and an adaptive model with dynamic λ_t (5) and temporal u_i (7).

The main indicator is the relative change in the cumulative weight of entity selection Δ_{atk} (2).

The data set was built by daily polling of the main networks for 60 days (September-October 2025). For this purpose, a special polling mechanism was designed, which was launched automatically via the workflow configuration in GitHub Actions. The script polled specific application interfaces of ten networks, including Solana, Avalanche, Cosmos Hub, Sui, and others, every 24 hours.

During this data collection process, each historical snapshot consisted of the current list of validators and their stake amounts, as well as a stable node identifier (usually the base address of the operator). A separate functional call was the rotation of consensus keys, which many networks allow for security reasons. The node history during a planned key change was preserved, since the main and immutable operator account was used. This allowed us to avoid false zeros and correctly reproduce the continuous age vector $\tau(t)$ for each sample participant.

The full dataset and simulation code are available in our accompanying repository [14].

5. Results of devising a method for dynamic consensus stabilization based on adaptive mixing of weights and time factor

5. 1. Result of constructing a mathematical model of adaptive control

The transition to a dynamic model is due to a compromise between distribution efficiency and consensus security. Fixed parameterization of λ in heterogeneous networks leads to suboptimal results. In networks with a low Gini level (diffuse stake), a high λ punishes large honest participants without a significant increase in security. On the other hand, in networks with a high Gini level, a moderate λ may be insufficient to reduce the influence of large players below critical thresholds.

To quantify the inequality of stakes, the Gini coefficient is used. For a non-negative weight vector $w(t) = (w_1(t), \ldots, w_{(Nt}(t))$ is defined as

$$G_t = \frac{\sum_{i=1}^{N_t}\sum_{j=1}^{N_t}\left|w_i(t) - w_j(t)\right|}{2N_t\sum_{i=1}^{N_t}w_i(t)}. \quad (3)$$

Since $\sum_i w_i(t) = 1$, the denominator is $2N_t$. A study of ten real networks revealed significant heterogeneity: new networks exhibit $G \approx 0.35$–0.40 while mature DPoS networks $G > 0.75$ (Fig. 1). The empirical Gini coefficients in the ten networks show significant differences. A static $\lambda = 0.3$ leads to regulatory imbalance. The adaptive model scales the intervention proportionally to the measured concentration.

Fig. 1 shows a comparison of the classical approach with a fixed mixing parameter ($\lambda = 0.3$, indicated by the red dotted line) and the adaptive mechanism. For ease of analysis, all networks in the plot are ranked by the level of centralization, from the lowest to the highest Gini coefficient value.

The left part of the visualization covers blockchains with a more or less even distribution of capital (such as Sui with $G \approx 0.36$ or Sei with $G \approx 0.37$). The use of rigid static rules in such ecosystems is counterproductive, as it creates excessive regulatory pressure that simply distorts the economic motivation of honest nodes without any security gain. The proposed dynamic model works differently. It recognizes the absence of a threat of monopolization, so the algorithmic intervention parameter λ_t naturally decreases to almost zero.

The opposite situation is observed on the right, where mature networks with stable oligopoly are collected (for example, Avalanche with $G \approx 0.82$ or Tezos with $G \approx 0.86$). Here, the basic constant $\lambda = 0.3$ is clearly not enough to restrain the dominance of the largest validators. In response to increasing concentration, the adaptive algorithm automatically increases λ_t. Thanks to this, the system receives exactly the level of protection that is needed to prevent centralized capture of the committee at the current moment.

To improve the efficiency of allocation in decentralized regimes and Sybil incentives under identity counting baselines, the static mixture (1) is extended to a state-dependent feedback control system. This structure combines λ_t with the observed concentration signal and defines the basic distribution as a function of stability.

The selection was modeled as a dynamic system, where the intensity of the adjustment λ_t adapts to the stake concentration. Exponential smoothing of the $\tilde{G}_t$ signal was used to filter the noise

$$\hat{G}_t = \alpha \cdot G_t + (1-\alpha)\cdot \hat{G}_{t-1}, \quad \alpha \in (0,1). \quad (4)$$

Let a proportional controller with a threshold G_{target} and saturation $\lambda_{\max}$

$$\lambda_t = \min\left(\lambda_{\max}, \max\left(0, K_p \cdot \left(\hat{G}_t - G_{target}\right)\right)\right), \quad (5)$$

where $K_p > 0$ is the gain factor. If $\tilde{G}_t \le G_{target}$, then $\lambda_t = 0$ (minimum intervention regime).

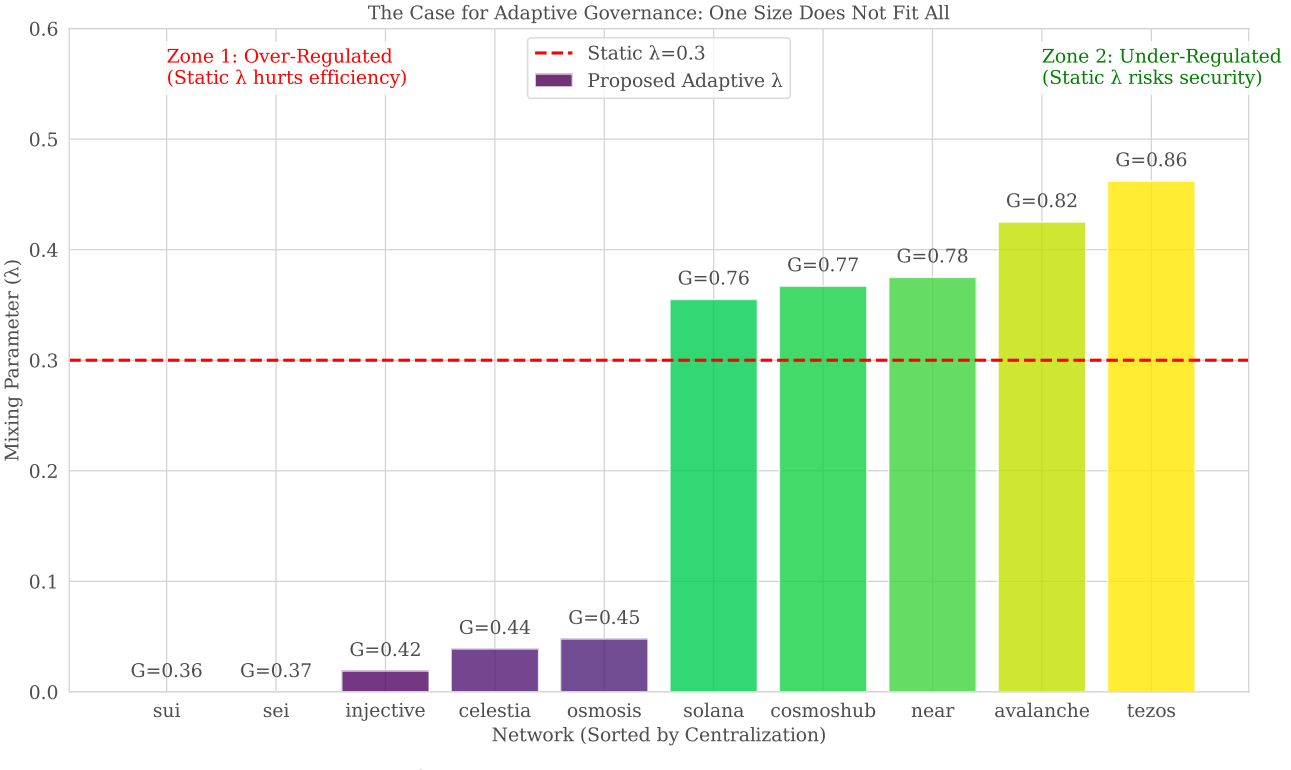

Fig. 1. Contextual adaptation based on concentration signal

To neutralize incentives to split the stake, the baseline is defined as a function of the length of participation. If $\tau_i(t) \geq 0$ is the age of identity i at epoch t, then the bounded concave transformation

$$\tilde{\tau}_i(t) = \min\left(\tau_{\max}, \tau_i(t)\right)^{\beta}, \quad \beta \in (0,1), \quad \tau_{\max} > 0 . \tag{6}$$

The simulations used epochs $\beta = 0.7$ and $\tau_{\max} = 730$ as fixed hyperparameters. All things being equal (the same constraint is present in the current operators), the new identifier reaches a base share coefficient of $(0.3)^{\wedge}\beta \approx 0.43$ after $0.3\tau_{\max}$ periods, and a level of $(0.5)^{\wedge}\beta \approx 0.62$ in about one year. The base distribution is defined as

$$u_i(t) = \frac{\tilde{\tau}_i(t)}{\sum_{j \in V_t} \tilde{\tau}_j(t)}, \tag{7}$$

and selection rule

$$q_i(t) = (1-\lambda_t) w_i(t) + \lambda_t u_i(t). \tag{8}$$

This selection rule is the basis of the mathematical model of adaptive control. The probability of a validator getting into the committee adapts and balances between the weight $w_i(t)$ and the accumulated historical reputation $u_i(t)$. The advantage of such a model is the rejection of rigid centralized protocol parameters. As soon as a dangerous increase in capital concentration is recorded, the λ_t parameter automatically changes the priority towards age participation. As a result, any attacker loses the opportunity to seize control over the network due to a sudden infusion of funds or the creation of hundreds of new clone nodes. To gain real influence, s/he needs to prove her/his conscientious presence for a long time.

The concentration signal is estimated by the Gini coefficient (G_t, a classic statistical indicator of the level of inequality), and exponential smoothing is used to reduce noise (a method in which the latest values are more important than the previous ones). The resulting smoothed indicator is used as the input of a controller with a deadband (a range in which the controller does not respond to changes in the input signal).

The proposed proportional controller with a threshold G_{target} turns on the intervention only when $\tilde{G}_t$ exceeds the target level. It also uses saturation up to $\lambda_{\max}$. This approach implements the principle of minimal intervention under low-concentration regimes. It strengthens the regulation as centralization increases. The contextual motivation for this mechanism is illustrated in Fig. 1. For comparison: a static $\lambda = 0.3$ leads to overregulation or underregulation depending on the network. In contrast, the adaptive model scales the intervention proportionally to the measured concentration.

5. 2. The result of the formation of a time-weighted base level and the proof of the anti-Sybil property

Lemma 1 (*Anti-Sybil property when the replacement type is rejected*). Consider the epoch t with a set of validators V_t and parameter $\lambda_t \in (0,1]$. Let $w_i(t)$ be the normalized stake weights, and $\tilde{\tau}_i(t) \geq 0$ be the base mass. The total base mass: $T(t) = \sum_{j \in V} \tilde{\tau}_j(t) > 0$. Let the entity E control the identity i with the stake $s \in (0,1)$ and the base mass $\tilde{\tau}_i > 0$. If E replaces i with $k \geq 2$ new identities $i_1,\ldots,i_k$ with the total stake $\sum_{l=1}^{k} w_{il}(t) = s$ preserved and the initial base mass $\tilde{\tau}_0(t) \geq 0$ for each, then under the condition

$$\tilde{\tau} > k\tilde{\tau}_0, \tag{9}$$

the total weight of the selection of entity E in epoch t strictly decreases

$$Q_E^{before}(t) > Q_E^{after}(t).$$

Proof. Due to the conservation of the stake, the proportional component of the weight is invariant

$$\sum_{\ell=1}^{k} (1-\lambda_t) w_{i_\ell}(t) = (1-\lambda_t) s.$$

The change in weight is due to the base component. Let $T_{-E} = T(t) - \tilde{\tau}$ be the base weight of the other participants. The share of the entity before replacement

$$B_{before} = \lambda_t \frac{\tilde{\tau}}{T_{-E} + \tilde{\tau}}.$$

After the replacement, the entity controls k identities with a total base mass $k\tilde{\tau}_0$. The new total mass is $T'(t) = T_{-E} + k\tilde{\tau}_0$. The fraction after the replacement

$$B_{after} = \lambda_t \frac{k\tilde{\tau}_0}{T_{-E} + k\tilde{\tau}_0}.$$

The inequality $B_{before} > B_{after}$ is equivalent to

$$\frac{\tilde{\tau}}{T_{-E} + \tilde{\tau}} > \frac{k\tilde{\tau}_0}{T_{-E} + k\tilde{\tau}_0},$$

which after simplification gives $\tilde{\tau} > k\tilde{\tau}_0$, i.e. (9).

In our model, newly created identities have $\tau_0 = 0$ (hence, $\tilde{\tau}_0 = 0$). Then condition (9) reduces to $\tilde{\tau}_i > 0$. This means that any mature validator loses weight when replaced by new identities, which creates an instant barrier to attack.

Therefore, with a one-time replacement "one mature identity → k new ones", the splitting strategy becomes strictly dominated when switching to a time basis. This means that the attacker has lost the accumulated reputation. It is not transferred to newly created Sybil identities.

5. 3. Results of comparative simulation on real network data

The algorithm in Fig. 2 sequentially processes each snapshot S_t for a specific network. Then, at each step, a set of active validators V_t is formed and their individual experience of continuous operation $\tau_V(t)$ is calculated. After that, the algorithm estimates the current level of stake capital concentration G_t. Next, this basic indicator is passed through an exponential moving average filter, forming a stable smoothed signal $\tilde{G}_t$ so that the regulator does not react to random spikes. This smoothing increases the stability of control. Due to this, the control parameter λ_t is saturated within $[0, \lambda_{\max}]$ and is activated only when the target G_{target} is exceeded. Thus, minimal intervention in a decentralized mode and increased regulation with increasing concentration are ensured.

At the time t_{atk}, a counterfactual scenario of a substitution attack is simulated. In this case, the largest validator by stake is replaced by k new identities with equal stake division and zero age $\tau = 0$. Thanks to the simulation of such a one-time deviation, all other background processes in the network are "frozen" and the net impact of identity division is identified. After that, the selection weights $qi(t)$ are

calculated according to rule (8) for the honest and attacked states. Comparing the total weight of the attacker before and after splitting Qhon and Qatk, the relative change Δ_{atk} in percent is obtained. This allows us to compare networks with different scales and different numbers of validators. For the reproducibility of the experiment, all key simulation parameters are given in Table 1.

Table 1

Adaptive control parameters

Parameters	Value	Substantiation
G_{target}	0.40	The threshold below which $\lambda_t = 0$ (minimal intervention)
K_p	2.0	Regulator gain
$\lambda_{\max}$	0.50	Saturation limit (min. 50% of the stake weight)
α	0.10	EMA smoothing factor

Comparative simulation for ten real blockchain networks was performed based on daily snapshots over 60 days (September-October 2025). Snapshots were formed through RPC polling, with a fixed set of validators, their stakes, and stable identifiers.

Algorithm 1 Modeling Adaptive Control and Injection of Sybil Attack

Require: Historical snapshots S, number of Sybil identities k, attack epoch t_{atk}, smoothing coefficient α

Initialize $\hat{G}_0 \leftarrow \text{Gini}(S_0)$
for $t = 1$ to $|S|$ **do**
  $V_t \leftarrow \text{ReadValidators}(S_t)$
  Step 1 – Update participation age.
  for each $v \in V_t$ **do**
    $\tau_v(t) \leftarrow \text{ComputeAge}(v)$
  end for
  $G_t \leftarrow \text{Gini}(V_t.\text{stake})$
  $\hat{G}_t \leftarrow \alpha \cdot G_t + (1-\alpha) \cdot \hat{G}_{t-1}$
  Step 2 —Compute control signal.
  $\lambda_t \leftarrow \min\left(\lambda_{\max}, \max\left(0, K_p \cdot (\hat{G}_t - G_{\text{target}})\right.\right.$
  Step 3 – Model counterfactual attack if $t = t_{\text{atk}}$.
  if $t = t_{\text{atk}}$ **then**
    $V_t^{\text{hon}} \leftarrow V_t$ {Save baseline (honest) state for comparison}
    $v_{\max} \leftarrow \arg\max(V_t.\text{стейк})$
    Remove $v_{\max}$ from V_t
    Add k identities with stake $w_{\max}/k$ and age $\tau = 0$
    $V_t^{\text{atk}} \leftarrow V_t$
  end if
  Step 4 – Compute selection weights.
  Compute $q_i(t)$ according to the adaptive rule (??)
  if $t = t_{\text{atk}}$ **then**
    Compute $Q^{\text{hon}} \leftarrow Q_E(V_t^{\text{hon}})$ та $Q^{\text{atk}} \leftarrow Q_E(V_t^{\text{atk}})$
    Record $\Delta_{\text{atk}} \leftarrow 100 \cdot \frac{Q^{\text{atk}} - Q^{\text{hon}}}{Q_{\text{hon}}}$
  end if
end for

Fig. 2. Adaptive control algorithm

Table 2 gives results of the attack simulation at the time $t_{atk} = 50$. The cumulative weight of the attacker in the base (honest) state Q_{honest}, in the static model Q_{static} and the adaptive model $Q_{adaptive}$ are compared.

The Q_{static} results show a significant incentive to attack when using a uniform base distribution. In the Solana network, the attacker's weight increases from 1.77% to 3.92% (+121%), confirming the multiplicative effect of splitting.

The $Q_{adaptiv}$ results show a negative Δ_{atk} value for all networks. The penalties range from −0.57% (Sui) to −8.06% (Avalanche). This confirms that the loss of experience (PoP) when splitting outweighs any benefits from the base component.

In networks with low concentration (Sui, Sei), the smoothed Gini signal $\tilde{G}_t$ is close to G_{target}, so $\lambda_t \approx 0$. This confirms the ability of the system to automatically switch to the minimal intervention mode.

Table 2

Quantitative impact of the substitution "dominant validator → Sybils" ($t = 50$)

Network	Gini before attack	Q_{honest} baseline	Q_{static} static	$Q_{adaptive}$ adaptive	Sybil effect Δ%
Avalanche	0.817	0.775%	2.663%	0.713%	−8.06%
Near	0.788	2.827%	7.654%	2.663%	−5.81%
Osmosis	0.468	5.586%	14.934%	5.350%	−4.22%
Injective	0.415	7.407%	20.913%	7.131%	−3.72%
Celestia	0.419	4.607%	13.745%	4.451%	−3.38%
Cosmos Hub	0.775	8.052%	16.941%	7.798%	−3.15%
Solana	0.754	1.773%	3.924%	1.718%	−3.10%
Tezos	0.855	8.015%	15.955%	7.813%	−2.53%
Sei	0.370	5.815%	21.010%	5.716%	−1.71%
Sui	0.366	2.810%	10.767%	2.794%	−0.57%

Fig. 3 shows the weight dynamics: at the moment of attack, the static model leads to a jump in the attacker's influence, while the adaptive one leads to its decline.

In the Avalanche network $G \approx 0.82$ static mixing strengthens the attacker, adaptive gives a penalty of −8.06%. For Solana $G \approx 0.75$ adaptive weighting leads to net weight losses for the attacker. Celestia $G \approx 0.42$ the transition zone and the effect of the penalty are preserved. Sui network $G \approx 0.36$ operate under the minimal interference mode

The use of an exponential moving average also protects against manipulation of the Gini index. Even if the attacker artificially lowers G_t by splitting, $\tilde{G}_t$ reacts inertially, keeping λ_t high and forcing the attacker to interact with the system under zero experience conditions. Sublinear transformation accelerates initial growth, and the cap prevents monopolization.

An illustrative convergence of the baseline with a constrained sublinear function of the participation age is shown in Fig. 4. The sublinear (concave) transformation accelerates early growth (relative to the raw linear age), while the constraint prevents unlimited accumulation.

Our results demonstrated that the maximum effect of the adaptive mechanism was manifested in the early post-attack window when new Sybil identities have not yet accumulated participation age. For static schemes with a uniform baseline, an artificial increase in the attacker's influence is observed due to the multiplication of identities. At the same time, under low-concentration regimes, the system retained minimal interference (small λ_t values), which corresponds to the principle of proportional regulation.

Thus, numerical experiments confirm the reduction of the benefit from stake-splitting and are consistent with the analytical motivation. This completes the solution to problem 3.

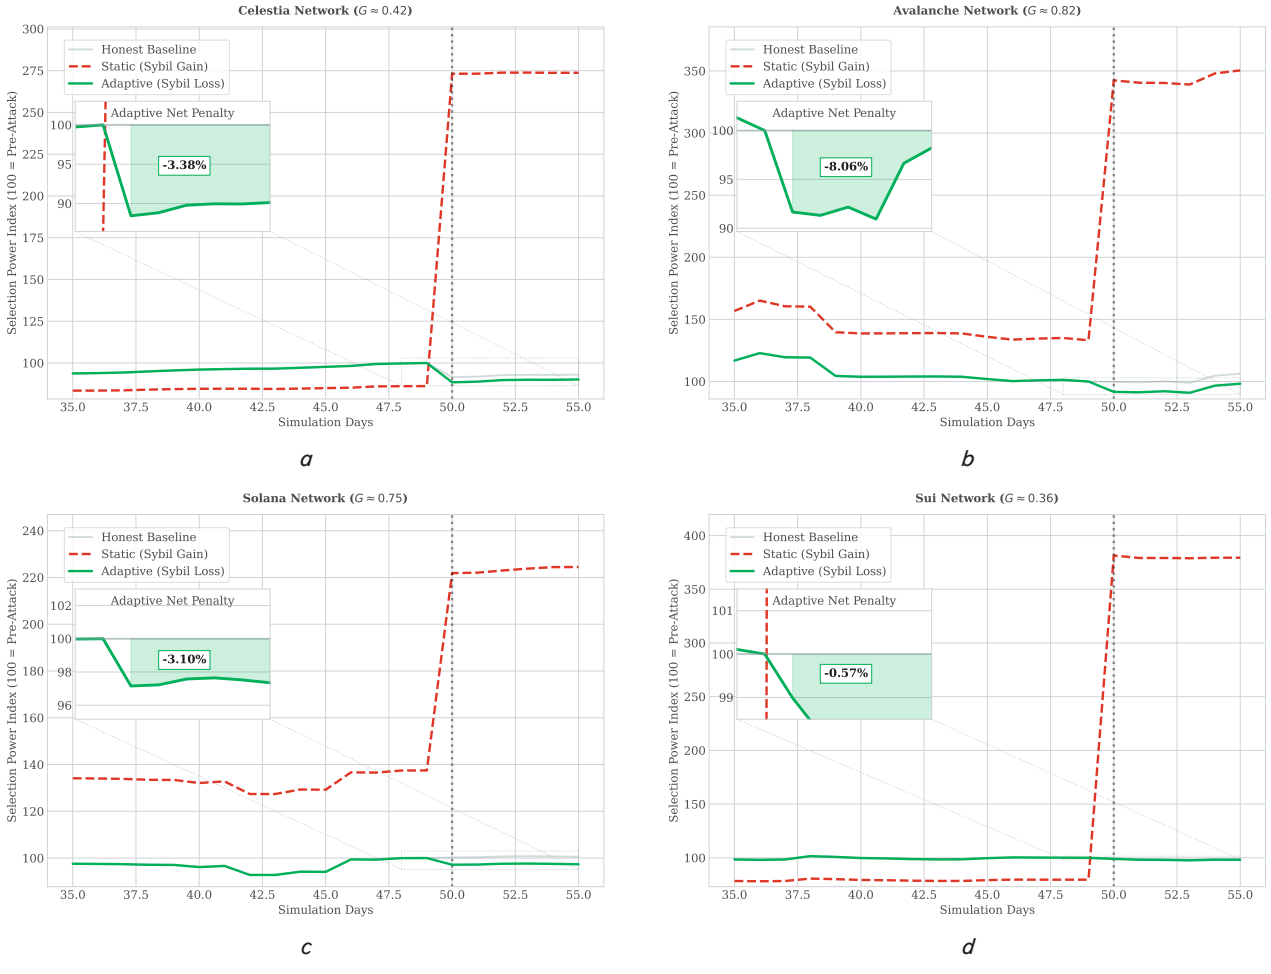

Fig. 3. Dynamics of the impact of Sybil substitution in different network topologies ($t = 50$): *a* — transition zone; *b* — static mixing; *c* — net losses; *d* — minimal interference mode

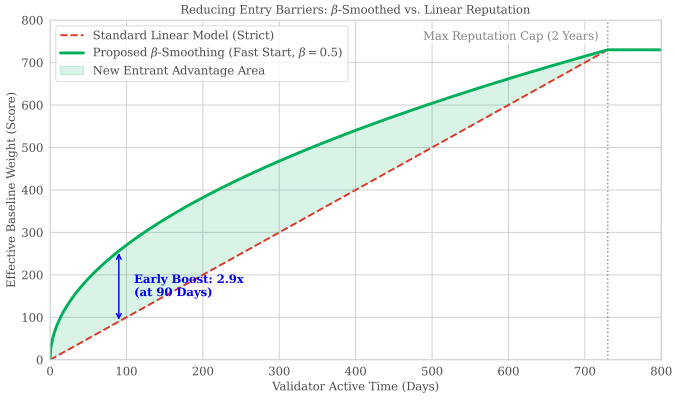

Fig. 4. Weight accumulation dynamics for new participants

6. Discussion of results based on investigating the dynamic consensus stabilization method

The proposed mathematical model of adaptive control (8) has the anti-Sybil property proven by Lemma 1. Our results, shown in Table 2, demonstrated a systematic decrease in the cumulative weight of the attacker $\Delta_{atk} < 0$ in all ten studied networks. This is explained by the architecture of the proposed selection rule described by formula (8). At the time of the attack simulation ($t = 50$), the new Sybil identities had zero experience $\tau = 0$. According to the concave transformation function (6), their contribution to the basic distribution was zero. As shown in the plots of the influence dynamics (Fig. 3), the adaptive controller (5) responded to the threat by shifting the priority towards reputation (Proof-of-Persistence). Therefore, the loss of the accumulated reputation weight strictly outweighed any benefits from splitting the stake.

A feature of the proposed method is a completely autonomous on-chain adaptation without external identification. Unlike static weight mixing methods [2, 11], which in decentralized regimes unintentionally created incentives for stake splitting, our method minimizes interference. This is clearly confirmed by Fig. 1, where for low-concentration networks (Sui, Sei) the parameter λ_t tends to zero. Reputation systems based on machine learning [1] or node clustering mechanisms [12] require complex off-chain calculations and trust in external sources. Our solution is based exclusively on deterministic mathematical indicators, which greatly simplifies its implementation.

The scope of our result's practical implementation is open permissionless PoS networks and protocol modules, where the committee selection procedure is probabilistic and depends on the stake, and identities can be created without verifying the individual. The method can be integrated either at the protocol core level (in the committee

draw rules) or at the smart contract layer level (if that is where the selection/delegation logic is implemented). The conditions for application include the possibility of continuous on-chain tracking of the validator's age and the presence of a stable identifier that is not reset due to technical operations such as key rotation. Without these conditions, Proof-of-Persistence cannot correctly perform the role of a protective factor.

The potentially expected effects of the implementation (which follow from the results in Table 2 and the logic of rule (8)) are as follows. First, the economic attractiveness of instant stake distribution as a Sybil strategy for committee capture is reduced. This creates a short-term anti-Sybil barrier. Second, in networks with diffuse stake distribution, algorithmic interference is minimized. This reduces the risk of "artificial" distortion of incentives in situations where the threat of centralization is low. Third, in highly concentrated networks, the role of the protective component increases. This is aimed at reducing the risk of oligopolistic dominance in committee selection.

The limitations inherent in this study relate to the conditions of application and the ranges of input data. First, the simulation (Table 2) and weight accumulation plots (Fig. 4) were obtained under the condition of fixed hyperparameters of the concave function $\beta = 0.7$, $\tau_{\max} = 730$ epochs. The reproducibility of the claimed effects is guaranteed precisely within these limits; changing the parameters will affect the rate of initial growth of the reputation of honest nodes. Second, the mathematical proof (Lemma 1) considers only the one-shot deviation model with total stake preservation. Third, the method is applicable only in those networks where the architecture allows for continuous monitoring of node age at the protocol level.

The disadvantage of this study is that the proposed model remains vulnerable to long-term planning of attacks (the "farming" strategy). An attacker can create a lot of "dormant" identities with a minimum stake in advance, accumulate experience for them, and then abruptly transfer the main capital there. In addition, in practice there is a risk of a shadow secondary market for the private keys of old validators. Such a strategy indicates that Proof-of-Persistence for full practical stability should be combined with economic responsibility mechanisms at the protocol level. These can be fines or rules that respond to anomalous movement of a large stake between identities. Thus, "reputation" would not be an asset that could be cheaply converted into committee control.

Further development of the method should be directed at a multifactorial control system. It should combine not only the Gini coefficient but also other indicators of concentration and "threshold vulnerability" (HHI, Nakamoto coefficient). The use of multiple metrics is consistent with current practice in measuring PoS decentralization. Different metrics highlight different types of centralization and different risk surfaces [13]. Methodologically, this requires the synthesis of a robust control law with multiple nonlinear state variables to avoid regulatory conflict.

7. Conclusions

1. A mathematical control law is formulated in which the weight mixing parameter λ_t is devoid of staticity. From now on, this value is calculated and adjusted dynamically, directly responding to current changes in the stake concentration. The resulting controller used the exponential smoothing method with a coefficient $\alpha = 0.1$, the maximum Gini threshold $G_{target} = 0.40$, the gain $K_p = 2.0$ and the saturation $\lambda_{\max}$. The peculiarity of the result is the contextual change in the intensity of the intervention, in contrast to the static $\lambda = 0.3$, which gives over-or underregulation in networks with different G. The resulting effect is explained by the fact that the proportional controller is activated only when the target concentration level is actually exceeded, and the signal smoothing creates a hysteresis effect that keeps the system from abrupt deregulation in the case of artificial manipulations of the stake by the attacker.

2. The concept of a time-weighted baseline reputation (Proof-of-Persistence) has been formulated, and Lemma 1 was analytically proven regarding its anti-Sybil properties under one-shot deviations of the substitution type. A qualitative indicator is fixed: the stake splitting strategy becomes strictly dominated under the condition that at $\tau = 0$ it is sufficient that $\tilde{\tau}_0 > 0$ for the mature identity to lose weight after substitution. The difference from known static approaches is that "time" becomes a scarce resource that cannot be instantly transferred to Sybil identities, while a uniform basis $u_i = 1/N_t$ can enhance the multiplication effect. This result is explained by the fact that new Sybil clones have zero experience at the time of creation, respectively, their total basic contribution is zeroed, and the instant loss of the accumulated reputation weight of the "mother" node mathematically outweighs any benefits from the split.

3. A comparative simulation of the effectiveness of the method was carried out on a historical data set of 10 real PoS networks with daily snapshots over 60 days (September–October 2025), which quantitatively confirmed the viability of the adaptive model. It was found that in all the studied topologies the proposed method generates a system penalty for the attacker (a decrease in its total weight from –0.57% in Sui to –8.06% in Avalanche). A feature of the result is its universality for different architectures without manual reconfiguration. Unlike the static model, where a similar attack on the Solana network would lead to a multiplicative increase in the attacker's influence by 121% (from 1.77% to 3.92%), the adaptive method blocks this possibility. Such quantitative indicators are explained by the synergistic action of the controller: at the moment of fixing the attack, the mixing parameter transfers the weight of the selection from financial capital to reputational experience, which the newly created fictitious identities of the attacker simply do not have. It can be implemented through smart contracts or at the core level of the protocol without involving centralized oracles for identity verification (KYC), which preserves the fundamental principles of openness in blockchain networks.

Conflicts of interest

The authors declare that they have no conflicts of interest in relation to the current study, including financial, personal, authorship, or any other, that could affect the study and the results reported in this paper.

Funding

The study was conducted without financial support.

Data availability

The data will be provided upon reasonable request.

Use of artificial intelligence

The work employed the openai-codex/gpt-5.3-codex model for language editing and translation of individual fragments of the manuscript. The use of AI was limited to editorial text processing (without generating scientific results, experimental data, and conclusions). All results, formulas, tables, and conclusions were manually checked by the authors; there is no influence of the AI tool on scientific conclusions.

Acknowledgments

The authors express their gratitude to reviewers for their constructive comments.

Authors' contributions

Ihor Solomka: Conceptualization, Methodology, Software, Validation, Writing – original draft; **Bohdan Liubinskyi:** Writing – review & editing, Supervision, Formal analysis.

References

1. Xu, J., Wang, C., Jia, X. (2023). A Survey of Blockchain Consensus Protocols. ACM Computing Surveys, 55 (13s), 1–35. https://doi.org/10.1145/3579845
2. Motepalli, S., Jacobsen, H.-A. (2024). How Does Stake Distribution Influence Consensus? Analyzing Blockchain Decentralization. 2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), 343–352. https://doi.org/10.1109/icbc59979.2024.10634400
3. Saleh, F. (2020). Blockchain without Waste: Proof-of-Stake. The Review of Financial Studies, 34 (3), 1156–1190. https://doi.org/10.1093/rfs/hhaa075
4. Roşu, I., Saleh, F. (2021). Evolution of Shares in a Proof-of-Stake Cryptocurrency. Management Science, 67 (2), 661–672. https://doi.org/10.1287/mnsc.2020.3791
5. Motepalli, S., Jacobsen, H.-A. (2025). Decentralization in PoS Blockchain Consensus: Quantification and Advancement. IEEE Transactions on Network and Service Management, 22 (4), 2930–2943. https://doi.org/10.1109/tnsm.2025.3561098
6. Cevallos, A., Stewart, A. (2021). A verifiably secure and proportional committee election rule. Proceedings of the 3rd ACM Conference on Advances in Financial Technologies, 29–42. https://doi.org/10.1145/3479722.3480988
7. Leporati, A., Rovida, L. (2024). Looking for stability in proof-of-stake based consensus mechanisms. Blockchain: Research and Applications, 5 (4), 100222. https://doi.org/10.1016/j.bcra.2024.100222
8. Huang, H., Peng, X., Lin, Y., Xu, M., Ye, G., Zheng, Z., Guo, S. (2023). Scheduling Most Valuable Committees for the Sharded Blockchain. IEEE/ACM Transactions on Networking, 31 (6), 3284–3299. https://doi.org/10.1109/tnet.2023.3278456
9. Windiatmaja, J. H., Hanggoro, D., Salman, M., Sari, R. F. (2023). PoIR: A Node Selection Mechanism in Reputation-Based Blockchain Consensus Using Bidirectional LSTM Regression Model. Computers, Materials & Continua, 77 (2), 2309–2339. https://doi.org/10.32604/cmc.2023.041152
10. Zhang, M., Liu, M., Ding, X., Wang, Y., Li, G. (2025). GPE-PoS: A Fair and Sybil-Resistant Proof of Stake Consensus. Journal of Internet Technology, 26 (4), 463–470. https://doi.org/10.70003/160792642025072604005
11. Castañeda, A., Hurault, A., Quéinnec, P., Roy, M. (2019). Tasks in Modular Proofs of Concurrent Algorithms. Stabilization, Safety, and Security of Distributed Systems, 69–83. https://doi.org/10.1007/978-3-030-34992-9_6
12. Iqbal, M., Matulevicius, R. (2021). Exploring Sybil and Double-Spending Risks in Blockchain Systems. IEEE Access, 9, 76153–76177. https://doi.org/10.1109/access.2021.3081998
13. Solomka, I. R., Liubinskyi, B. B., Peniak, B. O. (2025). Mixed-Weight Committee Selection in Proof-of-Stake: Tunable Stake-Baseline Mixing with Exponential Tail Guarantees and Incentive Compatibility. Mathematical Modeling and Computing, 12 (4), 1320–1332. https://doi.org/10.23939/mmc2025.04.1320
14. Ihoso. Mixweight-committee-consensus. Self-Stabilizing Consensus. Available at: https://github.com/ihosol/mixweight-committee-consensus/tree/feature/self-stabilized-consensus/self-stabilizing-consensus