future internet

MDPI

Article

A Hybrid ABAC–RpBAC Framework for Enhancing PoS Consensus Against Sybil Attacks

Mohammed Al Qurashi * and Ibtihaj Al Qarni

Department of Computer Science, Al-Baha University, Alaqiq 65779-7738, Saudi Arabia; ibtihaj.alqarni@gmail.com

* Correspondence: malqurashi@bu.edu.sa

Abstract

Sybil attacks remain a primary challenge for Proof-of-Stake (PoS) blockchain systems, as low-cost identity creation can distort validator participation and limit consensus reliability. This study proposes a hybrid participation–governance framework that integrates Attribute-Based Access Control (ABAC) and Reputation-Based Access Control (RpBAC) with a trust-based PoS workflow to reduce the influence of suspicious identities during validator selection and block validation. The proposed framework also incorporates graylisting and dynamic reward–penalty updates to support adaptive participation control. The strategy was evaluated in a simulation environment informed by Ethereum-derived block metadata, using network sizes ranging from 100 to 1000 nodes and Sybil attack ratios of 30%, 40%, and 50%. Its performance was compared with PoS-only and PoS + ABAC baselines using both security and performance indicators. The results show that the full ABAC + RpBAC configuration achieved the strongest and most stable security performance across the evaluated settings while introducing additional overhead at larger network sizes. These findings suggest that combining policy-based eligibility control with behavior-based reputation control strengthens the resilience against Sybil in PoS-like blockchain environments. However, this improvement requires a measurable trade-off between security and performance.

Keywords: blockchain; consensus security; proof-of-stake; Sybil attack; access control; ABAC; RpBAC

1. Introduction

Blockchain is a distributed ledger technology that enables multiple nodes to exchange and validate data in a decentralized peer-to-peer environment without relying on a central authority [1]. Its main advantages include decentralization, security, transparency, and immutability [2]. Initially associated with cryptocurrencies and financial applications, blockchain has since expanded into many other domains, including government, healthcare, the Internet of Things, and media, where it has demonstrated significant security and operational value [3]. Proof of Stake (PoS) has become one of the most important blockchain consensus mechanisms because it reduces the computational burden associated with Proof of Work while maintaining decentralized validation through stake-based participation [4]. This model is now used in major blockchain ecosystems such as Ethereum, where validators reach consensus by staking assets rather than relying on computational mining [5,6]. However, public and permissionless blockchain systems remain vulnerable to Sybil attacks, in which adversaries create multiple identities to increase their influence over participation

Academic Editor: Qiang Qu

Received: 8 March 2026
Revised: 12 May 2026
Accepted: 14 May 2026
Published: 22 May 2026

and local decision-making processes [7]. Although stake-based participation introduces an economic deterrent, recent studies indicate that ensuring the resilience against Sybil remains a persistent challenge in open blockchain environments. This has motivated the development of complementary trust-, reputation-, and identity-aware enhancements to PoS-like consensus designs [8–10].

In such environments, an attacker may create multiple low-cost identities to increase visibility in peer discovery, message propagation, or validator candidacy [9,10]. As a result, consensus security is influenced not only by the size of the stake held, but also by which identities are allowed to participate, under what conditions they remain eligible, and how suspicious behavior is handled once it emerges, creating an important gap in PoS-based systems [8,10]. Even when an adversary does not control the majority of the total stake, identity inflation can still be used to distort local network views, increase the presence of malicious candidates, or raise the likelihood of appearing in participation or selection processes [9,10]. In response to this gap, PoS security can be strengthened by integrating participation-control mechanisms that regulate eligibility and adapt influence according to policy and observed behavior. Attribute-Based Access Control (ABAC) introduces explicit rule-driven conditions for participation, such as minimum trust requirements, behavioral constraints, and contextual conditions [11,12]. Reputation-Based Access Control (RpBAC) complements this role using accumulated behavioral evidence to support adaptive participation privileges over time [8]. Rather than relying on static stake assumptions, this combination enables the system to limit eligibility, reduce the influence of suspicious identities, and dynamically restrict nodes that repeatedly violate expected consensus behavior while preserving PoS-style validation among eligible participants [8,10]. Existing studies have explored several strategies to improve blockchain resilience against Sybil attacks, including behavior monitoring, identity verification, reputation-based control, trust-oriented mechanisms, and hybrid consensus designs [7,8,13–16]. Although these approaches provide valuable insights, many of them focus on a single defensive dimension, rely on assumptions that are difficult to maintain in open blockchain environments, or are evaluated in limited or highly specialized settings.

In addition, prior studies have not sufficiently examined how policy-based participation control and behavior-based adaptive restriction can be jointly integrated in PoS-oriented environments to reduce Sybil influence while preserving blockchain performance [8,15,17]. This reveals a gap in the current literature regarding the use of complementary access control mechanisms to regulate validator eligibility and adaptive participation in support of a Sybil-resistant PoS consensus. While Delegated Proof of Stake (DPoS) represents a related consensus variant in which validator participation is mediated through delegated voting, the present study focuses specifically on a non-delegated PoS-oriented validator participation framework. Accordingly, this paper proposes a hybrid ABAC-RpBAC framework for strengthening the Proof-of-Stake consensus against Sybil attacks. The framework combines policy-based validator eligibility, behavior-aware adaptive participation, and dynamic restriction of suspicious nodes. It is evaluated across varying network densities and Sybil attack ratios to assess its effectiveness in improving security while preserving blockchain performance. The remainder of this paper is organized as follows. Section 2 reviews related studies on mitigating Sybil attacks and enhancing blockchain consensus. Section 3 presents the proposed framework and methodology. Section 4 reports and discusses the results. Finally, Section 5 concludes the paper and outlines future research directions.

2. Related Work

This section reviews prior studies on Sybil attack mitigation and consensus enhancement in blockchain systems. The literature is organized according to the main defensive strategies studied in previous work, including trust-based approaches, behavior monitoring, identity verification, reputation-oriented mechanisms, and hybrid consensus designs. This structure helps clarify the strengths and limitations of existing directions while highlighting the remaining gaps addressed in this study.

2.1. Sybil Attacks and the Limits of Stake-Only Participation

Sybil attacks exploit the low cost of identity creation in open networks, enabling adversaries to amplify their influence and disrupt protocol behavior by spawning multiple pseudonymous participants [18]. In permissionless blockchains, Sybil behavior can manifest both at the networking layer, by manipulating peer selection, interfering with message propagation, and exhibiting eclipse-style behaviors, and at the consensus layer by increasing the attacker's exposure to validator candidacy, committee processes, or local influence despite not controlling a majority of the stake [19,20]. Although Proof of Stake (PoS) reduces the computational waste of Proof of Work (PoW), it does not inherently guarantee strong Sybil resistance unless participation and influence are constrained beyond a mere identity count, especially when facing adaptive adversaries and heterogeneous network conditions [5,17]. In this context, recent research increasingly treats Sybil resistance as a cross-layer problem that requires a combination of governance, monitoring, and incentive mechanisms rather than relying on stake-only membership.

2.2. Trust-Graph and Federated Sybil Defenses

A prominent line of work has limited Sybil's influence by leveraging trust relationships or social graphs. Early defenses such as SybilGuard rely on the properties of social network graphs to bound Sybil penetration and isolate Sybil regions [21]. As another example, SybilInfer applies probabilistic inference over social networks to identify Sybil nodes [22]. In blockchain settings, SybilQuorum integrates trust-network assumptions using a ledger design to support open ledgers through trust networks and federated agreement, aiming to preserve quorum intersection when trust links are meaningful [23]. Despite their conceptual appeal, trust-graph approaches rely fundamentally on the existence, quality, and integrity of trust edges. In real-world scenarios, adversaries can leverage weak trust formation, and onboarding can become frictional when trust links are required, or social relations do not naturally exist, thereby limiting both scalability and openness [17,24]. These limitations motivate the use of complementary mechanisms that can still operate when trust edges are limited, noisy, or strategically manipulated.

2.3. Behavior Monitoring and Network/Protocol-Level Detection

Another line of defense against Sybil attacks is to monitor network- and protocol-level signals and then penalize or exclude suspicious nodes. Reference [7] employed distributed behavior monitoring, examining miners' forwarding behavior, and blacklisting any malicious nodes. A strong detection performance was reported in an Ethereum-like validation scenario. Related work in other anonymity and overlay settings has also shown that Sybil behaviors can be detected via characteristic topological or communication patterns. However, these approaches are often sensitive to threshold selection, churn, and adversarial adaptation [19,25]. A common limitation of monitoring-only systems is that they may yield a false positive in benign network variability or a false negative if strategic attackers mimic normal behavior. Furthermore, several of these designs mainly report detection results without jointly quantifying blockchain-level effects, such as throughput

and latency, under adversarial scaling [17]. This creates the need for frameworks that translate observed conduct into long-term participation restrictions rather than relying solely on reactive blacklisting.

2.4. Privacy-Preserving Uniqueness and Identity Verification

Identity-based defenses attempt to limit Sybil attacks by enforcing uniqueness per participant. Reference [13] introduces zk-PoI, a zero-knowledge proof-of-identity approach that aims to enhance Sybil resistance while preserving anonymity. More broadly, privacy-preserving identity and access control constructs in blockchain systems have been surveyed, highlighting both their promise and deployment challenges [26]. Reference [27] proposes verifiable anonymous identities and access control mechanisms for blockchains, emphasizing the need for anonymity and verifiability in controlled environments. While these approaches can decouple Sybil resistance from consensus, they usually depend on external trust anchors or assumptions that undermine truly permissionless decentralization. Most have been validated only conceptually or in limited experimental settings, creating uncertainty regarding the sustainability of large-scale robustness and security–performance trade-offs [17,26].

2.5. Reputation-Based and Hybrid Consensus Hardening

Reputation-centric strategies reduce the marginal benefit of creating new identities by tying influence to accumulated behavior rather than identity count. ReCon [8] couples external reputation systems with consensus to enhance Sybil resistance in large peer-to-peer networks. It has been reported to demonstrate strong resilience in simulations under specific assumptions. TrustChain [14] proposes a Sybil-resistant scalable blockchain using localized validation and a reputation mechanism derived from interaction graphs and evaluates its feasibility in practice. MeritRank similarly targets Sybil-tolerant reputations in tokenomics by emphasizing merit-based influence rather than identity proliferation [28]. Hybrid consensus designs also explore combinations of mechanisms, such as PoW/PoS blends or PoS variants, to improve both security and performance [15,29,30]. Reference [16] integrated a reputation mechanism into a PBFT-style consensus to mitigate Sybil attacks through reputation-weighted selection and decision processes. Although these studies demonstrate promising directions, they also introduce challenges in relation to the integrity of reputation, including collusion and gaming. They also often evaluate either security indicators or performance indicators in isolation rather than jointly under multiple Sybil ratios and network densities [23].

More recent studies have explored integrating PoS with dynamic reputation-based mechanisms to enhance validator selection and strengthen resistance to Sybil attacks in decentralized environments [31]. These approaches further emphasize the importance of combining consensus design with behavior-aware participation control rather than relying on static or single-dimension defenses.

2.6. Access Control for Participation Governance (ABAC/RpBAC) and Remaining Gaps

Access control has increasingly been reframed from data authorization to participation governance in decentralized systems: controlling who can join, propose, validate, vote, and relay can effectively govern state transitions. ABAC provides fine-grained, policy-driven authorization decisions based on subject, object, and context attributes, and has been deployed in blockchain settings through smart contracts to enable auditable and distributed policy enforcement [11,32,33]. Blockchain-based IoT access control systems similarly highlight the role of attribute governance, cross-domain requirements, and operational overhead [34,35]. Recent work has addressed the scalability of ABAC and challenges of decentralization through architectural techniques such as sharding-based and decentral-

ized enforcement [36]. Complementary to ABAC, reputation-driven participation control can be interpreted as a form of RpBAC, in which privileges adapt over time based on behavioral evidence. This reduces Sybil's payoff by making influence earned and revisable [8,28]. However, effective RpBAC requires safeguards against manipulation, including decay, verification, and anti-collusion mechanisms, and must be integrated with eligibility enforcement so that evidence of behavior meaningfully constrains validator participation [14,17]. In addition, recent studies on blockchain design in IoT-driven environments emphasize the importance of balancing security, flexibility, and system performance under dynamic and application-specific conditions [37]. These works highlight how introducing adaptive control mechanisms, while improving system robustness, may introduce additional operational overheads, reinforcing the need for carefully designed participation governance frameworks.

Overall, the literature indicates that strong Sybil resilience typically demands multi-mechanism designs. Nevertheless, a practical gap remains in approaches that jointly (i) enforce explicit policy-based eligibility through ABAC; (ii) support adaptive reputation-based restriction through RpBAC; and (iii) quantify resulting security–performance trade-offs under adversarial scaling using both detection-quality metrics and blockchain performance metrics.

3. Methodology

3.1. Study Design

This section presents a reproducible simulation-based methodology for evaluating a hybrid framework for Sybil mitigation in permissionless blockchain environments. The proposed approach integrates Attribute-Based Access Control (ABAC) and Reputation-Based Access Control (RpBAC) into a trust-based Proof-of-Stake (PoS) validator-selection workflow, supported by dynamic reward–penalty updates and a graylisting mechanism. The framework is designed to regulate validator eligibility and adaptive participation so that suspicious identities can be restricted before exerting undue influence on the consensus process. In addition, the methodology enables a joint assessment of security effectiveness and performance implications for various network sizes and Sybil ratios. The experimental design uses a quantitative simulation approach with controlled baselines and systematically varied attack conditions and network densities. This structure enables comparative evaluations across multiple scenarios and allows the proposed framework to be assessed in terms of both security-oriented and performance-oriented outcomes.

3.2. Benchmark Dataset

To support reproducibility and provide a realistic simulation basis, this study uses cryoto Ethereum [38], a preprocessed public dataset derived from the Ethereum blockchain dataset available on Kaggle [39]. The dataset contains block-level metadata, including block number, timestamp, block size, and transaction count. These attributes are used to inform the simulated blockchain environment and to preserve realistic block characteristics during experimental evaluation. It is important to note that the dataset provides block-level information only. Validator behavior, node participation, and Sybil attack scenarios are simulated within the experimental framework rather than being directly imported from the dataset. Accordingly, the dataset is used to approximate the structural and temporal characteristics of an Ethereum-like blockchain environment, rather than to reproduce the Ethereum network itself. A representative sample of the preprocessed dataset is presented in Table 1.

Table 1. Representative sample of the preprocessed Ethereum block dataset.

Block Number	Hash	Parent Hash	SHA3-Uncles	Timestamp	Transaction-Count	Size
900,000	0x388f34dd...	0xdbfa2da5...	0x1dcc4de8...	12/03/2018	9:59:35	95	41,360
900,001	0x8992ef39...	0x388f34dd...	0x1dcc4de8...	12/03/2018	11:11:01	119	34,751
900,002	0xd18199fd...	0x8992ef39...	0x1dcc4de8...	12/03/2018	19:20:22	69	16,998
900,003	0x1304ee0b...	0xd18199fd...	0x1dcc4de8...	12/03/2018	18:04:51	114	33,292
900,004	0xaa675351...	0x1304ee0b...	0x1dcc4de8...	01/03/2019	11:04:51	142	24,413

Note: Hash fields and other long strings have been truncated for readability only. This formatting does not affect the underlying data structure, model logic, or processing workflow.

3.3. The Experimental Design and Implementation Environment

The proposed framework was evaluated using a quantitative simulation design capable of producing measurable, reproducible, and comparable results across multiple experimental conditions. The evaluation was conducted using controlled baseline settings and systematically varied attack and network parameters in order to assess both security-related and performance-related outcomes. The controlled settings included an underlying consensus workflow, dataset-informed characteristics, simulation environment, and the workload configuration used across all scenarios to maintain consistency and comparability. The experimental variables included the network size, the Sybil attack ratio, and the access control configuration applied in each scenario. These variables were systematically altered to examine their effect on behavior and effectiveness in the proposed framework. Table 2 summarizes the controlled and experimental parameters used in the simulation.

Table 2. Summary of controlled and experimental parameters used in the simulation.

Parameters	Type	Description
Block characteristics	Controlled	Dataset-informed block properties were used consistently across experiments.
Consensus workflow	Controlled	The PoS-oriented workflow and its update logic were held constant across scenarios.
Simulation environment	Controlled	Transaction and block-processing settings were kept comparable across scenarios.
Network size	Experimental	Variations were performed using 100, 500, and 1000 nodes to assess scalability.
Sybil Node Ratio	Experimental	Variations were performed using 30%, 40%, and 50% to simulate different attack intensities.
Access Control configuration	Experimental	PoS-only, PoS + ABAC, and PoS + ABAC + RpBAC scenarios were compared.

All experiments were implemented in Python 3.10 using Google Colab as the development and execution environment. The simulation relied on commonly used libraries for data preprocessing, numerical computation, cryptographic hashing, network modeling, and result visualization. These tools supported dataset handling, blockchain process simulation, validator-selection experiments, and metric analysis. The main tools used in the implementation environment are summarized in Table 3.

Table 3. Main tools used in the implementation environment.

Tool/Library	Purpose
Python 3.10	Core implementation language for simulation and analysis.
Google Colab	Cloud-based development execution environment, accessed during 2025–2026.
Pandas 2.2.2	Data preprocessing and structured dataset handling.
NumPy 2.0.2	Numerical computation and probability-based calculation.
Matplotlib 3.10.0	Visualization of performance and security metrics.
NetworkX 3.6.1	Network modeling and structural analysis.
Hashlib	SHA-256-based hashing for transactions and blocks.
Random/NumPy.random (included in NumPy 2.0.2)	Controlled probabilistic selection and simulation variability.

3.4. Framework Implementation Phases

The proposed framework was implemented through five sequential phases, as illustrated in Figure 1.

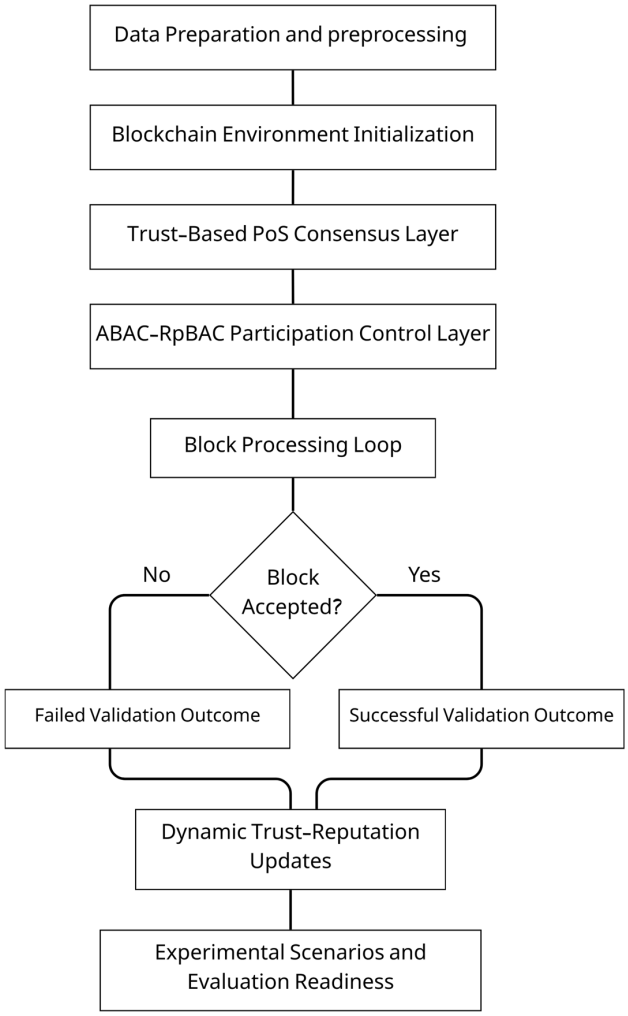

Figure 1. Workflow of the proposed ABAC-RpBAC framework for enhancing PoS consensus against Sybil attacks through five sequential phases.

Phase 1. Data Preparation and Preprocessing:

In the first phase, the benchmark dataset was cleaned and prepared for simulations. Preprocessing included data filtering, format standardization, and inspection of missing values using the Pandas library. Duplicated and incomplete records were removed to improve data consistency. To support a stable and representative simulation basis, block records with extreme sizes were excluded. Specifically, blocks smaller than 5000 bytes were removed because they are often empty or incomplete, whereas blocks larger than 50,000 bytes were excluded because they represent unusually large cases that may distort the experimental distribution. Accordingly, only blocks within the range 5000 < block size < 50,000 bytes were retained. After filtering, the dataset contained 439,334 records. To ensure consistency across experiments, a stratified random sampling procedure was applied to extract 1000 blocks using a fixed seed (random_state = 42). The sampling criteria required non-empty blocks, valid timestamps, and suitable transaction-count values. This phase established a clean and reproducible block-level dataset for the subsequent simulation stages.

Phase 2. Initialization of Blockchain Environment:

After data cleaning and preparation, a simulated blockchain environment was initialized to support evaluation of the proposed framework. First, a scalable set of nodes was generated, and each node was assigned a unique sequential identifier. Transactions were then created according to the transaction-count values derived from the benchmark dataset and hashed using SHA-256. The resulting transaction hashes were aggregated into a Merkle tree to compute a Merkle root for each block. A conceptual representation of this process is shown in Figure 2. Next, the blockchain was initialized with a genesis block, the parent hash of which was set to a null value. Each subsequent block was then hashed using SHA-256, and the resulting hash was embedded as the parent hash of the following block to preserve cryptographic linkage and chain continuity. The cleaned block records were iteratively used to populate the simulated chain based on key structural attributes such as block number, parent hash, timestamp, transaction count, and block size. These attributes are summarized in Table 4. It is important to note that the benchmark dataset provides block-level metadata only. This phase established the foundational blockchain structure required for subsequent consensus, access control, and Sybil mitigation experiments.

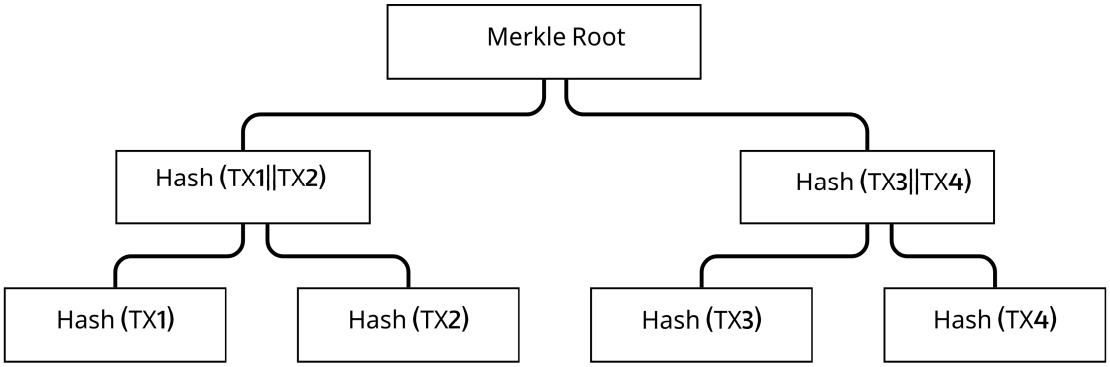

Figure 2. Conceptual representation of Merkle tree construction used to compute the Merkle root from block transactions in the simulated blockchain environment.

Table 4. Summarized block-level attributes used to construct the simulated blockchain records.

Attribute	Description
Block Number	Unique identifier, indicating the position of the block in the simulation chain.
Parent Hash	Hash of the preceding block, used to preserve chain continuity.
Hash	SHA-256 cryptographic hash of the current block.
Transaction Count	Number of transactions included in the block.
Block Size	Size of the block in bytes.
Timestamp	Time associated with block creation in the benchmark dataset.
Merkle Root	Cryptography roots derived from hashed block transaction.
Validator/Node ID	Identifier of the simulated node assigned to propose or validate the block.

Phase 3. Trust-Based PoS Consensus Layer:

After initializing the simulated blockchain environment, this phase establishes the baseline consensus layer used prior to the integration of access control. A trust-based Proof-of-Stake (PoS) workflow is adopted as the underlying consensus mechanism due to its wide use in contemporary blockchain systems and lower computational overhead compared to Proof of Work (PoW). It is also suitable for validator-based participation [4,5]. At the same time, prior studies indicate that PoS does not inherently guarantee strong Sybil resistance in permissionless environments when participation is not sufficiently constrained [17,20].

In this phase, validator selection is modeled through a trust-based weighted mechanism rather than a stake-only policy. All simulated nodes are initialized using a uniform trust range [0.4, 0.8], without differentiation between honest or malicious nodes. This ensures that no prior advantage is assigned at initialization, and that any behavioral differences emerge dynamically during the simulation process. Each node's trust score is updated based on observed validation outcomes, forming the basis for probabilistic selection. For each block, a candidate validator is selected according to the relative weight of its trust among eligible nodes, as follows:

$$P_i = \frac{T_i}{\sum_{j \in E} T_j}$$

where P_i denotes the probability of selecting node i as the validator; T_i is the current trust score of node i; and E represents the set of eligible nodes in the current round.

The selected validator attempts to corroborate and append the block to the chain. If the validation outcome is successful, the trust score of the node is increased according to the reward rule. If the outcome is unsuccessful or associated with suspicious behavior, then the trust score is reduced according to the penalty rule. These updates are applied iteratively to model the evolution of trust over time through the following equation:

$$T_i^{(t+1)} = T_i^{(t)} + \Delta T_i$$

where $T_i^{(t)}$ and $T_i^{(t+1)}$ denote the trust scores of node i before and after the current validation, respectively, and ΔT_i represents the adjustment of trust determined by the observed validation outcome.

To simulate adversarial conditions, each experimental scenario included a predefined proportion of nodes assigned malicious behavioral profiles associated with Sybil-oriented activity. These nodes are designed at scenario initialization to perform disruptive actions,

such as invalid parent hash insertion, block dropping, or inconsistent validation behavior. These behaviors are not triggered by explicit identity labels but are probabilistically activated during the simulation to emulate an adversarial condition. This setup reflects controlled adversarial modeling, where malicious behavior is introduced as part of the experimental design rather than inferred dynamically during its execution.

Algorithm 1 explicitly demonstrates that validator selection and trust updating are performed without access to ground-truth labels, thereby ensuring strict separation between the operational consensus process and the evaluation stage. Accordingly, the consensus workflow is driven exclusively by observable behavioral outcomes, including validation success, failure, and protocol-level inconsistencies. Ground-truth scenario assignments are maintained in a separate evaluation structure and are accessed only after simulation completion for post hoc computation of security metrics.

Algorithm 1. Label-Free Validator Selection and Trust Update in the Proposed PoS-Based Framework

Input: Set of nodes N with trust scores T and validation outcomes O
O denotes the observable validation outcome derived from
protocol-level behavior.
Output: Updated trust scores
1. Initialize nodes with operational attributes:
trust score, validation history, and behavioral state.
2. For each validation round:
Determine eligible nodes E based on:
- current trust score
- recent validation behavior
- consistency of participation
3. Select validator i ∈ E with probability:
$P_i = \frac{T_i}{\sum_{j \in E} T_j}$
4. The selected validator performs block validation.
5. Observe validation outcome O:
- successful validation
- failed validation
- inconsistent behavior
- block anomaly (e.g., invalid parent hash, dropped block)
6. Update trust score:
$T_i^{(t+1)} = T_i^{(t)} + \Delta T_i(O)$
7. Repeat for subsequent rounds.
8. After simulation terminates:
Use ground-truth labels only to compute evaluation metrics.
Note:
Ground-truth labels are not used in steps 2–6 and are excluded from the operational consensus workflow.

While this phase establishes a trust-driven baseline for validator selection, it does not explicitly regulate eligibility for participation. Therefore, the next phase extends this model by introducing policy-based and reputation-aware participation control through ABAC and RpBAC mechanisms.

Phase 4. ABAC, RpBAC Participation Control Layer:

After establishing the trust-based PoS consensus layer, the fourth phase introduces a participation control layer that integrates Attribute-Based Access Control (ABAC) and

Reputation-Based Access Control (RpBAC) into validator selection and block validation processes. The purpose of this layer is to regulate participation eligibility, reduce the influence of suspicious nodes, and support adaptive participation control under Sybil-oriented conditions.

In this phase, each node is associated with a set of attributes used by the participation control layer, including the current trust score, reputation score, recent validation outcomes, graylist status, and contextual participation behavior. ABAC enforces explicit participation conditions, including threshold-based eligibility and behavioral constraints, while RpBAC adapts participation privileges based on accumulated behavioral evidence. Together, these mechanisms extend validator selection beyond the concept of trust by incorporating both policy-based and behavior-aware control.

A graylisting mechanism is introduced as an intermediate filtering layer to identify nodes that require additional scrutiny. Rather than immediately excluding such nodes, the framework subjects them to an additional verification step before the block is accepted. Graylisting is triggered when trust, reputation, or behavioral consistency falls below predefined thresholds. In the current simulation, graylisting is triggered when one or more of the following conditions are met: (trust score < 0.50, reputation score < 0.45, or suspicion score > 0.60). These thresholds were selected to remain compatible with the initial trust range used in the simulation while preserving a balanced verification region between direct acceptance and exclusion. They were chosen to moderate two competing risks: overly strict graylisting may increase false positives and verification overhead, whereas overly relaxed graylisting may increase false negatives by allowing suspicious nodes to proceed with less scrutiny.

This graylist stage was intentionally introduced to improve the realism of the simulation and to avoid premature classification decisions. In practical blockchain environments, newly joined nodes may not exhibit sufficient behavioral evidence at the beginning of participation to justify immediate classification as either trustworthy or malicious. Therefore, instead of directly excluding such nodes, the framework places them under temporary observation and additional verification until sufficient behavioral evidence is accumulated. Once their behavior is validated and no suspicious pattern is confirmed, they may transition into normal participation. This design aims to provide a more realistic and fair representation of validator participation dynamics in open blockchain environments.

For each processed block, validator selection is performed probabilistically among eligible nodes according to a combined trust–reputation weighting rule, as follows:

$$P_i = \frac{W_i}{\sum_{j \in E} W_j}$$

where P_i denotes the probability of selecting node i as the validator; W_i is the composite participation weight of node I; and E represents the set of eligible nodes in the current round. In this framework, W_i is derived from the node's trust and reputation scores, as well as its eligibility status after policy-based filtering. Once selected, the validator attempts to corroborate and append the block to the chain. If the selected node is graylisted, the block undergoes an additional behavioral verification step based on recent trust, consistent reputation, and history of validation. Blocks that pass this verification step are accepted, whereas those that fail are rejected, and the corresponding node is subjected to stricter control in subsequent rounds.

After each validation round, trust and reputation scores are updated according to the observed validation outcomes, modeling the evolution of participation reliability over time as follows:

$$S_i^{(t+1)} = S_i^{(t)} + \Delta S_i$$

where $S_i^{(t)}$ and $S_i^{(t+1)}$ denote the current and updated score states of node i, respectively, and ΔS_i represents the adjustment induced by the validation outcome. Positive updates correspond to successful and consistent behavior, while negative updates correspond to failed or suspicious actions.

Consistent with the label-free execution principle illustrated in Algorithm 1, participation decisions in this phase are derived solely from dynamically updated trust scores, graylist status, and observable behavioral outcomes. Ground-truth labels remain separated from the operational workflow and are accessed only after simulation completion for post hoc evaluation.

Phase 5. Experimental Scenario and Evaluation of Readiness:

In the final phase, the simulation framework is executed under three controlled scenarios to support the comparative evaluation of the proposed model. Scenario 1 represents the baseline PoS-only setting without the proposed participation control. Scenario 2 incorporates PoS with ABAC, and Scenario 3 represents the full proposed framework integrating PoS, ABAC, and RpBAC. These scenarios enable a structured comparison between the baseline and enhanced configurations under different network sizes and Sybil attack ratios.

To ensure stability and reproducibility, each experimental configuration was executed 10 times under the same parameter settings. The reported values in Tables 5 and 6 are presented as Mean ± Standard Deviation (Mean ± Std) across these repeated runs. This statistical reporting was adopted to make run-to-run variability explicit and to provide a clearer assessment of result stability.

3.5. Evaluation Metrics

To evaluate the proposed framework under different network densities and Sybil attack ratios, a set of security and performance metrics was computed directly from the simulation logs.

Security metrics assess the ability of the framework to identify malicious participation. This include the True Positive Rate (TPR), False Negative Rate (FNR), and overall Accuracy and F1-score:

$$TPR = \frac{TP}{TP + FN} \tag{1}$$

TPR measures the proportion of malicious nodes that were correctly identified as Sybil nodes.

$$FNR = \frac{FN}{FN + TP} \tag{2}$$

FNR measures the proportion of malicious nodes that were not identified by the detection process.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{3}$$

Accuracy measures the proportion of correctly classified nodes, including both malicious and legitimate nodes. Here, TP denotes true positives, TN denotes true negatives, FP denotes false positives, and FN denotes false negatives.

$$F1\text{-}score = \frac{2TP}{2TP + FP + F\,N} \tag{4}$$

F1-score provides a harmonic summary of detection effectiveness and is particularly useful under class-imbalanced conditions, where overall accuracy alone may overestimate security performance.

A supplementary summary of the F1-score metric under different network sizes and Sybil attack ratios is provided in Appendix A Table A2.

Table 5. Security evaluation metrics across experimental scenarios under different network sizes and Sybil attack ratios (Mean ± Std over 10 runs).

Scenario	Nodes	TPR (%)			FNR (%)			Accuracy (%)		
		30%	40%	50%	30%	40%	50%	30%	40%	50%
PoS-Only	100	43.21 ± 0.50	66.66 ± 0.29	75.66 ± 0.30	56.69 ± 0.50	33.34 ± 0.29	24.24 ± 0.30	75.86 ± 0.25	88.78 ± 0.14	88.48 ± 0.15
	500	19.64 ± 0.35	47.23 ± 0.35	61.91 ± 0.41	80.36 ± 0.35	52.07 ± 0.35	37.99 ± 0.41	90.02 ± 0.17	89.22 ± 0.17	89.01 ± 0.20
	1000	25.81 ± 0.30	49.45 ± 0.36	41.34 ± 0.42	74.19 ± 0.30	50.55 ± 0.36	58.66 ± 0.42	93.01 ± 0.15	92.87 ± 0.18	86.52 ± 0.21
PoS + ABAC	100	48.41 ± 0.49	52.49 ± 0.54	64.47 ± 0.53	51.69 ± 0.49	40.41 ± 0.54	35.73 ± 0.53	66.95 ± 0.25	71.95 ± 0.27	81.59 ± 0.27
	500	80.20 ± 0.30	78.48 ± 0.47	81.70 ± 0.29	19.80 ± 0.30	21.52 ± 0.47	18.20 ± 0.29	94.05 ± 0.15	90.59 ± 0.23	91.65 ± 0.14
	1000	80.83 ± 0.27	79.64 ± 0.36	79.37 ± 0.37	19.07 ± 0.27	20.36 ± 0.36	21.13 ± 0.37	94.42 ± 0.14	91.52 ± 0.18	89.74 ± 0.19
PoS + ABAC + RpBAC	100	96.64 ± 0.31	94.95 ± 0.50	96.10 ± 0.45	3.39 ± 0.31	5.05 ± 0.50	3.90 ± 0.45	89.97 ± 0.16	93.97 ± 0.25	97.05 ± 0.22
	500	95.16 ± 0.29	96.55 ± 0.42	95.97 ± 0.47	4.84 ± 0.29	3.95 ± 0.42	4.03 ± 0.47	95.93 ± 0.15	96.53 ± 0.21	95.99 ± 0.24
	1000	94.94 ± 0.50	94.82 ± 0.41	96.83 ± 0.27	5.06 ± 0.50	5.28 ± 0.41	3.17 ± 0.27	95.27 ± 0.25	95.81 ± 0.21	95.82 ± 0.13

Note: Values are reported as Mean ± Standard Deviation over 10 repeated runs for each experimental configuration.

Table 6. Performance-oriented results across experimental scenarios under different network sizes and Sybil attack ratios (Mean ± Std over 10 runs).

Scenario	Nodes	TPS (Tx/s)			Time Latency (ms)		
		30%	40%	50%	30%	40%	50%
PoS-Only	100	81.11 ± 0.83	85.85 ± 0.66	80.00 ± 0.42	77.89 ± 0.83	76.95 ± 0.66	82.00 ± 0.42
	500	80.95 ± 0.42	77.96 ± 0.70	72.80 ± 0.94	76.05 ± 0.42	80.24 ± 0.70	78.70 ± 0.94
	1000	82.68 ± 0.53	71.10 ± 0.85	75.58 ± 0.43	79.62 ± 0.53	77.90 ± 0.85	77.72 ± 0.43
PoS + ABAC	100	83.10 ± 0.67	82.78 ± 0.81	70.94 ± 0.40	89.40 ± 0.67	81.72 ± 0.81	78.26 ± 0.40
	500	77.22 ± 0.61	82.12 ± 0.48	75.07 ± 0.59	89.78 ± 0.61	90.68 ± 0.48	89.63 ± 0.59
	1000	79.09 ± 0.53	71.35 ± 0.67	71.95 ± 0.50	90.61 ± 0.53	87.55 ± 0.67	90.05 ± 0.50
PoS + ABAC + RpBAC	100	82.26 ± 0.39	85.05 ± 0.72	59.86 ± 0.75	99.04 ± 0.39	98.95 ± 0.72	89.64 ± 0.75
	500	79.37 ± 0.50	81.68 ± 0.62	66.87 ± 0.55	97.03 ± 0.50	102.32 ± 0.62	103.13 ± 0.55
	1000	56.74 ± 0.58	40.85 ± 0.62	33.19 ± 0.73	101.26 ± 0.58	103.25 ± 0.62	104.81 ± 0.73

Note: Values are reported as Mean ± Standard Deviation over 10 repeated runs for each experimental configuration. As Blockchain Accuracy was used as a supplementary system-level metric, its detailed repeated-run summary is provided separately in Appendix A Table A1.

Performance metrics are used to evaluate the operational impact of the proposed framework:

$$TPS = \frac{N_{succ}}{T_{sim}} \tag{5}$$

Throughput (TPS) measures the number of successfully processed transactions per second during the simulation runtime, where N_{succ} is the number of validated transactions and T_{sim} is the total simulation time.

$$t_{tx}^{(i)} = t_{proc}^{(i)} + t_{prop}^{(i)} + t_{val}^{(i)} + t_{fin}^{(i)} \tag{6}$$

Transaction Latency represents the total time required for transaction i to be processed, propagated, validated, and finalized.

Blockchain Accuracy (Supplementary Metric):

$$Blockchain\ Accuracy = \left(\frac{VB}{VB + DB} \times 100\right) \times (1 - SF \times \alpha) + (T_{avg} \times \beta) \tag{7}$$

where VB denotes the number of successfully validated and appended blocks; DB denotes the number of dropped or rejected blocks; SF represents the dynamic Sybil-related penalty factor; and T_{avg} denotes the average trust-related stability term used in the simulation.

The weighting constants α and β were predefined and fixed across all scenarios prior to experimental execution. In this study, $\alpha = 0.6$ and $\beta = 6$. These values were selected to balance the influence of adversarial impact and trust stability within the composite metric. Specifically, α assigns a moderate to high penalty to Sybil-related influence, ensuring that adversarial behavior significantly affects the metrics without dominating it. β scales the trust component to ensure its effective contribution despite its smaller numerical range. These parameters were defined prior to experimentation based on normalization and balance considerations and were kept fixed across all scenarios without post hoc tuning.

A supplementary sensitivity analysis of the Blockchain Accuracy metric under alternative α and β settings is provided in Appendix A Figure A1.

The proposed metric is intended to reflect system-level blockchain reliability under adversarial conditions. It complements, rather than replaces, standard evaluation metrics by capturing aspects of system behavior not fully reflected in classification and performance metrics alone.

Overall, TPR, FNR, Accuracy, TPS, and Transaction Latency were treated as the primary evaluation metrics, while Blockchain Accuracy was used as a supplementary composite indicator to support scenario-level comparison.

4. Results

This section reports the evaluation results of the proposed framework under the three experimental scenarios described in Section 3.

The analysis focuses on both security effectiveness and performance implications across different network sizes and Sybil attack ratios. Security-related results are presented first, followed by performance-oriented analysis and comparative visualization.

4.1. Comparative Evaluation Across Scenarios

The proposed framework was evaluated under three experimental scenarios: PoS-only, PoS + ABAC, and PoS + ABAC + RpBAC. Then, a comparison was conducted across three network sizes (100, 500, and 1000 nodes) and three Sybil attack ratios (30%, 40%, and 50%).

Table 5 summarizes the security-oriented metrics, and Table 6 reports the corresponding performance-oriented metrics.

Table 5 shows that the PoS-only scenario exhibited unstable security performance across the evaluated settings, with relatively low detection capability (TPR) and high false negative rates (FNRs), particularly at medium and large network sizes under the same attack ratios. The PoS + ABAC scenario demonstrates a clear improvement in detection performance, especially with 500 and 1000 nodes, where TPR increases significantly, and FNR decreases accordingly. However, with a smaller network size (100 nodes), this improvement was less consistent, and in some cases, did not exceed the baseline of the PoS-only performance. This behavior can be attributed to the limited interaction volume in small-scale environments, where behavior-based differentiation is less effective, and the overhead of policy enforcement may restrict participation without providing sufficient evidence for accurate classification. In contrast, the full PoS + ABAC + RpBAC configuration achieves the most consistent and robust security performance across all evaluated conditions. The TPR remains above 94% in all network sizes, while the FNR is consistently maintained at low levels (below 5%). Accuracy also remains high and stable, indicating that the combined participation control layer significantly enhances Sybil detection and overall system reliability.

Table 6 indicates that the PoS-only scenario achieved the highest throughput and lowest latency across most settings; however, this was at the cost of weaker blockchain-level robustness under Sybil attack conditions. The PoS + ABAC scenario provides a more balanced profile, improving Blockchain Accuracy, particularly at moderate overhead. In contrast, the full scenario produced the highest Blockchain Accuracy values, particularly at moderate and large network sizes; however, this improvement is accompanied by lower throughput and higher latency due to the additional overhead associated with participation control.

Notably, performance degradation becomes more evident under higher network density and Sybil attack ratios, where throughput decreases significantly.

This observed throughput reduction is associated with the cumulative overhead of participation-control operations, including eligibility filtering, reputation-based evaluation, and additional graylist verification steps. More specifically, the main computational burden is primarily attributed to the RpBAC layer and the graylist verification stage rather than the ABAC filtering step itself. While ABAC mainly performs an initial eligibility screening, RpBAC requires repeated evaluation and updating of node reputation and trust-related behavior, which becomes increasingly expensive as the number of nodes grows. In addi-

tion, the graylist mechanism introduces intensive secondary verification for suspicious or borderline nodes before participation decisions are finalized. Under dense and highly adversarial settings, this additional inspection is triggered more frequently, making graylist processing a likely bottleneck and contributing substantially to the observed TPS reduction.

This behavior reflects the additional processing required for validating node behavior and filtering potentially malicious participants. Overall, these results confirm the presence of a clear security–performance trade-off, which becomes more pronounced as network size and adversarial conditions increase.

4.2. Visualization of Security–Performance Trends

Figures 3–5 collectively illustrate the overall behavior of the three evaluated scenarios across different network sizes and attack conditions. Figures 3–5 show that the integrated PoS + ABAC + RpBAC model consistently achieves the strongest security performance, as reflected by higher TPR and Accuracy values and lower FNRs across 100, 500, and 1000 nodes.

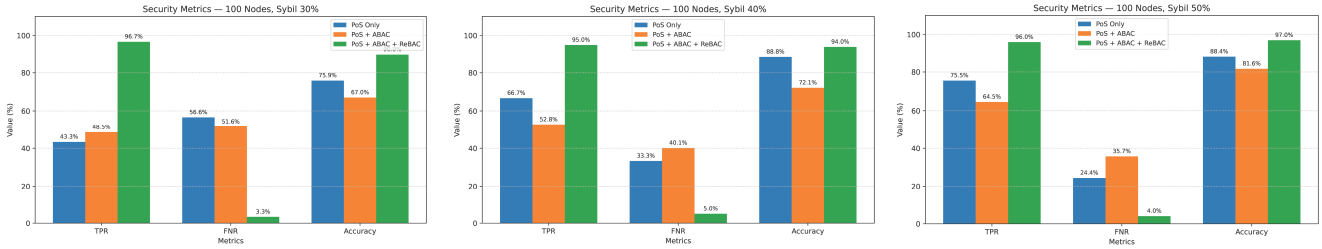

Figure 3. Comparative security metrics for the three evaluated scenarios under 100 nodes: (**a**) 30% Sybil ratio, (**b**) 40% Sybil ratio, and (**c**) 50% Sybil ratio.

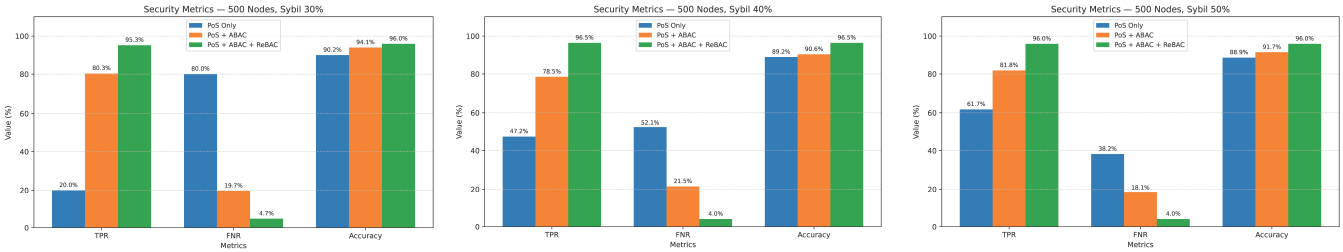

Figure 4. Comparative security metrics for the three evaluated scenarios under 500 nodes: (**a**) 30% Sybil ratio, (**b**) 40% Sybil ratio, and (**c**) 50% Sybil ratio.

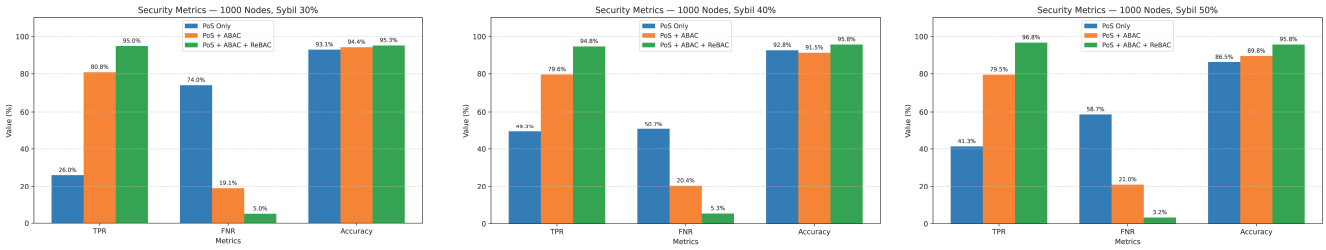

Figure 5. Comparative security metrics for the three evaluated scenarios under 1000 nodes: (**a**) 30% Sybil ratio, (**b**) 40% Sybil ratio, and (**c**) 50% Sybil ratio.

Figure 6 complements this analysis by showing the relationship between TPS and Blockchain Accuracy, highlighting the trade-off between security enhancement and transaction throughput. Together, these figures confirm that the proposed integrated model provides the most robust resistance against Sybil attacks while maintaining stable blockchain performance across different network densities.

The performance plots for 100, 500, and 1000 nodes support this combined interpretation of throughput and Blockchain Accuracy.

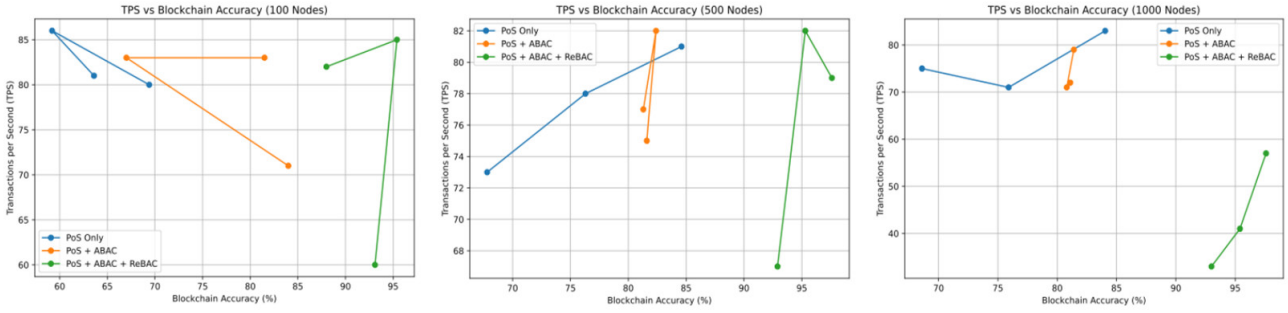

(**a**) (**b**) (**c**)

Figure 6. Comparative visualization of the trade-off between TPS and Blockchain Accuracy for the three evaluated scenarios under different network densities: (**a**) 100 nodes, (**b**) 500 nodes, and (**c**) 1000 nodes.

4.3. Analysis of Security–Performance Trade-off

The experimental results reveal a clear trade-off between security and performance across the evaluated scenarios. The PoS-only configuration achieves the highest throughput and lowest latency; however, it suffers from a limited Sybil detection capability and lower Blockchain Accuracy, particularly under high attack ratios and network densities.

The PoS + ABAC scenario improves security performance by filtering malicious nodes through policy-based constraints, resulting in higher TPRs and reduced FNRs compared to the baseline. This improvement is achieved with moderate overhead, leading to a slight reduction in throughput and an increase in latency.

In contrast, the proposed PoS + ABAC + RpBAC model consistently achieves the highest security performance, maintaining TPR values above 94% and low FNRs across all configurations. This enhanced security comes at the cost of reduced throughput and increased latency due to the additional overhead associated with participation control mechanisms. The impact becomes more pronounced in high-density networks and under elevated Sybil attack ratios. This reduction in the enhanced scenarios should not be attributed to a single mechanism in isolation. Rather, it results from the cumulative overhead of several coordinated participation-control operations. First, ABAC-based filtering restricts validator eligibility according to predefined policy conditions, thereby reducing the immediately available validator pool. Second, RpBAC introduces additional trust- and reputation-based evaluation steps before validator participation is confirmed. Third, the graylist mechanism adds an intermediate verification stage for nodes whose behavior has not yet been sufficiently established, requiring additional monitoring and validation before final participation decisions are made.

Despite this degradation in performance, the proposed model maintains high Blockchain Accuracy, indicating strong system-level reliability under adversarial conditions. Overall, the results confirm that the proposed framework effectively balances security and performance, prioritizing robust Sybil resistance while preserving acceptable operational efficiency.

Table 7 compares the proposed framework with representative prior studies that address Sybil mitigation through trust graphs, behavior monitoring, identity verification, reputation-based mechanisms, and hybrid consensus designs.

Table 7. Qualitative comparison of representative studies on Sybil mitigation in blockchain consensus.

Study	Year	Strategy	Implementation Environment	Evaluation Scale	Security/Performance Evaluation	Reported Outcomes
Sonnino and Danezis [23]	2019	Trust graph hybrid	Open distributed ledger through trust networks/FBAS evaluation on social graphs	Real-world social graphs; exact node count not explicitly reported	Security only	Empirical Sybil-thwarting/trust-network security evaluation
Swathi et al. [7]	2019	Behavior monitoring	Blockchain test bed/distributed behavior monitoring of miners	20–40 nodes	Both	95% TPR; 87.8% Accuracy
Sánchez [13]	2019	Identity verification (zk-PoI)	Theoretical model using zk-SNARKs and TEEs	No explicit blockchain node-scale evaluation reported	Security/theoretical only	Theoretical Sybil-resistance and anonymous authentication claims
Biryukov and Feher [8]	2020	Reputation-based strategy	Proof-of-concept simulator/reputation consensus over BFT-style committees	Simulation with variable node count and committee size as inputs; typical committee size discussed as 100	Security only	Strong Sybil resistance claimed 95% system robustness reported; compares favorably to maliciousness-detection methods
Otte et al. [14]	2020	Trust/reputation-based hybrid	P2P ledger/live-network experiment with TrustChain + NetFlow; additional throughput experiments	917-node live-network experiment; plus throughput tests on mobile devices and PCs	Both	Formal Sybil-resistance argument for NetFlow and throughput higher than traditional blockchain architectures reported
Siddiqui et al. [15]	2023	Hybrid PoS design (BlockPower + PoS)	Mathematical/theoretical protocol analysis with security proofs and performance analysis	No explicit node-count evaluation reported in the uploaded text	Both	Sybil-resistant block-power function claimed; improved TPS and time to finality reported
Wang and Tan [16]	2023	Reputation + PBFT hybrid	Python-based blockchain simulation of improved PBFT with reputation weighting	10 nodes	Both	Reduced attack success rate and increased attack cost reported; performance stated as not significantly decreased

Table 7. *Cont.*

Study	Year	Strategy	Implementation Environment	Evaluation Scale	Security/Performance Evaluation	Reported Outcomes
Umar et al. (2025) [31]	2025	Hybrid PoS + dynamic reputation scoring	Decentralized microgrid/energy system (application-oriented study)	Up to 1000 nodes	Both	99% Sybil-attack mitigation; consensus within 8 s for up to 1000 nodes; 12% peak-demand reduction; 83% load-shifting efficiency; 5.26% cost-savings improvement
This study	2026	ABAC + RpBAC with trust-based PoS	Python simulation in an Ethereum-like environment	100, 500, and 1000 nodes	Both	Up to 96% TPR; 97% Accuracy; TPS and latency jointly evaluated, with acceptable performance at 100–500 nodes and more noticeable overhead at 1000 nodes under higher Sybil ratios

Note: The table is intended to provide a qualitative comparison of study directions and the scope of evaluation rather than a strict quantitative benchmark. The direct numerical comparison between studies is limited due to differences in experimental settings, network scale, and threat models used.

Compared with these studies, the proposed framework differs in two main respects. First, it combines policy-based participation control and reputation-aware adaptive restriction within a PoS-oriented workflow rather than relying on a single defensive mechanism. Second, it is evaluated across multiple network sizes and Sybil attack ratios using both security and performance indicators. Several prior studies were either theoretical, limited to small-scale simulations, or did not explicitly report both detection-oriented and blockchain-level performance outcomes. Within the evaluated simulation setting of this study, the proposed framework achieved strong and stable security performance. The associated trade-off between security and performance was observed across different densities.

5. Discussion and Conclusions

This study investigated a hybrid consensus-hardening framework that integrates ABAC and RpBAC within a trust-oriented PoS workflow to mitigate Sybil attacks. The results demonstrate that regulating participation through policy-based eligibility and behavior-aware adaptation significantly improves Sybil resilience compared to a PoS-only baseline.

From a security perspective, the PoS-only configuration exhibited unstable detection performance when the network density and attack intensity increased, confirming the limitations of stake-based participation in open environments. The introduction of ABAC improved detection stability by enforcing explicit participation constraints, while the combined ABAC–RpBAC configuration achieved consistently high detection performance across all evaluated scenarios. These findings suggest that integrating policy-based control with behavior-driven adaptation provides more robust protection than relying on a single defensive mechanism.

From a performance perspective, the results reveal a clear trade-off between security and performance. While the PoS-only model achieved higher throughput and lower latency, it experienced limitations such as reduced blockchain robustness under adversarial conditions. In contrast, the proposed framework improved system reliability, as reflected in both standard metrics and the supplementary Blockchain Accuracy indicator, at the cost of additional computational overhead, particularly at larger network sizes.

The findings suggest that treating participation as a governed process rather than an unrestricted mechanism can enhance the security of permissionless blockchain systems. By combining ABAC, RpBAC, and graylisting within the consensus workflow, the proposed approach provides a scalable and practical method for enhancing Sybil resistance while maintaining acceptable performance under realistic operating conditions.

It is important to note that the initialization of trust scores within a fixed range [0.4, 0.8] represents a simplifying assumption adopted to ensure a neutral starting point without introducing prior bias among nodes. While this choice supports fair behavioral differentiation during the simulation, alternative initialization ranges may influence early-stage dynamics, particularly the convergence rate of trust adaptation and validator selection.

In addition, the experimental evaluation is conducted within a controlled simulation environment informed by Ethereum-like data. Although this setup supports reproducibility and comparative analysis, it may not fully capture real-world network dynamics, such as heterogeneous node behavior, variable network latency, and long-term adversarial adaptation over extended periods.

Furthermore, the proposed framework prioritizes security improvements through participation control, which introduces additional computational overhead. This overhead is mainly associated with the repeated reputation-based evaluation in the RpBAC layer and extra verification required by the graylist stage, especially under dense and highly

adversarial settings. While this trade-off is explicitly evaluated, its impact may vary under different deployment scenarios and network configurations.

Overall, these factors may influence the generalizability of the results and should be considered when interpreting the performance of the proposed framework in broader or real-world blockchain environments.

Building on these limitations, future work will explore the integration of intelligent prediction mechanisms to optimize participation control under varying network conditions. Machine learning-based approaches will be investigated to balance the overhead associated with participation control, especially in high-density networks and under elevated Sybil attack ratios, thereby improving throughput while maintaining strong security guarantees. In addition, future work will further investigate the graylisting mechanism through adaptive threshold tuning and sensitivity analysis under different initialization ranges, attack ratios, and network densities. Particular attention will be given to balancing fairness for newly joined benign nodes with the need to reduce unnecessary verification overhead and maintain effective Sybil detection.

Author Contributions: Conceptualization, M.A.Q. and I.A.Q.; methodology, I.A.Q.; software, I.A.Q.; validation, M.A.Q. and I.A.Q.; formal analysis, I.A.Q.; investigation, I.A.Q.; resources, I.A.Q.; data curation, I.A.Q.; writing—original draft preparation, I.A.Q.; writing—review and editing, M.A.Q. and I.A.Q.; visualization, M.A.Q. and I.A.Q.; supervision, M.A.Q. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The raw data used in this study are publicly available from the Ethereum blockchain dataset on Kaggle (BigQuery Ethereum Dataset) at https://www.kaggle.com/datasets/bigquery/ethereum-blockchain (accessed on 28 February 2026); the processed dataset generated during this study (cryptoEthereum) is available at https://drive.google.com/file/d/18UffKfUvskS6d36IpPWRWilTRoBhTDjl/view (accessed on 28 February 2026).

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

Table A1. Blockchain Accuracy across experimental scenarios under different network sizes and Sybil attack ratios (Mean ± Std over 10 runs).

Scenario	Nodes	Blockchain Accuracy 30%	Blockchain Accuracy 40%	Blockchain Accuracy 50%
PoS-Only	100	63.63 ± 0.21	59.16 ± 0.17	69.40 ± 0.10
PoS-Only	500	84.59 ± 0.10	76.29 ± 0.17	67.75 ± 0.24
PoS-Only	1000	83.97 ± 0.13	75.93 ± 0.21	68.74 ± 0.11
PoS + ABAC	100	81.52 ± 0.17	66.95 ± 0.20	83.99 ± 0.10
PoS + ABAC	500	81.36 ± 0.15	82.43 ± 0.12	81.62 ± 0.15
PoS + ABAC	1000	81.43 ± 0.13	80.89 ± 0.17	81.09 ± 0.13
PoS + ABAC + RpBAC	100	88.07 ± 0.10	95.41 ± 0.18	93.06 ± 0.19
PoS + ABAC + RpBAC	500	97.69 ± 0.13	95.22 ± 0.15	92.86 ± 0.14
PoS + ABAC + RpBAC	1000	97.54 ± 0.15	95.36 ± 0.15	93.05 ± 0.18

Note: Values are reported as Mean ± Standard Deviation over 10 repeated runs for each experimental configuration.

Table A2. F1-score across experimental scenarios under different network sizes and Sybil attack ratios (Mean ± Std over 10 runs).

Scenario	Nodes	F1-Score 30%	F1-Score 40%	F1-Score 50%
PoS-Only	100	51.79 ± 0.55	82.61 ± 0.25	86.78 ± 0.19
PoS-Only	500	0.00 ± 0.00	77.80 ± 0.41	84.92 ± 0.32
PoS-Only	1000	0.00 ± 0.00	84.73 ± 0.42	75.41 ± 0.47
PoS + ABAC	100	46.78 ± 0.44	59.95 ± 0.47	77.79 ± 0.39
PoS + ABAC	500	89.00 ± 0.29	86.96 ± 0.35	90.73 ± 0.18
PoS + ABAC	1000	89.68 ± 0.26	88.25 ± 0.26	88.55 ± 0.23
PoS + ABAC + RpBAC	100	85.25 ± 0.24	92.65 ± 0.32	97.02 ± 0.23
PoS + ABAC + RpBAC	500	93.35 ± 0.24	95.69 ± 0.27	95.98 ± 0.25
PoS + ABAC + RpBAC	1000	92.33 ± 0.41	94.76 ± 0.26	95.86 ± 0.14

Note: Values are reported as Mean ± Standard Deviation over 10 repeated runs for each experimental configuration.

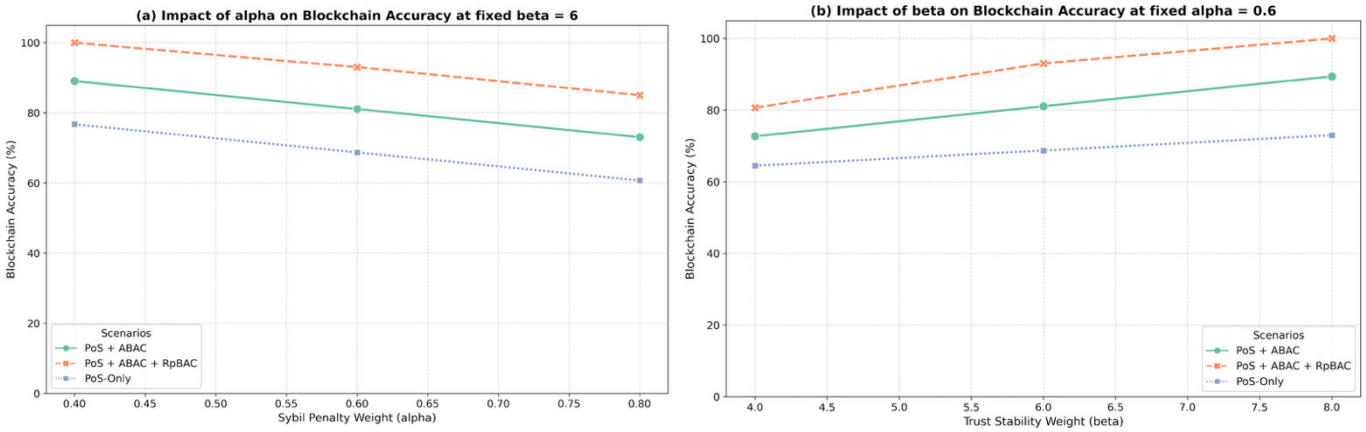

Figure A1. Sensitivity of the supplementary Blockchain Accuracy metric to variations in the weighting constants α and β under a representative high-stress configuration (1000 nodes, 50% Sybil ratio). Panel (**a**) shows the effect of varying α while fixing $\beta = 6$, and panel (**b**) shows the effect of varying β while fixing $\alpha = 0.6$. The comparative scenario-level trend remains consistent across the tested parameter ranges.

References

1. Zheng, Z.; Xie, S.; Dai, H.-N.; Chen, X.; Wang, H. Blockchain challenges and opportunities: A survey. *Int. J. Web Grid Serv.* **2018**, *14*, 352–375. [CrossRef]
2. Yli-Huumo, J.; Ko, D.; Choi, S.; Park, S.; Smolander, K. Where is current research on blockchain technology?—A systematic review. *PLoS ONE* **2016**, *11*, e0163477. [CrossRef] [PubMed]
3. Abou Jaoude, J.; Saade, R.B. Blockchain applications—Usage in different domains. *IEEE Access* **2019**, *7*, 45360–45381. [CrossRef]
4. Wang, W.; Hoang, D.; Hu, P.; Xiong, Z.; Niyato, D.; Wang, P.; Wen, Y.; Kim, D. A survey on consensus mechanisms and mining strategy management in blockchain networks. *IEEE Access* **2019**, *7*, 22328–22370. [CrossRef]
5. Lashkari, B.; Musilek, P. A comprehensive review of blockchain consensus mechanisms. *IEEE Access* **2021**, *9*, 43620–43652. [CrossRef]
6. Ethereum.org. Proof-of-Stake(PoS). Available online: https://ethereum.org/developers/docs/consensus-mechanisms/pos/ (accessed on 7 May 2025).
7. Swathi, P.; Modi, C.; Patel, D. Preventing Sybil attack in blockchain using distributed behavior monitoring of miners. In Proceedings of the 2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT), Kanpur, India, 6–8 July 2019.
8. Biryukov, A.; Feher, D. ReCon: Sybil-resistant consensus from reputation. *Pervasive Mob. Comput.* **2020**, *61*, 101109. [CrossRef]
9. Platt, M.; Platt, D.; McBurney, P. Sybil attack vulnerability trilemma. *Int. J. Parallel Emerg. Distrib. Syst.* **2024**, *39*, 446–460. [CrossRef]
10. Platt, M.; McBurney, P. Sybil attacks on identity-augmented Proof-of-Stake. *Comput. Netw.* **2021**, *199*, 108424. [CrossRef]
11. Rouhani, S.; Butterworth, R.; Culley, R.; Scrivens, N.; Deters, R. Distributed attribute-based access control system using permissioned blockchain. *World Wide Web* **2021**, *24*, 1617–1639. [CrossRef]

12. Malik, S.; Shah, M.A. Access control using blockchain: A taxonomy and review. In *Proceedings of the 2022 6th International Conference on Information System and Data Mining (ICISM 2022), Silicon Valley, CA, USA, 27–29 May 2022*; Association for Computing Machinery: New Your, NY, USA, 2022; pp. 46–54.
13. Sánchez, D.C. Zero-knowledge proof-of-identity: Sybil-resistant, anonymous authentication on permissionless blockchains and incentive compatible, strictly dominant cryptocurrencies. *arXiv* **2019**, arXiv:1902.02224. [CrossRef]
14. Otte, P.; de Vos, M.; Pouwelse, J. TrustChain: A Sybil-resistant scalable blockchain. *Future Gener. Comput. Syst.* **2020**, *107*, 770–780. [CrossRef]
15. Siddiqui, S.; Srivastava, V.; Maheshwari, R.; Gujar, S. QuickSync: A quickly synchronizing PoS-based blockchain protocol. In Proceedings of the 2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Dubai, United Arab Emirates, 1–5 May 2023; pp. 1–2.
16. Wang, Y.; Tan, M. Defense against Sybil attack in blockchain based on improved consensus algorithm. In Proceedings of the 2023 IEEE International Conference on Control, Electronics and Computer Technology (ICCECT), Jilin, China, 28–30 April 2023.
17. Platt, M.; McBurney, P. Sybil in the haystack: A comprehensive review of blockchain consensus mechanisms in search of strong Sybil attack resistance. *Algorithms* **2023**, *16*, 34. [CrossRef]
18. Douceur, J.R. The Sybil attack. In *Peer-to-Peer Systems: First International Workshop, IPTPS 2002, Cambridge, MA, USA, 7–8 March 2002*; Revised Papers; Springer: Berlin/Heidelberg, Germany, 2002.
19. Neudecker, T.; Hartenstein, H. Network layer aspects of permissionless blockchains. *IEEE Commun. Surv. Tutor.* **2018**, *21*, 838–857. [CrossRef]
20. Iqbal, M.; Matulevičius, R. Exploring Sybil and double-spending risks in blockchain systems. *IEEE Access* **2021**, *9*, 76153–76177. [CrossRef]
21. Yu, H.; Kaminsky, M.; Gibbons, P.B.; Flaxman, A. SybilGuard: Defending against Sybil attacks via social networks. In *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '06), Pisa, Italy, 11–15 September 2006*; Association for Computing Machinery: New York, NY, USA, 2006; pp. 267–278.
22. Danezis, G.; Mittal, P. SybilInfer: Detecting Sybil nodes using social networks. In *Proceedings of the Network and Distributed System Security Symposium (NDSS 2009), San Diego, CA, USA, 8–11 February 2009*; Internet Socity: Reston, VA, USA, 2009.
23. Sonnino, A.; Danezis, G. SybilQuorum: Open distributed ledgers through trust networks. *arXiv* **2019**, arXiv:1906.12237. [CrossRef]
24. Al-Qurishi, M.; Al-Rakhami, M.; Alamri, A.; Alrubaian, M.; Rahman, S.M.M.; Hossain, M.S. Sybil defense techniques in online social networks: A survey. *IEEE Access* **2017**, *5*, 1200–1219. [CrossRef]
25. Winter, P.; Ensafi, R.; Loesing, K.; Feamster, N. Identifying and characterizing Sybils in the Tor network. In Proceedings of the 25th USENIX Security Symposium (USENIX Security 16), Austin, TX, USA, 10–12 August 2016. Available online: https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_winter.pdf (accessed on 7 May 2025).
26. Bernabe, J.B.; Canovas, J.L.; Hernandez-Ramos, J.L.; Moreno, R.T.; Skarmeta, A. Privacy-preserving solutions for blockchain: Review and challenges. *IEEE Access* **2019**, *7*, 164908–164940. [CrossRef]
27. Hardjono, T.; Pentland, A. Verifiable anonymous identities and access control in permissioned blockchains. *arXiv* **2019**, arXiv:1903.04584. [CrossRef]
28. Nasrulin, B.; Ishmaev, G.; Pouwelse, J. MeritRank: Sybil tolerant reputation for merit-based tokenomics. In Proceedings of the 2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS), Paris, France, 27–30 September 2022.
29. Pass, R.; Shi, E. Hybrid consensus: Efficient consensus in the permissionless model. In *Proceedings of the 31st International Symposium on Distributed Computing (DISC 2017), Vienna, Austris, 16–20 October 2017*; Richa, A.W., Ed.; Leibniz International Proceedings in Informatics (LIPIcs); Schloss Dagstuhl–Leibniz-Zentrum für Informatik: Wadern, Germany, 2017; Volume 91, pp. 39:1–39:16.
30. Duong, T.; Fan, L.; Katz, J.; Thai, P.; Zhou, H.-S. 2-hop blockchain: Combining proof-of-work and proof-of-stake securely. In *Computer Security—ESORICS 2020*; Springer: Cham, Switzerland, 2020.
31. Umar, A.; Jamwal, P.K.; Kumar, D.; Gupta, N.; Gali, V.; Kumar, A. A sybil-resilient and privacy-aware blockchain architecture for dynamic demand response in decentralized microgrids. *Sustain. Energy Technol. Assess.* **2025**, *82*, 104540. [CrossRef]
32. Friedman, A.; Hu, V.C. Presentation 9. Attribute assurance for attribute based access control. In Proceedings of the 2014 IT Professional Conference, Gaithersburg, MD, USA, 22–22 May 2014.
33. Shammar, E.A.; Zahary, A.T.; Al-Shargabi, A.A. An attribute-based access control model for Internet of Things using Hyperledger Fabric blockchain. *Wirel. Commun. Mob. Comput.* **2022**, *2022*, 6926408. [CrossRef]
34. Sun, S.; Du, R.; Chen, S.; Li, W. Blockchain-based IoT access control system: Towards security, lightweight, and cross-domain. *IEEE Access* **2021**, *9*, 36868–36878. [CrossRef]
35. Shih, D.-H.; Wu, T.-W.; Shih, M.-H.; Chen, G.-W.; Yen, D.C. Hyperledger Fabric access control for industrial Internet of Things. *Appl. Sci.* **2022**, *12*, 3125. [CrossRef]

36. Ding, Y.; Wu, Z.; Miao, Y.; Ding, M. SharAcc: Enhancing scalability and security in attribute-based access control with sharding-based blockchain and full decentralization. *Comput. Netw.* **2025**, *257*, 110992. [CrossRef]
37. Javanmardi, S.; Scarpa, M.; Shojafar, M.; Distefano, S.; Merlino, G. Mutable blockchains in IoT-driven sustainable urban planning: Challenges, and analytical modeling. In *Proceedings of the 2025 IEEE International Conference on Smart Computing (SMARTCOMP), Cork, Ireland, 16–19 June 2025*; IEEE: New York, NY, USA, 2025; pp. 408–413.
38. cryotoEthereum Dataset. cryotoEthereum: Preprocessed Ethereum Block Dataset. Available online: https://drive.google.com/file/d/18UffKfUvskS6d36IpPWRWilTRoBhTDjl/view (accessed on 28 February 2026).
39. Google BigQuery. Ethereum Blockchain. Kaggle. Available online: https://www.kaggle.com/datasets/bigquery/ethereum-blockchain (accessed on 28 February 2026).