BLOCKCHAIN BASED STOCK EXCHANGE SYSTEM WITH REPUTATION DRIVEN PROOF-OF-STAKE

PHASE II REPORT

Submitted by

HARIPRIYA S	**22011103016**
KAVYA K	**22011103022**
OBULI DHARSHANTH S	**22011103037**

in partial fulfilment for the award of the degree of

BACHELOR OF TECHNOLOGY
IN
COMPUTER SCIENCE AND ENGINEERING
(CYBERSECURITY)

DEPARTMENT OF COMPUTER SCIENCE
AND ENGINEERING
SCHOOL OF ENGINEERING
SHIV NADAR UNIVERSITY CHENNAI

APRIL 2026

SHIV NADAR UNIVERSITY CHENNAI

BONAFIDE CERTIFICATE

Certified that this report titled "**BLOCKCHAIN BASED STOCK EXCHANGE SYSTEM WITH REPUTATION DRIVEN PROOF-OF-STAKE**" is the bonafide work of **HARIPRIYA S** (Reg. No: **22011103016**), **KAVYA K** (Reg. No: **22011103022**), **OBULI DHARSHANTH S** (Reg. No: **22011103037**), who carried out the work under my supervision. Certified further that to the best of my knowledge the work reported herein does not form part of any other thesis or dissertation on the basis of which a degree or award was conferred on an earlier occasion on this or any other candidate.

SIGNATURE

Dr. T. Nagarajan

Professor

Head of the Department

Department of Computer Science & Engineering

School of Engineering

Shiv Nadar University Chennai,

Kalavakkam -603110.

SIGNATURE

Dr. Rourab Paul

Assistant Professor

Supervisor

Department of Computer Science & Engineering

School of Engineering

Shiv Nadar University Chennai,

Kalavakkam -603110.

ABSTRACT

Traditional stock exchange systems rely on centralized intermediaries such as brokers, clearing houses, and regulatory authorities, which introduce settlement delays, higher transaction costs, and single points of failure. Although blockchain technology offers a decentralized alternative, existing Proof-of-Stake (PoS) mechanisms exhibit a major limitation: validator selection is driven solely by stake, resulting in wealth concentration and reduced emphasis on validator reliability and compliance.

This work proposes a Multi-Dimensional Proof-of-Stake with Reputation (PoS-R) consensus mechanism tailored for stock exchange transactions in the Indian market context. The model replaces linear stake-based selection with a composite weight function that integrates logarithmic stake normalization and a reputation-driven component. The reputation score combines multiple dimensions, including validation accuracy, latency aligned with real-time settlement requirements, integrity through collusion detection, regulatory compliance with SEBI norms, and consistency of validator behaviour over time.

The proposed system demonstrates improved detection of malicious validators and enhanced fairness in validator selection compared to traditional PoS models. It effectively identifies collusion patterns and prioritizes compliant and reliable validators over purely wealth-dominant participants. The inclusion of an integrity-focused metric further strengthens the system by addressing vulnerabilities similar to those observed in past stock exchange irregularities.

Keywords: Blockchain, Proof-of-Stake, Reputation Systems, Smart Contracts, Stock Exchange, Consensus Algorithm, SEBI, Validator Selection, Collusion Detection, Logarithmic Stake Normalisation

ACKNOWLEDGEMENT

We would like to express our sincere gratitude to everyone who provided guidance, support, and encouragement throughout the course of this project.

First and foremost, we are deeply grateful to **Dr. T. Nagarajan**, Professor and Head, Department of Computer Science and Engineering, School of Engineering, Shiv Nadar University Chennai, for his steadfast guidance and unwavering encouragement. His insights into the evolving landscape of distributed systems and secure computing have been instrumental in shaping the direction of this research.

We extend our heartfelt thanks to our project supervisor, **Dr. Rourab Paul**, Assistant Professor, Department of Computer Science and Engineering, School of Engineering, Shiv Nadar University Chennai, for his mentorship and consistent commitment to the success of this work. His expertise in the domains of Blockchain Technology and Cybersecurity provided us with the clarity and technical depth needed to tackle the complex challenges of consensus mechanism design and smart contract implementation.

We also acknowledge with gratitude the support of our families, whose understanding and belief in us have been a constant source of strength and motivation throughout this journey.

Haripriya S	**Kavya K**	**Obuli Dharshanth S**
(22011103016)	**(22011103022)**	**(22011103037)**

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS, ABBREVIATIONS

α	Stake weight parameter (default 0.75)
β	Reputation weight parameter (default 0.25)
R	Composite Reputation Score (0–100)
R_{acc}	Reputation sub-score: Accuracy
R_{comp}	Reputation sub-score: Compliance
R_{cons}	Reputation sub-score: Consistency
R_{intg}	Reputation sub-score: Integrity
R_{lat}	Reputation sub-score: Latency
S	Stake in ETH (Ether)
$W(v)$	Selection Weight of validator v
ABI	Application Binary Interface
AML	Anti-Money Laundering
bps	Basis points (1 bps = 0.01%)
BSE	Bombay Stock Exchange
DAO	Decentralised Autonomous Organisation
DeFi	Decentralised Finance
DLT	Distributed Ledger Technology
ETH	Ether — native cryptocurrency of Ethereum

EVM Ethereum Virtual Machine

Gini Gini coefficient — measure of statistical dispersion

KYC Know Your Customer

NSE National Stock Exchange of India

PoS Proof-of-Stake

PoS-R Proof-of-Stake with Reputation

SEBI Securities and Exchange Board of India

T+0 Same-day trade settlement

T+1 Next-day trade settlement

tx Transaction

ZK Zero-Knowledge (as in ZK-Proof)

CHAPTER 1: INTRODUCTION

1.1 Background

1.1.1 Concept

The financial markets are the heart of today's economy. In the Indian case, the two major exchanges, NSE and BSE, handle around 1.5 billion share trades each year. Behind each and every one of these trades, despite their seemingly instant execution, there lies a sequence of centrally managed entities that include stock brokers, clearing companies, depositories, custodians, and SEBI [1]. Such centralization brings along with itself the problem of settlement lags (currently at T+1, but headed by SEBI to become T+0), as well as transaction fees and vulnerabilities to systemic failure points. It does work, but it's all based on institutional rather than mathematical trust.

The blockchain system provides an entirely new approach to transaction processing. The blockchain system achieves this by storing all transactions in a tamper-proof, decentralized ledger managed by a distributed network of verifiers. This means that there is no requirement for trusted third parties at any point in the process[2]. Transactions can be cleared atomically without the presence of an escrow service provider, and the entire audit trail of every transaction is permanently stored in a completely visible manner within the blockchain system [3].

However, even with such a system, there is a need for a consensus protocol that will establish the identity of the entity allowed to authenticate and record transactions at any point in time. In today's blockchain systems, the common protocol used is called Proof of Stake (PoS), whereby the chances of being selected as a validator depend on the number of

coins locked by the entity [4]. While this protocol solves the issue of inefficiency that was associated with Proof of Work (PoW), it creates an architectural problem for the financial sector in which the richest entity has the highest probability of getting picked [5].

This project provides a solution for the given problem statement. In particular, this paper presents a consensus model called Multi-Dimensional Proof-of-Stake with Reputation (PoS-R), where the selection criterion for validators is defined by their economic contribution and an accumulated reputation score based on five different independent measures. This solution was implemented as a working model of a stock exchange based on the blockchain technology, where all logic for the consensus process is provided by Solidity smart contracts while a user interface was created using React.js technology [3].

Table 1.1 summarises the key distinctions between the current centralised model and the proposed blockchain-based approach.

Table 1.1: Comparison of Centralised vs. Blockchain-Based Exchange Systems

Attribute	**Centralised System**	**Blockchain-Based (PoS-R)**
Settlement speed	T+1 (targeting T+0)	Atomic (same block)
Intermediaries	Brokers, CCPs, depositories	Smart contracts only
Audit trail	Siloed, institution-controlled	Immutable, on-chain
Validator accountability	Regulatory obligation	On-chain reputation score
Fairness of validation	Structural privilege (co-location)	Merit-based selection via W(v)
Compliance enforcement	Manual, post-hoc	Embedded in consensus ($R_{compliance}$)

1.1.2 Motivation

The rationale behind the project arises due to three simultaneous and related trends within the Indian fintech ecosystem. Individually, each trend leads to one clear implication: the existing centralized exchange network is susceptible to flaws that can be effectively addressed

by a decentralized validation mechanism with reputation awareness.

The first and most immediate source of motivation is the event in 2015, where the co-location scandal emerged at the NSE [6]. The information came out that some high-frequency trading companies were given special priority to access the NSE trading server, and thus gain knowledge about trades happening on the order book seconds before anyone else could see them. Such informational inequality caused millions of dollars to be earned systematically on the back of normal traders. More importantly, this particular case is significant for the reason that the fraudulent behavior was done repeatedly as an ongoing practice. Under the current model, a repeating behavior of validating transactions coming from the same counterparty will be recognized through $R_{integrity}$ sub-score immediately.

The second driver comes from the policy stance of SEBI. The 2022 SEBI Discussion Paper [1] on Distributed Ledger Technology explicitly supports research into blockchain settlement, while emphasizing that such systems have to include KYC and AML compliance. Such compliance is regarded as an imperative condition for implementation in India. The $R_{compliance}$ sub-score that we have introduced directly implements the said imperative at the consensus level, as the weight given to each candidate-validator depends on its regulatory compliance.

The third driver has been SEBI's stance on favoring the T+0 settlement system. Fast settlement creates a demanding benchmark of performance, which is that of swift verification. This is why the $R_{latency}$ score has been used to create penalties for validators who are put under pressure of delayed confirmation.

In summary, all of the above factors contribute to creating clear system requirements. The research project makes an important contribution in the area of consensus algorithms research as well as the application of blockchain technology in finance.

1.2 Problem Statement

The main problem addressed in this research is the fact that existing Proof-of-Stake (PoS) systems base their selection process of validators solely on the amount of money staked. While this system works well in blockchain systems, it is not suitable for financial systems for three main reasons.

First, it rewards wealth over reliability. Validators with higher stake are selected more frequently regardless of behavioural performance.

Second, it lacks mechanisms for detecting behavioural fraud. Patterns such as repeated preferential treatment are invisible in stake-only systems.

Third, it cannot enforce regulatory requirements. Compliance attributes such as KYC and AML status are external to the consensus mechanism.

The research question is: *Can a multi-dimensional reputation score, integrated into the validator selection function of a Proof-of-Stake system, address these limitations while preserving security and liveness guarantees?*

1.3 Objectives

The objectives of this project are as follows.

1. To design a composite weight function combining stake and reputation.
2. To model validator behaviour through five independent reputation sub-scores.
3. To implement the system using Solidity smart contracts.

4. To develop a React.js dashboard for monitoring and interaction.

5. To evaluate the system through controlled simulations.

1.4 Scope and Limitations

This project provides an experimental implementation of a stock exchange platform that employs a blockchain infrastructure and performs its consensus operation on-chain. This solution handles the submission, validation, update of reputation, and settlement of transactions.

It is deployed in a private Hardhat blockchain for purposes of experimentation, with testnet deployment reserved for future endeavours.

There have been no audits of this project's implementation, and it currently has no mechanisms for security, privacy, such as zero-knowledge proofs, nor decentralised governance systems.

CHAPTER 2: LITERATURE SURVEY

2.1 Consensus Mechanisms and Proof-of-Stake

The foundational work on distributed consensus in the presence of Byzantine failures was established by Lamport, Shostak, and Pease [7]. Their work demonstrated that a distributed system can achieve agreement even when some participants behave maliciously, provided that a sufficient proportion of nodes remain honest. This principle forms the theoretical basis for modern blockchain consensus mechanisms.

Nakamoto [2] introduced Proof-of-Work (PoW) as a practical consensus mechanism for Bitcoin, where computational effort acts as a deterrent against malicious participation. While effective, PoW is energy-intensive and does not scale well for high-throughput financial systems. King and Nadal [4] later proposed Proof-of-Stake (PoS), where validation rights depend on economic stake rather than computational work. Ethereum's transition to PoS through "The Merge" [8] demonstrated its viability at scale.

Wang et al. [5] identified stake concentration as a key limitation of PoS systems. Their analysis showed that validator selection probability often reflects wealth distribution, leading to centralisation. This motivates the use of logarithmic stake normalisation in our model to reduce dominance by large stakeholders.

Delegated Proof-of-Stake (DPoS), introduced by Larimer [9], improves throughput by allowing token holders to elect a fixed set of validators. However, this approach can result in centralisation within the delegate group. In contrast, our approach incorporates behavioural reputation, ensuring that validator influence depends on performance rather than voting power alone.

The Ouroboros protocol proposed by Kiayias et al. [10] established the theoretical foundation for secure stake-based validator selection. Our model extends this by integrating a reputation component into the selection weight. The architecture of the proposed system is illustrated in Fig. 2.1.

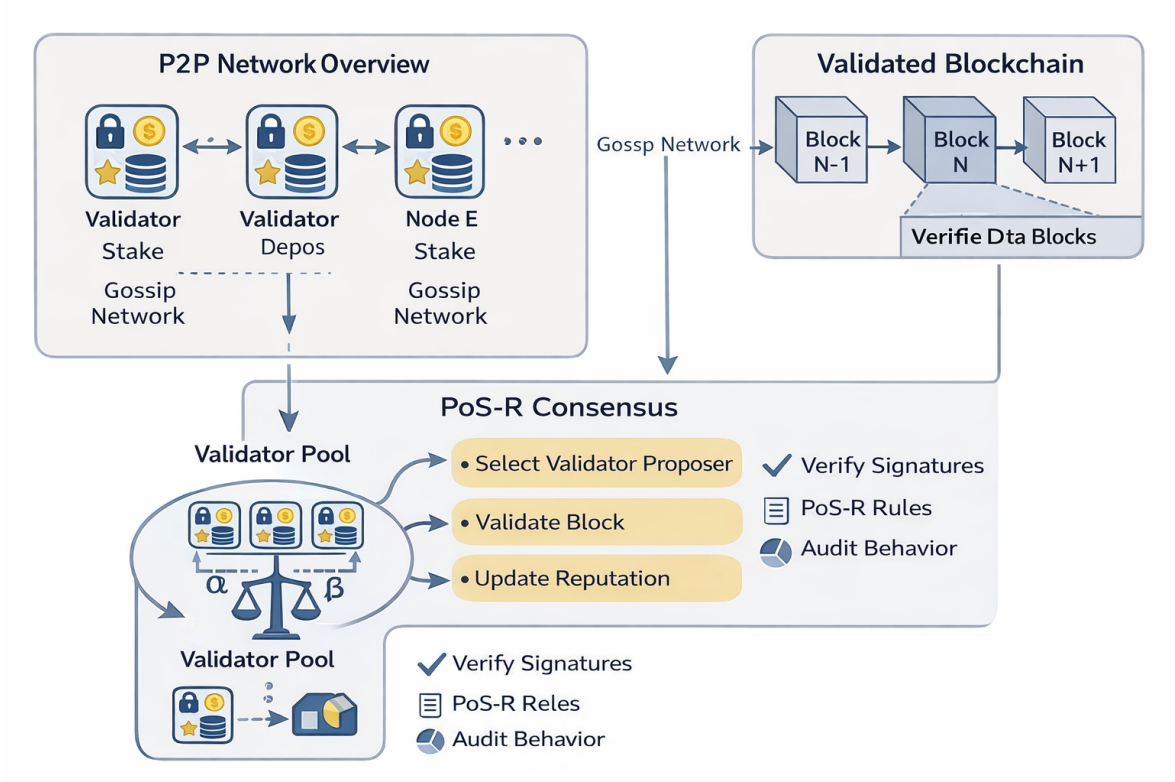

Figure 2.1: Conceptual Architecture of a Reputation-Driven Proof-of-Stake (PoS-R) Blockchain Network

2.2 Reputation Systems in Distributed Networks

2.2.1 Foundational Reputation Theory

Resnick et al. [11] introduced the concept of computational reputation systems, where past interactions are used to predict future behaviour. They identified key properties such as trust persistence, deterrence of dishonest behaviour, and adaptability over time. These principles directly influence our system design.

Mui et al. [12] defined reputation as a probabilistic estimate of cooperative behaviour.

However, a single-score representation does not distinguish between different types of failures. Our multi-dimensional approach addresses this limitation by separating behavioural attributes.

The EigenTrust algorithm proposed by Kamvar et al. [13] introduced transitive trust propagation in peer-to-peer networks. While effective, its computational complexity makes it unsuitable for on-chain implementation. Our model instead focuses on directly measurable behavioural metrics.

2.2.2 Reputation in Blockchain Systems

Li et al. [14] demonstrated that incorporating reputation into blockchain systems significantly reduces malicious behaviour. Their findings show that continuous evaluation discourages dishonest participation.

Zhang et al. [15] proposed a reputation-based consensus mechanism for distributed systems, where trust scores influence selection probability. Their results indicate improved resistance to attacks when reputation is included.

Lev-Ari et al. [16] introduced the FairLedger protocol and applied the Gini coefficient to measure fairness in validator selection. In our work, we adopt this metric while recognising that concentration can be acceptable when it reflects reliable performance.

2.3 Blockchain Applications in Financial Markets

Nakagawa et al. [17] demonstrated a blockchain-based stock settlement system capable of near real-time settlement. However, their implementation relied on a centralised validator, limiting decentralisation.

Pinna and Ruttenberg [18] identified validator accountability as a key challenge in dis-

tributed financial systems. This insight motivates the inclusion of compliance and integrity mechanisms in our model.

In the Indian context, the SEBI Discussion Paper [1] emphasised regulatory compliance, transparency, and auditability. Our model incorporates these requirements directly into the reputation system, ensuring alignment with regulatory expectations. Table 2.1 provides a comparative overview of existing approaches and the corresponding enhancements proposed in this work.

Table 2.1 - Summary of Related Work and Our Extensions

Author / Work	**Contribution**	**Limitation**	**Our Extension**
Resnick et al. (2000)	Reputation systems	Single-score model	Multi-dimensional scoring
Wang et al. (2019)	PoS fairness	Stake dominance	Logarithmic normalisation
Kiayias et al. (2017)	Secure PoS	Stake-only model	Composite weight function
Lev-Ari et al. (2019)	Fairness metric	No compliance focus	Compliance integration
Nakagawa et al. (2018)	Blockchain settlement	Centralised validation	Decentralised PoS-R
SEBI (2022)	Regulatory framework	No consensus model	Compliance-aware design

Table 2.1: Summary of Related Work and Our Extensions

CHAPTER 3: METHODOLOGY

3.1 System Overview

The proposed system consists of two interdependent components: a blockchain layer implemented as Solidity smart contracts, and an application layer implemented as a React.js frontend. The blockchain layer handles all consensus logic - validator registration, weight calculation, validator selection, trade verification, and reputation updates. The application layer provides a real-time interface for interacting with the deployed contracts through MetaMask, a browser-based Ethereum wallet.

The two core smart contracts are ValidatorRegistry.sol, which maintains all validator state (stake, sub-scores, compliance level, selection history), and StockExchange.sol, which handles trade submission, verification routing, and settlement. The exchange contract calls into the registry for validator selection and reputation updates, maintaining a clean separation of concerns between trading logic and consensus logic.

Figure 3.1 illustrates the architecture of the system. A trader interacts with the React dashboard to submit a trade. The StockExchange contract stores the trade as pending and routes the verification to the currently selected validator. The validator - identified by the highest composite weight $W(v)$ at the time of selection - calls verifyTrade(), which triggers a multi-dimensional reputation update in the ValidatorRegistry. All state changes are recorded as on-chain events, which the frontend subscribes to via WebSocket.

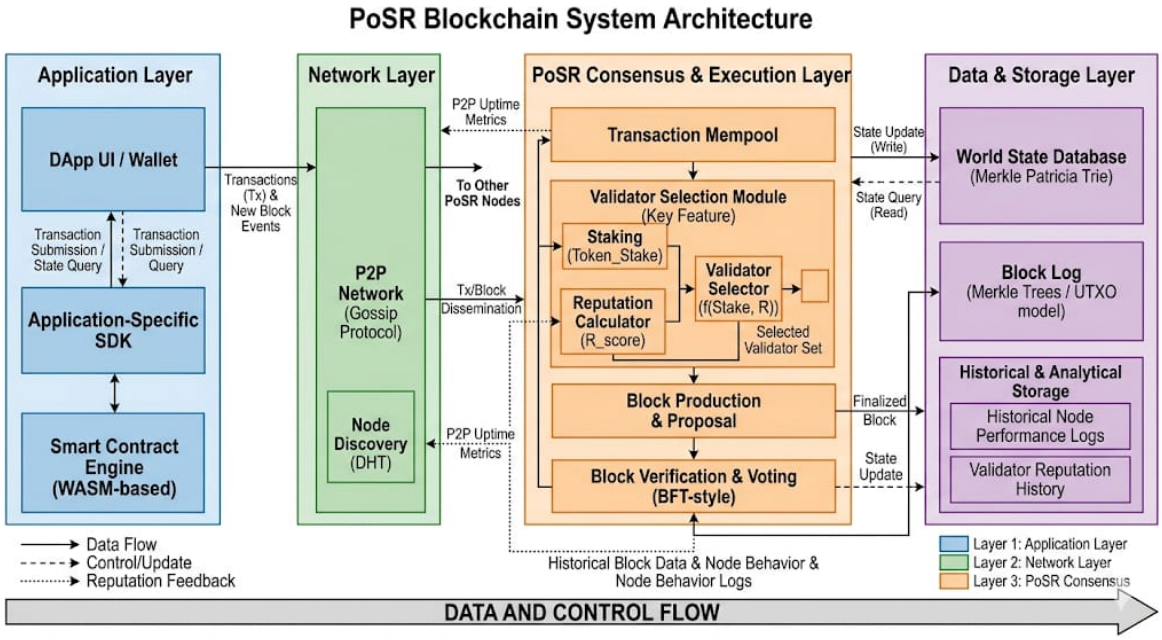

Figure 3.1: System Architecture

3.2 The Multi-Dimensional PoS-R Algorithm

3.2.1 Weight Function

The core innovation of this project is the composite weight function that determines the probability of selection of the validator. The function is as follows:

$$W(v) = \alpha \cdot \ln(S + 1) + \beta \cdot R_{\text{composite}}(v)$$

where S is the validator's stake in ETH, $R_{\text{composite}}$ is the composite reputation score (a value between 0 and 100), α is the stake weight parameter (default 0.75), and β is the reputation weight parameter (default 0.25). The constraint $\alpha + \beta = 1$ is enforced by the contract.

Two design choices in this function require justification. First, the use of $\ln(S + 1)$ rather than S. Linear stake creates a structural problem: a validator with 100 ETH has 10 times the selection weight of a validator with 10 ETH, regardless of their behavioral history. Logarithmic normalisation gives diminishing marginal returns on additional stake

- a validator with 100 ETH scores $\ln(101) \approx 4.62$, while a validator with 10 ETH scores $\ln(11) \approx 2.40$, a ratio of approximately 1.9× rather than 10×. Wealth still matters for Sybil resistance, but it yields progressively less selection advantage.

Second, the choice of $\alpha = 0.75$ and $\beta = 0.25$. These values were derived empirically through controlled experimentation rather than arbitrary assignment.

3.2.2 The Five Reputation Sub-Scores

$R_{\text{composite}}$ is computed as a weighted linear combination of five independently tracked sub-scores, each normalised to the range [0, 100]:

$$R_{\text{composite}} = 0.35R_{acc} + 0.25R_{lat} + 0.25R_{intg} + 0.10R_{comp} + 0.05R_{cons}$$

Table 3.1 - PoS-R Sub-Score Definitions and Weights

Sub-Score	**Weight**	**What it Measures**	**Basis**
$R_{accuracy}$	0.35	Correctness of trade verifications	Reputation theory
$R_{latency}$	0.25	Speed of verification	T+0 requirement
$R_{integrity}$	0.25	Collusion detection	NSE case
$R_{compliance}$	0.10	Regulatory compliance	SEBI
$R_{consistency}$	0.05	Stake stability	DPoS model

The definitions and corresponding weights of the PoS-R sub-scores are summarised in Table 3.1. Each sub-score captures a distinct dimension of validator trustworthiness and is updated based on specific triggers such as verification outcomes, latency, and compliance status.

3.2.3 The α Parameter Experiment

The choice of $\alpha = 0.75$ and $\beta = 0.25$ was validated through controlled experiments comparing validators with different stake and reputation profiles. The experimental results supporting the choice of α and β are presented in Table 3.2.

Table 3.2 - α Parameter Experiment Results

Config	α	β	Bob	Alice	Outcome
Stake-heavy	0.90	0.10	High	Low	Monopoly
Optimal	0.75	0.25	Balanced	Competitive	Best trade-off
Balanced	0.50	0.50	Moderate	High	Reputation dominates
Rep-heavy	0.25	0.75	Low	Very High	Sybil risk

3.2.4 Severity-Based Penalties and Slashing

The system defines multiple severity levels for validator behavior:

Table 3.3 - Severity Level Reputation Effects

Severity	**Impact**	**Meaning**
SUCCESS	+5	Correct verification
LAZY	-3	Slow or missed
NONE	-5	Incorrect
MALICIOUS	-15	Fraud + slashing

The severity levels and their corresponding impact on validator reputation are outlined in Table 3.3. Malicious validators are penalised through slashing, where a portion of their stake is removed and they are excluded from future participation.

3.3 Smart Contract Implementation

3.3.1 ValidatorRegistry.sol

ValidatorRegistry.sol maintains validator state including stake, reputation scores, compliance level, and activity history. It implements weight calculation, validator selection, and reputation updates.

3.3.2 StockExchange.sol

StockExchange.sol handles trade submission, validation routing, and settlement. It interacts with ValidatorRegistry to update reputation and enforce validation logic.

3.4 Frontend Implementation

The frontend is developed using React.js and interacts with the blockchain through the Ethers.js library. This enables the application to communicate with deployed smart contracts and perform operations such as validator registration, trade submission, and verification.

MetaMask is integrated for wallet connection and secure transaction signing. As shown in Figure 3.2, users are required to confirm transactions before they are executed on the blockchain. This ensures that all interactions are authenticated and prevents unauthorized operations.

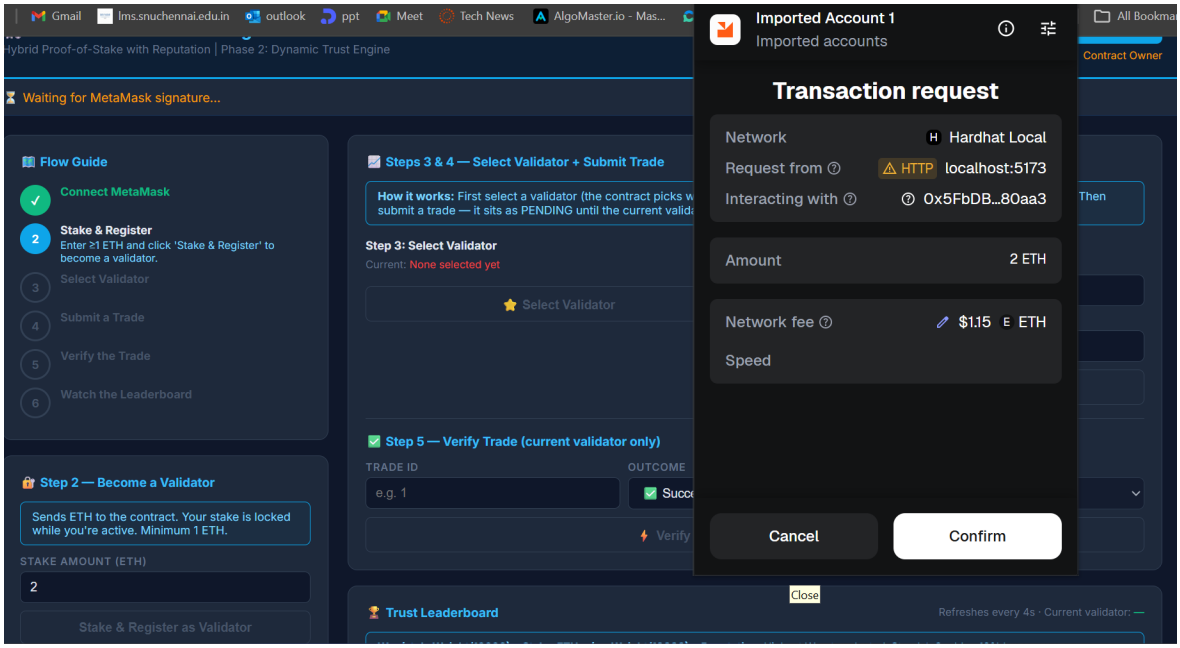

Figure 3.2: MetaMask transaction confirmation during validator registration and trade submission

The dashboard provides multiple interactive components including a validator leaderboard displaying stake, reputation, and weight values, a trade submission interface, and a real-time event log. These features allow users to monitor validator behaviour and track system activity effectively.

Additionally, the frontend reflects on-chain updates dynamically, ensuring that changes such as trade verification and reputation updates are visible instantly. This enhances transparency and improves user experience within the system.

CHAPTER 4: EXPERIMENTS

4.1 Simulation Design

The performance of the proposed Multi-Dimensional PoS-R mechanism is evaluated through simulation and compared with simpler consensus models.

The simulation is implemented entirely in JavaScript and models approximately 10,000 trade verification cycles. Validator characteristics remain fixed throughout the simulation to ensure consistent evaluation. Three different consensus approaches are tested independently to analyze their behaviour under identical conditions.

Validator selection in all models follows a probabilistic approach based on assigned weights. While validators with higher weights have a greater likelihood of being selected, randomness is still preserved in the selection process. This reflects real-world Proof-of-Stake systems, where selection is influenced by weight but not entirely deterministic. The comparative analysis of the proposed system is conducted against three consensus algorithms, described as follows:

The three algorithms compared are:

1. **Plain PoS:** Validator selection is based solely on stake. Reputation is not considered, and validators with higher stake have a proportionally higher chance of being selected. Past behaviour has no influence on the selection process.

2. **Basic PoS-R:** Validator selection depends on both stake and a single reputation score. This introduces a limited notion of trust, but relies on only one aggregated reputation metric.

3. **Multi-Dimensional PoS-R (Proposed):** Validator selection is based on a composite

weight function:

$$W(v) = \alpha \cdot \ln(S + 1) + \beta \cdot R_{\text{composite}}$$

where S represents stake and $R_{\text{composite}}$ is derived from multiple reputation dimensions. This approach balances stake influence with behavioural evaluation, enabling more robust and fair selection.

Table 4.1 - Validator Personality Profiles Used in Simulation

Validator	Stake	Type	Honest Rate	Latency	Behavioral Profile
Alice	3 ETH	HONEST	0.97	1–3	Fast, reliable
Bob	8 ETH	SLOW	0.95	15–35	Rich but slow
Charlie	4 ETH	COLLUDER	0.90	2–5	Same trader bias
Dave	3 ETH	LAZY	0.60	25–60	Frequent delays
Eve	5 ETH	MALICIOUS	0.45	3–8	Fraud approvals
Frank	4 ETH	HONEST	0.96	2–6	KYC validator
Grace	3 ETH	HONEST	0.98	1–4	Reliable
Heera	6 ETH	STAKE DROP	0.93	3–7	Reduces stake mid-run
Iyer	5 ETH	SEBI CLEAN	0.99	2–5	Compliance bonus
Jai	1 ETH	HONEST	0.95	1–2	Small but fast

The Gini coefficient is computed at every 100-trade snapshot as a measure of selection concentration.

4.2 Results and Analysis

4.2.1 Bad Actor Detection

The most striking result of the simulation is the difference in bad actor detection across the three modes.

In Plain PoS, malicious validators are never detected. They continue to participate regardless of fraudulent behavior.

In Basic PoS-R, malicious validators are detected after several trades, but collusion remains undetected.

In Multi-Dimensional PoS-R, malicious actors are detected significantly faster, and collusion is successfully identified a capability absent in other models, as illustrated in Figure 4.1.

Figure 4.1: Simulation Dashboard Showing Performance of Multi-Dimensional PoS-R

4.2.2 Gini Coefficient Analysis

The Gini coefficient results require careful interpretation. Plain PoS shows the lowest Gini value, indicating uniform distribution, but this includes malicious validators.

Multi-Dimensional PoS-R shows a higher Gini value, reflecting merit-based concentration where reliable validators are selected more frequently.

This is desirable in financial systems, where trust and reliability outweigh equal distribution. The detailed performance comparison across all evaluated algorithms is presented in Table 4.2.

Table 4.2 - Full Algorithm Comparison Results (10,000 Trades)

Metric	Plain PoS	Basic PoS-R	Multi PoS-R
Gini Coefficient	0.12	0.18	0.20
Eve detected	Never	Trade 3-22	Trade 9-11
Charlie detected	Never	Never	Trade 86-116
Dave penalised	No	Partial	Yes
Iyer selection	11%	12%	15%
Eve trade share	11%	0.01%	0.02%

4.2.3 Sub-Score Breakdown

Table 4.3 shows the final sub-score values for each validator after 10,000 trades under Multi-Dimensional PoS-R.

Table 4.3 - Final Sub-Score Breakdown

Validator	R_{acc}	R_{lat}	R_{intg}	R_{comp}	R_{cons}	Composite
Alice	100	100	100	50	100	95
Bob	100	0	100	50	100	70
Charlie	95-100	100	0	50	100	68-70
Dave	94-100	0	100	50	100	67-70
Eve	0	50	51	50	52	33
Frank	100	100	100	60	100	96
Grace	100	100	100	50	100	95
Heera	100	100	100	50	0-50	90-95
Iyer	100	100	100	80	100	98
Jai	100	100	100	50	100	95

CHAPTER 5: Conclusion

5.1 Summary of Contributions

This project presents the design, implementation, and evaluation of a Multi-Dimensional Proof-of-Stake with Reputation (PoS-R) consensus mechanism for a blockchain-based stock exchange system.

A key contribution of this work is the decomposition of traditional single-score reputation models into multiple dimensions. Instead of relying on a single reputation value, the proposed system introduces five distinct components:

- $R_{accuracy}$ – evaluates correctness of validation
- $R_{latency}$ – measures speed of transaction verification
- $R_{integrity}$ – detects collusion and conflicts of interest
- $R_{compliance}$ – ensures adherence to regulatory standards (e.g., SEBI)
- $R_{consistency}$ – tracks stable validator behaviour over time

These components collectively form a composite reputation score, enabling a more comprehensive and fair evaluation of validator behaviour. Unlike traditional systems, this approach allows detection of malicious and colluding validators that may otherwise go unnoticed.

The system also incorporates a weighted selection mechanism:

$$W(v) = \alpha \cdot \ln(S + 1) + \beta \cdot R_{\text{composite}}$$

which balances stake and reputation for validator selection.

Experimental evaluation using 10,000 simulated transactions demonstrates that the proposed model improves detection of malicious activities, enhances fairness in validator selection, and reduces vulnerabilities present in existing PoS-based systems.

5.2 Limitations

Despite its contributions, the system has certain limitations.

The current implementation is deployed using a local Hardhat environment. While suitable for testing, this setup does not fully replicate real-world blockchain conditions. As a result, performance metrics such as latency and gas cost may vary when deployed on public networks.

The system has not yet been tested on public testnets or mainnet environments such as Ethereum or Polygon. Therefore, behaviour under real network congestion and fluctuating gas prices remains uncertain.

Additionally, the compliance component currently depends on manual updates. This introduces reliance on a centralized authority, which may reduce decentralization and introduce potential risks. In a real-world deployment, this would require integration with automated oracles for real-time regulatory data.

Finally, validator behaviour in the simulation is predefined. In real-world scenarios, validator strategies may evolve dynamically, which could impact system performance in unforeseen ways.

5.3 Future Work

5.3.1 Immediate Next Steps

Future work includes deployment of the system on public test networks such as Polygon Mumbai or Ethereum Sepolia. This will allow evaluation under realistic conditions and

enable external participation.

Another immediate improvement is the integration of a compliance oracle to automate regulatory updates. This would eliminate manual intervention and improve system reliability.

5.3.2 Medium-Term Extensions

Further enhancements include introducing decentralized governance mechanisms. Validators could participate in voting to dynamically adjust parameters such as α and β.

Privacy-preserving techniques such as Zero-Knowledge Proofs may also be incorporated to protect sensitive transaction data while maintaining verifiability.

Additionally, layer-2 scaling solutions can be explored to improve system throughput and support high-frequency trading environments.

5.3.3 Research Directions

Future research may explore adaptive reputation models using machine learning techniques, where validator scores evolve based on observed behaviour over time.

Formal verification methods can also be applied to smart contracts to identify vulnerabilities and improve system security and reliability.

REFERENCES

[1] Securities and Exchange Board of India. *Discussion Paper on Distributed Ledger Technology*. 2022.

[2] Satoshi Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. URL: https://bitcoin.org/bitcoin.pdf.

[3] Gavin Wood. *Ethereum: A Secure Decentralised Generalised Transaction Ledger*. 2014.

[4] Sunny King and Scott Nadal. *PPCoin: Peer-to-Peer Crypto-Currency with Proof-of-Stake*. 2012.

[5] Wei Wang, Dinh Thai Hoang, and Peizhao Hu. "A Survey on Consensus Mechanisms and Mining Strategy Management in Blockchain Networks". In: *IEEE Access* 7 (2019), pp. 22328–22370.

[6] Central Bureau of Investigation. *NSE Co-location Case Investigation*. 2015.

[7] Leslie Lamport, Robert Shostak, and Marshall Pease. "The Byzantine Generals Problem". In: *ACM Transactions on Programming Languages and Systems* 4.3 (1982), pp. 382–401.

[8] Ethereum Foundation. *The Merge: Ethereum's Transition to Proof-of-Stake*. 2022. URL: https://ethereum.org/en/upgrades/merge/.

[9] Daniel Larimer. *Delegated Proof-of-Stake (DPoS)*. 2014.

[10] Aggelos Kiayias et al. "Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol". In: *Advances in Cryptology – CRYPTO 2017*. 2017.

[11] Paul Resnick et al. "Reputation Systems". In: *Communications of the ACM* 43.12 (2000), pp. 45–48.

[12] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. “A Computational Model of Trust and Reputation”. In: *Proceedings of the 35th Hawaii International Conference on System Sciences*. 2002.

[13] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. “The EigenTrust Algorithm for Reputation Management in P2P Networks”. In: *Proceedings of the 12th International World Wide Web Conference*. 2003.

[14] Zhen Li, Wei Wang, and Gongjun Liu. “Toward Blockchain-Based Fair and Anonymous Ad Dissemination in Vehicular Networks”. In: *IEEE Transactions on Vehicular Technology* (2020).

[15] Rui Zhang, Rui Xue, and Ling Liu. “Security and Privacy on Blockchain”. In: *ACM Computing Surveys* 52.3 (2019).

[16] Keren Lev-Ari et al. “FairLedger: A Fair Blockchain Protocol for Financial Institutions”. In: *OPODIS*. 2019.

[17] Hiroki Nakagawa, Takashi Nakagawa, and Shigeaki Yamamoto. “Applying Blockchain Technology to a Stock Exchange”. In: 2018.

[18] Andrea Pinna and Wiebe Ruttenberg. *Distributed Ledger Technologies in Securities Post-Trading*. 2016.

APPENDIX A: Smart Contract Design

This section outlines the main smart contracts used in the proposed PoS-R system. Each contract operates within the PoS-R framework, governing validator participation, transaction flow, and trust management. The system ensures that validator selection, trade validation, and reputation updates are performed in a decentralized and consistent manner.

A.1 ValidatorRegistry.sol

The `ValidatorRegistry` contract is the core component responsible for managing validators. It maintains and updates validator-related data in a structured and transparent way. The contract ensures that all validators are treated uniformly, and updates are applied only after proper verification.

It maintains the following data:

- Stake (in Ethereum)
- Reputation sub-scores (accuracy, latency, integrity, compliance, consistency)
- Composite reputation score
- Validator selection weight

The weight determines which validators are selected for verification.

The contract implements the following functionalities:

- Weight calculation using:

$$W(v) = \alpha \cdot \ln(S + 1) + \beta \cdot R_{\text{composite}}$$

- Reputation updates based on validation outcomes
- Validator selection logic based on computed weights

The logarithmic component ensures that stake influence grows in a controlled manner, while the reputation component introduces fairness by considering validator behaviour over time.

A.2 StockExchange.sol

The `StockExchange` contract manages the lifecycle of trades within the system. It ensures that each trade follows a structured process from submission to settlement.

Its key responsibilities include:

- Trade submission by users
- Routing trades to selected validators
- Trade verification process
- Settlement of ETH between participants

The contract interacts with the `ValidatorRegistry` to retrieve validator weights and update reputation scores based on validator behaviour.

Together, these smart contracts implement the core functionality of the PoS-R system, enabling decentralized validation, secure transactions, and transparent trust management.

APPENDIX B: Simulation Setup and Parameters

This appendix describes the simulation framework used to evaluate the performance of the proposed PoS-R system under controlled conditions.

B.1 Simulation Environment

The simulation was implemented using JavaScript and executed locally. It models 10,000 trade verification cycles under different consensus mechanisms to ensure consistent and comparable results.

Each transaction follows a predefined flow, while varying conditions are introduced to simulate realistic network behaviour.

B.2 Validator Profiles

Ten validators were defined with fixed behavioural characteristics to represent different types of participants in the network. These include:

- Honest rate
- Latency range
- Behavioural type (honest, slow, colluder, malicious, etc.)

Each validator behaves differently, allowing the system to be tested against various real-world scenarios, including malicious behaviour and collusion attempts.

B.3 Evaluation Metrics

The system performance was evaluated using the following metrics:

- Detection time of malicious validators
- Collusion detection capability
- Gini coefficient for fairness analysis
- Validator selection distribution

These metrics provide insights into the efficiency, fairness, and robustness of the Multi-Dimensional PoS-R model. They help evaluate system behaviour under both normal and adversarial conditions.