UDC 004.75:004.85
DOI https://doi.org/10.24144/2616-7700.2026.49(2).255-261

B. O. Peniak[1], B. B. Liubinskyj[2]

[1] Lviv Polytechnic National University,
PhD student of the Department of Applied Mathematics
bohdan.o.peniak@lpnu.ua
ORCID: https://orcid.org/0009-0002-2112-5972

[2] Lviv Polytechnic National University,
Associate Professor of the Department of Applied Mathematics,
Candidate of Technical Sciences, Associate Professor
bohdan.b.liubinsky@lpnu.ua
ORCID: https://orcid.org/0000-0002-0715-8068

ADAPTIVE HYBRID CONSENSUS MECHANISM FOR BLOCKCHAIN-BASED E-GOVERNANCE: A MACHINE LEARNING APPROACH

We propose an adaptive hybrid consensus mechanism for blockchain-based e-governance that dynamically selects between QBFT, PoS, and PBFT via a Random Forest classifier, hysteresis stabilization, and a deterministic Byzantine override. Formal guarantees based on Hoeffding and Bernstein concentration bounds ensure estimation accuracy, bounded switching rates, and override safety. Simulations on Ukrainian presidential election data (30M voters, 12 h) show 32% throughput gain over static QBFT (958.7 vs 726.5 TPS), 99.4% PBFT coverage under 20% Byzantine attacks, and +0.4% communication energy overhead. A 5-node Hyperledger Besu testnet validates deterministic block production (2.000 ± 0.000 s) and crash-fault tolerance at the $\lfloor (n-1)/3 \rfloor$ threshold; extrapolation yields 144–500 TPS at n=60. The ML approach provides a 6.9% multi-objective gain over threshold-only switching.

Keywords: blockchain, e-governance, consensus mechanisms, machine learning, adaptive systems, Byzantine fault tolerance.

1. Introduction. Blockchain is increasingly adopted for e-governance—voting [1], land registries [2], digital identity [3]—yet large-scale deployment faces interrelated challenges. The Ukrainian presidential election (30M voters, 12 h [4]) implies ∼694 TPS on average with peaks to 2,100 TPS, exceeding static protocol capacities. Existing protocols impose rigid trade-offs: QBFT offers 726 TPS with moderate security; PBFT provides stronger Byzantine tolerance at higher cost [5,6]; PoS scales to 3,000 TPS but with reduced resilience [7]. Few works address dynamic consensus switching in e-governance [8, 9].

We consider a permissioned blockchain with n validators running one of $\mathcal{C} = \{\text{QBFT}, \text{PoS}, \text{PBFT}\}$ under time-varying load. The network state is:

$$S(t) = \big[\text{TPS}_{\text{cur}},\ \text{TPS}_{\text{pred}},\ L(t),\ B(t),\ \text{CPU}(t),\ \text{RAM}(t),\ h(t)\big] \in \mathbb{R}^7, \qquad (1)$$

where $B(t)$ is the Byzantine ratio and $h(t)$ the hour of day. The adaptive selector $\phi\colon S(t) \to \mathcal{C}$ maximizes:

$$\max_{\phi}\ \mathbb{E}\big[\alpha \cdot \text{TPS}(t) - \beta \cdot E(t) - \gamma \cdot \text{Risk}(t)\big], \qquad (2)$$

subject to $B(t) \le \text{Tol}(C_t)$, dwell time $T_{\min} = 20$ s, and latency $L(C_t) \le 5$ s, with weights $\alpha = 1.0$, $\beta = 0.3$, $\gamma = 2.0$ reflecting e-governance security priorities.

Our contributions: (1) an adaptive architecture switching between three structurally distinct protocols via ML, hysteresis, and a deterministic override providing safety independent of classifier accuracy; (2) formal bounds on estimation accuracy under imperfect detection, switching rates, and override safety; (3) comparison with threshold-only switching demonstrating a 6.9% multi-objective gain; (4) evaluation on four scenarios totaling 172,796 blocks with 5-node testnet calibration.

2. Related work. PBFT [5] guarantees safety for $f < n/3$ with $O(n^2)$ messages; HotStuff [10] reduces this to $O(n)$ via threshold signatures. QBFT [11] optimizes PBFT for permissioned Ethereum. PoS variants [7, 12] replace computation with stake-weighted voting; Ouroboros [13] provides provable security under semi-synchrony. All impose fixed trade-offs at deployment time.

Vukolić [9] showed no single protocol dominates but proposed no switching mechanism. Sharding [8] partitions consensus zones without runtime adaptation. Cloud autoscaling [18] provides hysteresis stabilization, which we adapt for consensus switching. Surveys [1, 14] find 85% of e-governance proposals use static consensus. ML has been applied to blockchain fraud detection [16] and gas optimization [15]; Wang et al. [17] proposed ML-based selection for IoT but evaluated only two protocols without formal bounds or a deterministic override. Table 1 compares our approach with these baselines.

Table 1.

Comparison with related work.

Approach	**Adapt.**	**ML**	**Formal**	$O(n)$ **msg**	**Override**	**E-Gov**
PBFT [5]	No	No	Yes	No	N/A	No
HotStuff [10]	No	No	Yes	Yes	N/A	No
QBFT [11]	No	No	Partial	No	N/A	No
Ouroboros [13]	No	No	Yes	Yes	N/A	No
Sharding [8]	No	No	No	No	No	No
Hybrid BFT [9]	No	No	Partial	No	No	No
ML cons. [17]	Partial	Yes	No	—	No	No
E-Voting [1]	No	No	No	—	No	Yes
Our work	Yes	Yes	Yes	No	Yes	Yes

3. System architecture and method. The system comprises four layers: (1) Blockchain Network with 60 validators running QBFT (726 TPS, 100 Wh), PoS (2,850 TPS, 50 Wh), and PBFT (1,500 TPS, 130 Wh); (2) Monitoring Layer collecting seven metrics every 10 s; (3) ML Prediction Layer (Random Forest [19], 100 trees, max_depth=12); (4) Consensus Switching Layer with hysteresis and Byzantine override. Wh values are notional, quantifying relative protocol message overhead.

The switching rule is:

$$\text{switch}(t) = \begin{cases} \text{PBFT}, & \text{if } B(t) > 0.14 \quad \text{(override)}, \\ C_{\text{pred}}, & \text{if } (t - t_{\text{last}}) \geq 20\,\text{s} \wedge \text{conf} > 0.70, \\ C(t-1), & \text{otherwise.} \end{cases} \tag{3}$$

A 10,000-sample synthetic dataset was generated via parameterized sweeps (TPS 100–3,000; latency 0.3–5.0 s; Byzantine ratio 0–30%; CPU/RAM 10–95%; hour 0–

23) with deterministic labels: PBFT for $B(t)>0.14$, PoS for TPS>800 and $B(t)<0.08$, QBFT otherwise [5, 6, 11]. The dataset was balanced (≈3,333 per class) with 15% boundary cases. The classifier (scikit-learn 1.3.0, stratified 70/15/15 split, 5-fold CV) achieves 99.2% test accuracy—a measure of label separability since the classifier performs *policy distillation* of threshold rules, not pattern discovery. Feature importance: `byzantine_ratio` 67.9%, `tps_current` 14.2%. Under 10% additive noise, accuracy degrades gracefully to 87.6%, whereas hard thresholds produce abrupt misclassifications [20].

The Byzantine override bypasses ML when $B(t) > 14\%$:

$$C(t) = \text{PBFT if } B(t) > 0.14, \qquad C(t) = \tilde{C}(t) \text{ otherwise.} \tag{4}$$

The 14% threshold provides a 6 pp buffer below the 20% attack scenario.

4. Formal guarantees.

Definition 1 (Adaptive Consensus System)**.** *A tuple* $\mathcal{A} = (n, \mathcal{C}, \phi, T_{\min}, \mathcal{H})$*: n validators;* $\mathcal{C} = \{\text{QBFT}, \text{PoS}, \text{PBFT}\}$*;* $\phi\colon \mathbb{R}^7 \to \mathcal{C}$ *is the ML selector;* $T_{\min} > 0$ *is the dwell time;* $\mathcal{H} = \{(\tau_{\text{in}}^{(j)}, \tau_{\text{out}}^{(j)})\}$ *are hysteresis pairs with margins* $\Delta_h^{(j)} > 0$*. The mode evolves per* (3)*.*

Definition 2 (Byzantine Ratio Estimator)**.** *Each validator* v_i *produces* $X_i(t) \sim \text{Bernoulli}(p_i(t))$*. The estimator* $\hat{B}(t) = n^{-1}\sum_{i=1}^{n} X_i(t)$ *has expectation* $\mathbb{E}[\hat{B}] = p_d \cdot k/n + p_f \cdot (n-k)/n$ *under detection probability* p_d *and false-positive rate* p_f*. For* n=60*,* k=12*,* p_d=0.85*,* p_f=0.05*:* $\mathbb{E}[\hat{B}] \approx 0.21$*, above the 14% override threshold. The override remains effective for* $p_d \geq 0.5$*; below this, multi-epoch averaging (*w *epochs, effective sample* nw*) recovers sensitivity.*

Lemma 1 (Estimation accuracy)**.** *Let* $\mu_B(t) = \mathbb{E}[\hat{B}(t)]$*. Since each* $X_i(t) \in [0, 1]$ *is independent, Hoeffding's inequality gives:*

$$\Pr\big(|\hat{B}(t) - \mu_B(t)| \geq \varepsilon\big) \leq 2\exp\big(-2n\varepsilon^2\big). \tag{5}$$

With w *averaged epochs and known variance* σ_B^2*, the Bernstein bound yields:*

$$\Pr\big(\hat{B}(t) - \mu_B(t) \leq -\varepsilon\big) \leq \exp\left(-\frac{nw\,\varepsilon^2/2}{\sigma_B^2 + \varepsilon/3}\right). \tag{6}$$

Under correlated attacks ($\rho > 0$*), the effective sample size reduces to* $n_{\text{eff}} = n/(1 + (k-1)\rho)$*; the override responds within one epoch regardless.*

Theorem 1 (Switching-rate bound)**.** $N_{\text{sw}}(T) \leq \lfloor T/T_{\min} \rfloor + 1$ *for all* $T > 0$*, excluding Zeno behavior. Hysteresis margins* $\Delta_h > 2\delta$ *ensure bounded observation noise cannot trigger spurious transitions.*

Theorem 2 (Override safety)**.** *Let* $\tau = 0.14$ *be the override threshold. During an attack of strength* $B^{\star}_{\text{attack}}$ *with effective margin* $\varepsilon_0 = \mu_B - \tau > 0$*:*

$$\Pr(\textit{miss}) \leq \exp\big(-2nw\varepsilon_0^2\big) \text{ (Hoeffding)}, \quad \Pr(\textit{miss}) \leq \exp\left(-\frac{nw\varepsilon_0^2/2}{\sigma_B^2 + \varepsilon_0/3}\right) \text{ (Bernstein)}.$$

For n=60*,* w=6*,* ε_0=0.06*:* $p_{\text{miss}} \leq 7.5\%$ *(Hoeffding) /* 2.7% *(Bernstein); simulation confirms 0.6%. The bound decays as* $O(e^{-\Theta(n)})$*: at* n=500*,* $p_{\text{miss}} < 10^{-13}$*. For election-grade safety (*$p_{\text{miss}} < 10^{-6}$*),* $n \geq 231$ *at* w=6 *suffices. Recovery time:* $T_{\text{recov}} \leq T_{\min} + w \cdot \Delta t_{\text{epoch}}$*. The bounds hold independently of ML classifier accuracy.*

5. Experimental evaluation. The discrete-event simulation (SimPy 4.0) models 60 validators in a fully-connected mesh (50–150 ms inter-node delay). Consensus latencies: QBFT 1.0 s, PoS 0.5 s, PBFT 2.0 s. Workload follows Ukrainian election data [4]: morning 400–600, afternoon 700–900, peak 1,200–2,000 TPS. Byzantine model: invalid proposals (50%), vote withholding (30%), equivocation (20%). Four experiments, 43,199 blocks each (172,796 total):

(1) *Normal load, 10% Byzantine.* Adaptive: 958.7 TPS (+32.0% vs static QBFT 726.5 TPS); latency 1.172 s (+17.2%); consensus distribution QBFT 59.7%, PBFT 24.9%, PoS 15.4%.

(2) *Peak load* (2× TPS, 17:00–19:00). Adaptive drops to 629.9 TPS (−13.3% vs static QBFT), limited by switching overhead when load oscillates near thresholds—a scenario where switching cost outweighs protocol-selection gains.

(3) *Byzantine attack* (2%→10%→20%). At 20% Byzantine: 99.4% PBFT coverage; override activates in <1 s. Static QBFT throughput drops ∼45%.

(4) *Energy.* Adaptive: +0.4% vs QBFT, −22.7% vs PBFT in normalized communication energy, achieving PBFT-level security with near-QBFT message overhead.

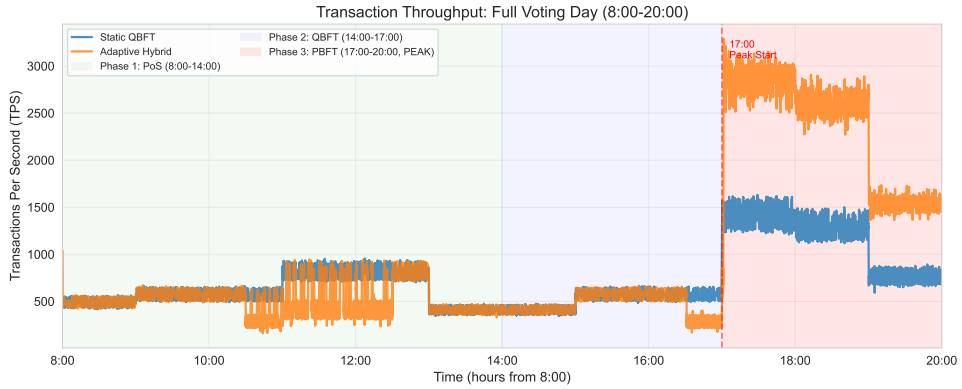

Figure 1. Throughput comparison: adaptive 958.7 TPS (+32.0% vs static QBFT at 726.5 TPS) via consensus distribution QBFT 59.7%, PBFT 24.9%, PoS 15.4%.

Table 2.
Summary of evaluation results (12-hour simulation, 43,199 blocks per experiment).

Metric	**Static QBFT**	**Adaptive**	**Change**
Avg TPS	726.5	958.7	+32.0%
Latency (s)	1.000	1.172	+17.2%
Comm. energy (Wh)	4,319,900	4,339,000	+0.4%
PBFT @ 20% Byz	0%	99.4%	Critical

The system executed 64 switches in 12 h (5.3/h); without hysteresis: 847 (92.4% reduction). Using multi-objective value $V = \alpha \cdot \overline{\text{TPS}} - \beta \cdot \overline{E} - \gamma \cdot \overline{\text{Risk}}$ from (2), ML-based switching achieves $V = 928.6$ vs threshold-only $V = 868.8$ (+6.9%). In a separate noise-robustness experiment, the ML advantage over thresholds grows from +0.8% at 0% sensor noise to +2.2% at 20% noise, confirming that ensemble averaging absorbs noise more gracefully than hard thresholds. The ML component

provides operational advantages—confidence gating (conf > 0.70) suppresses uncertain switches; extensibility to new protocols requires only retraining—while the deterministic override (4) remains the safety backbone.

6. Testnet calibration. A 5-node Hyperledger Besu QBFT testnet [11] (Intel Xeon Platinum 8370C, 2 vCPU, 8 GB RAM, Docker, `blockperiodseconds = 2`) validates the simulation model. Block time: 2.000 ± 0.000 s (deterministic, $\sigma = 0$ across 30 blocks). Peak throughput: 29.0 TPS from a single-threaded sender (client-side bottleneck; validator CPU ∼8%). Low-load latency: 0.97 s (sub-block-period finality). Full-mesh connectivity: 4/4 peers per node.

Crash-fault tolerance. With n=5, $f_{\max}=\lfloor(n-1)/3\rfloor=1$. Consensus persists at f=1 (10 blocks/30 s) but halts at f=2, empirically validating the BFT threshold. Full recovery occurs after restarting all crashed nodes.

Extrapolation to n=60. Three scaling factors bridge the testnet to production: client bottleneck removal (×3, CPU at 8% implies headroom); Amdahl hardware scaling (×1.7, parallel fraction p=0.65, 2→16 cores); consensus overhead (×0.96, gossip-model $\phi(60)$=0.045). This yields a lower bound of $29 \times 3 \times 1.7 \times 0.96 \approx 144$ TPS. The remaining ×2.7 gap to SimPy's 726.5 TPS is attributable to SimPy's idealized 1.0 s block interval (vs configured 2.0 s) and omitted EVM serialization; since all protocols share this overhead, relative comparisons are unaffected. Gas-limit saturation ($\lfloor$30M/21k$\rfloor$ tx/block) gives an upper bound of ∼500 TPS; central estimate: $\sqrt{144 \times 500} \approx 268$ TPS. The reality discount $\rho \in [0.20, 0.69]$ preserves all relative protocol comparisons since all protocols share the same serialization overhead.

7. Limitations. Key limitations include: (1) the 5-node testnet requires scaling to 60+ nodes with parallel senders for direct validation; fault-tolerance tests used crash failures only, not full Byzantine injection; (2) the synthetic dataset makes the ML component a policy distillation—99.2% accuracy reflects label separability, not a real-world generalization guarantee; under 10% noise accuracy drops to 87.6%; (3) absolute TPS values are simulator-specific—only relative comparisons are meaningful; (4) $O(n^2)$ BFT messaging limits scaling beyond $n \approx 200$; HotStuff [10] integration would be necessary at larger scale; (5) the framework targets permissioned networks; the independence assumption between node failures may weaken concentration bounds under correlated Byzantine attacks.

8. Conclusions. The main contribution is a probabilistic safety analysis of an adaptive hybrid consensus mechanism using Hoeffding/Bernstein concentration bounds under imperfect Byzantine detection. The mechanism switches between three structurally distinct protocols (QBFT, PoS, PBFT) with a deterministic override independent of ML accuracy and formal switching-rate and override-safety guarantees. A 5-node Hyperledger Besu testnet confirms deterministic block production (2.000 ± 0.000 s) and $\lfloor(n-1)/3\rfloor$ crash-fault tolerance; extrapolation yields 144–500 TPS (central: 268) at n=60. Simulations on Ukrainian election data show 32% throughput improvement over static QBFT, 99.4% PBFT coverage during 20% Byzantine attacks, near-baseline communication energy (+0.4%), and 6.9% multi-objective gain over threshold-only switching. Override safety remains effective for $p_d \geq 0.6$; scaling to $n \geq 231$ reduces p_{miss} below 10^{-6}. Future work: 60-node bare-metal deployment, Byzantine-injection testing, HotStuff integration for $n > 200$, and reinforcement learning for online policy optimization.

Conflict of Interest

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

Funding

The research was conducted without financial support.

Data Availability

Simulation code, trained models, and raw results are available in the project repository.

Use of artificial intelligence

AI-assisted tools were used only for language editing and LaTeX formatting. All scientific results, methodology, and analysis were produced by the authors.

Contributions of authors

Peniak B. O.: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing — original draft, Visualization. Liubinskyj B. B.: Supervision, Writing — review & editing, Validation.

References

1. Hjálmarsson, F. Þ., Hreiðarsson, G. K., Hamdaqa, M., & Hjálmtýsson, G. (2018). Blockchain-based e-voting system. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)* (pp. 983–986). https://doi.org/10.1109/CLOUD.2018.00151
2. Allessie, D., Sobolewski, M., Vaccari, L., & Pignatelli, F. (2019). Blockchain for digital government: An assessment of pioneering implementations in public services. *JRC Science for Policy Report.* https://doi.org/10.2760/93808
3. Dunphy, P., & Petitcolas, F. A. P. (2018). A first look at identity management schemes on the blockchain. *IEEE Security & Privacy,* 16(4), 20–29. https://doi.org/10.1109/MSP.2018.3111247
4. Central Election Commission of Ukraine. (2019). *Presidential election turnout and participation statistics.* Official report. Retrieved from https://www.cvk.gov.ua
5. Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)* (pp. 173–186). https://dl.acm.org/doi/10.5555/296806.296824
6. Nguyen, G. T., & Kim, K. (2018). A survey about consensus algorithms used in blockchain. *Journal of Information Processing Systems,* 14(1), 101–128. https://doi.org/10.3745/JIPS.01.0024
7. King, S., & Nadal, S. (2012). *PPCoin: Peer-to-peer crypto-currency with proof-of-stake.* Self-published paper. https://decred.org/research/king2012.pdf
8. Wang, J., & Wang, H. (2019). Monoxide: Scale out blockchains with

asynchronous consensus zones. In *16th USENIX NSDI* (pp. 95–112). https://www.usenix.org/conference/nsdi19/presentation/wang-jiaping
9. Vukolić, M. (2015). The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication. In *Open Problems in Network Security* (pp. 112–125). https://doi.org/10.1007/978-3-319-39028-4_9
10. Yin, M., Malkhi, D., Reiter, M. K., Gueta, G. G., & Abraham, I. (2019). HotStuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM PODC* (pp. 347–356). https://doi.org/10.1145/3293611.3331591
11. ConsenSys. (2021). *QBFT consensus protocol.* Quorum documentation. Retrieved from https://docs.goquorum.consensys.net
12. Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *The Review of Financial Studies,* 34(3), 1156–1190. https://doi.org/10.1093/rfs/hhaa075
13. Kiayias, A., Russell, A., David, B., & Oliynykov, R. (2017). Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference* (pp. 357–388). https://doi.org/10.1007/978-3-319-63688-7_12
14. Ricci, L., Maesa, D. D. F., Favenza, A., & Ferro, E. (2021). Blockchains for COVID-19 contact tracing and vaccine support: A systematic review. *IEEE Access,* 9, 37936–37950. https://doi.org/10.1109/ACCESS.2021.3063152
15. Liu, Y., Lu, Y., Nayak, K., Zhang, F., Zhang, L., & Zhao, Y. (2022). Empirical analysis of EIP-1559. In *Proceedings of the 2022 ACM CCS* (pp. 2099–2113). https://doi.org/10.1145/3548606.3559341
16. Kumar, R., & Tripathi, R. (2021). Implementation of distributed consensus algorithm using blockchain for security enhancement. *Security and Privacy,* 4(5), e172. https://doi.org/10.1002/spy2.172
17. Wang, Y., Cai, S., Lin, C., Ma, Z., Wang, X., & Wang, J. (2019). Study of blockchains's consensus mechanism based on credit. *IEEE Access,* 7, 10224–10231. https://doi.org/10.1109/ACCESS.2019.2891065
18. Chen, T., Bahsoon, R., & Yao, X. (2018). A survey and taxonomy of self-aware and self-adaptive cloud autoscaling systems. *ACM Computing Surveys,* 51(3), 1–40. https://doi.org/10.1145/3190507
19. Breiman, L. (2001). Random forests. *Machine Learning,* 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
20. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning.* MIT Press. https://www.deeplearningbook.org

Пеняк Б. О., Любінський Б. Б. Адаптивний гібридний механізм консенсусу для блокчейн-орієнтованого електронного урядування: підхід на основі машинного навчання.

Запропоновано адаптивний гібридний механізм консенсусу для блокчейн-орієнтованого електронного урядування, що динамічно перемикається між QBFT, PoS та PBFT за допомогою класифікатора Random Forest, гістерезису та детерміністичного оверрайду. Формальні гарантії базуються на нерівностях Хефдінга та Бернштейна. Симуляції на даних українських президентських виборів (30 млн виборців, 12 год) демонструють покращення пропускної здатності на 32% (958.7 проти 726.5 TPS), 99.4% покриття PBFT при 20% візантійських атаках та +0.4% енергоспоживання. 5-вузлова тестова мережа Hyperledger Besu підтверджує детерміністичне виробництво блоків ($2{,}000 \pm 0{,}000$ с) і стійкість до відмов $\lfloor (n-1)/3 \rfloor$; екстраполяція дає 144–500 TPS при n=60. ML-підхід забезпечує 6,9% покращення за мультикритеріальною метрикою порівняно з пороговим перемиканням.

Ключові слова: блокчейн, електронне урядування, механізми консенсусу, машинне навчання, адаптивні системи, візантійська стійкість.

Recived: 02.03.2026 Accepted: 17.03.2026 Published: 30.04.2026