Evaluating Blockchain Efficiency for Large-Scale Geospatial Dataset Distribution

Wera Dawa[1*] and Arun Kumar Singh[2]

[2] School of MCS, PNG University of Technology, Lae, Papua New Guinea, E-mail: 25984614weda@pgs.pnguot.ac.pg, aromoweradawa@gmail.com

[2] School of MCS, PNG University of Technology, Lae, Papua New Guinea, ORCID ID: 0000-0002-9974-4696, E-mail: arun.singh@pnguot.ac.pg, arunsinghiiita@gmail.com

***Corresponding author**

Received: 12 August 2025 | Revised: 25 November | Accepted: 26 November 2025

Abstract

The secure and reliable distribution of geospatial datasets is an emerging challenge for a country like Papua New Guinea (PNG), where responsible agencies increasingly depend on satellite imagery, cadastral surveys, and environmental monitoring data for the national development growth. Existing government systems are centralized systems which often encounter bottlenecks, risk of tampering, and outages due to single points of failure, which undermine data trust and availability. This study develops and evaluates a blockchain-driven simulation framework that is customized to the PNG's geospatial metadata distribution. We used synthetic datasets of varying sizes, model tests transaction creation, metadata logging, consensus validation, and block propagation under different consensus mechanisms. The performances were measured across throughput, latency, energy demand, and resilience against tampering. Results obtained from this simulation highlighted clear trade-offs between scalability and security across Proof of Work, Proof of Stake, Proof of Authority, and PBFT. The findings provide practical benchmarks for agencies such as the National Mapping Bureau, Department of Lands, and climate monitoring bodies in PNG, where transparency, low energy consumption, and efficiency are critical. The findings highlight how blockchain can be applied to build transparent and resilient data-sharing systems for national spatial data infrastructures operating under resource constraints.

Keywords: *Blockchain, Consensus Mechanisms, Data Integrity, Geospatial Data Sharing, Performance Evaluation, Scalability.*

1. Introduction

In Papua New Guinea (PNG), the demand for reliable geospatial data has increased significantly with the expansion of satellite coverage, climate monitoring programs, and land resource mapping initiatives. Geospatial metadata refers to information describing dataset origin, ownership, and accuracy that is critical for decision-making in land administration, forestry, and urban planning. For instance, cadastral boundaries used by provincial governments or floodplain maps for disaster response must remain accurate, transparent, and tamper-resistant. However, most metadata management systems in PNG are centralized, often dependent on single servers located in Port Moresby or Lae. Such systems are vulnerable to tampering, unauthorized access, and system outages, which undermine trust in geospatial data infrastructures (Zhang & Huang, 2024). Blockchain technology has emerged as a promising solution for addressing these vulnerabilities by offering decentralized, immutable, and transparent data management. Similar approaches have been applied in a blockchain-based spatial index verification method for remote sensing images using Hyperledger Fabric to improve retrieval efficiency and security in land management (Liu & Chang, 2024) and Spatial Planning Data Structure Based on Blockchain Technology (Tang et al., 2024). Embedding geospatial metadata within a blockchain ledger allows datasets to remain transparent, verifiable, and resistant to manipulation across their entire lifecycle (Liu & Chang, 2024). However, the performance of any blockchain platform is shaped by the consensus mechanism it relies on. Consensus protocols dictate how nodes agree on valid transactions and blocks, which in turn affects throughput, latency, scalability, and energy demand. Proof of Work (PoW), as the earliest protocol, offers strong security but comes at the cost of very high computational effort and energy consumption, making it poorly suited to low-resource contexts like PNG (Li, 2024). Alternatively, Proof of Stake (PoS) and Proof of Authority (PoA) have demonstrated promising scalability with lower energy use, while Practical Byzantine Fault Tolerance (PBFT) provides resilience in systems where low latency is demanded (Li, 2024; Zimba et al., 2025). Despite these advances, few studies have directly evaluated and compared how these consensus mechanisms perform in geospatial data-sharing environments. Much of the existing work remains conceptual, focusing on architectural proposals rather than systematic performance testing (Zhao et al., 2022). This study addresses that gap by benchmarking PoW, PoS, PoA, and PBFT under simulated conditions tailored to PNG's resource constraints. Objectives of this study were (i) to build and test a modular blockchain simulation framework for geospatial metadata using multi-scale synthetic datasets; (ii) to compare PoW, PoS, PoA, and PBFT in terms of throughput, latency, energy use, ledger growth, and tamper detection; and (iii) To identify the most efficient and sustainable consensus model for national geospatial metadata systems, with emphasis on Papua New Guinea.

2. Materials and Methods

The performance evaluation of different blockchain consensus mechanism carried out was looked in this section of the article. In order to obtain a strong evaluation performance, we designed a highly modular simulation framework. Following this section is the detailed architecture of the framework, the design of the synthetic geospatial datasets, plus the experimental customized configuration for executing simulations, and the steps taken to ensure repeatability and transparency.

2.1 Framework Architecture

The simulation framework for this study was built in Python and run inside a Jupyter Notebook using the Anaconda environment. To keep the work organized and easy to test, the whole system was broken into four major parts: creating transactions, generating blocks, running the consensus algorithms, and recording the results. This structure was to reflect how a real spatial data infrastructure (SDI) might integrate a blockchain layer behind the scenes (Figure 1).

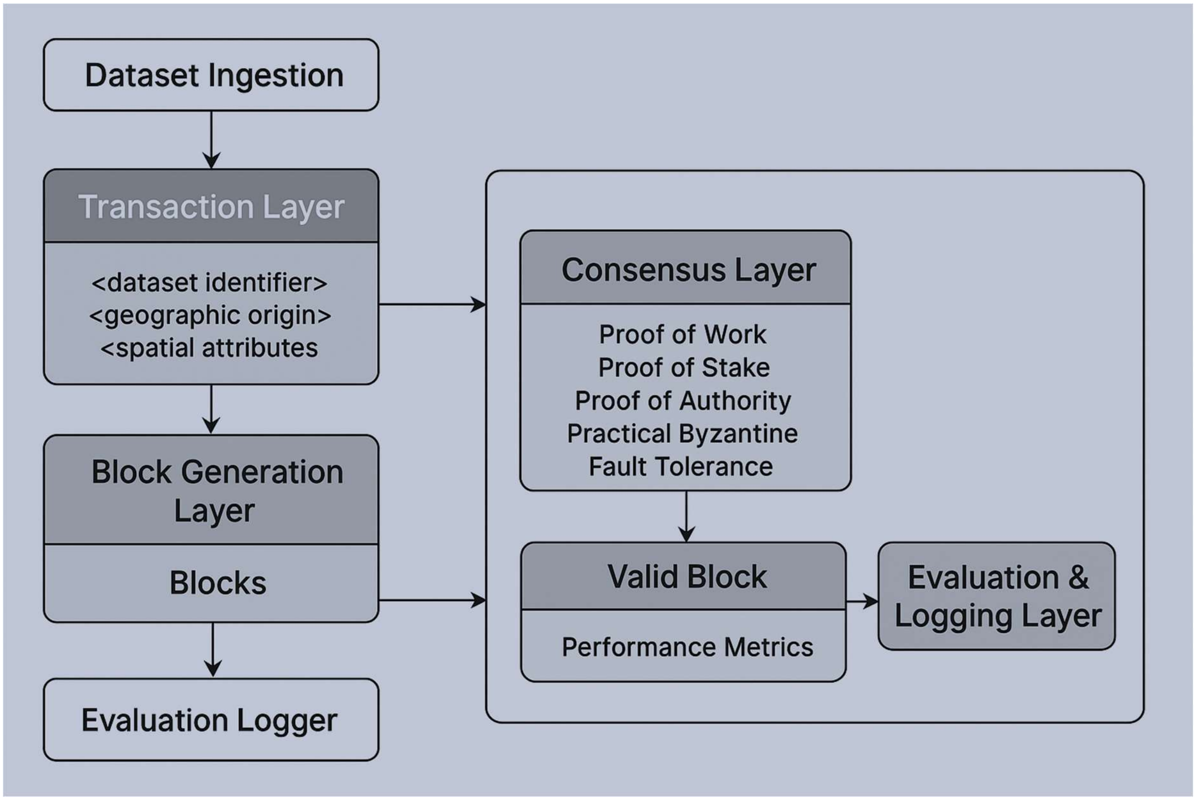

Fig. 1 The illustration of the Architectural flowchart of the system being designed

First, we start by creating the transactions, where geospatial metadata are turned into structured records, similar to what you would normally see in the national SDI/GIS platforms things like dataset IDs, coordinates (latitude and longitude), timestamps, and location identifiers such as provinces or regions. These entries reflect the kinds of data produced from cadastral surveys, satellite images, and environmental monitoring activities by the responsible bodies. We then passed each of the records through a SHA-256 hashing process, which produces a unique digital fingerprint. This fingerprint acts as the final transaction that gets stored on the ledger, ensuring the data cannot be changed once it is added.

The second part of the framework is where we assembled the blocks. After the system checks the incoming transections and once it is valid, it groups them into a block. Each block includes the list of hashed transactions, the time the block was created, and the hash of the previous block to keep the chain intact. Some consensus algorithms also require a nonce value, which the system adjusts until the block meets the necessary validation rules. One important design choice was to keep the block creation time fixed at 1.5 seconds for every algorithm. This way, any difference in performance comes from the consensus process itself, not from how fast the blocks were created.

Thirdly, the consensus algorithms are tested. Four different mechanisms were implemented Proof of Work (PoW), Proof of Stake (PoS), Proof of Authority (PoA), and Practical Byzantine Fault Tolerance (PBFT). In the PoW setup, miners repeatedly hash data to mimic solving a cryptographic puzzle, and energy use is estimated from how many hashing attempts are made. The PoS version doesn't rely on heavy computation but instead uses a probabilistic method to select validators, so its energy cost is tied to these selection cycles. To significantly reduce computational demands, PoA was assumes a set of trusted and permissioned validators such as government agencies. PBFT works differently again, relying on several rounds of communication between a fixed number of nodes, with its performance shaped by how many message exchanges were required. This combination of algorithms reflects recent work in the literature comparing the energy efficiency of modern blockchain consensus models.

The final part of the framework is the evaluation and logging system. We captured a range of performance indicators such as how many transactions were processed, how long it takes to validate a block, how much energy is consumed, how quickly the ledger grows, and whether any attempts to tamper with the chain are detected during each simulation run. All results were printed out in both CSV and JSON formats which were checked, visualized, or reused later. To ensure the experiment is fully traceable, every record also includes timestamps and the version of the simulation code used at the time.

2.2 Database and Dataset Design

To create a geospatial environment that felt as close to reality as possible, the simulation used three synthetic datasets that mimic the kinds of metadata handled at different government levels. The smallest dataset, roughly 50 MB, represents information you would usually find at the municipal level for example, zoning boundaries or basic cadastral maps. The medium dataset, at around 200 MB, reflects the types of records managed at the provincial scale, such as satellite-derived vegetation indices or forestry monitoring data. The largest dataset, reaching about 1,000 MB, simulates the heavy workloads of national systems, including climate datasets or full-country cadastral layers.

All the metadata used in these datasets was generated using NumPy’s pseudorandom functions, and the random seed was deliberately set to 2024. By fixing the seed, the system produces the same coordinates, timestamps, and region assignments every time the simulation is run. This consistency ensures that the datasets themselves don’t introduce unexpected variations into the results. Any performance differences that appear across the experiments can therefore be traced back to the consensus algorithms being tested, not to changes in the data. This approach follows common reproducibility standards in blockchain research.

In order to imitate how geospatial data would actually be fed into a blockchain customized for PNG practice, the transactions from these datasets were organized into blocks with three different densities. In the low-ingestion scenario, each block held only two transactions, reflecting occasional updates such as periodic survey results. The medium rate used five transactions per block, and the high rate used ten, capturing situations where data is generated more rapidly for instance, from sensors, satellite feeds, or continuous environmental monitoring systems.

2.3 Experimental Configuration

To test how different blockchain settings respond to various geospatial data conditions, we set up an experimental grid that systematically combined all the variables in the study (Figure 1). We looked at four consensus mechanisms PoW, PoS, PoA, and PBFT three dataset sizes (50 MB, 200 MB, and 1,000 MB), and three transaction densities (2, 5, and 10 transactions per block). Each setup was run twice to reduce the effect of randomness and improve the reliability of the findings. Altogether, this produced 72 separate simulation runs (Figure 2).

```
# -------------------------
# Experiment orchestration
# -------------------------
EXPERIMENT_CONFIG = {
    "consensus_list": ["PoW","PoA","PoS","PBFT"],
    "dataset_sizes_MB": [50, 200, 1000],   # three size points: small, medium, large
    "txs_per_block_list": [2, 5, 10],      # transactions per block scenarios
    "blocks_per_run": 6,
    "runs": 2,
    "seed_start": 100
}
OUTDIR = "./consensus_evaluation_outputs"
os.makedirs(OUTDIR, exist_ok=True)

all_results: List[EvalResult] = []
summary_rows = []

exp_counter = 0
total_runs = len(EXPERIMENT_CONFIG["consensus_list"]) * len(EXPERIMENT_CONFIG["dataset_sizes_MB"]) * len(EXPERIMENT_CONFIG["txs_per_block_list"]) * EXPER
print(f"Starting experiments: {total_runs} runs (each with {EXPERIMENT_CONFIG['blocks_per_run']} blocks)")

for consensus in EXPERIMENT_CONFIG["consensus_list"]:
    for size in EXPERIMENT_CONFIG["dataset_sizes_MB"]:
        for txs_per_block in EXPERIMENT_CONFIG["txs_per_block_list"]:
            for run_id in range(EXPERIMENT_CONFIG["runs"]):
                exp_counter += 1
                seed = EXPERIMENT_CONFIG["seed_start"] + exp_counter
                results, chain = evaluate_consensus(
                    consensus_name=consensus,
                    num_blocks=EXPERIMENT_CONFIG["blocks_per_run"],
                    txs_per_block=txs_per_block,
                    dataset_size_MB=size,
```

Fig.2 A snippet of experimental orchestration

All experiments were conducted under a fully controlled environment. The simulation was run in a Jupyter Notebook on a Windows 11 machine, and the computational setup never changed from one run to the next. Key parameters such as the hashing algorithm, the block design, and the fixed block creation time of 1.5 seconds were kept constant across every test. Keeping these elements, stable ensures that any differences we observed in performance were genuinely due to the consensus algorithms, not external factors or hardware variations. For each simulation, we recorded several performance metrics to understand how well each consensus model handled the geospatial data:

1. **Throughput** – We measured this as the average number of transactions processed per second. Focusing on the average value gives a clearer picture of sustained performance over time, rather than temporary bursts or dips.
2. **Block validation time** – This shows how long it took on average to confirm each block. We also examined how this timing changed when block size and transaction density increased, which helped us see how well each algorithm scales.
3. **Estimated energy consumption** – The energy models were tailored to match the characteristics of each consensus approach. For PoW, energy use was linked to hashing attempts; for PoS, it reflected validator selection cycles; PoA assumed a stable set of authority nodes; and PBFT's cost came from repeated message exchanges between nodes. These estimates are based on established modelling techniques and recent studies on energy-efficient blockchain protocols.

4. **Ledger storage growth** – After each run, we measured how much the blockchain grew, in megabytes. This gives a sense of how much storage a real-world deployment would require as more geospatial metadata is added over time.
5. **Tamper-detection performance** – To test integrity, we slipped a fabricated transaction into the first block and checked whether the system could detect it during the final verification stage. We recorded whether the tampering was caught, how long detection took, and whether this extra step affected the system's overall performance.

2.4 Reproducibility and Transparency

Reproducibility was a key focus of this study because it is essential in both blockchain and geospatial research. To keep the experiments consistent, all datasets were generated using a fixed random seed. This ensured that the metadata values were identical every time the simulations were run, removing one of the biggest sources of variation. With the data held constant, any changes in performance could be confidently attributed to the consensus mechanisms rather than random fluctuations. The simulation code itself was organized into clear, well-documented modules handling transaction creation, block assembly, consensus execution, and metric logging to make it easy for other researchers to review, replicate, or build on the work.

We also applied uniform timing conditions across all tests, including a fixed 1.5-second block creation window, and used the same cryptographic hash function (SHA-256) throughout. Keeping these factors, constant helps isolate the behavior of each consensus mechanism without interference from shifting system parameters. All outputs throughput, latency, estimated energy use, and ledger size were exported in both CSV and JSON formats and saved together with code-version details. We also logged the full Python runtime environment, including library versions and timestamps, so that the entire study can be replicated in the future or compared with other implementations.

To ensure the work remained grounded in current research, we drew on recent studies that address blockchain in geospatial applications and the growing emphasis on energy-efficient consensus models. (Zhao et al., 2022) highlight the rising need for decentralized and trustworthy systems for sharing geospatial metadata. (Tang et al., 2024) show how Hyperledger-based blockchain frameworks can support spatial planning by improving the integrity and traceability of spatial data. In the same vein, studies such as (Zimba et al., 2025) draw attention to the significant energy advantages of lighter protocols like Proof of Authority compared with more resource-intensive systems like Proof of Work. Incorporating these insights helped shape an experimental design that not only compares consensus algorithms but also aligns with emerging needs in geospatial data management.

3. Results and Discussion

This section presents the results from the seventy-two simulation runs carried out across the full combination of consensus mechanisms (PoW, PoS, PoA, and PBFT), dataset sizes (50 MB, 200 MB, and 1,000 MB), and transaction densities (2, 5, and 10 transactions per block). The analysis examines five core performance indicators: throughput, block validation time, estimated energy use, ledger growth, and tamper-detection behavior. The results are interpreted alongside established findings in blockchain research to help explain the computational, scalability, and security patterns that emerged from the experiments.

3.1 Throughput Performance

The throughput results made it clear that the consensus mechanism itself plays the most significant role in determining system performance. Proof of Authority recorded the highest throughput in almost every scenario, followed by PBFT, PoS, and finally PoW. This ordering mirrors explanations in recent literature, which note that PoA performs exceptionally well in permissioned blockchain environments because it eliminates mining competition and relies instead on a small, trusted set of validators. PoW, on the other hand, consistently produced the lowest throughput. Its mining process is computationally demanding, and even with a fixed block interval of 1.5 seconds, it slows down block confirmation.

Throughput dropped as dataset size and transaction density increased, although the degree of this decline differed across the consensus algorithms. PBFT showed a moderate sensitivity to scaling, largely because it relies on several rounds of communication among nodes. As blocks grow in size or contain more transactions, these message-passing rounds add overhead a pattern also reported in other studies of PBFT-based permissioned networks. PoS maintained better throughput than PoW as the workload increased, reflecting its lighter validator-selection mechanism, which avoids the heavy hashing race that defines PoW.

Transaction density produced its own clear trends. With only two transactions per block, throughput remained relatively high across all mechanisms except PoW. Lower-density blocks place less strain on the system, allowing most protocols to perform efficiently. As transaction density increased to five and then ten transactions per block, PoW throughput dropped sharply, showing again how the mining requirement becomes a bottleneck under heavier loads. PoA, by contrast, remained largely unaffected even at the highest densities. Its stability under heavier block conditions shows a strong capacity for handling continuous streams of geospatial metadata, which aligns with broader research describing PoA as one of the most efficient consensus mechanisms in trusted-validator settings (Figure 3).

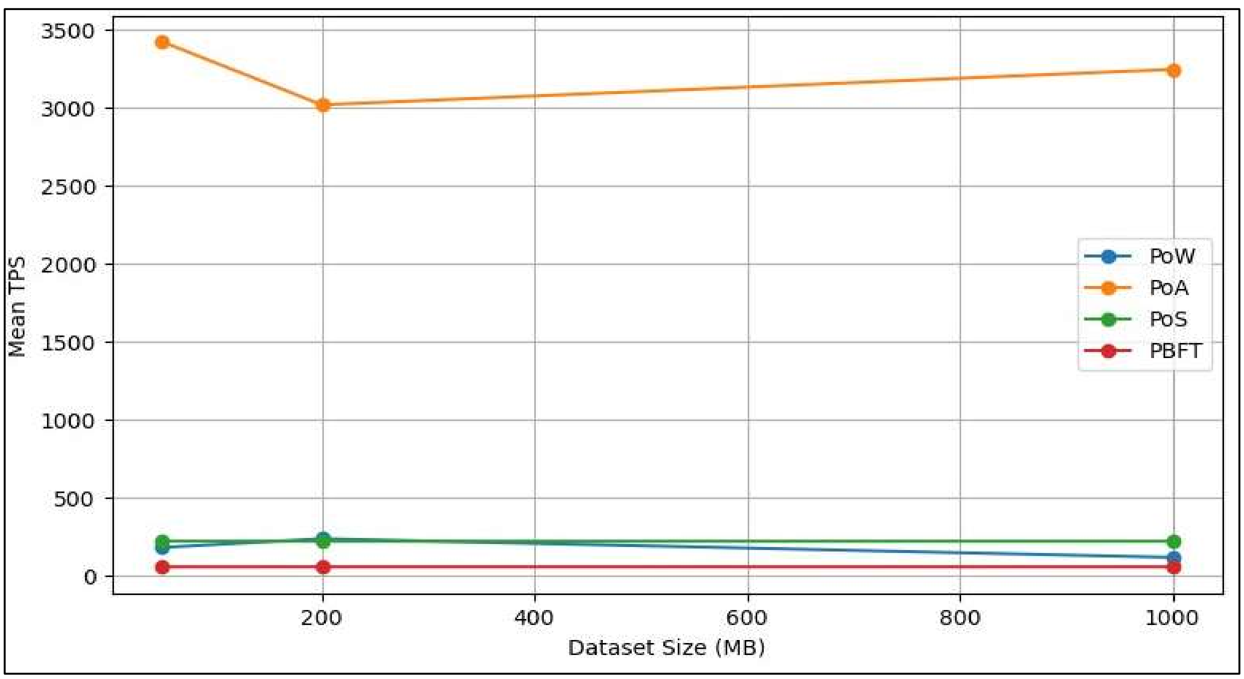

Fig. 3 Mean TPS vs dataset size (by Consensus)

Overall, the throughput findings reveal a clear performance hierarchy. PoA emerged as the most scalable option, while PoW consistently delivered the lowest performance. PoS held a steady middle ground with generally strong results, and PBFT proved reliable but noticeably more sensitive when transaction loads increased. These differences matter for geospatial metadata systems, which depend heavily on fast ingestion and synchronization of large, frequently updated spatial datasets.

3.2 Block Validation Time

The behavior of block validation time closely followed the throughput trends but also provided a deeper look into how each consensus mechanism handles confirmation delays. PoA consistently achieved the fastest validation times across all dataset sizes. This is expected, since PoA validators simply sign and endorse blocks without performing mining or undergoing probability-based selection processes. PBFT came next, followed by PoS, while PoW repeatedly recorded the slowest validation times (Figure 4).

Validation time increased as the datasets grew larger, largely because bigger metadata entries require more work during hashing and serialization. This effect was especially pronounced in PoW. Even small increases in block size significantly extended the time needed to find a valid proof-of-work, reinforcing previous findings that highlight PoW's limited scalability with data-heavy workloads.

PBFT showed stable performance when processing the small and medium datasets, but validation delays rose sharply under the largest dataset (1,000 MB). This slowdown stems from PBFT's dependence on multi-step voting and message exchanges processes that grow heavier as block payloads expand. While PBFT is computationally efficient, its communication overhead eventually becomes a bottleneck.

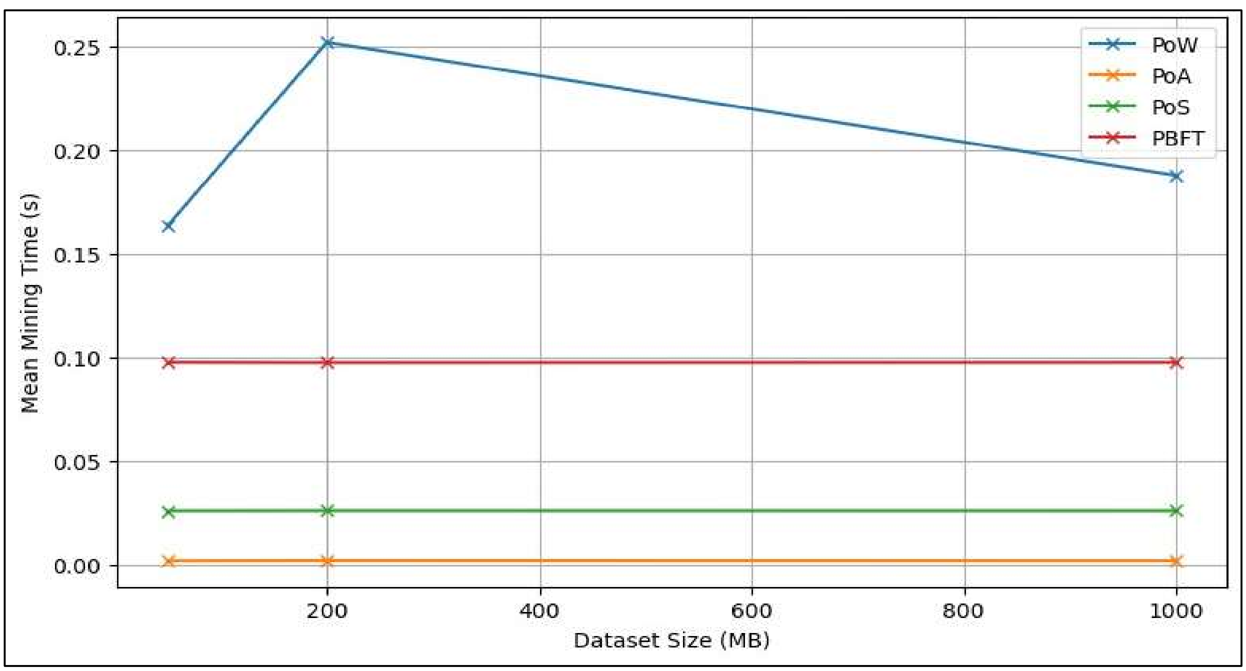

Fig. 4 Mean mining time vs dataset size (by consensus)

Transaction density produced similar patterns. With just two transactions per block, validation times stayed low across all mechanisms. But at ten transactions per block, PoW and PBFT experienced the greatest increases PoW due to rising computational difficulty, and PBFT because more communication rounds were required among validators. PoS and PoA handled higher densities more effectively, though PoS still showed a moderate increase linked to the extra coordination involved in selecting validators based on stake.

These validation time findings reinforce the broader performance pattern seen across the experiment. PoA and PoS remained the most stable and predictable under changing data loads, while PBFT stayed efficient but became increasingly sensitive to the communication overhead created by larger block sizes. PoW, on the other hand, clearly struggled to scale, making it an impractical choice for national geospatial systems that routinely handle large volumes of spatial data.

3.3 Energy Consumption

Energy consumption varied sharply across the four consensus mechanisms, reflecting trends widely reported in recent studies on blockchain efficiency. As expected, PoW recorded by far the highest estimated energy use several orders of magnitude above the other mechanisms. This happens because PoW depends on continuous hash computations to satisfy its difficulty requirement. Even with a fixed block time of 1.5 seconds, PoW still performs a large number of hashing attempts for every block, resulting in substantial computational and energy demand. This behavior is consistent with global assessments that repeatedly identify PoW as the most energy-intensive and environmentally costly consensus protocol.

In contrast, PoS and PoA consumed significantly less energy. PoS relies on a relatively lightweight process for selecting validators, while PoA uses a small, trusted group of authorities whose primary role is to sign and endorse blocks. Because PoA avoids both mining and random-selection cycles, its energy footprint is extremely low often considered negligible in comparison. These patterns align with existing literature, which frequently highlights PoA as one of the most energy-efficient choices for permissioned environments.

PBFT fell somewhere in between. Its energy cost stems mostly from communication overhead rather than computation. With validators exchanging multiple rounds of messages to agree on each block, PBFT uses more energy than PoA or PoS but remains vastly more efficient than PoW.

Dataset size and transaction density had minimal impact on the energy use of PoA and PoS (Figure 5). However, PoW reacted strongly to increased load. Under high-density conditions particularly when processing ten transactions per block in the largest dataset; PoW's estimated energy consumption rose sharply. This shows that larger metadata volumes intensify the mining burden, driving energy costs even higher. Taken together, these findings reinforce the growing consensus that PoW is not a feasible or sustainable option for large-scale or government-operated geospatial systems, where efficiency, affordability, and long-term environmental impact are crucial considerations.

The results show that ledger expansion increased predictably with both dataset size and transaction density. However, the choice of consensus mechanism also shaped the storage behavior of the blockchain system. PoW and PoS produced almost identical ledger sizes because both use comparable block structures and similar transaction serialization formats. By contrast, PoA and

PBFT generated slightly smaller ledgers owing to reduced metadata requirements such as fewer nonce values, simplified validation fields, and lower communication-log overhead. Although these reductions appear minor, they become meaningful in large-scale deployments where long-term storage capacity is a critical constraint.

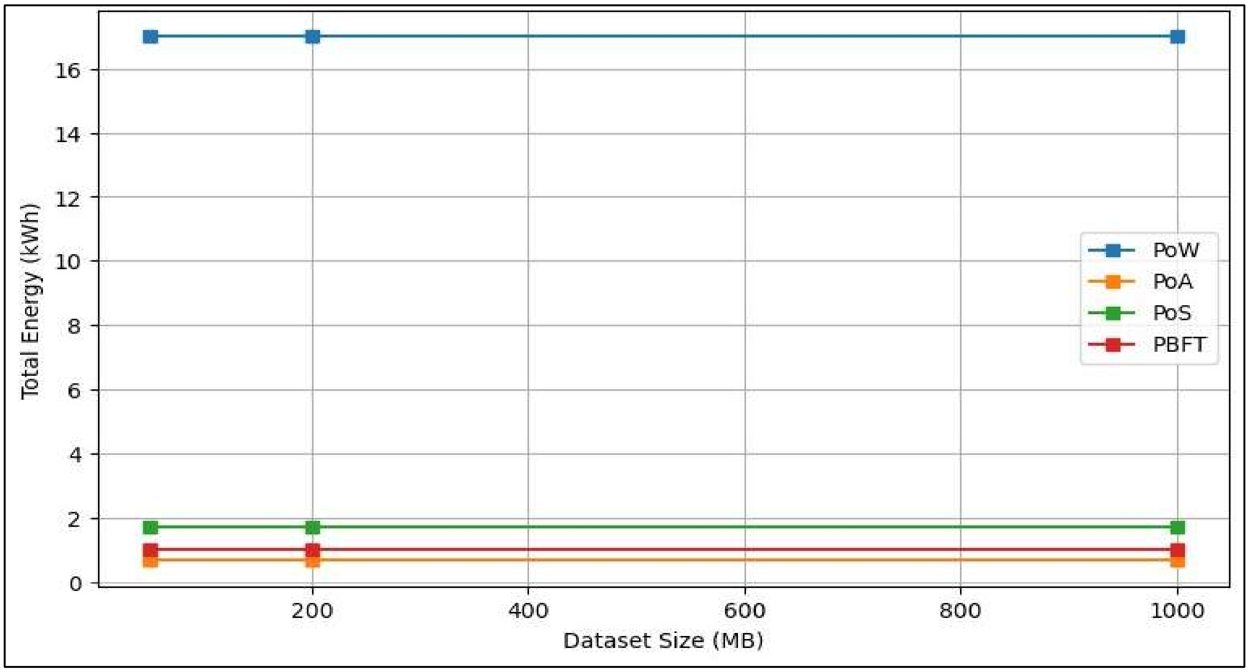

Fig. 5 Total energy vs dataset size (by consensus)

A notable pattern emerged in the high-density block configuration, where blocks contained ten transactions each. Increasing block density effectively reduced ledger bloat by packing metadata across fewer blocks, thereby lowering the total number of blocks generated for a given dataset. This finding is particularly relevant to geospatial infrastructures, where large datasets are common and storage optimization is an ongoing challenge. It also aligns with prior studies that emphasize block density as a key determinant of storage efficiency in blockchain-based spatial systems (Tang et al., 2024).

Nonetheless, even with these optimizations, the largest dataset (1,000 MB) produced significant ledger growth across all consensus mechanisms. This reinforces the need for strategic storage solutions such as off-chain data management, pruning, or decentralized storage systems like IPFS approaches widely recommended in geospatial blockchain literature (Zhao et al., 2022).

3.4 Tamper Detection Accuracy

Tamper detection achieved true in all scenarios, regardless of the consensus algorithm or dataset size (Table 1). This outcome was expected, as blockchain's linked-hash architecture inherently flags any modification to an existing transaction by invalidating all successive blocks. This characteristic remains one of blockchain's strongest advantages over traditional centralized databases and supports findings from previous geospatial metadata research (Zhao et al., 2022).

Table 1. Tamper detection in different scenarios

Consensus	Avg TPS (Max)	Avg Mining Time (s)	Energy Use (kWh)	Tamper Detection	Suitability
PoW	~ 180 – 850	0.04 – 0.73	6 – 30	True	Poor for SDI
PoA	~ 1173 – 5643	~ 0.0017	0.24 – 1.2	True	Excellent for SDI
PoS	~ 77 – 389	~ 0.025	0.6 – 3.0	True	Good for SDI
PBFT	~ 49 – 66	~ 0.04– 0.15	0.36 – 1.8	True	Moderate for SDI

While accuracy was uniform, detection speed varied. PBFT was consistently the fastest, benefiting from its message-driven verification model in which all nodes perform integrity checks in parallel. PoA and PoS followed, showing slightly slower but still efficient detection times. PoW performed the slowest due to its need to recompute proof-of-work hashes before inconsistencies can be identified, thereby increasing computational and time overhead. These results demonstrate that trust-based (PoA) and communication-coordinated (PBFT) consensus mechanisms provide far more efficient integrity verification than computation-intensive PoW.

3.5 Synthesis of Findings

When all performance metrics are considered together, clear behavioural patterns emerge:

- **PoA** demonstrated strong performance across throughput, validation time, energy use, storage behaviour, and tamper detection. Its predictable scalability and minimal resource requirements make it well-suited for large geospatial datasets.
- **PoS** delivered balanced, moderate performance, offering improved efficiency compared to PoW while maintaining decentralization principles.
- **PBFT** produced very low-latency results but encountered scalability constraints in larger datasets due to its extensive communication requirements.
- **PoW** performed the poorest across most metrics, reaffirming its unsuitability for large-scale, data-intensive geospatial systems.

These findings are closely aligned with broader research, which identifies PoA and PBFT as the most appropriate for permissioned, institutionally managed infrastructures such as national land registries, spatial planning systems, and mapping authorities (Rebelo Marcolino et al., 2025; Rukhiran et al., 2024). In Papua New Guinea where digital infrastructure, computational capacity, and energy resources remain limited PoA stands out as the most practical option. It combines excellent energy efficiency with strong performance under realistic geospatial workloads, making it particularly viable for government-led spatial data systems.

3.6 Limitations and Future Work

While the simulations provide valuable benchmarks, limitations include the use of synthetic metadata, exclusion of real-world network delays, and estimated energy models rather than hardware-specific testing. Future work should extend this study with real PNG datasets (e.g., provincial land surveys, forestry monitoring, and disaster-response imagery) and incorporate network constraints typical of rural and provincial connectivity. Exploring hybrid consensus protocols, such as combining PoA's efficiency with PBFT's fault tolerance, could also yield more tailored solutions for PNG's SDI environment.

4. Conclusion and Recommendation

This study compared four blockchain consensus mechanisms: PoW, PoA, PoS, and PBFT for geospatial metadata distribution. Results from 72 simulation runs showed that PoA consistently outperformed other protocols, delivering the highest throughput, lowest validation time, and best energy efficiency. PoS provided balanced performance, while PBFT was stable but limited by low throughput. PoW was the least suitable due to its high energy consumption and variable performance.

For PNG's Spatial Data Infrastructures (SDIs), where reliability, sustainability, and trust are critical, PoA is the most suitable choice, with PoS as a secondary option, and PoW should be avoided for collaborative provincial systems.

5. Acknowledgement

The authors would like to thank the School of Mathematics & Computer Science for providing the support to carry out this research. The author extends thanks to the organizing committee for their initiative and to the Papua New Guinea University of Technology for providing the academic environment that made this work possible.

6. References

1. Li, R. (2024). Comparative Analysis and Future Directions of Consensus Algorithms in Blockchain Technology. In Y. Wang (Ed.), *Proceedings of the 2024 2nd International Conference on Image, Algorithms and Artificial Intelligence (ICIAAI 2024)* (Vol. 115, pp. 333–344). Atlantis Press International BV. https://doi.org/10.2991/978-94-6463-540-9_34
2. Liu, Y., & Chang, Y. (2024). Blockchain-Based Method for Spatial Retrieval and Verification of Remote Sensing Images. *Sensors*, *24*(7), 2078. https://doi.org/10.3390/s24072078
3. Rebelo Marcolino, M., Reis Porto, T., Thompsen Primo, T., Targino, R., Ramos, V., Marques Queiroga, E., Munoz, R., & Cechinel, C. (2025). Student dropout prediction through machine learning optimization: Insights from moodle log data. *Scientific Reports*, *15*(1), 9840. https://doi.org/10.1038/s41598-025-93918-1
4. Rukhiran, M., Boonsong, S., & Netinant, P. (2024). Sustainable Optimizing Performance and Energy Efficiency in Proof of Work Blockchain: A Multilinear Regression Approach. *Sustainability*, *16*(4), 1519. https://doi.org/10.3390/su16041519
5. Tang, M., Dai, W., Yin, C., Hu, B., Chen, J., & Liu, H. (2024). Spatial Planning Data Structure Based on Blockchain Technology. *ISPRS International Journal of Geo-Information*, *13*(8), 290. https://doi.org/10.3390/ijgi13080290
6. Zhao, P., Cedeno Jimenez, J. R., Brovelli, M. A., & Mansourian, A. (2022). Towards geospatial blockchain: A review of research on blockchain technology applied to geospatial data. *AGILE: GIScience Series*, *3*, 1–6. https://doi.org/10.5194/agile-giss-3-71-2022

7. Zimba, A., Phiri, K. O., Mulenga, M., & Mukupa, G. (2025). A systematic literature review of blockchain technology and energy efficiency based on consensus mechanisms, architectural innovations, and sustainable solutions. *Discover Analytics*, *3*(1), 14. https://doi.org/10.1007/s44257-025-00041-6

Author's Bibliography

Wera Dawa, currently a Research Student (MPhil) Under the School of Mathematics and Computer Science at the Papua New Guinea University of Technology, Lae, Morobe Province, Papua New Guinea. He recently graduated with a Bachelor's Degree in Computer Science in April 2025. He is a Research Student who is passionate about being a Scholar. His main research work focusses on, Computer Networking, Deep Learning & Artificial Intelligent, Data Science.

Arun Kumar Singh is currently the Section Head of Computer Science at the School of Mathematics and Computer Science, PNG University of Technology, Lae, Papua New Guinea. He received B. E. (2002) degree in ECE from Agra University, Agra, M. Tech. (2005) Degree in IT-WCC from IIIT-Allahabad, Prayagraj, India and PhD (2013) degree from SU Meerut, India. He has published more than 50 research papers in reputed international journals and conferences including IEEE and it's also available online. His main research work focuses on Cybersecurity, AI, Sustainability, IoT and Computational Intelligence. He has 20 years of teaching experience and 15 years of research experience.