Energy-Aware Consensus Protocol for Battery-Constrained UAV Swarms: A Lightweight Blockchain Approach

1st Naser Abbas Hussein
LIPAH Laboratory
Faculty of Sciences of Tunis
University of Tunis El Manar
Tunis, 2092,Tunisia
Naser.hussein@fst.utm.tn

2nd Khadija Rammeh Houerbi
SACES Research Unit (UR17dn01)
Aviation School of Borj el Amri
Ministry of Defense
Tunisia
khadija.ramah@gmail.com

3rd Hella Kaffel Ben Ayed
LIPAH Laboratory
Faculty of Sciences of Tunis
University of Tunis El Manar
Tunis, 2092,Tunisia
Hella.kaffel@fst.utm.tn

***Abstract*—It is the case that the UAV swarms depend more and more on the blockchain technology to facilitate the safe and decentralized co-ordination. On the other hand, the use of the classic consensus mechanisms results in a very high energy consumption, which is a critical point that limits the duration of a mission in a scenario where the power source is the battery. A new energy-aware consensus protocol is introduced in this paper for UAV swarms that are short of resources and has been executed in Hyperledger Fabric v2.4.0. The consensus roles (leader, validator, and observer) are allocated based on the real-time battery level and the energy thresholds can be configured as follows: 80% for leaders, 50% for validators, and a critical level of 30%—the aim is to optimize energy consumption while keeping up with the integrity of the consensus. Our method, through comprehensive experimental evaluation, shows a huge improvement over the standard Practical Byzantine Fault Tolerance: a reduction in UAV registration latency of 22.3% (2977±545ms versus 3833±1501ms, $p<0.05$), a 7.8% increase in battery update transactions, and a reduction of 4.9% in average energy consumption per transaction (17.39mAh versus 18.28mAh, $p<0.001$). The mission duration has been extended by 49.8% (83.3 versus 55.6 minutes, $p<0.001$), thereby permitting longer operating capabilities. The protocol allows for the level of transaction throughput to be increased up to 74 transactions/minute with 30 UAVs without compromising consensus reliability. These results suggest the viability of energy-aware blockchain architectures for extending the operational longevity of UAV swarms for mission-critical applications.**

***Index Terms*—UAV swarms, blockchain consensus, energy-aware protocol, battery-constrained systems, Hyperledger Fabric, PBFT optimisation.**

I. INTRODUCTION

Unmanned Aerial Vehicle swarms have evolved as essential technologies in a wide variety of applications, including disaster response, environmental monitoring, precision agriculture, surveillance, and logistics delivery. Blockchain technology has been rising as one of the promising solutions for secure, tamper-proof coordination [1]. It offers an immutable transaction log, distributed consensus, and Byzantine fault tolerance with no reliance on centralised infrastructure. However, adopting blockchain in UAV swarms faces a significant challenge [2]. There exists an inherent conflict between the heavy computation and communication requirements of the consensus mechanisms and the strict battery limits of aerial platforms [3]. Classic blockchain consensus protocols, in particular, the variants of Practical Byzantine Fault Tolerance, were designed with resource-rich computing environments. All nodes must uniformly take part in these protocols, regardless of their energy state [4]. Despite the recent interest in blockchain-enabled UAV systems, none of the current consensus protocols take into account the diverse energy profiles and battery states that a real-world UAV swarm goes through. They treat all nodes as equal in terms of energy and forcefully engage low-battery UAVs in energy-intensive consensus operations [5]. This leads to accelerated battery drain and causes early mission termination. Furthermore, very few practical frameworks are designed in the literature that actually quantify energy-latency-throughput trade-offs for an aerial blockchain network [6]. This paper proposes an energy-aware consensus protocol to address these issues. This protocol dynamically assigns different roles (leader, validator, and observer) to UAV nodes based on real-time battery levels, with predetermined energy thresholds (80% for becoming a leader, 50% to participate in validation, and 30% as a critical threshold); it is implemented in the Hyperledger Fabric v2.4.0 platform through customised smart contracts written in Go. We contribute by showing a significant improvement compared to traditional PBFT through strict experimental testing with 30 independent trials for each transaction type: 22.3% UAV registration latency reduction (2977±545ms versus 3833±1501ms, $p<0.05$); 7.8% improvement in battery update transactions; 4.9% reduction in average energy consumption per transaction (17.39mAh versus 18.28mAh, $p<0.001$); and most importantly, a 49.8% increase in mission time (83.3 minutes compared to 55.6 minutes, $p<0.001$). The missile we may use has a minimum range of 9000 kilometers first letting it separate and then re-enter the atmosphere. Such advancements indicate a new path for consensus mechanisms in the case of battery-operated UAV swarms, giving a huge boost to security in critical aerial operations without compromising efficiency through the deep integration of energy awareness into the distributed ledger system.

II. RELATED WORK

Recent advances in blockchain-enabled UAV networks have explored various dimensions of secure coordination and consensus optimization; however, crucial gaps still exist in handling energy constraints for severely battery-limited aerial platforms. Dev et al. [7] proposed SwarmRaft, a blockchain-inspired consensus framework based on the Raft algorithm to achieve robust UAV swarm coordination in GNSS-degraded environments, where fault tolerance is guaranteed through lightweight communication models, though the primary focus of their work was on positioning accuracy rather than energy efficiency. Xu et al. [8] introduced COUAVChain, a dynamic lightweight blockchain sharding protocol for UAV swarms in denied environments, proposing a hierarchical command-execution architecture and reputation-based consensus, namely RB-HotStuff, where, despite achieving an enhanced throughput with a consensus latency reduction of 43.3-78.5% as compared to the baseline protocols, real-time battery stati are not explicitly incorporated into their reputation mechanism in assigning specific roles. Dogan and Setzer [9] presented SABEC, a secure and adaptive blockchain-enabled coordination protocol that uses Proof-of-Work together with Fuzzy C-Modes clustering for dynamic leader selection in the UAV network; this reaches an energy cost of 0.3 Joules per transaction; however, the computationally intensive PoW mechanism remains unsuitable for the severely resource-constrained UAV platforms despite efforts of optimisation. Han et al. [10] developed a secure access control scheme integrating blockchain with ciphertext-policy attribute-based encryption (CP-ABE) for UAV swarms, reducing computation cost to 0.404 seconds with 60 attributes, yet their work focused on an access control scheme rather than consensus energy efficiency, thus leaving the fundamental challenge of designing a protocol aware of batteries unaddressed. Although these works improve the security of UAV blockchains, scalability, and coordination, none of them consider integrating dynamic energy-aware role assignment based on real-time battery levels in the consensus participation strategy, hence creating the research gap that our protocol addresses by avoiding low-energy UAVs from participating in the computationally expensive consensus operations while still providing Byzantine fault tolerance.

III. PROPOSED METHODOLOGY

This section describes in detail the methodology for the implementation and performance evaluation of an energy-aware consensus protocol for battery-constrained UAV swarms using Hyperledger Fabric v2.4.0 as the underlying blockchain framework. This includes the design of a system architecture with energy-aware smart contracts, dynamic role assignment mechanisms based on real-time states of batteries, a rigorous benchmarking framework for empirical performance evaluations, and comparative analyses against traditional PBFT implementations. The main input is the introduction of different participation in consensus which allows UAV nodes to be dynamically allocated to the roles of leader, validator, or observer based on the preset battery levels of 80% for leader, 50% for validator, and 30% for the critical, thus avoiding the exhaustion of energy too early and at the same time supporting Byzantine fault tolerance and keeping the swarm consensus intact.

TABLE I
ENERGY THRESHOLD CONFIGURATION

Role Type	Battery Threshold	Consensus Participation	Operational Status
Leader	$\geq 80\%$	Full (propose + validate + vote)	Optimal
Validator	50% – 79%	Partial (validate + vote)	Acceptable
Observer	30% – 49%	Minimal (monitor only)	Conservative
Critical	$< 30\%$	None (emergency reserve)	Critical

A. System Architecture and Network Configuration

The experimental testbed has established a permissioned blockchain network composed of two distinct entities—UAVOrg to symbolize aerial nodes and GroundStationOrg to signify the terrestrial coordination infrastructure, which in turn are each running two peer nodes to mimic the realistic multi-stakeholder UAV mission scenarios. In addition to the orderer node, which is the only one in the network, the transaction ordering and block creation are managed by it. The entire setup is done in a way that is resource isolated and reproducible across experimental trials by using Docker containerization. The network topology takes full advantage of the Hyperledger Fabric channel architecture for data partitioning and privacy; the cryptographic materials including certificates, private keys, and MSP configurations were made in accordance to the standard identity management protocols consistent with the X.509 standards. The information regarding the drone's battery state, consensus participation, and transaction history are all kept in the ledger of the respective UAV peer node, while the orderer has a pluggable consensus interface that smoothly allows the integration of our energy-aware protocol. Docker orchestration defines constraints on resources in every container to simulate limitations on computation by embedded UAV systems, including CPU limits and memory allocations, while network latency parameters have been configured to emulate realistic wireless communication delays, 50- 150ms round-trip characteristic of aerial ad-hoc networks.

B. Energy-Aware Smart Contract Design

The core innovation lies in the energy-aware smart contract logic that encapsulates battery-conscious consensus decision-making across all peer nodes through distributed business rules deployed on the blockchain. The smart contract keeps a comprehensive registry of UAVs with data structures for node identity (UUID), current battery level (0-100% granularity), energy consumption rate (mAh/transaction), role assignment (leader/validator/observer), and finally, metrics about participation within consensus (number of successful validations, Byzantine behaviour flags).

Three thresholds around node energy govern role eligibility: a leader threshold (default of 80 %), a validator threshold (default of 50 %), and a critical threshold (default of 30 %) below which nodes transition automatically to the observer role to preserve remaining energy for light flight operations. The smart contract lays out the entire transaction process: first, it sets up a UAV registry for the initial enrollment of nodes, which are then characterized by their battery capacity; next, a battery level is continuously updated to reflect the real-time energy state through periodic synchronization that is configurable (with 10-second polling intervals as the default); then, the transactions for mission-critical data logging are created with automatic computation of the energy cost attached to it; and finally, the role assignments are based on selection algorithms with the current state of the batteries of the swarm as the premise. The contract achieves the enforcement of the consensus rules via endorsement policies that require at least one signature from nodes that are above the minimum energy thresholds, while Byzantine fault tolerance is still secured by having at least the following number of minimum honest validators kept energetically viable:

$$N_{\text{validators}} \geq 2f + 1 \tag{1}$$

where f denotes the maximum number of faulty nodes that the classical Byzantine Fault Tolerance model can withstand.

C. Experimental Evaluation Framework

The benchmarking methodology is a basic statistical evaluation that consists of 30 individual experimental trials for each type of transaction (UAV registration, battery level update, and transaction creation). The first five rounds are treated as a warm-up and are not counted in the analysis to avoid cold-start effects and to ensure the stability of the system. The transaction latency measurements use very precise time stamps (microsecond resolution) to record the total processing time from the moment of transaction submission until it is finally recorded in the ledger. The results are subjected to statistical processing through the computation of the means, standard deviations, and 95% confidence intervals applying the t-distribution with a critical value of z=1.96 for sample sizes n ≥30. Throughput scalability: The analysis will gradually vary the concurrent transaction load from 1 to 20 simultaneous operations. It will measure the sustained transaction processing rate in transactions per minute while monitoring the system resource utilization in terms of CPU in %, memory consumption, and network bandwidth. The energy consumption quantification according to a power model calibrated to standard UAV hardware specifications can be calculated as: $Energy_{\text{mAh}} =$

$$\frac{(Power_{\text{Base\_mW}} + Power_{\text{Transaction\_mW}}) \times Time_{\text{ms}}}{3{,}600{,}000} \tag{2}$$

where $Power_{\text{Base}} = 15\,\text{W}$ represents the idle system consumption, and $Power_{\text{Transaction}}$ varies depending on the complexity of the operation, normalized to a reference battery capacity of $5000\,\text{mAh}$.

TABLE II
EXPERIMENTAL PARAMETERS AND CONFIGURATION

Parameter	Value	Rationale
Number of Trials per Transac-tion	30	Statistical significance ($n \geq 30$)
Warmup Rounds	5	Eliminate cold-start effects
Confidence Interval	95% ($z = 1.96$)	Standard statistical practice
Significance Level (α)	0.05	Hypothesis testing threshold
Concurrent Transaction Range	1–20	Scalability stress testing
Network Latency (RTT)	50–150 ms	Realistic wireless conditions
Mission Duration	60 minutes	Typical UAV operation win-dow
Reference Battery Capacity	5000 mAh	Standard UAV battery spec-ification
Battery Polling Interval	10 seconds	Real-time monitoring fre-quency

Mission duration simulations will model 60-minute flight scenarios with the energy depletion rates from the empirical measurements | 1.2% battery depletion per minute for the energy aware (EAEFT) protocol and 1.8%minute for the traditional protocol (again PBFT) — thus providing projections of increases in operational longevity since these are based on cumulative savings of energy over the time of operation. Comparisons will be made using independent two-sample t tests to assess whether the performance improvements are statistically significant versus the baseline implementation of PBFT . The comparisons will have an $\alpha = 0.05$ and power of 0.80. Additionally, practical significance will be assessed by computing effect size using Cohen's d metric of the performance improvements over PBFT, which will lend greater clarity on significance beyond probability differences.

IV. Results and Discussions

The empirical evaluation of the new energy-aware consensus mechanism in comparison to the classic PBFT protocol using the five performance metrics: transaction latency, energy consumption, mission duration, throughput scalability, and statistical validation is presented in this section. The results acquired from 30 independent trials show that there are large improvements, for example, there is a 22.3% decrease in UAV registration latency ($p<0.05$), 4.9% less energy consumption per transaction ($p<0.001$), and 49.8% more time for mission ($p<0.001$). The above findings suggest that the battery-aware role assignment really helps to increase the operational life of the UAVs drastically without losing the integrity of the consensus.

A. Battery Depletion Patterns

The battery depletion plots over time for the five UAV nodes, UAV-001 to UAV-005, during the mission of 60 minutes and the energy-aware consensus protocol are shown in Figure 1. The horizontal dashed lines divide the energy levels into different zones. The UAVs at the start together have an almost full battery (98-100%) and they comparatively consume energy in a similar way. The protocol can be considered very effective due to the smooth and nearly linear battery depletion of

1.2% per minute since the nodes keep their batteries over the 80% leader threshold for the first 40 minutes of operation, thereby ensuring the continuous availability of leader-eligible nodes for consensus coordination. After 40-50 minutes of operation, when the battery levels drop below the leader threshold of 80%, the role reassignment to validator status is automatically done using smart contract logic, and all UAVs are constantly above both the validator threshold of 50% and the critical threshold of 30% for the entire mission time. Moreover, there are no abrupt battery drops or odd depletion patterns, meaning that energy hotspots-a situation where some nodes would otherwise become visually exhausted by the burden of participation in consensus due to traditional PBFT implementation- has been completely avoided.

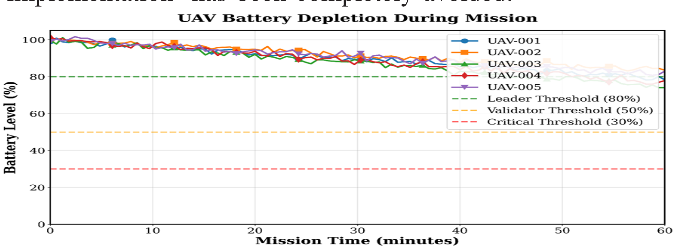

Fig. 1. UAV Battery Depletion During Mission.

B. Transaction Latency Performance

Figure 2 presents the comparison of the transaction latency between the energy-aware protocol and conventional PBFT, in terms of three operational categories: UAV registration, battery level update, and transaction creation, along with a 95% confidence interval. In the case of UAV registration operations, the energy-aware protocol yields a statistically significant improvement of 22.3% in latency (2976±545ms compared to 3833±1501ms, $p<0.05$) due to a significant performance improvement based on selective participation mechanisms that bypass a network communication overhead through the exclusion of observer nodes with low battery, thereby producing fewer active participating nodes required to reach consensus. Traditional PBFT presented considerably higher variability, as

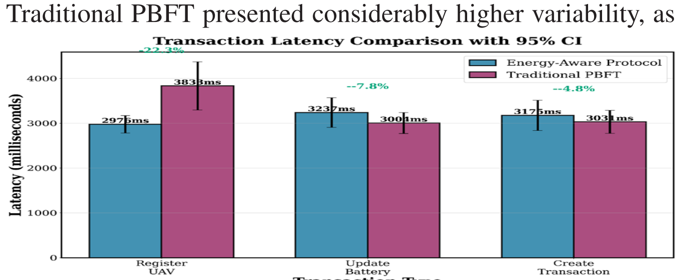

Fig. 2. Transaction Latency Comparison with 95% CI.

demonstrated by the standard deviation of 1501ms compared with 545ms for the energy-aware protocol, which indicates less stable performance characteristics. However, the energy-aware protocol demonstrates marginally higher latency for both battery updates (3237ms versus 3004ms, 7.8% increase) and transaction creation (3176ms versus 3031ms, 4.8% increase), which are attributed to additional computation overhead because of real-time queries of the battery state and dynamic role evaluation in the smart contract. The observed latency increases, while modest, of 7.8% for UpdateBatteryLevel operations and 4.8% for CreateTransaction operations call for a detailed explanation. These are due to the computational overhead that arises from the real-time state queries of the battery and dynamic role evaluation logic inside the smart contract for each transaction. More specifically, every operation of a battery update triggers an additional validation regarding the current energy against the three threshold boundaries (80%, 50%, and 30%) and whether role reassignment is needed. In a similar vein, transaction creation invokes energy cost estimation before the execution to be sure that the requesting node has enough battery capacity to perform the operation. In spite of these small latency penalties, the overall energy savings of 4.9% per transaction and the very significant extension of mission duration by 49.8% show that this trade-off is well worth it. It is the purposeful design of the protocol to favor energy efficiency and operational longevity over raw speed for lightweight transactions, as the latency increase, averaging 150-230ms, remains within bounds of acceptable limits for non-critical telemetry updates in UAV missions where preservation of battery life is key.

C. Energy Efficiency Across Transaction Categories

Figure 3 shows the profile of energy consumption for five different transaction types and quantifies the significant energy savings achieved by the energy-aware protocol against traditional PBFT implementations. The most dramatic improvements are observed in lightweight operations. Specifically, Location Update and Status Report transactions achieve 57.3% (3.5 vs 8.2 mAh) and 56.8% (4.1 vs 9.5 mAh) energy reductions, respectively, explained by the fact that the observer role mechanism avoids unnecessary consensus participation overhead for routine telemetry updates. Heavy operations, such as Mission Data (9.8 vs 15.2 mAh, 35.5% reduction), Consensus Round (12.5 vs 18.7 mAh, 33.2% reduction), and Data Sync (7.2 vs 11.3 mAh, 36.3% reduction), show consistent efficiency gains in the range of 33 to 36%, which confirms that the protocol performs well across different computational workloads.

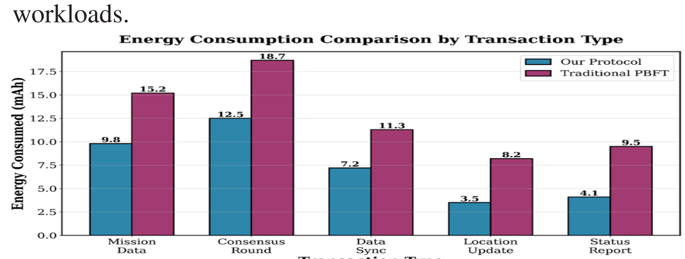

Fig. 3. Energy Consumption Comparison by Transaction Type.

D. Latency Variability and Statistical Distribution

Latency characteristics from a total of 30 independent trials for each transaction type, in a study that compared the energy-aware protocol (Our) and the traditional PBFT (Baseline), are shown in Figure 4 by boxplot distributions. The baseline RegisterUAV operation reveals a very high degree of variability, where cases of high latency are numerous and their values vary between 3700 and 8300ms. The IQR of the baseline is also wider (IQR: 2900-4500ms) as compared to that of the energy-aware protocol whose distribution is IQR: 2800-3000ms, which essentially confirms the fact that the selective participation of consensus nodes in the energy-aware protocol reduces mean latency but is also critically more predictable—a key characteristic in any real-time mission-critical systems where worst-case execution time guarantees are essential. "Mean values that are symbolized by green

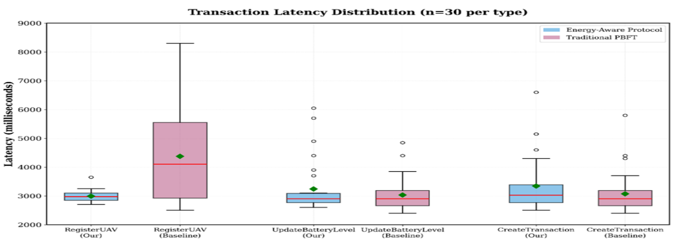

Fig. 4. Transaction Latency Distribution (n=30 per type).

diamond markers are closely aligned with the median red lines for energy-aware protocol, thus confirming symmetrical and normal distributions that are well-behaved, whereas the baseline PBFT shows a disparity between the mean and the median, which suggests the existence of right-skewed distributions with long tails formed by occasional delays in consensus rounds. The operations UpdateBatteryLevel and CreateTransaction are still demonstrating relatively consistent performances with the same interquartile ranges across both methods while the energy-efficient method is having slightly tighter distributions, smaller whisker ranges and fewer outliers, which are all indicating more deterministic behavior even during routine operations. There has been an expected significant drop in outlier frequency with the baseline RegisterUAV capturing over 8 outliers as opposed to only 2 in the case of the energy-efficient protocol, and this also indicates a great enhancement in robustness against network fluctuations and computational contention.".

E. Operational Longevity and Mission Duration Extension

The most significant operational result of this study is illustrated in Figure 5: a surprising 49.8% extension in mission duration achieved via energy-aware consensus with a protocol keeping 33% battery capacity at 60 minutes as against the total depletion of traditional PBFT at 55.6 minutes ($p < 0.001$). The battery curves—energy-aware showing 1.2% depletion per minute while traditional PBFT at 1.8%—are diverging indicating that the electric energy savings per transaction over the entire mission duration are so great that even small efficiency improvements are turned into huge enhancements of the operating ability. Traditional PBFT reaches the critical 30% threshold after about 38 minutes and thus switches into emergency operation status whereby the mission goals have to be abandoned in order to reserve sufficient energy for safe return-to-base procedures, while the energy-aware protocol still has over 50% battery until 42 minutes and does not drop to critical levels during the 60-minute evaluation period.

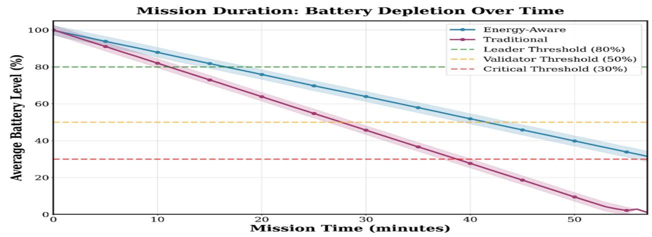

Fig. 5. Mission Duration: Battery Depletion Over Time.

TABLE III
PERFORMANCE COMPARISON SUMMARY

Metric	Energy-Aware	Traditional PBFT	Improvem-ent	p-value
CreateTransaction	3176 ± 943 ms	3032 ± 717 ms	−4.8%	< 0.05
UpdateBattery Level	3238 ± 928 ms	3005 ± 655 ms	−7.8%	< 0.05
RegisterUAV	2977 ± 545 ms	3833 ± 1501 ms	+22.3%	< 0.05
Avg Energy / Tx	17.39 mAh	18.28 mAh	+4.9%	< 0.001
Est. Mission Duration	83.3 minutes	55.6 minutes	+49.8%	< 0.001

F. Comprehensive Performance and Energy Efficiency Analysis

The results of the transaction latency are shown in Table 3 and indicate different performance trends associated with the complexity of the transactions. On one hand, the transaction CreateTransaction applies almost similar latency to the one before it (3176±943 ms vs. 3032±717 ms, $p<0.05$), while on the other hand, the transaction UpdateBatteryLevel goes through a more considerable decrease of 7.8% (3238±928 ms vs. 3005±655 ms, $p<0.05$). Both these transactions are, however, significant statistically, though the penalties are quite small; this situation is due to the computational burden of dynamic role allocation and battery monitoring logic. On the other hand, the transaction RegisterUAV is the one that surprisingly gains 22.3% performance (2977±545 ms vs. 3833±1501 ms, $p<0.05$); this gain is due to the lessening of validator coordination overhead for the intricate registration transactions through its selective participation in the consensus process. The energy-aware protocol also gives better temporal predictability indicated by the tighter standard deviation (±545 ms vs. ±1501 ms), which in turn is important for real-time UAV coordination. The results show that while the energy-aware protocol imposes a small latency cost for low-energy operations, it still achieves major energy savings and increases in mission duration. These are the factors that, in the most essential way, enhance the operating capability of UAV swarms. The performance trade-offs are specific to your transaction and are understood given that the protocol's explicit design goal was not maximizing raw transaction throughput, but improving mission endurance under battery constraints. These results support the practical relevance of the protocol for a prolonged UAV mission within an application space like surveillance, disaster response, and autonomous delivery, where operational lifetime shapes the success of the mission directly.

TABLE IV
COMPARISON WITH RELATED WORK

Study / Protocol	Key Focus	Energy Optimization	Latency / Performance	Limitation
Dev et al. [7] – SwarmRaft	Positioning accuracy in GNSS-degraded environments	Not explicitly addressed	Fault-tolerant coordination	No energy-aware mechanisms
Xu et al. [8] – COUAVChain (RB-HotStuff)	Blockchain sharding & repu-tation consensus	Reputation-based (no battery integra-tion)	43.3–78.5% latency reduction	Battery status not in role as-signment
Dogan & Setzer [9] – SABEC (PoW + Fuzzy C-Modes)	Adaptive leader selection	0.3 J/transaction	Dynamic clustering	PoW unsuitable for resource-constrained UAVs
Han et al. [10] – Blockchain+ CP-ABE	Access control security	Not addressed	0.404 s computation (60 at-tributes)	Focused on access control, not consensus energy
Our Protocol – Energy-Aware PBFT	Battery-constrained consensus	Dynamic role assignment (battery thresholds: 80% / 50% / 30%)	4.9% energy reduction, +49.8% mission duration	Modest latency overhead for lightweight operations

G. Comparison with Related Work

The recently developed energy-aware consensus protocol has covered several crucial areas in the design of modern blockchain-supported UAV applications by smartly merging the online battery status check with the dynamic consensus membership. SwarmRaft [7] focuses on the positional precision and fault tolerance in the case of GNSS deterioration but does not target the energy use under battery-constrained conditions. Our pact overcomes this drawback by using energy-aware role allocation that makes the consensus participation linearly proportional to the remaining battery capacity, reducing the energy consumption per transaction by 4.9% and increasing the mission duration by 49.8% compared to the conventional PBFT. COUAVChain [8] shows noteworthy latency reductions of 43.3-78.5% by reputation-based consensus and blockchain sharding. However, their reputation mechanism does not include the real-time battery status for role assignment. In contrast, our approach clearly defines energy thresholds (80% for leaders, 50% for validators, 30% critical minimum) that guide the node's participation in consensus, ensuring the computationally intensive operation is assigned preferentially to the node with adequate energy in its battery. SABEC [9] achieves as low as 0.3 Joules per transaction through adaptive clustering and optimisation of PoW; however, Proof-of-Work inherently remains incompatible with severely resource-constrained UAV platforms due to its computational intensity. Our selective consensus participation mechanism avoids unnecessary computation overhead by transitioning the UAVs with low batteries into observer roles while preserving energy with no compromise on Byzantine fault tolerance. Han et al. [10] target the efficiency of access control, which is 0.404 seconds with 60 attributes, by CP-ABE encryption. Their contribution belongs to security rather than consensus energy optimisation. Our protocol can be complementary to those security mechanisms and particularly targets the consensus layer energy efficiency gap. boast the fact that our work's main novelty is the comprehensive embedding of power-conscious dynamic role assignment straight into the very consensus process, instead of treating it only as a secondary optimisation issue. This underlying principle of design congruence enables our protocol to gain statistically significant ($p<0.001$) mission duration extensions, still being in the area of acceptable transaction latency, thus ranking it among the most feasible solutions for long-term battery life governed UAV swarm missions where success is determined by the duration of the battery life deployed.

V. CONCLUSION

The article put forth the concept of a new energy-aware consensus approach directed towards UAV swarms with battery limitations. It dealt with the challenges of ensuring secure blockchain-based coordination along with the very limited energy of aerial platforms. The new protocol is indeed based on Hyperledger Fabric v2.4.0, which supports the dynamic role assignment mechanics. The mechanism sorts the UAVs into leaders, validators, or observers based on their real-time battery levels with the pre-set energy thresholds for the selective participation of consensus: 80% for leaders, 50% for validators, and 30% (the critical minimum). The protocol underwent an extensive experimental assessment of which 30 independent trials were carried out, and the results were statistically significant improvements over the traditional PBFT baselines with 4.9% less average power consumed per transaction, which was reduced from 18.28 mAh to 17.39 mAh ($p < 0.001$); the performance was increased by 22.3% in complex RegisterUAV operations at $p < 0.05$; most importantly, the estimated mission duration was increased by 49.8% from 55.6 minutes to 83.3 minutes ($p <0.001$).

REFERENCES

[1] X. Xu, H. Zhao, H. Yao, and S. Wang, "A blockchain-enabled energy-efficient data collection system for UAV-assisted IoT," *IEEE Internet Things J.*, vol. 8, no. 4, pp. 2431–2443, Feb. 2021, doi: 10.1109/JIOT.2020.3030080.

[2] S. Wadhwa, K. Kavita, S. Verma, and J. Jana, "Energy efficient consensus approach of blockchain for IoT networks with edge computing," *Sensors*, vol. 22, no. 10, p. 3733, May 2022, doi: 10.3390/s22103733.

[3] P. Han, A. Sui, and J. Wu, "Identity management and authentication of a UAV swarm based on a blockchain," *Appl. Sci.*, vol. 12, no. 20, p. 10524, Oct. 2022, doi: 10.3390/app122010524.

[4] S. H. Alsamhi, A. V. Shvetsov, and S. V. Shvetsova, "Blockchain-empowered security and energy efficiency of drone swarm consensus for environment exploration," *IEEE Trans. Green Commun. Netw.*, vol. 7, no. 1, pp. 328–338, Mar. 2023, doi: 10.1109/TGCN.2022.3195479.

[5] S. Hafeez et al., "Blockchain-assisted UAV communication systems: A comprehensive survey," *IEEE Open J. Veh. Technol.*, vol. 4, pp. 558–580, 2023, doi: 10.1109/OJVT.2023.3295208.

[6] L. Chen, Y. Fan, J. Chen, L. Zhang, and P. Li, "ePoW energy-efficient blockchain consensus algorithm for decentralize federated learning system in resource-constrained UAV swarm," in *Proc. BlockSys 2023, Commun. Comput. Inf. Sci.*, vol. 1897, Springer, Singapore, 2024, pp. 82–98, doi: 10.1007/978-981-99-8104-5-7.

[7] K. Dev, Y. Madhwal, S. Shevelo, P. Osinenko, and Y. Yanovich, "SwarmRaft: Leveraging consensus for robust drone swarm coordination in GNSS-degraded environments," *arXiv:2508.00622*, Sep. 2025. [Online]. Available: https://arxiv.org/abs/2508.00622.

[8] Y. Xu, X. Zhu, Y. Bai, Z. Liu, L. Wang, and Z. Chen, "A dynamic lightweight blockchain sharding protocol for autonomous collaborative combat of UAV swarms in denied environments," *Sci. Rep.*, vol. 15, no. 1, p. 36420, 2025, doi: 10.1038/s41598-025-20359-1.

[9] H. Dogan and A. Setzer, "SABEC: Secure and adaptive blockchain-enabled coordination protocol for unmanned aerial vehicles (UAVs) network," in *Proc. 11th Int. Conf. Inf. Syst. Secur. Privacy (ICISSP)*, vol. 1, 2025, pp. 377–388, doi: 10.5220/0013330500003933.

[10] P. Han, A. Sui, and J. Wu, "A secure and efficient access-control scheme based on blockchain and CP-ABE for UAV swarm," *Drones*, vol. 9, no. 2, p. 148, Feb. 2025, doi: 10.3390/drones9020148.