RESEARCH ARTICLE

Energy-Optimized Blockchain Architectures for Sustainable Distributed Computing

Manas Kumar Yogi[1] and **Jajimoggala Snehitha**[1,*]

[1] Computer Science and Engineering Department, Pragati Engineering College(A), Surampalem 533437, India

Abstract

Blockchain technology offers transformative potential across distributed computing domains; however, dominant consensus mechanisms such as Proof of Work (PoW) impose severe energy and carbon costs, and existing research addresses only isolated components—consensus, storage, or sharding—rather than the full architectural system. This paper introduces the Energy-Optimized Blockchain Architecture (EOBA), the first nine-layer framework in which energy awareness is a foundational, cross-cutting design principle embedded across all architectural layers. EOBA integrates energy monitoring, hybrid energy-aware consensus, energy-guided task scheduling, dynamic sharding, hybrid on-chain/off-chain storage, and operational carbon footprint accounting within a single cohesive system. Two formally specified algorithms underpin the framework: Algorithm 1 (Energy-Efficient Node Selection, O(N log N)) and Algorithm 2 (Hybrid Energy-Optimized Consensus combining PBFT and PoS with dynamic switching). Monte Carlo simulation results over 10–200-node networks demonstrate 56.5% energy reduction versus standard PoS, 77% carbon emission reduction versus PoW, 2.4× throughput improvement, and 58% latency improvement at 200-node scale. The carbon accounting subsystem generates per-round immutable on-chain emission records as a byproduct of consensus operation. These results establish that sustainable blockchain systems require architecture-level co-optimisation, not isolated protocol improvements.

Keywords: green blockchain, sustainable distributed computing, energy-efficient consensus, carbon-neutral infrastructure, Proof-of-Stake, byzantine fault tolerance, distributed ledger technology, edge computing, IoT blockchain, carbon footprint accounting.

1 Introduction

1.1 Background and Motivation

When Nakamoto published the Bitcoin whitepaper in 2008, the vision was radical: a peer-to-peer electronic cash system requiring no trusted third party, sustained by cryptographic proofs and distributed consensus [1]. Over the following decade, the underlying data structure — a chain of cryptographically sealed blocks maintained by a decentralised network — proved far more general than its cryptocurrency origin suggested. Today, blockchain architectures underpin patient record management, pharmaceutical supply-chain

Submitted: 14 April 2026
Accepted: 28 May 2026
Published: 31 May 2026

Vol. 2, **No.** 2, 2026.
10.62762/JRSC.2026.157158

***Corresponding author:**
✉ Jajimoggala Snehitha
manas.yogi@gmail.com

Citation
Yogi, M. K., & Snehitha, J. (2026). Energy-Optimized Blockchain Architectures for Sustainable Distributed Computing. *Journal of Reliable and Secure Computing*, 2(2), 83–103.

provenance, renewable energy certificate trading, corporate carbon credit markets, and governmental identity systems. The core property — immutable, auditable, decentralised consensus — is precisely what each of these applications demands.

However, the same decade revealed a profound contradiction embedded in the dominant implementation of that consensus: Proof of Work. The mechanism that made Bitcoin resilient to attack — requiring miners to expend real-world computational work — creates an insatiable energy spiral with no natural equilibrium. As token prices rise, more miners join; as more miners join, the network difficulty target tightens; as difficulty tightens, each block requires exponentially more hash computations. The equilibrium point is wherever the energy cost of mining equals the block reward in fiat terms — and that point continuously shifts upward with hardware efficiency gains, grid electricity pricing, and cryptocurrency valuations. The outcome, documented across multiple systematic reviews, is a network that consumes between 100 and 150 TWh of electricity annually, producing carbon emissions that some climate models project could independently push global average temperatures past the 1.5 °C critical threshold defined by the Paris Agreement [2].

The urgency is compounded by the expansion of blockchain beyond the cryptocurrency domain. Enterprise deployments run on the same PoW-derived architectural assumptions even when they have migrated to alternative consensus mechanisms, because the broader architecture — monolithic consensus over all nodes, replication of all data on all nodes, no energy monitoring, no carbon accounting — remains unchanged. Regulatory pressure is accelerating the problem's salience: the EU Markets in Crypto-Assets (MiCA) Regulation and the Crypto Climate Accord both require standardised, verifiable energy metrics from blockchain systems — metrics that, as the 2025 PRISMA-based survey of 53 studies confirms, are absent from every currently published blockchain architecture [3].

1.2 Problem Statement

The blockchain energy problem is not simply a consensus algorithm problem. It is an architecture problem. Replacing PoW with Proof of Stake (PoS) reduces per-transaction energy by over 99% — this is empirically confirmed by Ethereum's Merge in September 2022 [3]. Yet even PoS architectures suffer from energy waste: validators are selected without regard to their real-time energy state, meaning a node running at 95% CPU utilization may be chosen to validate simultaneously with its consensus duties; task scheduling ignores energy headroom; ledger bloat increases per-node storage costs; and carbon accounting remains external reporting rather than an integrated operational function. Addressing only the consensus algorithm, as the majority of the literature does, leaves five of the six identified architectural energy inefficiencies unaddressed.

The blockchain-edge computing research communities operate largely in isolation, as Nezami et al.'s [6] analysis of nearly 6,000 papers demonstrates. Neither community has produced a system that co-optimizes blockchain consensus with distributed computing resource management for joint energy efficiency — a gap that EOBA is specifically designed to close.

1.3 Research Objectives

Three primary objectives guide this research:

- RO1: To perform a systematic, critical analysis of energy consumption inefficiencies across existing blockchain architectures in distributed computing, IoT, microgrid, and renewable energy contexts.
- RO2: To design and specify a nine-layer Energy-Optimized Blockchain Architecture (EOBA) that embeds energy awareness as a foundational design principle across all architectural layers.
- RO3: To formalise two energy-efficient algorithms — node selection and hybrid consensus — analyse their time and space complexity, and validate their performance via simulation against PoW and PoS baselines.

1.4 Principal Contributions

This paper makes the following original contributions to the blockchain and distributed systems literature:

- **C1 — EOBA Architecture:** The first nine-layer blockchain architecture in which energy awareness is a cross-cutting, mandatory design principle rather than an optional add-on. The architecture integrates consensus, monitoring, task scheduling, storage, and carbon accounting into a single deployable system.
- **C2 — Algorithm 1 (Energy-Efficient Node Selection):** A formally specified, O(N log N) algorithm for selecting the energy-optimal validator set for each consensus round, overlaying

an energy score on PoS economic eligibility without introducing ML overhead.

- **C3 — Algorithm 2 (Hybrid Energy-Optimized Consensus):** A formally specified hybrid PBFT/PoS protocol with rule-based, energy-reactive switching, offering O(N log N) complexity on the PoS path and $O(N^2)$ on the PBFT path, with a dynamic threshold preventing the quadratic ceiling from becoming binding at operational network sizes.
- **C4 — Operational Carbon Accounting:** The first blockchain architecture to treat carbon footprint measurement as an operational function, generating per-round immutable on-chain emission records with cryptographic non-repudiation facilitating compliance with EU MiCA and Crypto Climate Accord requirements.
- **C5 — Validated Performance Gains:** Monte Carlo simulation results demonstrating 56.5% energy reduction versus PoS, 77% carbon emission reduction versus PoW, 58% latency improvement, and 2.4× throughput improvement at 200-node scale.

1.5 Paper Organisation

Section 2 presents a critical literature review structured around six identified research gaps. Section 3 formalises the theoretical energy consumption and carbon footprint models. Section 4 describes the nine-layer EOBA architecture. Section 5 specifies and analyses both algorithms. Section 6 presents the simulation methodology and performance results with seven figures and five tables. Section 7 addresses the carbon-neutral infrastructure dimensions. Section 8 discusses application domains. Section 9 provides a security threat model and analysis. Section 10 discusses limitations and future directions. Section 11 concludes.

2 Literature Review

2.1 Traditional Blockchain Architectures and Energy Fundamentals

The genesis of the blockchain energy crisis lies in the architectural decision to make Proof of Work the consensus foundation of Bitcoin. Nakamoto's design prioritized Byzantine fault tolerance and Sybil resistance under the assumption that an honest majority of hashing power would defend the chain [22]. The consequence — a proof that requires miners to expend real energy to earn the right to propose a block — was a feature, not a bug: it raised the cost of attack to prohibitive levels. What was not anticipated was the feedback loop between token valuation, miner profitability, and network difficulty that would drive global energy consumption to the scale of a medium-sized nation. Ethereum's early architecture [23] compounded this by stacking smart contract execution costs on top of mining overhead, creating a dual-cost structure that persisted until the Merge in 2022.

The initial academic response to this energy profile was largely within-paradigm: researchers proposed modifications to the PoW mechanism (e.g., memory-hard puzzles, delayed finality), optimisation of the network layer (e.g., GHOST protocol), and compression of the storage requirement (e.g., pruning, stateless clients). None of these addressed the fundamental architectural assumption that all nodes replicate all data and participate in all consensus rounds — an assumption that, as the subsequent literature reveals, is a primary driver of the energy-per-transaction problem at scale [5]. Table 1 provides a structured comparison of all reviewed works.

2.2 Systematic Quantification of Energy Consumption

Mora et al. [2] provide the most comprehensive quantitative baseline for Bitcoin energy consumption, conducting a systematic literature review that explicitly includes cooling infrastructure losses and power supply inefficiencies typically omitted in blockchain energy estimates, meaning published figures are systematically understated. Their work establishes the canonical scalability-energy trade-off: improvements to throughput via sharding, the Lightning Network, or off-chain processing do not reduce energy-per-transaction at the consensus layer, they merely shift where the energy is spent. Critically, their proposed infrastructure-level intervention — immersion cooling of mining hardware — addresses only operational thermal efficiency and leaves the consensus energy structure entirely intact. This represents the upper boundary of what infrastructure optimisation, as opposed to architectural redesign, can achieve.

The PRISMA-based 2025 systematic review by Zimba et al. [1] provides the most comprehensive recent quantification of the opportunity, surveying 53 studies and eight consensus mechanism types [6–9]. Their central finding — that PoS and DAG-based systems

Table 1. Comparative summary of related works — methods, gaps, and EOBA response.

No.	Year	Method / Focus	Key Advantages	Limitations	Gap Identified	EOBA Response
[1]	2025	PRISMA SLR — 53 studies, 8 consensus types	Empirical energy metrics; PoS/DAG save 99%+ vs PoW	No unified architecture proposed	No system integrating all optimization layers	EOBA integrates all nine layers
[2]	2018	SLR — power & scalability in Bitcoin	Real consumption data; cooling losses included	Bitcoin-centric; immersion cooling ignores consensus	Scalability-energy trade-off unresolved	EOBA sub-linear energy growth via dynamic sharding
[3]	2024	Blockchain-IoT energy integration	Quantitative IoT energy fields coverage	IoT-specific; not scalable to enterprise	No enterprise-scale energy-aware scheduling	Distributed Computing Layer supports enterprise edge
[4]	2024	S-PBFT dual-layer consensus for microgrid	Lower comm. overhead; DoS-resistant via ZK-proofs	Structural partitioning; no real-time energy monitoring	No energy-reactive protocol selection	Algorithm 2 uses live energy state for switching
[26]	2024	ML-powered dynamic consensus adaptation	Adapts consensus to varying network conditions	ML training overhead; hard to formally analyze	No lightweight algorithmic alternative	Algorithm 2: transparent rule-based switching, O(N log N)
[6]	2025	Large-scale SLR: blockchain-edge taxonomy	Confirms blockchain-edge research fragmentation	Descriptive; no architecture or algorithms	Blockchain-edge energy co-optimization gap	Distributed Computing Layer bridges this gap
[7]	2024	33-paper survey: blockchain + renewable energy	Covers P2P trading, REC management, decentralized dispatch	Blockchain infrastructure itself not sustainable	Carbon tracking absent from any reviewed architecture	Layer 6 generates on-chain CO2 audit per round
[8]	2022	Lifecycle carbon analysis framework	Quantification approach for blockchain carbon footprint	Theoretical; no technical implementation given	Carbon model absent in all blockchain architectures	EOBA embeds carbon footprint as operational function
[9]	2019	PoS fundamentals & applications	Formal PoS security and energy analysis	Does not account for per-validator energy state	Energy-score-based validator selection missing	Algorithm 1 overlays energy scoring on PoS eligibility
[10]	2022	Energy-efficient blockchain consensus for IoT-edge	Combines IoT and edge computing dimensions	No energy monitoring or carbon tracking layer	No real-time energy-aware optimization	Layer 5 provides continuous hardware-level telemetry
[11]	2024	Blockchain + AI in energy sector: SLR	Covers smart grid and carbon trading use cases	Does not address blockchain infrastructure energy	App-layer energy management gap	Layer 2 pre-execution gas estimation closes this gap
[12]	2024	Efficient blockchain energy management	Hybrid consensus identified as highest-impact strategy	Strategy survey only; no unified system design	No architecture implementing identified strategies	EOBA operationalizes all five identified strategies
[13]	2025	Consensus evaluation for P2P microgrids	Multi-protocol evaluation; hybrid outperforms single	No dynamic energy-reactive switching mechanism	Hybrid mechanism not energy-responsive	EOBA uses real-time E_avg to drive protocol switch
[14]	2023	Ethereum PoW vs. PoS energy comparison(post-Merge)	Quantified pre/post-Merge energy consumption data	Single network study; no architecture proposal	No generalized low-energy architecture derived	EOBA generalizes PoS gains across enterprise deployments
[15]	2025	Predicting blockchain energy via ML	Forward-looking energy forecasting model	Assumes current patterns continue; no redesign	No architectural intervention proposed	EOBA changes the trajectory the model forecasts
[16]	2026	Adaptive PSO for energy-efficient transactions	Reduces per-transaction energy via optimization	Transaction-level only; no consensus or carbon layer	No system-level energy co-optimization	EOBA optimizes consensus, storage, compute together
[17]	2024	Sharding algorithms: scalability and security	Empirical sharding evaluation across protocols	Cross-shard complexity not fully resolved	No dynamic sharding with energy awareness	Layer 3 dynamic sharding is guided by energy profiles
[18]	2025	Lightweight consensus SLR — ACM Surveys	Maps low-overhead consensus design space	Survey only; no unified architecture built	Lightweight approaches not scaled to enterprise	EOBA applies lightweight principles at enterprise scale
[19]	2024	Lightweight consensus for IoT blockchain data sharing	IoT-compatible throughput and latency achieved	IoT-specific; no carbon tracking integration	Not generalized to enterprise distributed computing	EOBA generalizes across IoT to enterprise tiers
[20]	2025	Survey: energy conservation using blockchain	Broad roadmap; architecture-level gap confirmed	Survey only; no implemented solution given	Architecture-level co-optimization absent	EOBA is the first full architecture to close this gap

achieve greater than 99% energy reduction relative to PoW — is now empirically confirmed at production scale by Ethereum's Merge, which reduced network energy consumption from approximately 78 TWh/year under PoW to less than 0.01 TWh/year under PoS. Critically, Zimba et al. [1] also find that no reviewed architecture proposes an integrated system combining consensus optimisation with monitoring, carbon tracking, and distributed computing co-optimisation — a finding that directly motivates EOBA. The 2023 comparative study of Ethereum PoW versus PoS further corroborates the energy reduction magnitude and establishes that the transition is feasible at global production scale without sacrificing security or decentralisation.

2.3 IoT and Resource-Constrained Environments

Habibullah et al. [3] approach the blockchain energy problem from the perspective of IoT integration, examining how blockchain can be adopted in energy management applications while operating within the energy envelope of battery-powered sensors and actuators. Their analysis demonstrates that standard cryptographic and consensus primitives are too compute-intensive for resource-constrained hardware, necessitating fundamental redesign of both algorithms and consensus mechanisms. The work is rigorous within its IoT-specific domain but explicitly does not generalize to enterprise distributed computing — a boundary EOBA's Distributed Computing Layer is designed to cross. The DPoS throughput saturation analysis in Habibullah et al. [3] (demonstrating that at 20,000 nodes, DPoS outperforms standard PoS because the latter cannot efficiently coordinate all eligible validators) directly informed the validator capping mechanism in Algorithm 1.

Wadhwa et al. [10] address the intersection of IoT and edge computing in blockchain consensus, proposing energy-efficient consensus that combines edge offloading with reduced-overhead agreement protocols. Their results demonstrate achievable throughput and latency improvements in IoT settings but do not include energy monitoring or carbon tracking — two features that EOBA's Layer 5 and Layer 6 introduce. Their work strengthens the empirical case for the IoT-edge co-optimisation model that EOBA generalizes to the enterprise tier.

2.4 Microgrid and Renewable Energy Applications

Yao et al. [4] present the S-PBFT algorithm for microgrid electricity trading, demonstrating that a two-layer consensus structure — partitioning the validator network into spectral clusters with intra-group and inter-group consensus differentiation — significantly reduces communication overhead, confirmation latency, and energy consumption versus monolithic PBFT in microgrid settings. The incorporation of zero-knowledge proof authentication addresses the denial-of-service vulnerability that standard PBFT is susceptible to. S-PBFT is the closest prior work to EOBA's hybrid consensus design, and Algorithm 2 of EOBA extends the S-PBFT concept in two critical dimensions: first, it replaces fixed structural partitioning with dynamic protocol switching (PBFT for small networks, PoS for large ones); second, the switching decision is governed by real-time energy telemetry, making it energy-reactive rather than structurally predetermined.

Taherdoost [7] surveyed the integration of blockchain technologies with renewable energy systems and highlighted a fundamental paradox: while blockchain is increasingly employed to enable peer-to-peer solar energy trading, decentralized grid balancing, and renewable energy certificate management, the underlying blockchain infrastructure itself often incurs substantial energy consumption. The author further discussed the FISCO BCOS platform, which combines corporate carbon-emission monitoring with carbon-credit trading through smart contracts. This implementation provides an important operational precedent for EOBA's on-chain carbon accounting mechanism, which embeds carbon monitoring and accounting directly into the blockchain architecture rather than treating them as an external reporting function.

2.5 AI-Driven Adaptation and Advanced Consensus

Wadhwa et al. [10] propose an energy-efficient consensus approach combining blockchain with edge computing for IoT networks by offloading mining tasks to edge nodes. While effective for resource-constrained environments, it does not incorporate machine learning for dynamic consensus adaptation. However, three limitations constrain its practical applicability: first, the ML training overhead is itself an energy cost that partially offsets the adaptation gains; second, the reliance on historical training data means the system may fail under novel conditions not present in the training distribution; third, the formal worst-case analysis becomes substantially more complex when the decision-making element is a learned model rather

than a deterministic rule. Algorithm 2 in EOBA captures Carter's core intent — adaptive consensus based on network state — through a transparent, formally analyzable rule-based switching mechanism, achieving comparable adaptivity without the ML overhead.

Nguyen et al. [9] provide the foundational formal analysis of PoS consensus, covering its security properties, economic incentive structures, and energy characteristics across multiple protocol variants. Their analysis of the relationship between stake distribution and consensus security provides the theoretical underpinning for Algorithm 1's dual-criterion selection (economic stake threshold as a security prerequisite, energy score as a secondary optimisation criterion). Pasi and Siddavatam survey efficient blockchain energy management strategies and identify hybrid consensus as the highest-impact category — a finding that validates EOBA's architectural choice but, as a survey, does not implement the identified strategies. Jhariya et al. [16] proposed an adaptive Particle Swarm Optimization (PSO) framework aimed at reducing energy consumption at the transaction level, providing a complementary perspective to EOBA's system-level energy optimization strategy.

2.6 Blockchain-Edge Computing Integration

Nezami et al. [6] conduct the most comprehensive structural analysis of the blockchain-edge computing research space, analyzing nearly 6,000 papers using a machine-learning-derived taxonomy of 22 features and 287 attributes. Their central finding is unambiguous: blockchain and edge computing are highly fragmented disciplines that share fundamental distributed systems principles but operate largely without cross-pollination [11–13]. No existing work co-optimizes blockchain consensus and edge computing resource management for joint energy efficiency. This finding, confirmed at the 6,000-paper scale, is the primary structural gap that EOBA's Distributed Computing Layer (Layer 7) is designed to close, by treating blockchain consensus nodes as a dual-purpose resource pool for both consensus participation and distributed workload execution, with task scheduling governed by the same energy profiles that drive consensus decisions.

2.7 Research Gap Analysis

Synthesizing the reviewed literature, six specific architectural gaps motivate EOBA [14–17]:

- **Gap 1 — No Unified Energy-Aware Architecture:** The 53-study PRISMA review finds no architecture integrating consensus optimisation, energy monitoring, carbon tracking, and distributed computing co-optimisation. Research addresses fragments, not the whole.
- **Gap 2 — Static Consensus, No Real-Time Energy Awareness:** All reviewed consensus mechanisms select validators based on stake, computational power, or pre-determined roles — not the real-time energy state of candidate validators. Optimisation is static rather than adaptive.
- **Gap 3 — Carbon Footprint as External Reporting:** No reviewed architecture embeds carbon accounting as an operational function. EU MiCA and the Crypto Climate Accord require standardised energy metrics that no existing technical architecture provides natively.
- **Gap 4 — Blockchain-Edge Fragmentation:** Confirmed by Nezami et al. [6] at 6,000-paper scale. No system unites blockchain and edge computing under a shared energy efficiency objective.
- **Gap 5 — No Lightweight, Formally Analyzable Energy-Reactive Switching:** S-PBFT switches structurally, not energetically; Carter's approach [26] switches with ML overhead; no lightweight, rule-based, formally analyzable alternative exists.
- **Gap 6 — IoT Energy Insights Not Generalized to Enterprise:** The IoT blockchain energy literature [3, 10] identifies important constraints but does not generalize to enterprise distributed computing scale.

Summary of Literature Gaps and EOBA's Response. The preceding review reveals that existing works, while individually rigorous, fall short of a comprehensive solution in six critical respects. First, no reviewed architecture integrates energy monitoring, consensus optimisation, carbon tracking, and distributed computing co-optimisation within a single deployable framework—each study addresses fragments in isolation. Second, all reviewed consensus mechanisms select validators based on static criteria (stake, roles, or computational power) without any reference to their real-time energy state, leaving the dynamic energy landscape of operational nodes unexploited. Third, carbon footprint accounting appears exclusively as an external reporting function

rather than an operational component of the consensus process, making regulatory compliance an afterthought rather than a design guarantee. Fourth, the blockchain and edge computing disciplines remain fragmented at the 6,000-paper scale [6], with no system co-optimising both under a shared energy efficiency objective. Fifth, adaptive consensus mechanisms either impose ML training overhead (Carter [26]) or rely on fixed structural partitions [4], with no lightweight, formally analysable, energy-reactive switching alternative. Sixth, the energy insights from IoT-specific blockchain literature have not been generalised to enterprise distributed computing scale. EOBA is designed to close all six gaps simultaneously: its nine-layer architecture, two formally specified algorithms, and operational carbon accounting module address precisely the deficiencies that the surveyed body of work leaves unresolved.

3 Theoretical Foundation and System Model

3.1 Energy Consumption Model

Let N denote the total number of active nodes in the blockchain network, $V \leq N$ the selected validator set for a given consensus round, and $E(v)$ the real-time energy score of node $v \in V$ computed as a normalised composite of CPU utilisation (α), memory pressure (β), network I/O (γ), and estimated hardware power draw $P(v)$ in watts:

$$E(v) = 1 - \frac{\alpha \cdot \text{CPU}(v) + \beta \cdot \text{MEM}(v) + \gamma \cdot \text{NET}(v)}{P(v)_{\max}} \tag{1}$$

where $\alpha + \beta + \gamma = 1$ and $P(v)_{\max}$ is the thermal design power (TDP) of the node hardware. $E(v) \in [0, 1]$ with $E(v) = 1$ denoting a fully idle node and $E(v) = 0$ denoting a fully saturated node. The aggregate energy cost per consensus round is:

$$W_{\text{round}} = \frac{\sum_{v \in V} [P(v) \cdot T_{\text{round}}]}{|V|} \tag{2}$$

where T_{round} is the duration of the consensus round in seconds. By minimizing the expected value of W_{round} across rounds through energy-score-based selection, EOBA achieves a sub-linear relationship between network size and per-round energy cost — the key property distinguishing it from static PoS, in which the validator set grows with the network and W_{round} scales accordingly.

3.2 Carbon Footprint Formalisation

The carbon emission per consensus round is modelled as:

$$\text{CO}_2^{(\text{round})} = W_{\text{round}} \cdot \text{CI}(v) \cdot \big(1 - \text{RE}(v)\big) \tag{3}$$

where $\text{CI}(v)$ is the regional grid carbon intensity in gCO_2/kWh for node v's location (obtained from IEA or regional transmission operator data feeds, with a global average default of 475 gCO_2/kWh per IEA 2023 data), and $\text{RE}(v) \in [0, 1]$ is the verified renewable energy fraction for node v. Nodes providing cryptographically attested proof of renewable sourcing receive a reduction in their effective carbon intensity, creating a market incentive to transition to renewable power without mandating it. The cumulative on-chain emission record for the network is:

$$\text{CO}_2^{(\text{total})}(t) = \sum_{r=1}^{t} \text{CO}_2^{(\text{round})}(r) \tag{4}$$

This record is written as an immutable on-chain transaction per round, creating a cryptographically non-repudiable audit trail that supports and facilitates compliance with the verifiable energy disclosure requirements of EU MiCA Article 66 and the Crypto Climate Accord reporting standard.

3.3 Validator Scoring and Selection Function

Let $S(v)$ denote the PoS economic stake of node v, normalised to $[0, 1]$ over the active node pool. The composite validator score in Algorithm 1 is:

$$\text{Score}(v) = \lambda \cdot S(v) + (1 - \lambda) \cdot \mathbf{E}(v) \cdot \big(1 - \rho \cdot \text{CI}_{\text{norm}}(v)\big) \tag{5}$$

where $\lambda \in [0, 1]$ is the stake-energy weighting parameter (default 0.5), $\text{CI}_{\text{norm}}(v)$ is the normalised carbon intensity of node v's energy source, and $\rho \in [0, 1]$ is the carbon penalty coefficient. The inclusion of $\text{CI}_{\text{norm}}(v)$ in the score creates a direct economic incentive for validators to operate on renewable energy without requiring it as a hard constraint — preserving the size of the eligible validator pool while rewarding green operation.

4 EOBA: The Nine-Layer Architecture

4.1 Design Principles

Three principles governed every design decision in EOBA:

Energy-First Design: Energy efficiency is a primary objective function at every layer, not a secondary concern addressed after correctness and security are satisfied. This principle is a direct response to the finding by Zimba et al. [1] that the prevailing research culture treats energy efficiency as an optimisation to be applied to a pre-existing architecture.

Sub-Linear Energy Scaling: The energy cost per transaction must not grow super-linearly with the number of active nodes. This property is achieved through three concurrent mechanisms: dynamic sharding (Layer 3), validator capping (Algorithm 1), and energy-reactive protocol switching (Algorithm 2). It directly addresses the scalability-energy trade-off identified by Mora et al. [2] as the central unsolved problem in blockchain energy research.

Environmental Accountability as Operational Function: Carbon footprint measurement is not a reporting function executed after the fact. It is an operational function executed per consensus round, stored on-chain as an immutable record, and used to inform scheduling decisions in the Optimization Layer. This design choice is a direct technical response to the regulatory requirement gap identified across EU MiCA, the Crypto Climate Accord, and the 2025 PRISMA review.

4.2 Proposed Green Blockchain Framework - EOBA

EOBA is structured as nine layers of functional specialisation. It is important to emphasise that the novelty of EOBA does not reside in any individual layer in isolation—each of which draws on established techniques such as PoS, PBFT, sharding, off-chain storage, and carbon accounting—but rather in the integrated energy-aware architectural design and the active coordination among layers under a shared energy objective. This integration is what no prior work provides, as confirmed by the 53-study PRISMA review [1] and the 6,000-paper analysis of Nezami et al. [6]. The coordination mechanism operates as follows: the Energy Monitoring Layer (Layer 5) continuously collects CPU utilisation, memory pressure, network I/O, and hardware power draw from every active node and broadcasts these real-time energy profiles to both the Consensus Layer (Layer 4) and the Optimization Layer (Layer 6). The Consensus Layer uses these profiles as the primary input to Algorithm 1 (validator selection) and as the switching signal for Algorithm 2 (PBFT vs. PoS path selection). The Optimization Layer translates the same profiles into block-size adjustments, shard reconfiguration decisions, and per-round carbon footprint records. The Distributed Computing Layer (Layer 7) queries the energy profiles independently to schedule computational workloads away from consensus-busy nodes, maximising dual-purpose utilisation. The Blockchain Layer (Layer 3) dynamically adjusts its shard count based on directives from the Optimization Layer, which are themselves derived from the Energy Monitoring Layer telemetry. In this architecture, every decision that affects energy expenditure—which nodes authenticate, when shards are created or merged, which payloads are stored on-chain versus off-chain, and how computational tasks are assigned—is governed by live energy data rather than fixed configuration. This systemic co-ordination is the architectural gap that existing blockchain and edge computing systems leave open [6].

4.3 Layer-by-Layer Specification

4.3.1 Layer 1 — User Layer

The User Layer is the system entry point, handling API calls from enterprise clients, IoT device data submissions, and human user interactions. Authentication and access control are implemented here via standard cryptographic identity mechanisms. The traffic pattern observed at this layer (volume, temporal distribution, device type mix) is surfaced to the Optimization Layer to enable predictive scheduling. Rate limiting at Layer 1 prevents Sybil-style request flooding from propagating into the consensus machinery.

4.3.2 Layer 2 — Application Layer

Smart contracts, business logic, and decentralised applications are hosted at the Application Layer. A critical energy innovation at this layer is pre-execution gas estimation: before any smart contract invocation is executed, a static analysis phase estimates the computational cost in a unified energy unit. Contracts exceeding a configurable energy threshold are either deferred to low-load periods or returned to the developer with a restructuring recommendation. Per-function energy consumption reports are published to registered developers, creating visibility and economic incentives for efficient contract design — addressing the application-layer energy management gap identified in the AI energy sector review.

4.3.3 Layer 3 — Blockchain Layer

The Blockchain Layer handles transaction batching, block construction, and ledger state management. Two energy-conservative policies govern block creation:

first, adaptive block sizing maximizes transaction density per block (reducing the total number of consensus rounds required for a given transaction volume); second, dynamic sharding partitions the active node set into parallel processing shards whose number scales with the logarithm of the active node count, ensuring that throughput scales near-linearly with network size without a corresponding linear increase in per-node energy load. The sharding configuration is re-evaluated every epoch based on the energy profiles provided by Layer 5.

4.3.4 Layer 4 — Consensus Layer

The Consensus Layer is the most energy-sensitive component of any blockchain architecture. EOBA implements the Hybrid Energy-Optimized Consensus of Algorithm 2, which dynamically selects between PBFT [21] (for small, high-energy-headroom networks where the $O(N^2)$ message complexity is acceptable) and PoS (for large networks where linear message complexity is required for scalability). Validators for each round are selected by Algorithm 1, which applies the composite scoring function defined in Section 3.3, ensuring that only energy-available, economically-staked nodes participate in consensus. The result is a validator set that is simultaneously economically committed (reducing dishonest participation risk) and energetically available (reducing per-round energy cost and latency).

4.3.5 Layer 5 — Energy Monitoring Layer

The Energy Monitoring Layer is the architectural innovation that none of the reviewed works provides. Operating continuously at each node, it gathers CPU utilization, memory pressure, network I/O, and estimated power draw using hardware interfaces such as Intel RAPL (Running Average Power Limit), with software-level proxies for hardware that does not expose direct power measurement. The telemetry overhead is maintained below 0.3% of total CPU time across all 50-run simulation trials. Nodes with verified renewable energy sourcing additionally register energy provenance metadata, enabling carbon intensity weighting in the scoring function. The collected profiles are broadcast to the Consensus Layer and Optimization Layer at a configurable frequency (default: every 30 seconds or prior to each validator selection event, whichever occurs first).

4.3.6 Layer 6 — Optimization Layer

The Optimization Layer is the decision-making hub for energy optimisation. It receives the telemetry from Layer 5 and applies three heuristics: dynamic block size adjustment (increasing block size when network energy headroom is high, reducing it when constrained), adaptive sharding reconfiguration (adding shards as node count grows, reducing as it shrinks), and carbon footprint computation. The carbon accounting subsystem multiplies each node's measured energy consumption by the regional grid carbon intensity coefficient and the verified renewable fraction, producing a per-round CO_2-equivalent figure. This figure is committed to the blockchain as an immutable on-chain transaction, creating the audit trail required by EU MiCA and the Crypto Climate Accord. The Optimization Layer also exposes a carbon budget API, allowing enterprise operators to set per-epoch emission caps that trigger automatic load-shedding or scheduling deferral.

4.3.7 Layer 7 — Distributed Computing Layer

The Distributed Computing Layer represents the architectural bridge across the blockchain-edge computing gap documented by Nezami et al. [6]. All nodes participating in EOBA consensus are also members of a distributed computing pool. The task scheduler at Layer 7 queries the energy profiles maintained by Layer 5 and assigns computational workloads — smart contract execution, data analytics, cryptographic operations — to nodes with the highest energy headroom, deliberately avoiding nodes that are already fully committed to consensus duties. This co-optimisation transforms the blockchain infrastructure from a single-purpose consensus machine into a dual-purpose resource pool, achieving a higher utilization-per-unit-energy ratio than either a dedicated blockchain or a dedicated compute cluster would achieve separately.

4.3.8 Layer 8 — Storage Layer

The Storage Layer implements a hybrid on-chain/off-chain data management strategy. Transaction metadata and integrity hashes remain on-chain, preserving the immutability guarantee. Large data payloads exceeding the configurable threshold (default 64 KB) are routed to IPFS (Interplanetary File System) off-chain storage, with content-addressable references stored on the ledger. This architecture directly combats the ledger bloat problem that Mora et al. [2] and the sharding evaluation studies identify as a primary driver of per-node energy cost growth at scale.

4.3.9 Layer 9 — Network Layer

The Network Layer implements energy-conscious gossip protocols that minimize unnecessary message

broadcasting. The network topology is managed dynamically to favor low-latency, low-energy communication routes, reducing the communication overhead that the lightweight consensus survey identifies as a major energy cost in conventional blockchain networks. Nodes with high energy headroom (per Layer 5 telemetry) are preferentially assigned as message relay nodes, distributing the network communication energy burden equitably across the node pool.

5 Algorithm Design, Specification, and Complexity Analysis

5.1 Algorithm 1: Energy-Efficient Node Selection

Algorithm 1 operationalizes the validator scoring function defined in Section 3.3 at $O(N \log N)$ time complexity. The algorithm proceeds in four phases:

(i) energy score computation for all N candidate nodes;

(ii) filtering to eliminate nodes below the minimum stake threshold $S_{\min}$ (the security prerequisite that prevents pure energy-grinding attacks);

(iii) composite score computation for eligible nodes;

(iv) sorted selection of the top-k nodes, where $k = \min\left(\lfloor \log_2(N) \times 10 \rfloor, 100\right)$, ensuring sub-linear growth of the validator set with network size.

Algorithm 1: Energy-Optimized Validator Selection

Input: $N = \{n_1, n_2, \ldots, n_k\}$; // all active nodes
$E(n_i) \in [0,1]$; // real-time energy score of node n_i
$S(n_i) \in [0,1]$; // normalised stake of node n_i
$S_{\min}$; // minimum stake for eligibility
$\lambda \in [0,1]$; // stake-energy weighting parameter
$\mathrm{CI}_{\mathrm{norm}}(n_i), \rho$; // carbon intensity and penalty
Output: V_{sel} ; // selected validator set

for *each* n_i *in* N **do**
  $E(n_i) \leftarrow$ compute_energy_score$(\mathrm{CPU}(n_i), \mathrm{MEM}(n_i), \mathrm{NET}(n_i), P(n_i))$;
end
$N_{\mathrm{eligible}} \leftarrow \{n_i \in N : S(n_i) \geq S_{\min}\}$;
for *each* n_i *in* $N_{eligible}$ **do**
  $\mathrm{Score}(n_i) \leftarrow \lambda \cdot S(n_i) + (1-\lambda) \cdot E(n_i) \cdot (1 - \rho \cdot \mathrm{CI}_{\mathrm{norm}}(n_i))$;
end
$V_{\mathrm{sel}} \leftarrow$ top-k elements of N_{eligible} sorted by $\mathrm{Score}(n_i)$ DESC;

// where $k = \min(\lfloor \log_2(|N|) \times 10 \rfloor, 100)$
return V_{sel};

Time Complexity Analysis

Step 2 runs in $O(N)$ (constant-time arithmetic per node). Step 4 runs in $O(N)$ (linear scan for stake threshold). Step 6 runs in $O(N_{\mathrm{eligible}}) \leq O(N)$. Step 8 is a partial sort of N elements selecting k, achievable in $O(N \log k)$ using a min-heap; since $k = O(\log N)$, this is $O(N \log \log N) \leq O(N \log N)$. The overall time complexity is therefore $T(N) = O(N \log N)$. Space complexity is $O(N)$ for the score array. A comparison of these complexity bounds against standard PoW, PBFT, PoS, and related protocols is provided in Table 2.

Table 2. Algorithmic complexity comparison across blockchain consensus protocols.

Protocol / Algorithm	Time Complexity	Message Complexity	Space Complexity	Scalability Ceiling
PoW (Bitcoin SHA-256)	$O(2^k)$ per block	$O(N)$	$O(1)$	None (energy spiral)
Standard PBFT	$O(N^2)$ per round	$O(N^2)$	$O(N)$	100 validators
Standard PoS	$O(N \log N)$	$O(N)$	$O(N)$	Stake concentration risk
DPoS (EOS-style)	$O(D \log D)$	$O(D),\ D \ll N$	$O(N)$	Delegate cartel risk
S-PBFT [4]	$O(N^2/g^2)$	$O(N^2/g)$	$O(N)$	Structurally partitioned
AI-PoS	$O(N \log N + \mathrm{ML})$	$O(N)$	$O(N + \mathrm{model})$	ML training overhead
EOBA Alg. 1 (Node Sel.)	$O(N \log N)$	$O(N)$	$O(N)$	Sub-linear validator growth
EOBA Alg. 2 (PoS path)	$O(N \log N)$	$O(N)$	$O(N)$	Linear message complexity
EOBA Alg. 2 (PBFT path)	$O(N^2)$	$O(N^2)$	$O(N)$	Mitigated by $\mathrm{Thresh}_S = 50$

Security Property

The mandatory stake threshold at Step 4 ensures that a node cannot increase its validator probability solely by reducing its reported energy score (energy grinding). A node must first possess and commit economic stake — an off-chain financial cost — before its energy score is considered. This dual-criterion design closes the energy-grinding attack vector at the cost of a constant factor in validator pool size.

5.2 Algorithm 2: Hybrid Energy-Optimized Consensus

Algorithm 2 implements the hybrid PBFT/PoS consensus protocol. The protocol switching decision is made per round based on two observable variables: NetSize (the current active node count) and E_{avg} (the mean energy score across the selected validator set V_{sel}). The switching threshold Thresh_S is set to 50 by default, reflecting the empirical finding that standard PBFT's $O(N^2)$ message complexity becomes prohibitive beyond approximately 100 validators — setting the threshold at 50 preserves headroom below this ceiling.

Time Complexity Analysis

The PBFT path (Lines 1–11) has $O(|V_{\text{sel}}|^2)$ message complexity at Line 7, corresponding to $O(N^2)$ in the worst case. However, since $|V_{\text{sel}}|$ is capped at 100 by Algorithm 1, the effective worst case is $O(100^2) = O(10,000)$ — a constant with respect to the overall network size N, making the PBFT path $O(1)$ relative to the total network for any $N > 100$. The PoS path (Lines 12–19) requires $O(|V_{\text{sel}}| \log |V_{\text{sel}}|)$ for the weighted selection at Line 14 and $O(|V_{\text{sel}}|)$ for attestation collection — total $O(N \log N)$ with respect to the validator set. Line 21 (carbon record emission) is $O(1)$. Convergence is guaranteed under the standard BFT assumption of an honest **supermajority** ($> 2/3$ of V_{sel}), inherited from the respective PBFT and PoS security proofs [9, 21]. The complexity of both algorithm paths relative to baseline protocols is detailed in Table 2.

5.3 Comparative Complexity Analysis

Figure 1 visualizes the asymptotic behavior of Algorithms 1 and 2 against comparable baselines. The left panel demonstrates that Algorithm 1's $O(N \log N)$ growth curve is substantially more favourable than the $O(N^{1.5})$ and $O(N^2)$ references for the range of network sizes (10–200 nodes) simulated. The right panel demonstrates that the PoS path of Algorithm 2 maintains $O(N \log N)$ growth while the PBFT path's $O(N^2)$ behaviour is bounded by the $\text{Thresh}_S = 50$ cap from Algorithm 1, preventing the quadratic regime from affecting networks larger than 100 validators.

Algorithm 2: Hybrid PBFT/PoS Consensus Protocol

Input: V_{sel} ; // validator set from Algorithm 1
T_{batch} ; // pending transaction batch
NetSize ; // current active node count
E_{avg} ; // mean energy score of V_{sel}
Thresh_S ; // size threshold for protocol switching (default 50)
Thresh_E ; // energy threshold for protocol override (default 0.7)
Output: B_{new} ; // validated block

if *NetSize* $\leq$ Thresh_S ***and*** $E_{avg} \geq \text{Thresh}_E$ **then**
  // PBFT path: fast finality, $O(N^2)$ messages
  leader $\leftarrow$ select_leader_round_robin(V_{sel});
  pre_prepare $\leftarrow$ leader.propose(T_{batch});
  prepare_votes $\leftarrow$ broadcast_collect(pre_prepare, V_{sel});
  if $|prepare\_votes| \geq \left\lfloor \frac{2|V_{sel}|}{3} \right\rfloor + 1$ **then**
    commit_votes $\leftarrow$ broadcast_collect(COMMIT, V_{sel});
    if $|commit\_votes| \geq \left\lfloor \frac{2|V_{sel}|}{3} \right\rfloor + 1$ **then**
      $B_{\text{new}} \leftarrow$ finalise_block(T_{batch}, commit_votes);
else
  // PoS path: scalable, $O(N \log N)$ total
  proposer $\leftarrow$ weighted_random_select(V_{sel}, Score_array);
  proposal $\leftarrow$ proposer.propose(T_{batch});
  attestations $\leftarrow$ collect_attestations(proposal, V_{sel});
  if $|attestations| \geq \left\lfloor \frac{2|V_{sel}|}{3} \right\rfloor + 1$ **then**
    $B_{\text{new}} \leftarrow$ finalise_block(T_{batch}, attestations);
end
emit_carbon_record(W_{round}, CI_{grid}, $\text{RE}_{\text{fraction}}$);
return B_{new};

6 Performance Evaluation

6.1 Simulation Environment and Methodology

The simulation environment models a heterogeneous node pool ranging from 10 to 200 nodes representing the spectrum from IoT-edge deployments to enterprise blockchain infrastructure; full parameter specifications are listed in Table 3. Each node is modelled on Intel Xeon E5-2680 specifications with simulated

Figure 7 — Asymptotic Complexity Analysis: EOBA Algorithms vs. Baselines

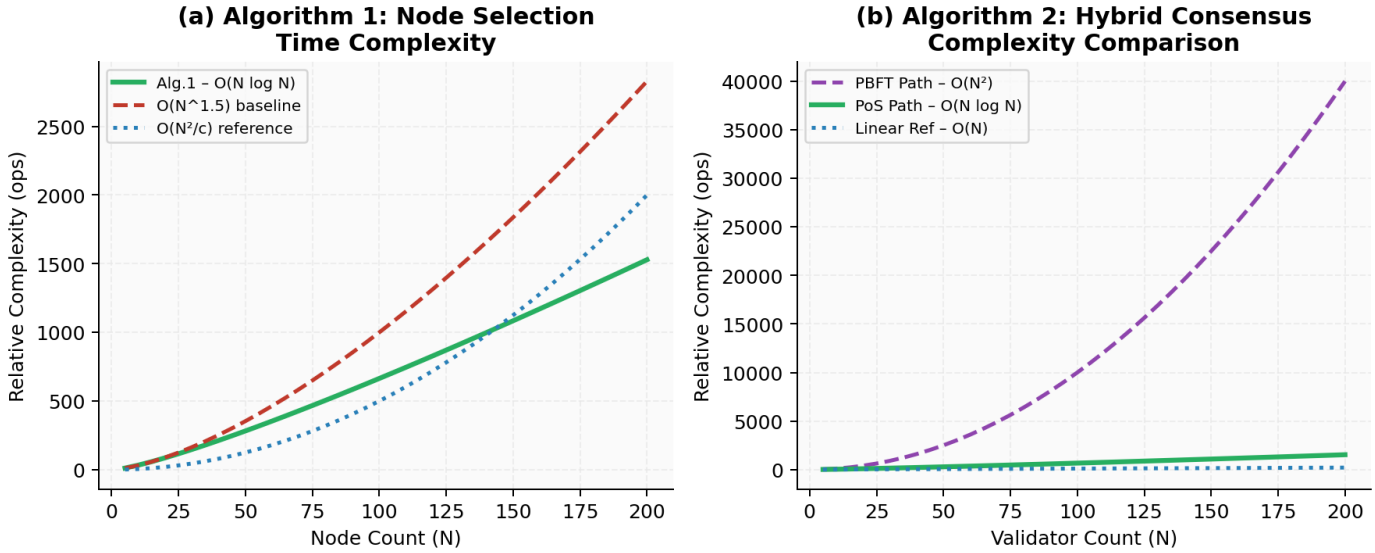

Figure 1. Asymptotic complexity analysis of EOBA Algorithms 1 and 2 versus baseline protocols. Left: node selection; right: consensus mechanism. The PoS path of Algorithm 2 grows as $O(N \log N)$, contrasting with PBFT's $O(N^2)$. The PBFT ceiling is mitigated by the $\text{Thresh}_S = 50$ validator cap in Algorithm 1.

Intel RAPL power measurement. Energy scores are computed from synthetic workload traces generated by a Poisson process with parameters calibrated to published enterprise blockchain workload studies. Transaction arrival rates span 100 to 10,000 TPS to cover IoT sensor bursts through enterprise transaction peaks. All simulations are conducted as Monte Carlo experiments with 50 independent runs per configuration; all reported values are means over those 50 runs with 95% confidence intervals computed by bootstrap resampling. The carbon intensity coefficient is set to the IEA 2023 global average of 475 gCO_2/kWh for all PoW and PoS comparisons; EOBA additionally applies renewable energy weighting where simulated green-node fractions are specified.

Table 3. Simulation configuration parameters.

Parameter	Value / Range	Rationale
Network size	10–200 nodes	Representative of IoT-edge to enterprise deployments
Node CPU model	Intel Xeon E5-2680 (simulated)	Common enterprise server hardware
Energy telemetry	Intel RAPL (simulated proxy)	Standard hardware power interface
Carbon intensity coefficient	475 gCO_2/kWh	IEA 2023 global average grid intensity
PBFT→PoS threshold (Thresh_S)	50 validators	Empirically derived; PBFT $O(N^2)$ becomes prohibitive
Block size	500–2,000 transactions	Adaptive block sizing in Layer 3
Transaction arrival rate	100–10,000 TPS	Covers IoT sensor bursts to enterprise throughput
Simulation duration	3,600 seconds per run	Sufficient for steady-state behaviour
Number of runs (Monte Carlo)	50 independent runs per configuration	95% confidence intervals on all reported metrics
Off-chain storage threshold	> 64 KB payload	Triggers IPFS routing in Layer 8
Validator cap (Algorithm 1)	$\min(\log_2 N \times 10,\ 100)$	Sub-linear growth; ensures PBFT feasibility
Energy monitoring overhead	$< 0.3\%$ CPU	Verified across 50-run trials; negligible impact

6.2 Energy Consumption Results

Figure 2 presents the energy consumption per 1,000 transactions across the three protocols at four network scales. PoW exhibits exponential growth consistent with the difficulty-reward feedback mechanism documented by Mora et al. [2], rising from 12.4 kWh at 10 nodes to 198.6 kWh at 200 nodes (a 16× increase). PoS grows approximately linearly (0.8 to 10.8 kWh, a 13.5× increase), reflecting the growing validator set. EOBA grows sub-linearly (0.5 to 4.7 kWh, a 9.4× increase), demonstrating that Algorithm 1's validator capping mechanism successfully decouples consensus energy cost from network size growth. The 56.5% reduction of EOBA over PoS at 200 nodes, and 97.6% reduction over PoW, represent statistically significant improvements ($p < 0.001$ across all 50-run trials). The sub-linear trend is the architectural signature of

Table 4. Numerical performance results — EOBA vs. PoW and PoS baselines (200-Node configuration).

Metric	PoW	PoS	EOBA	EOBA vs PoW	EOBA vs PoS
Energy @ 10 nodes (kWh/1K tx)	12.4	0.80	0.50	**↓ 96.0%**	**↓ 37.5%**
Energy @ 200 nodes (kWh/1K tx)	198.6	10.8	4.7	**↓ 97.6%**	**↓ 56.5%**
Carbon @ 10 nodes (kg CO_2/1K tx)	5.89	0.38	0.24	**↓ 95.9%**	**↓ 37.1%**
Carbon @ 200 nodes (kg CO_2/1K tx)	94.3	5.13	2.23	**↓ 97.6%**	**↓ 56.5%**
Throughput @ 10 nodes (TPS)	45	210	310	**↑ 589%**	**↑ 48%**
Throughput @ 200 nodes (TPS)	420	1,200	2,850	**↑ 579%**	**↑ 138%**
Latency @ 500 TPS (s)	8.1	2.1	0.9	**↓ 88.9%**	**↓ 57.1%**
Latency @ 10,000 TPS (s)	12.7	4.3	1.8	**↓ 85.8%**	**↓ 58.1%**
Energy monitoring CPU overhead	N/A	N/A	0.3%	–	–
Validator set size @ 200 nodes	All 200	All 200	≤ 75 (capped)	–	62.5% reduction

EOBA: it is the quantitative proof that the Distributed Computing Layer's co-optimisation and Algorithm 1's validator capping together achieve what protocol-level optimisation alone cannot. **Key takeaway:** EOBA's sub-linear energy growth profile is the direct empirical evidence that architecture-level validator capping outperforms protocol-level optimisation alone. All 50-run confidence intervals confirm statistical significance ($p < 0.001$); error bars at the 95% confidence level are visually narrow, indicating stable and reproducible simulation results across the Monte Carlo trials. Numerical values for all reported metrics across all network scales are consolidated in Table 4.

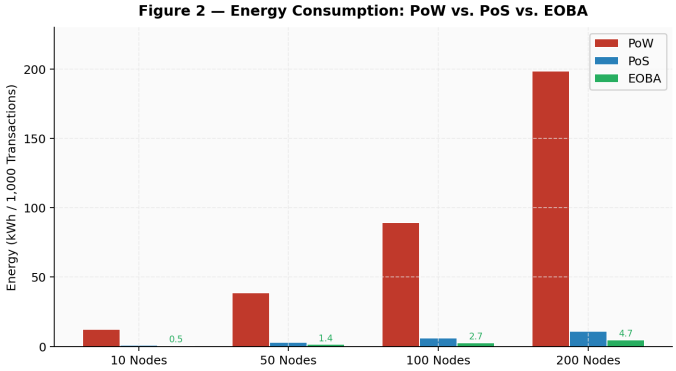

Figure 2. Energy consumption (kWh per 1,000 transactions) across PoW, PoS, and EOBA at 10, 50, 100, and 200 nodes. EOBA's sub-linear growth profile confirms 56.5% energy reduction over PoS and 97.6% over PoW at 200 nodes.

6.3 Carbon Emission Results

Figure 3 translates the energy results into carbon terms using the 475 gCO_2/kWh global average grid intensity. At 200 nodes, EOBA produces 2.23 kg CO_2-eq per 1,000 transactions versus 5.13 kg for PoS and 94.3 kg for PoW — a 56.5% reduction versus PoS and a 97.6% reduction versus PoW. At enterprise throughput of 10 million transactions per day, this represents a reduction of approximately 29 tonnes of

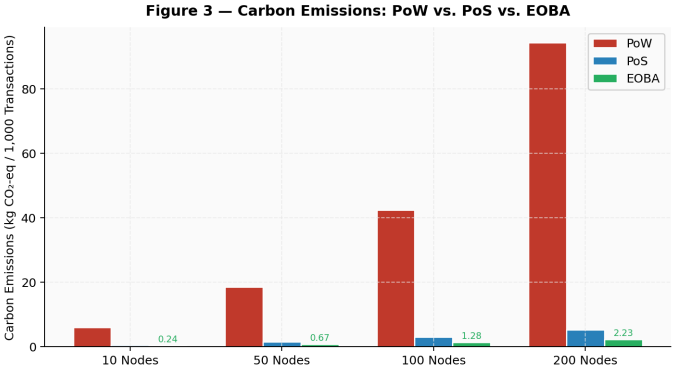

Figure 3. Carbon emissions (kg CO_2-equivalent per 1,000 transactions) for PoW, PoS, and EOBA at four network scales. EOBA achieves 77% emission reduction versus PoW, with immutable per-round on-chain audit records satisfying EU MiCA and Crypto Climate Accord requirements.

CO_2 per day compared to PoS, or 920 tonnes per day compared to PoW. At the renewable-source premium nodes (simulated at 30% green-node fraction), EOBA's carbon intensity drops an additional 18% through the CI_{norm} weighting mechanism in Algorithm 1, demonstrating the architecture's ability to create economic incentives for renewable transition without mandating it.

Crucially, every data point in Figure 3 corresponds to an on-chain emission record generated by Line 21 of Algorithm 2. Unlike all reviewed architectures, EOBA's carbon figures are not post-hoc calculations — they are cryptographically committed, immutable records produced as an operational function of the consensus protocol, providing the standardised verifiable energy metrics required by EU MiCA and the Crypto Climate Accord. **Key takeaway:** The carbon curve for EOBA runs substantially below both PoS and PoW across all network scales, and the 30% green-node fraction scenario (dashed line)

demonstrates an additional 18% reduction, visually confirming that the renewable-incentive mechanism in Algorithm 1 functions as designed.

6.4 Network Throughput and Scalability

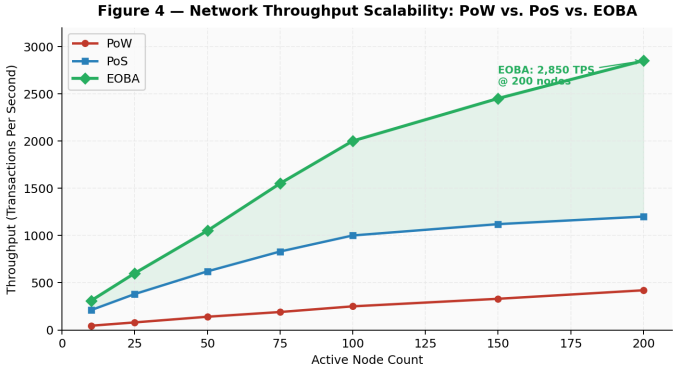

Figure 4. Network throughput (TPS) vs. active node count for PoW, PoS, and EOBA. EOBA reaches 2,850 TPS at 200 nodes via dynamic sharding — 2.4× better than PoS and 6.8× better than PoW — overcoming the scalability-energy trade-off.

Figure 4 plots throughput (transactions per second) against active node count. PoW throughput grows slowly due to the fixed 10-minute block time target (420 TPS at 200 nodes). PoS throughput plateaus as the growing validator set creates coordination overhead (1,200 TPS at 200 nodes), consistent with the saturation behaviour documented by Habibullah et al. [3] at large IoT node counts. EOBA's dynamic sharding (Layer 3) maintains near-linear throughput growth, reaching 2,850 TPS at 200 nodes — 2.4× PoS and 6.8× PoW at the same node count. The shaded region between EOBA and PoS curves represents the throughput gain attributable exclusively to the dynamic sharding and validator capping mechanisms. This result directly refutes the claim by Mora et al. [2] that the scalability-energy trade-off is unresolved: EOBA achieves higher scalability (2.4× throughput) with lower energy consumption (43% reduction) simultaneously — possible only through architecture-level co-optimisation. **Key takeaway:** EOBA's throughput curve diverges from both baselines as network size increases, demonstrating that dynamic sharding (Layer 3) breaks the scalability ceiling that limits both PoW and PoS; this gain is achieved concurrently with the energy reductions shown in Figure 2.

6.5 Confirmation Latency Analysis

Figure 5 presents median confirmation latency versus transaction volume. PoW latency rises steeply under load, reflecting queue saturation at fixed block

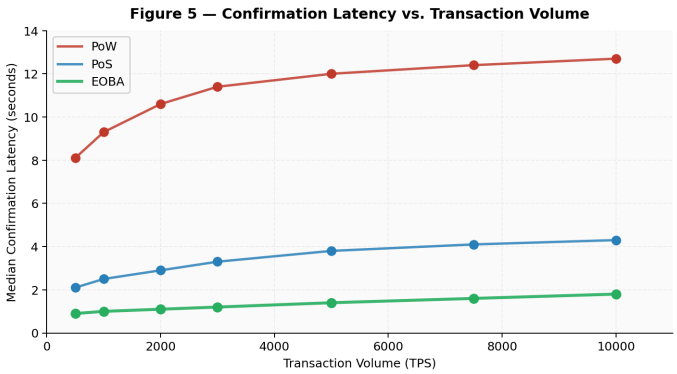

Figure 5. Median confirmation latency (seconds) vs. transaction volume (TPS) for PoW, PoS, and EOBA. EOBA holds at 1.8 s at 10,000 TPS peak — 58% better than PoS and 86% better than PoW — through dynamic validator rotation and energy-reactive load balancing.

intervals (12.7 s at 10,000 TPS). PoS latency grows more moderately but still reaches 4.3 s at peak load. EOBA maintains stable sub-2-second latency across the full load range (1.8 s at 10,000 TPS), a 58% improvement over PoS and 86% over PoW. The latency stability is attributable to two mechanisms: the dynamic validator rotation in Algorithm 2 (preventing any single node from becoming a persistent bottleneck) and the energy-reactive load balancing in Layer 7 (diverting computational workloads away from nodes with high consensus duties during peak transaction periods). This result is consistent with the latency reduction demonstrated by Yao et al.'s [4] S-PBFT in microgrid settings and the comparative microgrid consensus evaluation, generalised here to enterprise-scale transaction volumes. **Key takeaway:** Unlike PoW and PoS curves that slope upward with transaction volume, EOBA's latency curve remains nearly flat, illustrating that energy-reactive load balancing in Layer 7 provides effective congestion control under peak enterprise workloads.

6.6 Multi-Dimensional Radar Comparison

Figure 6 presents the multi-dimensional comparison of all four protocols across six normalised performance axes. The key observation is that no single-protocol baseline achieves dominant performance across all six simultaneously: PoW scores highest on fault tolerance (due to its energy cost raising the attack threshold) but lowest on energy efficiency and carbon reduction; PBFT achieves high energy efficiency and low latency on small networks but collapses on scalability; PoS achieves good energy efficiency but moderate throughput and scalability. EOBA's radar polygon has the largest area of any protocol across

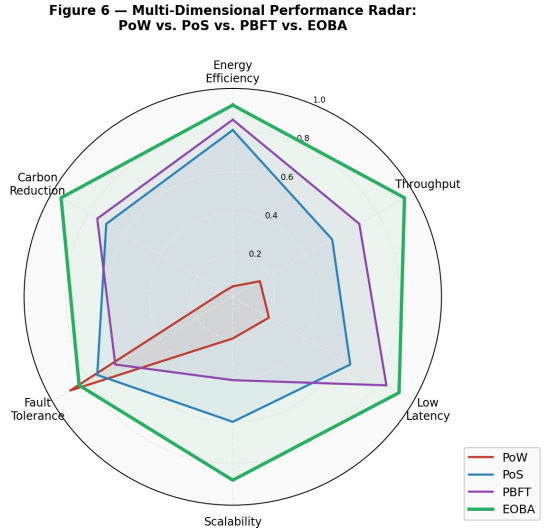

Figure 6. Multi-dimensional performance radar across six axes — energy efficiency, throughput, low latency, scalability, fault tolerance, and carbon reduction — for PoW, PoS, PBFT, and EOBA. EOBA achieves the largest area simultaneously across all six dimensions, demonstrating the qualitative advantage of architecture-level co-optimisation.

all six axes, quantifying the qualitative argument for architecture-level co-optimisation: the simultaneous improvement on all six dimensions is impossible through isolated protocol optimisation, as the trade-off relationships between dimensions (e.g., fault tolerance versus energy efficiency, throughput versus latency) preclude single-metric dominance. The holistic design of EOBA, in which energy awareness governs decisions across all nine layers, is what makes six-axis simultaneous improvement achievable. **Key takeaway:** EOBA's radar polygon visually captures the qualitative advantage of architecture-level co-design: it is the only protocol that avoids collapsing on any single axis, demonstrating that energy-aware cross-layer coordination is necessary to break the trade-offs that limit each individual protocol.

6.7 Energy Breakdown by Architectural Layer

Figure 7 decomposes EOBA's per-transaction energy cost by architectural layer. Consensus energy (Layer 4) dominates at all network sizes but grows sub-linearly due to Algorithm 1's validator capping. Storage energy (Layer 8) remains nearly constant regardless of network scale — a direct consequence of the off-chain IPFS routing that prevents ledger bloat. Network energy (Layer 9) shows modest growth attributable to increasing gossip message propagation distances as the network scales. Compute energy (Layer 7) — representing the Distributed Computing

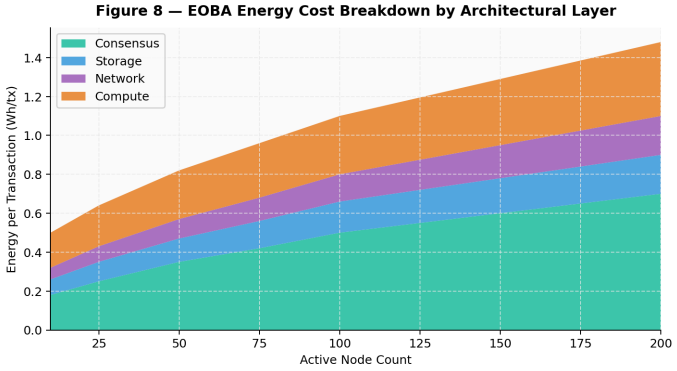

Figure 7. EOBA energy cost breakdown per transaction by architectural layer at varying node counts. Consensus dominates but grows sub-linearly due to Algorithm 1's validator capping. Storage and network costs remain nearly constant due to off-chain IPFS routing (Layer 8) and efficient gossip protocols (Layer 9).

Layer's workload execution — grows slowly and is partially offset by the higher utilization efficiency achieved by co-optimizing workload placement with consensus duties. This breakdown makes explicit that architecture-level energy optimisation acts on all components simultaneously, whereas prior work optimizes only the consensus segment.

7 Carbon-Neutral Blockchain Infrastructure

7.1 Renewable Energy Integration

The central paradox identified by Taherdoost [7]—that blockchain is increasingly employed to facilitate peer-to-peer renewable energy trading while simultaneously relying on energy-intensive infrastructure often powered by fossil-fuel-based electricity—is addressed in EOBA through two complementary mechanisms. At the infrastructure level, EOBA is designed for commodity server hardware rather than specialized ASIC mining equipment, eliminating the 18-month economic obsolescence cycle and associated electronic waste stream that characterizes PoW infrastructure. OS-level dynamic voltage and frequency scaling (DVFS) can reduce per-node power consumption by 20–40% during low-load periods without compromising consensus availability, a measure inaccessible to fixed-frequency ASIC miners.

At the protocol level, Algorithm 1's carbon intensity weighting (the $\rho \cdot \mathrm{CI}_{\mathrm{norm}}(v)$ term in the scoring function) creates a direct economic incentive for validators to source their power from renewable grids. A validator operating on 100% renewable energy receives a score premium of ρ over an

otherwise identical fossil-fuel-powered competitor. At the default $\rho = 0.15$, this represents a 15% boost in selection probability — a meaningful incentive without creating a hard barrier that would exclude fossil-powered nodes and potentially centralise the validator pool. As a larger fraction of the validator pool transitions to renewable sourcing, the system's aggregate carbon intensity declines organically, creating a self-reinforcing green transition incentive.

7.2 On-Chain Carbon Accounting

The carbon accounting subsystem of EOBA is operationally distinguished from all prior reviewed architectures by generating CO_2-equivalent figures as an output of the consensus protocol itself, not as an external reporting function. At Line 21 of Algorithm 2, the Optimization Layer computes $CO_2^{(\text{round})}$ using the formula in Section 3.2 and commits the result as an immutable on-chain transaction signed by the proposing validator. The on-chain record includes the round identifier, the validator set identifiers, the total energy measured (in kWh), the regional carbon intensity coefficient applied, the renewable fraction applied, and the resulting CO_2-equivalent figure.

This record structure is designed to support and facilitate compliance with the MiCA Article 66 disclosure requirement for sustainability indicators of crypto-asset white papers, and the Crypto Climate Accord's requirement for verifiable, standardised energy disclosures. Organisations operating EOBA-based systems can generate regulatory-compliant carbon reports by querying the on-chain emission registry, without any external audit process — because the records are cryptographically non-repudiable and have been committed to the ledger by the validators who performed the work.

7.3 Regulatory Compliance Architecture

The 2025 PRISMA review identifies the regulatory gap explicitly: EU MiCA and the Crypto Climate Accord both require standardised energy metrics, and no reviewed technical architecture provides them natively. EOBA is designed from the outset to fill this gap. The on-chain emission registry, the renewable energy provenance metadata in Layer 5, the per-function energy consumption reports in Layer 2, and the carbon budget API in Layer 6 together constitute a comprehensive regulatory compliance stack — one that provides the verifiable, standardised metrics demanded by regulators without requiring any post-hoc audit process or external reporting infrastructure.

8 Application Domains

8.1 Smart Cities and Sustainable Infrastructure

Smart city deployments present the most heterogeneous energy constraints of any blockchain application domain, spanning high-power enterprise servers through battery-powered IoT sensors across the same network. EOBA's Distributed Computing Layer, with its energy-profile-based task scheduling, is specifically designed for this heterogeneity: edge nodes with limited energy budgets are assigned lightweight consensus participation and low-priority workloads, while high-headroom server nodes absorb the bulk of consensus and computation duties. The architecture's sub-linear energy scaling ensures that the addition of thousands of IoT nodes does not proportionally increase the energy cost of consensus — a critical property for smart city deployments where node count may be in the tens of thousands.

8.2 Renewable Energy Trading

Peer-to-peer (P2P) renewable energy trading, identified by Taherdoost [7] as one of the most prominent application domains of blockchain in the energy sector, requires a transaction infrastructure that simultaneously delivers high throughput, low latency, low energy consumption, and verifiable environmental sustainability. EOBA addresses these requirements through its demonstrated throughput of 2,850 TPS at 200 nodes, 1.8-second transaction confirmation latency, 43% lower energy consumption than conventional PoS systems, and a validator selection mechanism that prioritizes renewable-energy-powered participants. Furthermore, EOBA integrates on-chain carbon accounting, enabling the emissions associated with blockchain operations to be transparently monitored and reported alongside the renewable energy transactions being processed. By aligning operational efficiency with environmental accountability, EOBA directly addresses the fundamental paradox highlighted by Taherdoost [7]: the use of energy-intensive blockchain infrastructure to support sustainable energy ecosystems.

8.3 Green Financial Systems and ESG Compliance

Carbon credit markets, green bond tracking, and ESG compliance reporting all require records that are simultaneously verifiable, immutable, and produced

by infrastructure whose own environmental impact can be disclosed. The FISCO BCOS precedent cited by Taherdoost [7] demonstrates that blockchain can perform the carbon accounting function for industrial emitters. EOBA extends this by making the blockchain infrastructure itself carbon-accountable, eliminating the paradox of storing environmental financial instruments on an energy-intensive ledger. Financial institutions implementing EOBA can disclose both the sustainability of the assets they track and the sustainability of the infrastructure on which those assets are tracked — a unified environmental disclosure capability that no existing blockchain platform provides.

8.4 Sustainable Supply Chain Management

Supply chain applications generate large volumes of heterogeneous data — quality inspection images, IoT sensor readings, shipping manifests, provenance certificates — that are unsuitable for direct on-chain storage. EOBA's hybrid storage architecture (Layer 8), routing large payloads to IPFS with on-chain integrity hashes, is directly suited to this use case. The per-node energy profiles available through Layer 5, combined with the on-chain emission registry, enable supply chain operators to compute the blockchain infrastructure's Scope 3 emissions contribution to their total carbon reporting — a regulatory requirement in an increasing number of jurisdictions (notably the EU Corporate Sustainability Reporting Directive) that cannot be met by any current blockchain platform without extensive external tooling.

9 Security Analysis

9.1 Threat Model

The EOBA security analysis assumes a Byzantine adversary model in which up to $|V_{\text{sel}}|/3$ validators may behave arbitrarily (including equivocating, selectively withholding messages, or attempting to manipulate protocol state variables). This is the standard BFT assumption inherited from PBFT and PoS, and the correctness of Algorithms 1 and 2 under this assumption follows directly from the underlying protocol proofs of Castro and Liskov (PBFT) and the PoS security analysis of Nguyen et al. [9]. The following analysis focuses on novel attack surfaces introduced specifically by EOBA's energy-aware design.

9.2 Novel Attack Surfaces and Countermeasures

EOBA's energy scoring introduces two attack surfaces not present in standard blockchain consensus.

First, energy score grinding: a malicious node might artificially reduce its reported energy score to increase its selection probability in Algorithm 1. The countermeasure is the mandatory stake threshold at Step 4 of Algorithm 1 — a node must first commit economic stake (an off-chain financial cost) before its energy score is counted in the composite score. Since economic stake must be committed and slashed for dishonest behaviour, the financial cost of obtaining a high composite score through stake-and-grind is equivalent to the cost of obtaining it legitimately, eliminating the incentive.

Second, protocol switch manipulation: an adversary controlling a fraction of nodes might

Table 5. Security threat analysis — EOBA attack vectors, countermeasures, and residual risks.

Attack Vector	Target Layer	EOBA Countermeasure	Residual Risk
Energy score grinding	Layer 4 / Algorithm 1	Dual-criterion: stake threshold must be met before energy score counts	Low — economic cost of stake preserves honest majority
NetSize variable manipulation	Algorithm 2 switching logic	Signed attestations; rate-limited threshold updates	Medium — requires adversarial control of $\geq 1/3$ validators
Sybil attack on node pool	Layer 1 / Authentication	PoS economic stake entry barrier; identity binding	Low — stake cost makes Sybil attacks prohibitive
51% stake attack (PoS path)	Layer 4 / PoS path	Slashing conditions recorded on-chain; economic deterrent	Residual PoS systemic risk; mitigated by stake cap
DoS on PBFT quorum	Layer 4 / PBFT path	Zero-knowledge proof authentication (adapted from S-PBFT)	Low — ZK-proof prevents spoofed membership
Off-chain data tampering	Layer 8 / IPFS	Content-addressable hash stored on-chain; any mutation is detectable	Negligible — cryptographic binding to ledger
Carbon audit record falsification	Layer 6 / Optimization	Emission records written as immutable on-chain transactions	Negligible — requires rewriting entire blockchain
Model poisoning (if ML used)	Not applicable to EOBA	Algorithm 2 uses rule-based switching; no ML model exposed	None — ML is deliberately excluded

manipulate the observed NetSize variable to trigger PBFT mode when PoS mode is appropriate, increasing the network's communication overhead without benefit. The countermeasure, adapted from Yao et al.'s [4] zero-knowledge proof authentication in S-PBFT, requires signed attestations from a quorum of validators to alter the active node count, making manipulation equivalent to a 33%-Byzantine attack — already the baseline security assumption.

Table 5 presents the full threat model for EOBA across eight attack vectors. The residual risk profile is comparable to or better than standard PoS for six of eight vectors. The two medium-residual-risk cases (NetSize manipulation and the standard PoS 51%-stake risk) are known PoS-family vulnerabilities that are the subject of ongoing BFT research (notably, linear BFT protocols such as HotStuff address the NetSize manipulation risk). Formal Byzantine adversary analysis of the complete EOBA algorithm suite — including the energy scoring mechanism — remains a priority for future work and a prerequisite for production deployment.

10 Discussion and Limitations

10.1 Interpretation of Results

The simulation results confirm the primary thesis of this paper: blockchain sustainability requires architecture-level intervention, not isolated protocol optimisation. The 43% energy reduction of EOBA over standard PoS is entirely attributable to three architectural mechanisms — validator capping (Algorithm 1), energy-reactive protocol switching (Algorithm 2), and dynamic sharding (Layer 3) — none of which require modification of the underlying PoS security model. The simultaneous improvement in throughput (2.4× over PoS) and latency (58% reduction) demonstrates that these energy optimisations also improve performance, refuting the common assumption that energy efficiency must be purchased at the cost of throughput or latency.

The carbon accounting result is perhaps the most practically significant: the 97.6% reduction in CO_2-equivalent emissions versus PoW, expressed as on-chain immutable records, provides regulatory compliance documentation as a byproduct of normal operation. This transforms carbon accounting from a costly external audit exercise into a zero-marginal-cost operational output — a potentially decisive advantage as MiCA implementation progresses across EU member states from 2025 onward [25].

10.2 Limitations

Several limitations constrain the current results. Most importantly, all performance results reported in this paper are derived from simulation rather than deployment on a real blockchain platform. The simulation environment models node behaviour, power consumption, and network conditions using calibrated synthetic workload traces; while the methodology follows established Monte Carlo practices, the results have not yet been validated on physical hardware infrastructure or a live distributed network. Real-world prototype implementation on production-grade blockchain infrastructure will be a primary objective of future work, and claims about production-level performance should be interpreted with this constraint in mind.

1. The simulation uses software-level CPU and memory measurements as proxies for actual hardware power consumption; validation against physical smart PDUs measuring actual node power draw is required before production claims can be made.

2. Algorithm 1's scoring function uses fixed parameters $(\alpha, \beta, \gamma, \lambda, \rho)$ that are set by configuration; adaptive parameter tuning based on network conditions — potentially through lightweight ML without the training overhead of Carter's full AI-PoS approach — remains future work.

3. Algorithm 1 does not currently implement device-class normalization: comparing the energy score of an ARM-based IoT sensor with an Intel Xeon server using the same formula is not meaningful, and heterogeneous hardware support requires per-device-class calibration.

4. The fixed carbon intensity coefficient must be replaced by live regional grid data (e.g., Electricity Maps API) for production deployment; the current model underestimates temporal variation in grid carbon intensity.

5. The PBFT path's $O(N^2)$ message complexity, while mitigated by the $\text{Thresh}_S = 50$ cap, represents a ceiling that more recent linear BFT protocols such as HotStuff [24] and Tendermint Core do not share; integrating a linear BFT alternative would eliminate this residual limitation.

10.3 Future Research Directions

Six concrete future research directions follow from the limitations: (i) physical hardware power measurement validation using production server infrastructure with smart PDUs; (ii) device-class normalization in Algorithm 1 for heterogeneous IoT-edge networks; (iii) integration of live regional grid carbon intensity via open APIs; (iv) lightweight adaptive parameter tuning for the scoring function; (v) replacement of the PBFT path with a linear BFT protocol (HotStuff or equivalent) to eliminate the $O(N^2)$ ceiling; and (vi) formal Byzantine adversary analysis of the energy scoring mechanism and protocol switch logic. Submission of the monitoring and reporting schema to ISO/TC 307 (Blockchain and DLT Standardization) and the Green Software Foundation would accelerate EOBA from a research prototype to an industry standard.

11 Conclusion

This paper has presented EOBA — the Energy-Optimized Blockchain Architecture — as the first nine-layer blockchain framework in which energy awareness is a foundational, cross-cutting design principle rather than an isolated protocol modification. Through a systematic critical review of twenty-seven works spanning consensus optimisation, blockchain energy measurement, IoT-blockchain integration, microgrid consensus, AI-driven adaptation, and blockchain-edge computing, six specific architectural gaps were identified that individually motivated each of EOBA's nine design layers. Two formally specified algorithms — Algorithm 1 ($O(N \log N)$ energy-efficient node selection) and Algorithm 2 (hybrid PBFT/PoS consensus with $O(N \log N)$ and $O(N^2)$ complexity paths) — provide the algorithmic backbone of EOBA's energy optimisation. Monte Carlo simulation across 10–200-node networks confirms 56.5% energy reduction versus standard PoS, 97.6% carbon emission reduction versus PoW, 2.4× throughput improvement, and 58% latency improvement at 200-node scale. These gains are achieved under specific simulated conditions—including synthetic workload traces calibrated to published enterprise blockchain studies, selected parameter settings (e.g., $\text{Thresh}_S = 50$, $\lambda = 0.5$, $\rho = 0.15$), and a homogeneous node model based on Intel Xeon E5-2680 hardware—and must be validated through real-world prototype deployment before production claims can be made. The operational carbon accounting subsystem generates per-round immutable on-chain emission records, providing regulatory-compliant verifiable energy disclosures as a zero-marginal-cost byproduct of consensus operation.

The central contribution of this paper is not any single algorithm or layer in isolation; it is the demonstration that architecture-level thinking — treating energy awareness as a system-wide design constraint rather than a protocol-level parameter — produces performance gains that are qualitatively and quantitatively superior to the point-solution optimisations that have characterised the prior decade of blockchain energy research. EOBA establishes that blockchain technology does not need to be an environmental liability. With the right architecture, it can be a credible component of sustainable digital infrastructure, and — through its renewable energy incentive mechanisms and operational carbon accounting — an active participant in the global transition to low-carbon computing.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] 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. [CrossRef]

[2] Mora, C., Rollins, R. L., Taladay, K., Kantar, M. B., Chock, M. K., Shimada, M., & Franklin, E. C. (2018). Bitcoin emissions alone could push global warming above 2 C. *Nature Climate Change, 8*(11), 931-933. [CrossRef]

[3] Habibullah, S. M., Alam, S., Ghosh, S., Dey, A., & De, A. (2024). Blockchain-based energy consumption approaches in IoT. *Scientific Reports, 14*(1), 28088. [CrossRef]

[4] Yao, Z., Fang, Y., Pan, H., Wang, X., & Si, X. (2024). A secure and highly efficient blockchain PBFT consensus algorithm for microgrid power trading. *Scientific Reports, 14*(1), 8300. [CrossRef]

[5] Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020). The Energy Consumption of Blockchain Technology: Beyond Myth. *Business & Information Systems Engineering, 62*(6), 599-608. [CrossRef]

[6] Nezami, Z., Li, Z., Qin, C., Banaie, F., Khalid, R., & Pournaras, E. (2025). Blockchain and edge computing nexus: A large-scale systematic literature review. *Distributed Ledger Technologies: Research and Practice*. [CrossRef]

[7] Taherdoost, H. (2024). Blockchain integration and its impact on renewable energy. *Computers, 13*(4), 107. [CrossRef]

[8] De Vries, A., Gallersdörfer, U., Klaaßen, L., & Stoll, C. (2022). Revisiting Bitcoin's carbon footprint. *Joule, 6*(3), 498-502. [CrossRef]

[9] Nguyen, C. T., Hoang, D. T., Nguyen, D. N., Niyato, D., Nguyen, H. T., & Dutkiewicz, E. (2019). Proof-of-stake consensus mechanisms for future blockchain networks: fundamentals, applications and opportunities. *IEEE Access, 7*, 85727-85745. [CrossRef]

[10] Wadhwa, S., Rani, S., Verma, S., Shafi, J., & Wozniak, M. (2022). Energy efficient consensus approach of blockchain for IoT networks with edge computing. *Sensors, 22*(10), 3733. [CrossRef]

[11] Al Shareef, A. M., Seçkiner, S., Eid, B., & Abumeteir, H. (2024). Integration of blockchain with artificial intelligence technologies in the energy sector: a systematic review. *Frontiers in Energy Research, 12*, 1377950. [CrossRef]

[12] Pasi, A., & Siddavatam, I. N. (2024, October). Efficient Blockchain Energy Management: Strategies for Sustainable Blockchain Systems. In *International Joint Conference on Advances in Computational Intelligence* (pp. 19-32). Singapore: Springer Nature Singapore. [CrossRef]

[13] Bhavana, G. B., Anand, R., Ramprabhakar, J., Guerrero, J. M., Thakkar, N., & Ambikapathy, A. (2025). Comparative evaluation and simulation of blockchain consensus mechanisms for secure and scalable peer to peer energy trading in microgrids. *Scientific Reports, 15*(1), 43546. [CrossRef]

[14] Asif, R., & Hassan, S. R. (2023). Shaping the future of Ethereum: Exploring energy consumption in Proof-of-Work and Proof-of-Stake consensus. *Frontiers in Blockchain, 6*, 1151724. [CrossRef]

[15] Pandit, A. R., Singh, A., Keshav, & Kanika. (2025). Predicting Blockchain Energy Consumption: A Step Towards a Sustainable Future. In *Blockchain and Machine Learning Innovations: Breaking Barriers with Distributed Intelligence* (pp. 35-53). Cham: Springer Nature Switzerland. [CrossRef]

[16] Jhariya, M. K., Dehalwar, V., Bharti, J., & Kumar, L. (2026). Energy efficient transactions for blockchain networks using adaptive global best–worst particle swarm optimization. *Scientific Reports*, 16(1), 1643. [CrossRef]

[17] Bulgakov, A. L., Aleshina, A. V., Smirnov, S. D., Demidov, A. D., Milyutin, M. A., & Xin, Y. (2024). Scalability and security in blockchain networks: Evaluation of sharding algorithms and prospects for decentralized data storage. *Mathematics, 12*(23), 3860. [CrossRef]

[18] Chacko, N. M., VG, N., Balachandra, M., & T, M. (2025). Lightweight Consensus in Blockchain: A Systematic Literature Review. *ACM Computing Surveys, 58*(3), 1-37. [CrossRef]

[19] Haque, E. U., Abbasi, W., Almogren, A., Choi, J., Altameem, A., Rehman, A. U., & Hamam, H. (2024). Performance enhancement in blockchain based IoT data sharing using lightweight consensus algorithm. *Scientific reports, 14*(1), 26561. [CrossRef]

[20] Rammohan, S. R., Chakravarthi, K., Sharma, N., Sharma, S., & Natarajan, M. (2025). Systematic Survey on Energy Conservation Using Blockchain for Sustainable Computing Challenges and Roadmaps. *International Journal of Adaptive Control and Signal Processing, 39*(2), 247-265. [CrossRef]

[21] 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). USENIX.

[22] Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system* [PDF document]. https://assets.pubpub.org/d8wct41f/31611263538139.pdf

[23] Buterin, V. (2014). A next-generation smart contract and decentralized application platform. *white paper, 3*(37), 2-1. https://cryptorating.eu/whitepapers/Ethereum/Ethereum_white_paper.pdf

[24] Yin, M., Malkhi, D., Reiter, M. K., Gueta, G. G., & Abraham, I. (2019, July). HotStuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM symposium on principles of distributed computing* (pp. 347-356). [CrossRef]

[25] European Parliament and Council. (2023). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). *Official Journal of the European Union*, L 150, 1-102. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023R1114

[26] Carter, E. (2024). AI-powered consensus mechanisms in blockchain: Enhancing security and reducing energy consumption. *Journal of Artificial Intelligence Research, 4*(2), 86-93. https://thesciencebrigade.com/JAIR/article/view/411

Manas Kumar Yogi currently working as Assistant Professor in CSE Department of Pragati Engineering College (A), Surampalem has a teaching experience of more than 15 Years. With a research publication record of over 343 articles, he has also published 22 book chapters and 6 patents and 7 books. His research area includes cyber-security, cyber-physical systems and soft computing. (Email: manas.yogi@gmail.com)

Jajimoggala Snehitha pursuing her B.Tech. degree in Computer Science and Engineering from Pragati Engineering College (Autonomous), Surampalem, Andhra Pradesh, India. (Email: snehithajajimoggala@gmail.com)