Scalable and Secure Blockchain Sharding for Decentralized V2V Energy Trading Networks

Yingsen Wang, Jun Zhao, Haonan Wang, YueYing Wang, Weihan Jiao, Weiyi Feng, Yu Miao, Guibin Wang, and Keqin Li, *Fellow, IEEE*

***Abstract*—Peer-to-Peer (P2P) energy trading represents a paradigm shift from traditional centralized energy transactions. Within the P2P energy trading, Vehicle-to-Vehicle (V2V) energy trading emerges as a quintessential example, underscoring the significance of exploring P2P energy trading in dynamic contexts. The investigation into V2V energy trading gains profound importance due to electric vehicles (EVs) acting as dynamic nodes within the network. This dynamism adds a layer of complexity and universality to the research. Blockchain technology offers a seamless match for the distributed nature of V2V energy trading. However, the majority of existing studies on blockchain-based V2V overlook the highly dynamic nature of EV nodes. Moreover, there is a notable scarcity of research focusing on the scalability of blockchain within the context of V2V energy trading. To address these challenges, this paper introduces sharding technique to explore its advantages and necessity in applying blockchain sharding to V2V energy trading, proposing a dynamic node sharding (DNS) strategy. The focus of our proposed method lies in the selection of sharding mechanisms, addressing the additional latency caused by cross-shard transactions, and mitigating the security vulnerabilities introduced by shard boundaries. We test our proposed method on the advanced BlockEmulator platform and deploy the V2V energy trading blockchain (ETB) on Hyperledger Fabric. The experiments demonstrate that our method significantly outperforms non-sharded V2V ETB approaches in terms of throughput and scalability. Furthermore, in comparison with other sharded V2V ETB frameworks, our method exhibits significant enhancements in both the management of cross-shard transaction validity and latency, and in the reinforcement of security protocols.**

***Index Terms*—Blockchain, energy trading, Peer-to-Peer, sharding, Vehicle-to-Vehicle.**

Received 12 March 2025; revised 18 June 2025 and 14 September 2025; accepted 16 October 2025. Date of publication 16 December 2025; date of current version 20 July 2026. This work is supported in part by the Technical Research Service Project of SGCC (Analysis of Operation Effectiveness of Power Market and Support of Platform Research, under Grant SGZB0000JYJS2501024, in part by State Grid Corporation Headquarters Technology Project under Grant 5400-202216167A-1-1-ZN, and in part by the Youth Foundation Project for Humanities and Social Sciences Research of the Ministry of Education under Grant 23YJCZH315. Paper 2025-ESC-0177.R2, presented at the 2024 IEEE Industry Applications Society Annual Meeting, Phoenix, AZ, USA, Oct. 20–24, and approved for publication in the IEEE TRANSACTIONS ON INDUSTRY APPLICATIONS by the Energy Systems Committee of the IEEE Industry Applications Society [DOI: 10.1109/IAS55788.2024.11023654]. *(Corresponding author: Yingsen Wang.)*

Yingsen Wang is with China Electric Power Research Institute Company Limited, Beijing 100085, China (e-mail: vjack9583@gmail.com).

Jun Zhao, YueYing Wang, Weiyi Feng, and Yu Miao are with the Department of Computer Science and Technology, Taiyuan University of Technology, Taiyuan 030024, China (e-mail: zhaojun0086@link.tyut.edu.cn; wangyueying0440@link.tyut.edu.cn; fengweiyi0416@link.tyut.edu.cn; miaoyu0433@link.tyut.edu.cn).

Haonan Wang is with the Beijing Electro-mechanical Engineering Institute, The Third Academy of China Aerospace Science and Industry Corporation, Beijing 100085, China (e-mail: wanghaonan@buaa.edu.cn).

Weihan Jiao is with the School of Electric and Electronic Engineering, North China Electric Power University, Beijing 102206, China (e-mail: 120212201610@ncepu.edu.cn).

Guibin Wang is with the Shandong Power Company of the State Grid, Rizhao 102206, China (e-mail: wgb369@sohu.com).

Keqin Li is with the Department of Computer Science, State University of New York, New Paltz, NY 12561 USA (e-mail: lik@newpaltz.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TIA.2025.3645002.

Digital Object Identifier 10.1109/TIA.2025.3645002

I. INTRODUCTION

TRADITIONAL centralized energy trading systems are often criticized for their inefficiency and opacity, leading to higher costs and limited consumer engagement [1], [2]. In contrast, P2P energy trading has emerged as a revolutionary alternative, providing greater efficiency, transparency, and empowerment for participants [3]. V2V energy trading represents a cutting-edge area of research in the P2P framework [4]. The inherent mobility of participants in V2V transactions adds a universal dimension, enabling flexible and spontaneous energy exchanges that can take place anytime and anywhere [5], [6]. V2V energy trading is paving the way for intelligent, autonomous energy systems, where EVs can seamlessly switch between roles as consumers and providers [7], [8], [9].

While V2V energy trading holds significant promise for transforming energy distribution systems, its decentralized structure presents notable challenges, especially concerning security and trust between participants [10]. Blockchain technology, known for its decentralized and distributed ledger features, provides an ideal solution to these issues, delivering a secure and transparent framework for conducting transactions [11].

However, numerous formidable challenges confront the V2V Energy Trading Blockchain (ETB)'s implementation:

① Existing research on V2V and P2P energy trading predominantly adopts blockchain frameworks, yet frequently fails to account for the inherent dynamism of EV nodes. Blockchain technology, initially designed for digital currencies in the financial sector, operates within static node environments. In contrast, EV nodes in V2V ETB are highly mobile, posing substantial challenges for information exchange and data synchronization among vehicles.

② Current V2V ETB research often employs static consensus mechanisms ill-suited for dynamic EV networks, leading to scalability bottlenecks like high latency and low throughput. The lack of dynamic shard reconfiguration exacerbates cross-shard

delays and security risks, highlighting the need for lightweight, topology-aware consensus designs tailored to V2V energy trading's latency-sensitive and resource-constrained nature.

③ To address scalability within V2V ETB, a minority of studies have employed sharding techniques. However, these efforts have primarily focused on simplistic node partitioning strategies, such as basic network sharding, without addressing the complexities introduced by node mobility or the latency and security vulnerabilities associated with cross-shard transactions.

Although blockchain technology has been integrated into V2V energy trading systems, there is a significant gap in addressing scalability challenges. Ensuring robust security, especially concerning data integrity and transaction validation, remains a paramount issue in these systems. Current frameworks frequently fall short in providing sufficient safeguards for cross-shard transaction security.

To address the aforementioned challenges, we propose a DNS technique. This technique concentrates on the selection of sharding methods, consensus mechanisms within shards, and the transaction efficiency between shards to enhance the system's scalability. The main contributions of this paper are as follows:

1) Nodes within the V2V ETB are categorized into three types: regulatory nodes (RN), active EV nodes (AN), and standard EV nodes (SN). Each category is assigned specific roles and responsibilities to mitigate challenges arising from node dynamism, as outlined in issue ①. Architecturally, this involves adopting two modes of sharding: network sharding for RN and state sharding for AN and SN.
2) To tackle issue ②, we deviate from conventional blockchain architectures and Byzantine Fault Tolerance (BFT) consensus protocols. Within DNS shards, a parallel accounting mechanism is employed to enhance throughput and minimize transaction latency. Drawing from prior work [1], cryptographic randomness is utilized to periodically select and refresh RN and some AN, ensuring the system maintains high security and fault tolerance after sharding.
3) In response to issue ③, we address transaction interruptions caused by cross-shard interactions and the delays associated with shard reconfiguration, as highlighted in ②. The proposed DNS approach is evaluated using the BlockEmulator simulator and implemented on the Hyperledger Fabric platform. Results demonstrate that the DNS method outperforms existing solutions in both intra-shard (In-DNS, IDNS) and inter-shard (Cross-DNS, CDNS) scenarios.

One of the emerging applications of blockchain lies in the real-time regulation and dispatch of renewable energy sources. Renewable generators such as solar and wind are intermittent and decentralized, which introduces uncertainty and coordination complexity into power systems. Blockchain provides a decentralized trust framework to coordinate multiple renewable assets in real time. For example, Chen et al. [12] proposed a Proof-of-Solution (PoSo) consensus mechanism that replaces conventional puzzles with real-time optimization problems for energy dispatch. The consensus directly aligns with the grid's operational needs, enabling real-time matching between supply and demand while securing the data exchange. This paradigm shows the potential of blockchain beyond simple recording, it becomes a computational layer for intelligent control. Projects like PowerLedger and Brooklyn Microgrid have demonstrated the feasibility of blockchain-based market settlement, using smart contracts to automate pricing and transaction execution without centralized utilities. Various consensus mechanisms such as Proof-of-Stake (PoS) and Delegated Byzantine Fault Tolerance (DBFT) have been explored to balance performance with trust in permissioned environments.

Building upon the content of the conference paper, this paper incorporates theoretical insights into the infinite scalability of the V2V ETB. Additionally, the experimental section has been expanded to include an analysis of how variations in shard size and block size impact system scalability. Furthermore, a detailed security analysis has been conducted to provide readers with a comprehensive understanding [13].

The remainder of the paper is organized as follows. Section II presents related works. Section III introduces our ETB model and the DNS method. Section IV provides the performance of the V2V ETB using our DNS method. The last section concludes our work and outlines future research.

II. Related Works

A. P2P or V2V Energy Trading Using Blockchain Consensus Mechanism

This paper presents the DNS method, which employs a BFT consensus mechanism with parallel ledger-keeping within shards and among RN. This section reviews existing blockchain consensus mechanisms used in energy trading blockchain. Reference [14] employs Proof of Work (PoW) to enhance data security and transaction efficiency, but its high computational demands hinder its suitability for sustainable energy systems. Reference [15] explores the application of the PoS consensus mechanism in P2P energy trading within a microgrid, focusing on its efficiency in recording transactions and managing payments without intermediaries. By restricting mining participation to prosumers within the microgrid, the approach seeks to lower mining complexity and energy usage. Nonetheless, stakeholders with larger holdings may exert undue influence over the network, potentially compromising the decentralized principles fundamental to blockchain-based energy trading systems. Reference [12] introduces PoSo, a blockchain consensus mechanism designed to integrate mathematical optimization into energy systems, replacing PoW's arbitrary puzzles with meaningful optimization problems. However, PoSo struggles with scalability and computational efficiency. The aforementioned methods can be categorized under the Proof-of-X (PoX) class of consensus algorithms. Applying the PoX consensus mechanism in P2P and V2V energy trading faces several drawbacks, such as potential scalability issues due to the complex and dynamic nature of energy networks.

Relevant research has explored the implementation of BFT-based consensus algorithms in energy trading. Reference [16] explores employing a Byzantine-based blockchain consensus framework to secure energy trading between EVs and the distribution network, particularly during peak load times.

Reference [17] presents an innovative consensus mechanism called SDMA-PBFT for permissioned blockchain systems. By moving from $O(N^2)$ to $O(nk\log_k n)$ communication cost, it enhances the scalability of the PBFT algorithm, making it more suitable for large-scale applications. Reference [18] introduced a scheme integrating voting rewards and punishments with a PBFT-based consistency protocol for credit evaluation. Reference [19] put forward a PBFT algorithm designed for interactive entities in multi-energy systems. [20] described a DPBFT model adapted to the fluctuating reputations within an energy blockchain framework.

Although existing research addresses node participation, persistent challenges such as substantial transaction latency and limited throughput remain unresolved. Current studies predominantly rely on conventional PBFT or its derivatives, which are insufficient to meet the scalability and efficiency demands of large-scale energy trading systems. To address these limitations, there is a critical need for energy trading platforms to innovate and optimize consensus mechanisms. Enhancing these protocols is essential to support the high-speed, high-volume transaction processing required in modern energy trading ecosystems, ensuring robust performance and scalability.

B. Blockchain Sharding Technique

Research on the application of blockchain sharding technology in the context of P2P and V2V energy trading remains limited. This section outlines key applications of blockchain sharding in energy transactions and explores its utilization in other fields. It concludes by addressing the appropriate implementation of blockchain sharding techniques within V2V ETB.

To address scalability challenges in blockchain systems, two prominent approaches are off-chain solutions and sharding techniques. Off-chain methods, however, are less suitable for V2V ETB due to the need for real-time transaction validation and the high trust and security requirements inherent in such systems. Elastico [21] is notable as the first sharding-based, permissionless blockchain protocol. It employs Proof of Work (PoW) for shard formation and Practical Byzantine Fault Tolerance (PBFT) for consensus within shards. Elastico enhances scalability by partitioning the network into smaller, independent shards, each capable of processing transactions autonomously. Elastico incorporates fundamental network sharding and transaction sharding techniques, and there are studies that have implemented more intricate state sharding models, such as Omniledger [22], Monoxide [23], and RapidChain [24]. In state shards, nodes maintain only a portion of the state data associated with their accounts. Omniledger is a protocol that focuses on maintaining the security and robustness of a decentralized ledger through sharding [22]. RapidChain is known for its resilience to Byzantine faults and for increasing the transaction throughput by tackling the communication overhead per transaction, which is seen as a significant bottleneck in previous sharding approaches [24]. Monoxide, on the other hand, introduces a unique way of handling shards by employing a different network structure that optimizes for efficiency and consensus procedures [23].

Despite these advancements, sharding technologies face several challenges. When a node transitions between shards, it must synchronize the entire blockchain history within the new shard, leading to significant latency in data synchronization. Periodic shard reshuffling, necessary for maintaining security, can cause transaction delays or temporary processing halts. Additionally, the dynamic nature of nodes in V2V ETB imposes stricter demands on sharding techniques. However, current research primarily employs static network sharding approaches, such as [25] and [26]. To be effective in V2V ETB, sharding mechanisms must adapt to the dynamic entry and exit of nodes, ensuring seamless integration and operation within the system.

C. Blockchain in Real-Time Control of Energy Systems

Ref. [27] proposes a novel consensus mechanism named Proof of Task to achieve real-time security control in renewable energy power systems. By replacing traditional mining with grid-stability-related computations, it significantly enhances blockchain's real-time performance. It improves both cybersecurity and operational reliability under dynamic conditions. Andoni et al. [28] provided a foundational systematic review, arguing that blockchain can act as a decentralized control layer for microgrids. They posit that smart contracts can automate grid services, such as frequency regulation and voltage support, by encoding control logic that executes upon receiving real-time data from grid sensors. It moves the concept beyond mere financial settlement and into the realm of physical grid management. Knirsch et al. [29] propose a privacy-preserving algorithm where dynamic tariff decisions are made and executed via a blockchain. The smart contract automatically adjusts charging schedules based on real-time grid congestion data, acting as a control loop that alleviates strain on the distribution network without compromising user data. It highlights a key control application: using blockchain to automate real-time, distributed load control in response to system conditions. Ref. [30] develops a blockchain-based framework for real-time operation and P2P energy trading in crowdsourced energy systems. It introduces a two-phase algorithm that integrates day-ahead scheduling with real-time control to manage distributed energy resources dynamically. The system is prototyped using IBM Hyperledger Fabric to enable secure and efficient transactional coordination between prosumers and utilities.

III. Proposed V2V Sharding Method

A. Selection of Sharding Patterns in V2V Energy Trading

In the realm of V2V energy trading, the choice of an optimal sharding strategy is critical to improving system scalability and operational efficiency. The primary sharding methodologies include network sharding, transaction sharding, and state sharding. A hybrid model integrating network and state sharding is considered particularly effective for this use case, offering a balanced solution to the unique demands of decentralized energy transactions.

The blockchain nodes within the V2V energy trading ecosystem are categorized into three distinct types: RN, AN and SN. RN primarily include government agencies, regulatory bodies, and a subset of AN, which are distinguished by their consistent online activity and involvement in block validation

within specific geographic areas. In contrast, SN demonstrate lower participation levels, primarily engaging in transactions relevant to their interests. Due to their significant contributions, RN and AN receive proportionally higher rewards.

The adoption of network sharding for RN and AN is supported by the relatively static nature of regulatory entities and the localized presence of AN, often operating within a single shard. Furthermore, the operational limitations of V2V energy trading systems, such as the restricted communication ranges imposed by Road Side Units (RSUs), reinforce the suitability of network sharding for these node types. On the other hand, SN exhibit high dynamism and frequent geographic mobility, often transitioning between shards. This necessitates the use of state sharding, a more advanced approach designed to accommodate the variability and unpredictable distribution of these nodes effectively.

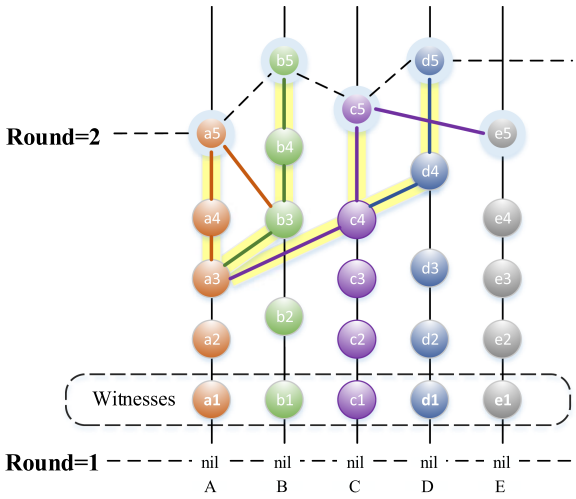

Fig. 1. The consensus with in shards.

B. The Consensus Within Shards

Building upon our prior work [1], which introduced an advanced Hashgraph-based parallel ledger-keeping framework for the V2V ETB, this study extends the framework by categorizing block synchronization into intra-shard and cross-shard mechanisms. We implement a Directed Acyclic Graph (DAG)-inspired parallel ledger-keeping model within individual shards. The DAG-based model also facilitates efficient handling of dynamic node participation, a critical requirement for V2V energy trading networks characterized by frequent node mobility and varying energy demands. It is noteworthy that most studies conventionally utilize PoW or PBFT consensus mechanisms, such as in classical approaches [23] and [31]. RN, along with highly active AN, form a committee that pre-agrees on atomic blocks to enhance the parallel processing capabilities of V2V ETB. Blocks are classified into two types: standard blocks and atomic blocks. Atomic blocks serve as the foundation for consensus within the committee, enabling parallel processing and significantly enhancing system throughput. Unlike the pre-validation phase for candidate primary nodes described in [1], our approach eliminates this step, streamlining the consensus process while maintaining robustness against malicious behavior. By restricting consensus activities to a dedicated committee, the system achieves both scalability and adaptability, aligning with the principles outlined in [1]. This approach not only ensures robust and tamper-proof consensus but also accommodates the fluid nature of decentralized automotive networks, making it a practical solution for real-world energy trading applications.

Atomic blocks consist of a set of energy transactions among EV nodes, timestamps, and hash references to two parent atomic blocks. Traditional blockchain systems typically link each block to a single predecessor, creating an immutable and sequential chain. However, this design often leads to performance bottlenecks, as adding a new block requires consensus across the entire network. In contrast, the consensus mechanism proposed in this study mandates that each atomic block references two parent atomic blocks: one from its own prior block and another from any other node within the shard. This dual-reference approach enhances parallelism and reduces bottlenecks, as illustrated in Fig. 1, while maintaining the integrity and efficiency of the consensus process within shards.

Within the committee, nodes implement a "second-order message propagation" algorithm, starting with a gossip protocol to distribute atomic blocks. At fixed intervals, such as every 1.0 s, each node broadcasts an atomic block, denoted as m, to k neighboring nodes, simulating the spread of a rumor. When a node receives m for the first time, it forwards m to k neighboring nodes, excluding the node from which it originally received m. This process continues iteratively across the network until all nodes have received m. BExpanding on the gossip mechanism, the second-order message propagation activates when an EV node, after receiving m via gossip, appends its digital signature to m before rebroadcasting it. This step authenticates the origin of the atomic block's validation. The EV node then encapsulates the signed m into a new atomic block, M, and randomly disseminates M other EV nodes. Each propagated atomic block includes signatures that validate its predecessor atomic blocks, m, ensuring the continuity and verifiability of the information across the network.

C. The Dynamic Node Sharding

The RN and a subset of AN form geographically distributed regulatory shards within the DNS framework. These regulatory shards are responsible for packaging and publishing blocks, ensuring their legality, and maintaining a complete ledger of all accounts within their shard. When a new EV node joins a shard, it relies on the regulatory shard to validate transactions and verify block legitimacy. While regulatory shards introduce a degree of centralization, this paper addresses the risk of collusion among byzantine nodes by integrating randomness through Verifiable Random Functions (VRF) [5], [32]. This periodic reshuffling of shards enhances the unpredictability of malicious activities. The process for electing shard committees is detailed in Algorithm 1. Furthermore, AN can also participate in regulatory shards, provided they have sufficient storage capacity and are willing to engage in the process, which offers greater rewards. The

Algorithm 1: Shard Committee Election using VRF.

Input: N: Set of all nodes in the V2V ETB; T: Threshold value for score comparison; E: Duration of an epoch
Output: Committee

$Committee \leftarrow \emptyset$
for *each node* $\in N$ **do**
  $(score, proof) \leftarrow$
    $VRF(node.privKey,$
    $Concatenate(EpochSeed, node.pubKey))$
  $isValid \leftarrow VerifyVRF(node.pubKey, score, proof)$
  if $isValid$ ***and*** $score < T$ **then**
    $Committee \leftarrow Committee \cup \{node\}$
  end
end
return *Committee*

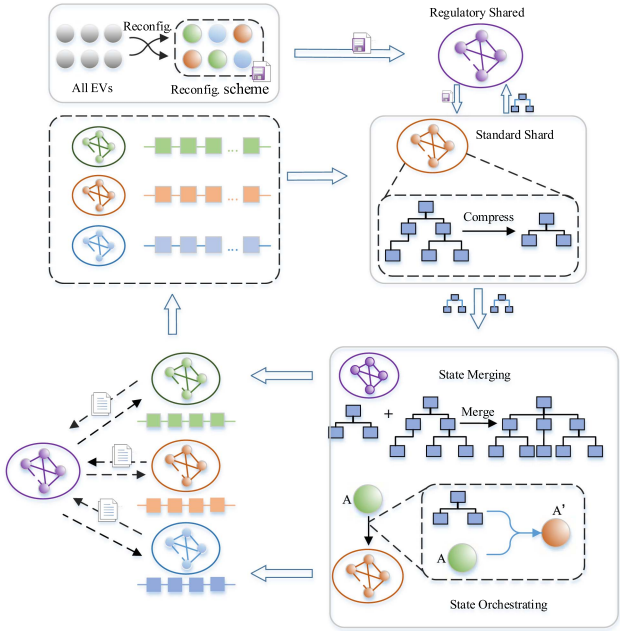

Fig. 2. The composition of the Energy Trading Blockchain.

DNS strategy optimizes computational resources by utilizing a BFT consensus mechanism with parallel ledger-keeping within shards, significantly reducing the computational demands compared to conventional blockchain systems.

The SN constitute ordinary shards responsible for validating transactions and blocks, as well as generating new blocks. Nodes within these ordinary shards, constrained by limited storage capacity or frequent changes in shard location, retain only blocks relevant to their own interests and those of active EV nodes. Additionally, these nodes maintain a record of several historical blocks they have previously validated. Since the blockchain operates by linking new blocks to previous ones, these historical blocks serve to validate the system's latest blocks and provide essential transaction information for nodes newly joining the shard. This selective retention of data ensures that ordinary shard nodes remain lightweight and efficient, while still contributing to the overall integrity of the blockchain. The DNS approach aims to strike a balance between efficiency and security in the shard-based blockchain system. By mitigating centralization risks and promoting adaptability, it fosters a resilient network structure capable of accommodating the dynamic nature of V2V energy trading environments. This design ensures that the system remains robust against malicious activities while maintaining high performance and scalability, even as the network evolves and expands. The integration of dynamic sharding and efficient data retention mechanisms further enhances the system's ability to handle the complexities of decentralized energy trading, making it well-suited for real-world applications in smart grids and autonomous vehicle networks.

To maintain the security and operational integrity of the V2V ETB system, we require periodic reconfiguration of sharding within the DNS framework at predetermined block intervals (specifically, every 100 blocks, as defined in this study). This reconfiguration ensures the accurate updating of account statuses for all EV nodes across both regulatory and ordinary shards. The DNS architecture is illustrated in Fig. 2. The reconfiguration process for regulatory shards is enabled through the use of Verifiable Random Functions (VRF), as previously mentioned, with all nodes within the shard collaboratively reaching consensus on the new sharding configuration. Once the proposed sharding architecture achieves unanimous approval across the entire network—encompassing both regulatory and ordinary shards—the system transitions into the reconfiguration phase. During this phase, an advanced Ethereum-based data structure [31] is employed to facilitate the process. This systematic approach to shard reconfiguration ensures the continuous adaptation and fortification of the V2V ETB system against potential security vulnerabilities while keeping participant account information current. The use of VRF in the restructuring process guarantees the randomness and impartiality of shard reassignment, thereby enhancing the network's resilience against manipulation. Furthermore, the requirement for network-wide consensus prior to reconfiguration aligns with the decentralized principles of blockchain technology, ensuring that any modifications to the network structure are democratically validated. The adoption of an enhanced Ethereum data structure during the reconfiguration phase highlights the commitment to leveraging advanced, secure, and efficient mechanisms for data storage and transaction processing. This reflects the ongoing evolution of blockchain systems to address the complexities of dynamic energy trading environments.

For a given EV node i within an ordinary shard x, the process involves traversing the state tree of its local account to remove nodes that have not participated in shard transactions for an extended period. This procedure results in an enhanced data structure, denoted as E_x^i. The enhanced data structure E_x^i symbolizes a pruned version of the shard's state that reflects a more accurate and dynamic representation of active participants, thereby improving the overall performance and security of the system. Nodes dispatch the enhanced data structure to peers

within both the new regulatory shards and ordinary shards. Upon receiving E_x^i, nodes within the regulatory shard amalgamate it with their own E_{y-1}^j, to derive their updated E_y^j, where $y-1$ signifies the previous term of the regulatory shard, and j epresents the identification number of that regulatory node. Conversely, nodes within ordinary shards retain their own E_x^i to facilitate transaction and block verification in subsequent processes. Ordinary shard nodes, by preserving E_x^i, maintain a streamlined dataset optimized for verifying new transactions and blocks, thereby reinforcing the system's integrity and efficiency.

Within regulatory shards, nodes update the enhanced data structures received from ordinary shards, denoted as $E_x^0, E_x^1, \ldots, E_x^N$. Subsequently, for all ordinary shards, regulatory nodes perform consensus on the root hash values of each E_x^i, achieving agreement on the account states across the network. Simultaneously, ordinary EV nodes reach consensus on E_x^i, signaling the completion of the shard's reconfiguration phase. This consensus among SN marks the transition into a new consensus phase for ordinary shards, wherein they commence the packaging and publishing of new blocks.

In the following sections, we explore cross-shard transactions, extending our earlier research that focused solely on inter-shard transaction patterns. This study broadens the scope to encompass not only transactions between shards but also those occurring across different phases. A "phase" is defined as a consensus period that concludes with the generation and formal addition of every 100 blocks to the chain. After each phase, the system enters the shard reshuffling phase, as previously outlined. Due to the high mobility of EV nodes, their entry times into the same shard often differ. Transactions that occur within the same shard but span multiple phases are categorized as cross-phase transactions. This expanded framework addresses the complexities introduced by the dynamic nature of EV nodes, ensuring the system remains robust and adaptable to varying network conditions. By incorporating cross-phase transactions, the model enhances its ability to handle temporal and spatial variability in energy trading, further improving the scalability and efficiency of the V2V energy trading network.

We denote a cross-shard transaction as Ψ_{tr}, involving a transaction (either tokens or energy) from user a in shard S_{in} to user b in shard S_{out}. The formulation of a cross-shard [23] transaction Ψ_{tr} is presented as $\Psi_{tr} = \langle a, b, \omega \rangle$, where ω represents the transfer amount. A cross-shard transaction, denoted as Ψ_{tr}, can be decomposed into sub-transactions pertaining to two distinct shards, which can be represented as $\Psi_{tr_1} = \langle \mu, a, \omega \rangle$ and $\Psi_{tr_2} = \langle \lambda, b, \omega, \phi \rangle$, where μ represents the amount or energy transferred from EV node a, correspondingly, λ represents the amount or energy received by EV node b, and ϕ signifies the on-chain legitimacy proof of Ψ_{tr_1}. Then the proof can be represented as $P = \langle \sigma, \Psi_{tr_1}, H(\Psi_{tr_1}), H(Merkle) \rangle$, where σ represents the phase number of Ψ_{tr_1}, $H(\Psi_{tr_1})$ is the hash value of the block that contains the Ψ_{tr_1}, and the $H(Merkle)$ denotes the Merkle proof of EV node a [31].

For the two cross-shard transaction modes, cross-shard and cross-phase transactions are primarily considered. If σ represents the current phase index of the node's location, this scenario is identified as a cross-shard transaction. Regulatory nodes within S_{in} can directly request hash value $H(\Psi_{tr_1})$ and Merkle proof $H(Merkle)$ from S_{out}, subsequently incorporating the cross-shard transaction into the newest block. Conversely, if σ denotes the phase index of a node's previous location, the situation is characterized as a cross-phase transaction. Regulatory nodes in S_{in} must request a list of nodes from the regulatory shard in S_{out}. Following this, regulatory nodes in S_{in} proceed to request $H(\Psi_{tr_1})$ and $H(Merkle)$ based on the acquired node list, ultimately publishing a block that contains the cross-shard transaction. The process of cross-shard transactions is detailed in Algorithm 2. Our DNS strategy divides the network into smaller shards, each managing transactions among specific subsets of EV nodes. DNS dynamically adjusts shard boundaries in response to real-time EV movement and energy demand, ensuring balanced shard workloads and reduced latency. To further minimize cross-shard transaction delays, we employ an efficient routing algorithm and batch processing techniques. Security is strengthened through a multi-layered protocol that integrates cryptographic methods, consensus mechanisms, and an intrusion detection system. The model also supports horizontal scalability, enabling the seamless addition of new shards and optimizing resource allocation to accommodate growing network demands. This approach ensures the system remains efficient, secure, and adaptable to the dynamic requirements of V2V energy trading environments.

Algorithm 2: Cross-Shard Transactions.

Input: Cross-shard transaction Ψ_{tr}: transfer from a in S_{in} to b in S_{out};
Ψ_{tr_1} and Ψ_{tr_2}: the decomposed transaction;
Transfer amount: ω; On-chain legitimacy proof: ϕ;
Proof P including: phase σ, hash value $H(\Psi_{tr_1})$, and Merkle proof $H(Merkle)$
Output: Committee

```
if σ is current then
    Request and incorporate H(Ψ_tr1) and H(Merkle) from S_out into new block
else
    Request node list from regulatory shard in S_out
    Request H(Ψ_tr1) and H(Merkle) based on node list
    Publish block containing Ψ_tr in S_in
end
return Succeed
```

We examine the fraction of nodes involved in the reconfiguration process under the partial reconfiguration scheme. To maintain the sharding network's integrity during reconfiguration, the ratio of malicious nodes in the temporary committee must remain below one-third [31]. Assume a shard initially contains d malicious nodes out of a total of e EVs. Let f represents the proportion of nodes participating in the reconfiguration process. The remaining $e(1-f)$ nodes constitute the temporary committee, tasked with block proposal and transaction validation. The probability that fewer than one-third of the nodes in the

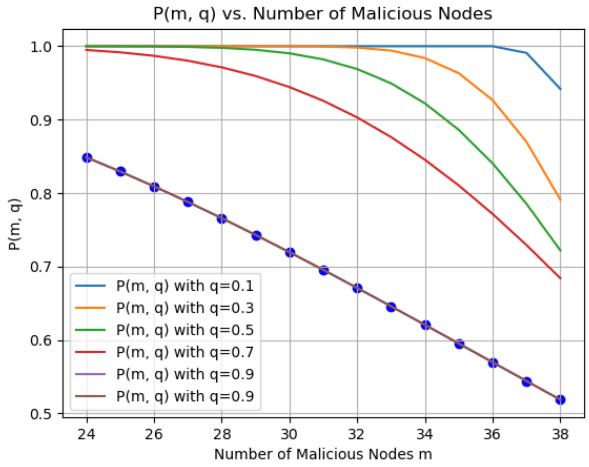

Fig. 3. The probability of secure and stable operation of the node committee under different proportions of malicious nodes.

temporary committee are malicious is given by:

$$P(d,f)=\sum_{x=0}^{\min\left\{\lfloor\frac{e(1-f)}{3}\rfloor,d\right\}}\frac{\binom{d}{x}\cdot\binom{e-d}{\lfloor e(1-f)\rfloor-x}}{\binom{e}{\lfloor e(1-f)\rfloor}} \tag{1}$$

Taking into account the possible values of d, the average likelihood of the temporary committee operating effectively can be determined as:

$$P(f)=\frac{3}{e}\sum_{d=0}^{\lfloor\frac{e}{3}\rfloor}\sum_{x=0}^{\min\left\{\lfloor\frac{e(1-f)}{3}\rfloor,d\right\}}\frac{\binom{d}{x}\cdot\binom{e-d}{\lfloor e(1-f)\rfloor-x}}{\binom{e}{\lfloor e(1-f)\rfloor}} \tag{2}$$

The probability of secure and stable operation of the node committee under different proportions of malicious nodes is shown in Fig. 3. In standard scenarios, the security of temporary committees can be upheld at a notably high level. By altering the fraction of malicious nodes within a shard and modifying the ratio of nodes involved in the shuffling process, denoted by d/e and f respectively, the average likelihood that the temporary committee contains fewer than one-third malicious nodes can be ascertained. When the proportion of shuffled nodes is less than 70% and the fraction of malicious nodes in the initial shard is below 30%, the temporary committee can attain a security level of 95%. For specific data, please refer to the experimental section.

D. Independent Blockchain and Main Chain

Our system architecture comprises three core components: independent chains for transaction recording, a main chain for maintaining a globally shared state, and a validation mechanism for transaction verification. Each node maintains an independent chain that logs transactions in a first-in-first-out sequence. These independent chains consist of an ordered series of blocks, each containing the hash of the preceding block and a set of transactions [33].

Blocks within the independent chains can vary in size. Periodically, nodes transmit nano block to the main chain. These atomic blocks encapsulate the essential information from the independent chains. The main chain serves to uphold the global state, while the validation mechanism ensures the integrity and validity of transactions. This design allows the system to maintain transaction records effectively while making concessions to the termination condition of BFT consensus protocols. The independent chain for node j comprises an ordered sequence of blocks $\{B_{(j,1)}, B_{(j,2)}, \ldots\}$, where each block $B_{j,k}$ is structured as $\{H(B_{j,k-1}), T_{j,k,1}, T_{j,k,2}\ldots\}$. Here, $T_{j,k,l}$ represents a transaction initiated by EV_j, which must satisfy conditions of valid sources, value consistency, and absence of double-spending. An initial value is assigned to each EV, analogous to a transaction without a source, where both the sender and receiver are the EV itself. A nano block of block $B_{j,k}$ is represented by $N_{j,k}=\langle j,k,H(B_{j,k}),Sig(j||k||H(B_{j,k}))\rangle$.

The main chain utilizes the DAG as its underlying structure, replacing the traditional PBFT consensus algorithm. The blocks in the main chain consist of nano blocks signed by the corresponding nodes. We assume that the abstracts of all genesis blocks are already included in the main chain. Since DAG-based systems have been demonstrated to achieve BFT consensus in our network model, we can consider the main chain as a reliable and secure primitive, with all abstracts included on the main chain reaching BFT consensus. Honest nodes will send nano blocks of their newest blocks to the main chain once they observe that their previous nano blocks have been successfully added to the chain.

In the independent blockchain, node transactions are not globally authenticated, lack node signatures, and are not tamper-proof. If a nano block associated with a block or subsequent nano blocks are included in the main chain, the transactions within these blocks become tamper-proof and contain the signatures of the nodes that issued the transactions. In this paper, such blocks are defined as deterministic blocks. A block $B_{j,k}$ is formally added to the main chain under two conditions: (1) any subsequent nano block associated with it, i.e., $N_{j,k'}, k'\geq k$, is included in the main chain, and (2) all nano blocks of EV_j, denoted by $A_{j,l}$ are present on the main chain, and $l\leq k'$. A transaction Tr_m can be denoted by $Tr_m=\langle Origin_m, s_m, r_m, tv_m, rv_m\rangle$, where $Origin_m$ is the set of transactions associated with Tr_m, s_m is the transaction sender, r_m is the receiver, tv_m is the transacted value, rv_m is the remaining value. A transaction Tr_m is confirmed if $Tr_m\in B_{j,k}, s_m=j$, and $B_{j,k}$ is confirmed. The $B_{j,k}$ and Tr_m are determined by the nano block $N_{j,k'}$. A transaction once confirmed becomes immutable, meaning it is officially recorded on the chain.

If $Tr_m=T_{j,k,l}$ is confirmed by the nano block $N_{j,k'}, k'\geq k$, Then, there cannot exist another confirmed chain containing a series of confirmed blocks $\{B'_{(j,1)}, B'_{(j,2)}, \ldots\}$ such that $T'_{j,k,l}\neq Tr_m$ while their corresponding hash values are equal. Once a transaction is verified, its position and content become immutable. Additionally, it is authenticated because the transaction's originator matches the signer of the summary, and the summary includes an unforgeable signature from the originator. It's important to note that a verified transaction in this context differs from those in systems like Bitcoin, as they have not yet undergone validation.

An EV_j needs to perform two tasks to validate transactions. First, EV_j collects nano blocks; any EV interested in a particular transaction can efficiently gather the nano blocks related to that transaction. Subsequently, they use a verification function [33] to determine whether the transaction is valid. Let $S_m = j$ be the sender of transaction Tr_m, where $Tr_m \in B_{j,k}$. Suppose there is a nano block $N_{j,k'}$ with $k' \geq k$ on the main chain. The validity nano proof $P(Tr_m)$ onsists of the collection of all blocks up to and including $B_{j,k'}$ and the proofs for all transactions in $Origin_i$. Formally, this can be expressed as $P(Tr_m) = \{B_{j,k''} | k'' \leq k'\} \cup \{B_{v,l} \in P(Tr_n), Tr_n \in Origin_m\}$. the validity proof for a transaction Tr_m within block $B_{j,k}$ encompasses the entire chain of node EV_j from its initial block up to a subsequent block $B_{j,k'}$, where $k' \geq k$ and a nano block of this block is present in the main chain. Additionally, it includes the chains of all source transactions of Tr_m, extending recursively to their respective sources until reaching the initial blocks.

If EV_j wishes to verify the validity of transaction Tr_m, it can request EV_v to provide the proof related to transaction Tr_m. If Tr_m is an unspent transaction, the recipient of Tr_m will supply the necessary proof. If Tr_m is a spent transaction, EV_j will only be interested in Tr_m if EV_j is also interested in an unspent transaction Tr_n, and the validity of Tr_m is essential for validating Tr_n. Since $P(Tr_m) \in P(Tr_n)$, EV_j can obtain the proof of Tr_n from the recipient of Tr_n. The proof of a transaction Tr_m can always be gathered reliably and efficiently.Efficiency implies that the collection process involves straightforward point-to-point communication, eliminating the need for complex reliable broadcast mechanisms to handle malicious activities [34], while our model assumes that the recipient of an unspent transaction is incentivized to provide the correct proof, the requester will not accept it without independent verification. If $Verify(P(Tr_m)) = success$, it indicates that $P(Tr_m)$ is a valid proof for transaction Tr_m, as the algorithm directly implements the definition of a validity proof. The transaction verification algorithm is shown in Algorithm 3.

Algorithm 3: Transaction Verification.

if $Ver(\mathcal{P}(tx_i)) \neq pass$ **then**
  return unknown
end
if $\sum(all\ remaining\ values\ from\ Origin_m) \neq tv_m + rv_m$ **then**
  return unknown
end
for $B_{u,m}, m = [1 : k]$ **do**
  for *All transactions* Tr_m *in* $B_{j,q}$ **do**
    if $Source_n \cap Source_m \neq \emptyset$ *and* $Tr_m \neq Tr_n$ **then**
      return unknow
    end
  end
end
for *all transactions* Tr_n *in* $Source_n$ **do**
  if $Ver(Tr_n, P(Tr_n)) \neq valid$ **then**
    return unknow
  end
end
return *valid*

IV. Performance Evaluation

We employ VIBES [35] and BlockEmulator [31] to validate the methodologies proposed in this paper, where VIBES is utilized for verifying consensus and data transmission efficiency among nodes within shards, and BlockEmulator is applied to assess the transaction efficiency across shards. The dataset employed for this validation comprises publicly available electricity trading and Ethereum datasets.

A. In-DNS Performance

We utilize two key metrics—blockchain length and block generation interval—to evaluate the performance of the proposed method, which accounts for the time required for random redistribution, against several established consensus mechanisms. Portions of the experimental data are sourced from our earlier studies [11].

The experimental results are shown in Figs. 4 and 5. The method employing improved PBFT (Reference [6]) shows an downward trend, indicating longer blockchain lengths as the number of EVs grows. This may indicate a robust system with potential scalability limitations. The "PoW" exhibits relative stability, implying consistent blockchain length growth irrespective of the number of EVs. IDNS, built on Hashgraph with periodic VRF-based node reconfiguration, shows a declining trend. While reconfiguration introduces some time overhead, the overall blockchain length remains shorter than other methods, suggesting an efficient consensus process despite the re-election time. The classical PBFT approach demonstrates a steady reduction in blockchain length as EV numbers rise, potentially signaling diminished efficiency or security with scalability. Conversely, Hashgraph maintains a stable and extended blockchain length across varying EV counts, underscoring its efficiency and scalability. IDNS effectively balances scalability and consensus efficiency, even with reconfiguration, outperforming traditional consensus mechanisms in sustaining blockchain length as node numbers grow.

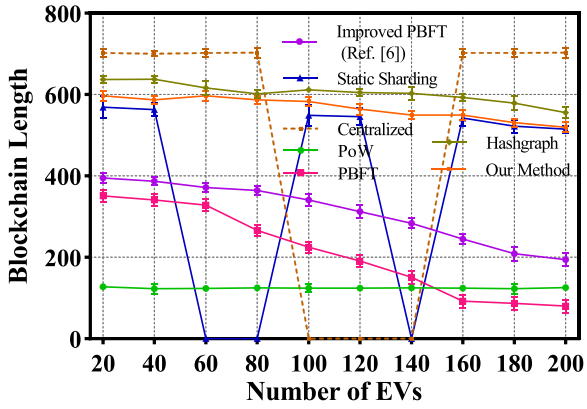

Fig. 4. The blockchain length.

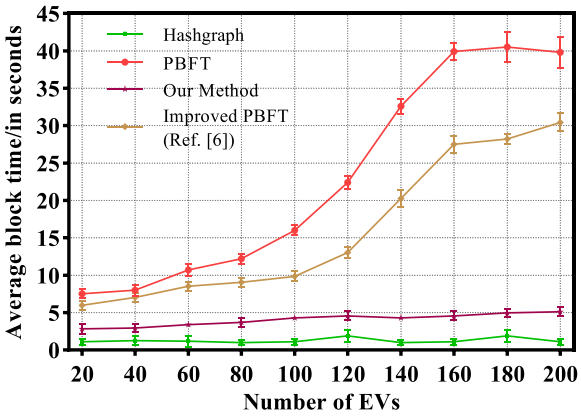

Fig. 5. The average block time.

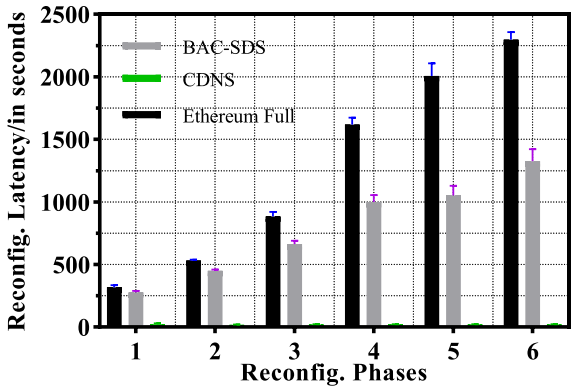

Fig. 6. The reconfiguration latency.

Hashgraph consistently achieves the lowest average block time, remaining stable regardless of EV count, highlighting its efficiency in block generation and scalability. PBFT's block time increases with more EVs, indicating strong performance in smaller networks but challenges as the network scales. IDNS, a Hashgraph variant with VRF-based reconfiguration, shows a slight upward trend but retains a lower average block time than standard PBFT. This implies that while reconfiguration introduces delays, Hashgraph's consensus efficiency remains largely unaffected. Improved PBFT (Reference [6]) experiences a sharp rise in block time with additional EVs, suggesting scalability constraints as consensus complexity grows with network size. The key reason behind this performance is the use of a DAG-based parallel ledger model and shard-level committee consensus, which enable concurrent block production and reduce the need for global coordination. Furthermore, our method integrates VRF-based committee selection to periodically reshuffle nodes, ensuring balanced participation and preventing overload in specific shards. These techniques allow our system to maintain fast block generation even in large-scale networks.

We have conducted comparative experiments between DNS and traditional static network sharding. We have introduced an additional experimental variable to simulate dynamic node behavior based on the proposed consensus mechanism, specifically modeling frequent node joining and leaving. The objective is to compare the blockchain growth curves of the proposed DNS scheme and traditional static sharding under dynamic network conditions to evaluate their stability. The traditional static sharding approach exhibits significant volatility in blockchain length under the same conditions. Without the ability to reconfigure shards in response to node dynamics, static sharding suffers from unbalanced workloads and potential quorum failures, particularly in shards that lose too many nodes. As a result, block generation becomes inconsistent, leading to periods of degraded throughput or even temporary consensus stalls. Static sharding assumes a fixed network topology, which is not realistic for V2V energy trading where EVs frequently join, leave, or relocate. DNS incorporates a DAG-based in-shard consensus with parallel ledger-keeping and committee-based validation, significantly reducing block generation latency. DNS periodically reconfigures shard composition using VRF, ensuring security while maintaining optimal load distribution. Unlike static sharding, which is unsuitable for latency-sensitive and geographically dispersed V2V scenarios, DNS is specifically tailored for such applications.

B. *Cross-DNS Performance*

Fig. 6 indicates that the latency for reconfiguring in the CDNS method remains relatively low and stable across different reconfiguration phases. This contrasts with BAC-SDS [11] and Ethereum Full sync, where latency increases with each phase. BAC-SDS's approach assigns data synchronization tasks to a committee, reducing the workload compared to Ethereum's full sync, where all nodes participate, resulting in lower latency. CDNS's efficiency stems from only transmitting the states of active accounts corresponding to the most recent phase, which keeps the volume of data transferred and the associated latency consistently low during reconfiguration. In Ethereum, "full sync" is a node synchronization process where every transaction and contract is downloaded and verified from the genesis block to the current state. This thorough method ensures complete accuracy of the blockchain's history, making it ideal for services that need the entire transaction history, like wallets and auditors. However, it is resource-intensive, as it requires significant storage and processing power, and can be time-consuming.

CDNS maintains consistently low and stable latency, ranging between 500 ms and 1,000 ms, demonstrating its efficiency and scalability. In contrast, BAC-SDS starts at approximately 500 ms in Phase 1 but increases steadily to 1,500 ms by Phase 6, indicating a scalability bottleneck despite its committee-based approach. Ethereum Full Sync begins at 1,500 ms in Phase 1 and escalates sharply to 2,000 ms by Phase 6, reflecting the resource-intensive nature of downloading and verifying the entire blockchain history. CDNS achieves its efficiency by transmitting only the states of active accounts corresponding to the most recent phase, minimizing data volume and computational overhead. BAC-SDS, while more efficient than Ethereum Full Sync, still incurs growing overhead as the system scales, highlighting its limitations for dynamic environments. Ethereum Full Sync's high and escalating latency underscores its unsuitability

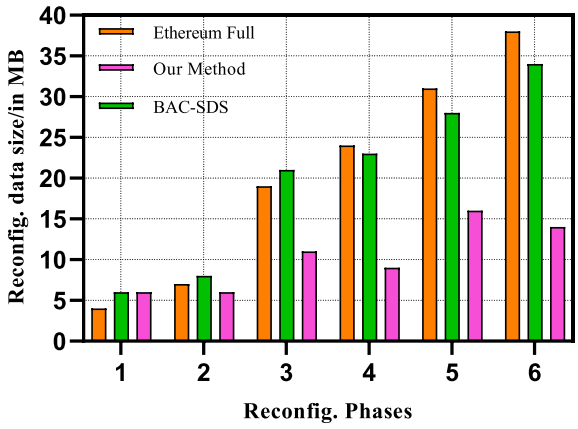

Fig. 7. The reconfiguration data size.

for real-time applications, making it more appropriate for tasks requiring complete historical accuracy, such as auditing.

Design implications include the use of dynamic sharding to maintain scalability and efficiency, as demonstrated by CDNS's performance. BAC-SDS could benefit from adaptive committee sizing or hierarchical synchronization to improve scalability, while Ethereum Full Sync might explore hybrid approaches combining full sync with selective state pruning to reduce latency. Future directions could focus on proximity-aware synchronization to prioritize nearby nodes, adaptive reconfiguration to adjust frequency and data volume based on network load, and hardware acceleration (e.g., GPU-based solutions) to reduce computational overhead. CDNS's low and stable latency across phases validates its suitability for real-time, latency-sensitive applications like V2V energy trading. By outperforming BAC-SDS and Ethereum Full Sync, CDNS demonstrates its efficiency and scalability, making it an ideal solution for dynamic networks. Future enhancements should focus on adaptive protocols and hardware acceleration to further optimize performance in large-scale deployments.

Fig. 7 shows the data size in megabytes required for reconfiguration across different phases. Ethereum Full sync shows a consistent upward trend in data size, indicating an increase in data requirements with each phase. Our Method seems to show variability, but generally, the data size appears to grow as the number of reconfiguration phases increases. BAC-SDS also shows an upward trend but with less steepness compared to Ethereum Full, suggesting it requires less data for reconfiguration.

CDNS's ability to maintain low and stable data size across phases demonstrates its suitability for large-scale, dynamic networks like V2V energy trading, where frequent reconfigurations are necessary. By focusing on active accounts, CDNS avoids the overhead of processing historical data, making it more efficient than BAC-SDS and Ethereum Full Sync. BAC-SDS's increasing data size suggests that committee-based approaches need further optimization, such as adaptive committee sizing or hierarchical synchronization, to improve scalability. Ethereum Full Sync's high and escalating data size underscores its unsuitability for real-time applications, making it more appropriate for tasks requiring complete historical accuracy, such as auditing or archival services.

C. Security Analysis of Dynamic Node Sharding

To optimize the performance of the sharding system while ensuring the security of each shard, we determine an appropriate reconfiguration interval by estimating the time until a single shard fails [31]. The emergence of malicious nodes within a shard is unpredictable and cannot be easily modeled using straightforward approaches. In this study, we treat the transformation of a honest node into a malicious one as a stochastic event to model the occurrence of malicious nodes in the shard. Let n represent the number of nodes in a single shard. The shard will fail when the count of malicious nodes surpasses $n/3$. The arrival time of the r^{th} malicious node is denoted as W_r. Given that the average arrival rate of malicious nodes is μ, the probability of the $n/3^{\text{th}}$ malicious node appearing at time t can be computed as:

$$f_{W_{\lfloor \frac{n}{3} \rfloor}}(t) = \frac{\mu^{\lfloor \frac{n}{3} \rfloor} t^{\lfloor \frac{n}{3} \rfloor - 1} e^{-\mu t}}{\left(\lfloor \frac{n}{3} \rfloor - 1\right)!} \tag{3}$$

Consider a particular shard composed of n nodes, with the arrival rate of malicious nodes being μ. The anticipated time until failure is:

$$E\left(W_{\lfloor \frac{n}{3} \rfloor}\right) = \frac{\lfloor \frac{n}{3} \rfloor}{\mu} \tag{4}$$

By calculating the expected value of W_n, we can determine the expected value of $W_{\lfloor \frac{n}{3} \rfloor}$:

$$E(W_n) = \int_0^\infty t \cdot \frac{\mu^n t^{n-1} e^{-\mu t}}{(n-1)!}\, dt \tag{5}$$

By substituting μt with θ, we obtain $d\theta = \mu dt$. Consequently, the expected value of $W_{\lfloor \frac{n}{3} \rfloor}$ can be computed as:

$$E(W_n) = \int_0^\infty \frac{1}{\mu} \cdot \frac{\theta^n e^{-\theta}}{(n-1)!}\, d\theta = \frac{1}{\mu(n-1)!} \cdot n! \int_0^\infty e^{-\theta}\, d\theta = \frac{n}{\mu} \tag{6}$$

Therefore,

$$E(W_{\lfloor \frac{n}{3} \rfloor}) = \frac{\lfloor \frac{n}{3} \rfloor}{\mu} \tag{7}$$

We can establish a suitable reconfiguration interval based on practical implementation, thereby maintaining the security of the sharded blockchain system.

D. Shard and Block Size Impact

The experimental results provide a comprehensive evaluation of the system's performance, focusing on message propagation latency, throughput, and the impact of block size and committee size, as shown in Figs. 8 and 9. The message propagation latency is analyzed across different committee sizes. The maximum latency increases from 350 ms (100 nodes) to 550 ms (250 nodes), reflecting the inherent challenge of synchronizing larger committees in a P2P network. The average latency rises

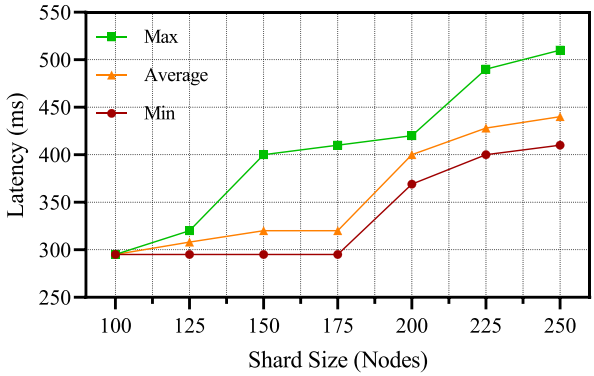

Fig. 8. Shard size impact.

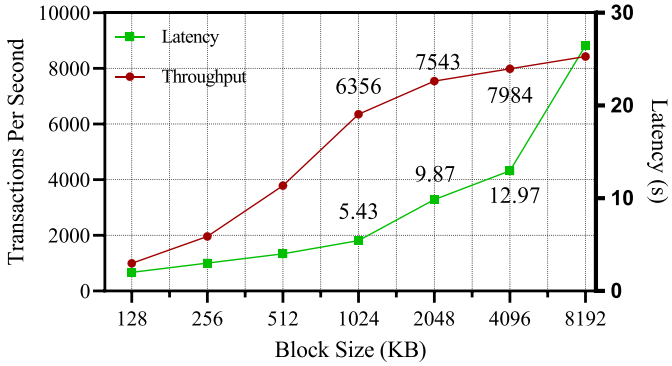

Fig. 9. Block size impact.

from 250 ms to 450 ms, indicating a near-linear relationship between committee size and communication overhead, while the minimum latency grows from 150 ms to250 ms, suggesting that even under ideal conditions, network dynamics impose a baseline delay. The conservative δ value of 600 ms (for 250 nodes) aligns with the observed maximum latency (550 ms), ensuring synchronous rounds remain feasible even at scale. This design choice prioritizes safety over liveness to prevent consensus failures due to delayed messages. However, the latency increase with committee size highlights the trade-off between fault tolerance (via redundancy) and communication efficiency. Larger committees introduce higher propagation delays, limiting real-time transaction finality, which underscores the need for dynamic sharding to cap committee sizes (e.g., $\leq$250 nodes). The linear latency trend suggests a lack of hierarchical routing or optimized peer selection, indicating that future work could explore proximity-aware gossip protocols to mitigate this. Our system ensures safe and timely message propagation. The average latency remains within acceptable bounds (around 400 ms), even for larger committees, thanks to our optimized second-order gossip protocol and shard-local consensus, which significantly reduce broadcast delays. Moreover, the use of localized message propagation and chunked data encoding further limits latency in larger shards.

The throughput peaks at 7,000tx/sec for a 2,048 KB block size, leveraging batch processing efficiency, while the latency stabilizes at 8.7 s for 2,048 KB blocks, balancing throughput gains with real-time requirements. The choice of 2,048 KB balances the trade-off between throughput (with larger blocks) and latency (due to chunk encoding/decoding overhead). Smaller blocks (e.g., 512 KB) reduce latency but sacrifice throughput, while larger blocks (e.g., 8,192 KB) introduce excessive delays ($>$15 s). The 128-way chunking of 2MB blocks ensures resilience to node failures but adds computational overhead during encoding/decoding, contributing to latency. The target latency of $<$10 sec aligns with mainstream payment systems, making the 2,048 KB block size pragmatically viable for V2V energy trading.

The synthesis of findings reveals a critical trade-off between committee size and block size. Smaller committees ($<$150 nodes) achieve lower latency ($<$300 ms) but limit scalability, while larger committees (250 nodes) enable broader participation but require optimized protocols (e.g., hierarchical gossip) to mitigate latency. Larger block sizes (2,048 KB) maximize throughput but demand efficient chunking/encoding to avoid latency spikes. Design implications include capping shard sizes to $\leq$250 nodes to align with the $\delta = 600$ ms threshold, ensuring consensus stability, and combining erasure codes with lightweight hashing (e.g., Merkle trees) for smaller, latency-sensitive transactions. Future directions could focus on integrating geolocation data to prioritize nearby nodes in message propagation, reducing average latency, and implementing runtime block size adjustments based on network load (e.g., smaller blocks during peak hours). Larger block sizes (e.g., 2048 KB) enable our system to process more transactions per second (peaking at 7543 TPS), while maintaining latency below the target threshold of 10 seconds. This efficiency results from our batch processing design and the use of deterministic nano blocks to encapsulate only essential transaction metadata, reducing consensus overhead. We also apply state pruning and selective data retention during shard reconfiguration, ensuring that block validation does not require full history download, which would otherwise increase delay. Notably, even at 4096 KB block size, our latency remains significantly lower than many static sharding methods, demonstrating the DNS framework's effectiveness in minimizing data-related delay.

The experimental results validate the DNS strategy's efficacy in balancing scalability, latency, and security. By coupling dynamic sharding with optimized block sizing and erasure coding, the system achieves 7,000tx/sec throughput and sub-10 sec latency, meeting the demands of decentralized V2V energy trading. Future enhancements should focus on adaptive protocols to further reduce latency in large-scale deployments.

V. Conclusion and the Road Ahead

This paper presents a DNS strategy for V2V energy trading, leveraging blockchain to enhance scalability and security. Our method significantly improves transaction throughput and reduces latency, especially in cross-shard transactions. Rigorous simulations on Hyperledger Fabric and BlockEmulator validate these benefits, demonstrating superior performance compared to traditional methods. The DNS strategy also strengthens security

against shard boundary vulnerabilities, ensuring reliable and efficient energy trading. Future work will refine this architecture for broader smart grid applications, with continuous improvements anticipated from ongoing tool development.

ACKNOWLEDGMENT

Our deepest gratitude goes to the anonymous reviewers. A preliminary version of the paper was presented in the 2024 IEEE Industry Applications Society Annual Meeting titled "Optimizing V2V Energy Trading via Blockchain Sharding in Decentralized Automotive Networks".

REFERENCES

[1] Y. Wang et al., "A fast and secured vehicle-to-vehicle energy trading based on blockchain consensus in the Internet of Electric Vehicles," *IEEE Trans. Veh. Technol.*, vol. 72, no. 6, pp. 7827–7843, Jun. 2023.

[2] Z. Zhang, H. Tang, J. Ren, Q. Huang, and W.-J. Lee, "Strategic prosumers-based peer-to-peer energy market design for community microgrids," *IEEE Trans. Ind. Appl.*, vol. 57, no. 3, pp. 2048–2057, May/Jun. 2021.

[3] H. Ma et al., "Optimal peer-to-peer energy transaction of distributed prosumers in high-penetrated renewable distribution systems," *IEEE Trans. Ind. Appl.*, vol. 60, no. 3, pp. 4622–4632, May/Jun. 2024.

[4] A. Esmat, M. D. Vos, Y. Ghiassi-Farrokhfal, P. Palensky, and D. Epema, "A novel decentralized platform for peer-to-peer energy trading market with blockchain technology," *Appl. Energy*, vol. 282, 2021, Art. no. 116123.

[5] Y. Wang et al., "An efficient, secured, and infinitely scalable consensus mechanism for peer-to-peer energy trading blockchain," *IEEE Trans. Ind. Appl.*, vol. 59, no. 5, pp. 5215–5229, Sep./Oct. 2023.

[6] H. N. Abishu, A. M. Seid, Y. H. Yacob, T. Ayall, G. Sun, and G. Liu, "Consensus mechanism for blockchain-enabled vehicle-to-vehicle energy trading in the Internet of Electric Vehicles," *IEEE Trans. Veh. Technol.*, vol. 71, no. 1, pp. 946–960, Jan. 2022.

[7] Z. Ding, F. Teng, P. Sarikprueck, and Z. Hu, "Technical review on advanced approaches for electric vehicle charging demand management, part II: Applications in transportation system coordination and infrastructure planning," *IEEE Trans. Ind. Appl.*, vol. 56, no. 5, pp. 5695–5703, Sep./Oct. 2020.

[8] M. Kim, J. Lee, J. Oh, K. Park, Y. Park, and K. Park, "Blockchain based energy trading scheme for vehicle-to-vehicle using decentralized identifiers," *Appl. Energy*, vol. 322, 2022, Art. no. 119445.

[9] M. Baza et al., "Privacy-preserving blockchain-based energy trading schemes for electric vehicles," *IEEE Trans. Veh. Technol.*, vol. 70, no. 9, pp. 9369–9384, Sep. 2021.

[10] Z. Liu, Y. Xu, C. Zhang, H. Elahi, and X. Zhou, "A blockchain-based trustworthy collaborative power trading scheme for 5G-enabled social Internet of Vehicles," *Digit. Commun. Netw.*, vol. 8, no. 6, pp. 976–983, 2022.

[11] Y. Wang et al., "Enhancing power grid resilience with blockchain-enabled vehicle-to-vehicle energy trading in renewable energy integration," *IEEE Trans. Ind. Appl.*, vol. 60, no. 2, pp. 2037–2052, Mar./Apr. 2024.

[12] S. Chen et al., "A blockchain consensus mechanism that uses proof of solution to optimize energy dispatch and trading," *Nature Energy*, vol. 7, no. 6, pp. 495–502, 2022.

[13] Y. Wang et al., "Optimizing V2V energy trading blockchain sharding in decentralized automotive networks," in *Proc. 2024 IEEE IAS Annu. Meeting*, 2024, pp. 1–8.

[14] V. K. Saini, C. S. Purohit, R. Kumar, and A. S. Al-Sumaiti, "Proof of work consensus based peer to peer energy trading in the indian residential community," *Energies*, vol. 16, no. 3, 2023, Art. no. 1253.

[15] J. Yang, A. Paudel, H. B. Gooi, and H. D. Nguyen, "A proof-of-stake public blockchain based pricing scheme for peer-to-peer energy trading," *Appl. Energy*, vol. 298, 2021, Art. no. 117154.

[16] A. Sheikh, V. Kamuni, A. Urooj, S. Wagh, N. Singh, and D. Patel, "Secured energy trading using Byzantine-based blockchain consensus," *IEEE Access*, vol. 8, pp. 8554–8571, 2020.

[17] L. Feng, H. Zhang, Y. Chen, and L. Lou, "Scalable dynamic multi-agent practical byzantine fault-tolerant consensus in permissioned blockchain," *Appl. Sci.*, vol. 8, no. 10, 2018, Art. no. 1919.

[18] Y. Wang et al., "Study of blockchains's consensus mechanism based on credit," *IEEE Access*, vol. 7, pp. 10224–10 231, 2019.

[19] D. Wang and L. Wang, "Consensus mechanism of multi-energy interactive subject based on practical Byzantine fault tolerance algorithm," *Automat. Electric Power Syst.*, vol. 43, no. 09, pp. 41–49, 2019.

[20] W. Cai, W. Jiang, K. Xie, Y. Zhu, Y. Liu, and T. Shen, "Dynamic reputation–based consensus mechanism: Real-time transactions for energy blockchain," *Int. J. Distrib. Sensor Netw.*, vol. 16, no. 3, 2020, Art. no. 1550147720907335.

[21] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proc. 2016 ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 17–30.

[22] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "Omniledger: A secure, scale-out, decentralized ledger via sharding," in *Proc. 2018 IEEE Symp. Secur. Privacy*, 2018, pp. 583–598.

[23] J. Wang and H. Wang, "Monoxide: Scale out blockchains with asynchronous consensus zones," in *Proc. 16th USENIX Symp. Networked Syst. Des. Implementation*, 2019, pp. 95–112.

[24] M. Zamani, M. Movahedi, and M. Raykova, "Rapidchain: Scaling blockchain via full sharding," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2018, pp. 931–948.

[25] H. Han, S. Chen, Z. Xu, X. Dong, and J. Zeng, "Trust management scheme of iov based on dynamic sharding blockchain," *Electronics*, vol. 13, no. 6, 2024, Art. no. 1016.

[26] V. S. Naresh, V. D. Allavarpu, and S. Reddi, "Blockchain iota sharding-based scalable secure group communication in large vanets," *IEEE Internet Things J.*, vol. 10, no. 6, pp. 5205–5213, Mar. 2023.

[27] Y. Yu, G.-P. Liu, Y. Huang, C. Y. Chung, and Y.-Z. Li, "A blockchain consensus mechanism for real-time regulation of renewable energy power systems," *Nature Commun.*, vol. 15, no. 1, 2024, Art. no. 10620.

[28] M. Andoni et al., "Blockchain technology in the energy sector: A systematic review of challenges and opportunities," *Renewable Sustain. Energy Rev.*, vol. 100, pp. 143–174, 2019.

[29] F. Knirsch, A. Unterweger, and D. Engel, "Privacy-preserving blockchain-based electric vehicle charging with dynamic tariff decisions," *Comput. Sci.- Res. Develop.*, vol. 33, no. 1, pp. 71–79, 2018.

[30] S. Wang, A. F. Taha, J. Wang, K. Kvaternik, and A. Hahn, "Energy crowdsourcing and peer-to-peer energy trading in blockchain-enabled smart grids," *IEEE Trans. Syst., Man, Cybern. Syst.*, vol. 49, no. 8, pp. 1612–1623, Aug. 2019.

[31] H. Huang, Y. Zhao, and Z. Zheng, "tmpt: Reconfiguration across blockchain shards via trimmed merkle patricia trie," in *Proc. IEEE/ACM 31st Int. Symp. Qual. Serv.*, 2023, pp. 1–10.

[32] S. Micali, M. Rabin, and S. Vadhan, "Verifiable random functions," in *Proc. 40th Annu. Symp. Found. Comput. Sci.*, 1999, pp. 120–130.

[33] Z. Ren, K. Cong, T. Aerts, B. D. Jonge, A. Morais, and Z. Erkin, "A scale-out blockchain for value transfer with spontaneous sharding," in *Proc. 2018 IEEE Crypto Valley Conf. Blockchain Technol.*, 2018, pp. 1–10.

[34] Z. Ren, K. Cong, J. Pouwelse, and Z. Erkin, "Implicit consensus: Blockchain with unbounded throughput," 2017, *arXiv:1705.11046*.

[35] L. Stoykov, K. Zhang, and H.-A. Jacobsen, "Vibes: Fast blockchain simulations for large-scale peer-to-peer networks," in *Proc. 18th ACM/IFIP/USENIX Middleware Conf.: Posters Demos*, 2017, pp. 19–20.