VRF-PoW: Proof of Work Consensus With Verifiable Random Function

Seungmin Kim, Haeung Choi, Minho Yoon, and Heung-No Lee, *Senior Member, IEEE*

***Abstract*—Proof of Work (PoW) is a fundamental and widely adopted consensus mechanism in blockchain systems that enables effective consensus in permissionless environments, solely through block propagation without additional message exchange between nodes. However, PoW can result in excessive energy consumption and network centralization. This paper proposes VRF-PoW, a novel consensus protocol that reduces energy consumption while preserving the decentralized and permissionless nature of PoW. The VRF-PoW method integrates Verifiable Random Functions (VRFs) into the PoW process, reducing network energy consumption. To prevent Sybil attacks, this work introduces a pre-PoW phase. Further, this work validates VRF-PoW via a theoretical analysis and numerical evaluation. The results demonstrate that VRF-PoW significantly reduces energy consumption compared to the traditional PoW, while maintaining decentralization, stable block generation, and network liveness.**

***Index Terms*—Blockchain, consensus algorithm, proof of work, verifiable random function, energy efficiency.**

I. Introduction

BLOCKCHAIN systems, such as Bitcoin [1], use Proof of Work (PoW) as a consensus mechanism in which the first node to solve a computational puzzle becomes the block proposer [2]. Blockchain data are organized into blocks, generated at regular intervals. To maintain a stable average block time, the puzzle difficulty is dynamically adjusted based on the total computational power of the network [3]. Increasing difficulty improves security by preventing the monopolization of block creation, but it leads to high energy consumption [4], [5], [6]. As of May 2025, Bitcoin's annual electricity consumption was estimated at 189.31 TWh, comparable to that of South Africa [7]. Moreover, PoW rewards only the miner that produces a valid block, causing unstable income for those with limited resources. In response, mining pools emerged to allow participants to share computational resources and split rewards proportionally [8]. However, when large pools gain excessive hash power, they can dominate block production, undermining decentralization [9].

To address the high energy consumption of PoW, Proof of Stake (PoS) has been proposed [10], which probabilistically selects block producers based on their stake [11], [12], [13], [14]. As block producers do not need to solve computational puzzles, PoS significantly reduces energy consumption. However, it introduces challenges, such as the nothing-at-stake problem and the tendency to concentrate wealth [15], [16], [17], [18]. Several alternatives to PoS have been proposed [19], [20], [21], [22], but they continue to suffer from network centralization. As an alternative, researchers have explored enhancements to the PoW mechanism. Proof-of-useful-work (PoUW) mechanisms repurpose mining computations for meaningful tasks, such as scientific or artificial intelligence (AI) problems [23], [24], [25], but face limitations due to their reliance on trusted execution environments. Structural PoW schemes aim to improve energy efficiency and fairness by modifying the traditional PoW [26], [27], [28], [29], but they face challenges, such as increased communication overhead and delayed block generation.

A verifiable random function (VRF) is a cryptographic primitive that enables a prover to generate pseudorandom output along with proof that the output was correctly computed using a secret key [30]. The verifiability of the VRF ensures that randomness can be reliably generated in trustless blockchain environments [31]. Primarily, the VRF has been adopted in PoS schemes to select block proposers randomly [32], [33], [34]. Recently, VRF-based PoW schemes have been proposed to improve fairness and decentralization by incorporating verifiable randomness into the miner selection process [35]. However, these schemes often suffer from high communication overhead. This work aims to reduce energy consumption while preserving the core advantages of PoW, including decentralization and minimal communication overhead.

This work introduces a novel consensus mechanism, the VRF-PoW, that integrates the VRF with the traditional PoW. At each block interval, miners run a VRF to determine probabilistically whether they are eligible to join the mining process. The VRF-PoW mechanism introduces a pre-PoW phase to integrate VRF into the PoW environment without requiring an additional message exchange. By allowing a select subset of miners to mine blocks, VRF-PoW reduces energy consumption while preserving essential properties, including permissionless participation and non-interactive consensus. The contributions of this research are as follows:

- *Verifiable Random Function Proof of Work:* This work proposes VRF-PoW, a novel consensus mechanism that

Received 27 June 2025; revised 19 December 2025; accepted 19 January 2026. Date of publication 23 January 2026; date of current version 4 February 2026. This work was supported in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) funded by the Korea Government (MSIT) under Grant IITP-2025-RS-2021-II210118, (Development of decentralized consensus composition technology for large-scale nodes), and in part by the IITP-Information Technology Research Center funded by the Korea Government (MSIT) under Grant IITP-2026-RS-2021-II211835. Recommended for acceptance by Dr. Liehuang Zhu. *(Corresponding author: Heung-No Lee.)*

The authors are with the School of Electrical Engineering and Computer Science, Gwangju Institute of Science and Technology, Gwangju 61005, South Korea (e-mail: seungminkim@gm.gist.ac.kr; haeung@gist.ac.kr; minho.yoon@gm.gist.ac.kr; heungno@gist.ac.kr).

Digital Object Identifier 10.1109/TNSE.2026.3657405

reduces energy consumption by enabling nodes to self-select for block generation through a VRF. To mitigate VRF-based Sybil attacks, VRF-PoW introduces a pre-PoW phase that requires miners to perform PoW before executing the VRF.

- *Signature-based Mining and Progressive Timeout Mechanism:* A signature-based mining process is adopted to discourage the sharing of answer hashes and the formation of mining pools. The progressive timeout mechanism ensures continuous block generation even when no miner is immediately self-selected, while preventing abrupt changes in network hash power.
- *Mathematical Modeling and Numerical Analysis:* This work presents a formal mathematical model developed to analyze energy consumption, block-time distribution, and Sybil attack resistance. Numerical simulations validate the analytical model, demonstrating that VRF-PoW consumes significantly less energy than the traditional PoW. The progressive timeout mechanism results in a more stable block time distribution compared to a fixed timeout.

The remainder of the paper is organized as follows. Section II reviews the related work on consensus mechanisms and their limitations. Next, Section III introduces the fundamental concepts of PoW and VRF. Then, Section IV details the VRF-PoW design and its energy-efficient, decentralized operation, and Section V provides a theoretical analysis of VRF-PoW, including probabilistic modeling and Sybil attack resistance. Section VI analyzes the security of the VRF-PoW. Further, Section VII reports numerical simulations that validate the analytical results and compares VRF-PoW with traditional PoW. Finally, Section VIII concludes the paper and outlines directions for future research.

II. Related Work

This section provides an overview of research on consensus protocols designed to address critical challenges in blockchain systems, including energy consumption and decentralization. These studies fall into three broad categories: PoS and its variants, which replace PoW with alternative selection mechanisms; PoUW mechanisms, which repurpose computational work for useful tasks; and structural PoW extensions that aim to enhance efficiency and decentralization. In addition, this section also examines recent efforts that incorporate VRF into consensus protocols.

A. Proof of Stake and its Variants

The PoS protocols can be broadly classified into chain-based and Byzantine Fault-Tolerance (BFT)-based approaches [36], [37]. Chain-based PoS uses a longest-chain rule with pseudorandom leader election and probabilistic finality [38], [39], [40]. However, BFT-based PoS employs a validator committee running BFT consensus to finalize blocks deterministically [12], [13], [14]. Although PoS improves energy efficiency compared to PoW, it introduces vulnerabilities, such as the nothing-at-stake problem, in which validators can exercise their stake across all forked chains [15], [16], [17]. To mitigate this problem, PoS systems often employ checkpoints; however, these checkpoints introduce external trust assumptions that compromise decentralization [41]. To address these limitations, several PoS-like alternatives have been proposed. Delegated PoS (DPoS) [42] elects a committee of validators using periodic voting by nodes that hold a sufficient stake. However, DPoS relies on a fixed-size committee, making it less adaptable to dynamic changes in the network scale. Qualified PoS (QPoS) [43] addresses centralization by incentivizing honest behavior and penalizing dishonest actions from voters and leaders. Proof of activity [19] combines PoW and PoS to address energy consumption and security concerns. Proof of contribution (PoC) [20] determines the stake based on the miner's honesty and adjusts the mining difficulty accordingly, encouraging honest behavior. Proof of reputation (PoR) [21] selects the node with the highest reputation as the leader to generate a block. However, these schemes restrict permissionless participation, limiting accessibility for new or resource-constrained participants [22].

B. Proof of Useful Work

Several alternative schemes have been proposed to redirect computational resources toward solving practical problems. Primecoin [44] introduced a PoW mechanism based on searching for chains of prime numbers. Miners search for prime numbers, generating additional potential scientific value during the mining process. Proof of exercise (PoX) [45] and proof of useful work (PoUW) [23] require miners to solve matrix-based scientific problems submitted by external sources. Moreover, PoUW extends the original concept by incorporating real-world nondeterministic polynomial hard (NP-hard) problems and demonstrates its feasibility through mathematical optimization modeling [24]. The REM [46] and proof of elapsed time (PoET) [47] rely on trusted hardware to perform useful tasks or remain idle, rather than executing the traditional PoW.

Recently, several studies have explored integrating AI training tasks into PoW mechanisms [25], [48], [49], [50]. For example, proof of federated learning (PoFL) [25], [50] redeploys PoW energy into federated learning by enabling distributed nodes to train AI models collaboratively. The proof of necessary work (PoNW) uses the energy that would be wasted generating valid proofs to perform useful blockchain verification computations. Despite their advantages in utility and energy efficiency, useful PoW schemes face scalability and deployment challenges due to their dependence on trusted hardware or centralized coordination. Compared to the simple hash functions in traditional PoW, integrating complex mathematical or AI-related tasks may introduce substantial computational overhead.

C. Structural Extensions of Proof of Work

Several studies have proposed structural extensions to the traditional PoW to reduce energy consumption and alleviate mining centralization. Sprints [26] applied trusted verifiable-delay-function (VDF) hardware to ensure uniform delays across miners, enabling intermittent mining. Sprints introduced proof of delay (PoD), which requires miners to prove they have waited a specific period before initiating mining. As PoD imposes minimal computational overhead, miners alternate between low-power PoD and high-power PoW. Sprints reports that allocating

5% of the block generation interval to PoW results in an energy reduction of approximately 92%. However, the practicality of PoD is limited by the lack of implementation of VDF that guarantees consistent delays across miners. The energy-efficient PoW (EPoW) and decentralized PoW (DPoW) [28], [51] suppress mining centralization by redesigning the reward mechanism to distribute the block reward between the first and second proposers. Moreover, EPoW allocates the entire block reward to the second proposer, whereas the first proposer receives only transaction fees. In contrast, DPoW dynamically distributes the reward between the two proposers. The protocol promotes lower hash rates in subsequent blocks to reduce mining costs, although this may compromise the overall blockchain security. Green-PoW [27] reduces energy consumption by up to 50% through merging the mining of two consecutive blocks into a single round. Only miners that submit valid but nonwinning blocks for the first block are eligible to compete for the second, limiting unnecessary competition. However, if no miners are eligible for the second block, a timeout may occur, delaying block generation. To overcome the risk of block delays, VRF-PoW introduces a progressive timeout mechanism that maintains stable block generation.

Several studies have explored two-phase PoW (2P-PoW) [52] mechanisms to enhance decentralization. The 2P-PoW mechanism was initially introduced, incorporating a cryptographic signing step into traditional PoW using private keys. In 2P-PoW, miners solve a standard hash-based PoW puzzle followed by a signature-based step in which the solution is signed with the private key. Only the private key holder can claim the block reward; hence, mining pool formation is discouraged as key sharing becomes economically unattractive. PieceWork [29] builds on the two-phase structure by formalizing the sign-to-mine mechanism to address block withholding attacks. Moreover, VRF-PoW extends the 2P-PoW approach by incorporating private-key signatures into the eligibility selection and block-generation processes.

D. VRF-Based Consensus Mechanisms

First introduced by Micali et al. [30], VRFs have since been standardized for use in decentralized systems [57], [58], [59] and are most commonly used in the consensus protocols of PoS blockchains [32], [33], [34], [60], [61], [62], [63]. Nodes are elected to the consensus committee via a VRF, with a selection probability that is proportional to the number of tokens held. Several recent studies have explored integrating the VRF into consensus mechanisms to enhance fairness, efficiency, and decentralization [35], [54], [55], [56]. In VRF-based mining, miners repeatedly execute a VRF bound to their private key rather than brute-force hash trials, deterring mining outsourcing by risking private key exposure [54]. However, this implementation merely employs the VRF output as a substitute for the hash function and does not employ it for miner selection. Proof of random leader (PoRL) [55] is a proof of authority-based consensus protocol that employs the VRF for leader selection in permissioned blockchains, enhancing randomness and throughput compared to Aura and Clique. The PoW with random selection (PoWR) [35] variant employs verifiable random selection to improve energy efficiency and prevent hash monopolization by application-specific integrated circuit (ASIC) miners. PoWR limits competition through randomized miner participation, reducing the energy consumption of ASIC-based mining by approximately 60–85%. However, PoWR requires additional consensus overhead, such as managing winner node tables and increased internode message exchange. Proof of verifiable functions (PoVF) proposes a fairness-oriented consensus mechanism that applies the verifiability and unpredictability of verifiable functions to select consensus participants randomly [56]. However, its implementation relies on trusted hardware to ensure synchronized timing for the VDF.

To highlight the distinctions between the VRF-PoW and related work, Table I compares the major related studies across the following eight criteria:

- *Useful Computation:* This criterion evaluates whether mining resources are repurposed for tasks beyond hash computations, such as scientific or AI-related workloads.
- *Energy Efficiency:* This criterion examines whether the protocol reduces the overall energy consumption by limiting the number or duration of active participants or by employing non-PoW consensus mechanisms.
- *Trustless Operation:* This criterion assesses whether consensus can be achieved without relying on centralized authorities or trusted hardware components.
- *Outsourcing Resistance:* This criterion evaluates whether the protocol binds block eligibility to individual miners, structurally discouraging mining outsourcing and pool participation.
- *Low Communication Overhead:* This criterion determines whether the consensus proceeds without additional message exchanges beyond a single block propagation, minimizing latency, forks, and network congestion.
- *ASIC Resistance:* This criterion assesses whether the protocol design limits the computational advantage of ASICs or makes their implementation difficult.
- *Committee Selection:* This criterion examines whether consensus involves a subset of nodes forming a committee, within which block proposers are selected.
- *No Extra State:* This criterion evaluates whether the consensus protocol operates without maintaining additional local or global state information.

Although VRF-PoW does not perform useful work during mining, it applies VRFs to select a subset of miners as a consensus committee probabilistically, reducing the number of active participants and enhancing the overall energy efficiency. The VRF execution is performed locally in each miner's environment; hence, no additional message exchanges or external state information are required, resulting in low communication overhead and stateless operation. By incorporating a pre-PoW phase that imposes a computational cost on VRF execution, the protocol maintains a fully decentralized trustless operation without relying on centralized authorities or trusted hardware. Moreover, VRF-PoW employs a mining process tied to private key signatures, inhibiting mining outsourcing and pool formation. Finally, although this study describes VRF-PoW using a PoW primitive based on a hash function, the VRF-PoW design can be readily extended by substituting the hash function with

TABLE I
COMPREHENSIVE COMPARISON OF RELATED WORK

Category	Protocol	Useful Comp.	Energy Eff.	Trustless Op.	Outsourcing Res.	Low Comm.	ASIC*-Res.	Committee Sel.	No Extra State
Useful PoW	PoUW [23]	●	○	●	○	◐	○	○	◐
	PoFL [25]	●	○	●	○	○	○	○	○
	PoNW [53]	●	○	●	●	●	○	○	●
Structural PoW	Sprints [26]	○	●	○	○	●	○	○	●
	Green-PoW [27]	○	●	●	●	○	○	●	○
	EPoW [28]	○	●	●	●	○	○	○	○
	2P-PoW [29]	○	○	●	●	●	○	○	●
	PoET [47]	●	○	○	●	○	●	○	○
VRF-based Consensus	PoWR [35]	○	●	○	●	○	●	●	○
	VRF-based mining [54]	○	○	●	●	●	◐	○	●
	PoRL [55]	○	●	○	○	○	●	●	○
	PoVF [56]	○	●	○	○	○	●	●	○
	VRF-PoW (Ours)	○	●	●	●	●	◐	●	●

○: without this feature; ●: with this feature; ◐: conditional/achievable via design.
*: application-specific integrated circuit

ASIC-resistant primitives, enhancing hardware fairness and sustainability [64], [65], [66].

III. PRELIMINARIES

This section outlines the fundamental components and operational principles of PoW and VRF, which underpin VRF-PoW.

A. Proof of Work Blockchain

A PoW blockchain network comprises a set of nodes operating in a peer-to-peer environment, where any node can participate without permission. Each node maintains its own local copy of the blockchain, structured as a Merkle tree of transaction blocks and their headers. The metadata for each block include the hash of the parent block, linking blocks in chronological order. Due to network latency or forks, nodes may temporarily hold different versions of the blockchain. However, the longest valid chain is recognized as the canonical chain [67].

Nodes compete to generate blocks to be included in the main chain at each block height. In PoW-based blockchains, miners repeatedly execute a hash function on the block header with varying nonce values to determine a hash that satisfies the required difficulty condition. Each attempt is independent, and the process terminates once a hash value is determined that is smaller than the target derived from the block difficulty. Bitcoin and many other PoW blockchains employ a double-SHA256 function to compute such hash values:

$$\mathsf{hash} = \mathsf{SHA256d}(\mathsf{header}\|\mathsf{nonce}). \tag{1}$$

The miner that finds a valid hash broadcasts the corresponding block to the network and, upon validation, receives a block reward.

PoW target difficulty represents the expected number of attempts required to determine a valid nonce. To maintain a consistent average block-generation time (BGT), the difficulty is periodically adjusted based on the observed block intervals. If the computational power of the network increases, the same difficulty results in blocks being determined more quickly, reducing the average block time. To address this problem, the difficulty is updated every ℓ blocks using the ratio of the observed BGT to the target BGT. Let t_i denote the timestamp of the ith block and t_{target} be the desired average block time. Then, the difficulty D_i of the ith block is defined as follows:

$$D_i = \begin{cases} D_{i-1} \cdot \left(\frac{t_i - t_{i-\ell}}{t_{target} \cdot \ell}\right), & \text{if } i \bmod \ell = 0, \\ D_{i-1}, & \text{otherwise.} \end{cases} \tag{2}$$

B. Verifiable Random Function

The VRF is a cryptographic primitive that enables a prover to generate pseudorandom output along with a proof that it was correctly computed using a secret key. Unlike conventional hash functions, the VRF provides verifiability of the output: the output and its proof can be verified using the corresponding public key. This property is advantageous in decentralized systems where trust assumptions are minimal. The VRF protocol comprises three principal phases: key generation, output and proof generation, and verification. In the key generation phase, the algorithm generates a public-private key pair using a standard public key cryptosystem, such as ECDSA or RSA [58]. The key pair $(\mathsf{pk}, \mathsf{sk})$ is generated as follows:

$$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(). \tag{3}$$

The KeyGen function is a randomized algorithm that takes no input and outputs a valid public–private key pair. In the proving phase, the prover computes a pseudorandom output β and a proof π using the private key sk and input α:

$$\pi \leftarrow \mathsf{Prove}(\mathsf{sk}, \alpha), \qquad \beta \leftarrow \mathsf{Hash}(\mathsf{sk}, \alpha), \tag{4}$$

$$\mathsf{VRF} : (\mathsf{sk}, \alpha) \longmapsto (\beta, \pi), \tag{5}$$

where VRF represents a deterministic function that takes a secret key sk and input α and returns a pseudorandom output β along with a proof π certifying its correctness. The output β must be publicly verifiable without access to the private key. To enable this method, VRFs define a public extraction function Hash_π that reconstructs β from the proof π alone:

$$\beta \leftarrow \mathsf{Hash}_\pi(\pi). \tag{6}$$

The proof π is considered valid if it is correctly generated using the private key corresponding to pk and the input α, and if it

satisfies the verification equation defined by the underlying VRF construction. The verifier checks the validity of π and confirms that the derived output β matches the expected value using the Verify function:

$$\text{Verify}(pk, \alpha, \pi) = \begin{cases} (\text{VALID}, \ \beta), & \text{if } \pi \text{ is valid,} \\ \text{INVALID}, & \text{otherwise.} \end{cases} \quad (7)$$

As β can be deterministically verified from π and pk, none of the participants can manipulate the output to their advantage, guaranteeing publicly verifiable and unbiased randomness, making VRFs suitable for trustless decentralized systems.

The reliability of the VRF mechanism is grounded in fundamental cryptographic security properties. Three properties define the security of a VRF:

- *Uniqueness:* For a given secret key and input, exactly one valid output and proof exists, preventing a prover from producing multiple valid outputs for the same input to bias the protocol.
- *Collision Resistance:* Determining two distinct inputs that produce the same VRF output is computationally infeasible, ensuring input-output binding.
- *Pseudorandomness:* Without the secret key, an adversary cannot predict the VRF output for any input, and the output is computationally indistinguishable from a uniformly random value.

In this paper, the VRF-PoW protocol uses the elliptic-curve VRF (ECVRF) specified in IETF RFC 9381 [58], which satisfies all three properties above. However, elliptic-curve cryptography may be vulnerable in the era of quantum computing due to Shor's algorithm, which can efficiently solve the discrete logarithm problem. As research on post-quantum VRFs progresses, ECVRF can be replaced with a post-quantum cryptography (PQC)-based construction to ensure long-term robustness [68], [69].

IV. Method

This work proposes VRF-PoW, an energy-efficient consensus mechanism that integrates PoW with a VRF. This mechanism is designed for large-scale, permissionless PoW networks where excessive energy consumption poses a significant challenge. The VRF-PoW mechanism probabilistically self-selects a subset of miners to participate in the PoW process, reducing network energy consumption while retaining the low communication overhead and fully stateless, trustless operation model of PoW. A node participating in the VRF-PoW consensus proceeds through three phases: pre-PoW, VRF-based self-selection, and post-PoW. In the pre-PoW phase, a node performs preliminary work to obtain an entry ticket, after which it executes its own VRF in the self-selection phase to determine eligibility for the post-PoW phase locally; only the self-selected nodes proceed to the post-PoW phase.

The critical challenges of integrating VRFs with PoW in a permissionless environment include that miner identities cannot be preverified. Moreover, empty rounds may occur when the VRF selects no eligible winner, and adversarial miners may attempt to exploit or reuse another miner's winning VRF output even if they were not legitimately selected. The VRF-PoW mechanism addresses these problems via three mechanisms: (i) *Pre-PoW*, which imposes a computational cost to execute the VRF, ensures that only nodes that satisfy a basic work requirement can execute it, suppressing unlimited identity replication and repeated VRF executions. (ii) *Progressive timeout*, which gradually increases the VRF winning probability when no block is generated for a certain period, ensures continuous block generation even during rounds without a selected miner and stabilizes the block interval distribution, which would increase sharply with fixed timeouts. (iii) *Key-bound mining*, which binds the private key used for VRF self-selection to the one used for block mining and signing, prevents unauthorized nodes from receiving or reusing another miner's winning VRF output.

When a node mines for a block, it first enters the pre-PoW phase. Nodes that solve the pre-PoW proceed to execute a VRF to generate pseudorandom output. A node becomes eligible for post-PoW participation if its VRF output is smaller than a threshold corresponding to the VRF winning probability p. The probability p gradually increases under the progressive timeout mechanism when no block is generated for an extended period. The selected miners enter the post-PoW phase, where they compete to determine a valid nonce that meets the difficulty target. Once found, the miner generates a new block including the corresponding VRF proof and broadcasts it to the network. Other nodes validate the block and the attached proof upon receipt. Only a subset of nodes participates in the post-PoW phase; thus, the protocol reduces network energy consumption. Fig. 1 illustrates the interactions between these three phases. This figure illustrates a miner's participation from b_i to b_{i+2}. The Post-PoW phase runs only for b_{i+1}, where the miner wins the VRF, while it is skipped for b_i and b_{i+2}. The Pre-PoW phase is executed for every block regardless of the VRF result.

A. Pre-Pow

The pre-PoW phase executed by all miners, which partially reduces the energy efficiency of VRF-PoW, is critical for defending against Sybil attacks. In PoW-based systems, computational effort serves as an assertion of identity. Without a nontrivial cost prior to VRF execution, an adversary could exploit the low computational burden of VRFs to generate multiple secret keys and repeatedly attempt self-selection, a strategy defined in this work as a VRF-PoW Sybil attack.

Such attacks are especially severe in permissionless networks where nodes can join anonymously, making it infeasible to detect whether multiple identities originate from the same entity. Although wallet-level inferences may provide limited insight, they are neither robust nor sufficient for defense. By requiring locally performed pre-PoW computations before each VRF invocation, VRF-PoW ensures that each self-selection attempt incurs a real cost, deterring identity spamming [70]. Honest miners incur this cost only once per block. In contrast, a Sybil attacker must perform multiple pre-PoW computations to make repeated VRF attempts, increasing their computational cost and limiting their advantage. A detailed analysis of the VRF-based Sybil attack is presented in Section VI-B.

The pre-PoW phase requires miners to search for a nonce that yields a hash value below a given difficulty target. However, an

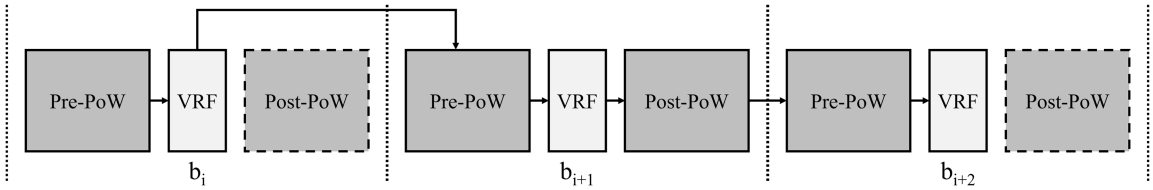

Fig. 1. VRF-PoW.

additional concern arises from the structure of VRF: the VRF output β depends on the input α and the private key sk. Thus, a miner that finds a valid pre-PoW solution could reuse the same input with multiple private keys to execute the VRF repeatedly without incurring additional computational cost.

To prevent repeated VRF executions using multiple private keys on the same pre-PoW solution, VRF-PoW binds the computation to the miner's identity by incorporating a digital signature into the hash input. The protocol also enforces that the private key for pre-PoW generation must be identical to that for VRF execution. Instead of directly using the block header as input to the hash function, VRF-PoW requires miners to sign the concatenation of the block header and a pre-PoW nonce, $\mathsf{nonce}_{\mathsf{pre}}$, using their private key. This mechanism was inspired by Eyal's 2P-PoW protocol, which was proposed to discourage selfish mining [29], [52]. The miner first generates a message by concatenating the block header with a pre-PoW nonce, denoted by $\mathsf{nonce}_{\mathsf{pre}}$. This message is signed using the miner's private key sk. The signature function, $\mathsf{Sign}_{sk}(m)$ takes the message $m = (\mathsf{header} \| \mathsf{nonce}_{\mathsf{pre}})$ as input and produces a cryptographic signature. This digital signature, denoted by σ_{pre}, is included in the block header as proof of the miner's identity and the integrity of the message:

$$\sigma_{\mathsf{pre}} = \mathsf{Sign}_{\mathsf{sk}}(\mathsf{header} \,\|\, \mathsf{nonce}_{\mathsf{pre}}), \tag{8}$$

$$\mathsf{hash}_{\mathsf{pre}} = \mathsf{SHA256d}(\sigma_{\mathsf{pre}}), \tag{9}$$

where SHA256d denotes the double SHA-256 function. The signature σ_{pre} ensures that the pre-PoW hash is cryptographically bound to the miner's identity and the selected nonce while guaranteeing key consistency. The generated hash $\mathsf{hash}_{\mathsf{pre}}$ serves as a verifiable output of the pre-PoW phase and is employed as the input to the VRF. If $\mathsf{hash}_{\mathsf{pre}}$ is less than the target value determined by the pre-PoW difficulty d_{pre}, it is a valid solution and proceeds to the VRF stage. Any attempt to reuse the same message with a different private key yields a different signature and different pre-PoW hash.

B. VRF Self-Selection

Miners that solve the pre-PoW puzzle execute the VRF function, determining their eligibility to participate in the post-PoW phase. The VRF takes the hash of the pre-PoW message as input, which is generated by signing the block header concatenated with a pre-PoW-specific nonce using the miner's private key. To ensure key consistency and prevent identity forgery, the same private key must be used for the pre-PoW signature and VRF execution. Miner eligibility is determined by a verifiable coin toss (VCT) function that probabilistically selects miners based on the VRF output. The VCT compares the VRF output against a threshold value derived from the predefined self-selection probability p. Therefore, only a fraction p of all miners in the network are expected to proceed to the post-PoW phase on average.

Due to the random selection property of the VCT, a small but nonzero probability exists that no miners are self-selected for the post-PoW phase. If no miner is self-selected, the post-PoW phase cannot proceed, potentially leading to a deadlock in which block generation halts. The network cannot determine whether no miners were self-selected or block generation was merely delayed. The VRF-PoW mechanism must ensure network continuity even when no one wins the VCT. A common solution is to introduce a timeout threshold, allowing all miners to continue mining if a block is not produced within the specified time limit [27], [28]. However, enabling a timeout can cause abrupt fluctuations in network hash power, potentially leading to chain forks and uncle blocks and reducing blockchain stability. Therefore, the timeout must be significantly longer than the target BGT.

This work employs an improved progressive timeout mechanism to ensure liveness without compromising stability. Instead of allowing all nodes to participate once the timeout is reached, the VRF self-selection probability $p(t)$ increases progressively from an initial value p_0 to 1. When a timeout occurs, miners with VRF outputs initially above the self-selection threshold but closer to the threshold at timeout are granted earlier opportunities to participate in the PoW process. The eligibility range gradually expands. By progressively implementing the timeout mechanism, VRF-PoW minimizes abrupt surges in network hash power, which could cause instability (e.g., chain forks or uncle blocks).

This work defines the time-dependent VRF self-selection probability $p(t)$, which increases with the elapsed time t since the last block, to formalize the mechanism. The self-selection probability $p(t)$ was designed to increase exponentially, as the influence of its growth depends on the current state of network hash power. Just after the timeout begins, even a small absolute increase in the self-selection probability can lead to a relatively large increase in total hash power, whereas the same increase at a later stage has a proportionally smaller effect. Let t_s be the timeout threshold, t be the elapsed time since the last block, and t_e be the time at which the self-selection probability reaches 1, marking the end of the progressive timeout window. Then, $p(t)$ is given by the following:

$$p(t) = \begin{cases} p_0, & \text{if } t < t_s, \\ p_0 \cdot e^{\frac{t-t_s}{\tau}}, & \text{if } t_s \le t < t_e, \\ 1, & \text{if } t \ge t_e. \end{cases} \tag{10}$$

The time constant τ is selected to satisfy $p(t_e) = 1$ and is defined as follows:

$$\tau = \frac{t_e - t_s}{\ln\left(\frac{1}{p_0}\right)}. \tag{11}$$

As time progresses from t_s to t_e, the probability $p(t)$ gradually increases, reaching a maximum value of 1 at t_e, resulting in a time-varying threshold $\theta(t)$ for the VRF output, where n represents the bit length of the VRF output:

$$\theta(t) = p(t) \cdot (2^n - 1). \tag{12}$$

If a miner's VRF output β is smaller than the threshold $\theta(t)$, the miner is eligible to participate in the post-PoW phase. The VCT is executed as follows:

$$(\beta, \pi, r) = \mathsf{VCT}(\mathsf{hash}_{\mathsf{pre}}, \mathsf{sk}, t),$$
$$\text{where } (\beta, \pi) = \mathsf{VRF}(\mathsf{sk}, \mathsf{hash}_{\mathsf{pre}}), \quad r = \begin{cases} 1, & \text{if } \beta \le \theta(t), \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

The miner's pre-PoW hash $\mathsf{hash}_{\mathsf{pre}}$ is the input message to the VRF. If $r = 1$, the miner proceeds to the post-PoW phase; otherwise, the miner enters an idle state. As $p(t)$ increases over time, the threshold $\theta(t)$ also rises, eventually allowing the miner to participate in the post-PoW phase by satisfying the VCT condition. The earliest possible submission time $t_{\text{sub}}(\beta)$ at which a miner with VRF output β can enter the post-PoW phase is defined as follows:

$$t_{\text{sub}}(\beta) = \begin{cases} 0, & \text{if } \beta \le \theta(t_s), \\ (t_e - t_s)\left(1 - \frac{\ln\beta - \ln(2^n - 1)}{\ln p_0}\right) + t_s, & \text{otherwise.} \end{cases} \tag{14}$$

The block header includes the VRF output and its corresponding proof. Upon receiving a new block, nodes verify the VRF output using the proof and associated public key, including verifying that the VRF public key matches the PoW key and that the output β was submitted within the acceptable time frame defined by $t_{\text{sub}}(\beta)$.

C. Post-Pow

Miners that pass the pre-PoW stage are self-selected via a VCT and proceed to the post-PoW phase. The input to the post-PoW phase includes the miner's pre-PoW hash, VRF output, and the corresponding proof. To ensure authenticity, the post-PoW input is also signed with the miner's private key, binding the resulting block to the miner possessing the associated secret key. The function of the private key in the post-PoW phase remains consistent with its role in the pre-PoW phase, establishing a secure linkage between the miner's identity and computational effort. The miner creates a digital signature σ_{post} for the concatenation of the block header and a random nonce for the post-PoW $\mathsf{nonce}_{\mathsf{post}}$, the pre-PoW hash $\mathsf{hash}_{\mathsf{pre}}$, and the VRF output β:

$$\sigma_{\mathsf{post}} = \mathsf{Sign}_{\mathsf{sk}}(\mathsf{header} \,\|\, \mathsf{nonce}_{\mathsf{post}} \| \mathsf{hash}_{\mathsf{pre}} \| \beta), \tag{15}$$

$$\mathsf{hash}_{\mathsf{post}} = \mathsf{SHA256d}(\sigma_{\mathsf{post}}). \tag{16}$$

The post-PoW hash $\mathsf{hash}_{\mathsf{post}}$ serves as the final VRF-PoW output and must satisfy the network's post-PoW difficulty target to be valid. The hash is derived from a signature that includes $\mathsf{hash}_{\mathsf{pre}}$; thus, it links the post-PoW phase to the pre-PoW result.

The post-PoW difficulty in VRF-PoW is dynamically adjusted, as in the Nakamoto consensus (NC), based on the time difference between the actual and target block intervals. The VRF-PoW mechanism maintains separate difficulty levels for pre and post-PoW, with the pre-PoW difficulty defined as a fixed fraction of the post-PoW difficulty. As pre-PoW is performed individually by all miners and post-PoW involves only a VRF self-selected subset of nodes, the pre-PoW difficulty must be lower to ensure feasibility and fairness. Let ℓ denote the difficulty adjustment interval, and let $k \in (0, 1)$ be the pre-PoW scaling factor applied to the post-PoW difficulty to define the pre-PoW difficulty. The difficulties are defined as follows:

$$D_{\text{post},i} = \begin{cases} D_{\text{post},i-\ell} \cdot \left(\frac{t_i - t_{i-\ell}}{t_{\text{target}} \cdot \ell}\right), & \text{if } i \bmod \ell = 0, \\ D_{\text{post},i-1}, & \text{otherwise,} \end{cases} \tag{17}$$

$$D_{\text{pre},i} = k \cdot D_{\text{post},i}. \tag{18}$$

Algorithm 1 presents the pseudocode for the VRF-PoW consensus mechanism, which operates in three phases: pre-PoW, VRF, and post-PoW. During execution, if a new block is received from another node, the local mining process is immediately aborted. The block header contains all relevant components, including the nonce, hash value, VRF output, and proof. The miner runs the VCT function to obtain the result r. If $r = 0$, the miner waits until the time $t_{\text{sub}}(\beta)$ is reached before proceeding to the post-PoW phase. Otherwise, the miner proceeds to the post-PoW phase immediately. When a valid block is generated or received, the node appends it to the local chain and proceeds to the next round.

The verification process for an incoming block involves three critical checks. First, the validity of the pre-PoW hash is confirmed by reconstructing the signed data using the block header and $\mathsf{nonce}_{\mathsf{pre}}$ and comparing it to the submitted hash. Second, the VRF output and its corresponding proof are validated using the miner's public key, ensuring that the output satisfies the expected threshold. Finally, the post-PoW result is verified by recalculating the hash with the extended block header and $\mathsf{nonce}_{\mathsf{post}}$, and confirming that it falls below the target difficulty. The block is appended to the local blockchain only upon the successful completion of all these checks.

D. Rationale for Using Pre-Proof of Work

Although pre-PoW is less energy-efficient than alternatives (e.g., decentralized identifier-based authentication or PoS), due to its reliance on solving cryptographic puzzles, it provides fundamental structural advantages [71], [72]. Decentralized Identifier (DID)-based and PoS mechanisms can mitigate Sybil attacks with lower energy consumption but cannot fully replicate the benefits of pre-PoW in a decentralized, permissionless network environment. The decision is driven by the need to ensure autonomous and independent block generation in such environments. Participation is open to anyone, without requiring prior identity registration or token ownership for validation. Alternative designs might group multiple blocks into epochs and perform collective preverification or replace pre-PoW entirely with group consensus mechanisms [27], [35]. However, these approaches introduce significant complexity. They require communication and synchronization to reach an agreement. These

Algorithm 1: VRF-PoW Consensus.

Initialize:
generate key pair $(\mathsf{sk}, \mathsf{pk})$;
initialize the blockchain as $blocks \leftarrow \{b_0, b_1, \ldots, b_h\}$;
retrieve timeout parameters (t_s, t_e) from the genesis block b_0;
for every block mining **do**
  extract $(p_0, k, D_{\mathsf{post},h})$ from b_h;
  compute $D_{\mathsf{pre},h}$;
  construct the initial header for new block b_{h+1};
  while $\mathsf{hash}_{\mathsf{pre}}$ does not satisfy $D_{\mathsf{pre},h}$ **do**
    select a random value $\mathsf{nonce}_{\mathsf{pre}}$;
    generate σ_{pre} by signing $(\mathsf{header} \| \mathsf{nonce}_{\mathsf{pre}})$ with sk;
    compute $\mathsf{hash}_{\mathsf{pre}}$ by applying $\mathsf{SHA256d}$ to σ_{pre};
  end while
  retrieve the current system time t;
  compute output (β, π, r) using VCT;
  if $r = 0$ **then**
    wait until $t_{\mathsf{sub}}(\beta)$ is reached;
  end if
  while $\mathsf{hash}_{\mathsf{post}}$ does not satisfy D_{post} **do**
    select a random value $\mathsf{nonce}_{\mathsf{post}}$;
    generate σ_{post} by signing:
      $(\mathsf{header} \| \mathsf{nonce}_{\mathsf{post}} \| \mathsf{hash}_{\mathsf{pre}} \| \beta)$ with sk;
    compute $\mathsf{hash}_{\mathsf{post}}$ by applying $\mathsf{SHA256d}$ to σ_{post};
  end while
  append the following fields to the block header header:
    $\mathsf{nonce}_{\mathsf{post}}$, $\mathsf{hash}_{\mathsf{post}}$, $\mathsf{nonce}_{\mathsf{pre}}$, $\mathsf{hash}_{\mathsf{pre}}$, β, π, and pk;
  generate a new block b_{h+1} with header;
  append b_{h+1} to $blocks$;
  broadcast b_{h+1} to the network;
  increment block height $h \leftarrow h + 1$;
end for

processes complicate the block-generation algorithm, increase implementation burden, and reduce system robustness and auditability. The increased communication overhead and coordination complexity make the approach unsuitable for large-scale decentralized networks.

In contrast, VRF-PoW enables consensus execution without requiring prior communication between nodes. Block producers can independently complete VRF and PoW computations before broadcasting the results. This method eliminates the need for premining, ensuring a permissionless, decentralized block-generation process. By preserving this structure, VRF-PoW maintains the integrity of the consensus protocol while ensuring the simplicity of implementation.

E. Decentralization and Scalability

The VRF-PoW approach employs private-key-based digital signatures in the pre and post-PoW phases to reinforce decentralization and security. This mechanism mitigates Sybil attacks by binding computational work to unique identities and discourages the formation of mining pools because mining privileges and block rewards cannot be outsourced or aggregated without compromising private key ownership. In traditional PoW, mining pools increase the probability of earning rewards by aggregating hash power, often leading to network centralization [9]. In contrast, collaborative mining in VRF-PoW is structurally infeasible, as reward authorization is strictly tied to the private key used in self-selection and block generation. Any attempt to share this key would entail a significant security risk, such as loss of wallet control, disincentivizing pool formation [29], [52].

The core constraint of VRF-PoW is that the outputs of all phases must be strictly bound to the same private key signature. This condition is the only one required; hence, the PoW primitive in VRF-PoW is not limited to hash-based functions; it can be replaced with any work function (e.g., Ethash-based PoW) or can be integrated with ASIC-resistant primitives (e.g., memory-hard or code-based computations) [64], [65]. This flexibility allows VRF-PoW to achieve ASIC resistance and narrow the performance gap between ASIC miners and commodity hardware, enhancing decentralization.

Scalability in blockchain networks refers to the ability to handle an increased number of transactions and users without compromising speed, security, or decentralization [73]. From a user perspective, VRF-PoW is designed as a permissionless protocol, meaning that an increase in the number of participants does not directly constrain its operation. However, the total network hash power will inevitably increase, and such growth can be mitigated by adjusting the VRF winning probability p to limit the proportion of active miners. From a transactional perspective, throughput can be improved only by increasing the block size or shortening the block interval, but both approaches increase the probability of forks. Shorter block intervals lead to more frequent chain reorganizations, whereas larger blocks take longer to propagate, increasing the likelihood of forks. Moreover, the fork probability is approximately proportional to the number of miners actively participating in the network [74]. The VRF-PoW mechanism reduces the fork probability and aggregate network power involved in consensus by limiting the fraction of miners participating in the post-PoW phase, providing more favorable scalability conditions than the conventional NC.

V. Theoretical Analysis

This section presents a theoretical analysis of the proposed VRF-PoW mechanism. First, this work derives the BGT distribution of the conventional NC as a baseline and models the BGT of VRF-PoW using a hypo-exponential distribution. Then, this work analyzes the difficulty adjustment, block success probability, and energy consumption. Finally, a VRF-based Sybil attack is examined, comparing the efficiency of honest and adversarial strategies.

A. Block Generation Time Distribution for the Nakamoto Consensus

To analyze key performance metrics, including difficulty adjustment, mining success probability, and energy consumption, this work first models the probability distribution of BGT in VRF-PoW. This work assumes the entire blockchain network has a total hash power H, with n nodes, each with a hash rate

of h_i, and the target average BGT is t_{target}. The total network hash power satisfies $\sum_{i=1}^{n} h_i = H$.

The preliminary step derives the BGT distribution of NC, which serves as a baseline for comparison with VRF-PoW. In NC, the BGT distribution for miner m_a follows the exponential distribution $X_a^p \sim \exp(\lambda_a)$ with $\lambda_a = h_a/D$. The cumulative distribution function (CDF) of the BGT $F_{X_a^p}(x)$ for this miner is given by [75]

$$F_{X_a^p}(x) = 1 - e^{-\lambda_a x}. \tag{19}$$

This work considers the random variable X^{p}, defined as the minimum of all X_a^{p} for $a = 1, \ldots, n$; hence, the network BGT can be defined as follows:

$$X^{\mathsf{p}} := \min\{X_1^{\mathsf{p}}, X_2^{\mathsf{p}}, \ldots, X_n^{\mathsf{p}}\}. \tag{20}$$

This work considers the survival function of the network BGT, the probability that the network BGT exceeds a given value x, or 1 minus the CDF, i.e., $S_{X^p}(x) = 1 - F_{X^p}(x)$. The survival function of an exponential distribution is $S_{X^{\mathsf{p}}}(x) = e^{-\lambda x}$. Thus, the survival function of the network BGT can be expressed as the product of the survival functions of the miners' BGT:

$$\begin{aligned} S_{X^p}(x) &= P(x < X^p) \\ &= \prod_{i=1}^{n} S_{X_i^p}(x) = e^{-\lambda x}, \end{aligned} \tag{21}$$

where $\lambda = \sum_{i=1}^{n} \lambda_i$. The probability density function (PDF) of the network BGT can be obtained by differentiating the CDF derived from the survival function:

$$f_{X^p}(x) = \frac{d}{dx}(1 - S_{X^p}(x)) = \lambda e^{-\lambda x}. \tag{22}$$

Thus, the network BGT in NC follows a simple exponential distribution.

B. Network BGT for VRF-PoW

This work derives the network BGT distribution for VRF-PoW by modeling the pre and post-PoW stages. As the VRF execution time is negligible compared to PoW, it is excluded from the analysis. Because timeout occurrences are rare, the VRF success probability p is assumed to be constant over time for analytical simplicity. Only self-selected miners contribute to block generation; hence, the BGT is modeled as a function of node being self-selected. In VRF-PoW, where a miner must sequentially complete the pre and post-PoW stages, the total BGT is modeled as the sum of two exponential phases.

The sum of two exponential distributions with different rates follows a hypo-exponential distribution, generalizing the Erlang distribution (which assumes equal rates) by allowing distinct rate parameters for each phase [76]. Accordingly, the BGT of a VRF-PoW miner m_a, denoted by X_a^v, follows a two-phase hypo-exponential distribution:

$$X_a^v \sim \text{Hypo}\left(\frac{\lambda_a}{k}, \lambda_a\right). \tag{23}$$

When $k = 1$, the two phases have identical rates, and the distribution reduces to the Erlang case. For $k \neq 1$, the rates differ, resulting in a special case of a phase-type distribution. The explicit forms of the PDF $f_{X_a^v}(x)$ and the survival function $S_{X_a^v}(x)$ for X_a^v, when $k \neq 1$ are given as follows [77]:

$$f_{X_a^v}(x) = \frac{\frac{\lambda_a}{k} \cdot \lambda_a}{\lambda_a - \frac{\lambda_a}{k}} \left(e^{-\frac{\lambda_a}{k}x} - e^{-\lambda_a x}\right), \tag{24}$$

$$S_{X_a^v}(x) = \frac{1}{\lambda_a - \frac{\lambda_a}{k}} \left(\lambda_a e^{-\frac{\lambda_a}{k}x} - \frac{\lambda_a}{k} e^{-\lambda_a x}\right). \tag{25}$$

The network BGT in VRF-PoW is determined by the minimum PoW time of the VRF self-selected miners from n nodes. The number of self-selected miners follows a binomial distribution, $M \sim \text{Bin}(n, p)$, where p represents the VRF self-selection probability. Although non-self-selected miners also perform the pre-PoW phase to verify eligibility, they do not broadcast blocks and do not affect the observed network BGT.

Let $X^v = \min\{X_1^v, \ldots, X_n^v\}$ denote the global BGT in VRF-PoW. Only self-selected miners can generate blocks; therefore the distribution of X^v depends on the random subset $\mathcal{S} \subseteq \{1, \ldots, n\}$ of selected miners. Each X_i^v follows a two-phase hypo-exponential distribution, as selected miners sequentially perform pre and post-PoW. The probability of selecting a specific subset $\mathcal{S}$ of size $|\mathcal{S}|$ is given by the binomial probability $p^{|\mathcal{S}|}(1-p)^{n-|\mathcal{S}|}$. The global survival function is expressed as a weighted sum over all such subsets:

$$S_{X^v}(x) = \sum_{\mathcal{S} \subseteq \{1,\ldots,n\}} p^{|\mathcal{S}|}(1-p)^{n-|\mathcal{S}|} \prod_{i \in \mathcal{S}} S_{X_i^v}(x), \tag{26}$$

where $S_{X_i^v}(x)$ denotes the survival function of the hypo-exponential BGT for miner i. The corresponding PDF is obtained by differentiating the CDF:

$$f_{X^v}(x) = \frac{d}{dx}\left(1 - S_{X^v}(x)\right). \tag{27}$$

C. Difficulty Adjustment Mechanism in VRF-PoW

Although it is well established that the difficulty adjustment algorithm in PoW blockchains ensures consistent BGT [75], VRF-PoW executes PoW twice, requiring an additional analysis of its influence on timing consistency. This section analyzes whether the average BGT in VRF-PoW remains consistent under its difficulty adjustment mechanism. The difficulty in VRF-PoW is updated based on the deviation between the actual and target BGT, with the previous difficulty level serving as a reference. This adjustment ensures that the expected BGT remains linearly proportional to the current difficulty level:

$$E[X^v] \propto D_{\text{post}}. \tag{28}$$

The expected BGT $E[X^v]$ is computed as follows:

$$\begin{aligned} E[X^v] &= \int_0^\infty x f_{X^v}(x)\, dx \\ &= [-x S_{X^v}(x)]_0^\infty - \int_0^\infty -S_{X^v}(x)\, dx. \end{aligned} \tag{29}$$

As $S_{X^v}(x) \to 0$ exponentially fast as $x \to \infty$, the boundary term $[-xS_{X^v}(x)]_0^\infty$ vanishes, and the expectation simplifies to

$$E[X^v] = \int_0^\infty S_{X^v}(x)\, dx. \tag{30}$$

For a baseline difficulty level of $D_{\text{post}} = 1$, the individual miner's rate parameter is denoted by λ'. The effective rate parameter is inversely proportional to the difficulty; thus, $\lambda = \lambda'/D_{\text{post}}$. The

variable transformation $u = x/D_{\text{post}}$ is applied to simplify the analysis, which implies that

$$dx = D_{\text{post}}\, du. \tag{31}$$

Under this change of variables, the integral for the expectation transforms accordingly. The expected value of X^v becomes:

$$E[X^v] = D_{\text{post}} \cdot E[X^v \mid D_{\text{post}} = 1], \tag{32}$$

where $E[X^v \mid D_{\text{post}} = 1]$ represents the expected BGT under unit difficulty. This result confirms that the expected BGT scales linearly with the post-PoW difficulty D_{post}, that is,

$$E[X^v] \propto D_{\text{post}}. \tag{33}$$

D. Mining Success Probability

In PoW, the probability that a node successfully mines a block is equal to the probability that the node's BGT is the minimum of all nodes. If $I = \min\{X_1, X_2, \ldots, X_n\}$, then the mining success probability for node a is given by

$$P(X_a = I) = \int_0^\infty f_{X_a}(x) \prod_{i \neq a}^{n} S_{X_i}(x)\, dx. \tag{34}$$

This expression computes the probability that node a generates a block at time x, whereas the BGTs of all other nodes exceed x. In NC, the survival function and mining success probability are proportional to the hash power:

$$P(X_a^p = I) = \int_0^\infty \lambda_a e^{-(\lambda_1+\lambda_2+\cdots+\lambda_n)x}\, dx = \frac{\lambda_a}{\lambda}. \tag{35}$$

Thus, the success probability is equal to the fraction of the hash power node a relative to the total network hash power. In the case of VRF-PoW, the overall survival function is computed by summing all subsets of VRF-selected nodes excluding the one under evaluation:

$$S_{X^v}^{(-a)}(x) = \sum_{\substack{\mathcal{S} \subseteq \{1,\ldots,n\} \\ a \notin \mathcal{S}}} p^{|\mathcal{S}|}(1-p)^{n-1-|\mathcal{S}|} \prod_{i \in \mathcal{S}} S_{X_i^v}(x), \tag{36}$$

where $\mathcal{S}$ denotes a subset of VRF-selected nodes excluding node a. The product computes the joint survival function of nodes in $\mathcal{S}$, and the summation weights each subset by its binomial probability. The mining success probability for node a in VRF-PoW is given by

$$P(X_a^v = I) = p \int_0^\infty f_{X_a^v}(x) \cdot S_{X^v}^{(-a)}(x)\, dx. \tag{37}$$

E. Expected Energy Consumption

Energy consumption is proportional to the duration of mining activity; therefore this work first models the actual mining time of each miner. In NC, miners continuously perform mining, so that the mining time can be modeled as X^v. In contrast, in VRF-PoW, mining participation varies with the VRF self-selection outcome, leading to heterogeneous mining durations across miners. This work models the mining time of each miner a using a random variable T_a, conditioned on each miner's VRF self-selection status. Let $T_a^{\text{pre}} \sim \text{Exp}(\mu_a)$ denote the pre-PoW time, where $\mu_a = \lambda_a/k$ and k represents a difficulty-scaling factor. The post-PoW duration is modeled as $T_a^{\text{post}} \sim \text{Exp}(\lambda_a)$. The total mining time is defined as

$$T_a = \begin{cases} X^v, & \text{if self-selected,} \\ T_a^u = \min(T_a^{\text{pre}}, X^v), & \text{otherwise,} \end{cases} \tag{38}$$

where X^v denotes the global BGT. If a miner is self-selected by the VRF, it performs pre and post-PoW. However, its hashing activity is terminated once a block is found in the network, which occurs after time X^v. Therefore, the total mining time for self-selected miners is truncated at X^v and modeled as $T_a = X^v$. In contrast, if a miner is not self-selected, it performs only the pre-PoW phase, which is interrupted upon completion of pre-PoW or earlier if a self-selected miner discovers a block. Hence, the work time is modeled as $T_a^u = \min(T_a^{pre}, X^v)$. The non-self-selected case is denoted by the superscript u. The mining time distribution is given by the following equation:

$$f_{T_a}(x) = p \cdot f_{X^v}(x) + (1-p) \cdot f_{T_a^u}(x), \tag{39}$$

where $f_{X^v}(x)$ and $f_{T_a^u}(x)$ represent the PDFs of the self and non-self-selected mining time, respectively. In particular,

$$f_{T_a^u}(x) = f_{T_a^{\text{pre}}}(x) \cdot S_{X^v}(x) + f_{X^v}(x) \cdot S_{T_a^{\text{pre}}}(x), \tag{40}$$

where $f_{T_a^{\text{pre}}}(x)$ and $S_{T_a^{\text{pre}}}(x)$ indicate the PDF and survival function of the pre-PoW phase. The distribution $f_{X^v}(x)$, as defined, models the BGT of self-selected miners.

The expected energy consumption for miner a is given by

$$E_a = \mathbb{E}[T_a] \cdot h_a, \tag{41}$$

where $\mathbb{E}[T_a]$ denotes the expected mining time derived from the mixture distribution $f_{T_a}(x)$, as defined. The energy savings can be interpreted as the product of the time not spent hashing and the miner's hash power. This work compares the expected work time of non-self-selected miners under two scenarios: (1) the scenario in which they continue hashing for the full block generation duration and (2) their actual behavior, when they terminate early when another miner completes the block. Let $\mathbb{E}[X^v]$ denote the expected global BGT and $\mathbb{E}[T_a^u]$ represent the expected truncated work time of non-self-selected miner a. The expected energy saved by miner a is

$$\Delta E_a = (1-p) \cdot (\mathbb{E}[X^v] - \mathbb{E}[T_a^u]) \cdot h_a, \tag{42}$$

where $1-p$ denotes the probability that miner a is not self-selected by the VRF.

After an expected time of $\mathbb{E}[T_a^u]$, miners terminate earlier upon completion of pre-PoW or when a block is found elsewhere in the network. The difference in expected work time corresponds to the duration of the avoided hashing effort. The corresponding energy savings can be quantified by multiplying the time gap by the miner's hash rate h_a. Only a fraction $1-p$ of miners are non-self-selected; hence, the expected savings are weighted accordingly.

VI. Security Analysis

This section presents a security analysis of the VRF-PoW protocol. First, this work briefly describes the fundamental security properties of PoW protocols and examines how these properties are preserved in VRF-PoW. This section explains why the VRF Sybil attack is crucial to attempting major attack vectors (e.g., double spending in VRF-PoW), modeling the attack, and analyzing its efficiency. Finally, this work analyzes the tradeoff

between security and energy efficiency determined by the key parameters p and k.

A. Security Properties and the VRF Sybil Attack

The NC assumes that the blockchain's safety and liveness are guaranteed if honest nodes collectively control more than half of the total network hash power. The following three fundamental properties [78] characterize the security of PoW protocols:

- *Common prefix:* After sufficient confirmations, the blockchains maintained by any two honest nodes share a common prefix, ensuring probabilistic finality and resistance to double-spending.
- *Chain quality:* The fraction of adversarial blocks in the chain is bounded, preventing long-term censorship or biased inclusion of transactions.
- *Chain growth:* The blockchain continues to extend steadily even under adversarial interference, ensuring liveness and continuous transaction processing.

These properties are satisfied under the honest-majority assumption [79] because, in NC, the probability that a miner generates a block is proportional to the fraction of its total network hash power, as presented in (35). If an adversary controls at least half of the hash power, its block-generation probability exceeds 50%, violating these properties. The VRF-PoW protocol inherits the same fundamental security assumptions as NC, but the block-generation probability is computed differently, as analyzed in Section V-D. As in (37), under honest participation, the block-generation probability of any single miner is limited to p. In an honest network, VRF-PoW satisfies the same core security properties.

However, because VRF-PoW also operates in a permissionless environment, an adversary node may generate and control numerous private keys. An adversary can exploit this vulnerability by repeatedly performing the pre-PoW and VRF self-selection steps across multiple identities, increasing its selection probability. This strategy is defined as the VRF Sybil attack. Under this strategy, the attacker's effective block-generation probability can exceed p. By applying multiple identities, such an adversary undermines the fairness of the self-selection process and gains the necessary preconditions to mount traditional attacks (e.g., double-spending). To assess the feasibility and potential advantage of the VRF Sybil attack, the following subsection defines a probabilistic model of the VRF Sybil strategy and analyzes its expected efficiency.

B. Modeling and Analysis of VRF Sybil Attack

Under the Sybil strategy, a malicious miner repeatedly performs pre-PoW until being self-selected by the VRF. Each pre-PoW attempt by miner a requires an exponentially distributed time $X_a^{(\text{pre},j)} \sim \text{Exp}(\mu_a)$, where $\mu_a = \lambda_a/k$ denotes the pre-PoW difficulty scaled by a factor k. The number of attempts until self-selection follows a geometric distribution $N \sim \text{Geo}(p)$. The total time spent by a miner until being self-selected, denoted as X_a^{sp}, is modeled as the sum of multiple pre-PoW durations. Each duration corresponds to an independent attempt, and the number of attempts N follows a geometric distribution. Formally,

$$X_a^{\text{sp}} = \sum_{j=1}^{N} X_a^{(\text{pre},j)}, \tag{43}$$

where each $X_a^{(\text{pre},j)} \sim \text{Exp}(\mu_a)$ and $N \sim \text{Geo}(p)$. the Laplace transform of X_a^{sp} is derived as:

$$\begin{aligned}\mathbb{E}[e^{-sX_a^{\text{sp}}}] &= \sum_{n=1}^{\infty}(1-p)^{n-1}p\left(\frac{\mu_a}{s+\mu_a}\right)^n \\ &= \frac{p\mu_a}{s+\mu_a}\cdot\frac{1}{1-\frac{(1-p)\mu_a}{s+\mu_a}} = \frac{p\cdot\frac{\lambda_a}{k}}{s+p\cdot\frac{\lambda_a}{k}}.\end{aligned} \tag{44}$$

This result corresponds to the Laplace transform of an exponential distribution with rate $\frac{p\lambda_a}{k}$; therefore,

$$X_a^{\text{sp}} \sim \text{Exp}\left(\frac{p\lambda_a}{k}\right). \tag{45}$$

The post-PoW time is modeled independently as $X_a^{\text{post}} \sim \text{Exp}(\lambda_a)$. The total Sybil mining time is defined as the sum of two independent exponential random variables with distinct rates, X_a^{sp} and X_a^{post}. By definition, such a sum follows a hypo-exponential distribution:

$$X_a^{\text{syb}} = X_a^{\text{sp}} + X_a^{\text{post}} \sim \text{Hypo}\left(\frac{p\lambda_a}{k}, \lambda_a\right). \tag{46}$$

The probability that a Sybil miner generates a block before all honest miners is expressed as follows:

$$P(X_a^{\text{syb}} = I) = \int_0^{\infty} f_{X_{\text{syb}}}(x)\cdot S_{X^v}^{(-a)}(x)\,dx. \tag{47}$$

The survival function $S_{X^v}^{(-a)}(x)$, defined in (36), represents the survival function of the minimum BGT of the honest self-selected nodes, excluding the Sybil miner.

This work defines v' as the VRF-PoW network that comprises the Sybil miner and subset of honest, self-selected miners. As the Sybil miner operates independently of the VRF self-selection process and participates in every round, its survival function must be considered jointly with that of the honest miners. The survival function for the BGT in v' is the product of the Sybil miner's survival function and that of the honest self-selected miners:

$$S_{X^{v'}}(x) = S_{X^v}^{(-a)}(x)\cdot S_{X^{\text{syb}}}(x). \tag{48}$$

The expected BGT in v' is given by

$$\mathbb{E}[X^{v'}] = \int_0^{\infty} S_{X^{v'}}(x)\,dx. \tag{49}$$

Thus, the expected energy expenditure of the Sybil miner is expressed as:

$$E_a^{\text{syb}} = \mathbb{E}[X^{v'}]\cdot h_a. \tag{50}$$

To compare the relative efficiency of honest versus Sybil strategies for a given miner a, across varying VRF self-selection probabilities p, the following efficiency ratio is defined:

$$\text{Efficiency Ratio}_a = \frac{P(X_a^v = I)/E_a}{P(X_a^{\text{syb}} = I)/E_a^{\text{syb}}}. \tag{51}$$

where $P(X = I)$ denotes the probability that miner a generates a block that is accepted into the blockchain under the given strategy. The terms E_a and E_a^{syb} represent the corresponding

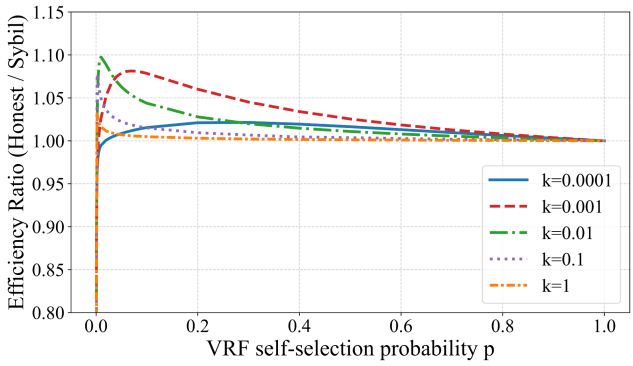

Fig. 2. Efficiency ratio between honest and Sybil mining strategies as a function of p and k, with $n = 10,000$.

expected energy expenditures. An efficiency ratio exceeding 1 implies that the honest strategy yields a higher block acceptance rate per unit of energy consumption than the Sybil strategy.

This work assumes a network of $n = 10000$ miners, each with identical hash power, to analyze these strategies under realistic network conditions. The identical hash power across miners enables the use of a generalized survival function, which yields a tractable, closed-form expression for network-wide BGT. The survival function $S_{X^v}^{(-a)}(x)$, derived via the probability generating function, simplifies to

$$S_{X^v}^{(-a)}(x) = (1 - p + p \cdot S_{X_0}(x))^n . \tag{52}$$

Fig. 2 illustrates the efficiency ratio between honest and Sybil mining strategies under varying self-selection probabilities p and pre-PoW difficulty scaling factors k. Small values of p give honest miners a significant efficiency advantage. As p increases, this advantage diminishes, and the efficiency ratio asymptotically approaches 1, indicating equivalent efficiency between the two strategies. When k is small, honest miners achieve higher energy efficiency, whereas larger k values reduce efficiency as honest participants still consume energy even when they are not selected. However, if k becomes excessively small, the pre-PoW workload becomes too light, which lowers the cost for Sybil attackers, which may reduce overall energy efficiency. Therefore, smaller values of p and k that still yield an efficiency ratio greater than 1 are desirable, balancing Sybil resistance and energy efficiency.

C. Security and Efficiency Under Parameters p and k

In most cases, the honest mining strategy is more efficient than the Sybil strategy. Nevertheless, miners may still attempt to launch Sybil attacks against the network, even at the cost of reduced personal efficiency. As p and k significantly influence the efficiency of the Sybil strategy and the overall mining probability and energy consumption of the network, this work analyzes how the adversarial hash power threshold required to dominate the chain and the network's total energy saving vary as functions of p and k. Under the Sybil strategy, the security threshold r_{adv} is defined as the minimum adversarial hash power

TABLE II
OPTIMAL PARAMETER CONFIGURATIONS ACROSS ADVERSARIAL THRESHOLD LEVELS ($n = 10000$)

Target r_{adv}	Optimal p	Optimal k	Expected $\eta(p,k)$
0.50	1	0	0
0.30	≈ 0.8	$\approx 10^{-4}$	≈ 0.12
0.20	≈ 0.4	$\approx 10^{-4}$	≈ 0.44
0.10	≈ 0.2	$\approx 10^{-4}$	≈ 0.67

required to achieve a block-generation probability of at least 0.5:

$$r_{\text{adv}} = \min\left\{ \frac{h_a}{H} : P(X_a^{\text{syb}} = I) \geq 0.5 \right\}. \tag{53}$$

The relative energy saving rate η of VRF-PoW with respect to NC is defined as the ratio of the total energy savings to the expected total mining effort, which varies with p and k:

$$\eta = \frac{\sum_{i=1}^{n} \Delta E_i}{\mathbb{E}[X^v] \cdot H}. \tag{54}$$

Fig. 3 illustrates how the parameters p and k affect r_{adv} and η when $n = 10000$. Fig. 3(a) reveals the security threshold r_{adv}. As p approaches 1, the threshold r_{adv} increases, converging toward the strong security guarantees of the conventional PoW. Conversely, when p decreases, fewer miners advance to the post-PoW phase, resulting in a lower threshold and increased vulnerability to Sybil attacks. Moreover, k also exerts noticeable influence; increasing k raises the security threshold r_{adv}, although less effectively than increasing p. However, when p is large, a smaller k can even enhance security. Fig. 3(b) presents the corresponding energy-saving rate η. As p decreases, the proportion of miners performing the full PoW phase decreases sharply, resulting in a substantial reduction in total energy consumption, which can reach 90% in certain regions. Larger k values increase the pre-PoW overhead, reducing η.

In addition, when selecting values for p and k, the timeout probability and efficiency of Sybil attacks must be considered. If k becomes excessively small, the efficiency gap between honest and Sybil mining strategies becomes negligible, allowing miners to gain incentives to adopt the Sybil strategy for personal benefit. Therefore, k should not be set below approximately 10^{-4}. Moreover, for stable network operation, the timeout probability (i.e., the probability that the VRF process selects no miner) must remain negligible. This probability depends on the self-selection probability p and the total number of miners n, and is expressed as $P_{\text{empty}} = (1-p)^n$. This work sets the condition $P_{\text{empty}} < 0.001$ to ensure network stability. Substituting $n = 10000$ yields a minimum self-selection probability of $p_{\min} = 1 - e^{\ln(0.001)/10000} \approx 6.9 \times 10^{-4}$.

Table II summarizes the optimal parameter configurations by adversarial threshold, representing the optimal combinations of p and k that minimize energy consumption while maintaining the required security threshold.

A higher value of p improves network security, whereas lower p values enhance energy efficiency at the cost of reduced robustness. The parameter k is selected to be the smallest value under which the honest mining strategy still outperforms the Sybil strategy in terms of efficiency. Therefore, the parameters

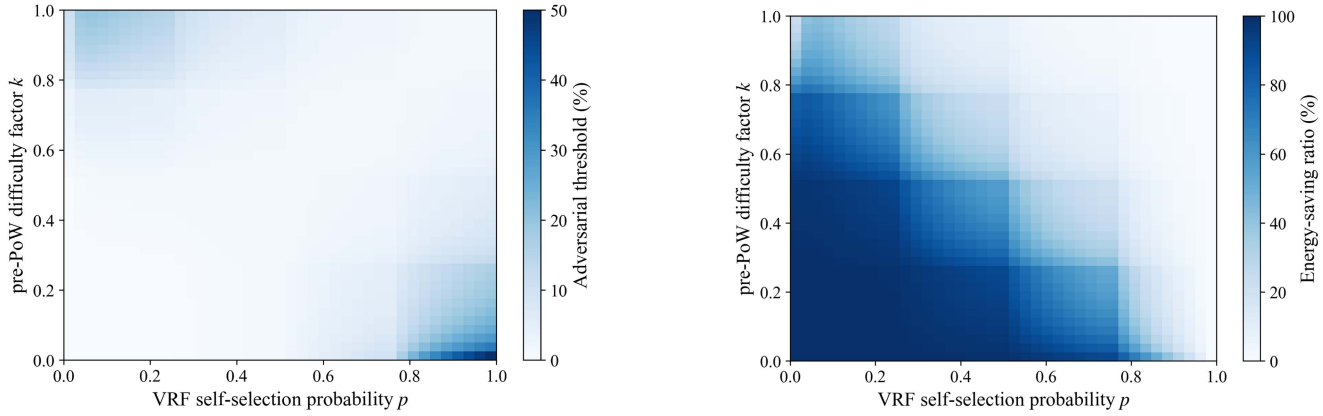

(a) Adversarial threshold r_{adv} (b) Energy Saving Rate η

Fig. 3. Variation of the adversarial threshold and energy efficiency as functions of p and k.

can be flexibly adjusted according to the network security requirements. For instance, in large-scale networks (e.g., Bitcoin), where it is practically impossible for a single entity to control more than 50% of the total hash power, the parameters can be tuned to reduce overall energy consumption. Conversely, in emerging or small-scale PoW networks that may face a higher risk of majority attacks, setting $p \approx 1$ and $k \approx 0$ is advisable to preserve the same security level as traditional PoW systems. Moreover, given the pool-suppression capability of VRF-PoW, such configurations can achieve greater security against attacks. In summary, the tunable parameters p and k allow VRF-PoW to balance energy efficiency and security dynamically, enabling adaptive optimization according to the scale and threat model of the network.

VII. Numerical Analysis

This section evaluates whether energy efficiency and decentralization can be enhanced without departing from the fundamental PoW paradigm, in which each node independently computes and propagates blocks. This work compares three protocols: the baseline NC, Green-PoW, and the proposed VRF-PoW. The EPoW, PoWR, and PoS protocols are excluded, because they operate under design principles that are fundamentally distinct from those of PoW. For instance, PoS follows a different paradigm, and EPoW limits the total hash power to save energy. Further, PoWR relies on trusted hardware, making direct comparison inappropriate. For the VRF parameters, this work adopts threshold settings of 10%, 20%, and 30%, as listed in Table II. Each configuration is denoted VRF-PoW@10, VRF-PoW@20, and VRF-PoW@30, corresponding to self-selection thresholds of 10%, 20%, and 30%, respectively. The timeout for VRF-PoW starts at approximately 2763 s ($\approx 600\log(10^2)$) and increases linearly up to 4144 s ($\approx 600\log(10^3)$), whereas Green-PoW uses a single runner-up block with a fixed timeout of 2763 s. Prior studies report that Bitcoin nodes typically maintain about 32 peer connections and a mean propagation delay of around 20 ms [80]. Accordingly, in the simulations, nodes are placed uniformly at random in a two-dimensional plane, each node connects to its k nearest neighbors ($k = 32$), and the edge latency is set to 20 ms. Hash power is assigned following a log-normal distribution, with σ set such that the top x % of miners account for roughly 1% to 50% of the total hash power. The simulation continues until 1 000 000 blocks are generated. As in Bitcoin, the difficulty is adjusted every 2016 blocks to maintain an average BGT of 600 s. The VRF execution time is negligible relative to the total BGT and is omitted from the computation. All experiments were conducted in Python 3.11 in a Linux environment. Table III summarizes the complete settings. This work verifies the consistency of the model with theoretical predictions, focusing on the BGT distribution under VRF selection and comparative energy consumption. Furthermore, this work assesses how the progressive timeout mechanism maintains liveness and fairness and compares it with the fixed timeout strategy of the Green-PoW.

TABLE III
Simulation Parameters and Settings

Parameter	Value / Setting
Nodes (N)	10,000
Simulated blocks	1,000,000
Difficulty adjustment	Every 2016 blocks (target 600 s)
Hash power distribution	Log-normal (top x% own 1% to 50%)
Network topology	k-NN (k=32) peers; latency 20 ms
Compared protocols	NC, Green-PoW, VRF-PoW@10,@20,@30

A. Results and Evaluation

Although the theoretical analysis provides mathematical evidence for the energy efficiency and security of VRF-PoW, these findings must be verified via numerical simulation. This work compares the theoretical BGT distribution and energy-saving ratio with simulation results to confirm the validity of the analytical model. Fig. 4 presents the theoretical and simulated BGT distributions for VRF-PoW, PoW and Green-PoW by hash power concentration. The simulated distributions closely follow the theoretical PDFs in terms of the mean and variance,

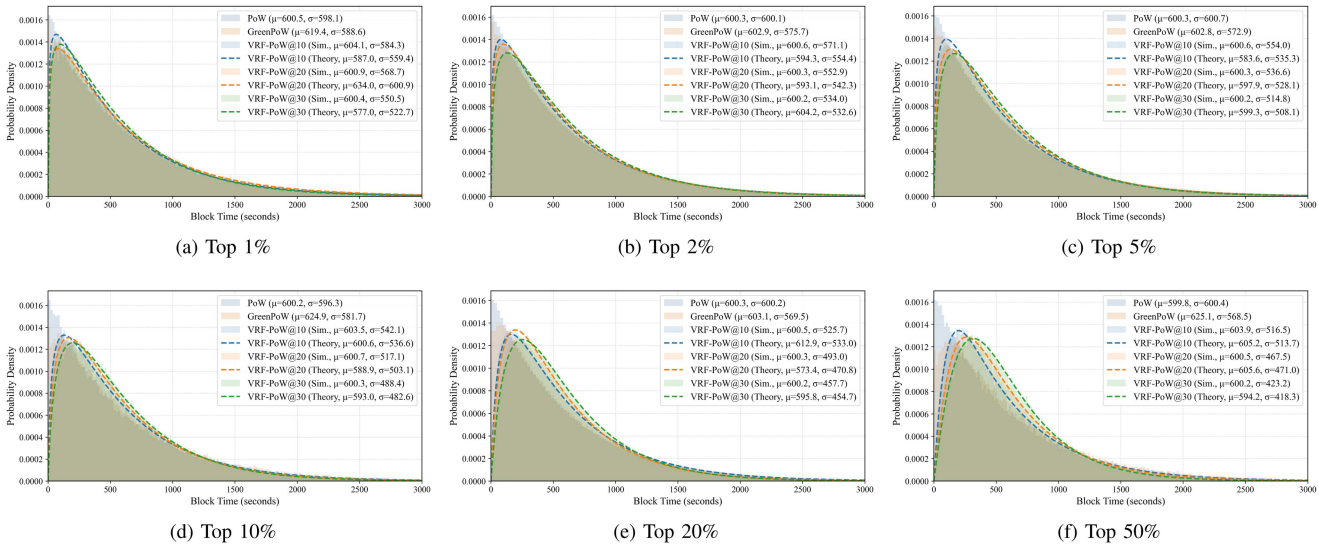

(a) Top 1% (b) Top 2% (c) Top 5%

(d) Top 10% (e) Top 20% (f) Top 50%

Fig. 4. Block generation time distributions for PoW, Green-PoW, and VRF-PoW under hash-power concentration scenarios (top $x\%$ miners hold $> 50\%$).

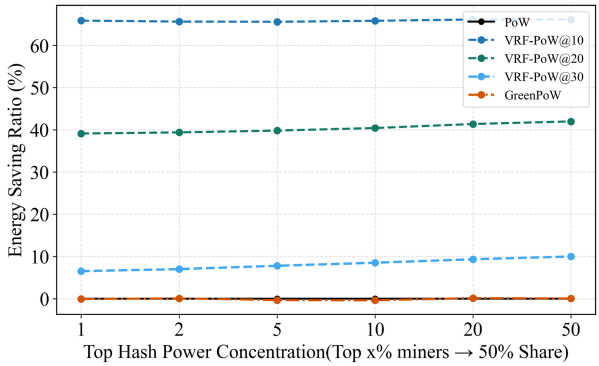

Fig. 5. Energy-saving ratio of each protocol relative to conventional PoW across hash-power concentration scenarios (top $x\%$ miners hold $> 50\%$).

supporting the accuracy of the analytical model. The VRF-PoW mechanism displays a smaller variance than PoW and Green-PoW, converging to a target BGT of about 600 s with reduced fluctuation. This lower variance minimizes the frequency of extremely short or long block times, enhancing propagation stability and network synchronization. Because of its two-phase structure, Green-PoW tends to reveal unstable difficulty adjustments, leading to longer average block times, higher variance, and irregular block intervals and delayed confirmations. As hash power concentration decreases from the top 1% to the top 50%, or as the VRF threshold increases, mining power becomes more evenly distributed among the selected nodes, further reducing the variance of BGT in VRF-PoW.

Fig. 5 compares the energy consumption of VRF-PoW and Green-PoW with that of conventional PoW. Under a decentralized hash power distribution, the observed energy savings are consistent with the theoretical efficiency η reported in Table II. As the network becomes more centralized, the energy-saving rate declines slightly, although the gap remains minor. Configured with parameters corresponding to a lower adversarial threshold, VRF-PoW achieves higher energy efficiency, whereas Green-PoW displays a much smaller reduction in energy consumption. The limited effectiveness of Green-PoW arises from its mechanism in which the uncle miner who solves the PoW puzzle second is granted the right to generate the next block. However, due to the memoryless property of the exponential distribution, the expected time until the uncle miner appears is essentially the same as the current BGT. During this period, all miners continue competing as in the standard PoW, resulting in only minimal energy savings. Therefore, Green-PoW's overall energy consumption remains nearly the same as that of NC across all hash power distributions. In contrast, VRF-PoW determines participation immediately after the VRF phase, so unnecessary competition stops much earlier, leading to substantially higher energy savings than Green-PoW. Consequently, VRF-PoW maintains stable energy efficiency across hash power concentrations, demonstrating robustness and adaptability to heterogeneous mining environments.

Fig. 6 illustrates the BGT distributions of the Green-PoW and VRF-PoW variants by hash power concentration. For a more precise visualization, kernel-density estimates with adaptive bandwidths are overlaid on the empirical distributions. Green-PoW, which applies a fixed timeout threshold, produces a distinct clustering of block-generation events near the timeout limit, yielding a sharp peak in the probability density. This clustering occurs because all miners become eligible simultaneously upon reaching the timeout, leading to multiple concurrent block discoveries and temporal crowding near the boundary. In contrast, VRF-PoW employs a progressive timeout mechanism that gradually extends block eligibility, resulting in a smoother, more continuous distribution without a sharp boundary spike. The progressive mechanism staggers miner participation, regularizing block intervals and removing the abrupt cutoff observed in Green-PoW.

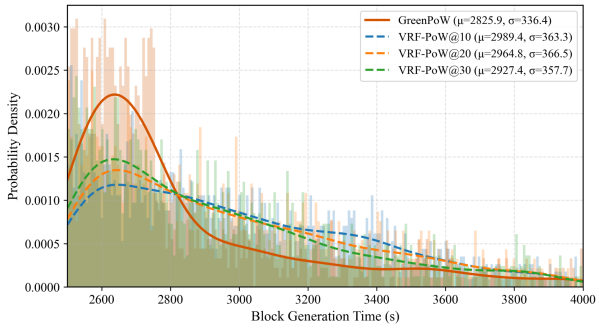

(a) Top 1%

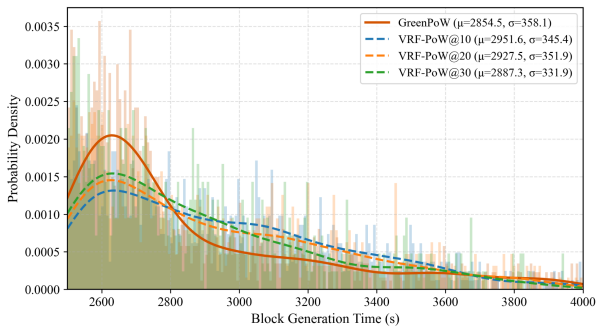

(b) Top 10%

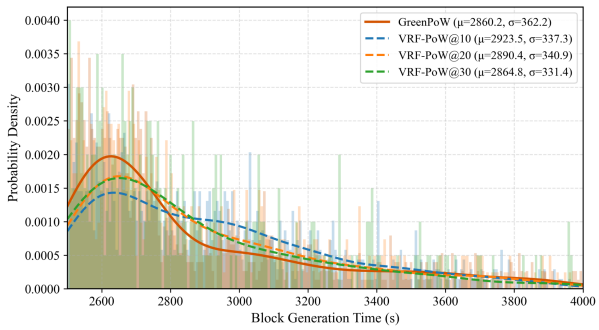

(c) Top 50%

Fig. 6. Block generation time distributions near the timeout region under fixed and progressive timeout mechanisms across hash-power concentration scenarios (top $x\%$ miners hold $> 50\%$).

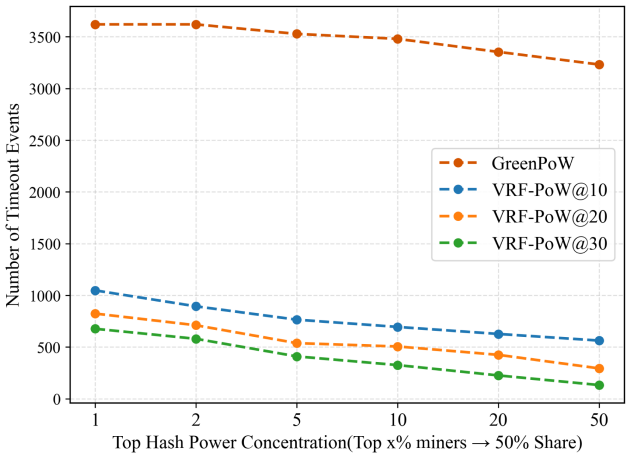

Fig. 7. Total number of timeout events under Green-PoW and VRF-PoW variants across hash-power concentration scenarios (top $x\%$ miners hold $> 50\%$), aggregated over 1,000,000 simulated blocks.

Fig. 7 compares the total number of timeout events in the Green-PoW and VRF-PoW variants. Overall, the number of timeout events decreases as hash power becomes more decentralized, because greater decentralization reduces the variance in hash power among selected miners. Green-PoW has the highest number of timeout events, exceeding 3500 even under decentralized conditions, and decreases only slightly as the hash power concentration increases. In contrast, VRF-PoW produces substantially fewer timeout events, and the count increases as the security threshold decreases, corresponding to a smaller VRF self-selection probability (p). These results indicate that VRF-PoW ensures more stable block generation than Green-PoW. In conclusion, the numerical results demonstrate that VRF-PoW achieves higher energy efficiency and greater stability in block generation than existing PoW variants.

VIII. Conclusions and Future Work

This paper proposes a novel consensus algorithm, VRF-PoW, which incorporates VRF into the traditional PoW framework to address the problem of excessive energy consumption in blockchain networks. The VRF-based miner self-selection mechanism significantly reduces energy consumption by limiting the number of miners per round. Moreover, this paper introduces a pre-PoW phase that mitigates Sybil attacks and reduces redundant message exchanges during block creation. Furthermore, implementing a progressive timeout mechanism stabilizes block-generation intervals and improves overall network reliability. The analysis indicates that VRF-PoW achieves superior energy efficiency compared with NC. Future work should focus on quantitatively evaluating the extent to which private key bound mining enhances decentralization and fairness in reward distribution. In addition, the authors plan to implement a prototype system and conduct real-world deployment tests to assess the scalability and robustness of this mechanism under realistic network conditions.

Acknowledgment

The authors used a generative AI tool (ChatGPT) to initially translate and refine the English phrasing in this manuscript.

References

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf
[2] C. Dwork and M. Naor, "Pricing via processing or combatting junk mail," in *Proc. Annu. Int. Cryptol. Conf.*, 1993, pp. 139–147.
[3] D. Fullmer and A. S. Morse, "Analysis of difficulty control in bitcoin and proof-of-work blockchains," in *Proc. IEEE Conf. Decis. Control*, 2018, pp. 5988–5992.
[4] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 3–16.
[5] C. Schinckus, "Proof-of-work based blockchain technology and anthropocene: An undermined situation?," *Renew. Sustain. Energy Rev.*, vol. 152, 2021, Art. no. 111682.
[6] S. Küfeoğlu and M. Özkuran, "Bitcoin mining: A global review of energy and power demand," *Energy Res. Social Sci.*, vol. 58, 2019, Art. no. 101273.
[7] Cambridge Centre for Alternative Finance, "Cambridge Bitcoin electricity consumption index: Comparisons," 2025. Accessed: May 23, 2025. [Online]. Available: https://ccaf.io/cbnsi/cbeci/comparisons

[8] N. Tovanich, N. Soulié, N. Heulot, and P. Isenberg, "The evolution of mining pools and miners' behaviors in the bitcoin blockchain," *IEEE Trans. Netw. Service Manag.*, vol. 19, no. 3, pp. 3633–3644, Sep. 2022.

[9] C. Wang, X. Chu, and Y. Qin, "Dissecting mining pools of bitcoin network: Measurement, analysis and modeling," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 1, pp. 398–412, Jan./Feb. 2023.

[10] C. Li, L. Wang, and H. Yang, "The optimal asset trading settlement based on proof-of-stake blockchains," *Decis. Support Syst.*, vol. 166, 2023, Art. no. 113909.

[11] S. King and S. Nadal, "PPcoin: Peer-to-peer crypto-currency with proof-of-stake," 2012. [Online]. Available: https://bitcoin.peryaudo.org/vendor/peercoin-paper.pdf

[12] V. Buterin and V. Griffith, "Casper the friendly finality gadget," 2019, *arXiv:1710.09437*.

[13] V. Buterin et al., "Combining GHOST and casper," 2020, *arXiv:2003.03052*.

[14] J. Kwon, "Tendermint: Consensus without mining," 2014. [Online]. Available: https://tendermint.com/static/docs/tendermint.pdf

[15] C. Schwarz-Schilling, J. Neu, B. Monnot, A. Asgaonkar, E. N. Tas, and D. Tse, "Three attacks on proof-of-stake ethereum," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2022, pp. 560–576.

[16] J. Brown-Cohen, A. Narayanan, A. Psomas, and S. M. Weinberg, "Formal barriers to longest-chain proof-of-stake protocols," in *Proc. ACM Conf. Econ. Comput.*, 2019, pp. 459–473.

[17] M. Platt and P. McBurney, "Sybil attacks on identity-augmented proof-of-stake," *Comput. Netw.*, vol. 199, 2021, Art. no. 108424.

[18] E. Deirmentzoglou, G. Papakyriakopoulos, and C. Patsakis, "A survey on long-range attacks for proof of stake protocols," *IEEE Access*, vol. 7, pp. 28712–28725, 2019.

[19] I. Bentov, C. Lee, A. Mizrahi, and M. Rosenfeld, "Proof of activity: Extending bitcoin's proof of work via proof of stake," *ACM SIGMETRICS Perform. Eval. Rev*, vol. 42, no. 3, pp. 34–37, 2014.

[20] T. Xue, Y. Yuan, Z. Ahmed, K. Moniz, G. Cao, and C. Wang, "Proof of contribution: A modification of proof of work to increase mining efficiency," in *Proc. IEEE Annu. Comput. Softw. Appli. Conf.*, 2018, pp. 636–644.

[21] Q. Zhuang, Y. Liu, L. Chen, and Z. Ai, "Proof of reputation: A reputation-based consensus protocol for blockchain based systems," in *Proc. Int. Electron. Commun. Conf.*, 2019, pp. 131–138.

[22] Q. Bao, B. Li, T. Hu, and X. Sun, "A survey of blockchain consensus safety and security: State-of-the-art, challenges, and future work," *J. Syst. Softw.*, vol. 196, 2023, Art. no. 111555.

[23] M. Ball, A. Rosen, M. Sabin, and P. N. Vasudevan, "Proofs of useful work," IACR Cryptol. ePrint Arch., Tech. Rep. 2017/203, 2017.

[24] M. Haouari, M. Mhiri, M. El-Masri, and K. Al-Yafi, "A novel proof of useful work for a blockchain storing transportation transactions," *Inf. Process. Manage.*, vol. 59, no. 1, 2022, Art. no. 102749.

[25] X. Qu, S. Wang, Q. Hu, and X. Cheng, "Proof of federated learning: A novel energy-recycling consensus algorithm," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 8, pp. 2074–2085, Aug. 2021.

[26] M. Mirkin, L. Zhou, I. Eyal, and F. Zhang, "Sprints: Intermittent blockchain PoW mining," in *Proc. USENIX Secur. Symp.*, 2024, pp. 6273–6289.

[27] N. Lasla, L. Al-Sahan, M. Abdallah, and M. Younis, "Green-PoW: An energy-efficient blockchain proof-of-work consensus algorithm," *Comput. Netw.*, vol. 214, 2022, Art. no. 109118.

[28] S. Yu, Y. Qiao, J. Bo, F. Yang, and S. Wang, "EPoW: Energy-efficient proof-of-work," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 6, pp. 6285–6297, Nov./Dec. 2024.

[29] P. Daian, I. Eyal, A. Juels, and E. G. Sirer, "(Short paper) piecework: Generalized outsourcing control for proofs of work," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2017, pp. 182–190.

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

[31] D. Galindo, J. Liu, M. Ordean, and J.-M. Wong, "Fully distributed verifiable random functions and their application to decentralised random beacons," in *Proc. IEEE Eur. Symp. Secur. Privacy*, 2021, pp. 88–102.

[32] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling Byzantine agreements for cryptocurrencies," in *Proc. Symp. Oper. Syst. Princ.*, 2017, pp. 51–68.

[33] Cardano, "Cardano documentation," 2025. Accessed: May 23, 2025. [Online]. Available: https://docs.cardano.org/about-cardano/introduction

[34] I. Abraham, D. Malkhi, K. Nayak, and L. Ren, "Dfinity consensus, explored," IACR Cryptol. ePrint Arch., Tech. Rep. 2018/1153, 2018.

[35] J. W. Jung, M. M. Islam, and H. P. In, "Proof of work with random selection (PoWR): An energy saving consensus algorithm with proof of work and the random selection function," *Sustainability*, vol. 16, no. 21, 2024, Art. no. 9342.

[36] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.

[37] X. Wang, S. Duan, J. Clavin, and H. Zhang, "BFT in blockchains: From protocols to use cases," *ACM Comput. Surv.*, vol. 54, no. 10s, pp. 1–37, 2022.

[38] D. Reijsbergen, P. Szalachowski, J. Ke, Z. Li, and J. Zhou, "LaKSA: A probabilistic proof-of-stake protocol," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2021, pp. 652–669.

[39] P. Daian, R. Pass, and E. Shi, "Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2019, pp. 23–41.

[40] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Proc. Annu. Int. Cryptol. Conf.*, 2017, pp. 357–388.

[41] E. N. Tas, D. Tse, F. Gai, S. Kannan, M. A. Maddah-Ali, and F. Yu, "Bitcoin-enhanced proof-of-stake security: Possibilities and impossibilities," in *Proc. 2023 IEEE Symp. Secur. Privacy*, 2023, pp. 126–145.

[42] D. Larimer, "Delegated proof of stake (DPoS)," 2014. Accessed: May 23, 2025. [Online]. Available: https://docs.bitshares.org/en/master/technology/dpos.html

[43] J. Mišić, V. B. Mišić, and X. Chang, "QPoS: Decentralized stake-based leader and voter selection in a pbft system with mobile voters," *IEEE Trans. Netw. Sci. Eng.*, vol. 12, no. 2, pp. 653–668, Mar./Apr. 2025.

[44] S. King, "Primecoin: Cryptocurrency with prime number proof-of-work," 2013. [Online]. Available: https://primecoin.io/primecoin-paper.pdf

[45] A. Shoker, "Sustainable blockchain through proof of exercise," in *Proc. IEEE Int. Symp. Netw. Comput. Appl.*, 2017, pp. 1–9.

[46] F. Zhang, I. Eyal, R. Escriva, A. Juels, and R. V. Renesse, "REM: Resource-efficient mining for blockchains," in *Proc. USENIX Secur. Symp.*, 2017, pp. 1427–1444.

[47] L. Chen, L. Xu, N. Shah, Z. Gao, Y. Lu, and W. Shi, "On security analysis of proof-of-elapsed-time (PoET)," in *Proc. Int. Symp. Stabilization Saf. Secur. Distrib. Syst.*, 2017, pp. 282–297.

[48] C. Chenli, B. Li, Y. Shi, and T. Jung, "Energy-recycling blockchain with proof-of-deep-learning," in *Proc. IEEE Int. Conf. Blockchain Cryptocurrency*, 2019, pp. 19–23.

[49] Y. Liu, Y. Lan, B. Li, C. Miao, and Z. Tian, "Proof of learning (PoLe): Empowering neural network training with consensus building on blockchains," *Comput. Netw.*, vol. 201, 2021, Art. no. 108594.

[50] Y. Wang, H. Peng, Z. Su, T. H. Luan, A. Benslimane, and Y. Wu, "A platform-free proof of federated learning consensus mechanism for sustainable blockchains," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 12, pp. 3305–3324, Dec. 2022.

[51] S. Yu, Y. Qiao, F. Yang, and J. Bo, "DPoW: A decentralized proof-of-work consensus mechanism for blockchain system," *Comput. Netw.*, vol. 270, 2025, Art. no. 111490.

[52] I. Eyal and E. G. Sirer, "How to disincentivize large bitcoin mining pools," 2014. Accessed: May 23, 2025. [Online]. Available: https://web.archive.org/web/20241212231011/https://hackingdistributed.com/2014/06/18/how-to-disincentivize-large-bitcoin-mining-pools/

[53] A. Kattis and J. Bonneau, "Proof of necessary work: Succinct state verification with fairness guarantees," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2024, pp. 18–35.

[54] R. Han, H. Lin, and J. Yu, "VRF-based mining simple non-outsourceable cryptocurrency mining," in *Proc. Int. Workshops Data Privacy Manage. Cryptocurrencies Blockchain Technol.*, 2020, pp. 287–304.

[55] M. M. Islam, M. M. Merlec, and H. P. IN, "Proof of random leader: A fast and manipulation-resistant proof-of-authority consensus algorithm for permissioned blockchains using verifiable random function," *IEEE Trans. Serv. Comput.*, vol. 18, no. 3, pp. 1655–1668, May/Jun. 2025.

[56] C. Xiong, T. Yang, Y. Wang, and B. Dong, "PoVF: Empowering decentralized blockchain systems with verifiable function consensus," *Comput. Netw.*, vol. 259, 2025, Art. no. 111092.

[57] Y. Dodis and A. Yampolskiy, "A verifiable random function with short proofs and keys," in *Proc. Int. Workshop Theory Pract. Public Key Cryptogr.*, 2005, pp. 416–431.

[58] S. Goldberg, L. Reyzin, D. Papadopoulos, and J. Vcelák, "Verifiable random functions (VRFs)," 2023, Accessed: May 23, 2025. [Online]. Available: https://datatracker.ietf.org/doc/rfc9381/

[59] R. Goyal, S. Hohenberger, V. Koppula, and B. Waters, "A generic approach to constructing and proving verifiable random functions," in *Proc. Theory Cryptogr.*, 2017, pp. 537–566.
[60] Polkadot, "Polkadot Wiki: Randomness," 2025, Accessed: May 23, 2025. [Online]. Available: https://wiki.polkadot.network/docs/learn-randomness
[61] B. David, P. Gaži, A. Kiayias, and A. Russell, "Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain," in *Proc. Int. Conf. Theory Appl. Cryptograph. Techn. Adv. Cryptol.*, 2018, pp. 66–98.
[62] C. Badertscher, P. Gaži, A. Kiayias, A. Russell, and V. Zikas, "Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2018, pp. 913–930.
[63] M. F. Esgin et al., "A new look at blockchain leader election: Simple, efficient, sustainable and post-quantum," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2023, pp. 623–637.
[64] H. Choi, S. Kim, and H.-N. Lee, "Error correction code verifiable computation consensus," *IEEE Trans. Inf. Forensics Secur.*, vol. 20, pp. 6678–6692, 2025.
[65] G. Wood et al., "Ethereum: A secure decentralised generalised transaction ledger," 2025. [Online]. Available: https://ethereum.github.io/yellowpaper/paper.pdf
[66] G. Colvin, A. Lanfranchi, and M. Carter, "EIP-1057: ProgPoW, a programmatic proof-of-work," 2018. Accessed: May 23, 2025. [Online]. Available: https://eips.ethereum.org/EIPS/eip-1057
[67] Z. Sharifian, H. Saidi, A. Fanian, and T. A. Gulliver, "A new approach to orphan blocks in the nakamoto consensus blockchain," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 2, pp. 1771–1784, Mar./Apr. 2024.
[68] M. F. Esgin et al., "Practical post-quantum few-time verifiable random function with applications to algorand," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, 2021, pp. 560–578.
[69] M. Buser et al., "Post-quantum verifiable random function from symmetric primitives in PoS blockchain," in *Proc. Eur. Symp. Res. Comput. Secur.*, 2022, pp. 25–45.
[70] M. Jakobsson and A. Juels, "Proofs of work and bread pudding protocols," in *Proc. IFIP TC6/TC11 Joint Work. Conf. Secure Inf. Netw. Commun. Multimedia Secur.*, 1999, pp. 258–272.
[71] T. Duong, L. Fan, J. Katz, P. Thai, and H.-S. Zhou, "2-HoP blockchain: Combining proof-of-work and proof-of-stake securely," in *Proc. Eur. Symp. Res. Comput. Secur.*, 2020, pp. 697–712.
[72] C. Mazzocca, A. Acar, S. Uluagac, R. Montanari, P. Bellavista, and M. Conti, "A survey on decentralized identifiers and verifiable credentials," *IEEE Commun. Surveys Tuts.*, vol. 27, no. 6, pp. 3641–3671, Dec. 2025.
[73] J. Xie, F. R. Yu, T. Huang, R. Xie, J. Liu, and Y. Liu, "A survey on the scalability of blockchain systems," *IEEE Netw.*, vol. 33, no. 5, pp. 166–173, Sep./Oct. 2019.
[74] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *Proc. IEEE Int. Conf. Peer-to-Peer Comput.*, 2013, pp. 1–10.
[75] D. Kraft, "Difficulty control for blockchain-based consensus systems," *Peer-to-Peer Netw. Appl.*, vol. 9, no. 2, pp. 397–413, 2016.
[76] S. Amari and R. Misra, "Closed-form expressions for distribution of sum of exponential random variables," *IEEE Trans. Rel.*, vol. 46, no. 4, pp. 519–522, Dec. 1997.
[77] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi, *Queueing Networks and Markov Chains: Modeling and Performance Evaluation With Computer Science Applications*. Hoboken, NJ, USA: Wiley, 2006.
[78] J. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol: Analysis and applications," *J. ACM*, vol. 71, no. 4, pp. 1–49, 2024.
[79] H. Zhu, X. Chang, J. Mišić, V. B. Mišić, L. Han, and Z. Chen, "Delay impact on stubborn mining attack severity in imperfect bitcoin network," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 3, pp. 2586–2595, May/Jun. 2024.
[80] Y. Shahsavari, K. Zhang, and C. Talhi, "A theoretical model for block propagation analysis in bitcoin network," *IEEE Trans. Eng. Manag.*, vol. 69, no. 4, pp. 1459–1476, Aug. 2022.

Seungmin Kim received the B.S. degree in computer education and blockchain security from Jeju National University, Jeju, South Korea, in 2022. He is currently currently working toward the Ph.D. degree with the School of Electrical Engineering and Computer Science, Gwangju Institute of Science and Technology (GIST), Gwangju, South Korea. His research interests include blockchain and cybersecurity.

Haeung Choi received the B.S. degree in electrical, electronics, and computer engineering from Kyungpook National University, Daegu, South Korea, in 2013, and the M.S. degree in electrical, electronics, and computer engineering in 2015 from the Gwangju Institute of Science and Technology, Gwangju, South Korea, where he is currently currently working toward the Ph.D. degree with the Gwangju Institute of Science and Technology. He is also a Researcher with Libervance Company Ltd. His research interests include blockchain and cybersecurity.

Minho Yoon received the B.S. degree in electrical engineering and computer science in 2025 from the Gwangju Institute of Science and Technology, Gwangju, South Korea, where he is currently working toward the Ph.D. degree with the Department of Electrical Engineering and Computer Science. His research interests include blockchain and artificial intelligence convergence.

Heung-No Lee (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of California at Los Angeles, CA, USA, in 1993, 1994, and 1999, respectively. From 1999 to 2002, he was a Research Staff Member with the HRL Laboratories, LLC, Malibu, CA, USA. From 2002 to 2008, he was an Assistant Professor with the University of Pittsburgh, PA, USA. In 2009, he joined the School of Electrical Engineering and Computer Science, Gwangju Institute of Science and Technology, Gwangju, South Korea. He is currently a Full Professor with the Gwangju Institute of Science and Technology. His research interests include information theory, signal processing theory, blockchain, communications/networking theory, and their application to wireless communications and networking, compressive sensing, future internet, and brain–computer interface. He was the recipient of the several prestigious national awards, including the Top 100 National Research and Development Award, in 2012, the Top 50 Achievements of Fundamental Research Award, in 2013, and the Science/Engineer of the Month, in January 2014.