SoK: Public Randomness

Alireza Kavousi
University College London

Zhipeng Wang
Imperial College London

Philipp Jovanovic
University College London

***Abstract*—Public randomness is a fundamental component in many cryptographic protocols and distributed systems and often plays a crucial role in ensuring their security, fairness, and transparency properties. Driven by the surge of interest in blockchain and cryptocurrency platforms and the usefulness of such a building block in those areas, designing secure protocols to generate public randomness in a distributed manner has received considerable attention in recent years. This paper presents a systematization of knowledge on the topic of public randomness with a focus on cryptographic tools providing public verifiability and key themes underlying these systems. We provide concrete insights on how state-of-the-art protocols achieve this task efficiently in an adversarial setting and present various research gaps that may be of interest for future research.**

1. Introduction

Public randomness is about the generation and distribution of random values that are publicly verifiable and accessible by anyone after a certain barrier point. This is in stark contrast to the notion of private randomness, which has strict confidentiality requirements and is not supposed to be shared widely, *e.g.*, as used for cryptographic key generation. While some of the earliest known use cases for public randomness date back to antiquity, the concept is also a critical building block for various modern applications that have a vital need for transparency and fairness, like lotteries [108], [132], online games [76], blockchain sharding [102], [147], timestamps [51], and more.

Public randomness can be obtained through various means with the most straightforward one being to simply gather it from a centralized entity [98], [115]. This approach, however, has obvious disadvantages that no one can check whether the provided values have actually been generated randomly or the entity can simply withhold an output. In this paper, we are concerned with randomness generation in a publicly verifiable manner and without relying on trusted third parties. Despite the apparent simplicity of the concept, it turns out that producing high-quality random values and ensuring their desired properties is a non-trivial challenge requiring deep insights from distributed systems and cryptography.

1.1. Applications

Having access to a trustworthy source of randomness is an essential part of many real-world systems. We now explore several application areas that have attracted more attention in recent years.

Lotteries and Streaming Games. The need for a reliable source of randomness is paramount in lotteries and streaming games [1], [78]. Lotteries, whether traditional [72] or blockchain-based [76], [108], [132], rely on randomness to determine the winning numbers or participants. Using proper randomness is crucial to prevent any manipulation or bias that could compromise the integrity of the process. In the context of streaming games [1], randomness plays a vital role in maintaining excitement and engagement by enabling randomly generated events, such as item drops, enemy spawns, or in-game challenges.

Single Secret Leader Election (SSLE). In proof-of-stake (PoS) blockchains, leader election plays a critical role in maintaining the security and liveness of the system. It turns out deploying a single secret leader election (SSLE) mechanism, where the unique elected leader gets to know about their leadership prior to others and can claim it in a publicly verifiable manner, is crucial [13], [34]. First, it protects the leaders from various attacks, particularly denial-of-service (DoS). Second, it obviates the need for a tie-breaking process to choose among possible set of leaders for a given time slot. The majority of the existing solutions for SSLE rely on public randomness to ensure security [94] or uniqueness [34].

Timelock Encryption. A critical characteristic for a proper source of randomness is to satisfy a regular timing for the release of its outputs [98]. Drand [4], employing the League of Entropy [2], precisely meets this requirement by currently producing an output every 30 seconds. The recent work of [75] showed how to take advantage of this feature to allow "encrypting to the future" by building a time-lock encryption [131] that does not involve sequential computation. In more detail, they observed that by viewing an output – which is a threshold VRF on an epoch number – as a secret key in an identity-based encryption scheme [35], it is possible to encrypt a message to the future with the public key being the epoch number and the corresponding secret key being the output.

Secure Aggregation. Public randomness has recently found applications for secure aggregation in the federated learning setting. Flamingo [110] is a multi-round secure aggregation protocol that deploys public randomness to establish a graph between clients. The neighbors of a client in the graph will help the server to obtain the outcome if the client becomes unavailable by unmasking its input collaboratively. In LiSA [144], public randomness is used to choose the set of clients contributing to the outcome at each aggregation round as well as the committee members with whom they derive shared keys as noise to mask their inputs.

As the need for trustworthy sources of public randomness continues to grow, a large volume of works utilizing different methods has appeared in academia and industry over the last decade or so. This paper contributes to the discussion by providing a comprehensive and systematic analysis of the notions, challenges, solutions, and techniques of state-of-the-art schemes. We acknowledge recent efforts on the topic [55], [128] and argue that this work presents complementary perspectives, distinguishing it from previous works in the following key aspects:

- First, our categorization according to primitives with public verifiability and the principal properties including security, liveness, and scalability allows for making appropriate connections across primitives and relevant properties. This allows us to observe and present various *novel insights and research gaps* as one of the main contributions of our work.
- Second, we approach the topic of public randomness from a more constructive yet less formal perspective than existing works and, in particular, elaborate on several important properties, such as *network assumptions*, *adaptiveness*, *responsiveness*, *guaranteed output delivery*, etc., that have a critical impact on the overall protocol design but that are covered either not at all or only superficially by previous works like [55]. For instance, giving up on the guaranteed output delivery allows using weaker primitives than broadcast channels [61].
- Third, our modular categorization not only facilitates the first two points but also offers a *bottom-up perspective* to readers in comparison to the top-down approach of [55]. This could make it easier, *e.g.*, for developers to decide which approach would be most suitable for their use case if they already have certain tools at their disposal or familiarity with them, instead of having to figure out the abstract concepts first before being able to decide on a technique to use.
- Fourth, we discuss various *applications* and *recent works* such as [16], [18], [20], [57], [62] that are not covered in previous SoKs.

Paper Organization. Section 2 introduces the required background, concepts, and primitives from cryptography and distributed systems. Section 3 presents the notion of distributed randomness beacon (DRB) and its relevant properties as a service for generating public randomness. Section 4 presents our systematization methodology and an overview of the four main tools to generate public randomness, which are publicly verifiable secret sharing (PVSS), verifiable random function (VRF), verifiable delay function (VDF), and public blockchain. These are then investigated in-depth in the subsequent sections. Finally, Section 9 concludes the paper with a discussion on further related properties and considerations. Also, throughout the paper, we include plenty of insights and potential research gaps labeled by "Insight" and "Gap", respectively.

2. Background

In the following we give a brief overview on the relevant preliminaries and primitives used in this paper.

Threat and System Models. We denote by n the total number of parties in a protocol and by f the number of Byzantine parties, which are those that an adversary has corrupted. Such parties are not required to follow the protocol steps and are free to act as they wish. In particular, they may collude with each other, provide wrong values, not answer at all, or follow any other malicious behaviour in an attempt to thwart the security guarantees of the protocol. An adversary may be *rushing* which means that it waits until all honest parties' messages have arrived before deciding on how to act next. We further denote by t a threshold parameter denoting the maximum number of corruptions with $f \leq t < n$. Protocols are often executed in discrete time intervals that we call *epochs*.

Network Model. We distinguish two message dissemination approaches, namely (1) via reliable authenticated channels which guarantee confidentiality and integrity of the exchanged messages and all communicating parties know each other and (2) via gossip where parties propagate messages to a random set of their peers and they do not necessarily know all other network participants. We distinguish three different communication models depending on the influence an adversary can have on the message delivery between (honest) parties. In the *synchronous* model, there is a finite known time-bound Δ for message delivery. The *partial synchronous* model relaxes this assumption by making the time-bound Δ unknown. The *asynchronous* model just makes the minimal assumption of eventual delivery without specifying any time-bound. It is worth noting that there is an equivalent treatment for partial synchrony, which divides the system into two periods; an initial period of asynchrony followed by synchrony after some unknown point in time.

Consensus. The fundamental problem of consensus [104], [122] deals with enabling a group of n parties, each with an initial input value v_i to reach agreement on a common value v despite adversarial behavior by a threshold t of them. A working consensus protocol must satisfy the two main properties: *safety/consistency*, ensuring that all honest parties agree on the same decision; and *liveness*, ensuring that the protocol eventually terminates with honest parties reaching a valid decision. Different variants of consensus problems exist (refer to [79] for more details). A well-studied one, relevant to the focus of this paper, is Byzantine fault-tolerant state machine replication (BFT-SMR) [135], which aims to maintain consistency among a group of parties while processing ever-growing inputs.

Publicly Verifiable Secret Sharing (PVSS). A (t, n) secret sharing scheme [137] allows a dealer to distribute shares of a secret s among n parties, such that any gathering of size at most a threshold t of shares reveals no information about the secret while it can be uniquely revealed by any larger subset. To enable *public verification*, where anyone (even an external party) can verify the sharing phase by the dealer and reconstruction phase by the parties, PVSS [47], [136] incorporates cryptographic primitives including encryption, commitment, and non-interactive zero-knowledge proofs.

Distributed Key Generation (DKG). To carry out cryptographic evaluations (*e.g.*, encryption, signing) jointly, a DKG protocol shares a uniformly distributed random secret sk among n parties. It provides each party with a partial secret key sk_i for performing partial evaluations, a

corresponding partial public key pk_i to verify the correctness of the partial evaluations, and a common public key pk to verify the correctness of the final evaluation. DKG is executed in a single-shot manner, meaning that it is only needed to run once to produce the required keys, which can then be used polynomially many times thereafter.

Verifiable Random Function (VRF). A VRF [116] allows one to produce a pseudorandom value along with a proof on an input using their own secret key sk such that anyone can verify its correctness using the corresponding public key pk. It can be considered as an asymmetric counterpart to a pseudorandom function (PRF) with an attached proof. With a DKG, it is possible to establish a threshold VRF, wherein only a proper subset of parties (*i.e.*, more than a threshold t) can jointly evaluate the function using their keys.

Verifiable Delay Function (VDF). A VDF [32], [126], [145] is an inherently sequential function that takes a predefined time T (*i.e.*, steps) to compute, even with a polynomial number of processors working in parallel. Given an input value x, it outputs a unique value y that can be verified efficiently by anyone in $poly(\log T, \lambda)$ time. Two well-known VDFs [126], [145] are based on repeated squaring, $y = x^{2^T}$, in a group of unknown order.

3. Distributed Randomness Beacon

The process of obtaining public randomness in a distributed manner is commonly captured in the literature via *distributed randomness beacon* (DRB), a service that produces a continuous series of randomness with public verifiability. It is straightforward, however, to see that a protocol producing public randomness can essentially turn into a DRB by executing multiple times. Although recent works such as [28], [61] have taken a formal approach to the DRB problem via game-based security definitions, a large portion of the existing literature has primarily relied on an informal treatment [133], [134], [142]. These works explore individual properties and make arguments about their fulfillment. Due to the purpose of this paper, we refrain from such formulations and refer the reader to [55] for further details.[1]

Distributed Randomness Beacon (DRB). A DRB protocol enables a group of n distrustful parties to jointly generate a series of random values with the following properties: (1) *unpredictability*, meaning the adversary cannot predict the future beacon outputs except with a negligible probability; (2) *unbiasability*, meaning the adversary cannot impact an output to its advantage; (3) *availability*, meaning the protocol should continue making progress and produce valid outputs; and (4) *public-verifiability*, meaning the correctness of the result should be verifiable even by an external party.

Following the above well-established properties, we conceptualize the problem using three rather *principal* challenges a DRB confronts. That is, *security*, capturing the secrecy of the produced values against adversarial behavior, *liveness*, capturing certain conditions that affect the progress of the protocol, and *scalability*, capturing the overhead of the protocol and its deployment at scale. We believe such distinction allows us to focus on the core concerns accordingly. Needless to say, the interconnected aspect of these challenges makes a perfect distinction difficult and some level of intersection is inevitable.

Security. Unpredictability and unbiasability are two vital security properties of a DRB protocol.

- Unpredictability: Arguably the most important property of a DRB is unpredictability, which prevents the adversary from predicting future beacon outputs according to its existing knowledge. Unpredictability can be satisfied in two flavors: *absolute* unpredictability and *probabilistic* unpredictability. The former places a hard bound $d \geq 1$ on the number of epochs when the outputs become fully unpredictable, whereas the latter only guarantees that the probability of obtaining the next outputs decreases exponentially with the number of epochs.
- Unbiasability: Ruling out the feasibility of any action by the adversary to influence the outputs to its advantage is captured by the unbiasability property. The adversarial impact on DRB can occur in different ways. One common type of such an impact is to withhold announcing contributions by an adversary after becoming aware of other parties' contributions in the hope of avoiding an undesirable output, known as *last actor attack* [91].

The baseline security for a randomness beacon is unpredictability, and not all the proposed solutions can satisfy unbiasability [27], [38], [65]. The core definition of unbiasability implies that the beacon output should be uniformly distributed across the set of possible outputs. Existing DRB protocols achieve this quality at two major types: *uniform randomness*, ensuring the beacon output is a truly random value, and *pseudorandomness*, ensuring the beacon output is computationally indistinguishable from a truly random value. A long-lasting approach for generating randomness is via the so-called *commit-reveal* method which each party commits to some local random value (*i.e.*, private randomness) to later open it and compute the resulting outcome by running some operation, *e.g.*, adding all the proposals. Although this approach provides unpredictability in its strongest type when only a *one* single honest contribution (*i.e.*, true randomness) is involved, it fails to guarantee unbiasability as the adversary may act as the last actor and decide on its action (*i.e.*, to open or not) based on the view of the protocol. Intuitively, *secure* constructions with a commit-reveal paradigm output uniform randomness that is of importance for specific applications such as PRG seed.

DRB protocols often operate in epochs, where one of the parties may be selected as a leader to coordinate and advance the protocol to the next epoch [61], [133], [134]. The leader election process is either *deterministic* when the sequence of leaders is known in advance,[2] or is *probabilistic* when the next leader is chosen randomly. Although most of the existing DRB protocols can produce randomness with absolute unpredictability, the ones with probabilistic leader election yield the weaker notion, as there is always a possibility that some corrupted party

1. However, not all the required properties have yet been defined formally, such as liveness/availability.

2. For instance, in a *Round-Robin* election the leader of epoch r is party $i = r \bmod n$.

is elected as the leader [29], [133], [134]. An immediate implication is the leader's ability to learn the next epoch's output prior to others, leading to achieving only d-unpredictability for $d \geq 2$ epochs in the future. One way to address this issue is to transform the role of the leader from a sole contributor (to the beacon output) to only a coordinator [61].

Liveness. Ensuring the continual operation of the protocol under certain conditions is a decisive aspect of a DRB. The liveness of a distributed protocol is typically evaluated based on two key parameters: *network assumption*, defining limits to what extent the adversary can cause delay in message delivery, and *fault tolerance*, indicating the maximum number of corrupted parties tolerated for successful protocol execution. The bulk of existing works on DRB focuses on the *synchronous* model. However, this cannot faithfully cover real-world scenarios as the protocol might face severe outages like denial-of-service (DoS) attacks. As a result, the study of DRB protocols in *non-synchronous* models, such as partial synchronous and asynchronous models, has gained attention [41], [61]. The partially synchronous model aims to strike a balance by offering the advantages of the others. It introduces a notion of time for convenient analysis while providing a robust network definition to handle occasional outages.

Significant progress has been made in the distributed system literature regarding fault tolerance under various network assumptions [70], [73]. In particular, the optimal fault tolerance for a BFT-SMR protocol is $t < n/2$ under the synchrony and is $t < n/3$ under the partial synchrony (and asynchrony). A DRB protocol with a set of parties on ever-growing inputs working towards establishing a totally ordered log of outputs resembles the consensus problem, particularly state machine replication (SMR) [135] with two ingredient properties of consistency and liveness.

Apart from the two principle aspects of network assumption and fault tolerance, there are two additional features relevant to liveness in the context of DRB protocols. Namely, *guaranteed output delivery* (GOD), ensuring that all honest parties obtain the beacon output at each epoch irrespective of the adversary's actions, and *responsiveness*, allowing the protocol to make progress at the actual network speed rather than under a conservative delay Δ. Although GOD is deemed to be pivotal for the proper functionality of a DRB that may need to consistently feed the end users even in the face of non-Byzantine failures (*e.g.*, temporary disconnection), not all the existing secure proposals can achieve it [61]. Due to the lack of (known) time-bound in their formulations, non-synchronous network models come with responsiveness that matters for getting higher throughput (*i.e.*, number of produced outputs per time unit). However, it has been shown that the synchronous model can only offer *optimistic* responsiveness under certain conditions, including the presence of an honest leader and $f < n/4$ [121].

Scalability. The level of reliability achieved in a distributed computing system consequently impacts the scalability. In the context of DRBs, it amounts to providing randomness at a reasonable cost depending on the number of participating parties. Asymptotically, this is typically measured using the following terms: (1) *communication complexity*, which quantifies the total number of messages exchanged among parties during protocol execution; (2) *computation complexity*, which measures the amount of local work performed by each party per output; and (3) *verification complexity*, which evaluates the work required by an external party to verify the output. Moreover, accommodating *dynamic* participation of parties also matters for a scalable system. The presence of an expensive setup phase as part of the protocol poses a serious challenge to achieving this goal.

4. Systematization Methodology

Upon reviewing the existing literature, we realize that a representative method to categorize existing works is based on their *underlying tools*. Public verifiability, as a default property, imposes constraints on the cryptographic tools employed in these constructions. This, in turn, enables us to pinpoint four of these, including PVSS, VRF, VDF, and public blockchain. The protocols deploying these tools can be further classified into *leaderless* and *leader-based* according to their communication patterns. The former requires all-to-all communication, whereas the latter designates a leader to communicate with others in a one-to-all (possibly together with an all-to-one) manner. For each category, we then evaluate different protocol designs with regard to security, liveness, and scalability. This categorization enables us to classify most of the literature coherently.[3] While we put our focus on illustrative peer-reviewed papers, where appropriate, we point to some compelling ideas presented in other works. We also present a comparison between different DRBs in Table 1. In the following, we provide a brief overview of each category before doing an in-depth analysis later.

Protocols Using PVSS. The majority of research around generating public randomness is focused on using PVSS schemes that yield the strongest quality (*i.e.*, uniform randomness). The idea of using PVSS is built upon the well-known commit-reveal paradigm [31], making it robust against adversarial attacks due to its inherent threshold security, and publicly verifiable due to the use of non-interactive zero-knowledge proof (NIZK). A major limitation of this approach is its intense communication cost due to the exchanging of large messages (*i.e.*, PVSS transcripts) in an all-to-all manner. To address this, various techniques have been proposed, including the appointment of a leader for coordination and transcript aggregation.

Protocols Using VRF. VRFs have gained popularity as a viable option for designing DRBs due to their simplicity and practicality. They are usually derived from signatures with uniqueness, a property ensuring that for each message and public key, there exists only one valid signature. The DRB's security properties are tied to the those of VRF, with DRB's unpredictability and unbiasability stemming from VRF's unforgeability and uniqueness, respectively. However, similar to the traditional commit-reveal mechanism, a VRF-based construction may be susceptible to bias if a party with prior knowledge of the beacon output chooses to abort the protocol. The common approach to deal with this issue is to make the construction

3. We highlight that our focus is merely on cryptographic solutions with public verifiability and do not consider other lines of work, *e.g.*, [66], [101].

thresholdize via deploying a distributed key generation (DKG) [123] in the setup phase. While this mitigates the bias concern, it introduces communication overhead and hampers efficient dynamic participation.

Protocols Using VDF. A promising way to produce public randomness is to deploy time-based cryptography and in particular verifiable delay functions that have efficient public verification. VDFs have found themselves useful in various blockchain-related applications since their introduction by Boneh et al. [32]. In this context, it has been proposed mainly as a method to protect against the last actor attack in a commit-reveal configuration by injecting a *computationally* guaranteed delay before releasing the output. Such a delay prevents the adversary from learning the output prior to the reveal phase. Interestingly, VDFs can also be used solely to generate public randomness, *e.g.*, through its continuous variant [71] or the version with trapdoor [133]. The security properties of the resulting DRB protocol stem from its inherently sequential characteristic and uniqueness.

Protocols Using Public Blockchain. A straightforward approach to obtaining public randomness is by extracting it from the available entropy in publicly available resources. One notable example of such resources is proof-of-work (PoW) blockchains, where parties can access high entropy values using the intrinsic randomness that lies in the mining process. Despite its simplicity, this approach has been shown to be vulnerable to different attacks that compromise the desired security properties [40].

5. Protocols Using PVSS

5.1. Security

The protocols in this category follow the commit-reveal paradigm used by the prominent work of [31], allowing two parties to jointly flip a coin and generate a uniformly random string. Due to the inherent weakness of this basic approach for guaranteeing unbiased randomness, numerous works [28], [47], [48], [61], [100], [134], [142] have employed PVSS to enable the recovery of any committed secret, relying on an honest majority assumption. In fact, it requires a quorum of parties (involving at least one honest party, *i.e.*, $f+1$) to execute the protocol by sharing some secret with uniform distribution. Since more than a threshold t of parties are involved in the randomness generation process, the existence of at least one honest contribution guarantees a *uniformly* distributed randomness.[4] Moreover, anyone can use the information posted on a public bulletin board, which is typically assumed to exist, to verify the correctness of the protocol.

Syta et al. [142] present a collection of randomness beacon protocols in an incremental way such that each one complements the previous in some respect. The first one, **RandShare** [142], runs in a single-shot manner and outputs a shared randomness to each party via collecting a proper agreed set of local randomness through VSS. The second one, **RandHound** [142], builds upon RandShare and introduces scalability improvements by randomly sharding parties into sub-groups. Within each sub-group, PVSS is executed to generate local contributions, which are then combined to produce a single beacon output. The third one, **RandHerd** [142] further enhances RandHound by augmenting it with collective signing [143] and threshold (Schnorr) signatures [141] to provide a continuous sequence of beacon outputs.

SCRAPE randomness beacon [47] follows the idea originally proposed in **Ouroboros** [100] and works by having each party run a PVSS for a randomly distributed secret and reconstruct an aggregated randomness from more than a threshold of contributions. Note that the public verification property of PVSS obviates the need for running a consensus among parties to determine the set of parties with correct sharing. In SCRAPE PVSS, parties use Lagrange interpolation in the exponent to reconstruct a group element of the form $S = h^s$, where $s \in Z_q$ and $h \in G$. **ALBATROSS** [48] takes the idea behind SCRAPE a step further to construct a protocol that generates a batch of beacon outputs at each epoch. They do so by deploying *packed Shamir secret sharing* [30] in SCRAPE PVSS. Packed Shamir secret sharing allows a dealer to share a vector of secrets of size l using a single polynomial of degree $t + l - 1$ as in the standard scheme. The parties share a tuple of l random secrets in the commit phase and compute a set of l^2 beacon outputs by locally applying a randomness extractor [56] on the correctly revealed secrets. If some parties refuse to open their secret tuples, others jointly do so. Moreover, the security properties of the randomness extractors imply that as long as there exists a proper set of uniform input values of size $n - t$, no information will be inferred about the output by the adversary choosing t input coordinates. This is in line with the honest majority assumption of the underlying PVSS scheme. One can observe that in such protocols each epoch's beacon output is determined once the set of parties with valid sharing is known. Subsequently, even if some parties in the set decide not to open their commitments in the hope of adversarially affecting the output of the protocol (*i.e.*, violating unbiasability), other parties can jointly reconstruct the secrets.

However, in some applications such as lottery, obtaining all the random values in one go is not desirable, as fresh randomness may be required. To address this issue, **GULL** [49] is introduced, which is a randomness beacon using PVSS that allows progressive release of sub-batch of beacon outputs at different steps depending on the need of the protocol. They do so by modifying the ALBATROSS in a way that instead of having parties reveal all their secrets in the reconstruction phase, they use threshold cryptography to encrypt parts of their secrets and open them gradually at further stages if needed. Thus, whenever some fresh uniformly random value is required, a new sub-batch will be revealed, which is significantly more efficient than re-executing the full ALBATROSS protocol.

Insight 1. *Unpredictability is the base level of security while pseudorandomness is strictly stronger and implies it. However, we may have pseudorandom values that are not unpredictable, if a batch of them is generated at once as in ALBATROSS [48].*

4. Note that the resulting beacon output has uniform distribution with computational security, due to the properties of PVSS.

Gap 1. *Producing a batch of uniform beacon outputs while maintaining their unpredictability as in GULL [49] is a subtle task. One alternative method to explore is designing packed PVSS with gradual release of secrets,* e.g.*, using techniques from [5].*

The authors in [134] propose **Hydrand**, a leader-based DRB where rather than having all parties perform PVSS per epoch, one single party (*i.e.*, leader) is designated to do so. In fact, either the leader opens their previously committed secret, or a threshold $t+1$ of the parties jointly reconstruct the secret using the corresponding shares. Adopting a leader-based style comes at the cost of weakening the unpredictability with two consequences. First, the leader gets to know in advance the output of the epoch in which they are selected as the leader. Second, if the adversary happens to control t consecutive leaders, it can pre-compute the next t random values ahead of all the honest parties. Therefore, Hydrand only achieves absolute unpredictability of the next $t+1$ epochs. To ensure an honest leader is indeed selected after $t+1$ successive epochs, Hydrand uses a leader election mechanism with the possibility of exempting the current leader to be selected in the next t epochs. **GRandPiper** roughly follows the same design model of Hydrand by having a leader derive the protocol and buffered PVSS shares to enable recovery in the presence of a Byzantine leader. The key idea behind **BrandPiper** [133] to improve upon GRandPiper in terms of achieving absolute unpredictability of the next epoch is to consume inputs from more than t parties for computing each beacon output. This ensures the presence of at least one honest contribution and takes away the opportunity of being the sole contributor from the leader.

The two recent works of **SPURT** [61] and **OptRand** [28] take advantage of the aggregation property of the PVSS transcript (for commitments and encrypted shares) to design a leader-based DRB, with the leader acting as an orchestrator instead of a sole contributor. This design rationale is promising as it offers absolute unpredictability of the next epoch due to the involvement of $t+1$ contributions in the beacon output. Moreover, even if a leader decides to abort, it will not affect the unbiasability property as this is a blind action without knowing the beacon output. Notice that a consensus protocol (*i.e.*, SMR) is required to get everyone to agree on a threshold set of contributions picked up by the leader.

Insight 2. *Corrupting leaders violates unbiasability if either the leader is the sole contributor with no recovery mechanism, or they can abort after observing the output prior to others. Also, corrupting the next t leaders violates unpredictability if the leader is the sole contributor.*

5.2. Liveness

The majority of works using PVSS are constructed over synchronous network [28], [29], [47], [100], [134], [142]. This is to ensure that sharing of parties' secrets arrives at the recipient within a certain time-bound to guarantee the completion of each epoch. Since randomness beacon is a continually-running service, [47] argues about the necessity of having the additional property of guaranteed output delivery (GOD) [127] as a *strong* notion of liveness. This property guarantees that (honest) parties always complete each epoch with a beacon output no matter what the adversary does. In [47], [48], [100], GOD is achieved via the use of a public bulletin board to post PVSS sharing by each party that essentially has the same effect as broadcast channels as a variant of consensus protocol [47]. These protocols can accomplish the fault tolerance of $t < n/2$ which is optimal in the synchronous setting.

In a leader-based construction where a single party advances the protocol at each epoch, achieving GOD becomes rather complicated as a (faulty) leader can withhold the beacon output or even abort the protocol. To address this challenge, [28], [29], [134] utilize the idea of *buffering* of secrets. The intuition behind buffering is that each party commits to a randomly chosen secret using PVSS in advance (*i.e.*, more than t epochs before) and updates their commitment with a fresh random value when they are selected as the leader again. The purpose of such buffering is twofold. First, it ensures GOD as no matter what the Byzantine leader decides to do, the honest parties in the protocol can jointly reconstruct the already committed secret and may also remove that leader from the protocol. Second, it ensures all the honest parties open the same committed value as the underlying consensus may require $t+1$ epochs to ensure consistency.[5] While [28], [29] are optimally resilient randomness beacons in synchrony, [134] can only satisfy fault tolerance of $t < n/3$ under the same network assumption. This limitation arises from their need for a *quorum intersection* as part of their consensus protocol to avoid equivocation by the leader.[6]

An immediate consequence of ensuring GOD is the deployment of broadcast channels [87], [127] that are only achievable in the synchronous network.[7] Recently, SPURT [61] proposed to relax the strong notion of GOD to just ensure that the adversary cannot abort the protocol *after* learning the beacon output at each epoch. This relaxation allows for a reduction in the reliance on broadcast channels and the possibility of replacing them with variants of SMR protocols [7], [146]. Roughly speaking, the leader of an epoch employs the SMR protocol to get parties to agree on a constant-sized message with the remaining heavy parts of the aggregated data being sent through private channels to each party individually to enable verification. Note that no buffering is used in SPURT as the protocol does not aim at achieving GOD.

Insight 3. *The notion of fairness in secure multiparty computation [60] is weaker than GOD. It states that if the adversary learns the output, it will not be able to prevent others from doing so, similar to SPURT [61].*

Adopting SMR has the benefit of implementing under a partial synchronous network [50], [146]. This, in turn, opens the door for *responsiveness*, a property allowing the randomness beacon protocol to proceed at the actual network speed and produce outputs faster than sticking to

5. BFT-SMR is weaker than Byzantine broadcast in the sense that it does not require committing an honest leader's proposal at each view [8].

6. Given a set of n elements, two subsets of size $n-t$ must intersect at $n-2t$ elements. So, $n > 3t$ is needed to derive a contradiction.

7. Note that this should not be confused with Byzantine reliable broadcast that can be realized in asynchrony [9].

the conservative synchrony condition. This is important for achieving higher throughput and enhancing resilience against DoS attacks. As responsiveness is a potential property of systems without strict time dependencies, the recent work of [28] aims at achieving optimal fault tolerance of $t < n/2$ in a synchronous setting while enjoying the feature of *optimistic responsiveness* [121], [140], obtaining responsiveness only in certain occasions when the leader is honest and $f < n/4$. This approach is realized through designing a BFT-SMR that supports responsive leader change by eliminating the timeout of $\Omega(\Delta)$ in BFT-SMR of [29] to detect equivocation before parties make a commit.

Insight 4. *The (partial) responsiveness in synchrony requires a reduction in the number of faults from $n/2$ to $n/4$, while (complete) responsiveness in partial synchrony can always tolerate $n/3$ faults.*

RandShare [143] is presented as an asynchronous protocol, targeting termination as a liveness property without making timing assumptions. However, as discussed in [134], it implicitly relies on a notion of time such that the protocol cannot guarantee liveness under asynchrony. The RandHound [143] protocol is driven by a single client which may abort a protocol run and enforce a restart.

5.3. Scalability

Despite offering uniform randomness and useful properties such as transcript aggregation, protocols using PVSS face a significant challenge in terms of the overhead imposed on participating parties. At a high level, every single party is supposed to commit to a secret value and broadcast the result to others. However, this broadcasting process entails sending a message of size $O(n)$, with a communication complexity of $O(n^3)$ based on the lower bound of $\Omega(n^2)$ [67]. As a result, the overall complexity becomes $O(n^4)$. Additionally, each PVSS sharing requires a computation and verification overhead of $O(n)$. A linear public verification cost is inevitable due to the need for verifying the transcript of protocol execution [61]. From a practical point of view, communication complexity is the main concern over the applicability of this line of randomness beacons, as they are inherently communication intense. The interesting work of SCRAPE [47] used the observation by [112] regarding the equivalence of (t, n) Shamir secret sharing and the Reed-Solomon Code [129] to introduce a PVSS scheme with linear computational complexity. However, their randomness beacon requires all parties to perform PVSS, resulting in an overall communication complexity of $O(n^4)$.

RandHound [142] makes use of sampling techniques to shard a system of n parties to the sub-groups of size c. Running a randomness beacon protocol within each subgroup and combining the results into the final outcome reduces the communication complexity from the total of $O(n^4)$ to $O(c^2 n)$. RandHerd [142] takes the previous approach a step further by taking advantage of collective signing [143] to generate a sequence of random values, reducing communication (and computation) complexity to $O(c^2 logn)$ and verification complexity to $O(1)$. Although sharding is a pivotal technique to reduce the overhead, it will cause overall fault tolerance reduction to ensure an honest majority for each shard [61]. Using packed secret sharing enables ALBATROSS [48] to produce a batch of $O(n^2)$ random values instead of one at each epoch, reducing the amortized communication and computation complexity to $O(n^2)$ and $O(n)$, respectively. Notice that, the resulting asymptotic boost in complexity comes with decreasing the fault tolerance as a natural trade-off, *i.e.*, $t \leq (n - l)/2$.

Recent works take advantage of their leader-based design to lower the overheads, particularly communication complexity [28], [29], [61], [134]. This is an effective approach to dramatically mitigate the use of broadcast channels by replacing all-to-all with an all-to-one/one-to-all communication pattern. In SPURT [61], the leader only broadcasts a constant-sized message (*i.e.*, cryptographic digest) and sends the bigger part of the data (*e.g.*, aggregated transcripts) over private channels, leading to a quadratic communication complexity. GRandPiper, RandPiper, and OptRand [28], [29] deploy a BFT-SMR protocol without any use of threshold signature to commit transcript of size $O(n)$, which could potentially lead to cubic communication complexity. To maintain the quadratic cost, they use error-correcting code [129] to encode the leader's proposal, and accumulators [118] to enable efficient equivocation checking. Utilizing a type of SMR that does not use threshold signature (and therefore DKG) is necessary to allow the system (efficiently) support dynamic change in the set of parties. They also present a concrete reconfiguration mechanism for their systems to maintain the resilience of the protocol after eliminating some Byzantine party.

Insight 5. *Supporting dynamic participation for underlying consensus may enable the leader-based DRB protocols using PVSS to deal with the reconfiguration procedure smoothly.*

6. Protocols using VRF

6.1. Security

Verifiable random function (VRF) [116], can be considered as an asymmetric variant of a pseudo-random function (PRF) [86], with public verifiability of the output. Intuitively, its security properties state that given a sequence of VRF values, $r_0, r_1, \ldots, r_{n-1}$, for any $n \geq 1$, it should not be computationally possible to distinguish r_n from the output of a random function. Such properties make VRF a promising fit for generating public randomness. Although the original work of [116] proposed a rather complex design, it was later shown a straightforward approach to build VRF is through applying a hash function modeled as random oracle to a unique signature scheme [65], [84]. The uniqueness of the signature matters to guarantee the security properties of VRF. Two of the most well-known unique signatures in the literature are RSA [139] and BLS [36].

Insight 6. *Applying a hash function to a signature results in a PRF in the random oracle model. When the signature is unique, it actually is a VRF with proof being the signature.*

DRBs using VRF follow a range of blueprint designs. **Algorand** [84] and **Ouroboros Praos** [65] deploy an integrated DRB protocol as part of their systems with a focus on achieving efficiency for large-scale usage rather than getting high-quality randomness. They do so by letting each party set up their own VRF key-pairs. Algorand has the taste of a leader-based protocol where the output of VRF evaluation by the leader at each epoch determines the epoch's beacon value and the next leader. Ouroboros Praos takes a slightly different approach by having the XOR of all submitted VRFs by parties to be the epoch's beacon output. Due to its randomized leader-based design, Algorand has probabilistic unpredictability. As observed by [77], in the constructions where parties individually generate VRF key-pairs, the security of the beacon output relies on an honest generation of keys. In fact, the adversary may generate a malicious key that affects the security of the output (*i.e.*, unbiasability). To address this issue, the authors in [49] design a VRF construction by adopting that of [65], guaranteeing the security of the VRF output under malicious key generation.

Unfortunately, these protocols do not satisfy unbiasability property as they are subject to the last actor attack. There are two typical options to resolve the issue. First, introducing a computational delay prior to the release of the beacon output as used in the design of **Harmony** [3]. Second, deploying VRF in a distributed setting with threshold security. When it comes to setting up a threshold cryptosystem, a distributed key generation (DKG) [123] is the first step that allows parties to obtain a common pubic key, a partial public key and its corresponding secret key. Motivated by the work of **Cachin et al.** [41], a number of protocols such as **Drand** [4], **DFINITY** [42], [91], **Glow** [77], and **GRandLine** [16] construct a DRB with continuously signing a common value (*e.g.*, epoch's number) in a threshold manner to create a series of randomness. To create a *chain* of randomness, the common value is typically considered to be the epoch number concatenated with the last beacon output so that each beacon value uniquely determines the chain of randomness all the way to the earliest one. In a threshold setting, the uniqueness property for an underlying signature scheme additionally requires any set of partial signatures of size above threshold results in the *same* signature. This consequently leads to the consistency of the produced randomness, removing the need to run some form of consensus among participating parties on the beacon output.

Insight 7. *The uniqueness of the threshold signature scheme is crucial with two main implications. First, it circumvents running a consensus. Second, it allows parties to multicast their partial signature instead of broadcasting which is a stronger requirement.*

Insight 8. *The chained and unchained DRB of [4] are different in terms of what gets signed in each epoch. However, due to the underlying threshold security both achieve the same security properties.*

The underlying threshold security guarantees no adversary controlling at most t parties neither can predict the future beacon outputs nor bias them. To instantiate a threshold VRF and thus ensure pseudorandomness, the beacon output is computed by applying a cryptographic hash function on the signature in the random oracle model [77], [91]. The basic security definition of a VRF can straightforwardly be extended to the threshold setting [77]. Due to the similarities in the concept of VRF and PRF, [77] adapts and revisits the definition of standard and strong pseudorandomness presented in [147] with respect to a threshold VRF. They provide a framework to build a DRB protocol from any threshold VRF instances in a secure way. At a high level, compared to the standard definition, the strong pseudorandomness preserves the security against an adversary with the additional power of getting partial evaluations on its challenges, selecting the corrupted parties' local secret keys, and influencing the public parameters computation. As observed by [77], all the constructions in [3], [65], [84] fail to satisfy two types of pseudorandomness.

Gap 2. *One caveat with using threshold VRF is that the initial seed cannot be used to generate unbiased randomness forever as the entropy is limited. It is worth exploring the process of quality degeneration over time and the way to handle it efficiently.*

Gap 3. *Adopting a leader-based design with VRF while handling the bias through another route helps to avoid a (DKG) setup phase. One possible approach would be deploying digital signatures with key extraction [11].*

Mt. Random [49] presents a DRB with three layers each one providing a different type of randomness in exchange for different security/performance trade-offs. More accurately, the first layer generates uniform randomness; the second layer generates pseudorandomness, and the last layer generates (bounded) biased randomness. They design such a construction by nicely combining different primitives including PVSS, threshold VRF, and VRF in each layer respectively. Moreover, each lower layer feeds the higher one with seeds to provide a consistent level of bias across the structures. Beaver et al. [22] proposed a DRB protocol called **STROBE**, where the beacon output at each epoch x_r is computed by repeated RSA decryption of the previous epoch's output x_{r-1} in a distributed way. In fact, after running a setup phase assuming a dealer who generates the RSA modulus N and Shamir decryption key shares sk_i, each party disseminates their contribution x^{sk_i} and the epoch's output is computed by aggregating a threshold number of valid contributions through performing a Lagrange interpolation in the exponent. This construction is an extension to the work of [23] which was the first distributed randomness beacon producing values by repeated squaring of a random seed. Distributed generation of RSA modulus [52] is an old but still active line of research that can be used to get rid of the trusted dealer in this protocol.

Gap 4. *Inspired by the construction in [49] that outputs different types of randomness, it would be interesting to build a construction that generates beacon outputs under various thresholds, where a higher threshold could imply better randomness quality. Adopting the "Multiverse" DKG proposed in [19] to generate a unique signature could essentially lead to such DRB.*

Gap 5. *One direction to explore is the effects of long-range attacks [14] on the DRBs using (threshold) VRF. What makes the problem sophisticated here is the fact that such attacks might be unobservable due to the pseudorandomness property of the beacon outputs.*

6.2. Liveness

The DRB protocols in [65], [84] are built on top of their underlying distributed ledger, with the former using a BFT-type consensus in partial-synchrony with $t < n/3$ and the latter having parties communicate in a synchronous peer-to-peer fashion with $t < n/2$.[8] In a stand-alone DRB protocol like [99], the fact that parties evaluate their VRFs individually necessitates deploying Byzantine reliable broadcast [9] that works in asynchrony.

Protocols deploying threshold VRF, however, do not need to use broadcast channels for producing their beacon outputs. This implicitly results in constructions that can be implemented in a *non-synchronous* setting while supporting the corresponding fault tolerance (*i.e.*, $t < 1/3$). Therefore, DFINITY [42], [91] is presented in partial-synchrony and Drand [4] in synchrony, despite having the same protocol flow. However, several existing DRB protocols using threshold VRF, such as [22], [49], [77], make synchronous assumptions due to the use of DKG in their setup phase [80]. The context of asynchronous DKG [6], [64] has recently started receiving more attention which has a direct effect on constructing asynchronous DRB protocol. The protocol proposed by Cachin et al. [41] works independently of network delay and therefore is suitable for an asynchronous setting.

Insight 9. *Designing a DRB protocol in asynchrony is tricky due to the fact that randomness itself is needed to get around FLP impossibility [73]. With an asynchronous DKG, however, it is possible to generate public randomness via computing a threshold VRF in a single-shot manner and then repeat the process multiple times.*

One important consideration in protocols with a chain of randomness is that producing beacon output at each epoch is necessary for initiating the next one. This means if the protocol fails to output at any epoch, the next epochs will also be influenced. Thus, satisfying GOD property is necessary for the correct operation of such protocols. Protocols like Algorand [84] and Ouroboros Praos [65] where parties individually evaluate their VRFs, cannot feature GOD due to the possibility of abortion.

6.3. Scalability

DRBs with individually set VRFs such as [84], [99], [100] are quite efficient in terms of computation and verification cost with a constant overhead of $O(1)$. However, they still need a quadratic communication complexity for broadcasting or deploying a public bulletin board.[9]

Running a DKG is the most expensive part of protocols using threshold VRF that dominates the communication and computation complexity. There are various DKG constructions in the literature among which [81], [88], [95], [103] are widely used. DFINITY uses the non-interactive distributed key generation of [88] with $O(n^3)$ communication complexity, while Drand employs [81] that incurs at least $O(n^3 \log n)$ overhead. As a popular option, generating key materials for threshold BLS over internet takes $O(n^3)$ communication complexity with an asynchronous DKG [6], [63]. Running such a setup phase is for once and when done, the parties can perform the rest of the process at a much lower cost. If needed, one can also use a publicly verifiable DKG [89] in the setup phase. The aggregation property of the underlying PVSS enables this DKG to offer lower communication and computation costs (with a gossiping approach for dissemination) and can be used to instantiate verifiable unpredictable function (VUF), which essentially works as a randomness beacon.[10] GRandLine [16] uses aggregation techniques to design a DKG with $O(n^2 \log n)$ communication cost. It achieves this by first separating the participant parties into two subsets and then running the protocol in each of the subsets recursively to produce two transcripts. Finally, all parties receive these transcripts and end up with a single aggregated PVSS transcript. Their techniques are inspired from the recursive Phase-King [107], where the protocol terminates in two phases instead of $t+1$ given the honest majority of at least one of the halves. Excluding the DKG setup, the DRB protocols using threshold VRF are way more efficient compared to the ones using PVSS. However, all the protocols of [22], [41], [91] still have a quadratic communication complexity of $O(n^2)$ due to multicasting a partial signature by each party to others at each epoch. Each party then needs to verify the received set of messages, leading to a linear cost. An external verifier just needs to verify the final beacon output against the common public key, providing optimum public verifiability of $O(1)$. In GRandLine [16], however, the final output is not efficiently verifiable and one needs to instead verify all of the partial signatures to achieve public verifiability. One major limitation of having DKG as the setup phase is an inability to support dynamic participation, due to the need for re-running the DKG whenever a new party joins or leaves.

Insight 10. *Achieving a sub-cubic cost in the pre-processing phase of GRandLine [16] sets it apart from Drand [4] as a randomness beacon. However, it comes at a higher cost of verification compared to Drand due to their locally verifiable threshold signature, establishing a trade-off.*

7. Protocols Using VDF

7.1. Security

Although time-based primitives have been around for decades [131], the evolution of VDF in recent years [32], [71], [145] as an interesting variant has opened new

8. In the original work of [65] the network was assumed to be "semi-synchronous". Later it was shown that a longest-chain consensus requires synchronous assumption to work safely [130].

9. Algorand [84] uses random committees of size c. This makes the communication complexity reduced to $O(cn)$, however, asymptotically remains quadratic if c depends on n.

10. The DKG of [89] outputs group element as secret key and is not suitable for BLS signature that requires a field element.

doors to take advantage of their promising features to produce public randomness [32], [71], [90], [133]. The most common VDF constructions are those based on repeated squaring in groups of unknown order, like an RSA group, to enforce a guaranteed computational delay [126], [145]. The properties of VDF including uniqueness, sequentially, and public verifiability make it a solid option to construct a randomness beacon by simply using its output iteratively evaluated on an initial seed at regular points in time. It is worth mentioning that the sequentially of VDF implies unpredictability and its uniqueness together with timed-dependent evaluation implie unbiasability. However, a VDF provides efficient verification for a complete invocation, and not for every iteration. That is the main motivation behind the design of the Continuous VDF (cVDF) [71] that addresses this issue by introducing a primitive that enables efficient verification of a VDF at each iteration, which is independent of the time parameter. Thus, starting with an unpredictable seed, cVDF is iteratively applied to produce the epoch's output which is also used as the seed for the next epoch. As with VRF, one needs to apply a cryptographic hash function on the resulting VDF output to turn beacon outputs pseudorandom [71].

Gap 6. *The construction using cVDF in [71] is a centralized randomness beacon. One could explore how to use such a primitive with iterative public verification to build a DRB. Using the notion of collaborative VDF [114] may be a direction to look into this further.*

The first concrete attempt towards constructing a DRB using VDF was due to the researchers at Ethereum Foundation [69]. Their proposal is based on applying a VDF on the aggregated local randomness from a set of parties through the commit-reveal process. The necessity of computing VDF prevents the adversary from learning the beacon output before revealing its input. **HeadStart** [105] is a recent randomness beacon protocol that follows a commit-reveal approach where a centralized organizer computes the VDF. To protect against a colluding organizer, it needs to compute the Merkle root of all the contributions and publishes the respective membership proofs prior to the release of the beacon output. Particularly noteworthy is the fact that having only one honest (*i.e.*, random) contribution is enough to ensure the security of the resulting beacon value. Very recently, Christ et al. [57] proposed **Cornucopia** that essentially follows the same design of HeadStart, but with a more generic approach (*i.e.*, using any accumulator) and elaborate security considerations.

RandRunner [133] constructs a DRB protocol by cleverly making use of *trapdoor* VDF, enabling fast computation of VDF with some additional knowledge. The core idea behind the construction is to have each party designated as the epoch's leader efficiently compute VDF on the last beacon output using their trapdoor and broadcast the result to others. The crucial part is to ensure that the knowledge of the trapdoor does not give the leader an opportunity to threaten the uniqueness of the VDF, and is merely used to accelerate the process of computation. To this end, given the security properties of [126] as the underlying VDF construction, a setup phase is required where each party must generate a NIZK [43] to show the correctness of the RSA modulus.[11] Observe that this procedure allows getting around an *interactive* setup phase for RSA modulus generation as each party just needs to *individually* set up their own VDF. An important consideration is the necessity of ensuring *strong* uniqueness for trapdoor VDF in case public parameters are adversarial generated, similar to the setting where parties set up their VRFs individually [77]. With a deterministic leader election, the protocol provides absolute unpredictability only after $d = \alpha \cdot t$ epochs which α denotes the adversary's computational advantage in evaluating VDFs compared to that of honest parties.

Gap 7. *As the only leader-based protocol using VDF, RandRunner [133] ensures absolute unpredictability of the next $t+1$ epochs. Further research is needed to reduce this preferably to one epoch.*

A recent attempt to design a DRB according to the commit-reveal paradigm is **Bicorn** [54], featuring an interesting approach to address the last actor attack efficiently. The subtle novelty of the approach is to enable the recovery of missed contributions just by doing *one* sequential computation, no matter how many parties abort. This is achieved by having each party commit to their contribution α_i as $c_i = g^{\alpha_i}$ with the resulting beacon output being $\prod_{i=1}^{n} h^{\alpha_i}$, where $h = g^{2^T}$ is a VDF evaluation with parameter T. In the event of abortion, this design allows anybody to compute the beacon output via $\prod_{i=1}^{n} (c_i)^{2^T}$, resembling a VDF computation on the product of commitment $\prod_{i=1}^{n} (c_i)$. Bicorn [54] has a similar flow as leaderless protocols using PVSS in the sense that they both follow a "commit-reveal-recover" paradigm [54], but the former handles the recovery with a slow computation and the latter does that with reconstruction due to its threshold security.

Insight 11. *Bicorn [54] shows the importance of the way beacon output is determined in a commit-reveal manner. Using an exponentiation operation results in an (optimistically) efficient protocol with optimal timed recovery, i.e., one slow computation for many aborts.*

7.2. Liveness

Time-based cryptography allows circumventing the result of [59], showing the impossibility of doing a secure coin flipping without an honest majority. That is, dishonest majority can be tolerated due to the possibility of timed recovery and security can be preserved due to the sequential nature of time-dependent computation, even against parallel processors. Although time-based cryptography can potentially relax the underlying network assumption and obtain the highest fault tolerance of $n-1$ [55], [113], there exists some hurdle in realizing these exciting properties. Working in a synchronous setting seems to be critical for the current DRB protocols using VDF [54], [69], [105], [133] for two reasons. First, the reliance on time/timeout may be needed to indicate different stages of the protocol [54], [69], [105]. Second, the use of a public bulletin board or broadcast channel is assumed in [54], [133] to ensure

11. It should be the product of two safe primes of form $2p+1$.

consistency or security.[12] More precisely, in Bicorn [54] parties should have access to a public bulletin board to consistently post their commitment in the commit phase. In RandRunner [133], the corrupted leader may try to violate the unpredictability of the protocol by *selectively* disseminating the beacon output to a portion of parties, forcing others to go through slow computation to catch up. Deploying broadcast channels/reliable broadcast in the system is a countermeasure that consequently implies a fault tolerance of $t < n/2$. Due to the strong uniqueness of the used VDF that ensures the equality of the VDF evaluation normally and with a trapdoor, the only aspect of the protocol suffered from periods of asynchrony (*i.e.*, network partitions) is the unpredictability as anybody can compute the missed beacon output by a slow VDF computation, offering unbiasability and GOD.

7.3. Scalability

The major issue with this type of DRB is the need for performing sequential computation which is a highly energy-consuming task, introducing latency and dampening throughput. The idea of trapdoor VDF in [133] allows a fast beacon computation for the leader and efficient verification of $O(1)$ for verifiers. The same is the case in [54], except in the leaderless fashion with $O(n)$ verification cost. The necessity of deploying broadcast or public bulletin board in the system leads to a cubic communication complexity for [54] and quadratic communication complexity for [133] due to its leader-based style. Gossiping is an alternative approach that has lower complexity of $O(n \log n)$ but increases the latency. While the set of parties is known and fixed in [133], the two works of [54], [105] support *public* participation, allowing parties to efficiently come and go without knowing the set of parties in advance.

8. Protocols Using Public Blockchain

8.1. Security

Some of the earlier works for building randomness beacons were based on using a source of information as high entropy data that is publicly available [27], [38], [58], [106], [125]. High-quality randomness is typically extracted by applying a randomness extraction function to parts of the corresponding data. Among the existing public data structures, proof-of-work (PoW) blockchain is considered the most suitable one as it is always available (unlike financial markets in [58]) and inherently comes with additional security properties like the underlying Nakamoto consensus [117]. The randomness used in the process of mining a block (*e.g.*, PoW puzzle) can be utilized to extract a large number of unpredictable random bits. However, as mentioned in [40], two types of manipulation attacks can be launched: (1) a miner could withhold proposing a valid block just because it does not lead to the desired randomness; [13] (2) due to the network latency it is possible that forks occur. One might also consider an attack scenario to affect the beacon output just by manipulating the network to prevent or delay the propagation of a particular block producing an undesirable output. Although imposing financial penalties or slashing is a well-known way to restrict miners from such manipulation attacks [12], [15],[14] the possible gain from attacks may be unbounded while any penalty is bounded. A detailed security analysis on the ability of malicious miners is carried out in [27], showing the impossibility of deriving even one single truly random bit when the attacker has a considerable fraction of the total computing power. On the other hand, when the attacker has a limited computing power that is not enough for block production, generating truly random bits is feasible. The authors in [119] present a model for uniform randomness extraction over a public blockchain that has each party output a public value together with secret values to the other parties in a multi-round fashion.

As common in the DRB literature, incorporating a delay function is a working method to protect against a malicious miner via imposing a delay period only after which the output is achievable. [40], [106] proposed two protocols for augmenting such delay functions based on computing modular square root and iterating a pseudorandom permutation (*e.g.*, block cipher) or hash function as compositionally-sequential functions. Notice that despite having a (randomized) leader-based style with the miner being the epoch's leader, these protocols can provide absolute unpredictability, even for the current epoch's beacon output, thanks to intrinsic randomness lied in the system (plus the use of delay). **RandChain** [90] combines security properties of Nakamoto consensus and delay functions to address two main issues in a PoW blockchain-based system including biasability, and unfairness, *i.e.*, parties with high computational power dominating the randomness generation. It introduces a primitive called SeqPoW which is a puzzle that, unlike the typical PoW puzzle, cannot be solved faster using multiple parallel processors. To put it another way, SeqPoW is a cryptographic puzzle that takes a random and unpredictable number of sequential steps to solve. Similar to VDFs, sequentiality in SeqPoW also implies unpredictability.

Insight 12. *Proof-of-work mechanism in Nakamoto consensus prevents equivocation by the leader (e.g., miner) via making it costly, somewhat resembling the use of (threshold) signature in BFT-type consensus.*

8.2. Liveness

As Nakamoto consensus works safely only in synchrony [130], the randomness beacon protocol built on top of it inherent the same network assumption [38], [40], [90]. More precisely, a non-synchronous network condition (partial-synchronous or asynchronous) leads to arbitrarily long network partitions, violating the consistency of the system. Moreover, these protocols can tolerate a (computationally) honest majority of $t < n/2$.

12. This requirement, however, could be lifted by deploying Byzantine reliable broadcast.

13. It is however costly, as they require to spend considerable time and computational efforts on this action.

14. We note that slashing is an established method for a proof-of-stake (PoS) setting where the actors have some stake already deposited in the system.

8.3. Scalability

The properties of DRBs using public blockchain rely on their underlying distributed ledger and cannot be treated in a stand-alone regime. In [38], [40], [90] the party who first finds the solution to a puzzle, publishes the result *globally* to the peer-to-peer network. In such a network, the message is delivered by a sender to a random subgroup of parties (*i.e.*, its peers) in different steps until all parties receive the message. Thus, the parties do not necessarily know each other in contrast to the network with point-to-point channels. Although the availability of a (public) distributed ledger facilitates the process, one may argue that posting to a blockchain could be more costly compared to implementing broadcast channels. That is, the former usually contains a large population of parties that are not necessarily involved in the randomness generation, but the latter is only concerned with a limited set of participating parties. Verification should be done efficiently just by checking the correctness of the solution, resulting in $O(1)$ cost. In these protocols, the participating parties may dynamically change over time, featuring public participation.

9. Discussion

We now look at some aspects of DRB protocols that are worthy enough to be highlighted independently from the systematization already provided. Due to the limitation in space, we provide further discussion in Appendix A.

Adaptive Security. There are two widely known strategies that an adversary can adopt to corrupt parties: *static corruption*, where the set of corrupted parties is fixed and known to the adversary before the protocol begins; and *adaptive corruption*, where the adversary can corrupt any party it wishes based on its view of the protocol. The adversary's capability is still limited to corrupt only a certain number of parties and they remain corrupted until the end of protocol [46], [109]. In [120], two levels of adaptiveness are considered for an adversary, namely *fully* and *mildly*. The former refers to a strong adversary that may corrupt the victims instantly while it takes some time for the latter to corrupt a new party. Most of the public randomness protocols in the literature consider the static adversary. In fact, providing proven security against an adversary enjoying adaptiveness is rather delicate and requires more work. RandRunner [133] offers absolute unpredictability of the next $t+1$ epochs for both static (worst-case) and adaptive adversary while it offers probabilistic unpredictability against a *mild* adaptive adversary since the protocol can make progress before the adversary gets a chance to corrupt new leaders.

Insight 13. *In a leader-based protocol with the leader solely contributing to the beacon output, an adaptive adversary can corrupt up to t consecutive leaders and violate the unpredictability. Therefore, the adaptive adversary essentially acts as a static adversary in a worst-case scenario.*

As mentioned in [99], the single or multi-round design of the protocol plays a critical role in providing adaptive security. A single-round protocol where each party only speaks once during each epoch can provide adaptive security as the best the adversary can do is to try and randomly corrupt parties before they contribute; otherwise, any corruption would be useless. This model is known and formalized under the notion of YOSO, you only speak once [26], [82]. Therefore, the important observation in designing an adaptively secure leader-based protocol is that the adversary should not be able to detect the next leader ahead of disseminating the result. Given this, Algorand [84] provides probabilistic unpredictability in the presence of an adaptive adversary. The work of **Kiayias et al.** [99] can be thought of as a parallelization of Algorand where instead of having a leader at each epoch, all parties individually compute the VRF on a common seed and broadcast it to others. Eventually, the hash of the least k submissions (*i.e.*, VRFs with smallest values) determines the beacon output. This approach eliminates the risk of a malicious adversary having complete control over all contributors, as the probability of the k-th smallest adversarial contribution being smaller than an honest one is bounded when the parameters are appropriately chosen. This ensures the inclusion of at least one honest contribution to secure the beacon output. Consequently, the protocol remains secure even when an adaptive adversary corrupts up to $t < n/2$ parties.

Gap 8. *One could explore how to construct an adaptively secure DRB protocol with absolute unpredictability, where the leader solely contributes to the beacon output at each epoch. Such protocol should likely resemble the Nakamoto-style structure, where by the time the adversary decides to corrupt a party, they have already revealed the contribution.*

Leaderless protocols using VRF, in particular those based on threshold signature, can be viewed as a single-round protocol excluding the one-time DKG phase. So, adaptiveness should be investigated in the combination of the two sub-protocols. As realized in [46], one major problem with adaptive security for threshold scheme is constructing a simulator to simulate the adversary's view in the real execution of the protocol as it is hard to predict which subset of parties is corrupted.[15] The authors in [10], [46] use some techniques such as erasing the secrets, rewinding, and zero-knowledge proof to achieve an adaptively secure threshold RSA signature on top of a static-secure one. A requirement of such transformation is the refreshment of partial secret keys after each signature generation that additionally leads to proactive security [93], [111], a level of security required when parties are under threat of losing/leaking their secret keys. Recently, Bacho and Loss [17] introduced an adaptive security proof in algebraic group model (AGM) [74] for threshold BLS signature which could essentially bring the respective randomness beacon protocols the joy of adaptiveness.

It is rather challenging to argue about the adaptiveness of protocols using PVSS despite not knowing concrete attack against them [61]. It is clear, however, to argue that in such multi-round protocols parties go through at least two logical steps of commit, and reveal. So, to provide security against an adaptive adversary using more than t contributions from distinct parties is inevitable. This

15. Unless we assume that the simulator knows all the private shares of parties which do not make sense.

TABLE 1: Comparison of distributed randomness beacons (DRBs).

Category	Protocol	Security		Liveness		Scalability			Setup	Dynamic	Adaptive	GOD	Responsive
		Unpre.	Unbias.	Net.	Faults	Comm.	Comp.	Veri.					
PVSS	**RandShare*** [142]	✓	✓	async.	$n/3$	$O(c^2n)$	$O(c^2n)$	$O(c^2n)$	CRS	✓	✓	✓	✓
	RandHound [142]	✓	✓	sync.	$n/3$	$O(c^2n)$	$O(c^2n)$	$O(c^2n)$	CRS	✓	✗	✗	✗
	RandHerd [142]	✓	✓	sync.	$n/3$	$O(c^2 \log n)$	$O(c^2 \log n)$	$O(1)$	DKG	✗	✗	✓	✗
	Ouroboros [100]	✓	✓	sync.	$n/2$	$O(n^4)$	$O(n^3)$	$O(n^3)$	CRS	✓	✗	✓	✗
	SCRAPE [47]	✓	✓	syn.	$n/2$	$O(n^4)$	$O(n^2)$	$O(n^2)$	CRS	✓	✗	✓	✗
	Hydrand [134]	$t+1$	✓	sync.	$n/3$	$O(n^2)$	$O(n)$	$O(n)$	CRS	✓	✗	✗	✗
	ALBATROSS [48]	✓	✓	sync.	$n/2$	$O(n^2)$	$O(n)$	$O(1)$	CRS	✓	✗	✓	✗
	GULL [49]	✓	✓	sync.	$n/2$	$O(n^4)$	$O(n^2)$	$O(n^2)$	DKG	✗	✗	✓	✗
	GRandPiper [29]	$t+1$	✓	sync.	$n/2$	$O(n^2)$	$O(n^2)$	$O(n)$	CRS††	✓	✗	✓	✓
	BRandPiper* [29]	✓	✓	sync.	$n/2$	$O(n^3)$	$O(n^2)$	$O(n^2)$	CRS††	✓	✓	✓	✓
	SPURT [61]	✓	✓	p.sync.	$n/3$	$O(n^2)$	$O(n)$	$O(n)$	CRS	✓	✓	✗	✓
	OptRand [28]	✓	✓	sync.	$n/2$	$O(n^2)$	$O(n)$	$O(n)$	CRS††	✓	✓	✓	✓
VRF	**Cachin et al.** [41]	✓	✓	async.	$n/3$	$O(n^2)$	$O(n)$	$O(1)$	DKG	✗	✗	✓	✓
	DFINITY [91]	✓	✓	p.sync.	$n/3$	$O(n^2)$	$O(n)$	$O(1)$	DKG	✗	✓	✓	✓
	Drand [4]	✓	✓	sync.	$n/2$	$O(n^2)$	$O(n)$	$O(1)$	DKG	✗	✓	✓	✗
	Algorand [84]	$\Omega(t)$	✗	p.sync.	$n/3$	$O(cn)$	$O(1)$	$O(1)$	CRS	✗	✓	✗	✓
	Ouroboros Praos [65]	$\Omega(t)$	✗	sync.	$n/2$	$O(n)^\dagger$	$O(n)$	$O(n)$	CRS	✓	✓	✗	✗
	Harmony [3]	$\Omega(t)$	✓	p.sync.	$n/3$	$O(n^2)$	VDF	$O(n)$	CRS	✓	✗	✗	✗
	Glow [77]	✓	✓	sync.	$n/2$	$O(n^2)$	$O(n)$	$O(1)$	DKG	✓	✗	✓	✗
	STROB [22]	✓	✓	sync.	$n/2$	$O(n^3)$	$O(n)$	$O(1)$	DKG	✗	✗	✓	✗
	Kiayias et al. [99]	✓	✓	sync.	$n/2$	$O(n^3)$	$O(n)$	$O(1)$	CRS	✓	✓	✗	✗
	GRandLine [16]	✓	✓	sync.	$n/2$	$O(n^2)$	$O(n)$	$O(n)$	DKG	✗	✓	✓	✓
VDF	**RandRunner** [133]	$t+1$	✓	sync.	$n/2$	$O(n^2)$	VDF	$O(1)$	CRS	✓	✗	✓	✓
	HeadStart [105]	✓	✓	sync.	n	$O(n)^\dagger$	VDF	$O(1)$	-	✓	✗	✗	✗
	Bicorn [54]	✓	✓	sync.	n	$O(n^3)$	VDF	$O(n)$	-	✓	✓	✓	✗
	Cornucopia [57]	✓	✓	sync.	n	$O(n)^\dagger$	VDF	$O(1)$	-	✓	✗	✗	✗
Blockchain	**Bitcoin** [38]	✓	✗	sync.	$n/2$	$O(1)^\dagger$	PoW	$O(1)$	CRS	✓	✓	✗	✗
	Proof-of-Delay [40]	✓	✓	sync.	$n/2$	$O(1)^\dagger$	VDF	$O(\log \lambda)$	CRS	✓	✓	✗	✗
	RandChain [90]	✓	✓	sync.	$n/2$	$O(1)^\dagger$	PoW	$O(1)$	CRS	✓	✓	✗	✗

* =: VSS is used; † =: Public bulletin board is assumed; †† =: Private setup is assumed.

is the way BrandPiper [29] achieves adaptive security, employing the communication-wise efficient VSS scheme of [97] with more than t contributions involved in producing the beacon output at each epoch. With this in mind, the authors in [18] very recently presented new security definition for *aggregatable* PVSS fulfilled against an adaptive adversary in the algebraic group model. Aggregation refers to homomorphic combination of multiple PVSS transcripts into one aggregated transcript that shares the sum of the corresponding secrets. This implies adaptive security for some of the existing DRBs in the literature such as [28], [61].

Gap 9. *Investigating the adaptive security for non-aggregatable PVSS and proving it for aggregatable PVSS under standard and less strong assumptions are two research questions. This, in turn, leads to analyzing the adaptiveness of resulting DRB protocols.*

History Generation. STROBE [22] put forth an interesting property for a DRB regarding the efficient (re)generation of the beacon history, given the current epoch's output. This novel feature is of importance in applications that require a high-throughput stream of randomness and are likely to suffer from occasional disconnections, *e.g.*, online gaming. Therefore, at epoch r, it is possible to generate all the previous beacon outputs $\{x_1, \ldots, x_{r-1}\}$ using x_r and public key. This also implies *self-verification*, enabling verification of each output against the previous ones with no need for NIZK proof. This property is the direct result of using a (threshold) RSA signature where the secret key is the multiplicative inverse of the public key, allowing to trace the chain of randomness back by iterative encryptions. Despite the possibility of generating the history, it gets more computationally expensive (*i.e.*, grows exponentially) as the number of epochs to (re)generate increases.[16] With threshold BLS, it is possible to create a unique chain of randomness, and with threshold RSA, it is possible to have continual back and forth on such a chain.

Insight 14. *In a (plain) secret sharing, the shares do not carry any self-verifying information. But, in STROBE [22] this is the case thanks to the underlying chain of RSA decryption.*

Insight 15. *As a trade-off, the history generation property diminishes randomness quality from being truly random or pseudorandom to just being unpredictable as it is not possible to get the full history from a fresh (truly/pseudo)random value.*

Cryptographic Assumption. Cryptographic assumptions are crucial in analyzing the security of the DRB protocols. One well-known assumption used in a range of randomness beacon protocols is random oracle [24]. Particularly, the beacons with pseudorandom outputs [4], [49], [65], [77], [91], [133] make use of such assumption to turn the unpredictable beacon values to the pseudorandom ones that are indistinguishable from a uniform distribution. This assumption is also used in PVSS schemes [47], [136], allowing more efficient constructions. Another popular assumption in the context of randomness beacons is decisional Diffie-Hellman (DDH) [37] commonly served as the standard assumption for underlying PVSS and DKG schemes [47]–[49], [53]. Some protocols may also rely on stronger assumptions that are less standard and their security properties are not studied substantially. For instance, SCRAPE in the plain model uses pairings, relying on the hardness assumption of decisional bilinear squar-

16. It is possible to deploy some techniques to decrease the proof size and the running time. For more details see [22].

ing (DBS) [92].[17] SPURT [61] makes modifications to the pairing-based PVSS of SCRAPE to turn its security assumption to the more standard decisional bilinear Diffie-Hellman (DBDH) assumption [37]. Also, repeated squaring (RSW) [131] is the common hardness assumption for DRB protocols using VDF.

Gap 10. *Although the leaderless DRB of [47] works with a version of PVSS in the random oracle model, all the existing leader-based protocols using PVSS need to rely on pairings. Designing a leader-based protocol without pairings may considerably boost performance.*

Post-quantum Security. The emergence of quantum computers could be a real threat to the existing protocols due to their cryptographic hardness assumptions. In particular, bilinear-map-based PVSS is susceptible to quantum attacks [138], making the corresponding DRB protocols insecure. Gentry, Halevi, and Lyubashevsky [83] recently proposed a proactive and non-interactive PVSS scheme in which the underlying encryption scheme is based on the learning with errors (LWE) problem [124]. Although the adoption of LWE encryption in their lattice-based scheme is primarily motivated by scaling PVSS in large-scale systems (*e.g.*, where committees may scale to hundreds or thousands of parties), the secrecy of their PVSS is preserved even against quantum attackers. HERB [53] builds a randomness beacon protocol with additive homomorphic threshold encryption where a group of parties encrypt their local randomness and any threshold of parties can retrieve the aggregated beacon value. A potential advantage of this construction is its ability to resist quantum attacks by replacing fully homomorphic lattice-based schemes, supporting DKG and threshold decryption [25]. HashRand [20] is a recent effort to design a post-quantum randomness beacon using hash-based VSS [68], where a Merkle tree of shares is used as a commitment to a (sharing) polynomial instead of conventional discrete log cryptography. To prove security against a polynomial-time quantum adversary the hash function is assumed to be a quantum random oracle [33].

10. Conclusion

In this paper we provide a comprehensive overview on the topic of public randomness and discuss its manifold aspects critically, presenting concrete insights and potential gaps using a modular categorization. We conclude with a summary of our analysis on each line of works considered in the paper.

Protocols Using PVSS. These protocols provide the highest quality of uniform randomness due to their commit-reveal structure and the fact that there exists at least one honest contribution in the outcome. However, they do not scale well due to the complexity of performing PVSS and the size of the transcript that has to be exchanged and verified. Adopting leader-based and/or aggregation techniques improve scalability but are often nevertheless insufficient for practical purposes [54]. Deploying batching techniques could be another venue to explore for scalability but requires further analysis to preserve the security of sub-batches (Gap 1). These protocols usually require some form of consensus to agree on the output, making their deployment in permissionless settings challenging. Another important direction to explore are weighted PVSS protocols in which contributions from participants might have different weights according to some predefined metric which would allow their integration with proof-of-stake blockchains [62].

Protocols Using VRF. These protocols do not provide uniform randomness, but pseudorandomness. However, this level of quality is suitable for many applications like the ones mentioned in Section 1.1. They are quite efficient and scale well when used in a leader-based fashion but may suffer from biasing. Using a setup phase (*i.e.*, DKG) eliminates this issue but impedes efficient dynamic participation. Utilizing time-based cryptography is a promising approach to address this challenge while remaining compatible with permissionless settings [3]. Further approaches to explore include using VRFs with key extraction (Gap 3) and supporting flexible thresholds (Gap 4).

Protocols Using VDF. These protocols provide pseudorandomness, similar to VRF-based ones. However, unlike VRF protocols adopting a leader-based approach with VDFs does not violate unbiasability due to the VDF properties but may affect the unpredictability (Gap 7). As discussed in [54], an important aspect to explore is the security considerations of VDFs and their correspondence to the wall clock and hardware speeds. Furthermore, due to the similarities between VRFs and VDFs, one could explore their combination in a modular fashion to produce randomness with different qualities, in a similar way to [49].

Protocols Using Public Blockchain. These protocols utilize the inherent randomness in the process of block construction that could be verified publicly. However, they cannot ensure unbiasability due to various possible attack vectors such as withholding the proposal or occurring fork. The use of shared randomness is of paramount importance for blockchain systems either for security reasons (*e.g.*, leader election) or running applications (*e.g.*, online gaming). Reliance on threshold security seems to be crucial to obtain secure and efficient randomness, either using off-chain committees (*e.g.*, randomness services) or on-chain committees (*e.g.*, validators) both of which have been explored recently in [96] and [62], respectively.

Acknowledgments. The authors would like to thank Renas Bacho, Kostas Chalkias, and Ewa Syta for insightful discussions.

17. SCRAPE PVSS scheme is proposed in two versions, one in the random oracle model under DDH assumption and one in the plain model under DBS assumption.

References

[1] Blockchain gaming optimized vrf. https://piratenation.medium.com/blockchain-gaming-optimized-vrf-5b9e67d45daf/.

[2] League of entropy. https://en.wikipedia.org/wiki/League_of_Entropy.

[3] Team harmony, technical whitepaper - version 2.0. https://harmony.one/whitepaper.pdf.

[4] Team drand, drand project website. https://drand.love, 2020.

[5] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, and Gilad Stern. Bingo: Adaptively secure packed asynchronous verifiable secret sharing and asynchronous distributed key generation. *Cryptology ePrint Archive*, 2022.

[6] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, and Gilad Stern. Bingo: Adaptivity and asynchrony in verifiable secret sharing and distributed key generation. In *Annual International Cryptology Conference*, pages 39–70. Springer, 2023.

[7] Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Maofan Yin. Sync hotstuff: Simple and practical synchronous state machine replication. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 106–118. IEEE, 2020.

[8] Ittai Abraham, Kartik Nayak, Ling Ren, and Zhuolun Xiang. Good-case latency of byzantine broadcast: A complete categorization. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 331–341, 2021.

[9] Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. Balanced byzantine reliable broadcast with near-optimal communication and improved computation. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 399–417, 2022.

[10] Jesús F Almansa, Ivan Damgård, and Jesper Buus Nielsen. Simplified threshold rsa with adaptive and proactive security. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 593–611. Springer, 2006.

[11] Orestis Alpos, Zhipeng Wang, Alireza Kavousi, Sze Yiu Chau, Duc Le, and Christian Cachin. Dske: Digital signature with key extraction. *Cryptology ePrint Archive*, 2022.

[12] Marcin Andrychowicz and Stefan Dziembowski. Distributed cryptography based on the proofs of work. *Cryptology ePrint Archive*, 2014.

[13] Sarah Azouvi and Daniele Cappelletti. Private attacks in longest chain proof-of-stake protocols with single secret leader elections. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 170–182, 2021.

[14] Sarah Azouvi, George Danezis, and Valeria Nikolaenko. Winkle: Foiling long-range attacks in proof-of-stake systems. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 189–201, 2020.

[15] Sarah Azouvi, Patrick McCorry, and Sarah Meiklejohn. Winning the caucus race: Continuous leader election via public randomness. *arXiv preprint arXiv:1801.07965*, 2018.

[16] Renas Bacho, Christoph Lenzen, Julian Loss, Simon Ochsenreither, and Dimitrios Papachristoudis. Grandline: Adaptively secure dkg and randomness beacon with (almost) quadratic communication complexity. *Cryptology ePrint Archive*, 2023.

[17] Renas Bacho and Julian Loss. On the adaptive security of the threshold bls signature scheme. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 193–207, 2022.

[18] Renas Bacho and Julian Loss. Adaptively secure (aggregatable) pvss and application to distributed randomness beacons. *Cryptology ePrint Archive*, 2023.

[19] Leemon Baird, Sanjam Garg, Abhishek Jain, Pratyay Mukherjee, Rohit Sinha, Mingyuan Wang, and Yinuo Zhang. Threshold signatures in the multiverse. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 2057–2073. IEEE Computer Society, 2022.

[20] Akhil Bandarupalli, Adithya Bhat, Saurabh Bagchi, Aniket Kate, and Michael Reiter. Hashrand: Efficient asynchronous random beacon without threshold cryptographic setup. *Cryptology ePrint Archive*, 2023.

[21] Carsten Baum, Bernardo David, Rafael Dowsley, Ravi Kishore, Jesper Buus Nielsen, and Sabine Oechsner. Craft: C omposable r andomness beacons and output-independent a bort mpc f rom t ime. In *Public-Key Cryptography–PKC 2023: 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7–10, 2023, Proceedings, Part I*, pages 439–470. Springer, 2023.

[22] Donald Beaver, Konstantinos Chalkias, Mahimna Kelkar, Lefteris Kokoris-Kogias, Kevin Lewi, Ladi de Naurois, Valeria Nikolaenko, Arnab Roy, and Alberto Sonnino. Strobe: Streaming threshold random beacons. In *5th Conference on Advances in Financial Technologies (AFT 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.

[23] Donald Beaver and Nicol So. Global, unpredictable bit generation without broadcast. In *Workshop on the Theory and Application of of Cryptographic Techniques*, pages 424–434. Springer, 1993.

[24] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.

[25] Rikke Bendlin and Ivan Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In *Theory of Cryptography: 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings 7*, pages 201–218. Springer, 2010.

[26] Fabrice Benhamouda, Craig Gentry, Sergey Gorbunov, Shai Halevi, Hugo Krawczyk, Chengyu Lin, Tal Rabin, and Leonid Reyzin. Can a public blockchain keep a secret? In *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part I 18*, pages 260–290. Springer, 2020.

[27] Iddo Bentov, Ariel Gabizon, and David Zuckerman. Bitcoin beacon. *arXiv preprint arXiv:1605.04559*, 2016.

[28] Adithya Bhat, Nibesh Shrestha, Aniket Kate, and Kartik Nayak. Optrand: Optimistically responsive reconfigurable distributed randomness. In *NDSS*, 2023.

[29] Adithya Bhat, Nibesh Shrestha, Zhongtang Luo, Aniket Kate, and Kartik Nayak. Randpiper–reconfiguration-friendly random beacons with quadratic communication. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 3502–3524, 2021.

[30] George Robert Blakley and Catherine Meadows. Security of ramp schemes. In *Workshop on the Theory and Application of Cryptographic Techniques*, pages 242–268. Springer, 1984.

[31] Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *ACM SIGACT News*, 15(1):23–27, 1983.

[32] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual international cryptology conference*, pages 757–788. Springer, 2018.

[33] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology–ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings 17*, pages 41–69. Springer, 2011.

[34] Dan Boneh, Saba Eskandarian, Lucjan Hanzlik, and Nicola Greco. Single secret leader election. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 12–24, 2020.

[35] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *Advances in Cryptology—CRYPTO 2001: 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19–23, 2001 Proceedings*, pages 213–229. Springer, 2001.

[36] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In *Advances in Cryptology—EUROCRYPT 2003: International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4–8, 2003 Proceedings 22*, pages 416–432. Springer, 2003.

[37] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. *Recuperado de https://crypto. stanford. edu/˜ dabo/cryptobook/BonehShoup_0_4. pdf*, 2017.

[38] Joseph Bonneau, Jeremy Clark, and Steven Goldfeder. On bitcoin as a public randomness source. *Cryptology ePrint Archive*, 2015.

[39] Johannes Buchmann and Hugh C. Williams. A key-exchange system based on imaginary quadratic fields. *Journal of Cryptology*, 1(2):107–118, 1988.

[40] Benedikt Bünz, Steven Goldfeder, and Joseph Bonneau. Proofs-of-delay and randomness beacons in ethereum. *IEEE Security and Privacy on the blockchain (IEEE S&B)*, 2017.

[41] Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.

[42] Jan Camenisch, Manu Drijvers, Timo Hanke, Yvonne-Anne Pignolet, Victor Shoup, and Dominic Williams. Internet computer consensus. *Cryptology ePrint Archive*, 2021.

[43] Jan Camenisch and Markus Michels. Proving in zero-knowledge that a number is the product of two safe primes. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 107–122. Springer, 1999.

[44] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of CRYPTOLOGY*, 13(1):143–202, 2000.

[45] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145. IEEE, 2001.

[46] Ran Canetti, Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Adaptive security for threshold cryptosystems. In *Annual International Cryptology Conference*, pages 98–116. Springer, 1999.

[47] Ignacio Cascudo and Bernardo David. Scrape: Scalable randomness attested by public entities. In *International Conference on Applied Cryptography and Network Security*, pages 537–556. Springer, 2017.

[48] Ignacio Cascudo and Bernardo David. Albatross: publicly attestable batched randomness based on secret sharing. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 311–341. Springer, 2020.

[49] Ignacio Cascudo, Bernardo David, Omer Shlomovits, and Denis Varlakov. Mt. random: Multi-tiered randomness beacons. In *Applied Cryptography and Network Security: 21st International Conference, ACNS 2023, Kyoto, Japan, June 19–22, 2023, Proceedings, Part II*, pages 645–674. Springer, 2023.

[50] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OsDI*, volume 99, pages 173–186, 1999.

[51] Panagiotis Chatzigiannis and Konstantinos Chalkias. Proof of assets in the diem blockchain. In *Applied Cryptography and Network Security Workshops: ACNS 2021 Satellite Workshops, AIBlock, AIHWS, AIoTS, CIMSS, Cloud S&P, SCI, SecMT, and SiMLA, Kamakura, Japan, June 21–24, 2021, Proceedings*, pages 27–41. Springer, 2021.

[52] Megan Chen, Carmit Hazay, Yuval Ishai, Yuriy Kashnikov, Daniele Micciancio, Tarik Riviere, Abhi Shelat, Muthu Venkitasubramaniam, and Ruihan Wang. Diogenes: lightweight scalable rsa modulus generation with a dishonest majority. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 590–607. IEEE, 2021.

[53] Alisa Cherniaeva, Ilia Shirobokov, and Omer Shlomovits. Homomorphic encryption random beacon. *Cryptology ePrint Archive*, 2019.

[54] Kevin Choi, Arasu Arun, Nirvan Tyagi, and Joseph Bonneau. Bicorn: An optimistically efficient distributed randomness beacon. *Cryptology ePrint Archive*, 2023.

[55] Kevin Choi, Aathira Manoj, and Joseph Bonneau. Sok: Distributed randomness beacons. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 75–92. IEEE Computer Society, 2023.

[56] Benny Chor, Oded Goldreich, Johan Hasted, Joel Freidmann, Steven Rudich, and Roman Smolensky. The bit extraction problem or t-resilient functions. In *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*, pages 396–407. IEEE, 1985.

[57] Miranda Christ, Kevin Choi, and Joseph Bonneau. Cornucopia: Distributed randomness beacons at scale. *Cryptology ePrint Archive*, 2023.

[58] Jeremy Clark and Urs Hengartner. On the use of financial data as a random beacon. In *2010 Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE 10)*, 2010.

[59] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 364–369, 1986.

[60] Ran Cohen and Yehuda Lindell. Fairness versus guaranteed output delivery in secure multiparty computation. *Journal of Cryptology*, 30(4):1157–1186, 2017.

[61] Sourav Das, Vinith Krishnan, Irene Miriam Isaac, and Ling Ren. Spurt: Scalable distributed randomness beacon with transparent setup. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 2502–2517. IEEE, 2022.

[62] Sourav Das, Benny Pinkas, Alin Tomescu, and Zhuolun Xiang. Distributed randomness using weighted vrfs. *Cryptology ePrint Archive*, 2024.

[63] Sourav Das, Zhuolun Xiang, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous high-threshold distributed key generation and distributed polynomial sampling. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5359–5376, 2023.

[64] Sourav Das, Thomas Yurek, Zhuolun Xiang, Andrew Miller, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous distributed key generation. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 2518–2534. IEEE, 2022.

[65] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 66–98. Springer, 2018.

[66] Luciano Freitas de Souza, Andrei Tonkikh, Sara Tucci-Piergiovanni, Renaud Sirdey, Oana Stan, Nicolas Quero, and Petr Kuznetsov. Randsolomon: Optimally resilient random number generator with deterministic termination. In *25th International Conference on Principles of Distributed Systems*, 2022.

[67] Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. *Journal of the ACM (JACM)*, 32(1):191–204, 1985.

[68] Shlomi Dolev and Ziyu Wang. Sodsbc/sodsbc++ & sodsmpc: Post-quantum asynchronous blockchain suite for consensus and smart contracts. In *Stabilization, Safety, and Security of Distributed Systems: 23rd International Symposium, SSS 2021, Virtual Event, November 17–20, 2021, Proceedings 23*, pages 510–515. Springer, 2021.

[69] Justin Drake. Minimal VDF randomness beacon - Sharding, 2018. https://ethresear.ch/t/minimal-vdf-randomness-beacon/3566.

[70] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, 1988.

[71] Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Continuous verifiable delay functions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 125–154. Springer, 2020.

[72] Torbjörn Fagerström. Lotteries in communities of sessile organisms. *Trends in Ecology & Evolution*, 3(11):303–306, 1988.

[73] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.

[74] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part II 38*, pages 33–62. Springer, 2018.

[75] Nicolas Gailly, Kelsey Melissaris, and Yolan Romailler. tlock: practical timelock encryption from threshold bls. *Cryptology ePrint Archive*, 2023.

[76] Sally M Gainsbury and Alex Blaszczynski. How blockchain and cryptocurrency technology could revolutionize online gambling. *Gaming Law Review*, 21(7):482–492, 2017.

[77] David Galindo, Jia Liu, Mihair Ordean, and Jin-Mann Wong. Fully distributed verifiable random functions and their application to decentralised random beacons. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 88–102. IEEE, 2021.

[78] Piotr Gałka and Artur Strzelecki. How randomness affects player ability to predict the chance to win at playerunknown's battlegrounds (pubg). *The Computer Games Journal*, 10:1–18, 2021.

[79] Juan Garay and Aggelos Kiayias. Sok: A consensus taxonomy in the blockchain era. In *Cryptographers' track at the RSA conference*, pages 284–318. Springer, 2020.

[80] Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *Advances in Cryptology—EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques Prague, Czech Republic, May 2–6, 1999 Proceedings 18*, pages 295–310. Springer, 1999.

[81] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. *Journal of Cryptology*, 20(1):51–83, 2007.

[82] Craig Gentry, Shai Halevi, Hugo Krawczyk, Bernardo Magri, Jesper Buus Nielsen, Tal Rabin, and Sophia Yakoubov. Yoso: You only speak once: Secure mpc with stateless ephemeral roles. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part II*, pages 64–93. Springer, 2021.

[83] Craig Gentry, Shai Halevi, and Vadim Lyubashevsky. Practical non-interactive publicly verifiable secret sharing with thousands of parties. In *Advances in Cryptology–EUROCRYPT 2022: 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30–June 3, 2022, Proceedings, Part I*, pages 458–487. Springer, 2022.

[84] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th symposium on operating systems principles*, pages 51–68, 2017.

[85] Oded Goldreich. *Foundations of Cryptography, Volume 2*. Cambridge university press Cambridge, 2004.

[86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.

[87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game, or a completeness theorem for protocols with honest majority. In *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 307–328. 2019.

[88] Jens Groth. Non-interactive distributed key generation and key resharing. *Cryptology ePrint Archive*, 2021.

[89] Kobi Gurkan, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Aggregatable distributed key generation. In *Advances in Cryptology–EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part I*, pages 147–176. Springer, 2021.

[90] Runchao Han, Jiangshan Yu, and Haoyu Lin. Randchain: Decentralised randomness beacon from sequential proof-of-work. *IACR Cryptol. ePrint Arch.*, 2020:1033, 2020.

[91] Timo Hanke, Mahnush Movahedi, and Dominic Williams. Dfinity technology overview series, consensus system. *arXiv preprint arXiv:1805.04548*, 2018.

[92] Somayeh Heidarvand and Jorge L Villar. Public verifiability from pairings in secret sharing schemes. In *International Workshop on Selected Areas in Cryptography*, pages 294–308. Springer, 2008.

[93] Amir Herzberg, Markus Jakobsson, Stanisław Jarecki, Hugo Krawczyk, and Moti Yung. Proactive public key and signature systems. In *Proceedings of the 4th ACM Conference on Computer and Communications Security*, pages 100–110, 1997.

[94] George Kadianakis. Whisk: A practical shuffle-based ssle protocol for ethereum. *Ethereum Research.(Jan. 13, 2022). Retrieved Sept*, 5:2022, 2022.

[95] Aniket Kate and Ian Goldberg. Distributed key generation for the internet. In *2009 29th IEEE International Conference on Distributed Computing Systems*, pages 119–128. IEEE, 2009.

[96] Aniket Kate, Easwar Vivek Mangipudi, Siva Maradana, and Pratyay Mukherjee. Flexirand: Output private (distributed) vrfs and application to blockchains. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 1776–1790, 2023.

[97] Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *International conference on the theory and application of cryptology and information security*, pages 177–194. Springer, 2010.

[98] John Kelsey, Luís TAN Brandão, Rene Peralta, and Harold Booth. A reference for randomness beacons: Format and protocol version 2. Technical report, National Institute of Standards and Technology, 2019.

[99] Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. Efficient random beacons with adaptive security for ungrindable blockchains. *Cryptology ePrint Archive*, 2021.

[100] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual international cryptology conference*, pages 357–388. Springer, 2017.

[101] Valerie King and Jared Saia. Byzantine agreement in expected polynomial time. *Journal of the ACM (JACM)*, 63(2):1–21, 2016.

[102] Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ewa Syta, and Bryan Ford. Omniledger: A secure, scale-out, decentralized ledger via sharding. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 583–598. IEEE, 2018.

[103] Eleftherios Kokoris Kogias, Dahlia Malkhi, and Alexander Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 1751–1767, 2020.

[104] LESLIE LAMPORT, ROBERT SHOSTAK, and MARSHALL PEASE. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

[105] Hsun Lee, Yuming Hsu, Jing-Jie Wang, Hao Cheng Yang, Yu-Heng Chen, Yih-Chun Hu, and Hsu-Chun Hsiao. Headstart: Efficiently verifiable and low-latency participatory randomness generation at scale. In *Network and Distributed System Security (NDSS) Symposium*, volume 2022, 2022.

[106] Arjen K Lenstra and Benjamin Wesolowski. Trustworthy public randomness with sloth, unicorn, and trx. *International Journal of Applied Cryptography*, 3(4):330–343, 2017.

[107] Christoph Lenzen and Sahar Sheikholeslami. A recursive early-stopping phase king protocol. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 60–69, 2022.

[108] Da-Yin Liao and Xuehong Wang. Design of a blockchain-based lottery system for smart cities applications. In *2017 IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*, pages 275–282. IEEE, 2017.

[109] Yehuda Lindell. Secure multiparty computation (mpc). *Cryptology ePrint Archive*, 2020.

[110] Yiping Ma, Jess Woods, Sebastian Angel, Antigoni Polychroniadou, and Tal Rabin. Flamingo: Multi-round single-server secure aggregation with applications to private federated learning. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 477–496. IEEE, 2023.

[111] Sai Krishna Deepak Maram, Fan Zhang, Lun Wang, Andrew Low, Yupeng Zhang, Ari Juels, and Dawn Song. Churp: dynamic-committee proactive secret sharing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 2369–2386, 2019.

[112] Robert J. McEliece and Dilip V. Sarwate. On sharing secrets and reed-solomon codes. *Communications of the ACM*, 24(9):583–584, 1981.

[113] Liam Medley, Angelique Faye Loe, and Elizabeth A Quaglia. Sok: Delay-based cryptography. *Cryptology ePrint Archive*, 2023.

[114] Liam Medley and Elizabeth A Quaglia. Collaborative verifiable delay functions. In *Information Security and Cryptology: 17th International Conference, Inscrypt 2021, Virtual Event, August 12–14, 2021, Revised Selected Papers 17*, pages 507–530. Springer, 2021.

[115] M.Haahr. Random.org: True random number service, 2010.

[116] Silvio Micali, Michael Rabin, and Salil Vadhan. Verifiable random functions. In *40th annual symposium on foundations of computer science (cat. No. 99CB37039)*, pages 120–130. IEEE, 1999.

[117] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.

[118] Lan Nguyen. Accumulators from bilinear pairings and applications. In *Cryptographers' track at the RSA conference*, pages 275–292. Springer, 2005.

[119] Jesper Buus Nielsen, João Ribeiro, and Maciej Obremski. Public randomness extraction with ephemeral roles and worst-case corruptions. In *Annual International Cryptology Conference*, pages 127–147. Springer, 2022.

[120] Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. *Cryptology ePrint Archive*, 2016.

[121] Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In *Advances in Cryptology–EUROCRYPT 2018: 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29-May 3, 2018 Proceedings, Part II 37*, pages 3–33. Springer, 2018.

[122] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)*, 27(2):228–234, 1980.

[123] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology—CRYPTO'91: Proceedings*, pages 129–140. Springer, 2001.

[124] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *Advances in Cryptology–CRYPTO 2008: 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings 28*, pages 554–571. Springer, 2008.

[125] Cécile Pierrot and Benjamin Wesolowski. Malleability of the blockchain's entropy. *Cryptography and Communications*, 10(1):211–233, 2018.

[126] Krzysztof Pietrzak. Simple verifiable delay functions. In *10th innovations in theoretical computer science conference (itcs 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[127] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 73–85, 1989.

[128] Mayank Raikwar and Danilo Gligoroski. Sok: Decentralized randomness beacon protocols. In *Information Security and Privacy: 27th Australasian Conference, ACISP 2022, Wollongong, NSW, Australia, November 28–30, 2022, Proceedings*, pages 420–446. Springer, 2022.

[129] Irving S Reed and Gustave Solomon. Polynomial codes over certain finite fields. *Journal of the society for industrial and applied mathematics*, 8(2):300–304, 1960.

[130] Ling Ren. Analysis of nakamoto consensus. *Cryptology ePrint Archive*, 2019.

[131] Ronald L Rivest, Adi Shamir, and David A Wagner. Time-lock puzzles and timed-release crypto. 1996.

[132] Kunal Sahitya and Bhavesh Borisaniya. D-lotto: the lottery dapp with verifiable randomness. In *Data Science and Intelligent Applications: Proceedings of ICDSIA 2020*, pages 33–41. Springer, 2021.

[133] Philipp Schindler, Aljosha Judmayer, Markus Hittmeir, Nicholas Stifter, and Edgar Weippl. Randrunner: Distributed randomness from trapdoor vdfs with strong uniqueness. 2021.

[134] Philipp Schindler, Aljosha Judmayer, Nicholas Stifter, and Edgar Weippl. Hydrand: Practical continuous distributed randomness. In *Proceedings of the 2020 IEEE Symposium on Security and Privacy*, 2020.

[135] Fred B Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys (CSUR)*, 22(4):299–319, 1990.

[136] Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In *Annual International Cryptology Conference*, pages 148–164. Springer, 1999.

[137] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[138] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[139] Victor Shoup. Practical threshold signatures. In *Advances in Cryptology—EUROCRYPT 2000: International Conference on the Theory and Application of Cryptographic Techniques Bruges, Belgium, May 14–18, 2000 Proceedings 19*, pages 207–220. Springer, 2000.

[140] Nibesh Shrestha, Ittai Abraham, Ling Ren, and Kartik Nayak. On the optimality of optimistic responsiveness. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 839–857, 2020.

[141] Douglas R Stinson and Reto Strobl. Provably secure distributed schnorr signatures and a (t, n) threshold scheme for implicit certificates. In *Australasian Conference on Information Security and Privacy*, pages 417–434. Springer, 2001.

[142] Ewa Syta, Philipp Jovanovic, Eleftherios Kokoris-Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J. Fischer, and Bryan Ford. Scalable bias-resistant distributed randomness. In *38th IEEE Symposium on Security and Privacy*, May 2017.

[143] Ewa Syta, Iulia Tamas, Dylan Visher, David Isaac Wolinsky, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ismail Khoffi, and Bryan Ford. Keeping authorities" honest or bust" with decentralized witness cosigning. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 526–545. Ieee, 2016.

[144] Elina van Kempen, Qifei Li, Giorgia Azzurra Marson, and Claudio Soriente. Lisa: Lightweight single-server secure aggregation with a public source of randomness. *arXiv preprint arXiv:2308.02208*, 2023.

[145] Benjamin Wesolowski. Efficient verifiable delay functions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 379–407. Springer, 2019.

[146] Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: Bft consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 347–356, 2019.

[147] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. Rapidchain: Scaling blockchain via full sharding. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 931–948, 2018.

A. Further Discussion

Setup Assumption. DRB protocols typically need to use some form of setup assumption for initializing the process. This can be either for efficiency purposes (*e.g.*, bootstrapping) or to enable the implementation of elaborate operations. One well-known example of such a setup is PKI which is a functionality F whose responsibility is to relay parties' public keys to others. Depending on whether functionality F outputs public or private values to parties, we can refer to it as a *public* or *private* setup. In settings

with a public setup, all parties receive the same value from the functionality. This can be some group specification (*e.g.*, group generator) or an initial seed, known as a common reference string (CRS). DRB protocols using PVSS and VDF can be implemented with such a public setup. We remark that deploying VDF also needs a setup phase for its underlying group of unknown order. However, a sidestep would be either using class groups of imaginary quadratic fields [39] or trapdoor VDF augmented with a NIZK [133]. On the other hand, protocols with a private setup need to either rely on a trusted party or run an MPC protocol (*e.g.*, DKG) to generate secret values that must be kept hidden during the protocol execution. This process is an efficiency bottleneck and makes the *re-configuration* problematic [29], *i.e.*, parties cannot be replaced easily once the setup gets executed.

Simulation-based Security. It is not difficult to see that a DRB protocol is actually a secure multi-party computation (MPC) [85], [109]. Treating protocols based on some specific properties, like the majority of existing works, poses the threat of not covering all the required ones. This consequently demands paying more attention to the grounded paradigm of *real/ideal simulation* which is a well-known security formulation in the context of secure computation. By defining an ideal functionality that acts as a trusted entity faithfully carrying out the computations, a secure system is defined by comparing the real-world execution of the protocol and the execution within the presence of the ideal functionality. This approach apart from having the advantage of capturing the security concerns, allows moving towards *composable security* [44], [45] which is an important but overlooked necessity for protocols producing public randomness. Such protocols, even when designed in a stand-alone manner, often are deployed within a larger system and may have interactions with other sub-protocols to provide required fresh randomness. To our knowledge, only recently a few rare attempts have been made in this direction [21], [48].

Insight 16. *Looking at DRB as a type of secure multiparty computation together with its resemblance to SMR, allows arguing about its properties formally by adopting a security definition from the former and a liveness definition from the latter.*