

HYDRAND

Practical Continuous Distributed Randomness

Philipp Schindler, Aljoshia Judmayer, Nicholas Stifter and Edgar Weippl

SBA Research, Vienna, Austria

{pschindler, ajudmayer, nstifter, eweippl}@sba-research.org

Abstract—A reliable source of randomness is not only an essential building block in various cryptographic, security, and distributed systems protocols, but also plays an integral part in the design of many new blockchain proposals. Consequently, the topic of publicly-verifiable, bias-resistant and unpredictable randomness has recently enjoyed increased attention in a variety of scientific contributions, as well as projects from the industry. In particular *random beacon protocols*, which are aimed at *continuous operation*, can be a vital component for many current Proof-of-Stake based distributed ledger proposals. We improve upon existing random beacon approaches by introducing *HydRand*, a novel distributed protocol based on publicly-verifiable secret sharing (PVSS) to ensure unpredictability, bias-resistance, and public-verifiability of a continuous sequence of random beacon values. Furthermore, *HydRand* is able to provide guaranteed output delivery of randomness at regular and predictable intervals in the presence of adversarial behavior. In comparison to existing PVSS based approaches, our solution improves scalability by lowering the communication complexity from $\mathcal{O}(n^3)$ to $\mathcal{O}(n^2)$. Furthermore, we are the first to present a comparison of recently described schemes in the area of random beacon protocols.

I. INTRODUCTION

The question of how to generate trustworthy random values among a set of mutually distrusting participants over a message passing network was first addressed by Blum in 1983, thereby introducing the notion of *coin tossing protocols* [5]. Lately, coin tossing protocols have received increased attention, in part because randomness is proving to be a vital component of most scalable distributed ledger design approaches (e.g. [4], [12], [15]) that do not require a computationally intensive *Proof-of-Work* (PoW) mechanism as found in Bitcoin [18] and similar cryptocurrencies. Specifically, *Proof-of-Stake* (PoS) blockchain proposals, which rely on *virtual resources* in the form of digital assets, call for manipulation resistant and unpredictable leader election as part of a secure protocol design. In this regard Kiayias et al. identified leader election as a fundamental problem of PoS based protocols, since any introduced entropy is subject to potential manipulation by an adversary [15]. The distributed generation of trustworthy random values can hence be considered a complementary problem to the development of such protocols.

Random beacon protocols aim to generate publicly-verifiable, bias-resistant and unpredictable randomness¹ in distributed environments. The concept of a random *beacon* was first formalized by Rabin, which proposed a service that emits a fresh random number at regular intervals [19]. In addition

to the scenario of leader election and establishing consensus in Proof-of-Stake (PoS) based distributed ledgers, *random beacons* are also useful in a variety of other scenarios: This includes gambling and lottery services, publicly-auditable selections such as soccer World Cup draws and the verifiable assignment of a limited number of resources. Syta et al. [21] lists additional use cases for randomness including Tor hidden services, generation of elliptic curve parameters, byzantine consensus and electronic voting. One prominent example from the domain of cryptocurrencies is the provision of randomness to *Smart Contracts*, which often rely on insecure sources (such as the hash of block headers which is subject to manipulation by miners) or trusted third parties (e.g. the NIST random beacon service) [2], [9].

For all mentioned scenarios the following properties, as outlined in [8], [21], are desiderata of a random beacon protocol:

- 1) **Availability/Liveness:** Any single participant or a colluding adversary should not be able to prevent progress.
- 2) **Unpredictability:** Correct and adversarial nodes should not be able to predict (precompute) future random beacon values.
- 3) **Bias-Resistance:** Any single participant or colluding adversary should not be able to influence future random beacon values to their advantage.
- 4) **Public-Verifiability:** Third parties, i.e. processes which are not directly partaking in the protocol, should also be able to verify generated values. As soon as a new random beacon value becomes available, all parties can verify the correctness of the new value using public information only.

Although not always explicitly stated, practical solutions should also achieve good **efficiency** in terms of computational resources as well as communication complexity. Furthermore, we suggest that **guaranteed output delivery**, i.e., the inability for an adversary to prevent correct nodes of the protocol from obtaining an output [11], can also be considered a valuable property in practical random beacon protocols.

Current random beacon protocols aim to provide solutions by employing different techniques, reaching from Proof-of-Delay [9] over publicly-verifiable secret sharing (PVSS) [15], [11], [21] and unique signatures [12], [13] to utilizing Bitcoin itself as a source of randomness [3], [8]. The diversity of these approaches, as well as the differences in their underlying assumptions and characteristics, make them difficult to compare and not equally suited for all use-cases. Moreover, some recently described protocols in this field tend to be closely

¹In the following we will simply refer to this by the term *randomness*.

coupled with their respective (PoS) blockchain schemes and are therefore not easily comparable or deployable in other settings, e.g. as a stand alone protocol.

A. Contribution

We present *HydRand*, a new PVSS based distributed random beacon protocol geared towards the continuous provision of randomness at regular intervals in a Byzantine failure setting. *HydRand* guarantees the generation of new, *bias-resistant* randomness in every round of the protocol, and ensures *unpredictability* with absolute certainty for every random output value that is produced after $f + 1$ rounds in the future. The protocol assumes a synchronous system model and $n = 3f + 1$ participants. In respect to previous approaches the communication complexity is hereby lowered from $\mathcal{O}(n^3)$ to $\mathcal{O}(n^2)$. Moreover, to the best of our knowledge, we are the first to provide a detailed comparative overview of recent random beacon protocols in this field.

B. Related Work

In recent years a substantial amount of research related to random beacon protocols for distributed ledgers has been published in academia as well as the industry:

Algorand [12] proposed by J. Chen and S. Micali builds a distributed ledger by combining (i) a randomness beacon based on unique signatures and hash functions with (ii) a newly proposed randomized Byzantine agreement protocol [17]. The protocol can be parameterized to achieve good probabilistic liveness guarantees that are sufficient for all intents and purposes. Although, strictly speaking the produced randomness is not completely bias-resistant, which is no problem in their usage scenario. *Ouroboros* [15] is a provably secure Proof-of-Stake blockchain protocol. It relies on a combination of publicly-verifiable secret sharing (PVSS) and other cryptographic primitives to obtain verifiable randomness, which is identified as a necessary component for the construction of such protocols. The agreed randomness is then used as the basis for the respective Proof-of-Stake algorithm, which is the main focus of *Ouroboros*. The protocol family *RandShare*, *RandHound* and *RandHerd* proposed in Syta et al. [21] also employs PVSS in combination with a Byzantine fault tolerant (BFT) consensus algorithm (*RandShare*, *RandHound*) and additionally with Collective Signing (*RandHerd*). For the scenario outlined by the authors, the more scalable protocols *RandHound* and *RandHerd* however operate with a failure probability of 0.08%.

In an orthogonal work, I. Cascudo and B. David [11] present *SCRAPE*, which introduces an optimized variant of Schoenmakers' secret sharing protocol [20] that can directly be used to reduce computation complexity in *Ouroboros*, *RandShare*, *RandHound* and *RandHerd* as well as the herein presented *HydRand* protocol. *Scrape* builds upon the assumption that a shared bulletin board, i.e., a distributed ledger, is available for exchanging information between participants, thereby necessitating some form of external blockchain or other consensus protocol if it is to be used as a stand-alone random beacon implementation. *HydRand* helps to close this gap by presenting a self-contained protocol that is focused towards a permissioned system model.

The *Dfinity* [14] project of the equally named foundation is aiming to build a decentralized verifiable random function as the key ingredient for reaching consensus among network nodes. Compared to the other schemes, they utilize BLS signatures for that purpose [13]. BLS provides signature uniqueness as well as support for signature aggregation [6], [7]. *Dfinity* combines both of these key properties to obtain a random beacon protocol. However, the security assumptions required for BLS signatures are less analysed when compared to traditional elliptic curve cryptography.

C. Structure of this paper

The paper is structured as follows. Section II gives a high level overview of our protocol. The required background is outlined starting with cryptographic primitives in section III and followed by our system model in section IV. Section V provides a blackbox construction, which outlines the main protocol design using a standard Byzantine Agreement (BA) protocol. The details of our protocol are described in section VI and an example execution of the protocol is provided in section VII. Proofs showing that the protocol achieves the desired properties are given in section VIII. Section IX compares the protocol to other related schemes and sections X and XI discuss and conclude the paper.

II. PROTOCOL OVERVIEW

The aim of the *HydRand* protocol is to provide a bias-resistant, publicly-verifiable and unpredictable stand-alone random beacon which emits random values at a regular interval. We target *HydRand* at a *permissioned* setting with a fixed set of participants and assume a known upper bound Δ on both computation and message transmission times.

During the *protocol setup*, all participants have to exchange their public keys and prepare an initial commitment. The protocol operation itself is separated into *rounds*, where each round consists of three distinct *phases*. In each round, the previously generated random value is used for uniquely selecting the current *leader* of this round. Generally speaking, the selected leader has two main choices: (i) The leader *reveals* the correct secret value he has committed himself to the last time he was leader (or during protocol setup) and attaches his next commitment. (ii) The leader does not reveal his secret value and therefore cannot attach another commitment. In the later case, this previously committed secret value will be *reconstructed* by $f + 1$ other nodes, including at least one correct participant. The properties of the underlying Schoenmakers' PVSS² scheme [20] ensure that the random beacon value obtained by reconstruction is equal to the value that would have been obtained if a leader has revealed his secret. Once the leader's previous commitment is reconstructed, the current leader is excluded from being eligible as leader in further rounds since he has not provided a new valid commitment.

If the leader is correct he constructs a new *dataset*, which (simply speaking) includes: (i) the revealed secret value he previously committed himself to, (ii) a new commitment to a randomly chosen value and (iii) a reference to the dataset of the previous round. The leader signs this dataset using his private key and broadcasts this message and signature to all

²See Section III for details.

other nodes in the network. After receiving and verifying the dataset, each node can compute a new random value.

In case a leader fails or purposely does not broadcast any data, other participants can collaborate to reconstruct the missing secret value, i.e. the value the leader has previously committed himself to in (ii). This reconstructed value can be used by each node to obtain a new random beacon value and thereby advance the protocol to the next round and hence to the next leader. This process is repeated until eventually a leader is selected that creates a new dataset that accounts for all reconstructed datasets in between.

To ensure that a correct node is selected as leader after (at most) $f + 1$ rounds, all previously selected leaders of the last f rounds are not allowed to become leader in the current round. Since malicious nodes do not know how a revealed or reconstructed commitment of a correct node influences future random beacon values, they cannot precompute future random values once a correct node has been selected. Moreover, correct participants agree on a single history after a correct node is selected as leader, because correct leaders are assumed to build on top of a single dataset and never sign different datasets in the same round. The correct node hence acts as a barrier, since all future values cannot be predicted by the attacker after this point. Therefore, *unpredictability* is ensured with certainty for any round after $f + 1$ rounds in the future. By leveraging the properties of the underlying PVSS scheme *public-verifiability* is established. Details and proofs regarding these properties are presented in the section VIII.

To bias the resulting sequence of random beacon values, a malicious leader could try to construct and send different commitments and hence different datasets to participating nodes or selectively withhold information. Such a construction necessitates some form of (Byzantine) *consensus protocol* for participants to reach agreement upon either the existence of a single, valid commitment or that the leader was faulty. In this respect HydRand leverages on its intended application as a continuous random beacon by reducing the communication overhead of Byzantine agreement (BA) that would be incurred at each round through bundling messages. Specifically, HydRand implements its own variation of a Byzantine agreement protocol that defers consensus decisions for up to $f + 1$ rounds and combines information from multiple instances of consensus that are executed with every consecutive new round in the HydRand protocol. Thereby, the overall communication (bit) complexity of comparable PVSS based random beacon schemes is reduced from $O(n^3)$ to $O(n^2)$ as HydRand only requires a single PVSS share distribution and potentially a single PVSS recovery per round. Still the protocol outputs a new random beacon value once per round, because these values are not dependent on immediate agreement on the protocol state.

III. CRYPTOGRAPHIC PRIMITIVES

In this section, we briefly outline cryptographic primitives that we rely on in the design of HydRand. Whenever we refer to the term hash function we imply a cryptographically secure variant. Merkle trees, and Merkle inclusion proofs, first described by R. Merkle [16], are used to cryptographically verify that a particular share was previously included in a commitment without having to resend the entire commitment.

A. Publicly-Verifiable Secret Sharing

We use publicly-verifiable secret sharing (PVSS) as a primary building block in the HydRand protocol. More specifically, we make use of Schoenmakers' PVSS scheme [20], which allows a node (dealer) to efficiently share a secret value $s \in \mathbb{Z}_q$ among a set of n recipients, such that any subset of size $\geq t$ of these nodes is able to recover / reconstruct the value $G^s \in G_q$. We follow Schoenmakers' notation, where the prime number q denotes the order of a group G_q in which the discrete logarithm is hard, whereas G and g are two independent generators of this group. The value of the reconstruction threshold t is set in a way that does not enable a colluding adversary to successfully recover a shared secret without requiring the collaboration of at least one correct node, i.e. $t = f + 1$. A key property of a *publicly-verifiable* secret sharing protocol is that, upon receiving the secret shares, not only the recipients but any third party with access to the public keys of the participants can verify the correctness of the shares prior to reconstruction of the secret. We use the term *PVSS commitment*, denoted by $Com(s_d)$, to refer to the result of the share distribution process of Schoenmakers' PVSS. To form a PVSS commitment, a dealer d provides:

- encrypted shares for a secret s_d , i.e. one encrypted share Y_i for each node i encrypted by the receiver's public key
- commitments C_0, C_1, \dots, C_{t-1} to the coefficients of the underlying polynomial
- a non-interactive zero-knowledge (NIZK) proof ensuring the correctness of the encrypted shares

For additional details we refer to [20].

IV. SYSTEM AND THREAT MODEL

We assume a fixed set of known participants, hereby referred to as nodes, of size $n = 3f + 1$, of which at most f nodes may exhibit Byzantine failures and can deviate arbitrarily from the specified protocol. A node is considered to be *correct* if it does not exhibit any incorrect behavior during the entirety of the protocol execution, else it is considered to be *faulty*. The terms *Byzantine* or *malicious* are used synonymous to refer to faulty nodes. The set of nodes is denoted by $\mathcal{P} = \{1, 2, \dots, n\}$ and each node $i \in \mathcal{P}$ is assumed to have a private / public key pair $\langle sk_i, pk_i \rangle$ known to all other participants. We assume a synchronous system model with a fully connected network of authenticated and reliable bidirectional point-to-point messaging channels.

V. BLACKBOX CONSTRUCTION

Before introducing a detailed description of the HydRand protocol, we highlight the general idea of the protocol design by using a black-box construction. Hereby, we rely on a Byzantine agreement (BA) protocol abstraction instead of our BA approach described later. Other synchronous BA protocols that fit well into our continuous agreement setting, such as the recently proposed protocol in Abraham et al. [1], can be considered as suitable instantiations³.

³When using the BA protocol from [1], HydRand can theoretically also work in a $n = 2f + 1$ setting. However, this leads to a trade-off between fault tolerance and communication (bit) complexity, as their BA protocol requires each node to broadcast a message of size $O(n)$ during the *notify* step, while in our protocol (in a $n = 3f + 1$ setting) messages that are sent from all nodes to all nodes are always of constant size.

By leveraging on a BA abstraction, we can obtain a basic HydRand variant as follows. For the setup phase we use BA to agree on an initial set of commitments of all nodes, as well as the initial leader. Each following protocol round consists of four steps.

- 1) The (agreed) random beacon value from the previous round deterministically selects the current rounds leader.
- 2) The leader reveals his previously shared secret and runs Schoenmakers' PVSS protocol to commit to a new randomly chosen secret. The leader broadcasts a message m with the shared secret as well as the new commitment to all nodes.
- 3) All nodes run an instance of BA to agree on the contents of the broadcasted message (i.e. on $v = H(m)$) or on the fact that the leader did not broadcast a correct message (i.e. on $m = \perp$).
- 4) If the nodes decide on some $v = H(m)$ the current leader's commitment is accepted and it can be selected as leader again after f rounds. Furthermore the new random beacon value is obtained by combining the revealed share with the previous random beacon value. Otherwise, the nodes start the recovery procedure according to Schoenmakers' PVSS and still obtain the random beacon value. In this case, the current leader must not be selected as leader in any future round.

The above protocol achieves agreement in each round due to the BA protocol used in step 3. It can be observed that a round's result does not (immediately) depend on whether or not a leader as revealed his previously shares secret $v = H(m)$ or not $v = \perp$, as the secret obtained by reconstruction is equivalent. This decision on whether a not a node has published a new commitment only effects the random beacon after f additional rounds. At the point the set of potential leaders, and thus the next selected leader, depends on whether or not to include the current leader into this set. Using this fact in combination with the stronger assumption of $n \geq 3f+1$, we in the following show how to build a more efficient protocol, which does not require a full BA instance at each round, but instead delays the decision for (at most) $f+1$ rounds.

VI. DETAILED PROTOCOL

The protocol proceeds in rounds. Each round $r \geq 1$ consists of three phases: *propose*, *acknowledge* and *vote*. Further, each round has an associated (randomly selected) leader $\ell_r \in \mathcal{P}$, denoted by ℓ if r is clear from the context.

In each round, ℓ_r is selected uniformly at random from the set of all nodes, which have not been selected as leader during the last $f+1$ rounds. The detailed leader selection mechanism is described in section VI-D. At the end of each round all nodes learn a new random beacon value R_r . For simplicity, we hereby assume that the correct nodes agree on the first random beacon value R_0 used to select the leader of round 0 as well as the set of initial commitments of all nodes. R_0 becomes public knowledge only after the set of initial commitments was defined during setup⁴.

To simplify our notation we assume that a node or leader, which broadcasts a message is also recipient of that message.

⁴In practice this initial random value can be obtained via *Proof-of-Delay* [9] or a *Proof-of-Work* [3].

TABLE I. USED VARIABLES AND SYMBOLS

Symbol	Description
f	Number of Byzantine nodes
n	Number of all nodes, defined as $n = 3f + 1$
t	Reconstruction threshold for PVSS, defined as $t = f + 1$
i	Some or any node as defined by context
$H(\cdot)$	Cryptographic hash function
$\langle sk_i, pk_i \rangle$	Private / public keypair of node i
$\langle m \rangle_i$	Some message m signed using the secret key sk_i of node i
\parallel	String / list concatenation
r, k, x	Some round as defined by context
D_x	Dataset of some round x
\hat{x}	Previous round of round x , such that there exists a valid dataset for round \hat{x}
$D_{\hat{x}}$	Previous dataset referenced in dataset D_x
$H(D_x)$	Hash of the <i>header</i> (D_x)
\mathcal{P}	Set of all nodes (processes), \mathcal{P} is of size n
\mathcal{P}_x	Set of available nodes for some round x , i.e., set of all nodes excluding recovered nodes till round x
\mathcal{L}_x	Set of potential leaders for some round x , i.e., set of all nodes excluding recovered nodes till round x and excluding nodes that have been selected as leader within the last f rounds
q	Order of group G_q according to Schoenmaker's PVSS
G_q	Some group in which the discrete log problem hard
g, G	Two independent generators for the group G_q
d	Some node acting as dealer for PVSS
s_d	Underlying secret value, a dealer wants to share with PVSS, $s_d \in \mathbb{Z}_q$
$Com(s_d)$	PVSS commitment to the value s_d , includes commitments to the coefficients of the underlying polynomial, encrypted shares and a NIZK correctness proof.
G^{s_d}	Result of the reconstruction process for a commitment $Com(s_d)$
C_0	Commitment to the value of the first coefficient $C_0 = g^{s_d}$, part of $Com(s_d)$, used for verification of s_d
Y_i	Encrypted share for node i , part of $Com(s_d)$
S_i	Decrypted share for node i , result of decrypting Y_i using i 's private key
ℓ	The leader of the current round r
ℓ_x	The leader of round x
s_ℓ	The current leader's previously committed secret value.
s_ℓ^*	The current leader's new randomly selected secret value.
$Com(s_\ell)$	The current leader's previous commitment
$Com(s_\ell^*)$	The current leader's new commitment
R_x	The randomness of round x , defined as $R_x \leftarrow H(R_{x-1} \parallel G^{s_{\ell_x}})$
$CC(D_x)$	The <i>commit certificate</i> of dataset D_x that contains at least $f+1$ valid <i>confirmation</i> messages.
$RC(x)$	The <i>recovery certificate</i> of round x that contains at least $f+1$ valid <i>recover</i> messages.
M_x	Root of a Merkle tree for the shares Y_1, Y_2, \dots, Y_n for ℓ_x 's commitment $Com(s_{\ell_x})$ in round x
$M_x[Y_i]$	Merkle branch for Y_i , showing that Y_i is under the Merkle root M_x

Similarly, the dealer in the PVSS protocol provides a share for himself. We denote a cryptographic signature on a message m by $\langle m \rangle_i$, where i denotes the node signing the message with its private key sk_i . We further assume, that all correct nodes discard invalidly signed messages and process only messages for the round and phase it is currently working on. For an overview of all used variables and symbols see table II.

A. Phase: Propose

During this phase the round's leader reveals his previously committed value s_ℓ and provides a new commitment $Com(s_\ell^*)$. For this purpose, it is the leader's task to propose a new dataset D_r for the current round r . A dataset D_r consists of two parts,

TABLE II. MESSAGE SUMMARY AND MESSAGE FORMATS

Message	Description
$\langle propose, \langle header(D_r) \rangle_\ell, body(D_r) \rangle_\ell$	The message that is broadcasted by correct leaders in the <i>propose</i> phase of each round.
$\langle \langle acknowledge, r, H(D_r) \rangle_i, \langle header(D_r) \rangle_\ell \rangle_i$	The message that is broadcasted by corrects nodes that received a valid <i>propose</i> messages from the leader of the current round. Broadcasting this messages ensures that the leader cannot equivocate.
$\langle confirm, r, H(D_r) \rangle_i$	The message that is broadcasted by corrects nodes that received $2f + 1$ valid <i>acknowledge</i> messages from other nodes during this round. Any node which received $f + 1$ of these messages can construct a valid <i>confirmation certificate</i> for round r .
$\langle \langle recover, r \rangle_i, s_\ell, Com(s_\ell)[S_i], Y_i, M_k[Y_i] \rangle_i$	The message that is broadcasted by correct nodes that did not receive a valid <i>propose</i> message from the leader at the beginning of this round. Any node which received $f + 1$ of these messages can reconstruct a valid <i>recovery certificate</i> for round r .

a header and a body, where the hash of the header is simply denoted by $H(D_r)$ and serves as a way to authenticate the integrity of a particular dataset. The header $header(D_r)$ of dataset D_r contains:

- the current round index r
- the rounds random beacon value R_r .
- the revealed secret value s_ℓ
- the most recent round index \hat{r} for which ℓ_r is in possession of a valid dataset $D_{\hat{r}}$, as well as a confirmation certificate $CC(D_{\hat{r}})$ for this dataset. In the first round of the protocol, this value is set to $\hat{r} = 0$ since no such dataset can exist
- the hash $H(D_{\hat{r}})$ for the referenced dataset $D_{\hat{r}}$ if $\hat{r} > 0$
- a list of random beacon values $\{R_k, R_{k+1}, \dots\}$ for all recovered rounds between \hat{r} and r such that $\hat{r} < k < r$
- coefficient C_0 of the new commitment $Com(s_\ell^*)$, which allows for later verification of s_ℓ^*
- the Merkle tree root hash M_r over all encrypted shares in the new commitment $Com(s_\ell^*)$ (see the definition of the body below)

The body $body(D_r)$ of dataset D_r contains:

- the commitment $Com(s_\ell^*)$ to a new randomly chosen secret s_ℓ^*
- a confirmation certificate $CC(D_{\hat{r}})$, which confirms that $D_{\hat{r}}$ was previously accepted as valid dataset
- a recovery certificate $RC(k)$ for all rounds $k \in \{\hat{r}+1, \hat{r}+2, \dots, r-1\}$, which confirms that there exists a recovery for all rounds between \hat{r} and r . If $\hat{r} = r-1$ then no such intermediate round exists and this value is omitted.

A correct leader ℓ broadcasts a signed *propose* message $\langle propose, \langle header(D_r) \rangle_\ell, body(D_r) \rangle_\ell$ to all nodes.

Each node i , that receives a dataset D_r from the leader before the end of the propose phase, checks the validity of this dataset. For this purpose i verifies that D_r is constructed as defined and properly signed. This includes a check that the revealed secret s_ℓ corresponds to the commitment $Com(s_\ell)$ submitted previously by the current leader.

Additionally the validity of the confirmation and recovery certificates is checked. A *confirmation certificate* for dataset $D_{\hat{r}}$ is valid iff it consists of $f + 1$ signed messages of the form $\langle confirm, \hat{r}, H(D_{\hat{r}}) \rangle_i$ from $f + 1$ different senders.

Similarly, a *recovery certificate* for some round k is a collection of $f + 1$ signed messages of the form $\langle recover, k \rangle_i$ from $f + 1$ different senders. The leader selects \hat{r} as the highest possible round index for which he is only in possession of a

valid confirmation certificate $CC(D_{\hat{r}})$ but does not know a recovery certificate $RC(q)$.

B. Phase: Acknowledge

If a node i receives a valid dataset D_r from the round's leader ℓ_r during the propose phase, it constructs and broadcasts a signed acknowledge message $\langle \langle acknowledge, r, H(D_r) \rangle_i, \langle header(D_r) \rangle_\ell \rangle_i$, thereby also forwarding the revealed secret value s_ℓ . Further, each node i collects and validates acknowledge messages from all nodes.

C. Phase: Vote

Each node i checks the following conditions:

- During the current propose phase a valid dataset D_r was received.
- During the current acknowledge phase $\geq 2f + 1$ acknowledge messages from different senders have been received.
- All of those messages acknowledge the received dataset's hash⁵ $H(D_r)$.

If all conditions are met, node i broadcasts a signed confirmation message $\langle confirm, r, H(D_r) \rangle_i$. Otherwise node i , broadcasts a recover message $\langle \langle recover, r \rangle_i, s_\ell, Com(s_\ell)[S_i], Y_i, M_k[Y_i] \rangle_i$. Here, $Com(s_\ell)[S_i]$ denotes i 's decrypted share S_i and its share decryption proof according to Schoenmakers' PVSS, which cryptographically proves that S_i is a valid decryption of Y_i under i 's secret key. Round k denotes the round in which ℓ has provided the commitment $Com(s_\ell)$ and a Merkle tree root hash M_k . The Merkle branch $M_k[Y_i]$ proofs that the encrypted share Y_i was previously distributed as part of $Com(s_\ell)$ and therefore also of D_k . The values Y_i and $M_k[Y_i]$ are required to enable nodes which are not in possession of $Com(s_\ell)$ to verify the share decryption proof for S_i .

Correct nodes always include values for s_ℓ and $Com(s_\ell)[S_i], Y_i, M_k[Y_i]$ if they are in possession of the required data. Otherwise the unknown value(s) are omitted. Upon receiving recovery messages from other nodes, correct nodes accept messages with omitted values. This is not a problem since the protocol ensures that there are always at least $f + 1$ correct nodes that have received the dataset with a valid confirmation certificate, and hence can provide the shares necessary for reconstructing the secret of the respective dataset. For an example see section VII.

⁵Valid acknowledge messages for more than one value of $H(D_r)$ form a cryptographic proof of leader equivocation. E.g. in a Proof-of-Stake setting, the protocol might be extended such that this equivocation proof is used to seize the security deposit of the leader.

At the end of this phase each node i can obtain the round's random beacon value R_r . We distinguish two cases: (i) node i already knows the secret value s_ℓ , because it received the dataset D_r or an acknowledge message for D_r , and (ii) node i has received at least $f + 1$ valid recover messages which include at least $f + 1$ decrypted secret shares for s_ℓ . In this case the reconstruction procedure of Schoenmakers' PVSS can be executed to produce the value G^{s_ℓ} . In both cases R_r is then obtained by computing:

$$R_r \leftarrow H(R_{r-1} \parallel G^{s_\ell}) \quad (1)$$

D. Leader selection

At the beginning of each round $r \geq 1$, a node i determines the round's leader ℓ_r based on the available local information it gathered so far. For this purpose node i uses the randomness R_{r-1} of the previous round to deterministically select ℓ_r from the set \mathcal{L}_r of potential leaders.

We denote the canonical representation of \mathcal{L}_r as $\langle l_0, l_1, \dots, l_{|\mathcal{L}_r|-1} \rangle$ and obtain ℓ_r as follows:

$$\ell_r \leftarrow l_{(R_{r-1} \bmod |\mathcal{L}_r|)} \quad (2)$$

Let $D_{\hat{r}}$ denote the most recent valid dataset, for which node i is *not* in possession of a corresponding recovery certificate $RC(\hat{r})$. If no such dataset exists⁶ we set $\hat{r} = 0$. Now we introduce a method to determine *recovered nodes* $rn(\cdot)$ as a component needed for the definition of \mathcal{L}_r . Intuitively, the set defined by $rn(\cdot)$ contains all nodes, which have not provided valid datasets for some round where the node has been selected as leader. We define the set of all leaders which have been recovered in some round up to a referenced dataset as follows:

$$rn(D_x) = \begin{cases} \emptyset & \text{if } \hat{x} = 0 \\ \{\ell_k \mid RC(k) \in D_x\} \cup rn(D_{\hat{x}}) & \text{otherwise} \end{cases} \quad (3)$$

Here $D_{\hat{x}}$ denotes the previous dataset referenced by D_x . This function is used to construct the set of available nodes \mathcal{P}_r for round r recursively by excluding all nodes which have been selected as leader in a round for which a valid reconstruction certificate exists:

$$\mathcal{P}_r = \mathcal{P} \setminus rn(D_{\hat{r}}) \quad (4)$$

Based on this notion, the definition of the set of potential leaders \mathcal{L}_r for round r follows:

$$\mathcal{L}_r = \mathcal{P}_r \setminus \{\ell_{r-f}, \ell_{r-f+1}, \dots, \ell_{r-1}\} \quad (5)$$

Intuitively, the set \mathcal{L}_r only includes nodes, which have not been selected as leader for at least f rounds in the past and have not been reconstructed in any previous round, i.e., distributed valid datasets for all rounds in which they have been selected as leader.

VII. EXAMPLE PROTOCOL EXECUTION

Figure 1 shows four rounds of an example execution of the HydRand protocol in a setting of $f = 2$ byzantine nodes. We assume that the leaders in this specific execution got selected randomly as described in section VI-D. The sequence of leaders in this example execution includes a worst case scenario, where f succinct leaders come from the set of byzantine nodes (nodes n_3 and n_4), followed by a correct node and then again the first byzantine node (n_4).

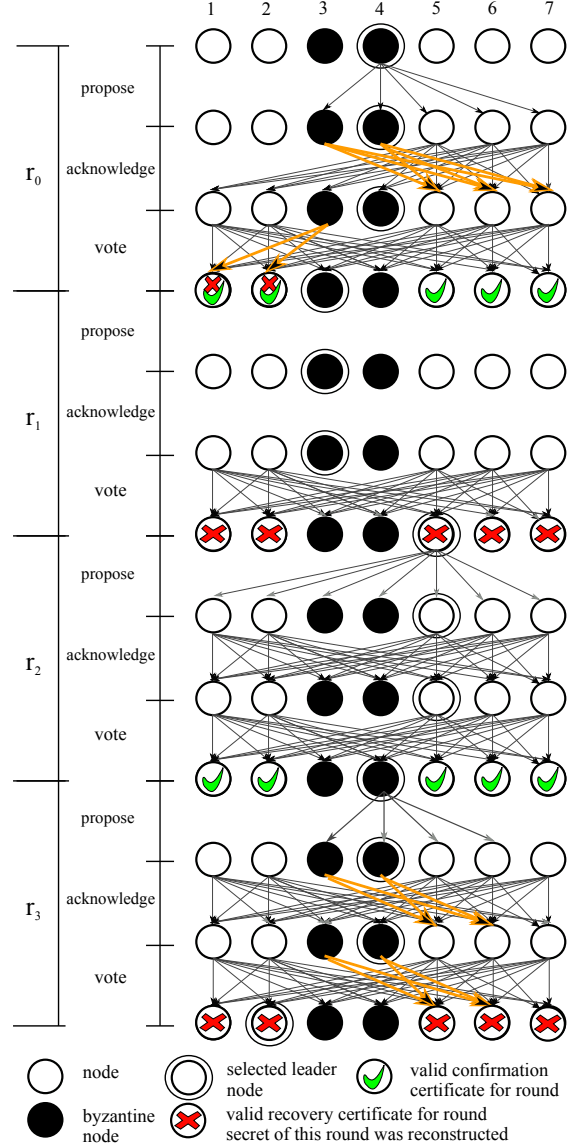


Fig. 1. Example execution of four rounds of the HydRand protocol with $n = 3f + 1 = 7$.

Round 0: In this execution the first node that is selected as the leader (i.e., node n_4) belongs to the set of byzantine nodes. This leader selectively sends a *propose* message only to a subset of correct nodes. In our case the nodes n_5 , n_6 and n_7 . Moreover, the byzantine nodes n_3 and n_4 only send *acknowledge* messages to nodes n_5 , n_6 and n_7 . After that phase, the byzantine node n_3 sends a *recover* message to the nodes n_1 and n_2 .

⁶In this scenario all rounds since protocol start can be recovered.

This leads to a situation where the correct nodes n_5 , n_6 and n_7 receive $f + 1$ acknowledge messages. Therefore, those nodes (n_5 , n_6 and n_7) broadcast *confirm* messages which together form a valid confirmation certificate of size $f + 1$ known to every node. Further, the nodes n_1 and n_2 as well as the adversary are in possession of a valid recovery certificate $RC(0)$, as nodes n_1 , n_2 and n_3 sent a *recover* messages.

Round 1: The next node (n_3) that gets selected as leader is also in the set of byzantine nodes and does not broadcast any message in this case. Therefore, the secret value of the rounds leader ℓ_1 gets reconstructed in the *vote* phase and all nodes are only in possession of a *reconstruction certificate* $RC(1)$ for this round.

Round 2: The leader of this round (n_5) belongs to the set of correct nodes and has received $f + 1$ *confirm* messages in round 0. However, node n_5 is not in possession of a valid recovery certificate since he has only received a *recover* message from node n_1 and n_2 but not from node n_3 . Therefore, it proposes a new dataset D_2 containing a valid *confirmation certificate* $CC(D_0)$ for round 0 as well as a *recovery certificate* $RC(D_1)$ for round 1 by broadcasting a corresponding message.

After receiving the *propose* message, all correct nodes, including n_1 and n_2 , are safe to assume that at least $f + 1$ correct nodes are in possession of dataset D_0 and hence accept this rounds new dataset D_2 containing $CC(D_0)$. This holds true, even for nodes n_1 and n_2 although they have not received dataset D_0 .

If node n_1 or n_2 would have been selected, then they would have constructed a dataset D_2 that contains a valid *recovery certificate* $RC(0)$ for round 0 as well. In that case the nodes n_5 , n_6 and n_7 would have thrown away their dataset $D(0)$.

Round 3: In this round node n_4 is again selected as leader. This is valid since $f + 2$ rounds have passed since this node has been selected as leader. Therefore, at least one correct node was selected as leader in between - in this case node n_5 . Since there is no *recovery certificate* $RC(2)$ for round 2 available. All further leaders have to include the *confirmation certificate* $CC(D_2)$ for round 2 to extend upon the chain of valid datasets. Otherwise their datasets would not be valid and rejected by all correct nodes. Therefore, all nodes including node n_4 , have to accept the view of node n_5 in this case.

In our example, node n_4 tries to stall the protocol by selectively releasing a new dataset D_3 only to the nodes n_3 , n_4 , n_5 and n_6 . But since those nodes are not able to reach the required number of $2f + 1$ *acknowledge* messages, no correct node will send a *confirmation* message in the last phase of this round. As a result all correct nodes will send *reconstruct* messages leading to a total of $2f + 1$ *reconstruct* messages, which is clearly enough for a *reconstruction certificate* and to reconstruct the secret $G^{s_{\ell_3}}$ of the rounds leader $\ell_3 = n_4$.

Note that, although possible, the PVSS reconstruction of the secret would not be necessary here, since in this example the leader selectively sent out a new dataset and therefore revealed the secret to at least one correct node. Per definition, correct nodes broadcast the revealed secret in their *acknowledge* messages. Therefore, all other correct nodes receive the revealed secret s_{ℓ_3} in this round even if they have not received

the dataset D_3 directly and therefore they send out *recover* messages in the *vote* phase.

VIII. PROTOCOL PROPERTIES

In this section, we show that HydRand achieves the desirable properties of a random beacon protocol as outlined in section I: *liveness*, *unpredictability*, *bias-resistance*, and *public-verifiability*. We furthermore show that our protocol also achieves *uniform agreement*.

Lemma 1. (*Possibility of construction of valid datasets*) For each round r a correct leader ℓ_r can construct a valid dataset D_r .

Proof: Since we are in a fully synchronous setting, we assume a correct leader always knows the round number r . Further, a correct leader is in possession of its own secret s_ℓ and thus knows R_r . Furthermore, the leader can always construct a new PVSS commitment for a new secret $Com(s_\ell^*)$ and is able to provide valid values for M_r and C_0 . Therefore, it only remains to show that each correct node is able to provide the required confirmation certificate $CC(\cdot)$ (and its round number) and recovery certificates $RC(\cdot)$. During the vote phase of all previous rounds, all correct nodes either broadcast a *recover* or *confirm* message. As there are at least $2f + 1$ correct nodes, each node receives at least $f + 1$ *recover* messages or at least $f + 1$ *confirm* messages (or both) for each of these rounds. As $f + 1$ *recover* messages form recovery certificate and $f + 1$ *confirm* messages form a confirmation certificate, each node is in possession of a recovery certificate or a confirmation certificate (or both) for each previous round, and is therefore able to provide the required certificates for D_r . ■

Lemma 2. (*No recovery of correct leaders*) If the leader ℓ_r is correct, there does not exist a node i , which is in possession of a valid recovery certificate $RC(r)$.

Proof: A correct leader ℓ_r sends valid proposal D_r to all nodes during the propose phase. By lemma 1, ℓ_r can always construct such a dataset. As all correct nodes consider D_r as valid, at least $2f + 1$ nodes broadcast *acknowledge* messages for D_r during the acknowledge phase. All $2f + 1$ correct nodes therefore receive $2f + 1$ valid *acknowledge* messages for D_r . As there cannot exist a valid *acknowledge* for a different dataset D'_r (because the leader only provided his signature for D_r) all correct nodes broadcast *confirm* messages during the vote phase. As correct nodes only broadcast either *confirm* or *recover* messages, there are at most f *recover* messages (from Byzantine nodes). A valid recovery certificate $RC(r)$ however requires at least $f + 1$ *recover* messages from different nodes, and therefore cannot exist. ■

Lemma 3. (*Availability of leaders*) For each round $r \geq 1$, the set of potential leaders \mathcal{L}_r contains at least $f + 1$ correct nodes.

Proof: We first show that for each round r , the set of available nodes \mathcal{P}_r contains at least $2f + 1$ correct nodes. By definitions 3 and 4, we have that only leaders ℓ_k for some round k , in which a recovery certificate $RC(k)$ exists, are excluded from the set \mathcal{P} to form \mathcal{P}_r . As we have shown in lemma 2 there are no recovery certificates for rounds with correct leaders. Therefore correct nodes cannot be excluded

from \mathcal{P} to form \mathcal{P}_r , and thus \mathcal{P}_r contains at least $2f + 1$ correct nodes.

Using the above result and definition 5, which excludes at most $f + 1$ nodes from \mathcal{P}_r to form \mathcal{L}_r , \mathcal{L}_r contains at least $f + 1$ correct nodes. ■

Lemma 4. *If a correct node knows the random beacon value R_{r-1} , it can output the random beacon value R_r by the end of round r (independent of the actions of the round's leader ℓ_r).*

Proof: Following lemma 3 we guarantee the existence of a leader ℓ_r . Since $\ell_r \in \mathcal{L}_r$ and $\mathcal{L}_r \subset \mathcal{P}_r$, we know that $\ell_r \in \mathcal{P}_r$. By applying definition 4 we get $\ell_r \notin rn(D_{\hat{r}})$. This means that there exists some history of datasets with head $D_{\hat{r}}$ in which there does not exist a recovery certificate $RC(k)$ for any round $k < \hat{r}$ in which ℓ_r was also leader. Such a history for any valid dataset D_k can only exist if at least one correct node confirmed that D_k was correctly distributed and acknowledged by $2f + 1$ nodes by providing a confirm message. Hence, at least $f + 1$ correct nodes know a common dataset D_k for all rounds k where ℓ_r was previously selected as leader. In addition all nodes know the shares for ℓ_r 's first commitment provided (and agreed upon) during the protocol setup. Thus at least $f + 1$ correct nodes can (and will) broadcast the decrypted share in case a recovery of the leader ℓ_r in round r is necessary. Hence all nodes learn the value G^{s_ℓ} corresponding to ℓ_r 's last commitment $Com(s_\ell)$, and thus obtain R_r using G^{s_ℓ} and R_{r-1} via definition 1. ■

Theorem 1. *(Liveness / Guaranteed Output Delivery) For each round r correct nodes output a new random beacon value R_r .*

Proof: We use lemmas 3 and 4 and proof the theorem by induction on the round index r . For the base case we have an agreed random beacon value R_0 as given by the protocol setup. For the induction step, we assume that R_{r-1} is known by all correct nodes. Lemma 3 ensures that the set of potential leaders \mathcal{L}_r contains at least $f + 1$ correct nodes. Therefore, definition 2 can always be applied to selected a leader ℓ_r using \mathcal{L}_r and R_{r+1} . Hence, we can use lemma 4, to show that by the end of round r each correct node outputs a value R_r . ■

Lemma 5. *(Selection of correct leaders) In each interval $\{k, k + 1, k + 2, \dots, k + f\}$ of $f + 1$ consecutive rounds there is at least one round $\hat{k} \in \{k, k + 1, k + 2, \dots, k + f\}$ such that the leader $\ell_{\hat{k}}$ of that round is correct.*

Proof: We assume that there is no correct leader in $\{\ell_k, \ell_{k+1}, \ell_{k+2}, \dots, \ell_{k+f}\}$ and derive a contradiction. We apply the definition of the set of potential leaders for round $k + f$:

$$\mathcal{L}_{k+f} = \mathcal{P}_{k+f} \setminus \{\ell_k, \ell_{k+1}, \dots, \ell_{k+f-1}\}$$

Notice that $\{\ell_k, \ell_{k+1}, \dots, \ell_{k+f-1}\}$ denotes a set of f Byzantine nodes. As there are only f Byzantine nodes in total, \mathcal{L}_{k+f} cannot contain any Byzantine nodes. However, the Byzantine node ℓ_{k+f} is leader of round $k + f$ and therefore $\ell_{k+f} \in \mathcal{L}_{k+f}$, which completes the contradiction. ■

Lemma 6. *(Agreement on potential leaders) If a node constructs a valid set of potential leaders \mathcal{L}_r in round r then every correct node constructs the same value for \mathcal{L}_r .*

Proof: Using lemma 5, for the interval $\{r - f - 1, r - f, \dots, r - 1\}$, we know that there is some round \hat{r} with a correct

leader $\ell_{\hat{r}}$ in this interval. Using lemma 1, we know that $\ell_{\hat{r}}$ is able to construct a valid dataset $D_{\hat{r}}$ in round \hat{r} . As $\ell_{\hat{r}}$ is correct, it has distributed this dataset to all nodes during the propose phase of round \hat{r} . All correct nodes therefore acknowledge $D_{\hat{r}}$ in the acknowledge phase of round \hat{r} . Since there are at least $2f + 1$ correct nodes, all correct nodes receive at least $2f + 1$ valid acknowledge messages for $D_{\hat{r}}$ by the end of the acknowledge phase. No node can receive a valid acknowledge for some different dataset $D'_{\hat{r}}$, because the correct leader $\ell_{\hat{r}}$ does not provide a signature for a different value. Therefore, all correct nodes broadcast confirm messages for $D_{\hat{r}}$. As all correct nodes broadcast either one confirm or one recovery message, there are at most f recover messages (by Byzantine nodes). Therefore, there is no valid recovery certificate $RC(\hat{r})$ for round \hat{r} . Thus, any valid future dataset needs to (indirectly) reference the common and unique dataset $D_{\hat{r}}$. Consequently, we established agreement on $D_{\hat{r}}$ and its common history provided by the references to the predecessor datasets.

As the set of available nodes $\mathcal{P}_{\hat{r}}$ for round \hat{r} is defined using only the agreed set of all nodes \mathcal{P} and $D_{\hat{r}}$, $\mathcal{P}_{\hat{r}}$ is also agreed upon. Since the definition of \mathcal{L}_r does not depend on whether or not leaders are recovered during the rounds $\{r - f, r - f + 1, \dots, r - 1\}$ and $\hat{r} \geq r - f - 1$ agreement on the set \mathcal{L}_r follows. ■

Theorem 2. *(Uniform Agreement) If a node outputs a valid random beacon value R_r in round r then every node that outputs a valid beacon value in round r outputs the same R_r .*

Proof: We proof the theorem by induction on the round index r . For the base case we have an agreed common random beacon value R_0 as given by the protocol setup.

For the induction step, we assume that every node that outputs a valid beacon value in round $r - 1$ output the same R_{r-1} . We have agreement on R_{r-1} by the induction hypothesis and shown agreement on the set of potential leaders \mathcal{L}_r in lemma 6. As the leader selection mechanism given in definition 2 only depends on those two argument, all correct nodes agree on a common unique leader ℓ_r . By applying lemma 4 we obtain that each correct nodes learns the leader's previously committed secret G^{s_ℓ} . By either checking the revealed value of s_ℓ against the leaders commitment⁷ or verifying the validity of the share decryption proof according to Schoenmakers' PVSS description [20], uniqueness of a valid G^{s_ℓ} and consequently R_r is ensured. ■

Theorem 3. *(Unpredictability) At the beginning of round r , no node can predict the outcome R_{r+f} of the random beacon protocol in round $r + f$.*

Proof: By applying lemma 5 we know that there is at least one correct leader during the interval of the $f + 1$ consecutive rounds $\{r, r + 1, r + 2, \dots, r + f\}$. Let k denote any round during this interval in which the leader ℓ_k is correct. As ℓ_k follows the protocol, it has not distributed the its secret value s_{ℓ_k} to any node at the beginning of round r . Additionally no correct nodes does provide a decrypted secret share, which could be

⁷It is sufficient to compare coefficient C_0 , a part of the PVSS commitment $Com(s_\ell)$, to the value of g^{s_ℓ} , where g denotes one of the generators used in the PVSS scheme by Schoenmakers. The value of s_ℓ is valid if and only if $C_0 = g^{s_\ell}$ holds.

used for the recovery process of the secret value. Therefore only f secret share are available to Byzantine nodes which try to recover the secret in order to compute R_k (and potentially consecutive random beacon values). However, the protocol defines the reconstruction threshold t used by the PVSS scheme to be $f + 1$. Therefore, an adversary cannot obtain the underlying secret before it is revealed or recovered during round k . Consequently, R_k and all consecutive random beacon values (including R_{r+f}) are unpredictable at the beginning of round r . ■

Theorem 4. (Bias-Resistance) *No node i can, for any round r , influence the value R_r of the random beacon protocol in a meaningful (i.e. predictable) way.*

Proof: This property follows from unpredictability and the fact that the protocol is constructed in a way that ensures that any action a (Byzantine) nodes takes in some round r , can only influence the value of the random beacon at round $r + f + 1$ or later. In theorem 3 we have shown that the random beacon value at round $r + f$ is unpredictable at the beginning of round r . Therefore, a (Byzantine) node cannot influence the random beacon values for rounds r to $r + f$, and may only influence values at round $r + f + 1$ or later in an unpredictable manner. ■

Theorem 5. (Public-Verifiability) *For each round r , an external verifier can check the correctness of the random beacon value R_r , at the end of round r .*

Proof: The external verifier v asks any correct node (i.e. at most $f + 1$ nodes) to provide its history up to and including round r . Then v can, by following the protocol rules, obtain the same random beacon value R_r if and only if the provided history is correct. Additionally, any dataset D_r and a its confirmation certificate $CC(D_r)$ allow an external verifier to obtain and check the random beacon value for round r and all rounds $k \in \{\hat{r} + 1, \hat{r} + 2, \dots, r - 1\}$. ■

IX. COMPARISON OF RANDOM BEACON PROTOCOLS

In this section, we provide an overview regarding the characteristics of various random beacon protocols. Thereby, we focus on designs that are suitable as building blocks in (PoS) blockchain protocols. For the sake of comparison, we are also include Proof of Work (PoW) and its iterated variant [10] in the given table. Simply speaking, a random beacon value via iterated PoW is computed by repeated hashing of a PoW block hash. Hereby, the number of iterations is set such that a miner cannot bias the random beacon value before the successor blocks are found with high probability.

The underlying models, assumptions, notations as well as the context may differ from protocol to protocol. Therefore, comparing existing approaches in this field is a non trivial task. We performed the hereby presented comparison to the best of our knowledge and explicitly state whenever we have not been able to pinpoint certain properties or had to estimate them. Table III presents a first step towards comparing current approaches. A property $prop$ is marked as uncertain using the notation $\sim prop$ if we have not been able to fully assess the property using the available information. For cells marked with '?' we cannot provide an adequate evaluation due to a lack of available information. The symbol \checkmark is used to

describe that a property is fulfilled, whereas \times refers to unfulfilled properties. Additionally, we use (\checkmark) to indicate that a property is achieved with probabilistic guaranties or over time. Further information on specific properties is indicated using the notation $prop_{(1-11)}$. For the complexity evaluations, n refers to the number of participants in the network, and c describes the size of the subset used in the specific protocol. Notice that c is different depending on the protocol. In the following, we provide additional details in regard to the assessment provided in table III:

- (1) In Algorand, a custom communication model is specified in great detail. Although synchrony is assumed to some extent by using time bounds, other protocols have stronger synchrony requirements.
- (2) The authors of the RandShare, RandHound and RandHerd protocols explicitly state asynchronous communication only for their RandShare protocol. However, the author's statement "The client chooses a subset of server inputs from each group, omitting servers that *did not respond on time* or with proper values [...]" [21] indicates that the communication model for RandHound is synchronous.
- (3) The exact probability is configurable as a protocol parameter. The given value represents a suggestion by the by the respective authors.
- (4) Liveness in the asynchronous communication model is only achieved after a barrier point. Whether or not this point is reached depends on the outcome of a Byzantine agreement protocol, which RandShare uses as a subprotocol [21].
- (5) Due to a lack of information, we can only estimate the communication complexity. Assuming that the only communication strictly necessary to produce the random beacon values is the broadcast of partial signatures, which each member of the correct group has to perform, the complexity $\mathcal{O}(cn)$ can be derived. This estimate excludes further potential messages exchange required for Dfinity's group setup.
- (6) Using the optimization of Schoenmakers' PVSS proposed by the authors of the Scrape protocol, the complexity can be further reduced by a factor of n .
- (7) Again using Scrape's optimization, the complexity can be reduced. Since the PVSS protocol is executed among a subset of participants, a reduction by a factor of c is possible.
- (8) The computation complexity is not dependent on the number of participants and therefore is $\mathcal{O}(1)$. However, as PoW is inherently computation intensive, the notation of $\mathcal{O}(1)$ would be misleading compared to other schemes.
- (9) HydRand reaches unpredictability with absolute certainty after $f + 1$ rounds in the future. Before that point, the protocol can only provide unpredictability with increasingly high probability.
- (10) For Algorand, bias-resistance is not achieved because the corresponding leader selection algorithm does not ensure leader uniqueness. Further, malicious leaders can selectively withhold values to bias the produced randomness.
- (11) According to our interpretation, RandHerd's communication complexity $\mathcal{O}(c^2 \log n)$ is stated per server only. Therefore, this value is not comparable to the other approaches, which consider the communication complexity of the overall system.

TABLE III. COMPARISON OF RANDOM BEACON PROTOCOLS

	Communication model	Liveness / Failure probability	Comm. complexity (overall system)	Unpredictability	Bias-Resistance	Comp. complexity (per participant)	Verification complexity (per verifier)	Characteristic cryptographic primitive(s)
Algorand	syn. ⁽¹⁾	1e-12 ⁽³⁾	?	(✓)	✗ ⁽¹⁰⁾	?	$\sim \mathcal{O}(1)$	unique signatures
Dfinity	?	1e-17	$\sim \mathcal{O}(cn)$ ⁽⁵⁾	✓	✓	?	$\sim \mathcal{O}(1)$	BLS signatures
RandShare	asyn.	✗ ⁽⁴⁾	$\mathcal{O}(n^3)$	✓	✓	$\mathcal{O}(n^3)$ ⁽⁶⁾	$\mathcal{O}(n^3)$ ⁽⁶⁾	PVSS
RandHound	\sim syn. ⁽²⁾	0.08%	$\sim \mathcal{O}(c^2n)$	✓	✓	$\sim \mathcal{O}(c^2n)$ ⁽⁷⁾	$\sim \mathcal{O}(c^2n)$ ⁽⁷⁾	PVSS & CoSi
RandHerd	\sim syn. ⁽²⁾	0.08%	? ⁽¹¹⁾	✓	✓	$\mathcal{O}(c^2 \log n)$	$\mathcal{O}(c^2 \log n)$	PVSS & CoSi
Ouroboros	syn.	✓	$\mathcal{O}(n^3)$	✓	✓	$\mathcal{O}(n^3)$ ⁽⁶⁾	$\mathcal{O}(n^3)$ ⁽⁶⁾	PVSS
Scrape	syn.	✓	$\mathcal{O}(n^3)$	✓	✓	$\mathcal{O}(n^2)$	$\mathcal{O}(n^2)$	PVSS
PoW	syn.	✓	$\mathcal{O}(n)$	(✓)	✗	very high ⁽⁸⁾	$\mathcal{O}(1)$	hash function
PoW iterated	syn.	✓	$\mathcal{O}(n)$	✓	✓	very high ⁽⁸⁾	very high ⁽⁸⁾	hash function
HydRand	syn.	✓	$\mathcal{O}(n^2)$	(✓) ⁽⁹⁾	✓	$\mathcal{O}(n^2)$ ⁽⁶⁾	$\mathcal{O}(n^2)$ ⁽⁶⁾	PVSS

X. DISCUSSION

For easier comparison and evaluation, HydRand is designed as stand-alone protocol, but with its applicability in the area of blockchains in mind. Potential areas of application are seen as part of future Proof-of-Stake approaches or permissioned blockchains. A desirable property for random beacons that our protocol can provide is *guaranteed output delivery*, i.e., a new random beacon value is guaranteed to be produced at each round regardless of the adversary’s actions. This is important for all application scenarios in which continuous operation is required and a benefit compared to other commitment schemes, such as collateral based approaches which cannot guarantee output deliver at every round.

The main benefit of HydRand, compared to other protocols in this field, is its low communication complexity of $\mathcal{O}(n^2)$ compared to $\mathcal{O}(n^3)$ of other related schemes. Based on the message format specified in table VI, we estimated the required amount of communication for the overall system for $n = 100$ and $n = 250$. For $n = 100$, a typical round without recovery results in an overall communication amount of ~ 5.4 MB, while a round with recovery leads to ~ 5.6 MB transmitted. In the scenario of $n = 250$, the respective values are ~ 34.0 MB and ~ 31.0 MB. This is an important improvement regarding the practicality of such approaches, and a further step towards their wide deployment in applications. Compared to the BA protocol from [1], our protocol makes a trade-off in favor better communication complexity while requiring a stronger fault tolerance assumption ($n \geq 3f + 1$). We achieve this by shifting the transmission of messages of size n to the leader and use cryptographically signed *conformation/recovery certificates* to converge on a history of datasets. Messages sent by all nodes are always of constant size. Moreover, the protocol does not force immediate agreement on whether or not a new commitment was submitted by a selected leader. Instead, the protocol guarantees that upon every correct node being selected as leader, all other nodes converge on one view, i.e., (at most) after $f + 1$ rounds, while in the meantime the protocols output is independent of the exact state of the nodes.

A. Future work

A promising approach for improving the scalability of HydRand to a larger number of participating nodes desirable in a permissionless setting, is to divide the overall number of nodes into different subsets/quorums. Although such an approach has been briefly indicated in [15], it was not further outlined and evaluated in the paper. For future work we consider a quorum based variant of HydRand, that only provides strong *probabilistic guarantees* but is able to scale to a very large number of nodes. The core modification is that leaders only provide shares for a subset of nodes. This subset is randomly selected for each leader based on the randomness produced by the protocol. Also during reconstruction, only members of the respective quorum (which was previously selected to receive the commits) are required to broadcast messages for reconstruction. Since the quorum sizes and reconstruction thresholds of the PVSS instances can potentially be much smaller than the total number of nodes (effective constant), the scalability and efficiency is hereby markedly improved. The design of such an approach would require an alternative agreement mechanism within the quorums and careful optimization of the protocols parameters for thresholds, quorum sizes and total number of nodes, such that liveness and unpredictability of random beacon values can be ensured with high probability.

XI. CONCLUSION

We propose HydRand, a synchronous Byzantine fault tolerant random beacon protocol that tolerates up to $f < 1/3$ Byzantine nodes and show that the protocol achieves *liveness*, *public-verifiability*, *bias resistance* and *unpredictability* for all output values after (at most) $f + 1$ rounds in the future. The presented protocol is designed for stand-alone usage, but it could also find utility in the context of current and future (PoS) and permissioned blockchain protocols. Furthermore, we provide the first step towards a systematization of novel random beacon proposals, which enables researchers to compare current as well as future designs objectively with each other. Thereby, we show that the HydRand protocol poses an improvement regarding performance and scalability in respect to comparable random beacon solutions.

REFERENCES

- [1] I. Abraham, S. Devadas, D. Dolev, K. Nayak, and L. Ren. Efficient synchronous byzantine consensus. Cryptology ePrint Archive, Report 2017/307, 2017. Accessed:2018-02-07.
- [2] N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on Ethereum smart contracts. In *International Conference on Principles of Security and Trust*, pages 164–186. Springer, 2017. Accessed: 2017-08-30.
- [3] I. Bentov, A. Gabizon, and D. Zuckerman. Bitcoin beacon. <https://arxiv.org/pdf/1605.04559v2>, 2016. Accessed: 2016-06-06.
- [4] I. Bentov, R. Pass, and E. Shi. Snow white: Provably secure proofs of stake. <https://eprint.iacr.org/2016/919.pdf>, 2016. Accessed: 2016-11-08.
- [5] M. Blum. Coin Flipping by Telephone A Protocol for Solving Impossible Problems. *ACM SIGACT News*, 15(1):23–27, 1983. Accessed: 2017-08-30.
- [6] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and Verifiably Encrypted Signatures from Bilinear Maps. In *Eurocrypt*, volume 2656, pages 416–432. Springer, 2003. Accessed: 2017-08-20.
- [7] D. Boneh, B. Lynn, and H. Shacham. Short Signatures from the Weil Pairing. *Advances in CryptologyASIACRYPT 2001*, pages 514–532, 2001. Accessed: 2017-08-20.
- [8] J. Bonneau, J. Clark, and S. Goldfeder. On Bitcoin as a public randomness source. *IACR Cryptology ePrint Archive*, 2015:1015, 2015. Accessed: 2017-08-22.
- [9] B. Bunz, S. Goldfeder, and J. Bonneau. Proofs-of-delay and randomness beacons in Ethereum. In *S&B '17: Proceedings of the 1st IEEE Security & Privacy on the Blockchain Workshop*, April 2017. Accessed: 2017-08-21.
- [10] B. Bünz, S. Goldfeder, and J. Bonneau. Proofs-of-delay and randomness beacons in ethereum. 2017.
- [11] I. Cascudo and B. David. Scrape: Scalable randomness attested by public entities. <http://eprint.iacr.org/2017/216>, 2017. Accessed: 2017-03-24.
- [12] J. Chen and S. Micali. Algorand. *arXiv preprint arXiv:1607.01341*, 2017. Accessed: 2017-08-20.
- [13] Dfinity Stiftung. Threshold Relay: How to Achieve Near-Instant Finality in Public Blockchains using a VRF, 2017. Accessed: 2017-08-20.
- [14] T. Hanke, M. Movahedi, and D. Williams. Dfinity technology overview series consensus system, Jan 2018.
- [15] A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol, 2016. Accessed: 2017-02-20.
- [16] R. C. Merkle. A digital signature based on a conventional encryption function. In C. Pomerance, editor, *Advances in Cryptology — CRYPTO '87*, pages 369–378, Berlin, Heidelberg, 1988. Springer Berlin Heidelberg.
- [17] S. Micali. Byzantine agreement, made trivial. <https://people.csail.mit.edu/silvio/SelectedApr> 2017. Accessed:2018-02-21.
- [18] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. Accessed: 2017-20-08.
- [19] M. O. Rabin. Transaction protection by beacons. *Journal of Computer and System Sciences*, 27(2):256–267, 1983.
- [20] B. Schoenmakers. A Simple Publicly Verifiable Secret Sharing Scheme and its Application to Electronic Voting. In *Annual International Cryptology Conference*, pages 148–164. Springer, 1999. Accessed: 2017-08-20.
- [21] E. Syta, P. Jovanovic, E. K. Kogias, N. Gailly, L. Gasser, I. Khoffi, M. J. Fischer, and B. Ford. Scalable Bias-Resistant Distributed Randomness. In *Security and Privacy (SP), 2017 IEEE Symposium on*, pages 444–460. IEEE, 2017. Accessed: 2017-08-20.