

VRaaS: Verifiable Randomness as a Service on Blockchains

Jacob Gorman
Supra Research
jacobgorman613@gmail.com

Lucjan Hanzlik
CISPA Helmholtz Center
for Information Security
hanzlik@cispa.de

Aniket Kate
Purdue University,
Supra Research
aniket@purdue.edu

Easwar Vivek Mangipudi
Supra Research
e.mangipudi@supraoracles.com

Pratyay Mukherjee
Supra Research
pratyay85@gmail.com

Pratik Sarkar
Supra Research
iampratiksarkar@gmail.com

Sri AravindaKrishnan
Thyagarajan
University of Sydney
t.srikrishnan@gmail.com

Abstract—Web3 applications, such as on-chain games, NFT minting, and leader elections necessitate access to unbiased, unpredictable, and publicly verifiable randomness. Despite its broad use cases and huge demand, there is a notable absence of comprehensive treatments of on-chain verifiable randomness services. To bridge this, we offer an extensive formal analysis of on-chain verifiable randomness services.

We present the *first* formalization of on-chain verifiable randomness in the blockchain setting by introducing the notion of Verifiable Randomness as a Service (VRaaS). We formally define VRaaS using an ideal functionality $\mathcal{F}_{\text{VRaaS}}$ in the Universal Composability model. Our definition not only captures the core features of randomness services, such as unbiasedness, unpredictability, and public verifiability, but also accounts for many other crucial nuances pertaining to different entities involved, such as smart contracts.

Within our framework we study a generic design of Verifiable Random Function (VRF)-based randomness service – where the randomness requester provides an input on which the randomness is evaluated as VRF output. We show that it does satisfy our formal VRaaS definition. Furthermore, we show that the generic protocol captures many real-world randomness services like Chainlink VRF and Supra dVRF.

We investigate whether our definition is minimalistic in terms of the desired security properties – towards that, we show that a couple of insecure constructions fall short of realizing our definition. Using our definition we also discover practical vulnerabilities in other designs such as Algorand beacon, Pyth VRF and Band VRF that offer on-chain verifiable randomness.

I. INTRODUCTION

Access to verifiable, secure randomness is increasingly critical in the blockchain domain, providing the foundation for liveness and safety guarantees in distributed systems [1]. This necessity extends to various Web3 applications [2, 3, 4], including online gaming, NFTs, finance, supply chain management, and many more. The rapidly expanding Web3 [5, 6] ecosystems have exploded

the need for randomness services. E.g. Chainlink VRF [7] has served 10.5 million requests in 2022 alone.

The randomness offered by these services is not only unpredictable but also bias-resistant, ensuring fairness among participants. Moreover, within the blockchain domain, achieving public verifiability of randomness is paramount. This ensures that participants, third parties, and auditors meticulously examine the legitimacy of the randomness generation process with conviction without requiring them to trust anyone, even in the future. Consequently, blockchain solutions often turn to on-chain verifiable random function services like Chainlink [7], Supra dVRF [8], or random-beacon services such as Drand [9]. These services are integrated into smart contract-enabled blockchains, allowing the contracts to access verifiable randomness for a fee.

However, we observe that current service architectures follow specific flows [7, 10, 8, 11] and interactions, which are not yet rigorously studied – this leaves significant scope for the inadvertent misuse of randomness, particularly while trying to optimize for low latency or gas cost (for example, in blockchain ecosystems like Ethereum). Such misuse may occur at the protocol-flow level without hurting the security of the underlying cryptographic primitives. This paper aims to address this significant gap between the theory and practice of on-chain verifiable randomness by formalizing the flows and analyzing the services.

To better understand the gap, let us look at existing randomness offerings. A typical randomness service (e.g. [12, 13]) constitutes a committee producing verifiable random function (VRF) outputs and uses a smart contract for bookkeeping. Upon receiving a randomness request from a requester, the smart contract computes a query input from the request and forwards the input to the VRF committee. The VRF committee generates a *publicly verifiable proof of correctness* of the generated

randomness. The smart contract verifies the proof and then returns the output to the requester via a callback^[1] function. The smart contract also records the payment and ensures the atomicity of the service offering against the payable fee. While appearing simplistic from a high-level, the details are much more intricate – as discussed next:

- Consider the vulnerabilities within systems where the service does not verify the *uniqueness* of the query input, especially when the output influences games like lotteries. This results in the same output for different randomness requests with the same query input value. An attacker, aiming to exploit the situation, might engage in front-running^[2] the game by examining multiple requests and strategically using a favorable random value.
- Services such as [11, 15] preemptively provide the output to the requester before publishing it on-chain. The requester possesses a secret value that is required for the output computation. The requester has the option to either continue the protocol execution by providing the secret value to the smart contract or abort the protocol. This lets a malicious requester selectively abort by observing the output and discontinue the protocol if the output is unfavorable.
- Some services [16, 17] crucially rely on “*properly implemented requester smart contracts*” to handle edge-cases. For example, the protocol from [17] relies on a periodic beacon to generate randomness, and the requester smart contract “*should*” handle situations when the beacon is unavailable or the beacon is discontinued. This incurs much higher responsibilities onto the smart contract designers – and each contract must also be audited carefully.
- Additionally, efforts to optimize gas costs, as discussed in [15], may lead a requester to obtain a private random seed (that is verifiable) from an output-private randomness service and then reuse the seed in a pseudorandom generator to generate multiple pseudorandom values across applications. However, given the seed, all these values are predictable and need to be verified together. Such nuances were not formally articulated in [15].

Given the wide plethora of attack vectors and venues and the intricacies of various design choices within randomness services, we delve into the formal analysis of such randomness services.

A. Our Contributions

We present our contributions in detail as follows:

^[1]A callback function [14] allows a caller to delegate the execution back to the caller. The requester receives the output randomness via the callback function once the randomness request is fulfilled.

^[2]An adversary creates a special transaction based on pending transactions. It manipulates the transactions’ ordering to execute their transaction first.

A formal framework for VRaaS. We introduce a formal framework for “*Verifiable Randomness as a Service*” (VRaaS) that rigorously captures the required properties from verifiable randomness on a blockchain. In Section IV, we define VRaaS via an ideal functionality $\mathcal{F}_{\text{VRaaS}}$ that formally captures the following properties – the output is (indistinguishable from) random, unique (even for the same requester inputs), unforgeable (i.e., unauthorized users cannot generate the output), and the randomness is generated and verified on-chain (i.e. the generated randomness is verified by a blockchain network). Our functionality is in the Universal Composability (UC) model [18], allowing seamless composition in Web3 applications.

VRaaS protocols based on VRF. We analyze the folklore VRF based generic protocol – the randomness requester provides an input to the service on which the randomness is evaluated as VRF output by a VRF committee using a keyed VRF computation. We show that it is indeed a VRaaS and it implements the $\mathcal{F}_{\text{VRaaS}}$ functionality. We present the VRF-based protocol in Section V. This protocol encompasses the most widely used randomness-service protocols, like Chainlink [7] and Supra dVRF [8], showcasing the wide applicability of our proposed ideal functionality $\mathcal{F}_{\text{VRaaS}}$.

Necessity of two transactions. We show that it is *necessary* to have two on-chain transactions to satisfy our definition. To see this, note that, under the hood, the service works as follows: a randomness requester posts a request on the blockchain via an on-chain transaction.^[3] In our framework, this request is read; a verifiable randomness output is generated in response, and finally, the output is posted on the blockchain via another on-chain transaction. The first transaction is necessary for authenticating and recording a request on-chain, and removing this transaction would allow a malicious requester to execute a front-running attack by presenting the already queried request as fresh. The second transaction is necessary to ensure that the output is available on-chain. Removing this transaction (and replacing it with a direct message to the requester) enables a malicious requester to deny the output if it is unfavorable and keep sampling until a favorable one is obtained. We capture the necessity of two transactions in Section VII by analyzing weaker functionalities (that post a single transaction instead of two) and show that they do not suffice for VRaaS.

Vulnerabilities in existing services. In our pursuit of analyzing existing randomness services, we also encounter shortcomings in certain protocols that lead to natural vulnerabilities. These include crucial dependence

^[3]These are transactions/activities that are verified by blockchain miners or validators, and once confirmed, they are permanently recorded on the blockchain.

(Algorand [17] and Band [10]) on the requester smart contract designs to guarantee output uniqueness for each request, reliance (PythVRF [11]) on honest requester behaviors to complete the protocol execution, and incorrect usage (DIA xRandom [16]) of the VRaaS output in Web3 applications. We discuss these in Section VI.

A new ledger functionality. Additionally, as part of our building blocks, we also propose a “simplified” ideal functionality \mathcal{F}_{LED} in Section III to model ledgers that support smart contracts. \mathcal{F}_{LED} extends the Gearbox ledger functionality [19] to enable running smart contracts since the original Gearbox functionality did not support smart contracts. In our VRaaS protocols, we formally capture blockchains and smart contracts using \mathcal{F}_{LED} , demonstrating its application in blockchain-based protocols. This functionality may be useful in future analysis of other blockchain-based protocols.

Responsible Disclosure. We communicated the vulnerabilities in Algorand beacon, DIA xRandom, and Band VRF to their respective teams in January 2024 and we had offered them at least four months to work on solutions/mitigations before making the vulnerabilities public. The DIA and Algorand team had responded to resolve the issues. We will keep the Band team informed about our progress with this paper.

B. Related Works

Next, we discuss the relevant literature in the randomness services.

Harmony [20] offers an on-chain verifiable randomness service protocol using distributed BLS-based VRF. It uses a verifiable delay function to prevent a malicious party in the server committee from determining whether to abort or continue after determining the output. Boba Network [21] constructs a publicly verifiable distributed randomness service from a distributed beacon implemented from a distributed VRF. Recently, Kate et al. [15] introduced output-private VRF called FlexiRand, which produces secret yet verifiable pseudorandom values based on a blinded variant of the BLS signature [22]. These protocols provide verifiable output and can be used to instantiate our VRF-based protocol.

Aptos Roll [23] is a novel on-chain randomness method introduced recently, providing instantaneous randomness within the Aptos blockchain. In this approach, Aptos validators produce a beacon value for each transaction block based on block content and a secret-shared key. When a client requests randomness in a transaction, validators/executors use the beacon value from the same block to generate randomness instantly by hashing it with client or smart contract data. This method is specific to Aptos’ blockchain design, relying on block ordering before transaction execution. It does not apply to other prominent blockchains such as Ethereum, where block

ordering and execution are not separated. In contrast, we focus on randomness services that utilize blockchain as a black box without depending on specific design features and only rely on the smart contract capabilities of the blockchain.

Another line of work considers verifiable randomness from verifiable delayed function (VDF). Chia [24] implemented a beacon using repeated-squaring VDFs in class groups. VeeDo [25] implemented a VDF-based beacon using SNARK-based VDFs. Cornucopia [26] proposed accumulator-based VDFs. CRAFT [27] proposed a UC modeling of Time-Lock Puzzles and VDFs. [28, 29] construct beacons from VDFs, proving them to be UC-secure by CRAFT. Given a VDF, they construct publicly verifiable beacons via time-lock puzzles. However, they incur significant computation overhead and evaluation delays due to the delay in the VDF computation.

II. PRELIMINARIES

We introduce the formal notations, recall the security models, and describe the necessary functionalities and primitives necessary for building our protocols here.

Notation. We use \mathbb{N} to denote the set of positive integers, \mathbb{Z} to denote the set of all integers, and $[n]$ to denote the set $\{1, 2, \dots, n\}$ (for $n \in \mathbb{N}$). We denote the security parameter by λ . We assume that every algorithm takes λ as an implicit input. We use $y := D(x)$ to denote the evaluation of a specifically deterministic algorithm D on input x to produce output y . We use $x := \text{val}$ to denote the assignment of a value val to the variable x . We use $x = y$ to check equality between x and y . We write $R(x) \rightarrow y$ or $y \leftarrow R(x)$ to denote the evaluation of a probabilistic algorithm R on input x to produce output y . We denote a randomized algorithm that runs in polynomial time as a probabilistic polynomial time (PPT) algorithm. We say a problem is computationally hard when given a problem instance, generated using the security parameter λ , for any probabilistic algorithm \mathcal{A} that runs in $O(\text{poly}(\lambda))$ time, the probability that \mathcal{A} can solve the given problem instance is upper bounded by $\text{negl}(\lambda)$, where $\text{negl}()$ is a negligible function in security parameter λ . For an algorithm A , we denote $A(x; y)$ as running the algorithm on input x and public parameter y where y might be omitted for notational simplicity. We denote an empty string as ε . A tuple of the form $(a, _)$ denotes the case where the first element is a and the second element could be any character/string/ ε .

Rand(u, w). The Rand function is defined over output distribution $\{0, 1\}^w$ and input $u \in \{0, 1\}^\lambda$ as follows, where T is initially an empty list:

$$\begin{aligned} \text{Rand}(u, w) &:= \text{If } \exists (u, r) \in T, \text{ return } r. \\ &:= \text{Else, sample } r \leftarrow \{0, 1\}^w, \\ &\quad T := T \cup (u, r), \text{ and return } r. \end{aligned}$$

Universal-Composability (UC) Model. We follow the Universal Composability Framework [18], in that a real-world multi-party protocol realizes an ideal functionality in the presence of an adversary. We assume the existence of a *default authenticated channel* in the real world between any two parties. This significantly simplifies our definitions and can be removed using an ideal authenticated channel functionality [30].

Random Oracle. A random oracle (RO) [31, 32] H is parameterized by an arbitrary domain and a specified range \mathcal{R} . An RO query on message m is denoted by $H(m)$. The plain random oracle assumption guarantees that $H(m)$ is indistinguishable from an element uniformly sampled from \mathcal{R} if m was not queried before. An observable RO additionally grants the simulator/reduction to observe (but not influence) the queries made, to H , by the adversary. A programmable RO [32] allows the simulator/reduction to program the RO to output a value y on $H(m) := y$ on a previously unqueried input message. The RO is used in the simulatable VRF below and we also discuss in Appendix A how to extend it to the global random oracle model [33].

Simulatable Verifiable Random Function. We recall the notion of simulatable VRF from the work of [34]. It consists of a tuple of four algorithms:

- $\text{Setup}(1^\lambda) \rightarrow (\text{crs}, \text{td})$: It is an optional algorithm that on input the security parameter λ outputs a setup string crs and a trapdoor td . The crs is internally used in the other algorithms.
- $\text{Gen}(1^\lambda; \text{crs}) \rightarrow (vk, sk)$: On input λ it outputs a public verification key vk and secret key sk .
- $\text{Eval}(sk, x; \text{crs}) \rightarrow (y, \pi)$: On input sk and input $x \in \{0, 1\}^\lambda$ it outputs a random string $y \in \mathcal{D}$ and a proof π .
- $\text{Verify}(vk, x, (y, \pi); \text{crs}) \rightarrow b \in \{0, 1\}$: It outputs a bit denoting whether proof verified or not.

A simulatable VRF ensures 1) *uniqueness* - for every x there exists a single (y, π) , and 2) *simulatability* - given an output $y \leftarrow \text{Rand}(x, \text{crs})$ on a previously unqueried input x the corresponding proof π can be simulated. The proof is obtained as $\pi \leftarrow \text{SimProve}(sk, \text{td}, y, x; \text{crs})$, where td is the setup trapdoor.

We need simulation-based security instead of game-based pseudorandomness since we use VRF to achieve simulation-secure randomness service where the random output will be randomly sampled and the proof will be simulated. We avoid mentioning crs explicitly in the algorithms for notational simplicity. We also demonstrate that the most widely known VRF protocols - BLS/GLOW-VRF [35], RSA-based [36], and the Elliptic-Curve based [36] VRFs are simulatable by modeling the hash function used in the protocols as a programmable random oracle. We refer to Appendix A for details.

III. SMART CONTRACT COMPATIBLE LEDGER FUNCTIONALITY

Gearbox [19] provides a simple easy-to-use timed ledger functionality for submitting messages on a global ledger; however, this functionality does not support smart contracts. On the other hand, the functionality in [37] considers smart contracts; however, it is too generalized as it captures both public and private ledgers. As we aim for a simpler variant that only captures the commonly employed public ledgers, we consider a simplification of the ideal functionality of [37] and merge that with the ledger functionality of [19]. This allows us to strengthen the simple timed ledger functionality of [19], allowing it to deploy and call smart contracts. We refer to Appendix D for a comparison between [19], [37] and our functionality.

We present our ledger functionality \mathcal{F}_{LED} in Fig. 1. \mathcal{F}_{LED} is initialized for a session sid with smart contract functions SC and an empty global ledger LOG^{sid} . The functionality allows parties to submit input transactions. Each input transaction can be of the following form:

- It can be a request to deploy a smart contract function $f()$. In this case, \mathcal{F}_{LED} parses the request as the function description and list of its initialized variables, denoted as st_f . Then, \mathcal{F}_{LED} runs a **Checker** algorithm on (f, st_f) to check the formatting. If the check passes and the smart contract function was not previously defined then $f()$ gets deployed.
- It can contain a valid call to a smart contract function $f()$ on input s . In this case, $f()$ is run on the input transaction, and the state st_f of the smart contract gets updated.

We allow the adversary to order the transactions and submit them to \mathcal{F}_{LED} to get appended on the public ledger. When a party P_i wants to read the ledger by running the **Read** command, the adversary returns a prefixed view of the ledger, denoted as $\text{LOG}_i^{\text{sid}}$ to the party. The functionality ensures that two properties - persistence and bounded timestamps, are always guaranteed. This is the same as Gearbox [19].

Security Properties. The adversary \mathcal{A}_{LED} can corrupt an arbitrary number of participating parties P_i and run on their behalf. Additionally, \mathcal{A}_{LED} has the power to order the submitted transactions for appending on the ledger. These transactions must satisfy the following two properties:

- 1) **Persistence.** Suppose parties P_i and P_j are honest, and $\text{LOG}_i^{\text{sid}}$ and $\text{LOG}_j^{\text{sid}}$ are outputs of **Read**(sid) obtained by P_i and P_j respectively at different times. Then either $\text{LOG}_i^{\text{sid}}$ is a prefix of $\text{LOG}_j^{\text{sid}}$ or vice versa. Also, $\text{LOG}_i^{\text{sid}}$ and $\text{LOG}_j^{\text{sid}}$ should be prefixes of the global ledger LOG^{sid} maintained by \mathcal{F}_{LED} . This is ensured since **Append** is an *atomic* operation inside

<p>Parties. Registered parties P_i. Adversary \mathcal{A}_{LED} corrupts the registered parties. \mathcal{F}_{LED} and parties have access to the current time via $\mathcal{F}_{\text{curr}}$.</p> <p>Parameters. \mathcal{F}_{LED} is parametrized by a hash function H, a smart contract code Checker algorithm, and a delay Δ that is unknown to all the parties but known to \mathcal{A}_{LED}.</p> <hr/> <p>Interface for party P_i:</p> <p>Session-Init(sid). If $\text{LOG}_i^{\text{sid}}$ exists then ignore this request. Else, initialize ledger $\text{LOG}_i^{\text{sid}} := \emptyset$ and $\text{HMap}[\text{sid}, 0] := 0$. Set internal transaction queue $\text{tque}_{\text{sid}} := \emptyset$. Initialize the internal memory.</p> <p>Submit(sid, tx). To submit a transaction tx: Invoke \mathcal{A}_{LED} on (tx, P_i) and add it to $\text{tque}_{\text{sid}} = \text{tque}_{\text{sid}} \cup (\text{tx}, P_i, t)$, where $t \leftarrow \mathcal{F}_{\text{curr}}$.</p> <p>Read(sid). Send (“Read”, P_i) to \mathcal{A}_{LED} to obtain $\text{LOG}_i^{\text{sid}}$. Return $\text{LOG}_i^{\text{sid}}$ to P_i.</p> <hr/> <p>Interface for the Adversary \mathcal{A}_{LED}:</p> <p>Append(sid, tno, t', tx). When \mathcal{A}_{LED} submits tx at time t':</p> <ol style="list-style-type: none"> 1) Verify that $\text{tno} = \text{tno}_{\text{last}} + 1$, where tno_{last} is the index of the latest entry on $\text{LOG}_i^{\text{sid}}$. 2) Verify that $(\text{tx}, P_i, t) \in \text{tque}_{\text{sid}}$ for some party P_i. If any of the above checks fail then return INVALID to \mathcal{A}_{LED}. Else, tx is valid, and \mathcal{F}_{LED} performs: <ol style="list-style-type: none"> a) If tx is of the form (“Deploy” SC_f): Parse $(f(), \text{st}_f) := SC_f$ where $f()$ is the function description and st_f is its initialized state. If $\text{Checker}(f(), \text{st}_f) \neq 1$ or $f()$ exists in internal memory then set $\text{txout} := \text{fail}$, otherwise set $\text{txout} := \text{Success}$ and store $(f(), \text{st}_f, P_i)$ in internal memory. b) If tx is of the form (“Call” f, s): If f is stored in internal memory, then execute $f()$ on input s to obtain output txout and updated state st_f: $(\text{txout}, \text{st}_f) := f(\text{tno}, t', s, \text{st}_f, \text{HMap}[\text{sid}, \cdot])$ <p>If function $f()$ is not deployed then set $\text{txout} := \text{fail}$.</p> 3) Set $h := H(\text{tno}, t', \text{tx}, \text{txout}, \text{HMap}[\text{sid}, \text{tno} - 1])$. Assign $\text{HMap}[\text{sid}, \text{tno}] := h$. 4) Append $(\text{tno}; t'; \text{tx}; \text{txout}; h)$ to $\text{LOG}_i^{\text{sid}}$. 5) Remove (tx, P_i, t) from tque_{sid}. <hr/> <p>At any time \mathcal{F}_{LED} enforces the following:</p> <p>Persistence. If P_i and P_j are honest, then either $\text{LOG}_i^{\text{sid}}$ is a prefix of $\text{LOG}_j^{\text{sid}}$ or $\text{LOG}_j^{\text{sid}}$ is a prefix of $\text{LOG}_i^{\text{sid}}$. And both $\text{LOG}_i^{\text{sid}}$ and $\text{LOG}_j^{\text{sid}}$ are prefixes of LOG^{sid}.</p> <p>Liveness and Bounded timestamps. For all transactions $(\text{tx}, P_i, t) \in \text{tque}_{\text{sid}}$, there is $t' \leq t + \Delta$ such that by time $t' + \Delta$, the tuple $(\text{tno}; t'; \text{tx}; \text{txout}; h) \in \text{LOG}_j^{\text{sid}}$ appears for all honest P_j.</p>

Fig. 1: Ideal Functionality \mathcal{F}_{LED} for Ledger

the ledger functionality and only a single transaction gets appended to the ledger one at a time. The property also prevents the adversary from appending garbage values to the party’s view of the ledger which is not present in the global ledger.

- 2) **Bounded timestamps** When a transaction is submitted to \mathcal{F}_{LED} at time t (obtained via the global time functionality $\mathcal{F}_{\text{curr}}$), the functionality guarantees that

it eventually gets appended to the ledger by the adversary by time $t' \leq t + \Delta$. Additionally, it appears on every honest party P_j ’s ledger $\text{LOG}_j^{\text{sid}}$ by time $t' + \Delta$ on issuing a Read command by P_j . This property ensures that an adversary can delay a transaction by at most Δ time where Δ is a bounded delay only known to the adversary. It also guarantees the liveness of the ledger as all submitted transactions will eventually get appended.

The above two security properties are inherited from the original ledger functionality of Gearbox [19]. Next, we discuss some of the design choices of our functionality and also how it captures existing smart contract services like on-chain payments. We also discuss how we model time for completeness.

Modeling Block Hash. \mathcal{F}_{LED} is parametrized by a hash function H . This is used to compute the block hash for every transaction/block being appended to the ledger. It captures block hash in real-world blockchains and we also use it in our VRaaS protocols later on. The details of the hash function can be modified based on the protocol implementing the ledger functionality. The functionality also stores a map, denoted as **HMap**, containing the block hash of each block and adds it to the block being added to the ledger.

Modeling Payments. Our functionality captures on-chain payments [38]. We assume that there is a smart-contract function $\text{Pay}()$ and its state st_{pay} maintains the account balances of each registered user. Whenever a (sender) user wants to pay a (receiver) user, the sender calls the $\text{Pay}()$ function with the amount and the receiver account details. The $\text{Pay}()$ function updates the account balance of the receiver and sender in its state and reflects the transaction. This approach is similar to the existing approaches [39, 38] for modeling payments.

Modeling Time and User Authentication. \mathcal{F}_{LED} and the participating parties crucially rely on a global clock $\mathcal{F}_{\text{curr}}$, which we realize with the TARDIS [40] model of time. In TARDIS, the ideal functionality modeling time is called a ticker. The ticker keeps track of time via time steps, allowing parties to perform any actions that they choose to perform between any two-time steps. Parties register to the ticker functionality and only the environment is allowed to progress time once it has received confirmation to proceed from all registered parties. We also need user authentication which is modeled using the ideal signature functionality of [30], similar to the Gearbox paper.

IV. VERIFIABLE RANDOMNESS SERVICE

A VRaaS framework enables a requester to generate random output by invoking VRaaS servers on requester input $x \in \{0, 1\}^\lambda \cup \varepsilon$, where ε is the empty string). The output is unique for every request (even on the same inputs). The output is accompanied by a cryptographic

proof. The proof attests to the integrity of the output corresponding to the registered server verification key and the requester’s input.

Once a requester initiates a randomness request, the request gets authenticated on the blockchain (via on-chain transactions) and the VRaaS server should only compute the VRaaS output and not be burdened with authenticating requests. This is especially important for the multi-blockchain setting, where requesters request randomness via multiple blockchains. This is performed via a request transaction through which a requester can request randomness. Once a request is made, the VRaaS reads it and fulfills it by generating the output and the proof. This output and proof are appended to the ledger so that they can be used for on-chain applications. This is modeled via a fulfillment transaction. We formalize VRaaS via our ideal functionality $\mathcal{F}_{\text{VRaaS}}$. The $\mathcal{F}_{\text{VRaaS}}$ functionality needs to read the global ledger and submit transactions to it. We model this by providing $\mathcal{F}_{\text{VRaaS}}$ access to the ledger functionality \mathcal{F}_{LED} . This approach is similar to the on-chain payment services in [39, 38] where the ideal functionality for payment has access to the ledger functionality.

$\mathcal{F}_{\text{SERV}}$ functionality. We capture the VRaaS servers via a separate ideal functionality $\mathcal{F}_{\text{SERV}}$ to allow modularity in the formalization of VRaaS. Abstracting out the servers as a separate functionality allows us to implement it using a centralized server protocol [7] or a distributed server protocol [8] without changing the protocol implementing $\mathcal{F}_{\text{VRaaS}}$. Our $\mathcal{F}_{\text{SERV}}$ functionality is adapted from the distributed server functionality of [15], except we had to customize it such that it works with the $\mathcal{F}_{\text{VRaaS}}$ functionality. It can also be used to capture the centralized server setting capturing randomness services offered by a single server, e.g. Chainlink. Our $\mathcal{F}_{\text{SERV}}$ is provided in Fig. 4. It is parametrized by a threshold t and $\mathcal{F}_{\text{SERV}}$ is assumed to be corrupt if more than t participating servers are corrupt.

We formally define the ideal functionality $\mathcal{F}_{\text{VRaaS}}$ for on-chain VRaaS in a smart contract-enabled blockchain environment in Fig. 2, 3. The blockchain is modeled as a ledger functionality \mathcal{F}_{LED} (described in Section III). We divide our functionality in two parts- setup phase and execution phase. The setup phase occurs once where the server verification key is registered, the ledger is initialized and the smart contracts for VRaaS are deployed. The execution phase models the randomness request process, fulfilling the request on-chain and finally verification of that fulfilled request.

Parties. $\mathcal{F}_{\text{VRaaS}}$ is run between the VRaaS server modeled as $\mathcal{F}_{\text{SERV}}$, multiple randomness requesters (Q_1, \dots, Q_n) and the ideal world adversary Sim. Other parties can assist in the fulfillment process of the protocol. We use the Sim algorithm to model the protocol steps

Parties. Server functionality $\mathcal{F}_{\text{SERV}}$, registered requesters (Q_1, \dots, Q_n) and other parties. The adversary, denoted by Sim, can corrupt $\mathcal{F}_{\text{SERV}}$, the requesters, and the other parties. We say $\mathcal{F}_{\text{SERV}}$ is corrupt when the adversary corrupts more than the (respective) threshold number of parties participating in $\mathcal{F}_{\text{SERV}}$. The above parties and $\mathcal{F}_{\text{VRaaS}}$ participate in \mathcal{F}_{LED} .

Notation. We denote each tuple $(\text{tno}; t'; \text{tx}; \text{txout}; h)$ by $(\text{tx}; \text{txout})$ for notational simplicity. We denote a variable as $_$ when its value is irrelevant.

Initialize VRaaS verification key $\text{VK} = \emptyset$, unique request counter $\text{reqCtr} := 0$, and internal memory of $\mathcal{F}_{\text{VRaaS}}$ as $\text{mem}_F := \emptyset$.

Key-Registration.

1. Upon receiving (“Key-Register”) from Sim forward message to $\mathcal{F}_{\text{SERV}}$. Upon receiving $(vk, \text{Verify}(\cdot))$ from $\mathcal{F}_{\text{SERV}}$: If $\text{VK} \neq \emptyset$, then ignore request.
2. Otherwise, set $\text{VK} = vk$ and store $\text{Verify}(\cdot)$ in mem_F .

Ledger-Initialization.

3. Upon receiving (“Init-Ledger”, sid) from Sim: Initialize the ledger for session sid by running $\mathcal{F}_{\text{LED}}.\text{Session-Init}(\text{sid})$. $\mathcal{F}_{\text{VRaaS}}$ will use ledger LOG^{sid} , through its local view LOG^{sid} .

Smart Contract Deployment.

4. Upon receiving (“Deploy”, $SC_{\text{Req-Rand}}$, $SC_{\text{Req-Fulf}}$, sid) from Sim: Deploy the smart contracts Req-Rand and Req-Fulf on \mathcal{F}_{LED} by running:

$$\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, \text{“Deploy” } SC_{\text{Req-Rand}}),$$

$$\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, \text{“Deploy” } SC_{\text{Req-Fulf}}).$$

Smart Contract Functions deployed on \mathcal{F}_{LED} :

- Req-Rand(x): Compute unique identifier qID from x and state of the smart contract. Append to LOG^{sid} :

$$(\text{REQ-RAND}, x; \text{qID}).$$

- Req-Fulf(qINFO, y, π): Generate qID and qINP from qINFO, and check $\text{Verify}(\text{VK}, \text{qINP}, y, \pi) = 1$, where VK is hardcoded. If verification succeeds then append to LOG^{sid} :

$$(\text{REQ-FULF}, \text{qINFO}, y, \pi; \text{qID}, \text{qINP}, \text{VK}).$$

Fig. 2: Ideal Functionality $\mathcal{F}_{\text{VRaaS}}$ (Setup Phase).

used to implement $\mathcal{F}_{\text{VRaaS}}$ (discussed later). $\mathcal{F}_{\text{VRaaS}}$ and Sim participates in \mathcal{F}_{LED} as parties so that they can access the ledger. $\mathcal{F}_{\text{VRaaS}}$ initializes the VRaaS public key as $\text{VK} := \emptyset$, an internal request counter (for keeping track of randomness requests) $\text{reqCtr} := 0$, and internal memory as $\text{mem}_F := \emptyset$.

VRaaS-Key-Registration. To register a server P , $\mathcal{F}_{\text{VRaaS}}$ receives a command (“Key-Register”) from Sim and forwards it to $\mathcal{F}_{\text{SERV}}$. Then $\mathcal{F}_{\text{SERV}}$ generates a verification key vk and the $\text{Verify}(\cdot)$ algorithm and returns it to $\mathcal{F}_{\text{VRaaS}}$. The $\text{Verify}(\cdot)$ algorithm will be used to verify the server output on-chain by smart-contracts without interaction with $\mathcal{F}_{\text{SERV}}$. $\mathcal{F}_{\text{VRaaS}}$ sets $\text{VK} = vk$ as the registered verification key. Allowing

The following three functions can be run in parallel by multiple participating parties. For notational simplicity, we assume that LOG^{sid} is used by the functions.

Request-Randomness.

5. If Q_i is honest: Upon receiving (“Req-Rand”, x) from Q_i , where $x \in \{0, 1\}^* \cup \varepsilon$, run the request transaction on x as:

$\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, (\text{“Call Req-Rand”}, x)).$

Read qID once the request gets confirmed.

If Q_i is corrupt: When Sim sends (qID, Q) ignore it if $\langle \text{REQ-RAND}, \text{qID}, _, _ \rangle \in \text{mem}_F$. Otherwise, proceed.

6. Store request in internal memory $\text{mem}_F := \text{mem}_F \cup \langle \text{REQ-RAND}, \text{qID}, \text{reqCtr}, Q_i \rangle$.
7. Update counter $\text{reqCtr} := \text{reqCtr} + 1$.

Fulfillment.

Upon receiving (“Eval-Req”, qID, w) from any party (where w denotes the number of random bits to be generated):

8. If the request identified by qID was marked fulfilled, i.e. $\langle \text{REQ-FULF}, \text{qID}, _, _, _, Q \rangle \in \text{mem}_F$, then ignore it.
9. Invoke $\text{Sim}(\text{“Req-FulInp”}, \text{qID}, w) \rightarrow (\text{qINP}, \text{qINFO})$ to obtain query input qINP and query information qINFO .
10. Fetch entry $\langle \text{REQ-RAND}, \text{qID}, \text{rCtr}, Q_i \rangle \leftarrow \text{mem}_F$ from memory corresponding to qID .
11. If $\mathcal{F}_{\text{SERV}}$ is honest: Set $y \leftarrow \text{Rand}(\text{VK}, \text{rCtr}, w)$ and obtain proof as $\pi := \mathcal{F}_{\text{SERV}}(\text{“Sim-Proof”}, \text{VK}, \text{qINP}, w, y)$.

If $\mathcal{F}_{\text{SERV}}$ is corrupt: Obtain output and proof as $(y, \pi) := \mathcal{F}_{\text{SERV}}(\text{“Eval”}, \text{VK}, \text{qINP}, w)$. If $\mathcal{F}_{\text{SERV}}$ return \perp then skip this fulfillment process.

12. Skip this fulfillment process if $\text{Verify}(\text{VK}, \text{qINP}, y, \pi) \neq 1$.
13. If $\mathcal{F}_{\text{SERV}}$ is honest or if $\mathcal{F}_{\text{SERV}}$ is corrupt and Sim instructs $\mathcal{F}_{\text{VRaaS}}$ (upon being queried by $\mathcal{F}_{\text{VRaaS}}$ as $\text{Sim}(\text{“Run Fulfillment?”}) \rightarrow \text{Yes/No}$) to run the fulfillment transaction then perform:

$\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, (\text{“Call Req-Fulf”}, \text{qINFO}, y, \pi)).$

Once the transaction output is appended on \mathcal{F}_{LED} mark qID fulfilled by storing it in memory as $\text{mem}_F := \text{mem}_F \cup \langle \text{REQ-FULF}, \text{qID}, \text{qINP}, \text{rCtr}, y, \pi, Q \rangle$.

Local Verification.

Upon receiving (“Verify-Local”, qINP, y, π, Q), from any party M :

14. If $\langle \text{REQ-FULF}, _, \text{qINP}, _, y, \pi, Q \rangle \in \text{mem}_F$ then send VERIFIED to party M . Otherwise, send \perp to M .

Fig. 3: Ideal Functionality $\mathcal{F}_{\text{VRaaS}}$ (Execution Phase).

$\mathcal{F}_{\text{SERV}}$ to generate vk allows us to limit the problem of key-registration inside $\mathcal{F}_{\text{SERV}}$. Looking ahead, this would also help us in simulation when $\mathcal{F}_{\text{VRaaS}}$ samples a random output y for a request, and then Sim invokes the adversary Sim_P to generate the simulated proof π by using the knowledge of sk .

Ledger-Initialization. To initialize the ledger Sim triggers $\mathcal{F}_{\text{VRaaS}}$ and then $\mathcal{F}_{\text{VRaaS}}$ initiates a new ledger LOG^{sid} corresponding to session sid by calling \mathcal{F}_{LED} . If a ledger already exists for session sid then \mathcal{F}_{LED} ignores this request.

Smart Contract Deployment. The functionality ini-

Parties. Registered parties $\mathbf{P} := (P_1, \dots, P_m)$. The adversary, denoted by Sim_P , can corrupt the parties.

Parameter. $\mathcal{F}_{\text{SERV}}$ is parametrized by corruption threshold t and a $\text{Verify}(\cdot)$ function. Initialize $T[\cdot, \cdot]$ and $T_{\text{par}}[\cdot, \cdot]$.

Key-Registration.

1. Upon receiving (“Key-Register”) from $\mathcal{F}_{\text{VRaaS}}$ forward it to Sim_P .
2. On $(\mathbf{P}, vk, \{vk_1, \dots, vk_m\})$ from Sim_P : Parse $\mathbf{P} := \{P_1, \dots, P_m\}$ and when vk is unique:
 - a) Define $\mathbf{P}_{\text{CORR}} \subset \mathbf{P}$ is the set of corrupt servers and $\mathbf{P}_{\text{HON}} := \mathbf{P} \setminus \mathbf{P}_{\text{CORR}}$ is the set of honest servers.
 - b) For each $P_i \in \mathbf{P}$ set $\text{Keys}[P_i] := vk_i$.
 - c) If $|\mathbf{P}_{\text{CORR}}| \geq t + 1$, then mark server set \mathbf{P} as “Corrupt”.
 - d) Send (\mathbf{P}, vk, vk_i) to each $P_i \in \mathbf{P}_{\text{HON}}$ and register (\mathbf{P}, vk) .
3. Send $(vk, \text{Verify}(\cdot))$ to $\mathcal{F}_{\text{VRaaS}}$.

Evaluation.

When $\mathcal{F}_{\text{VRaaS}}$ sends (“Sim-Proof”, $\text{VK}, \text{qINP}, w, y$) or (“Eval”, vk, qINP, w):

4. Send (“Eval”, qINP) to all servers in \mathbf{P} .
5. On (“Partial-Eval”, qINP, vk_j) from server P_j : Forward request to Sim_P . If Sim_P sends $(vk_j, \text{qINP}, j, y_j, \pi_j)$ then forward it to P_j and set $T_{\text{par}}[\text{qINP}, vk, P_j] := (y_j, \pi_j)$. If the same query is repeated then fetch (y_j, π_j) from T_{par} and return it to P_j .
6. On (“Aggregate”, $\text{qINP}, vk, \{(y_i, \pi_i)\}_{i \in [\ell]}$) from a registered party P : If $\ell < t + 1$, then return \perp . Otherwise perform the following based on the corruption of P :
 - If server set \mathbf{P} is honest then $\mathcal{F}_{\text{SERV}}$ was invoked on (“Sim-Proof”, $\text{VK}, \text{qINP}, w, y$). Perform:
 - a) Send $(\text{qINP}, w, y, vk, \{(y_i, \pi_i)\}_{i \in [\ell]})$ to Sim_P .
 - b) If Sim_P sends \perp then send it to $\mathcal{F}_{\text{VRaaS}}$ and exit. Otherwise, receive π from Sim_P and register (y, π) as $T[\text{qINP}, vk] := (y, \pi)$.
 - c) Send π to $\mathcal{F}_{\text{VRaaS}}$.
 - If server set \mathbf{P} is corrupt then $\mathcal{F}_{\text{SERV}}$ was invoked on (“Eval”, vk, qINP, w). Perform:
 - a) Send $(\text{qINP}, w, vk, \{(y_i, \pi_i)\}_{i \in [\ell]}, y)$ to Sim_P .
 - b) If Sim_P sends \perp then send it to P and exit. Otherwise, receive (y, π) from Sim_P and register (y, π) as $T[\text{qINP}, vk] := (y, \pi)$.
 - c) Send (y, π) to $\mathcal{F}_{\text{VRaaS}}$.

Fig. 4: Helper Ideal Functionality $\mathcal{F}_{\text{SERV}}$ for the Server in $\mathcal{F}_{\text{VRaaS}}$.

tializes two smart contracts - Req-Rand for requesting randomness and Req-Fulf for verifying and storing the fulfilled requests. The smart contract functions are defined by the concrete protocol that implements $\mathcal{F}_{\text{VRaaS}}$ and the $\text{Verify}(\cdot)$ function. So our abstraction supports such customizations by allowing Sim to send the concrete smart contract function details to $\mathcal{F}_{\text{VRaaS}}$. Details of the smart contract functions are:

- Req-Rand: Given a request input x it generates a unique request identifier qID and appends them to the ledger. Under the hood, the protocol implementing

$\mathcal{F}_{\text{VRaaS}}$ should ensure that Req-Rand generates a unique qID for every request (even for different requests on the same requester input x) since qID uniquely identifies each request on the ledger.

Looking ahead, in the fulfillment process a qINFO will be generated corresponding to each (x, qID) for a request. It basically contains the qID and some auxiliary information. This qINFO will be used to generate a qINP on which the $\mathcal{F}_{\text{SERV}}$ will generate the output y and proof π . qINFO contains additional information like the requester’s callback function, number of random bits requested, block hash etc, which are not included in qID and qINP .

- Req-Fulf : Given qINFO , output y and proof π it generates qID and qINP , and verifies the generation of y on qINP by verifying the proof. Once verification succeeds the output is appended to the ledger.

Request-Randomness. Upon receiving a randomness request on input x from any requester Q_i , $\mathcal{F}_{\text{VRaaS}}$ runs the VRaaS smart contract function Req-Rand on the requester input x to obtain a smart contract generated query ID qID_i . The request x and the qID gets appended to the ledger. $\mathcal{F}_{\text{VRaaS}}$ assigns the current value of reqCtr to the request and stores $\langle \text{REQ-RAND}, \text{qID}, x, \text{reqCtr}, Q_i \rangle$ in its memory. The randomness request counter is incremented, i.e. $\text{reqCtr} := \text{reqCtr} + 1$ for each request. The reqCtr is internal to $\mathcal{F}_{\text{VRaaS}}$ and ensures that each tuple is unique due to the uniqueness of reqCtr . This ensures that each request is stored in a unique format even though the requester input (x, Q_i) can be the same. Note that $\mathcal{F}_{\text{VRaaS}}$ allows requesting randomness on the same input multiple times as it is a VRaaS . Later, the randomness will be generated corresponding to each reqCtr value, and this assigning a unique reqCtr to each request ensures that the output randomness will be independently generated for each request.

Fulfillment. Upon receiving (“ Eval-Req ”, qID, w) as a randomness request from any party M : $\mathcal{F}_{\text{VRaaS}}$ verifies that the request with ID qID has not been fulfilled (looking ahead $\mathcal{F}_{\text{VRaaS}}$ registers each successful fulfillment request as fulfilled in its memory). This avoids double-fulfilling the same request, identified by qID . Once the checks pass, $\mathcal{F}_{\text{VRaaS}}$ invokes Sim on (qID, w) to obtain query information qINFO and query input qINP . The output and the proof will be generated on qINP and the fulfillment transaction will verify it. qINP depends on the protocol and state of the LOG^{sid} and so Sim is responsible for generating it. For example, qINP could include the block hash (i.e. hash of the block containing the request transaction). qINFO contains additional information that allows one to regenerate qID and qINP from it. qINFO acts as a placeholder to capture the details of the VRaaS protocol (implementing $\mathcal{F}_{\text{VRaaS}}$) which do not appear in the input to the VRaaS servers.

Once qINFO and qINP is generated by Sim , $\mathcal{F}_{\text{VRaaS}}$ generates the output and proof (y, π) as follows based on $\mathcal{F}_{\text{SERV}}$ ’s corruption.

- If VRaaS server P is honest, then $\mathcal{F}_{\text{VRaaS}}$ samples a random w -bits string by running $\text{Rand}(\text{VK}, \text{rCtr}, w)$ to generate the output y . The uniqueness of rCtr ensures that each request (VK, rCtr) uniquely corresponds to the session with verification VK . Once y is generated, $\mathcal{F}_{\text{VRaaS}}$ invokes $\mathcal{F}_{\text{SERV}}$ with input $(\text{VK}, \text{qINP}, w, y, Q)$ to obtain the simulated proof π that attests to the computation of y on qINP . In $\mathcal{F}_{\text{SERV}}$, the simulator algorithm Sim_P is run which returns the simulated proof to $\mathcal{F}_{\text{SERV}}$, and that is forwarded to $\mathcal{F}_{\text{VRaaS}}$. $\mathcal{F}_{\text{VRaaS}}$ sets the output as (y, π) and runs the smart contract function Req-Fulf on (qINFO, y, π) to register the output as fulfilled on LOG^{sid} . Once the fulfillment transaction gets appended to the ledger $\mathcal{F}_{\text{VRaaS}}$ registers the request as fulfilled by storing the entry $\langle \text{REQ-FULF}, \text{qID}, \text{qINP}, \text{rCtr}, y, \pi, Q \rangle$ ^[4] in its memory.
- Meanwhile, if VRaaS server P is corrupt then $\mathcal{F}_{\text{VRaaS}}$ invokes $\mathcal{F}_{\text{SERV}}$ with input $(\text{VK}, \text{qINP}, w)$ to obtain the output (y, π) . In this case, Sim is responsible for running the fulfillment transaction on behalf of the corrupt $\mathcal{F}_{\text{SERV}}$.

Local-Verification. This step is run by any party M that wants to verify an output (y, π) corresponding to a query input qINP . The verification process succeeds if the request is marked as fulfilled in the memory of $\mathcal{F}_{\text{VRaaS}}$.

Given the above functionality we discuss how $\mathcal{F}_{\text{VRaaS}}$ captures the following properties that are required for a VRaaS :

- 1) **Tackling Impersonation:** Each randomness request initiated by a requester Q_i is stored in the memory mem_F along with the requester ID Q_i in Step 6. It prevents a different requester $Q' \neq Q_i$ from requesting randomness on behalf of Q_i . Similarly, upon fulfillment of a request the functionality stores the query input qINP , verifiable out (y, π) , and the requester ID Q_i in mem_F in Step 13. It binds the output to Q_i . If Q' wants to use this output then verification fails in Step.14. as the output is not registered with Q' .
- 2) **Input Uniqueness:** For each randomness request with requester input x , the smart contract adds a tuple containing qID . Then $\mathcal{F}_{\text{VRaaS}}$ adds a unique nonce rCtr (maintained by the internal counter reqCtr) for each request in Step 6. The randomness request is stored in the memory as a tuple $\langle \text{REQ-RAND}, \text{qID}, \text{reqCtr}, Q_i \rangle$.

^[4]Storing qID and rCtr allows $\mathcal{F}_{\text{VRaaS}}$ to check that the qID was fulfilled by running the consistency check in Step 8. Storing qINP is essential for the output verification in the next step.

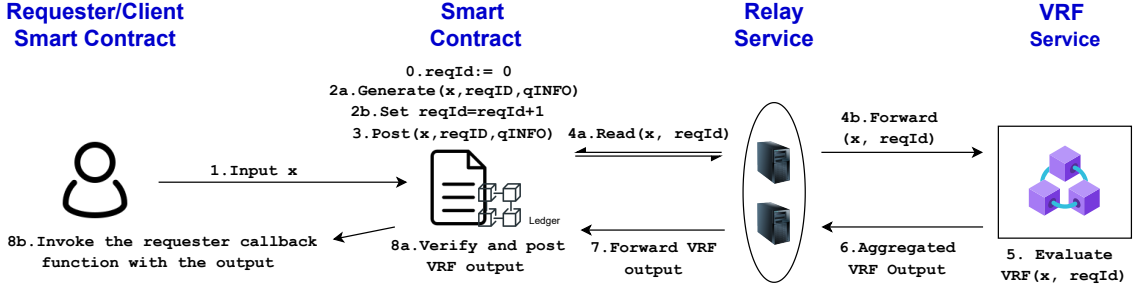


Fig. 5: Message flow in VRF based protocol π_{R-VRF} . Input x includes the callback function parameters.

- This tuple is unique due to the uniqueness of $rCtr$. Looking ahead, if a protocol implements \mathcal{F}_{VRaaS} then it should ensure that qID_i is unique for each request.
- 3) **Unbiasable Random Output:** When \mathcal{F}_{SERV} is honest and any party requests fulfillment of a query qID , \mathcal{F}_{VRaaS} generates a random w -bit output y on unique input $(VK, rCtr, w)$ by invoking $Rand$ function in Step 11. It is ensured the output is random and unique due to the uniqueness of $rCtr$ (as discussed above in “Input Uniqueness” paragraph). Looking ahead, if a protocol implements \mathcal{F}_{VRaaS} then it should ensure that $qINP$ is unique for each request.
 - 4) **Unforgeability:** Assuming \mathcal{F}_{SERV} is honest, an adversarial requester cannot forge a $VRaaS$ output (y, π) on $qINP$ without querying \mathcal{F}_{VRaaS} since the entry $\langle REQ-FULF, qID, qINP, rCtr, y, \pi, Q \rangle$ will not be registered in mem_F unless the entry was added (end of Step 13.) to mem_F by \mathcal{F}_{VRaaS} after successful execution of the Fulfillment step. This would lead to the verification process outputting \perp .
 - 5) **On-Chain Verifiability:** The request transaction and the fulfillment transaction can be verified on-chain as the output is appended on-chain. It can be further utilized for other on-chain applications.

We note that many existing randomness services [41] include the block hash h (of the block containing the request) inside $qINP$. This helps in ensuring unbiasedness of the output even when \mathcal{F}_{SERV} is corrupt. Since block hash (of the current block) is not available during $Req-Rand$, it cannot be included in qID . Instead, Sim includes the block hash inside $qINP$ and the block number (containing the $Req-Rand$ transaction) inside $qINFO$. Later, in the fulfillment phase, $Req-Fulf$ verifies the block hash of $Req-Rand$ in $qINP$ by matching it with $HMap$ in \mathcal{F}_{LED} corresponding to the block number in $qINFO$. We further discuss in Appendix. C on how to validate that $qINFO$ corresponds to the correct request x and its qID .

V. VRF-BASED PROTOCOL

In this section, we present a simple VRF-based randomness service π_{R-VRF} and prove that it implements \mathcal{F}_{VRaaS} functionality. Here, multiple requesters request randomness and there is a single $VRaaS$ server that responds to those requests. We assume that the $VRaaS$ server implements \mathcal{F}_{SERV} by running a VRF protocol. We assume that there is \mathcal{F}_{RLY} (Fig. 7) that acts as a relay between the randomness requester and the $VRaaS$ server. \mathcal{F}_{RLY} reads messages from the blockchain and invokes the $VRaaS$ server with it. Later, we show how to distribute the server by implementing \mathcal{F}_{SERV} using a server committee and \mathcal{F}_{RLY} using a relay committee.

Overview of π_{R-VRF} . We present the generic VRF-based randomness service protocol π_{R-VRF} in Fig. 6 and the message flow can be visualized in Fig. 5. We assume there is a single $VRaaS$ server that registers its public key vk on \mathcal{F}_{LED} . At a high level, a requester requests randomness on its input $x := param$, where $param$ are the parameters for its callback function (the requester receives the output randomness via the callback function once the randomness request is fulfilled).

The request gets uploaded on LOG^{sid} as $qID := (x, reqId)$ by running the $Req-Rand$ smart contract, where $reqId$ is a unique nonce for each request. \mathcal{F}_{RLY} reads the randomness requests from LOG^{sid} . \mathcal{F}_{RLY} verifies that the request was not processed by verifying that a corresponding output tuple is not present on LOG^{sid} . Next, to fulfill the request of the form $qID = (x, reqId)$ it invokes the $VRaaS$ server on input $qINP := qID$ to receive the verifiable output (y, π) . The \mathcal{F}_{RLY} then posts this output on LOG^{sid} by running the $Req-Fulf$ smart contract on the output. For on-chain verification, one needs to check that the output tuple $(REQ-FULF, qINFO, y, \pi; qID, qINP, vk)$ is present on LOG^{sid} (where $qINFO = qINP = qID$). The fact that the output tuple exists on-chain proves that the VRF output (y, π) verifies w.r.t. $(qINP, vk)$. To show that this protocol implements \mathcal{F}_{VRaaS} we need a secure simulatable VRF (Section II) since the \mathcal{F}_{VRaaS} functionality samples

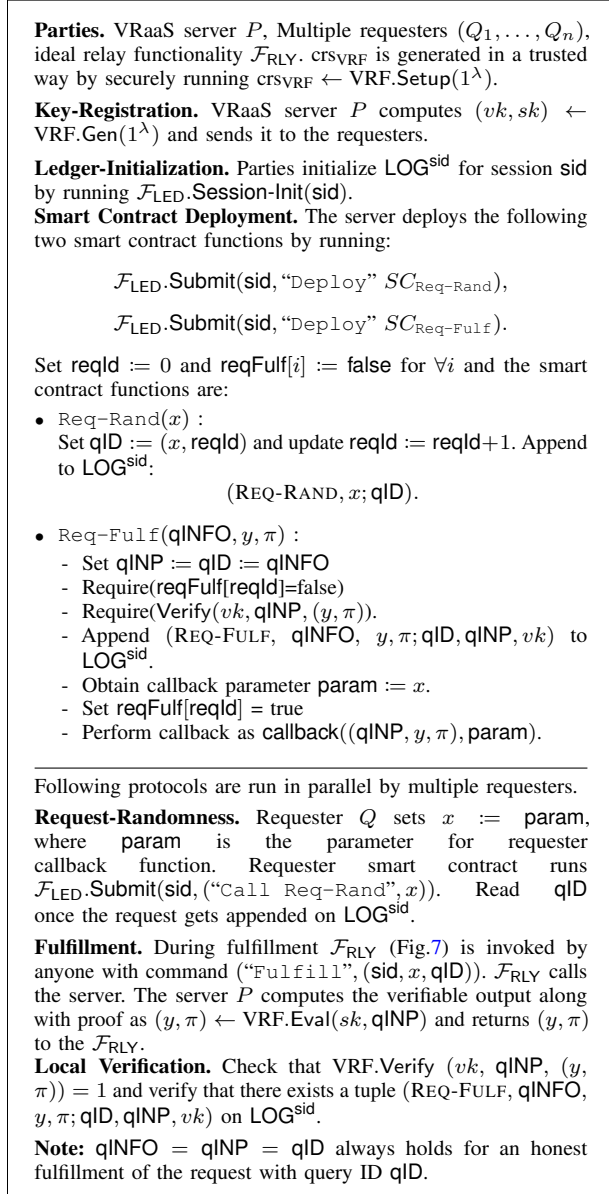


Fig. 6: VRF-Based Randomness Service $\pi_{\text{R-VRF}}$ in $(\mathcal{F}_{\text{LED}}, \mathcal{F}_{\text{RLY}})$ -model.

the random output in the ideal world and the simulator has to simulate the corresponding proof for it. Assuming such a VRF protocol, we prove that the protocol $\pi_{\text{R-VRF}}$ implements the randomness service by proving Thm. 1 in Appendix B.

Theorem 1. *Assume $\text{VRF} = (\text{Setup}, \text{Gen}, \text{Eval}, \text{Verify})$ is a simulatable verifiable random function. Then the protocol $\pi_{\text{R-VRF}}$ (Fig. 6) UC-securely implements $\mathcal{F}_{\text{VRaaS}}$ in the $(\mathcal{F}_{\text{LED}}, \mathcal{F}_{\text{RLY}})$ -model against malicious corruption of the requesters and $\mathcal{F}_{\text{SERV}}$ by a PPT adversary \mathcal{A} .*

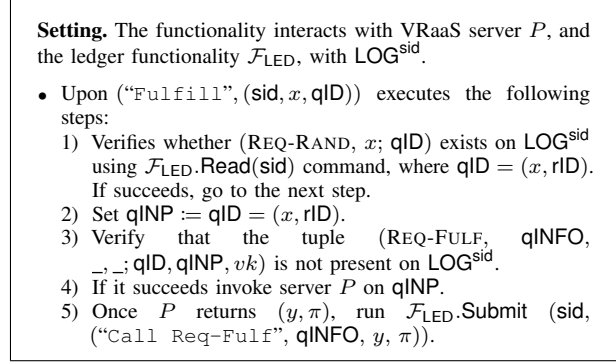


Fig. 7: Ideal Relay Functionality \mathcal{F}_{RLY} .

We show how to relax these two assumptions and discuss how to extend the protocol against adaptive corruption. We also revisit the need for blockchain awareness (i.e. participating in \mathcal{F}_{LED}) from \mathcal{F}_{RLY} and the server. We also argue that variants of $\pi_{\text{R-VRF}}$ capture the Chainlink VRF [41] and Supra dVRF VRF [13].

Distributing the VRaaS Server. The VRaaS server runs the VRF protocol, that consisting of three subprotocols - (Gen, Eval, Verify). These protocols can be distributed by running a distributed VRF [42, 13, 35] protocol based on BLS [22] or DDH-based VRFs. The server is replaced by a committee of n servers out of which at most t servers can be corrupt. The secret key is generated using a distributed key generation (DKG) algorithm [43]. Upon obtaining the secret key shares, each server outputs a partial evaluation of the input and a proof that the partial evaluation is correct. An aggregator party/requester aggregates the partial evaluations from the servers, verifies them, finds $(t + 1)$ correct partial evaluations, and computes the final output. Later, the Gen phase was securely implemented using distributed key generation (DKG) algorithms [43]. The most relevant work is by Galindo et al. [35] who formalized the security properties and analyzed three constructions. The first construction is a distributed pseudorandom function [44, 45], which is essentially a distributed counterpart of the Goldberg et al. [42] protocol with an appropriate zero-knowledge proofs and a specific DKG protocol (a variant of Gennaro et al. [43]) – this is termed as DDH-DVRF. While the computation is very efficient, the size of the final proof is proportional to the number of participants. The second construction they considered is the one that was proposed and also used by Dfinity [46] – this is similar to DDH-DVRF, but uses bilinear pairing to enable a compact proof. However, the use of bilinear groups comes with a cost over discrete log groups (as mentioned later). The construction is very similar to BLS signatures [22] and is used in many places [47, 48, 49, 50]. Their final

construction is called GLOW-DVRF – this was proposed in that paper. GLOW-DVRF uses bilinear pairing for final verification, but Schnorr’s proof of exponent for partial verification.

Adaptive Security of π_{R-VRF} . We briefly discuss the adaptive security of π_{R-VRF} where the adversary can adaptively corrupt servers in the VRaaS committee, nodes implementing \mathcal{F}_{RLY} , and the requesters. We note that adaptive corruption of requesters and the relay nodes does not provide the adversary with any additional benefit since the requester and the relay nodes do not possess any private inputs or private randomness. Hence, adaptive corruption of the relay nodes and the requesters can be trivially simulated, where the simulator simulates the relay nodes and requesters by running the static simulator, and upon adaptive corruption of the parties, the simulator provides the input and random tape as the simulated internal state. To obtain adaptive security for the server, the VRF.Gen has to be distributed using a DKG algorithm that is secure against adaptive corruption of parties. Similarly, the VRF.Eval algorithm also needs to be distributed using an adaptively secure protocol. The recent work of [51] provides an adaptively secure threshold BLS signature [52] scheme that relies on the hardness of DDH and co-CDH in asymmetric pairing group assuming random oracles. Previous protocols [53] proved adaptive security for the same, assuming One More Discrete Logarithm in the Algebraic Group Model [51] would suffice for our adaptively secure instantiation of the VRaaS server in π_{R-VRF} .

Realizing \mathcal{F}_{RLY} . \mathcal{F}_{RLY} can be realized by a committee of m nodes satisfying an honest majority. The nodes will have access to the \mathcal{F}_{LED} functionality and the adversary \mathcal{A} can corrupt at most $t' < \frac{m}{2}$ nodes. The VRaaS server/committee waits for $t' + 1$ evaluation requests that match and then run the VRaaS evaluation protocol. A strict-honest majority is required among the relay nodes to ensure that the evaluation request is valid and exists on LOG^{sid} since the VRaaS server/committee does not have access to \mathcal{F}_{LED} . Upon evaluating the VRF, the VRaaS server (or the committee implementing \mathcal{F}_{SERV}) sends the output to the relay nodes. We require that any of the relay nodes behave honestly and perform fulfillment to ensure liveness. This is the Anytrust [54] assumption.

Access to ledger functionality \mathcal{F}_{LED} . π_{R-VRF} assumes that the \mathcal{F}_{RLY} node, but not the VRaaS server, has access to \mathcal{F}_{LED} (i.e. blockchain-aware). In practical terms, this means that in a multi-blockchain environment, requesters can seek randomness from the VRaaS server through any blockchain. Each blockchain has an assigned \mathcal{F}_{RLY} that processes requester requests by forwarding them to the VRaaS server. The VRaaS server computes the output and sends it back to \mathcal{F}_{RLY} , which then posts it on the originating blockchain. Supra dVRF [13]

implements this. This approach simplifies scalability and allows efficient deployment of the VRaaS service on different blockchains by assigning a \mathcal{F}_{RLY} node to the blockchain, without modifying the server committee. However, it requires the \mathcal{F}_{RLY} to be implemented using a committee of nodes satisfying honest majority assumption. Another approach is to remove \mathcal{F}_{RLY} and allow the server to be blockchain-aware. The server reads the randomness requests from the blockchain, evaluates the output, and then posts it on the blockchain. Chainlink VRF [41] implements this. However, this is not preferable since the VRaaS server/committee has to keep track of different blockchains and it is not scalable as the number of blockchains increases. Onboarding a new VRaaS server/committee is also taxing as they have to be blockchain-aware for all the new blockchains.

Next, we argue that there cannot be a secure VRaaS protocol where 1) the VRaaS server does not participate in \mathcal{F}_{LED} , 2) there does not exist any party other than the requester and VRaaS server who participates in \mathcal{F}_{LED} , and 3) there are no other setup assumptions (like access to a broadcast channel) between the VRaaS server and requesters. We prove this by contradiction. Assume that such a protocol exists for on-chain verifiable randomness service between a VRaaS server and a requester where the requester is in charge of uploading the randomness requests and fulfillment/output transactions on \mathcal{F}_{LED} . Since the VRaaS server does not have access to \mathcal{F}_{LED} , it cannot read the randomness requests posted by the requester on \mathcal{F}_{LED} . The only way the VRaaS server can read the requests is when the requester directly sends the request to the VRaaS server. The server evaluates the output and has to send it directly to the requester since the VRaaS server does not have access to \mathcal{F}_{LED} . The requester is in charge of uploading this output as a transaction on \mathcal{F}_{LED} . Such a protocol allows a malicious requester to query the VRaaS server on multiple inputs, and obtain multiple outputs. Then the requester chooses the input-output pair that favors it the most and registers that pair as a transaction on \mathcal{F}_{LED} . The honest VRaaS server cannot detect this due to lack of access to \mathcal{F}_{LED} . This breaks unbiasedness of the randomness service.

A. Real-world Randomness Services

Finally, we conclude this section by showing that Chainlink VRF service and Supra dVRF service are captured via π_{R-VRF} protocol.

Analysis of Chainlink VRF Service. In Chainlink VRF [41], to request randomness the requester Q sets $x := (\text{khash}, \text{account}, \text{config}, \text{gas}, \text{len})$ where $\text{khash} = H(vk_i)$ specifies which trusted VRF server should fulfill the request, account is the client’s account information, config is the number of confirmation blocks the VRF server should wait, gas is the maximum gas

(on-chain instructions run by smart contract) cost to be executed in the callback function. The client smart contract generates a randomness request. The Chainlink smart contract verifies the authenticity of the request by checking the account information. It maintains a counter `reqId`. It computes `preseed = H(khash, cinfo, account, reqId)` and `qID = H(khash, preseed)`, and increments the counter `reqId`. It posts $(x; \text{qID})$ on the ledger and returns `qID`. It also maintains an on-chain commitment to store the hash of the above information for validation.

Once the transaction gets uploaded, the Chainlink VRF server reads $(x; \text{qID})$, and once `config` blocks have been confirmed on the ledger after the block containing the request, the Chainlink server computes `preseed` from $(x; \text{qID})$, verifies the hash of the request information against the on-chain commitment and computes the Goldberg-VRF [36] on $\text{qINP} := \text{H}(\text{preseed}, \text{bhash})$ to obtain a verifiable output (y, π) , where `bhash` is the block-hash (of the block containing $(x; \text{qID})$) preventing precomputation of the VRF output. The VRF server then calls the smart contract to register the fulfillment transaction on-chain. The smart contract validates the randomness y using proof π and performs the callback with (y, π) .

In the Chainlink VRF service, there are no relay nodes and individual VRF servers are blockchain-aware. It also relies on on-chain commitments to validate the information in `qINP` during fulfillment transactions. Next, we discuss that the output is unique for every randomness request. Each `qINP` is guaranteed to be unique, even for the same client inputs, due to the unique `reqId` added by the Chainlink smart contract in `preseed` (assuming `H` is collision-resistant). This translates to unique pseudorandom output y due to the pseudorandomness of the underlying VRF protocol, thus proving that the Chainlink VRF service implements the $\mathcal{F}_{\text{VRaaS}}$ functionality.

Analysis of Supra dVRF Service. In the Supra dVRF service [13], the VRaaS server is run by a clan of VRF server nodes that run the distributed version of BLS signatures. \mathcal{F}_{RLY} in $\pi_{\text{R-VRF}}$ is realized by a committee of relay nodes. To request randomness, the requester Q sets its input $x := (\text{param}, \text{config}, \text{cseed}, \text{account})$ where `param` is the callback function, `config` is the number of confirmation blocks the VRF Committee should wait, `cseed` (can be empty) is client provided entropy, and `account` is the account information (w.r.t. billing etc.). The client smart contract invokes the Supra smart contract with input x . The Supra Smart-contract checks that the client account has sufficient balance and it returns `qID := reqId`, where `reqId` is unique for every client request. The final request transaction that is posted is of the form $(x; \text{qID}, \text{aux})$ where `aux` contains some auxiliary information (the number of confirmation blocks, amount

of randomness to be generated, etc). `qINFO` contains the same information as (x, aux) and so Supra doesn't append it on-chain.

Then, the committee of relay nodes reads the request $(x; \text{qID}, \text{aux})$, and once `config` blocks have been confirmed on the ledger after `qID`, the relay nodes invoke the Supra dVRF clan on `qINP`. The VRF clan computes the VRF on $\text{qINP} = (x, \text{qID}, \text{bhash})$ to obtain a verifiable output (y, π) , where `bhash` is the block-hash (of the $N + \text{config}$ block where the request was confirmed on N th block) preventing precomputation of the VRF output. The relay nodes collect partial evaluations from the Supra dVRF clan, verify the partial evaluations, and then combine them to obtain each fulfillment output. Then they invoke the fulfillment transaction to post the output on-chain and send the output to the respective requester Q via the callback function in `param`.

The Supra dVRF clan has an honest-majority assumption and models the distributed VRaaS server. The Supra architecture can be considered as our $\pi_{\text{R-VRF}}$ protocol from Section. V where the relay committee is blockchain-aware (Approach 1 in Section V). It also reuses the trust assumption of the dVRF committee to validate that the `qINP` is valid and compute the VRF output on it after it is deemed to be valid. This is more efficient in practice as opposed to maintaining an on-chain commitment (what Chainlink does). To ensure that the Supra protocol securely implements $\mathcal{F}_{\text{VRaaS}}$ we only need to argue that each `qINP` $= (x, \text{qID}, \text{bhash})$ is unique. This is ensured since the `qID` is guaranteed to be unique due to the on-chain counter `reqId` maintained by the Supra dVRF smart contract. Hence, the Supra dVRF service implements the $\mathcal{F}_{\text{VRaaS}}$ functionality.

VI. VULNERABILITIES IN EXISTING PROTOCOLS

We present analyses of Algorand [17], Band-VRF [10] and Pyth-VRF [11] and show how they fail to implement $\mathcal{F}_{\text{VRaaS}}$. In Appendix E, we demonstrate vulnerabilities in the DIA xRandom smart contract code.

Algorand. The Algorand randomness service operates on a beacon model, where users rely on the Algorand randomness beacon for generating random values. Best practices for requester protocols are outlined in the Algorand developer documentation, including the use of requester smart contracts to implement these practices. One key practice involves the beacon smart contract storing values for a limited number of rounds before updates occur, and requester protocols should read beacon values before updates. Requester protocols need to handle scenarios such as beacon downtime, updates, or discontinuation of the beacon public key, which can be complex due to the need for proper auditing. An example scenario involves a client smart contract function, `fallbackrand()`, which uses the next beacon if the current

one is unavailable. However, a malicious user could exploit this by delaying transactions, causing the current beacon to be unavailable (since its value is unfavorable) until a later beacon is available. This way `fallbackrand()` uses the later beacon, allowing the party to effectively reroll its randomness, thereby breaking the protocol’s unbiasedness. In our $\mathcal{F}_{\text{VRaaS}}$ functionality, Step 11. cannot be simulated using this protocol as the requester can reroll y even when $\mathcal{F}_{\text{SERV}}$ is honest.

Band VRF Protocol. The Band VRF [10] protocol is a VRF-based protocol where the requester (denoted as the client in their protocol) smart contract calls the VRF server smart contract to request randomness. The Band VRF protocol operates across two blockchains via a bridge service and we refer to their protocol [10] for more details on this. We observe a vulnerability in the Band VRF provider code [55]. One of the input parameters to Band VRF provider (line 111 in [55]) is a client seed `clientseed`. However, the same `clientseed` cannot be repeated for a given client smart contract. If the same client smart contract makes two randomness requests with the same `clientseed`, only the first one will be fulfilled (lines 117-119 in [55]). An attacker monitors the requests containing the different `clientseed` values. These requests are public since they are on-chain. The attacker launches a denial-of-service attack by running the same client smart contract with the same `clientseed` value. If the attacker’s transaction gets confirmed first then the client’s request will be denied. In our $\mathcal{F}_{\text{VRaaS}}$ functionality, Step 5. cannot be simulated using this protocol as the Band VRF protocol rejects an honest client’s request on `clientseed` if the adversary has already made the same request previously. In the ideal world the honest client’s request will always be fulfilled allowing the adversary to distinguish.

Pyth-VRF. The Pyth-VRF protocol works in the commit-and-response paradigm where the server precomputes N random values (x_1, \dots, x_N) and commits to it on-chain. To request randomness, the requester samples a random number x_U and sends the hash $h_U = H(x_U)$ to the smart contract. The smart assigns a unique sequence number i to the request and stores the hash value $h_U = H(x_U)$ and sequence number on-chain. The smart contract increments the counter i so that every request gets assigned a unique sequence number. Upon receiving the sequence number i , the requester invokes the Pyth-VRF server with i to obtain x_i . Given x_i the requester computes the random value as $r := H(x_i, x_U)$ and invokes the smart contract with x_i to complete the fulfillment process. We describe an attack in Pyth-VRF as follows. The requester preemptively computes the output value r once it receives x_i from the server. If the malicious requester doesn’t like r then it does not complete the fulfillment phase. As a result, the output r is not generated on-chain. The smart contract or the server cannot generate the output since the requester’s

input x_U is secret. In our $\mathcal{F}_{\text{VRaaS}}$ functionality, Steps 11., 13. cannot be simulated using this protocol as the malicious requester can reroll r (denoted as y in $\mathcal{F}_{\text{VRaaS}}$) even when $\mathcal{F}_{\text{SERV}}$ is honest and prevent the fulfillment process if the output is unfavorable.

VII. IMPOSSIBILITY OF OBTAINING VRAAS USING ONE TRANSACTION

A VRaaS protocol/functionality using one online transaction would reduce latency and gas costs for the requester. Online transaction refers to the request randomness transaction, denoted `Submit(sid, (“Call Req-Rand”, x))` and request fulfillment transaction, denoted `Submit(sid, (“Call Req-Fulf”, $q\text{INFO}$, y , π))`, and not the smart-contract deployment transactions. However, we show that it is not possible to construct such a protocol satisfying our formalization. Below we provide some intuitions on how such protocols would be vulnerable to attacks in practice.

First, when the fulfillment transaction is not submitted, a malicious requester denies receiving an unfavorable output. In that case, it would not be possible to distinguish between the case whether the server aborted, or the requester is lying.

Second, when the request transaction is not submitted, a malicious requester may just run multiple request sessions with the VRF service in parallel, and then submit the fulfillment request for whichever is more favorable. This way, again a biased output can be obtained. For example, if the requester makes two simultaneous requests, with probability $3/4$ the first bit of the output is 0. Note that the VRF service may be able to link and subsequently stop the conflicting requests from the malicious requester assuming a public-key infrastructure (PKI) setup and signature-based authentication. However, a PKI setup between clients and VRF service is not practical if we expect the service to scale for a large number of requester clients.

Formally we prove the following theorem.

Theorem 2. *Unless there is a PKI setup, there does not exist a protocol that realizes the ideal functionality $\mathcal{F}_{\text{VRaaS}}$ using only one online transaction in the \mathcal{F}_{LED} model.*

Proof. We consider two cases. First, consider an arbitrary protocol that does not deploy the fulfillment transaction. For such protocol, we design an adversarial requester which works as follows

- Assume the protocol returns a bit output y . The requester makes a request and when it obtains the response (y, π) directly from the relay nodes it checks whether $y = 0$, if not then it denies receiving the output. Otherwise it publishes the value.

Clearly, for such an adversary, in the ideal world, the fulfillment transaction would be executed in Step 13. by the ideal functionality $\mathcal{F}_{\text{VRaaS}}$ – this implies that for any output (y, π) , a verification query made in the ideal world by any honest party would always return 1. However, in the real world, when $y = 1$, the verification would fail as the protocol was never completed. So, the adversary would be able to distinguish between the real and ideal world with probability $1/2$ (i.e. whenever the random output y is 1).

In the second case, assume that there is no PKI setup and no request transaction. A malicious requester’s requests can not be linked. Thus, the malicious requester may provide two requests with signatures with different public keys and the attack works as follows:

- Send two unlinkable requests to the servers. The server returns two values (y_1, π_1) and (y_2, π_2) . The servers have run the fulfillment transaction on both values.
- Now, the requester chooses the output whose value is 0 and links it. Essentially, the requester specifically chooses the output corresponding to its identity Q if the output is 0.

In the ideal world, both output pairs (y_1, π_1) and (y_2, π_2) would be linked to Q since a requester can only make queries from the identity that is registered with the functionality. So for any output, the probability of it being 0 would be $1/2$. Whereas, in the real world, the probability of the linked output to be 0 would be $3/4$, because, among four possible combinations, only one pair $(1, 1)$ would not result into a non-zero output. This will be distinguishable with probability $1/4$. \square

VIII. CONCLUSION

Our work performs a comprehensive analysis of on-chain verifiable randomness services. We first formalize the on-chain verifiable randomness in the blockchain setting by introducing the notion of Verifiable Randomness as a Service (VRaaS) – and define it via an ideal functionality. It features a smart contract on a ledger functionality and the VRF provider. We demonstrate that two blockchain transactions are necessary for a secure realization of VRaaS. Towards sufficiency, we show that the VRF based protocol is a VRaaS by proving that it implements our ideal functionality.

We rigorously analyzed existing randomness services in the Web3 ecosystem and demonstrated that our framework captures randomness services such as Chainlink VRF and Supra dVRF services. Our investigation also revealed susceptibility to attacks for three other designs namely, Band VRF, DIA xRandom, and Algorand randomness beacon, and we have responsibly disclosed the vulnerabilities to the teams.

Beyond analyzing existing systems, our VRaaS framework also offers several challenges to consider in the fu-

ture. These include securely employing random beacons, VDFs, secure randomness usage patterns in multi-player Web3 applications, and generating private randomness.

REFERENCES

- [1] “Filecoin: A decentralized storage network.” [Online]. Available: <https://filecoin.io/filecoin.pdf>
- [2] “Partnerships of algorand.” [Online]. Available: <https://algorandtechnologies.com/about/our-partners/>
- [3] “Partnerships of supra.” [Online]. Available: <https://supraoracles.com/partnerships/>
- [4] “Use cases of chainlink.” [Online]. Available: <https://chain.link/use-cases>
- [5] a16zcrypto, “2023 state of crypto report: Introducing the state of crypto index.” [Online]. Available: <https://a16zcrypto.com/posts/article/state-of-crypto-report-2023/>
- [6] —, “Introducing the 2022 state of crypto report.” [Online]. Available: <https://a16zcrypto.com/posts/article/state-of-crypto-report-a16z-2022/>
- [7] Chainlink, “Chainlink VRF: On-Chain Verifiable Randomness.” https://developer.wax.io/en/tutorials/create-wax-rng-smart-contract/rng_basics.html.
- [8] “Supra randomness service.” [Online]. Available: <https://supraoracles.com/>
- [9] “Drand: Distributed randomness beacon.” [Online]. Available: <https://drand.love/>
- [10] “Band vrf guaranteed integrity on the blockchain.” [Online]. Available: <https://www.bandprotocol.com/vrf>
- [11] “Pythvrf: Random number generation on pyth network.” [Online]. Available: <https://docs.pyth.network/documentation/entropy/protocol-design>
- [12] “Chainlink random number generation.” [Online]. Available: <https://chain.link/vrf>
- [13] “Supra whitepaper.” [Online]. Available: <https://supraoracles.com/docs/SupraOracles-VRF-Service-Whitepaper.pdf>
- [14] E. Albert, S. Grossman, N. Rinetzky, C. Rodríguez-Núñez, A. Rubio, and M. Sagiv, “Taming callbacks for smart contract modularity,” *Proc. ACM Program. Lang.*, vol. 4, no. OOPSLA, pp. 209:1–209:30, 2020.
- [15] A. Kate, E. V. Mangipudi, S. Maradana, and P. Mukherjee, “Flexirand: Output private (distributed) vrf’s and application to blockchains,” in *ACM CCS 2023*, 2023.
- [16] “Dia developer documents.” [Online]. Available: <https://docs.diadata.org/products/randomness-oracle/access-the-oracle>
- [17] O. Shem-Tov, “Usage and best practices for randomness beacon,” Sep 2022. [Online]. Available: <https://developer.algorand.org/articles/usage-and-best-practices-for-randomness-beacon/>
- [18] R. Canetti, “Universally composable security: A new paradigm for cryptographic protocols,” in *42nd FOCS*. IEEE Computer Society Press, Oct. 2001, pp. 136–145.
- [19] B. David, B. Magri, C. Matt, J. B. Nielsen, and D. Tschudi, “GearBox: Optimal-size shard committees by leveraging the safety-liveness dichotomy,” in *ACM CCS 2022*, H. Yin, A. Stavrou, C. Cremers, and E. Shi, Eds. ACM Press, Nov. 2022, pp. 683–696.
- [20] “Harmony randomness service.” [Online]. Available: <https://docs.harmony.one/home/>
- [21] “Boba network: Distributed randomness beacon.” [Online]. Available: <https://boba.network>
- [22] D. Boneh, B. Lynn, and H. Shacham, “Short signatures from the Weil pairing,” in *ASIACRYPT 2001*, ser. LNCS, C. Boyd, Ed., vol. 2248. Springer, Heidelberg, Dec. 2001, pp. 514–532.
- [23] Aptos, “Aip-41 - move apis for randomness generation.” [Online]. Available: <https://github.com/aptos-foundation/AIPs/blob/main/aips/aip-41.md>
- [24] “Chia.” [Online]. Available: <https://www.chia.net/>
- [25] “Veedo — stark-based verifiable delay function.” [Online]. Available: <https://github.com/starkware-libs/veedo>

- [26] M. Christ, K. Choi, and J. Bonneau, “Cornucopia: Distributed randomness beacons at scale,” *IACR Cryptol. ePrint Arch.*, p. 1554, 2023. [Online]. Available: <https://eprint.iacr.org/2023/1554>
- [27] C. Baum, B. David, R. Dowsley, R. Kishore, J. B. Nielsen, and S. Oechsner, “CRAFT: composable randomness beacons and output-independent abort MPC from time,” in *PKC’23*, 2023.
- [28] D. Boneh, J. Bonneau, B. Bünz, and B. Fisch, “Verifiable delay functions,” in *CRYPTO 2018, Part I*, ser. LNCS, H. Shacham and A. Boldyreva, Eds., vol. 10991. Springer, Heidelberg, Aug. 2018, pp. 757–788.
- [29] A. K. Lenstra and B. Wesolowski, “A random zoo: sloth, unicorn, and trx,” *IACR Cryptol. ePrint Arch.*, p. 366, 2015. [Online]. Available: <http://eprint.iacr.org/2015/366>
- [30] R. Canetti, “Universally composable signature, certification, and authentication,” in *17th IEEE Computer Security Foundations Workshop, (CSFW-17 2004), 28-30 June 2004, Pacific Grove, CA, USA*. IEEE Computer Society, 2004, p. 219.
- [31] R. Canetti, A. Jain, and A. Scafuro, “Practical UC security with a global random oracle,” in *ACM CCS 2014*, G.-J. Ahn, M. Yung, and N. Li, Eds. ACM Press, Nov. 2014, pp. 597–608.
- [32] R. Canetti, P. Sarkar, and X. Wang, “Efficient and round-optimal oblivious transfer and commitment with adaptive security,” in *ASIACRYPT 2020, Part III*, ser. LNCS, S. Moriai and H. Wang, Eds., vol. 12493. Springer, Heidelberg, Dec. 2020, pp. 277–308.
- [33] J. Camenisch, M. Drijvers, T. Gagliardoni, A. Lehmann, and G. Neven, “The wonderful world of global random oracles,” in *EUROCRYPT 2018, Part I*, ser. LNCS, J. B. Nielsen and V. Rijmen, Eds., vol. 10820. Springer, Heidelberg, Apr. / May 2018, pp. 280–312.
- [34] M. Chase and A. Lysyanskaya, “Simulatable VRFs with applications to multi-theorem NIZK,” in *CRYPTO 2007*, ser. LNCS, A. Menezes, Ed., vol. 4622. Springer, Heidelberg, Aug. 2007, pp. 303–322.
- [35] D. Galindo, J. Liu, M. Ordean, and J. Wong, “Fully distributed verifiable random functions and their application to decentralised random beacons,” in *IEEE EuroS&P 2021*, 2021.
- [36] D. Papadopoulos, D. Wessels, S. Huque, M. Naor, J. Včelák, L. Reyzin, and S. Goldberg, “Making nsec5 practical for dnssec,” *Cryptology ePrint Archive*, Paper 2017/099, 2017. [Online]. Available: <https://eprint.iacr.org/2017/099>
- [37] M. Graf, D. Rausch, V. Ronge, C. Egger, R. Küsters, and D. Schröder, “A security framework for distributed ledgers,” in *ACM CCS 2021*, G. Vigna and E. Shi, Eds. ACM Press, Nov. 2021, pp. 1043–1064.
- [38] R. Kumaresan, D. V. Le, M. Minaei, S. Raghuraman, Y. Yang, and M. Zamani, “Programmable payment channels,” in *ACNS 2024*, 2024.
- [39] S. Dziembowski, L. Eckey, S. Faust, J. Hesse, and K. Hostáková, “Multi-party virtual state channels,” in *EUROCRYPT 2019, Part I*, ser. LNCS, Y. Ishai and V. Rijmen, Eds., vol. 11476. Springer, Heidelberg, May 2019, pp. 625–656.
- [40] C. Baum, B. David, R. Dowsley, J. B. Nielsen, and S. Oechsner, “TARDIS: A foundation of time-lock puzzles in UC,” in *EUROCRYPT 2021, Part III*, ser. LNCS, A. Canteaut and F.-X. Standaert, Eds., vol. 12698. Springer, Heidelberg, Oct. 2021, pp. 429–459.
- [41] “Chainlink vrf: On-chain verifiable randomness.” [Online]. Available: <https://blog.chain.link/chainlink-vrf-on-chain-verifiable-randomness/>
- [42] S. Goldberg, J. Vcelak, D. Papadopoulos, and L. Reyzin, “Verifiable random functions (vrf),” <https://datatracker.ietf.org/doc/html/draft-goldbe-vrf-01>, 2018.
- [43] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Secure distributed key generation for discrete-log based cryptosystems,” *Journal of Cryptology*, vol. 20, no. 1, pp. 51–83, Jan. 2007.
- [44] S. Agrawal, P. Mohassel, P. Mukherjee, and P. Rindal, “DiSE: Distributed symmetric-key encryption,” in *ACM CCS 2018*, D. Lie, M. Mannan, M. Backes, and X. Wang, Eds. ACM Press, Oct. 2018, pp. 1993–2010.
- [45] M. Naor, B. Pinkas, and O. Reingold, “Distributed pseudo-random functions and KDCs,” in *EUROCRYPT’99*, ser. LNCS, J. Stern, Ed., vol. 1592. Springer, Heidelberg, May 1999, pp. 327–346.
- [46] T. Hanke, M. Movahedi, and D. Williams, “DFINITY technology overview series, consensus system,” *CoRR*, vol. abs/1805.04548, 2018. [Online]. Available: <http://arxiv.org/abs/1805.04548>
- [47] Cloudflare, “Decentralized Verifiable Randomness Beacon,” <https://developers.cloudflare.com/randomness-beacon/>.
- [48] Corestar, “Corestar Arcade: Tendermint-based Byzantine Fault Tolerant (BFT) middleware with an embedded BLS-based random beacon,” <https://github.com/corestar/tendermint>.
- [49] DAOBet (ex — DAO.Casino), “To Deliver On-Chain Random Beacon Based on BLS Cryptography.” <https://daobet.org/blog/on-chain-random-generator/>.
- [50] P. Schindler, A. Judmayer, N. Stifter, and E. Weippl, “ETHDKG: Distributed key generation with Ethereum smart contracts,” *Cryptology ePrint Archive*, Report 2019/985, 2019, <https://eprint.iacr.org/2019/985>.
- [51] S. Das and L. Ren, “Adaptively secure BLS threshold signatures from DDH and co-cdh,” *IACR Cryptol. ePrint Arch.*, p. 1553, 2023. [Online]. Available: <https://eprint.iacr.org/2023/1553>
- [52] A. Boldyreva, “Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme,” in *PKC 2003*, ser. LNCS, Y. Desmedt, Ed., vol. 2567. Springer, Heidelberg, Jan. 2003, pp. 31–46.
- [53] R. Bacho and J. Loss, “On the adaptive security of the threshold BLS signature scheme,” in *ACM CCS 2022*, H. Yin, A. Stavrou, C. Cremers, and E. Shi, Eds. ACM Press, Nov. 2022, pp. 193–207.
- [54] D. I. Wolinsky, H. Corrigan-Gibbs, and B. Ford, “Scalable anonymous group communication in the anytrust model.” [Online]. Available: <https://dedis.cs.yale.edu/dissent/papers/eurosec12-abs/>
- [55] “Band vrf provider code.” [Online]. Available: https://github.com/bandprotocol/vrf-and-bridge-contracts/blob/34df2ebb75355bfebf9ad24efb12c4c3e2c328e5/contracts/vrf/provider_v2/VRFProviderBaseV2.sol#L111
- [56] J. Camenisch, S. Krenn, R. Küsters, and D. Rausch, “iUC: Flexible universal composability made simple,” in *ASIACRYPT 2019, Part III*, ser. LNCS, S. D. Galbraith and S. Moriai, Eds., vol. 11923. Springer, Heidelberg, Dec. 2019, pp. 191–221.
- [57] C. Badertscher, P. Gazi, A. Kiayias, A. Russell, and V. Zikas, “Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability,” in *ACM CCS 2018*, D. Lie, M. Mannan, M. Backes, and X. Wang, Eds. ACM Press, Oct. 2018, pp. 913–930.
- [58] T. Kerber, A. Kiayias, M. Kohlweiss, and V. Zikas, “Ouroboros cryptsinous: Privacy-preserving proof-of-stake,” in *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2019, pp. 157–174.
- [59] R. G. Brown, “The corda platform: An introduction,” 2020. [Online]. Available: <https://www.r3.com/wp-content/uploads/2019/06/corda-platform-whitepaper.pdf>. (Accessed on 28/05/2020)
- [60] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, “OmniLedger: A secure, scale-out, decentralized ledger via sharding,” in *2018 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2018, pp. 583–598.
- [61] A. R. Choudhuri, V. Goyal, and A. Jain, “Founding secure computation on blockchains,” in *EUROCRYPT 2019, Part II*, ser. LNCS, Y. Ishai and V. Rijmen, Eds., vol. 11477. Springer, Heidelberg, May 2019, pp. 351–380.
- [62] I. Komargodski and Y. Tamir, “On distributed randomness generation in blockchains,” in *Cyber Security, Cryptology, and Machine Learning: 7th International Symposium, CSCML 2023, Be’er Sheva, Israel, June 29–30, 2023, Proceedings*. Springer-Verlag, 2023, p. 49–64.
- [63] R. Canetti, P. Sarkar, and X. Wang, “Triply adaptive UC NIZK,” in *ASIACRYPT 2022*, 2022.
- [64] T. Kerber, A. Kiayias, and M. Kohlweiss, “KACHINA - foundations of private smart contracts,” in *CSF 2021 Computer Security Foundations Symposium*, R. Küsters and D. Naumann, Eds. IEEE Computer Society Press, 2021, pp. 1–16.

[65] D. developer documentation, “Dia dice game.” [Online]. Available: <https://docs.diadata.org/products/randomness-oracle/access-the-oracle#example-dice-game>

APPENDIX A

SIMULATABLE VERIFIABLE RANDOM FUNCTION

We demonstrate that the most widely known VRF protocols - BLS/GLOW-VRF [35], RSA-based [36], and the Elliptic-Curve based [36] VRFs are simulatable by modeling the hash function used in the VRF protocols as a programmable random oracle.

- **BLS/GLOW-VRF [35].** The verification key is $vk = g_2^{sk}$ and secret key is sk . The input is x . The proof is $\pi = H_1(x)^{sk}$ and the output value is $y = H_2(\pi)$ where H_1, H_2 are random oracles. The verifier checks $e(\pi, g_2) \stackrel{?}{=} e(H_1(x), vk)$ and $y \stackrel{?}{=} H_2(\pi)$.

Simulatability. To simulate the VRF on unqueried input x to output y' , the SimProve programs H_2 s.t. $H_2(\pi) = y'$. The proof $\pi' = \pi$ remains unchanged. Simulatability follows from the unpredictability of $\pi = H_1(x)^{sk}$ without querying the VRF on x .

- **RSA-based [36].** The verification key is (n, e) and the secret key is d . The input is x . The proof is $\pi = H_1(x)^d \bmod n$ and the output value is $y = H_2(\pi)$ where H_1 is an IETF specified hash function and H_2 is a cryptographic hash function. The verification equation is $\pi^e \bmod n \stackrel{?}{=} H_1(x)$ and $y \stackrel{?}{=} H_2(\pi)$.

Simulatability. To simulate the VRF on unqueried input x to output y' , SimProve programs H_2 s.t. $H_2(\pi) = y'$. The proof $\pi' = \pi$ remains unchanged. Simulatability of the VRF follows from the unpredictability of $\pi = H_1(x)^d \bmod n$ without querying the VRF on x .

- **Elliptic-Curve-based [36].** This is used by Chainlink [41]. The secret key is $sk \in \mathbb{Z}_q$ and the public verification key is $vk = g^{sk}$. The input is x . Compute $h = H_1(x)$ and $\gamma = h^{sk}$. The output is $y = H_2(\gamma^f)$ for a public parameter f . To compute the VRF proof, compute a discrete log proof as: sample $k \leftarrow \mathbb{Z}_q$, set $c = H_3(g, h, vk, \gamma, g^k, h^k)$ and $s = k - cx \bmod q$. Set y as the output and $\pi = (\gamma, c, s)$ as the proof. To verify the output check that 1) $y \stackrel{?}{=} H_2(\gamma^f)$, and 2) compute $u = (vk)^c \cdot g^s$, $h = H_1(x)$, $v = \gamma^c \cdot h^s$ and check $c \stackrel{?}{=} H_3(g, h, vk, \gamma, u, v)$.

Simulatability. To simulate the VRF on unqueried input x to output y' , SimProve programs H_2 s.t. $H_2(\gamma^f) = y'$. The proof $\pi' = \pi$ remains unchanged. Simulatability follows from the unpredictability of γ without querying VRF on x .

- **Unforgeable VRF.** One can also build a simulatable VRF from a VRF protocol that satisfies unforgeability in the programmable random oracle model. Let $\text{VRF}' = (\text{Setup}', \text{Gen}', \text{Eval}', \text{Verify}')$

be a verifiable random function satisfying unforgeability. Then we construct a simulatable VRF = (Setup, Gen, Eval, Verify) assuming a programmable random oracle H as follows. Set $\text{Setup} := \text{Setup}'$ and $\text{Gen} := \text{Gen}'$. We set $\text{VRF.Eval}(sk, x; \text{crs}) : \text{Return output } (y, \pi) := (H(y'), (y', \pi'))$ where $(y', \pi') := \text{VRF}'.\text{Eval}'(sk, x; \text{crs})$. To verify (x, y, π) , parse (y', π') , check that 1) $y = H(y')$, and 2) $\text{Verify}'(vk, x, (y', \pi'); \text{crs})$.

Simulatability. To argue simulatability, the simulator programs H on y' to return the specific random y . The adversary cannot distinguish since y' is unforgeable.

These protocols are widely used by Dfinity [46], Chainlink VRF [41], and Supra dVRF [13]. Simulatability is a stronger notion than pseudorandomness [35] and simulatability is hard [34] to achieve in the non-programmable random oracle model.

Our simulatable VRF can be extended to the global random oracle (GRO) model of [33, 31] where a single instance of the random oracle is shared by multiple protocol sessions. We can use the random oracle $\mathcal{G}_{\text{rpoRO}}$ from Fig. 10 in [33] for this purpose. It provides restricted programmability and restricted observability, i.e. the simulator is allowed to program and observe GRO queries corresponding to the session it is participating in. In our simulatable VRF protocols described above, hash function H_2 will be modeled as the GRO with restricted programmability and restricted observability. The simulator can program H_2 without getting detected since the input to H_2 is unpredictable to an adversary and hence the adversary cannot prevent the programming.

APPENDIX B

PROOF OF THEOREM 1

We prove that $\pi_{\text{R-VRF}}$ implements $\mathcal{F}_{\text{VRaaS}}$ by proving Thm. 1. We assume that there exists a PPT environment \mathcal{Z} that selects the inputs for the honest parties and it instructs a PPT adversarial algorithm \mathcal{A} to corrupt a participating party in the protocol execution. In the real-world execution of the protocol, \mathcal{A} corrupts a party and interacts with the rest of the honest parties. At the end of the protocol execution, it forwards its view $\text{REAL}_{\pi_{\text{R-VRF}}, \mathcal{A}, \mathcal{Z}}(1^\lambda)$ to the environment \mathcal{Z} . In the ideal world execution of the protocol, we provide a PPT simulator Sim that given access to the adversarial algorithm \mathcal{A} and the functionality $\mathcal{F}_{\text{VRaaS}}$ produces the ideal world adversary view $\text{IDEAL}_{\mathcal{F}_{\text{VRaaS}}, \text{Sim}, \mathcal{Z}}(1^\lambda)$ and forwards it to the environment \mathcal{Z} . According to the UC definition, these two views should be indistinguishable.

Proof. Sim has access to the \mathcal{F}_{LED} functionality. The ideal world adversary \mathcal{A}_{LED} for \mathcal{F}_{LED} is invoked by Sim to setup \mathcal{F}_{LED} . We denote the ideal world adversary for $\mathcal{F}_{\text{SERV}}$ as Sim_p . We exhaustively consider all the

corruption cases where \mathcal{A} corrupts a combination of the VRaaS server and the client Q . We prove the security of our protocol in the \mathcal{F}_{RLY} -model where Sim simulates \mathcal{F}_{RLY} . We describe the simulation steps and argue indistinguishability between the real and ideal world execution for the two corruption cases as follows.

Server P is corrupt. The different steps are simulated by Sim and Sim_P are as follows:

- We assume the parties securely generate the VRF setup string $\text{crs}_{\text{VRF}} \leftarrow \text{VRF.Setup}(1^\lambda)$.
- **Setup Phase.** Initiate key-registration by invoking $\mathcal{F}_{\text{VRaaS}}$ with (“Key-Register”). $\mathcal{F}_{\text{VRaaS}}$ forwards this request to $\mathcal{F}_{\text{SERV}}$ (controlled by Sim via Sim_P). Upon obtaining $(vk, \text{Verify}(\cdot))$ from corrupt server P , verify $\text{Verify}(\cdot)$ and then send it to $\mathcal{F}_{\text{VRaaS}}$. Invoke $\mathcal{F}_{\text{VRaaS}}$ with command (“Init-Ledger”, sid) to initiate the ledger. Upon receiving Req-Rand and Req-Fulf from server P check it against $\text{Verify}(\cdot)$ algorithm and forward it to $\mathcal{F}_{\text{VRaaS}}$.
- **Request-Randomness.** If the requester Q is honest then Sim performs nothing, the honest requester directly interacts with $\mathcal{F}_{\text{VRaaS}}$ without the involvement of Sim . If the requester Q is corrupt and the requester smart contract invokes $\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, \text{“Call Req-Rand”}, x)$, then Sim reads qID once the request transaction gets appended. Sim returns (qID, Q) to $\mathcal{F}_{\text{VRaaS}}$.
- **Fulfillment.** Sim simulates \mathcal{F}_{RLY} by verifying $(\text{REQ-RAND}, x; \text{qID})$ exists on LOG^{sid} and sets $\text{qINFO} := \text{qINP} := \text{qID}$. Sim sends (“Eval-Req”, qID, x) to $\mathcal{F}_{\text{VRaaS}}$ and sends qINP to server P . When $\mathcal{F}_{\text{VRaaS}}$ invokes Sim to obtain query input just return $(\text{qINP}, \text{qINFO})$. When $\mathcal{F}_{\text{VRaaS}}$ invokes Sim_P with input (“Eval”, vk, qINP, w) forward it to server P . When P returns (y, π) then Sim_P forwards it to $\mathcal{F}_{\text{VRaaS}}$. When $\mathcal{F}_{\text{VRaaS}}$ queries Sim (“Run Fulfillment?”) respond with Yes.
- **Local Verification.** To verify an output, Sim invokes $\mathcal{F}_{\text{VRaaS}}$ with it and returns whatever $\mathcal{F}_{\text{VRaaS}}$ outputs.

Indistinguishability Argument. We provide our hybrid argument as follows:

- **Hyb₀:** Real-world execution of the protocol.
- **Hyb₁:** Ideal world execution of the protocol. This is the same as the real-world execution of the protocol except if the corrupt server responds with an invalid (y, π) s.t. $\text{Verify}(vk, \text{qINP}, (y, \pi)) = 0$, then simulated \mathcal{F}_{RLY} forwards it to $\mathcal{F}_{\text{VRaaS}}$ without running Req-Fulf by itself. In contrast, in **Hyb₀**, functionality \mathcal{F}_{RLY} invokes Req-Fulf with it and the verification is performed by the smart contract function Req-Fulf . The adversary successfully distinguishes between the two hybrids if the smart contract in Req-Fulf fails to

detect that $\text{Verify}(vk, \text{qINP}, (y, \pi)) = 0$ and uploads an output transaction containing this invalid output where $\text{qINP} = (x, \text{reqID})$. Whereas, in the ideal world, $\mathcal{F}_{\text{VRaaS}}$ detects this and rejects the fulfillment Step 12. A distinguisher distinguishing between the two hybrids breaks the security of \mathcal{F}_{LED} since the smart contract behaves incorrectly.

Client Q is corrupt and server P is honest. The different steps are simulated by Sim as follows:

- We assume the parties securely generate the VRF setup string $\text{crs}_{\text{VRF}} \leftarrow \text{VRF.Setup}(1^\lambda)$ where Sim and Sim_P knows td .
- **Setup Phase.** Initiate key-registration by invoking $\mathcal{F}_{\text{VRaaS}}$ with (“Key-Register”). $\mathcal{F}_{\text{VRaaS}}$ forwards this request to $\mathcal{F}_{\text{SERV}}$, which is forwarded to Sim_P . Sim_P generates $(vk, sk) \leftarrow \text{VRF.SimGen}(1^\lambda, \text{td})$ by invoking the simulator of VRF and returns $(vk, \text{Verify}(\cdot))$ to $\mathcal{F}_{\text{VRaaS}}$. Invoke $\mathcal{F}_{\text{VRaaS}}$ with command (“Init-Ledger”, sid) to initiate the ledger. Generate Req-Rand and Req-Fulf and forward it to $\mathcal{F}_{\text{VRaaS}}$.
- **Request-Randomness.** Requester Q is corrupt. If the requester smart contract invokes $\mathcal{F}_{\text{LED}}.\text{Submit}(\text{sid}, \text{“Call Req-Rand”}, x)$, then Sim reads qID once the request transaction gets appended. Sim returns (qID, Q) to $\mathcal{F}_{\text{VRaaS}}$.
- **Fulfillment.** The simulator algorithm Sim simulates \mathcal{F}_{RLY} . Sim reads the request $(\text{REQ-RAND}, x; \text{qID})$ on LOG^{sid} and invokes $\mathcal{F}_{\text{VRaaS}}$ with the command (“Eval-Req”, qID, w). Sim sets $\text{qINFO} := \text{qINP} := \text{qID}$. When $\mathcal{F}_{\text{VRaaS}}$ invokes Sim to obtain query input just return $(\text{qINP}, \text{qINFO})$. $\mathcal{F}_{\text{VRaaS}}$ samples $y' \leftarrow \text{Rand}(vk, \text{rCtr}, w)$ where rCtr is the internal request counter of $\mathcal{F}_{\text{VRaaS}}$ corresponding to qID . When $\mathcal{F}_{\text{VRaaS}}$ invokes $\mathcal{F}_{\text{SERV}}$ (“Req-Proof”, vk, qINP, w, y), the request is forwarded to Sim_P . Sim_P computes simulated proof $\pi' \leftarrow \text{SimProve}(sk, \text{td}, y', \text{qINP}; \text{crs}_{\text{VRF}})$. Sim returns π' to $\mathcal{F}_{\text{VRaaS}}$. In the simulated protocol, Sim_P runs the honest VRaaS server algorithm on qINP with sk to obtain the output (y', π') and sends it to the simulated \mathcal{F}_{RLY} .

• **Local Verification.** To verify an output, Sim invokes $\mathcal{F}_{\text{VRaaS}}$ with it and returns whatever $\mathcal{F}_{\text{VRaaS}}$ outputs.

Indistinguishability Argument. We provide our hybrid argument as follows:

- **Hyb₀:** Real-world execution of the protocol.
- **Hyb₁:** Same as **Hyb₀**, except $\mathcal{F}_{\text{VRaaS}}$ fulfills the request for party Q by setting $y' \leftarrow \text{Rand}(vk, \text{rCtr}, w)$ and Sim_P simulates the proof $\pi' \leftarrow \text{SimProve}(sk, \text{td}, y', \text{qINP}; \text{crs}_{\text{VRF}})$. Indistinguishability follows due to the simulatability property of the VRF. It also guarantees that the output of $\pi_{\text{R-VRF}}$ is pseudorandom.

- **Hyb₂**: Ideal world execution of the protocol. This is the same as **Hyb₁**, except local verification is performed by running the local verification steps of $\mathcal{F}_{\text{VRaaS}}$ on it. If the output (y, π) on **qINP** is not registered in the memory of $\mathcal{F}_{\text{VRaaS}}$ then $\mathcal{F}_{\text{VRaaS}}$ sends \perp during the verification process in the ideal world execution. An adversary distinguishing between the two worlds forges an output (y, π) on **qINP** and gets it registered on \mathcal{F}_{LED} , corresponding to vk , by submitting a tuple of the form $(\text{REQ-FULF}, \text{qINFO}, y, \pi; \text{qID}, \text{qINP}, vk)$ on LOG^{sid} , where $\text{qINFO} := \text{qID} := \text{qINP}$. This means that the VRaaS server was never queried in the ideal world. Such an adversary forges the VRF output on **qINP**, given vk , and thus it breaks the unforgeability of the VRF primitive, leading to an attack on the simulatability of VRF.

Client is honest and server P is honest. The simulator performs nothing since there are no corruptions.

Client Q and server P are corrupt. Sim acts as a forwarder for messages between the corrupt server and the corrupt client by stimulating \mathcal{F}_{RLY} .

Indistinguishability Argument. The adversary successfully distinguishes between the two hybrids if the smart contract in **Req-Fulf** fails to detect that $\text{Verify}(vk, \text{qINP}, (y, \pi)) = 0$ and uploads an output transaction containing this invalid output where $\text{qINP} = (x, \text{reqId})$. Whereas in the ideal world, $\mathcal{F}_{\text{VRaaS}}$ detects this and rejects the output in step 12. A distinguisher distinguishing between the two hybrids breaks the security of \mathcal{F}_{LED} since the smart contract behaves incorrectly. \square

APPENDIX C

VALIDATION OF **qINFO** BY **Req-Fulf**

The fulfillment transaction **Req-Fulf** in $\pi_{\text{R-VRF}}$ is initiated by \mathcal{F}_{RLY} or the VRaaS server/committee in Chainlink [41]. When invoked with the input **qINFO**, it generates **qID** and **qINP** and verifies **qINP** against the output (y, π) using the verification key **VK**. The smart contract implementing $\mathcal{F}_{\text{VRaaS}}$ must ensure the validity of **qINFO** supplied to **Req-Fulf** w.r.t x and **qID**. Failure to do so could lead to undesirable consequences, such as sending the verifiable output to the wrong requester, resulting in the incorrect account being charged for the output’s generation cost as these are provided inside **qINFO**. To address this, we suggest the following approaches to validate that **qINFO** in **Req-Fulf** is correct w.r.t. the specific **qID**.

- 1) *On-Chain Storage of **qINFO***: On-chain storage **reqs**[] can be used to store the mapping between **qID** and x during the randomness request phase as $\text{reqs}[\text{qID}] := x$. During the fulfillment phase, **Req-Fulf** generates **qID** from **qINFO**, obtains $x \leftarrow \text{reqs}[\text{qID}]$ and matches it with **qINFO**.

- 2) *On-Chain Commitment to **qINFO***: On-chain storage is expensive so it is undesirable to store all of x on-chain. Instead, a hash of x could be stored as $\text{reqs}[\text{qID}] := H(x)$ where H is a public hash function. During fulfillment, **Req-Fulf** generates **qID** and x from **qINFO** and checks that $\text{reqs}[\text{qID}] = H(x)$.

- 3) *Implicit Validation by VRaaS server+Relay Nodes*: The above solutions either require the smart contract to have access to the blockchain or maintain on-chain storage which can be expensive. We note that in real-world $\mathcal{F}_{\text{SERV}}$ will be honest and we can reuse the trust assumption to validate **qINFO**. Recall in $\pi_{\text{R-VRF}}$ (Fig. 6) the server generates **qINP** and it is either trusted or it is implemented using an honest majority server committee. The protocol can be modified to perform the validation. The relay node modifies **qINP** to $\text{qINP}' = (\text{qINP}, H(\text{qINFO}))$ where **qINP** is generated from **qINFO**. The relay node committee also has an honest majority assumption and so **qINP'** is correctly generated. The server computes (y, π) on **qINP'**. When **Req-Fulf** is invoked with **qINFO**, it generates **qINP'** and checks (y, π) w.r.t. to the key **VK**. If (y, π) verifies then it is guaranteed that **qINP'** was indeed generated by the server and hence **qINFO** was validated correctly. We call this technique *implicit validation* and it is more efficient in practice since the validation process reuses the implicit trust assumption of the server and the relay committee.

APPENDIX D

ADDITIONAL RELATED WORKS

Formal Modeling of Ledgers and Blockchains.

The work of [37] introduced a general framework for distributed ledgers that captures both private and public ledgers in the iUC model [56], along with support for smart contracts. They prove that the Bitcoin blockchain, the Ouroboros family [57, 58] of blockchains, non-blockchain protocols of Corda [59] and Omniledger [60] implement their functionality. The work of [61] introduced the notion of blockchain-active adversaries, where the adversary can understand that it is being rewound in the proof by posting its state on the blockchain. In this model, they prove various impossibilities and show that achieving UC security is impossible for general circuits without setup assumption. Meanwhile, we consider the randomness service function and prove UC security in the programmable random oracle model. The recent work of [62] constructed a privacy-preserving smart contract based protocol in the model of [61] by relying on universally composable non-interactive zero knowledge [63]. The work of [64] considers modeling of privacy-preserving smart contracts in the Universal-Composability model of [18] and introduces the Kachina protocol for deploying private smart contracts. However, a simpler

ledger functionality suffices for us since privacy is not required from smart contracts, in the context of on-chain randomness service. Gearbox [19] provides a simple timed ledger functionality. It is compatible with sharding but doesn't support smart contracts.

APPENDIX E
VULNERABILITY IN DIA xRANDOM SMART
CONTRACT

The DIA xRandom randomness service operates on a beacon model, where requesters utilize the DRand [9] randomness beacon to generate random values. We found a bug in the code provided in their official documentation [16]. Given the DIA xRandom randomness service, two players (say P_1 and P_2) want to roll a dice and the player with the higher integer value on the roll wins the game. To do so each P_b ($b \in \{1, 2\}$) samples a `seedb` and “commits” to it on-chain by sending it to the smart contract. Note that the commitment does not hide the value of `seedb`. Once the two seeds are stored on-chain the DIA xRandom oracle is used to obtain randomness r . The final roll value of each player P_b is considered to be $\text{roll}_b := (r + \text{seed}_b) \% 6$ ($\%$ denoting the remainder). If $(\text{roll}_1 = \text{roll}_2)$ then it is a draw. Otherwise, player P_1 is considered the winner if $\text{roll}_1 > \text{roll}_2$ and player P_2 is considered the winner if $\text{roll}_2 > \text{roll}_1$. In the above game, a malicious player P_2 chooses $\text{seed}_2 := (\text{seed}_1 + 1) \% 6$ after seeing `seed1` in the commitment phase. P_2 always wins the game with probability $\frac{5}{6}$. This occurs since the roll_1 and roll_2 values are linearly correlated. We propose $\text{roll}_b := H(r, \text{seed}_b, b) \% 6$ to solve this issue as it breaks the correlation. We also find another vulnerability in their smart contract [65]. There is no commitment to the start of the Dice game or the end of player entry. Consider if two parties P_1 and P_2 , both honest, wish to play the Dice Game. Some third-party Eve can perform a denial-of-service attack by continuously calling either the function `commitPlayer1` or `commitPlayer2`, which will continuously update the value of `latestRoundId`. If done frequently enough, no call to `RollDice` can ever be successfully executed as the game never reaches a state where the beacon for `_round=latestRoundId + 10` has been published, and this stalls the game.