

STROBE: Streaming Threshold Random Beacons

Donald Beaver ✉

Independent Scholar, Pittsburgh, PA, USA

Mahimna Kelkar ✉

Cornell University, New York City, NY, USA

Kevin Lewi ✉

Meta Platforms, Inc., Menlo Park, CA, USA

Valeria Nikolaenko ✉

a16z crypto, Palo Alto, CA, USA

Alberto Sonnino ✉

Mysten Labs, London, UK

University College London, UK

Konstantinos Chalkias ✉

Mysten Labs, Palo Alto, CA, USA

Lefteris Kokoris-Kogias ✉

Mysten Labs, London, UK

IST Austria, Klosterneuburg, Austria

Ladi de Naurois ✉

Washington DC, USA

Arnab Roy ✉

Mysten Labs, Palo Alto, CA, USA

Abstract

We revisit decentralized random beacons with a focus on practical distributed applications. Decentralized random beacons (Beaver and So, Eurocrypt'93) provide the functionality for n parties to generate an unpredictable sequence of bits in a way that cannot be biased, which is useful for any decentralized protocol requiring trusted randomness.

Existing beacon constructions are highly inefficient in practical settings where protocol parties need to rejoin after crashes or disconnections, and more significantly where smart contracts may rely on arbitrary index points in high-volume streams. For this, we introduce a new notion of **history-generating decentralized random beacons** (HGDRBs).

Roughly, the *history-generation* property of HGDRBs allows for previous beacon outputs to be efficiently generated knowing only the current value and the public key. At application layers, history-generation supports registering a sparser set of on-chain values if desired, so that apps like lotteries can utilize on-chain values without incurring high-frequency costs, enjoying all the benefits of DRBs implemented off-chain or with decoupled, special-purpose chains. Unlike rollups, HG is tailored specifically to recovering and verifying pseudorandom bit sequences and thus enjoys unique optimizations investigated in this work.

We introduce STROBE: an efficient HGDRB construction which generalizes the original squaring-based RSA approach of Beaver and So. STROBE enjoys several useful properties that make it suited for practical applications that use beacons:

1. **history-generating**: it can regenerate and verify high-throughput beacon streams, supporting sparse (thus cost-effective) ledger entries;
2. **concisely self-verifying**: NIZK-free, with state and validation employing a single ring element;
3. **eco-friendly**: stake-based rather than work based;
4. **unbounded**: refresh-free, addressing limitations of Beaver and So;
5. **delay-free**: results are immediately available.
6. **storage-efficient**: the last beacon suffices to derive all past outputs, thus $O(1)$ storage requirements for nodes serving the whole history.

2012 ACM Subject Classification Security and privacy → Public key encryption

Keywords and phrases decentralized randomness, beacons, consensus, blockchain, lottery

Digital Object Identifier 10.4230/LIPIcs.AFT.2023.7

Related Version *Full Version*: <https://eprint.iacr.org/2021/1643>

Acknowledgements Work done when all the authors were at Novi Research, Meta.



© Donald Beaver, Konstantinos Chalkias, Mahimna Kelkar, Lefteris Kokoris-Kogias, Kevin Lewi, Ladi de Naurois, Valeria Nikolaenko, Arnab Roy, and Alberto Sonnino;
licensed under Creative Commons License CC-BY 4.0

5th Conference on Advances in Financial Technologies (AFT 2023).

Editors: Joseph Bonneau and S. Matthew Weinberg; Article No. 7; pp. 7:1–7:16

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

A random beacon is a shared source of agreed-upon random bits. First introduced in 1981 by Rabin [44] in the context of digitally-signed documents, beacons use unpredictability to put adversarial strategies in doubt. Trusted beacons are increasingly recognized as a critical resource for Decentralized Finance (DeFi), blockchains, Byzantine Fault Tolerance (BFT), leader elections, and a range of Decentralized Applications (dApps) including lotteries.

At fundamental protocol layers such as consensus, beacons drastically reduce the time and effort needed to withstand protocol-defeating attacks. For example, Feldman and Micali [30] implemented a common coin to achieve 2-round probabilistic consensus, breaking the deterministic lower bound of $f + 1$ rounds [26]. As another example, most proof-of-stake blockchains use randomness to select leaders proposing blocks and, for increased efficiency, also use randomness to select subcommittees to confirm the blocks. To make sure the malicious actors have negligibly small window to corrupt the leaders or the subcommittees, it should be impossible for them to predict or bias the randomness used in the process, which creates strong demand for good randomness beacon protocols. Moreover, for the light-clients of these blockchains to synchronize the state, they need to verify the selection of leaders and subcommittees, implying that they have to verify the random beacon outputs.

At higher levels of abstraction, new and important properties and optimizations emerge. Crash failures, delays, and asynchrony in the interactions between processes or the execution of smart contracts in a ledger environment make it essential to have stronger and coordinated record-keeping, to make it easier to retrieve and verify past results.

To complicate matters, many high-level applications like lotteries or gaming require high-frequency streams. Frequencies and latencies at the scale of seconds are too slow. Registering streams bit-by-bit in real-time on a ledger is simply infeasible given transaction throughput and latencies, let alone enormously cost-prohibitive fees. Registering intermittent results is a feasible and cheaper workaround - as long as the intermittent results provide sufficient content and validation to enable optimistic fault detection and recourse for any on-chain relying parties. To date, such applications either run at very slow speeds (lotteries) or require centralized trust (gaming platforms).

While a naive approach might register a long sequence accompanied by aggregate signatures and rollups to prove step-by-step correctness, so that relying parties can check content and validity, we take a novel and alternative approach, employing **history generation** for streamlined communication size and speed. Our protocols are direct and simple, providing a concise element to *regenerate* the entire back sequence (effectively compressing it) and to *validate* it (obviating NIZKs, let alone aggregation).

Decentralized random beacons. Beaver and So presented the first decentralized beacon (DRB) [4], achieved by way of a threshold homomorphic secret sharing of future bits. Their approach capitalized on the Blum-Blum-Shub pseudorandom generator [7], in which successive squares in an RSA ring are produced. By reversing the sequence and employing tricks for threshold Lagrange interpolation, [4] gave a *self-certifying sequence* obviating any need for ZKPs.

History-generating DRBs. A key contribution of this work is the notion and implementation of *history generation*, a technique that is essentially equivalent to compressing long sequences and providing a maximally concise validity check. In the context of random streams produced by DRBs, we are able to take strong advantage of the sequential nature of the stream to achieve optimizations exceeding those available to general-purpose calculations-with-ZKPs.

Unbounded sequences. In the threshold homomorphic VSS approach of [4], there is a bound on the number of bits to be produced without engaging in some kind of refreshing state. The results are limited by stakeholders rather than VDF-style delays: there is no explicit notion of using work as a computationally-guaranteed limit on the cadence of a beacon. We address both limitations using a generalized construction without incurring proof-of-work.

Contributions. STROBE is a decentralized random beacon providing the following properties:

- **History-generating:** It can regenerate and verify high-throughput beacon streams, supporting sparse (thus cost-effective) ledger entries;
- **Concisely self-verifying:** It is NIZK-free, with state and validation employing a single ring element;
- **Eco-friendly:** It is stake-based rather than work-based;
- **Unbounded:** It is refresh-free, addressing limitations of Beaver and So;
- **Delay-free:** Results are immediately available.
- **Storage-efficient:** $O(1)$ storage for nodes serving the whole beacon history.

Roadmap. §2 describes intuitions of our constructions and how applications can benefit from the novel, history-generating property. We cover related work in §3, with particular attention to feature tradeoffs in Table 1. Details of the construction appear in §4, with security model in §5 and proofs in §6. We describe some extension in §7 and application details in §8.

2 Construction Intuition

Our starting point is to view RSA decryption as a trapdoor one-way function in reverse, which can be efficiently verified. The beacon output of an epoch is essentially the RSA decryption of the previous epoch’s output, and this carries on perpetually. The verification of an output is to just RSA encrypt it with the public key and see that it equals to the previous epoch’s output. In this sense, the beacon is *self-certifying*.

Step 1: Beacon setup. Viewed in a non-distributed setting, the setup of the beacon generates RSA modulus $N = pq$ along with a root-ing parameter s . Then the beacon proceeds perpetually as:

$$x \rightarrow x^{1/s} \rightarrow x^{1/s^2} \rightarrow \dots \rightarrow x^{1/s^T} \rightarrow \dots$$

To parties that just know about N as a public parameter, this provides some attractive properties: (1) The next value in the beacon is hard to predict given the earlier values. (2) It’s easy to verify a beacon value against the last value. In fact, we can check the value against any historical value, except that it gets progressively harder with the gap. (3) An especially tantalizing property is that any historical beacon value, can in fact be simply *generated* from the knowledge of the current value.

Step 2: Removing the trusted party. Of course, the problem with this is that taking roots in an RSA ring is hard without knowing the prime factors p and q . So we need a trusted party holding the primes to be generating all the beacon values. Recent advances

7:4 STROBE: Streaming Threshold Random Beacons

in distributed RSA modulus generation allows us to generate a public modulus N , with no single party knowing the factors. Imagine that in addition we also give secret shares of $s^{-1} \pmod{\phi(N)}$ to n distinct parties:

$$sk_1 + sk_2 + \dots + sk_n = s^{-1} \pmod{\phi(N)}$$

At the time epoch $T + 1$, the n parties can then output $x_T^{sk_i}$ each, where x_T is the output of the last epoch. On multiplying all the public shares, we get $x_{T+1} = x_T^{s^{-1}} \pmod{N}$. At the same time, observe that even $(n - 1)$ of the secret keys are effectively independently random. Also observe every public share is also self-certifying wrt the last epoch, in the sense $x_T^{sk_i} = (x_{T+1}^{sk_i})^s$.

Step 3: Adapting to the threshold setting. Adapting the n -of- n setting to a threshold t -of- n setting introduces additional challenges which do not affect known-order groups. Essentially, Shamir secret sharing involves fractional Lagrange interpolation coefficients which are efficient to compute to group elements if we know the order. However, this is not possible to do in the exponent of RSA group elements, as $\phi(N)$ is not public. We adapt and extend the techniques pioneered by [4] and also used by [48] to address this challenge. The core trick is to lift the Lagrange coefficients by a factor of $n!$, so that they are not fractional anymore - details in Section 4.

3 Related Work

Random Beacons. Practical bias-resistant random beacons producing regular series of random outputs typically follow one of three approaches: the first one uses publicly verifiable secret sharing (PVSS) mechanisms; the second uses verifiable random functions (VRFs), often in conjunction with threshold cryptography, and the third relies on verifiable delay functions (VDFs).

Beacons of the first type use the following blueprint design, they need 4 rounds for n participants to generate a random value. In the first round all nodes simultaneously secret-share a freshly generated random value s among the rest of the nodes. They do so by publishing the following values: n shares of s each encrypted under receiving party's public key, a commitment to the secret, and a non-interactive zero-knowledge proof that those were generated correctly. In the second round, the parties run some sort of consensus algorithm to agree which nodes did the sharing correctly and which failed. This can also be replaced with posting the shares on chain where anyone can verify the proofs independently, although this typically relies on consensus provided by a chain (rather than enabled by the beacon). In the third round, the parties reveal their secrets and in the fourth round for parties that withheld their secrets, shares of those secrets are broadcast for recovery. The resulting beacon's value is derived from the revealed or recovered secrets. The rounds can be pipelined to get regular random outputs at each round. Starting from the original proposal of Ouroboros [36] a sequence of papers (Scrape [17], Albatross [18], HydRand [47], RandHerd [49]), OptRand [6] has improved the communication, computation complexity per node and the verification complexity for beacon of this type. No centralized or trusted setup and standard cryptographic assumptions are the main advantages of those protocols. However, those protocols still remain communication-intensive, since they need to run the full protocol (including consensus) for every fresh random value and computation wise intensive, since every party needs to check that every secret-share has been correctly constructed. A

public verifier needs $O(n)$ messages to verify each beacon's value, making it communication-wise expensive to verify a series of random values. A full chain of our beacon, in contrast, can be verified using only the current beacon's output and the public parameters.

Beacons of the second type rely on a setup phase where a secret key is generated in a distributed manner, after which homomorphic VSS and/or threshold signatures can be employed. The use of homomorphic VSS was pioneered in the first DRB in 1993 [4], where successive values of a BBS generator are revealed. The shares of square roots are, homomorphically and similar in spirit to threshold RSA, square roots of shares. The shares and the reconstructed values each act as verifiers of previous values (*viz.* by way of squaring). This particular solution suffers from a need to refresh to new sequence values periodically; it is not unbounded as-is.

In numerous other generalized approaches of this second style, nodes can use a unique threshold signature scheme, such as threshold BLS [10, 11] (tBLS) in Dfinity [34] and drand [28], or as threshold RSA (tRSA) signatures [48] in Cachin et al. [15]) to sign an agreed-upon progression of values (either block-hashes, or round-numbers, or a combination of those). More generally, a verifiable random function is built in these construction to which unique signatures are a particular instantiation, but other constructions are also possible, e.g. NSEC5 VRF [41] used by Chainlink. The main advantage of those protocols is efficiency (each party only sends a single message per beacon's value) and ease of public verifiability (current beacon's value can be verified using only public parameters). The disadvantages are a complicated setup phase (though straightforward when a trusted party is assumed), for example to generate the threshold key for tBLS over the internet it still takes elaborate protocols with $O(n^4)$ communication complexity costs [35, 38] and $O(f)$ worst case run time. Alternatively, instead of a threshold key, the parties can use independently generated keys (as is currently being done in the Ethereum's RANDAO approach [12]), although unfortunately such beacons have some small degree of bias from the parties who may decide to withhold their messages, we therefore do not focus on such beacons here despite the fact that certain application (e.g. Ethereum's committees selection) may stay resilient to small bias of the beacon. We improve on the approach of tRSA by building a random beacon straight from the RSA assumption requiring no additional proofs for correctness of signature's shares. Our scheme naturally gives a novel property of history generation in contrast to existing approaches.

There is also a third approach that relies on a proof-of-delay mechanism [8]: either a block-hash is passed through a verifiable delay function (VDF), or the values submitted by the participants are passed through a VDF function to generate a random value. This approach makes it simpler to build unbiased beacons, as the participants will not have enough time to see how the bias on their contributions would affect the resulting value of the beacon. Most prominent systems are RANDAO w. VDF [27], continuous VDF [29] and RandRunner [46]. But those approaches are highly computationally intensive for the prover and require precise estimates of concrete complexity, which are hard to predict in practice (competitions with high-reward incentives were set-up by Chia Networks (chia.net) and Ethereum (ethereum.org) in partnership with Protocol Labs (protocol.ai) to get concrete estimates of VDFs' complexity). In the presence of a quantum adversary a quantum-resistant VDF could be used to produce a quantum-resistant random beacon, e.g. Veedo [50].

Recent surveys [22, 45] systematize knowledge of randomness beacons and discuss the complexity and practicality of the constructions in more details. In Table 1 we compare selected beacons to STROBE which is the only one that provides the property of history generation. Self-certifying beacons allow to inexpensively verify a beacon value against the

7:6 STROBE: Streaming Threshold Random Beacons

previous one; refresh-free beacons allow for generation of indefinite sequence of random values per single setup; BA-free beacons do not use Byzantine-Agreement protocols. Beacons based on the RSA assumption require a trusted setup or a distributed RSA modulus generation, our protocol additionally requires the generation of an inverse of a public exponent. The VDF-based beacons in RSA groups also do require a trusted setup.

■ **Table 1** Comparison among several beacon protocols.

	History gen	Self certifying	Single round	Refresh-free	BA-free	Setup-free	
Albatross [18], HydRand [47], OptRand [6], RandHerd [49]	✗	✗	✗	✓	✓	✓	PVSS
RandRunner [46], RANDAO++ [27], cVDF [29], Veedo [50]	✗	✗	✓	✓	✓	✗/✓	VDF
Dfinity [34], drand [28]	✗	✗	✓	✓	✓	DKG	tVRF
C03 [13], BS93 [4]	✗	✓	✓	✗	✓	mod	RSA
STROBE (this work)	✓	✓	✓	✓	✓	mod+ inv	RSA

Distributed generation of an RSA modulus and an inverse of a public exponent. The setup of our construction requires the generation of an RSA modulus N that is a product of two primes. The most common way of generating such a modulus in a centralized setting (a.k.a. with a trusted setup), is to randomly sample κ -bits integers running Miller-Rabin probabilistic primality tests on them [40, 43] until two primes are obtained with overwhelming probability, multiplying them gives a 2κ -bits bi-prime N . Since currently there is no known way to sample a bi-prime (using only public randomness) for which nobody knows the factors, the only way to alleviate the trusted setup is to distribute the generation of the bi-prime modulus via a dedicated multi-party computation protocol. Boneh and Franklin [9] initiated the study of distributed RSA modulus generation devising a protocol in a passive security model with honest majority. The follow-up protocol of Algesheimer, Camenisch and Shoup [1] devised a protocol for generation of N that is a product of two *safe* primes, passively secure with honest majority. The follow-up work has hardened the original Boneh-Franklin’s protocol to be secure in the presence of actively malicious parties and honest majority [31]. The works mentioned above also generate an inverse (RSA decryption key) in a distributed manner. An improvement to this part of the protocol was also made by Catalano et al. [19]. A promising approach of getting rid of the setup phase altogether, that was proposed in RandRunner [46] and in the work of Damgård and Koprowski [24], can potentially get applied to this work.

4 The STROBE Protocol

We now define the syntax of a History Generating Decentralized Random Beacon (HGDRB) and describe our STROBE construction.

► **Definition 1** (HGDRB). *A History Generating Decentralized Random Beacon (HGDRB) is a set of algorithms (Setup, Gen, Eval, VerifyShare, Combine, Verify, Back):*

Setup: $(\lambda, n, t) \rightarrow (pk, sk_1, \dots, sk_n)$. The Setup algorithm takes the security parameter λ and threshold parameters n and t . The scheme allows t -of- n reconstruction, with secret shares given to n parties. The output is a public key pk and secret shares sk_1, \dots, sk_n .

Gen: $pk \rightarrow x_0$. The (one-time) Gen algorithm samples an initial random value x_0 .

Eval: $(sk_i, x_T) \rightarrow x_{T+1,i}$. Each party takes the last epoch's (T) output x_T , computes and outputs a share of the next epoch's output $x_{T+1,i}$.

VerifyShare: $(pk, x_{T,i}, x_{T+1,i}) \rightarrow \{0, 1\}$. The VerifyShare algorithm checks the epoch $T+1$ shares against the corresponding epoch T shares.

Combine: $(pk, x_{T,P_1}, x_{T,P_2}, \dots, x_{T,P_t}) \rightarrow x_{T+1}$. Given t shares from epoch T that pass VerifyShare the Combine algorithm outputs the next epoch's beacon value x_{T+1} .

Verify: $(pk, x_T, x_{T+1}) \rightarrow \{0, 1\}$. The Verify algorithm checks the epoch $T+1$ beacon output against the epoch T beacon output.

Back: $(pk, x_T, k) \rightarrow x_{T-k}$: This outputs the beacon value at epoch $T - k$ given the epoch T beacon value x_T and $k < T$.

Informally, correctness asserts that honestly computed beacon values will pass verify checks with respect to previous beacon outputs. The same should hold for share outputs as well. The *Back* function allows to compute any historical beacon value efficiently (going back a polynomial number of epochs).

The STROBE protocol

We now provide our HGDRB construction, called *STROBE*. It is based on threshold inversion in RSA groups and its security follows from the RSA assumption.

Setup (λ, n, t) : The Setup algorithm takes the security parameter λ and samples an RSA modulus $N = pq$ with $\phi(N) = 4p'q'$.

Sample primes p, q such that $p - 1 = 2p', q - 1 = 2q'$ with p', q' also being primes. Pick a prime s , s.t. $\min(p', q') > s > n$. Let $N = pq$ and observe that $s \nmid \phi(N)$. Sample $a_1, \dots, a_{t-1} \leftarrow [1, N]$ and let $f(X) = v + a_1X + \dots + a_{t-1}X^{t-1}$, where $v = (n!s)^{-1} \pmod{p'q'}$.

Send secret shares $sk_i = f(i) \pmod{\phi(N)}$ to parties $i \in [1, n]$. Publish $pk = (N, s)$.

Gen (pk) : The Gen algorithm samples a seed value $\text{seed} \leftarrow [1, N]$. This is a public random value that can be computed by MPC or by taking a block hash. It then outputs $x_0 = \text{seed}^{4(n!)^2} \pmod{N}$.

The $4(n!)^2$ factor is an artifact of the security proof and will be explained in Section 6.

Eval (sk_i, x_T) : Each party takes the last epoch's (T) output x_T , computes and outputs a share of the next epoch's output $x_{T+1,i}$:

$$x_{T+1,i} = x_T^{sk_i} = x_T^{f(i)} \pmod{N}.$$

VerifyShare $(pk, x_{T,i}, x_{T+1,i})$: The epoch $T + 1$ shares can be self-verified against the corresponding epoch T shares, by checking that

$$x_{T+1,i}^s = x_{T,i} \pmod{N}.$$

Combine $(pk, x_{T,P_1}, x_{T,P_2}, \dots, x_{T,P_t})$: Given t epoch T shares, first check that each of them pass VerifyShare. Let γ denote the set of these indices $\{P_1, \dots, P_t\}$. The Combine algorithm then computes the combined epoch T beacon value x_T by computing the interpolation:

$$x_T = \prod_{i \in \gamma} x_{T,i}^{n!L_i(0)} \pmod{N},$$

7:8 STROBE: Streaming Threshold Random Beacons

where the Lagrange basis polynomials L_i 's are defined as:

$$L_i(X) = \prod_{j \in \gamma, j \neq i} \frac{X - j}{i - j}.$$

The polynomials satisfy the following property: for $\forall i, j \in \gamma : L_i(i) = 1$, and $L_i(j) = 0$ for $i \neq j$. The $(n!)$ factor is essential to clear the denominators of the interpolation coefficients, which makes sure there is no fractional exponent to compute.

Verify(pk, x_T, x_{T+1}): The epoch $T + 1$ beacon output can be self-verified against the corresponding epoch T beacon output, by checking that

$$x_{T+1}^s = x_T \pmod{N}.$$

Back(pk, x_T, k): This outputs the beacon value at epoch $T - k$ as:

$$x_{T-k} = x_T^{s^k} \pmod{N}.$$

We assume that each share carries the identity information (metadata) about which party generated it. Apart from the *Verify* (and *VerifyShare*) algorithms checking against the last epoch outputs, the *Back* algorithm provides an alternate way to check against any previous epoch output, all the way to epoch 0 (and 1). In fact if the shares from epoch 1 are also made part of the trusted pk , then this provides a way of checking the beacon value without accessing the past beacon values at all. The trade-off is that checking against beacon values in the past grows computationally expensive with the number of epochs elapsed. As an extension of the core protocol, we can also leverage techniques from popular VDF constructions, as described in Section 7. We now give the proof of correctness of the protocol.

4.1 Proof of Correctness of STROBE

We have $x_0 = \text{seed}^{4(n!)^2}$. Assume inductively, $x_{T-1} = \text{seed}^{4(n!)^2 s^{-T+1}}$. Correctness of the beacon outputs follows as below:

$$\begin{aligned} x_T &= \prod_{i \in \gamma} x_{T,i}^{n! L_i(0)} = x_{T-1}^{n! \sum_{i \in \gamma} f(i) L_i(0)} = \text{seed}^{4(n!)^2 s^{-T+1} n! \sum_{i \in \gamma} f(i) L_i(0)} \\ &= \text{seed}^{4(n!)^2 s^{-T+1} n! f(0)} = \text{seed}^{(n!)^2 s^{-T} 4(n!) v s} \end{aligned}$$

As $v = (n!s)^{-1} \pmod{p'q'}$, we have $4(n!)vs = 4 \pmod{\phi(N)}$. Therefore, $\text{seed}^{4(n!)vs} = \text{seed}^4 \pmod{N}$.

Substituting, we get:

$$x_T = \text{seed}^{4(n!)^2 s^{-T}}.$$

This carries the induction successfully forward, and also leads to successful verification: $x_T^s = x_{T-1}$. Similar steps apply to the individual shares as well.

5 Security Model

There are various flavors of security that we could require of a random beacon. To narrow down the syntax, we will focus on *stake-based*, *self-certifying*, *threshold* beacons. The baseline security we want is that an adversary should not be able to predict future beacon values based on seeing past values and corrupting less than a threshold number of participants.

Unpredictability vs. Pseudorandomness

We could require the next beacon value to be pseudorandom, instead of just unpredictable. We observe that we could essentially compile an unpredictable beacon into a pseudorandom one, either by applying a random oracle (similarly to what is described as “tick-tock” in [29]), or if we want to avoid the RO assumption, by extracting hardcore bit(s), as in [7]. There exist applications where unpredictability suffices [21], but in most of the cases, such as a decentralized lottery or leader election, unbiased randomness is essential [45]. We also note that, for a self-certifying and/or history generating beacon, there needs to be a part of the beacon output that cannot be pseudorandom, as otherwise it cannot be used for the certification assertion and/or historical value computation.

Active vs. Passive

A passive adversary just observes the transcript of an honest run of the protocol, including public shares and beacon outputs and then tries to predict the next beacon value. An active adversary, on the other hand, can actively modify the public share values and beacon outputs. In this paper we consider active adversaries. In our setting, the self-certification essentially ensures that there is only a unique value for each expected public share or beacon value that can pass onto the next phase. This makes any adversarial modifications immediately noticeable and subject to rejection.

Selective vs. Adaptive

Another dimension to specify the adversary is on whether we restrict it to corrupt parties that it declares upfront (selective), or allow the parties to be corrupted dynamically as it observes and interacts with the protocol (adaptive). Our core protocol only satisfies selective security in this respect. We leave it as an open problem to construct an adaptively secure protocol which retains the efficiency standard of our selective one. A generic (complexity-leveraging) approach can be used to get an adaptively secure scheme from a selectively secure one: the reduction simply needs to guess the set of corrupted parties and then run the selective reduction aborting if the selection of corrupted parties does not match the one requested by the adversary. This, unfortunately, leads to a loss in security that is exponential in the number of parties, n , thus limiting the number of parties to be at most logarithmic in the security parameter λ . There are other promising approaches in the works for threshold RSA cryptosystems. Canetti et al. [16] proposed a methodology for transforming a selectively-secure threshold scheme into an adaptively-secure one, where the protocol needs to be modified to carefully erase secrets and to use simple zero-knowledge proofs, the adversary is rewinded in the proof which also incurs a security loss although not as large as with the complexity-leveraging approach. A follow-up work of Almansa et al. [2] simplified this result for RSA, but at the cost of the secret’s re-sharing after every round, with an emergent benefit of making the scheme proactively secure. In a proactively secure scheme the adversary can corrupt at most t players in a time period determined by the protocol. Since the set of corrupted parties changes, each party can become corrupt at some point (i.e. leak its secrets), but if the party recovers from a compromise then a subsequent secret’s re-sharing will enable the party to be honest again.

Given the above discussion, we now formally define the security model that we consider for our core protocol: unpredictable, passive and selective.

7:10 STROBE: Streaming Threshold Random Beacons

► **Definition 2** (Selective-secure Unpredictability). *We say that an HGDRB is Selective-secure Unpredictable if the following adversary has negligible advantage:*

1. *The Challenger runs $Setup(\lambda, n, t)$ and outputs pk to the Adversary.*
2. *The Adversary selects a time epoch $S < poly(\lambda)$ and a set of parties $\gamma = \{P_1, P_2, \dots, P_{t-1}\}$ to corrupt.*
3. *The Challenger sends the transcript of the protocol till time epoch S to the Adversary, as well as the secret shares for parties in γ . This includes the outputs of Gen and those of $Eval$ for all $i \in [1, n]$ and $T \leq S$.*
4. *The Adversary outputs a quantity x' .*
5. *The Adversary wins if $Verify(pk, x_S, x')$ passes.*

In the next section, we formally prove that our construction satisfies Selective-Secure Unpredictability. We also claim that the construction protects against active adversaries, as the self-verification of beacon and share values ensure that only the unique correct values would not be rejected.

6 Proof of Security

In this section, we show that the STROBE protocol satisfies Selective-secure Unpredictability under the RSA assumption.

► **Definition 3** (RSA Assumption). *The Challenger samples RSA number $N = pq$ and picks a quantity s co-prime to $\phi(N)$. Then it randomly samples $z \leftarrow [1, N]$ and sends (N, s, z) to the Adversary. The Adversary outputs y . The RSA assumption states that the probability of $y^s = z \pmod{N}$ is negligible.*

► **Theorem 4** (Security). *The STROBE protocol is Selective-secure Unpredictable under the RSA Assumption.*

Proof. Let (N, s, z) be an RSA challenge for a prime $s > n$.

Setup. The Challenger outputs $pk = (N, s)$. Suppose the Adversary picks an epoch S and corrupts parties in the set $\gamma = \{P_1, \dots, P_{t-1}\}$. Sample $b_{P_1}, \dots, b_{P_{t-1}} \leftarrow [1, N]$. Send Adversary shares $sk_i = b_i$, for all $i \in \gamma$.

Eval. Consider an implicit polynomial $f(X)$, such that $f(0) = v$, where $v = (n!s)^{-1} \pmod{p'q'}$, and $f(i) = b_i$, for all $i \in \gamma$:

$$f(X) = vL_0(X) + \sum_{i \in \gamma} b_i L_i(X).$$

Here $L_i(X)$ are Lagrange basis polynomials of degree $t - 1$ each:

$$L_i(X) = \prod_{j \in \{\gamma \cup \{0\}\}, j \neq i} \frac{X - j}{i - j}.$$

The polynomials satisfy the following property: for $\forall i, j \in (\gamma \cup \{0\}) : L_i(i) = 1$, and $L_i(j) = 0$ for $i \neq j$. Note that it is possible to evaluate $z^{n!L_i(j)}$ for any $i, j \in [0, n]$, since the $(n!)$ factor eliminates the denominators of interpolation polynomials, making it possible to evaluate the exponentiation.

Set $x_0 = z^{4(n!)^2 s^S}$ and $x_{T,j} = z^{4f(j)(n!)^2 s^{S-T+1}}$. Now, for $j \in \gamma$, we can compute $x_{T,j}$ explicitly based on $f(j) = b_j$. We now show that for $j \notin \gamma$, we can compute $x_{T,j}$ without explicitly computing $f(j)$. Observe that $4v = 4(n!s)^{-1} \pmod{\phi(N)}$, therefore $z^{v4n!s} = z^4$ hence $x_{T,j}$ can be explicitly constructed as follows:

$$x_{T,j} = z^{4f(j)(n!)^2 s^{S-T+1}} = z^{(f(j)4n!s) \cdot (n!)s^{S-T}} = z^{(4L_0(j) + \sum_{i \in \gamma} b_i \cdot (4n!s) \cdot L_i(j)) \cdot (n!)s^{S-T}}$$

Note that the last $(n!)$ factor is essential to make sure we can clear the denominator of $L_0(j)$ and thus avoid any divisions in the exponent. Send x_0 and the $x_{T,j}$'s for all $T \in [1, S]$ and $j \in [1, n]$ to the adversary.

RSA response. Let's suppose the adversary comes up with x' as the next epoch candidate. Given the self-certification checks upto time epoch T , we inductively have $x_k = x_0^{s^{-k}} = z^{4(n!)^2 s^{S-k}}$, for $k \in [0, S]$. If the adversary response x' passes verification, then we should have $x'^s = x_S = z^{4(n!)^2}$. Let $w = (4(n!)^2)^{-1} \pmod{s}$ and let $w(4(n!)^2) = 1 + ks$ for some computable integer k . Then $x'^{sw} = z^{4(n!)^2 w} = z^{1+ks}$. Therefore, $z = (x'^w z^{-k})^s$. Hence $x'^w z^{-k}$ will be a winning response to the RSA challenge.

Distributions. Finally, we observe that as s^S is invertible modulo $\phi(N)$ and sampling uniformly from $[1, N]$ and $[1, \phi(N)]$ are statistically indistinguishable, the distribution of x_0 in the STROBE construction and this proof are statistically indistinguishable. Matching these distributions is the technical reason behind the extra $4(n!)^2$ factor in the *Gen* algorithm. ◀

On active adversaries

An active adversary is allowed to tamper with the intermediate shares and beacon values. Observe that, as s is invertible with respect to $\phi(N)$, the beacon values are deterministic in the past as well as the future, given a current value. Therefore, given the initial value x_0 , an adversarially tampered beacon value would be verifiably detected.

It is possible that intermediate shares could be tampered with without detection in certain circumstances. For example, if one of the parties is absent in all the epochs up to a certain point and then starts participating. The upshot of such attacks is essentially denial of service. The correct verifiability of beacon outputs is not affected.

7 Extensions

In this section we discuss some extensions of our core scheme to enable further functionalities and guarantees.

7.1 Dynamic Beacon Committees

STROBE can also easily handle dynamically changing the participants executing the beacon protocol. For example, this could involve rotating to a newly chosen committee (even of a different size) of parties after some fixed number of epochs. The new committee will then continue to generate future outputs of the beacon. In fact, the new committee can also be chosen based on the output of the distributed beacon. Dynamic participation can be accommodated by using the old committee to reshare the secret key to the new committee through existing literature on dynamic proactive secret sharing [5, 39]. Note however, that such protocols need to make the assumption that honest parties from previous committees

delete their shares so that an adversary cannot recover the secret key even by corrupting more than a threshold number of parties from a previous committee. The above scenario is also referred to as a “long range” attack (c.f., [21]).

7.2 Succinct Proofs of Beacon Validation

In earlier sections we observed that it is possible to check the beacon value x_T at epoch T against the seed value x_0 by checking $x_0 = x_T^{s^T}$. However, this takes sequential time T . One way to speed up verification is to exploit the RSA repeated powering structure of this check and use existing techniques like [42, 51] to add a proof in addition to the beacon value.

In particular, we can just apply Wesolowski’s [51] proof technique in reverse and produce the proof $\pi = x_T^{\lfloor s^T/\ell \rfloor}$, where ℓ is the result (a prime number) of a random oracle H applied to (x_0, x_T, T) . The verifier computes ℓ and r as the remainder on dividing s^T by ℓ and then checks $x_0 = \pi^\ell x_T^r$. However, a drawback of this method is that producing the proof takes about T time as well, which may not be ideal to do every epoch.

A better approach is to use a continuous VDF [29]. In a continuous VDF, it is efficient to publish and use intermediate proofs at every time epoch. The Ephraim et al [29] continuous VDF makes use of the recursive structure of Pietrzak’s VDF [42]. The high level idea is that they checkpoint a logarithmic number of past values and keep recursively merging an appropriate number of them in order to prevent growing the proof size. We can also similarly checkpoint and merge based on a similar recursion, while reversing the order: $x_0 = x_u^{s^{T/2}}$ and $x_u = x_T^{s^{T/2}}$, where $u = T/2$ is the midpoint.

We can use a continuous VDF at defined intervals as well, instead of every epoch. A verifier can sequentially compute the expected value until the last such interval and then just check the VDF proof.

7.3 Deeper Blockchain Integration

The design described in Section 4 relies on external authorities. However, incorporating STROBE within the infrastructure of a blockchain would enable the generation of randomness as a side effect of the normal system operations, taking no additional dependency on external authorities. Allowing efficient reconfiguration of STROBE for permissionless environments with highly dynamic set of nodes remains an open problem. Nevertheless, integration of STROBE into permissioned (BFT-based) blockchain platforms is straightforward. For instance, in Hyperledger Fabric [14], contracts run on private sets of computation nodes and use the Fabric protocols for cross-contract calls. In this setting, STROBE authorities can coincide with the Fabric smart contract authorities. Upon a contract setup, they perform a setup and key distribution, and then start generating randomness when authorized by the contract. For generating randomness, the only secrets maintained are the private STROBE authorities keys; all other operations of the contract can be logged and publicly verified. The threshold trust assumption – namely that of integrity and availability – is preserved.

Incentives for Randomness Generation

One open question is on managing the incentives of releasing shares for the randomness beacon. On a first sight the fault-tolerance of STROBE makes the problem a public goods game, since only a threshold needs to participate. One way to break this has been proposed in prior work [3, 37] where the shares are published on chain and rewards are given only to “useful” shares (*i.e.* the first t that make it on-chain). This breaks the public goods game and makes it a race between the authorities who are eager to release their share and get rewarded.

8 Applications

In this section, we discuss applications of random beacons. In particular, we argue how the novel history generation feature of STROBE can enable attractive technical advancements in many scenarios.

8.1 Blockchain Light-clients

Blockchain’s light client [20] is typically a resource-constrained node that performs limited verification of the blockchain by downloading block headers and verifying high-level state transitions following consensus decisions. In the context of proof-of-stake blockchains, a light-client would verify the validators’ signatures over block headers (in BFT-style protocols) or leader elections (in longest-chain-style protocols), or both (in hybrid-style blockchains), as well as track the changes to the validator set. Some protocols (e.g. Ethereum 2.0) do subsampling of the committees artificially lowering the number of parties in order to help drive the protocol to agreement faster. Both leader election and subsampling is done through the random-beacon protocols. This guarantees that the process is fair, transparent and verifiable, as well as that no malicious party can increase the probability of them being elected, and that no attacker can predict the next leaders or members of the subcommittees in advance preventing the launch of a targeted attack. All of these properties are provided by verifiability, uniformity, unbiasedness and unpredictability of a random beacon. However, a major challenge for light-clients is to follow and verify the beacon outputs in order to be able to follow and verify consensus decisions. A protocol that allows to generate historical values of the beacon, like ours, would allow the light-client to only verify the latest output and be able to generate the prior values locally simplifying the process of verifying the leaders and subcommittees.

8.2 Blockchain-based Gambling and Lotteries

Lotteries and gambling smart contract solutions are gaining in popularity, and a portion of them advertise transparency, unbiased randomness generation (uncontrollable from participants) and consumer privacy. It is well known that the original success of Bitcoin was partly due to gambling activities, especially via the Satoshi Dice which operated since 2012 and dominated the bitcoin transactions in its first years of operation [32]. There is a growing number of blockchain gambling contracts, and as of September 2021, according to statistics provided in [25], there exist at least 151 lottery, 129 casino, 70 poker and generally 600+ gambling smart contracts in Ethereum alone. Moreover, lotteries have also been proposed as an alternative reward scheme for miners by randomly recirculating lost coins and collecting gold dust [33].

One of the main advantages of STROBE in this setting is that the latest beacon can be automatically used to “verifiably” derive all of the previous random numbers down to the original genesis beacon. This property can be utilized by smart contract developers to minimize cost by skipping beacon epochs when required, but lotteries can still continuously run for every epoch. It also works as a DoS defense, especially when the beacon is provided by an external oracle service to the blockchain; if there is significant delay on updates, the latest beacon suffices to execute all of the pending lottery games.

8.3 High-throughput Beacon Streams

There exist applications requiring constant high-throughput of beacons, especially in the online gaming sector. Most games require a combination of “skills” and “luck”, and as video gaming has become a market where professional players participate in tournaments with real prizes, a fair beacon stream would enable new types of transparent gaming features. Due to potential internet speed issues, server overloading and other factors, reliability might be at stake [23] and thus, some game providers prefer UDP connections to offer greater flexibility by executing packets out of order or discarding non important ones [52]. STROBE’s history generation feature fits really well in streaming designs, and allows client software to generate game states by computing every missing beacon. Note that in these applications, VDF-based beacons are not good candidates due to the intrinsic delay and high-latency, while the proof of stake nature of STROBE offers fairness guarantees by not relying in the event organizer’s or software publisher’s honesty. Along the same lines, STROBE has an advantage in low or expensive bandwidth locations (i.e., remote IoT devices) by allowing reading a beacon infrequently and generating past randomness internally.

Finally, STROBE allows for flexible and more efficient database requirements on nodes serving the API for the derived random beacon per round. Typically, only the last beacon suffices, thus $O(1)$ storage, while one can implement a service that maintains a wisely selected number of checkpoints for faster past-beacon lookups as a trade-off between memory/storage and computations.

References

- 1 Joy Algesheimer, Jan Camenisch, and Victor Shoup. Efficient computation modulo a shared secret with application to the generation of shared safe-prime products. In *CRYPTO*, pages 417–432, 2002.
- 2 Jesús F Almansa, Ivan Damgård, and Jesper Buus Nielsen. Simplified threshold RSA with adaptive and proactive security. In *EUROCRYPT*, pages 593–611, 2006.
- 3 Zeta Avarikioti, EK Kogias, Roger Wattenhofer, and Dionysis Zindros. Brick: Asynchronous incentive-compatible payment channels. In *FC*, pages 209–230, 2021.
- 4 Donald Beaver and Nicol So. Global, unpredictable bit generation without broadcast. In *EUROCRYPT*, pages 424–434, 1993.
- 5 Fabrice Benhamouda, Craig Gentry, Sergey Gorbunov, Shai Halevi, Hugo Krawczyk, Chengyu Lin, Tal Rabin, and Leonid Reyzin. Can a blockchain keep a secret? In *TCC*, pages 260–290, 2020.
- 6 Adithya Bhat, Aniket Kate, Kartik Nayak, and Nibesh Shrestha. Optrand: Optimistically responsive distributed random beacons. In *NDSS*, 2023.
- 7 Lenore Blum, Manuel Blum, and Michael Shub. A simple unpredictable pseudo-random number generator. *SIAM J. Comput.*, 15:364–383, 1986.
- 8 Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *CRYPTO*, pages 757–788, 2018.
- 9 Dan Boneh and Matthew Franklin. Efficient generation of shared RSA keys. In *CRYPTO*, pages 425–439, 1997.
- 10 Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In *EUROCRYPT*, pages 416–432, 2003.
- 11 Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In *ASIACRYPT*, pages 514–532, 2001.
- 12 Vitalik Buterin. Vitalik’s annotated ethereum 2.0 spec. <https://github.com/ethereum/annotated-spec/blob/master/phase0/beacon-chain.md#randao>, 2020.

- 13 Christian Cachin. An asynchronous protocol for distributed computation of RSA inverses and its applications. In *PODC*, pages 153–162, 2003.
- 14 Christian Cachin. Architecture of the Hyperledger blockchain Fabric. In *DCCL*, 2016.
- 15 Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in Constantinople: Practical asynchronous Byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.
- 16 Ran Canetti, Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Adaptive security for threshold cryptosystems. In *CRYPTO*, pages 98–116, 1999.
- 17 Ignacio Cascudo and Bernardo David. SCRAPE: Scalable randomness attested by public entities. In *ACNS*, pages 537–556, 2017.
- 18 Ignacio Cascudo and Bernardo David. ALBATROSS: publicly attestable batched randomness based on secret sharing. In *ASIACRYPT*, pages 311–341, 2020.
- 19 Dario Catalano, Rosario Gennaro, and Shai Halevi. Computing inverses over a shared secret modulus. In *EUROCRYPT*, pages 190–206, 2000.
- 20 Panagiotis Chatzigiannis, Foteini Baldimtsi, and Konstantinos Chalkias. SoK: Blockchain light clients. In *FC*, 2022.
- 21 Panagiotis Chatzigiannis and Konstantinos Chalkias. Proof of assets in the Diem blockchain. In *ACNS*, pages 27–41, 2021.
- 22 Kevin Choi, Aathira Manoj, and Joseph Bonneau. SoK: Distributed randomness beacons. In *IEEE S&P*, pages 75–92, 2023.
- 23 Mia Consalvo. *Cheating: Gaining advantage in videogames*. Mit Press, 2009.
- 24 Ivan Damgård and Maciej Koprowski. Practical threshold RSA signatures without a trusted dealer. In *EUROCRYPT*, pages 152–165, 2001.
- 25 Dapp.com. List of gambling Ethereum smart contracts, 2021. URL: https://www.dapp.com/search_product?keyword=gambling.
- 26 Danny Dolev and H Strong. Polynomial algorithms for multiple processor agreement. *SIAM J Computing*, 12(4):656–666, 1982.
- 27 J. Drake. Minimal VDF randomness beacon. URL: <https://ethresear.ch/t/minimal-vmf-randomness-beacon/3566>.
- 28 Drand - a distributed randomness beacon daemon. URL: <https://drand.love/>.
- 29 Naomi Ephraim, Cody Freitag, Ilan Komargodski, and Rafael Pass. Continuous verifiable delay functions. In *EUROCRYPT*, pages 125–154, 2020.
- 30 Paul Feldman and Sylvio Micali. Byzantine agreement in constant expected time. In *FOCS*, pages 267–276, 1997.
- 31 Yair Frankel, Philip D MacKenzie, and Moti Yung. Robust efficient distributed RSA-key generation. In *STOC*, pages 663–672, 1998.
- 32 Sally M Gainsbury and Alex Blaszczyński. How blockchain and cryptocurrency technology could revolutionize online gambling. *Gaming Law Review*, 21(7):482–492, 2017.
- 33 Harald Gjermundrød, Konstantinos Chalkias, and Ioanna Dionysiou. Going beyond the coinbase transaction fee: Alternative reward schemes for miners in blockchain systems. In *PCI*, pages 1–4, 2016.
- 34 Timo Hanke, Mahnush Movahedi, and Dominic Williams. Dfinity technology overview series, consensus system. *arXiv preprint arXiv:1805.04548*, 2018.
- 35 Aniket Kate and Ian Goldberg. Distributed key generation for the internet. In *ICDCS*, pages 119–128, 2009.
- 36 Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *CRYPTO*, pages 357–388, 2017.
- 37 Eleftherios Kokoris-Kogias, Enis Ceyhun Alp, Linus Gasser, Philipp Jovanovic, Ewa Syta, and Bryan Ford. Calypso: Private data management for decentralized ledgers. *Proc. VLDB Endow.*, 14(4):586–599, 2020.

- 38 Eleftherios Kokoris Kogias, Dahlia Malkhi, and Alexander Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In *CCS*, pages 1751–1767, 2020.
- 39 Sai Krishna Deepak Maram, Fan Zhang, Lun Wang, Andrew Low, Yupeng Zhang, Ari Juels, and Dawn Song. Churp: Dynamic-committee proactive secret sharing. In *CCS*, pages 2369–2386, 2019.
- 40 Gary L Miller. Riemann’s hypothesis and tests for primality. *Journal of Computer and System Sciences*, 13(3):300–317, 1976.
- 41 Dimitrios Papadopoulos, Duane Wessels, Shumon Huque, Moni Naor, Jan Včelák, Leonid Reyzin, and Sharon Goldberg. Making nsec5 practical for dnssec. *Cryptology ePrint Archive*, 2017.
- 42 Krzysztof Pietrzak. Simple verifiable delay functions. In *ITCS*, pages 60:1–60:15, 2019.
- 43 Michael O Rabin. Probabilistic algorithm for testing primality. *Journal of Number Theory*, 12(1):128–138, 1980.
- 44 Michael O. Rabin. Transaction protection by beacons. *Journal of Computer and System Sciences*, 27(2):256–267, 1983.
- 45 Mayank Raikwar and Danilo Gligoroski. Sok: Decentralized randomness beacon protocols. *arXiv preprint arXiv:2205.13333*, 2022.
- 46 Philipp Schindler, Aljosha Judmayer, Markus Hittmeir, Nicholas Stifter, and Edgar Weippl. Randrunner: Distributed randomness from trapdoor VDFs with strong uniqueness. In *NDSS 2022*, 2021.
- 47 Philipp Schindler, Aljosha Judmayer, Nicholas Stifter, and Edgar Weippl. Hydrand: Efficient continuous distributed randomness. In *IEEE S&P*, pages 73–89, 2020.
- 48 Victor Shoup. Practical threshold signatures. In *EUROCRYPT 2000*, pages 207–220, 2000.
- 49 Ewa Syta, Philipp Jovanovic, Eleftherios Kokoris Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J Fischer, and Bryan Ford. Scalable bias-resistant distributed randomness. In *IEEE S&P*, pages 444–460, 2017.
- 50 VeeDo is a STARK-based verifiable delay function (VDF) service. URL: <https://github.com/starkware-libs/veedo>.
- 51 Benjamin Wesolowski. Efficient verifiable delay functions. In *EUROCRYPT*, pages 379–407, 2019.
- 52 Zheng Xue, Di Wu, Jian He, Xiaojun Hei, and Yong Liu. Playing high-end video games in the cloud: A measurement study. *IEEE Trans. Cir. and Sys. for Video Technol.*, 25(12):2013–2025, 2014.