

Bicorn:

An optimistically efficient distributed randomness beacon

Kevin Choi¹, Arasu Arun¹, Nirvan Tyagi², and Joseph Bonneau^{1,3}

¹ New York University

² Cornell University

³ a16z crypto research

Abstract. We introduce Bicorn, an optimistically efficient distributed randomness protocol with strong robustness under a dishonest majority. Bicorn is a “commit-reveal-recover” protocol. Each participant commits to a random value, which are combined to produce a random output. If any participants fail to open their commitment, recovery is possible via a single time-lock puzzle which can be solved by any party. In the optimistic case, Bicorn is a simple and efficient two-round protocol with no time-lock puzzle. In either case, Bicorn supports open, flexible participation, requires only a public bulletin board and no group-specific setup or PKI, and is guaranteed to produce random output assuming any single participant is honest. All communication and computation costs are (at most) linear in the number of participants with low concrete overhead.

1 Introduction

Distributed randomness beacons (DRBs) aim to enable a group of n participants to jointly compute a random output (which we denote Ω) such that no participant or coalition of participants can predict or influence the outcome. Among many other applications, they are useful for cryptographically verifiable lotteries or leader election in efficient distributed consensus protocols.

A classic approach is *commit-reveal* [9]. First, all participants publish a commitment $c_i = \text{Commit}(r_i)$ to a random value r_i . Next, participants reveal their r_i values and the result is $\Omega = \text{Combine}(r_1, \dots, r_n)$ for some suitable combination function (such as exclusive-or or a cryptographic hash). Commit-reveal protocols are simple, efficient and secure as long as one participant chooses a random r_i value—assuming all participants open their commitments. However, the output can be biased by the last participant to open their commitment (a so-called *last-revealer attack*), as that participant will know all other r_i values and can compute Ω early. If the last revealer doesn’t like the impending value of Ω , they can refuse to open, forcing the protocol to abort. Even if the last revealer is removed from subsequent protocol runs, this enables one bit of bias.

Related work. Several approaches exist to avoid last-revealer attacks. Commit-reveal-punish protocols impose a financial penalty on any participant who fails to open their commitment. This penalty can be automatically enforced using

modern cryptocurrencies [2,33], but this requires locking up capital and security relies on economic assumptions about the value of manipulation to the attacker.

Other protocols relax the security model of commit-reveal and assume an honest majority of participants. Many constructions enable a majority of participants to recover the input of a malicious minority of participants [7, 8, 19, 20, 24, 26, 28, 35, 36, 38], using cryptographic tools such as publicly verifiable secret sharing (PVSS). Typically, these constructions can tolerate some threshold t of malicious participants failing to complete the protocol, with the trade-off that any coalition of $t+1$ participants can (secretly) learn the impending output early and potentially bias the protocol, leading to a requirement that $t < \frac{n}{2}$ (honest majority). These protocols are also often quite complex, with communication and computation costs superlinear in n . Another approach is to rely on threshold cryptography for participants to jointly compute a cryptographic function which produces Ω , such as threshold signatures in Dfinity [18], threshold encryption [22], or threshold inversion in RSA groups [3,4]. The drand DRB [1], which uses a chain of threshold BLS signatures, is now deployed publicly with a group of 16 participating nodes producing a new random output every 30 seconds.

A very different approach to constructing DRBs uses time-based cryptography, specifically using *delay functions* to prevent manipulation. The simplest example is Unicorn [29], a one-round protocol in which participants directly publish (within a fixed time window) a random input r_i . The result is computed as $\Omega = \text{Delay}(\text{Combine}(r_1, \dots, r_n))$. By assumption, no party can compute the Delay function before the deadline to publish their contribution r_i and therefore cannot predict Ω or choose r_i in such a way as to influence it. This protocol retains the strong $n-1$ (dishonest majority) security model of commit-reveal, but with no last-revealer attacks. It is also simple and, using modern verifiable delay functions⁴ (VDFs) [11], the result can be efficiently verified. The downside is that a delay function must be computed for every run of the protocol.

Our approach. We introduce the Bicorn family of DRB protocols, which retain the advantages of Unicorn while enabling efficient computation of the result (with no delay) if all participants act honestly. The general structure is:

- Each of n participants chooses a random value r_i and publishes $c_i = \text{TCom}(r_i)$ using a timed commitment scheme [14] TCom before some deadline T_1 .
- In the optimistic case, every participant opens their commitment by publishing r_i . The DRB output is $\Omega = \text{Combine}(r_1, \dots, r_n)$. In this case the protocol is equivalent to a classic commit-reveal protocol.
- If any participant does not publish their r_i value, it can be recovered by computing $r_i = \text{ForceOpen}(c_i)$, a slow function requiring t steps of sequential work which cannot be evaluated quickly enough for a malicious coalition of participants to learn honest participant's committed values early. The result Ω is the same as in the optimistic case, even if all participants don't reveal their committed values.

⁴ The original Unicorn proposal used modular square roots in a prime-order group. We consider using a modern VDF instead.

This protocol structure was used in a recent proposal by Thyagarajan et al. [39]. They observe that by using a homomorphic commitment scheme, the commitments can be combined and only a single forced opening is required, instead of opening every withholding participants’ commitment separately. Asymptotically, their protocols require linear ($O(n)$) communication and computation costs when run with n participants.

However, Thyagarajan et al. use a general-purpose CCA-secure timed commitment scheme suitable for committing to arbitrary messages, which introduces significant practical complexity and overhead. Our key insight is that constructing a DRB does not require a general-purpose commitment scheme; it is sufficient to use a special restricted commitment scheme which only enables committing to a pseudorandom message. As a result, our protocols are considerably simpler and offer much better concrete performance.

Contributions. We introduce the Bicorn family of protocols, which comes in three flavors with slightly different security proofs and practical implications:

- Bicorn-ZK, which requires each participant to publish a zero-knowledge proof of knowledge of exponent. This imposes the highest practical overhead but offers the simplest security proof.
- Bicorn-PC, in which participants “pre-commit” their contribution before the protocol. This is the simplest version, though it adds an extra communication round (which can be amortized over multiple runs).
- Bicorn-RX, which utilizes a randomized exponent to prevent manipulation attacks. This is the most efficient version in practice, though the security proof relies on stronger assumptions.

In Section 3, we prove security of our constructions by reducing to the RSW assumption [34] in the algebraic group model (AGM) [25], except for Bicorn-ZK where we assume a zero-knowledge proof of knowledge of exponent (ZK-PoKE) exists. The Bicorn-RX variant assumes a random oracle. In Section 6, we report on concrete implementations of these protocols in Ethereum, showing that our constructions are practical and incur 3–8× increase in per-user cost compared to commit-reveal (but with no manipulation due to aborts) and 5–7× compared to Unicorn (but with no delay function required in the optimistic case).

2 Overview

2.1 Protocol Outline

We specify all three of our protocol variants in Protocol 1. Our protocols are initialized via a security parameter λ and a delay parameter t , and work over a *group of unknown order*, which we denote \mathbb{G} (see preliminaries in Section 3). In addition to the group \mathbb{G} , the public parameters include a pair (g, h) , where g is a generator of the group and $h = g^{2^t}$. If desired, a Wesolowski [42] or Pietrzak [31] proof of exponentiation can enable efficient verification that h was computed correctly. Note that this setup only needs be run once ever (for a

specific delay parameter t) and can be used repeatedly (and concurrently) by separate protocol instances; the number of participants does not need to be known and may dynamically change over time.

The common structure of Bicorn protocols is:

- Each of n participants chooses a random value α_i and publishes $c_i = g^{\alpha_i}$. The value c_i can be viewed as the input to a VDF whose output is $(c_i)^{2^t}$, with α_i serving as a trapdoor to quickly compute $(c_i)^{2^t} = (g^{\alpha_i})^{2^t} = (g^{2^t})^{\alpha_i} = h^{\alpha_i}$. Without knowledge of α_i this value is slow to compute. Depending on the security assumptions made, α_i can be sampled from different distributions. We abstract this choice by parameterizing by a uniform distribution \mathcal{B} from which α_i is sampled.
- Participants “open” their commitment c_i by revealing a value $\tilde{\alpha}_i$. It can be quickly verified that $\tilde{\alpha}_i$ is the correct α_i by verifying that $c_i = g^{\tilde{\alpha}_i}$.
- *Optimistic case:* Given all correct α_i values, the DRB output Ω is the product $\Omega = \prod_{i \in [n]} h^{\alpha_i}$, which is unpredictable as long as at least one of the values α_i was randomly chosen and is easy to compute if all values α_i are correctly revealed.
- *Pessimistic case:* If any participant withholds α_i (or chose c_i without knowledge of the corresponding α_i), then the missing value h^{α_i} can be recovered (slowly) by computing $h^{\alpha_i} = (c_i)^{2^t}$, equivalent to evaluating a VDF. If multiple participants withhold α_i , naively one must compute each missing value h^{α_i} individually. A more efficient approach (which works even all participants withhold α_i) is to first combine each participants’ contribution into the value $\omega = \prod_{i \in [n]} c_i$. The output can then be computed via a single slow computation as $\Omega = \omega^{2^t}$, which is identical to the output $\Omega = \prod_{i \in [n]} h^{\tilde{\alpha}_i}$ computed in the optimistic case.

By itself this protocol is insecure because a malicious participant need not choose c_i by choosing a value α_i and computing g^{α_i} . An adversary j who has pre-computed a desired output $\Omega_* = (\omega_*)^{2^t}$ and is able to publish last can compute a malicious contribution:

$$c_j = \omega_* \cdot \left(\prod_{i \in [n], i \neq j} c_i \right)^{-1} \quad (1)$$

This will cancel out every other participant’s contribution and force the output value Ω_* . There are three ways to prevent this attack, each leading to a protocol variant with slightly different properties, which we will present in the following subsections. We present the protocols combined for comparison in Protocol 1, with separate presentations in Appendix D.

2.2 Bicorn-ZK: Using Zero-Knowledge Proofs

The conceptually simplest fix is for each user to publish, along with their commitment c_i , a zero-knowledge proof-of-knowledge $\pi_i = \text{ZK-PoKE}(g, c_i, \alpha_i)$ of the

Setup (λ, t)			(run once for all protocol runs)
1. Run $(\mathbb{G}, g, A, B) \xleftarrow{\$} \text{GGen}(\lambda)$ to generate a group of unknown order 2. Compute $h \leftarrow g^{2^t}$, optionally with $\pi_h = \text{PoE}(g, h, 2^t)$ 3. Output $(\mathbb{G}, g, h, \pi_h, A, B)$			
Prepare ()			(run by each participant i)
$\alpha_i \xleftarrow{\$} \mathcal{B}$	$\alpha_i \xleftarrow{\$} \mathcal{B}$	$\alpha_i \xleftarrow{\$} \mathcal{B}$	
$c_i \leftarrow g^{\alpha_i}$	$c_i \leftarrow g^{\alpha_i}$	$c_i \leftarrow g^{\alpha_i}$	
$\pi_i \leftarrow \text{ZK-PoKE}(g, c_i, \alpha_i)$	$d_i \leftarrow H(c_i)$		
Precommit (d_i)			(run by each participant i)
–	Publish d_i	–	
..... <i>deadline</i> T_0			
Commit (c_i, π_i)			(run by each participant i)
Publish c_i, π_i	Publish c_i	Publish c_i	
..... <i>deadline</i> T_1			
Reveal (α_i)			(run by each participant i)
Publish α_i	Publish α_i	Publish α_i	
Finalize ($\{(\tilde{\alpha}_i, c_i, d_i, \pi_i)\}_{i=1}^n$)			(optimistic case, once per protocol run)
1. \forall_j Verify proof π_j – else: remove user j 2. \forall_j Verify $c_j = g^{\tilde{\alpha}_j}$ – else: go to Recover $\Omega = \prod_{i \in [n]} h^{\tilde{\alpha}_i}$	1. \forall_j Verify $d_j = H(c_j)$ – else: remove user j 2. \forall_j Verify $c_j = g^{\tilde{\alpha}_j}$ – else: go to Recover $\Omega = \prod_{i \in [n]} h^{\tilde{\alpha}_i}$	1. $b_* \leftarrow H(c_1 \dots c_n)$ 2. \forall_j Verify $c_j = g^{\tilde{\alpha}_j}$ – else: go to Recover $\Omega = \prod_{i \in [n]} \left(h^{H(c_i b_*)} \right)^{\tilde{\alpha}_i}$	
Recover ($\{(c_i, d_i, \pi_i)\}_{i=1}^n$)			(pessimistic case, once per protocol run)
$\Omega = \left(\prod_{i \in [n]} c_i \right)^{2^t}$	$\Omega = \left(\prod_{i \in [n]} c_i \right)^{2^t}$	$\Omega = \left(\prod_{i \in [n]} c_i^{H(c_i b_*)} \right)^{2^t}$	

Protocol 1: All Bicorn protocol variants: Bicorn-ZK (left column), Bicorn-PC (center column), and Bicorn-RX (right column). Each protocol is presented individually in Appendix D.

discrete logarithm of c_i to the base g_i (i.e. α_i). This version (Bicorn-ZK) is specified in Protocol 1 (left). This removes the attack above, as an adversary who computes c_j via Equation 1 will not know the discrete log of c_j to the base g . Such proofs can be done in groups of unknown order particularly efficiently in this case. The use of a fixed base g enables the simpler ZKPoKRep protocol of Boneh et al. [12] (possibly leveraging of their proof aggregation technique, e.g. the PoKCR protocol).

Participants publishing invalid proofs are removed and the protocol can continue and still produce output. Attempting to participate with an invalid proof is equivalent to not participating at all (though participants who do so might need to be blocked or penalized financially to deter denial-of-service attacks).

It might be tempting to optimize the protocol by not verifying each proof π_i in the optimistic case, instead checking directly that $c_i = g^{\alpha_i}$ using the revealed value α_i . However, this would introduce a subtle attack: a malicious participant could publish a correctly generated $(c_i, \tilde{\alpha}_i)$ pair but with an invalid proof $\tilde{\pi}_i$. Next, after all other participants have revealed their α values, the attacker can compute the impending result Ω with their own contribution included, as well as the alternative Ω' if they are removed. They could then choose which output is produced, introducing one bit of bias into the protocol: by publishing α_i , they will remain in the protocol (as $\tilde{\pi}_i$ is not checked) and Ω will result, whereas by withholding α_i they will force the pessimistic case, in which they will be removed on account of the faulty $\tilde{\pi}_i$ and Ω' will result. Thus, it is important to verify every participant's proof π_i in both cases to prevent this attack.

2.3 Bicorn-PC: Using Precommitment

Another approach to prevent manipulation is to add an initial precommitment round where participants publish $d_i = H(c_i)$, preventing them from choosing c_i in reaction to what others have chosen. This version (Bicorn-PC) is specified in Protocol 1 (center). Participants can decline to reveal their committed c_i , in which case they are removed and the protocol can continue safely. Because participants will not have time to compute the impending output before choosing whether to reveal, this does not introduce any opportunity for manipulation.

Note that the precommitted values d_i can be published at any point prior to T_0 (the point at which participants start revealing their actual commitment c_i). If the protocol is run iteratively, it is possible for participants to publish any number of precommitments d_i in advance (or a single commitment to a set of d_i values using a set commitment construction such as a Merkle Tree), making the protocol a two-round protocol on an amortized basis.

2.4 Bicorn-RX: Using Pseudorandom Exponents

Finally, we can prevent manipulation by raising each participant's contribution c_i to a unique (small) exponent which depends on all other participants' contributions. Specifically, we define b_* to be the hash of all c_i values: $b_* = H(c_1 || c_2 || \dots || c_n)$. We then raise each value c_i to the pseudorandom exponent

Protocol	Rounds	Communication	Assumptions
§2.2 Bicorn-ZK	2	$n(\langle \mathbb{G} \rangle + \langle \mathcal{B} \rangle + \pi)$	RSW, ZK-PoKE
§2.3 Bicorn-PC	3	$n(\langle \mathbb{G} \rangle + \langle \mathcal{B} \rangle + \lambda)$	RSW, AGM
§2.4 Bicorn-RX	2	$n(\langle \mathbb{G} \rangle + \langle \mathcal{B} \rangle)$	RSW, AGM, ROM

Table 1: A brief comparison of the Bicorn variants. See Figure 1 for notation ($\langle \mathbb{G} \rangle$ and $\langle \mathcal{B} \rangle$ are the sizes of elements from \mathbb{G} and \mathcal{B} , respectively) and Section 3 for a background on the RSW assumption, the algebraic group model (AGM), the random oracle model (ROM), and zero-knowledge proof of knowledge of exponent (ZK-PoKE).

$b_i = H(c_i \parallel b_*)$. The intuition is that modifying any contribution c_i will induce new exponents on each participant’s contribution which prevents an adversary from forcing the value $\omega = \prod_{i \in [n]} c_i^{H(c_i \parallel b_*)}$ to a fixed value. A similar technique was used by Boneh et al. [13] to prevent rogue-key attacks in BLS multi-signatures. This version (Bicorn-RX) is specified in Protocol 1 (right).

2.5 Comparison

Each of these leads to a secure protocol, albeit reducing to slightly different computational assumptions, as we will prove in Section 5. All of our protocols reduce to the RSW assumptions with Bicorn-PC and Bicorn-RX requiring the algebraic group model (AGM) for the security reductions and Bicorn-RX also assuming a random oracle. Bicorn-ZK doesn’t require the AGM explicitly but instead assumes a secure zero-knowledge proof of knowledge of exponent (ZK-PoKE) for which efficient existing protocols are proven secure only in the AGM [12].

Each protocol also offers slightly different performance trade-offs, though asymptotically all require $O(n)$ broadcast communication by participating nodes and $O(n)$ computation to verify the result. While Bicorn-PC incurs an extra round, Bicorn-ZK incurs extra computational overhead which may be significant in some scenarios (e.g. smart contracts). Bicorn-RX requires only two rounds and does not require the user to produce proofs but requires extra group exponentiations which incur slightly higher costs than Bicorn-PC.

3 Preliminaries

Algebraic group model. In some of our security proofs, we consider security against *algebraic* adversaries which we model using the algebraic group model, following the treatment of [25]. We call an algorithm \mathcal{A} *algebraic* if for all group elements Z that are output (either as final output or as input to oracles), \mathcal{A} additionally provides the representation of Z relative to all previously received group elements. The previous received group elements include both original inputs to the algorithm and outputs received from calls to oracles. More specifically, if $[X]_i$

$\mathcal{G}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}}(\lambda)$	$\mathcal{G}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}^e}(\lambda)$	$\mathcal{G}_{\mathcal{A},t,b,\text{GGen}}^{\text{D-RSW}}(\lambda)$
$(\mathbb{G}, g, A, B) \xleftarrow{\$} \text{GGen}(\lambda)$	$(\mathbb{G}, g, A, B) \xleftarrow{\$} \text{GGen}(\lambda)$	$(\mathbb{G}, g, A, B) \xleftarrow{\$} \text{GGen}(\lambda)$
$\sigma \leftarrow \mathcal{A}_0(\mathbb{G}, g, A, B)$	$\sigma \leftarrow \mathcal{A}_0(\mathbb{G}, g, A, B)$	$\sigma \leftarrow \mathcal{A}_0(\mathbb{G}, g, A, B)$
$x \xleftarrow{\$} \mathbb{G}$	$x \xleftarrow{\$} \mathbb{G}$	$x \xleftarrow{\$} \mathbb{G}; \tilde{y}_1 \leftarrow x^{2^t}; \tilde{y}_0 \xleftarrow{\$} \mathbb{G}$
$\tilde{y} \xleftarrow{\$} \mathcal{A}_1(\sigma, x)$	$(e, \tilde{y}) \xleftarrow{\$} \mathcal{A}_1(\sigma, x)$	$b' \xleftarrow{\$} \mathcal{A}_1(\sigma, x, \tilde{y}_b)$
Return $\tilde{y} = x^{2^t}$	Return $\tilde{y} = (x^e)^{2^t}$	Return $b = b'$

Fig. 1: Security games for the repeated squaring hardness assumptions: computational RSW (left), computational power-of-RSW (center), and decisional RSW (right).

is the list of group elements $[X_0, \dots, X_n] \in \mathbb{G}$ that \mathcal{A} has received so far, then, when producing group element Z , \mathcal{A} must also provide a list $[z]_i = [z_0, \dots, z_n]$ such that $Z = \prod_i X_i^{z_i}$.

Groups of unknown order and RSW assumptions. Our protocols will operate over cyclic groups of unknown order. We assume an efficient group generation algorithm $\text{GGen}(\lambda)$ that takes as input security parameter λ and outputs a group description \mathbb{G} , generator g , and range $[A, B]$ where A , B , and $B - A$ are all exponential in λ ; the group \mathbb{G} has order in range $[A, B]$. We assume efficient algorithms for sampling from the group ($g \xleftarrow{\$} \mathbb{G}$) and for testing membership.

There are a few currently known options with which to instantiate a group of unknown order. One option that requires only a transparent setup is through class groups of imaginary quadratic order [15]. However, class groups typically incur high concrete overheads. Instead, one may opt for more efficient RSA groups, which require a trusted setup or multiparty computation “ceremony” [21] to compute the modulus $N = pq$ without revealing safe primes p, q . Looking forward, we will require our group to additionally be cyclic and satisfy the low order assumption [10]. So instead we will use the group \mathbb{QR}_N^+ , the group of signed quadratic residues modulo N (we refer to Pietrzak for more details [31]).

The security of our constructions is based on the assumption, originally proposed by RSW [34], that, given a random element $x \in \mathbb{G}$, the fastest algorithm to compute $y = x^{(2^t)}$ takes t sequential steps. We use three RSW assumptions; we provide security games in Figure 1. Detailed descriptions of each assumption are provided in Appendix A.

Randomizing exponent sizes. We recall a useful lemma for randomizing group elements [30].

Lemma 1. *For any cyclic group \mathbb{G} and generator g , if $r \xleftarrow{\$} \mathcal{B}$ is chosen uniformly at random, then the statistical distance between g^r and the uniform distribution over \mathbb{G} is at most $\frac{|\mathbb{G}|}{2|\mathcal{B}|}$.*

Looking forward, we will use this lemma in our security proofs to replace a generator taken to the power of a large exponent of size $|\mathcal{B}| \approx 2^{2\lambda} \cdot |\mathbb{G}|$ with a random element. Alternatively, one may opt for the stronger *short exponent*

indistinguishability (SEI) assumption [23] which asserts that an adversary cannot computationally distinguish between a uniformly random element of \mathbb{G} and g^r for $r \xleftarrow{\$} [0, 2^{2\lambda}]$. The latter assumption enables significant efficiency gains in practice, with participants publishing 32-byte α values instead of 288 bytes.

Non-interactive zero-knowledge proofs. A *non-interactive proof system* for a relation \mathcal{R} over *statement-witness* pairs (x, w) enables producing a proof, $\pi \leftarrow \text{Prove}(pk, x, w)$, that convinces a verifier $\exists w : (x, w) \in \mathcal{R}$, $0/1 \leftarrow \text{Verify}(vk, \pi, x)$; pk and vk are proving and verification keys output by a setup, $(pk, vk) \leftarrow \text{Keygen}(\mathcal{R})$. A *non-interactive argument of knowledge* further convinces the verifier not only that the witness w exists but also that the prover *knows* w , and if proved in *zero-knowledge*, the verifier does not learn any additional information about w . We defer the formal security properties to Appendix B. In this work, we will make use of proof systems for two relations. First, we use PoE for the following relation for proofs of exponentiation in groups of unknown order [12, 31, 42]: $\{(x, y \in \mathbb{G}, \alpha \in \mathbb{Z}), \perp) : y = x^\alpha\}$. Second, we use ZK-PoKE (realized by ZKPoKRep from [12]) for zero-knowledge proofs of knowledge of exponent in groups of unknown order: $\{(x, y \in \mathbb{G}), \alpha \in \mathbb{Z}) : y = x^\alpha\}$.

4 Timed DRBs: Syntax and Security Definitions

We first define a timed DRB using a generalized syntax which captures all of our protocol variants. A timed DRB protocol DRB with time parameter t is a tuple of algorithms (**Setup**, **Commit**, **Verify**, **Finalize**, **Recover**). We describe them below for a run of the protocol with n participants:

- **Setup**(λ, t) $\xrightarrow{\$}$ **pp**: The setup algorithm takes as input a security parameter λ and a time parameter t and outputs a set of public parameters **pp**.
- **Prepare**(**pp**) $\xrightarrow{\$}$ $(\alpha_i, c_i, d_i, \pi_i)$: The prepare algorithm is run by each participant and outputs a tuple of opening, commitment, precommitment, and proof. The precommitment is contributed during the **Precommit** phase (of Protocol 1). The commitment and proof are contributed during the **Commit** phase, and the opening is contributed during the **Reveal** phase. The length of the **Commit** phase is dictated by the time parameter t .
- **Finalize**(**pp**, $\{(\alpha_i, c_i, d_i, \pi_i)\}_{i=1}^n$) $\rightarrow \Omega$: The finalize algorithm is run after the **Reveal** phase and verifies the contributions of participants to optimistically produce a final output Ω or returns \perp indicating the need to move to the pessimistic case.
- **Recover**(**pp**, $\{(c_i, d_i, \pi_i)\}_{i=1}^n$) $\rightarrow \Omega$: The recover algorithm performs the timed computation to recover the output Ω without any revealed α values.

We require **Finalize** to be a deterministic algorithm running in time $\text{polylog}(t)$ (the fast optimistic case), and **Recover** to be a deterministic algorithm running in time $(1 + \epsilon)t$ for some small ϵ . We also require the following security properties of a timed DRB (given in pseudocode in Figure 2):

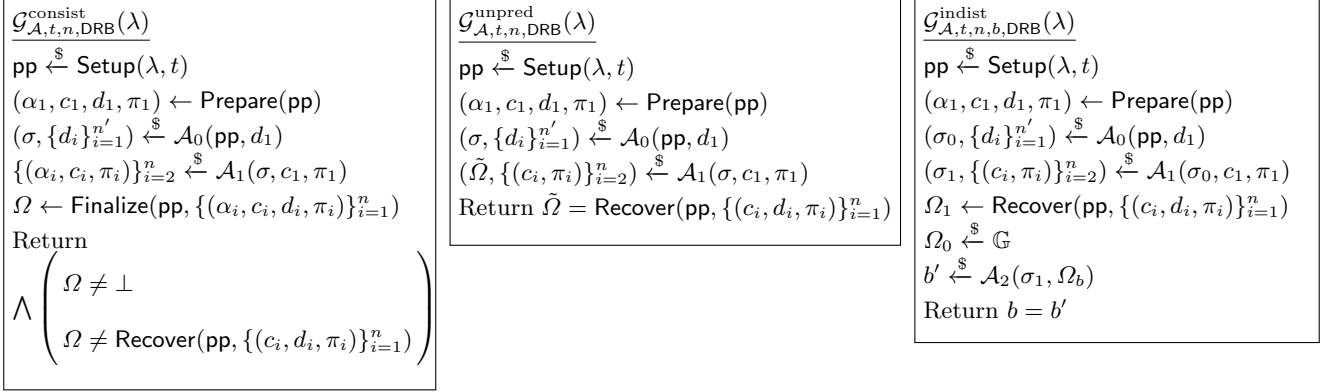


Fig. 2: Security games for our three main security properties: consistency (left), t -unpredictability (center), and t -indistinguishability (right).

Consistency. Our first security property is a form of correctness. We require that it is not possible for the optimistic and pessimistic paths to return different outputs. The adversary is tasked with providing an accepting set of contributions that results in different outputs from `Finalize` and `Recover`. We define the advantage of an adversary as $\text{Adv}_{\mathcal{A},t,n,\text{DRB}}^{\text{consist}}(\lambda) = \Pr \left[\mathcal{G}_{\mathcal{A},t,n,\text{DRB}}^{\text{consist}}(\lambda) = 1 \right]$.

t -Unpredictability. The t -unpredictability game tasks an adversary with predicting the final output Ω exactly, allowing it control of all but a single honest protocol participant (which publishes first). We define the advantage of an adversary as $\text{Adv}_{\mathcal{A},t,n,\text{DRB}}^{\text{unpred}}(\lambda) = \Pr \left[\mathcal{G}_{\mathcal{A},t,n,\text{DRB}}^{\text{unpred}}(\lambda) = 1 \right]$.

t -Indistinguishability. The t -unpredictability property does not guarantee the output is indistinguishable from random. For that, we provide a stronger t -indistinguishability property in which the adversary must distinguish an honest output from a random output, again allowing the adversary control of all but one participant. We define the advantage of an adversary as: $\text{Adv}_{\mathcal{A},t,n,\text{DRB}}^{\text{indist}}(\lambda) = \left| \Pr \left[\mathcal{G}_{\mathcal{A},t,n,1,\text{DRB}}^{\text{indist}}(\lambda) = 1 \right] - \Pr \left[\mathcal{G}_{\mathcal{A},t,n,0,\text{DRB}}^{\text{indist}}(\lambda) = 1 \right] \right|$. A timed DRB that satisfies t -unpredictability can be transformed generically into one with t -indistinguishability by applying a suitable randomness extractor [40, 41] or hash function (modeled as a random oracle) to the output. A nice feature of our DRBs is that they satisfy t -indistinguishability with respect to the group output space (without applying a randomness extractor) under the suitable decisional RSW assumption.

Discussion. In t -unpredictability and t -indistinguishability, the adversaries \mathcal{A}_1 and \mathcal{A}_2 are restricted to run in fewer than t sequential steps. This is a slight simplification of the (p, σ) -sequentiality assumption in VDFs [11], which is suitable for working in the AGM in which parallelism is not helpful in computing group operations.

Note that our syntax and security definitions encompass all three of our protocol variants. Except for Bicorn-ZK, the proofs π_i can be set to \perp and are ignored; except for Bicorn-PC, the precommitment values d_i can be set to \perp and are ignored. Also note that there are n' ($\geq n$) values of d_i output by the adversary; they have the option in Bicorn-PC to choose which to use in later steps. The implementation of `Recover` is unique to each protocol.

We observe that the consistency property holds unconditionally for all Bicorn variants, as `Finalize` and `Recover` are deterministic and algebraically equivalent. It remains to prove unpredictability and indistinguishability for each variant.

5 Security of Bicorn-RX

We present a proof of t -unpredictability for Bicorn-RX here, as it is representative of the techniques used for all other proofs; we defer the full security proofs for Bicorn-ZK, Bicorn-PC, and Bicorn-RX to Appendix C.

Theorem 1 (t -Unpredictability of Bicorn-RX). *Let $\mathcal{A}_{\text{brx}} = (\mathcal{A}_{\text{brx},0}, \mathcal{A}_{\text{brx},1})$ be an algebraic adversary against the t -unpredictability of BRX with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$ where hash function H is modeled as a random oracle. Then we construct an adversary $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{brx},t,n,\text{BRX}}}^{\text{unpred}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{C-RSW}^e}(\lambda) + \frac{2(q_{\text{ro}}^2 + n) + 1}{2^{2\lambda+1}} + \prod_{i=1}^{\ell} I_{\frac{1}{p_i}}(r_i, n),$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates the group of unknown order ($|\mathbb{G}| = \prod_{i=1}^{\ell} p_i^{r_i}$ for distinct primes p_1, \dots, p_{ℓ}) used by BRX, q_{ro} is the number of queries made to the random oracle, n is the number of participants, and $I_{\frac{1}{p}}(r, n) = (1 - \frac{1}{p})^n \sum_{j=r}^{\infty} \binom{n+r-1}{r} p^{-j}$ is the regularized beta function. The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{brx},0}) + 2t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{brx},1})$.

Proof. At a high level, our proof strategy will be to replace the initial commitment c_1 provided by the single honest participant with a random group element. If \mathcal{A}_{brx} can win with non-negligible probability, then we show that due to unpredictability of the random exponents applied in Bicorn-RX, it must be that a nontrivial large exponent of c_1 was computed which we can use to win the computational power-of-RSW game.

More specifically, we bound the advantage of \mathcal{A}_{brx} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G} = \mathcal{G}_{\mathcal{A}_{\text{brx},t,n,\text{BRX}}}^{\text{unpred}}(\lambda)$ and hybrids $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}(\lambda) = 1] - \Pr[\mathcal{G}_1(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_1(\lambda) = 1] - \Pr[\mathcal{G}_2(\lambda) = 1]| \leq \frac{q_{\text{ro}}^2}{2^{2\lambda}}$
- $|\Pr[\mathcal{G}_2(\lambda) = 1] - \Pr[\mathcal{G}_3(\lambda) = 1]| \leq \frac{n}{2^{2\lambda}} + \prod_{i=1}^{\ell} I_{\frac{1}{p_i}}(r_i, n)$
- $\Pr[\mathcal{G}_3(\lambda) = 1] = \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{C-RSW}^e}(\lambda)$

$\mathcal{G} \rightarrow \mathcal{G}_1$. Hybrid \mathcal{G}_1 is defined the same as \mathcal{G} except \mathcal{G}_1 samples c_1 in Prepare at random from \mathbb{G} instead of through an exponent sampled from \mathcal{B} . By Lemma 1, the statistical distance between \mathcal{G} and \mathcal{G}_1 is at most $1/2^{2\lambda+1}$.

We can view \mathcal{G}_1 as computing the beacon output Ω using the representations of $\{c_i\}_{i=2}^n$ provided by the algebraic adversary. Since \mathcal{A}_{brx} is algebraic, it will provide a representation for each c_i in terms of elements (c_1, g, h) . That is, the adversary outputs $[(e_{i,0}, e_{i,1}, e_{i,2})]_{i=2}^n$ such that $c_i = c_1^{e_{i,0}} g^{e_{i,1}} h^{e_{i,2}}$.

Given a value $\hat{h} = h^{2^t}$, we can compute Ω as follows. Consider the random exponents $b_i = H(c_i \parallel b_*)$ where $b_* = H(c_1 \parallel \dots \parallel c_n)$, and let $\mathbf{b} = (b_1, \dots, b_n)$. Using these, we have:

$$\begin{aligned} \Omega &= \left(\prod_{i=1}^n c_i^{b_i} \right)^{2^t} = \left(c_1^{b_1} \cdot \prod_{i=2}^n (c_1^{e_{i,0}} g^{e_{i,1}} h^{e_{i,2}})^{b_i} \right)^{2^t} \\ &= \left(c_1^{b_1 + \sum_{i=2}^n b_i e_{i,0}} g^{\sum_{i=2}^n b_i e_{i,1}} h^{\sum_{i=2}^n b_i e_{i,2}} \right)^{2^t} \\ \text{By letting } \mathbf{e} &= (1, e_{2,0}, \dots, e_{n,0}), m_1 = \sum_{i=2}^n b_i e_{i,1}, \text{ and } m_2 = \sum_{i=2}^n b_i e_{i,2}, \\ &= \left(c_1^{\langle \mathbf{b}, \mathbf{e} \rangle} g^{m_1} h^{m_2} \right)^{2^t} = (c_1^{2^t})^{\langle \mathbf{b}, \mathbf{e} \rangle} \cdot h^{m_1} \cdot \hat{h}^{m_2} \end{aligned}$$

Thus if \mathcal{A}_{brx} wins, i.e., $\tilde{\Omega} = \Omega$, then we have

$$(c_1^{2^t})^{\langle \mathbf{b}, \mathbf{e} \rangle} = \tilde{\Omega} \cdot h^{-m_1} \cdot \hat{h}^{-m_2}$$

and we build \mathcal{A}_{rsw} to win the computational power-of-RSW game by setting c_1 equal to challenge element x and returning this value along with $\langle \mathbf{b}, \mathbf{e} \rangle$. All that is left to show is that $\langle \mathbf{b}, \mathbf{e} \rangle \neq 0$ which we can do through an application of the Schwartz-Zippel lemma modulo a composite [17, 37, 44]. Define a non-zero polynomial $f(x_1, \dots, x_n) = x_1 + \sum_{i=2}^n x_i e_{i,0}$. Note that $f(\mathbf{b}) = \langle \mathbf{b}, \mathbf{e} \rangle$.

$\mathcal{G}_1 \rightarrow \mathcal{G}_2$. To apply the Schwartz-Zippel lemma modulo a composite, we must first have that the evaluation point \mathbf{b} does not coincide with values precomputed by the adversary. To do this, we step through \mathcal{G}_2 in which we disallow the output of the random oracle H from colliding with (the trailing substring of) any previous inputs to the random oracle. This ensures that the adversary has not made any previous queries that include b_* and ultimately ensures that the b_i values are chosen randomly *after* the polynomial is decided. We can apply a standard birthday analysis to bound the probability of collision among the \mathbf{q}_{ro} queries made to $\mathbf{q}_{\text{ro}}^2/2^{2\lambda}$, to bound the distinguishing advantage between \mathcal{G}_1 and \mathcal{G}_2 .

$\mathcal{G}_2 \rightarrow \mathcal{G}_3$. After we have that the evaluation point \mathbf{b} does not coincide with precomputed values, we transition to \mathcal{G}_3 which is identical to \mathcal{G}_2 except it aborts if $f(\mathbf{b}) = 0$. We bound the distinguishing advantage to probability $\frac{n}{2^{2\lambda}} + \prod_{i=1}^{\ell} I_{\frac{1}{p_i}}(r_i, n)$ by applying Schwartz-Zippel modulo a composite [17]. Adversary \mathcal{A}_{rsw} can simulate \mathcal{G}_3 perfectly, simulating the setup and computing \hat{h} with $2t$ work, and wins the RSW game with the same advantage as \mathcal{G}_3 . ■

Gas Costs ($\times 10^3$), Operations Involved

	Commit/user		Reveal/user		Recover		
Commit-Reveal	50	store $_{2\lambda}$	60	xor, hash	-		
[29] Unicorn	55	store $_{2\lambda}$		-	30n n-hash	}	
§2.2 Bicorn-ZK	2,950	zk-poke.v, store $_{\mathbb{G}}$	300	exp, mul	(negligible)		+2,330 poe.v
§2.3 Bicorn-PC	155; 180	mul, store $_{\mathbb{G}}$	300	exp, mul	(negligible)		
§2.4 Bicorn-RX	145	mul, store $_{\mathbb{G}}$	425	2-exp, mul	170n n-exp		

Table 2: Ethereum gas costs and main operations involved for each Bicorn variant as well as Unicorn and Commit-Reveal DRBs. For Bicorn-PC, the **Commit** cost is split to show **Precommit** and **Commit** costs. The operations are: store $_{\mathbb{G}/2\lambda}$, storing a group element or 2λ -bit value; mul, multiplication of two group elements; exp, raising a group element to a power of size 2λ bits; poe.v and zk-poke.v, verifying a proof of exponentiation and proof of knowledge of exponentiation, respectively. Concrete costs are given with $\mathbb{G} = \mathbb{QR}_N^+$ within an RSA-2048 group and $\lambda = 128$.

6 Implementation

We implemented all three variants of Bicorn in Solidity and measured the associated gas costs in Ethereum [43]. Our results are presented in Table 2. We instantiate \mathbb{G} as an RSA group with a 2048-bit modulus (specifically, it is the quadratic residue subgroup \mathbb{QR}_N^+ [31]). Multiplying two group elements costs $\sim 90,000$ gas and raising a group element to a power of size 32 bytes costs $\sim 150,000$ gas. As mentioned in Section 3, we use the short exponent indistinguishability (SEI) assumption [23] to reduce the size of the exponent required in practice from 256 to 32 bytes. The largest costs for each protocol are verifying a proof of exponentiation (PoE) for the VDF computation in the pessimistic Recover case and verifying a zero-knowledge proof of knowledge of exponent needed for each commitment in Bicorn-ZK. We implemented both proofs using non-interactive variants of Wesolowski proofs (ZKPoKRep from [12] for the latter), which requires a prime challenge to be sampled. Verifying this “hash-to-prime” operation costs between 2.3–4 million gas.⁵

Comparison to other DRBs. *Per-user Costs:* We find that the user operations for Bicorn-RX are practical on Ethereum with them costing $3\times$ for Commit and $7\times$ for Reveal when compared to the standard Commit-Reveal and Unicorn protocols. In total, the sum of these operations per user per run comes to under 600,000 gas, or \$6 USD when 1 Eth = \$1,000 USD and 1 gas = 10 Gwei.

⁵ Verifying “hash-to-prime” involves testing the primality of a number on-chain using Pocklington certificates. This costs between 2.3–4 million gas, depending on the size of the certificate. Table 2 reports costs with the smallest possible certificate.

Pessimistic Costs: In the pessimistic case, a single call to `Recover` is required in all versions of Bicorn, costing millions of gas. This pessimistic case is roughly equivalent to *every* run of Unicorn. As the number of users grows large and the chances of Bicorn’s optimistic case occurring decrease though, at some point it may make more sense to switch to Unicorn and avoid the overheads of `Commit` and `Reveal` that Bicorn protocols incur.

7 Discussion

Last revealer prediction. All Bicorn variants come with a fundamental security caveat: if participant j withholds their α_j value, but all others publish, then participant j will be able to simulate the optimistic case and learn Ω quickly, while the honest participants will need to execute the pessimistic case and compute the delay function to complete before learning Ω . Similarly, a coalition of malicious participants can share their α values and privately compute Ω . This issue appears fundamental; in any protocol with a fast optimistic case and a slow pessimistic case, a unified malicious coalition can simulate the optimistic case.

This does not undermine t -unpredictability or t -indistinguishability and does not allow an adversary to manipulate the outcome. As a result, any protocol built on top of Bicorn should consider the output Ω to be potentially available to adversaries as of the deadline T_1 , even if the result is not publicly known until $T_1 + t$ if the pessimistic case is triggered. For example, in a lottery application all wagers must be locked in before time T_1 .

Incentives and punishment. While all Bicorn variants ensure malicious participants cannot manipulate the output, they can waste resources by forcing the protocol into the more-expensive recovery mode. The protocol provides accountability as to which nodes published an incorrect α_i value or other minor deviations which lead to removal (i.e. publishing an incorrect c_i such that $H(c_i) \neq d_i$ in Bicorn-PC or publishing an incorrect π_i in Bicorn-ZK). If signatures are added to each message, efficient fraud proofs are possible. In a blockchain setting, financial penalties can be used to punish incorrect behavior.

Batch verification optimization. In the optimistic case, the n exponentiations required to verify that $c_i = g^{\alpha_i}$ for each participant can be streamlined via batch verification [5, 16]. The general idea is that $g^x = 1 \wedge g^y = 1$ can be verified more efficiently by checking $g^{r \cdot x + y} = 1$ for a random $r \xleftarrow{\$} \mathcal{R}$, as the latter equation implies the former with high probability given a large enough \mathcal{R} . In our case, to verify that $c_1 = g^{\alpha_1} \wedge c_2 = g^{\alpha_2} \wedge \dots \wedge c_n = g^{\alpha_n}$, we generate random values $r_i \xleftarrow{\$} \mathcal{R}$ and verify that $g^{\sum r_i \cdot \alpha_i} = \prod c_i^{r_i}$. Thus, instead of computing n exponentiations each with an exponent of size $|\mathcal{B}|$, verification requires only one exponentiation with an exponent of size $n|\mathcal{B}||\mathcal{R}|$ and one n -way multi-exponentiation [32].

Lowering costs with rollup proofs. Practical costs can become significant if all users must post data to the blockchain to participate. For example, each

run of Bicorn-RX costs about \$6 USD per user even in the optimistic case. An alternative solution is to perform Bicorn mediated via a *rollup server* (Rollup-Bicorn) which gathers every participant’s c_i value and publishes:

- A commitment $s = \text{SetCommitment}(C)$ to the set $C = \{c_1, \dots, c_n\}$ of all participant contributions. For example, s might be a Merkle Tree root.
- The value $c_* = \prod_{i \in [n]} c_i$, the product of all participant’s commitments.
 - For Bicorn-RX, c_* will be adjusted with each party’s exponent $H(c_i || b_*)$.
- A succinct proof (SNARK) $\pi_{\text{rollup-commit}}$ that c_* has been computed consistently with the set S . This proof does not need to be zero-knowledge.
 - For Bicorn-ZK, the proof must recursively check each proof π_i .
 - For Bicorn-PC, the proof must check c_i was correctly precommitted.
 - For Bicorn-RX, the proof must check c_i was raised to the power b_i .

In the optimistic case, if all participants reveal their private value α_i , then the rollup server can finalize the protocol by posting:

- The output Ω and a succinct proof (SNARK) $\pi_{\text{rollup-finalize}}$ that states that:
 - The prover knows a set $A = \{\alpha_1, \dots, \alpha_n\}$
 - For each $c_i \in C$, it holds that $c_i = g^{\alpha_i}$
 - The output Ω was computed correctly given the set A .

In the pessimistic case, if the rollup server goes offline without supplying the second proof (or some participants don’t publish α_i), anybody can still compute $\Omega = c_*^{(2^t)}$. A single proof could be used which is a disjunction of verifying the rollup server’s proof $\pi_{\text{rollup-finalize}}$ or verifying a PoE proof that $\Omega = c_*^{2^t}$. The end result is that Bicorn can be run with $O(1)$ cost for any number of participants.

Lowering cost with delegation. While the rollup approach requires only constant overhead on the blockchain regardless of the number of participants, the primary downside (in common with most rollup systems) is that the rollup server can *ensor* by refusing to include any participant’s c_i in the protocol. In the worst case, a malicious rollup server might only allow participants from a known cabal to participate, who are then able to manipulate the DRB output.

To achieve the best of both worlds (the efficiency of rollup servers for large protocol runs as well as robustness against censorship), we might design a *delegated* Bicorn protocol. In a delegated protocol, users can choose between multiple rollup servers or directly participate as an untrusted (possibly singleton) rollup server. This works like delegated proof-of-stake protocols: participants can delegate for efficiency if they want or participate individually if no server is considered trustworthy. This is straightforward for Bicorn-PC and Bicorn-ZK, as each rollup server can simply compute a partial product c_* which are multiplied together to obtain the final output Ω . Such a protocol for Bicorn-RX would require additional rounds of exponent randomization, to ensure each user’s exponent is randomized by contributions from users at other rollup servers.

Acknowledgments

Kevin Choi, Arasu Arun and Joseph Bonneau were supported by DARPA under Agreement No. HR00112020022. Nirvan Tyagi was supported via a Facebook Graduate Fellowship, and part of this work was done while he was a visiting student at Stanford University. Joseph Bonneau and Arasu Arun were also supported by a16z crypto research. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Government, DARPA, a16z, Facebook or any other supporting organization.

References

1. Drand. <https://drand.love/>
2. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure Multi-party Computations on Bitcoin. In: IEEE Security & Privacy (2014)
3. Beaver, D., Chalkias, K., Kelkar, M., Kogias, L.K., Lewi, K., de Naurois, L., Nicolaenko, V., Roy, A., Sonnino, A.: Strobe: Stake-based threshold random beacons. Cryptology ePrint Archive (2021)
4. Beaver, D., So, N.: Global, unpredictable bit generation without broadcast. In: Eurocrypt (1993)
5. Bellare, M., Garay, J.A., Rabin, T.: Fast batch verification for modular exponentiation and digital signatures. In: Eurocrypt (1998)
6. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 4004, pp. 409–426. Springer (2006)
7. Bhat, A., Kate, A., Nayak, K., Shrestha, N.: OptRand: Optimistically responsive distributed random beacons. Cryptology ePrint Archive, Paper 2022/193 (2022)
8. Bhat, A., Shrestha, N., Kate, A., Nayak, K.: RandPiper – Reconfiguration-Friendly Random Beacons with Quadratic Communication. Cryptology ePrint Archive, Paper 2020/1590 (2020)
9. Blum, M.: Coin flipping by telephone a protocol for solving impossible problems. ACM SIGACT News (1983)
10. Boneh, D., Bünz, B., Fisch, B.: A Survey of Two Verifiable Delay Functions. Cryptology ePrint Archive, Paper 2018/712 (2018)
11. Boneh, D., Bonneau, J., Bünz, B., Fisch, B.: Verifiable delay functions. In: CRYPTO (2018)
12. Boneh, D., Bünz, B., Fisch, B.: Batching techniques for accumulators with applications to IOPs and stateless blockchains. In: CRYPTO (2019)
13. Boneh, D., Drijvers, M., Neven, G.: Compact multi-signatures for smaller blockchains. In: Asiacrypt (2018)
14. Boneh, D., Naor, M.: Timed commitments. In: Annual international cryptology conference (2000)
15. Buchmann, J., Hamdy, S.: A survey on IQ cryptography. In: Public-Key Cryptography and Computational Number Theory (2011)
16. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: IEEE Security & Privacy (2018)
17. Bünz, B., Fisch, B.: Schwartz-zippel for multilinear polynomials mod n . Cryptology ePrint Archive, Paper 2022/458 (2022)
18. Camenisch, J., Drijvers, M., Hanke, T., Pignolet, Y.A., Shoup, V., Williams, D.: Internet computer consensus. In: ACM PODC (2022)
19. Cascudo, I., David, B.: Scrape: Scalable randomness attested by public entities. In: ACNS (2017)
20. Cascudo, I., David, B.: Albatross: publicly attestable batched randomness based on secret sharing. In: Asiacrypt (2020)
21. Chen, M., Hazay, C., Ishai, Y., Kashnikov, Y., Micciancio, D., Riviere, T., Shelat, A., Venkatasubramanian, M., Wang, R.: Diogenes: Lightweight Scalable RSA Modulus Generation with a Dishonest Majority. In: IEEE Security & Privacy (2021)
22. Cherniaeva, A., Shirobokov, I., Shlomovits, O.: Homomorphic encryption random beacon. Cryptology ePrint Archive, Paper 2019/1320 (2019)

23. Couteau, G., Kloof, M., Lin, H., Reichle, M.: Efficient range proofs with transparent setup from bounded integer commitments. In: Eurocrypt (2021)
24. Das, S., Krishnan, V., Isaac, I.M., Ren, L.: Spurt: Scalable distributed randomness beacon with transparent setup. Cryptology ePrint Archive, Paper 2021/100 (2021)
25. Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: CRYPTO (2018)
26. Guo, Z., Shi, L., Xu, M.: SecRand: A Secure Distributed Randomness Generation Protocol With High Practicality and Scalability. IEEE Access (2020)
27. Katz, J., Loss, J., Xu, J.: On the Security of Time-Lock Puzzles and Timed Commitments. In: TCC (2020)
28. Kiayias, A., Russell, A., David, B., Oliynykov, R.: Ouroboros: A provably secure proof-of-stake blockchain protocol. In: CRYPTO (2017)
29. Lenstra, A.K., Wesolowski, B.: A random zoo: sloth, unicorn, and trx. Cryptology ePrint Archive, Paper 2015/366 (2015)
30. Micciancio, D.: The RSA group is pseudo-free. In: CRYPTO (2005)
31. Pietrzak, K.: Simple Verifiable Delay Functions. In: ITCS (2018)
32. Pipenger, N.: On the evaluation of powers and monomials. SIAM Journal on Computing **9**(2), 230–250 (1980)
33. Qian, Y.: Randao: Verifiable random number generation (2017), https://randao.org/whitepaper/Randao_v0.85_en.pdf
34. Rivest, R.L., Shamir, A., Wagner, D.A.: Time-lock puzzles and timed-release crypto (1996)
35. Schindler, P., Judmayer, A., Stifter, N., Weippl, E.: Hydrand: Efficient continuous distributed randomness. In: IEEE Security & Privacy (2020)
36. Schoenmakers, B.: A simple publicly verifiable secret sharing scheme and its application to electronic voting. In: CRYPTO (1999)
37. Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. Journal of the ACM (JACM) **27**(4), 701–717 (1980)
38. Syta, E., Jovanovic, P., Kogias, E.K., Gailly, N., Gasser, L., Khoffi, I., Fischer, M.J., Ford, B.: Scalable bias-resistant distributed randomness. In: IEEE Security & Privacy (2017)
39. Thyagarajan, S.A.K., Castagnos, G., Laguillaumie, F., Malavolta, G.: Efficient CCA Timed Commitments in Class Groups. Cryptology ePrint Archive, Report 2021/1272 (2021)
40. Trevisan, L.: Extractors and pseudorandom generators. Journal of the ACM **48**(4) (2001)
41. Trevisan, L., Vadhan, S.: Extracting randomness from samplable distributions. In: FOCS (2000)
42. Wesolowski, B.: Efficient Verifiable Delay Functions. In: Eurocrypt (2019)
43. Wood, G., et al.: Ethereum: A secure decentralised generalised transaction ledger. Ethereum project yellow paper (2014)
44. Zippel, R.: Probabilistic algorithms for sparse polynomials. In: Symbolic and algebraic manipulation (1979)

A Repeated squaring (RSW) hardness assumptions

The security of our constructions is based on the assumption, originally proposed by RSW [34], that, given a random element $x \in \mathbb{G}$, the fastest algorithm to

compute $y = x^{(2^t)}$ takes t sequential steps. We follow the formalism of Katz et al. [27]. We use three RSW assumptions; we provide security games in Figure 1. *Computational RSW*. In the computational RSW game, the adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ is tasked with computing the 2^t -th power of a challenge element. The adversary acts in two stages. In the preprocessing stage, \mathcal{A}_0 is given a description of the group and outputs an intermediate state passed to \mathcal{A}_1 . In the challenge stage, \mathcal{A}_1 is given the challenge input x and intermediate state σ and attempts to output $y = x^{2^t}$. We define the advantage of an adversary as:

$$\text{Adv}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}}(\lambda) = \Pr [\mathcal{G}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}}(\lambda) = 1]$$

In this game and in the following, the advantage is only meaningful when the challenge stage adversary running time is limited to $< t$. In the Strong AGM (SAGM) [27] and when \mathbb{G} is \mathbb{QR}_N , it is shown that when $\mathcal{A}_0, \mathcal{A}_1$ run in fewer than $\text{poly}(\lambda), t$ steps, respectively, $\mathcal{G}^{\text{C-RSW}}$ can be won with only negligible probability assuming the hardness of factoring [27, Theorem 2].

Computational power-of-RSW. We introduce a stronger variant of computational RSW that we term computational “power-of-RSW.” In this game, the adversary need not output $y = x^{2^t}$ directly, rather the adversary may output (e, y_e) such that $y_e = (x^e)^{2^t}$. The hardness of computational power-of-RSW for time-bounded adversaries can be shown with a slight modification of the proof used to show the hardness of computational RSW in [27]. The SAGM adversary outputs d (alongside $e \neq 0$) with $|d| < 2^t$ such that $x^d = x^{e \cdot 2^t}$. Computing $4(e \cdot 2^t - d)$ then gives a multiple of $\phi(N)$, allowing us to factor N . We define the advantage of an adversary as:

$$\text{Adv}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}^e}(\lambda) = \Pr [\mathcal{G}_{\mathcal{A},t,\text{GGen}}^{\text{C-RSW}^e}(\lambda) = 1]$$

Decisional RSW. Finally, a stronger decisional assumption is that an attacker cannot even distinguish x^{2^t} from a random group element. There is no proof for this assumption, even in generic models. We define the advantage of an adversary as:

$$\text{Adv}_{\mathcal{A},t,\text{GGen}}^{\text{D-RSW}}(\lambda) = |\Pr [\mathcal{G}_{\mathcal{A},t,1,\text{GGen}}^{\text{D-RSW}}(\lambda) = 1] - \Pr [\mathcal{G}_{\mathcal{A},t,0,\text{GGen}}^{\text{D-RSW}}(\lambda) = 1]|$$

B Additional Zero-Knowledge Proof Preliminaries

Here we define the security properties for a non-interactive proof system.

Completeness. A proof system is *complete* if given a true statement, a prover with a witness can convince the verifier. We will make use of a proof systems with perfect completeness. A proof system has *perfect completeness* if for all $(x, w) \in \mathcal{R}$,

$$\Pr [\text{Verify}(x, \text{Prove}(w, x)) = 1] = 1 .$$

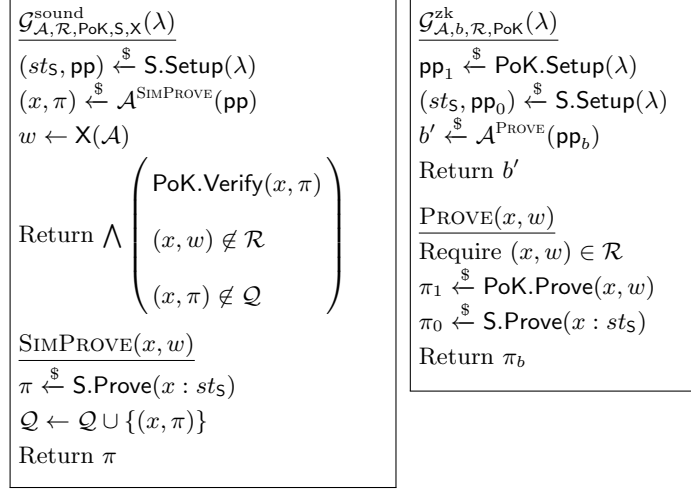


Fig. 3: Soundness (left) and zero knowledge (right) security games for non-interactive zero knowledge proof systems.

Knowledge soundness. A proof system is computationally *knowledge sound* if whenever a prover is able to produce a valid proof for a statement x , it is a true statement, i.e., there exists some witness w such that $(x, w) \in \mathcal{R}$. We require a stronger property to allow for simulating proofs for false statements. This scenario is common in security proofs and so it is desirable to have soundness even in the presence of simulated proofs. This stronger notion of knowledge soundness is known as *simulation-extractability* and is defined by the security game $\mathcal{G}_{\mathcal{A}, \mathcal{R}, \text{PoK}, \mathcal{S}, \mathcal{X}}^{\text{sound}}(\lambda)$ (Figure 3) in which an adversary is tasked with finding a verifying statement and proof where the statement is not in \mathcal{R} . The advantage of an adversary is defined as

$$\text{Adv}_{\mathcal{A}, \mathcal{R}, \text{PoK}, \mathcal{S}, \mathcal{X}}^{\text{sound}}(\lambda) = \Pr[\mathcal{G}_{\mathcal{A}, \mathcal{R}, \text{PoK}, \mathcal{S}, \mathcal{X}}^{\text{sound}}(\lambda) = 1] .$$

Zero knowledge. A proof system is computationally *zero-knowledge* if a proof does not leak any information besides the truth of a statement. Zero knowledge is defined by the security game $\mathcal{G}_{\mathcal{A}, b, \mathcal{R}, \text{PoK}, \mathcal{S}}^{\text{zk}}(\lambda)$ (Figure 3) in which an adversary is tasked with distinguishing between proofs generated from a valid witness and simulated proofs generated without a witness by simulator \mathcal{S} . The advantage of an adversary is defined as

$$\text{Adv}_{\mathcal{A}, \mathcal{R}, \text{PoK}}^{\text{zk}}(\lambda) = \left| \Pr[\mathcal{G}_{\mathcal{A}, 1, \mathcal{R}, \text{PoK}, \mathcal{S}}^{\text{zk}}(\lambda) = 1] - \Pr[\mathcal{G}_{\mathcal{A}, 0, \mathcal{R}, \text{PoK}, \mathcal{S}}^{\text{zk}}(\lambda) = 1] \right| .$$

C Additional security proofs

C.1 Security of Bicorn-RX

Theorem 2 (*t*-Indistinguishability of Bicorn-RX). *Let $\mathcal{A}_{\text{brx}} = (\mathcal{A}_{\text{brx},0}, \mathcal{A}_{\text{brx},1}, \mathcal{A}_{\text{brx},2})$ be an adversary against the *t*-indistinguishability of BRX with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$ where hash function H is modeled as a random oracle. Then we construct an adversary $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{brx},t,n,\text{BRX}}}^{\text{indist}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{D-RSW}}(\lambda) + \frac{2q_{\text{ro}}^2 + 1}{2^{2\lambda}} + 2(\tilde{p} + \tilde{q} - 1) \left(\frac{n}{|\text{im}(H)|} + \frac{n^2}{\tilde{p}\tilde{q}} \right),$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates a group of unknown order \mathbb{QR}_N^+ (whose order is $\tilde{p}\tilde{q}$ where $N = pq$ for $p = 2\tilde{p} + 1$, $q = 2\tilde{q} + 1$) used by BRX, q_{ro} is the number of queries made to the random oracle, n is the number of participants, and $\text{im}(H)$ is the image of H . The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{brx},0}) + 2t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{brx},1}) + T(\mathcal{A}_{\text{brx},2})$.

Proof. We bound the advantage of \mathcal{A}_{brx} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G}^b = \mathcal{G}_{\mathcal{A}_{\text{brx},t,n,b,\text{BRX}}}^{\text{indist}}(\lambda)$ and hybrids $\mathcal{G}_1^b, \mathcal{G}_2^b, \mathcal{G}_3^b$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}^b(\lambda) = 1] - \Pr[\mathcal{G}_1^b(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_1^b(\lambda) = 1] - \Pr[\mathcal{G}_2^b(\lambda) = 1]| \leq \frac{q_{\text{ro}}^2}{2^{2\lambda}}$
- $|\Pr[\mathcal{G}_2^b(\lambda) = 1] - \Pr[\mathcal{G}_3^b(\lambda) = 1]| \leq (\tilde{p} + \tilde{q} - 1) \cdot \left(\frac{n}{|\text{im}(H)|} + \frac{n^2}{\tilde{p}\tilde{q}} \right)$
- $|\Pr[\mathcal{G}_3^1(\lambda) = 1] - \Pr[\mathcal{G}_3^0(\lambda) = 1]| = \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{D-RSW}}(\lambda)$

Hybrid \mathcal{G}_1^b is defined the same as \mathcal{G}^b except \mathcal{G}_1^b samples c_1 in Prepare at random from \mathbb{G} instead of through an exponent sampled from \mathcal{B} . By Lemma 1, the statistical distance between \mathcal{G}^b and \mathcal{G}_1^b is at most $1/2^{2\lambda+1}$.

We can view \mathcal{G}_1^b as computing the beacon output Ω using the representations of $\{c_i\}_{i=2}^n$ provided by the algebraic adversary. Since \mathcal{A}_{brx} is algebraic, it will provide a representation for each c_i in terms of elements (c_1, g, h) . That is, the adversary outputs $[(e_{i,0}, e_{i,1}, e_{i,2})]_{i=2}^n$ such that $c_i = c_1^{e_{i,0}} g^{e_{i,1}} h^{e_{i,2}}$.

As in the unpredictability proof, given a value $\hat{h} = h^{2^t}$, we can compute Ω as follows. Consider the random exponents $b_i = H(c_i \parallel b_*)$ where $b_* = H(c_1 \parallel \dots \parallel c_n)$, and let $\mathbf{b} = (b_1, \dots, b_n)$. Using these, we have:

$$\begin{aligned} \Omega &= \left(\prod_{i=1}^n c_i^{b_i} \right)^{2^t} = \left(c_1^{b_1} \cdot \prod_{i=2}^n (c_1^{e_{i,0}} g^{e_{i,1}} h^{e_{i,2}})^{b_i} \right)^{2^t} \\ &= \left(c_1^{b_1 + \sum_{i=2}^n b_i e_{i,0}} g^{\sum_{i=2}^n b_i e_{i,1}} h^{\sum_{i=2}^n b_i e_{i,2}} \right)^{2^t} \end{aligned}$$

By letting $\mathbf{e} = (1, e_{2,0}, \dots, e_{n,0})$, $m_1 = \sum_{i=2}^n b_i e_{i,1}$, and $m_2 = \sum_{i=2}^n b_i e_{i,2}$,

$$= \left(c_1^{\langle \mathbf{b}, \mathbf{e} \rangle} g^{m_1} h^{m_2} \right)^{2^t} = (c_1^{2^t})^{\langle \mathbf{b}, \mathbf{e} \rangle} \cdot h^{m_1} \cdot \hat{h}^{m_2}$$

We consider a transition in which $c_1^{2^t}$ is replaced with a random group element following the decisional RSW game. We will want to show that the distinguishing advantage for this transition is equal to the advantage of \mathcal{A}_{rsw} . To do this, we will first need to set up the transition with a few more hybrids.

First, as in the unpredictability game, we transition through \mathcal{G}_2^b to disallow collisions in the random oracle in the way specified in the proof of Theorem 1. Next, we further define \mathcal{G}_3^b that is the same as \mathcal{G}_2^b but aborts if $g^{\langle \mathbf{b}, \mathbf{e} \rangle}$ is not a generator of \mathbb{QR}_N^+ . This happens if $\gcd(\langle \mathbf{b}, \mathbf{e} \rangle, |\mathbb{QR}_N^+|) \neq 1$, a condition that is equivalent (in modulo $|\mathbb{QR}_N^+| = \tilde{p}\tilde{q}$) to $f(\mathbf{b}) = \langle \mathbf{b}, \mathbf{e} \rangle = k$ for $k \in \{0, \tilde{p}, \tilde{q}, \dots, \tilde{p}\tilde{q} - \tilde{p}, \tilde{p}\tilde{q} - \tilde{q}\}$. As we can apply Schwartz-Zippel modulo a composite [17, Remark 3] to each k and there are $\tilde{p} + \tilde{q} - 1$ such k 's, we apply the union bound to bound the probability that $g^{\langle \mathbf{b}, \mathbf{e} \rangle}$ is not a generator of \mathbb{QR}_N^+ to $(\tilde{p} + \tilde{q} - 1) \cdot \left(\frac{n}{|\text{im}(H)|} + \frac{n^2}{\tilde{p}\tilde{q}} \right)$, where $|\text{im}(H)|$ denotes the size of the image of the random oracle. While a large $|\text{im}(H)|$ may be required for this theoretical bound, a remark here is that one may opt for stronger assumptions, such as the short exponent indistinguishability (SEI) assumption [23] due to the fact that H outputs occur only in the exponents, for efficiency gains in practice.

Now, we are set up to construct \mathcal{A}_{rsw} by bounding the distinguishing advantage between \mathcal{G}_3^1 and \mathcal{G}_3^0 where \mathcal{G}_3^1 computes the challenge Ω using the process above, while \mathcal{G}_3^0 computes the challenge Ω by replacing $c_1^{2^t}$ with a random group element. Note that in the case of \mathcal{G}_3^1 , Ω is computed to match the output of Recover. On the other hand, in the case of \mathcal{G}_3^0 , Ω is in fact a random group element. This can be seen as follows. The random group element that replaces $c_1^{2^t}$ can be written as g^r for some $r \xleftarrow{\$} [1, |\mathbb{G}|]$, and so we have:

$$\Omega = (g^r)^{\langle \mathbf{b}, \mathbf{e} \rangle} \cdot h^{m_1} \cdot \hat{h}^{m_2} = (g^{\langle \mathbf{b}, \mathbf{e} \rangle})^r \cdot h^{m_1} \cdot \hat{h}^{m_2}$$

Since $g^{\langle \mathbf{b}, \mathbf{e} \rangle}$ is a generator by the previous hybrid transition, we have that $(g^{\langle \mathbf{b}, \mathbf{e} \rangle})^r$ is a random group element, and so Ω is a random group element. Thus, \mathcal{A}_{rsw} perfectly simulates \mathcal{G}_3^b based on the challenge bit, simulating the setup and computing \hat{h} with $2t$ work, and wins the RSW game with the same advantage as the distinguishing advantage between \mathcal{G}_3^1 and \mathcal{G}_3^0 . ■

C.2 Security of Bicorn-ZK

Theorem 3 (t -Unpredictability of Bicorn-ZK). *Let $\mathcal{A}_{\text{bzk}} = (\mathcal{A}_{\text{bzk},0}, \mathcal{A}_{\text{bzk},1})$ be an adversary against the t -unpredictability of BZK with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$. Then we construct adversaries $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$, \mathcal{A}_{zk} , and $\mathcal{A}_{\text{sound}}$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{bzk}}, t, n, \text{BZK}}^{\text{unpred}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw}}, t, \text{GGen}}^{\text{C-RSW}}(\lambda) + \text{Adv}_{\mathcal{A}_{\text{zk}}, \text{ZK-PoKE}}^{\text{zk}}(\lambda) + \text{Adv}_{\mathcal{A}_{\text{sound}}, \text{ZK-PoKE}, \text{S}, \text{X}}^{\text{sound}}(\lambda) + \frac{1}{2^{2\lambda+1}},$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates the group of unknown order used by BZK, and S and X are the simulator and extractor for ZK-PoKE. The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{bzk},0}) + t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{bzk},1})$.

Proof. We bound the advantage of \mathcal{A}_{bzk} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G} = \mathcal{G}_{\mathcal{A}_{\text{bzk}}, t, n, \text{BZK}}^{\text{unpred}}(\lambda)$ and hybrids $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}(\lambda) = 1] - \Pr[\mathcal{G}_1(\lambda) = 1]| \leq \text{Adv}_{\mathcal{A}_{\text{zk}}, \text{ZK-PoKE}}^{\text{zk}}(\lambda)$
- $|\Pr[\mathcal{G}_1(\lambda) = 1] - \Pr[\mathcal{G}_2(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_2(\lambda) = 1] - \Pr[\mathcal{G}_3(\lambda) = 1]| \leq \text{Adv}_{\mathcal{A}_{\text{sound}}, \text{ZK-PoKE}, \text{S}, \text{X}}^{\text{sound}}(\lambda)$
- $\Pr[\mathcal{G}_3(\lambda) = 1] = \text{Adv}_{\mathcal{A}_{\text{rsw}}, t, \text{GGen}}^{\text{C-RSW}}(\lambda)$

Hybrid \mathcal{G}_1 is defined the same as \mathcal{G} except \mathcal{G}_1 simulates the zero-knowledge proof π_1 in **Prepare**. The distinguishing advantage is directly bounded by the zero-knowledge property of ZK-PoKE.

Hybrid \mathcal{G}_2 is defined the same as \mathcal{G}_1 except \mathcal{G}_2 samples c_1 in **Prepare** at random from \mathbb{G} instead of through an exponent sampled from \mathcal{B} . By Lemma 1, the statistical distance between \mathcal{G}_1 and \mathcal{G}_2 is at most $1/2^{2\lambda+1}$.

Hybrid \mathcal{G}_3 extracts the discrete log $\{\alpha_i\}_{i=2}^n$ from the adversary-provided $\{\pi_i\}_{i=2}^n$ using the extractor from the knowledge soundness property of ZK-PoKE. We bound the probability of any extraction failure using an adversary against the simulation-extractability soundness of ZK-PoKE.

Game \mathcal{G}_3 outputs 1 when $\tilde{\Omega} = \Omega = (c_1)^{2^t} \cdot \prod_{i=2}^n h^{\alpha_i}$. We build an adversary \mathcal{A}_{rsw} that wins the computational RSW game with the same advantage as \mathcal{G}_3 by replacing c_1 with challenge x and outputting $\tilde{y} = \frac{\tilde{\Omega}}{\prod_{i=2}^n h^{\alpha_i}}$. ■

Theorem 4 (*t*-Indistinguishability of Bicorner-ZK). *Let $\mathcal{A}_{\text{bzk}} = (\mathcal{A}_{\text{bzk},0}, \mathcal{A}_{\text{bzk},1}, \mathcal{A}_{\text{bzk},2})$ be an adversary against the *t*-indistinguishability of BZK with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$. Then we construct adversaries $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$, \mathcal{A}_{zk} , and $\mathcal{A}_{\text{sound}}$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{bzk}}, t, n, \text{BZK}}^{\text{indist}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw}}, t, \text{GGen}}^{\text{D-RSW}}(\lambda) + 2 \cdot \text{Adv}_{\mathcal{A}_{\text{zk}}, \text{ZK-PoKE}}^{\text{zk}}(\lambda) + 2 \cdot \text{Adv}_{\mathcal{A}_{\text{sound}}, \text{ZK-PoKE}, \text{S}, \text{X}}^{\text{sound}}(\lambda) + \frac{1}{2^{2\lambda}},$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates the group of unknown order used by BZK, and S and X are the simulator and extractor for ZK-PoKE. The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{bzk},0}) + t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{bzk},1}) + T(\mathcal{A}_{\text{bzk},2})$.

Proof. We bound the advantage of \mathcal{A}_{bzk} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G}^b = \mathcal{G}_{\mathcal{A}_{\text{bzk}}, t, n, b, \text{BZK}}^{\text{indist}}(\lambda)$ and hybrids $\mathcal{G}_1^b, \mathcal{G}_2^b, \mathcal{G}_3^b$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}^b(\lambda) = 1] - \Pr[\mathcal{G}_1^b(\lambda) = 1]| \leq \text{Adv}_{\mathcal{A}_{\text{zk}}, \text{ZK-PoKE}}^{\text{zk}}(\lambda)$
- $|\Pr[\mathcal{G}_1^b(\lambda) = 1] - \Pr[\mathcal{G}_2^b(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_2^b(\lambda) = 1] - \Pr[\mathcal{G}_3^b(\lambda) = 1]| \leq \text{Adv}_{\mathcal{A}_{\text{sound}}, \text{ZK-PoKE}, \text{S}, \text{X}}^{\text{sound}}(\lambda)$
- $|\Pr[\mathcal{G}_3^b(\lambda) = 1] - \Pr[\mathcal{G}_3^0(\lambda) = 1]| = \text{Adv}_{\mathcal{A}_{\text{rsw}}, t, \text{GGen}}^{\text{D-RSW}}(\lambda)$

Hybrids \mathcal{G}_1^b , \mathcal{G}_2^b and \mathcal{G}_3^b are defined as in the unpredictability proof for Bicorner-ZK, simulating π_1 , sampling a random c_1 , and extracting $\{\alpha_i\}_{i=2}^n$, respectively.

In \mathcal{G}_3^1 , the challenge output is computed to match the output of Recover as $\Omega = (c_1)^{2^t} \cdot \prod_{i=2}^n h^{\alpha_i}$. In \mathcal{G}_3^0 , the challenge output is computed in the same way but by replacing $(c_1)^{2^t}$ with a random group element resulting in Ω to be a random group element. Thus, \mathcal{A}_{rsw} perfectly simulates \mathcal{G}_3^b based on the challenge bit (by setting c_1 equal to challenge input x and replacing $c_1^{2^t}$ with challenge input y) and wins the RSW game with the same advantage as the distinguishing advantage between \mathcal{G}_3^1 and \mathcal{G}_3^0 . ■

C.3 Security of Bicorn-PC

Theorem 5 (*t*-Unpredictability of Bicorn-PC). *Let $\mathcal{A}_{\text{bpc}} = (\mathcal{A}_{\text{bpc},0}, \mathcal{A}_{\text{bpc},1})$ be an adversary against the *t*-unpredictability of BPC with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$ where hash function H is modeled as a random oracle. Then we construct an adversary $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{bpc},t,n,\text{BPC}}}^{\text{unpred}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{C-RSW}}(\lambda) + \frac{4n \cdot \mathbf{q}_{\text{ro}} + 1}{2^{2\lambda+1}},$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates the group of unknown order used by BPC, n is the number of participants, and \mathbf{q}_{ro} is the number of queries made to the random oracle. The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{bpc},0}) + 2t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{bpc},1})$.

Proof. We bound the advantage of \mathcal{A}_{bpc} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G} = \mathcal{G}_{\mathcal{A}_{\text{bpc},t,n,\text{BPC}}}^{\text{unpred}}(\lambda)$ and hybrids $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}(\lambda) = 1] - \Pr[\mathcal{G}_1(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_1(\lambda) = 1] - \Pr[\mathcal{G}_2(\lambda) = 1]| \leq \frac{n \cdot \mathbf{q}_{\text{ro}}}{2^{2\lambda}}$
- $|\Pr[\mathcal{G}_2(\lambda) = 1] - \Pr[\mathcal{G}_3(\lambda) = 1]| \leq \frac{n \cdot \mathbf{q}_{\text{ro}}}{2^{2\lambda}}$
- $\Pr[\mathcal{G}_3(\lambda) = 1] = \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{C-RSW}}(\lambda)$

Hybrid \mathcal{G}_1 is defined the same as \mathcal{G} except \mathcal{G}_1 samples c_1 in Prepare at random from \mathbb{G} instead of through an exponent sampled from \mathcal{B} . By Lemma 1, the statistical distance between \mathcal{G} and \mathcal{G}_1 is at most $1/2^{2\lambda+1}$.

Hybrid \mathcal{G}_2 is defined the same as \mathcal{G}_1 except we disallow collisions in the random oracle used for precommitments, i.e., we use sampling without replacement instead of sampling from $[2^{2\lambda}]$. We can apply a standard birthday analysis to bound the probability of collision among the n queries made to $n \cdot \mathbf{q}_{\text{ro}}/2^{2\lambda}$, which bounds the distinguishing advantage between \mathcal{G}_1 and \mathcal{G}_2 .

We can view \mathcal{G}_2 as computing the beacon output Ω using the representations of $\{c_i\}_{i=2}^n$ provided by the algebraic adversary. Since \mathcal{A}_{bpc} is algebraic, it will provide a representation for each c_i not in terms of (c_1, g, h) , but in terms of (g, h) . The reason, which is important to note, is that the adversary needs to precommit before being given c_1 in Bicorn-PC. Accordingly, we check if c_i was queried to the random oracle by $\mathcal{A}_{\text{bpc},0}$ for each $\{c_i\}_{i=2}^n$. Since we disallow collisions in the random oracle in a prior game hop, there is only one possible c_i

that maps to each d_i . If the random oracle was not queried on c_i then we do not have a representation for c_i in (g, h) . The contribution c_i will affect the output if the sampling of a new value for $H(c_i)$ matches d_i provided by the adversary. If it does, this is a “bad” case, and we can bound the probability of this occurring for all n by $n \cdot \mathbf{q}_{\text{ro}}/2^{2\lambda}$. We transition to \mathcal{G}_3 where this bad case does not occur.

Then we have that the adversary outputs $[(e_{i,1}, e_{i,2})]_{i=2}^n$ such that $c_i = g^{e_{i,1}} h^{e_{i,2}}$. Using this, and given a value $\hat{h} = h^{2^t}$, we can compute Ω as follows:

$$\begin{aligned} \Omega &= \left(\prod_{i=1}^n c_i \right)^{2^t} = \left(c_1 \cdot \prod_{i=2}^n g^{e_{i,1}} h^{e_{i,2}} \right)^{2^t} \\ &= \left(c_1 \cdot g^{\sum_{i=2}^n e_{i,1}} \cdot h^{\sum_{i=2}^n e_{i,2}} \right)^{2^t} \\ \text{By letting } m_1 &= \sum_{i=2}^n e_{i,1} \text{ and } m_2 = \sum_{i=2}^n e_{i,2}, \\ &= (c_1 \cdot g^{m_1} \cdot h^{m_2})^{2^t} = (c_1^{2^t}) \cdot h^{m_1} \cdot \hat{h}^{m_2} \end{aligned}$$

Thus if \mathcal{A}_{bpc} wins, i.e., $\tilde{\Omega} = \Omega$, then we have

$$(c_1^{2^t}) = \tilde{\Omega} \cdot h^{-m_1} \cdot \hat{h}^{-m_2},$$

and we build \mathcal{A}_{rsw} to win the computational RSW game by setting c_1 equal to challenge element x and returning this value. The simulation is perfect, with $2t$ work to perform setup and compute \hat{h} , and thus the advantage of \mathcal{A}_{rsw} matches the advantage of \mathcal{G}_3 . \blacksquare

Theorem 6 (t -Indistinguishability of Bicorn-PC). *Let $\mathcal{A}_{\text{bpc}} = (\mathcal{A}_{\text{bpc},0}, \mathcal{A}_{\text{bpc},1}, \mathcal{A}_{\text{bpc},2})$ be an adversary against the t -indistinguishability of BPC with random exponent space $\mathcal{B} = [2^{2\lambda} \cdot B]$ where hash function H is modeled as a random oracle. Then we construct an adversary $\mathcal{A}_{\text{rsw}} = (\mathcal{A}_{\text{rsw},0}, \mathcal{A}_{\text{rsw},1})$ such that*

$$\text{Adv}_{\mathcal{A}_{\text{bpc},t,n,\text{BPC}}}^{\text{indist}}(\lambda) \leq \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{D-RSW}}(\lambda) + \frac{4n \cdot \mathbf{q}_{\text{ro}} + 1}{2^{2\lambda}},$$

and where $\text{GGen} \xrightarrow{\$} (\mathbb{G}, g, A, B)$ generates the group of unknown order used by BPC, n is the number of participants, and \mathbf{q}_{ro} is the number of queries made to the random oracle. The running time of $T(\mathcal{A}_{\text{rsw},0}) \approx T(\mathcal{A}_{\text{bpc},0}) + 2t$ and $T(\mathcal{A}_{\text{rsw},1}) \approx T(\mathcal{A}_{\text{bpc},1}) + T(\mathcal{A}_{\text{bpc},2})$.

Proof. We bound the advantage of \mathcal{A}_{bpc} by bounding the advantage of a series of game hops, using the fundamental lemma of game playing and its identical-until-bad argument [6]. We define $\mathcal{G}^b = \mathcal{G}_{\mathcal{A}_{\text{bpc},t,n,\text{BPC}}}^{\text{indist}}(\lambda)$ and hybrids $\mathcal{G}_1^b, \mathcal{G}_2^b, \mathcal{G}_3^b$ for which we justify the following claims leading to the inequality above:

- $|\Pr[\mathcal{G}^b(\lambda) = 1] - \Pr[\mathcal{G}_1^b(\lambda) = 1]| \leq \frac{1}{2^{2\lambda+1}}$
- $|\Pr[\mathcal{G}_1^b(\lambda) = 1] - \Pr[\mathcal{G}_2^b(\lambda) = 1]| \leq \frac{n \cdot \mathbf{q}_{\text{ro}}}{2^{2\lambda}}$
- $|\Pr[\mathcal{G}_2^b(\lambda) = 1] - \Pr[\mathcal{G}_3^b(\lambda) = 1]| \leq \frac{n \cdot \mathbf{q}_{\text{ro}}}{2^{2\lambda}}$
- $|\Pr[\mathcal{G}_3^1(\lambda) = 1] - \Pr[\mathcal{G}_3^0(\lambda) = 1]| = \text{Adv}_{\mathcal{A}_{\text{rsw},t,\text{GGen}}}^{\text{D-RSW}}(\lambda)$

Bicorn Setup

Setup*input:* λ, t *output:* group \mathbb{G} , generators $g, h \in \mathbb{G}$, proof π_h , range $[A, B]$

1. Run $(\mathbb{G}, g, A, B) \xleftarrow{\$} \text{GGen}(\lambda)$ to generate a group of unknown order
2. Compute $h \leftarrow g^{2^t}$, optionally with $\pi_h = \text{PoE}(g, h, 2^t)$

Protocol 2: Bicorn setup routine (common to all protocol variants), where PoE is a proof of exponentiation [12].

Hybrids \mathcal{G}_1^b , \mathcal{G}_2^b and \mathcal{G}_3^b are defined as in the unpredictability proof for Bicorn-PC, sampling a random c_1 , disallowing random oracle collisions, and disallowing precommitments that do not provide a representation in (g, h) , respectively. In \mathcal{G}_3^1 , the challenge output is computed to match Recover as $\Omega = (c_1)^{2^t} \cdot h^{m_1} \cdot \hat{h}^{m_2}$. In \mathcal{G}_3^0 , the challenge output is computed in the same way but by replacing $(c_1)^{2^t}$ with a random group element resulting in Ω to be a random group element. Thus, \mathcal{A}_{RSW} perfectly simulates \mathcal{G}_3^b based on the challenge bit (by setting c_1 equal to challenge input x and replacing $c_1^{2^t}$ with challenge input y) and wins the RSW game with the same advantage as the distinguishing advantage between \mathcal{G}_3^1 and \mathcal{G}_3^0 . ■

D Individual protocol presentations

For reference, we present each protocol variant separately.

Bicorn-ZK

..... deadline T_0

Commit

Each participant i runs:

1. Sample $\alpha_i \xleftarrow{\$} \mathcal{B}$
2. Compute $c_i \leftarrow g^{\alpha_i}$
3. Compute $\pi_i \leftarrow \text{ZK-PoKE}(g, c_i, \alpha_i)$
4. Publish c_i, π_i

..... deadline T_1

Reveal

Each participant i runs:

1. Publish α_i

Finalize

input: $c_i, \pi_i, \tilde{\alpha}_i$ for $i \in [1, n]$

output: Ω

1. For all users i , verify π_i using c_i
 - (a) If verification fails for any π_i , remove participant i
2. Verify that $c_i = g^{\tilde{\alpha}_i}$ for all $i \in [1, n]$
 - (a) If so, output $\Omega = \prod_{i \in [n]} h^{\tilde{\alpha}_i}$ // optimistic case
3. Output $\Omega = \left(\prod_{i \in [n]} c_i \right)^{2^t}$ // pessimistic case
 - (a) Optionally, a proof π_Ω can be output to enable efficient verification of Ω

Protocol 3: Bicorn protocol with zero-knowledge proofs of knowledge of exponent (ZK-PoKE)

Bicorn-PC

Precommit

Each participant i runs:

1. Sample $\alpha_i \xleftarrow{\$} \mathcal{B}$
2. Compute $c_i \leftarrow g^{\alpha_i}$
3. Publish $d_i = H(c_i)$

..... *deadline* T_0

Commit

Each participant i runs:

1. Publish c_i

..... *deadline* T_1

Reveal

Each participant i runs:

1. Publish α_i

Finalize

input: $d_i, \tilde{c}_i, \tilde{\alpha}_i$ for $i \in [1, n]$

output: Ω

1. Verify that $d_i = H(\tilde{c}_i)$ for all $i \in [1, n]$
 - (a) If any $d_i \neq H(\tilde{c}_i)$ or \tilde{c}_i was not published by T_1 , remove participant i
2. Verify that $\tilde{c}_i = g^{\tilde{\alpha}_i}$ for all $i \in [1, n]$
 - (a) If so, output $\Omega = \prod_{i \in [n]} h^{\tilde{\alpha}_i}$ // optimistic case
3. Output $\Omega = \left(\prod_{i \in [n]} \tilde{c}_i \right)^{2^t}$ // pessimistic case
 - (a) Optionally, a proof π_Ω can be output to enable efficient verification of Ω

Protocol 4: Bicorn protocol with precommitment round

Bicorn-RX

..... *deadline* T_0

Commit

Each participant i runs:

1. Sample $\alpha_i \xleftarrow{\$} \mathcal{B}$
2. Compute $c_i \leftarrow g^{\alpha_i}$
3. Publish c_i

..... *deadline* T_1

Reveal

Each participant i runs:

1. Publish α_i

Finalize

input: $c_i, \tilde{\alpha}_i$ for $i \in [1, n]$

output: Ω

1. Compute $b_* = H(c_1 || c_2 || \dots || c_n)$
2. Verify that $c_i = g^{\tilde{\alpha}_i}$ for all $i \in [1, n]$
 - (a) If so, output $\Omega = \prod_{i \in [n]} \left(h^{H(c_i || b_*)} \right)^{\tilde{\alpha}_i}$ // optimistic case
3. Output $\Omega = \left(\prod_{i \in [n]} c_i^{H(c_i || b_*)} \right)^{2^t}$ // pessimistic case
 - (a) Optionally, a proof π_Ω can be output to enable efficient verification of Ω

Protocol 5: Bicorn protocol with randomized exponents using a random oracle H