

Powers of Tau in Asynchrony

Sourav Das*, Zhuolun Xiang[†], and Ling Ren*

*University of Illinois at Urbana-Champaign, [†]Aptos Labs

souravd2@illinois.edu, xiangzhuolun@gmail.com, renling@illinois.edu

Abstract—The q -Strong Diffie-Hellman (q -SDH) parameters are foundational to efficient constructions of many cryptographic primitives such as zero-knowledge succinct non-interactive arguments of knowledge, polynomial/vector commitments, verifiable secret sharing, and randomness beacon. The only existing method to generate these parameters securely is highly sequential, requires synchrony assumptions, and has very high communication and computation costs. For example, to generate parameters for any given q , each party incurs a communication cost of $\Omega(nq)$ and requires $\Omega(n)$ rounds. Here n is the number of parties in the secure multiparty computation protocol. Since q is typically large, i.e., on the order of billions, the cost is highly prohibitive.

In this paper, we present a distributed protocol to generate q -SDH parameters in an asynchronous network. In a network of n parties, our protocol tolerates up to one-third of malicious parties. Each party incurs a communication cost of $O(q + n^2 \log q)$ and the protocol finishes in $O(\log q + \log n)$ expected rounds. We provide a rigorous security analysis of our protocol. We implement our protocol and evaluate it with up to 128 geographically distributed parties. Our evaluation illustrates that our protocol is highly scalable and results in a 2-6 $\times$ better runtime and 4-13 $\times$ better per-party bandwidth usage compared to the state-of-the-art synchronous protocol for generating q -SDH parameters.

I. INTRODUCTION

The q -Strong Diffie Hellman assumption, or q -SDH assumption for short, refers to the cryptographic intractability problem of computing a group element of the form $(\tau + i)^{-1}g$ for any $i \neq -\tau \in \mathbb{F}$, given $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$. Here, g is a generator of a group $\mathbb{G}$ (typically an elliptic curve group), and τ is a random field element from the scalar field $\mathbb{F}$ of $\mathbb{G}$. The vector $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$ is referred to as the q -SDH parameters, also known as the *powers of τ* .

The q -SDH assumption is used in many applications. Boneh and Boyen [13] introduced the q -SDH assumption to design a short signature scheme that does not rely on a random oracle. Kate, Zaverucha, and Goldberg used the q -SDH assumption to design a constant size polynomial commitment scheme with constant size opening proofs [39]. This polynomial commitment scheme and many of its follow-up schemes have been used extensively in designing a variety of Succinct Non-interactive Argument of Knowledge (SNARK) protocol [44], [53], [19], [33]. The q -SDH parameters have also been used in designing efficient verifiable secret sharing [6], [5], cryptographic accumulators [38], vector commitments [52], distributed randomness beacon [11], etc.

For many applications, the degree q is typically very large. For example, in vector commitment schemes, q is the size of the committed vector, which can be very large [52]. In SNARKs, q is proportional to the size of the SNARK circuit, measured in the number of multiplication gates, which can range from a few million to hundreds of millions [44].

For the q -SDH assumption to hold, it is critical to keep τ hidden from an adversary $\mathcal{A}$. Otherwise, $\mathcal{A}$ can trivially break the q -SDH assumption and the security of the applications using it. One mechanism to generate such q -SDH parameters is to rely on a trusted third party. Specifically, a trusted party locally samples a random τ , computes and publishes $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$, and then deletes τ . However, this approach introduces a central trusted party and a single point of failure, which is undesirable.

This paper studies the problem of distributed secure and robust generation of q -SDH parameters in an asynchronous network.

Existing works. Existing protocols [10], [41], [14], [46] for generating q -SDH parameters follow the blueprint of [10]. These protocols assume a synchronous network plus a Byzantine broadcast channel [43] and proceed in a round-robin manner. Briefly, in these protocols, parties take turns to update existing q -SDH parameters with a randomly chosen value and broadcast the updated parameters to all other parties. Once every party updates the q -SDH parameters with its private randomness, the final output is obtained. Intuitively, as long as one honest party updates the parameter and the protocol terminates, the trapdoor τ remains hidden from the adversary.

Assuming the network is indeed synchronous and the existence of a broadcast channel (or blockchain) with external validity, the above protocols have several advantages. First, the q -SDH parameters remain hidden from the adversary as long as one honest party has contributed to it. Second, the set of participants does not have to be fixed in advance, i.e., a community can potentially decide on the fly to bring in more parties to contribute to the parameter generation.

However, without assuming a broadcast channel, an honest majority of parties are still needed in certain scenarios. More specifically, we would like to distinguish two formulations of the powers-of-tau: (i) all honest parties generate the q -SDH parameters to use among themselves, (ii) any external party can be convinced to use the generated q -SDH parameters. In formulation (i), the fault tolerance is $n - 1$. However, such formulation has limited practicality since all parties that require q -SDH parameters must participate in the entire powers-of-tau protocol. On the contrary, formulation (ii) is more suitable for generating the setup for any external entities but has a fault-tolerance of $n/2$; otherwise, an adversary corrupting

Table I: Comparison of protocols for generating secure q -SDH parameters. All of these protocols (including ours) generate updatable parameters in the sense of [44]. We provide a detailed breakdown of our cost in Table III. We measure the computation cost as the number of elliptic curve group multiplications.

	Network Model	Fault Tolerance	Communication Cost (Per Party)	Computation Cost (Per Party)	Total Round Complexity	Setup Assumption
[10] [*] , [41], [14], [46]	sync.	$n - 1$ [♣]	$\Omega(nq)$ [§]	$O(nq)$ [‡]	$\Omega(n)$	CRS, PKI [†]
This work	async.	$n/3$	$O(q + n^2 \log q)$	$O(q \log n + n^2 \log q)$ [‡]	$O(\log n + \log q)$	CRS PKI, RO

^{*} [10], describe an protocol to generate CRS of the form $C(\alpha)g$ for some random field element α and a circuit C from a family of F-arithmetic circuits. Generating q -SDH parameters is a special case of [10].

[†] Existing protocols require PKI to implement a broadcast channel.

[‡] State-of-the-art synchronous protocols [46] and our protocol require

parties to perform $O(n)$ and $O(n \log q)$ bilinear pairing, respectively.

[♣] The fault-tolerance $n - 1$ assumes a broadcast channel with external validity.

[§] The $O(nq)$ per party communication cost is a lower bound as each party needs to receive a message of length $O(nq)$.

the majority parties can generate identically distributed q -SDH parameters without using contributions from any honest party. Indeed, all real-world deployments [1], [45], [2], [32], [3] for generating q -SDH parameters assume an underlying blockchain (such as Ethereum) which requires an honest majority.

Synchrony assumption. The above protocols also have some significant drawbacks, most notably, the reliance on the network synchrony assumption. Briefly, if a party experiences temporary network asynchrony, then other parties will *time out*, skip that party, and move on. The fault tolerance is thus reduced. If the network remains asynchronous for an extended period of time, all honest parties may be skipped, and the trapdoor τ will be fully controlled and known to the adversary, as only adversarial parties will contribute randomness to the final output.

Inefficiency and insecure deployments. Existing protocols are inefficient and inherently sequential. They require running $O(n)$ sequential Byzantine broadcasts, once by each party, where the party needs to send $O(q)$ group elements through a costly broadcast channel. Each party must wait for all previous parties to update the q -SDH parameters. Moreover, each party also needs to verify updates by all previous parties before applying its own update.

To mitigate the inefficiencies, actual deployments of these protocols often cut corners on robustness or security. For example, one deployed version relies on a single party to act as a broadcast channel [3]. Moreover, in most deployed versions, parties skip verifying other parties' updates during the protocol and only verify the entire protocol transcript at the end. A consequence is that a single malicious party can now make the protocol produce invalid parameters by performing an invalid update. If that happens, the only recourse is to restart the entire protocol.

Even after cutting these corners on security and robustness, existing protocols perform poorly. For example, according to [32], to generate q -SDH parameters for $q = 2^{28}$, each party needs to perform 24 hours of computation. Hence, with n parties, the protocol would run for n days.

One might believe that a slow and costly q -SDH parameter generation protocol is acceptable because it needs to be run only once and can then be reused across all applications by all organizations. This is not always the case. One reason is that parameters generated by a set of parties may not be trusted by another set of parties. For example, two startups Aztec [2]

and Semaphore [3] repeated the powers-of-tau ceremony for the same elliptic curve BN254. Moreover, q -SDH parameters are specific to an elliptic curve group and thus must be generated from scratch if a system decides to adopt a new curve. For example, very recently, Ethereum [30] ran a powers-of-tau ceremony for the BLS12381 elliptic curve, although they have already run a powers-of-tau for the BN254 curve before [2].

One might try to adapt the existing synchronous protocols [10], [41], [14], [46] to asynchrony, where parties take turns to update the existing q -SDH parameters with their private randomness and then broadcast the result. We observe inherent difficulties in adapting these round-robin approaches to asynchrony for the following reasons. First, it is impossible to implement a Byzantine broadcast channel in asynchrony. Intuitively, it is impossible to distinguish a malicious broadcaster from an honest broadcaster with a slow network. The standard broadcast definition under asynchrony, namely reliable broadcast [15], does not require honest parties to output in case of a malicious sender, which is clearly insufficient for the previous constructions. Even for a stronger asynchronous broadcast notion named atomic broadcast [17], [16], it is impossible to output parties' contributions in a round-robin fashion due to the same reason that asynchrony makes it impossible to distinguish malice from slowness.

Our contributions. In this paper, we present the first protocol for the distributed generation of q -SDH parameters in asynchronous networks. Our protocol can tolerate up to t malicious parties out of $n \geq 3t + 1$ parties. Since the powers-of-tau output are agreed upon, implying consensus, one-third fault tolerance is optimal under asynchrony [29]. We provide a detailed summary of the properties of our protocol in Table I. The protocol finishes in $O(\log q + \log n)$ rounds. The per-party communication cost is $O(q + n^2 \log q)$ group elements, which improves the communication cost of prior best synchronous protocols by a factor $O(n)$. We also improve the per-party computation cost to $O(q \log n)$ group multiplications and $O(n \log q)$ pairings. Unlike existing protocols, our protocol is also *responsive*, i.e., it makes progress at a rate of actual network speed.

Evaluation. We implement our protocol in python with rust for cryptographic operations. Our implementation is *single-threaded*, supports the BLS12381 elliptic curve, and is publicly available^{*}. We evaluate our protocol with a network of up

^{*}<https://github.com/sourav1547/qsdh-py>

Table II: Notations used in the paper

Notation	Description
n	Total number of parties
t	Maximum number of malicious parties
κ	Security parameter
$\mathbb{G}$	Group of order p with hard Discrete Logarithm
$\mathbb{F}$	Scalar field of group $\mathbb{G}$
$\mathbf{g}, \mathbf{h}$	Random and independent generators of $\mathbb{G}$
q	Maximum degree of the q -SDH parameters
τ	q -SDH parameter trapdoor
$\llbracket z \rrbracket$	$(n, t + 1)$ Shamir secret sharing of z
$\llbracket z \rrbracket^{2t}$	$(n, 2t + 1)$ Shamir secret sharing of z
$\llbracket z \rrbracket_i, \llbracket z \rrbracket_i^{2t}$	Shares received by party i
$\llbracket z \rrbracket_{\mathbf{g}}$	The set $\{\llbracket z \rrbracket_{1\mathbf{g}}, \llbracket z \rrbracket_{2\mathbf{g}}, \dots, \llbracket z \rrbracket_{n\mathbf{g}}\}$
$[a, b]$	The set $\{a, a + 1, a + 2, \dots, b\}$
$[a]$	The set $\{1, 2, \dots, a\}$

to 128 geographically distributed parties. We also compare our protocol with the state-of-the-art synchronous protocol as the baseline. We only implement the computation part of the baseline and assume its networking to be instantaneous. Our evaluation illustrates that our protocol gives 2-6 $\times$ faster runtime and 4-13 $\times$ better per-party bandwidth usage. For example, with $n = 128$ and $q = 2^{14}$, our protocol takes 515 seconds to generate the q -SDH parameters, whereas the existing protocol takes at least 1805 seconds (3.5 $\times$). Similarly, in the same experiment, each party in our protocol needs to send 118.17 Megabytes of data, compared to 1536 Megabytes of data (13 $\times$) in the existing protocol.

Paper organization. We give notations, the system model, and an overview of our protocol in §II. In §III, we discuss the required preliminaries. We provide a detailed description of our protocol in §IV and §V. We analyze the correctness and security of our protocol in §VI. We present implementation and evaluation details in §VII. We discuss related works in §VIII and conclude with a discussion and open problems in §IX.

II. SYSTEM MODEL AND OVERVIEW

A. Notations and System Model

We use κ to denote the security parameter. For example, when we use a collision-resistant hash function, κ denotes the size of the hash function output. We use $|S|$ to denote the size of a set S . Let $\mathbb{G}$ be a group of prime order p with scalar field $\mathbb{F}$. We will be using the additive notation for elliptic curve operations. For any $x \in \mathbb{F}$ and any group element $\mathbf{g} \in \mathbb{G}$, we use $x\mathbf{g}$ to denote the group operations repeated x times. For any integer a , we use $[a]$ to denote the ordered set $\{1, 2, \dots, a\}$. Also, for two integers a and b where $a < b$, we use $[a, b]$ to denote the ordered set $\{a, a + 1, \dots, b\}$.

For any $x \in \mathbb{F}$, we use $\llbracket x \rrbracket$ to denote the $(n, t + 1)$ secret sharing of x , i.e., x is secret shared using a polynomial of degree t . Also, we use $\llbracket x \rrbracket_i$ to denote the share held by party i , and $\llbracket x \rrbracket_0$ to denote x . For a vector $\mathbf{x}$, we use $\llbracket \mathbf{x} \rrbracket$ to denote the element-wise secret sharing of $\mathbf{x}$. Similarly, we use $\llbracket \mathbf{x} \rrbracket_i$ to denote the element-wise share of $\mathbf{x}$ held by party i .

Threat model and network assumption. We consider a network of n parties where every pair of parties are connected via a pairwise authenticated channel. We consider the presence

of a probabilistic polynomial time (PPT) malicious adversary $\mathcal{A}$ that can corrupt up to t out of the $n \geq 3t + 1$ parties in the network. We assume the network is asynchronous, i.e., $\mathcal{A}$ can arbitrarily delay any message but must eventually deliver all messages sent between honest parties.

State-of-the-art solutions to two building blocks of our protocol, specifically, asynchronous distributed key generation and random double sharing, assume the existence of a public key infrastructure (PKI) for efficiency [26]. Both building blocks can be instantiated without PKI at higher costs [42], so the PKI assumption can be removed at the cost of lowering the efficiency of the protocol if the application calls for it.

B. Problem Definition

The q -Strong Diffie-Hellmann (q -SDH for short) refers to the cryptographic intractability problem defined below.

Definition 1 (q -SDH Hardness): Let κ be the security parameter. Let $\mathbb{G}$ and $\mathbb{F}$ be a group and field of size exponential in κ , respectively. Let $\mathbf{g} \leftarrow \text{Gen}(1^\kappa)$ be a uniform random generator of group $\mathbb{G}$ and $\tau \in \mathbb{F}$ be a trapdoor. For any given q , which is polynomially bounded in κ , given the vector $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$, q -SDH is assumed to be hard with respect to this vector if the following probability is negligible for all PPT adversary $\mathcal{A}$

$$\Pr[(c, (\tau + c)^{-1}\mathbf{g}), c \in \mathbb{F} \setminus \{-\tau\} : \mathcal{A}(\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}); \mathbf{g} \leftarrow \text{Gen}(1^\kappa), \tau \xleftarrow{\$} \mathbb{F}] \quad (1)$$

The vector $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$ is referred to as the q -SDH parameters.

The goal of this paper is to design a distributed protocol to implement $\mathcal{F}_{\text{qSDH}}$, as defined in Figure 2, i.e., to generate $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$ for $\tau \in \mathbb{F}$ at all honest parties given a uniformly random generator $\mathbf{g} \in \mathbb{G}$ as the common random string (CRS). The protocol is called t -secure for $t < n/3$ if the following *Correctness* and *Secrecy* properties hold in the presence of an adversary $\mathcal{A}$ that corrupts up to t parties.

- **Correctness.** If all honest parties start the protocol, every honest party will eventually terminate and output an identical vector $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$ for some $\tau \in \mathbb{F}$.
- **Security.** For every PPT adversary $\mathcal{A}$ that corrupts up to t parties, there exists a PPT simulator S_{qSDH} such that on input $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$, the q -SDH parameters output by the ideal functionality $\mathcal{F}_{\text{qSDH}}$, produces a view which is identical to the $\mathcal{A}$'s view of a run in the real protocol and ends with $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$ as the q -SDH parameters.

Our security definition immediately implies that if $\mathcal{A}$ can break the q -SDH hardness (definition 1) with respect to the parameters generated by our protocol, we can design a reduction adversary, using S_{qSDH} , that can use $\mathcal{A}$ to break the q -SDH assumption for parameters generated by $\mathcal{F}_{\text{qSDH}}$.

Remark. We want to note that since our protocol assumes $t < n/3$, the parties can trivially convince any external entity about the correct q -SDH parameters. One such approach using signatures is that the client only accepts q -SDH parameters that are signed by more than $1/3$ of the parties. It is also possible to achieve such guarantees without signatures using techniques from [27].




Figure 1: Overview of our protocol in a network of 4 parties where party 4 is malicious.

Functionality $\mathcal{F}_{q\text{SDH}}$

- Let $\mathbb{G}$ be an elliptic curve group with scalar field $\mathbb{F}$. Let g be a uniformly random generator of $\mathbb{G}$.
- Sample a uniformly random element $\tau \in \mathbb{F}$. Compute $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$ and send it to all parties.

Figure 2: The ideal functionality for generating q -SDH parameters

C. Overview of Our Protocol

One approach to generate q -SDH parameters in asynchrony is to use a generic asynchronous secure multiparty computation (MPC) protocol for a circuit $\mathcal{C}$ that outputs q -SDH parameters. However, this approach is very costly, primarily for the following reasons. First, the circuit $\mathcal{C}$ consists of $O(q)$ multiplication gates, which the parties need to evaluate using MPC multiplication units, such as multiplication triples or random double sharings. The best-known concretely efficient protocol for generating multiplication units in asynchrony has a per-party per-triple communication cost of $O(n^2)$ [26]. Thus, a generic MPC approach will result in a protocol with a per-party communication cost of $O(n^2 q)$, which is prohibitively expensive. Although asymptotically, these costs might be reduced using threshold FHE or threshold additively homomorphic encryptions. But that will require asynchronous distributed key generation (ADKG) protocol for such primitives, whose concretely efficient constructions are unknown. Second, the circuit $\mathcal{C}$ needs to resolve the discrepancy between the scalar and the base field of the underlying elliptic curve. We elaborate on this in §VIII.

This paper presents a new approach to distributed q -SDH parameter generation. Our first main idea is to securely compute τg for a uniformly random $\tau \in \mathbb{F}$. Moreover, we want to have τ secret shared among the parties using a $(n, t + 1)$ Shamir secret sharing. We will then find a way to use the public value τg and the secret shares of τ to compute the remaining powers of τ efficiently. We will first describe a naïve method that is incomplete and inefficient but demonstrates a core idea in our final protocol.

Naïve approach. Let $[\tau]_i$ be the secret share of party i . The protocol proceeds in rounds, where in the k -th round, parties generate $\tau^{k+1} g$ using $\tau^k g$. Thus, at the end of round k , parties generate the parameters $\{g, \tau g, \tau^2 g, \dots, \tau^{k+1} g\}$. At the start of k -th round, each party i locally computes $[\tau]_i \tau^k g$ and then multicasts $[\tau]_i \tau^k g$ to all parties. Also, party i upon receiving

$[\tau]_j \tau^k g$ from $t + 1$ distinct parties, computes $\tau^{k+1} g$ as

$$\tau^{k+1} g = \sum_{i \in T} \lambda_i [\tau]_i \tau^k g \quad (2)$$

where λ_i is the appropriate Lagrange coefficient. A approach similar to the above was used to generate q -Bilinear Diffie Hellman Exponentiation (q -BDHE) parameters [40] in a distributed manner.

This approach, however, has two major issues. First, the protocol is not robust, as a malicious party can send different and inconsistent messages to different parties, violating Correctness. Second, the protocol is very inefficient. It requires $O(q)$ rounds of interaction and $O(nq)$ per-party communication. Since q can be quite large in practice, say millions, this naïve approach is impractical.

Ensuring Correctness. Addressing the Correctness issue is relatively standard. Specifically, we need a mechanism for honest parties to validate the messages they receive from other parties. To achieve this, at the start of the protocol, we require that parties additionally hold $[\tau]g$, i.e., the vector $[[\tau]_1 g, [\tau]_2 g, \dots, [\tau]_n g]$. Then, each party i , when sending $[\tau]_i \tau^k g$, attaches a non-interactive zero knowledge (NIZK) proof π_i proving that $[\tau]_i \tau^k g$ is correctly computed from $\tau^k g$ and $[\tau]_i g$.

For this plan to work, the next natural question is how we establish the initial condition that parties agree on group elements τg and $[\tau]g$, and also hold their individual secret shares $[\tau]_i$. Our second observation is that the above initial condition exactly matches the output of an Asynchronous Distributed Key Generation (ADKG) protocol. An ADKG protocol generates a uniformly random secret key $\tau \in \mathbb{F}$ where each party i receives its share $[\tau]_i$ of the secret key, the public key τg and the threshold public keys $[\tau]g$.

Reducing round complexity. One approach to reducing the round complexity is once again generic secure multiparty computation (MPC). Let $\mathcal{C}$ be the arithmetic circuit that outputs all powers of τ in the exponent. Using repeated squaring, the depth of $\mathcal{C}$ will be $O(\log q)$. However, as we mentioned earlier, generic MPC has prohibitive communication costs.

We give an efficient method to reduce the round complexity to $O(\log q)$ by combining the idea of the naïve approach and the MPC approach. Our method matches the round complexity of the generic MPC without incurring high communication costs. At a high level, parties first use a customized MPC protocol to compute $\tau^{2^k} g$ and secret shares of $[\tau^{2^k}]$ for

each $k \in [\log q]$ (referred to as *powers-of-two*). For each $k \in [\log q]$, parties additionally output $\llbracket \tau^{2^k} \rrbracket \mathbf{g}$. Parties then use these values to efficiently compute the remaining powers using only $O(\log q)$ rounds of interaction.

Computing powers-of-two. Let $\mathcal{F}_{\text{sq}}$ be a secure MPC functionality for squaring defined as follows. $\mathcal{F}_{\text{sq}}$ takes as input the secret sharing of a , i.e., each party i inputs $\llbracket a \rrbracket_i$, and a publicly available list of group elements $\llbracket a \rrbracket \mathbf{g}$. $\mathcal{F}_{\text{sq}}$ then outputs to party i secret sharing of a^2 , i.e., $\llbracket a^2 \rrbracket_i$, and additionally outputs $\llbracket a^2 \rrbracket \mathbf{g}$ to all parties. We can then invoke $\mathcal{F}_{\text{sq}}$ sequentially $\log q$ times to compute the powers-of-two values.

Computing remaining powers. We generate the remaining powers of τ using the following ideas. We observe that given $\tau^\alpha \mathbf{g}$, secret shares of τ^β (i.e., $\llbracket \tau^\beta \rrbracket$), and public values $\llbracket \tau^\beta \rrbracket \mathbf{g}$, parties can compute $\tau^{\alpha+\beta} \mathbf{g}$, using a generalization of our naïve approach. Now consider any $a \in [q]$, we can write a as a sum of a subset of elements in $\{1, 2, 2^2, 2^3, \dots, 2^{\log(q)-1}\}$ according to its binary representation. Let S_a be the subset. Then, $\tau^a \mathbf{g} = \tau^{\sum_{k \in S_a} 2^k} \mathbf{g}$. It is easy to see that, using the idea of multiplication in the exponent, parties can compute the subset sum using $|S_a| \leq O(\log q)$ multiplications. Furthermore, parties can compute $\tau^a \mathbf{g}$ for every $a \in [q]$ in parallel.

Further optimizations. The method above incurs a per-party communication cost of $O(nq \log q)$ and a per-party computation cost of $O(nq \log q)$ group multiplications. In §V we discuss how we reduce the communication cost to $O(q + n^2 \log q)$ and the computation cost to $O(q \log n)$ group multiplications and $O(n \log q)$ bilinear pairings using *memoization* and *batching*.

III. PRELIMINARIES

A. Threshold Secret Sharing

A (n, k) threshold secret sharing scheme allows a secret $s \in \mathbb{F}$ to be shared among n parties such that any k of them can come together to recover the original secret, but any subset of $k - 1$ shares does not reveal any information about the secret [49], [12]. We use the common Shamir secret sharing [49] scheme, where the secret is embedded in a random degree $k - 1$ polynomial in the field $\mathbb{F}$. Specifically, to share a secret $s \in \mathbb{F}$, a polynomial $p(\cdot)$ of degree $k - 1$ is chosen such that $s = p(0)$. The remaining coefficients of $p(\cdot)$, $p_1, p_2, \dots, p_{k-1}$ are chosen uniformly randomly from $\mathbb{F}$. The resulting polynomial $p(x)$ is defined as:

$$p(x) = s + p_1 x + p_2 x^2 + \dots + p_{k-1} x^{k-1}$$

Each party is then given a single evaluation of $p(\cdot)$ evaluation. In particular, the i^{th} party is given $p(i)$, i.e., the polynomial evaluated at i . Observe that given k points on the polynomial $p(\cdot)$, one can efficiently reconstruct the polynomial using Lagrange Interpolation. Also, s is information-theoretically hidden from an adversary that knows any subset of $k - 1$ or fewer evaluation points on the polynomial other than $p(0)$.

B. Asynchronous Distributed Key Generation

Our protocol uses asynchronous distributed key generation (ADKG) functionality $\mathcal{F}_{\text{ADKG}}$, defined in Figure 10. Concretely, we run ADKG protocol from [26]. At the end of the ADKG protocol, parties output a $(n, t + 1)$ Shamir secret

sharing of a random value τ (i.e., $\llbracket \tau \rrbracket$), the ADKG public key $\tau \mathbf{g}$, and threshold public keys of every party, i.e., $\llbracket \tau \rrbracket \mathbf{g}$. The ADKG protocol of [26] assumes the hardness of Discrete Logarithm, has per-party communication cost of $O(n^2)$, and terminates in expected $O(\log n)$ rounds.

C. Asynchronous Double Sharing of Random Values

Our realization of $\mathcal{F}_{\text{sq}}$ uses double sharing of uniformly random field elements [23]. Specifically, we will use secret shares of a random field element z with both degree t and degree $2t$ polynomials, denoted as $\llbracket z \rrbracket$ and $\llbracket z \rrbracket^{2t}$, respectively. Here z is uniformly random and independent of τ .

Looking ahead, each invocation of $\mathcal{F}_{\text{sq}}$ will use double sharing of one random field element. Since we invoke $\mathcal{F}_{\text{sq}}$ $\log q$ times, we need $\log q$ double sharing of independent random field elements. Our realization of $\mathcal{F}_{\text{sq}}$ additionally require publicly available $\llbracket z \rrbracket \mathbf{g}$ and $\llbracket z \rrbracket^{2t} \mathbf{g}$ to publicly reconstruct $\llbracket \tau^{2^k} \rrbracket \mathbf{g}$ for each $k \in [\log q]$.

To generate double sharing of $\log q$ random elements and their corresponding public keys, we use the functionality $\mathcal{F}_{\text{Dou}}$ defined in Figure 8. We use the random double sharing protocol of [26, §6.1] but make minor modifications to facilitate a simulation-based security proof (described in Appendix C). Our modifications maintain their $O(n^2)$ per-party communication cost and $O(\log n)$ round complexity.

D. Bilinear Pairing

Definition 2 (Bilinear Pairing): Let $\mathbb{G}_1, \mathbb{G}_2$ and $\mathbb{G}_T$ be three cyclic groups of prime order p where $\mathbf{g}_1 \in \mathbb{G}_1$ and $\mathbf{g}_2 \in \mathbb{G}_2$ are generators. A pairing is an efficiently computable function $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ satisfying the following properties.

- 1) bilinear: For all $u, u' \in \mathbb{G}_1$ and $v, v' \in \mathbb{G}_2$ we have

$$\begin{aligned} e(u \cdot u', v) &= e(u, v) \cdot e(u', v), \text{ and} \\ e(u, v \cdot v') &= e(u, v) \cdot e(u, v') \end{aligned}$$

- 2) non-degenerate: $\mathbf{g}_T := e(\mathbf{g}_1, \mathbf{g}_2)$ is a generator of $\mathbb{G}_T$.

We refer to $\mathbb{G}_1$ and $\mathbb{G}_2$ as the pairing groups or source groups, and refer to $\mathbb{G}_T$ as the target group.

E. Equality of Discrete Logarithm

Our protocol requires parties to produce zero-knowledge proofs about the equality of discrete logarithms for a tuple of publicly known values. In particular, given a group $\mathbb{G}$ with scalar field $\mathbb{F}$ of prime order p , two elements $\mathbf{g}, \mathbf{h} \in \mathbb{G}$ with unknown discrete logarithm relations and a tuple $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b}) \in \mathbb{G}^4$, a prover $\mathcal{P}$ wants to prove to a probabilistic polynomial time verifier $\mathcal{V}$, in zero-knowledge, the knowledge of a witness $\alpha \in \mathbb{F}$ such that $\mathbf{a} = \alpha \mathbf{g}$ and $\mathbf{b} = \alpha \mathbf{h}$.

We use the Chaum-Pedersen "Σ-protocol" [18] (see Appendix B), which assumes the hardness of the Discrete Logarithm in $\mathbb{G}$. This protocol guarantees completeness, knowledge soundness, and zero knowledge. The knowledge soundness implies that if $\mathcal{P}$ convinces the $\mathcal{V}$ with non-negligible probability, there exists an efficient (polynomial time) extractor that can extract α from $\mathcal{P}$ non-negligible probability.

Algorithm 1 Our protocol Π_{qSDH} for party i

INPUT: $\mathbf{g}, sk_i, \{pk_j\}$ for each $j \in [n]$ OUTPUT: $\{\mathbf{g}, \tau\mathbf{g}, \tau^2\mathbf{g}, \dots, \tau^q\mathbf{g}\}$

SETUP PHASE:

// Run the ADKG and the Double sharing protocol

11: Let $\tau\mathbf{g}, \llbracket\tau\rrbracket_i \leftarrow \Pi_{\text{ADKG}}()$ 12: Let $\{\llbracket z_k \rrbracket \mathbf{g}, \llbracket z_k \rrbracket_i^{2t} \mathbf{g}, \llbracket z_k \rrbracket_i, \llbracket z_k \rrbracket_i^{2t}\}_{\forall k \in [\log q]} \leftarrow \Pi_{\text{Dou}}(\log q)$

POWERS-OF-TWO PHASE:

21: **for** each $k \in [\log q]$ **do**22: // Use double sharing of z_k to run Π_{sq} 23: Let $\llbracket\tau^{2^k}\rrbracket; \llbracket\tau^{2^k}\rrbracket \mathbf{g} = \Pi_{\text{sq}}(\llbracket\tau^{2^{k-1}}\rrbracket; \llbracket\tau^{2^{k-1}}\rrbracket \mathbf{g})$ 24: Compute $\tau^{2^k}\mathbf{g}$ by interpolating $\llbracket\tau^{2^k}\rrbracket \mathbf{g}$ and **output** $\tau^{2^k}\mathbf{g}$

ALL POWERS PHASE:

31: Compute all remaining powers as

 $\{\tau^a\mathbf{g}\}_{\forall a \in [q]} := \Pi_{\text{all}}(\{\llbracket\tau^{2^k}\rrbracket; \tau^{2^k}\mathbf{g}, \llbracket\tau^{2^k}\rrbracket \mathbf{g}\}_{\forall k \in [\log q]})$

The Chaum-Pedersen protocol can be made non-interactive in the Random Oracle model using the Fiat-Shamir heuristic [31]. We use the non-interactive variant of the protocol. For any given tuple $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b}) \in \mathbb{G}^4$ where $\mathbf{a} = \alpha\mathbf{g}$ and $\mathbf{b} = \alpha\mathbf{h}$, $\mathcal{P}$ uses $\text{dleq.Prove}(\alpha, \mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$ to generate the non-interactive zero-knowledge proof π . The proof π is $O(\kappa)$ bits long. Given a proof π and $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$, $\mathcal{V}$ uses the $\text{dleq.Verify}(\pi, \mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$ to verify the proof.

Remark. Alternatively, one could also use bilinear pairing to check the equality of discrete logarithms. The advantage of the bilinear pairing-based check is that it removes the need for the Random Oracle from our protocol, assuming the underlying ADKG does not require a Random Oracle. We use the Chaum-Pedersen “ Σ ”-protocol due to its better efficiency.

IV. GENERATING POWERS-OF-TWO

$\mathcal{F}_{\text{sq}}$ is the secure multiparty computation (MPC) functionality for squaring that takes as input $\llbracket a \rrbracket$, i.e., $(n, t+1)$ secret shares of a field element $a \in \mathbb{F}$ and outputs $\llbracket a^2 \rrbracket_i$ to party i . $\mathcal{F}_{\text{sq}}$ additionally takes the publicly available $\llbracket a \rrbracket \mathbf{g}$ as input and outputs threshold public keys of a^2 , i.e., $\llbracket a^2 \rrbracket \mathbf{g}$. Formally, we write the functionality $\mathcal{F}_{\text{sq}}$ as:

$$\llbracket a^2 \rrbracket; \llbracket a^2 \rrbracket \mathbf{g} = \mathcal{F}_{\text{sq}}(\llbracket a \rrbracket; \llbracket a \rrbracket \mathbf{g}) \quad (3)$$

Our protocol uses $\mathcal{F}_{\text{sq}}$ to compute $\tau^{2^k}\mathbf{g}$ for each $k \in [\log q]$. In this process, parties also receive secret shares of τ^{2^k} , i.e., $\llbracket\tau^{2^k}\rrbracket$, and $\llbracket\tau^{2^k}\rrbracket \mathbf{g}$, which they later use to compute $\tau^a\mathbf{g}$ for any arbitrary $a \in [q]$.

A. Design of $\mathcal{F}_{\text{sq}}$

We next describe our protocol Π_{sq} for realizing $\mathcal{F}_{\text{sq}}$. Its pseudocode is given in Algorithm 2. Π_{sq} could be designed using various techniques, such as multiplication triples [8], degree reduction [9], and random double sharing [24]. We adopt the random double sharing approach.

Π_{sq} assumes that parties hold double shares of a uniform random $z \in \mathbb{F}$ using $(n, t+1)$ and $(n, 2t+1)$ Shamir secret

Algorithm 2 Π_{sq} protocol at party i

INPUT: $\llbracket a \rrbracket_i$ and $\llbracket a \rrbracket \mathbf{g}$ SETUP: $\llbracket z \rrbracket_i, \llbracket z \rrbracket_i^{2t}, \llbracket z \rrbracket \mathbf{g}, \llbracket z \rrbracket^{2t} \mathbf{g}$ OUTPUT: $\llbracket a^2 \rrbracket_i$ and $\llbracket a^2 \rrbracket \mathbf{g}$

1: Let $a_i = \llbracket a \rrbracket_i \llbracket a \rrbracket_i + \llbracket z \rrbracket_i^{2t}$ 2: Let $\pi_i := \text{dleq.Prove}(\llbracket a \rrbracket_i, \mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g})$ 3: Send $\langle \text{SQ}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g}, a_i, \pi_i \rangle$ to all4: Let $K = \{\}$ 5: **upon** receiving $\langle \text{SQ}, \hat{a}_j \mathbf{g}, \tilde{a}_j, \pi_j \rangle$ from party j **do**6: Check π_j is a valid proof7: Check $\hat{a}_j \mathbf{g} + \llbracket z \rrbracket_j^{2t} \mathbf{g} = \tilde{a}_j \mathbf{g}$ 8: **if** both checks are successful **then**9: $K := K \cup \{j, \tilde{a}_j\}$ 10: **if** $|K| \geq 2t+1$ **then**11: Compute $a^2 + z$ by interpolating the values of K 12: Let $\llbracket a^2 \rrbracket_i := (a^2 + z) - \llbracket z \rrbracket_i$ 13: Let $\llbracket a^2 \rrbracket_j \mathbf{g} := (a^2 + z)\mathbf{g} - \llbracket z \rrbracket_j \mathbf{g}$ 14: Compute $\llbracket a^2 \rrbracket \mathbf{g}$ by interpolating $\llbracket a^2 \rrbracket_j \mathbf{g}$ for all $j \in K$.15: **output** $\llbracket a^2 \rrbracket_i$ and $\llbracket a^2 \rrbracket \mathbf{g}$

sharing, denoted as $\llbracket z \rrbracket_i$ and $\llbracket z \rrbracket_i^{2t}$, respectively. Also, Π_{sq} assumes that $\llbracket z \rrbracket \mathbf{g}$ and $\llbracket z \rrbracket^{2t} \mathbf{g}$ are public.

Each party i , locally multiplies its shares of a to get $\llbracket a \rrbracket_i \llbracket a \rrbracket_i$. Parties then publicly reconstruct the $a^2 + z$. In particular, each party i locally computes the non-interactive zero-knowledge (NIZK) proof π_i of equality of discrete logarithm (dleq) between $\{\mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g}\}$, i.e.,

$$\pi_i = \text{dleq.Prove}(\llbracket a \rrbracket_i, \mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g})$$

Party i then multicasts the message $\langle \text{SQ}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i + \llbracket z \rrbracket_i^{2t}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g}, \pi_i \rangle$ to every party. Upon receiving the message $\langle \text{SQ}, \tilde{a}_j, \mathbf{g}_j, \pi_j \rangle$, party i checks that $\tilde{a}_j$ is computed correctly, i.e.,

$$\text{dleq.Verify}(\mathbf{g}, \llbracket a \rrbracket_j \mathbf{g}, \llbracket a \rrbracket_j \mathbf{g}, \mathbf{g}_j, \pi_j) = 1; \text{ and} \\ \tilde{a}_j \mathbf{g} = \mathbf{g}_j + \llbracket z \rrbracket_j^{2t} \mathbf{g}$$

Upon receiving $2t+1$ valid SQ messages, parties reconstruct $a^2 + z$ by interpolating $\tilde{a}_j$, i.e.,

$$a^2 + z = \sum_j \lambda_j \tilde{a}_j \quad (4)$$

where λ_j is the appropriate Lagrange coefficients.

Upon reconstructing $a^2 + z$, party i computes its share of a^2 as $\llbracket a^2 \rrbracket_i = a^2 + z - \llbracket z \rrbracket_i$. Furthermore, for each $j \in [n]$, party i computes $\llbracket a^2 \rrbracket_j \mathbf{g}$ as:

$$\llbracket a^2 \rrbracket_j \mathbf{g} := (a^2 + z)\mathbf{g} - \llbracket z \rrbracket_j \mathbf{g} \quad (5)$$

In Appendix A, we prove that Π_{sq} securely realizes $\mathcal{F}_{\text{sq}}$ with a total communication cost of $O(n^2)$ per invocation.

B. Using $\mathcal{F}_{\text{sq}}$ for Generating Powers-of-two.

We now describe how parties use Π_{sq} to compute $\tau^{2^k}\mathbf{g}$ for every $k \in [\log q]$, i.e., the *powers-of-two*. While computing the powers-of-two, parties also output auxiliary values that they later use to compute the remaining powers.

Algorithm 3 Π_{all} protocol at party i

INPUT: $\{\llbracket \tau^{2^k} \rrbracket_i; \tau^{2^k} \mathbf{g}, \llbracket \tau^{2^k} \rrbracket \mathbf{g}\}$ for each $k \in [\log q]$

OUTPUT: $\{\tau^a \mathbf{g}\}$ for each $a \in [q]$

- 1: Create a binary tree with u_0^0 as its root and $\text{val}(u_0^0) = \tau^0 \mathbf{g}$.
 - 2: **for** each depth $d = 1, \dots, \log q$ **do**
 - 3: Create 2^d nodes labeled $u_0^d, u_1^d, \dots, u_{2^d-1}^d$
 - 4: Let $\{\text{val}(u_{2j}^d)\} := \{\text{val}(u_{2j-1}^{d-1})\}, \forall j \in [0, 2^{d-1} - 1]$.
 - 5: Let $\alpha := \tau^{2^{d-1}}$
 - 6: Let $\{\text{val}(u_{2j+1}^d)\} := \Pi_{\text{BatMul}}(\llbracket \alpha \rrbracket_i, \llbracket \alpha \rrbracket \mathbf{g}, \{\text{val}(u_{2j}^{d-1})\})$
 $\forall j \in [0, 2^{d-1} - 1]$.
 - 7: **output** $\text{val}(u_0^{\log q}), \text{val}(u_1^{\log q}), \dots, \text{val}(u_{q-1}^{\log q})$.
-

Parties start by running an ADKG protocol to secret share a uniformly random secret τ using a $(n, t + 1)$ Shamir secret sharing. In our implementation, we use the ADKG protocol from [26] with a reconstruction threshold of t . As a result, parties also output the *threshold* public keys $\llbracket \tau \rrbracket \mathbf{g}$.

While running the ADKG protocol, parties concurrently run Π_{Dou} , the protocol to generate random double sharing of $\log q$ uniform secrets $\{z_1, z_2, \dots, z_{\log q}\}$. Once the ADKG and random double sharing protocol terminate, parties compute powers-of-two by repeated invocations of the Π_{sq} protocol $\log q$ times in sequence, i.e.,

$$\begin{aligned} \llbracket \tau^2 \rrbracket; \llbracket \tau^2 \rrbracket \mathbf{g} &= \Pi_{\text{sq}}(\llbracket \tau \rrbracket; \llbracket \tau \rrbracket \mathbf{g}) \\ \llbracket \tau^4 \rrbracket; \llbracket \tau^4 \rrbracket \mathbf{g} &= \Pi_{\text{sq}}(\llbracket \tau^2 \rrbracket; \llbracket \tau^2 \rrbracket \mathbf{g}) \\ &\vdots \\ \llbracket \tau^{2^k} \rrbracket; \llbracket \tau^{2^k} \rrbracket \mathbf{g} &= \Pi_{\text{sq}}(\llbracket \tau^{2^{k-1}} \rrbracket; \llbracket \tau^{2^{k-1}} \rrbracket \mathbf{g}) \end{aligned}$$

V. GENERATING ALL POWERS

Using the protocol described in §IV, parties obtain secret shares of powers-of-two of τ , i.e., $\llbracket \tau, \tau^2, \tau^4, \dots, \tau^q \rrbracket$, as well as $\tau^{2^k} \mathbf{g}$ and $\llbracket \tau^{2^k} \rrbracket \mathbf{g}$ for each $k \in [\log q]$. In this section, we will describe how parties compute $\tau^a \mathbf{g}$ for all remaining $a \in [q]$. Formally, the interface of this functionality Π_{all} is:

$$\{\tau^a \mathbf{g}\}_{a \in [q]} = \Pi_{\text{all}}\left(\left\{\llbracket \tau^{2^k} \rrbracket; \tau^{2^k} \mathbf{g}, \llbracket \tau^{2^k} \rrbracket \mathbf{g}\right\}_{k \in [\log q]}\right)$$

A. Main Idea

Given the powers-of-two, for any given $a \in [q]$, we use the binary encoding of a to compute $\tau^a \mathbf{g}$. Let $a = \sum_{k=1}^{\log q} b_k 2^{k-1}$ with b_k being the k -th bit in the binary representation of a . Then, we can write $\tau^a \mathbf{g}$ as;

$$\tau^a \mathbf{g} = \left(\prod_{k \in [\log q]} \tau^{b_k 2^{k-1}} \right) \mathbf{g} \quad (6)$$

Next, we use the idea, referred to as the *multiplication in the exponent*, that given secret shares of τ^α , $\llbracket \tau^\alpha \rrbracket \mathbf{g}$, and $\tau^\beta \mathbf{g}$, parties can compute $\tau^{\alpha+\beta} \mathbf{g}$ using $O(n)$ per party communication cost and only one round of interaction. In particular, each party i , locally computes $\llbracket \tau^\alpha \rrbracket_i; \tau^\beta \mathbf{g}$ and the NIZK proof π_i of




Figure 3: The memoization protocol to compute all powers using a total $O(q)$ multiplications in the exponent.

its correctness. Party i then sends the tuple $\langle \llbracket \tau^\alpha \rrbracket_i; \tau^\beta \mathbf{g}, \pi_i \rangle$ to all parties. Also, upon receiving $\langle \mathbf{g}_j, \pi_j \rangle$ from party j , party i validates $\mathbf{g}_j$ for correctness using $\llbracket \tau^\alpha \rrbracket_i \mathbf{g}$, $\tau^\beta \mathbf{g}$ and π_j . Finally, upon receiving $t + 1$ such valid tuples T , it computes $\tau^{\alpha+\beta} \mathbf{g}$ as

$$\tau^{\alpha+\beta} \mathbf{g} = \tau^\alpha \tau^\beta \mathbf{g} = \sum_{i \in T} \lambda_i \mathbf{g}_i \quad (7)$$

Thus, for any given $a \in [q]$, parties can easily compute $\tau^a \mathbf{g}$ by repeating this technique for $\log q$ iterations, as per the bit representation of a and equation 6. This requires per-party communication cost of $O(n \log q)$ and $O(\log q)$ rounds of interactions. Hence, the per-party communication cost for computing $\tau^a \mathbf{g}$ for all $a \in [q]$ is $O(nq \log q)$ and each party performs $O(nq \log q)$ group multiplications.

We will next reduce the per-party communication cost to $O(nq)$ and computation cost to $O(nq)$ group multiplications using *memoization*. In §V-C, we will further reduce the communication cost to $O(q + n \log q)$ and computation cost to $O(q \log n)$ group multiplications and $O(n \log q)$ pairings using appropriate batching.

B. Memoization

We first use the standard memoization technique to reduce the per-party communication cost to $O(nq)$ and computation cost to $O(nq)$ group multiplications. Note that many of the $O(\log q)$ multiplications in the exponent are redundant. Consider two integers a, b that differ only in their most significant bit, i.e., $a = 0||s$ and $b = 1||s$ for some binary string s . Then, once we compute $\tau^a \mathbf{g}$ (which equals $\tau^s \mathbf{g}$), we can directly compute $\tau^b \mathbf{g}$ as $\tau^{2^{|s|}} \tau^a \mathbf{g}$, using only one more multiplication in the exponent.

We remove all such redundancies from our protocol based on a generalization of the above idea. Specifically, we create a binary tree of height $\log q$ where each node is associated with a binary string as illustrated in Figure 3. We place the binary representations of each $a = 0, 1, \dots, q$ at the leaves of the binary tree. Any path in the tree from the root r to a leaf node a consists of internal nodes $r = a_0, a_1, a_2, \dots$, and $a_{\log q} = a$. Stated differently, the left child and the right child of a node with the binary string s are binary strings $0||s$ and $1||s$, respectively. To each node in the tree with bit string s , we associate a value, $\text{val}(s) = \tau^s \mathbf{g}$. The root of the tree is initialized with the empty string and the value $\tau^0 \mathbf{g}$.

To compute $\tau^a \mathbf{g}$ for $a \in [q]$, we only need to traverse the tree from the root to the leaves (in parallel) and compute $\tau^x \mathbf{g}$ for each internal node with the binary string x . As mentioned, if a node x is the left child of a node s , then $\tau^x \mathbf{g} = \tau^s \mathbf{g}$;

Algorithm 4 Π_{BatMul} protocol for party i INPUT: $\mathbf{g}, [\alpha]_i, [\alpha]_i \mathbf{g}, \{\beta_1 \mathbf{g}, \beta_2 \mathbf{g}, \dots, \beta_m \mathbf{g}\}$ OUTPUT: $\{\alpha\beta_1 \mathbf{g}, \alpha\beta_2 \mathbf{g}, \dots, \alpha\beta_m \mathbf{g}\}$.

```

11: Let  $\ell = \lceil m/(n-t) \rceil$  be the number of batches.
12: Let  $\beta_1^{(1)}, \dots, \beta_{n-t}^{(1)}, \dots, \beta_1^{(\ell)}, \dots, \beta_{n-t}^{(\ell)}$  denote  $\beta_1, \beta_2, \dots, \beta_m$ 
13: for each batch  $k \in [\ell]$  do
14:   Let  $\beta^{(k)}(x) = \beta_1^{(k)} + \beta_2^{(k)}x + \dots + \beta_{n-t}^{(k)}x^{n-t-1}$ 
15:   Compute  $[\alpha]_i \beta^{(k)}(j) \mathbf{g}, \forall j \in [n]$  using NTT
  // Derive shared randomness using random oracle
16: Let  $\gamma_1, \gamma_2, \dots, \gamma_\ell = \text{RO}(\{\beta_1 \mathbf{g}, \beta_2 \mathbf{g}, \dots, \beta_m \mathbf{g}\})$ 
17: for each  $j \in [n]$  do
18:   Let  $\mathbf{a}_j := \sum_{k \in [\ell]} \gamma_k \beta^{(k)}(j)$ 
19:   Let  $\mathbf{b}_{i,j} = [\alpha]_i \mathbf{a}_j$ 
20:   Let  $\pi_{i,j} = \text{dleg.Prove}([\alpha]_i, \mathbf{g}, [\alpha]_i \mathbf{g}, \mathbf{a}_j, \mathbf{b}_{i,j})$   $\triangleright$  Batching
21: Send  $\langle \text{SHARE}, \{[\alpha]_i \beta^{(k)}(j) \mathbf{g}\}_{\forall k \in [\ell]}, \pi_{i,j} \rangle$  to party  $j$ 
22: Let  $T_k = \{\}$  for each  $k \in [\ell]$ 
23: upon receiving  $\langle \text{SHARE}, \{\tilde{\mathbf{g}}_j^{(k)}\}_{\forall k \in [\ell]}, \pi_{j,i} \rangle$  from party  $j$  do
24:   Let  $\tilde{\mathbf{b}}_{j,i} := \sum_{k \in [\ell]} \gamma_k \tilde{\mathbf{g}}_j^{(k)}$ 
25:   if  $\text{dleg.Verify}(\mathbf{g}, [\alpha]_j \mathbf{g}, \mathbf{a}_i, \tilde{\mathbf{b}}_{j,i}, \pi_{j,i})$  then
26:      $T_k := T_k \cup \{(j, \tilde{\mathbf{g}}_j^{(k)})\}$  for each  $k \in [\ell]$ 
27:   if  $|T_k| \geq t+1$  then
28:     Compute  $\alpha\beta^{(k)}(i) \mathbf{g}$  using Lagrange interpolation
29: Send  $\langle \text{EVAL}, \{\alpha\beta^{(k)}(i) \mathbf{g}\}_{\forall k \in [\ell]} \rangle$  to all parties
30: Let  $S_k = \{\}$  for each  $k \in [\ell]$ 
31: upon receiving  $\langle \text{EVAL}, \{\hat{\mathbf{g}}_j^{(k)}\}_{\forall k \in [\ell]} \rangle$  from party  $j$  do
32:   Locally sample  $\chi_k \in \mathbb{F}$  for each  $k \in [\ell]$ 
33:   if  $e(\sum_{k \in [\ell]} \chi_k \beta^{(k)}(j) \mathbf{g}, \alpha \mathbf{g}) = e(\sum_{k \in [\ell]} \chi_k \hat{\mathbf{g}}_j^{(k)}, \mathbf{g})$  then
34:      $S_k := S_k \cup \{(j, \hat{\mathbf{g}}_j^{(k)})\}$ , for each  $k \in [\ell]$ 
35:   if  $|S_k| \geq n-t, \forall k \in [\ell]$  then
36:     Compute  $\{\alpha\beta_a^{(k)} \mathbf{g}\}_{\forall a \in [n-t]}$  using NTT for  $k \in [\ell]$ 
37:   output  $\{\alpha\beta_a \mathbf{g}\}_{\forall a \in [m]}$ 

```

if x is the right child, then $\tau^x \mathbf{g} = \tau^{2^{|s|}} \tau^s \mathbf{g}$. Finally, the $\tau^a \mathbf{g}$ for $a \in [q]$ can be computed at all the leaves. Parties now compute one multiplication in the exponent per internal node of the tree, leading to a total of $O(q)$ multiplication in the exponent. The per-party communication and computation (in group multiplications) using this approach are both $O(nq)$.

C. Batched Multiplication in the Exponent

In this section, we will describe how to further reduce the communication cost of our protocol to $O(q + n \log q)$ using batching. Our batching uses an efficient solution to the following problem.

Problem definition and the naïve approach. We seek to design a protocol Π_{BatMul} where parties input $\{\beta_1 \mathbf{g}, \beta_2 \mathbf{g}, \dots, \beta_m \mathbf{g}\}$, $[\alpha]$, and publicly available $[\alpha] \mathbf{g}$, and each party outputs $\{\alpha\beta_1 \mathbf{g}, \alpha\beta_2 \mathbf{g}, \dots, \alpha\beta_m \mathbf{g}\}$. Formally,

$$\{\alpha\beta_k \mathbf{g}\}_{\forall k \in [m]} = \Pi_{\text{BatMul}}([\alpha]; [\alpha] \mathbf{g}, \{\beta_k \mathbf{g}\}_{\forall k \in [m]}) \quad (8)$$

One naïve approach is to compute $\alpha\beta_k \mathbf{g}$ for each $k \in [m]$ separately. This would result in a per-party communication cost of $O(nm)$ and a per-party computation cost of $O(nm)$ group multiplications.

Our approach. We will next describe protocol Π_{BatMul} , where each party incurs a communication cost of $O(n+m)$ and a per-party computation cost of $O(m \log n)$ group multiplications. It

does add a per-party computation cost of $O(n)$ (independent of m) bilinear pairing operations. We provide the pseudocode of Π_{BatMul} in Algorithm 4 and describe it next.

For simplicity, we will assume that $m = n - t$. For $m > n - t$, we can divide the inputs into batches of size $n - t$ and run Π_{BatMul} in parallel (with minor modifications) for each batch.

Let $\beta(\cdot)$ be the polynomial of degree $n - t - 1$ defined as

$$\beta(x) = \beta_1 + \beta_2 x + \beta_3 x^2 + \dots + \beta_{n-t} x^{n-t-1} \quad (9)$$

Given $\{\beta_k \mathbf{g}\}$ for $k \in [n - t]$, each party i locally computes $\beta(j) \mathbf{g}$ for each $j \in [n]$, i.e., evaluate the polynomial $\beta(\cdot)$ at all indices in $[n]$. Parties can compute these values using $O(n \log n)$ group multiplications using the Number Theoretic Transform (NTT). Additionally, party i locally computes $[\alpha] \beta(j) \mathbf{g}$ for each $j \in [n]$ along with the dleg proof $\pi_{i,j}$ given as

$$\pi_{i,j} = \text{dleg.Prove}([\alpha]_i, \mathbf{g}, [\alpha]_i \mathbf{g}, \beta(j) \mathbf{g}, [\alpha]_i \beta(j) \mathbf{g})$$

Party i then sends a message $\langle \text{SHARE}, [\alpha]_i \beta(j) \mathbf{g}, \pi_{i,j} \rangle$ to party j . Upon receiving $\langle \text{SHARE}, \mathbf{g}_j, \pi_{j,i} \rangle$ from party j , party i validates its correctness using dleg.Verify . Upon receiving $t + 1$ valid SHARE messages, party i computes $\alpha\beta(i) \mathbf{g}$ by interpolating in the exponent, i.e.,

$$\alpha\beta(i) \mathbf{g} = \sum_{j \in T} \lambda_j \mathbf{g}_j \quad (10)$$

where λ_j is the appropriate Lagrange coefficient.

Party i then sends a message $\langle \text{EVAL}, \alpha\beta(i) \mathbf{g} \rangle$ to all parties. Each party, upon receiving $\langle \text{EVAL}, \tilde{\mathbf{g}}_j \rangle$ from party j , validates its correctness by checking:

$$e(\beta(j) \mathbf{g}, \alpha \mathbf{g}) = e(\tilde{\mathbf{g}}_j, \mathbf{g}) \quad (11)$$

where $e(\cdot, \cdot)$ is the bilinear pairing operation. Note that every party can locally compute $\alpha \mathbf{g}$ and $\beta(j) \mathbf{g}$ using the inputs of Π_{BatMul} . Upon receiving $n - t$ valid EVAL messages, party i computes $\{\alpha\beta_1 \mathbf{g}, \alpha\beta_2 \mathbf{g}, \dots, \alpha\beta_{n-t} \mathbf{g}\}$ using inverse NTT.

Handling $m > n - t$. For $m > n - t$, parties divide the inputs into batches of size $n - t$ each. Let there be ℓ batches, i.e., $\ell = \lceil m/(n - t) \rceil$, where $\beta^{(k)}(\cdot)$ is the polynomial corresponding to the k -th batch. Parties then run Π_{BatMul} for each chunk in parallel with the following changes.

Batching dleg proofs of SHARE messages. Each party i , instead of sending ℓ dleg proofs to every party j , sends only one dleg proof $\pi_{i,j}$ attesting the correctness of SHARE messages for all batches by taking their random linear combination. More specifically, for each recipient j , party i computes $\mathbf{a}_j$ and $\mathbf{b}_{i,j}$ as

$$\mathbf{a}_j = \sum_{k \in [\ell]} \gamma_k \beta^{(k)}(j) \mathbf{g}; \text{ and } \mathbf{b}_{i,j} = [\alpha]_i \mathbf{a}_j \quad (12)$$

here γ_k are uniformly random elements in $\mathbb{F}$ generated by querying a random oracle on the input $\{\beta_1 \mathbf{g}, \beta_2 \mathbf{g}, \dots, \beta_m \mathbf{g}\}$. Let $\pi_{i,j}$ be the dleg proof given by

$$\pi_{i,j} = \text{dleg.Prove}([\alpha]_i, \mathbf{g}, [\alpha]_i \mathbf{g}, \mathbf{a}_j, \mathbf{b}_{i,j}) \quad (13)$$

Each party i then sends $\langle \text{SHARE}, \{[\alpha]_i \beta^{(k)}(j) \mathbf{g}\}_{k \in [\ell]}, \pi_{i,j} \rangle$ to party j as its SHARE message.

Party i , upon receiving $\langle \text{SHARE}, \{\tilde{\mathbf{g}}_j^{(k)}\}_{k \in [\ell]}, \pi_{i,j} \rangle$ message from party j , locally computes $\mathbf{a}_i$ using the publicly available information. Party i additionally computes $\mathbf{b}_{j,i}$ as:

$$\tilde{\mathbf{b}}_{j,i} = \sum_{k \in [\ell]} \gamma_k \tilde{\mathbf{g}}_j^{(k)}$$

Party i then validates the correctness of the SHARE message by checking that

$$\text{dlog.Verify}(\mathbf{g}, \llbracket \mathbf{a} \rrbracket_j \mathbf{g}, \mathbf{a}_i, \mathbf{b}_{j,i}, \pi_{j,i}) = 1 \quad (14)$$

Intuitively, in equation (13), we use the observation that the prover (party i) needs to prove the equality of discrete logarithm of different statements, one for each $\beta^{(k)}(j)\mathbf{g}$, that shares the same witness $\llbracket \alpha_i \rrbracket$. This enables us to batch all the statements into a single statement by taking their *random linear combination*. The check in equation (14) has a negligible error probability of $1/|\mathbb{F}|$ for each batch.

Batching checks of EVAL messages. Instead of checking equation (11) independently for every batch, parties combine them into a single check below:

$$e\left(\sum_{k \in [\ell]} \chi_k \beta^{(k)}(j)\mathbf{g}, \alpha \mathbf{g}\right) = e\left(\sum_{k \in [\ell]} \chi_k \tilde{\mathbf{g}}_j^{(k)}, \mathbf{g}\right) \quad (15)$$

Here, χ_k for each $k \in [\ell]$ are uniformly random field elements samples from $\mathbb{F}$.

Intuitively, in equation (15), party i , instead of individually checking the correctness of each EVAL message received from party j , takes a random linear combination of the values and checks them all at once. Similar to equation (14), the check in equation (15) also has a negligible error probability of $1/|\mathbb{F}|$ for each batch.

Analysis of Π_{BatMul} . For each batch of size $n - t$, each party sends a single SHARE and EVAL message to every other party. Hence, the per-party communication cost for a batch of size $n - t$ is $O(n)$. This implies that the per-party communication for a batch of size m is $O(n + m)$. Also, for each batch of size $n - t$, each party performs $O(n \log n)$ group multiplication due to NTT [51]. Hence, the total per-party computation cost for a batch of size m is $O(m \log n)$ group multiplications. Finally, due to the batch checking of all EVAL messages from a party, each party needs to perform only $O(n)$ pairing operations per batch. We reiterate that the number of pairing operations is independent of the batch size.

Using Π_{BatMul} to compute all powers. Recall from §V-B that, for any height h of the binary tree, we multiply with the identical τ^{2^h} to the values of every node at height h . Hence, our protocol computes them using Π_{BatMul} with $\alpha = \tau^{2^h}$ and the values at the nodes as $\{\beta_k \mathbf{g}\}$ for $k = 1, 2, \dots, 2^h$.

Since parties need to compute a total of q multiplications in the exponent, the total per-party communications cost is $O(q + n \log q)$. Also, each party performs $O(q \log n)$ group multiplications along with $O(n \log q)$ pairings in total.

D. Putting Things Together

In Algorithm 1, we present the full protocol for generating the q -SDH parameters. As mentioned in §III, the entire protocol consists of the Setup Phase where the parties run the ADKG protocol and the Double Sharing protocol, the Powers-of-two Phase where the parties run the MPC functionality $\mathcal{F}_{\text{sq}}$ to compute the powers-of-two parameters given the ADKG and Double Sharing outputs, and finally the All Powers Phase where the parties run the MPC functionality Π_{all} to generate the remaining powers-of- τ from the powers-of-two.

VI. ANALYSIS

We prove the security of our protocol Π_{qSDH} assuming Π_{sq} securely implements $\mathcal{F}_{\text{sq}}$. In Appendix A, we analyze the correctness and security of Π_{sq} .

A. Correctness

Lemma 1 (Correctness): If all honest parties start the protocol, then every honest party will output correct q -SDH parameters, i.e., there exists a $\tau \in \mathbb{F}$ such that parties output $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$, except with $\text{negl}(\kappa)$ probability.

Proof: We will prove correctness in $\mathcal{F}_{\text{ADKG}}$ and $\mathcal{F}_{\text{Dou}}$ hybrid model. $\mathcal{F}_{\text{ADKG}}$ guarantees that parties agree on a common public key $\tau \mathbf{g}, \llbracket \tau \rrbracket \mathbf{g}$, and each party i has $\llbracket \tau \rrbracket_i$. $\mathcal{F}_{\text{Dou}}$ guarantees that parties output double shares of $\log q$ random field elements $\{z_1, z_2, \dots, z_{\log q}\}$ along with $\llbracket z_k \rrbracket \mathbf{g}$ and $\llbracket z_k \rrbracket^{2t} \mathbf{g}$. With this setup, we will argue that all honest parties agree and output $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$.

In Appendix A, we prove that Π_{sq} securely realize $\mathcal{F}_{\text{sq}}$. This implies that during the squaring phase, parties output and agree on $\tau^{2^k} \mathbf{g}$ and $\llbracket \tau^{2^k} \rrbracket \mathbf{g}$ for each $k \in [\log q]$. Moreover, each party i outputs $\llbracket \tau^{2^k} \rrbracket_i$. Finally, while computing the remaining powers of τ , parties only accept valid SHARE and EVAL messages, except with probability $O(1/|\mathbb{F}|)$. Since $|\mathbb{F}|$ is super-polynomial in κ , the security parameter, this implies that all honest parties agree and output $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$, except with negligible probability. ■

B. Security

We prove the security of our protocol Π_{qSDH} using simulatability. Specifically, we prove that for every PPT static adversary $\mathcal{A}$ that corrupts up to t parties, there exists a PPT simulator $\mathcal{S}_{\text{qSDH}}$ such that on input $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$, the q -SDH parameters output by the ideal functionality $\mathcal{F}_{\text{qSDH}}$, produces a view which is identical to the $\mathcal{A}$'s view of a run in the real protocol and ends with $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$ as the q -SDH parameters. This immediately implies that if the q -SDH assumption holds for parameters generated by $\mathcal{F}_{\text{qSDH}}$, then the q -SDH assumption holds for parameters generated by Π_{qSDH} . We summarize the simulator $\mathcal{S}_{\text{qSDH}}$ in Figure 4 and describe it next.

Let $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$ be the parameters generated by $\mathcal{F}_{\text{qSDH}}$. $\mathcal{S}_{\text{qSDH}}$, upon receiving these parameters, simulates an execution of our protocol for $\mathcal{A}$ that outputs $\{\mathbf{g}, \tau \mathbf{g}, \tau^2 \mathbf{g}, \dots, \tau^q \mathbf{g}\}$ as the q -SDH parameters.

Let $\mathcal{C}$ be the set of parties corrupted by $\mathcal{A}$. For each $k \in [0, \log q]$, $\mathcal{S}_{\text{qSDH}}$ samples uniformly random shares $\llbracket \tau^{2^k} \rrbracket_i$ for

Simulator $\mathcal{S}_{qSDH}$	
Inputs. q -SDH parameters $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$	
1) Sample uniform random $[\tau]_i$ for each $i \in \mathcal{C}$ where $\mathcal{C}$ is the set of corrupted parties.	
2) For each $k \in [\log q]$, sample uniformly at random $[\tau^{2^k}]_i$ for each $i \in \mathcal{C}$.	
3) Run the input generation phase of $\mathcal{S}_{sq}$ on inputs $[\tau^{2^{k-1}}]_i, [\tau^{2^k}]_i$ for each $i \in \mathcal{C}$, and public values $[\tau^{2^{k-1}}]g, [\tau^{2^k}]g$ for every $k \in [\log q]$.	
4) Run the ADKG simulator $\mathcal{S}_{ADKG}$ on input $g, \tau g, [\tau]g$, and $\{[\tau]_i\}$ for each $i \in \mathcal{C}$.	
5) Run $\mathcal{S}_{Dou}$ to simulate a protocol to generate double sharing of $\log q$ values, with inputs computed during the input generation phase of $\mathcal{S}_{sq}$.	
6) While computing powers-of-two, run steps 2 and 3 of the simulation phase of $\mathcal{S}_{sq}$.	
7) During all powers phase, follow the honest protocol, except, whenever needed, generate the NIZK proof of equality of discrete logarithm using $\mathcal{S}_{dlog}$.	

Figure 4: Simulator for the protocol Π_{qSDH} simulating $\mathcal{F}_{qSDH}$

each $i \in \mathcal{C}$. $\mathcal{S}_{qSDH}$ then runs the input generation phase of $\mathcal{S}_{sq}$ on inputs $[\tau^{2^{k-1}}]_i, [\tau^{2^k}]_i$ for each $i \in \mathcal{C}$ and public values $[\tau^{2^{k-1}}]g$ and $[\tau^{2^k}]g$ for every $k \in [\log q]$. Intuitively, by doing so, $\mathcal{S}_{qSDH}$ generates the inputs of the $\mathcal{S}_{Dou}$.

$\mathcal{S}_{qSDH}$ then runs the ADKG simulator $\mathcal{S}_{ADKG}$ on input $g, \tau g$ and $[\tau]_i$ for each $i \in \mathcal{C}$. $\mathcal{S}_{ADKG}$ guarantees that parties output secret share of τ where adversarial shares matches the input to $\mathcal{S}_{ADKG}$. Simultaneously, $\mathcal{S}_{qSDH}$ runs $\mathcal{S}_{Dou}$ on inputs generated by the input generation phase of $\mathcal{S}_{sq}$. This concludes the simulation of the setup phase of our protocol. Next, while computing powers-of-two, $\mathcal{S}_{qSDH}$ simply runs steps (2) and (3) of $\mathcal{S}_{sq}$. Lastly, to compute all remaining powers of τ , $\mathcal{S}_{qSDH}$ follows the honest protocol, except, whenever needed, generates the NIZK proof of equality of discrete logarithm using the $\mathcal{S}_{dlog}$.

Remark. The fact that $\mathcal{A}$ outputs $\{g, \tau g, \tau^2 g, \dots, \tau^q g\}$ as a result of its interaction with $\mathcal{S}_{qSDH}$ and Lemma 2 immediately implies that if $\mathcal{A}$ outputs a valid tuple that breaks the q -SDH assumption with probability ε , $\mathcal{S}_{qSDH}$ breaks the q -SDH assumption with probability ε as well.

Lemma 2: For any PPT adversary $\mathcal{A}$ that corrupts up to t parties, the view of $\mathcal{A}$ during its interaction with $\mathcal{S}_{qSDH}$ is identically distributed to its view in the real protocol.

Proof: We will prove this by defining a series of hybrids, with hybrid 0 identical to the real protocol and hybrid 4 identical to the simulated protocol.

Hybrid 0. This corresponds to the real-world execution.

Hybrid 1. In this hybrid, we simulate the NIZK proofs of equality of discrete logarithms for each statement of the form $\{g, [a]_i g, [a]_i g, [a]_i [a]_i g\}$ in the all powers phase. Since the Chaum-Pedersen Σ protocol for proving the equality of discrete logarithm is perfect zero-knowledge, Hybrid 1 is identically distributed as Hybrid 0.

Hybrid 2. Sample a uniformly random polynomial $p(x)$ of degree t . Swap out the real execution of the ADKG protocol with the ADKG simulator $\mathcal{S}_{ADKG}$. The input to $\mathcal{S}_{ADKG}$ are $p(i)$

Table III: Cost of each phases in our protocol for any given n and q .

Protocol Phase	Communication (Per Party)	Computation (Per Party)	Expected Latency
Setup	$O(n^2 \log q)$	$O(n^2 \log q)$	$O(\log n)$
Powers-of-two	$O(n \log q)$	$O(n \log q)$	$O(\log q)$
All powers	$O(q + n \log q)$	$O(q \log n)$	$O(\log q)$
Overall	$O(q + n^2 \log q)$	$O(n^2 \log q + q \log n)$	$O(\log(nq))$

for all adversarial parties i and the commitments $p(j)g$ for all $j \in [n]$. The perfect simulatability of the $\mathcal{S}_{ADKG}$ guarantees that Hybrid 3 is identically distributed as Hybrid 2.

Hybrid 3. Replace the real execution of the Π_{sq} with the simulator $\mathcal{S}_{sq}$ with input consistent with $p(0)$ for the polynomial $p(x)$ sampled in Hybrid 2. The perfect simulatability of the $\mathcal{S}_{sq}$ guarantees that Hybrid 2 is identically distributed as Hybrid 1.

Hybrid 4. Change the ADKG public key part of the input to the $\mathcal{S}_{ADKG}$ to τg while keeping its shares of the adversarial parties the same. Compute the commitments to honest parties' shares by interpolating in the exponent. Accordingly, change the input to $\mathcal{S}_{sq}$ as per q -SDH parameters output by the $\mathcal{F}_{qSDH}$. Keep shares of malicious parties the same and update the commitment to honest parties' input. Since $\mathcal{F}_{qSDH}$ samples τ uniformly at random and Pedersen commitment is perfectly hiding, Hybrid 4 is identically distributed as Hybrid 3. ■

C. Performance

Round complexity. The ADKG protocol and the protocol for generating double sharing of random values takes expected $O(\log n)$ rounds of interaction [26]. Computing the powers-of-two takes $O(\log q)$ rounds. Finally, using memoization, computation of all remaining powers takes $O(\log q)$ additional rounds. Thus, the total expected round complexity of our protocol is $O(\log q + \log n)$.

Communication cost. The ADKG protocol generates secret shares of τ and publicly outputs τg and $[\tau]g$ with per party expected per-party communication cost of $O(n^2)$ [28]. The per-party communication cost of generating the double shares of $\log q$ random elements is $O(n^2 \log q)$. The communication cost of generating the powers-of-two is $O(n \log q)$ (ref. §IV). Finally, using memoization and batching (ref. §V), the communication cost of computing the remaining powers is $O(q + n \log q)$. Combining all the above, our protocol's total per-party communication cost is $O(q + n^2 \log q)$. When $q \gg n$, the total communication cost is $O(q)$, i.e., linear in the highest power of the q -SDH parameters. For example, with $n = 128$ and $q = 2^{20}$ i.e., about 1 Million, $q > n^2 \log q$.

Computation cost. We measure the computation cost as the number of group multiplications and pairing operations each party needs to perform in the entire protocol. During the ADKG protocol, each party performs $O(n^2)$ group multiplications [26]. Also, each party performs $O(n^2 \log q)$ group multiplications to compute the double-sharing of random values. While computing the powers-of-two, each party performs $O(n \log q)$ group multiplications in total. Finally, while computing all remaining powers of τ , for every height h in the memoization tree, each party performs $O(2^h \log n)$ group

multiplications and $O(n)$ pairings. This implies that the per-party computation cost in our protocol is $O(q \log n)$ group multiplications and $O(n \log q)$ pairings.

External verification. Let $\mathcal{V}$ be an external verifier that wishes to verify the correctness and security of the q -SDH parameters generated by our protocol. Since we assume PKI, this is relatively straightforward. Each party signs the output using its signing key and sends the signature to $\mathcal{V}$. The $\mathcal{V}$ waits for $t + 1$ valid signature on the matching output.

VII. IMPLEMENTATION

A. Implementation Details

We have implemented our protocol in python 3.7.13 on top of the open-source asynchronous DKG codebase of [26]. We use rust libraries for elliptic curve operations and asyncio for concurrency, though our prototype is single-threaded at each party. We implement the random double-sharing protocol we describe in Appendix C.

Our implementation uses the `bls12381` elliptic curve and the implementation from Zcash [34] (with a python wrapper) for primitive elliptic curve operations. Recall that a pairing-friendly curve involves three groups $\mathbb{G}_1, \mathbb{G}_2$, and $\mathbb{G}_T$. We generate the q -SDH parameters in $\mathbb{G}_1$ as it is more efficient than $\mathbb{G}_2$ [37]. The size of each $\mathbb{G}_1$ and $\mathbb{G}_2$ group element after point compression is 48 Bytes and 96 Bytes, respectively.

For equality of discrete logarithm, we use Chaum-Pedersen’s “ Σ ”-protocol. We choose the Chaum-Pedersen “ Σ ” protocol over the pairing-based check, as while benchmarking, we found that the Chaum-Pedersen’s Σ protocol is approximately $2.75\times$ computationally more efficient than the pairing-based check.

Recall from §V that parties need to compute coefficients of polynomials of degree $n - t - 1$, given any arbitrary subsets of $n - t$ points (EVAL messages) on the polynomial. One approach to compute these coefficients is to first compute the evaluations at all points and then apply the inverse NTT transform to get the coefficients. However, this requires each party to perform $O(n^2)$ group multiplications for every polynomial. An asymptotically superior method is the FNT-based NTT implementation due to Soro and Lacan [51]. The latter runs in $O(n \log n)$ time. We implement both approaches, and our microbenchmark illustrates that the quadratic approach performs better for a smaller number of parties.

B. Evaluation Setup

We evaluate our implementation with a varying number of parties: 16, 32, 64, and 128. We evaluate with different values of q : 2^{14} , 2^{16} , and 2^{18} . We run all parties on Amazon Web Services (AWS) *t3x.large* virtual machines with one party per virtual machine. Each virtual machine has 4 vCPUs and 16GB RAM and runs Ubuntu 20.04.

We place parties evenly across eight different AWS regions: Canada, Ireland, N. California, N. Virginia, Oregon, Ohio, Singapore, and Tokyo. We create an overlay network in which all parties are pair-wise connected, i.e., they form a complete graph.

Table IV: Runtime of different phases, for any (n, q) , of our protocol (in % of total runtime). The setup phase represents the combined runtime of both ADKG and the random double-sharing phase.

Protocol Phase	(16, 2^{14})	(16, 2^{18})	(128, 2^{14})	(128, 2^{18})
Setup	4.3	0.4	41.0	5.6
Powers-of-two	1.4	0.1	1.3	0.2
All powers	94.3	99.5	57.7	94.2

With this evaluation setup, we measure the *runtime* and per-party *bandwidth* usage. The runtime is measured from the start of the protocol to the time a party outputs the q -SDH parameters. Per-party bandwidth usage is the amount of data in Bytes sent by a party in the entire protocol.

Baselines. Our baseline is the sequential protocol with a synchronous broadcast channel we describe in §I. In the baseline, we implement the state-of-the-art update verification mechanism from [46], which reduces the verification cost of each update from $2q$ pairings to $2q$ group multiplications and 2 pairings.

The runtime of the baseline includes the time a sequence of n parties takes turns to update existing q -SDH parameters and verify all previous updates. Since the deployed versions skip verifications during the protocol, we implemented a pipelined verification step where every party verifies the update by party i as soon as party i finishes its update. This ensures the runtime of the baseline is the time it takes to perform $3nq$ group multiplications. Here, nq group multiplications for computing the q -SDH parameters and an additional $2nq$ group multiplications for verifying the parameters.

We approximate the bandwidth usage of the baseline as nq group elements. Note that this favors the baseline protocol, as in the synchronous protocol, each party will need to broadcast q group elements. Hence, each party will need to send at least nq group elements in the entire protocol.

Remark. The actual running time or bandwidth usage of the baseline will be higher than what we reported in the figures, as we only measure sub-components for the baseline.

C. Evaluation Results

Our evaluation demonstrates that our protocol scales well with the number of parties n and q and has reasonable runtime and bandwidth usage.

Runtime. We report the median runtime of our protocol and baseline, computed across the parties for a single run of each experiment, in Figure 6. The solid and dashed lines represent the runtime of our protocol and baseline, respectively. We also report the runtime breakdown by the phases of our protocol in Table IV.

Our evaluation results corroborate our analysis in §VI-C. Specifically, for any fixed q , the runtime of our protocol grows logarithmically with the number of parties. Also, for any given n , the runtime grows linearly with q . Note that although the runtime should grow logarithmically for $n = 16$ and $n = 32$, we see a linear growth because, for $n = 16$ and $n = 32$, we implement a protocol with asymptotic runtime of $O(nq)$, but with smaller constants.

Our evaluation illustrates that our protocol is significantly more efficient than the baseline protocol. For example, with




Figure 5: Per party median bandwidth usage (in Megabytes), measured as the amount of data sent by a party in the entire protocol.

64 parties and $q = 2^{16}$, our protocol takes approximately 1037 seconds, whereas the baseline protocol takes 3580 seconds ($3.4\times$ faster). Similarly, with 128 parties and $q = 2^{18}$, our protocol takes approximately 4721 seconds compared to 28883 seconds (not shown in the figure) taken by the baseline protocol, hence $6.1\times$ faster than the baseline protocol. Note that, for the baseline, we only measure the runtime without any networking component. Thus, if we also include the runtime of the synchronous broadcast, the benefits of our protocol will be even more significant.

In Table IV, we break down the runtime of different phases in our protocol for different choices of n and q . In Table V, the setup includes the runtime of both ADKG and random double-sharing protocol. We merge them as our implementation runs these phases such that they share some building blocks. Moreover, both ADKG and the random double-sharing phase use the same invocation of the consensus protocol. Based on Table IV, we conclude the following.

First, the powers-of-two phase takes less than 2% of the runtime in all experiments. Recall from Table III, during the powers-of-two phase, each party performs $O(n \log q)$ group multiplications and sends $O(n \log q)$ group elements. This is significantly smaller than the computation cost and bandwidth usage of setup and all powers phase. However, the powers-of-two phase requires $O(\log q)$ rounds of interaction, which is comparable to the setup and all powers phases. This illustrates that the number of communication rounds is not a bottleneck.

Second, for smaller n , the runtime of the all-powers phase contributes to almost all of the runtime of our protocol, especially for large q . This is as expected because each party needs to perform $O(n^2 \log q)$ and $O(q \log n)$ group multiplications during the setup phase and all powers phase, respectively. For larger q and smaller n , the latter is significantly larger than the former. However, with increasing n , as with the case of $n = 124, q = 2^{14}$, the computation cost of the setup phase also becomes significant. Finally, as expected, for $n = 128$, with larger $q = 2^{18}$, the runtime of the all powers phase again starts to dominate the total runtime.

Bandwidth usage. We report the per-party bandwidth usage, i.e., the amount of data (in Megabytes) each party sends during the entire protocol, in Figure 5. The solid and dashed lines represent the per-party bandwidth usage of our protocol and baseline, respectively. Consistent with the analysis from §VI,




Figure 6: Median runtime (in seconds), i.e., the time between the start of the protocol and the time parties output the q -SDH parameters.

Table V: Bandwidth usage of different phases of our protocol (in % of total bandwidth usages). The setup phase corresponds to the combined bandwidth usages of both ADKG and the random double-sharing phase. The tuple represents (n, q) .

Protocol Phase	$(16, 2^{14})$	$(16, 2^{18})$	$(128, 2^{14})$	$(128, 2^{18})$
Setup	34.4	4.1	95.2	68.8
Powers-of-two	2.2	0.3	0.9	0.7
All powers	63.4	95.6	3.9	30.5

the bandwidth usage in our protocol increases quadratically with the number of parties and linear in the q . We also report the breakdown of bandwidth usage across different phases of our protocol in Table V.

Our evaluation illustrates that parties in our protocol incur significantly less bandwidth usage than the baseline. For example, with 32 parties and $q = 2^{16}$, each party needs to send 13.57 Megabytes (MB) of data in our protocol, compared to 96 MB of data in the baseline protocol ($7\times$ reduction). Similarly, with 128 parties and $q = 2^{18}$, the bandwidth usage in our protocol is 118.17 MB, compared to 1536 MB in the baseline protocol; hence, $13\times$ less bandwidth usage than the baseline. Again, for the baseline, we only measure the bandwidth usage as $\kappa n q$ bytes, where $\kappa = 48$ is the size of each $\mathbb{G}_1$ element, and do not account for any bandwidth usage for the synchronous broadcast protocol. If we also include the bandwidth usage due to the synchronous broadcast channel, the benefits of our protocol will be even bigger.

Similar to the breakdown of runtime across different phases, in Table IV, we illustrate the breakdown of bandwidth usage across different phases of our protocol for different choices of n and q . Again, in Table V, the setup includes the runtime of both ADKG and the random double-sharing protocol. From Table V, we draw a few conclusions.

First, for the same reason as the runtime, the bandwidth usage of the powers-of-two phase takes less than 3% of the total bandwidth usage in all experiments.

Second, unlike runtime, the bandwidth usage in the setup phase is significant compared to the bandwidth usage of the all-powers phase. This is because the ratio between the bandwidth usage of the setup and the all-powers phase is greater, by a factor of $\log n$, than the ratio between the computation cost of the setup and the all-powers phase. Specifically, the former is approximately $O((n^2 \log q)/q)$ while the latter is

$O((n^2 \log q)/q \log n)$. For this same reason, for large $n = 128$ and small $q = 2^{14}$, the bandwidth usage of the setup phase is more than 95% of the total bandwidth usage.

Finally, we conclude that the computation cost is the primary bottleneck of our protocol. With $n = 128$, the setup phase, despite contributing 95.2% and 68.8% of the bandwidth usage for $q = 2^{14}$ and $q = 2^{18}$, respectively, contributes only 41% and 5.5% of the total runtime.

VIII. RELATED WORK

Ben-Sasson et al. [10] proposed the first distributed protocol for securely sampling arbitrary structured public parameters. The protocol of [10] lays the foundation for the round-robin-style protocols for generating secure structured public parameters. Briefly, in [10], parties take turns to update intermediate parameters with local randomness. Bowe et al. [14] adopt the approach of [10] and present a protocol for generating q -SDH parameters. Their protocol, however, relied on a publicly verifiable, unpredictable, and bias-resistant randomness beacon for security. Kohlweiss et al. [41] illustrated that the parameters generated by [14] are secure even without a randomness beacon. Very recently, Nikolaenko et al., [46] designed a protocol to generate q -SDH parameters using Ethereum as the underlying sequential broadcast channel. Their protocol demonstrates how to use a smart contract to eliminate the central coordinator in existing round-robin protocol and achieve censorship resistance.

All these protocols assume a synchronous network, have high communication and computation costs, and require n sequential broadcasts. We refer the reader to §I and Table I for a detailed comparison of these schemes with our construction. Very recently, Cohen et al. [22] presented a generic compiler to reduce the round complexity of these protocols to $O(\sqrt{n})$ sequential broadcasts. However, their construction is very theoretical and has very high constants. They also present a compiler with better constants but $O(\sqrt{n} \log q)$ sequential broadcasts.

Practical deployments. The round-robin style protocols for generating q -SDH parameters have been already deployed in practice [1], [45], [2], [32], [3]. All these deployments implement variants of [10]. However, as deployed, parties skip verifying the intermediate protocol transcript and verify the entire transcript only at the end. Despite these insecurities, they scale very poorly. For example, according to Semaphore [3], to generate q -SDH parameters for $q = 2^{28}$, each party needs to perform a 24-hour long computation. Hence, with n parties, the protocol would run for at least n days.

Comparison with generic MPC. An alternate approach to generate q -SDH parameters in asynchronous networks is to use generic MPC. However, this has many disadvantages.

Let $\mathcal{C}$ be the circuit that outputs the q -SDH parameters, then $\mathcal{C}$ will consist of $O(q)$ multiplication gates. For large q , evaluating $O(q)$ multiplication gates in an asynchronous MPC can be prohibitively expensive. For example, asynchronous MPC protocols that rely on pre-processing either require threshold additive homomorphic encryptions [35], [36], [20] or can tolerate only $n/4$ malicious parties [21]. We want to emphasize that while generating q -SDH parameters, we

also have to include the cost of the pre-processing phase in the cost of the overall protocol. Alternatively, protocols that do not rely on a pre-processing step require running $O(n)$ asynchronous complete secret sharing protocol for every multiplication gate [9], [4].

Another issue of using generic MPC is due to the difference between the scalar and based field of elliptic curve groups. Typically, MPC protocols are defined over a single finite field, whereas q -SDH parameter generation involves working with both the scalar field $\mathbb{F}$, from where τ is sampled, and the base field, which is used to define the elliptic curve group elements. Since the scalar $\mathbb{F}$ is different from the base field; the MPC protocol needs to support operations across two distinct fields, which can be prohibitively expensive [25].

Other related works. The multiplication in the exponent approach we adopt bears similarity with [50], [47]. However, there are crucial differences. Both [50] and [47] consider security-with-abort, albeit differently. For example, [47] does not verify shares from parties while computing the KZG evaluation proof and uses properties of the KZG polynomial commitment to check the correctness of the final proof. Similarly, [50] validates the reconstructed values using MACs and outputs a default value upon unsuccessful verification. Moreover, [50] assumes synchrony as it uses a broadcast channel to verify the MACs. Contrary to both, our approach ensures guaranteed output delivery in asynchrony with $1/3$ failures. Also, we use publicly available threshold public keys and bilinear pairings to verify the shares of each party efficiently.

IX. DISCUSSION AND CONCLUSION

In this paper, we presented our protocol, a distributed protocol to generate secure parameters for q -Strong Diffie-Hellman, also known as the q -SDH, interactability problem. In an asynchronous network of n parties, our protocol tolerates up to one-third of malicious parties. Our protocol is also very efficient. For any given q , the highest degree of the q -SDH parameters, each party incurs a communication cost of $O(q + n^2 \log q)$, which is optimal whenever $q \geq n^2 \log q$. Moreover, our protocol only requires $O(\log q + \log n)$ rounds of interaction. Furthermore, in the entire protocol, each party performs only $O(q \log n)$ group multiplications and $O(n \log q)$ bilinear pairings. We have implemented all parts of our protocol and evaluated it using up to 128 geographically distributed AWS instances. Our evaluation results corroborate the practicality of our approach and its significant improvement over the state-of-the-art protocol.

Extending our protocol with the round-robin approach. As we mention in §I, the existing round-robin synchronous protocol has several nice properties that our protocol does not have. To reiterate, in the round-robin protocol, the set of participants need not be known in advance, and parties can join on the fly. Additionally, the round-robin protocol has a you-only-speak-once flavor, as each party uses its secret only once. Some applications of powers-of-tau might seek to preserve these properties. Although we do not know how to achieve these properties in asynchrony, applications can potentially achieve these properties in practice by feeding the output of our protocol into a synchronous round-robin protocol. Intuitively, this provides the best of both world guarantees. We leave the rigorous security analysis of this approach as future work.

Open problems. In addition to improving the performance metrics, there are many other interesting future research directions to consider. One research direction is to design a protocol that achieves similar performance results with weaker assumptions, such as without bilinear pairing or random oracle. Another research direction is to generate the parameters for the multi-variate generalization of the q -SDH parameters [48]. Another interesting research direction is to design a network agnostic protocol for generating q -SDH parameters, i.e., a protocol that can gracefully tolerate $1/3$ failures in asynchrony and up to $1/2$ failures in synchrony. Very recently, Bacho et al. [7] presented such a protocol for distributed key generation.

ACKNOWLEDGMENTS

The authors would like to thank Amit Agarwal and Atsuki Momose for many helpful discussions related to the paper. This work is funded in part by a Chainlink Labs Ph.D. fellowship and the National Science Foundation award #2240976.

REFERENCES

- [1] “Plumo ceremony,” <https://celo.org/plumo>, 2017.
- [2] “Universal crs setup,” <https://docs.zksync.io/userdocs/security/#universal-crs-setup>, 2020.
- [3] “Perpetual powers of tau,” <https://github.com/weijiekoh/perpetualpowersoftau>, 2021.
- [4] I. Abraham, G. Asharov, and A. Yanai, “Efficient perfectly secure computation with optimal resilience,” in *Theory of Cryptography Conference*. Springer, 2021, pp. 66–96.
- [5] I. Abraham, P. Jovanovic, M. Maller, S. Meiklejohn, and G. Stern, “Bingo: Adaptively secure packed asynchronous verifiable secret sharing and asynchronous distributed key generation,” in *Annual International Cryptology Conference*. Springer, 2023.
- [6] N. Alhaddad, M. Varia, and H. Zhang, “High-threshold avss with optimal communication complexity,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2021, pp. 479–498.
- [7] R. Bacho, D. Collins, C.-D. Liu-Zhang, and J. Loss, “Network-agnostic security comes for free in dkg and mpc,” 2023.
- [8] D. Beaver, “Efficient multiparty protocols using circuit randomization,” in *Annual International Cryptology Conference*. Springer, 1991, pp. 420–432.
- [9] M. Ben-Or, S. Goldwasser, and A. Wigderson, “Completeness theorems for non-cryptographic fault-tolerant distributed computation,” in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, ser. STOC ’88, New York, NY, USA, 1988, p. 1–10.
- [10] E. Ben-Sasson, A. Chiesa, M. Green, E. Tromer, and M. Virza, “Secure sampling of public parameters for succinct zero knowledge proofs,” in *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 287–304.
- [11] A. Bhat, N. Shrestha, Z. Luo, A. Kate, and K. Nayak, “Randpiper—reconfiguration-friendly random beacons with quadratic communication,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 3502–3524.
- [12] G. R. Blakley, “Safeguarding cryptographic keys,” in *1979 International Workshop on Managing Requirements Knowledge (MARK)*. IEEE, 1979, pp. 313–318.
- [13] D. Boneh and X. Boyen, “Short signatures without random oracles and the sdh assumption in bilinear groups,” *Journal of cryptology*, vol. 21, no. 2, pp. 149–177, 2008.
- [14] S. Bowe, A. Gabizon, and I. Miers, “Scalable multi-party computation for zk-snark parameters in the random beacon model,” *Cryptology ePrint Archive*, 2017.
- [15] G. Bracha, “Asynchronous byzantine agreement protocols,” *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.
- [16] C. Cachin, K. Kursawe, F. Petzold, and V. Shoup, “Secure and efficient asynchronous broadcast protocols,” in *Annual International Cryptology Conference*. Springer, 2001, pp. 524–541.
- [17] T. D. Chandra and S. Toueg, “Unreliable failure detectors for reliable distributed systems,” *Journal of the ACM (JACM)*, vol. 43, no. 2, pp. 225–267, 1996.
- [18] D. Chaum and T. P. Pedersen, “Wallet databases with observers,” in *Annual International Cryptology Conference*. Springer, 1992, pp. 89–105.
- [19] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. Ward, “Marlin: preprocessing zk-snarks with universal and updatable srs,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2020, pp. 738–768.
- [20] A. Choudhury and A. Patra, “Optimally resilient asynchronous mpc with linear communication complexity,” in *Proceedings of the 2015 International Conference on Distributed Computing and Networking*, 2015, pp. 1–10.
- [21] —, “An efficient framework for unconditionally secure multiparty computation,” *IEEE Transactions on Information Theory*, vol. 63, no. 1, pp. 428–468, 2016.
- [22] R. Cohen, J. Doerner, Y. Kondi, and A. Shelat, “Guaranteed output in $O(\sqrt{n})$ rounds for round-robin sampling protocols,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2022, pp. 241–271.
- [23] I. Damgård, Y. Ishai, M. Krøigaard, J. B. Nielsen, and A. Smith, “Scalable multiparty computation with nearly optimal work and resilience,” in *Annual International Cryptology Conference*. Springer, 2008, pp. 241–261.
- [24] I. Damgård and J. B. Nielsen, “Scalable and unconditionally secure multiparty computation,” in *Annual International Cryptology Conference*. Springer, 2007, pp. 572–590.
- [25] I. Damgård and R. Thorbek, “Efficient conversion of secret-shared values between different fields,” *Cryptology ePrint Archive*, 2008.
- [26] S. Das, Z. Xiang, L. Kokoris-Kogias, and L. Ren, “Practical asynchronous high-threshold distributed key generation and distributed polynomial sampling,” in *To appear at the 32st USENIX Security Symposium (USENIX Security 23)*, 2023.
- [27] S. Das, Z. Xiang, and L. Ren, “Asynchronous data dissemination and its applications,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.
- [28] S. Das, T. Yurek, Z. Xiang, A. Miller, L. Kokoris-Kogias, and L. Ren, “Practical asynchronous distributed key generation,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022.
- [29] C. Dwork, N. Lynch, and L. Stockmeyer, “Consensus in the presence of partial synchrony,” *Journal of the ACM (JACM)*, vol. 35, no. 2, pp. 288–323, 1988.
- [30] Ethereum, “Powers of tau specification,” <https://github.com/ethereum/kzg-ceremony-specs>, 2022.
- [31] A. Fiat and A. Shamir, “How to prove yourself: Practical solutions to identification and signature problems,” in *Conference on the theory and application of cryptographic techniques*. Springer, 1986, pp. 186–194.
- [32] A. Gabizon, “Perpetual powers of tau (for bls381),” <https://github.com/arielgabizon/perpetualpowersoftau>, 2020.
- [33] A. Gabizon, Z. J. Williamson, and O. Ciobotaru, “Plonk: Permutations over lagrange-bases for ocumenical noninteractive arguments of knowledge,” *Cryptology ePrint Archive*, 2019.
- [34] J. Grigg and S. Bowe, “zkcrypto/pairing,” <https://github.com/zkcrypto/pairing>.
- [35] M. Hirt, J. B. Nielsen, and B. Przydatek, “Cryptographic asynchronous multi-party computation with optimal resilience,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2005, pp. 322–340.
- [36] —, “Asynchronous multi-party computation with quadratic communication,” in *International Colloquium on Automata, Languages, and Programming*. Springer, 2008, pp. 473–485.
- [37] Y. E. Housni, “Benchmarking pairing-friendly elliptic curves libraries,” <https://hackmd.io/@gnark/eccbench>, 2020.
- [38] I. Karantaidou and F. Baldimtsi, “Efficient constructions of pairing

based accumulators,” in *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. IEEE, 2021, pp. 1–16.

- [39] A. Kate, G. M. Zaverucha, and I. Goldberg, “Constant-size commitments to polynomials and their applications,” in *International conference on the theory and application of cryptology and information security*. Springer, 2010, pp. 177–194.
- [40] A. Kiayias, O. Oksuz, and Q. Tang, “Distributed parameter generation for bilinear diffie hellman exponentiation and applications,” in *Information Security: 18th International Conference, ISC 2015, Trondheim, Norway, September 9–11, 2015, Proceedings 18*. Springer, 2015, pp. 548–567.
- [41] M. Kohlweiss, M. Maller, J. Siim, and M. Volkhov, “Snarky ceremonies,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2021, pp. 98–127.
- [42] E. Kokoris Kogias, D. Malkhi, and A. Spiegelman, “Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 1751–1767.
- [43] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” in *Concurrency: the works of leslie lamport*, 2019, pp. 203–226.
- [44] M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn, “Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 2111–2128.
- [45] A. Miller and S. Bowe, “Announcing the world’s largest multi-party computation ceremony,” <https://zfn.org/announcing-the-worlds-largest-multi-party-computation-ceremony/>, 2017.
- [46] V. Nikolaenko, S. Ragsdale, J. Bonneau, and D. Boneh, “Powers-of-tau to the people: Decentralizing setup ceremonies,” *Cryptology ePrint Archive*, 2022.
- [47] A. Ozdemir and D. Boneh, “Experimenting with collaborative {zk-SNARKs}:[Zero-Knowledge] proofs for distributed secrets,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 4291–4308.
- [48] C. Papamanthou, E. Shi, and R. Tamassia, “Signatures of correct computation,” in *Theory of Cryptography Conference*. Springer, 2013, pp. 222–242.
- [49] A. Shamir, “How to share a secret,” *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [50] N. P. Smart and Y. Talibi Alaoui, “Distributing any elliptic curve based protocol,” in *Cryptography and Coding: 17th IMA International Conference, IMACC 2019, Oxford, UK, December 16–18, 2019, Proceedings*. Springer, 2019, pp. 342–366.
- [51] A. Soro and J. Lacan, “Fnt-based reed-solomon erasure codes,” in *2010 7th IEEE Consumer Communications and Networking Conference*. IEEE, 2010, pp. 1–5.
- [52] S. Srinivasan, A. Chepurnoy, C. Papamanthou, A. Tomescu, and Y. Zhang, “Hyperproofs: Aggregating and maintaining proofs in vector commitments,” in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, 2022, pp. 3001–3018.
- [53] T. Xie, J. Zhang, Y. Zhang, C. Papamanthou, and D. Song, “Libra: Succinct zero-knowledge proofs with optimal prover computation,” in *Annual International Cryptology Conference*. Springer, 2019, pp. 733–764.

APPENDIX A

ANALYSIS OF THE SQUARING PROTOCOL

In this subsection, we will assume that parties start with correct double sharing of uniform random value z .

Correctness. The verification of SQ messages ensures that honest parties only accept valid SQ messages. This implies that honest parties only interpolate correct shares and hence will correctly output $a^2 + z$. Finally, $a^2 + z - [z]_i$ is a valid

Simulator $\mathcal{S}_{\text{sq}}$

Inputs. Set of malicious parties $\mathcal{C}$, secret shares $[a]_i, [a^2]_i$ for all $i \in \mathcal{C}$, set of threshold public keys $[a]_g$ and $[a^2]_g$.

INPUT GENERATION PHASE:

- 1) Sample a random polynomial $r(\cdot)$ of degree $2t$.
- 2) Compute $[z]_i^{2t} = r(i) - [a]_i[a]_i, [z]_i = r(0) - [a^2]_i$ for each $i \in \mathcal{C}$. Also, compute $[z]_j^{2t} g = (r(j) - [a]_j[a]_j)g$ and $[z]_j g = (r(0) - [a^2]_j)g$ for each $j \in [0, n]$, where $[a]_0 = a$.

SIMULATION PHASE:

- 1) Run $\mathcal{S}_{\text{Dou}}$ on input $[z]_i, [z]_i^{2t}$ for each $i \in [t]$, and $[z]_g, [z]_g^{2t}$.
- 2) For each emulated honest party j , compute $[a]_j[a]_j g$ using equation (17). Let $\pi_j = \mathcal{S}_{\text{dlog}}(g, [a]_j g, [a]_j g, [a]_j[a]_j g)$ be the simulated proof of equality of discrete logarithm.
- 3) On behalf of each emulated honest party j , send $\langle \text{SQ}, [a]_j[a]_j g, r(j), \pi_j \rangle$ to every party.

Figure 7: Simulator for the protocol Π_{sq} for functionality $\mathcal{F}_{\text{sq}}$.

secret share of a^2 i.e., $[a^2]$ since

$$\sum_j \lambda_j (a^2 + z - [z]_j) = a^2 + z - \sum_j \lambda_j [z]_j = a^2 \quad (16)$$

Security. We will prove the security of Π_{sq} by showing its *simulatability*. Specifically, we will illustrate that, for any static PPT adversary $\mathcal{A}$, that corrupts up to t parties and additionally observes $[a]_g$ and $[a^2]_g$, there exists a simulator $\mathcal{S}_{\text{sq}}$, that takes as input only the adversarial shares and the publicly available $[a]_g$ and $[a^2]_g$, and simulates a view that is indistinguishable from $\mathcal{A}$ ’s view in the real execution of the protocol.

In our proof, we assume the existence of a simulator $\mathcal{S}_{\text{Dou}}$ for the protocol Π_{Dou} , that securely realizes the ideal functionality $\mathcal{F}_{\text{Dou}}$ shown in figure 8. $\mathcal{S}_{\text{Dou}}$ takes as inputs $[z]_i, [z]_i^{2t}$ for each corrupt party i , along with public values $zg, [z]_g$, and $[z]_g^{2t}$. $\mathcal{S}_{\text{Dou}}$ then simulates one invocation of Π_{Dou} such that at the end of the protocol parties output double shares of z where the shares of the adversarial parties matches the input to $\mathcal{S}_{\text{Dou}}$. The double sharing protocol of [26] does not immediately admits such a simulator. However, as we illustrate in Appendix C their protocol can be modified with minor overhead to admit such a simulator. We summarize our simulator $\mathcal{S}_{\text{sq}}$ in Figure 7 and describe it next.

Without loss of generality, we assume that $\mathcal{A}$ corrupts the first t parties. Thus, $\mathcal{S}_{\text{sq}}$ will receive the shares $[a]_i$ and $[a^2]_i$ for each $i \in [t]$, along with public values $[a]_g$ and $[a^2]_g$. With these inputs, $\mathcal{S}_{\text{sq}}$ emulates the remaining $n - t$ honest parties and generates the protocol transcript as follows.

$\mathcal{S}_{\text{sq}}$ first computes $[a]_j[a]_j g$ for every $j \in [n]$ using its knowledge of $[a]_i$ for each $i \in [t]$, $[a]_g$, and the following identity

$$[a]_j[a]_j = \left(\sum_{k=0}^t \mathcal{L}_k(j) [a]_k \right) \left(\sum_{k=0}^t \mathcal{L}_k(j) [a]_k \right) \quad (17)$$

$\mathcal{S}_{\text{sq}}$ then samples a random polynomial $r(\cdot)$ of degree $2t$

and computes:

$$\begin{aligned}\llbracket z \rrbracket_j^{2t} &= (r(j) - \llbracket a \rrbracket_j \llbracket a \rrbracket_j) \mathbf{g} \quad \forall j \in [0, n]; \\ \llbracket z \rrbracket_i^{2t} &= r(i) - \llbracket a \rrbracket_i \llbracket a \rrbracket_i \quad \forall i \in [t] \\ \llbracket z \rrbracket_i &= r(0) - \llbracket a^2 \rrbracket_i \quad \forall i \in [t]\end{aligned}$$

$\mathcal{S}_{\text{Sq}}$ then runs $\mathcal{S}_{\text{Dou}}$ using the values computed above. Also, for each emulated honest party j , $\mathcal{S}_{\text{Sq}}$ uses the NIZK simulator $\mathcal{S}_{\text{dleq}}$ of the dleq protocol to compute the simulated proof $\pi_j = \mathcal{S}_{\text{dleq}}(\mathbf{g}, \llbracket a \rrbracket_j \mathbf{g}, \llbracket a \rrbracket_j \mathbf{g}, \llbracket a \rrbracket_j \llbracket a \rrbracket_j \mathbf{g})$. Finally, on behalf of each emulated honest party j , $\mathcal{S}_{\text{Sq}}$ multicasts the message $\langle \text{SQ}, \llbracket a \rrbracket_j \llbracket a \rrbracket_j \mathbf{g}, r(j), \pi_j \rangle$ to every party.

We now prove that the view of $\mathcal{A}$ in the simulated protocol is identical to $\mathcal{A}$'s view in the real execution of the protocol. In particular, we define a series of hybrids, with Hybrid 0 identical to the real protocol and Hybrid 3 identical to the simulated protocol, to argue that no PPT distinguisher can distinguish between the real and ideal world.

Hybrid 0. This corresponds to the real-world execution.

Hybrid 1. In this hybrid, we simulate the NIZK proofs of equality of discrete logarithms for each statement of the form $\{\mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \mathbf{g}, \llbracket a \rrbracket_i \llbracket a \rrbracket_i \mathbf{g}\}$. Since Chaum-Pedersen Σ protocol for proving equality of discrete logarithm is perfect zero-knowledge Hybrid 1 is identically distributed as Hybrid 0.

Hybrid 2. Swap out the real execution of the Π_{Dou} with its simulator $\mathcal{S}_{\text{Dou}}$ whose inputs are adversarial shares of a double sharing of a random value z and the corresponding commitments $\llbracket z \rrbracket \mathbf{g}$ and $\llbracket z \rrbracket^{2t} \mathbf{g}$. The perfect simulatability of the $\mathcal{S}_{\text{Dou}}$ guarantees that Hybrid 2 is identically distributed as Hybrid 1.

Hybrid 3. Same as Hybrid 2 except the following changes. Sample a uniform random polynomial $r(\cdot)$ of degree $2t$ as per step 1 of INPUT GENERATION PHASE of Figure 7. Compute new input for $\mathcal{S}_{\text{Dou}}$ as step 2 of INPUT GENERATION PHASE of Figure 7. Run $\mathcal{S}_{\text{Dou}}$ on input $\llbracket z \rrbracket_i$ and $\llbracket z \rrbracket_i^{2t}$ for each malicious party i , and remaining inputs $\llbracket z \rrbracket \mathbf{g}$, and $\llbracket z \rrbracket^{2t} \mathbf{g}$. Also, for each honest party j , send $r(j)$ as a part of SQ message.

In Hybrid 3, since $r(\cdot)$ is a random polynomial of degree $2t$, the probability of $r(\cdot)$ taking a certain value is $1/|\mathbb{F}|^{2t+1}$, as there are $|\mathbb{F}|^{2t+1}$ such polynomials in total. Furthermore, in Hybrid 3, by construction, $r(\cdot)$ is consistent with the output of $\mathcal{S}_{\text{Dou}}$. In Hybrid 2, each honest party j reveals $\llbracket a \rrbracket_j \llbracket a \rrbracket_j + \llbracket z \rrbracket^{2t}$. Note that $\llbracket a \rrbracket_j$ is fixed for all $j \in [n]$, hence, the probability that $\llbracket a \rrbracket_j \llbracket a \rrbracket_j + \llbracket z \rrbracket^{2t} = r(j)$ is same as the probability of $\llbracket z \rrbracket^{2t} = r(j) - \llbracket a \rrbracket_j \llbracket a \rrbracket_j$. Since, $\llbracket z \rrbracket^{2t}$ in Hybrid 2 is a uniform random polynomial of degree $2t$, $\Pr[\llbracket z \rrbracket^{2t} = r(j) - \llbracket a \rrbracket_j \llbracket a \rrbracket_j] = 1/|\mathbb{F}|^{2t+1}$. This implies that Hybrid 3 is identically distributed as Hybrid 2.

APPENDIX B

ZERO KNOWLEDGE PROOF OF EQUALITY OF DISCRETE LOGARITHM

Given a group $\mathbb{G}$ with scalar field $\mathbb{F}$ of prime order p , two uniformly random generators $\mathbf{g}, \mathbf{h} \in \mathbb{G}$ and a tuple $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b}) \in \mathbb{G}^4$, a prover $\mathcal{P}$ wants to prove to a probabilistic polynomial time verifier $\mathcal{V}$, in zero-knowledge, the knowledge of a witness $\alpha \in \mathbb{F}$ such that $\mathbf{a} = \alpha \mathbf{g}$ and $\mathbf{b} = \alpha \mathbf{h}$.

Functionality $\mathcal{F}_{\text{Dou}}$

- Let $\mathbb{G}$ be an elliptic curve group with scalar field $\mathbb{F}$ and let $\mathbf{g}$ be a uniformly random generator of $\mathbb{G}$.
- Wait for $\mathcal{C}$, the set of adversarial parties and the (start, q) message from $\mathcal{A}$.
- Wait for (init, q) from all honest parties.
- Sample a uniformly random element $z \in \mathbb{F}$. Generate $(n, t+1)$ and $(n, 2t+1)$ shares of z , denoted with $\llbracket z \rrbracket$ and $\llbracket z \rrbracket^{2t}$, respectively.
- Compute $\llbracket z \rrbracket \mathbf{g}$ and $\llbracket z \rrbracket^{2t} \mathbf{g}$ and send the tuple $(\mathbf{g}, \llbracket z \rrbracket_i, \llbracket z \rrbracket_i^{2t}, \llbracket z \rrbracket \mathbf{g}, \llbracket z \rrbracket^{2t} \mathbf{g})$ to party i .

Figure 8: The functionality for random double sharing

Protocol for equality of discrete logarithm. We use the Chaum-Pedersen Σ -protocol [18], that assumes the hardness of the Discrete Logarithm in $\mathbb{G}$ and proceeds as follows.

- 1) $\mathcal{P}$ samples a random element $\beta \leftarrow \mathbb{F}$ and sends $(\mathbf{a}_1, \mathbf{a}_2)$ to $\mathcal{V}$ where $\mathbf{a}_1 = \beta \mathbf{g}$ and $\mathbf{a}_2 = \beta \mathbf{g}$.
- 2) $\mathcal{V}$ sends a challenge $e \leftarrow \mathbb{F}$.
- 3) $\mathcal{P}$ sends a response $z = \beta - \alpha e$ to $\mathcal{V}$.
- 4) $\mathcal{V}$ checks whether $\mathbf{a}_1 = z \mathbf{g} + e \mathbf{a}$ and $\mathbf{a}_2 = z \mathbf{h} + e \mathbf{b}$ and accepts if and only if both the equality holds.

The Chaum-Pedersen Σ -protocol guarantees completeness, knowledge soundness, and zero-knowledge. The knowledge soundness implies that if $\mathcal{P}$ convinces the $\mathcal{V}$ with non-negligible probability, there exists an efficient (polynomial time) extractor that can extract α from $\mathcal{P}$ non-negligible probability.

This protocol can be made non-interactive in the Random Oracle model using the Fiat-Shamir heuristic [31]. For any given tuple $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$ where $\mathbf{a} = \alpha \mathbf{g}$ and $\mathbf{b} = \alpha \mathbf{h}$, $\text{dleq.Prove}(\alpha, \mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$ generates the non-interactive zero proof π . The proof π is $O(\kappa)$ bits long. Given a proof π and $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$, $\text{dleq.Verify}(\pi, \mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b})$ verifies the proof.

Simulating a proof without the secret. We will use programmability of random oracle to generate an convincing NIZK proof without having access to the corresponding secret. Given the tuple $(\mathbf{g}, \mathbf{a}, \mathbf{h}, \mathbf{b}) \in \mathbb{G}^4$ where $\mathbf{a} = \alpha \mathbf{g}$ and $\mathbf{b} = \alpha \mathbf{h}$, the simulator works as follows.

- 1) Sample uniformly random $z, e \in \mathbb{F}$.
- 2) Compute $\mathbf{a}_1 = z \mathbf{g} + e \mathbf{a}$ and $\mathbf{a}_2 = z \mathbf{h} + e \mathbf{b}$.
- 3) Program the random oracle such that $\text{RO}(\mathbf{a}_1, \mathbf{a}_2) := e$.
- 4) Output $\pi = (\mathbf{a}_1, \mathbf{a}_2, z)$

where $\text{RO}(\cdot)$ denotes query to the random oracle.

Note that the distribution of the simulated proof is identical to the distribution of the proof generated by an honest prover.

APPENDIX C

ASYNCHRONOUS RANDOM DOUBLE SHARING

Our protocol uses $\log q$ double sharing of random values as per the ideal functionality $\mathcal{F}_{\text{Dou}}$ in Figure 8. We make minor modifications to the double sharing protocol of [26] to admit a simulation based security proof. We want to note that the double sharing protocol of [26] is secure in the context they

Simulator $\mathcal{S}_{\text{Dou}}$

Inputs. Set of adversarial parties $\mathcal{C}$, $\llbracket z \rrbracket_i, \llbracket z \rrbracket_i^{2t}$ for all $i \in \mathcal{C}$, and threshold public keys $\llbracket z \rrbracket_{\mathbf{g}}, \llbracket z \rrbracket_{\mathbf{g}}^{2t}$.

- 1) Sample a random $\alpha \in \mathbb{F}$ and let $\mathbf{h} = \alpha \mathbf{g}$.
- 2) Let $\mathcal{H} = [n] \setminus \mathcal{C}$ be the set of emulated honest parties. For each $j \in \mathcal{H}$, run the Sharing and Agreement phase as per the protocol specification.
- 3) Sample two random polynomials $\hat{z}(\cdot)$ and $z(\cdot)$ of degree t and $2t$, respectively, such that:
 - $\hat{z}(0)\mathbf{g} = z(0)\mathbf{g} = z\mathbf{g}$
 - $\hat{z}(i) = \llbracket z \rrbracket_i$ and $z(i) = \llbracket z \rrbracket_i^{2t}$, for each $i \in \mathcal{C}$,
 - $\llbracket z \rrbracket_j \mathbf{g} = \hat{z}(j)\mathbf{g}$ and $\llbracket z \rrbracket_j^{2t} \mathbf{g} = z(j)\mathbf{g}$
- 4) Use α and the extractability of the Sharing phase to emulate a Randomness extraction phase such that each malicious party $i \in \mathcal{C}$ outputs $\llbracket z \rrbracket_i$ and $\llbracket z \rrbracket_i^{2t}$ as its random double share.
- 5) Use $\mathcal{S}_{\text{dleg}}$ to compute the NIZK proofs required during the key derivation phase.

Figure 9: Simulator for the protocol for random double sharing.

consider in their paper. Let Π_{Dou} be our protocol. Before we describe our modifications, we provide a brief overview of the double sharing protocol of [26].

Double sharing protocol of [26]. The main observation in [26] is that double sharing of a random element z is equivalent to sampling two random polynomials of degree t and $2t$, respectively, that have the same constant term z . Let $\hat{z}(\cdot)$ and $z(\cdot)$ be the polynomials defined as:

$$\begin{aligned}\hat{z}(x) &= z + \hat{z}_1 x + \hat{z}_2 x^2 + \dots + \hat{z}_t x^t \\ z(x) &= z + z_1 x + z_2 x^2 + \dots + z_{2t} x^{2t}\end{aligned}$$

where all $z, \hat{z}_j, z_k$ are uniformly random element in $\mathbb{F}$. Each party i then receives $z(i)$ and $\hat{z}(i)$.

Difficulty in proving simulation security of [26]. The protocol of [26] samples the polynomials $\hat{z}(\cdot)$ and $z(\cdot)$ using different approaches. Their protocol has four phases: Sharing, Agreement and Randomness Extraction, and Key Derivation. They sample the polynomial $z(\cdot)$ in a way such that $\mathcal{A}$ learns no information about $z(i)$ for adversarial parties i until the Agreement phase finishes at least one honest party. However, for polynomial $\hat{z}(\cdot)$, $\mathcal{A}$ learns $\hat{z}(i)$ during the sharing phase.

Our approach. At a high-level, we modify the protocol such that both $\hat{z}(\cdot)$ and $z(\cdot)$ are sampled such that $\mathcal{A}$ learns no information its shares on both $\hat{z}(\cdot)$ and $z(\cdot)$ before the agreement phase terminates. This enables the simulator $\mathcal{S}_{\text{Dou}}$ to match the shares of the adversarial parties with the shares received from the $\mathcal{F}_{\text{Dou}}$. Specifically, we make the following changes:

- 1) During Sharing phase, each party secret shares three random values instead of two as in [26].
- 2) During the Randomness Extraction phase, in addition to $\llbracket z \rrbracket$ and $\llbracket z_j \rrbracket$ for $j \in [2t]$, parties also generate $\llbracket \hat{z}_k \rrbracket$ for each $k \in [t]$ using the random extractor. Parties assist each party i to additionally compute $\hat{z}(i)$.
- 3) During the Key Derivation phase, parties additionally output $\hat{z}(i)\mathbf{g}$ for all $i \in [n]$.

Analysis. We describe the simulator $\mathcal{S}_{\text{Dou}}$ in Figure 9, that on inputs adversarial double shares of a random value from $\mathcal{F}_{\text{Dou}}$ simulates the protocol Π_{Dou} .

We now prove that the simulated transcript is identically distributed to the real execution transcript. We will prove this by defining a series of hybrids, with hybrid 0 identical to the real protocol and hybrid 4 identical to the simulated protocol.

Hybrid 0. This corresponds to the real-world execution.

Hybrid 1. Same as Hybrid 0 except that the common random string element h is sampled as g^α for a known uniform random $\alpha \in \mathbb{F}$. Hybrid 1 is indistinguishable from Hybrid 0 as the distribution of h is identical in both hybrids.

Hybrid 2. Same as Hybrid 1, except we simulate the NIZK proofs of equality of discrete logarithms used during the key-derivation phase. Since Chaum-Pedersen Σ -protocol for proving the equality of discrete logarithm is perfect zero-knowledge Hybrid 2 is identically distributed as Hybrid 1.

Hybrid 3. Change the randomness extraction messages of all honest parties based on the output of $\mathcal{F}_{\text{Dou}}$ as per step 4 of Figure 9. Also, update the randomness of Pedersen's commitment of commitments of honest parties using knowledge of trapdoor α and knowledge of shares of malicious nodes computed during the randomness extraction phase.

Hybrid 3 is identically distributed as Hybrid 2 due to the perfect hiding of the Pedersen commitment scheme, the perfect secrecy of Shamir secret sharing and the fact that the output of the randomness extractor are uniformly random. The perfect hiding property of the Pedersen commitment scheme reveals no information about the underlying message. The security of the $(n, t + 1)$ Shamir secret sharing scheme ensures that less than or equal to t shares reveal no information about the remaining shares. Thus, when we change the output of the randomness extractor with uniformly random shares received from $\mathcal{F}_{\text{Dou}}$ while ensuring consistency of the randomness extraction messages, Hybrid 3 maintains the same distributed as Hybrid 2.

APPENDIX D

ASYNCHRONOUS DISTRIBUTED KEY GENERATION

Functionality $\mathcal{F}_{\text{ADKG}}$

- Let $n \geq 3t + 1$ be the total number of parties. Let $\mathbb{G}$ be a elliptic curve group with a random generator $\mathbf{g}$.
- Wait for $(d, \mathcal{C})$, the degree of the secret sharing and the set of adversarial parties and the (start, q) message from $\mathcal{A}$. Check that $d > |\mathcal{C}|$ and $|\mathcal{C}| \leq t$.
- Wait for (init, q) from all honest parties.
- Sample a uniformly random element $z \in \mathbb{F}$. Generate $(n, d + 1)$ shares of z , denoted with $\llbracket z \rrbracket$.
- Compute $\llbracket z \rrbracket \mathbf{g}$ and send $(\mathbf{g}, z\mathbf{g}, \llbracket z \rrbracket_i, \llbracket z \rrbracket_i \mathbf{g})$ to party i .

Figure 10: Asynchronous Distributed Key Generation functionality