

# EthDKG: Distributed Key Generation with Ethereum Smart Contracts

Philipp Schindler, Aljosha Judmayer, Nicholas Stifter, and Edgar Weippl  
SBA Research, University of Vienna, TU Wien

**Abstract**—Distributed key generation (DKG) is a fundamental building block for a variety of cryptographic schemes and protocols, such as threshold cryptography, multi-party coin tossing schemes, public randomness beacons and consensus protocols. More recently, the surge in interest for blockchain technologies, and in particular the quest for developing scalable protocol designs, has renewed and strengthened the need for efficient and practical DKG schemes. Surprisingly, given the broad range of applications and available body of research, fully functional and readily available DKG protocol implementations still remain limited. This paper hereby aims to close this gap by tailoring Gennaro et al.’s [1] well known protocol design towards being efficiently implementable within public cryptocurrency ecosystems such as Ethereum. Our theoretical improvements are supported by an open source, fully functional, well documented DKG implementation<sup>1</sup> that can employ any Ethereum Virtual Machine (EVM) compatible smart contract platform as a communication layer. We evaluate the efficiency of our protocol and demonstrate its practicability through the deployment and successful execution of our DKG contract in the Ethereum Ropsten testnet. Given the current Ethereum block gas limit, all steps required for the key generation process, even in demanding scenarios tested with up to 256 nodes, can be verified at the smart contract level.

## I. INTRODUCTION

Distributed key generation (DKG) protocols serve as a key building block for threshold cryptography. The goal of a DKG scheme is to agree on a common secret/public key pair such that the secret key is shared among a set of  $n$  participants. Only a subset of  $t + 1 \leq n$  parties can use or reveal the generated secret key, while  $t$  collaborating parties cannot learn any information about it. In this regard DKG is related to secret sharing protocols, as first introduced by Shamir [2] and Blakley [3]. However, in contrast to secret sharing, DKG protocols do not rely on a (trusted) dealer which generates, knows and distributes the secret key, and hence avoid this single point of failure. Instead, the key pair is generated using a multi-party computation in a way that no single party learns the secret that is being shared.

Distributed key generation has been studied and discussed for over two decades [4], [5], [1], [6], [7], [8], [9]. However, the extensive body of literature is currently not matched by a single clear, succinct, and practical protocol design template that reflects the state of the art and leverages on recent technical developments such as distributed ledgers. Moreover, real-world open source implementations of DKG protocols are still rare, and often not well documented.

<sup>1</sup>The source code, documentation, and logs of a successful execution in the Ropsten testnet are publicly available at <https://github.com/PhilippSchindler/EthDKG/>.

We aim to close this gap by providing and evaluating a lightweight, scalable, and well-documented protocol design and open source implementation of a DKG protocol. Our design is based on the Joint-Feldman DKG protocol [1] and incorporates the enhancements proposed by Neji et al. [9] to address biasing attacks [1], without requiring two distinct secret sharing rounds. Additionally, we describe and implement a new mechanism that handles disputes during the protocol execution more efficiently. The resulting protocol design is described in its generality for any discrete logarithm based cryptosystem, and we demonstrate that our protocol improvements enable the verification of the key generation process within Ethereum, and similar smart contract platforms.

Leveraging the capabilities provided by distributed ledger-based smart contract platforms, our DKG protocol allows the set of participating entities to be dynamically defined and can incentivize participation as well as penalize adversarial behavior. Further, we are able to ensure that any security deposits provided by participants following the protocol rules always remain safe, even if the DKG protocol itself is executed by a majority of adversarial participants. This design approach can help address the issue of Sybil nodes in settings where *open* participation for better decentralization [10] in the DKG is desirable.

### A. Structure of this Paper

We continue this paper by introducing and comparing related work to our approach in Section II. We describe our system model, including assumptions concerning the network infrastructure, the capabilities of the adversary as well as the security properties expected from DKG protocols, in Section III. Our generalized protocol design for discrete logarithm based cryptosystems is presented and analysed in sections IV and V. Section VI provides implementation specific details, while Section VII provides an overview of our evaluation results. Additional details regarding the evaluation are moved to appendices A, B and C due to space constraints. Finally, we discuss and conclude the paper in sections VIII and IX.

## II. RELATED WORK

The first protocol for DKG was introduced by Pedersen [4] in 1991, and was subsequently built upon within a wide range of publications in the field of threshold cryptography. A popular variant is the so called Joint-Feldman DKG protocol, introduced by Gennaro et al. [1] as a simplification of Pedersen’s work. The core idea of the Pedersen (and the Joint-Feldman) protocol is that each party executes Feldman’s [11] verifiable secret sharing (VSS) protocol, acting as a dealer in order to share a randomly chosen secret among all parties.

After a verification step, ensuring the participants shared their secrets correctly, the resulting group private key is defined by the sum of the properly shared secrets. This private key is unknown to the individual participants, but may be obtained by a collaborating group of parties. The corresponding public key can be computed using the commitments published during the sharing phase with Feldman’s VSS protocol and is the public result of executing the DKG protocol.

However, as described in great detail by the works of Gennaro et al. [1], [6], keys generated using (a wide range of variants of) the Pedersen protocol, are not guaranteed to be uniformly distributed over the respective keyspace. An adversary can bias bits of the resulting key by selectively denouncing the validity of shares of one or more of the parties it controls. Consequently, the set of parties which properly shared their secrets, and thus define the resulting key, is influenced as the denounced parties are excluded. The issue in this case is that honest parties have provided all the information required to compute the resulting public key *before* agreement on the set of shares that are used to create the master key is reached, allowing the adversary to influence the final outcome.<sup>2</sup>

Gennaro et al. [1] presents mitigation strategies against these kind of attacks. However, their approach adds complexity as it requires an additional secret sharing step using Pedersen’s VSS protocol [13]. Canetti et al. [14] extend the solution from Gennaro et al. to cope with adaptive adversaries, which may corrupt parties based on prior knowledge gathered during the protocol execution. More recently, Neji et al. [9] describe a different countermeasure which we adopt in this paper, avoiding these drawbacks.

Kate and Goldberg [7] were the first to study DKG in an asynchronous communication model, whereas synchronous message delivery was previously assumed. In order to support these weaker assumptions, they require a network of  $n \geq 3t + 2f + 1$  participants, out of which  $t$  are controlled by the adversary and thus considered Byzantine and  $f$  parties may fail in the crash-stop model. This is in contrast to works in the style of Gennaro et al. and our protocol, which require synchrony but can tolerate ( $n \geq 2t + 1$ ) Byzantine adversaries. In a subsequent extension of their work [8], Kate et al. provide an implementation, tested with up to 70 parties distributed over multiple continents. A crucial distinction between Kate and Goldberg’s work and the approach followed by Gennaro et al. and this paper, is that the former also implement a Byzantine agreement protocol alongside the DKG, whereas the consensus protocol is not part of the DKG specification in the latter. We outline the advantages and drawbacks of both design decisions in our discussion (see Section VIII).

To the best of our knowledge, the DKG protocol developed by the Orbs Network team [15] is the only publicly available protocol targeted at a similar deployment scenario, namely, an implementation of a DKG protocol using the Ethereum platform. However, the presented prototypical implementation appears to be incomplete and has not been updated since 5<sup>th</sup> August, 2018. A peer-reviewed publication outlining the details of the protocol is also not available at the time of

writing. Further, this approach, in comparison to the works of Gennaro et al., Kate and Goldberg, and our work, fails to guarantee liveness under adversarial behavior. It requires a protocol restart even if only a single adversarial participant sends an invalid share – a major drawback we can avoid.

### III. SYSTEM MODEL

Using our protocol, a set of  $n$  participants  $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$  wish to jointly generate a master secret/public key pair of the form  $mpk = g^{msk}$  for a discrete logarithm based threshold cryptosystem. We use  $g$  and  $h$  to denote two independently<sup>3</sup> selected generators of the group  $\mathbb{G}_q$  with prime order  $q$  and assume that computing discrete logarithms in  $\mathbb{G}_q$  is hard and the Computational Diffie-Hellman (CDH) assumption holds for this group. The master public key  $mpk$  is the (public) output of the protocol. The corresponding (virtual) secret key  $msk$  is shared among the participants, and may be obtained by pooling the shares from  $t + 1$  collaborating parties. Depending on the use case scenario, it may not be desirable or even necessary to ever obtain  $msk$ . For instance, by employing BLS threshold signatures [16], a signature verifying under the master public key  $mpk$  can be obtained by aggregating signature shares without recovering  $msk$  first.

#### A. Communication Model

We assume all parties can monitor and broadcast messages on a shared public and authenticated communication channel. Further, all participants are in agreement on a common view and ordering of these broadcast messages. We assume synchrony in the sense that, any message that is broadcast by a participant during some protocol phase is received by all other parties before the next phase starts. In this regard, our communication model is closely related to the notion of public bulletin boards [17].

In contrast to Gennaro et al. [1], we do not consider pairwise private communication links between parties. Instead, we assume that each participant  $P_i \in \mathcal{P}$  generates a *fresh* secret/public key pair  $\langle sk_i, pk_i \rangle$  of the form  $pk_i = g^{sk_i}$  prior to the protocol start and knows the public keys of all other participants.<sup>4</sup> Note that these keypairs are independent of the keys used to establish the authenticated communication channel, and are only used to derive a symmetric encryption key for each sender/receiver pair of nodes. These symmetric keys are then used *once* to ensure the secrecy of the key shares being transmitted in the sharing phase of the protocol.

Blockchain protocols, which allow inclusion of arbitrary data, and other BFT state machine replication and distributed ledger protocols present suitable candidates for such communication channels. In practice, we leverage the Ethereum blockchain as a public authenticated communication channel and consensus protocol, where agreement on message ordering is ensured through the common prefix property [18]. Together with our client software, which enforces appropriate stabilization times to ensure agreement with high probability, the desired guarantees can be achieved. We refer to Section VI-A

<sup>3</sup>I.e. the discrete logarithm  $dlog_g(h)$  between  $g$  and  $h$  is unknown.

<sup>4</sup>Instead of assuming a priori knowledge of the other parties’ keys, an additional registration phase (see Section VI-F) can be used to exchange the public keys.

<sup>2</sup>We refer to the works of Gennaro et al. [6], [12] for an in-depth discussion of the implications of a non-uniform distribution.

for additional details on how the communication channel is instantiated.

### B. Adversarial Model

To ensure secrecy of the generated secret key  $msk$ , we assume that an adversary controls at most  $t$  participants, whereas a collaboration of  $t + 1$  participants is required to derive  $msk$ . A node controlled by the adversary may deviate arbitrarily from the specified protocol. We consider an *adaptive* adversary, in the sense that it can decide which parties to corrupt based on prior observations. However, the adversary is *not mobile*, once a party is corrupted it is considered compromised for the entire protocol execution. To guarantee both, secrecy of the generated key as well as liveness, i.e., that the protocol completes successfully, the adversary must not control more than  $t < n/2$  parties. These are the optimal bounds one can hope to achieve in this setting [1].

### C. Security Properties

In the following, we reiterate on the security properties we aim for and expect from a DKG protocol. Hereby, we follow the definitions given by Gennaro et al. [1] and Neji et al. [9] for *correctness* and *secrecy* and refer to the corresponding works for a more formal definition. The *uniformity* property highlights a shortcoming identified by Gennaro et al. [1] that was not covered by the original Joint-Feldman protocol. Because recent DKG implementations appear to not consider this property, e.g., the Ethereum-based DKG implementation in [15], we use a distinct category to further emphasize this characteristic. *Robustness* ensures that a subset of parties, which want to recover the master secret key, is able to do so under adversarial influence. The definitions of secrecy, uniformity and robustness follow the correctness definitions C3 and C1' from Gennaro et al. We also add a definition for *liveness*, which was not explicitly stated in Gennaro et al.'s work.

*a) Secrecy.*: No information about the master secret key  $msk$  can be learned by the adversary except for what is implied by the value of the master public key  $mpk = h^{msk}$ .

*b) Correctness.*: All sets of  $t + 1$  correct key shares define the same unique master secret key  $msk$  and all honest parties agree on the common value of the master public key  $mpk = h^{msk}$ .

*c) Uniformity.*: The master secret key  $msk$  is uniformly distributed in  $\mathbb{Z}_q$ , and hence the master public key  $mpk$  is uniformly distributed in  $\mathbb{G}_q$ .

*d) Robustness.*: There is an efficient procedure that, on input of the public information of the DKG protocol and  $n$  submitted shares, outputs  $msk$ , even if up to  $t$  invalid shares have been submitted by malicious or faulty participants.

*e) Liveness.*: As long  $t + 1$  nodes are controlled by correct parties, an adversary cannot prevent the protocol from completing successfully.

## IV. PROTOCOL DESCRIPTION

In this section, we present our generalized DKG protocol design for discrete logarithm based cryptosystems. We start

by giving a brief overview of our three consecutive protocol phases, and then describe each phase in detail in sections IV-A, IV-B and IV-C. For implementation specific details we refer to Section VI.

*a) Sharing Phase.*: During the first phase, each participant in  $P_i \in \mathcal{P}$  selects a randomly chosen secret  $s_i \in_R \mathbb{Z}_q$  and subsequently uses Feldman's VSS to share this secret among all parties, such that  $t + 1$  collaborating parties can recover  $s_i$ , in case a malicious party withholds the required information during the key derivation phase. The verification procedure of Feldman's protocol enables the parties to check that received shares are indeed valid.

*b) Dispute Phase.*: During the dispute phase, each party that received one or more invalid shares in the previous phase uses a non-interactive proof technique to convince other parties about the fact that the issuer violated the protocol.

*c) Key Derivation Phase.*: At the beginning of the last phase, a set of qualified parties  $\mathcal{Q} \subseteq \mathcal{P}$  is formed. A party  $P_i$  is part of  $\mathcal{Q}$  if and only if it (i) broadcasted the required information during the sharing phase and (ii) no party broadcasted a valid dispute against  $P_i$  during the dispute phase. In other words, the set  $\mathcal{Q}$  contains all parties which correctly shared their secret and should thus contribute to form the master key pair  $\langle msk, mpk \rangle$ . Finally, for all parties  $P_i \in \mathcal{Q}$  the values  $h^{s_i}$ , related to the randomly chosen secrets  $s_i$ , are either revealed or recovered and used to derive the master public key  $mpk$ . Using Lagrange interpolation,  $msk$  can be computed after pooling the shares from  $t + 1$  parties. However, depending on the use case scenario, it may not be desirable or necessary to ever obtain  $msk$ .

### A. Sharing Phase

*a) Share Generation.*: At the beginning of the sharing phase, each party  $P_i \in \mathcal{P}$  executes the first step of the Joint-Feldman DKG protocol [1]. In order to share a randomly chosen secret  $s_i \in_R \mathbb{Z}_q$  among all<sup>5</sup> registered parties,  $P_i$  acts as the dealer in a  $(n, t)$  Feldman VSS protocol [11]. For this purpose it picks a secret polynomial  $f_i : \mathbb{Z}_q \rightarrow \mathbb{Z}_q$  with coefficients  $c_{i0} = s_i$  and  $c_{i1}, c_{i2}, \dots, c_{it}$  drawn uniformly at random from  $\mathbb{Z}_q$ :

$$f_i(x) = c_{i0} + c_{i1}x + c_{i2}x^2 + \dots + c_{it}x^t \pmod{q} \quad (1)$$

Then  $P_i$  computes the shares  $s_{i \rightarrow j} = f_i(j)$  for all  $P_j \in \mathcal{P}$ , and the commitments  $C_{i0} = g^{c_{i0}}, C_{i1} = g^{c_{i1}}, \dots, C_{it} = g^{c_{it}}$  to the coefficients of  $f_i(\cdot)$ . These commitments are used in the verification process for the shares and implicitly define  $P_i$ 's public polynomial  $F_i : \mathbb{Z}_q \rightarrow \mathbb{G}_q$ :

$$F_i(x) = C_{i0} \cdot C_{i1}^x \cdot C_{i2}^{x^2} \cdot \dots \cdot C_{it}^{x^t} \quad (2)$$

*b) Share Transmission.*: Next, each  $P_i$  has to securely send its shares  $s_{i \rightarrow j}$  to all other parties  $P_j \in \mathcal{P}$ . Contrary to the original description of the Joint-Feldman DKG, we do not assume access to private communication channels between parties, but rather realize the secure sending of the shares using encryption over our public broadcast channel. We use a symmetric key encryption algorithm  $\text{Enc}_{k_{ij}}(\cdot)$  to ensure

<sup>5</sup>For ease of exposition, we assume that  $P_i$  also provides one share for itself.

secrecy of a sent share from  $P_i$  to  $P_j$ . The corresponding encryption key  $k_{ij}$  can be derived non-interactively by both parties:

$$k_{ij} = pk_j^{sk_i} = pk_i^{sk_j} = g^{sk_i sk_j} \quad (3)$$

Notice that this approach is inspired by the techniques used in the Diffie Hellman key exchange protocol [19] and the ElGamal encryption scheme [20].

Finally,  $P_i$  broadcasts the encrypted shares  $\overline{s_{i \rightarrow j}} = \text{Enc}_{k_{ij}}(s_{i \rightarrow j})$  for all  $i \neq j$  as well as the commitments  $C_{i0}, C_{i1}, \dots, C_{it}$  from Feldman's VSS. Each party  $P_j$  monitors the communication channel for messages broadcasted by other participants. Upon receiving encrypted shares and commitments from  $P_i$ ,  $P_j$  decrypts its share to obtain  $s_{i \rightarrow j} = \text{Dec}_{k_{ij}}(\overline{s_{i \rightarrow j}})$ .

*c) Share Verification.*:  $P_j$  employs the verification procedure of Feldman's VSS to check the validity of each share  $s_{i \rightarrow j}$ . A share is valid if and only if the following share verification condition holds:

$$g^{s_{i \rightarrow j}} = F_i(j) \quad (4)$$

In case  $s_{i \rightarrow j}$  is found invalid, further actions are required in the dispute phase. As  $P_i$  only expects to receive a single message from each party, only the first message is processed, any additional messages from the same sender are ignored. In our smart contract based implementation (see Section VI), the smart contract itself ensures that parties can only broadcast a single message during the sharing phase.

## B. Dispute Phase

In case a party  $P_j$  notices that it received an invalid share  $s_{i \rightarrow j}$  from  $P_i$  in the previous phase,  $P_j$  must broadcast a dispute claim in order to ensure that  $P_i$  is excluded from further steps of the protocol execution. Intuitively,  $P_i$  must be excluded because its secret  $s_i$  may not be recoverable by a collaboration of  $t + 1$  correct parties.

In the original description of the Joint-Feldman DKG protocol, an adversarial  $P_j$  can always issue an (unsupported) claim stating that it received an invalid share from a correct  $P_i$ , requiring  $P_i$  to prove adherence to the protocol rules. We flip this notion in the sense that it is  $P_j$ 's obligation to show that  $P_i$  indeed violated the protocol. To accomplish this we use a non-interactive proof technique described below, and can consequently reduce the required number of interactions between parties.

*a) Issuing a Dispute Claim.*: The key idea how  $P_j$  is able to prove that  $P_i$  provided an invalid share  $s_{i \rightarrow j}$  is to publish the key  $k_{ij}$  used for encryption and decryption of the share. Using this key, other parties are able to decrypt the previously distributed share  $\overline{s_{i \rightarrow j}}$  and can, in the same way as  $P_j$  did, verify that  $s_{i \rightarrow j}$  is indeed invalid. To ensure that an adversarial  $P_j$  cannot just publish an invalid key  $k'_{ij}$ , which would again lead to a false accusation of  $P_i$ , it is required that  $P_j$  proves the correctness of  $k_{ij}$ . We use a common non-interactive zero-knowledge (NIZK) proof technique for showing the equality of the two discrete logarithms [21], [22] to show the correctness of  $k_{ij}$ . The corresponding proving and verification procedures are denoted by  $\text{DLEQ}(x_1, y_1, x_2, y_2, \alpha)$  and  $\text{DLEQ-verify}(x_1, y_1, x_2, y_2, \pi)$ .

*Procedure 1: DLEQ*( $x_1, y_1, x_2, y_2, \alpha$ ).

To show that  $\text{dlog}_{x_1}(y_1) = \text{dlog}_{x_2}(y_2)$  holds without revealing the discrete logarithm  $\alpha$ , a prover proceeds as follows:

- 1) compute  $t_1 = x_1^w$  adding  $t_2 = x_2^w$  for  $w \in_R \mathbb{Z}_q$
- 2) compute  $c = \text{H}(x_1, y_1, x_2, y_2, t_1, t_2)$
- 3) compute  $r = w - \alpha c \pmod{q}$
- 4) output  $\pi = \langle c, r \rangle$

Instantiating the above procedure,  $P_j$  can prove the correctness of the decryption key  $k_{ij}$  by providing  $\pi(k_{ij}) = \text{DLEQ}(g, pk_j, pk_i, k_{ij}, sk_j)$  in addition to  $k_{ij}$ .

*b) Verifying a Dispute Claim.*: Upon receiving a dispute claim  $\langle k_{ij}, \pi(k_{ij}) \rangle$  against  $P_i$ , issued by  $P_j$ , one can use  $\text{DLEQ-verify}(g, pk_j, pk_i, k_{ij}, \pi(k_{ij}))$  to check the validity of the received key  $k_{ij}$ .

*Procedure 2: DLEQ-verify*( $x_1, y_1, x_2, y_2, \pi$ ).

To check the correctness of a proof  $\pi = \langle c, r \rangle$ , showing that  $\text{dlog}_{x_1}(y_1) = \text{dlog}_{x_2}(y_2)$  holds, a verifier proceeds as follows:

- 1) compute  $t'_1 = x_1^r y_1^c$  and  $t'_2 = x_2^r y_2^c$
- 2) output **VALID** if  $c = \text{H}(x_1, y_1, x_2, y_2, t'_1, t'_2)$  holds  
output **INVALID** otherwise

If the key is found invalid, the dispute claim is invalid. Otherwise, the verification procedure continues by decrypting the corresponding share  $s_{i \rightarrow j} = \text{Dec}_{k_{ij}}(\overline{s_{i \rightarrow j}})$  and checking its correctness according to the share verification condition specified in Equation 4. The dispute is valid if and only if  $k_{ij}$  is found valid but the verification condition does not hold.

The protocol ensures that: (i) In case a correct participant received an invalid share from another party, the share issuer is considered disqualified by all (correct) parties at the end of the dispute phase. (ii) An adversary cannot wrongly accuse any correct party of providing it with an invalid share. (iii) The adversary does not gain any additional information when a party  $P_j$  reveals the values  $k_{ij}$  and  $\pi(k_{ij})$ , because the adversary can always compute (and therefore publish)  $k_{ij} = pk_j^{sk_i}$  using  $P_i$ 's secret key, and the NIZK proof  $\pi(k_{ij})$  does not reveal additional information apart from the correctness of the statement.

## C. Key Derivation

*a) Deriving the Set of Qualified Nodes.*: The first step in the key derivation phase is determining the set of qualified parties  $\mathcal{Q} \subseteq \mathcal{P}$ , describing which parties should contribute to the resulting key pair  $\langle msk, mpk \rangle$ . If we recall the current protocol state at the beginning of the key derivation phase, we observe that each  $P_i \in \mathcal{P}$  has either:

- 1) *correctly* shared its secret  $s_i$  with all other parties.
- 2) *incorrectly* shared its secret  $s_i$ .
- 3) did not share its secret  $s_i$  at all.

We say a secret was correctly shared by  $P_i$ , if and only if no valid dispute claim against  $P_i$  was filed during the dispute phase. Parties which incorrectly shared their secrets, or did not share their secrets at all, are disqualified and excluded from

the upcoming protocol steps. The remaining parties form the set  $\mathcal{Q}$ . In other words, a node  $P_i \in P$  is only part of  $\mathcal{Q}$  if (i) it published the values  $C_{i0}, C_{i1}, \dots, C_{it}$  and  $s_{i \rightarrow j}$  for all  $i \neq j$  during the sharing phase and (ii) no node  $P_j$  filed a valid dispute against  $P_i$  during the dispute phase.

*b) Bias when Computing the Keys Directly.:* Using this definition of the set  $\mathcal{Q}$ , the resulting group public key  $mpk$  could be derived by following the description of the Joint-Feldman protocol:

$$mpk = \prod_{P_i \in \mathcal{Q}} C_{i0} = \prod_{P_i \in \mathcal{Q}} g^{s_i} \quad (5)$$

However, as described in great detail by the works of Gennaro et al. [1], [6], the above approach does not ensure that the resulting key pair is uniformly distributed. An adversary can bias bits of the resulting key by selectively denouncing one or more of its nodes, which influences the set  $\mathcal{Q}$  and thus the resulting key. The critical<sup>6</sup> issue here is, that all information required to compute the resulting public key is known to the adversary *before* the set  $\mathcal{Q}$  is fixed.

*c) Protection against Biasing of the Generated Keys.:* We adopt a recent countermeasure described by Neji et al. [9] to ensure the resulting key is uniformly distributed. The key idea to ensure uniformity is to instead compute  $mpk$  as follows:

$$mpk = \prod_{P_i \in \mathcal{Q}} h^{s_i} \quad (6)$$

Here,  $h$  is used to denote an additional generator of the group  $\mathbb{G}_q$ , such that  $dlog_g(h)$  is unknown. The required values  $h^{s_i}$  used to compute  $mpk$  are published by the parties in  $\mathcal{Q}$  after this set is fixed. Each value  $P_i$  shows the correspondence between the values  $h^{s_i}$  and  $C_{i0} = g^{s_i}$  using the NIZK proof  $\pi(h^{s_i}) = \text{DLEQ}(g, g^{s_i}, h, h^{s_i}, s_i)$  as introduced in Section IV-B. In case any (adversarial) party  $P_i \in \mathcal{Q}$  does not reveal its value  $h^{s_i}$  and a valid proof  $\pi(h^{s_i})$  by the end of the key derivation phase, a set of  $t + 1$  correct parties is always able to use the recovery procedure of Feldman's VSS to obtain  $s_i$  and consequently  $h^{s_i}$  anyway. Without loss of generality, let  $\mathcal{R} \subseteq \mathcal{Q}$  denote a set of  $t + 1$  correct parties. Then,  $s_i$  is obtained via Lagrange interpolation:

$$s_i = \sum_{P_j \in \mathcal{R}} s_{i \rightarrow j} \prod_{\substack{P_k \in \mathcal{R} \\ j \neq k}} \frac{k}{k - j} \quad (7)$$

*d) Deriving the Keys.:* Finally, the common master public key  $mpk$  can be derived as specified in Equation 6 using the published or recovered values  $h^{s_i} \mid P_i \in \mathcal{Q}$ . Additionally, each  $P_j \in \mathcal{Q}$  can compute its individual group key pair  $\langle gsk_j, gpk_j \rangle$ :

$$gsk_j = \sum_{P_i \in \mathcal{Q}} s_{i \rightarrow j} \quad gpk_j = h^{gsk_j} \quad (8)$$

In order to enable a third party to verify  $gpk_j$ ,  $P_j$  provides the values  $g^{gsk_j}$  as well as a correctness proof  $\text{DLEQ}(g, g^{gsk_j}, h, gpk_j, gsk_j)$ . The verifier accepts  $gpk_j$  as valid if checking of the proof via  $\text{DLEQ-verify}(\cdot)$  succeeds,

and the verification of  $g^{gsk_j}$  using the previously committed public polynomials is successful:

$$g^{gsk_j} = \prod_{P_i \in \mathcal{Q}} F_i(j) \quad (9)$$

The corresponding master secret key  $msk$  is shared among all nodes in  $\mathcal{Q}$  and can be obtained as follows:

$$msk = \sum_{P_i \in \mathcal{Q}} s_i \quad (10)$$

In case  $P_i$  does not reveal its secret  $s_i$ , it can always be computed by  $t + 1$  collaborating parties, because each  $P_i \in \mathcal{Q}$  has correctly shared  $s_i$  among the parties during the first protocol phase. Alternatively, a set of  $t + 1$  collaborating parties, denoted by  $\mathcal{R}$ , can also derive the master secret key  $msk$  via Lagrange interpolation from their group secret keys:

$$msk = \sum_{P_j \in \mathcal{R}} gsk_j \prod_{\substack{P_k \in \mathcal{R} \\ j \neq k}} \frac{k}{k - j} \quad (11)$$

However, for many threshold cryptographic applications  $msk$  might never be computed at a single location. Considering, e.g. BLS threshold signatures,  $t + 1$  collaborating parties might produce a signature  $\sigma$  on message  $m$  which verifies under the public key  $mpk$ . For this purpose, each of these parties  $P_j$  uses its individual group signing key  $gsk_j$  to issue a partial signature for  $m$ , which upon aggregation form  $\sigma$ . There is no need to compute the master secret key  $msk$  in order to issue the signature in this scenario.

## V. SECURITY ANALYSIS

For brevity, we omit a detailed analysis of the guarantees in regard to *correctness* and *uniformity* in this paper, as the corresponding security proofs provided by Gennaro et al. [1] and Neji et al. [9] directly apply to our protocol. We hence refer the reader to the aforementioned publications for further details.

*a) Secrecy.:* In order to show that the original security proof regarding secrecy still applies, we show that the dispute process we introduce as alternative to the steps described by Neji et al. [9] does not provide the adversary with any additional information, and hence preserves secrecy. Specifically, any information a correct node  $P_i$  secretly transfers to another correct node  $P_j$  must remain hidden from the adversary to ensure it cannot reconstruct the master secret key  $msk$  from those messages. The only point in time when information is exchanged secretly, is the share transmission step (see Section IV-A). Here a correct party  $P_i$  always encrypts the share  $s_{i \rightarrow j}$  it sends to  $P_j$  using a symmetric key encryption algorithm  $\text{Enc}_{k_{ij}}(\cdot)$ . Under the Computational Diffie-Hellman assumption, the shared key  $k_{ij}$  used for en-/decryption can only be derived using secret information  $sk_i$  or  $sk_j$  from node  $P_i$  or  $P_j$ . However, neither  $P_i$  nor  $P_j$  reveal this information or  $k_{ij}$  itself during the protocol execution if they are both honest.

If we instead consider the case where  $P_i$  is honest but  $P_j$  is controlled by the adversary, the adversary also does not gain any additional information. In this case, the only point in time a honest node  $P_i$  would publish additional information, namely  $k_{ij}$  and the corresponding correctness

<sup>6</sup>See [6], [12] for an in-depth discussion on the implications of non-uniform distribution.

proof  $\pi(k_{ij}) = \text{DLEQ}(g, pk_i, pk_j, k_{ij}, sk_i)$ , is during the process of issuing a dispute claim (Section IV-B). However, being the intended communication partner, the adversary was already able to derive  $k_{ij} = pk_i^{sk_j}$  (and thus obtain  $s_{i \rightarrow j}$ ) as part of the protocol. Hence, no additional information is revealed when  $P_i$  publishes  $k_{ij}$ . Furthermore, e.g. as outlined by Camenisch and Stadler [22], the NIZK proof  $\pi(k_{ij})$  does not reveal any information in addition to correctness of  $k_{ij}$ , in particular does not reveal any information about  $sk_i$ .

*b) Robustness.:* Robustness requires an efficient procedure, that recovers the master secret key  $msk$  from a set of at least  $t + 1$  correct shares. However, this set may additionally contain up to  $t$  invalid shares provided by the adversary. We obtain such a procedure, by first checking the validity of a provided share  $gsk_i$  using the verification condition specified in Equation 9. Lagrange interpolation is then used to compute  $msk$  from any set of  $t + 1$  valid shares (see Equation 11).

*c) Liveness.:* In our synchronous system model, the protocol always reaches the beginning of the key derivation phase, as the sharing and dispute phases always end after a fixed amount of steps (the respective number of blocks per phase). Consequently, the completion of the key derivation phase (and thus the completion of the protocol), depends on the nodes' ability to gather all the information required to compute  $mpk$  from the values  $h^{s_i}$  provided by all  $P_i \in \mathcal{Q}$ . Each correct node in the set of qualified nodes  $\mathcal{Q}$ , publishes this value at the beginning of the phase. However, up to  $t$  adversarial nodes, which completed the sharing and dispute phase successfully, and are thus part of  $\mathcal{Q}$ , might not reveal the respective values. In this case, the correct parties obtain all missing values  $h^{s_i}$  by recovering  $s_i$  using Lagrange interpolation from their shares for  $s_i$  (see Section IV-C for additional details). This process requires the collaboration of at least  $t + 1$  correct nodes, and thus completes successfully for configurations where the adversary controls at most  $n - t - 1$  nodes.

## VI. IMPLEMENTATION

To highlight the feasibility and practicality of our approach, we present a prototype implementation. It consists of two parts: (i) an Ethereum smart contract serving as the communication and verification platform, and (ii) a client application written in Python and executed locally by each participant. Both implementations are open source and publicly available on Github <https://github.com/PhilippSchindler/EthDKG/>.

In the following, we describe the steps required to apply our generalized protocol description for the concrete use case of deriving key pairs to be used with the BLS signature scheme. Thereby, we outline (i) how our communication model can be realized, (ii) which techniques are necessary to efficiently implement the required cryptographic primitives, and (iii) how the protocol execution can be verified at the smart contract level, despite the limitations of the Ethereum platform. The BLS signature scheme was chosen not only because Ethereum has built-in support for a pairing friendly elliptic curve which can be used with BLS, but also due to the wide range of desirable properties this signature scheme provides for different application scenarios. These properties include short signature size, non-interactive aggregation capabilities as well as signature uniqueness. For additional details on BLS

signatures, their properties and use cases we refer the reader to the original descriptions [16], [23], [24].

When using our protocol for BLS signatures, a set of parties first executes our DKG protocol to compute a master BLS key pair  $\langle msk, mpk \rangle$ . The public key  $mpk$  is published and verified within the smart contract, whereas the (virtual) secret key  $msk$  is shared among the parties. Each party  $P_i$  is then capable of using its individual signing key  $gsk_i$  to sign messages with BLS. Any set of  $t + 1$  valid<sup>7</sup> signatures on a common message can be combined to form a threshold signature, which verifies under  $mpk$ , for that message. This aggregation process can be performed without necessitating on-chain transactions within Ethereum. Furthermore, the cost of verifying the resulting threshold signature within the smart contract does not depend on the number of participants or signers.

### A. Realizing our Communication Model

Revisiting the assumptions from our protocol description (see Section III-A), we require a shared agreed-upon authenticated broadcast channel and adherence to certain synchrony assumptions to separate the different protocol phases. These assumptions are realized as follows:

*a) Ethereum as a Broadcast Channel.:* In our implementation, each participant of the DKG protocol actively monitors the Ethereum blockchain. In particular, clients monitor all transactions to the address of the pre-deployed DKG contract. A message is broadcast by issuing a transaction that calls a function within the DKG smart contract when the transaction is mined within a block in the Ethereum network. Upon being called successfully, the contract triggers Ethereum events, which are processed by the client implementation.

*b) Agreement.:* After detecting the emission of a new event, the client software of each participant waits for a sufficient number<sup>8</sup> of confirming blocks. This ensures that all nodes agree on a common history of blocks, and consequently on the triggered events and their order w.h.p, before they react to the events. This requirement is a direct consequence of the fact that the Ethereum blockchain may fork and thus does not provide immediate agreement on newly mined blocks.

*c) Message Authentication.:* The requirements in regard to message authenticity are directly supported by Ethereum. In fact, Ethereum enforces that all transactions are cryptographically signed by the issuer in order to be processed.

*d) Synchrony Assumptions.:* Our synchrony assumptions can be realized by specifying the start and end of each protocol phase based on appropriate relative Ethereum block heights. Liveness, i.e. ensuring the protocol completes successfully even under adversarial conditions, critically depends on the ability of correct nodes to timely disseminate information. Consequently, it has to be ensured that any transaction a node issues at the beginning of a protocol phase is confirmed, and

<sup>7</sup>The process is robust in the sense that the validity of an individual signature can also be checked using the issuer's public key.

<sup>8</sup>For an in depth discussion on the required number of confirmations we refer to the works of Gervais et al. [25] and Sompolinsky and Zohar [26]. We furthermore provide concrete values for this manner in our evaluation (see Appendix C).

consequently received by all other correct nodes, by the beginning of the next phase. The required phase durations depend on a range of factors including: the number of participants, the state of the Ethereum network, and the amount of transaction fees the participants are willing to pay. Thus they need to be analysed on a case by case basis or selected conservatively. We provide an evaluation in regards to the required durations, considering the network conditions at the time of writing as well as a general description, in Section VII.

### B. Cryptographic Primitives

When leveraging a smart contract-based DKG implementation that is capable of performing the verification steps on-chain, an efficient implementation of the underlying cryptographic primitives can be crucial for a low cost protocol design. Within the Ethereum platform, only a limited range of so called pre-compiled contracts for elliptic curve cryptography are available currently. The supported operations target the groups  $\mathbb{G}_1$ ,  $\mathbb{G}_2$  and  $\mathbb{G}_T$  of prime order  $q$ , defined on the elliptic curve BN254 [27], [28] and include point/point addition ( $\mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_1$ ), point/scalar multiplication ( $\mathbb{G}_1 \times \mathbb{Z}_q \rightarrow \mathbb{G}_1$ ) and a verification procedure for the pairing  $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ . We rely upon these operations to efficiently implement the verification procedures for our DKG, targeting the generation of keys for the BLS signature scheme.

As BLS public keys reside in  $\mathbb{G}_2$ , most of the operations required for our protocol would use group  $\mathbb{G}_2$ , if we directly apply our general protocol description. However, as of the current Ethereum release, computations in  $\mathbb{G}_2$  are not natively supported, and implementing the required operations using available Ethereum Virtual Machine (EVM) opcodes would lead to very high gas consumption and thus render the approach inefficient.<sup>9</sup> Fortunately, the corresponding operations in group  $\mathbb{G}_1$  and a verification procedure for the pairing  $e$  exist as pre-compiled contracts in Ethereum [30], [31]. This allows us to efficiently perform operations in  $\mathbb{G}_1$  and verify the corresponding element in  $\mathbb{G}_2$  using the pairing check within the smart contract. In the following sections VI-C, VI-D and VI-E, we outline the details for incorporating this approach into our protocol design.

### C. Sharing Phase

During the sharing phase, each participant  $P_i \in \mathcal{P}$  proceeds as specified in our general protocol description (see Section IV-A). In particular,  $P_i$  shares a secret  $s_i \in_R \mathbb{Z}_q$  among all parties in  $\mathcal{P}$  using Feldman’s VSS protocol. The commitments  $C_{i0}, C_{i1}, \dots, C_{it}$  are group elements from  $\mathbb{G}_1$ :  $C_{ik} = g_1^{C_{ik}} \mid 0 \leq k \leq t$ , where  $g_1$  denotes a generator of  $\mathbb{G}_1$ . Because there are no primitives for symmetric encryption available within Ethereum, we realize the encryption and decryption algorithms  $\text{Enc}_{k_{ij}}(\cdot)$  and  $\text{Dec}_{k_{ij}}(\cdot)$  using a one time pad, where we derive a unique key from  $k_{ij}$  and  $j$  by using a cryptographic hash function<sup>10</sup>  $H(\cdot)$ :

$$\begin{aligned} \text{Enc}_{k_{ij}}(s_{i \rightarrow j}) &= s_{i \rightarrow j} \oplus H(k_{ij} \parallel j) \\ \text{Dec}_{k_{ij}}(\overline{s_{i \rightarrow j}}) &= \overline{s_{i \rightarrow j}} \oplus H(k_{ij} \parallel j) \end{aligned}$$

<sup>9</sup>A Solidity implementation of a single multiplication of a group element from  $\mathbb{G}_2$  with a 256 bit scalar requires approximately 2 000 000 gas [29],

<sup>10</sup>In our implementation, the value  $s_{i \rightarrow j}$  and the output of the used cryptographic hash function are 256 bits each.

To ensure that such a simple approach is secure in practice, it is crucial that (i) the pads used for encryption of messages between honest parties are indeed used only once, and (ii) the encrypted data is additionally protected against malleability. For two distinct honest parties  $P_i$  and  $P_j$ , the value of  $k_{ij}$  is defined by the values of the randomly generated private keys  $sk_i$  and  $sk_j$ , and is thus unique. Combining this unique value with the index of the share receiver  $j$  further ensures that the one time pads used to encrypt the single message from  $P_i$  to  $P_j$  and the single message from  $P_j$  to  $P_i$  are encrypted with different pads. Consequently, criterion (i) is met. Also criterion (ii) is fulfilled, as the encrypted values are transmitted as part of Ethereum transactions, which are signed and published on the broadcast channel and thus protected against malleability. To publish the required information, namely the encrypted shares  $\overline{s_{i \rightarrow j}}$  for all  $i \neq j$  and the commitments  $C_{i0}, C_{i1}, \dots, C_{it}$ , the client constructs and broadcasts the corresponding Ethereum transaction, invoking the pre-deployed smart contract.

The smart contract ensures that only *eligible* parties, i.e.  $P_i \in \mathcal{P}$  may provide a *single, well-formed* message. The set of eligible parties is either specified statically at the time of creation of the smart contract, or via a dynamic registration process as described in Section VI-F. A message is considered well-formed, if it contains exactly  $n - 1$  encrypted shares, and  $t + 1$  commitments to the coefficients of the secret sharing polynomial. Upon receiving a well-formed transaction from an eligible party, the smart contract notifies all other participants about the published information using an Ethereum event. The contents of the encrypted shares and the validity of the commitments are not verified at this point in time. Instead, the verification is only performed on demand, i.e. in case a dispute is submitted in the next protocol phase. In order to verify a potential dispute in the next phase, the smart contract stores a cryptographic hash of the message content. As we see in Section VI-D, the hash is sufficient to fully verify a potential dispute. It would also be possible to store the entire message instead of the digest. However, storing only the hash significantly reduces the amount of on-chain storage required, and thus lowers transactions fees, in particular for large  $n$ .

### D. Dispute Phase

In case a party  $P_j$  finds that  $P_i$  provided an invalid share for  $s_i$ ,  $P_j$  follows our general protocol description to publish a dispute. For this purpose, it constructs a transaction which, in addition to  $k_{ij}$  and  $\pi(k_{ij})$ , includes the message content sent by  $P_i$  in the previous protocol phase. This enables the smart contract to recompute and compare the hash of  $P_i$ ’s message with the stored value. If the hashes do not match, the dispute is found invalid and the smart contract aborts. Otherwise the smart contract has all information required to perform a full verification. In particular, it can verify that the encrypted share  $\overline{s_{i \rightarrow j}}$  present in the dispute transaction is indeed the share  $P_i$  previously distributed. The verification continues as stated in Section IV-A. The corresponding computations can efficiently be performed using the Ethereum pre-compiled contracts [30] for arithmetic in  $\mathbb{G}_1$ . If the dispute is considered valid, the share issuer is flagged as adversarial and thus excluded from the set  $\mathcal{Q}$  in the key derivation phase. Additionally, the smart contract triggers a corresponding event to notify all parties about the successful dispute. Optionally the issuer may be economically punished, and a security deposit could be used

to refund the disputer for its transaction fees. Similarly, an adversarial disputer could be penalized for submitting an invalid dispute. In either case, the contract may not process a dispute transaction against an already disqualified participant. In fact, in this scenario, our implementation of the smart contract aborts immediately in order to save transaction fees.

### E. Key Derivation Phase

Again, we closely follow our protocol specification from Section IV-C to implement the key derivation phase. Similar to the definition of  $h$ , we use  $h_1 \in \mathbb{G}_1$  and  $h_2 \in \mathbb{G}_2$  to denote independently selected generators for the groups  $\mathbb{G}_1$  and  $\mathbb{G}_2$ .

As a first step, each  $P_i \in \mathcal{Q}$  computes the values  $h_1^{s_i}$  and the corresponding NIZK proof  $\pi(h_1^{s_i})$  showing its correctness. The corresponding computations are performed in group  $\mathbb{G}_1$ . However, as the master public key  $mpk = h_2^{msk}$  is an element of  $\mathbb{G}_2$ ,  $P_i$  is also required to map its key share  $h_1^{s_i}$  to  $\mathbb{G}_2$ , i.e. compute  $h_2^{s_i}$ . Then,  $P_i$  crafts and publishes a transaction, containing  $h_1^{s_i}$ ,  $\pi(h_1^{s_i})$  and  $h_2^{s_i}$ . As described, a collaboration of  $t+1$  parties recovers  $s_i$  (and thus  $h_1^{s_i}$ ,  $\pi(h_1^{s_i})$  and  $h_2^{s_i}$ ) in case  $P_i$  does not publish the required information by the end of the key derivation phase. This recovery process can be performed either with or without interaction with the Ethereum platform. We opted for using Ethereum for this purpose, instead of adding complexity to the design by implementing an additional off-chain communication channel. After completing the recovery, any one of the involved parties can issue the corresponding transaction on behalf of  $P_i$ . Either way, it is ensured that  $h_1^{s_i}$ ,  $\pi(h_1^{s_i})$  and  $h_2^{s_i}$  become public and available for the smart contract for all  $P_i \in \mathcal{Q}$ . The smart contract can verify the provided information with the DLEQ-verify( $\cdot$ ) procedure and use the precompiled pairing contract [31] to check the validity of  $h_2^{s_i}$ . The value  $h_2^{s_i}$  is considered correct if  $e(h_1^{s_i}, h_2) = e(h_1, h_2^{s_i})$  holds.

Finally, any party can compute and publish the master public key  $mpk = \prod_{P_i \in \mathcal{Q}} h_2^{s_i}$  and  $mpk^* = \prod_{P_i \in \mathcal{Q}} h_1^{s_i}$ . The smart contract can recompute  $mpk^*$  and use the pairing  $e(\cdot)$  to verify the correctness of  $mpk$ .

### F. Dynamic Participation

The utilization of an open smart contract platform such as Ethereum also enables us to readily implement dynamic participation strategies. If the choice is made to employ this protocol feature, the set of participants  $\mathcal{P}$  which run the DKG protocol is not defined a priori, but rather obtained in an additional *registration phase*, executed at the beginning of the protocol. For this purpose, the creator of the corresponding smart contract specifies a set of participation rules at the time of contract creation. A participation rule specifies under which condition a particular Ethereum account is allowed to “join” the set  $\mathcal{P}$ . Within the limitations of the Ethereum platform, arbitrary smart contract code can be used to define participation rules. In the following, we provide basic examples for participation rules while more elaborate and robust schemes against adversarial behavior are left to future work.

- 1) *First come, first serve*: Only the first  $N$  parties to register are allowed to join the protocol.

- 2) *Security deposit*: Only parties, which provide a security deposit of at least  $X$  Ether are allowed to join the protocol.
- 3) *Highest bidding*: The  $N$  parties, which provided the highest amount of security deposit are allowed to join the protocol.

For conditions 1 and 2 the participation rules are checked as soon as a registration transaction is included in an Ethereum block. Only upon success is the issuer of the transaction added to the set  $\mathcal{P}$ , tracked within the smart contract. The implementation of condition 3 is rendered slightly more complex. In this case, the smart contract keeps track of the set  $\mathcal{P}$  consisting of up to  $N$  participants and their provided security deposits. Upon registration of party  $P_{N+1}$ , the registration is accepted if the deposit provided is bigger than the smallest deposit received so far. If this is the case, the registration is accepted by adding  $P_{N+1}$  to the set  $\mathcal{P}$  and removing the participant with the smallest deposit from  $\mathcal{P}$ . Otherwise the registration is rejected and  $\mathcal{P}$  remains unchanged.

## VII. EVALUATION

Due to space constraints, this section only provides a brief overview of our findings, while the detailed evaluation results are provided in appendices: A (computational costs for all interaction between the parties and the smart contract), B (communication complexity), and C (execution time).

Even in demanding scenarios (tested with up to 256 nodes), each step in our implementation the DKG protocol can be verified at the smart contract level well within the Ethereum block gas limit. However, the overall costs of running the DKG protocol depends on various highly fluctuating factors such as the exchange rate of ETH to, e.g., USD, or the gas price depending on the current network load. Therefore, it is difficult to provide accurate execution cost estimates. For example, at the time we initially performed our evaluation, Ethereum gas prices of 2 GWei<sup>11</sup> were recommended, whereas at the time of writing 50 GWei are common. Combining the increase in gas prices with the increase in the ETH to USD exchange rate, our approach, while technically feasible, is currently rendered costly for scenarios with a high number of nodes, whereas our original estimate of \$1.68 per participant in a 256 node scenario, was of little practical concern. We note that our solution can also be deployed on other EVM compatible ledgers that currently offer markedly lower transaction fees compared to Ethereum.

## VIII. DISCUSSION

*a) Model.*: In our DKG protocol, we follow the model described in the theoretical works of Gennaro et al. [1]. Consequently, we inherit three important characteristics for our protocol: (i) the synchronous communication model, (ii) the separation of the underlying consensus platform and the DKG protocol itself, (iii) the optimal threshold  $t$ , i.e. secrecy and liveness for all  $t < n/2$ . These are in contrast to the properties of the more recent works by Kate et al. [7], [8], which consider an asynchronous communication model. While these works still require a weak synchrony assumption [32] to

<sup>11</sup>1 GWei =  $10^{-9}$  ETH

ensure liveness, the protocol’s safety guarantees do not depend on timing assumptions of the underlying message delivery network. To mitigate this risk in a synchronous protocol design, the corresponding timings, i.e. the number blocks in each protocol phase for our protocol, have to be selected appropriately.

A drawback of moving to the asynchronous model, is a reduced resilience against Byzantine adversaries. In the hybrid failure model ( $n = 3t + 2f + 1$ ), described by Kate et al., the protocol can only tolerate less than  $1/3$  Byzantine parties ( $t$ ), and less than  $1/2$  crashed participants ( $f$ ). Here, our protocol design can prove advantageous as it ensures the desired security properties, in particular secrecy and liveness, with up to  $n = 2t + 1$  participants.

*b) Secrecy / Liveness Trade-off.*: Our protocol design enables the use of different values for the parameter  $t$ , specifying the threshold for the underlying secret sharing protocol, depending on the specific application scenario. The choice of  $t$  directly incurs a trade-off between liveness and secrecy. If an adversary controls at most  $t$  nodes, secrecy is ensured, whereas at least  $t + 1$  honest nodes are required to guarantee liveness. For example, setting  $t = n$ , ensures that as long as there is at least one honest participant, the master secret key  $msk$  cannot be learned by the adversary. On the contrary, even a single adversarial node can prevent successful completion of the protocol. In practice the choice of  $t$  is directly related to the application scenario. If we consider, for example, a synchronous BFT protocol in a setting with  $n = 2f + 1$  participants,  $t$  is set to equal  $f$ , whereas a typical requirement in asynchronous or particularly synchronous BFT protocols, i.e. that more than  $2/3$  of the parties have to sign a particular state or message, is supported by setting  $t = \lceil 2/3n \rceil - 1$ .

*c) Uniform Key Distribution.*: During the key derivation phase, we follow Neji et al. [9] to implement a protection mechanism, which prevents the adversary from biasing bits of the generated key pair. While the implemented countermeasure does not require a full additional secret sharing round, it requires up to two<sup>12</sup> additional transactions issued by all participants. To save these costs and reduce the protocol’s complexity, one might decide to omit the additional steps required to ensure uniform distribution of the key pair. Instead, each party  $P_i$  publishes a commitment  $H(C_{i0})$  to the value  $C_{i0}$  prior to the sharing phase. The values  $C_{i0}$ , published during the sharing phase, are only accepted if they match the corresponding commitment. During the key derivation phase, the master public key  $mpk$  is directly computed as described for the Joint-Feldman protocol (see Equation 5). Such a design decision may be useful e.g. in a deployment scenario, where we expect the DKG protocol to complete without any errors, i.e. in a scenario where we assume that it is very likely that all participants follow the protocol accordingly. However, as described in Appendix A, the additional costs required to achieve uniformity do not add much overhead to the overall protocol execution. Consequently, we recommend to use our protocol design without this modification for most practical scenarios.

*d) Ethereum as Communication Infrastructure.*: As described in Section III-A, a key component necessary for the implementation of our DKG protocol is a suitable communication layer. Using an existing distributed ledger that provides Byzantine fault tolerance and agreed upon total ordering of exchanged messages. Although our approach may also be used on top of traditional BFT protocols or other available blockchain platforms, we decided to use an existing blockchain platform, namely Ethereum, instead of deploying our own communication infrastructure. If we compare our solution to the protocol described by Kate et al. we observe a key difference in the design approach: whereas in our protocol, the core functions of the DKG protocol are separated from the underlying consensus mechanism, Kate et al. describe their protocol in a standalone setting, intertwining a custom BFT protocol with the DKG logic. We see advantages in both approaches, depending on the application scenario. While the technique we present can benefit from an easier deployment and a simplified protocol design due to the separation of concerns, the security of Kate et al.’s approach does not depend on an external consensus mechanism and can hence operate in a stand-alone setting.

*e) On-Chain Verification.*: While on-chain verification is not required for the core functionality of the protocol, it immediately provides a range of benefits: e.g. other applications on the Ethereum platform can be assured that the master public key was correctly computed, and can thus safely use this key to verify threshold signatures issued under the corresponding (shared) secret key. Furthermore, including monetary incentive mechanisms allows us to define a wide range of interesting dynamic and possibly open participation models. It is no longer required to define the set of parties  $\mathcal{P}$ , executing the protocol, prior to the protocol start. Instead, the smart contract logic can be used to specify under which conditions a party is allowed to join the protocol. When on-chain verification is not used, clients can still fully verify the protocol execution. However, the lack of on-chain verification also comes with the disadvantage, that seizing a security deposit becomes more difficult and potentially places honest clients at risk. It is no longer possible to seize the deposit automatically during the submission process of a dispute, as the smart contract does not perform the corresponding verification steps. A partial mitigation strategy is that a majority of the participants of the DKG verify a dispute off-chain and confirm its validity. However, this leads to the issue that a honest party’s security deposit may be seized if the DKG protocol is run by an adversarial majority. This is in contrast to the approach with on-chain verification, which always ensures that the deposit of correct party remains safe.

*f) Implementation and Scalability.*: To the best of our knowledge, there exist no implementations of a DKG protocol following Gennaro et al.’s design, despite the extensive theoretical research in this direction. Our protocol can be seen as a first realization of this theoretical line of research. It is implemented and evaluated using the Ethereum platform as a communication layer. Consequently, the scalability of our approach is limited by the computational capacities available and transaction fees required to execute transactions on Ethereum. Our measurements (see Appendix A) show that even in a demanding scenario with 256 participants, all transactions can be executed well within Ethereum’s current block gas limit.

<sup>12</sup>one transaction for publishing the key share  $h^{s_i}$  and proof  $\pi(h^{s_i})$ , and potentially an additional message for recovering any missing key shares

However, at the time of writing, the recent steep increase in gas cost in Ethereum due to its rise in popularity and price speculation has increased transaction the overall recommended network fees, introducing economic limitations especially for scenarios with a large number of nodes. Nevertheless, we expect fees to eventually return to lower levels as protocol improvements increase scalability and furthermore outline that our solution can be deployed on other EVM compatible ledgers with lower transaction fees.

The protocol design by Kate and Goldberg [7] was implemented and evaluated in subsequent work [8], performing tests of their implementation with up to 70 nodes on the PlanetLab platform. While in terms of execution time for small numbers of nodes, our solution is one order of magnitude slower than the completion times reported by Kate et al. [8], the parameters we use in our evaluation (see Appendix C) are selected conservatively, and only use 10% of Ethereum’s block capacity. Kate et al.’s protocol execution time increases sharply with an increasing number of nodes as the communication complexity of their protocol is  $\mathcal{O}(n^4)$ . Our evaluation shows that the communication complexity of our protocol is within  $\mathcal{O}(n^3)$ , while the amount of data processed on Ethereum is  $\mathcal{O}(n^2)$ . This leads to an approximate doubling in execution time when increasing the number of participants from 128 to 256.

## IX. CONCLUSION

We present EthDKG, a new state of the art protocol for distributed key generation, that demonstrates how to efficiently implement an improved variant of Gennaro et al.’s [1] theoretical work. Our enhancements include a new mechanism to resolve disputes, which arise if certain parties violate the protocol rules, as well as a range of techniques improving the performance of our implementation in practice. We outline that our tailored protocol design can readily be executed on existing blockchain infrastructures. In particular, we show that all verification steps required during the protocol execution can be performed efficiently within the constrained EVM environment of the Ethereum platform. By leveraging the Ethereum blockchain, or an alternative platforms with similar guarantees, we are able to decouple the implementation of the underlying consensus protocol and the cryptographic components at the core of the DKG protocol itself. This approach simplifies the protocol design and security analysis, while at the same time enabling novel features, such as dynamic participation and support for economic incentives, by utilizing the capabilities of the Ethereum smart contract platform. As such, our protocol provides a versatile building block for a range of designs within and beyond the Ethereum ecosystem.

## ACKNOWLEDGEMENTS

This material is based upon work partially supported by (1) the Christian-Doppler-Laboratory for Security and Quality Improvement in the Production System Lifecycle; The financial support by the Austrian Federal Ministry for Digital and Economic Affairs, the Nation Foundation for Research, Technology and Development and University of Vienna, Faculty of Computer Science, Security & Privacy Group is gratefully acknowledged; (2) SBA Research; the competence center SBA Research (SBA-K1) funded within the framework of COMET

Competence Centers for Excellent Technologies by BMVIT, BMDW, and the federal state of Vienna, managed by the FFG; (3) the FFG Industrial PhD projects 878835 and 878736; (4) the FFG Bridge 1 project 864738 PR4DLT; (5) the FFG ICT of the Future project 874019 dIdentity & dApps; (6) the European Union’s Horizon 2020 research and innovation programme under grant agreement No 826078 (FeatureCloud). We would also like to thank our anonymous reviewers for their valuable feedback and suggestions.

## REFERENCES

- [1] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Secure distributed key generation for discrete-log based cryptosystems,” in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 1999, pp. 295–310.
- [2] A. Shamir, “How to share a secret,” *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.
- [3] G. R. Blakley, “Safeguarding cryptographic keys,” *Proc. of the National Computer Conference*, vol. 48, pp. 313–317, 1979.
- [4] T. P. Pedersen, “A threshold cryptosystem without a trusted party,” in *Workshop on the Theory and Application of Cryptographic Techniques*. Springer, 1991, pp. 522–526.
- [5] D. Boneh and M. Franklin, “Efficient generation of shared rsa keys,” in *Annual International Cryptology Conference*. Springer, 1997, pp. 425–439. [Online]. Available: <https://link.springer.com/content/pdf/10.1007/BFb0052253.pdf>
- [6] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Revisiting the distributed key generation for discrete-log based cryptosystems,” *RSA Security’03*, pp. 89–104, 2003.
- [7] A. Kate and I. Goldberg, “Distributed key generation for the internet,” in *2009 29th IEEE International Conference on Distributed Computing Systems*. IEEE, 2009, pp. 119–128.
- [8] A. Kate, Y. Huang, and I. Goldberg, “Distributed key generation in the wild,” *IACR Cryptology ePrint Archive*, vol. 2012, p. 377, 2012.
- [9] W. Neji, K. Blibech, and N. Ben Rajeb, “Distributed key generation protocol with a new complaint management strategy,” *Security and communication networks*, vol. 9, no. 17, pp. 4585–4595, 2016.
- [10] C. Troncoso, G. Danezis, M. Isaakidis, and H. Halpin, “Systematizing decentralization and privacy: Lessons from 15 years of research and deployments,” in *Proceedings on Privacy Enhancing Technologies*, 2017, pp. 307–329. [Online]. Available: <https://petsymposium.org/2017/papers/issue4/paper87-2017-4-source.pdf>
- [11] P. Feldman, “A Practical Scheme for Non-interactive Verifiable Secret Sharing,” in *Foundations of Computer Science, 1987., 28th Annual Symposium on*. IEEE, 1987, pp. 427–438.
- [12] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Secure applications of pedersen’s distributed key generation protocol,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2003, pp. 373–390.
- [13] T. P. Pedersen, “Non-interactive and information-theoretic secure verifiable secret sharing,” in *Annual International Cryptology Conference*. Springer, 1991, pp. 129–140.
- [14] R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, “Adaptive security for threshold cryptosystems,” in *Annual International Cryptology Conference*. Springer, 1999, pp. 98–116.
- [15] Orbs Network, “DKG for BLS threshold signature scheme on the EVM using solidity,” <https://github.com/orbs-network/dkg-on-evm>, 2018, Accessed: 2021-05-11. [Online]. Available: <https://github.com/orbs-network/dkg-on-evm>
- [16] D. Boneh, B. Lynn, and H. Shacham, “Short signatures from the weil pairing,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2001, pp. 514–532.
- [17] A. R. Choudhuri, M. Green, A. Jain, G. Kaptchuk, and I. Miers, “Fairness in an unfair world: Fair multiparty computation from public bulletin boards,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2017, pp. 719–728. [Online]. Available: <http://delivery.acm.org/10.1145/3140000/3134092/p719-choudhuri.pdf>

- [18] J. Garay, A. Kiayias, and N. Leonardos, “The bitcoin backbone protocol: Analysis and applications,” in *Advances in Cryptology-EUROCRYPT 2015*. Springer, 2015, pp. 281–310. [Online]. Available: <http://courses.cs.washington.edu/courses/cse454/15wi/papers/bitcoin-765.pdf>
- [19] E. Rescorla, “Rfc 2631: Diffie-hellman key agreement method,” RFC, IETF, June, Tech. Rep., 1999.
- [20] T. ElGamal, “A public key cryptosystem and a signature scheme based on discrete logarithms,” *IEEE transactions on information theory*, vol. 31, no. 4, pp. 469–472, 1985.
- [21] D. Chaum and T. P. Pedersen, “Wallet databases with observers,” in *Annual International Cryptology Conference*. Springer, 1992, pp. 89–105.
- [22] J. Camenisch and M. Stadler, “Proof systems for general statements about discrete logarithms,” *Technical report/Dept. of Computer Science, ETH Zürich*, vol. 260, 1997.
- [23] D. Boneh, C. Gentry, B. Lynn, and H. Shacham, “Aggregate and verifiably encrypted signatures from bilinear maps,” in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2003, pp. 416–432.
- [24] D. Boneh, M. Drijvers, and G. Neven, “Compact multi-signatures for smaller blockchains,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2018, pp. 435–464. [Online]. Available: <https://eprint.iacr.org/2018/483.pdf>
- [25] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, “On the security and performance of proof of work blockchains,” in *Proceedings of the 2016 ACM SIGSAC*. ACM, 2016, pp. 3–16.
- [26] Y. Sompolinsky and A. Zohar, “Bitcoin’s security model revisited,” *arXiv preprint arXiv:1605.09193*, 2016. [Online]. Available: <http://arxiv.org/pdf/1605.09193.pdf>
- [27] J.-L. Beuchat, J. E. González-Díaz, S. Mitsunari, E. Okamoto, F. Rodríguez-Henríquez, and T. Teruya, “High-speed software implementation of the optimal ate pairing over barreto-naehrig curves,” in *International Conference on Pairing-Based Cryptography*. Springer, 2010, pp. 21–39. [Online]. Available: <https://eprint.iacr.org/2010/354.pdf>
- [28] D. F. Aranha, K. Karabina, P. Longa, C. H. Gebotys, and J. López, “Faster explicit formulas for computing pairings over ordinary curves,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2011, pp. 48–68. [Online]. Available: [https://link.springer.com/content/pdf/10.1007/978-3-642-20465-4\\_5.pdf](https://link.springer.com/content/pdf/10.1007/978-3-642-20465-4_5.pdf)
- [29] M. Al-Bassam, “Implementation of elliptic curve operations on G2 for alt\_bn128 in Solidity,” <https://github.com/musalbas/solidity-BN256G2>, 2019, Accessed: 2021-05-11. [Online]. Available: <https://github.com/musalbas/solidity-BN256G2>
- [30] C. Reitwiessner, “EIP 196: Precompiled contracts for addition and scalar multiplication on the elliptic curve alt\_bn128,” <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-196.md>, 2017, Accessed: 2021-05-11. [Online]. Available: <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-196.md>
- [31] V. Buterin and C. Reitwiessner, “EIP 197: Precompiled contracts for optimal ate pairing check on the elliptic curve alt\_bn128,” <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-197.md>, 2017, Accessed: 2021-05-11. [Online]. Available: <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-197.md>
- [32] M. Castro and B. Liskov, “Practical byzantine fault tolerance and proactive recovery,” *ACM Transactions on Computer Systems (TOCS)*, vol. 20, no. 4, pp. 398–461, 2002.
- [33] A. S. Cardozo and Z. Williamson, “EIP 1108: Reduce alt\_bn128 precompile gas costs,” <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1108.md>, 2018, Accessed: 2021-05-11. [Online]. Available: <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1108.md>
- [34] “Ethereum Network Status,” <https://ethstats.net/>, 2021, Accessed: 2021-08-22. [Online]. Available: <https://ethstats.net/>
- [35] J. Poon and V. Buterin, “Plasma: Scalable autonomous smart contracts,” Aug 2017, accessed: 2017-08-10. [Online]. Available: <https://plasma.io/plasma.pdf>
- [36] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten, “Arbitrum: Scalable, private smart contracts,” in *27th {USENIX} Security Symposium ({USENIX} Security 18)*, 2018, pp. 1353–1370.
- [37] “Matic Network - Scalable and instant blockchain transactions,” <https://matic.network/>, 2021, Accessed: 2021-05-11. [Online]. Available: <https://matic.network/>
- [38] “Matic Network | Documentation | Matic Gas Station,” <https://docs.matic.network/docs/develop/tools/matic-gas-station/>, 2021, Accessed: 2021-05-11. [Online]. Available: <https://docs.matic.network/docs/develop/tools/matic-gas-station/>
- [39] Vitalik Buterin, “How should I handle blockchain forks in my DApp?” <https://ethereum.stackexchange.com/questions/183/how-should-i-handle-blockchain-forks-in-my-dapp/203>, 2016, Accessed: 2021-05-11. [Online]. Available: <https://ethereum.stackexchange.com/questions/183/how-should-i-handle-blockchain-forks-in-my-dapp/203>
- [40] —, “On Slow and Fast Block Times,” <https://blog.ethereum.org/2015/09/14/on-slow-and-fast-block-times/>, 2015, Accessed: 2021-05-11. [Online]. Available: <https://blog.ethereum.org/2015/09/14/on-slow-and-fast-block-times/>

## APPENDIX A COMPUTATIONAL COSTS

Figure 1, provides the measured gas consumption per executed transaction for different numbers of parties participating in the DKG protocol. We observe that (i) gas costs for contract deployment (2 551 221), for registration (106 385), and (recovered) key share submission (222 510) do not depend on the number of participants, (ii) costs for recovery linearly depend on the number of recovered parties, whereas (iii) the costs for the other operations increase linearly with increasing numbers of participants. Figure 1 reports the measured costs for (ii) and (iii) in the worst case for different numbers of participants ( $n$ ). We use a setup with  $n = 2t + 1$  participants, where  $t$  parties are executing adversarial actions, i.e. they either provide invalid shares, handled by issuing dispute transactions, or they withhold the required values during the key generation phase, leading to a recovery of the missing information.

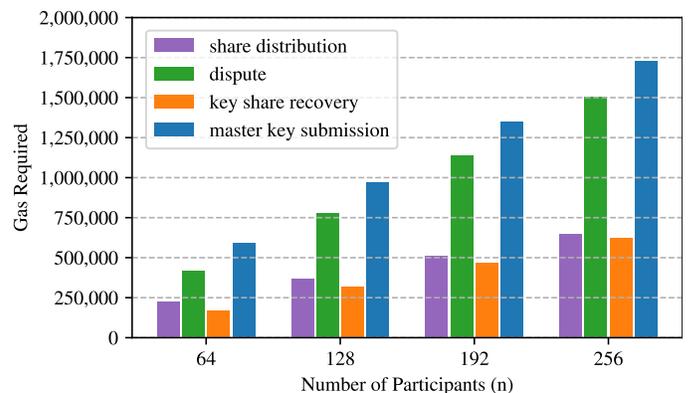


Figure 1. Computational costs, measured in gas per transaction, for the different types of interactions with the smart contract

The most critical operations in terms of gas consumption are the execution of a dispute transaction (potentially executed once per adversarial node), and the submission and verification of the master public key (once at protocol end). In the most demanding scenario with  $n = 256$  participants we evaluated, a dispute consumes approximately 1.5 million gas, whereas the master key submission requires around 1.7 million gas. In both cases, the costs are largely dominated by the internal verification procedures, relying on elliptic curve multiplications.

Compared to the evaluation of the previous version of this paper, the gas costs for these operations have been significantly reduced due to implementation of the EIP-1108 proposal [33] as part of Ethereum’s Istanbul hardfork. With these reductions, we tested our implementation with up to 256 participants and find that our protocol is able to perform all required operations well within the current Ethereum block gas limit of 15 000 000 gas [34].

In case all participants behave according to the protocol, no transactions for dispute, and key share recovery are executed. In the worst case, a dispute transaction has to be executed for each adversarial party in order to prove that the respective party violated the protocol rules. To avoid that correct parties need to cover the costs for the dispute transactions, a recommended mitigation strategy is to require security deposits during registration. The deposit from an adversarial party is then seized when a valid dispute is submitted, and used to refund the disputing party for the expenses incurred by publishing the dispute transaction. In case dispute transactions against the same party are issued concurrently, the fees for all but the first processed transaction are much cheaper, as the contract aborts prematurely. However, also in this case the additional costs may be covered by the adversary’s security deposit. A different mitigation strategy is to reduce the likelihood of concurrent submission of transactions by continuously monitoring for dispute transactions. In this case, dispute transactions are only issued on demand, i.e. in case there was no dispute against the specific party submitted yet, at randomized points in time within the bounds of the dispute phase. For many real world scenarios, in particular when the DKG is run between known entities, we expect the number of disputes to be very low if not zero. In this case, a high number of disputes would likely be addressed at an organization level and not within the protocol itself.

In order to keep costs for the share distribution low, we minimize the amount of data stored within the smart contract. In particular, we do not store the transaction data, i.e.  $n - 1$  encrypted shares and  $t + 1$  commitments to the secret sharing polynomial, in the smart contract. Instead, a cryptographic hash of the above information is stored, whereas triggering a corresponding Ethereum event renders the full data easily accessible to all clients. During the verification of a dispute, this cryptographic hash is recomputed and compared to the stored value to ensure that the disputer’s information is correct.

To further illustrate the costs in practice, tables I and II provide a costs overview for running our protocol on the Ethereum platform by converting the gas consumption into USD. Hereby we group the estimated base costs (covering the registration, share distribution and key share submission steps) each joining party has to cover, and list the costs for a dispute (per adversarial participant), a key share recovery (per failed participant after successful key sharing), as well as one-time costs (per DKG execution) for contract deployment and master key verification separately. Note that the costs are highly depended on the current gas price and Ethereum to USD exchange rate.

To reduce transaction fees, aside from choosing an alternative EVM compatible ledger with lower transaction fees, our protocol may also be adapted for layer 2 scaling solutions such as Plasma [35] and Arbitrum [36]. As a concrete example,

the Matic/Polygon network [37] (a already deployed Plasma variant), reports current gas prices of 1 GWei at the time of writing [38]. This leads to greatly reduced transaction fees compared to a native execution on Ethereum.

## APPENDIX B COMMUNICATION COMPLEXITY

Table III, describes the size, the number of invocations, the total amount of data processed within the Ethereum blockchain, as well as the total amount of data transferred through the network, for all the different transactions executed throughout a protocol execution. The reported values consider the worst case scenario, where the adversary sends invalid shares and fails to provide the required information during the key derivation phase. Overall, the communication complexity of our protocol is  $\mathcal{O}(n^3)$ , that is considering the network traffic generated by broadcasted all transactions. For the smart contract this is equivalent to a communication complexity of  $\mathcal{O}(n^2)$ , as the Ethereum client transparently handles the network communication. This distinction is crucial, as gas costs are only paid for the smart contract execution and are not dependent on the actual network traffic.

## APPENDIX C EXECUTION TIME

In the following, we estimate the total (worst case) execution time required to run our protocol. In practice, this execution time depends on a range of factors, including:

- 1) the number of confirmation required  $\Delta c$ , before a transaction is considered confirmed:  $\approx 12$  blocks [39]
- 2) Ethereum’s (average) time between two subsequent blocks  $\Delta b$ : 13 – 17 seconds [34], [40]
- 3) Ethereum’s block gas limit  $c_{block}$ :  $\approx 15\,000\,000$  gas [34]
- 4) the current load on the Ethereum network
- 5) the gas price participants are willing to pay
- 6) the number of parties executing the DKG protocol

Since all three protocol phases are executed subsequently, the total time required to execute the protocol  $T$  is the sum of the times required to execute each protocol phase. We use  $b_r$ ,  $b_s$ ,  $b_d$  and  $b_k$  to denote the number of blocks required to execute the protocol phases, *registration*, *sharing*, *dispute* and *key derivation* respectively. Consequently, we obtain  $T$  as follows:

$$T = \Delta b \cdot (b_r + b_s + b_d + b_k) \quad (12)$$

To compute the number of blocks required for each protocol phase, in particular  $b_r$ ,  $b_s$  and  $b_d$ , we consider a consensus stabilization period at the end of each phase ( $\Delta c$  blocks), a safe upper bound for the number of blocks to wait until a transaction is included in the Ethereum blockchain ( $\Delta i$ ), the capacity required to fit all transactions of the specific phase ( $c_r$ ,  $c_s$ ,  $c_d$ ), as well as the maximum capacity  $c_{max}$  the DKG protocol should use on the Ethereum platform during execution (e.g. 10% of the block gas limit).

$$b_x = \Delta c + \Delta i + \left\lceil \frac{c_x}{c_{max}} \right\rceil \quad x \in r, s, d \quad (13)$$

number of nodes	8	16	32	64	128	192	256
base	0.14 \$	0.15 \$	0.16 \$	0.18 \$	0.23 \$	0.28 \$	0.32 \$
dispute	0.03 \$	0.05 \$	0.08 \$	0.14 \$	0.26 \$	0.38 \$	0.50 \$
key share recovery	0.01 \$	0.01 \$	0.01 \$	0.01 \$	0.01 \$	0.01 \$	0.01 \$
master key verification	0.09 \$	0.10 \$	0.13 \$	0.20 \$	0.32 \$	0.45 \$	0.57 \$
deployment	0.84 \$	0.84 \$	0.84 \$	0.84 \$	0.84 \$	0.84 \$	0.84 \$

Table I. ESTIMATED TRANSACTION FEES AT THE TIME OF INITIAL EVALUATION (2020-04-12; GAS PRICE: 2 GWEI, EXCHANGE RATE: 165 \$ / ETH)

number of nodes	8	16	32	64	128	192	256
base	68.99 \$	71.76 \$	77.29 \$	88.40 \$	110.68 \$	133.06 \$	155.55 \$
dispute	16.77 \$	23.89 \$	38.10 \$	66.60 \$	123.79 \$	181.21 \$	238.95 \$
key share recovery	5.16 \$	5.16 \$	5.16 \$	5.16 \$	5.16 \$	5.16 \$	5.16 \$
master key verification	42.14 \$	49.60 \$	64.51 \$	94.37 \$	154.22 \$	214.23 \$	274.40 \$
deployment	405.01 \$	405.01 \$	405.01 \$	405.01 \$	405.01 \$	405.01 \$	405.01 \$

Table II. ESTIMATED TRANSACTION FEES AT THE TIME OF WRITING (2021-08-22; GAS PRICE: 50 GWEI; EXCHANGE RATE: 3 175 \$ / ETH)

	transaction size	number of invocations	communication complexity (smart contract)	communication complexity (broadcast)
register	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$
sharing	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$	$\mathcal{O}(n^3)$
dispute	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$	$\mathcal{O}(n^3)$
key share submission	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$
key share recovery	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$	$\mathcal{O}(n^3)$
recovered key share submission	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$
master key submission	$\mathcal{O}(n)$	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n^2)$

Table III. COMMUNICATION COMPLEXITY FOR THE DIFFERENT TYPES OF INTERACTIONS WITH THE SMART CONTRACT

number of participants	64	128	192	256
estimated lower bound	61 min	73 min	91 min	115 min
estimated upper bound	85 min	100 min	123 min	153 min

Table IV. PROTOCOL EXECUTION TIMES FOR DIFFERENT NUMBERS OF PARTICIPANTS

Here, the capacities  $c_x$  are derived from the required gas for the specific transaction type, as given in Appendix A and the number of transactions executed. Similarly, the values  $c_{k1}$ ,  $c_{k2}$  and  $c_{k3}$  used below represent the capacities for the key share submission, key share recovery, and recovered key share submission transactions, respectively. As the key derivation phase requires multiple steps,  $b_k$  is computed by considering:

- the number of blocks required for the submission of key shares:  $\Delta c + \Delta i + \left\lceil \frac{c_{k1}}{c_{max}} \right\rceil$
- the number of blocks required for a potential key share recovery:  $\Delta c \Delta i \left\lceil \frac{c_{k2}}{c_{max}} \right\rceil$
- the number of blocks required for submission of the recovered key shares:  $\Delta c \Delta i \left\lceil \frac{c_{k2}}{c_{max}} \right\rceil$
- as well as the number of blocks required for publishing the resulting master public key:  $\Delta c \Delta i$ .

recovered) to get and lower and upper bound for  $b_k$ :

$$b_{k,min} = 2\Delta c + 2\Delta i + \left\lceil \frac{c_{k1}}{c_{max}} \right\rceil \quad (14)$$

$$b_{k,max} = 4\Delta c + 4\Delta i + \left\lceil \frac{c_{k1}}{c_{max}} \right\rceil + \left\lceil \frac{c_{k2}}{c_{max}} \right\rceil + \left\lceil \frac{c_{k3}}{c_{max}} \right\rceil \quad (15)$$

If we consider a worst case scenario with  $n = 2t + 1$  participants, and select conservative values for the parameters above, i.e., we wait for  $\Delta c = 20$  confirmations ( $\approx 4.4$  minutes) before considered a transaction confirmed, assume the latency for transaction inclusion in a block is  $\Delta i = 30$  blocks ( $\approx 6.6$  minutes) and target a network load of 10% of Ethereum capacity ( $c_{max} = 15\,000\,000$ ), and use Ethereum current block interval of  $\Delta b = 13$  seconds [34] we obtain table IV, summarizing the estimated execution times for different numbers of nodes.

In the following, we distinguish between the optimal case (no recovery) and the worst case (49% of all nodes need to be