Synchronous Distributed Key Generation without Broadcasts

Nibesh Shrestha[a,1], Adithya Bhat[2], Aniket Kate[1,3] and Kartik Nayak[4]

[1] Supra Research, USA
[2] Visa Research, USA
[3] Purdue University, USA
[4] Duke University, USA

Abstract. Distributed key generation (DKG) is a key building block in developing many efficient threshold cryptosystems. This work initiates the study of communication complexity and round complexity of DKG protocols over a point-to-point (bounded) synchronous network. Our key result is the first synchronous DKG protocol for discrete log-based cryptosystems with $O(\kappa n^3)$ communication complexity (κ denotes a security parameter) that tolerates any $t < n/2$ Byzantine faults among n parties. We present two variants of the protocol: (i) a protocol with worst-case $O(\kappa n^3)$ communication and $O(t)$ rounds, and (ii) a protocol with expected $O(\kappa n^3)$ communication and expected constant rounds. In the process of achieving our results, we design (1) a novel weak gradecast protocol with a communication complexity of $O(\kappa n^2)$ for linear-sized inputs and constant rounds, (2) a primitive called "recoverable-set-of-shares" for ensuring recovery of shared secrets, (3) an oblivious leader election protocol with $O(\kappa n^3)$ communication and constant rounds, and (4) a multi-valued validated Byzantine agreement (MVBA) protocol with $O(\kappa n^3)$ communication complexity for linear-sized inputs and expected constant rounds. Each of these primitives is of independent interest.

Keywords: Distributed Key Generation · Synchrony · Multi-valued Validated Byzantine Agreement

1 Introduction

The problem of distributed key generation (DKG) is to set up a common public key and its corresponding secret keys among a set of participating parties without a trusted entity. DKG protocols are used to reduce the number of trust assumptions placed in cryptographic protocols such as threshold signatures [Bol03, Sho00] and threshold encryption schemes [DF90]. These threshold cryptosystems can themselves be used to implement random beacons [Dra, CKS00], reduce the complexity of consensus protocols [YMR+19, SARN20], in multiparty computation protocols [HNP05, HMQ04], or to outsource management of secrets to multiple, semi-trusted authorities [DS02, Lab21].

Given its widespread applications and their recent adoption in practice (e.g., [Dra]), we need efficient solutions for DKG. An ideal solution for DKG would have low communication complexity, low latency, optimal resilience, and provide uniform randomness of generated keys such that the generated keys can be useful in a wider class of cryptosystems while being secure. This work focuses on the synchronous network setting which has the advantage of tolerating up to a minority corruption (as against one-third corruption in other settings).

Table 1: **Comparison of related works on Distributed Key Generation**

	Net.Res.	Comm.	Round	Sim.	Dlog	Comp.	Setup	Crypto Assumption
Pedersen [Ped91]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✗	✓	$O(n^2)$	PKI	DL
GJKR [GJKR07]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✓	✓	$O(n^2)$	PKI	DL
CGJ+99[CGJ$^+$99]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✓	✓	$O(n^2)$	PKI	DL
Neji et al. [NBBR16]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✓	✓	$O(n^2)$	PKI	RO+CDH
ETHDKG [SJSW19]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✗	✓	$O(n^2)$	PKI	RO+CDH
GJM+21[GJM$^+$21]	sync.$\log n$	$O(n.\mathsf{B}(\kappa))+$ $O(\log n.\mathsf{B}(\kappa n))$	$O(R_n+ \log n)$	✗	✗	$O(n\log^2 n)$	PKI	RO+SXDH+$^{(1)}$
NIDKG [Gro21]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✓	✓	$O(n^2)$	PKI	RO+DDH+$^{(2)}$
CDSV[CDSV23]	sync. 1/2	$O(n.\mathsf{B}(\kappa n))$	$O(\mathcal{R}_n)$	✓	✓	$O(n^2)$	PKI	RO+DDH
Hybrid-DKG [KHG12]	psync. 1/3	$O(\kappa n^4)$	$O(t)$	✓	✓	$O(n^2)$	PKI	RO+DL
KKMS [KMS20]	async. 1/3	$O(\kappa n^4)$	$O(t)$	✗	✓	$O(n^3)$	PKI	RO+DDH
AJM+21[AJM$^+$21]	async. 1/3	$\tilde{O}(\kappa n^3)$	$E(O(1))$	✗	✗	$O(n^2)$	PKI	RO+SXDH
Das et al. [DYX$^+$22]	async. 1/3	$O(\kappa n^3)$	$E(O(\log n))$	✓	✓	$O(n^3)$	PKI	RO+DCR+DDH
Das et al. [DXKKR23]	async. 1/3	$O(\kappa n^3)$	$E(O(\log n))$	✓	✓	$O(n^2)$	PKI	RO+CDH
AJM+23[AJM$^+$23]	async. 1/3	$O(\kappa n^3)$	$E(O(1))$	✓	✓	$O(n^2)$	PKI+AGM	RO+OMDL+$^{(3)}$
Our work (§ 8.1)	sync. 1/2	$O(\kappa n^3)$	$E(O(1))$	✓	✓	$O(n^2)$	PKI+PoT	RO+CDH+q-SDH
Our work (§ 8.2)	sync. 1/2	$O(\kappa n^3)$	$O(t)$	✓	✓	$O(n^2)$	PKI+PoT	RO+q-SDH

We use $\mathsf{B}(\ell)$ to denote the communication cost of Byzantine consensus primitive for ℓ bit message and $\mathcal{R}_n$ to denote the round complexity of n parallel invocation of the Byzantine consensus primitive. κ is the security parameter. **Net.** refers to the network model. **Res.** refers to the number of Byzantine faults tolerated in the system. **Comm.** refers to the communication complexity. **Sim.** means the protocol maintains secrecy which can be proven via a simulator. **Dlog.** refers to the generation of keys for discrete log based cryptosystems. **Comp.** refers to the computational complexity per party. **PoT** refers to the power of tau setup required for bilinear accumulators. This setup can be removed by making use of Merkle trees at the cost of $\log n$ multiplicative communication overhead. **E(.)** implies "in expectation". (1) GJM+21 [GJM$^+$21] also assumes CBDH. (2) NIDKG [Gro21] assumes RO, rleaf-IND-CCA, DDH, Erasures, and one-more DH. (3) AJM+23 [AJM$^+$23] also assumes q-SDH assumption.

While a myriad of DKG protocols [Ped91, GJKR07, CGJ$^+$99, NBBR16, GJM$^+$21] have been proposed in this setting, existing solutions fall short in one way or the other. For example, Pedersen's DKG [Ped91] produces non-uniform keys in the presence of the adversary, the DKG protocol due to Gennaro et al. [GJKR07] has high latency as it requires additional secret sharing using Feldman's VSS [Fel87], and the protocol due to Gurkhan et al. [GJM$^+$21] does not generate keys for discrete log-based cryptosystems. A key concern with all the DKG protocols published in the synchronous model is that they assume the existence of a *broadcast channel* (that provides a consensus abstraction) and invoke $\Omega(n)$ broadcasts across two or more rounds [BKP11], where n is the number of parties. Such a broadcast channel does not exist in typical settings and must be realized from point-to-point communication channels. Can we achieve this efficiently?

Communication and round complexity of existing synchronous DKG protocols. A natural approach to instantiate the broadcast channels is via Byzantine consensus primitives such as Byzantine broadcast (BB) [DS83, TLP22] or Byzantine agreement (BA) [KK06]. To the best of our knowledge, all optimally resilient deterministic Byzantine consensus protocols incur $O(\kappa n^3)$ communication (κ is a security parameter) without threshold signatures and heavier cryptographic primitives[1] and require $O(t)$ rounds [DS83]. For randomized consensus protocols, the best known protocol with optimal resilience and without threshold signature setup is Katz and Koo [KK06] which incurs $O(\kappa n^4)$ communication and expected constant rounds.

Given that existing DKG protocols typically require $\Omega(n)$ broadcasts across two or more rounds, utilizing state-of-the-art deterministic Byzantine consensus protocols to instantiate the broadcast channels would result in a communication complexity of $O(\kappa n^4)$ and $O(t)$ rounds. In the context of randomized Byzantine consensus protocols, n

---

[1]Heavier cryptographic primitives such as zk-SNAKRS [Gro16] can be used to obtain threshold signatures of size $O(\kappa)$ without a prior DKG. When combined with BA protocol of Momose and Ren [MR21], this results in optimally resilient BA protocol with $O(\kappa n^2)$ communication and $O(t)$ rounds.

parallel BA protocols can be composed to achieve protocols with an expected constant number of rounds [BOEY03, KK06]. However, this approach leads to a significantly high communication complexity of $\Omega(\kappa n^4)$. This leaves us with the following open question: *Can we design a synchronous DKG protocol supporting a wide class of cryptosystems with $o(\kappa n^4)$ communication complexity, good latency, and tolerating a minority corruption?*

We answer this question positively by showing two DKG protocols for discrete log-based cryptosystems each with $O(\kappa n^3)$ communication complexity. The first DKG protocol is a randomized protocol which has expected $O(\kappa n^3)$ communication and expected constant rounds. Our second DKG protocol is deterministic and has the worst-cast $O(\kappa n^3)$ communication and $O(t)$ rounds.

1.1 Key Technical Ideas and Contributions

Our DKG protocols avoid the broadcast channel assumption and use only a single Byzantine consensus invocation (instead of n instances) to achieve $O(\kappa n^3)$ communication. While DKG protocols [KMS20, AJM+21, DYX+22, DXKKR23] without broadcast channel assumption have been explored in the asynchronous model, they either incur high communication [KMS20] or do not generate keys for discrete log-based cryptosystems [AJM+21] or use stronger cryptographic assumptions [DYX+22, AJM+23] and have longer latency [DYX+22, DXKKR23]. More importantly, protocols designed for asynchronous or partially-synchronous settings can only tolerate up to $t < n/3$ Byzantine failures, which is sub-optimal for many DKG applications such as random beacons [Dra]. In the synchronous model, we provide the first solutions to DKG without a broadcast channel with all the desirable properties with $O(\kappa n^3)$ communication.

Typical approach using broadcast channels. A common approach in existing works [GJKR07, Ped91] is to employ n parallel verifiable secret sharings (VSS) [Fel87, Ped92]. A VSS scheme allows a dealer to distribute shares of a secret among a group of parties such that (i) each party can validate the secret shares they receive to confirm their consistency with the shared secret and (ii) a specific number of parties can later reconstruct the secret if they combine their secret shares. The protocols typically assume the existence of a broadcast channel, which ensures that the dealer shares the same secret with all honest parties and that all honest parties receive correct secret shares. In the context of DKG, due to the reliance on the broadcast channel, all honest parties reach a consensus on a common set of qualified parties, denoted as QUAL, who have correctly performed the secret sharing. At this stage of the protocol, the assurance obtained is that all honest parties have acquired correct secret shares, which correspond to the secrets shared by the qualified parties in QUAL. Following this, the final public key and secret keys are computed from the secret shares of all parties in QUAL.

Our Approach

The use of broadcast channels by each of the n different parties is expensive from the standpoint of communication complexity and round complexity. Thus, in our protocols, we replace the use of n broadcast/consensus instances with a combination of weaker primitives such as gradecast [FM88, KK06] and use a single consensus instance. Unlike broadcast, where the honest parties must reach a unanimous decision, gradecast allows the honest parties to disagree by "a small amount". Specifically, parties output a grade alongside their output value; this grade can be perceived as the "confidence" of the party in its output. In the case of an honest sender, all honest parties will output a common value with the highest grade. However, if the sender is Byzantine, honest parties might output different values with different grades.

Following this approach, parties first perform secret sharing by using gradecast to identify a set of at least $n - t$ parties who correctly share their secrets, where t is the

fault tolerance. The challenge with this approach is that different honest parties may have different views regarding the acceptance of shared secrets. Consequently, different honest parties may obtain different sets of at least $n - t$ parties (say $\mathsf{AcceptList}_i$ for party P_i) who they accept to have performed secret sharing correctly. However, for DKG, it is required that all honest parties compute the final public key and secret keys from a common set of parties. Therefore, consensus on a common set of parties is crucial.

Thus, parties use a Byzantine consensus primitive to agree on one common set where the input is their individual $\mathsf{AcceptList}$ and the output is the input of one of the parties. Given that each party may input a different $\mathsf{AcceptList}$, the output of the Byzantine consensus primitive could correspond to the input of any party, including a Byzantine party that may input an arbitrary set. To safeguard against such malicious behavior, we ensure that parties produce their $\mathsf{AcceptList}$ with verifiable proof confirming that all parties in their set have correctly shared their secrets and these secrets are thus recoverable. This assurance is provided through a novel protocol, termed "Recoverable-set-of-shares".

It is important to note that the size of the set and its associated proof is linear in the number of parties. Given this, we need a consensus primitive that takes long messages as inputs, and outputs one of the "valid" input values. Such a primitive is called *multi-valued validated Byzantine agreement* (MVBA) [CKPS01] in the literature. Typically, deterministic consensus protocols are subject to $\Omega(t)$ rounds [DS83]. To circumvent this limitation, randomization is often employed. In the context of MVBA protocols, we employ a primitive known as oblivious leader election (OLE) to elect leaders uniformly at random and output a common honest leader with some constant probability. The common honest leader ensures that all honest parties agree on its proposed value and subsequently terminate.

Our protocols rely on a combination of primitives, including gradecast, recoverable-set-of-shares, MVBA, and oblivious leader election. However, using the current state-of-the-art for these primitives does not suffice from the standpoint of improving the communication complexity of our DKG protocol. Thus, our work contributes towards improving state-of-the-art in each of these primitives to design the resulting DKG protocol. In the following subsection, we provide a background on the state-of-the-art for these primitives. In Section 1.1.2, we detail our key contributions and results towards improving these primitives.

1.1.1   Background

1. Gradecast. As previously mentioned, our protocols utilize gradecast instead of the traditional broadcast channel. To the best of our knowledge, existing gradecast protocols with a resilience of at least $t < n/2$ incur a communication complexity of at least $O(n^3)$. For instance, the gradecast protocol by Katz and Koo [KK06] requires $O(\kappa n^3)$ communication for a single bit input in the absence of threshold signatures. In a similar vein, the gradecast protocol by Garay et al. [GKKO07] has a communication complexity of $O(g^*(\ell + \kappa)n^2)$ for an input of size ℓ, where g^* is the maximum supported grade. When $\ell = \Theta(n)$, their protocol also incurs $O(n^3)$ communication. This highlights the need for a more communication-efficient gradecast protocol, particularly when dealing with large inputs.
2. Multi-valued validated Byzantine agreement. MVBA was first formulated by Cachin et al. [CKPS01] to allow honest parties to decide on any externally valid values. Recent works [AMS19, LLTW20] have given communication efficient protocols for MVBA in the asynchronous model tolerating $t < n/3$ Byzantine faults. For long messages of size ℓ, the protocol due to Abraham et al. [AMS19] incurs $O((\ell + \kappa)n^2)$ communication and the protocol due to Lu et al. [LLTW20] incurs $O(n\ell + \kappa n^2)$. Both of these works assume a threshold setup. Without threshold setup assumptions, the communication blows up by a factor of n in all of the above protocols.

To the best of our knowledge, no MVBA protocols have been formulated in the

Table 2: **Comparison of related works on MVBA with ℓ-bit input**

		Network	Res.	Communication	Round	Setup
Cachin et al.	[CKPS01]	async.	1/3	$O(n^2\ell + \kappa n^2 + n^3)$	$E(O(1))$	Threshold setup
VABA	[AMS19]	async.	1/3	$O(n^2\ell + \kappa n^2)$	$E(O(1))$	Threshold setup
DUMBO-MVBA	[LLTW20]	async.	1/3	$O(n\ell + \kappa n^2)$	$E(O(1))$	Threshold setup
Our work		sync.	1/2	$O(n^2\ell + \kappa n^3)$	$E(O(1))$	PKI

Res. refers to resilience. **E(.)** implies "in expectation".

synchronous model tolerating $t < n/2$ faults. Recently, Nayak et al. [NRS$^+$20] provides an efficient BA protocol for long messages. However, since it is a BA protocol, they output a value only when all honest parties start with the same large input. This underscores the necessity for a synchronous MVBA protocol for long messages.

3. Oblivious leader election. In the absence of an existing threshold (DKG) setup, the OLE protocol was designed via n^2 parallel invocations of weaker VSS primitives such as graded VSS [FM88] or moderated VSS [KK06] which trivially incurs $\Omega(n^4)$ communication. While communication efficient OLE protocols tolerating $t < n/3$ Byzantine faults have been designed using Aggregatable PVSS [GJM$^+$21] for the asynchronous model, they required stronger cryptographic assumptions such as SXDH to achieve $O(\kappa n^3)$ communication [AJM$^+$21, GLL$^+$22]. Designing a communication-efficient OLE protocol with fewer cryptographic assumption is of an independent interest.

1.1.2   Our Results

We address the above challenges and make the following contributions in the paper.

1. Optimal weak gradecast with $O(n\ell + \kappa n^2)$ communication. We provide a communication optimal weak gradecast protocol satisfying the gradecast definition of Katz and Koo [KK06][2] in Section 4. Our weak gradecast protocol incurs $O(n\ell + \kappa n^2)$ communication for ℓ bit input and does not require use of threshold signatures. The reduction in communication is obtained via the use of efficient erasure coding schemes [RS60], cryptographic accumulators [Ngu05] and efficient detection of equivocation from the Byzantine sender. With q-Strong Diffe Hellman (q-SDH) [BB08] setup assumption, we show the following result:

Theorem 1. *Assuming a public-key infrastructure and a universal structured reference string under q-SDH assumption, there exists a gradecast protocol for an input of size ℓ bits with $O(n\ell + \kappa n^2)$ communication and 4 rounds tolerating $t < n/2$ Byzantine faults.*

2. Recoverable-set-of-shares using weak gradecast. We use the gradecast primitive to perform communication efficient secret sharing. In this protocol, each party P_i identifies a set of at least $n-t$ parties who P_i accepts to have shared secrets correctly. In this protocol, we ensure that for any set output by any party (including Byzantine parties), there is verifiable proof vouching that all parties in the set have correctly shared their secrets and these secrets are thus recoverable. In this regard, we call this protocol "Recoverable-set-of-shares". Using our communication optimal gradecast, our recoverable-set-of-shares protocol can be achieved in $O(\kappa n^3)$ communication and constant rounds.

Theorem 2. *Assuming a public-key infrastructure, random oracle and a universal structured reference string under q-SDH assumption, there exists an recoverable-set-of-shares protocol tolerating $t < n/2$ Byzantine faults with $O(\kappa n^3)$ communication and 10 rounds.*

3. Oblivious leader election with $O(\kappa n^3)$ communication. We design a communication efficient oblivious leader election (OLE) protocol with $O(\kappa n^3)$ communication

---

[2]This definition is slightly weaker than the one presented by Feldman and Micali [FM88].
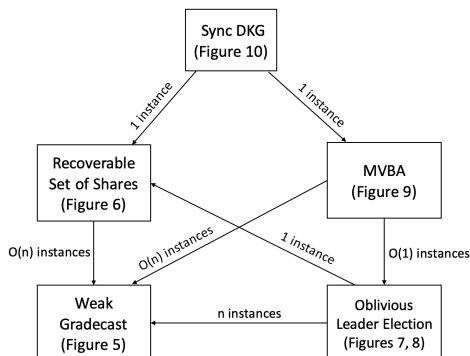
Figure 1: Overview of sub-protocols and their dependencies

and constant rounds in Section 6. Our OLE protocol uses only n weaker VSS instances and a non-interactive threshold signature scheme [CKS00] to generate randomness. To the best of our knowledge, we design the first OLE protocol that requires only n secret sharings. We note that our OLE protocol does not require a prior threshold (DKG) setup phase despite making use of threshold signatures. The security of our OLE protocol is based on the computational Diffie-Hellman (CDH) problem in the random oracle model. In particular, we show the following:

Theorem 3. *Assuming a public-key infrastructure, a universal structured reference string under q-SDH assumption, random oracle, and CDH, there exists an oblivious leader election protocol with $O(\kappa n^3)$ communication and $O(1)$ rounds tolerating $t < n/2$ Byzantine faults.*

4. Efficient multi-valued validated Byzantine agreement with $O(n^2\ell + \kappa n^3)$ communication. We construct the first MVBA protocol tolerating $t < n/2$ Byzantine faults in the synchronous setting without threshold setup in Section 7. The MVBA protocol relies on our OLE protocol to obtain expected constant number of rounds. Our MVBA protocol incurs expected $O(n^2\ell + \kappa n^3)$ communication for inputs of size ℓ bit and expected 36 rounds. Following our OLE protocol, the security of our MVBA protocol relies on the CDH problem in the random oracle model. Specifically, we show the following result:

Theorem 4. *Assuming a public-key infrastructure, random oracle, CDH, and a universal structured reference string under q-SDH assumption, there exists a multi-valued validated Byzantine agreement protocol for an input of size ℓ with expected $O(n^2\ell + \kappa n^3)$ communication and expected 36 rounds tolerating $t < n/2$ Byzantine faults.*

Secure DKG with two broadcast rounds. To provide the necessary background, we first introduce a secure DKG protocol that assumes the use of a broadcast channel in Section 3. This DKG protocol requires only 2 broadcast rounds, whereas prior work [GJKR07] necessitates three or more broadcast rounds. This DKG protocol is derived by selectively choosing components from existing literature [BSL+21, KHG12]. Subsequently, we demonstrate how to replace broadcast channels to achieve our communication and round-efficient DKG protocols.

Efficient distributed key generation. Using our recoverable-set-of-shares protocol where parties output different sets of size at least $n - t$ parties and our MVBA protocol, honest parties can agree on a common set from which the final public key and secret keys are computed. In particular, we obtain a DKG protocol with expected $O(\kappa n^3)$ communication and expected 47 rounds.

Theorem 5. *Assuming public-key infrastructure, random oracle, CDH and a universal structured reference string under q-SDH assumption, there exists a protocol that solves*

secure synchronous distributed key generation tolerating $t < n/2$ Byzantine faults with expected $O(\kappa n^3)$ communication and expected 47 rounds.

Although the DKG protocol terminates in constant expected time, it can take linear time in the worst case. In this case, the protocol incurs $O(\kappa n^4)$ communication. As an alternative, we provide a protocol that incurs $O(\kappa n^3)$ communication in the worst-case. RandPiper [BSL+21] provides a Byzantine fault tolerant state machine replication (BFT SMR) protocol with $O(\kappa n^2)$ communication per epoch even for $O(n)$-sized input. Here, an epoch is a period that incurs 7 rounds. In this protocol, we execute the BFT SMR protocol for $t + 1$ epochs with each epoch coordinated by a distinct leader. The leader proposes his set AcceptList along with the proof. Honest parties output the first committed set to compute the final public key and secret keys. In particular, we obtain the following result:

Theorem 6. *Assuming a public-key infrastructure, discrete-log assumption, random oracle and a universal structured reference string under q-SDH assumption there exists a protocol that solves secure synchronous distributed key generation tolerating $t < n/2$ Byzantine faults with $O(\kappa n^3)$ communication and $11 + 7(t + 1)$ rounds.*

Limitations. In this work, we assume that the adversary is static, similar to several DKGs [KG09, GJKR07, Ped91, NBBR16, SJSW19, GJM+21] in the literature. Canetti et al. [CGJ+99] show how to build adaptively secure DKG protocols and several of our techniques could be applicable in realizing their protocol in the point-to-point network setting. Bacho et al. [BL22] gave a relaxed definition of DKG and show that prior DKG protocols such as Gennaro et al [GJKR07] are adaptively-secure under this relaxed definition. It could be interesting to see if our protocols are adaptively-secure under their relaxed definition. In addition, our protocols make the q-SDH assumption. This assumption is only used for bilinear accumulators which could be replaced with Merkle tree accumulators resulting in a $\log n$ multiplicative overhead in the communication complexity.

2   Model and Preliminaries

We consider a system consisting of n parties $(P_1, \ldots, P_n)$ with pair-wise reliable, authenticated point-to-point channels, where up to $t < n/2$ parties can be Byzantine faulty. The model of corruption is static i.e., the adversary picks the corrupted parties before the start of protocol execution. The Byzantine parties may behave arbitrarily. A non-faulty party is said to be *honest* and executes the protocol as specified. We assume a synchronous communication model. Thus, if an honest party sends a message at the beginning of some round, the recipient receives the message by the end of that round.

Setup. Let p be a prime number that is $\mathsf{poly}(\kappa)$ bits long, and $\mathbb{G}$ be a group of order p such that it is computationally infeasible except with negligible probability in κ to compute discrete log. Let $\mathbb{Z}_p$ denote its scalar field. Moreover, let g and h denote the generators of $\mathbb{G}$ where $a \in \mathbb{Z}_p$ such that $g^a = h$ is not known to any t subset of the parties.

We make the standard computational assumption on the infeasibility to compute discrete logarithms called the discrete-log assumption [GJKR07]. In particular, we assume that the adversary is unable to compute discrete logarithms modulo large (based on the security parameter κ) primes.

We make use of digital signatures and PKI to prevent spoofing and replays and to validate messages. Message x sent by a party P_i is digitally signed by P_i's private key and is denoted by $\langle x \rangle_i$. We denote $H(x)$ to represent invocation of the random oracle H on input x. In addition, we use a hash function $H' : \mathbb{G} \to \{0, 1\}^\kappa$ in our leader election protocol.

Equivocation. Two or more messages of the same *type* but with different payload sent by a party is considered an equivocation. In order to facilitate efficient equivocation checks,

the sender sends the payload along with signed hash of the payload. When an equivocation is detected, broadcasting the signed hash suffices to prove equivocation by the sender.

2.1 Definitions

Distributed key generation. A DKG protocol for n parties $(P_1, \ldots, P_n)$ generates private outputs $(x_1, \ldots, x_n)$ called the *shares* and a public output y.

Definition 1 (Secure DKG for Dlog based cryptosystems [GJKR07])**.** A dlog based DKG protocol generates a uniformly random secret x, distributed among n parties through shares $(x_1, \ldots, x_n)$ where x_i is a share output to party P_i. The protocol is considered t-secure if, in the presence of an adversary that corrupts up to t parties, the following requirements for correctness and secrecy are satisfied.

> **Correctness.**
>
> > **C1.** All subsets of $t + 1$ shares provided by honest parties define the same unique secret key $x \in \mathbb{Z}_p$.
> >
> > **C2.** All honest parties have same value of public key $y = g^x \in \mathbb{G}$, where $x \in \mathbb{Z}_p$ is secret guaranteed by (C1).
> >
> > **C3.** x is uniformly distributed in $\mathbb{Z}_p$ (and hence y is uniformly distributed in $\mathbb{G}$).
>
> **Secrecy.** No information on x can be learned by the adversary except for what is implied by the value $y = g^x$.

More formally, the secrecy condition is expressed in terms of simulatability: for every (probabilistic polynomial-time) adversary $\mathcal{A}$ that corrupts up to t parties, there exists a (probabilistic polynomial-time) simulator $\mathcal{S}$, such that on input an element $y \in \mathbb{G}$, produces an output distribution which is polynomially indistinguishable from $\mathcal{A}$'s view of a run of the DKG protocol that ends with y as its public key output.

We acknowledge that our DKG definition is derived from Gennaro et al. [GJKR07], which is suitable for discrete logarithm access structures. Recently, Komlo et al. [KGS23] provided a generic definition for DKG that does not depend on any access structure. Satisfying these definitions could be an interesting avenue for future work.

Weak Gradecast. Gradecast is a relaxed version of broadcast introduced by Feldman and Micali [FM88] wherein parties output a value along with a grade; the grade serves as the "confidence" the party has in its output. Our weak gradecast protocol satisfies the weak gradecast definition of Katz and Koo [KK06].

Definition 2 (Weak Gradecast [KK06])**.** A protocol with a designated sender P_i holding an initial input v is a weak gradecast protocol tolerating $t < n/2$ Byzantine parties if the following conditions hold:

1. Each honest party P_j outputs a value v_j with a grade $g_j \in \{0, 1, 2\}$.

2. If the sender is honest, the output of every honest party P_i satisfies $v_i = v$ and $g_i = 2$.

3. If an honest party P_i outputs a value v_i with a grade of 2, then the output of every honest party P_j satisfies $v_j = v_i$ with a grade $g_j \geq 1$.

Observe that this definition allows honest parties to output different values with a grade of 1 when no honest party outputs a value with a grade of 2. This is in contrast to the definition of Feldman and Micali [FM88] that requires honest parties with a grade of 1 to output the same value. This weaker definition suffices for our purpose.

Oblivious leader election. An oblivious leader election protocol elects a common honest leader with some constant probability.

Definition 3 (Oblivious Leader Election [KK06]). *A protocol for parties $P_1, \ldots, P_n$ is an oblivious leader election protocol with fairness α tolerating t Byzantine failures if each honest party P_i outputs a value $v_i \in [n]$ and the following conditions holds with probability at least α:*

There exists a value $j \in [n]$ such that (i) each honest party P_i outputs $v_i = j$, and (ii) party P_j is honest.

Multi-valued validated Byzantine agreement. In an MVBA protocol, there is an external validity function ex-validation that every party has access to. Every honest party start with some externally valid input v_i, and on termination must output a value. An MVBA protocol has the following properties:

Definition 4 (Multi-valued Validated Byzantine Agreement [AMS19, LLTW20]). *A protocol solves multi-valued validated Byzantine agreement if it satisfies following properties except with negligible probability in the security parameter κ:*

- **Validity.** *If an honest party decides a value v, then ex-validation(v) = true.*

- **Agreement.** *No two honest parties decide on different values.*

- **Termination.** *If all honest parties start with externally valid values, all honest parties eventually decide.*

In addition to these properties, the MVBA definition of [AMS19, LLTW20] also include an additional property called "quality" where in the MVBA should decide a value proposed by an honest party with some constant probability. Our MVBA construction in Section 7 also satisfies the quality property with probability at least $\frac{1}{2}$. However, this property is not required for constructing our DKG protocol. Hence, we do not discuss the quality property.

2.2 Primitives

In this section, we present several primitives used in our protocols.

Linear erasure and error correcting codes. We use standard $(t + 1, n)$ Reed-Solomon (RS) codes [RS60]. This code encodes $t + 1$ data symbols into code words of n symbols using ENC function and can decode the $t + 1$ elements of code words to recover the original data using DEC function as explained below.

- **ENC.** Given inputs $m_1, \ldots, m_{t+1}$, an encoding function ENC computes $(s_1, \ldots, s_n) =$ ENC$(m_1, \ldots, m_{t+1})$, where $(s_1, \ldots, s_n)$ are code words of length n. A combination of any $t + 1$ elements of n code words uniquely determines the input message and the remaining of the code word.

- **DEC.** The function DEC computes $(m_1, \ldots, m_{t+1}) =$ DEC$(s_1, ..., s_n)$, and is capable of tolerating up to c errors and d erasures in code words $(s_1, \ldots, s_n)$, if and only if $t \geq 2c + d$.

Cryptographic accumulators. A cryptographic accumulator scheme constructs an accumulation value for a set of values using Eval function and produces a witness for each value in the set using CreateWit function. Given the accumulation value and a witness, any party can verify if a value is indeed in the set using Verify function. Formally, given a parameter k, and a set D of n values $d_1, \ldots, d_n$, an accumulator has the following components:

- Gen$(1^k, n)$: This algorithm takes a parameter k represented in unary form 1^k and an accumulation threshold n (an upper bound on the number of values that can be accumulated securely), returns an accumulator key a_k. The accumulator key a_k is part of the q-SDH setup [BB08] and therefore is public to all parties.

- Eval$(a_k, \mathcal{D})$: This deterministic algorithm takes an accumulator key a_k and a set $\mathcal{D}$ of values to be accumulated, returns an accumulation value z for the value set $\mathcal{D}$.

- CreateWit$(a_k, z, d_i, \mathcal{D})$: This algorithm takes an accumulator key a_k, an accumulation value z for $\mathcal{D}$ and a value d_i, returns $\perp$ if $d_i \notin \mathcal{D}$, and a witness w_i if $d_i \in \mathcal{D}$.

- Verify(a_k, z, w_i, d_i): This algorithm takes an accumulator key a_k, an accumulation value z for $\mathcal{D}$, a witness w_i and a value d_i, returns true if w_i is the witness for $d_i \in \mathcal{D}$, and false otherwise.

In this paper, we use *collision resistant bilinear accumulators* from Nguyen [Ngu05] as cryptographic accumulators which generates constant sized witness, but requires q-SDH assumption. Alternatively, we can use Merkle trees [Mer88] (and avoid q-SDH assumption) at the expense of $O(\log n)$ multiplicative communication. The bilinear accumulator from Nyugen [Ngu05] satisfies the following property:

Lemma 1 (Collision resistant accumulator [Ngu05])**.** *The cryptographic accumulator proposed in [Ngu05] (named collision-free bilinear accumulator) satisfies the following property. For any set of size n and a probabilistic polynomial-time adversary $\mathcal{A}$, the following function is negligible in κ:*

$$\mathsf{P}r \left[\begin{array}{c} ak \leftarrow \mathsf{Gen}(1^\kappa, n), \\ (\{d_1, \ldots, d_n\}, d', w') \leftarrow \\ \mathcal{A}(1^\kappa, n, ak), \\ z \leftarrow \mathsf{Eval}(ak, \{d_1, \ldots, d_n\}) \end{array} \middle| \begin{array}{c} (d' \notin \{d_1, ..., d_n\}) \wedge \\ (\mathsf{Verify}(ak, z, w', d') = 1) \end{array} \right]$$

Non-interactive threshold signature scheme. We use (t, n) *non-interactive threshold signature scheme* of Cachin et al. [CKS00] in one of our protocols. The threshold signature scheme is secure against static adversary. The signature scheme consists of the following efficient algorithms:

- The randomized *key generation* algorithm $\mathsf{KeyGen_{TS}}$ that takes a security parameter κ as input and outputs a tuple $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ of secret keys, a tuple $(\mathsf{pk}_1, \ldots, \mathsf{pk}_n)$ and a common public key pk.

- The deterministic signing algorithm $\mathsf{Sign_{TS}}$ that takes as input sk_i and a message m and outputs a signature σ_i on m.

- The deterministic *share verification* algorithm $\mathsf{ShareVerify_{TS}}$ that takes as input public key pk_i, a signature share σ_i and tuple (i, m). It outputs a bit $b \in \{0, 1\}$ indicating whether σ_i is a valid signature share on m under secret key sk_i.

- The deterministic *combining* $\mathsf{Combine_{TS}}$ takes as input a tuple of public keys $(\mathsf{pk}_1, \ldots, \mathsf{pk}_n)$, a message m, and a list of $t + 1$ pairs (i, σ_i). It outputs either a signature σ on m or $\perp$, if (i, σ_i) contains ill-formed signature shares.

- The deterministic *verification* algorithm $\mathsf{Verify_{TS}}$ takes as input a signature σ, a message m and a common public key pk. It outputs a bit $b \in \{0, 1\}$ indicating whether σ is a valid signature on m.

The threshold signature scheme satisfies robustness (i.e., it is computationally infeasible for an adversary to produce $t + 1$ valid signature shares such that the output of the share combining algorithm is not a valid signature) and unforgeability (i.e., it is computationally infeasible for the adversary to output a valid signature on a message m given t signature shares on m).

Non-Interactive Proof-of-Equivalence of commitments [KHG12]. Given two commitments $\mathcal{C}_{\langle g \rangle}(s) = g^s$ and $\mathcal{C}_{\langle g, h \rangle}(s, r) = g^s h^r$ to the same value s for generators

$g, h \in \mathbb{G}$ and $s, r \in \mathbb{Z}_p$, a prover proves that she knows s and r such that $\mathcal{C}_{\langle g \rangle}(s) = g^s$ and $\mathcal{C}_{\langle g,h \rangle}(s, r) = g^s h^r$. We denote it by $\mathsf{NIZKPK}_{\equiv \mathsf{Com}}(s, r, g, h, \mathcal{C}_{\langle g \rangle}(s), \mathcal{C}_{\langle g,h \rangle}(s,r)) = \pi_{\equiv \mathsf{Com}} \in \mathbb{Z}_p^3$. A full construction of $\mathsf{NIZKPK}_{\equiv \mathsf{Com}}$ is as follows:

- Pick $v_1, v_2 \in_R \mathbb{Z}_p$, and let $t_1 = g^{v_1}$ and $t_2 = h^{v_2}$.
- Compute hash $c = \mathsf{H}_{\equiv \mathsf{Com}}(g, h, \mathcal{C}_{\langle g \rangle}(s), \mathcal{C}_{\langle g,h \rangle}(s,r), t_1, t_2)$, where $\mathsf{H}_{\equiv \mathsf{Com}} : \mathbb{G}^6 \to \mathbb{Z}_p$ is a random oracle hash function.
- Let $u_1 = v_1 - c \cdot s$ and $u_2 = v_2 - c \cdot r$.
- Send the proof $\pi_{\equiv \mathsf{Com}} = (c, u_1, u_2)$ along with $\mathcal{C}_{\langle g \rangle}(s)$ and $\mathcal{C}_{\langle g,h \rangle}(s,r)$.

The verifier checks this proof (given $\pi_{\equiv \mathsf{Com}}$, g, h, $\mathcal{C}_{\langle g \rangle}(s)$, $\mathcal{C}_{\langle g,h \rangle}(s,r)$) as follows:

- Let $t_1' = g^{u_1} \mathcal{C}_{\langle g \rangle}(s)^c$ and $t_2' = h^{u_2}(\frac{\mathcal{C}_{\langle g,h \rangle}(s,r)}{\mathcal{C}_{\langle g \rangle}(s)})^c$.
- Accept the proof as valid if $c = \mathsf{H}_{\equiv \mathsf{Com}}(g, h, \mathcal{C}_{\langle g \rangle}(s), \mathcal{C}_{\langle g,h \rangle}(s,r), t_1', t_2')$.

Normalizing the length of cryptographic building blocks. Let λ denote the security parameter, $\kappa_h = \kappa_h(\lambda)$ denote the hash size, $\kappa_a = \kappa_a(\lambda)$ denote the size of the accumulation value and witness of the accumulator and $\kappa_v = \kappa_v(\lambda)$ denote the size of secret share and witness of a secret. Further, let $\kappa = \mathsf{max}(\kappa_h, \kappa_a, \kappa_v)$; we assume $\kappa = \Theta(\kappa_h) = \Theta(\kappa_v) = \Theta(\kappa_a) = \Theta(\lambda)$. Throughout the paper, we can use the same parameter κ to denote the hash size, signature size, accumulator size and secret share size for convenience.

3 Secure DKG with Two Broadcast Rounds

We first present a secure DKG protocol assuming a broadcast channel motivated from Gennaro et al. DKG [GJKR07]. The presented DKG reduces the number of required rounds with broadcast to two, which is a significant improvement over [GJKR07] requiring three broadcast rounds in the best case and five broadcast rounds otherwise.[3] In later sections, we replace the broadcast channel with a novel consensus primitives to design communication-efficient DKG protocols.

Gennaro et al. [GJKR07] presented a secure DKG protocol that produces uniform public keys based on Pedersen's VSS [Ped92]. In their protocol, each party, as a dealer, selects a secret uniformly at random and shares the secret using Pedersen's VSS protocol. Since Pedersen's VSS provides information theoretic secrecy guarantees, the adversary has no information about the public key and hence cannot bias it. At the end of the secret sharing, a set of qualified parties QUAL who correctly shared their secret is defined. Once the set QUAL is fixed, parties in set QUAL invoke an additional round of secret sharing using Feldman's VSS [Fel87] to generate the final public key. While this approach ensures generation of uniform keys and maintains secrecy, it adds additional overhead as it incurs more latency and communication to perform additional secret sharing. In addition to the above overhead, Pedersen VSS requires three broadcast rounds. In particular, parties post the commitment, complaints and secret shares corresponding to the complaints on to the broadcast channel during the sharing phase.

The protocol in Figure 2 improves upon the DKG protocol of Gennaro et al. [GJKR07] in the following ways.

Improving latency in the sharing phase. We improve latency by reducing information posted on the broadcast channel by using improved eVSS (iVSS) protocol [BSL+21] which requires only 2 broadcast rounds.[4] Reducing the broadcast rounds greatly improves latency as broadcast channels are generally instantiated using Byzantine broadcast or Byzantine agreement protocols which have worst-case linear round complexity.

---

[3] Using NIZK similar to us, the number of rounds for Gennaro et al. DKG [GJKR07] can be reduced to two in the best case and three otherwise in a rather straightforward manner; however, reducing to two broadcast rounds in all situations is the key challenge here.

[4] Alternatively, we can use broadcast optimal VSS protocol of Backes et al. [BKP11] which has 2 broadcast rounds. We prefer iVSS protocol for its simplicity.
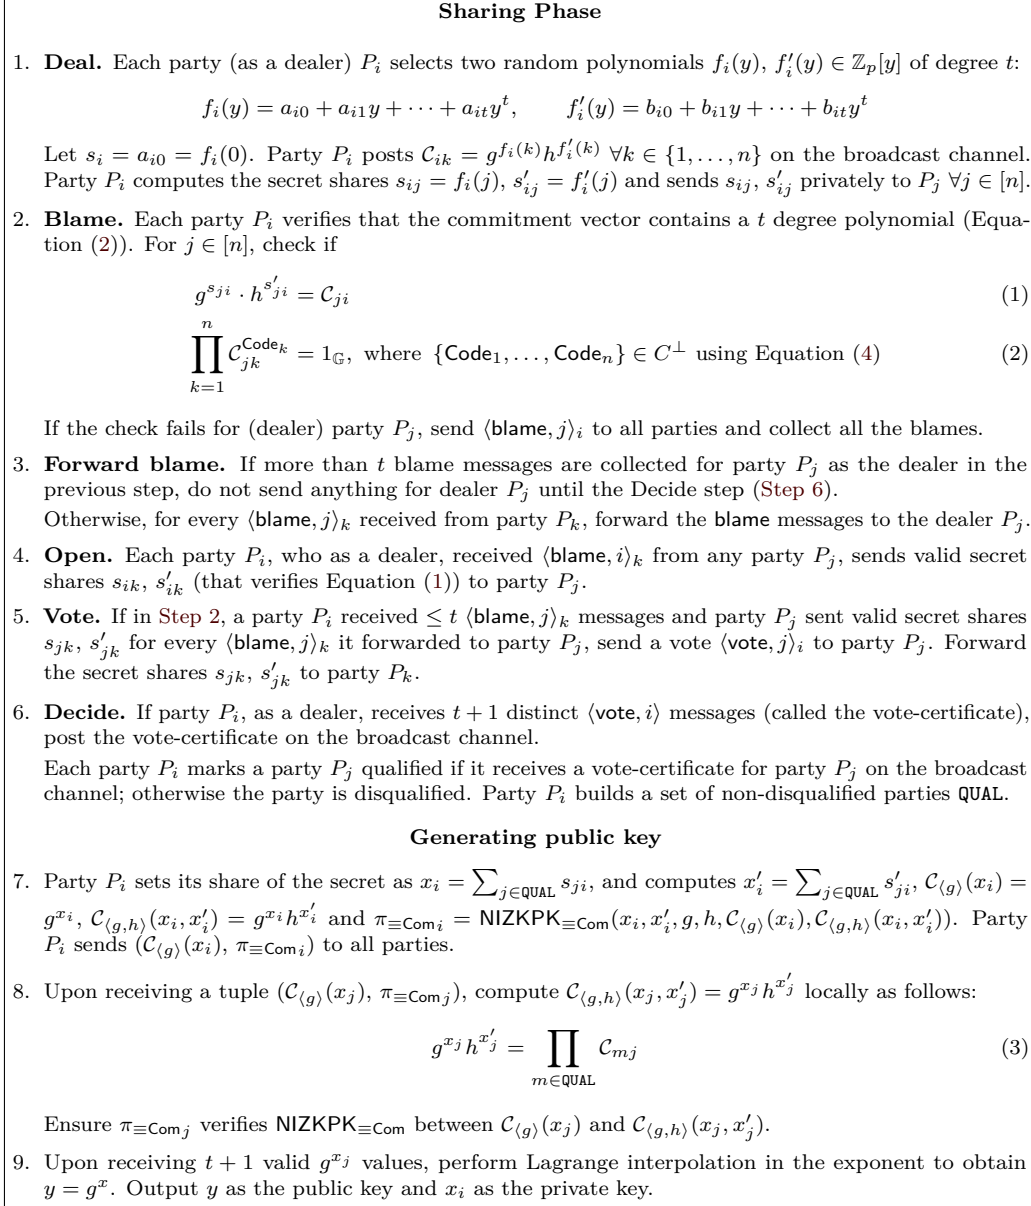

Sharing Phase

1. **Deal.** Each party (as a dealer) P_i selects two random polynomials $f_i(y)$, $f_i'(y) \in \mathbb{Z}_p[y]$ of degree t:

$$f_i(y) = a_{i0} + a_{i1}y + \cdots + a_{it}y^t, \qquad f_i'(y) = b_{i0} + b_{i1}y + \cdots + b_{it}y^t$$

   Let $s_i = a_{i0} = f_i(0)$. Party P_i posts $\mathcal{C}_{ik} = g^{f_i(k)}h^{f_i'(k)} \ \forall k \in \{1, \ldots, n\}$ on the broadcast channel. Party P_i computes the secret shares $s_{ij} = f_i(j)$, $s_{ij}' = f_i'(j)$ and sends s_{ij}, s_{ij}' privately to P_j $\forall j \in [n]$.

2. **Blame.** Each party P_i verifies that the commitment vector contains a t degree polynomial (Equation (2)). For $j \in [n]$, check if

$$g^{s_{ji}} \cdot h^{s_{ji}'} = \mathcal{C}_{ji} \tag{1}$$

$$\prod_{k=1}^{n} \mathcal{C}_{jk}^{\mathsf{Code}_k} = 1_{\mathbb{G}}, \text{ where } \{\mathsf{Code}_1, \ldots, \mathsf{Code}_n\} \in C^{\perp} \text{ using Equation (4)} \tag{2}$$

   If the check fails for (dealer) party P_j, send $\langle \mathsf{blame}, j \rangle_i$ to all parties and collect all the blames.

3. **Forward blame.** If more than t blame messages are collected for party P_j as the dealer in the previous step, do not send anything for dealer P_j until the Decide step (Step 6).

   Otherwise, for every $\langle \mathsf{blame}, j \rangle_k$ received from party P_k, forward the blame messages to the dealer P_j.

4. **Open.** Each party P_i, who as a dealer, received $\langle \mathsf{blame}, i \rangle_k$ from any party P_j, sends valid secret shares s_{ik}, s_{ik}' (that verifies Equation (1)) to party P_j.

5. **Vote.** If in Step 2, a party P_i received $\leq t$ $\langle \mathsf{blame}, j \rangle_k$ messages and party P_j sent valid secret shares s_{jk}, s_{jk}' for every $\langle \mathsf{blame}, j \rangle_k$ it forwarded to party P_j, send a vote $\langle \mathsf{vote}, j \rangle_i$ to party P_j. Forward the secret shares s_{jk}, s_{jk}' to party P_k.

6. **Decide.** If party P_i, as a dealer, receives $t + 1$ distinct $\langle \mathsf{vote}, i \rangle$ messages (called the vote-certificate), post the vote-certificate on the broadcast channel.

   Each party P_i marks a party P_j qualified if it receives a vote-certificate for party P_j on the broadcast channel; otherwise the party is disqualified. Party P_i builds a set of non-disqualified parties QUAL.

Generating public key

7. Party P_i sets its share of the secret as $x_i = \sum_{j \in \mathtt{QUAL}} s_{ji}$, and computes $x_i' = \sum_{j \in \mathtt{QUAL}} s_{ji}'$, $\mathcal{C}_{\langle g \rangle}(x_i) = g^{x_i}$, $\mathcal{C}_{\langle g,h \rangle}(x_i, x_i') = g^{x_i}h^{x_i'}$ and $\pi_{\equiv \mathsf{Com}_i} = \mathsf{NIZKPK}_{\equiv \mathsf{Com}}(x_i, x_i', g, h, \mathcal{C}_{\langle g \rangle}(x_i), \mathcal{C}_{\langle g,h \rangle}(x_i, x_i'))$. Party P_i sends $(\mathcal{C}_{\langle g \rangle}(x_i), \pi_{\equiv \mathsf{Com}_i})$ to all parties.

8. Upon receiving a tuple $(\mathcal{C}_{\langle g \rangle}(x_j), \pi_{\equiv \mathsf{Com}_j})$, compute $\mathcal{C}_{\langle g,h \rangle}(x_j, x_j') = g^{x_j}h^{x_j'}$ locally as follows:

$$g^{x_j}h^{x_j'} = \prod_{m \in \mathtt{QUAL}} \mathcal{C}_{mj} \tag{3}$$

   Ensure $\pi_{\equiv \mathsf{Com}_j}$ verifies $\mathsf{NIZKPK}_{\equiv \mathsf{Com}}$ between $\mathcal{C}_{\langle g \rangle}(x_j)$ and $\mathcal{C}_{\langle g,h \rangle}(x_j, x_j')$.

9. Upon receiving $t + 1$ valid g^{x_j} values, perform Lagrange interpolation in the exponent to obtain $y = g^x$. Output y as the public key and x_i as the private key.

Figure 2: **Secure distributed key generation in dlog-based cryptosystems**

In iVSS, the dealer posts commitments on the broadcast channel and privately sends the secret shares to each party. Instead of posting the complaints on the broadcast channel, parties multicast blame message if they receive invalid secret shares or receive no secret shares at all. Parties then forward all blame messages to the dealer[5]. The dealer is expected to send secret shares corresponding to the blame messages (i.e., secret shares s_{ij}, s'_{ij} if a P_j sent blame message against dealer P_i). If the dealer sends all secret shares corresponding to the blame message it forwarded, a party sends a vote message to the dealer. Upon receiving $t + 1$ vote messages, the dealer posts a vote-certificate containing $t + 1$ vote messages. Honest parties consider the dealer to be honest if they see the vote-certificate on the broadcast channel.

Observe that using iVSS scheme, the dealer posts only the commitment and vote-certificate on the broadcast channel. This improves the sharing phase by one broadcast round.

Using commitments to evaluations instead of commitments to coefficients. In VSS such as Pedersen's VSS and Feldman's VSS and thus in [GJKR07], commitments to the secret share are commitments to the coefficients of a t-degree polynomial, which imply verifying a share requires $O(t)$ computations. This results in $O(nt)$ computations per VSS instance in the complaint stage (where every party verifies opening of up to t complaints) and during reconstruction. SCRAPE [CD17, Section 2.1] showed how to commit (using discrete log commitments) to evaluations instead of coefficients of the polynomial and verify that the committed evaluations are of a degree t polynomial by using the property of coding schemes: if C is the code space for an (n, t) sharing, then the following vector

$$C^\perp := \{\mathsf{Code}_1, \ldots, \mathsf{Code}_n; \mathsf{Code}_i = \mathsf{poly}(i) \prod_{j=1, j \neq i}^{n} 1/(i - j)$$

$$\mathsf{poly}(x) \text{ is a random polynomial of degree } n - t + 1\} \quad (4)$$

is orthogonal to C. We can check that the Pedersen's commitments to the evaluations are an (n, t) sharing (see Equation (1)). If λ is $\log_g h$, then commitments to evaluations form a polynomial $g^f h^{f'} = g^{f + \lambda f'}$ which is another (n, t) polynomial thereby allowing to use the coding technique. This is an information-theoretic technique and therefore does not affect the security of the underlying VSS.

Removing additional secret sharing while generating public key. We remove the additional secret sharing performed using Feldman's VSS by taking an alternate approach [KHG12]. Instead of executing an additional secret sharing, assuming random oracle, we make use of the NIZK proof of equivalence of commitments NIZKPK$_{\equiv\mathsf{Com}}$ to generate the public key. This approach does not require additional secret sharing via Feldman's VSS. Once the sharing phase is completed, a set of qualified parties QUAL is finalized. Then, each party P_i computes its share of the shared secrets i.e., $x_i = \sum_{P_j \in \mathtt{QUAL}} s_{ji}$ and $x'_i = \sum_{P_j \in \mathtt{QUAL}} s'_{ji}$ along with commitments $\mathcal{C}_{\langle g \rangle}(x_i)$, $\mathcal{C}_{\langle g, h \rangle}(x_i, x'_i)$. It then multicasts commitment of its share $\mathcal{C}_{\langle g \rangle}(x_i)$ and the corresponding NIZKPK$_{\equiv\mathsf{Com}}$ proof $\pi_{\equiv\mathsf{Com}i}$ to prove P_i knows x_i and x'_i.

All parties can compute the commitment $\mathcal{C}_{\langle g, h \rangle}(x_i, x'_i)$ locally as shown in Equation (3) and verify the correctness of commitment $\mathcal{C}_{\langle g \rangle}(x_i)$ using $\pi_{\equiv\mathsf{Com}i}$. The final public key Y is computed via Lagrange interpolation in the exponent using $t + 1$ distinct commitments $\mathcal{C}_{\langle g \rangle}(x_i)$.

3.1   Analysis of Secure DKG

We rely on the following Lemma of [Ped92].

---

[5]In an implementation, we can only forward up to t blames instead of all the blames.

Lemma 2 ([Ped92]). *Under the discrete-log assumption, Pedersen's VSS satisfies the following properties in the presence of a polynomially bounded adversary that corrupts up to t parties.*

(i) If the dealer is not disqualified during the sharing phase, then all honest parties hold secret shares that interpolates to unique polynomial of degree t. In particular, any $t + 1$ of these shares suffice to reconstruct the secret σ.

(ii) The protocol produces information (i.e., commitments $\mathcal{C}_k$ and secret shares σ_i) that can be used at reconstruction time to test for the correctness of each secret share; thus, reconstruction is possible, even in the presence of malicious parties, from any subset of shares containing at least $t + 1$ correct secret shares.

(iii) The view of the adversary is independent of the value of the secret σ, and therefore the secrecy of σ is unconditional.

Note that Lemma 2 also holds when using evaluations instead of coefficients as discussed in Section 8. The coding check (see Equation (2)) ensures that the shared commitments to evaluations are indeed a t degree polynomial except with $1/p$ probability in $\mathbb{Z}_p$. Since p is sufficiently large ($\mathsf{poly}(\kappa)$), the probability of the check failing is negligible in the security parameter.

Fact 7. *If a dealer P_i receives a vote-certificate, all honest parties must have received their corresponding secret shares s_{ij}, s'_{ij}.*

Proof. Suppose a dealer P_i receives a vote-certificate i.e, $t + 1$ vote messages. At least one of the vote messages is sent by an honest party (say P_j). An honest party P_j sends a vote message only when it receives no blame messages or receives up to t blame messages and dealer P_i sent secret shares s_{ik}, s'_{ik} for every $\langle \mathsf{blame}, i \rangle_k$ message it forwarded.

If party P_j received no blame messages, all honest parties must have received their corresponding secret shares s_{ij}, s'_{ij}; otherwise honest parties would have sent blame messages. On the other hand, if party P_j received $f \leq t$ blame messages, $n - t - f$ honest parties must have received their corresponding secret shares; otherwise, these honest parties would have sent blame messages and party P_j would have received more than f blame messages. Since party P_j forwards secret shares s_{ik}, s'_{ik} to party P_k for every $\langle \mathsf{blame}, i \rangle_k$ message it received, all honest parties must have received corresponding secret shares. $\square$

Theorem 8. *Under discrete-log assumption and random oracle, the protocol in Figure 2 is a secure protocol for distributed key generation in dlog-based cryptosystem tolerating $t < n/2$ Byzantine faults.*

Proof. We first prove correctness of the protocol. Observe that all honest parties build the same set of non-disqualified parties `QUAL` in Step 6. This is true because the commitment to the shared polynomials and vote-certificates are posted on the broadcast channel and broadcast channel ensures all honest parties output a common value.

Note that if a party $P_j \in$ `QUAL`, it must have posted its commitment and vote-certificate on the broadcast channel. By Fact 7, all honest parties have received secret shares shared by party P_j. This implies party P_j is not disqualified during the sharing phase. By part (i) of Lemma 2, all honest parties hold correct secret shares and any $t + 1$ of these secret shares suffices to reconstruct the secret s_j. This is true for all parties $P_j \in$ `QUAL`. Since, the secret key x is sum of individual secret s_j contributed by $P_j \in$ `QUAL` and each secret s_j can be reconstructed using Lagrange interpolation via a combination of $t + 1$ secret shares provided by honest parties, the secret key x can be reconstructed via $t + 1$ shares provided by honest parties. This proves property C1 of a secure DKG protocol.

By part (ii) of Lemma 2, there exists information (i.e., commitments) that can be used to verify correctness of each secret share. Observe that each honest party P_j sends
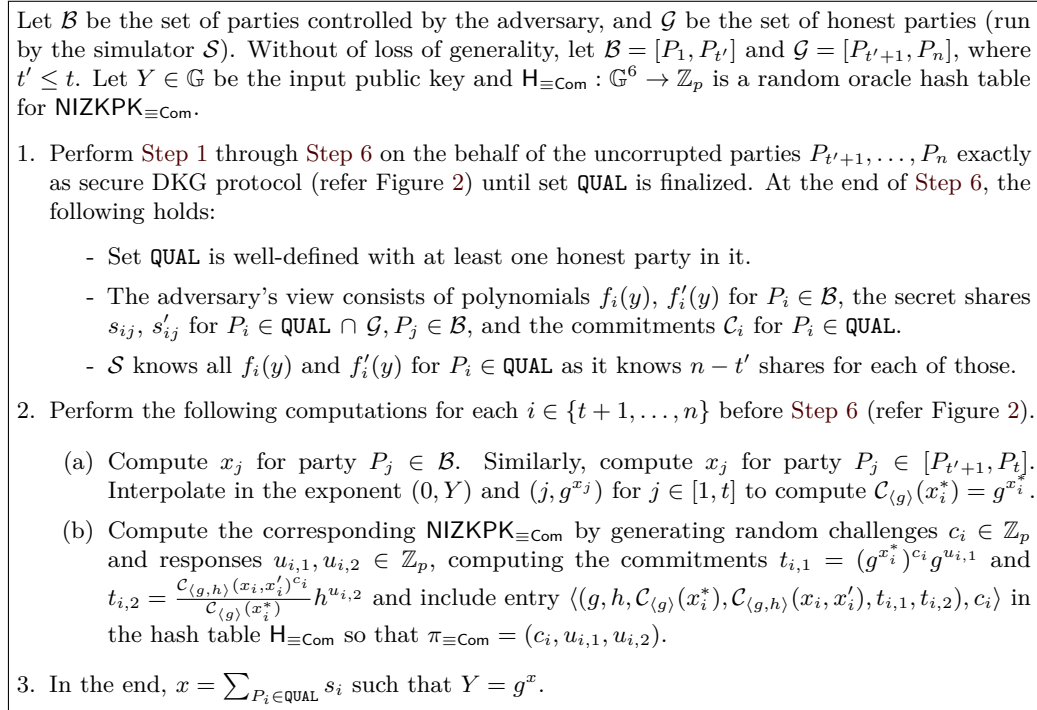

Let $\mathcal{B}$ be the set of parties controlled by the adversary, and $\mathcal{G}$ be the set of honest parties (run by the simulator $\mathcal{S}$). Without of loss of generality, let $\mathcal{B} = [P_1, P_{t'}]$ and $\mathcal{G} = [P_{t'+1}, P_n]$, where $t' \leq t$. Let $Y \in \mathbb{G}$ be the input public key and $\mathsf{H}_{\equiv\mathsf{Com}} : \mathbb{G}^6 \to \mathbb{Z}_p$ is a random oracle hash table for $\mathsf{NIZKPK}_{\equiv\mathsf{Com}}$.

1. Perform Step 1 through Step 6 on the behalf of the uncorrupted parties $P_{t'+1}, \ldots, P_n$ exactly as secure DKG protocol (refer Figure 2) until set QUAL is finalized. At the end of Step 6, the following holds:

   - Set QUAL is well-defined with at least one honest party in it.

   - The adversary's view consists of polynomials $f_i(y)$, $f_i'(y)$ for $P_i \in \mathcal{B}$, the secret shares s_{ij}, s_{ij}' for $P_i \in$ QUAL $\cap$ $\mathcal{G}, P_j \in \mathcal{B}$, and the commitments $\mathcal{C}_i$ for $P_i \in$ QUAL.

   - $\mathcal{S}$ knows all $f_i(y)$ and $f_i'(y)$ for $P_i \in$ QUAL as it knows $n - t'$ shares for each of those.

2. Perform the following computations for each $i \in \{t+1, \ldots, n\}$ before Step 6 (refer Figure 2).

   (a) Compute x_j for party $P_j \in \mathcal{B}$. Similarly, compute x_j for party $P_j \in [P_{t'+1}, P_t]$. Interpolate in the exponent $(0, Y)$ and (j, g^{x_j}) for $j \in [1, t]$ to compute $\mathcal{C}_{\langle g \rangle}(x_i^*) = g^{x_i^*}$.

   (b) Compute the corresponding $\mathsf{NIZKPK}_{\equiv\mathsf{Com}}$ by generating random challenges $c_i \in \mathbb{Z}_p$ and responses $u_{i,1}, u_{i,2} \in \mathbb{Z}_p$, computing the commitments $t_{i,1} = (g^{x_i^*})^{c_i} g^{u_{i,1}}$ and $t_{i,2} = \frac{\mathcal{C}_{\langle g, h \rangle}(x_i, x_i')^{c_i}}{\mathcal{C}_{\langle g \rangle}(x_i^*)} h^{u_{i,2}}$ and include entry $\langle (g, h, \mathcal{C}_{\langle g \rangle}(x_i^*), \mathcal{C}_{\langle g, h \rangle}(x_i, x_i'), t_{i,1}, t_{i,2}), c_i \rangle$ in the hash table $\mathsf{H}_{\equiv\mathsf{Com}}$ so that $\pi_{\equiv\mathsf{Com}} = (c_i, u_{i,1}, u_{i,2})$.

3. In the end, $x = \sum_{P_i \in \mathsf{QUAL}} s_i$ such that $Y = g^x$.

Figure 3: Simulator for Secure DKG

g^{x_j} and $\mathsf{NIZKPK}_{\equiv\mathsf{Com}}$ proof $\pi_{\equiv\mathsf{Com}j}$ at the end of sharing phase. Each party P_i can verify correctness of $\mathcal{C}_{\langle g \rangle}(x_j)$ by checking Equation (3). A valid $\mathsf{NIZKPK}_{\equiv\mathsf{Com}}$ proof $\pi_{\equiv\mathsf{Com}j}$ proves in zero knowledge that party P_j knows x_j and x_j' thus proving the correctness of g^{x_j}. By using $t+1$ valid g^{x_j}, honest parties can compute the same g^x via Lagrange interpolation in the exponent which is the public key. This proves property C2 of a secure DKG protocol.

Observe that the secret key x is the sum of secrets shared by parties in QUAL which contains at least one honest party and honest parties select their secret uniformly at random. This suffices to prove property C3 of a secure DKG protocol.

We now prove secrecy. Our proof of secrecy is based on the proof of secrecy in earlier works [GJKR07, KHG12]. We provide a simulator $\mathcal{S}$ for our secure DKG protocol in Figure 3. Without loss of generality, we assume the adversary $\mathcal{A}$ compromises parties $P_1, \ldots, P_{t'}$, where $t' \leq t$, denoted by set $\mathcal{B}$. The rest of the parties $P_{t'+1}, \ldots, P_n$, denoted by set $\mathcal{G}$ are controlled by the simulator.

Informally, the simulator $\mathcal{S}$ with input Y runs as follows. $\mathcal{S}$ will run on the behalf of the honest parties $\mathcal{G}$ Step 1 until Step 6 following exactly the instructions. At this point, the set QUAL is well-defined and $\mathcal{S}$ knows all $f_i(y)$ and $f_i'(y)$ for $P_i \in$ QUAL as it knows $n - t'$ shares for each of those. Observe that the view of adversary $\mathcal{A}$ that interacts with $\mathcal{S}$ is identical to the view of $\mathcal{A}$ that interacts with honest parties in a regular run of the protocol. In particular, $\mathcal{A}$ sees the following distribution of data:

- Polynomials $f_i(y)$, $f_i'(y)$ for $P_i \in \mathcal{B}$

- Values $f_i(j)$, $f_i'(j)$ for $i \in \mathcal{G}$, $j \in \mathcal{B}$ and values $\mathcal{C}_i$ for $P_i \in$ QUAL

$\mathcal{S}$ will then change the secret shared by one honest party (say P_n) to "hit" the desired public key Y such that the above data distribution observed by $\mathcal{A}$ remains identical. For parties $P_i \in (\mathcal{G} \setminus \{P_n\})$, the input polynomial $f_i(y)$ and $f_i'(y)$ remains identical. Thus, their data distribution remains identical. For party P_n, the input polynomial is modified such that $g^{f_n^*(0)} = g^{s_n^*} = \frac{Y}{\prod_{P_j \in \mathsf{QUAL} \setminus \{P_n\}} g^{s_i}}$ and $f_n^*(j) = s_{nj}$ for $j \in [1, t]$. Define $f'^*(y)$

such that $f_n^*(y) + \lambda f_n'^*(y) = f_n(y) + \lambda f_n'(y)$, where $\lambda = \log_g(h)$. Observe that for these polynomials, the evaluations and commitments seen by parties in $\mathcal{B}$ is identical to the real run of the protocol.

Simulator $\mathcal{S}$ will then compute g^{x_j} for party $P_j \in [P_1, P_t]$ and interpolate in the exponent $(0, Y)$ and (j, g^{x_j}) for $j \in [1, t]$ to compute $\mathcal{C}_{\langle g \rangle}(x_i^*) = g^{x_i^*}$ and the corresponding $\mathsf{NIZKPK}_{\equiv \mathsf{Com}}$ and publish these values. Observe that these values pass the verification in the real run of protocol.

It remains to be shown that polynomials $f_i^*(y)$ and $f_i'^*(y)$ belong to the right distribution. For $\mathsf{QUAL} \setminus (\mathcal{G} \setminus \{P_n\})$, this is trivially true as they are defined identically to $f_i(y)$ and $f_i'(y)$ which were chosen uniformly at random. For f_n^*, the polynomial evaluates to random values $f_n(j)$ at $j \in [1, t]$ and evaluates to $\log_g(s_n^*)$ required to hit Y. Finally, $f_n'^*(y)$ is defined as $f_n^*(y) + \lambda f_n'^*(y) = f_n(y) + \lambda f_n'(y)$, and since $f_n'(y)$ is chosen to be random, so is $f_n'^*(y)$. $\qquad\square$

4   Communication Optimal Weak Gradecast

One of the main tools in the design of our communication efficient protocols is our communication optimal *weak gradecast* protocol. Our weak gradecast protocol has a communication complexity of $O(n\ell + \kappa n^2)$ for ℓ bit input and incurs 4 rounds. In Section 10, we show a quadratic lower bound on the communication complexity of weak gradecast for completeness.
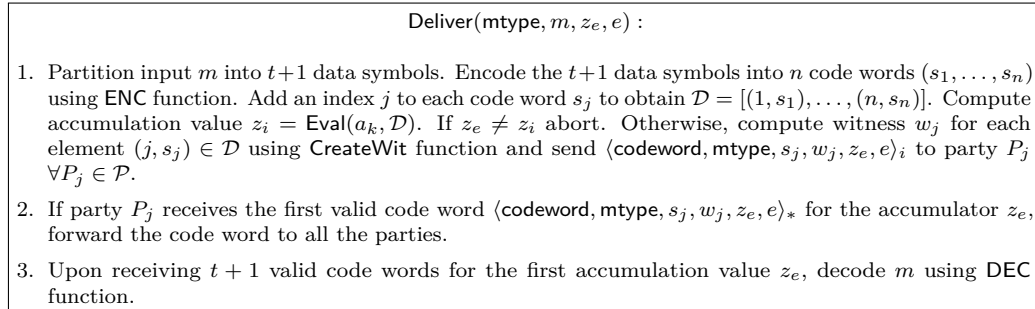

$\mathsf{Deliver}(\mathsf{mtype}, m, z_e, e)$:

1. Partition input m into $t+1$ data symbols. Encode the $t+1$ data symbols into n code words $(s_1, \ldots, s_n)$ using ENC function. Add an index j to each code word s_j to obtain $\mathcal{D} = [(1, s_1), \ldots, (n, s_n)]$. Compute accumulation value $z_i = \mathsf{Eval}(a_k, \mathcal{D})$. If $z_e \neq z_i$ abort. Otherwise, compute witness w_j for each element $(j, s_j) \in \mathcal{D}$ using $\mathsf{CreateWit}$ function and send $\langle \mathsf{codeword}, \mathsf{mtype}, s_j, w_j, z_e, e \rangle_i$ to party P_j $\forall P_j \in \mathcal{P}$.

2. If party P_j receives the first valid code word $\langle \mathsf{codeword}, \mathsf{mtype}, s_j, w_j, z_e, e \rangle_*$ for the accumulator z_e, forward the code word to all the parties.

3. Upon receiving $t + 1$ valid code words for the first accumulation value z_e, decode m using DEC function.

Figure 4: **Deliver function**

Deliver. As a building block, we first present a $\mathsf{Deliver}$ function (refer Figure 4) used by an honest party to efficiently propagate long messages. This function is adapted from RandPiper [BSL+21] where linear-sized messages are propagated among all honest parties with $O(\kappa n^2)$ communication cost. The $\mathsf{Deliver}$ function enables efficient propagation of long messages using erasure coding techniques and cryptographic accumulators. The input parameters to the function are a keyword mtype, long message m, accumulation value z_e corresponding to message m and epoch e in which $\mathsf{Deliver}$ function is invoked. The input keyword mtype corresponds to message *type* containing long message m sent by its sender. In order to facilitate efficient equivocation by the sender, the input keyword mtype, hash of long message m, accumulation value z_e, and epoch e are signed by the sender of message m. We omit epoch parameter when the $\mathsf{Deliver}$ function is not invoked within an epoch. The $\mathsf{Deliver}$ function incurs 2 rounds.

The gradecast protocol is presented in Figure 5. In round 1, the designated sender P_j sends value v by multicasting $\langle \mathsf{gcast}, v, z \rangle_j$ where z is the accumulation value for value v. We note that the size of input value v can be large. To facilitate efficient equivocation checks, the sender P_j signs $\langle \mathsf{gcast}, H(v), z \rangle$ and sends v separately. Whenever an equivocation by the sender is detected, multicasting signed hashes suffices to prove equivocation by the sender. All-to-all multicasting of κ-sized signed hashes incurs only

---

Set $o_i = \perp$ and $g_i = \perp$. Each party P_i performs the following operations:

- **Round 1:** If party P_j is the designated sender, then it multicasts its input value v in the form of $\langle \mathsf{gcast}, v, z \rangle_j$ where z is the accumulation value of v.

- **Round 2:** If party P_i receives $pr := \langle \mathsf{gcast}, v, z \rangle_j$ for the first time, then invoke $\mathsf{Deliver}(\mathsf{gcast}, pr, z)$.

- **Round 4:** If party P_i invoked $\mathsf{Deliver}$ in round 2 and no party P_j equivocation has been detected so far, set $o_i = v$ and $g_i = 2$. Let v_i be the first value received. If $v_i = \perp$, set $o_i = \perp$ and $g_i = 0$, else if $o_i = \perp$, set $o_i = v_i$ and $g_i = 1$. Output (o_i, g_i).

- **At any round:** If equivocating hashes signed by party P_j are detected, multicast the equivocating hashes.

Figure 5: **Weak Gradecast with $O(n\ell + (\kappa + w)n^2)$ communication.**

$O(\kappa n^2)$ in communication. The reduction in communication is obtained via the use of efficient erasure coding schemes [RS60], cryptographic accumulators [BP97] and multicast of equivocating hashes (if any).

In round 2, if party P_i receives $\langle \mathsf{gcast}, v, z \rangle_j$, it invokes $\mathsf{Deliver}$ to propagate long message v. Note that $\mathsf{Deliver}$ function requires 2 rounds. Round 3 accommodates steps of $\mathsf{Deliver}$ function invoked in rounds 2. In round 4, each party P_i sets its output value and grade as follows. If party P_i received $\langle \mathsf{gcast}, v, z \rangle_j$ in round 2 and did not detect any equivocation so far, it outputs value v with a grade of 2. Otherwise, party P_i outputs the first value it received with a grade of 1. If no value has been received, P_i outputs $\perp$ with a grade of 0. Note that if party P_i receives $\langle \mathsf{gcast}, v, z \rangle_j$ in round 2, but detects an equivocation by round 4, it outputs value v with a grade of 1.

The first value. If party P_i receives a valid code word corresponding to value v or $\langle \mathsf{gcast}, v, z \rangle_j$ before receiving any other values (i.e, $\langle \mathsf{gcast}, v', z' \rangle_j$ or a valid code word for any other value), then value v is the first value for P_i. Hereafter, party P_i only decodes value v when it receives $t + 1$ valid code words for it.

4.1   Analysis of Gradecast

Lemma 3. *Suppose party P_j is the designated sender. If an honest party invokes $\mathsf{Deliver}$ in round r for a value m sent by party P_j and no honest party has detected a party P_j equivocation by round $r + 1$, then all honest parties will receive value m by round $r + 2$.*

Proof. Suppose an honest party P_i invokes $\mathsf{Deliver}$ at round r for a value m sent by party P_j. Party P_i must have sent valid code words and witness $\langle \mathsf{codeword}, \mathsf{mtype}, s_k, w_k, z_e, e \rangle_i$ computed from value m to every party P_k $\forall k \in [n]$ at round r. The code words and witness arrive at all honest parties by round $r + 1$.

Since no honest party has detected a party P_j equivocation by round $r + 1$, it must be that either honest parties will forward their code word $\langle \mathsf{codeword}, \mathsf{mtype}, s_k, w_k, z_e, e \rangle$ when they receive the code words sent by party P_i or they already sent the corresponding code word when they either invoked $\mathsf{Deliver}$ for value m or received the code word from some other party. In any case, all honest parties will forward their code word corresponding to value m by round $r + 1$. Thus, all honest parties will have received $t + 1$ valid code words for a common accumulation value z_e by round $r + 2$ sufficient to decode value m. □

Theorem 9. *The protocol in Figure 5 is a gradecast protocol satisfying Definition 2.*

Proof. Suppose party P_j is the designated sender with its input value v.

We first consider the case when an honest party P_i outputs value v with a grade $g_i = 2$. Honest party P_i must have invoked $\mathsf{Deliver}$ for value v by round 2 and did not detect a party P_j equivocation by round 4. This implies no honest party detected a party P_j equivocation by round 3. By Lemma 3, all honest parties receive value v by round 4. In addition, since party P_i invoked $\mathsf{Deliver}$ for value v by round 2, all honest parties receive a

code word for value v by round 3. Thus, value v is the first value received by all honest parties. Since $v \neq \perp$, all honest parties will output value v with a grade ≥ 1.

Next, we consider the case when the designated sender is honest. Since, the sender is honest, it sends its input value v to all honest parties such that all honest parties receive value v in round 2. Thus, all honest parties invoke Deliver to propagate value v in round 2. Moreover, the honest sender does not equivocate. Thus, all honest parties output value v with a grade of 2 in round 4.

The case where each honest party outputs a value with a grade $\in \{0, 1, 2\}$ is trivial by design.   □

Lemma 4 (Communication Complexity)**.** *Let ℓ be the size of the input, κ be the size of accumulator, and w be the size of witness. The communication complexity of the protocol in Figure 5 is $O(n\ell + (\kappa + w)n^2)$.*

Proof. At the start of the protocol, the sender multicasts its value of size ℓ to all party P_j $\forall j \in [n]$ along with κ sized accumulator. This step incurs $O(n\ell + \kappa n)$. Invoking Deliver on an object of size ℓ incurs $O(n\ell + (\kappa + w)n^2)$, since each party multicasts a code word of size $O(\ell/n)$, a witness of size w and an accumulator of size κ. Thus, the overall communication complexity is $O(n\ell + (\kappa + w)n^2)$.   □

5 Recoverable-Set-of-Shares

In Section 3, we presented a secure DKG protocol by assuming broadcast channels. In this section, we present a slightly weaker sharing protocol by appropriately replacing the broadcast channel with multicast and our weak gradecast. This protocol completes in constant rounds and acts as a building block towards constructing the DKG. We call this protocol *Recoverable-set-of-shares*.

In the sharing phase of our secure DKG protocol with broadcast channels (Figure 2), each honest party outputs a common set QUAL consisting of size at least $n - t$ parties such that the secrets shared by parties in set QUAL can be reconstructed. In more detail, honest parties have a common decision on which parties correctly shared their secret at the end of the sharing phase. Requiring this agreement was free in the presence of broadcast channels; however, under a point-to-point network, it blows up communication complexity.

In our protocol, we utilize gradecast for secret sharing. Given that gradecast does not ensure a unanimous output, each honest party P_i may have a different view regarding the acceptance of the shared secret. Thus, each honest party P_i outputs a possibly different set AcceptList$_i$ of at least $n - t$ parties which they accept to have shared the secret correctly; i.e., P_i observes the secrets shared by parties in AcceptList$_i$ can be reconstructed. It is in this regard, we call our protocol *recoverable-set-of-shares* as the secret shared by parties in AcceptList$_i$ can be reconstructed. We stress that in recoverable-set-of-shares protocol, honest parties need not agree on a common set and may output a different set of at least $n - t$ parties which they believe have shared the secret properly. To ensure that the final keys for DKG are generated from a common set, parties need to agree on one such set. In the following sections, we present a multi-valued validated Byzantine agreement protocol to agree on a common set.

We call an AcceptList *certified* if it is accompanied by a set of signatures from at least $t + 1$ parties. The set of $t + 1$ signatures on AcceptList forms the certificate for AcceptList and denoted as $\mathcal{AC}(\text{AcceptList})$.

Definition 5 (Recoverable-set-of-shares)**.** Each party P_i, as a dealer, secret shares a uniformly random input s_i. Each honest party outputs an n element certifed list AcceptList$_i$ with an entry corresponding to each party as a dealer such that AcceptList$_i[j] \in \{0, 1, 2\}$ $\forall j \in [n]$. A recoverable-set-of-shares protocol tolerating t Byzantine failures satisfies the following properties:
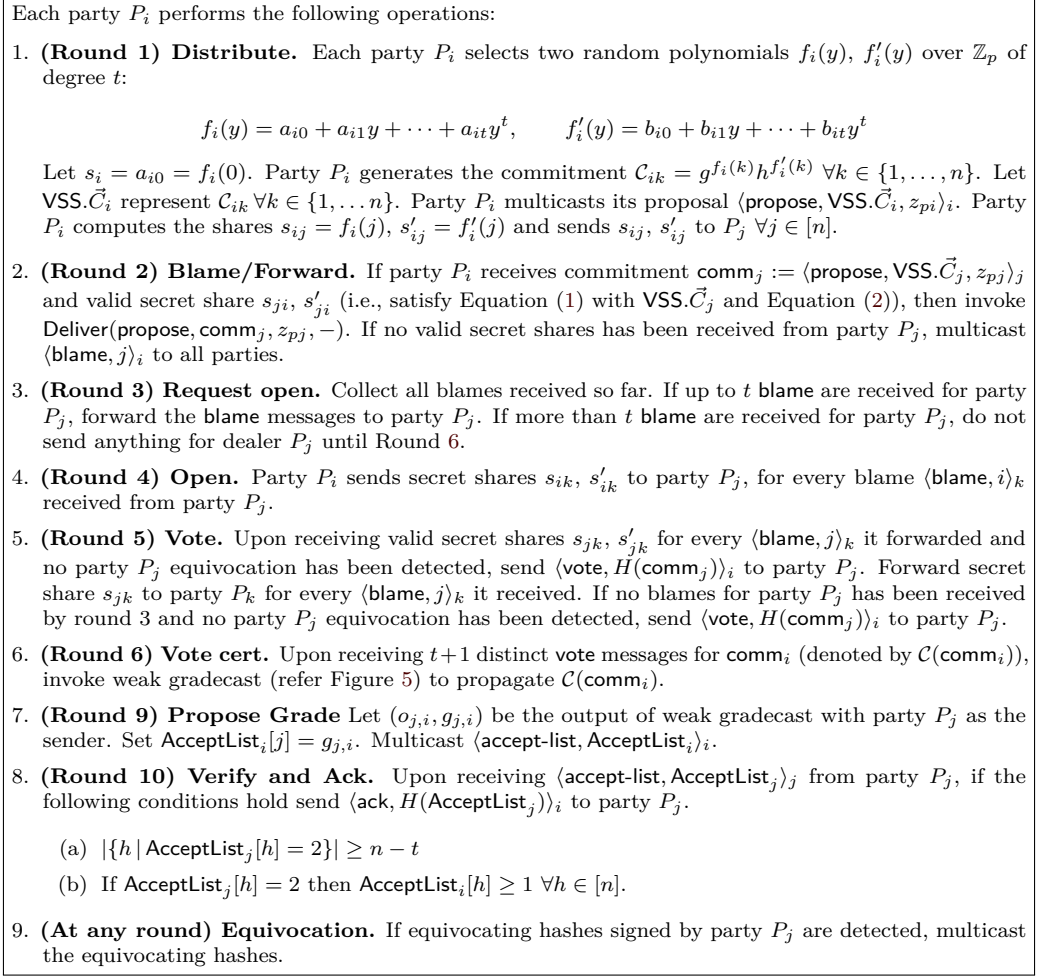

Each party P_i performs the following operations:

1. **(Round 1) Distribute.** Each party P_i selects two random polynomials $f_i(y)$, $f_i'(y)$ over $\mathbb{Z}_p$ of degree t:

$$f_i(y) = a_{i0} + a_{i1}y + \cdots + a_{it}y^t, \qquad f_i'(y) = b_{i0} + b_{i1}y + \cdots + b_{it}y^t$$

   Let $s_i = a_{i0} = f_i(0)$. Party P_i generates the commitment $\mathcal{C}_{ik} = g^{f_i(k)}h^{f_i'(k)} \ \forall k \in \{1, \ldots, n\}$. Let $\mathsf{VSS}.\vec{C}_i$ represent $\mathcal{C}_{ik} \ \forall k \in \{1, \ldots n\}$. Party P_i multicasts its proposal $\langle \mathsf{propose}, \mathsf{VSS}.\vec{C}_i, z_{pi} \rangle_i$. Party P_i computes the shares $s_{ij} = f_i(j)$, $s_{ij}' = f_i'(j)$ and sends s_{ij}, s_{ij}' to $P_j \ \forall j \in [n]$.

2. **(Round 2) Blame/Forward.** If party P_i receives commitment $\mathsf{comm}_j := \langle \mathsf{propose}, \mathsf{VSS}.\vec{C}_j, z_{pj} \rangle_j$ and valid secret share s_{ji}, s_{ji}' (i.e., satisfy Equation (1) with $\mathsf{VSS}.\vec{C}_j$ and Equation (2)), then invoke $\mathsf{Deliver}(\mathsf{propose}, \mathsf{comm}_j, z_{pj}, -)$. If no valid secret shares has been received from party P_j, multicast $\langle \mathsf{blame}, j \rangle_i$ to all parties.

3. **(Round 3) Request open.** Collect all blames received so far. If up to t blame are received for party P_j, forward the blame messages to party P_j. If more than t blame are received for party P_j, do not send anything for dealer P_j until Round 6.

4. **(Round 4) Open.** Party P_i sends secret shares s_{ik}, s_{ik}' to party P_j, for every blame $\langle \mathsf{blame}, i \rangle_k$ received from party P_j.

5. **(Round 5) Vote.** Upon receiving valid secret shares s_{jk}, s_{jk}' for every $\langle \mathsf{blame}, j \rangle_k$ it forwarded and no party P_j equivocation has been detected, send $\langle \mathsf{vote}, H(\mathsf{comm}_j) \rangle_i$ to party P_j. Forward secret share s_{jk} to party P_k for every $\langle \mathsf{blame}, j \rangle_k$ it received. If no blames for party P_j has been received by round 3 and no party P_j equivocation has been detected, send $\langle \mathsf{vote}, H(\mathsf{comm}_j) \rangle_i$ to party P_j.

6. **(Round 6) Vote cert.** Upon receiving $t+1$ distinct vote messages for comm_i (denoted by $\mathcal{C}(\mathsf{comm}_i)$), invoke weak gradecast (refer Figure 5) to propagate $\mathcal{C}(\mathsf{comm}_i)$.

7. **(Round 9) Propose Grade** Let $(o_{j,i}, g_{j,i})$ be the output of weak gradecast with party P_j as the sender. Set $\mathsf{AcceptList}_i[j] = g_{j,i}$. Multicast $\langle \mathsf{accept\text{-}list}, \mathsf{AcceptList}_i \rangle_i$.

8. **(Round 10) Verify and Ack.** Upon receiving $\langle \mathsf{accept\text{-}list}, \mathsf{AcceptList}_j \rangle_j$ from party P_j, if the following conditions hold send $\langle \mathsf{ack}, H(\mathsf{AcceptList}_j) \rangle_i$ to party P_j.

   (a) $|\{h \,|\, \mathsf{AcceptList}_j[h] = 2\}| \geq n - t$

   (b) If $\mathsf{AcceptList}_j[h] = 2$ then $\mathsf{AcceptList}_i[h] \geq 1 \ \forall h \in [n]$.

9. **(At any round) Equivocation.** If equivocating hashes signed by party P_j are detected, multicast the equivocating hashes.

Figure 6: **Recoverable-set-of-shares protocol**

1. If dealer P_j is honest, then each honest party P_i outputs $\mathsf{AcceptList}_i[j] = 2$.

2. A certified $\mathsf{AcceptList}_i$ must have $|\{h \,|\, \mathsf{AcceptList}_i[h] = 2\}| \geq n - t$.

3. If $\mathsf{AcceptList}_i$ is certified and $\mathsf{AcceptList}_i[j] = 2$, then secret s_j can be recovered from the secret shares s_{ji} received by each honest party P_i.

Protocol details. At the start of the protocol (refer Figure 6), each honest party P_i selects two random t degree polynomials $f_i(y) = \sum_k a_{ik}y^k$ over $\mathbb{Z}_p$ and $f_i'(y) = \sum_k b_{ik}y^k$ over $\mathbb{Z}_p$ such that $f_i(0) = s_i$ and $f_i'(0) = s_i'$. Party P_i generates the commitment $\mathcal{C}_{ik} = g^{f_i(k)}h^{f_i'(k)} \ \forall k \in \{1, \ldots, n\}$. Let $\mathsf{VSS}.\vec{C}_i$ represent $\mathcal{C}_{ik} \ \forall k \in \{1, \ldots n.\}$. Party P_i multicasts the commitment in the form of a proposal $\langle \mathsf{propose}, \mathsf{VSS}.\vec{C}_i, z_{pi} \rangle_i$ where z_{pi} is the accumulation value of $\mathsf{VSS}.\vec{C}_i$. In order to facilitate efficient equivocation checks, party P_i signs $\langle \mathsf{propose}, H(\mathsf{VSS}.\vec{C}_i), z_{pi} \rangle$ separately and sends $\mathsf{VSS}.\vec{C}_i$ separately. Party P_i also privately sends secret share s_{ij}, s_{ij}' to party $P_j \ \forall j \in [n]$.

If a party P_j receives valid secret share s_{ij}, s_{ij}' along with the proposal $\mathsf{comm}_i := \langle \mathsf{propose}, \mathsf{VSS}.\vec{C}_i, z_{pi} \rangle_i$ by the start of round 2, it invokes $\mathsf{Deliver}(\mathsf{propose}, \mathsf{comm}_i, z_{pi}, -)$ to propagate the commitment $\mathsf{VSS}.\vec{C}_i$; otherwise party P_j multicasts $\langle \mathsf{blame}, i \rangle_j$. Observe that we ignore the epoch e parameter in $\mathsf{Deliver}$ as the current protocol is not executed in an epoch.

Party P_j waits to collect any blame messages sent by other parties. If up to t blame messages are received for P_i, P_j forwards the blame messages to party P_i. Party P_i then privately sends secret shares s_{ik}, s'_{ik} to party P_j, for every blame $\langle \mathsf{blame}, i \rangle_k$ received from party P_j. Upon receiving valid secret shares for all $\langle \mathsf{blame}, i \rangle_k$ it forwarded, party P_j sends a vote $\langle \mathsf{vote}, H(\mathsf{comm}_i) \rangle$ to party P_i and also forwards secret shares s_{ik}, s'_{ik} to party P_k if no party P_i has been detected by round 5. Additionally, if no blame messages are received for P_i by round 3, party P_j sends $\langle \mathsf{vote}, H(\mathsf{comm}_i) \rangle$ to party P_i at round 5.

Party P_i then waits to collect $t + 1$ vote messages for $H(\mathsf{comm}_i)$, denoted by $\mathcal{C}(\mathsf{comm}_i)$. A certificate on the comm_i implies that secret s_i shared by party P_i can be reconstructed later. Party P_i then gradecasts $\mathcal{C}(\mathsf{comm}_i)$. Invocation of gradecast on $\mathcal{C}(\mathsf{comm}_i)$ ensures that if the party P_i is honest, all honest parties output a common $\mathcal{C}(\mathsf{comm}_i)$ with a grade of 2 and if an honest party P_k output $\mathcal{C}(\mathsf{comm}_i)$ with a grade of 2, all other honest parties output the certificate with a grade ≥ 1.

Note that all parties (at least all honest parties) are executing the secret sharing phase. Thus, at the end of gradecast step, each honest party outputs at least $n - t$ certificates with a grade of 2 and outputs at most t values with a grade ≤ 2. We call the list of grades for party P_j as $\mathsf{AcceptList}_j$. This list is a set of parties which party P_j observes to have shared their secret properly and each secret can be reconstructed. Party P_j then multicasts its $\mathsf{AcceptList}_j$ to all other parties. Party P_k then checks the validity of $\mathsf{AcceptList}_j$ by checking if (i) $|\{h \mid \mathsf{AcceptList}_j[h] = 2\}| \geq n - t$, and (ii) if $\mathsf{AcceptList}_j[h] = 2$ then $\mathsf{AcceptList}_k[h] \geq 1 \ \forall h \in [n]$. The first check ensures that $\mathsf{AcceptList}_j$ contains at least $n - t$ entries with $\mathsf{AcceptList}_j[h] = 2$. This check trivially satisfies for $\mathsf{AcceptList}$ sent by an honest party as each honest party receives at least $n - t$ certificates with a grade of 2. Later, the DKG protocols use secrets from parties in $\mathsf{AcceptList}_j$ such that $\mathsf{AcceptList}_j[h] = 2$ to compute the final keys. This is required to ensure security of DKG protocol. The second check ensures that all the secrets corresponding to $\mathsf{AcceptList}_j[h] = 2$ are recoverable; observe that if $\mathsf{AcceptList}_j[h] = 2$ then $\mathsf{AcceptList}_k[h] \geq 1$ due to weak gradecast properties. This implies party P_k has received a $\mathcal{C}(\mathsf{comm}_h)$ from party P_h and $\mathcal{C}(\mathsf{comm}_h)$ implies the secret shared by party P_h can be reconstructed. If the checks pass, party P_k sends $\langle \mathsf{ack}, H(\mathsf{AcceptList}_j) \rangle_k$ to party P_j. A set of $t + 1$ ack (ack-cert) messages for $\mathsf{AcceptList}_j$ (denoted by $\mathcal{AC}(\mathsf{AcceptList}_j)$) implies at least one honest party has verified that all the secrets corresponding to $\mathsf{AcceptList}_j[h] = 2$ can be recovered.

The idea of using gradecast to perform secret sharing has been explored before in the works of Feldman and Micali [FM88, FM97] to generate common source of randomness. Compared to their work, our protocols work in authenticated model with $t < n/2$ resilience and invoke a single gradecast per secret sharing. Their protocols work in *unauthenticated* model without PKI with $t < n/4$ [FM88] and $t < n/3$ [FM97] resilience and involved multiple invocation of gradecast per secret sharing.

5.1   Analysis of Recoverable-set-of-shares protocol

Lemma 5. *If an honest party sends* vote *for a commitment* comm*, then (i) all honest parties receive* comm*, (ii) all honest parties receive their valid secret shares corresponding to commitment* comm*.*

Proof. Suppose an honest party P_i sends a vote for commitment $\mathsf{comm}_k := \langle \mathsf{propose}, \mathsf{VSS}.\vec{C}_k, z_{pk} \rangle_k$ at round 5. Party P_i must have received up to t blame messages for party P_k. This implies at least one honest party P_j received valid secret shares $s_{k,j}$, $s'_{k,j}$ and commitment comm_k and invoked $\mathsf{Deliver}(\mathsf{propose}, \mathsf{comm}_k, z_{pk}, -)$ at round 2. Moreover, party P_i did not detect party P_k equivocation by round 5. This implies no honest party detected party P_k equivocation by round 3. By Lemma 3, all honest parties receive the commitment comm_k by round 4. This proves part (i) of the Lemma.

For part (ii), party P_i can send vote message on two occasions: (a) when it does not

detect a $\langle\mathsf{blame}, k\rangle$ by round 3 and party k equivocation by round 5, and (b) when party k sent valid secret shares for every $\langle\mathsf{blame}, k\rangle$ message it forwarded and does not detect any party k equivocation by round 5.

In case (a), party P_i did not detect a party k equivocation by round 5 and $\langle\mathsf{blame}, k\rangle$ by round 3. Observe that all honest parties must have received valid secret shares corresponding to the commitment comm_k; otherwise party P_i must have received $\langle\mathsf{blame}, k\rangle$ by round 3 (since honest parties send $\langle\mathsf{blame}, k\rangle$ if no valid secret shares are received at round 2). Thus, all honest parties receive valid secret shares corresponding to commitment comm_k.

In case (b), party P_i receives valid secret shares from party P_k for every $\langle\mathsf{blame}, k\rangle$ (up to t blame) messages it forwarded and detected no party k equivocation by round 5. Observe that party P_i received $f \leq t$ $\langle\mathsf{blame}, k\rangle$ messages and received valid secret shares for every $\langle\mathsf{blame}, k\rangle$ message it forwarded. This implies at least $n - t - f$ honest parties have received valid shares for commitment comm_k from party P_k; otherwise, party P_i would have received more than f $\langle\mathsf{blame}, k\rangle$ message by round 3. Since, party P_i forwards f received secret shares corresponding to f received $\langle\mathsf{blame}, k\rangle$, all honest parties receive valid secret shares corresponding to commitment comm_k. $\qquad\square$

Lemma 6. *If an honest party sends an* ack *for a grade list* $\mathsf{AcceptList}_j$, *then all honest parties have valid secret shares corresponding to* comm_h *for all* h *such that* $\mathsf{AcceptList}_j[h] = 2$.

Proof. Suppose an honest party P_i sends an ack for a grade list $\mathsf{AcceptList}_j$. Then, it must be that if $\mathsf{AcceptList}_j[h] = 2$ then $\mathsf{AcceptList}_i[h] \geq 1$ $\forall h \in [n]$. Party P_i sets $\mathsf{AcceptList}_i[h] \geq 1$ when it receives a vote certificate $\mathcal{C}(\mathsf{comm}_h)$. If there is a vote certificate $\mathcal{C}(\mathsf{comm}_h)$ for value comm_h, then at least one honest party (say party P_k) must have voted for comm_h. By Lemma 5 part (ii), all honest parties have valid secret shares corresponding to commitment comm_h. Thus, all honest parties have valid secret shares corresponding to comm_h for all h such that $\mathsf{AcceptList}_j[h] = 2$. $\qquad\square$

Lemma 7 (Liveness)**.** *Each honest party P_i will receive an* ack-cert *for its grade list* $\mathsf{AcceptList}_i$.

Proof. Consider an honest party P_i. Party P_i will send valid commitment $\mathsf{VSS}.\vec{C}_i$ and secret shares s_{ij}, s'_{ij} to party P_j $\forall j \in [n]$ in round 1. All honest parties will receive their valid secret shares s_{ij}, s'_{ij} and commitment comm_i in round 2. Thus, no honest party will send $\langle\mathsf{blame}, i\rangle$ for party P_i.

Observe that up to t Byzantine parties can always send $\langle\mathsf{blame}, i\rangle$. Honest parties wait until round 3 to collect blame messages for any party. Honest parties forward $\langle\mathsf{blame}, i\rangle$ to party P_i which party P_i receives by round 4. Party P_i forwards valid secret shares to party P_j for every $\langle\mathsf{blame}, i\rangle$ message it received from party P_j which party P_j receives by round 5. Thus, party P_j will send vote for party P_i which party P_i receives by round 6. This implies party P_i collects $t + 1$ distinct vote messages by round 6.

Party P_i invokes weak gradecast to propagate $\mathcal{C}(comm_i)$ which completes by round 9. Due to the properties of weak gradecast, for an honest party P_i, all honest parties set $\mathsf{AcceptList}[i]$ to 2. Thus, for any honest party P_j, all honest parties set $\mathsf{AcceptList}[j]$ to 2. This implies all honest parties will have $|\{h \,|\, \mathsf{AcceptList}_j[h] = 2\}| \geq n - t$.

Next, we consider the case when an honest party sets $\mathsf{AcceptList}_i[l] = 2$ for a Byzantine party P_l and receive $\mathcal{C}(\mathsf{comm}_l)$. Due to the properties of weak gradecast, all honest parties receive $\mathcal{C}(\mathsf{comm}_l)$ and set $\mathsf{AcceptList}[l] \geq 1$. Thus, for every $\mathsf{AcceptList}_i[h] = 2$ then $\mathsf{AcceptList}[h] \geq 1$ for all honest parties.

Party P_i multicasts its $\mathsf{AcceptList}_i$ in round 9. Since, $\mathsf{AcceptList}_i$ satisfies both the conditions $|\{h \,|\, \mathsf{AcceptList}_i[h] = 2\}| \geq n - t$ and $\mathsf{AcceptList}_i[h] = 2$ then $\mathsf{AcceptList}[h] \geq 1$,

all honest parties will send ack for $\mathsf{AcceptList}_i$ proposed by party P_i and party P_i will receive ack-cert for $\mathsf{AcceptList}_i$ the end of round 10.

$\square$

Theorem 10. *The protocol in Figure 6 is a recoverable-set-of-shares protocol satisfying Definition 5.*

Proof. Straight forward from Lemma 5, Lemma 6 and Lemma 7 $\square$

Lemma 8 (Communication Complexity). *Let ℓ be the size of commitment* comm, *κ be the size of secret share and accumulator, and w be the size of witness. The communication complexity of the protocol is $O(n^2\ell + (\kappa + w)n^3)$ bits per epoch.*

Proof. At the start of the protocol, each party P_i multicasts comm_i of size ℓ to all party P_j $\forall j \in [n]$ and sends secret share $s_{i,j}$ to party P_j $\forall j \in [n]$. This step incurs $O(n^2\ell + \kappa n^3)$. In the Forward step, parties invoke Deliver for the first comm_j from party P_j for $j \in [n]$. Invoking Deliver on an object of size ℓ incurs $O(n\ell + (\kappa + w)n^2)$, since each party multicasts a code word of size $O(\ell/n)$, a witness of size w and an accumulator of size κ. Thus, invoking Deliver on n commitments incurs $O(n^2\ell + (\kappa + w)n^3)$.

In the Blame step, honest parties may blame up to t Byzantine parties if they do not receive valid secret shares. Multicast of t blame from each party incurs $O(\kappa t n^2)$ communication. In addition, t Byzantine parties always can blame honest parties. Honest parties forward up to t $\langle \mathsf{blame}, j \rangle$ messages to party P_j. This incurs $O(\kappa t n^2)$ communication.

In the Private open step each party can send up to t secret shares to all other parties. This incurs $O(\kappa t n^2)$ for all parties. In the Vote cert step, each party multicasts $O(n)$-sized vote-cert to all other parties which incurs $O(\kappa n^3)$ in communication. Invoking Deliver on an $O(n)$-sized certificate incurs $O(n^2 + (\kappa + w)n^2)$. For n certificate, this step incurs $O(n^3 + (\kappa + w)n^3)$.

In the Propose grade step, each party multicast their grade list of size $O(n)$. Multicast of $O(n)$-sized grade list by n parties incurs $O(n^3)$ communication. Thus, the total communication complexity is $O(n^2\ell + (\kappa + w)n^3)$ bits. $\square$

6 Oblivious Leader Election

In this section, we construct a communication efficient oblivious leader election (OLE) protocol that outputs a common honest leader with some constant probability called the *fairness*. Our OLE protocol uses only n parallel invocations of weaker VSS primitives and a non-interactive threshold signature scheme [CKS00]. Importantly, our OLE protocol does not require a prior threshold (DKG) setup phase despite making use of threshold signatures. The security of our OLE protocol is based on the computational Diffie-Hellman (CDH) problem in the random oracle model. The resulting protocol incurs a communication complexity of $O(\kappa n^3)$ and constant rounds.

Construction. The starting point of our construction is the threshold coin-tossing scheme of Cachin et al. [CKS00] which makes use of non-interactive threshold signature scheme. The threshold signature scheme requires a prior threshold setup which is essentially a DKG. The threshold setup establishes a tuple $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ of secret keys, a tuple $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ of verification keys. After the threshold setup phase, each party signs a common message (e.g., an epoch number) with its threshold secret key to obtain a threshold share. A combination of any $t + 1$ valid threshold shares is then used to obtain a unique and random threshold signature σ. A random oracle $H'' : \mathbb{G} \to \{0,1\}$ is then used to generate an unbiased and unpredictable random bit from the threshold signature σ.

Note that the threshold signature scheme requires a prior threshold setup to establish a tuple $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ of secret keys, a tuple $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ of verification keys. We fulfill

---

Each party P_i performs the following operations:

1. **Round 1:** Each party P_i invokes recoverable-set-of-shares protocol (refer Figure 6). Each party P_i outputs (AcceptList2$_i$ $\mathcal{AC}$(AcceptList2$_i$)).

2. **Round 11:** Each party P_i invokes weak gradecast to propagate (AcceptList2$_i$, $\mathcal{AC}$(AcceptList2$_i$)).

3. **Round 15:** Let $(o_{j,i}, \mathsf{grade}_i[j])$ be the output with party P_j as dealer. Let $o_{j,i}$ contains AcceptList2$_j$. If $\mathsf{grade}_i[j] \geq 1$, set local-dkg$_i[j]$ = AcceptList2$_j$, local-dkg-grade$_i[j]$ = $\mathsf{grade}_i[j]$.

   - Set $\mathsf{sk}_{j,i} = \sum_{m \in \mathsf{AcceptList2}_j | \mathsf{AcceptList2}_j[m]=2} s_{mi}$, $\mathsf{vk}_{j,i} = g^{\mathsf{sk}_{j,i}}$, and $\mathsf{sk}'_{j,i} = \sum_{m \in \mathsf{AcceptList2}_j | \mathsf{AcceptList2}_j[m]=2} s'_{mi}$.

   - Compute $\mathcal{C}_{\langle g \rangle}(\mathsf{sk}_{j,i})$, $\mathcal{C}_{\langle g,h \rangle}(\mathsf{sk}_{j,i}, \mathsf{sk}'_{j,i}))$ and $\pi_{\equiv \mathsf{Com}\, j,i} = $ NIZKPK$_{\equiv \mathsf{Com}}(\mathsf{sk}_{j,i}, \mathsf{sk}'_{j,i}, g, h, \mathcal{C}_{\langle g \rangle}(\mathsf{sk}_{j,i}), \mathcal{C}_{\langle g,h \rangle}(\mathsf{sk}_{j,i}, \mathsf{sk}'_{j,i}))$. Multicast $(\mathsf{vk}_{j,i}, \pi_{\equiv \mathsf{Com}\, j,i})$ to all parties.

Figure 7: Threshold setup protocol

this requirement by using the output of recoverable-set-of-shares protocol (from Section 5) to establish a local threshold setup corresponding to each party. In the recoverable-set-of-shares protocol, each party P_i outputs an AcceptList$_i$ along with $\mathcal{AC}$(AcceptList$_i$). An AcceptList$_i$ (accompanied by $\mathcal{AC}$(AcceptList$_i$)) consists of at least $n - t$ entries with grades of 2 and all honest parties must have received secret shares shared by parties in AcceptList$_i$ whose grades are 2. Thus, each party P_j uses secret shares shared by parties in an AcceptList$_i$ with grades of 2 to compute its secret key $\mathsf{sk}_{i,j}$ and verification key $\mathsf{vk}_{i,j} = g^{\mathsf{sk}_{i,j}}$ to establish local DKG setup local-dkg$[i]$ corresponding to party P_i.

In order to setup local-dkg$[i]$, party P_i first invokes weak gradecast to propagate its AcceptList$_i$ (along with $\mathcal{AC}$(AcceptList$_i$)). If party P_j outputs (AcceptList$_i$, $\mathcal{AC}$(AcceptList$_i$)) with a grade of ≥ 1, it uses AcceptList$_i$ to compute its secret key $\mathsf{sk}_{i,j}$ and verification key $\mathsf{vk}_{i,j} = g^{\mathsf{sk}_{i,j}}$ to establish local DKG setup local-dkg$[i]$ corresponding to party P_i. Note that this establishes a separate threshold setup for each party P_i. With local DKG setup local-dkg$[i]$ as the threshold setup for party P_i, parties then sign a common message to generate a unique and random threshold signature σ_i. Parties then use a random oracle $H' : \mathbb{G} \to \{0,1\}^\kappa$ to generate κ bit random coin value assigned to party P_i. Each party P_j uses the random coin value assigned to party P_i if it outputs (AcceptList$_i$, $\mathcal{AC}$(AcceptList$_i$)) with a grade of 2. From the set of parties for which party P_j outputs (AcceptList$_k$, $\mathcal{AC}$(AcceptList$_k$)) with a grade of 2, it selects the party with highest (or lowest) coin value as its leader. With probability at least $\frac{1}{2}$, all honest parties select a common honest leader using this approach.

Note that threshold coin-tossing scheme of Cachin et al. [CKS00] produces a single bit output. However, it can also be used to generate κ bit strings using κ-bit hash function [CKPS01]. Looking ahead, the final DKG is also computed from one of the valid AcceptList output from the recoverable-set-of-shares protocol. Making use of the secret shares in an AcceptList output from the recoverable-set-of-shares protocol during this local DKG setup phase will leak the final public key before the final DKG is decided. Note that the final public key can be computed from $t + 1$ verification keys. This allows the adversary ability to force the final DKG to have certain final public key. To circumvent this issue, we execute two separate instances of recoverable-set-of-shares protocol in parallel; one instance to setup local DKG instances and the other to setup the final DKG instance. To remove this ambiguity, we call the accept list output from the recoverable-set-of-shares protocol executed for local DKG as AcceptList2 i.e. each party P_i outputs an AcceptList2$_i$ along with $\mathcal{AC}$(AcceptList2$_i$).

Protocol details. The setup phase of the protocol is presented in Figure 7. Each party P_i invokes recoverable-set-of-shares protocol and outputs AcceptList2$_i$ (along with $\mathcal{AC}$(AcceptList2$_i$)). Each party P_i then invokes weak gradecast to propagate (AcceptList2$_i$, $\mathcal{AC}$(AcceptList2$_i$)). At the end of the setup phase, each party P_i sets up the local DKG

instance for each party P_j (i.e., local-dkg$_i[j]$) as AcceptList2$_j$ if local-dkg-grade$_i[j] \geq 1$. If local-dkg-grade$_i[j] = 2$, due to weak gradecast properties, all honest parties have a common local DKG instance for party P_j (i.e., local-dkg$[j]$). In addition, for an honest party P_j, all honest parties will have a common local DKG instance local-dkg$[j]$. Each party P_i also computes required secret keys sk$_{j,i}$, verification keys vk$_{j,i}$ for local DKG instance local-dkg$_i[j]$ computed from local-dkg$_i[j]$ as shown in Figure 7.

Let sid be the input of party P_i.
Set $\mathcal{X}_i \leftarrow \emptyset$. Each party P_i performs following operations:

1. Perform $\sigma_{j,i} = \mathsf{Sign}_{\mathsf{TS}}(\mathsf{sk}_{j,i}, (j, \mathsf{sid}))$ and multicast $\sigma_{j,i}$ if local-dkg-grade$_i[j] \geq 1 \; \forall j \in [n]$.

2. Upon receiving a set S of $t+1$ valid signature shares for party P_j, compute $\sigma_j = \mathsf{Combine}_{\mathsf{TS}}(\mathsf{pk}, \mathsf{sid}, S)$ and $\mathcal{X}_i[j] \leftarrow H'(\sigma_j)$.

3. Perform $\ell \leftarrow argmax_h\{\mathcal{X}_i[h] \| \text{local-dkg-grade}_i[h] = 2\}$. Output P_ℓ.

Figure 8: Oblivious Leader Election

The OLE protocol is presented in Figure 8. The input to the protocol is a sequence id sid. Once the local DKG instances are setup, each party P_i uses its secret key sk$_{j,i}$ to sign a common message i.e., (j, sid) (for party P_j) if local-dkg-grade$_i[j] \geq 1$ to obtain a threshold share $\sigma_{j,i}$. A set of $t+1$ valid signature shares corresponding to local-dkg$[j]$ is combined to form a single threshold signature σ_j and a hash $H'(\sigma_j)$ generates κ bit coin value for party P_j. We note that two or more parties could output the same grade list (i.e, AcceptList2) in the recoverable-set-of-shares protocol; hence their local DKG might be same. However, parties sign a distinct message e.g. (j, sid) for party P_j. Such generated threshold signatures are unique and random regardless of their local DKG instance being common; hence the coin value assigned to each party is also random. Honest parties consider coin values for party P_j only if local-dkg-grade$_i[j] = 2$. Note that if local-dkg-grade$_i[j] = 2$, a threshold signature σ_j will exist for party P_j. This is because all honest parties will have local-dkg-grade$[j] \geq 1$ and a common local-dkg$[j]$ due to weak gradecast properties and each honest party P_i will send their signature share $\sigma_{j,i}$. A coin value is then computed as $H'(\sigma_j)$. The party P_ℓ with highest coin value is elected as leader.

Round complexity and communication complexity. The threshold setup phase has a latency of 15 rounds to invoke recoverable-set-of-shares protocol, n parallel instances of weak-gradecast and distribute verification keys. The OLE protocol requires only 1 round to generate threshold signatures. The threshold setup phase invokes recoverable-set-of-shares protocol, n parallel weak-gradecasts with an input of size $O(\kappa n)$ and sharing verification keys. This incurs $O(\kappa n^3)$ communication. The threshold signature generation incurs $O(\kappa n^3)$ communication.

Remark. While OLE protocols without threshold setup have been formulated for the asynchronous model in the literature [AJM+21, GLL+22, AJM+23] with a communication complexity of $O(\kappa n^3)$, these protocols require stronger cryptographic assumptions such as SXDH and OMDL. Compared to these constructions, our OLE protocol tolerates $t < n/2$ Byzantine faults and relies on CDH assumption. It is an interesting direction to explore asynchronous OLE protocols with lesser cryptographic assumptions. Such an OLE protocol can be used to obtain threshold setup-free asynchronous consensus protocols with lesser cryptographic assumptions.

6.1   Analysis of OLE protocol

Our coin generation protocol is similar to the threshold coin-tossing scheme of [CKS00]. In Cachin et al. [CKS00], the coin value is a single bit computed from the threshold signature using $H'' : \mathbb{G} \rightarrow \{0, 1\}$. In our scheme, the coin value is a κ bit string computed from the threshold signature using a κ bit hash function $H' : \mathbb{G} \rightarrow \{0, 1\}^\kappa$ which is also secure [CKPS01]. We rely on the following Lemma of [CKS00].

Lemma 9 ([CKS00])**.** *In the random oracle model, the coin-tossing scheme of Cachin et al. [CKS00] is secure i.e., satisfies robustness and unpredictability under CDH assumption.*

Theorem 11. *Assuming public-key infrastructure, random oracle and CDH assumption, the protocol in Figure 8 is an oblivious leader election protocol with fairness at least $\frac{1}{2}$.*

Proof. We first show termination i.e., honest party P_i will obtain a threshold signature σ_j (and coin value for party P_j) if local-dkg-grade$_i[j] = 2$. This is because all honest parties will have local-dkg-grade$[j] \geq 1$ and a common local-dkg$[j]$ due to weak gradecast properties. Thus, each honest party P_k will send their signature share $\sigma_{j,k}$ i.e., a set of $t + 1$ valid signature shares will be available sufficient to obtain threshold signature σ_j (and coin value $H'(\sigma_j)$).

By Lemma 9 the threshold signature generation protocol satisfies robustness and unpredictability. Thus, the coin value generated from threshold signature is robust and unpredictable.

Observe that each party P_i signs a distinct message (i.e, (j, sid)) for each part P_j. Thus, the threshold signature σ_j for each party P_j is unique and random even if two or more parties have the same local DKG instance; hence each party P_j will be assigned random coin value $(H'(\sigma_j))$. Since, the coin value assigned to a party is random, the coin value assigned to an honest party will be a global maximum with probability at least $\frac{n-t}{n}$. The probability that coin values of any two parties can be maximum is bounded by $\frac{1}{2^\kappa}$. Thus, all honest parties select the coin value corresponding to a common honest leader with probability $\frac{n-t}{n} - \frac{1}{2^\kappa} \geq \frac{1}{2}$ when $\kappa = 2 \log n$. □

7   Multi-valued Validated Byzantine Agreement

In this section, we present a synchronous MVBA protocol tolerating $t < n/2$ Byzantine faults with $O(n^2\ell + \kappa n^3)$ communication for inputs of size ℓ bits and expected constant rounds. We extend the Binary Byzantine agreement (BBA) protocol of Katz and Koo [KK06] to obtain MVBA for large ($\ell = \Theta(n)$) input. The BBA protocol of Katz and Koo [KK06] tolerates $t < n/2$ Byzantine faults and terminates in expected 4 epochs. Their protocol involves invoking n parallel gradecasts; with each gradecast propagating small sized input. As mentioned before, their gradecast protocol incurs $O(\kappa n^3)$ communication for a single bit input; thus, their protocol trivially incurs $O(\kappa n^4)$ communication. We replace their gradecast protocol with our communication optimal gradecast protocol from Section 4. Our gradecast protocol incurs only $O(n\ell + \kappa n^2)$ communication while propagating ℓ-bit input. Using our gradecast protocol allows BBA protocol of Katz and Koo [KZG10] to handle large input while simultaneously reducing the communication to $O(n^2\ell + \kappa n^3)$.

To circumvent the linear round lower bound for a deterministic BA protocol [DS83], BA protocols use a common source of randomness called *common coin* to achieve agreement in constant expected rounds. The common coin is *weak* if all honest parties obtain a common honest leader with some constant probability (and with the remaining probability either the common leader is Byzantine or honest parties may disagree on the leader). In Katz and Koo BBA, the weak common coin was obtained by invoking n^2 moderated VSS instances which incurs $\Omega(\kappa n^4)$ communication and blows up the communication complexity. In this work, we replace their weak common coin protocol with our communication efficient leader election protocol from Section 6 which outputs a common honest leader with probability at least $\frac{1}{2}$. Our OLE protocol incurs $O(\kappa n^3)$ communication and a single round after an initial setup phase (refer Figure 7) which incurs 15 rounds.

Protocol details. Our MVBA protocol in presented in Figure 9. The underlying consensus mechanism is identical to the BBA protocol of Katz and Koo [KK06]. The protocol progresses through a sequence of numbered *epochs*, with all parties starting in

Let v_i be party P_i's input and e be the current epoch. Each party P_i sets $\mathsf{lock}_i \leftarrow \bot$. Each party P_i performs following operations.

1. **(Round 1) Propose.** Each party P_i invokes weak gradecast to propagate v_i.

2. **(Round 4) Update.** Let $(v_{j,i}, \mathsf{grade}_i[j])$ be the output with party P_j as the dealer. Let $\mathcal{S}_i^v := \{j : v_{j,i} = v \wedge \mathsf{grade}_i[j] = 2\}$ and $\tilde{\mathcal{S}}_i^v := \{j : v_{j,i} = v \wedge \mathsf{grade}_i[j] \geq 1\}$. If $\mathsf{lock}_i = \bot$, then:

   (a) If $|\tilde{\mathcal{S}}_i^v| > t$, update $v_i \leftarrow v$.

   (b) If $|\mathcal{S}_i^v| > t$, set $\mathsf{lock}_i \leftarrow 1$.

   Invoke weak gradecast (refer Figure 5) to propagate v_i.

3. **(Round 7) Update2.** Again, let $(v_{j,i}, \mathsf{grade}_i[j])$ be the output with party P_j as the dealer. Define $\mathcal{S}_i^v$ and $\tilde{\mathcal{S}}_i^v$ as above. If $\mathsf{lock}_i = \bot$ and $|\tilde{\mathcal{S}}_i^v| > t$, set $v_i \leftarrow v$. Multicast v_i.

4. **(Round 8) Leader election.** Invoke OLE protocol with input e.

5. **(Round 9) Terminate/Advance Epoch.** Let P_ℓ be the output of leader election protocol.

   (a) If $\mathsf{lock}_i = 0$, output v_i and terminate.

   (b) If $\mathsf{lock}_i = 1$, set $\mathsf{lock}_i = 0$. If $\mathsf{lock}_i = \bot$ and $|\mathcal{S}_i^v| \leq t$, $v_{\ell,i} \neq \bot$ and $\mathsf{ex\text{-}validation}(v_{\ell,i}) = \mathsf{true}$, update $v_i \leftarrow v_{\ell,i}$. Advance to epoch $e + 1$.

6. **(At any round) Equivocation.** If equivocating hashes signed by party P_j are detected, multicast the equivocating hashes.

Figure 9: **MVBA** with $O(n^2\ell + \kappa n^3)$ communication and expected 4 epochs.

epoch 1 and progressing to higher epochs as the protocol continues. At the start of the protocol execution, each party sets its lock_i to $\bot$. We explain the protocol steps for an epoch e.

In round 1, each party P_i invokes weak gradecast protocol to propagate its input v_i. Our weak gradecast protocol incurs 4 rounds. Rounds 2 and 3 accommodates the steps of the weak gradecast protocol. In round 4, each party P_i outputs $(v_{j,i}, \mathsf{grade}_i[j])$ for the weak gradecast corresponding to party P_j as the dealer. If party P_i observes a common value v with a grade of at least 1 from at least $t + 1$ parties (i.e., $\tilde{\mathcal{S}}_i^v \geq t$), party P_i adopts value v (i.e., sets its input v_i to value v). In addition, if P_i observes a common value v with a grade of 2 from at least $t + 1$ parties (i.e., $\mathcal{S}_i^v \geq t$), it updates lock_i to 1 (i.e., locks to value v). In this case, due to weak gradecast properties, all other honest parties will have $\tilde{\mathcal{S}}_i^v \geq t$ and update their inputs to value v.

Again in round 4, each party P_i invokes weak gradecast protocol to propagate its updated input v_i. Rounds 5 and 6 accommodates the steps of the weak gradecast protocol. In round 7, each party P_i outputs $(v_{j,i}, \mathsf{grade}_i[j])$ for the weak gradecast corresponding with party P_j as the dealer. $\mathcal{S}_i^v$ and $\tilde{\mathcal{S}}_i^v$ are defined similar to round 4. In round 7, P_i updates its input v_i to v if $\mathsf{lock}_i = \bot$ and $\tilde{\mathcal{S}}_i^v \geq t$ (we explain the intuition behind this update step later in round 9). At the end of round 7, each party P_i multicasts its input v_i.

In round 8, parties invoke the OLE protocol to elect a leader. The OLE protocol outputs a common honest leader with probability at least $\frac{1}{2}$. Let P_ℓ be the output of leader election protocol for party P_i. In round 9, if $\mathsf{lock}_i = 0$, party P_i outputs v_i and terminates the protocol. On the other hand, if $\mathsf{lock}_i = 0$, party P_i sets $\mathsf{lock}_i = 1$ and continues to the next epoch. This ensures party P_i participates in the next epoch $e + 1$ and all honest parties receive input from P_i in the next epoch $e + 1$. If an honest party P_i updates its lock_i to 1 in epoch e, our protocol ensures that all honest parties terminate the protocol in epoch $e + 2$.

If $\mathsf{lock}_i = \bot$ and $\mathcal{S}_i^v \leq t$, P_i updates v_i to the externally valid $v_{\ell,i}$ when $v_{\ell,i} \neq \bot$. When the OLE outputs different leaders or when the common elected leader is Byzantine, different honest parties can update their input to different values and continue to epoch $e + 1$. However, when the OLE outputs a common honest leader P_ℓ (which happens with probability at least $\frac{1}{2}$), all honest parties will adopt a common value which is the value

sent by P_ℓ in round 7. This property holds even when some honest party adopts the leader P_ℓ's value and while some honest party do not adopt the leader P_ℓ's value. Observe that an honest party P_j does not adopt the leader's value when $\mathcal{S}_j^v \geq t$. Due to weak gradecast properties, all other honest parties must have $\tilde{\mathcal{S}}_i^v \geq t$ in round 7. Thus, all honest parties (including the common and honest elected leader) will adopt value v in round 7 and multicast value v. Thus, each honest party with $\mathcal{S}_j^v \leq t$ will pick value v and advance to epoch $e + 1$. This ensures all honest parties will update their input to a common value v and advance to epoch $e + 1$ when the elected leader is common and honest.

Round complexity. By Theorem 11, a common honest leader is selected with probability at least $\frac{1}{2}$ and all honest parties terminate in the next 2 epochs. Thus, the expected number of epochs required is 4 epochs.

Remark on quality property. By Theorem 11, a common honest leader is selected with probability at least $\frac{1}{2}$. When a common honest leader is elected, all honest parties decide on the leader's proposed value. Thus, our MVBA protocol also satisfies the quality property with probability at least $\frac{1}{2}$.

7.1   Analysis of MVBA

Lemma 10. *If an honest party sets* lock *to 1 with a value v in epoch e, then all honest parties adopt value v in epoch e.*

Proof. Suppose an honest party P_i sets lock_i to 1 in epoch e. Party P_i must have received value v from a set Q of at least $t + 1$ parties such that $|\mathcal{S}_i^v| > t$. By the properties of weak gradecast, all other honest parties receive value v corresponding to parties in Q with a grade ≥ 1 (i.e., all other honest parties have $\mathsf{grade}[j] \geq 1 \ \forall j \in Q$) and $|\tilde{\mathcal{S}}^v| > t$ for all other honest parties and all honest parties adopt value v in the Update step.

Once all honest parties adopt value v in the Update step, they invoke weak-gradecast to propagate value v at the end of the Update step. Since, honest parties do not equivocate and send value v in a timely manner, all honest parties output value v such that $\mathsf{grade}[j]$ to 2. Thus, $|\tilde{\mathcal{S}}_i^v| > t$ and $|\mathcal{S}_i^v| > t$ in the Update2 step. Since, $|\mathcal{S}_i^v| > t$, no honest party will adopt value v_ℓ selected from the proposal election protocol. Thus, all honest parties adopt value v in epoch e.   □

Lemma 11. *If all honest parties start an epoch e with same input v, then all honest parties decide value v and terminate by the end of epoch $e + 1$.*

Proof. Suppose all honest parties start an epoch e with the same input v. All honest parties invoke weak-gradecast with value v in the Propose step. By the properties of weak gradecast, for an honest dealer, all honest parties output a grade of 2. Thus, all honest parties will set $\mathsf{grade}[j] = 2$ for all other honest parties. Thus, for value v, all honest parties have $|\mathcal{S}_i^v| > t$ and $|\tilde{\mathcal{S}}_i^v| > t$ If $\mathsf{lock} = \bot$, honest parties set lock to 1.

Similarly, all honest parties invoke weak-gradecast with value v in the Update2 step. By similar argument, all honest parties will set $\mathsf{grade}[j] = 2$ for all other honest parties i.e., $|\mathcal{S}_i^v| > t$ and $|\tilde{\mathcal{S}}_i^v| > t$ for all honest parties at the of Update 2 step. Moreover, no honest party will adopt the value output from the proposal election protocol.

Honest parties with $\mathsf{lock} = 0$, output v and terminate in epoch e. All the remaining honest parties with $\mathsf{lock} = 1$, set $\mathsf{lock} = 0$ and advances to epoch $e + 1$. In the next epoch, all the remaining honest parties have $\mathsf{lock} = 1$ and will not update its value and stick to value v. At the end of epoch $e + 1$, they set their lock $\mathsf{lock} = 0$, output value v and terminate. Thus, all honest parties output v and terminate by the end of epoch $e + 1$.   □

Theorem 12. *The protocol in Figure 9 solves MVBA.*

Proof. We first consider external validity i.e., if an honest party decides a value v, then ex-validation$(v) = \text{true}$. Observe that an honest party P_i decides a value v only when its sets $\text{lock}_i = \text{true}$. An honest party sets $\text{lock}_i = \text{true}$ only when it observes $|\mathcal{S}_i^v| > t$. Thus, at least one honest party P_j must have sent value v in Propose step. Honest party P_j sends value v either when its input at the start of the protocol execution is v in which case ex-validation$(v) = \text{true}$, or when its updates its value v_j to v at the end of an epoch. In the latter case, party P_j checks if ex-validation$(v) = \text{true}$.

Next, we consider agreement. Consider an epoch e and let P_ℓ be the common leader in epoch e elected via OLE protocol. There are two cases to consider.

<u>Case I.</u> $\text{lock}_i = 1$ for at least one honest party P_i with a value v in epoch e. By Lemma 10, all honest party adopt value v in epoch e and enter epoch $e + 1$ with same value v. By Lemma 11, all honest parties output value v and terminate by epoch $e + 2$.

<u>Case II.</u> $\text{lock}_i = \perp$ for all honest parties in epoch e. If leader P_ℓ is honest, leader P_ℓ sends the same value v_ℓ to all parties. If $|\mathcal{S}_i^v| \leq t$ for all honest parties, then all honest parties adopt the value v_ℓ in epoch e. By Lemma 11, all honest parties output value v_ℓ and terminate in epoch $e + 2$.

If $|\mathcal{S}_i^v| > t$ for at least one honest party P_i in the Update2 step, by the properties of weak-gradecast, $|\tilde{\mathcal{S}}^v| > t$ for all honest parties. Thus, all honest parties including leader P_ℓ adopt value v in the Update2 step. If the leader P_ℓ is honest, it sends the same value v to all parties. Honest parties with $|\mathcal{S}_i^v| \leq t$ adopt value v_ℓ which is the same value adopted by party P_i with $|\mathcal{S}_i^v| > t$. Thus, all honest parties have value v at the end of epoch e. By Lemma 11, all honest parties output value v and terminate by epoch $e + 2$. $\square$

Lemma 12 (Communication Complexity). *Let ℓ be the size of input v for each party, κ be the size of accumulator and w be the size of witness. The communication complexity of the protocol is $O(n^2\ell + (\kappa + w)n^3)$ bits per epoch.*

Proof. At the start of the protocol, each party P_i invokes weak gradecast with $O(\ell)$-sized proposal. By Lemma 4, this step incurs $O(n^2\ell + (\kappa + w)n^3)$. Similarly, in the Update2 step, each party invokes weak gradecast with $O(\ell)$-sized proposal. By Lemma 4, this step also incurs $O(n^2\ell + (\kappa + w)n^3)$. The proposal election protocol has a communication complexity of $O(\kappa n^3)$. Thus, the total communication complexity of the protocol is $O(n^2\ell + (\kappa + w)n^3)$ bits per epoch. $\square$

8   Distributed Key Generation

Finally, we present two communication efficient DKG protocols with $O(\kappa n^3)$ communication. The first protocol incurs expected $O(\kappa n^3)$ communication and terminates in expected constant rounds while the second protocol incurs $O(\kappa n^3)$ communication in the worst case and terminates in $t + 1$ epochs. The DKG protocols in this section differs from the secure DKG protocol of Section 3 in the following ways. First, we replace the broadcast channel with weaker consensus primitives and use a single invocation of consensus instance. Second, in the secure DKG protocol, the final public key and secret keys are computed from the secret shares of all honest parties. In particular, all honest parties belong to set QUAL and the public key and secret keys are computed from parties in QUAL. In contrast, the DKG protocols in this section compute the final public key and secret keys from a common set of at least $n - t$ parties where at least $n - 2t$ parties are honest (i.e., at least one honest party when $n = 2t + 1$). This suffices to ensure construction of a secure DKG protocol.

8.1   DKG with $O(\kappa n^3)$ communication and expected $O(1)$ rounds

The DKG protocol uses recoverable-set-of-shares protocol (refer Figure 6) to perform secret sharing. The threshold setup protocol (refer Figure 7) is also executed at the start of

---

1. **Deal/Setup.** Each party P_i invokes recoverable-set-of-shares protocol (refer Figure 6). Each party P_i outputs a set $\mathsf{AcceptList}_i$ with an ack-cert for $\mathsf{AcceptList}_i$ (i.e., $\mathcal{AC}(\mathsf{AcceptList}_i)$). Each party P_i also invokes threshold setup phase (refer Figure 7) in parallel.

2. **MVBA.** Each party P_i invokes MVBA (Figure 9) with input $(\mathsf{AcceptList}_i, \mathcal{AC}(\mathsf{AcceptList}_i))$. Let $\mathsf{AcceptList}_k$ be the output of all honest parties.

3. **Generating keys.** Let $x_i = \sum_{j \in \mathsf{AcceptList}_k | \mathsf{AcceptList}_k[j]=2} s_{ji}$ and $x'_i = \sum_{j \in \mathsf{AcceptList}_k | \mathsf{AcceptList}_k[j]=2} s'_{ji}$ be the sum of secret shares in $\mathsf{AcceptList}_k$. Compute $\mathcal{C}_{\langle g \rangle}(x_i)$, $\mathcal{C}_{\langle g,h \rangle}(x_i, x'_i)$ and $\pi_{\equiv \mathsf{Com}\, i} = \mathsf{NIZKPK}_{\equiv \mathsf{Com}}(x_i, x'_i, g, h, \mathcal{C}_{\langle g \rangle}(x_i), \mathcal{C}_{\langle g,h \rangle}(x_i, x'_i))$.

   - Multicast $(\mathcal{C}_{\langle g \rangle}(x_i), \pi_{\equiv \mathsf{Com}\, i})$ to all parties.
   - Verify the received $(\mathcal{C}_{\langle g \rangle}(x_i), \pi_{\equiv \mathsf{Com}\, j})$ as shown in Equation (3).
   - Upon receiving $t+1$ valid $\mathcal{C}_{\langle g \rangle}(x_i)$, interpolate them to obtain $y = g^x$. Set y as the public key and x_i as the private key.

Figure 10: DKG with expected $O(\kappa n^3)$ communication and expected $O(1)$ rounds

the execution. At the end of the recoverable-set-of-shares protocol, each honest party P_i outputs a (possibly different) set of at least $n-t$ parties ($\mathsf{AcceptList}_i$) which they observe to have correctly shared their secret along with an ack-cert for $\mathsf{AcceptList}_i$ ($\mathcal{AC}(\mathsf{AcceptList}_i)$). The ack-cert for $\mathsf{AcceptList}_i$ serves an external validity function to the MVBA protocol i.e., if there is an $\mathcal{AC}(\mathsf{AcceptList}_i)$ for $\mathsf{AcceptList}_i$, then ex-validation($\mathsf{AcceptList}_i$) = true. Note that both $\mathsf{AcceptList}_i$ and $\mathcal{AC}(\mathsf{AcceptList}_i)$ are linear sized. Each honest party P_i then invokes MVBA protocol with $(\mathsf{AcceptList}_i, \mathcal{AC}(\mathsf{AcceptList}_i))$ as input. At the end of MVBA protocol, each honest party outputs a common set $\mathsf{AcceptList}_k$. The final secret key and public key is then computed using secret shares shared by parties h such that $\mathsf{AcceptList}_k[h] = 2$ using the reconstruction protocol in Figure 2.

Latency and communication complexity. The recoverable-set-of-shares protocol has a round complexity of 10 rounds and $O((\kappa + w)n^3)$ communication. The threshold setup protocol incurs a communication of $O((\kappa + w)n^3)$ and 15 rounds; but is executed in parallel and completes before the OLE protocol is invoked in the MVBA protocol. Thus, it does not increase overall round complexity of the protocol. The MVBA protocol incurs expected 4 epochs (with each epoch being 9 rounds) and $O((\kappa + w)n^3)$ communication where the size of input is $O(\kappa n)$. The reconstruction phase requires $O(\kappa n^2)$ communication and a single round. Thus, the protocol incurs $O((\kappa + w)n^3)$ communication and expected 47 rounds.

Theorem 13. *Assuming public-key infrastructure, random oracle, CDH and a universal structured reference string under q-SDH assumption, the protocol in Figure 10 is a secure protocol for distributed key generation in dlog-based cryptosystem tolerating $t < n/2$ Byzantine faults.*

Proof. The MVBA protocol ensures that all honest parties agree on the same set of accepted parties $\mathsf{AcceptList}_k$ at the end of Step 2 in Figure 10.

Given the common set of accepted parties $\mathsf{AcceptList}_k$, the proof of correctness and secrecy properties of a secure DKG protocol remains identical to the proof of Theorem 8. □

8.2 DKG with worst-case $O(\kappa n^3)$ communication and $O(t)$ rounds

While the above protocol terminates in expected 4 epochs in the best case, it has probabilistic termination and may require a linear number of epochs in the worst case with a communication of $O(\kappa n^4)$. As an alternate solution, we present a DKG protocol with guaranteed termination in $t+1$ epochs with $O(\kappa n^3)$ communication in the worst case. The protocol is presented in Figure 11. In the protocol, honest parties execute the recoverable-set-of-shares protocol and each honest party P_i outputs a (possibly different) set of at least $n-t$ parties ($\mathsf{AcceptList}_i$) which they observe to have correctly shared their secret along

with an ack-cert for $\mathsf{AcceptList}_i$ ($\mathcal{AC}(\mathsf{AcceptList}_i)$). The tuple ($\mathsf{AcceptList}_i$, $\mathcal{AC}(\mathsf{AcceptList}_i)$) is input into a leader-based Byzantine fault tolerant state machine replication (BFT SMR) protocol of RandPiper [BSL+21] to agree on a common set. We present a brief overview of the BFT SMR.

---

1. **Deal.** Each party P_i invokes recoverable-set-of-shares protocol (refer Figure 6). Each party P_i output a set $\mathsf{AcceptList}_i$ with an ack-cert for $\mathsf{AcceptList}_i$.

2. **BFT SMR.** Each party P_i participates in BFT SMR [BSL+21] with input $\mathsf{AcceptList}_i$ and $\mathcal{AC}(\mathsf{AcceptList}_i)$. The BFT SMR protocol is executed in round-robin manner with first $t + 1$ leaders. Let $\mathsf{AcceptList}_k$ be the first committed value of all honest parties.

3. **Generating keys.** Let $x_i = \sum_{j \in \mathsf{AcceptList}_k | \mathsf{AcceptList}_k[j]=2} s_{ji}$ and $x'_i = \sum_{j \in \mathsf{AcceptList}_k | \mathsf{AcceptList}_k[j]=2} s'_{ji}$ be the sum of secret shares in $\mathsf{AcceptList}_k$. Compute $\mathcal{C}_{\langle g \rangle}(x_i)$, $\mathcal{C}_{\langle g,h \rangle}(x_i, x'_i))$ and $\pi_{\equiv \mathsf{Com}_i} = \mathsf{NIZKPK}_{\equiv \mathsf{Com}}(x_i, x'_i, g, h, \mathcal{C}_{\langle g \rangle}(x_i), \mathcal{C}_{\langle g,h \rangle}(x_i, x'_i))$.

   - Multicast ($\mathcal{C}_{\langle g \rangle}(x_i), \pi_{\equiv \mathsf{Com}_i}$) to all parties.
   - Verify the received ($\mathcal{C}_{\langle g \rangle}(x_i), \pi_{\equiv \mathsf{Com}_j}$) as shown in Equation (3).
   - Upon receiving $t + 1$ valid $\mathcal{C}_{\langle g \rangle}(x_i)$, interpolate them to obtain $y = g^x$. Set y as the public key and x_i as the private key.

Figure 11: DKG with worst-case $O(\kappa n^3)$ communication and $t + 1$ epochs

BFT SMR of RandPiper [BSL+21]. The BFT SMR protocol of RandPiper [BSL+21] is a communication efficient rotating-leader SMR protocol with $O(\kappa n^2)$ communication per epoch even for $O(n)$-sized input. The BFT SMR protocol has optimal resilience i.e., tolerates $t < n/2$ Byzantine faults. The leaders are rotated in each epoch; in their protocol, an epoch is a duration of 7 rounds. When the leader of an epoch is honest, all honest parties commit the proposed value in the same epoch, whereas, when the leader of the epoch is Byzantine, some honest parties may require linear number of epochs to commit the proposed value. The BFT SMR utilizes the "block-chaining" paradigm i.e., each proposal is represented in the form of a block which explicitly extends a block B proposed earlier by including hash of previous block B. In this paradigm, when a block B is committed, all its ancestors are also committed. We refer the readers to the RandPiper [BSL+21] for more details.

In this DKG protocol, we execute the BFT SMR protocol for $t + 1$ epochs. In each epoch, the epoch leader is expected to propose its ($\mathsf{AcceptList}$, $\mathcal{AC}(\mathsf{AcceptList})$). If the epoch leader is honest, all honest parties commit the proposed set in the same epoch; otherwise honest parties may require linear number of epochs when the leader is Byzantine to commit the proposed value or commit no value at all if the Byzantine leader does not propose. Since the BFT SMR protocol is executed for $t + 1$ epochs, there will be at least one honest leader; thus all honest parties commit at least one set. Honest parties output the first committed set and perform reconstruction using this set to generate the final secret key and public key.

Latency and communication complexity. The recoverable-set-of-shares protocol incurs a latency of 10 rounds and $O(\kappa n^3)$ communication. The BFT SMR protocol incurs $O(\kappa n^2)$ communication per epoch; $O(\kappa n^3)$ communication for $t + 1$ epochs. The length of each epoch is 7 rounds. The reconstruction phase requires $O(\kappa n^2)$ communication and a single round . Thus, the protocol incurs $O(\kappa n^3)$ communication in the worst-case and $11 + 7 * (t + 1)$ rounds.

Theorem 14. *Assuming public-key infrastructure, discrete-log assumption, random oracle, a universal structured reference string under q-SDH assumption, the protocol in Figure 11 is a secure protocol for distributed key generation in dlog-based cryptosystem tolerating $t < n/2$ Byzantine faults.*

Proof. The underlying BFT SMR protocol [BSL+21] ensures that all honest parties agree on a common set of accepted parties $\mathsf{AcceptList}_k$ at the end of Step 2 in Figure 11.

Given the common set of accepted parties $\mathsf{AcceptList}_k$, the proof of correctness and secrecy properties of a secure DKG protocol remains identical to the proof of Theorem 8. $\qquad\square$

9   Related Work

9.1   Related Works in Distributed Key Generation Literature

We review the most recent and closely related DKG protocols. An overview of the closely related work is provided in Table 1. While a myriad of DKG protocols [Ped91, GJKR07, CGJ+99, NBBR16, GJM+21, SJSW19, EFR21, Gro21, CDSV23, Kat23] have been proposed in the synchronous model, all of these protocols assume a broadcast channel. All of these protocols invoke $\Omega(n)$ parallel broadcasts over two or more rounds. A natural choice to instantiate the broadcast channels is via Byzantine consensus primitives such as Byzantine Broadcast [DS83, ADD+19] or Byzantine agreement [KK06]. To the best of our knowledge, all optimally resilient deterministic Byzantine consensus protocols incur $O(\kappa n^3)$ communication without threshold signatures and $t + 1$ rounds [DS83]. For randomized consensus protocols, the best known protocol with optimal resilience in this setting is Katz and Koo [KK06] which incurs $O(\kappa n^4)$ communication. Although, randomized consensus protocols terminate in expected constant rounds, n parallel instances of randomized consensus requires expected $O(\log n)$ rounds to terminate [BOEY03]. While there are known techniques [BOEY03, KK06] to obtain expected constant round protocols via parallel composition of the BA protocol that terminate in constant expected rounds, the communication complexity of the resulting protocol is high ($\Omega(\kappa n^4)$). Compared to all prior DKG protocols, our protocols do not use broadcast channel and use a single instance of Byzantine consensus protocols. Our protocols incur either expected $O(\kappa n^3)$ communication and expected $O(1)$ rounds or worst-case $O(\kappa n^3)$ communication and $O(t)$ rounds. Our protocols are secure against static failures and generate uniform keys for discrete logarithm based cryptosystems.

We also argue that the protocols by Momose and Ren [MR21] and Tsimos et al. [TLP22] are relevant but not sufficient to achieve our goals. Momose and Ren [MR21] gave a deterministic BA protocol with $O(\kappa n^2)$ communication with sub-optimal resilience of $t < (1 - \epsilon)n/2$ for a small constant ϵ. Using their BA protocol to instantiate broadcast channels will result in DKG protocols with $O(\kappa n^3)$ communication but with *sub-optimal* resilience and linear round complexity. Similarly, Tsimos et al. [TLP22] present a communication-efficient broadcast protocol RandomBroadcast in the bulletin PKI setting. It works with $t < (1 - \epsilon)n$ resilience, $O(\kappa^2 n^2)$ communication, linear round complexity, and negligible error probability. Using RandomBroadcast to instantiate broadcast channels will result in DKG protocols with optimal resilience, $O(\kappa^2 n^3)$ communication, linear round complexity and negligible error probability. In contrast, our protocols have optimal resilience, $O(\kappa n^3)$ communication and expected constant rounds (or $O(t)$ rounds).

Pedersen [Ped91] introduced the first efficient DKG protocol for discrete log cryptosystems in the synchronous setting. Their protocol is based on n parallel invocations of Feldman VSS [Fel87]. Gennaro et al. [GJKR07] showed that Pedersen's DKG protocol can be biased by an adversary to generate non-uniform keys. To remove the bias, they proposed a new DKG protocol that requires additional secret sharing rounds; hence, is less efficient. Canneti et al. [CGJ+99] extended Gennaro et al.'s DKG to handle adaptive corruptions.

Neji et al. [NBBR16] presented an efficient DKG protocol to remove the bias without the additional secret sharing round. However, in their protocol, honest parties still need to agree on whether to perform reconstruction for a secret shared by a party which requires additional consensus invocation.

Gurkhan et al. [GJM+21] presented DKG protocol without a complaint phase by using

publicly verifiable secret sharing (PVSS) [CD17] scheme. However, they tolerate only $\log n$ Byzantine faults and do not generate keys for discrete-logarithms based cryptosystems; reducing its usefulness.

Groth [Gro21] presents a non-interactive DKG protocol with a refresh procedure that allows refreshing the secret key shares to a new committee. Erwig et al. [EFR21] considers large scale non-interactive DKG protocol and handles mobile Byzantine faults. Cascudo et al. [CDSV23] presented a DKG protocol using PVSS [CD17] scheme which generates field elements as the secret keys. All of the above protocols assume broadcast channels.

Several other works tackle the DKG problem from different angles. Kate et al. [KHG12] reduced the size of input to the broadcast channel from $O(n)$ to $O(1)$ by using polynomial commitments [KZG10]. Tomescu et al. [TCZ$^+$20] reduce the computational cost of dealings in Kate et al. [KHG12] at the cost of a logarithmic increase in communication cost. Schindler et al. [SJSW19] instantiate the broadcast channel with the Ethereum blockchain. In Table 1, we replaced the Ethereum blockchain with Byzantine consensus primitives for fair comparison.

Kate et al. [KHG12] gave the first practical DKG protocol in the partially synchronous communication model which requires $3t + 2f + 1$ parties to tolerate t Byzantine faults and f crash faults. Kokoris-Kogias et al. [KMS20] gave the first DKG protocol in asynchronous communication model with optimal resilience ($t < n/3$). Their protocol has $O(\kappa n^4)$ communication and $O(t)$ rounds overhead. Abraham et al. [AJM$^+$21] gave an improved DKG protocol with $O(\kappa n^3)$ communication and expected $O(1)$ round complexity. However, their protocol does not generate keys for dlog-based cryptosystems and requires stronger cryptographic assumptions such as SXDH. Das et al. [DYX$^+$22] gave the dlog-based DKG protocol with $O(\kappa n^3)$ communication and optimal resilience in the asynchronous model. However, their protocol incurs expected $O(\log n)$ round complexity and requires stronger Decisional Composite Residuosity (DCR) assumption. Recently, Das et al. [DXKKR23] gave the dlog-based DKG protocol with expected $O(\kappa n^3)$ communication and optimal resilience in the asynchronous model with CDH assumption. However, their construction still incurs expected $O(\log n)$ round complexity. More recently, Abraham et al. [AJM$^+$23] gave an adaptively-secure dlog-based DKG protocol with expected $O(\kappa n^3)$ communication and expected $O(1)$ constant rounds. However, the security of their protocol relies on stronger cryptographic assumptions such as one-more discrete logarithm (OMDL) and non-standard AGM model.

Subsequent work. We note some recent works that achieves DKG protocols with linear round complexity. Bacho et al. [BCLZL23] presented a DKG protocol with $O(\kappa n^3)$ communication and $O(t)$ rounds matching our DKG protocol with linear round complexity. Very recently, Bacho et al. [BLL$^+$23] presented a DKG protocol with sub-cubic communication and linear round complexity. However, their protocol generates group elements as secret keys, uses heavier cryptographic primitives such as SNARKs and relies on idealized and non-standard algebraic group model(AGM) [FKL18]. Additionally, Feng et al. [FLT24] have also designed a DKG protocol with sub-cubic communication and linear round complexity, where the secret keys are field elements. Nevertheless, their protocol also relies on heavier cryptographic primitives such as SNARKs.

Concrete round complexity. All prior synchronous DKG protocols invoke $\Omega(n)$ broadcasts over two or more rounds. Invoking broadcast channels with the state-of-the art Byzantine consensus protocols would require at least $2t + 2$ rounds. Our expected constant round DKG protocol requires only 47 rounds in expectation. When $t > 23$, the concrete round complexity of our protocol is better than prior work.

Concrete communication complexity. The concrete communication complexity of our expected constant round DKG protocol is $46\kappa n^3 + 5\kappa n^2$. Similarly, the concrete communication complexity for our protocol with linear round complexity is $14\kappa n^3 + 5\kappa n^2$. For protocols that assume a broadcast channel, their concrete communication complexity

depends on how the broadcast channel is instantiated. If the state-of-the-art optimally resilient Byzantine broadcast or Byzantine agreement is used, a single instance of it incurs $O(\kappa n^3)$ communication resulting in a total communication cost of $O(\kappa n^4)$ since they invoke broadcast channels $\Omega(n)$ times. Even when using the RandomBroadcast of Tsimos et al. [TLP22], it would incur $O(\kappa^2 n^3)$ which is higher than our constructions when $\kappa > 46$. Moreover, these DKG protocols would have linear round complexity when Tsimos et al. [TLP22] is used.

9.2 Related Works in Byzantine Agreement Literature

There has been a long line of work in improving communication and round complexity of consensus protocols [KK06, FM97, ADD+19, YMR+19, AMS19, MR21, SARN20]. We review the most recent and closely related works.

Multi-valued validated Byzantine agreement was first introduced by Cachin et al. [CKPS01] to allow honest parties to agree on any externally valid values. Their protocol works in asynchronous communication model and has optimal resilience ($t < n/3$) with $O(n^2\ell + \kappa n^2 + n^3)$ communication for input of size ℓ. Later, Abraham et al. [AMS19] gave an MVBA protocol with optimal resilience and $O((\ell + \kappa)n^2)$ communication in the same asynchronous setting. Lu et al. [LLTW20] extended the work of Abraham et al. [AMS19] to handle long messages of size ℓ with a communication complexity of $O(n\ell + \kappa n^2)$. All of these protocols assumed threshold setup. In the absence of threshold setup, the communication complexity blows up by a factor of n in all of these protocols.

Recently, Gao et al. [GLL+22] designed an asynchronous validated BA protocol without threshold setup. Their construction incurs expected $O(\kappa n^3)$ communication and expected $O(1)$ rounds. Their techniques can be used to obtain threshold setup-free asynchronous MVBA protocols. However, their construction relies on aggregatable PVSS [GJM+21] which requires stronger cryptographic assumptions (such as SXDH).

To the best of our knowledge, no MVBA protocol has been formulated in the synchronous setting tolerating $t < n/2$ Byzantine faults. Our MVBA protocol incurs $O(n^2\ell + \kappa n^3)$ for inputs of size ℓ and does not assume threshold setup and terminates in expected constant rounds. The security of our MVBA protocol relies on CDH problem in the random oracle model.

Our MVBA protocol can also be used for binary inputs as a Binary Byzantine Agreement (BBA) protocol tolerating $t < n/2$ Byzantine faults and terminating in expected $O(1)$ rounds. Feldman and Micali [FM97] gave a BBA protocol that terminates in constant expected rounds. Their protocol works in plain authenticated model without PKI and tolerates $t < n/3$ Byzantine faults (which is optimal). In the authenticated setting, Fitzi an Garay [FG03] gave the first BBA protocol with optimal resilience $t < n/2$ and expected constant round complexity that does not require any trusted dealer. However, their protocol requires specific number-theoretic assumptions (essentially, some appropriately-homomorphic public-key encryption scheme). Katz and Koo [KK06] gave a BBA protocol tolerating $t < n/2$ Byzantine faults terminating in expected constant rounds. Their protocol incurs $O(\kappa n^4)$ communication and terminates in expected 4 epochs. We extend the BBA protocol of Katz and Koo [KK06] and reduce its communication by linear factor while handling multi-valued input by designing a communication optimal gradecast protocol and an OLE protocol with $O(\kappa n^3)$ communication, but requires CDH assumptions. A simple and efficient BBA tolerating $t < n/3$ Byzantine faults in the authenticated model was given by Micali [Mic16]. Abraham et al. [ADD+19] reduced the round complexity of BBA protocol to expected 10 rounds. However, their protocol requires a threshold setup to generate a perfect common coin; a perfect common coin ensures all honest parties output the same random value. Compared to their work, our work does not require a threshold setup and executes with a weak common coin.

Recently, Abraham et al. [AAPP22] gave a BBA protocol in the plain authenticated

model without PKI and digital signatures tolerating $t < n/3$ Byzantine faults. Their protocol has an expected communication complexity of $O(n^4 \log n)$ and expected constant rounds.

9.3   Related Work in the Gradecast Literature

Gradecast was originally introduced by Feldman and Micali [FM88] in plain authenticated model without PKI and digital signatures. They gave a protocol tolerating $t < n/3$ Byzantine faults. Katz and Koo [KK06] gave a slightly relaxed definition of gradecast which allows honest parties to output any value with a grade of 1 when no honest party outputs a value with grade of 2. They gave a gradecast protocol tolerating $t < n/2$ Byzantine faults in the authenticated model with PKI and digital signatures. Without threshold signature setup, their gradecast protocol incurs $O(\kappa n^3)$ communication even for a single bit. We also recall the gradecast protocol with multiple grades introduced by Garay et al. [GKKO07] and later improved by [FLZL21]. Their gradecast protocol supports arbitrary number of grades. Their protocol works in the authenticated model with PKI and digital signatures and has a communication complexity of $O(g^*(\ell + \kappa)n^2)$ for input of size ℓ bit where g^* is the maximum grade supported. Compared to all these works, our gradecast protocol satisfies the definition of Katz and Koo [KK06] and tolerates $t < n/2$ Byzantine faults and incurs $O(n\ell + \kappa n^2)$ communication for inputs of size ℓ in authenticated model with PKI and digital signatures.

10   A Lower Bound on the Communication Complexity of Weak Gradecast

In this section, we show a quadratic communication lower bound for the weak gradecast protocol. The proof of this lower bound is a trivial extension of the communication lower bound for Byzantine broadcast by Dolev and Reischuk [DR82].

Lemma 13. *There does not exist a protocol for weak gradecast tolerating t Byzantine parties with a communication complexity of at most $t^2/4$ messages.*

Proof. Suppose for the sake of contradiction, there exists such a protocol. Consider the parties being partitioned into the following two sets: A: a set of $\lceil t/2 \rceil$ parties, and B: all remaining parties which includes the designated sender r.

We consider two executions W1 and W2 where the third property of weak gradecast (i.e., if an honest party outputs a value v with a grade of 2, all other honest parties output value v with a grade ≥ 1) is violated in the W2. In the first execution (W1), all parties in A are Byzantine. Parties in A do not communicate with each other. Towards B, parties in A execute honestly except they ignore the first $\lceil t/2 \rceil$ messages from parties in B. The designated sender $r \in A$ sends value v to all parties. Since, the maximum faults in W1 is $\lceil t/2 \rceil$ and the designated sender is honest, all honest parties decide value v with a grade of 2.

Since the communication complexity of the protocol is at most $t^2/4$, there must exist a party (say s) in A that receives at most $t/2$ messages from parties in B; otherwise the communication complexity will be more than $t^2/4$. Let B_s be the set of all parties that send messages to party s in W1.

In the second execution (W2), all parties in $A \setminus \{s\}$ are Byzantine and all parties in B_s are Byzantine which includes the designated sender r. The total number of Byzantine parties is $(\lceil t/2 \rceil - 1) + \lceil t/2 \rceil \leq t$ which is within allowed fault threshold t. The designated sender r sends value v. The parties in B_s execute the protocol in the same way as in W1 except they do not send any messages to party s. Parties in $A \setminus \{s\}$ execute the protocol in the same way as in W1. Party s in W1 behave as an honest party which did not receive

the first $\lceil t/2 \rceil$ messages which is similar to party s in W2 which receives no messages. Thus, parties in $B \setminus B_s$ cannot distinguish W1 and W2. Thus, they decide value v with a grade of 2. Since, party s does not receive any messages in W2, it does not decide v with a grade of ≥ 1. This violates the third property of weak gradecast where if an honest party outputs a value v with a grade of 2, then all honest parties need to output a value v with a grade ≥ 1. A contradiction. □

Theorem 15. *Let $\mathcal{CC}(\ell)$ be the communication complexity of weak gradecast for ℓ bit input. Then $\mathcal{CC}(\ell) = \Omega(n\ell + n^2)$*

Proof. Since each party must learn ℓ bit input, the protocol needs $\Omega(n\ell)$ bits (The argument follows from [FH06]). From Lemma 13, weak gradecast requires $\Omega(n^2)$ even for a single bit input. Thus, $\mathcal{CC}(\ell) = \Omega(n\ell + n^2)$ for ℓ bit input. □

Acknowledgements

References

[AAPP22]   Ittai Abraham, Gilad Asharov, Shravani Patil, and Arpita Patra. Asymptotically free broadcast in constant expected time via packed VSS. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022: 20th Theory of Cryptography Conference, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 384–414, Chicago, IL, USA, November 7–10, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-22318-1_14.

[ADD+19]   Ittai Abraham, Srinivas Devadas, Danny Dolev, Kartik Nayak, and Ling Ren. Synchronous byzantine agreement with expected $O(1)$ rounds, expected $O(n^2)$ communication, and optimal resilience. In Ian Goldberg and Tyler Moore, editors, *FC 2019: 23rd International Conference on Financial Cryptography and Data Security*, volume 11598 of *Lecture Notes in Computer Science*, pages 320–334, Frigate Bay, St. Kitts and Nevis, February 18–22, 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-32101-7_20.

[AJM+21]   Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Reaching consensus for asynchronous distributed key generation. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, PODC'21, page 363–373, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465084.3467914.

[AJM+23]   Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, and Gilad Stern. Bingo: Adaptivity and asynchrony in verifiable secret sharing and distributed key generation. In *Advances in Cryptology – CRYPTO 2023: 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20–24, 2023, Proceedings, Part I*, page 39–70, Berlin, Heidelberg, 2023. Springer-Verlag. doi:10.1007/978-3-031-38557-5_2.

[AMS19]   Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. In Peter Robinson and Faith Ellen, editors, *38th ACM Symposium Annual on Principles of*

Distributed Computing, pages 337–346, Toronto, ON, Canada, July 29 – August 2, 2019. Association for Computing Machinery. `doi:10.1145/3293 611.3331612`.

[BB08]   Dan Boneh and Xavier Boyen. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, 21(2):149–177, April 2008. `doi:10.1007/s00145-007-9005-7`.

[BCLZL23]   Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, and Julian Loss. Network-agnostic security comes (almost) for free in dkg and mpc. In *Advances in Cryptology – CRYPTO 2023: 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20–24, 2023, Proceedings, Part I*, page 71–106, Berlin, Heidelberg, 2023. Springer-Verlag. `doi:10.1007/978-3-031-38557-5_3`.

[BKP11]   Michael Backes, Aniket Kate, and Arpita Patra. Computational verifiable secret sharing revisited. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 590–609, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany. `doi:10.1007/978-3-642-25385-0_32`.

[BL22]   Renas Bacho and Julian Loss. On the adaptive security of the threshold bls signature scheme. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, page 193–207, New York, NY, USA, 2022. Association for Computing Machinery. `doi:10.1145/3548 606.3560656`.

[BLL+23]   Renas Bacho, Christoph Lenzen, Julian Loss, Simon Ochsenreither, and Dimitrios Papachristoudis. GRandLine: Adaptively secure DKG and randomness beacon with (almost) quadratic communication complexity. *IACR Cryptol. ePrint Arch.*, 2023. `https://eprint.iacr.org/2023/1887`. URL: `https://eprint.iacr.org/2023/1887`.

[BOEY03]   Michael Ben-Or and Ran El-Yaniv. Resilient-optimal interactive consistency in constant time. *Distributed Computing*, 16(4):249–262, 2003. `doi:10.100 7/s00446-002-0083-3`.

[Bol03]   Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46, Miami, FL, USA, January 6–8, 2003. Springer, Heidelberg, Germany. `doi:10.1007/3-540-36288-6_3`.

[BP97]   Niko Bari and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In Walter Fumy, editor, *Advances in Cryptology – EUROCRYPT'97*, volume 1233 of *Lecture Notes in Computer Science*, pages 480–494, Konstanz, Germany, May 11–15, 1997. Springer, Heidelberg, Germany. `doi:10.1007/3-540-69053-0_33`.

[BSL+21]   Adithya Bhat, Nibesh Shrestha, Zhongtang Luo, Aniket Kate, and Kartik Nayak. RandPiper - reconfiguration-friendly random beacons with quadratic communication. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 3502–3524, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press. `doi:10.1145/3460120.3484574`.

[CD17]   Ignacio Cascudo and Bernardo David. SCRAPE: Scalable randomness at-
            tested by public entities. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki
            Kikuchi, editors, *ACNS 17: 15th International Conference on Applied Cryp-
            tography and Network Security*, volume 10355 of *Lecture Notes in Computer
            Science*, pages 537–556, Kanazawa, Japan, July 10–12, 2017. Springer, Hei-
            delberg, Germany. `doi:10.1007/978-3-319-61204-1_27`.

[CDSV23]   Ignacio Cascudo, Bernardo David, Omer Shlomovits, and Denis Varlakov.
            Mt. random: Multi-tiered randomness beacons. In *Applied Cryptography and
            Network Security: 21st International Conference, ACNS 2023, Kyoto, Japan,
            June 19–22, 2023, Proceedings, Part II*, page 645–674, Berlin, Heidelberg,
            2023. Springer-Verlag. `doi:10.1007/978-3-031-33491-7_24`.

[CGJ+99]   Ran Canetti, Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal
            Rabin. Adaptive security for threshold cryptosystems. In Michael J. Wiener,
            editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes
            in Computer Science*, pages 98–115, Santa Barbara, CA, USA, August 15–19,
            1999. Springer, Heidelberg, Germany. `doi:10.1007/3-540-48405-1_7`.

[CKPS01]   Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure
            and efficient asynchronous broadcast protocols. In Joe Kilian, editor, *Advances
            in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer
            Science*, pages 524–541, Santa Barbara, CA, USA, August 19–23, 2001.
            Springer, Heidelberg, Germany. `doi:10.1007/3-540-44647-8_31`.

[CKS00]   Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles
            in constantipole: practical asynchronous byzantine agreement using cryp-
            tography (extended abstract). In Gil Neiger, editor, *19th ACM Sympo-
            sium Annual on Principles of Distributed Computing*, pages 123–132, Port-
            land, OR, USA, July 16–19, 2000. Association for Computing Machinery.
            `doi:10.1145/343477.343531`.

[DF90]   Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard,
            editor, *Advances in Cryptology – CRYPTO'89*, volume 435 of *Lecture Notes
            in Computer Science*, pages 307–315, Santa Barbara, CA, USA, August 20–24,
            1990. Springer, Heidelberg, Germany. `doi:10.1007/0-387-34805-0_28`.

[DR82]   Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for
            byzantine agreement. In Robert L. Probert, Michael J. Fischer, and Nicola
            Santoro, editors, *1st ACM Symposium Annual on Principles of Distributed
            Computing*, pages 132–140, Ottawa, Canada, August 18–20, 1982. Association
            for Computing Machinery. `doi:10.1145/800220.806690`.

[Dra]   Drand. Drand - a distributed randomness beacon daemon. URL: `https:
            //github.com/drand/drand`.

[DS83]   Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine
            agreement. In *SIAM Journal on Computing*, volume 12, pages 656–666. SIAM,
            1983. `doi:10.1137/0212045`.

[DS02]   Paolo D'Arco and Douglas R. Stinson. On unconditionally secure robust
            distributed key distribution centers. In Yuliang Zheng, editor, *Advances in
            Cryptology – ASIACRYPT 2002*, volume 2501 of *Lecture Notes in Computer
            Science*, pages 346–363, Queenstown, New Zealand, December 1–5, 2002.
            Springer, Heidelberg, Germany. `doi:10.1007/3-540-36178-2_22`.

[DXKKR23] Sourav Das, Zhuolun Xiang, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous high-threshold distributed key generation and distributed polynomial sampling. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5359–5376, Anaheim, CA, August 2023. USENIX Association. URL: https://www.usenix.org/conference/usenixsecurity23/presentation/das.

[DYX⁺22] Sourav Das, Thomas Yurek, Zhuolun Xiang, Andrew K. Miller, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous distributed key generation. In *2022 IEEE Symposium on Security and Privacy*, pages 2518–2534, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press. doi:10.1109/SP46214.2022.9833584.

[EFR21] Andreas Erwig, Sebastian Faust, and Siavash Riahi. Large-scale non-interactive threshold cryptosystems in the YOSO model. *IACR Cryptol. ePrint Arch.*, 2021. https://eprint.iacr.org/2021/1290. URL: https://eprint.iacr.org/2021/1290.

[Fel87] Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science*, pages 427–437, Los Angeles, CA, USA, October 12–14, 1987. IEEE Computer Society Press. doi:10.1109/SFCS.1987.4.

[FG03] Matthias Fitzi and Juan A Garay. Efficient player-optimal protocols for strong and differential consensus. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing (PODC'03)*, pages 211–220, 2003. doi:10.1145/872035.872066.

[FH06] Matthias Fitzi and Martin Hirt. Optimally efficient multi-valued Byzantine agreement. In Eric Ruppert and Dahlia Malkhi, editors, *25th ACM Symposium Annual on Principles of Distributed Computing*, pages 163–168, Denver, CO, USA, July 23–26, 2006. Association for Computing Machinery. doi:10.1145/1146381.1146407.

[FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In *Advances in Cryptology (CRYPTO'18): 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part II 38*, pages 33–62. Springer, 2018. doi:10.1007/978-3-319-96881-0_2.

[FLT24] Hanwen Feng, Zhenliang Lu, and Qiang Tang. Breaking the cubic barrier: Distributed key and randomness generation through deterministic sharding. *Cryptology ePrint Archive*, 2024. https://eprint.iacr.org/2024/168. URL: https://eprint.iacr.org/2024/168.

[FLZL21] Matthias Fitzi, Chen-Da Liu-Zhang, and Julian Loss. A new way to achieve round-efficient byzantine agreement. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, PODC'21, page 355–362, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465084.3467907.

[FM88] Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *20th Annual ACM Symposium on Theory of Computing*, pages 148–161, Chicago, IL, USA, May 2–4, 1988. ACM Press. doi:10.1145/62212.62225.

[FM97]   Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, 1997. doi:10.1137/S0097539790187084.

[GJKR07]   Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. *Journal of Cryptology*, pages 51–83, 2007. doi:10.1007/s00145-006-0347-3.

[GJM+21]   Kobi Gurkan, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Aggregatable distributed key generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'21)*, pages 147–176. Springer, 2021. doi:10.1007/978-3-030-77870-5_6.

[GKKO07]   Juan A. Garay, Jonathan Katz, Chiu-Yuen Koo, and Rafail Ostrovsky. Round complexity of authenticated broadcast with a dishonest majority. In *48th Annual Symposium on Foundations of Computer Science*, pages 658–668, Providence, RI, USA, October 20–23, 2007. IEEE Computer Society Press. doi:10.1109/FOCS.2007.61.

[GLL+22]   Yingzi Gao, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Efficient asynchronous byzantine agreement without private setups. In *2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS'22)*, pages 246–257. IEEE, 2022. doi:10.1109/ICDCS54860.2022.00032.

[Gro16]   Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 305–326. Springer, 2016. doi:10.1007/978-3-662-49896-5_11.

[Gro21]   Jens Groth. Non-interactive distributed key generation and key resharing. *IACR Cryptol. ePrint Arch.*, 2021:339, 2021. https://eprint.iacr.org/2021/339. URL: https://eprint.iacr.org/2021/339.

[HMQ04]   Dennis Hofheinz and Joern Mueller-Quade. A synchronous model for multi-party computation and the incompleteness of oblivious transfer. *IACR Cryptol. ePrint Arch.*, 2004. https://eprint.iacr.org/2004/016. URL: https://eprint.iacr.org/2004/016.

[HNP05]   Martin Hirt, Jesper Buus Nielsen, and Bartosz Przydatek. Cryptographic asynchronous multi-party computation with optimal resilience (extended abstract). In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 322–340, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany. doi:10.1007/11426639_19.

[Kat23]   Jonathan Katz. Round optimal fully secure distributed key generation. *IACR Cryptol. ePrint Arch.*, 2023. https://eprint.iacr.org/2023/1094. URL: https://eprint.iacr.org/2023/1094.

[KG09]   Aniket Kate and Ian Goldberg. Distributed key generation for the internet. In *29th IEEE International Conference on Distributed Computing Systems–ICDCS'09*, pages 119–128, 2009. doi:10.1109/ICDCS.2009.21.

[KGS23]   Chelsea Komlo, Ian Goldberg, and Douglas Stebila. A formal treatment of distributed key generation, and new constructions. *IACR Cryptol. ePrint Arch.*, 2023. https://eprint.iacr.org/2023/292. URL: https://eprint.iacr.org/2023/292.

[KHG12]   Aniket Kate, Yizhou Huang, and Ian Goldberg. Distributed key generation in the wild. *IACR Cryptol. ePrint Arch.*, 2012:377, 2012. https://eprint.iacr.org/2012/377. URL: https://eprint.iacr.org/2012/377.

[KK06]   Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 445–462, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany. doi:10.1007/11818175_27.

[KMS20]   Eleftherios Kokoris-Kogias, Dahlia Malkhi, and Alexander Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1751–1767, Virtual Event, USA, November 9–13, 2020. ACM Press. doi:10.1145/3372297.3423364.

[KZG10]   Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194, Singapore, December 5–9, 2010. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-17373-8_11.

[Lab21]   Torus Lab. Torus: Globally accessible public key infrastructure for everyone. https://tor.us/, 2021.

[LLTW20]   Yuan Lu, Zhenliang Lu, Qiang Tang, and Guiling Wang. Dumbo-MVBA: Optimal multi-valued validated asynchronous byzantine agreement, revisited. In Yuval Emek and Christian Cachin, editors, *39th ACM Symposium Annual on Principles of Distributed Computing*, pages 129–138, Virtual Event, Italy, August 3–7, 2020. Association for Computing Machinery. doi:10.1145/3382734.3405707.

[Mer88]   Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology – CRYPTO'87*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378, Santa Barbara, CA, USA, August 16–20, 1988. Springer, Heidelberg, Germany. doi:10.1007/3-540-48184-2_32.

[Mic16]   Silvio Micali. Byzantine agreement, made trivial, 2016. URL: https://api.semanticscholar.org/CorpusID:10040011.

[MR21]   Atsuki Momose and Ling Ren. Optimal communication complexity of authenticated byzantine agreement. In *35th International Symposium on Distributed Computing (DISC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.DISC.2021.32.

[NBBR16]   Wafa Neji, Kaouther Blibech, and Narjes Ben Rajeb. Distributed key generation protocol with a new complaint management strategy. *Security and communication networks*, 9(17):4585–4595, 2016. doi:10.1002/sec.1651.

[Ngu05]   Lan Nguyen. Accumulators from bilinear pairings and applications. In Alfred Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 275–292, San Francisco, CA, USA, February 14–18, 2005. Springer, Heidelberg, Germany. `doi:10.1007/978-3-540-30574-3_19`.

[NRS+20]   Kartik Nayak, Ling Ren, Elaine Shi, Nitin H Vaidya, and Zhuolun Xiang. Improved extension protocols for byzantine broadcast and agreement. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.DISC.2020.28`.

[Ped91]   Torben P. Pedersen. A threshold cryptosystem without a trusted party. In Donald W. Davies, editor, *Advances in Cryptology – EUROCRYPT'91*, volume 547 of *Lecture Notes in Computer Science*, pages 522–526, Brighton, UK, April 8–11, 1991. Springer, Heidelberg, Germany. `doi:10.1007/3-540-46416-6_47`.

[Ped92]   Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140, Santa Barbara, CA, USA, August 11–15, 1992. Springer, Heidelberg, Germany. `doi:10.1007/3-540-46766-1_9`.

[RS60]   Irving S Reed and Gustave Solomon. Polynomial codes over certain finite fields. *Journal of the society for industrial and applied mathematics*, 8(2):300–304, 1960. `doi:10.1137/0108018`.

[SARN20]   Nibesh Shrestha, Ittai Abraham, Ling Ren, and Kartik Nayak. On the optimality of optimistic responsiveness. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 839–857, Virtual Event, USA, November 9–13, 2020. ACM Press. `doi:10.1145/3372297.3417284`.

[Sho00]   Victor Shoup. Practical threshold signatures. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 207–220, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany. `doi:10.1007/3-540-45539-6_15`.

[SJSW19]   Philipp Schindler, Aljosha Judmayer, Nicholas Stifter, and Edgar Weippl. ETHDKG: Distributed key generation with Ethereum smart contracts. *IACR Cryptol. ePrint Arch.*, 2019. `https://eprint.iacr.org/2019/985`. URL: `https://eprint.iacr.org/2019/985`.

[TCZ+20]   Alin Tomescu, Robert Chen, Yiming Zheng, Ittai Abraham, Benny Pinkas, Guy Golan-Gueta, and Srinivas Devadas. Towards scalable threshold cryptosystems. In *2020 IEEE Symposium on Security and Privacy*, pages 877–893, San Francisco, CA, USA, May 18–21, 2020. IEEE Computer Society Press. `doi:10.1109/SP40000.2020.00059`.

[TLP22]   Georgios Tsimos, Julian Loss, and Charalampos Papamanthou. Gossiping for communication-efficient broadcast. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part III*, volume 13509 of *Lecture Notes in Computer Science*, pages 439–469, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany. `doi:10.1007/978-3-031-15982-4_15`.

[YMR+19]   Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai
           Abraham. HotStuff: BFT consensus with linearity and responsiveness. In
           Peter Robinson and Faith Ellen, editors, *38th ACM Symposium Annual on
           Principles of Distributed Computing*, pages 347–356, Toronto, ON, Canada,
           July 29 – August 2, 2019. Association for Computing Machinery. doi:
           10.1145/3293611.3331591.