

Breaking the Cubic Barrier: Distributed Key and Randomness Generation through Deterministic Sharding

Hanwen Feng^{*}, Zhenliang Lu^{*}, and Qiang Tang^{*}

^{*} School of Computer Science,
The University of Sydney, Australia
{hanwen.feng, zhenliang.lu, qiang.tang}@sydney.edu.au

Abstract. There are long line of researches on the fundamental distributed key generation (DKG) protocols. Unfortunately, all of them suffer from a large cubic total communication, due to the fact that $O(n)$ participants need to *broadcast* to all n participants.

We introduce the first two DKG protocols, both achieving optimal resilience, with sub-cubic total communication and computation. The first DKG generates a secret key within an Elliptic Curve group, incurring $\tilde{O}(n^{2.5}\lambda)$ total communication and computation. The second DKG, while slightly increasing communication and computation by a factor of the statistical security parameter, generates a secret key as a field element. This property makes it directly compatible with various off-the-shelf DLog-based threshold cryptographic systems. Additionally, both DKG protocols straightforwardly imply an improved (single-shot) common coin protocol.

At the core of our techniques, we develop a simple-yet-effective methodology via deterministic sharding that *arbitrarily* group nodes into shards; and a new primitive called consortium-dealer secret sharing, to enable a shard of nodes to securely contribute a secret to the whole population only at the cost of one-dealer. We also formalize simulation-based security for publicly verifiable secret sharing (PVSS), making it possible for a modular analysis for DKG. Those might be of independent interest.

1 Introduction

Distributed Key Generation. Distributed key generation (DKG) [18, 25, 33, 34, 36, 42, 49, 52] enables a group of participants to collaboratively create a public key, while each member obtains a secret share of the secret key. This decentralized approach eliminates the need of relying on a trusted party to do the key generation process, serving as one indispensable building block in many distributed protocols, such as secure multiparty computation [22], Byzantine consensus [17, 41], threshold cryptography, and various emerging applications in blockchain (cryptocurrency wallets [35], securing proof of stake against long-range attacks via checkpointing [5], and more).

Despite its fundamental importance, existing DKG protocols involving n participants suffer from a total communication ¹ cost of $\Omega(n^3\lambda)$, ² which is prohibitive even for moderate-scale deployments, particularly when DKG is required to be continuously run in several settings, e.g., [5]. The cubic complexity in DKG arises from the common design method of collectively executing verifiable secret sharing (VSS) [21, 24, 52]. In a nutshell, among n participants where up to t could be adversarial, each participant P_i picks a t -degree polynomial f_i to define $sk^{(i)} = f_i(0)$. They then *deliver* the share $sk_j^{(i)} = f_i(j)$ to other P_j and *broadcast* a commitment, com_i , for the polynomial $f_i(X)$. Participants validate received shares and collectively engage in a *complaint* phase, identifying the set $\mathbb{J} \in [n]$

¹ Throughout the paper, communication cost refers to the *total* communicated bits of honest participants during one execution.

² The only exception is a very recent concurrent work [6], that we will explain the difference at the end of Sec.1.2.

ensuring all transmitted secret shares are valid. The final secret share for P_i is $sk_i = \sum_{j \in \mathbb{J}} sk_i^{(j)}$, and the aggregate secret key is $sk = \sum_{j \in \mathbb{J}} sk^{(j)}$.

Unfortunately, both *delivery* and *complaint*³ phases currently involve n *broadcast* instances via Byzantine broadcast (BB) protocols [29]. Without a common coin [17], deterministic BB protocols generate $\Omega(t^2)$ messages [28].⁴ Ensuring participation from at least $t + 1$ participants is vital for sk 's secrecy, causing the DKG's cubic communication for any $t = \Theta(n)$ ⁵.

If assuming a common coin, the cubic barrier of DKG may be circumvented, for example, by sampling a committee to reduce the number of VSS instances [9], or by using a sub-quadratic randomized BB protocol [2, 44]. However, this workaround raises another crucial question: how can a common coin be established within the group (by those participants collectively) to begin with?

Distributed Common Coin. A common coin protocol allows a group of participants/nodes to produce an unbiased and unpredictable common randomness, which is paramount to many applications, such as lottery [14], committee sampling in distributed protocols [60], and asynchronous consensus [32]. In history, the line of common coin study is closely related to, yet in parallel with DKG.

A major approach to common coin is called *commit-then-reveal*, where each participant P_i first commits a randomness r_i to all others and then reveals r_i such that the coin $r = \sum r_i$. To prevent an attacker from withholding a commitment (after seeing other r_i) to bias the coin, we will need a commitment scheme supporting forced opening, such as publicly verifiable secret sharing (PVSS) [19, 33] and time-locked commitment [56]. This approach similarly requires each participant to *broadcast* its commitment, those $O(n)$ broadcast instances immediately cause $\Omega(n^3\lambda)$ communication cost again when implementing with deterministic BB protocols (now, there is no coin to use).

A recent line of research [10, 11, 23] discards the expensive broadcast procedures and uses a leader node to coordinate the communication; with an honest leader, the group can produce a common coin at the communication cost of $O(n^2\lambda)$. However, as leaders are switched in the Round-Robin manner, only after the $t + 1$ leader election can we guarantee an honest leader. Therefore, for a single-shot coin generation, the leader-based approach still incurs $O(n^3\lambda)$ communication.

Another natural solution for the common coin is letting the group execute a DKG protocol and recover (the hash of) the secret key as the coin. Indeed, as suggested in Cachin et al.'s pioneering work [17] and adopted by many real world projects like *drand* [1], we may use (the hash of) a unique threshold signature as the common coin, after a DKG or trusted key generation for setting up the signing keys. Unfortunately, these approaches based on DKG creates a *circular* problem. In this paper, we are addressing the following long standing problem:

Can we develop a DKG and/or Common Coin protocol with sub-cubic communication complexity?

³ The complaint phase may be obviated using publicly verifiable secret sharing [19, 33].

⁴ Broadcast needs to ensure agreement among all honest receiver outputs, thus causing a higher communication cost than a simple multicast.

⁵ While we focus on synchronous DKG, asynchronous DKGs similarly let each participant invoke an asynchronous variant of broadcast protocol whose communication cost is $\Omega(n^2)$ [15], and consequently, all those asynchronous DKGs [3, 25, 34] also require $O(n^3\lambda)$ communication cost.

In this paper, we answer this question affirmatively. We first present the first (two) DKG protocols with sub-cubic communication complexity, in the standard PKI model with synchronous network. These also lead to the first sub-cubic common coin protocol (without a strong setup) ⁶.

1.1 Our contributions and techniques

We compare our results with the state-of-the-art efficient DKGs in Table.1 and review more related works in Sect.1.2.

A DKG with sub-cubic communication for group-element secrets. Our first DKG protocol has the communication complexity of

$$O(n^{2.5} \cdot |w| + n \cdot \text{BB}_{\sqrt{n}}(n\lambda) + \sqrt{n} \cdot \text{BA}_n(n\lambda)), \quad (1)$$

where n is the number of participants, $|w|$ is the size of membership witness of a cryptographic accumulator scheme, and $\text{BB}_k(\ell)$ (or $\text{BA}_k(\ell)$) represents the communication cost of Byzantine broadcast (or Byzantine agreement (BA)) with k nodes on input of ℓ bits.

Using optimal BB/BA [47, 48]⁷ and an accumulator with $|w| = O(\lambda)$ [50] (which assumes a CRS), the communication cost is $O(n^{2.5}\lambda)$. If CRS is not preferable, we also have a transparent instantiation with the communication cost of $O(n^{2.5} \log n\lambda)$. In contrast, previous constructions incurs communication cost of $O(n \cdot \text{BB}_n(n\lambda))$ or $O(n^3\lambda)$.

Moreover, this DKG can tolerate up to $n/2$ Byzantine nodes (which is optimal resilience), has *publicly verifiable* transcripts (*i.e.*, not need a complaint phase), and produces a secret key in a pairing-friendly Elliptic Curve group (given currently available instantiations). Moreover, it only requires each participant to perform $O(n^{1.5})$ group operations, which is superior to all prior work that needs $O(n^2)$ group operations (when considering optimal resilience).

THE METHODOLOGY: DETERMINISTIC SHARDING + CDSS. While the cubic communication cost of existing DKGs comes from that every participant needs to *broadcast* to the whole population. To get around this, our central methodology is *deterministic sharding*, which simply partitions the whole population into small-sized shards, by following an arbitrarily pre-defined rule; together with CDSS, a new secret sharing primitive Consortium-Dealer Secret Sharing we formulated (to be explained soon), which can enable each shard to jointly contribute one secret, but at the cost of only one single-sender BB.

On rough terms, we let each individual only broadcast to the shard it belongs to, and essentially only one population-level broadcast is needed from each shard (“viewing” as one participant, still jointly executed by the shard members). In this way, we can hope to reduce the communication (say there are \sqrt{n} shards each with \sqrt{n} nodes) to $O(n\text{BB}_{\sqrt{n}}(\cdot) + \sqrt{n}\text{BB}_n(\cdot))$. It is conceptually similar to a “decentralized representative system” where opinions of the whole population are formed only among representatives of each community, but the “representatives” are jointly emulated by all members of the corresponding community, without incurring much communication overhead.

⁶ Common coin protocols that rely on external sources (like stock price [13] or blockchain states [14, 16]) or an initial common coin [26, 38, 58] (which is strictly stronger than common reference string (CRS) model, as the coin has to be independently generated after the PKI setup) do not fit into our setting. These strong setups are unfavorable and may even contradict the primary objective of DKG (or a single-shot coin).

⁷ We discuss the instantiations of BA/BB in detail in Sect.2.1 and here we assume we have optimal BA/BB elsewhere, *i.e.*, $\text{BA}_k(\ell) = \text{BB}_k(\ell) = k\ell + k^2\lambda$.

Table 1. Comparison with the state-of-the-art synchronous DKG protocols.

Schemes	Resi.	Field?	PV?	Comm. Cost (total)		Comp. Cost (per node)	Round
				w. oracles	opt. imp.		
Pedersen [52]	1/2	✓	✗	$O(n \cdot \text{BB}_n(n\lambda))$	$O(n^3\lambda)$	$O(n^2)$	$O(\Delta_{\text{BB}(n)})$
KZG [43](G.)	1/2	✓	✗	$O(n \cdot \text{BB}_n(\lambda))$	$O(n^3\lambda)$	$O(n \log n)$	$O(\Delta_{\text{BB}(n)})$
KZG [43](B.)				$O(n \cdot \text{BB}_n(n\lambda))$	$O(n^3\lambda)$		
FS [33]	1/2	✓	✓	$O(n \cdot \text{BB}_n(n\lambda))$	$O(n^3\lambda)$	$O(n^2)$	$O(\Delta_{\text{BB}(n)})$
GJM+ [42]	$\frac{\log n}{n}$	✗	✓	$O(n \text{BB}_n(\lambda) + \log n \cdot \text{BB}_n(n\lambda))$	$O(n^3\lambda)$	$O(n \log^2 n)$	$O(\Delta_{\text{BB}(n)})$
SBKN [55]	1/2	✓	✓	$O(n^3\lambda)$		$O(n^2)$	$O(n)$
Ours Sect.4	1/2	✗	✓	$O(n \cdot \text{BB}_{\sqrt{n}}(n\lambda) + \sqrt{n} \cdot \text{BB}_n(n\lambda))$	$O(n^{2.5}\lambda)$	$O(n^{1.5})$	$O(\Delta_{\text{BB}(n)})$
Ours Sect.5	1/2	✓	✗	$O(n \cdot \text{BB}_{\sqrt{n}}(n\lambda) + \sqrt{n} \cdot \text{BB}_n(n\lambda\kappa))$	$O(n^{2.5}\lambda\kappa)$	$O(n^{1.5}\kappa)$	$O(\Delta_{\text{BB}(n)})$

λ : computational security parameter; κ : statistical security parameter

Resi.: the maximal fraction of byzantine nodes the protocol can tolerate.

Field? asks if the secret key is in a finite field \mathbb{Z}_p for some prime p .

PV? asks if the transcripts of the protocol are publicly verifiable.

Round: $\Delta_{\text{BB}(n)}$ is the number of rounds of a Byzantine broadcast among n parties.

Comp. Cost measures the number of group exponent operations performed by each node. We assume the optimization from [19] has been applied whenever applicable.

Comm. Cost measures the number of bits sent by all honest nodes. In **w. oracles**, the cost is analyzed with oracle calls to BB/BA.

In **opt. imp.**, we calculate the cost with optimal BB and BA protocols. Some caveat exists, nevertheless, our claim of *sub-cubic* communication still holds, and we discussed in detail in Sect.2.1.

In KZG (G.), we analyze the cost of the optimized KZG protocol [61] when there is no complaint; In KZG (B.), we analyze the cost when there are $O(n^2)$ complaints.

The idea of dividing participants into groups is not new, conventional sharding techniques (originated from the database community) exist in the literature [27, 60] for consensus and more. However, our deterministic sharding is fundamentally different: conventional sharding for consensus protocols usually aims to improve the throughput, thus letting each shard concurrently output transactions, and hoping all shard transactions can be collected and confirmed. This in turn requires almost *every* shard to function properly, thus relying on common coins (which we don't have for DKG and of course for common coin itself) to randomly divide the population to ensure the corruption ratio in each of the shards, while ours is deterministic and arbitrary.

Nonetheless, we observe that for DKG, we may not need the above strong guarantee for each shard: recall the secret key sk is the form of $\sum sk^{(i)}$, where $sk^{(i)}$ is contributed by shard- i . Having one honestly contributed is sufficient for the secrecy of sk . We further observe that even with a deterministic sharding, there is always *at least one* shard that contains an honest majority.

THE CORE COMPONENTS: CONSORTIUM-DEALER SECRET SHARING. After deterministically dividing population into small-size shards, our main technical tool is to make each shard “behave” like a single node to contribute one secret to the whole population, in terms of both security and efficiency.

Specifically, the secret contributed by a good shard should remain *secret* to the adversary, and the total communication cost should be close to that of a single-sender broadcast. Trivial approaches have to sacrifice one of the two: If all nodes in the shard distribute their secrets independently to all n nodes, the cost is blown up to cubic again; if only one node in each shard contributes the secret,

we may lose secrecy (as we have no way to pick an honest one). We formulate a new secret sharing primitive CDSS to capture this functionality, where a group of dealers (in one shard) form a dealer consortium and jointly share one secret with a larger population. If the dealer consortium has an honest-majority, it is as good as a conventional secret sharing with an honest dealer: all receivers have correct shares while the adversary does not know the secret. If not, we still have robustness which ensures all receivers obtain a correct share (or abort). We then show a DKG can be built by concurrently executing multiple CDSS instances where each shard acts as a dealer consortium. Please see Sect.4.1 for details.

Constructing CDSS efficiently requires conquering the following challenges: first, similar to conventional DKG, when nodes (now in each shard) jointly contribute one secret $sk^{(i)}$, they need to distribute each of their shares to the whole population for future reconstruction; but at the same time, they cannot directly reveal those shares — they together can be used to reconstruct $sk^{(i)}$, which might be the only “good” secret (thus adversary can recover the final sk without enough shares). We leverage the *aggregatable* PVSS [19, 42]⁸ to handle the secrecy concerns of the shard secret. PVSS produces a sequence of encrypted shares under receivers’ public keys with proof of well-formedness of ciphertexts; aggregatable PVSS allows to compression of many transcripts (including the ciphertexts and proofs) into one. We let each node in the shard generate a PVSS transcript for the whole population but only broadcast it *within the shard*. Then, everyone in the shard aggregates the received PVSS transcripts into the *same* (because of the public verifiability) one transcript ($sk^{(i)}$ is not leaked at all), which is to be sent to the whole population.

Next is to deliver the (aggregated) shard transcript to the whole population. The security-efficiency dilemma strikes again: letting everyone broadcast ruins our efforts to reduce communication, while if the task is assigned to only a small number of nodes, they could all be malicious. To tackle this issue, we formulate a new primitive called consortium-sender byzantine broadcast (CSBB) for a consortium of senders to broadcast to a larger population at the cost of one-sender broadcast. BB protocols usually consist of a delivery phase, where the sender sends the message to receivers; and an agreement phase, where all receivers agree on the same message. Inspired by (but different with) BA extension protocols [46], we utilize error-correction code [12] to realize a conceptual “deduplicable storage system”: all shard members “store” their transcript into it, and receivers can retrieve the most frequently appeared one. While the “storage system” can serve as an efficient delivery phase, we can have a secure broadcast protocol by following only one agreement phase. We also use a cryptographic accumulator [50] to help error-correction. Please see Section.4.2 for details.

A simple corollary: a sub-cubic common coin protocol. Our DKG protocol gives rise to a sub-cubic common coin protocol. Concretely, to produce a common coin, a group of n participants initiates our DKG protocol. Upon the DKG’s completion, participants exchange their secret shares amongst themselves. Subsequently, they can reconstruct the secret key and derive the common coin by hashing this key. Given that the adversary lacks prior knowledge of the secret key, the

⁸ We remark aggregatable PVSS has been leveraged recently to reduce the cost of DKG by Gurkan et al. [42] from $nBB_n(n\lambda)$ to $nBB_n(\lambda) + \log n \cdot BB_n(n\lambda)$ (but it is still $O(n^3\lambda)$, and with the price of lowering resilience to $O(\log n/n)$). We take a different approach to use aggregatable PVSS together with deterministic sharding and our customized CDBB (to be explained), finally breaking the cubic barrier.

coin remains unpredictable and unbiased in the random oracle model. Table.2 ⁹ contrasts this construction with other common coin protocols that don't rely on a strong setup.

Table 2. Comparison with common coin rotocols without a strong setup

Techniques	Schemes	Resi.	Comm. Cost total	Comp. Cost (per node)
Time-lock Puzzl.	TCLM [56]	$\frac{n-1}{n}$	$O(n \cdot \text{BB}_n(\lambda))$	Time-lock Puzzl.
PVSS	Scrape [19]	1/2	$O(n \cdot \text{BB}_n(n\lambda))$	$O(n^2)$
Leader-PVSS	OptRand [10]	1/2	$O(n^3\lambda)$	$O(n^2)$
DKG	Ours	1/2	$O(n \cdot \text{BB}_{\sqrt{n}}(n\lambda) + \sqrt{n} \cdot \text{BA}_n(n\lambda))$	$O(n^{1.5})$

Extension to DKG for field-element secrets. We then present a DKG protocol whose secret key is in the finite field \mathbb{Z}_p for some secure prime p , while it enjoys the same complexity and security as our group-element DKG does. This protocol can be a drop-in replacement for ElGamal encryption, BLS signature [7], Schnorr signature [54], and more [40].¹⁰

The technical challenge arises from the current incompatibility of aggregatable PVSS with field-element secrets. To address this, we adopt a more general approach, constructing a CDSS from a conventional (non-aggregatable) PVSS scheme. In our CDSS with group-element secrets, aggregatable PVSS is employed to reduce the message size broadcasted by the dealer group while ensuring the secrecy of the corresponding secret key against adversaries. We observe that both goals can be achieved using a common coin within the dealer group. After the dealers broadcast their PVSS transcripts within the group, a common coin is employed to determine a small number (denoted by κ , linear to the statistical security parameter) of valid transcripts. This ensures that at least one of them is from an honest dealer with overwhelming probability. The dealer group then broadcasts all selected transcripts to all receivers via CSBB. Finally, the receiver decrypts all transcripts and computes the sum of the decrypted values as the received share. As at least one transcript is from an honest dealer, the secret key remains unknown to the adversary. However, since the dealer group now needs to broadcast κ PVSS transcripts, the communication cost is higher than that of the aggregatable PVSS-based approach by a factor of κ . Please refer to Section 5 for detailed information.

PVSS with simulation-based security. While our DKGs use (aggregatable) PVSS as building blocks, we find it challenging to provide a modular security proof due to the gap between the secrecy definitions of DKG and PVSS. Specifically, DKG requires simulation-based secrecy since it should simulate a trusted key generation process of, e.g., threshold cryptosystems. In contrast, existing works formalize PVSS's secrecy with a game-based indistinguishability definition. This approach cannot capture the secret key distribution from an adversarial view and fails to address potential malleability issues. Given these gaps, PVSS-based DKG constructions like [42] must reduce their security to underlying hard problems without using the abstraction of PVSS.

⁹ Notably, we omit various randomness beacon schemes, such as those based on PoW/VDF [58], as they pre-suppose an initial coin, making them unsuitable for single-shot common coin generation.

¹⁰ [40] showed that key-expressible DKG suffices for rekeyable primitives including ElGamal and BLS. While Schnorr signature is not rekeyable, its key generation algorithm can also be securely instantiated with a key-expressible DKG, as recently discussed in [40]

We close this gap by studying simulation-based security for PVSS, similar to zero-knowledge proof definitions, including: (1) *Secrecy*. A PVSS transcript w.r.t. a uniform public key can be simulated without knowing the secret key. (2) *Soundness*. Any adversary cannot produce a valid transcript that will be decrypted to a set of inconsistent shares. (3) *Simulation soundness*. Some form of soundness must be preserved, even the adversary has seen a simulated transcript.

Simulation soundness provides a certain guarantee of non-malleability that is indispensable for proving simulation based security for DKG. To facilitate a formal definition, we further let each transcript carry a tag or its creator identity (CID). Only transcripts with different CIDs than the simulated one will be considered as a successful attack against simulation soundness, excluding the trivial attack of duplicating the simulated one. For aggregatable PVSS, we view the aggregated transcript in the same form as the original ones, such that the above security applies to aggregatable PVSS. Please see Sect.3 for more details.

Regarding constructions, it's not surprising to see some existing PVSS schemes already satisfy our definitions. We demonstrate a PVSS satisfying our definitions in Appendix.A by slightly modifying existing constructions like [33, 37]. We showcase a simplified variant of Gurkan et al.'s aggregatable PVSS [42] in Appendix.B.

1.2 Related Works

Distributed Key Generation (DKG) has emerged as a significant research domain over the decades. Pedersen [52] laid the groundwork in this area by presenting the first efficient protocol for Dlog-based cryptosystems, which is built upon Feldman's Verifiable Secret Sharing (VSS) [30]. In Pedersen's approach, all users collaboratively execute n instances of Feldman's VSS, with each user acting as the dealer for one instance.

Within Feldman's VSS framework, the dealer must broadcast a commitment to a polynomial while privately dispatching the shares to all other users. Given that the commitment size is $O(n\lambda)$, the communication overhead stands at $O(n\text{BB}_n(n\lambda))$, where $\text{BB}_n(\ell)$ represents the communication cost for executing a Byzantine broadcast protocol among n nodes with an ℓ -bit input. Furthermore, Pedersen's DKG necessitates a complaint phase where users broadcast their complaints against dishonest dealers. Considering that a user might broadcast multiple complaints simultaneously, the communication overhead for this phase also matches $O(n\text{BB}_n(n\lambda))$. It's crucial to note, however, that during this complaint phase, each user might verify up to $O(n^2)$ shares. For Feldman's VSS, the computational cost associated with verifying a single share amounts to $O(n)$ group operations. Therefore, the per-node computational overhead before the complaint phase is $O(n^2)$, which can escalate to $O(n^3)$ in the complaint phase.

Most DKG constructions adhere to the joint-VSS paradigm. Essentially, every novel VSS protocol can be transformed into a new DKG protocol. Moreover, since VSS can be established on polynomial commitments, every polynomial commitment scheme can also culminate in a VSS and, ultimately, a DKG. Kate et al. [43] proposed the first polynomial commitment (denoted as KZG) with an $O(\lambda)$ commitment size. Consequently, prior to the complaint phase, the communication cost can be tapered down to $O(n\text{BB}_n(\lambda))$. This, however, is still asymptotically $O(n^3\lambda)$, even when paired with an optimal broadcast protocol. An added advantage of the KZG polynomial commitment is its efficiency in share verification; verifying a single share incurs a mere $O(1)$ cost. This implies that the per-node computational cost for verification before the complaint phase is a mere $O(n)$ in group operations, though it can inflate to $O(n^2)$ during the complaint phase. While the

computational overhead for generating the polynomial commitment was believed to be $O(n^2)$ [57], a novel study by Zhang et al. [61] demonstrated that the computational overhead for generating a KZG commitment can be optimized to $O(n \log n)$. Though KZG mandates a CRS setup, other endeavors [59, 61] focusing on efficient polynomial commitments without a trusted setup have been explored, but they fall short of KZG’s efficiency.

Fouque and Stern [33] circumvented the need for a complaint phase by integrating publicly verifiable secret sharing (PVSS). Given that a PVSS transcript encompasses $O(n)$ ciphertexts, the communication cost invariably aligns with $O(n \text{BB}_n(n\lambda))$, should all users opt to broadcast the transcript. Traditionally, verifying a PVSS transcript demanded an $O(n^2)$ overhead, implying that the per-node computational cost in DKG could reach $O(n^3)$. This hurdle was overcome by Cascudo and David in Scrape [19], which introduced a PVSS scheme that limits verification time to $O(n)$. Notably, Scrape’s methodology is versatile and can be applied to enhance several preceding techniques, including Pedersen’s, ensuring that the computational overhead during the complaint phase remains at $O(n^2)$ rather than surging to $O(n^3)$. A dedicated line of research, evident in works like [37], has aimed to refine the concrete performance of PVSS.

Gurkan et al. [42] harnessed an aggregatable PVSS in tandem with gossip protocols to devise a publicly verifiable DKG. Their communication cost is $n \text{BB}_n(\lambda) + \log n \cdot \text{BB}_n(n\lambda)$ (still $O(n^3 \lambda)$) instead of $n \text{BB}_n(n\lambda)$, and their per-node communication cost is $O(n \log^2 n)$. However, their scheme can only tolerate $O(\log n)$ Byzantine nodes. Shrestha et al. [55] charted a different path, presenting a DKG without resorting to BB protocols. Instead, they employed an MVBA [46] protocol to facilitate agreement, which culminates in an $O(n^3 \lambda)$ communication overhead. They posited that achieving optimal resilience with BB without a private setup should require $O(n^3)$ communication cost. Yet, in light of recent advancements in transparent threshold signatures [4], this hypothesis might need reevaluation. We delve deeper into the intricacies of BB/BA in Sect. 2.1, while elsewhere, we assume the existence of an optimal BA/BB.

Beyond these efforts directed at augmenting the efficiency of DKG, there have been other studies addressing this challenge from distinct criteria. Gennaro et al. [36] discerned that the secret key distribution in Pedersen’s DKG could be influenced by adversarial entities. They rectified this shortcoming, achieving full secrecy but at a higher computational cost. Gurkan et al. [42] introduced a weaker version of secrecy termed “key-expressability”, which postulates that adversaries might influence key distribution but within confined parameters. They argued that a key-expressable DKG suffices for a plethora of applications, with numerous DKG frameworks, including those of Pedersen [52], Fouque-Stern [33], and our own, aligning with this definition. Canetti et al. [18] contributed a DKG protocol with adaptive security, a contrast to our model and several others that ensure security against only static adversaries. Recent contributions by Bacho and Loss [7] introduced an oracle-aided adaptive definition and verified that multiple protocols align with this definition in the algebraic group model.

Lastly, a few recent endeavors [3, 25, 34] have shifted the focus towards DKG in asynchronous networks. These constructions adopt the joint-VSS framework and rely on an asynchronous broadcast protocol termed “reliable broadcast” [15] to ensure verifiability, consequently still facing the cubic computational barrier. Notably, Das et al. [25] presented the pioneering asynchronous DKG with an $O(n^3 \lambda)$ communication overhead for field-element secrets, while Abraham et al. [3] delivered an adaptively secure asynchronous DKG with same complexity.

Concurrent Work: In a very recent concurrent work [6]¹¹, Bacho et al. also introduced a DKG with sub-cubic communication complexity. Here, we provide a comparison between two works.

Regarding techniques, the one introduced in [6] and ours differ significantly. Their DKG employs dedicated consensus to agree on aggregated PVSS transcripts, while ours leveraging arbitrary grouping together with consortium dealer secret sharing (and broadcast), which can be built upon any BB/BA in a blackbox manner. These underlying techniques may find different applications beyond DKGs. Also, these different techniques may lead to different efficiency, security, functionality trade-offs as briefly elaborated below.

In terms of communication complexity, their DKG (with a communication complexity of $\mathcal{O}(n^2 \log n\lambda)$) asymptotically outperforms our current constructions. We didn't do further optimization when getting communication down to sub-cubic; in principle, our techniques maybe applied recursively and further bring down the communication complexity.

Regarding functionality, we provide sub-cubic DKGs for both group-element secrets and field-element secrets, while [6] only demonstrates a DKG with group-element secrets. It is worth noting that it is possible to use the DKG in [6] as a common coin to construct a communication-efficient DKG with field-element secrets, for instance, by sampling a committee of dealers [31]. However, the resulting scheme cannot be strongly adaptively secure [2], even assuming an adaptively secure PVSS with field-element secrets. In contrast, our field-element DKG can be strongly adaptively secure, assuming an adaptively secure PVSS and memory erasures.

In terms of security, the two results are generally incomparable. While their primary application is a randomness beacon, the DKG in [6] is proven to be *unpredictable* against *adaptive* adversaries in the algebraic group model. In contrast, our focus is on applications to threshold cryptography, and we prove that our DKG satisfies *key-expressibility* against *static* adversaries in the standard model (while our instantiations may assume a random oracle). However, we believe there is no significant gap between the security guarantees provided by the two works. Our DKG schemes can be adaptively secure if the underlying components, particularly the PVSS scheme, are adaptively secure. By utilizing the adaptive PVSS scheme introduced in [8], our group-element DKG can achieve the same security guarantee as the DKG in [6]. Conversely, if focusing on static security, their DKG may achieve the stronger *key-expressibility*, although further analysis is required.

2 Model, Preliminaries, and Protocol Composition

Notations. Throughout the paper: We use λ to represent the security parameter. The notation $[i, n]$ represents the set $\{i, i+1, \dots, n\}$, where i and n are integers with $i < n$. We might abbreviate $[1, n]$ simply as $[n]$. For a set $\{x_1, x_2, \dots, x_n\}$ and a sequence (x_1, x_2, \dots, x_n) , we may abbreviate them as $\{x_i\}_{i \in [n]}$ and $(x_i)_{i \in [n]}$, respectively. A function $f(n)$ is deemed negligible in n , denoted by $f(n) \leq \text{negl}(n)$, if for every positive integer c , there exists an n_0 such that for all $n > n_0$, $f(n) < n^{-c}$. Conversely, a non-negligible function is denoted as $f(n) > \text{negl}(n)$. For a set \mathbb{X} , the notation $x \leftarrow_{\$} \mathbb{X}$ signifies sampling x uniformly from \mathbb{X} . Given a distribution X , $x \leftarrow X$ denotes sampling x from X . For a probabilistic algorithm A , $A(x_1, x_2, \dots; r)$ represents the result of running A with inputs x_1, x_2, \dots and random coins r . We use $y \leftarrow A(x_1, x_2, \dots)$ to represent choosing r randomly and obtaining $y = A(x_1, x_2, \dots; r)$. If \odot represents a group operation in \mathbb{G} , then $g_1 \odot g_2 \cdots \odot g_n$ is denoted as $\bigodot_{i \in [n]} g_i$ where each $g_i \in \mathbb{G}$. An execution of the protocol Π

¹¹ A preliminary version of our paper was submitted to Eurocrypt 2024 in October 2023, while [6] was available online in December 2023.

involving n participants P_i , each with inputs v_i , is represented by $\Pi\langle\{P_i(v_i)\}_{i\in[n]}\rangle$. Adversaries are assumed to be probabilistic polynomial time (PPT).

Communication and threat model. We assume the network is synchronous and the protocol proceeds in rounds. The network is fully connected, meaning there is a communication channel between each pair of nodes. We assume the channel is authenticated since every message is signed by its sender using an existentially-unforgeable signature scheme. Every message sent by an honest node is guaranteed to be received by an honest recipient before the beginning of the next round. We measure communication complexity by the number of bits sent by honest nodes.

We assume an initial phase that optionally generates a common reference string (CRS) and sets up PKI for every participant. We consider a *static* adversary \mathcal{A} , which, given the CRS (if any), specifies the set of corrupted nodes, and generates the public keys on behalf of the corrupted nodes after seeing public keys of honest nodes. After the initial phase, all public keys and CRS are accessible to every participant, and \mathcal{A} can arbitrarily control the behaviour of corrupted nodes. Allowing \mathcal{A} to determine the corrupted set after observing the CRS aligns with many scenarios where the CRS was universally established and reused across multiple executions.

2.1 Byzantine Agreement and Broadcast

Byzantine Agreement. In an (n, t, ℓ) -Byzantine Agreement (BA) protocol, there are n parties P_1, \dots, P_n , each P_i having an ℓ -bit initial input v_i , denoted as $\text{BA}\langle P_i(v_i) \rangle$. Against any adversary \mathcal{A} that corrupts up to t parties, a secure BA ensures the following properties:

- **Validity.** If all honest parties share the same input v , they all output v .
- **Agreement.** All honest parties output the same message.
- **Termination.** All honest parties produce an output message.

Byzantine Broadcast. In an (n, t, ℓ) -Byzantine Broadcast (BB) protocol, there is a sender P_s with ℓ -bit input message msg and a set of receivers $\mathcal{P} = \{P_1, \dots, P_n\}$, denoted as $\text{BB}\langle P_s(\text{msg}), \mathcal{P} \rangle$. A secure BB has the same agreement and termination guarantee as BA does, while it concerns the following validity.

- **Validity.** If the sender is honest, all honest receivers output the sender’s input message msg .

Instantiations of BA and BB. We first examine the candidates of BA protocols. For ease of reference, we use $\text{BA}_n(\ell)$ to denote the communication complexity of a BA protocol among n parties with ℓ -bit input. We summarize the status in Table.3, and elaborate on them in the following.

Table 3. Summary of BA Protocols

Candidates	$\text{BA}_n(\ell)$	Resilience	Setup	Computation
B1 + E1	$O(n\ell + n^2\lambda)$	$t < (1/2 - \epsilon)n$	CRS	light
B1 + E2	$O(n\ell + n^2 \log n\lambda)$	$t < (1/2 - \epsilon)n$	Transparent	light
B2 + E2	$O(n\ell + n^2 \log n\lambda)$	$t < n/2$	Transparent	moderate
B3 + E1	$O(n\ell + n^2\lambda)$	$t < n/2$	CRS	heavy

In scenarios where $\ell = O(\lambda)$, Momose and Ren [47] have presented the state-of-the-art protocols. These include:

- A BA protocol (referred to as B1) with $\text{BA}_n(\lambda) = O(n^2\lambda)$, capable of tolerating $t < (1/2 - \epsilon)n$ corruptions, for any positive constant $\epsilon \in (0, 1/2)$. This protocol requires only conventional digital signatures.
- Another BA protocol with $\text{BA}_n(\lambda) = O(n^2 \cdot \text{thrSigSize})$, tolerating $t < n/2$ corruptions, where thrSigSize is the size of a threshold signature. Traditional threshold signatures require either trusted key generation or a DKG phase, which is unfavorable for our purpose. However, recent progress [4] provides a threshold signature of size $O(\lambda \log n)$ with a transparent setup (referred to as B2). Additionally, a constant-size zk-SNARK [39] can also yield a threshold signature of $O(\lambda)$ size (referred to as B3).

For cases where $\ell > k$, Nayak et al. [48] have provided extension protocols. These result in $\text{BA}_n(\ell) = O(n \cdot \ell + \text{BA}_n(\lambda) + n^2\lambda)$ if a CRS for a pairing-based accumulator is allowed (referred to as E1). Alternatively, if the CRS is not allowed, the complexity becomes $\text{BA}_n(\ell) = O(n \cdot \ell + \text{BA}_n(\lambda) + n^2 \log n\lambda)$ (referred to as E2).

It has been suggested that $\text{BA}_n(\ell) = O(n\ell + n^2\lambda)$ could be optimal for $\ell > \lambda$ in the PKI setting [48], despite the proved lower bound being $O(n\ell + n^2)$ [28]. Combining B3 and E1 could yield a BA protocol with optimal communication complexity and optimal resilience $t < \frac{n}{2}$, although it necessitates a CRS setup and potentially intensive computation. Other combinations may alleviate concerns about setup or computation, offering sub-optimal communication complexity or resilience.

For $t < n/2$, a BB protocol can be considered where the sender first multicasts its input to all receivers, and then the receivers execute a BA protocol to finalize their output. Thus, $\text{BB}_n(\ell) = O(\text{BA}_n(\ell))$. Suitable instantiations follow our discussion about BA protocols

Nevertheless, the study of optimal BA/BB protocols is a rapidly evolving field and largely unrelated to our primary focus. Hence, for simplicity, in the remainder of this paper, we assume a BA protocol with optimal communication complexity and resilience and do not account for potential setup requirements and computational overhead.

2.2 Distributed Key Generation

Homomorphic Key Structure. We consider a generic homomorphic key structure. Let \mathcal{SK} represent the group of secrets with the operation \oplus , and let \mathcal{PK} denote the group of public keys with the operation \otimes . The structure includes a PPT algorithm KeyGen and a relation $\text{Rela} \subset (\mathcal{PK}, \mathcal{SK})$ such that

$$\Pr[(pk, sk) \leftarrow \text{KeyGen}(1^\lambda) : (pk, sk) \in \text{Rela}] = 1.$$

We say that $(\mathcal{PK}, \mathcal{SK})$ forms a homomorphic key structure if Rela is homomorphic. That is, for any $pk_1, pk_2 \in \mathcal{PK}$ and $sk_1, sk_2 \in \mathcal{SK}$ such that $(pk_i, sk_i) \in \text{Rela}$ ($i = 1, 2$) and any integers $\alpha, \beta \in \mathbb{N}$, it holds that

$$(\alpha pk_1 \otimes \beta pk_2, \alpha sk_1 \oplus \beta sk_2) \in \text{Rela}.$$

The key generation algorithm KeyGen should satisfy specific properties to be useful in cryptography. However, these specificities are not central to our discussion on securely decentralizing the algorithm.

Instantiation. Throughout this paper, we focus on two key structures. The first is the standard key structure in DLog-based cryptography. Here, the key generation algorithm yields $sk \leftarrow \mathcal{SK} := \mathbb{Z}_p$ and $pk = g^{sk} \in \mathcal{PK} := \mathbb{G}$, where p is a secure prime and g is a generator of a cyclic group \mathbb{G} of order p . The relationship $(pk, sk) \in \text{Rela}$ holds iff $pk = g^{sk}$.

The second structure is a pairing-based one [42]. In this structure, the key generation algorithm creates $pk = (g^s, u^s) \in \mathbb{G}_1 \times \mathbb{G}_2$ and $sk = h^s \in \mathbb{G}_2$, where $s \leftarrow \mathbb{Z}_p$ for a particular secure prime p . \mathbb{G}_1 (with generator g) and \mathbb{G}_2 (with generators u and h) are cyclic groups of order p . A bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ exists. The relation $(pk = (pk', pk''), sk) \in \text{Rela}$ holds iff $e(pk', h) = e(g, sk)$ and $e(pk', u) = e(g, pk'')$. Here, $pk_1 \oplus pk_2 = (pk'_1 \cdot pk'_2, pk''_1 \cdot pk''_2)$.

DKG. An (n, t) -DKG protocol for $\{\mathcal{PK}, \mathcal{SK}\}$, denoted as Π_{DKG} , involves n parties, $\mathcal{P} = (P_1, \dots, P_n)$. After execution, each P_i outputs a public key $pk \in \mathcal{PK}$, a list of public key shares $(pk_i)_{i \in [n]} \in \mathcal{PK}^n$, and a secret key share $sk_i \in \mathcal{SK}$. The protocol includes an initial phase and a reconstruction algorithm:

- **Init** $(1^\lambda, n)$. It generates a CRS crs and establishes the PKI for \mathcal{P} .
- **Rec** $((i, sk_i)_{i \in \mathbb{I}})$. Given a set of $t + 1$ secret key shares as input, it outputs the secret key sk for pk .

In this paper, we address Π_{DKG} meeting both *robustness* and *key-expressability*, the latter being a weaker form of *secrecy*.

- **Robustness:** Π_{DKG} is robust if, even when up to t parties are compromised, every honest P_i outputs the same $(pk, (pk_i)_{i \in [n]})$, and its sk_i satisfies $(pk_i, sk_i) \in \text{Rela}$. Additionally, for any two sets, \mathbb{I}_1 and \mathbb{I}_2 , with $t + 1$ honest participants each, a unique secret key sk can be reconstructed from their secret shares, as described by the equation below:

$$\Pr \left[\begin{array}{l} \text{Rec}(pk, (pk_i)_{i \in [n]}, \{sk_i\}_{i \in \mathbb{I}_1}) = \text{Rec}(pk, (pk_i)_{i \in [n]}, \{sk_i\}_{i \in \mathbb{I}_2}) \\ \wedge (pk, \text{Rec}(\{(i, sk_i)\}_{i \in \mathbb{I}_1})) \in \text{Rela} \end{array} \right] = 1.$$

- **Key expressability:** Π_{DKG} is *key expressible* if, for any PPT adversary \mathcal{A} that compromises up to t nodes, a PPT simulator $\text{Sim}^{\mathcal{A}}$ exists such that the following equation holds for any PPT distinguisher \mathcal{A}' :

$$\left| \Pr \left[\begin{array}{l} \Pi_{\text{DKG}}^{\mathcal{A}}(\{P_i(1^\lambda)\}_{i \in [n]}) \\ \rightarrow (pk, \text{view}_{\mathcal{A}}) : \\ \mathcal{A}'(pk, \text{view}_{\mathcal{A}}) = 1 \end{array} \right] - \Pr \left[\begin{array}{l} \text{KeyGen}(1^\lambda) \rightarrow (pk, sk), \\ \text{Sim}^{\mathcal{A}}(pk) \rightarrow (sk', pk'), \\ \alpha \in \mathbb{Z}^+, \text{sview}_{\mathcal{A}} : \\ \mathcal{A}'(\alpha pk \otimes pk_1, \text{sview}_{\mathcal{A}}) = 1 \\ \wedge (sk', pk') \in \text{Rela} \end{array} \right] \right| \leq \text{negl}(\lambda),$$

where the notation $\Pi_{\text{DKG}}^{\mathcal{A}}(\{P_i(1^\lambda)\}_{i \in [n]}) \rightarrow (pk, \text{view}_{\mathcal{A}})$ represents an execution of Π_{DKG} involving the adversary \mathcal{A} , including the initial phase. Here, pk is the public key generated by the execution, and $\text{view}_{\mathcal{A}}$ represents the adversary's view during the execution, including all public messages and its internal states; KeyGen is the default key generation algorithm for $(\mathcal{PK}, \mathcal{SK})$.

The concept of key-expressability is weaker than *full secrecy* as described in [36]. The latter ensures that a computationally bounded adversary cannot deduce any information about the secret sk beyond what's publicly available. However, as demonstrated in [42], key-expressability is sufficient for instantiating the key generation algorithm for a variety of cryptosystems which are "re-keyable", including BLS signatures, ElGamal encryption, etc. Moreover, as shown in [54], it also suffices for Schnorr signatures. On the other hand, some applications such as random coins may only require the DKG to be **unpredictable**, i.e., for any PPT adversary \mathcal{A} , it holds that

$$\Pr[\Pi_{\text{DKG}}^{\mathcal{A}}(\{P_i(1^\lambda)\}_{i \in [n]}) \rightarrow (pk, \text{view}_{\mathcal{A}}), \mathcal{A}(\text{view}_{\mathcal{A}}) \rightarrow sk : (pk, sk) \in \text{Rela}] \leq \text{negl}(\lambda).$$

Unpredictability is clearly implied by key expressability.

2.3 Other Cryptography Primitives

Cryptographic accumulators. A cryptographic accumulator provides a succinct representation of a set while ensuring efficient membership verification. Formally, such an accumulator scheme, denoted as Acc , comprises the following four algorithms: (1) $\text{Gen}(1^\lambda, n)$ outputs an accumulator key ak . (2) $\text{Eval}(ak, \mathcal{S})$ on inputs ak and a set \mathcal{S} to be accumulated, it returns an accumulated value u for the set \mathcal{S} . (3) $\text{Wit}(ak, u, \mathcal{S}, s_i)$ on inputs ak, u for the set \mathcal{S} , and an element $s_i \in \mathcal{S}$, it returns a membership witness w_i for s_i . (4) $\text{Vrfy}(ak, u, s_i, w_i)$ decides if s_i is an element accumulated into u , using the witness w_i .

An accumulator scheme is *correct*, if for $ak \leftarrow \text{Gen}(1^\lambda, n)$, any set $\mathcal{S} = \{s_i\}_{i \in [n]}$, $u \leftarrow \text{Eval}(ak, \mathcal{S})$, and any $w_i \leftarrow \text{Wit}(ak, u, s_i)$, it holds that $\Pr[\text{Vrfy}(ak, u, s_i, w_i) = 1] = 1$. An accumulator scheme is *unforgeable*, if for an honestly generated ak , and any PPT adversary,

$$\Pr[(\mathcal{S}, s^*, w^*) \leftarrow \mathcal{A}(ak) : s^* \notin \mathcal{S} \wedge \text{Vrfy}(ak, \text{Eval}(ak, \mathcal{S}), s^*, w^*) = 1] \leq \text{negl}(\lambda).$$

For simplicity, throughout this paper, we consider an accumulator scheme whose Eval and Wit are deterministic.

Instantiation. We primarily consider the pairing-based accumulator [50] which requires a CRS and has $O(\lambda)$ -sized witness. Merkle tree is also a candidate featured by its transparent setup, although the witness size is $O(\lambda \log n)$.

Erasur code scheme. A (k, n) -erasure code scheme [12] consists of two deterministic algorithms Encode and Decode . The Encode algorithm maps any vector $\mathbf{m} = (m_1, \dots, m_k)$ of k data fragments into an vector $\mathbf{c} = (c_1, \dots, c_n)$ of n coded fragments, such that any k elements in the code vector \mathbf{c} is enough to reconstruct \mathbf{m} due to the Decode algorithm. I.e., for any $\mathbf{m} \in \mathcal{B}^k$ and any $\mathbb{I} \subset [n]$ that $|\mathbb{I}| = k$, we have

$$\Pr[\text{Decode}(\{(i, c_i)\}_{i \in \mathbb{I}}) = \mathbf{m} \mid \mathbf{c} := (c_1, \dots, c_n) \leftarrow \text{Encode}(\mathbf{m})] = 1.$$

Instantiation. Throughout the paper, we consider a $(t+1, n)$ -erasure code where $2t+1 = n$. Additionally, it's important to note that this erasure code scheme will implicitly select an appropriate \mathcal{B} based on the actual length of each element in \mathbf{v} . This ensures that the encoding results in only a constant size increase.

NIZK. A non-interactive zero-knowledge (NIZK) proof system Π , for an NP language L , enables the prover, who holds a witness of an instance $x \in L$, to convince the verifier that $x \in L$ via a single proof. Typically, it can be described by the following a triple of probabilistic polynomial-time (PPT) algorithms:

- $\sigma \leftarrow \text{Setup}(1^\lambda)$. The setup algorithm outputs a CRS σ .
- $\pi \leftarrow \text{Prove}(\sigma, x, w)$. The prover algorithm takes as inputs the CRS σ , an instance $x \in L$ with its witness $w \in R_L(x)$, and outputs a string π called a proof.
- $b \leftarrow \text{Verify}(\sigma, x, \pi)$. The verifier algorithm takes as inputs σ , an instance x and a proof π , and outputs either 1 accepting it or 0 rejecting it.

We consider a NIZK satisfying *completeness*, *zero-knowledge*, and *simulation soundness*.

1. **COMPLETENESS:** For all security parameters $\lambda \in \mathbb{N}$ and for all $x \in L_\lambda$ and $w \in R_L(x)$,

$$\Pr[\sigma \leftarrow \text{Setup}(1^\lambda); \pi \leftarrow \text{Prove}(\sigma, x, w) : \text{Verify}(\sigma, x, \pi) = 1] = 1.$$

2. ZERO KNOWLEDGE: There is a PPT simulator $(\text{SimSetup}, \text{SimProve})$, *s.t.* for every PPT adversary \mathcal{A} , we have

$$\Pr[\sigma \leftarrow \text{Setup}(1^\lambda) : 1 \leftarrow \mathcal{A}^{\mathcal{O}_1(\sigma, \cdot)}(\sigma)] - \Pr[(\sigma, \tau \leftarrow \text{SimSetup}(1^\lambda) : 1 \leftarrow \mathcal{A}^{\mathcal{O}_2(\sigma, \tau, \cdot)}(\sigma)] \leq \text{negl}(\lambda).$$

Both the oracles \mathcal{O}_1 and \mathcal{O}_2 take as input a pair $(x, w) \in R_L(x)$. While \mathcal{O}_1 returns $\pi \leftarrow \text{Prove}(\sigma, x, w)$, \mathcal{O}_2 returns $\pi \leftarrow \text{SimProve}(\sigma, \tau, x)$.

3. SIMULATION SOUNDNESS: For any PPT adversary \mathcal{A} , it holds that

$$\Pr \left[(\sigma, \tau) \leftarrow \text{SimSetup}(1^\lambda); (x^*, \pi^*) \leftarrow \mathcal{A}^{\mathcal{O}_2(\sigma, \tau, \cdot)}(\sigma) : \text{Verify}(\sigma, x^*, \pi^*) = 1 \wedge x^* \notin L \right] \leq \text{negl}(\lambda).$$

We highlight that NIZK can also be constructed in the random oracle model, without assuming a CRS.

2.4 Unique Identifier Model

Our DKG constructions are built upon multiple sub-protocols, which could run concurrently. To ensure the security of our DKGs, each sub-protocol must remain secure during simultaneous operations. Lindell *et al.* [45] highlighted that many BA protocols retain their security in concurrent settings if each is given a *unique identifier*. We define this concept as follows.

Definition 1 (Unique identifier model.) *Protocol Π uses a signature scheme $\Sigma = (\text{Gen}, \text{Sign}, \text{Vrfy})$ to sign/verify its messages. In the unique identifier model, every instance of Π gets a distinct identifier id , leading to a modified protocol Π_{id} . Π_{id} is like Π , but it utilizes $\Sigma_{id} = (\text{Gen}, \text{Sign}_{id}, \text{Vrfy}_{id})$, where $\text{Sign}_{id}(sk, m) = \text{Sign}(sk, id||m)$ and $\text{Vrfy}_{id}(vk, m, \sigma) = \text{Verify}(vk, id||m, \sigma)$. Different instances must have prefix-free id strings, ensuring one id isn't a prefix of another.*

Simply put, protocols can maintain concurrent security in this model by disregarding messages with different identifiers, ensuring security akin to isolated settings. Most consensus protocols should be concurrently secure in this model.

3 Simulation-based Definitions for PVSS

In this section, we provide simulation-based definitions for publicly verifiable secret sharing (PVSS), making it a valuable tool in the realm of DKG. With the new definitions, we will be able to analyze our DKGs in a modular way.

Brief Overview. An (n, t) -secret sharing (SS) scheme allows a dealer to distribute a secret s among n participants. Any group of $t + 1$ honest parties can reconstruct s , yet any smaller group (up to t parties) remains oblivious to s . Whereas SS assumes a trustworthy dealer, verifiable secret sharing (VSS) addresses the possibility of a malicious dealer by letting a receiver validate the consistency of its share with a public commitment. PVSS takes VSS a step further: all encrypted shares, complete with a verification proof, are placed on a public channel for universal validation. An *aggregatable* PVSS can compress multiple PVSS transcripts into one single publicly verifiable transcript.

The Need for Simulation-based Definitions. One of DKG’s core goals is to act as a stand-in for the trusted key generation phase of threshold cryptosystems. Given this, DKG should be able to *emulate* standard key generation to serve a wide range of distributed cryptography applications, thus DKG usually takes a simulation based security modeling.

However, PVSS was mostly formulated using a game-based indistinguishability definition, termed IND-secrecy [19]. This definition doesn’t fully capture the essence of key distribution from an adversary’s perspective. It overlooks potential malleability challenges which, in a DKG setting, could allow adversaries to arbitrarily sway key distribution. This makes a security reduction for DKG using PVSS as a blackbox inherently challenging.

To bridge this gap, we put forth simulation-based definitions for PVSS. Notably, we enrich PVSS’s syntax to include the set of creator identities (CID) in each transcript. This prevents adversaries from merely replicating a simulated transcript. It’s worth mentioning that recent research by Bacho and Loss [8] also formalized aggregatable PVSS and incorporated the ID into its syntax. However, their primary application was to randomness beacons, and their definitions did not adopt a simulation-based approach.

The Syntax. We describe aggregatable PVSS with the following eight algorithms/phases. For simplicity, we assume that the “native” transcripts (produced by Deal) and the aggregated transcripts are in the same form (though they differ by the size of their CID set), and thus all algorithms and properties apply to both types of transcripts. The syntax and definitions for a (non-aggregatable) PVSS can be obtained by removing the algorithm **Agg**.

- **Init**($1^\lambda, n$): In the initial phase, a CRS \mathbf{crs} is generated, and the encryption/decryption keys $\{(ek_i, dk_i)\}_{i \in [n]}$ for all participants are set up. \mathbf{crs} is an implicit input for all other algorithms.
- **Deal**($((ek_i)_{i \in [n]}), \mathbf{cid}$) \rightarrow (**Trans**, sk): It produces a secret $sk \in \mathcal{SK}$ and a transcript **Trans**, consisting of a commitment \mathbf{com} to the secret sk , ciphertexts $(c_i)_{i \in [n]}$, a proof π of validity, and the CID set $\{\mathbf{cid}\}$.
- **Agg**($\{(\mathbf{Trans}_i)\}_{i \in [m]}, (ek_i)_{i \in [n]}$) \rightarrow **Trans**: It outputs an aggregated transcript **Trans** whose CID set is $\{\mathbf{cid}_i\}_{i \in [m]}$, where \mathbf{cid}_i is from \mathbf{Trans}_i .
- **PubVrfy**($((ek_i)_{i \in [n]}), \mathbf{Trans}$) $\rightarrow b$: It checks if **Trans** is valid.
- **getCID**(**Trans**) $\rightarrow \{\mathbf{cid}_i\}_{i \in [m]}$: It returns the CID set.
- **PubDriv**(**Trans**) $\rightarrow (pk, (pk_i)_{i \in [n]})$: It derives the public key (shares).
- **Dec**($(ek_i, dk_i), \mathbf{Trans}$) $\rightarrow sk_i$: One can decrypt the ciphertext c_i in **Trans** and obtain the secret share sk_i .
- **Rec**($\{(i, sk_i)\}_{i \in \mathbb{I}}$) $\rightarrow sk$: It first determines coefficients $\{\alpha_i\}_{i \in \mathbb{I}}$, where $\alpha_i \in \mathbb{N}$ based on \mathbb{I} and reconstructs the committed secret sk as $\bigoplus_{i \in \mathbb{I}} \alpha_i sk_i$ from any subset $\mathbb{I} \subset [n]$ and $|\mathbb{I}| = t + 1$.

Security. A PVSS scheme should satisfy *correctness*, *soundness*, *secrecy*, and *simulation soundness*. In the following definitions, we use $\mathbf{Init}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\mathbf{crs}, \mathcal{C}, (ek_i, dk_i)_{i \notin \mathcal{C}}, (ek_i)_{i \in \mathcal{C}}, st_{\mathcal{A}})$ to denote an execution of the initial phase involving the adversary \mathcal{A} , where \mathcal{C} is the set of corrupted nodes, and $st_{\mathcal{A}}$ is the state of the adversary.

- **Correctness:** Assume $\mathbf{Init}_{\mathcal{A}}(1^\lambda, n)$ has been done. For $(\mathbf{Trans}, sk) \leftarrow \mathbf{Deal}(ek_1, ek_2, \dots, ek_n)$, the transcript can always be verified, *i.e.*,

$$\Pr[\mathbf{PubVrfy}((ek_i)_{i \in [n]}, \mathbf{Trans}) = 1] = 1, \tag{2}$$

Assume $\{\text{Trans}_j\}_{j \in [m]}$ are valid “native” transcripts, and $\text{PubDriv}(\text{Trans}_j) \rightarrow (pk^{(j)}, (pk_i^{(j)}))$. For $\text{Agg}(\{\text{Trans}_j\}_{j \in [m]}, (ek_i)_{i \in [n]}) \rightarrow \text{Trans}$, it holds that $\text{PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}) = 1$, and

$$\text{PubDriv}(\text{Trans}) = \left(\bigotimes_{j \in [m]} pk^{(j)}, \left(\bigotimes_{j \in [m]} pk_i^{(j)} \right)_{i \in [n]} \right).$$

- **Soundness:** Any adversary cannot produce a valid transcript while it will be decrypted to a set of inconsistent shares. Formally, assume $\text{Init}_{\mathcal{A}}(1^\lambda, n)$ has been done. If a transcript is verified, *i.e.*, $\text{PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}) = 1$, then for any two subsets \mathbb{I}_1 and \mathbb{I}_2 of $t + 1$ uncorrupted participants, the secret recovered from the transcript is unique, *i.e.*,

$$\Pr [\text{Rec}(\{\text{Dec}(ek_i, dk_i, \text{Trans})\}_{i \in \mathbb{I}_1}) = \text{Rec}(\{\text{Dec}(ek_i, dk_i, \text{Trans})\}_{i \in \mathbb{I}_2})] = 1. \quad (3)$$

- **Secrecy:** There is a triple of PPT simulators $\{\text{SInit}, \text{SDeal}, \text{SRec}\}$. Such that, for any static PPT adversary \mathcal{A} which corrupts up to t nodes, a PVSS transcript by an honest dealer does not leak sk beyond its public information, *i.e.*, for any cid ,

$$\left| \begin{array}{l} \Pr \left[\begin{array}{l} \text{Init}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i, dk_i)_{i \notin \mathcal{C}}, (ek_i)_{i \in \mathcal{C}}, st_{\mathcal{A}}) \\ \text{Deal}((ek_i)_{i \in [n]}, \text{cid}) \rightarrow (\text{Trans}, sk) : \\ \mathcal{A}(\text{crs}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{Trans}) = 1 \end{array} \right] \\ - \Pr \left[\begin{array}{l} \text{KeyGen}(1^\lambda) \rightarrow (pk, sk), \\ \text{SInit}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{tk}), \\ \text{SDeal}((ek_i)_{i \in [n]}, pk, \text{tk}, \text{cid}) \rightarrow \text{Trans} : \\ \mathcal{A}(\text{crs}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{Trans}) = 1 \\ \wedge \text{PubDriv}(\text{Trans}) = (pk, \cdot). \end{array} \right] \end{array} \right| \leq \text{negl}(\lambda), \quad (4)$$

where KeyGen is the default key generation algorithm of $(\mathcal{PK}, \mathcal{SK})$.

- **Simulation soundness.** Some form of soundness must be preserved, even after the adversary seeing a simulated transcript. Formally, for any static PPT adversary, it holds that

$$\Pr \left[\begin{array}{l} \text{KeyGen}(1^\lambda) \rightarrow (pk, sk), \\ \text{SInit}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{tk}), \\ \text{SDeal}((ek_i)_{i \in [n]}, pk, \text{tk}, \text{cid}) \rightarrow \text{Trans}, \\ \mathcal{A}(\text{crs}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{Trans}) \rightarrow \text{Trans}^* \\ \text{SRec}(\text{tk}, \text{Trans}^*) \rightarrow sk^*, \text{PubDriv}(\text{Trans}^*) \rightarrow (pk^*, \cdot) : \\ (pk^*, sk^*) \in \text{Rela} \wedge \text{cid} \notin \text{getCID}(\text{Trans}^*) \\ \wedge \text{PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}^*) = 1 \end{array} \right] \leq \text{negl}(\lambda). \quad (5)$$

Remark that we need both soundness and simulation soundness. The former does not directly imply the latter due to different ways of extraction.

PVSS instantiation. Conventional PVSS schemes [33, 37] with minor enhancements can meet our definitions. In Appendix.A we present a secure PVSS scheme for the standard key structure in DLog-based cryptography, *i.e.*, $sk \in \mathbb{Z}_p$ and $pk = g^{sk} \in \mathbb{G}$. This scheme is obtained by following the general “encrypt-and-proof” paradigm, and additionally applying a signature of knowledge (SoK) to embed a creator ID into its transcript. We further explicitly employ the Scrape’s technique [19] to improve its verification time. While the PVSS scheme in Appendix.A could have various instantiations, we summarize the result about an LWE-based instantiation (which can be seen as a variant of [37]) as its concrete performance stands out.

Lemma 1. *Under the LWE assumption and DL assumption [37], there is a PVSS scheme for the standard key structure in DLog-based cryptography, satisfying correctness, soundness, secrecy, and simulation soundness. Particularly, the transcript size of \mathbf{Trans} is $O((n)\lambda)$. Both \mathbf{Deal} and $\mathbf{PubVrfy}$ require $O(n)$ group operations. The computational costs for other functions are minor, approximately $O(1)$ group operations.*

Aggregatable PVSS instantiation. All known aggregatable PVSS are variants of Scrape PVSS [19, 42] for the pairing-based key structure. For completeness, we present a variant of Scrape PVSS in Appendix.B which meets our definitions. We summarize the result about the instantiation in the following lemma.

Lemma 2. *Under the SXDH and BDG assumption [42], there is an aggregatable PVSS scheme for the pairing-based key structure, satisfying correctness, soundness, secrecy, and simulation soundness. Particularly, the transcript size of \mathbf{Trans} is $O((n + m)\lambda)$, where m represents the number of transcripts aggregated into \mathbf{Trans} . Both \mathbf{Deal} and $\mathbf{PubVrfy}$ require $O(n)$ group operations, whereas \mathbf{Agg} demands $O(n \log m)$ group operations. The computational costs for other functions are minor, approximately $O(1)$ group operations.*

4 Sharded DKG with Sub-Cubic Communication

Deterministic Sharding. As discussed in Introduction, conventional sharding relies on common randomness to ensure all shards to have adequate honest participants. This creates a circular issue for DKG/Coin, which is meant to establish such randomness. Instead, we propose a deterministic sharding, which divides the population into shards using a deterministic predefined rule. While this method may not offer a strong guarantee, it still ensures that at least one group maintains the honesty ratio, which can be leveraged together with CDSS.

Lemma 3 (Any-good Partition). *For a population $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ and a partition $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_m\}$ over $[n]$. If there are t corrupted nodes, denoted as $\{P_i\}_{i \in \mathcal{C}}$ where $\mathcal{C} \subset [n]$ and $|\mathcal{C}| = t$, then there exists a subset \mathcal{S}_j such that the proportion of corrupted nodes in $\{P_i\}_{i \in \mathcal{S}_j}$, given by $\frac{|\mathcal{C} \cap \mathcal{S}_j|}{|\mathcal{S}_j|}$ is at most $\frac{t}{n}$.*

Proof. Suppose that no subset \mathcal{S}_j satisfies the condition, i.e., $\forall j \in [m], |\mathcal{C} \cap \mathcal{S}_j| > \frac{t|\mathcal{S}_j|}{n}$. It follows $\sum_{j \in [m]} |\mathcal{C} \cap \mathcal{S}_j| > t$. However, we also have $\sum_{j \in [m]} |\mathcal{C} \cap \mathcal{S}_j| \leq |\mathcal{C}| \leq t$, which contradicts our assumption. \square

The ratio-preservation of deterministic sharding introduces an avenue for more efficient DKG. Existing DKG designs, rooted in VSS, derive the secret key from the aggregation of all qualified secret sharings. The key’s secrecy is intact as long as one honest participant is involved. Conventional designs require the completion of at least $t + 1$ VSS, leading to $\Omega(n^3)$ communication complexity for $t = \Theta(n)$. With deterministic sharding ensuring an honest majority in at least one group, we have the opportunity to treat each group as a “single entity” to share a secret, thereby reducing required secret sharings while retaining secrecy.

In this section, we introduce a new secret sharing paradigm called *Consortium-Dealer Secret Sharing* (CDSS) to realize the idea of treating a participant group as a singular dealer in terms of efficiency and security. We then outline a DKG framework built on CDSS and present a CDSS construction for better DKG.

4.1 DKG from Consortium-Dealer Secret Sharing

Consortium-Dealer Secret Sharing (CDSS). We formalize the notion of CDSS, which enables a consortium of participants to distribute shares of a *random value* to a large population.

Syntax. An (n, η, t, τ) -CDSS scheme for $(\mathcal{PK}, \mathcal{SK})$ involves n participants $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ with a special subset $\mathcal{D} = \{D_1, D_2, \dots, D_\eta\} \subset \mathcal{P}$ acting as a dealer consortium. It consists of an initialization phase Init , a deal protocol $\text{Deal}\langle \mathcal{D}, \mathcal{P} \rangle$, and a reconstruction algorithm Rec .

- $\text{Init}(1^\lambda, n)$. This sets up the PKI and generates a CRS crs .
- $\text{Deal}\langle \mathcal{D}, \mathcal{P} \rangle$. At the end of the protocol, each receiver P_i outputs a public key $pk \in \mathcal{PK}$, a sequence of public key shares $(pk_i)_{i \in [n]} \in \mathcal{PK}^n$, and a secret key share $sk_i \in \mathcal{SK}$.
- $\text{Rec}(\{(i, sk_i)\}_{i \in \mathbb{I}})$. It reconstructs the secret key sk for pk . We require Rec to be linear, i.e., it first determines coefficients $\{\alpha_i\}_{i \in \mathbb{I}}$ where $\alpha_i \in \mathbb{N}$ and reconstructs the secret $sk = \bigoplus_{i \in \mathbb{I}} \alpha_i sk_i$ for any subset $\mathbb{I} \in [n]$ with $|\mathbb{I}| = t + 1$.

We consider the *robustness* and *key-expressability* of CDSS in the *multi-instance* setting. Assume that after an honest initialization phase $\text{Init}(1^\lambda, n)$, there are m instances $\{\text{Deal}\langle \mathcal{D}^{(j)}, \mathcal{P} \rangle\}_{j \in [m]}$ running concurrently in the unique identifier model (cf. Def. 1). Assume there is a PPT adversary \mathcal{A} that corrupts $\{P_i\}_{i \in \mathcal{C}}$ for $|\mathcal{C}| \leq t$. We detail each property below.

- **Multi-instance robustness.** For any \mathcal{A} and any integer m polynomial in λ , we have the following guarantees for each instance:
 - For the j -th instance where $|\{P_i\}_{i \in \mathcal{C}} \cap \mathcal{D}^{(j)}| \leq \tau$, all honest P_i 's output properly. That is, every P_i outputs the same public tuple $(pk, (pk_i)_{i \in [n]})$ and its secret share sk_i such that $\text{Rela}(pk_i, sk_i) = 1$. Meanwhile, for any two subsets $\mathbb{I}_1, \mathbb{I}_2 \subset [n]$ and $|\mathbb{I}_1| = |\mathbb{I}_2| = t + 1$, it follows that the same sk is reconstructed from $\{sk_i\}_{i \in \mathbb{I}_1}$ and $\{sk_i\}_{i \in \mathbb{I}_2}$, and $(pk, sk) \in \text{Rela}$.
 - For the j -th instance where $|\{P_i\}_{i \in \mathcal{C}} \cap \mathcal{D}^{(j)}| > \tau$, all honest receivers either return \perp or output properly.
- **Multi-instance key-expressability.** For any \mathcal{A} and any integer m polynomial in λ such that $\exists j \in [m], |\{P_i\}_{i \in \mathcal{C}} \cap \mathcal{D}^{(j)}| \leq \tau$, there is a PPT simulator algorithm $\text{Sim}^{\mathcal{A}}$. For any PPT distinguisher \mathcal{A}' , it holds that

$$\left| \begin{array}{l} \Pr \left[\begin{array}{l} \{\{\text{Deal}_{id_j}^{\mathcal{A}}\langle \mathcal{D}^{(j)}, \mathcal{P} \rangle\}_{j \in [m]}\} \rightarrow ((\text{out}_j)_{j \in [m]}, \text{view}_{\mathcal{A}}), \\ s.t. \text{out}_j = pk^{(j)} \text{ or } \perp: \mathcal{A}'((\text{out}_j)_{j \in [m]}, \text{view}_{\mathcal{A}}) = 1 \end{array} \right] \\ \text{KeyGen}(1^\lambda) \rightarrow (pk, sk), \text{Sim}^{\mathcal{A}}(pk) \rightarrow (\{\text{tup}\}_{j \in [m]}, \\ \text{sview}_{\mathcal{A}}, s.t. \text{tup}_j = (sk'^{(j)}, pk'^{(j)}, \alpha^{(j)}) \text{ or } \perp, \\ - \Pr \left[\begin{array}{l} \text{out}_j \leftarrow \alpha^{(j)} \cdot pk \otimes pk'^{(j)}, \text{ or } \text{out}_j \leftarrow \perp \text{ if } \text{tup}_j = \perp: \\ (\text{if } \text{tup}_j \neq \perp, \text{ then } (sk'^{(j)}, pk'^{(j)}) \in \text{Rela}) \\ \wedge (\exists j^*, \alpha^{(j^*)} \neq 0) \wedge \mathcal{A}'((\text{out}_j)_{j \in [m]}, \text{sview}_{\mathcal{A}}) = 1 \end{array} \right] \end{array} \right] \leq \text{negl}(\lambda).$$

$\langle \{\text{Deal}_{id_j}^{\mathcal{A}}\langle \mathcal{D}^{(j)}, \mathcal{P} \rangle\}_{j \in [m]}\rangle \rightarrow ((\text{out}_j)_{j \in [m]}, \text{view}_{\mathcal{A}})$ denotes a concurrent execution of m CDSS instances $\{\text{Deal}_{id_j}\}_{j \in [m]}$ involving \mathcal{A} following an execution of the same initialization. out_j is the public key $pk^{(j)}$ that an honest P_i outputs in the instance or \perp if it aborts; $\text{view}_{\mathcal{A}}$ is the view of the adversary \mathcal{A} , including all public messages and its internal states. KeyGen is the default key generation algorithm for $(\mathcal{PK}, \mathcal{SK})$.

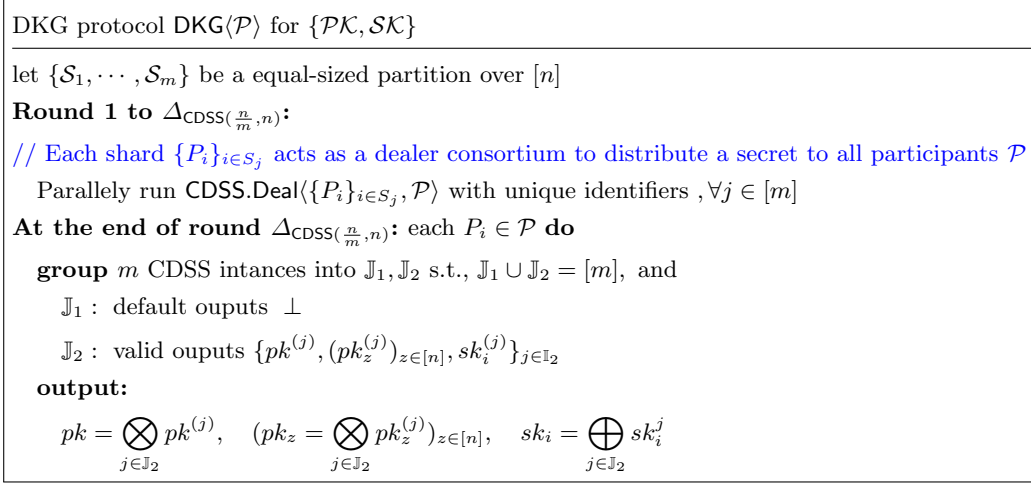


Fig. 1. DKG from CDSS. $\Delta_{\text{CDSS}(\eta, n)}$ is the number of rounds needed for running the CDSS's deal protocol CDSS.Deal with a dealer consortium of η members and a receiver set of n nodes.

DKG from CDSS. Let $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ be the whole population. Let $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_m\}$ be an arbitrary partition of $[n]$. W.l.o.g, we assume for every \mathcal{S}_i and \mathcal{S}_j it holds that $|\mathcal{S}_i| = |\mathcal{S}_j| = \eta = n/m$. Then, we can have a DKG for the key structure $\{\mathcal{PK}, \mathcal{SK}\}$ over \mathcal{P} , by parallely invoking m CDSS instances, where each shard $\{P_i\}_{i \in \mathcal{S}_j}$ acts like a dealer consortium. We present the protocol description in Fig.1, while its initialization phase is CDSS.Init and its reconstruction algorithm is CDSS.Rec .

The security of our DKG inherently stems from the multi-instance security of the CDSS, given the fact that a least one shard maintains an honest majority.

Theorem 1. *The DKG protocol in Fig.1 is a secure (n, t) -DKG, which satisfies robustness and key-expressability against static adversaries, if the underlying CDSS is a secure (η, τ) -CDSS for some integer $\tau \geq \frac{\eta t}{n}$, satisfying multi-instance robustness and multi-instance key-expressability.*

Proof. The robustness of the DKG is implied by the multi-instance robustness of the underlying CDSS. Assume that the adversary \mathcal{A} corrupts $\{P_i\}_{i \in \mathcal{C}}$ for $|\mathcal{C}| = t$. By Lemma.3, there exists at least one $j^* \in [m]$, such that $|\{P_i\}_{i \in \mathcal{C}} \cap \mathcal{D}^{(j^*)}| \leq \tau$. Then, by multi-instance robustness, all honest participants should output \perp for all $j \in \mathbb{J}_1$ and output properly for all $j \in \mathbb{J}_2$; Moreover, \mathbb{J}_2 is non-empty, as $j^* \in \mathbb{J}_2$. For each $j \in \mathbb{J}_2$, we have $(pk_i^{(j)}, sk_i^{(j)}) \in \text{Rela}$ for $i \in [n]$. By the homomorphism of Rela , it follows $(sk_i = \bigoplus_{j \in \mathbb{J}_2} sk_i^j, pk_i = \bigotimes_{j \in \mathbb{J}_2} pk_i^{(j)}) \in \text{Rela}$. Then, we argue the secret key sk reconstructed from any subset of $t + 1$ secret shares will be identical and satisfy $(pk, sk) \in \text{Rela}$. For $\mathbb{I} \subset [n]$ s.t. $|\mathbb{I}| = t + 1$, the reconstruction algorithm of CDSS will determine $\{\alpha_i\}_{i \in \mathbb{I}}$. Then, the unique secret key $sk^{(j)}$ of j -th CDSS is $\bigoplus_{i \in \mathbb{I}} \alpha_i sk_i^{(j)}$ for every $j \in \mathbb{J}_2$. By description, the reconstruction algorithm of DKG is also CDSS.Rec , which means the final secret key reconstructed from $\{(i, sk_i)\}_{i \in \mathbb{I}}$ is $sk = \bigoplus_{i \in \mathbb{I}_1} \alpha_i (\bigoplus_{j \in \mathbb{J}_2} sk_i^{(j)}) = \bigoplus_{j \in \mathbb{J}_2} sk^{(j)}$, which is determined by $\{sk^{(j)}\}$ and independent of \mathbb{I} .

The key-expressability is implied by the multi-instance key-expressability of the CDSS. We proceed the proof by constructing the simulator $\text{Sim}_{\text{DKG}}^{\mathcal{A}}$ for any PPT adversary \mathcal{A} which corrupts

$\{P_i\}_{i \in \mathcal{C}}$ for some $|\mathcal{C}| = t$. $\text{Sim}_{\text{DKG}}^A$ directly invokes the simulator $\text{Sim}_{\text{CDSS}}^A$ and "aggregates" its outputs.

For clarity, we write the code of the simulator below.

```

 $\text{Sim}_{\text{DKG}}^A(pk)$ 


---


Invoke  $\text{Sim}_{\text{CDSS}}^A(pk) \rightarrow (\{\mathbf{tup}\}_{j \in [m]}, \mathbf{sview}_{\mathcal{A}})$ , s.t.  $\mathbf{tup}_j = (sk^{(j)}, pk^{(j)}, \alpha^{(j)})$  or  $\perp$ ,
Compute  $sk' = \bigoplus_{j \in \mathbb{J}} sk^{(j)}$ ,  $pk' = \bigotimes_{j \in \mathbb{J}} pk^{(j)}$ ,  $\alpha = \sum_{j \in \mathbb{J}} \alpha^{(j)}$ , for  $\mathbb{J} = \{j : \mathbf{tup}_j \neq \perp\}$ 
return  $(sk', pk', \alpha, \mathbf{sview}_{\mathcal{A}})$ 

```

By the definition of multi-instance key-expressability, no PPT distinguisher \mathcal{A}'' can distinguish $\{(\alpha^{(j)} \cdot pk \otimes pk^{(j)})_{j \in \mathbb{J}}, \mathbf{sview}_{\mathcal{A}}\}$ and $((pk^{(j)})_{j \in \mathbb{J}_2}, \mathbf{view}_{\mathcal{A}})$ from the real execution. Notice that

$$\alpha \cdot pk \otimes pk' = \bigoplus_{j \in \mathbb{J}} (\alpha^{(j)} \cdot pk \otimes pk^{(j)}) \quad \text{and} \quad \bigoplus_{j \in \mathbb{J}_2} (pk^{(j)})$$

are obtained by applying the same operations on the two tuples. Therefore, there is no PPT distinguisher \mathcal{A}' can distinguish

$$(\alpha \cdot pk \otimes pk', \mathbf{sview}_{\mathcal{A}}) \quad \text{and} \quad (\bigoplus_{j \in \mathbb{J}_2} (pk^{(j)}), \mathbf{view}_{\mathcal{A}})$$

with non-negligible advantage. Thus, the DKG satisfies key-expressability. \square

CDSS construction, intuitions. A straightforward yet non-succinct construction could let every dealer in the consortium \mathcal{D} run a complete secret sharing to all receivers. To reduce the communication cost, our idea is to let the receiver receive one "aggregated" and valid secret share, instead of to send multiple shares to be aggregated. In particular, we leverage the aggregatable PVSS, which enables us to delegate the share aggregation without harming the secrecy.

Assuming the dealer consortium has properly decided on one aggregated PVSS, the next step is to broadcast this aggregated transcript to all receivers. However, if we simply let every dealer broadcast the message, the communication cost will be blown up to $\Omega(\eta n^2 \lambda)$ again. In Section.4.2, we first introduce and construct a new broadcast primitive termed by *consortium-sender byzantine broadcast* (CSBB), which allows the dealer consortium to broadcast the aggregated message at the cost of $\Theta(n^2 \lambda)$. Then, in Section.4.3, we use this new broadcast primitive as a building block and construct our CDSS scheme.

4.2 Consortium-Sender Byzantine Broadcast

Definition. An (n, η, t, τ, ℓ) Consortium-Sender Byzantine Broadcast (CSBB) involves n participants $\mathcal{P} = \{P_1, \dots, P_n\}$ that has a subset $\mathcal{D} = \{D_1, \dots, D_\eta\} \subset \mathcal{P}$ acting as the sender consortium. The honest senders have the same ℓ -bit input \mathbf{msg} . We denote an instance of CSBB by $\text{CSBB}(\mathcal{D}(\mathbf{msg}), \mathcal{P})$. A CSBB is secure, if it satisfies the following properties against any PPT adversary \mathcal{A} corrupting up to t parties in \mathcal{P} .

- **Validity.** If all honest nodes in \mathcal{D} have the same valid input \mathbf{msg} , and the adversary corrupts up to τ parties in \mathcal{D} , all honest parties in \mathcal{P} output \mathbf{msg} .

- **Agreement.** All honest parties output the same message.
- **Termination.** All honest parties output a message.

A CSBB may have an initialization phase $\text{Init}(1^\lambda, n)$ for PKI setup and generating CRS. We require the CSBB remains all security properties even when there are polynomial many instances running **concurrently** in the unique identifier model (cf. Def.1) after the same initialization.

Building Blocks. We use a cryptographic accumulator Acc , an error correction code EC , and a Byzantine agreement BA as building blocks. Particularly, Acc provides a succinct representation of a set while ensuring efficient membership verification. It incorporates the algorithms Gen for accumulator key generation, Eval to accumulate a set \mathcal{S} into a value u , Wit to generate a witness w_i for an element $s_i \in \mathcal{S}$, and Vrfy to verify if s_i is in the set represented by u using w_i . EC includes deterministic algorithms Encode , which encodes a message into n code blocks $(c_i)_{i \in [n]}$, and Decode to reconstruct a message from any $t + 1$ code blocks. Formal definitions about the two primitives are recalled in the preliminaries.

Constructing CSBB. We give a construction for CSBB in Fig.2. A typical construction for BB is through the multicast-then-BA paradigm: BA guarantees all receiver output the same value. We follow a similar approach for CSBB, while making necessary changes to the multicast phase to keep it efficient. Particularly, multicasting a value of ℓ bits to a population of n nodes incurs $O(n\ell)$ communication cost; if all η senders in the consortium perform the multicast, the cost will be $O(\eta n\ell)$, not better than independently invoking BB for η times when $\ell = O(n\lambda)$. To reduce the communication cost, we utilize the erasure code [12] which is a common trick in distributed protocols. More specifically, we let each sender in the consortium deterministically encode the $O(n\lambda)$ -sized transcript into n blocks (c_1, \dots, c_n) each having $O(\lambda)$ bits, and send each c_i to P_i . P_i should receive multiple copies of c_i from the senders. However, it only multicasts the block which appears most frequently to all other receivers. By doing so, the communication cost in the phase becomes $O(n\ell)$ again. When the sender consortium has an honest majority, P_i will only relay the correct block of the message. We also follow recent works to use a cryptographic accumulator [50] to help decode the erasure code in the presence of up to $n/2$ malicious blocks, such that the receiver should reconstruct the correct message sent by the sender consortium.

Analysis. We analyze the performance and security of our CSBB. At round 1, each sender in \mathcal{D} sends out (u, c_j, w_j) whose size is $O(|w| + \ell/n)$ to every $P_j \in \mathcal{P}$. The communication cost of this round is $O(\eta(\ell + n|w|))$. At round 2, the cost is $O(n(\ell + n|w|))$. Adding them together with the cost of BA, we have the total cost of $O((n + \eta)(\ell + n|w|)) + \text{BA}_n(\ell)$. Regarding computation, each sender needs to generate n witness, and each receiver needs to verify $O(n)$ witnesses w.r.t. Acc . We assume without loss of generality that the per-node computation cost is $O(n)$ group operations.

Regarding security, at a high level, the properties of agreement and termination are derived directly from the security of the BA protocol. Concerning validity, each P_j in the second round will yield the code c_j corresponding to that input. As a consequence, all P_j participants in the third round can reconstruct the initial input message. By the validity of BA, these participants should output the original input. The concurrent security essentially follows our intuition that honest nodes can ignore messages with different identifiers. Formally, we have the following theorem.

Theorem 2. *The protocol in Fig.2 is a concurrently secure (n, η, t, τ, ℓ) -CSBB for any $t < \frac{n}{2}$ and $\tau < \frac{\eta}{2}$ in the unique identifier model, assuming the underlying accumulator scheme is secure, and the BA protocol is concurrently secure in the unique identifier model against adversary corrupting up to t nodes.*

Init ($1^\lambda, n$)	
$(\text{Sig}.vk_i, \text{Sig}.sk_i) \leftarrow \Sigma.\text{KeyGen}(1^\lambda)$, generate signing keys for every $P_i \in \mathcal{P}$ $\text{Acc.Gen}(1^\lambda, n) \rightarrow ak$ // Publish the accumulator key as a CRS.	
<hr/> CSBB ($\mathcal{D}(\text{msg}), \mathcal{P}$)	
<hr/> Round 1: each $D_i \in \mathcal{D}$ do $(c_1, \dots, c_n) \leftarrow \text{Encode}(\text{msg})$ $(u, \{w_j\}_{j \in [n]}) \leftarrow \text{Accumulate}(ak, (c_1, \dots, c_n))$ send (u, c_j, w_j) to $P_j \in \mathcal{P}, \forall j \in [n]$	
Round 2: each $P_j \in \mathcal{P}$ do receive: $\{(u^{(i)}, c_j^{(i)}, w_j^{(i)})\}$ from all S_i in round 1. if $\exists I \subset [n] \wedge I \geq \eta - \tau$, s.t. $(u^{(i)}, c_j^{(i)}, w_j^{(i)}) = (u^*, c_j^*, w_j^*) \forall i \in I$ send (c_j^*, w_j^*) to all $P_{j'} \in \mathcal{P}$, and store u^* else send \perp to all $P_{j'} \in \mathcal{P}$, and store $u^* = \perp$	
End of Round 2 each $P_j \in \mathcal{P}$ do receive: $\{(c_{j'}^*, w_{j'}^*)\}$ from all $P_{j'}$ in round 2. if $u^* = \perp$, then $\text{msg}_j = \perp$ else set $C = \emptyset$ for $j' \in [n]$ do if $\text{Acc.Vrfy}(ak, u^*, (c_{j'}^*), w_{j'}^*) = 1$, then $C = C \cup \{(j', c_{j'}^*)\}$ if $ C = t + 1$, then $\text{msg}_j \leftarrow \text{EC.Decode}(C)$, break .	
Round 3 to $2 + \Delta_{BA}$: // All participants $P_j \in \mathcal{P}$ run a BA to decide the output message $\text{BA}(P_j(\text{msg}_j)) \rightarrow \text{msg}$ each P_j outputs msg .	
<hr/> Algorithm $\text{Encode}(\text{msg}, n, t)$	<hr/> Algorithm $\text{Accumulate}(ak, (c_1, \dots, c_n))$
Divide msg into (m_1, \dots, m_t) $ m_i = m_j , \forall i, j \in [t]$ Compute (c_1, \dots, c_n) $\leftarrow \text{EC.Encode}(m_1, \dots, m_t)$, return (c_1, \dots, c_n)	Define $C = \{(i, c_i)\}_{i \in [n]}$ Compute $u \leftarrow \text{Acc.Eval}(ak, C)$ for $i \in [n]$ Compute $w_i \leftarrow \text{Acc.Wit}(ak, C, u, (i, c_i))$ return $(u, \{w_i\}_{i \in [n]})$

Fig. 2. The CSBB protocol. $\Delta_{BA(n)}$ denotes the number of rounds needed for the BA protocol for n participants. Σ is the signature scheme for authenticating messages.

We first give a general definition of concurrent security in the unique identifier model. Intuitively, a protocol could achieve concurrent security in the unique identifier model if an instance can aptly “ignore” messages with different identifiers, preserving its security as in the standalone setting. Though a message sent by P_i in an instance with identifier id can typically be crafted with access

to the signing oracle $\text{Sign}_{id}(sk_i, \cdot)$, we formalize the intuition as security against cross-instance signing queries (or CIS-Security).

Definition 2 (CIS security). *Protocol Π uses signature scheme Σ . In its variant Π_{id} , it employs Σ_{id} . If Π_{id} retains its security properties for any id , even when facing adversaries with signing oracles $\text{Sign}_{id'}(sk_i, \cdot)$ for any other id' not prefixed by id and any party's key sk_i , it's termed CIS-secure.*

Now we show our CSBB is *concurrently secure* in the unique identifier model. We start with the simple fact that our CSBB transcripts are perfectly simulatable with the help of a signing oracle, since during the protocol execution honest parties do not use any private input beyond the signing keys.

Lemma 4. *Let \mathcal{A} be a PPT adversary that corrupts an arbitrary number of nodes in \mathcal{D} and \mathcal{P} in an instance of CSBB with identifier id after $\text{Init}(1^\lambda, n)$. There is a simulator $S^{\mathcal{A}}((\text{Sig}.vk_i)_{i \in [n]})$ with access to signing oracles $\text{Sign}_{id}(\text{Sig}.sk_i, \cdot)$ for any $i \in [n]$, outputting the view $\text{view}_{\mathcal{A}}$ whose distribution is identical to the distribution of the view of \mathcal{A} in a real execution of this instance.*

Proof. We let the simulator $S^{\mathcal{A}}$ run CSBB with the adversary \mathcal{A} by acting on behalf of all honest nodes. Specifically, $S^{\mathcal{A}}$ follows the protocol specification to generate all protocol messages, which is feasible because no message in our CSBB protocol requires secret input. Before sending a message msg on behalf of an honest node P_i , $S^{\mathcal{A}}$ queries the oracle $\text{Sign}_{id}(\text{Sig}.sk_i, \cdot)$ with msg , and then sends out msg along with the signature. At the point of \mathcal{A} 's view, the execution simulated by $S^{\mathcal{A}}$ is identical to the real execution, and thus the distribution of the simulated view and that of the real view should be identical. \square

Then, we present a generic result that shows CIS-security (cf. Def. 2) will imply concurrent security, when the protocol transcripts can be perfectly simulatable by using access to signing oracles. It can be seen as a generalization to Lindell *et al.*'s result on BA protocols [45].

Lemma 5. *Let Π be a protocol that uses a signature scheme Σ in a blackbox manner. Let Π_{id} be a protocol which is obtained by replacing Σ with Σ_{id} . If for any PPT adversary \mathcal{A} , its view in an execution of Π_{id} can be perfectly simulated by a simulator $S^{\mathcal{A}}$ with access to signing oracles $\text{Sign}_{id}(\text{Sig}.sk_i, \cdot)$, and Π satisfies the CIS security, then Π maintains its security even when polynomially many instances are executed concurrently in the unique identifier model.*

Proof. Assuming there are m instances of Π running concurrently with prefix-free identifiers $\{id_1, \dots, id_m\}$, we argue the j -th instance maintains all the security properties, for an arbitrary $j \in [m]$. Since Π is CIS secure, Π_{id_j} shall be secure against any PPT adversary having access to signing oracles $\text{Sign}_{id_{j'}}()$ for $j' \neq j$. However, if there exists a PPT adversary \mathcal{A} that involves in all the m instances of Π and breaks the security of the j -th instance, we have a PPT adversary \mathcal{B} with access to signing oracles $\text{Sign}_{id_{j'}}()$ for $j' \neq j$ breaking the security of Π_{id_j} . The strategy of \mathcal{B} is simple: it invokes \mathcal{A} as a subroutine, forwards all messages between \mathcal{A} and honest parties in Π_{id_j} , and simulates all other instances by using the signing oracles. From \mathcal{A} 's point of view, the environment simulated by \mathcal{B} is identical to that of a real execution. Therefore, \mathcal{B} can break the security of Π_{id_j} if \mathcal{A} can break the security of j -th instance in the concurrent execution, which contradicts the CIS security. \square

By Lemma.5 and 4, it would be sufficient for showing its concurrent security by showing its CIS-security.

Theorem 3. *The protocol in Fig.2 is a CIS-secure (n, η, t, τ, ℓ) -CSBB for any $t < \frac{n}{2}$ and $\tau < \frac{\eta}{2}$, assuming the underlying accumulator scheme is secure, and the BA protocol is CIS-secure against adversary corrupting up to t nodes. Moreover, the communication complexity of the CSBB is*

$$O((n + \eta) \cdot (\ell + n \cdot |w|)) + BA_n(\ell),$$

where $|w|$ is the size of a membership witness in the accumulator scheme, and $BA_n(\ell)$ is the communication of BA among n participants with ℓ -bit inputs.

Proof. Termination and agreement follow directly from the underlying BA protocol. To see validity, we analyze the status after each round, where an adversary \mathcal{A} corrupts up to τ parties in \mathcal{D} and up to t parties in \mathcal{P} , and there are $\eta - \tau$ honest parties having the same input \mathbf{msg} . At the end of round 1, every honest P_j will receive the same (u^*, c_j^*, w_j^*) from the at least $\eta - \tau$ honest parties in \mathcal{D} , as Encode and Accumulate are deterministic, and thus will relay (c_j^*, w_j^*) to all other $P_i \in \mathcal{P}$. At the end of round 2, honest P_j receives $\{(c_i^*, w_i^*)\}$ from other P_i 's which contains at least $n - t$ honest pairs that pass the verification. On the other hand, by the unforgeability of the accumulator, if $\text{Acc.Vrfy}(ak, u^*, (i, c_i^*), w_i^*) = 1$, c_i^* must be the correct code of \mathbf{msg} at i -th position. Therefore, every honest party P_j should reconstruct the same message \mathbf{msg} . Then, by the validity of the underlying BA protocol, all honest parties output \mathbf{msg} .

Regarding CIS security, we let honest parties ignore any message that is invalid under the current identifier. The CIS security of CSBB then follows the CIS security of BA and the fact that an adversary cannot forge a signature under the current identifier by leveraging signing oracles under other identifiers.

We then analyze its communication cost. At round 1, each sender in \mathcal{D} sends out (u, c_j, w_j) whose size is $O(|w| + \ell/n)$ to every $P_j \in \mathcal{P}$. The communication cost of this round is $O(\eta(\ell + n|w|))$. At round 2, the cost is $O(n(\ell + n|w|))$. Adding them together with the cost of BA, we have the total cost of $O((n + \eta)(\ell + n|w|)) + BA_n(\ell)$. \square

4.3 The CDSS Construction

Recall our intuition about CDSS construction, the dealer consortium needs to agree on one aggregated PVSS before broadcasting it via our CSBB. For secrecy, it is crucial to ensure that the secret w.r.t. the aggregated PVSS remains unknown to an adversary which may corrupt τ out of η dealers and t out of n receivers. We guarantee the secrecy by letting each dealer generate an (n, t) -PVSS transcript under the public keys of receivers, broadcast to the *dealer consortium*, and then aggregate all valid PVSS transcripts. In particular, note that the final secret key is the sum of secret keys shared by all dealers, which means the adversary cannot know the final secret key unless it corrupts all dealers. Meanwhile, by the definition of PVSS, an adversary corrupting t receivers cannot learn the secret key from the decrypted shares. Moreover, as each dealer sends its PVSS transcript via a BB protocol, it ensures all dealers have the same view of valid transcripts and thus obtain the same aggregated transcript.

Formally, assume an aggregatable PVSS for $(\mathcal{PK}, \mathcal{SK})$, a Byzantine Broadcast protocol BB, and a CSBB protocol CSBB. We delineate the deal protocol of CDSS in Fig.3, while its reconstruction algorithm is the same as PVSS.Rec. The initialization algorithm $\text{Init}(1^\lambda, n)$ is as follows: It invokes


```

Deal( $\mathcal{D}, \mathcal{P}$ )


---


Round 1 to  $\Delta_{\text{BB}(\eta)}$ : each  $D_j \in \mathcal{D}$  do
  PVSS.Deal( $(ek_i)_{i \in [n]}, \text{cid}_{D_j}$ )  $\rightarrow$  ( $\text{Trans}_j, sk_j$ )
  broadcast: BB( $D_j(\text{Trans}_j), \mathcal{D}$ ) with an unique identifier  $id_j^{\text{BB}}$ 
End of Round  $\Delta_{\text{BB}(\eta)}$ : each  $D_j \in \mathcal{D}$ 
  receive broadcast messages  $\{\text{Trans}_{j'}\}_{j' \in [\eta]}$  from all  $D_{j'} \in \mathcal{D}$ 
  set TRANS =  $\emptyset$ 
  for  $j' \in [\eta]$ 
    if PVSS.PubVrfy( $(ek_i)_{i \in [n]}, \text{Trans}_{j'}$ ) = 1 and PVSS.getCID( $\text{Trans}_{j'}$ ) =  $\{\text{cid}_{D_{j'}}\}$ 
      then TRANS = TRANS  $\cup$   $\{\text{Trans}_{j'}\}$ 
  PVSS.Agg(TRANS,  $(ek_i)_{i \in [n]}$ )  $\rightarrow$  Trans
Round  $\Delta_{\text{BB}(\eta)} + 1$  to  $\Delta_{\text{BB}(\eta)} + \Delta_{\text{CSBB}(\eta, n)}$ :
// The dealer consortium  $\mathcal{D}$  broadcasts the aggregated transcript Trans to all receivers  $\mathcal{P}$  via CSBB
CSBB( $\mathcal{D}(\text{Trans}), \mathcal{P}$ )
End of Round  $\Delta_{\text{BB}(\eta)} + \Delta_{\text{CSBB}(\eta, n)}$ : each  $P_i$  do
  receive broadcast message Trans from all  $\mathcal{D}$ , and let CID  $\leftarrow$  PVSS.getCID(Trans)
  if PVSS.PubVerify( $(ek_i)_{i \in [n]}, \text{Trans}$ ) = 0  $\vee$   $\neg(\text{CID} \subset \{\text{cid}_{D_j}\}_{j \in [\eta]} \wedge |\text{CID}| \geq \tau + 1)$ 
    then output  $\perp$ 
  else  $(pk, pk_1, \dots, pk_n) \leftarrow$  PVSS.PubDerive(Trans),  $sk_i \leftarrow$  PVSS.Dec( $ek_i, dk_i, \text{Trans}$ )
    output  $(pk, pk_1, \dots, pk_n, sk_i)$ 

```

Fig. 3. Deal(\mathcal{D}, \mathcal{P}) Protocol of complete CDSS. $\Delta_{\text{BB}(\eta)}$ (or $\Delta_{\text{CSBB}(\eta, n)}$) is the number of rounds need for running BB with n parties (or CSBB with η senders and n receivers). We assume every P_i has its publicly known CID, denoted by cid_{P_i} .

CSBB.Init($1^\lambda, n$), which includes a PKI setup for a digital signature scheme Σ , and PVSS.Init($1^\lambda, n$) which generates $(ek_i, dk_i)_{i \in [n]}$ for the PVSS scheme.

4.4 Analysis

Communication complexity. We first analyze the communication complexity of our CDSS construction in Fig.3. All η dealers broadcast a PVSS transcript of size $O(n\lambda)$ bits, which incurs bit complexity of $\eta \text{BB}_\eta(n\lambda)$ in total; and dealers and receivers invoke a CSBB protocol to disseminate the aggregated PVSS transcript, which incurs bit complexity of $O(n^2 \cdot w) + \text{BA}_n(n\lambda)$, assuming $\eta < n$ and the witness size w . The communication complexity of CDSS is

$$O(n^2 \cdot w) + \eta \text{BB}_\eta(n\lambda) + \text{BA}_n(n\lambda), \quad (6)$$

where $\text{BB}_z(\ell)$ (or $\text{BA}_z(\ell)$) is the communication cost of Byzantine Broadcast BB (or Byzantien Agreement BA) among z participants with ℓ -bit input.

Computation complexity. In our design, each dealer creates one PVSS transcript, verifies η transcripts, and aggregates η transcripts; each receiver verifies one transcript. They invoke CSBB once, costing $O(n)$ group operations per node. With the PVSS scheme in Appendix.B, the per-node computation cost is $O(n)$.

On complexity of the DKG. The bit communication complexity of the DKG (Fig.1) is equal to m (the number of shards) times the complexity of the CDSS construction. Therefore, with CDSS in Fig.3, the bit communication complexity of our DKG is $O(mn^2 \cdot |w|) + n\text{BB}_\eta(n\lambda) + m\text{BA}_n(\lambda)$, while $n = \eta m$.

Now we discuss the best sharding parameters for the smallest communication. Assuming we are using the optimal BA and BB, *i.e.*, $\text{BA}_z(\ell) = \text{BB}_z(\ell) = O(z\ell + z^2\lambda)$, and the accumulator with witness size $|w| = O(\lambda)$, we notice that $\eta = m = \sqrt{n}$ yields a communication cost of DKG which is $O(n^{2.5}\lambda)$. Regarding computation cost, with the PVSS in Appendix.B, the per-node computation cost of the DKG is $O(n^{1.5})$ group operations.

Security analysis. We establish the security of our CDSS in the following.

Theorem 4. *Assuming that the underlying PVSS is secure, and BB and CSBB are concurrently secure in the unique identifier model, the CDSS protocol in Fig.3 satisfies the multi-instance robustness and the multi-instance key-expressability.*

We prove the multi-instance robustness in Lemma.6 and the multi-instance key-expressability in Lemma.7

Lemma 6. *The CDSS protocol satisfies multi-instance robustness, assuming concurrent security of BB and CSBB, and correctness and soundness of PVSS.*

Proof. First, we argue that for any instance j , all honest nodes either return \perp or output properly, *i.e.*, they have the same view of public information $(pk^{(j)}, (pk_i)_{i \in [n]}^{(j)})$ and obtain a correct secret share, despite there is a PPT adversary \mathcal{A} corrupting $\{P_i\}_{i \in \mathcal{C}}$ for $|\mathcal{C}| \leq t$. By agreement of CSBB, all honest receivers P_i 's will receive the same message and will return \perp if the message is not a valid PVSS transcript or its CID is not consistent with its dealer consortium. Then, by the soundness of PVSS, when the PVSS transcript is valid, every honest receiver can obtain the correct secret share by decrypting the transcript.

Then, we show that for the instance j where $|\{P_i\}_{i \in \mathcal{C}} \cap \mathcal{D}^{(j)}| \leq \tau$, the honest parties must output properly. Since at most τ nodes in $\mathcal{D}^{(j)}$ are corrupted, the BB instances within \mathcal{D} have all the security guarantees. By the agreement of BB, all honest nodes in \mathcal{D} will receive the message from each instance of BB. Moreover, by the correctness of PVSS and the validity of BB, an honest node in \mathcal{D} broadcasts a valid PVSS transcript that will be received by all honest nodes in \mathcal{D} . Therefore, all honest nodes in \mathcal{D} can receive at least $\eta - \tau$ valid PVSS transcripts and can aggregate them into one. By the *validity* of CSBB, all nodes in \mathcal{P} receive the same valid transcript. Then by the soundness of PVSS, all P_i 's can obtain a valid secret share and derive the same public key (shares).

Lemma 7. *The CDSS protocol satisfies the multi-instance key-expressability, assuming the concurrent security of BB and CSBB, and the correctness, soundness, secrecy, and simulation soundness of PVSS.*

Proof. We proceed with this proof by constructing the simulator algorithm $\text{Sim}_{\text{CDSS}}^A$ which leverages the simulator algorithms $\{\text{PVSS.SInit}, \text{PVSS.SDeal}, \text{PVSS.SRec}\}$ of the PVSS. For clarity, we write down the pseudo-code of $\text{Sim}_{\text{CDSS}}^A$ below, which takes as input pk which is generated by the default key generation algorithm $\text{KeyGen}(1^\lambda) \rightarrow (pk, sk)$.

<p>$\text{Sim}_{\text{CDSS}}^{\mathcal{A}}(pk)$</p> <hr/> <p>Initialize: $\text{CSBB.Init}_{\mathcal{A}}(1^\lambda, n)$ and $\text{PVSS.SInit}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \mathbf{tk})$</p> <p>Run the m instances $\{\text{Deal}_{id_j}^{\mathcal{A}}\}_{j \in [m]}$ with \mathcal{A}, as below:</p> <p>Select $j^* \in [m]$ and $i^* \in [n]$ s.t. $\mathcal{D}^{(j^*)} \cap \{P_i\}_{i \in \mathcal{C}} \leq \tau$, and $D_{i^*}^{(j^*)} \in \mathcal{D}^{(j^*)} \setminus \{P_i\}_{i \in \mathcal{C}}$</p> <p>Round 1 to $\Delta_{\text{BB}(\eta)}$:</p> <p style="padding-left: 20px;">Run $\text{SDeal}((ek_i)_{i \in [n]}, pk, \mathbf{tk}, D_{i^*}^{(j^*)}) \rightarrow \text{Trans}^*, \text{BB}(D_{i^*}^{(j^*)}(\text{simTrans}), \mathcal{D}^{(j^*)})$</p> <p style="padding-left: 20px;">Honestly execute the code for every $D_i^{(j)} \notin \{P_i\}_{i \in \mathcal{C}}$ and $(i, j) \neq (i^*, j^*)$</p> <p style="padding-left: 20px;">Honestly execute the remaining rounds, besides performing the tasks below:</p> <p style="padding-left: 20px;">At the end of round $\Delta_{\text{BB}(\eta)}$</p> <p style="padding-left: 40px;">Let TRANS^* be the set of valid PVSS transcripts in $\mathcal{D}^{(j^*)}$</p> <p style="padding-left: 40px;">for $\text{Trans}_i \in \text{TRANS}^*$ and $\text{Trans}_i \neq \text{Trans}^*$</p> <p style="padding-left: 80px;">$\text{PVSS.SRec}(\mathbf{tk}, \text{Trans}_i) \rightarrow sk'_i$, and $\text{PVSS.PubDriv}(\text{Trans}_i) \rightarrow pk'_i$</p> <p style="padding-left: 40px;">Let $pk'^{(j^*)} = \bigotimes pk'_i, sk'^{(j^*)} = \bigoplus sk'_i, \alpha^{(j^*)} = 1$</p> <p style="padding-left: 20px;">At the end of round $\Delta_{\text{BB}(\eta)} + \Delta_{\text{CSBB}(\eta, n)} + 1$</p> <p style="padding-left: 40px;">for $j \in [m]$ and $j \neq j^*$</p> <p style="padding-left: 80px;">Receive $\text{Trans}^{(j)}$ from CSBB by $\mathcal{D}^{(j)}, \forall j \in [m]; \text{CID}^{(j)} \leftarrow \text{PVSS.getCID}(\text{Trans}^{(j)})$</p> <p style="padding-left: 80px;">if $\text{PVSS.PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}^{(j)}) = 1 \wedge \text{CID}^{(j)} \cap \{\text{cid}_{D_i^{(j)}}\}_{i \in [n]} \geq \tau + 1$</p> <p style="padding-left: 120px;">then $\text{PVSS.SRec}(\mathbf{tk}, \text{Trans}^{(j)}) \rightarrow sk'^{(j)}, \text{PVSS.PubDriv} \rightarrow (pk'^{(j)}, \cdot), \alpha^{(j)} = 0$</p> <p style="padding-left: 120px;">$\text{tup}_j = (sk'^{(j)}, pk'^{(j)}, \alpha^{(j)});$ else $\text{tup}_j = \perp$</p> <p>return $(\{\text{tup}_j\}_{j \in [m]}, \text{view}_{\mathcal{A}})$.</p>
--

Recall the multi-instance key-expressability definition, the PPT distinguisher \mathcal{A}' is required to distinguish $(\{\alpha^{(j)} \cdot pk \otimes pk'^{(j)}\}_{j \in \mathbb{J}'}, \text{view}_{\mathcal{A}})$ (provided by the above simulator) and $(\{pk^{(j)}\}_{j \in \mathbb{J}}, \text{view}_{\mathcal{A}})$ (from a real execution), where $\mathbb{J}' = \{j \in [m] : \text{Trans}_j \text{ is valid}\}$ in the simulated execution and $\mathbb{J} = \{j \in [m], \text{Trans}_j \text{ is valid}\}$ in the real execution. Note that $\{\alpha^{(j)} \cdot pk \otimes pk'^{(j)} = pk^{(j)}\}_{j \in \mathbb{J}'}$ can be derived from \mathcal{A}' 's view. At the point of \mathcal{A}' 's view, the only difference between the execution simulated by $\text{Sim}_{\text{CDSS}}^{\mathcal{A}}$ and the real execution is about how PVSS.Init is executed and how Trans^* is generated. By the secrecy of PVSS, the distinguisher \mathcal{A}' cannot distinguish the two tuples with a non-negligible advantage. Note that the CID of $\text{Trans}^{(j)}$ does not contain $\text{cid}_{D_{i^*}^{(j^*)}}$ for $j \neq j^*$.

By the *simulation soundness* of PVSS, PVSS.SRec can obtain $sk^{(j)}$ which is the valid secret key of $pk^{(j)}$. Moreover, we have $\alpha = \alpha^{(j^*)} = 1$. \square

5 DKG with Field-Element Secret

In DLog-based cryptography, most conventional cryptographic schemes possess a secret key within the finite field \mathbb{Z}_p . Our earlier construction utilized aggregatable PVSS, which is primarily aligned with group-element secrets (unless using generic zkSNARK [39]). In this section, we present a CDSS construction that can be built upon a conventional PVSS without aggregation, which could naturally give rise to a subcubic DKG for field-element secrets.

CDSS For Field-Element Secret. The aggregatable PVSS in the CDSS construction in Sect.4.3 is employed for reducing communication cost while maintaining security. Particularly, in comparison

with a naive solution where the consortium should broadcast all valid PVSS transcripts, aggregating them effectively reduces the size of the broadcast message from $O(n\eta\lambda)$ to $O(n\lambda)$ (assuming n the number of receivers and η is the consortium size), which is vital for achieve subcubic DKG. On the other hand, if we let the consortium just broadcast one of these PVSS transcripts, the corresponding secret key may be known to the adversary, while the secret key of an aggregated transcript is guaranteed to be hidden whenever it has a contribution from an honest party.

In this section, we introduce a different path for achieving these goals, by utilizing a common coin “within” the dealer consortium. Like the CDSS in Sect.4.3, all dealers first broadcast their PVSS transcripts within the consortium. However, after this step, the dealers do not aggregate these valid transcripts; Instead, they invoke a common coin to randomly pick κ PVSS transcripts, which will be sent out together via the CSBB protocol. Here, κ is the statistical security parameter that is independent of n or η . A receiver can obtain its share by decrypting all the κ transcripts and adding the decrypted values together. With a high probability $(1 - (\frac{\tau}{\eta})^\kappa)$, at least one of the selected transcripts is from an honest dealer, which guarantees the secrecy of this scheme. Particularly, we can implement the coin protocol by letting the consortium first run a DKG protocol, then reconstruct the secret key, and finally apply a hash function (which is modeled as a random oracle) to the secret key.

Formally, assume a PVSS for $(\mathcal{PK}, \mathcal{SK})$, a Byzantine Broadcast protocol BB, a CSBB protocol CSBB, a DKG protocol DKG^{unp} which we only assumes its secret key is unpredictable, and a hash function Hash which maps $\{0, 1\}^*$ to κ indexes in the range of $[\eta]$. We elucidate the deal phase of CDSS scheme in Fig.4, while its reconstruction algorithm is the same as PVSS.Rec. The initialization algorithm $\text{Init}(1^\lambda, n)$ is as follows: It invokes $\text{CSBB.Init}(1^\lambda, n)$, which includes a PKI setup for a digital signature scheme Σ , $\text{PVSS.Init}(1^\lambda, n)$ which generates $(ek_i, dk_i)_{i \in [n]}$ for the PVSS scheme, and the setup for DKG^{unp} .

Security analysis. We establish the security of the CDSS scheme in Fig.4 in the following.

Theorem 5. *Assuming that the underlying PVSS is secure, the DKG^{unp} is unpredictable and robust, and BB and CSBB are concurrently secure in the unique identifier model, the CDSS in Fig.4 satisfies the multi-instance robustness and the multi-instance key-expressability in the random oracle model.*

The proof largely resembles to the proof for Theorem 4, except that we need to leverage the fact that the selected κ transcripts include one from an honest dealer. We discuss the multi-instance robustness in Lemma 8 and the multi-instance key-expressability in Lemma 9, respectively.

Lemma 8. *The CDSS protocol satisfies multi-instance robustness in the random oracle model, assuming concurrent security of BB and CSBB, the correctness and soundness of PVSS, and that the DKG^{unp} is unpredictable and robust.*

Proof. The agreement of CSBB ensures that all honest receivers P_i 's will receive the same message, which could be a PVSS transcript or \perp . The soundness of PVSS ensures that when the PVSS transcript is valid, every honest receiver can obtain the correct share by decrypting the transcript. Therefore, we have that for any instance j , all honest nodes either return \perp or output properly.

We then argue that for the instance j where $|\{P_j\}_{j \in \mathcal{C}} \cap \mathcal{D}^{(j)}| \leq \tau$, the honest parties output properly. By the robustness of DKG^{unp} , the set of dealers $\mathcal{D}^{(j)}$ will obtain a set of consistent secret shares at the end of Round $\Delta_{\text{DKG}(\eta)}$. By the unpredictability of DKG^{unp} , the output of Hash is

```

Deal( $\mathcal{D}, \mathcal{P}$ )
Round 1 to  $\Delta_{\text{DKG}(\eta)}$  //  $\mathcal{D}$  run a DKG protocol, and each  $D_i$  obtains a secret key share  $sk_{\mathcal{D},i}$ 
   $\text{DKG}^{\text{unp}}(\mathcal{D}) \rightarrow (pk_{\mathcal{D}}, sk_{\mathcal{D},1}, \dots, sk_{\mathcal{D},\eta})$ 
Round 1 +  $\Delta_{\text{DKG}(\eta)}$  to  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$ : each  $D_j \in \mathcal{D}$  do
   $\text{PVSS.Deal}((ek_i)_{i \in [n]}, \text{cid}_{D_j}) \rightarrow (\text{Trans}_j, sk_j)$ 
  broadcast:  $\text{BB}(D_j(\text{Trans}_j), \mathcal{D})$  with a unique identifier  $id_j^{\text{BB}}$ 
Round  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  + 1 to  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  + 2: each  $D_i \in \mathcal{D}$  do
  receive broadcast messages  $\{\text{Trans}_{j'}\}_{j' \in [\eta]}$  from all  $D_{j'} \in \mathcal{D}$ 
  set  $\text{TRANS} = \emptyset$ 
  for  $j' \in [\eta]$ 
    if  $\text{PVSS.PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}_{j'}) = 1$  and  $\text{PVSS.getCID}(\text{Trans}_{j'}) = \{\text{cid}_{D_{j'}}\}$ 
      then  $\text{TRANS} = \text{TRANS} \cup \{\text{Trans}_{j'}\}$ 
    send  $sk_{\mathcal{D},i}$  to all  $D_j \in \mathcal{D}$ 
End of Round  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  + 2:
  receive secret shares  $\{sk_{\mathcal{D},j}\}$ , reconstruct  $sk_{\mathcal{D}}$ , and compute  $\text{Hash}(sk_{\mathcal{D}}) \rightarrow (j_1, \dots, j_\kappa)$ 
  set  $\text{OutTRANS} = \{\text{Trans}_j \in \text{TRANS} : j \in \{j_1, \dots, j_\kappa\}\}$ 
Round  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  + 3 to  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  +  $\Delta_{\text{CSBB}(\eta,n)}$  + 2:
  // The dealer consortium  $\mathcal{D}$  broadcasts the selected transcript set  $\text{OutTRANS}$  to all receivers  $\mathcal{P}$  via CSBB
   $\text{CSBB}(\mathcal{D}(\text{OutTRANS}), \mathcal{P})$ 
End of Round  $\Delta_{\text{BB}(\eta)}$  +  $\Delta_{\text{DKG}(\eta)}$  +  $\Delta_{\text{CSBB}(\eta,n)}$  + 2: each  $P_i$  do
  receive broadcast message  $\text{OutTRANS}$  from all  $\mathcal{D}$ , and parse  $\text{OutTRANS} = (\text{Trans}_{j_1}, \dots, \text{Trans}_{j_\kappa})$ 
  if  $\exists \text{Trans}_{j_z} \in \text{OutTRANS}$ , s.t.,  $\text{PVSS.PubVerify}((ek_i)_{i \in [n]}, \text{Trans}_{j_z}) = 0 \vee \text{PVSS.getCID}(\text{Trans}_{j_z}) \not\subseteq \{\text{cid}_{D_j}\}_{j \in [\eta]}$ 
    then output  $\perp$ 
  for  $\text{Trans}_{j_z} \in \text{OutTRANS}$  do
     $(pk^{(j_z)}, pk_1^{(j_z)}, \dots, pk_n^{(j_z)}) \leftarrow \text{PVSS.PubDerive}(\text{Trans}_{j_z}), sk_i^{(j_z)} \leftarrow \text{PVSS.Dec}(ek_i, dk_i, \text{Trans}_{j_z})$ 
  output  $(pk = \bigotimes_{i \in [\kappa]} pk^{(j_z)}, pk_1 = \bigotimes_{i \in [\kappa]} pk_1^{(j_z)}, \dots, pk_n = \bigotimes_{i \in [\kappa]} pk_n^{(j_z)}, sk_i = \bigoplus_{i \in [\kappa]} sk_i^{(j_z)})$ 

```

Fig. 4. Deal(\mathcal{D}, \mathcal{P}) Protocol of CDSS for field-element secrets. $\Delta_{\text{BB}(\eta)}$ (or $\Delta_{\text{CSBB}(\eta,n)}$, or $\Delta_{\text{DKG}(\eta)}$) is the number of rounds need for running BB with n parties (or CSBB with η senders and n receivers, or DKG with η parties). We assume every P_i has its publicly known CID, denoted by cid_{P_i} .

distinguishable with κ uniformly sampled indexes from $[\eta]$. Therefore, the probability of no honest dealer being selected is bounded by $(\frac{\tau}{\eta})^\kappa + \text{negl}(\lambda)$, which is negligibly small. By the validity of BB, the selected honest node should have correctly broadcasted a valid PVSS to the dealer consortium. It follows that the dealer consortium will least broadcast one valid PVSS transcript, such that all receivers must output properly for the dealer consortium. \square

Lemma 9. *The CDSS protocol satisfies multi-instance key-expressability in the random oracle model, assuming the concurrent security of BB and CSBB, the correctness, soundness, secrecy, and simulation soundness of PVSS, and the robustness and unpredictability of DKG^{unp} .*

Proof. We proceed with this proof by constructing the simulator algorithm $\text{Sim}_{\text{CDSS}}^A$ which leverages the simulator algorithms $\{\text{PVSS.SInit}, \text{PVSS.SDeal}, \text{PVSS.SRec}\}$ of the PVSS. For clarity, we write

down the pseudo-code of $\text{Sim}_{\text{CDSS}}^{\mathcal{A}}(pk)$ below, which takes as input pk which is generated by the default key generation algorithm $\text{KeyGen}(1^\lambda) \rightarrow (pk, sk)$.

Sim_{CDSS}^A(pk)

Initialize: $\text{CSBB.Init}_{\mathcal{A}}(1^\lambda, n)$, $\text{PVSS.SInit}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{tk})$, and setup for DKG^{unp}

Run the m instances $\{\text{Deal}_{id_j}^{\mathcal{A}}\}_{j \in [m]}$ with \mathcal{A} , as below:

Round 1 to $\Delta_{\text{DKG}(\eta)}$:

Honestly execute the DKG protocols on behalf of honest nodes

Round 1 + $\Delta_{\text{DKG}(\eta)}$ to $\Delta_{\text{BB}(\eta)}$ + $\Delta_{\text{DKG}(\eta)}$

Select $j^* \in [m]$ s.t. $|\mathcal{D}^{(j^*)} \cap \{P_i\}_{i \in \mathcal{C}}| \leq \tau$

Reconstruct the secret key $sk_{\mathcal{D}^{(j^*)}}$ using the honest parties's shares, and compute $\text{Hash}(sk_{\mathcal{D}^{(j^*)}}) \rightarrow \{j_1, \dots, j_\kappa\}$

Select $i^* \in \{j_1, \dots, j_\kappa\}$ s.t. $D_{i^*}^{(j^*)} \in \mathcal{D}^{(j^*)} \setminus \{P_i\}_{i \in \mathcal{C}}$ // It can find such i^* with overwhelming probability

Round 1 to $\Delta_{\text{BB}(\eta)}$:

Run $\text{SDeal}((ek_i)_{i \in [n]}, pk, \text{tk}, D_{i^*}^{(j^*)}) \rightarrow \text{Trans}^*, \text{BB}(D_{i^*}^{(j^*)})(\text{simTrans}), \mathcal{D}^{(j^*)}$

Honestly execute the code for every $D_i^{(j)} \notin \{P_i\}_{i \in \mathcal{C}}$ and $(i, j) \neq (i^*, j^*)$

Honestly execute the remaining rounds, besides performing the tasks below:

At the end of round $\Delta_{\text{BB}(\eta)}$

Let TRANS^* be the set of valid PVSS transcripts produced by $D_{j_1}, \dots, D_{j_\kappa} \in \mathcal{D}^{(j^*)}$

for $\text{Trans}_i \in \text{TRANS}^*$ and $\text{Trans}_i \neq \text{Trans}^*$

$\text{PVSS.SRec}(\text{tk}, \text{Trans}_i) \rightarrow sk'_i$, and $\text{PVSS.PubDriv}(\text{Trans}_i) \rightarrow pk'_i$

Let $pk'^{(j^*)} = \bigotimes pk'_i, sk'^{(j^*)} = \bigoplus sk'_i, \alpha^{(j^*)} = 1$

At the end of round $\Delta_{\text{BB}(\eta)} + \Delta_{\text{CSBB}(\eta, n)} + 1$

for $j \in [m]$ and $j \neq j^*$

Receive $\text{OutTrans}^{(j)}$ from CSBB by $\mathcal{D}^{(j)}$, and parse $\text{OutTrans}^{(j)} = \{\text{Trans}_{i_1}^{(j)}, \dots, \text{Trans}_{i_z}^{(j)}\}, \forall j \in [m]$;

if $\exists i' \in \{i_1, \dots, i_z\}$ s.t. $\text{PVSS.PubVrfy}((ek_i)_{i \in [n]}, \text{Trans}_{i'}^{(j)}) = 0 \vee \text{PVSS.getCID}(\text{Trans}_{i'}^{(j)}) \notin \{\text{cid}_{D_i^{(j)}}\}_{i \in [n]}$

continue

else

then $\text{PVSS.SRec}(\text{tk}, \text{Trans}_{i_a}^{(j)}) \rightarrow sk'_a{}^{(j)}, \text{PVSS.PubDriv} \rightarrow (pk'_a{}^{(j)}, \cdot), \alpha_a^{(j)} = 0, \forall a \in [z]$

tup_j = $(\bigoplus_{a \in [z]} sk'_a{}^{(j)}, \bigotimes_{a \in [z]} pk'_a{}^{(j)}, \alpha^{(j)} = \sum_{a \in [z]} \alpha_a^{(j)})$; **else** **tup_j** = \perp

return $(\{\text{tup}_j\}_{j \in [m]}, \text{view}_{\mathcal{A}})$.

Recall the multi-instance key-expressability definition, the PPT distinguisher \mathcal{A}' is required to distinguish $(\{\alpha^{(j)} \cdot pk \otimes pk'^{(j)}\}_{j \in \mathbb{J}'}, \text{view}_{\mathcal{A}})$ (provided by the above simulator) and $(\{pk^{(j)}\}_{j \in \mathbb{J}}, \text{view}_{\mathcal{A}})$ (from a real execution), where $\mathbb{J}' = \{j \in [m] : \text{Trans}_j \text{ is valid}\}$ in the simulated execution and $\mathbb{J} = \{j \in [m], \text{Trans}_j \text{ is valid}\}$ in the real execution. By the robustness and unpredictability of DKG^{unp} , the simulator $\text{Sim}_{\text{CDSS}}^{\mathcal{A}}$ can finish the above simulation with an overwhelming probability. The following arguments are similar to that for Lemma 7. Particularly, under the condition that the simulator could finish the above simulation, at the point of \mathcal{A}' 's view, the only difference between the execution simulated by $\text{Sim}_{\text{CDSS}}^{\mathcal{A}}$ and the real execution is about how PVSS.Init is executed and how Trans^* is generated. By the secrecy of PVSS, the distinguisher \mathcal{A}' cannot distinguish the two tuples with a non-negligible advantage. Note that the CID of $\text{Trans}^{(j)}$ does not contain $\text{cid}_{D_{i^*}^{(j^*)}}$ for

$j \neq j^*$. By the *simulation soundness* of PVSS, PVSS.SRec can obtain $sk^{(j)}$ which is the valid secret key of $pk^{(j)}$. Moreover, we have $\alpha = \alpha^{(j^*)} = 1$. \square

Performance analysis. With the PVSS scheme outlined in Lemma 1, the PVSS transcript size is $\mathcal{O}(n\lambda)$. Therefore, the communication cost incurred by all η dealers broadcast their PVSS transcripts within the consortium is $\eta\text{BB}_\eta(n\lambda)$. The communication cost for broadcasting κ PVSS transcripts via CSBB is $\mathcal{O}(n^2\lambda\kappa \cdot |w|) + \text{BA}_n(n\lambda\kappa)$, where $|w|$ is size of an accumulator witness. With optimal BA/BB and an accumulator with constant-sized witness, the above communication cost will be $\mathcal{O}(n^2\lambda\kappa)$. While the dealer consortium needs to perform a DKG within the consortium, the communication cost for the DKG is at most $\mathcal{O}(\eta^3\lambda)$ without using the sub-cubic DKG in Sect.4, which is not the dominating term for either communication or computation cost. For computation cost, since each receiver needs to process κ transcripts, the per-node computation cost is $\mathcal{O}(n\kappa)$ group operations.

Application to DKG. As the CDSS in Fig.4 satisfies both multi-instance robustness and multi-instance key-expressability, it can be plugged into the DKG construction in Fig.1. The security of the resulting DKG follows Theorem 1. Regarding performance, the DKG parallelly invokes $m = \sqrt{n}$ instances of CDSS, and thus the total communication complexity will be $\mathcal{O}(n^{2.5}\lambda\kappa)$, and the per-node computation cost will be $\mathcal{O}(n^{1.5}\kappa)$ group operations.

References

1. drand. <https://drand.love/>
2. Abraham, I., Chan, T.H., Dolev, D., Nayak, K., Pass, R., Ren, L., Shi, E.: Communication complexity of byzantine agreement, revisited. *Distributed Comput.* **36**(1), 3–28 (2023)
3. Abraham, I., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G.: Bingo: Adaptivity and asynchrony in verifiable secret sharing and distributed key generation. In: *CRYPTO (1)*. Lecture Notes in Computer Science, vol. 14081, pp. 39–70. Springer (2023)
4. Attema, T., Cramer, R., Rambaud, M.: Compressed Σ -protocols for bilinear group arithmetic circuits and application to logarithmic transparent threshold signatures. In: *ASIACRYPT (4)*. Lecture Notes in Computer Science, vol. 13093, pp. 526–556. Springer (2021)
5. Azouvi, S., Vukolic, M.: Pikachu: Securing pos blockchains from long-range attacks by checkpointing into bitcoin pow using taproot. In: *Proceedings of the 2022 ACM Workshop on Developments in Consensus*. pp. 53–65 (2022)
6. Bacho, R., Lenzen, C., Loss, J., Ochsenreither, S., Papachristoudis, D.: Grandline: Adaptively secure DKG and randomness beacon with (almost) quadratic communication complexity. Tech. rep., Cryptology ePrint Archive (2023), <https://eprint.iacr.org/2023/1887>
7. Bacho, R., Loss, J.: On the adaptive security of the threshold BLS signature scheme. In: *CCS*. pp. 193–207. ACM (2022)
8. Bacho, R., Loss, J.: Adaptively secure (aggregatable) pvss and application to distributed randomness beacons. In: *CCS*. ACM (2023)
9. Benhamouda, F., Halevi, S., Krawczyk, H., Miao, A., Rabin, T.: Threshold cryptography as a service (in the multiserver and YOSO models). In: *CCS*. pp. 323–336. ACM (2022)
10. Bhat, A., Kate, A., Nayak, K., Shrestha, N.: Optrand: Optimistically responsive distributed random beacons. In: *NDSS* (2023)
11. Bhat, A., Shrestha, N., Luo, Z., Kate, A., Nayak, K.: Randpiper - reconfiguration-friendly random beacons with quadratic communication. In: *CCS*. pp. 3502–3524. ACM (2021)
12. Blahut, R.E.: Theory and practice of error control codes. Addison-Wesley (1983)
13. Boneh, D., Bonneau, J., Büinz, B., Fisch, B.: Verifiable delay functions. In: *CRYPTO (1)*. Lecture Notes in Computer Science, vol. 10991, pp. 757–788. Springer (2018)
14. Bonneau, J., Clark, J., Goldfeder, S.: On bitcoin as a public randomness source. Tech. rep., Cryptology ePrint Archive (2015), <https://eprint.iacr.org/2015/1015>
15. Bracha, G.: Asynchronous byzantine agreement protocols. *Information and Computation* **75**(2), 130–143 (1987)
16. Büinz, B., Goldfeder, S., Bonneau, J.: Proofs-of-delay and randomness beacons in ethereum. *IEEE Security and Privacy on the blockchain (IEEE S&B)* (2017)
17. Cachin, C., Kursawe, K., Shoup, V.: Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography (extended abstract). In: *PODC*. pp. 123–132. ACM (2000)
18. Canetti, R., Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Adaptive security for threshold cryptosystems. In: *CRYPTO*. Lecture Notes in Computer Science, vol. 1666, pp. 98–115. Springer (1999)
19. Cascudo, I., David, B.: SCRAPE: scalable randomness attested by public entities. In: *ACNS*. Lecture Notes in Computer Science, vol. 10355, pp. 537–556. Springer (2017)
20. Chase, M., Lysyanskaya, A.: On signatures of knowledge. In: *CRYPTO*. Lecture Notes in Computer Science, vol. 4117, pp. 78–96. Springer (2006)
21. Chor, B., Goldwasser, S., Micali, S., Awerbuch, B.: Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In: *FOCS*. pp. 383–395. IEEE Computer Society (1985)
22. Cramer, R., Damgård, I., Nielsen, J.B.: Multiparty computation from threshold homomorphic encryption. In: *EUROCRYPT*. Lecture Notes in Computer Science, vol. 2045, pp. 280–299. Springer (2001)
23. Das, S., Krishnan, V., Isaac, I.M., Ren, L.: Spurt: Scalable distributed randomness beacon with transparent setup. In: *SP*. pp. 2502–2517. IEEE (2022)
24. Das, S., Xiang, Z., Tomescu, A., Spiegelman, A., Pinkas, B., Ren, L.: A new paradigm for verifiable secret sharing. Tech. rep., Cryptology ePrint Archive (2023), <https://eprint.iacr.org/2023/1196>
25. Das, S., Yurek, T., Xiang, Z., Miller, A., Kokoris-Kogias, L., Ren, L.: Practical asynchronous distributed key generation. In: *SP*. pp. 2518–2534. IEEE (2022)
26. David, B., Gazi, P., Kiayias, A., Russell, A.: Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In: *EUROCRYPT (2)*. Lecture Notes in Computer Science, vol. 10821, pp. 66–98. Springer (2018)

27. David, B., Magri, B., Matt, C., Nielsen, J.B., Tschudi, D.: Gearbox: Optimal-size shard committees by leveraging the safety-liveness dichotomy. In: CCS. pp. 683–696. ACM (2022)
28. Dolev, D., Reischuk, R.: Bounds on information exchange for byzantine agreement. In: PODC. pp. 132–140. ACM (1982)
29. Dolev, D., Strong, H.R.: Authenticated algorithms for byzantine agreement. SIAM J. Comput. **12**(4), 656–666 (1983)
30. Feldman, P.: A practical scheme for non-interactive verifiable secret sharing. In: FOCS. pp. 427–437. IEEE Computer Society (1987)
31. Feng, H., Mai, T., Tang, Q.: Scalable and adaptively secure any-trust distributed key generation and all-hands checkpointing. Tech. rep., Cryptology ePrint Archive (2023), <https://eprint.iacr.org/2023/1773>
32. Fischer, M.J., Lynch, N.A., Paterson, M.: Impossibility of distributed consensus with one faulty process. J. ACM **32**(2), 374–382 (1985)
33. Fouque, P., Stern, J.: One round threshold discrete-log key generation without private channels. In: Public Key Cryptography. Lecture Notes in Computer Science, vol. 1992, pp. 300–316. Springer (2001)
34. Gao, Y., Lu, Y., Lu, Z., Tang, Q., Xu, J., Zhang, Z.: Efficient asynchronous byzantine agreement without private setups. In: ICDCS. pp. 246–257. IEEE (2022)
35. Gennaro, R., Goldfeder, S., Narayanan, A.: Threshold-optimal DSA/ECDSA signatures and an application to bitcoin wallet security. In: ACNS. Lecture Notes in Computer Science, vol. 9696, pp. 156–174. Springer (2016)
36. Gennaro, R., Jarecki, S., Krawczyk, H., Rabin, T.: Secure distributed key generation for discrete-log based cryptosystems. J. Cryptol. **20**(1), 51–83 (2007)
37. Gentry, C., Halevi, S., Lyubashevsky, V.: Practical non-interactive publicly verifiable secret sharing with thousands of parties. In: EUROCRYPT (1). Lecture Notes in Computer Science, vol. 13275, pp. 458–487. Springer (2022)
38. Gilad, Y., Hemo, R., Micali, S., Vlachos, G., Zeldovich, N.: Algorand: Scaling byzantine agreements for cryptocurrencies. In: SOSP. pp. 51–68. ACM (2017)
39. Groth, J.: On the size of pairing-based non-interactive arguments. In: EUROCRYPT (2). Lecture Notes in Computer Science, vol. 9666, pp. 305–326. Springer (2016)
40. Groth, J., Shoup, V.: Fast batched asynchronous distributed key generation. Tech. rep., Cryptology ePrint Archive (2023), <https://eprint.iacr.org/2023/1175>
41. Guo, B., Lu, Z., Tang, Q., Xu, J., Zhang, Z.: Dumbo: Faster asynchronous BFT protocols. In: CCS. pp. 803–818. ACM (2020)
42. Gurkan, K., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G., Tomescu, A.: Aggregatable distributed key generation. In: EUROCRYPT (1). Lecture Notes in Computer Science, vol. 12696, pp. 147–176. Springer (2021)
43. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 6477, pp. 177–194. Springer (2010)
44. King, V., Saia, J.: Breaking the $O(n^2)$ bit barrier: scalable byzantine agreement with an adaptive adversary. In: PODC. pp. 420–429. ACM (2010)
45. Lindell, Y., Lysyanskaya, A., Rabin, T.: On the composition of authenticated byzantine agreement. In: STOC. pp. 514–523. ACM (2002)
46. Lu, Y., Lu, Z., Tang, Q., Wang, G.: Dumbo-mvba: Optimal multi-valued validated asynchronous byzantine agreement, revisited. In: PODC. pp. 129–138. ACM (2020)
47. Momose, A., Ren, L.: Optimal communication complexity of authenticated byzantine agreement. In: DISC. LIPIcs, vol. 209, pp. 32:1–32:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021)
48. Nayak, K., Ren, L., Shi, E., Vaidya, N.H., Xiang, Z.: Improved extension protocols for byzantine broadcast and agreement. In: DISC. LIPIcs, vol. 179, pp. 28:1–28:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
49. Neji, W., Sinaoui, K.B., Rajeb, N.B.: Distributed key generation protocol with a new complaint management strategy. Secur. Commun. Networks **9**(17), 4585–4595 (2016)
50. Nguyen, L.: Accumulators from bilinear pairings and applications. In: CT-RSA. Lecture Notes in Computer Science, vol. 3376, pp. 275–292. Springer (2005)
51. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 1592, pp. 223–238. Springer (1999)
52. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: CRYPTO. Lecture Notes in Computer Science, vol. 576, pp. 129–140. Springer (1991)
53. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: CRYPTO. Lecture Notes in Computer Science, vol. 5157, pp. 554–571. Springer (2008)

54. Shoup, V.: The many faces of schnorr. Tech. rep., Cryptology ePrint Archive (2023), <https://eprint.iacr.org/2023/1019>
55. Shrestha, N., Bhat, A., Kate, A., Nayak, K.: Synchronous distributed key generation without broadcasts. Tech. rep., Cryptology ePrint Archive (2021), <https://eprint.iacr.org/2021/1635>
56. Thyagarajan, S.A.K., Castagnos, G., Laguillaumie, F., Malavolta, G.: Efficient CCA timed commitments in class groups. In: CCS. pp. 2663–2684. ACM (2021)
57. Tomescu, A., Chen, R., Zheng, Y., Abraham, I., Pinkas, B., Golan-Gueta, G., Devadas, S.: Towards scalable threshold cryptosystems. In: IEEE Symposium on Security and Privacy. pp. 877–893. IEEE (2020)
58. Wang, G., Nixon, M.: Randchain: Practical scalable decentralized randomness attested by blockchain. In: Blockchain. pp. 442–449. IEEE (2020)
59. Yurek, T., Luo, L., Fairoze, J., Kate, A., Miller, A.: hbackss: How to robustly share many secrets. In: NDSS. The Internet Society (2022)
60. Zamani, M., Movahedi, M., Raykova, M.: Rapidchain: Scaling blockchain via full sharding. In: CCS. pp. 931–948. ACM (2018)
61. Zhang, J., Xie, T., Hoang, T., Shi, E., Zhang, Y.: Polynomial commitment with a one-to-many prover and applications. In: USENIX Security Symposium. pp. 2965–2982. USENIX Association (2022)

A A PVSS Scheme with Field-element Secrets

Public key encryption. A PKE scheme Σ_{enc} consists of KeyGen , Enc , Dec , satisfying the standard IND-CPA security.

Signature of Knowledge. A signature of knowledge (SoK) scheme SoK for an NP language L consists of three algorithms. Setup generates a CRS which is an implicit input of other algorithms. $\text{Sign}(x, w, m)$ on inputs a statement x , a witness w , and a message m , produces a signature σ on m . $\text{Vrfy}(x, m, \sigma)$ verifies the signature. A SoK scheme should satisfy the Sim-Ext security [20].

NIZK. We need a NIZK Π_{zk} for showing the well-formedness of the ciphertexts. It consists of Setup , Prove , and Vrfy , for a language L_{pvss} which becomes apparent in the description of construction. We require Π_{zk} to satisfy completeness, simulation soundness, and zero knowledge against PPT adversaries.

The construction.

- $\text{Init}(1^\lambda, n)$. (1) A generator g for the group \mathbb{G} of order p ; (2) For every $i \in [n]$, $(ek_i, dk_i) \leftarrow \Sigma_{\text{enc}}.\text{KeyGen}(1^\lambda)$; (3) A setup for SoK ; (4) A setup for NIZK Π_{nizk} .
- $\text{Deal}((ek_i)_{i \in [n]}, \text{cid})$. Sample $(a_0, a_1, \dots, a_t) \leftarrow_{\$} \mathbb{Z}_p^{t+1}$, define $f(X) = \sum_{i=0}^t a_i X^i$, and compute $(A_i = g^{f(i)})_{i \in [0, n]}$, and $(c_i = \Sigma_{\text{enc}}.\text{Enc}(ek_i, f(i)); r_i)_{i \in [n]}$ where r_i is a fresh randomness. Sign cid with the knowledge of $f(0)$ w.r.t. A_0 : $\text{SoK}.\text{Sign}(A_0, f(0), \text{cid}) \rightarrow \sigma$. Prove the well-formedness of ciphertexts and obtain a proof π_{zk} : Π_{zk} is for the following language

$$\begin{aligned} & \{\text{statement}: (A_i, ek_i, c_i)_{i \in [n]}; \text{witness}: (f(i), r_i)_{i \in [n]} : \\ & A_i = g^{f(i)} \wedge c_i = \Sigma_{\text{enc}}.\text{Enc}(ek_i, f(i); r_i) \wedge f(i) < p \} \end{aligned} \quad (7)$$

$\text{Trans} = ((A_i)_{i \in [0, n]}, (c_i)_{i \in [n]}, \pi_{\text{zk}}, \{\sigma\}, \{\text{cid}\})$.

- $\text{PubVrfy}((ek_i)_{i \in [n]}, \text{Trans})$. It first checks whether

$$\text{SoK}.\text{Vrfy}(A_0^{(j)}, \text{cid}^{(j)}, \sigma^{(j)}) = 1, \text{ for all } j \in [m]. \quad (8)$$

Next, it randomly samples a $(n - t)$ -degree polynomial $q(x) \in \mathbb{Z}_p[x]$, and computes the dual code

$$(\text{code}_0^\perp, \dots, \text{code}_n^\perp), \text{ where } \text{code}_i^\perp = \frac{q(i)}{\prod_{j=0, j \neq i}^n (i - j)}. \quad (9)$$

The dual code is used for checking the validity of Scrape's polynomial commitment. See [19]. In our case, it computes $A_0 = \prod_j A_0^{(j)}$ and checks whether

$$\prod_{\tau=0}^n A_\tau^{\text{code}_\tau^\perp} = 1. \quad (10)$$

If the above check passes, it confirms us that the exponents of $(A_i)_{i \in [0, n]}$ is from a t -degree polynomial $f(X)$, and $A_i = g_1^{f(i)}$. Then, it checks if π_{zk} is valid proof. It returns 1 if all checks pass; otherwise, it returns 0.

- $\text{getCID}(\text{Trans})$. It returns $\{\text{cid}^{(j)}\}_{j \in [m]}$.
- $\text{PubDrv}(\text{Trans})$. It returns $pk = \prod_j A_0^{(j)}$ and $(pk_i = A_i)_{i \in [n]}$.
- $\text{Dec}(ek_i, dk_i, \text{Trans})$. It returns $sk_i = \text{Paillier}.\text{Dec}(ek_i, dk_i, c_i) \bmod p$.
- $\text{Rec}(\{(i, sk_i)\}_{i \in \mathbb{I}})$. It first computes the Lagrange coefficients $\{\lambda_i\}_{i \in \mathbb{I}}$ based on \mathbb{I} , and then $sk = \sum \lambda_i sk_i \bmod p$.

A.1 Analysis

Security analysis. The correctness is easy to follow. We establish the other security properties via the following lemmas.

Lemma 10. *Assuming the soundness of Π_{zk} for Eq. 7, our PVSS satisfies the soundness.*

Proof. In the soundness definition of PVSS, we require that every transcript passing the public verification can be decrypted to a consistent set of secret shares. According to the public verification algorithm's description, a valid transcript will pass the checking step in Eq. 10. This step guarantees, with overwhelming probability, that A_0, \dots, A_n commit to $f(0), \dots, f(n)$ for a polynomial f with a degree of up to t , as analyzed in [19]. Subsequently, by the soundness of Π_{zk} , the ciphertexts c_1, \dots, c_n encrypt to $f(1), \dots, f(n)$, ensuring that the decrypted values constitute a consistent set of secret shares. Similarly, for a weakly aggregated transcript wTrans

□

Lemma 11. *Assuming the IND-CPA security of the underlying PKE Σ_{enc} , the zero-knowledgness of Π_{zk} , and the security of SoK, the PVSS satisfies the strengthened secrecy and the simulation soundness.*

Proof. For clarity, we outline most simulator algorithms in the following.

- $\text{SInit}_{\mathcal{A}}(1^\lambda, n)$. It invokes the simulated setup algorithm of Π_{zk} and SoK to generate $(\text{crs}_{\text{zk}}, \text{tk}_{\text{zk}})$ and $(\text{crs}_{\text{sok}}, \text{tk}_{\text{sok}})$, respectively. Then, the CRS $\text{crs} = (\text{crs}_{\text{zk}}, \text{crs}_{\text{sok}})$ is provided to the adversary \mathcal{A} . After receiving the set of corrupted parties \mathcal{C} and the encryption keys $\{ek_i\}_{i \in \mathcal{C}}$ from \mathcal{A} , the simulator generates the encryption/decryption key pairs $(ek_i, dk_i)_{i \in [n] \setminus \mathcal{C}}$ for all uncorrupted users. It publishes all encryption keys $(ek_i)_{i \in [n]}$, and sets the trapdoor as $\text{tk} = (\text{tk}_{\text{zk}}, \text{tk}_{\text{sok}})$.
- $\text{SDeal}((ek_i)_{i \in [n]}, pk, \text{tk}, \text{cid})$. First, it sets $A_0 = pk$. For $i \in \mathcal{C}$, it samples $y_i \leftarrow_{\$} \mathbb{Z}_p$, computes $A_i = g^{y_i}$, and encrypts y_i under ek_i : $c_i \leftarrow \Sigma_{\text{enc}}.\text{Enc}(ek_i, y_i)$. For $i \in [n] \setminus \mathcal{C}$, it computes A_i via Lagrange interpolation using A_0 and $(A_i)_{i \in \mathcal{C}}$, and $c_i \leftarrow \Sigma_{\text{enc}}.\text{Enc}(ek_i, 0)$. Next, it invokes the simulated signing algorithm of SoK to sign cid and obtains the simulated signature σ . Finally, it invokes the simulated prover algorithm of Π_{zk} to generate the proof π_{zk} , and returns the transcript $\text{Trans} = ((A_i)_{i \in [0, n]}, (c_i)_{i \in [n]}, \pi_{\text{zk}}, \{\sigma\}, \{\text{cid}\})$.
- $\text{SRec}(\text{tk}, \text{Trans})$. It parses A_0, σ from Trans . Then, it runs the extractor algorithm of SoK with the trapdoor key tk_{sok} to extract y^* from σ , such that $A_0 = g^{y^*}$.

We argue the strengthened secrecy through the following hybrid arguments.

Hybrid 0. It runs $\text{Init}_{\mathcal{A}}(1^\lambda, n) \rightarrow (\text{crs}, \mathcal{C}, (ek_i, dk_i)_{i \notin \mathcal{C}}, (ek_i)_{i \in \mathcal{C}}, st_{\mathcal{A}})$ and $\text{Deal}((ek_i)_{i \in [n]}, \text{cid}) \rightarrow (\text{Trans}, sk)$. The adversary \mathcal{A} is provided with $(\text{crs}, (ek_i)_{i \in [n]}, st_{\mathcal{A}}, \text{Trans})$ as inputs.

Hybrid 1. It is almost identical to Hybrid 0, except that during the initial phase, the setup for Π_{zk} is replaced with the simulated setup algorithm of Π_{zk} , and that the prover algorithm of Π_{zk} in the Deal algorithm is replaced with the simulated prover algorithm.

Hybrid 2. It is almost identical to Hybrid 1, except that during the initial phase, the setup for SoK is replaced with the simulated setup algorithm of SoK, and that the signature of knowledge σ in Trans is generated with the simulated signing algorithm.

Hybrid 3. It is almost identical to Hybrid 2, except that it generates Trans in the following way. (1) Sample $y_i \leftarrow_{\$} \mathbb{Z}_p$ for all $i \in \mathcal{C}$, and sample a t -degree polynomial f , such that $f(i) = y_i$ for $i \in \mathcal{C}$. (2)

Compute $(A_i = g^{f(i)})_{i \in [0, n]}$, and $(c_i = \Sigma_{\text{enc}}.\text{Enc}(ek_i, f(i)); r_i)_{i \in [n]}$ where r_i is a fresh randomness. (3) Generate the signature of knowledge σ using the simulated signing algorithm. (4) Prove the well-formedness of ciphertexts via the simulated prover algorithm and obtain a proof π_{zk} .

Hybrid 4. It is almost identical to Hybrid 4, except that during generating **Trans**, the ciphertexts c_i for all $i \in [n] \setminus \mathcal{C}$ is generated as $c_i \leftarrow \Sigma_{\text{enc}}.\text{Enc}(ek_i, 0)$.

Hybrid 5. It is almost identical to Hybrid 4, except that in generating **Trans**, it firstly samples $A_0 \leftarrow \mathbb{G}$ and $y_i \leftarrow \mathbb{Z}_p$ for $i \in \mathcal{C}$, interpolates a t -degree polynomial in the exponent based on A_0 and $(A_i = g^{y_i})_{i \in \mathcal{C}}$, and obtains A_i for $i \in [n] \setminus \mathcal{C}$. After that, the remaining procedures for generating transcript **Trans** are identical to that in Hybrid 4.

For each hybrid $k \in [0, 5]$, we denote the probability that the adversary \mathcal{A} outputs 1 by $\mathbb{P}_{\mathcal{A}}^{(k)}$. Notice that, the PVSS scheme satisfies the strengthened secrecy, if

$$|\mathbb{P}_{\mathcal{A}}^{(0)} - \mathbb{P}_{\mathcal{A}}^{(5)}| \leq \text{negl}(\lambda). \quad (11)$$

It is easy to see that $|\mathbb{P}_{\mathcal{A}}^{(0)} - \mathbb{P}_{\mathcal{A}}^{(1)}| \leq \text{negl}(\lambda)$, due to zero-knowledgeness of Π_{zk} . Similarly, we have $|\mathbb{P}_{\mathcal{A}}^{(1)} - \mathbb{P}_{\mathcal{A}}^{(2)}| \leq \text{negl}(\lambda)$, from the security of SoK. Also, we have $\mathbb{P}_{\mathcal{A}}^{(2)} = \mathbb{P}_{\mathcal{A}}^{(3)}$ (and $\mathbb{P}_{\mathcal{A}}^{(4)} = \mathbb{P}_{\mathcal{A}}^{(5)}$), as the adversary's views in Hybrid 2 and 3 (resp. Hybrid 4 and 5) are essentially identical. Moreover, it holds that $|\mathbb{P}_{\mathcal{A}}^{(3)} - \mathbb{P}_{\mathcal{A}}^{(4)}| \leq \text{negl}(\lambda)$, from the IND-CPA security of the encryption scheme.

Claim. Assuming the security of SoK, our PVSS satisfies the simulation soundness.

Proof. Recall the simulation soundness definition, where we require that the adversary can only issue a challenge transcript that does not contain the CID included in the simulated transcript. Therefore, the signature of knowledge σ in the challenge transcript must be different from the simulated signature of knowledge, and thus we can extract the witness y^* from the signature. \square

Instantiation and performance analysis. The underlying PKE can be instantiated by the Paillier encryption [51], the LWE-based PVW encryption [53], or any other encryption scheme which is accompanied with efficient NIZK Π_{zk} for the language specified in Eq. 7. See [37] for a comprehensive overview on the candidate constructions. Regarding the SoK, we can essentially apply Schnorr signature, while viewing A_0 as the verification key and a_0 as the signing key, respectively. As demonstrated by Gentry et al. [37], the PVW encryption and the associated NIZK enjoy better concrete performance.

Then, assuming the proof size, prover time, and verification time of Π_{zk} are linear to n (which is true of the instantiations discussed above), the transcript size of **Trans** is $O(n\lambda)$. Both **Deal** and **PubVrfy** necessitate $O(n)$ group operations. The computational overhead for the remaining operations is relatively insubstantial, approximating $O(1)$ group operations.

B An Aggregatable PVSS

We present an aggregatable PVSS scheme in this section, which is a simplified variant of the scheme in [42].

Building blocks.

Pairing. There is an efficient deterministic algorithm **GroupGen** which outputs the description of the pairing groups, including $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$, where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are groups of order p , g_1 is the generator of \mathbb{G}_1 , \hat{h}_1 is the generator of \mathbb{G}_2 , and $e : \mathbb{G}_1 \times \mathbb{G}_2$ is a bilinear map.

Signature of Knowledge. A signature of knowledge (SoK) scheme SoK for an NP language L consists of three algorithms. Setup generates a CRS which is an implicit input of other algorithms. $\text{Sign}(x, w, m)$ on inputs a statement x , a witness w , and a message m , produces a signature σ on m . $\text{Vrfy}(x, m, \sigma)$ verifies the signature. A SoK scheme should satisfy the Sim-Ext security [20].

The construction.

- $\text{Init}(1^\lambda, n)$. (1) $\text{GroupGen} \rightarrow (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$; (2) a group element $\hat{u}_1 \leftarrow \mathbb{G}_2$; (3) the setup for $\text{SoK.Setup} \rightarrow \text{crs}_{\text{soK}}$. Define $\text{crs} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1, \hat{u}_1, \text{crs}_{\text{soK}})$. (4) For uncorrupted $i \in [n]$, $dk_i \leftarrow \mathbb{Z}_p$, and $ek_i = \hat{h}_1^{dk_i}$;
- $\text{Deal}((ek_i)_{i \in [n]}, \text{cid})$. Sample $(a_0, a_1, \dots, a_t) \leftarrow \mathbb{Z}_p^{t+1}$, define $f(X) = \sum_{i=0}^t a_i X^i$, and compute $\hat{u}_2 = \hat{u}_1^{a_0}$, $(A_i = g^{f(i)})_{i \in [0, n]}$, and $(\hat{Y}_i = ek_i^{f(i)})_{i \in [1, n]}$. Sign cid with the knowledge of $f(0)$ w.r.t. A_0 : $\text{SoK.Sign}(A_0, f(0), \text{cid}) \rightarrow \sigma$. $\text{Trans} = ((A_i)_{i \in [0, n]}, (\hat{Y}_i)_{i \in [n]}, \hat{u}_2, \{\sigma\}, \{\text{cid}\})$.
- $\text{Agg}(\{\text{Trans}_j\}_{j \in [m]}, (ek_i)_{i \in [n]})$. Parse $\text{Trans}_j = ((A_i^{(j)})_{i \in [0, n]}, (\hat{Y}_i^{(j)})_{i \in [n]}, \hat{u}_2^{(j)}, \sigma^{(j)}, \{\text{cid}^{(j)}\})$. Compute $(A_i = \prod_j A_i^{(j)})_{i \in [0, n]}$, $(\hat{Y}_i = \prod_j \hat{Y}_i^{(j)})_{i \in [n]}$, and $\hat{u}_2 = \prod_j \hat{u}_2^{(j)}$. Return $\text{Trans} = ((A_i)_{i \in [n]}, (\hat{Y}_i)_{i \in [n]}, \hat{u}_2, \{A_0^{(j)}\}_{j \in [m]}, \{\sigma^{(j)}\}_{j \in [m]}, \{\text{cid}^{(j)}\}_{j \in [m]})$.
- $\text{PubVrfy}((ek_i)_{i \in [n]}, \text{Trans})$. It first checks $\text{SoK.Vrfy}(A_0^{(j)}, \text{cid}^{(j)}, \sigma^{(j)})$ for $j \in [m]$. Next, it randomly samples a $(n - t)$ -degree polynomial $q(x) \in \mathbb{Z}_p[x]$, and computes the dual code

$$(\text{code}_0^\perp, \dots, \text{code}_n^\perp), \text{ where } \text{code}_i^\perp = \frac{q(i)}{\prod_{j=0, j \neq i}^n (i - j)}.$$

The dual code is used for checking the validity of Scrape's polynomial commitment. See [19]. In our case, it computes $A_0 = \prod_j A_0^{(j)}$ and checks whether

$$\prod_{\tau=0}^n A_\tau^{\text{code}_\tau^\perp} = 1. \quad (12)$$

If the above check passes, it confirms us that the exponents of $(A_i)_{i \in [0, n]}$ is from a t -degree polynomial $f(X)$, and $A_i = g_1^{f(i)}$. Then, it checks if $e(A_0, \hat{u}_1) = e(g_1, \hat{u}_2)$, and if $e(g_1, \hat{Y}_i) = e(A_i, ek_i)$ for $i \in [n]$. It returns 1 if all checks pass; otherwise, it returns 0.

- $\text{getCID}(\text{Trans})$. It returns $\{\text{cid}^{(j)}\}_{j \in [m]}$.
- $\text{PubDriv}(\text{Trans})$. It returns $pk = (\prod_j A_0^{(j)}, \hat{u}_2)$ and $(pk_i = A_i)_{i \in [n]}$.
- $\text{Dec}(ek_i, dk_i, \text{Trans})$. It returns $sk_i = \hat{Y}_i^{\frac{1}{dk_i}}$.
- $\text{Rec}(\{(i, sk_i)\}_{i \in \mathbb{I}})$. It first computes the Lagrange coefficients $\{\lambda_i\}_{i \in \mathbb{I}}$ based on \mathbb{I} , and then $sk = \prod sk_i^{\lambda_i}$.

Performance analysis. The transcript size of Trans is $O((n + m)\lambda)$, where m represents the number of transcripts aggregated into Trans . Both Deal and PubVrfy require $O(n)$ group operations, whereas Agg demands $O(n \log m)$ group operations. The computational costs for other functions are minor, approximately $O(1)$ group operations.

B.1 Security Analysis

The security follows the proofs in [42], which can be reduced to SXDH and BDH assumptions. For clarity, we describe the simulators SInit , SDeal , and SRec in the following, which are implicitly given the proof in [42, Theorem 2].

- $\text{SInit}_{\mathcal{A}}(1^\lambda, n)$. Invoke (1) $\text{GroupGen} \rightarrow (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1)$, sample (2) a group element $\hat{u}_1 \leftarrow \mathbb{G}_2$, and run (3) the *simulated* setup for SoK which produces crs_{soK} and tk_{soK} . Define $\text{crs} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, \hat{h}_1, \hat{u}_1, \text{crs}_{\text{soK}})$, provide it to the adversary \mathcal{A} , and wait \mathcal{A} to specify the set of corrupted nodes \mathcal{C} . (4) for every $i \in [n] \setminus \mathcal{C}$, sample $\mu_i \leftarrow \mathbb{Z}_p$, and define $ek_i = \hat{u}_1^{\mu_i}$. Define $\text{tk} = (\text{tk}_{\text{soK}}, \{\mu_i\}_{i \in [n] \setminus \mathcal{C}})$.
- $\text{SDeal}((ek_i)_{i \in [n]}, pk, \text{tk}, \text{cid})$. Parse $pk = (g_0, h_0)$. Then, (1) for $i \in \mathcal{C}$, sample $a_i \leftarrow \mathbb{Z}_p$, set $A_i = g^{a_i}$, and $\hat{Y}_i = ek_i^{a_i}$. We assume with out loss of generality that $|\mathcal{C}| = t$. (2) For $i \notin \mathcal{C}$, let $A_i = g_0^{\lambda_0(i)} \prod_{j \in \mathcal{C}} A_j^{\lambda_j(i)}$, and $\hat{Y}_i = (h_0^{\lambda_0(i)} \prod_{j \in \mathcal{C}} \hat{u}_1^{a_j \lambda_j(i)})^{\mu_i}$, where $\lambda_j(i) = \prod_{k \in \mathcal{C}, k \neq j} \frac{i-k}{j-k}$. (3) Run the simulation signer algorithm of SoK to generate a simulated signature σ for cid . Output $\text{Trans} = ((A_i)_{i \in [0, n]}, (\hat{Y}_i)_{i \in [n]}, \hat{u}_2 = h_0, \sigma, \{\text{cid}\})$.
- $\text{SRec}(\text{tk}, \text{Trans}) \rightarrow sk$. Parse Trans , and obtain σ . Run the extraction algorithm of SoK with σ and tk_{soK} , and obtain $s \in \mathbb{Z}_p$ such that $A_0 = g_1^s$. Output $sk = \hat{h}_1^s$.