

Asynchronous Distributed Private-Key Generators for Identity-Based Cryptography

Aniket Kate Ian Goldberg
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
{akate,iang}@cs.uwaterloo.ca

Abstract

Identity-based cryptography can greatly reduce the complexity of sending encrypted messages over the Internet. However, it necessarily requires a private-key generator (PKG), which can create private keys for clients, and so can passively eavesdrop on all encrypted communications. Although a distributed PKG has been suggested as a way to mitigate this problem for Boneh and Franklin’s identity-based encryption (IBE) scheme, the solution previously proposed does not work over the asynchronous Internet. Further, a distributed PKG has not been considered for any other IBE scheme.

In this paper, we design distributed PKG setup and private key extraction protocols in an asynchronous communication model for three important IBE schemes; namely, Boneh and Franklin’s IBE, Sakai and Kasahara’s IBE, and Boneh and Boyen’s BB_1 -IBE. To establish the efficiency and the reliability of our protocols, we implement and test the asynchronous distributed key generator used in these protocols on the PlanetLab platform. We also perform a comparative study of the three distributed PKGs and present recommendations for their use.

1 Introduction

In 1984, Shamir [53] introduced the notion of identity-based cryptography (IBC) as an approach to simplify public-key and certificate management in a public-key infrastructure (PKI) and presented an open problem to provide an identity-based encryption (IBE) scheme. After seventeen years, Boneh and Franklin [8] proposed the first practical and secure IBE scheme (BF-IBE) using bilinear maps. After this seminal work, in the last few years, significant progress has been made in IBC (for details, refer a recent book on IBC [34] and references therein).

In an IBC system, a client chooses an arbitrary string such as her e-mail address to be her public key. Consequently, with a standardized public-key string format, an IBC scheme completely eliminates the need for public-key certificates. As an example, in an IBE scheme, a sender can encrypt a message for a receiver knowing just the identity of the receiver and importantly, without obtaining and verifying the receiver’s public-key certificate. Naturally, in such a system, a client herself is not capable of generating a private key for her identity. There is a trusted party called a *private-key generator* (PKG) which performs the system setup, generates a secret called *master key* and provides private keys to clients using it. As the PKG computes a private key for a client, it can decrypt all (her) messages passively. This inherent *key escrow* property asks for complete trust in the PKG, which is difficult to find in many realistic scenarios.

Need for the Distributed PKG Importantly, the amount of trust placed in the holder of an IBC master key is far greater than that placed in the holder of the private key of a certifying authority (CA) in a PKI. In a PKI, in order to attack a client, the CA has to actively generate a fake certificate for the client containing a fake public-key/private-key pair. In this case, it is often possible for the client to detect and prove the malicious behaviour of the CA. The CA cannot perform any passive attack; specifically, it cannot decrypt a message encrypted for the client using a client-generated public key and it cannot sign some document for the client, if the verifier gets a correct certificate from the client. On the other hand, in IBC,

- knowing the master key, the PKG can decrypt or sign the messages for any client, without any active attack and consequent detection (key escrow),

- the PKG can make clients’ private keys public without any possible detection, and
- in a validity period-based key revocation system [8], bringing down the PKG is sufficient to bring the system to a complete halt (*single point of failure*).

Therefore, the PKG in IBC needs to be far more trusted than the CA in a PKI. This has been considered as a reason for the slow adoption of IBC schemes outside of closed organizational settings.

Boneh and Franklin [8] suggest the use of a distributed PKG to solve these problems. In an (n, t) -distributed PKG, the master key of the IBC system is distributed among n PKG nodes such that a set of nodes of size t or smaller cannot compute the master key, while a client extracts her private key by obtaining private-key shares from any $t + 1$ or more nodes; she can then use the system’s public key to verify the correctness of her thus-extracted key. Boneh and Franklin [8] propose *verifiable secret sharing* (VSS) of the master key among multiple PKGs using Shamir secret sharing with a *dealer* [52] to design a distributed PKG and also hint towards a completely distributed approach using the distributed (shared) key generation (DKG) schemes of Gennaro et al. [27]. However, none of the IBE schemes defined after [8] consider the need and design of a distributed PKG. On the system side, the DKG schemes [27] suggested in [8] to design a distributed PKG are not advisable for use over the Internet. These DKG schemes are defined for the *synchronous communication model*, having bounded message delivery delays and processor speeds, and do not provide *safety* (the protocol does not fail or produce incorrect results) and *liveness* (the protocol eventually terminates) over the asynchronous Internet, having no bounds on message transfer delays or processor speeds.

As a whole, although various proposed practical applications using IBE, such as key distribution in ad-hoc networks [37] or pairing-based onion routing [36], require a distributed PKG as a fundamental need, there is no distributed PKG available for use over the Internet yet. Defining efficient distributed PKGs for various IBE schemes which can correctly function over the Internet has been an open problem for some time. This practical need for distributed PKGs for IBC schemes that can function over the Internet forms the motivation of this work.

Contributions In this paper, we present distributed PKGs for all three important IBE frameworks: namely, full-domain-hash IBEs, exponent-inversion IBEs and commutative-blinding IBEs. We propose distributed PKG setups and distributed private-key extraction protocols for Boneh and Franklin’s IBE (BF-IBE) [8], Sakai and Kasahara’s IBE (SK-IBE) [49], and Boneh and Boyen’s (modified) BB_1 -IBE [10] schemes for use over the Internet. In the process, we also design practical protocols for distributed multiplication and inverse computation tasks, which have their own applications. Observing that a distributed (shared) key generator (DKG) is the single most important component of distributed PKG, we implement a recently devised asynchronous DKG protocol [35] and demonstrate its efficiency and reliability with extensive testing over the PlanetLab platform [47]. Using our implementation results, operation counts, key sizes, and possible pairing types, we compare the performance of the distributed PKGs we define.

In §2, we compare various techniques suggested to solve the key escrow and single point of failure problems in IBC. We also discuss previous work related to DKG protocols. In §3, we describe a realistic asynchronous system model over the Internet and justify the choices made, while we define and describe cryptographic tools in our model in §4. With this background, in §5, we define and analyze distributed PKG protocols for the BF-IBE, SK-IBE and BB_1 -IBE schemes. We then implement a practical DKG protocol, and test its performance over the PlanetLab platform in §6. We also compare the IBE schemes based on their distributed PKGs and touch upon proactive security and group modification protocols for the system.

2 Related Work

We divide the related work into two parts. Distributed (shared) key generation is the most important component for distributed private-key generation in identity-based cryptography. We first discuss the existing work towards distributed key generation. As designing distributed PKGs is our main goal in this work, we concentrate on protocols in computational (as opposed to unconditional / information-theoretic) settings. Although somewhat ignored, there have been some efforts to mitigate the single point of failure and the key escrow issues in IBC systems; in the latter part of this section, we compare these alternatives with distributed PKG.

Although we are defining protocols for IBE schemes, as we are concentrating on distributed cryptographic protocols and due to space constraints, we do not include a comprehensive account of IBE here. We refer readers to [10] for a detailed discussion on the various IBE schemes defined in the literature. As a take-away message from this survey, we work in the random oracle model for efficiency and consequently practicality reasons.

Distributed Key Generation The notion of secret sharing was introduced independently by Shamir [52] and Blakley [6] in 1979. Since then, it has remained an important topic in security research. Significantly, Chor et al. [19] introduced verifiability in secret sharing. Feldman [22] developed the first efficient and non-interactive VSS protocol and Pedersen [45] presented a modification to it. However, these VSS are defined assuming a synchronous communication model. For an *asynchronous communication model*, Cachin et al. (AVSS) [12], Zhou et al. (APSS) [56], and Schultz et al. (MPSS) [51] defined VSS schemes in the computational setting. Of these, the APSS protocol is impractical for any reasonable system size, as it has an exponential $\binom{n}{t}$ factor in the message complexity (number of messages transferred), while MPSS is developed for a more mobile setting where set of the system nodes has to change completely between two consecutive phases. AVSS by Cachin et al. with its seemingly optimal *communication complexity* (number of bits transferred) is certainly a suitable choice for a distributed PKG system.

Pedersen [46] introduced the concept of distributed key generation and developed a DKG, where each node runs a variation of Feldman’s VSS, and distributed shares are added at the end to generate a combined shared secret without a dealer. Gennaro et al. [27] presented a simplification using just the original Feldman VSS called the Joint Feldman DKG (JF-DKG). Further, they found that DKGs based on the Feldman VSS (or using Feldman commitments [22]) do not guarantee uniformly random secret keys and define a new DKG combining Feldman and Pedersen commitments [46] which increases the *latency* (number of communication rounds) by one. However, in [28], they observed that DKGs based on Feldman commitments produce hard instances of discrete logarithm problems (DLPs), which may be sufficient for the security of some threshold cryptographic schemes.

To the best of our knowledge, the first DKG scheme in an asynchronous setting was only defined recently by Kate and Goldberg [35]. This protocol modifies the AVSS protocol to a more realistic hybrid model and performs leader-based agreement with a leader-changing mechanism to decide which of the nodes’ VSS will be included in the DKG calculation; that is, whereas in synchronous DKG schemes such as Pedersen’s above, all of the successful VSSs can be added at the end of the protocol to determine the final master key shares, in the asynchronous setting, some global consensus must be reached in order to find a sufficiently large set of VSSs which all honest nodes have completed. We implement this DKG protocol and verify its efficiency and reliability. Consequently, this DKG system forms the basis of our distributed PKG protocols. The original asynchronous DKG protocol uses Feldman commitments and consequently does not guarantee uniform randomness of the key. However, we observe that, in the random oracle model, using non-interactive zero-knowledge proofs of knowledge based on the Fiat-Shamir methodology [23], if required, it is possible to achieve uniform randomness in their scheme. In such a scheme, Feldman commitments are initially replaced by Pedersen commitments; the Feldman commitments are introduced only at the end of the protocol to obtain the required private key. The zero-knowledge proofs are used to show that the Feldman and Pedersen commitments both commit to the same values.

All of the above schemes are proved secure only against a static adversary, which can only choose its t compromisable nodes before a protocol run. They are not considered secure against an adaptive adversary because their simulation-based security proofs do not go through when the adversary can corrupt nodes adaptively.[29, §4.4] Feldman claimed [22, §9.3] that his VSS protocol is also secure against adaptive adversaries even though his simulation-based security proof did not work out. Canetti et al. [14] presented a scheme provably secure against adaptive adversaries with at least two more communication rounds as compared to JF-DKG and with interactive zero-knowledge proofs. Due to the inefficiency of adaptive (provably) secure DKG protocols, we stick to protocols provably secure only against a static adversary, though they have remained unattacked by an adaptive adversary for the last 22 years.

Alternatives to a Distributed PKG Although none of the IBE schemes except BF-IBE considered distributed PKG setup and key extraction in order to solve the inherent key escrow and single point of failure issues, there have been a few other efforts in the literature to counter those.

Al-Riyami and Paterson [1] introduce *certificateless public key cryptography* (CL-PKC) to address the key escrow problem by combining IBC with public-key cryptography (PKC). Their elegant approach, however, does not address the single point of failure problem. Although it is possible to solve the problem by distributing their PKG using a VSS (which employs a trusted dealer to generate and distribute the key shares), which is inherently cheaper than a DKG-based PKG by a linear factor, it is impossible to stop a dealer’s active attacks without completely distributed master-key generation. Further, as private-key extractions are less frequent than encryptions, it is certainly advisable to use more efficient options during encryption rather than private-key extraction. Finally, with the requirement of online access to the receiver’s public key, CL-PKC becomes ineffective for systems without continuous network access, where IBC is considered to be an important tool.

Lee et al. [39] and Gangishetti et al. [26] propose variants of the distributed PKG involving a more trustworthy

key generation centre (KGC) and other key privacy authorities (KPAs). As observed by Chunxiang et al. [20] for [39], these approaches are, in general, vulnerable to passive attack by the KGC. In addition, the trust guarantees required by a KGC can be unattainable in practice.

Recently, Goyal [31] reduces the required trust in the PKG by restricting its ability to distribute a client’s private key. This does not solve the problem of single point of failure. Further, the PKG in his system still can decrypt the clients’ messages passively, which leaves a secure and practical implementation of a generic distributed PKG wanting.

3 System Model and Assumptions

In this section, we discuss the assumptions and the system model for our distributed PKG system, giving special attention to its practicality over the Internet. We follow the system model of [35], which closely depicts the Internet, and as their DKG forms the basis of our distributed PKGs.

3.1 Communication Model

In the theoretical sense, distributed protocols designed with a synchronous or a *partially synchronous* (bounded message delivery delays and processor speeds, but the bounds are unknown and eventual [21]) communication assumption tend to be more efficient in terms of latency and message complexity than their counterparts designed with an asynchronous communication assumption. However, protocols defined in the synchronous or partially synchronous communication model invariably use some time bounds in their definition. An adversary, knowing those bounds, may slow down the protocol by appropriately delaying its messages, which makes deciding the time bounds correctly a difficult problem to solve. On the other hand, protocols defined for the asynchronous communication model use only numbers and types of messages and guarantee to finish quickly with only honest nodes communicating promptly. Therefore, we assume an asynchronous communication model.

Weak Synchrony (only for liveness) Generating true randomness in a completely distributed (dealerless) asynchronous setting efficiently, without using a DKG, although not impossible [15], is a difficult task to perform; the known computational threshold coin-tossing algorithms [13] require a dealer or a synchronous communications assumption. As observed in [35], asynchronous DKG requires a protocol to solve the *agreement on a set* problem [4], which needs distributed randomness or a synchrony assumption [24]. In the absence of an efficient randomization procedure, [35] uses a *weak synchrony* assumption by Castro and Liskov [16] for liveness, but not safety. According to this assumption, a function $delay(t)$, defining the message transmission delay of a message sent at time t , does not grow faster than t indefinitely. Assuming that network faults are eventually repaired and DoS attacks eventually stop, this assumption is valid in practice. We further discuss this assumption in §6.1.

3.2 Hybrid Adversary Model

Instead of using a standard t -Byzantine adversary in a system with n nodes P_1, P_2, \dots, P_n , we use a *hybrid* adversary introduced in [2], having another f non-Byzantine crashes, and modified in [35] to include network link failures.

For the standard t -Byzantine adversary, t nodes compromised or crashed by the adversary remain compromised forever. This does not depict the adversary model over the Internet accurately. Along with arbitrary behaviour by t Byzantine nodes, some nodes can just crash silently without showing malicious behaviour or just get disconnected from the system due to network failure or partitioning. As the adversary does not capture these f nodes or their secret parameters, it is not computationally and communicationally optimal to consider these nodes as Byzantine. It also gives rise to a sub-optimal resilience of $n \geq 3(t + f) + 1$ instead of the $n \geq 3t + 2f + 1$ bound effected by treating crashes and link failures separately from the Byzantine adversary.

In this *hybrid adversary model*, crashes and link failures belong to the same set of f nodes, as from a perspective of any other node of the system a crashed node behaves exactly same as a node whose link with it is broken. We recover secrets at these f nodes immediately after their trusted rebooting, which gives us the assumption that all non-Byzantine nodes may crash and recover repeatedly with a maximum f crashed nodes at any instant. If two nodes cannot communicate, then we treat at least one of two nodes as being either Byzantine or one of the currently crashed nodes. That is, following the standard asynchronous communication model literature, we assume that the adversary controls the network, but faithfully delivers all the messages between two honest uncrashed nodes.

3.3 Cryptographic Background

Bilinear Pairings IBC extensively utilizes bilinear pairings over elliptic curves. For three cyclic groups \mathbb{G} , $\hat{\mathbb{G}}$, and \mathbb{G}_T (all of which we shall write multiplicatively) of the same prime order p , a *bilinear pairing* e is a map $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$ with following properties.

- **Bilinearity:** For $g \in \mathbb{G}$, $\hat{g} \in \hat{\mathbb{G}}$ and $a, b \in \mathbb{Z}_p$, $e(g^a, \hat{g}^b) = e(g, \hat{g})^{ab}$.
- **Non-degeneracy:** The map does not send all pairs in $\mathbb{G} \times \hat{\mathbb{G}}$ to unity in \mathbb{G}_T .

If there is an efficient algorithm to compute $e(g, \hat{g})$ for any $g \in \mathbb{G}$ and $\hat{g} \in \hat{\mathbb{G}}$, the pairing e is called *admissible*. We also expect that it is not feasible to invert a pairing and come back to \mathbb{G} or $\hat{\mathbb{G}}$. All pairings considered in this paper are admissible and infeasible to invert. We call such groups \mathbb{G} and $\hat{\mathbb{G}}$ *pairing-friendly* groups. We refer readers to [5, Chap. IX and X] for a detailed mathematical discussion of bilinear pairings.

Following [25], we consider three types of pairings: namely, type 1, 2, and 3. In *type 1* pairings, an isomorphism $\phi : \hat{\mathbb{G}} \rightarrow \mathbb{G}$ as well as its inverse ϕ^{-1} are efficiently computable. These are also called *symmetric pairings* as for such pairings $e(g, \hat{g}) = e(\phi(\hat{g}), \phi^{-1}(g))$ for any $g \in \mathbb{G}$ and $\hat{g} \in \hat{\mathbb{G}}$. In *type 2* pairings, only the isomorphism ϕ , but not ϕ^{-1} , is efficiently computable. Finally in *type 3* pairings, neither of ϕ nor ϕ^{-1} can be efficiently computed. The efficiency of the pairing computation improves from type 1 to type 2 to type 3 pairings. For a detailed discussion of the performance aspects of pairings we refer the reader to a survey by Galbraith et al. [25].

Assumptions As mentioned in §2, for efficiency reasons, we assume the random oracle framework. Further, our adversary is computationally bounded with a security parameter κ . We assume an instance of a pairing infrastructure of multiplicative groups \mathbb{G} , $\hat{\mathbb{G}}$ and \mathbb{G}_T , whose common order p is a κ -bit prime. For commitments and proofs of knowledge, we use the *discrete logarithm* (DLog) [42, Chap. 3] assumption. For the security of the IBE schemes, we use the *bilinear Diffie-Hellman* (BDH) [32] and *bilinear Diffie-Hellman inversion* (BDHI) [43, 7] assumptions. For definitions of asymmetric versions of the latter two assumptions, we refer readers to Boyen’s recent survey [10].

4 Cryptographic Tools

In this section, we describe important cryptographic tools required to design distributed PKGs in the hybrid model having an asynchronous network of $n \geq 3t + 2f + 1$ nodes with a t -limited Byzantine adversary and f -limited crashes and network failures. Note that these tools are also useful in other asynchronous computational multiparty settings.

4.1 Homomorphic Commitments over \mathbb{Z}_p

A verification mechanism for a consistent dealing is fundamental to VSS. It is achieved using distributed computing techniques in the unconditional setting. In the computational setting, *homomorphic commitments* provide an efficient alternative. Let $\mathcal{C}(\alpha, [r]) \in \mathbb{G}$ be a homomorphic commitment to $\alpha \in \mathbb{Z}_p$, where r is an optional randomness parameter and \mathbb{G} is a (multiplicative) group. For such a homomorphic commitment, given $C_1 = \mathcal{C}(\alpha_1, [r_1])$ and $C_2 = \mathcal{C}(\alpha_2, [r_2])$, we have $C_1 \cdot C_2 = \mathcal{C}(\alpha_1 + \alpha_2, [r])$.

VSS protocols utilize two forms of commitments. Let g and h be two random generators of \mathbb{G} . Feldman, for his VSS protocol [22], used a commitment scheme of the form $\mathcal{C}_{(g)}(\alpha) = g^\alpha$ with computational security under the DLog assumption and unconditional share integrity. Pedersen [46] presented another commitment of the form $\mathcal{C}_{(g,h)}(\alpha, r) = g^\alpha h^r$ with unconditional security but computational integrity under the DLog assumption. In PKC based on computational assumptions, with adversarial access to the public key, unconditional security of the secret (private key or master key) is impossible. Further, in VSS schemes based on Pedersen commitments, in order to randomly select the generator h , an additional round of communication is required during bootstrapping. Consequently, in our scheme, we use simple and efficient Feldman commitments, except during a special case described in the DKG discussion below.

In their VSSs, Feldman and Pedersen use commitments of *coefficients* of shared polynomials. However, following the computational multiparty computation protocol by Gennaro et al. [30] and AVSS by Cachin et al. [12], we instead use commitments of *evaluations* of shared polynomials. This reduces the communication complexity (the total bit length of messages exchanged) of AVSS by a linear factor and makes verifications of shares’ products easier in the distributed multiplication protocol of [30]. To that end, we define the *Feldman commitment vector*

$\mathcal{C}_{(g)}^{(s)} = [g^s, g^{\varphi(1)}, \dots, g^{\varphi(n)}]$ where φ is a randomly selected polynomial of degree t over \mathbb{Z}_p with $\varphi(0) = s$. Similarly, the *Pedersen commitment vector* $\mathcal{C}_{(g,h)}^{(s,s')} = [g^s h^{s'}, g^{\varphi(1)} h^{\psi(1)}, \dots, g^{\varphi(n)} h^{\psi(n)}]$ where φ is as above, and ψ is similar, but with $\psi(0) = s'$. The j^{th} element of a Feldman commitment vector (counting from 0) will be denoted by $(\mathcal{C}_{(g)}^{(s)})_j$ (and similarly for Pedersen commitment vectors).

4.2 Non-interactive Proofs of Knowledge

As we assume the random oracle model in this paper, we can use non-interactive zero-knowledge proofs of knowledge (NIZKPK) based on the Fiat-Shamir methodology [23]. In particular, we use a variant of NIZKPK of a discrete logarithm and one for proof of equality of two discrete logarithms.

We employ a variant of NIZKPK of a discrete logarithm where given a Feldman commitment $\mathcal{C}_{(g)}(s)$ and a Pedersen commitment $\mathcal{C}_{(g,h)}(s, r)$ for $s, r \in \mathbb{Z}_p$, a prover proves that she knows s and r such that $\mathcal{C}_{(g)}(s) = g^s$ and $\mathcal{C}_{(g,h)}(s, r) = g^s h^r$. That is, the prover proves that the Feldman commitment and the Pedersen commitment are to the same value s . We denote this proof as

$$\text{NIZKPK}_{\equiv \text{Com}}(s, r, \mathcal{C}_{(g)}(s), \mathcal{C}_{(g,h)}(s, r)) = \pi_{\equiv \text{Com}} \in \mathbb{Z}_p^3. \quad (1)$$

We describe it in detail in Appendix A; it is nearly equivalent to proving knowledge of two discrete logarithms separately.

We also use another NIZKPK (proof of equality) of discrete logarithms [17] such that given two Feldman commitments $\mathcal{C}_{(g)}(s) = g^s$ and $\mathcal{C}_{(h)}(s) = h^s$, a prover proves equality of the associated discrete logarithms. We denote this proof as

$$\text{NIZKPK}_{\equiv \text{DLog}}(s, \mathcal{C}_{(g)}(s), \mathcal{C}_{(h)}(s)) = \pi_{\equiv \text{DLog}} \in \mathbb{Z}_p^2. \quad (2)$$

and refer readers to Appendix A for details. There exists an easier way to prove this equality of discrete logarithms if a pairing between the groups generated by g and h is available. Using a technique due to Joux and Nguyen [33] to solve the DDH problem over pairing-friendly groups, given g^x and $h^{x'}$ the verifier checks if $e(g, h^{x'}) \stackrel{?}{=} e(g^x, h)$. However, when using a type 3 pairing, in the absence of an efficient isomorphism between \mathbb{G} and $\hat{\mathbb{G}}$, if both g and h belong to the same group (say \mathbb{G} without loss of generality), then the pairing-based verification scheme does not work. In such a situation, the above NIZKPK provides a less efficient but completely practical alternative.

4.3 DKG over \mathbb{Z}_p

In an (n, t) -DKG protocol over \mathbb{Z}_p , a set of n nodes generates an element $s \in \mathbb{Z}_p$ in a distributed fashion with its shares $s_i \in \mathbb{Z}_p$ spread over the n nodes such that any subset of size greater than a threshold t can reveal or use the shared secret, while smaller subsets cannot. A DKG protocol consists of a *sharing* (DKG-Sh) phase and a *reconstruction* (DKG-Rec) phase. In the DKG-Sh phase, a distributed secret $s \in \mathbb{Z}_p$ is generated among n nodes such that each node P_i holds a share s_i and a commitment vector $\mathcal{C}^{(s)}$ of s and all of its shares. During the DKG-Rec phase, each node P_i reveals its share s_i and reconstructs s using verified revealed shares.

For our hybrid model having an asynchronous network of $n \geq 3t + 2f + 1$ nodes with a t -limited Byzantine adversary and f -limited crashes and network failures, We use a DKG protocol defined in [35] satisfying the following conditions:

Liveness: Once protocol DKG-Sh starts, all honest finally up nodes complete the protocol, except with negligible probability.

Agreement: If some honest node completes protocol DKG-Sh then, except with negligible probability, all honest finally up nodes will eventually complete protocol DKG-Sh.

Consistency: Once an honest node completes protocol DKG-Sh then there exists a fixed value $s \in \mathbb{Z}_p$ such that, if an honest node P_i reconstructs $z_i \in \mathbb{Z}_p$ during DKG-Rec, then $z_i = s$.

Privacy: If no honest node has started protocol DKG-Rec then, except with negligible probability, an adversary cannot compute the shared secret s .

A closer look at the privacy property suggests that in the presence of an adversary, the shared secret in the above DKG may not be *uniformly* random; this is a direct effect of using only Feldman commitments.[29, §3] However, in many cases, we do not need a uniformly random secret key; the security of these schemes relies on the assumption that the adversary cannot compute the secret. Most of the schemes in this paper similarly only require the assumption that it is infeasible to compute the secret given public parameters and we stick with Feldman commitments those cases. However, we do indeed need a uniformly random shared secret in the protocol in §4.6, which computes shares of the inverse of a shared secret. In that case, we use Pedersen commitments, but we do not employ the methodology defined by Gennaro et al. [29], which increases the latency in the system. We observe instead that with the random oracle assumption at our disposal, the communicationally demanding technique by Gennaro et al. can be replaced with the much simpler computational non-interactive zero-knowledge proof of equality of committed values $\text{NIZKPK}_{\equiv \text{Com}}$ described in Eq. 1.

We represent DKG protocols using Feldman commitments and Pedersen commitments as DKG_{Feld} and DKG_{Ped} respectively. For node P_i , the corresponding DKG-Sh and DKG-Rec schemes are defined as follows.

$$\left(\mathcal{C}_{\langle g, h \rangle}^{(s, s')}, [\mathcal{C}_{\langle g \rangle}^{(s)}, \text{NIZKPK}_{\equiv \text{Com}}], s_i, s'_i \right) = \text{DKG-Sh}_{\text{Ped}}(n, t, f, t', g, h, \alpha_i, \alpha'_i) \quad (3)$$

$$\left(\mathcal{C}_{\langle g \rangle}^{(s)}, s_i \right) = \text{DKG-Sh}_{\text{Feld}}(n, t, f, t', g, \alpha_i) \quad (4)$$

$$s = \text{DKG-Rec}_{\text{Ped}}(t, \mathcal{C}_{\langle g, h \rangle}^{(s, s')}, s_i, s'_i) \quad (5)$$

$$s = \text{DKG-Rec}_{\text{Feld}}(t, \mathcal{C}_{\langle g \rangle}^{(s)}, s_i) \quad (6)$$

Here, t' is the number of VSS instances to be chosen ($t < t' \leq 2t + 1$), $g, h \in \mathbb{G}$ are commitment generators, $\alpha_i, \alpha'_i \in \mathbb{Z}_p$ are respectively a secret and randomness shared by P_i , and $\mathcal{C}_{\langle g \rangle}^{(s)}$ and $\mathcal{C}_{\langle g, h \rangle}^{(s, s')}$ are respectively the Feldman and Pedersen commitment vectors described in §4.1. The optional $\text{NIZKPK}_{\equiv \text{Com}}$ is a vector of zero-knowledge proofs of knowledge that the corresponding entries of $\mathcal{C}_{\langle g \rangle}^{(s)}$ and $\mathcal{C}_{\langle g, h \rangle}^{(s, s')}$ commit to the same values. (The polynomial φ for the two types of commitments will be the same in this case.)

The worst-case message and communication complexities of protocol DKG-Sh [35] are $O(tdn^2(n + d))$ and $O(\kappa tdn^3(n + d))$ respectively, while those of protocol DKG-Rec are $O(n^2)$ and $O(\kappa n^2)$ respectively. Here, the function $d(\cdot)$ bounds the number of crashes that the adversary is allowed to perform.

Distributed Random Sharing over \mathbb{Z}_p This protocol generates shares of a secret z chosen jointly at random from \mathbb{Z}_p . Every node generates a random $r_i \in \mathbb{Z}_p$ and shares that using the DKG-Sh protocol with Feldman or Pedersen commitments as $\text{DKG-Sh}(n, t, f, t' = t + 1, g, [h], r_i, [r'_i])$ where the generator h and randomness r'_i are only required if Pedersen commitments are used. Liveness, agreement, consistency, privacy and message and communication complexities remain the same as those of the DKG-Sh protocol. We represent the corresponding protocols as follows:

$$\left(\mathcal{C}_{\langle g \rangle}^{(z)}, z_i \right) = \text{Random}_{\text{Feld}}(n, t, f, g) \quad (7)$$

$$\left(\mathcal{C}_{\langle g, h \rangle}^{(z, z')}, z_i, z'_i \right) = \text{Random}_{\text{Ped}}(n, t, f, g, h). \quad (8)$$

4.4 Distributed Addition over \mathbb{Z}_p

Let $\alpha, \beta \in \mathbb{Z}_p$ be two secrets shared among n nodes using the DKG-Sh protocol. Let polynomials $f(x), g(x) \in \mathbb{Z}_p[x]$ be the respectively associated degree- t polynomials and let $c \in \mathbb{Z}_p$ be a non-zero constant. Due to the linearity of Shamir's secret sharing [52], a node P_i with shares α_i and β_i can locally generate shares of $\alpha + \beta$ and $c\alpha$ by computing $\alpha_i + \beta_i$ and $c\alpha_i$, where $f(x) + g(x)$ and $cf(x)$ are the respective polynomials. $f(x) + g(x)$ is random if either one of $f(x)$ or $g(x)$ is, and $cf(x)$ is random if $f(x)$ is. Commitment entries for the resultant shares respectively are $\left(\mathcal{C}_{\langle g \rangle}^{(\alpha + \beta)} \right)_i = \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)} \right)_i + \left(\mathcal{C}_{\langle g \rangle}^{(\beta)} \right)_i$ and $\left(\mathcal{C}_{\langle g \rangle}^{(c\alpha)} \right)_i = \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)} \right)_i^c$.

4.5 Distributed Multiplication over \mathbb{Z}_p

Unlike addition, local distributed multiplication of two shared secret α and β looks unlikely. We use a distributed multiplication protocol against a computational adversary by Gennaro et al. [30, §4]. However, instead of their in-

teractive zero-knowledge proof, we utilize a pairing-based DDH problem solving technique [33] to verify the correctness of the product value shared by a node non-interactively. For shares α_i and β_i with Feldman commitments g^{α_i} and \hat{g}^{β_i} , given a commitment $g^{\alpha_i\beta_i}$ of the shared product, other nodes can verify its correctness by checking if $e(g^{\alpha_i}, \hat{g}^{\beta_i}) \stackrel{?}{=} e(g^{\alpha_i\beta_i}, \hat{g})$ provided the groups of g and \hat{g} are pairing-friendly. We observe that it is also possible to perform this verification when one of the involved commitments is a Pedersen commitment. However, if both commitments are Pedersen commitments, then we have to compute Feldman commitments for one of the values and employ $\text{NIZKPK}_{\equiv \text{Com}}$ to prove its correctness in addition to using the pairing-based verification. In such a case, the choice between the latter technique and the non-interactive version of zero-knowledge proof suggested by Gennaro et al. depends upon implementation efficiencies of the group operation and pairing computations.

In our IBC schemes, we always use the multiplication protocol with at least one Feldman commitment. We denote the multiplication protocol involving two Feldman commitments as Mul_{Feld} and the one involving a combination of the two types of commitments as Mul_{Ped} .

$$\left(\mathcal{C}_{\langle g^* \rangle}^{(\alpha\beta)}, (\alpha\beta)_i \right) = \text{Mul}_{\text{Feld}}(n, t, f, g^*, \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i \right), \left(\mathcal{C}_{\langle \hat{g} \rangle}^{(\beta)}, \beta_i \right)) \quad (9)$$

$$\left(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(\alpha\beta, \alpha\beta')}, (\alpha\beta)_i, (\alpha\beta')_i \right) = \text{Mul}_{\text{Ped}}(n, t, f, \hat{g}, \hat{h}, \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i \right), \left(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(\beta, \beta')}, \beta_i, \beta'_i \right)) \quad (10)$$

For Mul_{Feld} , $g^* = g$ or \hat{g} . For Mul_{Ped} , without loss of generality, we assume that β is distributed with the Pedersen commitment. If instead α uses Pedersen commitment, then the Pedersen commitment groups for $(\alpha\beta)$ change to g and h instead of \hat{g} and \hat{h} .

Briefly, the protocol works as follows. Every honest node runs appropriate $\text{DKG-Sh}(n, t, f, t' = 2t+1, \hat{g}, [\hat{h}], \alpha_i\beta_i, [\alpha_i\beta'_i])$ from Eq. 3 or 4. As discussed above, pairing-based DDH solving is used to verify that the shared value is equal to the product of α_i and β_i .¹ At the end of the DKG-Sh protocol, instead of adding the subshares of the selected VSS instances, we interpolate them at index 0 to get the new share $(\alpha\beta)_i$ of $\alpha\beta$.

When DKG-Sh runs for α and β are completed before the protocol starts, this protocol can be seen as an execution of DKG-Sh with Feldman or Pedersen commitments, and liveness, agreement, and consistency directly follow from the underlying DKG-Sh. For privacy, along with privacy of DKG-Sh we also need to prove that for an honest node, during DKG-Rec of $(\alpha\beta)_i$, the adversary cannot derive the original α_i or β_i . We obtain that by modifying Theorem 4 in [30] defined for a synchronous model, but leave that to the next version of this paper. The message and communication complexities are the same as those of the DKG protocol. Further, as the distributed addition can be performed locally, the above Mul protocols can be seamlessly extended for distributed computation of any expression having binary products. For ℓ shared secrets x_1, x_2, \dots, x_ℓ , and their corresponding Feldman commitments $\mathcal{C}_{\langle g \rangle}^{(x_1)}, \mathcal{C}_{\langle g \rangle}^{(x_2)}, \dots, \mathcal{C}_{\langle g \rangle}^{(x_\ell)}$, shares of any binary product $x' = \sum_{i=1}^m k_i x_{a_i} x_{b_i}$ with known constants k_i and indices a_i, b_i can be easily computed by extending the protocol in Eq. 9. We denote this generalization as follows.

$$\left(\mathcal{C}_{\langle g^* \rangle}^{(x')}, x'_i \right) = \text{Mul}_{\text{BP}}(n, t, f, g^*, \{(k_i, a_i, b_i)\}, \left(\mathcal{C}_{\langle g \rangle}^{(x_1)}, (x_1)_i \right), \left(\mathcal{C}_{\langle g \rangle}^{(x_2)}, (x_2)_i \right), \dots, \left(\mathcal{C}_{\langle g \rangle}^{(x_\ell)}, (x_\ell)_i \right)) \quad (11)$$

Node P_j shares $\sum_i k_i (x_{a_i})_j (x_{b_i})_j$. For a type 1 pairing, verification of the correctness of the sharing is done by other nodes as follows.

$$e(g^{\sum_i k_i (x_{a_i})_j (x_{b_i})_j}, g) \stackrel{?}{=} \prod_i e((g^{(x_{a_i})_j})^{k_i}, g^{(x_{b_i})_j})$$

For type 2 and 3 pairings, $\text{NIZKPK}_{\equiv \text{DLog}}$ is used to provide Feldman commitments to the $(x_{b_i})_j$ with generator \hat{g} , and then a pairing computation like the above is used.

We use this protocol in Eq. 11 during distributed private-key extraction in Boneh and Boyen's BB_1 -IBE scheme in §5.4.

4.6 Sharing the Inverse of a Shared Secret

Given an (n, t, f) -distributed secret α , computing shares of its inverse α^{-1} in distributed manner (without reconstructing α) can be done trivially but inefficiently using a distributed computation of α^{p-1} . It involves $O(\log p)$ distributed multiplications. However, using a technique by Bar-Ilan and Beaver [3] this can be done using just one Random, one Mul and one DKG-Rec protocol.

¹For type 3 pairings, a careful selection of commitment generators is required to make the pairing-based verification possible.

This protocol involves a **DKG-Rec** which outputs the product of the shared secret α with a distributed random element z . If z is created using Feldman commitments and is not uniformly random, the product αz may leak some information about α . We avoid this by using Pedersen commitments while generating z .

$$\left(\mathcal{C}_{\langle g^* \rangle}^{(\alpha^{-1})}, (\alpha^{-1})_i\right) = \text{Inverse}(n, t, f, \hat{g}, \hat{h}, \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i\right)) \quad (12)$$

Here g^* belongs to any group of order p . Assuming a distributed secret $\left(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i\right)$, the protocol works as follows: every node P_i :

1. runs $\left(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}, z_i, z'_i\right) = \text{Random}_{\text{Ped}}(n, t, f, \hat{g}, \hat{h})$;
2. computes shares of $(w, w') = (\alpha z, \alpha z')$ as $\left(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(w, w')}, w_i, w'_i\right) = \text{Mul}_{\text{Ped}}(n, t, f, \hat{g}, \hat{h}, \left(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i\right), \left(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}, z_i, z'_i\right))$;
3. then reconstructs $w = \text{DKG-Rec}_{\text{Ped}}(t, \mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(w, w')}, w_i, w'_i)$. If $w = 0$, repeats the above two steps, else locally computes $(\alpha^{-1})_i = w^{-1} z_i$;
4. finally, computes the commitment $\mathcal{C}_{\langle g^* \rangle}^{(\alpha^{-1})}$ using w^{-1} , $\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}$, and if required, any of the NIZKPK techniques.

As with the **Mul** protocol, the liveness, agreement and consistency of the **Inverse** protocol follow directly from the corresponding properties of **DKG-Sh**. To prove the privacy, we employ the privacy property of the **DKG-Sh** protocol, Lemma 6 in [3] and the uniform randomness of the w computed in public; the details are in the next version of this paper. Although in practice this protocol is at least two times as expensive as a **DKG-Sh** instance, in theoretical sense, it has same asymptotic message and communication complexities as those of the **DKG-Sh** protocol.

5 Distributed PKG for IBE

We present distributed PKG setup and private key extraction protocols for three IBE schemes: namely, Boneh and Franklin's IBE (BF-IBE) [8], Sakai and Kasahara's IBE (SK-IBE) [49], and Boneh and Boyen's IBE (BB₁-IBE) [10]. Each of these schemes represents a distinct important category of an IBE classification defined by Boyen [9]. They respectively belong to *full-domain-hash* IBE schemes, *exponent-inversion* IBE schemes, and *commutative-blinding* IBE schemes. Note that the distributed PKG architectures that we develop for each of the three schemes apply to every scheme in their respective categories. Our above choice of IBE schemes is influenced by a recent identity-based cryptography standard (IBCS) [11] and also a comparative study by Boyen [10], which finds the above three schemes to be the most practical IBE schemes in their respective categories. In his classification, Boyen [9] also includes another category for quadratic-residuosity-based IBE schemes; however, none of the known schemes in this category are practical enough to consider here.

The role of a PKG in an IBE scheme ends with a user's private-key extraction. The distributed form of the PKG does not affect the encryption and decryption steps of IBE. Consequently, we concentrate only the distributed PKG setup and private-key extraction steps of the three IBE schemes under consideration. Note that during private-key extractions, we insist on minimal interaction between clients and PKG nodes—transferring identity credentials from the client at the start and private-key shares from the nodes at the end. We start by describing a bootstrapping procedure required by all IBE schemes.

5.1 Bootstrapping Procedure

Each of the IBE schemes under consideration here requires the following three bootstrapping steps.

1. Determine the group size n , the security threshold t and the crashed-nodes threshold f such that $n \geq 3t + 2f + 1$ and form a group of n PKG nodes.
2. Choose the pairing type to be used and compute three groups \mathbb{G} , $\hat{\mathbb{G}}$, and \mathbb{G}_T of prime order p such that there exists a bilinear pairing e of the decided type with $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$. The group order p is determined by the security parameter κ . We will write all of the groups multiplicatively.

3. Choose two generators $g \in \mathbb{G}$ and $\hat{g} \in \hat{\mathbb{G}}$ required to generate public parameters as well as the commitments.

Any untrusted entity can perform these offline tasks. Honest DKG nodes can verify the correctness of the tuple (n, t, f) and confirm the group choices \mathbb{G} , $\hat{\mathbb{G}}$, and \mathbb{G}_T as the first step of their distributed PKG setup. If unsatisfied, they may decline to proceed.

5.2 Boneh and Franklin's BF-IBE

BF-IBE [8] belongs to the full-domain-hash IBE family. In a BF-IBE setup, a PKG generates a master key $s \in \mathbb{Z}_p$ and an associated public key $g^s \in \mathbb{G}$, and derives private keys ($d \in \hat{\mathbb{G}}$) for clients using their well-known identities (ID) and s . A client with identity ID receives the private key $d_{\text{ID}} = H(\text{ID})^s \in \hat{\mathbb{G}}$, where $H : \{0, 1\}^* \rightarrow \hat{\mathbb{G}}^*$ is a full-domain cryptographic hash function. ($\hat{\mathbb{G}}^*$ denotes the set of all elements in $\hat{\mathbb{G}}$ except the identity.) The security of BF-IBE is based on the BDH assumption.

Distributed PKG Setup The distributed PKG setup involves generation of the system master key and the associated system public-key tuple in the (n, t, f) -distributed form among n nodes. Each node P_i participates in a DKG over \mathbb{Z}_p to generate its share $s_i \in \mathbb{Z}_p$ of the distributed master key s . The system public-key tuple is of the form $\mathcal{C}_{(g)}^{(s)} = [g^s, g^{s_1}, \dots, g^{s_n}]$. We represent this using our $\text{Random}_{\text{Feld}}$ protocol from Eq. 7.

$$\left(\mathcal{C}_{(g)}^{(s)}, s_i\right) = \text{Random}_{\text{Feld}}(n, t, f, g)$$

Private-key Extraction After a successful setup, PKG nodes are ready to extract private keys for clients. As a client needs $t + 1$ correct shares, it is sufficient for the client to contact any $2t + f + 1$ nodes (say set \mathcal{Q}). The private-key extraction protocol works as follows.

1. A client with identity ID contacts every node from the set \mathcal{Q} .
2. Every honest node $P_i \in \mathcal{Q}$ verifies the client's identity and returns a private-key share $H(\text{ID})^{s_i} \in \hat{\mathbb{G}}$ over a secure and authenticated channel.
3. Upon receiving $t+1$ valid shares, the client can construct her private key d_{ID} as $d_{\text{ID}} = \prod_{P_i \in \mathcal{Q}} (H(\text{ID})^{s_i})^{\lambda_i} \in \hat{\mathbb{G}}$, where the Lagrange coefficient $\lambda_i = \prod_{P_j \in \mathcal{Q} \setminus \{i\}} \frac{j}{j-i}$.
4. The client can verify the correctness of the computed private key d_{ID} by checking $e(g, d_{\text{ID}}) \stackrel{?}{=} e(g^s, H(\text{ID}))$. If unsuccessful, she can verify the correctness of each received $H(\text{ID})^{s_i}$ by checking if $e(g, H(\text{ID})^{s_i}) \stackrel{?}{=} e(g^{s_i}, H(\text{ID}))$. An equality proves the correctness of the share, while an inequality indicates misbehaviour by the node P_i and its consequential removal from \mathcal{Q} .

In asymmetric pairings, elements of \mathbb{G} generally have a shorter representation than those of $\hat{\mathbb{G}}$. Therefore, we put the more frequently accessed system public-key shares in \mathbb{G} , while the occasionally transferred client private-key shares belong to $\hat{\mathbb{G}}$. This also leads to a reduction in the ciphertext size. However, for type 2 pairings, an efficient hash-to- $\hat{\mathbb{G}}$ is not available for the group $\hat{\mathbb{G}}$ [25]; in that case we compute the system public key shares in $\hat{\mathbb{G}}$ and use the more feasible group \mathbb{G} for the private key shares.

5.3 Sakai and Kasahara's SK-IBE

SK-IBE [49] belongs to the exponent-inversion IBE family. The PKG setup here remains exactly same as BF-IBE and the PKG generates a master key $s \in \mathbb{Z}_p$ and an associated public key $g^s \in \mathbb{G}$ just as in BF-IBE. However, the key-extraction differs significantly. Here, a client with identity ID receives the private key $d_{\text{ID}} = \hat{g}^{\frac{1}{s+H(\text{ID})}} \in \hat{\mathbb{G}}$, where $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$. Chen and Cheng [18] prove the security of SK-IBE based on the BDHI.

Distributed PKG Setup The distributed PKG setup remains the exactly same as that of BF-IBE, where $s_i \in \mathbb{Z}_p$ is the master-key share for node P_i and $\mathcal{C}_{(g)}^{(s)} = [g^s, g^{s_1}, \dots, g^{s_n}]$ is the system public-key tuple.

Private-key Extraction The private-key extraction for SK-IBE is not as straightforward as that for BF-IBE. We slightly modify the `Inverse` protocol described in §4.6; specifically, here a private-key extracting client receives w_i from the node in step 3 and instead of PKG nodes, the *client* performs the interpolation step of `DKG-Rec`. In step 4, instead of publishing, PKG nodes forward \hat{g}^{z_i} and the associated $\text{NIZKPK}_{\equiv \text{Com}}$ directly to the client, which computes \hat{g}^z and then $d_{\text{ID}} = (\hat{g}^z)^{w^{-1}}$. The reason behind this is to avoid possible key-escrow if the node computes both \hat{g}^z and w . Further, the nodes precompute another generator $\hat{h} \in \hat{\mathbb{G}}$ for Pedersen commitments using $(\mathcal{C}_{\langle \hat{g} \rangle}^{(r)}, r_i) = \text{Random}_{\text{Feld}}(n, t, f, \hat{g})$, and set $\hat{h} = (\mathcal{C}_{\langle \hat{g} \rangle}^{(r)})_0 = \hat{g}^r$.

1. A client with identity ID contacts all n nodes the system.
2. Node P_i verifies the client's identity, runs $(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}, z_i, z'_i) = \text{Random}_{\text{Ped}}(n, t, f, \hat{g}, \hat{h})$ and computes $s_i^{\text{ID}} = s_i + H(\text{ID})$ and for $0 \leq j \leq n$, $(\mathcal{C}_{\langle g \rangle}^{(s^{\text{ID}})})_j = (\mathcal{C}_{\langle g \rangle}^{(s)})_j \cdot g^{H(\text{ID})} = g^{s_j + H(\text{ID})}$.
3. P_i performs $(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(w, w')}, w_i, w'_i) = \text{Mul}_{\text{Ped}}(n, t, f, \hat{g}, \hat{h}, (\mathcal{C}_{\langle g \rangle}^{(s^{\text{ID}})}, s_i^{\text{ID}}), (\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}, z_i, z'_i))$, where $w = s^{\text{ID}}z = (s + H(\text{ID}))z$ and $w' = (s + H(\text{ID}))z'$ and sends $(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(w, w')}, w_i, w'_i)$ to the client, which upon receiving $t + 1$ verifiably correct shares (w_i, w'_i) reconstructs $w \in \mathbb{Z}_p$ using Lagrange-interpolation. If $w \neq 0$, then it computes w^{-1} or else starts again from step 1.
4. Node P_i computes $(\mathcal{C}_{\langle \hat{g} \rangle}^{(z)})_i = \hat{g}^{z_i}$ and sends that along with $\text{NIZKPK}_{\equiv \text{Com}}(z_i, z'_i, (\mathcal{C}_{\langle \hat{g} \rangle}^{(z)})_i, (\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')})_i)$ from Eq. 1.
5. The client Lagrange-interpolates $t + 1$ valid \hat{g}^{z_i} to compute \hat{g}^z and derives her private key $(\hat{g}^z)^{w^{-1}} = \hat{g}^{\frac{1}{(s+H(\text{ID}))}}$.

This protocol can be used without any modification with any type of pairing. Further, online execution of the $\text{Random}_{\text{Ped}}$ computation can be eliminated using batch precomputation of distributed random elements $(\mathcal{C}_{\langle \hat{g}, \hat{h} \rangle}^{(z, z')}, z_i, z'_i)$.

5.4 Boneh and Boyen's BB_1 -IBE

BB_1 -IBE belongs to the commutative-blinding IBE family. Boneh and Boyen [7] proposed the original scheme with a security reduction to the DBDH assumption in the standard model against selective-identity attacks. However, with a practical requirement of security against adaptive-identity chosen-ciphertext attacks (IND-ID-CCA), in the recent IBCS standard [11], Boyen and Martin proposed a modified version of BB_1 , which is IND-ID-CCA secure in the random oracle model under the BDH assumption. In [10], Boyen rightly claims that for practical applications, it would be preferable to rely on the random-oracle assumption rather than using a less efficient IBE scheme with a stronger security assumption or a weaker attack model. Here, we consider the modified BB_1 -IBE scheme as described in [10] and [11].

In the BB_1 -IBE setup, the PKG generates a master-key triplet $(\alpha, \beta, \gamma) \in \mathbb{Z}_p^3$ and an associated public key tuple $(g^\alpha, g^\gamma, e(g, \hat{g})^{\alpha\beta})$. A client with identity ID receives the private key tuple $d_{\text{ID}} = (\hat{g}^{\alpha\beta + (\alpha H(\text{ID}) + \gamma)r}, \hat{g}^r) \in \hat{\mathbb{G}}^2$, where $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$.

Distributed PKG Setup In [10], Boyen does not include the parameters \hat{g} and \hat{g}^β from the original BB_1 scheme [7] in his public key, as they are not required during key extraction, encryption or decryption (they are not omitted for security reasons). In the distributed setting, we in fact need those parameters to be public for efficiency reasons; a verifiable distributed computation of $e(g, \hat{g})^{\alpha\beta}$ becomes inefficient otherwise. To avoid key escrow of clients' private-key components (\hat{g}^r) , we also need \hat{h} and $\mathcal{C}_{\langle \hat{h} \rangle}^{(\beta)}$; otherwise, parts of clients' private-keys would appear in public commitment vectors. As in SK-IBE in §5.3, this extra generator $\hat{h} \in \hat{\mathbb{G}}$ is precomputed using the $\text{Random}_{\text{Feld}}$ protocol.

Distributed PKG setup of BB_1 involves distributed generation of the master-key tuple (α, β, γ) . Distributed PKG node P_i achieves this using the following three $\text{Random}_{\text{Feld}}$ protocol invocations.

$$\begin{aligned} (\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i) &= \text{Random}_{\text{Feld}}(n, t, f, g) \\ (\mathcal{C}_{\langle \hat{g} \rangle}^{(\beta)}, \beta_i) &= \text{Random}_{\text{Feld}}(n, t, f, \hat{g}) \\ (\mathcal{C}_{\langle g \rangle}^{(\gamma)}, \gamma_i) &= \text{Random}_{\text{Feld}}(n, t, f, g) \end{aligned}$$

Here, $(\alpha_i, \beta_i, \gamma_i)$ is the tuple of master-key shares for node P_i . We also need $\mathcal{C}_{\langle \hat{h} \rangle}^{(\beta)}$; each node P_i provides this by publishing $(\mathcal{C}_{\langle \hat{h} \rangle}^{(\beta)})_i = \hat{h}^{\beta_i}$ and the associated NIZKPK $_{\equiv DLog}(\beta_i, \hat{g}^{\beta_i}, \hat{h}^{\beta_i})$. The tuple $(\mathcal{C}_{\langle g \rangle}^{(\alpha)}, e(g, \hat{g})^{\alpha\beta}, \mathcal{C}_{\langle g \rangle}^{(\gamma)}, \mathcal{C}_{\langle \hat{h} \rangle}^{(\beta)})$ forms the system public key, where $e(g, \hat{g})^{\alpha\beta}$ can be computed from the public commitment entries. The vector $\mathcal{C}_{\langle \hat{g} \rangle}^{(\beta)}$, although available publicly, is not required for any further computation.

Private-key Extraction The most obvious way to compute a BB₁ private key seems to be for P_i to compute $\alpha_i\beta_i + (\alpha_i H(\text{ID}) + \gamma_i)r_i$ and provide the corresponding $\hat{g}^{\alpha_i\beta_i + (\alpha_i H(\text{ID}) + \gamma_i)r_i}$, \hat{g}^{r_i} to the client, who now needs $2t + 1$ valid shares to obtain her private key. However, $\alpha_i\beta_i + (\alpha_i H(\text{ID}) + \gamma_i)r_i$ here is not a share of a random degree- $2t$ polynomial. The possible availability of \hat{g}^{r_i} to the adversary creates a suspicion about privacy of the master-key share with this method.

For private-key extraction in BB₁-IBE with a distributed PKG, we instead use the Mul_{BP} protocol in which the client is provided with $\hat{g}^{(\alpha\beta + (\alpha H(\text{ID}) + \gamma)r)_i}$, where $(\alpha\beta + (\alpha H(\text{ID}) + \gamma)r)_i$ is a share of random degree t polynomial. The protocol works as follows.

1. A client with identity ID contacts all n nodes the system.
2. Node P_i verifies the client's identity and runs $(\mathcal{C}_{\langle \hat{h} \rangle}^{(r)}, r_i) = \text{Random}_{\text{Feld}}(n, t, f, \hat{h})$
3. P_i computes its share $w_i = (\alpha\beta + (\alpha H(\text{ID}) + \gamma)r)_i$ of $w = \alpha\beta + (\alpha H(\text{ID}) + \gamma)r$ using protocol Mul_{BP} in Eq. 11.

$$(\mathcal{C}_{\langle g^* \rangle}^{(w)}, w_i) = \text{Mul}_{\text{BP}}(n, t, f, g^*, desc, (\mathcal{C}_{\langle g \rangle}^{(\alpha)}, \alpha_i), (\mathcal{C}_{\langle \hat{h} \rangle}^{(\beta)}, \beta_i), (\mathcal{C}_{\langle g \rangle}^{(\gamma)}, \gamma_i), (\mathcal{C}_{\langle \hat{h} \rangle}^{(r)}, r_i))$$

where $desc = \{(1, 1, 2), (H(\text{ID}), 1, 4), (1, 3, 4)\}$ is the description of the required binary product under the ordering $(\alpha, \beta, \gamma, r)$ of secrets.

To justify our choices of commitment generators, we present the verification of the value shared in the Mul_{BP} protocol:

$$e(g^{\alpha_i\beta_i + (\alpha_i H(\text{ID}) + \gamma_i)r_i}, \hat{h}) \stackrel{?}{=} e(g^{\alpha_i}, \hat{h}^{\beta_i})e((g^{\alpha_i})^{H(\text{ID})}g^{\gamma_i}, \hat{h}^{r_i})$$

For type 2 and 3 pairings, $g^* = g$, as there is no efficient isomorphism from \mathbb{G} to $\hat{\mathbb{G}}$. However, for type 1 pairings, we use $g^* = \hat{h} = \phi(h)$. Otherwise, the resultant commitments for w (which is public) will contain the private-key part $g^{\alpha\beta + (\alpha H(\text{ID}) + \gamma)r}$.

4. Once the Mul_{BP} protocol has succeeded, Node P_i generates \hat{g}^{w_i} and \hat{g}^{r_i} and sends those to the client over a secure and authenticated channel.
5. The client Lagrange-interpolates the valid received shares to generate her private key $(\hat{g}^{\alpha\beta + (\alpha H(\text{ID}) + \gamma)r}, \hat{g}^r)$.
6. For type 1 and type 2 pairings, the client can use the pairing-based DDH solving to check the validity of the shares. However, for type 3 pairings, without an efficient mapping from $\hat{\mathbb{G}}$ to \mathbb{G} , pairing-based DDH solving can only be employed to verify \hat{g}^{w_i} . As a verification of \hat{g}^{r_i} , node P_i includes a NIZKPK $_{\equiv DLog}(r_i, \hat{h}^{r_i}, \hat{g}^{r_i})$ along with \hat{g}^{w_i} and \hat{g}^{r_i} .

As in SK-IBE in §5.3, online execution of the Random_{Feld} computation can be eliminated using batch precomputation of distributed random elements $(\mathcal{C}_{\langle \hat{h} \rangle}^{(r)}, r_i)$.

5.5 Security Analysis

Next we touch upon the security analysis of the three distributed PKG setup and key extraction protocols defined above. We only provide rough proof ideas here and leave detailed proofs for the next version of the paper.

The distributed PKG setup for all three IBE schemes only involve the instantiation of protocol Random_{Feld} (n, t, f, g^*) from Eq. 7, where $g^* = g$ or \hat{g} . Consequently, the distributed PKG setups obtain liveness and agreement and the master keys obtain consistency and privacy directly from protocol DKG-Sh.

For the private-key extraction protocols, liveness and consistency are easy to prove. A client always requests a sufficient number of nodes ($2t + f + 1$ in BF-IBE and n in SK-IBE and BB₁-IBE) so that malicious and crashed nodes do not hamper the completion of the protocols. BF-IBE achieves consistency using the consistency property

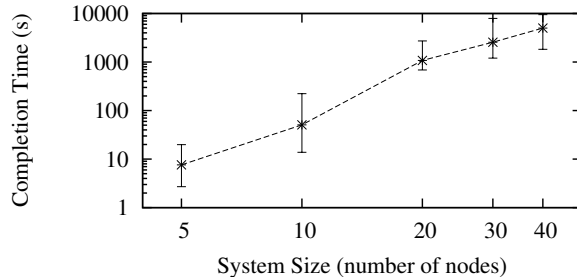


Figure 1: Completion Time (with min/max bars) vs System Size (log-log plot)

of the underlying DKG-Sh and pairing-based verification technique, while for SK-IBE and BB_1 -IBE, we also require consistency of protocol Mul.

To prove the privacy property of private-key extraction protocols against a polynomially bounded static adversary \mathcal{A} , which has access to t nodes in the system and can crash f others at any point during the protocol, we use the simulation game technique. A simulator \mathcal{B} wants to solve an instance of an appropriate hard problem (BDH problem for BF-IBE and BB_1 -IBE, and BDHI problem for SK-IBE) using the adversary algorithm \mathcal{A} . It runs all $n - t$ non-Byzantine nodes, sets one of the problem parameters as a part of one of its nodes' (say P_n 's) commitments, and completes the DKG-Sh, keeping P_n 's VSS instance in the set of selected VSSs. In the key extraction part, it simulates the random oracle and answers key extraction queries from the adversary using the available information. It finally sets the challenge for the adversary from the remaining parameters of the hard problem, and uses the answer from \mathcal{A} to answer the hard problem. With no interaction between the PKG nodes, proving the privacy of BF-IBE is straightforward using the above approach, while for SK-IBE and BB_1 -IBE, we also need to use the privacy property of protocol Mul.

6 System Aspects

In this section, we discuss the system aspects of distributed PKGs. As DKG is by far the most important component of our distributed PKGs, we first implement and test the DKG protocol [35] that we use in our distributed PKGs. In the process, we propose several system-level optimizations for this DKG. We also analyze practical aspects of our distributed PKGs and present a comparative study. Finally, we mention proactive security and group modification protocols for our distributed PKGs.

Note that two distributed CAs for PKC, Ω [48] and Cornell Online Certification Authority (COCA) [55], have been designed previously. However, with their focus on CAs, the protocols they provide are mismatched to the requirements of a distributed PKG. As a result, we do not design our distributed PKGs using these solutions.

6.1 DKG Implementation on PlanetLab

We design our DKG nodes as state machines (using the state machine replication approach [38, 50]), where nodes move from one state to another based on messages received. Messages are categorized into three types: operator messages, network messages and timer messages. The operator messages define interactions between nodes and their operators, the network messages realize protocol flows between nodes, and the timer messages implement the weak synchrony assumption described in §3.1.

We aim at building a distributed PKG for IBE schemes. Therefore, we develop our object-oriented C++ implementation over the PBC library [40] for the underlying elliptic-curve and finite-field operations and a PKI infrastructure with DSA signatures based on GnuTLS [41] for confidentiality and message authentication. (Note that *nodes* have TLS PKI certificates, which does not conflict with the goal of providing IBE private keys to *clients*.) In order to examine its realistic performance, we test our DKG implementation on the PlanetLab platform [47].

Table 1: Operation count and key sizes for distributed PKG setups and distributed private-key extractors (per private key)

	BF-IBE		SK-IBE		BB ₁ -IBE	
	Setup	Extraction	Setup	Extraction	Setup	Extraction
Operation Count						
Generator h or \hat{h}	X		√		√	
DKG-Sh ^a						
(precomputed)	-	0	-	1 ^P	-	1 ^F
(online)	1 ^F	0	1 ^F	1 ^P	3 ^F	1 ^F
Pairings						
@PKG Node	0	0	0	2 n	1 ^b	2 n
@Client	-	2(2 t + 2)	-	0	-	2 n ^b
NIZKPK	0	0	0	n	n ^b	n ^b
Interpolations	0	1	0	2	1	2
Key Sizes						
PKG Public Key	$(n + 2)\mathbb{G}^c$		$(n + 3)\mathbb{G}$		$(2n + 3)\mathbb{G}, (n + 2)\hat{\mathbb{G}}, (1)\mathbb{G}_T$	
Private-key Shares	$(2t + 1)\hat{\mathbb{G}}^c$		$(3n)\mathbb{Z}_p, (3n + 1)\hat{\mathbb{G}}$		$(2n)\mathbb{Z}_p^b, (2n)\hat{\mathbb{G}}$	

^aFor DKG-Sh F indicates use of Feldman commitments, while P indicates Pedersen commitments.

^bFor type 1 and 2 pairings, n NIZKPKs can be replaced by $2n$ extra pairings and the $2n \mathbb{Z}_p$ elements are omitted from the private-key shares.

^cFor type 2 pairings, the groups used for the PKG public key and the private-key shares are interchanged.

Performance Analysis We test the performance of our DKG implementation for systems of up to 40 nodes and we observe an expected approximately cubic growth in the average completion time.² Figure 1 presents our results in graphical form. In practical applications such as [36], these values, ranging from seconds to a little over an hour, are small as compared to DKG phase sizes (in days). Importantly, the use of dedicated high-performance servers instead of unreliable resource-shared PlanetLab nodes can drastically improve the performance. We also measure minimum and maximum completion times for the experiments. Big gaps between those values demonstrate the robustness of the DKG system against the Internet’s asynchronous nature and varied resource levels of the PlanetLab nodes.

To check the applicability of the weak synchrony assumption [16] that we use in DKG, we also tested the system with crashed leaders. In such scenarios, the DKG protocol successfully completed after a few leader changes. However, we observe that the average completion time of a system critically varies with the choice of $delay(t)$ functions and we suggest that this should only be finalized for a system after rigorous testing.

While implementing this system, we also found two system-level optimizations for this DKG.

- To the original DKG protocol, we add a new **shared** network message from a node to a leader having $2t + f + 1$ signed **ready** messages for a completed VSS. The leader can then include this VSS instance in its DKG **send** without completion of the VSS instance at its own machine.
- During our experiments, we observed that the VSS instances are more resource consuming than the agreement required at the end. Except during the Mul protocol, we only need $t + 1$ VSS instances to succeed. Assuming $t + f$ VSS instances might fail during a DKG, it is sufficient to start VSSs at just $2t + f + 1$ nodes instead of at all n nodes. Nodes that do not start a VSS initially may utilize the weak synchrony assumption to determine when to start a VSS instance if required.

6.2 Comparing Distributed PKGs

In this section, we concentrate on the performance of the setup and key extraction procedures of the three distributed PKGs defined in §5. For a detailed comparison of the encryption and decryption algorithms of BF-IBE, SK-IBE and BB₁-IBE, we refer readers to the survey by Boyen [10]. The general recommendations from this survey are to avoid SK-IBE and other exponent-inversion IBEs due to their reliance on the strong BDHI assumption, and that BB₁-IBE and BF-IBE both are good, but BB₁-IBE can be a better choice due to BF-IBE’s less efficient encryption.

Table 1 provides a detailed operation count and key size comparison of our three distributed PKGs. We count DKG-Sh instances, pairings, NIZKPKs, interpolations and public and private key sizes. We leave aside comparatively small exponentiations and other group operations. As mentioned in §5.4, for BB₁-IBE, with curves of type 1 and

²With cubic message complexity, larger distributed systems ($n > 50$) are not practical for the Internet.

2, there is a choice that can be made between using n NIZKPKs and $2n$ pairing computations. The table shows the NIZKPK choice (the only option for type 3 pairings), and footnote *b* shows where NIZKPKs can be traded off for pairings. As discussed in §5.2, for curves with type 2 pairings, an efficient algorithm for hash-to- $\hat{\mathbb{G}}$ is not available and we have to interchange the groups used for the system public key shares and client private-key shares. Footnote *c* indicates how that affects the key sizes.

In Table 1, we observe that the distributed PKG setup and the distributed private-key extraction protocols for BF-IBE are significantly more efficient than those for SK-IBE and BB_1 -IBE. Importantly, for BF-IBE, distributed PKG nodes can extract a key for a client without interacting with each other, which is not possible in the other two schemes; both BB_1 -IBE and SK-IBE require at least one DKG instance for every private-key extraction; the second required instance can be batch-precomputed. Therefore, for IBE applications in the random oracle model, we suggest the use of the BF-IBE scheme, except in situations where private-key extractions are rare and efficiency of the encryption step is critical to the system. For such applications, we suggest BB_1 -IBE as the small efficiency gains in the distributed PKG setup and extraction protocols of SK-IBE do not well compensate for the strong security assumption required. Further, BB_1 -IBE can also be proved secure in the standard model with selective-identity attacks. For applications demanding security in the standard model, our distributed PKG for BB_1 -IBE also provides a solution to the key escrow and single point of failure problems, using pairings of type 1 or 2.

6.3 Proactive Security & Group Modification

With an endless supply of software and network security flaws, system attacks not only are prevalent but have also been growing. The distributed nature of our protocols mitigates the effects of those attacks to some extent, but their time-independence makes them vulnerable to a *gradual break-in* by a *mobile attacker* breaking into system nodes one by one. The concept of proactive security [44] has been introduced to counter these attacks. Further, on a long-term basis, the set of PKG nodes will need to be modified, which can also cause changes to the system’s security threshold t and the crash-limit f . Therefore, for our distributed PKG systems, we need proactive security and group modification protocols.

We observe that the proactive security and group modification protocols defined in [35], for the DKG protocol used in our distributed PKGs, are directly applicable to our distributed PKGs. We suggest the use of these protocols to achieve proactive security of our master keys and group modification of our PKGs. Note that this is possible only due to the nature of the master keys for the three IBE schemes that we use. All master key elements in these three schemes belong to \mathbb{Z}_p , which is also the output domain for the DKG protocol. In contrast to the three IBEs that we consider, we leave as an open problem the possibility of providing proactive security and group modification protocols to the master keys for IBE schemes such the original BB_1 -IBE [7] or Waters’ IBE [54].

7 Conclusion

In this paper, we designed and compared distributed PKG setup and private key extraction protocols for Boneh and Franklin’s BF-IBE, Sakai and Kasahara’s SK-IBE, and Boneh and Boyen’s BB_1 -IBE. We observe that the distributed PKG implementation for BF-IBE is the most simple and efficient among all and we suggest its use when the system can support its relatively costly encryption step. For systems requiring a faster encryption, we suggest use BB_1 -IBE instead. However, during every distributed private key extraction, it requires a DKG and consequently, interaction among PKG nodes. That being said, during private-key extractions, we successfully avoid any interaction between clients and PKG nodes except the necessary identity at the start and key share transfers at the end. Further, each of the above three schemes represents a separate category of IBE schemes and our designs can be applied to other schemes in those category as well.

While developing our distributed PKGs, we also developed asynchronous computational protocols for distributed multiplication and distributed inverse computation, which may have their own applications. To confirm the feasibility of a distributed PKG in the asynchronous communication model, we also implemented and verified the efficiency and the reliability of an asynchronous DKG protocol using extensive testing over the PlanetLab platform. We also suggested proactive security and group modification protocols for our distributed PKGs and in the future, we would like add those features to our system.

References

- [1] S. S. Al-Riyami and K. G. Paterson. Certificateless Public Key Cryptography. In *ASIACRYPT'03*, pages 452–473, 2003.
- [2] M. Backes and C. Cachin. Reliable Broadcast in a Computational Hybrid Model with Byzantine Faults, Crashes, and Recoveries. In *DSN'03*, pages 37–46, 2003.
- [3] J. Bar-Ilan and D. Beaver. Non-Cryptographic Fault-Tolerant Computing in Constant Number of Rounds of Interaction. In *PODC'89*, pages 201–209, 1989.
- [4] M. Ben-Or, R. Canetti, and O. Goldreich. Asynchronous Secure Computation. In *STOC'93*, pages 52–61, 1993.
- [5] I. Blake, G. Seroussi, and N. P. Smart, editors. *Advances in Elliptic Curve Cryptography*. Number 317 in London Mathematical Society Lecture Note Series. Cambridge University Press, 2005. 183–252.
- [6] G. R. Blakley. Safeguarding Cryptographic Keys. In *the National Computer Conference*, pages 313–317, 1979.
- [7] D. Boneh and X. Boyen. Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In *EUROCRYPT'04*, pages 223–238, 2004.
- [8] D. Boneh and M. K. Franklin. Identity-Based Encryption from the Weil Pairing. In *CRYPTO'01*, pages 213–229, 2001.
- [9] X. Boyen. General *Ad Hoc* Encryption from Exponent Inversion IBE. In *EUROCRYPT'07*, pages 394–411, 2007.
- [10] X. Boyen. A Tapestry of Identity-based Encryption: Practical Frameworks Compared. *IJACT*, 1(1):3–21, 2008.
- [11] X. Boyen and L. Martin. Identity-Based Cryptography Standard (IBCS) (Version 1), Request for Comments (RFC) 5091. <http://www.ietf.org/rfc/rfc5091.txt>, 2007.
- [12] C. Cachin, K. Kursawe, A. Lysyanskaya, and R. Strohli. Asynchronous Verifiable Secret Sharing and Proactive Cryptosystems. In *ACM CCS'02*, pages 88–97, 2002.
- [13] C. Cachin, K. Kursawe, and V. Shoup. Random Oracles in Constantinople: Practical Asynchronous Byzantine Agreement Using Cryptography. In *PODC'00*, pages 123–132, 2000.
- [14] R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Adaptive Security for Threshold Cryptosystems. In *CRYPTO'99*, pages 98–115, 1999.
- [15] R. Canetti and T. Rabin. Fast Asynchronous Byzantine Agreement with Optimal Resilience. In *STOC'93*, pages 42–51, 1993.
- [16] M. Castro and B. Liskov. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, 2002.
- [17] D. Chaum and T. P. Pedersen. Wallet Databases with Observers. In *CRYPTO'92*, pages 89–105, 1992.
- [18] L. Chen and Z. Cheng. Security Proof of Sakai-Kasahara's Identity-Based Encryption Scheme. Cryptology ePrint Archive, Report 2005/226, 2005.
- [19] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch. Verifiable Secret Sharing and Achieving Simultaneity in the Presence of Faults (Extended Abstract). In *FOCS'85*, pages 383–395, 1985.
- [20] X. Chunxiang, Z. Junhui, and Q. Zhiguang. A Note on Secure Key Issuing in ID-based Cryptography. Technical report, 2005. <http://eprint.iacr.org/2005/180>.
- [21] C. Dwork, N. A. Lynch, and L. J. Stockmeyer. Consensus in the Presence of Partial Synchrony. *Journal of ACM*, 35(2):288–323, 1988.
- [22] P. Feldman. A Practical Scheme for Non-interactive Verifiable Secret Sharing. In *FOCS'87*, pages 427–437, 1987.
- [23] A. Fiat and A. Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO'86*, pages 186–194, 1986.
- [24] M. J. Fischer, N. A. Lynch, and M. Paterson. Impossibility of Distributed Consensus with One Faulty Process. *Journal of ACM*, 32(2):374–382, 1985.
- [25] S. D. Galbraith, K. G. Paterson, and N. P. Smart. Pairings for cryptographers. *Discrete Applied Mathematics*, 156(16):3113–3121, 2008.
- [26] R. Gangishetti, M. C. Gorantla, M. Das, and A. Saxena. Threshold key issuing in identity-based cryptosystems. *Computer Standards & Interfaces*, 29(2):260–264, 2007.
- [27] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. In *EUROCRYPT'99*, pages 295–310, 1999.
- [28] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure Applications of Pedersen's Distributed Key Generation Protocol. In *CT-RSA'03*, pages 373–390, 2003.

- [29] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure Distributed Key Generation for Discrete-Log Based Cryptosystems. *Journal of Cryptology*, 20(1):51–83, 2007.
- [30] R. Gennaro, M. O. Rabin, and T. Rabin. Simplified VSS and Fact-Track Multiparty Computations with Applications to Threshold Cryptography. In *PODC'98*, pages 101–111, 1998.
- [31] V. Goyal. Reducing Trust in the PKG in Identity Based Cryptosystems. In *CRYPTO'07*, pages 430–447, 2007.
- [32] A. Joux. A One Round Protocol for Tripartite Diffie-Hellman. In *ANTS*, pages 385–394, 2000.
- [33] A. Joux and K. Nguyen. Separating Decision Diffie-Hellman from Computational Diffie-Hellman in Cryptographic Groups. *Journal of Cryptology*, 16(4):239–247, 2003.
- [34] M. Joye and G. Neven. *Identity-Based Cryptography - Volume 2 Cryptology and Information Security Series*. IOS Press, Amsterdam, The Netherlands, The Netherlands, 2008.
- [35] A. Kate and I. Goldberg. Distributed Key Generation for the Internet. Technical Report 2008-25, CACR, University of Waterloo, 2008. To appear at ICDCS'09. Available as <http://www.cacr.math.uwaterloo.ca/techreports/2008/cacr2008-25.pdf>.
- [36] A. Kate, G. M. Zaverucha, and I. Goldberg. Pairing-Based Onion Routing. In *7th Privacy Enhancing Technologies Symposium (PET)*, pages 95–112, 2007.
- [37] A. Khalili, J. Katz, and W. Arbaugh. Toward Secure Key Distribution in Truly Ad-Hoc Networks. In *IEEE Workshop on Security and Assurance in Ad-Hoc Networks 2003*, pages 342–346, 2003.
- [38] L. Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM*, 21(7):558–565, 1978.
- [39] B. Lee, C. Boyd, E. Dawson, K. Kim, J. Yang, and S. Yoo. Secure key issuing in ID-based cryptography. In *ACSW Frontiers'04*, pages 69–74, 2004.
- [40] B. Lynn. PBC Library – The Pairing-Based Cryptography Library. <http://crypto.stanford.edu/pbc/>, 2009. Accessed April 2009.
- [41] N. Mavroyanopoulos, F. Fiorina, T. Schulz, A. McDonald, and S. Josefsson. The GNU Transport Layer Security Library. <http://www.gnu.org/software/gnutls/>, 2009. Accessed April 2009.
- [42] A. Menezes, P. V. Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1st edition, 1997.
- [43] S. Mitsunari, R. Sakai, and M. Kasahara. A New Traitor Tracing. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E85-A(2):481–484, 2002.
- [44] R. Ostrovsky and M. Yung. How to Withstand Mobile Virus Attacks (Extended Abstract). In *PODC'91*, pages 51–59, 1991.
- [45] T. P. Pedersen. A Threshold Cryptosystem without a Trusted Party. In *Eurocrypt'91*, pages 522–526. Springer-Verlag, 1991.
- [46] T. P. Pedersen. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *CRYPTO'91*, pages 129–140, 1991.
- [47] L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A Blueprint for Introducing Disruptive Technology into the Internet. *SIGCOMM Comput. Commun. Rev.*, 33(1):59–64, 2003.
- [48] M. K. Reiter, M. K. Franklin, J. B. Lacy, and R. N. Wright. The Omega Key Management Service. *Journal of Computer Security*, 4(4):267–288, 1996.
- [49] R. Sakai and M. Kasahara. ID based Cryptosystems with Pairing on Elliptic Curve. Cryptology ePrint Archive, Report 2003/054, 2003.
- [50] F. B. Schneider. Implementing Fault-Tolerant Services Using the State Machine Approach: A Tutorial. *ACM Comput. Surv.*, 22(4):299–319, 1990.
- [51] D. A. Schultz, B. Liskov, and M. Liskov. Mobile Proactive Secret Sharing. In *PODC'08*, page 458, 2008. (Extended Draft).
- [52] A. Shamir. How to Share a Secret. *Commun. ACM*, 22(11):612–613, 1979.
- [53] A. Shamir. Identity-Based Cryptosystems and Signature Schemes. In *CRYPTO'84*, pages 47–53, 1984.
- [54] B. Waters. Efficient Identity-Based Encryption Without Random Oracles. In *EUROCRYPT'05*, pages 114–127, 2005.
- [55] L. Zhou, F. B. Schneider, and R. van Renesse. COCA: A secure distributed online certification authority. *ACM Trans. Comput. Syst.*, 20(4):329–368, 2002.
- [56] L. Zhou, F. B. Schneider, and R. van Renesse. APSS: proactive secret sharing in asynchronous systems. *ACM Trans. Inf. Syst. Secur.(TISSec)*, 8(3):259–286, 2005.

A Non-interactive Zero-knowledge Proofs

We now present the details of the non-interactive zero-knowledge proofs of knowledge (NIZKPKs) introduced in §4.2. Here, H is a hash function modelled by a random oracle.

The first proof (which to our knowledge has not appeared before in the literature, but is straightforward) is that a Feldman commitment $F = \mathcal{C}_{\langle g \rangle}(s) = g^s$ and a Pedersen commitment $P = \mathcal{C}_{\langle g, h \rangle}(s, r) = g^s h^r$ are both committing to the same value s . We denote this by $\text{NIZKPK}_{\equiv \text{Com}}(s, r, F, P)$.

The proof is generated as follows:

- Pick $v_1, v_2 \in_R Z_p$
- Let $t_1 = g^{v_1}, t_2 = h^{v_2}$
- Let $c = H(g, h, F, P, t_1, t_2)$
- Let $w_1 = v_1 - c \cdot s \pmod{p}, w_2 = v_2 - c \cdot r \pmod{p}$
- The proof is $\pi_{\equiv \text{Com}} = (c, w_1, w_2)$

The verifier checks this proof (given $\pi_{\equiv \text{Com}}, g, h, F, P$) as follows:

- Let $t'_1 = g^{w_1} F^c, t'_2 = h^{w_2} (P/F)^c$
- Accept the proof as valid if $c = H(g, h, F, P, t'_1, t'_2)$

The second proof is that two Feldman commitments $F_1 = \mathcal{C}_{\langle g \rangle}(s) = g^s$ and $F_2 = \mathcal{C}_{\langle h \rangle}(s) = h^s$ commit to the same value; that is, the discrete logs of F_1 and F_2 to the bases of g and h respectively are equal. We denote this by $\text{NIZKPK}_{\equiv \text{DLog}}(s, F_1, F_2)$. The proof is standard [17]:

The proof is generated as follows:

- Pick $v \in_R Z_p$
- Let $t_1 = g^v, t_2 = h^v$
- Let $c = H(g, h, F_1, F_2, t_1, t_2)$
- Let $w = v - c \cdot s \pmod{p}$
- The proof is $\pi_{\equiv \text{DLog}} = (c, w)$

The verifier checks this proof (given $\pi_{\equiv \text{DLog}}, g, h, F_1, F_2$) as follows:

- Let $t'_1 = g^w F_1^c, t'_2 = h^w F_2^c$
- Accept the proof as valid if $c = H(g, h, F_1, F_2, t'_1, t'_2)$