Accountable Fine-grained Blockchain Rewriting in the Permissionless Setting

Yangguang Tian
Osaka University
Japan

Bowen Liu
Singapore University of
Technology and Design
Singapore
bowen_liu@mymail.sutd.edu.sg

Yingjiu Li
University of Oregon
USA
yingjiul@uoregon.edu

Pawel Szalachowski
Singapore University of
Technology and Design
Singapore
pjszal@gmail.com

Jianying Zhou
Singapore University of
Technology and Design
Singapore
jianying_zhou@sutd.edu.sg

ABSTRACT

Blockchain rewriting with fine-grained access control allows a user to create a transaction associated with a set of attributes, while another user (or modifier) who possesses enough rewriting privileges from a trusted authority satisfying the attribute set can rewrite the transaction. However, it lacks accountability and is not designed for open blockchains that require no trust assumptions. In this work, we introduce accountable fine-grained blockchain rewriting in a permissionless setting. The property of accountability allows the modifier's identity and her rewriting privileges to be held accountable for the modified transactions in case of malicious rewriting (e.g., modify the registered content from good to bad). We first present a generic framework to secure blockchain rewriting in the permissionless setting. Second, we present an instantiation of our approach and show its practicality through evaluation analysis. Last, we demonstrate that our proof-of-concept implementation can be effectively integrated into open blockchains.

KEYWORDS

Blockchain Rewriting, Accountability, Open Blockchains

1 INTRODUCTION

Blockchains have received tremendous attention from research communities and industries in recent years. The concept was first introduced in the context of Bitcoin [42], where all payment transactions are appended in a public ledger, and each transaction is ordered and verified by network nodes in a peer-to-peer manner. Blockchain ledgers grow by one block at a time, where the new block in the chain is decided by a consensus mechanism (e.g., Proof-of-Work in Bitcoin [31]) executed by the network nodes. Usually, blockchains deploy hash-chains as an append-only structure, where the hash of a block is linked to the next block in the chain. Each block includes a set of valid transactions which are accumulated into a single hash value using the Merkle tree [41], and each transaction contains certain content which needs to be registered in the blockchain.

Blockchains are designed to be immutable, such that the registered content cannot be modified once they are appended. However, blockchain rewriting is often required in practice, or even legally necessary in data regulation laws such as GDPR in Europe [4]. Since the platform is open, it is possible some users append transactions into a chain containing illicit content such as sensitive information, stolen private keys, and inappropriate videos [39, 40]. The existence of illicit content in the chain could pose a challenge to law enforcement agencies like Interpol [53].

Blockchain rewriting is usually realized by replacing a standard hash function, used for generating transaction hash in the blockchain, by a trapdoor-based chameleon hash [34]. Then the users, who have the same privilege (i.e., hold the trapdoor), can modify a transaction. In other words, the same transaction can be modified by users with the same privileges only. Nonetheless, for most real-life blockchain applications, blockchain rewriting with fine-grained access control is desired so that various rewriting privileges can be granted to different modifiers and that the same transaction can be modified by users with different privileges. For blockchain rewriting with fine-grained access control, a user first associates his transaction with a set of attributes. Then, any user possessing a chameleon trapdoor can modify the transaction if her access privilege corresponding to the chameleon trapdoor satisfies the embedded attribute set. It is possible that multiple users with different access privileges can modify a same transaction if their privileges satisfy the set of attributes associated with the transaction.

Motivation. Blockchain rewriting with fine-grained access control has been recently studied in the permissioned setting [19]; however, the proposed solution is not suitable for permissionless (i.e., open) blockchains for two reasons: 1) It requires a trusted authority to distribute rewriting privileges. 2) It lacks accountability as modifiers may rewrite the blockchain without being identified. The main motivation of this work is to extend blockchain rewriting with fine-grained access control to open blockchains such as Bitcoin [1] and Ethereum [2], which assume no trusted authority for managing access privileges.

Blockchain rewriting with fine-grained access control requires accountability especially in the permissionless setting. It is more critical for open blockchains, as there is no trusted authority to identify which modifier is responsible for any maliciously modified transaction (e.g., a modifier may rewrite the registered content in a transaction from good to bad). Besides, if modifiers attempt to generate an access device (or blackbox) that accumulates various kinds of rewriting privileges, and distribute or sell it to the public, it is more challenging to identify the responsible rewriting privileges given a modified transaction. This is because any user may rewrite a transaction successfully if he holds that access blackbox. Therefore, public accountability in this work encompasses two aspects: 1) identify a modifier in case of malicious modification on a transaction, and 2) identify the modifier's rewriting privileges even if an access blackbox is used.

Let us consider an open blockchain system that includes multiple committees, and each committee contains multiple users appointed for granting rewriting privileges. For ease of exposition, we assume there exists only one committee and one modifier for modifying mutable transactions during an epoch (i.e., a fixed interval of time). We have two blockchain rewriting cases: 1) a modifier rewrites a transaction using her access privilege granted from a committee, and the modifier may have various access privileges by joining different committees at different epochs. 2) an unauthorized user has no rewriting privilege from any committee; he may still rewrite a transaction using an access blackbox that accumulates various modifiers' access privileges. In such cases, it is desired to hold the rewriting privileges accountable for the modified transactions. Overall, if a transaction is maliciously modified, any public user can identify the responsible rewriting privilege that links to either a designated modifier or an unauthorized user with an access blackbox. Public accountability allows the modifiers (including the unauthorized users with access blackboxes) and the rewriting privileges to be held accountable for the modified transactions in open blockchain systems.

This Work. We introduce a new framework of accountable fine-grained blockchain rewriting, which is used to secure blockchain rewriting in a permissionless setting. The proposed framework achieves strong security without trusted authority and achieves public accountability simultaneously. Strong security without trusted authority indicates that the proposed framework remains secure even when attackers can compromise a threshold number of users in any committee. Public accountability means that any user in the public can link a modified transaction to a modifier and the responsible rewriting privilege.

We now explain our key technical insights. First, we rely on dynamic proactive secret sharing (DPSS) [38] to remove the trusted authority and achieve strong security. We replace the trusted authority by a committee of multiple users, where each user holds a share of the trust. We allow any user to join in and leave from a committee in any time epoch. Since we use key-policy attribute-based encryption (KP-ABE) [29] to ensure fine-grained access control, we highlight the following points: 1) The master secret key is split into multiple key shares so that each user in a committee holds a single key share. 2) A certain number of shareholders in a committee collaboratively recover the master secret key and distribute access privileges to modifiers. 3) Any user can freely join/leave a committee, and the master secret key remains fixed across different committees. We achieve strong security because the master secret key remains secure even if the attackers are allowed to compromise a threshold number of shareholders in any committee.

Second, we achieve public accountability using a digital signature scheme, a commitment scheme, and KP-ABE with public traceability (ABET). We show the purpose of using those primitives as follows: 1) The digital signature helps the public link a modified transaction to a modifier (or modifier's public key), as she signs the modified transaction using her signing key, and the signed transaction is publicly verifiable. 2) The commitment scheme helps the public link modifiers' public keys with committees. The intention is to show a modifier is indeed obtained a rewriting privilege from a committee. 3) If an unauthorized user holds an access blackbox, ABET helps the public to obtain a set of rewriting privileges from interacting with the access blackbox due to ABET's public traceability. Since there is no existing ABET can be applied to this work, we propose a new ABET scheme, which is suitable for decentralized systems such as open blockchains.

Our Contributions. The major contributions of this work are summarized as follows.

- *Generic Framework.* We introduce a new generic framework of accountable fine-grained blockchain rewriting, which is based on the chameleon hash function. A unique feature of this framework is that it allows the fine-grained blockchain rewriting to be performed in the permissionless setting.
- *Public Accountability.* We introduce a new notion called public accountability, such that modifiers' public keys and rewriting privileges are held accountable for the modified transactions, which is essential to blockchain rewriting because it helps thwart malicious rewriting.
- *Practical Instantiation.* We present a practical instantiation, and our evaluation analysis validates its practicality. We present an efficient ABET, which is of independent interest. The proposed ABET scheme is the first KP-ABE scheme with public traceability designed for decentralized systems.
- *Integration to Open Blockchains.* The proof-of-concept implementation shows that blockchain rewriting based on our approach incurs almost no overhead to chain validation when compared to the immutable blockchain.

2 PRELIMINARY

In this section, we present the complexity assumptions and the building blocks, which are used in our proposed generic construction and instantiation.

2.1 Complexity Assumptions

Bilinear Maps. Let (g, h) denote two group generators, which takes a security parameter λ as input and outputs a description of a group $\mathbb{G}, \mathbb{H}$. We define the output of (g, h) as $(q, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, \hat{e})$, where q is a prime number, $\mathbb{G}, \mathbb{H}$ and $\mathbb{G}_T$ are cyclic groups of order q, and $\hat{e} : \mathbb{G} \times \mathbb{H} \to \mathbb{G}_T$ is a bilinear map such that: (1) Bilinearity: $\forall g, h \in \mathbb{G}$ and $a, b \in \mathbb{Z}_q$, we have $\hat{e}(g^a, h^b) = e(g, h)^{ab}$; (2) Non-degeneracy: $\exists g \in \mathbb{G}$ such that $\hat{e}(g, h)$ has order q in $\mathbb{G}_T$. We assume that group operations in $\mathbb{G}, \mathbb{H}$ and $\mathbb{G}_T$ and bilinear map $\hat{e}$ are computable in polynomial time with respect to λ. We refer to $\mathbb{G}$ and $\mathbb{H}$ as the source groups and $\mathbb{G}_T$ as the target group.

We introduce a new assumption below, which is used to prove the semantic security of the proposed ABET scheme. The new assumption is proven secure in the generic group model [50]. We underline specific elements to show the differences between the new assumption and the original one.

Definition 2.1 (Extended q-type Assumption). Given group generators $g \in \mathbb{G}$ and $h \in \mathbb{H}$, define the following distribution:

$$
\begin{aligned}
\mathsf{Adv}_{\mathcal{A}}^{q'} &= |\Pr[\mathsf{Adv}(1^\lambda, \mathsf{par}, D, T_0) = 1] \\
&\quad - \Pr[\mathsf{Adv}(1^\lambda, \mathsf{par}, D, T_1) = 1]|, where \\
\mathsf{par} &= (q, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, \hat{e}, g, h) \leftarrow \mathsf{GroupGen}(1^\lambda) \\
a, b, c, d &\leftarrow \mathbb{Z}_q^*, s, \{z\} \leftarrow \mathbb{Z}_q; D = (g^a, \underline{h^b}, g^c, \\
&g^{(ac)^2}, \underline{g^{abd}, g^{d/ab}, h^{abd}, h^{abcd}, h^{d/ab}}, h^c, h^{cd/ab}, \\
&g^{z_i}, g^{acz_i}, g^{ac/z_i}, g^{a^2cz_i}, g^{b/z_i^2}, g^{b^2/z_i^2}, \forall i \in [q] \\
&g^{acz_i/z_j}, g^{bz_i/z_j^2}, g^{abcz_i/z_j}, g^{(ac)^2 z_i/z_j}, \forall i, j \in [q], \\
&i \neq j); T_0 = g^{abc}, T_1 = g^s.
\end{aligned}
$$

The extended q-type (or q'-type) assumption is secure if $\mathsf{Adv}_{\mathcal{A}}(\lambda)$ is negligible λ.

The detailed theorem and proof are shown in Appendix A. We also present an Extended Decisional Diffie-Hellman Assumption [52], which is used to prove the ciphertext anonymity of the proposed ABET scheme.

Definition 2.2 (Extended Decisional Diffie-Hellman (eDDH)). Given group generators $g \in \mathbb{G}$ and $h \in \mathbb{H}$, define the following distribution:

$$
\begin{aligned}
\mathsf{Adv}_{\mathcal{A}}^{\mathsf{eDDH}} &= |\Pr[\mathsf{Adv}(1^\lambda, \mathsf{par}, D, T_0) = 1] \\
&\quad - \Pr[\mathsf{Adv}(1^\lambda, \mathsf{par}, D, T_1) = 1]|, where \\
\mathsf{par} &= (q, \mathbb{G}, \mathbb{H}, \mathbb{G}_T, \hat{e}, g, h) \leftarrow \mathsf{GroupGen}(1^\lambda) \\
a, b, c &\leftarrow \mathbb{Z}_q^*, s \leftarrow \mathbb{Z}_q; D = (g^a, g^b, g^{ab}, h^c, h^{ab}, \\
&h^{1/ab}, h^{abc}); T_0 = h^{c/ab}, T_1 = h^s.
\end{aligned}
$$

The eDDH assumption is secure if $\mathsf{Adv}_{\mathcal{A}}(\lambda)$ is negligible in λ.

2.2 Attribute-based Encryption

Access Structure. Let $\mathcal{U}$ be an attribute universe. An access structure Λ is a collection of non-empty subsets of $\mathcal{U}$ (i.e., $\Lambda \subseteq 2^{\mathcal{U}} \setminus \{\phi\}$). It is called monotone if $\forall B, C$: if $B \in \Lambda$ and $B \subseteq C$ then $C \in \Lambda$.

Monotone Span Program (MSP). A secret-sharing scheme $\prod$ with domain of secrets realizing access structure is called linear over $\mathbb{Z}_q$ if: 1) The shares of a secret $s \in \mathbb{Z}_q$ for each attribute form a vector over $\mathbb{Z}_q$; 2) For each access structure Λ on δ, there exists a matrix $\mathbf{M}$ with n_1 rows and n_2 columns called the share-generating matrix for $\prod$. For $\mu = 1, ..., n_1$, we define a function π labels row μ of $\mathbf{M}$ with attribute $\pi(\mu)$ from the attribute universe $\mathcal{U}$. We consider the column vector $\vec{v} = (s, r_2, ..., r_{n_2})^\top$, where $s \in \mathbb{Z}_q$ is the secret to be shared and $r_2, ..., r_{n_2} \in \mathbb{Z}_q$ are chosen at random. Then $\mathbf{M}\vec{v} \in \mathbb{Z}_q^{n_1 \times 1}$ is the vector of n_1 shares of the secret s according to $\prod$. The share $(\mathbf{M}\vec{v})_\mu$ belongs to attribute $\pi(\mu)$, where $\mu \in [n_1]$.

According to [12], every linear secret-sharing scheme has the linear reconstruction property, which is defined as follows: we assume that $\prod$ is a MSP for the access structure Λ, $\delta' \in \Lambda$ is an authorized set and let $I \subset \{1, 2, ..., n_1\}$ be defined as $I = \{\mu \in [n_1] \wedge \pi(\mu) \in \delta'\}$. There exists the constants $\{\gamma_\mu \in \mathbb{Z}_q\}_{\mu \in I}$ such that for any valid share $\{\lambda_\mu = (\mathbf{M}\vec{v})_\mu\}_{\mu \in I}$ of a secret s according to $\prod$, $\sum_{\mu \in I} \gamma_\mu \lambda_\mu = s$. Meanwhile, such constants $\{\gamma_\mu\}_{\mu \in I}$ can be found in time polynomial in the size of the share-generating matrix $\mathbf{M}$. For any unauthorized set δ'', no such $\{\gamma_\mu\}$ exist.

Attribute-Based Encryption with Public Traceability. It consists of the following algorithms. We assume an index space $\{1, \cdots, k\}$, where k denotes the maximal number of the index. This definition is inspired by [36].

- Setup(1^λ): It takes a security parameter λ as input, outputs a master key pair (msk, mpk).
- KeyGen(msk, Λ): It takes the master secret key msk, an access policy as input, outputs a decryption key ssk_{Λ_i}, which is assigned by a unique index i.
- Enc(mpk, m, δ, j): It takes the master public key mpk, a message m, a set of attributes $\delta \in \mathcal{U}$, and an index $j \in \{1, k + 1\}$ as input, outputs a ciphertext C. Note that C contains δ, not index j.
- Dec(mpk, C, ssk_{Λ_i}): It takes the master public key mpk, a ciphertext C, and the decryption key ssk_{Λ_i} as input, outputs the message m if $1 = \Lambda_i(\delta) \wedge j \leq i$.
- Trace(mpk, $\mathcal{D}, \epsilon$): It takes master public key pair mpk, a policy-specific decryption device/blackbox $\mathcal{D}$, and a parameter $\epsilon > 0$ as input, outputs a set of indexes $\mathbb{K}_T \in \{1, \cdots, k\}$, where ϵ is polynomially related to λ, and $\mathbb{K}_T$ denotes the index set of the accused decryption keys.

Public Traceability. Given a policy-specific decryption device/blackbox that includes a set of decryption keys, the tracing algorithm, which treats the decryption blackbox as an oracle, can identify the accused decryption keys that have been used in constructing the decryption blackbox. The decryption blackbox is associated with a specific access policy $\Lambda_{\mathcal{D}}$. Informally speaking, any public user can generate a ciphertext on a message under a set of attributes that satisfies $\Lambda_{\mathcal{D}}$, and an index $j \in \{1, \cdots, k + 1\}$. Then, the public sends the ciphertext to the decryption blackbox and checks whether the decryption is successful. If decryption succeeds, the public outputs the accused index j; otherwise, the public generates a new ciphertext

under another attribute set and index. The public continues this process until finding a set of accused indexes $\mathbb{K}_T \in \{1, \cdots, k\}$. The tracing algorithm's formal definition is referred to [15, 16], which is analogous to the traitor tracing algorithm used in the broadcast encryption.

The ABET scheme requires that the encryptor generates a ciphertext on a message associated with a set of attributes and a hidden index j. The decryptor can decrypt the message if the set of attributes is satisfied by her access policy, and $j \leq i$. We stress that the hidden index (or index-hiding) is critical to ABET. On the one hand, the index-hiding ensures that a ciphertext generated by the encryptor using an index j reveals no information about j. On the other hand, the public user can pick a possible accused index $j \in \{1, \cdots, k+1\}$ in generating ciphertext for public tracing. In this work, we call index-hiding as ciphertext anonymity. We denote policy-specific decryption blackbox as access blackbox because it accumulates various rewriting privileges for blockchain rewriting.

2.3 Digital Signature

A digital signature scheme Σ = (Setup, KeyGen, Sign, Verify) is homomorphic, if the following conditions are held.

- *Simple Key Generation.* It requires $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(pp)$ and $pp \leftarrow \mathsf{Setup}(1^\lambda)$, where pk is derived from sk via a deterministic algorithm $\mathsf{pk} \leftarrow \mathsf{KeyGen}'(pp, \mathsf{sk})$.
- *Linearity of Keys.* It requires $\mathsf{KeyGen}'(pp, \mathsf{sk}+\Delta(\mathsf{sk})) = M_{\mathsf{pk}}(pp, \mathsf{KeyGen}'(pp, \mathsf{sk}), \Delta(\mathsf{sk}))$, where M_{pk} denotes a deterministic algorithm which takes pp, a verification key pk and a "shifted" value $\Delta(\mathsf{sk})$ as input, outputs a "shifted" verification key pk'. Δ denotes the difference or shift between two keys.
- *Linearity of Signatures.* Two distributions are identical: $\{\sigma' \leftarrow \mathsf{Sign}(pp, \mathsf{sk} + \Delta(\mathsf{sk}), m)\}$ and $\{\sigma' \leftarrow M_\Sigma(pp, \mathsf{pk}, m, \sigma, \Delta(\mathsf{sk}))\}$, where $\sigma \leftarrow \mathsf{Sign}(pp, \mathsf{sk}, m)$, and M_Σ denotes a deterministic algorithm which takes pp, a verification key pk, a message-signature pair (m, σ) and a "shifted" value $\Delta(\mathsf{sk})$ as input, outputs a "shifted" signature σ'.
- *Linearity of Verifications.* It requires $\mathsf{Verify}(pp, M_{\mathsf{pk}}(pp, \mathsf{pk}, \Delta(\mathsf{sk})), m, M_\Sigma(pp, \mathsf{pk}, m, \sigma, \Delta(\mathsf{sk}))) = 1$, and $\mathsf{Verify}(pp, \mathsf{pk}, m, \sigma) = 1$.

The Schnorr signature scheme [47] satisfies the homomorphic properties regarding keys and signatures. We rely on this homomorphic property to find the connection between a transaction and its modified versions.

2.4 Dynamic Proactive Secret Sharing

A dynamic proactive secret sharing DPSS consists of Share, Redistribute, and Open [11] protocols. It allows a dealer to share a secret s among a group of n_0 users such that the secret is secure against a *mobile* adversary, and allow any group of n_0-t users to recover the secret, where t denotes a threshold. The proactive security means that the execution of the protocol is divided into phases (or epochs) [44], and a mobile adversary is allowed to corrupt users across all epochs, under the condition that no more than a threshold number of users are corrupted in any given epoch. The Redistribute protocol prevents the mobile adversary from disclosing or destroying the secret and allows the set of the users and the threshold to change. Assuming that for each epoch i, no more than t users are corrupted during epoch i, the following three properties hold:

- *Termination:* All honest users engaged in the protocol complete each execution of Share, Redistribute, and Open.
- *Correctness:* All honest users output a secret s' upon completing of Open, such that $s' = s$ if the dealer was honest during the execution of Share.
- *Secrecy:* If the dealer is honest, then s leaks no information to the adversary.

The definition described in [11] is for information-theoretically (or perfectly) secure protocols. We merely require DPSS to be computationally secure due to the instantiation used in this work has computational security. Dynamic allows the set of users in a group (or committee) to be dynamically changed, which is useful in the permissionless blockchains. The Redistribute protocol has two processes: resharing the key shares to change the committee membership and threshold, and updating the key shares across epochs to tackle mobile adversary.

- *Resharing the Key Shares [20].* We rely on a bivariate polynomial to share a secret s: $f(x, y) = \underline{s} + a_{0,1}x + a_{1,0}y + a_{1,1}xy + \cdots + a_{t_x,t_y}x^{t_x}y^{t_y}$, where t_x, t_y denote different thresholds. So there are two ways to share the secret s:
  (1) If we fix $y = 0$, then the key shares include $\{f(i_0, 0), f(i_1, 0), \cdots, f(i_{t_x}, 0)\}$;
  (2) If we fix $x = 0$, then the key shares include $\{f(0, j_0), f(0, j_1), \cdots, f(0, j_{t_y})\}$.
  We show how to transfer the ownership of the shareholders from committee A to committee B. First, we distribute key shares $\{f(i, y)\}$ to all users in committee A. Second, each user in committee A generates a set of temporary shares by running a secret sharing scheme (SSS) (e.g., Shamir's [49]) on his own key share. In other words, his key share is the secret for SSS. Third, users in committee A send those temporary shares to users in committee B. Now, users in the committee B accumulate the received temporary shares and obtain another form of key shares $\{f(x, j)\}$ via interpolation of t_y temporary shares. To this end, the transfer between the two committees is successful. Note that either key shares $\{f(i, y)\}$ or $\{f(x, j)\}$ can be used to recover the secret s.
- *Updating the Key Shares [30].* Suppose that a bivariate polynomial is used to share the secret s: $f(x, y) = \underline{s} + a_{0,1}x + a_{1,0}y + a_{1,1}xy + a_{0,2}x^2 + a_{2,0}y^2 + a_{2,2}x^2y^2 + \cdots + a_{t_x,t_y}x^{t_x}y^{t_y}$. To update $f(x, y)$, we need another bivariate polynomial: $f'(x, y) = \underline{0} + a'_{0,1}x + a'_{1,0}y + a'_{1,1}xy + \cdots + a'_{t_x,t_y}x^{t_x}y^{t_y}$, which takes 0 as the secret. The reason is that the secret s in $f(x, y)$ will not be changed after updating by $f'(x, y)$. A crucial point is, we allow users in a new committee to collaboratively generate a polynomial $f'(x, y)$, thus the shareholders between old and new committees become independent (to ensure proactive security). Note that t_x may not equal to t_y because the threshold between committees can be different, and we call it asymmetric bivariate polynomial.

2.5 Polynomial Commitments

A simplified version of polynomial commitment scheme [32] is shown as follows.

- $\mathsf{Setup}(1^\lambda, t)$: It takes a security parameter λ and t as input, outputs a key pair (msk, mpk), where $\mathsf{msk} = \alpha$, $\mathsf{mpk} = (g, g^\alpha, \cdots, g^{\alpha^t}, h, h^\alpha, \hat{e})$.

- Commit(mpk, $f(x)$): It takes the public key mpk, and a polynomial $f(x) = a_0 + a_1x + a_2x^2 + \cdots + a_tx^t$ as input, outputs $C = \prod_{j=0}^{t}(g^{\alpha^j})^{a_j}$ as the commitment to $f(x)$.
- CreateWitness(mpk, $C, f(x)$): It takes the public key mpk, and the polynomial $f(x)$ as input, outputs a tuple $(i, f(i), w_i)$. Specifically, it computes a polynomial $\frac{f(x)-f(i)}{x-i}$ (note that the coefficients of the resulting quotient polynomial are $(\widehat{a_0}, \widehat{a_1}, \cdots, \widehat{a_t})$), and a witness $w_i = \prod_{j=0}^{t}(g^{\alpha^j})^{\widehat{a_j}}$.
- VerifyEval(mpk, $C, i, f(i), w_i$): It takes the public key mpk, a commitment C, and the tuple $(i, f(i), w_i)$ as input, outputs 1 if $\grave{e}(C/g^{f(i)}, h) = \grave{e}(w_i, h^{\alpha}/h^i)$.

The witness w_i proves that $f(i)$ is a correct evaluation at $i \in \mathbb{Z}_q$, without revealing the polynomial $f(x)$. The binding property is based on the t-Strong Diffie-Hellman assumption [32], while the hiding property is based on the Discrete Logarithm (DL) assumption. If the KZG commitment scheme is used in DPSS [38], we can hold users accountable in a committee. In particular, the KZG commitment scheme is publicly verifiable if we append users' commitments and witnesses to blockchain. We assume that they are confirmed in the blockchain using Proof-of-Work (PoW) consensus (it is not difficult to extend this assumption to other consensus like Proof of Stake [6], Proof of Space [24]). In this work, a polynomial commitment scheme is used in DPSS.

3 MODELS AND DEFINITIONS

3.1 System Model

The system model involves three types of entities: user, modifier, and miner, in which the entities can intersect, such as a user can be a modifier and/or a miner. The communication model considers both on-chain and off-chain settings. The on-chain setting is the permissionless blockchain, where *read* is public, but *write* is granted to anyone who can show PoW. The off-chain setting assumes that every user has a point-to-point (P2P) channel with every other users. One may use Tor or transaction ghosting to establish a P2P channel [38], and further detail is given in Appendix B. Such P2P channel works in a synchronous model, i.e., any message sent via this channel is received within a known bounded time-period.

The system proceeds in fixed time periods called epochs. In the first epoch, a committee election protocol (e.g., Algorand's $BA*$ protocol [28], or other methods [37, 56]) is executed, so that a set of users can agree on an initial committee with Byzantine fault tolerance (e.g., up to 1/3 malicious members). The secret s in the initial committee can be generated by an honest user (e.g., committee leader) or in a distributed fashion [27]. The secret s is shared among the committee members. Similarly, the setup of the commitment scheme can be performed by an honest user in the initial committee.

In Figure 1, a blockchain is generated by users who append their hashed contents to the blockchain. Later, modifiers with enough rewriting privileges are required to rewrite the hashed contents. We stress that the link of hash-chain remains intact after rewriting, and the secret remains fixed across different committees. We assume at most n users (i.e., protocol participants) exist in each epoch. We consider k dynamic committees, each of which has a varying number of committee members, and we denote $n_0 \leq n$
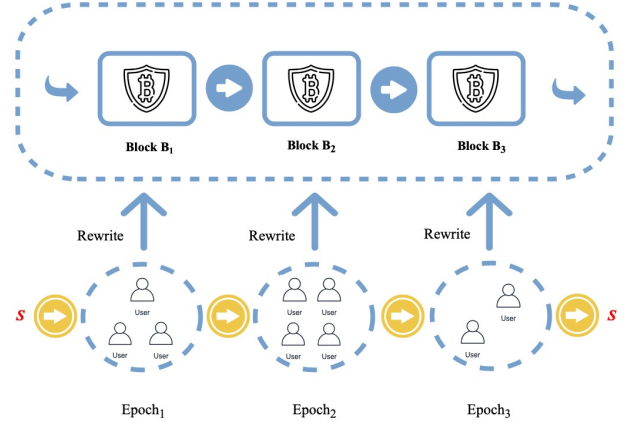

Figure 1: Blockchain rewriting with dynamic committees. Users may join in or leave from a committee, and a designated modifier in a committee can rewrite the blockchain. The secret s remains fixed across different committees.

as a committee's size. The parameters n and k are independent. We also consider dynamic churn (i.e., join/leave) of the protocol participants. In particular, we do not assume that k committees exist in each different epoch (or we allow several committees to exist in the same epoch).

Remark. To prevent a malicious user from controlling a committee by launching Sybil attacks [22], we rely on the PoW-based identity generation mechanism [9, 37]. The mechanism allows all users to establish their identities in a committee, yet limiting the number of Sybil identities created by a malicious user. In Elastico [37], each user locally generates/establishes an identity consisting of a public key, an IP address, and a PoW solution. The user must solve a PoW puzzle which has publicly verifiable solutions to generate the final component of the identity. A PoW solution also allows other committee members to verify and accept the identity of a user. Because solving PoW requires computation, the number of identities that the malicious user can create is limited by a fraction of malicious computational power. One can refer to [37, 55, 56] for the detailed discussion on Byzantine fault resiliency.

3.2 Definition

An accountable fine-grained chameleon hash with dynamic committees consists of the following algorithms.

- Setup(1^{λ}): It takes a security parameter λ as input, outputs a master key pair (msk, mpk). Note that msk is shared by committee members.
- KeyGen(C_i^e, Λ): It takes a committee C_i^e, and a policy Λ as input, outputs a secret key sk_{Λ_i}. The committee index $i \in \{1, \cdots, k\}$, where k denotes the total number of committees.
- Hash(mpk, m, δ, j): It takes the master public key mpk, a message $m \in \mathcal{M}$, a set of attributes $\delta \in \mathcal{U}$, and an index $j \in \{1, \cdots, k+1\}$ as input, outputs a chameleon hash h, a randomness r, and a signature σ. Note that $\mathcal{M} = \{0, 1\}^*$ denotes a general message space.

$$
\begin{aligned}
&\text{Experiment } \mathsf{Exp}_{\mathcal{A}}^{\mathsf{IND}}(\lambda)\\
&(C_i^e, \mathsf{mpk}) \leftarrow \mathsf{Setup}(1^\lambda), b \leftarrow \{0,1\}\\
&b^* \leftarrow \mathcal{A}^{\mathsf{HashOrAdapt}(C_i^e, \cdots, b)}(\mathsf{msk})\\
&\quad \text{where } \mathsf{HashOrAdapt}(C_i^e, \cdots, b) \text{ on input } m, m', \delta, \Lambda, j, j':\\
&\quad\quad \mathsf{sk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}(C_i^e, \Lambda)\\
&\quad\quad (h_0, r_0, \sigma_0) \leftarrow \mathsf{Hash}(\mathsf{mpk}, m', \delta, j')\\
&\quad\quad (h_1, r_1, \sigma_1) \leftarrow \mathsf{Hash}(\mathsf{mpk}, m, \delta, j)\\
&\quad\quad r_1 \leftarrow \mathsf{Adapt}(\mathsf{sk}_{\Lambda_i}, m, m', h_1, r_1, \sigma_1)\\
&\quad\quad \text{return } \perp \text{ if } r_0 = \perp \vee r_1 = \perp\\
&\quad\quad \text{return } (h_b, r_b, \sigma_b)\\
&\text{return } 1, \text{ if } b^* = b; \text{ else, return } 0.
\end{aligned}
$$

Figure 2: Indistinguishability.

- Verify($\mathsf{mpk}, h, m, r, \sigma$): It takes the master public key mpk, chameleon hash h, message m, randomness r, signature σ as input, output a bit $b \in \{0, 1\}$.
- Adapt($\mathsf{sk}_{\Lambda_i}, h, m, m', r, \sigma$): It takes the secret key sk_{Λ_i}, chameleon hash h, messages m and m', randomness r, and signature σ as input, outputs r' and σ' if $1 = \Lambda(\delta)$ and $i \leq j$.
- Judge(mpk, T'): It takes the master public key mpk, a modified transaction T' as input, outputs a linked transaction-committee pair (T', C_i^e), where $T' = (h, m', r', \sigma')$. It means a user with a rewriting privilege from committee C_i^e has modified transaction $T = (h, m, r, \sigma)$.

Correctness. The definition is *correct* if for all security parameters λ, for all $\delta \in \mathcal{U}$, all keys $(\mathsf{msk}, \mathsf{mpk}) \leftarrow \mathsf{Setup}(1^\lambda)$, for all $\delta \in \Lambda$, for all $i \leq j$, for all $\mathsf{sk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}(C_i^e, \Lambda)$, for all $m \in \mathcal{M}$, for all $(h, r, \sigma) \leftarrow \mathsf{Hash}(\mathsf{mpk}, m, \delta, j)$, for all $m' \in \mathcal{M}$, for all $r' \leftarrow \mathsf{Adapt}(\mathsf{sk}_{\Lambda_i}, m, m', h, r, \sigma)$, we have $1 = \mathsf{Verify}(\mathsf{mpk}, h, m, r, \sigma) = \mathsf{Verify}(\mathsf{mpk}, h, m', r', \sigma')$.

3.3 Security Model

We consider three security guarantees, including indistinguishability, adaptive collision-resistance, and accountability.

- *Indistinguishability.* Informally, an adversary cannot decide whether for a chameleon hash its randomness was freshly generated using Hash algorithm or was created using Adapt algorithm even if the secret key is known. We define a formal experiment between an adversary $\mathcal{A}$ and a simulator $\mathcal{S}$ in Figure 2. The security experiment allows $\mathcal{A}$ to access a left-or-right HashOrAdapt oracle, which ensures that the randomness does not reveal whether it was obtained from Hash or Adapt algorithm. The hashed messages are adaptively chosen from the same message space $\mathcal{M}$ by $\mathcal{A}$.
We require $1 = \Lambda(\delta)$ and $\mathsf{Verify}(\mathsf{mpk}, m', h_0, r_0, \sigma_0) = \mathsf{Verify}(\mathsf{mpk}, m, h_1, r_1, \sigma_1) = 1$. Note that msk is shared by committee members, such that $C_i^e = \{s_i^e\}^{n_0}$. We define the advantage of the adversary as

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{IND}}(\lambda) = |\Pr[\mathsf{Exp}_{\mathcal{A}}^{\mathsf{IND}}(1^\lambda) \rightarrow 1] - 1/2|.$$

Definition 3.1. The proposed generic framework is indistinguishable if for any probabilistic polynomial-time (PPT) $\mathcal{A}$, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{IND}}(\lambda)$ is negligible in λ.

- *Adaptive Collision-Resistance.* Informally, a mobile adversary can find collisions for a chameleon hash if she possesses a secret key satisfies an attribute set associated with the chameleon hash (this condition is modelled by KeyGen' oracle). We define a formal

$$
\begin{aligned}
&\text{Experiment } \mathsf{Exp}_{\mathcal{A}}^{\mathsf{ACR}}(\lambda)\\
&(C_i^e, \mathsf{mpk}) \leftarrow \mathsf{Setup}(1^\lambda), Q_1, Q_2, Q_3 \leftarrow \emptyset\\
&(m^*, r^*, m^{*'}, r^{*'}, h^*, \sigma^*, \sigma^{*'}) \leftarrow \mathcal{A}^O(\mathsf{mpk})\\
&\quad \text{where } O \leftarrow \{\mathsf{KeyGen}', \mathsf{KeyGen}'', \mathsf{Hash}', \mathsf{Adapt}', \mathsf{Corrupt}\}\\
&\quad \text{and } \mathsf{KeyGen}'(C_i^e, \cdot) \text{ on input } \Lambda:\\
&\quad\quad \mathsf{msk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}(C_i^e, \Lambda)\\
&\quad\quad Q_1 \leftarrow Q_1 \cup \{\Lambda\}\\
&\quad\quad \text{return } \mathsf{msk}_{\Lambda_i}\\
&\quad \text{and } \mathsf{KeyGen}''(C_i^e, \cdot) \text{ on input } \Lambda:\\
&\quad\quad \mathsf{msk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}(C_i^e, \Lambda)\\
&\quad\quad Q_2 \cup \{(i, \mathsf{msk}_{\Lambda_i})\}\\
&\quad\quad i \leftarrow i + 1\\
&\quad \text{and } \mathsf{Hash}'(\mathsf{mpk}, \cdots) \text{ on input } m, \delta, j:\\
&\quad\quad (h, r, \sigma) \leftarrow \mathsf{Hash}(\mathsf{mpk}, m, \delta, j)\\
&\quad\quad Q_3 \leftarrow Q_3 \cup \{(h, m, \delta)\}\\
&\quad\quad \text{return } (h, r, \sigma)\\
&\quad \text{and } \mathsf{Adapt}'(\mathsf{mpk}, \cdots) \text{ on input } m, m', h, r, i, \sigma:\\
&\quad\quad \text{return } \perp, \text{ if } (i, \mathsf{msk}_{\Lambda_i}) \notin Q_2 \text{ for some } \mathsf{msk}_{\Lambda_i}\\
&\quad\quad r' \leftarrow \mathsf{Adapt}(\mathsf{msk}_{\Lambda_i}, m, m', h, r, \sigma)\\
&\quad\quad \text{if } (h, m, \delta) \in Q_3 \text{ for some } \delta,\\
&\quad\quad \text{let } Q_3 \leftarrow Q_3 \cup \{(h, m', \delta)\}\\
&\quad\quad \text{return } r'\\
&\quad \text{and } \mathsf{Corrupt}(\mathsf{mpk}, \cdot) \text{ on input } C_i^e:\\
&\quad\quad \text{return } \{s_i^e\}^t\\
&\text{return } 1, \text{ if}\\
&1 = \mathsf{Verify}(\mathsf{mpk}, m^*, h^*, r^*, \sigma^*) = \mathsf{Verify}(\mathsf{mpk}, m^{*'}, h^*, r^{*'}, \sigma^{*'})\\
&\wedge (h^*, \cdot, \delta) \in Q_3, \text{ for some } \delta \wedge m^* \neq m^{*'}\\
&\wedge \delta \cap Q_1 = \emptyset \wedge (h^*, m^*, \cdot) \notin Q_3\\
&\text{else, return } 0.
\end{aligned}
$$

Figure 3: Adaptive Collision-Resistance.

$$
\begin{aligned}
&\text{Experiment } \mathsf{Exp}_{\mathcal{A}}^{\mathsf{ACT}}(\lambda)\\
&(C_i^e, \mathsf{mpk}) \leftarrow \mathsf{Setup}(1^\lambda), Q \leftarrow \emptyset\\
&T^* \leftarrow \mathcal{A}^{\mathsf{Judge}'(C_i^e, \cdots)}(\mathsf{mpk})\\
&\quad \text{where } \mathsf{Judge}'(C_i^e, \cdots) \text{ on input } T, \Lambda, m':\\
&\quad\quad \mathsf{msk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}(C_i^e, \Lambda)\\
&\quad\quad r' \leftarrow \mathsf{Adapt}(\mathsf{msk}_{\Lambda_i}, T, m')\\
&\quad\quad Q \leftarrow Q \cup \{(T, T')\}\\
&\quad\quad (T', C_i^e) \leftarrow \mathsf{Judge}(C_i^e, T)\\
&\quad\quad \text{return } (T', C_i^e)\\
&\text{return } 1, \text{ if } (T^*, C^*) \wedge (T^* \notin Q \vee \mathsf{pk}^* \notin C^*)\\
&\text{else, return } 0.
\end{aligned}
$$

Figure 4: Accountability.

experiment in Figure 3. We allow $\mathcal{A}$ to see collisions for arbitrary access policies (i.e., KeyGen" and Adapt' oracles). We also allow $\mathcal{A}$ to corrupt a threshold number of shareholders (i.e., Corrupt oracle) in a committee. Note that the key shares can be transferred between committees while msk is fixed.

$\mathcal{A}$ is not allowed to corrupt more than a threshold number of shareholders in any committee. We define the advantage of the adversary as

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{ACR}}(\lambda) = \Pr[\mathsf{Exp}_{\mathcal{A}}^{\mathsf{ACR}}(1^\lambda) \rightarrow 1].$$

Definition 3.2. The proposed generic framework is adaptively collision-resistant if for any PPT $\mathcal{A}$, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{ACR}}(\lambda)$ is negligible in λ.

- *Accountability.* Informally, an adversary cannot generate a bogus message-signature pair for a chameleon hash, which links a user to an accused committee, but the user has never participated in the accused committee. We define a formal experiment in Figure 4. We allow $\mathcal{A}$ to see whether a modified transaction links to a committee (i.e., Judge' oracle). Let set Q record the transactions produced by the Judge' oracle.

We denote $T = (h, m, r, \sigma)$ and $T' = (h, m', r', \sigma')$ as original and modified transactions with respect to chameleon hash h. We also denote the linked transaction-committee pair as (T', C_i^e). We define the advantage of the adversary as

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{ACT}}(\lambda) = \Pr[\mathsf{Exp}_{\mathcal{A}}^{\mathsf{ACT}}(1^\lambda) \to 1].$$

Definition 3.3. The proposed generic framework is accountable if for any PPT $\mathcal{A}$, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{ACT}}(\lambda)$ is negligible in λ.

4 GENERIC CONSTRUCTION

The proposed generic construction consists of the following building blocks.

- A chameleon hash scheme CH = (Setup, KeyGen, Hash, Verify, Adapt).
- An attribute-based encryption scheme with public traceability ABET = (Setup, KeyGen, Enc, Dec, Trace).
- A dynamic proactive secret sharing scheme DPSS = (Share, Redistribute, Open).
- A digital signature scheme Σ = (Setup, KeyGen, Sign, Verify).

High-level Description. We assume that every user has a key pair $(\mathsf{sk}, \mathsf{pk})$ and that users' public keys (i.e., users' identities) are known to all users in a committee. Meanwhile, each user possesses a set of attributes. In particular, more than a threshold number of users in a committee can collectively grant a rewriting privilege to a modifier based on her attribute set. A user with pk creates a transaction T that includes a chameleon hash, a ciphertext under his attribute set, and a signature (i.e., signs T using his secret key sk). A modifier with pk', who is granted the rewriting privilege from a committee, is allowed to rewrite the transaction T, and signs the modified transaction using her secret key sk'. We assume the t-out-of-n_0 DPSS scheme to be executed over off-chain P2P channels and let all k committees have the same parameters (t, n_0). The proposed construction is shown below.

- Setup(1^λ): A user takes a security parameter λ as input, outputs a public parameter PP = $(\mathsf{mpk}_{\mathsf{ABET}}, \mathsf{PP}_\Sigma, \mathsf{PP}_{\mathsf{CH}})$, and a secret key $\mathsf{msk}_{\mathsf{ABET}}$, where $(\mathsf{msk}_{\mathsf{ABET}}, \mathsf{mpk}_{\mathsf{ABET}}) \leftarrow \mathsf{Setup}_{\mathsf{ABET}}(1^\lambda)$, $\mathsf{PP}_\Sigma \leftarrow \mathsf{Setup}_\Sigma(1^\lambda)$, $\mathsf{PP}_{\mathsf{CH}} \leftarrow \mathsf{Setup}_{\mathsf{CH}}(1^\lambda)$. The key shares $\{s_0\}^{n_0} \leftarrow \mathsf{Share}_{\mathsf{DPSS}}(\mathsf{msk}_{\mathsf{ABET}})$ are distributed to users within committee C^0, where each user holds a key share, and a key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}_\Sigma(\mathsf{PP}_\Sigma)$.
- KeyGen(C_i^e, Λ): A group of $t+1$ users in committee C_i^e take their secret shares $\{s_e\}^{t+1}$ and a policy Λ as input, output a secret key sk_{Λ_i} for a modifier, where $\mathsf{sk}_{\Lambda_i} \leftarrow \mathsf{KeyGen}_{\mathsf{ABET}}(\mathsf{msk}_{\mathsf{ABET}}, \Lambda)$, $\mathsf{msk}_{\mathsf{ABET}} \leftarrow \mathsf{Open}_{\mathsf{DPSS}}(\{s_e\}^{t+1})$, and secret shares $\{s_e\}^{n_0} \leftarrow \mathsf{Redistribute}_{\mathsf{DPSS}}(\{s_{e-1}\}^{n_0})$. Note that one of $t+1$ users can be the modifier.
- Hash(PP, m, δ, j): A user appends a message m, a set of attributes δ, and an index j to the blockchain, performs the following operations
  (1) generate a chameleon hash $(h_{\mathsf{CH}}, r) \leftarrow \mathsf{Hash}_{\mathsf{CH}}(\mathsf{pk}_{\mathsf{CH}}, m)$, where $(\mathsf{sk}_{\mathsf{CH}}, \mathsf{pk}_{\mathsf{CH}}) \leftarrow \mathsf{KeyGen}_{\mathsf{CH}}(\mathsf{PP}_{\mathsf{CH}})$.
  (2) generate a ciphertext $C \leftarrow \mathsf{Enc}_{\mathsf{ABET}}(\mathsf{mpk}_{\mathsf{ABET}}, \mathsf{sk}_{\mathsf{CH}}, \delta, j)$, where $\mathsf{sk}_{\mathsf{CH}}$ denotes the encrypted message.
  (3) generate a message-signature pair (c, σ_Σ), where $\sigma_\Sigma \leftarrow \mathsf{Sign}_\Sigma(\mathsf{sk}, c)$, and message c is derived from sk and $\mathsf{sk}_{\mathsf{CH}}$.

  (4) output (h, m, r, σ), where $h \leftarrow (h_{\mathsf{CH}}, \mathsf{pk}_{\mathsf{CH}}, C)$, and $\sigma \leftarrow (c, \sigma_\Sigma)$.
- Verify(PP, h, m, r, σ): It outputs 1 if $1 \leftarrow \mathsf{Verify}_{\mathsf{CH}}(\mathsf{pk}_{\mathsf{CH}}, m, h_{\mathsf{CH}}, r)$ and $1 \leftarrow \mathsf{Verify}_\Sigma(\mathsf{pk}, c, \sigma_\Sigma)$, and 0 otherwise.
- Adapt($\mathsf{sk}_{\Lambda_i}, h, m, m', r, \sigma$): A modifier with a secret key sk_{Λ_i} and a new message m', performs the following operations
  (1) check $1 \overset{?}{=} \mathsf{Verify}(\mathsf{PP}, h, m, r, \sigma)$.
  (2) compute $\mathsf{sk}_{\mathsf{CH}} \leftarrow \mathsf{Dec}_{\mathsf{ABET}}(\mathsf{mpk}_{\mathsf{ABET}}, C, \mathsf{sk}_{\Lambda_i})$.
  (3) compute a new randomness $r' \leftarrow \mathsf{Adapt}_{\mathsf{CH}}(\mathsf{sk}_{\mathsf{CH}}, m, m', h, r)$.
  (4) generate a ciphertext $C' \leftarrow \mathsf{Enc}_{\mathsf{ABET}}(\mathsf{mpk}_{\mathsf{ABET}}, \mathsf{sk}_{\mathsf{CH}}, \delta, j)$.
  (5) generate a message-signature pair (c', σ'_Σ), where c' is derived from sk' and $\mathsf{sk}_{\mathsf{CH}}$.
  (6) output (h, m', r', σ'), where $h \leftarrow (h_{\mathsf{CH}}, \mathsf{pk}_{\mathsf{CH}}, C')$, and $\sigma' \leftarrow (c', \sigma'_\Sigma)$.
- Judge(PP, T'): It takes the public parameter PP, and a modified transaction T' as input, outputs a transaction-committee pair (T', C_i^e) if the modified transaction T' links to a committee C_i^e, where $T' = (h, m', r', \sigma')$.

Correctness. The Judge algorithm allows any public user to identify the responsible modifiers and committees given a modified transaction. The modifier (or modifier's public key) is publicly known because a digital signature is used in the construction. Below, we explain the judge process in detail.

First, any public user verifies a connection between a transaction T and its modified version T'. The connection can be established, as both message-signature pair (c, σ_Σ) in T and message-signature pair (c', σ'_Σ) in T', are derived from the same chameleon trapdoor $\mathsf{sk}_{\mathsf{CH}}$. Since different modifiers may modify the same transaction, the chameleon trapdoor $\mathsf{sk}_{\mathsf{CH}}$ is used in many modified versions of a transaction. Here, we consider a single modified transaction T' for simplicity.

Second, any public user obtains a set of accused committees from interacting with an access blackbox O, such that $\{C\} \leftarrow \mathsf{Trace}_{\mathsf{ABET}}(\mathsf{mpk}_{\mathsf{ABET}}, O, \epsilon)$. Specifically, the public sends a ciphertext encrypted a message under a set of attributes (that satisfies the access privilege involved in O) and a committee index $j \in \{1, \cdots, k+1\}$ to O. Then, the public outputs the committee index j (we call it accused committee) if decryption succeeds. The public repeats this tracing procedure until output all accused committees.

Third, if a user with pk' acts as a modifier in an accused committee, the public outputs (T', C). It means that a transaction T' is indeed modified by the user pk' whose rewriting privilege is granted from committee C. Because we allow the commitment scheme to be used in DPSS, the user pk' is held accountable in a committee. More specifically, user pk' joins in committee C by showing a commitment on his key share to other committee members, and further detail is given in the instantiation. If user pk' acts as modifiers for many accused committees, the public outputs $(T', \{C\})$. However, if user pk' did not join in any accused committees, the public still outputs the indexes of the accused committees. This is the second case of blockchain rewriting: an unauthorized user has no granted rewriting privileges from any committee but holds an access blackbox.

To conclude, we achieve public accountability via three steps: 1) Verify a modified transaction; 2) Find an accused committee; 3) Link the modified transaction to the accused committee. We also consider a committee has multiple modifiers with different rewriting privileges, but they should have the same committee index. In this case, the public still identifies the modifiers (i.e., holding different rewriting privileges) in the same committee as the modifiers sign the modified transactions using their secret keys.

Remark. One may notice that a modifier pk' can assign a new set of attributes δ' to a modified transaction T' in the Adapt algorithm. In other words, the attribute set associated with a mutable transaction can be dynamically updated to satisfy different security requirements in case blockchain system evolves. Therefore, we remark that the proposed generic construction supports fine-grained and flexible blockchain rewriting. Besides, the modifier can use a new index to create T'. One may also notice that such flexibility could be misused. The modifier pk' may intend to rewrite the transaction T with malicious content and change the rewriting privileges to disallow others to modify the transaction. We argue that more than a threshold number of committee members can collectively reset the transaction T's attribute set if such malicious behavior happens. If flexibility is not desired in accountable blockchain rewriting, one can let the transaction owner sign the embedded attribute set so that the modifier rewrites the transaction T only without changing its rewriting privilege.

The second remark is that we only allow blockchain rewriting that does not affect a transaction's past and future events. If a modifier removes a transaction entirely or changes spendable data of a transaction, it may lead to serious transaction inconsistencies in the chain [21]. The last remark is, the committee members can reconstruct the secret $\mathsf{msk}_{\mathsf{ABET}}$ as DPSS includes the Open algorithm. If any committee member rewrites blockchain maliciously, he/she is held accountable because the modified transaction is signed by him/her.

Security Analysis. We show the security result of our proposed construction, and the detailed proofs are referred to Appendix C.

THEOREM 4.1. *The proposed generic framework is indistinguishable if the* CH *scheme is indistinguishable.*

THEOREM 4.2. *The proposed generic framework is adaptively collision-resistant if the* ABET *scheme is semantically secure, the* CH *scheme is collision-resistant, and the* DPSS *scheme has secrecy.*

THEOREM 4.3. *The proposed generic framework is accountable if the* Σ *scheme is existential unforgeability under chosen message attack (EUF-CMA) secure, and the* DPSS *scheme has correctness.*

5 INSTANTIATION

5.1 The Proposed ABET Scheme

For constructing a practical ABET, we require that the underlying KP-ABE scheme should have minimal number of components in master secret key, while the size of the ciphertext is constant (i.e., independent of the number of committees). Therefore, we rely on a KP-ABE scheme [46] and a hierarchy identity-based encryption (HIBE) scheme [14]. The KP-ABE [46] can be viewed as the stepping stone to construct ABET. It has a single component in master secret

key msk, which requires a single execution of the DPSS. It works in prime-order group, and its security is based on q-type assumption in the standard model. One may use more efficient ABE schemes such as [7]. However, the master secret key msk has several components, which requires multiple executions of the DPSS. The HIBE [14] has constant-size ciphertext. The ciphertext has just three group elements, and the decryption requires only two pairing operations. In particular, HIBE has one master secret key, which can be shared with KP-ABE. The ABET has been studied in [35, 36, 43, 52]. Their schemes are based on cipher-policy ABE. They cannot be applied to open blockchains with decentralized access control (i.e., a threshold number of committee members grant decryption keys to users). Our proposed ABET scheme makes it possible based on the KP-ABE scheme.

The intertwined ABET scheme is not anonymous because its ciphertext reveals user's committee information to the public. We extend the intertwined ABET scheme into an anonymous one using asymmetric pairings, i.e., $\hat{\mathsf{e}} : \mathbb{G} \times \mathbb{H} \to \mathbb{G}_T$ (as described in [23]). The basic idea is, the index-based elements in a modifier's decryption key belong to group $\mathbb{G}$. The index-based elements in a ciphertext belong to group $\mathbb{H}$ so that the ciphertext can conceal the committee's index if the master secret key is unknown. The concrete construction of ABET is embedded in the instantiation. Below, we present Theorem 5.1 to show the proposed ABET scheme has semantic security and ciphertext anonymity. The security analysis is referred to Appendix D.

THEOREM 5.1. *The proposed ABET scheme achieves semantic security and ciphertext anonymity, if the* q'-*type assumption and eDDH assumptions hold in the asymmetric pairing groups.*

5.2 Instantiation

First, we use the proposed ABET scheme to construct our instantiation. Specifically, the Setup and KeyGen algorithms in ABET are directly used in the instantiation. The Enc and Dec algorithms in ABET are part of Hash and Adapt, respectively. Second, we rely on a recent work [38] to initiate the DPSS scheme. We particularly show an instantiation of DPSS with pessimistic case, which allows users to be held accountable in a committee using KZG commitment [32].

We denote an index space as $\{I_1, \cdots, I_k\} \in (\mathbb{Z}_q)^k$, which is associated with k committees. We define a hierarchy as follows: index i is close to the root node k, and index j is close to the leaf node. We assume each committee has n_0 users, and the threshold is t, where $t < n_0/2$ according to [38]. Let $\mathsf{H} : \{0, 1\}^* \to \mathbb{Z}_q$ be a hash function, and the size of hash output H is assumed to be l. Let $\hat{\mathsf{e}} : \mathbb{G} \times \mathbb{H} \to \mathbb{G}_T$ be a bilinear pairing. The concrete instantiation is shown below.

- Setup(1^λ): It takes a security parameter λ as input, outputs a master public key $\mathsf{mpk} = (g, u, v, w, h, \hat{\mathsf{e}}(g, h)^\alpha, \{g_1^\alpha, \cdots g_k^\alpha\}, \{h_1^\alpha, \cdots h_k^\alpha\}, g^\beta, h^{1/\alpha}, h^{\beta/\alpha}, \hat{\mathsf{e}}(g, h)^{\theta/\alpha})$, and a master secret key $\mathsf{msk} = (\alpha, \beta, \theta)$, where $(\alpha, \beta, \theta) \in \mathbb{Z}_q^*$ $\{z_1, \cdots, z_k\} \in \mathbb{Z}_q$, g is generator of group $\mathbb{G}$, h is generator of group $\mathbb{H}$, $(u, v, w) \in \mathbb{G}$, $\{g_1, \cdots, g_k\} = \{g^{z_1}, \cdots, g^{z_k}\}$, $\{h_1, \cdots, h_k\} = \{h^{z_1}, \cdots, h^{z_k}\}$. Note that the key shares of α and θ are distributed to users in committee C^0, and these key shares can be redistributed between dynamic committees (see correctness below).

- KeyGen(C_i^e, $(\mathbf{M}, \pi)$): It inputs a committee C_i^e with index $(I_1, \cdots, I_i)$, and an access structure $(\mathbf{M}, \pi)$ ($\mathbf{M}$ has n_1 rows and n_2 columns), outputs a secret key sk_{Λ_i} for a modifier. Specifically, a group of $t+1$ users in the committee C_i^e first recover secrets α and θ. Then, they pick $\{t_1, \cdots, t_{n_1}\} \in \mathbb{Z}_q$, for all $\tau \in [n_1]$, compute $sk_{(\tau,1)} = g^{s_\tau} w^{t_\tau}$, $sk_{(\tau,2)} = (u^{\pi(\tau)} v)^{-t_\tau}$, $sk_{(\tau,3)} = h^{t_\tau}$, where s_τ is a key share from α. Eventually, they pick $\{r_1, \cdots, r_{n_1}\} \in \mathbb{Z}_q$, compute $sk_0 = (g^{t^*/\alpha}, g^{r^*})$, $sk_1 = g^\theta \cdot \widehat{i}^{t^*} \cdot g^{\beta \cdot r^*}$, $sk_2 = \{g_{i-1}^{\alpha \cdot t^*}, \cdots g_1^{\alpha \cdot t^*}\}$, where $t^* = \sum_{\tau \in |n_1|}(t_\tau)$, $r^* = \sum_{\tau \in |n_1|}(r_\tau)$, and $\widehat{i} = g_k^{\alpha I_1} \cdots g_i^{\alpha I_i} \cdot g \in \mathbb{G}$ is associated with a committee C_i^e with index $(I_1, \cdots, I_i)$. The secret key is $sk_{\Lambda_i} = (\{sk_\tau\}_{\tau \in [n_1]}, sk_0, sk_1, sk_2)$.
- Hash(mpk, m, δ, j): To hash a message $m \in \mathbb{Z}_q$ under a set of attributes δ, and an index $(I_1, \cdots I_j)$, a user performs the following operations
  (1) choose a randomness $r \in \mathbb{Z}_q^*$, and a trapdoor R, compute a chameleon hash $b = g^m \cdot p'^r$ where $p' = g^e$, $e = H(R)$. Note that R denotes a short bit-string.
  (2) generate a ciphertext on message $M = R$ under a set of attributes $\delta = \{A_1, \cdots, A_{|\delta|}\}$ and index $(I_1, \cdots I_j)$. It first picks $s, r_1, r_2, \cdots, r_{|\delta|} \in \mathbb{Z}_q$, for $\tau \in |\delta|$ computes $ct_{(\tau,1)} = h^{r_\tau}$ and $ct_{(\tau,2)} = (u^{A_\tau} v)^{r_\tau} w^{-s}$. Then, it computes $ct = (R||0^{l-|R|}) \oplus H_2(\hat{e}(g,h)^{\alpha s} || \hat{e}(g,h)^{\theta s/\alpha})$, $ct_0 = (h^s, h^{s/\alpha}, h^{\beta \cdot s/\alpha})$, and $ct_1 = \widehat{j}^s$, where $\widehat{j} = h_k^{\alpha I_1} \cdots h_j^{\alpha I_j} \cdot h \in \mathbb{H}$. Eventually, it sets $C = (ct, \{ct_{(\tau,1)}, ct_{(\tau,2)}\}_{\tau \in [\delta]}, ct_0, ct_1)$.
  (3) generate a signature $epk = g^{esk}$, $\sigma = esk + sk \cdot H(epk||c)$, where (esk, epk) denotes an ephemeral key pair, and $c = g^{sk+(R||0^{l-|R|})}$ denotes a signed message.
  (4) output $(m, p', b, r, C, c, epk, \sigma)$.
- Verify(mpk, m, p', b, r, c, epk, σ): Any public user can verify whether a given hash (b, p') is valid, it outputs 1 if $b = g^m \cdot p'^r$, and $g^\sigma = epk \cdot pk^{H(epk||c)}$.
- Adapt(sk_{Λ_i}, m, m', p', b, r, C, c, epk, σ): A modifier with a secret key sk_{Λ_i}, and a new message $m' \in \mathbb{Z}_q$, performs the following operations
  (1) check $1 \stackrel{?}{=} $Verify(mpk, m, p', b, r, c, epk, σ).
  (2) run the following steps to decrypt trapdoor R:
     (a) generate a delegated key w.r.t an index $(I_1, \cdots I_{i+1})$. It picks $t' \in \mathbb{Z}_q$, computes $sk_0 = (g^{(t^*+t')/\alpha}, g^{r^*})$, $sk_1 = g^\theta \cdot \widehat{i}^{t^*} \cdot g^{\beta \cdot r^*} \cdot (g_{i-1}^{\alpha \cdot t^*})^{I_{i+1}} \cdot (g_k^{\alpha \cdot I_1} \cdots g_{i-1}^{\alpha \cdot I_{i+1}} \cdot g)^{t'}$, $sk_2 = \{g_{i-2}^{\alpha \cdot t^*} \cdot g_{i-2}^{\alpha \cdot t'}, \cdots g_1^{\alpha \cdot t^*} \cdot g_1^{\alpha \cdot t'}\}$. The delegated secret key is $sk_{\Lambda_{i+1}} = (\{sk_\tau\}_{\tau \in [n_1]}, sk_0, sk_1, sk_2)$.
     (b) if the attribute set δ involved in the ciphertext satisfies the policy MSP $(\mathbf{M}, \pi)$, then there exists constants $\{\gamma_\mu\}_{\mu \in I}$ that satisfy the equation in Section 2.2. It computes B as follows.

$$
\begin{aligned}
B &= \prod_{\mu \in I} (\hat{e}(sk_{(\mu,1)}, ct_{(0,1)}) \hat{e}(sk_{(\mu,2)}, ct_{(\mu,1)}) \\
&\quad \hat{e}(ct_{(\mu,2)}, sk_{(\mu,3)}))^{\gamma_\mu} \\
&= \hat{e}(g,h)^{s \sum_{\mu \in I} \gamma_\mu s_\mu} = \hat{e}(g,h)^{\alpha s}, \\
&\quad where \sum_{\mu \in I} \gamma_\mu s_\mu = \alpha.
\end{aligned}
$$

(c) check $(R||0^{l-|R|}) \stackrel{?}{=} ct \oplus H(B||A)$, where $A = \frac{\hat{e}(sk_1, ct_{(0,2)})}{\hat{e}(sk_{(0,1)}, ct_1)\hat{e}(sk_{(0,2)}, ct_{(0,3)})}$. The format "$||0^{l-|R|}$" is used to check when the encrypted value R is decrypted successfully with certainty $1 - 2^{l-|R|}$. If the encrypted value is decrypted, then the delegation procedure terminates (note that B is computed once). Note that $sk_{(0,1)}, sk_{(0,2)}$ denote the first, and second element of sk_0, and the same rule applies to ct_0.

(3) compute a new randomness $r' = r + (m - m')/e$, where $e = H(R)$.

(4) generate a new ciphertext C' on the same message $M = R$ using the attribute set δ and index $(I_1, \cdots, I_j)$.

(5) generate a signature $epk' = g^{esk'}$, $\sigma' = esk' + sk' \cdot H(epk'||c')$, where $c' = g^{sk'+(R||0^{l-|R|})}$.

(6) output $(m', p', b, r', C', c', epk', \sigma')$.

Correctness of DPSS scheme. Two secrets need to be distributed: (α, θ). We specifically show users in committee C^{e-1} securely *hand-off* their key shares of secret α to users in committee C^e. According to the DPSS scheme in [38], an asymmetric bivariate polynomial is used: $f(x, y) = \underline{\alpha} + a_{0,1}x + a_{1,0}y + a_{1,1}xy + a_{1,2}xy^2 + \cdots + a_{t,2t}x^t y^{2t}$. Each user in committee C^{e-1} holds a *full* key share after running Share protocol. For example, a user with pk holds a key share $f(i, y)$, which is a polynomial with dimension t. Overall, the handoff (i.e., Redistribute protocol) includes three phases: share reduction, proactivization, and full-share distribution.

- *Share Reduction.* It requires each user in committee C^{e-1} reshares its full key share. For example, user pk derives a set of *reduced* shares $\{f(i, j)\}_{j \in [1, n_0]}$ from its key share $f(i, y)$ using SSS. Then, each user distributes the reduced shares to users in committee C^e, which includes a user with pk'. As a result, each user in C^e obtains a reduced share $f(x, j)$ by interpolating the received shares $\{f(i, j)\}_{i \in [1, t]}$. Note that the dimension of $f(x, j)$ is $2t$, and $2t+1$ of these reduced key shares $\{f(x, j)\}_{j \in [1, 2t+1]}$ can recover α (see Section 2.4). The goal of this dimension-switching (from t to $2t$) is to achieve optimal communication overhead, such that only $2t+1$ users in committee C^e are required to update $f(x, j)$.
- *Proactivization.* It requires $F(x, j) = f(x, j) + f'(x, j)$, where $f'(x, y)$ is a new asymmetric bivariate polynomial with dimension $(t, 2t)$ and $f'(0, 0) = 0$. We provide more details of $f'(x, y)$ later.
- *Full-share Distribution.* It requires each user in committee C^e to recover its full key share with dimension t. For example, a full key share $F(i, y)$ is recovered by interpolating the reduced shares $\{F(i, j)\}_{j \in [1, 2t+1]}$ in committee C^e. This full key share $F(i, y)$ belongs to user pk', and $t+1$ of these full key shares can recover α.

Now we show the generation of an asymmetric bivariate polynomial $f'(x, y)$ with dimension $(t, 2t)$ such that $f'(0, 0) = 0$, which is used to update the reduced key shares $f(x, j)$ during proactivization. We denote a subset of C^e as $\mathcal{U}'$, which includes $2t+1$ users. The generation of $f'(x, y)$ requires two steps: univariate zero share, and bivariate zero share.

- *Univariate Zero Share.* It requires each user in $\mathcal{U}'$ to generate a key share $f_j'(y)$ from a common univariate polynomial with dimension $2t$. First, each user i generates a univariate polynomial

$f'_i(y) = 0 + a'_1 y + a'_2 y^2 + \cdots + a'_{2t} y^{2t}$, and broadcasts it to all users in $\mathcal{U}'$. Second, each user in $\mathcal{U}'$ generates a common univariate polynomial $f'(y) = \sum_{i \in [1,2t+1]} f'_i(y)$ by combining all received polynomials, and obtains a key share $f'_j(y)$ from $f'(y)$.

- *Bivariate Zero Share.* It requires each user in committee C^e to generate a key share $f'(x,j)$ from a common bivariate polynomial with dimension $(t, 2t)$. First, each user in $\mathcal{U}'$ generates a set of reduced shares $\{f'(i,y)\}_{i \in [1,n_0]}$ with dimension t from its key share $f'_j(y)$ (i.e., resharing process), where $f'(i,y) = 0 + a'_{1,0} y + a'_{2,0} y^2 + \cdots + a'_{2t,0} y^{2t}$. Since the reduced shares are distributed to all users in committee C^e, a common bivariate polynomial with dimension $(t, 2t)$ is established: $f'(x,y) = 0 + a'_{0,1} x + a'_{1,0} y + a'_{1,1} xy + a'_{1,2} xy^2 + \cdots + a'_{t,2t} x^t y^{2t}$. Second, each user in committee C^e obtains a reduced key share $f'(x,j)$ by interpolating the received shares $\{f'(i,j)\}_{j \in [1,2t+1]}$. The key share $f'(x,j) = 0 + a'_{0,1} x + a'_{0,2} x^2 + \cdots + a'_{0,t} x^t$ is used to update $f(x,j)$ in the proactivization.

The asymmetric bivariate polynomial $f'(x,y)$ can be reused in another proactivization when sharing secret θ. In other words, multiple handoff protocols with respect to different secrets can be updated using the same bivariate polynomial, with the condition that these handoff protocols are executed within the same committee.

Correctness of Judge **algorithm.** We show the judge process in detail. First, any public user verifies the connection between a transaction and its modified version, and this connection is publicly verifiable. For example, given two chameleon hash outputs: $(m, m', b, p', C, C', c, c', epk, \sigma, epk', \sigma')$, the public performs the following operations

- verify chameleon hash $b = g^m \cdot p'^r = g^m \cdot p'^{r'}$.
- verify message-signature pair (c, σ) under (epk, pk), and message-signature pair (c', σ') under (epk', pk').
- verify $pk' = pk \cdot \Delta(sk)$, where $\Delta(sk) = c'/c = g^{sk'-sk}$ (the meaning of $\Delta(sk)$ is referred to Section 2.3). Note that (c, c') are derived from the same chameleon trapdoor R.

Second, any public user obtains a set of accused committees from interacting with an access blackbox O. We note that the modifier's delegated secret keys are disallowed to be used in generating O. If some modifiers use their delegated secret keys to generate O, the public cannot identify the accused committees correctly because the delegated secret keys may share the same committee index. We argue that it is challenging to prevent such malicious modifiers from using their delegated secret keys to generate O, as some ABET schemes [35, 36] (including our proposed one) support a delegation process. The delegation allows a user to obtain some delegated decryption keys by re-randomizing the given decryption key.

Eventually, the public outputs a transaction-committee pair (T', C^e). In particular, we rely on the KZG commitment and PoW consensus to hold a modifier pk' accountable in an accused committee C^e. Now we provide more details.

- *Share Reduction.* We require user pk' in committee C^e to generate a commitment $C_{f(x,j)}$, which is a KZG commitment to the reduced key shares $\{f(i,j)\}_{j \in [1,2t+1]}$, and a set of witnesses $\{w_{f(i,j)}\}_{j \in [1,2t+1]}$. A witness $w_{f(i,j)}$ means the witness to evaluation of $f(x,j)$ at i. Note that $i \in [1, 2t+1]$ indicates the order

of user pk''s public key in committee C^e (we order nodes lexicographically by users' public keys and choose the first $2t + 1$).
- *Full-share Distribution.* We require user pk' in committee C^e to generate a commitment $C_{F(x,j)}$, which is a KZG commitment to the reduced key shares $\{F(i,j)\}_{j \in [1,2t+1]}$, and a set of witnesses $w_{F(i,j)}$. A witness $w_{F(i,j)}$ means the witness to evaluation of $F(x,j)$ at i.
- *PoW Consensus.* We require user pk' to hash the KZG commitment and the set of witnesses, store them to an immutable transaction, and put them on-chain for PoW consensus.

Overall, the commitment and witness can ensure the correctness of handoff. Specifically, new committee members can verify the correctness of reduced shares from old committee members, thus the correctness of dimension-switching. Meanwhile, the proof of correctness is publicly verifiable, such that any public user can verify that $f(i,j)$ (or $F(i,j)$) is the correct evaluation at i (i.e., user pk') of the polynomial committed by $C_{f(x,j)}$ (or $C_{F(x,j)}$) in committee C^e.

6 IMPLEMENTATION AND EVALUATION

In this section, we evaluate the performance of the proposed solution based on a proof-of-concept implementation in Python and Flask framework [3]. We create a mutable open blockchain system with basic functionalities and a PoW consensus mechanism. The simulated open blockchain system is "healthy", satisfying the properties of persistence and liveness [26]. The system is specifically designed to include ten blocks, each block includes 100 transactions. Please note, that our implementation can easily extend it to real-world applications such as a block containing 3500 transactions. We simulate ten nodes in a peer-to-peer network, each of them is implemented as a lightweight blockchain node. They can also be regarded as the users in a committee. A chain of blocks is established with PoW mechanism by consolidating transactions broadcast by the ten nodes. The implementation code is available on GitHub [5].

First, if users append mutable transactions to blockchain, they use the proposed solution to hash the registered message m. Later, a miner uses the conventional hash function SHA-256 H to hash the chameleon hash output h and validates $H(h)$ using a Merkle tree. Note that the non-hashed components such as randomness r, are parts of a mutable transaction $T = (pk_{CH}, m, h, r, C, c, \sigma)$. As a consequence, a modifier can replace T by $T' = (pk_{CH}, m', h, r', C', c', \sigma')$ without changing the hash output $H(h)$.

Second, we mimic a dynamic committee that includes five users, we split the master secret key into five key shares so that each user in a committee holds a key share. We simulate the basic functionality of DPSS, including resharing and updating key shares. Any user can join in or leave from a committee by transmitting those key shares between committee members. In particular, we simulate three users in a committee can collaboratively recover the master secret key and grant access privileges to the modifiers.

Now, we conclude that: 1) The proposed solution incurs no overhead to chain validation. This is because, rewrite the message in T has no effect on the PoW mechanism, as the chameleon hash output h is used for computing the transaction hash for Merkle tree leaves. 2) The proof-of-concept implementation indicates that the proposed solution can act as an additional layer on top of any

open blockchains to perform accountable rewriting. Specifically, we append mutable transactions using the proposed solution to the blockchain, and we allow dynamic committees to grant access privileges for rewriting those mutable transactions. 3) The proposed solution is compatible with existing blockchain systems for the following reasons. The only change which CH-based approach requires is to replace the standard the hash function H by CH for generating a chameleon hash value h before validating the transactions in each block [10]. Besides, DPSS is designed for open blockchains, and decentralized systems [38].

6.1 Evaluation

We implement our proposed solution using the Charm framework [8] and evaluate its performance on a PC with Intel Core i5 (2.7GHz×2) and 8GB RAM. We use Multiple Precision Arithmetic Library, Pairing-Based Cryptography (PBC) Library, and we choose MNT224 curve for pairing, which is the best Type-III paring in PBC. We instantiate the hash function and the pseudo-random generator with the corresponding standard interfaces provided by the Charm framework.

First, the Setup algorithm takes about 0.52 seconds (s). The running time of KeyGen, Hash, and Adapt algorithms are measured and shown in Figure 5 (a-c). The performance of these algorithms is linear to the number of attributes or the size of policies. Specifically, the run-time of KeyGen takes only 2.37s, even if the size of the policy is 100. We discover that the dominating operation is parsing the access policy $(\mathbf{M}, \pi)$, and we argue that a better designed parsing interface can reduce the overall cost of KeyGen. Moreover, it only takes 2.44s and 3.87s respectively for Hash and Adapt to handle 100 attributes. The run-time cost of such algorithms mainly comes from processing the attributes list and access policy $(\mathbf{M}, \pi)$, i.e., the coefficient calculation of every attribute and the cost of determining whether a given attribute set satisfies the access policy.

Second, we evaluate the running time of a t-out-of-n_0 DPSS protocol, where n_0 indicates the number of users in a committee and t is the threshold. Let $t < n_0/2$ be a safe threshold. The overhead includes the distribution cost between committee members, and the polynomial calculation cost. Figure 5 (d) shows that the overall running time is linear to the square number of users n^2 in a committee, since more shares need to be distributed and more polynomials need to calculated among n_0 users. Our implementation can scale up to larger committees. For example, the running times for $t = 20$ and $t = 30$ are about 3.09s and 7.34s, respectively.

To conclude, the implementation performs the resharing twice and updating once regarding two shared secrets. Besides, the number of updating process is constant in a committee, independent of the number of shared secrets used in ABET. Since only two secrets are needed to be shared and recovered, we argue that the proposed ABET scheme is the most practical one. It is suitable for decentralized systems when applying DPSS to the proposed ABET scheme. On the security-front, because every committee has at most $\frac{n_0}{3}$ malicious members [37] and $\frac{n_0}{2}+1$ committee members recover the shared secrets [38], the malicious committee members cannot dictate the committee and control the rewriting privileges. For the storage cost, we mention that the number of mutable transactions ranges from 2% to 10% inside a block [21]. Each mutable transaction needs to store $T =$

$(\mathsf{pk}_{\mathsf{CH}}, m, h, \mathsf{r}, C, c, \sigma)$. The storage cost of a mutable transaction includes: 1) $2\mathcal{L}_{\mathbb{Z}_q} + 3\mathcal{L}_{\mathbb{G}}$ regarding DL-based chameleon hash; 2) $\mathcal{L}_{\mathbb{Z}_q} + |\delta| \times \mathcal{L}_{\mathbb{G}} + (|\delta| + 4) \times \mathcal{L}_{\mathbb{H}}$ regarding ABET; 3) $\mathcal{L}_{\mathbb{Z}_q} + 2\mathcal{L}_{\mathbb{G}}$ regarding digital signature. The committee's on-chain storage cost regarding DPSS [38] is $2(t + 1) \times [\mathcal{L}_{\mathbb{G}} + (2t + 1)(\mathcal{L}_{\mathbb{Z}_q} + \mathcal{L}_{\mathbb{H}})]$.

7 RELATED WORK

Blockchain Rewriting. Blockchain rewriting was first introduced by Ateniese et al. [10]. They propose to replace the regular SHA256 hash function with a chameleon hash (CH) in blockchain generation [34]. The hashing of CH is parametrized by a public key pk, and CH behaves like a collision-resistant hash function if the chameleon secret key sk (or trapdoor) is unknown. A trapdoor holder (or modifier) can find collisions and output a new message-randomness pair without changing the hash value.

Camenisch et al. [17] introduced a new cryptographic primitive: chameleon hash with ephemeral trapdoor (CHET). CHET requires that a modifier must have two trapdoors to find collisions: one trapdoor sk is associated with the public key pk; the other one is an ephemeral trapdoor etd chosen by the party who initially computed the hash value. CHET provides more control in rewriting in the sense that the party, who computed the hash value, can decide whether the holder of sk shall be able to rewrite the hash by providing or withholding the ephemeral trapdoor etd.

Derler et al. [19] proposed policy-based chameleon hash (PCH) to achieve fine-grained rewriting in the blockchain. The proposed PCH replaces the public key encryption scheme in CHET by a ciphertext-policy ABE scheme, such that a modifier must satisfy a policy to find collisions given a hash value. Recently, Tian et al. proposed an accountable PCH for blockchain rewriting (PCHBA) [52]. The proposed PCHBA enables the modifiers of transactions to be held accountable for the modified transactions. In particular, PCHBA allows a third party (e.g., key generation center) to resolve any dispute over modified transactions.

In another work, Puddu et al. [45] proposed μchain: a mutable blockchain. A transaction owner introduces a set of transactions, including an active transaction and multiple inactive transactions, where the inactive transactions are possible versions of the transaction data (namely, mutations) encrypted by the transaction owner, and the decryption keys are distributed among miners using Shamir's SSS [49]. The transaction owner enforces access control policies to define who is allowed to trigger mutations in which context. Upon receiving a mutation-trigger request, a set of miners runs a Multi Party Computation (MPC) protocol to recover the decryption key, decrypt the appropriate version of the transaction and publish it as an active transaction. μchain incurs considerable overhead due to the use of MPC protocols across multiple miners. It works at both permissioned and permissionless blockchains.

Deuber et al. [21] introduced an efficient redactable blockchain in the permissionless setting. The proposed protocol relies on a consensus-based e-voting system [33], such that the modification is executed in the chain if a modification request from any public user gathers enough votes from miners (we call it V-CH for convenience). In a follow-up work, Thyagarajan et al. [51] introduced a protocol called Reparo to repair blockchains, which acts as a publicly verifiable layer on top of any permissionless blockchain. The
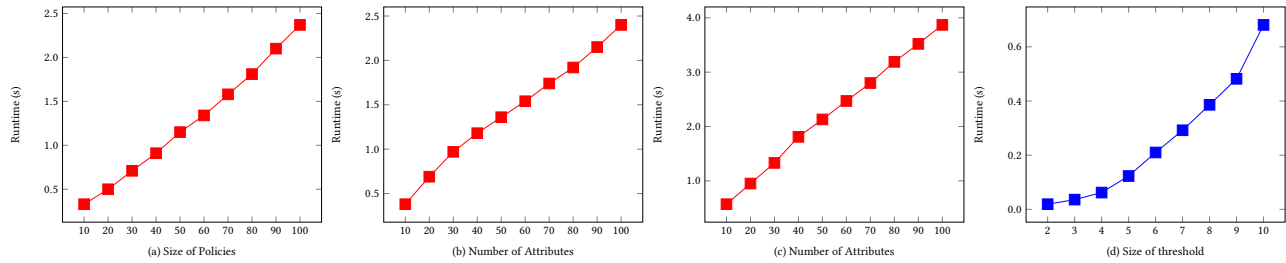
Figure 5: Run-time of KeyGen, Hash, Adapt **algorithms, and** DPSS **scheme.**

Table 1: The comparison between various blockchain rewriting solutions. CH-based indicates CH-based blockchain rewriting. Fine-grained means that each mutable transaction is associated with an access policy such that the transaction can be modified by anyone whose rewriting privilege satisfy the policy.

	CH [10]	μchain [45]	PCH [19]	V-CH [21]	PCHBA [52]	Ours
CH-based	✓	✗	✓	✗	✓	✓
Permissionless	✓	✓	✗	✓	✗	✓
Fine-grained	✗	✗	✓	✗	✓	✓
Accountability	✓	✓	✗	✓	✓	✓

unique feature of Reparo is that it is immediately integrable into open blockchains in a backward compatible fashion (i.e., any existing blockchains already containing illicit contents can be redacted using Reparo).

There are mainly two types of blockchain rewritings in the literature: CH-based [10, 17–19, 52], and non CH-based [21, 45, 51]. CH-based blockchain rewritings allow one or more trusted modifiers to rewrite blockchain. The non-CH-based solution requires a threshold number of parties (or miners) to rewrite the blockchain. We stress that both of them aim to rewrite blockchains securely and efficiently. One can apply both of them to redactable blockchains.

Table 1 shows a comparison between blockchain rewriting related solutions. In this work, we use chameleon hash cryptographic primitive to secure the blockchain rewriting. Our proposed solution supports a fine-grained and controlled rewriting for open blockchains. It holds both the modifiers' public keys and the rewriting privileges accountable for the modified transactions. Overall, this work can be viewed as a step forward from PCH and PCHBA.

Dynamic Proactive Secret Sharing. Proactive security was first introduced by Ostrovsky and Yung [44], which is refreshing secrets to withstand compromise. Later, Herzberg et al. [30] introduced proactive secret sharing (PSS). The PSS allows the distributed key shares in a SSS to be updated periodically, so that the secret remains secure even if an attacker compromises a threshold number of shareholders in each epoch. However, it did not support dynamic committees because users may join in or leave from a committee dynamically. Desmedt and Jajodia [20] introduced a scheme that redistributes secret shares to new access structure (or new committee). Specifically, a resharing technique is used to change the committee and threshold in PSS. However, the scheme is not verifiable, which disallows PSS to identify the faulty (or malicious)

users. The property of verifiability is essential to PSS (i.e., verifiable secret sharing such as Feldman [25]), which holds malicious users accountable. So, the dynamic proactive secret sharing (DPSS) we considered in this work includes verifiability.

There exist several DPSS schemes in the literature. Wong et al. [54] introduced a verifiable secret redistribution protocol that supports dynamic committee. The proposed protocol allows new shareholders to verify the validity of their shares after redistribution between different committees. Zhou et al. [57] introduced an APSS, a PSS protocol for asynchronous systems that tolerate denial-of-service attacks. Schultz et al. [48] introduced a resharing protocol called MPSS. The MPSS supports mobility, which means the group of shareholders can change during resharing. Baron et al. [11] introduced a DPSS protocol that achieves a constant amortized communication overhead per secret share. In CCS'19, Maram et al. [38] presented a practical DPSS: CHURP. CHURP is designed for open blockchains, and it has very low communication overhead per epoch compared to the existing schemes [11, 48, 54, 57]. Specifically, the total number of bits transmitted between all committee members in an epoch is substantially lower than in existing schemes. Recently, Benhamouda et al. [13] introduced anonymous secret redistribution. The benefit is to ensure sharing and resharing of secrets among small dynamic committees.

DPSS can be used to secure blockchain rewriting, such as μchain [45]. μchain relies on encryption with secret sharing (ESS) to hide illegal content, as certain use-cases aim to prevent distribution of illegal content (e.g., child pornography) via the blockchain. ESS allows all the mutable transactions containing illegal content to be encrypted using transaction-specific keys. The transaction-specific keys are split into shares using DPSS [11], and these resulting shares are distributed to a number of miners, which then reshare the keys among all online miners dynamically. In this work, we use KP-ABE with DPSS to ensure blockchain rewiring with fine-grained access control. The master secret key in KP-ABE is split into key shares, and these key shares are distributed to all users in a committee. The key shares can be securely redistributed across dynamic committees. To the best of our knowledge, ours is the first attempt to distribute the master secret key in KP-ABE for decentralized systems.

8 CONCLUSION

In this paper, we proposed a new framework of accountable fine-grained blockchain rewriting. The proposed framework is designed for open blockchains that require no trust assumptions. Besides,

the proposed framework achieves public accountability, which can thwart the malicious rewriting of blockchain. Specifically, public accountability allows the modifiers' public keys and the rewriting privileges to be held accountable for the modified transactions. We presented a practical instantiation, and showed that the proposed solution is suitable for open blockchain applications. In particular, the proof-of-concept implementation demonstrated that our proposed solution can be easily integrated into the existing open blockchains.

9 ACKNOWLEDGMENTS

REFERENCES

[1] [n.d.]. Bitcoin. https://bitcoin.org/en/.
[2] [n.d.]. Ethereum. https://ethereum.org/en/.
[3] [n.d.]. Flask Framework. https://flask.palletsprojects.com/en/1.1.x/.
[4] [n.d.]. General Data Protection Regulation. https://eugdpr.org.
[5] [n.d.]. Our Source Code. https://github.com/lbwtorino/Fine-Grained-Blockchain-Rewriting-in-Permissionless-Setting.
[6] [n.d.]. Proof of Stake. https://en.wikipedia.org/wiki/Proof_of_stake.
[7] Shashank Agrawal and Melissa Chase. 2017. FAME: fast attribute-based message encryption. In CCS. 665–682.
[8] Joseph A Akinyele, Christina Garman, Ian Miers, Matthew W Pagano, Michael Rushanan, Matthew Green, and Aviel D Rubin. 2013. Charm: a framework for rapidly prototyping cryptosystems. Journal of Cryptographic Engineering 3, 2 (2013), 111–128.
[9] Marcin Andrychowicz and Stefan Dziembowski. 2015. Pow-based distributed cryptography with no trusted setup. In CRYPTO. 379–399.
[10] Giuseppe Ateniese, Bernardo Magri, Daniele Venturi, and Ewerton Andrade. 2017. Redactable blockchain–or–rewriting history in bitcoin and friends. In EuroS&P. 111–126.
[11] Joshua Baron, Karim El Defrawy, Joshua Lampkins, and Rafail Ostrovsky. 2015. Communication-optimal proactive secret sharing for dynamic groups. In ACNS. 23–41.
[12] Amos Beimel. 1996. Secure schemes for secret sharing and key distribution. Ph.D. Dissertation. PhD thesis, Israel Institute of Technology, Technion, Haifa, Israel.
[13] Fabrice Benhamouda, Craig Gentry, Sergey Gorbunov, Shai Halevi, Hugo Krawczyk, Chengyu Lin, Tal Rabin, and Leonid Reyzin. 2020. Can a Public Blockchain Keep a Secret?. In TCC. 260–290.
[14] Dan Boneh, Xavier Boyen, and Eu-Jin Goh. 2005. Hierarchical identity based encryption with constant size ciphertext. In CRYPTO. 440–456.
[15] Dan Boneh, Amit Sahai, and Brent Waters. 2006. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In EUROCRYPT. 573–592.
[16] Dan Boneh and Brent Waters. 2006. A fully collusion resistant broadcast, trace, and revoke system. In CCS. 211–220.
[17] Jan Camenisch, David Derler, Stephan Krenn, Henrich C Pöhls, Kai Samelin, and Daniel Slamanig. 2017. Chameleon-hashes with ephemeral trapdoors. In PKC. 152–182.
[18] David Derler, Kai Samelin, and Daniel Slamanig. 2020. Bringing Order to Chaos: The Case of Collision-Resistant Chameleon-Hashes. In PKC. 462–492.
[19] David Derler, Kai Samelin, Daniel Slamanig, and Christoph Striecks. 2019. Fine-Grained and Controlled Rewriting in Blockchains: Chameleon-Hashing Gone Attribute-Based. In NDSS.
[20] Yvo Desmedt and Sushil Jajodia. 1997. Redistributing secret shares to new access structures and its applications. Technical Report.
[21] Dominic Deuber, Bernardo Magri, and Sri Aravinda Krishnan Thyagarajan. 2019. Redactable blockchain in the permissionless setting. In IEEE S&P. 124–138.
[22] John R Douceur. 2002. The sybil attack. In International workshop on peer-to-peer systems. 251–260.
[23] Léo Ducas. 2010. Anonymity from asymmetry: New constructions for anonymous HIBE. In CT-RSA. 148–164.
[24] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. 2015. Proofs of space. In CRYPTO. 585–605.
[25] Paul Feldman. 1987. A practical scheme for non-interactive verifiable secret sharing. In FOCS. 427–438.
[26] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The bitcoin backbone protocol: Analysis and applications. In EUROCRYPT. 281–310.

[27] Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. 1999. Secure distributed key generation for discrete-log based cryptosystems. In EUROCRYPT. 295–310.
[28] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017. Algorand: Scaling byzantine agreements for cryptocurrencies. In Proceedings of the 26th Symposium on Operating Systems Principles. 51–68.
[29] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. 2006. Attribute-based encryption for fine-grained access control of encrypted data. In CCS. 89–98.
[30] Amir Herzberg, Stanisław Jarecki, Hugo Krawczyk, and Moti Yung. 1995. Proactive secret sharing or: How to cope with perpetual leakage. In CRYPTO. 339–352.
[31] Markus Jakobsson and Ari Juels. 1999. Proofs of work and bread pudding protocols. In Secure information networks. 258–272.
[32] Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. 2010. Constant-size commitments to polynomials and their applications. In ASIACRYPT. 177–194.
[33] Tadayoshi Kohno, Adam Stubblefield, Aviel D Rubin, and Dan S Wallach. 2004. Analysis of an electronic voting system. In IEEE S&P. 27–40.
[34] Hugo Krawczyk and Tal Rabin. 2000. Chameleon Signatures. In NDSS.
[35] Junzuo Lai and Qiang Tang. 2018. Making any attribute-based encryption accountable, efficiently. In ESORICS. 527–547.
[36] Zhen Liu, Zhenfu Cao, and Duncan S Wong. 2013. Blackbox traceable CP-ABE: how to catch people leaking their keys by selling decryption devices on ebay. In CCS. 475–486.
[37] Loi Luu, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. 2016. A secure sharding protocol for open blockchains. In CCS. 17–30.
[38] Sai Krishna Deepak Maram, Fan Zhang, Lun Wang, Andrew Low, Yupeng Zhang, Ari Juels, and Dawn Song. 2019. Churp: Dynamic-committee proactive secret sharing. In CCS. 2369–2386.
[39] Roman Matzutt, Jens Hiller, Martin Henze, Jan Henrik Ziegeldorf, Dirk Müllmann, Oliver Hohlfeld, and Klaus Wehrle. 2018. A quantitative analysis of the impact of arbitrary blockchain content on bitcoin. In FC. 420–438.
[40] Roman Matzutt, Oliver Hohlfeld, Martin Henze, Robin Rawiel, Jan Henrik Ziegeldorf, and Klaus Wehrle. 2016. Poster: I don't want that content! on the risks of exploiting bitcoin's blockchain as a content store. In CCS. 1769–1771.
[41] Ralph C Merkle. 1989. A certified digital signature. In CRYPTO. 218–238.
[42] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
[43] Jianting Ning, Zhenfu Cao, Xiaolei Dong, Junqing Gong, and Jie Chen. 2016. Traceable CP-ABE with short ciphertexts: How to catch people selling decryption devices on ebay efficiently. In ESORICS. 551–569.
[44] Rafail Ostrovsky and Moti Yung. 1991. How to withstand mobile virus attacks. In ACM PODC. 51–59.
[45] Ivan Puddu, Alexandra Dmitrienko, and Srdjan Capkun. 2017. μchain: How to Forget without Hard Forks. IACR Cryptology ePrint Archive 2017 (2017), 106.
[46] Yannis Rouselakis and Brent Waters. 2013. Practical constructions and new proof methods for large universe attribute-based encryption. In CCS. 463–474.
[47] Claus-Peter Schnorr. 1991. Efficient signature generation by smart cards. Journal of cryptology 4, 3 (1991), 161–174.
[48] David A Schultz, Barbara Liskov, and Moses Liskov. 2008. Mobile proactive secret sharing. In ACM PODC. 458–458.
[49] Adi Shamir. 1979. How to share a secret. Commun. ACM 22, 11 (1979), 612–613.
[50] Victor Shoup. 1997. Lower bounds for discrete logarithms and related problems. In EUROCRYPT. 256–266.
[51] Sri Aravinda Krishnan Thyagarajan, Adithya Bhat, Bernardo Magri, Daniel Tschudi, and Aniket Kate. 2020. Reparo: Publicly Verifiable Layer to Repair Blockchains. arXiv preprint arXiv:2001.00486 (2020).
[52] Yangguang Tian, Nan Li, Yingjiu Li, Pawel Szalachowski, and Jianying Zhou. 2020. Policy-based Chameleon Hash for Blockchain Rewriting with Black-box Accountability. In ACSAC. 813–828.
[53] Giannis Tziakouris. 2018. Cryptocurrencies—a forensic challenge or opportunity for law enforcement? an interpol perspective. IEEE S&P 16, 4 (2018), 92–94.
[54] Theodore M Wong, Chenxi Wang, and Jeannette M Wing. 2002. Verifiable secret redistribution for archive systems. In First International IEEE Security in Storage Workshop, 2002. Proceedings. 94–105.
[55] Haifeng Yu, Ivica Nikolić, Ruomu Hou, and Prateek Saxena. 2020. Ohie: Blockchain scaling made simple. In IEEE (S&P). 90–105.
[56] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. 2018. Rapidchain: Scaling blockchain via full sharding. In CCS. 931–948.
[57] Lidong Zhou, Fred B Schneider, and Robbert Van Renesse. 2005. APSS: Proactive secret sharing in asynchronous systems. ACM (TISSEC) 8, 3 (2005), 259–286.

A SECURITY ANALYSIS OF NEW ASSUMPTION

THEOREM A.1. Let $(\epsilon_1, \epsilon_2, \epsilon_T) : \mathbb{Z}_q \rightarrow \{0, 1\}^*$ be three random encodings (injective functions) where $\mathbb{Z}_q$ is a prime field. ϵ_1 maps all $a \in \mathbb{Z}_q$ to the string representation $\epsilon_1(g^a)$ of $g^a \in \mathbb{G}$. Similarly, ϵ_2

for $\mathbb{H}$ and ϵ_T for $\mathbb{G}_T$. If $(a, b, c, d, \{z_i\}_{i \in [1,q']}) \xleftarrow{R} \mathbb{Z}_q$ and encodings $\epsilon_1, \epsilon_2, \epsilon_T$ are randomly chosen, then we define the advantage of the adversary in solving the q'-type with at most Q queries to the group operation oracles O_1, O_2, O_T and the bilinear pairing $\hat{e}$ as

$$
\begin{aligned}
\mathsf{Adv}_{\mathcal{A}}^{q'\text{-}type}(\lambda) =\ & |\Pr[\mathcal{A}(q, \epsilon_1(1), \epsilon_1(a), \epsilon_1(c), \epsilon_1((ac)^2), \\
& \epsilon_1(abd), \epsilon_1(d/ab), \epsilon_1(z_i), \epsilon_1(acz_i), \\
& \epsilon_1(ac/z_i), \epsilon_1(a^2cz_i), \epsilon_1(b/z_i^2), \\
& \epsilon_1(b^2/z_i^2), \epsilon_1(acz_i/z_j), \epsilon_1(bz_i/z_j^2), \\
& \epsilon_1(abcz_i/z_j), \epsilon_1((ac)^2 z_i/z_j), \\
& \epsilon_2(1), \epsilon_2(b), \epsilon_2(abd), \epsilon_2(abcd), \\
& \epsilon_2(d/ab), \epsilon_2(c), \epsilon_2(cd/ab), \\
=\ & b : (a, b, c, d, \{z_i\}_{i \in [1,q']}, s \xleftarrow{R} \mathbb{Z}_q, b \in (0, 1), \\
& t_b = abc, t_{1-b} = s)] \\
& -1/2| \le \frac{16(Q + q' + 22)^2}{q}
\end{aligned}
$$

PROOF. Let $\mathcal{S}$ play the following game for $\mathcal{A}$. $\mathcal{S}$ maintains three polynomial sized dynamic lists: $L_1 = \{(p_i, \epsilon_{1,i})\}, L_2 = \{(q_i, \epsilon_{2,i})\}, L_T = \{(t_i, \epsilon_{T,i})\}$, the $p_i \in \mathbb{Z}_q[A, B, C, D, Z_i, Z_j, T_0, T_1]$ are 8-variate polynomials over $\mathbb{Z}_q$ (note that $i \ne j$), such that $p_0 = 1, p_1 = A, p_2 = C, p_3 = (AC)^2, p_4 = ABD, p_5 = D/AB, p_6 = Z_i, p_7 = ACZ_i, p_8 = AC/Z_i, p_9 = A^2CZ_i, p_{10} = B/Z_i^2, p_{11} = B^2/Z_i^2, p_{12} = ACZ_i/Z_j, p_{13} = BZ_i/Z_j^2, p_{14} = ABCZ_i/Z_j, p_{15} = (AC)^2 Z_i/Z_j, q_0 = 1, q_1 = B, q_2 = ABD, q_3 = ABCD, q_4 = D/AB, q_5 = C, q_6 = CD/AB, p_{16} = T_0, p_{17} = T_1, t_0 = 1$, and $(\{\epsilon_{1,i}\}_{i=0}^{16} \in \{0,1\}^*, \{\epsilon_{2,i}\}_{i=0}^5 \in \{0,1\}^*, \{\epsilon_{T,0}\} \in \{0,1\}^*)$ are arbitrary distinct strings. Therefore, the three lists are initialized as $L_1 = \{(p_i, \epsilon_{1,i})\}_{i=0}^{17}, L_2 = \{(q_i, \epsilon_{2,i})\}_{i=0}^6, L_T = (t_0, \epsilon_{T,0})$.

At the beginning of the game, $\mathcal{S}$ sends the encoding strings $(\{\epsilon_{1,i}\}_{i=0,\cdots,17}, \{\epsilon_{2,i}\}_{i=0,\cdots,6}, \epsilon_{T,0})$ to $\mathcal{A}$, which includes $q'+26$ strings. Note that the number of encoding string $\epsilon_{1,i}$ is linear to the parameter q'. After this, $\mathcal{S}$ simulates the group operation oracles O_1, O_2, O_T and the bilinear pairing $\hat{e}$. We assume that all requested operands are obtained from $\mathcal{S}$.

- O_1: The group operation involves two operands $\epsilon_{1,i}, \epsilon_{1,j}$. Based on these operands, $\mathcal{S}$ searches the list L_1 for the corresponding polynomials p_i and p_j. Then $\mathcal{S}$ performs the polynomial addition or subtraction $p_l = p_i \pm p_j$ depending on whether multiplication or division is requested. If p_l is in the list L_1, then $\mathcal{S}$ returns the corresponding ϵ_l to $\mathcal{A}$. Otherwise, $\mathcal{S}$ uniformly chooses $\epsilon_{1,l} \in \{0,1\}^*$, where $\epsilon_{1,l}$ is unique in the encoding string L_1, and appends the pair $(p_l, \epsilon_{1,l})$ into the list L_1. Finally, $\mathcal{S}$ returns $\epsilon_{1,l}$ to $\mathcal{A}$ as the answer. Group operation queries in O_2, O_T are treated similarly.

- $\hat{e}$: The group operation involves two operands $\epsilon_{T,i}, \epsilon_{T,j}$. Based on these operands, $\mathcal{S}$ searches the list L_T for the corresponding polynomials t_i and t_j. Then $\mathcal{S}$ performs the polynomial multiplication $t_l = t_i \cdot t_j$. If t_l is in the list L_T, then $\mathcal{S}$ returns the corresponding $\epsilon_{T,l}$ to $\mathcal{A}$. Otherwise, $\mathcal{S}$ uniformly chooses $\epsilon_{T,l} \in \{0,1\}^*$, where $\epsilon_{T,l}$ is unique in the encoding string L_T, and appends the pair $(t_l, \epsilon_{T,l})$ into the list L_T. Finally, $\mathcal{S}$ returns $\epsilon_{T,l}$ to $\mathcal{A}$ as the answer.

After querying at most Q times of corresponding oracles, $\mathcal{A}$ terminates and outputs a guess $b' = \{0, 1\}$. At this point, $\mathcal{S}$ chooses

random $a, b, c, d, z_i, z_j, s \in \mathbb{Z}_q$ and $t_b = abc$ and $t_{1-b} = s$. $\mathcal{S}$ sets $A = a, B = b, C = c, D = d, Z_i = z_i, Z_j = z_j, T_0 = t_b, T_1 = t_{1-b}$. The simulation by $\mathcal{S}$ is perfect (and reveal nothing to $\mathcal{A}$ about b) unless the abort event happens. Thus, we bound the probability of event abort by analyzing the following cases:

(1) $p_i(a, b, c, d, z_i, z_j, t_0, t_1) = p_j(a, b, c, d, z_i, z_j, t_0, t_1)$: The polynomial $p_i \ne p_j$ due to the construction method of L_1, and $(p_i - p_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ is a non-zero polynomial of degree $[0, 6]$, or q-2 (q-2 is produced by Z_j^{q-2}). Since $Z_j \cdot Z_j^{q-2} = Z_j^{q-1} \equiv 1 (\mod q)$, we have $(AC)^2 Z_i Z_j \cdot Z_j^{q-2} \equiv (AC)^2 Z_i Z_j (\mod q)$. By using Lemma 1 in [50], we have $\Pr[(p_i - p_j)(a, b, c, d, z_i, z_j, t_0, t_1) = 0] \le \frac{6}{q}$ because the maximum degree of $(AC)^2 Z_i/Z_j(p_i - p_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ is 6. So, we have $\Pr[p_i(a, b, c, d, z_i, z_j, t_0, t_1) = p_j(a, b, c, d, z_i, z_j, t_0, t_1)] \le \frac{6}{q}$, and the abort probability is $\Pr[\mathsf{abort}_1] \le \frac{6}{q}$.

(2) $q_i(a, b, c, d, z_i, z_j, t_0, t_1) = q_j(a, b, c, d, z_i, z_j, t_0, t_1)$: The polynomial $q_i \ne q_j$ due to the construction method of L_2, and $(q_i - q_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ is a non-zero polynomial of degree $[0, 4]$, or q-2 (q-2 is produced by $(AB)^{q-2}$). Since $AB \cdot (AB)^{q-2} = (AB)^{q-1} \equiv 1 (\mod q)$, we have $CDAB \cdot (AB)^{q-2} \equiv CDAB(\mod q)$. The maximum degree of $CD/AB(q_i - q_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ is 4, so the abort probability is $\Pr[\mathsf{abort}_2] \le \frac{4}{q}$.

(3) $t_i(a, b, c, d, z_i, z_j, t_0, t_1) = t_j(a, b, c, d, z_i, z_j, t_0, t_1)$: The polynomial $p_i \ne p_j$ due to the construction method of L_1, and $(p_i - p_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ is a non-zero polynomial of degree $[0, 6]$, or q-2. Since $(AC)^2 Z_i \cdot Z_j^{q-2}(t_i - t_j)(a, b, c, d, z_i, z_j, t_0, t_1)$ has degree 6, we have $\Pr[(p_i - p_j)(a, b, c, d, z_i, z_j, t_0, t_1) = 0] \le \frac{6}{q}$. The abort probability is $\Pr[\mathsf{abort}_3] \le \frac{6}{q}$.

By summing over all valid pairs (i, j) in each case (i.e., at most $\binom{Q_{\epsilon_1}+18}{2} + \binom{Q_{\epsilon_2}+7}{2} + \binom{Q_{\epsilon_T}+1}{2}$ pairs), and $Q_{\epsilon_1} + Q_{\epsilon_2} + Q_{\epsilon_T} = Q + q' + 26$, we have the abort probability is

$$
\begin{aligned}
\Pr[\mathsf{abort}] &= \Pr[\mathsf{abort}_1] + \Pr[\mathsf{abort}_2] + \Pr[\mathsf{abort}_3] \\
&\le [\binom{Q_{\epsilon_1} + 18}{2} + \binom{Q_{\epsilon_2} + 7}{2} + \binom{Q_{\epsilon_T} + 1}{2}] \\
&\quad \cdot (\frac{4}{q} + 2\frac{6}{q}) \le \frac{16(Q + q' + 26)^2}{q}.
\end{aligned}
$$

□

B P2P COMMUNICATION TECHNIQUE [38]

P2P channels can be implemented in different ways depending on the deployment environment. In a permissionless setting, establishing a direct off-chain connection between users is undesirable, as it compromises users' anonymity. Revealing network-layer identities is also dangerous, as it may lead to targeted attacks. Here, we list two approaches. The first approach is to use anonymizing overlay networks such as Tor, at the cost of considerable additional setup and engineering complexity. The second approach uses transaction ghosting, a technique for P2P messaging on a blockchain, which is an overlay on existing blockchain infrastructure. The key idea is to overwrite transactions so that they are broadcast but subsequently dropped by the network. Most of these transactions are broadcast for free.

Now, we use the Ethereum P2P network as an example; a similar technique can apply to other blockchains such as Bitcoin. Suppose Alice creates a transaction T and sends it to network peers. Note that T includes a message (i.e., payload). The network peers add T to their pool of unconfirmed transactions, known as *mempool*. They propagate T so that it can be included in all peers' views of the *mempool*. T remains in the *mempool* until a miner includes it in a block, at which point it is removed, and a transaction fee is transferred from Alice to the miner. The key observation is, until T is mined, Alice can overwrite it with another transaction T' (embed empty message). When this happens, T is dropped from the *mempool*. Thus, both T and T' are propagated to all users, but Alice only pays for T'. One can see that T is broadcast for free, and we denote T as pending transaction before overwrite. To conclude, transaction ghosting guarantees that a sender and a receiver can efficiently establish an Ethereum P2P channel via pending transactions such as T. The core idea is that the sender can transmit messages to the receiver by embedding them in pending transactions.

C SECURITY ANALYSIS OF GENERIC FRAMEWORK

In this section, we present the security analysis of the proposed generic framework, including indistinguishability, adaptive collision-resistance, and accountability.

C.1 Proof of Theorem C.1

THEOREM C.1. *The proposed generic framework is indistinguishable if the* CH *scheme is indistinguishable.*

PROOF. The reduction is executed between an adversary $\mathcal{A}$ and a simulator $\mathcal{S}$. Assume that $\mathcal{A}$ activates at most $n(\lambda)$ chameleon hashes. Let $\mathcal{S}$ denote a distinguisher against CH, who is given a chameleon public key pk^* and a HashOrAdapt oracle, aims to break the indistinguishability of CH. In particular, $\mathcal{S}$ is allowed to access the chameleon trapdoor $Dlog(\mathsf{pk}^*)$ [19]. $\mathcal{S}$ randomly chooses $g \in [1, n(\lambda)]$ as a guess for the index of the HashOrAdapt query. In the g-th query, $\mathcal{S}$'s challenger directly hashes a message $(h, \mathsf{r}) \leftarrow \mathsf{Hash}(\mathsf{pk}^*, m)$, instead of calculating the chameleon hash and randomness (h, r) using Adapt algorithm.

$\mathcal{S}$ sets up the game for $\mathcal{A}$ by distributing a master secret key to a group of users in a committee. $\mathcal{S}$ can honestly generate secret keys for any modifier associated with an access privilege Λ. If $\mathcal{A}$ submits a tuple (m_0, m_1, δ) in the g-th query, then $\mathcal{S}$ first obtains a chameleon hash (h_b, r_b) from his HashOrAdapt oracle. Then, $\mathcal{S}$ simulates a message-signature pair (c, σ), and a ciphertext C on message $Dlog(\mathsf{pk}^*)$ according to the protocol specification. Eventually, $\mathcal{S}$ returns $(h_b, \mathsf{r}_b, C, c, \sigma)$ to $\mathcal{A}$. $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs. If $\mathcal{A}$ guesses the random bit correctly, then $\mathcal{S}$ can break the indistinguishability of CH. □

C.2 Proof of Theorem C.2

THEOREM C.2. *The proposed generic framework is adaptively collision-resistant if the* ABET *scheme is semantically secure, the* CH *scheme is collision-resistant, and the* DPSS *scheme has secrecy.*

PROOF. We define a sequence of games $\mathbb{G}_i$, $i = 0, \cdots, 4$ and let Adv_i^{GF} denote the advantage of the adversary in game $\mathbb{G}_i$. Assume that $\mathcal{A}$ issues at most $n(\lambda)$ queries to Hash oracle.

- $\mathbb{G}_0$: This is the original game for adaptive collision-resistance.
- $\mathbb{G}_1$: This game is identical to game $\mathbb{G}_1$ except that $\mathcal{S}$ will output a random bit if $\mathcal{A}$ outputs a correct master secret key when no more than t users in a committee are corrupted, and the setup is honest (or the dealer is honest). The difference between $\mathbb{G}_0$ and $\mathbb{G}_1$ is negligible if DPSS scheme has secrecy.

$$\left| \mathsf{Adv}_0^{GF} - \mathsf{Adv}_1^{GF} \right| \leq \mathsf{Adv}_{\mathcal{S}}^{\mathsf{DPSS}}(\lambda). \tag{1}$$

- $\mathbb{G}_2$: This game is identical to game $\mathbb{G}_1$ except the following difference: $\mathcal{S}$ randomly chooses $g \in [1, n(\lambda)]$ as a guess for the index of the Hash′ oracle which returns the chameleon hash $(h^*, m^*, \mathsf{r}^*, C^*, c^*, \sigma^*)$. $\mathcal{S}$ will output a random bit if $\mathcal{A}$'s attacking query does not occur in the g-th query. Therefore, we have

$$\mathsf{Adv}_1^{GF} = n(\lambda) \cdot \mathsf{Adv}_2^{GF} \tag{2}$$

- $\mathbb{G}_3$: This game is identical to game $\mathbb{G}_2$ except that in the g-th query, the encrypted message $\mathsf{sk}_{\mathsf{CH}}$ in C^* is replaced by "$\perp$" (i.e., an empty value). Below we show that the difference between $\mathbb{G}_2$ and $\mathbb{G}_3$ is negligible if ABET scheme is semantically secure.
  Let $\mathcal{S}$ denote an attacker against ABET with semantic security, who is given a public key pk^* and a key generation oracle, aims to distinguish between encryptions of M_0 and M_1 associated with a challenge index j^* and a challenge set of attributes δ^*, which are predetermined at the beginning of the game for semantic security. $\mathcal{S}$ simulates the game for $\mathcal{A}$ as follows.
  - $\mathcal{S}$ sets up $\mathsf{mpk}_{\mathsf{ABET}} = \mathsf{pk}^*$ and completes the remainder of Setup honestly, which includes user's key pairs and chameleon key pairs for hashing in CH. $\mathcal{S}$ returns all public information to $\mathcal{A}$.
  - $\mathcal{S}$ can honestly answer the queries made by $\mathcal{A}$ regarding decryption keys using his given oracle, such that $\Lambda_i(\delta^*) \neq 1, i \neq j^*$. In the g-th query, upon receiving a hash query w.r.t., a hashed message m^* from $\mathcal{A}$. $\mathcal{S}$ first submits two messages (i.e., $[M_0 = \mathsf{sk}_{\mathsf{CH}}, M_1 = \perp]$) to his challenger, and obtains a challenge ciphertext C^* under index j^* and δ^*. Then, $\mathcal{S}$ returns the tuple $(h^*, m^*, \mathsf{r}^*, C^*, c^*, \sigma^*)$ to $\mathcal{A}$. Note that $\mathcal{S}$ can simulate the message-signature pair (c^*, σ^*) honestly using user's key pairs. Besides, $\mathcal{S}$ can simulate the adapt query successfully using $\mathsf{sk}_{\mathsf{CH}}$.
  If the encrypted message in C^* is $\mathsf{sk}_{\mathsf{CH}}$, then the simulation is consistent with $\mathbb{G}_2$; Otherwise, the simulation is consistent with $\mathbb{G}_3$. Therefore, if the advantage of $\mathcal{A}$ is significantly different in $\mathbb{G}_2$ and $\mathbb{G}_3$, $\mathcal{S}$ can break the semantic security of the ABET. Hence, we have

$$\left| \mathsf{Adv}_2^{GF} - \mathsf{Adv}_3^{GF} \right| \leq \mathsf{Adv}_{\mathcal{S}}^{\mathsf{ABET}}(\lambda). \tag{3}$$

- $\mathbb{G}_4$: This game is identical to game $\mathbb{G}_3$ except that in the g-th query, $\mathcal{S}$ outputs a random bit if $\mathcal{A}$ outputs a valid collision $(h^*, m^{*\prime}, \mathsf{r}^{*\prime}, C^{*\prime}, c^{*\prime}, \sigma^{*\prime})$, and it was not previously returned by the Adapt′ oracle. Below we show that the difference between $\mathbb{G}_3$ and $\mathbb{G}_4$ is negligible if CH is collision-resistant.
  Let $\mathcal{S}$ denote an attacker against CH with collision-resistant, who is given a chameleon public key pk^* and an Adapt′ oracle, aims

to find a collision which was not simulated by the Adapt' oracle. $\mathcal{S}$ simulates the game for $\mathcal{A}$ as follows.

– $\mathcal{S}$ sets up $\mathsf{pk}_{\mathsf{CH}} = \mathsf{pk}^*$ for the g-th hash query, and completes the remainder of Setup honestly, which includes user's key pairs and master key pair in ABET. $\mathcal{S}$ returns all public information to $\mathcal{A}$.

– $\mathcal{S}$ can simulate all queries made by $\mathcal{A}$ except adapt queries. If $\mathcal{A}$ submits an adapt query in the form of $(h, m, \mathsf{r}, C, c, \sigma, m')$, then $\mathcal{S}$ obtains a randomness r' from his Adapt' oracle, and returns $(h, m', \mathsf{r}', C', c', \sigma')$ to $\mathcal{A}$. In particular, $\mathcal{S}$ simulates the g-th hash query as $(h^*, m^*, \mathsf{r}^*, C^*, c^*, \sigma^*)$ w.r.t. a hashed message m^*, where $C^* \leftarrow \mathsf{Enc}_{\mathsf{ABET}}(\mathsf{mpk}^*, \perp, \delta^*, j^*)$, c^* is derived from $\perp$ because $Dlog(\mathsf{pk}^*)$ is unknown.

– If $\mathcal{A}$ outputs a collision $(h^*, m^*, \mathsf{r}^*, C^*, c^*, \sigma^*, m^{*'}, \mathsf{r}^{*'}, C^{*'}, c^{*'}, \sigma^{*'})$ with respect to the g-th query, and all relevant checks are succeed, then $\mathcal{S}$ output $(h^*, m^{*'}, \mathsf{r}^{*'}, C^{*'}, c^{*'}, \sigma^{*'})$ as a valid collision to CH; Otherwise, $\mathcal{S}$ aborts the game. Therefore, we have

$$\left| \mathsf{Adv}_3^{GF} - \mathsf{Adv}_4^{GF} \right| \le \mathsf{Adv}_{\mathcal{S}}^{\mathsf{CH}}(\lambda). \tag{4}$$

Combining the above results together, we have

$$\begin{aligned} \mathsf{Adv}_{\mathcal{A}}^{GF}(\lambda) \quad \le \quad & n(\lambda) \cdot (\mathsf{Adv}_{\mathcal{S}}^{\mathsf{DPSS}}(\lambda) + \mathsf{Adv}_{\mathcal{S}}^{\mathsf{ABET}}(\lambda) \\ & + \mathsf{Adv}_{\mathcal{S}}^{\mathsf{CH}}(\lambda)). \end{aligned}$$

□

C.3 Proof of Theorem C.3

THEOREM C.3. *The proposed generic framework is accountable if the Σ scheme is EUF-CMA secure, and the DPSS scheme has correctness.*

PROOF. We define a sequence of games $\mathbb{G}_i$, $i = 0, \cdots, 2$ and let Adv_i^{GF} denote the advantage of the adversary in game $\mathbb{G}_i$.

• $\mathbb{G}_0$: This is the original game for accountability.

• $\mathbb{G}_1$: This game is identical to game $\mathbb{G}_1$ except that $\mathcal{S}$ will output a random bit if all honest users in a committee outputs a master secret key msk' such that $\mathsf{msk}' \ne \mathsf{msk}$, and the setup is honest. The difference between $\mathbb{G}_0$ and $\mathbb{G}_1$ is negligible if DPSS scheme has correctness.

$$\left| \mathsf{Adv}_0^{GF} - \mathsf{Adv}_1^{GF} \right| \le \mathsf{Adv}_{\mathcal{S}}^{\mathsf{DPSS}}(\lambda). \tag{5}$$

• $\mathbb{G}_2$: This game is identical to game $\mathbb{G}_1$ except that $\mathcal{S}$ will output a random bit if $\mathcal{A}$ outputs a valid forgery σ^*, where σ^* was not previously simulated by $\mathcal{S}$ and the user is honest. The difference between $\mathbb{G}_1$ and $\mathbb{G}_2$ is negligible if Σ is EUF-CMA secure. Let $\mathcal{F}$ denote a forger against Σ, who is given a public key pk^* and a signing oracle O^{Sign}, aims to break the EUF-CMA security of Σ. Assume that $\mathcal{A}$ activates at most n users in the system.

– $\mathcal{F}$ randomly chooses a user in the system and sets up its public key as pk^*. $\mathcal{F}$ completes the remainder of Setup honestly. Below we mainly focus on user pk^* only.

– To simulate a chameleon hash for message m, $\mathcal{F}$ first obtains a signature σ from his signing oracle O^{Sign}. Then, $\mathcal{F}$ generates chameleon hash and ciphertext honestly because $\mathcal{F}$ chooses the chameleon secret key $\mathsf{sk}_{\mathsf{CH}}$, and returns $(m, h, \mathsf{r}, C, c, \sigma)$ to $\mathcal{A}$. Besides, the message-signature pairs and collisions can be

perfectly simulated by $\mathcal{F}$ for any adapt query. $\mathcal{F}$ records all the simulated message-signature pairs by including them to a set Q.

– When forging attack occurs, i.e., $\mathcal{A}$ outputs $(m^*, h^*, \mathsf{r}^*, C^*, c^*, \sigma^*)$, $\mathcal{F}$ checks whether:

* the forging attack happens to user pk^*;

* the ciphertext C^* encrypts the chameleon trapdoor $\mathsf{sk}_{\mathsf{CH}}$;

* the message-signature pair (c^*, σ^*) is derived from the chameleon trapdoor $\mathsf{sk}_{\mathsf{CH}}$;

* the message-signature pair $(c^*, \sigma^*) \notin Q$;

* $1 \leftarrow \Sigma.\mathsf{Verify}(\mathsf{pk}^*, c^*, \sigma^*)$ and $1 \leftarrow \mathsf{Verify}(\mathsf{pk}_{\mathsf{CH}}, m^*, h^*, \mathsf{r}^*)$.

If all the above conditions hold, $\mathcal{F}$ confirms that it as a successful forgery from $\mathcal{A}$, then $\mathcal{F}$ extracts the forgery via $\sigma \leftarrow M_{\Sigma}(\mathsf{PP}, \mathsf{pk}^*, \sigma^*, \Delta(\mathsf{sk}))$ due to the homomorphic property of Σ (regarding keys and signatures), where $\Delta(\mathsf{sk})$ is derived from (c, c^*). To this end, $\mathcal{F}$ outputs σ as its own forgery; Otherwise, $\mathcal{F}$ aborts the game.

$$\left| \mathsf{Adv}_1^{GF} - \mathsf{Adv}_2^{GF} \right| \le n \cdot \mathsf{Adv}_{\mathcal{F}}^{\Sigma}(\lambda). \tag{6}$$

Combining the above results together, we have

$$\mathsf{Adv}_{\mathcal{A}}^{GF}(\lambda) \quad \le \quad \mathsf{Adv}_{\mathcal{S}}^{\mathsf{DPSS}}(\lambda) + n \cdot \mathsf{Adv}_{\mathcal{F}}^{\Sigma}(\lambda).$$

□

D SECURITY ANALYSIS OF ABET

In this section, we present the security analysis of the proposed ABET scheme, including semantic security and ciphertext anonymity.

D.1 Semantic Security

Informally, an ABE scheme is secure against chosen plaintext attacks if no group of colluding users can distinguish between encryption of M_0 and M_1 under an index and a set of attributes δ^* of an attacker's choice as long as no member of the group is authorized to decrypt on her own. The selective security is defined as an index, as well as a set of attributes δ^*, are chosen by attackers at the beginning security experiment. The semantic security model here is based on the selective security model defined in [46].

THEOREM D.1. *The proposed ABET scheme is semantically secure in the standard model if the q'-type assumption is held in the asymmetric pairing groups.*

PROOF. Let $\mathcal{S}$ denote a q'-type problem attacker, who is given the terms from the assumption, aims to distinguish g^{abc} and g^s. The reduction is performed as follows.

• $\mathcal{S}$ simulates master public key $\mathsf{mpk} = (g, u, v, w, h, \hat{e}(g, h)^{\alpha}, \{g_1^{\alpha}, \cdots, g_k^{\alpha}\}, \{h_1^{\alpha}, \cdots h_k^{\alpha}\}, g^{\beta}, h^{1/\alpha}, h^{\beta/\alpha}, \hat{e}(g, h)^{\theta/\alpha})$ as follows: $\hat{e}(g, h)^{\alpha} = \hat{e}(g^a, h^b)$, $\{g_1^{\alpha}, \cdots g_k^{\alpha}\} = \{g^{abdz_1}, \cdots, g^{abdz_k}\}$, $\{h_1^{\alpha}, \cdots h_k^{\alpha}\} = \{h^{abdz_1}, \cdots, h^{abdz_k}\}$, $h^{1/\alpha} = h^{d/ab}$, $h^{\beta/\alpha} = h^{\beta d/ab}$, $\hat{e}(g, h)^{\theta/\alpha} = \hat{e}(g^{\theta}, h^{d/ab})$, and (u, v, w) are simulated using the same method described in [46]. Note that $(\beta, \theta, \{z_1, \cdots, z_k\})$ are randomly chosen by $\mathcal{S}$, and α (or $1/\alpha$) is implicitly assigned as ab (or d/ab) from the q'-type assumption. $\mathcal{A}$ submits a challenge index $j^* = \{I_1, \cdots, I_j\}$ and a set of attributes δ^* to $\mathcal{S}$.

- $\mathcal{S}$ simulates decryption keys $\mathsf{sk}_{\Lambda_i} = (\{\mathsf{sk}_\tau\}_{\tau \in [n_1]}, \mathsf{sk}_0, \mathsf{sk}_1, \mathsf{sk}_2)$ (note that $\Lambda_i(\delta^*) \neq 1$) as follows: $\mathsf{sk}_0 = (g^{dt/ab}, g^r)$, $\mathsf{sk}_1 = g^\theta \cdot \widehat{i}^t \cdot g^{\beta \cdot r}$, $\mathsf{sk}_2 = \{g_{i-1}^{abdt}, \cdots, g_1^{abdt}\}$, where $\widehat{i} = g_k^{abdI_1} \cdots g_i^{abdI_i}$ · g. Note that $\{\mathsf{sk}_\tau\}_{\tau \in [n_1]}$ is simulated using the same method described in [46], and (t, r) are randomly chosen by $\mathcal{S}$.
- $\mathcal{S}$ simulates challenge ciphertext $C^* = (ct, \{ct_{\tau,1}, ct_{\tau,2}\}_{\tau \in [\delta]}, ct_0, ct_1)$ as follows: $ct = M_b \oplus \mathsf{H}(T || \hat{e}(g, h)^{\theta cd/ab})$, $ct_0 = (h^c, h^{cd/ab}, h^{\beta cd/ab})$, and $ct_1 = \widehat{j}^c = h_k^{abcdI_1} \cdots h_j^{abcdI_j} \cdot h$. Note that $\{ct_{\tau,1}, ct_{\tau,2}\}_{\tau \in [\delta]}$ are simulated using the same method described in [46], s is implicitly assigned as c from the q'-type assumption, and T can be either $\hat{e}(g, h)^{abc}$ or $\hat{e}(g, h)^s$.

Finally, $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs. If $\mathcal{A}$ guesses the random bit correctly, then $\mathcal{S}$ can break the q'-type problem.

□

D.2 Ciphertext Anonymity

Informally, ciphertext anonymity requires that any third party cannot distinguish the encryption of a chosen message for a first chosen index from the encryption of the same message for a second chosen index. In other words, the attacker cannot decide whether a ciphertext was encrypted for a chosen index or a random index. We prove the ABET scheme has selective ciphertext anonymity (i.e., the index is chosen prior to the security experiment).

THEOREM D.2. *The proposed ABET scheme is anonymous if the eDDH assumption is held in the asymmetric pairing groups.*

PROOF. Let $\mathcal{S}$ denote an eDDH problem distinguisher, who is given terms from the assumption, aims to distinguish $h^{c/ab}$ and h^s. The reduction is performed as follows.

- $\mathcal{S}$ simulates master public key $\mathsf{mpk} = (g, u, v, w, h, \hat{e}(g, h)^\alpha, \{g_1^\alpha, \cdots g_k^\alpha\}, \{h_1^\alpha, \cdots h_k^\alpha\}, g^\beta, h^{1/\alpha}, h^{\beta/\alpha}, \hat{e}(g, h)^{\theta/\alpha})$ as follows: $\hat{e}(g, h)^\alpha = \hat{e}(g, h)^{ab}$, $\{g_1^\alpha, \cdots, g_k^\alpha\} = \{g_1^{ab}, \cdots, g_k^{ab}\}$, $\{h_1^\alpha, \cdots, h_k^\alpha\} = \{h_1^{ab}, \cdots, h_k^{ab}\}$, $h^{1/\alpha} = h^{1/ab}$, $h^{\beta/\alpha} = h^{\beta/ab}$, $\hat{e}(g, h)^{\theta/\alpha} = \hat{e}(g, h)^{\theta/ab}$. Note that $\mathcal{S}$ randomly chooses (u, v, w) and $(\beta, \theta, \{z_i\})$, and implicitly sets $\alpha = ab$. $\mathcal{A}$ submits a challenge index $j^* = \{I_1, \cdots, I_j\}$ to $\mathcal{S}$.
- $\mathcal{S}$ simulates a decryption key $\mathsf{sk}_{\Lambda_i} = (\{\mathsf{sk}_\tau\}_{\tau \in [n_1]}, \mathsf{sk}_0, \mathsf{sk}_1, \mathsf{sk}_2)$ with respect to index $i = \{I_1, \cdots, I_i\}$ as follows: $\mathsf{sk}_0 = (g^t, g^{\beta \cdot r} \cdot g^{-\Sigma(z_i I_i)t \cdot b})$, $\mathsf{sk}_1 = g^\theta \cdot g^{a \cdot r}$, where $(t, r) \in \mathbb{Z}_q$ are randomly chosen by $\mathcal{S}$. The components $(\{\mathsf{sk}_\tau\}_{\tau \in [n_1]}, \mathsf{sk}_2)$ are honestly simulated by $\mathcal{S}$. The simulated components $(g^t, g^{\beta \cdot r} \cdot g^{-\Sigma(z_i I_i)t \cdot b}, g^\theta \cdot g^{a \cdot r})$ are correctly distributed, because $g^\theta \cdot g^{a \cdot r} = g^\theta \cdot \widehat{i}^t \cdot g^{a[\beta \cdot r - \Sigma(z_i I_i)t \cdot b]} = g^\theta \cdot \widehat{i}^t \cdot g^{a \cdot \bar{r}}$, where $\bar{r} = \beta \cdot r - \Sigma(z_i I_i)t \cdot b$, and $\widehat{i} = g_k^{abI_1} \cdots g_i^{abI_i} \cdot g$. So, the simulated components $(g^t, g^{\bar{r}}, g^\theta \cdot \widehat{i}^t \cdot g^{a \cdot \bar{r}})$ match the real distribution.
- $\mathcal{S}$ simulates the challenge ciphertext $C^* = (ct, \{ct_{\tau,1}, ct_{\tau,2}\}_{\tau \in [\delta]}, ct_0, ct_1)$ with respect to index j^* as follows: $ct = M_b \oplus \mathsf{H}(\hat{e}(g, h)^{abc} || \hat{e}(g, T)^\beta)$, $ct_0 = (h^c, T, T^\beta)$, and $ct_1 = h_k^{abcI_1} \cdots h_j^{abcI_j} \cdot h^c$. Note that $\mathcal{S}$ simulates $\{ct_{\tau,1}, ct_{\tau,2}\}_{\tau \in [\delta]}$ honestly, and T can be either $h^{c/ab}$ or h^s.

Finally, $\mathcal{S}$ outputs whatever $\mathcal{A}$ outputs. If $\mathcal{A}$ guesses the random bit correctly, then $\mathcal{S}$ can break the eDDH problem.

□