Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

HERMES: Provably Secure Private Accounts for Account-Model Blockchains via Commitment-Coherent Encryption

Haofeng Guo

71255902022@stu.ecnu.edu.cn

East China Normal University https://orcid.org/0009-0009-0447-0361

Research Article

Keywords: account model, commitment-coherent encryption, Halo2, privacy-preserving transactions, zk-Rollup

Posted Date: October 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10608242/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

HERMES: Provably Secure Private Accounts for Account-Model Blockchains via Commitment-Coherent Encryption

Haofeng Guo

***Abstract*—Account-model blockchains record every balance, transfer amount, and counterparty in public. Existing privacy designs close only part of this gap: schemes encrypting balances leave sender–receiver relations observable, while systems hiding those relations rely on off-chain notification. We present HERMES, a privacy-preserving transaction scheme for account-model blockchains providing balance confidentiality, amount indistinguishability, and sender–receiver unlinkability at once. Our contributions are twofold. First, commitment-coherent encryption (CCE), whose ciphertext is simultaneously a Pedersen commitment, so that range and validity statements are proved directly over ciphertexts with no bridging protocol. Second, the HERMES design realizing it: a triple-key account structure separating these roles so that any single key compromise degrades at most one guarantee, three transaction types covering the full lifecycle, constant-cost membership proofs over an anonymity set, view tags for non-interactive payment discovery, and unlinkable nullifiers against double spending. We implement HERMES over BN254 with Halo2 and no trusted setup, and reduce the three guarantees to the decisional Diffie–Hellman assumption with explicit tightness. The evaluation shows a transfer circuit several times smaller than a non-CCE baseline, recursive aggregation amortizing on-chain verification to tens of bytes per transaction at over one hundred transactions per second, and verification costs orders of magnitude below prior work.**

***Index Terms*—account model, commitment-coherent encryption, Halo2, privacy-preserving transactions, zk-Rollup**

I. Introduction

A. Background and Motivation

Account-model public blockchains (Ethereum [1] and the EVM ecosystem) bind a persistent on-chain identity to a mutable balance, exposing transacting parties and transferred value to any observer. This transparency is a liability for DeFi users, whose front-running susceptibility scales with mempool visibility [2], for stablecoin and central-bank-digital-currency issuers who need retail balances confidential, and for enterprise operators whose flows would be exposed through on-chain relationship graphs [3].

Bringing privacy to the account model is structurally harder than to the UTXO model. In UTXO chains (Zcash [4], Monero [5]), each note is a freshly issued output disjoint from prior history, so anonymity sets can be assembled at the granularity of individual notes. In account chains, by contrast, a single persistent identifier accumulates every interaction, and the privacy requirement decomposes into *three* simultaneous properties:

(P1) *Balance confidentiality*: an observer learns nothing about the current account balance;

(P2) *Transaction-amount indistinguishability*: an observer cannot tell which of two candidate transfer amounts a transaction carries;

(P3) *Sender-receiver unlinkability*: an observer cannot tell who paid whom.

These three properties stand in tension. Hiding amounts via decryptable encryption invites algebraic-leakage attacks between successive ciphertexts (e.g., when a shared randomizer correlates balance updates). Hiding addresses via anonymity sets balloons proof size. Verifying both on-chain demands a zero-knowledge proof system whose verifier cost stays constant in the anonymity-set size. No existing work reconciles all three properties under a single *formal* framework with explicit reductions to a standard hardness assumption.

B. Related Work

Prior confidential-payment designs each tackle a strict subset of the three properties. We analyze the design space along four orthogonal axes that any privacy account model must address, and position existing work within this framework.

The choice of encryption primitive determines whether balance ciphertexts can be decrypted on-chain without off-chain side channels. *Pedersen commitments* [6] are hiding but not decryptable, so BlockMaze [7] requires off-chain notifications (violating autonomous discovery). *Standard ElGamal* [8] is decryptable but its second component is not a commitment, forcing a bridging Σ-protocol for range proofs (Zether [9]). *Twisted ElGamal* (underlying PGC [10], Solana CT [11], and cWETH [12]) restructures the ciphertext as $(\mathsf{pk}^r, g^r \cdot h^m)$ so that one component is structurally a Pedersen commitment under the fixed bases (g, h), eliminating the bridging step. HERMES follows this line, adopting the related *dual-generator exponential* form, and formalizes the underlying coherence as a standalone primitive (§III).

Sender–receiver unlinkability hides the input/output mapping. UTXO systems (e.g., Monero [5]) use ring signatures or per-transaction mixing but expose no persistent account balances for composability. Account-model approaches are either *explicit-visibility* (PGC, cWETH: sender/receiver keys in cleartext, no unlinkability) or *anonymity-set/SNARK-based* (Zether, Aleo [13], Penumbra [14]). HERMES follows the anonymity-set approach (§IV).

H. Guo (ORCID: 0009-0009-0447-0361) is with the School of Software Engineering, East China Normal University, Shanghai, China. E-mail: 71255902022@stu.ecnu.edu.cn

A critical gap is the absence of a *unified* formal model. Zether argues properties in prose without game-based definitions; PGC proves amount privacy and balance confidentiality individually but not unlinkability (its account model has no recipient anonymity); cWETH's model targets a specific ERC-20 instantiation. No prior work jointly covers all three properties in one game-based framework with explicit reduction tightness—HERMES fills this with a unified DDH-based framework (§V) reporting tightness per reduction.

The proof system sets the trust model. Groth16 and universal-setup SNARKs (Sonic [15]) require (per-circuit or universal) trusted setup; folding schemes (Nova [16]) inherit it. Only IPA-based systems (Halo2 [17]) are fully transparent. HERMES defaults to Halo2's IPA backend for transparency (§II); Aztec [18] shows recursive-SNARK rollups are feasible but relies on trusted setup.

Across these four axes (systematic comparison deferred to Supplementary Appendix G), HERMES is the only construction combining all four properties: a commitment-coherent primitive enabling zero-bridge circuits, unlinkability via anonymity sets, all three properties reduced to DDH in one game-based model, and no proof-system-layer trusted setup (Halo2 IPA); a fuller property and efficiency comparison appears in §VI-G.

C. Contributions

This paper makes three contributions:

C1: Commitment-coherent encryption (CCE), a new cryptographic primitive with a complete formal theory (§III). We introduce CCE and endow it with standalone syntax, correctness conditions (C1)–(C4), and security notions (IND-CPA, composition, simulation extractability), all reduced to the standard DDH assumption in a single game-based model with explicit tightness. We prove an algebraic characterization (Theorem III.5): any account-model privacy scheme satisfying balance confidentiality, homomorphic balance conservation, and representation unity is necessarily equivalent to dual-generator exponential ElGamal up to affine reparameterization. We further establish a group-element lower bound that is matched exactly, a scalar-multiplication lower bound in the Generic Group Model matched up to fixed-base precomputation (Theorems III.6–III.7), zero-bridging optimality in ZK circuits (Theorem III.10), and the uniqueness of per-ciphertext independent randomization for full leakage closure under DDH (Theorem V.6). CCE is primitive-level and proof-system-agnostic, applying to any account-model construction beyond HERMES itself.

C2: HERMES, a CCE-based privacy-preserving transaction scheme for the account model (§IV). The scheme realizes full-ciphertext balances in which per-ciphertext independent randomizers eliminate the algebraic-leakage surface of reused-randomizer designs, and a triple-key separation architecture (amount key, view key, address) that provably bounds the compromise radius of any single key to its own role (Proposition IV.3). It defines three transaction types—Transfer, AccountOpen, and Withdraw—covering the complete lifecycle with nullifier-based double-spend prevention, and achieves non-interactive incoming-payment discovery via view tags, removing the off-chain notification dependency of BlockMaze-style schemes. Anonymity-set membership proofs hide sender and receiver within a set of size N while on-chain verification stays constant-cost.

C3: A five-requirement ZK circuit design with zero bridging, and its empirical validation (§IV-E, §VI). We specify the proof obligations R1–R5 (balance validity, transaction correctness, range constraints, state membership, and anonymity-set indistinguishability) and show that CCE coherence lets range proofs act directly on ciphertexts, eliminating the bridging Σ-protocol that non-CCE schemes require—an empirically confirmed saving of $4{,}072$ constraints (80%) at $N{=}8$. A BN254 prototype corroborates practicality: all cryptographic-layer primitives complete in the microsecond-to-millisecond regime, with batch-amortized on-chain verification.

D. Paper Organization

§II fixes notation and recalls standard building blocks. §III formalizes the CCE primitive and its security model. §IV presents the HERMES scheme. §V gives the game-based security analysis. §VI reports the prototype evaluation. §VII concludes and outlines future work.

II. Preliminaries

We fix notation and recall the standard building blocks used in the sequel: DDH-hard groups, ElGamal encryption and Pedersen commitments, and PLONKish zero-knowledge proofs.

A. Notation and Assumptions

All algorithms are parameterized by a security parameter λ. HERMES operates in a prime-order cyclic group $\mathbb{G}_1$ of order q (instantiated on BN254), with scalar operations in $\mathbb{F}_q$; $\epsilon \leftarrow \mathcal{A}(x)$ denotes $\mathcal{A}$ executed on x with fresh randomness.

Definition II.1 (Decisional Diffie-Hellman Problem (DDH))**.** *Given $(g, g^a, g^b, g^c) \in \mathbb{G}_1^4$, determine whether $c \stackrel{?}{=} ab \pmod q$. The DDH assumption states that for any probabilistic polynomial-time (PPT) algorithm $\mathcal{A}$:*

$$\left|\Pr[\mathcal{A}(g, g^a, g^b, g^{ab})=1] - \Pr[\mathcal{A}(g, g^a, g^b, g^c)=1]\right| \leq \mathsf{negl}(\lambda) \tag{1}$$

where $a, b, c \xleftarrow{\$} \mathbb{Z}_q$.

B. ElGamal Encryption and Pedersen Commitments

ElGamal [8] is a public-key encryption scheme whose semantic security reduces to DDH. Over a group $\mathbb{G} = \langle g \rangle$ of order q with public key $\mathsf{pk} = g^{\mathsf{sk}}$, *exponential ElGamal* places the message in the exponent: encryption of $m \in \mathbb{Z}_q$ under randomness r outputs $c = (c_1, c_2) = (g^r, \mathsf{pk}^r \cdot g^m)$, and decryption recovers $g^m = c_2 \cdot c_1^{-\mathsf{sk}}$. The encoding is additively homomorphic, $\mathsf{Enc}(m_1; r_1) \cdot \mathsf{Enc}(m_2; r_2) = \mathsf{Enc}(m_1{+}m_2;\, r_1{+}r_2)$, at the cost that exponent recovery is bounded-range: for $m \in [0, \mathsf{MAX})$ it takes $O(\sqrt{\mathsf{MAX}})$ group operations via Baby-Step Giant-Step (BSGS) [19]. A *Pedersen commitment* [6] to m with randomness ρ over independent generators (g, h) with unknown $\log_g h$ is

$$\mathsf{PedCom}(m; \rho) = g^\rho \cdot h^m; \tag{2}$$

it is perfectly hiding, computationally binding under the discrete-logarithm assumption, and additively homomorphic in (m, ρ).

C. Zero-Knowledge Proofs and Halo2

A zero-knowledge proof lets a prover convince a verifier of a statement's validity without revealing more; formal definitions follow [20].

Definition II.2 (zkSNARK)**.** *A zero-knowledge succinct non-interactive argument of knowledge (zkSNARK) is a triple* $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$*:* $\mathsf{Setup}(1^\lambda, C) \to (pk, vk)$*,* $\mathsf{Prove}(pk, x, w) \to \pi$*,* $\mathsf{Verify}(vk, x, \pi) \to \{0, 1\}$*. It satisfies* ***Completeness*** *(valid proofs verify),* ***Knowledge Soundness*** *(proofs require witness knowledge), and* ***Zero-Knowledge*** *(proofs reveal nothing beyond validity).*

PLONKish Arithmetization. PLONK [21], extended by Halo2 [17], supports three constraint mechanisms: custom gates (algebraic relations within a row), permutation constraints (identical values across cells), and lookup arguments (cell values in a precomputed table, e.g., for range proofs).

Halo2 and Recursive Aggregation. Halo2's Inner Product Argument (IPA) backend provides transparent setup and native recursion, and its Kate–Zaverucha–Goldberg (KZG) backend aggregates M proofs into one whose size and verification time are independent of M [22].

III. Commitment-Coherent Encryption: Formal Framework

This section introduces *commitment-coherent encryption* (CCE), a cryptographic primitive in which encrypting a message simultaneously commits to it: the ciphertext contains a component that is itself a commitment to the plaintext, so zero-knowledge proofs attach to ciphertexts directly, eliminating the commitment-conversion overhead in zero-knowledge circuits. Although this ciphertext–commitment coherence has been exploited implicitly in prior constructions (notably PGC [10]), CCE elevates it into a standalone primitive with its own syntax, correctness conditions, and security notions. We give the formal definition, security model, and instantiation (§III-A–§III-B), then prove that this instantiation is the unique coherent one up to reparameterization (§III-C) and state its lower bounds, composition guarantees, and failure modes (§III-D–§III-F); complete proofs are deferred to the Supplementary Material.

A. Formal Definition and Security Model

a) Algorithm specification: A *commitment-coherent encryption scheme* Π_{CCE} over a prime-order group $\mathbb{G} = \langle g \rangle$ of order q with an independent generator $h \in \mathbb{G}$ ($\log_g h$ unknown) consists of seven PPT algorithms $\Pi_{\mathrm{CCE}} = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Com}, \mathsf{HomAdd}, \mathsf{ReRand})$, where the message space $\mathcal{M}$, randomness space $\mathcal{R}$, ciphertext space $\mathcal{C}$, and commitment space $\mathcal{P}$ are determined by the public parameters pp:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{pp}$: samples $\mathsf{pp} = (\mathbb{G}, q, g, h)$, where $h = \mathsf{Hash\text{-}to\text{-}Curve}(g)$ is derived deterministically with *no secret parameter* ($\log_g h$ unknown by hash unbiasedness and discrete-log hardness); no trusted setup is introduced.
- $\mathsf{KeyGen}(\mathsf{pp}) \to (pk, sk)$: samples $\mathsf{sk} \xleftarrow{\$} \mathbb{Z}_q$, computes $\mathsf{pk} = g^{\mathsf{sk}}$.
- $\mathsf{Enc}(\mathsf{pp}, pk, m; r) \to c$: samples $r \xleftarrow{\$} \mathbb{Z}_q$, outputs $c = (c_1, c_2) = (g^r, \mathsf{pk}^r h^m)$.
- $\mathsf{Dec}(\mathsf{pp}, sk, c) \to m$ **or** $\perp$: computes $t = c_2 \cdot c_1^{-\mathsf{sk}} = h^m$, then recovers m via BSGS over $[0, \mathsf{MAX})$; returns $\perp$ if no solution exists in range (illegal ciphertext or out-of-range plaintext).
- $\mathsf{Com}(c) \to \mathrm{com} \in \mathcal{P}$: deterministically extracts the commitment component $\mathrm{com} = c_2$.
- $\mathsf{HomAdd}(c, c') \to c''$: component-wise multiplication under the *same* public key; **closed only for ciphertexts under a common** pk (in the account model this is without loss, since balance updates always occur under one account key pk_a).
- $\mathsf{ReRand}(c; s) \to c'$: rerandomizes to $\mathsf{Enc}(\mathsf{pp}, pk, m; r{+}s)$ preserving the plaintext.

b) Correctness conditions: Let $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda)$, $(pk, sk) \leftarrow \mathsf{KeyGen}(\mathsf{pp})$. A valid CCE scheme satisfies four correctness conditions:

- **(C1) Decryption correctness.** For all $m \in \mathcal{M}$ and $r \in \mathcal{R}$, $\Pr[\mathsf{Dec}(sk, \mathsf{Enc}(pk, m; r)) = m] = 1$.
- **(C2) Homomorphic correctness.** For all $m_1, m_2 \in \mathcal{M}$ with $m_1 + m_2 \in \mathcal{M}$ and $r_1, r_2 \in \mathcal{R}$:

$$\mathsf{Dec}\big(sk, \mathsf{Enc}(pk, m_1{+}m_2;\, r_1{+}r_2)\big) = m_1 + m_2.$$

- **(C3) Rerandomization correctness.** For all $m \in \mathcal{M}$, $r, s \in \mathcal{R}$: $\mathsf{Dec}\big(sk,\ \mathsf{ReRand}(\mathsf{Enc}(pk, m; r); s)\big) = m$, and $\mathsf{ReRand}\big(\mathsf{Enc}(pk, m; r); s\big)$ is distributed identically to $\mathsf{Enc}(pk, m; r + s)$.
- **(C4) Commitment coherence.** There exist a deterministic map ϕ and a randomness function ρ such that for every legal ciphertext $c = \mathsf{Enc}(\mathsf{pp}, pk, m; r)$: $\phi\big(\mathsf{Com}(c)\big) = \mathsf{PedCom}(m;\ \rho(r))$, where PedCom is the standard Pedersen commitment. That is, the extracted component *is itself* a Pedersen commitment, so range proofs can be built directly on it without bridging.

c) Commitment coherence and its three sub-properties:

Definition III.1 (Commitment-Coherent Encryption (CCE))**.** *Scheme* Π *satisfies* ***commitment coherence*** *if it satisfies correctness (C1)–(C4) and, for the instantiation with* $\rho(\mathsf{sk}, r) = \mathsf{sk} \cdot r$*, the commitment component of every legal ciphertext is a Pedersen commitment that is simultaneously* extractable*,* hiding*, and* binding *(Definition III.2).*

Definition III.2 (Three Sub-Properties)**.**

(a) ***Extractability.*** *A public deterministic map* $\mathsf{Com}\colon \mathcal{C} \to \mathcal{P}$ *extracts* $\mathsf{Com}(c)$ *from every legal ciphertext* $c = \mathsf{Enc}(pk, m; r)$ *such that* $\phi(\mathsf{Com}(c)) = \mathsf{PedCom}(m; \rho(r))$*, where* ρ *is the deterministic randomness function determined by* (pk, r) *(in the instantiation,* $\rho(r) = \mathsf{sk} \cdot r$*).*

(b) ***Hiding (information-theoretic).*** *For all* $m_0, m_1 \in \mathcal{M}$*, the marginal distributions of the extracted commitment component are* identical*:* $\mathsf{Com}(\mathsf{Enc}(pk, m_0)) \equiv \mathsf{Com}(\mathsf{Enc}(pk, m_1))$ *over the uniform randomizer. This is* perfect *hiding, independent of any computational*

assumption (Theorem V.2); when the full ciphertext c (not just $\mathsf{Com}(c)$) is included in the tuple, the property degrades to computational IND-CPA under DDH (§V-B2).

(c) ***Binding (computational, with opening).*** *For any PPT adversary $\mathcal{A}$, the probability that $\mathcal{A}$ outputs two legal, decodable ciphertexts with openings (c_0, m_0, ρ_0) and (c_1, m_1, ρ_1) such that $\mathsf{Com}(c_0) = \mathsf{Com}(c_1)$ but $(m_0, \rho_0) \neq (m_1, \rho_1)$ is negligible (Theorem V.3). Requiring the adversary to output the opening is standard for Pedersen binding and is what makes the discrete-log reduction computable.*

d) Opening semantics: The commitment component $\mathsf{Com}(c) = \mathsf{pk}^r \cdot h^m$, with $\mathsf{pk} = g^{\mathsf{sk}}$, equals $g^{\mathsf{sk}\cdot r} \cdot h^m = \mathsf{PedCom}(m;\ \rho(r))$ with $\rho(r) = \mathsf{sk}\cdot r$ (Definition III.1), and its *opening* is the pair (m, ρ). The private-key holder decrypts c with sk to recover m and can therefore provide the complete opening (m, ρ); because $\rho(r) = \mathsf{sk}\cdot r$ is bound to the recipient public key, this commitment is *key-bound*: it is not publicly verifiable, and only the private-key holder can produce a valid opening.

e) CCE-IND-CPA security game:

Definition III.3 ($\mathsf{Game}_{\mathsf{CCE}}$: CCE-IND-CPA)**.** *Let λ be the security parameter. Challenger $\mathcal{C}$ and adversary $\mathcal{A}$ execute:*

1) ***Setup.*** *$\mathcal{C}$ runs $\mathsf{Setup}(1^\lambda)$, generates $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(\mathsf{pp})$, and sends $(\mathsf{pp}, \mathsf{pk})$ to $\mathcal{A}$.*
2) ***Challenge.*** *$\mathcal{A}$ outputs (m_0, m_1) with $m_0, m_1 \in [0, \mathsf{MAX})$. $\mathcal{C}$ samples $\beta \xleftarrow{\$} \{0,1\}$ and $r \xleftarrow{\$} \mathbb{Z}_q$, computes $(C_1^*, C_2^*) \leftarrow \mathsf{Enc}(\mathsf{pp}, \mathsf{pk}, m_\beta; r)$, and sends (C_1^*, C_2^*) to $\mathcal{A}$.*
3) ***Guess.*** *$\mathcal{A}$ outputs $\beta' \in \{0,1\}$.*

$\mathcal{A}$'s advantage is $\mathsf{Adv}^{\mathsf{CCE\text{-}IND\text{-}CPA}}(\mathcal{A}) = \left|\Pr[\beta' = \beta] - \frac{1}{2}\right|$.

$\mathsf{Game}_{\mathsf{CCE}}$ is syntactically identical to standard ElGamal IND-CPA; the distinction is semantic: because $C_2 = \mathsf{Com}(m; \mathsf{sk}\cdot r)$ is a Pedersen commitment, an adversary might attempt to exploit the commitment structure. Theorem V.4 (§V-B2) shows this provides no advantage beyond DDH, with advantage bound $\mathsf{Adv}^{\mathsf{CCE\text{-}IND\text{-}CPA}}(\mathcal{A}) = \mathsf{Adv}^{\mathsf{DDH}}(\mathcal{B})$.

IND-CPA is the declared security boundary of the CCE primitive. Because HomAdd and ReRand endow ciphertexts with structural malleability, a decryption oracle would let an adversary rerandomize (or homomorphically combine) a submitted ciphertext to bypass the challenge restriction and read the plaintext directly, so the primitive is *not* CCA1/CCA2-secure—matching homomorphic schemes such as Zether, PGC, and Paillier. Ciphertext integrity is instead provided at the system layer by transaction authentication and range proofs (§IV-A).

B. Instantiation: Dual-Generator Exponential ElGamal

The specification of §III-A is realized concretely as *dual-generator exponential ElGamal*: encryption outputs $(Y_1, Y_2) = (g^r,\ \mathsf{pk}^r \cdot h^m)$ over a prime-order group with secondary generator $h = \mathsf{Hash\text{-}to\text{-}Curve}(g)$, and decryption computes $h^m = Y_2 \cdot Y_1^{-\mathsf{sk}}$ followed by BSGS recovery over $[0, \mathsf{MAX})$ with $\mathsf{MAX} = 2^{32}$ ($O(2^{16})$ operations, millisecond-level); Paillier [23] lacks the group structure needed for efficient ZK integration. The decisive structural fact is that $Y_2 = \mathsf{pk}^r \cdot h^m = g^{\mathsf{sk}\cdot r} \cdot h^m$ is itself a Pedersen commitment under the fixed bases (g, h) with opening randomness $\rho = \mathsf{sk}\cdot r$ (condition C4), so range proofs attach to it directly with zero bridging; standard ElGamal instead requires a separate bridging protocol in ZK circuits. Setting $m = 0$ yields a zero encryption (g^r, pk^r) indistinguishable from $\mathsf{Enc}(m; r')$ under DDH. That this instantiation is not an arbitrary choice but the unique one (up to reparameterization) is the content of §III-C.

C. CCE Algebraic Characterization

We establish that CCE is not merely *sufficient* for private account-model transactions, but *algebraically necessary*: any scheme satisfying three independently motivated properties is provably equivalent to dual-generator exponential ElGamal up to affine reparameterization.

Definition III.4 (Structural Properties of Account-Model Privacy Schemes)**.** *An account-model privacy scheme Π over a prime-order group $\mathbb{G}$ must satisfy:*

- ***(S1) Balance Confidentiality.*** *Transaction amounts and account balances are computationally indistinguishable from uniformly random to any PPT adversary without the view key. Formally, Π satisfies IND-CPA security (Definition III.3).*
- ***(S2) Homomorphic Balance Conservation.*** *Ciphertexts support component-wise additive homomorphism: there exists an efficiently computable operation $\oplus$ such that $\mathsf{Enc}(m_1; r_1) \oplus \mathsf{Enc}(m_2; r_2) = \mathsf{Enc}(m_1 + m_2; r_1 + r_2)$, enabling zero-knowledge balance verification ($\sum$ inputs $=$ $\sum$ outputs) without revealing amounts.*
- ***(S3) Representation Unity.*** *The ciphertext component that enables homomorphic balance verification (S2) is the* same algebraic object *from which the plaintext is deterministically recovered via $\mathsf{Dec}(\mathsf{sk}, \cdot)$. That is, no separate "commitment" and "encryption" objects coexist for the same plaintext; a single group element C_2 serves both functions simultaneously.*

Remark on S3. S3 does *not* assume that C_2 is a Pedersen commitment, only that balance verification and plaintext recovery share a component; that this component *necessarily* has Pedersen structure is the content of Theorem III.5.

Theorem III.5 (CCE Algebraic Characterization)**.** *Let Π be an account-model privacy scheme over a prime-order group $\mathbb{G} = \langle g \rangle$ of order q with public parameters $\mathsf{pp} = (\mathbb{G}, q, g, h)$, where $\log_g h$ is unknown. If Π satisfies S1–S3 (Definition III.4), then, up to affine reparameterization of generators $(g \mapsto g' = g^a, h \mapsto h' = h^b)$ and message encoding $(m \mapsto b \cdot m)$, Π's encryption is algebraically equivalent to dual-generator exponential ElGamal:*

$$\mathsf{Enc}(\mathsf{pp}, \mathsf{pk}, m; r) = \big((g')^r,\ \mathsf{pk}'^{\,r} \cdot (h')^m\big), \quad \mathsf{pk}' = (g')^{\mathsf{sk}}. \tag{3}$$

In particular, C_2 is necessarily a Pedersen commitment $\mathsf{Com}(m; \rho(\mathsf{sk}, r))$ with $\rho(\mathsf{sk}, r) = \mathsf{sk}\cdot r$.

Proof sketch. The argument eliminates one degree of freedom per step: S2 forces any homomorphic ciphertext into the form

$(g_1^r h_1^m, \ldots, g_k^r h_k^m)$, S1 forces $k \geq 2$ (one component must mask r independently of m), and S3 pins the message base and randomness function, forcing $g_2 = g_1^{\mathsf{sk}} = \mathsf{pk}'$ and $h_2 = h$ and yielding $\mathsf{Enc} = (g_1^r, \mathsf{pk}'^r h^m)$ with $\rho(\mathsf{sk}, r) = \mathsf{sk} \cdot r$. Full proof in Supplementary Appendix E.1. ■

Remark. PGC's Twisted ElGamal [10], $C = (\mathsf{pk}^r,\ g^r \cdot h^m)$, belongs to the same equivalence class via $\mathsf{sk}' = \mathsf{sk}^{-1}$, $r' = \mathsf{sk} \cdot r$; its commitment component is therefore likewise directly consumable by range proofs. The separation from HERMES is operational rather than algebraic (§IV-C).

D. Lower Bounds and Optimality

We establish a group-element lower bound that HERMES matches exactly, and a scalar-multiplication lower bound in the Generic Group Model (GGM) that HERMES matches up to fixed-base precomputation.

Theorem III.6 (CCE Group-Element Lower Bound)**.** *Any CCE-coherent encryption scheme over a prime-order group $\mathbb{G}$ of order q requires at least 2 group elements per ciphertext: one to hide the randomness (indistinguishability under S1), one to bind the message (computational binding under S3).*

Proof. With a single element $C = g^{f(r,m)}$, IND-CPA (S1) forces $r \mapsto f(r, m)$ injective for each fixed m, while binding (S3) forces $m \mapsto f(r, m)$ computationally injective for each fixed r; but one element has only q values for $q \cdot \mathsf{MAX}$ inputs, so pigeonhole collisions violate one injectivity. Hence $k \geq 2$ (Supplementary Appendix E.2). ■

Theorem III.7 (CCE Scalar-Multiplication Lower Bound in GGM)**.** *In the Generic Group Model, any CCE-coherent encryption scheme requires at least 2 scalar multiplications per ciphertext.*

Proof. In the GGM, producing g^x from known generators requires one scalar multiplication; by the group-element lower bound $k \geq 2$, and C_2's dependency on m is information-theoretically invisible in C_1 (IND-CPA), so a second multiplication is unavoidable (Supplementary Appendix E.3). ■

Corollary III.8 (HERMES Cost versus the Lower Bounds)**.** *HERMES matches the group-element bound with exactly 2 group elements per ciphertext, and encrypts with three scalar multiplications (g^r, pk^r, h^m): two are fixed-base and admit precomputation, leaving a single variable-base multiplication (pk^r) as in standard ElGamal.*

The practical value of CCE lies in its circuit-level implications. We formalize the *bridging overhead*: the additional constraints required when encryption and commitment are separate objects versus when unified by CCE coherence.

Definition III.9 (Bridging Overhead)**.** *Let* $\mathrm{cost}(\mathcal{C})$ *denote the number of PLONKish constraints in circuit $\mathcal{C}$. The* bridging overhead *of a scheme is:*

$$\Delta_{\mathsf{bridge}} = \mathrm{cost}(\mathcal{C}_{\mathsf{separate}}) - \mathrm{cost}(\mathcal{C}_{\mathsf{unified}}), \tag{4}$$

where $\mathcal{C}_{\mathsf{separate}}$ verifies ciphertext validity, range constraints, and balance conservation using independent *encryption and commitment objects, and $\mathcal{C}_{\mathsf{unified}}$ uses a single CCE-coherent ciphertext component.*

Theorem III.10 (Zero-Bridging: CCE Optimality)**.** *Let $\mathcal{C}_{\mathsf{cce}}$ be the PLONKish circuit for a CCE-coherent scheme and $\mathcal{C}_{\mathsf{non\text{-}cce}}$ be the corresponding circuit for a non-CCE scheme. Then:*

$$\mathrm{cost}(\mathcal{C}_{\mathsf{non\text{-}cce}}) - \mathrm{cost}(\mathcal{C}_{\mathsf{cce}}) \ \geq\ N \cdot \big(\mathrm{cost}_{\mathsf{Com}} + \mathrm{cost}_{\mathsf{eq}}\big), \tag{5}$$

where $\mathrm{cost}_{\mathsf{Com}} = 2\lceil \log_2 q \rceil$ is the constraint cost of one Pedersen commitment computation and $\mathrm{cost}_{\mathsf{eq}} = 1$ is the cost of one copy constraint linking the plaintext witness.

Intuition. In a non-CCE scheme, each anonymity-set member requires an independent commitment computation ($2\lceil \log_2 q \rceil$ constraints) plus a copy constraint to link the plaintext across encryption and commitment sub-circuits; CCE eliminates both since C_2 *is* the commitment. For HERMES (N=8, BN254), $\Delta_{\mathsf{bridge}} \geq 4{,}072$ constraints; the complete proof appears in §V-B3.

E. Composition and Simulation Extractability

The single-ciphertext IND-CPA guarantee extends to HERMES's many-ciphertext homomorphic composition: in the adaptive CCE-Composition game (N challenge queries with *independent* randomizers), the adversary's advantage is at most $N \cdot \mathsf{Adv}^{\mathsf{DDH}}$ via a standard hybrid argument—a shared randomizer breaks the argument, since the ratio of two commitment components collapses to $h^{\Delta m}$ (statement and proof in Supplementary Appendix B). CCE ciphertexts also appear as witnesses inside ZK proofs: under (SE1) deterministic group encoding, (SE2) an independent randomness witness, and (SE3) a binding range proof, the CCE-encryption relation in PLONKish arithmetization satisfies simulation extractability [24] (proof in Supplementary Appendix C), so any valid forged transaction proof must "know" the underlying balances and randomizers.

F. Failure Modes and the Three-Layer Relationship

CCE holds precisely when *the algebraic structure of a ciphertext component equals that of the intended commitment scheme* and *the ciphertexts use mutually independent randomizers*. Standard exponential ElGamal merges m and r on the same base and cannot separate them; Zether's second component carries no message; Paillier leaves the discrete-log setting—all three break extractability; Supplementary Appendix E.5 classifies these failure modes in detail.

CCE operates on three distinct layers, which clarifies the novelty claim: the **CCE primitive** (this section), its **instantiation** (§III-B), and its **application** (§IV).

IV. The HERMES Scheme

This section presents the HERMES scheme. Following the CCE framework of §III, we instantiate all components via dual-generator exponential ElGamal (Theorem III.5 establishes this is the unique CCE-coherent choice). We describe the architecture, threat model, key management, transaction protocols, and ZK circuit instantiation.

A. Architecture Overview and Threat Model

HERMES is designed around four objectives: full-ciphertext balance privacy; sender–receiver anonymity via anonymity sets and ZK proofs; minimal trust (Halo2's transparent IPA backend needs no trusted setup); and fully on-chain closure,

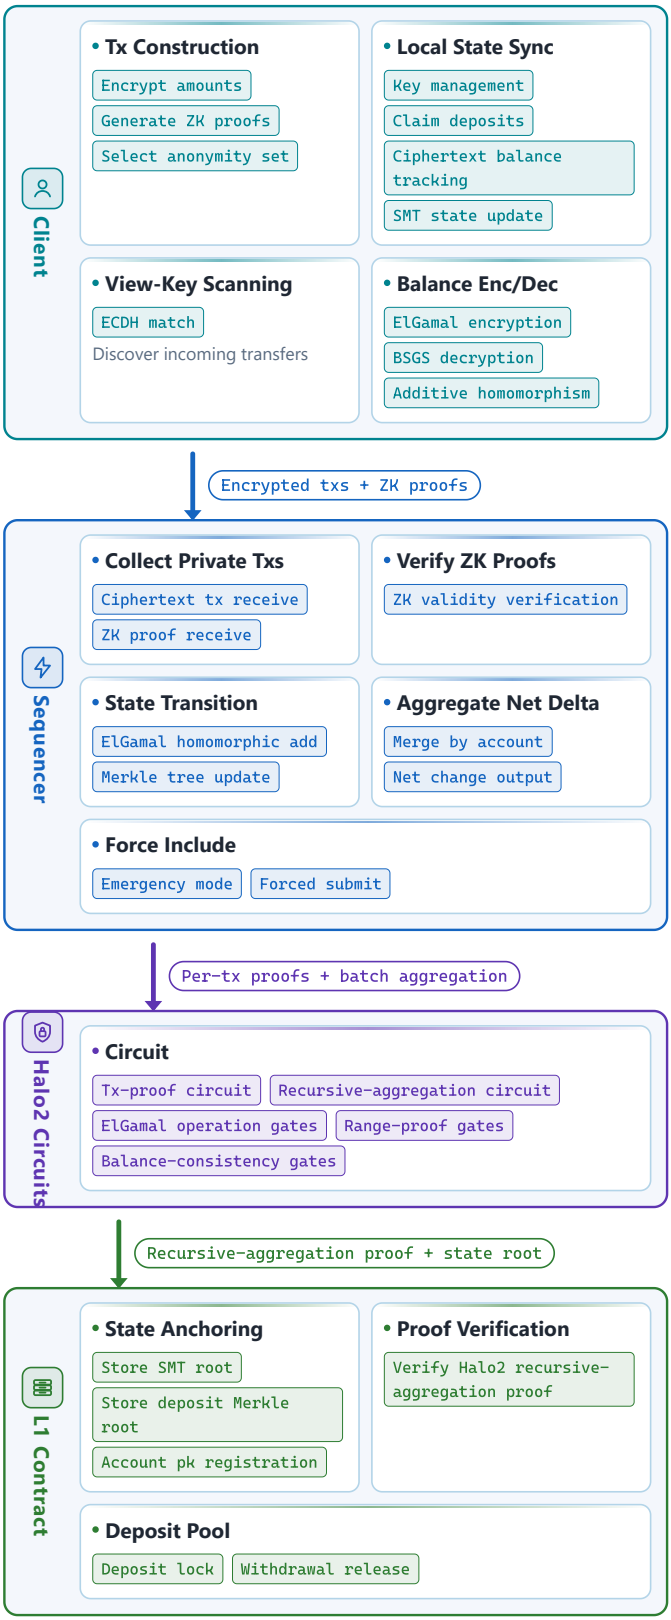

Fig. 1. HERMES architecture: the off-chain Client and Sequencer, the Halo2 proof circuits, and the L1 contract, with the data flow and trust boundaries among them.

letting recipients discover incoming transfers from on-chain data alone.

Layer 1 / Layer 2 Architecture. HERMES moves privacy computation and proof generation off-chain, with L1 serving solely for verification and state anchoring. The off-chain side comprises the Client (key management and transaction construction), the Sequencer (proof verification and batch processing), and the Halo2 proof circuits (transfer and recursive-aggregation circuits); the on-chain side is the L1 Contract (state anchoring and proof verification). Figure 1 summarizes these components and their trust boundaries.

Threat Model. HERMES considers three adversary levels: *honest-but-curious* adversaries who follow protocol rules but infer private information from observed data (defended by ElGamal semantic security under DDH); *malicious* adversaries who forge proofs or launch double-spend attacks (prevented by ZK knowledge soundness); and *adaptive-corruption* adversaries who corrupt arbitrary users at any time (remaining users protected by anonymity sets and independent encryption).

a) Adversary capabilities.: $\mathcal{A}$ has on-chain data access, can intercept protocol messages (subject to Sequencer honesty), adaptively corrupt users (obtaining $(\mathsf{sk}_a^{(i)}, \mathsf{sk}_v^{(i)})$), and attempts proof forgery. We exclude L1 consensus compromise, unbounded computation, and side channels.

b) Corruption models.:

Definition IV.1 (Static vs. Adaptive Corruption)**.** *In* static corruption*, $\mathcal{A}$ selects corruption targets $\mathcal{C} \subset \{1, \ldots, N\}$ before protocol execution. In* adaptive corruption*, $\mathcal{A}$ may corrupt any user i at any point, with decisions depending on all information observed so far.*

c) Security boundaries.: With N users and t corrupted, uncorrupted users' ciphertexts stay semantically secure and the remaining $N-t$ candidates form an indistinguishable set; Sequencer corruption causes only a system-level, batch-metadata leak (§V-E2).

B. Key Architecture and Address Derivation

The key architecture employs triple separation:

- **Amount key** $(\mathsf{sk}_a, \mathsf{pk}_a = g^{\mathsf{sk}_a})$: encrypts/decrypts balance and delta values.
- **View key** $(\mathsf{sk}_v, \mathsf{pk}_v = g^{\mathsf{sk}_v})$: enables scanning for incoming transfers; no signing capability.
- **Address** $\mathsf{addr} = \mathsf{Poseidon}(\mathsf{pk}_a, \mathsf{pk}_v)$: one-way image of both public keys.

All three are computationally independent under DLP: $\mathsf{sk}_a, \mathsf{sk}_v \xleftarrow{\$} \mathbb{Z}_q$ are sampled independently; the address inherits Poseidon's one-wayness. Figure 2 depicts the derivation graph.

Definition IV.2 (Compromise Radius)**.** *The* compromise radius *of key component k is*

$$\mathsf{CR}(k) = \{P \in \mathcal{P} \mid k \textit{ violates } P\},$$
$$\mathcal{P} = \{P_{\mathsf{bal}}, P_{\mathsf{disc}}, P_{\mathsf{addr}}\}.$$

Proposition IV.3 (Compromise Independence)**.** *Under DDH and DLP:* $\mathsf{CR}(\mathsf{sk}_a) = \{\textit{Balance Privacy}\}$, $\mathsf{CR}(\mathsf{sk}_v) = \{\textit{Autonomous Discovery}\}$, $\mathsf{CR}(\mathsf{addr}) = \{\textit{Address Unlinkability}\}$*, and compromising k_i does not reduce the security of k_j for $k_i \neq k_j$.*

The proof appears in Supplementary Appendix D. HERMES's key separation provides *strict fault isolation*: compro-

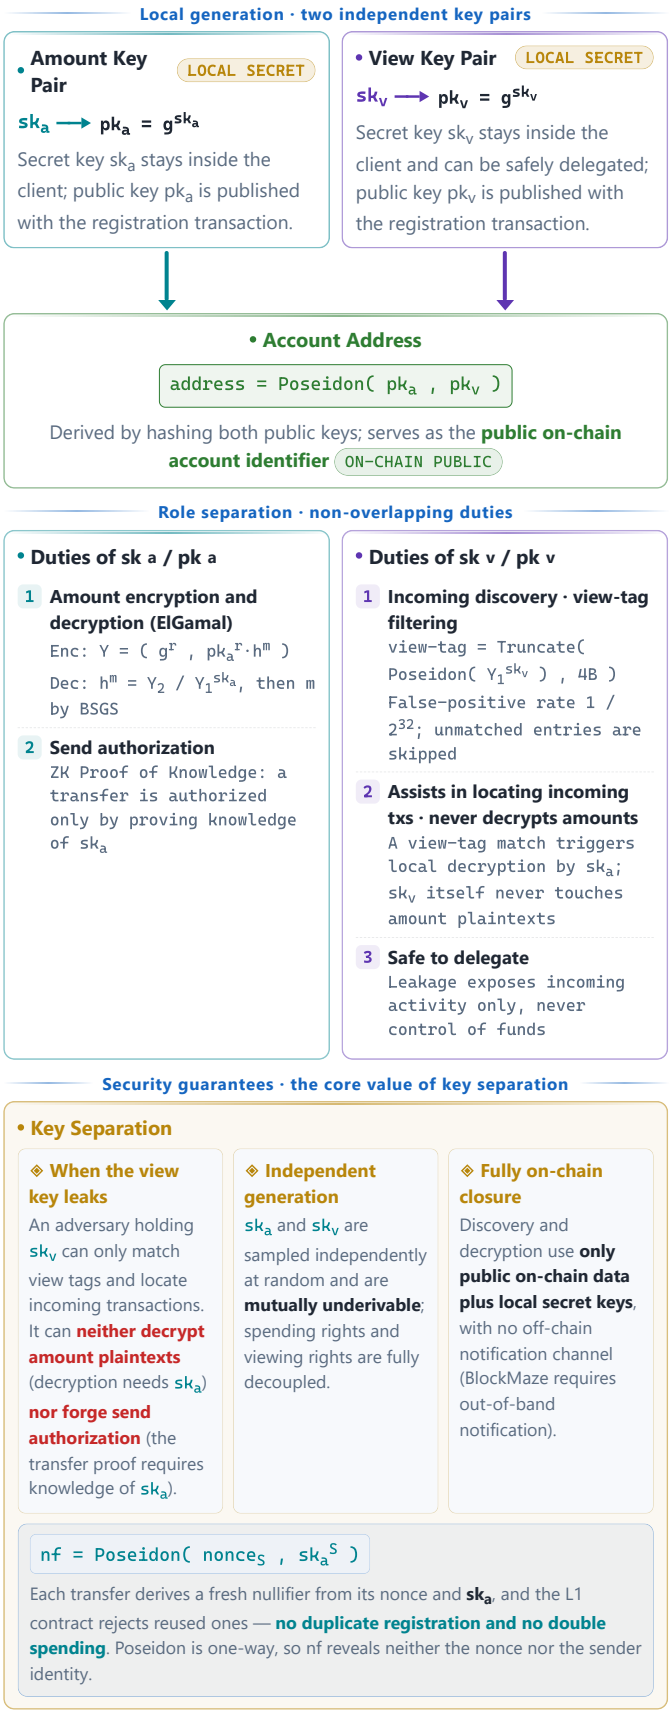

Fig. 2. Key architecture: amount key, view key, and one-way address derivation, with the compromise radius of each component.

mising any single key reduces neither the other keys' security nor the guarantees of uncompromised peers.

Independent Randomizers. If a construction reuses the same encryption randomizer r across sender and receiver delta ciphertexts encoding $\mp t$, the ratio $Y_{2,D}/Y_{2,S}$ collapses to h^{2t}, enabling direct amount recovery. HERMES assigns *independent* randomizers per ciphertext: the term $\mathsf{pk}_a^{r_S - r_D}$ constitutes an unremovable DDH mask.

C. Dual-Generator Exponential ElGamal Ciphertext Balance Model

HERMES adopts the dual-generator exponential ElGamal instantiation of §III-B: each ciphertext $(Y_1, Y_2) = (g^r,\ \mathsf{pk}_a^r \cdot h^m)$ provides complete encryption/decryption, and CCE coherence makes Y_2 a Pedersen commitment that range proofs use directly without bridging conversion. Each ciphertext k uses a unique randomizer $r_k \xleftarrow{\$} \mathbb{Z}_q$: for $r_1 \neq r_2$, the ratio $Y_{2,1}/Y_{2,2} = \mathsf{pk}_a^{r_1 - r_2} \cdot h^{m_1 - m_2}$ retains the unpredictable mask $\mathsf{pk}_a^{r_1 - r_2}$, indistinguishable from random under DDH.

View Key Scanning. Recipients compute $\mathsf{view\text{-}tag} = \mathsf{Truncate}(\mathsf{Poseidon}(Y_1^{\mathsf{sk}_v}), 4\text{ bytes})$ for fast filtering; matching tags trigger full BSGS decryption ($m < 2^{32}$, $O(2^{16})$ operations). The 32-bit truncation yields false-positive rate $\approx 2^{-32}$ per non-matching ciphertext, and false positives incur only extra decryption cost. View tags are an established wallet-scanning mechanism (e.g., Zcash Orchard [25]) that HERMES uses as a building block, eliminating the off-chain notification dependency of BlockMaze-style schemes.

D. Transaction Protocols

HERMES defines three transaction types covering the complete lifecycle. Transfer and Withdraw hide the amount in an encrypted delta under an anonymity set of size N with nullifier $\mathsf{Poseidon}(\mathsf{nonce}_S, \mathsf{sk}_a^S)$ (Withdraw alone reveals the public amount t at the L1 payout), while AccountOpen encrypts the deposit in C_0 under a one-way-hash address.

Transfer. On input $(S, R, t, \mathcal{A}_N)$, the sender assembles the anonymity set $\mathcal{A} = \{S\} \cup \text{decoys} \cup \{R\}$, samples independent $r_k \xleftarrow{\$} \mathbb{Z}_q$ for all k, sets $\delta_S = -t$, $\delta_R = +t$, $\delta_k = 0$ with $C_k = (g^{r_k}, \mathsf{pk}_{a,k}^{r_k} h^{\delta_k})$, computes $\mathsf{view\text{-}tag}_k = \mathsf{Trunc}(\mathsf{Poseidon}(\mathsf{pk}_{v,k}^{r_k}), 4\mathrm{B})$ and $\mathsf{nf} = \mathsf{Poseidon}(\mathsf{nonce}_S, \mathsf{sk}_a^S)$, proves in π the key relations, $\sum \delta_k = 0$, $0 \le t < 2^{32}$, $\mathsf{bal}_S \ge t$, membership, and the state successor, and emits $\mathsf{tx} = (\mathcal{A}, \{C_k\}, \{\mathsf{view\text{-}tag}_k\}, \pi, \mathsf{root}_{\mathrm{old}}, \mathsf{nf})$.

Transfer analysis. Balance conservation holds: $\sum_k \delta_k = (-t) + t + 0 = 0$. The nullifier $\mathsf{nf} = \mathsf{Poseidon}(\mathsf{nonce}_S, \mathsf{sk}_a^S)$ prevents double-spending without revealing nonce_S: since Poseidon is a random oracle and sk_a^S is unknown, nf reveals neither nonce_S nor the sender's identity.

AccountOpen. On deposit of amount, the user generates $(\mathsf{sk}_a, \mathsf{pk}_a)$ and $(\mathsf{sk}_v, \mathsf{pk}_v)$, derives $\mathsf{addr} = \mathsf{Poseidon}(\mathsf{pk}_a, \mathsf{pk}_v)$, samples $\mathsf{nullifier}, \mathsf{secret}, r_0 \xleftarrow{\$} \mathbb{Z}_q$, sets $\mathsf{comm} = \mathsf{Poseidon}(\mathsf{nullifier}, \mathsf{secret}, \mathsf{amount})$, $C_0 = \mathsf{Enc}(\mathsf{amount}; r_0)$, and $\mathsf{view\text{-}tag}_0 = \mathsf{Trunc}(\mathsf{Poseidon}(\mathsf{pk}_v^{r_0}), 4\mathrm{B})$, proves knowledge of the opening, correct comm, consistency, and $0 \le \mathsf{amount} < 2^{32}$, and emits $\mathsf{tx} = (\mathsf{addr}, \mathsf{comm}, C_0, \mathsf{view\text{-}tag}_0, \pi)$.

Withdraw. On input $(t, \mathsf{addr}_{\mathrm{L1}})$, the sender assembles $N-1$ decoys with independent r_k, sets $\delta_S = -t$ and $C_k = (g^{r_k}, \mathsf{pk}_{a,k}^{r_k} h^{\delta_k})$, computes $\mathsf{nf} = \mathsf{Poseidon}(\mathsf{nonce}_S, \mathsf{sk}_a^S)$,

proves that t matches the delta, $\mathsf{bal}_S - t \geq 0$, membership, and the state successor, and emits $\mathsf{tx} = (\mathcal{A}, \{C_k\}, \{\mathsf{view\text{-}tag}_k\}, \pi, \mathsf{root}_{\mathrm{old}}, \mathsf{nf}, \mathsf{addr}_{\mathrm{L1}}, t)$.

E. ZK Argument Instantiation

HERMES instantiates its ZK arguments in Halo2 [17] over BN254, using PLONKish arithmetization. Per transaction, the proof system establishes five circuit requirements:

- **R1 Balance Validity**: encrypted balance is a well-formed dual-generator exponential ElGamal ciphertext.
- **R2 Transaction Correctness**: delta ciphertext encodes signed amount difference; sender has sufficient balance.
- **R3 Range Constraints**: $0 \leq t < 2^{32}$.
- **R4 State Membership**: sender and all anonymity-set members exist in the current state SMT.
- **R5 Anonymity-Set Indistinguishability**: true sender hidden within set of size N via One-out-of-Many proofs [26], [27].

The CCE property lets (R1) and (R3) share the same commitment component Y_2, eliminating the bridging Σ-protocol. Requirements R1–R5 are proof-system-independent, realized via a knowledge-of-balance proof, a range proof on t, a One-out-of-Many membership proof, and a view-tag derivation proof.

V. Security Analysis

This section establishes HERMES's security guarantees under the model of §IV-A: the three security games (§V-A), the CCE structural properties (§V-B), the system-level theorems (§V-C), and adaptive adversaries with concrete bounds (§V-D–§V-E).

A. Security Model and Game Definitions

Corruption oracle. Each game Game_i ($i \in \{1,2,3\}$) is equipped with a corruption oracle $\mathcal{O}^{\mathsf{corrupt}}$ (§IV-A): $\mathcal{A}$ may query $\mathcal{O}^{\mathsf{corrupt}}(j)$ to obtain $(\mathsf{sk}_a^{(j)}, \mathsf{sk}_v^{(j)})$. *Restriction:* the oracle must not be queried on challenge users. If $\mathcal{A}$ does so, the game aborts. Corruption reveals only $(\mathsf{sk}_a, \mathsf{sk}_v)$; it does *not* reveal the nonce counter nonce_S.

1) Game_1: Balance Privacy

Game Game_1 (Balance Privacy). $\mathcal{C}$ runs $\mathsf{Setup}(\lambda)$ and generates $(\mathsf{sk}_a, \mathsf{pk}_a)$ with $\mathsf{pk}_a = g^{\mathsf{sk}_a}$. $\mathcal{A}$ selects $b_0, b_1 \in [0, \mathsf{MAX})$; $\mathcal{C}$ samples $\beta \xleftarrow{\$} \{0,1\}$ and $r \xleftarrow{\$} \mathbb{Z}_q$, and sends $C^* = \mathsf{Enc}(b_\beta; r) = (g^r, \mathsf{pk}_a^r \cdot h^{b_\beta})$ to $\mathcal{A}$. $\mathcal{A}$ outputs β', and $\mathsf{Adv}_1(\mathcal{A}) = \left|\Pr[\beta' = \beta] - \frac{1}{2}\right|$.

2) Game_2: Transaction Amount Indistinguishability

Game Game_2 (Transaction Amount Indistinguishability). $\mathcal{C}$ generates two key pairs $(\mathsf{sk}_a^S, \mathsf{pk}_a^S)$ and $(\mathsf{sk}_a^D, \mathsf{pk}_a^D)$. $\mathcal{A}$ selects $v_0, v_1 \in [0, \mathsf{MAX})$; $\mathcal{C}$ samples $\beta \xleftarrow{\$} \{0,1\}$ and *independent* randomizers $r_S, r_D \xleftarrow{\$} \mathbb{Z}_q$, computes $C_S^* = \mathsf{Enc}(v_\beta; r_S) = (g^{r_S}, \mathsf{pk}_a^{S r_S} \cdot h^{v_\beta})$ and $C_D^* = \mathsf{Enc}(v_\beta; r_D) = (g^{r_D}, \mathsf{pk}_a^{D r_D} \cdot h^{v_\beta})$. $\mathcal{A}$ outputs β', and $\mathsf{Adv}_2(\mathcal{A}) = \left|\Pr[\beta' = \beta] - \frac{1}{2}\right|$.

3) Game_3: Sender-Receiver Unlinkability

Game Game_3 (Sender-Receiver Unlinkability). $\mathcal{C}$ generates keys for N users, selects sender i^* and receiver j^* ($i^* \neq j^*$), and constructs a privacy transfer embedded in an anonymity set of N keys with independent randomizers. $\mathcal{A}$ receives the batch output (balance ciphertexts, delta ciphertexts, view tags $\{\mathsf{view\text{-}tag}_k\}$, nullifier nf, state root, proof π) and outputs (i', j'), with $\mathsf{Adv}_3(\mathcal{A}) = \left|\Pr[(i', j') = (i^*, j^*)] - \frac{1}{N(N-1)}\right|$.

B. CCE Property Proofs

This section proves the four CCE properties stated in §III. As a standalone primitive whose ciphertext doubles as a key-bound commitment, CCE's security is established by the four self-contained results below (correctness, perfect hiding, computational binding) and the IND-CPA reduction that follows.

1) Correctness and Commitment Coherence (C1–C4)

Theorem V.1 (CCE Correctness)**.** *The CCE construction satisfies conditions C1–C4: (a)* decryption $\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, m; r)) = m$*; (b)* homomorphic $\mathsf{Dec}\big(\mathsf{sk}, \mathsf{HomAdd}(\mathsf{Enc}(\mathsf{pk}, m_1; r_1), \mathsf{Enc}(\mathsf{pk}, m_2; r_2))\big) = m_1 \oplus m_2$*; (c)* rerandomization $\mathsf{Dec}(\mathsf{sk}, \mathsf{ReRand}(\mathsf{Enc}(\mathsf{pk}, m; r); s)) = m$ *and* ReRand *preserves the ciphertext distribution; (d)* commitment coherence $\mathsf{Com}\big(\mathsf{Enc}(\mathsf{pk}, m; r)\big)$ *is the Pedersen commitment* $\mathsf{PedCom}(m; \mathsf{sk} \cdot r)$.

Proof. All four conditions follow by direct verification (full derivations in Supplementary Appendix A): (a) $\mathsf{pk}^r h^m / (\mathsf{pk})^r = h^m$ decodes to m; (b) the component-wise product gives $\mathsf{Enc}(\mathsf{pk}, m_1{+}m_2; r_1{+}r_2)$; (c) ReRand yields $\mathsf{Enc}(\mathsf{pk}, m; r{+}s)$; (d) $\mathsf{Com}(c) = \mathsf{pk}^r h^m = \mathsf{PedCom}(m; \mathsf{sk}{\cdot}r)$, matching Definition III.1. ∎

Theorem V.2 (Perfect Hiding)**.** *For the CCE construction, the commitment component* $\mathsf{Com}(\mathsf{Enc}(\mathsf{pk}, m; r))$ *is* perfectly *(information-theoretically) hiding: for every two messages m_0, m_1 and every ciphertext (without opening),* $\mathsf{Com}(\mathsf{Enc}(\mathsf{pk}, m_0; r_0))$ *and* $\mathsf{Com}(\mathsf{Enc}(\mathsf{pk}, m_1; r_1))$ *are identically distributed over the uniform randomizers r_0, r_1.*

Proof. For uniform r_j, $r_j \mapsto \mathsf{pk}^{r_j} = g^{\mathsf{sk} r_j}$ is a bijection on $\mathbb{G}$ ($\mathsf{sk} \not\equiv 0$), so $\mathsf{pk}^{r_j} h^{m_j}$ is uniform and message-independent: $\Pr[\mathsf{Com}(c_{m_0}) = X] = \Pr[\mathsf{Com}(c_{m_1}) = X] = 1/p$ for every X, with no computational assumption (full proof in Appendix A). ∎

Theorem V.3 (Computational Binding)**.** *Under DL hardness in $\mathbb{G}$, the CCE commitment is computationally binding: no PPT adversary outputs two decodable ciphertexts and their openings $(c_0, m_0, \rho_0), (c_1, m_1, \rho_1)$ with $\mathsf{Com}(c_0) = \mathsf{Com}(c_1)$ and $(m_0, \rho_0) \neq (m_1, \rho_1)$ except with negligible probability.*

Proof. A collision $\mathsf{pk}^{\rho_0} h^{m_0} = \mathsf{pk}^{\rho_1} h^{m_1}$ with distinct openings solves $\alpha = \log_g h$ via $\mathsf{sk}(\rho_0 - \rho_1) = \alpha(m_1 - m_0)$ when $m_1 \not\equiv m_0$, or forces $\rho_0 = \rho_1$ when $m_1 = m_0$—both contradicting distinct openings; requiring the adversary to output the opening (m, ρ) makes the reduction computable (Appendix A). ∎

2) CCE-IND-CPA Security

Theorem V.4 (CCE-IND-CPA $\leq$ DDH: Tight Reduction)**.** *Assuming DDH is hard in $\mathbb{G}$, for any PPT adversary $\mathcal{A}$ against* $\mathsf{Game}_{\mathsf{CCE}}$*, there exists a PPT DDH distinguisher $\mathcal{B}$ such that* $\mathsf{Adv}^{\mathsf{CCE\text{-}IND\text{-}CPA}}(\mathcal{A}) = \mathsf{Adv}^{\mathsf{DDH}}(\mathcal{B})$*. The reduction is* tight*: no multiplicative loss.*

Proof. $\mathcal{B}$ receives $(g, A{=}g^a, B{=}g^b, T)$ with $T = g^{ab}$ or g^c, sets $\mathsf{pk} = A$, and constructs $C_1^* = B$, $C_2^* = T \cdot h^{m_\beta}$. If $T =$

g^{ab}, then $C_2^* = \mathsf{pk}^b \cdot h^{m_\beta} = \mathsf{Com}(m_\beta; a \cdot b)$, distributed identically to $\mathsf{Game}_{\mathsf{CCE}}$, so $\mathcal{A}$ keeps advantage $\mathsf{Adv}^{\mathsf{CCE\text{-}IND\text{-}CPA}}(\mathcal{A})$. If $T = g^c$ for uniform c, then $C_2^* = g^{c+\alpha m_\beta}$ ($\alpha = \log_g h$) is uniform regardless of β, so $\mathcal{A}$ has advantage 0. Combining, $\mathsf{Adv}^{\mathsf{DDH}}(\mathcal{B}) = \mathsf{Adv}^{\mathsf{CCE\text{-}IND\text{-}CPA}}(\mathcal{A})$. ■

Corollary V.5 (Commitment Structure Does Not Degrade Security)**.** *The CCE coherence property introduces no additional attack surface beyond standard ElGamal IND-CPA.*

3) Zero-Bridging, Classification, and Leakage Characterization

We prove three structural results stated in §III: the zero-bridging optimality theorem, the algebraic classification/uniqueness theorem, and the leakage-closure uniqueness theorem. The first two have complete proofs in Supplementary Appendix E (E.4 and E.1, respectively); we state the third here.

Theorem V.6 (Leakage-Closure Uniqueness)**.** *Within the DDH framework, the only strategy preventing all three algebraic-leakage surfaces from exposing plaintext information is* per-ciphertext independent randomization*: every C_k uses a fresh $r_k \xleftarrow{\$} \mathbb{Z}_q$ not reused across distinct plaintexts.*

Proof of Theorem V.6. The three leakage surfaces are: (1) *balance ciphertexts*—distinct keys give computational independence by DDH regardless of randomization; (2) *same-account delta stream*—shared r makes $Y_{2,a}/Y_{2,b} = h^{\delta_a - \delta_b}$ recoverable by lookup, so any reused randomizer yields such a collision (independence is *necessary*), while independent r_k keeps $\mathsf{pk}_a^{r_a - r_b}$ uniform and the ratio indistinguishable from random by DDH (independence is *sufficient*); (3) *sender/receiver paired deltas*—the $\delta = -t$ special case of surface (2). Hence per-ciphertext independent randomization is both necessary and sufficient for full leakage closure under DDH. ■

These three results together establish that: CCE coherence is necessary and sufficient for zero-bridge circuits, dual-generator ElGamal is the unique CCE-coherent scheme, and independent randomization is the unique leakage-closure strategy. The compositional IND-CPA bound and the conditional simulation extractability of the CCE-encryption relation (§III-E; proofs in Supplementary Appendices B and C) complete the CCE property layer.

C. System Security Proofs

1) Game_1 Balance Privacy

Theorem V.7 (Balance Privacy)**.** *For any PPT $\mathcal{A}$ not corrupting the challenge user:* $\mathsf{Adv}_1(\mathcal{A}) \leq \mathsf{Adv}_{\mathbb{G}_1}^{DDH}$.

Proof sketch. In the DDH reduction, $\mathcal{B}$ embeds the challenge into $C^* = (g^b, T \cdot h^{b_\beta})$; the real/twin-DDH cases give advantage Adv_1 and 0, yielding a *tight* bound. Full proof: Supplementary Appendix F.1. ■

2) Game_2 Amount Indistinguishability

Theorem V.8 (Transaction Amount Indistinguishability)**.** *For any PPT $\mathcal{A}$ not corrupting challenge users S or D:* $\mathsf{Adv}_2(\mathcal{A}) \leq 2 \cdot \mathsf{Adv}_{\mathbb{G}_1}^{DDH}$.

Proof sketch. A two-step hybrid replaces C_S^*, then C_D^*, with encryptions of zero; each step costs $\leq \mathsf{Adv}^{\mathrm{DDH}}$, giving $\mathsf{Adv}_2 \leq 2\mathsf{Adv}^{\mathrm{DDH}}$ (at most one bit lost). This captures the *ratio attack*: PGC commits under the fixed base g and reuses one r, so dividing the two components cancels g^r and leaks h^{2t}. HERMES blocks it via the account-key base pk_a and independent r_k, keeping $\mathsf{pk}_a^{r_D - r_S}$ as an unremovable DDH mask. Full proof: Supplementary Appendix F.2. ■

3) Game_3 Unlinkability

Theorem V.9 (Sender-Receiver Unlinkability)**.** *For any PPT $\mathcal{A}$ not corrupting challenge users i^* or j^*:*

$$\mathsf{Adv}_3(\mathcal{A}) \leq \mathsf{Adv}_{S_{\mathsf{ZK}}}^{\mathsf{ZK}} + 2N \cdot \mathsf{Adv}_{\mathbb{G}_1}^{DDH}. \tag{6}$$

Proof sketch. A three-step hybrid: $H_{0.5}$ simulates the proof (cost $\mathsf{Adv}^{\mathsf{ZK}}$); H_1 replaces all N balance ciphertexts with $\mathsf{Enc}(0;\cdot)$ (cost $N \cdot \mathsf{Adv}^{\mathrm{DDH}}$); H_2 masks each delta ciphertext with $\mathsf{Enc}(0;\rho_k)$ under fresh randomness (cost $N \cdot \mathsf{Adv}^{\mathrm{DDH}}$). In H_2 every observation is independent of (i^*, j^*), so $\mathcal{A}$ guesses at $1/(N(N-1))$; view tags and the one-way nullifier image do not weaken this. Total $\mathsf{Adv}_3 \leq \mathsf{Adv}^{\mathsf{ZK}} + 2N\mathsf{Adv}^{\mathrm{DDH}}$. Full proof: Supplementary Appendix F.3. ■

D. Adaptive Adversary and Concrete Security Bounds

a) Adaptive corruption.:

Lemma V.10 (Static-to-Adaptive Reduction)**.** *For any PPT adaptive adversary $\mathcal{A}_{\mathrm{adapt}}$ corrupting at most $t < N$ users:*

$$\mathsf{Adv}_{i,\mathrm{adapt}}(\mathcal{A}_{\mathrm{adapt}}) \leq \binom{N}{t} \cdot \mathsf{Adv}_{i,\mathrm{stat}}(\mathcal{A}_{\mathrm{stat}}). \tag{7}$$

When t is constant (HERMES's regime, $t \leq N-2$ for $N = 8$), $\mathsf{Adv}_{i,\mathrm{adapt}}$ remains negligible.

Proof Sketch. The static reduction guesses the corruption subset $\mathcal{C}^\star$ of size t (correct with probability $1/\binom{N}{t}$). Conditioned on a correct guess, the view is identical to the real game. Full proof: Supplementary Appendix F.4. ■

b) Concrete security bounds.: On BN254 ($\lambda_{\mathsf{DDH}} \approx 100$), balance privacy retains 100 bits ($\mathsf{Adv}_1 \leq \mathsf{Adv}^{\mathrm{DDH}}$), amount indistinguishability 99 bits ($\mathsf{Adv}_2 \leq 2\mathsf{Adv}^{\mathrm{DDH}}$), and unlinkability 96 bits at N=8 and 93 bits at N=64 ($\mathsf{Adv}_3 \leq 2N\mathsf{Adv}^{\mathrm{DDH}} + \mathsf{Adv}^{\mathsf{ZK}}$, with $\mathsf{Adv}^{\mathsf{ZK}}$ negligible). The first two properties are essentially tight; unlinkability degrades linearly with N. Even at $q \leq 2^{40}$ queries, the Game_3 advantage stays $\leq 2^{40} \cdot 2^{-96} = 2^{-56} \ll 1$—a large concrete margin.

E. Security Discussion

1) Cryptographic Assumptions and Limitations

HERMES depends on: (1) **DDH** (all three reductions); (2) **DLP** (key security); (3) **hash collision resistance** (Poseidon for addresses/commitments); (4) **generator independence** ($\log_g h$ unknown; h derived via hash-to-curve). Deployment-level limitations (N=8 default, single-Sequencer liveness) are discussed in §VII.

2) Sequencer Near-Full-Privacy

The Sequencer processes only ciphertexts and proofs, never plaintext. A corrupted Sequencer leaks current-batch metadata (anonymity-set correlations, nullifier frequencies) but not amounts or directions; historical batches are unaffected.

TABLE I
HERMES CRYPTOGRAPHIC-LAYER BENCHMARKS ON BN254 (SINGLE CORE, MEAN OF 100 ROUNDS).

Metric	Mean	StdDev	Notes
Dual-Gen ElGamal Enc	320 μs	14 μs	$(g^r,\ \mathsf{pk}^r \cdot h^m)$
BSGS Decryption (2^{16} table)	1,500 μs	62 μs	$m < 2^{32}$
Homomorphic Addition	40 μs	2 μs	componentwise group mult.
View-Key Scanning	600 μs	25 μs	per ciphertext
Peak proving memory (IPA)	$\approx$ 976 MiB	—	Halo2 IPA
BSGS table memory (2^{16})	$\approx$ 2 MiB	—	built once

3) Security Comparison with Existing Schemes

The property–mechanism correspondence is direct: balance privacy rests on DDH via dual-generator ElGamal; amount indistinguishability on DDH plus independent randomizers; unlinkability on DDH via the anonymity set, delta aggregation, and independent randomizers; transaction correctness on DLP and hash collision resistance via ZK proofs and nullifiers; decryptable balances on small-exponent DLP via ElGamal with BSGS and range proofs; and on-chain closure on DDH via the view key and view tags. The differentiators are: (i) one DDH framework for all three privacy properties, (ii) formal leakage closure via independent randomizers (Theorem V.6), and (iii) decryptable balances with full on-chain closure.

VI. CRYPTOGRAPHIC-LAYER EVALUATION

This section evaluates HERMES empirically at two layers: §VI-B measures the four cryptographic-layer primitives (encryption, decryption, homomorphic addition, View-Key scanning), and §VI-C–§VI-F measure the ZK proof-system cost, its scalability with the anonymity-set size N, and the resulting throughput and memory overhead.

A. SETUP AND GOALS

We address four questions: **G1** does the construction sustain sub-millisecond latency for the four core operations; **G2** how does HERMES's cryptographic layer compare with representative schemes; **G3** does CCE (§III) translate into measurable overhead reductions versus a non-CCE baseline; **G4** how do proof-system cost and proof size scale as N grows from 8 to 64.

Setup. Single-core AMD Ryzen 7 5800X at 3.8 GHz, 32 GiB DDR4; Rust 1.82 (release, thin LTO), `halo2_proofs` for the proof system, `ark-bn254` for curve arithmetic; fixed-budget microbenchmark harness, mean of 100 rounds with $\geq 1{,}000$ inner iterations per round.

B. PRIMITIVE BENCHMARKS

The per-primitive benchmarks (Table I) answer G1: every operation completes within ≤ 1.5 ms. Encryption and homomorphic addition operate at 10^1–10^2 μs; BSGS decryption at 1.5 ms is paid at most once per incoming transfer (after View-Key filtering at ≈ 0.6 ms). Figure 3 summarizes HERMES's own measurements (top row) against a quantitative cross-scheme comparison (bottom row).

C. ZK PROOF-SYSTEM BENCHMARKS

Single IPA proof (full Transfer): Prove ≈ 5.0 s, verify ≈ 10 ms, proof size ≈ 1.46 KB. The Halo2/KZG backend yields $\approx 3{,}424$ B for on-chain use (§VI-F). The circuit-level constraint breakdown (Supplementary Appendix G) answers G3: the CCE-driven design eliminates the bridging Σ-protocol that non-CCE schemes require for range proofs on Y_2.

D. BRIDGING OVERHEAD: EMPIRICAL VALIDATION

Theorem III.10 (§III-D) establishes that non-CCE schemes incur at least $N \cdot (2\lceil \log_2 q \rceil + 1)$ additional constraints. For HERMES (N=8, BN254, $q \approx 2^{254}$), the baseline adds $2\lceil \log_2 q \rceil = 508$ commitment constraints plus 1 copy constraint per member, totaling 4,072 constraints—an 80% saving relative to the 5,094-constraint non-CCE baseline (the 1,022-constraint HERMES Transfer circuit plus these bridging constraints). CCE achieves the theoretical minimum: C_2 *is* the commitment, so range proofs act directly on ciphertexts.

E. SCALABILITY VERSUS ANONYMITY-SET SIZE

The One-out-of-Many module (R5) dominates with $O(N \log N)$ constraints: from the N=8 baseline (1,022 constraints; 5,000 ms prove; 1.46 KB), projected totals grow to 1,038/5,100 ms at N=16, 1,134/5,800 ms at N=32, and 1,390/7,700 ms at N=64, with proof size stable at ≈ 1.5 KB—a practical range throughout.

F. SYSTEM OVERHEADS

Native throughput: ≈ 0.2 txs/s (prove ≈ 5 s + verify ≈ 10 ms per Transfer); batch aggregation addresses this at the system layer. BSGS table (2^{16}, ≈ 2 MiB) is built once and kept resident. Per-account state is 160 B (64 B ciphertext, 32 B view-key public key, 64 B SMT leaf), i.e., $\sim$160 MiB at 10^6 accounts.

Per transaction, the on-chain footprint comprises the anonymity-set keys (8×32 B), balance ciphertexts (8×64 B), view tags (8×4 B), one nullifier and one old state root (2×32 B), and the Halo2/KZG proof ($\sim$3,424 B: a $\sim$48 B $\mathbb{G}_1$ commitment plus evaluation openings), totaling $\sim$4,288 B. Under recursive batch composition the proof size stays $\sim$3,424 B regardless of batch size M; at M=64 the amortized cost is $\sim$53 B/tx.

G. COMPARISON WITH PEER-REVIEWED SCHEMES

TABLE II
COMPARISON OF PEER-REVIEWED PRIVACY TRANSACTION SCHEMES (TS? = TRUSTED SETUP; ✓= PROVIDED, × = NOT, ○ = PARTIAL, RING = RING-BASED UTXO). "UNIFIED DDH" = ALL THREE PROPERTIES REDUCED TO DDH IN ONE GAME-BASED MODEL.

Scheme (Venue)	Model	Encryption	Proof	TS?	Bal.	Amt.	Unlink.	Unified DDH
Zether (FC'20) [9]	Acct	Std ElGamal	Σ+Bullets	No	✓	✓	×	Partial
Anon. Zether (S&P'21) [26]	Acct	Std ElGamal	Σ+Bullets	No	✓	✓	✓	Partial
QuisQuis (ASIACRYPT'19) [29]	Hybrid	ElGamal+Com	Σ+Bullets	No	✓	✓	○	Partial
PGC (ESORICS'20) [10]	Acct	Twisted ElGamal	Σ+Bullets	No	✓	✓	×	Partial
BlockMaze (TDSC'22) [7]	Acct	Pedersen Com	Groth16	Yes	✓	✓	○	No
LedgerMaze (TC'23) [30]	Acct	ElGamal+Com	NIZK	Yes	✓	✓	✓	No
PriDe CT (S&P'24) [31]	Acct	ElGamal/Com	zkSNARK	Yes	✓	✓	○	No
MatRiCT+ (S&P'22) [32]	UTXO	Ring-LWE	Lattice ZK	No	✓	✓	Ring	No (PQ)
HERMES (this)	**Acct**	**Dual-Gen EG**	**Halo2**	**No**	✓	✓	✓	**Yes**

A proof-system throughput comparison under KZG batching (M=64) is deferred to Supplementary Appendix G; HERMES attains 133 batch-amortized TPS with a transparent setup.

VII. CONCLUSION AND FUTURE WORK

Summary. We presented HERMES, a permissionless account-model privacy-transaction scheme whose security argument reduces to the standard DDH assumption via three game-based experiments: balance privacy (tight reduction, $\mathsf{Adv}_1 \leq \mathsf{Adv}^{\mathsf{DDH}}$), transaction-amount indistinguishability ($\mathsf{Adv}_2 \leq 2 \cdot \mathsf{Adv}^{\mathsf{DDH}}$), and sender–receiver unlinkability

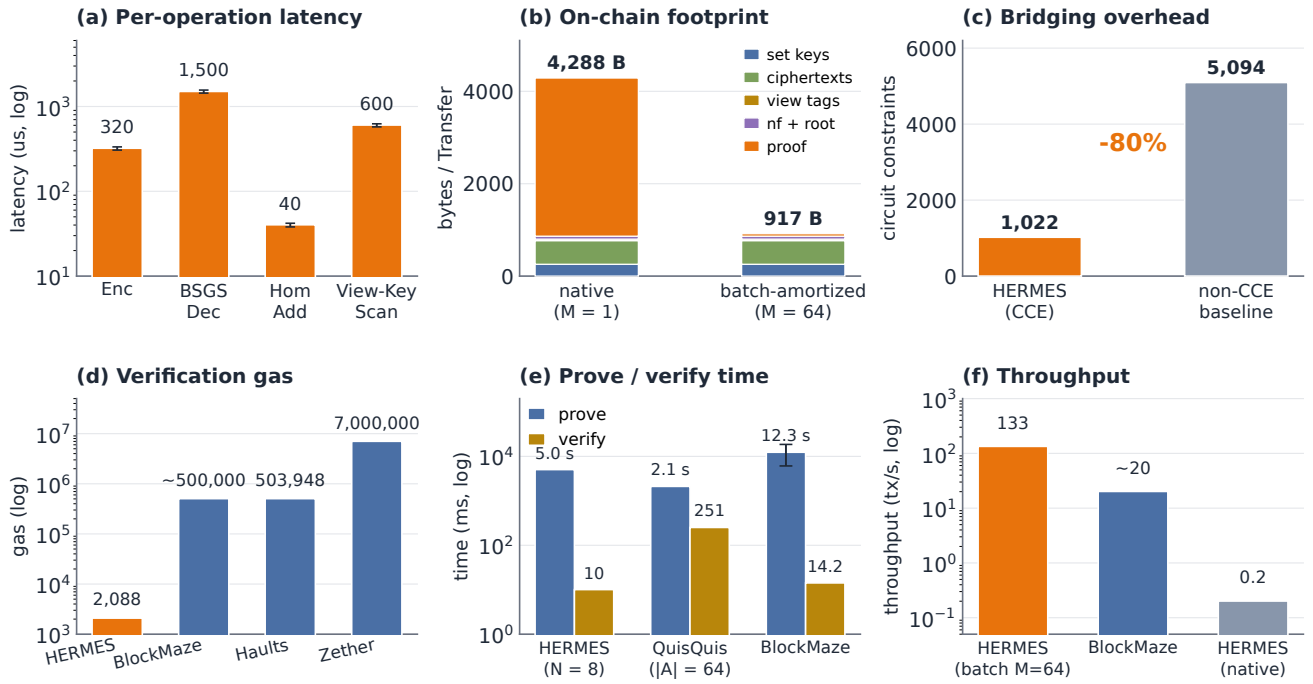

Fig. 3. Evaluation summary. Top row, HERMES measurements on BN254: (a) per-operation latency (log scale; error bars: ± 1 std dev over 100 rounds); (b) on-chain footprint per Transfer, native versus batch-amortized (M=64); (c) bridging-overhead validation at N=8 (§VI-D). Bottom row, cross-scheme comparison as reported in the respective papers (hardware, proof systems, and EVM versions differ; values are indicative rather than apples-to-apples): (d) per-transaction proof-verification gas—HERMES measured on Sepolia, amortized over M=64; BlockMaze an order-of-magnitude estimate [7]; Haults [28]; Zether pre-EIP-1108 [9]; (e) proof generation and verification time (BlockMaze proving ranges 6.1–18.6 s [7]; QuisQuis at $|A|$=64 [29]); and (f) throughput—HERMES batch-amortized (M=64, §VI-F) versus BlockMaze ($\approx$ 20 TPS [7]) and HERMES native (no batching).

($\mathsf{Adv}_3 \leq \mathsf{Adv}^{\mathsf{ZK}} + 2N \cdot \mathsf{Adv}^{\mathsf{DDH}}$). The construction rests on three co-designed components: (i) Commitment-Coherent Encryption (CCE, §III), a new primitive that fuses Pedersen commitments and exponential ElGamal into a single algebraic object and is proven to be essentially unique among additively homomorphic, decryptable schemes (Theorem III.5); (ii) independent per-ciphertext randomizers that close all algebraic-leakage surfaces under DDH (Theorem V.6); and (iii) a five-requirement ZK circuit specification (R1–R5) translating cryptographic invariants into verifiable constraints.

Limitations. The present instantiation carries bounded limitations, none of which affects confidentiality: the default anonymity set $N = 8$ may be small against adversaries with strong auxiliary information; the Sequencer is a single liveness party that can stall or censor but cannot violate confidentiality; and the formal proofs target static adversaries, with the static-to-adaptive lemma (§V-D) closing the gap at a $\binom{N}{t}$-factor cost.

Future Work. Two directions extend this work beyond its present scope. First, *post-quantum migration*: the algebraic characterization of §III-C isolates the exact structure that any coherent instantiation must realize, and an LWE-based analogue of CCE is the natural migration target toward quantum-resistant account-model privacy. Second, *auditable privacy*: practical deployments of confidential transactions increasingly require selective disclosure to designated parties, and HERMES's key separation offers a principled starting point—a proof-based auditing interface in which compliance is verifiable without weakening the privacy of non-audited users would bridge provable privacy and regulation.

References

[1] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," 2014. [Online]. Available: https://ethereum.org/en/whitepaper/

[2] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *IEEE Symposium on Security and Privacy (S&P)*, 2020, pp. 910–927.

[3] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, "A fistful of bitcoins: Characterizing payments among men with no names," *Communications of the ACM*, vol. 59, no. 4, pp. 86–93, 2016.

[4] E. B. Sasson, A. Chiesa, and M. e. a. Green, "Zerocash: Decentralized anonymous payments from bitcoin," in *IEEE S&P 2014*, 2014, pp. 459–474.

[5] Monero Community, "Monero: Private digital currency," 2015. [Online]. Available: https://www.getmonero.org/

[6] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *CRYPTO 1991*, 1992, pp. 129–140.

[7] Z. Guan, Z. Wan, and Y. e. a. Yang, "Blockmaze: An efficient privacy-preserving account-model blockchain based on zk-snarks," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, pp. 1446–1463, 2022.

[8] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Transactions on Information Theory*, vol. 31, no. 4, pp. 469–472, 1985.

[9] B. Bünz, S. Agrawal, M. Zamani, and D. Boneh, "Zether: Towards privacy in a smart contract world," in *FC 2020*, 2020, pp. 423–443. [Online]. Available: https://eprint.iacr.org/2019/191

[10] Y. Chen, X. Ma, C. Tang, and M. H. Au, "PGC: Decentralized confidential payment system with auditability," in *European Symposium on Research in Computer Security (ESORICS)*, ser. Lecture Notes in Computer Science, vol. 12308, 2020, pp. 591–610.

[11] S. Labs, "Spl token-2022: Confidential transfer extension," 2022. [Online]. Available: https://spl.solana.com/confidential-transfers

[12] A. Chystiakov and M. Zhvanko, "Confidential wrapped ethereum," Cryptology ePrint Archive, Report 2507.09231, 2025, preprint / engineering whitepaper, no formal security analysis. [Online]. Available: https://arxiv.org/abs/2507.09231

[13] A. S. Inc., "Aleo: Zero-knowledge programming platform," Project documentation, 2023. [Online]. Available: https://aleo.org

[14] P. Labs, "Penumbra: A shielded, cross-chain network for cosmos," Protocol specification, 2023. [Online]. Available: https://protocol.penumbra.zone/

[15] M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn, "Sonic: Zero-knowledge snarks from linear-size universal and updateable structured reference strings," in *CCS 2019*, 2019, pp. 2111–2128. [Online]. Available: https://eprint.iacr.org/2019/099

[16] A. Kothapalli, S. Setty, and I. Tzialla, "Nova: Recursive zero-knowledge arguments from folding schemes," in *CRYPTO 2022*, 2022, pp. 359–388. [Online]. Available: https://eprint.iacr.org/2021/370

[17] Z. Foundation, "Halo2: A proof system supporting recursive composition," 2022. [Online]. Available: https://github.com/zcash/halo2

[18] A. Team, "Aztec network technical whitepaper," 2024. [Online]. Available: https://aztec.network

[19] D. Shanks, "Algorithms for finding discrete logarithms in $gf(p)$," in *Proceedings of the 2nd Southeastern Conference on Combinatorics, Graph Theory and Computing*, 1971, pp. 51–57.

[20] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, 1989.

[21] A. Gabizon, Z. J. Williamson, and O. Ciobotaru, "Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge," in *Cryptology ePrint Archive, Report 2019/953*, 2019. [Online]. Available: https://eprint.iacr.org/2019/953

[22] P. Valiant, "Incrementally verifiable computation or proofs of knowledge imply time/space efficiency," in *TCC 2008*, 2008, pp. 1–18.

[23] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT 1999*, 1999, pp. 223–238.

[24] M. Fischlin, "Short proofs of knowledge," in *Topics in Cryptology – CT-RSA 2001*, ser. Lecture Notes in Computer Science, vol. 2020, 2001, pp. 112–126.

[25] D. Hopwood, S. Bowe, T. Hornby, and N. Wilcox, "Zcash protocol specification: Sapling (version 2018.0)," Electric Coin Company Technical Report, 2018. [Online]. Available: https://zips.z.cash/protocol/sapling.pdf

[26] B. E. Diamond, "Many-out-of-many proofs and applications to anonymous zether," in *IEEE S&P 2021*, 2021, pp. 1220–1237. [Online]. Available: https://eprint.iacr.org/2020/293

[27] J. Groth and M. Kohlweiss, "One-out-of-many proofs with applications to anonymous zether," in *EUROCRYPT 2015, Part II*, 2015, pp. 253–277. [Online]. Available: https://eprint.iacr.org/2014/764

[28] X. Salleras, "Homomorphic encryption-based vaults for anonymous balances on VM-enabled blockchains," arXiv:2511.17842, 2025.

[29] P. Fauzi, S. Meiklejohn, R. Mercer, and C. Orlandi, "Quisquis: A new design for anonymous cryptocurrencies," in *Advances in Cryptology – ASIACRYPT 2019*, ser. Lecture Notes in Computer Science, vol. 11476, 2019, pp. 649–678. [Online]. Available: https://eprint.iacr.org/2018/990

[30] Z. Bao, D. He, W. Wei, C. Peng, and X. Huang, "Ledgermaze: An efficient privacy-preserving noninteractive zero-knowledge scheme over account-model blockchain," *IEEE Transactions on Computers*, vol. 72, no. 12, pp. 3489–3502, 2023.

[31] Y. Guo, H. Karthikeyan, A. Polychroniadou, and C. Huussin, "PriDe CT: Towards public consensus, private transactions, and forward secrecy in decentralized payments," in *IEEE Symposium on Security and Privacy (S&P)*, 2024, pp. 3904–3922.

[32] M. F. Esgin, R. Steinfeld, and R. K. Zhao, "MatRiCT+: More efficient post-quantum private blockchain payments," in *IEEE Symposium on Security and Privacy (S&P)*, 2022.