Post-Quantum Signatureless Authorization

Ignotus Nemo
dev@parano1d.org

29 October 2026

Abstract

Digital signatures authorize transactions, commands and other requests. Quantum attacks threaten widely used signature schemes. Their post-quantum replacements can be costly to verify inside proofs of execution. Signatureless authorization binds knowledge of an authority secret to a canonical request through a freshly randomized proof capsule. The public statement consists of a request identifier and a 32-byte authority commitment. Application rules enforce context, policy and freshness. The interface supports standalone requests and recursive proof composition.

The construction combines a masked degree-ten authority reduction over $\mathbb{F}_{2^{128}} \subset \mathbb{F}_{2^{256}}$ with a bank-and-companion code proof. Source restoration gives expected polynomial-time public-coin extraction with error at most $5^{-65} + 701202001642/2^{255}$.

The native model fixes the public authority circuit and its address input slice, then samples a uniform permutation completion elsewhere. Under explicit finite orbit-profile conditions, classical authorization requests and coherent forward and inverse queries admit the bound $\Pr[\mathsf{Forge}_{\mathrm{Q}}] \leq \mathrm{Adv}^{\mathrm{Q}}_{\mathrm{pre}} + \varepsilon^{\mathrm{Q}}_{\mathrm{sim}} + \varepsilon^{\mathrm{Q}}_{\mathrm{ext}} + \varepsilon^{\mathrm{Q}}_{\mathrm{id}}$. An adaptive witness-free simulator and expected polynomial-time source extractor give explicit error terms for the shared native transcript and commitment permutation. The source bank proves the fixed authority circuit's relation. For the ideal 256-bit authority-preimage task, constant-success generic quantum search has query scale 2^{128}.

Parano1d supplies a blockchain implementation and recursive verifier. It uses a 2,048-cell bank and 65 code queries, with canonical capsules bounded by 92,696 bytes. On four CPU threads, median proving takes 0.81 to 0.89 seconds and verification takes approximately 25.6 milliseconds for two request fixtures.

1 Introduction

Quantum algorithms for factoring and discrete logarithms threaten authorization based on these problems [Sho95]. NIST identifies affected standards and post-quantum replacements [MPR+26]. The ML-DSA parameter sets in FIPS 204 use 1,312 to 2,592-byte public keys and 2,420 to 4,627-byte signatures [Nat24a]. SLH-DSA signatures in FIPS 205 occupy 7,856 to 49,856 bytes [Nat24b]. Applications that certify both authorization and execution must also represent authorization verification in their proof arithmetic. Post-quantum Privacy Pass studies this coupling through STARK-oriented Dilithium [PWFHW23].

This motivates the following design question: how to bind each request to its authority within the same arithmetic used to prove execution, while protecting a reusable secret across repeated authorizations. The resulting relation must also support recursive verification and a stable boundary between authority and public execution. We therefore consider a direct proof of knowledge of an authority secret.

Signatureless authorization binds freshly randomized proofs of knowledge of an authority secret to canonical requests. For intent I, identifier $m = \mathsf{ID}(I)$ and required 32-byte authority commitment a, the public statement and authority relation are

$$x = (m, a), \qquad \mathcal{R}_{\mathrm{own}}(a, s) = [\mathsf{Addr}(s) = a]. \tag{1}$$

The adapter binds context and authority. Application rules check execution, policy and freshness. For fixed intent and authority, the same capsule remains verifiable when the public execution witness is updated within the intent's constraints.

The construction reduces the address relation to eleven linear functionals of a 2,048-cell bank. The code argument is an interactive oracle proof (IOP). Exact masking ranks and constant Owner fibers give causal simulation. Finite-list restoration gives round-by-round knowledge. A native compiler preserves complete response groups, the Owner bridge and the shared transcript and commitment permutation. Complete simulation and extraction establish quantum authorization in the address-compatible random-completion model under finite slice-profile conditions, with coherent forward and inverse access. Parano1d supplies the blockchain and recursive-verification case study.

Related work includes signatures of knowledge [CL06] and post-quantum signatures built from symmetric primitives [CDG+17]. The algebra uses binary-field permutations [GKK+25], multilinear reductions [LFKN92, GKR08], Libra masking [XZZ+19] and binary-tower code commitments [ZCF23, DP24].

2 Statements, authority and execution

For $s = (s_0, s_1) \in \mathbb{F}_{2^{128}}^2$, fixed address capacity (ι_0, ι_1) and the width-four Poseidon2b permutation,

$$\mathsf{Addr}(s) = \mathrm{first}_2\, P(s_0, s_1, \iota_0, \iota_1) \in \mathbb{F}_{2^{128}}^2. \tag{2}$$

The witness stores the initial linear layer and 66 nonlinear rounds. Two linear boundaries fix its initial capacity and two its final address.

Lemma 2.1 (Address witness)**.** *With exact arithmetic and authenticated permutation constants, a state table satisfying the* $66 \cdot 4$ *transitions and four boundaries determines a secret* s *with* $\mathsf{Addr}(s) = a$.

Proof. Invert the initial linear layer. Its first two coordinates give s and the other two are fixed by the initial boundaries. Induction identifies the prescribed trace. Its final boundaries give a. □

Intents. The application defines domain, operation, context, action data and execution constraints in

$$m = \mathsf{ID}(I) = H_{\mathrm{intent}}(\mathsf{domain}, \mathrm{encode}(I)). \tag{3}$$

Canonical encoding fixes lengths, field order and unused records. The authority adapter derives a from this checked intent and policy. Services bind audience, operation, parameters and nonce. Devices bind target, command, counter and validity. Artifact approvals bind digest, intended use and policy. Ledgers bind context, resources and action terms. Execution checks the applicable authority, policy and freshness. A freshness condition consumes a nonce, counter or resource reference.

Algorithms. $\mathsf{Statement}(I, c)$ validates intent and authority context, then returns $x = (\mathsf{ID}(I), a)$. $\mathsf{Authorize}(s, x)$ checks the address and uses fresh randomness to produce the capsule. $\mathsf{CheckAuth}(x, \pi)$ checks shape and replays transcript, reduction and code openings. $\mathsf{Execute}(I, \mathsf{st}, c)$ checks state, context, authority, policy and freshness, then returns the permitted successor state. The public vector is

$$(m_0, m_1, a_0, a_1) \in \mathbb{F}_{2^{128}}^4. \tag{4}$$

The verifier derives all challenges and ordered claims itself.

Proposition 2.2 (Stable authorization boundary)**.** *For fixed I, a and authenticated capsule parameters,* $\mathsf{CheckAuth}((\mathsf{ID}(I), a), \pi)$ *is identical for every public state. Execution acceptance also requires* $\mathsf{Execute}$ *to succeed in that state and context.*

Proof. Authorization inputs are exactly the parameters, four public fields and capsule bytes. State-dependent predicates belong to execution. □

3 The authorization capsule

Fields and challenges. The bank uses $\mathbb{F}_{2^{128}}$. Messages use the quadratic extension

$$\mathbb{F}_{2^{256}} = \mathbb{F}_{2^{128}}[X]/(X^2 + X + \tau), \qquad \mathrm{Tr}_{\mathbb{F}_{2^{128}}/\mathbb{F}_2}(\tau) = 1. \tag{5}$$

Raw lanes $u, v \in \mathbb{F}_{2^{128}}$ give

$$u + (v^2 + v + \tau)X. \tag{6}$$

Independent uniform raw lanes give uniform challenges on

$$\mathcal{S} = \{u + wX : \mathrm{Tr}(w) = 1\}, \qquad |\mathcal{S}| = 2^{255}. \tag{7}$$

These elements are nonzero and outside $\mathbb{F}_{2^{128}}$. Public-coin bounds use independent draws and the degree-d root bound $d/|\mathcal{S}|$.

Bank. $B \in \mathbb{F}_{2^{128}}^{2048}$ has state cells 0 through 511, Libra buffer 512 through 943, terminal pads 944 through 953, fresh padding 954 through 1023, and source coins 1024 through 2047. State rows 0 through 66 hold the four-lane address trace. Other state rows are zero. All 1,536 remaining cells are fresh. Active Libra coefficients use 242 cells. The companion $C \in \mathbb{F}_{2^{256}}^{2048}$ is independently fresh. Affine Reed–Solomon encoding gives length 2^{16} and rate $1/32$. A source leaf contains eight bank and eight companion symbols. The source commitment is an ordered cap of eight 256-bit subtree roots, published before the Owner challenges.

Address polynomial. Boolean positions have two lane, seven row and two high bits. A selects zero high bits and rows below 66. Operands $N_B, S_{B,0}, \ldots, S_{B,3}$ select next-row and input-lane trace values on active positions. At Boolean positions 2047 and 2046, the five operand tables read their paired pads from bank cells 944 through 953. Other inactive entries are zero. For linear-layer tables M_j, nonlinear flags σ_j and round constants c_j, tildes denote multilinear extension and

$$\begin{aligned} Q_B(z) = \widetilde{A}(z)\Big[\widetilde{N}_B(z) + \sum_{j=0}^{3} \widetilde{M}_j(z)\big(\widetilde{\sigma}_j(z)(\widetilde{S}_{B,j}(z) + \widetilde{c}_j(z))^7 \\ + (1 + \widetilde{\sigma}_j(z))\widetilde{S}_{B,j}(z)\big)\Big]. \end{aligned} \tag{8}$$

Its individual degree is at most ten. Its Boolean table is exactly the transition residual and is zero off the active set. Pads therefore preserve transitions and sparse address boundaries.

Owner. The framed channel absorbs (m, a) and the source cap, then draws $\rho \in \mathcal{S}^{11}$. The prover sends $\mu = \mathcal{M}[g](\rho)$ for its separable degree-ten mask, before $\lambda \in \mathcal{S}$. Here $\mathcal{M}[f]$ is the MLE of f's Boolean table. The eleven high-to-low MLE-check rounds reduce

$$\mathcal{M}[Q_B + \lambda g](\rho) = \lambda\mu. \tag{9}$$

Each round sends ten nonconstant coefficients of R_i. For running claim c, the verifier reconstructs

$$(1+\rho_i)R_i(0)+\rho_i R_i(1)=c, \qquad a_0=c+\rho_i\sum_{k=1}^{10}a_k, \tag{10}$$

then draws $r_i\in\mathcal{S}$ and updates to $R_i(r_i)$. The final point is $(r_0,\dots,r_{10})$. Five terminal operands and $g(r)$ supply the final check $Q_B(r)+\lambda g(r)$. Only afterward is η drawn.

Bank claims. The ordered linear claims (w_j,v_j) are five operands, four boundaries, μ and $g(r)$. Form

$$t=\sum_{j=0}^{10}\eta^j w_j, \qquad b=\sum_{j=0}^{10}\eta^j v_j. \tag{11}$$

Batching links the Owner claims to the source bank through $\langle B,t\rangle=b$. Every weight is verifier-derived. Owner closes into four base-field bridge lanes, absorbed by Main under its fixed handoff domain.

Main and openings. The prover sends $\sigma=\langle C,t\rangle$ before γ, then forms

$$U_\gamma=(1+\gamma)B+\gamma C. \tag{12}$$

Phase A is eleven degree-two sumcheck rounds for

$$\sum_{z\in\{0,1\}^{11}}\widetilde{U}_\gamma(z)\widetilde{t}(z)=(1+\gamma)b+\gamma\sigma. \tag{13}$$

At its point s, Phase B authenticates $v=\widetilde{U}_\gamma(s)$. It supplies an upper table of 256 extension-field elements, three source folds, a middle cap, four middle folds, sixteen tail coefficients and one local fold. Each challenge follows the messages it constrains. A 64-bit nonce is accepted when the low sixteen bits of the grinding lane are zero. Seven base-field seed lanes give 65 positions from disjoint 13-bit windows. Ordered source and middle openings check both caps, the fold equations and the Phase-A evaluation. Appendix A gives the causal schedule and encoding.

4 Masking algebra

The following ranks fix admissible challenges and independent uniform masks. Appendix C gives their causal coupling.

At $r\in\mathcal{S}^{11}$, the weights of positions 2047 and 2046 are

$$t_r=\prod_{i=1}^{10}r_i, \qquad k_0=t_r r_0, \qquad k_1=t_r(1+r_0). \tag{14}$$

The five terminal operands are

$$Y_j=A_j(r)+k_0p_j+k_1q_j, \qquad p_j,q_j\in\mathbb{F}_{2^{128}}, \quad j<5. \tag{15}$$

Here $A_j(r)$ depends only on the address state table.

Theorem 4.1 (Terminal operand masking)**.** *For fixed $r \in \mathcal{S}^{11}$ and any state table, independent uniform pads $(p_j, q_j)_{j<5}$ make $(Y_j)_{j<5}$ uniform in $\mathbb{F}_{2^{256}}^5$. The map has rank ten over $\mathbb{F}_{2^{128}}$. Pads preserve Boolean transitions and address boundaries.*

Proof. Write $r_0 = u + vX$, with $v \neq 0$. Then

$$\det_{\mathbb{F}_{2^{128}}}(k_0, k_1) = \mathrm{Norm}_{\mathbb{F}_{2^{256}}/\mathbb{F}_{2^{128}}}(t_r)\, v \neq 0. \tag{16}$$

Indeed $(r_0, 1 + r_0)$ has determinant v and multiplication by t_r has determinant its norm. Each pair bijects $\mathbb{F}_{2^{128}}^2$ to $\mathbb{F}_{2^{256}}$. Five disjoint pairs give rank ten and uniform outputs. Their support is inactive in Q_B and absent from the boundaries. □

Libra. The mask uses 121 independent extension-field coefficients,

$$g(z) = \sum_{i=0}^{10} \sum_{k=0}^{10} a_{i,k} z_i^k, \qquad a_{i,k} = \ell_{i,k} + X h_{i,k}. \tag{17}$$

Coefficient $a_{i,k}$ uses base-field cells $512 + 16i + k$ and $768 + 16i + k$ for its low and high limbs. Other buffer cells remain fresh. The observations are μ, the 110 contributions $\lambda a_{i,k}$ for $k > 0$, and $g(r)$, with

$$\mu = \sum_{i=0}^{10} \left(a_{i,0} + \rho_i \sum_{k=1}^{10} a_{i,k} \right). \tag{18}$$

Theorem 4.2 (Libra observation rank)**.** *For fixed ρ, r and $\lambda \neq 0$, these 112 extension-field observations have rank 111 over $\mathbb{F}_{2^{256}}$. Their sole relation is the endpoint telescope. The base-field rank is 222, leaving 20 active base-field degrees of freedom.*

Proof. The first 110 observations recover all nonconstant coefficients. Then μ fixes the independent sum $\sum_i a_{i,0}$, while

$$g(r) = \mu + \sum_{i=0}^{10} \sum_{k=1}^{10} a_{i,k}(r_i^k + \rho_i). \tag{19}$$

Thus the first 111 observations are independent and the last adds this relation. Doubling the extension-field rank gives 222, with kernel dimension $242 - 222 = 20$. □

Source coins. For the encoding basis,

$$X_{1024+j}(z) = W_{10}(z) X_j(z), \qquad j < 1024, \tag{20}$$

and the affine evaluation domain avoids the zeros of W_{10}.

Lemma 4.3 (Source-opening rank)**.** *At any $q \leq 1024$ distinct source positions, the dedicated coefficient block has evaluation rank q over $\mathbb{F}_{2^{128}}$. Its bank symbols are uniform conditional on all other cells.*

Proof. The first 1024 basis elements span degree-below-1024 polynomials. Interpolation gives rank q. Nonzero row factors W_{10} preserve it, as does translation from the other cells. □

The 65 leaves expose at most 520 distinct bank positions.

Proposition 4.4 (Disjoint masking blocks)**.** *For fixed admissible challenges and $q \le 520$ distinct source positions, source-bank, terminal and Libra observations have joint rank $q+232$ over $\mathbb{F}_{2^{128}}$. Their $q+234$ base-field coordinates satisfy two Libra telescope relations.*

Proof. Select a rank-q source minor, ten pad columns and 222 independent Libra columns. Other observations vanish on the source block. The pad and Libra blocks are invertible, giving a block triangular minor of size $q+10+222$. The two relations give the matching upper bound. □

Conditioned companion. The scalar $\sigma = \langle C, t\rangle$ precedes γ. For $\langle B, t\rangle = b$, define

$$\mathcal{C}_{t,\sigma} = \{C : \langle C, t\rangle = \sigma\}, \tag{21}$$

$$\mathcal{U}_{t,b,\sigma,\gamma} = \{U : \langle U, t\rangle = (1+\gamma)b + \gamma\sigma\}. \tag{22}$$

Lemma 4.5 (Conditioned companion bijection)**.** *For $\gamma \neq 0$, $C \mapsto (1+\gamma)B + \gamma C$ bijects these fibers and preserves uniformity. Their dimension is 2047 over $\mathbb{F}_{2^{256}}$ when $t \neq 0$, and 2048 when $t = 0$ and the claims are zero.*

Proof. The inverse $C = \gamma^{-1}(U + (1+\gamma)B)$ satisfies the source equation. A nonzero functional has rank one. The zero functional with zero claim imposes no constraint. □

Every linear L therefore satisfies

$$L(C) = \gamma^{-1}\big(L(U_\gamma) + (1+\gamma)L(B)\big). \tag{23}$$

5 Authorization security

For banks fixed before the coins, incorrect post-claims give a nonzero degree-ten discrepancy in η. With $e_B = b - \langle B, t\rangle$ and $e_C = \sigma - \langle C, t\rangle$, the blend error is

$$e_B + \gamma(e_B + e_C). \tag{24}$$

It is nonzero of degree at most one for any incorrect required input.

Theorem 5.1 (Algebraic reduction)**.** *Fix $x = (m, a)$ and banks $B \in \mathbb{F}_{2^{128}}^{2048}$ and $C \in \mathbb{F}_{2^{256}}^{2048}$ before the verifier's coins. Use the transcript order of Section 3 with independent challenges from $\mathcal{S}$. The probability that the algebraic verifier accepts, the Phase-B value is the correct $\widetilde{U}_\gamma(s)$, and B violates an address transition or boundary equation is at most*

$$\varepsilon_{\mathrm{alg}} = \frac{11 + 110 + 10 + 1 + 22}{2^{255}} = \frac{154}{2^{255}}. \tag{25}$$

The protocol may use arbitrary adaptive prover messages.

Proof. A nonzero Boolean transition table has a nonzero degree-eleven MLE, so its evaluation vanishes at ρ with probability at most $11/|\mathcal{S}|$. Conditional on correct bank claims, eleven false Owner round differences of degree ten cost at most $110/|\mathcal{S}|$. A false post-claim vector batches correctly only at a root of its nonzero degree-ten polynomial in η, costing $10/|\mathcal{S}|$. A false bank or companion claim gives the nonzero degree-one blend error in Equation (24). Finally, a false Phase-A target reaches a correct terminal product only through a degree-two round coincidence, costing at most $22/|\mathcal{S}|$.

Outside this union, the correct Phase-B evaluation propagates backward through Phase A, the blend and the batched bank claims to the Owner terminal equation. Its round checks then give the correct address transition table and all four boundaries. This contradicts the stated invalidity. Each polynomial is fixed before its independent challenge, so the argument permits adaptive prover messages. □

Corollary 5.2 (Reduction with a code-opening bound)**.** *Suppose a commitment argument supplies banks fixed before the algebraic coins and authenticates the linked evaluation $\widetilde{U}_\gamma(s)$, with joint failure probability at most $\varepsilon_{\mathrm{pcs}}$. Then acceptance with an invalid address bank has probability at most $\varepsilon_{\mathrm{pcs}} + 154/2^{255}$ in that public-coin experiment.*

Proof. On a correct commitment/evaluation event apply Theorem 5.1. Add the commitment failure event once. □

Source restoration and round-by-round knowledge are proved in Theorems D.3 and F.3. Complete simulation and fresh extraction are proved in Appendices E and G.

5.1 Fresh-intent forgery game

The authority samples $s \leftarrow \mathbb{F}_{2^{128}}^2$ and publishes $a = \mathsf{Addr}(s)$. An adversary can request complete authorizations on adaptively chosen canonical intents for this authority. Each request uses fresh independent prover randomness. The adversary wins by producing $(I_\star, \pi_\star)$ such that the required authority is a, the capsule verifies, and the canonical intent was not authorized by a previous oracle request.

Policy freshness refers to the authorized policy statement and its bound resources. Satisfying execution choices remain within that policy.

Let $\varepsilon_{\mathrm{id}}$ bound the event that a semantically fresh intent receives an identifier already authorized in the same authority context. This event includes a relevant hash collision or a failure of the canonical semantic encoding. The public adapters and domain grammar are part of this condition.

5.2 Quantum chosen-request authorization

Requests, responses and final outputs are classical. The adversary has quantum workspace and coherent forward and inverse access to the native permutation within its resource envelope.

Let P_0 be the fixed public authority circuit of Equation (2). In flat rate/capacity coordinates, put

$$A = \{(u, c_{\mathrm{addr}}) : u \in \{0,1\}^{256}\}, \qquad F = P_0|_A.$$

The native oracle P is a uniform permutation conditioned on $P|_A = F$. The bank proves the P_0 transition and boundary equations, so its validated secret satisfies $\mathsf{Addr}(s) = \mathrm{rate}\, P_0(s, c_{\mathrm{addr}}) = a$. The native address interface agrees with this same circuit on A. The finite orbit profiles of $g(u) = \mathrm{rate}\, P_0(u, c_{\mathrm{addr}}) \oplus u$ and their resource conditions are specified in Appendix H. One permutation supplies the native transcript, source and middle sponges and shared feed-forward nodes.

Let $\mathrm{Adv}^{\mathrm{Q}}_{\mathrm{pre}}(\mathcal{R}')$ be the best probability of recovering a preimage of an address sampled as $\mathsf{Addr}(s)$ with the same public primitive data and oracle interfaces in the reduction's resource envelope $\mathcal{R}'$.

Theorem 5.3 (Native quantum fresh-intent authorization)**.** *In this random-completion model, under the finite profile and resource conditions of Appendices E–H and with independently fresh prover randomness, fresh-intent forgery after S authorization attempts and T charged base calls satisfies*

$$\Pr[\mathsf{Forge}_{\mathrm{Q}}] \le \mathrm{Adv}^{\mathrm{Q}}_{\mathrm{pre}}(\mathcal{R}') + \varepsilon^{\mathrm{Q}}_{\mathrm{sim}}(S,T) + \varepsilon^{\mathrm{Q}}_{\mathrm{ext}}(S,T) + \varepsilon^{\mathrm{Q}}_{\mathrm{id}}. \tag{26}$$

The explicit simulation and extraction losses are Equations (41) *and* (51)*. Simulation and extraction share one compressed base state and its classical input-relabeling table. All constructors, inverse calls, verification, recording rechecks and nonce-search work belong to T and $\mathcal{R}'$. In the sampled ideal-slice setup, $\varepsilon^{\mathrm{Q}}_{\mathrm{ext}}$ also includes the finite-profile setup failure once.*

Proof. Given a challenge address a, answer requests with the complete native simulator of Theorem E.2. Its accessible-view loss is at most $\varepsilon^{\mathrm{Q}}_{\mathrm{sim}}$. Outside intent-binding failure, a fresh semantic intent has an identifier absent from the simulator's requested set. The one-run extractor of Theorem G.2 verifies the capsule, rechecks its base pairs in order, and decodes the native source word. Apart from $\varepsilon^{\mathrm{Q}}_{\mathrm{ext}}$, it returns a bank whose validated input lanes satisfy $\mathsf{Addr}(s) = a$. Return those lanes in the preimage game. Its success is bounded by $\mathrm{Adv}^{\mathrm{Q}}_{\mathrm{pre}}(\mathcal{R}')$. Adding the displayed losses proves the bound. □

The simulation covers repeated authorizations, companion claims and adaptive source positions. Every attempt uses fresh independent masks, including abandoned attempts. The native extraction proof preserves all 30 response groups and fixes each causal source snapshot through later authorization, verification and recheck calls.

For $S \le T$, the combined simulation and extraction loss is below 2^{-67} when $T \le 2^{24}$ and below 2^{-41} when $T \le 2^{32}$. These bounds include the uniform-slice setup failure once. The artifact gives exact certificates for these resource envelopes.

For the ideal 256-bit authority-preimage task, constant-success generic quantum search has query scale 2^{128} [Gro96]. Equation (26) quantifies the complete authorization probability with its simulation, extraction and intent-binding terms.

6 Blockchain and recursive verification

Parano1d's wallet produces the capsule for (m, a) while its block prover supplies the public execution witness. For fixed intent and authority, the same capsule remains verifiable when this witness is updated within the intent's constraints. Paged payments bind ordered pages, owner, epoch, fee, input incarnations and outputs. Execution consumes those incarnations and allocates fresh slots. Asset adapters bind operations and opening commitments. Trades bind policy and ordered funding, constraining settlement limits, recipients and fees.

The arithmetic verifier shares the Owner point, operands, masks and η with the post-claim relation, which derives b and $\widetilde{t}(s)$. The wrapped tile contributes 16,536 incremental rows, including 1,852 for Owner, post-claim and Phase A. Shared hash regions and execution predicates belong to the enclosing relation. The artifact records component counts.

Proposition 6.1 (Recursive composition)**.** *Assume exact binary-field arithmetic and complete tile coverage. The transcript and hash regions are authenticated and statement and capsule bindings are canonical. A satisfying recursive witness then enforces each selected capsule's native equations. The included execution predicates enforce the intent's state and policy constraints.*

Proof. Statement pins fix the four fields. Duplex regions, the challenge map and four-lane bridge fix the coins and handoff. Shared expressions identify Owner, post-claim and Phase-A values. Phase B links its evaluation to typed source and middle openings and ordered caps. Tile coverage supplies these checks for every required slot. Execution predicates impose the intent's state and policy rules. □

Production preparation checks canonical bytes and binds transcript material to region assembly through a private capability. Finalization requires all tiles and shared regions. The recursive proof uses its own soundness and authenticated boundary.

7 Implementation and evaluation

The artifact pins 35 production files and reproduces the wallet entry point, encoding, native verification and algebra. Its manifest records revision, file lengths and SHA-256 digests. The

capsule ceiling is 92,696 bytes, including both caps, messages, nonce, payloads, shared siblings and six bounded-vector prefixes. Sibling ceilings are 453 for source and 193 for middle. Each attempt samples 1,536 fresh base-field bank cells and 2,048 fresh extension-field companion cells from operating-system entropy. The wallet's authority-witness type derives `Zeroize` and `ZeroizeOnDrop`.

Measurements use an Intel Core i7-1365U, Linux x86-64, Rust 1.96.0, release mode, CPU indices zero through three and four Rayon threads. Each fixture has one warmup and five fresh proofs, each verified three times. Verification statistics are medians of per-proof medians.

Intent	Prove (ms)	Verify (ms)	Capsule bytes
Payment (1 input, 1 output)	893.8	25.59	88,760 to 90,168
Consolidation (8 inputs, 2 outputs)	814.1	25.60	88,888 to 90,104

Table 1: Measured medians and observed capsule-size ranges.

Proving includes address trace, entropy, source commitment, proof, nonce search and self-verification. Serialization and decoding precede the separate verification timer. Each fixture authorizes one owner and one atomic page. All five sampled proofs in each fixture have distinct encodings. Verification rejects both tested changes in each fixture: increasing the fee and decreasing the first output amount by one, or adding $1 \in \mathbb{F}_{2^{256}}$ to the final mask value.

Production-field checks verify pad inversion, unchanged Boolean and boundary relations, Libra rank 222 and conditioned companion transport. Raw samples, matrices, inverse/opening checks, full wire/row ledgers and commands accompany the source.

8 Security model and primitive conditions

Theorem 5.3 establishes native chosen-request authorization in the address-compatible random-completion model. The public circuit P_0 determines the authority relation and the source bank's transition equations. Its address slice $F = P_0|_A$ is fixed, and the native permutation is a uniform completion of F with coherent forward and inverse access. The source, middle and duplex capacity roots are distinct and lie outside the address slice.

The finite-profile model fixes the slice's node branch g with $\Lambda_{10} \le 2^{15}$, $R_{10} < 2^{25}$ and $D_{10} < 2^{24}$, as defined in Appendix H. These conditions control inverse orbit fibers, short returns and shifted orbit intersections through the maximum authentication height. The declared resource envelopes in Appendices E–H ensure the stated sampling and compressed-oracle bounds apply. In the uniformly injected-slice idealization, the quantified setup failure is included once in $\varepsilon_{\mathrm{ext}}^{\mathrm{Q}}$.

Authority-preimage hardness and canonical intent binding for the selected Poseidon2b parameters enter Equation (26) through their separate terms. Application adapters enforce the semantics, policy and freshness in Section 2. Recursive composition uses the public bindings of Proposition 6.1. The source revision and measured production behavior are recorded in Section 7 and the accompanying artifact.

A Transcript and canonical capsule

The transcript uses fixed framing, domain tags and ordered field encodings. Each challenge follows the messages that it constrains. Table 2 lists the causal schedule. Algebraic challenges use Equation (6).

The Owner channel produces 24 extension-field challenges. Main produces twenty extension-field algebraic challenges, one base-field grinding value and seven base-field query seeds. The query index stream consists of 65 disjoint 13-bit windows in the 896 seed bits. Its 845 used bits select source leaves among 8,192 positions. Queries may repeat.

Channel	Absorbed or supplied material	Next challenge
Owner	Protocol domain, (m, a), source cap	$\rho \in \mathcal{S}^{11}$
Owner	Mask MLE value μ	$\lambda \in \mathcal{S}$
Owner	Ten nonconstant coefficients per round	$r_{10}, \ldots, r_0$
Owner	Final mask value and five operand values	$\eta \in \mathcal{S}$
Owner	Fixed close framing	Four-lane bridge
Main	Handoff domain, Owner bridge, σ	$\gamma \in \mathcal{S}$
Main	Two coefficients per Phase-A round	$s_{10}, \ldots, s_0$
Main	Claimed v and 256-field upper table	Three source-fold challenges
Main	Eight-digest middle cap	Four middle-fold challenges
Main	Sixteen tail coefficients	One local-fold challenge
Main	64-bit nonce	16-bit predicate and seed lanes
Opening	Derived source and middle positions	Authenticated leaf checks

Table 2: The complete causal authorization schedule.

The source leaves expose eight bank symbols and eight companion symbols per queried position. Each middle leaf exposes sixteen extension-field symbols. The canonical wire representation flattens a source leaf into 24 base-field limbs. Thus the source and middle payloads hold 1,560 base-field limbs and 1,040 extension-field symbols respectively. Canonical multiproofs share siblings across the derived query positions. The high index bits choose the ordered cap entry. The remaining bits select a path of height ten for source leaves or six for middle leaves.

The canonical ceiling is 92,696 bytes. Source and middle sibling ceilings are 453 and 193. The exact component sum is recorded in artifacts/encoding-costs.json. The decoder checks the bounded shape before deserialization, uses fixed-integer little-endian encoding and requires complete byte consumption. Measurements in Table 1 use this capsule boundary.

B Finite-list decoding

The encoding basis spans ordinary degree-bounded polynomials. For $0 \le j \le 7$, put

$$N_j = 65536/2^j, \quad K_j = 2048/2^j, \quad k_j = K_j - 1, \quad A_j = \lceil N_j/5 \rceil. \tag{27}$$

Both codes are analyzed over $\mathbb{F}_{2^{256}}$.

Theorem B.1 (Seventeen-candidate list)**.** *For any selected layer and distinct evaluation points, at most seventeen polynomials of degree at most k_j agree with a word in at least A_j positions. They are recoverable in randomized expected time polynomial in N_j and $\log |\mathbb{F}_{2^{256}}|$.*

Proof. Apply multiplicity interpolation [GS99] with multiplicity three, Y degree seventeen and weighted degree $D = 3A_j - 1$. Since $D - 17k_j > 0$, there are

$$U_j = 54A_j - 153k_j \tag{28}$$

coefficients against at most $6N_j$ Hasse constraints, with

$$U_j - 6N_j \ge 3N_j/160 + 153 > 0. \tag{29}$$

A nonzero interpolant exists. A candidate gives $Q(X, p(X))$ of degree at most D with $3A_j > D$ zeros counted with multiplicity. Thus $Y - p(X)$ is a factor, and the Y degree bounds the list by seventeen. The same decoder's polynomial-root step recovers and filters the candidates in expected polynomial time. □

Lemma B.2 (Base-field bank candidates)**.** *Bank words and evaluation points in* $\mathbb{F}_{2^{128}}$ *give candidate coefficients in* $\mathbb{F}_{2^{128}}$*. Companion coefficients remain in* $\mathbb{F}_{2^{256}}$*.*

Proof. For $q = |\mathbb{F}_{2^{128}}|$, let $p^{(q)}$ apply Frobenius to each coefficient. Then $p^{(q)} - p$ vanishes at all $A_j > k_j$ agreement points. It is identically zero, so every coefficient is fixed by Frobenius. □

The initial lists contain at most $17^2 = 289$ pairs. The exact layer ledger is in the artifact.

C Causal field simulation

In this field experiment, the source cap is sampled independently of the witness and fresh bank masks, with the prescribed shape. Challenges come from causal kernels on public history and private coins independent of the witness and fresh bank masks. Later kernels may depend on every preceding public message.

Lemma C.1 (Nonzero post-claim relation)**.** *For the selected geometry and every* $r \in \mathcal{S}^{11}$*,* $t \neq 0$*.*

Proof. Cell 944 is the first operand's dedicated pad and occurs in no other operand, boundary or Libra functional. Its coefficient in t is $k_0 = \prod_{i=0}^{10} r_i \neq 0$, since its batching power is one. □

Lemma C.2 (Constant Owner fibers)**.** *Fix a valid address table, an admissible causal challenge path and an accepting Owner view. Exactly* $|\mathbb{F}_{2^{128}}|^{1304}$ *assignments to the 1536 fresh bank cells produce this view. The view is parameterized by* μ*, 110 nonconstant round coefficients and five terminal operands, totaling 116 extension-field elements. The final mask value follows from the terminal identity.*

Proof. Five operands determine all ten pads by Theorem 4.1, fixing Q_B and its round polynomials. Each transmitted nonconstant coefficient then determines its Libra coefficient because $\lambda \neq 0$. The value μ fixes the sum of eleven constant coefficients, leaving ten extension-field constants. The valid table has $\mathcal{M}[Q_B](\rho) = 0$, so the telescope also gives the terminal mask identity without another constraint. The free cells are 20 constant-mask coordinates, 1024 source coins and 260 inactive/padding cells. Their total is 1304, independently of the valid table and transcript. □

Theorem C.3 (Causal field-view simulator)**.** *In this experiment, a witness-free simulator exactly reproduces all capsule field messages and at most 1024 distinct adaptive source-bank symbols. The selected 65-leaf opening exposes at most 520. The result holds for adaptive causal challenges and finite sequences of authorizations with independently fresh masks and adaptive statements.*

Proof. Sample μ uniformly in $\mathbb{F}_{2^{256}}$. After λ, sample the ten nonconstant coefficients in each round uniformly in $\mathbb{F}_{2^{256}}$, receiving its challenge after its message. After eleven rounds sample five uniform terminal operands and set

$$g_{\mathrm{fin}} = \lambda^{-1}(c_{\mathrm{fin}} + Q_{\mathrm{pub}}(r, Y)). \tag{30}$$

The public polynomial uses only r, Y and public data. The real and simulated probabilities of any complete path are both its causal-kernel probability times $|\mathbb{F}_{2^{128}}|^{1304}/|\mathbb{F}_{2^{128}}|^{1536} = |\mathbb{F}_{2^{256}}|^{-116}$. Summing equal path probabilities proves adaptive Owner equality.

After η, form t, b. Since $t \neq 0$, a uniform companion gives a uniform scalar σ. Sample it before γ. Then sample U uniformly from

$$\{U : \langle U, t\rangle = (1+\gamma)b + \gamma\sigma\}. \tag{31}$$

A nonzero coordinate supplies a pivot. The companion bijection makes this exactly the real conditional law for every bank in its Owner fiber, independent of that bank. Run the honest Main and middle-code algorithms on U with the causal challenge kernels. Their entire field view has the real law even if all of U is exposed.

The dedicated source coefficients remain independent uniform after Owner and U. Evaluation rank makes every new source symbol uniform conditional on all previously revealed distinct positions, up to 1024. Store repeated positions. Set companion symbols by

$$C(x) = \gamma^{-1}\big(U(x) + (1+\gamma)B(x)\big). \tag{32}$$

Invertible source normalization preserves equality. Middle symbols come from the codeword already computed from U. Finally condition on the prior public history at each authorization. Fresh masks let the same equality apply inductively. □

Appendix E realizes this independent-cap field view with the native common permutation and bounds quantum distinguishing error.

D Restoration through grouped folds

Let $\mathcal{C}_j$ be the layer-j code from Equation (27). For an adjacent pair and public twiddle ω, inverse butterfly and fold are

$$d = y_1 + y_0, \quad c = y_0 + \omega d, \quad F_\beta(y_0, y_1) = (1+\beta)c + \beta d. \tag{33}$$

For codewords, c, d encode even/odd message coordinates in $\mathcal{C}_{j+1}$. Conversely $y_0 = c + \omega d$, $y_1 = y_0 + d$ restore $\mathcal{C}_j$. Twiddles lie in $\mathbb{F}_{2^{128}} \subset \mathbb{F}_{2^{256}}$. Invertible within-leaf normalization commutes with the blend and recovers raw coordinates. Eight-symbol source and sixteen-symbol middle leaves therefore implement the consecutive folds below.

For length n and $w_i \in [0,1]$, put $\mu_w(A) = n^{-1}\sum_{i\in A} w_i$. Child weights $w'_i = (w_{2i} + w_{2i+1})/2$ preserve the mass when pulling a child set back to both parent points.

Lemma D.1 (One-fold restoration)**.** *For any fixed parent word W at layer $j < 7$ and weights w, outside at most E_{j+1} challenges, every child candidate $p \in \mathcal{C}_{j+1}$ whose full agreement set has $\mu_{w'}$ mass at least $1/5$ comes from a $P \in \mathcal{C}_j$ with $F_\beta(P) = p$. It agrees with W above the whole agreement set and preserves its mass, where*

$$E_j = \left\lceil N_j \frac{2h^5 + 3h(4/5)\rho_j}{3\rho_j^{3/2}} + \frac{h}{\sqrt{\rho_j}} \right\rceil, \quad \rho_j = \frac{K_j - 1}{N_j}, \quad h = \frac{17}{2}. \tag{34}$$

Rational upper bounds may replace the square-root expressions.

Proof. Apply list correlated agreement [BSCH+25, Theorem 4.6] to $c + \beta(c+d)$ after the inverse butterfly. Its conditions $\rho_{j+1} < 1/25$ and $9\sqrt{\rho_{j+1}} \leq 8/5$ permit multiplicity at most eight, bounded by the stated expression. Weighted mass at least $1/5$ gives ordinary density at least $1/5$. Outside the universal exception set, two child codewords agree with c, d on the same full set A. Their blend equals p since $|A|$ exceeds the child degree. The inverse butterfly restores P. The child-weight identity preserves mass. The exception set works simultaneously for every A, p, including choices made after β. □

The same theorem for $W_B + \gamma(W_B + W_C)$ gives E_0. Every candidate of density at least $1/5$ then comes from a source pair agreeing on the same set, and uniqueness identifies its blend.

Fix source words $W_B \in \mathbb{F}_{2^{128}}^{65536}, W_C \in \mathbb{F}_{2^{256}}^{65536}$ before Owner. Let $V_0 = (1+\gamma)W_B + \gamma W_C$ and V_3 its three folds. The prover fixes $M \in \mathbb{F}_{2^{256}}^{8192}$ before its four middle folds G_7. The sixteen revealed tail coefficients encode $T \in \mathcal{C}_7$. For source index i,

$$\mathsf{Good}(i) \iff V_3(i) = M(i), \quad G_7(\lfloor i/16 \rfloor) = T(\lfloor i/16 \rfloor). \tag{35}$$

Theorem D.2 (One source origin)**.** *With independent source-blend and seven-fold challenges from $\mathcal{S}$, outside probability*

$$\varepsilon_{\mathrm{prox}} \le 701201954824/2^{255}, \tag{36}$$

every completed prefix with good-index density greater than $1/5$ has an initial pair (P_B, P_C) whose blend folds to a middle codeword and then exactly to the tail. The bank, companion and middle candidates agree with their received words on density greater than $1/5$. There are at most 289 initial pairs.

Proof. Before middle challenges fix $w_i = \mathbf{1}[V_3(i) = M(i)]$. Averaging over four successive fibers makes the G_7, T agreement mass exactly the density of good source indices. Four backward applications of Lemma D.1 give P_M folding to T and agreeing with M on complete fibers. Its agreement with V_3 has density greater than $1/5$. Three more applications with uniform weights restore P_0 agreeing with V_0. Source correlated agreement gives $P_0 = (1+\gamma)P_B + \gamma P_C$.

Within each group its received word is fixed by the preceding scalar prefix and no prover move enters. Universal exceptional sets thus permit the final candidate and agreement-set choices. The union cost is $\sum_{j=0}^{7} E_j/|\mathcal{S}|$. Replace each denominator square root by $2^{-48}\lfloor 2^{48}\sqrt{\rho_j}\rfloor$ and round upward. The eight rational bounds sum to 701201954824, as recorded in the artifact. The finite-list and subfield lemmas give 289 pairs with the bank candidate over $\mathbb{F}_{2^{128}}$. □

Theorem D.3 (Public-coin source extractor)**.** *For arbitrary adaptive prover messages and independent public coins, an expected polynomial-time extractor with access to the complete fixed bank and companion source words recovers a preimage of a with joint error*

$$\begin{aligned} \Pr[\mathsf{Verify}\ \textit{accepts and extraction fails}] &\le 5^{-65} + \frac{701201954824 + 289(154+8)}{2^{255}} \\ &= 5^{-65} + \frac{701202001642}{2^{255}}. \end{aligned} \tag{37}$$

Its work is polynomial in 65536 *and* $\log|\mathbb{F}_{2^{256}}|$.

Proof. Decode the bank list, convert to message coordinates, validate address transitions and boundaries, and return a valid candidate's preimage. This is expected polynomial time. Failure makes all bank candidates invalid. Charge proximity exceptions. A good-index density at most $1/5$ passes 65 independent queries with probability at most 5^{-65}.

Otherwise fix a restored pair from the at most 289 initial pairs. Let s be Phase A's point, H the upper table and $H_U(z) = \widetilde{U}(z, s_8, s_9, s_{10})$. They are fixed before the source-fold group. Exact restoration and the upper-to-tail equality imply $(\widetilde{H} - H_U)(\beta_0, \ldots, \beta_7) = 0$. If this polynomial is nonzero, its total degree eight costs $8/|\mathcal{S}|$. Otherwise $\widetilde{H}(s_0, \ldots, s_7) = v$ gives $v = \widetilde{U}(s)$. The algebraic reduction then costs $154/|\mathcal{S}|$ for an invalid candidate. It allows adaptive messages and applies to the fixed candidate regardless of received-word errors outside its agreement set. Union over the initial pairs gives the stated bound. □

E Native common-permutation simulation

Write permutation states in native flat coordinates. Public field and wire conversions are fixed bijections. Let P_0 be the public width-four permutation implemented by the authority circuit, let $A = \{(r, c_{\mathrm{addr}}) : r \in \{0,1\}^{256}\}$, and fix its injection $F = P_0|_A$. The native oracle P is a uniform permutation completion of F. The authority circuit and spending witness are fixed independently of this completion. Source leaves, middle leaves, nodes and both transcript channels use this one oracle.

Put $s = 2^{256}$, $N = s^2$ and $M = N - s$. The representation $P = P_0 \circ R$, where R fixes A and is uniform on its complement, gives exactly the $M!$ completions. Membership of an output in $F(A)$ is tested by P_0^{-1}. Each free native call costs one query to R with public coordinate gates. This also specifies coherent inverse access. Ordinary evaluation of the authority circuit is unchanged.

Parameters. The source has $n = 8192$ leaves and $u = n - 8 = 8184$ nodes below its ordered cap. Set $v = n + u = 16376$, $m = 11n + u = 98296$, $L = 128$, $\overline{L} = 129$ and $e = 2^{-210}$. Define

$$\overline{C} = (\overline{L} + 1)\binom{m}{2} + \overline{L}m^2 + (\overline{L} + 1)mn, \tag{38}$$

$$\delta = n\left(\frac{129}{s} + \frac{20}{N}\right) + \frac{2u}{N}, \tag{39}$$

$$h_A = \frac{128n}{\sqrt{s}} + \frac{21n + 2u}{s}. \tag{40}$$

Here $\overline{C} = 1979123020424$, $\overline{C}/s < 2^{-214}$ and $h_A < 2^{-107}$. The global capacity event used below has failure probability at most e.

Lemma E.1 (Capacity loads and source entropy)**.** *Except with probability e over the completion, for each input capacity $r \neq c_{\mathrm{addr}}$ and output capacity t, at most 128 permutation pairs have these two capacities. On this event, conditional on the field view, full U and at most 520 opened raw bank positions, an unopened leaf's second native input has maximum point probability $128/s$.*

Proof. For $k = 129$, the capacity-load union bound is at most

$$N\binom{s}{k}\frac{(s)_k}{(M)_k} \leq \frac{N}{k!}\left(\frac{N}{M - k + 1}\right)^k \leq \frac{2N}{129!} < 2^{-210}.$$

The last inequalities follow from $129! > 2^{723}$ and $N^{129} \leq 2(N - s - 128)^{129}$.

Include all eight raw positions of the unopened leaf. Evaluation rank on $520+8 = 528 \leq 1024$ positions and invertible within-leaf normalization give uniform independent normalized bank coordinates B_0, B_1. The companion equation commutes with normalization, so the first two rate blocks are $(B_0, \ell(B_0))$ and $(h(B_0), B_1)$ for fixed functions ℓ, h. A prescribed second input capacity permits at most 128 first inputs. Each fixes B_0, and a prescribed second rate then fixes B_1. There are at most 128 successful pairs among s. For two leaves use all sixteen raw positions, giving rank $536 \leq 1024$. □

Exact refresh. For a uniform completion of a partial permutation h, a new input x can be answered by drawing y uniformly outside its recorded image and replacing P by $(P(x)\ y) \circ P$. The result remains uniform subject to h and $x \mapsto y$: every resulting completion has exactly $N - |h|$ equiprobable predecessors. Existing inputs retain their recorded values. Starting with $h = F$ gives the prescribed-slice version of this identity.

Sample v independent 256-bit source vertex labels, including the eight cap labels. At every leaf retain its true first call. Propose ten fresh full outputs for calls two through eleven and the assigned leaf label with a fresh capacity for call twelve. These outputs fix all eleven suffix inputs. At a node with child labels a, b and parent label z, propose

$$x = (a, b \oplus v_{\text{node}}), \qquad y = (z \oplus a, c), \qquad c \leftarrow \{0,1\}^{256}.$$

Its native feed-forward digest is z. Each bottom-up parent label is first used at this step, so its proposed full output is uniform. The exact refresh identity therefore gives the honest joint forest and permutation law whenever the proposals pass the finite fences.

For all m proposed pairs require distinct inputs and outputs, preservation of the first source pairs, and $P(x_i) \neq y_j$ for all i, j. The core constructor queries only $z_i = P^{-1}(y_i)$. These fences are the equivalent disjointness tests on $X = \{x_i\}$ and $Z = \{z_i\}$. On success,

$$\rho = \prod_i (x_i \ z_i), \qquad P_* = P \circ \rho$$

realizes every proposal and preserves both directions of each protected pair. Its core uses one inverse query per pair. Consistency rechecks are charged separately. Require $X \cup Z$ to avoid A. Then the entire address slice is preserved.

Finite estimates. On the capacity event, active-input pairs cost at most $\overline{L}\binom{m}{2}/s$, output pairs cost $\binom{m}{2}/s$, cross-image pairs cost $\overline{L}m^2/s$, and comparisons with first source calls cost $(\overline{L}+1)mn/s$. The second-source case uses Lemma E.1. For a leaf's second input to equal its own first input, the capacity return permits at most 128 values of B_0, and the second rate fixes B_1. Later inputs have fresh full predecessor states. Distinct nodes have disjoint child variables. For an adjacent source equality, the uniform full predecessor has at most $128s$ admissible off-address-row states and at most s address-row states. Its probability is therefore at most $129/s$. These counts sum to $\overline{C}/s$.

A first output visits the address capacity with probability at most $128/\sqrt{s}$ after field conditioning. The remaining slice tests have fresh capacities or full states, and cost at most $(21n + 2u)/s$. This gives h_A. These tests also exclude prescribed outputs in $F(A)$.

For a fixed point, one source suffix contributes at most $129/s + 20/N$ to either direction's changed support: its second input costs $128/s$, its terminal proposed output costs $1/s$, and its ten other inputs and ten internal outputs each cost $1/N$. A node contributes $2/N$. Fixed permutation bijections preserve these point bounds. Their union gives δ.

Before field exposure, B_0, B_1 and both companion limbs are independent. The first input rate is uniform on s points, and a specified second input fixes another independent pair of limbs. Thus its point mass is at most $128/N$. The pre-field support bound $[n(L + 21) + 2u]/N$ is at most δ. This justifies the prefix comparison before conditioning on its later field view.

The retained simulator. At request start mark its identifier in J. Sample the source labels and run the field-view simulator with the exact native channels, middle tree and nonce search under the current base interface. After the positions are fixed, sample the required raw bank symbols, store repetitions, and derive companion symbols from U. Construct only their native suffixes and authentication paths. At most $11 \cdot 65 + 10 \cdot 65 = 1365$ active pairs suffice.

The local fences use retained data and protect all observed public computation pairs. Failure aborts the complete experiment. Consistent constructor pairs are rechecked before installation. The simulator emits the canonical capsule and retains its public pairs and input adapter. Later requests use the composed adapter. All parsed duplex prefixes bind the current request identifier. Commitment records are handled by their native types and dependency guards.

Theorem E.2 (Adaptive native capsule simulation)**.** *For S adaptive classical attempts and a common bound T on charged forward and inverse calls, the witness-free simulator has event error*

$$\varepsilon_{\mathrm{sim}}^{\mathrm{Q}}(S,T) \le \min\left\{1, 4ST\sqrt{\delta} + 4S\left(\overline{C}/s + T\delta + e + h_A\right)\right\}. \tag{41}$$

Its work is expected polynomial in the code length and charged proving and nonce-search work.

Proof. Replace attempts in reverse order, giving each selected attempt a real prefix and an arbitrary continuation containing later simulations. Presample its source coins before that prefix. The refresh identity and fences couple its full forest to the honest law with error $\overline{C}/s + e + h_A$.

Replace the prefix's programmed oracle by its base oracle. Telescoping with unedited prefixes costs $2T\sqrt{\delta}$: for any fixed query state, the average mass on the hidden support is at most δ. The request and prefix history now use the base oracle before field conditioning. Protect the subsequent field, middle and transcript calls. Their collision probability is at most $T\delta$.

Retain the opened suffixes, path nodes and frontier labels. On every raw tape define the auxiliary permutation $P_R \circ \rho_U$, including failed full fences. On full success its disjoint swaps equal P_*. Conditional on the base, full U, field view, opened values and retained states, both child labels of every omitted node remain uniform. An unopened leaf uses rank 528. Two-leaf comparisons use rank 536. The same δ bounds all unexposed endpoints. Deleting their swaps costs another $2T\sqrt{\delta}$, with the finite failure terms charged separately. The retained guard depends only on retained data, so this argument uses unnormalized raw tapes throughout.

After deletion, omitted first calls are unused deterministic reads under the fixed base and can be erased. The remaining view and at most 520 opened bank values have exactly the field simulator's law by Theorem C.3. No unopened bank or witness value is needed. The displayed constants cover both refresh and pruning endpoints, protected calls and slice tests. Sum over S. □

F Round-by-round source knowledge

The source IOP has knowledge error

$$\kappa = \max\{5^{-65}, 701202001931/2^{255}\}. \tag{42}$$

Its 30 verifier moves use independent public coins. The ρ, source-fold and middle-fold groups are atomic, as is the final query-seed vector. Grinding belongs to the non-interactive compiler. Source and middle oracles become immutable at their prover moves. Missing leaves have canonical zero completion and queried missing leaves reject.

For round-by-round knowledge [CMS19, Definitions 8.3,8.5], a state function starts at zero, cannot change zero to one at a prover move, and is one on accepting complete transcripts. At any zero-state verifier prefix, transition probability above κ must make its polynomial-time prefix extractor return a witness. The state function may be inefficient.

Decode the immutable initial lists $\mathcal{J} = \mathcal{L}_B \times \mathcal{L}_C$, with $|\mathcal{J}| \le 289$. At each prefix containing the source oracle, the extractor returns the first address-valid bank's preimage. Each candidate's message coordinates, Q_B and g are fixed.

Record only the following verifier-step flags:

1. At ρ, an address-valid bank occurs in $\mathcal{L}_B$.
2. For any initial pair, a nonzero algebraic discrepancy vanishes at its fresh challenge. These are the discrepancies of Theorem 5.1, plus the λ discrepancy $\mathcal{M}[Q_B](\rho) + \lambda(\mathcal{M}[g](\rho) - \mu)$.
3. At γ or a fold challenge, membership in its universal list-correlated-agreement exception set.

4. For a pair, the nonzero upper discrepancy $D(z) = \widetilde{H}(z) - \widetilde{U}_\gamma(z, s_8, s_9, s_{10})$ becomes identically zero after a 3, 4 or 1 variable group.

5. The complete query-seed move makes the verifier accept.

State is one after its first flag. Permanently malformed transcripts have state zero and reject every continuation.

Per-pair discrepancy	Degree budget
ρ transition table	11
λ masked target	1
Eleven Owner rounds	110
η post-claims	10
γ blend	1
Eleven Phase-A rounds	22
Upper collapses in $3 + 4 + 1$ groups	8
Total	163

Table 3: Nonzero-discrepancy root budgets.

Lemma F.1 (State-function properties)**.** *This function is zero initially, cannot change zero to one on a prover move, and is one on every complete accepting transcript.*

Proof. Flags use only data already fixed at their verifier step and remain recorded afterward. Universal source exceptions are independent of later middle messages. No prover move can create a flag. The last flag records acceptance. □

Lemma F.2 (Zero-state transitions)**.** *At any well-formed zero-state verifier prefix, the extractor returns a preimage or the chance of state one is at most κ.*

Proof. At ρ, a valid bank is extracted. Otherwise its first flag has probability zero. Later zero-state prefixes also contain no valid bank, since it would have flagged at ρ. Each nonzero discrepancy has the stated univariate or total degree. For an upper collapse, choose one nonzero coefficient in its remaining-variable expansion. Its dependence on the next group has degree at most that group's size. Identical collapse requires a root of this coefficient, giving $3 + 4 + 1 = 8$. Thus each initial pair costs at most $163/|\mathcal{S}|$.

At source-line and fold moves, universal list-correlated-agreement sets apply to future candidates and weights. Within atomic groups, condition on each scalar prefix. Summing the field budgets bounds any one verifier move by

$$\Pr[\text{first flag at this field move}] \le \frac{701201954824 + 289 \cdot 163}{|\mathcal{S}|} = \frac{701202001931}{2^{255}}. \tag{43}$$

Before the last move both oracles, fields and tail are fixed. Failed field checks reject. If good-index density exceeded $1/5$, source restoration would give an exact initial-pair chain. The upper check forces $D(\beta) = 0$, and absence of a collapse flag forces $D = 0$. Thus Phase B supplies the true candidate value. The backward argument of Theorem 5.1, with all its round, blend, batching and ρ coincidences excluded by these flags, forces an address-valid candidate. This contradicts zero state. The good density is therefore at most $1/5$ and all 65 independent queries pass with probability at most 5^{-65}. □

Theorem F.3 (Round-by-round knowledge)**.** *The selected IOP has round-by-round knowledge error κ and an expected polynomial-time source extractor.*

Proof. The preceding two lemmas give the state-function and extraction conditions. Candidate decoding and validation are polynomial time. □

G Native fresh-request extraction

The native simulator and extractor share one compressed base permutation and the classical input permutation σ of Appendix E. The visible oracle is $P \circ \sigma$. Its address slice stays fixed. Use s, N, M, H, c, g and the finite profiles of Appendix H, and the full-forest bounds n, u, v, m and endpoint bound δ of Appendix E. An identifier enters the requested set before its first simulator call. Blank or inconsistent classical answers and constructor failures abort the shadow experiment. All dictionary rules use effective native coordinates.

Let T bound all adversary, simulator and verifier base calls, including inverse constructors, pair rechecks and each squeeze's advance. Private lookahead moves the actual response circuit and adds no calls. Assume $3T - 1 \le M$ and define

$$\begin{aligned} a_H &= (H+1)\Lambda_H, \qquad K_H = (H+1)^2\Lambda_H, \\ \beta_T &= 2T + 2c + 3 + 3a_H(5T+2) + 2a_H + \Sigma_H(T), \\ m_0 &= M - T - \beta_T s, \qquad c_Y = v(2K_H + H + 1), \\ m_* &= m_0 - c_Y s \ge N/2, \qquad r_* = \sqrt{2/m_*}, \\ \epsilon_T &= 4\sqrt{\beta_T s/(M-T)}, \qquad \epsilon_Y = 4\sqrt{c_Y s/m_0}, \\ \chi_T &= \frac{8}{M-T-1} + 8\sqrt{\frac{2}{M-T}}, \qquad X_T = \chi_T + 2\epsilon_T. \end{aligned} \tag{44}$$

These inequalities hold for the stated uniform-slice profiles and $T \le 2^{220}$. All probabilities below are truncated at one.

Protected queries and causal words. For a row-deleted database h, compression swaps absence with the uniform available extensions. Replace this reflection by the one using only allowed extensions and fixing forbidden extensions. For a bad static base impose no further static restriction. Monotonicity of the static predicate makes its good and bad sectors invariant.

The additional rules exclude a new rooted forward output capacity at a root, the address capacity, an old reachable vertex or an existing edge's input capacity. Every new rooted inverse input is excluded. Form $W(h)$ from zero and each record's two node children, node key, rate key and decoded rooted absorb block. Include one externally supplied last absorb block. Then $|W| \le 5|h| + 2$. Node keys and rooted rate keys avoid $W^{-H} = \{a : g^j(a) \in W \text{ for some } j \le H\}$, whose size is at most $a_H|W|$. Rooted output capacities also avoid this set. Every rooted forward query additionally protects its own current decoded block against g^j of its node and rate keys. The latter rule costs at most $2a_H s$ extensions. The static exclusions cost at most $s\Sigma_H(T)$. Counting rate or capacity slices bounds all forbidden extensions by $\beta_T s$ in either direction.

If a fraction p is forbidden, the old and allowed uniform vectors have inner product $\sqrt{1-p}$. Their compression reflections differ by at most $2\sqrt{p}$, and their two query sandwiches by at most $4\sqrt{p}$. Thus each protected query differs from an ordinary query by at most ϵ_T, including coherent input and direction controls.

Before a candidate measurement use the global rules independent of that candidate. Its external block and capacity shield are added only after measurement. The bound β_T covers both versions. For a measured logical challenge input x, protect its capacity x_C. Let E_x contain

the parsed prefix, all committed source and middle words, upper and tail messages, and their finite dependency records and missing pointers. The external block decoded from x_R belongs to W. Its predecessor edge is compression-fixed by the protected capacity. Earlier edges are fixed by descendants. Cap and child pointers fix existing node rows through W^{-H} and forbid new preimages. Rooted leaf ancestors freeze recursively. Both literal Owner bridge blocks fix its close path. Hence every spectral projector of E_x commutes with every protected query. Accessible gates commute as well.

The current-block rule handles a challenge row that supplies its own last cap through a fixed-slice orbit. At exponent zero a node alias would force its rate to equal the predecessor rate, and a leaf alias would force its node key to that rate. Positive exponents are excluded by the current-block orbit tests. Thus removing that challenge row preserves its earlier completion. The typed guard permits literal same-pair leaf/node roles throughout this argument.

Trusted updates. Put $d = 11n$ and define the following finite fresh-tape bounds:

$$\begin{aligned} s\eta_Y &= vR_H + \big(\tbinom{v}{2} + 2Tv\big)K_H + v(T+5)a_H, \\ s\eta_{\mathrm{sec}} &= dR_H + \big(d(2T+v) + \tbinom{d}{2}\big)K_H + d(T+5)a_H, \\ s\eta_{\mathrm{cap}} &= d(7T+7+d+u), \\ s\eta_Z &= d(2T+12n+u)(H+1), \\ s\eta_{\mathrm{move}} &= T(3n+m) + T(10n+u)/s, \\ s\eta_{\mathrm{ptr}} &= (12n+u)a_H(5T+2), \qquad \eta_{\mathrm{end}} = 12\delta. \end{aligned} \tag{45}$$

All terms are evaluated before restricting to guard success. Set

$$\begin{aligned} b_{\mathrm{stat}} &= 2\big(\sqrt{\eta_Y} + \sqrt{\eta_{\mathrm{sec}}} + \sqrt{\eta_{\mathrm{cap}}} + \sqrt{\eta_Z} + \sqrt{\eta_{\mathrm{move}}}\big) + 2mr_*, \\ b_{\mathrm{pref}} &= b_{\mathrm{stat}} + 2\sqrt{\eta_{\mathrm{ptr}}} + 2\sqrt{\eta_{\mathrm{end}}}, \\ B_{\mathrm{stat}}(S,T) &= T(\epsilon_Y + 2r_*) + Sb_{\mathrm{stat}}, \\ B_{\mathrm{pref}}(S,T) &= 32T\epsilon_Y + 22TX_T + 2Tr_* + Sb_{\mathrm{pref}}. \end{aligned} \tag{46}$$

Lemma G.1 (Trusted static and causal preservation)**.** *The static bad-component amplitude increases by at most B_{stat} over all trusted updates. Preserving one measured E_x, its physical candidate and at most eleven private response inputs costs at most B_{pref} in state distance, in addition to ordinary protected-query replacement and one ordinary private-query crossing.*

Proof. On an allowed fiber, write a vector as $\alpha|\ \perp\rangle + \gamma$. For at most k nonblank answers, compression gives the effect bound

$$CQ_kC \preceq 2Q_k + 4k/m_*\, I :$$

its selected component is $Q_k\gamma + (\alpha - \langle \mathrm{unif}|\gamma\rangle)Q_k|\mathrm{unif}\rangle$. This proves the inequality by squared triangle inequality. After the field view and U, the first two normalized source bank coordinates are independent uniform 128-bit values by the rank lemma. Their blocks are $(B_0, \ell(B_0))$ and $(h(B_0), B_1)$ for fixed functions. The first coordinate selects at most one old rooted first row for a prescribed capacity. The second remains unread. Thus the second source input has point mass at most $2/s + 4\sqrt{s}/m_* \le 3/s$. An inverse constructor output has a fresh capacity at its first use. Its input point mass is at most $2/s + 4/m_* \le 3/s$. The bounds hold for arbitrary compressed states and masked sets of old rows.

Fresh source labels have at most η_Y orbit and old-target shock. During ordinary calls, additionally excluding anchors whose forward orbits meet their label closure costs ϵ_Y per call. Rooted output capacities also avoid its translates by the node IV. These remove at most $c_Y s$

extensions. Static-bad bases add no such restriction. The resulting good sector is invariant. Internal source rates are independent uniform values. Their eleven successive node keys form an affine chain in which any pair has rank two, including the last key with its fixed terminal label. Self returns, pair intersections and old targets give η_{sec}. Independent capacities give $\eta_{\mathrm{cap}}, \eta_Z$. These inequalities follow respectively by counting inverse fibers with Λ_H, returns with R_H and orbit intersections with K_H. The displayed formulas retain every term.

Before a forward call remove its input from an old-row test, and before an inverse call remove its output image. The test shrinks and the query preserves its remaining row-deleted or image-deleted fiber. This carries the initial η_{move} bound through cap-dependent calls. Desired images are freshly sampled after ordinary first calls. Delete those images in the old-row test during construction. Its owned inverse rows are then separate from old collisions. On explicit records, the guarded input swaps insert only the active pairs. A measured nonblank answer has correct-row probability $(1-1/m_a)^2$, where $m_a \geq m_*$. Recording every ordinary forward pair, measured constructor inverse answer and owned private pair recheck costs at most $(q_S+m)r_*$ per attempt. Inconsistent repeats abort. The factors two in (46) cover both comparison directions.

The frozen cone of E_x makes each primary-label equality test commute with later ordinary calls. Secondary keys are tested at their first fresh-rate sampling. This gives η_{ptr}. The point bounds above and fresh later source states give $\eta_{\mathrm{end}} \leq 12\delta$ for the candidate and response inputs, conditional on their exposure. Outside these tests the update fixes their physical base inputs and the full present and missing dependency cone.

For the late private response circuit L, let B be an attempt's q_S ordinary protected calls and $\delta_B = \ell q_S X_T$. On endpoint-good private outcomes, exact relabel covariance intertwines L and the update T_ρ. Their bad norms are equal by isometry. The endpoint bound applies in the early order BL. Transfer to LB costs δ_B. Hence $\|[L, T_\rho B]\| \leq 3\delta_B + 2\sqrt{\eta_{\mathrm{end}}}$ on the compared state. The compiler counts one crossing. The $22TX_T$ term pays the other two for $\ell \leq 11$, with slack for one early protected L reference. The endpoint argument uses the base protected calls. The all-label shadow is required for the static and cone certificates. Its two-order cost is bounded by $32T\epsilon_Y$. Summing attempts and recording costs proves the stated envelopes.

A fresh candidate measured inside an attempt also satisfies this timing. On the explicit forward-record component, source and middle sponge inputs have their own rooted types. Middle node inputs are unrooted: their right child is an anchor avoiding translated reachable capacities under the static guard. Inverse constructor pairs and their owned rechecks are unrooted outside η_{move}. Only simulator duplex calls can parse a valid challenge prefix, and their identifier is already requested. Their fresh terminal branch is zero. Thus every continuing fresh candidate precedes the sampling of the later attempts' coins. □

Measured native responses. Let $\mathcal{A}$ be the forger, its requests and native prefix replay, stopping before a response group's first permutation call. For physical labels $u, v \notin A$, let T transpose database rows u, v. Let M transpose these labels in a forward input register or an inverse answer register, fixing the blank symbol. Put

$$K_u = (|v\rangle - |u\rangle)\langle u|/\sqrt{2}, \qquad K_v = (|u\rangle - |v\rangle)\langle v|/\sqrt{2}.$$

For the q-query computation $\psi = \mathcal{A}|0, \varnothing\rangle$, let ψ_i precede query O_i and $\mathcal{A}_{>i}$ follow it. Exact compressed covariance gives $\mathsf{T}O_i\mathsf{T} = \mathsf{M}O_i\mathsf{M}$. Since $\mathsf{M}K_h = -K_h$, commutator telescoping gives

$$\begin{aligned} \phi_{i,0,h} &= -\mathcal{A}_{>i}O_i\mathsf{T}K_h\psi_i, \qquad \phi_{i,1,h} = \mathcal{A}_{>i}\mathsf{T}K_hO_i\psi_i, \\ \psi &= \mathsf{T}\psi + \sqrt{2}\sum_{i=1}^{q}\sum_{h\in\{u,v\}}\sum_{e=0}^{1}\phi_{i,e,h}. \end{aligned} \tag{47}$$

Indeed $[O_i, \mathsf{T}] = \mathsf{T}(\mathsf{M}O_i(\mathsf{M}-I)+(\mathsf{M}-I)O_i)$ and $\mathsf{M}-I = \sqrt{2}(K_u+K_v)$. This identity holds before protected-query replacement. The empty initial database is row-relabel invariant. Multiplying by

any terminal success contraction, including a database-dependent effect, and applying Cauchy–Schwarz gives coefficient $8q + 1$.

A hit measures u before or after the selected query and then samples v uniformly. A miss measures v and then samples u uniformly. Their Kraus squares sum to at most the identity after uniform averaging. Thus all replacement costs are direct-sum bounds without an M factor. Freeze the current table and set $x = \sigma_{\text{past}}^{-1}(u)$. For hits and subsequent whole attempts, the trusted envelope preserves $\sigma_{\text{future}}^{-1}(u) = x$. The current-partial-attempt miss is rebased at completion as described below. A hit relabeling preserves E_x except on squared mass

$$\iota_T = \frac{T + [2T + c + 1 + a_H(5T + 2) + \Sigma_H(T)]s}{M}. \tag{48}$$

The count excludes an old row at v, old source/reachable capacities, a new inverse-orbit preimage for a frozen pointer, and a static bad inverse extension at the old output after deleting the irrelevant row. Capacity and orbit tests use $\sigma_{\text{past}}^{-1}(v)$, uniform on the free inputs because the past table is a fixed bijection. Its copy distance is $2\sqrt{\iota_T}$. When row v is absent, row u is absent after the swap. Move the actual private response left across the suffix, stopping after this swap. It starts cold at u.

In a miss the past inverse table is a fixed bijection before the fresh physical draw. Logical x_C is reachable or an old source capacity with probability at most $\mu_T = (2T + c + 1)s/M$. Outside this set the swap moves at most one row to an unrooted input and removes its old row. The rooted graph loses edges. The capacity shield and trusted updates keep this prefix absent. This argument uses the fixed past table, while the shield fixes its future value.

A miss inside an unfinished attempt uses its current label shadow. Instrument each classical simulator call with a retained intended-input copy and reject a changed query-input wire. Oracle queries preserve that wire, so the instrumented real experiment is identical. A forward cat's wrong-input component rejects, including owned rechecks. An inverse constructor starts with a blank answer, so its before-query term is zero. In its after-query term, the stored answer can be u. The explicit owned row is then $u \mapsto Y_i$, and installation maps that physical target to the active source or unrooted-node input X_i.

The current labels Y precede this draw. Moving the selected row to fresh u adds a static shock of at most $s\Sigma_H(T)/M$ and a label-orbit shock of at most vK_Hs/M. Their copy costs are bounded by $\epsilon_T + \epsilon_Y$. Restoring the label guard for the current remainder costs at most $T\epsilon_Y$. The two copy orderings require at most four whole-circuit label-shadow replacements, so this extra $2T\epsilon_Y$ fits the $32T\epsilon_Y$ envelope. The static shock fits the unused miss stream reserve $22TX_T$.

A stale continuation may become unrooted. The target capacity shield and exclusion of reverse attachment prevent that continuation from rooting this particular candidate. Keep the capacity shield during ordinary remaining calls. The guarded program adds source-typed or unrooted source rows and unrooted nodes, never duplex-rooted rows, without rooting an old tail. If it touches the physical target, the resulting logical input has source or owned-unrooted type. Otherwise its duplex prefix remains absent or has the current requested identifier. At completion, rebase to $\sigma_{\text{end}}^{-1}(u)$ and reset its capacity shield. This role argument also covers a moved owned recheck row and a malformed current source capsule. No fresh duplex branch results. Only subsequent whole attempts use the fresh-coin preservation argument.

For a response of $1 \le \ell \le 11$ raw blocks, define

$$\begin{aligned} d_\ell &= M - T - \ell + 1, \qquad \Gamma_\ell = (N/d_\ell)^\ell, \\ \eta_\ell &= T/M + (\ell-1)(T+s+\ell+1)/d_\ell, \\ A_\ell &= (T+\ell)\epsilon_T + \ell(T+1)X_T + 2\sqrt{\iota_T} + B_{\mathrm{pref}}(S,T) \\ &\quad + \sqrt{\ell\eta_\ell} + \ell\sqrt{2/(M-T)}, \\ h_\ell(k) &= (\sqrt{\Gamma_\ell k} + A_\ell)^2, \\ z_T &= (\sqrt{\mu_T} + T\epsilon_T + B_{\mathrm{pref}}(S,T))^2, \\ E_\ell(k) &= (8T+1)\big((2T+1)h_\ell(k) + 2Tz_T\big). \end{aligned} \tag{49}$$

Lemma H.6 supplies the full joint density and unflagged cold-query cost on arbitrary post-measurement database states. Lemma H.5 gives X_T for each ordinary private/protected crossing. The first term of A_ℓ pays all ordinary replacements and private switches. The trusted envelope pays the extra update crossings. State distances are added before squaring. Equation (47) has one output term, $2q$ hits and $2q$ misses. Applying their bounds and $q \le T$ proves $E_\ell(k)$, including the output case with the same cold density.

Extraction and the first flag. Put $b = 701202001931/2^{255}$ and $p = 5^{-65}$, and define

$$\begin{aligned} \mathcal{U}(S,T) &= 26E_1(b) + E_{11}(b) + E_3(b) + E_4(b) + E_4(p), \\ G(S,T) &= \left(T\sqrt{\frac{8s\Sigma_H(T)}{M-T}} + T\epsilon_T + B_{\mathrm{stat}}(S,T)\right)^2, \\ D_M(T) &= \frac{120T}{\sqrt{M-T+1}} + \frac{T(T+1)}{M-2T+1}. \end{aligned} \tag{50}$$

Let $\delta_F = 0$ for a fixed primitive satisfying the finite profiles. For the uniformly injected-slice setup, δ_F is the exceptional probability in Theorem H.4.

Theorem G.2 (Native fresh-request extraction)**.** *From one fresh accepted capsule in the complete native simulated experiment, an expected polynomial-time compressed-oracle extractor recovers an authority preimage with joint loss at most*

$$\varepsilon^{\mathrm{Q}}_{\mathrm{ext}}(S,T) \le \min\left\{1, \delta_F + D_M(T) + \frac{2T^2}{M-T} + G(S,T) + \mathcal{U}(S,T)\right\}. \tag{51}$$

More precisely, ordinary native fresh acceptance is at most compressed preimage-recovery probability plus this bound. The simulator and extractor share one compressed state and its table history.

Proof. Run the complete simulator, native verification and ordered pair rechecks, then measure the base database once and apply the final input relabeling. Lemma H.1 constructs both typed words by finite forward closure and fixed-height lookup. Give missing raw positions canonical zero completion, decode the source word's two lists and validate the candidate banks. Each list has at most seventeen entries. A validated bank's input lanes give $\mathsf{Addr}(s) = a$. This gives the stated expected polynomial-time extractor.

The accessible comparison with ordinary random-completion queries costs $D_M(T)$ [CM26, Corollary 4.3]. Let q include verification and let r off-slice pairs be rechecked, so $q + r \le T$. For their truth and record projectors M_i, Q_i, the ordered estimate [CM26, Equation (83)] gives

$$\|(M_r \cdots M_1 - Q_r \cdots Q_1)L_{\le q}\| \le r\sqrt{2/(M-q)}.$$

Truth projectors may be noncommuting. Accepted matching rechecks with a missing explicit pair cost at most $2r^2/(M-q) \le 2T^2/(M-T)$. The finite static bad-amplitude sum is at most $T\sqrt{8s\Sigma_H(T)/(M-T)}$. Its base protected-history comparison costs $T\epsilon_T$ and trusted updates add B_{stat}, yielding G. On the remaining event every checked opening has the value of its reconstructed final word.

Apply the source-IOP state function to the accepted transcript and these final words. Failed bank extraction makes its initial valid-bank flag false. Its first zero-to-one flag occurs at one of the 30 verifier moves. This is a post hoc partition. The extractor need not compute the possibly inefficient flag tests.

For each terminal first-flag effect, use (47). In its protected experiment, E_x is copied before the fresh response and stays fixed through all later requests, verification and pair rechecks. Source and companion words are included from ρ onward. The middle word is included from its commitment onward. Earlier flags use universal exception sets independent of future middle words. Thus the committed final words in this experiment equal the causal snapshot. Late verifier calls cannot fill its missing pointers. The ordinary-to-protected and trusted-update distances are precisely those in (49).

At a field move the zero-state transition density is at most b by Theorem F.3. The native two-to-one trace-one decoder preserves the uniform $\mathcal{S}$ product distribution. At the final move, both words, upper and tail are fixed. Accepted openings force every index into their good-index set. On failed extraction and zero state, restoration bounds its size by 1,638 of 8,192. The 65 disjoint 13-bit windows therefore have raw independent-input density at most p. Grinding only removes outputs. The cold lemma transfers these complete group densities. There are 26 scalar field moves, groups of eleven, three and four field outputs, and the final four-block query group. Their first-flag union is exactly bounded by $\mathcal{U}$.

Adding accessible comparison, records, finite guard and causal union proves the result. The primitive setup event δ_F is charged once. Every recovered secret is checked against the authority circuit. Combining this extractor with the native simulator and semantic intent binding gives Theorem 5.3. □

H Native words and compressed permutation queries

Use the flat coordinates and fixed slice F of Appendix E. The three sponge capacities for `CAPSMIX_`, `CAPS256_` and `KSCH256_` are distinct and outside the address slice. Their number is $c = 3$. The native node is

$$N(a,b) = \operatorname{rate} P(a, b \oplus v_{\mathrm{node}}) \oplus a. \tag{52}$$

Put $b_* = c_{\mathrm{addr}} \oplus v_{\mathrm{node}}$ and $g(a) = \operatorname{rate} F(a, c_{\mathrm{addr}}) \oplus a$. Then $N(a, b_*) = g(a)$. Let $H = 10$ and define the finite profiles

$$\Lambda_H = \max_{j \le H, z} |(g^j)^{-1}(z)|, \tag{53}$$

$$R_H = |\{a : \exists\, 0 \le i < j \le H,\ g^i(a) = g^j(a)\}|, \tag{54}$$

$$D_H = \max_{w \ne 0} |\{a : \exists\, i, j \le H,\ g^i(a) = g^j(a \oplus w)\}|, \tag{55}$$

$$\Sigma_H(t) = 16(t+3)(H+1)^2\Lambda_H + 2R_H + D_H. \tag{56}$$

These quantities depend on the public primitive slice and fixed height.

H.1 A finite typed dictionary

A free permutation record (x, y) gives the capacity edge $x_C \to y_C$ and node anchor $y_R \oplus x_R$, with children $(x_R, x_C \oplus v_{\mathrm{node}})$. Reachability starts at the three sponge capacities. Exclude a reachable address capacity, a reachable edge ending at a root, and two distinct reachable edges

with the same endpoint. Every reachable vertex then has a unique rooted path. A twelve-edge source path or sixteen-edge middle path supplies its typed leaf anchor y_R and recovers rate blocks by

$$m_j = (x_j)_R \oplus (y_{j-1})_R, \qquad (y_0)_R = 0.$$

Undo the public coordinate maps and within-leaf normalization to recover its raw symbols.

Each row supplies a node anchor and at most one typed leaf anchor. Identify the two roles of the same row when $x_R = 0$. This is the literal equality between its leaf digest and its feed-forward digest. Let $\mathcal{A}(D)$ be the resulting at most $2|D|$ anchor classes. The finite guard additionally requires the map

$$(a, j) \longmapsto g^j(a), \qquad (a, j) \in \mathcal{A}(D) \times \{0, \dots, H\},$$

to be injective, and its image to avoid

$$Z_D = \{0, b_*\} \cup \{r \oplus v_{\text{node}} : r \text{ is a reachable capacity}\}.$$

Its failure predicate B_{stat} is increasing. Added rows add anchors and reachable paths. The fixed same-row identification merges a newly supplied role with that row's existing class.

Lemma H.1 (Typed native reconstruction)**.** *On a guarded database, a cap and index determine at most one source payload and at most one middle payload. Every accepted opening whose free pairs occur in that database has the reconstructed payload. The dictionaries and all zero-completed words are computable in polynomial time using forward evaluations of g.*

Proof. Every checked sponge path stays outside the address slice, since a first visit would make its capacity reachable. Its terminal rate is its typed leaf anchor. Along an authentication path each digest is an anchor or a forward iterate of one. A fixed-slice right step requires current digest b_*, excluded by the guard. A fixed-slice left step applies g.

Reconstruct downward from the cap entry selected by the high index bits. An exponent $j > 0$ forces one left step to $g^{j-1}(a)$. At exponent zero the unique node row supplies both children. An alternative fixed-slice predecessor on a supported path would give a second anchor/exponent representation within the height bound. The remaining index bits therefore force each step. At leaf height, the typed rooted path gives its symbols.

In a shared leaf/node row, $x_R = 0$. Its two node children are 0 and a reachable capacity translated by v_{node}. Both lie in Z_D, so neither supports a competing positive-height path. This permits the literal role equality. A missing lookup makes the position undefined. There are at most $2(H + 1)|D|$ closure keys. Canonical zero completion gives the raw source and middle words. □

Lemma H.2 (Two-sided guard sparsity)**.** *From a guarded database of degree at most t, at most $s\Sigma_H(t)$ legal extensions violate the guard at any fixed forward input or inverse output.*

Proof. For a forward input $x = (w, r)$, first charge at most $(2t + 4)s$ outputs that attach an old capacity tail or end at an old reachable capacity, a root or the address capacity. Otherwise at most one fresh rooted endpoint and one new leaf anchor appear. The new node and possible leaf anchors are $y_R \oplus w$ and y_R. Their self returns cost $2R_H s$. Intersections with old anchor orbits cost at most $4t(H + 1)^2\Lambda_H s$. For $w \neq 0$ their mutual orbit intersection costs $D_H s$. For $w = 0$ their roles are identified.

Hits on old forbidden targets cost at most $2(t + 5)(H + 1)\Lambda_H s$. A new reachable capacity translated into an old closure costs $2t(H+1)s$, and its intersection with either new closure costs $2(H + 1)s$. These terms are bounded by $s\Sigma_H(t)$.

For a fixed inverse output, charge all $(t+3)s$ rooted input choices. Otherwise reachability stays fixed and only the node anchor $y_R \oplus x_R$ is added. Its rate is bijective in x_R. The same return, old-orbit and forbidden-target counts give at most $s[R_H + 2t(H+1)^2\Lambda_H + (t+5)(H+1)\Lambda_H]$ more choices. The stated bound dominates both directions. Restriction to free unused inputs and outputs decreases these counts. □

For an immutable completion, the compressed search estimate [CM26, Lemma 5.2] gives

$$\Pr[B_{\mathrm{stat}}] \le \min\left\{1, \frac{8q^2 s\Sigma_H(q)}{M-q}\right\}. \tag{57}$$

Let q include verification, recheck its ℓ free pairs and put $T = q + \ell$. Define

$$D_M(T) = \frac{120T}{\sqrt{M-T+1}} + \frac{T(T+1)}{M-2T+1}.$$

Theorem H.3 (Extraction from native openings)**.** *For $3T-1 \le M$, extraction with consistent payload recovery has probability at least real verification acceptance minus*

$$\min\left\{1, D_M(T) + \frac{8T^2 s\Sigma_H(T)}{M-T} + \frac{2\ell^2}{M-q}\right\}. \tag{58}$$

Proof. Real pairs pass their ordered rechecks. The compressed approximation costs $D_M(T)$ by Corollary 4.3 of [CM26]. For truth projectors M_i and explicit-record projectors Q_i, [CM26, Equation (83)] gives

$$\|(M_\ell \cdots M_1 - Q_\ell \cdots Q_1)L_{\le q}\| \le \ell\sqrt{2/(M-q)}.$$

This ordered estimate accommodates noncommuting truth checks. Outside the finite guard failure, explicit records give consistent reconstruction. The missing-record and finite-bad branches are orthogonal diagonal events, so their probabilities add. □

H.2 Profiles in the uniform-slice idealization

Theorem H.4 (Uniform-slice profiles)**.** *A uniform injection from the s-point address slice to $N = s^2$ outputs satisfies*

$$\Lambda_{10} \le 2^{15}, \qquad R_{10} < 2^{25}, \qquad D_{10} < 2^{24}$$

except with probability $\delta_F < 2^{-214}$. Consequently $\Sigma_{10}(t) \le 2^{26}(t+5)$ on this event.

Proof. First sample independent full outputs for every slice input. Then g is an independent uniform random map on s vertices. Full-output injectivity has probability greater than $1/2$. Conditioning on it gives exactly the uniform injective slice and increases any failure probability by less than two.

Explore a backward ball through depth ten, stopping at $B = 2^{15}$ vertices. Before stopping its offspring are dominated by Binomial$(s, 1/(s-B))$, hence by Poisson$(101/100)$. For its total population, the moment generating recurrence at $\theta = 1/100$ is

$$m_0 = e^\theta, \qquad m_{j+1} = \exp(\theta + (101/100)(m_j - 1)), \qquad m_{10} < 6/5.$$

A union over roots bounds ball failure by $(6/5)se^{-B/100} < 2^{-216}$, giving $\Lambda_{10} \le B$. The number Z of cycles of length at most ten has $\mathbb{E}(Z)_k \le 3^k$. Thus $\Pr[Z \ge 96] \le (3e/96)^{96} < 2^{-341}$ and $R_{10} \le 10 \cdot 96B < 2^{25}$.

Join distinct vertices whose first eleven iterates intersect. The graph has degree at most $d = 121B$. Conditional on its unlabelled shape, its labels are uniform. For a fixed nonzero XOR shift, its pairs pull back to a uniform perfect matching. If K graph edges occur in that matching,

$$\mathbb{E}\binom{K}{k} \le \frac{(sd/2)^k}{k!(s-2k)^k}, \qquad \Pr[K \ge 2d] < (7/10)^{2d}.$$

The shift correlation is $2K$. Union over shifts costs at most $s(7/10)^{2d} < 2^{-1000000}$, and $4d < 2^{24}$. Doubling the three failure bounds is below 2^{-214}. Substitution gives $\Sigma_{10}(t) \le 2^{26}(t+5)$. The numerical inequalities have exact rational certificates in the artifact. For a fixed concrete slice, the authorization theorem uses its stated finite profiles. □

H.3 Private queries and complete response groups

Let O be a forward or inverse compressed swap query with blank symbol $\perp$. On its deleted-row fiber it is CSC, where C swaps absence with the uniform extension vector and S performs lookup. Fixed-slice calls are public gates. All degree bounds below cover the complete reordered circuits.

Lemma H.5 (Private-query commutator)**.** *For $T < M - 1$, two queries on separate accessible registers, coherently controlled in either direction, obey*

$$\|[O^a, O^b]\|_{\le T} \le c_T := \frac{8}{M-T-1} + 8\sqrt{\frac{2}{M-T}}. \tag{59}$$

Proof. Write each compression reflection as $I - 2P$. For two projection frames with overlap singular value σ, their reflection commutator has norm $4\sigma\sqrt{1-\sigma^2}$. For two distinct forward rows with m available values the frame overlap is

$$\frac{1}{2}\begin{pmatrix} 1 & -\mathbf{1}^{\mathsf{T}}/\sqrt{m} \\ -\mathbf{1}/\sqrt{m} & (J-I)/(m-1) \end{pmatrix}.$$

Its eigenvalues are $1, 0, -1/(2(m-1))$. For one forward row and one inverse image the corresponding overlap is

$$\begin{pmatrix} (m+1)/(2m) & -\mathbf{1}^{\mathsf{T}}/(2\sqrt{m}) \\ -\mathbf{1}/(2\sqrt{m}) & J/(2(m-1)) \end{pmatrix},$$

with eigenvalues $1, 1/(2m), 0$. The compression commutators therefore cost at most $2/(m-1)$ and $2/m$. A forward lookup and inverse compression cross only at their common pair. Its extension diagonal is $1 - 1/m$, giving cross norm at most $2\sqrt{2/m}$. Reordering the two CSC queries uses four compression crossings and four such lookup crossings, proving the bound. Same-direction lookups commute with the other row's compression. Equal query points commute exactly. Direct sums give coherent controls. □

Lemma H.6 (A coherent cold response group)**.** *Fix a measured native prefix observable e before an independent uniform free label z is drawn. Its response predicate has density at most κ on ℓ independent full states. For an arbitrary post-measurement database state of degree at most T, a native ℓ-call response group starting at the fresh relabelled input has successful amplitude at most*

$$\sqrt{\Gamma_\ell \kappa} + \sqrt{\ell \eta_\ell} + \ell\sqrt{2/(M-T)}, \quad \begin{cases} \Gamma_\ell = (N/(M-T-\ell+1))^\ell, \\ \eta_\ell = T/M + \dfrac{(\ell-1)(T+s+\ell+1)}{M-T-\ell+1}. \end{cases} \tag{60}$$

Proof. For an absent input w, the ideal extension isometry is

$$J_w|h\rangle = \frac{1}{\sqrt{M-|h|}} \sum_{y \in Y \setminus \operatorname{im} h} |y, h[w \mapsto y]\rangle, \qquad Y = \Omega \setminus F(A).$$

The exact compressed query differs by operator norm $\sqrt{2/(M-|h|)}$: on a deleted-row fiber, compression subtracts $m^{-1}\sum_y |h_y\rangle$ and adds $m^{-1/2}|h\rangle$ to each extension term, giving squared difference $2/m$. Distinct deleted-row bases occupy orthogonal fibers.

Apply ℓ ideal extensions, keeping full private outputs. Stop on a fixed-slice input or an old row. The first failure has mass at most T/M. Later failures are bounded by the second term of η_ℓ, also allowing the two exchanged labels. Every continuing output tuple has mass at most $(M-T-\ell+1)^{-\ell}$, and at most $N^\ell\kappa$ tuples satisfy the predicate. For a fixed tuple and fresh label, deleting its new input rows recovers the initial database uniquely. Thus distinct initial bases have orthogonal final states, and this density bound holds on arbitrary coherent superpositions. Telescoping the cold projections costs at most $\sum_j \sqrt{\eta_j} \le \sqrt{\ell\eta_\ell}$. The ordinary-query comparisons cost $\ell\sqrt{2/(M-T)}$. A fixed past input adapter is a known bijection, so the same deletion and counting argument applies in its native coordinates. □

References

[BSCH+25] Eli Ben-Sasson, Dan Carmon, Ulrich Haböck, Swastik Kopparty, and Shubhangi Saraf. On proximity gaps for Reed–Solomon codes. Technical Report TR25-169, Electronic Colloquium on Computational Complexity, 2025. https://eccc.weizmann.ac.il/report/2025/169/.

[CDG+17] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. Cryptology ePrint Archive, Paper 2017/279, 2017. https://eprint.iacr.org/2017/279.

[CL06] Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. Cryptology ePrint Archive, Paper 2006/184, 2006. https://eprint.iacr.org/2006/184.

[CM26] Joseph Carolan and Christian Majenz. Compressed permutation oracles revisited. arXiv:2609.28469v1, 2026. https://arxiv.org/abs/2609.28469.

[CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. Cryptology ePrint Archive, Paper 2019/834, 2019. https://eprint.iacr.org/2019/834.

[DP24] Benjamin E. Diamond and Jim Posen. Polylogarithmic proofs for multilinears over binary towers. Cryptology ePrint Archive, Paper 2024/504, 2024. https://eprint.iacr.org/2024/504.

[GKK+25] Lorenzo Grassi, Dmitry Khovratovich, Katharina Koschatko, Christian Rechberger, Markus Schofnegger, Verena Schröppel, and Zhuo Wu. Poseidon(2)b: Binary field versions of Poseidon/Poseidon2. Cryptology ePrint Archive, Paper 2025/1893, 2025. https://eprint.iacr.org/2025/1893.

[GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. In *Proceedings of STOC 2008*, pages 113–122, 2008. https://doi.org/10.1145/1374376.1374391.

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of STOC 1996*, pages 212–219, 1996. https://arxiv.org/abs/quant-ph/9605043.

[GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of Reed–Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. https://doi.org/10.1109/18.782097.

[LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992. https://doi.org/10.1145/146585.146605.

[MPR+26] Dustin Moody, Ray Perlner, Andrew Regenscheid, Angela Robinson, and David Cooper. Transition to post-quantum cryptography standards. Technical Report NIST IR 8547, National Institute of Standards and Technology, 2026. https://csrc.nist.gov/pubs/ir/8547/final.

[Nat24a] National Institute of Standards and Technology. Module-lattice-based digital signature standard. Technical Report FIPS 204, National Institute of Standards and Technology, 2024. https://doi.org/10.6028/NIST.FIPS.204.

[Nat24b] National Institute of Standards and Technology. Stateless hash-based digital signature standard. Technical Report FIPS 205, National Institute of Standards and Technology, 2024. https://doi.org/10.6028/NIST.FIPS.205.

[PWFHW23] Guru-Vamsi Policharla, Bas Westerbaan, Armando Faz-Hernández, and Christopher A. Wood. Post-quantum privacy pass via post-quantum anonymous credentials. Cryptology ePrint Archive, Paper 2023/414, 2023. https://eprint.iacr.org/2023/414.

[Sho95] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. arXiv:quant-ph/9508027, 1995. https://arxiv.org/abs/quant-ph/9508027.

[XZZ+19] Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. Cryptology ePrint Archive, Paper 2019/317, 2019. https://eprint.iacr.org/2019/317.

[ZCF23] Hadas Zeilberger, Binyi Chen, and Ben Fisch. BaseFold: Efficient field-agnostic polynomial commitment schemes from foldable codes. Cryptology ePrint Archive, Paper 2023/1705, 2023. https://eprint.iacr.org/2023/1705.