Vela and Carina: Fast Pairing-Based Multilinear Polynomial Commitments from Reciprocal Polynomials

Yuncong Zhang

School of Cyber Science and Technology, Shandong University
yuncong@sdu.edu.cn

Abstract. Pairing-based multilinear polynomial commitments offer succinct verification for sum-check-based proof systems, but existing schemes trade prover work against proof size and verifier cost. Building on Mercury's representation of a multilinear evaluation as the constant coefficient of a (reciprocal) Laurent polynomial, we construct two schemes with different performance trade-offs. Vela uses inversion symmetry to derive a half-length auxiliary polynomial and opens two polynomials at z and z^{-1} with one univariate KZG proof. Vela has a $2\mathbb{G}_1 + 3\mathbb{F}$ proof, a two-term pairing check, and approximately $2N$ opening MSM scalars, at the cost of $O(N \log N)$ field operations on the prover side. Carina applies the same constant-term reduction once in each coordinate and jointly proves evaluation claims for the committed bivariate polynomial and two auxiliary polynomials with one opening at four bivariate points. The four points form a grid, allowing Carina to restructure the opening equation to save verifier works and proof size. Carina has a $4\mathbb{G}_1 + 8\mathbb{F}$ proof, a three-term pairing check, and $O(N)$ field operations plus exactly $N + 2\sqrt{N} - 6$ MSM works on the prover side.
We prove perfect evaluation completeness, computational commitment binding, and knowledge soundness in the AGM and ROM, with explicit random-oracle error bounds under a falsifiable bilinear q-type DLOG assumption. Our curve-generic Rust implementation also provides common implementations of six pairing-based baselines. On BLS12-381 at $\mu = 20$, Vela gives the smallest measured proof (200 bytes, 1.88x smaller than the next smallest) and the lowest verification median, 2.06 ms. Carina's opening median remains in the same leading tier as the prover-oriented mKZG and CHOPIN implementations, while its 452-byte proof and 2.67 ms verification median improve on both. Compared with the verifier-oriented Mercury, Carina's opening median is 1.96x faster at the cost of a 0.48 ms verification slowdown.

Keywords: Polynomial Commitments, Multilinear Polynomials, KZG, Laurent Polynomials, SNARKs

1 Introduction

Polynomial commitment schemes (PCSs) [20] allow a prover to bind to a polynomial and later certify evaluations with communication sublinear in the poly-

nomial size. They are a standard compilation layer for turning polynomial-oracle protocols into succinct arguments [11, 16, 26]. Consequently, the opening prover, proof size, and verifier of a PCS often determine the corresponding costs of the resulting argument. Motivated by the recent explosion of research in SNARKs [10, 11, 16, 24, 26, 28, 29, 32], PCS have attracted significant effort to reduce their opening costs, proof sizes and verifier costs [1–4,6,9,13,17–19,21,23].

This performance dependence on PCS is especially pronounced for protocols based on the sum-check protocol [7, 25]. Such protocols work natively with multilinear extensions of vectors [10, 22, 29] and can process a circuit of N gates in $O(N)$ field operations [10, 29, 30]. In comparison, a univariate arithmetization [11, 16, 26] may instead require basis conversion, FFTs, or extra quotient work before the same claim can be committed; HyperPlonk identifies this conversion as a concrete prover bottleneck [10]. A multilinear PCS (MLPCS) [6, 13, 21, 27] authenticates the sum-check claim $f(r) = y$ in its native representation and can therefore preserve the promised linear prover.

Pairing-based commitments [20, 27] occupy an important point in this design space. They require a structured reference string which necessitates a trusted setup, but offer algebraic commitments, constant-size group elements, and verification by a few pairing equations [20]. These properties prioritize communication and verifier succinctness, two costs that remain visible even when the surrounding sum-check prover is linear. Yet known pairing-based MLPCSs [3, 6, 13, 17, 21, 27] have a persistent three-way tension. Direct multivariate KZG, commonly called mKZG [27], opens with nearly one size-N MSM, but its proof and verifier pairing works grow with the number of variables. Univariate compilers such as Gemini [6] and Zeromorph [21] obtain succinct verification at the cost of folding and quotient commitments. Constant-size schemes due to Mercury [13] and the concurrent Samaritan [17], followed by CHOPIN [3], reduce communication, but still choose between roughly two full-size MSMs, more pairing equations, or stronger bivariate assumptions. Thus neither "constant proof" nor "linear prover" alone identifies the most efficient construction in practice.

This work builds on Mercury's representation of a multilinear evaluation as the constant coefficient of a Laurent product [13]. From this representation, we derive an auxiliary polynomial of length $N - 1$ using the symmetry under $X \mapsto X^{-1}$. Vela opens this polynomial and the committed polynomial at z and z^{-1} with one batched KZG quotient [20]. It minimizes communication and verifier work, while accepting $O(N \log N)$ structured field work. Carina instead applies the constant-term reduction once along each coordinate of a $\sqrt{N} \times \sqrt{N}$ tensor embedding and uses one bivariate KZG opening to verify evaluations of the tensor polynomial and two auxiliary polynomials. Its proof remains constant size, while its prover work is close to the single size-N MSM used by mKZG [27].

The resulting trade-off is concrete. Vela has a $2\mathbb{G}_1+3\mathbb{F}$ proof and one two-term pairing check. Carina has a $4\mathbb{G}_1 + 8\mathbb{F}$ proof, one three-term pairing check, $O(N)$ field work, and exactly $N + 2\sqrt{N} - 6$ opening MSM scalars. Both constructions avoid FFTs and use only a constant number of second-group SRS elements. Our security analysis proves perfect evaluation completeness, computational commit-

ment binding, and knowledge soundness in the algebraic group model (AGM) [15] and the random-oracle model (ROM) used by Fiat–Shamir [14], with explicit error bounds. The reductions use a falsifiable bilinear q-type DLOG assumption from prior AGM analyses [13, 15]; Carina uses a random affine embedding to simulate its two independent trapdoors from one q-DLOG challenge.

1.1 Contributions

Our contributions are as follows.

- We construct Vela from Mercury's constant-term encoding [13]. Inversion symmetry yields a degree-$(N-2)$ auxiliary polynomial and a single batched two-point KZG opening [20]. Vela has proof size $2\mathbb{G}_1 + 3\mathbb{F}$, two pairing terms, one dynamic $\mathbb{G}_2$ scalar multiplication, and three $\mathbb{G}_2$ SRS elements. Its opening prover performs two dense MSMs and $O(N \log N)$ field operations. We prove correctness, perfect evaluation completeness, computational commitment binding, AGM+ROM knowledge soundness with an explicit error bound, and evaluation binding against algebraic adversaries in the ROM.
- We construct Carina, an MLPCS with proof size $4\mathbb{G}_1 + 8\mathbb{F}$, three pairing terms, two dynamic $\mathbb{G}_2$ scalar multiplications, and five $\mathbb{G}_2$ SRS elements. Its opening prover uses $O(N)$ field operations and $N + 2\sqrt{N} - 6$ MSM scalars. We prove correctness, perfect evaluation completeness, computational commitment binding, AGM+ROM knowledge soundness with an explicit error bound, and evaluation binding against algebraic adversaries in the ROM under the falsifiable bilinear q-type DLOG assumption [13, 15].
- We give curve-generic Rust implementations of both schemes and common implementations of mKZG [27], Gemini [6], Zeromorph [21], Samaritan [17], Mercury [13], and CHOPIN [3], including schemes for which no reusable implementation was available, and evaluate every scheme through the same PCS interface. BLS12-381 experiments at $\mu = 20$ show that Vela's 200-byte proof is 1.88x smaller than the next smallest and that its 2.06 ms verification median is the lowest measured in our comparison. Carina joins mKZG [27] and CHOPIN [3] in the leading prover-performance tier, while improving their proof size and verification time. Against the verifier-oriented Mercury [13], Carina's opening median is 1.96x faster and its proof is smaller, while verification is 0.48 ms slower.

1.2 Related Works

The history of pairing-based PCSs begins with the univariate KZG commitment of Kate, Zaverucha, and Goldberg [20]. The direct multivariate extension of Papamanthou, et al. [27], named mKZG, commits under a tensor-product key and produces one quotient-witness commitment per variable, whose corresponding MSM lengths form the geometric sequence $N/2, N/4, \ldots, 1$, whose sum is $N-1$. This approach therefore minimizes aggregate prover group work, but its proof size and verifier work remain linear in $\mu = \log N$.

A second approach compiles multilinear evaluation into univariate openings. Gemini [6] recursively folds the multilinear polynomial and batches the resulting claims, while Zeromorph [21] combines multilinear quotients into one univariate relation. These schemes reduce verification cost but introduce fold commitments, quotient commitments, or larger SRS slices.

The most closely related work targets constant-size proofs with linear-time field work. Mercury [13] represented multilinear evaluation as the constant term of a Laurent polynomial and achieved constant proof size without prover FFTs, at the cost of $2N+O(\sqrt{N})$ scalar multiplications. The concurrent Samaritan [17] transforms a univariate PCS into a multilinear PCS and obtains constant proof size using bivariate KZG. CHOPIN [3] also reduces the large quotient commitment with bivariate KZG. Its prover approaches one size-N MSM but uses more proof elements and pairing work.

Chmel et al. [12] introduce the concept and AGM-clarity and show that, for AGM-clear PCSs, AGM knowledge soundness is closely tied to evaluation binding through weak interpolation knowledge soundness (WIKS). The linear commitments of Vela and Carina are typical AGM-clear instances. This observation calibrates the model-internal interpretation of our security theorems.

Our two schemes optimize different costs. Vela reduces proof size and verifier work; Carina reduces prover group work to one size-N MSM plus lower-order terms. Their knowledge-soundness analyses use the falsifiable bilinear q-type DLOG assumption from prior AGM work; Carina additionally incurs a statistical loss from embedding its two trapdoors into one random affine q-DLOG instance.

Transparent alternatives use several different techniques. The encode-and-hash approach begins with the Reed–Solomon proximity protocol FRI [4]. Its descendant STIR [1] improves query complexity through rate-improving folds. WHIR [2] is another FRI descendant and extends testing to constrained RS codes. FRI requires structured domains: its analysis treats binary additive cosets and smooth multiplicative subgroups. STIR and WHIR likewise require multiplicative cosets of power-of-two order. Thus these transparent schemes are not field agnostic. BaseFold [31] instead generalizes the FRI folding paradigm to foldable linear codes and gives a field-agnostic instantiation. Brakedown [18] follows the tensor-code line and obtains a linear-time prover over arbitrary sufficiently large fields, at the cost of substantially larger proofs.

Transparent group-based systems make different choices. Bulletproofs [8] uses discrete-logarithm inner-product arguments; Dory [23] uses bilinear groups and pairings but has a transparent setup; and DARK [9] uses groups of unknown order. These systems have proof-size, verifier, and recursion profiles different from the structured pairing-based setting studied here, in which a few group elements and pairing terms are the primary optimization targets.

2 Technical Overview

This section gives a self-contained account of the two constructions. Both start from Mercury's representation of a multilinear evaluation as a constant coeffi-

cient [13], but they authenticate the resulting identity differently. Vela keeps the representation univariate and minimizes the proof and verifier. Carina moves to a bivariate commitment, removes an intermediate polynomial commitment, and minimizes the prover's large group operations.

Let $\mathbf{a} = (a_0, \ldots, a_{N-1})$ be the hypercube evaluation vector of a multilinear polynomial f, where $N = 2^\mu$. For an opening point $r \in \mathbb{F}^\mu$, let $\mathbf{e}_r = (\mathrm{eq}(0; r), \ldots, \mathrm{eq}(N-1; r))$ be the multilinear evaluation weight vector. The opening claim is the inner-product statement

$$f(r) = \langle \mathbf{a}, \mathbf{e}_r \rangle = \sum_{i=0}^{N-1} a_i \mathrm{eq}(i; r) = y.$$

The goal is to authenticate this statement while controlling three different costs: full-size prover MSMs, proof group elements, and verifier pairing terms.

2.1 From an Evaluation to a Symmetric Laurent Identity

Encode the entries of $\mathbf{a}$ and $\mathbf{e}_r$ as coefficients of two univariate polynomials:

$$f_v(X) = \sum_{i=0}^{N-1} a_i X^i, \qquad T_r(X) = \sum_{i=0}^{N-1} \mathrm{eq}(i; r) X^i = \prod_{k=0}^{\mu-1} \big((1 - r_k) + r_k X^{2^k}\big).$$

In the Laurent product $H(X) = f_v(X) T_r(X^{-1})$, a term contributes to the constant coefficient exactly when the two coefficient indices agree. Consequently, $[X^0]H = f(r)$, where $[X^0]$ stands for the constant term of the polynomial. The tensor product formula for T_r also lets the verifier evaluate $T_r(z)$ at any point using only $O(\mu)$ field operations.

Mercury proves a constant-coefficient statement by exploiting the positive and negative parts of this Laurent product [13]. We use the inversion-symmetric residual

$$D(X) = H(X) + H(X^{-1}) - 2y.$$

Its constant coefficient is $2(f(r) - y)$, which is zero exactly when the claim is correct because $\mathbb{F}$ has odd characteristic. Moreover, $D(X) = D(X^{-1})$. If the claim is correct, write

$$D(X) = \sum_{k=1}^{N-1} d_k (X^k + X^{-k})$$

and define the ordinary polynomial

$$\bar{h}(X) = \sum_{k=1}^{N-1} d_k X^{k-1}.$$

Then $\deg \bar{h} \le N - 2$ and

$$D(X) = X\bar{h}(X) + X^{-1}\bar{h}(X^{-1}).$$

Conversely, any Laurent polynomial of this form has zero constant coefficient. Thus the original multilinear evaluation is equivalent to the existence of $\bar{h}$ satisfying

$$X\bar{h}(X) + X^{-1}\bar{h}(X^{-1}) + 2y = f_v(X)T_r(X^{-1}) + f_v(X^{-1})T_r(X).$$

The prover does not need a generic polynomial multiplication to construct $\bar{h}$. Each factor of T_r has two nonzero coefficients, so multiplying by the μ tensor factors can be implemented as structured shift-and-add passes. The total cost is $O(N \log N)$ field operations and linear memory, with no FFT. Once a commitment to $\bar{h}$ is fixed, evaluating the identity at a random $z \notin \{0, 1, -1\}$ gives the scalar check

$$z\bar{h}(z) + z^{-1}\bar{h}(z^{-1}) + 2y = f_v(z)T_r(z^{-1}) + f_v(z^{-1})T_r(z).$$

If $y \neq f(r)$, the cleared Laurent residual is a nonzero polynomial of degree at most $2N$, so a random z detects the false claim except with the usual Schwartz–Zippel error. What remains is to authenticate the polynomial evaluations appearing in this scalar identity.

2.2 Vela: One Batched Two-Point Opening

Vela applies the preceding identity directly to a univariate KZG commitment [20]. Setup samples τ and the commitment is $C_f = [f_v(\tau)]_1$. To open at r, the prover first constructs $\bar{h}$ and sends $C_h = [\bar{h}(\tau)]_1$. This commitment is absorbed before the challenge z is sampled, preventing the prover from choosing the auxiliary polynomial after seeing the test point. The prover then sends

$$v_0 = f_v(z), \quad v_1 = f_v(z^{-1}), \quad w_0 = \bar{h}(z),$$

and the verifier recovers the fourth value $w_1 = \bar{h}(z^{-1})$ from the scalar identity above,

$$w_1 = z\big(v_0 T_r(z^{-1}) + v_1 T_r(z) - z w_0 - 2y\big).$$

The recovery is always defined because the protocol guarantees $z \neq 0$. The same compression cannot omit v_0 or v_1 instead: solving the identity for them requires dividing by $T_r(z)$ or $T_r(z^{-1})$, either of which may vanish.

The four evaluations are authenticated by one degree-two KZG opening. Let R_f be the degree-one polynomial interpolating (z, v_0) and (z^{-1}, v_1), and define R_h analogously from w_0, w_1. After all four values have been absorbed, the transcript samples a batching challenge α. With

$$Z(X) = (X - z)(X - z^{-1}) = X^2 - (z + z^{-1})X + 1,$$

the prover forms

$$q(X) = \frac{f_v(X) + \alpha\bar{h}(X) - R_f(X) - \alpha R_h(X)}{Z(X)}$$

and sends $\pi = [q(\tau)]_1$. If all four evaluations are correct, the numerator vanishes at both roots of Z and q is a polynomial. Conversely, the random α prevents incorrect evaluations of f_v and $\bar{h}$ from cancelling in their linear combination, except with inverse-field probability.

The verifier reconstructs $R = R_f + \alpha R_h$ and its low-degree commitment C_R, and checks

$$e(C_f + \alpha C_h - C_R, [1]_2) = e\big(\pi, [Z(\tau)]_2\big).$$

Only three fixed second-group SRS elements are needed because

$$[Z(\tau)]_2 = [\tau^2]_2 - (z + z^{-1})[\tau]_2 + [1]_2.$$

No independent scalar-identity check remains: the recovery formula makes the claimed tuple satisfy the Laurent identity by construction, while the pairing equation certifies that the three transmitted evaluations and the recovered fourth one are genuine evaluations of f_v and $\bar{h}$, linking them to the original claim $f(r) = y$. The transcript order is essential: C_h precedes z, and α is sampled only after w_1 has been recovered and bound into the transcript. In the formal analysis, AGM representations reveal the committed polynomials, and $(N+1)$-DLOG lifts the accepted KZG equation at the hidden point τ to the required polynomial identity. This gives AGM+ROM knowledge soundness. The remaining errors are the random-linear-combination error for α and the Schwartz–Zippel error for z.

The proof is $(C_h, \pi, v_0, v_1, w_0)$, namely $2\mathbb{G}_1 + 3\mathbb{F}$. Excluding the initial commitment of f itself, the two prover MSMs contain $(N-1) + (N-2) = 2N-3$ scalars. The verifier performs one two-term pairing product, one dynamic scalar multiplication in $\mathbb{G}_2$, a constant-size $\mathbb{G}_1$ linear combination, and $O(\mu)$ field operations. The SRS contains N elements in $\mathbb{G}_1$ and three elements in $\mathbb{G}_2$.

The compact proof deliberately costs more field arithmetic. Constructing $\bar{h}$ takes $O(N \log N)$ field operations rather than the $O(N)$ field work of prover-oriented schemes. The extra logarithmic factor consists of regular shift-and-add passes; it does not add an MSM or a proof element. In our repeated measurements at $\mu = 20$, Vela's opening median is within 0.7% of Zeromorph's and is 1.21x and 1.23x faster than Gemini and Samaritan, respectively, while producing the smallest proof and the lowest verification median in the comparison. Thus the asymptotic field-operation disadvantage produces a moderate measured overhead rather than making Vela impractical.

2.3 Carina: Opening an Implicitly Fixed Restriction

Carina follows Mercury's approach of reducing the $O(N \log N)$ prover work by folding the polynomial to size $O(\sqrt{N})$ before applying the Laurent equation. The key innovation of Carina originates from an avoidable cost in Mercury's route [13]. Assume first that μ is even and set $b = 2^{\mu/2} = \sqrt{N}$. Reshape the evaluation vector into a $b \times b$ matrix F and commit under a tensor KZG key to

$$F(X, Y) = \sum_{i,j<b} F_{ij} X^i Y^j, \qquad C_f = [F(\tau, \sigma)]_1.$$

Split the opening point as (u, v), with $u, v \in \mathbb{F}^{\mu/2}$. The univariate encoding of the restricted polynomial $f(X, v)$ gives

$$g(X) = \sum_{i<b} \left(\sum_{j<b} F_{ij} \mathrm{eq}(j; v) \right) X^i,$$

Once C_f is binding and v is fixed, the polynomial g is uniquely determined. Mercury commits to this intermediate polynomial and proves its consistency with the original commitment. This intermediate commitment provides a convenient verification handle, but it carries no additional freedom in the committed statement since g is completely determined by f and v. Carina asks whether its evaluations can instead be authenticated directly from C_f.

Carina follows the same route of Mercury without committing to g. Define

$$T_u(X) = \prod_{k<\mu/2} \big((1 - u_k) + u_k X^{2^k}\big).$$

The first application of Laurent polynomial equation constructs a degree-at-most-$(b-2)$ polynomial S_0 satisfying

$$g(X)T_u(X^{-1}) + g(X^{-1})T_u(X) = 2y + XS_0(X) + X^{-1}S_0(X^{-1}).$$

The prover sends $C_0 = [S_0(\tau)]_1$ before the transcript samples $\rho \notin \{0, 1, -1\}$. It then sends

$$a_0 = g(\rho), \qquad a_1 = g(\rho^{-1}), \qquad b_0 = S_0(\rho).$$

At $X = \rho$, the displayed identity determines the omitted value $S_0(\rho^{-1})$ from a_0, a_1, b_0 and public data. This first random test reduces the claim $f(u, v) = y$ to authenticating a_0 and a_1 as against the committed bivariate polynomial. In Mecury, this is achieved via directly opening g. Without committing to g, Carina instead exploits the definition of g as follows.

After a_0, a_1 are fixed, a nonzero challenge λ batches those two claims. Define

$$\phi(Y) = F(\rho, Y) + \lambda F(\rho^{-1}, Y), \qquad y_1 = a_0 + \lambda a_1.$$

If both claims are correct, then the coefficients of ϕ encodes the evaluations of a multilinear polynomial whose evaluation at v is y_1. To authenticate this, applying the same identity in the second coordinate, with

$$T_v(Y) = \prod_{k<\mu/2} \big((1 - v_k) + v_k Y^{2^k}\big),$$

gives an auxiliary polynomial S_1 such that

$$\phi(Y)T_v(Y^{-1}) + \phi(Y^{-1})T_v(Y) = 2y_1 + YS_1(Y) + Y^{-1}S_1(Y^{-1}).$$

The prover sends $C_1 = [S_1(\sigma)]_1$ before a second point $\zeta \notin \{0, 1, -1\}$ is sampled. It sends the four values

$$c_{ij} = F\big(\rho^{(-1)^i}, \zeta^{(-1)^j}\big), \qquad i, j \in \{0, 1\},$$

together with $b_1 = S_1(\zeta)$. The c_{ij} determine $\phi(\zeta)$ and $\phi(\zeta^{-1})$, and the second scalar identity recovers $S_1(\zeta^{-1})$. At this point the original evaluation has been reduced to claims about F, S_0, and S_1 on the common Cartesian grid

$$A \times B = \{\rho, \rho^{-1}\} \times \{\zeta, \zeta^{-1}\}.$$

The final step authenticates this grid without introducing a commitment to g. Let $I(X, Y)$ interpolate the four claimed values of F on $A \times B$, let $L_0(X)$ interpolate S_0 on A, and let $L_1(Y)$ interpolate S_1 on B. Regard $S_0 - L_0$ as a bivariate polynomial that depends only on X, and $S_1 - L_1$ as one that depends only on Y. After all claimed values are fixed, a challenge η combines the three residuals:

$$E_\eta(X, Y) = (F - I) + \eta(S_0 - L_0) + \eta^2(S_1 - L_1).$$

If the claims are correct, E_η vanishes on $A \times B$. For

$$Z_A(X) = (X - \rho)(X - \rho^{-1}), \qquad Z_B(Y) = (Y - \zeta)(Y - \zeta^{-1}),$$

canonical division first in X and then in Y gives

$$E_\eta(X, Y) = Z_A(X)U_\eta(X, Y) + Z_B(Y)V_\eta(X, Y).$$

The prover sends $\mathit{\Pi}_X = [U_\eta(\tau, \sigma)]_1$ and $\mathit{\Pi}_Y = [V_\eta(\tau, \sigma)]_1$. From C_f, C_0, C_1 and the low-degree interpolants, the verifier reconstructs the commitment C_E to E_η and as

$$C_E = C_f + \eta C_0 + \eta^2 C_1 - [I(\tau, \sigma) + \eta L_0(\tau) + \eta^2 L_1(\sigma)]_1,$$

and checks the single product equation

$$e(C_E, [1]_2) = e\big(\mathit{\Pi}_X, [Z_A(\tau)]_2\big)e\big(\mathit{\Pi}_Y, [Z_B(\sigma)]_2\big).$$

The two second-group elements on the right are constructed from the five fixed elements $[1]_2, [\tau]_2, [\tau^2]_2, [\sigma]_2, [\sigma^2]_2$.

This sequence explains how the omitted commitment is replaced. The first random test reduces $f(u, v) = y$ to two evaluations of g; λ binds those two values to the restriction of F tested in the second coordinate; the second random test reduces that relation to four evaluations of F; and the grid opening authenticates the four evaluations against the original commitment C_f. In the security proof, a falsifiable bilinear q-type DLOG assumption used in prior AGM analyses [13,15] lifts the accepted pairing equation at the hidden point (τ, σ) to formal membership in the grid vanishing ideal. The reduction embeds both trapdoors as independently blinded affine functions of one q-DLOG challenge. The simulated pair remains independent and uniform, while any nonzero bivariate relation specializes to a nonzero univariate relation except with a Schwartz–Zippel loss. As for Vela, the Fiat–Shamir transcript is analyzed in the ROM and algebraic commitments are analyzed in the AGM. The resulting analysis establishes AGM+ROM knowledge soundness.

The cost follows from the canonical quotient supports. The commitment to each S_i uses $b - 1$ scalars. The X-quotient has support $(b - 2) \times b$, and the

remaining Y-quotient has support $2 \times (b-2)$, for $N-4$ quotient scalars in total. Hence opening uses exactly

$$2(b-1) + (N-4) = N + 2\sqrt{N} - 6$$

MSM scalars and $O(N)$ field operations. The proof consists of C_0, C_1, Π_X, Π_Y and the eight transmitted field values, for $4\mathbb{G}_1 + 8\mathbb{F}$. The verifier performs one three-term pairing product, two dynamic scalar multiplications in $\mathbb{G}_2$, a constant-size $\mathbb{G}_1$ linear combination, and $O(\mu)$ field operations. The SRS contains N tensor elements in $\mathbb{G}_1$ and five elements in $\mathbb{G}_2$. Relative to Mercury, this reduces the proof from $8\mathbb{G}_1 + 6\mathbb{F}$ to $4\mathbb{G}_1 + 8\mathbb{F}$ and the opening group work from $2N + O(\sqrt{N})$ to one size-N MSM plus lower-order terms.

The two constructions therefore optimize opposite sides of the same constant-term test. Vela retains a univariate commitment and spends structured field work to minimize communication and verification. Carina uses a tensor commitment and two coordinate reductions to eliminate the intermediate restriction commitment and approach the minimum prover group work. Sections 5 and 6 gives the complete protocols and security arguments, while Section 7 compares their exact costs with prior MLPCSs.

3 Preliminaries

3.1 Notation and Conventions

The security parameter is $\lambda \in \mathbb{N}$. An algorithm is efficient if it is probabilistic polynomial time (PPT) in λ, and $\mathsf{negl}(\lambda)$ denotes a negligible function. For a finite set S, the notation $x \xleftarrow{\$} S$ means that x is sampled uniformly from S. Vectors are written in bold; their coordinates are indexed from zero. For a polynomial p, $[X^i]p$ denotes the coefficient of X^i.

Let $\mathbb{F}$ be a prime field of odd characteristic approximately 2^λ. Let $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ be cyclic groups of prime order $|\mathbb{F}|$ with generators G_1, G_2 and a nondegenerate bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$. We write $[a]_1 = aG_1$ and $[a]_2 = aG_2$. All algorithms and assumptions are parameterized by this group family.

Let $\mu \geq 1$ be the number of variables and $N = 2^\mu$ the evaluation-table size. The theoretical presentation of Carina assumes even μ and sets

$$b = 2^{\mu/2} = \sqrt{N}.$$

The implementation supports odd μ through the rectangular split described in Section 7.2.

For $i = \sum_{k=0}^{\mu-1} i_k 2^k$ with $i_k \in \{0, 1\}$, define

$$\mathrm{eq}(i; r) = \prod_{k=0}^{\mu-1} \big(i_k r_k + (1 - i_k)(1 - r_k)\big).$$

The multilinear extension of a vector $\mathbf{a} = (a_0, \ldots, a_{N-1})$ is

$$f(X_0, \ldots, X_{\mu-1}) = \sum_{i=0}^{N-1} a_i \mathrm{eq}(i; X), \qquad f(r) = \sum_{i=0}^{N-1} a_i \mathrm{eq}(i; r).$$

3.2 Interactive Proofs and Arguments

Let $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$ be an efficiently decidable relation and let $L_{\mathcal{R}} = \{x : \exists w\ (x,w) \in \mathcal{R}\}$. For a prover $\mathcal{P}$ and verifier $\mathcal{V}$, write $b \leftarrow \langle \mathcal{P}(x,w), \mathcal{V}(x) \rangle$ for their interaction and the verifier's output $b \in \{0,1\}$.

Definition 1 (Interactive proof and argument). *An interactive proof for $L_{\mathcal{R}}$ is a pair $(\mathcal{P}, \mathcal{V})$ in which $\mathcal{P}$ may be computationally unbounded and $\mathcal{V}$ is PPT. The prover takes input x, w and the verifier inputs x. The parties interact with each other. After receiving the last message from the prover, the verifier outputs a bit $b \in \{0,1\}$. This process is denoted by $b \leftarrow \langle \mathcal{P}(x,w), \mathcal{V}(x) \rangle$. The proof satisfies the following properties for completeness error ε_c and soundness error ε_s:*

- Completeness. *For every $(x,w) \in \mathcal{R}$,*

$$\Pr[b = 1 \mid b \leftarrow \langle \mathcal{P}(x,w), \mathcal{V}(x) \rangle] \geq 1 - \varepsilon_c(\lambda).$$

- Soundness. *For every $x \notin L_{\mathcal{R}}$ and every prover $\mathcal{P}^*$,*

$$\Pr[b = 1 \mid b \leftarrow \langle \mathcal{P}^*(x), \mathcal{V}(x) \rangle] \leq \varepsilon_s(\lambda).$$

An interactive argument is defined identically except that soundness is required only against PPT provers $\mathcal{P}^*$.

Definition 2 (Argument of knowledge). *An interactive argument $(\mathcal{P}, \mathcal{V})$ for $\mathcal{R}$ is an argument of knowledge with knowledge error κ if there exists an expected-PPT extractor $\mathcal{E}$ such that, for every PPT prover $\mathcal{P}^*$,*

$$\Pr\left[b = 1 \ \wedge\ (x,w) \notin \mathcal{R} \;\middle|\; \begin{array}{l} b \leftarrow \langle \mathcal{P}^*(x), \mathcal{V}(x) \rangle, \\ w \leftarrow \mathcal{E}^{\mathsf{nextmsg}[\mathcal{P}^*]}(x) \end{array} \right] \leq \kappa(\lambda).$$

The extractor has oracle access to the next-message function of $\mathcal{P}^*$, and the probability is over all protocol and extractor coins.

Our interactive opening arguments are public coin. The Fiat–Shamir transform [14] replaces each verifier message by a domain-separated hash of the statement and the complete preceding transcript. A detailed specification is given in Appendix A.

3.3 Multilinear Polynomial Commitments

Let $\mathcal{P}_\mu$ be the multilinear polynomials in μ variables over $\mathbb{F}$. Both commitments in this paper are deterministic, so the committed polynomial itself is the witness in the evaluation relation.

Definition 3 (Multilinear polynomial commitment scheme). *An MLPCS consists of algorithms* $(\mathsf{Setup}, \mathsf{Commit}, \mathsf{Open}, \mathsf{Verify})$ *with the following syntax:*

– $\mathsf{Setup}(1^\lambda, 1^\mu)$ *returns a commitment key* ck *and a verification key* vk.
– $\mathsf{Commit}(\mathsf{ck}, f) \to C$ *deterministically commits to* $f \in \mathcal{P}_\mu$.
– Open *takes* $(\mathsf{ck}, \mathsf{vk}, f, C, r, y)$ *and returns an evaluation proof* π. *The honest algorithm is invoked with* $C = \mathsf{Commit}(\mathsf{ck}, f)$ *and* $y = f(r)$.
– $\mathsf{Verify}(\mathsf{vk}, C, r, y, \pi) \to \{0, 1\}$ *verifies the claimed evaluation.*

The displayed Open *and* Verify *algorithms are the Fiat–Shamir versions of the public-coin opening arguments specified later.*

The relation proved by an evaluation proof is

$$\mathcal{R}_{\mathsf{eval}} = \{((\mathsf{ck}, \mathsf{vk}, C, r, y), f) : f \in \mathcal{P}_\mu \ \wedge \ C = \mathsf{Commit}(\mathsf{ck}, f) \ \wedge \ f(r) = y\}. \quad (1)$$

Definition 4 (Evaluation completeness and commitment binding). *An MLPCS has perfect evaluation completeness if, for every supported* f *and* $r \in \mathbb{F}^\mu$,

$$\Pr\left[\mathsf{Verify}(\mathsf{vk}, C, r, y, \pi) = 1 \;\middle|\; \begin{array}{l}(\mathsf{ck}, \mathsf{vk}) \leftarrow \mathsf{Setup}(1^\lambda, 1^\mu),\\ C \leftarrow \mathsf{Commit}(\mathsf{ck}, f), \quad y = f(r),\\ \pi \leftarrow \mathsf{Open}(\mathsf{ck}, \mathsf{vk}, f, C, r, y)\end{array}\right] = 1.$$

It is computationally binding if every PPT adversary has negligible probability in

$$\Pr\left[\begin{array}{c}f_0, f_1 \in \mathcal{P}_\mu, \quad f_0 \neq f_1,\\ \mathsf{Commit}(\mathsf{ck}, f_0) = \mathsf{Commit}(\mathsf{ck}, f_1)\end{array} \;\middle|\; \begin{array}{l}(\mathsf{ck}, \mathsf{vk}) \leftarrow \mathsf{Setup}(1^\lambda, 1^\mu),\\ (f_0, f_1) \leftarrow \mathcal{A}(\mathsf{ck}, \mathsf{vk})\end{array}\right].$$

Definition 5 (Evaluation binding). *Fix an adversary model* $\mathcal{M}$ *and oracle access* $\mathcal{O}$. *An MLPCS is evaluation binding in* $(\mathcal{M}, \mathcal{O})$ *with error* $\varepsilon_{\mathsf{eb}}$ *if every PPT adversary* $\mathcal{A}^{\mathcal{O}}$ *in* $\mathcal{M}$ *wins the following experiment with probability at most* $\varepsilon_{\mathsf{eb}}(\lambda)$*:*

$$\begin{aligned}(\mathsf{ck}, \mathsf{vk}) &\leftarrow \mathsf{Setup}(1^\lambda, 1^\mu),\\ (C, r, y_0, \pi_0, y_1, \pi_1) &\leftarrow \mathcal{A}^{\mathcal{O}}(\mathsf{ck}, \mathsf{vk}),\\ b_i &\leftarrow \mathsf{Verify}(\mathsf{vk}, C, r, y_i, \pi_i), \qquad i \in \{0, 1\}.\end{aligned}$$

The adversary wins if $y_0 \neq y_1$ *and* $b_0 = b_1 = 1$. *We omit the error when it is negligible. The theorems in this paper instantiate* $(\mathcal{M}, \mathcal{O}) = (\mathsf{AGM}, \mathsf{ROM})$ *and include the random oracle in the probability space.*

Evaluation binding forbids two accepted values for one commitment and point. The AGM knowledge-soundness notion used by our main theorems is defined after the AGM in Section 3.5.

3.4 Pairing-Based Polynomial Commitments

The univariate KZG scheme of Kate, Zaverucha, and Goldberg [20] commits to $p(X) = \sum_{i=0}^{d} p_i X^i$ as $[p(\tau)]_1 = \sum_i p_i [\tau^i]_1$. For $S = \{x_1, \ldots, x_k\}$, let $Z_S(X) =$

$\prod_{x \in S}(X - x)$ and let R interpolate the claimed values of p on S. The proof $\pi = [q(\tau)]_1$ for $q = (p - R)/Z_S$ is verified by

$$e\big([p(\tau)]_1 - [R(\tau)]_1, [1]_2\big) = e\big(\pi, [Z_S(\tau)]_2\big).$$

Several polynomials opened on the same set are combined with a verifier challenge before applying this pairing check.

The bivariate KZG construction derived from mKZG [27] samples independent trapdoors τ, σ and publishes the balanced tensor key

$$\big\{[\tau^i \sigma^j]_1 : 0 \le i, j < b\big\}.$$

It commits to $P(X, Y) = \sum_{i,j<b} p_{ij} X^i Y^j$ as $[P(\tau, \sigma)]_1$. If P vanishes on $A \times B$, then

$$P(X, Y) = Z_A(X) U(X, Y) + Z_B(Y) V(X, Y),$$

where $Z_A = \prod_{a \in A}(X - a)$ and $Z_B = \prod_{b \in B}(Y - b)$. Commitments to U and V verify through one three-term pairing product. We refer to this technique as the grid opening of bivariate KZG. This is the final step of Carina.

3.5 Security Models and Assumptions

We use the algebraic group model (AGM) of Fuchsbauer et al. [15]. Every group element output by an adversary is accompanied by coefficients relative to the input group elements. These coefficients are an algebraic explanation of the output. In particular, an explanation of a commitment relative to a monomial SRS may reveal coefficient data for a polynomial.

Definition 6 (AGM knowledge soundness in the ROM). *The Fiat–Shamir opening is AGM knowledge sound in the ROM with knowledge error κ if there exists a single PPT extractor $\mathcal{E}$ such that, for every PPT algebraic adversary $\mathcal{A}^H$,*

$$\Pr\left[b = 1 \ \wedge \ ((\mathsf{ck}, \mathsf{vk}, C, r, y), f) \notin \mathcal{R}_{\mathsf{eval}}\right] \le \kappa(\lambda),$$

where

$$\begin{aligned} (\mathsf{ck}, \mathsf{vk}) &\leftarrow \mathsf{Setup}(1^\lambda, 1^\mu), \\ ((C, \gamma_C), r, y, \pi) &\leftarrow \mathcal{A}^H(\mathsf{ck}, \mathsf{vk}), \\ f &\leftarrow \mathcal{E}(\mathsf{ck}, \mathsf{vk}, (C, \gamma_C)), \\ b &\leftarrow \mathsf{Verify}(\mathsf{vk}, C, r, y, \pi). \end{aligned}$$

Here γ_C is the algebraic explanation of the commitment relative to the input SRS. The adversary supplies explanations for all group elements that it outputs, as required by the AGM, but the extractor above receives only the public parameters, the commitment, and its explanation. It does not receive unrestricted prover state. The probability is over setup, the random oracle H, and all algorithmic coins. The quantifier order is $\exists \mathcal{E}\, \forall \mathcal{A}^H$: the extractor is fixed for the scheme, not selected separately for each adversary.

We use the bilinear Q-DLOG assumption introduced in AGM analyses [15] and stated for asymmetric source groups by Eagen and Gabizon [13].

Definition 7 (Bilinear Q-DLOG). *Let $x \xleftarrow{\$} \mathbb{F}^*$. Give an adversary*

$$\left(\left([x^i]_1\right)_{i=0}^{Q}, \left([x^i]_2\right)_{i=0}^{Q} \right).$$

The bilinear Q-DLOG assumption states that every PPT adversary outputs x with negligible probability. We write $\varepsilon_Q^{\mathsf{qdlog}}(\mathcal{B})$ for the success probability of an algorithm $\mathcal{B}$, and $\varepsilon_Q^{\mathsf{qdlog}}$ for the maximum such advantage over all PPT algorithms in the reduction class.

This is a falsifiable bilinear q-type discrete-logarithm assumption for the two source groups, not an SRS-specific relation assumption. The Vela reduction uses $Q = N + 1$ and reveals only the powers required by its SRS. The Carina reduction uses $Q = 2b - 2$: a random affine embedding of one Q-DLOG challenge simulates its two independent tensor trapdoors exactly. The resulting statistical specialization loss is proved in Section 6 and Appendix B.2.

4 Laurent Constant-Term Reduction

4.1 Encoding and Symmetrization

Let f be represented by its hypercube values $\mathbf{a} = (a_0, \ldots, a_{N-1})$, and encode the vector as

$$f_v(X) = \sum_{i=0}^{N-1} a_i X^i.$$

For an opening point $r \in \mathbb{F}^\mu$, define the tensor polynomial

$$T_r(X) = \prod_{k=0}^{\mu-1} \left((1 - r_k) + r_k X^{2^k}\right) = \sum_{i=0}^{N-1} \mathrm{eq}(i; r) X^i.$$

The equality follows by expanding the product and using the binary representation of i. Consequently, for the Laurent polynomial

$$h(X) = f_v(X) T_r(X^{-1}),$$

the constant coefficient is $[X^0]h = \sum_i a_i \mathrm{eq}(i; r) = f(r)$. This is the Laurent-polynomial constant-term encoding first introduced by Bootle et al. [5] and later exploited by Mercury [13].

For a claimed value y, symmetrize h as

$$L(X) = h(X) + h(X^{-1}) - 2y.$$

Its coefficients satisfy $L_i = L_{-i}$ and $L_0 = 2(f(r) - y)$. Define

$$\bar{h}(X) = \sum_{i=0}^{N-2} L_{i+1} X^i.$$

Then

$$L(X) - L_0 = X\bar{h}(X) + X^{-1}\bar{h}(X^{-1}).$$

Since the field has odd characteristic, $L_0 = 0$ exactly when $f(r) = y$. For an honest claim, evaluating this decomposition at a nonzero z gives

$$z\bar{h}(z) + z^{-1}\bar{h}(z^{-1}) + 2y = f_v(z)T_r(z^{-1}) + f_v(z^{-1})T_r(z). \tag{2}$$

Conversely, if the two Laurent polynomials in (2) agree identically, then their constant coefficients imply $f(r) = y$. Vela tests (2) once, while Carina instantiates it once for each tensor coordinate. In both cases, soundness follows by applying Schwartz–Zippel after clearing negative exponents.

4.2 Structured Witness Computation

The prover can compute h without general polynomial multiplication. Each factor of $T_r(X^{-1})$ performs a weighted copy and a right shift by 2^k. Algorithm 1 records this procedure.

Algorithm 1 Structured computation of the auxiliary polynomial

Require: values $a_0, \ldots, a_{N-1}$ and point $r = (r_0, \ldots, r_{\mu-1})$
Ensure: $y = f(r)$ and coefficients of $\bar{h}$
Initialize a Laurent array $H[-(N-1), \ldots, N-1]$ to zero
for $i = 0$ to $N - 1$ **do**
    $H[i] \leftarrow a_i$
for $k = 0$ to $\mu - 1$ **do**
    $s \leftarrow 2^k$ and initialize a zero array H' on the same range
    for each represented exponent d **do**
        $H'[d] \leftarrow H'[d] + (1 - r_k)H[d]$
        $H'[d - s] \leftarrow H'[d - s] + r_k H[d]$
    $H \leftarrow H'$
$y \leftarrow H[0]$
for $i = 0$ to $N - 2$ **do**
    $\hbar_i \leftarrow H[i + 1] + H[-(i + 1)]$
return y and $\bar{h}(X) = \sum_{i=0}^{N-2} \hbar_i X^i$

The algorithm uses $O(N\mu) = O(N \log N)$ field operations and $O(N)$ memory. It requires neither roots of unity nor FFT-friendly field structure. Vela runs Algorithm 1 once at size N. Carina runs the same structured computation on the two $\sqrt{N}$-size univariate restrictions derived from its bivariate representation.

5 Vela: Compact Two-Point Openings

5.1 Protocol Description

Vela commits to the flat encoding f_v defined in Section 4. To prove $f(r) = y$, the prover commits to the auxiliary polynomial $\bar{h}$ and opens both f_v and $\bar{h}$ at z, z^{-1}. Of the four evaluations, three are transmitted and the fourth is recovered from the Laurent identity. A random linear combination reduces the four evaluations to one two-point KZG proof. Protocol 1 gives the complete Fiat–Shamir protocol, where the DeriveChallenge is specified in Appendix A.

Protocol 1 (The Vela multilinear PCS). Let $N = 2^\mu$ and let $a_0, \dots, a_{N-1}$ be the hypercube evaluations of f.

Setup. On input $(1^\lambda, 1^\mu)$, sample $\tau \xleftarrow{\$} \mathbb{F}$. Output a commitment key containing $\{[\tau^i]_1\}_{i=0}^{N-1}$ and a verification key containing $[1]_2, [\tau]_2, [\tau^2]_2$.

Commit. Form $f_v(X) = \sum_i a_i X^i$ and output $C_f = [f_v(\tau)]_1$.

Open. On input $(\mathsf{ck}, \mathsf{vk}, f, C_f, r, y)$ with $C_f = \mathsf{Commit}(\mathsf{ck}, f)$ and $y = f(r)$, proceed as follows.

(a) Compute $\bar{h}$ using Algorithm 1 and send $C_h = [\bar{h}(\tau)]_1$. Define the first transcript prefix $T_0 = (\mathsf{ppid}, \mu, C_f, r, y, C_h)$.

(b) Derive $(z, c_z) = \mathsf{DeriveChallenge}(\texttt{"vela-z"}, T_0, \{0, 1, -1\})$.

(c) Send $v_0 = f_v(z)$, $v_1 = f_v(z^{-1})$, and $w_0 = \bar{h}(z)$. Recover

$$w_1 = z\big(v_0 T_r(z^{-1}) + v_1 T_r(z) - z w_0 - 2y\big), \tag{3}$$

which equals $\bar{h}(z^{-1})$ by the Laurent evaluation identity and is defined because $z \neq 0$. Define the ordered prefix $T_1 = (T_0, c_z, z, v_0, v_1, w_0, w_1)$, which still contains the recovered w_1, and derive $(\alpha, c_\alpha) = \mathsf{DeriveChallenge}(\texttt{"vela-alpha"}, T_1, \varnothing)$.

(d) Let R_f and R_h interpolate the two values of f_v and $\bar{h}$, respectively. For $Z(X) = (X - z)(X - z^{-1})$, compute

$$q(X) = \frac{f_v(X) + \alpha \bar{h}(X) - R_f(X) - \alpha R_h(X)}{Z(X)}.$$

Send $\pi = [q(\tau)]_1$. The proof is $(C_h, v_0, v_1, w_0, \pi)$.

Verify. On input $(\mathsf{vk}, C_f, r, y, \pi)$, reconstruct T_0 from the statement and C_h, recompute (z, c_z), and reject malformed data. Compute $T_r(z), T_r(z^{-1})$ and recover w_1 via (3). Rebuild T_1, including c_z before z and the recovered w_1, and recompute (α, c_α). Reconstruct $R = R_f + \alpha R_h$ from v_0, v_1, w_0 and the recovered w_1, form $C_R = [R(\tau)]_1$ and

$$Z_2 = [\tau^2]_2 - (z + z^{-1})[\tau]_2 + [1]_2,$$

and check

$$e(C_f + \alpha C_h - C_R, [1]_2) = e(\pi, Z_2). \tag{4}$$

Accept if and only if the KZG pairing equation holds: the recovery relation (3) replaces the independent scalar equality check, since it forces the claimed tuple onto the Laurent identity by definition.

The proof contains C_h, π, and three field elements; the fourth evaluation w_1 is recovered by both sides via (3) and bound into the transcript before α. The verifier constructs the quadratic KZG divisor in $\mathbb{G}_2$ from three fixed SRS elements, so the second-group SRS is independent of μ.

5.2 Security Analysis

We analyze the security notions of Definitions 4 to 6, allowing the algebraic adversary to produce C_f with its AGM explanation. From $\gamma_C = (\gamma_0, \ldots, \gamma_{N-1})$, the extractor recovers the polynomial $f_v(X) = \sum_i \gamma_i X^i$ and returns the unique multilinear polynomial whose evaluation vector is $(\gamma_0, \ldots, \gamma_{N-1})$; validity of the explanation gives $C_f = \sum_i \gamma_i [\tau^i]_1 = [f_v(\tau)]_1$. Fiat–Shamir challenges are interpreted as independent domain-separated random-oracle calls.

Theorem 1 (Security of Vela). *The* Vela *MLPCS is perfectly evaluation-complete in the sense of Definition 4. Under $(N+1)$-DLOG, its commitment is computationally binding with advantage at most ε_{N+1}. Its Fiat–Shamir opening is AGM knowledge sound in the ROM in the sense of Definition 6. An algebraic adversary making at most Q candidate random-oracle queries under the two displayed domains, including rejected counters, has knowledge error at most*

$$\delta_{\mathsf{Vela}}(Q) = Q\left(\frac{2N}{|\mathbb{F}| - 3} + \frac{2}{|\mathbb{F}|}\right) + \varepsilon_{N+1} + \mathsf{negl}(\lambda). \tag{5}$$

Consequently, its evaluation-binding advantage against algebraic adversaries in the ROM is at most $2\delta_{\mathsf{Vela}}(Q) + \varepsilon_{N+1}$ in the sense of Definition 5.

Proof (Sketch). For honest inputs, the Laurent identity makes the value w_1 recovered via (3) equal to $\bar{h}(z^{-1})$, so R_h is the correct interpolant and the numerator defining q vanishes at both roots of Z. Thus the quotient is a polynomial and the KZG equation holds, proving perfect completeness.

For commitment binding, two distinct encodings with the same commitment give a known nonzero polynomial Δ of degree below N with $\Delta(\tau) = 0$. Factoring Δ and testing its roots against $[\tau]_1$ recovers τ, contradicting $(N+1)$-DLOG except with advantage ε_{N+1}.

For knowledge soundness, the commitment explanation yields f_v, while the remaining AGM representations yield $\bar{h}$ and q. Equation (4) gives $W(\tau) = 0$ for $W = f_v + \alpha\bar{h} - R - qZ$ and $\deg W \le N+1$. If W is nonzero, the same root-recovery reduction breaks $(N+1)$-DLOG. Hence $W = 0$, except with advantage ε_{N+1}. The proof transmits three evaluations v_0, v_1, w_0; the fourth claimed value is the recovered w_1 of (3), which is uniquely determined by $(z, y, v_0, v_1, w_0, T_r(z), T_r(z^{-1}))$ and therefore fixed and absorbed before α. Since all four claimed evaluations—three transmitted, one recovered—precede α, an incorrect evaluation survives their random linear combination at either point in $\{z, z^{-1}\}$ with probability at most $2/|\mathbb{F}|$.

With genuine evaluations, the recovery relation makes the claimed tuple satisfy the Laurent defect

$$D(X) = X\bar{h}(X) + X^{-1}\bar{h}(X^{-1}) + 2y - f_v(X)T_r(X^{-1}) - f_v(X^{-1})T_r(X)$$

at $X = z$ by definition, so the authenticated genuine evaluations give $D(z) = 0$ for the actual defect. Its support lies in $[-N, N]$ and its constant coefficient is $2(y - f(r))$. For a false claim, $X^N D(X)$ is nonzero of degree at most $2N$. Since C_h precedes z, Schwartz–Zippel contributes $2N/(|\mathbb{F}| - 3)$. A union bound over ROM queries gives (5). The extracted f_v is a witness for (1). If two distinct values both verify, at least one knowledge-soundness failure occurs or the two extracted polynomials violate commitment binding. A union bound proves the stated evaluation-binding bound. Appendix B.1 gives the complete reduction.

5.3 Efficiency

Excluding the initial commitment, the prover computes $\bar{h}$ in $O(N \log N)$ field operations, commits to its $N-1$ coefficients, and commits to a quotient with at most $N-2$ coefficients. Opening therefore uses $2N-3$ MSM scalars, up to lower-order interpolation and evaluation work. Recovering w_1 instead of evaluating $\bar{h}$ at z^{-1} also saves the honest prover one $O(N)$ Horner pass, a low-order optimization that changes neither the $O(N \log N)$ field complexity nor the two opening MSM lengths. The verifier uses one product check with two pairing terms, one dynamic $\mathbb{G}_2$ scalar multiplication, a constant-size $\mathbb{G}_1$ linear combination, and $O(\mu)$ field operations.

The proof is $2\mathbb{G}_1 + 3\mathbb{F}$: 192 bytes of cryptographic payload and 200 serialized bytes in our implementation. The SRS contains N elements in $\mathbb{G}_1$ and three elements in $\mathbb{G}_2$. A detailed comparison appears in Section 7.1.

6 Carina: Grid-Based Openings

6.1 Protocol Description

Assume even μ and let $b = \sqrt{N}$. Carina reshapes the N values of f into a $b \times b$ matrix and commits to

$$F(X, Y) = \sum_{i=0}^{b-1} \sum_{j=0}^{b-1} F_{ij} X^i Y^j$$

using the tensor SRS $\{[\tau^i \sigma^j]_1\}_{i,j<b}$. Split the opening point as (u, v) with $u, v \in \mathbb{F}^{\mu/2}$. Restricting the second coordinate gives

$$g(X) = \sum_{i=0}^{b-1} \left(\sum_{j=0}^{b-1} F_{ij} \mathrm{eq}(j; v) \right) X^i, \tag{6}$$

whose coefficients encodes the evaluations of the partially evaluated polynomial $f(\mathbf{X}, v)$. The prover does not commit to g. Instead, the constant-term identity reduces $f(u, v) = y$ to evaluations of g at ρ, ρ^{-1}, and a second application proves that $g(\rho)$ and $g(\rho^{-1})$ agree with restrictions of the already committed F. One

Cartesian-grid KZG proof then jointly verifies the claimed evaluations of F, S_0, and S_1. Protocol 2 specifies the complete Fiat–Shamir transcript.

The missing commitment to g is deliberate. Sending such a commitment would authenticate only an independently chosen univariate polynomial unless the prover also proved that it equals the definition. Carina instead uses the definition of g itself: after the first test samples ρ, the prover must justify both $g(\rho)$ and $g(\rho^{-1})$ as evaluations of F. A random challenge λ combines these two claims before the next witness is fixed, so a false restriction can survive with probability at most one over the challenge space.

The second application uses $\phi(Y) = F(\rho, Y) + \lambda F(\rho^{-1}, Y)$. Its random point ζ gives the four points in the grid $\{\rho, \rho^{-1}\} \times \{\zeta, \zeta^{-1}\}$. The protocol then has evaluation claims for F, S_0, and S_1 on the same grid. A final challenge η combines their residuals, preventing cancellation between independently false claims.

The last step is algebraic rather than recursive. A polynomial vanishing on the four-point Cartesian grid belongs to the ideal generated by the two quadratic vanishing polynomials. Canonical division in X and then Y produces two quotient commitments whose supports contain $N-4$ coefficients in total. This grid-ideal proof removes the need for a second size-N MSM.

Canonical division first divides in X. The quotient U_η has support $(b-2) \times b$; the remainder has X-degree below two, and its Y-quotient V_η has support $2 \times (b-2)$. This decomposition, first introduced in Mercury, reduces the large group work to one size-N MSM plus lower-order terms.

6.2 Security Analysis

The pairing equation initially gives equality only at the hidden tensor point (τ, σ). We lift this equality to formal membership in the grid vanishing ideal under the falsifiable bilinear q-type DLOG assumption of Definition 7, with $Q = 2b - 2$. The reduction uses the random affine embedding technique of Fuchsbauer, Kiltz, and Loss [15], also used by Mercury's AGM analysis [13]. From an explanation $(\gamma_{ij})_{i,j<b}$ of C_f, the extractor recovers the tensor polynomial $F(X,Y) = \sum_{i,j<b} \gamma_{ij} X^i Y^j$ and the corresponding multilinear polynomial under the fixed reshaping; validity gives $C_f = \sum_{i,j} \gamma_{ij} [\tau^i \sigma^j]_1 = [F(\tau, \sigma)]_1$.

Theorem 2 (Security of Carina). *Let $Q = 2b - 2$. The Carina MLPCS is perfectly evaluation-complete in the sense of Definition 4. In the AGM under bilinear Q-DLOG, its commitment is computationally binding with advantage at most*

$$\varepsilon_Q^{\mathsf{qdlog}} + \frac{2b-2}{|\mathbb{F}|-1}.$$

Its Fiat–Shamir opening is AGM knowledge sound in the ROM in the sense of Definition 6. Its knowledge error is at most

$$\delta_{\mathsf{Carina}} = \varepsilon_Q^{\mathsf{qdlog}} + \frac{2b}{|\mathbb{F}|-1} + \frac{2bQ_\rho + (3b-1)Q_\zeta}{|\mathbb{F}|-3} + \frac{Q_\lambda + 2Q_\eta}{|\mathbb{F}|-1} + \mathsf{negl}(\lambda), \tag{8}$$

Protocol 2 (The Carina multilinear PCS). Let $N = b^2$, and define $A = \{\rho, \rho^{-1}\}$ and $B = \{\zeta, \zeta^{-1}\}$ after the challenges are sampled.

Setup. On input $(1^\lambda, 1^\mu)$, sample independent $\tau, \sigma \xleftarrow{\$} \mathbb{F}$. Output $\{[\tau^i\sigma^j]_1\}_{i,j<b}$ and $[1]_2, [\tau]_2, [\tau^2]_2, [\sigma]_2, [\sigma^2]_2$.

Commit. Form $F(X,Y) = \sum_{i,j<b} F_{ij} X^i Y^j$ and output $C_f = [F(\tau,\sigma)]_1$.

Open. On input $(\mathsf{ck}, \mathsf{vk}, F, C_f, (u,v), y)$ with $C_f = \mathsf{Commit}(\mathsf{ck}, F)$ and $y = F(u,v)$, proceed as follows.

(a) Form g as in (6). Compute the auxiliary polynomial $S_0(X)$ for the following constant-term identity associated with $g(u) = y$

$$g(X)T_u(X^{-1}) + g(X^{-1})T_u(X) = 2y + XS_0(X) + X^{-1}S_0(X^{-1}).$$

and send $C_0 = [S_0(\tau)]_1$. Define $T_0 = (\mathsf{ppid}, \mu, C_f, u, v, y, C_0)$ and derive $(\rho, c_\rho) = \mathsf{DeriveChallenge}(\texttt{"carina-r"}, T_0, \{0, 1, -1\})$.

(b) Send $a_0 = g(\rho)$, $a_1 = g(\rho^{-1})$, and $b_0 = S_0(\rho)$. Define $T_1 = (T_0, c_\rho, \rho, a_0, a_1, b_0)$ and derive $(\lambda, c_\lambda) = \mathsf{DeriveChallenge}(\texttt{"carina-lambda"}, T_1, \{0\})$.

(c) Set $\phi(Y) = F(\rho, Y) + \lambda F(\rho^{-1}, Y)$ and $y_1 = a_0 + \lambda a_1$. Compute the corresponding auxiliary polynomial $S_1(Y)$ satisfying

$$\phi(Y)T_v(Y^{-1}) + \phi(Y^{-1})T_v(Y) = 2y_1 + YS_1(Y) + Y^{-1}S_1(Y^{-1})$$

and send $C_1 = [S_1(\sigma)]_1$. Define $T_2 = (T_1, c_\lambda, \lambda, C_1)$ and derive $(\zeta, c_\zeta) = \mathsf{DeriveChallenge}(\texttt{"carina-s"}, T_2, \{0, 1, -1\})$.

(d) Send $c_{ij} = F(\rho^{(-1)^i}, \zeta^{(-1)^j})$ for $i, j \in \{0, 1\}$ and $b_1 = S_1(\zeta)$. Define $T_3 = (T_2, c_\zeta, \zeta, c_{00}, c_{01}, c_{10}, c_{11}, b_1)$ and derive $(\eta, c_\eta) = \mathsf{DeriveChallenge}(\texttt{"carina-eta"}, T_3, \{0\})$. Let $T_4 = (T_3, c_\eta, \eta)$.

(e) Recover $S_0(\rho^{-1})$ and $S_1(\zeta^{-1})$ by the above two equations. Interpolate I from the c_{ij} on $A \times B$, L_0 from S_0 on A, and L_1 from S_1 on B. Regard $S_0 - L_0$ as a bivariate polynomial that depends only on X, and $S_1 - L_1$ as one that depends only on Y. Set

$$E_\eta = (F - I) + \eta(S_0 - L_0) + \eta^2(S_1 - L_1).$$

Divide canonically $E_\eta = Z_A(X)U_\eta(X,Y) + Z_B(Y)V_\eta(X,Y)$ and send $\mathit{\Pi}_X = [U_\eta(\tau,\sigma)]_1$ and $\mathit{\Pi}_Y = [V_\eta(\tau,\sigma)]_1$. The proof is $(C_0, C_1, a_0, a_1, b_0, c_{00}, c_{01}, c_{10}, c_{11}, b_1, \mathit{\Pi}_X, \mathit{\Pi}_Y)$.

Verify. On input $(\mathsf{vk}, C_f, (u,v), y, \pi)$, reconstruct T_0, T_1, T_2, T_3, T_4 from the statement and proof. Recompute the four challenge–counter pairs with the displayed $\mathsf{DeriveChallenge}$ calls and rejection sets, appending each counter before its challenge, and reject malformed data. Recover the two omitted values $S_0(\rho^{-1})$ and $S_1(\zeta^{-1})$ using the above two Laurent equations. Reconstruct I, L_0, L_1 and

$$C_E = C_f + \eta C_0 + \eta^2 C_1 - [I(\tau,\sigma) + \eta L_0(\tau) + \eta^2 L_1(\sigma)]_1.$$

Form $Z_{A,2} = [\tau^2]_2 - (\rho + \rho^{-1})[\tau]_2 + [1]_2$ and $Z_{B,2} = [\sigma^2]_2 - (\zeta + \zeta^{-1})[\sigma]_2 + [1]_2$. Accept if and only if the scalar identities hold and

$$e(C_E, [1]_2) = e(\mathit{\Pi}_X, Z_{A,2})e(\mathit{\Pi}_Y, Z_{B,2}). \tag{7}$$

where each Q_* counts candidate random-oracle queries under the corresponding domain, including rejected counters. Its evaluation-binding advantage against algebraic adversaries in the ROM is at most

$$2\delta_{\mathsf{Carina}} + \varepsilon_Q^{\mathsf{qdlog}} + \frac{2b-2}{|\mathbb{F}|-1}.$$

Proof (Sketch). The honest polynomials S_0 and S_1 satisfy the two instances of (2) associated with $g(u) = y$ and $\phi(v) = y_1$. Their interpolated residuals vanish on $A \times B$, so canonical division exists and (7) holds. This establishes perfect evaluation completeness.

For the reduction, receive the bilinear Q-DLOG powers of a hidden x and sample $a, c \xleftarrow{\$} \mathbb{F}^*$ and $b, d \xleftarrow{\$} \mathbb{F}$. Set implicitly $\tau = ax + b$ and $\sigma = cx + d$. Every tensor monomial $\tau^i \sigma^j$ in the $\mathbb{G}_1$ SRS has degree at most $i + j \le 2b - 2 = Q$ as a polynomial in x, and the five $\mathbb{G}_2$ elements have degree at most two. The reduction therefore computes the complete Carina SRS from the Q-DLOG instance. The random translations b, d make (τ, σ) independent and uniform, so the simulation is exact.

Two distinct tensor-supported polynomials that yield the same commitment give a known nonzero difference Δ with $\Delta(\tau, \sigma) = 0$ and total degree at most $2b-2$. The univariate specialization $\Delta(aZ+b, cZ+d)$ has x as a root. Its leading coefficient is the highest homogeneous component of Δ evaluated at (a, c), and hence it is nonzero except with probability $(2b-2)/(|\mathbb{F}|-1)$. Otherwise factoring the specialization and testing its roots against $[x]_1$ recovers x, breaking bilinear Q-DLOG.

For knowledge soundness, the commitment explanation yields F, while the remaining AGM representations yield S_0, S_1, U_η, V_η. The accepted pairing equation implies

$$D(\tau, \sigma) = 0, \qquad D = E_\eta - Z_A U_\eta - Z_B V_\eta.$$

The degree bounds give total degree at most $2b$ for D. If D is formally nonzero, the same affine specialization has x as a root and is nonzero except with probability $2b/(|\mathbb{F}|-1)$. Factoring it otherwise breaks bilinear Q-DLOG. Consequently $D = 0$ formally, except with the first two error terms of (8), and $E_\eta \in \langle Z_A, Z_B \rangle$. In the quotient ring $\mathcal{Q}_{A,B} = \mathbb{F}[X, Y]/\langle Z_A, Z_B \rangle \cong \mathbb{F}^4$, this gives

$$[F - I] + \eta[S_0 - L_0] + \eta^2[S_1 - L_1] = 0.$$

Unless all three residuals vanish on the grid, a nonzero coordinate has at most two admissible values of η.

If $g(u) = y$ is false, its Laurent defect has degree at most $2b$ after clearing negative exponents. The defect for $\phi(v) = y_1$, after separation by λ, contributes degree at most $3b - 1$. Schwartz–Zippel in transcript order, followed by a query-tree union bound, gives (8). The extracted F determines the unique multilinear polynomial in (1); two accepted values force either a knowledge-soundness failure or a commitment-binding violation. Appendix B.2 gives the complete reduction.

6.3 Efficiency

Computing $F(X, v)$ and $F(\rho, Y) + \lambda F(\rho^{-1}, Y)$, together with canonical grid division, uses $O(N)$ field operations. Computing S_0 and S_1 adds $O(\sqrt{N} \log N)$ work. To compute the MSM scale, their commitments use $2b-2$ scalars, and the two quotient supports contain $N-4$ scalars in total. Thus opening uses exactly $N + 2\sqrt{N} - 6$ MSM scalars.

The proof is $4\mathbb{G}_1 + 8\mathbb{F}$: 448 bytes of cryptographic payload and 452 serialized bytes in our implementation. The verifier performs one three-term pairing product, two dynamic $\mathbb{G}_2$ scalar multiplications, a constant-size $\mathbb{G}_1$ linear combination, and $O(\mu)$ field operations. The SRS contains N tensor elements in $\mathbb{G}_1$ and five elements in $\mathbb{G}_2$. Section 7.1 compares these costs with prior work.

7 Implementation and Evaluation

7.1 Theoretical Comparison

We first compare single openings independently of machine-level effects. Because all MLPCS enjoy the same sum-check-based batching technique that incurs independent cost, the single opening comparison is sufficient to predict their performance in batched contexts. Table 1 counts the protocol variants implemented in our benchmark. Opening excludes the initial commitment to f. "MSM scalars" is the sum of the scalar-vector lengths of all opening MSMs. The verifier column counts pairing terms and dynamic second-group scalar multiplications; Samaritan and CHOPIN finalize those terms in two independent product equations. Constant-size interpolation MSMs and serialization metadata are omitted.

Table 1. Theoretical costs of the evaluated single-opening implementations. Here P denotes a pairing term and M_2 a dynamic $\mathbb{G}_2$ multiplication. For Samaritan, $m = 2^{\kappa} \approx \mu$ and $\ell = N/m$.

Scheme	Proof	MSM scalars	Field ops	Verifier
mKZG [27]	$\mu\mathbb{G}_1$	$N-1$	$O(N)$	$(\mu+1)P + \mu M_2$
Gemini [6]	$(\mu+1)\mathbb{G}_1 + 2\mu\mathbb{F}$	$3N + O(1)$	$O(N)$	$2P + M_2$
Zeromorph [21]	$(\mu+2)\mathbb{G}_1$	$3N-2$	$O(N)$	$2P + M_2$
Samaritan [17]	$7\mathbb{G}_1 + \mathbb{F}$	$3N + 2\ell + 2m + O(1)$	$O(N)$	$4P + M_2$
Mercury [13]	$8\mathbb{G}_1 + 6\mathbb{F}$	$2N + O(\sqrt{N})$	$O(N)$	$2P$
CHOPIN [3]	$7\mathbb{G}_1 + 7\mathbb{F}$	$N + 6\sqrt{N} + O(1)$	$O(N)$	$5P + 2M_2$
Carina	$4\mathbb{G}_1 + 8\mathbb{F}$	$N + 2\sqrt{N} - 6$	$O(N)$	$3P + 2M_2$
Vela	$2\mathbb{G}_1 + 3\mathbb{F}$	$2N-3$	$O(N \log N)$	$2P + M_2$

The aggregate counts conceal different MSM shapes. mKZG uses a geometric sequence of quotient MSMs of lengths $N/2, N/4, \ldots, 1$. Gemini commits to folds of total length $N-2$ and performs two Shplonk MSMs of lengths approximately

N and $N-1$. Zeromorph combines quotient commitments whose lengths sum to $N-1$ with size-N and size-$(N-1)$ MSMs. The evaluated Samaritan path has three size-N commitments plus lower-order commitments. Mercury has two leading size-N operations. By contrast, Vela has exactly two opening MSMs, of lengths $N-1$ and $N-2$.

The four Carina opening MSMs, used for C_0, C_1, Π_X, and Π_Y, have lengths

$$\sqrt{N}-1, \quad \sqrt{N}-1, \quad (\sqrt{N}-2)\sqrt{N}, \quad 2(\sqrt{N}-2),$$

which sum to $N+2\sqrt{N}-6$. Thus mKZG has the least prover group work but has an $O(\mu)$ proof and verifier; CHOPIN and Carina add only lower-order work to one size-N MSM; Mercury and Vela require roughly two size-N MSMs; and Gemini, Zeromorph, and Samaritan are near three size-N MSMs in the evaluated variants.

Field work separates schemes with similar group counts. Mercury, Samaritan, CHOPIN, and Carina use a constant number of full-vector passes. Computing $F(X,v)$, $F(\rho,Y)+\lambda F(\rho^{-1},Y)$, and the canonical grid quotients costs $O(N)$ in Carina, while computing S_0 and S_1 adds $O(\sqrt{N}\log N)$ lower-order work. Vela constructs its symmetric witness through μ structured shift-and-add passes and therefore uses $O(N\log N)$ field operations. This explains why its group work is comparable to Mercury but its large-instance opening is slower. Equal aggregate MSM lengths can still have different timings because the MSMs have different shapes, SRS layouts, and scalar orders.

Verifier costs follow a different ordering. mKZG has $\mu+1$ pairing terms and μ evaluation-dependent $\mathbb{G}_2$ multiplications. Gemini, Zeromorph, Mercury, and Vela each reduce to one two-term product. Samaritan checks a KZG opening and a separate shift relation, while CHOPIN finalizes independent three-term and two-term products. Carina has one three-term product and two dynamic $\mathbb{G}_2$ multiplications. It should therefore verify slower than Mercury and Vela, but faster than constructions requiring two independent pairing products. Finally, Vela minimizes communication with $2\mathbb{G}_1+3\mathbb{F}$, while Carina spends two additional group and five field elements to reduce prover work from roughly two size-N MSMs to roughly one.

7.2 Implementation and Methodology

We implement all schemes in curve-generic Rust using arkworks 0.4.2. The benchmark instantiates BLS12-381 and uses the same dense multilinear representation, variable-base MSM backend, and Merlin Fiat–Shamir transcript for every backend. Group and field elements use canonical compressed arkworks serialization. The executable specification in (9) uses unbiased rejection sampling with an exhaustive 64-bit counter. The benchmark prototype instead instantiates Fiat–Shamir through Merlin's challenge interface: Carina retries inadmissible values at most 64 times, whereas Vela aborts if its first point challenge is inadmissible. Thus the prototype is not a byte-for-byte implementation of (9); the security theorems apply to the specification in the paper, while the implementation uses Merlin as a separate random-oracle instantiation. Only the probability

of hitting an inadmissible field element is negligible, and neither retry limit nor abort condition was reached in our experiments. We implement Vela and Carina as separate PCS backends behind the same curve-generic interface, with the common Setup/Commit/Open/Verify API defined in Definition 3. For odd μ, the Carina implementation uses $M_L = 2^{\lceil \mu/2 \rceil}$ and $M_R = 2^{\lfloor \mu/2 \rfloor}$ with $M_L M_R = N$; this is an engineering generalization of the balanced $\sqrt{N} \times \sqrt{N}$ protocol used in the theory.

The comparison includes mKZG [27], Gemini [6], Zeromorph [21], Samaritan [17], Mercury [13], and CHOPIN [3]. We implemented missing or non-reusable baseline components in the same framework and tested honest openings, malformed proofs, transcript tampering, and small-instance polynomial identities. The machine is an Apple M2 Pro with 12 CPU cores and 16 GB of memory, running macOS, a release build of `rustc 1.97.0`, and 12 Rayon threads. A fixed seed, 20260713, generates the same polynomial and point for every backend at each parameter.

Each backend and parameter pair runs in a fresh process, and backends are run serially. Within that process, SRS generation, trim, commitment, and opening are each measured five times. The benchmark records every sample and reports the median. Verification is invoked 100 times and summarized in the same way. Five repetitions keep the full matrix practical because the heavy phases at $\mu = 20$ take seconds, while providing a robust estimate of typical runtime. The opening timer starts with an existing commitment and uses a commitment-aware API, excluding both the initial commitment and a redundant statement recommitment. Proof size is compressed canonical serialization, including backend framing.

7.3 Empirical Results

Table 2 reports all measured phases at $\mu = 20$. Figure 1 shows the complete range $\mu \in \{8, 10, 12, 14, 16, 18, 20\}$. SRS cost varies because the schemes publish different SRS shapes; in particular, Zeromorph generates substantially more second-group material. Commitment times cluster closely because every backend ultimately performs one dense size-N MSM.

The measurements support the theoretical trade-off. Vela has the smallest proof at 200 bytes, 1.88x smaller than Samaritan's 376-byte proof, the next smallest. Its 2.058 ms verification median is the lowest measured at $\mu = 20$: 2.67x faster than mKZG, 2.00x faster than CHOPIN, and 1.07x faster than Mercury. Its 3840.611 ms opening median is not prover-optimal: the $O(N \log N)$ witness computation makes it slower than Mercury and the one-MSM schemes. It nevertheless remains 1.21x faster than Gemini and 1.23x faster than Samaritan; its 0.7% difference from Zeromorph is too small to interpret as a stable ordering.

Carina, mKZG, and CHOPIN form the leading opening-performance tier at $\mu = 20$, with medians of 1485.955 ms, 1471.259 ms, and 1490.272 ms, respectively. Their spread is at most 1.3%, so we do not interpret it as a meaningful ordering from five samples. Within this tier, Carina has a 452-byte proof and a 2.667 ms verification median, compared with 968 bytes and 5.495 ms for mKZG

Table 2. Single-opening comparison at $\mu = 20$. Entries are medians of five repetitions for SRS generation, commitment, and opening, and 100 repetitions for verification. Arithmetic means and raw samples are retained in the experiment artifact.

Backend	SRS (ms)	Commit (ms)	Open (ms)	Verify (ms)	Proof (B)
mKZG	3606.344	1355.781	1471.259	5.495	968
Gemini	1242.622	1343.425	4649.623	2.404	2312
Zeromorph	10196.858	1388.552	3816.024	2.145	1064
Samaritan	1333.952	1374.730	4732.237	3.302	376
Mercury	1377.427	1395.822	2914.251	2.192	584
CHOPIN	1363.588	1411.106	1490.272	4.121	564
Carina	1366.936	1403.123	1485.955	2.667	452
Vela	1349.441	1397.884	3840.611	2.058	200

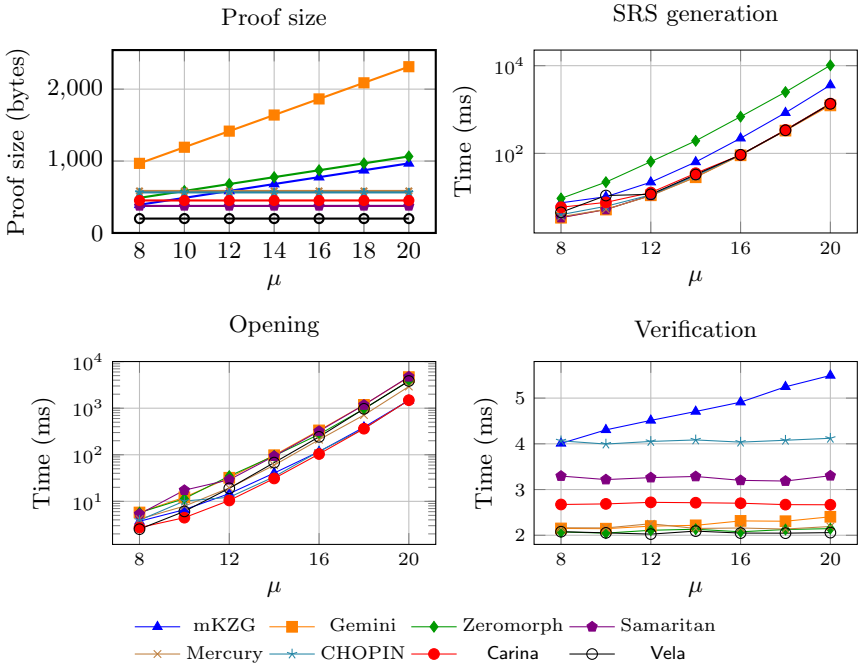


Fig. 1. Proof size and median SRS-generation, opening, and verification times over the evaluated parameter range. SRS generation and opening use logarithmic axes; proof size and verification use linear axes. SRS generation and opening use five repetitions; verification uses 100.

and 564 bytes and 4.121 ms for CHOPIN. Thus Carina combines prover performance comparable to the most prover-efficient baselines with better communication and verification in this experiment.

The comparison with Mercury exposes the complementary trade-off more clearly. Carina's 1485.955 ms opening median is 1.96x faster than Mercury's 2914.251 ms, whereas its verification median is 0.475 ms slower because it adds one pairing term and two dynamic $\mathbb{G}_2$ multiplications. Its proof is also smaller, 452 bytes rather than 584 bytes. The broader separation from the two- and three-MSM prover designs likewise matches the predicted group work.

8 Conclusion

We developed two multilinear polynomial commitments from Mercury's Laurent-polynomial constant-term representation. Vela combines a half-length auxiliary polynomial with a two-point KZG opening to minimize proof size and verifier work. Carina applies the constant-term reduction once in each coordinate and uses a bivariate grid opening to approach the prover cost of direct multilinear commitments while preserving a constant-size proof. Their measured performance follows these theoretical trade-offs: Carina matches the leading prover-performance tier represented by mKZG and CHOPIN while improving their proof size and verifier time, and it roughly halves Mercury's opening time at the cost of a modest verifier slowdown.

The constructions also leave clear limitations. Vela performs $O(N \log N)$ field work, and the security analyses for both schemes are in the AGM and ROM under a falsifiable bilinear q-type DLOG assumption. Proving standard-model knowledge soundness for the exact Vela and Carina protocols remains open. Neither scheme is hiding or post-quantum secure. Adding zero knowledge without changing the cost profile, proving the exact protocols in the standard model, and optimizing the dynamic second-group operations in recursive verification remain directions for further work.

References

1. Arnon, G., Chiesa, A., Fenzi, G., Yogev, E.: Stir: Reed–solomon proximity testing with fewer queries. In: Advances in Cryptology – CRYPTO 2024, Part X. vol. 14929, p. 380–413. Springer International Publishing (2024). https://doi.org/10.1007/978-3-031-68403-6_13, https://eprint.iacr.org/2024/390
2. Arnon, G., Chiesa, A., Fenzi, G., Yogev, E.: Whir: Reed–solomon proximity testing with super-fast verification. In: Advances in Cryptology – EUROCRYPT 2025, Part IV. vol. 15604, p. 214–243. Springer International Publishing (2025), https://eprint.iacr.org/2024/1586
3. Belohorec, J., Hubáček, P., Kalsta, A., Mašková, K.: Chopin: Optimal pairing-based multilinear polynomial commitments from bivariate kzg (2026), https://eprint.iacr.org/2026/480
4. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Fast reed-solomon interactive oracle proofs of proximity. In: 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018). vol. 107, pp. 14:1–14:17. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2018). https://doi.org/10.4230/LIPIcs.ICALP.2018.14

5. Bootle, J., Cerulli, A., Chaidos, P., Groth, J., Petit, C.: Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In: Advances in Cryptology – CRYPTO 2016, Part II. vol. 9815, p. 327–357. Springer Berlin Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_12, https://eprint.iacr.org/2016/263
6. Bootle, J., Chiesa, A., Hu, Y., Orrù, M.: Gemini: Elastic snarks for diverse environments. In: Advances in Cryptology – EUROCRYPT 2022, Part II. vol. 13276, p. 427–457. Springer International Publishing (2022). https://doi.org/10.1007/978-3-031-07085-3_15, http://eprint.iacr.org/2022/420
7. Bootle, J., Chiesa, A., Sotiraki, K.: Sumcheck arguments and their applications. In: Advances in Cryptology – CRYPTO 2021, Part I. vol. 12825, p. 742–773. Springer International Publishing (2021). https://doi.org/10.1007/978-3-030-84242-0_26, http://eprint.iacr.org/2021/333
8. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA. p. 315–334. IEEE Computer Society (2018). https://doi.org/10.1109/SP.2018.00020
9. Bünz, B., Fisch, B., Szepieniec, A.: Transparent snarks from dark compilers. In: Advances in Cryptology – EUROCRYPT 2020. pp. 677–706. Springer International Publishing (2020). https://doi.org/10.1007/978-3-030-45721-1_24
10. Chen, B., Bünz, B., Boneh, D., Zhang, Z.: Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In: Advances in Cryptology – EUROCRYPT 2023, Part II. p. 499–530. Springer International Publishing (2023). https://doi.org/10.1007/978-3-031-30617-4_17, https://eprint.iacr.org/2022/1355
11. Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.: Marlin: Preprocessing zksnarks with universal and updatable srs. In: Advances in Cryptology – EUROCRYPT 2020. pp. 738–768. Springer International Publishing (2020). https://doi.org/10.1007/978-3-030-45721-1_26
12. Chmel, P., Hubáček, P., Stejskal, D.: Knowledge soundness of polynomial commitments in the algebraic group model does not guarantee extractability. IACR Cryptology ePrint Archive (2026), https://eprint.iacr.org/2026/284
13. Eagen, L., Gabizon, A.: Mercury: A multilinear polynomial commitment scheme with constant proof size and no prover ffts (2025), https://eprint.iacr.org/2025/385
14. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings. vol. 263, pp. 186–194. Springer (1986). https://doi.org/10.1007/3-540-47721-7_12
15. Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications (2017), https://eprint.iacr.org/2017/620
16. Gabizon, A., Williamson, Z.J., Ciobotaru, O.: Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge (2019), https://eprint.iacr.org/2019/953
17. Ganesh, C., Patranabis, S., Singh, N.: Samaritan: linear-time prover snark from new multilinear polynomial commitments. In: Advances in Cryptology – ASIACRYPT 2025. vol. 16249, p. 456–489. Springer International Publishing (2025), https://eprint.iacr.org/2025/419
18. Golovnev, A., Lee, J., Setty, S., Thaler, J., Wahby, R.S.: Brakedown: Linear-time and field-agnostic snarks for r1cs. In: Advances in Cryptology – CRYPTO 2023.

vol. 14084, p. 193–226. Springer International Publishing (2023). https://doi.org/10.1007/978-3-031-38551-3_7, https://eprint.iacr.org/2021/1043
19. Haböck, U.: Basefold in the list decoding regime (2024), https://eprint.iacr.org/2024/1571
20. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: Advances in Cryptology – ASIACRYPT 2010. vol. 6477, p. 177–194. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_11
21. Kohrita, T., Towa, P.: Zeromorph: Zero-knowledge multilinear-evaluation proofs from homomorphic univariate commitments (2023), https://eprint.iacr.org/2023/917
22. Kothapalli, A., Setty, S.: Hypernova: Recursive arguments for customizable constraint systems. In: Advances in Cryptology – CRYPTO 2024, Part X. vol. 14929, p. 345–379. Springer International Publishing (2024). https://doi.org/10.1007/978-3-031-68403-6_11, https://eprint.iacr.org/2023/573
23. Lee, J.: Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In: Theory of Cryptography – TCC 2021. vol. 13043, p. 1–34. Springer International Publishing (2021). https://doi.org/10.1007/978-3-030-90453-1_1, http://eprint.iacr.org/2020/1274
24. Liu, T., Zhang, Z., Zhang, Y., Hu, W., Zhang, Y.: Ceno: non-uniform, segment and parallel zero-knowledge virtual machine. Journal of Cryptology **38**(2), 17 (2025). https://doi.org/10.1007/s00145-024-09533-2, https://doi.org/10.1007/s00145-024-09533-2
25. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. In: 31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I. p. 2–10. IEEE Computer Society (1990). https://doi.org/10.1109/FSCS.1990.89518
26. Maller, M., Bowe, S., Kohlweiss, M., Meiklejohn, S.: Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019. p. 2111–2128. ACM (2019). https://doi.org/10.1145/3319535.3339817
27. Papamanthou, C., Shi, E., Tamassia, R.: Signatures of correct computation. Lecture Notes in Computer Science (2013). https://doi.org/10.1007/978-3-642-36594-2_13
28. Ràfols, C., Zapico, A.: An algebraic framework for universal and updatable snarks. In: Annual International Cryptology Conference. p. 774–804. Springer (2021)
29. Setty, S.: Spartan: Efficient and general-purpose zksnarks without trusted setup. In: Annual International Cryptology Conference. p. 704–737. Springer (2020)
30. Xie, T., Zhang, J., Zhang, Y., Papamanthou, C., Song, D.: Libra: Succinct zero-knowledge proofs with optimal prover computation. In: Advances in Cryptology – CRYPTO 2019, Part III. vol. 11694, p. 733–764. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-26954-8_24
31. Zeilberger, H., Chen, B., Fisch, B.: Basefold: Efficient field-agnostic polynomial commitment schemes from foldable codes. In: Advances in Cryptology – CRYPTO 2024, Part X. vol. 14929, p. 138–169. Springer International Publishing (2024). https://doi.org/10.1007/978-3-031-68403-6_5, https://eprint.iacr.org/2023/1705
32. Zhang, Y., Szepieniec, A., Zhang, R., Sun, S.F., Wang, G., Gu, D.: Voproof: Efficient zksnarks from vector oracle compilers. In: Proceedings of the

2022 ACM SIGSAC Conference on Computer and Communications Security. p. 3195–3208. Association for Computing Machinery (2022). https://doi.org/10.1145/3548606.3559387, https://doi.org/10.1145/3548606.3559387

A Fiat-Shamir Specifications

We use the following single challenge-generation procedure in both schemes.

Let enc be an injective canonical byte encoding. Field elements use their fixed-width canonical integer encoding, group elements use canonical compressed encoding, vectors and byte strings are length-prefixed, tuples are encoded by concatenating their length-prefixed components, and public integers use unsigned big-endian encoding. For $c \in \{0, \dots, 2^{64} - 1\}$, let $\mathsf{ctr}(c)$ be its fixed-width eight-byte unsigned big-endian encoding.

Fix a byte hash H with output length h bytes. For a byte string m and output length L, define $\mathsf{Expand}_H(m, L)$ as the first L bytes of

$$H(\texttt{0x00} \| m \| \mathsf{ctr}(0)) \| H(\texttt{0x00} \| m \| \mathsf{ctr}(1)) \| \cdots .$$

Let $p = |\mathbb{F}|$ and $L = \lceil \log_2(p)/8 \rceil$. For an ASCII domain string d, an ordered transcript prefix T, and an explicit counter $c \in \mathbb{N}$, define

$$\begin{aligned} \mathsf{HashToField}(d, T, c) : \quad & m_c = \mathsf{enc}(\texttt{"VC-PCS-v1"}, d, T, \mathsf{ctr}(c)); \\ & w_c = \mathsf{OS2IP}(\mathsf{Expand}_H(m_c, L)); \\ & \text{return } w_c \in \mathbb{F} \text{ if } w_c < p, \text{ and return } \perp \text{ otherwise.} \end{aligned} \tag{9}$$

Here $\mathsf{OS2IP}$ interprets a byte string as an unsigned big-endian integer. The integer $w_c < p$ is identified with its canonical element of $\mathbb{F}$; returning $\perp$ rather than reducing modulo p avoids bias. For a rejection set $R \subseteq \mathbb{F}$, define

$$\mathsf{DeriveChallenge}(d, T, R) = (u, c),$$

where $c \in \{0, \dots, 2^{64}-1\}$ is the least integer for which $u = \mathsf{HashToField}(d, T, c) \neq \perp$ and $u \notin R$. Prover and verifier therefore both start at counter zero and abort if the counter space is exhausted. A returned challenge record is encoded as the counter followed by the field element, and that record is appended to the transcript before any subsequent prover message. Define the fixed-size public-parameter identifier

$$\mathsf{ppid} = H\big(\mathsf{enc}(\texttt{"VC-PCS-v1/parameters"}, \mathsf{vk})\big).$$

Hashing this identifier, rather than the full SRS, binds each proof to its verification parameters without adding $O(N)$ hashing to every challenge. Thus the ordered tuples displayed in Protocols 1 and 2 determine every hash input byte-for-byte; counters need not be transmitted because the verifier recomputes them. In the analysis, H is a classical random oracle, so the accepted output is uniform over $\mathbb{F} \setminus R$.

B Complete Security Proofs

B.1 Proof of Theorem 1

For perfect completeness, let C_f commit to the honest encoding f_v and let $y = f(r)$. The constant coefficient of $f_v(X)T_r(X^{-1}) + f_v(X^{-1})T_r(X) - 2y$ is zero. Its invariance under $X \mapsto X^{-1}$ therefore gives the polynomial $\bar{h}$ used by the prover, and the Laurent identity

$$z\bar{h}(z) + z^{-1}\bar{h}(z^{-1}) + 2y = f_v(z)T_r(z^{-1}) + f_v(z^{-1})T_r(z)$$

holds at every admissible z. Hence the value w_1 recovered via (3) equals $\bar{h}(z^{-1})$, so R_h is the correct interpolant of $\bar{h}$ at $\{z, z^{-1}\}$. The polynomial $f_v + \alpha\bar{h} - R_f - \alpha R_h$ vanishes at z and z^{-1}, so it is divisible by $Z = (X - z)(X - z^{-1})$. The KZG equation then proves (4), and the honest verifier accepts with probability one.

For commitment binding, suppose an algebraic adversary outputs distinct f_v, f_v' of degree below N with the same commitment. Then the known nonzero polynomial $\Delta = f_v - f_v'$ satisfies $\Delta(\tau) = 0$. Factoring Δ produces at most $N - 1$ candidates; testing their first-group encodings against $[\tau]_1$ recovers the trapdoor. The reduction embeds the Vela key in an $(N+1)$-DLOG instance, so the binding advantage is at most ε_{N+1}.

Consider an algebraic prover in the AGM+ROM knowledge-soundness experiment of Definition 6. The commitment explanation recovers f_v, and flattening the remaining AGM representations to the public SRS reveals $\bar{h}, q \in \mathbb{F}[X]$ for C_h, π, each of degree at most $N-1$. Let R_f and R_h be the degree-one interpolants formed from the three transmitted evaluations v_0, v_1, w_0 and the recovered value w_1 of (3), and set $R = R_f + \alpha R_h$ and $Z = (X - z)(X - z^{-1})$. The pairing check implies

$$W(\tau) = 0, \qquad W(X) = f_v(X) + \alpha\bar{h}(X) - R(X) - q(X)Z(X),$$

with $\deg W \le N + 1$.

Suppose W is not the zero polynomial. Its coefficients are known to the reduction, which factors W over $\mathbb{F}$ and enumerates its at most $N + 1$ roots. The unique candidate t satisfying $[t]_1 = [\tau]_1$ reveals τ. A reduction given an $(N+1)$-DLOG instance can provide the adversary with the subset of powers used by Vela, so this event has probability at most ε_{N+1}. Outside that event, $W = 0$ formally.

At each $p \in \{z, z^{-1}\}$, evaluating $W = 0$ gives

$$\big(f_v(p) - \widehat{v}_p\big) + \alpha\big(\bar{h}(p) - \widehat{w}_p\big) = 0,$$

where the hats denote the claimed values: three transmitted values and, for $\widehat{w}_{z^{-1}}$, the recovered value of (3). The recovered value is uniquely determined by $(z, y, v_0, v_1, w_0, T_r(z), T_r(z^{-1}))$, so all four claimed values are fixed before α is sampled. If either coefficient is nonzero, this linear equation in α has at most one solution. A union bound over z and z^{-1} gives batching error at most $2/|\mathbb{F}|$.

Hence, outside these two bad events, the claimed values of f_v and $\bar{h}$ at z and z^{-1} are genuine.

The recovery relation makes the claimed tuple a root of the Laurent defect

$$D(X) = X\bar{h}(X) + X^{-1}\bar{h}(X^{-1}) + 2y - f_v(X)T_r(X^{-1}) - f_v(X^{-1})T_r(X)$$

at $X = z$ by definition: substituting $w_1 = z(v_0T_r(z^{-1}) + v_1T_r(z) - zw_0 - 2y)$ into $zw_0 + z^{-1}w_1 + 2y - v_0T_r(z^{-1}) - v_1T_r(z)$ gives zero identically. Since the batched KZG authentication certifies the three transmitted evaluations and the recovered one as the genuine evaluations of f_v and $\bar{h}$, the actual defect satisfies $D(z) = 0$. Its support is contained in $[-N, N]$: the two witness terms contribute only nonzero exponents, while each Laurent product has support $[-(N-1), N-1]$. Its constant coefficient is

$$[X^0]D = 2y - 2f(r).$$

For a false claim, odd characteristic therefore makes D nonzero. The ordinary polynomial $P(X) = X^N D(X)$ has degree at most $2N$ and is fixed before z is sampled because C_h precedes the challenge. Sampling z uniformly from $\mathbb{F} \setminus \{0, 1, -1\}$ gives

$$\Pr[D(z) = 0 \mid y \neq f(r)] \leq \frac{2N}{|\mathbb{F}| - 3}.$$

This proves the interactive bound. For Fiat–Shamir, associate to every well-formed candidate hash query, including a rejected counter, the roots of its fixed nonzero residual. A classical adversary making Q such queries hits one of these bad answers with probability at most Q times the statistical bound. Revelation from the AGM explanations is straight-line, so no forking step is required. This proves (5), and the extracted f_v is a witness for (1). If two proofs for one (C_f, r) and distinct y, y' are accepted, either one knowledge-soundness failure occurs or the extracted polynomials are distinct and have the same commitment. A union bound over the two knowledge errors and commitment binding gives AGM+ROM evaluation-binding advantage at most $2\delta_{\mathsf{Vela}}(Q) + \varepsilon_{N+1}$.

B.2 Proof of Theorem 2

For perfect completeness, S_0 and S_1 satisfy the two instances of (2) used for $g(u) = y$ and $\phi(v) = y_1$. The interpolants I, L_0, L_1 agree with F, S_0, S_1 on $A \times B$ after their natural bivariate embeddings. Hence E_η vanishes on the grid. Canonical division by Z_A and then Z_B therefore yields U_η, V_η satisfying the grid-ideal decomposition, and bilinearity gives (7). The honest verifier accepts with probability one.

We first reduce the bivariate pairing relation to the falsifiable bilinear q-type DLOG assumption. Let $Q = 2b - 2$ and let a reduction $\mathcal{B}$ receive a bilinear Q-DLOG instance

$$\left(([x^k]_1)_{k=0}^{Q}, ([x^k]_2)_{k=0}^{Q}\right)$$

for hidden $x \xleftarrow{\$} \mathbb{F}^*$. It samples $r_X, r_Y \xleftarrow{\$} \mathbb{F}^*$ and $s_X, s_Y \xleftarrow{\$} \mathbb{F}$, and implicitly sets

$$\tau = r_X x + s_X, \qquad \sigma = r_Y x + s_Y. \tag{10}$$

For $0 \le i, j < b$, the polynomial $(r_X Z + s_X)^i (r_Y Z + s_Y)^j$ has degree at most $i + j \le 2b - 2 = Q$. Expanding it in the monomial basis lets $\mathcal{B}$ compute $[\tau^i \sigma^j]_1$ from its first-group Q-DLOG powers. It likewise computes $[1]_2, [\tau]_2, [\tau^2]_2, [\sigma]_2, [\sigma^2]_2$ from the first three second-group powers. For every fixed x, the uniform translations s_X, s_Y make τ, σ independent and uniform in $\mathbb{F}$. Thus $\mathcal{B}$ simulates the Carina SRS perfectly.

Lemma 1 (Random affine specialization). *Let $H \in \mathbb{F}[X, Y]$ be a nonzero polynomial of total degree d. Suppose H is determined by an algebraic adversary's view of the simulated SRS and its coins. Under (10),*

$$\Pr\big[H(r_X Z + s_X, r_Y Z + s_Y) \equiv 0\big] \le \frac{d}{|\mathbb{F}| - 1}.$$

Proof. Condition on x, the simulated values (τ, σ), the adversary's coins, and its complete view. For each fixed nonzero (r_X, r_Y), the translations compatible with these values are uniquely $s_X = \tau - r_X x$ and $s_Y = \sigma - r_Y x$. Hence (r_X, r_Y) remains uniform in $(\mathbb{F}^*)^2$ after conditioning on the view; in particular the revealed H is independent of the slopes.

Let H_d be the nonzero homogeneous component of H of highest total degree d. The coefficient of Z^d in $H(r_X Z + s_X, r_Y Z + s_Y)$ is $H_d(r_X, r_Y)$. Schwartz–Zippel over $(\mathbb{F}^*)^2$ bounds the probability that this coefficient vanishes by $d/(|\mathbb{F}| - 1)$. Outside this event the univariate specialization is nonzero.

For commitment binding, suppose distinct tensor-supported polynomials F and F' yield the same C_f. Then $\Delta = F - F'$ is nonzero, has total degree at most $2b - 2$, and satisfies $\Delta(\tau, \sigma) = 0$. Define

$$P_\Delta(Z) = \Delta(r_X Z + s_X, r_Y Z + s_Y).$$

By Lemma 1, P_Δ is nonzero except with probability $(2b-2)/(|\mathbb{F}| - 1)$, while the commitment equality gives $P_\Delta(x) = 0$. The reduction factors P_Δ, compares each root z with the supplied challenge element $[x]_1$ by testing $[z]_1 = [x]_1$, and outputs the matching root. Therefore, for an efficient bilinear Q-DLOG adversary $\mathcal{B}_{\mathsf{bind}}$,

$$\mathrm{Adv}^{\mathsf{bind}}_{\mathsf{Carina}} \le \varepsilon^{\mathsf{qdlog}}_Q(\mathcal{B}_{\mathsf{bind}}) + \frac{2b - 2}{|\mathbb{F}| - 1}. \tag{11}$$

Write $\mathcal{R}_b = \mathbb{F}[X, Y]_{<b,<b}$. A malicious prover may use the entire tensor key for every group message. Its AGM explanations therefore reveal

$$F, A_0, A_1, U_\eta, V_\eta \in \mathcal{R}_b$$

for the corresponding group elements $C_f, C_0, C_1, \Pi_X, \Pi_Y$. The honest prover uses $A_0(X, Y) = S_0(X)$ and $A_1(X, Y) = S_1(Y)$, but the proof does not assume these forms. Transcript order ensures that F, A_0 are fixed before ρ, A_1 is fixed before ζ, and all scalar claims are fixed before η.

Let I interpolate the four claimed values of F on $A \times B$, where $A = \{\rho, \rho^{-1}\}$ and $B = \{\zeta, \zeta^{-1}\}$. Let $L_0(X)$ interpolate the claimed values of S_0 on A, and let

$L_1(Y)$ interpolate the claimed values of S_1 on B; each pair includes one value recovered by the verifier. Define

$$R_f = F - I, \qquad R_0 = A_0 - L_0, \qquad R_1 = A_1 - L_1.$$

The accepted pairing equation is equivalent to

$$D(\tau, \sigma) = 0, \quad D = R_f + \eta R_0 + \eta^2 R_1 - Z_A U_\eta - Z_B V_\eta.$$

Each residual lies in $\mathcal{R}_b$. Multiplication by the quadratic Z_A adds at most two to the X degree, and multiplication by Z_B adds at most two to the Y degree. Thus D has total degree at most $2b$.

Apply the same bilinear Q-DLOG simulation to the AGM+ROM knowledge-soundness adversary, including an honest simulation of its random oracle. The adversary sees a perfectly distributed Carina SRS, and its AGM representations determine D. Pairing acceptance and nondegeneracy give

$$P_D(x) = 0, \qquad P_D(Z) = D(r_X Z + s_X, r_Y Z + s_Y).$$

If $D \neq 0$, Lemma 1 says that P_D is nonzero except with probability $2b/(|\mathbb{F}| - 1)$. In that case, factoring P_D and comparing its roots with $[x]_1$ solves bilinear Q-DLOG. Hence, for an efficient Q-DLOG adversary $\mathcal{B}_{\mathsf{open}}$,

$$\Pr[\text{pairing accepts and } D \neq 0] \leq \varepsilon_Q^{\mathsf{qdlog}}(\mathcal{B}_{\mathsf{open}}) + \frac{2b}{|\mathbb{F}| - 1}. \tag{12}$$

Except for this event, acceptance therefore gives

$$R_f + \eta R_0 + \eta^2 R_1 = Z_A U_\eta + Z_B V_\eta.$$

Reduce the polynomial identity $R_f + \eta R_0 + \eta^2 R_1 = Z_A U_\eta + Z_B V_\eta$ modulo $J_{A,B} = \langle Z_A, Z_B \rangle$. The Chinese remainder theorem gives $\mathbb{F}[X, Y]/J_{A,B} \cong \mathbb{F}^{A \times B}$. If one of $[R_f], [R_0], [R_1]$ is nonzero, select a coordinate on which the vector polynomial

$$[R_f] + \eta[R_0] + \eta^2[R_1]$$

is nonzero. It has at most two roots in $\mathbb{F}^*$, while its coefficients are fixed before η. Hence, except with probability $2/(|\mathbb{F}| - 1)$, all three residuals belong to $J_{A,B}$. In particular, the four claimed values of F and the claimed values represented by A_0, A_1 are authentic on the four-point grid $A \times B$.

It remains to show that authenticated grid values imply the original multilinear evaluation. Let

$$h_+(Y) = F(\rho, Y), \qquad h_-(Y) = F(\rho^{-1}, Y), \qquad H_\lambda = h_+ + \lambda h_-,$$

and let a_+, a_- be the two transmitted values claimed for $g(\rho), g(\rho^{-1})$. Put $A_\lambda = a_+ + \lambda a_-$. After ρ and λ are fixed but before ζ is sampled, the polynomial $T_1(Y) = A_1(\rho, Y)$ is fixed. The inner Laurent residual is

$$\begin{aligned} D_1(Y) = H_\lambda(Y) T_v(Y^{-1}) + H_\lambda(Y^{-1}) T_v(Y) - 2A_\lambda \\ - Y T_1(Y) - Y^{-1} T_1(Y^{-1}). \end{aligned}$$

Its constant coefficient is $2(H_\lambda(v) - A_\lambda)$. If $H_\lambda(v) \neq A_\lambda$, $Y^b D_1(Y)$ is nonzero with degree at most $2b$. The authenticated grid relations and the instance of (2) for $(\phi, T_v, S_1, v, y_1, \zeta)$ imply $D_1(\zeta) = 0$ in every accepting transcript. This event has probability at most $2b/(|\mathbb{F}| - 3)$.

The random λ separates the claims $F(\rho, v) = a_+$ and $F(\rho^{-1}, v) = a_-$. If

$$e_+ = F(\rho, v) - a_+, \qquad e_- = F(\rho^{-1}, v) - a_-$$

are not both zero, then $e_+ + \lambda e_- = 0$ for at most one $\lambda \in \mathbb{F}^*$. This contributes $1/(|\mathbb{F}| - 1)$. Outside this event, $a_+ = g(\rho)$ and $a_- = g(\rho^{-1})$ for the polynomial g defined in (6).

Finally, A_0 is fixed before ρ. Define the bivariate Laurent defect

$$\begin{aligned} D_0(X, Y) = g(X)T_u(X^{-1}) + g(X^{-1})T_u(X) - 2y \\ - XA_0(X, Y) - X^{-1}A_0(X^{-1}, Y). \end{aligned}$$

If $g(u) \neq y$, its constant coefficient in X is nonzero. Therefore $P_0(X, Y) = X^b D_0(X, Y)$ is a nonzero polynomial with $\deg_X P_0 \leq 2b$ and $\deg_Y P_0 \leq b - 1$. Authenticated grid values and the recovered value $S_0(\rho^{-1})$ imply $P_0(\rho, \zeta) = 0$ in every accepting transcript. Since P_0 was fixed before ρ and (ρ, ζ) is uniform over $(\mathbb{F} \setminus \{0, 1, -1\})^2$, Schwartz–Zippel bounds this event by $(3b - 1)/(|\mathbb{F}| - 3)$.

A false opening must therefore cause at least one of: a bilinear Q-DLOG break, a degenerate affine specialization, a two-root η collision, a root of D_1, a one-root λ collision, or a root of P_0. Their interactive probability is at most

$$\varepsilon_Q^{\mathsf{qdlog}} + \frac{2b}{|\mathbb{F}| - 1} + \frac{5b - 1}{|\mathbb{F}| - 3} + \frac{3}{|\mathbb{F}| - 1}.$$

For Fiat–Shamir, define local bad-answer sets at each canonical transcript prefix. Their sizes are at most $2b$, 1, $3b - 1$, and 2 for $\rho, \lambda, \zeta, \eta$, respectively. Each set is fixed before the corresponding challenge and the domains are separated. A query-tree union bound gives (8). The extracted F defines the unique multilinear polynomial witnessing (1). If two distinct values at one point are accepted for the same commitment, either one knowledge-soundness failure occurs or the extracted polynomials violate commitment binding. A union bound gives the AGM+ROM evaluation-binding advantage stated in Theorem 2 and completes the proof.