Notes and Proofs for Divisor Techniques

Mathias Hall-Andersen Diego F. Aranha

April 27, 2026 · 681d28ac

ZKSECURITY

We revisit the protocol for proving Elliptic Curve Inner Products (ECIPs) by Eagen [Eag22]. The goal of this document is to supersede and summarize prior material and discussions by Bassa [Bas24b] [Bas24a] [Bas25] and Goodell et al. [GSSS25b] [GSSS25a]. In this document, we provide a self-contained explanation of the techniques, from the required algebraic geometry, to the interactive proof and its soundness, to the composition with a simulation extractable non-interactive proof and the efficiency of expressing the verifier as an R1CS circuit, resulting in the concrete gadget of Parker [Par24b]. Along the way, we clarify proof details, including simplifying the soundness argument for the interactive protocol, and correct other minor issues in the prior works. We aim for this document to be readable (and verifiable) by an audience with preexisting knowledge of elliptic curves, but not wider theoretical results in algebraic geometry or Galois theory. Additionally, we assume an understanding of the concepts of interactive proofs, but provide the required formal definitions along the way.

1 Background on Elliptic Curve Divisors

Throughout this document we adopt the notation of Bassa, distinct from the original paper of Eagen [Eag22]. We lift some basic definitions and background material from [Maa04].

Let q be a prime and let E be an elliptic curve in short Weierstrass form over $\mathbb{F}_q$, given by $y^2 = x^3 + Ax + B$. Let p be the order of a fixed prime-order subgroup $\mathbb{G} \subseteq E(\mathbb{F}_q)$. Define the function $r \in \mathbb{F}_q[x, y]$ by $r(x, y) = y^2 - (x^3 + Ax + B)$, the *coordinate ring* of the affine curve by $\mathbb{F}_q[E] = \mathbb{F}_q[x, y]/\langle r \rangle$, and the point at infinity by $\mathcal{O}$. Elements of $\mathbb{F}_q[E]$ can be represented in *canonical form* as $l(x, y) = a(x) - b(x) \cdot y$ with $a(x), b(x) \in \mathbb{F}_q[x]$ by using the curve equation to knock down exponents of y greater than or equal to 2.

By taking the field of fractions of $\mathbb{F}_q[E]$, we obtain the function field $\mathbb{F}_q(E)$, such that elements f_1/g_1 and f_2/g_2 represent the same element if $f_1 g_2 = f_2 g_1$. A non-zero rational function $f \in \mathbb{F}_q(E)^*$ is said to be defined at a point $P \in E \setminus \{\mathcal{O}\}$ if f can be written as $f = g/h$ for $g, h \in \mathbb{F}_q[E]$ with $h(P) \neq 0$. Additionally, f is said to have a *zero* at P if $f(P) = 0$; and f is said to have a *pole* at P (denoted by $f(P) = \infty$) if f is not defined at P. Equivalently, if $f = g/h$ is written in reduced form (so that g and h share no common factor), then f has a pole at P if and only if $h(P) = 0$.

To determine the value of a rational function $f = g/h \in \mathbb{F}_q(E)$ at $\mathcal{O}$, we compare the degrees of the denominator and numerator by weighting x and y by 2 and 3, respectively, such that substituting the curve equation leaves the degrees unchanged. For non-zero $f \in \mathbb{F}_q[E]$ in canonical form, we define *degree* as:

$$\deg_{\mathrm{E}}(f) \coloneqq \max\{2 \cdot \deg(a(x)),\ 3 + 2 \cdot \deg(b(x))\}$$

If $\deg(g) < \deg(h)$, then $f(\mathcal{O}) = 0$; and if $\deg(g) > \deg(h)$, then f has a pole at $\mathcal{O}$. If $\deg(g) = \deg(h)$, then $f(\mathcal{O}) = \mathrm{lc}(g)/\mathrm{lc}(h)$, where lc returns the coefficient of the leading term of g and h in canonical form, respectively.

For every point $P \in E$, there exists a rational function u with $u(P) = 0$ such that every non-zero rational function $f \in \mathbb{F}_q(E)^*$ can be written as $f = u^d s$ for $s \in \mathbb{F}_q(E)$, $s(P) \neq 0, \infty$, and integer d. The function u is said to be a *uniformizing parameter* for P, and the value of d does not depend on the choice of u. Then the *order* of f at P equals d, written as $v_P(f)$. If P is a zero of f, then $d > 0$ and the zero is said to have *multiplicity* d. If P is a pole, the pole is said to have multiplicity $-d$.

A rational function on an elliptic curve is determined by its finite zeros and poles, up to a constant multiple. A *divisor* is a formal sum $\sum n_P(P)$ that records the zeros and poles of a function, with coefficients $n_P = v_P(f) \in \mathbb{Z}$ (positive at zeros, negative at poles). The divisors form a group denoted $\mathrm{Div}(E)$, where addition is given by adding the coefficients for the same points. The *support* is given by the set $\mathrm{Supp}(D) = \{P \in E | n_P \neq 0\}$. The notation $(D)_0$ denotes the positive part of the divisor D, or $\sum_{n_P > 0} n_P(P)$.

A divisor $D \in \mathrm{Div}(E)$ is called *principal* if $D = \mathrm{div}(f)$ for some rational function f. Because of this one-to-one correspondence between principal divisors and rational functions, it is common to use D as both the divisor and the associated rational function. We use the following characterization:

Theorem 1 ([Sil86, Corollary 3.5])**.** *Let E be an elliptic curve and let $D = \sum n_P(P) \in \mathrm{Div}(E)$. Then D is a principal divisor if and only if both:*

$$\sum_{P \in E} n_P = 0 \quad \textit{and} \quad \sum_{P \in E} [n_P]P = \mathcal{O}$$

The first condition defines the *degree* of the divisor as 0, and the second establishes that the points sum to zero under the group law. A rational function f can be evaluated on a divisor $D = \sum n_P(P) \in \mathrm{Div}(E)$ that satisfies $\mathrm{Supp}(\mathrm{div}(f)) \cap \mathrm{Supp}(D) = \emptyset$ as:

$$f(D) = \prod_{P \in \mathrm{Supp}(D)} f(P)^{n_P}$$

As an illustrating example, a line $L(x, y) = y - \lambda x - \mu$ with slope λ is a rational function on E which meets the curve in 3 points. Therefore, it has 3 zeros and must have 3 poles, and the only pole is at $\mathcal{O}$. So the divisor is $\mathrm{div}(L) = (P_1) + (P_2) + (P_3) - 3(\mathcal{O})$. The coefficient of -3 at $\mathcal{O}$ comes from the fact that the line function, viewed as a rational function on the curve, has a pole of multiplicity 3 at $\mathcal{O}$. To see this concretely, define $u(x, y) = x/y$ as a uniformizing parameter for $\mathcal{O}$. One can express x and y in terms of u as $x = u^{-2}s_1$ and $y = u^{-3}s_2$, where s_1, s_2 are rational functions obtained by substituting u into the equation for E and solving iteratively. Note that the weighting above reflects the pole multiplicities of x and y at $\mathcal{O}$, and that in general $\deg_{\mathrm{E}}(D) = -v_{\mathcal{O}}(D)$, i.e. the degree equals the pole multiplicity at $\mathcal{O}$.

The soundness of the protocol [Eag22] will rely on sampling random points from $E(\mathbb{F}_q)$, so we need to know that the group is large. The following theorem makes this precise.

Theorem 2 (Hasse bound [Sil86, Theorem V.1.1])**.** *Let E be an elliptic curve over $\mathbb{F}_q$. Then:*

$$|\#E(\mathbb{F}_q) - (q + 1)| \leq 2\sqrt{q}$$

The following two lemmas from Bassa bound the probability of "bad" evaluation points when sampling A_0, A_1 uniformly from $E(\mathbb{F}_q)$.

Lemma 1 (Slope Distribution [Bas24b, Lemma 4])**.** *Let $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ be chosen uniformly at random. Let L be the affine line passing through A_0, A_1, with the convention that L is the tangent line at A_0 when $A_0 = A_1$. For a given $\lambda \in \mathbb{F}_q$:*

$$\Pr[\mathrm{Slope}(L) = \lambda] \leq \frac{2}{\#E(\mathbb{F}_q) - 1}$$

Proof. Consider the q affine lines $L_\mu(X, Y) = Y - \lambda X - \mu$ of slope λ. Classify them by how they meet E:

(i) N_0 lines with at most one affine rational intersection point.

(ii) N_1 lines tangent to E at a rational 3-torsion point.

(iii) N_2 lines meeting E in two affine rational points and tangent at one of them.

(iv) N_3 lines meeting E in three distinct affine rational points.

For a fixed λ, every affine rational point lies on a unique line L_μ, and a line of type (ii), (iii), or (iv) contains exactly 1, 2, or 3 such points, respectively. Observe that:

$$N_1 + 2N_2 + 3N_3 \leq \#E(\mathbb{F}_q) - 1$$

Note that N_0 is omitted, which can only make the left quantity smaller; and any sampled line will have at least two intersection points with E, A_0 and A_1, so we don't care about bounding it. Now, for a fixed $\mu \in \mathbb{F}_q$, the number of ordered pairs $(A_0, A_1) \in (E(\mathbb{F}_q) \setminus \{\mathcal{O}\})^2$ that result in L_μ as the line passing through them is: 0 in case (i), 1 in case (ii), 3 in case (iii), and 6 in case (iv): in case (ii) only (T, T) works, where T is the rational 3-torsion point; in case (iii), if T is the tangency point and P the other rational point, the possibilities are (T, T), (T, P), and (P, T); and in case (iv) the possibilities are the six ordered pairs of distinct rational points on L_μ. Therefore:

$$\Pr[\text{Slope}(L) = \lambda] = \frac{N_1 + 3N_2 + 6N_3}{(\#E(\mathbb{F}_q) - 1)^2}$$

And since:

$$N_1 + 3N_2 + 6N_3 \leq 2 \cdot (N_1 + 2N_2 + 3N_3) \leq 2 \cdot (\#E(\mathbb{F}_q) - 1)$$

We conclude:

$$\Pr[\text{Slope}(L) = \lambda] \leq \frac{2}{\#E(\mathbb{F}_q) - 1}$$ □

> ***Comments On Prior Work.***
>
> The original lemma in [Bas24b] states the bound as $2/(\#E(\mathbb{F}_q) - 2)$, but the proof yields the *marginally* tighter $2/(\#E(\mathbb{F}_q) - 1)$.

Whereas Lemma 1 controls the slope of the line passing through A_0, A_1, the next lemma controls the points themselves: when D has few zeros, a random challenge triple $\{A_0, A_1, A_2\}$ avoids the support of D with high probability. This is what allows the verifier's evaluations $D(A_i)$ to be well-defined and nonzero in an honest evaluation of the protocol.

Lemma 2 (Support Disjointness [Bas24b, Lemma 5])**.** *Let $D \in \mathbb{F}_q[E]$ with $-v_\mathcal{O}(D) = N$ and let $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ be chosen uniformly at random, $A_2 = -(A_0 + A_1)$. Assuming $3 \cdot (N + 1) \leq 2 \cdot \#E(\mathbb{F}_q)$, the probability that the support of (D) is not disjoint from $\{A_0, A_1, A_2\}$ is at most $3(N + 1)/\#E(\mathbb{F}_q)$.*

Proof. D has N zeros and one pole at $\mathcal{O}$, so $|\text{Supp}(D)| = N + 1$. Since $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$, neither can hit the pole; $A_2 = \mathcal{O}$ requires $A_1 = -A_0$, which occurs with probability $1/(\#E(\mathbb{F}_q) - 1)$. For zeros: A_0 and A_1 are each uniform on $\#E(\mathbb{F}_q) - 1$ affine points, hitting one of the (at most) N rational zeros with probability at most $N/(\#E(\mathbb{F}_q) - 1)$. For fixed A_0, the point $A_2 = -(A_0 + A_1)$ is uniform on $E(\mathbb{F}_q) \setminus \{-A_0\}$, so it hits a zero with probability at most $N/(\#E(\mathbb{F}_q) - 1)$. By the union bound: $3N/(\#E(\mathbb{F}_q) - 1) + 1/(\#E(\mathbb{F}_q) - 1) = (3N + 1)/(\#E(\mathbb{F}_q) - 1) \leq 3(N + 1)/\#E(\mathbb{F}_q)$, which holds when $3(N + 1) \leq 2 \cdot \#E(\mathbb{F}_q)$. □

Comments On Prior Work.

The original lemma in [Bas24b] misses the condition $3(N+1) \leq 2 \cdot \#E(\mathbb{F}_q)$ in the statement.

With the probability of a bad support collision under control, we turn to a related failure mode: even when the support of D is disjoint from $\{A_0, A_1, A_2\}$, the norm $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)$ may still vanish at an auxiliary evaluation point $L(P)$ simply because the random line L happens to identify P with one of the zeros Q_i of D. The next theorem bounds exactly this event, and its proof is a direct application of Lem. 1.

Let's first recall the norm definition for our function fields of interest. Let $L(x, y) = y - \lambda x - \mu \in \mathbb{F}_q(E)$ be a non-vertical line, and let $\mathbb{F}_q(L) \subseteq \mathbb{F}_q(E)$ be the subfield of the function field generated by L. The element $L \in \mathbb{F}_q(E)$ is clearly non-constant and thus transcendental over $\mathbb{F}_q$. Hence $\mathbb{F}_q(L) \cong \mathbb{F}_q(t)$, identifying t with L, and it is a rational function field in one variable. Furthermore, L defines a map $L : E \to \mathbb{P}^1$, mapping (x, y) to $y - \lambda x - \mu$ and $\mathcal{O} \in E$ to $\mathcal{O} \in \mathbb{P}^1$. Observe that $L^{-1}(0)$ consists of the three intersection points of the line L with E. In general, the cardinality of the fibers over the algebraic closure, and hence the degree of the map, is $[\mathbb{F}_q(E) : \mathbb{F}_q(L)] = 3$. Additionally, the extension is separable: observe that $\mathbb{F}_q(E)$ is generated by x over $\mathbb{F}_q(L)$, since we can recover $y = \lambda x + \mu + t$, and the minimal polynomial of x has distinct roots. To see this, substitute $y = \lambda x + \mu + t$ into $y^2 = x^3 + Ax + B$ and rearrange as a polynomial in x:

$$\begin{aligned} m(X) &= X^3 + AX + B - (\lambda X + \mu + t)^2 \\ &= X^3 - \lambda^2 X^2 + (A - 2\lambda(\mu + t))\, X + (B - (\mu + t)^2) \in \mathbb{F}_q(t)[X] \end{aligned}$$

Its roots are the x-coordinates of the three intersection points between E and L. Viewed as a polynomial in t, its discriminant has leading term $-27\, t^4$ (from the $-27(B-(\mu+t)^2)^2$ summand; every other summand has t-degree at most 3), so the discriminant is non-zero in $\mathbb{F}_q(t)$ provided $\mathrm{char}(\mathbb{F}_q) \neq 3$. Observe that this is the place where the non-verticality of L is applied: a vertical line L intersects E *twice* in the *same* x-coordinate, so the extension would not be separable.

Since $m(X)$ is separable of degree 3, it has three distinct roots x_0, x_1, x_2 in the algebraic closure $\overline{\mathbb{F}_q(L)}$: the x-coordinates of the three intersection points of L with E. Each root determines an $\mathbb{F}_q(L)$-embedding $\sigma_i : \mathbb{F}_q(E) \hookrightarrow \overline{\mathbb{F}_q(L)}$ by $x \mapsto x_i$, and these are all such embeddings. It is not generally true that the other intersection points x_1, x_2 lie in $\mathbb{F}_q(E)$ once x_0 does, so the extension $\mathbb{F}_q(E)/\mathbb{F}_q(L)$ is not generally *normal*, i.e. need not be Galois.

Example 1. A concrete example is $E : y^2 = x^3 - x$ with $L = y$, where $m(X) = X^3 - X - y^2$ has discriminant $4 - 3x^2$, which is not a square in $\mathbb{F}_q(E)$ (in characteristic $\neq 2, 3$).

For $f \in \mathbb{F}_q(E)$, the *field norm* from $\mathbb{F}_q(E)$ to $\mathbb{F}_q(L)$ is:

$$\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(f) \coloneqq \prod_{i=0}^{2} \sigma_i(f)$$

A priori this lies in $\overline{\mathbb{F}_q(L)}$, but the absolute Galois group of $\mathbb{F}_q(L)$ permutes the roots x_i, and hence the embeddings σ_i, so it fixes the product: $\mathrm{N}(f) \in \mathbb{F}_q(L)$. The norm is multiplicative and vanishes only at $f = 0$, so it restricts to a homomorphism $\mathbb{F}_q(E)^* \to \mathbb{F}_q(L)^*$.

For $f \in \mathbb{F}_q[E]$ with poles only at $\mathcal{O}$ and $(f)_0 = \sum_k \beta_k \cdot (Q_k)$, the norm lies in the polynomial ring $\mathbb{F}_q[t] \subseteq \mathbb{F}_q(t) = \mathbb{F}_q(L)$ and has the explicit form:

$$\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(f) = \mathrm{lc}(f)^3 \cdot \prod_k (t - L(Q_k))^{\beta_k}$$

Where $L(Q_k) = y(Q_k) - \lambda\, x(Q_k) - \mu \in \mathbb{F}_q$ is the value of the line function at Q_k. When the line L is clear from context we abbreviate $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(f)$ by $\mathrm{N}(f)$.

Theorem 3 ([Bas24b, Theorem 6]). *Suppose D vanishes at $Q_1, \ldots, Q_N$ and suppose $P \neq Q_i$ for $i = 1, \ldots, N$. Choose random $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ with $A_0 \neq -A_1$ and consider the line L passing through them. Then the probability that $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)(L(P)) = 0$ is at most $2N/(\#E(\mathbb{F}_q) - 1)$.*

Proof. Since $A_0, A_1 \neq \mathcal{O}$ and $A_0 \neq -A_1$, the line L is not vertical, so $L = y - \lambda x - \mu$ for some $\lambda, \mu \in \mathbb{F}_q$. The function $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)$ has zeros exactly at $L(Q_i) = y(Q_i) - \lambda \cdot x(Q_i) - \mu$ for $i = 1, \ldots, N$. Hence $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)(L(P)) = 0$ iff $L(P) = L(Q_i)$ for some i, equivalently:

$$y(P) - \lambda \cdot x(P) = y(Q_i) - \lambda \cdot x(Q_i)$$

If $x(Q_i) = x(P)$ then $P \neq Q_i$ forces $y(Q_i) \neq y(P)$, so no λ satisfies the equation. Otherwise we need $\lambda = \frac{y(P)-y(Q_i)}{x(P)-x(Q_i)}$. There are at most N such values of λ, each occurring as the slope of L with probability at most $2/(\#E(\mathbb{F}_q) - 1)$ by Lem. 1. By a union bound the result follows. □

We now assemble the main soundness machinery for the norm-based check. Rather than reasoning directly about the elliptic-curve function D, the strategy is to package the discrepancy between the zeros of D and a claimed multiset $(P_1), \ldots, (P_N)$ into a single polynomial f in four ring variables (X_0, Y_0, X_1, Y_1), and then bound its rational zero set on $E \times E$ via a Schwartz–Zippel-type estimate. The next three results carry this out in stages: first we check that f really is a polynomial with coefficients in $\mathbb{F}_q$ (Lem. 3), then that it does not vanish identically on $E \times E$ (Lem. 4), and finally we invoke a variety-bound (Thm. 4) to count its $\mathbb{F}_q$-rational zeros.

Lemma 3 ([Bas24b, Lemma 7]). *Let $D \in \mathbb{F}_q[E]$ with $(D)_0 = \sum_{i=1}^{N}(Q_i)$ and let $P_1, \ldots, P_N \in E(\mathbb{F}_q)$. Define:*

$$\begin{aligned} f(X_0, Y_0, X_1, Y_1) = &\prod_{i=1}^{N} ((y(Q_i) - Y_0) \cdot (X_1 - X_0) - (x(Q_i) - X_0) \cdot (Y_1 - Y_0)) \\ &- \prod_{i=1}^{N} ((y(P_i) - Y_0) \cdot (X_1 - X_0) - (x(P_i) - X_0) \cdot (Y_1 - Y_0)) \end{aligned}$$

As a function on $E \times E$. Then $f(X_0, Y_0, X_1, Y_1) \in \mathbb{F}_q[X_0, Y_0, X_1, Y_1]$.

Proof. Since $P_i \in E(\mathbb{F}_q)$, we have $x(P_i), y(P_i) \in \mathbb{F}_q$, so the second product is in $\mathbb{F}_q[X_0, Y_0, X_1, Y_1]$. The Q_i are zeros of $D \in \mathbb{F}_q[E]$, so they form full Galois orbits over $\mathbb{F}_q$. Hence for each factor $(y(Q_i)-Y_0)\cdot(X_1-X_0)-(x(Q_i)-X_0)\cdot(Y_1-Y_0)$ in the first product, all of its Galois conjugates also appear as factors. The first product is therefore Galois-invariant and lies in $\mathbb{F}_q[X_0, Y_0, X_1, Y_1]$. □

Rationality alone is not enough: to eventually apply a variety bound we also need f to be genuinely nonzero as a function on $E \times E$, for otherwise every pair (A_0, A_1) would satisfy $f(A_0, A_1) = 0$ and the soundness statement would be vacuous. The following lemma shows that non-vanishing is guaranteed as soon as the two zero multisets, the divisors $\sum(Q_i)$ and $\sum(P_i)$, disagree.

Lemma 4 ([Bas24b, Lemma 8]). *Assume $\sum(Q_i) \neq \sum(P_i)$ (as divisors) and $3N \leq q - 2\sqrt{q}$. Then $f(X_0, Y_0, X_1, Y_1)$ does not vanish on $E \times E$.*

Proof. As $E \times E$ is irreducible, it suffices to show that f is not identically zero on $E \times E$. Since $\sum(Q_i) \neq \sum(P_i)$, by reordering we may assume $v_{P_1}(\sum(Q_i)) \neq v_{P_1}(\sum(P_i))$ and, without loss of generality, that the coefficient of P_1 in $\sum(P_i)$ exceeds that in $\sum(Q_i)$ (exchanging the roles of $\{P_i\}$ and $\{Q_i\}$ if necessary). Cancelling suitable powers of the quadratic factor $(y(P_1) - Y_0)(X_1 - X_0) - (x(P_1) - X_0)(Y_1 - Y_0)$, which does not vanish on $E \times E$, we may assume $P_1 \notin \mathrm{Supp}\,(\sum(Q_i))$. Substituting $X_0 = x(P_1)$, $Y_0 = y(P_1)$ makes the $i = 1$ factor of the second product zero, so:

$$f(x(P_1), y(P_1), X_1, Y_1) = \prod_{i=1}^{N} ((y(Q_i) - y(P_1))(X_1 - x(P_1)) - (x(Q_i) - x(P_1))(Y_1 - y(P_1)))$$

Since $P_1 \notin \mathrm{Supp}(\sum(Q_i))$, each factor is a nonzero linear form in (X_1, Y_1). Its zero set on E has at most 3 points, so the product vanishes on at most $3N$ points of $E(\mathbb{F}_q)$. By Thm. 2, $\#E(\mathbb{F}_q) \geq q + 1 - 2\sqrt{q}$, so the hypothesis $3N \leq q - 2\sqrt{q}$ gives $3N \leq \#E(\mathbb{F}_q) - 1 < \#E(\mathbb{F}_q)$, and f does not vanish on $E(\mathbb{F}_q) \times \{(x(P_1), y(P_1))\}$. □

The last ingredient is a counting tool that translates "f does not vanish identically" into a concrete bound on the fraction of pairs $(A_0, A_1) \in E(\mathbb{F}_q) \times E(\mathbb{F}_q)$ that do satisfy $f(A_0, A_1) = 0$. We use the affine version of the generalized Schwartz–Zippel lemma.

Theorem 4 (Generalized Schwartz–Zippel [DKL14, Claim 7.2])**.** *Let V be an affine variety of dimension n and degree d. Then:*

$$\#V(\mathbb{F}_q) \leq d \cdot q^n$$

> ***Comments On Prior Work.***
>
> [Bas24b, Theorem 9] states this bound for *projective* varieties, citing both [DKL14, Claim 7.2] and [EOT10, Lemma A.3]. The former gives $d \cdot q^n$ only for *affine* varieties; the latter gives $d \cdot (q+1)^n$ for projective varieties. The bound $d \cdot q^n$ fails projectively: a line in $\mathbb{P}^2$ has $q + 1 > q$ rational points. The affine version is what the proofs require.

Combining the three preceding results yields the first concrete soundness bound for the norm-based check. Intuitively, f captures the discrepancy between the zero divisor of D and the claimed multiset $\sum(P_i)$; Lem. 3 places f in $\mathbb{F}_q[X_0, Y_0, X_1, Y_1]$, Lem. 4 ensures it is not identically zero on $E \times E$ whenever the claim is wrong, and Thm. 4 bounds the number of $\mathbb{F}_q$-rational pairs (A_0, A_1) on which f can vanish. The following theorem is the direct consequence for a monic D.

Theorem 5 ([Bas24b, Theorem 10])**.** *Given $D \in \mathbb{F}_q[E]$ with leading coefficient 1 and $-v_{\mathcal{O}}(D) = N$ and $P_1, \ldots, P_N \in E(\mathbb{F}_q)$. Suppose that $(D)_0 \neq (P_1) + \cdots + (P_N)$. Choose random $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$, with $A_0 \neq \pm A_1$. Let $L = y - \lambda \cdot x - \mu$ be the line passing through A_0, A_1. Then:*

$$\begin{aligned}\Pr\left[\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)(0) = \prod_{i=1}^{N}(-L(P_i))\right] &\leq \frac{18Nq}{(\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3)} \\ &\leq \frac{18Nq}{(q - 2\sqrt{q})(q - 2\sqrt{q} - 2)} \leq \frac{96N}{q}\end{aligned}$$

Where the middle bound follows from the Hasse bound $\#E(\mathbb{F}_q) \geq q + 1 - 2\sqrt{q}$ (Thm. 2), and the rightmost bound holds for all $q \geq 16$ (any cryptographic instantiation).

Proof. For a random choice of A_0, A_1, the norm equality holds exactly when $f(A_0, A_1) = 0$, where f is the function from Lem. 3. Writing $(D)_0 = \sum_{i=1}^{N}(Q_i)$ for the zero divisor of D, the hypothesis $(D)_0 \neq (P_1) + \cdots + (P_N)$ is exactly $\sum(Q_i) \neq \sum(P_i)$ as divisors, so Lem. 4 applies and f does not vanish identically on $E \times E$. The surface $E \times E$ is irreducible of degree $3 \cdot 3 = 9$ and f has bi-degree (N, N); their intersection is a curve of degree $2N \cdot 9 = 18N$. By Thm. 4, the number of $\mathbb{F}_q$-rational zeros of f on $E \times E$ is at most $18Nq$. The number of pairs $(A_0, A_1) \in (E(\mathbb{F}_q) \setminus \{\mathcal{O}\})^2$ with $A_0 \neq \pm A_1$ is $(\#E(\mathbb{F}_q) - 1)^2 - 2(\#E(\mathbb{F}_q) - 1) = (\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3)$, yielding the probability bound $18Nq/((\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3))$. By Thm. 2, $\#E(\mathbb{F}_q) \geq q + 1 - 2\sqrt{q}$, so $(\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3) \geq (q - 2\sqrt{q})(q - 2\sqrt{q} - 2)$ and the probability is at most $18Nq/((q - 2\sqrt{q})(q - 2\sqrt{q} - 2))$. For $q \geq 16$ this can be bounded above by $96N/q$. □

> ***Comments On Prior Work.***
>
> Theorem 10 of [Bas24b] states the bound as $18Nq/((\#E(\mathbb{F}_q) - 1)^2 - 2(\#E(\mathbb{F}_q) - 1)) \approx 18N/q$, and Theorem 11 merely as "negligible" under $q \gg N$. We chain Hasse (Thm. 2) through to the explicit $96N/q$ for $q \geq 16$, likewise for Thm. 6 below. Because the techniques in these works deal with

exponential quantities, we do this to avoid, for instance, accidentally doing a union bound over an exponential number of negligible terms and having it result in something non-negligible. Furthermore, the *concrete bounds* throughout help compute parameters for *concrete security* levels, e.g. the function t^4/q where t is the running time of the adversary, is negligible; but would substantially affect the concrete q which achieve 128-bits of security. All the components/analysis in this paper have a loss *linear in the adversary's running time* and an elliptic curve of ≈ 256-bits achieves 128-bits of security. Observe that only the composition in Section 3 actually requires discrete log to be intractable on $E(\mathbb{F}_q)$.

The following theorem due to Bassa bounds the probability that the norm-based check passes for a "wrong" witness function D. It generalizes Thm. 5 along two axes: D need not be monic (its leading coefficient may be any element of $\mathbb{F}_q$, and the monic case is recovered when $\mathrm{lc}(D) = 1$), and the claimed multiset $(P_1), \ldots, (P_{N_1})$ may have a different cardinality N_1 from the pole order $N_2 = v_{\mathcal{O}}(D)$. The three cases (i)–(iii) below enumerate the possible ways the check can still accept a wrong witness, and each is reduced to an instance of the Lem. 3–Thm. 4 machinery developed above.

Theorem 6 (Schwartz–Zippel on $E \times E$ [Bas24b, Theorem 11])**.** *Let $D \in \mathbb{F}_q[E]$ and let $P_1, \ldots, P_{N_1} \in E(\mathbb{F}_q)$. Choose random $A_0, A_1 \in E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ with $A_0 \neq \pm A_1$. Let $L = y - \lambda \cdot x - \mu$ be the line passing through A_0, A_1. Write $\Delta \coloneqq (\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3)$ for the number of valid pairs (A_0, A_1); by Thm. 2, $\Delta \geq (q - 2\sqrt{q})(q - 2\sqrt{q} - 2)$. Then, over the random choice of A_0, A_1, the probability that:*

$$\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)(0) = \prod_{i=1}^{N_1}(-L(P_i))$$

The probability is bounded in each of the following cases.

- *(i)* $D = 0$*: probability* $\leq 18N_1q/\Delta \leq 96N_1/q$ *for* $q \geq 16$.
- *(ii)* $D \neq 0$, $(D)_0 = \sum_{i=1}^{N_2}(Q_i)$, $\sum P_i \neq \sum Q_i$*: probability* $\leq 18Nq/\Delta \leq 96N/q$ *for* $q \geq 16$*, where* $N = \max\{N_1, N_2\}$.
- *(iii)* $D \neq 0$, $(D)_0 = \sum_{i=1}^{N_2}(Q_i)$, $\sum P_i = \sum Q_i$*, but* $\mathrm{lc}(D)^3 \neq 1$ *(where* $\mathrm{lc}(D)$ *denotes the leading coefficient of* D*): probability* $\leq 18N_2q/\Delta \leq 96N_2/q$ *for* $q \geq 16$.

The rightmost simplifications hold for all $q \geq 16$, which is satisfied in any cryptographic instantiation.

Proof. We follow Bassa's argument [Bas24b, Theorem 11] in each case, adapted to the labels of the present document.

Case (i). If $D = 0$, then:

$$D(A_0) \cdot D(A_1) \cdot D(A_2) = \mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)\big|_{L=0} = 0$$

The LHS of the claimed norm equality vanishes automatically; it remains to bound the probability that the RHS vanishes as well. Consider the following function on $E \times E$:

$$g(X_0, Y_0, X_1, Y_1) = \prod_{i=1}^{N_1}((y(P_i) - Y_0)(X_1 - X_0) - (x(P_i) - X_0)(Y_1 - Y_0))$$

For any $(A, B) \in E(\mathbb{F}_q) \setminus \{P_1, \ldots, P_{N_1}\}$, the equation $g(A, B, X_1, Y_1) = 0$ defines a union of N_1 lines, each intersecting $\{(A, B)\} \times E$ in at most 3 points, so g does not vanish identically on $E \times E$. As in the proof of Thm. 5, the intersection of $g = 0$ with the surface $E \times E$ is a curve of degree at most $18N_1$, so by Thm. 4 the number of $\mathbb{F}_q$-rational zeros of g on $E \times E$ is at most $18N_1q$. The number of valid pairs (A_0, A_1) is $\Delta = (\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3)$, so the probability is bounded by $18N_1q/\Delta$. By Thm. 2, $\Delta \geq (q - 2\sqrt{q})(q - 2\sqrt{q} - 2)$, which for $q \geq 16$ gives $18N_1q/\Delta \leq 96N_1/q$.

Case (ii). If $\sum P_i \neq \sum Q_i$, then:

$$D(A_0) \cdot D(A_1) \cdot D(A_2) = \mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)\big|_{L=0} = c \cdot \prod_{i=1}^{N_2} \big(-(y(Q_i) - \lambda \cdot x(Q_i) - \mu)\big)$$

We need to show that with high probability:

$$c \cdot \prod_{i=1}^{N_2} \big(-(y(Q_i) - \lambda \cdot x(Q_i) - \mu)\big) \neq \prod_{i=1}^{N_1} \big(-(y(P_i) - \lambda \cdot x(P_i) - \mu)\big)$$

Let $N = \max\{N_1, N_2\}$. Expressing λ, μ in terms of A_0, A_1 and clearing denominators, this reduces to showing that:

$$(x_1 - x_0)^{\epsilon_1} \cdot c \cdot \prod_{i=1}^{N_2} \big((y(Q_i) - y_0)(x_1 - x_0) - (x(Q_i) - x_0)(y_1 - y_0)\big)$$
$$- (x_1 - x_0)^{\epsilon_2} \cdot \prod_{i=1}^{N_1} \big((y(P_i) - y_0)(x_1 - x_0) - (x(P_i) - x_0)(y_1 - y_0)\big) \neq 0$$

Where $\epsilon_1 = N - N_2$ if $N_1 > N_2$ and 0 otherwise, and symmetrically $\epsilon_2 = N - N_1$ if $N_2 > N_1$ and 0 otherwise. Invoking Lem. 4 (non-vanishing of f on $E \times E$ under the multiset condition) and the counting argument of Thm. 5, the probability is bounded by $18Nq/\Delta$, which is at most $96N/q$ for $q \geq 16$.

Case (iii). If $(D)_0 = \sum_{i=1}^{N_2}(Q_i)$ agrees with $\sum P_i$ as multisets but $\mathrm{lc}(D)^3 \neq 1$, then the norm computation yields a nontrivial constant $c = \mathrm{lc}(D)^3 \neq 1$, and we need:

$$c \cdot \prod_{i=1}^{N_2} \big(-(y(Q_i) - \lambda \cdot x(Q_i) - \mu)\big) \neq \prod_{i=1}^{N_2} \big(-(y(Q_i) - \lambda \cdot x(Q_i) - \mu)\big)$$

Rewriting:

$$(c - 1) \cdot \prod_{i=1}^{N_2} \big(-(y(Q_i) - \lambda \cdot x(Q_i) - \mu)\big) \neq 0$$

This reduces to case (i) applied to the product on the right (using the points Q_i in place of P_i). Hence the probability is bounded by $18N_2q/\Delta$, which is at most $96N_2/q$ for $q \geq 16$. □

The results up to this point analyze the *norm*-based check, which tests an equality between two products of linear evaluations. We now turn to a second, complementary check based on the *logarithmic derivative*, introduced in [Bas24a, Bas25]. Its soundness analysis rests on two function-field facts: first, that a rational function of low degree is determined by its logarithmic derivative up to a constant (Lem. 5), and second, that the logarithmic derivative of the norm $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)$ admits a clean closed-form expression in terms of the three line points (Lem. 6). Together these yield a Schwartz–Zippel bound for the log-derivative check (Lem. 7) that parallels Thm. 6.

Lemma 5 ([Bas24a, Lemma 1])**.** *Let F/K be a function field with K perfect and the full constant field of F. Let δ be a derivation of F/K and let $\mathcal{L}\colon F^\times \to F$ denote the logarithmic derivative $f \mapsto \delta(f)/f$. Suppose $t \in \ker(\mathcal{L})$ and $\deg_F t < p$, where p is the characteristic. Then $t \in K$, i.e. the only elements in the kernel of the logarithmic derivative of low degree are constants. In particular for nonzero f, g of degree $< p$ we have $\mathcal{L}(f) = \mathcal{L}(g)$ if and only if $f = c \cdot g$ for some constant $c \in K$.*

The previous lemma tells us *what* a logarithmic-derivative equality implies; the next lemma gives us an explicit formula for the quantity that actually appears in the protocol, namely the logarithmic derivative of a norm $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)$ evaluated at the zero of L. The resulting expression is a sum indexed by the three points A_0, A_1, A_2 cut out by the line, which is what makes it amenable to a Schwartz–Zippel argument in the subsequent lemma.

Lemma 6 ([Bas25, Lemma 5]; proved in [Bas24a, Section 4])**.** *Let $D(x,y) = a(x) - y \cdot b(x) \in \mathbb{F}_q[E] \setminus \{0\}$ be a nonzero rational function on E having only poles at $\mathcal{O}$, let $L = y - \lambda \cdot x - \mu$ and consider the subfield $\mathbb{F}_q(L)$ of the function field $\mathbb{F}_q(E)$. Denote by N the field norm from $\mathbb{F}_q(E)$ to $\mathbb{F}_q(L)$. Let $\mathcal{L}$ denote the logarithmic derivative in $\mathbb{F}_q(L)$ with respect to L, i.e. $\mathcal{L}\colon \mathbb{F}_q(L)^\times \to \mathbb{F}_q(L)$ is given by $f \mapsto \delta(f)/f$, where δ denotes the derivation with respect to L. Let $(L = 0)$ be the zero of L in $\mathbb{F}_q(L)$. Then, writing $x_i = x(A_i)$ and $y_i = y(A_i)$:*

$$\mathcal{L}\Big(\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)\Big)|_{(L=0)} = \sum_{i=0}^{2} \frac{a'(x_i) - \frac{3x_i^2+A}{2y_i} b(x_i) - y_i\, b'(x_i)}{a(x_i) - y_i\, b(x_i)} \cdot \frac{2y_i}{3x_i^2 + A - 2\lambda\, y_i}$$

The log-derivative check in the protocol tests an identity of the form:

$$\sum_{i=0}^{2} \frac{D'(A_i)}{D(A_i)} \cdot \frac{dx(A_i)}{dz} = \sum_{j=1}^{M} \frac{m_j}{L(R_j)}$$

By Lem. 6, this is exactly the evaluation of the logarithmic derivative of $\mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)$ at $(L = 0)$, equated against a rational function whose poles encode the claimed coefficients m_j at the points R_j. The next lemma bounds the soundness error of this check by a Schwartz–Zippel argument that closely mirrors the one used for Thm. 5, but with the polynomial G now derived from the log-derivative identity.

> ***Comments On Prior Work.***
>
> Bassa [Bas25] does not bound the soundness error of the log-derivative check by a Schwartz–Zippel argument; the argument instead proceeds via a k-special-soundness extractor that recovers a witness from any set of at most $13kq$ accepting transcripts. For reasons explained later, the lemma below replaces that combinatorial route with a direct Schwartz–Zippel bound, using Bassa's own ingredients: the closed form for the logarithmic derivative of the norm (Lem. 6, from [Bas24a]) and the Schwartz–Zippel machinery on $E \times E$ developed in [Bas24b] for the norm check (Thm. 6). Bassa [Bas24a] explicitly suggests this adaptation, pointing to the surface $E \times E$ and the Dvir–Kollár–Lovett variety bounds [DKL14] as the route to a Schwartz–Zippel analogue, but does not formally carry out this argument. We need it, so we do that here.

Lemma 7 (Schwartz–Zippel for Log-Derivative of Norm Check)**.** *Let $D \in \mathbb{F}_q[E] \setminus \{0\}$, let $R_1, \ldots, R_M \in E(\mathbb{F}_q)$, and let $m_1, \ldots, m_M \in \mathbb{F}_q$. For $(Q_0, Q_1) \in E \times E$ with $Q_2 = -(Q_0 + Q_1)$ and L_Q the line through Q_0, Q_1, Q_2, define:*

$$f(Q_0, Q_1) := \sum_{i=0}^{2} \frac{D'(Q_i)}{D(Q_i)} \cdot \frac{dx(Q_i)}{dz} - \sum_{j=1}^{M} \frac{m_j}{L_Q(R_j)}$$

If f is not identically zero on $E \times E$, then for $A_0, A_1 \leftarrow\!\$\; E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$:

$$\Pr[f(A_0, A_1) = 0] \le \frac{18\big(\deg_{\mathrm{E}}(D) + M - 1\big) \cdot q}{\big(\#E(\mathbb{F}_q) - 1\big)\big(\#E(\mathbb{F}_q) - 3\big)}$$

Proof. Write $(D)_0 = \sum_k \beta_k \cdot (Q_k)$ with Q_k distinct, $\deg_{\mathrm{E}}(D) = \sum \beta_k$, and set $d = \#\{Q_k\}$.

Set $z \coloneqq y - \lambda x$ (so that $L = z - \mu$ and the place $(L = 0)$ corresponds to $z = \mu$). By Lem. 6:

$$f(Q_0, Q_1) = \mathcal{L}\big(\mathrm{N}(D)\big)\big|_{(L_Q=0)} - \sum_{j=1}^{M} \frac{m_j}{L_Q(R_j)}$$

Since $D \in \mathbb{F}_q[E]$ has only poles at $\mathcal{O}$, the norm $\mathrm{N}(D) = \mathrm{N}_{\mathbb{F}_q(E)/\mathbb{F}_q(z)}(D)$ lies in $\mathbb{F}_q[z]$ and inherits its zeros from those of D: concretely $\mathrm{N}(D) = \mathrm{lc}(D)^3 \prod_k (z - z(Q_k))^{\beta_k}$, where $z(Q_k) = y(Q_k) - \lambda\, x(Q_k)$. Taking the logarithmic derivative (the constant $\mathrm{lc}(D)^3$ drops out):

$$\mathcal{L}\big(\mathrm{N}(D)\big) = \sum_k \frac{\beta_k}{z - z(Q_k)}$$

Evaluating at $z = \mu$ and noting $\mu - z(P) = -L_Q(P)$:

$$f = 0 \quad \Longleftrightarrow \quad \sum_k \frac{\beta_k}{L_Q(Q_k)} + \sum_{j=1}^{M} \frac{m_j}{L_Q(R_j)} = 0$$

For any $P \in E$, define the bilinear form $\ell_P \coloneqq (y(P) - Y_0)(X_1 - X_0) - (x(P) - X_0)(Y_1 - Y_0)$, a polynomial of degree 2 in (X_0, Y_0, X_1, Y_1) satisfying $\ell_P = L_Q(P) \cdot (X_1 - X_0)$. On the valid challenge space (where $X_1 \neq X_0$ and all ℓ_{Q_k}, ℓ_{R_j} are nonzero), clearing denominators gives $f = H \cdot G$, where G is:

$$G \coloneqq \sum_k \beta_k \prod_{k' \neq k} \ell_{Q_{k'}} \cdot \prod_j \ell_{R_j} + \sum_j m_j \prod_k \ell_{Q_k} \cdot \prod_{j' \neq j} \ell_{R_{j'}}$$

And H is a nonzero polynomial (a product of the ℓ_{Q_k}, ℓ_{R_j}, and a power of $(X_1 - X_0)$) that does not vanish identically on $E \times E$. Hence $f \not\equiv 0$ on $E \times E$ implies $G \not\equiv 0$ on $E \times E$.

Each summand of G has degree $2(d - 1 + M) = 2(d + M - 1)$ in (X_0, Y_0, X_1, Y_1), so $\deg G \leq 2(d + M - 1) \leq 2(\deg_{\mathrm{E}}(D) + M - 1)$ (using $d \leq \deg_{\mathrm{E}}(D)$). The degree bound is attained (no leading-term cancellation): the k_0-th summand of the first sum is not divisible by $\ell_{Q_{k_0}}$, while every other summand of G is.

As an affine subvariety of $\mathbb{A}^4$, the product $E \times E$ is cut out by two cubics in disjoint variable sets (X_0, Y_0) and (X_1, Y_1), so it has dimension 2 and degree $3 \cdot 3 = 9$. By Theorem 4 applied to the subvariety $\{G = 0\} \cap (E \times E)$, which has dimension ≤ 1 and degree $\leq 9 \cdot 2(\deg_{\mathrm{E}}(D) + M - 1) = 18(\deg_{\mathrm{E}}(D) + M - 1)$, the number of $\mathbb{F}_q$-rational zeros of G on $E \times E$ is at most $18(\deg_{\mathrm{E}}(D) + M - 1) \cdot q$. Therefore:

$$\Pr[f(A_0, A_1) = 0] \leq \frac{18(\deg_{\mathrm{E}}(D) + M - 1) \cdot q}{(\#E(\mathbb{F}_q) - 1)(\#E(\mathbb{F}_q) - 3)}$$

□

Remark 1 (Non-Vanishing). *Suppose $f \equiv 0$ on $E \times E$ and $\deg_{\mathrm{E}}(D) < q$. For each slope λ, valid challenges realize $\geq q/3$ distinct evaluation points μ; since $\sum_k \beta_k/(z - z_\lambda(Q_k)) + \sum_j m_j/(z - z_\lambda(R_j))$ is a rational function of z with numerator degree $< q$ vanishing at all these μ, it is identically zero in $\mathbb{F}_q(z)$. By uniqueness of partial fractions, the residues match at every pole. For all but finitely many λ $(\leq \binom{d+M}{2})$, the z_λ-projections of the combined support are distinct, forcing a bijection $Q_k = R_{\sigma(k)}$ with $\beta_k \equiv -m_{\sigma(k)} \pmod q$, and $m_j = 0$ for unmatched R_j.*

2 Public-Coin Interactive Proofs for Discrete Log

In this section we formalize how the techniques of Section 1 give rise to 1-round and 3-round interactive proofs for showing discrete logs.

2.1 Preliminaries

A *public-coin interactive proof* (IP) [GMR85] is a protocol between an efficient verifier V and a possibly unbounded prover P in which V's messages are uniform coins and V eventually accepts or rejects.

Definition 1 (Public-Coin Interactive Proof [GMR85])**.** *A* public-coin interactive proof *(IP) for a relation $\mathcal{R}$ is a pair $\Pi = (\mathsf{P}, \mathsf{V})$ of PPT interactive algorithms in which every message sent by V is a fresh uniformly random string. On common input $\mathbb{x}$, with P additionally holding $\mathbb{w}$ and sampling a random tape $r \leftarrow\!\$\ \{0,1\}^*$, the parties run $\ell \geq 1$ rounds producing a transcript $\tau = (\pi_1, c_1, \pi_2, c_2, \ldots, c_{\ell-1}, \pi_\ell)$, where in round i the prover sends $\pi_i \leftarrow \mathsf{P}(\mathbb{x}, \mathbb{w}, c_1, \ldots, c_{i-1}; r)$ and, if $i < \ell$, the verifier samples $c_i \leftarrow\!\$\ \{0,1\}^*$ and sends it. After round ℓ the verifier outputs $\mathsf{V}(\mathbb{x}, \tau) \in \{\texttt{accept}, \texttt{reject}\}$ using fresh coins; we write $\langle \mathsf{P}(\mathbb{w}), \mathsf{V}\rangle(\mathbb{x})$ for this output. Π has* completeness error ϵ_{c} *if for every $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$:*

$$\Pr[\langle \mathsf{P}(\mathbb{w}), \mathsf{V}\rangle(\mathbb{x}) = \texttt{accept}] \geq 1 - \epsilon_{\mathsf{c}}(\kappa)$$

Knowledge soundness for an interactive proof states that there exists a polynomial-time algorithm that, given black-box rewinding access to any cheating prover P^*, recovers a witness for the relation with *almost* the same probability with which the prover P^* makes the verifier accept. Observe that occasionally the *relation for which the IP is knowledge sound might not coincide with the relation for which it's complete*: the extractor may recover a witness from a "relaxed" relation.

Definition 2 (Black-Box Knowledge Soundness [BG93])**.** *A ℓ-round public-coin IP $\Pi = (\mathsf{P}, \mathsf{V})$ is* knowledge sound *for $\mathcal{R}$, with* knowledge error ϵ_{s} *if there exists a PPT extractor $\mathcal{E}$ such that, for every (possibly unbounded) prover P^* and every statement $\mathbb{x}$, $\mathcal{E}^{\mathsf{P}^*}(\mathbb{x})$ outputs either $\perp$ or a string $\mathbb{w}$ satisfying:*

1. Validity: *whenever $\mathcal{E}^{\mathsf{P}^*}(\mathbb{x}) \neq \perp$, its output $\mathbb{w}$ satisfies $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$.*
2. Success:
$$\Pr\left[\mathcal{E}^{\mathsf{P}^*}(\mathbb{x}) \neq \perp\right] \geq \epsilon^*(\mathbb{x}) - \epsilon_{\mathsf{s}}(\kappa)$$
Where $\epsilon^*(\mathbb{x}) \coloneqq \Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(\mathbb{x}) = \texttt{accept}]$ is the acceptance probability of P^* on statement $\mathbb{x}$, taken over the coins of P^* and V.

The extractor is given black-box rewinding access to P^*: concretely, $\mathcal{E}$ queries the next-message function $\mathsf{P}^*(\mathbb{x}, c_1, \ldots, c_{i-1}; r)$ on any random tape $r \in \{0,1\}^*$ and any prefix of challenges $(c_1, \ldots, c_{i-1}) \in (\{0,1\}^*)^{i-1}$ of its choice, for any $i \in \{1, \ldots, \ell\}$.

2.2 Relations

Before formalizing the protocol, we define two relations: $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ and $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}]$. We will show a protocol that is complete for $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}]$, but knowledge sound for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$.

Relation (Knowledge Soundness). The protocol is parameterized by a fixed public basis $\vec{B} = (B_1, \ldots, B_k) \in \mathbb{G}^k$ with $B_i \neq B_j$ for all $i \neq j$, and will be *knowledge-sound* for the relation:

$$\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}] \coloneqq \left\{ \left(\mathbb{x} = P \in \mathbb{G}, \mathbb{w} = (\vec{n}) \in \mathbb{F}_p^k \right) \mid P = \sum_{i=1}^{k} [n_i] \cdot B_i \right\} \tag{1}$$

It's easy to see that without loss of generality we can take the B_j pairwise distinct: if $B_i = B_j$ for some $i \neq j$, then $[n_i]B_i + [n_j]B_j = [n_i + n_j]B_i$.

Relation (Completeness). The protocol is additionally parameterized by an *admissible set* $\mathcal{S} \subseteq \mathbb{F}_q[x]^2$ with $(0,0) \notin \mathcal{S}$. The protocol is *complete* for the *honest* subset:

$$\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}] := \left\{ \left(\mathbb{x} = P, \mathbb{w} = (\vec{n}) \in \mathbb{N}^k_{\geq 0}\right) \;\middle|\; \begin{array}{l} (\mathbb{x},\mathbb{w}) \in \mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}] \\ \exists\,(a(x),b(x)) \in \mathcal{S}: \\ \quad \deg_{\mathrm{E}}(a(x) - b(x)y) \leq \mathsf{d} \\ \quad (a(x) - b(x)y)_0 = (-P) + \sum_i n_i \cdot (B_i) \end{array} \right\} \tag{2}$$

Observe that $\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}] \subseteq \mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$, as expected.

> ***Comments On Prior Work.***
>
> We observe that this distinction between what an honest prover can demonstrate (completeness) and what the extractor can recover (knowledge soundness), is also present in prior work by Bassa / Eagen, but not made explicit. In particular, in the work of Bassa, the relation is defined as:
>
> $$\mathcal{R}_{DL} = \{(\mathbf{s} = (s_0,\ldots,s_k) \in \mathbb{Z}^{k+1}, 0 \leq s_i < p; P \in E) | P = \sum_{i=0}^{k} [s_i] \cdot B_i\}$$
>
> This is equivalent to $\mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$. Observe that in the relation the witness may be any integer $0 \leq s_i < p$, however, it's trivial to see that the protocol is not complete for witnesses for which $\sum s_i$ is large: the divisor would be too large to commit to. We capture this soundness/completeness "gap" explicitly using $\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$ and $\mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$.

In practice, the gap between $\mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$ and $\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$ can be eliminated by *ensuring* that there always exists a witness in $\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$; this will be the case in the concrete applications considered. The following example is also outlined by Bassa, and we show it slightly more formally with our more explicit notation:

Example 2 (Discrete Log via Binary Decomposition)**.** Consider the following relation:

$$\mathcal{R}^{\text{dlog}}_{G} := \{(P \in E(\mathbb{F}_q),\ x \in \mathbb{F}_p) \mid P = [x] \cdot G\} \subseteq E(\mathbb{F}_q) \times \mathbb{F}_p$$

We can construct a knowledge-sound and complete argument for $\mathcal{R}^{\text{dlog}}_{G}$ using any IP with knowledge soundness for $\mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$ and completeness for $\mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$ with $\mathsf{d} = \lceil \log_2(p) \rceil + 1$, where $\vec{B}$ is chosen as the binary-decomposition basis:

- Let $B_j := \left[2^{j-1}\right] \cdot G$ for $j \in [k]$ with $2^k \geq p$ (so $\vec{B}$ is fixed by G),
- Write $\varphi \colon \mathbb{F}_p^k \to \mathbb{F}_p$ for the map $\varphi(\vec{n}) = \sum_{j=1}^{k} n_j \cdot 2^{j-1}$

Then for the same $\mathcal{R}^{\text{dlog}}_{G}$ we get:

- *Completeness.* For every $x \in \mathbb{F}_p$, it is efficient to find $\vec{n} \in \{0,1\}^k$ st. $\varphi(\vec{n}) = x$ and:

$$(P,x) \in \mathcal{R}^{\text{dlog}}_{G} \implies (P,\vec{n}) \in \mathcal{R}^{\text{dlog-honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$$

- *Soundness.* For $(\mathbb{x} = P, \mathbb{w} = \vec{n})$ observe that $\varphi(\vec{n})$ is efficiently computable and:

$$(P,\vec{n}) \in \mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}] \implies (P,\varphi(\vec{n})) \in \mathcal{R}^{\text{dlog}}_{G}$$

The protocol outlined by Eagen for $\mathcal{R}^{\text{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$ is a one-round interactive proof (Def. 1) shown in Fig. 1, i.e. the communication pattern is: 1. the prover sends a message, 2. the verifier flips some coins and decides to `accept`/`reject`based on the message and its random tape. Motivated by the concrete composition with a CP-NIZKAoK we write the protocol in a form where all $\mathsf{P} \to \mathsf{V}$ messages are vectors over $\mathbb{F}_q$.

Protocol $\Pi^{\mathsf{MA}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$

Prover $\mathsf{P}(\mathbb{x}, \mathbb{w})$		**Verifier** $\mathsf{V}(\mathbb{x})$
// Compute Divisor Function		
Let $D(x,y) = a(x) - b(x)y \in \mathbb{F}_q[x,y]$ st. $(D(x,y))_0 = (-P) + \sum n_i \cdot (B_i)$		
// Reduce Witness mod q		
Let $\vec{m} = \vec{n} \bmod q \in \mathbb{F}_q^k$		
	$\vec{m} \in \mathbb{F}_q^k$ $a(x), b(x) \in \mathbb{F}_q[x]$ $\longrightarrow$	
		// Check that P is Well-Formed
		assert $P \in E(\mathbb{F}_q)$
		// Check Degree Bound
		assert $\deg_{\mathrm{E}}(a(x) - b(x)y) \leq \mathsf{d}$
		// Ensure $D \neq 0$
		assert $(a(x), b(x)) \in \mathcal{S}$
		// Sample Challenge Points
		$A_0, A_1 \leftarrow\!\!\$\ E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ with $A_0 \neq \pm A_1$
		$A_2 \leftarrow -(A_0 + A_1)$
		Let L be the line through A_0, A_1, A_2
		// Weil Reciprocity Check
		$\mathsf{lhs} \leftarrow \sum_{i=0}^{2} \frac{D'(A_i)}{D(A_i)} \cdot \frac{dx(A_i)}{dz}$
		$\mathsf{rhs} \leftarrow \frac{-1}{L(-P)} + \sum_{i=1}^{k} \frac{-1}{L(B_i)} \cdot m_i$
		assert $\mathsf{lhs} = \mathsf{rhs}$

Figure 1: Merlin-Arthur Protocol $\Pi^{\mathsf{MA}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]/\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$

Construction 1 (One-Round Proof)**.** Let $\mathsf{d} = \mathsf{poly}(\kappa)$ be a protocol parameter affecting the knowledge error of the protocol and the completeness relation. We formalize the one-round protocol $\Pi^{\mathsf{MA}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ implicit in the work of Eagen [Eag22] in Fig. 1.

Observation 1 (Polynomial Degree Bound)**.** *Any* $\mathsf{d} = \mathsf{poly}(\kappa)$ *suffices for a field of exponential size in the security parameter* κ*: soundness degrades only as* $O(\mathsf{d}/q)$ *which remains negligible. In particular, upper bounding the degree by the adversary's running time suffices. In the concrete application* d *is going to be a constant, defined by the circuit which checks the execution of the verifier, see Section 3.*

Observation 2 (Accepting Transcript with $-P \in \{B_1, \ldots, B_k\}$)**.** *Let* $B_1 = -P$*, all other* B_j *can be arbitrary,* $m_1 = q-1$*,* $m_j = 0$ *for* $j \geq 2$*, and* $D = 1$ *i.e.* $a(x) = 1$*,* $b(x) = 0$*. Now observe:*

- ***Degree Check:*** $\deg_{\mathrm{E}}(D) = 0 \leq \mathsf{d}$.
- ***Equality Check:*** $D' = 0$ *so* $\mathsf{lhs} = 0$*; and using* $B_1 = -P$ *and* $-(q-1) \equiv 1 \pmod{q}$*:*

$$\mathsf{rhs} = \tfrac{-1}{L(-P)} + \tfrac{-(q-1)}{L(B_1)} = \tfrac{-1}{L(-P)} + \tfrac{1}{L(-P)} = 0$$

Hence the verifier accepts with probability 1*, however a naive unsigned lift* $n_i \coloneqq \mathsf{lift}(m_i) \in \mathbb{N}$ *gives* $\vec{n} = (q-1, 0, \ldots, 0)$*, which fails the relation's inner-product check on the curve:* $[q-1]B_1 = [q-1](-P) \neq P$ *(since* $p \neq q$*).*

Observation 3 (Accepting Transcript with $D \equiv 0$)**.** *Let* $a(x) = 0$*,* $b(x) = 0$ *(so* $D \equiv 0$*) and* $\vec{m} \in \mathbb{F}_q^k$ *arbitrary. In* $\Pi^{\mathsf{IP}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$*, the hint check* $h_i \cdot D(A_i) = D'(A_i)$ *reduces to* $0 = 0$ *and imposes no constraint on* $\vec{h} \in \mathbb{F}_q^3$*; the remaining equation reduces to:*

$$\sum_{i=0}^{2} h_i \cdot \tfrac{dx(A_i)}{dz} = g + \sum_{j=1}^{k} \tfrac{-m_j}{L(B_j)}$$

This is one linear equation in three unknowns, always solvable. Hence the verifier accepts any $\vec{m}$*, while the extractor of Thm. 7 returns* $\perp$ *whenever* $P \neq \sum_i [\mathsf{lift}(m_i)] \cdot B_i$*. The soundness argument implicitly assumes* $D \neq 0$ *via Lems. 6 and 7; the verifier check* $(a(x), b(x)) \in \mathcal{S}$ *in Figs. 1 and 2 discharges this hypothesis whenever* $(0,0) \notin \mathcal{S}$*.*

Instantiations of $\mathcal{S}$. Soundness of $\Pi^{\mathsf{MA}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}/\Pi^{\mathsf{IP}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ depends on $(0,0) \notin \mathcal{S}$: any $(a,b) \in \mathcal{S}$ satisfies $D = a(x) - b(x)y \neq 0$ in $\mathbb{F}_q[E]$. The relation $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S}}[\vec{B}]$ depends on the choice of $\mathcal{S}$, here are some possible constructions:

- *Maximal:* $\mathcal{S}_{\mathsf{Max}} \coloneqq \mathbb{F}_q[x]^2 \setminus \{(0,0)\}$.
- *Parker* [Par24a, Par24b]: $\mathcal{S}_{\mathsf{Parker}} \coloneqq \{(a,b) : a_1 = 1\}$ (coefficient of x in $a(x)$ is 1).
- *Eagen* [Eag22, §3.4]: $\mathcal{S}_{\mathsf{Eagen}} \coloneqq \{(a,b) : a_0 = 1\}$ (constant coefficient in $a(x)$ is 1).
- *Hash:* $\mathcal{S}_{\mathsf{Hash},\vec{r}} \coloneqq \{(a,b) : \langle \vec{r},\, a \| b \rangle \neq 0\}$ for $\vec{r} \in \mathbb{F}_q^n$ ($a \| b$ concatenates the coefficients).

For the composition with Bulletproofs in Section 3 we need the check $(a,b) \in \mathcal{S}$ to have an efficient R1CS relation over the committed coefficients. The Maximal construction is the largest possible admissible set, but checking $(a,b) \neq (0,0)$ over a vector of coefficients does not admit a cheap R1CS relation; Parker's and Eagen's constructions impose a single equality constraint instead. The Hash construction is slightly less efficient than Parker's and Eagen's, but has the advantage that completeness holds for *every divisor*: the verifier (or prover) can sample a fresh $\vec{r}$ for each run of the protocol, whereas the others have a negligible subset of divisors which can never be proven. The universal hash can be checked in-circuit by witnessing inv and showing $\langle \vec{r},\, a \| b \rangle \cdot \mathsf{inv} = 1$, a single R1CS constraint.

Comments On Prior Work.

The extractor of Bassa [Bas25] fails to account for the case outlined in Observation 2; it's implicitly assumed to not occur. The argument outlined does not hold when $-P \in \{B_1, \ldots, B_k\}$. The wraparound of coefficients from $\mathbb{Z}$ to $\mathbb{F}_p$ is flagged in [Bas24b, §6] under the clause "we can assume with high probability that the $z(P_i)$ are distinct"; a genericity claim about the random slope λ that silently presupposes the underlying P_i are distinct. The verifier in [Bas24a, Fig. 1] has no check ruling this out, and the review in [GSSS25a, §3.2] does not flag it. However, in application, it is crucial: in the concrete use case the prover can pick P and the $B_1, \ldots, B_k$ are fixed, hence we want a protocol which is sound, even when $-P \in \{B_1, \ldots, B_k\}$. Luckily, the fix is very simple and does not require changes to the protocol: whenever $-P \in \{B_1, \ldots, B_k\}$ we can *trivially* compute a witness (without even running the prover).

Let $\mathfrak{E}_{\mathsf{deg}}$ denote the event, over $(A_0, A_1) \leftarrow\!\$\ E(\mathbb{F}_q)^2$, that some denominator in the verifier's field expression vanishes ($D(A_i) = 0$, $L(-P) = 0$, $L(B_j) = 0$, or dx/dz denominator at A_i). On $\mathfrak{E}_{\mathsf{deg}}$ the verifier's output is implementation-defined; errors below absorb it.

Comments On Prior Work.

Bassa [Bas24b] proposes to enforce support disjointness ($D(A_i) \neq 0$) via verifier-side resampling:

> "Alternatively, the verifier can just sample A_0, A_1 from the set $E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$, compute $A_2 = -A_0 - A_1$ and check that $x(A_0) \neq x(A_1)$ and $D(A_i) \neq 0$ for $i = 0, 1, 2$, resampling if necessary. Hence we can assume that the support of D and L are disjoint."

We note that in the composition (see Section 3), it's not possible for the verifier to check this condition: the challenge is sampled out-of-circuit and the out-of-circuit verifier (crucially) does not see D. Rejecting any line L which intersects with D in-circuit would be very expensive and defeat the purpose of the optimization. The later Bassa paper for the interactive protocol [Bas25] does not adopt verifier-side resampling.

Theorem 7. *Assume* $\mathsf{d} < \min\{p, q\}$. $\Pi^{\mathsf{MA}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ *(Construction 1) is a 1-round IP with:*

- ***Completeness*** *for* $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}]$*, with error:*

$$\epsilon_{\mathsf{c}} := \frac{3(\mathsf{d}+1)}{q+1-2\sqrt{q}} \leq \frac{6(\mathsf{d}+1)}{q}$$

- ***Knowledge Soundness*** *for* $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ *with* ***straight-line extraction*** *and error:*

$$\epsilon_{\mathsf{s}} := \frac{18(\mathsf{d}+k)\cdot q}{(q-2\sqrt{q})(q-2\sqrt{q}-2)} \leq \frac{96(\mathsf{d}+k)}{q}$$

Both bounds include $\Pr[\mathfrak{E}_{\mathsf{deg}}]$ *plus the residual on* $\neg\mathfrak{E}_{\mathsf{deg}}$*. The simplified (rightmost) bounds hold for all* $q \geq 16$*, which is satisfied in any cryptographic instantiation.*

Proof. The completeness error follows immediately by Thm. 1 and Lem. 2. One detail of note is that the divisor can be efficiently computed, and the original paper [Eag22] demonstrates how to achieve this. Since $(0,0) \notin \mathcal{S}$, the check $(a,b) \in \mathcal{S}$ implies $(a,b) \neq (0,0)$, hence $D = a(x) - b(x)y \neq 0$ as an element of $\mathbb{F}_q[E]$; satisfying the $D \neq 0$ hypothesis of Lem. 7 and Lem. 6 (see also Observation 3). We define an extractor $\mathcal{E}^{\mathsf{P}^*}(\mathbb{x})$ for any (potentially unbounded) Turing machine P^* as follows:

Extractor $\mathcal{E}^{\mathsf{P}^*}(\mathbb{x})$

// Special Case: $-P \in \{B_j\}$
if $\exists\, j^* \in [k] : B_{j^*} = -P$:
  $\vec{n} \leftarrow \vec{0} \in \mathbb{F}_p^k, n_{j^*} \leftarrow -1$
  return $\vec{n}$
// General Case: $-P \notin \{B_j\}$
$r \leftarrow\!\$\ \{0,1\}^*$
$(\vec{m}, a(x), b(x)) \leftarrow \mathsf{P}^*(\mathbb{x}; r)$
if $\mathsf{lift}(\vec{m}) \notin [0, \mathsf{d}]^k$**return** $\perp$
$\vec{n} \leftarrow \mathsf{lift}(\vec{m}) \bmod p \in \mathbb{F}_p^k$
if $P \neq \sum_i [n_i] \cdot B_i$**return** $\perp$
return $\vec{n}$

Writing $\mathsf{lift}(z) : \mathbb{F}_q \mapsto [0, q) \subseteq \mathbb{N}_{\geq 0}$ for taking the canonical integer representative of $z \in \mathbb{F}_q$. The special case handles $-P \in \{B_1, \ldots, B_k\}$: the trivial witness $n_{j^*} = -1$ satisfies $[-1] \cdot B_{j^*} = [-1](-P) = P$ unconditionally, so the extractor can return it without even reading P^*'s message. In the general case ($-P \notin \{B_j\}$) the extractor takes the naive lift of each m_j and checks it lies in the degree bound. If the extractor does not return $\perp$, the final check $P = \sum_i [n_i] \cdot B_i$ ensures its output is a valid witness. It remains to bound the probability that the extractor returns $\perp$ conditioned on the verifier accepting. We define two bad events and bound them separately.

Define the function $f \colon E \times E \dashrightarrow \mathbb{F}_q$:

$$f(Q_0, Q_1) \coloneqq \left(\sum_{i=0}^{2} \frac{D'(Q_i)}{D(Q_i)} \cdot \frac{dx(Q_i)}{dz} \right) - \left(\frac{-1}{L_Q(-P)} + \sum_{j=1}^{k} \frac{-1}{L_Q(B_j)} \cdot m_j \right)$$

Where $Q_2 = -(Q_0 + Q_1)$ and L_Q is the line through Q_0, Q_1, Q_2. Recall that the verifier accepts iff $f(A_0, A_1) = 0$ and the degree bound on D holds.

- **Event $\mathfrak{E}_{\mathsf{NotEq}}$**: $f \not\equiv 0$ on $E \times E$, yet $f(A_0, A_1) = 0$. Since (A_0, A_1) is sampled uniformly from $E(\mathbb{F}_q)^2$, Lem. 7 (instantiated with the $k+1$ evaluation points $\{-P, B_1, \ldots, B_k\}$ and $\deg_{\mathrm{E}}(D) \leq \mathsf{d}$), combined with the Hasse bound Thm. 2 ($\#E(\mathbb{F}_q) \geq q + 1 - 2\sqrt{q}$), gives:

$$\Pr[\mathfrak{E}_{\mathsf{NotEq}}] \leq \frac{18(\mathsf{d}+k) \cdot q}{(q - 2\sqrt{q})(q - 2\sqrt{q} - 2)} = \epsilon_{\mathsf{s}}$$

- **Event $\mathfrak{E}_{\mathsf{BadWit}}$**: the verifier accepts, but the extractor returns $\perp$. We claim:

$$\Pr[\mathfrak{E}_{\mathsf{BadWit}} \mid \neg\, \mathfrak{E}_{\mathsf{NotEq}}] = 0$$

  Conditioned on $\neg\, \mathfrak{E}_{\mathsf{NotEq}}$ and the verifier accepting, we must have $f \equiv 0$ on $E \times E$: indeed, acceptance requires $f(A_0, A_1) = 0$, and if additionally $f \not\equiv 0$ then $\mathfrak{E}_{\mathsf{NotEq}}$ holds, contradicting the conditioning. It therefore suffices to show that $f \equiv 0$ (together with the verifier's degree check) implies that the extractor succeeds.

  Suppose $f \equiv 0$ as a rational function on $E \times E$. Write $L = L_{\lambda,\mu} = y - \lambda \cdot x - \mu$ for the line determined by (A_0, A_1), equivalently by its slope λ and intercept μ. We show in the following steps that $f \equiv 0$ forces the divisor of zeros of D to have the claimed form and the extractor to return a valid witness.

  Step 1. *Identify the LHS of f as a log-derivative.* By Lem. 6, for each fixed line L:

$$\sum_{i=0}^{2} \frac{D'(A_i)}{D(A_i)} \cdot \frac{dx(A_i)}{dz} = \mathcal{L}_L\big(N_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)\big)\big((L = 0)\big)$$

Where $\mathcal{L}_L$ is the logarithmic derivative with respect to L in the degree-3 extension $\mathbb{F}_q(E)/\mathbb{F}_q(L)$.

Step 2. *Identify the RHS of f as a log-derivative.* Let $(D)_0 = \sum_\alpha n_\alpha \cdot (Q_\alpha)$ be the actual divisor of zeros of D, with integer multiplicities $n_\alpha \geq 1$ and distinct points $Q_\alpha \in E(\overline{\mathbb{F}_q})$. Since $D \in \mathbb{F}_q[E]$ has poles only at $\mathcal{O}$ with order $\deg_{\mathrm{E}}(D) = \sum_\alpha n_\alpha$, the norm is a polynomial in t:

$$N_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)(t) = c_L \cdot \prod_\alpha (t - L(Q_\alpha))^{n_\alpha}$$

Where c_L is some constant depending on D and L. Taking the log-derivative (in characteristic q, multiplicities reduce mod q) and evaluating at $t = 0$:

$$\mathcal{L}_L\big(N_{\mathbb{F}_q(E)/\mathbb{F}_q(L)}(D)\big)(0) = -\sum_\alpha \frac{n_\alpha \bmod q}{L(Q_\alpha)}$$

Similarly, the RHS of f is the log-derivative at $t = 0$ of the "expected norm" $\widetilde{N}_L(t) = (t - L(-P)) \cdot \prod_{j=1}^{k} (t - L(B_j))^{m_j}$, giving:

$$-\sum_\alpha \frac{n_\alpha \bmod q}{L(Q_\alpha)} = \frac{-1}{L(-P)} + \sum_{j=1}^{k} \frac{-m_j}{L(B_j)} \tag{3}$$

Step 3. *View* (3) *as an identity in* (λ, μ). The hypothesis $f \equiv 0$ on $E \times E$ means (3) holds as an identity of rational functions in (λ, μ), where each term $1/L(X) = 1/(y(X) - \lambda \cdot x(X) - \mu)$ is a rational function of the line parameters.

Step 4. *Match poles in μ for fixed generic λ.* Fix $\lambda \in \mathbb{F}_q$ and view (3) as a rational function of μ. Each $1/L(X)$ has a simple pole in μ at $\mu = \mu_X(\lambda) \coloneqq y(X) - \lambda \cdot x(X)$ with residue -1. For (3) to hold as a rational function in μ, the two sides must have the same poles with matching residues.

Step 5. *Separation of points by generic lines.* For any two distinct points $R, S \in E(\overline{\mathbb{F}_q})$, $\mu_R(\lambda) = \mu_S(\lambda)$ holds for at most one value of λ:

- If $x(R) \neq x(S)$: $\mu_R = \mu_S$ gives $\lambda = (y(R) - y(S))/(x(R) - x(S))$, a single value.
- If $x(R) = x(S)$: then $\mu_R = \mu_S$ iff $y(R) = y(S)$. Since R, S lie on E, are distinct, and $x(R) = x(S)$, we have $y(R) = -y(S) \neq y(S)$ (assuming $\mathrm{char} \neq 2$), so $\mu_R(\lambda) \neq \mu_S(\lambda)$ for every $\lambda \in \mathbb{F}_q$.

Hence for all but finitely many $\lambda \in \mathbb{F}_q$, the values $\{\mu_{Q_\alpha}(\lambda)\}_\alpha$ and $\{\mu_{-P}(\lambda)\} \cup \{\mu_{B_j}(\lambda)\}_j$ each consist of distinct poles.

Step 6. *Match residues to get the combined-residue identity.* For such generic λ, the two sides of (3) must have the same poles with matching residues (in $\mathbb{F}_q$). For each distinct $R \in E(\overline{\mathbb{F}_q})$:

$$n_R \equiv \mathbb{1}[R = -P] + \sum_{j\,:\,B_j = R} m_j \pmod q \tag{4}$$

Where $n_R \in \mathbb{Z}_{\geq 0}$ denotes the coefficient of R in $(D)_0$ (i.e., the multiplicity of R as a zero of D when $R \in \{Q_\alpha\}$, and 0 otherwise).

Step 7. *Lift combined residues to integer multiplicities.* Each n_R is a non-negative integer satisfying $n_R \leq \deg_{\mathrm{E}}(D) \leq \mathsf{d} < q$, so n_R and its reduction $n_R \bmod q$ agree as elements of $[0, q)$. The RHS of (4) is an element of $\mathbb{F}_q$; applying $\mathsf{lift}(\cdot) \in [0, q)$

to take its canonical integer representative, the congruence lifts to the integer equality:

$$n_R = \mathsf{lift}(\mathbb{1}[R = -P] + \textstyle\sum_{j\,:\,B_j = R} m_j) \in [0, \mathsf{d}] \quad \text{for every distinct } R \qquad (5)$$

By distinctness of $\vec{B}$, the m_j-sum reduces to m_j when $R = B_j$ and to 0 when $R \notin \{B_1, \ldots, B_k\}$. For $R \notin \{-P, B_1, \ldots, B_k\}$ the RHS is 0 in $\mathbb{F}_q$, so $n_R = 0$; hence $\mathrm{supp}((D)_0) \subseteq \{-P, B_1, \ldots, B_k\}$.

Step 8. *The extractor returns a valid $\mathbb{F}_p$-witness.* We split on whether $-P \in \{B_1, \ldots, B_k\}$.

Special case, $-P = B_{j^*}$ for some j^*. The extractor's special-case branch returns $n_{j^*} = -1 \in \mathbb{F}_p$ and $n_j = 0$ for $j \neq j^*$, *without inspecting* $\vec{m}$. The witness is valid unconditionally: $\sum_i [n_i] \cdot B_i = [-1] \cdot B_{j^*} = [-1](-P) = P$.

General case, $-P \notin \{B_1, \ldots, B_k\}$. Then (5) at $R = -P$ has an empty m_j-sum, so $n_{-P} = 1$. For each $j \in [k]$, (5) at $R = B_j$ gives $\mathsf{lift}(m_j) = n_{B_j} \in [0, \mathsf{d}]$, so the extractor's check $\mathsf{lift}(m_j) \leq \mathsf{d}$ passes (under $f \equiv 0$), and the extractor assigns $n_j = \mathsf{lift}(m_j) \bmod p = n_{B_j}$ (using $n_{B_j} \leq \mathsf{d} \ll p$, so reduction mod p is trivial).

By Thm. 1, since $(D) = (D)_0 - \deg_{\mathrm{E}}(D) \cdot (\mathcal{O})$ is principal, $\sum_R [n_R] \cdot R = \mathcal{O}$ in $E(\mathbb{F}_q)$. Using $n_{-P} = 1$ and $\mathrm{supp}((D)_0) \subseteq \{-P, B_1, \ldots, B_k\}$:

$$[1] \cdot (-P) + \sum_{j=1}^{k} [n_{B_j}] \cdot B_j = \mathcal{O}, \qquad \text{i.e. } P = \sum_{j=1}^{k} [n_j] \cdot B_j$$

So the final relation check passes.

In both cases the extractor returns $\vec{n} \in \mathbb{F}_p^k$, a valid witness for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$; $\mathfrak{E}_{\mathsf{BadWit}}$ does not occur. Therefore $\Pr[\mathfrak{E}_{\mathsf{BadWit}} \mid \neg\, \mathfrak{E}_{\mathsf{NotEq}}] = 0$.

We conclude by the conditional decomposition:

$$\begin{aligned} \Pr\left[\mathcal{E}^{\mathsf{P}^*}(\mathbb{x}) = \bot \wedge \mathtt{accept}\right] &= \Pr[\mathfrak{E}_{\mathsf{BadWit}}] \leq \Pr[\mathfrak{E}_{\mathsf{NotEq}}] + \Pr[\mathfrak{E}_{\mathsf{BadWit}} \mid \neg\, \mathfrak{E}_{\mathsf{NotEq}}] \\ &\leq \epsilon_{\mathsf{s}} + 0 = \epsilon_{\mathsf{s}} \end{aligned}$$

□

Instead of a one-round Merlin-Arthur protocol, the prior work of Bassa [Bas25] describes a *three-round interactive proof* (IP) which maps better onto an optimized composition with a CP-NIZKAoK: an additional message (after the verifier challenge) enables the prover to "help" the verifier by providing him with field inverses. From a theoretical perspective, we note that *this transformation is generic*: any MA protocol can be transformed into a *three-round IP* in which the *verifier can be expressed as a depth 1 circuit* as follows: letting the last round consist of the *wire assignment of an accepting verifier* and changing the IP verifier to check the correctness of all wire assignments in parallel. We include a variant of the Bassa [Bas25] diagram for the three-round IP in Fig. 2 for completeness and prove that it is a knowledge-sound IP for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ by reducing to the soundness of the MA protocol.

Construction 2 (Three-Round IP)**.** We formalize the three-round protocol $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ in Fig. 2.

Theorem 8. $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (*Construction 2*) *is a three-round IP with completeness for* $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}]$ *and knowledge soundness for* $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ *via a* straight-line *extractor; both errors match* $\Pi^{\mathsf{MA}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (*Thm. 7*)*:* $\epsilon_{\mathsf{c}} = \epsilon^{\mathsf{MA}}_{\mathsf{c}}$ *and* $\epsilon_{\mathsf{s}} = \epsilon^{\mathsf{MA}}_{\mathsf{s}}$*, where* $\epsilon^{\mathsf{MA}}_{\mathsf{c}}, \epsilon^{\mathsf{MA}}_{\mathsf{s}}$ *are the errors of* $\Pi^{\mathsf{MA}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (*Thm. 7*).

Protocol $\Pi^{\mathsf{IP}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$

Prover $\mathsf{P}(\mathbb{x}, \mathbb{w})$		**Verifier** $\mathsf{V}(\mathbb{x})$
// Compute Divisor Function		
Let $D(x,y) = a(x) - b(x)y \in \mathbb{F}_q[x,y]$ st. $(D(x,y))_0 = (-P) + \sum n_i \cdot (B_i)$		
// Reduce Witness mod q		
Let $\vec{m} = \vec{n} \bmod q \in \mathbb{F}_q^k$		
	$\vec{m} \in \mathbb{F}_q^k,\ a(x), b(x) \in \mathbb{F}_q[x]$ $\longrightarrow$	
		// Check that P is Well-Formed
		assert $P \in E(\mathbb{F}_q)$
		// Check Degree Bound
		assert $\deg_{\mathrm{E}}(a(x) - b(x)y) \leq \mathsf{d}$
		// Ensure $D(x,y) \neq 0$
		assert $(a(x), b(x)) \in \mathcal{S}$
		// Sample Challenge Points
		$A_0, A_1 \leftarrow\$\ E(\mathbb{F}_q) \setminus \{\mathcal{O}\}$ with $A_0 \neq \pm A_1$
	$\longleftarrow$ $A_0, A_1 \in E(\mathbb{F}_q)$	
// Compute Helpers for Verifier		
$A_2 \leftarrow -(A_0 + A_1)$		
Let L be the line through A_0, A_1, A_2		
$h_0 \leftarrow D'(A_0)/D(A_0)$		
$h_1 \leftarrow D'(A_1)/D(A_1)$		
$h_2 \leftarrow D'(A_2)/D(A_2)$		
$g \leftarrow -1/L(-P)$		
	$h_0, h_1, h_2, g \in \mathbb{F}_q$ $\longrightarrow$	
		$A_2 \leftarrow -(A_0 + A_1)$
		Let L be the line through A_0, A_1, A_2
		// Weil Reciprocity Check
		$\mathsf{lhs} \leftarrow \sum_{i=0}^{2} h_i \cdot \frac{dx(A_i)}{dz}$
		$\mathsf{rhs} \leftarrow g + \sum_{i=1}^{k} \frac{-1}{L(B_i)} \cdot m_i$
		assert $\mathsf{lhs} = \mathsf{rhs}$
		// Verify Hints
		assert $h_0 \cdot D(A_0) = D'(A_0)$
		assert $h_1 \cdot D(A_1) = D'(A_1)$
		assert $h_2 \cdot D(A_2) = D'(A_2)$
		assert $g \cdot L(-P) = -1$

Figure 2: Three-Round Interactive Protocol $\Pi^{\mathsf{IP}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q,E(\mathbb{F}_q)}[\vec{B}]$; conversion of $\Pi^{\mathsf{MA}}_{\mathbb{F}_q,E(\mathbb{F}_q),\mathsf{d},\mathcal{S},\vec{B}}$ with a "helping" prover to improve the verifier's efficiency.

Proof. Completeness is immediate: the honest IP prover sets $h_i = D'(A_i)/D(A_i)$ for $i \in \{0,1,2\}$ and $g = -1/L(-P)$, which satisfies the hint checks; substituting these into the IP Weil reciprocity check yields the MA Weil reciprocity check. Hence $\epsilon_{\mathsf{c}} = \epsilon_{\mathsf{c}}^{\mathsf{MA}}$.

We reduce knowledge soundness of the IP to knowledge soundness of the MA. Fix any first round $(\vec{m}, a, b)$ and challenge (A_0, A_1), and suppose the IP verifier accepts on some third-round message (h_0, h_1, h_2, g). On $\neg\mathfrak{E}_{\mathsf{deg}}$ we show the MA verifier accepts on $((\vec{m}, a, b), (A_0, A_1))$: the IP hint checks $h_i \cdot D(A_i) = D'(A_i)$ force $h_i = D'(A_i)/D(A_i)$, and $g \cdot L(-P) = -1$ forces $g = -1/L(-P)$. Substituting these values into the IP Weil reciprocity check recovers the MA Weil reciprocity check, so the MA verifier accepts. The $\mathfrak{E}_{\mathsf{deg}}$ case is absorbed into $\epsilon_{\mathsf{s}}^{\mathsf{MA}}$. Hence IP accepts $\implies$ MA accepts on the same $((\vec{m}, a, b), (A_0, A_1))$. The IP first-round message $(\vec{m}, a, b)$ coincides with the MA first-round message. Define the IP extractor as the MA extractor of Thm. 7 applied to this message. By the reduction, its success probability is at least that of the MA extractor conditioned on MA acceptance; hence $\epsilon_{\mathsf{s}} = \epsilon_{\mathsf{s}}^{\mathsf{MA}}$. □

> ***Comments On Prior Work.***
>
> In a prior work of Bassa [Bas25], knowledge soundness of the protocol in Fig. 2 is proved by showing that it is $13kq$-special sound: recovering the witness from any set of $13kq$ accepting transcripts sharing the first round message. The argument is correct (except for the fixable edge-cases discussed earlier), but *it does not lead to a polynomial-time extractor*: even though there is an *exponential gap* between the $13kq$-special soundness and the size of the challenge space ($(\#E(\mathbb{F}_q) - 1)^2$), $13kq$ is *exponential* in the security parameter, which means *the extractor will need to rewind the prover an exponential number of times* to obtain sufficiently many transcripts to recover a witness using the $13kq$-special soundness extractor. In particular, for the Fiat-Shamir compiled protocol, the extraction strategy of e.g. Attema, Fehr and Klooß [AFK23] is (expected) polynomial time only if the tree of transcripts is also of polynomial size; it is *not sufficient that the set of bad challenges has inverse exponential density* for the general knowledge extraction strategies to apply.
>
> As briefly noted in the subsequent Cypher Stack review [GSSS25a], all this seems intuitively unnecessary: the protocol has the prover *send the purported witness* in the first round. A working extractor can therefore be obtained more easily, by simply recovering the witness from the first round message. All that remains is to argue that this extracted witness is valid with all-but-negligible probability. With this approach, it is not an issue that there is an exponential number of "bad" challenges as long as the subset is exponentially sparse in the challenge space. The result is a knowledge-sound 3-round IP, not a special-sound Σ-protocol, and suffices for the tight composition of Section 3: the straight-line extractor feeds directly into Thm. 10, paying only the unavoidable $Q \cdot \epsilon_{\mathsf{s}}$ factor of any Fiat–Shamir compilation (a union bound over the malicious prover's Q candidate first-round messages).

3 Compositions of Interactive Proofs with Non-Interactive Arguments

We show that composing a 3-move interactive protocol with a commit-and-prove NIZKAoK in the ROM (random oracle model) yields a zero-knowledge and extractable argument (soundness is computational) for the committed relation. We furthermore observe that if the NIZKAoK is simulation extractable, then *the composition is also simulation extractable*. Prior works showing that Bulletproofs is simulation extractable [DG23, GOP+21] therefore apply and show that the composition of the IP of Eagen/Bassa [Eag22] [Bas24a] with Bulletproofs is simulation extractable.

Definition 3 (Commitment Scheme)**.** *A **commitment scheme** is a pair of PPT algorithms* (Setup, Com)*:*

1. $\mathsf{Setup}(1^\kappa) \to \mathsf{pp}$ *generates public parameters.*
2. $\mathsf{Com}(\mathsf{pp}, \mathsf{m}; r) \to \mathsf{cm}$ *maps a message* m *and randomness* r *to a commitment* cm*.*

Definition 4 (Perfectly Hiding)**.** *A commitment scheme* $(\mathsf{Setup}, \mathsf{Com})$ *is* ***perfectly hiding*** *if for every (possibly unbounded) adversary* $\mathcal{A}$ *and* $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$*: given any two messages* $\mathsf{m}_0, \mathsf{m}_1$ *in the message space, the distributions* $\mathsf{Com}(\mathsf{pp}, \mathsf{m}_0; r)$ *and* $\mathsf{Com}(\mathsf{pp}, \mathsf{m}_1; r)$ *are identical over* r*.*

Definition 5 (Computationally Binding)**.** *A commitment scheme* $(\mathsf{Setup}, \mathsf{Com})$ *is* ***computationally binding*** *if for every PPT adversary* $\mathcal{A}$ *and* $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$*:*

$$\Pr[\mathsf{Com}(\mathsf{pp}, \mathsf{m}_0; r_0) = \mathsf{Com}(\mathsf{pp}, \mathsf{m}_1; r_1) \wedge \mathsf{m}_0 \neq \mathsf{m}_1] \leq \mathsf{negl}(\kappa)$$

Where $(\mathsf{m}_0, r_0, \mathsf{m}_1, r_1) \leftarrow \mathcal{A}(\mathsf{pp})$*.*

Construction 3 (Pedersen Commitment)**.** Let E_{out} be an elliptic curve over $\mathbb{F}_{\mathsf{out}}$ with a prime-order subgroup $\mathbb{G}_{\mathsf{out}} \subseteq E_{\mathsf{out}}(\mathbb{F}_{\mathsf{out}})$ of order q. The Pedersen commitment scheme for message space $\mathbb{F}_q^n$ and randomness space $\mathbb{F}_q$ is defined as:

- $\mathsf{Setup}(1^\kappa)$: sample $G_1, \ldots, G_n, H \leftarrow\!\!\$\ \mathbb{G}_{\mathsf{out}}$ and output $\mathsf{pp} = (G_1, \ldots, G_n, H)$.
- $\mathsf{Com}(\mathsf{pp}, \vec{\mathsf{m}}; r)$: output $\mathsf{cm} = [r] \cdot H + \sum_{i=1}^{n} [\mathsf{m}_i] \cdot G_i \in \mathbb{G}_{\mathsf{out}}$.

The scheme is perfectly hiding (Def. 4) and computationally binding (Def. 5) under the discrete logarithm assumption in $\mathbb{G}_{\mathsf{out}}$.

Definition 6 (CP-NIZKAoK in the ROM)**.** *Let* $(\mathsf{Setup}, \mathsf{Com})$ *be a commitment scheme and let* $\mathcal{R}$ *be a relation. A* ***Commit-and-Prove Non-Interactive Zero-Knowledge Argument of Knowledge*** *(CP-NIZKAoK) for* $\mathcal{R}$ *with respect to* Com *in the random oracle model is a pair of PPT algorithms* $(\mathsf{Prove}, \mathsf{Verify})$ *with oracle access to a random oracle* H*. On public parameters* $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$*:*

- $\mathsf{Prove}^H(\mathsf{pp}, \mathbb{x}, \mathbb{w}, \vec{\mathsf{cm}}, \vec{\mathsf{m}}, \vec{r}) \to \pi$.
- $\mathsf{Verify}^H(\mathsf{pp}, \mathbb{x}, \vec{\mathsf{cm}}, \sigma) \to \{\mathtt{accept}, \mathtt{reject}\}$

Define the ***commit-and-prove relation****:*

$$\mathcal{R}_{\mathsf{cp}} := \{((\mathbb{x}, \vec{\mathsf{cm}}), (\mathbb{w}, \vec{\mathsf{m}}, \vec{r})) | \vec{\mathsf{cm}} = \mathsf{Com}(\mathsf{pp}, \vec{\mathsf{m}}; \vec{r}) \wedge (\mathbb{x}, \mathbb{w}, \vec{\mathsf{m}}) \in \mathcal{R}\}$$

We require ***completeness****, i.e. for every* $((\mathbb{x}, \vec{\mathsf{cm}}), (\mathbb{w}, \vec{\mathsf{m}}, \vec{r})) \in \mathcal{R}_{\mathsf{cp}}$*:*

$$\Pr\left[\mathsf{Verify}^H\left(\mathsf{pp}, \mathbb{x}, \vec{\mathsf{cm}}, \mathsf{Prove}^H(\mathsf{pp}, \mathbb{x}, \mathbb{w}, \vec{\mathsf{cm}}, \vec{\mathsf{m}}, \vec{r})\right) = \mathtt{accept}\right] \geq 1 - \mathsf{negl}(\kappa)$$

The probability is taken over the choice of H *and prover/verifier randomness.*

Definition 7 (Random Oracle Reprogramming)**.** *Let* H *be a random oracle modeled as a lazily-sampled table. The reprogramming oracle* RePro*, on input* (a, b)*, sets* $H(a) := b$*, overwriting any previously-defined value. All subsequent queries to* H *(by any party) return the reprogrammed value.*

Definition 8 (Zero-Knowledge)**.** *A CP-NIZKAoK* $(\mathsf{Prove}, \mathsf{Verify})$ *(Def. 6) is* ***zero-knowledge*** *if a PPT simulator* Sim*, given only* $(\mathsf{pp}, \mathbb{x}, \vec{\mathsf{cm}})$ *and* RePro*-access to* H *(Def. 7), produces proofs indistinguishable from real ones: for every PPT* $\mathcal{A}$ *and every* $((\mathbb{x}, \vec{\mathsf{cm}}), (\mathbb{w}, \vec{\mathsf{m}}, \vec{r})) \in \mathcal{R}_{\mathsf{cp}}$*:*

$$\left|\Pr\left[\mathcal{A}^H\left(\mathsf{Prove}^H(\mathsf{pp}, \mathbb{x}, \mathbb{w}, \vec{\mathsf{cm}}, \vec{\mathsf{m}}, \vec{r})\right) = 1\right] - \Pr\left[\mathcal{A}^H\left(\mathsf{Sim}^{H, \mathsf{RePro}}(\mathsf{pp}, \mathbb{x}, \vec{\mathsf{cm}})\right) = 1\right]\right| \leq \mathsf{negl}(\kappa)$$

Here $\mathcal{A}$ *sees the* H *as reprogrammed by* Sim *in the right-hand experiment.*

Definition 9 (Simulation Extractability)**.** *A CP-NIZKAoK* $(\mathsf{Prove}, \mathsf{Verify})$ *(Def. 6) with ZK simulator* Sim *(Def. 8) is* ***simulation extractable*** *if a PPT extractor* $\mathcal{E}$*, given black-box rewinding access to* $\mathcal{A}$ *(including control over* H *and the simulation oracle), recovers a valid witness from any forgery* $\mathcal{A}$ *produces after seeing simulated proofs on statements of its choice. Formally, for every PPT* $\mathcal{A}$*: sample* $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\kappa)$*, let* H *be a random oracle, and give* $\mathcal{A}$ *oracle access to* H *and to a simulation oracle*

Sim' *that on query* $(\mathbb{x}_i, \vec{\mathsf{cm}}_i)$ *returns* $\sigma_i \leftarrow \mathsf{Sim}^{H,\mathsf{RePro}}(\mathsf{pp}, \mathbb{x}_i, \vec{\mathsf{cm}}_i)$ (*so* Sim *may reprogram* H). *Let* $Q_{\mathsf{Sim}} = \{(\mathbb{x}_i, \vec{\mathsf{cm}}_i, \sigma_i)\}$ *collect the simulated triples;* $\mathcal{A}$ *then outputs a forgery* $(\mathbb{x}^*, \vec{\mathsf{cm}}^*, \sigma^*)$. *We require:*

$$\Pr\begin{bmatrix}\mathsf{Verify}^H(\mathsf{pp}, \mathbb{x}^*, \vec{\mathsf{cm}}^*, \sigma^*) = \mathtt{accept} \\ \wedge (\mathbb{x}^*, \vec{\mathsf{cm}}^*, \sigma^*) \notin Q_{\mathsf{Sim}} \\ \wedge \mathcal{E}^{\mathcal{A}}(\mathsf{pp}, \mathbb{x}^*, \vec{\mathsf{cm}}^*, \sigma^*) = \bot\end{bmatrix} \leq \mathsf{negl}(\kappa)$$

On a non-$\bot$ *output,* $\mathcal{E}$ *produces* $(\mathbb{w}, \vec{\mathsf{m}}, \vec{r})$ *with* $((\mathbb{x}^*, \vec{\mathsf{cm}}^*), (\mathbb{w}, \vec{\mathsf{m}}, \vec{r})) \in \mathcal{R}_{\mathsf{cp}}$.

Observation 4 (Simulation Extractability (Def. 9) Implies Non-Malleability)**.** *An adversary who observes a valid proof* σ *for* $(\mathbb{x}, \vec{\mathsf{cm}})$ *cannot produce a fresh accepting proof* $\sigma' \neq \sigma$ *for the same or a related statement: if a simulated proof (indistinguishable from a real proof) could be mauled to produce a new proof string, then the extractor must recover a witness for the relation.*

Theorem 9 ([DG23, GOP+22])**.** *Bulletproofs (viewed as a CP-NIZKAoK with Pedersen commitments) is simulation extractable (Def. 9) in the random oracle model under the discrete logarithm assumption.*

The simulation extractability results of [DG23, GOP+22] are specific to the standard Bulletproofs protocol; they do not immediately extend to Bulletproofs++ [EKRN24] or the Generalized Bulletproofs construction [Fei24]. However, both constructions likely have $(k_1, \ldots, k_{\log})$-special soundness and unique response. Dao and Grubbs [DG23] show that multi-round Fiat–Shamir compiled protocols satisfying computational special soundness and weak unique response are simulation extractable in the ROM; hence it is highly likely that Generalized Bulletproofs is also simulation extractable, though a formal proof has not been written.

Construction 4 (Composing a 3-round IP with a CP-NIZKAoK)**.** Let $\Pi = (\mathsf{P}, \mathsf{V})$ be a 3-round public-coin IP knowledge-sound for a relation $\mathcal{R}$ and complete for an **honest** sub-relation $\mathcal{R}^{\mathsf{hon}} \subseteq \mathcal{R}$, with transcript (π_1, c, π_3). Let $(\mathsf{Setup}, \mathsf{Com})$ be a commitment scheme and let $(\mathsf{Prove}, \mathsf{Verify})$ be a CP-NIZKAoK w.r.t. Com for the **committed verifier relation**:

$$\mathcal{R}_{\mathsf{V}} \coloneqq \left\{ \Big(\underbrace{c}_{\mathbb{x}}, \underbrace{\pi_3}_{\mathbb{w}}, \underbrace{(\tilde{\mathbb{x}}, \pi_1)}_{\vec{\mathsf{m}}} \Big) \,\middle|\, \mathsf{V}(\tilde{\mathbb{x}}, (\pi_1, c, \pi_3)) = \mathtt{accept} \right\}$$

Matching the triple $(\mathbb{x}, \mathbb{w}, \vec{\mathsf{m}})$ of Def. 6: the CP-NIZKAoK's public statement is the challenge c, its non-committed witness is the IP's third message π_3, and the committed message is the pair $(\tilde{\mathbb{x}}, \pi_1)$ consisting of the IP statement $\tilde{\mathbb{x}}$ and first-round message π_1. The committed witness (opening of cm) is therefore $(\tilde{\mathbb{x}}, \pi_1, r)$. The composed NIZKAoK $\Pi' = (\mathsf{P}', \mathsf{V}')$ defined in Fig. 3 is a NIZKAoK in the ROM, knowledge-sound for the **committed relation**:

$$\mathcal{R}^{\mathsf{cm}} \coloneqq \{(\mathsf{cm}, (\mathbb{x}, \pi_1, \mathbb{w}, r)) | \mathsf{cm} = \mathsf{Com}(\mathsf{pp}, (\mathbb{x}, \pi_1); r) \wedge (\mathbb{x}, \mathbb{w}) \in \mathcal{R}\}$$

And complete for the **committed honest relation** $\mathcal{R}^{\mathsf{cm,hon}} \subseteq \mathcal{R}^{\mathsf{cm}}$:

$$\mathcal{R}^{\mathsf{cm,hon}} \coloneqq \left\{ (\mathsf{cm}, (\mathbb{x}, \pi_1, \mathbb{w}, r)) \,\middle|\, \mathsf{cm} = \mathsf{Com}(\mathsf{pp}, (\mathbb{x}, \pi_1); r) \wedge (\mathbb{x}, \mathbb{w}) \in \mathcal{R}^{\mathsf{hon}} \right\}$$

Theorem 10. *Let* $\Pi = (\mathsf{P}, \mathsf{V})$ *be a 3-round public-coin IP knowledge-sound for* $\mathcal{R}$ *with knowledge error* ϵ_{s} *and* ***straight-line*** *extractor* $\mathcal{E}_\Pi$. *Let* Π *be complete for* $\mathcal{R}^{\mathsf{hon}} \subseteq \mathcal{R}$ *with completeness error* ϵ_{c}. *Let* $(\mathsf{Setup}, \mathsf{Com})$ *be a perfectly hiding commitment scheme whose commitment output is uniform over its range. Let* $(\mathsf{Prove}, \mathsf{Verify})$ *be a CP-NIZKAoK for* $\mathcal{R}_{\mathsf{V}}$ *w.r.t.* Com *that is zero-knowledge (Def. 8) and simulation extractable (Def. 9) with extraction advantage* ϵ_{se}. *Then the composed protocol* Π' *(Construction 4) is a NIZKAoK in the ROM, knowledge-sound for* $\mathcal{R}^{\mathsf{cm}}$ *with knowledge error:*

$$Q \cdot \epsilon_{\mathsf{s}} + \epsilon_{\mathsf{se}}$$

$\mathsf{P}'(\mathsf{pp}, \mathbb{x}, \mathbb{w})$

// First-Round Message
$\rho \leftarrow\$ \{0,1\}^*$
$\pi_1 \leftarrow \mathsf{P}_1(\mathbb{x}, \mathbb{w}; \rho)$
// Commit and Derive Challenge
$r \leftarrow\$ \{0,1\}^*$
$\mathsf{cm} \leftarrow \mathsf{Com}(\mathsf{pp}, (\mathbb{x}, \pi_1); r)$
$c \leftarrow H(\mathsf{cm})$
// Third-Round Message
$\pi_3 \leftarrow \mathsf{P}_3(\mathbb{x}, \mathbb{w}, c; \rho)$
// CP-NIZKAoK Proof
$\sigma \leftarrow \mathsf{Prove}^H(\mathsf{pp}, c, \pi_3, \mathsf{cm}, (\mathbb{x}, \pi_1), r)$
return (cm, σ)

$\mathsf{V}'(\mathsf{pp}, \mathsf{cm}, \sigma)$

// Derive Challenge
$c \leftarrow H(\mathsf{cm})$
// Verify CP-NIZKAoK Proof
return $\mathsf{Verify}^H(\mathsf{pp}, c, \mathsf{cm}, \sigma)$

Figure 3: Composed NIZKAoK from a 3-round public-coin IP and a CP-NIZKAoK.

Where Q is the number of queries that the malicious prover $\mathcal{A}$ makes to H. The $Q \cdot \epsilon_{\mathsf{s}}$ term is a union bound over $\mathcal{A}$'s Q candidate first-round messages. The composition is complete for $\mathcal{R}^{\mathsf{cm},\mathsf{hon}}$ with error $\epsilon_{\mathsf{c}} + \mathsf{negl}(\kappa)$ (assuming the CP-NIZKAoK is complete), and zero-knowledge w.r.t. real proofs of $\mathcal{R}^{\mathsf{cm},\mathsf{hon}}$ with simulator-distance at most $\epsilon_{\mathsf{c}} + \delta_{\mathsf{zk}} + Q/|\mathcal{C}_{\mathsf{cm}}|$, where δ_{zk} is the maximum distinguishing advantage against the CP-NIZKAoK simulator (Def. 8) and $|\mathcal{C}_{\mathsf{cm}}|$ is the size of the commitment range.

Proof of Simulation Extractability. We construct a simulator and extractor for Π'.

Simulator. The simulator Sim' for Π', on input pp, works as follows: sample $\mathsf{cm} \leftarrow \mathsf{Com}(\mathsf{pp}, \vec{0}; r)$ for fresh r (a commitment to a zero message of the appropriate length, covering both the first-round message and IP-statement slots), sample a fresh challenge c, program $\mathsf{RePro}(\mathsf{cm}, c)$ so that subsequent queries return $H(\mathsf{cm}) = c$, and invoke the CP-NIZKAoK simulator $\sigma \leftarrow \mathsf{Sim}^{H,\mathsf{RePro}}(\mathsf{pp}, c, \mathsf{cm})$, which itself may further reprogram H. Output (cm, σ).

Simulator Distance. Fix $(\mathsf{cm}, (\mathbb{x}, \pi_1, \mathbb{w}, r)) \in \mathcal{R}^{\mathsf{cm},\mathsf{hon}}$ used to generate the distribution of real proofs. We bound the statistical/computational distance between the real prover output and $\mathsf{Sim}'(\mathsf{pp})$ via a sequence of hybrids. Let δ_{zk} denote the maximum distinguishing advantage against the CP-NIZKAoK simulator (Def. 8).

- $\mathcal{H}_0$: the real execution of P' (Fig. 3). Namely, $\pi_1 \leftarrow \mathsf{P}_1(\mathbb{x}, \mathbb{w}; \rho)$, $\mathsf{cm} \leftarrow \mathsf{Com}(\mathsf{pp}, (\mathbb{x}, \pi_1); r)$, $c \leftarrow H(\mathsf{cm})$, $\pi_3 \leftarrow \mathsf{P}_3(\mathbb{x}, \mathbb{w}, c; \rho)$, $\sigma \leftarrow \mathsf{Prove}^H(\mathsf{pp}, c, \pi_3, \mathsf{cm}, (\mathbb{x}, \pi_1), r)$; output (cm, σ).
- $\mathcal{H}_1$: as $\mathcal{H}_0$, but replace Prove with the CP-NIZKAoK simulator: $\sigma \leftarrow \mathsf{Sim}^{H,\mathsf{RePro}}(\mathsf{pp}, c, \mathsf{cm})$. Let E be the event that the IP transcript (π_1, c, π_3) is accepting under V; since $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}^{\mathsf{hon}}$, $\Pr[\neg E] \leq \epsilon_{\mathsf{c}}$. Conditioned on E, the tuple $((c, \mathsf{cm}), (\pi_3, (\mathbb{x}, \pi_1), r)) \in \mathcal{R}_{\mathsf{cp}}$, so zero-knowledge of the CP-NIZKAoK (Def. 8) bounds the conditional advantage by δ_{zk}. Hence the statistical distance between $\mathcal{H}_0$ and $\mathcal{H}_1$ is at most $\epsilon_{\mathsf{c}} + \delta_{\mathsf{zk}}$.
- $\mathcal{H}_2$: as $\mathcal{H}_1$, but commit to zero: $\mathsf{cm} \leftarrow \mathsf{Com}(\mathsf{pp}, \vec{0}; r)$. By perfect hiding (Def. 4), $\mathsf{Com}(\mathsf{pp}, (\mathbb{x}, \pi_1); r)$ and $\mathsf{Com}(\mathsf{pp}, \vec{0}; r)$ are identically distributed over r; hence the joint distribution of (cm, σ) is unchanged, and $\mathcal{H}_1$ and $\mathcal{H}_2$ are identically distributed.
- $\mathcal{H}_3 = \mathsf{Sim}'(\mathsf{pp})$: as $\mathcal{H}_2$, but instead of $c \leftarrow H(\mathsf{cm})$, sample c fresh and program $\mathsf{RePro}(\mathsf{cm}, c)$. Since cm is a freshly sampled commitment independent of $\mathcal{A}$'s view, $H(\mathsf{cm})$ is uniformly random and unqueried except when cm collides with a prior H query; over uniform cm this occurs with probability at most $Q/|\mathcal{C}_{\mathsf{cm}}|$, where Q bounds the total number of H queries (by the distinguisher and the CP-NIZKAoK simulator) and $|\mathcal{C}_{\mathsf{cm}}|$ is the size of

the commitment range. Conditioned on no collision, the two ways of obtaining c induce the same distribution on (c, H), so the statistical distance between $\mathcal{H}_2$ and $\mathcal{H}_3$ is at most $Q/|\mathcal{C}_{\mathsf{cm}}|$.

The total distance between the real view and $\mathsf{Sim}'(\mathsf{pp})$ is therefore at most $\epsilon_{\mathsf{c}} + \delta_{\mathsf{zk}} + Q/|\mathcal{C}_{\mathsf{cm}}|$.

Extractor. Suppose a PPT adversary $\mathcal{A}$ making at most Q queries to H, with oracle access to Sim' and H, outputs a fresh accepting proof $(\mathsf{cm}^*, \sigma^*)$ for Π', where $c^* = H(\mathsf{cm}^*)$, with probability $\epsilon_{\mathcal{A}}$. Since Sim' is implemented using the CP-NIZKAoK simulator Sim (which has RePro-access to H), $\mathcal{A}$ is a valid adversary in the CP-NIZKAoK simulation extractability game (Def. 9). Run $\mathcal{A}$ once to obtain an accepting $(\mathsf{cm}^*, \sigma^*)$ and invoke the CP-NIZKAoK extractor on σ^* to obtain $(\mathbb{x}^*, \pi_1^*, \pi_3^*, r^*)$ satisfying:

$$\mathsf{cm}^* = \mathsf{Com}(\mathsf{pp}, (\mathbb{x}^*, \pi_1^*); r^*) \quad \text{and} \quad \mathsf{V}(\mathbb{x}^*, (\pi_1^*, c^*, \pi_3^*)) = \mathtt{accept}$$

Except with probability ϵ_{se}. Feed the single accepting transcript (π_1^*, c^*, π_3^*) into the straight-line IP extractor $\mathcal{E}_\Pi$ to obtain $\mathbb{w}^*$ with $(\mathbb{x}^*, \mathbb{w}^*) \in \mathcal{R}$, except with the IP knowledge error. Union-bounding over $\mathcal{A}$'s Q candidate commitments cm_j (each fixing a first-round message via its CP-NIZKAoK opening), the extractor outputs $(\mathbb{x}^*, \pi_1^*, \mathbb{w}^*, r^*)$ as the witness for $\mathcal{R}^{\mathsf{cm}}$, with knowledge error $Q \cdot \epsilon_{\mathsf{s}} + \epsilon_{\mathsf{se}}$. □

The factor Q is inherent to any Fiat–Shamir compilation: $\mathcal{A}$ can try any of its Q candidate first-round messages. This is asymptotically tight: there is a concrete adversary matching it, which simply queries the oracle Q times on Q commitments, hoping to trigger one of the bad events in the original MA protocol.

The interactive proof Π need not be zero-knowledge or non-malleable; in fact, in our application Π is neither. This is because the outer NIZK and commitment provide zero-knowledge, by hiding the transcript and providing a simulation trapdoor. The commitment scheme also does not need to be non-malleable; in the concrete instantiation with Pedersen commitments (Construction 3), it is not. Intuitively this is because the commitment is made non-malleable by the fact that it is provided along with a simulation extractable NIZK proving knowledge of its opening.

> ***Comments On Prior Work.***
>
> Prior works have reflected upon the "mauling" of valid proofs/witnesses for the MA/IP of Eagen/Bassa [Eag22] [Bas25]. We observe that such reflections are not needed to conclude anything about the non-malleability (or simulation extractability) of the resulting composition as shown above: the outer NIZK (e.g. Bulletproofs) eliminates any malleability of the inner proof/commitments.

Construction 5 (NIZKAoK for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$)**.** We construct a NIZKAoK knowledge-sound for the discrete log relation $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ and complete for $\mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}] \subseteq \mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$ (Section 2), by instantiating Construction 4 with $\mathcal{R} := \mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$, $\mathcal{R}^{\mathsf{hon}} := \mathcal{R}^{\mathsf{dlog\text{-}honest}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}}[\vec{B}]$, the three-round IP $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (Fig. 2), Pedersen commitments over $\mathbb{G}_{\mathsf{out}}$ (Construction 3), and Bulletproofs as the CP-NIZKAoK. Motivated by the notes of Parker [Par24b] we fix $\mathcal{S} = \mathcal{S}_{\mathsf{Parker}}$ for concreteness. Recall that the first-round message is $\pi_1 = (\vec{m},\ a(x),\ b(x)) \in \mathbb{F}_q^k \times \mathbb{F}_q[x]^2$, the verifier challenge is $c = (A_0, A_1) \in E(\mathbb{F}_q)^2$, and the third-round message is $\pi_3 = (h_0, h_1, h_2, g) \in \mathbb{F}_q^4$. The degree bound $\deg_{\mathrm{E}}(D) \le \mathsf{d}$ fixes the maximum degrees of $a(x)$ and $b(x)$; their coefficient vectors are zero-padded to this fixed length so that the Pedersen commitment always commits to a vector in $\mathbb{F}_q^n$ for a fixed n determined by d and k. The concrete committed verifier relation $\mathcal{R}_{\mathsf{V}}$ (Construction 4) is obtained by expanding the checks of V in $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (Fig. 2). Writing $D(x, y) = a(x) - b(x)y$, $A_2 = -(A_0 + A_1)$, and L for the line through A_0, A_1, A_2; observe that since the A_i are derived from the (public) challenge and the B_j are protocol parameters, the verifier

can compute $\alpha_i = \frac{dx(A_i)}{dz}$ and $\beta_j = \frac{-1}{L(B_j)}$ outside the NIZK. The relation $\mathcal{R}_{\mathsf{V}}$ is parameterized by these quantities and the Bulletproofs circuit only needs to check the linear and multiplicative relations over the committed values, including the IP statement $P = (x(P), y(P)) \in \mathbb{F}_q^2$, which is committed alongside the first-round message:

$$\mathcal{R}_{\mathsf{V}} = \left\{ \left(\underbrace{(\vec{\alpha}, \vec{\beta}, c)}_{\mathbb{x}}, \underbrace{(h_0, h_1, h_2, g)}_{\mathbb{w}}, \underbrace{(P, \vec{m}, a(x), b(x))}_{\vec{\mathsf{m}}} \right) \;\middle|\; \begin{array}{l} y(P)^2 = x(P)^3 + A \cdot x(P) + B \\ \sum_{i=0}^{2} h_i \cdot \alpha_i = g + \sum_{j=1}^{k} \beta_j \cdot m_j \\ h_i \cdot D(A_i) = D'(A_i) \quad \text{for } i = 0, 1, 2 \\ g \cdot L(-P) = -1 \\ a_1 = 1 \end{array} \right\}$$

The composed protocol is shown in Fig. 4.

Observation 5. *It is important that an implementation absorb the out-of-circuit quantities $\vec{\alpha}$ and $\vec{\beta}$ into the Bulletproofs transcript for provable soundness: they define part of the relation.*

With $\mathcal{S} = \mathcal{S}_{\mathsf{Parker}}$, the committed verifier relation $\mathcal{R}_{\mathsf{V}}$ of Construction 5 coincides, up to notation, with the "$\texttt{DiscreteLog}_G$" gadget of Parker [Par24b, §5.1.4]. The cost of proving an elliptic curve inner product with a fixed basis is 7 *multiplicative constraints*, broken down in Fig. 5. Additionally the prover must commit to $\mathsf{d} + k + 2$ field elements, which for a 256-bit curve/field with $k = 256$ and $\mathsf{d} = k + 1 = 257$ as in the binary decomposition of Example 2, becomes 515.

References

[AFK23] Thomas Attema, Serge Fehr, and Michael Klooß. Fiat-Shamir transformation of multi-round interactive proofs (extended version). *Journal of Cryptology*, 36(4):36, October 2023.

[Bas24a] Alp Bassa. On the use of logarithmic derivatives in Eagen's proof of sums of points. Veridise Technical Report, 2024. Link.

[Bas24b] Alp Bassa. Soundness proof for Eagen's proof of sums of points. Veridise Technical Report, 2024. Link.

[Bas25] Alp Bassa. Soundness proof for an interactive protocol for the discrete logarithm relation. Veridise Technical Report, 2025. Link.

[BG93] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 390–420. Springer, Berlin, Heidelberg, August 1993.

[DG23] Quang Dao and Paul Grubbs. Spartan and bulletproofs are simulation-extractable (for free!). In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 531–562. Springer, Cham, April 2023.

[DKL14] Zeev Dvir, János Kollár, and Shachar Lovett. Variety evasive sets. *Computational Complexity*, 23(4):509–529, 2014.

[Eag22] Liam Eagen. Zero knowledge proofs of elliptic curve inner products from principal divisors and weil reciprocity. Cryptology ePrint Archive, Report 2022/596, 2022.

[EKRN24] Liam Eagen, Sanket Kanjalkar, Tim Ruffing, and Jonas Nick. Bulletproofs++: Next generation confidential transactions via reciprocal set membership arguments. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part V*, volume 14655 of *LNCS*, pages 249–279. Springer, Cham, May 2024.

$\mathsf{P}'(\mathsf{pp}, \mathbb{x} = P, \mathbb{w} = \vec{n})$

// First-Round Message

Let $D(x, y) = a(x) - b(x)y \in \mathbb{F}_q[x, y]$ st.

$(D(x, y))_0 = (-P) + \sum n_i \cdot (B_i)$

$\vec{m} \leftarrow \vec{n} \bmod q \in \mathbb{F}_q^k$

Let $\vec{\mathsf{m}} = (x(P), y(P), \vec{m}, \text{coeffs}(a), \text{coeffs}(b)) \in \mathbb{F}_q^{n+2}$

// Commit and Derive Challenge

$r \leftarrow\!\$\, \mathbb{F}_q$

$\mathsf{cm} \leftarrow [r] \cdot H + \sum_{i=1}^{n+2} [\mathsf{m}_i] \cdot G_i \in \mathbb{G}_{\mathsf{out}}$

$(A_0, A_1) \leftarrow H(\mathsf{cm})$

// Third-Round Message

$A_2 \leftarrow -(A_0 + A_1)$

Let L be the line through A_0, A_1, A_2

$h_0 \leftarrow D'(A_0)/D(A_0)$

$h_1 \leftarrow D'(A_1)/D(A_1)$

$h_2 \leftarrow D'(A_2)/D(A_2)$

$g \leftarrow -1/L(-P)$

// Compute Out-of-Circuit Quantities

$\alpha_i \leftarrow \frac{dx(A_i)}{dz}$ for $i = 0, 1, 2$

$\beta_j \leftarrow \frac{-1}{L(B_j)}$ for $j = 1, \ldots, k$

// Bulletproofs Proof

$\sigma \leftarrow \mathsf{BP.Prove}^H($

$\mathsf{pp},$

$\mathbb{x}' = (\vec{\alpha}, \vec{\beta}, (A_0, A_1)),$

$\mathbb{w}' = (h_0, h_1, h_2, g),$

$\mathsf{cm}, \vec{\mathsf{m}}, r$

$)$

return (cm, σ)

$\mathsf{V}'(\mathsf{pp}, \mathsf{cm}, \sigma)$

// Derive Challenge

$(A_0, A_1) \leftarrow H(\mathsf{cm})$

$A_2 \leftarrow -(A_0 + A_1)$

Let L be the line through A_0, A_1, A_2

// Compute Out-of-Circuit Quantities

$\alpha_i \leftarrow \frac{dx(A_i)}{dz}$ for $i = 0, 1, 2$

$\beta_j \leftarrow \frac{-1}{L(B_j)}$ for $j = 1, \ldots, k$

// Verify Bulletproofs Proof

return $\mathsf{BP.Verify}^H($

$\mathsf{pp}, (\vec{\alpha}, \vec{\beta}, (A_0, A_1)), \mathsf{cm}, \sigma$

$)$

Figure 4: Concrete NIZKAoK for $\mathcal{R}^{\mathsf{dlog}}_{\mathbb{F}_q, E(\mathbb{F}_q)}[\vec{B}]$: composition of $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$ (Fig. 2) with Bulletproofs and Pedersen commitments (Construction 3). The statement P is committed alongside the first-round message, so the composed verifier V' takes only (cm, σ) as input. The proof σ establishes that the committed IP statement and first-round message, together with the helpers (h_0, h_1, h_2, g), satisfy all verifier checks of $\Pi^{\mathsf{IP}}_{\mathbb{F}_q, E(\mathbb{F}_q), \mathsf{d}, \mathcal{S}, \vec{B}}$.

Constraint	Mul. Count
$y(P)^2 = x(P)^3 + A \cdot x(P) + B$	3
$h_i \cdot D(A_i) = D'(A_i)$ for $i = 0, 1, 2$	3
$g \cdot L(-P) = -1$	1
$a_1 = 1$	0
$\sum_{i=0}^{2} h_i \cdot \alpha_i = g + \sum_{j=1}^{k} \beta_j \cdot m_j$	0
Total	**7**

Figure 5: Breakdown of the R1CS constraint costs of Construction 5.

[EOT10] Jordan S. Ellenberg, Richard Oberlin, and Terence Tao. The Kakeya set and maximal conjectures for algebraic varieties over finite fields. *Mathematika*, 56(1):1–25, 2010.

[Fei24] Aaron Feickert. Generalized bulletproofs. Technical Report, 2024. Link.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.

[GOP+21] Chaya Ganesh, Claudio Orlandi, Mahak Pancholi, Akira Takahashi, and Daniel Tschudi. Fiat–shamir bulletproofs are non-malleable (in the algebraic group model). Cryptology ePrint Archive, Report 2021/1393, 2021.

[GOP+22] Chaya Ganesh, Claudio Orlandi, Mahak Pancholi, Akira Takahashi, and Daniel Tschudi. Fiat-Shamir bulletproofs are non-malleable (in the algebraic group model). In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 397–426. Springer, Cham, May / June 2022.

[GSSS25a] Brandon Goodell, Rigo Salazar, Freeman Slaughter, and Luke Szramowski. A further review of the DL gadget of interest. Cypher Stack Technical Report, 2025. Link.

[GSSS25b] Brandon Goodell, Rigo Salazar, Freeman Slaughter, and Luke Szramowski. A review of the Veridise discrete logarithm relation gadget. Cypher Stack Technical Report, 2025. Link.

[Maa04] Martijn Maas. Pairing-based cryptography. Master's thesis, Technische Universiteit Eindhoven, January 2004.

[Par24a] Luke "Kayaba" Parker. FCMP++. manuscript, May 2024. May 8, 2024. Link.

[Par24b] Luke "Kayaba" Parker. A R1CS gadget for a 2^k-bit scaling of a fixed generator in 7 multiplicative constraints. manuscript, May 2024. May 9, 2024. Link.

[Sil86] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1986.