Improved GGM Witness-Encryption Bypassing Rank Gap

Charanjit S. Jutla
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

Abstract

Hair and Sahai construct witness encryption in the classical prime-order generic-group model using a designed matrix space with a growing rank gap. We retain their encryption mechanism while replacing the minimum-rank requirement with a lower bound on the min-entropy of Mr for nonzero matrices M that do not permit witness extraction, where r is the short random vector used in encryption. A matrix with substantial concentration has a column factorization from a finite family of short rational patterns. One list of linear forms evaluated at powers of a public random seed preserves all these patterns simultaneously, except with exponentially small probability, and therefore suffices in place of their multiple folded lists.

For a quadratic verifier encoding with N Boolean variables, using $t = \lceil \log N \rceil$ forms gives one-bit ciphertexts of $\widetilde{O}(N^5)$ bits and optimized decryption using $\widetilde{O}(N^4)$ group queries. For sufficiently large N, distinguishing advantage on false instances is at most $N^{-\log N}$ against classical generic adversaries making at most $N^{\log N}$ group queries. Correctness error is $2^{-\Omega(N)}$, simultaneously for all valid witnesses of a ciphertext. At these same security and correctness guarantees, a beneficial reparameterization of Hair–Sahai's deterministic construction gives ciphertexts of $\widetilde{O}(N^6)$ bits and, with the same decryption optimizations, $\widetilde{O}(N^5)$ group queries.

Keywords: Witness encryption; generic-group model; anti-concentration; rank condensers; Boolean moments; Nullstellensatz.

Contents

1 Introduction **3**
1.1 Weighted tables and their moment representation 4
1.2 Why concentration makes one list sufficient 5
1.3 Result and relation to prior work 6

2 Witness encryption and the generic-group model **7**
2.1 Fresh prime-order generic groups 8
2.2 Witness-encryption definition 8

3 Weighted tables and their moments **9**
3.1 Quadratic tests and weighted assignment tables 9
3.2 One list evaluated at powers of the seed 10
3.3 Available moments and their identities 11
3.4 Efficient public basis computation 12

4 Preserved columns give a representation on satisfying assignments **12**
4.1 The representation theorem 12
4.2 Finding a witness without computing the representation 17

5 Concentration and a simultaneous good seed **18**
5.1 Collision probability beyond rank . . . 18
5.2 A determinant for evaluations at powers of the seed . . . 20
5.3 Simultaneous preservation of short-coordinate patterns . . . 20

6 Witness extraction or small collision probability **21**

7 Encryption, decryption, and correctness **22**
7.1 The algorithms . . . 22
7.2 Simultaneous correctness . . . 23

8 Security without a minimum-rank bound **23**
8.1 Longer messages . . . 24

9 From affine-source extraction to witness encryption **25**

10 Efficiency and comparison with Hair–Sahai **27**
10.1 Ciphertext size and decryption queries . . . 27
10.2 What changes in the deterministic construction . . . 28
10.3 A smaller field for the deterministic baseline . . . 29
10.4 A matched comparison at stronger security . . . 30
10.5 High-order folding bases and smaller fields . . . 30
10.6 The role of the weight family . . . 32

11 Acknowledgment **32**

A Computing a basis without enumerating assignments **34**
A.1 Extend, reduce, retain . . . 34
A.2 Why every common extension preserves dependencies . . . 35
A.3 Imposing the verifier tests . . . 35
A.4 The polynomial-evaluation viewpoint . . . 36

B Compressed moments, representation, and Nullstellensatz **36**
B.1 Why the available degrees suffice . . . 37
B.2 Changing the compressed coordinates . . . 37
B.3 Why the original coefficients need not have valid support . . . 38
B.4 How weighted verifier equations constrain support . . . 38
B.5 The dual Nullstellensatz viewpoint . . . 39

C Bounded sampling **39**

D The generic-group security proof **40**
D.1 A common symbolic execution . . . 40
D.2 Bounding every symbolic collision . . . 41

E Parameters for the stronger security comparison **42**
E.1 Algebra and the public seed . . . 42
E.2 Sampling and correctness . . . 43
E.3 Security for quasipolynomially many queries . . . 43
E.4 Costs . . . 44

E.5 Smaller fields with per-witness correctness 44

F Efficient witness decryption using systematic coordinates 45
F.1 The cost of direct basis reconstruction 45
F.2 Read the certificate at the pivot positions 46
F.3 Group equal weights before exponentiating 47
F.4 The bounded search by baby-step/giant-step 47
F.5 Optimized complexity and comparison 47

G One-variable rank-preserving maps 49
G.1 Proof of the power-evaluation lemma 49
G.2 A criterion and maps with smaller determinant degree 50

1 Introduction

Witness encryption associates a ciphertext with a public NP statement: a witness decrypts it, while ciphertexts for false statements hide the message [1]. Hair and Sahai construct witness encryption in the classical prime-order generic-group model [2]. They turn a verifier into a designed public matrix space. A satisfying assignment supplies a rank-one matrix with a Boolean right factor, which enables bit decryption by a short search. Their false-instance security proof uses a growing minimum rank for nonzero matrices in that space.

We retain this encryption mechanism and simplify the weighted matrix construction. The probability of a generic-group equality depends on Mr, where M is a linear combination of public basis matrices and r is the secret vector of small integer coordinates used in encryption. Rank bounds that probability, but is not its exact measure. Controlling the concentration of Mr lets us use one random list of weights in place of Hair–Sahai's many deterministic lists. Our algebraic explanation reformulates their Boolean interpolation argument using a symmetric moment tensor. Column span preservation makes tensor compression reversible; Boolean moment identities and weighted verifier tests then represent the entire table by compatible satisfying assignments. We then identify which (weight) tables need this column span preservation guarantee.

Our argument also connects witness encryption to min-entropy and rank condensation. For a uniform affine source, preserving the dimension of its direction space preserves its min-entropy. We use the rank-preserving map underlying Gabizon–Raz's deterministic affine-source extractor [8] to preserve simultaneously the short coordinate patterns forced by insufficient min-entropy; the Boolean moment constraints then yield a witness. On a good public seed, false instances therefore force every nonzero matrix image Mr to have enough min-entropy to make generic-group collisions unlikely. The distribution Mr is an image of a short integer grid and need not itself be an affine source. Section 9 makes the connection precise: Theorem 9.1 converts an explicit linear extractor for affine sources with short bases into witness encryption, with ciphertext size determined by the extractor's output dimension and field size.

For a quadratic verifier encoding with N Boolean variables, using $t = \lceil \log N \rceil$ forms gives one-bit ciphertexts of $\widetilde{O}(N^5)$ bits, compared with $\widetilde{O}(N^6)$ for the beneficially reparameterized Hair–Sahai construction. Applying the same decryption optimizations to both schemes, honest decryption uses $\widetilde{O}(N^2)$ field operations and $\widetilde{O}(N^4)$ group queries for our construction, compared with $\widetilde{O}(N^3)$ field operations and $\widetilde{O}(N^5)$ group queries for Hair–Sahai. The corresponding local bit-work bounds are $\widetilde{O}(N^5)$ and $\widetilde{O}(N^6)$, including reading the ciphertext; neither optimized decryptor requires a prepublished matrix basis. These comparisons match distinguishing advantage $N^{-\log N}$ against $N^{\log N}$ generic-group queries and correctness error $2^{-\Omega(N)}$, simultaneously for all valid witnesses of

a ciphertext. Here $\widetilde{O}$ suppresses polylogarithmic factors. Section 10.4 gives the size comparison, and Appendix F gives the optimized decryption algorithms and their costs.

1.1 Weighted tables and their moment representation

Encode the verifier by quadratic equations $q_1, \dots, q_J$ in N Boolean wire variables. Put $D = N + 1$, $b_0 = 1$, and $v(b) = (1, b_1, \dots, b_N)^\top$. Choose a public $\tau \in \mathbb{F}_p$ and use the t forms

$$\ell_s(b) = \sum_{i=0}^{N} \tau^{is} b_i, \qquad 1 \le s \le t.$$

For every $\alpha \in \mathbb{Z}^t_{\ge 0}$ with $|\alpha| \le t$, let $h_\alpha(b) = \prod_s \ell_s(b)^{\alpha_s}$. Stack the weighted quadratic tables to obtain

$$A_\tau(b) = \big(h_\alpha(b) v(b) v(b)^\top\big)_\alpha = u_\tau(b) v(b)^\top.$$

There are $W = \binom{2t}{t}$ weights and $m = DW$ rows. The space $\mathcal{V}_\tau$ spans these tables over all Boolean assignments. The instance-dependent subspace $\mathcal{S}_{x,\tau}$ imposes, in every weight block, the linear tests corresponding to every q_j. Thus every satisfying assignment supplies a nonzero rank-one matrix in $\mathcal{S}_{x,\tau}$.

The central algebraic statement concerns a matrix $0 \ne M \in \mathcal{S}_{x,\tau}$ and the coefficient matrix $F_t = (\tau^{is})_{0 \le i \le N,\, 1 \le s \le t}$ of the forms. Suppose that the forms preserve the *whole* column space:

$$\operatorname{col}(MF_t) = \operatorname{col}(M).$$

Then $\rho = \operatorname{rank} M \le t$. Theorem 4.1 proves that

$$M = \sum_e \omega_e A_\tau(b_e),$$

with at most 2^ρ nonzero terms, where every b_e is satisfying and $v(b_e)$ lies in the row space of M. The coefficients are field elements; neither positivity nor a normalization of the constant moment is required. This is a representation of the stored table. An arbitrary representation of M in the original Boolean assignment span need not itself have satisfying support.

The proof has three stages.

1. *Column preservation gives reversible tensor compression.* Choose actual independent columns and write $M = PR$, with R the matrix of column coordinates. Preservation lets us represent these selected coordinates by linear combinations of the available forms. Adjoin the constant form and project both endpoint factors to form coordinates, obtaining a symmetric tensor of order $t + 2$. The column relations constrain its last slot; symmetry extends the same constraint to every slot. Thus compression to a moment tensor in ρ coordinates is reversible and recovers the full table.

2. *Boolean identities give interpolation on 2^ρ points.* The identities $b_i^2 = b_i$ transfer to the selected coordinates. Reduction by these identities expresses the moments using the Boolean interpolation weights $\chi_e(y) = \prod_{j:e_j=1} y_j \prod_{j:e_j=0}(1 - y_j)$. Their degrees are at most $\rho \le t$, so the complete weight list provides the required moments. The resulting coefficients ω_e need not sum to one, and their sum can be zero even when the table is nonzero.

3. *Weighted verifier tests restrict the support.* Interpolation isolates each candidate. The tests weighted by its interpolation polynomial force $\omega_e q_j(b_e) = 0$ for every verifier equation. Together with the Boolean and constant-coordinate identities, this leaves only compatible satisfying assignments with nonzero coefficients.

This last stage explains the role of testing the equations at all available weights. Unweighted equations could allow contributions from different candidates to cancel. Weighted interpolation tests isolate those contributions individually.

This argument uses Hair–Sahai's underlying Boolean interpolation [2, Section 5.3]. Their work also identifies the dual connection to bounded-degree Nullstellensatz and polynomial calculus. The tensor formulation makes explicit how column span preservation permits reversible compression, Boolean interpolation gives a representation on a small cube, and verifier equations restrict its support. Since the representation of a nonzero table has a nonzero coefficient, enumerating the at most 2^ρ compatible Boolean candidates finds a witness.

The universal span $\mathcal{V}_\tau$ depends only on N, p, t, τ, so its public basis can be computed once and reused for all instances with those parameters. The verifier tests defining $\mathcal{S}_{x,\tau}$ still depend on x. The algorithms in this paper use a basis of that instance-dependent space; reusing the universal span does not by itself eliminate this further linear algebra. Appendix A explains the finite basis computation without enumerating all Boolean assignments.

1.2 Why concentration makes one list sufficient

For a basis $M_1, \dots, M_k \in \mathbb{F}_p^{m \times D}$, encryption of zero publishes generic-group labels for the exponents

$$s_1, \dots, s_m, \qquad s^\top M_1 r, \dots, s^\top M_k r, \qquad s \leftarrow \mathbb{F}_p^m, \quad r \leftarrow [B]^D,$$

where $[B] = \{0, \dots, B-1\}$ and B is polynomially bounded. Encryption of one replaces the latter k exponents by independent uniform field elements. A generic equality test for bit zero has the form

$$c + s^\top (a + Mr) = 0, \qquad M \in \operatorname{span}\{M_j\}.$$

Unless $Mr = -a$, this is a nonconstant affine equation in uniform s. Its vanishing probability is therefore at most $\kappa_B(M) + p^{-1}$, where

$$\kappa_B(M) = \max_z \Pr_{r \leftarrow [B]^D}[Mr = z], \qquad H_B(M) = -\log \kappa_B(M).$$

The quantity $H_B(M)$ is the min-entropy of Mr. Rank ρ gives $H_B(M) \geq \rho \log B$, as in [2, Lemma 3.4]. Low rank can nevertheless have much higher entropy: the rank-one row $(1, B, \dots, B^{D-1})$ maps $[B]^D$ injectively when $p > B^D$, and hence has entropy $D \log B$. This illustrates the probabilistic distinction; it is not a claim that this row belongs to the designed space.

The opposite problem explains why unweighted quadratic tables do not suffice. If $i, j \geq 1$ are distinct and $b_i b_j$ never occurs in the verifier equations, the symmetric matrix $E_{ij} + E_{ji}$ belongs to the unweighted Boolean table span and passes every verifier test. The corresponding group element for bit zero is $X_i^{r_j} X_j^{r_i}$. Enumerating the B^2 possible pairs finds a match, whereas an independent uniform group element matches with probability at most B^2/p. Linear terms in the verifier may still rule out assignments involving these variables. Thus sparse matrices can reveal short coordinates without supplying a witness. The needed condition is that *every matrix with substantial concentration must yield a witness.*

Hair–Sahai's deterministic construction uses $2Nt + 1$ folded lists. Every nonzero matrix of rank at most t has a list that both preserves its whole column space and contains a nonzero block [2, Sections 5.2–5.3]. The interpolation argument then supplies a witness. Our construction keeps only one list. Every block belongs to that list, so the nonzero-block condition is automatic. We require column preservation only for matrices whose concentration is too large.

If a matrix has $t+1$ columns with no nonzero integer relation whose coefficients have absolute value less than B, then $\kappa_B(M) \le B^{-(t+1)}$. Otherwise a maximal set of such columns has size at most t and expresses all remaining columns using fractions a/b with $|a| < B$ and $1 \le b < B$. This gives a factorization through one of finitely many short rational coordinate patterns. The selected columns need not be independent over the field; the preservation argument accommodates this distinction.

A determinant root bound shows that one random τ, over a sufficiently large field, preserves all these patterns simultaneously. The guarantee applies to every left factor and every instance with the same parameters, including matrices chosen after the seed is published. Combining this event with the moment representation gives

$$\text{a valid witness can be extracted from } M \quad \text{or} \quad \kappa_B(M) \le B^{-(t+1)}.$$

On a false statement, the first alternative is impossible. This is the collision bound needed by the generic-group proof, without requiring a growing minimum rank on false instances.

1.3 Result and relation to prior work

Theorem 1.1 (Main result, informal)**.** *Every NP relation admits witness encryption in the classical prime-order generic-group model with the following parameters. For a padded quadratic verifier encoding with N Boolean variables, choose $t = \lceil \log N \rceil$ and $B = N^3$. The algorithms run in polynomial time, and one-bit ciphertexts have $\widetilde{O}(N^5)$ bits. Optimized decryption uses $\widetilde{O}(N^2)$ field operations and $\widetilde{O}(N^4)$ group queries, with $\widetilde{O}(N^5)$ local bit work and no prepublished matrix basis. Correctness error is $2^{-\Omega(N)}$, simultaneously for all valid witnesses of a ciphertext. For sufficiently large N, distinguishing advantage on false statements is at most $N^{-\log N}$ against every classical generic adversary making at most $N^{\log N}$ group queries.*

These probability bounds average over the encryption coins and fresh group encoding. Correctness holds for every seed apart from the stated sampling and decryption error; security includes an exceptional-seed probability at most 2^{-2N}. Appendix E gives the field choice and derives these guarantees from the quantitative bound in Section 8, whose symbolic-label proof appears in Appendix D. The algebraic witness representation establishes false-instance concentration; the security experiment does not run an extractor.

At the same security and correctness target, the reparameterized deterministic Hair–Sahai construction has one-bit ciphertexts of $\widetilde{O}(N^6)$ bits. With the same decryption optimizations, it uses $\widetilde{O}(N^3)$ field operations and $\widetilde{O}(N^5)$ group queries. Section 10.4 compares these explicit upper bounds using the shared grid $B = N^3$ and separately chosen fields. Appendix F gives the optimized algorithms: the encryptor uses a systematic basis and sends pivot positions, from which the decryptor reads witness coefficients without reconstructing the basis. It also quantifies the field arithmetic and group queries of straightforward decryption.

For a fair comparison, we let the deterministic construction choose its field size independently of its row count. Its matrix conditions and witness-encryption proof permit $\log p_{\mathrm{HS}} = \Theta(N)$ with exponentially small simultaneous correctness error.

The smaller parameter choice used initially in Sections 3–8 sets $t = \lceil \log \log N \rceil$ and $B = m = (N+1)\binom{2t}{t}$. It gives one-bit ciphertexts of $O(N^3(\log N)^4)$ bits and straightforward decryption using $O(N^3(\log N)^4)$ group queries. Simultaneous correctness error is $2^{-\Omega(m)}$, and distinguishing advantage is $N^{-\Omega(\log \log N)}$ for every polynomial query bound. At this weaker threshold, the deterministic baseline has ciphertexts of $\widetilde{O}(N^4)$ bits against our $\widetilde{O}(N^3)$ bits; Section 10 checks the parameters. The $\log \log N$ threshold itself is permitted by [2, Theorem 3.1]; it is not a separate

improvement over their allowable rank bounds. Here t counts forms and bounds weight degree, rather than specifying a minimum nonzero rank.

The correctness target matters to the comparison. If error $N^{-\log N}$ is required only for each fixed witness, a high-order folding base permits smaller fields for the deterministic baseline. At the logarithmic threshold both constructions then have $\widetilde{O}(N^5)$-bit ciphertexts, with a factor of $\log N$ between the explicit size bounds. Section 10.5 explains this alternative and the role of the folding base.

The encryption mechanism, weighted assignment tables, Boolean interpolation, and finite basis algorithm originate in Hair–Sahai. Their paper also explains the relevant closure conditions and the bounded-degree Nullstellensatz viewpoint [2, Sections 5.2–5.4 and 6.2]. Our symmetric-tensor proof makes explicit how the entire table is recovered from compressed moments and how each family of equations constrains its representation. The additional probabilistic argument restricts column preservation to short rational patterns, allowing one random list instead of all deterministic lists. We use the power-evaluation map of Gabizon–Raz [8], whose determinant property requires no multiplicative-order assumption.

The weighted space is defined in Section 3. Section 4 proves the moment representation; Section 5 establishes the finite-pattern and good-seed bounds; Section 6 combines them into the extraction-or-concentration guarantee. Section 9 states the general affine-extractor implication and its field-size constraint. The appendices give the basis algorithm, further algebraic and Nullstellensatz details, and the bounded sampling implementation. One can also formulate separate criteria on weight families for efficient basis computation, witness extraction, and collision bounds; we do not develop that general characterization here.

2 Witness encryption and the generic-group model

Fix a polynomial-time, polynomially balanced relation $R(x, w)$ and its language $\mathcal{L} = \{x : \exists w\, R(x, w) = 1\}$. All logarithms have base two. Write $\mathbb{F}_p$ for the field of prime order p, and $[B] = \{0, \dots, B-1\}$. Statistical distance is denoted by SD. The security parameter is λ; instance and witness lengths are polynomially bounded in λ in the usual asymptotic experiments.

Boolean quadratic encoding. Compile the verifier on x into a Boolean circuit, retaining the original witness bits and a length encoding when necessary. Introduce a variable for each input and wire value. An AND gate contributes $z - ab = 0$, a NOT gate contributes $z - (1 - a) = 0$, and the output is constrained to be one. Fixed bits and padding variables are constrained to their prescribed values. This gives quadratic polynomials

$$q_1, \dots, q_J \in \mathbb{F}_p[b_1, \dots, b_N]$$

whose common Boolean zeros correspond to accepting computations. Pad so that $N \geq \max\{\lambda, |x|, 4\}$; both N and J are polynomial in $\lambda + |x|$. Every valid witness has a fixed canonical extension to a satisfying assignment, and every satisfying assignment can be decoded into a valid witness. The integer gate equations are interpreted in $\mathbb{F}_p$ for each chosen prime. On Boolean assignments their meaning is unchanged for every prime used below. Put $b_0 = 1$ and

$$v(b) = (b_0, b_1, \dots, b_N)^\top \in \{0,1\}^{N+1}.$$

The coordinate b_0 represents constants and linear terms in the quadratic tests. Booleanity of $b_1, \dots, b_N$ is built into the assignment domain; its polynomial identities remain available in the moment arguments.

2.1 Fresh prime-order generic groups

For each encryption the model supplies a fresh uniformly random injection

$$\sigma : \mathbb{F}_p \longrightarrow \{0,1\}^{\ell_g}, \qquad \ell_g = 2\lceil \log p \rceil. \tag{1}$$

Write $g^a = \sigma(a)$ and $g = \sigma(1)$. The associated context Γ provides addition and subtraction in the exponent:

$$(\sigma(a), \sigma(b)) \longmapsto \sigma(a+b), \qquad (\sigma(a), \sigma(b)) \longmapsto \sigma(a-b).$$

An invalid operand produces $\perp$. Equality is string equality, and exponentiation by a known scalar uses $O(\log p)$ group queries. There is no operation multiplying two unknown exponents or revealing an exponent. This is the classical prime-order model used by Hair–Sahai [2, Section 2].

The adversary sees label bits and may perform arbitrary polynomial-time ordinary computation and query arbitrary strings. Every call to Γ, including calls made by executing the public decryption algorithm, counts toward its group-query budget. Its advice, auxiliary input, and initial randomness are independent of the fresh encoding and encryption coins conditional on the public instance. The same auxiliary input is used in the two security experiments. Different ciphertext contexts have independent encodings. The encoding table is part of the model: a ciphertext transmits finite public data together with access to its context, rather than a description of that table.

2.2 Witness-encryption definition

For a bit β, encryption returns public data and a fresh context,

$$(C, \Gamma) \leftarrow \mathsf{Enc}(1^\lambda, x, \beta).$$

Decryption is the fixed public oracle algorithm $\mathsf{Dec}^\Gamma(1^\lambda, x, C, w)$. It first checks $R(x,w)$ and returns $\perp$ on invalid witnesses. Encryption, decryption, and ciphertext length are polynomially bounded in $\lambda + |x|$, counting group queries as operations.

Definition 2.1 (Witness encryption in the generic-group model)**.** *The pair* $(\mathsf{Enc}, \mathsf{Dec})$ *is witness encryption for* R *if the following conditions hold.*

1. Correctness. *For every* x, β *and every* w *with* $R(x,w) = 1$*, decryption returns* β *except with negligible probability over encryption and the fresh context.*

2. Security on false instances. *For every efficient classical adversary* A*, every* $x \notin \mathcal{L}$*, and every polynomially bounded auxiliary input* ξ *fixed before encryption,*

$$\begin{aligned} \mathrm{Adv}_A(x,\xi) = \big| &\Pr[A^{\Gamma_0}(1^\lambda, x, \xi, C_0) = 1] \\ &- \Pr[A^{\Gamma_1}(1^\lambda, x, \xi, C_1) = 1] \big| \le \mathrm{negl}(\lambda), \end{aligned} \tag{2}$$

   where $(C_\beta, \Gamma_\beta) \leftarrow \mathsf{Enc}(1^\lambda, x, \beta)$ *in the corresponding experiment.*

Our correctness proof also gives one negligible bound for failure on *any* valid witness for a fixed ciphertext. The security theorem compares the entire finite adversary views statistically within a specified group-query budget, and hence implies the output condition in (2).

3 Weighted tables and their moments

We use Hair–Sahai's assignment tables [2, Section 5] with one complete list of polynomial weights. The table stores moments of Boolean assignments, and the verifier imposes linear conditions on those moments. Section 4 proves that column preservation forces a nonzero table to have a representation on compatible satisfying assignments. Section 5 then shows why one public random list preserves every table whose short-vector concentration is too large.

3.1 Quadratic tests and weighted assignment tables

Let $q_1, \dots, q_J$ be the Boolean quadratic encoding of the verifier from Section 2, on $N \geq 4$ variables. Write $D = N + 1$, $b_0 = 1$, and $v(b) = (b_0, b_1, \dots, b_N)^\top$. For a quadratic polynomial written as

$$q(b) = a_0 + \sum_{i=1}^{N} a_i b_i + \sum_{1 \leq i \leq j \leq N} a_{ij} b_i b_j,$$

define the linear test on D-by-D matrices

$$T_q(C) = a_0 C_{00} + \sum_{i=1}^{N} a_i C_{0i} + \sum_{1 \leq i \leq j \leq N} a_{ij} C_{ij}.$$

For an arbitrary vector $z \in \mathbb{F}_p^D$, define the homogeneous quadratic

$$\widetilde{q}(z) = a_0 z_0^2 + \sum_{i=1}^{N} a_i z_0 z_i + \sum_{1 \leq i \leq j \leq N} a_{ij} z_i z_j = T_q(zz^\top). \tag{3}$$

Thus $\widetilde{q}(v(b)) = T_q(v(b)v(b)^\top) = q(b)$. For example, the AND equation $b_k - b_i b_j = 0$ becomes $C_{0k} - C_{ij} = 0$.

For any indexed family $\mathcal{H} = (h_1, \dots, h_K)$ of functions on the Boolean cube, form the assignment table

$$A_{\mathcal{H}}(b) = \begin{pmatrix} h_1(b)v(b)v(b)^\top \\ \vdots \\ h_K(b)v(b)v(b)^\top \end{pmatrix} = \begin{pmatrix} h_1(b)v(b) \\ \vdots \\ h_K(b)v(b) \end{pmatrix} v(b)^\top. \tag{4}$$

Let $\mathcal{V}_{\mathcal{H}}$ be the span of all these tables for $b \in \{0,1\}^N$. If $C_h(M)$ denotes the D-by-D block of M indexed by h, define

$$\mathcal{S}_{x,\mathcal{H}} = \{M \in \mathcal{V}_{\mathcal{H}} : T_{q_s}(C_h(M)) = 0 \text{ for every } h \in \mathcal{H},\ 1 \leq s \leq J\}. \tag{5}$$

The span is taken over all Boolean assignments before the verifier tests are imposed; computing it does not require finding a witness.

For a representation $M = \sum_b \lambda_b A_{\mathcal{H}}(b)$, put

$$\mathcal{L}_\lambda(f) = \sum_{b \in \{0,1\}^N} \lambda_b f(b), \qquad \text{and, thus} \quad M_{(h,i),j} = \mathcal{L}_\lambda(h b_i b_j). \tag{6}$$

The term *moment* is borrowed from probability, where expectations of products of coordinates are called joint moments. For example, $\mathcal{L}_\lambda(b_i b_j) = \sum_b \lambda_b b_i b_j$ has the same algebraic form as a joint second moment. Here the coefficients λ_b are arbitrary elements of $\mathbb{F}_p$ and need not sum to one, so

these are *algebraic moments*, with no probability interpretation required. A table entry $\mathcal{L}_\lambda(hb_ib_j)$ is a weighted moment; the moment equations below are linear relations among these sums.

The weighted verifier test $T_{q_s}(C_h(M)) = 0$ is exactly $\mathcal{L}_\lambda(hq_s) = 0$. Its value is independent of the representation because it is the fixed linear test T_{q_s} of a block of M. The coefficients λ_b need never be computed.

Moment equations and pointwise constraints. A column relation $M\theta = 0$ says

$$\sum_b \lambda_b \cdot \big(h\, v(b) \cdot\ (\theta^\top v(b))\big) = 0 \qquad \text{for every } h \in \mathcal{H},$$

or, more compactly, treating $\mathcal{H}$ as a column of functions,

$$\sum_b \lambda_b \cdot \big(\mathcal{H} \otimes v(b) \cdot\ (\theta^\top v(b))\big) = 0.$$

For the complete weight list in (8), write $\ell(b) = (1, \ell_1(b), \dots, \ell_t(b))^\top$. In ordered tensor notation the same equations say

$$\sum_b \lambda_b\big(\ell(b)^{\otimes t} \otimes v(b)\big)(\theta^\top v(b)) = 0.$$

These are projected moment equations and can admit cancellation; they do not imply $\lambda_b\theta^\top v(b) = 0$ for each b. Moreover, these equations do not immediately imply $q_s(b) = 0$. The representation theorem will construct new coefficients on a small Boolean cube, where the available polynomial weights span every function and then we can make this conclusion along with satisfying q_s.

3.2 One list evaluated at powers of the seed

Set

$$\begin{aligned} &t = \lceil \log\log N \rceil, \qquad W = \binom{2t}{t}, \qquad m = DW, \\ &B = m, \qquad A_B = 2B^2, \\ &\ell_p = m + \lceil \log(Nt^2) \rceil + tD\lceil \log A_B \rceil. \end{aligned} \tag{7}$$

Take a prime $p \in [2^{\ell_p}, 2^{\ell_p+1})$ and a uniform public seed $\tau \in \mathbb{F}_p$. In particular, $p > B$; no factorization of $p-1$ or multiplicative-order condition on τ is required. Evaluate the assignment polynomial $\sum_{i=0}^N b_i X^i$ at $\tau, \tau^2, \dots, \tau^t$. The resulting linear forms and their complete degree-t weight list are

$$\ell_j(b) = \sum_{i=0}^{N} \tau^{ij} b_i, \qquad h_\alpha(b) = \prod_{j=1}^{t} \ell_j(b)^{\alpha_j}, \quad \alpha \in \mathbb{Z}^t_{\geq 0}, \quad |\alpha| \leq t. \tag{8}$$

If we also include $\ell_0(b) = b_0 = 1$, then we can restrict to α s.t. $|\alpha| = t$, i.e.

$$h_\alpha(b) = \prod_{j=0}^{t} \ell_j(b)^{\alpha_j}, \quad \alpha \in \mathbb{Z}^{t+1}_{\geq 0}, \quad |\alpha| = t. \tag{9}$$

There are W indices α. We keep distinct indices even if they define the same function, and interpret all zeroth powers as one, including at $\tau = 0$. In particular, $h_{t,0,\dots,0} = 1$. Distinguish the t power columns from the augmented matrix of forms:

$$\begin{aligned} &F_t(\tau) = (\tau^{ij})_{0\leq i\leq N,\, 1\leq j\leq t} \in \mathbb{F}_p^{D\times t}, \\ &F(\tau) = [\mathbf{e}_0 \mid F_t(\tau)] \in \mathbb{F}_p^{D\times(t+1)}, \qquad \ell(b) = (1, \ell_1(b), \dots, \ell_t(b))^\top = F(\tau)^\top v(b). \end{aligned} \tag{10}$$

Here $\mathbf{e}_0$ is the zeroth coordinate vector. Column preservation will be established for F_t; the constant column in F lets one write all weights as homogeneous degree-t products. In tensor notation, we have

$$h_\alpha(b) = \prod_{j=0}^{t} \Big(F(\tau)^{(j)} \cdot v(b)\Big)^{\alpha_j} \;=\; \left(\bigotimes_{j=0}^{t} (F(\tau)^{(j)})^{\otimes \alpha_j}\right) \cdot v(b)^{\otimes t} \tag{11}$$

For this family write

$$A_\tau(b)_{(\alpha,i),j} = h_\alpha(b) b_i b_j, \qquad A_\tau(b) = u_\tau(b) v(b)^\top, \qquad u_\tau(b) = (h_\alpha(b) v(b))_\alpha. \tag{12}$$

These are nonzero m-by-D rank-one matrices. Abbreviate the two spaces by $\mathcal{V}_\tau$ and $\mathcal{S}_{x,\tau}$, so

$$\mathcal{S}_{x,\tau} = \{M \in \mathcal{V}_\tau : T_{q_s}(C_\alpha(M)) = 0 \text{ for all } \alpha, s\}, \qquad k := \dim \mathcal{S}_{x,\tau} \le mD. \tag{13}$$

We keep matrices rectangular throughout; this avoids storing unused zero columns or sampling unused coordinates of the short vector.

3.3 Available moments and their identities

Let $M \in \mathcal{V}_\tau$ have representation

$$M = \sum_b \lambda_b A_\tau(b).$$

Its block indexed by α is

$$C_\alpha(M) = \sum_b \lambda_b h_\alpha(b) v(b) v(b)^\top.$$

For a weight $f(b) = \sum_\alpha a_\alpha h_\alpha(b)$, define

$$C_f(M) := \sum_\alpha a_\alpha C_\alpha(M) = \sum_b \lambda_b f(b) v(b) v(b)^\top. \tag{14}$$

Thus adding weights corresponds to adding blocks. For example, when $t = 2$, the homogeneous weight $\ell_1(\ell_0 - \ell_2)$ gives the block $C_{\ell_0\ell_1} - C_{\ell_1\ell_2}$. For $M \in \mathcal{S}_{x,\tau}$, every available C_f satisfies

$$C_f = C_f^\top, \qquad (C_f)_{ii} = (C_f)_{0i}, \qquad T_{q_s}(C_f) = 0, \qquad \ker M \subseteq \ker C_f, \tag{15}$$

since each α-th block of M satisfies these properties; the last inclusion (on right kernel) holds because C_f is a linear combination of the stacked blocks of M. These identities provide, respectively, symmetric moments, Boolean reduction, weighted verifier equations, and the transfer of column relations. The two factors $b_i b_j$ supply two moment degrees in addition to the weight degree t.

Every satisfying assignment has $A_\tau(b) \in \mathcal{S}_{x,\tau}$ for every seed. Once a public basis $M_1, \ldots, M_k$ is fixed, its coordinates give the certificate

$$\sum_{j=1}^{k} a_j M_j = u_\tau(b) v(b)^\top \neq 0, \qquad v(b) \in \{0,1\}^D, \qquad \mathrm{wt}(v(b)) \le D. \tag{16}$$

The Boolean right factor is the property used for decryption.

3.4 Efficient public basis computation

Proposition 3.1 (Efficient public basis)**.** *For every choice of public parameters x, p, τ, a deterministic algorithm computes an independent basis $M_1, \dots, M_k$ of $\mathcal{S}_{x,\tau}$ in time* $\mathrm{poly}(N, m, \log p)$*. Given a satisfying assignment, it computes the certificate in* (16) *within the same time bound.*

Proof sketch. Assigning one more bit translates each linear form by a known constant. The complete degree-t weight list is closed under these translations, so each bit extension acts linearly on the table vectors. Dependencies are therefore preserved. At each of the N stages, extend the retained assignments by both bit values and keep a basis of their tables; at most mD assignments survive. Their final tables span $\mathcal{V}_\tau$. Impose the weighted verifier tests by a homogeneous linear system, then solve for the coordinates of an honest table when a witness is given. Fixed orders make the algorithm deterministic. Appendix A gives the full *extend, reduce, retain* algorithm from [2, Section 5.4], the bit-update identity, and the invariant that justifies pruning. □

The basis of $\mathcal{V}_\tau$ depends only on N, p, t, τ and can be public and reused across instances of the same padded size. Imposing the verifier tests is the instance-dependent step; its result can also be cached for repeated encryptions of the same x. Appendix A distinguishes these preprocessing costs and explains why the simultaneous seed guarantee permits reuse.

4 Preserved columns give a representation on satisfying assignments

The central algebraic statement in this work extends Hair–Sahai's witness-extraction argument to a representation of the whole table. A table whose column space is preserved by the chosen linear map can be represented using at most 2^ρ compatible satisfying assignments, where ρ is the rank of the table. We first compress the table to a symmetric moment tensor, from which the entire table can be recovered. Boolean moment identities then replace the original sum over assignments by a sum over a ρ-dimensional Boolean cube, and weighted verifier equations restrict its support to satisfying assignments. This uses the Boolean interpolation underlying Hair–Sahai's extraction argument [2, Section 5.3], with tensor symmetry making reconstruction explicit.

4.1 The representation theorem

Fix $t \geq 1$ and the complete degree-t weight list (8). Recall the two matrices from (10): $F_t(\tau)$ consists of the t power columns, and $F(\tau) = [\mathbf{e}_0 \mid F_t(\tau)]$ includes the constant form. The colSpan preservation hypothesis below concerns F_t; the *augmented* matrix F is used to write every weight as a homogeneous degree-t product. Column-span preservation is a substantive condition even when $\operatorname{rank} M \leq t$. Looking ahead, Lemma 5.4 will establish column span preservation for all matrices with substantial concentration. The argument in this section works over every field and for any t affine-linear forms in place of the power forms, with their constant augmentation.

Theorem 4.1 (Representation on compatible satisfying assignments)**.** *Let $0 \neq M \in \mathcal{S}_{x,\tau}$ satisfy*

$$\mathrm{colSpan}(MF_t(\tau)) = \mathrm{colSpan}(M). \tag{17}$$

Put $\rho = \operatorname{rank} M$, choose ρ independent actual columns as P, and write $M = PR$. The corresponding columns of $R \in \mathbb{F}_p^{\rho \times D}$ form I_ρ. Then $\rho \leq t$, and there is a nonempty set $\mathcal{E} \subseteq \{0,1\}^\rho$, and coefficients

$\omega_e \in \mathbb{F}_p \setminus \{0\}$ *such that*

$$M = \sum_{e \in \mathcal{E}} \omega_e A_\tau(b_e), \qquad v(b_e) = R^\top e, \qquad q_s(b_e) = 0 \quad (1 \le s \le J). \tag{18}$$

Every $v(b_e)$ is an augmented Boolean vector in the row space of M, and $|\mathcal{E}| \le 2^\rho$.

Proof overview. With the column basis and compression fixed, the proof follows the diagram below. Its arrows describe constructions justified in the proof. Here μ is the compressed moment functional, Z its tensor of degree-$(t+2)$ moments, and $\omega = (\omega_e)_{e \in \{0,1\}^\rho}$ its Boolean interpolation coefficients.

$$\begin{array}{ccccc} M & \xmapsto{\text{compression}} & \mu & \xmapsto{\text{interpolation}} & \omega \\ \| & & \downarrow{\scriptstyle \text{top moments}} & & \downarrow{\scriptstyle \text{weighted tests}} \\ M_{\text{rec}} & \xmapsto{\text{reconstruction}} & Z & \xmapsto{\text{moments}} & \omega|_{\mathcal{E}} \end{array} \tag{19}$$

The weighted tests show that every e in $\mathcal{E} = \{e : \omega_e \neq 0\}$ gives a compatible satisfying assignment b_e. The right square says that both routes give the same tensor Z; the left says that reconstruction returns $M_{\text{rec}} = \sum_{e \in \mathcal{E}} \omega_e A_\tau(b_e) = M$. In particular, nonvanishing of the original table forces a nonempty satisfying support.

Proof. Abbreviate $F_t = F_t(\tau)$ and $F = F(\tau)$. Choose any representation $M = \sum_b \lambda_b A_\tau(b)$ and put $\mathcal{L}(f) = \sum_b \lambda_b f(b)$. These original coefficients λ_b need not vanish on incompatible or unsatisfying assignments b.

Column preservation and the rank bound. Column-span preservation implies that RF_t has row rank ρ: otherwise PRF_t could not have the same column span as P. Indeed, multiplication by the full-column-rank matrix P preserves rank, so

$$\operatorname{rank}(RF_t) = \operatorname{rank}(MF_t) = \operatorname{rank} M = \rho \le t,$$

where the last inequality uses the t columns of F_t. The augmented matrix $Q := RF$ also has row rank ρ. Set $a = t+1$, choose a right inverse $W \in \mathbb{F}_p^{a \times \rho}$ so that $QW = I_\rho$, and define

$$U = FW, \qquad RU = I_\rho, \qquad w(b) = U^\top v(b) = W^\top \ell(b), \qquad L_i(y) = (R^\top y)_i, \tag{20}$$

where $\ell(b) = F^\top v(b)$ includes the constant coordinate $\ell_0(b) = 1$. The linear map $U^\top : \mathbb{F}_p^D \longrightarrow \mathbb{F}_p^\rho$ compresses the augmented Boolean vector $v(b) \in \{1\} \times \{0,1\}^N$ to $w(b)$.

Every available block $C = C_h(M)$ is symmetric and annihilates $\ker M$. Since $MUR = PRUR = PR = M$, it follows that

$$C = CUR = R^\top (U^\top C U) R.$$

If $i_1, \dots, i_\rho$ are the selected column indices (for P), then $MU_j = M\mathbf{e}_{i_j}$ and hence $CU_j = C\mathbf{e}_{i_j}$. For a polynomial f over $\mathbb{F}_p$ in ρ variables and of total degree at most $t+2$, define the scalar functional

$$\mu(f) = \mathcal{L}(f(w(b))) \quad (\deg f \le t+2).$$

For a polynomial g in the same variables, further define the weighted block

$$D_g = C_{g(w)}(M) \quad (\deg g \le t).$$

Since $w(b) = W^\top \ell(b)$, the function $g(w(b))$ is a polynomial of degree at most t in the available forms. Expanding it gives a linear combination of the listed weights; lower degrees are padded by the constant form. In particular,

$$\begin{aligned} D_g &= \sum_b \lambda_b g(w(b)) v(b) v(b)^\top, \\ U^\top D_g U &= \sum_b \lambda_b g(w(b)) w(b) w(b)^\top = \mu\big(g(y) y y^\top\big). \end{aligned} \tag{21}$$

Here, y is a vector of ρ variables and μ is applied entrywise to polynomial matrices. Combining this identity with the block relations above, i.e. $CU_j = C\mathbf{e}_{i_j}$ and $C = CUR$, gives

$$D_g U_j = D_g \mathbf{e}_{i_j}, \qquad D_g = R^\top \mu\big(g(y) y y^\top\big) R. \tag{22}$$

All moments through degree $t+2$ are available: every monomial splits into a weight of degree at most t and at most two further factors supplied by the endpoint vectors. For lower degrees, selecting $\mathbf{e}_0$ in an endpoint slot supplies the constant factor $\mathbf{e}_0^\top v(b) = 1$. Thus these moments are determined by M and the chosen compression, independently of the initial representation λ.

A symmetric tensor and its reversible compression. For the proof, expand the homogeneous monomial list from (9) into ordered t-tuples of forms. Each ordering contributes a copy of the corresponding block, so this expansion only duplicates blocks, without introducing multinomial coefficients. Denote the resulting stacked table by $\widetilde{M}$. Put

$$\begin{aligned} \widetilde{M} &= \sum_b \lambda_b \big(\ell(b)^{\otimes t} \otimes v(b)\big) v(b)^\top, \\ G &= (I_{a^t} \otimes F^\top) \widetilde{M} F = \sum_b \lambda_b \, \ell(b)^{\otimes(t+1)} \ell(b)^\top. \end{aligned} \tag{23}$$

Thus G, viewed as a tensor of order $t+2$, is symmetric under every permutation of its slots. These ordered tensors are a notational device for the proof; the algorithms keep the more compact monomial list.

Since $\widetilde{M}UR = \widetilde{M}$, we have $G = GWQ$. Equivalently, the projection $\Pi = Q^\top W^\top$ fixes the last tensor slot of G. Symmetry implies that it fixes every slot. Applying these commuting projections in all slots gives

$$\begin{aligned} Z &= (W^\top)^{\otimes(t+1)} G W = \sum_b \lambda_b \, w(b)^{\otimes(t+1)} w(b)^\top, \\ G &= (Q^\top)^{\otimes(t+1)} Z Q. \end{aligned} \tag{24}$$

The first line is multiplication on the left by $(W^\top)^{\otimes t} \otimes W^\top$ and on the right by W. Although W can be rectangular, the second line recovers G. Both directions can be followed in this diagram of linear maps:

$$\begin{array}{ccc} \mathbb{F}_p^a & \xrightarrow{\quad G \quad} & (\mathbb{F}_p^a)^{\otimes(t+1)} \\ {\scriptstyle Q}\downarrow \quad \uparrow{\scriptstyle W} & & (W^\top)^{\otimes(t+1)}\downarrow \quad \uparrow (Q^\top)^{\otimes(t+1)} \\ \mathbb{F}_p^\rho & \xrightarrow[\quad Z \quad]{} & (\mathbb{F}_p^\rho)^{\otimes(t+1)} \end{array} \tag{25}$$

Starting with $\xi \in \mathbb{F}_p^a$ or $z \in \mathbb{F}_p^\rho$, respectively, the downward and upward chases give

$$(W^\top)^{\otimes(t+1)} G\xi = (W^\top)^{\otimes(t+1)} GWQ\xi = ZQ\xi,$$
$$GWz = (Q^\top)^{\otimes(t+1)} ZQWz = (Q^\top)^{\otimes(t+1)} Zz.$$

The lower round trips are identities because $QW = I_\rho$. The upper round trips are projections: WQ on the input and $(Q^\top W^\top)^{\otimes(t+1)}$ on the output. Preservation and symmetry make these projections fix G, which is why the compression loses no information here.

The endpoint projections also lose no information. Applying $C = R^\top U^\top CUR$ to every block and using $U = FW$ gives

$$\begin{aligned} \widetilde{M} &= (I_{a^t} \otimes R^\top W^\top) GWR \\ &= \big((Q^\top)^{\otimes t} \otimes R^\top\big) ZR. \end{aligned} \tag{26}$$

This recovers every original weight block from the same compressed tensor. In particular, $M \neq 0$ implies $Z \neq 0$. No pointwise replacement on the original λ_b-support has been asserted. Appendix B also gives this transformation with an additional invertible change of compressed coordinates.

Replacing the sum over b by a sum over Boolean e. The values $w(b)$ need not be Boolean. Nevertheless, column span preservation transfers the blocks' Boolean identities to the moments of $w(b)$. Consequently, every polynomial of degree at most $t+2$ has the same moment as its multilinear reduction modulo $y_j^2 - y_j$, as we show next. For $\deg g \leq t$,

$$\begin{aligned} \mu(gy_j^2) &= U_j^\top D_g U_j = \mathbf{e}_{i_j}^\top D_g \mathbf{e}_{i_j} \\ &= \mathbf{e}_0^\top D_g \mathbf{e}_{i_j} = \mathbf{e}_0^\top D_g U_j = \mu(gy_j). \end{aligned}$$

The middle equality is the Boolean diagonal identity for D_g. Consequently,

$$\mu\big(g(y)(y_j^2 - y_j)\big) = 0 \qquad (\deg g \leq t). \tag{27}$$

Every polynomial of degree at most $t+2$ can therefore be reduced inside μ inductively to its multilinear representative: each reduction uses a Boolean equation with a multiplier of degree at most t. For each $e \in \{0,1\}^\rho$, define

$$\chi_e(y) = \prod_{j:e_j=1} y_j \prod_{j:e_j=0} (1 - y_j), \qquad \omega_e = \mu(\chi_e).$$

These polynomials have degree at most $\rho \leq t$ and satisfy $\chi_e(e') = \delta_{e,e'}$. They form the interpolation basis for multilinear polynomials in ρ variables. For a polynomial f of degree at most $t+2$, let $\bar{f}$ be its multilinear representative. Boolean reduction gives $\mu(f) = \mu(\bar{f})$, and interpolation gives

$$\bar{f}(y) = \sum_{e \in \{0,1\}^\rho} f(e)\chi_e(y).$$

This passage can be followed without making any pointwise claim about the original $w(b)$. Write $\mathcal{P}_{\leq k} = \mathbb{F}_p[y_1, \ldots, y_\rho]_{\leq k}$, let $\mathcal{P}_{\mathrm{ml}}$ be the space of multilinear polynomials, and put $\Omega = \{0,1\}^\rho$. Boolean reduction and evaluation give

$$\begin{array}{ccccc} \mathcal{P}_{\leq t+2} & \xrightarrow{f \mapsto \bar{f}} & \mathcal{P}_{\mathrm{ml}} & \xrightarrow[\simeq]{\mathrm{ev}} & \mathbb{F}_p^\Omega \\ & {\scriptstyle\mu}\searrow & \downarrow{\scriptstyle\mu} & \swarrow{\scriptstyle\langle\omega,\cdot\rangle} & \\ & & \mathbb{F}_p & & \end{array} \tag{28}$$

Here $\mathrm{ev}(f) = (f(e))_{e\in\Omega}$ and $\langle \omega, z\rangle = \sum_e \omega_e z_e$. The middle arrow restricts μ to multilinear polynomials, whose degrees are at most $\rho \le t$. Starting with f, the direct path gives $\mu(f)$, the middle path gives $\mu(\bar{f})$, and the right path gives $\sum_e \omega_e f(e)$. Boolean reduction and interpolation show that these are equal:

$$\mu(f) = \mu(\bar{f}) = \sum_{e\in\{0,1\}^\rho} \omega_e f(e) \qquad (\deg f \le t+2). \tag{29}$$

In particular, the same coefficients represent the whole tensor:

$$Z = \sum_{e\in\{0,1\}^\rho} \omega_e \, e^{\otimes(t+1)} e^\top. \tag{30}$$

Moreover, $\omega_e = \mathbf{e}_0^\top D_{\chi_e} \mathbf{e}_0$, so they are determined by M and the chosen column basis P, independently of the initial representation λ. Neither positivity nor $\mu(1) = 1$ is required.

Restricting the support and recovering assignment tables. For $\deg g \le t$, the weighted-sum identity (21) and $\mathbf{e}_0^\top v(b) = 1$ give

$$\begin{aligned}
\mathbf{e}_0^\top D_g \mathbf{e}_0 &= \sum_b \lambda_b g(w(b)) = \mu(g), \\
\mathbf{e}_0^\top D_g U R \mathbf{e}_0 &= \sum_b \lambda_b g(w(b)) w(b)^\top R \mathbf{e}_0 = \mu(gL_0).
\end{aligned}$$

Here $w(b)^\top R\mathbf{e}_0 = L_0(w(b))$ by (20). The constant-coordinate relation $\mathbf{e}_0 - UR\mathbf{e}_0 \in \ker M \subseteq \ker D_g$ therefore gives

$$0 = \mathbf{e}_0^\top D_g(\mathbf{e}_0 - UR\mathbf{e}_0) = \mu\big(g(y)(1 - L_0(y))\big). \tag{31}$$

To transfer the remaining Boolean identities, apply (22) entrywise. Since $L_i(y) = y^\top R\mathbf{e}_i$ and μ is linear,

$$\begin{aligned}
(D_g)_{ij} &= \mathbf{e}_i^\top R^\top \mu\big(g(y) y y^\top\big) R\mathbf{e}_j \\
&= \mu\big(g(y) L_i(y) L_j(y)\big) \qquad (0 \le i,j \le N).
\end{aligned}$$

In particular, $(D_g)_{ii} = \mu(gL_i^2)$ and $(D_g)_{0i} = \mu(gL_0L_i)$. The original Boolean diagonal identity $(D_g)_{ii} = (D_g)_{0i}$ therefore gives the first line below by subtraction. This does not assume that $L_i(y)$ is Boolean for arbitrary y. The verifier tests similarly give the second line, using (3):

$$\begin{aligned}
0 &= \mu\big(g(y)(L_i(y)^2 - L_0(y)L_i(y))\big) \\
&\quad (1 \le i \le N), \\
0 &= T_{q_s}(D_g) = \mu\big(g(y)\widetilde{q}_s(R^\top y)\big) \\
&\quad (1 \le s \le J).
\end{aligned} \tag{32}$$

All these identities (31), (32) hold for $\deg g \le t$. To see how they constrain each coefficient separately, let c be any of the transferred polynomials $1 - L_0$, $L_i^2 - L_0L_i$, or $\widetilde{q}_s(R^\top y)$, and set

$$\mu_c(g) = \mu(gc), \qquad \Lambda_c(z) = \sum_{e\in\Omega} \omega_e c(e) z_e.$$

The weighted identities and interpolation give a commuting square:

$$\begin{array}{ccc} \mathcal{P}_{\le t} & \xrightarrow{\mu_c=0} & \mathbb{F}_p \\ {\scriptstyle \mathrm{ev}}\big\downarrow & & \big\| \\ \mathbb{F}_p^{\Omega} & \xrightarrow[\Lambda_c]{} & \mathbb{F}_p \end{array} \tag{33}$$

Both paths send g to $\mu(gc) = \sum_e \omega_e g(e)c(e)$; these moments are available because $\deg(gc) \le t+2$. Evaluation is onto since each selector χ_e has degree at most $\rho \le t$. Chasing χ_e, whose evaluation is the coordinate vector δ_e, gives

$$0 = \mu(\chi_e c) = \Lambda_c(\delta_e) = \omega_e c(e).$$

The selector product has degree at most $\rho + 2 \le t + 2$. For the displayed constraints, this says

$$\begin{aligned} \omega_e(1 - L_0(e)) &= 0, \\ \omega_e(L_i(e)^2 - L_0(e)L_i(e)) &= 0, \\ \omega_e \widetilde{q}_s(R^\top e) &= 0. \end{aligned}$$

Thus every e with nonzero coefficient ω_e has $L_0(e) = 1$, all $L_i(e)$ Boolean, and $q_s(b_e) = 0$ for $v(b_e) = R^\top e$. The weighted tests are what impose these conditions separately at each point; an unweighted equation would allow cancellation.

Finally substitute (30) into (26):

$$\widetilde{M} = \sum_e \omega_e \big((Q^\top e)^{\otimes t} \otimes R^\top e\big)(R^\top e)^\top.$$

For every nonzero coefficient, $Q^\top e = F^\top R^\top e = F^\top v(b_e)$. Removing the duplicated blocks therefore proves (18). Equivalently, for every polynomial weight H of degree at most t,

$$C_{H(\ell)}(M) = \sum_{e:\omega_e \neq 0} \omega_e H(Q^\top e)(R^\top e)(R^\top e)^\top. \tag{34}$$

Some ω_e is nonzero because $Z \neq 0$. Taking $\mathcal{E} = \{e : \omega_e \neq 0\}$ completes the proof. □

4.2 Finding a witness without computing the representation

The representation theorem proves that the existing column-basis search succeeds. The algorithm does not compute U, μ, or the coefficients ω_e.

Given $M \in \mathcal{S}_{x,\tau}$, return $\perp$ if $M = 0$ or $\operatorname{rank} M > t$. Otherwise choose independent actual columns P, write $M = PR$, and set $L_i(y) = (R^\top y)_i$. Enumerate $e \in \{0,1\}^\rho$ and test

$$L_0(e) = 1, \qquad L_i(e) \in \{0,1\}\ (1 \le i \le N), \qquad q_s(L_1(e), \dots, L_N(e)) = 0 \quad (1 \le s \le J). \tag{35}$$

Decode and validate a witness from the first passing assignment; if no candidate passes, return $\perp$. Denote this deterministic algorithm by $\mathsf{Ext}_{x,\tau}$, with fixed column and candidate orders. Its running time is $2^t \operatorname{poly}(N, m, \log p)$, including unsuccessful searches.

An assignment is *compatible* with M when its augmented vector lies in the row space of M. Its entries at the selected column indices form an $e \in \{0,1\}^\rho$, so the search includes every compatible Boolean assignment.

Lemma 4.2 (Conditional extraction)**.** *If $0 \neq M \in \mathcal{S}_{x,\tau}$ satisfies* (17)*, then $\mathsf{Ext}_{x,\tau}(M)$ returns a valid witness.*

Proof. Theorem 4.1 supplies a nonempty support of compatible satisfying assignments. Each occurs in the candidate list, so the search finds one. □

For $t = \lceil \log \log N \rceil$, at most $2 \log N$ candidates are tested; for $t = \lceil \log N \rceil$, at most $2N$ are tested. Honest decryption starts with a supplied witness and does not run this search. The theorem is used to prove security on false instances.

5 Concentration and a simultaneous good seed

Theorem 4.1 explains what column preservation buys: a nonzero table must have satisfying support. We now show that preservation is needed only for matrices with substantial short-vector concentration. Such a matrix has a column factorization from a finite family of short rational patterns, all preserved simultaneously by one random seed. The scalar-moment proof itself uses no bound on these rational coordinates.

5.1 Collision probability beyond rank

For $[B] = \{0, \ldots, B-1\}$ and $M \in \mathbb{F}_p^{m \times D}$, define

$$\kappa_B(M) = \max_{y \in \mathbb{F}_p^m} \Pr_{r \leftarrow [B]^D}[Mr = y], \qquad H_B(M) = -\log \kappa_B(M). \tag{36}$$

The next identity refines the conditioning argument in Hair–Sahai's fixed-linear-test lemma [2, Lemma 3.4].

Lemma 5.1 (Exact bilinear collision probability)**.** *For independent uniform $s \in \mathbb{F}_p^m$ and $r \in [B]^D$,*

$$\max_{c \in \mathbb{F}_p,\, \alpha \in \mathbb{F}_p^m} \Pr[c + s^\top(\alpha + Mr) = 0] = p^{-1} + (1 - p^{-1})\kappa_B(M). \tag{37}$$

In particular, each fixed pair (c, α) has collision probability at most $\kappa_B(M) + p^{-1} = 2^{-H_B(M)} + p^{-1}$. If $M = 0$ and $(c, \alpha) \neq (0, 0)$, the probability is at most p^{-1}.

Proof. Put $a = \Pr[Mr = -\alpha]$. On this event the expression is c; off it, the expression is a nonconstant affine-linear function of uniform s and vanishes with probability $1/p$. Its zero probability is therefore $a + (1-a)/p$ for $c = 0$ and $(1-a)/p$ for $c \neq 0$. Maximizing proves the identity. For $M = 0$, a nonzero linear part gives probability $1/p$, while a nonzero constant with zero linear part never vanishes. □

Rank ρ implies $H_B(M) \geq \rho \log B$ by conditioning on the coordinates outside ρ independent columns. Rank is only a sufficient certificate: the rank-one row $(1, B, \ldots, B^{D-1})$ has $H_B = D \log B$ when $p > B^D$. The next definition captures the injectivity used in the conditioning argument without requiring $\mathbb{F}_p$-independence.

Call columns $c_1, \ldots, c_\nu$ *B-dissociated* if

$$\sum_{i=1}^{\nu} a_i c_i = 0 \text{ with } \; a_i \in \{-(B-1), \ldots, B-1\} \quad \Longrightarrow \quad a_1 = \cdots = a_\nu = 0.$$

Coefficients are integers interpreted in $\mathbb{F}_p$. $\mathbb{F}_p$-independent columns are clearly B-dissociated, but the converse need not be true.

Lemma 5.2 (Concentration or short column coordinates)**.** *If M contains ν B-dissociated columns, then*

$$H_B(M) \geq \nu \log B, \qquad \textit{equivalently} \qquad \kappa_B(M) \leq B^{-\nu}. \tag{38}$$

If $M \neq 0$ contains no $t+1$ such columns, then for some $1 \leq \nu \leq t$ it has a factorization

$$\begin{aligned} M = PQ, \qquad P \in \mathbb{F}_p^{m\times\nu}, \qquad Q \in \mathcal{R}_B^{\nu\times D}, \qquad \operatorname{rank} Q = \nu, \\ \mathcal{R}_B = \{a/b \bmod p : |a| < B,\ 1 \leq b < B\}, \qquad |\mathcal{R}_B| < A_B = 2B^2. \end{aligned} \tag{39}$$

The columns of P need not be independent over $\mathbb{F}_p$.

Proof. Condition on coordinates of r outside a dissociated set. Two choices of the remaining coordinates with the same image would give a forbidden short relation. The map on these B^ν choices is injective, proving (38).

For the factorization, choose a maximal (by subset relation) dissociated set of $\nu \leq t$ columns and place it in P. It is nonempty because $M \neq 0$ and $p > B$. Maximality implies that every other column c satisfies

$$ac + \sum_{i=1}^{\nu} a_i P_i = 0, \qquad 0 < |a| < B, \qquad |a_i| < B.$$

The coefficient of c must be nonzero since the selected columns are dissociated. Change all signs if necessary to make $a > 0$. Then c has coordinates $-a_i/a \in \mathcal{R}_B$. At the selected columns choose the corresponding columns of Q to be I_ν. Thus Q has row rank ν even if P is $\mathbb{F}_p$-dependent. The number of possible fractions is at most $(2B-1)(B-1) < 2B^2$. □

This factorization is only a proof device. Encryption, decryption, and the extractor below do not search for dissociated sets or short integer relations.

Relation to Littlewood–Offord theory. Classical Littlewood–Offord theory studies the concentration of random signed sums. Erdős's sharp bound [11] implies that, for nonzero real coefficients $a_1, \ldots, a_n$ and independent uniform signs $\varepsilon_i \in \{-1, 1\}$,

$$\max_{z\in\mathbb{R}} \Pr\left[\sum_{i=1}^{n} a_i\varepsilon_i = z\right] \leq 2^{-n}\binom{n}{\lfloor n/2\rfloor} = O(n^{-1/2}).$$

Our random sum $Mr = \sum_i r_i c_i$ instead has columns $c_i \in \mathbb{F}_p^m$ and independent uniform coefficients $r_i \in [B]$. The proof above establishes its concentration bound directly by conditioning on the coordinates outside a dissociated set and using injectivity on the remaining grid; it does not invoke the classical bound for real coefficients.

The structural consequence has an elementary inverse Littlewood–Offord interpretation: substantial concentration forces additive structure. Specifically, for $M \neq 0$,

$$H_B(M) < (t+1)\log B \quad \Longrightarrow \quad M = PQ \text{ as in (39), with } 1 \leq \nu \leq t.$$

Indeed, the concentration bound rules out $t+1$ dissociated columns, and an inclusion-maximal dissociated set expresses every remaining column using fractions a/b, with $|a| < B$ and $1 \leq b < B$, interpreted modulo p. This reflects the principle that concentration forces structure in inverse Littlewood–Offord theory [12], but follows directly from the maximality argument in Lemma 5.2; no general inverse theorem is used. For our application, the useful consequence is the finite family of possible coordinate matrices Q. It permits the union bound in Lemma 5.4, ensuring that one public seed preserves all relevant patterns simultaneously, including those used by matrices chosen after the seed is known.

5.2 A determinant for evaluations at powers of the seed

The power columns $j = 1, \dots, t$ of the form matrix in (10) have polynomial entries

$$F_t(T)_{i,j} = T^{ij}, \qquad 0 \le i \le N, \quad 1 \le j \le t.$$

These columns are the coefficients of the forms ℓ_j when $T = \tau$. Write $F_\nu(T)$ for the first ν columns of $F_t(T)$. Gabizon–Raz use the following rank-preservation lemma, stated here with our indexing, as the linear seeded component of their construction of deterministic extractors for affine sources. Hair–Sahai use the geometric-folding refinement of Forbes–Shpilka [2, Lemma 5.2], [9, Theorem 4.1].

These maps are examples of *seeded rank condensers* [10]. Appendix G gives a reformulation of the proof and describes other univariate polynomial matrices $F(T)$ with the same preservation property and smaller determinant degree, including bounds of $O(N\nu)$. For the power map used here, the first evaluation point is T: using $T^0 = 1$ would give a zero first column for the nonzero polynomial $X - 1$.

Lemma 5.3 (Gabizon–Raz [8, Lemma 6.1])**.** *Over any field $\mathbb{F}_p$, for every full-row-rank $Q \in \mathbb{F}_p^{\nu \times D}$ with $1 \le \nu \le t$, the determinant of $QF_\nu(T)$ is a nonzero polynomial of degree at most $N\nu(\nu+1)/2$.*

5.3 Simultaneous preservation of short-coordinate patterns

Lemma 5.4 (One seed preserves every short-coordinate pattern)**.** *Except with probability*

$$\delta_\tau \le \frac{N}{p} \sum_{\nu=1}^{t} \frac{\nu(\nu+1)}{2} A_B^{\nu D} \le \frac{N t^2 A_B^{tD}}{p} \le 2^{-m}, \tag{40}$$

a uniform seed τ satisfies

$$\operatorname{col}(MF_t(\tau)) = \operatorname{col}(M)$$

simultaneously for every M that has a factorization as in (39)*, regardless of left factor P.*

Proof. Fix a full-row-rank $Q \in \mathcal{R}_B^{\nu \times D}$ for $1 \le \nu \le t$. Lemma 5.3 and the root bound for a nonzero polynomial give

$$\Pr_\tau[\operatorname{rank}(QF_t(\tau)) < \nu] \le \frac{N\nu(\nu+1)}{2p}.$$

There are at most $A_B^{\nu D}$ possible such matrices Q. Union-bound over these matrices and then over ν to obtain the first inequality of (40). For the second, put $R = A_B^D \ge 2$. Induction on t gives

$$\sum_{\nu=1}^{t} \frac{\nu(\nu+1)}{2} R^\nu \le t^2 R^t :$$

the case $t = 1$ is equality, and the normalized sum at step t is at most

$$\frac{t(t+1)}{2} + \frac{(t-1)^2}{R} \le t^2 - \frac{t-1}{2} \le t^2.$$

Thus the larger individual determinant degrees still permit the stated aggregate bound. The definition of ℓ_p gives $p \ge 2^m N t^2 A_B^{tD}$, yielding the last inequality.

Outside this exceptional event both Q and $QF_t(\tau)$ map (surjectively) onto $\mathbb{F}_p^\nu$. Consequently, for every P their compositions have the same column space, namely $\operatorname{col}(P)$. □

The exceptional event is defined using only the finite family of short-coordinate matrices Q. It places no restriction on P: the left factor may depend on the field, the public seed, the instance, and the adversary's previous computation. Hence the preservation statement holds simultaneously for adaptively chosen matrices and for all verifier instances with the same N.

The matrix Q in this counting argument need not be the coordinate matrix R used in Theorem 4.1. Here P may have dependent columns, while that theorem uses an actual independent column basis. The common conclusion is (17); once it holds, the representation theorem chooses its own column coordinates.

The seed can also be chosen deterministically by bounding the integer coefficients of the preservation determinants and using a larger field. This removes the exceptional-seed event; we do not explore that deterministic variant here.

6 Witness extraction or small collision probability

The representation theorem and the simultaneous seed lemma now meet. A matrix with too much concentration has a short-coordinate factorization. On a good seed its columns are preserved, so its moments have a nonzero representation on satisfying assignments. This gives the interface needed by witness encryption.

Theorem 6.1 (Extraction or small collision probability)**.** *For every seed outside the exceptional event of Lemma 5.4, every nonzero $M \in \mathcal{S}_{x,\tau}$ satisfies at least one of:*

1. *M admits the satisfying representation* (18)*, and $\mathsf{Ext}_{x,\tau}(M)$ returns a valid witness for x;*
2. *$\kappa_B(M) \le B^{-(t+1)}$.*

The event is simultaneous over all such matrices and all instances with the same N, including true instances.

Proof. If M has $t+1$ dissociated columns, apply Lemma 5.2. Otherwise that lemma gives a short-coordinate factorization, the good-seed property gives (17), and Theorem 4.1 supplies the representation. Lemma 4.2 then supplies a witness. □

Corollary 6.2 (False-instance concentration)**.** *On a good seed, for every $x \notin \mathcal{L}$ and every nonzero $M \in \mathcal{S}_{x,\tau}$, $\kappa_B(M) \le B^{-(t+1)}$. More generally the same bound holds whenever the validated extractor returns $\perp$. Consequently, for such a matrix, every fixed c, α and independent uniform $s \in \mathbb{F}_p^m$, $r \in [B]^D$ satisfy*

$$\Pr[c + s^\top(\alpha + Mr) = 0] \le B^{-(t+1)} + p^{-1}. \tag{41}$$

Proof. A validated extractor cannot output a witness on a false instance. Apply Theorem 6.1 and Lemma 5.1. □

The theorem makes no minimum-rank assertion. Low-rank matrices can remain in the space if their short-vector concentration is small. This is sufficient for the witness-encryption proof: on a false instance every nonzero symbolic difference matrix has the required collision bound. With $B = m$ polynomial in N and $t = \lceil \log\log N \rceil$, the term $B^{-(t+1)}$ is $N^{-\Omega(\log\log N)}$ and stays negligible after any fixed polynomial query factor.

7 Encryption, decryption, and correctness

We use Hair–Sahai's encryption and decryption formulas with the single-list space $\mathcal{S}_{x,\tau}$ constructed above. The representation theorem is used to analyze false instances: a table with substantial concentration would give satisfying support. Honest decryption only needs the rank-one table supplied directly by its witness.

7.1 The algorithms

Use the parameters D, t, W, m, B, ℓ_p from (7). Choose a prime $p \in [2^{\ell_p}, 2^{\ell_p+1})$, sample a uniform public seed $\tau \in \mathbb{F}_p$, and compute the fixed-order basis $M_1, \dots, M_k \in \mathbb{F}_p^{m\times D}$ of $\mathcal{S}_{x,\tau}$. Independently sample

$$s \leftarrow \mathbb{F}_p^m, \qquad r \leftarrow [B]^D, \qquad \eta \leftarrow \mathbb{F}_p^k.$$

In a fresh generic-group context publish

$$\begin{aligned} C_{x,\beta} &= (p, \tau, g, X_1, \dots, X_m, Y_1, \dots, Y_k), \\ X_i = g^{s_i}, \qquad Y_j &= \begin{cases} g^{s^\top M_j r}, & \beta = 0, \\ g^{\eta_j}, & \beta = 1. \end{cases} \end{aligned} \tag{42}$$

The derived dimensions and basis are computable from $1^\lambda, x, p, \tau$ and need not be transmitted. Both bits use the same sampling schedule, including unused samples, as specified below.

Given a valid witness w, decryption computes its canonical satisfying assignment and the certificate (a, u, v) of (16). Thus $\sum_{j=1}^k a_j M_j = uv^\top \neq 0$, with $v \in \{0,1\}^D$. Compute

$$H = \prod_{i=1}^m X_i^{u_i}, \qquad Z = \prod_{j=1}^k Y_j^{a_j}, \qquad L_w = (B-1)\,\mathrm{wt}(v). \tag{43}$$

Return zero precisely when $Z \in \{1, H, H^2, \dots, H^{L_w}\}$, and return one otherwise. The identity label is obtained by subtracting a group element from itself. Successive group operations generate the tested list, whose length is at most $(B-1)D+1$. Invalid witnesses and malformed ciphertexts are rejected according to a fixed public convention.

Sampling convention. Both bits use the bounded, seed-independent schedule specified in Appendix C: sample one coordinate for τ, m for s, D for r, and mD for η, discarding the η coordinates past k. On any sampling failure publish a common abort symbol, for which decryption returns zero on valid witnesses and rejects invalid ones. The failure probability is $\delta_{\mathrm{samp}} = 2^{-\Omega(m)}$ as bounded in (59). Conditional on an accepted prime and successful sampling, τ is exactly uniform and the retained numerical samples have their independent uniform laws, also after conditioning on τ. The accepted-prime and abort laws are the same for both payload bits.

Encryption and decryption take polynomial time. The ciphertext contains $m+k+1$ group labels and the two field-sized public values p, τ. Decryption uses $O((m+k)\log p + BD)$ group queries. These bounds count the ciphertext and the group queries of honest decryption. The field arithmetic for reconstructing the basis and finding the witness certificate is additional. Put $E = mD$. With no cached basis, classical Gaussian elimination in Appendix A, followed by the certificate solve, costs at most $O(NE^3)$ field operations for the $J = O(N)$ sparse gate equations. Even if the basis is supplied, a fresh dense certificate solve costs $O(Ek^2)$ operations.

For the two thresholds $t = \lceil \log\log N \rceil$ and $t = \lceil \log N \rceil$, this direct $O(NE^3)$ implementation gives $\widetilde{O}(N^7)$ and $\widetilde{O}(N^{13})$ field operations for our space, respectively. The same count for the D

active columns of Hair–Sahai's all-list space gives $\widetilde{O}(N^{10})$ and $\widetilde{O}(N^{16})$. These are implementation upper bounds, not lower bounds. Hair–Sahai [2, Sections 3.1 and 5.4] prove polynomial time by an iterative Gaussian-elimination algorithm and a final linear solve, without an explicit field-operation count.

Appendix F gives the detailed count and a systematic-basis variant that avoids this linear algebra in decryption. With pivot positions appended to the ciphertext, it uses $O(Nt + W)$ field operations and $O(m + k + W \log p + \sqrt{BD})$ group queries, plus the ordinary bit work of scanning the ciphertext and metadata. No basis must be published in advance. The appendix gives the local bit costs and applies the same optimizations to Hair–Sahai for comparison.

7.2 Simultaneous correctness

Proposition 7.1 (Correctness for all witnesses of one ciphertext)**.** *For every x and β, the probability that decryption fails on some valid witness is at most*

$$\delta_{\mathrm{samp}} + 2^N(mD+1)2^{-\ell_p} = 2^{-\Omega(m)}. \tag{44}$$

Conditional on any accepted prime p and successful sampling, the corresponding bound is $2^N(mD+1)/p$, for every public seed τ.

Proof. For bit zero, the rank-one certificate gives

$$Z = g^{s^\top(\sum_j a_j M_j)r} = g^{(s^\top u)(v^\top r)} = H^{v^\top r}.$$

Since v is Boolean, the integer $v^\top r$ lies in $[0, L_w]$. Thus decryption returns zero for every valid witness, including when $H = 1$.

For bit one and a fixed valid witness, $a \neq 0$, so $a^\top \eta$ is uniform in $\mathbb{F}_p$ and independent of s. Conditional on s, the tested list contains at most $L_w + 1 \leq mD + 1$ group elements. The probability of incorrectly returning zero is at most $(mD+1)/p$.

The circuit encoding retains the original witness and its length, and uses a fixed canonical extension and certificate procedure. There are at most 2^N such satisfying assignments. A union bound therefore proves the simultaneous bound for every seed, regardless of whether that seed satisfies the concentration guarantee. Add the common sampling failure probability and use $p \geq 2^{\ell_p}$. Since $m = (N+1)\binom{2t}{t} \geq 2(N+1)$ and $\ell_p \geq m$, the result is $2^{-\Omega(m)}$. □

In particular, an exceptional seed in Lemma 5.4 does not cause an additional correctness error. The seed exception is used only in the security analysis.

8 Security without a minimum-rank bound

On a false instance, no nonzero table can have the satisfying representation of Theorem 4.1. The concentration alternative of Theorem 6.1 therefore holds for every nonzero matrix in $\mathcal{S}_{x,\tau}$ whenever the public seed is good. This gives a direct generic-group proof of witness encryption: every nontrivial comparison made by the adversary has small collision probability. No extractor is executed in the security experiment.

Theorem 8.1 (Generic-group witness encryption)**.** *The algorithms of Section 7 satisfy Definition 2.1. Fix any false instance x, an auxiliary input independent of the fresh encryption coins,*

and a classical adversary making at most q group queries. Conditional on an accepted prime p and successful sampling, the statistical distance between its views for the two bits is at most

$$\epsilon(q,p) \le \delta_\tau + O\big(H_q^2(B^{-(t+1)} + p^{-1})\big), \qquad H_q = q + m + mD + 2, \tag{45}$$

where $\delta_\tau \le 2^{-m}$ is the exceptional-seed probability in (40)*. Unconditionally, the same bound holds with δ_τ replaced by 2^{-m} and p^{-1} by $2^{-\ell_p}$. Common sampling aborts add no security error. For every polynomial query bound, the advantage is $N^{-\Omega(\log\log N)}$ and hence negligible in λ.*

Proof sketch. Fix a good seed and the public basis. In the symbolic execution, each exponent is represented by a formal affine expression $c + \alpha^\top S + \gamma^\top Z$, where S_i, Z_j stand for the exponents of X_i, Y_j. Distinct formal vectors receive uniform distinct labels. Fixing the adversary and label tapes fixes the list of at most H_q formal vectors independently of the hidden samples s, r, η.

For a nonzero difference with $\gamma \neq 0$, basis independence gives $M_\gamma = \sum_j \gamma_j M_j \neq 0$. Since x is false, Corollary 6.2 gives $\kappa_B(M_\gamma) \le B^{-(t+1)}$. For payload zero the difference evaluates to $c + s^\top(\alpha + M_\gamma r)$, whose vanishing probability is at most $B^{-(t+1)} + p^{-1}$ by Lemma 5.1. Differences with $\gamma = 0$, and all nonzero differences for payload one, vanish with probability at most p^{-1}.

Union-bound over pairs in the fixed symbolic list. Each numerical execution agrees with the common symbolic view until such a collision occurs, so their total distance is $O(H_q^2(B^{-(t+1)} + p^{-1}))$. Arbitrary-string queries add $O(q/p)$, which is absorbed in this bound. The good-seed guarantee is simultaneous over all matrices, including those chosen after seeing τ; averaging over the seed therefore adds δ_τ once, outside the collision union bound. Common sampling aborts agree exactly. The resulting error is negligible for polynomial q by the parameter choice. Appendix D gives the full coupling, arbitrary-string analysis, and proof. □

The proof follows the symbolic-label argument of [2, Theorem 3.5], using the generic-group proof method of [7]. The changed input to that argument is the uniform small-ball estimate for $M_\gamma r$: their proof obtains it from a rank lower bound, whereas ours obtains it from Corollary 6.2. The ciphertext and group operations remain those of Hair–Sahai.

8.1 Longer messages

A polynomial-length message is encrypted one bit at a time using independent encryption coins and independent generic-group contexts. Decrypt each coordinate with the same witness. For a message of length L, a union bound multiplies the simultaneous correctness error by at most L. A hybrid replacing one encrypted bit at a time gives false-instance distinguishing advantage at most

$$L\left[2^{-m} + O\big(H_q^2(B^{-(t+1)} + 2^{-\ell_p})\big)\right]$$

against an adversary making at most q group queries in total. The other contexts have independent initial randomness in each hybrid; adaptive use of their labels is covered by the same arbitrary-string interface. Thus polynomial message length preserves negligible error.

Several ciphertexts with one public seed. The setup variant in Appendix A reuses p, τ and the basis of $\mathcal{V}_\tau$. Conditional on a good seed, the same hybrid argument applies with fresh secret coins and fresh generic-group contexts for every ciphertext. The good-seed event already covers all verifier instances of the fixed size, including ones chosen after seeing τ. For L ciphertexts the resulting bound is

$$2^{-m} + L\,O\big(H_q^2(B^{-(t+1)} + 2^{-\ell_p})\big).$$

Thus seed failure is a setup event charged once. Correctness holds for every seed and its per-ciphertext errors are still union-bounded.

9 From affine-source extraction to witness encryption

The power-evaluation map enters our construction through its preservation of short coordinate patterns. This section states the resulting connection to affine-source extraction independently of that particular map. The premise concerns an explicit linear extractor for a restricted source family; the conclusion gives ciphertext size directly in terms of its output dimension and field size.

Put $D = N + 1$ and recall

$$\mathcal{R}_B = \{a/b \bmod p : |a| < B,\ 1 \le b < B\}.$$

Write U_t for the uniform distribution on $\mathbb{F}_p^t$. The affine sources considered here have the form

$$X = a + Q^\top U_t, \qquad a \in \mathbb{F}_p^D, \quad Q \in \mathcal{R}_B^{t\times D}, \quad \operatorname{rank} Q = t.$$

Each source has min-entropy $t \log p$. The restriction is on the existence of a short basis for its direction space; the affine offset a is unrestricted. A single extractor must work for this entire family without being given Q or a.

Theorem 9.1 (Linear affine extractors give witness encryption)**.** *Let $N \ge 4$, $D = N+1$, $1 \le t \le D$, and $2 \le B < p$, where p is a public prime. Suppose a deterministic algorithm, given (N, t, B, p), computes $G \in \mathbb{F}_p^{D\times t}$ in time polynomial in N and $\log p$, such that the linear map $L_G(z) = G^\top z$ satisfies*

$$\mathrm{SD}\big(L_G(a + Q^\top U_t), U_t\big) \le \frac{1}{3} \tag{46}$$

for every $a \in \mathbb{F}_p^D$ and every full-row-rank $Q \in \mathcal{R}_B^{t\times D}$. Set $W = \binom{2t}{t}$ and $m = DW$. For an NP verifier with a quadratic encoding on N Boolean variables, the construction below has one-bit ciphertext length

$$|C| = O\left(D^2 \binom{2t}{t} \log p\right) \quad \textit{bits.} \tag{47}$$

On a false instance, a classical generic adversary making at most q queries has distinguishing advantage at most

$$O\Big(H^2\big(B^{-(t+1)} + p^{-1}\big)\Big), \qquad H = q + DW + D^2W + 2. \tag{48}$$

Conditional on successful numerical sampling, correctness error is at most $(BD+1)/p$ for each fixed valid witness and at most $2^N(BD+1)/p$ simultaneously for all valid witnesses of one ciphertext. A common sampling abort adds its probability to these correctness bounds and adds no distinguishing advantage.

The algorithms run in polynomial time when B, W, and $\log p$ are polynomially bounded in N. In parameter regimes where the stated correctness and security bounds are negligible, this gives witness encryption in the classical prime-order generic-group model. The same public map G works for all verifier instances of the given size.

Proof. First, the constant-error premise forces exact rank preservation. The distribution $L_G(a + Q^\top U_t)$ is uniform on an affine subspace of dimension $r = \operatorname{rank}(QG)$. Its statistical distance from U_t is $1 - p^{r-t}$. If $r < t$, this distance is at least $1 - 1/p > 1/3$, contradicting (46). Hence QG is invertible for every short t-row matrix in the premise.

This also preserves every full-row-rank $Q \in \mathcal{R}_B^{\nu\times D}$ with $1 \le \nu < t$. Successively append standard coordinate rows outside the current row space until the matrix has t independent rows.

All appended entries lie in $\mathcal{R}_B$. Invertibility of the extended matrix times G implies independence of the images of its original ν rows. Thus, for every full-row-rank $Q \in \mathcal{R}_B^{\nu\times D}$ with $1 \le \nu \le t$,

$$\operatorname{rank}(QG) = \nu. \tag{49}$$

Use the t linear forms and their complete monomial list

$$\ell_j(b) = (G^\top v(b))_j, \qquad h_\alpha(b) = \prod_{j=1}^{t} \ell_j(b)^{\alpha_j}, \qquad \alpha \in \mathbb{Z}^t_{\ge 0}, \quad |\alpha| \le t.$$

The degree-zero monomial supplies the constant weight 1. Construct the weighted assignment tables and impose the verifier tests as in (4) and (5), obtaining a space $\mathcal{S}_{x,G} \subseteq \mathbb{F}_p^{m\times D}$ of dimension $k \le mD$.

The public basis is efficiently computable for these arbitrary forms. Initially $\ell_j(0) = G_{0j}$, and a bit extension gives

$$\ell_j(b + \varepsilon e_i) = \ell_j(b) + \varepsilon G_{ij}.$$

The binomial expansion of every translated weight uses only weights of degree at most t. The dependency-preservation proof and the extend, reduce, retain algorithm of Appendix A therefore apply. Gaussian elimination then imposes the verifier tests and finds the rank-one certificate for an honest witness.

Let $0 \ne M \in \mathcal{S}_{x,G}$ have $H_B(M) < (t+1)\log B$. Lemma 5.2 gives $M = PQ$ with a full-row-rank short Q having $\nu \le t$ rows. By (49), both Q and QG map onto $\mathbb{F}_p^\nu$. Consequently

$$\operatorname{col}(MG) = \operatorname{col}(P) = \operatorname{col}(M), \qquad \operatorname{rank} M \le \nu \le t.$$

Apply Theorem 4.1 with G in place of F_t and $[\mathbf{e}_0 \mid G]$ as the augmented form matrix. Its symmetric-tensor proof works for any coefficient matrix of the t forms and therefore yields a satisfying witness. On a false instance this is impossible, so every nonzero M has $\kappa_B(M) \le B^{-(t+1)}$. The symbolic-label proof of Theorem 8.1 gives (48), with no exceptional-seed term.

Apply the encryption and decryption formulas of Section 7 to this basis. The ciphertext has $m+k+1 \le DW + D^2W + 1$ group labels. Even if G is transmitted explicitly, its Dt field entries are absorbed by (47). The short decryption list has at most $(B-1)D+1$ elements, giving the stated correctness bounds by the proof of Proposition 7.1. Its length is polynomial when B is polynomial. The remaining basis and certificate computations are polynomial in $N, W, \log p$. With a supplied prime, the bounded coordinate samplers of Appendix C can make the common abort probability arbitrarily small using polynomially many attempts whenever the logarithm of its reciprocal is polynomially bounded. □

Extractors already supplied by the construction. The preservation results also have a direct extractor consequence, independent of witness encryption. For $0 < \delta < 1$, the proof of Lemma 5.4 shows that, if

$$p \ge \delta^{-1} N t^2 (2B^2)^{tD},$$

then with probability at least $1-\delta$, one uniform public seed τ makes $F_t(\tau)^\top$ extract all t field elements *exactly*, simultaneously for every short-basis affine source in (46). Here F_t denotes only the t power columns. This is a random setup of a map that is deterministic once fixed; the setup exception is separate from the zero extraction error of every good map.

For every good setup map, extraction has zero error and retains all $t \log p$ bits of source entropy. This guarantee concerns a restricted source family over a larger field, with an exceptional setup

probability. Gabizon–Raz's deterministic extractor covers arbitrary affine sources over fields of polynomial size and extracts almost all the entropy with small error [8, Theorems 1–2]; their seeded extractor already outputs all t field elements [8, Theorem 5]. Our simultaneous guarantee for one fixed map uses the restriction on source directions and the larger field.

The public seed can also be fixed deterministically over a larger field by clearing denominators and bounding the integer coefficients of the preservation determinants. This yields an explicit perfect extractor for the restricted family, but we do not develop that deterministic variant or its field-size bounds here.

Field size and a possible parameter improvement. The extractor premise imposes a necessary field-size constraint:

$$t \log p \geq D \log B. \tag{50}$$

Indeed, if distinct $r, r' \in [B]^D$ had $G^\top r = G^\top r'$, then $(r - r')^\top$ would be a nonzero short one-row matrix annihilated by G, contradicting (49). Thus L_G is injective on $[B]^D$, and $p^t \geq B^D$.

For $t = \lceil \log N \rceil$ and $B = N^3$, a construction meeting the premise over a field with $\log p = \Theta(N)$ would give

$$|C| = O\left(\frac{N^5}{\sqrt{\log N}}\right), \tag{51}$$

since $W = \Theta(N^2/\sqrt{\log N})$. The lower bound (50) already forces $\log p \geq (3 - o(1))N$ in this regime. Choose the sampling-abort probability to be $2^{-\Omega(N)}$. Then the simultaneous correctness error is $2^{-\Omega(N)}$, and (48) is at most $N^{-\log N}$ for $q \leq N^{\log N}$ and sufficiently large N. Compared with the present one-list choice $\ell_1 = \Theta(N(\log N)^2)$ in (62), this would save a factor $\Theta((\log N)^2)$ in the explicit ciphertext upper bound. The existence of an efficient extractor at this field size is an additional hypothesis; the theorem does not supply one.

Scope of the connection. The premise is stronger than a guarantee that most seeds work for each fixed source. It specifies one efficiently constructed map that works simultaneously for the entire short-basis family. If a randomized setup instead produces such a map except with probability δ, the same argument adds δ to (48). Gabizon–Raz's affine-source construction uses a rank-preserving map with a per-source good-seed guarantee [8, Lemma 6.1 and Theorem 5]; our simultaneous short-pattern lemma supplies the stronger setup guarantee needed here, at the field size stated in the main construction. Linearity and preservation of the full source dimension are the properties used to reconstruct column coordinates and to obtain translation-closed weights. A general nonlinear affine extractor, or one that loses source dimension, does not by itself provide these properties. The distribution Mr governing witness-encryption security remains an image of a short product distribution; it need not be an affine source.

10 Efficiency and comparison with Hair–Sahai

10.1 Ciphertext size and decryption queries

For $t = \lceil \log \log N \rceil$ and $N \geq 4$,

$$W = \binom{2t}{t} = \Theta(4^t/\sqrt{t}), \qquad m = (N+1)W = O\big(N(\log N)^2\big).$$

The prime-length parameter in (7) satisfies

$$\ell_p = m + O(tN \log m) = \Theta(m),$$

since $tN \log m = o(m)$. Thus the field size needed to make the public seed simultaneously good still gives $O(m)$-bit group labels. The seed itself costs $O(m)$ bits.

The rectangular matrix space has dimension $k \le mD$, where $D = N + 1$. The encryptor transmits $m + k + 1 = O(mD)$ labels and the public parameters p, τ. The basis is computed from x, p, τ and is not transmitted. The decryption products use $O((m+k)\log p)$ group queries, and the search uses at most $(B-1)D = O(mD)$ further queries. Consequently

$$\begin{aligned} \#\text{ciphertext labels} &= O\big(N^2(\log N)^2\big), \\ |C_{x,\beta}|, q_{\mathrm{dec}} &= O(Nm^2) = O\big(N^3(\log N)^4\big). \end{aligned} \tag{52}$$

The statement x is a public input; including it in the ciphertext does not change these bounds. They count ciphertext bits and decryption group queries. Computing the matrix basis, forming the ciphertext exponents, and finding the witness certificate all have polynomial bit complexity, but are not asserted to satisfy the same cubic bound.

The tables in this section use the direct decryption algorithm of Section 7. Appendix F quantifies its basis-reconstruction cost and compares optimized decryption for both constructions, using systematic coordinates, grouping of repeated weights, and baby-step/giant-step. That variant adds pivot metadata and avoids reconstructing a basis in decryption.

A further variant can use a public random mixing of the forms and replace many full blocks by vector rows, reducing the row count and ciphertext size. We do not explore that variant here; the bounds below describe the complete-block construction.

These preprocessing costs can be amortized. At fixed N, p, t, τ, the basis of $\mathcal{V}_\tau$ is independent of the instance and may be public and reused across encryptions. The verifier tests and basis of $\mathcal{S}_{x,\tau}$ depend on x, but can also be cached when the same instance is reused. Thus a computation from only the numerical parameters should be distinguished from decryption with cached public preprocessing. Appendix A explains why one public random seed suffices for this reuse. The same amortization of the verifier-independent span applies to Hair–Sahai.

10.2 What changes in the deterministic construction

The direct comparison is with [2, Proposition 5.3 and Sections 5.2–5.4]. At the same threshold t, their deterministic construction uses $2Nt + 1$ folded lists and hence

$$m_{\mathrm{HS}} = (N+1)(2Nt+1)\binom{2t}{t} = (2Nt+1)m$$

rows. These samples guarantee that every nonzero matrix of rank at most t has a list preserving both its full column span and a nonzero block. The Boolean moment argument yields a witness, giving minimum nonzero rank greater than t on a false instance for every permitted prime.

Our construction retains their complete degree-bounded weight list, but uses the evaluation points $\tau, \tau^2, \ldots, \tau^t$ at one random parameter. Lemmas 5.2 and 5.4 show that it suffices to preserve the patterns with short rational coordinates; every other matrix already has small concentration. The large field makes the exceptional-seed probability exponentially small. Theorem 6.1 gives extraction or small concentration, without asserting a growing minimum rank.

Theorem 4.1 makes the algebraic step explicit: when the single list preserves the column space of $M \in \mathcal{S}_{x,\tau}$ of rank $\rho \le t$, the matrix is a linear combination of at most 2^ρ satisfying assignment tables. The proof constructs scalar moments in ρ coordinates, derives their Boolean and verifier identities, and applies Boolean interpolation. Tensor symmetry extends the column relations to every slot and reconstructs all stored blocks from these same moments. This organizes the separation

argument underlying Hair–Sahai's proof as an explicit moment representation. For their multiple-list construction the same argument reconstructs the selected list's blocks, whose nonvanishing suffices for extraction.

10.3 A smaller field for the deterministic baseline

The field size should be chosen to meet each construction's correctness and security conditions. In particular, the deterministic construction need not use field elements whose bit length is proportional to its row count. Hair–Sahai's Theorem 3.1 fixes $p \in [2^{m_{\mathrm{HS}}}, 2^{m_{\mathrm{HS}}+1})$, but the underlying matrix construction only requires

$$p_{\mathrm{HS}} > \max\{2^N, 2Nt\}$$

at threshold t [2, Equation (10) and Proposition 5.3]. We therefore reparameterize their construction using the correctness and collision proofs in [2, Sections 3.2–3.3].

Keep the source's grid size $B_{\mathrm{HS}} = m_{\mathrm{HS}}$ and choose

$$\begin{aligned} \ell_{\mathrm{HS}} &= 2N + \lceil \log(m_{\mathrm{HS}} D + 2Nt + 2) \rceil , \\ p_{\mathrm{HS}} &\in [2^{\ell_{\mathrm{HS}}}, 2^{\ell_{\mathrm{HS}}+1}). \end{aligned} \tag{53}$$

Since m_{HS} is polynomial in N, this gives $\ell_{\mathrm{HS}} = \Theta(N)$. The chosen prime satisfies the matrix conditions and $p_{\mathrm{HS}} > B_{\mathrm{HS}}$. As throughout this paper, use the D potentially nonzero columns and the Boolean right factor of weight at most D. Decryption therefore searches at most $B_{\mathrm{HS}} D + 1$ group elements. Conditional on successful sampling, the probability of failure for any valid witness of one ciphertext is at most

$$\frac{2^N (B_{\mathrm{HS}} D + 1)}{p_{\mathrm{HS}}} \le 2^{-N}.$$

The bounded prime and coordinate samplers add $2^{-\Omega(N)}$ failure probability by the argument in Appendix C, with ℓ_{HS} in place of ℓ_p.

For a false instance, every nonzero matrix has rank at least $t + 1$. Writing $H_{\mathrm{HS}} = q + m_{\mathrm{HS}} + m_{\mathrm{HS}} D + 2$, the generic collision proof gives distinguishing advantage at most

$$O\left(H_{\mathrm{HS}}^2 \left(B_{\mathrm{HS}}^{-(t+1)} + p_{\mathrm{HS}}^{-1}\right)\right) = N^{-\Omega(\log \log N)}$$

for every polynomial q. Common sampling aborts are independent of the message. Thus the smaller prime retains both WE correctness and negligible distinguishing advantage. This conclusion uses the underlying proofs with the new field choice, rather than the fixed prime range in the statement of Hair–Sahai's Theorem 3.1.

The basis dimension satisfies $k_{\mathrm{HS}} \le m_{\mathrm{HS}} D$. Consequently ciphertext bits and decryption group queries are bounded by

$$\begin{aligned} O\big((m_{\mathrm{HS}} + k_{\mathrm{HS}})\ell_{\mathrm{HS}} + B_{\mathrm{HS}} D\big) &= O(m_{\mathrm{HS}} D N) \\ &= O\big(N^4 t (\log N)^2\big) = \widetilde{O}(N^4). \end{aligned} \tag{54}$$

The ciphertext bound omits the search term, which is smaller. Thus the comparison at threshold $t = \lceil \log \log N \rceil$ is between upper bounds $\widetilde{O}(N^4)$ and $\widetilde{O}(N^3)$. The improvement comes from using one list in place of $2Nt + 1$ lists. It does not charge the deterministic baseline an additional field-length factor proportional to its row count.

The following table uses the smaller field for both deterministic rows and our stated field choice for the random-list row. In the first row, replace t by $\lfloor \log N \rfloor$ in (53); its field length is still $\Theta(N)$. The notation $\widetilde{O}$ suppresses factors polynomial in $\log N$.

Construction	Rows	Field bits	Ciphertext bits / decryption queries
All lists, $\lfloor \log N \rfloor$	$\widetilde{O}(N^4)$	$\Theta(N)$	$\widetilde{O}(N^6)$
All lists, $\lceil \log \log N \rceil$	$\widetilde{O}(N^2)$	$\Theta(N)$	$\widetilde{O}(N^4)$
One random list, $\lceil \log \log N \rceil$	$\widetilde{O}(N)$	$\widetilde{O}(N)$	$\widetilde{O}(N^3)$

These are upper bounds for the stated one-list construction and reparameterizations of the Section 5 deterministic family in the supplied STOC draft. The first row's logarithmic threshold still supports a stronger quasipolynomial query regime: the p_{HS}^{-1} term remains exponentially small in N. Their Section 6 and Theorem 6.1 provide arbitrary fixed polylogarithmic rank gaps. The two lower rows target negligible advantage against every polynomial query bound, with bounds of the form $N^{-\Omega(\log \log N)}$. Both give exponentially small correctness error in N; our stated parameters give the stronger $2^{-\Omega(m)}$ bound. Section 10.5 also considers smaller fields when this correctness target is relaxed.

10.4 A matched comparison at stronger security

The one-list construction also supports a stronger quantitative guarantee. For a matched comparison, require distinguishing advantage at most $N^{-\log N}$ against every classical generic adversary making at most $N^{\log N}$ group queries. Set

$$t = \lceil \log N \rceil, \qquad D = N + 1, \qquad W = \binom{2t}{t}, \qquad B = N^3$$

in both constructions. This is a separate parameter choice from (7): the grid size and field length are chosen independently of the row count. The one-list construction has $m_1 = DW$ rows, and the deterministic Section 5 construction of Hair–Sahai has $m_{\mathrm{HS}} = (2Nt + 1)m_1$ rows.

Appendix E gives explicit primes and proves the following bounds. Both rows have the same security target and simultaneous correctness error $2^{-\Omega(N)}$.

Construction	Rows	Field bits	Ciphertext bits / decryption queries
Hair–Sahai, all lists	$\widetilde{O}(N^4)$	$\Theta(N)$	$\widetilde{O}(N^6)$
One random list	$\widetilde{O}(N^3)$	$\Theta(N(\log N)^2)$	$\widetilde{O}(N^5)$

For both rows the collision term is $B^{-(t+1)}$. Even after the quadratic loss for $N^{\log N}$ queries, it contributes only $O(N^{-\log N - 3})$; the field terms and our exceptional-seed probability are exponentially small in N. Thus the advantage is at most $N^{-\log N}$ for sufficiently large N. All honest algorithms remain polynomial time. The algebraic candidate search also remains polynomial time since $2^t \le 2N$.

This comparison shows that the saving from one list persists at stronger security while retaining simultaneous correctness error $2^{-\Omega(N)}$. These are explicit matched upper bounds for our one-list family and the deterministic Section 5 family, not an optimization over all choices of t, B or over all constructions in Hair–Sahai. The costs count ciphertext bits and decryption group queries; the additional linear algebra has polynomial bit complexity as before.

10.5 High-order folding bases and smaller fields

Suppose a prime p is supplied with an element $\gamma \in \mathbb{F}_p^*$ whose multiplicative order exceeds N. Such an element can, for example, be obtained from a generator when the factorization of $p-1$ is known.

One may then use geometric folding,

$$\ell_j^{(\gamma)}(b) = \sum_{i=0}^{N} (\gamma^{j-1}\tau)^i b_i.$$

The leading coefficient in the folded determinant is a nonzero Vandermonde determinant in the distinct values γ^{e_i}, as in [2, Lemma 5.2]. Its degree in τ is at most $N\nu$, compared with $N\nu(\nu+1)/2$ for powers of the seed. For one random list this gives the seed bound

$$\delta_\tau^{(\gamma)} \le \frac{N}{p} \sum_{\nu=1}^{t} \nu A_B^{\nu D}.$$

The corresponding bound for powers of τ is larger by at most $(t+1)/2$, so optimizing this prefactor can save only $O(\log t)$ field bits. The common bound in (40) and both preceding one-list field choices cover either formulation. For these one-list primes, $\gamma = 2$ already has order greater than N, so this prefactor improvement itself needs no special prime.

For the deterministic all-list construction, a known high-order γ replaces the algebraic condition $p > 2^N$ imposed by the choice $\gamma = 2$. The same $2Nt + 1$ sample points suffice, provided $p > 2Nt$. This removes an exponential field restriction from the matrix construction itself. The Gabizon–Raz power map used here already avoids that restriction without an order condition. At the exponentially small correctness targets used above, however, the correctness and one-list seed estimates still require the stated field lengths; there is no further improvement in their polynomial cost exponents.

In fact, finding an element of order greater than N does not require a special prime or factoring $p - 1$. For $p > N(N+1)/2 + 1$, test the first $N(N+1)/2 + 1$ nonzero field elements, rejecting a candidate if one of its first N powers is one. At most $\sum_{d=1}^{N} d = N(N+1)/2$ elements can fail, so a suitable γ is found in $O(N^3)$ field operations. This is additional polynomial-time arithmetic; it uses no group queries.

The gain with ordinary per-witness correctness. Definition 2.1 requires correctness for each fixed valid witness. If its error target is $N^{-\log N}$, rather than the exponentially small simultaneous bound used above, geometric folding allows a smaller-field deterministic baseline. At $t = \lceil \log N \rceil$ and $B = N^3$, choose

$$\ell_\gamma = \left\lceil 4(\log N)^2 \right\rceil, \qquad p_\gamma \in [2^{\ell_\gamma}, 2^{\ell_\gamma+1}),$$

and find γ by the preceding search. This variant has distinguishing advantage at most $N^{-\log N}$ against $N^{\log N}$ group queries and correctness error at most $N^{-\log N}$ for each fixed valid witness, for sufficiently large N. Appendix E.5 proves these claims, including bounded sampling.

The all-list ciphertext and decryption bounds then become $\widetilde{O}(N^5)$, matching the polynomial exponent of our one-list bound. More precisely, writing $L = \log N$, the explicit parameter choices give

One random list:	$O(N^5 L^{3/2})$,
All lists with a high-order base:	$O(N^5 L^{5/2})$

for both ciphertext bits and decryption group queries. The saving in these upper bounds is logarithmic. At the weaker threshold $t = \lceil \log \log N \rceil$, keeping $B_{\mathrm{HS}} = m_{\mathrm{HS}}$ and using the same polylogarithmic field length similarly gives $\widetilde{O}(N^3)$ all-list costs, with negligible per-witness correctness error and security against every polynomial query bound. Thus the factor-N comparisons above depend on the stronger correctness benchmark. Merely dropping simultaneity while retaining correctness error $2^{-\Omega(N)}$ does not yield these smaller fields under the correctness estimate used here.

10.6 The role of the weight family

The efficient basis computation still uses the fact that assigning one more bit translates each linear form by a known constant. Complete degree-bounded lists are closed under these translations. The scalar-moment proof uses the same completeness to obtain all polynomials of the required degree in the reduced coordinates. Boolean interpolation expresses the scalar functional as a linear combination of evaluations at satisfying candidates, and the common column relations reconstruct the weighted blocks. These properties do not require multiple samples. The multiple samples in the deterministic construction address the separate requirement of preserving every low-rank column space and a nonzero block at the same sample.

The algebraic proof interpretation is also shared with the source. Hair–Sahai explicitly connect the multiplied equations and their dual weighted sums to bounded-degree Nullstellensatz and polynomial calculus. Our scalar functional need not be positive or normalized; its Boolean interpolation is an algebraic identity over the field. Neither argument assumes that arbitrary false instances have low-degree refutations in the original variables. Hair–Sahai also credit earlier rank-preserving maps for their folded-minor lemma. We use the power-evaluation map of Gabizon–Raz [8], with a reformulation of their leading-degree argument; Appendix G describes other univariate seeded rank condensers. Our use of concentration changes which column spaces must be preserved. These separate roles can also be formulated as general criteria on weight families for basis computation, extraction, and concentration. We do not pursue that broader characterization here.

11 Acknowledgment

This paper is a result of extensive collaboration between the author and GPT6-Astra for over a month.

References

[1] S. Garg, C. Gentry, A. Sahai, and B. Waters. *Witness Encryption and its Applications*. STOC 2013, pp. 467–476. `https://eprint.iacr.org/2013/258`.

[2] I. M. Hair and A. Sahai. *Witness Encryption via Prime-Order Generic Groups*. Revised manuscript, 2026; section numbering follows the STOC draft. Public preprint: arXiv:2609.18275. `https://arxiv.org/abs/2609.18275`.

[3] P. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, and P. Pudlák. *Lower Bounds on Hilbert's Nullstellensatz and Propositional Proofs*. Proceedings of the London Mathematical Society 73(1), pp. 1–26, 1996. `https://doi.org/10.1112/plms/s3-73.1.1`.

[4] M. Clegg, J. Edmonds, and R. Impagliazzo. *Using the Groebner Basis Algorithm to Find Proofs of Unsatisfiability*. STOC 1996, pp. 174–183. `https://doi.org/10.1145/237814.237860`.

[5] B. Buchberger and H. M. Möller. *The Construction of Multivariate Polynomials with Preassigned Zeros*. In J. Calmet (ed.), *Computer Algebra (EUROCAM '82)*, Lecture Notes in Computer Science 144, pp. 24–31. Springer, 1982.

[6] S. R. Buss. *Lower Bounds on Nullstellensatz Proofs via Designs*. In P. Beame and S. R. Buss (eds.), *Proof Complexity and Feasible Arithmetics*, DIMACS Series 39, pp. 59–71. American Mathematical Society, 1998. `https://doi.org/10.1090/dimacs/039/04`.

[7] V. Shoup. *Lower Bounds for Discrete Logarithms and Related Problems*. Advances in Cryptology—EUROCRYPT 1997, pp. 256–266. https://www.shoup.net/papers/dlbounds1.pdf.

[8] A. Gabizon and R. Raz. *Deterministic Extractors for Affine Sources over Large Fields*. Combinatorica 28(4), pp. 415–440, 2008. Preliminary version in FOCS 2005. https://doi.org/10.1007/s00493-008-2259-3.

[9] M. A. Forbes and A. Shpilka. *On Identity Testing of Tensors, Low-Rank Recovery and Compressed Sensing*. STOC 2012, pp. 163–172. https://doi.org/10.1145/2213977.2213995.

[10] M. A. Forbes and V. Guruswami. *Dimension Expanders via Rank Condensers*. APPROX/RANDOM 2015, LIPIcs 40, pp. 800–814. https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2015.800.

[11] P. Erdős. *On a Lemma of Littlewood and Offord*. Bulletin of the American Mathematical Society 51, pp. 898–902, 1945.

[12] T. Tao and V. Vu. *Inverse Littlewood–Offord Theorems and the Condition Number of Random Discrete Matrices*. Annals of Mathematics 169(2), pp. 595–632, 2009.

A Computing a basis without enumerating assignments

We present the basis algorithm of Hair–Sahai [2, Section 5.4] as three repeated steps: *extend, reduce, retain.* Fix p and τ. The powers-of-τ formulation makes the extension rule explicit, and the algorithm works for every seed. It first computes $\mathcal{V}_\tau$, the span of all assignment tables; the verifier tests are imposed only afterward.

Recall that $b_0 = 1$, $D = N + 1$, and

$$\ell_j(b) = \sum_{r=0}^{N} b_r \tau^{rj}, \qquad h_\alpha(b) = \prod_{j=1}^{t} \ell_j(b)^{\alpha_j}, \qquad |\alpha| \le t.$$

Let $\mathbf{h}(b) \in \mathbb{F}_p^W$ list these weights. Listing the entries of the assignment table in block order, row by row, gives

$$\begin{aligned} \Phi(b) := \operatorname{vec} A_\tau(b) &= (h_\alpha(b) b_a b_c)_{\alpha,a,c} \\ &= \mathbf{h}(b) \otimes v(b) \otimes v(b) \in \mathbb{F}_p^{mD}. \end{aligned} \tag{55}$$

The algorithm evaluates these coordinates directly. We will retain actual assignments whose table vectors form a basis.

A.1 Extend, reduce, retain

Represent a prefix by an N-bit vector with all unassigned bits zero. Start with the single all-zero assignment, whose table is nonzero because $b_0 = 1$ and $h_0 = 1$; initially every ℓ_j equals one. At stage $i = 1, \dots, N$, start an empty new assignment list and an empty auxiliary echelon basis, and process candidates in a fixed order.

1. *Extend.* For each prefix b retained from the preceding stage, form both $b' = b + \varepsilon e_i$, with $\varepsilon \in \{0, 1\}$. Here $b_i = 0$ before extension. Update the forms using

   $$\ell_j(b') = \ell_j(b) + \varepsilon \tau^{ij}, \qquad 1 \le j \le t,$$

   and evaluate the new table vector $\Phi(b')$.

2. *Reduce.* Use Gaussian elimination over $\mathbb{F}_p$ to reduce $\Phi(b')$ against the candidates already accepted at this stage. A zero remainder means that the candidate contributes no new direction, so discard it.

3. *Retain.* For a nonzero remainder, retain the original assignment b' and its original table $\Phi(b')$. Add the normalized remainder to the auxiliary echelon basis used for subsequent reductions.

The echelon vectors serve only to test independence. The next stage extends the retained assignments, whose original tables need not be in echelon form.

At most mD assignments survive each stage. Thus each stage handles at most $2mD$ candidate vectors, each of length mD. After N stages their tables form the desired basis. The following argument explains why discarding a dependent table cannot lose a direction that would appear only after assigning later bits.

A.2 Why every common extension preserves dependencies

For fixed i, ε, the increments

$$c_j = \varepsilon \tau^{ij}$$

are independent of the prefix. The binomial formula gives

$$h_\alpha(b') = \sum_{\beta \le \alpha} \left(\prod_{j=1}^{t} \binom{\alpha_j}{\beta_j} c_j^{\alpha_j - \beta_j} \right) h_\beta(b).$$

All weights on the right are stored because $\beta \le \alpha$ implies $|\beta| \le t$. For example, when $t \ge 2$ and the new bit is one,

$$(\ell_1 + \tau^i)(\ell_2 + \tau^{2i}) = \ell_1\ell_2 + \tau^{2i}\ell_1 + \tau^i\ell_2 + \tau^{3i}.$$

Every term is an available weight, including the constant term.

The other two table factors cause no difficulty: expanding $b'_a b'_c = (b_a + \varepsilon\delta_{a,i})(b_c + \varepsilon\delta_{c,i})$ introduces only $b_a b_c$, b_a, b_c, and 1. The linear and constant factors are already present among the table coordinates because $b_0 = 1$. Consequently every coordinate of $\Phi(b')$ is a fixed linear combination of coordinates of $\Phi(b)$, with coefficients depending only on i, ε, p, τ.

In matrix notation, let $C_{i,\varepsilon}$ be the weight-update matrix given by the binomial formula, and put $R_{i,\varepsilon} = I_D + \varepsilon e_i e_0^\top$. The corresponding table update is

$$\begin{aligned} \Phi(b + \varepsilon e_i) &= U_{i,\varepsilon}\Phi(b), \\ U_{i,\varepsilon} &= C_{i,\varepsilon} \otimes R_{i,\varepsilon} \otimes R_{i,\varepsilon}. \end{aligned} \tag{56}$$

These matrices justify the update; the algorithm need not construct them. In particular, a relation among prefixes propagates with exactly the same coefficients:

$$\Phi(b) = \sum_a \lambda_a \Phi(b^{(a)}) \quad \Longrightarrow \quad \Phi(b + \varepsilon e_i) = \sum_a \lambda_a \Phi(b^{(a)} + \varepsilon e_i).$$

This applies when bit i is still unassigned in all these prefixes. Applying it successively proves the same assertion for every common suffix. This is the reason that a dependent prefix can be discarded.

Formally, after stage i the retained tables form a basis of

$$\mathcal{V}_i = \operatorname{span}\{A_\tau(b_1, \ldots, b_i, 0, \ldots, 0) : (b_1, \ldots, b_i) \in \{0,1\}^i\}. \tag{57}$$

The invariant holds at stage zero. If it holds at stage $i-1$, express any prefix table in the retained basis. Dependency preservation expresses either extension in the corresponding retained extensions. The candidates therefore span $\mathcal{V}_i$, and Gaussian elimination selects a basis of that span. At stage N, $\mathcal{V}_N = \mathcal{V}_\tau$.

A.3 Imposing the verifier tests

The preceding computation is independent of x. For fixed N, p, t, τ, its basis of $\mathcal{V}_\tau$ can be published or cached once and reused across all verifier instances padded to the same number of variables. Only the following step depends on the verifier equations.

Finally, let $V_1, \ldots, V_s$ be this basis and form

$$(E_x)_{(\alpha,j),a} = T_{q_j}(C_\alpha(V_a)).$$

Then $\sum_a z_a V_a \in \mathcal{S}_{x,\tau}$ if and only if $E_x z = 0$. For a basis $z^{(1)}, \ldots, z^{(k)}$ of $\ker E_x$, set $M_j = \sum_a z_a^{(j)} V_a$. These matrices are an independent basis of $\mathcal{S}_{x,\tau}$. Direct evaluations and Gaussian elimination take time polynomial in $N, m, \log p$; fixed assignment orders and pivot rules make the output determined by x, p, τ. Given a satisfying assignment, evaluate $A_\tau(b)$ and solve for its coordinates in this basis to obtain (16). For repeated encryptions of the same x, this verifier-dependent basis and its coordinate-solving data can also be cached. Appendix F gives explicit field-operation counts and a systematic-basis variant whose honest decryptor reads the certificate from pivot coordinates instead of repeating these computations.

Reusing the public seed. Section 7 samples p, τ in each encryption, but they may instead be generated once as reusable public parameters for fixed N, t, B. Lemma 5.4 holds simultaneously for all verifier instances with these parameters, including instances chosen after seeing τ. Publishing the deterministically computed basis of $\mathcal{V}_\tau$ introduces no additional randomness or secret information. With fresh s, r, η and a fresh generic-group context for each ciphertext, the security proof applies conditional on this one good-setup event. Thus a multiple-ciphertext bound charges the exceptional-seed probability once; the conditional collision errors are combined by the usual hybrid argument. This reuse does not require detecting whether the sampled seed is good. The verifier-independent span in Hair–Sahai's deterministic construction admits the same basis preprocessing and reuse.

A.4 The polynomial-evaluation viewpoint

There is a useful connection to the linear-algebra step of the Buchberger–Möller algorithm [5]. For fixed p, τ, consider the Boolean coordinate ring and the space of stored polynomial functions

$$\mathcal{A} = \mathbb{F}_p[b_1, \ldots, b_N]/\langle b_1^2 - b_1, \ldots, b_N^2 - b_N \rangle,$$
$$\mathcal{W}_\tau = \mathrm{span}_{\mathbb{F}_p}\{h_\alpha b_a b_c : |\alpha| \le t,\ 0 \le a, c \le N\} \subseteq \mathcal{A}.$$

The vector $\Phi(b)$ records evaluation at b on this generating list. These evaluations span $\mathcal{W}_\tau^*$: a function in $\mathcal{W}_\tau$ annihilated by every Boolean evaluation is zero. Thus $\mathcal{V}_\tau$ identifies with $\mathcal{W}_\tau^*$, and the algorithm selects a basis of evaluation functionals.

At the evaluation-matrix level, Buchberger–Möller fixes points and selects independent monomial-evaluation columns. Here the polynomial list is fixed, and we select independent point-evaluation rows. The known bit-update maps let us generate those rows incrementally without enumerating the Boolean cube. The reduction step is Gaussian elimination; the algorithm does not compute an ideal's Gröbner basis or form S-polynomials. In particular, $\mathcal{W}_\tau$ is generally not closed under multiplication, so it should not be treated as a quotient algebra of dimension at most mD. The Boolean algebra of dimension 2^ρ used in Theorem 4.1 arises only after column preservation compresses the moments of a particular matrix. It is not an algebra structure assumed for this universal basis computation.

The powers of τ make the increments $\varepsilon\tau^{ij}$ explicit, but τ is a fixed field element during this computation. The same dependency argument works for any fixed list of linear forms with the complete degree-t weight list, including the forms of Hair–Sahai. The viewpoint explains the existing polynomial-time basis algorithm; it does not assert a further asymptotic speedup.

B Compressed moments, representation, and Nullstellensatz

The full proof of Theorem 4.1 is in the main text. This appendix explains its degree bounds, the distinction between original and reconstructed coefficients, and the role of the weighted verifier equations.

B.1 Why the available degrees suffice

The table stores weights of degree at most t in ℓ and two additional factors from $v(b)v(b)^\top$. After column compression, it therefore determines the scalar functional μ through degree $t+2$. The compressed forms w_j are linear combinations of ℓ; an arbitrary linear form in all N bits is not automatically an available weight.

The degree accounting in the proof is as follows:

operation	largest weight degree	largest scalar moment degree
Boolean reduction $g(y)(y_j^2 - y_j)$	t	$t+2$
Boolean interpolation coefficient $\mu(\chi_e)$	ρ	ρ
normalization multiplied by χ_e	ρ	$\rho+1$
Boolean or verifier equation multiplied by χ_e	ρ	$\rho+2$
symmetric tensor projection and reconstruction	t	$t+2$

Every entry is within the stored moments because $\rho \le t$. The argument works also when $\rho = t = 1$ and uses no divisions by factorials or assumptions on the characteristic.

For reconstruction, adjoining the constant form $\ell_0 = 1$ writes every weight homogeneously at degree t. The ordered tensor table $\widetilde{M}$ in (23) duplicates the stored monomial blocks; it introduces no multinomial coefficients. Projecting its two remaining slots by $F^\top$ gives a fully symmetric tensor G of order $t+2$. The common column relations fix its last slot under a projection, and symmetry fixes every other slot under the same projection. Thus (24) and (26) reconstruct every stored block using only the degree-$(t+2)$ compressed tensor Z. These linear transformations do not increase the moment degree. The ordered expansion is proof notation; the construction retains the compact monomial list.

The scalar functional μ is still defined through every degree up to $t+2$ using the original table. In particular, its lower moments and the coefficients $\mu(\chi_e)$ are available even though Z directly records homogeneous degree-$(t+2)$ moments. Reconstruction gives $M \neq 0 \Longrightarrow Z \neq 0$, so Boolean interpolation has nonempty support; $\mu(1)$ alone may be zero.

B.2 Changing the compressed coordinates

Use the notation of the proof: $a = t+1$, $Q = RF$, $QW = I_\rho$, and $U = FW$. For any invertible $\bar{F} \in \mathbb{F}_p^{\rho\times\rho}$, put

$$V = W\bar{F}^\top \in \mathbb{F}_p^{a\times\rho}.$$

Then

$$\bar{F}U^\top = \bar{F}W^\top F^\top = V^\top F^\top, \qquad z(b) = \bar{F}w(b) = V^\top \ell(b).$$

Every polynomial weight of degree at most t in z is therefore available from the original complete list. This identity uses no inverse of the generally rectangular matrix W.

The same $V^\top$ can be applied in all tensor slots after both endpoints have been projected into form coordinates:

$$G = (I_{a^t} \otimes F^\top)\widetilde{M}F,$$
$$\bar{Z} = (V^\top)^{\otimes(t+1)}GV = \bar{F}^{\otimes(t+1)}Z\bar{F}^\top.$$

Thus $\bar{Z}$ is the order-$(t+2)$ moment tensor of z. Writing $\bar{Q} = \bar{F}^{-\top}Q$, the reconstruction identities are

$$G = (\bar{Q}^\top)^{\otimes(t+1)}\bar{Z}\bar{Q},$$
$$\widetilde{M} = \big((\bar{Q}^\top)^{\otimes t} \otimes R^\top\bar{F}^{-1}\big)\bar{Z}\bar{F}^{-\top}R.$$

An invertible Vandermonde matrix is one possible choice of $\bar{F}$; the identity matrix also suffices. Such a change preserves the polynomial degrees. The Boolean moment identities give a reconstructed representation supported on $\{0,1\}^\rho$ in w-coordinates; in z-coordinates its support lies in $\bar{F}\{0,1\}^\rho$.

B.3 Why the original coefficients need not have valid support

Suppose $M = \sum_b \lambda_b A_\tau(b)$ and $M\theta = 0$. The resulting equations

$$\sum_b \lambda_b h_\alpha(b) v(b)\, \theta^\top v(b) = 0 \qquad \text{for all } \alpha$$

are moment equations. They do not imply $\lambda_b \theta^\top v(b) = 0$ separately for each assignment. The map taking a full tensor of assignment moments to the stored weighted moments is generally not injective. Column preservation makes compression reversible on the relevant row space of M; it does not invert this map on all tensors or all coefficient lists.

Here is an explicit cancellation within our weight family. Assume $N \geq t+3$, vary the first $t+3$ bits, and fix the remaining bits to zero. On this subcube set

$$\kappa_b = (-1)^{b_1+\cdots+b_{t+3}},$$

and put $\kappa_b = 0$ outside it. Every table entry $h_\alpha(b) b_i b_j$ has Boolean degree at most $t+2$. Summing against κ annihilates each entry, since each multilinear monomial omits at least one of the varied bits. Thus

$$\sum_b \kappa_b A_\tau(b) = 0$$

for every seed. The rank-one matrix $M = A_\tau(0)$ consequently also has coefficients $\lambda_b = \mathbf{1}_{b=0} + \kappa_b$. Its columns $1, 2, 3$ vanish, but at $b = e_1$ the coefficient is -1 and $b_1 + b_2 + b_3 = 1$. Moreover $MF_t(\tau)$ has the same column space as M, so preservation does not repair the pointwise assertion about these original coefficients.

The representation theorem instead supplies

$$\omega_e = \sum_b \lambda_b \chi_e(w(b)), \qquad M = \sum_{e:\omega_e \neq 0} \omega_e A_\tau(b_e).$$

These are interpolation coefficients of the stored moments. They are not generally the marginals $\sum_{b:(b_{i_1},\ldots,b_{i_\rho})=e} \lambda_b$: the forms $w(b)$ need not be Boolean, or equal to those selected bits, on the original support. The weighted moment identities, rather than pointwise substitutions in that support, justify the new representation. All column relations and verifier equations do hold on its nonzero support.

B.4 How weighted verifier equations constrain support

The representation obtained from Boolean interpolation is a linear combination over $\mathbb{F}_p$, not a probability distribution. A single equation

$$\sum_e \omega_e q_s(b_e) = 0$$

allows cancellation among invalid assignments. Our space imposes $T_{q_s}(C_{g(w)}) = 0$ for every polynomial g of degree at most t. In compressed coordinates these are exactly

$$\sum_e \omega_e g(e) \widetilde{q}_s(R^\top e) = 0.$$

Every function on the ρ-dimensional Boolean cube has a multilinear representative of degree at most $\rho \le t$. The equations therefore imply $\omega_e \widetilde{q}_s(R^\top e) = 0$ individually.

For an AND gate $b_k - b_i b_j = 0$, the compressed equation is $L_0 L_k - L_i L_j = 0$. The normalization equations first force $L_0(e) = 1$ at every nonzero coefficient. The weighted gate tests then force $L_k(e) = L_i(e) L_j(e)$ there. This step uses the verifier tests independently of the proof of Booleanity.

One may express this separation by the interpolation polynomials χ_e, or by invertibility of the Boolean evaluation matrix. Both express the Boolean interpolation principle used in Hair–Sahai [2, Section 5.3].

B.5 The dual Nullstellensatz viewpoint

A Boolean Nullstellensatz refutation is an identity

$$1 = \sum_{s=1}^{J} f_s(b) q_s(b) + \sum_{i=1}^{N} g_i(b)(b_i^2 - b_i). \tag{58}$$

Its degree bounds the degree of each product on the right. Polynomial calculus uses the related operations of taking linear combinations and multiplying equations by variables. Hair–Sahai explicitly connect their multiplied equations to these proof systems [2, Sections 5.3 and 6.2], following [3, 4, 6].

The table is a dual object: it stores moments of a linear functional annihilating permitted multiples of verifier equations. Compression makes this viewpoint particularly concrete. Boolean reduction identifies the relevant functional with one on the finite algebra

$$\mathcal{B}_\rho = \mathbb{F}_p[y_1, \dots, y_\rho] / \langle y_j^2 - y_j : 1 \le j \le \rho \rangle \simeq \prod_{e \in \{0,1\}^\rho} \mathbb{F}_p.$$

Every class has a multilinear representative of degree at most ρ. The transferred constraints are

$$1 - L_0, \qquad L_i^2 - L_0 L_i \ (1 \le i \le N), \qquad \widetilde{q}_s(R^\top y) \ (1 \le s \le J).$$

Because their degrees are at most two, all their products with multilinear representatives are covered by the moments through degree $\rho + 2 \le t + 2$. Thus μ annihilates the ideal generated by these constraints in $\mathcal{B}_\rho$. Under the displayed product decomposition, this says exactly that ω_e vanishes at every point violating a constraint.

If there is no compatible satisfying assignment, that ideal is the whole algebra, and μ is the zero functional. Reconstructing the stored blocks then gives $M = 0$. The contradiction uses the nonzero table, not an assumed normalization of its constant moment. This is the dual Nullstellensatz argument expressed through compressed scalar moments.

The algebra $\mathcal{B}_\rho$ should be distinguished from the original span of stored polynomial functions used for basis computation. That original span need not be closed under multiplication. The finite Boolean algebra here becomes available only after preservation and the moment identities have been proved for the particular matrix under consideration.

C Bounded sampling

This appendix specifies the bounded samplers used by Section 7 and records their exact conditional laws and failure probability.

Test at most ℓ_p^2 independent uniform candidates in $[2^{\ell_p}, 2^{\ell_p+1})$ with a deterministic polynomial-time primality test, accepting the first prime. For each subsequent field or grid coordinate, use at most ℓ_p rejection attempts with the minimum sufficient bit length. Sample one coordinate for τ, m for s, D for r, and mD for η; discard the η coordinates past k. The number of coordinates is fixed before τ is chosen, even when k varies with the seed. This is the bounded-sampling implementation of [2, Lemma 3.2 and Section 7], with our parameters.

The prime density is $\Omega(1/\ell_p)$, and each coordinate attempt succeeds with probability at least $1/2$. Consequently

$$\delta_{\mathrm{samp}} \le \exp(-\Omega(\ell_p)) + (1 + m + D + mD)2^{-\ell_p} = 2^{-\Omega(m)}. \tag{59}$$

On any failure publish a common abort symbol. Decryption still rejects invalid witnesses and returns zero on valid witnesses for that symbol. The abort probability and accepted-prime law are independent of β. Conditional on an accepted prime and a successful schedule, τ is exactly uniform, and the retained numerical samples have their independent uniform distributions, also after conditioning on τ. The fixed schedule prevents conditioning on successful sampling from biasing the public seed.

D The generic-group security proof

We give the detailed proof of Theorem 8.1, including the treatment of label bits and arbitrary-string queries. The public sampling laws are those of Appendix C. Write $H_q = q + m + mD + 2$ as in the theorem.

D.1 A common symbolic execution

Fix x, p, τ and the resulting public basis. Start with formal symbols

$$1, 0, S_1, \ldots, S_m, Z_1, \ldots, Z_k.$$

Every form produced by addition and subtraction has a unique coefficient representation

$$f = c + \alpha^\top S + \gamma^\top Z, \qquad (c, \alpha, \gamma) \in \mathbb{F}_p^{1+m+k}. \tag{60}$$

Identical coefficient vectors receive the same label. Distinct vectors receive uniform distinct labels of length ℓ_g, and the oracle adds or subtracts vectors. Reserve a label for zero at initialization; it need not be disclosed separately. At most $m + k + q + 2 \le H_q$ forms receive labels.

An operand string that has never received a label is rejected and excluded from all subsequent fresh-label choices. Wrong-length strings are also rejected. This defines a symbolic view that is independent of s, r, η and of the payload bit. In particular, after fixing the adversary's random tape and the symbolic label tape, the entire symbolic execution, including decisions based on label bits, is fixed independently of the numerical samples.

For each bit define a numerical execution by substituting

$$(S, Z) = \begin{cases} (s, (s^\top M_j r)_{j=1}^k), & \beta = 0, \\ (s, \eta), & \beta = 1, \end{cases} \tag{61}$$

and indexing the label table by actual field values. Numerically equal forms share a label even if their formal vectors differ. Use the same rejection and exclusion convention for unassigned operand strings.

Lemma D.1 (Arbitrary-string queries)**.** *For either bit, the real execution and its numerical execution have statistical distance at most $O(q/p)$ when $H_q \le p$.*

Proof. Couple the executions until an unassigned operand is a valid unseen encoding in the real group. Conditional on no earlier such event, the remaining valid encodings are exchangeable among unassigned, unexcluded strings. At most H_q labels have been assigned and at most $2q$ operand strings tested. Each new candidate has conditional probability at most

$$\frac{p}{2^{\ell_g} - H_q - 2q} = O(1/p)$$

of being valid, because $2^{\ell_g} \ge p^2$. A union bound over at most $2q$ candidates proves the claim. We declare unseen validity of either operand a bad event even if the other operand would already cause rejection. This covers an oracle that returns a single undifferentiated invalid-input answer. □

D.2 Bounding every symbolic collision

Proof of Theorem 8.1. Correctness is Proposition 7.1, and efficiency follows from the polynomial basis algorithm and the algorithms in Section 7. It remains to prove secrecy. We may assume $H_q \le p$, since the stated bound is already vacuous otherwise.

First fix a good seed. For any distinct pair of symbolic forms write its coefficient difference as (c, α, γ) and put

$$M_\gamma = \sum_{j=1}^{k} \gamma_j M_j.$$

If $\gamma \neq 0$, basis independence gives $M_\gamma \neq 0$. Because x is false, Corollary 6.2 gives

$$\max_z \Pr_{r \leftarrow [B]^D}[M_\gamma r = z] \le B^{-(t+1)}.$$

For payload zero, the difference evaluates to $c + s^\top(\alpha + M_\gamma r)$, so Lemma 5.1 bounds its probability of vanishing by $B^{-(t+1)} + p^{-1}$. This estimate holds for every nonzero M_γ, whatever its rank. If $\gamma = 0$, the difference is a nonzero affine form in uniform s and vanishes with probability at most p^{-1}. For payload one, every nonzero symbolic difference is a nonzero affine form in uniform (s, η) and has vanishing probability at most p^{-1}.

Fix the entire symbolic execution using its ordinary and label tapes. Its list of at most H_q forms is independent of the hidden numerical samples. Union-bound over all pairs in this list, including initial forms. The collision probabilities for bits zero and one are at most

$$\binom{H_q}{2}(B^{-(t+1)} + p^{-1}) \quad \text{and} \quad \binom{H_q}{2}p^{-1},$$

respectively. These are unconditional estimates over the numerical samples for the fixed symbolic list; no conditioning on the absence of earlier collisions is required.

Couple each numerical execution to the common symbolic execution by using the same ordinary and label tapes. Until two distinct formal vectors have the same numerical value, every assigned label and oracle reply agrees. The preceding union bounds therefore bound the statistical distances to the common symbolic view. Applying Lemma D.1 to each bit adds $O(q/p)$, which is absorbed in (45). This argument covers adaptive queries and arbitrary inspection of label bits.

The seed lemma is simultaneous over all matrices that could arise after the seed is known. Consequently the fixed-good-seed argument applies to every adaptive symbolic list, even though

that list depends on τ. Average over the uniform seed. On the exceptional event use the trivial distance bound one; thus δ_τ is added once, outside the collision union bound.

Finally, the two encryption experiments have identical prime and sampling-success laws, and their common abort outputs agree exactly. Averaging over accepted primes proves the unconditional claim. Since $B = m \geq N$, $m = N^{O(1)}$, and $t = \lceil \log \log N \rceil$, a polynomial q gives

$$H_q^2 B^{-(t+1)} = N^{-\Omega(\log \log N)}.$$

The terms 2^{-m} and $H_q^2 2^{-\ell_p}$ are smaller. Our padding has $N \geq \lambda$, and N is polynomial in λ on the instance families of Definition 2.1, so this is negligible in λ. □

E Parameters for the stronger security comparison

We justify the matched comparison in Section 10.4. Throughout this appendix, use the following parameters in place of (7) and (53):

$$\begin{aligned} &L = \log N, \qquad t = \lceil L \rceil, \qquad D = N + 1, \qquad W = \binom{2t}{t}, \qquad B = N^3, \\ &m_1 = DW, \qquad m_{\mathrm{HS}} = (2Nt + 1)m_1, \\ &\ell_0 = 2N + \lceil \log(BD + 2Nt + 2) \rceil, \\ &\ell_{\mathrm{HS}} = \ell_0, \qquad \ell_1 = \ell_0 + \lceil \log(Nt^2) \rceil + tD \lceil \log(2B^2) \rceil. \end{aligned} \tag{62}$$

For $i \in \{1, \mathrm{HS}\}$, choose a prime $p_i \in [2^{\ell_i}, 2^{\ell_i + 1})$ and use the shared short-vector grid $[B]^D$. The one-list construction samples one uniform $\tau \in \mathbb{F}_{p_1}$ and evaluates at $\tau, \ldots, \tau^t$; the deterministic construction retains the base 2 and uses all $2Nt + 1$ samples from [2, Section 5]. In each case let $k_i \leq m_i D$ be the dimension of the resulting matrix space.

E.1 Algebra and the public seed

The basis, determinant, and scalar-moment arguments apply to these parameters. Both chosen primes exceed 2^N, $2Nt$, and B. The condition $p > 2^N$ is used by the base-2 deterministic family; Lemma 5.3 for powers of the seed has no such requirement. The number of weights W is polynomial in N, and the extractor tests at most $2^t \leq 2N$ candidates. Hence all honest algorithms and the algebraic extractor have polynomial bit complexity.

For the one-list construction, the first two inequalities of (40) give

$$\delta_\tau \leq \frac{Nt^2 (2B^2)^{tD}}{p_1} \leq 2^{-\ell_0} \leq 2^{-2N}. \tag{63}$$

The simultaneous preservation statement and the moment representation of Theorem 4.1 therefore give $\kappa_B(M) \leq B^{-(t+1)}$ for every nonzero matrix on a false instance, outside this seed event. For the deterministic construction, the Section 5 argument of Hair–Sahai gives minimum nonzero rank at least $t + 1$ and hence the same concentration bound, with no seed exception.

These arguments do not require $B = m_i$ or a field length proportional to m_i. The grid enters the short-coordinate count and collision bound; the field must make the seed and field errors small and satisfy the applicable determinant and deterministic-sample conditions. In particular, this specialization uses (63), rather than the default bound $\delta_\tau \leq 2^{-m}$.

E.2 Sampling and correctness

Use the bounded sampling procedure of Appendix C with ℓ_i as the bit-length parameter and m_i as the row count. The all-list construction needs no random seed coordinate; both constructions use a fixed sampling schedule independent of the payload and of the computed basis dimension. In each case,

$$\delta_{\mathrm{samp},i} \le \exp(-\Omega(\ell_i)) + (1 + m_i + D + m_i D)2^{-\ell_i} = 2^{-\Omega(N)}.$$

Conditional on an accepted prime and successful sampling, all retained coordinates have the required independent uniform laws, including the one-list public seed. Common sampling aborts add no distinguishing error.

For either construction, the proof of Proposition 7.1 bounds the simultaneous decryption error, conditional on successful sampling and any accepted prime, by

$$\frac{2^N(BD+1)}{p_i} \le \frac{2^N(BD+1)}{2^{\ell_0}} \le 2^{-N}.$$

Here the search length is at most $(B-1)D+1$, independently of the row count. Adding sampling failures gives unconditional correctness error $2^{-\Omega(N)}$ for all witnesses of one ciphertext. An exceptional one-list seed affects security only.

E.3 Security for quasipolynomially many queries

The symbolic-label proof of Theorem 8.1 applies with the new parameters. It bounds the number of formal label pairs as a function of q and does not require the adversary's running time or q to be polynomial. Thus it also gives a quantitative guarantee beyond the efficient adversaries in Definition 2.1, under the same independence conditions on auxiliary input, the fresh encoding, and the encryption coins. Write

$$H_i = q + m_i + m_i D + 2, \qquad \delta_1 = 2^{-\ell_0}, \qquad \delta_{\mathrm{HS}} = 0.$$

For every false instance, the unconditional distinguishing advantage in construction i is at most

$$\mathrm{Adv}_i(q) \le \delta_i + O\left(H_i^2\left(B^{-(t+1)} + 2^{-\ell_i}\right)\right). \tag{64}$$

For the deterministic construction this is the same generic-group proof using its rank-based concentration bound; it reuses the underlying proofs of [2, Sections 3 and 5], with the new field and grid sizes.

Since $m_i D$ is polynomial in N, for all sufficiently large N and every $q \le N^L$ we have $H_i \le 2N^L < p_i$, as required by the arbitrary-string analysis in Lemma D.1. Consequently

$$H_i^2 B^{-(t+1)} \le 4N^{2L-3(t+1)} \le 4N^{-L-3}.$$

Also,

$$H_i^2 2^{-\ell_i} \le 4\,2^{2L^2-2N} = 2^{-\Omega(N)}, \qquad \delta_i \le 2^{-2N}.$$

Substituting in (64) and absorbing its absolute constant yields, for all sufficiently large N,

$$q \le N^{\log N} \quad \Longrightarrow \quad \mathrm{Adv}_i(q) \le N^{-\log N}.$$

The bound holds for the statistical distance of the two views and therefore also for distinguishing advantage.

E.4 Costs

The central binomial coefficient gives

$$W = \Theta(4^t/\sqrt{t}) = \widetilde{O}(N^2), \qquad m_1 = \widetilde{O}(N^3), \qquad m_{\mathrm{HS}} = \widetilde{O}(N^4).$$

The field lengths in (62) satisfy

$$\ell_1 = \Theta(N(\log N)^2), \qquad \ell_{\mathrm{HS}} = \Theta(N).$$

With the public basis recomputed, ciphertext size is $O((m_i+k_i)\ell_i)$ bits and decryption uses $O((m_i+k_i)\ell_i+BD)$ group queries. Since $k_i \le m_i D$ and $BD = O(N^4)$, both costs are $\widetilde{O}(N^5)$ for one random list and $\widetilde{O}(N^6)$ for all lists. The larger one-list field accounts for simultaneous preservation of all short-coordinate patterns, while the smaller row count retains a factor-N saving up to logarithmic factors in these upper bounds.

E.5 Smaller fields with per-witness correctness

Here we relax the correctness target to error at most $N^{-\log N}$ for each fixed valid witness. Retain $L = \log N$, $t = \lceil L \rceil$, $B = N^3$, $D = N+1$, and the same row counts, but use the all-list construction with

$$\ell_\gamma = \lceil 4L^2 \rceil, \qquad p_\gamma \in [2^{\ell_\gamma}, 2^{\ell_\gamma+1}), \qquad \mathrm{ord}_{p_\gamma}(\gamma) > N.$$

Every such prime exceeds $N(N+1)/2+1$, $2Nt$, and B for sufficiently large N. The deterministic search in Section 10.5 supplies γ from the accepted prime. Both are public. Use the $2Nt+1$ lists at parameters $0, \ldots, 2Nt$, with evaluation points $\gamma^{j-1}\tau$ in each list.

For a nonzero matrix of rank at most t, the proof of [2, Proposition 5.3] applies with this base: the determinant degree is at most Nt and a nonzero block entry has degree at most Nt in the list parameter. Among the $2Nt+1$ distinct parameters, some list therefore preserves the column space and has a nonzero block, yielding a witness. On a false instance the minimum nonzero rank is at least $t+1$, exactly as before.

For bounded sampling, test at most ℓ_γ^3 prime candidates, and use ℓ_γ rejection attempts per coordinate in a fixed schedule for both payloads. The prime density and coordinate acceptance bounds from Appendix C give

$$\delta_{\mathrm{samp}} \le \exp(-\Omega(\ell_\gamma^2)) + (m_{\mathrm{HS}} + D + m_{\mathrm{HS}}D)2^{-\ell_\gamma} = \exp(-\Omega(L^4)) + N^{O(1)-4L}.$$

The order search is deterministic and succeeds for every accepted prime. For each fixed valid witness, the decryption proof now yields

$$\Pr[\text{decryption error}] \le \delta_{\mathrm{samp}} + \frac{BD+1}{p_\gamma} \le N^{-L}$$

for sufficiently large N. This estimate has no union bound over the 2^N possible assignments and does not claim simultaneous correctness for all witnesses of a ciphertext.

For $q \le N^L$, put $H = q + m_{\mathrm{HS}} + m_{\mathrm{HS}}D + 2 \le 2N^L$, asymptotically. The rank-based collision proof bounds advantage by

$$O\left(H^2(B^{-(t+1)} + p_\gamma^{-1})\right) = O\left(N^{-L-3} + N^{-2L}\right) \le N^{-L}.$$

The bound is uniform over accepted primes, and common sampling aborts add no distinguishing error.

Finally, $W = \Theta(N^2/\sqrt{L})$, $m_1 = \Theta(N^3/\sqrt{L})$, and $m_{\mathrm{HS}} = \Theta(N^4\sqrt{L})$. The one-list parameters of (62) have $\ell_1 = \Theta(NL^2)$, whereas $\ell_\gamma = \Theta(L^2)$. Using $k_i \le m_i D$ in the ciphertext and decryption bounds gives $O(N^5L^{3/2})$ for one list and $O(N^5L^{5/2})$ for all lists; the search term $BD = O(N^4)$ is smaller. Both are $\widetilde{O}(N^5)$ under this per-witness correctness target.

For the weaker threshold $t = \lceil \log\log N \rceil$, take $B_{\mathrm{HS}} = m_{\mathrm{HS}}$ and the same $\ell_\gamma = \lceil 4L^2 \rceil$. Then $m_{\mathrm{HS}} = \widetilde{O}(N^2)$, so ciphertext bits and decryption queries are $\widetilde{O}(N^3)$. The same sampling and fixed-witness correctness estimates are negligible, and $B_{\mathrm{HS}}^{-(t+1)}$ yields negligible advantage against every polynomial query bound.

F Efficient witness decryption using systematic coordinates

The witness supplies an explicit assignment table, but the direct decryption algorithm must still express that table in the published group labels' matrix basis. This appendix first quantifies the cost of rebuilding that basis from public parameters. It then gives a variant in which the encryptor uses a systematic basis and transmits its pivot positions. Decryption reads its coefficients directly from the witness, groups repeated weights before exponentiation, and uses baby-step/giant-step for the final bounded search. No matrix basis needs to be published in advance or reconstructed by the decryptor.

Write

$$D = N + 1, \qquad W = \binom{2t}{t}, \qquad m = DW, \qquad E = mD = D^2W, \qquad \ell = \lceil \log p \rceil.$$

Field operations mean arithmetic in $\mathbb{F}_p$; group-oracle queries and ordinary bit operations are counted separately. Verifying the witness and extending it to the circuit assignment takes the usual verifier time. The gate encoding has $J = O(N)$ equations of constant sparsity. When J is left general, we display its contribution explicitly.

F.1 The cost of direct basis reconstruction

Put $d = \dim \mathcal{V}_\tau$ and $k = \dim \mathcal{S}_{x,\tau}$, so $k \le d \le E$. In Appendix A, every prefix span is contained in $\mathcal{V}_\tau$. Each of the N stages therefore reduces at most $2d$ candidate vectors of length E against at most d retained vectors. Classical Gaussian elimination costs $O(Ed^2)$ field operations per stage, or $O(NEd^2)$ in total. Evaluating the candidate tables fits within this bound: update the t forms, evaluate the weights, and fill the E entries. Precomputing the coefficients τ^{ij} also fits within it.

The verifier-test matrix has WJ rows and d columns. For sparse gate equations it takes $O(WJd)$ operations to form, and at most $O(WJd^2)$ to compute its kernel by classical elimination. Forming the resulting k explicit basis tables costs $O(Edk)$. Given these tables without coordinate-solving data, a fresh dense solve for the witness coefficients costs $O(Ek^2)$. Thus the complete arithmetic cost of direct decryption from only public numerical parameters is bounded by

$$\begin{aligned} \mathsf{F}_{\mathrm{direct}} &= O\big(NEd^2 + WJd^2 + Edk + Ek^2\big) \\ &= O\big((NE + WJ)d^2\big) \le O\big(NE^3 + WJE^2\big). \end{aligned} \tag{65}$$

For $J = O(N)$ this is $O(NE^3)$ field operations, with $O(E^2)$ field elements of working storage. Fast field arithmetic converts it to $\widetilde{O}(NE^3\ell)$ arithmetic bit operations; this is in addition to the group-query work. These are bounds for the stated dense implementation, not lower bounds on decryption.

Hair–Sahai [2, Section 3.1] ask decryption to compute the factors and basis coefficients promised by the reduction. Their Section 5.4 supplies the iterative span algorithm and says that the witness coordinates follow from "one further linear solve." It proves polynomial running time without giving an explicit field-operation bound. The count above makes that implementation quantitative. For their $2Nt+1$ lists, replace W by $K=(2Nt+1)W$ and E by $E_{\mathrm{HS}}=D^2K$ throughout. Their Section 5.4 explicitly works with the D active columns before padding with zeros, so the comparison also uses this rectangle.

With no cached basis or coordinate data, the resulting upper bounds are as follows; $\widetilde{O}$ suppresses powers of logarithms.

Threshold	One list	Hair–Sahai, all lists
$t=\lceil\log\log N\rceil$	$\widetilde{O}(N^7)$	$\widetilde{O}(N^{10})$
$t=\lceil\log N\rceil$	$\widetilde{O}(N^{13})$	$\widetilde{O}(N^{16})$

All entries count field operations, independently of field bit length. They do not include an algebraic witness search: the decryptor already has its witness.

F.2 Read the certificate at the pivot positions

Flatten the matrices of $\mathcal{S}_{x,\tau}$ into row vectors of length E, in the fixed order (α,i,c) used in Appendix A. The encryptor replaces its ordered basis by the nonzero rows of its reduced row echelon form. Reshape these rows as $M_1,\dots,M_k$, and let $\pi_1<\cdots<\pi_k$ be their pivot positions. This is a canonical basis of the space for the chosen coordinate order, and

$$M_j[\pi_i]=\delta_{ij},\qquad A=\sum_{j=1}^{k}A[\pi_j]M_j\quad\text{for every } A\in\mathcal{S}_{x,\tau}. \tag{66}$$

Indeed, if $A=\sum_j a_jM_j$, its entry at π_i is a_i.

Use this basis in the encryption formulas (42), and append the ordered pivot positions to the ciphertext. They require $O(k\log E)$ bits as a list, or E bits as a mask whose set bits are read in increasing order. Choose the shorter encoding, with a fixed format indicator. The matrix entries are not transmitted. The encryptor still computes the basis; normalizing it costs at most $O(Ek^2)$ additional field operations, already within the preceding bound.

For a valid witness, put $v=v(b)$ and write $\pi_j=(\alpha_j,i_j,c_j)$. The coefficients and left factor are now explicit:

$$a_j=h_{\alpha_j}(b)v_{i_j}v_{c_j},\qquad u=(h_\alpha(b)v)_\alpha. \tag{67}$$

Equation (66) gives $\sum_j a_jM_j=A_\tau(b)=uv^\top\neq 0$. Reading a coefficient requires a weight lookup and two Boolean tests; it does not require a linear solve or an entry of the matrix basis. The coefficients need not be stored explicitly.

Correctness and security. The same rank-one certificate proves correctness. For payload one, sample independent uniform exponents in this systematic basis, as in the original construction; the fixed schedule may still reserve E such samples and discard those beyond k. The systematic basis and its pivots are deterministic functions of the public parameters and instance. The original security proof already permits the adversary to compute them. Its linear independence and concentration arguments therefore apply unchanged, including for a matrix chosen after the metadata is seen. For malformed metadata, reject invalid coordinate ranges, unordered or repeated pivots, and incompatible lengths. Honest decryption does not recompute the basis to certify that metadata: the claimed correctness and security concern ciphertexts generated by the encryption algorithm.

F.3 Group equal weights before exponentiating

Evaluate the t forms $\ell_s(b) = \sum_i v_i \tau^{is}$ in $O(Nt)$ field operations: compute τ^s successively and evaluate each polynomial by Horner's rule. Fix, for every nonzero exponent tuple α, an index $s(\alpha)$ with $\alpha_{s(\alpha)} > 0$. In increasing total degree, compute

$$h_0 = 1, \qquad h_\alpha = h_{\alpha - e_{s(\alpha)}} \ell_{s(\alpha)}.$$

All W weights then take $O(W)$ multiplications. This uses no division and remains valid when a form or weight vanishes.

Let $I = \{i : v_i = 1\}$, which contains zero. Index the labels X_i from (42) more explicitly by (α, i). Form the group elements

$$P_\alpha = \prod_{i \in I} X_{\alpha,i}, \qquad Q_\alpha = \prod_{\substack{j:\alpha_j = \alpha \\ i_j, c_j \in I}} Y_j.$$

An empty product is the identity. The decryption products become

$$H = \prod_\alpha P_\alpha^{h_\alpha(b)}, \qquad Z = \prod_\alpha Q_\alpha^{h_\alpha(b)}. \tag{68}$$

These are exactly the products in (43), by (67). Forming the inner products uses $O(m + k)$ group queries. There are at most two scalar exponentiations per weight, costing $O(W\ell)$ queries in total, instead of exponentiating each of the $m + k$ labels separately. Pivot order is block order, so the products can be formed while scanning the ciphertext once.

F.4 The bounded search by baby-step/giant-step

It remains to test whether $Z = H^z$ for some integer $0 \le z \le L_w = (B-1)\,\mathrm{wt}(v)$. If $H = 1$, test $Z = 1$ directly. Otherwise H has order p. Put $L' = \min\{L_w, p-1\}$ and $s_0 = \lceil \sqrt{L'+1} \rceil$. Store the baby steps

$$H^j, \qquad 0 \le j < s_0,$$

with their indices. Compare each giant step

$$Z(H^{-s_0})^i, \qquad 0 \le i \le \lfloor L'/s_0 \rfloor,$$

against this table. A match at j is accepted precisely when $is_0 + j \le L'$. Each integer in $[0, L']$ has such a quotient and remainder, so this is the exact original acceptance test. The truncation at $p-1$ also makes it valid if later parameters allow $L_w \ge p-1$; all stated parameters have $L_w < p$.

Successive group operations generate both lists in $O(\sqrt{L'+1}) = O(\sqrt{BD})$ queries, with that many stored labels. Sorting the baby labels and using binary search gives deterministic lookup with only logarithmic overhead in local bit work. No discrete logarithm outside the prescribed interval is computed.

F.5 Optimized complexity and comparison

Proposition F.1 (Decryption without reconstructing a basis)**.** *For the systematic variant just described, honest decryption uses*

$$\begin{aligned} \mathsf{F}_{\mathrm{dec}} &= O(Nt + W) && \textit{field operations}, \\ \mathsf{G}_{\mathrm{dec}} &= O(m + k + W\ell + \sqrt{BD}) && \textit{group queries}. \end{aligned} \tag{69}$$

There is also a scan of the labels and pivot metadata, Boolean endpoint tests, and the work of verifying the witness. These are ordinary bit operations, rather than additional field arithmetic. With fast modular arithmetic, the local bit work is bounded by

$$\widetilde{O}\Big(\ell(Nt + W + m + k + W\ell + \sqrt{BD}) + k \log E + N + J\Big). \tag{70}$$

This counts reading labels and manipulating their bit strings but does not assign an internal implementation cost to the generic-group oracle. Apart from the input ciphertext, storage is $O(W+\sqrt{BD})$ field elements or group labels and their indices. Storing the circuit assignment additionally takes $O(N)$ bits.

Proof. The forms and the weight recurrence give the field-operation count. Selecting coefficients uses only copies, zeros, and Boolean tests. Equation (68) and the bounded search give the query count. Field arithmetic has quasi-linear bit cost in ℓ, up to logarithmic factors, and each group query handles $O(\ell)$-bit labels. The latter contributes the $W\ell^2$ term when all weight exponentiations are charged at field length. Scanning the pivot encoding costs at most $O(k \log E)$ bits; the alternative mask is used only if it is shorter. Label sorting and index bookkeeping introduce only the suppressed logarithms. □

The ciphertext still has $m+k+1$ group labels, with the additional $O(\min\{k \log E, E\})$ metadata bits. This does not change the ciphertext-size bounds in either parameter regime. The field-operation count in (69) does not mean that decryption can avoid reading a large ciphertext; that cost is included in (70).

Apply the same optimizations to Hair–Sahai. Their assignment tables have the same Boolean endpoint factors. For $r_{\text{lists}} = 2Nt + 1$ lists and $K = r_{\text{lists}}W$ weights, the systematic certificate and grouping apply with the combined weight index (list, α). Evaluating all forms and weights costs

$$O(r_{\text{lists}}Nt + K)$$

field operations. The group bound becomes $O(m_{\text{HS}} + k_{\text{HS}} + K\ell_{\text{HS}} + \sqrt{B_{\text{HS}}D})$. Thus the comparison credits both constructions with the same decryption improvements.

The following are worst-case upper bounds using $k \le E$ and $k_{\text{HS}} \le E_{\text{HS}}$. The one-list rows use our stated fields; the Hair–Sahai rows use the $\Theta(N)$-bit fields justified in Section 10. For the logarithmic threshold use the matched parameters of Appendix E, including $B = N^3$ in both constructions. For the smaller threshold use each construction's stated grid. All rows retain exponentially small simultaneous correctness error in N.

Threshold	Construction	Field ops.	Group queries	Local bits
$\log \log N$	One list	$\widetilde{O}(N)$	$\widetilde{O}(N^2)$	$\widetilde{O}(N^3)$
$\log \log N$	Hair–Sahai	$\widetilde{O}(N^2)$	$\widetilde{O}(N^3)$	$\widetilde{O}(N^4)$
$\log N$	One list	$\widetilde{O}(N^2)$	$\widetilde{O}(N^4)$	$\widetilde{O}(N^5)$
$\log N$	Hair–Sahai	$\widetilde{O}(N^3)$	$\widetilde{O}(N^5)$	$\widetilde{O}(N^6)$

Here ceilings are understood in the thresholds. These are bounds for implementations of the two spaces, not claims that Hair–Sahai already state these optimizations. They do not charge the baseline for its original, larger field length. With the polylogarithmic field and per-witness correctness target of Appendix E.5, the optimized Hair–Sahai local bit bounds become $\widetilde{O}(N^3)$ and $\widetilde{O}(N^5)$, respectively; its field-operation and group-query bounds stay as displayed. Supply the high-order folding base γ with the public parameters so that decryption need not repeat its search.

What is preprocessed, and what is avoided. The systematic variant needs no advance publication of a basis. Its encryptor performs the basis computation and sends only the pivots along with the labels. If N, p, t, τ are reused, the universal span $\mathcal{V}_\tau$ can be cached once as described in Appendix A. The verifier-dependent basis and pivots can also be cached for repeated encryptions of the same instance. Publishing the universal basis alone does not remove the instance-dependent kernel computation from the direct method. The optimized decryptor reads neither dense basis: it uses the pivots and its own witness. This improves decryption without claiming that the encryptor's remaining linear algebra has the same cost. The arithmetic bounds here also do not apply to the algebraic extractor, which has a different input and purpose.

The systematic method can also be extended to deterministic seed choices and to smaller hybrid tables. We do not analyze these extensions here.

G One-variable rank-preserving maps

The map in Lemma 5.3 is one choice of a polynomial matrix $F(T)$ that preserves each fixed input rank for most seeds. Here $F(T)$ denotes a $D \times t$ matrix, corresponding to $F_t(T)$ for the power map. The power map is due to Gabizon–Raz [8]; Forbes–Shpilka [9] give its geometric-folding refinement. Forbes–Guruswami [10] discuss the general framework of seeded rank condensers. After transposition, the maps below are *lossless seeded rank condensers*: $F(\tau)^\top$ preserves the rank of a fixed matrix $Q^\top$ except at a bounded number of seeds. All constructions here use a single scalar variable T. We first record the proof of Lemma 5.3, then give a common sufficient condition and alternatives with smaller determinant degree.

G.1 Proof of the power-evaluation lemma

The following proof reformulates the leading-degree argument of Gabizon–Raz [8, Lemma 6.1].

Proof of Lemma 5.3. Associate to row i the polynomial $f_i(T) = \sum_{a=0}^{N} Q_{i,a} T^a$. Full row rank of Q makes these polynomials linearly independent over $\mathbb{F}_p$. Constant invertible row operations put their highest-degree coefficients in echelon form, so we may assume

$$d_1 := \deg f_1 < \cdots < d_\nu := \deg f_\nu \le N.$$

These operations multiply the determinant under consideration by a nonzero field element. It therefore suffices to prove that

$$E(T) = (f_i(T^j))_{1 \le i,j \le \nu}$$

has nonzero determinant.

A dependence would require too many leading-degree ties. Suppose otherwise. Over the field $\mathbb{F}_p(T)$, there is a nonzero row dependence. Clearing denominators gives polynomials $a_i(T) \in \mathbb{F}_p[T]$, not all zero, such that

$$\sum_{i=1}^{\nu} a_i(T) f_i(T^j) = 0 \qquad (j = 1, \ldots, \nu).$$

Omit the zero coefficients and reindex the remaining rows in increasing degree order. There are $r \le \nu$ of them, with $d_1 < \cdots < d_r$. Put $\delta_i = \deg a_i$. In column j, the degree of the ith nonzero summand is exactly

$$\delta_i + j d_i.$$

The largest such degree must be attained at least twice; otherwise its nonzero leading coefficient cannot cancel.

As column index j increases, the row indices attaining the maximum can only increase. Indeed, if i maximizes at j and k maximizes at $j' > j$, then

$$\delta_i + jd_i \geq \delta_k + jd_k, \qquad \delta_k + j'd_k \geq \delta_i + j'd_i.$$

Adding gives $(j' - j)(d_k - d_i) \geq 0$, so $i \leq k$. For each column let l_j and h_j be the smallest and largest maximizing (row) indices, respectively, since we are claiming that there are at least two maximizing indices (for each column). Cancellation in all ν columns would imply

$$l_1 < h_1 \leq l_2 < h_2 \leq \cdots \leq l_\nu < h_\nu.$$

There are ν strict increases, requiring at least $\nu + 1$ distinct indices among only $r \leq \nu$ rows. This contradiction proves that the determinant is nonzero.

Degree of the determinant. Every entry in column j has degree at most Nj. Each product in the determinant expansion uses exactly one entry from every column, and hence has degree at most $\sum_{j=1}^{\nu} Nj$. Consequently

$$\deg \det E(T) \leq N \sum_{j=1}^{\nu} j = \frac{N\nu(\nu+1)}{2}. \tag{71}$$

Undoing the constant row operations does not change this degree. The argument works in every characteristic; it uses no condition such as $p > 2^N$ or a restriction on the factorization of $p - 1$. □

G.2 A criterion and maps with smaller determinant degree

Put $D = N + 1$ and index rows by $0, \ldots, N$. For an increasing ν-element index set $I = (i_1 < \cdots < i_\nu)$, write F_I for the corresponding square row submatrix of $F \in \mathbb{F}_p[T]^{D \times \nu}$. Write $I \preceq J$ if $i_k \leq j_k$ for every k.

Proposition G.1 (A sufficient condition on minor degrees)**.** *Suppose every $\Delta_I(T) = \det F_I(T)$ is nonzero, and put $d(I) = \deg \Delta_I$. Assume*

$$I \preceq J, \quad I \neq J \quad \Longrightarrow \quad d(I) < d(J).$$

Then, for every fixed full-row-rank $Q \in \mathbb{F}_p^{\nu \times D}$, $\det(QF(T))$ is nonzero and has degree at most $\max_I d(I)$.

Proof. Constant invertible row operations put Q in echelon form from the right. Order the resulting rows so that their rightmost nonzero positions satisfy $e_1 < \cdots < e_\nu$. Thus row k is supported on columns at most e_k. The minor on $E = (e_1, \ldots, e_\nu)$ is triangular with nonzero diagonal, so $\det Q_E \neq 0$; row operations preserve which maximal minors are nonzero.

Every I with $\det Q_I \neq 0$ satisfies $I \preceq E$. Indeed, if $i_k > e_k$, at most $k - 1$ selected columns lie at most e_k, whereas the first k echelon rows are supported there. The corresponding square minor would therefore be singular. Now Cauchy–Binet gives

$$\det(QF(T)) = \sum_{|I|=\nu} \det Q_I \, \Delta_I(T).$$

The summand indexed by E has strictly larger degree than every other nonzero summand. Its leading coefficient is nonzero, so the sum is nonzero and has degree exactly $d(E)$. □

This is a sufficient condition, not a classification of rank condensers. The proof abstracts the pivot and leading-degree arguments used for the maps above; the proposition is our formulation of that common argument. For a $D \times t$ map, write F_ν for its first ν columns and apply the criterion separately for each $1 \le \nu \le t \le D$.

Weighted Vandermonde maps. Choose any distinct $\lambda_0, \dots, \lambda_N \in \mathbb{F}_p$ and set

$$F^{\mathrm{V}}_{i,j}(T) = \lambda_i^{j-1} T^i, \qquad 0 \le i \le N, \quad 1 \le j \le t, \tag{72}$$

with zeroth powers interpreted as one. Factoring T^{i_k} from each selected row gives

$$\det(F^{\mathrm{V}}_\nu(T))_I = T^{\sum_k i_k} \prod_{k<\ell} (\lambda_{i_\ell} - \lambda_{i_k}).$$

The coefficient is nonzero and $d(I) = \sum_k i_k$ is strictly increasing under coordinatewise increase. Consequently

$$\deg \det(QF^{\mathrm{V}}_\nu(T)) \le N\nu - \binom{\nu}{2} \le N\nu. \tag{73}$$

The choice $\lambda_i = \gamma^i$, for an element γ of order greater than N, gives the geometric folding used by Hair–Sahai [2] and in Section 10.5. For prime $p > N$, the simpler choice $\lambda_i = i$ also works and requires no multiplicative-order search. Applied to a row polynomial $f(T)$, its columns are $f(T), (T\frac{d}{dT})f(T), \dots, (T\frac{d}{dT})^{t-1} f(T)$.

Ordinary derivative Wronskians. Assume $p > N$ and use

$$F^{\mathrm{W}}_{i,j}(T) = \begin{cases} (i)_{j-1} T^{i-j+1}, & i \ge j-1, \\ 0, & i < j-1, \end{cases} \tag{74}$$

where $(i)_r = i(i-1)\cdots(i-r+1)$ and $(i)_0 = 1$. Thus QF^{W}_ν contains each row polynomial and its first $\nu - 1$ ordinary derivatives. Since the falling-factorial polynomials are monic of successive degrees, their evaluation determinant is a Vandermonde determinant. It follows that

$$\det(F^{\mathrm{W}}_\nu(T))_I = T^{\sum_k i_k - \binom{\nu}{2}} \prod_{k<\ell} (i_\ell - i_k).$$

One may derive the identity by factoring in $\mathbb{F}_p(T)$; the exponent is nonnegative because $i_k \ge k-1$. All displayed differences are nonzero when $p > N$. Proposition G.1 therefore gives

$$\deg \det(QF^{\mathrm{W}}_\nu(T)) \le N\nu - \nu(\nu-1) = \nu(N - \nu + 1). \tag{75}$$

The characteristic condition matters for ordinary derivatives; for example, the derivative of T^p vanishes in characteristic p.

The power-evaluation map. For $F^{\mathrm{P}}_{i,j}(T) = T^{ij}$, the increasing ordering of $i_1, \dots, i_\nu$ and $1, \dots, \nu$ makes the diagonal permutation the unique highest-degree term of each minor. Its coefficient is one and its degree is $\sum_k k i_k$. The criterion thus recovers Lemma 5.3, with the stated bound $N\nu(\nu+1)/2$, over every characteristic. The two preceding maps reduce this degree to $O(N\nu)$ under their stated field conditions.

Random seeds and the short-pattern family. If d_ν is any of the degree bounds above, the polynomial root bound gives, for each fixed full-row-rank Q,

$$\Pr_{\tau \leftarrow \mathbb{F}_p}[\operatorname{rank}(QF_\nu(\tau)) < \nu] \le d_\nu/p.$$

For any prescribed finite families $\mathcal{Q}_\nu$ of such matrices, a union bound makes the same seed work for all of them except with probability at most

$$\frac{1}{p}\sum_{\nu=1}^{t} d_\nu |\mathcal{Q}_\nu|.$$

In particular, one may use $|\mathcal{Q}_\nu| \le A_B^{\nu D}$ as in Lemma 5.4, replacing its determinant-degree factor by the smaller applicable bound. These are guarantees for a fixed matrix or a prescribed finite family; they do not assert that one numerical seed preserves every subspace over $\mathbb{F}_p$. As in Section 10.5, reducing this degree prefactor saves only $O(\log t)$ field bits in the present regimes.