

A Note on Extractable Witness Encryption in Generic Groups

Matteo Campanelli

Offchain Labs & University of Tartu, Estonia
binarywhalesinternaryseas@gmail.com

October 7, 2026

Abstract. Hair and Sahai [HS26] have recently proposed a construction of witness encryption for NP satisfying semantic security. In this short note we observe that this recent witness-encryption construction is unconditionally *extractable* in the generic group model. This yields a construction of extractable witness encryption without assuming subexponential soundness of SNARGs, in contrast to the recent work of Jin [Jin26], which requires this assumption. The result in this note, moreover, complements the concurrent work of Chen and Vaikuntanathan [CV26], who show the same construction to be public-coin extractable under a knowledge-style assumption; our argument is unconditional, but in the GGM.

1 Introduction

Hair and Sahai [HS26] have recently proposed the first construction of witness encryption (WE) for the NP-complete problem SAT that is “direct” (it does not go through indistinguishability obfuscation) and does not require unproven conjectures. They prove their construction to be semantically secure in the generic-group model (GGM). Semantic security is the “baseline” security notion of witness encryption and states that, for any false statement x , an encryption of the bit 0 with respect to x is computationally indistinguishable from an encryption of the bit 1 with respect to the same statement x .

We show that their construction already satisfies, unconditionally and in the GGM, a stronger property: extractability. In the context of SAT, this intuitively means that the ability to distinguish encrypted messages implies knowledge of a satisfying assignment.

Comparison with recent work. The recent work of Jin [Jin26] gives a construction of extractable WE for NP in the GGM but requires SNARGs for NP with subexponential soundness and additional efficiency properties. Another recent paper, the concurrent and independent work of Chen and Vaikuntanathan [CV26]¹, shows that Hair and Sahai’s construction is public-coin extractable under a knowledge-style assumption, namely the Knowledge of Orthogonality Assumption (KOALA). In contrast to Jin’s work, our result does not require additional

¹ This work appeared on the date this note was submitted to the IACR ePrint Archive.

assumptions on SNARG security. Our result complements that of Chen and Vaikuntanathan, who require a specific assumption but can dispense with generic groups. We note that their model differs slightly from ours. Finally, their assumption guarantees the existence of a potentially *nonuniform* extractor, whereas we provide a single explicit, uniform extraction procedure that runs the adversary and observes its generic-group queries, without requiring its code.

1.1 The WE Construction of Hair and Sahai

The NP problem we reduce SAT to. Hair and Sahai reduce a SAT instance φ of size N to an instance of the promise problem *Boolean-factor MinRank*. The latter consists of a basis of matrices $M_1, \dots, M_k$ whose span is $S = \left\{ \sum_{j=1}^k a_j M_j : \mathbf{a} \in \mathbb{F}_p^k \right\}$ and whose properties depend on whether φ is satisfiable or not:

- **YES instances:** if φ is satisfiable then S contains a nonzero matrix $\mathbf{u}\mathbf{v}^\top$, with $\mathbf{v}$ Boolean and $\mathbf{u} \in \mathbb{F}_p^m$. Moreover, given a satisfying assignment w of φ , it is possible to efficiently find $\mathbf{u}, \mathbf{v}, \mathbf{a}$ such that $\mathbf{u}\mathbf{v}^\top = \sum_j a_j M_j \in S$.
- **NO instances:** if φ is unsatisfiable then every nonzero matrix in S has rank at least $1 + \lfloor \log N \rfloor$.

The WE scheme. We work in a generic group $\mathbb{G}$ (see Section 2) of prime order p . We denote its generator by $[1]$ and use additive notation for group operations. Given $x \in \mathbb{F}_p$, we write $[x] = x[1]$. For a vector $\mathbf{x} = (x_1, \dots, x_\ell)$, we define $[\mathbf{x}] = ([x_1], \dots, [x_\ell])$.

Encryption. In order to encrypt a bit $\beta \in \{0, 1\}$, we:

- sample a uniform vector $\mathbf{s} \in \mathbb{F}_p^m$, where m is the side length of each basis matrix M_j ;
- independently sample $\mathbf{r} = (r_1, \dots, r_m)$ with each r_j uniform in $\{0, \dots, m-1\}$;
- produce a ciphertext $\text{ct} = (\mathbf{X}, \mathbf{Y})$ with $\mathbf{X} = [\mathbf{s}]$. If $\beta = 1$, the entries of $\mathbf{Y}$ are independent uniform elements of $\mathbb{G}$, also independent of $\mathbf{s}$ and $\mathbf{r}$; otherwise $\mathbf{Y} = ([\mathbf{s}^\top M_1 \mathbf{r}], \dots, [\mathbf{s}^\top M_k \mathbf{r}])$, one group element per basis matrix.

Decryption. The decryption algorithm recognizes whether the ciphertext $\text{ct} = (\mathbf{X}, \mathbf{Y})$ contains a random $\mathbf{Y}$ (encrypting 1) or a “structured” $\mathbf{Y}$ (encrypting 0). It tests for the latter case by building two group elements, H from $\mathbf{X}$ and Z from $\mathbf{Y}$, and searching for a relation between them. If such a relation holds, decryption returns 0; otherwise it returns 1.

Recall that, given a satisfying assignment w of φ , it is possible to efficiently find $\mathbf{u}, \mathbf{v}, \mathbf{a}$ such that $\mathbf{u}\mathbf{v}^\top = \sum_j a_j M_j \in S$. We compute $Z = \sum_i a_i Y_i$. By the

property above and by inspecting what we did during encryption, whenever the encrypted message is the bit 0, this holds:

$$Z = \sum_i a_i Y_i = \underbrace{[\mathbf{s}^\top \mathbf{u}]_f}_{f} \underbrace{(\mathbf{v}^\top \mathbf{r})}_{f'},$$

that is, Z is the group encoding of the product of two factors, f and f' , defined as above.

What the decryptor now does is to use $\mathbf{X}$ to check whether this is actually the shape of the Z it computed. It proceeds in two steps, *i*) first by building the group encoding of the first factor f and then *ii*) by searching for the second factor f' :

i) it computes $H = \sum_j u_j X_j$. Clearly $H = [\mathbf{s}^\top \mathbf{u}] = [f]$.

ii) it checks whether there exists $\hat{f}' \in \{0, 1, \dots, B\}$, where $B = (m-1) \text{wt}(\mathbf{v})$, such that $Z = \hat{f}' H$. If so, decryption returns 0; otherwise it returns 1.

We can search efficiently for $f' = \mathbf{v}^\top \mathbf{r}$ because it is small: as an integer, it lies in $\{0, \dots, B\}$ with $B = O(m^2)$, since $\mathbf{v}$ is Boolean and the entries of $\mathbf{r}$ lie in $\{0, \dots, m-1\}$.

On correctness error. Notice that this test can fail, though only in one direction. When the encrypted bit is 0 the element Z has exactly the shape above and the search always succeeds; when it is 1, however, Z is a uniformly random group element, and nothing prevents it from accidentally coinciding with $\hat{f}' H$ for one of the $B+1$ candidates. By a union bound, this happens with probability at most $(B+1)/p$, which is negligible.

On extractability. Extractability guarantees that, given the description of a distinguishing adversary, we can efficiently find a satisfying assignment with probability at least the adversary's advantage, up to negligible error. In the GGM, the adversary manipulates ciphertext group elements through an oracle (we adopt Maurer's GGM [Mau05]). It sees only placeholders for these elements ("handles") and can build new handles from old ones. Every handle it builds represents an affine expression

$$c + \boldsymbol{\alpha}^\top \mathbb{X} + \boldsymbol{\gamma}^\top \mathbb{Y}$$

where $\mathbb{X}$ and $\mathbb{Y}$ are formal variables representing the exponents of the ciphertext elements $\mathbf{X}$ and $\mathbf{Y}$, respectively. The coefficients can be tracked from the adversary's queries.

One key intuition, common in GGM-based extractors, is that the extractor can act as the group oracle and observe the adversary's queries. Creating new handles gives the adversary no information about the encrypted bit. Distinguishing requires a successful zero test on a nonzero affine form: the evaluated exponent is zero even though its coefficients are not all zero.

Our extraction approach proceeds as follows: first, the extractor runs the adversary on an encryption of 0 and records the coefficients of its queries. This is the informative case for a good distinguisher, since only encryptions of 0 contain “structured” group elements; see Lemma 4. With probability at least the adversary’s advantage minus negligible error, the first successful zero test on a nonzero affine form yields a nonzero matrix $M_\gamma = \sum_j \gamma_j M_j$ of rank at most $R = \lfloor \log N \rfloor$. From such a matrix, we can efficiently construct a polynomial-size list containing a satisfying assignment and test its candidates. Hair and Sahai use a similar observation in contrapositive form to establish semantic security for false statements; here it provides the witness we need for extraction.

2 Preliminaries

All logarithms have base two. For a prime p , let $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$. Vectors are columns, and all matrix ranks and linear spans are over $\mathbb{F}_p$. For $\mathbf{v} \in \{0, 1\}^m$, let $\text{wt}(\mathbf{v})$ denote its Hamming weight. A function $\nu: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is negligible if, for every positive polynomial q , one has $\nu(\lambda) < 1/q(\lambda)$ for all sufficiently large λ . An algorithm is PPT if it is probabilistic polynomial time.

Let $\mathcal{R}$ be an NP witness relation, that is, one that is decidable in polynomial time and satisfies $|w| \leq \text{poly}(|x|)$ whenever $(x, w) \in \mathcal{R}$, and let

$$L_{\mathcal{R}} = \{x : \exists w, (x, w) \in \mathcal{R}\}.$$

It suffices to give witness encryption for Boolean circuit satisfiability. A verifier for $\mathcal{R}$ compiles in polynomial time into a circuit φ_x , preserving witnesses both ways, so a scheme for circuit satisfiability yields one for $\mathcal{R}$, and an extractor that outputs a satisfying assignment of φ_x outputs a witness for x after the same compilation. From here on the statement is a Boolean circuit φ and a *satisfying assignment* of φ is an input w with $\varphi(w) = 1$.

2.1 Generic Cyclic Groups

We work in Maurer’s formulation [Mau05] of the generic group model. For the prime p , the oracle keeps a table T of field elements indexed by *handles*, initially empty, together with a counter, initially 0. A handle holding $x \in \mathbb{F}_p$ stands for the group element with discrete logarithm x , written $[x]$, so that $[1]$ is the generator and $[0]$ the identity, and we use additive notation throughout. Every party interacts with T through three queries.

- *Assignment*. On $x \in \mathbb{F}_p$, the oracle stores x at the handle given by the counter, returns that handle, and increments the counter.
- *Group operation*. On $(i_1, i_2, a_1, a_2) \in \mathbb{Z}^2 \times \mathbb{F}_p^2$, if handles i_1 and i_2 hold x_1 and x_2 the oracle stores $a_1 x_1 + a_2 x_2$ at the handle given by the counter, returns that handle, and increments the counter; if either of i_1, i_2 is unassigned it returns $\perp$ and the counter does not move.

- *Zero test.* On $i \in \mathbb{Z}$, if handle i holds x the party receives 1 when $x = 0$ and 0 otherwise; if handle i is unassigned it receives $\perp$.

We write Q for the total number of assignment and group-operation queries a party makes. Zero tests do not count toward Q . Group elements are never handed to a party: they live in T , and a party names one by its handle. A single group-operation query already forms an arbitrary linear combination of two handles, so scalar multiples and negations each require one group-operation query.

Handles are thus issued in the order of the assignment and group-operation queries, and a handle, once assigned, keeps its value. By convention the encryptor queries first, so that the generator receives handle 0 and the ciphertext handles $1, \dots, m + k$, and that layout is public. Advice may depend on the statement, but not on the encryption randomness.

2.2 Boolean-Factor Matrix Reduction

Let φ be a Boolean circuit. Fix its Boolean quadratic encoding

$$q_1(z) = \dots = q_J(z) = 0, \quad z \in \{0, 1\}^{t_\varphi},$$

where the equations impose gate consistency and acceptance. An AND gate gives $z_k - z_i z_j = 0$, a NOT gate gives $z_k + z_i - 1 = 0$, and acceptance gives $z_{\text{out}} - 1 = 0$. A satisfying assignment of φ extends efficiently to a solution z of these equations, and every solution yields one.

Theorem 1 (Boolean-factor matrix reduction [HS26]). *Let N be an integer with $N \geq \max\{2, |\varphi|, t_\varphi\}$, and define*

$$R = \lfloor \log N \rfloor, \quad m = (N + 1)(2NR + 1) \binom{2R}{R}.$$

For every prime $p > \max\{2^N, 2NR\}$, a deterministic algorithm

$$(M_1, \dots, M_k) \leftarrow \text{Reduce}(\varphi, 1^N, p)$$

computes an ordered basis of a space $\mathcal{S}_{\varphi, N, p} \subseteq \mathbb{F}_p^{m \times m}$ in time $\text{poly}(N, \log p)$. The following properties hold.

- (i) *For every satisfying assignment w of φ , a deterministic polynomial-time algorithm*

$$(\mathbf{a}, \mathbf{u}, \mathbf{v}) \leftarrow \text{Lift}(\varphi, 1^N, p, w)$$

outputs $\mathbf{a} \in \mathbb{F}_p^k$ and $\mathbf{u}, \mathbf{v} \in \mathbb{F}_p^m$ such that

$$\sum_{j=1}^k a_j M_j = \mathbf{u} \mathbf{v}^\top \neq 0, \quad \mathbf{v} \in \{0, 1\}^m, \quad \text{wt}(\mathbf{v}) \leq N + 1.$$

- (ii) *If φ is unsatisfiable, every nonzero $M \in \mathcal{S}_{\varphi, N, p}$ has $\text{rank}(M) \geq R + 1$.*

In particular, $m = O(N^4 \log N)$, $k \leq m^2$, and $2^R \leq N$.

3 The Witness Encryption Construction in [HS26]

The construction is that of Hair and Sahai [HS26, Section 3]. For an input $(1^\lambda, \varphi)$, define

$$n = \max\{2, \lambda, |\varphi|\}, \quad N = \max\{n, t_\varphi\},$$

$$R = \lfloor \log N \rfloor, \quad m = (N + 1)(2NR + 1) \binom{2R}{R}.$$

Pad the quadratic encoding with unused variables to obtain N variables. The prime p is a parameter of the scheme and is required to satisfy $2^m \leq p < 2^{m+1}$.

Encryption $\text{Enc}(1^\lambda, \varphi, b)$. On input a message $b \in \{0, 1\}$, execute the following steps.

1. For the given prime p , compute

$$(M_1, \dots, M_k) \leftarrow \text{Reduce}(\varphi, 1^N, p).$$

2. Independently sample

$$\mathbf{s} \leftarrow \mathbb{F}_p^m, \quad \mathbf{r} \leftarrow \{0, \dots, m-1\}^m, \quad \boldsymbol{\eta} \leftarrow \mathbb{F}_p^k,$$

drawing all three for either message, including the one the branch below does not use.

3. By assignment queries, in this order, assign 1, then s_i for $i \in \{1, \dots, m\}$, and then, for $j \in \{1, \dots, k\}$, the value

$$\begin{cases} \mathbf{s}^\top M_j \mathbf{r}, & b = 0, \\ \eta_j, & b = 1, \end{cases}$$

receiving handles $0, 1, \dots, m+k$ in return. The ciphertext ct is the resulting state of T ; we write X_i for the group element at handle i and Y_j for the one at handle $m+j$.

Decryption $\text{Dec}(1^\lambda, \varphi, w, \text{ct})$. Execute the following steps.

1. If w is not a satisfying assignment of φ , return $\perp$. If ct is malformed, return $\perp$.
2. Use the handles of $X_1, \dots, X_m, Y_1, \dots, Y_k$, namely $1, \dots, m+k$, and compute

$$(M_1, \dots, M_k) \leftarrow \text{Reduce}(\varphi, 1^N, p), \quad (\mathbf{a}, \mathbf{u}, \mathbf{v}) \leftarrow \text{Lift}(\varphi, 1^N, p, w).$$

3. Compute

$$H = \sum_{i=1}^m u_i X_i, \quad Z = \sum_{j=1}^k a_j Y_j, \quad B = (m-1) \text{wt}(\mathbf{v}).$$

4. For $t \in \{0, 1, \dots, B\}$ form $Z - tH$ and zero-test it. Return 0 if any of these tests succeeds, and 1 otherwise.

Theorem 2 (Correctness and false-statement security [HS26, Proposition 3.3 and Theorem 3.5]). *For every satisfying assignment w of φ and every $b \in \{0, 1\}$,*

$$\Pr[\text{Dec}(1^\lambda, \varphi, w, \text{Enc}(1^\lambda, \varphi, b)) \neq b] \leq (m(N+1) + 1)2^{-m} = 2^{-\Omega(m)}.$$

For every unsatisfiable φ and every generic adversary $\mathcal{A}$ making at most Q assignment and group-operation queries in total,

$$\begin{aligned} & |\Pr[\mathcal{A}(1^\lambda, \varphi, \text{Enc}(1^\lambda, \varphi, 0)) = 1] - \Pr[\mathcal{A}(1^\lambda, \varphi, \text{Enc}(1^\lambda, \varphi, 1)) = 1]| \\ & \leq O\left((Q + m + k)^2 \left(m^{-(R+1)} + 2^{-m}\right)\right). \end{aligned}$$

The probabilities include the encryption randomness and the coins of the algorithms.

4 Extractable Witness Encryption

A witness-encryption scheme for Boolean circuit satisfiability is a pair of uniform probabilistic polynomial-time algorithms

$$\text{ct} \leftarrow \text{Enc}(1^\lambda, \varphi, b), \quad b' \leftarrow \text{Dec}(1^\lambda, \varphi, w, \text{ct}),$$

where $b \in \{0, 1\}$ and $b' \in \{0, 1, \perp\}$. We take $|\varphi| \leq \lambda$, which loses nothing because the security parameter may be padded to $\max\{2, \lambda, |\varphi|\}$. An *adversary* is a non-uniform probabilistic polynomial-time algorithm, that is, a family of circuits of size polynomial in λ ; its advice, its total number of assignment and group-operation queries, and its input length are all bounded by that size. Write

$$\text{Adv}_{\mathcal{A}}(\lambda, \varphi) = |\Pr[\mathcal{A}(1^\lambda, \varphi, \text{Enc}(1^\lambda, \varphi, 0)) = 1] - \Pr[\mathcal{A}(1^\lambda, \varphi, \text{Enc}(1^\lambda, \varphi, 1)) = 1]|.$$

Definition 1 (Correctness). *The scheme is correct if there is a negligible function κ such that for every λ , every circuit φ with $|\varphi| \leq \lambda$, every satisfying assignment w of φ and every $b \in \{0, 1\}$,*

$$\Pr[\text{Dec}(1^\lambda, \varphi, w, \text{Enc}(1^\lambda, \varphi, b)) = b] \geq 1 - \kappa(\lambda).$$

Definition 2 (Extractability). *The scheme is extractable if for every efficient (possibly non-uniform) adversary $\mathcal{A}$ there are an efficient extractor E and a negligible function $\nu_{\mathcal{A}}$ such that*

$$\Pr[\varphi(w) = 1 \mid w \leftarrow E(1^\lambda, \varphi)] \geq \text{Adv}_{\mathcal{A}}(\lambda, \varphi) - \nu_{\mathcal{A}}(\lambda) \quad (1)$$

for every $\lambda \in \mathbb{N}$ and every circuit φ with $|\varphi| \leq \lambda$.

Extractor $E(1^\lambda, \varphi)$, with $\mathcal{A}_\lambda$ hardwired

1. Compute N, R, m as in Section 3.
2. $(M_1, \dots, M_k) \leftarrow \text{Reduce}(\varphi, 1^N, p)$ for the given prime p .
3. Sample $\mathbf{s} \leftarrow \mathbb{F}_p^m$ and $\mathbf{r} \leftarrow \{0, \dots, m-1\}^m$ and assign an encryption of 0: handle 0 gets 1, handle i gets s_i , and handle $m+j$ gets $\mathbf{s}^\top M_j \mathbf{r}$.
4. Run $\mathcal{A}_\lambda$ on $(1^\lambda, \varphi)$, answering its queries as the oracle would, and record alongside each handle the coefficient vector of the affine form it holds, handles $0, \dots, m+k$ starting from $1, \mathbb{X}_1, \dots, \mathbb{X}_m, \mathbb{Y}_1, \dots, \mathbb{Y}_k$.
5. For every zero test that returns 1 on a handle whose form $F = c + \boldsymbol{\alpha}^\top \mathbb{X} + \boldsymbol{\gamma}^\top \mathbb{Y}$ is nonzero:
 - (a) set $M_\gamma = \sum_{j=1}^k \gamma_j M_j$;
 - (b) if $\gamma = 0$ or $\text{rank}(M_\gamma) \geq R+1$, continue;
 - (c) $w \leftarrow \text{Recover}(\varphi, 1^N, p, M_\gamma)$ via Lemma 1;
 - (d) if $\varphi(w) = 1$, return w .
6. Return $\perp$.

Fig. 1. The extractor of Theorem 3. Affine forms are defined in its proof. The recovery algorithm `Recover` is provided by Lemma 1.

Remark 1. In the generic-group variant, the probabilities in the extractability definition also range over the fresh encryption randomness, and E answers the table queries of the adversary it runs from a table of its own.

Remark 2 (Uniformity of the extractor). Although Definition 2 asks only for $\forall \mathcal{A} \exists E$, we give a uniform polynomial-time procedure that maps a description of $\mathcal{A}_\lambda$ to a description of the corresponding extractor E_λ . The extractor uses the adversary’s code and advice; the procedure constructing it is the same for every adversary.

5 Unconditional Extractability of the WE in [HS26] in the GGM

Theorem 3 (Extractability). *The construction of Section 3 is correct in the sense of Definition 1, with error at most $2^{-\Omega(m)}$, and extractable in the sense of Definition 2, with Figure 1 as the extractor E . If $\mathcal{A}$ makes at most Q assignment and group-operation queries in total, then for each circuit φ the extraction error is bounded by*

$$\nu_{\mathcal{A}}(\lambda, \varphi) = O\left((Q + m + k) \left(m^{-(R+1)} + 2^{-m}\right)\right). \quad (2)$$

Proof. Correctness follows from Theorem 2: its error bound is negligible uniformly over $|\varphi| \leq \lambda$, since $m \geq 2N^2 \geq 2\lambda^2$. We prove extractability. The argument revolves around two main facts:

1. on any M_γ (see Fig. 1) such that $\gamma \neq 0$ and $\text{rank}(M_\gamma) \leq R$, the call `Recover`($\varphi, 1^N, p, M_\gamma$) outputs a satisfying assignment of φ ;

2. by recording the coefficients of the adversary's group oracle queries, E finds such a γ with probability at least $\text{Adv}_{\mathcal{A}}(\lambda, \varphi) - \nu_{\mathcal{A}}(\lambda, \varphi)$.

The first fact follows from Lemma 1 applied to the span of the basis, an observation also used in [HS26], albeit in contrapositive form, to prove standard semantic security. We use the same observation but with the explicit goal of building an extractor. The second fact relies on the generic group model and basic linear algebra: a distinguishing adversary must make some zero test succeed on a handle it built as a combination whose coefficients are not all zero, and by Lemma 2 those coefficients implicitly describe a matrix M_{γ} of rank at most R , except with negligible probability. The extractor records the coefficients of each handle as it answers the queries, so it can read off the corresponding γ when such a zero test succeeds.

Fix a circuit φ and the parameters N, R, m, k of Section 3. Let p be the given prime, $(M_1, \dots, M_k) \leftarrow \text{Reduce}(\varphi, 1^N, p)$ the basis, and $\mathbf{s}, \mathbf{r}$ the encryption randomness used by the $b = 0$ branch. Let $h = m + k + Q + 2$.

Let $\mathbb{X} = (\mathbb{X}_1, \dots, \mathbb{X}_m)$ and $\mathbb{Y} = (\mathbb{Y}_1, \dots, \mathbb{Y}_k)$ be formal variables for the exponents of $X_1, \dots, X_m$ and $Y_1, \dots, Y_k$. An *affine form* is a vector $F = (c, \alpha, \gamma) \in \mathbb{F}_p \times \mathbb{F}_p^m \times \mathbb{F}_p^k$, written $F = c + \alpha^\top \mathbb{X} + \gamma^\top \mathbb{Y}$. It records how, through group oracle queries, a handle was built out of handle 0 and the ciphertext, and the party that built the handle knows its coefficients. Writing $M_{\gamma} = \sum_{j=1}^k \gamma_j M_j$, the exponent such a handle carries when the encrypted message is 0 is

$$\text{ev}_0(F) = c + \mathbf{s}^\top \alpha + \mathbf{s}^\top M_{\gamma} \mathbf{r}.$$

Let $\mathcal{T} = (F_0, F_1, \dots)$ be the list of affine forms of Lemma 3. Up to the first successful zero test on a nonzero affine form, F_i records how handle i was built in the real execution.

Fact 1: we can obtain a witness from a low-rank M_{γ} . Let $\gamma \in \mathbb{F}_p^k$ satisfy $\gamma \neq 0$ and $\text{rank}(M_{\gamma}) \leq R$. The basis $M_1, \dots, M_k$ is linearly independent, so $M_{\gamma} \neq 0$, and it lies in $\mathcal{S}_{\varphi, N, p}$ by definition of that span. It is therefore a nonzero element of $\mathcal{S}_{\varphi, N, p}$ of rank at most R . We can thus apply Lemma 1, and use the algorithm $\text{Recover}(\varphi, 1^N, p, M_{\gamma})$ to obtain a satisfying assignment of φ . So whenever step 5 of Fig. 1 reaches such a γ , the test in step 5(d) passes and E returns a satisfying assignment.

Fact 2: E finds a suitable γ whenever $\mathcal{A}$ is a good distinguisher. Write Succ for the event in which $E(1^\lambda, \varphi)$ successfully returns a w with $\varphi(w) = 1$. We prove

$$\Pr[\text{Succ}] \geq \text{Adv}_{\mathcal{A}}(\lambda, \varphi) - \nu_{\mathcal{A}}(\lambda, \varphi), \quad \nu_{\mathcal{A}}(\lambda, \varphi) = h(m^{-(R+1)} + 2p^{-1}). \quad (3)$$

We first show that a successful zero test with a suitable γ occurs with the required probability, then check that E records this γ and passes it to step 5. We define two main events:

- a “vanishing” event (under encryption of 0), one where the adversary is effectively obtaining some new information through a successful zero test:

$$\text{Van}_0 : \quad \mathcal{A} \text{ zero-tests a handle } i \text{ with } F_i \neq 0 \text{ and } \text{ev}_0(F_i) = 0.$$

We call the form tested at the first such successful query the *witnessing form*. This is the event of Lemma 3, so the witnessing form lies in $\mathcal{T}$ and the union bound below may range over $\mathcal{T}$.

- an “accident”, one where the vanishing happens for a reason that carries no low-rank matrix with it:

$$\text{Acc} : \quad \text{Van}_0 \wedge (\gamma = 0 \text{ or } \text{rank}(M_\gamma) \geq R + 1).$$

The case $\gamma = 0$ is a form that involves no $\mathbb{Y}$ coordinate, so $M_\gamma = 0$; the case $\text{rank}(M_\gamma) \geq R + 1$ is one whose matrix sits above the rank threshold. Neither is of use to Fact 1.

What we want is a vanishing that is not an accident. The target is therefore $\text{Van}_0 \wedge \neg \text{Acc}$, on which the witnessing form has a γ meeting the hypothesis of Fact 1.

We obtain the above from two bounds. The first is Lemma 4, which says that distinguishing forces a success in the $b = 0$ world: $\Pr[\text{Van}_0] \geq \text{Adv}_{\mathcal{A}}(\lambda, \varphi) - hp^{-1}$.

The second bound is $\Pr[\text{Acc}] \leq h(m^{-(R+1)} + p^{-1})$. To prove it, we need to apply Lemma 2 to fixed forms. The forms queried in the real execution are chosen adaptively: the adversary’s choices depend on oracle answers, which in turn depend on $(\mathbf{s}, \mathbf{r})$. Instead, we use the symbolic list $\mathcal{T}$ from Lemma 3. This list is independent of the encryption randomness and contains the form witnessing Van_0 . Thus, if Acc occurs, some nonzero form in $\mathcal{T}$ with $\gamma = 0$ or $\text{rank}(M_\gamma) \geq R + 1$ must vanish.

Condition on $\mathcal{T} = T$. Every form in T is now fixed, while $(\mathbf{s}, \mathbf{r})$ retain their original independent distributions. For each nonzero $F = (c, \alpha, \gamma) \in T$ of the above kind, Lemma 2 gives

$$\Pr[\text{ev}_0(F) = 0 \mid \mathcal{T} = T] \leq \begin{cases} p^{-1}, & \gamma = 0, \\ m^{-(R+1)} + p^{-1}, & \text{rank}(M_\gamma) \geq R + 1. \end{cases}$$

Indeed, in the first case $M_\gamma = 0$ and $(c, \alpha) \neq (0, 0)$, so the second clause of the lemma applies. In the second case its first clause gives $m^{-\text{rank}(M_\gamma)} + p^{-1} \leq m^{-(R+1)} + p^{-1}$. A union bound over the at most h forms in T therefore yields

$$\Pr[\text{Acc} \mid \mathcal{T} = T] \leq h(m^{-(R+1)} + p^{-1}).$$

This bound is the same for every T , so averaging over $\mathcal{T}$ gives the claimed unconditional bound.

Combining the two bounds, we obtain

$$\Pr[\text{Van}_0 \wedge \neg \text{Acc}] \geq \Pr[\text{Van}_0] - \Pr[\text{Acc}] \geq \text{Adv}_{\mathcal{A}}(\lambda, \varphi) - \nu_{\mathcal{A}}(\lambda, \varphi),$$

where $\nu_{\mathcal{A}}(\lambda, \varphi) = h(m^{-(R+1)} + 2p^{-1})$: one hp^{-1} term comes from Lemma 4, and the remaining terms come from the accident bound. On this event, $\mathcal{A}$ queries a form whose γ satisfies the hypothesis of Fact 1.

Since $h = O(Q + m + k)$ and $p \geq 2^m$, this error has the form stated in (2). To see that it is negligible, recall that h is polynomial in λ . Also, $m \geq 2N^2$ and $R = \lfloor \log N \rfloor$ imply

$$m^{-(R+1)} \leq N^{-2(R+1)} \leq 2^{-2 \log^2 N}.$$

Because $N \geq \lambda$, both this term and $p^{-1} \leq 2^{-m}$ are negligible in λ , even after multiplication by h . For parameter values where $\nu_{\mathcal{A}}(\lambda, \varphi)$ exceeds one, the probability lower bound is vacuous.

Steps 3 and 4 of Fig. 1 run the real encryption-of-0 experiment. By Lemma 3, the coefficients E records agree with $\mathcal{T}$ up to the query witnessing Van_0 . Thus E can read γ directly from the coefficient vector $F_i = (c, \alpha, \gamma)$ of the tested handle. On $\text{Van}_0 \wedge \neg \text{Acc}$, step 5 processes this vector and Fact 1 gives a satisfying assignment. Hence $\Pr[\text{Succ}] \geq \Pr[\text{Van}_0 \wedge \neg \text{Acc}]$, which proves (3).

The error bound is negligible uniformly over $|\varphi| \leq \lambda$. Taking its maximum over these circuits and all admissible primes gives a negligible function $\nu_{\mathcal{A}}(\lambda)$ independent of φ , as required by Definition 2.

The extractor E runs $\mathcal{A}_\lambda$ once and performs polynomial work per query. Each call to **Recover** takes $2^R \text{poly}(N, \log p) \leq N \text{poly}(N, \log p)$ time by Lemma 1. Thus E runs in polynomial time. $\square$

5.1 Supporting Lemmas

Fix φ, N, R, m, p and the basis $M_1, \dots, M_k$ from Theorem 1. Affine forms, M_γ and ev_0 are as in the proof of Theorem 3, and

$$\text{ev}_1(F) = c + \alpha^\top s + \gamma^\top \eta$$

is the exponent a handle carries when the encrypted message is 1. The properties below hold for arbitrary φ , including satisfiable instances.

Lemma 1 (Recovery from a low-rank matrix). *There is a deterministic algorithm **Recover** such that, for every circuit φ and every nonzero $M \in \mathcal{S}_{\varphi, N, p}$ of rank $\rho \leq R$,*

$$w \leftarrow \text{Recover}(\varphi, 1^N, p, M)$$

is a satisfying assignment of φ , in time $2^\rho \text{poly}(N, \log p)$, which is $\text{poly}(N, \log p)$ since $2^\rho \leq 2^R \leq N$.

Proof. This fact is provided by the explicit candidate-list conclusion at the end of [HS26, Section 4.5, p. 18]. The approach there is as follows. The reduction of Theorem 1 builds matrices with $N + 1$ columns and appends $m - (N + 1)$ zero columns to make them square, which leaves the rank unchanged [HS26, Section 4.4]. Discard those columns of M and let $\mathbf{c}_0, \dots, \mathbf{c}_N$ be the ones that remain, in the ordering the reduction fixes. Choose a column basis $\mathbf{c}_{i_1}, \dots, \mathbf{c}_{i_\rho}$ and write

$$\mathbf{c}_i = \sum_{j=1}^{\rho} \gamma_{ij} \mathbf{c}_{i_j}, \quad L_i(\mathbf{y}) = \sum_{j=1}^{\rho} \gamma_{ij} y_j \quad (0 \leq i \leq N),$$

both obtained from a single Gaussian elimination, which also yields ρ . The algorithm $\text{List}(\varphi, 1^N, p, M)$ outputs

$$\mathcal{L}(M) = \{(L_1(\mathbf{e}), \dots, L_N(\mathbf{e})) : \mathbf{e} \in \{0, 1\}^\rho\} \cap \{0, 1\}^N,$$

a list of at most 2^ρ Boolean vectors, and [HS26] shows that it contains a solution of the padded quadratic encoding of φ . Testing the list against that encoding and keeping the circuit-input coordinates of the first vector that passes yields a satisfying assignment of φ , which is the output of **Recover**. The enumeration dominates the cost and runs for $2^\rho \leq 2^R \leq N$ iterations. $\square$

Lemma 2 (Rank-sensitive collision bound). *Let $M \in \mathbb{F}_p^{m \times m}$ have rank ρ , and fix $c \in \mathbb{F}_p$ and $\alpha \in \mathbb{F}_p^m$. For independent uniform $\mathbf{s} \leftarrow \mathbb{F}_p^m$ and $\mathbf{r} \leftarrow \{0, \dots, m-1\}^m$, where $m < p$,*

$$\Pr[c + \mathbf{s}^\top \alpha + \mathbf{s}^\top M \mathbf{r} = 0] \leq m^{-\rho} + p^{-1}.$$

If $M = 0$ and $(c, \alpha) \neq (0, 0)$, the probability is at most p^{-1} .

Proof. This is [HS26, Lemma 3.4, p. 8]. Its proof picks ρ independent columns of M and fixes the coordinates of $\mathbf{r}$ outside them; at most one choice of the remaining coordinates gives $M\mathbf{r} = -\alpha$, and their grid values are distinct field elements because $m < p$, so that event has probability at most $m^{-\rho}$. Off it the expression is a nonconstant affine form in uniform $\mathbf{s}$ and vanishes with probability p^{-1} . $\square$

The following lemma lets us analyze the adversary using a symbolic execution that does not depend on the encryption randomness. This execution agrees with the real one until a nonzero affine form passes a zero test (so any distinguishing advantage must come from such a vanishing event).

Lemma 3. *Let $\mathcal{A}$ make at most Q assignment and group-operation queries in total and let $h = m + k + Q + 2$. Consider a symbolic execution of $\mathcal{A}$ in which handles store affine forms in $\mathbb{X}, \mathbb{Y}$ instead of their evaluated exponents. The initial handles store $1, \mathbb{X}_1, \dots, \mathbb{X}_m$, and $\mathbb{Y}_1, \dots, \mathbb{Y}_k$, in that order; assignment queries store constant forms, group operations combine forms, and a zero test returns 1 exactly when the tested form has all coefficients zero.*

Let $\mathcal{T} = (F_0, F_1, \dots)$ list the forms in this execution, in handle order. It has at most h entries and, for fixed public inputs, depends only on the coins of $\mathcal{A}$. In particular, $\mathcal{T}$ is independent of $(\mathbf{s}, \mathbf{r}, \boldsymbol{\eta})$.

For each $b \in \{0, 1\}$, run $\mathcal{A}$ on a real encryption of b using the same adversary coins as in the symbolic execution. The two executions make the same queries and issue the same handles until the first zero test on a handle i with $F_i \neq 0$ but $\text{ev}_b(F_i) = 0$. Up to and including this query, every real handle i holds $\text{ev}_b(F_i)$. At this query the real zero test returns 1 while the symbolic zero test returns 0. Write

$\text{Van}_b :$ *such a zero-test query occurs in the real execution.*

If Van_b does not occur, the two executions have identical queries, answers and outputs. Consequently,

$$\text{Adv}_{\mathcal{A}}(\lambda, \varphi) \leq \Pr[\text{Van}_0] + \Pr[\text{Van}_1], \quad (4)$$

Proof. First consider the symbolic execution alone. It starts with $m + k + 1$ forms, and each assignment or group-operation query adds at most one form. Thus $|\mathcal{T}| \leq m + k + Q + 1 \leq h$. All symbolic answers are determined by the coefficient vectors and the order in which handles are issued. Once the public inputs and the coins of $\mathcal{A}$ are fixed, these answers fix its subsequent queries, even though $\mathcal{A}$ is adaptive. The entire list $\mathcal{T}$ is therefore fixed without using any encryption randomness, proving its independence of $(\mathbf{s}, \mathbf{r}, \boldsymbol{\eta})$.

Now fix b , the coins of $\mathcal{A}$, and the encryption randomness, and compare the symbolic and real executions. Initially, the real handle corresponding to each form F_i holds $\text{ev}_b(F_i)$ by the definition of encryption. As long as the answers agree, $\mathcal{A}$ makes the same next query in both executions. We check that this agreement and the correspondence of handle values are preserved:

- An assignment of x stores the constant form x in the symbolic table and the exponent $x = \text{ev}_b(x)$ in the real table. Both return the same next handle.
- A group operation with coefficients a_1, a_2 stores $a_1 F_{i_1} + a_2 F_{i_2}$ in the symbolic table. Its evaluation is $a_1 \text{ev}_b(F_{i_1}) + a_2 \text{ev}_b(F_{i_2})$, exactly the exponent stored in the real table, and both return the same next handle.
- A zero test on an assigned handle returns 1 symbolically if $F_i = 0$, in which case $\text{ev}_b(F_i) = 0$ and the real answer is also 1. If $F_i \neq 0$, both answers are 0 unless $\text{ev}_b(F_i) = 0$. In that case the symbolic answer is 0 and the real answer is 1: this is precisely the query witnessing Van_b .

Queries involving unassigned handles return $\perp$ in both executions, since the same handles have been issued. Induction on the queries proves the claimed agreement up to the first differing zero-test answer, including the value of the tested handle. If Van_b never occurs, the executions receive identical answers throughout and produce the same output.

Finally, let π be the probability that the symbolic execution outputs 1, and let π_b be that probability in the real encryption-of- b experiment. The symbolic execution does not use b , so the same π applies to both messages. With the shared adversary coins above, the outputs can differ only on Van_b , giving

$$|\pi_b - \pi| \leq \Pr[\text{the two executions have different outputs}] \leq \Pr[\text{Van}_b].$$

The triangle inequality now yields

$$\begin{aligned} \text{Adv}_{\mathcal{A}}(\lambda, \varphi) &= |\pi_0 - \pi_1| \\ &\leq |\pi_0 - \pi| + |\pi_1 - \pi| \leq \Pr[\text{Van}_0] + \Pr[\text{Van}_1]. \end{aligned}$$

A similar comparison with a symbolic execution appears in [HS26, proof of Theorem 3.5, pp. 8–9]. $\square$

Lemma 4 (Distinguishing forces a vanishing event in the $b = 0$ world).
In the setting of Lemma 3,

$$\text{Adv}_{\mathcal{A}}(\lambda, \varphi) \leq \Pr[\text{Van}_0] + hp^{-1}.$$

Proof. By Lemma 3,

$$\text{Adv}_{\mathcal{A}}(\lambda, \varphi) \leq \Pr[\text{Van}_0] + \Pr[\text{Van}_1].$$

It therefore suffices to prove $\Pr[\text{Van}_1] \leq hp^{-1}$.

To handle the adversary's adaptive queries, condition on the symbolic list $\mathcal{T} = T$ of Lemma 3. This fixes all forms in T . Since $\mathcal{T}$ is independent of the encryption randomness, $(\mathbf{s}, \boldsymbol{\eta})$ remains uniform on $\mathbb{F}_p^{m+k}$ after conditioning.

For any nonzero $F = (c, \boldsymbol{\alpha}, \boldsymbol{\gamma}) \in T$, we have

$$\text{ev}_1(F) = c + \boldsymbol{\alpha}^\top \mathbf{s} + \boldsymbol{\gamma}^\top \boldsymbol{\eta}.$$

If some entry of $\boldsymbol{\alpha}$ or $\boldsymbol{\gamma}$ is nonzero, fix all coordinates of $(\mathbf{s}, \boldsymbol{\eta})$ except the corresponding one. Exactly one value of that remaining uniform coordinate makes $\text{ev}_1(F) = 0$, so the probability is p^{-1} . If both vectors are zero, then $c \neq 0$ because $F \neq 0$, and the form never vanishes. In either case,

$$\Pr[\text{ev}_1(F) = 0 \mid \mathcal{T} = T] \leq p^{-1}.$$

If Van_1 occurs, Lemma 3 ensures that the form tested at the first differing answer belongs to T . Thus some nonzero form in T must evaluate to zero. A union bound over its at most h forms gives

$$\Pr[\text{Van}_1 \mid \mathcal{T} = T] \leq \sum_{\substack{F \in T \\ F \neq 0}} \Pr[\text{ev}_1(F) = 0 \mid \mathcal{T} = T] \leq hp^{-1}.$$

Averaging over $\mathcal{T}$ yields $\Pr[\text{Van}_1] \leq hp^{-1}$, which proves the claim when substituted into the first inequality. $\square$

References

- CV26. Liyan Chen and Vinod Vaikuntanathan. Public-coin differing-input obfuscation from KOALA and iO. Cryptology ePrint Archive, Paper 2026/2346, 2026. <https://eprint.iacr.org/2026/2346>.
- HS26. Isaac M. Hair and Amit Sahai. Witness encryption via prime-order generic groups. arXiv:2609.18275v1, 2026. <https://arxiv.org/abs/2609.18275v1>.
- Jin26. Zhengzhong Jin. Witness encryption for NP from SNARGs and groups. Cryptology ePrint Archive, Paper 2026/2063, 2026. <https://eprint.iacr.org/2026/2063>.
- Mau05. Ueli Maurer. Abstract models of computation in cryptography. In *Cryptography and Coding, 10th IMA International Conference*, volume 3796 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2005.