A Practical Key-Recovery Attack on GRAFHEN

Jules Dumezy

University of Paris-Saclay, CEA-List, France
jules.dumezy@cea.fr

Abstract. We give a structural cryptanalysis of GRAFHEN, a proposed group-based fully homomorphic encryption scheme whose public rewriting rules hide a group representation used for decryption. Under the designers' isomorphism hypothesis, the compatible representations in a finite target T form one free $\mathrm{Aut}(T)$-orbit. Thus every representative induces the same zero predicate. For a fixed publication, the published key-class expression is the expected trial count of uniform verified sampling rather than a lower bound on reconstruction work. More generally, a homomorphism ψ gives the correct zero predicate exactly when $\ker\psi\cap\Gamma\subseteq\Lambda$, where Γ is the ciphertext set of the presented group and Λ its zero-ciphertext subgroup. After reconstruction, we give a finite-sample distinguishing bound for independent zero queries and an independent challenge. It requires neither uniform sampling nor generation of the full zero subgroup. For matched query and challenge distributions, the target alone puts the advantage above $1/2$ after 14 queries at degree 7 and 26 at degree 11. Our attack, HEnbane, combines automorphism-reduced search over complete group elements with algebraic forcing from short public consequences. At the level of a generic algorithmic schema, an uncapped whole-element branching fallback makes reconstruction complete for every finite effectively enumerable target with enumerable element-orbit representatives. The released software instantiates this schema for its supported symmetric-group targets and the fixed target $\mathrm{PSL}_2(343)$. The two-seed force chains used by all supplied instances are a fast path rather than a completeness assumption. It decrypts all five released symmetric challenges, with a slowest ten-run mean of 25.4 seconds. It also solves the author-supplied $\mathrm{SL}_2(343)$ challenge for the revised parameters by reconstructing in the $\mathrm{PSL}_2(343)$ quotient and recovering the ten central lift signs over $\mathbb{F}_2$, averaging 37.7 seconds end to end without exact-order hints.

Keywords: Cryptanalysis · Group-Based Cryptography · Homomorphic Encryption

1 Introduction

GRAFHEN is a proposed fully homomorphic encryption scheme without noise [8]. It encodes plaintexts in a finite group and represents ciphertexts by words in a public rewriting system. The secret key is a tuple of generators of a known finite group. For the originally recommended parameters, the designers report

approximately $11!^4 \approx 2^{101}$ equivalence classes of keys and cite this cardinality as evidence of security [8, Section 8]. Two keys are equivalent when their generator tuples induce the same full rule set, and the estimate counts these classes under the hypothesis that the secret-key map is an isomorphism.

Appendix B of [8] studies an unsuccessful computational attack based on Todd–Coxeter enumeration. It uses monoid-mode Todd–Coxeter to approximate the least one-sided congruence containing $(w_0, 1)$ for every supplied zero word w_0 and classifies ciphertexts through the resulting identity-state loops rather than recovering a key. The reported experiments do not break a semidirect instance with $n > 7$, and the appendix states that key recovery appears to be a significantly harder problem.

Our contribution. We attack the key directly, and we determine what a recovered key implies for the stated security goal.

Our first contribution is structural. Under the same isomorphism hypothesis, GRAFHEN's Lemma 27 already proves, for an arbitrary finite target T, that full-rule-equivalent generating tuples form free $\mathrm{Aut}(T)$-orbits [8, Lemma 27]. Lemma 1 sharpens this for a fixed publication: every generating tuple satisfying the published relations lies in the secret tuple's orbit and recognizes the same zero words. The designers' expression approximates the pre-publication class count. For a fixed publication, that same expression is the expected number of uniform verified trials, not a lower bound or an average-case estimate for relation-guided reconstruction. We also characterize the more general maps that decrypt: a homomorphism ψ gives the correct zero predicate exactly when $\ker \psi \cap \Gamma \subseteq \Lambda$ (Proposition 2). Theorem 1 turns such a map and fresh zero encryptions into a finite-sample IND-CPA distinguisher without word statistics or generation of the full zero subgroup.

Our second contribution is computational. HEnbane searches complete target-group elements modulo automorphisms and uses short public consequences to force the remaining generators. If forcing reaches a fixed point, it branches on another complete target element and resumes propagation. An uncapped traversal is therefore complete without assuming that a two-seed force chain exists. For S_7 and S_{11} it works with whole permutations, whereas for the author-supplied $\mathrm{SL}_2(343)$ challenge it reconstructs in $\mathrm{PSL}_2(343)$ and lifts the aligned pair exactly. It recovers equivalent keys and decrypts all five released symmetric challenges, including the originally recommended S_{11} semidirect construction and its public-key variant. The $\mathrm{SL}_2(343)$ challenge yields representation and label recovery, and Section 2 explains why its central labels do not instantiate a valid GRAFHEN encoding. These are fixed-instance results and do not determine the running-time distribution over arbitrary generated keys. The implementation and experimental artifacts are available at `https://github.com/jdumezy/henbane`.

Related work. Geraud-Stewart gives a chosen-plaintext distinguishing attack that seeks statistical bias in reduced-word lengths and letter frequencies, without first recovering a representation [5]. The designers dispute that analysis, and we neither rely on nor resolve that dispute. Subsequent independent work by

Geraud-Stewart presents *Maverick*, a different equivalent-key recovery method that also recovers every released challenge [6]. For the symmetric instances, Maverick obtains regular actions by group-mode Todd–Coxeter, recognizes their natural degree-n actions, and aligns the two copies through the mixed rules. HEnbane constructs neither a coset or congruence action nor a regular action: it searches directly in compact target representations and forces generators from sound public consequences. The methods also use labeled samples differently: our finite-sample guarantee removes Maverick's assumption that the zero samples generate the transported zero subgroup (Sections 4.5 and 3.3). The structural results of Section 3 hold whichever reconstruction method is used.

The July revision of GRAFHEN keeps the construction and proposes $\mathrm{PSL}_2(343)$ as the target group [9]. Its complexity discussion estimates our earlier point-table prototype, described in Appendix A, by substituting the natural projective-line degree 344 into that algorithm's worst-case bound. The reconstruction challenge supplied by the authors uses exact words in $\mathrm{SL}_2(343)$, whose simple central quotient is $\mathrm{PSL}_2(343)$. Section 6.2 examines that revision and the reconstruction problem it poses.

2 The GRAFHEN construction

We recall only the parts of GRAFHEN needed for the attack and follow the notation of [8]. Let T be a finite evaluation group generated by d elements, and let $A = \{x_1, \dots, x_d\}$ be an alphabet. A rewriting rule is a pair of positive words (u, v) over A that evaluate to the same element of T. The public relation set R defines the presented group

$$\mathcal{P} = \langle A \mid u = v \text{ for } (u, v) \in R \rangle.$$

The secret tuple is $\mathbf{s} = (s_1, \dots, s_d)$ and its evaluation homomorphism is

$$\sigma_{\mathbf{s}} \colon \mathcal{P} \longrightarrow T, \qquad \sigma_{\mathbf{s}}(x_i) = s_i.$$

The tuple generates T, so $\sigma_{\mathbf{s}}$ is surjective. As in GRAFHEN's key-equivalence analysis, our structural results state explicitly when they assume that $\sigma_{\mathbf{s}}$ is an isomorphism.

A plaintext ring K is homomorphically encoded in a group E by an injective map Enc together with two polynomial maps of E^2 into E that realize the ring operations on $\mathrm{Enc}(K)$ [8]. Following GRAFHEN we normalize $\mathrm{Enc}(0) = 1_E$. The scheme chooses a subgroup $L \leq T$ and a surjective homomorphism $\pi \colon L \to E$. We write $N = \ker \pi$ for the masking kernel, which GRAFHEN denotes by Z. The group elements that encrypt a fixed encoded plaintext form a coset of N, and the encryptions of zero are exactly the elements of N. We write $C = \pi^{-1}(\mathrm{Enc}(K))$ for the set of ciphers. Since Enc is only required to be injective, this definition does not force $\mathrm{Enc}(K)$ to be a subgroup of E, and GRAFHEN accordingly calls C the group of ciphers only in the examples where it is one. It is one whenever addition is realized by the group multiplication, since Enc is then a monomorphism of

$(K, +)$ into E. That is the case for the encodings the design paper constructs explicitly, and for the one it reports using in its own degree-11 benchmark, but its originally recommended parameters leave the choice of encoding open. None of our arguments assumes that C is a subgroup.

Write $\Gamma = \sigma_{\mathbf{s}}^{-1}(C)$ and $\Lambda = \sigma_{\mathbf{s}}^{-1}(N)$. Then Λ is a subgroup of $\mathcal{P}$, and Γ is a union of Λ-cosets because $NC = C$. Every element of Γ normalizes Λ, because N is normal in L. A ciphertext word decrypts to zero exactly when its class lies in Λ, and decryption factors through the set of left (equivalently, right) Λ-cosets contained in Γ, which is in bijection with $\mathrm{Enc}(K)$. When C is a subgroup of L, these are the usual subgroup, normal subgroup, and quotient group. Neither definition requires $\sigma_{\mathbf{s}}$ to be injective. With the secret permutation representation, decryption consists of evaluating a ciphertext word and recognizing the appropriate subgroup or coset.

The originally recommended single-copy construction takes $T = S_{11}$ with $d = 5$ generators. Its rules are admissible in the sense of Section 7.1 of [8]: both sides use the whole alphabet, both sides meet the prescribed minimum length, and the two sides have neither a common prefix nor a common suffix. The rewriting system is non-confluent, and the designers report about twenty million rules for one copy.

The originally recommended parameters combine two copies of a base group G by the semidirect construction of Section 7.2 of [8], so the full evaluation group is now $T = G \rtimes G$. Let A and B be disjoint alphabets for the two copies. The second copy acts on the first by conjugation, giving $G \rtimes G \cong G \times G$ [8, Lemma 33]. The construction uses the following homomorphism.

$$f \colon G \rtimes G \longrightarrow G, \qquad f(g, h) = gh.$$

Besides the relations internal to A and B, the public rewriting system contains a rule $b\,a \longrightarrow w_{b,a}\,b$ for every $a \in A$ and $b \in B$, where $w_{b,a}$ is an A-word. Writing $\alpha_a, \beta_b \in G$ for the two secret base-group generators, the rule is chosen so that $w_{b,a}$ evaluates to $\beta_b \alpha_a \beta_b^{-1}$. It therefore imposes

$$\beta_b \alpha_a \beta_b^{-1} = \alpha(w_{b,a}). \tag{1}$$

The designers report just under forty million rules for their originally recommended key. These design-paper figures are compared with the exact released-challenge counts in Table 2 and Section 6.1.

GRAFHEN also describes a public-key conversion [8]. For a fixed underlying key, this conversion changes the production of ciphertexts but not the secret generator tuple or its rewriting relations. The representation-recovery attack therefore applies to both modes.

The author-supplied $\mathrm{SL}_2(343)$ *challenge.* The July revision proposes $G_{\mathrm{PSL}} = \mathrm{PSL}_2(343)$ as the new base group and its subgroup $E = \mathrm{PSL}_2(7) \cong \mathrm{PSL}_3(2)$ as the encoding group, equipped with the noncentral homomorphic encoding described in the design paper [9]. The GRAFHEN authors separately supplied a challenge that does not instantiate that encoding directly. They described it as an

instance using $\mathrm{PSL}_2(7^3)$ in place of S_{11} and explained that, for implementation reasons, its rewriting equalities are exact in

$$G_{\mathrm{SL}} = \mathrm{SL}_2(\mathbb{F}_{7^3}).$$

We call it the author-supplied G_{SL} challenge. They specified[1] that every supplied class-0 word evaluates to I and every supplied class-1 word evaluates to the central scalar matrix

$$-I = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The two labeled evaluations are therefore distinct in G_{SL} but coincide in

$$G_{\mathrm{PSL}} = G_{\mathrm{SL}}/Z(G_{\mathrm{SL}}) = \mathrm{PSL}_2(343).$$

This is why quotient reconstruction must be followed by an exact central lift.

These labels do not define a valid GRAFHEN homomorphic encoding of $\mathbb{F}_2$. They coincide in the proposed target G_{PSL}. Retaining the sign in G_{SL} instead makes the nonidentity class representative central in any encoding-group image. GRAFHEN proves that such a representative cannot be central [8, Lemma 12]. The supplied instance nevertheless defines well-posed representation-recovery and central-label-recovery problems, which we analyze under the authors' terminology "challenge".

3 The structural break

Throughout this section T is an arbitrary finite target group published as part of the parameters. We call T effectively represented when its elements have polynomial-size encodings and its group operations and subgroup membership are polynomial-time computable. The symmetric targets satisfy this definition in their natural actions [13]. For the G_{SL} challenge, group operations use the compact matrix representation and subgroup tests, when needed, use its explicitly computable degree-344 action. We make no uniform complexity claim here for a family in which the field size grows.

3.1 Compatible representations and aligned copies

Definition 1 (Reconstruction over a target group). $\mathsf{REP}(T, A, R)$ *is the problem of finding a tuple* $(\alpha_1, \ldots, \alpha_d) \in T^d$ *that generates* T *and satisfies* $\alpha(u) = \alpha(v)$ *for every* $(u, v) \in R$.

Definition 2 (Two-copy reconstruction). *Given a base group* G*, alphabets* A *and* B*, each of size* d*, internal relation sets* R_A *and* R_B*, and mixed rules*

[1] The GRAFHEN authors supplied the challenge and stated its label semantics in private correspondence on 4 August 2026.

$V\,u = u'\,V'$ with $V, V' \in B^*$ and $u, u' \in A^*$, find generating tuples $\alpha, \beta \in G^d$ satisfying the internal relations and

$$\beta(V)\alpha(u) = \alpha(u')\beta(V') \tag{2}$$

for every mixed rule. The cross rules of Equation (1) *have $V = V' = b$, $u = a$, and $u' = w_{b,a}$.*

Under the isomorphism hypothesis, GRAFHEN proves that the action is free and that its orbits are the full-rule equivalence classes [8, Lemma 27]. The distinction in the next lemma is that $\mathsf{REP}(T, A, R)$ imposes only the published relations R.

Lemma 1 (Fixed-presentation orbit). *Assume that the secret evaluation map $\sigma_{\mathbf{s}} \colon \mathcal{P} \to T$ is an isomorphism. Then the solution set of $\mathsf{REP}(T, A, R)$ is exactly the orbit of $\mathbf{s}$ under the componentwise action of $\mathrm{Aut}(T)$ on T^d. This action is free on generating tuples, so the solution set has exactly $|\mathrm{Aut}(T)|$ elements.*

Proof. The secret key is a solution, since every published rule holds under evaluation. Conversely let α be a solution. Because α satisfies R, it induces a homomorphism $\psi_\alpha \colon \mathcal{P} \to T$, and ψ_α is surjective because the components of α generate T. The hypothesis on $\sigma_{\mathbf{s}}$ gives $|\mathcal{P}| = |T| < \infty$, so ψ_α is an isomorphism and $\theta = \psi_\alpha \circ \sigma_{\mathbf{s}}^{-1}$ is an automorphism of T with $\alpha_i = \theta(s_i)$ for every i. In the other direction every automorphism carries a solution to a solution, since automorphisms preserve every relation and preserve generation. The action is free because an automorphism fixing each s_i fixes the subgroup they generate, which is T. □

Every solution in Lemma 1 is an equivalent key under this definition, because it is obtained from the secret tuple by an automorphism and therefore induces the same full rule set. Let $\mathbf{SK}$ and $\mathbf{R}$ denote the random secret tuple and relation set output by key generation. For every publication R_0 in the support of $\mathbf{R}$ for which the isomorphism hypothesis holds, the posterior support is contained in this orbit, so $H_\infty(\mathbf{SK} \mid \mathbf{R} = R_0) \leq \log_2 |\mathrm{Aut}(T)|$. Equality holds exactly when the posterior is uniform on that orbit, which is ensured, for example, by a uniform prior together with an automorphism-invariant publication distribution. Let $g_d(T)$ be the number of generating d-tuples in T^d. The exact pre-publication number of full-rule equivalence classes is $g_d(T)/|\mathrm{Aut}(T)|$. The designers' expression $|T|^d/|\mathrm{Aut}(T)|$ uses the approximation $g_d(T) \approx |T|^d$. By contrast, a fixed presentation satisfying the hypothesis has exactly $|\mathrm{Aut}(T)|$ solutions among all $|T|^d$ tuples. For targets with efficient uniform sampling and candidate verification, Lemma 1 also gives a baseline Las Vegas search: a uniform tuple in T^d is a solution with probability $|\mathrm{Aut}(T)|/|T|^d$, so independent verified trials succeed after $|T|^d/|\mathrm{Aut}(T)|$ trials in expectation. This is an expected upper bound for an unstructured search, not a lower bound or a prediction of the cost of relation-guided reconstruction.

Proposition 1 (Alignment of the two copies). *Assume that G is centerless, that the compatible generating tuples for each internal relation set form one $\mathrm{Aut}(G)$-orbit, and that the mixed rules include the cross rules $b\,a = w_{b,a}\,b$ of Section 2. Then the compatible solutions (α, β) form one free orbit under the diagonal action of $\mathrm{Aut}(G)$, and either internal tuple uniquely determines the other.*

Proof. Suppose first that the A-tuple is fixed. For each $b \in B$, Equation (1) asks for a simultaneous conjugator from $(\alpha_a)_{a\in A}$ to $(\alpha(w_{b,a}))_{a\in A}$. Two such conjugators differ by an element centralizing every α_a. Because these elements generate the centerless group G, the conjugator is unique when it exists.

Conversely, fix β and let α and α' be compatible solutions. By the internal-orbit hypothesis, there is $\theta \in \mathrm{Aut}(G)$ with $\alpha' = \theta \circ \alpha$. Applying θ to the cross relation for α and comparing it with the cross relation for α' shows that $\theta(\beta_b)^{-1}\beta_b$ centralizes every $\theta(\alpha_a)$. These elements generate G, so centerlessness gives $\theta(\beta_b) = \beta_b$ for every b. The β_b generate G, hence θ is the identity and $\alpha' = \alpha$.

Finally, for any two compatible pairs, the internal-orbit hypothesis supplies an automorphism carrying the second component of one pair to that of the other. After applying its inverse diagonally, the second components agree, so the uniqueness just proved makes the first components agree as well. The action is free because either tuple generates G. □

Remark 36 of [8] anticipates the A-first direction. The B-first direction is what HEnbane uses in Section 4.5. Here Proposition 1 characterizes the full set of compatible pairs as one diagonal orbit, while Maverick turns the same centerless uniqueness principle into an algorithmic template for aligning the two natural-degree actions returned by its recognition stage [6].

3.2 Which representations decrypt

For a general recovered homomorphism, equality on the published rules alone need not prove equivalence in GRAFHEN's full-rule sense. The following criterion identifies exactly when such a map nevertheless gives the correct zero predicate.

Proposition 2 (Exact kernel criterion). *Let Λ be a subgroup of $\mathcal{P}$ and let Γ be a union of left Λ-cosets, as in Section 2. Let $\psi\colon \mathcal{P} \to Q$ be a homomorphism and put $H_\psi = \psi(\Lambda)$. Then*

$$\big(\forall w \in \Gamma,\quad \psi(w) \in H_\psi \iff w \in \Lambda\big) \quad\iff\quad \ker\psi \cap \Gamma \subseteq \Lambda. \tag{3}$$

If Q is effectively represented and generators of H_ψ are known, this gives a polynomial-time zero test for every ciphertext.

Proof. Since $\psi^{-1}(H_\psi) = \Lambda \ker\psi$, the kernel condition implies $(\Lambda\ker\psi)\cap\Gamma = \Lambda$: if $w = \lambda k \in \Gamma$ with $\lambda \in \Lambda$ and $k \in \ker\psi$, then $k = \lambda^{-1}w$ lies in $\Lambda\Gamma = \Gamma$, so

$k \in \ker\psi \cap \Gamma \subseteq \Lambda$. Thus $\psi(w) \in H_\psi$ if and only if $w \in \Lambda$ for $w \in \Gamma$. Conversely, any $k \in (\ker\psi \cap \Gamma) \setminus \Lambda$ maps to $1 \in H_\psi$ and violates this equivalence. The computational statement follows from subgroup membership in Q. □

For a single-copy instance under the isomorphism hypothesis, Lemma 1 makes every reconstruction solution an equivalent key whose associated map has trivial kernel. Under the hypotheses of Proposition 1, a compatible semidirect pair (α, β) is obtained from the secret pair by one common automorphism θ. The map $\Theta(g,h) = (\theta(g), \theta(h))$ is an automorphism of $G \rtimes G$, and the pair canonically lifts to the key map

$$\widehat{\psi}(a) = (\alpha_a, 1), \qquad \widehat{\psi}(b) = (1, \beta_b)$$

into $G \rtimes G$. Let $\mathcal{P}_{A,B}$ denote the group presented by all internal and mixed public relations, and let $\widehat{\sigma}_{\mathbf{s}} \colon \mathcal{P}_{A,B} \to G \rtimes G$ be its full secret evaluation map. More precisely, $\widehat{\psi} = \Theta \circ \widehat{\sigma}_{\mathbf{s}}$, so it induces the same full rule set and is an equivalent key. Decryption uses the associated map $\psi = f \circ \widehat{\psi}$ into G. Since $f \circ \Theta = \theta \circ f$, one has $\psi = \theta \circ f \circ \widehat{\sigma}_{\mathbf{s}}$. In the originally recommended construction, $C \subseteq L \leq G \rtimes G$, $N \leq L$, and π factors through f [8]. Hence, with $\Gamma = \widehat{\sigma}_{\mathbf{s}}^{-1}(C)$ and $\Lambda = \widehat{\sigma}_{\mathbf{s}}^{-1}(N)$,

$$\ker\psi \cap \Gamma = \widehat{\sigma}_{\mathbf{s}}^{-1}(\ker f \cap C) \subseteq \widehat{\sigma}_{\mathbf{s}}^{-1}(N) = \Lambda.$$

Therefore Proposition 2 applies although ψ is noninjective. Consequently, HEnbane recovers an equivalent semidirect key together with its noninjective associated decryption map. When generators of $\psi(\Lambda)$ are known, subgroup membership decrypts the bit-encryption mode exactly. The conclusion also applies to the public-key conversion, which changes ciphertext production but not the underlying key or rewriting relations.

3.3 CPA security after reconstruction

GRAFHEN adopts CPA security as its goal [8]. For group-homomorphic encryption, earlier work characterizes IND-CPA security through subgroup membership [7,1]. Applying their full equivalences to GRAFHEN requires care because ciphertexts are words in a non-confluent rewriting magma and the word sampler is nonuniform. The direct implication below avoids those additional assumptions.

For a finite group J, define its subgroup-chain height by

$$h(J) = \max\{\ell : 1 = J_0 < J_1 < \cdots < J_\ell \leq J\}.$$

The arithmetic function Ω counts prime factors with multiplicity.

Remark 1 (Advantage normalization). Let $\mathcal{D}_b$ denote the adversary's view when the challenge encrypts b. We use max-one distinguishing advantage $\mathrm{Adv}_{\mathrm{dist}} = |\Pr[\mathcal{A}(\mathcal{D}_0) = 1] - \Pr[\mathcal{A}(\mathcal{D}_1) = 1]|$. GRAFHEN Section 2.1 states the prediction goal without fixing a numerical normalization. For a uniform challenge bit, $\mathrm{Adv}_{\mathrm{dist}}$ is twice the excess guessing probability $\Pr[\mathcal{A}(\mathcal{D}_b) = b] - 1/2$.

Theorem 1 (Finite-sample CPA distinguishing after reconstruction). *Let Q be finite and effectively represented, and let $\psi\colon \mathcal{P} \to Q$ be a homomorphism available through efficiently evaluable images of the public generators, fixed before the encryption queries, and satisfying*

$$\ker\psi \cap \Gamma \subseteq \Lambda.$$

Assume that the zero-encryption queries are independent with common image distribution μ under ψ, and that a zero challenge is independent of them with image distribution ν. Put $J_\mu = \langle \operatorname{supp}\mu \rangle$ and

$$\Delta = \Delta_{\mathrm{TV}}(\mu,\nu) = \max_{\mathcal{U}\subseteq Q} |\mu(\mathcal{U}) - \nu(\mathcal{U})|.$$

After m independent encryptions of zero, there is a polynomial-time IND-CPA adversary against the bit-encryption mode with max-one distinguishing advantage at least

$$\max\left\{0, 1 - \frac{h(J_\mu)}{m+1} - \Delta\right\}. \tag{4}$$

Moreover,

$$h(J_\mu) \le h(\psi(\Lambda)) \le h(Q) \le \Omega(|Q|) \le \lfloor \log_2 |Q| \rfloor.$$

Proof. Let $\xi_1, \dots, \xi_m$ be the images under ψ of the m zero encryptions, let ξ_{m+1} be one further independent sample from μ, and put $D_i = \langle \xi_1, \dots, \xi_i \rangle$, with $D_0 = 1$. The last sample is an auxiliary device in the analysis and is not a query. Whenever $\xi_i \notin D_{i-1}$, the strict subgroup inclusion $D_{i-1} < D_i$ has index at least two. Every D_i lies in J_μ, so on every sample path,

$$\sum_{i=1}^{m+1} \mathbf{1}[\xi_i \notin D_{i-1}] \le h(J_\mu) =: h_\mu.$$

Write $p_i = \Pr[\xi_i \notin D_{i-1}]$. Independence gives $p_i = \mathbb{E}[\mu(Q \setminus D_{i-1})]$. The subgroups D_i increase along every sample path, so the sequence (p_i) is non-increasing. It follows that

$$(m+1)p_{m+1} \le \sum_{i=1}^{m+1} p_i \le h_\mu.$$

The adversary submits $(0,1)$, forms D_m from its zero-encryption queries, and outputs 0 exactly when $\psi(c^*) \in D_m$. For a zero challenge its error probability is $\mathbb{E}[\nu(Q \setminus D_m)]$, the expectation being over the queries. The challenge is independent of D_m, so comparing ν with μ on the set $Q \setminus D_m$ path by path gives

$$\mathbb{E}[\nu(Q \setminus D_m)] \le \mathbb{E}[\mu(Q \setminus D_m)] + \Delta = p_{m+1} + \Delta.$$

For a challenge encrypting one, Proposition 2 gives $\psi(c^*) \notin \psi(\Lambda)$, hence $\psi(c^*) \notin D_m$, so the adversary is always correct. Its advantage is therefore at least $1 - p_{m+1} - \Delta$, which proves Equation (4). Neither a uniformity assumption on μ or ν nor generation of all of $\psi(\Lambda)$ is used, and matched distributions give $\Delta = 0$.

The support of μ lies in $\psi(\Lambda)$, which proves the first two height inequalities. For any strict chain $1 = U_0 < \cdots < U_k \leq U$, the indices $[U_i : U_{i-1}]$ are at least two and their product is $|U_k|$, which divides $|U|$. Consequently $k \leq \Omega(|U|)$, and $\Omega(|U|) \leq \lfloor \log_2 |U| \rfloor$ because every prime factor is at least two. Applying this observation to $U = Q$ completes the displayed chain of bounds. □

Corollary 1 (Composing reconstruction and distinguishing). *Suppose a time-τ reconstruction algorithm returns a certified homomorphism satisfying the kernel criterion with probability p_τ over key generation and its internal coins. Compose it with the distinguisher of Theorem 1, and have the resulting adversary guess fairly when reconstruction fails. Let $h_{\max}$ be a deterministic upper bound on $h(J_\mu)$, and δ one on Δ, for every successful reconstruction outcome. Under the hypotheses of Theorem 1, the resulting adversary has advantage at least*

$$p_\tau \max\left\{0, 1 - \frac{h_{\max}}{m+1} - \delta\right\}$$

and runs in time $\tau + \mathrm{poly}(|\mathsf{pub}|, m)$, where $|\mathsf{pub}|$ is the encoded length of the published instance. If τ and $h_{\max}$ are polynomial in the security parameter and δ is a constant below one, choosing $m \geq 2h_{\max}/(1-\delta)$ gives advantage at least $p_\tau(1-\delta)/2$, so non-negligible p_τ violates IND-CPA for the corresponding key-generation family. For an already reconstructed fixed key one may set $p_\tau = 1$, obtaining a concrete distinguisher. For a fixed successful reconstruction, if the sampled subgroup is certified to equal $\psi(\Lambda)$, the distinguisher has advantage exactly one.

Proof. For every fixed successful outcome, with output one denoting nonmembership, the signed gap is

$$\Pr[\mathcal{A}(\mathcal{D}_1) = 1] - \Pr[\mathcal{A}(\mathcal{D}_0) = 1] = 1 - \mathbb{E}[\nu(Q \setminus D_m)] \geq 0.$$

On failure, the independent fair guess has signed gap zero. Thus averaging over keys and reconstruction coins cannot cancel advantages of opposite signs and gives the factor p_τ. The stated running time is immediate. If reconstruction itself uses labeled ciphertexts, independent sample splitting fixes ψ before the m queries used by the theorem. □

For each of the five released symmetric instances, the associated practical map lands in $Q = S_n$. The Cameron–Solomon–Turull formula gives

$$h(S_n) = \left\lceil \frac{3n}{2} \right\rceil - \mathrm{wt}_2(n) - 1,$$

where $\mathrm{wt}_2(n)$ is the number of ones in the binary expansion of n [3]. Consequently, without any masking information,

$$h(J_\mu) \leq h(S_7) = 7 \quad \text{or} \quad h(J_\mu) \leq h(S_{11}) = 13$$

at degrees 7 and 11, respectively. Thus using $m = 14$ fresh zero queries gives $\mathrm{Adv}_{\mathrm{dist}} \geq 8/15 > 1/2$ for every reconstructed degree-7 key, while $m = 26$

gives $\mathrm{Adv}_{\mathrm{dist}} \geq 14/27 > 1/2$ at degree 11. These target-only bounds require no information about the masking kernel N, and they hold under the fixed-map and independence hypotheses of Theorem 1 with matched query and challenge distributions. At degree 11, the exact height improves both the order-factor bound $\Omega(11!) = 16$ and the generic value $\lfloor \log_2(11!) \rfloor = 25$.

The degree-11 benchmark sampler uses an encoding for which $\mathrm{Enc}(0) = 1_E$ and writes the masked image as the product of the encoded message with a random $r_{\mathrm{mask}} \in S_5$ [8]. For a zero message the encoded-message factor is therefore the identity, so the masked image is r_{mask} and the associated recovered G-valued map $\psi = \theta \circ f \circ \widehat{\sigma}_{\mathbf{s}}$ has $J_\mu \leq \theta(S_5)$. The chain

$$1 < \langle (12)(34) \rangle < V_4 < A_4 < S_4 < S_5$$

and the upper bound $\Omega(120) = 5$ show that $h(S_5) = 5$. Thus $m = 10$ fresh zero queries give $\mathrm{Adv}_{\mathrm{dist}} \geq 1 - 5/11 = 6/11 > 1/2$ for the stated benchmark sampler, provided the challenge is drawn by the same procedure as the queries, so that $\Delta = 0$. Equivalently, for a uniform challenge bit the success probability is at least $17/22$ and the excess guessing probability is at least $3/11$.

4 HEnbane reconstruction and decryption

We instantiate the reconstruction problem of Definition 1 in two target families. For S_7 and S_{11}, the unknowns are complete permutations considered modulo simultaneous conjugation. For the revised target, the unknowns are projective 2×2 matrices over $\mathbb{F}_{343}$ and target automorphisms reduce the seed search. In both cases, short public consequences determine generators algebraically whenever possible. At a forcing fixed point, the complete algorithm branches on another whole target element and resumes forcing. The released instances finish on the faster two-seed force-chain path, while uncapped branching makes termination and completeness independent of that instance property. The search therefore takes place on group elements rather than on individual images in a permutation action.

4.1 Sound relation preparation and forcing

The bounded closure is a streaming, deliberately incomplete consequence extractor. Its corpus $\mathcal{W}$ consists of the positive words that have actually been inserted: short sides of internal public relations, words exposed by cancellation, cyclic rotations of relators exposed by such cancellation, and the empty word. The implementation admits relation sides of length at most 13 and emits consequences of total side length at most 18. It maintains an equivalence relation on $\mathcal{W}$ initialized by the admitted public equalities and applies the following group-valid deductions whenever the relevant words are present:

$$u \equiv v \Rightarrow xu \equiv xv, ux \equiv vx, \qquad pus \equiv pvs \Rightarrow u \equiv v, \qquad uv \equiv 1 \Rightarrow vu \equiv 1.$$

On every class merge it compares members from the two old classes, strips their maximal common prefix and suffix, and links extensions already visible in the corpus. When this cancellation exposes a positive relator equal to the empty word, it also inserts its cyclic rotations. The current implementation makes one streaming pass and does not replay merges after later extensions arrive. Its output can therefore depend on rule order, affecting completeness and performance but not soundness.

A separate streaming selector retains up to 5,000 short public relations per internal copy. It also retains up to 2,000 shortest mixed relations having the coupling form $Vu = uV'$. The driver combines these relations with the closure output, cancels common contexts once more, preserves a shortest two-seed force chain when one exists (otherwise a longest force prefix), and fills the remaining slots by increasing total length. The experimental search portfolio contains at most 3,500 equalities.

Lemma 2 (Soundness of consequence extraction). *Every equality emitted by the bounded closure holds in every group satisfying the public relations.*

Proof. The initial equalities are public relations. Transitivity and multiplication on either side preserve equality. Cancellation is valid because left and right multiplication are bijections in a group. Finally, $uv = 1$ gives $v = u^{-1}$ and hence $vu = 1$. The streaming restrictions can omit consequences but cannot create a false one. □

Write $u = u_1 \cdots u_{\ell_u}$ and $v = v_1 \cdots v_{\ell_v}$, and form the signed relator $u_1 \cdots u_{\ell_u} v_{\ell_v}^{-1} \cdots v_1^{-1}$. Suppose that its only occurrence of an unassigned generator is x_j^{ϵ} with $\epsilon \in \{-1, 1\}$ and that the known prefix and suffix products are Π_0 and Σ_0. Evaluation gives

$$\Pi_0 x_j^{\epsilon} \Sigma_0 = 1, \qquad x_j^{\epsilon} = \Pi_0^{-1} \Sigma_0^{-1}.$$

Thus one multiplication and inversion calculation determines the complete group element assigned to x_j. A force chain applies this identity repeatedly until all generators are assigned. Any relation that then fails rejects the seed pair.

The closure also yields sound order multiples. The equality $x_i^k = 1$ proves only $\mathrm{ord}(x_i) \mid k$, while $x_i^{k_1} = x_i^{k_2}$ with $k_1 < k_2$ proves $\mathrm{ord}(x_i) \mid (k_2 - k_1)$. Neither proves an exact order. The implementation retains the smallest observed positive multiple and tests candidate orders only for divisibility. For symmetric targets, applying the sign homomorphism to the retained relations gives a homogeneous linear system over $\mathbb{F}_2$. An exact parity pattern is imposed only when its nullspace is one-dimensional: generation of S_n excludes the all-even solution and leaves the unique nonzero vector. All these filters are consequences of the selected equalities and the target-generation requirement.

4.2 Complete whole-element branch and force

Let $c_{\mathrm{Aut}}(T) = |\mathrm{Aut}(T)\backslash T|$ be the number of automorphism orbits on individual target elements. The complete traversal enumerates one representative of each

such orbit for a first generator, applies sound filtering and forcing to a fixed point, and, if generators remain unassigned, branches over every element of T for one of them. It resumes forcing after every branch. A complete assignment is accepted only after it is proved to generate T and is checked against the entire public relation set R, not merely the bounded search portfolio R_{sel}.

Theorem 2 (Complete whole-element reconstruction). *Let T be a finite effectively enumerable group, let $d \geq 1$, and suppose representatives of the $\mathrm{Aut}(T)$-orbits on T can be enumerated. Let R_{sel} consist of sound consequences of the public relations, and suppose exact relation and generation tests in the chosen target representation are effective. The uncapped depth-first whole-element branch-and-force traversal terminates. It returns a certified solution if and only if $\mathsf{REP}(T, A, R)$ has one, and it visits at most*

$$c_{\mathrm{Aut}}(T) \sum_{j=0}^{d-1} |T|^j$$

partial whole-element assignments. For $|T| \geq 2$, this is $O\big(c_{\mathrm{Aut}}(T)|T|^{d-1}\big)$. For the trivial group, the displayed sum equals d. Its search stack stores $O(d)$ target elements in addition to the relation storage and the workspaces of the target enumerator and certifier.

Proof. Finiteness gives the displayed geometric bound and termination. Soundness follows from accepting only complete tuples that pass generation and all of R. For completeness, let $\mathbf{s} = (s_1, \ldots, s_d)$ be a solution. The first-generator enumeration contains $t_1 = \theta(s_1)$ for some $\theta \in \mathrm{Aut}(T)$. Follow that branch. Whenever the traversal branches, it enumerates $\theta(s_i)$. Whenever a single occurrence forces x_i, evaluation of group terms commutes with θ, so the uniquely forced value is $\theta(s_i)$. Sound order, parity, membership, and relation filters cannot remove this path. It therefore reaches $\theta(\mathbf{s})$, which passes exact certification. Depth-first traversal stores only the current assignment and recursion trail. □

Let E_T bound the cost of producing the next target element or first-generator orbit representative, including the corresponding enumeration bookkeeping. If F_{sel} bounds forcing and filtering work per visited state and V_R bounds complete relation and generation certification, the corresponding time bound is

$$O\left(c_{\mathrm{Aut}}(T) \sum_{j=0}^{d-1} |T|^j \, (E_T + F_{\mathrm{sel}}) + c_{\mathrm{Aut}}(T)|T|^{d-1} \, V_R\right).$$

This separates the combinatorial state count from target-specific enumeration and certification costs. In particular, R_{sel} may be empty. This removes propagation but does not affect termination or completeness. The production driver retains wall, node, and survivor limits and may report *inconclusive*. A single absolute deadline covers preprocessing, candidate certification, coupling, alignment, and nested reconstruction, while native backend failures are reported separately from exhaustive termination. Its explicit exhaustive mode removes those limits and implements the traversal of Theorem 2.

4.3 Whole-permutation search for S_7 and S_{11}

For $n \in \{7, 11\}$, every automorphism of S_n is inner. Consequently every solution orbit has a representative in which the first seed g_1 is the canonical permutation of its cycle type. For each cycle type compatible with the parity and order restrictions, the algorithm constructs $C_{S_n}(g_1)$ and enumerates a second seed g_2. It retains g_2 exactly when g_2 is lexicographically minimal in its $C_{S_n}(g_1)$ conjugacy orbit. Fixing g_1 and quotienting by its centralizer therefore removes simultaneous conjugation exactly.

Proposition 3 (Seed-orbit coverage). *Suppose a selected set R_{sel} of sound public consequences contains a two-seed force chain. A traversal with no resource limit and no early termination contains the seed pair of every S_n-conjugacy orbit of generating tuples satisfying R_{sel}, and hence of every solution to $\mathsf{REP}(S_n, A, R)$. It contains at most*

$$\kappa_2(S_n) = \frac{1}{n!} \sum_{g \in S_n} |C_{S_n}(g)|^2 = \sum_{\lambda \vdash n} z_\lambda$$

forced seed-pair trials, where z_λ is the centralizer order of cycle type λ. For $n = 7$ and $n = 11$, these bounds are 5,579 *and* 40,853,520*, respectively, close to $|S_n|^2/|\mathrm{Aut}(S_n)| = |S_n|$.*

Proof. Every first seed is conjugate to exactly one canonical cycle-type representative g_1. After fixing g_1, simultaneous conjugations that preserve it are precisely the action of $C_{S_n}(g_1)$ on the second seed. The lexicographic test retains exactly one representative of each such orbit. A force chain determines every remaining generator uniquely, and all order and parity filters are consequences of the relations. Thus no solution orbit of R_{sel} is discarded by the gauge or by forcing. Burnside's lemma gives the displayed number of simultaneous-conjugacy orbits on S_n^2. The relation filters can only reduce this count. □

The trial bound counts pairs that enter forcing. The present enumerator scans partner permutations and tests minimality against $C_{S_n}(g_1)$, which costs at most $|C_{S_n}(g_1)|$ conjugations per scanned partner. For the successful gauges, the cycle types and centralizer orders are $(4, 2, 1)$ and 8 for the single-copy S_7 instance, $(4, 3)$ and 12 for the semidirect S_7 instance, and $(10, 1)$ and 10 for all three S_{11} instances.

For $d \geq 2$, the pair gauge remains useful when forcing is incomplete. Let b be the maximum number of additional whole-generator branch decisions below a seed pair in the traversed tree. Thus $0 \leq b \leq d - 2$, and a complete traversal visits at most

$$\kappa_2(S_n) \sum_{j=0}^{b} (n!)^j = O\big(\kappa_2(S_n)(n!)^b\big) \tag{5}$$

states after the pair gauge. This is unconditional because at a forcing fixed point the implementation selects an unassigned generator, branches over all of S_n,

and resumes forcing. Equation (5) specializes to Proposition 3's pair trials when $b = 0$.

For a simpler target-only bound, suppose $n \neq 6$ and let $p(n)$ be the partition function. The $\mathrm{Aut}(S_n)$-orbits on individual elements are then precisely the cycle types, so Theorem 2 gives

$$p(n) \sum_{j=0}^{d-1} (n!)^j = O\big(p(n)(n!)^{d-1}\big) . \tag{6}$$

Using the partition asymptotic and Stirling's formula, the leading leaf envelope is

$$p(n)(n!)^{d-1} = 2^{(d-1)n \log_2 n - (d-1)(\log_2 e)n + O(d \log n + \sqrt{n})}.$$

No force-chain hypothesis occurs in either complete bound.

The S_7 instances have two generators, so a seed pair is already a complete candidate. Each S_{11} instance has a three-relation force chain. Table 1 gives the exact chains obtained from bounded closure. The letters shown belong to the copy reconstructed by the orbit search.

Table 1. Whole-permutation force chains for the released S_{11} instances.

instance	seeds	forced equality	
single	(a, d)	c	`aaaadd = dddadaadc`
		b	`bcaccd = dcddac`
		e	`acbbe = dacbbb`
semidirect	(A, B)	C	`ABAAABB = BABAAABC`
		D	`CBBABC = ACBABACD`
		E	`ADADC = EAAACAB`
variant	(A, E)	D	`AEAAEE = DAEEEEEEEA`
		B	`AEEDE = DDDABADA`
		C	`AB = CDBABDBEA`

After forcing, the search evaluates all relations in the portfolio. A native kernel can stream every portfolio survivor. The driver continues past a tuple that fails the separate public-relation pool, generation of S_n by Schreier–Sims [13], or complete-corpus certification. In exhaustive mode, a tuple is accepted only after a streaming verifier checks the complete internal public relation set. Labeled words are checked separately as the experimental decryption certificate. Consequently, Proposition 3 is a coverage statement about the underlying seed traversal, not a claim that a portfolio survivor is already a solution to the full presentation. For every reported experiment, the first kernel survivor also passed the pool, generation, and complete-verifier checks. The bounded portfolio therefore accelerates search but does not serve as the final certificate. Only resource exhaustion in the production configuration yields an inconclusive result.

4.4 Projective-matrix search for the revised target

Projected solution set. The proposed target group is $G_{\mathrm{PSL}} = \mathrm{PSL}_2(343)$ and no particular permutation realization is part of its definition. The exact challenge supplied by the authors uses the central cover $G_{\mathrm{SL}} = \mathrm{SL}_2(343)$, whose center is $Z(G_{\mathrm{SL}}) = \{I, -I\}$ and whose quotient is G_{PSL}. The expensive search takes place in G_{PSL}.

Let $\rho\colon G_{\mathrm{SL}} \to G_{\mathrm{PSL}}$ be the quotient map. Under the isomorphism hypothesis for either internal presentation in the G_{SL} challenge, any generating tuple in G_{PSL} that satisfies its projected relations induces a surjection $G_{\mathrm{SL}} \to G_{\mathrm{PSL}}$. Its kernel has order two. In odd characteristic, a matrix of order two in SL_2 is diagonalizable with eigenvalues in $\{1, -1\}$, and determinant one leaves only the identity and $-I$. Thus $-I$ is the unique involution of G_{SL}, and the kernel equals $Z(G_{\mathrm{SL}})$. Consequently the induced map is $\theta \circ \rho$ for some $\theta \in \mathrm{Aut}(G_{\mathrm{PSL}})$, so the quotient solutions form one free $\mathrm{Aut}(G_{\mathrm{PSL}})$-orbit. This conclusion is conditional on the presentation hypothesis, not on the experimental verification of a candidate.

Orbit-reduced matrix search. Write $q = 343$ and let $F = \mathbb{F}_7[\omega]/(\omega^3 - 2)$, which is a field because 2 is not a cube in $\mathbb{F}_7$. An element of G_{PSL} is represented by a nonsingular matrix over F modulo nonzero scalar multiplication. A canonical representative is obtained by scaling its first nonzero entry to one. Membership in G_{PSL} is equivalent to the determinant being a nonzero square in F. Multiplication, inversion, equality, and relation evaluation are performed on these four field coordinates. The recovered matrices induce the familiar faithful action of G_{PSL} on the 344 points of the projective line. Passing to this action is optional and is not part of reconstruction. In the experiments it is materialized only after recovery for an independent generation test.

The automorphism group used for reduction is

$$\mathrm{Aut}(G_{\mathrm{PSL}}) = \mathrm{P\Gamma L}_2(q) = \mathrm{PGL}_2(q) \rtimes \mathrm{Gal}(F/\mathbb{F}_7).$$

For a nonidentity first seed, the algorithm enumerates companion matrices

$$X_\chi = \begin{pmatrix} 0 & -1 \\ 1 & \chi \end{pmatrix}.$$

The element type (split, nonsplit, or unipotent) and the projective trace invariant $\mathrm{tr}(X_\chi)^2/\det(X_\chi)$ identify its $\mathrm{PGL}_2(F)$ class. Taking the minimum over the three Frobenius iterates merges these classes under the field automorphisms (with repetitions when the invariant lies in a proper fixed field). The search tries both seed orientations, so a generating solution is still covered if exactly one of its seed values is the identity.

Once X is fixed, the implementation uses the subgroup

$$C_X^\circ = \{\eta I + \zeta X : \det(\eta I + \zeta X) \neq 0\}/F^*.$$

It is contained in $C_{\mathrm{PGL}_2(F)}([X])$ and consists of the projective classes having a representative that commutes with X exactly. The two groups agree unless

$[X]$ is an involution. In that exceptional case the full projective centralizer also has a coset represented by matrices that anticommute with X. Using C_X° rather than that full centralizer cannot remove a seed orbit, although it may leave duplicates. Partners having a representative that commutes with X lie in C_X°. This group is abelian, so each such projective class is fixed under conjugation by C_X°. The implementation enumerates all of them that pass the membership and order-multiple filters. Because X is nonscalar, the linear map $Y \mapsto XY - YX$ on $M_2(F)$ has kernel $FI + FX$ and a two-dimensional image. For every remaining partner, the direction of $XY - YX$ therefore lies on a projective line. The group C_X° acts on these $q+1$ directions. Fix a nonzero representative K of one such direction and write

$$\mathcal{F}_K = \{Y \in M_2(F) : XY - YX = K\} = Y_0 + FI + FX.$$

If $c \in C_X^\circ$ stabilizes the direction of K, there is a unique $\lambda_c \in F^*$ such that $cKc^{-1} = \lambda_c K$. Ordinary conjugation need not preserve the exact fiber $\mathcal{F}_K$. The induced action on that fiber is

$$c \star Y = \lambda_c^{-1} cYc^{-1}.$$

Indeed, $X(c \star Y) - (c \star Y)X = K$. The direction stabilizer acts on the one-dimensional space FK through the multiplicative character $c \mapsto \lambda_c$, so

$$\lambda_{c_1 c_2} = \lambda_{c_1} \lambda_{c_2}.$$

Consequently,

$$(c_1 c_2) \star Y = c_1 \star (c_2 \star Y),$$

and $\star$ is an action of the direction stabilizer on $\mathcal{F}_K$. The normalization by λ_c^{-1} does not change the projective class of cYc^{-1}. The implementation quotients $\mathcal{F}_K$ by this normalized action of the direction stabilizer, retains the lexicographically least projective representative in each orbit, and filters for membership in G_{PSL} and the sound order multiples. This enumerates complete matrices without storing the degree-344 action. It removes conjugation by C_X° and the first-seed Frobenius symmetry. A Frobenius automorphism stabilizing the chosen first seed, and the extra projective-centralizer coset when $[X]$ is an involution, may leave duplicate second-seed representatives but cannot remove a solution.

Proposition 4 (Projective seed coverage). *Ignoring resource limits and early termination, the two seed orientations visit a representative of every* $\mathrm{Aut}(G_{\mathrm{PSL}})$*-orbit of ordered seed pairs in which at least one seed is nonidentity. The membership and order-multiple filters cannot remove a seed pair belonging to a solution of the selected sound relations. Residual Frobenius stabilizers and, for an involutory first seed, the unused coset of* C_X° *in the full projective centralizer may cause an orbit to be visited more than once.*

Proof. Every nonidentity element of G_{PSL} is $\mathrm{PGL}_2(F)$-conjugate to a companion matrix X_χ. Its element type and projective trace invariant determine that

conjugacy class, and taking Frobenius orbits of the invariant accounts for the field-automorphism part of $\mathrm{Aut}(G_{\mathrm{PSL}})$. After the first seed is fixed to $[X]$, its remaining inner stabilizer is $C_{\mathrm{PGL}_2(F)}([X])$. The implementation quotients by C_X°, which is the full inner stabilizer except in the involutory case, where using an index-two subgroup only leaves duplicate representatives. If Y has a representative commuting with X, then $[Y] \in C_X^\circ$. Since C_X° is abelian, the implementation covers its orbit by enumerating $[Y]$ directly. Otherwise, the nonzero commutator direction is first reduced to a C_X°-orbit representative. After selecting a representative K for that direction, all exact preimages are the affine fiber $\mathcal{F}_K = Y_0 + FI + FX$. If $cKc^{-1} = \lambda_c K$ for an element of the direction stabilizer, then $Y \mapsto \lambda_c^{-1} cYc^{-1}$ preserves $\mathcal{F}_K$ and induces ordinary projective conjugation. Quotienting by this normalized action therefore gives every remaining C_X°-orbit. Trying both orientations covers a pair with exactly one identity seed. Finally, membership in G_{PSL} and divisibility by sound order multiples are necessary conditions for a solution. □

Together with a sound force chain, Proposition 4 gives the fast pair-gauge path for quotient solutions of the selected relations. When no complete two-seed chain is available, the backend instead invokes the generic traversal of Theorem 2: its first-seed list includes the identity orbit, and a stalled fixed point branches over a duplicate-free streaming enumeration of all 20,176,632 canonical projective matrices in G_{PSL}. The standard class count gives 174 conjugacy classes in G_{PSL} [10], while automorphism orbits can only merge conjugacy classes. Consequently $c_{\mathrm{Aut}}(G_{\mathrm{PSL}}) \le 174$, and for the supplied value $d = 5$ the coarse complete state bound is

$$174 \sum_{j=0}^{4} |G_{\mathrm{PSL}}|^j < 2^{104.6}. \tag{7}$$

This deliberately loose fallback bound is a termination guarantee, not a practical cost prediction and not a claim about the distribution of generated instances.

The production driver imposes time, node, and survivor limits and can therefore return inconclusive. The exhaustive configuration removes these limits, continues past portfolio survivors rejected by later checks, and certifies a quotient tuple against generation and the complete projected internal relation corpus. For the successful uppercase copy, the seeds are (B, D) and the exact force chain is

$$\begin{aligned} \texttt{BBDBBDBDE} &= \texttt{DDDDDB} \quad (E), \\ \texttt{BCE} &= \texttt{EBEBBEEEBB} \quad (C), \\ \texttt{BBCCECD} &= \texttt{EEAE} \quad (A). \end{aligned}$$

The surviving quotient tuple is found without supplied exact orders. It is then checked against the separate public-relation pool and proved to generate G_{PSL}. No generic whole-matrix branch is entered on the supplied challenge, so its observed additional branch depth is $b = 0$.

4.5 Coupling, central lifting, and decryption

Coupling. For a two-copy instance, first reconstruct the uppercase tuple β and then use mixed relations

$$Vu = uV'$$

to recover the lowercase tuple α. Such a relation imposes

$$\beta(V)\alpha(u) = \alpha(u)\beta(V'). \tag{8}$$

For $u = a_i$, the possible values of α_i form a simultaneous conjugator coset. In S_n, a nonempty domain is enumerated by matching cycles of equal length and choosing their relative phases. For $u = a_i a_j$, the constraint $\alpha_i\alpha_j$ lies in the corresponding conjugator coset. The symmetric backend enforces arc consistency by exact set intersection. After propagation it enumerates the Cartesian product of the generator domains and checks every retained mixed relation, the lowercase internal-relation pool, and generation of S_n. This is the practical fast path, with a configured domain-product cap of 2,000,000 tuples. If a domain remains unresolved, the product exceeds that cap, or all domain candidates fail certification, the complete mode reconstructs the lowercase tuple independently and enumerates the inner automorphisms of S_n to align it with the uppercase tuple.

More generally, let $C_G(d)$ denote the cost of one complete internal reconstruction, including certification, let $E_{\mathrm{Aut}}(d)$ bound the cost of enumerating and applying one automorphism to a d-tuple, and let V_{mix} bound the cost of checking all mixed relations for one aligned pair. Under the internal single-orbit hypothesis of Proposition 1, independent reconstruction followed by automorphism alignment is complete and costs at most

$$2C_G(d) + |\mathrm{Aut}(G)|\big(E_{\mathrm{Aut}}(d) + V_{\mathrm{mix}}\big). \tag{9}$$

Indeed, write the independently recovered tuples as $\alpha = \theta_A(\mathbf{s}_A)$ and $\beta = \theta_B(\mathbf{s}_B)$. The enumeration contains $\theta = \theta_B\theta_A^{-1}$, for which $\theta(\alpha) = \theta_B(\mathbf{s}_A)$ is aligned with β and satisfies every mixed relation. When G is centerless and the hypotheses of Proposition 1 hold, this alignment is unique. Thus the domain solver and automorphism enumeration are accelerated and fallback paths of one algorithm, not separate notions of reconstruction.

For the G_{SL} challenge, coupling instead takes place after projection to the centerless group G_{PSL}. Proposition 1 applies there, and the aligned tuples are lifted to G_{SL} only afterward. In G_{PSL}, put $g_V = \beta(V)$ and $g_{V'} = \beta(V')$, and choose matrix representatives $P_V, P_{V'}$ for these two elements. Using ordinary matrix multiplication and a representative M of c, the conjugator equation $g_V c = c g_{V'}$ lifts to

$$P_V M = \lambda M P_{V'}, \qquad \lambda^2 = \frac{\det(P_V)}{\det(P_{V'})}.$$

Because both projective elements lie in G_{PSL}, the determinant ratio is a nonzero square, so there are two possible values of λ in the odd-characteristic field F. For each value, this is a four-variable homogeneous linear system over F in the

entries of M. Its nonsingular projective solutions are filtered for membership in G_{PSL}. The implementation composes transformations from left to right. Its stored matrices use the corresponding opposite-product convention, under which the same constraint appears internally with the products reversed. For the G_{SL} challenge, singleton constraints alone reduce all five lowercase domains to size one. The unique candidate is then checked against every retained mixed relation, the lowercase internal-relation pool, and the quotient-generation condition, so no second orbit search is required. If these domains did not certify a pair, the complete fallback would reconstruct the lowercase quotient tuple independently and enumerate $\mathrm{Aut}(G_{\mathrm{PSL}}) = \mathrm{P\Gamma L}_2(343)$, using projective-linear conjugation followed by one of the three Frobenius maps. Its $6|G_{\mathrm{PSL}}|$ alignment trials are negligible compared with the coarse single-copy fallback in Equation (7).

Central lifting. Projection to G_{PSL} loses the signs of the ten G_{SL} generators. Choose a determinant-one lift of every reconstructed projective matrix and associate a bit ε_i to multiplication by $-I$. Each exact public equality or labeled central word gives a linear equation over $\mathbb{F}_2$. Its right-hand side records whether the chosen base lifts evaluate to I or $-I$. Its coefficient of ε_i is the parity of the total number of occurrences of generator i on the two sides, or in the labeled word. If the resulting consistent system has nullity r, the complete lift step enumerates its affine family of 2^r sign vectors and checks each exact lift against the complete relation corpus. When the reconstruction task also includes labeled examples, those are checked separately for every surviving lift. An inconsistent system rejects the projective pair, while rejection of one affine solution continues with the next. For the G_{SL} challenge, 10,000 labeled class-0 words and 10,000 labeled class-1 words give rank ten and determine all signs uniquely without any public-relation equation. Exact evaluation in G_{SL} then classifies the held-out words. Although the lift solver stops adding relation equations once the sign system has full rank, the separate verifier checks the reported exact lifts against every internal and mixed public relation. Let H_A and H_B be the subgroups generated by the five lifted matrices in the two internal copies. The separate quotient generation tests give $\rho(H_A) = \rho(H_B) = G_{\mathrm{PSL}}$. For either subgroup H, one has $|H| = |G_{\mathrm{PSL}}|\,|H \cap Z(G_{\mathrm{SL}})|$. If H were proper, then $H \cap Z(G_{\mathrm{SL}}) = 1$ and H would have index two in G_{SL}. This is impossible because G_{SL} is perfect, so both exact internal tuples generate G_{SL}. Under the isomorphism hypothesis, Lemma 1 now expresses the two lifted tuples as automorphic images of the two secret tuples. Their induced quotient automorphisms agree by projective alignment, and the natural map $\mathrm{Aut}(G_{\mathrm{SL}}) \to \mathrm{Aut}(G_{\mathrm{PSL}})$ is injective: an element of its kernel would be a central automorphism, which is trivial because G_{SL} is perfect. They therefore arise from one common automorphism and form an equivalent key for the supplied two-copy presentation.

Symmetric-instance decryption. For the symmetric instances, a semidirect word reduces to uV and the recovered G-valued map is

$$\widehat{\psi}(uV) = (\alpha(u), \beta(V)), \qquad \psi(uV) = \alpha(u)\beta(V).$$

The known-zero images generate a subgroup and the known-one images select a disjoint coset. The challenge is classified by subgroup or coset membership. Maverick's equivalent-decryptor result likewise assumes that the zero samples generate the transported zero subgroup [6]. The finite-sample attack of Theorem 1 needs neither this full generation certificate nor a known-one sample.

5 Experiments

Table 2 lists the five released symmetric challenges solved and decrypted by HEnbane [4]. The sixth instance is the author-supplied G_{SL} challenge, reconstructed exactly through the proposed quotient G_{PSL}. As explained in Section 2, its central labels do not constitute a valid homomorphic encoding, so "solved" in the last row means representation and label recovery rather than decryption of a complete FHE instantiation.

Table 2. Per-instance reconstruction results.

target	construction	public relations	result
S_7, $d = 2$	single	42,137	solved
S_7, $d = 2$	semidirect	83,586	solved
S_{11}, $d = 5$	single	22,000,002	solved
S_{11}, $d = 5$	semidirect	44,608,403	solved
S_{11}, $d = 5$	public-key variant	44,608,403	solved
G_{SL}, $d = 5$ per copy	author-supplied challenge	39,285,006	solved

Measurements used an AMD Ryzen 9 9950X with 30 search workers for the symmetric instances and one worker for the G_{SL} challenge. Each row of Table 3 was measured in ten sequential runs with the persistent preprocessing cache disabled. No externally supplied exact generator orders were used. The search still uses sound exponent multiples derived from public relations.

The end-to-end column includes reading the public instance, selecting relations, bounded closure, orbit search, coupled-copy recovery, central lifting when applicable, and writing the solution. The end-to-end measurements exclude the independent complete-corpus verification pass described below. We measured that pass separately over ten runs per instance. Peak resident memory was measured in a separate ten-run repetition by sampling the aggregate RSS of the driver and all descendant processes every 0.5 milliseconds. The table reports the largest observed value rather than an average of per-run peaks. The kernel column isolates orbit search on the resulting bounded portfolio.

The search makes no random choices and uses no randomized restarts. Each worker traverses a fixed range, so the reported deviations measure wall-clock, scheduling, and input-processing noise rather than a heavy-tailed search distribution. Each entry reports the arithmetic mean and sample standard deviation in seconds. Every single-copy orbit search used by the six supplied instances finishes on the two-seed path and reports additional whole-generator branch depth

$b = 0$. The generic fallback is exercised by the artifact's small no-chain regression instances, not by these timing rows. The measurements use the resource-bounded production configuration, whereas the completeness statements concern its uncapped exhaustive mode.

Table 3. Measured resource use on the development machine with the application cache disabled. Wall times are mean ± sample standard deviation over ten runs. Peak RSS is the maximum aggregate process-tree value observed over ten runs.

instance	full reconstruction	search kernel	peak RSS	complete verifier
S_7 single	0.117 ± 0.001	0.002 ± 0.0003	61.7 MiB	0.180 ± 0.005
S_7 semidirect	0.520 ± 0.009	0.002 ± 0.0003	79.4 MiB	0.209 ± 0.004
S_{11} single	22.410 ± 0.242	0.255 ± 0.004	1666.4 MiB	3.485 ± 0.026
S_{11} semidirect	25.413 ± 0.373	0.257 ± 0.002	1703.6 MiB	7.875 ± 0.020
S_{11} variant	25.440 ± 0.162	0.249 ± 0.003	1682.3 MiB	7.891 ± 0.068
G_{SL} challenge	37.660 ± 0.128	3.448 ± 0.026	1685.1 MiB	8.821 ± 0.047

For each symmetric instance, an independent verification proves generation of S_n, checks every published internal and mixed relation, and checks the labeled ciphertexts. For the G_{SL} challenge, an independent finite-field verification checks all 39,285,006 exact relations, the 20,000 public labeled words, and the classification of all 10,000 held-out words. The matrix search without exact-order hints finds the quotient orbit at reduced seed-pair candidate 3,484,145. The exact projective orders $(43, 171, 57, 172, 86)$ are computed from the recovered tuple and are not inputs to these measurements.

6 Security implications and repairs

6.1 Evidence and parameter pressure

Success on the originally recommended S_{11} semidirect parameters shows that 2^{101} key classes do not imply a 2^{101} attack cost [8]. It does not determine how reconstruction scales with the target, generator count, selected relations, or their distribution, so we do not extrapolate a replacement security level beyond the released keys. Corollary 1 makes this fixed-key versus key-generation distinction explicit. For the symmetric targets, Schreier–Sims gives polynomial-time subgroup membership in the permutation degree. For the G_{SL} challenge, the explicit degree-344 action gives a practical test. On the concrete instances studied here, the decisive step is therefore finding a representation compatible with the public presentation. Maverick independently recovers every released challenge and twelve generated S_{11} keys with a different algorithm [6]. This independently corroborates the practical break of the five released symmetric instances. Neither work gives an unconditional asymptotic success or runtime bound for GRAFHEN's filtered key-generation distribution.

For the symmetric target, the private generator tuple has an $O(dn \log n)$-bit image representation, but evaluation uses an explicit public rewriting database whose size is at least linear in the number and total length of its rules [8]. GRAFHEN stops its Froidure–Pin computation early because complete enumeration may be impractical and itself warns that increasing n or d makes rule production expensive. The July revision itself reports no generated key for the proposed projective-linear target. Its benchmark still uses the originally recommended S_{11} construction and measures only post-generation evaluation, so it provides neither a key-generation cost nor a rule count or public-key size for the revised target [9, Section 8]. The G_{SL} challenge gives one concrete reconstruction measurement, but it does not characterize the filtered key-generation distribution or supply a valid homomorphic encoding. Between the released $(n, d) = (7, 2)$ and $(11, 5)$ instances of Table 2, the single-copy relation count grows by a factor of about 522 and the semidirect count by about 534. Because both parameters change and reconstruction cost is presentation-dependent, these figures establish neither an effect attributable to n alone nor an absolute practical limit. A revised parameter claim therefore needs controlled measurements of rule generation, encoded public size, evaluation throughput, and reconstruction success, rather than an exhaustive bound for the earlier point-table algorithm.

Changing the target, degree, or number of generators may raise reconstruction cost but leaves Lemma 1 intact. Admissibility hides short published rules, not necessarily their short consequences, as the force chains in Table 1 illustrate. Additional rules may strengthen propagation as well as evaluation. On centerless targets satisfying the internal-orbit hypothesis of Proposition 1, the semidirect wrapper aligns its copies. For the supplied central lift, the same argument is applied only after projection to the centerless quotient. Changing the masking sampler does not by itself prevent the attack: for every query distribution μ it induces, Theorem 1 gives $h(J_\mu) \le h(\psi(\Lambda)) \le h(S_n)$, so the target-only bounds of 14 queries at degree 7 and 26 at degree 11 remain valid for matched query and challenge sampling under the theorem's fixed-map and independence hypotheses. If the query and challenge distributions differ, the same theorem incurs the additive penalty Δ. Within the isomorphic-presentation design, security therefore requires a falsifiable average-case reconstruction assumption together with practical key-generation and evaluation evidence.

6.2 The $\mathrm{PSL}_2(343)$ proposal and the supplied $\mathrm{SL}_2(343)$ challenge

The July revision proposes the abstract target group $G_{\mathrm{PSL}} = \mathrm{PSL}_2(343)$ [9]. Its analysis of our earlier point-table algorithm uses the natural degree-344 action on the projective line and estimates $O(2^{1,720})$ work by substituting $d = 5$ and $n = 344$ into that algorithm's worst-case bound. Sections 2 and 4.4 distinguish this proposal from the author-supplied G_{SL} challenge, whose central labels do not instantiate a valid GRAFHEN encoding. The experiment therefore measures representation and label recovery for that challenge, not the security of a complete revised FHE scheme.

The quotient argument of Section 4.4 and the lift construction of Section 4.5 use the following group data:

$$\begin{aligned} |G_{\mathrm{PSL}}| &= \frac{343(343^2-1)}{2} = 20{,}176{,}632, \\ |G_{\mathrm{SL}}| &= 2|G_{\mathrm{PSL}}| = 40{,}353{,}264, \\ \mathrm{Aut}(G_{\mathrm{SL}}) \cong \mathrm{Aut}(G_{\mathrm{PSL}}) &\cong \mathrm{P\Gamma L}_2(7^3), \\ |\mathrm{Aut}(G_{\mathrm{SL}})| = |\mathrm{Aut}(G_{\mathrm{PSL}})| &= 6|G_{\mathrm{PSL}}|, \end{aligned}$$

where the factor six is the index-two extension from PSL to PGL and the three field automorphisms [10]. For $d = 5$, under the isomorphism hypothesis for either supplied internal G_{SL} presentation, the quotient argument of Section 4.4 shows that the generating tuples in G_{PSL}^5 satisfying its projected relations form exactly one free $\mathrm{Aut}(G_{\mathrm{PSL}})$-orbit. Uniform verified sampling in G_{PSL}^5 against those projected relations therefore succeeds after

$$\frac{|G_{\mathrm{PSL}}|^5}{|\mathrm{Aut}(G_{\mathrm{PSL}})|} = \frac{|G_{\mathrm{PSL}}|^4}{6} \approx 2^{94.48}$$

trials in expectation. This is also the designers' key-count expression at the proposed target. Direct uniform search against a fixed isomorphic G_{SL} presentation instead has expectation

$$\frac{|G_{\mathrm{SL}}|^5}{|\mathrm{Aut}(G_{\mathrm{SL}})|} = 32\frac{|G_{\mathrm{PSL}}|^5}{|\mathrm{Aut}(G_{\mathrm{PSL}})|} \approx 2^{99.48}.$$

These are unstructured-search baselines, not costs paid by the quotient-and-lift algorithm.

A projective matrix occupies $O(\log |G_{\mathrm{PSL}}|)$ bits, and multiplication, equality, and inversion reduce to field operations on its four entries. If $\|R\|$ denotes the total encoded relation length, Equation (7) and complete certification give the state and relation-processing envelope

$$O\big(174|G_{\mathrm{PSL}}|^4 \,\mathrm{poly}(\|R\|, \log |G_{\mathrm{PSL}}|)\big)$$

for $d = 5$, independently of any permutation degree. This is not a complete bit-complexity bound: target-element enumeration and target-specific generation tests remain additional costs. The actual relations make exhaustive enumeration unnecessary: after automorphism reduction, three matrices are forced from two seeds and the ten central signs are solved linearly. The measurements in Table 3, certified against the complete relation corpus, therefore show that the earlier degree-based point-table bound does not describe this instance.

For comparison, consider a valid GRAFHEN bit instantiation after reconstruction of the proposed target G_{PSL}.

$$\begin{aligned} |G_{\mathrm{PSL}}| &= 2^3 \cdot 3^2 \cdot 7^3 \cdot 19 \cdot 43, \\ |G_{\mathrm{SL}}| &= 2^4 \cdot 3^2 \cdot 7^3 \cdot 19 \cdot 43. \end{aligned}$$

Theorem 1 bounds $h(J_\mu)$ by 10 for maps into G_{PSL} and by 11 for maps into G_{SL}. Under its fixed-map and independence hypotheses, matched query and challenge sampling therefore gives advantage above $1/2$ after 20 or 22 fresh zero queries, respectively. These are target-only guarantees for valid instantiations, not an interpretation of the supplied central labels.

Maverick reports end-to-end recovery through $\mathrm{PSL}_2(169)$ and recognition and coordinatization at $q = 343$ from a generated regular action [6]. Our result reconstructs the author-supplied G_{SL} challenge directly from its relations and completes the exact lift.

6.3 What a structural repair must prevent

Giving $\sigma_{\mathbf{s}}$ a nontrivial kernel invalidates the single-orbit conclusion and is the genuine structural direction noted in Remark 29 of [8]. It is not enough, however, to prevent recovery of the literal secret map. By Proposition 2, any efficiently evaluable homomorphism $\psi\colon \mathcal{P} \to Q$ into a finite effectively represented group satisfying

$$\ker \psi \cap \Gamma \subseteq \Lambda$$

gives the exact zero predicate when generators of $\psi(\Lambda)$ are known, and Theorem 1 gives a sampled distributional test without that full description. A faithful representation into such a group is a sufficient special case because its kernel is trivial, although constructing one from the public presentation may itself be hard.

Pseudo-boundedness, GRAFHEN's randomized empirical test for bounded reduced-word behavior, neither proves that $\mathcal{P}$ is finite or effectively representable nor bounds all reduced-word lengths. GRAFHEN exhibits zero-ciphertext powers of unbounded reduced length [8]. An infinite presented group is also compatible with correctness because decryption only requires the quotient map $\sigma_{\mathbf{s}}$. Thus a repair with $\ker \sigma_{\mathbf{s}} \neq 1$ must make every suitable quotient representation hard to construct, not merely enlarge $\mathcal{P}$.

This requirement creates a practical tension with the present construction method. A complete Froidure–Pin computation for either internal presentation stores a representative for every enumerated element of the base group, here G_{SL} of order 40,353,264. The full two-copy challenge contains 39,285,006 public relations. Its semidirect evaluation group has order $|G_{\mathrm{SL}}|^2 = 1{,}628{,}385{,}915{,}453{,}696$. Making either internal presented group much larger is expected to raise generation, storage, and evaluation costs. Product constructions are more economical, but when the internal solution sets are single automorphism orbits, their mixing relations expose the alignment of Proposition 1 on centerless targets or applicable centerless quotients.

A repaired proposal must therefore rule out, for its key-generation and publication distribution, efficient construction of any homomorphism into a finite effectively represented group with $\ker \psi \cap \Gamma \subseteq \Lambda$, even given polynomially many labeled ciphertexts. This may be established by a reduction or stated as a falsifiable average-case assumption. GRAFHEN presently gives neither. Doing so

while retaining practical rule generation and homomorphic evaluation is the central open problem left by the attack.

7 Conclusion

Under GRAFHEN's isomorphism hypothesis, a fixed public presentation selects one free $\mathrm{Aut}(T)$-orbit of compatible generating representations, all recognizing the same zero words. The published class count is therefore an unstructured-search baseline, not a lower bound on reconstruction from the relations. HEnbane exploits those relations with automorphism-reduced whole-element search and algebraic forcing, branching on another complete target element whenever forcing stalls. The uncapped traversal terminates with either a certified solution or exhaustive failure for every finite effectively enumerable target covered by Theorem 2. A two-seed force chain improves its instance-sensitive bound but is not required for completeness. It recovers equivalent keys and decrypts all five released symmetric challenges. In the semidirect construction, the mixed rules align the two copies rather than preserving independent ambiguity.

The exact kernel criterion extends the attack beyond equivalent keys, and Theorem 1 turns any suitable recovered map into a finite-sample CPA distinguisher without first generating the full zero subgroup. On the author-supplied G_{SL} challenge, reconstruction in G_{PSL} followed by linear central lifting recovers and verifies the exact representation without using the degree-344 action during search. Because the supplied central labels are not a valid GRAFHEN encoding, this is a representation-and-label-recovery result rather than a break of a correct revised FHE instantiation. All experimental conclusions remain fixed-instance results. A structural repair must rule out efficient construction of every finite, effectively represented homomorphic image satisfying the kernel criterion while retaining practical rule generation and evaluation. GRAFHEN currently gives no such assumption.

Acknowledgments. We thank Aymen Boudguiga for reading and providing feedback on an earlier version of this work. We also thank the GRAFHEN authors for discussions and for supplying the G_{SL} challenge analyzed here.

References

1. Armknecht, F., Katzenbeisser, S., Peter, A.: Group homomorphic encryption: characterizations, impossibility results, and applications. Designs, Codes and Cryptography **67**(2), 209–232 (2013). https://doi.org/10.1007/s10623-011-9601-2
2. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence. pp. 399–404 (2009)
3. Cameron, P.J., Solomon, R., Turull, A.: Chains of subgroups in symmetric groups. Journal of Algebra **127**(2), 340–352 (1989). https://doi.org/10.1016/0021-8693(89)90256-1

4. Dumezy, J.: Henbane. Source code and experimental artifacts (2026), https://github.com/jdumezy/henbane
5. Geraud-Stewart, R.: GRAFHEN is not IND-CPA secure. Cryptology ePrint Archive, Paper 2026/700 (2026), https://eprint.iacr.org/2026/700, version of 29 April 2026
6. Geraud-Stewart, R.: Practical equivalent-key recovery in GRAFHEN. Cryptology ePrint Archive, Paper 2026/1518 (2026), https://eprint.iacr.org/2026/1518, received 24 July 2026; approved 27 July 2026
7. Gjøsteen, K.: Symmetric subgroup membership problems. In: Public Key Cryptography (PKC 2005). Lecture Notes in Computer Science, vol. 3386, pp. 104–119. Springer (2005)
8. Guillot, P., Duc, A.H., Koskas, M., Méhats, F.: Introducing GRAFHEN: GRoup-bAsed fully homomorphic encryption without noise. Cryptology ePrint Archive, Paper 2025/1907 (2025), https://eprint.iacr.org/archive/2025/1907/20260515:083231, revision of 15 May 2026
9. Guillot, P., Duc, A.H., Koskas, M., Méhats, F.: Introducing GRAFHEN: GRoup-bAsed fully homomorphic encryption without noise. Cryptology ePrint Archive, Paper 2025/1907 (2025), https://eprint.iacr.org/archive/2025/1907/20260722:084520, revision of 22 July 2026
10. Huppert, B.: Endliche Gruppen I, Grundlehren der mathematischen Wissenschaften, vol. 134. Springer (1967)
11. Luby, M., Sinclair, A., Zuckerman, D.: Optimal speedup of Las Vegas algorithms. Information Processing Letters **47**(4), 173–180 (1993)
12. Marques-Silva, J.P., Sakallah, K.A.: GRASP: A search algorithm for propositional satisfiability. IEEE Transactions on Computers **48**(5), 506–521 (1999)
13. Seress, Á.: Permutation Group Algorithms, Cambridge Tracts in Mathematics, vol. 152. Cambridge University Press, Cambridge (2003). https://doi.org/10.1017/CBO9780511546549
14. Zhang, L., Madigan, C.F., Moskewicz, M.H., Malik, S.: Efficient conflict driven learning in a Boolean satisfiability solver. In: Proceedings of the 2001 IEEE/ACM International Conference on Computer-Aided Design. pp. 279–285. ICCAD '01, IEEE Press (2001). https://doi.org/10.1109/ICCAD.2001.968634

A Historical note on the earlier **HEnbane** solver

The 22 July 2026 revision of GRAFHEN cites the earlier version of this work and analyzes its point-table search through the permutation degree, giving an $O(2^{dn})$ worst-case estimate and hence $O(2^{1,720})$ at $d = 5$ and degree 344 [9]. That estimate concerns the prototype described here, not the whole-element backends used for the present experiments.

The earlier symmetric-group solver combined bounded closure with direct cancellation mining of public rules having a common side. It represented each generator by a partial permutation table and propagated relations from both ends of their point walks, thereby reconstructing a permutation action one image at a time. Conflicts between point images produced learned nogoods by first-UIP analysis. The solver used watched literals, nonchronological backtracking, Luby restarts, and limited exchange of low-LBD clauses between workers [12,14,11,2].

That solver recovered the released symmetric challenges, but its runtime was randomized and strongly instance-dependent. The current algorithm instead enumerates complete permutations modulo the exact automorphism action, forces generators algebraically from short relations whenever possible, and branches on complete permutations otherwise. It uses neither point-image clauses nor restarts, and its projective backend works directly with compact matrices rather than a degree-344 point table. For comparison, the earlier table-state proposition used

$$P_n^d, \qquad P_n = \sum_{k=0}^{n} \binom{n}{k}^2 k!,$$

as its partial-state envelope, in addition to complete-tuple certification. The present first-orbit leaf envelope for $n \neq 6$ is $p(n)(n!)^{d-1}$ from Equation (6). Its exact geometric state sum is less than twice this quantity for $n \geq 2$. Since $P_n \geq n!$, the ratio between the old partial-state envelope and the new exact state envelope is therefore at least

$$\frac{n!}{2p(n)} = 2^{n \log_2 n - O(n)}.$$

For the supported degrees 7 and 11 this factor is already greater than one, so the new proved state envelope is strictly smaller. Comparing the old partial-state envelope with only the leading term of the new state sum removes the factor two. This compares two proved upper envelopes. It is neither a lower-bound separation between the algorithms nor a runtime prediction for generated instances. The earlier solver remains archived at commit `06a2993` of the artifact repository [4].