CONJURA · OPEN PROBLEM

The Optimal Bound for Non-Adaptive DDH with Preprocessing

The square-root term in the proved bound should not be there — but only for DDH, and not for its square variant

Status. Statement: AI-written from Itai Dinur, Nathan Keller and Avichai Marmor, *Non-Adaptive Cryptanalytic Time-Space Lower Bounds via a Shearer-like Inequality for Permutations*, IACR ePrint 2025/783; STOC 2026. Not yet checked by a human. Open, and asserted rather than asked: the source proves an upper bound on the success probability, states in two places that it conjectures the bound is not tight, and names the value it believes is right. Its companion bound for square-DDH, proved by the same theorem, *is* sharp, which is what makes the DDH case a real question rather than a suspicion.

Category. *Idealized Models, Public Key.*

Conjecture (informal). *Preprocessing helps an attacker distinguish Diffie–Hellman keys from random. How much depends on whether the attacker's queries may react to each other. The source paper proves that a non-adaptive attacker with S bits of advice and T queries does no better than $\frac{1}{2} + \widetilde{O}(T^2/N + \sqrt{ST/N})$, and shows that for the square variant of the problem this is exactly right — there is a matching attack. For plain DDH it believes its own square-root term is an artefact of combining two bounds, and that the truth is $\frac{1}{2} + \widetilde{O}(T^2/N + ST/N)$. That is the conjecture. It is the statement that DDH under preprocessing behaves like discrete logarithm does once adaptivity is removed, and that the resemblance to the square variant stops here.*

1 The setting

Three problems that usually move together. In Shoup's generic group model [Sho97] a group of prime order N is presented through a random encoding σ, and an algorithm may only ask for encodings of linear (or, for the decisional problems, low degree) combinations

of what it has been given. Without preprocessing, all three of DLOG, DDH and square-DDH have the same T^2/N bound [Sho97, CDG18]. With preprocessing they also moved together for years: Corrigan-Gibbs and Kogan [CK18] proved an upper bound of $\frac{1}{2}+\widetilde{O}(\sqrt{ST^2/N})$ for both decisional problems, and that bound is attained by their algorithm for square-DDH but not for DDH. The gap for DDH closed only recently, when Akshima, Besselman, Guo, Xie and Ye [ABGXY24] proved the sharp $\frac{1}{2}+\widetilde{\Theta}(ST^2/N)$ for *adaptive* algorithms.

What the source proves, and where it stops. The source paper asks what happens without adaptivity. Its notion is the natural one for these problems: the *pre-translated* query is a tuple of coefficients chosen in advance, while the oracle query it becomes still depends on the secret. For DDH the coefficients are (a_1, a_2, a_3, b) and the query is $\sigma(a_1d_1 + a_2d_2 + a_3d_3 + b)$ in the random case and $\sigma(a_1d_1 + a_2d_2 + a_3d_1d_2 + b)$ in the real one, so a non-adaptive algorithm fixes a set of low-degree polynomials and asks for their values. Theorem 1 is what the source's machinery gives.

Theorem 1 (The source's Theorem 1.2). *Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be a non-adaptive (S, T)-algorithm for DDH (respectively square-DDH) in the generic group model over a group of prime order N. Then the success probability of $\mathcal{A}$ is at most*

$$\frac{1}{2} + \frac{2T^2}{N} + \sqrt{\frac{2\log_e(2)\,ST}{N}}.$$

The square-root term is what this page is about. The source is explicit that it is not an artefact for the square variant: "For the sqDDH problem, the theorem is sharp, as a simple non-adaptive variant of the adaptive algorithm of [CK18] sketched in Appendix B attains its success probability bound." For DDH no such attack is known, and the source says why it thinks none exists — the bound "is obtained by combining two different bounds", a combination that is lossy for DDH but tight for square-DDH.

The conjecture as the source states it. Twice. In its results section: "For the DDH problem, we conjecture that the bound on the success probability is not sharp, and the 'right' bound is" the

value in Conjecture 1. And in its open-problems paragraph: "Another open problem, mentioned above, is to close the gap between upper and lower bounds for non-adaptive (S,T)-algorithms for the DDH problem. As written, we conjecture that the bound of Theorem 1.2 is not tight, and the optimal bound is" — the same value. It also names the route it expects to work: "Possibly, the techniques used in the recent result [ABG+24] which determined the maximal success rate of adaptive algorithms can be combined with our techniques to show this improved bound in the non-adaptive setting."

What this page fixes. The source says "the optimal bound is $\frac{1}{2} + \widetilde{O}(T^2/N + ST/N)$", which asserts both directions: that no non-adaptive algorithm does better, and that this is attained. Conjecture 1 states the upper bound only, because that is the direction the source's own sentence about closing "the gap between upper and lower bounds" identifies as missing — a matching non-adaptive algorithm in this regime is a simple variant of the ones the source already cites, and Remark 2 records that reading rather than folding it into the statement. A reviewer who reads the source's "optimal" as asserting both halves in one breath would state a two-sided conjecture, and would be making a defensible choice this page does not make.

2 Notation and parameters

- N is prime and $\sigma : \mathbb{Z}_N \to [N]$ a uniformly random bijection; the group element g^j is presented as $\sigma(j)$.
- The DDH secret is (d_1, d_2, d_3, k) with $d_1, d_2, d_3 \leftarrow \mathbb{Z}_N$ and $k \leftarrow \{0,1\}$. The algorithm must output k.
- An outer query is $(a_1, a_2, a_3, b) \in \mathbb{Z}_N^4$, excluding the case $a_1 = a_2 = a_3 = N$, and is answered with σ applied to $a_1 d_1 + a_2 d_2 + a_3 d_3 + b \bmod N$ if $k = 0$, and to $a_1 d_1 + a_2 d_2 + a_3 (d_1 d_2) + b \bmod N$ if $k = 1$. So $k = 0$ presents $\sigma(d_1), \sigma(d_2), \sigma(d_3)$ and $k = 1$ presents $\sigma(d_1), \sigma(d_2), \sigma(d_1 d_2)$, which is DDH.
- S bounds the advice in bits and T the number of queries; $\widetilde{O}$ hides factors polylogarithmic in N.

Definition 1 (Non-adaptive (S,T)-algorithm). A pair $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ where $\mathcal{A}_0$ sees σ and outputs advice z of at most S bits, and $\mathcal{A}_1$ receives z and makes at most T queries, all of whose pre-translated forms are determined by z alone — in the source's words, "$\mathcal{A}_1$ is non-adaptive if given any z, its queries are fixed and do not (further) depend on σ". The success probability is taken over uniform σ and a uniform secret.

3 The conjecture

Conjecture 1 (Optimal non-adaptive DDH bound). *There is a function $c(\cdot)$, polylogarithmic in N, such that for every prime N and every non-adaptive (S,T)-algorithm $\mathcal{A}$ for the DDH problem in the generic group model over a group of order N, the success probability of $\mathcal{A}$ is at most*

$$\frac{1}{2} + c(N) \cdot \left(\frac{T^2}{N} + \frac{ST}{N} \right).$$

Remark 1 (The strength, in one line). Conjecture 1 is strictly stronger than Theorem 1 exactly when ST/N is small, which is the regime of interest: there $ST/N \ll \sqrt{ST/N}$, so replacing the square root by the linear term is a real improvement. When $ST \geq N$ both bounds are vacuous.

Remark 2 (The lower bound half, kept out of the statement). The source calls its conjectured value "the optimal bound", which also claims a matching algorithm. That half is not in Conjecture 1, for the reason given in Section 1: the source frames the missing work as closing a gap from above, and its own Appendix B already sketches how a non-adaptive variant of the [CK18] algorithm attains the square-DDH bound. Anyone settling the conjecture should check whether the corresponding DDH attack achieves ST/N, since if it does not the two-sided reading of the source's sentence is false even if Conjecture 1 is true.

Remark 3 (Why square-DDH is the interesting control). Both bounds come out of the same theorem, and one of them is sharp. In square-DDH the secret is (d_1, d_2, k) and the real case presents $\sigma(d_1), \sigma(d_1^2)$, so a query is a polynomial of degree at most 2 in a

single unknown; in DDH it is a polynomial in three unknowns, of degree at most 2, with the quadratic term restricted to $d_1 d_2$. The source's proof treats both through the same $(N/2)$-uniformity of the translation function, so whatever distinguishes them is invisible to it. That is a precise statement of what a solution has to find: the feature of the DDH query structure that the uniformity parameter does not see.

4 Bibliography

[ABGXY24] Akshima, Tyler Besselman, Siyao Guo, Zhiye Xie and Yuping Ye. *Tight time-space tradeoffs for the decisional Diffie–Hellman problem.* STOC 2024, pp. 1739–1749. ACM, 2024. The sharp $\frac{1}{2} + \widetilde{\Theta}(ST^2/N)$ for adaptive algorithms, and the techniques the source expects to be combinable with its own.

[CDG18] Sandro Coretti, Yevgeniy Dodis and Siyao Guo. *Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models.* CRYPTO 2018, LNCS 10991, pp. 693–721. Springer, 2018. Cited by the source for the $\frac{1}{2} + T^2/N$ bound for square-DDH without preprocessing.

[CK18] Henry Corrigan-Gibbs and Dmitry Kogan. *The discrete-logarithm problem with preprocessing.* EUROCRYPT 2018, LNCS 10821, pp. 415–447. Springer, 2018. The earlier $\frac{1}{2} + \widetilde{O}(\sqrt{ST^2/N})$ upper bound for both decisional problems, and the algorithm whose non-adaptive variant attains the square-DDH bound.

[DKM25] Itai Dinur, Nathan Keller and Avichai Marmor. *Non-adaptive cryptanalytic time-space lower bounds via a Shearer-like inequality for permutations.* IACR ePrint 2025/783; STOC 2026. The source. Theorem 1.2 and the first statement of the conjecture are on printed p. 5 (PDF p. 6); the open-problems restatement is on printed p. 10 (PDF p. 11); the DDH permutation-challenge game is in Section 3.3, printed p. 17 (PDF p. 18); the proof is Section 5.2.

[Sho97] Victor Shoup. *Lower bounds for discrete logarithms and related problems.* EUROCRYPT 1997, LNCS 1233, pp. 256–266. Springer, 1997. The generic group model and the T^2/N bounds without preprocessing.