AICR · OPEN PROBLEM

Discrete Logarithm with Preprocessing and r Rounds of Adaptivity

One formula that should interpolate between the baby-step giant-step bound at one round and the Corrigan-Gibbs–Kogan bound at T rounds

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: the source states it as a conjecture in its open-problems paragraph, says it would be sharp, and proves only the $r = 1$ case. The definition of "r rounds of adaptivity" is supplied by this page, since the source uses the phrase without defining it.

Category. *Idealized Models, Public Key.*

Conjecture (informal). *An attacker who is allowed unlimited precomputation on a group, and then S bits of notes and T online queries, can compute discrete logarithms far better than one who cannot precompute — but only if its online queries may depend on each other. That is the source paper's result: with all queries fixed in advance, preprocessing buys nothing beyond the classical baby-step giant-step bound, while with fully adaptive queries the known algorithms achieve $ST^2 \approx N$. This conjecture is the formula in between. Allow the attacker r rounds of adaptivity — r batches of queries, each batch chosen after seeing the previous answers — and its success probability should be $\widetilde{O}(T^2/N + rST/N)$: one round recovers the paper's own theorem, T rounds recovers the known adaptive bound, and every r in between is matched by an attack that builds T/r chains of length r instead of one chain of length T.*

1 The setting

Preprocessing, and why adaptivity is the question. In the preprocessing (auxiliary-input) model an algorithm is a pair $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$: an unbounded $\mathcal{A}_0$ sees the whole primitive and writes down S bits of advice, and an online $\mathcal{A}_1$ receives the advice and the

challenge and makes T queries. This is called an (S,T)-algorithm, and it is the right model for an attacker willing to amortize a one-time computation over many instances. For generic discrete logarithm the model has a known answer in the *adaptive* case: the algorithms of Mihalcik [Mih10], Bernstein and Lange [BL13] and Corrigan-Gibbs and Kogan [CK18] achieve $ST^2 \le \widetilde{O}(N)$, and [CK18] proves that this is best possible.

What was missing is the role of adaptivity. Almost every good cryptanalytic tradeoff algorithm is adaptive — Pollard's Rho for DLOG being the canonical example — and, as the source paper observes, there was no formal statement justifying that. The source supplies one: in the generic group model with preprocessing, non-adaptive algorithms for DLOG cannot beat baby-step giant-step at all.

The model, and the paper's theorem. Section 2 fixes the generic group model with preprocessing. The source's notion of non-adaptivity is the natural one for this problem and is worth stating carefully: a DLOG query is specified by a pair (a,b) and is sent to the oracle as $\sigma(a \cdot d + b \bmod N)$, so the *pre-translated* query (a,b) may be fixed in advance while the actual oracle query still depends on the secret d. The source's definition is that "$\mathcal{A}_1$ is non-adaptive if given any z, its queries are fixed and do not (further) depend on σ". This is what makes baby-step giant-step non-adaptive — its pairs are $(0,j)$ and $(1,-im)$, chosen before the challenge — and it is why the notion is not vacuous, unlike in function inversion, where challenge-independent queries trivialize the problem.

Theorem 1 (The source's Theorem 1.1, the $r = 1$ case). *Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be a non-adaptive (S,T)-algorithm for DLOG in the generic group model over a group of prime order N. Then the success probability of $\mathcal{A}$ is at most*

$$\frac{3T^2}{N} + \frac{4\log_e(2)\,ST}{N}.$$

So for $S \le \widetilde{O}(\sqrt{N})$ the online time cannot drop below $\widetilde{O}(\sqrt{N})$, however much preprocessing is done: preprocessing is worthless without adaptivity, and worth a cube-root speedup with it.

The question, and what it would settle. The source's closing paragraph: "More generally, a natural goal to pursue, in view of our results, is to obtain time-space lower bounds for adaptive attacks with preprocessing, as a function of the number of adaptivity rounds for the problems considered in this paper." It then states the conjecture recorded here, and adds why it believes the shape: "This matches our result for non-adaptive algorithms (i.e., 1 round of adaptivity), as well as the bounds of [CK18] for adaptive algorithms (i.e., T rounds of adaptivity)", and "it would be sharp, as for any $1 \le r \le T$, it is matched by a variant of the adaptive algorithm of [BL13, CK18, Mih10], in which instead of constructing one chain of length T one constructs multiple chains of length r".

What this page supplies and the source does not. The source uses "r rounds of adaptivity" without defining it; the phrase occurs once, in the conjecture. Definition 2 below is therefore this page's, and it is the first thing a reviewer should challenge. Two constraints pin it down almost completely: the source says $r = 1$ must be its own non-adaptive notion and $r = T$ must be the unrestricted adaptive one, and Definition 2 is the batching notion that does both. Note in particular that the pre-translated queries within a batch must be chosen together, and that they may depend on the advice, on the challenge's *encoding* as received, and on all answers from earlier batches.

2 Notation and parameters

- N is a prime, G a group of order N with generator g, and $\sigma : \mathbb{Z}_N \to [N]$ a uniformly random bijection that encodes g^j as $\sigma(j)$. This is Shoup's generic group model [Sho97].
- The secret is $d \leftarrow \mathbb{Z}_N$; the algorithm receives $\sigma(1)$ and $\sigma(d)$ and must output d.
- An *outer query* is a pair $(a, b) \in \mathbb{Z}_N^2$, sent to the oracle as $\sigma(a \cdot d + b \bmod N)$. An *inner query* is a direct query $\sigma(i)$.
- S bounds the advice length in bits, T the total number of queries the online algorithm makes. $\widetilde{O}$ hides factors polylogarithmic in N.

Definition 1 ((S,T)-algorithm with preprocessing). A pair $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ where $\mathcal{A}_0$ has direct access to σ and outputs an advice string $z = \mathcal{A}_0(\sigma)$ of at most S bits, and $\mathcal{A}_1$ receives z, the encodings $\sigma(1), \sigma(d)$, and makes at most T queries as above. Both may be assumed deterministic without loss of generality. The success probability is the probability, over uniform σ and d, that $\mathcal{A}_1$ outputs d.

Definition 2 (r rounds of adaptivity — *this page's definition*). For $1 \le r \le T$, an (S,T)-algorithm has r *rounds of adaptivity* if $\mathcal{A}_1$'s queries are partitioned into r batches $B_1, \ldots, B_r$, issued in order, such that the outer and inner queries in batch B_i are a function of the advice z, the encodings $\sigma(1), \sigma(d)$, and the answers to the queries in $B_1, \ldots, B_{i-1}$ only — and in particular do not depend on the answers to queries in B_i itself. The batch sizes may be arbitrary and need not be fixed in advance.

Remark 1 (The two endpoints are the ones the source names). At $r = 1$ every query is a function of z and the encodings alone, which is the source's non-adaptive notion, so Theorem 1 is exactly the $r = 1$ case of Conjecture 1. At $r = T$ each batch may be a single query and Definition 2 is no restriction at all, so Conjecture 1 reads $\widetilde{O}(T^2/N + ST^2/N) = \widetilde{O}(ST^2/N)$, which is the bound proved in [CK18].

3 The conjecture

Conjecture 1 (r-round DLOG tradeoff). *There is a function $c(\cdot)$, polylogarithmic in N, such that for every prime N, every $1 \le r \le T$ and every (S,T)-algorithm $\mathcal{A}$ for DLOG in the generic group model over a group of order N having r rounds of adaptivity in the sense of Definition 2, the success probability of $\mathcal{A}$ is at most*

$$c(N) \cdot \left(\frac{T^2}{N} + \frac{r\,S\,T}{N} \right).$$

Remark 2 (Why the bound is the right guess, and what makes it sharp). The upper bound is matched, for every r, by an algo-

rithm the source describes in one line: run the adaptive algorithm of [BL13, CK18, Mih10] but build T/r chains of length r rather than one chain of length T. Reading the conjecture as an equality up to polylogarithmic factors, the interesting content is the linear dependence on r: each extra round of adaptivity buys the attacker a factor in the advice term ST/N and nothing in the query term T^2/N. Settling it would replace the current binary picture — adaptive algorithms are stronger than non-adaptive ones — with a quantitative one, and would say exactly how much of Pollard-style sequentiality an attacker really needs.

Remark 3 (The obstruction the source names). The source's proof runs through a Shearer-like inequality for permutations, and the restriction to one round is where that inequality binds: "The independence of the queries from the challenge appears inherent to our proof strategy." Known Shearer-type bounds control the information a random function reveals on a set of input–output pairs that is fixed independently of the function; allowing later queries to depend on earlier answers produces composed inputs such as $f(d + f(d))$, which are not of that form. The source's assessment is that "extending our approach beyond this restriction would require new Shearer-type inequalities or alternative techniques that can handle such challenge-dependent queries", and it notes that the one inequality in the literature allowing some dependence [RS18] is too weak here.

Remark 4 (Where a refutation would come from). An attack, and the natural place to look is the term the conjecture claims is independent of r. If some algorithm with a constant number of rounds achieved, say, $\widetilde{O}(ST^2/N)$ for S in a middle range, the conjecture is false and the linear-in-r shape is wrong. Nothing rules that out: the only proved points are $r = 1$ and $r = T$.

4 Bibliography

[BL13] Daniel J. Bernstein and Tanja Lange. *Non-uniform cracks in the concrete: the power of free precomputation.* ASIACRYPT 2013, LNCS 8270, pp. 321–340. Springer, 2013.

[CK18] Henry Corrigan-Gibbs and Dmitry Kogan. *The discrete-logarithm problem with preprocessing.* EUROCRYPT 2018, LNCS 10821, pp. 415–

447. Springer, 2018. The adaptive tradeoff $ST^2 = \widetilde{O}(N)$, proved optimal, which the conjecture must reproduce at $r = T$.

[CK19] Henry Corrigan-Gibbs and Dmitry Kogan. *The function-inversion problem: barriers and opportunities.* TCC 2019, LNCS 11891, pp. 393–421. Springer, 2019. Where the question of the power of adaptivity was first raised, for function inversion.

[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. The conjecture is in the "Open problems" paragraph on printed p. 10 (PDF p. 11); Theorem 1.1 is on printed p. 4 (PDF p. 6) and reproved in Section 5.1; the model and the non-adaptivity definition are in Section 3 on printed pp. 14–15 (PDF pp. 16–17).

[Mih10] Joseph P. Mihalcik. *An analysis of algorithms for solving discrete logarithms in fixed groups.* Master's thesis, Naval Postgraduate School, 2010.

[RS18] Anup Rao and Makrand Sinha. *Simplified separation of information and communication.* Theory of Computing, 14(1):1–29, 2018. The Shearer-type inequality allowing some dependence that the source cites as too weak for its setting.

[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 bound without preprocessing.