CONJURA · OPEN PROBLEM

A Quadratic Attack on Three-Party NIKE in Shoup's Generic Group Model

Generic groups with explicit labels, query complexity only

Status. Statement: AI-written from Abtin Afshar, Geoffroy Couteau, Mohammad Mahmoody and Elahe Sadeghi, *Fine-Grained Non-Interactive Key-Exchange: Constructions and Lower Bounds*, IACR Cryptology ePrint Archive 2023, not yet checked by a human. The corresponding statement in Maurer's generic group model is settled by the paper itself, for every $K \geq 3$ parties and including imperfect correctness, while the statement below — the same attack in Shoup's generic group model, where group elements carry explicit labels — is open, and the paper's linear-structure technique gives nothing there.

Category. *Idealized Models & Non-Uniformity.*

Conjecture (informal). *In a generic group model, parties can do group operations but nothing else clever with group elements. Merkle's puzzles show that two parties who each do n units of work can agree on a key that costs an eavesdropper about n^2 work to recover, and a classical theorem says n^2 is exactly the right answer when the only idealized object is a hash function. The source paper asks the same question when the idealized object is a group. It proves the n^2 attack for three parties in Maurer's version of the generic group model, where a group element lives inside a private array and a party can only add elements and test whether a combination is zero. But its own positive result — a four-party protocol with an n^2 security gap — lives in Shoup's version, where every group element has an explicit bit-string label that parties may hash, sort, or run a discrete-logarithm algorithm on. The conjecture is that the attack survives the move to the labelled model: for every three-party non-interactive key exchange in Shoup's generic group model in which each honest party makes at most n oracle queries, there is an eavesdropper making $O(n^2)$ queries that recovers the shared key with high probability. Settling it would say that generic algebraic structure, without pairings, buys no more than a quadratic gap for three parties — still strictly more than the $n^{1.5}$ the paper's random-oracle 3-NIKE achieves, and matching the quadratic bound Barak–Mahmoody proved optimal for two-party key agreement from random oracles.*

1 The setting

Non-interactive key exchange for three or more parties is one of the few elementary primitives with no construction from standard assumptions outside of pairings [Jou00] and obfuscation [BZ14]; two-party NIKE goes back to Diffie and Hellman [DH76]. Fine-grained cryptography relaxes the question: rather than security against all polynomial-time adversaries, one asks for a fixed polynomial gap, where honest parties run in time n and security holds against adversaries running in time n^c for some $c > 1$. Merkle's puzzles [Mer78] give a two-party protocol of this kind in the random oracle model with c approaching 2, and Barak and Mahmoody [BM09, BM17] proved that $c = 2$ is optimal there: every n-query key agreement in the random oracle model is broken by an $O(n^2)$-query attacker.

The source paper asks what happens when the idealized object carries algebraic structure instead. On the positive side it constructs a four-party NIKE in Shoup's generic group model [Sho97] with a quadratic gap, and on the negative side it gives an $O(n^2)$-query attack on every n-query three-party NIKE in Maurer's generic group model [Mau05] (its proof is written for three parties and, by the paper's own footnote, implies the same for every $K \geq 3$). The two models differ in what a group element *is*. In Shoup's model, group elements are explicit labels that a party may hash, sort, or feed to a generic discrete-logarithm algorithm. In Maurer's model, group elements live in a private array and are reachable only through addition and zero-test queries; a party never sees a representation. The attack exploits exactly this: in Maurer's model, every element a party writes is a fixed linear function of the elements it was handed at the start [FKL17], and that structural handle is what drives the randomness-switching lemma and, after zero tests are compiled away, the whole attack.

No such handle exists in Shoup's model, and the paper is careful about what its Maurer-model theorem therefore establishes: negative results in the MGGM "are generally weaker than those in the SGGM and should be interpreted cautiously" [Zha22, DHH21] (p. 3). Proving the same quadratic attack in Shoup's model would be the generic-group analogue of the Barak–Mahmoody theorem: it would say that generic algebraic structure, without pairings, buys no more than a quadratic gap for three parties — still strictly more than the $n^{1.5}$ gap the paper's own random-oracle 3-NIKE achieves, and matching the quadratic bound Barak–Mahmoody proved optimal for two-party key agreement from random oracles. Together with the paper's four-party construction, it would pin the maximum achievable gap in Shoup's model at quadratic across $K \geq 3$ collectively — achieved at $K = 4$ — while the best gap achievable for three parties would remain open between $n^{1.5}$ and n^2.

2 Notation and parameters

Throughout, λ is the security parameter, $p = p(\lambda)$ is a prime, and $\mathbb{Z}_p$ is the additive group of integers modulo p. We write $[k] = \{1, \ldots, k\}$. A function $\mathrm{negl}(\lambda)$ is negligible if it is $\lambda^{-\omega(1)}$. The quantity $n = n(\lambda)$ is the query budget of each honest party, and is the cost measure

throughout: as in the source setting, in an idealized model the number of oracle queries stands in for running time, so oracle queries are counted and local computation is left unmetered. Leaving local computation unmetered is a modelling choice of this record rather than the source's own security definition, which asks all algorithms to be efficient and to submit at most $\mathrm{poly}(\lambda)$ queries; it is, however, what the attacker of the source's Maurer-model theorem needs, since that attacker rejection-samples random seeds compatible with the transcript. The value $\delta = \delta(\lambda) \in [0,1]$ is the completeness error of the protocol and $\delta' \in (0,1)$ is a constant appearing in the attacker's success probability. The label space is S, the encoding map is enc, the oracle is O_{enc}, the parties' randomness is $\mathsf{r}_1, \mathsf{r}_2, \mathsf{r}_3$, their broadcast messages are $\mathsf{m}_1, \mathsf{m}_2, \mathsf{m}_3$ with $\mathsf{tran} = (\mathsf{m}_1, \mathsf{m}_2, \mathsf{m}_3)$, and their outputs are $\mathsf{key}_1, \mathsf{key}_2, \mathsf{key}_3$. The eavesdropper is E and the protocol is $\Pi = (\mathsf{Msg}, \mathsf{Key})$.

The parameters of the statement are:

- λ, the security parameter.
- n, the query budget of each honest party; the fine-grained cost measure.
- p, the prime order of the generic group.
- δ, the completeness error of the protocol, i.e. the probability that the three parties fail to agree.
- δ', the slack constant in the attacker's success probability; it determines the constant hidden in $O(n^2)$.

3 The conjecture

Definition 1 (Shoup's generic group model). Let p be a prime and let S be a label space with $|S| \geq p$. Sample a uniformly random injective map $\mathsf{enc} : \mathbb{Z}_p \to S$. Every algorithm — honest party and adversary alike — is given access to one oracle O_{enc} supporting two kinds of query:

- $\mathsf{enc}(x)$, for $x \in \mathbb{Z}_p$, which returns the label $\mathsf{enc}(x) \in S$;
- $\mathsf{Add}(c_1, c_2, \sigma_1, \sigma_2)$, for $c_1, c_2 \in \mathbb{Z}_p$ and $\sigma_1, \sigma_2 \in S$, which

returns $\mathsf{enc}(c_1x_1 + c_2x_2)$ when $\sigma_1 = \mathsf{enc}(x_1)$ and $\sigma_2 = \mathsf{enc}(x_2)$ for some $x_1, x_2 \in \mathbb{Z}_p$, and $\perp$ otherwise.

Labels are explicit bit strings: an algorithm may perform arbitrary (unbounded) local computation on the labels it holds. The cost of an algorithm is its number of oracle queries; local computation is free.

Definition 2 (Three-party NIKE in an idealized model). A three-party non-interactive key-exchange protocol relative to an oracle O is a pair of oracle algorithms $\Pi = (\mathsf{Msg}, \mathsf{Key})$. Each party $i \in [3]$ samples private randomness r_i and broadcasts a single message $\mathsf{m}_i \leftarrow \mathsf{Msg}^O(1^\lambda, i, \mathsf{r}_i)$; all three messages are sent simultaneously, and $\mathsf{tran} := (\mathsf{m}_1, \mathsf{m}_2, \mathsf{m}_3)$. Without any further interaction, party i outputs $\mathsf{key}_i \leftarrow \mathsf{Key}^O(i, \mathsf{r}_i, \mathsf{tran})$. The protocol is *$n$-query* if, for each i, the total number of oracle queries made by $\mathsf{Msg}^O(1^\lambda, i, \cdot)$ and $\mathsf{Key}^O(i, \cdot, \cdot)$ together is at most n. It has *completeness error* δ if

$$\Pr[\mathsf{key}_1 = \mathsf{key}_2 = \mathsf{key}_3] \;\geq\; 1 - \delta,$$

the probability being over $\mathsf{r}_1, \mathsf{r}_2, \mathsf{r}_3$, the oracle, and the algorithms' coins.

Conjecture 1 (Quadratic attack on three-party NIKE in Shoup's model). *For every constant $\delta' \in (0,1)$ there is a constant $c = c(\delta') > 0$ such that the following holds. Let $n = n(\lambda)$ and let $p = p(\lambda)$ be a prime, and let $\Pi = (\mathsf{Msg}, \mathsf{Key})$ be any n-query three-party NIKE protocol in Shoup's generic group model over $\mathbb{Z}_p$ (Definitions 1 and 2), with completeness error $\delta = \delta(\lambda)$. Then there exists a computationally unbounded eavesdropper E which, on input 1^λ and the transcript tran, and with access to the same oracle O_{enc}, makes at most $c \cdot n(\lambda)^2$ oracle queries and satisfies*

$$\Pr\left[\, \mathsf{E}^{O_{\mathsf{enc}}}(1^\lambda, \mathsf{tran}) = \mathsf{key}_1 \,\right] \;\geq\; 1 - 4\delta(\lambda) - \delta'$$

for all sufficiently large λ, where the probability is over $\mathsf{r}_1, \mathsf{r}_2, \mathsf{r}_3$, the choice of enc, and the coins of E.

Remark 1 (What is settled). The paper proves exactly this statement in Maurer's generic group model. The proof runs in two stages. For protocols that ask no zero tests, a randomness-switching lemma (Lemma 19) shows that the first party's key is recoverable from three keys computed under resampled randomness, which gives an attack with $O(\alpha)$ Add queries, in the paper's own notation (Theorem 20). For general protocols, a learning phase discovers all ε-heavy pure linear constraints on the broadcast group elements and compiles the zero tests out of the protocol, reducing to the previous case (Construction 23, Theorem 21); the corollary is that n-query parties agreeing with probability at least 0.99 are broken with probability at least 0.95 by $O(n^2)$ queries. Both stages rest on Lemma 17, that in Maurer's model every element a party writes is a fixed linear function of the elements copied into its array at the start. Nothing in the paper carries this to Shoup's model. The paper also records that before this work, no polynomial-query attack on three-party NIKE in Maurer's model was known at all: "Prior to our work, it was open to break 3-NIKE protocols in Maurer's model with any polynomial number of queries" (p. 1).

Remark 2 (Openness). The paper states the gap and names closing it as a question it finds interesting: "In our third contribution, we prove our lower bound in Maurer's generic group model, whereas our positive result holds in Shoup's generic group model, which is more flexible: this leaves a gap between our positive and negative results" (p. 3); "We view as an interesting question the goal of closing the gap between our positive and negative results, either by building a 4-NIKE protocol with quadratic security in Maurer's generic group model, or by extending our impossibility result to Shoup's generic group model" (p. 3). Conjecture 1 is the second of those two routes. What is proved is the Maurer-model form: "In particular, we prove that in Maurer's generic group model (MGGM), where the access to the group is further limited through an oracle who does all the calculations, one cannot achieve more than quadratic gap between the honest parties and the adversary's query complexity" (p. 6). The paper describes its result as "a natural first step towards proving a stronger negative result for a basic question of whether 3-NIKE can be based merely on simple algebraic assumptions without pairing" (p. 3).

Remark 3 (Why it would matter). The statement is the generic-group analogue of the Barak–Mahmoody theorem, and it answers the paper's motivating question in the direction that matters. If true, no amount of pairing-free generic algebraic structure gets a three-party NIKE past a quadratic gap, so the search for fine-grained multiparty NIKE beyond quadratic must leave the generic group model entirely. If false, the refutation is a protocol that beats quadratic using labels in an essential way, which would be a genuinely new fine-grained construction. The statement itself needs only the two definitions above, both of which fit in a paragraph; the cost model is a query count, with no computational assumptions anywhere, so the claim is unconditional and information-theoretic, and both the target model and the target bound are already fixed by the paper's own theorem in the weaker model.

Remark 4 (Caveats for a reviewer). Two things deserve the hardest look. First, resolution: the source is a 2023 ePrint in an active line — fine-grained NIKE and generic-group lower bounds — and the papers citing it have not been checked; if a follow-up has proved or refuted the Shoup-model statement, this record is void. Second, the exact quantitative form: the $1 - 4\delta - \delta'$ success bound is transported verbatim from the paper's Maurer-model Theorem 21, on the grounds that "extending our impossibility result" names exactly that theorem. A reviewer might reasonably prefer the cruder form the abstract uses — any n-query 3-NIKE is broken by an $O(n^2)$-query attacker — which is the same claim up to constants for constant completeness error. Note finally that the honest parties here are computationally unbounded and only their queries are counted; this follows the paper's convention that in an idealized model query count stands in for running time (Definition 2, p. 8: "When the protocol is in an idealized model, we use the number of queries by the algorithms to the oracle as the measure of their running time"), but the paper's own security definition (Definition 14, p. 15) asks all algorithms to be efficient and to submit at most $\mathrm{poly}(\lambda)$ queries, so the unmetered-computation reading is this record's, and it makes the statement a strong one: in Shoup's model, unbounded local computation on labels is a real power.

4 Bibliography

[BM09] B. Barak and M. Mahmoody-Ghidary. *Merkle puzzles are optimal — an $O(n^2)$-query attack on any key exchange from a random oracle.* CRYPTO 2009, LNCS 5677, pages 374–390, Springer, 2009.

[BM17] B. Barak and M. Mahmoody-Ghidary. *Merkle's key agreement protocol is optimal: An $O(n^2)$ attack on any key agreement from random oracles.* Journal of Cryptology, 30(3):699–734, 2017.

[BZ14] D. Boneh and M. Zhandry. *Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation.* CRYPTO 2014, Part I, LNCS 8616, pages 480–499, Springer, 2014.

[DH76] W. Diffie and M. E. Hellman. *New directions in cryptography.* IEEE Transactions on Information Theory, 22(6):644–654, 1976.

[DHH21] N. Dottling, D. Hartmann, D. Hofheinz, E. Kiltz, S. Schage, and B. Ursu. *On the impossibility of purely algebraic signatures.* TCC 2021, Part III, LNCS 13044, pages 317–349, Springer, 2021.

[FKL17] G. Fuchsbauer, E. Kiltz, and J. Loss. *The algebraic group model and its applications.* Cryptology ePrint Archive, Paper 2017/620, 2017.

[Joo00] A. Joux. *A one round protocol for tripartite Diffie–Hellman.* International Algorithmic Number Theory Symposium, pages 385–393, Springer, 2000.

[Mau05] U. M. Maurer. *Abstract models of computation in cryptography (invited paper).* 10th IMA International Conference on Cryptography and Coding, LNCS 3796, pages 1–12, Springer, 2005.

[Mer78] R. C. Merkle. *Secure communications over insecure channels.* Communications of the ACM, 21(4):294–299, 1978.

[Sho97] V. Shoup. *Lower bounds for discrete logarithms and related problems.* EUROCRYPT'97, LNCS 1233, pages 256–266, Springer, 1997.

[Zha22] M. Zhandry. *To label, or not to label (in generic groups).* CRYPTO'22, 2022.