AICR · OPEN PROBLEM

Round-Optimal Blind Signatures in the Generic Group and Random Oracle Model

Beyond the logarithmic bound on user-side and verification hash queries

Status. Statement: AI-written from Marian Dietz, Julia Kastner, Stefano Tessaro, *On the Impossibility of Round-Optimal Pairing-Free Blind Signatures in the ROM*, Cryptology ePrint Archive 2026, not yet checked by a human. Proved for $q \leq O(\log \lambda)$ (Theorem 3.1 of the source), for random oracles with bitstring output (Definition 2.2); Section 3.10 sketches, but does not carry out in full, the extension to random oracles whose outputs also contain group elements. Open for q superlogarithmic and polynomially bounded; the superpolynomial message space hypothesis is retained, not relaxed.

Category. *Idealized Models & Non-Uniformity.*

Conjecture (informal). *A blind signature lets a user obtain a signer's signature on a message without the signer learning the message or being able to link the resulting signature to the session. Every known scheme built purely from ordinary elliptic curves (no pairings) needs at least three messages of interaction, which forces the signer to keep state; round-optimal schemes — one message each way — are only known from RSA, lattices, pairings, or from generic transformations that compile a signature verification circuit into a zero-knowledge proof, which is expensive. The source paper asks whether this three-move barrier is real for schemes that touch the group only through generic operations (add two elements, test two elements for equality) and treat the hash function as a random oracle. It proves that in this model no round-optimal scheme can be simultaneously correct, blind, and one-more unforgeable, provided the message space is superpolynomially large and the user's second algorithm together with the verification algorithm make at most logarithmically many random oracle queries. The logarithmic restriction is an artifact of the attack: the attacker learns linear relations among group elements by re-running the user's finalization and verification with re-randomized hash answers, and to do that it must decide, for each hash query, whether the key generator or the signer already made that query in this session; it cannot tell, so it guesses, and all guesses are right only with probability exponentially small in the number of queries. Conjectured here is that the restriction is not real: the same impossibility should hold when the user's final step and verification together make any polynomial number of random oracle queries. Settling it in the affirmative would close the question for the entire natural class of schemes that use the group as a black box; refuting it would mean exhibiting a round-optimal pairing-free blind signature in this model, which would be a significant construction.*

1 The setting

Blind signatures [Cha82] let a user obtain a signature on a message without the signer learning the message or being able to link a signature to the session that produced it. Practical deployments overwhelmingly use RSA-based schemes, largely because they are *round optimal*: one message from the user, one from the signer, so the signer keeps no per-session state. Round-optimal schemes

are also known from pairings and from lattices, and generically from Fischlin's transformation [Fis06], which has the user prove in zero knowledge that it knows a signature on a commitment to the message — a non-black-box use of the underlying signature scheme, and expensive. What is conspicuously missing is a round-optimal scheme from *pairing-free* groups that uses the group as a black box: every such scheme known needs at least three moves.

The source paper [DKT26] asks whether this three-move barrier is inherent, in a model it formalizes as GGM+ROM: Maurer's generic group model [Mau05] — where algorithms manipulate group elements only through an addition oracle and an equality oracle, and in particular cannot exploit the bit representation of group elements — combined with the random oracle model [BR93], where the hash may take group elements as input. This model captures all existing black-box-group blind signatures and excludes Fischlin-style circuit-level proofs, while still permitting fully algebraic Σ-protocols. The starting point is the impossibility of purely algebraic signatures of Döttling et al. [DHHKSU], whose attack maintains a space of learned linear constraints on the verification key and, whenever the constraints already imply the existence of a signature, produces one without asking the signer. Extending this to arbitrary (non-purely-algebraic) verification follows a generalization in the spirit of [CFGG22], and extending it across a random oracle is the paper's main technical contribution: the attacker simulates verification's group-equality checks and hash queries for free against a guessed constraint space, and it learns new constraints by re-running the user's finalization and verification many times with re-randomized hash answers on the *same* signer message. Blindness is what makes this possible: a hash query made both by the signer and by verification would link sessions, so up to "irrelevant" queries the two query sets are disjoint.

The result proved is that no scheme in this model can be correct, blind, and one-more unforgeable, under two hypotheses: the message space is superpolynomial, and User_2 and Verify together make $O(\log \lambda)$ hash queries. The second is where the argument stalls, and the paper says why: to re-randomize only the queries that are fresh in the current session, the attacker would have to know which queries KeyGen or Sign already made, and it cannot tell, so it guesses for every query, succeeding with probability 2^{-q}.

The factor 2^q is not confined to that one step — it reappears in the number of re-randomization rounds, in the thresholds of the paper's bad events, and in its oracle-fixing lemma, which bounds by $2^q \cdot p$ the failure probability of a q-query procedure after an arbitrary set of hash outputs has been fixed. The conjecture is that all of this is an artifact: that the impossibility holds for any polynomial number of user-side and verification hash queries. Nothing in the paper suggests a scheme could exploit many hash queries; the paper's own framing is that folklore expects round-optimal pairing-free blind signatures to be impossible outright, and that both of its hypotheses hold for every existing black-box-group construction. Outside the model the picture is different and instructive: Katz, Schröder and Yerukhimovich [KSY11] rule out blind signatures from one-way permutations by intersection-query techniques that the paper reports it does not know how to use here; Kastner, Nguyen and Reichle [KNR24] give a round-optimal pairing-free scheme that additionally assumes RSA and therefore escapes the model; and concurrent work [KTZ26] instantiates Fischlin's transform over pairing-free curves, using NIZKs about the correct evaluation of a conversion function from group elements to scalars — a non-black-box technique that falls outside the model — in a scheme based on variants of the Nyberg-Rueppel signature scheme. Settling the conjecture in the affirmative would close the black-box-group question completely; refuting it would produce a round-optimal pairing-free blind signature of a shape nobody currently knows how to build.

Progress to date. The source proves the full statement for $q \leq O(\log \lambda)$: a computationally unbounded but query-efficient GGM+ROM adversary that learns linear relations on a growing pile of group elements during a learning phase and then, in a forging phase, searches over simulated signer messages for one whose simulated User_2+Verify execution accepts with probability at least $1/4$. Blindness is converted into the statement that the verifier essentially never repeats a signer hash query except on "irrelevant" queries, of which there are few and which the attacker can harvest. A technical lemma of independent interest bounds by $2^q \cdot p$ the probability that a q-query procedure with failure probability p can be made to fail after an adversarially chosen set of random oracle outputs is fixed. Section 3.10 extends everything to random

oracles whose outputs contain group elements. No progress toward superlogarithmic q is reported.

2 Notation and parameters

- λ — the security parameter. $\mathrm{negl}(\lambda)$ denotes a function that is $\lambda^{-\omega(1)}$; a quantity is *polynomial* if it is $O(\lambda^c)$ for some constant c and *superpolynomial* if it is $\lambda^{\omega(1)}$. *ppt* abbreviates probabilistic polynomial time.
- $\mathbb{G} = \mathbb{G}(\lambda)$, $p = p(\lambda)$ — a family of groups of prime order p, written additively, with neutral element $\mathfrak{o}$ and fixed generator $\mathfrak{1}$; group elements are written in fraktur. The group is accessed only generically.
- O_{eq}, O_{grp} — the generic-group equality and group-operation oracles (Definition 1).
- $\mathcal{H}$, $\overline{\mathcal{H}}$ — the random oracle, taking n^{H} group elements and a bitstring and returning an r-bit string, and the interface through which a GGM algorithm queries it.
- n^{H}, r — the number of group elements in a random oracle input and the bitlength of its output.
- $\overline{A}$, A — a GGM algorithm and its transformation into an algorithm on actual group elements; O_{Sign} (resp. $\overline{O_{\mathsf{Sign}}}$) is the signing oracle (resp. its GGM interface).
- $\mathsf{BS} = (\mathsf{KeyGen}, \mathsf{User}_1, \mathsf{Sign}, \mathsf{User}_2, \mathsf{Verify})$ — a round-optimal blind signature scheme with message space $\mathcal{M} = \mathcal{M}(\lambda)$, which is assumed of superpolynomial size.
- q_A — an upper bound on the number of $\mathcal{H}$-queries made by an algorithm A; $q := q_{\mathsf{User}_2} + q_{\mathsf{Verify}}$. The known result assumes $q \le O(\log \lambda)$; the conjecture allows any polynomial q.
- ℓ — the number of signing-oracle calls made by the one-more-unforgeability adversary.
- $\mathrm{adv}^{\mathbf{blind}}_{\mathsf{BS},\mathcal{S}}$, $\mathrm{adv}^{\mathbf{omuf}}_{\mathsf{BS},\mathcal{U}}$ — the blindness and one-more-unforgeability advantages (Definitions 3 and 4).
- $[n]$ denotes $\{1, \ldots, n\}$ and $\oplus$ denotes exclusive or.

3 The conjecture

Definition 1 (GGM+ROM). Let $\mathbb{G} = \mathbb{G}(\lambda)$ be a family of groups of prime order $p = p(\lambda)$, written additively, with neutral element $\mathtt{0}$ and a fixed generator $\mathtt{1}$. In Maurer's generic group model [Mau05] a *GGM algorithm* $\overline{A}$ never sees elements of $\mathbb{G}$: it manipulates indices (labels) into a table $\mathtt{T}$ of group elements which is initialized with $\mathtt{1}$ followed by $\overline{A}$'s group-element inputs. It has access to two group oracles: $O_{\mathsf{eq}}(i,j)$ returns 1 if $\mathtt{T}[i] = \mathtt{T}[j]$ and 0 otherwise, and $O_{\mathsf{grp}}(i,j)$ appends $\mathtt{T}[i] + \mathtt{T}[j]$ to $\mathtt{T}$ and returns its index. The output of $\overline{A}$ consists of a list of indices together with a bitstring. Every GGM algorithm $\overline{A}$ induces an ordinary algorithm A, called its *transformation*, which takes and returns actual group elements and maintains the table $\mathtt{T}$ externally.
A *random oracle for groups* with n^{H} input group elements and output length $r = r(\lambda)$ is a uniformly random function $\mathcal{H} : \mathbb{G}^{n^{\mathsf{H}}} \times \{0,1\}^* \to \{0,1\}^r$. A GGM algorithm queries it through an oracle $\overline{\mathcal{H}}$ by supplying n^{H} indices and a bitstring, which are translated into the corresponding group elements before $\mathcal{H}$ is evaluated; the answer is a bitstring. The *GGM+ROM* is the model in which all algorithms and adversaries are GGM algorithms with access to O_{eq}, O_{grp} and $\overline{\mathcal{H}}$.

Definition 2 (Round-optimal blind signature and correctness). A *round-optimal blind signature scheme* in the ROM with message space $\mathcal{M} = \mathcal{M}(\lambda)$ is a tuple $\mathsf{BS} = (\mathsf{KeyGen}, \mathsf{User}_1, \mathsf{Sign}, \mathsf{User}_2, \mathsf{Verify})$ of ppt algorithms with access to $\mathcal{H}$, where for fixed polynomials $n^{\mathsf{vk}}, n^{\mathsf{st}}, n^{\mathsf{msgU}}, n^{\mathsf{msgS}}, n^{\mathsf{sig}}$ in λ:

- $\mathsf{KeyGen}^{\mathcal{H}}(1^\lambda)$ outputs $(\mathsf{sk}, \mathsf{vk})$ with $\mathsf{sk} \in \{0,1\}^*$ and $\mathsf{vk} \in \mathbb{G}^{n^{\mathsf{vk}}} \times \{0,1\}^*$;
- $\mathsf{User}_1^{\mathcal{H}}(\mathsf{vk}, m)$, for $m \in \mathcal{M}$, outputs a user message $\mathsf{msgU} \in \mathbb{G}^{n^{\mathsf{msgU}}} \times \{0,1\}^*$ and a state $\mathsf{stU} \in \mathbb{G}^{n^{\mathsf{st}}} \times \{0,1\}^*$;
- $\mathsf{Sign}^{\mathcal{H}}(\mathsf{sk}, \mathsf{msgU})$ outputs a signer message $\mathsf{msgS} \in$

$\mathbb{G}^{n^{\mathsf{msgS}}} \times \{0,1\}^*$;

- $\mathsf{User}_2^{\mathcal{H}}(\mathsf{stU}, \mathsf{msgS})$ outputs a signature $\sigma \in \mathbb{G}^{n^{\mathsf{sig}}} \times \{0,1\}^*$;
- $\mathsf{Verify}^{\mathcal{H}}(\mathsf{vk}, m, \sigma)$ deterministically outputs a bit.

We write $\langle \mathsf{User}^{\mathcal{H}}(\mathsf{vk}, m), \mathsf{Sign}^{\mathcal{H}}(\mathsf{sk}) \rangle$ for the sequential execution $(\mathsf{msgU}, \mathsf{stU}) \leftarrow \mathsf{User}_1^{\mathcal{H}}(\mathsf{vk}, m)$, $\mathsf{msgS} \leftarrow \mathsf{Sign}^{\mathcal{H}}(\mathsf{sk}, \mathsf{msgU})$, $\sigma \leftarrow \mathsf{User}_2^{\mathcal{H}}(\mathsf{stU}, \mathsf{msgS})$, whose output is σ. For an algorithm A we write $q_A = q_A(\lambda)$ for an upper bound on the number of calls to $\mathcal{H}$ made by A.
BS is *correct* if for every $m \in \mathcal{M}(\lambda)$,

$$\Pr\Big[\, (\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}^{\mathcal{H}}(1^\lambda), \\ \sigma \leftarrow \langle \mathsf{User}^{\mathcal{H}}(\mathsf{vk}, m), \mathsf{Sign}^{\mathcal{H}}(\mathsf{sk}) \rangle : \\ \mathsf{Verify}^{\mathcal{H}}(\mathsf{vk}, m, \sigma) = 1 \,\Big] \;\geq\; 1 - \mathsf{negl}(\lambda),$$

the probability being over $\mathcal{H}$ and the coins of the algorithms.

Definition 3 (Blindness). For a scheme BS and an adversary $\mathcal{S}$, the game $\mathbf{blind}_{\mathsf{BS},\mathcal{S}}$ samples $b \leftarrow_\$ \{0,1\}$ and a random oracle $\mathcal{H}$ and runs

$$b' \leftarrow \mathcal{S}^{\mathcal{H}, O_{\mathsf{init}}, O_{\mathsf{User}_1,0}, O_{\mathsf{User}_1,1}, O_{\mathsf{User}_2,0}, O_{\mathsf{User}_2,1}}(1^\lambda),$$

where

- $O_{\mathsf{init}}(\mathsf{vk}, m_0, m_1)$ may be called at most once and stores vk, m_0, m_1;
- $O_{\mathsf{User}_1,c}()$, for $c \in \{0,1\}$, may be called at most once each and only after O_{init}, runs $(\mathsf{stU}_c, \mathsf{msgU}_c) \leftarrow \mathsf{User}_1^{\mathcal{H}}(\mathsf{vk}, m_{b\oplus c})$ and returns msgU_c;
- $O_{\mathsf{User}_2,c}(\mathsf{msgS}_c)$, for $c \in \{0,1\}$, may be called at most once each and only after $O_{\mathsf{User}_1,c}$, computes and stores $\sigma_{b\oplus c} \leftarrow \mathsf{User}_2^{\mathcal{H}}(\mathsf{stU}_c, \mathsf{msgS}_c)$; if $O_{\mathsf{User}_2,1\oplus c}$ was already called and $\mathsf{Verify}^{\mathcal{H}}(\mathsf{vk}, m_c, \sigma_c) = 1$ for both $c = 0,1$, it returns (m_0, σ_0) and (m_1, σ_1), and otherwise $\perp$.

The advantage is $\mathrm{adv}^{\mathbf{blind}}_{\mathsf{BS},\mathcal{S}} := 2 \cdot \left|\Pr[b = b'] - \frac{1}{2}\right|$.

Definition 4 (One-more unforgeability). For a scheme BS and an adversary $\mathcal{U}$, the game $\mathbf{omuf}_{\mathsf{BS},\mathcal{U}}$ samples a random oracle $\mathcal{H}$ and $(\mathsf{sk},\mathsf{vk}) \leftarrow \mathsf{KeyGen}^{\mathcal{H}}(1^\lambda)$, runs $(m_i,\sigma_i)_{i\in[\ell+1]} \leftarrow \mathcal{U}^{\mathcal{H},\mathcal{O}_{\mathsf{Sign}}}(1^\lambda,\mathsf{vk})$ where $\mathcal{O}_{\mathsf{Sign}}(\cdot) := \mathsf{Sign}^{\mathcal{H}}(\mathsf{sk},\cdot)$ and ℓ is the number of calls $\mathcal{U}$ made to $\mathcal{O}_{\mathsf{Sign}}$, and outputs 1 if $m_i \neq m_j$ for all $i \neq j$ and $\mathsf{Verify}^{\mathcal{H}}(\mathsf{vk},m_i,\sigma_i) = 1$ for all $i \in [\ell+1]$. The advantage is $\mathrm{adv}^{\mathbf{omuf}}_{\mathsf{BS},\mathcal{U}} := \Pr[\mathbf{omuf}_{\mathsf{BS},\mathcal{U}} = 1]$.

Conjecture 1 (no round-optimal blind signature in the GGM+ROM). *Fix a family of groups $\mathbb{G} = \mathbb{G}(\lambda)$ of prime order $p = p(\lambda)$ and a random oracle $\mathcal{H}$ for groups as in Definition 1. Then there does not exist a round-optimal blind signature scheme $\mathsf{BS} = (\mathsf{KeyGen},\mathsf{User}_1,\mathsf{Sign},\mathsf{User}_2,\mathsf{Verify})$ in the ROM with message space $\mathcal{M} = \mathcal{M}(\lambda)$ (Definition 2) for which all of the following hold simultaneously:*

1. BS *is correct (Definition 2);*
2. *the algorithms User_1, User_2 and Verify are transformations of GGM algorithms $\overline{\mathsf{User}_1}$, $\overline{\mathsf{User}_2}$ and $\overline{\mathsf{Verify}}$ in the sense of Definition 1;*
3. *$|\mathcal{M}(\lambda)|$ is superpolynomial in λ, i.e. $|\mathcal{M}(\lambda)| = \lambda^{\omega(1)}$;*
4. *for every ppt GGM blindness adversary $\overline{\mathcal{S}}$, the advantage $\mathrm{adv}^{\mathbf{blind}}_{\mathsf{BS},\mathcal{S}}(\lambda)$ (Definition 3) is negligible in λ;*
5. *for every computationally unbounded GGM one-more-unforgeability adversary $\overline{\mathcal{U}}$ that makes at most a polynomial number of calls to its oracles $\mathcal{O}_{\mathsf{grp}}$, $\overline{\mathcal{H}}$ and $\overline{\mathcal{O}_{\mathsf{Sign}}}$, the advantage $\mathrm{adv}^{\mathbf{omuf}}_{\mathsf{BS},\mathcal{U}}(\lambda)$ (Definition 4) is negligible in λ.*

Equivalently, writing $q := q_{\mathsf{User}_2} + q_{\mathsf{Verify}}$ for the total number of $\mathcal{H}$-queries made by User_2 and Verify: for every polynomial Q there is no scheme as above with $q(\lambda) \le Q(\lambda)$.

Remark 1 (what is already a theorem). The case $q(\lambda) \leq O(\log \lambda)$ is a theorem [DKT26]; the conjecture is that the restriction $q \leq O(\log \lambda)$ may be dropped, no other hypothesis being changed.

Remark 2 (no query bound on the other algorithms). Note that no bound on q_{KeyGen}, q_{User_1} or q_{Sign} is imposed beyond the polynomial bound implied by these algorithms being ppt.

Remark 3 (what settling it would give). The conjecture is the exact remaining gap in the only formal evidence that the three-move barrier for pairing-free blind signatures is real. A proof would say that anyone who wants a round-optimal pairing-free blind signature must leave the black-box-group world entirely — which is precisely what the concurrent Fischlin-instantiation line of work is doing, at substantial cost — and it would do so for the whole class of schemes, not just those that hash a handful of times on the user's side. A refutation would be a construction of a kind nobody currently knows how to build, and it would tell us something surprising: that superlogarithmically many hash calls during finalization and verification are not a convenience but a resource that defeats generic-group extraction. Either way the answer is not a constant or an exposition improvement.

4 Bibliography

[BR93] M. Bellare, P. Rogaway. *Random oracles are practical: A paradigm for designing efficient protocols.* ACM CCS 93, pages 62–73, ACM Press, 1993.

[CFGG22] D. Catalano, D. Fiore, R. Gennaro, E. Giunta. *On the impossibility of algebraic vector commitments in pairing-free groups.* TCC 2022, Part II, LNCS volume 13748, pages 274–299, Springer, 2022.

[Cha82] D. Chaum. *Blind signatures for untraceable payments.* CRYPTO'82, pages 199–203, Plenum Press, 1982.

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

[DKT26] M. Dietz, J. Kastner, S. Tessaro. *On the Impossibility of Round-Optimal Pairing-Free Blind Signatures in the ROM.* Cryptology ePrint Archive, Paper 2026/090, 2026. [UNVERIFIED]

[Fis06] M. Fischlin. *Round-optimal composable blind signatures in the common reference string model.* CRYPTO 2006, LNCS volume 4117, pages 60–77, Springer, 2006.

[KNR24] J. Kastner, K. Nguyen, M. Reichle. *Pairing-free blind signatures from standard assumptions in the ROM.* CRYPTO 2024, Part I, LNCS volume 14920, pages 210–245, Springer, 2024.

[KSY11] J. Katz, D. Schröder, A. Yerukhimovich. *Impossibility of blind signatures from one-way permutations.* TCC 2011, LNCS volume 6597, pages 615–629, Springer, 2011.

[KTZ26] J. Kastner, S. Tessaro, G. Zaverucha. *Round-optimal pairing-free blind signatures.* Cryptology ePrint Archive, Paper 2026/091, 2026.

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