AICR · OPEN PROBLEM

Three-Move Pairing-Free Blind Signatures from Discrete Log

Black-box in the group, random oracle model only, unbounded concurrency

Status. Statement: AI-written from Rutchathon Chairattana-Apirom, Michael Reichle, Stefano Tessaro, *Playing Tag with Okamoto–Schnorr: Three-Move Pairing-Free Blind Signatures from DDH*, Cryptology ePrint Archive 2026 (full version of an article to appear at CRYPTO 2026), not yet checked by a human. Settled are three moves under DL but only in the algebraic group model on top of the random oracle model, three moves under DDH in the random oracle model alone (the source, and concurrently Chen), four moves under DL in the random oracle model alone, and the impossibility of two moves for schemes black-box in the group; open is the remaining cell — three moves, black-box in the group, random oracle model only, DL only.

Category. *Public-Key Cryptography & Assumptions.*

Conjecture (informal). *A blind signature lets a user get a signature on a message the signer never learns, and the signature cannot afterwards be linked back to the session that produced it. The security property that matters is one-more unforgeability: a user who completes some number of signing sessions must not walk away with even one signature more than that, no matter how many sessions are open at the same time. In ordinary elliptic-curve groups with no pairing, the cheapest known protocols exchange three messages, but until recently every three-message scheme with full concurrent security needed an idealized model of the group (the algebraic group model) on top of the random oracle. The source paper removes that: it gives a three-message scheme whose one-more unforgeability is proved in the random oracle model alone, but under decisional Diffie–Hellman rather than under the plain discrete logarithm assumption. Discrete log is the weakest and oldest assumption available in these groups, and three messages is known to be the fewest possible for schemes that treat the group as a black box, so one cell of the table is still empty. The conjecture is that this cell can be filled: that some three-message pairing-free blind signature, using the group only through generic group operations, is blind and one-more unforgeable in the random oracle model against adversaries running an unbounded polynomial number of concurrent sessions, assuming only that discrete logarithms are hard. The paper calls this the dream result of the area, names two specific reasons why gluing together the known techniques does not get there, and leaves it for future work.*

1 The setting

Blind signatures were introduced by Chaum [Cha82], and the security notion that has organised the area since Juels, Luby and Ostrovsky [JLO97] is *one-more unforgeability*: a malicious user who completes ℓ signing sessions, possibly all concurrently, must not obtain $\ell + 1$ signatures on distinct messages. Building such schemes over pairing-free groups is the practically relevant case, because these are the curves that cryptographic libraries actually support, and the theoretically interesting one, because such groups are close to the least structured algebraic setting in which efficient public-key cryptography lives.

The obstruction is concurrency. Blind Schnorr and blind Okamoto–Schnorr are three-move protocols that are broken by the ROS attack of Benhamouda et al. [BLL21], and the classical analyses [PS00] only cover a polylogarithmic number of concurrent sessions. Three-move schemes with *full* one-more unforgeability do exist [Abe01, TZ22], but their proofs invoke the Algebraic Group Model [FKL18] on top of the random oracle model. Removing the AGM was achieved only recently, at the cost of a fourth move and usually of an assumption stronger than discrete log: four-move schemes from CDH [CATZ24] and from DDH [KR25], and a four-move scheme from DL [KLR25].

The source paper closes the round gap. It gives the first three-move pairing-free blind signature whose one-more unforgeability is proved in the random oracle model alone, under DDH, by tagging the Okamoto–Schnorr nonce with an algebraic MAC and having the user blind the tag together with the transcript; the scheme is in fact one-more *strongly* unforgeable. The same round complexity, also under DDH, was obtained concurrently by Chen [Che26].

Three moves is optimal for this class of schemes: Dietz, Kastner and Tessaro [DKT26] show that no round-optimal (two-move) construction that uses the group only as a black box can be proved secure in the combined generic-group [Sho97, Mau05] and random-oracle model. The black-box restriction also matters on the assumption side: using the group non-black-box, via Fischlin's paradigm [Fis06] with generic NIZKs for group-dependent statements, one obtains two-move schemes in the random oracle model from a one-way function alone [CMV26], so the interesting question is what the group buys when it is used *as a group*.

What remains is a single cell of the table, which the source paper calls the "dream result": three moves, black-box in the group, random oracle model only, and discrete log as the sole assumption. The paper names two concrete obstructions to reaching it by combining the known techniques. First, its own construction needs an algebraic MAC and a linearly-homomorphic extractable commitment; both are known from DDH, and neither is known from DL in a black-box manner. Second, in the four-move DL scheme of [KLR25], signatures are Fiat–Shamir proofs for a *message-dependent* statement that the reduction punctures at the forgery message; because the statement depends on the message, the user must send the message in the first

move, and the remaining three moves are consumed establishing the proof. The paper explicitly leaves the reconciliation of the two techniques open.

Progress to date. What the source achieves is three moves, pairing-free, black-box in the group, in the random oracle model alone, with full (unbounded-concurrency) one-more strong unforgeability — but under DDH, in two instantiations: one with computational blindness and signatures of 15 group elements plus 22 scalars, and one with statistical blindness and signatures of 21 group elements plus 35 scalars; communication is 17 group elements plus 14 scalars in both. The concurrent work of Chen [Che26] achieves three moves in the random oracle model from DDH by a different route, using one-time signatures together with encrypted-message blind signatures. From DL alone, the best known in the random oracle model is four moves [KLR25]. From DL at three moves, the best known needs the AGM on top of the random oracle model [Abe01, TZ22]. Two moves is impossible for group-black-box schemes in the combined generic-group and random-oracle model [DKT26], so three is the floor.

2 Notation and parameters

- $\lambda \in \mathbb{N}$ is the security parameter, supplied to all algorithms in unary as 1^λ. A function $\nu : \mathbb{N} \to [0,1]$ is *negligible* if $\nu(\lambda) = \lambda^{-\omega(1)}$. *PPT* means probabilistic polynomial time in λ. We write $x \leftarrow_\$ S$ for sampling x uniformly from a finite set S.

- A *pairing-free group generator* GGen is a PPT algorithm that on input 1^λ outputs $(\mathbb{G}, p, G)$, where $\mathbb{G}$ is a cyclic group of prime order $p = p(\lambda)$ with $\lceil \log_2 p \rceil = \Theta(\lambda)$, and G generates $\mathbb{G}$. We write $\mathbb{G}$ additively, so xG denotes the x-fold sum of G for $x \in \mathbb{Z}_p$. No bilinear map on $\mathbb{G}$ is assumed, provided, or used.

- The *discrete logarithm advantage* of an algorithm $\mathcal{B}$ is

$$\mathsf{Adv}^{\mathrm{dl}}_{\mathsf{GGen}}(\mathcal{B},\lambda) := \Pr\big[\, x' = x \;:\; (\mathbb{G},p,G) \leftarrow \mathsf{GGen}(1^\lambda),\; x \leftarrow_{\$} \mathbb{Z}_p,\; x' \leftarrow \mathcal{B}(\mathbb{G},p,G,xG) \,\big].$$

  The *discrete logarithm (*DL*) assumption* holds relative to GGen if this quantity is negligible for every PPT $\mathcal{B}$.

- H denotes a random oracle, modelled as a uniformly random function from $\{0,1\}^*$ to a range fixed by the scheme; every algorithm and every adversary below has oracle access to H unless stated otherwise. It is the only idealized object allowed: in particular the algebraic group model and the generic group model may not be used in the security proof.

- $\mu \in \{0,1\}^*$ denotes a message to be signed, σ a signature, pp public parameters, and $(\mathsf{sk},\mathsf{pk})$ a signing key pair.

- $\mathsf{Adv}^{\mathrm{blind}}_{\mathsf{BS}}(\mathcal{A},\lambda)$ and $\mathsf{Adv}^{\mathrm{omuf}}_{\mathsf{BS}}(\mathcal{A},\lambda)$ are the blindness and one-more unforgeability advantages of $\mathcal{A}$ against a scheme BS, defined in Definitions 3 and 4.

- ℓ is the number of signing sessions the adversary completes in the one-more unforgeability game; it must not be able to produce more than that many signatures on distinct messages.

3 The conjecture

Definition 1 (Three-move blind signature scheme). A *three-move blind signature scheme* is a tuple

$$\mathsf{BS} = (\mathsf{Setup},\mathsf{KeyGen},\mathsf{Sign}_1,\mathsf{U}_1,\mathsf{Sign}_2,\mathsf{U}_2,\mathsf{Ver})$$

of PPT algorithms, each taking a group description $(\mathbb{G},p,G)$ as an implicit input and each with oracle access to H:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{pp}$;

- $\mathsf{KeyGen}(\mathsf{pp}) \to (\mathsf{sk}, \mathsf{pk})$;
- $\mathsf{Sign}_1(\mathsf{pp}, \mathsf{sk}) \to (\mathsf{smsg}_1, \mathsf{st}_S)$;
- $\mathsf{U}_1(\mathsf{pp}, \mathsf{pk}, \mu, \mathsf{smsg}_1) \to (\mathsf{umsg}, \mathsf{st}_U)$;
- $\mathsf{Sign}_2(\mathsf{umsg}, \mathsf{st}_S) \to \mathsf{smsg}_2$;
- $\mathsf{U}_2(\mathsf{smsg}_2, \mathsf{st}_U) \to \sigma \in \{0,1\}^* \cup \{\bot\}$;
- $\mathsf{Ver}(\mathsf{pp}, \mathsf{pk}, \mu, \sigma) \to b \in \{0,1\}$, deterministic.

A signing session on a message μ consists of exactly three transmitted messages: smsg_1 from signer to user, then umsg from user to signer, then smsg_2 from signer to user. No further communication is permitted.

Definition 2 (Correctness). BS is *correct* if for every λ, every $\mathsf{pp} \in \mathsf{Setup}(1^\lambda)$, every $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{KeyGen}(\mathsf{pp})$ and every $\mu \in \{0,1\}^*$, an honest execution

$$\begin{aligned} (\mathsf{smsg}_1, \mathsf{st}_S) &\leftarrow \mathsf{Sign}_1(\mathsf{pp}, \mathsf{sk}), \\ (\mathsf{umsg}, \mathsf{st}_U) &\leftarrow \mathsf{U}_1(\mathsf{pp}, \mathsf{pk}, \mu, \mathsf{smsg}_1), \\ \mathsf{smsg}_2 &\leftarrow \mathsf{Sign}_2(\mathsf{umsg}, \mathsf{st}_S), \\ \sigma &\leftarrow \mathsf{U}_2(\mathsf{smsg}_2, \mathsf{st}_U) \end{aligned}$$

satisfies $\mathsf{Ver}(\mathsf{pp}, \mathsf{pk}, \mu, \sigma) = 1$ with probability 1, the probability being over the coins of the four algorithms and over H.

Definition 3 (Blindness). For an adversary $\mathcal{A}$ acting as a malicious signer, consider the game: sample $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda)$ and run $\mathcal{A}(\mathsf{pp})$, which outputs a public key pk, two messages $\mu_0, \mu_1 \in \{0,1\}^*$ and a state. Sample $b \leftarrow_\$ \{0,1\}$. The game then runs two user instances, one on μ_b and one on μ_{1-b}, each executing U_1 and then U_2 with $(\mathsf{pp}, \mathsf{pk})$; $\mathcal{A}$ supplies all signer messages to both instances and chooses their interleaving. Let σ_0 and σ_1 be the outputs of the instances run on μ_0 and on μ_1 respectively. If $\sigma_0 = \bot$ or $\sigma_1 = \bot$, then $\mathcal{A}$ is given $(\bot, \bot)$;

otherwise $\mathcal{A}$ is given (σ_0, σ_1). Finally $\mathcal{A}$ outputs $b' \in \{0,1\}$. Set

$$\mathsf{Adv}^{\mathrm{blind}}_{\mathsf{BS}}(\mathcal{A},\lambda) := \left|\Pr[b' = b] - \tfrac{1}{2}\right|.$$

BS is *blind* if this is negligible for every PPT $\mathcal{A}$.

Definition 4 (One-more unforgeability). For an adversary $\mathcal{A}$ acting as a malicious user, consider the game: sample $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda)$ and $(\mathsf{sk},\mathsf{pk}) \leftarrow \mathsf{KeyGen}(\mathsf{pp})$, initialise a counter $\mathsf{cnt} \leftarrow 0$, a closed-session counter $\ell \leftarrow 0$ and a set $\mathcal{S} \leftarrow \emptyset$, and run $\mathcal{A}(\mathsf{pp},\mathsf{pk})$ with two oracles:

- $\mathsf{S}_1()$: set $\mathsf{cnt} \leftarrow \mathsf{cnt}+1$, compute $(\mathsf{smsg}_1, \mathsf{st}_\mathsf{S}^{(\mathsf{cnt})}) \leftarrow \mathsf{Sign}_1(\mathsf{pp},\mathsf{sk})$, and return $(\mathsf{cnt},\mathsf{smsg}_1)$;
- $\mathsf{S}_2(\mathsf{sid},\mathsf{umsg})$: if $\mathsf{sid} \notin \{1,\ldots,\mathsf{cnt}\}$ or $\mathsf{sid} \in \mathcal{S}$, return $\perp$; otherwise set $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathsf{sid}\}$, set $\ell \leftarrow \ell+1$, and return $\mathsf{Sign}_2(\mathsf{umsg}, \mathsf{st}_\mathsf{S}^{(\mathsf{sid})})$.

$\mathcal{A}$ finally outputs pairs $(\mu^{(1)},\sigma^{(1)}),\ldots,(\mu^{(k)},\sigma^{(k)})$ and wins if $k \geq \ell+1$, the messages $\mu^{(1)},\ldots,\mu^{(k)}$ are pairwise distinct, and $\mathsf{Ver}(\mathsf{pp},\mathsf{pk},\mu^{(i)},\sigma^{(i)}) = 1$ for all $i \in \{1,\ldots,k\}$. Set $\mathsf{Adv}^{\mathrm{omuf}}_{\mathsf{BS}}(\mathcal{A},\lambda) := \Pr[\mathcal{A}\text{ wins}]$. BS is *one-more unforgeable* if this is negligible for every PPT $\mathcal{A}$. Note that $\mathcal{A}$ may interleave its S_1 and S_2 queries arbitrarily, so any polynomial number of sessions may be open at once.

Definition 5 (Black-box use of the group). BS *makes black-box use of the group* if there is a single tuple of oracle PPT algorithms such that, for every group description $(\mathbb{G},p,G)$, the algorithms of BS on that description are obtained by giving those oracle algorithms the following interface, and nothing else: (i) the order p and an opaque handle for the generator G; (ii) an oracle that, given handles for $X,Y \in \mathbb{G}$, returns a handle for $X+Y$, and one that, given a handle for X, returns a handle for $-X$; (iii) an oracle that, given handles for X,Y, returns 1 if $X=Y$ and 0 otherwise; (iv) an oracle returning a handle for a uniformly random element of $\mathbb{G}$. The algorithms never read

the bit representation of a group element, and group elements appearing in pp, pk, protocol messages and signatures are transmitted as group elements. Equivalently, BS is well defined as a scheme in the generic group model of Shoup [Sho97] and Maurer [Mau05]. Adversaries are *not* restricted in this way: in Definitions 3 and 4 the adversary receives the real group description and may exploit the representation of group elements arbitrarily.

Conjecture 1 (three moves, black-box in the group, from DL**).** *There exist oracle PPT algorithms*

$$\mathsf{BS} = (\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{U}_1, \mathsf{Sign}_2, \mathsf{U}_2, \mathsf{Ver})$$

which form a three-move blind signature scheme in the sense of Definition 1, which make black-box use of the group in the sense of Definition 5, and which have access to a random oracle H*, such that the following holds for* every *pairing-free group generator* GGen*. Write* BS[GGen] *for the scheme obtained by sampling* $(\mathbb{G}, p, G) \leftarrow \mathsf{GGen}(1^\lambda)$ *and running the algorithms of* BS *on the group description* $(\mathbb{G}, p, G)$*. Then:*

1. (Correctness.) BS[GGen] *is correct in the sense of Definition 2.*
2. (Blindness from DL.) *If the discrete logarithm assumption holds relative to* GGen*, then for every PPT adversary* $\mathcal{A}$ *with access to* H*,* $\mathsf{Adv}^{\mathrm{blind}}_{\mathsf{BS}[\mathsf{GGen}]}(\mathcal{A}, \lambda)$ *as in Definition 3 is negligible in* λ*.*
3. (One-more unforgeability from DL.) *If the discrete logarithm assumption holds relative to* GGen*, i.e.* $\mathsf{Adv}^{\mathrm{dl}}_{\mathsf{GGen}}(\mathcal{B}, \lambda)$ *is negligible for every PPT* $\mathcal{B}$*, then* $\mathsf{Adv}^{\mathrm{omuf}}_{\mathsf{BS}[\mathsf{GGen}]}(\mathcal{A}, \lambda)$ *as in Definition 4 is negligible for every PPT adversary* $\mathcal{A}$ *with access to* H*. Here* $\mathcal{A}$ *is an arbitrary (not necessarily group-generic, not necessarily algebraic) PPT algorithm which receives the actual group description* $(\mathbb{G}, p, G)$*, and there is no a priori bound on the number of signing sessions* $\mathcal{A}$ *opens, nor on how*

> *many of them it keeps open concurrently, beyond $\mathcal{A}$ running in polynomial time.*

Remark 1 (the black-box clause is a formalisation). Definition 5 is read off the source paper's framing rather than copied from a definition it states, and it is what makes Conjecture 1 open. Without it the statement is already settled: Fischlin's paradigm [Fis06] with generic NIZKs and non-black-box use of the group yields two-move schemes in the random oracle model from a one-way function alone [CMV26], which DL implies. The source restricts its comparison table to schemes that make black-box use of the group, and invokes [DKT26] for the optimality of its own round complexity precisely because it makes only generic use of the group. A reviewer should check that Definition 5 is a faithful and workable rendering; competing formalisations (algebraic algorithms, Maurer-style generic algorithms with or without an equality oracle or a random-element oracle) differ at the margins, and one of them could conceivably make the question easier or harder.

Remark 2 (blindness is conditioned on DL). Clause 2 of Conjecture 1 asks for blindness only under the same hypothesis as clause 3, since an unconditional demand would also constrain generators relative to which discrete log is easy, which is more than "a three-move blind signature from DL" means. The unconditional variant is not vacuous — one of the source's two instantiations achieves statistical blindness — but it is not what the source poses.

Remark 3 (asymptotic, not concrete). Clause 3 is stated asymptotically rather than as a concrete reduction on purpose. The source's own DDH proof loses a cube root in the DL term of one of its cases, so demanding a linear-loss reduction would be strictly stronger than what "from DL" means in this literature.

Remark 4 (direction of the statement). The conjecture is stated positively because that is the direction the source points. Refutation would require an impossibility result for three-move group-black-box schemes under DL, which is a different kind of object from a construction; a reviewer may prefer a neutral "decide" phrasing.

Remark 5 (what settling it would give). Discrete log is the minimal standard assumption in a pairing-free group, the random oracle

model alone is the weakest idealisation the community accepts for these schemes, and three moves is provably optimal for group-black-box constructions. So this is not an incremental improvement: it is the last remaining cell of a table whose other cells have all been filled over the last two years, and it asks whether the three-move round complexity that all practically deployed pairing-free blind signature candidates use can be justified from the same assumption that Schnorr signatures themselves rest on. A positive answer would make three-move pairing-free blind signatures rest on nothing more than the hardness of discrete logarithms; a negative answer would be the first separation showing that concurrent blind signing genuinely costs more than plain signing in terms of assumptions, at fixed optimal round complexity.

Remark 6 (a fast-moving area). The source's reference list contains five works dated 2025 or 2026, two of them to appear at CRYPTO 2026, so it is possible that a three-move DL construction has appeared since. No such claim appears in the source, and the concurrent work of Chen [Che26] also stops at DDH.

4 Bibliography

[Abe01] M. Abe. *A secure three-move blind signature scheme for polynomially many signatures.* EUROCRYPT 2001, LNCS 2045, pages 136–151, Springer, 2001.

[BLL21] F. Benhamouda, T. Lepoint, J. Loss, M. Orrù, M. Raykova. *On the (in)security of ROS.* EUROCRYPT 2021, Part I, LNCS 12696, pages 33–53, Springer, 2021.

[CATZ24] R. Chairattana-Apirom, S. Tessaro, C. Zhu. *Pairing-free blind signatures from CDH assumptions.* CRYPTO 2024, Part I, LNCS 14920, pages 174–209, Springer, 2024.

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

[Che26] Y. Chen. *Three-move blind signatures in pairing-free groups.* To appear at CRYPTO 2026.

[CMV26] M. Ciampi, P. Della Monica, I. Visconti. *Round-optimal GUC-secure blind signatures from minimal computational and setup assumptions.* Cryptology ePrint Archive, Paper 2026/145, 2026.

[DKT26] M. Dietz, J. Kastner, S. Tessaro. *On the impossibility of round-optimal pairing-free blind signatures in the ROM.* To appear at CRYPTO 2026.

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

[FKL18] G. Fuchsbauer, E. Kiltz, J. Loss. *The algebraic group model and its applications.* CRYPTO 2018, Part II, LNCS 10992, pages 33–62, Springer, 2018.

[JLO97] A. Juels, M. Luby, R. Ostrovsky. *Security of blind digital signatures (extended abstract).* CRYPTO'97, LNCS 1294, pages 150–164, Springer, 1997.

[KLR25] M. Klooß, R. W. F. Lai, M. Reichle. *Blind signatures from arguments of inequality.* Cryptology ePrint Archive, Report 2025/1886, 2025.

[KR25] M. Klooß, M. Reichle. *Blind signatures from proofs of inequality.* CRYPTO 2025, Part VI, LNCS 16005, pages 157–189, Springer, 2025.

[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.

[PS00] D. Pointcheval, J. Stern. *Security arguments for digital signatures and blind signatures.* Journal of Cryptology, 13(3):361–396, 2000.

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

[TZ22] S. Tessaro, C. Zhu. *Short pairing-free blind signatures with exponential security.* EUROCRYPT 2022, Part II, LNCS 13276, pages 782–811, Springer, 2022.