AICR · OPEN PROBLEM

A Two-Round Adaptively Secure Threshold Signature With a Message-Independent First Round

Lattice assumptions, optimal corruption threshold $T - 1$, random oracle model

Status. Statement: AI-written from Kaijie Jiang, Stefano Tessaro, Hoeteck Wee, Chenzhi Zhu, *Tweed: Adaptively Secure Lattice-Based Two-Round Threshold Signatures*, Cryptology ePrint Archive 2026, not yet checked by a human. Settled: two-round adaptive security with both rounds message-dependent ($0 + 2$), from MLWE and MSIS, is achieved by this paper; $1 + 1$ lattice schemes exist but are proved only statically secure; and adaptive security with an offline round is achieved at three rounds in the pairing-free setting. Open: any scheme that is $1 + 1$ (two rounds total, message-independent first round) and adaptively secure against $T - 1$ corruptions, in the lattice setting, and — per the paper's Figure 1, which covers lattices and pairing-free groups without the algebraic group model — among pairing-free group-based schemes as well; the paper says nothing about pairing-friendly or algebraic-group-model constructions.

Category. *Public-Key Cryptography & Assumptions.*

Conjecture (informal). *A threshold signature lets any large-enough subset of a group of signers jointly sign a message under a single public key. Because signing is interactive, one counts rounds, and one distinguishes rounds that can be run before the message is known — those can be precomputed offline — from rounds that need the message. Separately one asks how strong the corruption model is: the conservative choice lets the adversary decide which signers to corrupt while protocols are already running, learning their key shares and all of their in-progress session state, up to one fewer than the threshold. These two goals have not so far been met together at two rounds. The efficient two-round lattice schemes do precompute their first round, but are only proved secure against corruptions fixed in advance, because their proofs rewind the adversary, and rewinding wrecks the simulation of corruptions once corruptions may happen after the forking point. The schemes proved secure against adaptive corruptions either need more than two rounds, or make both of their two rounds depend on the message. The conjecture is that the gap can be closed: there is a two-round threshold signature whose first round does not depend on the message and which is provably unforgeable against adaptive corruption of up to one fewer than the threshold number of signers, under standard lattice assumptions, in the random oracle model.*

1 The setting

A T-out-of-N threshold signature lets any T of N signers jointly produce a signature under a public key whose secret key is shared among them, while fewer than T of them cannot. Practical schemes are interactive, and the resource to minimise is the number of communication rounds. It is standard to split those rounds into an *offline* part, whose messages do not depend on the message μ to be signed and can therefore be precomputed, and an *online* part, which does depend on μ; a scheme with X offline and Y online rounds is written $X + Y$. A message-independent first round can be preprocessed, and signing can then be completed later, in a single round, once the message becomes known, so that such schemes "promise to be nearly as good as round-optimal ones". The paper immediately qualifies this: the feature "is likely less beneficial than one may think, as ensuring that parts of the preprocessed first-

round communication are not reused adds significant complexity from an engineering standpoint, and such reuse typically leads to a complete loss of security". The offline round is thus recorded both as something the paper does not achieve and as something whose practical value the paper questions.

A second axis is the corruption model. Static security fixes the corrupted set before the protocol runs. Adaptive security lets the adversary choose whom to corrupt during execution, on the basis of what it has already seen, and hands it the corrupted signer's key share together with the secret state of every one of that signer's open signing sessions. Adaptive security is the conservative model, NIST has stated its interest in it, and the folklore heuristic that a statically secure scheme is also adaptively secure has recently been undermined [CS25].

In the lattice setting the two axes have so far been incompatible at two rounds. The efficient two-round lattice schemes [EKT24, CATZ24, ZT25, BKL25] are all $1+1$: their first round is message-independent, so once μ arrives a signature costs one online round. All of them are proved only statically secure. The obstruction the paper names is rewinding: their proofs fork the adversary, and a reduction running two correlated executions cannot guarantee that fewer than T corruptions occur across both once corruptions may be made after the forking point. The one previously known adaptively secure lattice scheme [KRT24] needs five rounds.

The paper builds Tweed, a $0+2$ scheme — two rounds, but both message-dependent — that is adaptively secure against $T-1$ corruptions under MLWE and MSIS in the random oracle model. It follows the "the reduction knows the secret key, and so can simulate corruptions for free" strategy of Dazzle [Che25a] and Twinkle [BLT24], which sidesteps rewinding in the corruption simulation, and its underlying signature follows Lyubashevsky's framework [Lyu12] with noise flooding in place of rejection sampling.

The same gap is open over pairing-free groups, and the paper's own comparison table says so: the adaptively secure group-based schemes that do have an offline round need three rounds in total (for instance Twinkle [BLT24] at $2+1$ and FROST-Mask [Che25b] at $2+1$), while the two-round adaptively secure ones (Dazzle [Che25a], HBTS-Mask [Che25b], and Tweed) are all $0+2$. So in neither setting is there a scheme that is simultaneously two rounds, offline in its

first round, and adaptively secure at the optimal threshold $T - 1$. The authors record this as something Tweed does not achieve, and say their techniques are unlikely to reach it.

Progress to date. The paper delivers the $0 + 2$ adaptive lattice scheme, so the remaining gap is exactly the offline/online split. It also isolates why the existing $1 + 1$ schemes resist an adaptive proof: their reductions rewind, and rewinding forces the simulator to handle two correlated executions in which more than $T - 1$ corruptions may occur in total, all of which can be requested after the forking point. The paper's own proof shows that rewinding is not entirely forbidden in this style of argument — the MSIS branch of its single-signer proof does rewind — but the paper's own explanation for why this is safe is only that its reduction knows the secret key throughout and can therefore simulate corruptions without the forked run needing to succeed, and that its unforgeability proof "largely avoids rewinding" to begin with. The paper does not itself say why this suffices where the rewinding in the $1 + 1$ schemes does not; that comparison is left to the reader, not asserted by the authors.

Why it is worth settling. It is the last gap in a table the community has been filling in for three years: round count against corruption model. Closing it would give a post-quantum threshold signature where the only online cost, once a message arrives, is a single broadcast, with no weakening of the corruption threshold, no erasures, and no ideal group model. Refuting it — a separation showing that two rounds with an offline first round forces static security, or at least a $T/2$ threshold — would be the first genuine round/adaptivity separation for threshold signatures, and would explain a pattern that currently rests only on the observation that nobody has done it.

2 Notation and parameters

$\kappa \in \mathbb{N}$ is the security parameter; all algorithms take 1^κ as an implicit input, "PPT" means probabilistic polynomial time in κ, and $\mathrm{negl}(\kappa)$ denotes an unspecified function that is $\kappa^{-\omega(1)}$. For $N \in \mathbb{N}$ write $[N] := \{1, \ldots, N\}$.

N is the number of signers and T the threshold, with $1 \le T \le N$; both are polynomially bounded functions of κ. A *signer set* is a subset

$\mathsf{SS} \subseteq [N]$ with $|\mathsf{SS}| \geq T$. Messages are strings $\mu \in \{0,1\}^*$.

Let d be a power of two and q an odd prime, and set $\mathcal{R} := \mathbb{Z}[X]/(X^d+1)$ and $\mathcal{R}_q := \mathcal{R}/q\mathcal{R}$. Representatives of elements of $\mathcal{R}_q$ are taken with coefficients in $(-q/2, q/2]$. For $v \in \mathcal{R}$ with coefficient vector $(v_0, \ldots, v_{d-1}) \in \mathbb{Z}^d$, $\|v\|_\infty := \max_j |v_j|$, and for $\beta > 0$ set $\mathcal{B}_\beta := \{v \in \mathcal{R} : \|v\|_\infty \leq \beta\}$. For $\mathbf{x} \in \mathcal{R}^m$, $\|\mathbf{x}\|$ denotes the Euclidean norm of the concatenated coefficient vector of $\mathbf{x}$ in $\mathbb{Z}^{md}$. $\mathbf{I}_k$ is the $k \times k$ identity matrix over $\mathcal{R}_q$, and $k, m \in \mathbb{N}$ are module dimensions with $k < m$.

The parameters of the statement below are as follows.

- κ: security parameter.
- N: number of signers.
- T: threshold; any T signers can sign, and the adversary may adaptively corrupt up to $T-1$ of them.
- SS: the signer set of a signing session, a subset of $[N]$ of size at least T.
- d: degree of the cyclotomic ring $\mathcal{R} = \mathbb{Z}[X]/(X^d+1)$, a power of two.
- q: modulus of $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$, an odd prime.
- $k_1, m_1, \beta_{\mathsf{sis}}$: the MSIS parameters — number of rows and of columns of the matrix over $\mathcal{R}_q$, with $k_1 < m_1$, and the Euclidean-norm bound on solutions.
- $k_2, m_2, \beta_{\mathsf{lwe}}$: the MLWE parameters — number of rows and of columns of the matrix over $\mathcal{R}_q$, with $k_2 < m_2$, and the infinity-norm bound on secrets.

The two problems are parametrised independently: the norm bound tolerated by the MSIS branch of a proof of the kind envisaged here is a signature-norm bound, and is of a very different magnitude from the small secret bound of the MLWE branch, so no single tuple (k, m, β) should be made to serve both.

3 The conjecture

Definition 1 (two-round threshold signature with offline first round). A *two-round threshold signature scheme with offline first round* for N signers and threshold T is a tuple $\mathsf{TS} = (\mathsf{KeyGen}, \mathsf{Sign}_1, \mathsf{Sign}_2, \mathsf{Combine}, \mathsf{Ver})$ of polynomial-time algorithms, all of which take the public key pk implicitly as input:

- $\mathsf{KeyGen}(1^\kappa) \to (\mathsf{pk}, \{\mathsf{sk}_i\}_{i\in[N]})$.
- $\mathsf{Sign}_1(i, \mathsf{sk}_i, \mathsf{SS}) \to (\mathsf{st}, \mathit{psig}_i^{(1)})$, defined for $\mathsf{SS} \subseteq [N]$ with $|\mathsf{SS}| \geq T$ and $i \in \mathsf{SS}$, producing a secret state st and a first-round protocol message $\mathit{psig}_i^{(1)}$. *The message μ is not an input to Sign_1; this is what makes the first round offline.*
- $\mathsf{Sign}_2(i, \mathsf{sk}_i, \mathsf{st}, \mathsf{SS}, \mu, M) \to \mathit{psig}_i^{(2)}$, where $\mu \in \{0,1\}^*$ and $M = \{\mathit{psig}_j^{(1)}\}_{j\in\mathsf{SS}}$.
- $\mathsf{Combine}(\mathsf{SS}, \mu, M, M') \to \mathit{sig}$, where $M' = \{\mathit{psig}_j^{(2)}\}_{j\in\mathsf{SS}}$. Combine takes no secret input and may be run by any party.
- $\mathsf{Ver}(\mathsf{pk}, \mu, \mathit{sig}) \to 0/1$.

TS has *correctness error* $\varepsilon_{\mathsf{cor}}$ if for every $\mu \in \{0,1\}^*$ and every $\mathsf{SS} \subseteq [N]$ with $|\mathsf{SS}| \geq T$, running KeyGen, then Sign_1 for every $i \in \mathsf{SS}$, then Sign_2 for every $i \in \mathsf{SS}$ on the resulting M, then Combine, yields sig with $\mathsf{Ver}(\mathsf{pk}, \mu, \mathit{sig}) = 1$ except with probability at most $\varepsilon_{\mathsf{cor}}(\kappa)$.

Definition 2 (adaptive unforgeability, adp-TSUF). Let TS be as in Definition 1 and let H be a random oracle. The game $\text{adp-TSUF}^{\mathcal{A}}_{\mathsf{TS}[N,T]}(\kappa)$ initialises $S \leftarrow \emptyset$, $\mathsf{CS} \leftarrow \emptyset$, samples $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i\in[N]}) \leftarrow \mathsf{KeyGen}(1^\kappa)$, and for each $i \in [N]$ sets a counter $\mathsf{st}_i.\mathsf{cnt} \leftarrow 0$, an empty table $\mathsf{st}_i.\mathsf{P}$ of per-session states, and $\mathsf{st}_i.\mathsf{sk} \leftarrow \mathsf{sk}_i$. It runs $\mathcal{A}(\mathsf{pk})$ with access to:

- $\mathsf{oSign}_1(i, \mathsf{SS})$: requires $i \in \mathsf{SS} \setminus \mathsf{CS}$, $\mathsf{SS} \subseteq [N]$, $|\mathsf{SS}| \geq T$. Sets $\mathsf{st}_i.\mathsf{cnt} \leftarrow \mathsf{st}_i.\mathsf{cnt} + 1$, runs $(\mathit{psig}^{(1)}, \mathsf{st}) \leftarrow \mathsf{Sign}_1(i, \mathsf{sk}_i, \mathsf{SS})$,

stores $\mathsf{st}_i.\mathsf{P}(\mathsf{st}_i.\mathsf{cnt}) \leftarrow (0, \mathsf{st}, \mathsf{SS})$, and returns $\mathit{psig}^{(1)}$.

- $\mathrm{oSign}_2(i, \mathit{sid}, \mathsf{SS}, \mu, M)$: requires $i \in \mathsf{SS} \setminus \mathsf{CS}$, $\mathit{sid} \leq \mathsf{st}_i.\mathsf{cnt}$, $\mathsf{SS} \subseteq [N]$, $|\mathsf{SS}| \geq T$. Parses $(\mathit{isUsed}, \mathsf{st}, \mathsf{SS}') \leftarrow \mathsf{st}_i.\mathsf{P}(\mathit{sid})$ and returns $\perp$ if $\mathit{isUsed} = 1$ or $\mathsf{SS}' \neq \mathsf{SS}$. Otherwise it sets $S \leftarrow S \cup \{\mu\}$ and $\mathsf{st}_i.\mathsf{P}(\mathit{sid}) \leftarrow (1, \mathsf{st}, \mathsf{SS})$, and returns $\mathit{psig}^{(2)} \leftarrow \mathrm{Sign}_2(i, \mathsf{sk}_i, \mathsf{st}, \mathsf{SS}, \mu, M)$.
- $\mathrm{Cor}(i)$: requires $|\mathsf{CS}| < T - 1$. Sets $\mathsf{CS} \leftarrow \mathsf{CS} \cup \{i\}$ and returns the whole of st_i, i.e. the secret-key share sk_i together with the stored state of *every* session of signer i, used or unused. (There are no erasures, and signers keep no state beyond sk_i and these per-session states.)
- H.

When $\mathcal{A}$ halts with output (μ, sig), the game returns 1 if $|\mathsf{CS}| \leq T - 1$, $\mu \notin S$ and $\mathrm{Ver}(\mathsf{pk}, \mu, \mathit{sig}) = 1$, and 0 otherwise. Write $\mathrm{Adv}^{\text{adp-tsuf}}_{\mathsf{TS}[N,T]}(\mathcal{A}, \kappa) := \Pr[\text{adp-TSUF}^{\mathcal{A}}_{\mathsf{TS}[N,T]}(\kappa) = 1]$.

Two features of this game are adaptations of the paper's own, forced by removing μ from the first round. The message μ is added to S in oSign_2 rather than in oSign_1, since the first-round oracle does not see it. And oSign_2 rejects when $\mathsf{SS}' \neq \mathsf{SS}$: the signer set is now the only thing the first round commits to, so the second round re-checks it. The paper's own game stores only $(\mathit{isUsed}, \mathsf{st})$ and makes neither check.

Definition 3 (MLWE and MSIS). Fix d, q, k, m, β and let $\mathcal{R}_q$, $\mathcal{B}_\beta$ be as in the notation. The two problems below are parametrised separately; a statement that uses both quantifies over one tuple for each.

- $\mathrm{MSIS}_{q,d,k,m,\beta}$: sample $\mathbf{A} \leftarrow \mathcal{R}_q^{k \times m}$ and give it to C, which outputs $\mathbf{x} \in \mathcal{R}^m$; C wins if $\mathbf{x} \neq \mathbf{0}$, $\|\mathbf{x}\| \leq \beta$ and $\mathbf{A}\mathbf{x} = \mathbf{0} \bmod q$. Write $\mathrm{Adv}^{\text{msis}}_{q,d,k,m,\beta}(C, \kappa)$ for the winning probability.
- $\mathrm{MLWE}_{q,d,k,m,\beta}$: sample $\mathbf{A}' \leftarrow \mathcal{R}_q^{k \times (m-k)}$, set $\mathbf{A} := [\mathbf{A}' \mid \mathbf{I}_k]$, sample $\mathbf{s} \leftarrow \mathcal{B}_\beta^m$, set $\mathbf{t}_0 := \mathbf{A}\mathbf{s}$, sample $\mathbf{t}_1 \leftarrow \mathcal{R}_q^k$ and $b \leftarrow$

$\{0,1\}$, and give $(\mathbf{A}, \mathbf{t}_b)$ to $\mathcal{B}$, which outputs a guess $\tilde{b}$. Write $\mathrm{Adv}^{\mathrm{mlwe}}_{q,d,k,m,\beta}(\mathcal{B},\kappa) := |\Pr[\tilde{b} = b] - 1/2|$.

Conjecture 1 (a 1+1 adaptively secure lattice threshold signature). *For every pair of polynomially bounded functions $N = N(\kappa)$ and $T = T(\kappa)$ with $1 \le T(\kappa) \le N(\kappa)$ for all κ, there exist polynomially bounded parameters $d = d(\kappa)$ (a power of two), $q = q(\kappa)$ (an odd prime with $\log q = \mathrm{poly}(\kappa)$), and two independent tuples $(k_1, m_1, \beta_{\mathsf{sis}})$ and $(k_2, m_2, \beta_{\mathsf{lwe}})$ of polynomially bounded functions of κ, with $k_1 < m_1$ and $k_2 < m_2$, and a two-round threshold signature scheme TS with offline first round for N signers and threshold T in the random oracle model (Definition 1), such that both of the following hold.*

(i) Correctness. *TS has correctness error $\mathrm{negl}(\kappa)$ in the sense of Definition 1.*

(ii) Adaptive unforgeability from MLWE and MSIS. *Suppose that $\mathrm{MSIS}_{q,d,k_1,m_1,\beta_{\mathsf{sis}}}$ and $\mathrm{MLWE}_{q,d,k_2,m_2,\beta_{\mathsf{lwe}}}$ are hard, in the sense that every PPT $\mathcal{C}$ has $\mathrm{Adv}^{\mathrm{msis}}_{q,d,k_1,m_1,\beta_{\mathsf{sis}}}(\mathcal{C},\kappa) = \mathrm{negl}(\kappa)$ and every PPT $\mathcal{B}$ has $\mathrm{Adv}^{\mathrm{mlwe}}_{q,d,k_2,m_2,\beta_{\mathsf{lwe}}}(\mathcal{B},\kappa) = \mathrm{negl}(\kappa)$, with the two problems as in Definition 3. Then for every PPT adversary $\mathcal{A}$ making at most $\mathrm{poly}(\kappa)$ queries to each of oSign_1, oSign_2, Cor and the random oracle,*

$$\mathrm{Adv}^{\text{adp-tsuf}}_{\mathsf{TS}[N,T]}(\mathcal{A},\kappa) = \mathrm{negl}(\kappa),$$

where the advantage is as in Definition 2.

The two clauses are to be read together: what is asserted is the existence of a single scheme that is simultaneously (a) two rounds in total, (b) message-independent in its first round, (c) adaptively secure at the optimal corruption threshold $T-1$ with no erasures and no state carried across sessions, and (d) proved so under MLWE and MSIS in the random oracle model.

Remark 1 (the shape of a concrete bound). Clause (ii) is stated asymptotically on purpose. The reductions the paper itself achieves are not linear-loss: they are either/or statements, one branch reducing to MSIS with a cube-root (in one place, a square-root) loss and a quadratic dependence on the number Q_h of random oracle queries,

the other reducing to MLWE with a polynomial factor. Suppressing parameter subscripts, a concrete version of clause (ii) in the paper's own style would assert that for every PPT $\mathcal{A}$ there are $\mathcal{B}$, C of comparable running time and a constant c such that either

$$\mathrm{Adv}^{\text{adp-tsuf}} \;\le\; c\, Q_h^2 \big(\mathrm{Adv}^{\text{msis}} + 2^{-\Omega(\kappa)}\big)^{1/3}$$

or

$$\mathrm{Adv}^{\text{adp-tsuf}} \;\le\; \mathrm{poly}(\kappa) \cdot \mathrm{Adv}^{\text{mlwe}} + \mathrm{negl}(\kappa).$$

Demanding instead a single bound linear in $\mathrm{Adv}^{\text{mlwe}} + \mathrm{Adv}^{\text{msis}}$ would be strictly stronger than anything the paper achieves or asks for, and a scheme proved to the paper's own standard would fail it.

Remark 2 (readings fixed by this write-up). Several points are the page author's rather than the paper's. The paper never writes the words "open problem" here: it heads the paragraph "What we do not achieve", states the limitation, and says its techniques are unlikely to overcome the barrier; a reviewer should judge whether that is enough to count as the authors leaving it open. Relatedly, the paper takes no position on direction — an impossibility result is just as consistent with what it says — so the existence phrasing of Conjecture 1 is a choice about how to state the open problem, not something the authors assert. The paper also argues that a preprocessing round is worth less in practice than one might think, because reuse of preprocessed material is catastrophic and preventing it is an engineering burden; a reviewer may weigh that against the value of settling this. On the formalisation: the paper's own adp-TSUF game adds μ to the set S of signed messages in the first-round oracle, which is impossible when the first round does not see μ, so Definition 2 moves that bookkeeping to the second-round oracle and adds the signer-set check; a reviewer should confirm that this is the intended adaptation and does not weaken the game. Finally, the paper's comparison table is a snapshot; a reviewer should check whether any 2025–2026 follow-up — in particular further work along the lines of the two-round group-based scheme the paper attributes to Chen — already gives a $1 + 1$ adaptively secure scheme.

Remark 3 (why the statement is clean). The target is a single object with four crisply checkable properties: exactly two rounds, no μ in the first-round algorithm, negligible advantage in a fully

specified adaptive unforgeability game with $T - 1$ corruptions and no erasures, and a reduction to MLWE and MSIS. All four are pinned down by the definitions above, and the whole statement fits on one page. Someone can attack it with nothing beyond the standard lattice toolkit; no unpublished machinery is involved.

4 Bibliography

[BKL25] C. Boschini, D. Kaviani, R. W. F. Lai, G. Malavolta, A. Takahashi, M. Tibouchi. *Ringtail: Practical two-round threshold signatures from learning with errors.* 2025 IEEE Symposium on Security and Privacy, pages 149–164, IEEE Computer Society Press, 2025.

[BLT24] R. Bacho, J. Loss, S. Tessaro, B. Wagner, C. Zhu. *Twinkle: Threshold signatures from DDH with full adaptive security.* EUROCRYPT 2024, Part I, volume 14651 of LNCS, pages 429–459, Springer, 2024.

[CATZ24] R. Chairattana-Apirom, S. Tessaro, C. Zhu. *Partially non-interactive two-round lattice-based threshold signatures.* ASIACRYPT 2024, Part IV, volume 15487 of LNCS, pages 268–302, Springer, 2024.

[Che25a] Y. Chen. *Dazzle: Improved adaptive threshold signatures from DDH.* PKC 2025, Part III, volume 15676 of LNCS, pages 233–261, Springer, 2025.

[Che25b] Y. Chen. *Round-efficient adaptively secure threshold signatures with rewinding.* IACR Communications in Cryptology, 2(2), 2025.

[CS25] E. C. Crites, A. Stewart. *A plausible attack on the adaptive security of threshold Schnorr signatures.* CRYPTO 2025, Part VI, volume 16005 of LNCS, pages 457–479, Springer, 2025.

[EKT24] T. Espitau, S. Katsumata, K. Takemure. *Two-round threshold signature from algebraic one-more learning with errors.* CRYPTO 2024, Part VII, volume 14926 of LNCS, pages 387–424, Springer, 2024.

[KRT24] S. Katsumata, M. Reichle, K. Takemure. *Adaptively secure 5 round threshold signatures from MLWE/MSIS and DL with rewinding.* CRYPTO 2024, Part VII, volume 14926 of LNCS, pages 459–491, Springer, 2024.

[Lyu12] V. Lyubashevsky. *Lattice signatures without trapdoors.* EUROCRYPT 2012, volume 7237 of LNCS, pages 738–755, Springer, 2012.

[ZT25] C. Zhu, S. Tessaro. *The algebraic one-more MISIS problem and applications to threshold signatures.* CRYPTO 2025, Part I, volume 16000 of LNCS, pages 548–581, Springer, 2025.