AICR · OPEN PROBLEM

Tight Security of the Single-Session Lattice Chevallier-Mames Signature

One challenge per query instead of two, from search Module-LWE

Status. Statement: AI-written from Rutchathon Chairattana-Apirom, Nico Döttling, Julian Loss, Stefano Tessaro, Benedikt Wagner, *Tight Lattice-Based Signatures without Trapdoors from Search LWE*, Cryptology ePrint Archive, full version of an article to appear at CRYPTO 2026, 2026, not yet checked by a human. The two-session scheme SS_{R} is settled — the paper's Theorem 4.2 proves tight strong unforgeability for it from search Module-LWE — while the single-session scheme SS^1 below is open: for it the paper sketches only a reduction that resamples challenges and therefore loses a factor about $1/\varepsilon$, and states that it does not know how to prove tight security.

Category. *Public-Key Cryptography & Assumptions.*

Conjecture (informal). *There is a lattice signature scheme built by translating the Chevallier-Mames CDH-based signature into the Module-LWE world: the public key is a noisy linear sample, the signer derives a second matrix from a hash of its own commitment, publishes a second noisy sample under that matrix with the same secret, and proves that both samples are close to the row span of their respective matrices with the same secret. A security reduction wants to take a forgery and read the secret out of it, and it can: from two different accepting challenges for the same query it recovers the noise term and then the secret. The trouble is the arithmetic of how often that happens. If a single challenge is accepting with probability epsilon, then a pair of accepting challenges shows up only with probability epsilon squared, so the reduction has to keep resampling challenges, and the resulting proof loses a polynomial factor. The paper's fix is to run two independent copies of the protocol side by side under one public key with two independent challenges; a short inclusion-exclusion argument then shows that a pair of accepting challenges for one of the two copies is at least as likely as a forgery, which makes the reduction tight, at the cost of doubling signature size and signing time. The question left open is whether the doubling is really needed: is the single-copy scheme itself tightly secure under search Module-LWE, or is the parallel repetition inherent to this style of proof?*

1 The setting

A security reduction for a signature scheme is *tight* when a forger running in time t with success probability ε yields an algorithm for the underlying problem running in time $t' \approx t$ with success probability $\varepsilon' \approx \varepsilon$. Tightness matters concretely: a lossy reduction forces larger parameters to absorb the loss. For lattices, the hash-and-sign approach of Gentry, Peikert and Vaikuntanathan [GPV08] has a tight reduction to SIS but requires trapdoor preimage sampling, while the trapdoor-free Fiat–Shamir route of Lyubashevsky [Lyu09, Lyu12] is proved by rewinding and is not amenable to a tight reduction to SIS; tight proofs for Fiat–Shamir lattice signatures, including Dilithium's [KLS18], instead go through lossy identification and therefore rely on *decisional* LWE. Since known search-to-decision reductions for LWE lose a factor polynomial in $1/\varepsilon$, a proof from

search LWE is quantitatively preferable, and obtaining one for an efficient Fiat–Shamir lattice signature is the paper's stated goal.

The paper's starting point is the Chevallier-Mames signature [Che05], abstracted as a five-move identification protocol in [KLP17], which has a tight reduction to CDH: a signature is a Σ-protocol transcript proving that $X = g^x$ and $Z = h^x$ share a discrete logarithm, where h is a hash of the commitment and the message. The naive lattice translation is the scheme SS^1 of Definition 2: the public key is a Module-LWE sample $\mathbf{t} = \mathbf{sA} + \mathbf{e}$, the signer derives $\mathbf{H} = \mathsf{H}_1(\mathbf{u}, \mu)$ from its own commitment, publishes $\bar{\mathbf{t}} = \mathbf{sH} + \bar{\mathbf{e}}$, and the verification equations assert simultaneous proximity of $\mathbf{t}$ to $\mathbf{sA}$ and of $\bar{\mathbf{t}}$ to $\mathbf{sH}$. Signing is simulatable without the secret key by rejection sampling [Lyu12] together with a rerandomisation of $(\mathbf{A}, \mathbf{t})$ into $(\mathbf{H}, \bar{\mathbf{t}})$ [GMPW20]; hardness of sMLWE for a uniform secret and a Gaussian error of this width follows from worst-case module lattice problems [LS15], provided $\sigma_{\mathsf{sk}} > \sqrt{n} \cdot \omega(\sqrt{\log N})$.

The obstruction is on the extraction side. Unlike $\mathbf{e}$, the term $\bar{\mathbf{e}}$ is adversarially chosen, and an adversary can pick a non-short $\bar{\mathbf{e}}$ for which $c \cdot \bar{\mathbf{e}}$ is nonetheless short with inverse polynomial probability over the challenge c: the proximity proof is not sound, so decoding $\mathbf{s}$ from $\bar{\mathbf{t}}$ with a lattice trapdoor for $\mathbf{H}$ [MP12] fails. What survives is a shortcut: once c is fixed, the response $\mathbf{z}$ is statistically determined, so a reduction holding a trapdoor for $\mathbf{H}$ can decide from an H_2 query alone whether the query can lead to a forgery, without waiting for the forgery. From two accepting challenges $c \neq c'$ for the same query one gets

$$(c - c')\mathbf{t} = (\mathbf{z} - \mathbf{z}')\mathbf{A} + \mathbf{f} - \mathbf{f}',$$

whence $\mathbf{e} = (c - c')^{-1}(\mathbf{f} - \mathbf{f}')$ computed in the number field, and then $\mathbf{s}$ by linear algebra. But if a single challenge is accepting with probability ε, a *pair* of accepting challenges appears only with probability ε^2, so the reduction must resample about $1/\varepsilon$ times, and the proof is not tight.

The paper's fix is a two-fold parallel repetition: two sessions share one public key and one $\mathbf{H}$, and H_2 returns two independent challenges (c_0, c_1). Writing ε_b for the per-session acceptance probability, the inclusion-exclusion identity

$$\varepsilon_0^2 + \varepsilon_1^2 - \varepsilon_0^2\varepsilon_1^2 \;\geq\; \varepsilon_0\varepsilon_1$$

says that a pair of accepting challenges for *one* of the two sessions is at least as likely as an accepting pair (c_0, c_1) for both, i.e. at least as likely as the adversary's own success. This "polarization", combined with a balanced sampler that keeps the reduction's *worst-case* rather than expected running time under control [JT20] and a monotone coupling argument lifting the per-query guarantee to the whole game, yields a tight reduction from sMLWE to strong unforgeability of the two-session scheme, with polynomial modulus. That proof factors through key-only unforgeability: strong unforgeability under chosen-message attack follows from key-only unforgeability together with sMLWE (the paper's Section 4.3), and the extraction argument above is what carries the key-only step. The price of the repetition is exactly a factor two in signature size and signing time. The question recorded here is whether that price is necessary.

Progress to date. The paper gives, in its technical overview, a non-tight reduction for the single-session scheme: the reduction generates $\mathbf{H}$ with a lattice trapdoor, and because the response is statistically determined once the challenge is fixed, it can decide from each H_2 query whether the query is accepting; on an accepting query it resamples challenges until a second accepting one appears, which costs about $1/\varepsilon$ attempts, and then recovers $\mathbf{e}$ and $\mathbf{s}$ from the two transcripts. It also shows precisely what the two-session variant buys: for two independently sampled challenges the probability of an accepting pair within one session is at least the probability of an accepting pair across the two sessions, which is the adversary's advantage. No such identity is available with a single challenge per query, where the pair probability is ε^2.

2 Notation and parameters

Throughout, κ is the security parameter, $N = N(\kappa)$ is a power of two, $q = q(\kappa)$ is an odd prime, and

$$\mathcal{R} := \mathbb{Z}[\mathsf{X}]/\langle \mathsf{X}^N + 1\rangle, \qquad \mathcal{R}_q := \mathcal{R}/q\mathcal{R}.$$

We write $\mathrm{poly}(\kappa)$ and $\mathrm{negl}(\kappa)$ for the classes of polynomially bounded and of negligible functions of κ. All vectors are row vectors. An element $a = \sum_{i=0}^{N-1} a_i \mathsf{X}^i \in \mathcal{R}$ is identified with its coefficient vector $(a_0, \ldots, a_{N-1}) \in \mathbb{Z}^N$, and a vector $\mathbf{a} \in \mathcal{R}^m$ with the concatenation of the coefficient vectors of its entries, an element

of $\mathbb{Z}^{mN}$; $\|\cdot\|$, $\|\cdot\|_1$ and $\|\cdot\|_\infty$ denote the ℓ_2, ℓ_1 and ℓ_∞ norms of that integer vector. For $a \in \mathcal{R}$, $\mathsf{Rot}(a) \in \mathbb{Z}^{N\times N}$ is the negacirculant matrix of multiplication by a on $\mathcal{R}$, and for $\mathbf{e} = (e_1, \ldots, e_m) \in \mathcal{R}^m$ we set

$$\mathsf{Rot}(\mathbf{e}) := (\mathsf{Rot}(e_1) \mid \cdots \mid \mathsf{Rot}(e_m)) \in \mathbb{Z}^{N\times mN}.$$

For $\sigma > 0$ and $\mathbf{v} \in \mathcal{R}^m$, $\mathcal{D}_{\mathcal{R}^m,\mathbf{v},\sigma}$ denotes the discrete Gaussian on $\mathcal{R}^m$ (through the above embedding) with standard deviation σ centred at $\mathbf{v}$, i.e. the distribution with probability mass proportional to $\exp(-\pi\|\mathbf{x}-\mathbf{v}\|^2/\sigma^2)$; we write $\mathcal{D}_{\mathcal{R}^m,\sigma} := \mathcal{D}_{\mathcal{R}^m,\mathbf{0},\sigma}$ and denote by $\mathcal{D}_{\mathcal{R}^m,\mathbf{v},\sigma}(\mathbf{x})$ the probability it assigns to $\mathbf{x}$. For a symmetric positive definite $\mathbf{\Sigma} \in \mathbb{R}^{mN\times mN}$, $\mathcal{D}_{\mathcal{R}^m,\sqrt{\mathbf{\Sigma}}}$ is the discrete Gaussian on $\mathcal{R}^m$ centred at $\mathbf{0}$ with covariance $\mathbf{\Sigma}$, i.e. with probability mass proportional to $\exp(-\pi\,\mathbf{x}\mathbf{\Sigma}^{-1}\mathbf{x}^\top)$. We write $\mathbb{I}_k$ for the $k\times k$ identity matrix and $\otimes$ for the Kronecker product.

For $\gamma_c \in \mathbb{N}$, the challenge set is

$$C := \{\, c \in \mathcal{R} \;:\; \|c\|_\infty = 1 \;\wedge\; \|c\|_1 = \gamma_c \,\},$$

the set of ring elements with coefficients in $\{-1,0,1\}$ of which exactly γ_c are nonzero. We write $x \leftarrow S$ for sampling x uniformly from a finite set S or according to a distribution S.

The parameters of the statement below are as follows.

- κ: security parameter.
- N: degree of the cyclotomic ring, a power of two.
- q: odd prime modulus of the ring $\mathcal{R}_q$.
- n, m: row and column dimensions of the public matrix $\mathbf{A} \in \mathcal{R}_q^{n\times m}$; also the dimensions of the search Module-LWE instance.
- m': number of columns of the hash-derived matrix $\mathbf{H} \in \mathcal{R}_q^{n\times m'}$.
- σ_{sk}: standard deviation of the discrete Gaussian secret-key error $\mathbf{e}$.
- $\sigma_r, \bar{\sigma}_r$: standard deviations of the commitment noise $\mathbf{d}$ and $\bar{\mathbf{d}}$, and the target widths of the rejection sampling.
- σ_x, σ_y: parameters of the covariance $\bar{\mathbf{\Sigma}} = \sigma_x^2\mathsf{Rot}(\mathbf{e})\mathsf{Rot}(\mathbf{e})^\top + \sigma_y^2\mathbb{I}_N$ of the second error $\bar{\mathbf{e}}$.

- $\beta_{\mathbf{z}}, \bar{\beta}_{\mathbf{z}}$: verification norm bounds for the two proximity checks.
- γ_c: number of nonzero coefficients of a challenge, i.e. the ℓ_1 norm of every $c \in C$.
- M: rejection sampling parameter, $M = \exp(t/\alpha + 2/\alpha^2)$; for SS^1 the per-iteration acceptance probability is $1/M^2$, so the expected number of signing iterations is about M^2 (the paper's two-session scheme has $1/M^4$).
- α, t: the two rejection-sampling slack parameters entering M, both $\omega(\sqrt{\log(mN)})$.
- ℓ: maximum number of rejection sampling iterations before signing gives up.
- $\varepsilon_{\mathsf{sm}}$: the smoothing-parameter slack appearing in the bounds on σ_{sk} and η (written ε in the paper's Table 1; renamed here to keep ε for a forger's success probability).
- η: the Gaussian width below which the paper's rerandomisation and trapdoor arguments apply.
- b, σ_{td}: the gadget base and the trapdoor Gaussian width of the trapdoor sampler used by the reduction; they constrain m' and q only.
- Q_{S}, Q_1, Q_2: numbers of signing, H_1 and H_2 queries made by the forger.
- C_0, p: the multiplicative advantage loss and the per-query time overhead polynomial that make the reduction tight.

3 The conjecture

Definition 1 (Search Module-LWE). Let n, m, N, q be as above and $\sigma > 0$. For an algorithm $\mathcal{B}$ define

$$\mathsf{Adv}^{\mathrm{smlwe}}_{n,m,N,q,\sigma}(\mathcal{B}, \kappa) := \Pr[\, \mathbf{s} = \mathbf{s}' \,],$$

where the probability is over $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathcal{R}_q^n$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathcal{R}^m, \sigma}$ and $\mathbf{s}' \leftarrow \mathcal{B}(\mathbf{A}, \mathbf{sA} + \mathbf{e})$, with $\mathbf{A}$ and $\mathbf{s}$ uniform and all arith-

metic over $\mathcal{R}_q$. The assumption $\mathsf{sMLWE}_{n,m,N,q,\sigma}$ states that this advantage is negligible for every probabilistic polynomial-time $\mathcal{B}$.

Definition 2 (The single-session scheme SS^1). Let $n, m, m', \ell, \gamma_c \in \mathbb{N}$, $M \geq 1$, and $\sigma_{\mathsf{sk}}, \sigma_r, \bar{\sigma}_r, \sigma_x, \sigma_y, \beta_{\mathbf{z}}, \bar{\beta}_{\mathbf{z}} > 0$, and let $\mathsf{H}_1 : \{0,1\}^* \to \mathcal{R}_q^{n \times m'}$ and $\mathsf{H}_2 : \{0,1\}^* \to C$ be hash functions. The signature scheme

$$\mathsf{SS}^1 = \mathsf{SS}^1[n, m, m', N, q, \ell, M, \sigma_{\mathsf{sk}}, \sigma_r, \bar{\sigma}_r, \sigma_x, \sigma_y, \beta_{\mathbf{z}}, \bar{\beta}_{\mathbf{z}}, \gamma_c, \mathsf{H}_1, \mathsf{H}_2]$$

consists of the following algorithms, with all arithmetic over $\mathcal{R}_q$.

$\mathsf{KeyGen}(1^\kappa)$: sample $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathcal{R}_q^n$ and $\mathbf{e} \leftarrow \mathcal{D}_{\mathcal{R}^m, \sigma_{\mathsf{sk}}}$, set $\mathbf{t} \leftarrow \mathbf{sA} + \mathbf{e}$, and return $\mathsf{pk} := (\mathbf{A}, \mathbf{t})$ and $\mathsf{sk} := (\mathsf{pk}, \mathbf{s}, \mathbf{e})$.

$\mathsf{Sign}(\mathsf{sk}, \mu)$ for $\mu \in \{0,1\}^*$: set

$$\bar{\mathbf{\Sigma}} := \sigma_x^2 \, \mathsf{Rot}(\mathbf{e}) \mathsf{Rot}(\mathbf{e})^\top + \sigma_y^2 \, \mathbb{I}_N \in \mathbb{R}^{N \times N}.$$

Repeat the following at most ℓ times; if no signature is output, return $\perp$.

1. Sample $\mathbf{r} \leftarrow \mathcal{R}_q^n$, $\mathbf{d} \leftarrow \mathcal{D}_{\mathcal{R}^m, \sigma_r}$ and $\bar{\mathbf{d}} \leftarrow \mathcal{D}_{\mathcal{R}^{m'}, \bar{\sigma}_r}$, and set $\mathbf{u} \leftarrow \mathbf{rA} + \mathbf{d} \in \mathcal{R}_q^m$.
2. Set $\mathbf{H} \leftarrow \mathsf{H}_1(\mathbf{u}, \mu) \in \mathcal{R}_q^{n \times m'}$.
3. Sample $\bar{\mathbf{e}} \leftarrow \mathcal{D}_{\mathcal{R}^{m'}, \sqrt{\mathbb{I}_{m'} \otimes \bar{\mathbf{\Sigma}}}}$ and set $\bar{\mathbf{t}} \leftarrow \mathbf{sH} + \bar{\mathbf{e}} \in \mathcal{R}_q^{m'}$ and $\bar{\mathbf{u}} \leftarrow \mathbf{rH} + \bar{\mathbf{d}} \in \mathcal{R}_q^{m'}$.
4. Set $c \leftarrow \mathsf{H}_2(\mathbf{u}, \bar{\mathbf{u}}, \mathbf{H}, \bar{\mathbf{t}}, \mu) \in C$.
5. Set $\mathbf{z} \leftarrow \mathbf{r} + c\,\mathbf{s}$, $\mathbf{f} \leftarrow \mathbf{d} + c\,\mathbf{e}$ and $\bar{\mathbf{f}} \leftarrow \bar{\mathbf{d}} + c\,\bar{\mathbf{e}}$ (over $\mathcal{R}$).
6. With probability $\rho \cdot \bar{\rho}$ return $\sigma := (\bar{\mathbf{t}}, \mathbf{u}, \bar{\mathbf{u}}, \mathbf{z})$; otherwise go to the next iteration. Here

$$\rho := \min\left\{1, \frac{\mathcal{D}_{\mathcal{R}^m, \sigma_r}(\mathbf{f})}{M \, \mathcal{D}_{\mathcal{R}^m, c\mathbf{e}, \sigma_r}(\mathbf{f})}\right\}$$

and

$$\bar{\rho} := \min\left\{1, \frac{\mathcal{D}_{\mathcal{R}^{m'},\bar{\sigma}_T}(\bar{\mathbf{f}})}{M\,\mathcal{D}_{\mathcal{R}^{m'},\mathbf{c}\bar{\mathbf{e}},\bar{\sigma}_T}(\bar{\mathbf{f}})}\right\}.$$

$\mathsf{Ver}(\mathsf{pk}, \mu, \sigma)$: parse

$$\sigma = (\bar{\mathbf{t}}, \mathbf{u}, \bar{\mathbf{u}}, \mathbf{z}) \in \mathcal{R}_q^{m'} \times \mathcal{R}_q^{m} \times \mathcal{R}_q^{m'} \times \mathcal{R}_q^{n},$$

set $\mathbf{H} \leftarrow \mathsf{H}_1(\mathbf{u}, \mu)$ and $c \leftarrow \mathsf{H}_2(\mathbf{u}, \bar{\mathbf{u}}, \mathbf{H}, \bar{\mathbf{t}}, \mu)$, and return 1 if and only if both

$$\|\mathbf{u} + c \cdot \mathbf{t} - \mathbf{z}\mathbf{A}\| \le \beta_{\mathbf{z}},$$
$$\|\bar{\mathbf{u}} + c \cdot \bar{\mathbf{t}} - \mathbf{z}\mathbf{H}\| \le \bar{\beta}_{\mathbf{z}},$$

where the two vectors are taken with representatives in $\{-\lfloor q/2 \rfloor, \ldots, \lfloor q/2 \rfloor\}$ coefficientwise.

Definition 3 (Correctness and SUF-CMA security). A signature scheme SS has correctness error at most ε if for every $\mu \in \{0,1\}^*$, $\Pr[\mathsf{SS.Ver}(\mathsf{pk}, \mu, \sigma) = 0] \le \varepsilon$, over $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{SS.KeyGen}(1^\kappa)$ and $\sigma \leftarrow \mathsf{SS.Sign}(\mathsf{sk}, \mu)$ (an output of $\perp$ counts as a failure). For an adversary $\mathcal{A}$ with access to a signing oracle S that on input μ returns $\sigma \leftarrow \mathsf{SS.Sign}(\mathsf{sk}, \mu)$ and records the pair (μ, σ) in a set Sigs, define

$$\mathsf{Adv}_{\mathsf{SS}}^{\mathrm{sufcma}}(\mathcal{A}, \kappa) := \Pr\left[\, (\mu^*, \sigma^*) \notin \mathsf{Sigs} \ \wedge\ \mathsf{Win} \,\right],$$

where Win is the event $\mathsf{SS.Ver}(\mathsf{pk}, \mu^*, \sigma^*) = 1$ and the probability is over $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{SS.KeyGen}(1^\kappa)$ and $(\mu^*, \sigma^*) \leftarrow \mathcal{A}^{\mathsf{S}}(\mathsf{pk})$. In the random oracle model $\mathcal{A}$ additionally has oracle access to H_1 and H_2, drawn uniformly from the corresponding function spaces, and so do Sign and Ver.

Definition 4 (Admissible parameters). An *admissible param-*

eter family is a tuple of functions

$$\Pi = (n, m, m', N, q, \ell, \gamma_c, b, M, \sigma_{\mathsf{sk}}, \sigma_r, \bar{\sigma}_r, \sigma_x, \sigma_y, \beta_{\mathbf{z}}, \bar{\beta}_{\mathbf{z}}, \sigma_{\mathsf{td}}, \alpha, t, \varepsilon_{\mathsf{sm}})$$

of κ, the first eight taking values in $\mathbb{N}$ and the rest in $\mathbb{R}_{>0}$, with $\varepsilon_{\mathsf{sm}} \in (0,1)$, all computable in time $\mathrm{poly}(\kappa)$, such that N is a power of two, q is an odd prime, $b \geq 2$, $M \geq 1$, all of $n, m, m', \ell, \gamma_c, \log q$ are bounded by $\mathrm{poly}(\kappa)$, and such that for every κ the following hold, where

$$\eta := \tfrac{8}{\sqrt{\pi}}\, q^{n/m} \sqrt{N \ln\bigl(2mN(1+1/\varepsilon_{\mathsf{sm}})\bigr)}.$$

- $m > n$ and $m' \geq m + n \log_b q$;
- $\sigma_{\mathsf{sk}} \geq \sqrt{n \log N}$ and $\sigma_{\mathsf{sk}} \geq \sqrt{\log(2mN(1+1/\varepsilon_{\mathsf{sm}}))/\pi}$;
- $\sigma_x = \sqrt{2}\,\eta$ and $\sigma_y = \sigma_x\, \sigma_{\mathsf{sk}}\, N\sqrt{m}$;
- $\sigma_r = \alpha\, \gamma_c\, \sigma_{\mathsf{sk}} \sqrt{mN}$ and $\bar{\sigma}_r = \alpha\, \gamma_c\, \sigma_y \sqrt{2m'N}$;
- $\beta_{\mathbf{z}} = \sigma_r \sqrt{mN}$ and $\bar{\beta}_{\mathbf{z}} = \bar{\sigma}_r \sqrt{m'N}$;
- $M = \exp(t/\alpha + 2/\alpha^2)$, with $\alpha = \omega(\sqrt{\log(mN)})$ and $t = \omega(\sqrt{\log(mN)})$;
- $|C| \geq 2^{2\kappa}$;
- $\sigma_{\mathsf{td}} \geq \eta$ and $q \geq 4\, b\, \sigma_{\mathsf{td}}\, \bar{\beta}_{\mathbf{z}} \sqrt{m'mN}$.

These are the paper's Table 1 constraints with the parallel repetition removed.

Conjecture 1 (tight security of SS^1). *Let Π be an admissible parameter family as in Definition 4, and let H_1 and H_2 be modelled as random oracles. Then the scheme $\mathsf{SS}^1 = \mathsf{SS}^1[\Pi, \mathsf{H}_1, \mathsf{H}_2]$ of Definition 2 satisfies both of the following.*

(i) (Correctness.) *The correctness error of* SS^1 *is* $\mathrm{negl}(\kappa)$.

(ii) (Tight security.) *There exist a constant* $C_0 \geq 1$ *and a polynomial* p *such that for every adversary* $\mathcal{A}$ *against* SUF-CMA *security of* SS^1 *in the random oracle model, running in time* $\mathsf{Time}(\mathcal{A})$ *and making at most* Q_S *signing queries and at most* Q_1 *and* Q_2 *queries to* H_1 *and* H_2*, there is an adversary* $\mathcal{B}$ *against* $\mathsf{sMLWE}_{n,m,N,q,\sigma_{\mathsf{sk}}}$ *with*

$$\mathsf{Time}(\mathcal{B}) \;\leq\; \mathsf{Time}(\mathcal{A}) + \big(Q_S \ell + Q_1 + Q_2\big) \cdot p(\kappa)$$

and

$$\mathsf{Adv}^{\mathrm{sufcma}}_{\mathsf{SS}^1}(\mathcal{A},\kappa) \;\leq C_0 \cdot \mathsf{Adv}^{\mathrm{smlwe}}_{n,m,N,q,\sigma_{\mathsf{sk}}}(\mathcal{B},\kappa) + \mathrm{negl}(\kappa).$$

In particular, C_0 *and* p *do not depend on* $\mathcal{A}$*, and the loss is independent of* Q_S, Q_1, Q_2 *and of* $\mathsf{Adv}^{\mathrm{sufcma}}_{\mathsf{SS}^1}(\mathcal{A},\kappa)$.

Remark 1 (the constant C_0). The paper's own bound for the two-session scheme SS_{R} has the two-adversary form $\mathsf{Adv}^{\mathrm{smlwe}}(\mathcal{B}) + \mathsf{Adv}^{\mathrm{smlwe}}(\mathcal{B}') + \delta_{\mathrm{negl}}$, from its Lemmas 4.4 and 4.5, so $C_0 = 2$ is the faithful constant to aim at; the statement above leaves C_0 an unspecified constant because that is what tightness requires.

Remark 2 (readings fixed by this write-up). The sentence the paper leaves this open with refers to "this protocol", meaning the overview-level Attempt 1 on page 4, which omits the rejection sampling and all parameter choices. It is instantiated here as the one-session specialisation of the fully specified scheme in Figure 2, which is the faithful reading (the overview says the missing details are a routine modification), but a reader should check that specialisation against Figure 2 on page 19. Two further deliberate choices. The conjecture does not require a polynomial modulus, which the paper does achieve for the two-session scheme, since that is not part of what the paper says is open; and the formulation "there exists a reduction with constant loss" is the standard theorem shape in this literature and is not formalised by the paper itself, so anyone attacking the negative direction will need to fix a meta-reduction model. Finally, it has not been checked whether follow-up work has since settled this; the paper is dated 2026 and only its own reference

list was available.

Remark 3 (why the statement is clean). The object is the paper's own published scheme with one bit of the construction changed: run one session instead of two. Both the scheme and the assumption are fully specified in the paper, so the statement can be written from scratch in a page, and the target is a single well-defined implication rather than a research programme. The obstruction is named and quantitative: with one challenge per query, an accepting pair has probability ε^2 rather than ε, and the paper's inclusion-exclusion identity, which repairs this, has no single-session analogue.

Remark 4 (why it is interesting). The question isolates a clean structural point: is two-fold parallel repetition inherent to tight lattice Fiat–Shamir proofs from search assumptions, or an artifact of this particular argument? A positive answer halves the signature size and signing time of the first efficient lattice Fiat–Shamir signature with a tight reduction to a search assumption, on a scheme whose parameters are already large. A negative answer, in the form of a meta-reduction ruling out tight reductions for the single-session scheme, would be the first evidence that the repetition is necessary and would sit naturally beside the known impossibility results for tight reductions for Schnorr-type signatures. Either way the answer says something about the polarization technique the paper introduces, which is its main conceptual contribution and is claimed to be of independent interest.

4 Bibliography

[Che05] B. Chevallier-Mames. *An efficient CDH-based signature scheme with a tight security reduction.* CRYPTO 2005, LNCS 3621, pages 511–526, Springer, 2005.

[GMPW20] N. Genise, D. Micciancio, C. Peikert, M. Walter. *Improved discrete gaussian and subgaussian analysis for lattice cryptography.* PKC 2020, Part I, LNCS 12110, pages 623–651, Springer, 2020.

[GPV08] C. Gentry, C. Peikert, V. Vaikuntanathan. *Trapdoors for hard lattices and new cryptographic constructions.* 40th ACM STOC, pages 197–206, ACM Press, 2008.

[JT20] J. Jaeger, S. Tessaro. *Expected-time cryptography: Generic techniques and applications to concrete soundness.* TCC 2020, Part III, LNCS 12552, pages 414–443, Springer, 2020.

[KLP17] E. Kiltz, J. Loss, J. Pan. *Tightly-secure signatures from five-move identification protocols.* ASIACRYPT 2017, Part III, LNCS 10626, pages 68–94, Springer, 2017.

[KLS18] E. Kiltz, V. Lyubashevsky, C. Schaffner. *A concrete treatment of Fiat–Shamir signatures in the quantum random-oracle model.* EUROCRYPT 2018, Part III, LNCS 10822, pages 552–586, Springer, 2018.

[LS15] A. Langlois, D. Stehlé. *Worst-case to average-case reductions for module lattices.* Designs, Codes and Cryptography 75(3):565–599, 2015.

[Lyu09] V. Lyubashevsky. *Fiat–Shamir with aborts: Applications to lattice and factoring-based signatures.* ASIACRYPT 2009, LNCS 5912, pages 598–616, Springer, 2009.

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

[MP12] D. Micciancio, C. Peikert. *Trapdoors for lattices: Simpler, tighter, faster, smaller.* EUROCRYPT 2012, LNCS 7237, pages 700–718, Springer, 2012.