On the Quantum Security of HAWK

Serge Fehr[1,2] and Yu-Hsuan Huang[1]

[1] Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[2] Mathematical Institute, Leiden University, Leiden, The Netherlands
{serge.fehr, yhh}@cwi.nl

Abstract. In this paper, we prove the quantum security of the signature scheme HAWK, proposed by Ducas, Postlethwaite, Pulles and van Woerden (ASIACRYPT 2022). More precisely, we reduce its strong unforgeability in the quantum random oracle model (QROM) to the hardness of the one-more SVP problem, which is the computational problem on which also the classical security analysis of HAWK relies. Our security proof deals with the quantum aspects in a rather black-box way, making it accessible also to non-quantum-experts.

Keywords: quantum security· HAWK· digital signature· random oracle model

1 Introduction

Background. The discovery of Shor's algorithm has rendered most of the currently deployed public-key cryptosystems vulnerable to quantum attacks. As of 2016, the US National Institute of Standards and Technology (NIST) initiated the standardization process for post-quantum cryptography in the scope of key-encapsulation mechanism (KEM) and signature schemes. In 2022, the 3rd round winners were announced, but the process is still ongoing with the alternative KEM candidates and with a new call for signature schemes (see below). The selected signature schemes consist of Falcon [FHK+18] and Dilithium [DKL+18], which are lattice-based, and of SPHINCS+ [BHH+15], which is hash-based. For the sake of diversity, NIST has launched a new standardization process with a call for additional post-quantum signatures.

In 2022, [DvW22] introduced a new cryptographic framework based on the *lattice-isomorphism problem (LIP)*. The framework can be used to build various post-quantum cryptographic schemes, including KEMs and signatures. One particularly interesting scheme is the signature scheme HAWK, proposed in [DPPW22]. It uses the simple lattice $\mathbb{Z}^{2n}$, endowed with the (module) structure of cyclotomic ring $\mathbb{Z}[x]/(x^n + 1)$ for competitiveness. Due to this choice, the discrete Gaussian sampling (DGS), which is often the efficiency bottleneck, becomes much simpler and efficient. Indeed, by [DPPW22, Table 1], HAWK generally outperforms Falcon, which is considered to be one of the most efficient post-quantum signature schemes. It is an "open secret" that HAWK will be submitted to the new NIST standardization process mentioned above.

The classical security of HAWK has been analyzed and rigorously proven (in the random oracle model), via a security reduction to the considered — and believed-to-be (quantum) hard — one-more SVP problem [DPPW22]. Especially in the light of being a candidate in the new NIST post-quantum competition, the *quantum* security of HAWK is of particular interest.

As is common for security proofs in the random oracle model, the classical security proof for HAWK from [DPPW22] does not carry over to the quantum setting, where the attacker can make quantum superposition queries to the random oracle. Also, HAWK does not follow a standard construction design, for which one could apply an off-the-shelf quantum-security result (like [DFMS19,LZ19]). As a matter of fact, HAWK follows some non-standard randomized variant of the hash-and-sign paradigm, where first the hash $h := H(m, r)$ of the to-be-signed message m and some randomness r is computed, where r is then part of the signature $sig = (r, s)$, and the other part s is then sampled according to some distribution, which depends on the public key and on h, and that can be efficiently sampled if and (as far as we know) only if the secret key is given. The verification works by checking if some specific deterministic function of s and $H(m, r)$ satisfies some property (namely, is a non-zero short vector). Previous quantum analyses of generic hash-and-sign signature schemes, including the randomized variants considered in [BDF+11,Zha15], which rely on preimage samplable functions, do not apply to HAWK (independent of whether one considers classical or quantum attacks). Thus, an explicit quantum security proof for HAWK is necessary to establish provable quantum security.

Contribution. In this work, we analyze the quantum security of HAWK in the random oracle model. In particular, we prove strong unforgeability against chosen-message attacks in the quantum random oracle model (QROM), where the quantum attacker is given superposition access to the random oracle. Our proof is in the form of a security reduction to the (same) one-more SVP problem, with an explicit security loss. Our result positively confirms that the claimed quantum security of HAWK lies on a firm theoretical foundation.

Our quantum security proof for HAWK recycles some elements of the classical proof from [DPPW22], but requires some new elements to deal with the quantum aspect. For example, we invoke the adaptive reprogramming technique from [GHHM21] as well as the optimality of Grover for preimage search. Our proof is rather modular, and from our quantum security proof one can easily extract a variant of the classical security proof as well (with a slightly improved bound compared to the original classical proof in [GHHM21]), simply by replacing the adaptive reprogramming and the preimage search parts by their classical counterparts. In particular, our quantum security proof is meant to be accessible to a large extent also to non-quantum-experts.

2 Preliminary

2.1 Setting Up the Stage

Let $\mathcal{R}$ be the cyclotomic ring $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$, which is isomorphic to $\mathbb{Z}^n$ as a $\mathbb{Z}$-module; later on, n will be restricted to be an integer power of 2. Furthermore, we fix the obvious inclusion map $\mathcal{R}/2\mathcal{R} \hookrightarrow \mathcal{R}$, and thus consider the reduction mod 2 as a map $\mathcal{R} \to \mathcal{R}/2\mathcal{R} \subset \mathcal{R}$.

In order to abstract away the technical details of the property sym-break from [DPPW22], we consider the function $\langle \cdot \rangle : \mathcal{R}^2 \to \mathcal{R}^2$ defined by

$$v = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \mapsto \langle v \rangle := \begin{cases} -v \text{ if } \mathsf{sym\text{-}break}(v_1) = 1, \\ v \text{ otherwise,} \end{cases}$$

where, as defined in [DPPW22], $\mathsf{sym\text{-}break}(v_1) = 1$ if and only if $v_1 \neq 0$ with the first nonzero coefficient being positive. It is convenient to think of $\langle v \rangle$ as a representation of the equivalence class $\{v, -v\}$, with the representation being unique if $v_1 \neq 0$. Indeed, what will be relevant is that

$$\langle v \rangle = \langle -v \rangle \in \{v, -v\} \quad \forall\, v \notin \{0\} \times \mathcal{R}\,, \tag{1}$$

while $\langle v \rangle = v$ for $v \in \{0\} \times \mathcal{R}$.

Let $\mathbf{B} \in \mathsf{GL}_2(\mathcal{R})$ be an invertible 2×2 matrix over $\mathcal{R}$, and let $\mathbf{Q} = \mathbf{B}^*\mathbf{B}$. Looking ahead, $\mathbf{B}$ will form the secret key and $\mathbf{Q}$ the public key in HAWK. Such a Hermitian matrix $\mathbf{Q}$ induces the norm $\|v\|_{\mathbf{Q}} := \mathsf{tr}\,(v^*\mathbf{Q}v)\,/n$. As the name suggests, this is a norm in the $\mathbb{Z}$-module $\mathcal{R}^2$, meaning $\|u + v\|_{\mathbf{Q}} \leq \|u\|_{\mathbf{Q}} + \|v\|_{\mathbf{Q}}$, $\|av\|_{\mathbf{Q}} = |a|\,\|v\|_{\mathbf{Q}}$ and $\|v\|_{\mathbf{Q}} = 0 \Rightarrow v = 0$ for all $u, v \in \mathcal{R}$ and $a \in \mathbb{Z}$.

For every $\sigma > 0$ and $h \in \mathcal{R}^2/2\mathcal{R}^2 \subset \mathcal{R}^2$, consider the following $\mathbf{Q}$-dependent distributions $D_\sigma^{\mathbf{B}}$, $\widetilde{D}_\sigma^{\mathbf{B}}\,[\,h\,]$ and $\widetilde{D}_\sigma^{\mathbf{B}}$. They can be efficiently sampled if the matrix $\mathbf{B}$ is known, which motivates the superscript-$\mathbf{B}$ notation.

Definition 1. *For every $\sigma > 0$ and $\mathbf{Q} = \mathbf{B}^*\mathbf{B}$ as specified above, define*

- *$D_\sigma^{\mathbf{B}}$: the discrete (σ-deviated) Gaussian distribution under $\mathbf{Q}$-norm centered at 0 and supported at $\mathcal{R}^2$.*
- *$\widetilde{D}_\sigma^{\mathbf{B}}\,[\,h\,]$: the discrete ($\sigma$-deviated) Gaussian distribution under $\mathbf{Q}$-norm centered at 0 and supported at $h + 2\mathcal{R}^2 \subset \mathcal{R}^2$.*
- *$\widetilde{D}_\sigma^{\mathbf{B}}$: with a random choice of $h \leftarrow \mathcal{R}^2/2\mathcal{R}^2$, sampling $D'_\sigma\,[\,h\,]$.*

Note that for every $h \in \mathcal{R}^2/2\mathcal{R}^2 \subset \mathcal{R}^2$ we have

$$v \bmod 2 = h \quad \forall\, v \in \mathsf{supp}\left(\widetilde{D}_\sigma^{\mathbf{B}}\,[\,h\,]\right) = h + 2\mathcal{R}^2\,. \tag{2}$$

Furthermore, by Lemma 9 in the full version of [DPPW22], for every $\epsilon > 0$ and $\sigma \geq \eta_\epsilon(\mathbb{Z}^n)$, the statistical distance

$$\delta\left(D_\sigma^{\mathbf{B}}, \widetilde{D}_\sigma^{\mathbf{B}}\right) \leq \epsilon/(1 - \epsilon)\,, \tag{3}$$

where $\eta_\epsilon(\,\cdot\,)$ is as defined in [DPPW22, Definition 3]. This also implies that the distribution of $h := v \bmod 2$ for $v \leftarrow D_\sigma^{\mathbf{B}}$ is close to uniformly random in $\mathcal{R}^2/2\mathcal{R}^2$, with statistical distance at most $\epsilon/(1-\epsilon)$. Furthermore, as a direct consequence of [DPPW22, Lemma 3], the guessing probability of $v \bmod 2$ for $v \leftarrow D_\sigma^{\mathbf{B}}$ is bounded by

$$\mathrm{guess}(v \bmod 2) := \max_{h^\circ \in \mathcal{R}^2/2} \Pr[v \bmod 2 = h^\circ] \le 2^{-2n} \cdot \frac{1+\epsilon}{1-\epsilon}\,. \tag{4}$$

2.2 Geometric Units

Define the set of geometric units as $\mu_{\mathbb{K}} := \{x^1, \ldots, x^{2n}\} \subseteq \mathcal{R}$. The following Lemma 1 is recycled from the proof of Lemma 10 in the full version of [DPPW22], and will be useful in later analysis.

Lemma 1. *Let n be a power of 2, $\epsilon > 0$, $\sigma \ge \eta_\epsilon(\mathbb{Z}^n)$, and $h^\circ \in \mathcal{R}^2/2\mathcal{R}^2$. Consider $v \leftarrow D_\sigma$ and set $h := v \bmod 2$. Then*

$$\Pr\left[\exists\, \alpha \in \mu_{\mathbb{K}} \setminus \{\pm 1\} :\ \tfrac{1}{2}(h + \alpha v) \in \mathcal{R}^2\right] \le 2^{-n} \cdot \frac{1+\epsilon}{1-\epsilon}\,, \text{ and} \tag{5}$$

$$\Pr\left[\exists\, \alpha \in \mu_{\mathbb{K}} :\ \tfrac{1}{2}(h^\circ + \alpha v) \in \mathcal{R}^2\right] \le n \cdot 2^{-2n} \cdot \frac{1+\epsilon}{1-\epsilon}\,. \tag{6}$$

Proof. For (5), note that $\frac{1}{2}(h + \alpha v) \in \mathcal{R}^2$ implies that $h + \alpha v \in 2\mathcal{R}^2$, and thus $\alpha v \equiv h \equiv v \pmod 2$. Furthermore, any $\alpha \in \mu_{\mathbb{K}} \setminus \{\pm 1\}$ satisfies $\alpha \equiv x^i \pmod 2$ for some $1 \le i < n$, and so we have $x^i v \equiv v \pmod 2$ and thus by repeated application

$$x^{ki} v = x^{(k-1)i} x^i v \equiv x^{(k-1)i} v = \cdots \equiv v \pmod 2 \tag{7}$$

for any positive integer k. Furthermore, exploiting that, by the choice of n as a power of 2, $n/2$ must be a multiple of $\gcd(i, n)$ and thus can be written as $n/2 = ki + \ell n$ (where one may choose k to be positive), and using that $x^n \equiv 1 \pmod 2$, we obtain that

$$x^{n/2} v \equiv v \pmod 2\,.$$

Thus, we conclude (5) by

$$\begin{aligned}
&\Pr\left[\exists\, \alpha \in \mu_{\mathbb{K}} \setminus \{\pm 1\} :\ \tfrac{1}{2}(h + \alpha v) \in \mathcal{R}^2\right] \\
&\le \Pr\left[x^{n/2} v \equiv v \pmod 2\right] \\
&\le \mathrm{guess}(v \bmod 2) \cdot \#\left\{v^\circ \in \mathcal{R}^2/2\mathcal{R}^2 \;\middle|\; x^{n/2} v^\circ \equiv v^\circ \pmod 2\right\} \\
&\le 2^{-2n} \cdot \frac{1+\epsilon}{1-\epsilon} \cdot 2^n \\
&\le 2^{-n} \cdot \frac{1+\epsilon}{1-\epsilon}\,,
\end{aligned}$$

where we exploited (4).

Similarly, for (6), note that $\frac{1}{2}(h^\circ + \alpha v) \in \mathcal{R}^2$ implies that $v \equiv \alpha^{-1} h^\circ \pmod 2$, and furthermore $\#\{\alpha^{-1} h^\circ \bmod 2 \mid \alpha \in \mu_{\mathbb{K}}\} \le n$. Together with (4), we conclude the claimed bound (6). □

2.3 Adaptive Reprogramming Lemma

The following reprogramming lemma adapts from [GHHM21, Theorem 1], with the overall loss slightly improved. Intuitively, it states that, if the location x of a reprogramming is hard to guess prior to when it is taking place, then such a reprogramming is hard to notice.

Lemma 2 (Slight modification of [GHHM21, Theorem 1]). *Let $H : \mathcal{X} \to \mathcal{Y}$ be a random oracle, $\epsilon > 0$ and Ω be a family of distributions on $\mathcal{X}$ where every $\mathcal{D} \in \Omega$ is with guessing probability guess$(\mathcal{D}) := \max_{x^\circ} \Pr_{x \leftarrow \mathcal{D}}[x = x^\circ] \leq \epsilon$. Define the reprogramming oracle Repro_b for $b \in \{0,1\}$ that, on input (a suitable representation of) $\mathcal{D} \in \Omega$, works as below:*

$\mathsf{Repro}_0(\mathcal{D})$
1: $x \leftarrow \mathcal{D}$
2: $y := H(x)$
3: **return** (x, y)

$\mathsf{Repro}_1(\mathcal{D})$
1: $x \leftarrow \mathcal{D}$
2: $H(x) := y \leftarrow \mathcal{Y}$
3: **return** (x, y)

Suppose $\mathcal{A}^{\mathsf{Repro}_b, H}$ for $b \in \{0,1\}$ makes at most q_R queries to the reprogramming oracle Repro_b, and at most q_H quantum queries to H before the last reprogramming query. Then,

$$\left|\Pr\left[1 \leftarrow \mathcal{A}^{\mathsf{Repro}_0, H}\right] - \Pr\left[1 \leftarrow \mathcal{A}^{\mathsf{Repro}_1, H}\right]\right| \leq 2q_R\sqrt{(q_H + q_R) \cdot \epsilon}\,.$$

The intuition is quite simple: $\mathcal{A}$ can notice whether $H(x)$ gets reprogrammed or not only if it has queried x beforehand, which is unlikely the case since it is chosen with high entropy. The compressed oracle technique allows to make this line of reasoning rigorous, even when the queries to H are quantum: Before every Repro query we measure the compressed oracle to check whether x has been queried, we argue that the measurement outcome is "no" with overwhelming probability due to the high entropy in x, we conclude that the measurement caused little disturbance due to the gentle measurement lemma, and we observe that in case of a "no" outcome there is no difference between Repro_0 and Repro_1.

The full proof is based on the compressed oracle technique, but is rather standard. We refer readers to Appendix A for a detailed proof.

3 Brief Recap on HAWK and the One-More SVP

In the scope of HAWK, we take it as understood that the degree n of the cyclotomic ring $\mathcal{R}$ is a power of 2. Let $H : \mathcal{X} \to \mathcal{R}^2/2\mathcal{R}^2 \subseteq \mathcal{R}^2$ be a hash function, modelled as a random oracle, and let the parameters $\sigma_{\mathsf{pk}}, \sigma_{\mathsf{sign}}, \sigma_{\mathsf{ver}}, \mathsf{saltlen}$ be as specified in Lemma 10 in the full version of [DPPW22]. In particular, it holds that $2\sigma_{\mathsf{sign}} \leq \eta_\epsilon(\mathbb{Z}^n)$ for some negligible $\epsilon > 0$. We write *Gen* for the (σ_{pk}-dependent) key generation procedure specified in [DPPW22], producing a public-secret key pair $(\mathbf{Q}, \mathbf{B})$ with $\mathbf{B} \in \mathsf{GL}_2(\mathcal{R})$ and $\mathbf{Q} = \mathbf{B}^*\mathbf{B}$

The signing $Sign_{\mathbf{B}}$ and verification $Vrfy_{\mathbf{Q}}$ of HAWK works as follows.

$Sign_{\mathbf{B}}(m)$:
1: $r \leftarrow \{0,1\}^{\mathsf{saltlen}}$
2: $h := H(m, r)$
3: $v \leftarrow \widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}[\,h\,]$
4: $s := \frac{1}{2}(h + \langle v \rangle) \in \mathcal{R}^2$
5: **return** $sig := (r, s)$

$Vrfy_{\mathbf{Q}}(m, sig \in \{0,1\}^{\mathsf{saltlen}} \times \mathcal{R}^2)$:
1: $(r, s) := sig$
2: $h := H(m, r)$
3: $v := 2s - h$
4: **check** $v = \langle v \rangle$ and $v \notin \{0\} \times \mathcal{R}$
5: **check** $\|v\|_{\mathbf{Q}} \leq 2\sigma_{\mathsf{ver}} \cdot \sqrt{2n}$
6: **return** 1 if all check pass

Fig. 1. $Sign_{\mathbf{B}}$ and $Vrfy_{\mathbf{Q}}$ of HAWK

Remark 1. We take it as understood that $Vrfy_{\mathbf{Q}}$ implicitly checks that the signature $sig = (r, s)$ is well-formed, i.e. $r \in \{0,1\}^{\mathsf{saltlen}}$ and $s \in \mathcal{R}^2$.

Remark 2. The description in Fig. 1 matches the specification of HAWK (see [DPPW22, Algorithm 2 and Algorithm 3]) up to some small changes in the presentation (only). In particular, v as specified above coincides with $\frac{1}{2}\mathbf{B}^{-1}\mathbf{x}$ in the original specification of HAWK, and our definition of $s := \frac{1}{2}(h + \langle v \rangle)$ captures that $s := \frac{1}{2}h \mp \mathbf{B}^{-1}\mathbf{x}$, where the choice of the sign depends on $\mathsf{sym\text{-}break}(h_1 - 2s_2)$. Finally, the check $v = \langle v \rangle$ and $v \notin \{0\} \times \mathcal{R}$ is equivalent to checking $\mathsf{sym\text{-}break}(h_1 - 2s_1)$ as in the specification of HAWK.

Remark 3. Here we only concerns the uncompressed version of HAWK, while in practice an additional layer of compression is deployed for optimization. Nevertheless, it suffices to analyze the security of uncompressed HAWK because, according to [DPPW22, Section 3.2], the security of compressed HAWK follows immediately after.

Below, we describe the one-more SVP problem, as considered in [DPPW22], which considers an oracle algorithm $\mathcal{A}$ that makes at most q_S queries to the distribution $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$. We stress that when considering $\mathcal{A}$ to be a quantum algorithm, the queries to the oracle/distribution $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$ are restricted to be classical.

Definition 2. *Consider the one-more SVP game $G^{omSVP}_{\mathcal{A}}$, defined as follows:*

1: $(\mathbf{Q}, \mathbf{B}) \leftarrow Gen$
2: $v^* \leftarrow \mathcal{A}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}(\mathbf{Q})}$ *// Write $v_1, \ldots, v_{q'_S}$ for the responses given by $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$.*
3: **return** 1 *if and only if* $0 < \|v^*\|_{\mathbf{Q}} \leq 2\sigma_{\mathsf{ver}}\sqrt{2n}$ *and*

$$v^* \notin \{\alpha \cdot v_i \mid (i, \alpha) \in [q'_S] \times \mu_{\mathbb{K}}\} \ . \tag{8}$$

The advantage $adv^{omSVP}_{\mathcal{A}}$ of winning the one-more SVP game is then defined as $adv^{omSVP}_{\mathcal{A}} := \Pr\left[1 \leftarrow G^{omSVP}_{\mathcal{A}}\right]$.

4 Quantum Security of HAWK

4.1 Warming Up: NMA Security

As a warm up, let $\mathcal{A}$ be an NMA attacker against HAWK, which on input the public key $\mathbf{Q}$ outputs a message-forgery pair (m^*, sig^*) with $sig^* = (r^*, s^*) \in$

$\{0,1\}^{\mathsf{saltlen}} \times \mathcal{R}^2$. Furthermore, consider the algorithm $\mathcal{E}$ that on input such a message-forgery pair (m^*, sig^*) computes

$$h^* := H(m^*, r^*) \qquad \text{and} \qquad v^* := 2s^* - h^* \tag{9}$$

and outputs v^*. Then $Vrfy_{\mathbf{Q}}(m^*, sig^*) = 1$ only if $0 < \|v^*\|_{\mathbf{Q}} \leq 2\sigma_{\mathsf{ver}} \cdot \sqrt{2n}$ by the definition of $Vrfy_{\mathbf{Q}}$. Thus, if $\mathcal{A}$ succeeds in forging a signature then $\mathcal{B} := \mathcal{E} \circ \mathcal{A}$ succeeds in finding a non-zero short vector. Formally,

$$adv_{\mathcal{A}}^{\mathsf{NMA}} := \Pr\left[Vrfy_{\mathbf{Q}}(m^*, sig^*) = 1 \,\middle|\, \begin{matrix} (\mathbf{Q}, \mathbf{B}) \leftarrow Gen \\ (m^*, sig^*) \leftarrow \mathcal{A}(\mathbf{Q}) \end{matrix} \right] \leq adv_{\mathcal{B}}^{\mathsf{omSVP}}.$$

We note that the above reasoning holds in the plain model with H being an arbitrary hash function, as well as in the random oracle model.

Remark 4. The reduction algorithm $\mathcal{B}$ here from NMA to one-more SVP does not make any query to $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$.

4.2 Full CMA Quantum Security

Consider a CMA attacker $\mathcal{A}^{Sign_{\mathbf{B}},H}(\mathbf{Q})$ against HAWK in the random oracle model, which on input the public key $\mathbf{Q}$ makes at most q_H queries to the random oracle H and at most q_S queries to the signing oracle $Sign_{\mathbf{B}}$, and eventually outputs a message-forgery pair (m^*, sig^*) with $sig^* = (r^*, s^*) \in \{0,1\}^{\mathsf{saltlen}} \times \mathcal{R}^2$. Without loss of generality, we assume $\mathcal{A}$ makes exactly q_H, q_S queries to $H, Sign_{\mathbf{B}}$ respectively.[3] The goal is to turn $\mathcal{A}$ into an algorithm $\mathcal{B}$ that solves the one-more-SVP problem.

Theorem 1 (Quantum Security of HAWK). *Let HAWK be as specified in Section 3, and let $\mathcal{A}^{Sign_{\mathbf{B}},H}(\mathbf{Q})$ be a chosen-message attack making at most q_S queries to $Sign_{\mathbf{B}}$ and at most q_H quantum queries to H respectively. Then there exists an algorithm $\mathcal{B}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}$ making q_S queries to solve one-more SVP, with running time $\mathsf{TIME}(\mathcal{B}) \approx \mathsf{TIME}(\mathcal{A}) + \mathsf{Overhead}(q_S, q_H)$ consisting of an additional overhead $\mathsf{Overhead}(q_S, q_H)$ of simulating q_H, q_S queries to H and $Sim^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}$ (specified in Fig. 2), such that*

$$\begin{aligned} adv_{\mathcal{A}}^{sUF\text{-}CMA} \leq\ & adv_{\mathcal{B}}^{omSVP} + \frac{q_S \epsilon}{1-\epsilon} + 2q_S\sqrt{q_H + q_S} \cdot 2^{-\mathsf{saltlen}/2} \\ & + q_S\left(2^{-n} + (q_S - 1) \cdot n \cdot 2^{-2n}\right) \cdot \frac{1+\epsilon}{1-\epsilon} + O\big(q_H^2 \cdot n \cdot q_S / 2^{2n}\big), \end{aligned}$$

where the CMA advantage is defined as below:

$$adv_{\mathcal{A}}^{sUF\text{-}CMA} := \Pr\left[\begin{matrix} Vrfy_{\mathbf{Q}}(m^*, sig^*) = 1 \\ \forall i \in [q_S] : (m^*, r^*) \neq (m_i, sig_i) \end{matrix} \middle| \begin{matrix} (\mathbf{Q}, \mathbf{B}) \leftarrow Gen \\ (m^*, sig^*) \leftarrow \mathcal{A}^{Sign_{\mathbf{B}},H}(\mathbf{Q}) \end{matrix} \right],$$

with (m_i, sig_i) in the probability being the transcript produced at the ith signing query.

[3] Otherwise, we let $\mathcal{A}$ make dummy queries to H and $Sign_{\mathbf{B}}$ respectively, with the dummy queries to $Sign_{\mathbf{B}}$ being on messages different from m^*, so that they do not affect the freshness of a forgery.

Simulating the signing queries. First, we show that we can replace the signing oracle *Sign* by a particular simulator *Sim* that does not (explicitly) hold the secret key $\mathbf{B}$, but instead has access to the $\mathbf{Q}$-dependent distribution $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$, and that can reprogram the random oracle H; see Fig. 2 (right) below. Towards this goal, we also consider the oracle $\mathit{Trans}_{\mathbf{B}}$ as specified in Fig. 2 (left), and we show that

$$\mathcal{A}^{\mathit{Sign}_{\mathbf{B}},H}(\mathbf{Q}) \approx \mathcal{A}^{\mathit{Trans}_{\mathbf{Q}},H}(\mathbf{Q}) \approx \mathcal{A}^{\mathit{Sim}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}},H}(\mathbf{Q})\,.$$

We have used subscript $\mathbf{Q}, \mathbf{B}$ to indicate that the oracle's execution depends on the keys, but for later convenience, we may also omit those subscripts based on the relevance of the context.

$\mathit{Trans}_{\mathbf{B}}(m)$:
1: $r \leftarrow \{0,1\}^{\mathsf{saltlen}}$
2: $H(m,r) := h \leftarrow (\mathcal{R}/2)^2$
3: $v \leftarrow \widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}[\,h\,]$
4: $s := \frac{1}{2}(h + \langle v \rangle) \in \mathcal{R}^2$
5: **return** $\mathit{sig} := (r,s)$

$\mathit{Sim}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}(m)$:
1: $r \leftarrow \{0,1\}^{\mathsf{saltlen}}$
2: $v \leftarrow D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$
3: $H(m,r) := h := v \mod 2$
4: $s := \frac{1}{2}(h + \langle v \rangle) \in \mathcal{R}^2$
5: **return** $\mathit{sig} := (r,s)$

Fig. 2. Oracles $\mathit{Trans}_{\mathbf{B}}$ and $\mathit{Sim}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}$

Note that, the only difference between *Sign* and *Trans*, is that the former computes $h := H(m,r)$ while the latter replaces it with reprogramming $H(m,r) := h \leftarrow \mathcal{R}^2/2\mathcal{R}^2$ for a freshly chosen h. It follows therefore directly from Lemma 2, with $\epsilon = 2^{-\mathsf{saltlen}}$ and $q_R = q_S$, that

$$\Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{\mathit{Sign},H}\right] - \Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{\mathit{Trans},H}\right] \leq 2q_S\sqrt{q_H + q_S}/2^{\mathsf{saltlen}/2}\,, \tag{10}$$

where it is understood that the verification Vrfy^H is performed using the possibly reprogrammed H. Furthermore, by the closeness of the distributions $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$ and $\widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$ (see Lemma 3 for the detailed reasoning), replacing the calls to *Trans* one-by-one by calls to *Sim*, one obtains

$$\Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{\mathit{Trans},H}\right] - \Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{\mathit{Sim},H}\right] \leq \frac{q_S \cdot \epsilon}{1-\epsilon}\,.$$

We thus conclude that the *validity* of a forgery is preserved when replacing the signing oracle *Sign* by *Sim*, up to the sum of the two above probabilities.

Furthermore, the *freshness* of a forgery is also preserved, in that we can assume without loss of generality that $\mathcal{A}$ never outputs a forgery (m^*, sig^*) that matches the response of a signing query.

Lemma 3. *Let $\epsilon > 0$ and $\sigma_{\mathsf{sign}} \geq \eta_\epsilon(\mathbb{Z}^n)$. Then the respective distributions of (r,h,s,v) in an execution of Sim and of Trans have statistical distance at most $\epsilon/(1-\epsilon)$.*

Proof. First, we note that in *Trans*, right after the choice of $v \leftarrow \widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}[\,h\,]$ in line 3, we can redefine $H(m,r) := h := v \bmod 2$ with no effect, since $v \bmod 2 = h$ for $v \leftarrow \widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}[\,h\,]$ by (2). But now, the only difference between *Trans* and *Sim* is that in the former v is sampled by $\widetilde{D}^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$ and in the latter by $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$. The claim thus follows from (3). □

Extracting a fresh short vector. Slightly abusing notation, we now consider the algorithm $\mathcal{B}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}} := \mathcal{E}^{H} \circ \mathcal{A}^{Sim^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}},H}$ where, as before, $\mathcal{E}$ computes h^* and v^* as in (9) and outputs v^*, and where we take it as understood that the random oracle H is simulated by $\mathcal{B}$. As in the NMA case, it follows that if $\mathcal{A}^{Sim,H}$ succeeds in producing a valid forgery then $\mathcal{B}$'s output v^* is a short non-zero vector, i.e., $0 < \|v^*\|_{\mathbf{Q}} \leq 2\sigma_{\mathsf{ver}} \cdot \sqrt{2n}$. It remains to show that v^* is fresh as well.

To show that this holds (almost with certainty), we assume that (m^*, sig^*) is a valid and fresh forgery (where the latter is without loss of generality), yet $v^* = \alpha v_j$ for some $(j,\alpha) \in [q_S] \times \mu_{\mathbb{K}}$, and we show that this implies an event that has negligible probability. For this purpose, let m_i, r_i, h_i, s_i, v_i be the transcripts m, r, h, s, v produced at the ith query to *Sim*, and let $sig^* = (r^*, s^*)$ be the signature output by $\mathcal{B}$. We distinguish between the following two cases.

Case $(m^*, r^*) = (m_i, r_i)$ for some $i \in [q_S]$, where we consider i to be maximal such that the equality holds.[4] Then $h^* = h_i$, and so

$$\mathcal{R}^2 \ni s^* = \frac{1}{2}(h^* + v^*) = \frac{1}{2}(h_i + \alpha v_j)\,.$$

However, if $i \neq j$ then for any fixed choice of h_i, the probability over the choice of v_j of there being an $\alpha \in \mu_{\mathbb{K}}$ as above, is at most $n \cdot 2^{-2n} \cdot \frac{1+\epsilon}{1-\epsilon}$ by (6). On the other hand, if $i = j$ then we get that

$$\mathcal{R}^2 \ni s^* = \frac{1}{2}(h^* + v^*) = \frac{1}{2}(h_i + \alpha v_i)\,.$$

Furthermore, $\alpha \neq \pm 1$ then; indeed, otherwise, $\langle v_i \rangle = \langle v^* \rangle = v^*$, where the second equality holds by the validity of sig^* and the first follows from $\{0\} \times \mathcal{R} \not\ni v^* = \pm v_i$ and (1), and so $s^* = \frac{1}{2}(h_i + \langle v_i \rangle) = s_i$ which contradicts the freshness of sig^*. However, the probability over the choice of v_i of there being an $\alpha \in \mu_{\mathbb{K}} \setminus \{\pm 1\}$ as above, is at most $2^{-n} \cdot \frac{1+\epsilon}{1-\epsilon}$ by (5).

In case $(m^*, r^*) \neq (m_i, r_i)$ for every $i \in [q_S]$, we must have that

$$\mathcal{R}^2 \ni s^* = \frac{1}{2}(h^* + v^*) = \frac{1}{2}(h^* + \alpha v_j)\,,$$

and so $H(m^*, r^*) = h^* = \alpha v_j \bmod 2$; furthermore, H has not been reprogrammed throughout the execution at the location (m^*, r^*). Hence

$$H_{\mathsf{init}}(m^*, r^*) = H(m^*, r^*) \in \{\alpha v_j \bmod 2 \mid (j,\alpha) \in [q_S] \times \mu_{\mathbb{K}}\} =: S\,, \qquad (11)$$

[4] If i is not the largest, it can be $(m^*, r^*) = (m_i, r_i)$ yet $h^* \neq h_i$ because h^* is computed via the possibly reprogramed H.

where H_{init} is the initial H before being reprogrammed. Thus, parsing $\mathcal{A}^{Sim,H}$ as $\mathcal{C}^{H_{\mathsf{init}}}$, which runs the calls to *Sim* (for arbitrary but fixed samples $v_1, \ldots, v_{q_S}$ of $D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}$) and the reprogramming of H internally, we obtain a preimage-finding algorithm that finds a preimage under H_{init} of an element in S, making q_H queries to H_{init}. Given that $\#S \leq n \cdot q_S$, such an algorithm can succeed with probability at most $O\big(q_H^2 \cdot n \cdot q_S / 2^{2n}\big)$ via the standard preimage finding bound.

Collecting all the different error terms, we obtain that

$$\begin{aligned} adv^{\text{sUF-CMA}}_{\mathcal{A}} \leq\ & adv^{\text{omSVP}}_{\mathcal{B}} + \frac{q_S \epsilon}{1-\epsilon} + 2q_S\sqrt{q_H + q_S} \cdot 2^{-\mathsf{saltlen}/2} \\ & + q_S\left(2^{-n} + (q_S - 1) \cdot n \cdot 2^{-2n}\right) \cdot \frac{1+\epsilon}{1-\epsilon} + O\big(q_H^2 \cdot n \cdot q_S / 2^{2n}\big)\ , \end{aligned}$$

which concludes Theorem 1.

4.3 Classical Security

As our proof is modular, one may substitute certain part of the proof of Theorem 1 to obtain better bounds when considering the attacker $\mathcal{A}$ that only makes classical queries to H.

In (10) where the closeness between *Sign* and *Trans*, one with and one without reprogramming, is argued, we may substitute the advantage by

$$\Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{Sign,H}\right] - \Pr\left[1 \leftarrow \mathit{Vrfy}^H \circ \mathcal{A}^{Trans,H}\right] \leq 2q_S(q_H + q_S)/2^{\mathsf{saltlen}}\ .$$

Moreover, to control the event (11) of finding a preimage of at most $n \cdot q_S$ elements, there is a better classical bound as well:

$$\Pr\left[H_{\mathsf{init}}(m^*, r^*) \in S\right] \leq (q_H + 1) \cdot n \cdot q_S / 2^{2n}\ .$$

Putting things together, we obtain the classical security of HAWK as follows.

Theorem 2 (Classical Security of HAWK). *Let HAWK be as specified in Section 3, and let $\mathcal{A}^{Sign_{\mathbf{B}},H}(\mathbf{Q})$ be a chosen-message attack making at most q_S queries to $Sign_{\mathbf{B}}$ and at most q_H classical queries to H respectively. Then there exists an algorithm $\mathcal{B}^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}$ making q_S queries to solve one-more SVP, with running time $\mathsf{TIME}(\mathcal{B}) \approx \mathsf{TIME}(\mathcal{A}) + \mathsf{Overhead}(q_S, q_H)$ consisting of an additional overhead $\mathsf{Overhead}(q_S, q_H)$ of respectively simulating q_H, q_S queries to H and $Sim^{D^{\mathbf{B}}_{2\sigma_{\mathsf{sign}}}}$ (specified in Fig. 2), such that*

$$\begin{aligned} adv^{sUF\text{-}CMA}_{\mathcal{A}} \leq\ & adv^{omSVP}_{\mathcal{B}} + \frac{q_S \epsilon}{1-\epsilon} + 2q_S(q_H + q_S)/2^{\mathsf{saltlen}} \\ & + q_S\left(2^{-n} + (q_S - 1) \cdot n \cdot 2^{-2n}\right) \cdot \frac{1+\epsilon}{1-\epsilon} + (q_H + 1) \cdot n \cdot q_S / 2^{2n}\ . \end{aligned}$$

Acknowledgement

The authors thank Jelle Don and Eamonn W. Postlethwaite, Ludo N. Pulles for their useful discussions. Yu-Hsuan Huang is supported by the Dutch Research Agenda (NWA) project HAPKIDO (Project No. NWA.1215.18.002), which is financed by the Dutch Research Council (NWO).

References

BDF+11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology–ASIACRYPT 2011: 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings 17*, pages 41–69. Springer, 2011.

BHH+15. Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O'Hearn. Sphincs: Practical stateless hash-based signatures. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015*, pages 368–397, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

DFMS19. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 356–383, Cham, 2019. Springer International Publishing.

DKL+18. Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 238–268, 2018.

DPPW22. Léo Ducas, Eamonn W. Postlethwaite, Ludo N. Pulles, and Wessel van Woerden. Hawk: Module lip makes lattice signatures fast, compact and simple. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022*, pages 65–94, Cham, 2022. Springer Nature Switzerland.

DvW22. Léo Ducas and Wessel van Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 643–673, Cham, 2022. Springer International Publishing.

FHK+18. Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, Zhenfei Zhang, et al. Falcon: Fast-fourier lattice-based compact signatures over ntru. *Submission to the NIST's post-quantum cryptography standardization process*, 36(5), 2018.

GHHM21. Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the qrom. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 637–667, Cham, 2021. Springer International Publishing.

LZ19. Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019*, pages 326–355, Cham, 2019. Springer International Publishing.

Zha15. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. *International Journal of Quantum Information*, 13(04):1550014, 2015.

A More Proofs

Proof of Lemma 2. Without loss of generality, assume $\mathcal{A}$ makes exactly q_R queries to the reprogramming oracle Repro_b by doing additional dummy queries if otherwise. Define a sequence of hybrid games $\mathcal{G}_i$ that replaces the first i reprogramming query of $\mathcal{A}^{\mathsf{Repro}_1,H}$ to querying Repro_0, where by definition $\mathcal{G}_0$ and $\mathcal{G}_{q_R}$ run as $\mathcal{A}^{\mathsf{Repro}_1,H}$ and $\mathcal{A}^{\mathsf{Repro}_0,H}$ respectively.

It suffices to show the closeness $\mathcal{G}_i \approx \mathcal{G}_{i+1}$ for every $0 \le i < q_R$, where we refer to the only query that differs as the crucial query. For the sake of analysis we consider the random oracle H to be (perfectly) simulated via compressed oracle in a designated database register D, which, within the crucial query before $y := H(x)$ or $H(x) := y \leftarrow \mathcal{Y}$, is decompressed and meausred in the computational basis to obtain the oracle H to be used later.

Define $\mathcal{G}', \mathcal{G}''$ to respectively run as $\mathcal{G}_i, \mathcal{G}_{i+1}$ except additionally doing a binary measurement $\{M_0, M_1\}$ where $M_1 := \sum_{D(x)=\perp} |D\rangle\langle D|_{\mathsf{D}}$ after $x \leftarrow \mathcal{D}$ being sampled but before $y := H(x)$ or $H(x) := y \leftarrow \mathcal{Y}$, and abort if the outcome does not match M_1. $\mathcal{G}'$ and $\mathcal{G}''$ behaves identically because on non-abort, the database register D collapses into $|\perp\rangle_{\mathsf{D}(x)}$, for which the reprogramming $H(x) := y \leftarrow \mathcal{Y}$ do not affect the decompressed-and-measured distribution of $\mathsf{D}(x)$. The closeness of $\mathcal{G}' \approx \mathcal{G}_i$ and $\mathcal{G}'' \approx \mathcal{G}_{i+1}$ follows from the gentle-measurement lemma, together with the fact that there has been $\bar{q}_i + q_S$ queries of interaction with H prior to the crucial query, so $\Pr[\mathcal{G}' \text{ aborts}] = \Pr[\mathcal{G}'' \text{ aborts}] \le (q_H + q_R)\cdot\epsilon$. This concludes the proof, which can be summarized by the following chain of closeness

$$\mathcal{G}_i \overset{\sqrt{(q_H+q_R)\epsilon}}{\approx} \mathcal{G}' \overset{0}{\approx} \mathcal{G}'' \overset{\sqrt{(q_H+q_R)\epsilon}}{\approx} \mathcal{G}_{i+1} \ .$$

□