

Transitioning to a Quantum-Resistant Public Key Infrastructure

Nina Bindel¹, Udyani Herath², Matthew McKague², and Douglas Stebila³

¹ Technische Universität Darmstadt, Darmstadt, Germany

² Queensland University of Technology, Brisbane, Australia

³ McMaster University, Hamilton, Ontario, Canada

Abstract. To ensure uninterrupted cryptographic security, it is important to begin planning the transition to post-quantum cryptography. In addition to creating post-quantum primitives, we must also plan how to adapt the cryptographic infrastructure for the transition, especially in scenarios such as public key infrastructures (PKIs) with many participants. The use of hybrids—multiple algorithms in parallel—will likely play a role during the transition for two reasons: “hedging our bets” when the security of newer primitives is not yet certain but the security of older primitives is already in question; and to achieve security and functionality both in post-quantum-aware and in a backwards-compatible way with not-yet-upgraded software.

In this paper, we investigate the use of hybrid digital signature schemes. We consider several methods for combining signature schemes, and give conditions on when the resulting hybrid signature scheme is unforgeable. Additionally we address a new notion about the inability of an adversary to separate a hybrid signature into its components. For both unforgeability and non-separability, we give a novel security hierarchy based on how quantum the attack is. We then turn to three real-world standards involving digital signatures and PKI: certificates (X.509), secure channels (TLS), and email (S/MIME). We identify possible approaches to supporting hybrid signatures in these standards while retaining backwards compatibility, which we test in popular cryptographic libraries and implementations, noting especially the inability of some software to handle larger certificates.

1 Introduction

Since the initial advent of modern symmetric and public key cryptography in the 1970s, there have only been a handful of transitions from one widely deployed algorithm to another. These include: from DES and Triple-DES to AES; from MD5 and SHA-1 to the SHA-2 family; from RSA key transport and finite field Diffie–Hellman to elliptic curve Diffie–Hellman key exchange; and from RSA and DSA certificates to ECDSA certificates. Some of these transitions have gone well: AES is nearly ubiquitous today, and modern communication protocols predominantly use ECDH key exchange. Transitions involving public key infrastructure have a more mixed record: browser vendors and CAs have had

a long transition period from SHA-1 to SHA-2 in certificates, with repeated delay of deadlines; the transition to elliptic curve certificates has been even slower, and still today the vast majority of certificates issued for the web use RSA.

In the medium-term, we are likely to see another transition to *post-quantum* public key cryptography. Some aspects of the post-quantum transition will be straightforward: using post-quantum key exchange in protocols that support negotiation such as the Transport Layer has already been demonstrated [9,10], and can be adopted piecewise. Other migrations will be harder, especially when it is difficult for old and new configurations to operate simultaneously. A recent whitepaper [12] discusses some of these issues at a high level.

The transition to post-quantum cryptography is further complicated by the relative immaturity of some of the underlying mathematical assumptions in current candidates: because they have not been studied for very long, there is a higher risk that they might be insecure. This motivates *hybrid* operation, in which both a traditional algorithm and one or more post-quantum algorithms are used in parallel: as long as one of them remains unbroken, confidentiality or authenticity can be ensured. This leads us to three research questions:

1. What are the appropriate security properties for hybrid digital signatures?
2. How should we combine signature schemes to construct hybrid signatures?
3. How can hybrid signatures be realized in popular standards and software, ideally in a backwards-compatible way?

1. Security notions for hybrid digital signatures. The widely accepted security notion for digital signatures is *unforgeability under chosen message attack* (EUF-CMA): the adversary interacts with a signing oracle to obtain signatures on any desired messages, and then must output a forgery on a new message. Hybrid signatures should retain that property. Boneh and Zhandry [8] first studied security notions for digital signature schemes against quantum adversaries, and gave a quantum analogue of EUF-CMA in which a quantum adversary is able to interact with a quantum signing oracle and thereby obtain signatures on quantum states of its choosing (which may be in superposition).

As we transition to post-quantum digital signatures, Boneh and Zhandry's definition might be overly strong: for example, we might be using a signature scheme for the next five years, and during this period we are confident that no adversary has a quantum computer, and moreover we are definitely not signing anything in superposition; but later, the adversary may eventually be able to use a quantum computer. We describe several security notions depending on how quantum the adversary is. We use the notation X^yZ to denote the adversary's type with respect to three options:

- X : whether the adversary is classical ($X = C$) or quantum ($X = Q$) *during* the period in which it can interact with the signing oracle;
- y : whether the adversary can interact with the signing oracle classically ($y = c$) or quantumly ($y = q$); and
- Z : whether the adversary is classical ($Z = C$) or quantum ($Z = Q$) *after* the period in which it can interact with the signing oracle.

Table 1. Combiners for constructing hybrid signatures using schemes Σ_1 and Σ_2 ; $\max\{X^Y Z, U^V W\}$ denotes the stronger unforgeability notion with respect to the natural hierarchy of security notions.

Combiner	Combined signature $\sigma = (\sigma_1, \sigma_2)$	Unforgeability	Non-separability
<i>Single-message combiners</i>			
C_{\parallel}	$\sigma_1 \leftarrow \text{Sign}_1(m); \sigma_2 \leftarrow \text{Sign}_2(m)$	$\max\{X^Y Z, U^V W\}$	no
$C_{\text{weak-nest}}$	$\sigma_1 \leftarrow \text{Sign}_1(m); \sigma_2 \leftarrow \text{Sign}_2(\sigma_1)$	$X^Y Z$	$U^C W-2$
$C_{\text{str-nest}}$	$\sigma_1 \leftarrow \text{Sign}_1(m); \sigma_2 \leftarrow \text{Sign}_2((m, \sigma_1))$	$\max\{X^Y Z, U^V W\}$	$U^C W-2$
<i>Dual-message combiners</i>			
D_{nest}	$\sigma_1 \leftarrow \text{Sign}_1(m_1); \sigma_2 \leftarrow \text{Sign}_2((m_1, \sigma_1, m_2))$	$U^V W, X^Y Z^*$	$U^C W-2$

Unforgeability: If Σ_1 is $X^Y Z$ -eufcma and Σ_2 is $U^V W$ -eufcma, then $C(\Sigma_1, \Sigma_2)$ is \dots -eufcma.

Non-separability: If Σ_1 is $X^Y Z$ -eufcma and Σ_2 is $U^V W$ -eufcma, then $C(\Sigma_1, \Sigma_2)$ is \dots -nonsep.

* Unforgeability of $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ under $X^Y Z$ -eufcma-security of Σ_1 in a restricted sense.

These security notions form a natural hierarchy ($Q^Q \implies Q^C Q \implies C^C Q \implies C^C C$) with separations between each level of the hierarchy.

We describe a second security property specifically related to hybrid signatures, called *non-separability*: for a hybrid signature involving two (or more) signature schemes, is it possible for an adversary to separate the hybrid signature into a valid signature in any of the component signature schemes? This security property is interesting in the context of a transition. Suppose a signer issues hybrid signatures during a transition. Suppose further that there is a verifier who can understand both hybrid signatures and single-scheme signatures, but possibly acts upon them differently. The goal of non-separability is to prevent an attacker from taking a hybrid signature and turning it into something that the verifier accepts as coming from a single-scheme signature—thereby misrepresenting the signer’s original intention. Specifically, if $\Sigma' = C(\Sigma_1, \Sigma_2)$ is the hybrid signature scheme from using combiner C to combine signature schemes Σ_1 and Σ_2 , then we say Σ' is $X^Y Z$ - τ -nonsep if it is hard for an $X^Y Z$ -adversary to construct a valid Σ'_τ signature given access to a signing oracle for Σ' . Our notions are aided by a *recognizer* algorithm, which a verifier can apply to a signature to attempt to help distinguish separated hybrid signatures.

2. Signature combiners. Having laid out security properties for hybrid signatures, we proceed to investigate how to construct hybrid schemes using *combiners*. Table 1 shows the combiners we consider: concatenation and three forms of nested signatures. These particular combiners are motivated by two factors: that they are fairly natural constructions, and three of them arise in our applications.

3. Hybrid signatures in standards and software. Our goal is to provide a guide for how to transition to hybrid post-quantum digital signatures in various standards and software. We consider three standards: X.509 for certificates [14], the Transport Layer Security (TLS) protocol for secure channels [15], and Cryptographic Message Syntax (CMS) [20] as part of Secure/Multipurpose Internet Mail Extensions (S/MIME) [22] for secure email. For each, we ask:

- 3.a) How can hybrid / multiple signature schemes be used in the standard?

- 3.b) Is this approach backwards-compatible with old software?
- 3.c) Are there potential problems involving large public keys or signatures?

We identify promising techniques for hybrid X.509 certificates, and hybrid S/MIME signed messages; using multiple signature algorithms in TLS does not seem immediately possible, though one mechanism in the current draft of TLS 1.3 seems to allow multiple client authentications and a recent proposal could allow multiple server authentications. The software we tested had no problems with certificates or extensions up to ~ 10 kB (accommodating ideal lattice schemes), and some software up to ~ 80 kB (accommodating hash-based schemes), but none could handle the megabyte-plus size public keys of the largest lattice-based scheme; details appear in Section 5.

2 Signature schemes and unforgeability

Definition 1 (Signature scheme). A digital signature scheme Σ is a tuple $\Sigma = (\Sigma.\text{KeyGen}, \Sigma.\text{Sign}, \Sigma.\text{Verify})$ of algorithms:

- $\Sigma.\text{KeyGen}() \xrightarrow{s} (sk, vk)$: The probabilistic key generation algorithm that returns a secret or signing key sk and public or verification key $vk \in \mathcal{VK}_\Sigma$.
- $\Sigma.\text{Sign}(sk, m) \xrightarrow{s} \sigma$: The probabilistic signature generation algorithm which takes as input a signing key sk and a message $m \in \mathcal{M}_\Sigma$, and outputs a signature $\sigma \in \mathcal{S}_\Sigma$. The space of random coins is \mathcal{R}_Σ .
- $\Sigma.\text{Verify}(vk, m, \sigma) \rightarrow 0$ or 1 : The verification algorithm which takes as input a verification key vk , a message m , and a signature σ , and returns a bit $b \in \{0, 1\}$. If $b = 1$, we say that the algorithm accepts, otherwise we say that it rejects the signature σ for message m .

If a proof for Σ is being given in the random oracle model, we use \mathcal{H}_Σ to denote the space of functions from which the random hash function is randomly sampled. We say that Σ is ϵ -correct if, for every message m in the message space, we have that $\Pr[\text{Verify}(vk, m, \sigma) = 1 : (sk, vk) \xleftarrow{s} \text{KeyGen}(), \sigma \xleftarrow{s} \text{Sign}(sk, m)] \geq 1 - \epsilon$ where the probability is taken over the randomness of the probabilistic algorithms.

2.1 Unforgeability security definitions

The standard definition of security for signature schemes is *existential unforgeability under chosen message attack* (EUF-CMA). In the traditional formulation of unforgeability dating back to Goldwasser, Micali, and Rivest [18], the adversary can obtain q_S signatures via signing oracle queries, and must output one valid signature on a message not queried to the oracle. This cannot be directly quantized: if the adversary is allowed to query oracles in superposition, we cannot restrain the adversary’s forgery to be on a *new* message since the experiment cannot keep a copy of messages queried to the signing oracle for later checking.

An equivalent formulation in the classical setting demands the adversary output $q_S + 1$ valid signatures on distinct messages. Boneh and Zhandry [8] use

this formulation to give a quantum analogue of EUF-CMA.⁴ That notion involves a fully quantum adversary throughout so is quite strong. We envision a hierarchy of intermediate notions, distinguishing between whether the honest parties are signing classical or quantum messages (i.e., does the adversary have access to a classical or quantum signing oracle?) and whether the adversary is classical or quantum during the period it has access to the signing oracle (i.e., are we concerned about a quantum adversary only in the future, or also now?).

In our notation X^yZ , X denotes the type of adversary ($X = C$ for classical or Q for quantum) while the adversary is able to interact with the signing oracle; y denotes the type of access the adversary has to the signing oracle; and Z denotes the type of adversary after it no longer has access to its signing oracle. Combinatorially, there are $2^3 = 8$ possibilities, but some do not make sense, such as C^qZ or Q^yC . Figure 1 shows our unified definition for EUF-CMA parameterized for any of the four types of adversaries in the standard model, i.e., without access to a random (or hash) oracle. It follows the EUF-CMA formulation of Boneh and Zhandry [8] but separates out the adversary to be a two-stage adversary $(\mathcal{A}_1, \mathcal{A}_2)$, where \mathcal{A}_1 (of type X) interacts with either a signing oracle (of type y) and outputs an intermediate state st (of type X), which \mathcal{A}_2 (of type Z) then processes. The input to \mathcal{A}_1 and the output of \mathcal{A}_2 are always classical.

Figure 2 shows how the experiment is altered in the classical or quantum random oracle model: at the start of the experiment, a random function H is sampled uniformly from the space of all such functions \mathcal{H}_Σ . (In the classical setting, it is common to formulate the random oracle using lazy sampling, but such a formulation does not work in the quantum setting.) For simplicity, we assume the adversary has quantum random oracle access whenever it is quantum.

We define advantage as $\text{Adv}_\Sigma^{X^yZ\text{-eufcma}}(\mathcal{A}) = \Pr \left[\text{Expt}_\Sigma^{X^yZ\text{-eufcma}}(\mathcal{A}) = 1 \right]$.

2.2 Separations and implications

Our family of notions C^cC -, C^cQ -, Q^cQ -, Q^qQ -eufcma form a natural hierarchy. These implications induce an ordering on security notions, so we sometimes write $C^cC \leq C^cQ$ etc. Note, the stronger the security notion the smaller the advantage of an adversary \mathcal{A} breaking a signature scheme of corresponding security. For example, let the signature scheme Σ be C^cQ -secure. C^cQ is stronger than C^cC , i.e., $C^cQ \geq C^cC$. Hence, $\text{Adv}_\Sigma^{C^cQ\text{-eufcma}}(\mathcal{A}) \leq \text{Adv}_\Sigma^{C^cC\text{-eufcma}}(\mathcal{A})$. Similarly we use $\max\{\cdot, \cdot\}$ based on this ordering.

Due to space constraints, we defer details on the results to Appendix B. The implications are straightforward. Each of the separations $A \not\Rightarrow B$ follows from a common technique: from an A -secure scheme Σ , construct a (degenerate) A -secure scheme Σ' that is not B -secure, because the additional powers available to a B -adversary allow it to recover the secret signing key of Σ that was cleverly embedded somewhere in Σ' .

- $C^cC\text{-eufcma} \not\Rightarrow C^cQ\text{-eufcma}$: In the public key for Σ' , include a copy of the signing secret key encrypted using an RSA-based public key encryption

⁴ A brief overview of notation for quantum computing appears in Appendix A.

<u>$\text{Expt}_{\Sigma}^{\text{X}^{\text{Y}}\text{Z-eufcma}}(\mathcal{A}_1, \mathcal{A}_2)$:</u>	<u>Classical signing oracle $\mathcal{O}_S(m)$:</u>
1 $q_S \leftarrow 0$	8 $q_S \leftarrow q_S + 1$
2 $(sk, vk) \leftarrow_{\$} \Sigma.\text{KeyGen}()$	9 $\sigma \leftarrow_{\$} \Sigma.\text{Sign}(sk, m)$
3 $st \leftarrow_{\$} \mathcal{A}_1^{\mathcal{O}_S(\cdot)}(vk)$	10 Return σ to \mathcal{A}
4 $((m_1^*, \sigma_1^*), \dots, (m_{q_S+1}^*, \sigma_{q_S+1}^*)) \leftarrow_{\$} \mathcal{A}_2(st)$	
5 If $(\Sigma.\text{Verify}(vk, m_i^*, \sigma_i^*) = 1 \forall i \in [1, q_S + 1])$ $\wedge (m_i^* \neq m_j^* \forall i \neq j)$:	
6 Return 1	
7 Else return 0	
<u>Quantum signing oracle $\mathcal{O}_S(\sum_{m,t,z} \psi_{m,t,z} m, t, z\rangle)$:</u>	
11 $q_S \leftarrow q_S + 1$	
12 $r \leftarrow_{\$} \mathcal{R}_{\Sigma}$	
13 Return state $\sum_{m,t,z} \psi_{m,t,z} m, t \oplus \Sigma.\text{Sign}(sk, m; r), z\rangle$ to \mathcal{A}	

Fig. 1. Unified security experiment for $\text{X}^{\text{Y}}\text{Z-eufcma}$ in the standard model: existential unforgeability under chosen-message attack of a signature scheme Σ for a two-stage adversary \mathcal{A}_1 (of type X), \mathcal{A}_2 (of type Z) with signing oracle of type y; if $y = c$ then \mathcal{A}_1 has classical access to the signing oracle, otherwise quantum access.

scheme. Assuming breaking RSA is classically hard, the encrypted signing key is useless to a C^cC -adversary, but a C^cQ -adversary will be able to break the public key encryption, recover the signing key, and forge signatures.

- $\text{C}^c\text{Q-eufcma} \not\Rightarrow \text{Q}^c\text{Q-eufcma}$: In the public key for Σ' , include an RSA-encrypted random challenge string, and redefine the $\Sigma'.$ Sign so that, if the adversary queries the signing oracle on the random challenge string, the signing key is returned. Assuming breaking RSA is hard for a classical algorithm, a C^cQ -adversary will not be able to recover the challenge while it has access to the signing oracle, and thus cannot make use of the degeneracy to recover the signing key; a Q^cQ adversary can.
- $\text{Q}^c\text{Q-eufcma} \not\Rightarrow \text{Q}^q\text{Q-eufcma}$: Here we hide the secret using a query-complexity problem that can be solved with just a few queries by a quantum algorithm making queries in superposition, but takes exponential queries when asking classical queries. The specific problem we use is a variant of the *hidden linear structure* problem [4].

3 Separability of hybrid signatures

In Section 4, we will investigate combiners for constructing one signature scheme from two. Before looking at specific combiners, an interesting security property arises in general for combined signature schemes: is it possible for a signature in a combined signature scheme to be separated out into valid signatures for either of its individual component schemes?

Let C be a combiner, and let Σ_1, Σ_2 be signature schemes. Let $\Sigma' = C(\Sigma_1, \Sigma_2)$. The initial idea for the security notion for non-separability is based on the standard

<u>Expt$_{\Sigma}^{\mathcal{X}^Y\mathcal{Z}\text{-eufcma}}(\mathcal{A}_1, \mathcal{A}_2)$:</u>	<u>Classical random oracle $\mathcal{O}_H(x)$:</u>
0 $H \leftarrow_{\$} \mathcal{H}_{\Sigma}$	14 $q_H \leftarrow q_H + 1$
1 $q_H \leftarrow 0, q_S \leftarrow 0$	15 Return $H(x)$
2 $(sk, vk) \leftarrow_{\$} \Sigma.\text{KeyGen}()$	Quantum random oracle
3 $st \leftarrow_{\$} \mathcal{A}_1^{\mathcal{O}_S(\cdot), \mathcal{O}_H(\cdot)}(vk)$	<u>$\mathcal{O}_H(\sum_{x,t,z} \psi_{x,t,z} x, t, z\rangle)$:</u>
4 $((m_1^*, \sigma_1^*), \dots, (m_{q_S+1}^*, \sigma_{q_S+1}^*))$ $\leftarrow_{\$} \mathcal{A}_2^{\mathcal{O}_H(\cdot)}(st)$	16 $q_H \leftarrow q_H + 1$
5 // continues as in Figure 1	17 Return state $\sum_{x,t,z} \psi_{x,t,z} x, t \oplus H(x), z\rangle$

Fig. 2. $\mathcal{X}^Y\mathcal{Z}$ -eufcma experiment in the classical and quantum random oracle models; if $\mathcal{X} = \mathcal{C}$, then \mathcal{A}_1 has classical access to the random oracle, otherwise quantum access; similarly for \mathcal{Z} and \mathcal{A}_2 .

<u>Expt$_{C, \Sigma_1, \Sigma_2, C(\Sigma_1, \Sigma_2).R}^{\mathcal{X}^Y\mathcal{Z}\text{-}\tau\text{-nonsep}}(\mathcal{A}_1, \mathcal{A}_2)$:</u>	<u>\mathcal{O}_S:</u>
1 $q_S \leftarrow 0$	8 If $y = \mathcal{c}$, use classical signing oracle
2 $(sk', vk') \leftarrow_{\$} C(\Sigma_1, \Sigma_2).\text{KeyGen}()$	\mathcal{O}_S from Figure 1 for $C(\Sigma_1, \Sigma_2)$.
3 $st \leftarrow_{\$} \mathcal{A}_1^{\mathcal{O}_S(\cdot)}(vk')$	9 If $y = \mathcal{q}$, use quantum signing oracle
4 $(m^*, \sigma^*) \leftarrow_{\$} \mathcal{A}_2(st)$	\mathcal{O}_S from Figure 1 for $C(\Sigma_1, \Sigma_2)$.
5 If $(\Sigma_{\tau}.\text{Verify}((vk')_{\tau}, m^*, \sigma^*) = 1)$ $\wedge (C(\Sigma_1, \Sigma_2).R(m^*) = 0)$	
6 Return 1	
7 Else return 0	

Fig. 3. Unified security experiment for $\mathcal{X}^Y\mathcal{Z}\text{-}\tau\text{-nonsep}$: τ -non-separability of a combiner C with signature schemes Σ_1, Σ_2 with respect to a recognizer $C(\Sigma_1, \Sigma_2).R$ for a two-stage adversary \mathcal{A}_1 (of type \mathcal{X}), \mathcal{A}_2 (of type \mathcal{Z}) with signing oracle of type y . $(vk')_{\tau}$ denotes the projection (extraction) of the public key associated with scheme Σ_{τ} from the combined scheme's public key vk' , which we assume is possible.

EUFCMA experiment: given a signing oracle that produces signatures for the combined scheme Σ' , it should be hard for an adversary to produce a valid signature for Σ_1 (“1-non-separability”) or Σ_2 (“2-non-separability”).

However, this approach needs some refinement. In all of the combiners we consider in Section 4, the combined signature contains subcomponents which are valid in the underlying schemes. This makes it impossible to satisfy the naive version of non-separability. For example, suppose we have a signer issuing signatures from a combined signature scheme. Suppose further we have a verifier for one of the underlying signature schemes *who is also aware of the combined signature scheme*. Given signatures from the combined scheme, it should be hard to make a signature that is valid in one of the underlying signature schemes *and which is not recognized as being from the combined signature scheme*. In sum: separability is about downgrading.

Figure 3 shows the τ -non-separability security experiment, $\mathcal{X}^Y\mathcal{Z}\text{-}\tau\text{-nonsep}$, with $\tau \in \{1, 2\}$ for signature scheme Σ_1 or Σ_2 . It checks the ability of a two-stage $\mathcal{X}^Y\mathcal{Z}$ -adversary $(\mathcal{A}_1, \mathcal{A}_2)$ to create a valid Σ_{τ} signature; the adversary's pair has

to “fool” the recognizer algorithm $\Sigma.R$ which is supposed to recognize values used in a combined signature scheme Σ .

Formally, a recognizer related to a combined signature scheme $\Sigma = C(\Sigma_1, \Sigma_2)$ is a function $\Sigma.R$ that takes one input and outputs a single bit. For a signature scheme Σ_τ with message space $\mathcal{M}_{\Sigma_\tau}$, it may be that $\Sigma.R$ yields 1 on some elements of $\mathcal{M}_{\Sigma_\tau}$ and 0 on others: the purpose of R is to recognize whether certain inputs are associated with a *second* signature scheme.

Because the $X^Y Z$ - τ -nonsep experiment is parameterized by the recognizer algorithm R , each R gives rise to a different security notion. If R is vacuous, it can lead to a useless security notion. We can make statements of the form “If $\langle \text{some assumption on } \Sigma_1 \text{ and } \Sigma_2 \rangle$ and R is the algorithm \dots , then $C(\Sigma_1, \Sigma_2)$ is $X^Y Z$ -1-non-separable with respect to recognizer algorithm R .” However, a statement of the form “ C is not 1-non-separable” is more difficult: one must quantify over all (or some class of) recognizer algorithms R . We will say informally that “ C is not τ -non-separable” if the way that Σ_τ is used in C is to effectively sign the message without any modification, since the only recognizer that would recognize all signatures from C would cover the entire message space of Σ_τ .

One can view a recognizer algorithm as a generalization of the long-standing technique of “domain separation”. Starting from scratch, it might be preferable to build combiners that explicitly include domain separation in their construction, thereby eliminating the need for recognizer algorithms. Since our combiners are motivated by existing protocol constraints, we do not have that luxury.

4 Combiners

We now examine several methods of using two signature schemes Σ_1 and Σ_2 to produce hybrid signatures. For all our combiners, the key generation of the combined scheme will simply be the concatenation of the two schemes’ keys: $sk' \leftarrow (sk_1, sk_2)$; $vk' \leftarrow (vk_1, vk_2)$. The verification algorithm in each case is defined in the natural way. Proofs / proof sketches appear in Appendix C.

4.1 C_{\parallel} : Concatenation

This “combiner” is the trivial combiner, which just places independent signatures from the two schemes side-by-side:

- $C_{\parallel}(\Sigma_1, \Sigma_2).\text{Sign}(sk', m)$: $\sigma_1 \leftarrow_{\$} \Sigma_1.\text{Sign}(sk_1, m)$, $\sigma_2 \leftarrow_{\$} \Sigma_2.\text{Sign}(sk_2, m)$. Return $\sigma' \leftarrow \sigma_1 \parallel \sigma_2$.

Theorem 1 (Unforgeability of C_{\parallel}). *If either Σ_1 or Σ_2 is unforgeable in the classical (or quantum) random oracle model, then $\Sigma' = C_{\parallel}(\Sigma_1, \Sigma_2)$ is unforgeable in the classical (or quantum, respectively) random oracle model. More precisely, if Σ_1 is $X^Y Z$ -eufcma-secure or Σ_2 is $U^V W$ -eufcma-secure, then $\Sigma' = C_{\parallel}(\Sigma_1, \Sigma_2)$ is $\max\{X^Y Z, U^V W\}$ -eufcma-secure.*

Clearly, C_{\parallel} is neither 1-non-separable nor 2-non-separable: σ_1 is immediately a Σ_1 -signature for m , with no way of recognizing this as being different from typical Σ_1 signatures. Similarly for σ_2 .

4.2 $C_{\text{str-nest}}$: Strong nesting

For this combiner, the second signature scheme signs both the message and the signature from the first signature scheme:

- $C_{\text{str-nest}}(\Sigma_1, \Sigma_2).\text{Sign}(sk', m)$: $\sigma_1 \leftarrow^{\$} \Sigma_1.\text{Sign}(sk_1, m)$, $\sigma_2 \leftarrow^{\$} \Sigma_2.\text{Sign}(sk_2, (m, \sigma_1))$. Return $\sigma' \leftarrow (\sigma_1, \sigma_2)$.

Theorem 2 (Unforgeability of $C_{\text{str-nest}}$). *If either Σ_1 or Σ_2 is unforgeable in the classical (or quantum) random oracle model, then $\Sigma' = C_{\text{str-nest}}(\Sigma_1, \Sigma_2)$ is unforgeable in the classical (or quantum, respectively) random oracle model. More precisely, if Σ_1 is $X^Y Z$ -eufcma-secure or Σ_2 is $U^V W$ -eufcma-secure, then $\Sigma' = C_{\text{str-nest}}(\Sigma_1, \Sigma_2)$ is $\max\{X^Y Z, U^V W\}$ -eufcma-secure.*

$C_{\text{str-nest}}$ is not 1-non-separable: σ_1 is immediately a Σ_1 -signature for m , with no way of recognizing this as being different from typical Σ_1 signatures. However, since the inputs to Σ_2 in $C_{\text{str-nest}}$ have a particular form, we can recognize those and achieve 2-non-separability:

Theorem 3 (2-non-separability of $C_{\text{str-nest}}$). *If Σ_2 is $X^Y Z$ -eufcma-secure, then $\Sigma' = C_{\text{str-nest}}(\Sigma_1, \Sigma_2)$ is $X^c Z$ -2-nonsep with recognizer $R(m) = (m \in \{0, 1\}^* \times \mathcal{S}_{\Sigma_1})$.*

4.3 D_{nest} : Dual message combiner using nesting

Some of our applications in Section 5 require a combiner for two (possibly related) messages signed with two signature schemes. For example, in our X.509 certificates application, we generate one certificate signed with Σ_1 , then embed that certificate as an extension inside a second certificate signed with Σ_2 .

- $D_{\text{nest}}(\Sigma_1, \Sigma_2).\text{Sign}(sk', (m_1, m_2))$: $\sigma_1 \leftarrow^{\$} \Sigma_1.\text{Sign}(sk_1, m_1)$, $\sigma_2 \leftarrow^{\$} \Sigma_2.\text{Sign}(sk_2, (m_1, \sigma_1, m_2))$. Return $\sigma' \leftarrow (\sigma_1, \sigma_2)$.

This dual-message combiner is not designed to give unforgeability of *both* messages under *either* signature scheme, though it does preserve unforgeability of each message under its corresponding signature scheme, as well as give unforgeability of both messages under the outer signature scheme Σ_2 .

Theorem 4 (Unforgeability of D_{nest}). *If either Σ_1 or Σ_2 is unforgeable in the classical (or quantum) random oracle model, then $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is unforgeable (in a certain sense) in the classical (or quantum, respectively) random oracle model. More precisely, if Σ_1 is $X^Y Z$ -eufcma-secure, then the combined scheme $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is $X^Y Z$ -eufcma-secure with respect to its first message component only. If Σ_2 is $U^V W$ -eufcma-secure, then $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is $U^V W$ -eufcma-secure.*

D_{nest} is not 1-non-separable: σ_1 is immediately a Σ_1 -signature for m , with no way of recognizing this as being different from typical Σ_1 signatures. However, since the inputs to Σ_2 in D_{nest} have a particular form, we can recognize those and achieve 2-non-separability:

Theorem 5 (2-non-separability of D_{nest}). *If Σ_2 is $X^Y Z$ -eufcma-secure, then $\Sigma' = D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is $X^c Z$ -2-nonsep with respect to recognizer algorithm $R(m) = (m \in \{0, 1\}^* \times \mathcal{S}_{\Sigma_1} \times \{0, 1\}^*)$.*

5 Hybrid signatures in standards

We now examine three standards which make significant use of digital signatures—X.509 for certificates, TLS for secure channels, and S/MIME for secure email—to identify how hybrid signatures might be used in PKI standards, and evaluate backwards-compatibility of various approaches with existing software. Source code for generating the hybrid certificates and messages we used for testing, as well as scripts for running the tests, are available online at <https://www.douglas.stebila.ca/code/pq-pki-tests/>.

5.1 X.509v3 certificates

The X.509 standard version 3 [14] specifies a widely used format for public key certificates, as well as mechanisms for managing and revoking certificates.

The structure of an X.509v3 certificate is as follows. The body of the certificate (called a `tbsCertificate`) contains the name of the certificate authority as well as information about the subject, including the distinguished name of the subject, the subject’s public key (including an algorithm identifier), and optionally some extensions, each of which consists of an extension identifier, value, and a flag whether the extension is to be considered critical. (If a critical extension can not be recognized or processed, the system must reject it; a non-critical extension should be processed if it is recognized and may ignored if it is not.) The `tbsCertificate` is followed by CA’s signature over the `tbsCertificate` signed using the CA’s private key, along with an algorithm identifier.

Our goal will be to construct a hybrid certificate which somehow includes two public keys for the subject (e.g., one traditional and one post-quantum algorithm) and two CA signatures. Notably, the X.509v3 standard says that a certificate can contain exactly one `tbsCertificate`, which can contain exactly one subject public key, and all of this can be signed using exactly one CA signature. This makes creating backwards-compatible hybrid certificates challenging.

Approach 1: Dual certificates. The simplest approach is of course to create separate certificates: one for the traditional algorithm, and the other for the post-quantum algorithms. (This would be a dual-message analogue of the concatenation combiner C_{\parallel} from Section 4.1.) This approach leaves the task of conveying the “hybrid” certificate (actually, two certificates) to the application, which will suffice in some settings (e.g., in S/MIME and some TLS settings, see the following sections), but is unsatisfactory in others. (This approach and the next both require assigning additional object identifiers (OIDs) for each post-quantum algorithms, but this can be easily done.)

Approach 2: Second certificate in extension. Since X.509v3 does not provide any direct way of putting two public keys or two signatures in the same certificate, one option is to use the standard’s extension mechanism. Let c_1 be the certificate obtained by the CA signing `tbsCertificate` m_1 (containing subject public key vk_1) using signature scheme Σ_1 . Construct certificate c_2 by the CA signing `tbsCertificate` m_2 (containing subject public key vk_2 as well as (an

Table 2. Compatibility of hybrid X.509v3 certificates containing large extensions.

	Extension size (and corresponding example signature scheme)				
	1.5 KiB (RSA)	3.5 KiB (GLP [19])	9.0 KiB (BLISS [16])	43.0 KiB (SPHINCS [6])	1333.0 KiB (TESLA-416 [2])
GnuTLS 3.5.11	✓	✓	✓	✓	✓
Java SE 1.8.0_131	✓	✓	✓	✓	✓
mbedTLS 2.4.2	✓	✓	✓	✓	✓
NSS 3.29.1	✓	✓	✓	✓	✓
OpenSSL 1.0.2k	✓	✓	✓	✓	✓

encoding of) c_1 as an extension in m_2) using signature scheme Σ_2 . The extension containing c_1 would use a distinct extension identifier saying “this is an additional certificate” and would be marked as non-critical. This is an instantiation of the dual message nested combiner D_{nest} from Section 4.3. (Alternatively, the extension could contain a subset of fields, such as just the public key and CA’s signature, rather than a whole certificate.)

By marking the “additional certificate” extension as non-critical, existing software (not aware of the hybrid structure) *should* ignore the unrecognized extension and continue validating the certificate and using it in applications without change. Is this really the case—is this approach backwards-compatible with old software, and do large public keys or signatures cause problems?

Experimental evaluation of approach 2. We constructed hybrid certificates following approach 2. The “outside” certificate c_2 contains a 2048-bit RSA public key, and is signed by a CA using 2048-bit RSA key. The extension for embedding c_1 in c_2 is identified by a distinct and previously unused algorithm identifier (OID), and is marked as non-critical. Because post-quantum public keys and signatures vary substantially in size, we use a range of extension sizes to simulate the expected size of an embedded certificate for various post-quantum signature algorithms; the extension sizes we use are across the columns of Table 2, derived from public key and signature sizes summarized in Table 5 in the Appendix. For our purposes of evaluating backwards compatibility, it does not matter whether the extension actually *contains* a valid post-quantum certificate, just that it is the *size* of such a certificate. The hybrid certificates were created using a custom-written Java program using the BouncyCastle library.

Table 2 shows the results of using command-line certificate verification programs in various libraries; all libraries we tested were able to parse and verify X.509v3 certificates containing unrecognized extensions of all sizes we tested.

5.2 TLS

TLSv1.2 [15] is the currently standardized version and is widely deployed. Ciphersuites with digital signatures allow servers and (optionally) clients to authenticate each other by presenting their public key and signing certain messages. While the parties can negotiate which signature algorithms to use, which public key

Table 3. Compatibility of TLS connections using hybrid X.509v3 certificates containing large extensions.

	Extension size in KiB				
	1.5	3.5	9.0	43.0	1333.0
<i>Libraries</i> (library’s command-line client talking to library’s command-line server)					
GnuTLS 3.5.11	✓	✓	✓	✓	×
Java SE 1.8.0_131	✓	✓	✓	✓	✓
mbedtls 2.4.2	✓	✓	✓	×	×
NSS 3.29.1	✓	✓	✓	✓	×
OpenSSL 1.0.2k	✓	✓	✓	✓	×
<i>Web browsers</i> (talking to OpenSSL’s command-line server)					
Apple Safari 10.1 (12603.1.30.0.34)	✓	✓	✓	✓	✓
Google Chrome 58.0.3029.81	✓	✓	✓	✓	×
Microsoft Edge 38.14393.1066.0	✓	✓	✓	×	×
Microsoft IE 11.1066.14393.0	✓	✓	✓	×	×
Mozilla Firefox 53.0	✓	✓	✓	✓	×
Opera 44.0.2510.1218	✓	✓	✓	✓	×

formats to use, and which CAs to trust, once having done so they can each use only a single public key and signature algorithm to authenticate.

The current draft of TLSv1.3 [23] does not change how server authentication works. However, it has a “post-handshake authentication” mode for clients [23, §4.5.2], where clients can be requested to (further) authenticate using a certificate for a given algorithm. This would allow client authentication using two (or more) signature schemes. This is an example of the concatenation combiner $C_{||}$ from Section 4.1, since each client signature is over the same handshake context data structure. A proposal for “exported authenticators” [24] is currently before the TLS working group and would allow a similar approach for server authentication, although it envisions that this takes place out-of-band (e.g., at the application layer). Neither would require hybrid certificates as in Section 5.1.

TLS data structures allow certificates of size up to 2^{24} bytes = 16 MiB, which would accommodate even very large post-quantum algorithms. However, TLS record layer fragments can be at most 16 KiB; TLS messages can be split across multiple fragments, but this increases the risk of incompatibility with poorly implemented software and can be problematic with datagram transport (UDP).

Experimental evaluation of hybrid certificates in TLS. Table 3 shows the results of testing the compatibility of a variety of TLS libraries and web browsers when using the hybrid certificates from Approach 2 of Section 5.1.

In the top half of the table, we test whether popular TLS libraries can be used to establish a TLS 1.2 connection using an RSA certificate with an extension of the given size. In each case, the experiment is carried out between that library’s own TLS server and TLS client command-line programs (in the case of Java, we wrote a simple HTTPS server and client using built-in libraries). Only Java

completes connections with extensions the size of a TESLA-416 [2] certificate (1.3 MiB), and mbedTLS cannot handle certificates with extensions the size of a SPHINCS [6] certificate (43 KiB). (For GnuTLS and OpenSSL, we found they could handle an 80 KiB extension but not a 90 KiB extension).

In the bottom half of the table, we test whether popular web browsers can be used to establish a TLS 1.2 connection to a TLS server run using the OpenSSL command-line `s_server` program using an RSA certificate with an extension of the given size. Microsoft browsers on Windows 10 cannot handle SPHINCS-sized extensions, and no browser except Safari could handle TESLA-416-sized extensions (1.3 MiB). (Curiously, Safari was able to handle a 1.3 MiB extension with an OpenSSL command-line server despite OpenSSL’s own command-line client not being able to handle it.)

5.3 CMS and S/MIME

Cryptographic Message Syntax (CMS) [20] is the main cryptographic component of S/MIME [22], which enables public key encryption and digital signatures for email. In an S/MIME signed email, a header is used to specify the algorithms used, and then the body of the email is divided into chunks: one chunk is the body to be signed, and the other is a Base-64 encoding of a CMS `SignedData` object. The `SignedData` object contains several fields, including a set of certificates and a set of `SignerInfo` objects. Each `SignerInfo` object contains a signer identifier, algorithm identifier, signature, and optional signed and unsigned attributes.

Approach 1: Parallel SignerInfos. To construct a standards-compliant hybrid signature in S/MIME, one could put the certificate for each algorithm in the `SignedData` object’s set of certificates (with no need for hybrid certificates from Section 5.1), and then include `SignerInfo` objects for the signature from each algorithm. This is an example of the concatenation combiner C_{\parallel} from Section 4.1.

Approach 2: Nested signature in SignerInfo attributes. For an alternative and still standards-compliant approach, we could use the optional attributes in the `SignerInfo` object to embed a second signature. We need to convey the certificate for the second algorithm, as well as the signature (using the second algorithm) for the message. There are several options based on the standards:

- 2.a) Put a second certificate in the set of certificates, and put a second `SignerInfo` in an attribute of the first `SignerInfo`.
- 2.b) Put a hybrid certificate in the set of certificates, and put a second `SignerInfo` in an attribute of the first `SignerInfo`.
- 2.c) Put a second `SignedData` in an attribute of the first `SignerInfo`.

These approaches require defining a new attribute type, but this is easily done. The CMS standard indicates that verifiers can accept signatures with unrecognized attributes, so this approach results in backwards-compatible signatures that should be accepted by existing software.)

If the extra data is put in the *signed* attribute of the first `SignerInfo`, then we are using the strong nesting combiner $C_{\text{str-nest}}$ from Section 4.2; if the extra

Table 4. Compatibility of hybrid S/MIME approaches.

	Approach								
	0	1	2.a	2.b (attribute size in KiB)					2.c
				1.5	3.5	9.0	43.0	1333.0	
Apple Mail 10.2 (3259)	✓	✓	✓	✓	✓	✓	✓	✓	✓
BouncyCastle 1.56 with Java SE 1.8.0_131	✓	–	✓	✓	✓	✓	✓	✓	✓
Microsoft Outlook 2016 16.0.7870.2031	✓	×	✓	✓	✓	✓	✓	✓	✓
Mozilla Thunderbird 45.7.1	✓	×	✓	✓	✓	✓	✓	×	✓
OpenSSL 1.0.2k	✓	×	✓	✓	✓	✓	✓	✓	✓

Approach 0: both RSA-2048.

Approach 1: one RSA-2048, one with unknown algorithm of similar key and signature size; unable to test Approach 1 on BouncyCastle.

Approach 2.a, 2.c: both RSA-2048.

Approach 2.b: outer RSA-2048, inner random binary extension of given size.

data is put in the *unsigned* attribute of the first **SignerInfo**, then we are using the concatenation combiner C_{\parallel} from Section 4.1.

Experimental evaluation of CMS and S/MIME approaches. We tested five S/MIME libraries/applications for acceptance of S/MIME messages from each approach above. The configurations and results appear in Table 4.

Regarding approach 1, the S/MIME and CMS standards appear to be silent on how to validate multiple **SignerInfo** objects: should a signed message be considered valid if *any* of the **SignerInfo** objects is valid, or only if *all* of them are? Apple Mail accepted in this case, whereas the three others we were able to test rejected, so approach 1 is not fully backwards-compatible. In principle all the tested libraries support approaches 2.a–2.c. We only tested multiple attribute sizes in one of these three approaches (2.b), but the results should generalize to 2.a and 2.c. Only Thunderbird struggled with very large attributes.

5.4 Discussion

Summarizing our experimental observations, backwards compatibility is most easily maintained when post-quantum objects can be placed as non-critical extensions (attributes in the S/MIME case) of pre-quantum objects. These constructions end up leading to one of our nested combiners, either D_{nest} for X.509 certificate extensions or $C_{\text{str-nest}}$ for S/MIME and CMS. Both these combiners offer unforgeability under the assumption that either scheme is unforgeable. For non-separability, the pre-quantum algorithm would be the “outside” signature Σ_2 in both D_{nest} and $C_{\text{str-nest}}$, and so we get 2-non-separability under unforgeability of the pre-quantum scheme, not the post-quantum scheme. Extensions up to 43.0 KiB were mostly (but not entirely) handled successfully, which covers many but not all post-quantum schemes; the largest schemes such as TESLA-416 would be more problematic to use in hybrid modes with existing software.

Acknowledgements

NB acknowledges support by the German Research Foundation (DFG) as part of project P1 within the CRC 1119 CROSSING. DS acknowledges support from Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grant RGPIN-2016-05146.

References

1. Akleylek, S., Bindel, N., Buchmann, J.A., Krämer, J., Marson, G.A.: An efficient lattice-based signature scheme with provably secure instantiation. In: AFRICACRYPT 16. LNCS, vol. 9646, pp. 44–60. Springer (Apr 2016)
2. Alkim, E., Bindel, N., Buchmann, J., Dagdelen, Ö.: TESLA: Tightly-secure efficient signatures from standard lattices. Cryptology ePrint Archive, Report 2015/755 (2015)
3. Barreto, P.S.L.M., Longa, P., Naehrig, M., Ricardini, J.E., Zanon, G.: Sharper ring-LWE signatures. Cryptology ePrint Archive, Report 2016/1026 (2016)
4. de Beaudrap, N., Cleve, R., Watrous, J.: Sharp quantum versus classical query complexity separations. *Algorithmica* 34(4), 449–461 (2002)
5. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: EUROCRYPT’94. LNCS, vol. 950, pp. 92–111. Springer (May 1995)
6. Bernstein, D.J., Hopwood, D., Hülsing, A., Lange, T., Niederhagen, R., Papachristodoulou, L., Schneider, M., Schwabe, P., Wilcox-O’Hearn, Z.: SPHINCS: Practical stateless hash-based signatures. In: EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 368–397. Springer (Apr 2015)
7. Bindel, N., Herath, U., McKague, M., Stebila, D.: Transitioning to a quantum-resistant public key infrastructure (full version). Cryptology ePrint Archive (April 2017)
8. Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 361–379. Springer (Aug 2013)
9. Bos, J.W., Costello, C., Naehrig, M., Stebila, D.: Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In: 2015 IEEE Symposium on Security and Privacy. pp. 553–570. IEEE Computer Society Press (May 2015)
10. Braithwaite, M.: Google Security Blog: Experimenting with post-quantum cryptography (Jul 2016), <https://security.googleblog.com/2016/07/experimenting-with-post-quantum.html>
11. Buchmann, J., Dahmen, E., Hülsing, A.: XMSS - a practical forward secure signature scheme based on minimal security assumptions. In: PQCrypto 2011. LNCS, vol. 7071, pp. 117–129. Springer (2011)
12. Campagna, M., et al.: Quantum safe cryptography and security: An introduction, benefits, enablers and challengers. Tech. rep., ETSI (European Telecommunications Standards Institute) (June 2015), <http://www.etsi.org/images/files/ETSIWhitePapers/QuantumSafeWhitepaper.pdf>
13. Chen, A.I.T., Chen, M.S., Chen, T.R., Cheng, C.M., Ding, J., Kuo, E.L.H., Lee, F.Y.S., Yang, B.Y.: SSE implementation of multivariate PKCs on modern x86 CPUs. In: CHES 2009. LNCS, vol. 5747, pp. 33–48. Springer (Sep 2009)

14. Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., Polk, W.: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280 (May 2008)
15. Dierks, T., Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (Aug 2008)
16. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal Gaussians. In: CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 40–56. Springer (Aug 2013)
17. Fischlin, R., Schnorr, C.P.: Stronger security proofs for RSA and Rabin bits. *J. Cryptology* 13(2), 221–244 (2000)
18. Goldwasser, S., Micali, S., Rivest, R.L.: A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Computing* 17(2), 281–308 (Apr 1988)
19. Güneysu, T., Lyubashevsky, V., Pöppelmann, T.: Practical lattice-based cryptography: A signature scheme for embedded systems. In: CHES 2012. LNCS, vol. 7428, pp. 530–547. Springer (Sep 2012)
20. Housley, R.: Cryptographic Message Syntax (CMS). RFC 5652 (Sep 2009)
21. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press (2000)
22. Ramsdell, B., Turner, S.: Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification. RFC 5751 (Jan 2010)
23. Rescorla, E.: The Transport Layer Security (TLS) protocol version 1.3, draft 19 (March 2017), <https://tools.ietf.org/html/draft-ietf-tls-tls13-19>
24. Sullivan, N.: Exported authenticators in TLS, draft 01 (March 2017), <https://tools.ietf.org/html/draft-sullivan-tls-exported-authenticator-01>

A A brief review of quantum computation

A full explanation of quantum computation is beyond the scope of this paper; see a standard text such as Nielsen and Chuang [21]. We can rely on a subset of quantum computation knowledge.

A quantum system is a complex Hilbert space \mathcal{H} with an inner product. Vectors in \mathcal{H} are typically denoted using “ket” notation, such as $|x\rangle$, and the complex conjugate transpose of $|y\rangle$ is denoted by $\langle y|$, so that their inner product of $|x\rangle$ and $|y\rangle$ is given by $\langle y|x\rangle$. A quantum state is a vector in \mathcal{H} of norm 1. For two quantum systems \mathcal{H}_1 and \mathcal{H}_2 , the joint quantum system is given by the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$; for two states $|x_1\rangle \in \mathcal{H}_1$ and $|x_2\rangle \in \mathcal{H}_2$, the joint state is denoted by $|x_1\rangle|x_2\rangle$, or more compactly as $|x_1, x_2\rangle$.

Some quantum states can be represented as superpositions of other quantum states, such as $|x\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$. More generally, if $\{|x\rangle\}_x$ is a basis for \mathcal{H} , then we can write any superposition in the form $|y\rangle = \sum_x \psi_x |x\rangle$ where ψ_x are complex numbers such that $|y\rangle$ has norm 1.

Quantum operations on \mathcal{H} can be represented by unitary transformations \mathbf{U} . A side effect of the fact that quantum operations are unitary transformations is that quantum computation (prior to measurement) is reversible, imposing some constraints on how we quantize classical computations. In particular, suppose we want to quantize a classical algorithm A which takes an input $x \in \{0, 1\}^a$ and gives an output $y \in \{0, 1\}^b$. First, we would imagine the classical reversible mapping

$\{0, 1\}^a \times \{0, 1\}^b \rightarrow \{0, 1\}^a \times \{0, 1\}^b : (x, t) \mapsto (x, t \oplus A(x))$. Then we construct the corresponding unitary transformation \mathbf{A} which acts linearly on superpositions of such states: $\mathbf{A} : \sum_{x,t} \psi_{x,t} |x, t\rangle \mapsto \sum_{x,t} \psi_{x,t} |x, t \oplus A(x)\rangle$. For full generality, we may allow a workspace register alongside the input and output registers, and thus we in fact use $\mathbf{A} : \sum_{x,t,z} \psi_{x,t,z} |x, t, z\rangle \mapsto \sum_{x,t,z} \psi_{x,t,z} |x, t \oplus A(x), z\rangle$.

B Unforgeability separations and implications

Theorem 6 ($\text{Q}^{\text{qQ}} \implies \text{Q}^{\text{cQ}} \implies \text{C}^{\text{cQ}} \implies \text{C}^{\text{cC}}$). *If Σ is a Q^{qQ} -eufcma-secure signature scheme, then Σ is also Q^{cQ} -eufcma-secure. If Σ is a Q^{cQ} -eufcma-secure signature scheme, then Σ is also C^{cQ} -eufcma-secure. If Σ is a C^{cQ} -eufcma-secure signature scheme, then Σ is also C^{cC} -eufcma-secure.*

Theorem 7 ($\text{C}^{\text{cC}} \not\implies \text{C}^{\text{cQ}}$). *If the RSA problem is hard for classical computers and there exists a signature scheme Σ that is C^{cC} -eufcma-secure, then there exists a signature scheme Σ' that is C^{cC} -eufcma-secure but not C^{cQ} -eufcma-secure.*

Proof. Let Π be a public key encryption scheme that is IND-CPA-secure against classical adversaries and whose security relies on the hardness of the RSA problem, e.g., [17] or OAEP [5]. However, a quantum adversary could use Shor's algorithm to factor the modulus and decrypt ciphertexts encrypted using Π . We construct a scheme Σ' that is based on Σ , but the public key of Σ' includes a Π -encrypted copy of the Σ secret key:

- $\Sigma'.\text{KeyGen}()$: $(sk, vk) \leftarrow_{\$} \Sigma.\text{KeyGen}()$. $(dk, ek) \leftarrow_{\$} \Pi.\text{KeyGen}()$.
 $c \leftarrow_{\$} \Pi.\text{Enc}(ek, sk)$. $vk' \leftarrow (vk, ek, c)$. Return (sk, vk') .
- $\Sigma'.\text{Sign}(sk, m)$: Return $\Sigma.\text{Sign}(sk, m)$.
- $\Sigma'.\text{Verify}(vk' = (vk, ek, c), m, \sigma)$: Return $\Sigma.\text{Verify}(vk, m, \sigma)$.

The theorem then follows as a consequence of the following two claims, the proofs of which are immediate. \square

Claim. If Π is IND-CPA-secure against a classical adversary and Σ is C^{cC} -eufcma-secure, then Σ' is C^{cC} -eufcma-secure.

Claim. If there exists an efficient quantum adversary \mathcal{A} against the message recovery of Π , then Σ' is not C^{cQ} -eufcma-secure.

Theorem 8 ($\text{C}^{\text{cQ}} \not\implies \text{Q}^{\text{cQ}}$). *If the RSA problem is hard for classical computers and there exists a signature scheme Σ that is C^{cQ} -eufcma-secure, then there exists a signature scheme Σ' that is C^{cQ} -eufcma-secure but not Q^{cQ} -eufcma-secure.*

Since the basic idea for the proof of Theorem 8 is similar to that of Theorem 7, and due to space constraints, we leave details to the full version [7]. Briefly, the idea of the construction of the scheme for the separation is as follows. Here, we put an encrypted *random challenge* in the public verification key, and if the adversary asks for that challenge to be signed, we have the signing oracle return the signing key. Intuitively, only an adversary that can break the challenge while it has access to the signing oracle (i.e., a quantum stage-1 adversary) can solve the challenge. The scheme Σ' is shown below.

- $\Sigma'.\text{KeyGen}(): (sk, vk) \leftarrow_{\$} \Sigma.\text{KeyGen}(). (dk, ek) \leftarrow_{\$} \Pi.\text{KeyGen}(). s^* \leftarrow_{\$} \{0, 1\}^{256}.$
 $ch \leftarrow \Pi.\text{Enc}(ek, s^*). vk' \leftarrow (vk, ek, ch). sk' \leftarrow (sk, s^*). \text{Return } (sk', vk').$
- $\Sigma'.\text{Sign}(sk' = (sk, s^*), m):$ If $m = s^*$, return sk . Else, return $\Sigma.\text{Sign}(sk, m)$.
- $\Sigma'.\text{Verify}(vk' = (vk, ek, ch), m, \sigma):$ Return $\Sigma.\text{Verify}(vk, m, \sigma)$.

Theorem 9 ($\text{Q}^c\text{Q} \not\Rightarrow \text{Q}^q\text{Q}$). *Assuming there exists a quantum-secure pseudo-random family of permutations, and a signature scheme Σ that is Q^cQ -eufcma-secure, then there exists a signature scheme Σ' that is Q^cQ -eufcma-secure but not Q^qQ -eufcma-secure.*

Similar to Theorem 8, we will construct a signature scheme where the secret key is hidden behind a problem which is hard for some adversaries and easy for others. Here the hidden problem will be on oracle problem where a small number of queries suffices to retrieve a secret string when the oracle is queried in superposition, but a large number of queries is required if the oracle is queried classically. We will use the *hidden linear structure* problem [4].

Definition 2 ([4]). *The hidden linear structure problem is as follows: given oracle access to $\mathcal{B}_{s,\pi}(x,y) = (x, \pi(y \oplus sx))$, where $x, y, s \in GF(2^n)$ and $\pi \in \text{Perm}(\{0, 1\}^n)$ with s and π chosen uniformly at random, determine s . (Here, $\text{Perm}(S)$ denotes the set of all permutations on a set S .)*

The hidden linear structure problem requires 2^b classical queries to solve with probability 2^{2b-n+1} (i.e. $O(2^{n/2})$ queries to solve with a constant probability), and one query to solve with quantum queries [4]. Unfortunately, describing π requires an exponential number of bits in n , but we can replace the random permutation π with a family of quantum-safe pseudo random permutation with a short key. This results in an oracle with a short description. Supposing that the PRP is indistinguishable from a random permutation in time $c_{\mathcal{P}}$ except with advantage $p_{\mathcal{P}}$, the resulting restricted oracle problem is indistinguishable from the hidden linear structure problem except with advantage $p_{\mathcal{P}}$. From now on we assume that π is implemented by a PRP.

Our construction starts with a Q^cQ -eufcma-secure signature scheme Σ . For our purposes, we will need $\Sigma.\text{Sign}$ to be deterministic. That is, for a particular message and signing key the signature should always be the same. If this is not the case, then we can use standard techniques to make it so, for example by providing randomness through a quantum-secure PRF applied to the signing key and the message. Let us suppose that it takes at least time c_{Σ} for an adversary to win the Q^cQ -eufcma security game with probability at least p_{Σ} .

We will need to address several parts of messages for signing. For a message m we will define $m.x, m.y, m.z$ to be bits 1 to 256, bits 257 to 512, and bits 513 to 768 of m , respectively. In particular, m must be at least 768 bits long. Bits beyond 768 will play no special role in the signing algorithm, but remain part of the message. Also let $\delta_{a,b}$ be the Kronecker delta, which is 1 when $a = b$ and 0 otherwise.

We now define our signature scheme Σ' as follows:

- $\Sigma'.\text{KeyGen}()$: $(sk, vk) \leftarrow_s \Sigma.\text{KeyGen}()$. $s \leftarrow_s \{0, 1\}^{256}$. $t \leftarrow_s \{0, 1\}^{256}$. $vk' \leftarrow (vk)$. $sk' \leftarrow (sk, s, t)$. Return (sk', vk') .
- $\Sigma'.\text{Sign}(sk', m)$: Return $(\Sigma.\text{Sign}(sk, m), \mathcal{B}_{s,t}(m.x, m.y), sk \cdot \delta_{s,m.z})$.
- $\Sigma'.\text{Verify}(vk', m, (\sigma, u, v, w))$: If $\Sigma.\text{Verify}(vk, m, \sigma)$ accepts, $(u, v) = \mathcal{B}_{s,t}(m.x, m.y)$ and $w = sk \cdot \delta_{m.z,s}$ then accept, otherwise reject.

Since we are interested in the case of quantum access, we define the quantum version of the signing oracle by $U_{\Sigma',sk}$, which has the action

$$U_{\Sigma',sk} |m, a, b, c, d\rangle = |m, a \oplus \sigma, b \oplus u, c \oplus v, d \oplus w\rangle$$

where $\sigma = \Sigma.\text{Sign}(sk, m)$, $(u, v) = \mathcal{B}_{s,t}(m.x, m.y)$, and $w = sk \cdot \delta_{s,m.z}$. Note that $U_{\Sigma',sk}$ is its own inverse.

Lemma 1. *Suppose that, with classical queries, at least $c_{\mathcal{B}}$ queries to $\mathcal{B}_{s,t}$ are required to determine s with probability $p_{\mathcal{B}}$, and that it takes at least time c_{Σ} for an adversary to win the $\text{Q}^c\text{Q-eufcma}$ security game for Σ with probability at least p_{Σ} . If a (possibly quantum) adversary \mathcal{A} with classical access to a $\Sigma'.\text{Sign}$ oracle and vk runs for time $c < \min\{c_{\mathcal{B}}, c_{\Sigma}\}$, then \mathcal{A} wins the $\text{Q}^c\text{Q-eufcma}$ security game for Σ' with probability at most $p \leq p_{\mathcal{B}} + p_{\Sigma} + 2^{-256}c$.*

The lemma can be proven by noting $\mathcal{B}_{s,t}$ and Σ are not related, so we can basically add the probabilities of determining s through $\mathcal{B}_{s,t}$, producing valid signatures without s , and guessing s directly.

Lemma 2. *Suppose $\Sigma.\text{Sign}$ is deterministic. If, given quantum query access to $\mathcal{B}_{s,t}$ it is possible to recover s with 1 query, then 3 quantum queries to $U_{\Sigma',sk}$ suffice to efficiently generate any polynomial number of valid signatures for Σ' .*

The basic mechanism here is to use a standard technique in quantum computing called uncomputing to construct a quantum oracle for $\mathcal{B}_{s,t}(x, y)$ out of two calls to $U_{\Sigma',sk}$. Then it is possible to determine s and recover sk with one more call to $U_{\Sigma',sk}$.

We are now in a position to prove Theorem 9.

Proof (Proof of Theorem 9). We use Σ' as defined earlier, with $\mathcal{B}_{s,t}$ being the oracle for a quantum safe hidden linear structure problem, which exists by the existence of \mathcal{P} . By Lemma 2, Σ' is not $\text{Q}^q\text{Q-eufcma}$ -secure since a quantum adversary allowed quantum oracle access to $\Sigma'.\text{Sign}$ can efficiently generate a polynomial number of signatures using a constant number of oracle queries.

Now suppose we have a quantum adversary \mathcal{A} which has classical oracle access to $\Sigma'.\text{Sign}$ and runs in time $2^b < \max\{2^{n/2-2}, c_{\Sigma}\}$. \mathcal{A} obtains s through classical oracle access to \mathcal{B} with probability at most $2^{2b-n+1} + p_{\mathcal{P}}$. Then we can set $p_{\mathcal{B}} = 2^{2b-n+1} + p_{\mathcal{P}}$ and apply Lemma 1 to find that \mathcal{A} breaks unforgeability of Σ' with probability at most $p_{\Sigma} + 2^{2b-n+1} + \delta + 2^{b-256}$. If \mathcal{A} runs in polynomial time, then $b \in O(\log(\text{poly}(n)))$ and hence Σ' is $\text{Q}^c\text{Q-eufcma}$ -secure. \square

C Proofs for combiners

C.1 C_{\parallel} : Concatenation

Proof (Proof of Theorem 1 – unforgeability of C_{\parallel}). Suppose \mathcal{A} is an R^{ST} -adversary that finds a forgery in $\Sigma' = C_{\parallel}(\Sigma_1, \Sigma_2)$ — in other words, it outputs $q_S + 1$ valid signatures under Σ' on distinct messages. We can construct an R^{ST} algorithm \mathcal{B}_1 that finds a forgery in Σ_1 . \mathcal{B}_1 interacts with an R^{ST} challenger for Σ_1 which provides a public key vk_1 . \mathcal{B}_1 generates a key pair $(sk_2, vk_2) \leftarrow \Sigma_2.\text{KeyGen}()$ and sets the public key for Σ' to be (vk_1, vk_2) . When \mathcal{A} asks for $\sum_{m,t,z} \psi_{m,t,z} |m, t, z\rangle$ to be signed using Σ' , we treat t as consisting of two registers $t_1 \| t_2$, \mathcal{B}_1 proceeds by passing the m , t_1 , and z registers to its signing oracle for Σ_1 , then runs the quantum signing operation from Figure 1 for $\Sigma_2.\text{Sign}$ on the m , t_2 , and z registers. There is a one-to-one correspondence between \mathcal{A} 's queries to its signing oracle and \mathcal{B}_1 's queries to its signing oracle.

If Σ_1 is proven to be secure in the random oracle (rather than standard) model, then this proof of $C_{\parallel}(\Sigma_1, \Sigma_2)$ also proceeds in the random oracle model: \mathcal{B}_1 relays \mathcal{A} 's hash oracle queries directly to its oracle, giving a one-to-one correspondence between \mathcal{A} 's queries to its hash oracle and \mathcal{B}_1 's queries to its hash oracle. This holds in either the classical or quantum random oracle model.

If \mathcal{A} wins the $\text{R}^{\text{ST-eufcma}}$ game, then it has returned $q_S + 1$ valid signatures $\sigma'_i = (\sigma'_{i,1}, \sigma'_{i,2})$ on distinct messages m_i such that $\Sigma_1.\text{Verify}(vk_1, m_i, \sigma'_{i,1}) = 1$ and $\Sigma_2.\text{Verify}(vk_2, m_i, \sigma'_{i,2}) = 1$. \mathcal{B}_1 can extract from this $q_S + 1$ valid signatures under Σ_1 on distinct messages. Thus, $\text{Adv}_{\Sigma'}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{A}) \leq \text{Adv}_{\Sigma_1}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_1)$. Similarly it holds for Σ_2 : $\text{Adv}_{\Sigma'}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{A}) \leq \text{Adv}_{\Sigma_2}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_2)$.

It follows that $\text{Adv}_{\Sigma'}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{A}) \leq \min\{\text{Adv}_{\Sigma_1}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_1), \text{Adv}_{\Sigma_2}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_2)\}$. Thus, if either $\text{Adv}_{\Sigma_1}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_1)$ or $\text{Adv}_{\Sigma_2}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{B}_2)$ is small, then so too is $\text{Adv}_{\Sigma'}^{\text{R}^{\text{ST-eufcma}}}(\mathcal{A})$.

C.2 $C_{\text{str-nest}}$: Strong nesting

Proof (Proof of Theorem 2 – unforgeability of $C_{\text{str-nest}}$). This proof follows the same approach as the proof of unforgeability for C_{\parallel} (Theorem 1). Details appear in the full version [7]. \square

Proof (Proof sketch of Theorem 3 – 2-non-separability of $C_{\text{str-nest}}$). We can construct a reduction \mathcal{B}_2 which is an $\text{X}^{\text{cZ-eufcma}}$ adversary for Σ_2 . \mathcal{B}_2 generates a keypair (vk_1, sk_1) for Σ_1 , and interacts with an $\text{X}^{\text{cZ-eufcma}}$ challenge for Σ_2 . When \mathcal{A} classically queries its signing oracle to obtain a signature under Σ' of m_i , \mathcal{B}_2 signs m_i with Σ_1 to obtain $\sigma_{i,1}$. Afterwards, \mathcal{B}_2 passes $(m, \sigma_{i,1})$ to its Σ_2 signing oracle and returns the resulting $\sigma_{i,2}$ to \mathcal{A} . Eventually, \mathcal{A} returns (μ^*, σ^*) such that $\Sigma_2.\text{Verify}(vk_2, \mu^*, \sigma^*) = 1$ but $\Sigma'.R(\mu^*) = 0$, i.e., $\mu^* \notin \{0, 1\}^* \times \mathcal{S}_{\Sigma_1}$. This means in particular that $\mu^* \neq (m_i, \sigma_{1,i})$ for all i . Moreover, all the $(m_i, \sigma_{1,i})$ are distinct, since all m_i are distinct. This means we have $q_S + 1$ valid message-signature pairs under Σ_2 , yielding a successful forgery for the $\text{X}^{\text{cZ-eufcma}}$ experiment for Σ_2 . Thus, $\Pr[S_1] \leq \text{Adv}_{\Sigma_2}^{\text{X}^{\text{cZ-eufcma}}}(\mathcal{B}_2)$. \square

C.3 D_{nest} : Dual message combiner using nesting

Proof (Proof sketch of Theorem 4 – unforgeability of D_{nest}). This theorem contains two statements. The first statement is: If Σ_1 is $\text{X}^{\text{Y}}\text{Z}$ -eufcma-secure, then $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is $\text{X}^{\text{Y}}\text{Z}$ -eufcma-secure with respect to its first message component only. $D_{\text{nest}}(\Sigma_1, \Sigma_2)$, when restricted to its first message component only, is just Σ_1 , so the first statement follows vacuously.

Now consider the second statement: $D_{\text{nest}}(\Sigma_1, \Sigma_2)$ is $\text{U}^{\text{V}}\text{W}$ -eufcma-secure if Σ_2 is $\text{U}^{\text{V}}\text{W}$ -eufcma-secure. Suppose \mathcal{A} is a $\text{U}^{\text{V}}\text{W}$ algorithm that outputs a forgery for $\Sigma' = D_{\text{nest}}(\Sigma_1, \Sigma_2)$ — in other words, it outputs $q_S + 1$ valid signatures under Σ' on distinct messages. We can construct an $\text{U}^{\text{V}}\text{W}$ algorithm \mathcal{B}_2 that finds a forgery in Σ_2 . \mathcal{B}_2 interacts with an $\text{U}^{\text{V}}\text{W}$ challenger for Σ_2 which provides a public key vk_2 . \mathcal{B}_2 generates a key pair $(sk_1, vk_1) \leftarrow \Sigma_1.\text{KeyGen}()$ and sets the public key for Σ' to be (vk_1, vk_2) . When \mathcal{A} asks for $\sum_{m,t,z} \psi_{m,t,z} |m, t, z\rangle$ to be signed using Σ' , we treat t as consisting of two registers $t_1 || t_2$, \mathcal{B}_2 proceeds by passing the m , t_2 , and z registers to its signing oracle for Σ_2 , then runs the quantum signing operation from Figure 1 for $\Sigma_1.\text{Sign}$ on the m , t_1 , and z registers. There is a one-to-one correspondence between \mathcal{A} 's queries to its oracle and \mathcal{B}_2 's queries to its oracle. As before in the proof of Theorem 1, if Σ_1 is proven to be secure in the random oracle model, then this proof of $C_{\text{weak-nest}}(\Sigma_1, \Sigma_2)$ also proceeds in the random oracle model: \mathcal{B}_2 relays \mathcal{A} 's hash oracle queries directly to its hash oracle, giving a one-to-one correspondence between \mathcal{A} 's queries to its (classical or quantum) hash oracle and \mathcal{B}_2 's queries to its (classical or quantum, respectively) hash oracle.

If \mathcal{A} wins the $\text{U}^{\text{V}}\text{W}$ -eufcma game, then it has returned $q_S + 1$ distinct tuples $(m_{1,i}, m_{2,i}, \sigma_{1,i}, \sigma_{2,i})$ such that $\Sigma_1.\text{Verify}(vk_1, m_{1,i}, \sigma_{1,i}) = 1$ and $\Sigma_2.\text{Verify}(vk_2, (m_{1,i}, \sigma_{1,i}, m_{2,i}), \sigma_{2,i}) = 1$.

Hence, \mathcal{B}_2 can extract $q_S + 1$ valid signatures under Σ_2 and thus it holds that $\text{Adv}_{\Sigma'}^{\text{R}^{\text{S}}\text{T-eufcma}}(\mathcal{A}) \leq \text{Adv}_{\Sigma_2}^{\text{R}^{\text{S}}\text{T-eufcma}}(\mathcal{B}_2)$. \square

Table 5. Post-quantum signature schemes; keys and signature sizes, estimated certificate sizes, and claimed security level

Scheme	Size (bytes)				Security (bits)
	public key	secret key	signature	certificate	
<i>Lattice-based</i>					
GLP [19]	1 536	256	1 186	3.0 KiB	100
Ring-TESLA-II [1]	3 328	1 920	1 488	5.1 KiB	128
TESLA#-I [3]	3 328	2 112	1 616	5.2 KiB	128
BLISS [16]	7 168	2 048	1 559	9.0 KiB	128
TESLA-416 [2]	1 331 200	1 011 744	1 280	1 332.8 KiB	128
<i>Hash-based</i>					
XMSS [11]	912	19	2 451	3.6 KiB	82
SPHINCS [6]	1,056	1,088	41,000	42.3 KiB	>128
Rainbow [13]	44 160	86 240	37	44.5 KiB	80