

Mitigating Multi-Target Attacks in Hash-based Signatures

Andreas Hülsing¹, Joost Rijneveld², and Fang Song³

¹ Department of Mathematics and Computer Science
Technische Universiteit Eindhoven
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
andreas@huelising.net

² Radboud University, Digital Security Group,
P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
joost@joostrijneveld.nl

³ Department of Combinatorics & Optimization,
and Institute for Quantum Computing
University of Waterloo
fang.song@uwaterloo.ca

Abstract. This work introduces XMSS-T, a new stateful hash-based signature scheme with tight security. Previous hash-based signatures are facing a loss of security, linear in performance parameters such as the total tree height. Our new scheme can achieve the same security level but using hash functions with a smaller output length, which immediately leads to a smaller signature size. The same techniques also apply directly to the recent stateless hash-based signature scheme SPHINCS (Eurocrypt 2015), and the signature size is reduced as well.

Being a little more specific and technical, the tight security stems from new multi-target notions of hash-function properties which we define and analyze. We show precise complexity for breaking these security properties under both classical and quantum generic attacks, thus establishing a reliable estimate for the quantum security of XMSS-T. Especially, we prove quantum query complexity tailored for cryptographic applications, which overcome some limitations of standard techniques in quantum query complexity such as they usually only consider worst-case complexity. Our proof techniques may be useful elsewhere.

We also implement XMSS-T and compare its performance to that of XMSS (PQCrypto 2011), the most recent stateful hash-based signature scheme before our work.

Keywords: post-quantum cryptography, hash-based signatures, hash function security, multi-target attacks, quantum query complexity

Full version of this paper including missing proofs is available at: <http://ia.cr/2015/1256>. This work was supported by European Commission through the ICT program under contract ICT-645622 (PQCRYPTO). Part of this work was done while the first author was visiting IQC. F.S. acknowledges support by Cryptoworks21, Canada's NSERC and ORF.

1 Introduction

Hash-based signatures are considered the most promising post-quantum alternative to existing schemes RSA and ECDSA which are vulnerable to quantum attacks. This is especially because the security of cryptographic hash functions has been well understood under intensive scrutinization. In addition, there are exact reductionist proofs relating the hardness of breaking the schemes to the hardness of breaking security properties of the hash functions used in the schemes. This allows precise estimation on the security of specific parameter sets.

Traditionally, the security of hash-based signature schemes was related to collision-resistance of the used hash function. In recent years several works focused on basing security on milder assumptions [15,11,12,21,23,5], such as second-preimage resistance and one-wayness. There are two fundamental reasons driving this trend. On the one hand, the attacks against the collision-resistance of SHA1 and MD5 motivated researchers to develop collision-resilient signature schemes [19,27]. On the other hand, collision resistance is subject to birthday attacks while (second-)preimage resistance is not. Hence, to reach a security level of λ bits, a hash function with $n = 2\lambda$ bit digests is needed if collision resistance is required whereas for (second-)preimage resistance only $n = \lambda$ bit digests are needed. Halving the output size of the used hash function immediately halves the signature and key sizes of hash-based signatures.

Multi-target attacks. The above statement is only half the truth because it bears on the implicit assumption that a hash function is used only once. Clearly, for many cryptographic constructions this is not the case. Consider for example preimage resistance (aka. one-wayness). For many cryptographic constructions, an adversary will be able to learn a magnitude of function values and security breach may occur once he finds a preimage for just one of them. More specifically, suppose that a hash function with n bit outputs is used d times in a cryptographic construction. If it suffices to invert the hash function on any one out of the d outputs to break the security of the scheme, then the attack complexity is downgraded to $\mathcal{O}(2^n/d)$ instead of $\mathcal{O}(2^n)$. Intuitively this is because every input value that an adversary tries has probability $d/2^n$ of being a solution instead of $1/2^n$, if we treat the hash function as a random function. For theoretical (asymptotic) security this worries nobody as d is normally at most polynomial in n . However, when choosing parameters in practice this can easily cause serious consequences.

This issue is indeed very pertinent to hash-based signatures. Consider for example the hash-based signature scheme XMSS [12] and its multi-tree version XMSS^{MT} [23] (see Section 4) with parameters that allow to use a keypair for a virtually unlimited amount of signatures (e.g. a total tree height of $h = 60$). In this case, an attacker can learn about 2^{66} images under the same hash function and will succeed in forging a signature if he finds a single preimage for any one of the 2^{66} values. Consequently, to achieve for example security of 256 bits one cannot use a 256 bit hash function but has to use one with output length 322.

This does not only imply the use of a hash function with a bigger output length (and hence a slowdown), it also increases the signature size by roughly 25%.

This work. In this work we introduce a new hash-based signature scheme XMSS-T that is not vulnerable to multi-target attacks. Towards this end, we propose two new *multi-target* notions for preimage and second-preimage resistance. We then analyze the generic security of hash functions with regard to these new properties against classical and quantum adversaries, proving upper and lower bounds on the query complexity of generic attacks. More specifically, the first type of notions (single-function multi-target) models a notion that is implicitly used by recent collision-resilient hash-based signature schemes like XMSS, XMSS^{MT} and SPHINCS [12,23,5]. In these notions, an adversary \mathcal{A} receives p target values and a random function from the hash function family. Then, \mathcal{A} is asked to find a preimage (or second-preimage, respectively) for one of the target values under the given function. We prove that compared to standard (second-)preimage resistance, the query complexity of generic attacks drops by a factor p for classical and \sqrt{p} for quantum adversaries. Then we introduce multi-function multi-target notions of preimage and second-preimage resistance. For these notions, \mathcal{A} is given multiple pairs of function and target value, drawn independently at random. It is now \mathcal{A} 's goal to find a preimage (or second-preimage, respectively) for one of the target values under the associated function. We prove that in this case the query complexity of generic attacks is exactly the same as for the standard (single-function, single-target) notions.

Given that multi-function multi-target notions are as hard as the standard notions of preimage and second-preimage resistance we construct a new hash-based signature scheme with security based on these new notions. As the basic construction follows that of XMSS, we call the new scheme XMSS-T, indicating XMSS with tightened security. While XMSS loses in the bit security an amount linear in several parameters including the total tree height, XMSS-T loses only two bits, independent of any parameters. The differences between XMSS^{MT} and XMSS-T are a different hash tree and one-time signature scheme construction such that the security can be based on the multi-target multi-function properties. The basic change is that for every hash function call within a hash tree or a hash chain, a different hash function key and different bitmasks are used. Note that XMSS-T is stateful and it may be not suitable in some practical use cases. The good news is that we can make similar changes to the stateless hash-based signature scheme SPHINCS easily. Roughly speaking, it amounts to replacing the used hash trees and one-time signatures by the ones described in this work.

Finally, we present an implementation of XMSS-T and compare it to XMSS and XMSS^{MT}. We show that the applied changes only have marginal performance implications (a factor 3 loss in speed for all algorithms). Our code is available at https://joostrijneveld.nl/papers/multitarget_xmss.

Remarks on proving quantum generic security. At first sight the tasks of breaking the various security properties for hash functions seem similar to some standard problems studied in quantum query complexity. However due to some limitations, existing results such as techniques for proving quantum query lower

bounds [3,2] cannot be applied directly. For example, there are famous works showing upper and lower bounds on finding collisions in r -to-1 functions [9,1]. Nonetheless, random functions, whose properties our work studies, are very unlikely to be r -to-1. More generally, quantum query complexity usually considers *worst-case* complexity only, whereas in cryptographic settings we care about average-case complexity. Another issue is that, as observed by Zhandry [30], quantum query lower bounds often have implications about quantum algorithms with *high* success probability only. For cryptographic applications however, an attacker with small but noticeable chance of breaking a scheme is still relevant. Therefore, a complete lower bound would be bounding the success probability of any algorithm making a specified number of queries. It might be possible to find fixes by digging into existing works, the situation is yet unclear. We expect that techniques developed in this work can find useful in other cryptographic settings as well.

Organization. We introduce and discuss the new security notions for hash function families in Section 2, where detailed analysis for quantum generic security is presented in Section 3. In Section 4 we present XMSS-T and discuss its security in Section 5. Finally, we present our implementation results in Section 6.

Notation. We write $x \xleftarrow{\$} \mathcal{X}$ if x is randomly chosen from the set \mathcal{X} using the uniform distribution. We further write \log for \log_2 . We denote the uniform distribution over bit strings of length n by \mathcal{U}_n . We write $m = \text{poly}(n)$ to denote that m is a function, polynomial in n . We call a function $\epsilon(n) : \mathbb{N} \rightarrow [0, 1]$ negligible and write $\epsilon(n) = \text{negl}(n)$ if for any $c \in \mathbb{N}, c > 0$ there exists a $n_c \in \mathbb{N}$ s.t.h. $\epsilon(n) < n^{-c}$ for all $n > n_c$.

2 New security notions for hash function families

In this section, we recall some known and define several new security notions for (hash) function families and discuss their security against both classical and quantum generic attacks. In the following we restrict ourselves to function families that operate on bit strings and have a fixed input size, as this is the case in our constructions. However, the definitions are the same for the more general case. In the following let $n \in \mathbb{N}$ be the security parameter, $m = \text{poly}(n), k = \text{poly}(n)$, and $\mathcal{H}_n = \{H_K : \{0, 1\}^m \rightarrow \{0, 1\}^n\}_{K \in \{0, 1\}^k}$ be a family of functions. We say a function family \mathcal{H}_n is efficient if there exists a probabilistic polynomial time (PPT) algorithm that evaluates $H_K(M)$ for any $M \in \{0, 1\}^m$ and $K \in \{0, 1\}^k$. We require all used functions to be efficient, unless we state otherwise. For hash-based signatures we are mainly interested in functions with $m, k \geq n$. However, we try to keep our results as general as possible and make it explicit whenever we are relying on $m, k \geq n$.

2.1 Defining the security notions

Preimage-resistance (ow). Let's revisit the standard notion of preimage resistance (a.k.a. one-wayness). We define the success probability of an adversary

\mathcal{A} against the preimage resistance of a hash function family \mathcal{H}_n as

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{ow}}(\mathcal{A}) &= \Pr[K \xleftarrow{\$} \{0, 1\}^k; M \xleftarrow{\$} \{0, 1\}^m, Y \leftarrow \text{H}_K(M); \\ &M' \xleftarrow{\$} \mathcal{A}(K, Y) : Y = \text{H}_K(M')] . \end{aligned} \quad (1)$$

Single-function, multi-target preimage resistance (SM-OW). We now define the success probability of an adversary against SM-OW. This is the basic multi-target notion of preimage resistance implicitly used by previous collision resilient hash-based signature schemes like XMSS. We show in Section 3 that this notion is significantly easier to attack than standard preimage resistance. The definition takes another parameter p defining the number of targets.

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n, p}^{\text{SM-ow}}(\mathcal{A}) &= \Pr[K \xleftarrow{\$} \{0, 1\}^k; M_i \xleftarrow{\$} \{0, 1\}^m, Y_i \leftarrow \text{H}_K(M_i), 0 < i \leq p; \\ &M' \xleftarrow{\$} \mathcal{A}(K, (Y_1, \dots, Y_p)) : \exists 0 < i \leq p, Y_i = \text{H}_K(M')] . \end{aligned} \quad (2)$$

Multi-function, multi-target preimage resistance (MM-OW). Next we define the success probability of an adversary \mathcal{A} against MM-OW. This is the notion we are aiming for with XMSS-T as it is as hard to break as standard preimage resistance, as we will show below. Again the definition is parameterized by the number of targets:

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n, p}^{\text{MM-ow}}(\mathcal{A}) &= \Pr[K_i \xleftarrow{\$} \{0, 1\}^k, M_i \xleftarrow{\$} \{0, 1\}^m, Y_i \leftarrow \text{H}_{K_i}(M_i), 0 < i \leq p; \\ &(j, M') \xleftarrow{\$} \mathcal{A}((K_1, Y_1), \dots, (K_p, Y_p)) : Y_j = \text{H}_{K_j}(M')] . \end{aligned} \quad (3)$$

The difference between these two new definitions is that for SM-OW all targets are for the same function while for MM-OW each target has an associated random function from the family. We decided that \mathcal{A} has to output the associated index i in case of MM-OW as otherwise any reduction would have to search for i and \mathcal{A} knows i for any attack that does better than guessing.

Second-preimage resistance (SPR). After presenting the multi-target notions for one-wayness, we now turn to second-preimage resistance. We start revisiting the standard notion of second-preimage resistance. We define the success probability of an adversary \mathcal{A} against the second-preimage resistance (SPR) of a hash function family \mathcal{H}_n as

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{SPR}}(\mathcal{A}) &= \Pr[K \xleftarrow{\$} \{0, 1\}^k; M \xleftarrow{\$} \{0, 1\}^m; \\ &M' \xleftarrow{\$} \mathcal{A}(K, M) : M' \neq M \wedge \text{H}_K(M) = \text{H}_K(M')] . \end{aligned} \quad (4)$$

Note that in this definition the adversary is not promised to receive an M that actually has a second-preimage. Hence, especially for families \mathcal{H}_n with $m = n$, i.e. same size of domain and co-domain, the adversaries success probability is

largely influenced by the probability that a random M actually has a second-preimage.

Single-function, multi-target second-preimage resistance (SM-SPR). As for one-wayness, we define two multi-target notions: single-function multi-target second-preimage resistance (SM-SPR) and multi-function multi-target second-preimage resistance (MM-SPR). The first one (SM-SPR) is the notion implicitly used in XMSS. The latter is the notion we aim for with XMSS-T that is as hard to break as standard second-preimage resistance, as we will prove below. We start defining the success probability of an adversary against SM-SPR. The definition again takes another parameter p defining the number of targets:

$$\begin{aligned} \text{Succ}_{\mathcal{H}_{n,p}}^{\text{SM-SPR}}(\mathcal{A}) &= \Pr[K \xleftarrow{\$} \{0,1\}^k; M_i \xleftarrow{\$} \{0,1\}^m, 0 < i \leq p; \\ &M' \xleftarrow{\$} \mathcal{A}(K, (M_1, \dots, M_p)) : \\ &\exists 0 < i \leq p : M' \neq M_i \wedge H_K(M_i) = H_K(M')] . \end{aligned} \quad (5)$$

Multi-function, multi-target second-preimage resistance (MM-SPR).

Next we define the success probability of an adversary \mathcal{A} against MM-SPR. Again the definition is parameterized by the number of targets:

$$\begin{aligned} \text{Succ}_{\mathcal{H}_{n,p}}^{\text{MM-SPR}}(\mathcal{A}) &= \Pr[K_i \xleftarrow{\$} \{0,1\}^k, M_i \xleftarrow{\$} \{0,1\}^m, 0 < i \leq p; \\ &(j, M') \xleftarrow{\$} \mathcal{A}((K_1, M_1), \dots, (K_p, M_p)) : \\ &M' \neq M_j \wedge H_{K_j}(M_j) = H_{K_j}(M')] . \end{aligned} \quad (6)$$

Extended target collision resistance (ETCR). In [20] Halevi and Krawczyk introduced extended target collision resistance (eTCR) as a hash function property that is close to target collision resistance. In the classical target-collision resistance game, the adversary is allowed to choose a target message M . Afterwards he learns a function (by learning a key K) and has to find a collision for the M under this function H_K . While the setup of the eTCR game is exactly the same, the adversary wins if he can present a new message M' and a (possibly new) key K' such that $H_K(M) = H_{K'}(M')$. Formally, the success probability of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, where \mathcal{A}_1 and \mathcal{A}_2 have shared memory, against ETCR is defined as follows:

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{ETCR}}(\mathcal{A}) &= \Pr[M \xleftarrow{\$} \mathcal{A}_1(1^n); K \xleftarrow{\$} \{0,1\}^k; (M', K') \xleftarrow{\$} \mathcal{A}_2(K, M) : \\ &M' \neq M \wedge H_K(M) = H_{K'}(M')] . \end{aligned} \quad (7)$$

Multi-target extended target collision resistance (M-ETCR). We can also define a multi target version (ETCR is inherently multi function anyway). To keep the definition readable we use a challenge oracle $\text{Box}(\cdot)$ that on input of a

message outputs a uniformly random function key. This oracle models the ability of \mathcal{A} to adaptively obtain p eTCR challenges for the same function family. We denote by (M_i, K_i) the i th query-answer pair of $\text{Box}(\cdot)$. The success probability of an adversary \mathcal{A} against m-eTCR that makes no more than p queries to $\text{Box}(\cdot)$ is defined as:

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{M-eTCR}}(\mathcal{A}) = \Pr[(M', K', i) \xleftarrow{\$} \mathcal{A}^{\text{Box}(\cdot)}(1^n) : M' \neq M_i \wedge H_{K_i}(M_i) = H_{K'}(M')] . \quad (8)$$

2.2 Generic security

To determine secure parameters for hash function families or constructions based on them, their security against generic attacks is analyzed. Generic attacks show which security level is achievable at all for a given property as they do not take any possibly existing function specific weaknesses into account. A hash function family is considered broken if the security level for one property is (significantly) lower than the generic security.

Classical generic security. The standard way to analyze the complexity of generic attacks against a security property of hash function families is analyzing the success probability of an adversary \mathcal{A} against a random function family to which it is given black box access. The classical security is well understood in the literature. The security of the new notions we defined can be easily established as well. For completeness, we give brief justifications in Appendix A. Table 1 summarizes the classical and quantum generic security.

Quantum generic security. When we analyze the properties of hash functions under generic quantum attacks, we treat any hash function as a random function and the adversary can issue quantum superposition queries to the function. Namely, we are essentially working under the quantum random-oracle model [6]. When there are multiple functions, we assume they are independent random functions and the adversary can query them jointly in superposition. Namely, queries in the form of

$$\sum_{K,M,z} \alpha_{K,M,z} |K, M, z\rangle \mapsto \sum_{K,M,z} \alpha_{K,M,z} |K, M, z + H_K(M)\rangle ,$$

are permitted⁴. This choice is meant to capture the fact that in reality all hash functions are public, and a quantum adversary can certainly evaluate them jointly in superposition. This is in contrast to the classical setting, where each query must specify an index, and hence the adversary only gets one value of *one* function per query. One can define a similar model in the quantum setting (i.e., each query must specify one and only one function index K) and study all the security properties therein. We stress that this model seems weaker than the

⁴ Alternatively, one can think of it as a global random function $(K, M) \mapsto O(K, M)$.

one we choose, and in particular our lower bounds results are hence stronger. Namely, they hold against stronger quantum attacks. It is an interesting theoretical question as to determining whether the two models are indeed different.

We prove our results regarding quantum generic security in Sect. 3. Our findings are summarized in Table 1. Please note that the constant hidden in the Θ is small, i.e. 16 for the lower bounds.

	OW, MM-OW, SPR, MM-SPR	SM-OW, SM-SPR	E _T TCR	M-E _T TCR
Classical	$\frac{q+1}{2^n}$	$\frac{(q+1)p}{2^n}$	$\frac{(q+1)}{2^n} + \frac{q}{2^k}$	$\frac{(q+1)p}{2^n} + \frac{qp}{2^k}$
Quantum	$\Theta\left(\frac{(q+1)^2}{2^n}\right)$	$\Theta\left(\frac{(q+1)^2 p}{2^n}\right)$	$\Theta\left(\frac{(q+1)^2}{2^n} + \frac{q^2}{2^k}\right)$	$\Theta\left(\frac{(q+1)^2 p}{2^n} + \frac{q^2 p}{2^k}\right)$

Table 1. Security against generic classical and quantum attacks. Entries represent the success probability of a q -query adversary (upper and lower bound).

3 Analyzing Quantum Generic Security

In the following we establish the generic security of hash function families against quantum attacks on the defined properties. For each security property, we give attacks and analyze their success probabilities. All attacks are based on Grover’s quantum search algorithm, but we will need to analyze the complexity for random problem instances. More importantly, we establish matching lower bounds for all cases. The proofs of lower bounds follow a unified structure. Specifically, we first define a family of distributional search problems and bound the success probability of quantum algorithms against these problems. Then, we reduce various instances of the search problem to the task of breaking each of the security properties we care about. The hardness of the distributional search problems hence implies the generic security of hash functions for these security properties.

3.1 Toolbox

(Generalized) Grover’s quantum search algorithm. One of the most useful algorithmic tools in quantum computing is Grover’s quantum search algorithm and its many generalizations (e.g., [18,7,10,8] to name a few). Here we just need a simple version for searching a universe with multiple marked items. We state it in the following Lemma.

Lemma 1. *Let $f : \mathcal{X} \rightarrow \{0, 1\}$ be an oracle function and let $\mathcal{X}_f = \{x \in \mathcal{X} : f(x) = 1\}$. Then there is a quantum algorithm QSEARCH with q queries that finds an $x \in \mathcal{X}_f$ with success probability $\Omega\left(q^2 \frac{|\mathcal{X}_f|}{|\mathcal{X}|}\right)$.*

Most of the attacks we describe later will apply QSEARCH in a straightforward way. However, since our problem instances are generated randomly, we will need to give a new analysis of the average-case performance.

A hard average-case search problem. It is well known that Grover’s search algorithm is also optimal [4]. Namely, adopting notations from Lemma 1, any q -query algorithm can find a marked item with probability at most $O(q^2 \frac{|X_f|}{|X|})$. However since the security notions we defined all refer to average-case problems, the worst-case lower bound of Grover’s search is not very useful. Here we introduce a distributional search problem, and prove a stringent hardness result.

Definition 1. Let $\mathcal{F} := \{f : \{0, 1\}^m \rightarrow \{0, 1\}\}$ be the collection of all boolean functions on $\{0, 1\}^m$. Let $\lambda \in [0, 1]$ and $\varepsilon > 0$. Define a family of distributions D_λ on \mathcal{F} such that $f \leftarrow_R D_\lambda$ satisfies

$$f : x \mapsto \begin{cases} 1 & \text{with prob. } \lambda, \\ 0 & \text{with prob. } 1 - \lambda \end{cases}$$

for any $x \in \{0, 1\}^m$.

We define $\text{Avg-Search}_\lambda$ to be the problem that given oracle access to $f \leftarrow D_\lambda$, finds an x such that $f(x) = 1$. For any quantum algorithm \mathcal{A} that makes q queries, we define

$$\text{Succ}_\lambda^q(\mathcal{A}) := \Pr_{f \leftarrow D_\lambda} [f(x) = 1 : x \leftarrow \mathcal{A}^f(\cdot)].$$

Theorem 1. $\text{Succ}_\lambda^q(\mathcal{A}) \leq 8\lambda(q+1)^2$ holds for any quantum algorithm \mathcal{A} with q queries.

Note that this theorem matches the intuitive argument that for $f \leftarrow D_\lambda$, there are $2^m \lambda$ marked items on average and hence any quantum algorithm needs $\Theta(\sqrt{2^m / (2^m \lambda)}) = \Theta(1/\sqrt{\lambda})$ queries. We defer its proof to the full version.

Simulating random functions. In our reductions to show lower bounds, we usually assume we have access to some random function $f : \mathcal{X} \rightarrow \mathcal{Y}$. Ultimately, we will need to simulate f efficiently so that any algorithm with q queries cannot notice a difference. Fortunately, the following claim allows us to do so by sampling uniformly from a $2q$ -wise independent hash function family \mathcal{H} .

Lemma 2. [29, Theorem 6.1] For any quantum adversary that makes no more than q queries to either a truly random function or a function drawn uniformly from \mathcal{H} , the final states are identical.

There exists a vast literature on efficient constructions of t -wise independent hash functions. Interested readers are referred to, e.g., [24,14,25]. There is a technical subtlety though. Most constructions of \mathcal{H} consider output space \mathcal{Y} with size being a prime or a prime power. We need one with $\mathcal{Y} = [N]$, $N = 2^n - 1$. A natural approach is to pick a prime $M \gg N$ and construct a $2q$ -wise independent family $\mathcal{H}_0 : \mathcal{X} \rightarrow [M]$. Then we would expect that $\mathcal{H} : x \mapsto \mathcal{H}_0(x) \bmod N$ will suffice for our purpose, modulo a tiny error. However we were unable to identify a rigorous proof in the literature for the correctness of this

“mod” construction, especially with respect to quantum attacks. We give a formal proof in the full version.

Sometimes we need a random function f that excludes some output $y \in \mathcal{Y}$. This is easy to realize as follows. We take a random function $g : \mathcal{X} \rightarrow [k]$ where $k = |\mathcal{Y}| - 1$. Then $f(x)$ will be obtained by applying g on x and then mapping the outcome to $\mathcal{Y} \setminus y$ according to some canonical isomorphism (e.g., any thing smaller than y remains unchanged, and anything else is incremented by 1.).

3.2 Hardness of breaking the security

We analyze in this section the hardness of generic quantum attacks on the various notions of hash functions. We give upper bounds on the success probabilities of any quantum adversary making at most q queries. Basically, we reduce $\text{Avg-Search}_\lambda$ with various λ to the task of breaking the security notion generically. The hardness of Avg-Search then implies the security against generic quantum attacks. The bounds for OW, SPR and their variants are given in Propositions 1 and 2. While the proofs are quite similar we have to deal with a restriction for the OW notions that we did not figure out how to circumvent. Namely, we require that $2^m \gg 2^n$ (e.g. $m = 2n$) and $p \ll 2^n$, which is the case for most relevant hash function families. The complexity for ETCR and M-ETCR involves additional technical difficulty concerning programming a random oracle, and we analyze them in Proposition 3.

Proposition 1. *Let $m = cn$ for a positive real constant $c > 1$ and $p = o(n)$. For any quantum adversary with q queries, it holds that*

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{OW}}(\mathcal{A}) &= O((q+1)^2/2^n), \text{Succ}_{\mathcal{H}_n}^{\text{SM-OW}}(\mathcal{A}) = O((q+1)^2 p/2^n), \\ \text{Succ}_{\mathcal{H}_n}^{\text{MM-OW}}(\mathcal{A}) &= O((q+1)^2/2^n). \end{aligned}$$

The proof is given in the full version.

Proposition 2. *For any quantum adversary with q queries, it holds that*

$$\begin{aligned} \text{Succ}_{\mathcal{H}}^{\text{SPR}}(\mathcal{A}) &= O((q+1)^2/2^n), \text{Succ}_{\mathcal{H}_n}^{\text{SM-SPR}}(\mathcal{A}) = O((q+1)^2 p/2^n), \\ \text{Succ}_{\mathcal{H}_n}^{\text{MM-SPR}}(\mathcal{A}) &= O((q+1)^2/2^n). \end{aligned}$$

We give the proof for MM-SPR. The others can be proven analogously and are deferred to the full version.

Proof (Hardness of MM-SPR). Given an Avg-Search instance, we construct an instance of MM-OW in Figure 1:

Note that the way that f is generated ensures that each constructed \tilde{H}_i is distributed identically to a uniformly random function $H : \{0, 1\}^m \rightarrow \{0, 1\}^n$. Therefore the output instance in the reduction is valid according to the definition

Given: $f \leftarrow D_\lambda : \{0, 1\}^m \rightarrow \{0, 1\}^n$, $\lambda = p/2^n$.

1. For $i = 1, \dots, p$, sample $M_i \leftarrow \{0, 1\}^m$ and $y_i \leftarrow \{0, 1\}^n$ independently and uniformly at random. Denote $S = \{M_i\}$ and $T = \{y_i\}$.
2. For $i = 1, \dots, p$, let $g_i : \{0, 1\}^m \rightarrow \{0, 1\}^n \setminus T$ be random function. We construct $\tilde{H}_i : \{0, 1\}^m \rightarrow \{0, 1\}^n$ as follows: for any $x \in \{0, 1\}^m$

$$x \mapsto \begin{cases} y_i & \text{if } x = M_i \\ y_i & \text{if } x \neq M_i \wedge f(i||x) = 1 \\ g_i(x) & \text{otherwise.} \end{cases}$$

Output: MM-SPR instance $(S, \{\tilde{H}_i\}_{i=1}^p)$. Namely an adversary is given M_i and oracle access to \tilde{H}_i , and the goal is to find (i^*, x^*) such that $x^* \neq M_{i^*}$ and $\tilde{H}_{i^*}(x^*) = y_{i^*}$.

Fig. 1. Reducing Avg-Search to MM-SPR.

Eq. 6. This implies that any q -query attacker solving MM-SPR will give rise to a $2q$ -query algorithm for Avg-Search $_\lambda$. As a consequence

$$\text{Succ}_{\mathcal{H}_n}^{\text{MM-SPR}}(\mathcal{A}) \leq \text{ADV}_{\mathcal{A}}^{2q}(\lambda) \leq 16(q+1)^2/2^n,$$

follows by Theorem 1. We remark that, as mentioned in Sect. 3.1, \tilde{H}_i can be implemented efficiently.

Proposition 3. Let $\varepsilon = 8(q+1)^2/2^n$ and $\delta = 4q^2/2^k$. For any quantum adversary with q queries, it holds that

$$\text{Succ}_{\mathcal{H}_n}^{\text{E}^{\text{TCR}}}(\mathcal{A}) \leq \varepsilon + 2\delta, \quad \text{Succ}_{\mathcal{H}_n}^{\text{M-E}^{\text{TCR}}}(\mathcal{A}) \leq p(\varepsilon + 2\delta).$$

To prove the proposition, we need a lemma that allows us to adaptively program a quantum random oracle. The proof follows standard techniques (see similar analyses for different scenarios in [28, 16]). Let \mathcal{A} be an arbitrary quantum algorithm and let $H : \{0, 1\}^m \times \{0, 1\}^k \rightarrow \{0, 1\}^n$ be a random function. Consider two games as follows:

- Game G_0 : \mathcal{A} gets access to H . In phase 1, after making at most q_1 queries to H , \mathcal{A} outputs a message $M \in \{0, 1\}^m$. Then a random $\hat{K} \in_R \{0, 1\}^k$ is sampled and $(\hat{K}, H_{\hat{K}}(M))$ is handed to \mathcal{A} . \mathcal{A} continues to the second phase and makes at most q_2 queries. \mathcal{A} outputs $b \in \{0, 1\}$ at the end.
- Game G_1 : \mathcal{A} gets access to H . After making at most q_1 queries to H , \mathcal{A} outputs a message $M \in \{0, 1\}^m$. Then a random $\hat{K} \in_R \{0, 1\}^k$ is sampled as well as a random range element $y \in_R \{0, 1\}^n$. Program $H_{\hat{K}}(M) = y$ and call the new oracle H' . \mathcal{A} receives $(\hat{K}, y = H'_{\hat{K}}(M))$ and proceeds to the second phase. After making at most q_2 queries, \mathcal{A} outputs $b \in \{0, 1\}$ at the end.

Lemma 3. $|\Pr[\mathcal{A}(G_0) = 1] - \Pr[\mathcal{A}(G_1) = 1]| \leq 2\delta$, with $\delta = 4q^2/2^k$.

The proof of the Lemma is described in the full version. Using the Lemma, we can prove Proposition 3.

Proof (Proof of Proposition 3). We give a reduction from Avg-Search to breaking eTCR. Assume that there is \mathcal{A} that breaks eTCR with probability η . We construct an adversary \mathcal{A}' that solves Avg-Search with probability $\eta - 2\delta$. Note that as long as \mathcal{A} does not notice that we reprogrammed H , its view would be identical to that of the standard eTCR game, and by assumption \mathcal{A} wins with probability at least η . By Lemma 3, reprogramming only incurs an additive error $2\delta = 4q^2/2^k$. We claim that $\Pr[f(K^*, M^*) = 1] \geq \eta - 2\delta$. But we know that the success probability of Avg-Search is at most $\varepsilon := 8(q+1)^2/2^n$ by Theorem 1. Therefore $\eta \leq \varepsilon + 2\delta$ and this proves Proposition 3. We can generalize the arguments above to the multi-target case easily.

Given: $f \leftarrow D_\lambda : \mathcal{X} := \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^n, \lambda = 1/2^n$.

1. Sample $y \leftarrow \{0, 1\}^n$ uniformly at random.
2. Let $g : \mathcal{X} \rightarrow \{0, 1\}^n \setminus \{y\}$ be a random function. Construct $H : \mathcal{X} \rightarrow \{0, 1\}^n$ as follows: for any $x \in \mathcal{X}$

$$x \mapsto \begin{cases} y & \text{if } f(x) = 1, \\ g(x) & \text{otherwise.} \end{cases}$$

3. \mathcal{A} accesses H and issues q_1 queries. \mathcal{A} outputs M at end of Phase 1.
4. Sample $K \leftarrow \{0, 1\}^k$. Program $H_K(M) = y$. Denote the new oracle H' . Send (K, y) to \mathcal{A} .
5. \mathcal{A} makes q_2 queries to H' . Outputs (K^*, M^*) .

Output: (K^*, M^*) .

Fig. 2. Reducing Avg-Search to eTCR

3.3 Quantum attacks

In this section, we apply quantum search algorithm QSEARCH to attack the various notions generically. In most cases, we get bounds on success probabilities matching the hardness results we have shown in Sect. 3.2.

Proposition 4. *There exist quantum adversaries $\mathcal{A}_1, \dots, \mathcal{A}_8$ all of which making $\Theta(q)$ queries, such that*

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{OW}}(\mathcal{A}_1) &= \Omega(q^2/2^n), \text{Succ}_{\mathcal{H}_n}^{\text{SM-OW}}(\mathcal{A}_2) = \Omega(q^2p/2^n), \\ \text{Succ}_{\mathcal{H}_n}^{\text{MM-OW}}(\mathcal{A}_3) &= \Omega(q^2/2^n); \text{Succ}_{\mathcal{H}_n}^{\text{SPR}}(\mathcal{A}_4) = \Omega(q^2/2^n), \\ \text{Succ}_{\mathcal{H}_n}^{\text{SM-SPR}}(\mathcal{A}_5) &= \Omega(q^2p/2^n), \text{Succ}_{\mathcal{H}_n}^{\text{MM-SPR}}(\mathcal{A}_6) = \Omega(q^2/2^n); \\ \text{Succ}_{\mathcal{H}_n}^{\text{eTCR}}(\mathcal{A}_7) &= \Omega(q^2/2^n), \text{Succ}_{\mathcal{H}_n}^{\text{M-eTCR}}(\mathcal{A}_8) = \Omega(q^2p/2^n). \end{aligned}$$

The proof is adapting standard analysis to the average case. We illustrate the basic idea by proving the case of preimage-resistance. The others are left to the full version.

Proof (Quantum attack on OW). We describe a $O(q)$ -query attacker \mathcal{A}_1 as follows. Given y and oracle access to H , \mathcal{A}_1 will apply QSEARCH to search for x such that $H(x) = y$. More specifically, \mathcal{A}_1 constructs $g_H : \{0, 1\}^m \rightarrow \{0, 1\}$ such that $g_H(x) = 1$ iff. $H(x) = y$. Each evaluation on g_H can be realized efficiently by two queries to h . For any $h \in \mathcal{H}$, let $p_h := \Pr_{H \leftarrow \mathcal{H}}[H = h]$ and let $X_H = |H^{-1}(y)|$ be the random variable representing the preimage size of y . Then by Lemma 1 we can see that

$$\begin{aligned} \text{Succ}_{\mathcal{H}_n}^{\text{ow}}(\mathcal{A}_1) &= \sum_h p_h \cdot \Omega\left(q^2 \frac{X_h}{2^m}\right) = \Omega\left(\frac{q^2}{2^m} \sum_h p_h X_h\right) \\ &= \Omega\left(\frac{q^2}{2^m} \cdot \mathbb{E}(X_H)\right) = \Omega\left(\frac{q^2}{2^n}\right). \end{aligned}$$

In the last step we observe, by linearity of expectation, that

$$\mathbb{E}(X_H) = \sum_{x \in X} \Pr_{H \leftarrow \mathcal{H}}[H(x) = y] = 2^m / 2^n.$$

4 XMSS-T

The eXtended Merkle Signature Scheme (XMSS) was proposed by Buchmann, Dahmen, and Hülsing in [12]. The original proposal for XMSS essentially combines a collision-resilient version of the Winternitz one-time signature scheme (WOTS) from [11] with the collision-resilient hash tree construction from [15] and adds two different kinds of pseudorandom key generation, one leading an EU-CMA-secure and one a forward-secure signature scheme. Under the name XMSS^{MT} Hülsing, Rausch, and Buchmann [23] later proposed a multi-tree version of XMSS.

In this work we introduce XMSS-T, XMSS with tightened security. In contrast to XMSS, XMSS-T avoids multi-target attacks. To this end, XMSS-T uses a new hash tree construction and a new WOTS variant WOTS-T. XMSS-T is based on XMSS^{MT}. The main difference in the construction of XMSS^{MT} and XMSS-T is the use of independent function keys and bitmasks for *every* call to a hash function inside of the hash trees or WOTS-T. XMSS^{MT} used a single fixed key per function family and the same bitmask per internal tree level or chain position. The function keys and bitmasks used by XMSS-T are needed for verification. To keep the public key small these values are generated pseudorandomly, using a hash-based pseudorandom function family and a seed value that becomes part of the public key. In the following we describe XMSS-T.

Parameters. XMSS-T uses several parameters and several functions. The main security parameter is $n \in \mathbb{N}$, the message digest length $m \in \text{poly}(n)$. The functions include two keyed, short-input cryptographic hash functions $F : \{0, 1\}^n \times$

$\{0,1\}^n \rightarrow \{0,1\}^n$ and $H : \{0,1\}^n \times \{0,1\}^{2n} \rightarrow \{0,1\}^n$; one arbitrary-input randomized hash function $\mathcal{H} : \{0,1\}^m \times \{0,1\}^* \rightarrow \{0,1\}^m$; and two ensembles of pseudorandom function families $\mathcal{F}_n : \{0,1\}^n \times \{0,1\}^* \rightarrow \{0,1\}^n$, $\mathcal{F}_m : \{0,1\}^n \times \{0,1\}^* \rightarrow \{0,1\}^m$, where we denote by $\{0,1\}^*$ the ability to handle arbitrary input lengths up to some practical limit (e.g. 2^{64} bits as in the case of the SHA family). Of course, these functions can all be built from a single cryptographic hash function, but the security analysis gets easier separating the functions according to the required properties.

XMSS-T uses a hyper-tree (a tree of trees) of total height $h \in \mathbb{N}$, where h is a multiple of d and the hyper-tree consists of d layers of trees, each having height h/d . WOTS allows for a space-time trade-off using the Winternitz parameter $w \in \mathbb{N}, w > 1$. The Winternitz parameter w and the length of the bit string that is signed λ determine ℓ the number of function chains for WOTS:

$$\ell_{1,\lambda} = \left\lceil \frac{\lambda}{\log(w)} \right\rceil, \quad \ell_{2,\lambda} = \left\lceil \frac{\log(\ell_1(w-1))}{\log(w)} \right\rceil + 1, \quad \ell_\lambda = \ell_1 + \ell_2.$$

The bit strings signed using WOTS are the m -bit message digests on the lowest layer and the n -bit root nodes of the layer below on all other layers.

As a running example we present concrete numbers for XMSS-T-256; the choices are explained in Section 6. For XMSS-T-256 we use $n = 256, m = 316, h = 60, d = 3, w = 16$ which leads to $\ell_n = 67$ and $\ell_m = 82$.

Addressing scheme. XMSS-T requires an addressing scheme for hash function calls. Every addressing scheme that assigns to every call to either F or H within the virtual structure of a XMSS-T hyper-tree a unique address can be used (e.g. numbering all the calls in some order). We suggest to use a recursive addressing scheme that numbers sub-structures (e.g. a OTS key pair) inside a structure (e.g. a tree). The addressing scheme generates an address for a substructure, taking the address of the structure and appending the index of the substructure. For trees, which contain three different kinds of substructures (OTS key pairs, L-trees, and nodes), an additional identifier for the type of substructure is added. Below we assume that a function $\text{GENADDR}(\mathbf{a}_s, \text{index})$ exists that takes the address of the structure and the index of the substructure and outputs a unique address for this substructure within an XMSS-T key pair. The advantage of this addressing scheme is that it only uses information that is available when the hash call is executed.

The addressing scheme is publicly known and the same addresses can be used for all XMSS-T key pairs. The resulting addresses are used as inputs to PRF \mathcal{F}_n to pseudorandomly generate function keys and bitmasks.

WOTS-T. We now describe the new WOTS version. The construction differs from [21] in that it uses fresh keys and bitmasks for each hash function call. We denote the message length by $\lambda \in \{n, m\}$ and to improve readability we write ℓ, ℓ_1 , and ℓ_2 instead of $\ell_\lambda, \ell_{1,\lambda}$, and $\ell_{2,\lambda}$. We include pseudorandom key generation, meaning that a seed value takes the place of a secret key in our description. We describe the algorithms as used by XMSS-T, hence, they take

global secret and public information. For a standalone version, this information would have to be generated during key generation.

The difference between all WOTS variants is in the way the so called chaining function is constructed. WOTS-T uses the function F to construct the following chaining function:

Chaining function $c^{i,j}(x, \mathbf{a}_C, \text{SEED})$: On input of value $x \in \{0, 1\}^n$, iteration counter $i \in \mathbb{N}$, start index $j \in \mathbb{N}$, chain address \mathbf{a}_C , and (public) seed SEED , the chaining function works the following way. In case $i = 0$, c returns x , i.e., $c^{0,j}(x, \mathbf{a}_C, \text{SEED}) = x$. For $i > 0$ we define c recursively as

$$c^{i,j}(x, \mathbf{a}_C, \text{SEED}) = F(k_{i,j}, c^{i-1,j}(x, \mathbf{a}_C, \text{SEED}) \oplus r_{i,j}),$$

where key $k_{i,j} = \mathcal{F}_n(\text{SEED}, \text{GENADDR}(\mathbf{a}_C, 2 \cdot (j+i)))$ and bitmask $r_{i,j} = \mathcal{F}_n(\text{SEED}, \text{GENADDR}(\mathbf{a}_C, 2 \cdot (j+i) + 1))$. I.e. in every round, the function first takes the bitwise xor of the previous value $c^{i-1,j}(x, \mathbf{a}_C, \text{SEED})$ and bitmask $r_{i,j}$ and evaluates F with key $k_{i,j}$ on the result.

Now we describe the three algorithms of WOTS-T.

Key Generation Algorithm $((\text{sk}, \text{pk}) \leftarrow \text{WOTS.kg}(\mathcal{S}, \mathbf{a}_{\text{OTS}}, \text{SEED}))$: On input of a global secret key seed $\mathcal{S} \in \{0, 1\}^n$ (used for every WOTS-T keypair within a XMSS-T keypair), the address of the WOTS-T keypair within a tree \mathbf{a}_{OTS} , and public seed SEED , the key generation algorithm computes the internal secret key $\text{sk} = (\text{sk}_1, \dots, \text{sk}_\ell)$ as $\text{sk}_i \leftarrow \mathcal{F}_n(\mathcal{S}, \text{GENADDR}(\mathbf{a}_{\text{OTS}}, i))$, i.e., the ℓ n bit secret key elements are derived from the secret key seed using the address of the chain they are contained in. The public key pk is computed as

$$\text{pk} = (\text{pk}_1, \dots, \text{pk}_\ell) = (c^{w-1,0}(\text{sk}_1, \mathbf{a}_{C_1}, \text{SEED}), \dots, c^{w-1,0}(\text{sk}_\ell, \mathbf{a}_{C_\ell}, \text{SEED})),$$

where $\mathbf{a}_{C_i} = \text{GENADDR}(\mathbf{a}_{\text{OTS}}, i)$. Note that \mathcal{S} requires less storage than sk ; thus we generate sk and pk on the fly when necessary.

Signature Algorithm $(\sigma \leftarrow \text{WOTS.sign}(M, \mathcal{S}, \mathbf{a}_{\text{OTS}}, \text{SEED}))$: On input of a λ -bit message M , the global secret key seed $\mathcal{S} \in \{0, 1\}^n$, the address of the WOTS-T keypair within a tree \mathbf{a}_{OTS} , and public seed SEED , the signature algorithm first computes a base- w representation of M : $M = (M_1 \dots M_{\ell_1})$, $M_i \in \{0, \dots, w-1\}$. That is, M is treated as the binary representation of a natural number x and then the w -ary representation of x is computed. Next it computes the checksum $C = \sum_{i=1}^{\ell_1} (w-1 - M_i)$ and its base w representation $C = (C_1, \dots, C_{\ell_2})$. The length of the base w representation of C is at most ℓ_2 since $C \leq \ell_1(w-1)$. We set $B = (b_1, \dots, b_\ell) = M \parallel C$, the concatenation of the base w representations of M and C . Then the internal secret key is generated using $\text{sk}_i \leftarrow \mathcal{F}_n(\mathcal{S}, \text{GENADDR}(\mathbf{a}_{\text{OTS}}, i))$ the same way as during key generation. The signature is computed as

$$\sigma = (\sigma_1, \dots, \sigma_\ell) = (c^{b_1,0}(\text{sk}_1, \mathbf{a}_{C_1}, \text{SEED}), \dots, c^{b_\ell,0}(\text{sk}_\ell, \mathbf{a}_{C_\ell}, \text{SEED})),$$

where $\mathbf{a}_{C_i} = \text{GENADDR}(\mathbf{a}_{\text{OTS}}, i)$ as above.

Input: Leaf index i , leaf L_i , authentication path $\text{Auth}_i = (A_0, \dots, A_{h-1})$ for L_i .

Output: Root node ROOT of the tree that contains L_i .

Set $P_0 \leftarrow L_i$;

for $j \leftarrow 1$ **up to** h **do**

Set $i' = \lfloor i/2^j \rfloor$;

$P_j = \begin{cases} \text{H}_{k_{i',j}}((P_{j-1} \| A_{j-1}) \oplus r_{i',j}), & \text{if } \lfloor i/2^{j-1} \rfloor \equiv 0 \pmod{2}; \\ \text{H}_{k_{i',j}}(A_{j-1} \| P_{j-1}) \oplus r_{i',j}), & \text{if } \lfloor i/2^{j-1} \rfloor \equiv 1 \pmod{2}; \end{cases}$

end

return P_h

Algorithm 1: Root Computation

the tree is constructed as described above but with one difference: A left node that has no right sibling is lifted to a higher level of the L-Tree until it becomes the right sibling of another node. Apart from this the computations work the same as for binary trees. The L-Trees have height $\lceil \log \ell_\lambda \rceil$.

4.1 XMSS-T

Given all of the above we can finally describe the algorithms of the XMSS-T construction. An XMSS-T keypair completely defines a hyper-tree of height h that consists of d layers of trees of height h/d . Each of these trees looks as follows. The leaves of a tree are $2^{h/d}$ L-Tree root nodes that each compress the public key of a WOTS-T key pair. Hence, a tree can be viewed as a key pair that can be used to sign $2^{h/d}$ messages. The hyper-tree is structured into d layers. On layer $d - 1$ it has a single tree. On layer $d - 2$ it has $2^{h/d}$ trees. The roots of these trees are signed using the WOTS-T key pairs of the tree on layer $d - 1$. In general, layer i consists of $2^{(d-1-i)(h/d)}$ trees and the roots of these trees are signed using the WOTS-T key pairs of the trees on layer $i + 1$. Finally, on layer 0 the WOTS-T key pairs are used to sign the message digests.

To improve readability, we only give a functional description of the algorithms of XMSS-T. To obtain a practical scheme, this has to be combined with the distributed signature generation method from [23] which in turn makes use of the BDS algorithm [13] for efficient tree traversal.

Key Generation Algorithm $((\text{SK}, \text{PK}) \leftarrow \text{kg}(1^n))$: The key generation algorithm first samples two secret values $(\text{SK}_1, \text{SK}_2) \in \{0, 1\}^n \times \{0, 1\}^n$. The value $\text{SK}_1 = \mathcal{S}$ is the seed used for pseudorandom key generation in WOTS-T. The value SK_2 is used to generate pseudorandom values to randomize the message hash in `sign`. Also, the public seed $\text{SEED} \xleftarrow{\mathcal{S}} \{0, 1\}^n$ is sampled as a uniform random value.

The remaining part of `kg` consists of generating the root node of the tree on layer $d - 1$. Towards this end the WOTS-T key pairs for the single tree on layer $d - 1$ are generated using SK_1 as \mathcal{S} . The i th leaf L_i of the tree is the root of an L-Tree that compresses pk_i . Finally, a binary hash tree is built using the constructed leaves and its root node becomes PK_1 .

Besides the secret values and SEED, the secret key also contains the index i of the next WOTS-T key pair to use for message signing. The index takes h bits and is initialized with the all 0 bit string. The XMSS-T secret key is $\text{SK} = (i = 0^h, \text{SK}_1, \text{SK}_2, \text{SEED})$, the public key is $\text{PK} = (\text{PK}_1, \text{SEED})$. kg returns the key pair $((\text{SK}_1, \text{SK}_2, \text{SEED}), (\text{PK}_1, \text{SEED}))$.

Signature Algorithm $((\Sigma, \text{SK}) \leftarrow \text{sign}(M, \text{SK}))$: On input of a message $M \in \{0, 1\}^*$ and secret key $\text{SK} = (i, \text{SK}_1, \text{SK}_2, \text{SEED})$, sign computes a randomized message digest $D \in \{0, 1\}^m$: First, a pseudorandom $R \in \{0, 1\}^m$ is computed as $R \leftarrow \mathcal{F}_m(\text{SK}_2, M)$. Then, $D \leftarrow \mathcal{H}(R, M)$ is computed as the randomized hash of M using R as randomness. Note that signing is deterministic, i.e., we need no real randomness as all required ‘randomness’ is pseudorandomly generated using PRF \mathcal{F}_m .

Given index i , the i th WOTS-T key pair on layer $d = 0$ is used to sign D . More specifically, this is the i_0 th WOTS-T keypair in the i'_0 th tree on layer 0, where i_0 is given by the last h/d bits of i and i'_0 by the remaining $(d-1)h/d$ bits of i . Next, the authentication path Auth_{i_0} for the i_0 th leaf of the i'_0 th tree is computed as well as the root of that tree. Now, for every layer $1 \leq \delta \leq d-1$ the same procedure is repeated with the difference that $i = i'_{\delta-1}$ and the root computed on layer $\delta-1$ is signed. So, to sign the root from layer $\delta-1$, the i_δ th WOTS-T keypair in the i'_δ th tree on layer δ is used, where i_δ is given by the last h/d bits of $i'_{\delta-1}$ and i'_δ by the remaining $(d-1)h/d$ bits of $i'_{\delta-1}$. Then the authentication path Auth_{i_δ} for the i_δ th leaf of the i'_δ th tree is computed as well as the root of that tree. The XMSS-T signature $\Sigma = (i, R, \sigma_{W,0}, \text{Auth}_{i_0}, \dots, \sigma_{W,d-1}, \text{Auth}_{i_{d-1}})$ contains the used index i , randomness R and one WOTS-T signature – authentication path pair $\sigma_{W,j}, \text{Auth}_{i_j}, j \in [d-1]$ per layer.

Finally, sign updates the secret key SK setting $i = i + 1$ and outputs the pair (Σ, SK) .

Verification Algorithm $(b \leftarrow \text{vf}(M, \Sigma, \text{PK}))$: On input of a message $M \in \{0, 1\}^*$, a signature Σ , and a public key PK , the algorithm computes the message digest $D \leftarrow \mathcal{H}(R, M)$ using the randomness R contained in the signature. Using i , the indices i_δ, i'_δ are computed for $0 \leq \delta \leq d-1$. The message digest D and the SEED from PK are used to compute the first WOTS-T public key $\text{pk}_{W,0} \leftarrow \text{WOTS.vf}(D, \sigma_{W,0}, \mathbf{a}_{\text{OTS}_0}, \text{SEED})$, where $\mathbf{a}_{\text{OTS}_0}$ is the address of the i_0 th WOTS-T keypair in the i'_0 th tree on layer 0. An L-Tree is used to compute L_{i_0} , the leaf corresponding to $\text{pk}_{W,0}$. Then, the root ROOT_0 of the respective tree is computed using Algorithm 1 with index i_0 , leaf L_{i_0} and authentication path Auth_{i_0} .

Then, this procedure gets repeated for layers 1 to $d-1$ with the following two differences. First, on layer $1 \leq \delta \leq d-1$ the root of the previously processed tree $\text{ROOT}_{\delta-1}$ is used to compute the WOTS-T public key $\text{pk}_{W,\delta}$. Second, the leaf computed from $\text{pk}_{W,\delta}$ using an L-Tree is L_{i_δ} . The result of the final repetition on layer $d-1$ is a value ROOT_{d-1} for the root node of the single tree on the top layer. This value is compared to the first element of the public key, i.e., $\text{PK}_1 \stackrel{?}{=} \text{ROOT}_{d-1}$. If the comparison holds, vf returns true, otherwise false.

5 Security

In the following we give a security reduction for XMSS-T. First, we review the required security definitions. Afterwards we give a security reduction for XMSS-T.

Existential Unforgeability under Adaptive Chosen Message Attacks.

The standard security notion for digital signature schemes is existential unforgeability under adaptive chosen message attacks (EU-CMA) [17] which is defined using the following experiment. By $\text{DSS}(1^n)$ we denote a signature scheme with security parameter n .

Experiment $\text{Exp}_{\text{DSS}(1^n)}^{\text{EU-CMA}}(\mathcal{A})$

$(\text{sk}, \text{pk}) \leftarrow \text{kg}(1^n)$

$(\text{Msg}^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{pk})$

Let $\{(\text{Msg}_i, \sigma_i)\}_1^q$ be the query-answer pairs of $\text{sign}(\text{sk}, \cdot)$.

Return 1 iff $\text{vf}(\text{pk}, \text{Msg}^*, \sigma^*) = 1$ and $\text{Msg}^* \notin \{\text{Msg}_i\}_1^q$.

For the success probability of an adversary \mathcal{A} in the above experiment we write

$$\text{Succ}_{\text{DSS}(1^n)}^{\text{EU-CMA}}(\mathcal{A}) = \Pr \left[\text{Exp}_{\text{DSS}(1^n)}^{\text{EU-CMA}}(\mathcal{A}) = 1 \right].$$

A signature scheme is called EU-CMA-secure if any PPT adversary has only negligible success probability:

Definition 2 (EU-CMA). *Let $n \in \mathbb{N}$, DSS a digital signature scheme as defined above. We call DSS EU-CMA-secure if for all $q, t = \text{poly}(n)$ the maximum success probability $\text{InSec}^{\text{EU-CMA}}(\text{DSS}(1^n); t, q)$ of all possibly probabilistic adversaries \mathcal{A} running in time $\leq t$, making at most q queries to Sign in the above experiment, is negligible in n :*

$$\text{InSec}^{\text{EU-CMA}}(\text{DSS}(1^n); t, q) \stackrel{\text{def}}{=} \max_{\mathcal{A}} \{ \text{Succ}_{\text{DSS}(1^n)}^{\text{EU-CMA}}(\mathcal{A}) \} = \text{negl}(n).$$

To be precise, XMSS-T is a so-called key-evolving signature scheme which automatically updates the secret key after each signature. We capture this, assuming that the oracle $\text{sign}(\text{sk}, \cdot)$ in the above experiment replaces the secret key sk after each signature with the one returned by XMSS-T.sign and that it returns the empty string when $i \geq 2^h$, i.e. when the maximum number of signatures were done.

Pseudorandom Function Families. In the following we give the missing definition for the properties of (hash) function families that we use, namely pseudorandomness. In our definition we use the definition of (hash) function families from Section 2. In the definition of the success probability of an adversary against pseudorandomness (PRF) the adversary gets black-box access to an oracle Box . Box is either initialized with a function from \mathcal{H}_n or a function from the set $\mathcal{G}(m, n)$ of all functions with domain $\{0, 1\}^m$ and range $\{0, 1\}^n$. The goal of the

adversary is to distinguish both cases:

$$\text{Succ}_{\mathcal{H}_n}^{\text{PRF}}(\mathcal{A}) = \left| \Pr[\text{Box} \stackrel{\$}{\leftarrow} \mathcal{H}_n : \mathcal{A}^{\text{Box}(\cdot)} = 1] - \Pr[\text{Box} \stackrel{\$}{\leftarrow} \mathcal{G}(m, n) : \mathcal{A}^{\text{Box}(\cdot)} = 1] \right|. \quad (9)$$

Using this success probability, we define a pseudorandom function family the following way.

Definition 3 (PRF). *Let \mathcal{H}_n be defined as above. We call \mathcal{H}_n a pseudorandom function family, if it is efficient and for all $t = \text{poly}(n)$ the maximum success probability $\text{InSec}^{\text{PRF}}(\mathcal{H}_n; t)$ of all possibly probabilistic adversaries \mathcal{A} , running in time $\leq t$, is negligible in n :*

$$\text{InSec}^{\text{PRF}}(\mathcal{H}_n; t) \stackrel{\text{def}}{=} \max_{\mathcal{A}} \{\text{Succ}_{\mathcal{H}_n}^{\text{PRF}}(\mathcal{A})\} = \text{negl}(n).$$

5.1 Security reduction

We now proof the security of XMSS-T. We will base the security of the core scheme on the multi-function multi-target second-preimage resistance of F, H , the pseudorandomness of \mathcal{F}_n , the multi-target extended target collision resistance of \mathcal{H} and a functional requirement on F defined below in the random oracle model. Please note that the random oracle model is only required to show that we can hand out the seed SEED used to generate the public function keys and bitmasks. Towards this end, we have to split the use of \mathcal{F}_n into two parts. Assume two functions \mathcal{F}_n^1 and \mathcal{F}_n^2 . We assume \mathcal{F}_n^1 is used in place of \mathcal{F}_n for pseudorandom (secret) key generation and generation of the message hash randomness. For \mathcal{F}_n^1 we require standard model pseudorandomness. On the other hand, \mathcal{F}_n^2 is used to replace \mathcal{F}_n when generating the hash keys $k_{i,j}$ and bitmasks $r_{i,j}$. In the proof, *only* \mathcal{F}_n^2 is modeled as random oracle (using the concatenation of key and input as input to the RO).

As mentioned above we need an additional requirement on F . Informally we require that every element in the image of F has at least two preimages, i.e.,

$$(\forall k \in \{0, 1\}^n)(\forall y \in \text{IMG}(F_k))(\exists x, x' \in \{0, 1\}^n) : x \neq x' \wedge F_k(x) = F_k(x'). \quad (10)$$

Please note that this requirement meets the expectation for a random function. This additional requirement is needed to not having to use the one-wayness of F . If we had to use the one-wayness of F , we still would have to guess the messages an adversary sends to the oracle. The reason is that plugging a challenge image into a chain means not knowing any previous value of the chain. Hence, we could not answer a query where the signature contains such a previous value of a chain. This would imply a security loss of roughly h bits. Given the above property we can instead extract a second preimage if \mathcal{A} inverts F with probability $1/2$, loosing only 1 bit in the security level.

Now we got everything needed for the security reduction. We proof the following theorem:

Theorem 2. *XMSS-T is existentially unforgeable under adaptive chosen message attacks with respect to the random oracle model if*

- F and H are multi-function multi-target second-preimage resistant function families,
- F fulfills the requirement of Eqn. 10,
- $\mathcal{F}_n^1, \mathcal{F}_m$ are pseudorandom function families,
- \mathcal{F}_n^2 is modeled as a random oracle, and
- \mathcal{H} is an multi-target extend target collision resistant hash function family.

More specifically, the insecurity function $\text{InSec}^{\text{EU-CMA}}(\text{XMSS-T}; \xi, 2^h)$ describing the maximum success probability over all adversaries running in time $\leq \xi$ against the EU-CMA security of XMSS-T is bounded by

$$\begin{aligned} \text{InSec}^{\text{EU-CMA}}(\text{XMSS-T}; \xi) &\leq \text{InSec}^{\text{PRF}}(\mathcal{F}_n^1; \xi) + \text{InSec}^{\text{PRF}}(\mathcal{F}_m; \xi) \\ &\quad + \max\{\text{InSec}^{\text{M-eTCR}}(\mathcal{H}; \xi), 2\text{InSec}^{\text{MM-SPR}}(F; \xi), \text{InSec}^{\text{MM-SPR}}(H; \xi)\} \end{aligned}$$

The general idea of the proof follows that of previous hash-based schemes. There are a few mutually exclusive cases what could have happened if an adversary succeeded. First, the attacker could have broken the m-eTCR property of \mathcal{H} . This case is easily detected and can be handled in a straight-forward manner. Otherwise, the message digests have to differ. In this case the adversary has found a second preimage or a preimage for F or H with high probability. To extract a second preimage in this case, the reduction takes one MM-SPR challenge (M, K) per hash function call (H and F). Then, for this call the function is keyed with K and the bitmask is selected such that the input to the hash function is M . This means, if the input before the XOR with the bitmask is X , we use $X \oplus M$ as bitmask. Then the RO is programmed such that it generates this bitmask and key for this hash function call. This programming is done adaptively, only when the adversary queries the RO for a value. Now, any second preimage in the scheme will be a valid solution for MM-SPR of either H or F . The full proof can be found in Appendix B.

6 Implementation

In Section 4, we described XMSS-T, which builds on XMSS^{MT} , altering the functions that are used to construct WOTS chains and hash trees. We now examine the cost of this change in terms of computation time. In order to measure the cost of the additional bitmasks and keys that are required for each application of the functions F and H , we have implemented and benchmarked XMSS-T and XMSS^{MT} . We use the BDS tree traversal algorithm [13] to speed up the authentication path computation, making the scheme practical.

We examine the scheme for two parameter configurations from the current Internet Draft for XMSS^{MT} [22], obtaining measurements for both a single-tree and a multi-tree set-up. For both settings, we use $w = 16$ and $m = n = 256$.

For the first benchmark, we set $h = 20, d = 1$. We use the same subtree height for the second configuration, setting $h = 60, d = 3$ to construct three layers of subtrees with a height of twenty nodes each. We set $k = 2$ as the BDS parameter for both parameter sets, we rely on the SHA 256 function to construct F and H, and use ChaCha20 as the pseudorandom generator. These choices are also in accordance with [22]. For more parameter sets see [22].

For XMSS these parameters lead to a security level of 190 bits classical and 95 bits quantum for $h = 60$ (230 and 115 for $h = 20$). Following the security analysis in the last section and the lower bounds in Section 2, these parameters have a security level of more than 256 bits classical, and 128 bits quantum (assuming that each hash query requires more than 4 bit operations) for XMSS-T (without the message digest). With the message digest we get approximately 190 bits classical and 95 bits quantum like for XMSS as M-ETCR is still vulnerable to multi-target attacks. However, this can be fixed by increasing just the message digest size to $m = 276$ for $h = 20$ and $m = 316$ for $h = 60$. With this change we get 256 bit classical and 128 bit quantum security. However, to get an insight into the effects of the changes made in XMSS-T, we decided to run the experiments for the exact same parameters ($m = 256$). For more benchmarks with different parameters see the full version.

To carry out these benchmarks, we have used a single core of an Intel Core i7-4770K CPU, running at 3.5 GHz, although the implementation was not optimized specifically for this platform. In the first setting, with $h = 20$, we measure an average signing time of 12 488 458 clock cycles per signature for XMSS, and 34 862 033 clock cycles for XMSS-T. For the multi-tree version ($h = 60$), the scheme takes 13 014 401 clock cycles per signature for XMSS, and 37 025 552 cycles for XMSS-T.

This difference is quite significant. However, this was to be expected as the running time of the scheme is largely dominated by applications of F and H – precisely the functions that are changed for XMSS-T. For plain XMSS with the aforementioned parameters, these functions merely consist of calls to SHA-256 with inputs of 256 and 512 bits, respectively. Each of these inputs fits within the internal block size of SHA-256 (512 bits). When considering the Merkle-Damgård construction [26] that defines the structure of SHA-256, this implies a single application of the internal compression function. When transforming F and H into keyed hash functions, the input length increases. To ensure that the key and the input are in separate blocks, the key is prefixed with 256 zero-bits. This results in inputs of 768 and 1024 bits, respectively, implying the need for two blocks, as well as two applications of the compression function. The straightforward calls to SHA-256 for F and H run in 1 072 and 1 924 cycles, while the keyed variants take 1 932 and 2 812 cycles, respectively.

A bigger factor weighing down F and H is the time needed to generate the keys and bitmasks pseudorandomly. Both these values require calls to the pseudorandom generator. For F, we require two output blocks of 256 bits each; H requires three. At an expense of 560 cycles per output block, generating randomness for the masks and keys carries a significant cost.

Altogether, the experiments show that the tightened security comes at the cost of a factor less than 3 increase in the runtime.

References

1. Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004. [4](#)
2. Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. [4](#)
3. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald De Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001. [4](#)
4. Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. [9](#)
5. Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Peter Schwabe, and Zooko Wilcox-O’Hearn. SPHINCS: practical stateless hash-based signatures. In Marc Fischlin and Elisabeth Oswald, editors, *EUROCRYPT 2015*, volume 9056 of *LNCS*, pages 368–397. Springer Berlin Heidelberg, 2015. [2](#), [3](#)
6. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In DongHoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer Berlin Heidelberg, 2011. [7](#)
7. Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *arXiv preprint quant-ph/9605034*, 1996. [8](#)
8. Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002. [8](#)
9. Gilles Brassard, Peter Hoyer, and Alain Tapp. Quantum algorithm for the collision problem. *arXiv preprint quant-ph/9705002*, 1997. [4](#)
10. Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum counting. In *Automata, Languages and Programming*, pages 820–831. Springer, 1998. [8](#)
11. Johannes Buchmann, Erik Dahmen, Sarah Ereth, Andreas Hülsing, and Markus Rückert. On the security of the Winternitz one-time signature scheme. In A. Nitaj and D. Pointcheval, editors, *Africacrypt 2011*, volume 6737 of *LNCS*, pages 363–378. Springer Berlin / Heidelberg, 2011. [2](#), [13](#)
12. Johannes Buchmann, Erik Dahmen, and Andreas Hülsing. XMSS - a practical forward secure signature scheme based on minimal security assumptions. In Bo-Yin Yang, editor, *Post-Quantum Cryptography 2011*, volume 7071 of *LNCS*, pages 117–129. Springer Berlin / Heidelberg, 2011. [2](#), [3](#), [13](#)
13. Johannes Buchmann, Erik Dahmen, and Michael Schneider. Merkle tree traversal revisited. In Johannes Buchmann and Jintai Ding, editors, *Post-Quantum Cryptography*, volume 5299 of *LNCS*, pages 63–78. Springer Berlin / Heidelberg, 2008. [17](#), [21](#)
14. J Lawrence Carter and Mark N Wegman. Universal classes of hash functions. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 106–112. ACM, 1977. [9](#)
15. Erik Dahmen, Katsuyuki Okeya, Tsuyoshi Takagi, and Camille Vuillaume. Digital signatures out of second-preimage resistant hash functions. In Johannes Buchmann

- and Jintai Ding, editors, *Post-Quantum Cryptography 2008*, volume 5299 of *LNCS*, pages 109–123. Springer, 2008. [2](#), [13](#), [16](#)
16. Edward Eaton and Fang Song. Making existential-unforgeable signatures strongly unforgeable in the quantum random-oracle model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015, May 20-22, 2015, Brussels, Belgium*, pages 147–162, 2015. [11](#)
 17. Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988. [19](#)
 18. Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996. [8](#)
 19. Shai Halevi and Hugo Krawczyk. Strengthening digital signatures via randomized hashing. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 41–59. Springer Berlin / Heidelberg, 2006. [2](#)
 20. Shai Halevi and Hugo Krawczyk. Strengthening digital signatures via randomized hashing. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 41–59. Springer Berlin / Heidelberg, 2006. [6](#)
 21. Andreas Hülsing. W-OTS+ shorter signatures for hash-based signature schemes. In Amr Youssef, Abderrahmane Nitaj, and AboulElla Hassanien, editors, *Africacrypt 2013*, volume 7918 of *LNCS*, pages 173–188. Springer, 2013. [2](#), [14](#)
 22. Andreas Hülsing, D. Butin, S. Gazdag, and A. Mohaisen. Xmss: Extended hash-based signatures draft-irtf-cfrg-xmss-hash-based-signatures-01. Crypto Forum Research Group Internet-Draft, 2015. <https://tools.ietf.org/html/draft-irtf-cfrg-xmss-hash-based-signatures-01>. [21](#), [22](#)
 23. Andreas Hülsing, Lea Rausch, and Johannes Buchmann. Optimal parameters for XMSS^{MT}. In Alfredo Cuzzocrea, Christian Kittl, Dimitris E. Simos, Edgar Weippl, and Lida Xu, editors, *Security Engineering and Intelligence Informatics*, volume 8128 of *LNCS*, pages 194–208. Springer Berlin Heidelberg, 2013. [2](#), [3](#), [13](#), [17](#)
 24. A Joffe et al. On a set of almost deterministic k -independent random variables. *The Annals of Probability*, 2(1):161–162, 1974. [9](#)
 25. Howard Karloff and Yishay Mansour. On construction of k -wise independent random variables. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 564–573. ACM, 1994. [9](#)
 26. Ralph Charles Merkle. *Secrecy, authentication, and public key systems*. PhD thesis, Stanford University, 1979. [22](#)
 27. Ilya Mironov. Collision-resistant no more: Hash-and-sign paradigm revisited. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography - PKC 2006*, volume 3958 of *LNCS*, pages 140–156. Springer Berlin / Heidelberg, 2006. [2](#)
 28. Dominique Unruh. Quantum position verification in the random oracle model. In *CRYPTO 2014*, volume 8617 of *LNCS*, pages 1–18. Springer, August 2014. Preprint on IACR ePrint 2014/118. [11](#)
 29. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer Berlin Heidelberg, 2012. [9](#)
 30. Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Information and Computation*, 15(7&8), 2015. [4](#)

A Classical generic security

Preimage resistance. For preimage resistance, analysis shows that the success probability of any classical \mathcal{A} that makes q queries to its oracle is

$$\text{Succ}_{\mathcal{H}_n}^{\text{OW}}(\mathcal{A}) = \left(\frac{q+1}{2^n} \right), \quad (11)$$

where the probability is taken over the internal coins of the oracle and the random choices of K and M . An attacker that makes no query but simply outputs a random domain element has success probability 2^{-n} of hitting the target Y . An attacker that makes one query can verify the first guess. If that one did not hit Y he can make another guess which he now can not verify anymore. Together this gives a success probability of $2/2^n$. Iterating this gives the above bound. Consequently, an attacker needs $\mathcal{O}(2^n)$ queries to reach a success probability of at least 0.5.

Single-function multi-target preimage resistance. For SM-OW a similar analysis shows a bound of

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{SM-OW}}(\mathcal{A}) = \left(\frac{(q+1)p}{2^n} \right), \quad (12)$$

The reason is that the success probability of a single guess is now $p/2^n$. Otherwise, the argument follows along the lines of the above argument. Consequently, the query complexity of a successful attack is $\mathcal{O}(2^n/p)$. Please note, we also can get this result using a reduction from OW. In this case, we replace a random Y_i by the given Y from the OW game. The reduction loses a factor $1/p$.

Multi-function multi-target preimage resistance. While the previous cases are more or less known results, for MM-OW we are not aware of any such results. The difference to SM-OW is that the adversary now basically plays p independent OW games at once. In contrast to the OW game \mathcal{A} can not use a query he made to attack Y_i for any other Y_j for $j \neq i$. The reason is that different functions are associated to the different Y_i . So, in the classical case we get a query bound of

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{MM-OW}}(\mathcal{A}) = \left(\frac{q+1}{2^n} \right), \quad (13)$$

The reason is that a guess has success probability $1/2^n$. As one also has to guess (K_i, Y_i) , every verification query can only check if a given M fulfills $Y_i = H_{K_i}(M)$ for a single i . Viewed differently, each query has to fix K_i in advance and outputs the associated Y_i only with probability 2^{-n} . Consequently, we get the same query bound $\mathcal{O}(2^n)$ as for OW.

Second-preimage resistance. In the case of second-preimage resistance, the success probability of any \mathcal{A} that makes q queries to its oracle is

$$\text{Succ}_{\mathcal{H}_n}^{\text{SPR}}(\mathcal{A}) = \left(\frac{q+1}{2^n} \right), \quad (14)$$

where the probability is taken over the internal coins of the oracle and the random choices of K and M . The bound can easily be derived following the analysis for one-wayness. Consequently, an attacker needs $\mathcal{O}(2^n)$ queries to reach a success probability of at least 0.5.

Single-function multi-target second-preimage resistance. For SM-SPR a similar analysis shows a bound of

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{SM-SPR}}(\mathcal{A}) = \left(\frac{(q+1)p}{2^n} \right). \quad (15)$$

Again, the analysis follows along the lines of the respective analysis for SM-OW. Consequently, the query complexity of a successful attack is $\mathcal{O}(2^n/p)$.

Multi-function multi-target preimage resistance. While the previous cases are more or less known results, for MM-SPR we are not aware of any such results. The difference to SM-SPR is that the adversary now basically plays p independent SPR games at once. As for MM-OW, in the classical case, we get a query bound of

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{MM-SPR}}(\mathcal{A}) = \left(\frac{q+1}{2^n} \right). \quad (16)$$

Again, the analysis follows along the lines of the respective analysis for MM-OW. Consequently, the query complexity of a successful attack is $\mathcal{O}(2^n)$.

Extended target collision resistance. For eTCR, analysis shows that the success probability of any adversary \mathcal{A} that makes no more than q queries to its oracle is

$$\text{Succ}_{\mathcal{H}_n}^{\text{eTCR}}(\mathcal{A}) \leq \left(\frac{q+1}{2^n} + \frac{q}{2^k} \right), \quad (17)$$

where the probability is taken over the internal coins of the oracle and the random choice of K .

Consider an arbitrary adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ attacking eTCR. \mathcal{A}_1 makes q_1 queries and outputs a message M . Afterwards \mathcal{A}_2 obtains K and makes q_2 queries, with $q_1 + q_2 = q$. Without loss of generality, we assume that \mathcal{A}_1 stores all his query results in the shared memory. When \mathcal{A}_2 receives K , we can distinguish two mutually exclusive cases:

Case 1: \mathcal{A}_1 already queried the oracle for $H_K(M)$. To simplify analysis, we consider this a success for \mathcal{A} . As K is a random key and \mathcal{A}_1 made queries for no more than q_1 different keys, this case occurs with probability

$$\epsilon_1 \leq \frac{q_1}{2^k}.$$

Case 2: \mathcal{A}_1 did not query $H_K(M)$ before. In this case, every query made by \mathcal{A}_1 and every query made by \mathcal{A}_2 has probability 2^{-n} to hit $H_K(M)$ and hence to be a solution. If \mathcal{A}_2 does not find a solution using all query results, he

can make another guess that has the same success probability as the queries before. Hence, the success probability in this case is exactly

$$\epsilon_2 = \frac{q_1 + q_2}{2^n}.$$

The sum of the two bounds $\epsilon_1 + \epsilon_2$ takes its maximum for $q_1 = q$. This gives the claimed bound. Note that the analysis for case 1 above is very rough and could be tightened (This is only a success if \mathcal{A}_1 already found a pseudo-collision for (K, M)). However, in all relevant cases we know $k \gtrsim n$ and hence tightening is of little use.

Multi-target-ETCR. Now, switching to M-ETCR the complexities for the two cases change as follows: In both cases a factor of q is lost. In Case 1 this is caused by the fact that there are now q keys returned (over the game) that might hit a previously queried one. In Case 2 this is caused by the fact that each query works for all q targets (as in the case of MM-SPR). This leads to the bound

$$\text{Succ}_{\mathcal{H}_{n,p}}^{\text{M-ETCR}}(\mathcal{A}) = \frac{(q+1)p}{2^n} + \frac{qp}{2^k}.$$

B Proof of Theorem 2

In the following we omit indices for the MM-SPR challenge pairs to preserve readability. Assume that there exists an adversary \mathcal{A} running in time ξ that breaks the EU-CMA security of XMSS with probability $\epsilon_{\mathcal{A}}$. In the following we will prove that $\epsilon_{\mathcal{A}} \leq \text{InSec}^{\text{EU-CMA}}(\text{XMSS-T}; \xi, 2^h)$. First, consider the following two games:

Game 1. This is the original game.

Game 2. This is the same as Game 1 but instead of using random elements from \mathcal{F}_n^1 and \mathcal{F}_m (by sampling \mathcal{S} and SK_2 from the key space), two truly random functions $G_n : \{0, 1\}^n \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $G_m : \{0, 1\}^n \times \{0, 1\}^* \rightarrow \{0, 1\}^m$ are used.

The difference in the success probability of \mathcal{A} playing one of these games must be bound by $\text{InSec}^{\text{PRF}}(\mathcal{F}_n^1; \xi) + \text{InSec}^{\text{PRF}}(\mathcal{F}_m; \xi)$, otherwise we could use \mathcal{A} to distinguish \mathcal{F}_n^1 or \mathcal{F}_m from a truly random function, breaking the pseudorandomness which would contradict the assumption. Hence, it suffices to analyze the success probability of \mathcal{A} in Game 2. Towards this end, we construct an oracle machine $\mathcal{M}^{\mathcal{A}}$ that breaks either the multi-function multi-target second-preimage resistance of F or of H, or the extended target collision resistance of \mathcal{H} . $\mathcal{M}^{\mathcal{A}}$ takes q_1 challenge pairs $\{(K_i, M_i)\}_1^{q_1}$ for F and q_2 challenge pairs $\{(K'_i, M'_i)\}_1^{q_2}$ for H where q_1 (q_2) is the number of calls to F (H) for the XMSS-T key pair. Each of these challenge pairs gets associated with one specific call to F or H within the XMSS-T key pair.

$\mathcal{M}^{\mathcal{A}}$ first samples a random seed $\text{SEED} \xleftarrow{\$} \{0, 1\}^n$. The XMSS-T public key PK becomes $(\text{PK}_1 = \text{H}_{K'_j}(M'_j), \text{SEED})$ for (K'_j, M'_j) – the pair associated with the

call to H that computes the root. Now \mathcal{A} is run on this PK. When \mathcal{A} makes his i th query using some message Msg_i , $\mathcal{M}^{\mathcal{A}}$ first sends Msg_i to the m-ETCR challenger, receiving back a function key R_i . Then it computes the message digest as $D_i = \mathcal{H}(R_i, \text{Msg}_i)$ and sets the signature index to i . The next steps are the same for each tree involved in the signature. First, $\mathcal{M}^{\mathcal{A}}$ computes the chain indices b . The WOTS⁺ signature is collected by selecting the challenge (K, M) for the b_j th call to F in the j th chain and computing the j th signature element as $F_K(M)$. Similarly, the authentication path for the WOTS⁺ key pair is generated by figuring out the nodes that are required. Then, these nodes are calculated as $H_K(M)$ where (K, M) is the challenge pair associated with the call to H that computes this authentication path node. The same is done for the root node. Afterwards, the whole procedure is repeated for the parent tree, until the top tree is done. Then the XMSS-T signature Σ_i is sent back to \mathcal{A} .

If \mathcal{A} queries the RO, the RO is adaptively programmed such that it outputs the right bitmasks and keys. W.l.o.g., we assume that \mathcal{A} makes each query only once and stores the result. For any query $(X_1 || X_2) \in \{0, 1\}^n \times \{0, 1\}^*$ where $X_1 \neq \text{SEED}$ or X_2 is no valid address of a bitmask or a key for a hash function call, the RO simply outputs a random value. Otherwise, assume for simplicity (again w.l.o.g.) that the adversary always queries the RO for the bitmask(s) and the key of a hash function call at once. Then the RO is programmed as follows: Assume (K, M) is the challenge pair for the associated hash call. If the address refers to the first hash call in a WOTS⁺ function chain, the start value of that chain is generated as $X = G(\text{SK}_1, \mathbf{a})$ where \mathbf{a} is the address of that function chain. Next, the RO is programmed to return bitmask $r = X \oplus M$ and key K . If the address refers to a hash call in a WOTS⁺ function chain that is not the first in that chain, let (K', M') be the challenge template associated with the previous hash call in that chain. Then the RO is programmed to return bitmask $r = F_{K'}(M') \oplus M$ and key K . Finally, if the address refers to a hash call to H, let $(K_1, M_1), (K_2, M_2)$ be the challenge templates associated with the hash calls computing its two child nodes. Then the RO is programmed to return bitmask $r = (H_{K_1}(M_1) || H_{K_2}(M_2)) \oplus M$ and key K (Note that the challenge templates might be associated with calls to F if the address is associate with the computation of a node on level 1 in an L-tree. In this case H is replaced by F in the computation of r).

When \mathcal{A} outputs a forgery (Msg, Σ) with $\Sigma = (i, R, \sigma_{W,0}, \text{Auth}_{i_0}, \dots, \sigma_{W,d-1}, \text{Auth}_{i_{d-1}})$, $\mathcal{M}^{\mathcal{A}}$ runs the verification algorithm on (Msg, Σ) and (Σ_i, Msg_i) . If the forgery is invalid, $\mathcal{M}^{\mathcal{A}}$ returns \perp . Otherwise, three mutually exclusive cases can occur. $\mathcal{M}^{\mathcal{A}}$ compares the values computed during the two verification runs in order of computation.

Case 1: If $D = \mathcal{H}(R, \text{Msg}) = \mathcal{H}(R_i, \text{Msg}_i) = D_i$, i.e., if the digests of the i th query is the same as that of the forgery, $\mathcal{M}^{\mathcal{A}}$ broke m-eTCR and returns (i, R, Msg) . Hence, the probability that \mathcal{A} outputs a case 1 forgery must be upper bounded by $\text{InSec}^{\text{m-eTCR}}(\mathcal{H}; \xi)$ per assumption.

If the digests are different, the corresponding b_i are also different and hence, parts of the data computed by the two verification runs must also differ. Now,

$\mathcal{M}^{\mathcal{A}}$ only compares the computed WOTS⁺ public keys and the computed root values. By the pigeonhole principle, the signatures have to agree on one of these for the first time as they lead to the same root of the top tree.

Case 2: If the data generated verifying the two signatures first agrees on a WOTS⁺ public key, the message digests or the root nodes signed with this WOTS⁺ keypair where different. Hence, we got a WOTS⁺ forgery. In this case, by the construction of the checksum there must be one chain j in this WOTS⁺ keypair such that $b_j < (b_j)_i$, i.e. the j th signature value of the forgery belongs to an *earlier* hash call than the one of the answer to the i th query. As both chains end in the same public key value, they must collide at the output of some call to F. If this point is not the $(b_j)_i$ th call it has to be a later one. In this case, the input to the colliding call to F computed from the forgery is a second preimage for the challenge template associated with that call to F and $\mathcal{M}^{\mathcal{A}}$ outputs it, breaking MM-SPR of F. Otherwise, the two chains collide on the output of the $(b_j)_i$ th call to F, i.e., on $(\sigma_j)_i$, the j th value of the original signature. Let (K, M) be the challenge pair associated with the call to F that produced $(\sigma_j)_i$. According to Eqn. 10, $(\sigma_j)_i$ has at least two preimages under F_k . As \mathcal{A} has no information about the preimage, the value X that can be computed from the forgery and that leads $f_K(X) = (\sigma_j)_i$ is unequal to M with at least probability 1/2. In that case, $\mathcal{M}^{\mathcal{A}}$ found a second preimage of M under F_K and outputs it. Otherwise it returns \perp . Consequently, the probability that \mathcal{A} outputs a case 2 forgery must be upper bounded by $2\text{InSec}^{\text{MM-SPR}}(\text{F}; \xi)$ per assumption.

Case 3: If the data generated verifying the two signatures first agrees on a root node, the WOTS⁺ public keys that are used to compute this root node have to differ. A third time by the pigeonhole principle, there must be one call to H between the WOTS⁺ public key and the root node where the output for the forgery and the correct signature agree for the first time. As the input data depends on previously computed outputs of H (or F), it must differ. Hence, for challenge pair (K, M) , the input to this call to H_K is a second preimage for M , that $\mathcal{M}^{\mathcal{A}}$ returns breaking MM-SPR of H. Hence, the probability that \mathcal{A} outputs a case 3 forgery must be upper bounded by $\text{InSec}^{\text{MM-SPR}}(\text{H}; \xi)$.

Combining the upper bounds from the three cases shows that the success probability $\epsilon_{\mathcal{A}}$ of \mathcal{A} winning in Game 2 must be upper bounded by

$$\epsilon_{\mathcal{A}} \leq \max\{\text{InSec}^{\text{M-ETCR}}(\mathcal{H}; \xi), 2\text{InSec}^{\text{MM-SPR}}(\text{F}; \xi), \text{InSec}^{\text{MM-SPR}}(\text{H}; \xi)\}.$$

Combining this with the result that the difference in \mathcal{A} 's success probability between playing in Game 1 and playing in Game 2 must be upper bounded by $\text{InSec}^{\text{PRF}}(\mathcal{F}_n^2; \xi)$, we get the claimed bound on the success probability of any adversary \mathcal{A} running in time ξ :

$$\begin{aligned} \text{Succ}_{\text{XMSS-T}}^{\text{EU-CMA}}(\mathcal{A}) &\leq \text{InSec}^{\text{PRF}}(\mathcal{F}_n^1; \xi) + \text{InSec}^{\text{PRF}}(\mathcal{F}_m; \xi) \\ &\quad + \max\{\text{InSec}^{\text{M-ETCR}}(\mathcal{H}; \xi), 2\text{InSec}^{\text{MM-SPR}}(\text{F}; \xi), \text{InSec}^{\text{MM-SPR}}(\text{H}; \xi)\} \end{aligned}$$

□