Optimal Parameters for XMSS^{MT}

Andreas Hülsing[⋆], Lea Rausch, and Johannes Buchmann

Cryptography and Computeralgebra
Department of Computer Science
TU Darmstadt, Germany
huelsing@cdc.informatik.tu-darmstadt.de

Abstract. We introduce Multi Tree XMSS (XMSS^{MT}), a hash-based signature scheme that can be used to sign a virtually unlimited number of messages. It is provably forward and hence EU-CMA secure in the standard model and improves key and signature generation times compared to previous schemes. XMSS^{MT} has — like all practical hash-based signature schemes — a lot of parameters that control different trade-offs between security, runtimes and sizes. Using linear optimization, we show how to select provably optimal parameter sets for different use cases.

Keywords hash-based signatures, parameter selection, linear optimization, forward secure signatures, implementation

1 Introduction

Digital signatures are among the most important cryptographic primitives in practice. They are used to secure communication protocols like SSL/TLS and SSH, software updates, or to replace handwritten signatures. Hash-based signature schemes are an interesting alternative to the signature schemes used today. Not only because they are assumed to resist quantum computer aided attacks, but also because of their fast signature generation and verification times as well as their strong security guarantees. The latest hash-based signature schemes [BDH11, HBB13] come with a standard model security proof and outperform RSA in many settings regarding runtimes.

Practical hash-based signature schemes have a lot of parameters that control several trade-offs between runtimes and sizes. Hence, an important open problem on the way of making hash-based signatures practical is parameter selection. In previous works [BDH11, HBB13, BDK+07, BGD+06, DOTV08] it was shown how to select secure parameters using the exact security reductions, but this still leaves too many possible parameter choices to apply a brute-force search. Even applying reasonable restrictions on the parameters, the search space can easily grow to the order of 2^{80} and more. In this work we introduce Multi Tree XMSS (XMSS^{MT}), a new hash-based signature scheme that contains the latest hash-based signature schemes XMSS and XMSS$^+$ as special cases and propose a systematic way to select the optimal parameter set for a given scenario for XMSS^{MT}.

XMSS^{MT} improves XMSS in a way, that it allows for a virtually unlimited number of signatures, while it preserves its desirable properties. XMSS^{MT} has minimal security requirements, i.e. it can be proven EU-CMA secure in the standard model, when instantiated with a pseudorandom function family and a second preimage resistant hash function. Moreover, XMSS^{MT} is forward secure [BM99] and introduces a new trade-off between signature and key generation time on the one hand and signature size on the other hand. XMSS^{MT} can be viewed as applying the tree

---

chaining idea introduced in [BGD+06] to XMSS and combine it with the improved distributed signature generation proposed in [HBB13].

For XMSS^{MT}, we show how to model the parameter selection problem as a linear optimization problem. A straight forward approach would lead many non-linear constraints. To overcome this problem, we use the generalized lambda method [Mor07] for linearization. The resulting model can then be solved using the Simplex algorithm [Dan63]. As the general lambda method is exact, the Simplex algorithm outputs a provably optimal solution under the given constraints. We made the model flexible, such that it can be used for different use cases with different requirements using different inputs. We present results for demonstrative use cases and compare them to possible parameter choices from previous work. However, the code is available on the corresponding authors home page and can be used to select optimal parameters for other use cases. Linear optimization was used before in [BDK+07] to find optimal parameters. Unfortunately, the results given there do not carry over to the new hash-based signature schemes, as less parameters were taken into account and the authors of [BDK+07] do not provide details about how to model the problem.

Organization. We start with an introduction to XMSS^{MT} in Section 2, where we also discuss theoretical runtimes and parameter sizes as well as security and correctness of the scheme. In Section 3 we show how to model the parameter selection problem for XMSS^{MT} as a linear optimization problem and present optimal parameters for exemplary use cases in Section 4. We discuss our results and draw a conclusion in Section 5.

2 Multi Tree XMSS - XMSS^{MT}

In this section we introduce XMSS^{MT}. For a better understanding, we first give a brief description of the scheme. A hash-based signature scheme starts with a one-time signature scheme (OTS). To obtain a many time signature scheme, many instances of the OTS are used and their public key are authenticated using a binary hash tree, called Merkle-Tree. The root of the tree is the public key of the scheme and the secret key is the seed of a pseudorandom generator that is used to generate the OTS secret keys. A signature contains an index, the OTS signature, and the so called authentication path which contains the siblings of the nodes on the way from the used OTS key pair to the root. For a picture of an authentication path see Figure 1. For XMSS^{MT} we extend this using many layers of trees. The trees on the lowest layer are used to sign the messages and the trees on higher layers are used to sign the roots of the trees on the layer below. The public key contains only the root of the tree on the top layer. A signature consists of all the signatures on the way to the highest tree. A graphical representation of the scheme can be found in Appendix B. In the following we describe the construction in detail, starting with the description of some building blocks. Afterwards we describe the algorithms of the scheme and present an theoretical analysis of its performance, security and correctness.

Parameters. For security parameter $n \in \mathbb{N}$, the basic building blocks of XMSS^{MT} are a pseudorandom function family $\mathcal{F}_n = \{F_K : \{0,1\}^n \to \{0,1\}^n | K \in \{0,1\}^n\}$, and a second preimage resistant hash function $H : \{0,1\}^{2n} \to \{0,1\}^n$. Further parameters are the number of layers $d \in \mathbb{N}$, the binary message length m and one parameter set per layer. For a layer $0 \leq i \leq d-1$ a parameter set contains the tree height $h_i \in \mathbb{N}$, the so called BDS parameter $k_i \in \mathbb{N}$ with the restrictions $k_i < h_i$ and $h_i - k_i$ is even and the Winternitz parameter $w_i \in \mathbb{N}, w_i \geq 2$. To enable improved distributed signature generation we require $(h_{i+1} - k_{i+1} + 6)/2 \leq 2^{h_i - k_i + 1}$ for $0 \leq i < d-1$, as

**Fig. 1.** The authentication path for a leaf i.

well as $(h_0 - k_0)/2 \geq d - 1$. A XMSS^{MT} key pair can be used to sign $h = \sum_{i=0}^{d-1} h_i$ messages of m bits. The remaining parameters define the various trade-offs explained below. These parameters are publicly known.

Winternitz OTS. XMSS^{MT} uses the Winternitz-OTS (W-OTS) from [BDE^{+}11]. W-OTS uses the function family $\mathcal{F}_n$ and a value $X \in \{0,1\}^n$ that is chosen during XMSS key generation. Besides message length m and Winternitz parameter w it has another parameter $\ell = \ell(m, w)$ that is a function of m and w. The W-OTS secret key, public key and signatures consist of ℓ bit strings of length n. An important property of W-OTS is that the public key can be computed from a valid signature. The sizes of signature, public, and secret key are ℓn bits. The runtimes for key generation, signature generation and signature verification are all bounded by $\ell(w-1)$ evaluations of elements from $\mathcal{F}_n$. The Winternitz parameter w controls a time - space trade-off, as ℓ shrinks logarithmically in w. For more detailed information see Appendix A.

XMSS Tree. XMSS^{MT} uses the XMSS tree construction, originally proposed in [DOTV08]. An XMSS tree of height h is a binary tree with $h+1$ levels. The leaves are on level 0, the root on level h. The nodes on level j, are denoted by $N_{i,j}$, for $0 \leq i < 2^{h-j}, 0 \leq j \leq h$. To construct the tree, h bit masks $B_j \in \{0,1\}^{2n}$, $0 < j \leq h$ and the hash function H, are used. $N_{i,j}$, for $0 < j \leq h$, is computed as $N_{i,j} = \mathrm{H}((N_{2i,j-1}||N_{2i+1,j-1}) \oplus B_j)$. Figure 2 shows the construction of the XMSS tree.

Fig. 2. The XMSS tree construction

Leaf Construction. The leaves of a XMSS tree are the hash values of W-OTS public keys. To avoid the need of a collision resistant hash function, another XMSS tree is used, called L-tree. The ℓ leaves of an L-tree are the ℓ bit strings $(\mathsf{pk}_0, \ldots, \mathsf{pk}_\ell)$ from the corresponding public key. If ℓ is not a power of 2, 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. For the L-tree, the same hash function as above but $\lceil \log \ell \rceil$ new bitmasks are used. These bitmasks are the same for all L-trees.

Pseudorandom Generator. The W-OTS key pairs belonging to one XMSS tree are generated using two pseudorandom generators (PRG) build using $\mathcal{F}$. The stateful forward secure PRG FsGen : $\{0,1\}^n \to \{0,1\}^n \times \{0,1\}^n$ is used to generate one seed value per W-OTS keypair. Then the seed is expanded to the ℓ W-OTS secret key bit strings using $\mathcal{F}_n$. FsGen starts from a uniformly random state $S_0 \xleftarrow{\$} \{0,1\}^n$. On input of a state S_i, FsGen generates a new state S_{i+1} and a pseudorandom output R_i: $\mathrm{FsGen}(S_i) = (S_{i+1}||R_i) = (\mathrm{F}_{S_i}(0)||\mathrm{F}_{S_i}(1))$. The output R_i is used to generate the ith W-OTS secret key $(\mathsf{sk}_1, \ldots, \mathsf{sk}_\ell)$ as $\mathsf{sk}_j = \mathrm{F}_{R_i}(j-1), 1 \le j \le \ell$.

Tree Traversal - The BDS *Algorithm.* To compute the authentication paths for the W-OTS key pairs of one XMSS tree, XMSS^{MT} uses the BDS algorithm [BDS08]. The choice of this algorithm is already a time - memory trade-off. Instead of computing every node of the authentication path when needed, BDS uses some storage to decrease the worst case runtime. The BDS algorithm uses the $\mathrm{TreeHash}$ algorithm (see Algorithm 1) as a subroutine. The BDS algorithm reduces the

---

Algorithm 1 $\mathrm{TreeHash}$

Input: Stack Stack, node N_1
Output: Updated stack Stack

1. **While** top node on Stack has same height as N_1 **do**
   (a) $t \longleftarrow N_1.height() + 1$
   (b) $N_1 \longleftarrow \mathrm{H}\left((\mathsf{Stack}.pop()||N_1) \oplus B_t\right)$
2. $\mathsf{Stack}.push(N_1)$
3. **Return** Stack

---

worst case signing time from $2^h - 1$ to $(h-k)/2$ leaf computations and evaluations of $\mathrm{TreeHash}$. More specifically, the BDS algorithm does three things. First, it uses the fact that a node that is a left child can be computed from values that occurred in an authentication path before, spending only one evaluation of H. Second, it stores the right nodes from the top k levels of the tree during key generation. Third, it distributes the computations for right child nodes among previous authentication path computations. It schedules one instance of $\mathrm{TreeHash}$ per tree level. The computation of the next right node on a level starts, when the last computed right node becomes part of the authentication path. BDS uses a state $\mathsf{State}_{\mathrm{BDS}}$ of $2(h-k)$ states of FsGen and at most $\left(3h + \lfloor \frac{h}{2} \rfloor - 3k - 2 + 2^k\right)$ tree nodes that is initialized during key generation. To compute the authentication paths, the BDS algorithm spends $(h-k)/2$ leaf computations and evaluations of $\mathrm{TreeHash}$ as well as $(h-k)$ calls to F to update its state per signature. For more details see [BDS08].

Key Generation. The XMSS^{MT} key generation algorithm takes as input all of the above parameters. First, the $\max_{0 \le i \le d-1}\{h_i + \lceil \log \ell_i \rceil\}$ bitmasks and X are chosen uniformly at random. The same

4

bitmasks and X are used for all layers. Then, the root of the first XMSS tree on each layer is computed. This is done in an ordered way, starting from layer 0. For the tree TREE_i on layer i the initial state of FsGen, $S_{0,i}$ is chosen uniformly at random and a copy of it is stored as part of the secret key SK. The tree is constructed using the TREEHASH algorithm above. Starting with an empty stack Stack_i and $S_{0,i}$, all 2^{h_i} leaves are successively generated and used as input to the TREEHASH algorithm to update Stack_i. This is done updating S_i and generating the W-OTS key pairs using its output. The W-OTS public key is then used to compute the corresponding leaf using an L-tree which in turn is used to update the current Stack_i using TREEHASH. Then the W-OTS key pair is deleted. After all 2^{h_i} leaves were processed by TREEHASH, the only value on Stack_i is the root ROOT_i of TREE_i. When ROOT_i $0 \leq i < d-1$, is computed, it is signed using the first W-OTS key pair of TREE_{i+1}, which is computed next. This signature σ_{i+1} can be extracted while TREE_i is generated and hence does not need any additional computation. Then σ_{i+1} is stored as part of SK. If the highest layer $d-1$ is reached, ROOT_{d-1} is stored in the public key PK. During the computation of ROOT_i, the state of the BDS algorithm $\text{State}_{\text{BDS},i}$ is initialized. For details see [BDS08].

Finally, the data structures for the next trees are initialized: For the next tree NEXT_i on each layer $0 \leq i < d-1$ a FsGen state $S_{n,i}$ is chosen uniformly at random and a new TREEHASH stack $\text{Stack}_{next,i}$ is initialized. Summing up, SK consists of the states $(S_{0,i}, \text{State}_{\text{BDS},i}), 0 \leq i \leq d-1$ and the $d-1$ signatures σ_i, $0 < i \leq d-1$. Additionally, it contains $d-1$ FsGen states $S_{n,i}$, $d-1$ TREEHASH stacks $\text{Stack}_{next,i}$ and $d-1$ BDS states $\text{State}_{\text{BDS},n,i}$ for the next trees NEXT_i on layer $0 \leq i < d-1$. The public key PK consists of the $\max_{0 \leq i \leq d-1}\{h_i + \lceil \log \ell_i \rceil\}$ bitmasks, the value X and ROOT_{d-1}.

Signature generation. The signature generation algorithm takes as input a m bit message M, the secret key SK, and the index i, indicating that this is the ith message signed with this keypair. The signature generation algorithm consists of two phases. First, M is signed. A XMSS^{MT} signature $\Sigma = (i, \sigma_0, \text{Auth}_0, \sigma_1, \text{Auth}_1, \ldots, \sigma_d, \text{Auth}_d)$ contains the index i, the W-OTS signature σ_0 on the message M, the corresponding authentication path for TREE_0 and the W-OTS signatures on the roots of the currently used trees together with the corresponding authentication paths. The only thing that has to be computed is σ_0 — the W-OTS signature on message M using the ith W-OTS key pair on the lowest layer. All authentications paths and the W-OTS signatures on higher layers are already part of the current secret key.

The second phase is used to update the secret key. Therefor BDS is initialized with $\text{State}_{\text{BDS},0}$ and receives $(h_0 - k_0)/2$ updates. If not all of these updates are needed to update $\text{State}_{\text{BDS},0}$, i.e. all scheduled TREEHASH instances are finished and there are still updates left, the remaining updates are used for the upper trees. On the upper layers, not only the BDS state has to be updated, but while one leaf is used, one leaf in the next tree must be computed and $h_i - k_i$ FsGen states must be updated. This means that remaining updates from layer zero are first used to update $\text{State}_{\text{BDS},1}$. If all scheduled TREEHASH instances in $\text{State}_{\text{BDS},1}$ are finished, one update is used to compute a new leaf for NEXT_1 and to update $\text{Stack}_{next,1}$ afterwards. The next update is used to for the FsGen states. If layer 1 does not need anymore updates, the remaining updates are forwarded to layer 2 and so on, until either all updates are used or all tasks are done. Finally, one leaf of the next tree is computed and used as input for TREEHASH to update $\text{Stack}_{next,0}$.

A special case occurs if $i \mod 2^{h_0} = 2^{h_0} - 1$. In this case, the last W-OTS key pair of the current TREE_0 was used. This means that for the next signature a new tree is needed on every layer j with $i \mod 2^{h_j} = 2^{h_j} - 1$. For all these layers, $\text{Stack}_{next,j}$ already contains the root of

NEXT$_j$. So, TREE$_{j+1}$ is used to sign ROOT$_j$. Each signature is counted as one update. In case not all updates are needed, remaining updates are forwarded to the first layer that did not get a new tree. In SK, State$_{\text{BDS},j}$, S_j, and Σ_j are replaced by the newly computed data. Afterwards, new data structures for the next tree on layer j are initialized and used to replace the ones in SK. Finally, the signature SIG, the updated secret key SK and $i+1$ are returned.

Signature verification. The signature verification algorithm takes as input a signature $\Sigma = (i, \sigma_0, \text{Auth}_0, \sigma_1, \text{Auth}_1, \ldots, \sigma_d, \text{Auth}_d)$, the message M and the public key PK. To verify the signature, σ_0 and M are used to compute the corresponding W-OTS public key. The corresponding leaf $N_{0,j}$ of TREE$_0$ is constructed and used together with Auth$_0$ to compute the path $(P_0, \ldots, P_{h_0})$ to the root of TREE$_0$, where $P_0 = N_{0,j}$, $j = i \mod 2^{h_0}$ and

$$
P_c = \begin{cases} \text{H}((P_{c-1}||\text{Auth}_{c-1,0}) \oplus B_c), & \text{if } \lfloor j/2^c \rfloor \equiv 0 \mod 2 \\ \text{H}((\text{Auth}_{c-1,0}||P_{c-1}) \oplus B_c), & \text{if } \lfloor j/2^c \rfloor \equiv 1 \mod 2 \end{cases}
$$

for $0 \le c \le h_0$. This process is then iterated for $1 \le a \le d-1$, using the output of the last iteration $P_{h_{a-1}} = \text{ROOT}_{a-1}$ as message, σ_a, Auth$_a$ and $j = \left\lfloor i/2^{\sum_{b=0}^{a-1} h_b} \right\rfloor \mod 2^{h_a}$. If the output of the last iteration $P_{h_{d-1}}$ equals the root value contained in PK, ROOT$_{d-1}$, the signature is assumed to be valid and the algorithm returns 1. In any other case it returns 0.

2.1 Analysis

In the following we provide an analysis of XMSS^{MT}. We discuss correctness, provable security and bit security. But first, we discuss key and signature sizes as well as the theoretical runtimes of the algorithms.

Theoretical Performance. First we look at the sizes. A signature contains the index and d pairs of W-OTS signature and authentication path. Hence a signature takes $24 + n \cdot \sum_{i=0}^{d-1} (\ell_i + h_i)$ bits, assuming we reserve three bytes for the index. The public key contains the bitmasks, X and ROOT$_{d-1}$. Thus, the public key size is $n \cdot (\max_{0 \le i \le d-1}\{h_i + \lceil \log \ell_i \rceil\} + 2)$ bits. The secret key contains one FSGEN state and one BDS state consisting of $2(h_i - k_i)$ FSGEN states and no more than $(3h_i + \lfloor \frac{h_i}{2} \rfloor - 3k_i - 2 + 2^{k_i})$ tree nodes [BDS08] per currently used tree TREE$_i$. In addition it contains the $d-1$ W-OTS signatures $\sigma_1, \ldots, \sigma_{d-1}$ which have a total size of $n \cdot \sum_{i=1}^{d-1} \ell_i$ bits and the structures for upcoming trees. As observed by the authors of [HBB13], these structures do not require a full BDS state, as some of the structures are not filled during initialization and the space to store the k top levels of nodes can be shared with the current tree. Thus, these structures require only $(h_i - k_i + 1)$ FSGEN states (one for building the tree and the remaining as storage for the BDS state) and no more than $3h_i - k_i + 1$ tree nodes per TREE$_i$, $0 \le i < d-1$. The secret key size is

$$
n \cdot \left(\sum_{i=0}^{d-1} \left[\left(5h_i + \left\lfloor \frac{h_i}{2} \right\rfloor - 5k_i - 2 + 2^{k_i} \right) + 1 \right] + \sum_{i=0}^{d-2} (\ell_{i+1} + 4h_i - 2k_i + 2) \right)
$$

bits. For the runtimes we only look at the worst case times and get the following. During key generation, the first tree on each layer has to be computed. This means, that each of the 2^h W-OTS key pairs has to be generated, including the execution of the PRGs. Further, the leaves of the trees have to be computed and all internal nodes of the tree, to obtain the root. If key

generation generates the trees in order, starting from the first one, the W-OTS signatures on the roots of lower trees need no additional computation, as the signature can be extracted while the corresponding W-OTS key pair is generated. The key generation time is

$$
\mathsf{t_{Kg}} \le t_\mathrm{H} \left(\sum_{i=0}^{d-1} \left(2^{h_i}(\ell_i + 1) \right) \right) + t_\mathrm{F} \left(\sum_{i=1}^{d-1} \left(2^{h_i}(2 + \ell_i(w_i + 1)) \right) \right),
$$

where t_H and t_F denote the runtimes of one evaluation of H and F, respectively. During one call to Sign, a W-OTS signature on the message must be generated, including generation of the key $(t_\mathrm{F}(2 + \ell_0(w_0 + 1)))$, the BDS algorithm receives $(h_0 - k_0)/2$ updates, one leaf of the next tree on layer 0 must be computed and the BDS algorithm updates $h_0 - k_0$ upcoming seeds $(t_\mathrm{F}(h_0 - k_0))$. The worst case signing time is bounded by

$$
\mathsf{t_{Sign}} \le \max_{i \in [0, d-1]} \left\{ \begin{array}{l} t_\mathrm{H} \left(\frac{h_0 - k_0 + 2}{2} \cdot (h_i - k_i + \ell_i) + h_0 \right) \\ + t_\mathrm{F} \left(\frac{h_0 - k_0 + 4}{2} \cdot (\ell_i(w_i + 1)) + h_0 - k_0 \right) \end{array} \right\}
$$

Signature verification consists of computing d W-OTS public keys and the corresponding leafs plus hashing to the root. Summing up verification takes $\mathsf{t_{Vf}} \le \sum_{i=0}^{d-1} (t_\mathrm{H}(\ell_i + h_i) + t_\mathrm{F}(\ell_i w_i))$ in the worst case.

Correctness. In the following we argue that the BDS state on every layer receives its $(h - k)/2$ updates between two signatures. The authors of [HBB13] showed that there is a gap between the updates a XMSS key pair receives over its lifetime and the updates it really uses. Namely, for $2 \le k \le h$ there are 2^{h-k+1} unused updates. For every tree on a layer other than 0, $h - k + 3$ updates are needed. The first $h - k$ updates are needed to update the TREEHASH instances in the BDS state. Further one update is needed to compute a node in the next tree, to update the FsGEN states in the BDS state and to sign the root of the next tree on the lower layer, respectively. Between two signatures of a tree on layer $i + 1$, a whole tree on layer i is used, so if we ensure that $(h_{i+1} - k_{i+1} + 6)/2 \le 2^{h_i - k_i + 1}$, as we do, the $(h_i - k_i)/2$ updates TREE$_i$ receives over its lifetime leave enough unused updates such that TREE$_{i+1}$ receives the required updates. There would still occur a problem, if the number of updates per signature would be smaller than $d - 1$. The reason is, that this would mean that the roots of the new trees on different layers could not always be signed using the updates of the last signature before the change. In this case the private storage would grow, as we would need some intermediate storage for the new signatures. This is the reason why the second condition $((h_0 - k_0)/2 \ge d - 1)$ is needed.

Security. For the security, the argument is similar to that of [HBB13]. In [BDH11], an exact proof is given which shows that XMSS is forward secure, if $\mathcal{F}$ is a pseudorandom function family and $\mathcal{H}$ a second preimage resistant hash function family. The tree chaining technique corresponds to the product composition from [MMM02]. In [MMM02] the authors give an exact proof for the forward security of the product composition if the underlying signature schemes are forward secure. It is straight forward to combine both security proofs to obtain an exact proof for the forward security of XMSS^{MT}. In contrast to the case of XMSS$^+$, for XMSS^{MT} the product composition is applied not only once, but $d - 1$ times. As forward security implies EU-CMA security, this also shows that XMSS^{MT} is EU-CMA secure.

Security Level. We compute the security level in the sense of [Len04]. This allows a comparison of the security of XMSS^{MT} with the security of a symmetric primitive like a block cipher for given security parameters. Following [Len04], we say that XMSS^{MT} has security level b if a successful attack on the scheme can be expected to require approximately 2^{b-1} evaluations of functions from $\mathcal{F}_n$ and $\mathcal{H}_n$. Following the reasoning in [Len04], we only take into account generic attacks on $\mathcal{H}_n$ and $\mathcal{F}_n$. A lower bound for the security level of XMSS was computed in [BDH11]. For XMSS^{MT}, we combined the exact security proofs from [BDH11] and [MMM02]. Following the computation in [BDH11], we can lower bound the security level b by

$$b \geq \min_{0 \leq i \leq d-1} \left\{ n - 4 - w_i - 2log(\ell_i w_i) - \sum_{j=i}^{d-1} h_j \right\}$$

for the used parameter sets.

3   Optimization

Given the theoretical formulas for runtimes and sizes from the last section, we now show how to use them to model the parameter selection problem as linear optimization problem. There are parameters which control different trade-offs. The BDS parameters $k_i \in \mathbb{N}$ control a trade-off between signature time and secret key size, the Winternitz parameters $w_i \in \mathbb{N}$ control a trade-off between runtimes and signature size and the number of layers d determines a trade-off between key generation and signature time on the one hand and signature size on the other hand. Moreover there are the different tree heights h_i that do not define any obvious trade-off, but influence the security as well as the performance of the scheme. The goal of the optimization is to choose these parameters. The function families $\mathcal{F}_n$ and H can be instantiated, either using a cryptographic hash function or a block cipher. Hence the security parameter n is restricted to the output size of such functions. We choose 128 and 256 bit corresponding to AES and SHA-2 for our optimization, respectively.

Optimization Model. To find good parameter choices, we use linear optimization. In the following we discuss how we model the problem of optimal parameter choices as a linear optimization problem. As objective function of our problem we chose a weighted sum of all runtimes and sizes that should be minimized. Using the weights it is possible to control the importance of minimizing a certain parameter and thereby using the model for different scenarios. We further allow to apply absolute bounds on the runtimes and sizes. The formulas for runtimes and sizes are modeled as constraints as well as the parameter restrictions and the formula for bit security. The input to the model are the runtimes of $\mathcal{F}$ and $\mathcal{H}$ for $n = 128$ and $n = 256$, the overall height h, the message length m and a value b as lower bound on the bit security.

As many of our initial constraints are not linear, we have to linearize all functions and restrictions. This is done using the generalized lambda method [Mor07]. In addition, we split the problem into sub problems each having some decision variables fixed. The optimization problem contains the parameters of the scheme (d, n, the h_i, k_i, w_i for all layer) as decisions variables which are determined by solving the optimization. Further the message length m_i on each layer has to be modelled as a decision variable. Since solving the optimization problem takes much time and memory, we split the problem into sub problems by fixing the decision variables n and d. Therefore, we receive one sub problem for each combination of possible values of n and d. The resulting sub problems

are solved independently and the best of their solutions is chosen as global solution of the original optimization problem.

The next step is to linearize the remaining sub problems by using the generalized lambda method. Therefore, we introduce a grid point for each possible combination of the remaining variables h, k, w and m on each layer i. For each grid point we have a binary variable $\lambda_{h,k,w,m,i}$ which takes value 1 if the combination of h, k, w, m is chosen on layer i. Otherwise, it takes value 0. Since we need one choice of h, k, w, m for each $i \in [0, d-1]$, d λ's must be chosen.

We use those lambdas to calculate the functions describing the problem. Thus, before optimizing we determine the values of the functions for each possible values of their variables. To make this feasible, we have to introduce bounds on the decision variables. We bound the tree height per layer layer by 24. As $k \leq h$ this bound also applies to k. For w we chose 255. These bounds are reasonable for the scenarios of the next section. For different scenarios they might have to be changed. Then, to model the needed space of signatures $\sum_{i=0}^{d-1} (\ell_i + h_i) \cdot n$, we formulate the constraint

$$SpaceSig == \sum_{i=0}^{d-1} \sum_{h=1}^{24} \sum_{k=1}^{24} \sum_{w=2}^{512} \sum_{m \in \{128, 256\}} \lambda_{h,k,w,m,i} \cdot \underbrace{f_{SpaceSig}(w, h, m)}_{pre-calculated}$$

in the optimizing model, where $f_{SpaceSig}(w, h, m) = (\ell + h)n$. Thus, $SpaceSig$ gives the exact value of the needed space of signatures for the choice of lambda's and can be used in constraints and objective function.

To linearize a condition containing the maximum of some terms, such as the public key size $n \cdot (\max_{0 \leq i \leq d-1}\{h_i + \lceil \log \ell_i \rceil\} + 2)$, we write the following constraint:

$$\forall i \in \{1, ..., T\}$$
$$SpacePK \geq \sum_{h=1}^{24} \sum_{k=1}^{24} \sum_{w=2}^{512} \sum_{m \in \{128, 256\}} \lambda_{h,k,w,m,i} \cdot (\underbrace{f_{\lceil \log \ell \rceil}(w, m)}_{pre-calculated} + h + 2)n$$

Hence, $SpacePK$ gives the public key size for the choice of lambda's. This constraint pushes the value of $SpacePK$ up high and due to the objective function the value will be pushed down as low as possible, so that in the end it takes the exact value.

4 Optimization Results

In this section we present optimal parameters for two exemplary use cases. More results will be included in the full version of this paper. To solve the optimization problem, we used the IBM Cplex solver [IBM], that implements the Simplex algorithm [Dan63] with some improvements. The linearization described in the last section is exact. Thus, there is no loss of information or error. Therefore, it can be proven that the solution found by linear optimization based on the Simplex algorithm is the best possible solution. In the following we present the results and compare them with the results for parameter sets proposed in [BDH11] and [HBB13]. We choose a message length of 256 bits for all use cases assuming that the message is the output of a collision resistant hash function. Moreover we use 80 bits as lower bound for the provable bit security. This seems reasonable, as the used bit security represents a provable lower bound on the security of the scheme and is not related to any known attacks. We used the instantiations for $\mathcal{F}$ and $\mathcal{H}$ proposed in

[BDH11] with AES and SHA2 for $n = 128$ and $n = 256$, respectively and measured the resulting runtimes on a Laptop with Intel(R) Core(TM) i5-2520M CPU @2.5 GHz and 8 GB RAM. We got $t_F = 0.000225ms$ and $t_H = 0.00045ms$ for $n = 128$ as well as $t_F = 0.00169ms$ and $t_H = 0.000845ms$ for $n = 256$.

The first use case we look at meets the requirements of a document or code signature. We assume that the most important parameters are signature size and verification time and try to minimize them, while keeping reasonable bounds on the remaining parameters. We used the bounds $t_{Sign} < 1000ms, t_{Vf} < 1000ms, t_{Kg} < 60s, sk < 25kB, pk < 1.25kB, \sigma < 100kB$ and the weights $t_{Sign} = 0.00000001, t_{Vf} = 0.00090000, t_{Kg} = 0.00000001, sk = 0.00000001, \sigma = 0.99909996, pk = 0.00000001$ for $h = 20$. We chose different weights for t_{Vf} and σ, because the optimization internally counts in bits and milliseconds. We set the weights such that $1ms$ costs the same as 1000 bit. The remaining weights are not set to zero but to $1.0e - 8$, the smallest possible value that we allow. This is necessary to ensure that our implementation of inequalities in the model works. This also ensures that within the optimal solutions regarding t_{Vf} and σ, the best one regarding the remaining parameters is chosen. It turns out, that the optimization can be solved for $d \geq 2$. For $d = 1$ the bound on the key generation time cannot be achieved for the required height. If we relax this bound to be $t_{Kg} < 600s$, i.e. 10 minutes, the problem can be solved for $n = 128$ using AES. For $d \geq 2$ this relaxation does not change the results. The optimal parameters for this setting are $n = 128, d = 2, h_0 = 17, k_0 = 5, w_0 = 5$ and $h_1 = 3, k_1 = 3, w_1 = 22$. For comparison we use a parameter set from [HBB13] that matches the bound on the bit security ($n = 128, d = 2, h_0 = h_1 = 10, k_0 = k_1 = 4, w_0 = w_1 = 4$). The resulting runtimes and sizes are shown in Table 1. The results show that it is possible to reduce the signature size by almost one kilo byte, changing the other parameters within their bounds and increasing the second important parameter, the verification time, by 0.08 milliseconds.

As a second use case we take a total tree height of 80 and aim for a balanced performance over all parameters. This use cases corresponds to the use in a communication protocol. Again, we choose the weights such that $1ms$ costs the same as 1000 bit, but this time we use the same weights for all runtimes and for all sizes. For comparison we use parameters from [BDK$^+$07] ($d = 4$, $h_0 = h_1 = h_2 = h_3 = 20$, $w_0 = 5$, $w_1 = w_2 = w_3 = 8$, $k_0 = k_1 = k_2 = k_3 = 4$). As they do not use a BDS parameter, we choose $k = 4$ on all layers. To make a fair comparison, we limited our optimization also to four layers. The optimal parameters returned by the optimization are $h_0 = h_1 = h_2 = h_3 = 20$, $w_0 = 14$, $w_1 = w_2 = w_3 = 24$ and $k_0 = k_1 = k_2 = 4, k_3 = 2$. The results are shown in Table 1. It turns out, that again, by trading some runtime, the signature size can be significantly reduced.

Table 1. Runtimes and sizes for optimized parameters and parameters proposed in previous works.

	Runtimes (ms)			Sizes (bit)		
Use case	t_{Kg}	t_{Sign}	t_{Vf}	σ	PK	SK
$UC1$ optimal	27251	1.65	0.36	21376	6144	25472
$UC1$ from [HBB13]	326	1.00	0.28	28288	4608	25856
$UC2$ optimal	166min	25.55	9.13	83968	13824	209152
$UC2$ from [BDK$^+$07]	98min	14.53	5.01	119040	13824	233472

5 Conclusion

With XMSS^{MT} we presented a new hash-based signature scheme, that allows to use a key pair for 2^{80} signatures, which is a virtually unlimited number of signatures in practice. The new scheme is highly flexible and can be parameterized to meet the requirements of any use case for digital signatures. In order to get the maximum benefit from this parameterization, we showed how to use linear optimization to obtain provably optimal parameter sets. We have shown the strength and functionality of our approach by presenting the parameter sets for two different use cases. Comparing our results to parameter sets from other works, it turns out that there is a lot of space for optimization.

References

[BDE$^+$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 *Lecture Notes in Computer Science*, pages 363–378. Springer Berlin / Heidelberg, 2011.

[BDH11] 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*, volume 7071 of *Lecture Notes in Computer Science*, pages 117–129. Springer Berlin / Heidelberg, 2011.

[BDK$^+$07] Johannes Buchmann, Erik Dahmen, Elena Klintsevich, Katsuyuki Okeya, and Camille Vuillaume. Merkle signatures with virtually unlimited signature capacity. In Jonathan Katz and Moti Yung, editors, *Applied Cryptography and Network Security*, volume 4521 of *Lecture Notes in Computer Science*, pages 31–45. Springer Berlin / Heidelberg, 2007.

[BDS08] Johannes Buchmann, Erik Dahmen, and Michael Schneider. Merkle tree traversal revisited. In Johannes Buchmann and Jintai Ding, editors, *Post-Quantum Cryptography*, volume 5299 of *Lecture Notes in Computer Science*, pages 63–78. Springer Berlin / Heidelberg, 2008.

[BGD$^+$06] Johannes Buchmann, L. C. Coronado García, Erik Dahmen, Martin Döring, and Elena Klintsevich. CMSS - an improved Merkle signature scheme. In *INDOCRYPT*, volume 4329 of *Lecture Notes in Computer Science*, pages 349–363. Springer, 2006.

[BM99] Mihir Bellare and Sara Miner. A forward-secure digital signature scheme. In Michael Wiener, editor, *Advances in Cryptology — CRYPTO' 99*, volume 1666 of *Lecture Notes in Computer Science*, pages 786–786. Springer Berlin / Heidelberg, 1999.

[Dan63] George B. Dantzig. *Linear Programming And Extensions*. Princeton University Press, 1963.

[DOTV08] 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*, volume 5299 of *Lecture Notes in Computer Science*, pages 109–123. Springer Berlin / Heidelberg, 2008.

[HBB13] Andreas Hülsing, Christoph Busold, and Johannes Buchmann. Forward secure signatures on smart cards. In Lars R. Knudsen and Huapeng Wu, editors, *Selected Areas in Cryptography*, volume 7707 of *Lecture Notes in Computer Science*, pages 66–80. Springer Berlin Heidelberg, 2013.

[IBM] IBM. IBM ILOG CPLEX Optimizer. `http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/`. Accessed Januray 2013.

[Len04] Arjen K. Lenstra. Key lengths. Contribution to The Handbook of Information Security, 2004.

[MMM02] Tal Malkin, Daniele Micciancio, and Sara Miner. Efficient generic forward-secure signatures with an unbounded number of time periods. In Lars Knudsen, editor, *Advances in Cryptology — EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 400–417. Springer Berlin / Heidelberg, 2002.

[Mor07] Susanne Moritz. *A Mixed Integer Approach for the Transient Case of Gas Network Optimization*. PhD thesis, TU Darmstadt, Februar 2007.

A Winternitz OTS.

XMSS^{MT} uses the Winternitz-OTS (W-OTS) from [BDE$^+$11]. W-OTS uses the function family $\mathcal{F}_n$ and a value $X \in \{0,1\}^n$ that is chosen during XMSS key generation. For $K, X \in \{0,1\}^n$, $e \in \mathbb{N}$, and

$F_K \in \mathcal{F}_n$ we define $\mathrm{F}_K^e(X)$ as follows. We set $\mathrm{F}_K^0(X) = K$ and for $e > 0$ we define $K' = \mathrm{F}_K^{e-1}(X)$ and $\mathrm{F}_K^e(X) = \mathrm{F}_{K'}(X)$. Also, define

$$\ell_1 = \left\lceil \frac{m}{\log(w)} \right\rceil, \quad \ell_2 = \left\lfloor \frac{\log(\ell_1(w-1))}{\log(w)} \right\rfloor + 1, \quad \ell = \ell_1 + \ell_2.$$

The secret signature key of W-OTS consists of ℓ n-bit strings sk_i, $1 \le i \le \ell$. The generation of the sk_i will be explained later. The public verification key is computed as

$$\mathsf{pk} = (\mathsf{pk}_1, \ldots, \mathsf{pk}_\ell) = (\mathrm{F}_{\mathsf{sk}_1}^{w-1}(X), \ldots, \mathrm{F}_{\mathsf{sk}_\ell}^{w-1}(X)),$$

with F^{w-1} as defined above. W-OTS signs messages of binary length m. They are processed in base w representation. They are of the form $M = (M_1 \ldots M_{\ell_1})$, $M_i \in \{0, \ldots, w-1\}$. The checksum $C = \sum_{i=1}^{\ell_1}(w-1-M_i)$ in base w representation is appended to M. It is of length ℓ_2. The result is a sequence of ℓ base w numbers, denoted by $(T_1, \ldots, T_\ell)$. The signature of M is

$$\sigma = (\sigma_1, \ldots, \sigma_\ell) = (\mathrm{F}_{\mathsf{sk}_1}^{T_1}(X), \ldots, \mathrm{F}_{\mathsf{sk}_\ell}^{T_\ell}(X)).$$

It is verified by constructing $(T_1 \ldots, T_\ell)$ and checking

$$(\mathrm{F}_{\sigma_1}^{w-1-T_1}(X), \ldots, \mathrm{F}_{\sigma_\ell}^{w-1-T_\ell}(X)) \overset{?}{=} (\mathsf{pk}_1, \ldots, \mathsf{pk}_\ell).$$

The sizes of signature, public, and secret key are ℓn bits. The runtimes for key generation, signature generation and signature verification are all bounded by $\ell(w-1)$ evaluations of elements from $\mathcal{F}_n$. The Winternitz parameter w controls a time - space trade-off, as ℓ shrinks logarithmically in w. For more detailed information see [BDE$^+$11].

B   XMSS^{MT} in Graphics

Figure 3 shows a schematic representation of a XMSS^{MT} instance with four layers.

Fig. 3. A schematic representation of a XMSS^{MT} instance with four layers

13