Unified MEDS Accelerator

Sanjay Deshpande[1], Yongseok Lee[2], Mamuri Nawan[3], Kashif Nawaz[3],
Ruben Niederhagen[4,5], Yunheung Paek[2], and Jakub Szefer[6]

[1] CASLAB, DEE, Yale University, New Haven, US, sanjay.deshpande@yale.edu
[2] ECE and ISRC, Seoul National University, Seoul, South Korea,
yslee@sor.snu.ac.kr,ypaek@snu.ac.kr
[3] CRC, Technology Innovation Institute, Abu Dhabi, UAE,
mamuri@tii.ae, kashif.nawaz@tii.ae
[4] IIS, Academia Sinica, Taipei, Taiwan, ruben@polycephaly.org
[5] IMADA, University of Southern Denmark, Odense, Denmark
[6] Northwestern University, Evanston, US, jakub.szefer@northwestern.edu

Abstract. The Matrix Equivalence Digital Signature (MEDS) scheme, a code-based candidate in the first round of NIST's Post-Quantum Cryptography (PQC) standardization process, offers competitively small signature sizes but incurs high computational costs for signing and verification. This work explores how a high-performance FPGA-based hardware implementation can enhance MEDS performance by leveraging the inherent parallelism of its computations, while examining the trade-offs between performance gains and resource costs. This work in particular proposes a unified hardware architecture capable of efficiently performing both signing and verification operations within a single combined design. The architecture jointly supports all security parameters, including the dynamic, run-time handling of different prime fields without the need to re-configure the FPGA. This work also evaluates the resource overhead of supporting different prime fields in a single design, which is relevant not only for MEDS but also for other cryptographic schemes requiring similar flexibility. This work demonstrates that custom hardware for PQC signature schemes can flexibly support different prime fields with limited resource overhead. For example, for NIST security Level I, our implementation achieves signing times of 4.5 ms to 65.2 ms and verification times of 4.2 ms to 64.5 ms utilizing 22k to 72k LUTs and 66 to 273 DSPs depending on design variant and optimization goal.

Keywords: Post-Quantum Cryptography · Digital Signature Algorithm · Matrix Equivalence Digital Signature (MEDS)

1 Introduction

Matrix Equivalence Digital Signature (MEDS) is a code-based digital signature scheme based on the hardness assumption of the Matrix Code Equivalence (MCE) problem. It was submitted to the National Institute of Standards and Technology (NIST) Post-Quantum Cryptography (PQC) Signature standardization process launched in 2023, but it did not advance to the second round due to its comparatively low performance (e.g., according to data from the PQM4 project[7], when evaluated in software on an embedded evaluation board, MEDS is 25×, 244×, 856× slower in key generation, signing, and verification times respectively when compared to NIST standard ML-DSA Dilithium [KKPY24] on security level I) and the novelty of its security assumptions. Given the importance of ensuring the practicality of the cryptographic schemes, this paper focuses on developing a unified high-performance hardware implementation of the MEDS signing and verification operation. We present a joint design that supports all security parameter sets, selectable at runtime, and evaluate the overhead of accommodating arithmetic for different finite fields in a single design. Our results on the overhead generalize well and can be transferred to other schemes with similar parameter set specifications such as the NIST PQC Signature Round 2 submission CROSS [BBB+24a], which has parameter sets for $\mathbb{F}_{127}$ and $\mathbb{F}_{509}$. Other related schemes such as LESS [BBB+24b] and PERK [ABB+24] that currently are using a single small binary field might benefit from investigating trade-offs of offering parameter sets with different fields.

By addressing the computational challenges associated with MEDS, we aim to contribute to the broader goal of making post-quantum cryptography viable for widespread adoption in the face of advancing quantum computing technologies. To the best of our knowledge, this work presents the first hardware implementation of MEDS. Our implementation operates in constant time (i.e., there are no timing variations depending on secret data), though it does not include additional side-channel protections.

Related Work. Gaussian elimination is an important step in MEDS as well as many other PQC schemes. Early work [HQR89] used processor arrays as large as the matrix being processed. The work in [SWM+10] presented one of the first hardware accelerators for the code-based Niederreiter cryptosystem. The work used a systolic processor array for performing Gaussian systemization. Later work on Classic McEliece also used a systemizer similar to [SWM+10] with smaller processor array. For example, in [WSN16,WSN17] the authors present a hardware design of a key-generation module for the Niederreiter cryptosystem. In [CCD+22], the authors implemented a complete design of Classic McEliece compliant with the specification submitted to NIST standardization process including binary field systemizers. We are using a similar systemizer in our work.

Although there are similar computations in regard to matrix systemization between MEDS and Classic McEliece, compared to Classic McEliece, instead of working with binary fields, MEDS requires $\mathbb{F}_q$ with prime q. As a result, the latencies (in overall wall-clock time) of finite field operations are larger compared to when using $\mathbb{F}_2$ or a small binary extension filed $\mathbb{F}_q^m$. This difference necessitates significant design effort for the systemizer design, e.g., a different control logic and a more complex pipelining process.

Within current NIST standardization process for digital signatures schemes, in [SMA+24,HSK+23] the authors develop hardware implementations of the

[7] https://github.com/mupq/pqm4

MAYO digital signature scheme. In [DHSY24], authors present the first hardware implementation of the SDitH digital signature scheme. [dPRS23] describes hardware FPGA implementations of the Raccoon digital signature scheme.

In [WJW$^+$19], the authors present several hardware accelerators that work with a RISC-V core to accelerate the XMSS signature scheme. In [BCH$^+$23], the authors present hardware implementation of the Oil and Vinegar (OV) signature scheme. In [TYD$^+$11] authors presented an FPGA-based implementation of the multivariate-based signature scheme Rainbow. Later, in [FG18] authors presented high-speed, FPGA-based implementation of Rainbow with updated parameters for NIST's first round of PQC standardization process. Both designs are operating on small binary fields using a processor array similar to the code-based designs discussed above. Among ASIC designs, in [ZZL$^+$23] authors present a processor designed specifically for PQC algorithms, which can support schemes such as Saber, Kyber, Dilithium, NTRU, McEliece, and Rainbow. This design also includes a module for the systemization of matrices over small binary fields, but instead of a processor array composed of several vector units, the authors here are using only a single a vector unit.

Hardware architectures for algorithms already being standardized by NIST include [SAW$^+$23], where the authors implement a hardware design of the FALCON scheme and [LSG21,BNG21] where the authors explore hardware implementations of CRYSTALS-Dilithium (being standardized under the name "ML-DSA"). [ALCZ20,Saa24,DLK$^+$25] explores implementations of SPHINCS+ (being standardized under the name "SLH-DSA").

We are not aware of any other hardware implementation of MEDS and our work presents a hardware design of MEDS that provides a competitive performance and resource utilization compared hardware implementations of other signature schemes.

Contributions. This paper introduces the first hardware implementation of the sign and verify operations for the PQC signature scheme MEDS. Our implementation combines both operations into a single design sharing most resources between sign and verify. We further provide joint support of all security parameter sets in one single design selectable at runtime.

We aim to address the following research questions:

Q1: Does MEDS provide sufficient inherent parallelism to accelerate the sign and verify operations to a competitive level? What are the associated resource costs for achieving such a speed-up?

Q2: To what extent can resources be shared in a combined design for the sign and verify operations?

Q3: What is the overhead of supporting all security parameters at runtime, particularly concerning arithmetic in two distinct finite fields?

The source code of our hardware designs is available under an open source license at `https://github.com/caslab-code/pqc-hw-meds`.

2 Preliminaries

We first introduce our notation in Section 2.1 and then explain the relevant details of the MEDS specification in Section 2.2.

2.1 Notation

We are following the notation of the MEDS specification document [CNP$^+$23a] and denote matrices with bold capital letters, e.g., $\mathbf{M}$, $\mathbf{A}$, and $\mathbf{B}$. We follow the MEDS specification in denoting submatrices with square brackets, e.g., $\mathbf{M}[a, b; c, d]$ denotes the submatrix of the intersection of rows a to b and columns c to d of matrix $\mathbf{M}$ [CNP$^+$23a, Section 2.1]. If no row or column range is provided, all rows or columns are included. $\mathbb{F}_q$ is a finite field with q elements. We define an $[m \times n, k]$ matrix rank metric code over $\mathbb{F}_q$ as k-dimensional subspace of $\mathbb{F}_q^{m \times n}$. Here, m and n are the codeword sizes and k the code dimension.

The operation $\mathsf{SF}(\mathbf{M})$ returns the systematic form of a matrix $\mathbf{M}$ if it exists or $\perp$ if not. The operation $\pi_{\mathbf{A},\mathbf{B}}(\mathbf{M})$ with $\mathbf{A} \in \mathbb{F}_q^{m \times m}$, $\mathbf{B} \in \mathbb{F}_q^{n \times n}$, and $\mathbf{M} \in \mathbb{F}_q^{k \times mn}$ first maps each row $i \in \{0, \ldots, k-1\}$ of $\mathbf{M}$ to a matrix $\mathbf{P}_i \in \mathbb{F}_q^{m \times n}$ such that $\mathbf{P}_i[\lfloor j/n \rfloor; j \bmod n] = \mathbf{M}[i, j]$ for $j \in \{0, \ldots, mn-1\}$. It then computes $\mathbf{P}'_i \in \mathbb{F}_q^{m \times n}$ as $\mathbf{P}'_i = \mathbf{A}\mathbf{P}_i\mathbf{B}$, maps the $\mathbf{P}'_i$ back to the rows i of a matrix $\mathbf{M}' \in \mathbb{F}_q^{k \times mn}$, and finally returns $\mathbf{M}'$ as result.

2.2 MEDS

The MEDS scheme [CNP$^+$23a,CNP$^+$23b] was a submissions to the on-ramp to the NIST PQC signature standardization process[8] but it did not advance to the second round. It is based on the notion of Matrix Code Equivalence (MCE):

Definition 1 (Matrix Code Equivalence). *Let $\mathcal{C}$ and D be two $[m \times n, k]$ matrix codes over $\mathbb{F}_q$. We say that C and $\mathcal{D}$ are* equivalent *if there exist two invertible matrices $\mathbf{A} \in \mathbb{F}_q^{m \times m}$ and $\mathbf{B} \in \mathbb{F}_q^{n \times n}$ such that $\mathcal{D} = \mathbf{A} \cdot \mathcal{C} \cdot \mathbf{B}$, i.e., for all $\mathbf{C} \in \mathcal{C}$, $\mathbf{A} \cdot \mathbf{C} \cdot \mathbf{B} \in \mathcal{D}$.*

This gives rise to the MCE problem:

Problem 1 (Matrix Code Equivalence Problem). Given two k-dimensional matrix codes $\mathcal{C}, \mathcal{D} \subset \mathbb{F}_q^{m \times n}$, find two invertible matrices $\mathbf{A} \in \mathbb{F}_q^{m \times m}$, $\mathbf{B} \in \mathbb{F}_q^{n \times n}$ such that $\mathcal{D} = \mathbf{A} \cdot \mathcal{C} \cdot \mathbf{B}$.

The MCE problem is at least as hard as the linear code equivalence problem and as hard as the isomorphism of polynomials problem [BFV13], and as hard as the alternating trilinear form equivalence problem [GQT21,TDJ$^+$22]. Please refer to [CNP$^+$23a,CNP$^+$23b] for a concrete security analysis.

MEDS uses the MCE problem to construct a signature scheme from an interactive Σ-protocol using the Fiat-Shamir (FS) transform [FS87] with t rounds and by applying some tricks to improve signature size and performance, i.e., by using s matrix codes in the public key, by using challenges with a fixed weight w, by seeding 0-responses, and by generating these seeds from a seed tree.

[8] `https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures`

Table 1: MEDS parameter sets.

Level	Parameter Set	q	n	m	k	s	t	w	pk (byte)	sig (byte)
MEDS Parameter Sets [CNP+23a, Table 2]:										
Level I	MEDS-9923	4093	14	14	14	4	1152	14	9923	9896
Level I	MEDS-13220	4093	14	14	14	5	192	20	13 220	12 976
Level III	MEDS-41711	4093	22	22	22	4	608	26	41 711	41 080
Level III	MEDS-69497	4093	22	22	22	5	160	36	55 604	54 736
Level V	MEDS-134180	2039	30	30	30	5	192	52	134 180	132 528
Level V	MEDS-167717	2039	30	30	30	6	112	66	167 717	165 464
New Parameter Sets May 2024 (signature bytes with seed tree)[9]:										
Level I		4093	26	25	25	2	144	48	21 595	5456
Level III		4093	35	34	34	2	208	75	55 520	10 786
Level V		4093	45	44	44	2	272	103	122 000	21 052

MEDS parameters. Table 1 shows the parameter sets of MEDS from the first round of the NIST PQC signature scheme standardization. In May 2024, the MEDS submission team announced new parameter sets in the NIST PQC mailing list[9] as reaction to refined attacks [NQT24]. However, they suggest to combine these new parameter sets with signature-size optimization techniques introduced in [CNRS24]. These optimizations require algorithmic changes compared to the MEDS Round 1 specification document [CNP+23a]. We decided to provide a hardware implementation that follows the MEDS Round 1 specification and therefore do not implement the optimization from [CNRS24] since there is no concrete specification for that MEDS variant and since we are interested in analyzing the effect of supporting different prime fields in a joint design.

MEDS key generation. To generate a key pair, first chose a random $\mathbf{G}_0 \in \mathbb{F}_q^{k\times mn}$. Then choose random invertible matrices $\mathbf{A}_1, \ldots, \mathbf{A}_{s-1} \in \mathbb{F}_q^{m\times m}$ and $\mathbf{B}_1, \ldots, \mathbf{B}_{s-1} \in \mathbb{F}_q^{n\times n}$. Finally, compute the matrices $\mathbf{G}_1, \ldots, \mathbf{G}_{s-1} \in \mathbb{F}_q^{k\times mn}$ as $\mathbf{G}_i = \mathsf{SF}(\pi_{\mathbf{A}_i,\mathbf{B}_i}(\mathbf{G}_0))$, $i \in \{1, \ldots, s-1\}$. The public key of MEDS consists of the seed for randomly generating $\mathbf{G}_0 \in \mathbb{F}_q^{k\times mn}$ and $s-1$ matrices $\mathbf{G}_1, \ldots, \mathbf{G}_{s-1} \in \mathbb{F}_q^{k\times mn}$. The secret key of $s-1$ pairs of invertible matrices $(\mathbf{A}_1^{-1}, \mathbf{B}_1^{-1}), \ldots, (\mathbf{A}_{s-1}^{-1}, \mathbf{B}_{s-1}^{-1}) \in \mathbb{F}_q^{m\times m} \times \mathbb{F}_q^{n\times n}$.

MEDS signing. During signing, the signer commits to t matrices $\tilde{\mathbf{G}}_i \in \mathbb{F}_q^{k\times mn}$, $i \in \{0, \ldots, t-1\}$ using t maps of pairs of random invertible matrices $(\tilde{\mathbf{A}}_0, \tilde{\mathbf{B}}_0), \ldots, (\tilde{\mathbf{A}}_{t-1}, \tilde{\mathbf{B}}_{t-1}) \in \mathbb{F}_q^{m\times m} \times \mathbb{F}_q^{n\times n}$ such that $\tilde{\mathbf{G}}_i = \mathsf{SF}(\pi_{\tilde{\mathbf{A}}_i,\tilde{\mathbf{B}}_i}(\tilde{\mathbf{G}}_0))$, $i \in \{0, \ldots, t-1\}$. The signer then hashes the $\tilde{\mathbf{G}}_i$ to the commitment hash d and parses d to the

[9] `https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/pbT_DnPrc2A/m/ZPrIVSmFCQAJ`

challenge vector $h_0, \ldots, h_{t-1}$ of weight w with $h_i \in \{0, s-1\}$. The signature finally is composed by providing d as well as the seeds $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ for all i where $h_i = 0$ and $\mu_i = \tilde{\mathbf{A}}_i \cdot \mathbf{A}_{h_i}^{-1}$ and $\nu_i = \mathbf{A}_{h_i}^{-1} \cdot \tilde{\mathbf{B}}_i$ for all i where $h_i > 0$.

MEDS verification. For verification, the verifier parses the challenge vector $h_0, \ldots, h_{t-1}$ from d. The verifier then computes $\hat{\mathbf{G}}_i \in \mathbb{F}_q^{k \times mn},\ i \in \{0, \ldots, t-1\}$ as $\hat{\mathbf{G}}_i = \mathsf{SF}(\pi_{\mu_i,\nu_i}(\mathbf{G}_{h_i}))$, where the matrices $\mu_i = \tilde{\mathbf{A}}_i$ and $\nu_i = \tilde{\mathbf{B}}_i$ are regenerated from the seed-tree seeds for all i with $h_i = 0$, and $\mu_i = \tilde{\mathbf{A}}_i \cdot \mathbf{A}_{h_i}^{-1}$ and $\nu_i = \mathbf{A}_{h_i}^{-1} \cdot \tilde{\mathbf{B}}_i$ are parsed from the signature for all i with $h_i > 0$. If the hash of the $\hat{\mathbf{G}}_i$ equals the commitment hash d, the verification is successful.

Cost. The most expensive sub-operations in MEDS are SF with a cost of $\mathrm{O}(mnk^2)$ finite field operations for performing Gaussian elimination on a matrix of size $k \times mn$ and π with a cost of $\mathrm{O}(kmn^2 + km^2n)$ for performing k matrix products of sizes $\mathbb{F}_q^{m \times m}$ times $\mathbb{F}_q^{m \times n}$ times $\mathbb{F}_q^{n \times n}$. Hence, the computational cost of SF and π is quartic in the security parameters with $k \approx m \approx n$.

For key generation, we need to compute SF and π each only s times, where $s \in \{4, 5, 6\}$ — but for signing and verification, we need to perform these costly operations t times with $t \in \{112, 160, 192, 608, 1152\}$. Therefore, signing and verification are 18.6 to 288 times more expensive than key generation and hence benefit the most from hardware acceleration. Since verification performs almost the same computations as signing with only little difference in the control flow, we combine both operations in a single design, sharing as many resources between these operations as possible.

3 Unified MEDS Design

We call our hardware design of MEDS a *unified* design, since it *combines* both the MEDS signing and verification operations into a single design, and it provides *joint* support for all parameter sets selectable at runtime, while sharing hardware resources between both operations and all parameter sets as much as possible.

Algorithm 1 and Algorithm 2 in the Appendix show the signing and verification operations in detail [CNP+23a, Algorithms 11 and 12]. The inputs to the signing operation are the secret key sk and the message msg that needs to be signed. The inputs to the verify operation are the public key pk and the signed message to be verified. Comparing the main loop in line 14 to line 22 in Algorithm 1 with the main loop in line 12 to line 27 in Algorithm 2 shows that the main operations of obtaining invertible matrices, applying the π and SF operations, and computing a commitment hash are structurally identical. Hence, in our combined design, resources for performing these computations can be shared efficiently between the sign and the verify operation with little control overhead.

3.1 Top-down Overview of Our MEDS Hardware Design

Figure 1 shows the operation flow of our combined hardware implementation. Data flow unique to the sign operation is shown in red and unique to the verify

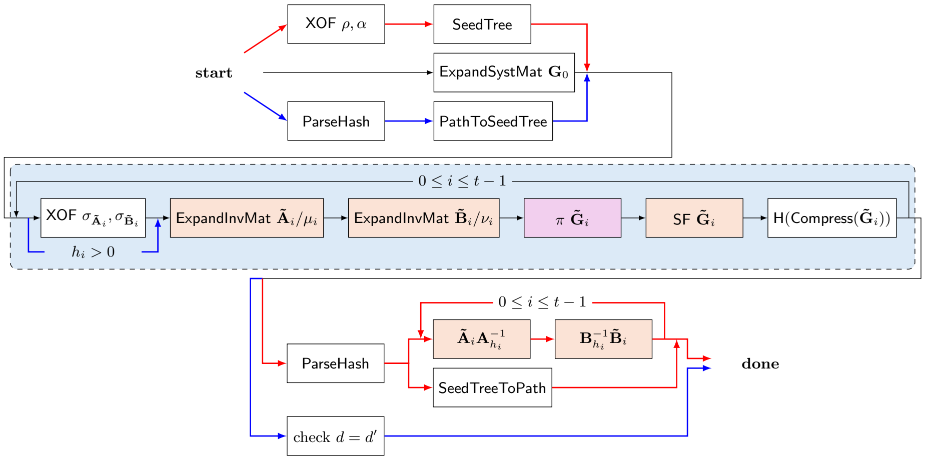
XOF ρ, α
SeedTree
start
ExpandSystMat $\mathbf{G}_0$
ParseHash
PathToSeedTree
$0 \leq i \leq t-1$
XOF $\sigma_{\tilde{\mathbf{A}}_i}, \sigma_{\tilde{\mathbf{B}}_i}$
ExpandInvMat $\tilde{\mathbf{A}}_i/\mu_i$
ExpandInvMat $\tilde{\mathbf{B}}_i/\nu_i$
π $\tilde{\mathbf{G}}_i$
SF $\tilde{\mathbf{G}}_i$
H(Compress($\tilde{\mathbf{G}}_i$))
$h_i > 0$
$0 \leq i \leq t-1$
ParseHash
$\tilde{\mathbf{A}}_i \mathbf{A}_{h_i}^{-1}$
$\mathbf{B}_{h_i}^{-1} \tilde{\mathbf{B}}_i$
SeedTreeToPath
done
check $d = d'$

Fig. 1: Overview of the operational flow of our combined sign and verify module, red lines for sign only, blue lines for verify only.

operation in blue. First, both operations expand $\mathbf{G}_0$ from seed $\sigma_{\mathbf{G}_0}$, which is parsed from `sk` or `pk` respectively, using the module `ExpandSystMat`.

Signature generation starts with the generation of the seed δ of length $\ell_{\mathsf{sec\_seed}}$ from a randomness source. We assume that the seed δ will be initialized by a top-level hardware module, e.g., by interfacing to a True Random Number Generator (TRNG) or by implementing the NIST Known Answer Test (KAT) Pseudo Random Number Generator (PRNG). The seed δ is then expanded using the module `XOF` to the root seed ρ of length $\ell_{\mathsf{tree\_seed}}$ of the seed tree and a salt α of length ℓ_{salt}. In our hardware design, we accomplish this again using a `SHAKE256` module via a `XOF` wrapper. Then the next step is the seed tree generation using the root seed ρ and the salt α with the `SeedTree` module. Verification instead parses the hash and then generates the partial seed tree from the signature. We implement this in the modules `ParseHash` and `PathToSeedTree`.

For the main loop in sign and verify we need to expand the seeds σ_i from the seed tree leaves into seeds $\sigma_{\tilde{\mathbf{A}}_i}$ and $\sigma_{\tilde{\mathbf{B}}_i}$ and expand these seeds to $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ (i.e., μ_i and ν_i respectively for verify) using a `ExpandInvMat` module. For signing, we store each σ_i in a memory which is later used to recompute $\sigma_{\tilde{\mathbf{A}}_i}$ and $\sigma_{\tilde{\mathbf{B}}_i}$ and then expand these seeds to $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ for computation of μ_i and ν_i. We then feed the matrices into a `Pi` module together with the matrix $\mathbf{G}_0$, compute the systematic form $\tilde{\mathbf{G}}_i$ of the result using the module `SF`. We pull the absorb-function of the hash function H into the loop to avoid the need to store all $\tilde{\mathbf{G}}_i$. We will describe the implementation of this loop in detail in Section 4.5.

The steps after the main loop are different for sign and verify. For sign, we also absorb the message `msg` into the hash state and finalize H to obtain the challenge vector. Then, to compute the final loop, we recompute $\sigma_{\tilde{\mathbf{A}}_i}$ and $\sigma_{\tilde{\mathbf{B}}_i}$ using σ_i, and then expand them to generate $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ for the computation of μ_i

and ν_i using a matrix-matrix multiplication module. Eventually, we compute the seed tree path for the signature using `SeedTreeToPath`. For verify, we only need to check that $d = d'$ and set the return value of the module correspondingly.

3.2 Joint Security Parameter Design

One of the main questions we want to investigate in this work is the hardware overhead of supporting several prime fields for different security parameter sets selectable at runtime in one joint hardware design. For modules that require both 11-bit arithmetic in $\mathbb{F}_{2039}$ and 12-bit arithmetic in $\mathbb{F}_{4093}$, we are using an input signal `i_sel_4093` that is pulled high for $\mathbb{F}_{4093}$ and low for $\mathbb{F}_{2039}$. The bus width for field elements is set to 12 bit for the joint design, with the most significant bit set to zero when $\mathbb{F}_{2039}$ is being selected. Therefore, the cost for memory storing field elements is defined by the larger field in the joint design. A similar design approach of constructing field arithmetic modules capable of operating over two distinct prime fields was adopted in [AMI$^+$23,MR24], where the authors integrated CRYSTALS-Kyber (which operates over a 12-bit prime field polynomial coefficients) and CRYSTALS-Dilithium (which utilizes a 23-bit prime field polynomial coefficients). In contrast, MEDS features prime fields that differ by only a single bit, enabling support for multiple fields with minimal overhead.

If the Verilog code is synthesised for one specific field only, we are fixing `i_sel_4093` to constant 0 or 1 respectively so that the optimization step during synthesis can remove unnecessary logic automatically. This enables us to get precised resource estimates for hardware designs supporting only one respective field and for the joint design with otherwise identical logic.

The remaining security parameters besides the prime field are matrix dimensions n, m, and k, the number s of `pk` matrices, as well as t and s for the challenge vector. All these other parameters mostly affect the upper limit of loops and can easily be selected using multiplexers at runtime at little resource overhead. We pre-compute related pre-defined constant values (e.g., counter widths) and provide them as macros and module parameters during synthesis (and simulation).

4 Implementation

For the description of our implementation of the MEDS signature scheme, we take a bottom-up approach. The different building blocks and sub-modules are described in the following including finite field arithmetic, matrix operations, and the complete sign and verify module.

4.1 Finite Field Arithmetic

As described in specified in Section 2.2, all underlying arithmetic operations multiplication, addition, and inversion in the MEDS signature scheme are performed in the prime field $\mathbb{F}_q$ where q is either the 11-bit prime $2039 = 2^{11} - 9$

Table 2: Resource consumption and performance of the $\mathbb{F}_q$ arithmetic functions targeting an Xilinx Artix 7 (`xc7a200t`) FPGA.

		Resources							
Operation	q	**Area**		**Memory**		**Cycles**	**Freq.**	**Time**	**Time×Area**
		(LUT)	(DSP)	(FF)	(BRAM)	(cyc.)	(MHz)	(us)	
Addition	2 039	27	0	23	—	1	244	0.004	0.110×10^3
	4 093	28	0	25	—	1	221	0.005	0.127×10^3
	both	33	0	25	—	1	201	0.005	0.164×10^3
Multiplication	2 039	30	1	40	—	4	361	0.011	0.333×10^3
	4 093	44	1	41	—	4	357	0.011	0.494×10^3
	both	46	1	45	—	4	343	0.012	0.536×10^3
Inverse 2× parallel	2 039	0	0	1	1.0	1	392	0.003	—
	4 093	0	0	1	1.5	1	392	0.003	—
	both	12	0	1	2.5	1	269	0.004	0.045×10^3

or the 12-bit prime $4093 = 2^{12} - 3$. We implement these operations as integer arithmetic modulo the prime q. An overview of the resource consumption and performance for each module is provided in Table 2.

Modular Multiplication. Since the size of each element in $\mathbb{F}_q$ is 11 or 12 bits, we can use the DSPs available on the target FPGA (AMD Artix 7 `xc7a200t-3`) to perform integer multiplication followed by full reduction. Using the DSPs helps to avoid long critical paths in the FPGA design. For modular reduction after integer multiplication, we design a specific modular reduction unit targeting the q value. As q is in a pseudo-Mersenne form ($q = 2^{11} - 9 = 2039$ or $q = 2^{12} - 3 = 4093$), we use the folding technique to perform the modular reduction [HGG07]. As the size of the multiplied output is up to $2\lceil \log_2(q) \rceil$ bits (i.e., 22 resp. 24 bits), we take the most-significant $\lceil \log_2(q) \rceil$ bits (i.e., 11 resp. 12 bits), of the output, multiply them by $2^{\lceil \log_2(q) \rceil} - q$ (i.e., $9 = 1001_b$ or $3 = 11_b$ — since the binary representation of the factor is sparse, we simple shift-and-add instead of multiplying), and add this to the least significant $\lceil \log_2(q) \rceil$ bits (i.e., 11 resp. 12 bits). As the result of the addition may have more than $\lceil \log_2(q) \rceil$ bits, we repeat this folding process until all most-significant bits are folded and until there is no more carry generated from addition after folding. Eventually, we use a multiplexer at the intermediate folded result r to identify corner cases $r \geq q$ and in case we detect one of these cases, we subtract the modulus q from r.

Many computations in the joint design are the same for both fields. The only difference in the folding reduction step is that for $q = 2039$ we need to multiply the top bits by 9 and for $q = 4093$ by 3, which only requires an additional multiplexer. Also the final conditional subtraction just requires a multiplexer to select the prime.

The input-independent latency for our modular multiplication unit for all variants is four clock cycles. One register is placed after the multiplication unit followed by two registers to pipeline the reduction unit efficiently to reduce the overall critical path. The output is not registered by default.

Modular Inversion. This operation is mostly required for matrix systemization (described in detail in Section 4.3). To perform that computation efficiently, we require low-latency field inversion. To achieve this and since the size of the finite field $\mathbb{F}_q$ with 11 and 12 bits is relatively small, we decided not to use expensive, high-latency inversion algorithms like the extended Euclidean algorithm, Fermat's little theorem or Montgomery inversion for our modular inversion module. Instead, we trade computation for memory and precompute the inverse of all elements in $\mathbb{F}_q$ at synthesis time and store them as look-up table in Block RAM (BRAM). Hence, we have a input-independent latency of only one clock-cycle for performing modular inversion. To use the memory resources as efficiently as possible, we are exploiting the dual-ported interface of the BRAM to perform two independent inversions in parallel, hence halving the effective resource cost.

The joint design for this operation provides look-up tables for both primes and some additional logic for selecting the output of the requested field. The resource requirements are shown on Table 2.

Summary. The resource evaluation in Table 2 demonstrates that hardware support for multiple fields imposes only a small additional cost, provided the fields are "related," meaning they share the same reduction algorithm. Similar works that support two distinct prime fields [AMI$^+$23,MR24] do not present a fine-grained analysis of the overhead introduced at the modular arithmetic level due to dual-field support; instead, they primarily emphasize the associated overhead at the level of polynomial multiplication units. In our reduction module, when employing the folding technique, efficient resource sharing can be achieved under the condition that the primes have a low Hamming distance. Consequently, supporting different prime fields in hardware is not prohibitively resource-intensive if the fields are selected carefully. This can make using different prime fields a viable design choice for cryptographic primitives when using different fields offers meaningful benefits, such as reduced signature or ciphertext sizes.

4.2 Matrix Multiplication

Matrix multiplication is one of the most used operations in the MEDS It is used in the π operation and for computing the μ_i and ν_i during signing. Consequently, we take care to design an efficient matrix multiplication unit that is parameterizable at synthesis time to enable performance trade-offs. As mentioned in Section 3.1, matrices in MEDS are specified to be sampled in row-wise order.

The general task of this module is to compute $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$. Instead of traditional schoolbook matrix multiplication with limited parallelism, we are using a vector-based approach similar to standard SIMD techniques in our matrix multiplication design.

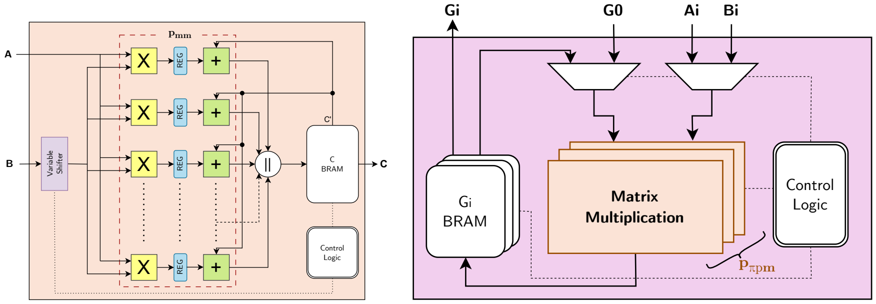

(a) matrix multiplication module.

(b) π module.

Fig. 2: Hardware design of the matrix multiplication and π modules.

Asymptotically more efficient matrix multiplication algorithms, e.g., Strassen or Coppersmith–Winograd, do not apply to MEDS due to the relatively small matrix dimensions. Additionally, techniques such as the trick by Arlazarov, Dinic, Kronrod, and Faradžev [ADKF70], which are not asymptotically faster and are most effective for small fields when the matrix dimension exceeds the field size (e.g., GF(16) with matrix dimension up to 133 in [BCC$^+$24]), are also not applicable in the context of MEDS due to the relatively large fields $\mathbb{F}_{2039}$ and $\mathbb{F}_{4093}$ and small matrix dimension of up to 30.

This approach allows us to take advantage of the pipelining of the $\mathbb{F}_q$ multiplier and also allows us to perform single-element additions in each iteration. This design also allows us to stream the inputs to the multiplication unit and to make best use of all pipelining registers. The latency l_{mm} of our matrix multiplication unit can be computed from the matrix dimensions d_1, d_2, and d_3, the number of pipeline stages l_{mmpipe}, and the performance parameter p_{mm} (i.e., the column-block width) as $l_{\text{mm}} = l_{\text{mmpipe}} + \frac{d_1 \cdot d_2 \cdot d_3}{p_{\text{mm}}}$.

Performance and resource demands for the different parameter sets and design variants are shown in Table 3. Since the arithmetic latencies of the field-specific variants and the joint design are identical, the number of cycles depends only on the matrix dimensions d_1, d_2, and d_3 as well as the degree of parallelization p_{mm}. The maximum frequency is in about the same range for all cases with about 10% variance. The joint design, however, requires more computational logic resources for the joint field arithmetic (see Section 4.1) and for supporting different matrix dimensions in the same design.

4.3 Matrix Systemization

Our hardware design for matrix systemization follows the state-of-the art such as [SWM$^+$10,WSN16,WSN17,CCD$^+$22,ZZC$^+$23,BCH$^+$23]: We are using a processor array of a quadratic shape that processes column-blocks of the input matrix in several rounds consisting of several steps. For a detailed description of the operation of the systemizer, we refer to [CCD$^+$22, Section 3].

Table 3: Comparison of the time and area for our Matrix Multiplication module targeting Xilinx Artix 7 (`xc7a200t`) FPGA.

			Resources							
d_1, d_2, d_3	p_{mm}	q	**Area**		**Memory**		**Cycles**	**Freq.**	**Time**	**Time×Area**
			(LUT)	(DSP)	(FF)	(BRAM)	(cyc.)	(MHz)	(us)	
Parameter-set specific designs:										
$14, 14, 14$	2	4 093	279	2	197	0.5	1 380	218	6.34	2×10^3
$22, 22, 22$	2	4 093	294	2	206	0.5	5 332	218	24.50	7×10^3
$30, 30, 30$	2	2 039	298	2	206	0.5	13 508	195	69.20	21×10^3
$14, 14, 14$	4	4 093	383	4	317	1.0	792	218	3.64	1×10^3
$22, 22, 22$	4	4 093	396	4	332	1.0	2 912	218	13.38	5×10^3
$30, 30, 30$	4	2 039	387	4	320	1.0	7 208	218	33.11	13×10^3
$14, 14, 14$	8	4 093	598	8	565	1.5	400	218	1.84	1×10^3
$22, 22, 22$	8	4 093	625	8	580	1.5	1 460	199	7.32	5×10^3
$30, 30, 30$	8	2 039	606	8	556	1.5	3 608	199	18.09	11×10^3
Joint designs:										
$14, 14, 14$	2	both	335	2	220	0.5	1 380	199	6.93	2×10^3
$22, 22, 22$							5 332		26.79	9×10^3
$30, 30, 30$							13 508		67.88	23×10^3
$14, 14, 14$	4	both	460	4	348	1.0	792	217	3.65	2×10^3
$22, 22, 22$							2 912		13.42	6×10^3
$30, 30, 30$							7 208		33.23	15×10^3
$14, 14, 14$	8	both	753	8	612	1.5	400	196	2.04	2×10^3
$22, 22, 22$							1 460		7.43	6×10^3
$30, 30, 30$							3 608		18.36	14×10^3

The design in [BCH+23] combines matrix systemization and matrix multiplication in one design with the goal to share resources between these two operations. However, we are using a pipelined overall design to increase signing and verification performance and hence we need individual modules for matrix systemization and multiplication. Due to the pipelined overall design, all modules are under load most of the time (as shown in Table 7) and resources are in use in parallel — and hence do not need to be shared to achieve efficiency.

The design in [ZZC+23] reduces memory cost by streaming out parts of the systemized matrix as complete columns become available during computation without storing the entire matrix on-chip. This optimization does not apply to MEDS (except if external memory is used), since the large $k \times nm$ matrices $\tilde{\mathbf{G}}_\mathbf{i}$ and $\hat{\mathbf{G}}_\mathbf{i}$ need to be fed row-wise into the hash function.

In contrast to earlier designs aiming at the systemization of matrices over $\mathbb{F}_2$ as they appear, e.g., in binary codes used for the code-based McEliece cryptosystem [SWM+10,WSN16,CCD+22] or matrices over binary fields as used, e.g., in code-based [WSN17] as well as multivariate cryptosystems [FG18,BCH+23], we are dealing with matrices over medium-sized prime fields, which means that we are dealing with a longer arithmetic latency: while binary field arithmetic

Table 4: Comparison of the time and area for our Systemizer module targeting Xilinx Artix 7 (`xc7a200t`) FPGA.

			Resources							
Matrix Size	p_{mm}	q	**Area**		**Memory**		**Cycles**	**Freq.**	**Time**	**Time×Area**
rows×cols			(LUT)	(DSP)	(FF)	(BRAM)	(cyc.)	(MHz)	(us)	
					Parameter-set specific designs:					
14×196	2	4 093	943	6	1 076	1.5	9 345	185	50.52	48×10^3
22×484	2	4 093	1 067	6	1 334	1.5	57 393	177	324.21	346×10^3
30×900	2	2 039	1 178	6	1 511	1.0	199 393	199	1 002.15	$1\,181 \times 10^3$
14×196	4	4 093	2 172	20	2 670	3.0	2 718	177	15.39	33×10^3
22×484	4	4 093	2 288	20	3 155	3.0	15 702	183	85.94	197×10^3
30×900	4	2 039	2 388	20	3 408	2.0	53 222	173	308.32	736×10^3
14×196	8	4 093	6 442	72	7 030	7.5	794	193	4.12	27×10^3
22×484	8	4 093	6 645	72	8 192	7.5	4 069	166	24.54	163×10^3
30×900	8	2 039	6 783	72	8 939	5.0	13 490	170	79.24	537×10^3
					Joint designs:					
14×196	2	both	1 263	6	1 595	2.5	9 345	177	52.67	67×10^3
22×484							57 393		323.47	409×10^3
30×900							199 393		1 123.78	$1\,419 \times 10^3$
14×196	4	both	2 778	20	3 685	5.0	2 718	173	15.72	44×10^3
22×484							15 702		90.82	252×10^3
30×900							53 222		307.84	855×10^3
14×196	8	both	8 333	72	9 138	10.0	794	169	4.69	39×10^3
22×484							4 069		24.04	200×10^3
30×900							13 490		79.71	664×10^3

in earlier work has a latency of one cycle, we decided to deal with the larger arithmetic latency by using several clock cycles for multiplication and addition. This does not affect the overall process of the systemization, but we need to take the latencies into consideration to avoid race conditions during the parallel processing of that data. For example, our design needs to wait for different amounts of overhead cycles between consecutive rounds to wait for data of the previous round being written back to memory depending on the number of column blocks that are being processed in between. Depending on the size of the processor array, the arithmetic latencies add up and we quickly get an overall latency of the processor array of dozens of cycles. However, in contrast to earlier designs [WSN16,WSN17,CCD+22] we do pipeline computation between consecutive rounds as much as possible.

Similar to [BCH+23], we are not too concerned about systemization failures, which happen only very rarely for our field sizes: The probability that a random $n \times m$ matrix ($m \geq n$) over a finite field $\mathbb{F}_q$ has a systematic form is the same as that of an $n \times n$ matrix over $\mathbb{F}_q$ being full rank. For large n and q, this probability is approximately $1 - \frac{1}{q}$. Hence, for $q = 4093$, we have a probability of success for the SF operation in line 20 of about 99.98% and for $q = 2039$ of about 99.96%.

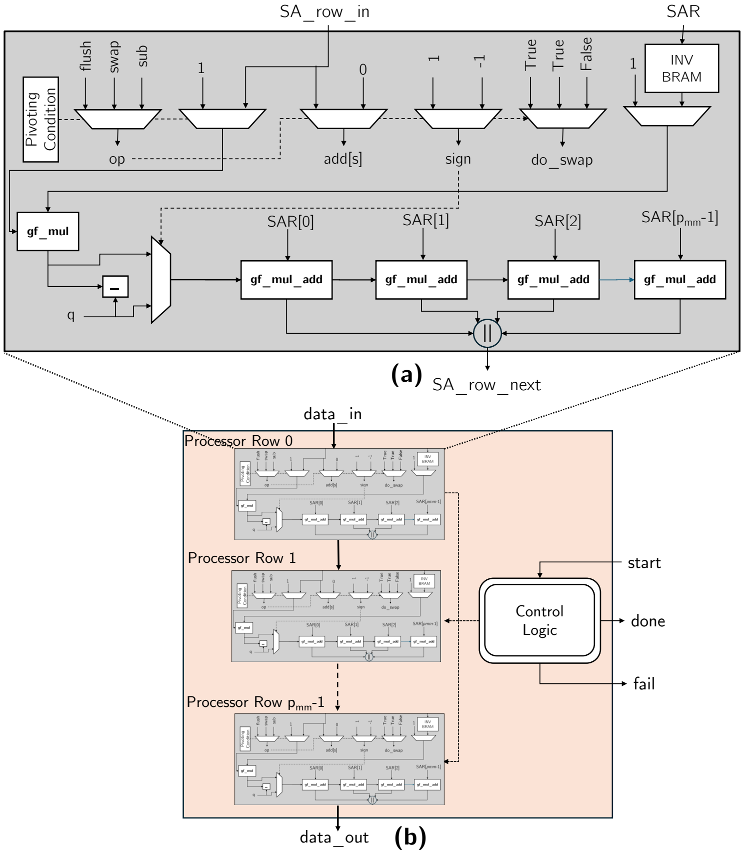

Fig. 3: Systemizer module with p_{sys} vector ALUs of width p_{sys}.

In other words, we expect a failure to systemize a matrix over $\mathbb{F}_q$ roughly every q attempts — for $q = 4093$ we expect one failure roughly every 4093 and for $q = 2039$ roughly every 2039 systemization attempts. Therefore, we omit some of the tweaks described in [CCD+22]: As we do not require a mechanism to detect failure as early as possible but only at the end of systemization, we do not need store all operations that are generated in pivoting steps in-place in the data memory for later use. Instead, we can use a relatively small additional memory as in [WSN16,WSN17,BCH+23] to store the operations only of the current pivoting step, overwriting those of previous steps. Similar to these designs, we detect systemization failure at the end of the computation in the final pivoting step.

Figure 3 shows an overview of the systemizer module. We use the performance parameter p_{sys} to specify the number of both rows and columns of the processor array. The number of columns corresponds to the column-block width of the corresponding BRAM that stores the matrix data. We are grouping the processor elements of each row into a vector Arithmetic Logic Units (ALUs) of width p_{sys}.

The design operates by feeding matrix data in column blocks row-by-row into the processor array, one row in each cycle. Once a pivot matrix row reaches the corresponding processor row, it is stored in registers `SAR[i,j]`. The pivoting rows need to be normalized, which requires a scalar inversion and a scalar-vector multiplication. Later in the process, incoming rows are reduced by the pivoting row, which requires a scalar multiplication and a subtraction.

Now, as mentioned before, the prime-filed arithmetic operations have a considerable latency. Since processor array operates in a pipelined fashion, the result of the normalization is already required for the next input. Overcoming the latencies of the normalization would require several buffer stages. Also, we would like to use only one vector unit computing scalar-vector multiplication and vector addition, not two for normalization and reduction.

Our solution to this problem is to delay the normalization of pivoting rows until they are read out of the processor array. This means that a new pivot row can be stored in `SAR[i,j]` unmodified and is available without further latency in the next cycle. For the normalization and reduction we now can use just one single vector ALU, but we need to prepare the respective inputs at the cost of several cycles of latency. The vector ALU now always multiplies `SAR[i,j]` by a scalar factor either for normalization or for reduction and adds an input vector to the result which is either all zero for normalization or input matrix data for reduction. We need one additional scalar-scalar multiplication per processor row for normalizing the pivot row towards reducing follow-up rows.

The top part of Figure 3 shows the pipelined modules that prepare the inputs for and control the behavior of the vector unit and the bottom part shows the vector unit itself with p_{sys} multiply-and-add units. The overall latency of each processor row is defined by the sequential operation of the top part and the parallel latency of bottom part as the latencies of one multiplication, one addition, and one multiply-and-add operation plus some additional latencies to optimize data paths of control logic.

The performance of our systemizer design is shown in Table 4. The maximum frequency ranges between 166 MHz and 199 MHz. The joint design requires the same number of cycles but its maximum frequency is rather at the low end of the range.

4.4 Pi Module

The algorithm for the π operation is given in [CNP+23a, Algorithm 7] and was introduced in Section 2.1. To recap, the $\pi_{\mathbf{A},\mathbf{B}}(\mathbf{G})$ operation for $\mathbf{A} \in \mathbb{F}_q^{m\times m}$, $\mathbf{B} \in \mathbb{F}_q^{n\times n}$, and $\mathbf{G} \in \mathbb{F}_q^{k\times mn}$ is defined as follows: arrange the rows of matrix $\mathbf{G}$ into a k separate $m\times n$ matrices $\{\mathbf{P}_0, ..., \mathbf{P}_{k-1}\}$ and then compute $\mathbf{P}'_i = \mathbf{A}\cdot\mathbf{P}_i\cdot\mathbf{B}$. Finally, we get the result as matrix $\mathbf{G}' \in \mathbb{F}_q^{k\times mn}$ by packing $\mathbf{P}'_i$ into row i of $\mathbf{G}'$.

In our implementation of the π operation, we avoid the first step of arranging the rows of matrix $\mathbf{G}$ into k separate $m \times n$ matrices. Instead, while generating matrix $\mathbf{G}$ using the `ExpandSystMat` module, the matrix is already arranged in k separate BRAM units. Therefore, we can directly perform the $\mathbf{A}\cdot\mathbf{P}_i\cdot\mathbf{B}$ operation

Table 5: Comparison of the time and area for our Pi module targeting Xilinx Artix 7 (`xc7a200t`) FPGA for $p_{\pi mm} = 4$.

			Resources							
k, m, n	$p_{\pi pm}$	q	**Area**		**Memory**		**Cycles**	**Freq.**	**Time**	**Time×Area**
			(LUT)	(DSP)	(FF)	(BRAM)	(cyc.)	(MHz)	(us)	
Parameter-set specific designs:										
$14, 14, 14$	2	4 093	1 329	8	660	2.0	25 474	218	117.03	155.0×10^3
$22, 22, 22$	2	4 093	1 610	8	706	2.0	91 354	218	419.68	675.0×10^3
$30, 30, 30$	2	2 039	1 655	8	688	2.0	223 474	218	1 026.64	$1\,697.0 \times 10^3$
$14, 14, 14$	4	4 093	2 172	16	1 248	4.0	13 588	218	62.42	135.0×10^3
$22, 22, 22$	4	4 093	2 204	16	1 313	4.0	48 724	196	249.17	548.0×10^3
$30, 30, 30$	4	2 039	2 457	16	1 269	4.0	119 188	196	609.53	$1\,494.0 \times 10^3$
$14, 14, 14$	8	4 093	3 350	32	2 438	8.0	6 796	218	31.22	104.0×10^3
$22, 22, 22$	8	4 093	3 526	32	2 536	8.0	24 364	218	111.93	394.0×10^3
$30, 30, 30$	8	2 039	4 028	32	2 437	8.0	59 596	218	273.78	$1\,101.0 \times 10^3$
Joint designs:										
$14, 14, 14$	2	both	1 891	8	751	2.0	25 474	217	117.44	197.0×10^3
$22, 22, 22$							91 354		421.14	705.0×10^3
$30, 30, 30$							223 474		1 030.22	$1\,725.0 \times 10^3$
$14, 14, 14$	4	both	2 856	16	1 392	4.0	13 588	217	62.64	190.0×10^3
$22, 22, 22$							48 724		224.62	682.0×10^3
$30, 30, 30$							119 188		549.46	$1\,668.0 \times 10^3$
$14, 14, 14$	8	both	4 732	32	2 661	8.0	6 796	217	31.33	174.0×10^3
$22, 22, 22$							24 364		112.32	623.0×10^3
$30, 30, 30$							59 596		274.74	$1\,523.0 \times 10^3$

using our matrix multiplication unit (described in Section 4.2). We first perform $\mathbf{T} = \mathbf{A} \cdot \mathbf{P}$ and then $\mathbf{P}' = \mathbf{T} \cdot \mathbf{B}$. Figure 2b shows the design of our π module.

We provide a performance parameter $p_{\pi\text{pm}}$ for the k matrix multiplications such that we can select number of matrix multiplications to be performed in parallel (the "pm" strands for parallel matrix multiplications). We control the number of parallel finite field operations in the matrix multiplication units of the π operation individually by setting the performance parameter p_{mm} of the matrix multiplication unit (see Section 4.2) to $p_{\pi\text{mm}}$ for the π operation. This gives us two different performance parameters for our π design: first, the number $p_{\pi\text{mm}}$ of prime field ALUs per matrix multiplication and second, the number $p_{\pi\text{pm}}$ of parallel matrix multiplications. The latency l_π of our π operation is $l_\pi = \frac{k}{p_{\pi\text{pm}}}\left(2l_{\text{mmpipe}} + \frac{mmn+mnn}{p_{\pi\text{mm}}}\right)$ with l_{mm} and l_{mmpipe} as defined in Section 4.2.

The synthesis results for our `Pi` module for different configurations are shown in Table 5. The maximum frequency is similar for all variants. The larger matrix dimensions for $k, m, m = 30$ seem to outweigh the smaller resource requirement of $\mathbb{F}_{2039}$ arithmetic. The resource overhead of the joint design compared to the field-specific variants is moderate.

4.5 Signing and Verification

Figure 4 shows the block diagram of the overall hardware design combining both MEDS sign and verify algorithms detailed in Algorithm 1 and Algorithm 2 respectively by sharing hardware modules between both operations.

In sign, we assume that byte deserialization of the input data for the secret key sk into the seed $\sigma_{\mathbf{G}_0}$ and the matrices $\mathbf{A}_i^{-1}$ and $\mathbf{B}_i^{-1}$ using the `Decompress` module are performed on a higher layer. We also leave the composition of the signed message $\mathsf{msg}_{\mathsf{sig}}$ including the serialization of output data μ_i and ν_i using the `Compress` module to a higher layer and just provide read ports to its components as output ports to the design. We also omit the implementation of the random number generator $\mathsf{Randombytes}$ and instead assume that the random seed δ is provided from a higher level.

Similarly, in verification, we assume that byte deserialization of the input data for the public key pk into the seed $\sigma_{\mathbf{G}_0}$ and deserialization of signed message $\mathsf{msg}_{\mathsf{sig}}$ into p, d, α, and msg are performed at a higher layer.

We note that, in our hardware design, the operations `XOF`, `ExpandInvMat`, `SeedTree`, and `PathToSeedTree` share the same SHAKE256 module as shown in Figure 4 but that we use a separate SHAKE256 module that is shared between `ExpandSystMat`, `ParseHash` for the combined `Compress` and `Hash` operation (line 23 in Algorithm 1 and line 28 in Algorithm 2). The reason for this will be described in the discussion of the loop pipeline later on in this section.

The most expensive operations π and `SF` are in the main for-loops of the sign and verify operations i.e., line 14 to line 22 in Algorithm 1 and line 12 to line 27 in Algorithm 2 respectively. Initially, we had considered to construct a large finite field vector ALU for the use in both π and `SF`. However, the required multiplexing and control logic likely would have been inefficient, complex, error prone, and hard to maintain. Also, π requires about two times more finite field multiplications compared to `SF` but no finite field inversions and the different matrices that are systemized in `SF` and `ExpandInvMat` have different dimensions. A joint large finite field vector ALU would hence lead to a less fine-grained control over performance vs. area trade-offs between those operations. Therefore, to obtain high performance and high efficiency, we decided to implement this loop in a pipelined fashion with four pipeline stages.

Stage 1: This stage consists of operations involved in lines 15 to 18 of Algorithm 1 and lines 20 to 23 and lines 14 to 18 of Algorithm 2. In signing, first a leaf seed (σ_i) concatenated with the salt (α) and the iteration number (i) is extended into two seeds $\sigma_{\tilde{A}_i}$, $\sigma_{\tilde{B}_i}$ using the `XOF` module. Then $\sigma_{\tilde{A}_i}$ and $\sigma_{\tilde{B}_i}$ are used to generate invertible matrices $\tilde{A}_i$ and $\tilde{B}_i$. In our hardware design, we accomplish the generation of $\sigma_{\tilde{A}_i}$ and $\sigma_{\tilde{B}_i}$ using the `SHAKE256` module and the `XOF` interface module by selecting the "XOF mode". The expanded seeds are stored in `sigmaAB` BRAM. Then, the `ExpandInvMat` module uses the seed $\sigma_{\tilde{A}_i}$ and the `XOF` interface and `SHAKE256` module by selecting the "ExpandInvMat" mode. This generates data for matrix $\tilde{A}_i$ and the `Systemizer` module inside the `ExpandInvMat` module is used to check if the matrix is invertible or not.

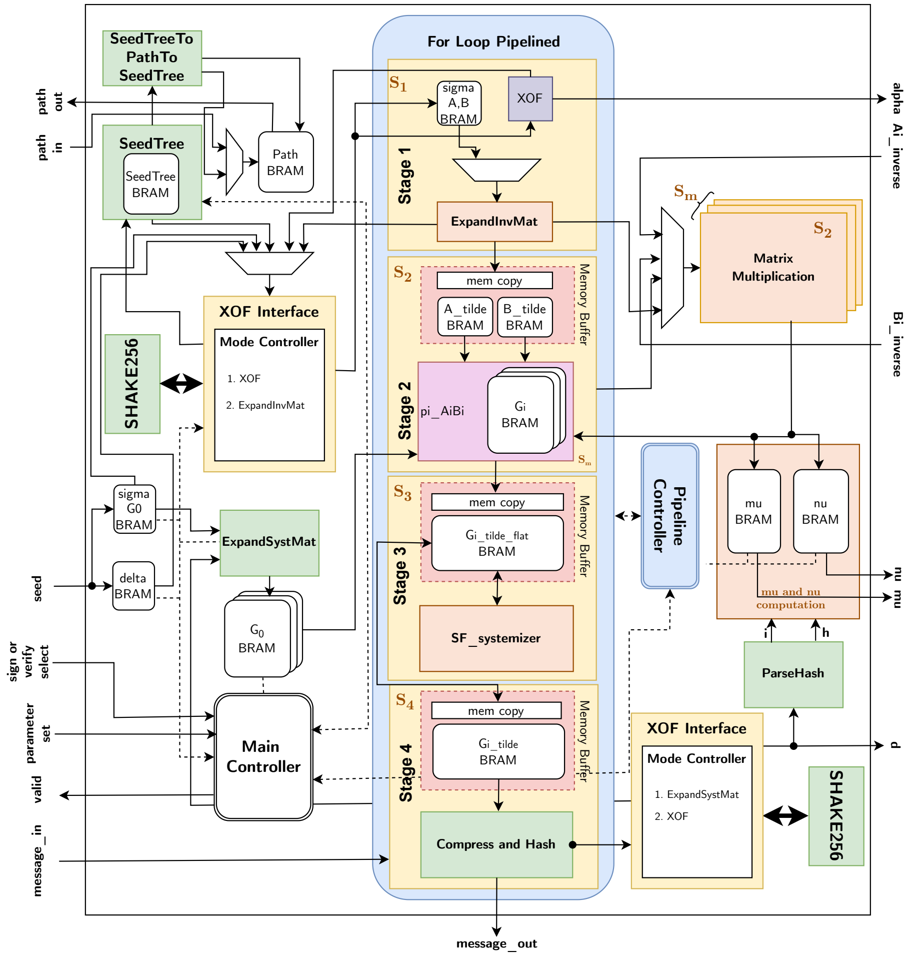

Fig. 4: High-level design overview of the MEDS signing operation.

If it is not invertible then, the module draws another matrix from the current SHAKE256 state and performs another check for invertability of the matrix. This process is repeated until an invertible matrix is found. Once it is ensured that the matrix is invertible then $\tilde{B}_i$ is generated in similar fashion. We note that, for each iteration i, we store the related σ_i value in a BRAM inside Stage 1 for later use in computing μ_i and ν_i matrices.

In verification, if the value of $h_i = 0$ then all operations in stage 1 are similar to signing except here we use $\sigma_{\hat{A}_i}$ and $\sigma_{\hat{B}_i}$ and generate invertible matrices μ_i and ν_i. And in the case of $h_i > 0$, we skip the sampling part inside the ExpandInvMat and instead copy the μ_i and ν_i from the input into the BRAM where sampled μ_i and ν_i are stored and check for their invertability. In case the μ_i or ν_i are non-invertible, then a global "systemization fail" signal is generated, which will trigger "invalid signature" output.

Stage 2: This stage consists of the π operation (i.e., line 19 of Algorithm 1 and line 24 of Algorithm 2). In signing, this operation involves matrix multiplication of rows of $\mathbf{G}_0 \in \mathbb{F}_q^{k\times mn}$ with matrix $\tilde{\mathbf{A}}_i \in \mathbb{F}_q^{m\times m}$ and with $\tilde{\mathbf{B}}_i \in \mathbb{F}_q^{n\times n}$. As shown in Figure 4, in our hardware design, this operation is accomplished by the `Pi` module using $\mathbf{G}_0$ from `G0 BRAM` and $\tilde{\mathbf{A}}_i$, $\tilde{\mathbf{B}}_i$ (generated in Stage 1). For verification, similar operations are performed as for signing except that $\mathbf{G}_{hi}$, μ_i, and ν_i are being processed in place of $\mathbf{G}_0$, $\tilde{\mathbf{A}}_i$, and $\tilde{\mathbf{B}}_i$.

Stage 3: This stage consists of operation SF and the checking if the $\tilde{\mathbf{G}}_i$ has systematic form (i.e., line 20 of Algorithm 1 or line 25 of Algorithm 2). We use a `Systemizer` module (described in Section 4.3) to compute the systematic form. As described in Section 4.3, there is a possibility that the input matrix does not have a systematic form. We describe how we handle this situation in detail below. This step is the same for sign and verify with corresponding inputs.

Stage 4: This stage consists of operations for compressing and hashing matrices $\tilde{\mathbf{G}}_i$ together with the message msg (i.e., line 23 of Algorithm 1). Although the compress and hash operation are not part of the for loop in the algorithm, in our hardware design we move the compression and the hash-absorb operation inside the for loop. This avoids the need to store all $\tilde{\mathbf{G}}_i$ values in the memory, which would require a significant amount of BRAM storage on the FPGA. For this, we use a dedicated `SHAKE256` module for Stage 4. This `SHAKE256` module digests each $\tilde{\mathbf{G}}_i$ and maintains the hash state. Hence, we only need to store one $\tilde{\mathbf{G}}_i$ in the memory at a given time. This step is the same for sign and verify.

After finishing all t iterations of the loop, the message input is absorbed into the state as well and finally we squeeze the hash state to compute the hash value d.

Pipeline Control and Memory Buffering. The pipelined part of our hardware design is highlighted in blue in Figure 4. The pipelining in our hardware design varies from the traditional register-based pipelining: Since the data we move from one stage to another are relatively large matrices, we use BRAMs as buffers between the pipeline stages. These memory buffers are shown as blush pink blocks in Figure 4. The `MemCopy` modules copy the data from the internal memories of each stage and move the data to the memory buffers, which can then be consumed in the next stage.

As the workload of each stage is different, the number of clock cycles taken by each stage is also different. Therefore, to balance the clock cycles taken by each stage, we introduce the following five performance parameters:

S_1: Stage 1 performance parameter that corresponds to the performance parameter p_{sys} (see Section 4.3) of the matrix systemizer for checking the invertability of $\tilde{A}_i$ and $\tilde{B}_i$ in the sign operation respective μ_i and ν_i in verify. S_1 controls the number of rows and columns in the processor array used inside the matrix systemizer.

S_2: Stage 2 performance parameter that corresponds to the performance parameter $p_{\pi\mathrm{mm}}$ (see Section 4.4) of the matrix multiplication unit inside the `Pi` module. S_2 controls the vector width inside the matrix multiplication unit.

S_{mat}: Stage 2 performance parameter that corresponds to the performance parameter $p_{\pi\mathrm{pm}}$ of the `Pi` module used inside stage 2. S_{mat} controls number of matrix multiplications performed in parallel.

S_3: Stage 3 performance parameter that corresponds to the performance parameter p_{sys} (see Section 4.3) of the large $k \times m \cdot n$ matrix systemizer used inside stage 3 for computing the systematic form of $\tilde{G}_i$ in sign respective $\hat{G}_i$ in verify. S_3 controls the number of rows and columns in the processor array used inside the matrix systemizer.

S_4: Stage 4 performance parameter that corresponds to the shifter width used inside the compress and hash module. Larger S_4 leads to fewer memory accesses while compressing and hashing the matrices $\tilde{G}_i$ respective $\hat{G}_i$.

Choosing suitable performance parameters. Recall that the goal behind choosing the performance parameters is to balance the pipeline stages so that all the modules in the pipeline are busy with their respective workloads. Let us look at the methodology we followed using the example (shown in Table 6) for the Security Level I parameter sets: Firstly, we fix the value of $S_{mat} = 1$ and then we assign all possible values to S_1, S_2, S_3, and S_4 based on the MEDS parameters up to $m = n = k = 14$ as shown in Table 6. To further lower the latency of stage 2, the S_{mat} parameter can be tweaked. The latencies reported in Table 6 also include the clock cycles required for buffering data from one stage to another. Based on these latencies, we chose a value for each S_1, S_2, S_3, and S_4 to meet our optimization goal (performance or resource consumption),

As an example for parameters with moderate resource requirement, we can select $S_1 = 1, S_2 = 7, S_3 = 6$, and $S_4 = 6$ (marked in Table 6) such that the cycles counts are roughly equal. Setting $S_{mat} = 4$ in this example brings down the cycle count of stage 2 to about $12\,132/4 = 3033$ (plus the unaffected overhead for data movement between the stages).

Based on this method, we propose different configurations for specific parameter sets as well as for a unified design that aim at balancing the pipeline stages in all parameter sets as shown in Table 7. We propose two variants of configurations: 1) a balanced design choice and 2) a high performance design choice as shown in Table 7. The choice to parameters is not just limited to the combinations provided in Table 7. Several such parameter combinations can be constructed based on resources available on the target device. We note that when selecting the parameters for a unified design, to avoid resource wastage, we limit the values of parameters based on m, n, and k values of smallest parameter set.

Data flow. To cope with the different ready-times of each stage that results from the different cycle counts, we resort to a handshake-based control mechanism to control the flow of data between the stages. After each stage is done with its computation, it sends a handshake signal to the next stage to indicate that the output data is ready for consumption and waits for the response from the

Table 6: Latencies of all pipeline stages for the Security Level 1 parameter sets with performance parameter $S_{mat} = 1$.

$S_1 = S_2 = S_3 = S_4$	Stage 1	Stage 2	Stage 3	Stage 4
1	3 534	83 364	48 191	16 599
2	1 432	41 812	14 877	8 521
3	1 142	29 926	8 350	5 883
4	1 104	23 990	5 702	4 503
5	1 042	18 054	4 020	3 751
6	1 120	18 054	3 557	3 178
7	980	12 132	2 477	2 782
8	1 032	12 118	2 294	2 522
9	1 084	12 118	2 153	2 484
10	1 136	12 118	2 054	2 371
11	1 188	12 118	1 955	2 652
12	1 240	12 118	1 912	2 197
13	1 292	12 118	1 869	2 137
14	975	6 195	1 643	1 796

Table 7: Configuration table showing selected configurations for $S_1, S_2, S_3, S_4, S_{mat}$ and idle percentages for the different stages for each configuration. In case of the unified design the idle percentages are average of different security levels.

Design Choice	**Parameter Set**	q	**Perf. Metrics**					**Stage-wise Idle %**			
			S_1	S_2	S_3	S_4	S_{mat}	**1**	**2**	**3**	**4**
Balanced	Level I	4093	1	7	6	6	3	3.6%	0.0%	3.4%	6.1%
	Level III	4093	1	11	8	9	5	0.5%	1.7%	0.0%	0.5%
	Level V	2039	1	10	10	12	5	0.0%	4.4%	2.8%	5.8%
	Unified Level I–V	both	1	10	10	12	5	3.4%	5.5%	15.4%	15.7%
High Perf.	Level I	4093	2	14	14	14	4	5.7%	3.8%	2.4%	0.0%
	Level III	4093	2	11	22	22	8	1.5%	0.3%	0.0%	3.9%
	Level V	2039	2	15	14	15	8	0.2%	2.7%	0.6%	0.0%
	Unified Level I–V	both	2	14	14	14	4	18.4%	1.3%	17.5%	13.9%

next stage. This handshake mechanism is handled by the `PipelineController` module shown in Figure 4.

As mentioned above in Stage 3, an additional challenge in the pipeline design is the handling of possible systemization failures in some iteration i in Stage 3. In this case, we need to repeat the operations of the previous stages starting again at iteration i. We flush all the data from Stages 1 and 2, which are working on data that belongs to iterations $i + 1$ and $i + 2$, restart Stage 1 from iteration i, and step by step refill the pipeline. Stage 4 is not affected by this as it is working on data that belongs to iteration $i - 1$. The reason we restart from iteration i is because of the on-the-fly `CompressHash` operation in Stage 4. The data fed into the `SHAKE256` module has to be in sequential order to produce the correct hash value. This means that we would need memory buffers after Stage 3 to store data related to iterations $i + 1$ and $i + 2$ and wait until data from iteration i

becomes ready to be loaded into module `CompressHash`. We avoid this expensive memory buffering by simply restarting the pipeline. As the failure probability of the systemizer is low (as specified in Section 4.3), the overhead of restarting the pipeline is marginal. To restart Stage 1 with iteration i, Stage 1 requires to backup the seed for i, i.e., σ_i generated in line 16 of Algorithm 1. After Stage 3 has completed successfully, the stored σ_i is discarded. The logic related to flushing and restaring is also handled by the `PipelineController` module.

Other operations. After the loop is finished, the `CompressHash` module requests the message `msg` as data stream from input port `message_in` and the message is loaded into the `SHAKE256` module to compute the hash value d. Following that, for verification, this d value is compared against the input d value to ensure the verfication of the signature. Where as for signing, the value d is then loaded into the `ParseHash` module shown in Figure 4, which parses the hash value and generates the vector h. The `ParseHash` module also captures i values where $h_i > 0$ while generating t. This is useful for the following computation of μ_i and ν_i. Rather than iterating over all t values, we only iterate over the w indices where $h_i > 0$. To compute matrices μ_i and ν_i, We first generate the matrices $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ using the σ_i seeds stored in Stage 1 and by reusing the logic from Stage 1. The reason for regenerating the matrices $\tilde{\mathbf{A}}_i$ and $\tilde{\mathbf{B}}_i$ is to avoid storing the t - $n \times n$ matrices and t - $m \times m$ matrices, which would otherwise consume a significant amount of BRAM. Therefore, we opted for regenerating the matrices instead. Then, we reuse the `MatMul` units of the `Pi` module. The result is then stored in BRAMs `mu` and `nu` respectively. If the `Pi` module is configured to use multiple matrix multiplication units, the `MuNuComputation` module computes μ_i and ν_i matrices in parallel. In parallel to the computation of μ_i and ν_i, the h vector along with ρ and α are loaded into the `SeedTreeToPath` module, which generates the seed path. The μ_i, ν_i, p, d, α values finally can be accessed by a top-level module through the output ports `mu`, `nu`, `path`, `d`, and `alpha` respectively as shown in Figure 4.

5 Evaluation

Table 8 shows the time and area results for our MEDS signing and verification hardware design for all parameter sets targeting an AMD `xc7a200t-3` FPGA. We note that the maximum clock frequency is limited by two different factors: 1) In case of MEDS-9923, the critical path lies inside the duplicate detection logic inside the `ParseHash` module. We use a bit-vector mapping technique to perform the duplicate detection and since the value t is large in case of MEDS-9923, the fully combinatorial variable shifter and comparator is quite large. 2) In all other cases, the critical path lies inside the sampling unit of `XOF` interface.

As discussed in Section 4.5, we propose two different design variants for our implementation: "Balanced" and "High Performance". The results for these choices are presented in Table 8. While the high performance design requires only 4.5 ms to 29.5 ms for signing and a similar amount of time for verification,

MEDS Parameter Set	**Resources**				**Time**					**Improvment**		
	Area		**Memory**		**Freq.**	**Sign**		**Verify**		**cyc**	**t**	**T×A**
	LUT $\times 10^3$	DSP	FF $\times 10^3$	BRAM	MHz	Mcyc	ms	Mcyc	ms			$\times 10^6$
Balanced – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{1, 7, 6, 6, 3}**												
MEDS-9923	24.3	66	20.5	112	116	7.5	65	7.4	64	69	4.1	3.1
MEDS-13220	21.6	66	19.5	70	130	1.3	10	1.2	10	68	4.6	0.4
High Performance – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{2, 14, 14, 14, 3}**												
MEDS-9923	46.8	259	44.8	193	115	3.3	29	3.3	29	154	9.3	2.7
MEDS-13220	44.1	259	43.8	151	132	0.6	4	0.6	4	154	10.6	0.4
Balanced – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{1, 11, 8, 9, 5}**												
MEDS-41711	34.0	130	31.2	206	123	9.8	80	9.5	78	151	9.8	5.3
MEDS-69497	33.1	130	30.7	187	125	2.9	23	2.5	20	141	9.2	1.4
High Performance – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{2, 11, 22, 22, 8}**												
MEDS-41711	80.5	601	91.0	205	129	6.1	47	5.9	46	242	16.4	7.6
MEDS-69497	79.6	601	90.5	186	129	1.8	14	1.5	12	227	15.4	2.1
Balanced – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{1, 10, 10, 12, 5}**												
MEDS-134180	38.3	163	39.5	285	130	11.2	87	9.0	70	160	10.9	5.9
MEDS-167717	38.0	163	39.4	328	133	8.1	61	5.3	40	141	9.8	3.9
High Performance – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{2, 15, 14, 15, 8}**												
MEDS-134180	61.5	337	60.9	290	127	5.4	43	4.5	36	322	21.5	4.9
MEDS-167717	60.6	337	60.9	288	131	3.9	29	2.7	21	290	19.9	3.0
Unified Design												
Balanced – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{1, 10, 10, 12, 5}**												
MEDS-9923	44.8	163	43.1	344.5	113	5.2	46	5.1	46	100	5.9	4.1
MEDS-13220						0.9	8	0.8	8	98	5.8	0.7
MEDS-41711						13.4	119	13.1	116	110	6.5	10.6
MEDS-69497						3.8	34	3.4	31	105	6.2	2.9
MEDS-134180						11.2	100	9.0	80	160	9.4	8.0
MEDS-167717						8.1	72	5.3	47	141	8.3	5.3
High Performance – $\{S_1, S_2, S_3, S_4, S_{mat}\} =$ **{2, 14, 14, 14, 4}**												
MEDS-9923	72.8	273	63.4	386.5	113	3.3	33	3.3	33	154	8.2	4.2
MEDS-13220						0.5	6	0.5	5	154	8.2	0.7
MEDS-41711						10.2	100	10.1	99	144	7.7	12.9
MEDS-69497						2.7	27	2.6	25	145	7.7	3.3
MEDS-134180						11.1	109	10.7	104	149	7.9	13.7
MEDS-167717						6.7	67	6.2	61	146	7.8	8.2
Reference Software Implementation on an AMD Ryzen 7 PRO 5850U CPU [CNP+23a]												
MEDS-9923	—	—	—	—	1 900	518.0	273	515.5	271	—	—	—
MEDS-13220	—	—	—	—	1 900	88.9	47	87.4	46	—	—	—
MEDS-41711	—	—	—	—	1 900	1 467.0	772	1 461.9	769	—	—	—
MEDS-55604	—	—	—	—	1 900	387.2	204	380.7	200	—	—	—
MEDS-134180	—	—	—	—	1 900	1 629.8	858	1 612.5	849	—	—	—
MEDS-167717	—	—	—	—	1 900	961.8	506	938.8	494	—	—	—

Table 8: Performance data for our `Sign` and `Verify` modules targeting a Xilinx Artix 7 (`xc7a200t`) FPGA. Improvements are calculated as $T \times A = (t_{sign} + t_{verify}) \times LUT$ $cyc = \frac{(Sign+Verify\ Cycles\ for\ Software)}{(Sign+Verify\ Cycles\ for\ Hardware)}$, $t = \frac{(Sign+Verify\ time\ for\ Software)}{(Sign+Verify\ time\ for\ Hardware)}$.

it also takes up significant resources, mainly in terms of BRAM utilization. We note that while the storage required for the data remains the same in both the design choices, due to the wider memory word width in the High Performance variant, the synthesis tool needs to use multiple BRAM units.

Additionally, for each of the design variants, the design allows further flexibility to be synthesized for a specific parameter set or for a unified design that allows us to choose any of the six different parameter choices at run-time. While the unified design does not impact the cycles taken for signing and verification compared to specific designs, it does have an impact on resource utilization. However, it can be seen from Table 8 that the resource utilization of the unified design is close to the utilization of the parameter set specific designs for the biggest parameter sets (i.e., MEDS-134180 and MEDS-167717).

We note that to the best of our knowledge this is the first and only MEDS hardware implementation. Therefore, our primary comparison is with the optimized software implementation provided along with the MEDS specification [CNP+23a]. We note that while our design is running at a frequency range of 115 MHz to 133 MHz, the software implementations results are reported for a CPU running at a frequency of 1.9 GHz. We note that both of our design choices outperform the optimized software implementation by a significant margin as shown in Table 8. Our hardware implementation achieves a speed-up of 69× to 322× in sign + verify cycles compared to the software implementation. Similarly, it achieves a speed-up of 4.1× to 21.5× in sign + verify time, despite operating at a significantly lower frequency than the processor used in the software implementation. This answers Research Question Q1 positively, showing that MEDS algorithms provide inherent parallelism which can be exploited through hardware implementation to provide substantial speed-up as demonstrated by our results (in Table 8).

Comparison to Related Work. Table 9 presents a comparison to a number of other PQC signature schemes that have been presented in literature. We are unaware of other MEDS hardware implementations, so no other MEDS works are included in the table. Many works present data for the same XC7A100T FPGA that we use and they generally report similar frequencies as we. Specific comparisons with these schemes to our implementation is difficult as different schemes are based on different mathematical problems. Further, some related works use high-end FPGAs, which generally give much better performance.

Compared to the hardware (co-)designs of some other PQC schemes shown in Table 9, our hardware implementation of MEDS (results shown in Table 8) makes MEDS a competitive choice when the time area product ($\mathbf{T}\times\mathbf{A}$) is considered. Some of the other designs are much more efficient than our implementation, however this is due to the high computational cost inherent to the MEDS specification in particular compared to lattice-based schemes. Nevertheless, the performance of our implementation is comparable to several other PQC schemes such as SPHINCS+/SLH-DSA, SDitH, and Raccoon. We believe that this shows that the quality of our implementation is on par with other work in this field.

Design	Algorithm	FPGA	Freq.	Sign	Verify	Resources	T×A
			(MHz)	(ms)	(ms)	(LUT/FF/DSP/BR)	
[BCH$^+$23]	ov-Ip-pkc	XC7A200T	100	0.08	0.69	37k/25k/2/81	0.03
[BCH$^+$23]	ov-V-pkc+skc	XC7A200T	100	28.57	3.93	83k/41k/4/359	2.70
[SMA$^+$24]	MAYO 1	Arm/Zynq-7020	100	28.60	—	21k/13k/11/129	0.60
[HSK$^+$23]	MAYO 1	Artix-7	75	0.43	0.05	106k/38k/2/45.5	0.05
[DHSY24]	SDitH L1 GF256	XC7A200T	164	41.00	52.90	17k/9k/0/164.5	1.60
[DHSY24]	SDitH L3 GF251	XC7A200T	164	276.10	183.60	34k/31k/472/521.5	15.63
[dPRS23]	Raccoon-128 2 shares	RISC-V/XC7A100T	78	30.70	18.40	10k/4k/3/—	0.49
[dPRS23]	Raccoon-128 32 shares	RISC-V/XC7A100T	78	284.10	17.86		3.02
[WJW$^+$19]	XMSS SHA256 $h = 10$	RISC-V/Cyclone V	145	9.95	5.80	7k/10k/—/145	0.11
[LSG21]	Dilithium-III F	XC7A100T	145	0.85	0.23	30k/11k/45/21	0.03
[BNG21]	Dilithium-V	Artix-7	116	0.21	0.12	53k/28k/16/29	0.02
[Saa24]	SLH-DSA-SHAKE-128f	RISC-V/XC7A100T	100	49.00	4.40	14k/—/—/—	0.75
[Saa24]	SLH-DSA-SHA2-256s	RISC-V/XC7A100T	100	69 620.10	8.90		974.81
[ALCZ20]	SPHINCS+-128f-simple	XC7A100T	250	1.01	0.16	48k/73k/1/11.5	0.06
[ALCZ20]	SPHINCS+-256s-robust	XC7A100T	250	36.10	0.20	50k/76k/1/30	1.82
[SAW$^+$23]	Falcon-512	ZCU104	188	4.20	—	23k/26k/101/23	0.17
[SAW$^+$23]	Falcon-512	ZCU104	214	—	0.62	12k/8k/15/13	
[SAW$^+$23]	Falcon-1024	ZCU104	188	8.70	—	45k/41k/182/37	0.58
[SAW$^+$23]	Falcon-1024	ZCU104	214	—	1.30	13k/9k/2/4	

Table 9: Existing FPGA-based hardware implementations of various PQC signature schemes. For the listed related work, if the prior work implemented different variants of an algorithm, the fastest design is listed. The "—" indicates that the corresponding parameter was not specified or not implemented. $\mathbf{T}\times\mathbf{A} = \frac{(t_{sign}+t_{verify})\times LUT}{10^6}$

Our results emphasize that MEDS is computationally expensive but that FS-schemes with a large round number can be parallelized efficiently and can achieve high performance (at high resoruce cost). Possible future works include the implementation of key generation and of recent proposals for the optimization of MEDS as well as the exploration of supporting different fields in one joint design for other applicable schemes such as CROSS [BBB$^+$24a], LESS [BBB$^+$24b], and PERK [ABB$^+$24].

6 Conclusion

In the introduction in Section 1, we raised three research questions:

Research Question Q1 asks if there is sufficient inherent parallelism in MEDS to speed up the sign and verify operations. Given that there is plenty parallelism on the low level multiplying and systemizing matrices as well as on the high level iterating over t independent computations, there is indeed ample opportunity to accelerate MEDS. We provide performance parameters to control the low level parallelization and we pipeline the main loop, achieving a significant speed up for MEDS signing and verification. For e.g., for security level I, our balanced design achieves a speed-up of 4-5× for both signing and verification times, whereas our high-performance design achieves a speed-up of 9-10× for both signing and verification times when compared to the optimized software reference implementation. Notably, these gains are achieved while our design operates in the frequency range of 115 MHz to 132 MHz, whereas the optimized software implementation runs at 1.9 GHz. However, Q1 also asks about the resource cost of accelerating MEDS. Since the computational cost of MEDS sign and verification is significant, selecting large performance parameters for a high-performance design results in high resource cost.

Research Question Q2 asks to what extent resources can be shared between the sign and verify operations in a shared design implementing both operations. Since the main loop operates very similar in both sign and verify, the resources of sign can be reused by verify with only little control logic overhead. Sign requires some additional computations at the end to compute the responses for the signature, which can simply be skipped by verify.

Research Question Q3 asks about the overhead of supporting all parameter sets selectable at runtime in a single joint design. The overhead for additional control logic for supporting different matrix dimensions and challenge lengths is marginal. To our surprise, supposing arithmetic in multiple prime fields also comes at only a small overhead as for the fields specified in MEDS, most resources can be shared between both fields. Only slightly more logic than needed for the larger field is required to support both fields. We conclude that — if the fields are chosen carefully — supporting multiple fields in a joint hardware implementation is very much feasible. Specifically our results show that arithmetic for pseudo-Mersenne primes with small Hamming distance can be implemented jointly with small overhead.

Acknowledgements

This research has been partially supported by the US government through NSF grant 2332406, by the Korean government through the BK21 FOUR Education and Research Program for Future ICT Pioneers, the IITP grant IITP-2023-RS-2023-00256081, and the NRF grant RS-2023-00277326, and by the Taiwanese government through the NSTC grant 113-2221-E-001-024-MY3.

References

ABB$^+$24. Najwa Aaraj, Slim Bettaieb, Loïc Bidoux, Alessandro Budroni, Victor Dyseryn, Andre Esser, Thibauld Feneuil, Philippe Gaborit, Mukul Kulkarni, Victor Mateu, Marco Palumbi, Lucas Perin, Matthieu Rivain, Jean-Pierre Tillich, and Keita Xagawa. PERK. Technical report, National Institute of Standards and Technology, 2024. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures`.

ADKF70. V. Arlazarov, E. Dinic, M. Kronrod, and I. Faradzev. On economical construction of the transitive closure of a directed graph. *Soviet Mathematics Doklady*, 11(5):1209–1210, 1970.

ALCZ20. Dorian Amiet, Lukas Leuenberger, Andreas Curiger, and Paul Zbinden. FPGA-based SPHINCS$^+$ Implementations: Mind the Glitch. In *23rd Euromicro Conference on Digital System Design, DSD 2020, Kranj, Slovenia, August 26-28, 2020*, pages 229–237. IEEE, 2020.

AMI$^+$23. Aikata Aikata, Ahmet Can Mert, Malik Imran, Samuel Pagliarini, and Sujoy Sinha Roy. KaLi: A Crystal for Post-Quantum Security Using Kyber and Dilithium. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 70(2):747–758, 2023.

BBB$^+$24a. Marco Baldi, Alessandro Barenghi, Michele Battagliola, Sebastian Bitzer, Marco Gianvecchio, Patrick Karl, Felice Manganiello, Alessio Pavoni, Gerardo Pelosi, Paolo Santini, Jonas Schupp, Edoardo Signorini, Freeman Slaughter, Antonia Wachter-Zeh, and Violetta Weger. CROSS — Codes and Restricted Objects Signature Scheme. Technical report, National Institute of Standards and Technology, 2024. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures`.

BBB$^+$24b. Marco Baldi, Alessandro Barenghi, Luke Beckwith, Jean-François Biasse, Tung Chou, Andre Esser, Kris Gaj, Patrick Karl, Kamyar Mohajerani, Gerardo Pelosi, Edoardo Persichetti, Markku-Juhani O. Saarinen, Paolo Santini, Robert Wallace, and Floyd Zweydinger. LESS — Linear Equivalence Signature Scheme. Technical report, National Institute of Standards and Technology, 2024. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures`.

BCC$^+$24. Ward Beullens, Fabio Campos, Sofía Celi, Basil Hess, and Matthias J. Kannwischer. Nibbling MAYO: Optimized Implementations for AVX2 and Cortex-M4. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2024(2):252–275, Mar. 2024.

BCH$^+$23. Ward Beullens, Ming-Shing Chen, Shih-Hao Hung, Matthias J. Kannwischer, Bo-Yuan Peng, Cheng-Jhih Shih, and Bo-Yin Yang. Oil and vinegar: Modern parameters and implementations. *IACR TCHES*, 2023(3):321–365, 2023.

BFV13. Charles Bouillaguet, Pierre-Alain Fouque, and Amandine Véber. Graph-Theoretic Algorithms for the "Isomorphism of Polynomials" Problem. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 211–227. Springer, Heidelberg, May 2013.

BNG21. Luke Beckwith, Duc Tri Nguyen, and Kris Gaj. High-Performance Hardware Implementation of CRYSTALS-Dilithium. In *International Conference on Field-Programmable Technology, (IC)FPT 2021, Auckland, New Zealand, December 6-10, 2021*, pages 1–10. IEEE, 2021.

CCD$^+$22. Po-Jen Chen, Tung Chou, Sanjay Deshpande, Norman Lahr, Ruben Niederhagen, Jakub Szefer, and Wen Wang. Complete and Improved FPGA Implementation of Classic McEliece. *IACR TCHES*, 2022(3):71–113, 2022.

CNP$^+$23a. Tung Chou, Ruben Niederhagen, Edoardo Persichetti, Lars Ran, Tovohery Hajatiana Randrianarisoa, Krijn Reijnders, Simona Samardjiska, and Monika Trimoska. MEDS — Matrix Equivalence Digital Signature. Technical report, National Institute of Standards and Technology, 2023. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures`.

CNP$^+$23b. Tung Chou, Ruben Niederhagen, Edoardo Persichetti, Tovohery Hajatiana Randrianarisoa, Krijn Reijnders, Simona Samardjiska, and Monika Trimoska. Take Your MEDS: Digital Signatures from Matrix Code Equivalence. In Nadia El Mrabet, Luca De Feo, and Sylvain Duquesne, editors, *AFRICACRYPT 23*, volume 14064 of *LNCS*, pages 28–52. Springer Nature, July 2023.

CNRS24. Tung Chou, Ruben Niederhagen, Lars Ran, and Simona Samardjiska. Reducing Signature Size of Matrix-Code-Based Signature Schemes. In Markku-Juhani O. Saarinen and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Oxford, UK, June 12-14, 2024, Proceedings, Part I*, volume 14771 of *Lecture Notes in Computer Science*, pages 107–134. Springer, 2024.

DHSY24. Sanjay Deshpande, James Howe, Jakub Szefer, and Dongze Yue. SDitH in Hardware. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2024(2):215–251, 2024.

DLK$^+$25. Sanjay Deshpande, Yongseok Lee, Cansu Karakuzu, Jakub Szefer, and Yunheung Paek. SPHINCSLET: An Area-Efficient Accelerator for the Full SPHINCS+ Digital Signature Algorithm. *ACM Trans. Embed. Comput. Syst.*, April 2025. Just Accepted.

dPRS23. Rafaël del Pino, Thomas Prest, Mélissa Rossi, and Markku-Juhani O. Saarinen. High-Order Masking of Lattice Signatures in Quasilinear Time. In *2023 IEEE Symposium on Security and Privacy*, pages 1168–1185. IEEE Computer Society Press, May 2023.

FG18. Ahmed Ferozpuri and Kris Gaj. High-speed FPGA Implementation of the NIST Round 1 Rainbow Signature Scheme. In *2018 International Conference on ReConFigurable Computing and FPGAs (ReConFig)*, pages 1–8, 2018.

FS87. Amos Fiat and Adi Shamir. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

GQT21. Joshua A. Grochow, Youming Qiao, and Gang Tang. Average-Case Algorithms for Testing Isomorphism of Polynomials, Algebras, and multilinear forms. In Markus Bläser and Benjamin Monmege, editors, *38th*

International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference), volume 187 of *LIPIcs*, pages 38:1–38:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

HGG07. William Hasenplaugh, Gunnar Gaubatz, and Vinodh Gopal. Fast Modular Reduction. In *18th IEEE Symposium on Computer Arithmetic (ARITH-18 2007), 25-27 June 2007, Montpellier, France*, pages 225–229. IEEE Computer Society, 2007.

HQR89. Bertrand Hochet, Patrice Quinton, and Yves Robert. Systolic Gaussian Elimination over GF(p) with Partial Pivoting. *IEEE Trans. Computers*, 38(9):1321–1324, 1989.

HSK$^+$23. Florian Hirner, Michael Streibl, Florian Krieger, Ahmet Can Mert, and Sujoy Sinha Roy. Whipping the MAYO Signature Scheme using Hardware Platforms. Cryptology ePrint Archive, Paper 2023/1267, 2023. `https://eprint.iacr.org/2023/1267`.

KKPY24. Matthias J. Kannwischer, Markus Krausz, Richard Petri, and Shang-Yi Yang. pqm4: Benchmarking NIST Additional Post-Quantum Signature Schemes on Microcontrollers. Cryptology ePrint Archive, Paper 2024/112, 2024.

LSG21. Georg Land, Pascal Sasdrich, and Tim Güneysu. A Hard Crystal - Implementing Dilithium on Reconfigurable Hardware. In Vincent Grosso and Thomas Pöppelmann, editors, *Smart Card Research and Advanced Applications - 20th International Conference, CARDIS 2021, Lübeck, Germany, November 11-12, 2021, Revised Selected Papers*, volume 13173 of *Lecture Notes in Computer Science*, pages 210–230. Springer, 2021.

MR24. Suraj Mandal and Debapriya Basu Roy. KiD: A Hardware Design Framework Targeting Unified NTT Multiplication for CRYSTALS-Kyber and CRYSTALS-Dilithium on FPGA. In *2024 37th International Conference on VLSI Design and 2024 23rd International Conference on Embedded Systems (VLSID)*, pages 455–460, 2024.

NQT24. Anand Kumar Narayanan, Youming Qiao, and Gang Tang. Algorithms for Matrix Code and Alternating Trilinear Form Equivalences via New Isomorphism Invariants. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part III*, volume 14653 of *Lecture Notes in Computer Science*, pages 160–187. Springer, 2024.

Saa24. Markku-Juhani O. Saarinen. Accelerating SLH-DSA by Two Orders of Magnitude with a Single Hash Unit. *IACR Cryptol. ePrint Arch.*, page 367, 2024. To appear in CRYPTO 2024, August, 2024.

SAW$^+$23. Michael Schmid, Dorian Amiet, Jan Wendler, Paul Zbinden, and Tao Wei. Falcon Takes Off - A Hardware Implementation of the Falcon Signature Scheme. *IACR Cryptol. ePrint Arch.*, page 1885, 2023.

SMA$^+$24. Oussama Sayari, Soundes Marzougui, Thomas Aulbach, Juliane Krämer, and Jean-Pierre Seifert. HaMAYO: A Fault-Tolerant Reconfigurable Hardware Implementation of the MAYO Signature Scheme. In Romain Wacquez and Naofumi Homma, editors, *Constructive Side-Channel Analysis and Secure Design - 15th International Workshop, COSADE 2024, Gardanne, France, April 9-10, 2024, Proceedings*, volume 14595 of *Lecture Notes in Computer Science*, pages 240–259. Springer, 2024.

SWM$^+$10. Abdulhadi Shoufan, Thorsten Wink, H. Gregor Molter, Sorin A. Huss, and Eike Kohnert. A Novel Cryptoprocessor Architecture for the McEliece Public-Key Cryptosystem. *IEEE Trans. Computers*, 59(11):1533–1546, 2010.

TDJ$^+$22. Gang Tang, Dung Hoang Duong, Antoine Joux, Thomas Plantard, Youming Qiao, and Willy Susilo. Practical post-quantum signature schemes from isomorphism problems of trilinear forms. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 582–612. Springer, Heidelberg, May / June 2022.

TYD$^+$11. Shaohua Tang, Haibo Yi, Jintai Ding, Huan Chen, and Guomin Chen. High-Speed Hardware Implementation of Rainbow Signature on FPGAs. In Bo-Yin Yang, editor, *Post-Quantum Cryptography*, pages 228–243, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

WJW$^+$19. Wen Wang, Bernhard Jungk, Julian Wälde, Shuwen Deng, Naina Gupta, Jakub Szefer, and Ruben Niederhagen. XMSS and Embedded Systems. In Kenneth G. Paterson and Douglas Stebila, editors, *SAC 2019*, volume 11959 of *LNCS*, pages 523–550. Springer, Heidelberg, August 2019.

WSN16. Wen Wang, Jakub Szefer, and Ruben Niederhagen. Solving large systems of linear equations over GF(2) on FPGAs. In Peter M. Athanas, René Cumplido, Claudia Feregrino, and Ron Sass, editors, *International Conference on ReConFigurable Computing and FPGAs, ReConFig 2016, Cancun, Mexico, November 30 - Dec. 2, 2016*, pages 1–7. IEEE, 2016.

WSN17. Wen Wang, Jakub Szefer, and Ruben Niederhagen. FPGA-based Key Generator for the Niederreiter Cryptosystem Using Binary Goppa Codes. In Wieland Fischer and Naofumi Homma, editors, *CHES 2017*, volume 10529 of *LNCS*, pages 253–274. Springer, Heidelberg, September 2017.

ZZC$^+$23. Yihong Zhu, Wenping Zhu, Chen Chen, Min Zhu, Zhengdong Li, Shaojun Wei, and Leibo Liu. Mckeycutter: A High-throughput Key Generator of Classic McEliece on Hardware. In *60th ACM/IEEE Design Automation Conference, DAC 2023, San Francisco, CA, USA, July 9-13, 2023*, pages 1–6. IEEE, 2023.

ZZL$^+$23. Yihong Zhu, Wenping Zhu, Chongyang Li, Min Zhu, Chenchen Deng, Chen Chen, Shuying Yin, Shouyi Yin, Shaojun Wei, and Leibo Liu. RePQC: A 3.4-uJ/Op 48-kOPS Post-Quantum Crypto-Processor for Multiple-Mathematical Problems. *IEEE Journal of Solid-State Circuits*, 58(1):124–140, 2023.

Appendix

Algorithm 1: MEDS sign (from [CNP+23a])

```
   Input: secret key $\mathsf{sk} \in \mathcal{B}^{\ell_{\mathsf{sk}}}$, message $\mathsf{msg} \in \mathcal{B}^{\ell_{\mathsf{msg}}}$
   Output: signed message $\mathsf{msg}_{\mathsf{sig}} \in \mathcal{B}^{\ell_{\mathsf{sig}}+\ell_{\mathsf{msg}}}$
1  $f_{\mathsf{sk}} \leftarrow \ell_{\mathsf{sec\_seed}}$;
2  $\sigma_{\mathbf{G}_0} \leftarrow \mathsf{sk}[f_{\mathsf{sk}}, f_{\mathsf{sk}} + \ell_{\mathsf{pub\_seed}} - 1]$;
3  $f_{\mathsf{sk}} \leftarrow f_{\mathsf{sk}} + \ell_{\mathsf{pub\_seed}}$;
4  $\mathbf{G}_0 \in \mathbb{F}_q^{k \times mn} \leftarrow \mathsf{ExpandSystMat}(\sigma_{\mathbf{G}_0})$;
5  forall $i \in \{1, \ldots, s-1\}$ do
6      $\mathbf{A}_i^{-1} \in \mathbb{F}_q^{m \times m} \leftarrow \mathsf{Decompress}(\mathsf{sk}[f_{\mathsf{sk}}, f_{\mathsf{sk}} + \ell_{\mathbb{F}_q^{m \times m}}], m, m)$;
7      $f_{\mathsf{sk}} \leftarrow f_{\mathsf{sk}} + \ell_{\mathbb{F}_q^{m \times m}}$;
8  forall $i \in \{1, \ldots, s-1\}$ do
9      $\mathbf{B}_i^{-1} \in \mathbb{F}_q^{n \times n} \leftarrow \mathsf{Decompress}(\mathsf{sk}[f_{\mathsf{sk}}, f_{\mathsf{sk}} + \ell_{\mathbb{F}_q^{n \times n}}], n, n)$;
10     $f_{\mathsf{sk}} \leftarrow f_{\mathsf{sk}} + \ell_{\mathbb{F}_q^{n \times n}}$;
11 $\delta \in \mathcal{B}^{\ell_{\mathsf{sec\_seed}}} \leftarrow \mathsf{Randombytes}(\ell_{\mathsf{sec\_seed}})$;
12 $\rho \in \mathcal{B}^{\ell_{\mathsf{tree\_seed}}}, \alpha \in \mathcal{B}^{\ell_{\mathsf{salt}}} \leftarrow \mathsf{XOF}(\delta, \ell_{\mathsf{tree\_seed}}, \ell_{\mathsf{salt}})$;
13 $\sigma_0, \ldots, \sigma_{t-1} \in \mathcal{B}^{\ell_{\mathsf{tree\_seed}}} \leftarrow \mathsf{SeedTree}_t(\rho, \alpha)$;
14 forall $i \in \{0, \ldots, t-1\}$ do
15     $\sigma_i' \in \mathcal{B}^{\ell_{\mathsf{salt}}+\ell_{\mathsf{tree\_seed}}+4} \leftarrow \big(\alpha|\sigma_i|\mathsf{ToBytes}(2^{1+\lceil \log_2(t) \rceil} + i, 4)\big)$;
16     $\sigma_{\tilde{\mathbf{A}}_i}, \sigma_{\tilde{\mathbf{B}}_i} \in \mathcal{B}^{\ell_{\mathsf{pub\_seed}}}, \sigma_i \in \mathcal{B}^{\ell_{\mathsf{tree\_seed}}} \leftarrow \mathsf{XOF}(\sigma_i', \ell_{\mathsf{pub\_seed}}, \ell_{\mathsf{pub\_seed}}, \ell_{\mathsf{tree\_seed}})$;
17     $\tilde{\mathbf{A}}_i \in \mathrm{GL}_m(q) \leftarrow \mathsf{ExpandInvMat}(\sigma_{\tilde{\mathbf{A}}_i}, m)$;
18     $\tilde{\mathbf{B}}_i \in \mathrm{GL}_n(q) \leftarrow \mathsf{ExpandInvMat}(\sigma_{\tilde{\mathbf{B}}_i}, n)$;
19     $\tilde{\mathbf{G}}_i \in \mathbb{F}_q^{k \times mn} \leftarrow \pi_{\tilde{\mathbf{A}}_i, \tilde{\mathbf{B}}_i}(\mathbf{G}_0)$;
20     $\tilde{\mathbf{G}}_i \in \mathbb{F}_q^{k \times mn} \cup \{\perp\} \leftarrow \mathsf{SF}(\tilde{\mathbf{G}}_i)$;
21     if $\tilde{\mathbf{G}}_i = \perp$ then
22         goto line 15;
23 $d \in \mathcal{B}^{\ell_{\mathsf{digest}}} \leftarrow$
     $\mathsf{H}(\mathsf{Compress}(\tilde{\mathbf{G}}_0[;k, mn-1]) \,|\, \ldots \,|\, \mathsf{Compress}(\tilde{\mathbf{G}}_{t-1}[;k, mn-1]) \,|\, \mathsf{msg})$;
24 $h_0, \ldots, h_{t-1} \in \{0, \ldots, s-1\} \leftarrow \mathsf{ParseHash}_{s,t,w}(d)$;
25 $f_v \leftarrow 0$;
26 forall $i \in \{0, \ldots, t-1\}$ do
27     if $h_i > 0$ then
28         $\mu_i \in \mathbb{F}_q^{m \times m} \leftarrow \tilde{\mathbf{A}}_i \cdot \mathbf{A}_{h_i}^{-1}$;
29         $\nu_i \in \mathbb{F}_q^{n \times n} \leftarrow \mathbf{B}_{h_i}^{-1} \cdot \tilde{\mathbf{B}}_i$;
30         $v_{f_v} \in \mathcal{B}^{\ell_{\mathbb{F}_q^{m \times m}} + \ell_{\mathbb{F}_q^{n \times n}}} \leftarrow (\mathsf{Compress}(\mu_i) \,|\, \mathsf{Compress}(\nu_i))$;
31         $f_v \leftarrow f_v + 1$;
32 $p \in \mathcal{B}^{\ell_{\mathsf{path}}} \leftarrow \mathsf{SeedTreeToPath}_t(h_0, \ldots, h_{t-1}, \rho, \alpha)$;
33 return $\mathsf{msg}_{\mathsf{sig}} \in \mathcal{B}^{w(\ell_{\mathbb{F}_q^{m \times m}} + \ell_{\mathbb{F}_q^{n \times n}}) + \ell_{\mathsf{path}} + \ell_{\mathsf{digest}} + \ell_{\mathsf{salt}} + \ell_{\mathsf{msg}} = \ell_{\mathsf{sig}} + \ell_{\mathsf{msg}}} =$
     $(v_0 \,|\, \ldots \,|\, v_{w-1} \,|\, p \,|\, d \,|\, \alpha \,|\, \mathsf{msg})$;
```

Algorithm 2: MEDS verify (adapted from [CNP+23a])

```
Input: public key $\mathsf{pk} \in \mathcal{B}^{\ell_{\mathsf{pk}}}$, signed message $\mathsf{msg_{sig}} \in \mathcal{B}^{\ell_{\mathsf{sig}}+\ell_m}$
Output: message $\mathsf{msg} \in \mathcal{B}^{\ell}_{m}$ or $\perp$
$\sigma_{\mathbf{G}_0} \leftarrow \mathsf{pk}[0, \ell_{\mathsf{pub\_seed}} - 1]$;
$\mathbf{G}_0 \in \mathbb{F}_q^{k \times mn} \leftarrow \mathsf{ExpandSystMat}(\sigma_{\mathbf{G}_0})$;
$f_{\mathsf{pk}} \leftarrow \ell_{\mathsf{pub\_seed}}$;
forall $i \in \{1, \ldots, s-1\}$ do
    $\mathbf{G}_i \in \mathbb{F}_q^{k \times mn} \leftarrow \mathsf{DecompressG}(\mathsf{pk}[f_{\mathsf{pk}}, f_{\mathsf{pk}} + \ell_{G_i}])$;
    $f_{\mathsf{pk}} \leftarrow f_{\mathsf{pk}} + \ell_{G_i}$;
$p \in \mathcal{B}^{\ell_{\mathsf{path}}} \leftarrow \mathsf{msg_{sig}}[\ell_{\mathsf{sig}} - \ell_{\mathsf{digest}} - \ell_{\mathsf{salt}} - \ell_{\mathsf{path}}, \ell_{\mathsf{sig}} - \ell_{\mathsf{digest}} - \ell_{\mathsf{salt}} - 1]$;
$d \in \mathcal{B}^{\ell_{\mathsf{digest}}}, \alpha \in \mathcal{B}^{\ell_{\mathsf{salt}}}, \mathsf{msg} \in \mathcal{B}^{*} \leftarrow$
  $\mathsf{msg_{sig}}[\ell_{\mathsf{sig}} - \ell_{\mathsf{digest}} - \ell_{\mathsf{salt}}, \ell_{\mathsf{sig}} - \ell_{\mathsf{salt}} - 1], \mathsf{msg_{sig}}[\ell_{\mathsf{sig}} - \ell_{\mathsf{salt}}, \ell_{\mathsf{sig}} - 1], \mathsf{msg_{sig}}[\ell_{\mathsf{sig}}, ]$;
$h_0, \ldots, h_{t-1} \in \{0, \ldots, s-1\} \leftarrow \mathsf{ParseHash}_{s,t,w}(d)$;
$\sigma_0, \ldots, \sigma_{t-1} \in \mathcal{B}^{\ell_{\mathsf{tree\_seed}}} \leftarrow \mathsf{PathToSeedTree}_t(h_0, \ldots, h_{t-1}, p, \alpha)$;
$f_{m_s} \leftarrow 0$;
forall $i \in \{0, \ldots, t-1\}$ do
    if $h_i > 0$ then
        $\mu_i \in \mathbb{F}_q^{m \times m} \leftarrow \mathsf{Decompress}(\mathsf{msg_{sig}}[f_{m_s}, f_{m_s} + \ell_{\mathbb{F}_q^{m \times m}} - 1], m, m)$;
        $\nu_i \in \mathbb{F}_q^{n \times n} \leftarrow$
          $\mathsf{Decompress}(\mathsf{msg_{sig}}[f_{m_s} + \ell_{\mathbb{F}_q^{m \times m}}, f_{m_s} + \ell_{\mathbb{F}_q^{m \times m}} + \ell_{\mathbb{F}_q^{n \times n}} - 1], n, n)$;
        $f_{m_s} \leftarrow f_{m_s} + \ell_{\mathbb{F}_q^{m \times m}} + \ell_{\mathbb{F}_q^{n \times n}}$;
        if $\mu_i \notin \mathrm{GL}_m(q)$ or $\nu_i \notin \mathrm{GL}_n(q)$ then
            return $\perp$;
    else
        $\sigma'_i \in \mathcal{B}^{\ell_{\mathsf{salt}} + \ell_{\mathsf{tree\_seed}} + 4} \leftarrow \Big(\alpha | \sigma_i | \mathsf{ToBytes}(2^{1+\lceil \log_2(t) \rceil} + i, 4)\Big)$;
        $\sigma_{\hat{\mathbf{A}}_i}, \sigma_{\hat{\mathbf{B}}_i} \in \mathcal{B}^{\ell_{\mathsf{pub\_seed}}}, \sigma_i \in \mathcal{B}^{\ell_{\mathsf{tree\_seed}}} \leftarrow \mathsf{XOF}(\sigma'_i, \ell_{\mathsf{pub\_seed}}, \ell_{\mathsf{pub\_seed}}, \ell_{\mathsf{tree\_seed}})$;
        $\mu_i \in \mathrm{GL}_m(q) \leftarrow \mathsf{ExpandInvMat}(\sigma_{\hat{\mathbf{A}}_i}, m)$;
        $\nu_i \in \mathrm{GL}_n(q) \leftarrow \mathsf{ExpandInvMat}(\sigma_{\hat{\mathbf{B}}_i}, n)$;
    $\hat{\mathbf{G}}_i \in \mathbb{F}_q^{k \times mn} \leftarrow \pi_{\mu_i, \nu_i}(\mathbf{G}_{h_i})$;
    $\hat{\mathbf{G}}_i \in \mathbb{F}_q^{k \times mn} \cup \{\perp\} \leftarrow \mathsf{SF}(\hat{\mathbf{G}}_i)$;
    if $\hat{\mathbf{G}}_i = \perp$ then
        return $\perp$;
$d' \in \mathcal{B}^{\ell_{\mathsf{digest}}} \leftarrow$
  $\mathsf{H}(\mathsf{Compress}(\hat{\mathbf{G}}_0[; k, mn-1]) \mid \ldots \mid \mathsf{Compress}(\hat{\mathbf{G}}_{t-1}[; k, mn-1]) \mid \mathsf{msg})$;
if $d = d'$ then
    return $\mathsf{msg}$;
else
    return $\perp$;
```