

LightCROSS: A Secure and Memory Optimized Post-Quantum Digital Signature CROSS

-Authors' Version-

Puja Mondal¹, Suparna Kundu², Supriya Adhikary¹ and Angshuman
Karmakar¹

¹ Department of Computer Science and Engineering, IIT Kanpur, India

{pujamondal,adhikarys,angshuman}@cse.iitk.ac.in

² COSIC, KU Leuven, Kasteelpark Arenberg 10, Bus 2452, B-3001 Leuven-Heverlee, Belgium

suparna.kundu@esat.kuleuven.be

Abstract. CROSS is a code-based post-quantum digital signature scheme based on a zero-knowledge (ZK) framework. It is a second-round candidate of the National Institute of Standards and Technology's additional call for standardizing post-quantum digital signatures. The memory footprint of this scheme is prohibitively large, especially for small embedded devices. In this work, we propose various techniques to reduce the memory footprint of the key generation, signature generation, and verification by as much as 50%, 52%, and 74%, respectively, on an ARM Cortex-M4 device. Moreover, our memory-optimized implementations adapt the countermeasure against the recently proposed (ASIACRYPT-24) fault attacks against the ZK-based signature schemes.

Keywords: Post-quantum cryptography · Digital signatures · Code-based digital signatures · CROSS · Software implementation · ARM Cortex-M4 · Countermeasure

1 Introduction

In August 2022, the National Institute of Standards and Technology (NIST) released its first set of standards (FIPS 203-5) [NIS24a, NIS24b, NIS24c] for post-quantum (PQ) key-encapsulation mechanisms and digital signature (DS) schemes. One observation that immediately stands out is that these standard schemes are heavily dependent on hard lattice problems. Therefore, one major cryptanalysis breakthrough may jeopardize the whole migration plan from classical to PQC. Further, the current standard DS schemes *i.e.* Dilithium [DLL⁺17], Falcon [FHK⁺20], and SPHINCS+ [BHK⁺19] have very large signatures. This is very problematic for many real-world scenarios, e.g. chain-of-trust-based authentication mechanisms. NIST has also acknowledged these issues and initiated another standardization procedure for additional post-quantum DS schemes [NIS23]. It aims to diversify the portfolio of digital signature schemes and to develop digital signatures with small signature sizes and fast verification.

CROSS [BBB⁺24] is a promising new code-based scheme currently in the second round of the NIST's additional DS standardization procedure [NIS24d]. It is based on the hard restricted syndrome decoding problem (RSDP) [BBC⁺21] in an interactive zero-knowledge framework. One major drawback of this scheme is the huge memory footprint. It is prohibitively large for small microcontrollers, and some of its parameters do not even fit on these platforms. Therefore, reducing the memory footprint is a crucial step for widespread real-world deployment. It is also essential for further progress in the standardization

procedure, as evident from the previous NIST standardization procedures. Therefore, addressing this issue is our major focus in this work. In particular, our contributions are as follows.

- **Memory optimized key generation:** We have analyzed the key generation routine and observed that two big matrices (\mathbf{V} and \mathbf{W}) are the most significant contributors to the large memory footprint consuming almost 2118 – 15150 bytes¹ of memory to generate a vector $\boldsymbol{\eta}$. We have shown that storing the full matrices at a time is unnecessary. Instead, $\boldsymbol{\eta}$ can be generated on the fly. With this modification, we have reduced memory requirement 43% – 50% on an ARM Cortex-M4 microcontroller.
- **Memory efficient signature generation and verification:** In several applications of DS, the key generation is usually performed once and, if required, can be done outside the device. On the other hand, signature generation and verification are performed more frequently and must be performed on-device. Therefore, memory-efficient implementations of these two procedures are even more important. First, in CROSS, these procedures use a large Merkle tree of size $2t - 1$ (9760 – 127424 bytes) and an array of *commitments* (4896 – 63744 bytes) to generate the Merkle root and `MerkleProof` consuming significant memory. We show a memory-efficient technique of Merkle root generation without storing the commitments. Second, in some cases, the input of `HASH` is stored in a big array before applying the `HASH` (e.g., the digests d_1 and d_b generation). We have exploited the sponge construction of `SHAKE` [KjCP16]. We have streamlined the input generation and `xof_shake_update` function to reduce the memory requirement of the input array to a fraction of the original requirement. Third, we removed many redundant memory usages through a thorough execution flow analysis.
- **Countermeasure against ZKFault:** Recently, a fault attack ZKFault [MAKK24], has been shown on zero-knowledge based signature schemes. The authors have shown that their attack can recover the full signing key of CROSS from a single faulted signature. We have observed that this attack is also valid for the latest version of CROSS-v1.2. The paper [MAKK24] also proposed a secure implementation of LESS [BMPS20] against the fault attack. In this work, we have adapted their countermeasure technique for CROSS-v1.2. Further, we have also integrated our memory reduction techniques into this countermeasure. Finally, our implementations of the signature generation and verification algorithm reduce the memory consumption 40% – 52% and 52% – 74% compared to the original implementations with minuscule performance overhead even after the addition of the countermeasure.

We want to note that all our implementations² are constant-time, and we do not introduce any secret dependent branches or any other new side-channel vulnerability. Also, our implementations are fully compatible with the original implementation of CROSS.

2 CROSS and its memory footprint

In this section, we briefly describe the CROSS digital signature. We also describe the primary factors responsible for its large memory requirements. We begin with the notions used in the text. \mathbb{F}_p and \mathbb{F}_p^* represent the finite field of prime order p and the corresponding multiplicative group. The set of all size k vectors over the field \mathbb{F}_p is denoted by \mathbb{F}_p^k and all $k \times n$ matrices over the field \mathbb{F}_p is indicated by $\mathbb{F}_p^{k \times n}$. We denote elements of \mathbb{F}_p or

¹Since the CROSS scheme has several parameter sets, we present the smallest and largest values of the corresponding parameter sets. We follow this notation throughout the text.

²Our code is available at <https://github.com/s-adhikary/pqm4-CROSS>

scalars with lowercase letters, elements of \mathbb{F}_p^k and $\mathbb{F}_p^{k \times n}$ with bold lowercase letters and bold uppercase letters. \mathcal{I}_m and \mathbf{I}_m denote the set $\{0, 1, \dots, m-1\}$ and identity matrix of dimension m , respectively. We denote $\mathbf{a}[i]$ and $\mathbf{A}[i, j]$ to represent the i -th element of the vector \mathbf{a} and (i, j) -th element of the matrix \mathbf{A} respectively. Let $\mathcal{J} \subset \mathcal{I}_n$ be an ordered set, then the notation $\mathbf{a}_{(\mathcal{J})}$ and $\mathbf{A}_{(*, \mathcal{J})}$ (corresponding $\mathbf{A}_{(\mathcal{J}, *)}$) represents the subvector of \mathbf{a} and submatrix of \mathbf{A} formed by selecting the entries of the vector \mathbf{a} and columns (rows) of \mathbf{A} with indices specified in the set \mathcal{J} .

2.1 CROSS signature scheme

Let g be a generator of the multiplicative group \mathbb{F}_p^* of order z . Given the matrix $\mathbf{H} \in \mathbb{F}_p^{n \times (n-k)}$ and $\mathbf{s} \in \mathbb{F}_p^{(n-k)}$, the RSDP problem is to decide if there exists a vector $\mathbf{e} \in \mathbb{E}^n (= \{g^i : i \in \mathcal{I}_z\}^n)$ such that it satisfies $\mathbf{s} = \mathbf{e}\mathbf{H}$. CROSS has another compact variant called RSDPG [BBP⁺24]. Here, the problem is to decide if there exists a vector $\mathbf{e} \in G$ such that the equality $\mathbf{s} = \mathbf{e}\mathbf{H}$ holds where G is a subgroup of \mathbb{E}^n with cardinality $z^m < z^n$. For each of these two variants, CROSS provides three primary parameters satisfying three NIST security levels of 1, 3, and 5. For each of these primary parameters, CROSS further has three optimization corners, namely (i) **fast**, (ii) **balanced**, and (iii) **small**. The **fast** version aims to attain speed, the **small** version targets to reduce the signature size, and the **balanced** version strikes a trade-off between them.

In this section, we discuss the key generation and signature generation algorithms and explain some portions of these algorithms that are required for our work. However, the verification algorithm is quite similar to the signature generation algorithm. Due to the page restriction, we do not describe the verification algorithm in our paper. For more details, we refer to the CROSS documentation [BBB⁺24].

2.1.1 Key generation algorithm (Alg. 1)

The parity check matrix in `CROSS_KeyGen()` is $\mathbf{H} = \begin{bmatrix} \mathbf{V} \\ \mathbf{I}_{n-k} \end{bmatrix}$, where the matrix \mathbf{V} can be generated from the `Seedpk` using `CSPRNG` function. On the other hand, the restricted vector \mathbf{e} is generated from `Seede` using `CSPRNG`. However, for the RSDPG variant, a vector $\boldsymbol{\zeta}$ and a matrix \mathbf{G} are generated from `Seede` and `Seedpk` respectively. Then the matrix $\mathbf{M}_G = [\mathbf{W} | \mathbf{I}_m]$ and the vector $\boldsymbol{\zeta}$ are used to compute the restricted vector \mathbf{e} . The `Seedpk` and $\mathbf{s} = \mathbf{e}\mathbf{H}$ is set as public key and `Seedsk` is our secret key.

Memory footprint of `CROSS_KeyGen()`: Here, we have observed that the matrices \mathbf{V} of order $k \times (n-k)$ and \mathbf{W} of order $m \times (n-m)$ occupy 1368 – 15150 bytes and 750 – 2784 bytes of memory, respectively. Therefore, for RSDP variant 1368 – 15150 bytes (for matrix \mathbf{V}) and for RSDPG, 2118 – 15150 bytes (for \mathbf{V} and \mathbf{W}) are required in `CROSS_KenGen()` Alg. 1.

This algorithm mostly uses `CSPRNG` function. The above two matrices are generated using `CSPRNG` function, which will take `seed` as input and return a random element from the domain \mathcal{D} as outputs. Fig. 1 shows the procedure of generating the matrices \mathbf{V} and \mathbf{W} for both *RSDP* and *RSDPG*. The algorithm `initialized_csprng` first initializes a state. Then, `csprng_randombytes` generates a buffer say `buf` of length r from the state, where r is a fixed value depending on the domain \mathcal{D} and the parameter sets. Next, the `CSPRNG_mat` function generates each element of the domain using the buffer. All of these functions are explained in the documentation of CROSS [BBB⁺24].

$state \leftarrow \text{initialize_csprng}(seed_{pk}, len)$ RSDP	RSDPG
$buf \leftarrow \text{csprng_randombytes}(state, r)$	
$V \leftarrow \text{CSPRNG_mat}(buf, \mathbb{F}_p^{k \times (n-k)})$	
$buf_1 \leftarrow \text{csprng_randombytes}(state, r')$	
$W \leftarrow \text{CSPRNG_mat}(buf_1, \mathbb{F}_z^{m \times (n-m)})$	

Figure 1: Generation of V and W from the seed $seed_{pk}$ **Algorithm 1** CROSS_KeyGen() [BBB⁺24]**Input:** Security parameter λ **Output:** Public key: $(Seed_{pk}, s)$, Private key: $Seed_{sk}$

- 1: $Seed_{sk} \xleftarrow{\$} \{0, 1\}^\lambda$
- 2: $(Seede, Seed_{pk}) \leftarrow \text{CSPRNG}(Seed_{sk}, \{0, 1\}^\lambda \times \{0, 1\}^\lambda)$
- 3: **if** RSDP **then**
- 4: $V \leftarrow \text{CSPRNG}(Seed_{pk}, \mathbb{F}_p^{k \times (n-k)})$
- 5: $\eta \leftarrow \text{CSPRNG}(Seede, \mathbb{F}_z^n)$
- 6: **if** RSDPG **then**
- 7: $(V, W) \leftarrow \text{CSPRNG}(Seed_{pk}, \mathbb{F}_p^{k \times (n-k)} \times \mathbb{F}_z^{m \times (n-m)})$
- 8: $M_G \leftarrow [W | I_m]$
- 9: $\zeta \leftarrow \text{CSPRNG}(Seede, \mathbb{F}_z^m)$
- 10: $\eta \leftarrow \zeta M_G$
- 11: $e_{(\mathcal{I}_n)} \leftarrow g^{\eta_{(\mathcal{I}_n)}}$
- 12: $H \leftarrow \begin{bmatrix} V \\ I_{n-k} \end{bmatrix}$
- 13: $s \leftarrow eH$
- 14: **return** Public key: $(Seed_{pk}, s)$, Private key: $Seed_{sk}$

2.1.2 Signature generation algorithm (Alg. 2)

The signature generation algorithm first constructs t many ephemeral seeds $Seed[0], \dots, Seed[t-1]$ from the master seed $MSeed$ and a random salt $Salt$. The seeds are generated in different ways depending on the parameter sets of the CROSS signature. For “fast” parameter sets of CROSS, the method does not use the Goldreich-Goldwasser-Micali (GGM) [GGM86] tree construction or Merkle tree construction. In this case, using $MSeed \in \{0, 1\}^\lambda$, $Salt \in \{0, 1\}^\lambda$ as input seeds we generate all the $Seeds$ using CSPRNG function. The remaining parameter sets of CROSS construct a GGM tree called $SeedTree$ to generate t many random ephemeral seeds. These seeds and the $Salt$ are utilized to compute u' and η' , which are also used to generate two sets of commitments $cmt_0[0], \dots, cmt_0[t-1]$ (5216 – 63744 bytes) and $cmt_1[0], \dots, cmt_1[t-1]$ (5216 – 63744 bytes) utilizing the $SeedTree$, and has a good contribution in memory-footprint.

After the $SeedTree$ generation, the Merkle tree T is constructed with commitments $cmt_1[i], i \in \mathcal{I}_t$. Then $Hash$ is applied on the commitments cmt_0 and cmt_1 . The output is stored in d_{01} and sent as part of the signature. Next, the first challenge β is generated using the message and commitments. After that, η' , u' and β are used to generate y , which is utilized to construct the second challenge $\mathbf{b} = (\mathbf{b}[0], \dots, \mathbf{b}[t-1])$. Here, $\mathbf{b}[i] \in \{0, 1\}, \forall i \in \mathcal{I}_t$ and let us assume $J = \{i \in \mathcal{I}_t : \mathbf{b}[i] = 1\}$. Then, to ensure that all the seed tree leaves $Seed[i], i \in J$ can be recomputed, $SeedPath$ is calculated by

assembling intermediate seeds from GGM tree. Here, another complete binary tree \mathbf{x} , called **ReferenceTree** is generated from the digest \mathbf{b} , which decides which seed node will be published. Then, the intermediate seed $\text{SeedTree}[c - \text{missing_nodes}[h']]$ will be published if all the leaf nodes corresponding to the subtree rooted as $\mathbf{x}[c]$ have to be published. **SeedPath** is also added as part of the signature together with $(\mathbf{y}_i, \boldsymbol{\sigma}_i)$ (or $(\mathbf{y}_i, \boldsymbol{\delta}_i)$), $\text{cmt}_1[i]$, $i \in J$. Lastly, for the remaining leaves of the seed tree $i \notin J$, **MerkleProofs** is constructed with $\text{cmt}_0[i]$, $i \notin J$ and sent as part of the signature.

Memory footprint of CROSS_Sign(): In the signature generation algorithm, when hash functions are computed, the hash inputs are large and are loaded into the main memory. We have shown that we do not need to load the complete input for hash computation. Rather, we can compute small portions of hash in part. Also, The variables in the signature algorithm are allocated memory when computations are needed, but once the computations are complete, the memory is not reclaimed, even if the data is no longer required for subsequent processes. Finally, in the Merkle proof generation process, the cmt_0 array is copied to the leaf nodes of the Merkle tree. After the tree generation, we compute the Merkle proof; thus, storing cmt_0 array requires considerably large memory (5216 – 63744 bytes). We have discussed our technique to avoid this extra memory consumption.

3 Our memory optimization of CROSS

This section describes our techniques to reduce the memory footprint of CROSS key generation, signature generation, and verification algorithms.

3.1 Our memory optimization of CROSS key generation

As we described earlier, the method **CSPRNG_mat** generate each of the matrix elements of the matrices $\mathbf{V} \in \mathbb{F}_p^{k \times (n-k)}$, $\mathbf{W} \in \mathbb{F}_z^{m \times (n-m)}$ sequentially and stores it. The vectors $\boldsymbol{\eta}$, $\mathbf{e}(=g^\boldsymbol{\eta})$ are also computed and stored. The public vector \mathbf{s} is computed by multiplying the stored matrix \mathbf{V} and vector \mathbf{e} . Similarly, for the **RSDPG** variant, the private vector $\boldsymbol{\zeta}$ is stored and multiplied with \mathbf{W} . The results of these multiplications are as follows:

$$\mathbf{s}[j] = \mathbf{e}[k+j] + \sum_{i=0}^k \mathbf{e}[i] \mathbf{V}[i, j], \text{ for each } 0 \leq j \leq n-k-1.$$

$$\boldsymbol{\eta}[j] = \boldsymbol{\zeta}[m+j] + \sum_{i=0}^m \boldsymbol{\zeta}[i] \mathbf{W}[i, j], \text{ where } 0 \leq j \leq n-m-1.$$

Observe that, each entry $\mathbf{V}[i, j]$ and $\mathbf{W}[i, j]$ is required only once to generate the vectors \mathbf{s} and $\boldsymbol{\eta}$ respectively. Therefore, we compute $\boldsymbol{\zeta}$ before the computation of \mathbf{W} and initialized each $\boldsymbol{\eta}[j]$ by $\boldsymbol{\zeta}[m+j]$. After that we generate (i, j) -th element of \mathbf{W} say $w := \mathbf{W}[i, j]$, we can update $\boldsymbol{\eta}[j]$ as $\boldsymbol{\eta}[j] \leftarrow \boldsymbol{\eta}[j] + \boldsymbol{\zeta}[i] \cdot w$. Therefore, we do not need to store the complete matrix \mathbf{W} at all. The final vector $\boldsymbol{\eta}$ is still the multiplication of matrix \mathbf{W} and vector $\boldsymbol{\zeta}$. Similarly, multiplication of \mathbf{V} and \mathbf{e} is possible without storing \mathbf{V} . In Alg. 3, we have shown the matrix generation technique the authors of CROSS have used. We have modified the method to directly compute multiplication at the time of matrix generation. We have struck through the parts that we have omitted in our updated method, and the parts that we have included instead are coloured blue. As we discussed, to merge the generation process of matrix \mathbf{V} and the computation of \mathbf{s} , only modification of the function **CSPRNG_mat** is not enough as we need to initialize \mathbf{s} to some values of $\boldsymbol{\eta}$. We need to generate the vector \mathbf{s} before the generation of \mathbf{V} . This step will not hamper the security, as the generation of $\boldsymbol{\eta}$ and \mathbf{V} are independent of each other. Similarly, we also need to re-organize the part for the **RSDPG** variant. So, we have modified the steps of

Algorithm 2 CROSS_Sign(Msg, Seed_{sk})**Input:** Message: Msg and private key: Seed_{sk}.**Output:** Signature: signature

```

1: if RSDP then  $(\eta, \mathbf{H}) \leftarrow \text{ExpandPrivateSeed}(\text{Seed}_{\text{sk}})$ 
2: else  $(\eta, \zeta, \mathbf{H}, M_G) \leftarrow \text{ExpandPrivateSeed}(\text{Seed}_{\text{sk}})$ 
3: MSeed  $\xleftarrow{\$} \{0, 1\}^\lambda$ , Salt  $\xleftarrow{\$} \{0, 1\}^{2\lambda}$ 
4: if No_Tree then
5:   (Seed[0], ..., Seed[t - 1])  $\leftarrow \text{ComputeRoundSeed}(\text{Mseed}, \text{Salt})$ 
6: else
7:   Seed_Tree  $\leftarrow \text{SeedTreeGen}(\text{Mseed}, \text{Salt})$ 
8:   (Seed[0], ..., Seed[t - 1])  $\leftarrow \text{SeedTreeLeaves}(\text{Seed\_Tree})$ 
9: for  $i = 0$  to  $t - 1$  do
10:  (Seedu', Seede')  $\leftarrow \text{CSPRNG}(\text{Seed}[i], \{0, 1\}^\lambda \times \{0, 1\}^\lambda)$ 
11:  if RSDP then  $\eta'_i, \mathbf{u}'_i \leftarrow \text{CSPRNG}(\text{Seed}[i] \parallel \text{Salt} \parallel i + c, \mathbb{F}_z^n \times \mathbb{F}_p^n)$ 
12:  else
13:     $\zeta', \mathbf{u}'_i \leftarrow \text{CSPRNG}(\text{Seed}[i] \parallel \text{Salt} \parallel i + c, \mathbb{F}_z^m \times \mathbb{F}_p^n)$ 
14:     $\delta_i \leftarrow \zeta - \zeta', \eta'_i \leftarrow \zeta' M_G$ 
15:     $\sigma_i \leftarrow \eta - \eta'_i$ 
16:     $\mathbf{v}_{(\mathcal{J}_n)} = g^{\sigma_i(\mathcal{J}_n)}$ 
17:     $\mathbf{u} = \mathbf{v} * \mathbf{u}'_i$  // * represents component-wise product
18:     $\tilde{\mathbf{s}} \leftarrow \mathbf{u} \mathbf{H}$ 
19:    if RSDP then  $\text{cmt}_0[i] \leftarrow \text{Hash}(\tilde{\mathbf{s}} \parallel \sigma_i \parallel \text{Salt} \parallel i + c + \text{dsc})$ 
20:    else  $\text{cmt}_0[i] \leftarrow \text{Hash}(\tilde{\mathbf{s}} \parallel \delta_i \parallel \text{Salt} \parallel i + c + \text{dsc})$ 
21:     $\text{cmt}_1[i] \leftarrow \text{Hash}(\text{Seed}[i] \parallel \text{Salt} \parallel i + c + \text{dsc})$ 
22:  $d_0, \text{Merkle\_Tree} \leftarrow \text{MerkleRoot}(\text{cmt}_0[0], \dots, \text{cmt}_0[t - 1])$ 
23:  $d_1 \leftarrow \text{Hash}(\text{cmt}_1[0] \parallel \dots \parallel \text{cmt}_1[t - 1])$ 
24:  $d_{01} \leftarrow \text{Hash}(d_0 \parallel d_1)$ ,  $d_m \leftarrow \text{Hash}(\text{Msg})$ 
25:  $d_\beta \leftarrow \text{Hash}(d_m \parallel d_{01} \parallel \text{Salt})$ 
26:  $\text{beta} \leftarrow \text{CSPRNG}(d_\beta, (\mathbb{F}_p^*)^t)$ 
27: for  $i = 0$  to  $t - 1$  do
28:   $\mathbf{e}'_{(\mathcal{J}_n)} \leftarrow g^{\eta'_i(\mathcal{J}_n)}$ 
29:   $\mathbf{y}_i \leftarrow \mathbf{u}'_i + \text{beta}[i] \mathbf{e}'_i$ 
30:  $d_b \leftarrow \text{Hash}(\mathbf{y}_0 \parallel \dots \parallel \mathbf{y}_{t-1} \parallel d_\beta)$ 
31:  $\mathbf{b} \leftarrow \text{CSPRNG}(d_b, \mathcal{B}_w^t)$ 
32: if No_Tree then
33:  MerkleProofs  $\leftarrow \text{MerkleProof}(\text{cmt}_0[0], \dots, \text{cmt}_0[t - 1], \mathbf{b})$ 
34:  SeedPath  $\leftarrow \text{SeedTreePaths}(\text{Seed}, \mathbf{b})$ 
35: else
36:  MerkleProofs  $\leftarrow \text{MerkleProof}(\text{Merkle\_Tree}, \mathbf{b})$ 
37:  SeedPath  $\leftarrow \text{SeedTreePaths}(\text{Seed\_Tree}, \mathbf{b})$ 
38:  $j \leftarrow 0$ 
39: for  $i = 0$  to  $t - 1$  do
40:  if  $\mathbf{b}[i] = 0$  then
41:    if RSDP then  $\text{rsp}_0[j] \leftarrow (\mathbf{y}_i, \sigma_i)$ 
42:    else  $\text{rsp}_0[j] \leftarrow (\mathbf{y}_i, \delta_i)$ 
43:     $\text{rsp}_1[j] \leftarrow \text{cmt}_1[i]$ 
44:     $j \leftarrow j + 1$ 
45: return signature  $\leftarrow \text{Salt} \parallel d_{01} \parallel d_b \parallel \text{MerkleProofs} \parallel \text{SeedPath} \parallel \text{rsp}_0 \parallel \text{rsp}_1$ 

```

Algorithm 3 CSPRNG_mat_update(buf, $\mathcal{D}(= \mathbb{F}_q^{k \times n})$, \mathbf{a})

Input: A buffer buf, a domain $\mathcal{D} = \mathbb{F}_q^{k \times n}$ of the matrix that is generated from the buffer, a vector \mathbf{a} .

Output: The matrix A is generated from buf. The resultant multiplication $\mathbf{b} = \mathbf{a}A$

```

1: mask =  $2^{\log q} - 1$ 
2: placed = 0,  $i = 8$ 
3:  $j = 64$ ,  $i' = 0$ ,  $j' = 0$ 
4: s_buf = buf[0]
5: while placed <  $kn$  do
6:   if  $j \leq 32$  then
7:     r_buf = buf[ $i$ ]
8:      $i = i + 4$ 
9:     s_buf = (s_buf  $\vee$  r_buf)  $\times 2^j$ 
10:     $j = j + 32$ 
11:    v[placed, =]s_buf  $\wedge$  mask ▷ This value is equivalent to  $A[i', j']$ 
12:    if v[placed] <  $q$  (s_buf  $\wedge$  mask) <  $q$  then
13:      s[ $j'$ ] = s[ $j'$ ] + ( $\eta[i'] \times$  (s_buf  $\wedge$  mask))
14:       $j' = j' + 1$ 
15:      if  $j' = n - k$  then  $j' = 0$ ,  $i' = i' + 1$ 
16:      placed = placed + 1
17:      s_buf = s_buf /  $2^{\log q}$ 
18:       $j = j - \log q$ 
19:    else
20:      s_buf = s_buf / 2,  $j = j - 1$ 

```

CROSS_KeyGen in Alg. 1 and provided the modified version as Updated_CROSS_KeyGen in Alg. 4.

3.2 Optimizing signature generation and verification

This section explains our techniques for memory-efficient signature generation and verification algorithms. Although this work explains the strategy for memory optimization of the signature generation procedure, we have used similar techniques in the verification.

3.2.1 Just-in-time strategy for HASH function

The latest CROSS (v1.2) [BBB⁺24] uses SHAKE as the HASH function. It takes an input buffer (message), say buf, of arbitrary length, say len, and returns a fixed length string as output. However, in the CROSS signature, the hash input is first stored separately e.g., the arrays of cmt₁ (5216 – 63744 bytes) and \mathbf{y} (16830 – 242968 bytes). However, we do not have to generate the complete hash input at once. Instead, we can only generate the first few commitments and proceed with hash generation. After that, the next batch of commitments is generated and fed to the hash generation procedure. This process is repeated until all commitments are passed.

SHAKE has a parameter r , called the rate of the hash. SHAKE takes an input message buf of length len. Then, it divides the whole buffer into blocks of r bits and updates the state. It works as follows:

- First, state is initialized using the function xof_shake_init. After initialization, the function xof_shake_update takes the buffer buf and the length of the buffer

Algorithm 4 Updated_CROSS_KeyGen()**Input:** None**Output:** Public key: $(\text{Seed}_{\text{pk}}, \mathbf{s})$, Private key: Seed_{sk}

```

1:  $\text{Seed}_{\text{sk}} \xleftarrow{\$} \{0, 1\}^\lambda$ 
2:  $(\text{Seede}, \text{Seed}_{\text{pk}}) \leftarrow \text{CSPRNG}(\text{Seed}_{\text{sk}}, \{0, 1\}^\lambda \times \{0, 1\}^\lambda)$ 
3: if RSDP then
4:    $\eta \leftarrow \text{CSPRNG}(\text{Seede}, \mathbb{F}_z^n)$ 
5:    $\text{state} \leftarrow \text{initialized\_csprng}(\text{seed}_{\text{pk}}, \text{len})$ 
6:    $\text{buffer} \leftarrow \text{csprng\_randombytes}(\text{state}, r)$ 
7: if RSDPG then
8:    $\zeta \leftarrow \text{CSPRNG}(\text{Seede}, \mathbb{F}_z^m)$ 
9:    $\text{state} \leftarrow \text{initialized\_csprng}(\text{seed}_{\text{pk}}, \text{len})$ 
10:   $\text{buffer} \leftarrow \text{csprng\_randombytes}(\text{state}, r)$ 
11:   $\text{buffer1} \leftarrow \text{csprng\_randombytes}(\text{state}, r')$ 
12:   $\boldsymbol{\eta}_{(\mathcal{J}_{n-m})} = 0$ 
13:   $\boldsymbol{\eta}_{(\{n-m, \dots, n-1\})} = \zeta_{(\mathcal{J}_m)}$ 
14:   $\boldsymbol{\eta} \leftarrow \text{CSPRNG\_mat\_update}(\text{buffer1}, \mathbb{F}_z^n, \zeta)$ 
15:   $\mathbf{e}_{(\mathcal{J}_n)} = g^{\boldsymbol{\eta}_{(\mathcal{J}_n)}}$ 
16:   $\mathbf{s}_{(\mathcal{J}_{n-k})} = \mathbf{e}_{(\{k, \dots, n-1\})}$ 
17:   $\mathbf{s} \leftarrow \text{CSPRNG\_mat\_update}(\text{buffer}, \mathbb{F}_p^{k \times (n-k)}, \boldsymbol{\eta})$ 
18: return Public key:  $(\text{Seed}_{\text{pk}}, \mathbf{s})$ , Private key:  $\text{Seed}_{\text{sk}}$ 

```

len as inputs and modifies the state . For each message, this function takes the block of r bits and XORs the first r bits of state with it. Then, it permutes the updated state to introduce a pseudorandomness. We continue this process until the remaining bits in the buffer are less than r bits, after which the remaining buffer is only XOR-ed with the state .

- Next, the function `xof_shake_final` function finalize the value of state . Finally, the function `xof_shake_extract` is used to extract the fixed length digest value from the state state .

We have observed that the functions `xof_shake_init`, `xof_shake_final` and `xof_shake_extract` take equal memory for different sizes of input buffers. Therefore, we keep these functions unchanged. However, the function `xof_shake_update` takes the complete buffer as input, but what can be observed is that only a small part of the buffer is used at each step of the hash computation. Therefore, we do not need to generate the complete buffer at once. We have used a technique that only generates a small amount of the buffer and operates on it before generating the next part of the buffer. This approach helps us reduce the memory required to store the complete buffer. We have explained the technique in detail.

- Step 1: We initialize state .
- Step 2: We generate the first few commitments and store them until the total length is $\geq r$. A small part of the generated commitments may not fit in the buffer. Therefore, we fix the buffer size to $\ell = r + \text{digest length}$ bits to store the excess bits.
- Step 3: If the loaded elements in buffer ℓ_1 bits, and $\ell_1 \geq r$ in the previous step, then we set our flag "Flag" to 0. If $\ell_1 = 0$ then we go to Step 6. Otherwise, we set "Flag" to 1.

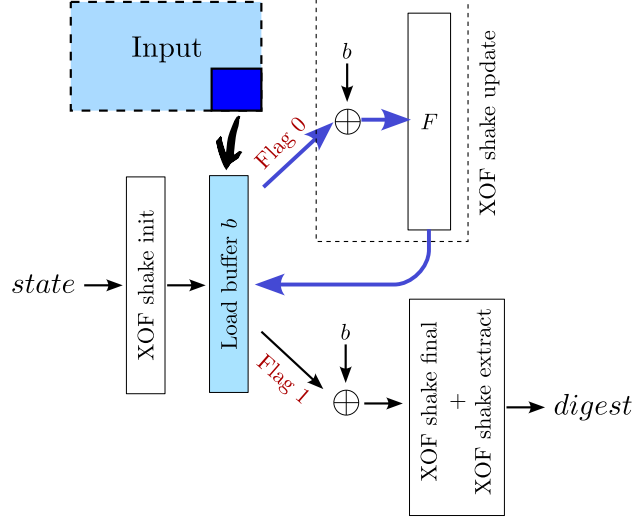


Figure 2: Our streamlined integration of Hash function with input array. Here F denotes the `KeccakF1600_StatePermute` function.

- Step 4: If the Flag is 0, then we use the method `xof_shake_update` and update the `state`. Next put the stored excess part from the previous step at the beginning of the buffer and generate the next few commitments to fill the rest of the buffer and repeat the process from step 2.
- Step 5: If the Flag is 1 then the input of the buffer is XORed with the `state`.
- Step 6: After all the commitments are generated and went through the above process we apply the final `xof_shake_final` and `xof_shake_extract` function to get the digest.

Fig. 2 illustrates the method we have implemented. Our technique does not generate a new hash digest but generates the same digest as the original method.

3.2.2 Efficient memory utilization through execution flow analysis

We performed a thorough analysis of the execution flow of the signature generation algorithm to remove the redundant memory usage leading to a compact implementation. Note that we need to keep a value stored in the memory only if it is going to be used by any operation in subsequent steps. Otherwise, we can overwrite it with any different value without affecting the correctness. We observed that in `CROSS_Sign` (Alg 2), the value $\mathbf{y}_i \in \mathbb{F}_p^n$ is computed as $\mathbf{y}_i = \mathbf{u}'_i + \text{beta}[i] \mathbf{e}'_i$ (line: 29 of Alg 2), where $\mathbf{u}'_i, \mathbf{e}'_i \in \mathbb{F}_p^n$, $\text{beta}[i] \in \mathbb{F}_p^*$, for each $i \in \mathcal{I}_n$. Also, the value \mathbf{u}'_i is not used after the above computation. Therefore, we can overwrite \mathbf{u}'_i with the value of \mathbf{y}_i and the resultant computation would be the same. Therefore, replacing the variables \mathbf{y}_i by \mathbf{u}'_i saves us $16830 - 242968$ bytes (size $tn \log p$) memory as $\mathbf{y} = (\mathbf{y}_0, \dots, \mathbf{y}_{t-1}) \in (\mathbb{F}_p^n)^t$.

In Alg 2, for each $i \in \mathcal{I}_t$ the vector σ_i is computed as $\sigma_i = \eta - \eta_i$ (line: 15). The vector σ_i is used only to compute the commitment $\text{cmt}_0[i]$ for RSDP parameter set and to compute \mathbf{v} . Therefore, we do not need to store the vector σ_i once we have generated vector \mathbf{v} and commitment $\text{cmt}_0[i]$. However, later we need to store the i -th value σ_i for the output response component $\text{rsp}_0[i]$ for RSDP parameter set which satisfies $\mathbf{b}[i] = 0$, where \mathbf{b} is a fixed length random digest vector. For this reason, all the vectors $\sigma_i, \forall i \in \mathcal{I}_t$ are

stored. However, we observe that the cardinality of the set $\{i \in \mathcal{S}_t : \mathbf{b}[i] = 0\}$ rather small (usually in the range 22 – 158). So, if we re-compute the σ_i corresponding to $\mathbf{b}[i] = 0$, then some computations needed to be add-on in lieu of storing the vector σ . This way, we can save 16720 – 242717 bytes of (size $(t - 1)n \log p$) memory by incurring small computation costs.

3.2.3 Re-formulate the Merkle root generation

We have found that two of the most memory-intensive functions in the `CROSS_Sign` algorithm are `Merkle_Root` and `Merkle_Proof` generation. However, this Merkle tree generation is not used for “fast” parameter sets of CROSS. Therefore, the techniques described in this section do not contribute towards memory reduction for the “fast” parameters.

The function `Merkle_Root` uses a `Merkle_Tree` and the commitment array `cmt0` to compute the root. Array `cmt0` and the Merkle tree occupy $t \times \text{digest length}$ (5216 – 63744 bytes) and $(2t - 1) \times \text{digest length}$ (10400 – 127424 bytes) amount of memory, respectively. Also, the function `Merkle_Proof` uses the `cmt0` array and Merkle tree with another binary flag tree with $2t - 1$ nodes. These two functions (`Merkle_Root` and `Merkle_Proof`) use a huge amount of memory. Therefore, the memory reduction in this part will significantly reduce the scheme’s memory consumption.

The `Merkle_Root` function copies each element of the vector `cmt0` one by one to the leaf nodes of the `Merkle_Tree` in an intermediate step. This is the only location where the stored commitment vector `cmt0` is used. We can do the same work by copying the i -th element `cmt0[i]` of `cmt0` to the i -th leaf node of `Merkle_tree` just after `cmt0[i]` generation. To maintain the execution flow of the algorithm, we have to re-organize the steps of the `Merkle_Root` function. We first allocate the required memory for the Merkle tree and then start generating the commitments `cmt0`. After each element of `cmt0` is generated, it is stored in the corresponding leaf node of the Merkle tree. This way, we do not need to allocate memory to store `cmt0`. Finally, we can generate the Merkle root using the remaining parts of the Merkle root generation after the storing process of the leaf nodes of `Merkle_Tree`.

3.3 Countermeasure

The paper [MAKK24] proposed a countermeasure strategy for LESS against ZKFault and claimed that it can be applied to the initial version of CROSS. However, recently, the CROSS signature scheme’s algorithms and implementation have been modified significantly in the updated version. Unfortunately, the prior fault attack is still possible on the new version of CROSS. Therefore, we adapted the countermeasure strategy against the ZKFault attack for the current version of CROSS, along with the memory-optimized implementation in this work. We elaborate on this below.

3.3.1 ZKFault attack on CROSS

The binary vector \mathbf{b} has been checked twice in Alg. 2 (lines: 34/37 depending on RSDP/RSDPG variant and 40). First, whether $\mathbf{b}[i] = 1$ is checked, and if satisfies then the information of `Seed[i]` is published using the function `SeedTreePaths`. Second, whether $\mathbf{b}[i] = 0$ is checked, and if satisfies then `rsp0` is computed (lines:41-42 in Alg. 2). Let us assume that for some i -th location, $\mathbf{b}[i]$ is zero. Then, depending on the RSDP or RSDPG parameter set, the pair (\mathbf{y}_i, σ_i) or (\mathbf{y}_i, δ_i) is stored in `rsp0`. The `Seed[i]` information will not be public in this case. However, if a fault is injected in `SeedTreePaths` algorithm and `Seed[i]`

information is retrieved then η'_i or ζ' can be calculated from $\text{Seed}[i]$. These values help us to find the secret vector η or ζ by computing $\eta = \sigma_i + \eta'_i$ for RSDP version (or $\zeta = \delta_i + \zeta'$ for RSDPG). The security of the scheme will be compromised as the attacker can generate a valid signature of any message from recovered η or ζ .

In the CROSS version v1.2 [BBB⁺24], `SeedTreePaths` is defined for two cases: one is for “fast” parameters, which does not use any tree construction, and the remaining parameter sets that use `SeedTree`. We explain the attack for both cases as follows:

Attack on “fast” parameters: For each $i \in \mathcal{S}_t$, this function `SeedTreePaths` checks whether $\mathbf{b}[i] = 1$. If $\mathbf{b}[i] = 1$ for some $i \in \mathcal{S}_t$, then it stores $\text{Seed}[i]$ in `SeedPath`, which is a component of signature. In this case, if we skip the checking condition for some i which satisfies $\mathbf{b}[i] = 0$, then we will get corresponding seed $\text{Seed}[i]$. The next checking condition, we will get $(\mathbf{y}_i, \sigma'_i)$. As mentioned earlier if we can get both $\text{Seed}[i]$ and the corresponding pair $(\mathbf{y}_i, \sigma'_i)$ (or $(\mathbf{y}_i, \delta'_i)$), then the ZKFault attack will work.

Attack on other parameters: These schemes use a technique that compactly store the `Seed`'s in `SeedPath`. The function `SeedTreePaths` computes a complete binary tree called as `Reference_Tree x` using another function `compute_seeds_to_publish`. The leaf nodes of \mathbf{x} are assigned by the digest vector \mathbf{b} and the remaining nodes are computed as $\mathbf{x}[i] = \mathbf{x}[\text{Left_child}(i)] \wedge \mathbf{x}[\text{Right_child}(i)]$. After the computation of \mathbf{x} , for each node $\mathbf{x}[i]$ (from top to bottom), it checks whether $\mathbf{x}[i] = 1$ and $\mathbf{x}[\text{Parent}(i)] = 0$. If satisfies, then `SeedPath` stores `Seed_Tree[i - missing_nodes[h']]`, where h' is the height of the node $\mathbf{x}[i]$. Let $\mathbf{x}[i] = \mathbf{b}[i - l + 1] = 0$ for some i -th leaf node of the of the tree \mathbf{x} . But, if we change the value of the node $\mathbf{x}[i]$ to 1 by injecting ZKFault, then the `SeedTreePaths` stores the information of the seed $\text{Seed}[i - l + 1] = \text{Seed_Tree}[i - \text{missing_nodes}[\log t]]$ in `SeedPath`. Therefore, the attack will work.

3.3.2 Countermeasure against ZKFault attack

Now, we describe the countermeasure strategy against ZKFault for the fast and other parameter sets of CROSS.

Countermeasure for “fast” parameters: To counter this attack, we will check each value of \mathbf{b} once. For each $i \in \mathcal{S}_t$, if we observe that $\mathbf{b}[i] = 0$, then we will store (\mathbf{y}_i, σ_i) or (\mathbf{y}_i, δ_i) in `rsp0`, otherwise we will store $\text{Seed}[i]$ in `SeedPath`. Since for each digest value $\mathbf{b}[i]$ either (\mathbf{y}_i, σ_i) (or corresponding (\mathbf{y}_i, δ_i)) or $\text{Seed}[i]$ will be published; therefore the attacker can not get both values at a time. Therefore, this attack can be prevented.

Countermeasure for other parameters: To prevent the ZKFault attack in this scenario, we have to compute `SeedPath` and two responses by checking each digest element only once. We can observe that the path of each leaf node of \mathbf{x} toward the root will be of the form $1^r 0^s$ due to the computation process of the tree \mathbf{x} . We check this path of each leaf node starting from the leftmost node of \mathbf{x} ($\mathbf{x}[l - 1]$). Then, we must find the location of the last 1, which contains the index c' and the node height (h') that satisfies the storage condition. Suppose that there is no 1 in this path for node $\mathbf{x}[i]$. In that case, the corresponding element of the digest vector $\mathbf{d}[i - l + 1]$ is zero. We store the corresponding `rsp0` and `rsp1`. Otherwise, we will store `Seed_Tree[c' - missing_nodes[h']]` in `SeedPath` and ignore the next $2^{\log t - h'}$ many nodes. The algorithm `CROSS_Countermeasure_Others` in Alg. 5 shows the exact process of the `rsp0`, `rsp1` and `SeedPath` generation.

Let a value of a node $\mathbf{x}[i]$ be changed from 0 to 1 by the ZKFault attack. Let $\mathbf{x}[i'']$ be the leftmost leaf node corresponding to the subtree rooted as $\mathbf{x}[i]$ and $\mathbf{x}[i']$ be the highest ancestor of node $\mathbf{x}[i]$ (with height h) having the value 1. Therefore, `SeedPath` must store the seed `Seed_Tree[i' - missing_nodes[h]]` and the next $2^{\log t - h}$ leaf nodes will be skipped. Therefore, we cannot get both the values (\mathbf{y}_i, σ_i) and $\text{Seed}[i]$ from this fault preventing the attack.

Algorithm 5 CROSS_Countermeasure_Others()

Input: \mathbf{b} : the digest vector, For RSDP parameter set, $(\mathbf{y}_i, \boldsymbol{\sigma}_i)$, and $(\mathbf{y}_i, \boldsymbol{\sigma}_i)$ otherwise for all $0 \leq i \leq t - 1$.

Output: Compute rsp_0 , rsp_1 , SeedPath

```

1:  $\mathbf{x} \leftarrow \text{compute\_seeds\_to\_publish}(\mathbf{b})$ 
2:  $i = 0, j = 0, j' = 0$ 
3: while  $i < t$  do
4:    $c = l - 1 + i, c' = c, h' = \log t$ 
5:    $\text{flag} = 0, h = \log t$ 
6:   while  $c \neq 0$  do
7:     if  $\mathbf{x}[c] = 1$  then
8:        $c' = c, h' = h, \text{flag} = 1$ 
9:        $c = \text{Parent}(c), h = h - 1$ 
10:  if  $\text{flag} = 0$  then
11:    if RSDP then  $\text{rsp}_0[j] \leftarrow (\mathbf{y}_i, \boldsymbol{\sigma}_i)$ 
12:    else  $\text{rsp}_0[j] \leftarrow (\mathbf{y}_i, \boldsymbol{\delta}_i)$ 
13:     $\text{rsp}_1[j] \leftarrow \text{cmt}_1[i]$ 
14:     $j \leftarrow j + 1, i \leftarrow i + 1$ 
15:  else
16:     $\text{SeedPath}[j'] = \text{Seed\_Tree}[c' - \text{missing\_nodes\_before}[h']$ 
17:     $j' \leftarrow j' + 1, i \leftarrow i + 2^{\log t - h'}$ 
18: Return  $\text{rsp}_0, \text{rsp}_1, \text{SeedPath}$ 

```

4 Results

In this section, we present the results of our implementation of CROSS on the STM32 Nucleo-144 development board with STM32L4R5ZI MCU, featuring an ARM Cortex-M4 chip. For our implementation, we have used the popular post-quantum cryptographic framework PQM4 [KRSS], along with the arm-none-eabi-gcc compiler, version 10.3.1.

We provide the improvement in memory consumption of CROSS RSDPG-1 after applying the optimizations described in this paper in Table 1. S-1 represents the key generation optimization steps described in Sec. 3.1 and the signature generation and verification optimization steps illustrated in Sec. 3.2.1. S-2 and S-3 denote the optimization strategy for the signature generation and verification explained in Sec. 3.2.2 and Sec. 3.2.3, respectively. CM represents the countermeasure of ZKFault described in Sec. 3.3. We also provide the implementation results of the original CROSS-v1.2, along with our improvement for comparison. We observe that S-1 provides 43% improvements in memory consumption for key generation in all the versions. Signature generation improves memory consumption by 27%/15%/16% for **fast/balanced/small** versions in S-1. Verification improves memory consumption by 35%/10%/11% for **fast/balanced/small** versions in S-1. Signature generation and verification improves memory consumption by 37%/32%/33% and 56%/41%/44%, respectively, for **fast/balanced/small** versions in S-2 compared to S-1. In S-3, signature generation and verification improves memory consumption by 9%/9% and 19%/21%, respectively, for **balanced/small** versions. Finally, after including the countermeasure, the overall memory requirements of our CROSS implementation with respect to the original one decreased by 43% for the key generation, 54%/47%/49% for signature generation in **fast/balanced/small** versions, and 71%/57%/60% for the verification in **fast/balanced/small** versions with negligible performance overhead.

In Table 2, we present the memory consumption of our secure implementations of all the versions of CROSS for all the security categories and compare them with the original

Table 1: Improvement steps of our implementations for CROSS RSDPG-1 on ARM Cortex-M4. *bal.* denotes the *balanced* version.

SDPG I	Opt.	Steps	Performance (kcycles)			Memory (bytes)		
			KG	Sign	Verify	KG	Sign	Verify
This work	<i>fast</i>	S-1	248	13,586	7,958	2,216	67,600	30,184
		S-2	248	13,544	7,961	2,216	42,408	13,344
		S-3	248	13,544	7,960	2,216	42,408	13,344
		CM	248	13,548	7,960	2,216	42,408	13,344
v1.2[BBB ⁺ 24]			182	12,710	7,958	3,920	93,096	46,772
This work	<i>bal.</i>	S-1	248	25,021	14,972	2,216	123,792	64,996
		S-2	248	25,107	14,954	2,216	83,784	38,252
		S-3	248	25,069	14,954	2,216	76,568	31,120
		CM	248	25,101	14,961	2,216	76,864	31,120
v1.2[BBB ⁺ 24]			182	24,624	15,030	3,920	145,484	71,860
This work	<i>small</i>	S-1	248	88,592	53,674	2,216	431,220	220,132
		S-2	248	87,905	53,543	2,216	287,668	124,316
		S-3	248	88,238	53,429	2,216	261,620	98,660
		CM	248	88,297	53,409	2,216	263,356	98,668
v1.2[BBB ⁺ 24]	<i>small</i>		182	86,124	53,618	3,920	512,284	247,092

Table 2: Comparison of memory consumption (in bytes) between our implementations of different versions of CROSS with the state-of-the-art implementation on ARM Cortex-M4. KG denotes the key generation algorithm, *bal.* denotes the *balanced* version.

Algo. & sec.	Opt.	CROSS-v1.2[BBB ⁺ 24]			This work		
		KG	Sign	Verify	KG	Sign	Verify
RSDP 1	<i>fast</i>	7,984	120,416	58,296	4,240	59,136	15,220
	<i>bal.</i>	7,984	203,436	92,524	4,240	99,368	33,480
	<i>small</i>	-	-	-	4,240	352,004	108,772
RSDP 3	<i>fast</i>	16,616	264,160	127,024	8,352	158,024	61,100
	<i>bal.</i>	16,616	471,268	211,444	8,352	227,552	73,480
	<i>small</i>	-	-	-	8,352	514,144	157,796
RSDP 5	<i>fast</i>	-	-	-	14,368	227,864	51,648
	<i>bal.</i>	-	-	-	14,368	386,680	121,688
	<i>small</i>	-	-	-	-	-	-
RSDPG 1	<i>fast</i>	3,920	80,664	43,568	2,216	42,408	13,344
	<i>bal.</i>	3,920	145,484	71,860	2,216	76,864	31,120
	<i>small</i>	3,920	512,284	247,092	2,216	263,356	98,660
RSDPG 3	<i>fast</i>	7,336	174,680	93,368	3,936	90,144	26,528
	<i>bal.</i>	7,336	223,656	110,604	3,936	118,320	48,344
	<i>small</i>	-	-	-	3,936	417,164	155,284
RSDPG 5	<i>fast</i>	11,312	306,368	162,140	6,296	156,096	42,076
	<i>bal.</i>	11,312	412,636	203,212	6,296	215,384	84,352
	<i>small</i>	-	-	-	6,296	577,468	214,108

CROSS-v1.2 [BBB⁺24] implementations on the PQM4 framework [KKPY24]. Thanks to our memory optimizations, we could run all parameter sets except one on the board, whereas only 11 out of 18 parameter sets of CROSS-v1.2 fit on the board. Finally, we show 43-50% decrease in memory requirement for the key generation, 40-52% for the signature generation, and 52-74% for the verification compared to the original CROSS.

5 Conclusion

CROSS uses a novel zero-knowledge framework with a quantum hard problem to design a PQ DS scheme. Compared to the other DS schemes such types of constructions are relatively newer has not been explored much. Many practical aspects of such designs, such as efficient and secure implementations, side-channel resistance, and countermeasures, etc., have not been investigated properly yet. Due to their structures, these schemes often have huge memory requirements. We would like to note that our optimization methods are not limited to CROSS and can be applied to other code-based schemes that use similar ZK frameworks, like LESS [BMPS20]. It is another second-round candidate in the NIST additional signature standardization competition. Due to its massive dynamic memory allocation of the commitment array, the implementation of LESS can not run on the board³. We believe that our memory optimization process will help resolve LESS's memory issues. As memory-optimized implementations are essential for real-world applications, we believe this work will benefit the overall advancement of PQC DS schemes.

Acknowledgements

This work was partially supported by Horizon 2020 ERC Advanced Grant (101020005 Belfort), CyberSecurity Research Flanders with reference number VR20192203, BE QCI: Belgian-QCI (3E230370) (see beqci.eu), Intel Corporation, Secure Implementation of Post-Quantum Cryptosystems (SECPQC) DST-India, BELSPO, Google India Research and IIT Kanpur initiation grant. Puja Mondal is supported by C3iHub, IIT Kanpur. Supriya Adhikary is supported by the Prime Minister's Research Fellowship (PMRF), India.

References

- [BBB⁺24] Marco Baldi, Alessandro Barenghi, Sebastian Bitzer, Patrick Karl, Felice Manganiello, Alessio Pavoni, Gerardo Pelosi, Paolo Santini, Jonas Schupp, Freeman Slaughter, Antonia Wachter-Zeh, and Violetta Weger. CROSS: Codes and Restricted Objects Signature Scheme - Specification Document, February 2024.
- [BBC⁺21] Marco Baldi, Massimo Battaglioni, Franco Chiaraluce, Anna-Lena Horlemann-Trautmann, Edoardo Persichetti, Paolo Santini, and Violetta Weger. A new path to code-based signatures via identification schemes with restricted errors, 2021.
- [BBP⁺24] Marco Baldi, Sebastian Bitzer, Alessio Pavoni, Paolo Santini, Antonia Wachter-Zeh, and Violetta Weger. Zero knowledge protocols and signatures from the restricted syndrome decoding problem. In Qiang Tang and Vanessa Teague, editors, Public-Key Cryptography - PKC 2024 - 27th IACR International Conference on Practice and Theory of Public-Key Cryptography, Sydney, NSW, Australia, April 15-17, 2024, Proceedings, Part II, volume 14602 of Lecture Notes in Computer Science, pages 243–274. Springer, 2024.
- [BHK⁺19] Daniel J. Bernstein, Andreas Hülsing, Stefan Kölbl, Ruben Niederhagen, Joost Rijneveld, and Peter Schwabe. The SPHINCS+ Signature Framework. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS '19, page 2129–2146, New York, NY, USA, 2019. Association for Computing Machinery.

³<https://github.com/mupq/pqm4/issues/278>

- [BMPS20] Jean-François Biasse, Giacomo Micheli, Edoardo Persichetti, and Paolo Santini. LESS is More: Code-Based Signatures Without Syndromes. In Abderrahmane Nitaj and Amr Youssef, editors, *Progress in Cryptology - AFRICACRYPT 2020*, pages 45–65, Cham, 2020. Springer International Publishing.
- [DLL⁺17] Leo Ducas, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehle. CRYSTALS – Dilithium: Digital Signatures from Module Lattices. *Cryptology ePrint Archive*, Paper 2017/633, 2017. <https://eprint.iacr.org/2017/633>.
- [FHK⁺20] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU, 2020. <https://falcon-sign.info/falcon.pdf>.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, August 1986.
- [KjCP16] John Kelsey, Shu jen Chang, and Ray Perlner. SHA-3 Derived Functions: cSHAKE, KMAC, TupleHash and ParallelHash, 2016-12-22 00:12:00 2016.
- [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.
- [KRSS] Matthias J. Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. PQM4: Post-quantum crypto library for the ARM Cortex-M4. <https://github.com/mupq/pqm4>.
- [MAKK24] Puja Mondal, Supriya Adhikary, Suparna Kundu, and Angshuman Karmakar. ZKFault: Fault attack analysis on zero-knowledge based post-quantum digital signature schemes. In *Advances in Cryptology - ASIACRYPT, 2024*. just accepted.
- [NIS23] NIST. NIST Announces Additional Digital Signature Candidates for the PQC Standardization Process. Online. Accessed 10th November, 2024, 2023.
- [NIS24a] NIST. FIPS 203 Module-Lattice-Based Key-Encapsulation Mechanism Standard. Online. Accessed 10th November, 2024, 2024.
- [NIS24b] NIST. FIPS 204 Module-Lattice-Based Digital Signature Standard. Online. Accessed 10th November, 2024, 2024.
- [NIS24c] NIST. FIPS 205 Stateless Hash-Based Digital Signature Standard. Online. Accessed 10th November, 2024, 2024.
- [NIS24d] NIST. PQC digital signature second round announcement. Online. Accessed 10th November, 2024, 2024.