

HIGH-ORDER AND CORTEX-M4 FIRST-ORDER IMPLEMENTATIONS OF MASKED FRODOKEM



François Gérard¹ and Morgane Guerreau^{2*}

¹University of Luxembourg, ²PQShield



Introduction

The recent standardization of post-quantum schemes following the NIST standardization process has shown that the field is maturing and calls for further evolutions on the practical side. In particular, side-channel secure implementations have to be developed on various platforms. Here, we will focus on the masking countermeasure for **FRODOKEM**, a KEM recommended by European agencies ANSSI and BSI that is a plain LWE sibling of the standard ML-KEM. While the masking techniques required to protect different parts of the scheme have already appeared in the literature, we make the first step toward securely deploying the scheme by proposing a high-order generic C implementation. Furthermore, we specialize the code at order 1 for Cortex-M4, by rewriting in ARM assembly the basic masking gadgets used by the generic implementation, in order to thwart (micro-)architectural leakage. Our work is validated by performing TVLA on a ChipWhisperer-Lite.

FrodoKEM Decaps

Input: $c = c_1 || c_2 || salt$, $sk = s || seed_A || b || S^T$
Output: ss
 $B' \leftarrow \text{Unpack}(c_1, \bar{n}, n)$
 $C \leftarrow \text{Unpack}(c_2, \bar{n}, \bar{n})$
 $M \leftarrow C - B'S$
 $u' \leftarrow \text{Decode}(M)$
 $seed_{SE'} || k' \leftarrow \text{SHAKE}(u' || salt)$
 $S', E', E'' \leftarrow \text{SampleMatrix}(seed_{SE'})$
 $A \leftarrow \text{Gen}(seed_A)$
 $B'' \leftarrow S'A + E'$
 $B \leftarrow \text{Unpack}(b, n, \bar{n})$
 $V \leftarrow S'B + E''$
 $C' \leftarrow V + \text{Encode}(u')$
 $\bar{k} \leftarrow k' \text{ if } B' || C = B'' || C' \text{ else } \bar{k} \leftarrow s$
 $ss \leftarrow \text{SHAKE}(c_1 || c_2 || salt || \bar{k})$
return ss


What to mask and how?

- Matrix operations: Use the fact that AS can be computed as $AS_0, AS_1, \dots, AS_n$.
- Hashing: Required for the re-encryption through the FO transform. Out of scope of this work ☹️
- Encoding and decoding: Encode maps $\bar{n}^2$ B -bit substrings to values in $\mathbb{Z}_q$ by multiplying them by $q/2^B$. Decode is the inverse. Since q is a power of two, these are shifts, and we use A2B/B2A conversions to compute them.
- Gaussian sampling: Go through a cumulative distribution table. Most expensive operation in masked form.
- Comparison: To verify if a is equal to b , we compute $a - b$ and perform a masked zero test.

Micro-architectural Leakage


While implementing a theoretically validated masking scheme is a good first step toward a secure implementation, it is known that micro-architectural effects can introduce additional sources of leakage by, for example, manipulating shares on a common resource (register, bus). It is thus required to work at low-level to aim at limiting those effects.

```
ldrh rx0, [px], #2
ldrh r, [pool]
ldrh rx1, [px], #2
```




```
ldrh rx0, [px], #2
ldrh r, [pool]
ldrh rx1, [px], #2


strh t1, [pool]
```



```
ldrh rx0, [px], #2
ldrh r, [pool]
ldrh rx1, [px], #2
```



```
eor t2, t2, t2
strh t1, [pool]
```



Experimental validation

TVLA for naive (resp. hardened) ASM versions using 5,000 (resp. 100,000) traces.




Fig. 1: t -tests for SECADD (naive)




Fig. 2: t -tests for SECADD (hardened)




Fig. 3: t -tests for SECZEROTEST (naive)




Fig. 4: t -tests for SECZEROTEST (hardened)

Overhead on Cortex-M4

	C	ASM	ASM _h
SECAND	49	51	68 (+39%)
SECADD	206	149	248 (+20%)
BOOLEANToARITHMETIC	33	47	54 (+64%)
ARITHMETICToBOOLEAN	154	125	222 (+44%)
SECZEROTEST	144	126	223 (+55%)

Tab. 1: Number of cycles on Cortex-M4

Manually hardening the gadgets induces a performance overhead ranging from +20% to +64%. This performance loss remains better than what can be obtained with automated tools.

Performances on x64

Order	1	2	3	4	5	6	7
Key encode	2	152	219	392	596	788	994
Key decode	4	76	113	234	368	483	602
Compare	408	6312	9370	19,358	31,125	40,369	49,756
Sampler	9647	94,967	144,200	223,058	318,535	422,430	531,283

Tab. 2: Benchmarks on x64 of the large gadgets used in FrodoKEM (in kilocycles)

Order	1	2	3	4	5	6	7
FRODO-640	57,395	293,703	498,200	809,345	1,138,708	1,523,870	2,010,644
FRODO-976	93,498	446,945	781,157	1,258,737	1,844,603	2,490,802	3,143,712
FRODO-1344	117,860	522,497	903,360	1,550,615	2,215,093	2,962,697	3,874,540

Tab. 3: Benchmarks on x64 of all versions of FrodoKEM (in kilocycles).

Randomness usage

Order	1	2	3	4	5	6	7
FRODO-640	855	9388	18,632	31,628	47,673	66,720	88,816
FRODO-976	1113	12,524	24,861	42,336	63,884	89,434	119,058
FRODO-1344	1013	12,392	24,611	42,342	64,121	89,850	119,628

Tab. 4: Number of random 16-bit integers required during decaps ($\times 1000$)

Even if they have different parameters, **FRODO-1344** and **FRODO-976** actually need roughly the same amount of randomness because the CDT of **FRODO-1344** is smaller.

ePrint



ia.cr/2025/1065

GitHub



fragerar/masked_Frodo