






Efficient isochronous fixed-weight sampling with applications to NTRU

Décio Luiz Gazzoni Filho^{1,2} , Tomás S. R. Silva³  and Julio López¹ 

¹ Universidade Estadual de Campinas (UNICAMP), Instituto de Computação, Campinas, Brazil

² State University of Londrina, Department of Electrical Engineering, Londrina, Brazil

³ Universidade Estadual de Campinas (UNICAMP), Instituto de Matemática, Estatística e Computação Científica, Campinas, Brazil

Abstract. We present a solution to the open problem of designing a linear-time, unbiased and timing attack-resistant shuffling algorithm for fixed-weight sampling. Although it can be implemented without timing leakages of secret data in any architecture, we illustrate with ARMv7-M and ARMv8-A implementations; for the latter, we take advantage of architectural features such as NEON and conditional instructions, which are representative of features available on architectures targeting similar systems, such as Intel. Our proposed algorithm improves asymptotically upon the current approach based on constant-time sorting networks ($O(n)$ versus $O(n \log^2 n)$), and an implementation of the new algorithm applied to NTRU is also faster in practice, by a factor of up to 6.91 (591%) on ARMv8-A cores and 12.89 (1189%) on the Cortex-M4; it also requires fewer uniform random bits. This translates into performance improvements for NTRU encapsulation, compared to state-of-the-art implementations, of up to 50% on ARMv8-A cores and 72% on the Cortex-M4, and small improvements to key generation (up to 2.7% on ARMv8-A cores and 6.1% on the Cortex-M4), with negligible impact on code size and a slight improvement in RAM usage for the Cortex-M4.

Keywords: Post-quantum cryptography · NTRU · Sampling · ARM

1 Introduction

In the late 1990s, the rise of quantum algorithms for database search and factorization [Gro96, Sho97] posed a threat to public-key cryptosystems based on integer factorization and/or discrete logarithms. Even though quantum computers capable of efficiently performing such computations do not exist yet, growing concern within the community led to seeking alternative cryptographic primitives capable of resisting attacks from quantum algorithms. Thus, *Post-Quantum Cryptography* (PQC) arises as an attempt to counter these attacks by developing new public-key cryptographic algorithms built on problems known to be resistant to quantum attacks, such as lattice-based problems.

One of the oldest lattice-based cryptosystems is NTRU, first presented in the rump session of CRYPTO '96 [HPS96]. It remains relevant, as shown by advancing to the third round of the NIST PQC standardization process [CDH⁺20, Nat17], and its standardization in other forums [Ins09, Ame17]. A performance bottleneck of NTRU is fixed-weight sampling of polynomials, i.e. those with a prescribed number of randomly permuted -1 , 0 and 1 coefficients, employed in key generation and encapsulation. Unless carefully optimized, this sampling incurs a significant runtime cost, particularly to encapsulation.

E-mail: decio.gazzoni@ic.unicamp.br (Décio Luiz Gazzoni Filho), tomas@ime.unicamp.br (Tomás S. R. Silva), jlopez@ic.unicamp.br (Julio López)



Shuffling algorithms appear perfectly suited to solve the problem of fixed-weight sampling; however, there is no known linear-time algorithm that is resistant to timing attacks for this problem. Instead, constant-time sorting is used to generate random permutations, as mandated by the NTRU submission to the NIST PQC contest [CDH⁺20]. We propose a new, timing attack-resistant shuffling algorithm to replace the sorting-based approach, with improved asymptotic running time and large performance improvements in actual implementations, especially for embedded architectures.

Prior to our work, the main proposal to avoid the cost of fixed-weight sampling for NTRU is NTRU-HRSS [HRSS17]. Their technique was later merged into the NTRU proposal for NIST’s PQC standardization process [CDH⁺20]. Due to larger key and ciphertext sizes, it was adopted for only one out of the four suggested parameter sets.

There exist many shuffling algorithms, such as Fisher–Yates [FY38, Dur64, Knu97], Rao–Sandelius [Rao61, San62] and MERGESHUFFLE [BBHL18]. Algorithms in the coin tossing model, aimed at minimizing the consumption of random bits, are reviewed in [BBHT17]. However, none of these are designed to resist side-channel attacks. Indeed, [Dan19] remarks that Fisher–Yates is the most straightforward implementation of fixed-weight sampling, but cautions that “implementing Fisher–Yates in such a way that there is no side channel is difficult.” They opt for a constant-time sorting network [Bat68], as proposed by Bernstein, Chou and Schwabe for use with the McEliece [BCS13] cryptosystem, and Bernstein et al. for the NTRU Prime [BCLv17] cryptosystem.

Sendrier [Sen21] proposed a constant-time Fisher–Yates variant for binary vectors for use with the BIKE cryptosystem. Its running time is $O(w^2)$, where w is the weight (count of non-zero elements) of an n -element vector; since the weight in BIKE grows as $w = O(\sqrt{n})$, its performance behaves as $O((\sqrt{n})^2) = O(n)$. In NTRU, ternary vectors are used, and also, a dense vector is sampled, i.e. $w = O(n)$; thus, Sendrier’s algorithm cannot compete with the $O(n \log^2 n)$ running time of the fastest practical sorting networks.

Our contributions. In §3, we solve the open problem of designing an unbiased linear-time shuffling algorithm resistant to timing attacks for the NTRU fixed-weight sampling problem. It is a drop-in replacement for NTRU’s current sampling-by-sorting approach, improving the running time from $O(n \log^2 n)$ for the best practical sorting networks to $O(n)$, without impacting security. We also discuss its application to other post-quantum cryptosystems. In §4, we discuss implementation aspects. We show in §5 that an implementation of our proposed approach is considerably faster for the fixed-weight sampling step, by factors of up to 6.91 (591%) on ARMv8-A cores and 12.89 (1189%) on the Cortex-M4. This translates into considerable improvements for the KEM encapsulation operation (up to 50% on ARMv8-A cores and 72% on the Cortex-M4) and smaller improvements for key generation (up to 2.7% and 6.1% on the same respective platforms), with little effect on code size, and small gains in memory usage, for embedded architectures. We illustrate how to implement its main operations efficiently in the ARMv8-A and ARMv7-M architectures, as well as generic operations suitable for any architecture, and discuss possible implementations for Intel architectures. Our implementations are available under an open source license in two different GitHub repositories: <https://github.com/dgazzone/NTRU-sampling> for a reference version with ARMv8-A optimizations, and <https://github.com/dgazzone/NTRU-sampling-pqm4> for a Cortex-M4 implementation.

2 Preliminaries

2.1 NTRU random sampling

NTRU is a post-quantum public-key cryptosystem whose security relies on the difficulty of finding short vectors in high-dimensional lattices [Ajt96, MG02]. It is based on a polynomial

ring over a finite field, and some of its parameters are random ternary polynomials¹, i.e., with coefficients in $\{-1, 0, +1\}$. A subset of these are restricted to being *fixed-weight*, i.e. with a prescribed number of non-zero coefficients. Indeed, NTRU’s specification [CDH⁺20] requires an even stronger property, defining $\mathcal{T}(d)$ as the set of ternary “polynomials that have exactly $d/2$ coefficients equal to $+1$ and $d/2$ coefficients equal to -1 ”.

The straightforward approach to sample from $\mathcal{T}(d)$ is to fix a representative of $\mathcal{T}(d)$ (e.g., -1 for the first $d/2$ coefficients, $+1$ for the next $d/2$ coefficients, and 0 for the remaining ones), and randomly permute its coefficients using a shuffling algorithm. However, known shuffling algorithms are not timing-attack resistant [Dan19]. The usual alternative, mandated by the NTRU specification [CDH⁺20], is based on constant-time sorting. Briefly, an array of key-value pairs is created, using uniformly random samples as keys, while values are coefficients from the chosen fixed representative of $\mathcal{T}(d)$. Sorting the random keys induces a random permutation of the coefficients. This approach is illustrated in Algorithm 1. While most classical sorting algorithms are variable-time, sorting networks [Bat68], [Knu98, §5.3.4] are constant-time and shown to be efficient in practice [BCS13, BCLv17].

Algorithm 1 SAMPLEFIXEDTYPE: fixed-weight sampling by sorting [CDH⁺20, §1.10.5]

Input: $(b_0, b_1, \dots, b_{l-1})$ (random bit string of length $l = 30(n-1)$)

Output: v (an $(n-1)$ -coefficient polynomial in $\mathcal{T}(q/16-1)$)

Notes: We denote by $\text{INT}(x_0, \dots, x_{k-1})$ the unsigned integer with x_j ($0 \leq j \leq k-1$) at the j -th bit of its binary representation.

```

1: a  $\leftarrow [0, 0, \dots, 0]$  ▷ Array of  $n-1$  zeros
2:  $v \leftarrow 0$  ▷ The zero polynomial
3:  $i \leftarrow 0$ 
4: for  $i = 0$  to  $q/16 - 2$  do
5:    $\mathbf{a}[i] \leftarrow 1 + 4 \cdot \text{INT}(b_{30i}, \dots, b_{30i+29})$ 
6: for  $i = q/16 - 1$  to  $q/8 - 3$  do
7:    $\mathbf{a}[i] \leftarrow 2 + 4 \cdot \text{INT}(b_{30i}, \dots, b_{30i+29})$ 
8: for  $i = q/8 - 2$  to  $n - 2$  do
9:    $\mathbf{a}[i] \leftarrow 0 + 4 \cdot \text{INT}(b_{30i}, \dots, b_{30i+29})$ 
10: Sort a in constant time
11: for  $i = 0$  to  $n - 2$  do
12:    $v \leftarrow v + (\mathbf{a}[i] \bmod 4)x^i$ 
13: return  $v$ 

```

2.2 Shuffling algorithms

Fisher–Yates. The Fisher–Yates shuffle algorithm, also known as Knuth’s shuffle [FY38, Dur64, Knu97], is a classical technique for randomly and unbiasedly permuting elements in a collection. It is displayed in Algorithm 2.

Algorithm 2 FISHER–YATES(**a**, n)

Input: An array **a** of n elements

Output: A random permutation of **a**

```

1: for  $i = n - 1$  downto  $1$  do
2:    $j \xleftarrow{\$} \{0, 1, \dots, i\}$ 
3:   Exchange  $\mathbf{a}[j]$  and  $\mathbf{a}[i]$ 
4: return a

```

¹Given that these are ternary polynomials, the coefficient 2 may be used interchangeably with -1 .

Fisher–Yates has favorable performance characteristics: $O(n)$ running time with small constants. However, array accesses indexed by secret data are susceptible to timing attacks, due to variabilities induced by the presence or absence of data in CPU caches. This was observed in [Koc96], and demonstrated in practice in [TTMM02, TSS⁺03, Ber04] and many other works. This principle can be applied to Algorithm 2 to recover the indices j in the accesses to $\mathbf{a}[j]$ in line 3, allowing an attacker to reconstruct the permutation.

Rao–Sandelius. A relevant shuffling algorithm is Rao–Sandelius, independently proposed in the 1960s by [Rao61] and [San62]. It relies on a divide-and-conquer strategy.

Algorithm 3 RS(\mathbf{a}, n): Rao–Sandelius shuffle

Input: An array \mathbf{a} of n elements

Output: A random permutation of \mathbf{a}

```

1: if  $n \leq 1$  then
2:   return  $\mathbf{a}$ 
3: if  $n = 2$  then
4:   if  $\text{rand-bit} = 1$  then
5:     return  $[\mathbf{a}[1], \mathbf{a}[0]]$ 
6:   else
7:     return  $[\mathbf{a}[0], \mathbf{a}[1]]$ 
8: Let  $\mathbf{A}_0$  and  $\mathbf{A}_1$  be two empty arrays
9: for  $i = 0$  to  $n$  do
10:  Add  $\mathbf{a}[i]$  into  $\mathbf{A}_{\text{rand-bit}}$ 
11: return RS( $\mathbf{A}_0, |\mathbf{A}_0|$ ) || RS( $\mathbf{A}_1, |\mathbf{A}_1|$ )

```

The case $n = 2$ can be made constant-time using standard techniques. Line 10 directs each element $\mathbf{a}[i]$ to a different array depending on a random bit; by evicting both arrays from the cache for later probing, an attacker can find which array was written to. This can be countered by writing to both arrays regardless of the random bit drawn, but only incrementing the correct pointer. However, the random choice of array for assignment may lead to uneven growth of the arrays. We are unaware of any concrete analyses in the literature, but conjecture that this leaks enough data to mount a cache timing attack.

MergeShuffle. Finally, MERGESHUFFLE, introduced in [BBHL18], “is an (easy to implement) extremely efficient algorithm to generate random permutations (or to randomly permute an existing array)”. As with the Rao–Sandelius algorithm, MERGESHUFFLE uses a divide-and-conquer strategy and is amenable to a parallel implementation.

Let k be a cut-off threshold to switch to Fisher–Yates. MERGESHUFFLE splits an input array $(a_0, a_1, \dots, a_{n-1})$ into 2^k blocks to be shuffled using Fisher–Yates (Algorithm 2), and then merges the resulting permutations as presented in Algorithm 4. The merging procedure is similar in spirit to that of e.g. mergesort, but it is performed in-place and uses a random bit to choose whether to swap elements from the two input arrays.

The use of Fisher–Yates as a subroutine of MERGESHUFFLE renders it equally susceptible to cache timing attacks. It is also unclear whether the merging step can be vectorized, to attain competitive performance, and implemented in constant-time.

3 Fixed-weight sampling by constant-time shuffling

As just discussed, while shuffling is the natural solution to the fixed-weight sampling problem in NTRU, we are unaware of any shuffling algorithm resistant to side-channel attacks. In this section, we propose a linear-time, unbiased and timing attack-resistant

Algorithm 4 MERGESHUFFLE(\mathbf{a}, k)**Input:** An array \mathbf{a} of n elements**Output:** A random permutation of \mathbf{a}

- 1: Divide \mathbf{a} into 2^k blocks of roughly the same size
- 2: Shuffle each block independently using Fisher–Yates
- 3: $p \leftarrow k$
- 4: **repeat**
- 5: Merge adjacent blocks of size 2^p into new blocks of size 2^{p+1} ▷ See text
- 6: $p \leftarrow p + 1$
- 7: **until** \mathbf{a} consists of a single block
- 8: **return** \mathbf{a}

shuffling algorithm suitable for NTRU fixed-weight sampling. Throughout this section, n is defined as in the NTRU specification and assumes values of either 509, 677 or 821.

We first describe a subroutine (Algorithm 5) to generate an array of random integers \mathbf{si} such that $\mathbf{si}[i] \sim \mathcal{U}(0, n-1-i)$ for $0 \leq i < n-1$. It is a slightly modified version of [Lem19, Algorithm 5]. While other approaches exist to achieve the same result, some of which are discussed in the same paper, this method achieves the best performance among all methods we experimented with, while restricting costly (and, in all CPUs we are familiar with, variable-time) divisions by non-power-of-two integers to a pre-computation step.

Algorithm 5 REJSAMPLINGMOD(n): Generation of an array of unbiased uniformly random integers modulo $n-1, n-2, \dots, 2$ **Input:** n **Output:** \mathbf{si} (output array of $n-1$ integers, s.t. $\mathbf{si}[i] \sim \mathcal{U}(0, n-1-i)$)

- 1: **for** $i = 0$ **to** $n-2$ **do** ▷ Precomputation
- 2: $\mathbf{t}[i] = 2^L \bmod (n-1-i)$
- 3: **for** $i = 0$ **to** $n-2$ **do**
- 4: **repeat**
- 5: $x \xleftarrow{\$} \{0, 1, \dots, 2^L - 1\}$ ▷ Sample a uniformly random integer of L bits
- 6: $s \leftarrow n-1-i$
- 7: $m \leftarrow x \cdot s$
- 8: $l \leftarrow m \bmod 2^L$ ▷ Reduce using a bitmask
- 9: **until** $l \geq \mathbf{t}[i]$
- 10: $\mathbf{si}[i] \leftarrow \lfloor m/2^L \rfloor$ ▷ Divide using right-shift by L bits
- 11: **return** \mathbf{si}

Lemma 1 (Correctness and unbiasedness of Algorithm 5 [Lem19]). *Let $L \in \mathbb{N}^* = \mathbb{N} \setminus \{0\}$. Then, $\forall s \in [0, 2^L)$ and $\forall y \in [0, s)$, with s, y integers, there are $\lfloor 2^L/s \rfloor$ values of $x \in [0, 2^L)$ such that $\lfloor (x \cdot s)/2^L \rfloor = y$ and $l = (x \cdot s) \bmod 2^L \geq 2^L \bmod s$.*

Proof. Take $s \in [0, 2^L)$ and $x \in [0, 2^L)$. Write $y \cdot 2^L + l = x \cdot s$. Thus, $x \cdot s$ lies in the interval $[y \cdot 2^L, (y+1) \cdot 2^L)$ of length 2^L . As we also require that $l = (x \cdot s) \bmod 2^L \geq 2^L \bmod s$, the interval is narrowed down to $[y \cdot 2^L + (2^L \bmod s), (y+1) \cdot 2^L)$, of length $2^L - (2^L \bmod s)$, which is a multiple of s . There are $(2^L - (2^L \bmod s))/s = \lfloor 2^L/s \rfloor$ representatives of each residue class modulo s in this interval, and in particular multiples of s of the form $x \cdot s$. \square

Lemma 1 implies that rejecting values such that $x \cdot s \bmod 2^L < 2^L \bmod s$ ensures that $\lfloor x \cdot s/2^L \rfloor$ is unbiasedly sampled from $[0, s)$. We discuss issues of timing attack resistance, as well as the choice of the performance-critical parameter L in §4. Algorithm 6 is our proposed shuffling approach for ternary fixed-weight sampling.

Algorithm 6 SHUFFLE(n, c_0, c_1): Ternary fixed-weight sampling by shuffling

Input: n **Input:** c_0, c_1 (prescribed number of coefficients equal to 0, resp. 1)**Output:** \mathbf{v} (output array of $n - 1$ integers)

```

1:  $\mathbf{si} \leftarrow \text{REJSAMPLINGMOD}(n)$ 
2: for  $i = 0$  to  $n - 2$  do
3:   if  $\mathbf{si}[i] < c_0$  then           ▷ See text for discussion of constant-time implementation
4:      $\mathbf{v}[i] \leftarrow 0$ 
5:      $c_0 \leftarrow c_0 - 1$ 
6:   else if  $\mathbf{si}[i] < c_0 + c_1$  then
7:      $\mathbf{v}[i] \leftarrow 1$ 
8:      $c_1 \leftarrow c_1 - 1$ 
9:   else
10:     $\mathbf{v}[i] \leftarrow -1$ 
11: return  $\mathbf{v}$ 

```

Firstly, we note that, while NTRU prescribes an identical count of 1 and -1 elements in the sampled vector, Algorithm 6 imposes no such restriction: the user is free to choose the parameters c_0 and c_1 (and implicitly, c_{-1} from $n - 1 = c_0 + c_1 + c_{-1}$).

Evidently, the main loop of Algorithm 6, as presented, does not execute in constant time due to the use of branches. However, architecture-agnostic standard techniques, as well as architecture-specific conditional instructions, can be used to obtain a branchless, constant-time implementation; see §4. Moreover, all accesses to the arrays \mathbf{si} and \mathbf{v} are performed sequentially. We exploit the fact that $O(1)$ distinct values need to be shuffled (indeed, only 3: $-1, 0, 1$), a situation not considered in the usual shuffling algorithms. Intuitively, one could draw an analogy between Fisher–Yates shuffling and selection sort, and by replacing the latter with counting sort, arrive at our proposed algorithm.

Lemma 2 (Correctness and unbiasedness of Algorithm 6). *Let n be an integer and c_i for $i \in \{-1, 0, 1\}$ be the number of coefficients equal to i in the output polynomial, so that $c_{-1} + c_0 + c_1 = n - 1$, and write $\Sigma = (\Sigma_1, \dots, \Sigma_{n-1}) = \{-1\}^{c_{-1}} \{0\}^{c_0} \{1\}^{c_1}$, i.e. a concatenation of c_{-1} , c_0 and c_1 repetitions of -1 , 0 and 1 , respectively. Then, Algorithm 6 produces an array $\mathbf{v} = \sigma(\Sigma)$, where $\sigma \in \text{Perm}(\Sigma; c_{-1}, c_0, c_1)$ is a uniformly drawn permutation of Σ .*

Proof. It is easily seen that Algorithm 6 outputs $\mathbf{v} = (\sigma(\Sigma_1), \dots, \sigma(\Sigma_{n-1}))$ for some $\sigma \in \text{Perm}(\Sigma; c_{-1}, c_0, c_1)$. To show unbiasedness, we proceed by directly computing $P(v = (\sigma(\Sigma_1), \dots, \sigma(\Sigma_{n-1})))$ for an arbitrary σ . Define $c_{i,k}$ for $i \in \{-1, 0, 1\}$ as the count of elements equal to i yet to be output by the algorithm at the start of the k -th iteration; thus, $c_{i,1} = c_i$. Since $\mathbf{si}[k - 1] \sim \mathcal{U}(0, n - k)$ for $1 \leq k \leq n - 1$, the probability that $v_k = i$, given all previously drawn values, is $c_{i,k}/(n - k)$. We have:

$$\begin{aligned}
P(v = (\sigma(\Sigma_1), \dots, \sigma(\Sigma_{n-1}))) &= P(v_1 = \sigma(\Sigma_1), \dots, v_{n-1} = \sigma(\Sigma_{n-1})) \\
&= P(v_1 = \sigma(\Sigma_1)) \prod_{k=2}^{n-1} P(v_k = \sigma(\Sigma_k) | v_1 = \sigma(\Sigma_1), \dots, v_{k-1} = \sigma(\Sigma_{k-1})) \\
&= \frac{c_{\sigma(\Sigma_1),1}}{n-1} \cdot \frac{c_{\sigma(\Sigma_2),2}}{n-2} \cdot \frac{c_{\sigma(\Sigma_3),3}}{n-3} \dots \frac{c_{\sigma(\Sigma_{n-1}),n-1}}{1}.
\end{aligned}$$

The numerators in the last equation are in bijection with the values $c_{-1}, (c_{-1} - 1), \dots, 1, c_0, (c_0 - 1), \dots, 1, c_1, (c_1 - 1), \dots, 1$, permuted in a specific order. It follows from the commutativity of multiplication that:

$$P(v = (\sigma(\Sigma_1), \dots, \sigma(\Sigma_{n-1}))) = \frac{c_{-1}! c_0! c_1!}{(n-1)!} = \frac{1}{|\text{Perm}(\Sigma; c_{-1}, c_0, c_1)|}.$$

Then Algorithm 6 covers all possible permutations $\sigma \in \text{Perm}(\Sigma; c_{-1}, c_0, c_1)$ with uniform probability. \square

Lemma 3. *Algorithm 6 executes in time $O(n)$ on average, assuming $2^L > 2n$.*

Proof. The loop of line 2 clearly executes in time $O(n)$. Thus, the remaining work consists in analyzing Algorithm 5. The outer loop (line 3) consists of $n - 1$ iterations. Noticing that $t[i] < n, \forall i$, the condition in line 9 will be satisfied with probability $1 - \frac{n}{2^L} > 1/2$. Thus, the expected number of iterations in the inner loop (line 4) is less than 2, so Algorithm 5 also executes in time $O(n)$ on average. \square

We remark that the $O(n)$ running time of Algorithm 6 improves upon the $O(n \log^2 n)$ running time of sorting networks typically used for constant-time sorting implementations [BCS13, BCLv17], such as Batchers odd-even merge sort [Bat68, Knu98].

The algorithm necessarily consumes at least $n - 1$ random L -bit integers and may, in principle, consume an infinite number of them due to rejections; however, in §4, we show that, for $L = 16$, generating just 4% to 5.5% extra random integers is sufficient in practice.

Application to other PQC cryptosystems. While we have directed our efforts to a single scheme (NTRU) to showcase the potential of our approach through highly optimized code, many other PQC schemes also rely on fixed-weight sampling. We briefly mention some relevant schemes, commenting on the applicability of our techniques to them; note that this is not an exhaustive list of potential applications.

For lattice-based schemes, we are only aware of applications of fixed-weight sampling to NTRU variants, such as NTRU LPrime and Streamlined NTRU Prime [BCLv17, BBC⁺20], also submitted to the NIST PQC contest. As in NTRU, these variants sample vectors in $\{-1, 0, 1\}$ with a fixed weight, i.e. count of non-zero elements; however, we recall that, in NTRU, the count of elements equal to 1 and -1 in a sampled vector is identical, a constraint which is lifted in these variants². For concreteness, we consider the parameter sets `sntrup761` and `ntrulpr761`; in the notation of Algorithm 6, we have $n = 761$ and $c_0 = 761 - 286 = 475$ for `sntrup761` and $c_0 = 761 - 250 = 511$ for `ntrulpr761`. Whereas NTRU fixes $c_1 = (n - c_0)/2$, in these variants c_1 is freely chosen, although there are attacks for the case that c_1 that is too small [BCLv19, §6]. Consider the process of uniform sampling of ternary vectors of size 761 with c_0 elements equal to zero. Then, c_1 follows a binomial distribution with $n - c_0$ trials and identical success/failure probabilities of 0.5 (mapping wlog “success” and “failure” to sampling 1 and -1 , respectively). We can efficiently sample from this distribution by a simple procedure: sample an array of $n - c_0$ random bits and set c_1 to the population count of this random bit array³. We believe that this extra step does not add too much cost, and that these NTRU variants should still benefit considerably from our approach; we encourage others to investigate this possibility.

Another class of PQC schemes that require fixed-weight sampling are code-based cryptosystems, such as Classic McEliece [ABC⁺22], BIKE [ABB⁺22] and HQC [AAB⁺22], all of which are contenders in the fourth round of NIST’s PQC contest. They require binary, rather than ternary, fixed-weight vectors; Algorithm 6 is easily modified to handle this case. Nevertheless, our approach is unlikely to benefit these schemes, due to their use of low weight (sparse) vectors: typically $w = O(\sqrt{n})$. These schemes perform fixed-weight sampling using distinct approaches: [Sen21, Algorithm 5] for BIKE and HQC, and rejection sampling of indices of non-zero elements (rejecting in the case of index collisions) for McEliece. In both cases, the running time, memory usage and randomness consumption

²We note that Algorithm 6, as previously discussed, is not restricted to identical counts of 1 and -1 elements either.

³Intel and ARMv8-A architectures provide population count instructions; for architectures lacking such instructions, such as the Cortex-M4, there are efficient algorithms for its computation using universally available arithmetic, logic and shift instructions [War12].

are a function of the weight w rather than the vector length n , whereas these parameters are a function of n in Algorithm 6. D. J. Bernstein, in private communication with the authors, indicates that he has devised a new algorithm with favorable performance characteristics for the sparse setting⁴.

WAVE [BCC⁺23] is a code-based signature scheme submitted to NIST’s additional digital signatures contest, which calls for random permutations of fixed-weight ternary vectors. Despite similarities to the fixed-weight sampling problem treated in this paper, an important distinction is a requirement for computing the explicit permutation mapping the input vector to the randomly permuted output vector. We have superficially investigated a modified version of Algorithm 6 to fit this scenario, by keeping track of the indices of each element during the shuffling process, but this version departs significantly from the efficiency and constant-time characteristics of Algorithm 6. Lacking further progress, it appears that sorting remains better suited to this scenario for now.

4 Implementation aspects

Architectural guarantees regarding constant-time execution. Both ARMv8-A and Intel architectures have recently introduced hardware flags that, when set, guarantee constant-time execution of a subset of CPU instructions, which should generally be sufficient to implement most cryptographic algorithms: FEAT_DIT for ARMv8-A [ARM23, §A2.6.1, B1.3.6, C5.2.4] and DOIT for Intel [Int23a, Int23b]. We verified that all instructions handling secret data in our ARMv8-A implementations are included in the affected subset.

These new features do not imply that CPUs launched prior to the introduction of these flags execute these instructions in variable time. Indeed, ARM claims to be unaware of older CPUs with variable timing for instructions now covered by FEAT_DIT [ARM]⁵; and Intel advises developers to assume older microarchitectures behave as if DOIT is enabled [Int23a].

This issue has garnered attention at the beginning of 2024, as Apple ARMv8-A cores (which are designed by Apple and not ARM) are subject to a microarchitectural attack called GoFetch [CWS⁺24]; setting the FEAT_DIT bit on the M3 disables the data memory-dependent prefetchers targeted by the attack, rendering it ineffective, while for the M1 there appears to be no countermeasure at the hardware level.

Resistance against timing attacks of Algorithm 5. There are some possible sources of timing leaks in Algorithm 5, which we enumerate and analyze.

The integer multiplication in line 7 must execute in constant-time, which is the norm in modern CPUs⁶, although there are rare exceptions such as the ARM Cortex-M3 for $32 \times 32 = 64$ -bit multiplications; however, $32 \times 32 = 32$ -bit multiplication suffices for the purpose of this algorithm, and there is evidence that it executes in constant time in the Cortex-M3 [dG15, Por18].

Array accesses in line 9 use sequential indices; thus, secret data is not leaked. The loop in lines 4 to 9 performs rejection sampling based on public data, precomputed in line 2: the remainder of 2^L divided by integers in the sequence $n - 1, n - 2, \dots, 1$, where L and n are public parameters. Nevertheless, given the attack of Guo et al. [GHJ⁺22] targeting rejection sampling in fixed-weight sampling algorithms for BIKE and HQC code-based cryptosystems, it is worth analyzing whether a similar attack could apply here. We note that their attack relies on two key assumptions:

⁴Prototype source code available at <https://cr.yp.to/2024/gcww-20240512.py> and <https://cr.yp.to/2024/insertionseries-20240515.py>.

⁵However, [GS23] measured an increase in execution time for integer multiplication when both operands are $> 2^{32}$ in the ARM Cortex-A53, as found in the Raspberry Pi 3B single-board computer.

⁶Note that multiplication instructions are covered by ARMv8-A’s FEAT_DIT and Intel’s DOIT flags.

1. A high rejection rate, leading to multiple calls to the `seedexpander` routine (equivalently in our case, the `randombytes` routine) which creates a timing distinguisher. As discussed next, the rejection rate for our chosen parameter $L = 16$ is sufficiently small that e.g. a full run of Algorithm 5 in the case $n = 509$ has $> 40\%$ probability of no rejections at all. Due to this low rejection rate, and using a batch random number generation idea discussed later, we sample enough uniform random integers from the outset so that the probability that extra samples are required is deemed negligible ($< 2^{-74}$, which can be further reduced at little extra cost). This allows calling `randombytes` only once, generating a fixed number of random integers, while introducing a negligible overhead.
2. Derivation of the random seed for fixed-weight sampling from secret data – namely, the output of decryption from the reencryption step of decapsulation, as required by the Fujisaki-Okamoto transform for IND-CCA security of the KEM. The attack starts by trial encrypting many candidate messages until finding an m that requires multiple calls to `seedexpander`, which gives rise to a timing distinguisher (a possibility ruled out by the first point above). Carefully constructed perturbations of the resulting ciphertext c are fed to the decapsulation procedure, while using the timing distinguisher to determine whether the decryption step of reencryption outputs the same m or a different message, allowing the attacker to learn information about the secret key. Repeated application of this procedure extracts the vast majority of key material, and the remaining bits are easily found. However, we note that NTRU does not require reencryption due to the rigidity of the NTRU DPKE [CDH⁺20, Figures 9 and 10]; indeed, the fixed-weight sampling algorithm is not executed at all during either decapsulation or decryption.

Thus, we conclude that Algorithm 5 does not render NTRU vulnerable to the attack of Guo et al [GHJ⁺22].

Choosing the parameter L . The choice of L in Algorithm 5 is a tradeoff between the cost of random number generation and the frequency of rejections; the latter lead to branch mispredictions and costly pipeline flushes in modern, highly-pipelined superscalar CPUs such as some of the ARMv8-A cores considered in this work. If samples are rarely rejected, a SIMD implementation of the algorithm becomes feasible; one can keep track of which lanes were rejected and resample them later (usually with scalar code). To minimize rejection, one must choose L such that $2^L \gg n - i$, but this translates into added cost for random number generation, and thus L should not be unreasonably larger than $n - i$.

We propose $L = 16$ as a natural choice, supported by all scalar and SIMD instruction sets we are aware of. The next smaller size, 8 bits, is insufficient for half or more of the values to be sampled in the standard NTRU parameter sets, and for most of the intervals where it is sufficient, it would lead to a high rejection rate, running counter to the SIMD philosophy. By exactly matching an available lane size, no bit shifts/masks/permutations are required to load random integers into SIMD registers, further improving performance. It is also the natural choice for storing the 11- or 12-bit NTRU polynomial coefficients; indeed, it is the representation used by the reference code and the state-of-the-art implementations we chose for performance comparisons, requiring no size conversions.

Finally and most importantly, rejections are relatively rare: a block of 16 samples is fully accepted (zero rejections) with probability at least 94.2%, 91.6% and 90.1% for $n = 509, 677$ and 821 , respectively. These are minimum figures, and as $n - i$ decreases, the acceptance probability increases even further. Furthermore, the probability of accepting all $n - 1$ samples (i.e., no rejections at all during a complete execution of the algorithm) is 40.2%, 18.9% and 8.6% for $n = 509, 677$ and 821 respectively. These figures are obtained by modeling the number of required samples as a sum of geometric random variables and are displayed in a Jupyter notebook accompanying the source code of our implementation.

Due to the low rejection probability, it is sufficient to generate just a few extra random integers over the lower bound of n . For each n , we computed the cumulative distribution function $P(x \leq k)$ and sought the minimum k such that $1 - P(x \leq k) < 2^{-74}$, enough to sample $2^{10} > n$ integers for each of 2^{64} key exchanges. For $L = 16$, and rounding up to the next multiple of 8 (the number of 16-bit lanes in a NEON register), we find that 536, 704 and 856 random 16-bit integers are sufficient (i.e., an overhead of 5.5%, 4.1% and 4.4%) for $n = 509, 677$ and 821 , respectively. This calculation is included in the aforementioned Jupyter notebook, which can be adapted to other choices of L if desired.

One might argue that $L = 16$ is a “wasteful” choice, as it requires 123%, 109% and 103% more bits than the (unattainable) lower bound of $\log_2(n!)$ bits for $n = 509, 677$ and 821 , respectively. Still, we note this is slightly more than half as many random bits as the approach dictated by the NTRU specification [CDH⁺20], which calls for $30 \times n$ bits.

Taking $L > 16$ appears counterproductive, e.g. due to reduced computational throughput from using larger SIMD lanes. On the other hand, in scenarios where pseudo-random number generation is expensive, SIMD is not available and pipeline flushes have less performance impact (i.e. deeply embedded cores such as the Cortex-M4), choosing $L < 16$ (say, 12 or 10) may result in better overall performance. One might even conceive of an adaptive choice, decreasing L along with $n - i$, although this results in more complex code.

D. J. Bernstein, in private communication with the authors, pointed us to an unpublished paper of his [Ber18] which quantifies the security loss due to foregoing rejection sampling in Algorithm 5, with obvious performance benefits due to the avoidance of branches. For the choice $L = 16$, it leads to 1.52, 2.61 and 3.84 bits of security loss, respectively, for $n = 509, 677$ and 821 . Alternatively, a larger value of L can be chosen to reduce the security loss, but then the performance gains due to branchless code must be weighted against extra randomness consumption and reduced arithmetic throughput. We have experimented with $L = 20$, which reduce security by 0.09, 0.16 and 0.24 bits, respectively, for $n = 509, 677$ and 821 . Benchmarks show a speedup of 9.7% to 11.5% in the performance of fixed-weight sampling on the Cortex-M4; however, due to the increased randomness consumption, encapsulation performance actually decreased by 0.6% to 0.7%. Preliminary experiments on a NEON version with $L = 22$ resulted in a small slowdown, but we cannot rule out a crossover to a small improvement through careful optimization.

Batch random number generation in Algorithm 5. For ease of understanding, we displayed uniform random numbers as being sampled online in line 5 of Algorithm 5. However, this leads to high overheads in actual implementations. Instead, we sample the required count of uniform random numbers (see previous discussion about the choice of L) into an array ahead of time, and consume the randomness directly from this array.

SIMD implementation of Algorithm 5. To minimize the execution time of Algorithm 6, we seek to implement Algorithm 5 using SIMD instructions. At first glance, it is unsuitable for SIMD, as some lanes may be rejected while others are accepted during sampling. However, it is possible to sample a whole SIMD register and take note of which lanes, if any, were rejected, to be fixed up later using scalar code (recall that an adequate choice of L ensures that rejections occur with low probability, so the performance impact of this fixup procedure is limited.) However, when the batch random number generation idea is used, rejections induce an irregular access pattern to the array of random numbers. We avoid this issue by using disjoint ranges of the array for SIMD sampling (indices 0 to $n - 2$) and the fixup procedure ($n - 1$ onwards). These ideas are captured in Algorithm 7.

In addition to previously discussed issues of timing attack resistance of Algorithm 5, we note that any non-sequential accesses to the array `rnd` arise from switching between the ranges of indices $0 \leq i + k < n - 1$ and $j \geq n - 1$, that is, they are due to rejections and thus do not leak secret data; accesses within each range are sequential.

Algorithm 7 SIMD-REJSAMPLINGMOD(n, \mathbf{rnd}): SIMD version of Algorithm 5.

Input: n **Input:** \mathbf{rnd} (array of random L -bit integers; refer to previous discussion about its length)**Output:** \mathbf{si} (output array of $(W+1)\lfloor(n-1)/W\rfloor$ integer elements, of which only the first $n-1$ entries are valid.)

```

1: for  $i = 0$  to  $n - 2$  do ▷ Precomputation
2:    $\mathbf{t}[i] = 2^L \bmod (n - 1 - i)$ 
3:  $j \leftarrow n - 1$ 
4: for  $i = 0, W, 2W, \dots, W\lfloor(n-1)/W\rfloor$  do
5:    $\mathit{mask} \leftarrow 0$ 
6:   for  $k = 0$  to  $W - 1$  do ▷ Loop body implemented using SIMD code
7:      $\mathbf{m}[k] \leftarrow \mathbf{rnd}[i + k] \times (n - 1 - (i + k))$ 
8:      $\mathbf{l}[k] \leftarrow \mathbf{m}[k] \bmod 2^L$ 
9:      $\mathbf{si}[i + k] \leftarrow \lfloor \mathbf{m}[k] / 2^L \rfloor$ 
10:    if  $\mathbf{l}[k] < \mathbf{t}[i + k]$  then
11:       $\mathit{mask}_k \leftarrow 1$  ▷  $\mathit{mask}_k$  denotes the  $k$ -th bit of  $\mathit{mask}$ 
12:    else
13:       $\mathit{mask}_k \leftarrow 0$ 
14:    while  $\mathit{mask} \neq 0$  do ▷ Loop body implemented using scalar code
15:       $k = \text{COUNTTRAILINGZEROS}(\mathit{mask})$ 
16:      repeat
17:         $\mathbf{m}' \leftarrow \mathbf{rnd}[j] \times (n - 1 - (i + k))$ 
18:         $j \leftarrow j + 1$ 
19:         $\mathbf{l}' \leftarrow \mathbf{m}' \bmod 2^L$ 
20:        until  $\mathbf{l}' \geq \mathbf{t}[i + k]$ 
21:         $\mathbf{si}[i + k] \leftarrow \lfloor \mathbf{m}' / 2^L \rfloor$ 
22:         $\mathit{mask}_k \leftarrow 0$ 
23: return  $\mathbf{si}$ 

```

Line 10 should use SIMD comparison instructions (e.g. NEON’s CMHI or AVX2’s VPCMPGT). These create a mask with all bits set or clear in the corresponding lane, while Algorithm 7 as written calls for setting and clearing individual bits, a choice made purely for ease of exposition. Actual implementations are advised to tweak the representation to employ groups of bits instead, so as to achieve an efficient implementation of the inner loop of line 6. For instance, VPMOVMASKB is a natural choice in AVX2, resulting in 2-bit mask groups for 16-bit lanes. In NEON, we extract 8-bit masks with UZP1, and reduce them to 4-bit masks using SHRN by 4. NEON’s 128-bit registers suggest a choice of $W = 8$ if $L = 16$. However, we achieved better performance by taking $W = 16$, implemented as an unrolled 2-iteration loop processing 8-element vectors. We attribute this to the fact that converting a mask with UZP1 and SHRN costs the same for 8 or 16 values.

Constant-time implementation of Algorithm 6. We now discuss how to implement Algorithm 6 in constant-time. First, we rewrite it using the C language’s ternary operator, as shown in Algorithm 8, and then discuss strategies to implement this operator in constant time, firstly as an architecture-agnostic solution, and then consider conditional instructions present in the ARMv8-A, ARMv7-M and Intel architectures. Note that this version replaces -1 coefficients by 2; this is not an issue as the sampled polynomial has coefficients in $\mathbb{Z}/3\mathbb{Z}$, and indeed, the reference NTRU code employs the same representation.

Expressions of the form $(x < y) ? -1 : 0$, in lines 4 and 5 of Algorithm 8, can be made constant-time by noticing that, in two’s complement integer arithmetic (used in nearly all modern architectures), -1 and 0 have all bits set and cleared, respectively. The sign (most

Algorithm 8 CT-SHUFFLE($n, c_0, c_1, \mathbf{rnd}$): Fixed-weight sampling by shuffling, implemented in constant-time

Input: n

Input: c_0, c_1 (prescribed number of coefficients equal to 0, resp. 1)

Input: \mathbf{rnd} (array of random L -bit integers; refer to previous discussion about its length)

Output: \mathbf{v} (output array of $n - 1$ integers)

Notes: We employ the C language ternary operator $?$ to denote constant-time selection between two values based on a condition. See text for implementation possibilities.

```

1:  $\mathbf{si} \leftarrow \text{SIMD-REJSAMPLINGMOD}(n, \mathbf{rnd})$ 
2:  $c_{01} \leftarrow c_0 + c_1$  ▷ Note this invariant is maintained in the loop body
3: for  $i = 0$  to  $n - 2$  do
4:    $t_0 \leftarrow (\mathbf{si}[i] < c_0) ? -1 : 0$ 
5:    $t_1 \leftarrow (\mathbf{si}[i] < c_{01}) ? -1 : 0$ 
6:    $c_0 \leftarrow c_0 + t_0$ 
7:    $c_{01} \leftarrow c_{01} + t_1$ 
8:    $\mathbf{v}[i] \leftarrow 2 + t_0 + t_1$ 
9: return  $\mathbf{v}$ 

```

significant) bit of $x - y$ is 1 if $x < y$ and 0 otherwise; an arithmetic right shift by $w - 1$ bits, where w is the word size, replicates the sign bit across the entire word. Concretely, the following C code implements line 4 for 16-bit signed integer variables:

```
t0 = (si[i] - c0) >> 15;
```

While already efficient, better performance is achievable. To that end, we analyze the critical path of the main loop of Algorithm 8, shown in Figure 1. We disregard memory loads and stores, which can be removed from the critical path by proper scheduling. For any mobile-, desktop- or server-class modern CPU, one can assume at least a 2-way superscalar pipeline and single-cycle latency for all used operations, in which case the critical path of lines 4 and 5 from one iteration to the next (the bold arrows in the figure) takes 3 cycles.

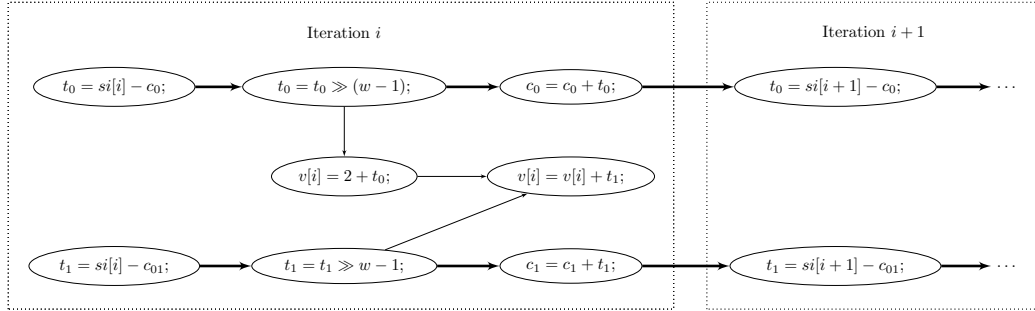


Figure 1: Critical path of the main loop of Algorithm 8.

In ARMv8-A, arithmetic instructions can be encoded so that one of the input operands is shifted; thus, a single instruction can compute both $t_0 = t_0 \gg (w - 1)$ and $c_0 = c_0 + t_0$. Unfortunately, ARMv8-A CPUs considered in this work, such as the Apple M1 [Joh22] and Cortex-A72 [ARM15], execute these instructions with a 2-cycle latency, offering no gain in performance (but a slight reduction in code size).

By employing ARMv8-A conditional instructions such as CINC and CSET, it is possible to reduce the critical path to 2 cycles. However, Algorithm 8 calls for decrementing c_0 and c_{01} , and there is no CDEC instruction in ARMv8-A; we modify the algorithm to use

negative values for c_0 , c_{01} and $si[i]$, so that we can increment c_0 and c_{01} using CINC instead. Thus, we arrive at the code of Listing 4.1 for the algorithm’s main loop.

```

cmp      c0, si[i]
cinc     c0,  c0,  lt
cset     v[i], ge
cmp      c01, si[i]
cinc     c01, c01, lt
cinc     v[i], v[i], ge

```

Listing 4.1: Main loop of Algorithm 8 in the ARMv8-A architecture.

There are two critical paths: one from `cmp c0, r` to `cinc c0, c0, lt` to the next iteration’s `cmp c0, r`; and the second for the same instructions involving c_{01} . In all the considered ARM CPUs, all instructions in the code fragment above have single-cycle latency, and thus the loop has the potential to execute in 2 cycles/iteration.

Unfortunately, we run into throughput issues: in the Apple M1, reverse engineering efforts [Joh22] indicate that, although it is capable of executing 6 scalar instructions/cycle, only 3 execution units can execute flag-setting and conditional instructions, i.e. all instructions in the above code fragment. While theoretically sufficient to run the code at maximum throughput, we have observed instruction scheduling issues while attempting to software-pipeline Algorithms 7 and 8, preventing execution at maximum throughput. The instruction sequence of Listing 4.2 requires more μops , but performs better in the M1.

```

subs     tmp,  c0, si[i]
cinc     c0,  c0,  lt
add      v[i], two, tmp, asr #31
subs     tmp,  c01, si[i]
cinc     c01, c01, lt
add      v[i], v[i], tmp, asr #31

```

Listing 4.2: Main loop of Algorithm 8 optimized for the Apple M1.

We use 32-bit registers (`w0`, `w1`, etc.) and initialize `two` with the constant 2. It is also advantageous for the Cortex-A72, since the `add` instruction with shifted argument executes in the M pipeline, whereas all other instructions execute in the I0/I1 pipelines. While other bottlenecks come into play in the Cortex-A72, notably its 3-wide instruction decoder, this alternative instruction sequence performs better than the original.

Intel has conditional instructions for conditional moves (`CMOVcc`) and sets (`SETcc`), where `cc` are condition codes, but no conditional increments or decrements. For positive values of c_0 and c_{01} , as in the original version of Algorithm 8, an alternative is to decrement c_0 and c_{01} and use `CMOV` to select between original and decremented values; decrements can execute in parallel with comparisons, thus the critical path is not lengthened.

Unfortunately, Intel instructions do not offer the three-operand form of ARMv8-A and other RISC architectures, so an extra `MOV` is required to create a copy prior to decrementing in order to avoid overwriting the original values; this doesn’t necessarily increase the critical path, due to `MOV` elimination [Fog22], but it does increase front-end pressure. Implementers are advised to keep in mind the achievable performance given the critical path, to benchmark and analyze compiler-generated code if employing a high-level language, and to consider inline assembly (or a full assembly language implementation) to emit instructions that are well-matched to the decoder restrictions.

For the ARMv7-M architecture, a straightforward implementation of Algorithm 8, implementing lines 4 and 5 using the arithmetic right shift trick, works really well; this is aided by the ability to shift one of the input operands to data processing (logical and

arithmetic) instructions. We have experimented with ARMv7-M’s conditional execution (IT) instruction, but we were unable to improve performance compared to the straightforward implementation. However, we did find an especially compact instruction sequence devoid of IT instructions to implement the main loop of Algorithm 8, shown in Listing 4.3.

```

cmp    si[i], c0
sbc    c0, #0
sbc    v[i], one
cmp    si[i], c01
sbc    c01, #0
sbc    v[i], #-1

```

Listing 4.3: Main loop of Algorithm 8 optimized for the Cortex-M4.

We set `one` to the constant 1. As the straightforward implementation is already efficient, this alternative saves one clock cycle per loop iteration, i.e. < 1000 cycles for the full algorithm. As fixed-weight sampling is performed only once during key generation and encapsulation, the speedup is just $< 0.02\%$ for the former and $\approx 0.15\%$ for the latter.

Software pipelining of Algorithms 7 and 8. Modern superscalar CPUs use distinct execution units for scalar and SIMD instructions. Most of the execution time of Algorithm 7 is spent in SIMD code, while Algorithm 8 is strictly scalar. This is amenable to software pipelining [Lam88]. In the best-case scenario, one can achieve execution time close to the maximum, rather than the sum, of the execution times of Algorithms 7 and 8.

Concretely, we inline Algorithm 7 into Algorithm 8, strip-mine the main loop of the latter, and then fuse the outer loops of both algorithms, processing W entries at a time. With this approach, we were able to achieve, in the Apple M1, execution times only $\approx 12\%$ slower than the lower bound (2 cycles/iteration) for the main loop of Algorithm 8 alone. This includes all overhead such as function calls and returns, prologue and epilogue, initialization, and of course, the execution of Algorithm 7 itself, as seen in Table 4. The narrow (3-wide) decoder of the Cortex-A72 precludes achieving a similar result as the M1, but by interleaving instructions of both algorithms to improve scheduling, we achieved results not far from the limit dictated by the decoder bandwidth bottleneck.

Known Answer Tests. We note that the Known Answer Tests (KATs) in NTRU’s specification [CDH⁺20] are tightly coupled to the fixed-weight sampling by sorting approach mandated there. Therefore, an implementation employing Algorithm 6 will fail these KATs for key generation and encapsulation. However, our sampled polynomials meet the fixed-weight requirement imposed by NTRU and are in principle indistinguishable from those generated by the existing approach. Thus, keys generated using our algorithm are valid, and the result of an encapsulation employing our algorithm will produce a correct decapsulation even by an unmodified implementation of the current NTRU proposal.

Given the simplicity and improved performance and code size characteristics of Algorithm 6, we suggest that future standardization attempts of NTRU specify our approach instead of sampling by sorting, and generate KATs accordingly. Implementers attempting to replicate our results, whether on ARMv8-A or other architectures, can use unofficial KATs generated by us, included in our source code package.

5 Experimental results

We now present experimental results for implementations of our proposed approach for various 64-bit ARMv8-A cores, as well as the 32-bit ARMv7-M Cortex-M4 core.

5.1 Methodology

We implemented reference versions of Algorithms 5 and 6, and optimized versions for ARMv7-M and ARMv8-A by replacing Algorithm 6 with Algorithm 8; for ARMv8-A specifically, we replaced Algorithm 5 by a NEON version of Algorithm 7. We integrated the reference and optimized implementations with existing state-of-the-art implementations of NTRU: `pqm4` [KRSS19]⁷ for ARMv7-M and [GFBL24, NG21, CCHY24] for ARMv8-A. KATs were generated using the reference implementation and compared against the optimized implementations; we added tests to ensure interoperability between a conventional implementation (using sampling by sorting) and our proposed approach.

Testbeds and measurement methods. Our testbeds for performance measurement, with their corresponding CPU cores, are:

- Apple M1 P-core at 3200 MHz in an Apple MacBook Air laptop running macOS;
- Apple M3 P-core at 4064 MHz in an Apple MacBook Pro laptop running macOS;
- Cortex-A72 at 1500 MHz in a Raspberry Pi 4 single-board computer running Linux;
- Cortex-A57 at 1430 MHz in an Nvidia Jetson Nano single-board computer running Linux;
- Cortex-A53 at 1400 MHz in a Raspberry Pi 3 single-board computer running Linux;
- Cortex-M4 at 24 MHz in an STM32F4DISCOVERY development board.

Save for the ARMv7-M Cortex-M4 core, the remaining testbeds are ARMv8-A, running in 64-bit mode. Of these, the Apple M1, M3 and Cortex-A57 cores feature ARMv8-A Cryptographic Extensions, but the Cortex-A72 and the Cortex-A53 do not.

Our ARMv8-A performance measurements use the cycle counting routines originally introduced in [NG21]. Each routine is executed for 1,024 times and the average cycle count is reported. ARMv7-M measurements employ the `pqm4` [KRSS19] benchmarking harness, which counts cycles using the Cortex-M4 SysTick timer. The number of iterations is set to 10, and the mean of results are reported; although this is a small number, the Cortex-M4 core is much simpler and more deterministic than the large out-of-order ARMv8-A cores, thus exhibiting little run-to-run variability.

While, to a first approximation, cycle counts are not influenced by CPU clock speed, there may be second-order effects such as the decoupling of CPU and bus/RAM/cache clocks. Thus, we take precautions to maximize the likelihood that benchmarks are performed at the nominal clock speeds quoted above. For Linux systems, we use the `performance` scaling governor. In Apple systems, as far as we aware, there is no control over clock speeds, and there is no TurboBoost-like feature. In both cases, we try to avoid thermal throttling by inserting delays between benchmark runs to allow systems to cool down. The Cortex-M4 core does not automatically boost/throttle clock speeds; `pqm4` configures it to 24 MHz at startup, ensuring all benchmarks run at that fixed clock speed.

ARMv8-A binaries were compiled with Apple clang 15.0 (Apple M1 and M3), clang 17.0 (Cortex-A72 and Cortex-A53), and clang 10.0 (Cortex-A57), with `-O3` and core-specific `-mcpu` optimization flags. ARMv7-M binaries were compiled with gcc 13.2.1, passing the `-O speed` flag to the `pqm4` benchmark script. We enable the FEAT_DIT bit on ARMv8-A cores where it is available (in the case of our testbeds, only the Apple M1 and M3).

⁷Although NTRU was removed from the most recent version of `pqm4`, after Kyber was selected in the NIST post-quantum standardization process, we used the most recent version prior to NTRU's removal.

ARMv8-A implementation. Our implementation is based on the source code provided by [GFBL24], which contains their AMX implementation and the NEON implementations of [CCHY24, NG21]. As [CCHY24] is the state-of-the-art in NEON implementations, but targets only the HPS2048677 and HRSS701 parameter sets, [NG21] is included to display HPS2048509 and HPS4096821 results. Importantly, [GFBL24] backports optimized auxiliary routines of [CCHY24] to [NG21] (in particular a NEON implementation of constant-time sorting) and provides an optimized implementation of NIST’s `randombytes()` AES-CTR-DRBG pseudo-random number generator (PRNG), using ARMv8-A Cryptographic Extensions. These routines are critical to the performance of fixed-weight sampling.

For CPUs that do not feature ARMv8-A Cryptographic Extensions, the ChaCha20 PRNG of [CCHY24] is used. As KATs are incompatible across different PRNGs, we supply two KAT sets for validation, using ChaCha20 and AES-CTR-DRBG generators. We ensure that latter matches those provided in the NTRU specification, which uses the same PRNG.

ARMv7-M implementation. `pqm4` [KRSS19] is the gold standard for Cortex-M4 implementations of PQC schemes. While its NTRU implementation has highly optimized polynomial multiplication and inversion routines, the constant-time sorting routine in use is the `portable3` variant of `djsort` [Ber19], using an architecture-agnostic implementation of the core minimum/maximum operation of the sorting network. Inspection of the binaries reveals that the minimum/maximum idiom was not recognized by the compiler, thus generating suboptimal code without using e.g. conditional instructions. We performed some optimization work on this routine, so as to avoid casting our proposed approach in an excessively favorable light. We switched to the more efficient `portable4` variant of `djsort`, wrote inline assembly versions of the core minimum/maximum operation using conditional operations and a reduced number of memory accesses, and replaced all `long long` (64-bit) variables by 32-bit `long` variables to avoid unnecessary use of multi-precision arithmetic, given that ARMv7-M is a 32-bit architecture. This range reduction does not present an issue in NTRU due to the small lengths (hundreds of elements) of the arrays to be sorted. While it is certainly possible to further optimize this routine, further experiments by us resulted in code size increases, which are undesirable in deeply embedded environments.

Table 1 compares the performance, code size and stack memory usage of encapsulation in the existing version of NTRU (using sampling by sorting), for the original `pqm4` implementation and our optimized version in our STM32F4DISCOVERY testbed; we denote these as “[KRSS19] original” and “[KRSS19] optimized”, respectively, in Table 1. It is seen that our optimizations result in large speedups (43–47%) with negligible effect on code size and none at all on stack usage. While we omit corresponding figures for key generation, our optimizations also outperformed the original implementation, although by smaller amounts (5.4–6.0%); code size and stack usage differences are similar. Results for decapsulation and for the HRSS701 parameter set are not shown, as they do not call the constant-time sorting routine.

5.2 Performance figures and analysis

We present performance figures for NTRU KEM key generation and encapsulation in Tables 2 (for Apple SoCs) and 3 (for ARM Cortex cores); decapsulation does not employ fixed-weight sampling, thus its performance is unaffected by our proposed approach. We present NEON results from the implementations of [NG21] for the HPS2048509 and HPS4096821 parameter sets, and [CCHY24] for the HPS2048677 and HRSS701 parameter sets. AMX results are from the implementation of [GFBL24]. We emphasize that all ARMv8-A implementations use the NEON optimized constant-time sorting routine of [CCHY24]. For the Cortex-M4 core, we use the implementation of [KRSS19], incorporating our optimizations for constant-time sorting. We present performance results as cycle counts, calculating speedups as $c_{\text{sorting}}/c_{\text{shuffling}} - 1$.

Table 1: Comparison of the original pqm4 [KRSS19] NTRU implementation and our optimized version for encapsulation. Code size and stack usage are in bytes. For differences, positive values denote an increase in the optimized version relative to the original one.

Parameter set	Work	Cycle count	Code size	Stack usage
HPS2048509	[KRSS19] original	557 976	191 760	14 068
	[KRSS19] optimized	390 729	191 824	14 068
	Speedup	43%		
	Difference		+0.03%	0%
HPS2048677	[KRSS19] original	801 993	281 504	19 980
	[KRSS19] optimized	552 700	281 568	19 980
	Speedup	45%		
	Difference		+0.02%	0%
HPS4096821	[KRSS19] original	998 761	370 008	23 420
	[KRSS19] optimized	681 402	370 076	23 420
	Speedup	47%		
	Difference		+0.02%	0%

Results for the shuffling approach consist in replacing the `sample_fixed_type` routine by our proposed algorithms, and adjusting the amount of uniform random bits to match the requirements of the shuffling algorithms, as discussed in §4.

We also present performance figures for fixed-weight sampling, by measuring calls to the `sample_fixed_type` routine, whose results are presented in Table 4. Finally, we present code size (Flash) and stack (RAM) usage figures for the Cortex-M4 in Table 5.

Key generation and encapsulation. Our proposed approach achieved performance improvements across the board, for both key generation and encapsulation, save for a few outliers in the former. For Cortex-M4, these improvements come at a negligible cost to code size (Flash), and even a slight improvement in stack (RAM) usage, as seen in Table 5.

With regards to key generation, we see improvements of up to 2.7% for ARMv8-A cores and 6.1% for the Cortex-M4. We recall that NTRU key generation is computationally expensive; disregarding simpler operations, it requires a modulo- q inversion (usually realized by a modulo-2 inversion followed by 8 multiplications), a modulo-3 inversion, 5 extra multiplications, 2 different types of sampling (including `sample_fixed_type`) and pseudo-random number generation. Therefore, it is not surprising that optimizing a single sampling routine results in limited performance improvements.

Results are more significant for encapsulation, which are arguably of more interest than key generation, seeing as, for most cryptographic applications, the former will be run far more often than the latter. We see improvements of up to 44% and 50% for NEON and AMX implementations in ARMv8-A, respectively, and 72% for the Cortex-M4. Improvements correlate well with polynomial multiplication performance, which is fastest for NEON in the HPS2048677 parameter set (based on the faster TMVP approach of [CCHY24]) and in AMX implementations; this is expected due to Amdahl’s law.

Fixed-weight sampling. Table 4 shows that our shuffling approach significantly improves performance of fixed-weight sampling compared to the sampling by sorting approach of previous works. We see very significant speedups for all platforms: factors of up to 6.91 (591%) in ARMv8-A cores and 12.89 (1189%) in the Cortex-M4. Measurements do not include the cost of pseudo-random number generation (i.e. the `randombytes` routine),

Table 2: Cycle counts (in kilocycles) for NTRU KEM key generation (**KG**) and encapsulation (**Enc.**) in the Apple M1 and M3 SoCs.

Param. Set	Sampling	Apple M1				Apple M3			
		NEON		AMX		NEON		AMX	
		KG	Enc.	KG	Enc.	KG	Enc.	KG	Enc.
509	Sorting	218	16.1	170	12.5	214	15.5	164	11.7
	Shuffling	214	12.5	167	8.90	211	12.0	160	8.23
	Speedup	1.7%	29%	2.1%	40%	1.7%	29%	2.3%	43%
677	Sorting	307	20.6	283	17.1	296	19.4	266	16.0
	Shuffling	309	14.9	278	11.9	296	14.1	261	11.0
	Speedup	-0.8%	39%	1.9%	44%	-0.1%	37%	1.9%	46%
821	Sorting	498	28.0	384	19.4	491	27.2	371	18.1
	Shuffling	491	21.8	378	13.1	485	21.2	365	12.1
	Speedup	1.3%	28%	1.7%	48%	1.2%	28%	1.8%	50%
701	N/A	323	14.6	287	11.5	309	13.9	269	10.5
	Slowdown vs. 677 sorting	5.4%	-29%	1.5%	-33%	4.4%	-28%	1.4%	-35%
	Slowdown vs. 677 shuffling	4.6%	-1.7%	3.4%	-3.9%	4.2%	-1.7%	3.3%	-4.4%

which is highly platform-dependent; recall that our approach requires slightly more than half as many pseudo-random bytes as sampling by sorting.

Effect of the FEAT_DIT bit. Recall that, out of all our ARMv8-A testbeds, only the Apple M1 and M3 SoCs implement the FEAT_DIT bit for data-independent timing. In order to ascertain the performance impact of setting this bit, we have also ran benchmarks with FEAT_DIT unset on those Apple SoCs. Although we omit the full dataset, it is available as part of our GitHub repository.

Briefly, for fixed-weight sampling alone (i.e. benchmarks equivalent to those of Table 4, across all parameter sets, both SoCs and both the sorting and shuffling approaches, we see differences of 0 to at most 3 clock cycles (corresponding to at most a 0.14% difference), in both directions: faster with FEAT_DIT set or unset. On average, the difference is 0.01% in favor of FEAT_DIT unset. Given the magnitude and varying direction of the difference, we conjecture that this is merely due to measurement noise, which is to be expected when running benchmarks in such a complex core.

Protocol-level benchmarks display similar behavior, but with a slightly greater magnitude (from 0.14% faster to 0.37% slower with FEAT_DIT set), averaging 0.004% slower across all parameters sets and different SoCs, implementations (AMX/NEON) and shuffling/sorting approaches. There is no clear performance trend either in favor or against setting FEAT_DIT, and the increased magnitude appears to be due to benchmarking more complex code, which further taxes the cache/memory subsystem and AMX coprocessor.

We thus conclude that, at least for Apple SoCs, setting FEAT_DIT has at most a negligible impact on performance, and most likely no impact at all.

Comparison with NTRU-HRSS. It is instructive to compare NTRU-HPS2048677 to NTRU-HRSS701, as both are designed to the same NIST security level. Fortunately, the state-of-the-art NEON implementation of [CCHY24] implements both parameter sets, allowing for a fair comparison. Tables 2 and 3 include rows marked “Slowdown vs. 677 sorting” and “Slowdown vs. 677 shuffling”, computed as $c_{701}/c_{677} - 1$; thus, positive values

Table 3: Cycle counts (in kilocycles) for NTRU KEM key generation (**KG**) and encapsulation (**Enc.**) in ARM Cortex cores.

Param. Set	Sampling	Cortex-A72		Cortex-A57		Cortex-A53		Cortex-M4	
		KG	Enc.	KG	Enc.	KG	Enc.	KG	Enc.
509	Sorting	884	71.2	884	61.5	1243	111	2690	391
	Shuffling	860	54.1	892	51.8	1213	86.7	2536	235
	Speedup	2.7%	32%	-0.8%	19%	2.4%	28%	6.1%	66%
677	Sorting	1140	77.4	1146	65.1	1636	121	4332	553
	Shuffling	1123	55.3	1123	45.3	1600	84.9	4111	327
	Speedup	1.5%	40%	2.1%	44%	2.3%	43%	5.4%	69%
821	Sorting	2156	135	2131	121	2993	197	5740	681
	Shuffling	2111	109	2102	95.2	2951	156	5449	396
	Speedup	2.1%	24%	1.4%	27%	1.4%	26%	5.3%	72%
701	N/A	1200	58.0	1200	53.6	1579	77.7	4210	370
	Slowdown vs. 677 sorting	5.2%	-25%	4.7%	-18%	-3.5%	-36%	-2.8%	-33%
	Slowdown vs. 677 shuffling	6.8%	4.8%	6.9%	18%	-1.3%	-8.5%	2.4%	13%

Table 4: Cycle counts (in kilocycles) for fixed-weight sampling, excluding the cost of uniform random number generation.

Param. Set	Sampling	Apple M1	Apple M3	Cortex-A72	Cortex-A57	Cortex-A53	Cortex-M4
509	Sorting	4.57	4.36	13.0	14.0	23.2	156
	Shuffling	1.11	1.06	2.28	2.30	4.97	13.3
	Speedup	4.12×	4.12×	5.71×	6.11×	4.67×	11.73×
677	Sorting	6.43	6.24	19.5	21.2	35.1	227
	Shuffling	1.48	1.41	3.11	3.07	6.79	18.4
	Speedup	4.34×	4.42×	6.25×	6.91×	5.18×	12.36×
821	Sorting	7.65	7.49	22.6	24.9	40.9	287
	Shuffling	1.79	1.71	3.75	3.77	8.12	22.3
	Speedup	4.28×	4.37×	6.03×	6.61×	5.03×	12.89×

indicate that HRSS701 is slower than HPS2048677, and the contrary for negative values.

Even with the sampling by sorting approach, HPS2048677 is usually faster than HRSS701 for key generation, with the exception of the Cortex-A53 and Cortex-M4 cores; with the shuffling approach, HPS2048677 key generation also outperforms HRSS701 in the Cortex-M4, and reduces the gap in the Cortex-A53. As for encapsulation, HPS2048677 using sampling by sorting was significantly slower than HRSS701 in all cases, by up to 35%, 36% and 33% in Apple SoCs, ARMv8-A Cortex cores and the Cortex-M4, respectively. The shuffling approach closes this gap, with HPS2048677 slower by at most 4.4% in Apple SoCs, and 8.5% in the Cortex-A53; for other ARMv8-A cores, HPS2048677 is actually faster, by up to 18%, and in the Cortex-M4, it is also faster by 13%.

6 Conclusion

In this work, we showed that timing attack-resistant fixed-weight sampling can be performed without using constant-time sorting. We have proposed a new algorithm (Algorithm 6)

Table 5: Code size (Flash) and stack (RAM) usage, in bytes, for ARMv7-M binaries. Statically allocated data (`.data` and `.bss` sections) were reported as zero in all cases. “**Diff.**” refers to the percentual difference between implementations; positive values denote an increase in our version relative to [KRSS19].

Param. set	Work	Code size	Stack usage	
			Key gen.	Encaps.
HPS 2048509	[KRSS19]	191 824	21 360	14 068
	Ours	192 280	20 528	13 228
	Diff.	+0.2%	−3.9%	−6.0%
HPS 2048677	[KRSS19]	281 568	28 472	19 980
	Ours	282 364	27 344	18 852
	Diff.	+0.3%	−4.0%	−5.6%
HPS 4096821	[KRSS19]	370 076	35 224	23 420
	Ours	371 152	33 864	22 060
	Diff.	+0.3%	−3.9%	−5.8%
HRSS701	[KRSS19]	264 492	27 528	18 316

which achieves a running time of $O(n)$, an improvement over $O(n \log^2(n))$ for previous, sorting network-based approaches. This results in performance improvements in actual implementations across a range of different platforms, from deeply embedded to high-performance laptop CPUs. Additionally, the amount of random data needed for sampling is reduced by almost half, which is advantageous for architectures without instructions to accelerate cryptographically secure PRNGs. Moreover, our proposed method may be simpler to implement in an optimized fashion than constant-time sorting networks.

This solves a long-standing open problem: to date, the best alternative was the NTRU-HRSS variant, which also seeks to eliminate the cost of constant-time sorting required for sampling fixed-weight polynomials. As discussed in §5, a modified NTRU-HPS2048677, using our proposed approach, nearly closes the performance gap to NTRU-HRSS701 in some cases, and actually outperforms it in others (recalling that both are designed to the same NIST security level). We also note that key and ciphertext sizes for NTRU-HPS2048677 are smaller: 930 (resp. 1138) bytes for the public key and ciphertext, and 1234 (resp. 1450) bytes for the private key, for NTRU-HPS2048677 (resp. NTRU-HRSS701). Finally, the need to support both NTRU-HPS and NTRU-HRSS to achieve different security levels results in increased implementation complexity, e.g. due to the HRSS-specific version of `Lift` [CDH⁺20, §1.9.3] and the additional `Ternary_Plus` sampling routine [CDH⁺20, §1.10.4]. In light of these arguments, we call into question the need for a separate NTRU-HRSS parameter set.

Future work. Although NTRU is no longer being considered by NIST, we recall that it has been standardized in other forums [Ins09, Ame17]. Since our proposed Algorithm 6 improves upon the existing fixed-weight sampling by sorting approach mandated by the NTRU specification submitted to NIST [CDH⁺20], we suggest amending NTRU specifications to use Algorithm 6, and incorporating it into any future standardization efforts (for instance, we note that FrodoKEM [BCD⁺16] is also no longer under consideration by NIST, but is being considered for standardization by ISO [Int23c]). Similarly, as discussed in §3, other PQC schemes such as Streamlined NTRU Prime and NTRU LPrime are likely to benefit from our approach, and we encourage investigations of its use in these schemes. We also suggest developing implementations for other widely-used architectures, in particular, Intel (using AVX2 and AVX-512 SIMD extensions) and the recently released

ARMv8.1-M Helium SIMD instruction set for deeply embedded systems [Dir19].

Algorithm 8, as stated, is not amenable to vectorization, due to a loop-carried dependency between iterations of its main loop. Using a similar idea as the initial step of MERGESHUFFLE (Algorithm 4), vectorization becomes possible; we developed a prototype implementation that confirms its potential for large speedups, especially on wide CPUs such as the M1 and M3. However, without applying the remaining steps of MERGESHUFFLE, the resulting permutation is biased, which may create an avenue of attack. An alternative we envisioned involves sampling from the hypergeometric distribution; however, this is an uncommon distribution in cryptography, and we were unable to find any efficient, constant-time algorithms. We invite future work into either modifying MERGESHUFFLE to be constant-time, or to propose efficient, constant-time hypergeometric sampling algorithms.

While Algorithm 5 already consumes fewer random bits than the approach mandated by the NTRU specification, there is room for improvement, as discussed in §4. This is particularly relevant for architectures without instructions to accelerate cryptographically-secure PRNGs, in which case the performance bottleneck may be the generation of uniformly random samples consumed by Algorithm 5. We encourage further research into parameter choices and new approaches that may reduce randomness requirements.

Acknowledgements

The first and third author thank the Cryptography Research Centre at the Technology Innovation Institute for financial support during the development of this work. The second author is supported by São Paulo Research Foundation (FAPESP) grant 2022/09891-4.

We thank the anonymous reviewers for helpful suggestions.

References

- [AAB⁺22] Carlos Aguilar-Melchor, Nicolas Aragon, Slim Bettaieb, Loïc Bidoux, Olivier Blazy, Jean-Christophe Deneuville, Philippe Gaborit, Edoardo Persichetti, Gilles Zémor, Jurjen Bos, Arnaud Dion, Jerome Lacan, Jean-Marc Robert, and Pascal Veron. HQC. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions>.
- [ABB⁺22] Nicolas Aragon, Paulo Barreto, Slim Bettaieb, Loic Bidoux, Olivier Blazy, Jean-Christophe Deneuville, Phillipe Gaborit, Shay Gueron, Tim Guneysu, Carlos Aguilar-Melchor, Rafael Misoczki, Edoardo Persichetti, Nicolas Sendrier, Jean-Pierre Tillich, Gilles Zémor, Valentin Vasseur, Santosh Ghosh, and Jan Richter-Brokmann. BIKE. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions>.
- [ABC⁺22] Martin R. Albrecht, Daniel J. Bernstein, Tung Chou, Carlos Cid, Jan Gilcher, Tanja Lange, Varun Maram, Ingo von Maurich, Rafael Misoczki, Ruben Niederhagen, Kenneth G. Paterson, Edoardo Persichetti, Christiane Peters, Peter Schwabe, Nicolas Sendrier, Jakub Szefer, Cen Jung Tjhai, Martin Tomlinson, and Wen Wang. Classic McEliece. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-4-submissions>.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th Annual ACM Symposium on Theory of Computing*, pages 99–108. ACM Press, May 1996. doi:10.1145/237814.237838.

- [Ame17] American National Standards Institute. Lattice-based polynomial public key establishment algorithm for the financial services industry. ASC X9.98-2010 (R2017), 2017.
- [ARM] ARM Limited. How is instruction timing affected by the FEAT_DIT architectural feature? URL: <https://developer.arm.com/documentation/ka005181/latest/>.
- [ARM15] ARM Limited. Cortex®-A72 software optimization guide, 2015. URL: <https://developer.arm.com/documentation/uan0016/a/>.
- [ARM23] ARM Limited. Arm® architecture reference manual for A-profile architecture, 2023. URL: <https://developer.arm.com/documentation/ka005181/1-0/>.
- [Bat68] K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference, AFIPS '68 (Spring)*, page 307–314, New York, NY, USA, 1968. Association for Computing Machinery. doi:10.1145/1468075.1468121.
- [BBC+20] Daniel J. Bernstein, Billy Bob Brumley, Ming-Shing Chen, Chitchanok Chuengsatiansup, Tanja Lange, Adrian Marotzke, Bo-Yuan Peng, Nicola Taveri, Christine van Vredendaal, and Bo-Yin Yang. NTRU Prime. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.
- [BBHL18] Axel Bacher, Olivier Bodini, Alexandros Hollender, and Jérémie O. Lumbroso. MergeShuffle: a very fast, parallel random permutation algorithm. In Luca Ferrari and Malvina Vamvakari, editors, *Proceedings of the 11th International Conference on Random and Exhaustive Generation of Combinatorial Structures, GASCom 2018, Athens, Greece, June 18-20, 2018*, volume 2113 of *CEUR Workshop Proceedings*, pages 43–52, Aachen, Germany, 2018. CEUR-WS.org. URL: <http://ceur-ws.org/Vol-2113/paper3.pdf>.
- [BBHT17] Axel Bacher, Olivier Bodini, Hsien-Kuei Hwang, and Tsung-Hsi Tsai. Generating random permutations by coin tossing: Classical algorithms, new analysis, and modern implementation. *ACM Trans. Algorithms*, 13(2), feb 2017. doi:10.1145/3009909.
- [BCC+23] Gustavo Banegas, Kévin Carrier, André Chailloux, Alain Couvreur, Thomas Debris-Alazard, Philippe Gaborit, Pierre Karpman, Johanna Loyer, Ruben Niederhagen, Nicolas Sendrier, Benjamin Smith, and Jean-Pierre Tillich. Wave. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- [BCD+16] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 1006–1018. ACM Press, October 2016. doi:10.1145/2976749.2978425.
- [BCLv17] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU prime: Reducing attack surface at low cost. In Carlisle

- Adams and Jan Camenisch, editors, *SAC 2017: 24th Annual International Workshop on Selected Areas in Cryptography*, volume 10719 of *Lecture Notes in Computer Science*, pages 235–260. Springer, Heidelberg, August 2017. doi: [10.1007/978-3-319-72565-9_12](https://doi.org/10.1007/978-3-319-72565-9_12).
- [BCLv19] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU Prime. Technical report, National Institute of Standards and Technology, 2019. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-2-submissions>.
- [BCS13] Daniel J. Bernstein, Tung Chou, and Peter Schwabe. McBits: Fast constant-time code-based cryptography. In Guido Bertoni and Jean-Sébastien Coron, editors, *Cryptographic Hardware and Embedded Systems – CHES 2013*, volume 8086 of *Lecture Notes in Computer Science*, pages 250–272. Springer, Heidelberg, August 2013. doi: [10.1007/978-3-642-40349-1_15](https://doi.org/10.1007/978-3-642-40349-1_15).
- [Ber04] Daniel J. Bernstein. Cache-timing attacks on AES. <http://cr.yp.to/papers.html#cachetiming>, 2004.
- [Ber18] Daniel J. Bernstein. Divergence bounds for random fixed-weight vectors obtained by sorting, April 2018. URL: <https://ntruprime.cr.yp.to/divergence-20180430.pdf>.
- [Ber19] Daniel J. Bernstein. djbsort. <https://sorting.cr.yp.to>, 2019.
- [CCHY24] Han-Ting Chen, Yi-Hua Chung, Vincent Hwang, and Bo-Yin Yang. Algorithmic views of vectorized polynomial multipliers – NTRU. In Anupam Chattopadhyay, Shivam Bhasin, Stjepan Picek, and Chester Rebeiro, editors, *Progress in Cryptology – INDOCRYPT 2023*, pages 177–196, Cham, 2024. Springer Nature Switzerland. doi: [10.1007/978-3-031-56235-8_9](https://doi.org/10.1007/978-3-031-56235-8_9).
- [CDH⁺20] Cong Chen, Oussama Danba, Jeffrey Hoffstein, Andreas Hülsing, Joost Rijneveld, John M. Schanck, Peter Schwabe, William Whyte, Zhenfei Zhang, Tsunekazu Saito, Takashi Yamakawa, and Keita Xagawa. NTRU. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.
- [CWS⁺24] Boru Chen, Yingchen Wang, Pradyumna Shome, Christopher W. Fletcher, David Kohlbrenner, Riccardo Paccagnella, and Daniel Genkin. GoFetch: Breaking constant-time cryptographic implementations using data memory-dependent prefetchers. In *USENIX Security*, 2024. URL: <https://gofetch.fail/files/gofetch.pdf>.
- [Dan19] Oussama Danba. Optimizing NTRU using AVX2. Master’s thesis, Radboud University, 2019. URL: https://www.ru.nl/publish/pages/769526/y_oussama_danba.pdf.
- [dG15] Wouter de Groot. A performance study of X25519 on Cortex-M3 and M4. Master’s thesis, Eindhoven University of Technology, 2015. URL: <https://research.tue.nl/files/47038543/800603-1.pdf>.
- [Dir19] Rhonda Dirvin. Next-generation Armv8.1-M architecture: Delivering enhanced machine learning and signal processing for the smallest embedded devices. <https://www.arm.com/company/news/2019/02/next-generation-armv8-1-m-architecture>, 2019.

- [Dur64] Richard Durstenfeld. Algorithm 235: Random permutation. *Commun. ACM*, 7(7):420, jul 1964. doi:10.1145/364520.364540.
- [Fog22] Agner Fog. The microarchitecture of Intel, AMD, and VIA CPUs. <https://www.agner.org/optimize/microarchitecture.pdf>, 2022.
- [FY38] R. A. Fisher and F. Yates. *Statistical tables for biological, agricultural and medical research*. Oliver & Boyd, Oxford, England, 3rd edition, 1938.
- [GFBL24] Décio Luiz Gazzoni Filho, Guilherme Brandão, and Julio López. Fast polynomial multiplication using matrix multiplication accelerators with applications to NTRU on Apple M1/M3 SoCs. *IACR Communications in Cryptology*, 1(1), 2024. doi:10.62056/a3txommol.
- [GHJ⁺22] Qian Guo, Clemens Hlauschek, Thomas Johansson, Norman Lahr, Alexander Nilsson, and Robin Leander Schröder. Don't reject this: Key-recovery timing attacks due to rejection-sampling in HQC and BIKE. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2022(3):223–263, 2022. doi:10.46586/tches.v2022.i3.223-263.
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery. URL: <https://doi.org/10.1145/237814.237866>.
- [GS23] Garrett Gu and Hovav Shacham. Constant-time wasmtime, for real this time: End-to-end verified zero-overhead constant-time programming for the web and beyond, 2023. URL: <https://arxiv.org/abs/2311.14246>.
- [HPS96] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: a new high speed public key cryptosystem. CRYPTO '96 rump session, 1996. <https://web.securityinnovation.com/hubfs/files/ntru-orig.pdf>.
- [HRSS17] Andreas Hülsing, Joost Rijneveld, John M. Schanck, and Peter Schwabe. High-speed key encapsulation from NTRU. In Wieland Fischer and Naofumi Homma, editors, *Cryptographic Hardware and Embedded Systems – CHES 2017*, volume 10529 of *Lecture Notes in Computer Science*, pages 232–252. Springer, Heidelberg, September 2017. doi:10.1007/978-3-319-66787-4_12.
- [Ins09] Institute of Electrical and Electronics Engineers. IEEE standard specification for public key cryptographic techniques based on hard problems over lattices. IEEE Std 1363.1-2008, 2009. doi:10.1109/IEEESTD.2009.4800404.
- [Int23a] Intel Corporation. Data operand independent timing instruction set architecture (ISA) guidance, 2023. URL: <https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/resources/data-operand-independent-timing-instructions.html>.
- [Int23b] Intel Corporation. Data operand independent timing instructions, 2023. URL: <https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/best-practices/data-operand-independent-timing-isa-guidance.html>.
- [Int23c] International Organization for Standardization. FrodoKEM: Learning with errors key encapsulation preliminary draft standard, 2023. URL: <https://frodokem.org/files/FrodoKEM-ISO-20230314.pdf>.

- [Joh22] Dougall Johnson. Apple M1 microarchitecture research. <https://dougallj.github.io/applecpu/firestorm.html>, 2022.
- [Knu97] Donald E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley, Boston, third edition, 1997.
- [Knu98] Donald E. Knuth. *The Art of Computer Programming, Volume 3: Sorting and Searching*. Addison-Wesley Longman Publishing Co., Inc., USA, 2nd edition, 1998.
- [Koc96] Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO’96*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113. Springer, Heidelberg, August 1996. doi:10.1007/3-540-68697-5_9.
- [KRSS19] Matthias J. Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. pqm4: Testing and benchmarking NIST PQC on ARM Cortex-M4. Workshop Record of the Second PQC Standardization Conference, 2019. URL: <https://csrc.nist.gov/CSRC/media/Events/Second-PQC-Standardization-Conference/documents/accepted-papers/kannwischer-pqm4.pdf>.
- [Lam88] M. Lam. Software pipelining: An effective scheduling technique for VLIW machines. *SIGPLAN Not.*, 23(7):318–328, jun 1988. doi:10.1145/960116.54022.
- [Lem19] Daniel Lemire. Fast random integer generation in an interval. *ACM Trans. Model. Comput. Simul.*, 29(1), jan 2019. doi:10.1145/3230636.
- [MG02] Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*, volume 671 of *The Springer International Series in Engineering and Computer Science*. Springer, New York, NY, first edition, 2002. doi:10.1007/978-1-4615-0897-7.
- [Nat17] National Institute of Standards and Technology. Post-quantum cryptography standardization: Call for proposals announcement, 2017. <https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization>.
- [NG21] Duc Tri Nguyen and Kris Gaj. Fast NEON-based multiplication for lattice-based NIST post-quantum cryptography finalists. In Jung Hee Cheon and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 12th International Workshop, PQCrypto 2021*, pages 234–254. Springer, Heidelberg, 2021. doi:10.1007/978-3-030-81293-5_13.
- [Por18] Thomas Pornin. Constant-time multiplication. <https://www.bearssl.org/ctmul.html>, 2018.
- [Rao61] C. Radhakrishna Rao. Generation of random permutations of given number of elements using random sampling numbers. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)*, 23(3):305–307, 1961. URL: <http://www.jstor.org/stable/25049166>.
- [San62] Martin Sandelius. A simple randomization procedure. *Journal of the Royal Statistical Society. Series B (Methodological)*, 24(2):472–481, 1962. URL: <http://www.jstor.org/stable/2984238>.

- [Sen21] Nicolas Sendrier. Secure sampling of constant-weight words – application to BIKE. Cryptology ePrint Archive, Report 2021/1631, 2021. <https://eprint.iacr.org/2021/1631>.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997. doi:10.1137/S0097539795293172.
- [TSS⁺03] Yukiyasu Tsunoo, Teruo Saito, Tomoyasu Suzaki, Maki Shigeri, and Hiroshi Miyachi. Cryptanalysis of DES implemented on computers with cache. In Colin D. Walter, Çetin Kaya Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2003*, volume 2779 of *Lecture Notes in Computer Science*, pages 62–76. Springer, Heidelberg, September 2003. doi:10.1007/978-3-540-45238-6_6.
- [TTMM02] Yukiyasu Tsunoo, Etsuko Tsujihara, Kazuhiko Minematsu, and Hiroshi Miyachi. Cryptanalysis of block ciphers implemented on computers with cache. In *Proceedings of the International Symposium on Information Theory and Its Applications, ISITA 2002*, pages 803–806, 2002.
- [War12] Henry S. Warren. *Hacker’s Delight*. Addison-Wesley Professional, 2nd edition, 2012.