



# Fast polynomial multiplication using matrix multiplication accelerators with applications to NTRU on Apple M1/M3 SoCs

Décio Luiz Gazzoni Filho<sup>1,2</sup> , Guilherme Brandão<sup>3</sup> and Julio López<sup>1</sup>

<sup>1</sup> Universidade Estadual de Campinas (UNICAMP), Instituto de Computação, Campinas, Brazil

<sup>2</sup> State University of Londrina, Department of Electrical Engineering, Londrina, Brazil

<sup>3</sup> Independent Researcher, Londrina, Brazil

**Abstract.** Efficient polynomial multiplication routines are critical to the performance of lattice-based post-quantum cryptography (PQC). As PQC standards only recently started to emerge, CPUs still lack specialized instructions to accelerate such routines. Meanwhile, deep learning has grown immeasurably in importance. Its workloads call for teraflops-level of processing power for linear algebra operations, mainly matrix multiplication. Computer architects have responded by introducing ISA extensions, coprocessors and special-purpose cores to accelerate such operations. In particular, Apple ships an undocumented matrix-multiplication coprocessor, AMX, in hundreds of millions of mobile phones, tablets and personal computers. Our work repurposes AMX to implement polynomial multiplication and applies it to the NTRU cryptosystem, setting new speed records on the Apple M1 and M3 systems-on-chip (SoCs): polynomial multiplication, key generation, encapsulation and decapsulation are sped up by  $1.54\text{--}3.07\times$ ,  $1.08\text{--}1.33\times$ ,  $1.11\text{--}1.50\times$  and  $1.20\text{--}1.98\times$ , respectively, over the previous state-of-the-art.

**Keywords:** PQC · NTRU · Apple Silicon · Accelerators

## 1 Introduction

In the 1990s, Shor [Sho97] described an efficient quantum algorithm to solve hard problems (integer factorization and discrete logarithms) for classical computers, compromising many existing cryptosystems. *Post-quantum cryptography* (PQC) seeks to develop new public-key cryptosystems based on hard computational problems resistant to quantum attacks, such as those based on lattices. Research into efficient implementation flourished as multiprecision integer arithmetic gave way to polynomial multiplication modulo “small” (CPU word size) integers for lattice-based schemes, introducing new challenges and performance tradeoffs.

In their Turing lecture [HP19], computer architecture pioneers Hennessy and Patterson claim that “(a)n era without Dennard scaling, along with reduced Moore’s Law and Amdahl’s Law in full effect means inefficiency limits improvement in performance to only a few percent per year”, suggesting that further hardware improvements must come from *domain-specific architectures* tailored to specific applications. This is seen in the addition of SIMD extensions and special-purpose instructions (such as for symmetric cryptography and binary polynomial multiplication) to CPUs, as well as increasingly popular accelerators such as programmable graphics processing units (GPUs), FPGAs and others. Such developments have opened up new research agendas for efficient implementation of cryptosystems.

---

E-mail: [decio.gazzoni@ic.unicamp.br](mailto:decio.gazzoni@ic.unicamp.br) (Décio Luiz Gazzoni Filho), [brandaogbs@gmail.com](mailto:brandaogbs@gmail.com) (Guilherme Brandão), [jlopez@ic.unicamp.br](mailto:jlopez@ic.unicamp.br) (Julio López)



The meteoric rise of artificial intelligence, machine learning and deep learning in the 2010s, and their demand for teraflops-level linear algebra performance, led to an introduction of domain-specific architectures such as tensor cores for NVIDIA’s GPUs [MCL<sup>+</sup>18], IBM Power ISA’s Matrix-Multiply Assist (MMA) [MBB<sup>+</sup>21], Intel’s Advanced Matrix Extensions (AMX) [Int22], and ARMv9-A’s Scalable Matrix Extensions (SME) [WMS22]. The first three have launched in production hardware in 2017, 2021 and 2023, respectively, whereas the latter, to our knowledge, hasn’t shipped as of 2023. Although less publicized, a matrix multiplication accelerator has shipped in hundreds of millions of mobile phones, tablets and personal computers, specifically those made by Apple, starting with the 2019 A13 system-on-chip (SoC) used in the iPhone 11 [Rod20, Section 7.6]. It is a coprocessor named AMX (no relationship to Intel AMX), also present in newer generations of Apple SoCs, such as the M-series powering ARM-based Apple personal computers.

Our work investigates repurposing AMX to accelerate the core polynomial multiplication routine of lattice-based post-quantum cryptosystems, by implementing NTRU [HPS98], specifically the proposal [CDH<sup>+</sup>20] that advanced to round 3 of NIST’s PQC standardization process [Nat17], for Apple’s M1 and M3 SoCs. Despite NIST’s preference for Kyber [SAB<sup>+</sup>20], NTRU is standardized by IEEE Std 1363.1 [Ins09] and ASC X9.98 [Ame17]. It is also representative of a wider class of lattice-based schemes based on NTT-unfriendly rings such as SABER [DKR<sup>+</sup>20] and FrodoKEM [BCD<sup>+</sup>16, NAB<sup>+</sup>20]; the latter is recommended by the German BSI [Bun21] and under standardization by ISO [Int23c]. In case of cryptanalytic attacks on Kyber, NTRU may be the target of renewed interest.

A potential roadblock is that AMX is undocumented by Apple; programmers are intended to access it via calls to Apple’s Accelerate framework [App23], which provides accelerated versions of the industry-standard BLAS library [LHK79] and the proprietary BNNS neural network API. Fortunately, extensive reverse-engineering efforts [Joh22b, Han23, Caw24] document AMX’s instruction set and performance characteristics, allowing us to develop a full NTRU implementation with AMX-accelerated polynomial multiplication routines, as presented in this paper. The authors were not directly involved in these reverse-engineering efforts, working only with the information made public in these references; in the interest of full disclosure, the first author has corresponded with one of the reverse-engineering researchers (Peter Cawley), in particular providing him access to M3 hardware to investigate functional changes in its AMX unit over previous generation SoCs.

**Related works.** Google’s Tensor Processing Unit is described in [JYP<sup>+</sup>17]. It achieves a peak throughput of 92 TOPS/s and 15–30× faster neural network inference than CPUs and GPUs of that era. [MBB<sup>+</sup>21] reports a 2× speedup by using MMA over regular vector code for double-precision matrix multiplication in the POWER10 CPU.

Some works demonstrate speedups from NVIDIA tensor cores, including for PQC schemes [MCL<sup>+</sup>18, WZF<sup>+</sup>22, LSH<sup>+</sup>22, LSZH22]. However, GPU architectures, being high-throughput but high-latency, need aggressive batching (thousands to tens of thousands of parallel KEM/signature operations) to realize their performance potential, severely limiting applicable use cases. We view direct comparisons of batch and online implementations (for CPUs and, as we shall show, AMX) as unfair, and refrain from doing so in this paper.

We found no references to Apple AMX in the scientific literature, and Apple does not officially document AMX functionality or performance characteristics. We resort to reverse engineering efforts of [Joh22b, Han23, Caw24], which we summarize in Section 3.

Lastly, we consider ARMv8-A implementations of NTRU. Karatsuba and Toom-Cook algorithms are implemented in [NG21], achieving speedups on the M1 of up to 6.68× and 8.49× for encapsulation and decapsulation, respectively, versus the reference implementation. [CCHY23] set new speed records for the HPS2048677 and HRSS701 parameter sets using the TMVP algorithm. They also optimized polynomial inversion and constant-time sorting, speeding up key generation, encapsulation and decapsulation by

7.67×, 2.48× and 1.77×, respectively, compared to [NG21] on the Cortex-A72 core.

**Our contributions.** We present the first (to our knowledge) implementation targeting AMX in the scientific literature, and also the first cryptographic implementation on a CPU-coupled matrix multiplication accelerator; previous works targeted GPU tensor cores.

Unlike GPUs, CPU-coupled matrix multiplication accelerators (coprocessors or instruction set extensions) can achieve high throughput without sacrificing latency. Indeed, we set new NTRU speed records on Apple M1/M3 SoCs, outperforming previous state-of-the-art NEON implementations in polynomial multiplication, key generation, encapsulation and decapsulation by 1.54–3.07×, 1.08–1.33×, 1.11–1.50× and 1.20–1.98×, respectively. Notably, we achieve such speedups under conditions previously seen in CPU implementations only: latency (cycle count) for a *single* execution of the scheme’s operations, with *no batching at all*. Usually, accelerator (e.g. GPU) implementations require enormous batching levels to achieve claimed speedups, with high latencies for a single execution.

SIMD extensions required a paradigm shift from a scalar to a vector (1D) view of computation; matrix multiplication accelerators further shift towards a matrix (2D) view. AMX’s peak throughput is only achieved for its main matrix (outer product<sup>1</sup>) operation; vector operation throughput is similar to NEON. Also, its instruction set is rather limited, seemingly designed to implement only a restricted set of algorithms. Our main contribution is repurposing AMX for a new application (polynomial multiplication) while realizing its performance potential; this required a deep rethinking of polynomial multiplier architectures to expose two-dimensional parallelism, maximizing the ratio of matrix to vector operations to reap as much of its throughput as possible. We conjecture that the techniques we propose for AMX may be applicable to many similar accelerators soon to reach the market.

Our implementation is available at <https://github.com/dgazzoni/NTRU-AMX> under various licenses (such as Apache and MIT) for existing code reused in our implementation, and Creative Commons CC0 (“No Rights Reserved”) for the authors’ own code. It also includes an optimized implementation of NIST’s `randombytes()` function based on AES-CTR-DRBG, using ARMv8 Cryptographic Extensions’ AES instructions, which may be of independent interest for other NTRU implementations and PQC schemes.

## 2 Preliminaries

**Notation.** Let  $a(x) = a_0 + a_1x + \dots a_{n-1}x^{n-1}$ . A “slice” of the  $i$ -th through  $j$ -th coefficients ( $j \geq i$ ) of  $a(x)$  is written as  $a_{i:j}(x) = a_i x^i + a_{i+1} x^{i+1} + \dots + a_j x^j$ . Boldface variable names refer to an associated row vector representation:  $\mathbf{a}_{i:j} = [a_i, a_{i+1}, \dots, a_j]$ . Let  $\mathbf{x}_{i:j} = [x^i, x^{i+1}, \dots, x^j]$ ; then  $a_{i:j}(x) = \mathbf{a}_{i:j} \mathbf{x}_{i:j}^\top$ . Finally,  $i : k : j$  represents ranges with a non-unit step of  $k$ ; e.g.  $\mathbf{a}_{i:k:j} = [a_i, a_{i+k}, a_{i+2k}, \dots, a_j]$ . We also combine slice notation with C’s array indexing notation, i.e.  $\mathbf{X}[i : j]$  selects elements of index  $i, \dots, j$  of  $\mathbf{X}$ .

**The NTRU cryptosystem.** Based on the pioneering work of [HPS98], NTRU [CDH<sup>+</sup>20] is a key encapsulation mechanism (KEM) merging the previous NTRUEncrypt [ZCH<sup>+</sup>19] and NTRU-HRSS-KEM [SHRS17] proposals, providing the suggested NTRU-HPS and NTRU-HRSS parameter sets, respectively. Based on structured lattices, it uses polynomial arithmetic modulo  $x^n - 1$  for prime  $n$  with coefficients reduced either modulo 3 or a power of two,  $q$ . The original cryptosystem [HPS98] is a partially correct, probabilistic public-key encryption (PKE) scheme. The NIST submission uses the techniques of [HPS96] to obtain a perfectly correct, deterministic PKE, and constructs a KEM using a generic transformation from this PKE [HHK17] which is IND-CCA2 secure in the random oracle model (ROM), and in the quantum-accessible ROM under a non-standard assumption.

<sup>1</sup>Also known as a rank-1 update, implemented e.g. by the BLAS [LHKK79] level 2 subroutine `xGER`.

Let  $\Phi_1 = x - 1$  and  $\Phi_n = x^{n-1} + x^{n-2} + \dots + 1$ , so that  $\Phi_1\Phi_n = x^n - 1$ . Let  $S_3 = \mathbb{Z}_3/\Phi_n$ ; we define the canonical representative of an element in  $S_3$ , denoted as  $\underline{S}_3$ , as the polynomial of degree at most  $n - 2$  with coefficients in  $\{-1, 0, 1\}$ .

Algorithms 2.1, 2.2 and 2.3 define the NTRU PKE, where most of the execution time is spent; for a full description of the scheme, including the `Samplef,g` routine, see [CDH<sup>+</sup>20]. `Lift(m)`, for a polynomial  $m(x)$ , is defined as  $\underline{S}_3(m)$  for HPS and  $\Phi_1 \cdot \underline{S}_3(m/\Phi_1)$  for HRSS parameter sets. Algorithms 2.2 and 2.3 require 1 and 3 polynomial multiplications (respectively) in the HPS parameter sets, and `Lift` adds an extra multiplication for HRSS. At first glance, Algorithm 2.1 uses 5 polynomial multiplications, but polynomial inversion modulo  $q$  in line 3 is usually realized as modulo-2 inversion followed by 4 Newton iterations to lift the result to  $\mathbb{Z}_q/(\Phi_1\Phi_n)$ , which require an extra 8 multiplications, for a total of 13.

---

**Algorithm 2.1** PKE.KeyGen: PKE keypair generation

---

**Input:** seed

**Output:** (sk =  $(f, f_p, h_q)$ , pk =  $h$ )

- 1:  $(f, g) \leftarrow \text{Sample}_{f,g}(\text{seed})$
  - 2:  $f_p \leftarrow (1/f) \pmod{3, \Phi_n}$
  - 3:  $v_1 \leftarrow 1/(3gf) \pmod{q, \Phi_n}$
  - 4:  $h \leftarrow (v_1(3g)^2) \pmod{q, \Phi_1\Phi_n}$
  - 5:  $h_q \leftarrow (v_1f^2) \pmod{q, \Phi_1\Phi_n}$
  - 6: **return** (sk =  $(f, f_p, h_q)$ , pk =  $h$ )
- 

**Algorithm 2.2** PKE.Enc: PKE encryption

---

**Input:**  $h, (r, m)$

**Output:**  $c$

- 1:  $m' \leftarrow \text{Lift}(m)$
  - 2:  $c \leftarrow r \cdot h + m' \pmod{q, \Phi_1\Phi_n}$
  - 3: **return**  $c$
- 

---

**Algorithm 2.3** PKE.Dec: PKE decryption

---

**Input:**  $((f, f_q, h_q), c)$

**Output:**  $(r, m, \text{fail})$

- 1:  $v_1 \leftarrow cf \pmod{q, \Phi_1\Phi_n}$
  - 2:  $m_0 \leftarrow v_1f_p \pmod{3, \Phi_n}$
  - 3:  $m_1 \leftarrow \text{Lift}(m_0)$
  - 4:  $r \leftarrow (c - m_1)h_q \pmod{q, \Phi_n}$
  - 5: **if**  $(r, m_0) \in \mathcal{L}_r \times \mathcal{L}_m$  **then**  
     ▷ See [CDH<sup>+</sup>20, Sections 1.2 and 1.3] for definitions of  $\mathcal{L}_r$  and  $\mathcal{L}_m$
  - 6:     **return**  $(r, m_0, 0)$
  - 7: **else**
  - 8:     **return**  $(r, m_0, 1)$
- 

**Table 1:** NTRU parameters and sizes (in bytes).

Parameter	HPS2048509	HPS2048677	HPS4096821	HRSS701
$(n, q)$	(509, 2048)	(677, 2048)	(821, 4096)	(701, 8192)
Key sizes (public, private)	699, 935	930, 1234	1230, 1590	1138, 1450
Ciphertext sizes	699	930	1230	1138

The submission includes four parameter sets: HPS2048509, HPS2048677, HPS4096821, and HRSS701. The first three use fixed-weight sample spaces proposed by Hoffstein, Pipher, and Silverman [HPS96, HPS98]; the last uses arbitrary-weight sample spaces proposed by Hülsing, Rijneveld, Schanck, and Schwabe [HRSS17]. Table 1 lists parameters and sizes.

### 3 Apple’s AMX coprocessor

In this section, we compile and summarize information about AMX from reverse engineering efforts and analyses of Apple patents by [Joh22b, Han23, Caw24]. Such information is, by its nature, conjectural, although backed by extensive functional and performance experiments. This section seeks to make this paper as self-contained as possible, and is also the first (to our knowledge) presentation of AMX in the scientific literature.

We do not lay claim to any of these discoveries; our main contribution lies in cataloguing techniques to repurpose AMX to efficiently perform polynomial multiplication modulo “small” integers, and its application to NTRU – although other lattice-based PQC cryptosystems should also benefit, a task left as future work in Section 6.

AMX is a coprocessor from Apple to accelerate matrix multiplication operations, first introduced in the Apple A13 SoC powering the iPhone 11 [Rod20]. Functionally, the M1 and M3 AMX units are almost identical; there are two added features in the M3’s which we did not find useful for our implementation. Performance-wise, our results of Section 5 indicate that M3 delivers improvements over M1.

### 3.1 Programmer’s model

AMX exposes 80 64-byte registers, split into eight **X** and eight **Y** registers ( $X_0$  to  $X_7$  and  $Y_0$  to  $Y_7$ ), and 64 **Z** registers better viewed as rows of a matrix, denoted as  $Z[0], \dots, Z[63]$ . Figure 1, redrawn from an Apple patent [SBG<sup>+</sup>16, Figure 2], is a visualization of intended register file organization. Some instructions can concatenate together either the **X** or **Y** registers for bitwise addressing as 512-byte circular buffers, for which we use the slice notation defined in Section 2. Figure 2 shows different **X**/**Y** register addressing notations.

	$X[0]$	$X[1]$	$\dots$	$X[n]$
$Y[0]$	$Z[0][0] += Y[0]X[0]$	$Z[0][1] += Y[0]X[1]$	$\dots$	$Z[0][n] += Y[0]X[n]$
$Y[1]$	$Z[1][0] += Y[1]X[0]$	$Z[1][1] += Y[1]X[1]$	$\dots$	$Z[1][n] += Y[1]X[n]$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
$Y[n]$	$Z[n][0] += Y[n]X[0]$	$Z[n][1] += Y[n]X[1]$	$\dots$	$Z[n][n] += Y[n]X[n]$

**Figure 1:** AMX register file organization.

Bytes	0	$\dots$	63	64	$\dots$	127	128	$\dots$	447	448	$\dots$	511
Registers	$X_0, Y_0$			$X_1, Y_1$			$\dots$			$X_7, Y_7$		
Slice notation	$X[0 : 31]$			$X[32 : 63]$			$\dots$			$X[224 : 255]$		
(16-bit elements)	$Y[0 : 31]$			$Y[32 : 63]$			$\dots$			$Y[224 : 255]$		

**Figure 2:** Relationship between bytes, registers and slices in **X** and **Y** AMX registers.


**X** and **Y** are inputs and **Z** are outputs for most instructions, with a few exceptions. Data cannot be moved between CPU and AMX registers directly; it must go through memory.

Different data types can be operated on (8-, 16- or 32-bit integers, or 16-, 32- or 64-bit floating-point values). Input/output lane widths may be identical or mixed in specific combinations [Caw24], but we only use 16-bit integer inputs and outputs. The inputs to outer product operations ( $X_i$  and  $Y_i$ ) are 32-element vectors of 16 bits each, resulting in a  $32 \times 32$  output matrix of 16-bit elements; this is smaller than **Z**’s available storage space (16,384 out of an available 32,768 bits). As such, each row of the output matrix is mapped to either the even or odd rows of **Z**, which are fully populated with output coefficients, as the size of a row of **Z** (512 bits) matches the size of a row of the output matrix.

CPUs in a cluster can share an AMX unit through per-CPU replication of architectural state [Joh22a, Han23]. Instructions are tagged with their source CPU to identify their copy of the state; thus, multiple CPUs may interleave instruction execution [Han23]. Special instructions (`set/clr`) control AMX’s enabled/disabled status for each thread [Caw24].

### 3.2 Programming interface

AMX instructions are inserted into the CPU’s instruction stream, and once no longer speculative, are dispatched to AMX via the CPU’s store units [Han23]. They are encoded as A64 instructions within a reserved opcode space [Caw24]; given A64’s fixed 32-bit instruction encoding, only 10 bits remain, 5 of which encode the instruction’s opcode. The remaining 5 bits either encode the index of a scalar 64-bit register through which parameters are passed, or an immediate. This layout is shown in Figure 3. [Caw24] provides preprocessor macros to emit AMX instructions, with an intrinsics-like syntax.



**Figure 3:** AMX instruction encoding.

### 3.3 Instruction set

We now list all known AMX instructions, reviewing those used by our polynomial multiplication implementation, as well as parameters of interest; for an exhaustive specification, see [Caw24]. We propose a taxonomy of instructions in Table 2, categorized by functionality.

**Table 2:** A taxonomy of AMX instructions. Opcodes given in base 10.

Instruction Type	Mnemonics(Opcode)	Description
Loads and stores	<code>ldx(0)</code> , <code>ldy(1)</code> , <code>stx(2)</code> , <code>sty(3)</code> , <code>ldz(4)</code> , <code>stz(5)</code> , <code>ldzi(6)</code> , <code>stzi(7)</code>	Data movement between memory and AMX registers
Extract	<code>extrh(8)</code> , <code>extrx(8)</code> , <code>extrv(9)</code> , <code>extry(9)</code>	Data movement within the AMX register file
First generation matrix/vector	<code>fma64(10)</code> , <code>fms64(11)</code> , <code>fma32(12)</code> , <code>fms32(13)</code> , <code>mac16(14)</code> , <code>fma16(15)</code> , <code>fms16(16)</code>	Outer or pointwise products with accumulation/subtraction
Second generation matrix and vector	<code>vecint(18)</code> , <code>vecfp(19)</code> , <code>matint(20)</code> , <code>matfp(21)</code>	Outer or pointwise products with accumulation/subtraction
Miscellaneous	<code>set(17)</code> , <code>clr(17)</code> , <code>genlut(22)</code>	Context switching, lookup tables

We first consider loads and stores. The least-significant 56 bits of their 64-bit argument are treated as a pointer to the source/target memory address; remaining bits encode parameters as per Table 3. While many other instructions can address X/Y registers at arbitrary byte positions, loads and stores require aligning to the start of a register (64-byte boundary); register indices are encoded in bits 58–56 of the argument (or 61–56 for Z rows). Memory accesses can be 64 bytes (bit 62 = 0) or 128 bytes (bit 62 = 1) wide.

**Table 3:** Parameters for AMX load and store instructions.

Bits	ldx, ldy, stx, sty	ldz, stz
63		
62	Load/store single register (0) or pair of consecutive registers (1)	
61–59		Z-row index (0 to 63)
58–56	Register index (0 to 7)	
55–0	Least significant 56 bits of pointer	



We illustrate the macros of [Caw24] with an example: loading from an array with base address  $v$  to register  $X_4$ . We show a version using just their macros, a second using helper macros defined by us, and a simplified notation used in the algorithms of Section 4:

C notation	C notation, our helper macros	Algorithm notation
<code>AMX_LDX((uint64_t)v   4 &lt;&lt; 56);</code>	<code>AMX_LDX(AMX_PTR(v)   LDX_REG(4));</code>	$X_4 \leftarrow \text{ldx}(v[0 : 31])$

We briefly review functionalities of interest for other instructions. Due to the default role of  $X$  and  $Y$  as input and  $Z$  as output registers, data may need to be moved from  $Z$  to  $X$  or  $Y$ , for which we use the `extrh` instruction. Its basic operation is copying a row of  $Z$  to a byte-addressable 64-byte slice of either  $X$  or  $Y$ . For a visualization of `extrh` in use, refer to Figure 4. Our algorithmic notation for an example operation of extracting  $Z[63]$  to  $Y_0$  is:

$$Y_0 \leftarrow \text{extrh}(Z[63])$$

The main AMX arithmetic instruction we use is `mac16`, which supports *vector* and *matrix* modes. We start by reviewing the vector-mode version. Let  $\mathbf{x}, \mathbf{y}, \mathbf{z}$  be 32-element vectors of 16-bit integers. *Vector-mode* `mac16` computes pointwise multiply-accumulates (MACs):  $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{x} \circ \mathbf{y}$ , where  $+$  is vector addition and  $\circ$  is the Hadamard (pointwise) product. A row of  $Z$  is used for  $\mathbf{z}$ , and  $\mathbf{x}$  and  $\mathbf{y}$  are 32-element slices of  $X$  and  $Y$ . Each of  $\mathbf{x}, \mathbf{y}, \mathbf{z}$  can be optionally “skipped” to realize different operations: vector addition ( $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{x}$  or  $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{y}$ ), Hadamard product without accumulation ( $\mathbf{z} \leftarrow \mathbf{x} \circ \mathbf{y}$ ), or copying  $\mathbf{x}$  or  $\mathbf{y}$  to  $\mathbf{z}$ , i.e. the `extrh` instruction in reverse. We now define an algorithmic notation, assuming for the sake of example the computation of vector addition  $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{x}$  (thus, the  $\mathbf{y}$  input is skipped), in which we choose  $\mathbf{z}$  to be the first row (i.e. that of index 0) of  $Z$  and  $\mathbf{x}$  to be elements 64 to 95 of  $X$  (i.e. bytes 128 to 191, or equivalently, the full  $X_2$  register):

$$Z[0] \leftarrow \text{mac16}(Z[0] + X[64 : 95])$$

We now turn to matrix-mode `mac16`. Let  $\mathbf{Z}$  be a  $32 \times 32$  matrix, and  $\mathbf{x}$  and  $\mathbf{y}$  be 32-element (row) vectors. *Matrix-mode* `mac16` realizes an outer product operation with accumulation, as in Figure 1:  $\mathbf{Z} \leftarrow \mathbf{Z} + \mathbf{y}^T \mathbf{x}$ . Concretely, the operands  $\mathbf{x}, \mathbf{y}, \mathbf{Z}$  are mapped to slices of  $X$  and  $Y$ , and even or odd rows of  $Z$ , respectively. As with vector-mode `mac16`, operands can be skipped, although skipping  $\mathbf{Z}$  appears to be the only sensible possibility, which disables accumulation. It is also possible to update only a subset of output rows and columns, effectively realizing outer products of smaller dimensions. Consider an example operation of computing the outer product  $\mathbf{Y}_0^T \mathbf{X}_0$  (i.e. elements 0 to 31 of each of these registers), which are accumulated with, and saved to, the even rows of  $Z$ , while updating only columns 0 to 15 of the output matrix. Our algorithmic notation for this operation is:

$$Z[0 : 2 : 62] \leftarrow \text{mac16}(Z[0 : 2 : 62] + \mathbf{Y}_0^T \mathbf{X}_0, \text{columns} = 0 : 15)$$

When the “columns” parameter is omitted, it is understood that all columns of the output matrix are updated. Our notation does not distinguish between vector- and matrix-mode `mac16` directly, but rather, the mode is implied from the operands of the instruction. A caveat: although AMX’s  $Y$  registers are best understood as column vectors, our notation assumes vectors (say  $\mathbf{v}$ ) to be row vectors, while column vectors are denoted  $\mathbf{v}^T$ . We extend this convention to AMX’s  $X$  and  $Y$  registers, writing outer products as  $\mathbf{Y}_i^T \mathbf{X}_j$ .

Finally, `vecint` is a vector (pointwise) instruction which generalizes vector-mode `mac16`; in particular, a single instruction can compute  $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{x} + \mathbf{y}$ , attaining twice the throughput as vector-mode `mac16`, a fact which we’ve used to speed up some of our algorithms. As an example, this is how we denote an example operation of 16-bit integer addition of elements 16 to 47 of the  $X$  register (i.e., bytes 32 to 95 of  $X$ , incorporating elements of both  $X_0$  and  $X_1$ ) with elements 0 to 31 of the  $Y$  register (i.e.  $Y_0$ ), accumulating to  $Z[1]$ :

$$Z[1] \leftarrow \text{vecint}(Z[1] + X[16 : 47] + Y[0 : 31])$$

### 3.4 Performance considerations

The M1 has two AMX coprocessors, one per core cluster (performance and efficiency), implementing a second generation of AMX instructions [Caw24]. It runs at the same clock speed as the cores and accesses memory via the L2 cache [Joh22b]. It supports out-of-order execution independently of the CPU, with buffers for 28 to 32 operations [Joh22b].

Outer product latency and throughput is discussed in [Joh22b] for floating-point operations. The performance cluster’s AMX unit has an array of pipelined MAC units with 4-cycle latency, capable of executing one 32- or 64-bit outer product operation per cycle, or a 16-bit operation every two cycles. To achieve maximum performance, inter-instruction data dependencies must be avoided; either by using multiple *accumulators*, i.e. different subsets of  $Z$  as a destination (such as even and odd rows), or by issuing instructions from different CPUs, due to independent per-CPU AMX register files (see Section 3.1).

AMX microbenchmarking code is supplied in [Caw24]. We have run this code on an M1-based laptop, and report selected throughput figures, restricted to single-threaded results, presumably running on the AMX performance-cluster unit. Multiply-accumulates (MACs) count as 2 operations, and we report best-case scenarios, often using multiple accumulators; performance may degrade, significantly in certain cases, if fewer accumulators are used.

Outer products achieve up to 3.05 TFLOPS/s throughput with 16-bit floating-point or 8-bit integer arithmetic; throughput drops to 1.53 TOPS/s for 16-bit integer arithmetic. In vector mode, up to 381 GFLOPS/s is achievable with floating-point or 8-bit integer MACs, dropping to 191 GOPS/s for 16-bit integer arithmetic. Replacing MACs with additions or multiplications alone further halves throughput (due to performing half as many operations per instruction). As a comparison, peak NEON throughput on M1 performance cores for 16-bit integer data is 204.8 GOPS/s for MACs and 102.4 GOPS/s for additions or multiplications alone, based on the microarchitectural investigations of [Joh22a].

Thus, AMX’s peak throughput is  $\approx 8\times$  higher for matrix (outer product) operations versus vector operations; the latter’s peak performance approaches NEON’s. That is, *outer products are much cheaper than vector operations*. Throughout the paper, we consistently strive to reduce vector operation count, which is the bottleneck of our computations.

While our implementation takes the usual constant-time precautions (foregoing branches and memory accesses conditional/indexed on secret data), it implicitly relies on constant-time execution of AMX instructions. Recently, ARM and Intel have started guaranteeing data-independent timing for some instructions [ARM, Int23b, Int23a]. Given AMX’s lack of official documentation, it is impossible to obtain such assurances, although experiments in Section 5.1 suggest AMX displays data-independent timing. The previous lack of constant-time guarantees for CPUs hasn’t prevented implementers from striving for, and achieving in practice, implementations with apparent constant-time characteristics, and may have played a key part in incentivizing manufacturers to provide such guarantees. Our paper is a first step in that direction for CPU-coupled matrix-multiplication accelerators.

## 4 Implementing polynomial multiplication on AMX

We now describe our AMX implementation of schoolbook multiplication in  $\mathbb{Z}_{2^{16}}/(x^n - 1)$ .

### 4.1 Challenges and strategies for efficient implementation

We start by summarizing the strategies that were evaluated for efficient implementation of polynomial multiplication on AMX, why we discarded some of them, and how we evolved and optimized our implementation to eventually achieve new speed records on Apple SoCs.

Many techniques for efficient polynomial multiplication in ARMv8 are based on vector (1D) operations, a natural match for the NEON SIMD instruction set. However, as stated in Section 3.4, vector operations achieve a fraction of AMX’s peak throughput,



comparable to or even slower than the CPU’s NEON units. There is little hope for AMX implementations to outperform NEON if restricted to vector operations; matrix (2D) operations must be used to approach AMX’s peak theoretical performance.

The Toeplitz Matrix-Vector Product (TMVP) algorithm [CCHY23] is the current state-of-the-art in NEON performance. By casting polynomial multiplication in the language of linear algebra, it may appear perfectly suited to matrix accelerators such as AMX. However, its main operation is a matrix-vector product. AMX’s outer product instruction computes matrix-matrix products very efficiently; however, the cost of a  $32 \times 32$  by  $32 \times 1$  matrix-vector product (the natural AMX block size for 16-bit integers) using this instruction is essentially identical to that of a  $32 \times 32$  by  $32 \times 32$  matrix-matrix product. We thoroughly investigated other approaches, but found no strategies that outperform NEON for matrix-vector products. Further evidence comes from our benchmarks of (floating-point) matrix-vector products using Apple’s BLAS library, whose performance does not exceed the peak theoretical throughput of NEON. As such, we also disregard the use of TMVP.

Our key observation leading to efficient polynomial multiplication in AMX was the similarity between outer products and operand-scanning multiplication, resulting in the algorithm of Section 4.2 for multiplying 32-coefficient slices of polynomials, using the  $32 \times 32$  natural AMX block size. However, for each outer product operation, dozens of vector operations are used to shift partial products and perform sum-reduction to arrive at the polynomial product; as the throughput of matrix AMX operations is  $\approx 8\times$  higher than vector operations (see Section 3.4), it is imperative to increase the balance of matrix to vector operations to approach the peak AMX theoretical throughput. Extrapolating the performance of a basic block ( $\geq 90$  clock cycles) to polynomial degrees used in NTRU shows that this technique alone is not competitive with state-of-the-art NEON implementations.

Next, we shifted to a size-32 blocked view of polynomial multiplication. We noticed that a blockwise product scanning approach allows computing and accumulating all partial products in each block column, deferring all vector operations for shifting and sum-reduction of that block column to the end. This reduces vector operation count asymptotically from  $O(n^2)$  to  $O(n)$ , improving matrix/vector instruction balance and thus AMX throughput. This strategy and the resulting algorithms are described in Section 4.3.

The product of degree- $n$  polynomials, as in the previous paragraph, results in a degree- $2n$  polynomial; however, in NTRU and other schemes, a reduction is applied to obtain a degree- $n$  output. In Section 4.4, we show how to perform reduction concurrently, in an integrated fashion, with blockwise product scanning; this reduces, by approximately half, the number of vector operations for shifts and sum-reduction compared to the strategy of the previous paragraph, again improving matrix/vector instruction balance. The importance of this optimization is highlighted by the results of Table 4, comparing its performance to polynomial multiplication without reduction (Section 4.3) – for NTRU, of course, polynomial reduction must be implemented, further widening the performance gap.

**Table 4:** Cycle counts for AMX polynomial multiplication without reduction (Section 4.3) and polynomial multiplication with integrated reduction (Section 4.4).

Parameter set	Section 4.3		Section 4.4		Speedup ( $\times$ )	
	M1	M3	M1	M3	M1	M3
HPS2048509	3534	2963	2365	2117	<b>1.49</b>	<b>1.40</b>
HPS2048677	5430	4694	3790	3417	<b>1.43</b>	<b>1.37</b>
HPS4096821	6854	6007	4870	4497	<b>1.41</b>	<b>1.34</b>
HRSS701	5431	4691	3769	3456	<b>1.44</b>	<b>1.36</b>

While the previous paragraph’s strategy achieves competitive and even record-setting

performance, we noted a claim in [Han23] of performance degradation from concurrent accesses by AMX and the CPU to the same memory page. In Section 4.5, we investigate allocation of input arrays in separate memory pages via the POSIX `mmap()` function. Performance results are summarized in Table 5, for both AMX and NEON, from geometric means of measurements presented later in Table 6. Briefly, per-page memory allocation benefits AMX considerably (M1 more so than M3), while NEON displays mixed results regardless of SoC. This also suggests that concurrent processing of the same dataset with NEON and AMX may, counterintuitively, have detrimental performance effects.

**Table 5:** Geometric mean of speedups ( $\times$ ) for `mmap()` compared to stack memory allocation.

Instruction set	Keygen		Encaps		Decaps		Poly. mul.	
	M1	M3	M1	M3	M1	M3	M1	M3
NEON	1.00	1.00	1.02	1.01	1.03	1.02	0.96	0.98
AMX	1.19	1.05	1.38	1.08	1.33	1.33	1.34	1.15

All implementations of Section 4 are based on the  $O(n^2)$  schoolbook method. In Section 5.1, we investigate the performance of subquadratic algorithms, showing experimental evidence that the Karatsuba algorithm with a single recursion level does not outperform schoolbook in AMX until well outside the range of polynomial degrees of cryptographic interest. We refer to that section for hypotheses to explain this somewhat counterintuitive result, and evidence that other strategies based on subquadratic algorithms are also not expected to outperform schoolbook in AMX, for the polynomial degrees used in NTRU.

## 4.2 Basic block: multiplication of 32-coefficient slices

The basic block of our implementation multiplies slices of 32 coefficients from each input polynomial. We define a specific notation for polynomial slices starting at indices that are a multiple of 32:  $a_{32k:32k+31}(x) = \mathbf{a}_{\mathbf{k}}(x)$  (note boldface  $\mathbf{k}$ ). We also define a notation for the product of 32-coefficient polynomial slices:  $\mathbf{c}_{\mathbf{k},\mathbf{l}}(x) = \mathbf{a}_{\mathbf{k}}(x)\mathbf{b}_{\mathbf{l}}(x)$ . To clarify, we have:

$$\begin{aligned} \mathbf{c}_{\mathbf{k},\mathbf{l}}(x) &= (a_{32k}x^{32k} + \dots + a_{32k+31}x^{32k+31})(b_{32l}x^{32l} + \dots + b_{32l+31}x^{32l+31}) \\ &= c_{32(k+l)}x^{32(k+l)} + \dots + c_{32(k+l)+62}x^{32(k+l)+62}, \end{aligned}$$

where  $c_m = \sum_{i+j=m} a_i b_j$ . As before,  $\mathbf{c}_{\mathbf{k},\mathbf{l}}$  is the associated row vector representation. Throughout the rest of Section 4, we resort to examples using a hypothetical 1/8-size AMX unit for space reasons. Boldface indices will then refer to 4-coefficient slices as appropriate for these examples, rather than 32-coefficient slices for full AMX; thus,  $\mathbf{a}_{\mathbf{k}}$  would refer to  $\mathbf{a}_{4k:4k+3}$ . It will be clear from the context whether we refer to 4- or 32-coefficient slices.

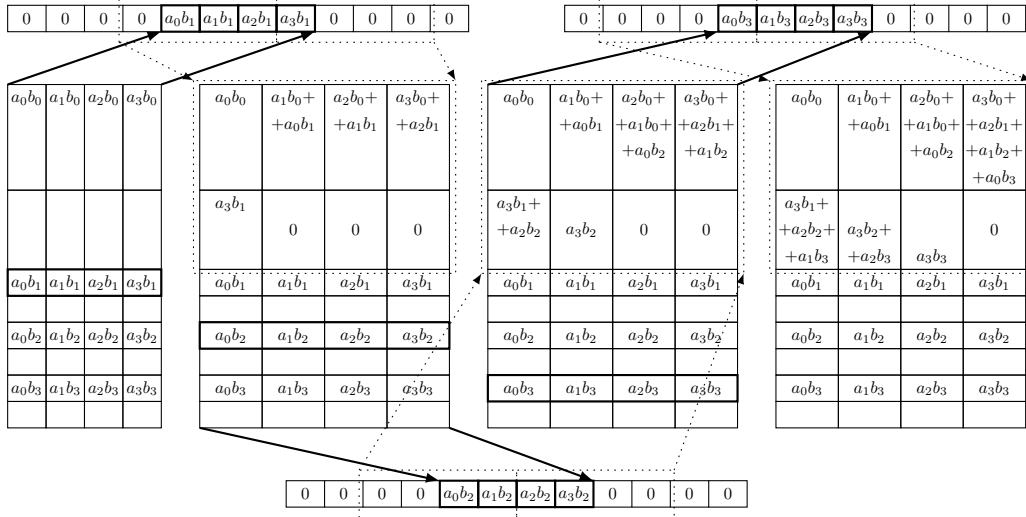
We first note that the outer product  $\mathbf{b}_{\mathbf{l}}^T \mathbf{a}_{\mathbf{k}}$  generates the required partial products; shifting each row produces the usual multiplication parallelogram. We then perform sum-reduction of partial products in each column (an operation which we refer throughout the rest of Section 4 as *flattening* to avoid confusion with polynomial reduction) to obtain  $\mathbf{c}_{\mathbf{k},\mathbf{l}}$ . This can be visualized by considering an analogous operation in a hypothetical 1/8-size AMX unit, operating on 4-coefficient slices starting at the constant coefficient:

$b_0$	$a_3 b_0$	$a_2 b_0$	$a_1 b_0$	$a_0 b_0$	Left shift $i$ -th row by $i$ positions $\rightarrow$					
$b_1$	$a_3 b_1$	$a_2 b_1$	$a_1 b_1$	$a_0 b_1$						
$b_2$	$a_3 b_2$	$a_2 b_2$	$a_1 b_2$	$a_0 b_2$						
$b_3$	$a_3 b_3$	$a_2 b_3$	$a_1 b_3$	$a_0 b_3$						
							$a_3 b_3$	$a_2 b_3$	$a_1 b_3$	$a_0 b_3$

We then flatten the parallelogram by sum-reduction of columns and multiplication by the corresponding power of  $x$  to obtain  $a_3b_3x^6 + (a_2b_3 + a_3b_2)x^5 + \dots + (a_0b_1 + a_1b_0)x + a_0b_0$ .

On AMX, outer products are easily realized by loading  $\mathbf{a}_k$  and  $\mathbf{b}_l$  to slices of the  $\mathbf{X}$  or  $\mathbf{Y}$  registers, and executing the matrix-mode `mac16` instruction. By default, results are accumulated with the current values of  $\mathbf{Z}$ , although there is a flag to skip current contents of  $\mathbf{Z}$ , so as to realize multiplication without accumulation. The result occupies half of the  $\mathbf{Z}$  register file, either the odd or even rows of  $\mathbf{Z}$ ; we assume the latter for the sake of example.

To realize shifts and flattenings, we initialize the two ends of a 3-register-wide slice of  $\mathbf{X}$  (say  $\mathbf{X}_0$  and  $\mathbf{X}_2$ ) as well as  $\mathbf{Z}[1]$ , the first odd row of  $\mathbf{Z}$ , with zeros. For each even row of  $\mathbf{Z}$  except the first, we use the `extrh` instruction to extract it to the remaining  $\mathbf{X}$  register of the slice ( $\mathbf{X}_1$  in this example). Recall that AMX allows addressing any 64-byte slice of the full 512-byte  $\mathbf{X}$  register; we use this feature to select slices containing the appropriate number of zero 16-bit values to realize the desired shifting. These slices are used as input to a pair of `vecint` or vector-mode `mac16` instructions, performing additions only, in order to accumulate the least significant coefficients to  $\mathbf{Z}[0]$ , and the most significant coefficients to  $\mathbf{Z}[1]$ . Figure 4 illustrates this operation, again on a hypothetical 1/8-size AMX unit.



**Figure 4:** Full multiplication process of 4-coefficient polynomial slices on a hypothetical 1/8-size AMX unit. The sequence of operations starts from the leftmost state of the  $\mathbf{Z}$  registers (result of the outer product operation) and follows the arrows, representing different states of  $\mathbf{Z}$  and the relevant part of  $\mathbf{X}$ . Thick arrows indicate an `extrh` operation of the thick-bordered row of  $\mathbf{Z}$  to  $\mathbf{X}$ , while dotted arrows indicate an accumulation of each dotted slice of  $\mathbf{X}$  with a corresponding row (0 or 1) of the target  $\mathbf{Z}$  registers.

Performance is further improved by working on pairs of even rows of  $\mathbf{Z}$ , extracting one to  $\mathbf{X}$  and the other to  $\mathbf{Y}$ . As discussed in Section 3.3, a single `vecint` can accumulate slices of both  $\mathbf{X}$  and  $\mathbf{Y}$  registers to the desired  $\mathbf{Z}$  row. Microbenchmarks indicate that `vecint`, on its own, achieves twice the arithmetic throughput of addition with  $\mathbf{X}$  or  $\mathbf{Y}$  alone; however, this improvement is not fully realized as there appears to be contention with resources used to execute `extrh`, and moreover, microbenchmarks suggest there is a penalty for unaligned accesses to the  $\mathbf{X}$  and  $\mathbf{Y}$  registers, as required to shift the input operands.

Algorithm 4.1 formalizes this procedure for polynomial multiplication of 32-coefficient slices. The reader may wish to review notations in Sections 2 and 3, in particular Figure 2.

---

**Algorithm 4.1** POLYMUL $32 \times 32(\mathbf{c}_{k,1}, \mathbf{a}_k, \mathbf{b}_j)$ : modulo- $2^{16}$  multiplication of two 32-coefficient polynomials using AMX.

---

**Input:**  $\mathbf{a}_k, \mathbf{b}_j$  (arrays of 32-coefficient slices of  $\mathbf{a}$  and  $\mathbf{b}$ )

**Output:**  $\mathbf{c}_{k,1}$  (array of 63 coefficients of  $c_{k,1}(x) = a_k(x) \times b_j(x)$ )

**Notes:** since AMX stores 64 bytes (in this case, 32 coefficients of 16 bits each) at a time, output array must be allocated with enough space for 64 coefficients

```

1:  $\mathbf{X}_0, \mathbf{X}_2, \mathbf{Y}_0, \mathbf{Y}_2 \leftarrow \text{ldx}([0 \ 0 \ \dots \ 0])$   $\triangleright$  load array of 32 zeros to each indicated register
2:  $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}_k), \text{ldy}(\mathbf{b}_j)$ 
3:  $\mathbf{Z}[0 : 2 : 62] \leftarrow \text{mac16}(\mathbf{Y}_1^\top \mathbf{X}_1)$ 
4:  $\mathbf{X}_1 \leftarrow \text{extrh}(\mathbf{Z}[2])$ 
5:  $\mathbf{Z}[0] \leftarrow \text{vecint}(\mathbf{Z}[0] + \mathbf{X}[31 : 62])$ 
6:  $\mathbf{Z}[1] \leftarrow \mathbf{X}[63 : 94]$   $\triangleright$  implement using e.g. mac16 with  $\mathbf{Y}$  and  $\mathbf{Z}$  skipped
7: for  $i \in \{2, 4, 6, \dots, 30\}$  do
8:    $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{extrh}(\mathbf{Z}[2i]), \text{extrh}(\mathbf{Z}[2(i+1)])$ 
9:    $\mathbf{Z}[0] \leftarrow \text{vecint}(\mathbf{Z}[0] + \mathbf{X}[32-i : 63-i] + \mathbf{Y}[31-i : 62-i])$ 
10:   $\mathbf{Z}[1] \leftarrow \text{vecint}(\mathbf{Z}[1] + \mathbf{X}[64-i : 95-i] + \mathbf{Y}[63-i : 94-i])$ 
11:  $\mathbf{c}_{k,1}[0 : 31], \mathbf{c}_{k,1}[32 : 63] \leftarrow \text{stz}(\mathbf{Z}[0]), \text{stz}(\mathbf{Z}[1])$ 

```

---

### 4.3 Polynomial multiplication by product-scanning of basic blocks

We turn to the issue of realizing a full multiplication of  $n$ -coefficient inputs. Since NTRU requires  $n$  to be prime, and the basic block of Section 4.2 works on 32-coefficient slices, we work around this issue by zero-padding input polynomials to  $n' = 32 \lceil n/32 \rceil$  coefficients.

In principle, we could apply Algorithm 4.1 to each slice and accumulate results using vector instructions to realize a full-size multiplication. However, we propose a different architecture, based on blockwise product scanning with lazy sum-reduction, to increase the balance of faster matrix (outer product) to slower vector (extraction/addition) operations. We again illustrate with a small example for a hypothetical 1/8-size AMX unit, multiplying 8-coefficient input polynomials split into 4-coefficient slices as in Section 4.2.

The parallelogram of Figure 5 consists of 4 outer products:  $\mathbf{b}_0^\top \mathbf{a}_0$ ,  $\mathbf{b}_1^\top \mathbf{a}_0$ ,  $\mathbf{b}_1^\top \mathbf{a}_0$  and  $\mathbf{b}_1^\top \mathbf{a}_1$ . It is seen in Figure 5(b) that  $\mathbf{b}_1^\top \mathbf{a}_0$  and  $\mathbf{b}_0^\top \mathbf{a}_1$  are aligned columnwise and can be accumulated (by matrix addition) prior to shifting and flattening; that is, shifting and flattening are performed lazily. In general, outer products of the form  $\mathbf{b}_j^\top \mathbf{a}_i$  and  $\mathbf{b}_l^\top \mathbf{a}_k$  are aligned, and can be accumulated, if  $i + j = k + l$ . Note that matrix addition does not need to be performed explicitly in AMX, as the outer product instructions accumulate with the existing contents of the  $\mathbf{Z}$  registers at no extra cost. The accumulation procedure is shown as Algorithm 4.2, a subroutine of our polynomial multiplication algorithm given later.

---

**Algorithm 4.2** ACCUMULATEOUTERPRODUCTS( $\mathbf{a}, \mathbf{b}, j, r$ ): accumulate outer products  $\mathbf{b}_l^\top \mathbf{a}_k$  such that  $k + l = j$  to  $\mathbf{Z}[r : 2 : 62 + r]$ , where  $r = 0$  or 1.

---

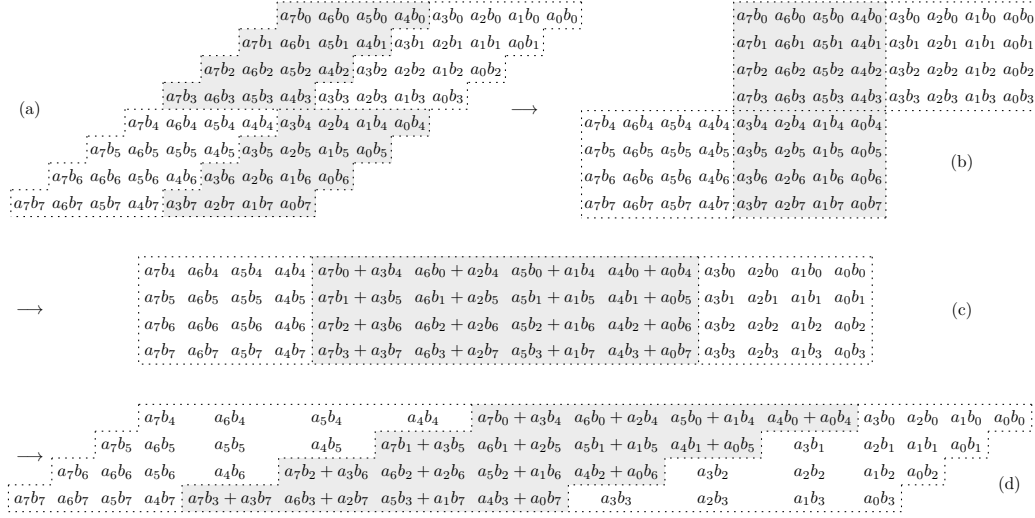
```

1:  $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}[32j : 32j + 31]), \text{ldy}(\mathbf{b}[0 : 31])$ 
2:  $\mathbf{Z}[r : 2 : 62 + r] \leftarrow \text{mac16}(\mathbf{Y}_1^\top \mathbf{X}_1)$ 
3: for  $1 \leq l \leq j$  do
4:    $k \leftarrow j - l$   $\triangleright$  Thus:  $k + l = j$ 
5:    $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}[32k : 32k + 31]), \text{ldy}(\mathbf{b}[32l : 32l + 31])$ 
6:    $\mathbf{Z}[r : 2 : 62 + r] \leftarrow \text{mac16}(\mathbf{Z}[r : 2 : 62 + r] + \mathbf{Y}_1^\top \mathbf{X}_1)$ 

```

---

We now consider a full-size AMX unit and realistic polynomial sizes, say  $n'$  coefficients. The naïve approach suggested at the beginning of this section would call for  $(n'/32)^2 = O(n^2)$  applications of Algorithm 4.1, that is  $O(n^2)$  outer products and shifts/flattening, plus extra operations to combine the results to obtain the desired polynomial multiplication



**Figure 5:** Transformations of the multiplication parallelogram for 8-coefficient polynomials for efficient implementation on a hypothetical 1/8-size AMX unit. Dotted lines enclose each outer product of 4-coefficient slices. Starting from the regular parallelogram (a), the  $i$ -th row is shifted right by  $i \bmod 4$  positions (i.e. the reverse of the shifting step of Figure 4) to obtain (b), revealing how two outer products are aligned. In (c), these two outer products are summed prior to left-shifting the  $i$ -th row by  $i$  positions in the last state shown, (d). The final result is obtained by flattening each column (not shown).

result. The advantage of lazy shifting/flattening is clear from extrapolation of Figure 5: we see that it results in  $2(n'/32) - 1 = O(n)$  matrices of dimension  $32 \times 32$  in the third step of the process (Figure 5(c)), which are shifted (Figure 5(d)) and flattened (not shown) to obtain the final polynomial multiplication result. While we still compute the same amount of outer products,  $(n'/32)^2 = O(n^2)$ , they are considerably cheaper than shifts and flattenings realized by AMX vector instructions, of which only  $O(n)$  are required.

Moving on to practical implementation issues, the Z register, where outer products accumulate, can only store two  $32 \times 32$  matrices, one in the even rows and another in the odd rows. Careful sequencing of shift and flattening operations is needed to avoid spills and reloads of Z's contents. Consider the flattened parallelogram in Figure 5(d), which is split into three sub-parallelograms, each corresponding to one of three matrices of Figure 5(c), formed from accumulating outer products; we denote them, from right to left, as  $M_0$ ,  $M_1$  and  $M_2$ , respectively. Working through the final flattening step (i.e. sum-reduction of columns, not shown in the figure), from the rightmost column towards the leftmost one, we see that the first columns (0 to 3) contain only elements from  $M_0$ . As we move left, columns 4 to 6 contain elements from  $M_0$  and  $M_1$ ; column 7 from  $M_1$  only; columns 8 to 10 from  $M_1$  and  $M_2$ ; and columns 11 to 14 from  $M_2$  only. A general pattern emerges: each column contains either elements from  $M_j$  only, or from  $M_j$  and  $M_{j+1}$ .

This suggests a structure for our polynomial multiplication algorithm, described in terms of full-size ( $n'$  coefficients) input polynomials. Formalizing the concept from the previous paragraph, we define  $M_j = \sum_{k+l=j} \mathbf{b}_l^T \mathbf{a}_k$ ; the computation of each  $M_j$  is performed by Algorithm 4.2. We first compute  $M_0$  and  $M_1$ , storing them in the even and odd rows of Z, respectively, as needed to perform shifts and flattenings to obtain the first 64 coefficients  $c_0, \dots, c_{63}$  of  $c(x) = a(x)b(x)$ , a procedure we formalize as Algorithm 4.3. We now compute  $M_2$  overwriting the even rows of Z, and use both  $M_1$  (still available in the odd rows of Z) and  $M_2$  to obtain the next 32 coefficients of  $c(x)$ , i.e.  $c_{64}, \dots, c_{95}$ , by the procedure of

Algorithm 4.4. The pattern continues for  $j = 3, \dots, 2(n/32) - 3$  with  $M_j$  overwriting even or odd rows of  $Z$ , according to whether  $j$  is even or odd, respectively, and using  $M_{j-1}$  and  $M_j$  to compute  $c_{32j}, \dots, c_{32j+31}$  using Algorithm 4.4. Finally, the case  $j = 2(n'/32) - 2$  is handled by a third procedure, Algorithm 4.5, which computes the final batch of coefficients of  $c(x)$  from  $M_{j-1}$  and  $M_j$ . We formalize the complete procedure as Algorithm 4.6.

---

**Algorithm 4.3** FLATTENFIRSTTWOBLOCKS( $\mathbf{c}$ ): sum-reduction of columns from the two least significant multiplication sub-parallelgrams.

---

**Input:**  $Z[0 : 2 : 62] = \mathbf{b}_0^T \mathbf{a}_0$  and  $Z[1 : 2 : 63] = \mathbf{b}_0^T \mathbf{a}_1 + \mathbf{b}_1^T \mathbf{a}_0$  (implicitly)

**Output:**  $\mathbf{c}[0 : 63]$  (coefficients of  $x^0$  through  $x^{63}$  of the polynomial multiplication result)

**Notes:** assumes  $\mathbf{X}_0 = \mathbf{Y}_0 = [0, 0, \dots, 0]$

```

1:  $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(Z[2]), \text{extrh}(Z[3])$ 
2:  $Z[0], Z[1] \leftarrow \text{mac16}(Z[0] + \mathbf{X}[31 : 62]), \text{mac16}(Z[1] + \mathbf{X}[63 : 94])$ 
3: for  $i \in \{2, 4, 6, \dots, 30\}$  do
4:    $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(Z[2i]), \text{extrh}(Z[2i + 1])$ 
5:    $\mathbf{Y}_1, \mathbf{Y}_2 \leftarrow \text{extrh}(Z[2i + 2]), \text{extrh}(Z[2i + 3])$ 
6:    $Z[0] \leftarrow \text{vecint}(Z[0] + \mathbf{X}[32 - i : 63 - i] + \mathbf{Y}[31 - i : 62 - i])$ 
7:    $Z[1] \leftarrow \text{vecint}(Z[1] + \mathbf{X}[64 - i : 95 - i] + \mathbf{Y}[63 - i : 94 - i])$ 
8:  $\mathbf{c}[0 : 31], \mathbf{c}[32 : 63] \leftarrow \text{stz}(Z[0]), \text{stz}(Z[1])$ 
```

---

#### 4.4 Integrated reduction modulo $x^n - 1$

The algorithms of Section 4.3, culminating in Algorithm 4.6, multiply polynomials of  $n \times n$  coefficients and return the full result with  $2n - 1$  coefficients. In NTRU, the result must be reduced modulo  $x^n - 1$ , a special form which allows efficient implementation as a post-processing step using e.g. NEON instructions, but also allowing reduction to be merged with multiplication in an integrated procedure, reducing the number of shifting/flattening operations by  $\approx 50\%$  compared to Algorithm 4.6, improving performance considerably.

To exemplify the integrated procedure, we again resort to a hypothetical 1/8-size AMX unit, now performing polynomial multiplication modulo  $x^{10} - 1$ , which illustrates the main issues faced with larger polynomials on a full-size AMX unit. Although 10 coefficients are enough to represent polynomial inputs and outputs for this case, note that each block in this hypothetical 1/8-size AMX unit performs a  $4 \times 4$  outer-product operation. Thus, we work with the next multiple of 4 coefficients, which is 12. We refer to Figure 6 as we work through the procedure, and assume input polynomials have  $a_i = b_j = 0$  for  $i, j \geq 10$ .

The results of all required outer product calculations are shown explicitly in Figure 6(a). The actual procedure accumulates  $4 \times 4$  outer product results of a given color (white, light gray or gray) into a single per-color  $4 \times 4$  matrix, as in Figure 5(c); for didactic reasons, we omit this step in the course of the explanation. Zero partial products are struck out in the figure; for some of these, this is a natural result of having  $a_i$  or  $b_j$  as input for  $i, j \geq 10$ , and for the others, we explicitly disable some columns during outer product computations.

Viewing Figure 6(a) as a block matrix of  $4 \times 4$  submatrices, note that submatrices in and above the secondary (block) diagonal are constructed identically to the example of Figure 5, as formalized in Algorithm 4.2. Notice the pattern followed by partial products  $a_i b_j$  in each row: those to its left and right are of the form  $a_{i+1} b_j$  and  $a_{i-1} b_j$ , respectively. The rightmost columns of each submatrix in the secondary (block) diagonal have  $i = 0$ , and the pattern continues to their right, with  $i - 1$  taken modulo  $n$  (10 in the example). For each partial product  $a_i b_j$ , with  $k$  its row index modulo 4 and  $l$  its column index, as indicated to the left and below the matrix in Figure 6(a), we have  $k + l \equiv i + j \pmod{n}$ .

Submatrices in and above the secondary (block) diagonal are outer products of the form  $\mathbf{b}_{4l:4l+3}^T \mathbf{a}_{4k:4k+3}$  (no relation to  $k$  and  $l$  in the previous paragraph), while submatrices



---

**Algorithm 4.4** FLATTENMIDDLEBLOCK( $\mathbf{c}, j, r$ ): sum-reduction of columns from multiplication sub-parallelgrams.

---

**Input:**  $j$  (index of current sub-parallelgram)

**Input:**  $r \in \{0, 1\}$  (indicates relative order of even and odd rows of  $\mathbf{Z}$ )

**Output:**  $\mathbf{c}[32j : 32j + 31]$  (coefficients of  $x^{32j}$  through  $x^{32j+31}$  of the polynomial multiplication result)

**Notes:** assumes that  $\mathbf{Z}[1 - r : 2 : 63 - r]$  and  $\mathbf{Z}[r : 2 : 62 + r]$  contain the output of Algorithm 4.2 for sub-parallelgram indices  $j - 1$  and  $j$ , respectively

- 1:  $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(\mathbf{Z}[3 - r]), \text{extrh}(\mathbf{Z}[2 + r])$
  - 2:  $\mathbf{Z}[r] \leftarrow \text{mac16}(\mathbf{Z}[r] + \mathbf{X}[63 : 94])$
  - 3: **for**  $i \in \{2, 4, 6, \dots, 30\}$  **do**
  - 4:    $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(\mathbf{Z}[2i + 1 - r]), \text{extrh}(\mathbf{Z}[2i + r])$
  - 5:    $\mathbf{Y}_1, \mathbf{Y}_2 \leftarrow \text{extrh}(\mathbf{Z}[2i + 3 - r]), \text{extrh}(\mathbf{Z}[2i + 2 + r])$
  - 6:    $\mathbf{Z}[r] \leftarrow \text{vecint}(\mathbf{Z}[r] + \mathbf{X}[64 - i : 95 - i] + \mathbf{Y}[63 - i : 94 - i])$
  - 7:  $\mathbf{c}[32j : 32j + 31] \leftarrow \text{stz}(\mathbf{Z}[r])$
- 

---

**Algorithm 4.5** FLATTENLASTTWOBLOCKS( $\mathbf{c}, j$ ): sum-reduction of columns from the two most significant multiplication sub-parallelgrams.

---

**Input:**  $j$  (index of current block)

**Output:**  $\mathbf{c}[32j : 32j + 62]$  (coefficients of  $x^{32j}$  through  $x^{32j+62}$  of  $c(x) = a(x) \times b(x)$ )

**Notes:** assumes  $\mathbf{X}_3 = \mathbf{Y}_3 = [0, 0, \dots, 0]$ , and that  $\mathbf{Z}[1 : 2 : 63]$  and  $\mathbf{Z}[0 : 2 : 62]$  contain the output of Algorithm 4.2 for sub-parallelgram indices  $j - 1$  and  $j$ , respectively; one extra coefficient written at  $\mathbf{c}[32j + 63]$  since AMX writes 64 bytes at a time

- 1:  $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(\mathbf{Z}[3]), \text{extrh}(\mathbf{Z}[2])$
  - 2:  $\mathbf{Z}[0], \mathbf{Z}[1] \leftarrow \text{mac16}(\mathbf{Z}[0] + \mathbf{X}[63 : 94]), \text{mac16}(\mathbf{Z}[1] + \mathbf{X}[95 : 126])$
  - 3: **for**  $i \in \{2, 4, 6, \dots, 30\}$  **do**
  - 4:    $\mathbf{X}_2, \mathbf{X}_1 \leftarrow \text{extrh}(\mathbf{Z}[2i]), \text{extrh}(\mathbf{Z}[2i + 1])$
  - 5:    $\mathbf{Y}_2, \mathbf{Y}_1 \leftarrow \text{extrh}(\mathbf{Z}[2i + 2]), \text{extrh}(\mathbf{Z}[2i + 3])$
  - 6:    $\mathbf{Z}[0] \leftarrow \text{vecint}(\mathbf{Z}[0] + \mathbf{X}[32 - i : 63 - i] + \mathbf{Y}[31 - i : 62 - i])$
  - 7:    $\mathbf{Z}[1] \leftarrow \text{vecint}(\mathbf{Z}[1] + \mathbf{X}[64 - i : 95 - i] + \mathbf{Y}[63 - i : 94 - i])$
  - 8:  $\mathbf{c}[32j : 32j + 31], \mathbf{c}[32j + 32 : 32j + 63] \leftarrow \text{stz}(\mathbf{Z}[0]), \text{stz}(\mathbf{Z}[1])$
- 

---

**Algorithm 4.6** POLYMUL( $\mathbf{c}, \mathbf{a}, \mathbf{b}$ ): modulo- $2^{16}$  multiplication of two  $n$ -coefficient polynomials using AMX. Assumes  $n$  a multiple of 32; input must be zero-padded otherwise.

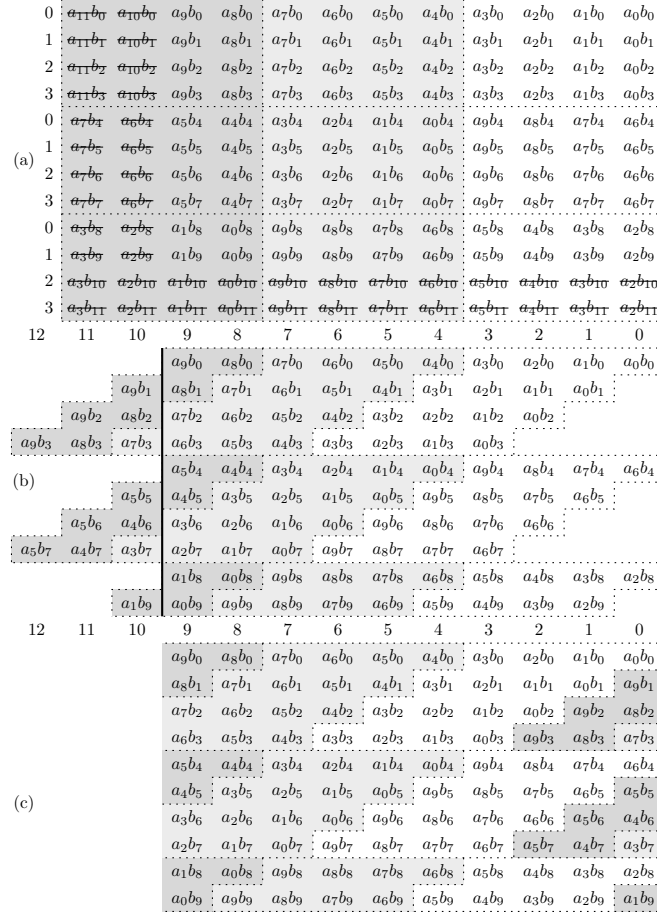
---

**Input:**  $\mathbf{a}, \mathbf{b}$  (arrays of  $n$  coefficients each of  $a(x)$  and  $b(x)$ )

**Output:**  $\mathbf{c}$  (array of  $2n - 1$  coefficients of  $c(x) = a(x)b(x)$ )

**Notes:** since AMX stores 64 bytes (in this case, 32 coefficients of 16 bits each) at a time, output array must be allocated with  $2n$  coefficients

- 1:  $\mathbf{X}_0, \mathbf{X}_3, \mathbf{Y}_0, \mathbf{Y}_3 \leftarrow \text{ldx}([0, 0, \dots, 0])$   $\triangleright$  load array of 32 zeros to each indicated register
  - 2: ACCUMULATEOUTERPRODUCTS( $\mathbf{a}, \mathbf{b}, 0, 0$ )
  - 3: ACCUMULATEOUTERPRODUCTS( $\mathbf{a}, \mathbf{b}, 1, 1$ )
  - 4: FLATTENFIRSTTWOBLOCKS( $\mathbf{c}$ )
  - 5: **for**  $i = 2$  **to**  $2(n/32) - 3$  **do**
  - 6:   ACCUMULATEOUTERPRODUCTS( $\mathbf{a}, \mathbf{b}, i, i \bmod 2$ )
  - 7:   FLATTENMIDDLEBLOCK( $\mathbf{c}, i, i \bmod 2$ )
  - 8: ACCUMULATEOUTERPRODUCTS( $\mathbf{a}, \mathbf{b}, 2(n/32) - 2, 0$ )
  - 9: FLATTENLASTTWOBLOCKS( $\mathbf{c}, 2(n/32) - 2$ )
-



**Figure 6:** Sequence of operations to implement integrated polynomial multiplication and reduction of 12-coefficient polynomials modulo  $x^{10} - 1$  on a hypothetical 1/8-size AMX unit. See text for a description of transformations from each subfigure to the next.

below this diagonal are obtained from outer products  $\mathbf{b}_{4l:4l+3}^\top \mathbf{a}_{4k+2:4k+5}$ ; note how the coefficients of  $\mathbf{a}$  are “unaligned”, i.e. the starting index is not a multiple of 4 (32 for a full-size AMX unit). The offset 2 in the example corresponds in general to  $n \bmod 4$ , due to the modulo  $x^n - 1$  reduction when  $n$  is not a multiple of 4 (32 for a full-size AMX unit).

We now present Algorithm 4.7, a counterpart of Algorithm 4.2 for the integrated polynomial multiplication and reduction procedure. It follows Algorithm 4.2 in accumulating outer products with matching columns (i.e. the submatrices of the same color in Figure 6(a)), a step which we recall was omitted from the example for ease of exposition.

We transform Figure 6(a) into Figure 6(b) by removing struck-out partial products and shifting rows of Figure 6(a) by the row index modulo 4; indices  $i, j$  of partial products  $a_i b_j$  in each column add exactly to the column index modulo 10. As we are working modulo  $x^{10} - 1$ , the final result must be restricted to indices from 0 to 9; we mark this boundary with a thick vertical line in the figure. Some partial products remain in columns with index  $k \geq 10$ , i.e. to the left of this line. We move them to the column with index  $k' = k \bmod 10$  in the same row to arrive at Figure 6(c); note that  $i + j \equiv k \pmod{10}$  always holds.

Note parallelogram widths in Figure 5(c)/(d) and Figure 6(c):  $2n - 1$  and  $n$  coefficients, respectively, validating the earlier claim that the number of sum-reductions (shifts and flattenings) is cut approximately in half by integrating multiplication and reduction.

Algorithms 4.3 and 4.4 perform shifts and flattenings for columns 0 to 9 of the

---

**Algorithm 4.7** ACCUMULATEOUTERPRODUCTSREDUCTION( $\mathbf{a}, \mathbf{b}, j, r, n, m$ ): accumulate outer products for polynomial multiplication with reduction modulo  $x^n - 1$  to  $\mathbf{Z}[r : 2 : 62 + r]$ , where  $r = 0$  or  $1$ ; only the first  $m \leq 32$  columns are computed.

---

```

1:  $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}[32j : 32j + 31]), \text{ldy}(\mathbf{b}[0 : 31])$ 
2:  $\mathbf{Z}[r : 2 : 62 + r] \leftarrow \text{mac16}(\mathbf{Y}_1^\top \mathbf{X}_1, \text{columns} = 0 : m - 1)$ 
3: for  $1 \leq l \leq j$  do
4:    $k \leftarrow j - l$   $\triangleright$  Thus:  $k + l = j$ 
5:    $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}[32k : 32k + 31]), \text{ldy}(\mathbf{b}[32l : 32l + 31])$ 
6:    $\mathbf{Z}[r : 2 : 62 + r] \leftarrow \text{mac16}(\mathbf{Z}[r : 2 : 62 + r] + \mathbf{Y}_1^\top \mathbf{X}_1, \text{columns} = 0 : m - 1)$ 
7: for  $j < l \leq \lfloor (n - 1) / 32 \rfloor$  do
8:    $\mathbf{X}_1, \mathbf{Y}_1 \leftarrow \text{ldx}(\mathbf{a}[n - 32(l - 1) : n - 32(l - 1) + 31]), \text{ldy}(\mathbf{b}[32l : 32l + 31])$ 
9:    $\mathbf{Z}[r : 2 : 62 + r] \leftarrow \text{mac16}(\mathbf{Z}[r : 2 : 62 + r] + \mathbf{Y}_1^\top \mathbf{X}_1, \text{columns} = 0 : m - 1)$ 

```

---

parallelogram; in this example, Algorithm 4.3 is applied to the first two blocks (columns 0 to 3 and 4 to 7 of Figure 6(b)), and Algorithm 4.4 to columns 8 and 9 (which also includes columns 10 and 11 in the result, an issue we must work around). Handling columns 10, 11 and 12, i.e. merging them with the main matrix as in Figure 6(c) (in this case, with columns 0, 1 and 2, respectively) requires a new procedure, formalized as Algorithm 4.8.

We leverage the algorithms of this section and of Section 4.3 in a procedure for integrated polynomial multiplication and reduction, formalized as Algorithm 4.9, which we use as the main polynomial multiplication routine in NTRU, i.e. the `poly_Rq_mul()` function.

---

**Algorithm 4.8** MERGEFIRSTANDLASTBLOCKS( $\mathbf{c}, j, r, n$ ): sum-reduction of columns from the two most significant multiplication sub-parallelograms.

---

**Input:**  $j$  (index of current sub-parallelogram)  
**Input:**  $r \in \{0, 1\}$  (indicates relative order of even and odd rows of  $\mathbf{Z}$ )  
**Input:**  $n$  (reduction polynomial is  $x^n - 1$ )  
**Output:**  $\mathbf{c}[0 : 31]$  (coefficients of  $x^0$  through  $x^{31}$  of the polynomial multiplication result)  
**Notes:** assumes that  $\mathbf{Z}[1 - r : 2 : 63 - r]$  and  $\mathbf{Z}[r : 2 : 62 + r]$  contain the output of Algorithm 4.7 for sub-parallelogram indices  $j - 1$  and  $j$ , respectively

```

1:  $\delta \leftarrow (-n) \bmod 32$ 
2:  $\mathbf{Z}[r] \leftarrow \text{ldz}(\mathbf{c}[0 : 31])$ 
3:  $\mathbf{X}_2 \leftarrow \text{extrh}(\mathbf{Z}[3 - r])$ 
4:  $\mathbf{Z}[r] \leftarrow \text{mac16}(\mathbf{Z}[r] + \mathbf{X}[95 - \delta : 126 - \delta])$ 
5:  $i \leftarrow 1$ 
6: while  $i < 31 - \delta$  do
7:    $\mathbf{X}_2, \mathbf{Y}_2 \leftarrow \text{extrh}(\mathbf{Z}[2i + 3 - r]), \text{extrh}(\mathbf{Z}[2i + 5 - r])$ 
8:    $\mathbf{Z}[r] \leftarrow \text{vecint}(\mathbf{Z}[r] + \mathbf{X}[95 - \delta - i : 126 - \delta - i] + \mathbf{Y}[94 - \delta - i : 125 - \delta - i])$ 
9:    $i \leftarrow i + 2$ 
10: while  $i < 31$  do
11:    $\mathbf{X}_1, \mathbf{X}_2 \leftarrow \text{extrh}(\mathbf{Z}[2i + 2 + r]), \text{extrh}(\mathbf{Z}[2i + 3 - r])$ 
12:    $\mathbf{Y}_1, \mathbf{Y}_2 \leftarrow \text{extrh}(\mathbf{Z}[2i + 4 + r]), \text{extrh}(\mathbf{Z}[2i + 5 - r])$ 
13:    $\mathbf{Z}[r] \leftarrow \text{vecint}(\mathbf{Z}[r] + \mathbf{X}[95 - \delta - i : 126 - \delta - i] + \mathbf{Y}[94 - \delta - i : 125 - \delta - i])$ 
14:    $i \leftarrow i + 2$ 
15:  $\mathbf{c}[0 : 31] \leftarrow \text{stz}(\mathbf{Z}[r])$ 

```

---

## 4.5 Working around memory access slowdowns

An initial implementation of polynomial multiplication modulo  $x^n - 1$ , applying all optimizations mentioned in Sections 4.3 and 4.4, outperforms the analogous routines from the

---

**Algorithm 4.9** `POLYMODMUL(c, a, b, n)`: modulo- $2^{16}$  multiplication of two  $n$ -coefficient polynomials using AMX, reduced modulo  $x^n - 1$ .

---

**Input:** `a, b` (arrays of  $n$  coefficients each of  $a(x)$  and  $b(x)$ )

**Output:** `c` (array of  $n$  coefficients of  $c(x) = a(x)b(x)$ )

**Notes:** since AMX stores 64 bytes (in this case, 32 coefficients of 16 bits each) at a time, output array must be allocated with  $32\lceil n/32 \rceil$  coefficients

```

1:  $n' \leftarrow \lfloor (n-1)/32 \rfloor$ 
2:  $X_0, X_3, Y_0, Y_3 \leftarrow \text{ldx}([0, 0, \dots, 0])$  ▷ load zeros to all indicated registers
3: ACCUMULATEOUTERPRODUCTSREDUCTION(a, b, 0, 0, n, 32)
4: ACCUMULATEOUTERPRODUCTSREDUCTION(a, b, 1, 1, n, 32)
5: FLATTENFIRSTTWOBLOCKS(c)
6: for  $i = 2$  to  $n' - 2$  do
7:   ACCUMULATEOUTERPRODUCTSREDUCTION(a, b, i, i mod 2, n, 32)
8:   FLATTENMIDDLEBLOCK(c, i, i mod 2)
9: ACCUMULATEOUTERPRODUCTSREDUCTION(a, b, n' - 1, (n' - 1) mod 2, n, n mod 32)
10: FLATTENMIDDLEBLOCK(c, n' - 1, (n' - 1) mod 2)
11: MERGEFIRSTANDLASTBLOCKS(c, n' - 1, (n' - 1) mod 2, n)

```

---

implementations of [NG21] and [CCHY23] when benchmarked in isolation. However, improvements were modest for the operations of the full scheme (key generation, encapsulation and decapsulation), with [CCHY23] still outperforming our NTRU implementation.

Handley [Han23] claims that load/store queues track AMX memory accesses at a granularity of memory pages, marking them as “interesting” for AMX, and concurrent CPU and AMX accesses to the same page will block. Inspection of the `kem.c`, `owcpa.c` and `poly.c` source files shows that polynomial coefficient arrays are allocated on the stack, next to other variables used exclusively by CPU code, and thus stored in adjacent addresses in memory, very likely in the same memory page. We conjectured this was the root cause of our performance issues, and switched to allocating one polynomial per memory page using the `mmap()` function. Benchmarks using this new memory allocation scheme led to our implementation outperforming those of [NG21] and [CCHY23] in all tested scenarios.

Johnson [Joh22b] suggests using CPU prefetch instructions, which (unlike AMX instructions) are executed out of order by the CPU, triggering memory accesses in advance of AMX load/store instructions and reducing latency. However, we saw no performance improvements from this. It appears this suggestion is targeted towards typical AMX applications with working sets too large to fit CPU caches, which is not the case for NTRU.

As the CPU’s NEON units are independent from AMX, one might conjecture that splitting tasks between both can achieve even higher throughput. However, AMX/CPU memory conflicts incur performance penalties which may invalidate this strategy, unless NEON and AMX work on distinct data, stored on different memory pages.

## 5 Experimental results

We benchmark our implementation on a 2020 Apple MacBook Air laptop with the M1 SoC (P-core clock speed of 3200 MHz), and a 2023 Apple MacBook Pro laptop with the M3 Max SoC (P-core clock speed of 4064 MHz). All implementations are compiled with Apple clang 15 using the `-O3` optimization flag. Following guidance by ARM [ARM], and taking into account the recent GoFetch microarchitectural attack on Apple M-series SoCs [CWS<sup>+</sup>24], we measure execution times with the DIT (data-independent timing) bit set.

We compare our implementation to the state-of-the-art ones of [CCHY23], which implements only HPS2048677 and HRSS701 (and holds the current speed record for

these sets), and [NG21], which implements all parameter sets. We backport optimized polynomial inversion and constant-time sorting routines from [CCHY23] to [NG21], seeking to highlight differences in polynomial multiplication performance only; we include this modified implementation in our GitHub repository. We also optimize NIST’s AES-256 CTR-DRBG `randombytes()` function with AES instructions from ARMv8-A’s Cryptographic Extensions, using software tests to verify that outputs are bit-identical to the NIST version.

All of these optimizations (polynomial inversion, constant-time sorting and random number generation) are also applied to our AMX implementation. We ensure Known Answer Tests for all implementations match those provided in the NTRU reference code.

Recall from Section 4.5 that our implementation uses a specific memory allocation strategy for arrays of polynomial coefficients. To ensure a fair comparison, we create stack and `mmap()` allocation variants for every implementation and report performance figures for both, computing speedups by comparing the fastest strategy for each implementation.

Performance measurements use the same cycle-counting harness found in [NG21] and [CCHY23]. To reduce timing variability, each routine is run for 1,024 times in a loop, measuring the cycle counts of each run individually, and the average value is reported.

## 5.1 Performance measurements

We present NTRU performance results in Table 6. We omit the `tc` variant of HPS2048677 included in [CCHY23], as the `tmvp` variant always outperforms it.

**Table 6:** Cycle counts for NTRU KEM operations and polynomial multiplication. Speedups computed as a ratio of the best cycle counts for previous implementations and our implementation (in both cases, across different memory allocations types).

Par. set	Mem. alloc.	Work	Operation							
			Key gen.		Encaps.		Decaps.		Poly. mult.	
			M1	M3	M1	M3	M1	M3	M1	M3
509	mmap	[NG21]	218189	214775	16085	15560	29617	28336	6160	6141
		Ours	170217	163870	12452	11759	18821	17099	2364	2118
	stack	[NG21]	217970	214179	16091	15532	29604	28331	6157	6136
		Ours	181195	171333	19407	14258	26585	23792	3752	3622
	Speedup (×)		<b>1.28</b>	<b>1.31</b>	<b>1.29</b>	<b>1.32</b>	<b>1.57</b>	<b>1.66</b>	<b>2.60</b>	<b>2.90</b>
677	mmap	[NG21]	363495	348191	23460	22481	44713	42788	10056	9794
		[CCHY23]	307340	295968	19079	18362	31467	30407	7375	6813
		Ours	282943	265971	17157	16028	26314	23896	3780	3418
	stack	[NG21]	363323	347883	22508	22668	46359	44633	10040	9838
		[CCHY23]	306893	295643	20634	19347	35241	32715	5832	5893
		Ours	315413	291748	27591	16300	35505	32504	4770	3452
	Speedup (×)		<b>1.08</b>	<b>1.11</b>	<b>1.11</b>	<b>1.15</b>	<b>1.20</b>	<b>1.27</b>	<b>1.54</b>	<b>1.72</b>
821	mmap	[NG21]	497443	492582	28002	27330	56576	55332	13782	13835
		Ours	384345	371084	19369	18119	30827	27955	4868	4497
	stack	[NG21]	498127	492015	28044	27257	56792	55312	13793	13815
		Ours	608741	382194	27950	19185	41121	37968	6108	4497
	Speedup (×)		<b>1.29</b>	<b>1.33</b>	<b>1.45</b>	<b>1.50</b>	<b>1.84</b>	<b>1.98</b>	<b>2.83</b>	<b>3.07</b>
701	mmap	[NG21]	393214	380440	20378	19619	53837	52306	12407	12506
		[CCHY23]	323395	308329	14549	13874	37201	35171	6772	6681
		Ours	287182	268797	11451	10493	28902	26333	3768	3456
	stack	[NG21]	397073	384903	20545	20262	55017	53323	12147	12454
		[CCHY23]	323501	307870	14610	13944	37121	35218	6752	6684
		Ours	305125	272861	11570	10744	35851	31928	4774	3489
	Speedup (×)		<b>1.13</b>	<b>1.15</b>	<b>1.27</b>	<b>1.32</b>	<b>1.28</b>	<b>1.34</b>	<b>1.79</b>	<b>1.93</b>


In general, NEON implementations perform similarly regardless of memory allocation

strategy (with a few outliers), but AMX consistently favors `mmap()` allocation, more so in the M1; recall that speedup calculations pick the best performing memory allocation method for each implementation. The M3 AMX unit is faster than the M1’s; in particular, microbenchmarks show that the M3 executes some vector AMX operations at a rate of 2 instructions/cycle, a phenomenon we didn’t observe on the M1.

The most representative comparisons are with HPS2048677 and HRSS701 TMVP NEON implementations [CCHY23]. AMX polynomial multiplication achieves significant speedups of  $1.54\times$  (M1)/ $1.72\times$  (M3) for the former and  $1.79\times$  (M1)/ $1.93\times$  (M3) for the latter. These translate into smaller speedups for KEM operations ( $1.08\times$  to  $1.28\times$  for the M1 and  $1.11\times$  to  $1.34\times$  for the M3), since the time spent in other routines (such as polynomial inversion for key generation, and constant-time sorting for key generation and encapsulation) remain identical between our implementation and previous ones. Nevertheless, our implementation consistently outperforms all others. For other parameter sets, our gains are even more pronounced, but we expect the gap would shrink if TMVP were applied to them.

**Scaling and subquadratic algorithms.** In Figure 7(a), we provide polynomial multiplication scaling results as a function of polynomial degree, either with or without reduction modulo  $x^n - 1$ . We also perform a least-squares fit to an equation of the form  $ax^b$ , and note that in this range, the algorithm displays subquadratic scaling; this can be attributed to  $O(n^2)$  scaling of outer product operations (cheaper) versus  $O(n)$  scaling of shifts and flattenings (more expensive). We note that the version with reduction outperforms the one without, as it requires only about half as many slower shift/flatten operations.

We also implemented Karatsuba on AMX, for a single recursion level with our routine of Section 4.3 as a basecase, and compare it to the routine of Section 4.4 in Figure 7(b). Karatsuba only outperforms schoolbook for  $n \geq 4480$ , far outside the range of cryptographic interest. We give numerical results for all NTRU parameter sets in Table 7, and also compute lower bounds (accounting for the cost of pointwise multiplications only) for single-level Karatsuba and subquadratic algorithms across all possible multi-level recursion strategies, using either schoolbook, Karatsuba, Toom-3 or Toom-4 at each level (the bounds are computed by a Jupyter notebook included in our GitHub repository). We conclude that schoolbook outperforms any subquadratic strategy for all NTRU parameter sets.



**Figure 7:** Cycle counts for AMX polynomial multiplication routines on the M3 as a function of  $n$ : (a) schoolbook algorithm with and without reduction modulo  $x^n - 1$ ; (b) schoolbook and (single recursion level) Karatsuba algorithms with reduction.

**Constant-time experiments.** We empirically tested whether AMX instructions execute with data-independent timing, by benchmarking polynomial multiplication routines with either zero or random polynomials as inputs; the former were chosen as the most likely case for hypothetical data-dependent hardware optimizations. Experiments were run with the same combinations of parameter set, memory allocation and implementation as in Table 6, including NEON implementations (presumed constant-time) for comparison purposes.

Our experimental design uses paired measurements of cycle counts for multiplying either zero polynomials or randomly generated polynomials (generated once at the beginning,




**Table 7:** Cycle counts for schoolbook and Karatsuba polynomial multiplication in the M3, and a lower bound computed across all possible subquadratic recursion strategies.

Parameter set	Schoolbook performance	Subquadratic lower bound	Single level Karatsuba	
			lower bound	performance
HPS2048509	2117	2968	3444	3650
HPS2048677/HRSS701	3455	4696	5268	5599
HPS4096821	4496	6004	6636	6963

and reused for every measurement). We randomize the order (zero or random polynomials first), benchmark 64 runs and report the average cycle count. The process is repeated for 1,000,000 iterations, discarding the 0.5% smallest and 0.5% largest measurements. As with performance measurements, we enable the DIT (data-independent timing) bit.

We show a representative subset of results in Figure 8. In Figure 8(a), we see a multimodal distribution; we conjecture it arises from CPU/AMX memory subsystem interactions when stack allocation is used, as discussed in Section 4.5. Figure 8(b) shows that very stable results are obtained if `mmap` allocation is used with AMX. In Figure 8(c), we note a larger spread for NEON implementations; we conjecture this is due to scheduling variabilities in the CPU’s out-of-order engine across different iterations. In all cases, the measurements of zero and random polynomials in the left-hand-side histograms nearly coincide, while the middle and right-hand-side plots show that cycle count densities for zero and random inputs (respectively) over time are highly correlated.



**Figure 8:** Constant time experiment results on the M3, for selected choices of parameter set, memory allocation and implementation: (a) HPS2048509, stack, ours; (b) HPS2048509, `mmap`, ours; (c) HPS2048509, `mmap`, [NG21]. The left-hand-side graphs are histograms of cycle counts for multiplication with both zero and random inputs. The middle and right-hand-side are 2D histograms, plotting sample density over the duration of execution of each experiment, for zero and random inputs, respectively.

This suggests that results are indistinguishable between zero and random inputs, which we seek to confirm using statistical hypothesis tests. Conventional approaches, such as

Student’s  $t$ -test, evaluate the null hypothesis that the means of both distributions is identical, i.e. the difference of cycle counts between paired observations is zero. Given the large sample size of 1,000,000 iterations, even very small deviations from zero, attributable to measurement noise, may lead to a spurious rejection of the null hypothesis. Additionally, even if the null hypothesis is not rejected, one cannot conclude that there is no effect; in other words, “absence of evidence is not evidence of absence.” [AB95]

*Equivalence tests* [Lak17] are based on the concept of an equivalence range, such that any effects within that range are considered negligible or of no practical interest. The null hypothesis is that there is an effect outside of the equivalence range, and thus, a rejection of the null hypothesis can be properly interpreted as evidence for equivalence. We have chosen to use such a test, in particular the TOST (Two One-Sided Tests) method using Student’s  $t$ -test as the underlying hypothesis test, with an equivalence range of  $\pm 1$  clock cycle, which is the meaningful effect threshold in a clocked digital circuit such as an SoC.

A potential objection to the use of Student’s  $t$ -test is that it is parametric and thus assumes a normal distribution, which is not the case for some of our measurements, as seen in Figure 8. However, it is more accurate to say that the test assumes that the outcome variable (i.e. the mean of the difference between paired measurements) is normally distributed, which follows for large sample sizes from the central limit theorem [LDEC02].

In Table 8, we report confidence intervals and  $p$ -values at a significance level  $\alpha = 10^{-6}$ . In all cases (whether NEON or AMX implementations), the confidence intervals are contained in the equivalence range  $[-1, 1]$ , and  $p < \alpha$ . We conclude that all implementations deviate by  $< 1$  clock cycle between zero and random inputs, and thus are constant-time.

**Table 8:** Confidence intervals and  $p$ -values for the statistical hypothesis test of equivalence of cycle counts between zero and random inputs.

Par. set	Mem. alloc.	Work	M1		M3	
			Conf. interval	$p$ -value	Conf. interval	$p$ -value
509	mmap	[NG21]	$[-0.00483, 0.00468]$	0	$[-0.00456, 0.00495]$	0
		Ours	$[-0.00490, 0.00461]$	0	$[-0.00515, 0.00436]$	0
	stack	[NG21]	$[-0.00386, 0.00565]$	0	$[-0.00477, 0.00474]$	0
		Ours	$[-0.00255, 0.00696]$	0	$[-0.00446, 0.00505]$	0
677	mmap	[NG21]	$[-0.00502, 0.00449]$	$< 10^{-95}$	$[-0.00430, 0.00520]$	$< 10^{-59}$
		[CCHY23]	$[-0.00485, 0.00465]$	0	$[-0.01262, -0.00312]$	0
		Ours	$[-0.00582, 0.00368]$	0	$[-0.00463, 0.00488]$	0
	stack	[NG21]	$[-0.00565, 0.00386]$	$< 10^{-93}$	$[-0.00518, 0.00433]$	$< 10^{-56}$
		[CCHY23]	$[-0.00423, 0.00528]$	0	$[0.07198, 0.08150]$	0
		Ours	$[-0.00509, 0.00442]$	0	$[-0.00463, 0.00488]$	0
821	mmap	[NG21]	$[-0.00217, 0.00733]$	$< 10^{-162}$	$[-0.00458, 0.00492]$	$< 10^{-118}$
		Ours	$[-0.00571, 0.00380]$	0	$[-0.00308, 0.00642]$	0
	stack	[NG21]	$[-0.01076, -0.00125]$	$< 10^{-117}$	$[-0.00475, 0.00476]$	$< 10^{-119}$
		Ours	$[-0.00588, 0.00363]$	0	$[-0.00442, 0.00509]$	0
701	mmap	[NG21]	$[-0.00726, 0.00224]$	$< 10^{-91}$	$[-0.00523, 0.00427]$	$< 10^{-17}$
		[CCHY23]	$[-0.00692, 0.00259]$	0	$[-0.00471, 0.00480]$	0
		Ours	$[-0.00349, 0.00602]$	0	$[-0.00774, 0.00176]$	0
	stack	[NG21]	$[-0.01237, -0.00286]$	$< 10^{-8}$	$[-0.00438, 0.00512]$	$< 10^{-52}$
		[CCHY23]	$[-0.00751, 0.00200]$	0	$[0.00406, 0.01357]$	0
		Ours	$[-0.00583, 0.00368]$	0	$[-0.00585, 0.00366]$	0

We also measure the latency of matrix-mode `mac16` for both zero and random inputs. In a loop, we run `mac16` and extract the results back from `Z` to `X` and `Y`, creating an inter-iteration data dependency which serializes all operations, so the code is latency-bound. Running this on the M1 (M3) for 100,000 iterations, we get 2,699,699 (2,299,566) cycles for zero inputs and 2,699,688 (2,299,666) cycles for random inputs, a difference of 11 (100) cycles. A variable-time ALU would show at least 1 cycle difference in latency between each case, so one would expect a minimum difference of 1 cycle/iteration, i.e., 100,000 cycles.

In summary, we see no evidence of data-dependent timing in AMX.

## 6 Conclusion

We report on an implementation of polynomial multiplication, applied to the NTRU lattice-based PQC scheme, leveraging a matrix-multiplication coprocessor which, as of 2024, is shipping on hundreds of millions of devices. Our implementation sets new speed records on two representative devices featuring this coprocessor, the Apple M1 and M3 SoCs, outperforming state-of-the-art conventional implementations for the same platforms.

CPU-coupled matrix multiplication accelerators will soon flood the market, bringing about huge leaps in processing power, but implementers must adapt to their new architectural paradigms to harness their added performance. Our most surprising conclusion stems from multiplication (2D outer product) being much cheaper in AMX than vector (1D) operations; the latter must be applied per-row/column, increasing instruction count. As subquadratic algorithms trade off multiplication count for extra vector operations, AMX’s architectural tradeoffs favor schoolbook for polynomial degrees of cryptographic interest.

**Future work.** We suggest several avenues for future work based on this paper’s results.

The GoFetch attack [CWS<sup>+</sup>24] requires expensive countermeasures for the M1 and M2 SoCs, but we note AMX may be unaffected by this attack. Should this be confirmed, AMX implementations may be used as a countermeasure and even bring performance gains.

Other schemes based on NTT-unfriendly rings may see similar speedups from AMX as NTRU. It is unclear whether NTT-based schemes such as Kyber [SAB<sup>+</sup>20] and Dilithium [LDK<sup>+</sup>22] are a good match for AMX, but it merits investigation.

While we have argued that schoolbook should outperform other subquadratic algorithms for NTRU, we assume that recursion is performed in a “black-box” fashion, calling self-contained functions to compute multiplications of smaller degree. We encourage investigating if an integrated approach, as in Sections 4.3 and 4.4, can improve performance.

Finally, we are witnessing the initial generations of matrix-multiplication accelerators. As they evolve, implementation techniques must equally adapt: faster vector operations could render subquadratic algorithms feasible, and the introduction of instructions to accelerate matrix-vector products could favor the use of TMVP. In this vein, other accelerators such as Intel AMX and ARMv9-A SME undoubtedly differ to some degree in available instructions and performance characteristics from Apple’s AMX, and warrant a careful study aiming to exploit their performance potential for post-quantum cryptography.

## Acknowledgments

The first and third author thank the Cryptography Research Centre at the Technology Innovation Institute for financial support during the development of this work.

We thank Gora Adj for finding some typos in the algorithms; Jacques Wainer and Carlos A. A. Trujillo for discussions about the statistical analysis of the constant-time behavior of AMX; and the anonymous reviewers for helpful suggestions.

## References

- [AB95] D. G. Altman and J. M. Bland. Statistics notes: Absence of evidence is not evidence of absence. *BMJ*, 311(7003):485–485, August 1995. doi:[10.1136/bmj.311.7003.485](https://doi.org/10.1136/bmj.311.7003.485).
- [Ame17] American National Standards Institute. Lattice-based polynomial public key establishment algorithm for the financial services industry. ASC X9.98-2010 (R2017), 2017.
- [App23] Apple Inc. Accelerate: Make large-scale mathematical computations and image calculations, optimized for high performance and low energy consumption, 2023. <https://developer.apple.com/documentation/accelerate>.
- [ARM] ARM Limited. How is instruction timing affected by the FEAT\_DIT architectural feature? URL: <https://developer.arm.com/documentation/ka005181/latest/>.
- [BCD<sup>+</sup>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](https://doi.org/10.1145/2976749.2978425).
- [Bun21] Bundesamt für Sicherheit in der Informationstechnik. Migration to post quantum cryptography: Recommendations for action by the BSI, 2021. URL: [https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/Migration\\_to\\_Post\\_Quantum\\_Cryptography.pdf?\\_\\_blob=publicationFile&v=2](https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/Migration_to_Post_Quantum_Cryptography.pdf?__blob=publicationFile&v=2).
- [Caw24] Peter Cawley. Apple AMX instruction set, 2024. <https://github.com/cor-six/amx/>.
- [CCHY23] Han-Ting Chen, Yi-Hua Chung, Vincent Hwang, and Bo-Yin Yang. Algorithmic views of vectorized polynomial multipliers – NTRU. Cryptology ePrint Archive, Report 2023/1637, 2023. <https://ia.cr/2023/1637>. To appear at INDOCRYPT 2023.
- [CDH<sup>+</sup>20] Cong Chen, Oussama Danba, Jeffrey Hoffstein, Andreas Hulsing, 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<sup>+</sup>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>.
- [DKR<sup>+</sup>20] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, Frederik Vercauteren, Jose Maria Bermudo Mera, Michiel Van Beirendonck, and Andrea Basso. SABER. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography>.

- [cryptography/post-quantum-cryptography-standardization/round-3-submissions](#).
- [Han23] Maynard Handley. AArch64-Explore: Exploration of Apple CPUs – volume 3: SoC, 2023. <https://github.com/name99-org/AArch64-Explore>.
- [HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017: 15th Theory of Cryptography Conference, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 341–371. Springer, Heidelberg, November 2017. doi:10.1007/978-3-319-70500-2\_12.
- [HP19] John L. Hennessy and David A. Patterson. A new golden age for computer architecture. *Communications of the ACM*, 62(2):48–60, January 2019. doi:10.1145/3282307.
- [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>.
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In Joe P. Buhler, editor, *Algorithmic Number Theory*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg. doi:10.1007/BFb0054868.
- [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.
- [Int22] Intel Corporation. Intel® architecture instruction set extensions and future features: Programming reference (revision 047), December 2022. <https://cdrdv2-public.intel.com/671368/architecture-instruction-set-extensions-programming-reference.pdf>.
- [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>.
- [Joh22a] Dougall Johnson. Apple M1 microarchitecture research. <https://dougallj.github.io/applecpu/firestorm.html>, 2022.

- [Joh22b] Dougall Johnson. IDA (disassembler) and Hex-Rays (decompiler) plugin for Apple AMX, 2022. <https://gist.github.com/dougallj/7a75a3be1ec69ca550e7c36dc75e0d6f>.
- [JYP<sup>+</sup>17] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, ISCA '17, page 1–12, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3079856.3080246.
- [Lak17] Daniël Lakens. Equivalence tests: A practical primer for  $t$  tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, 8(4):355–362, May 2017. doi:10.1177/1948550617697177.
- [LDEC02] Thomas Lumley, Paula Diehr, Scott Emerson, and Lu Chen. The importance of the normality assumption in large public health data sets. *Annual Review of Public Health*, 23(1):151–169, May 2002. doi:10.1146/annurev.publhealth.23.100901.140546.
- [LDK<sup>+</sup>22] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [LHKK79] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh. Basic linear algebra subprograms for Fortran usage. *ACM Trans. Math. Softw.*, 5(3):308–323, September 1979. doi:10.1145/355841.355847.
- [LSH<sup>+</sup>22] Wai-Kong Lee, Hwajeong Seo, Seong Oun Hwang, Ramachandra Achar, Angshuman Karmakar, and Jose Maria Bermudo Mera. DPCrypto: Acceleration of post-quantum cryptography using dot-product instructions on GPUs. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 69(9):3591–3604, 2022. doi:10.1109/TCSI.2022.3176966.
- [LSZH22] Wai-Kong Lee, Hwajeong Seo, Zhenfei Zhang, and Seong Oun Hwang. TensorCrypto: High throughput acceleration of lattice-based cryptography using tensor core on GPU. *IEEE Access*, 10:20616–20632, 2022. doi:10.1109/ACCESS.2022.3152217.
- [MBB<sup>+</sup>21] José E. Moreira, Kit Barton, Steven Battle, Peter Bergner, Ramon Bertran, Puneeth Bhat, Pedro Caldeira, David Edelsohn, Gordon C. Fossum, Brad Frey,



- Nemanja Ivanovic, Chip Kerchner, Vincent Lim, Shakti Kapoor, Tulio Machado Filho, Silvia Melitta Mueller, Brett Olsson, Satish Sadasivam, Baptiste Saleil, Bill Schmidt, Rajalakshmi Srinivasaraghavan, Shricharan Srivatsan, Brian W. Thompto, Andreas Wagner, and Nelson Wu. A matrix math facility for Power ISA(TM) processors, 2021. doi:10.48550/arXiv.2104.03142.
- [MCL<sup>+</sup>18] S. Markidis, S. Chien, E. Laure, I. Peng, and J. S. Vetter. NVIDIA tensor core programmability, performance & precision. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 522–531, Los Alamitos, CA, USA, May 2018. IEEE Computer Society. doi:10.1109/IPDPSW.2018.00091.
- [NAB<sup>+</sup>20] Michael Naehrig, Erdem Alkim, Joppe Bos, Léo Ducas, Karen Easterbrook, Brian LaMacchia, Patrick Longa, Ilya Mironov, Valeria Nikolaenko, Christopher Peikert, Ananth Raghunathan, and Douglas Stebila. FrodoKEM. 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>.
- [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, July 2021. doi:10.1007/978-3-030-81293-5\_13.
- [Rod20] Andres Rodriguez. *Deep Learning Systems: Algorithms, Compilers, and Processors for Large-Scale Production*. Synthesis Lectures on Computer Architecture. Springer Cham, October 2020. doi:10.1007/978-3-031-01769-8.
- [SAB<sup>+</sup>20] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, and Damien Stehlé. CRYSTALS-KYBER. 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>.
- [SBG<sup>+</sup>16] Ali Sazegari, Eric Bainville, Jeffery E. Gunion, Gerard R. Williams III, and Andrew J. Beaumont-Smith. Outer product engine. US patent US2018/074824 A1, 2016. URL: <https://patents.google.com/patent/US20180074824A1/en>.
- [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.
- [SHRS17] John M. Schanck, Andreas Hulsing, Joost Rijneveld, and Peter Schwabe. NTRU-HRSS-KEM. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions>.

- [WMS22] Finn Wilkinson and Simon McIntosh-Smith. An initial evaluation of Arm’s Scalable Matrix Extension. In *2022 IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pages 135–140, 2022. doi:[10.1109/PMBS56514.2022.00018](https://doi.org/10.1109/PMBS56514.2022.00018).
- [WZF<sup>+</sup>22] Lipeng Wan, Fangyu Zheng, Guang Fan, Rong Wei, Lili Gao, Yuewu Wang, Jingqiang Lin, and Jiankuo Dong. A novel high-performance implementation of CRYSTALS-Kyber with AI accelerator. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *ESORICS 2022: 27th European Symposium on Research in Computer Security, Part III*, volume 13556 of *Lecture Notes in Computer Science*, pages 514–534. Springer, Heidelberg, September 2022. doi:[10.1007/978-3-031-17143-7\\_25](https://doi.org/10.1007/978-3-031-17143-7_25).
- [ZCH<sup>+</sup>19] Zhenfei Zhang, Cong Chen, Jeffrey Hoffstein, William Whyte, John M. Schanck, Andreas Hulsing, Joost Rijneveld, Peter Schwabe, and Oussama Danba. NTRU-Encrypt. 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>.