

Faster Kyber and Dilithium on the Cortex-M4

Amin Abdulrahman^{1,2} Vincent Hwang^{3,4} Matthias J. Kannwischer³ Amber
Sprenkels⁵

¹Ruhr University Bochum, Germany

²Max Planck Institute for Security and Privacy, Bochum, Germany

³Academia Sinica, Taipei, Taiwan

⁴National Taiwan University, Taipei, Taiwan

⁵Digital Security Group, Radboud University, Nijmegen, The Netherlands

23rd June 2022

Section 1

Introduction

- ▶ Kyber, Dilithium
- ▶ Part of CRYSTALS
- ▶ NIST PQC round 3 finalists
- ▶ Lattice-based

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
$$\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
$$\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
$$\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
$$\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Kyber

- ▶ IND-CCA2 secure KEM
- ▶ Based on MLWE
- ▶ Operates on
 $\mathcal{R}_{3329} = \mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

Dilithium

- ▶ Signature scheme that is strongly secure under CMA
- ▶ Based on Fiat-Shamir with Aborts, MSIS, and MLWE
- ▶ Operates on
 $\mathcal{R}_{8380417} = \mathbb{Z}_{8380417}[X]/(X^{256} + 1)$
- ▶ Built to profit from NTT

NTT in general

- ▶ Variant of the DFT defined over finite fields
- ▶ Negacyclic NTT $\hat{=}$ Evaluation of polynomial at powers of primitive n -th root of unity ζ_n for \mathcal{R}_q followed by twisting with powers of $2n$ -th root of unity ζ_{2n} .

$$\text{NTT}(a) = \hat{a} = \sum_{i=0}^{n-1} \hat{a}_i X^i \quad \text{with} \quad \hat{a}_i = \sum_{j=0}^{n-1} a_j \zeta_{2n}^j \zeta_n^{ij}$$

$$\text{iNTT}(\hat{a}) = a = \sum_{i=0}^{n-1} a_i X^i \quad \text{with} \quad a_i = n^{-1} \zeta_{2n}^{-i} \sum_{j=0}^{n-1} \hat{a}_j \zeta_n^{-ij}$$

- ▶ Efficient NTT using Cooley–Tukey or Gentleman–Sande FFT algorithms
- ▶ Fast polynomial multiplication: Let $f, g \in \mathcal{R}_q$ and \circ be base multiplication in \mathcal{R}_q

$$f \circ b = \text{iNTT}(\text{NTT}(f) \circ \text{NTT}(g))$$

NTT in general

- ▶ Variant of the DFT defined over finite fields
- ▶ Negacyclic NTT $\hat{=}$ Evaluation of polynomial at powers of primitive n -th root of unity ζ_n for \mathcal{R}_q followed by twisting with powers of $2n$ -th root of unity ζ_{2n} .

$$\text{NTT}(a) = \hat{a} = \sum_{i=0}^{n-1} \hat{a}_i X^i \quad \text{with} \quad \hat{a}_i = \sum_{j=0}^{n-1} a_j \zeta_{2n}^j \zeta_n^{ij}$$
$$\text{iNTT}(\hat{a}) = a = \sum_{i=0}^{n-1} a_i X^i \quad \text{with} \quad a_i = n^{-1} \zeta_{2n}^{-i} \sum_{j=0}^{n-1} \hat{a}_j \zeta_n^{-ij}$$

- ▶ Efficient NTT using Cooley–Tukey or Gentleman–Sande FFT algorithms
- ▶ Fast polynomial multiplication: Let $f, g \in \mathcal{R}_q$ and \circ be base multiplication in \mathcal{R}_q

$$f \circ b = \text{iNTT}(\text{NTT}(f) \circ \text{NTT}(g))$$

NTT in general

- ▶ Variant of the DFT defined over finite fields
- ▶ Negacyclic NTT $\hat{=}$ Evaluation of polynomial at powers of primitive n -th root of unity ζ_n for \mathcal{R}_q followed by twisting with powers of $2n$ -th root of unity ζ_{2n} .

$$\text{NTT}(a) = \hat{a} = \sum_{i=0}^{n-1} \hat{a}_i X^i \quad \text{with} \quad \hat{a}_i = \sum_{j=0}^{n-1} a_j \zeta_{2n}^j \zeta_n^{ij}$$
$$\text{iNTT}(\hat{a}) = a = \sum_{i=0}^{n-1} a_i X^i \quad \text{with} \quad a_i = n^{-1} \zeta_{2n}^{-i} \sum_{j=0}^{n-1} \hat{a}_j \zeta_n^{-ij}$$

- ▶ Efficient NTT using Cooley–Tukey or Gentleman–Sande FFT algorithms
- ▶ Fast polynomial multiplication: Let $f, g \in \mathcal{R}_q$ and \circ be base multiplication in \mathcal{R}_q

$$f \circ b = \text{iNTT}(\text{NTT}(f) \circ \text{NTT}(g))$$

NTT in general

- ▶ Variant of the DFT defined over finite fields
- ▶ Negacyclic NTT $\hat{=}$ Evaluation of polynomial at powers of primitive n -th root of unity ζ_n for \mathcal{R}_q followed by twisting with powers of $2n$ -th root of unity ζ_{2n} .

$$\text{NTT}(a) = \hat{a} = \sum_{i=0}^{n-1} \hat{a}_i X^i \quad \text{with} \quad \hat{a}_i = \sum_{j=0}^{n-1} a_j \zeta_{2n}^j \zeta_n^{ij}$$
$$\text{iNTT}(\hat{a}) = a = \sum_{i=0}^{n-1} a_i X^i \quad \text{with} \quad a_i = n^{-1} \zeta_{2n}^{-i} \sum_{j=0}^{n-1} \hat{a}_j \zeta_n^{-ij}$$

- ▶ Efficient NTT using Cooley–Tukey or Gentleman–Sande FFT algorithms
- ▶ Fast polynomial multiplication: Let $f, g \in \mathcal{R}_q$ and \circ be base multiplication in \mathcal{R}_q

$$f \circ b = \text{iNTT}(\text{NTT}(f) \circ \text{NTT}(g))$$

- ▶ Special case of the NTT with modulus a Fermat number $F_t := 2^{2^t} + 1$
- ▶ For F_t prime: Cyclic transformations up to $n = 2^{2^t} = F_t - 1$, negacyclic transformations up to $n = 2^{2^t - 1}$
 - ⇒ Twiddles on first t layers are powers of two
 - ⇒ No multiplication, only shifting
- ▶ Prime Fermat numbers: $F_0 = 3$, $F_1 = 5$, $F_2 = 17$, $F_3 = 257$, $F_4 = 65537$

- ▶ Special case of the NTT with modulus a Fermat number $F_t := 2^{2^t} + 1$
- ▶ For F_t prime: Cyclic transformations up to $n = 2^{2^t} = F_t - 1$, negacyclic transformations up to $n = 2^{2^t - 1}$
 - ⇒ Twiddles on first t layers are powers of two
 - ⇒ No multiplication, only shifting
- ▶ Prime Fermat numbers: $F_0 = 3, F_1 = 5, F_2 = 17, F_3 = 257, F_4 = 65\,537$

- ▶ Special case of the NTT with modulus a Fermat number $F_t := 2^{2^t} + 1$
- ▶ For F_t prime: Cyclic transformations up to $n = 2^{2^t} = F_t - 1$, negacyclic transformations up to $n = 2^{2^t - 1}$
 - ⇒ Twiddles on first t layers are powers of two
 - ⇒ No multiplication, only shifting
- ▶ Prime Fermat numbers: $F_0 = 3, F_1 = 5, F_2 = 17, F_3 = 257, F_4 = 65\,537$

- ▶ Special case of the NTT with modulus a Fermat number $F_t := 2^{2^t} + 1$
- ▶ For F_t prime: Cyclic transformations up to $n = 2^{2^t} = F_t - 1$, negacyclic transformations up to $n = 2^{2^t - 1}$
 - ⇒ Twiddles on first t layers are powers of two
 - ⇒ No multiplication, only shifting
- ▶ Prime Fermat numbers: $F_0 = 3, F_1 = 5, F_2 = 17, F_3 = 257, F_4 = 65\,537$

- ▶ Special case of the NTT with modulus a Fermat number $F_t := 2^{2^t} + 1$
- ▶ For F_t prime: Cyclic transformations up to $n = 2^{2^t} = F_t - 1$, negacyclic transformations up to $n = 2^{2^t - 1}$
 - ⇒ Twiddles on first t layers are powers of two
 - ⇒ No multiplication, only shifting
- ▶ Prime Fermat numbers: $F_0 = 3$, $F_1 = 5$, $F_2 = 17$, $F_3 = 257$, $F_4 = 65537$

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

- ▶ Hardware target: STM32F407-DISCOVERY with STM32-F407VG MCU
- ▶ 1 MiB flash, 192 KiB
- ▶ Based on Armv7E-M
- ▶ 14 usable general purpose registers
- ▶ 32 single-precision floating-point registers
- ▶ Powerful DSP with useful SIMD instructions taking one cycle
 - ▶ e.g., `smul{b,t}{b,t}`, `smlad{,x}`

Section 2

Kyber

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

- ▶ Three different parameter sets: Kyber-512, Kyber-768, and Kyber-1024.
- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ No $2n$ -th but only n -th primitive root of unity
 - ⇒ 7 layer incomplete NTT
 - ⇒ 2×2 schoolbook multiplications modulo $(X^2 - \omega)$

Algorithm: Kyber PKE key gen.

Output: public key: $pk = (\hat{\mathbf{t}}, \rho)$

Output: secret key: $sk = (\hat{\mathbf{s}})$

- 1 $\rho, \sigma \in \{0, 1\}^{256} \leftarrow \text{sampleUniform}()$
 - 2 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 3 $\mathbf{s}, \mathbf{e} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{m_1}(\sigma)$
 - 4 $\hat{\mathbf{t}} \leftarrow \hat{\mathbf{A}} \circ \text{NTT}(\mathbf{s}) + \text{NTT}(\mathbf{e})$
 - 5 **return** (pk, sk)
-

Algorithm: Kyber PKE decryption

Input : secret key: $sk = (\hat{\mathbf{s}})$

Input : compressed ciphertext: $(\mathbf{u}', \mathbf{v}')$

Output: message $m \in \mathcal{R}_q$

- 1 $\mathbf{u} \leftarrow \text{Decompress}(\mathbf{u}')$
 - 2 $\mathbf{v} \leftarrow \text{Decompress}(\mathbf{v}')$
 - 3 **return** $m \leftarrow \mathbf{v} - \text{iNTT}(\hat{\mathbf{s}}^T \circ \text{NTT}(\mathbf{u}))$
-

Algorithm: Kyber PKE encryption

Input : public key: $pk = (\hat{\mathbf{t}}, \rho)$

Input : message: $m \in \mathcal{R}_q$

Input : random coins: $\mu \in \{0, 1\}^{256}$

Output: ciphertext $(\mathbf{u}', \mathbf{v}')$

- 1 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 2 $\mathbf{r} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{m_1}(\mu)$
 - 3 $\mathbf{e}_1 \in \mathcal{R}_q^{k \times 1}, \mathbf{e}_2 \in \mathcal{R}_q \leftarrow \text{sampleCBD}^{m_2}(\mu)$
 - 4 $\hat{\mathbf{r}} \leftarrow \text{NTT}(\mathbf{r})$
 - 5 $\mathbf{u} \leftarrow \text{iNTT}(\hat{\mathbf{A}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_1$
 - 6 $\mathbf{v} \leftarrow \text{iNTT}(\hat{\mathbf{t}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_2 + m$
 - 7 **return** $(\text{Compress}(\mathbf{u}), \text{Compress}(\mathbf{v}))$
-

Optimization: NTT and inverse NTT

- ▶ Caching in FPU registers: Store reusable values in floating point registers to avoid loading from memory
- ▶ CT-Butterflies for i NTT: Avoid intermediate reductions
- ▶ Better layer merging: Merge layers 7–4, 3–1 instead of 7–5, 4–2, and computing layer 1 separately

Optimization: NTT and inverse NTT

- ▶ Caching in FPU registers: Store reusable values in floating point registers to avoid loading from memory
- ▶ CT-Butterflies for i NTT: Avoid intermediate reductions
- ▶ Better layer merging: Merge layers 7–4, 3–1 instead of 7–5, 4–2, and computing layer 1 separately

Optimization: NTT and inverse NTT

- ▶ Caching in FPU registers: Store reusable values in floating point registers to avoid loading from memory
- ▶ CT-Butterflies for i NTT: Avoid intermediate reductions
- ▶ Better layer merging: Merge layers 7–4, 3–1 instead of 7–5, 4–2, and computing layer 1 separately

4 Layer Merge

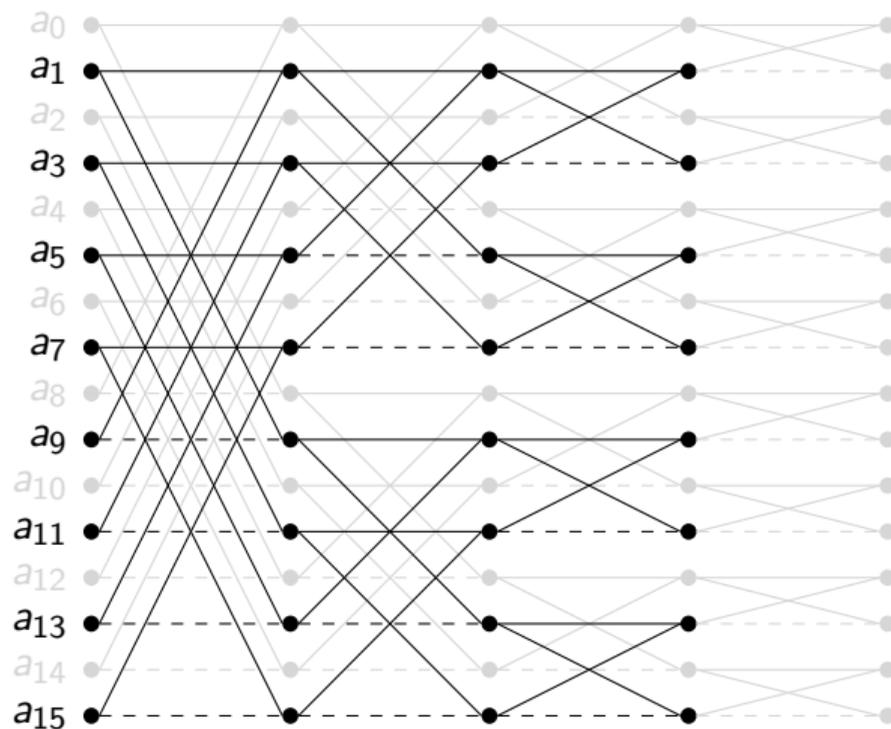


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

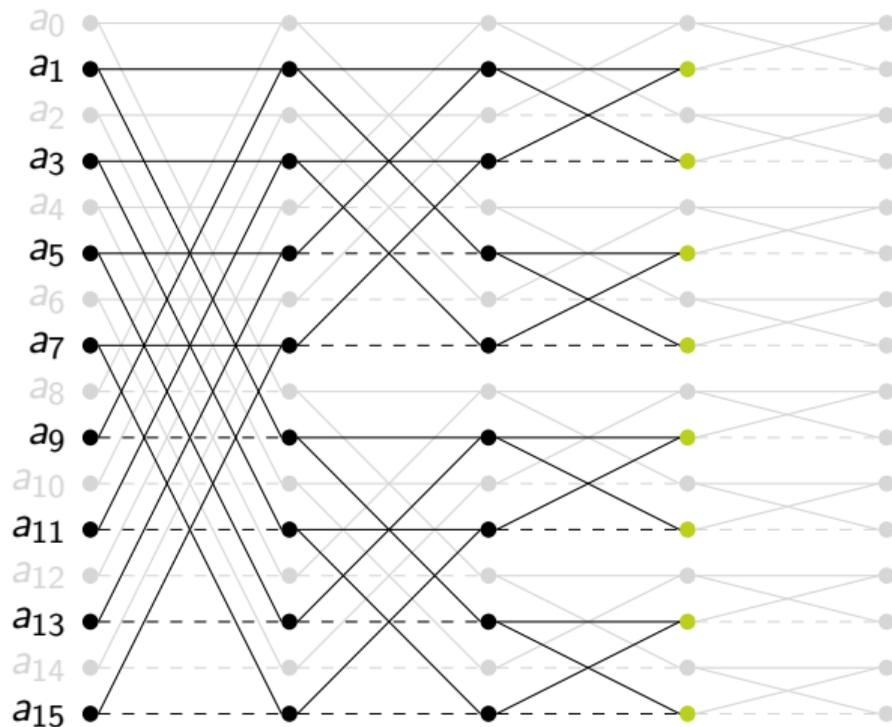


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

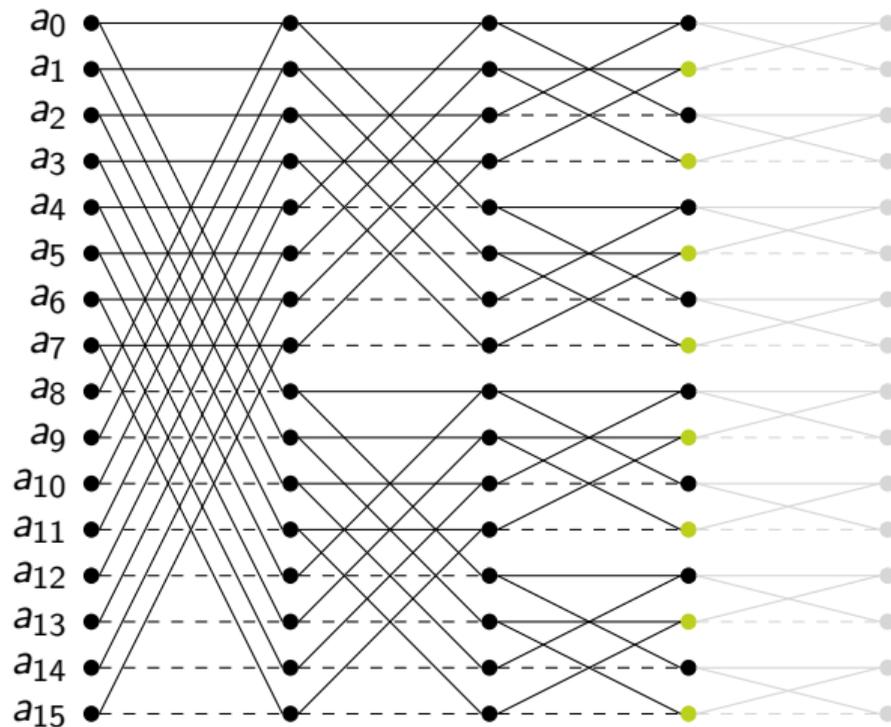


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

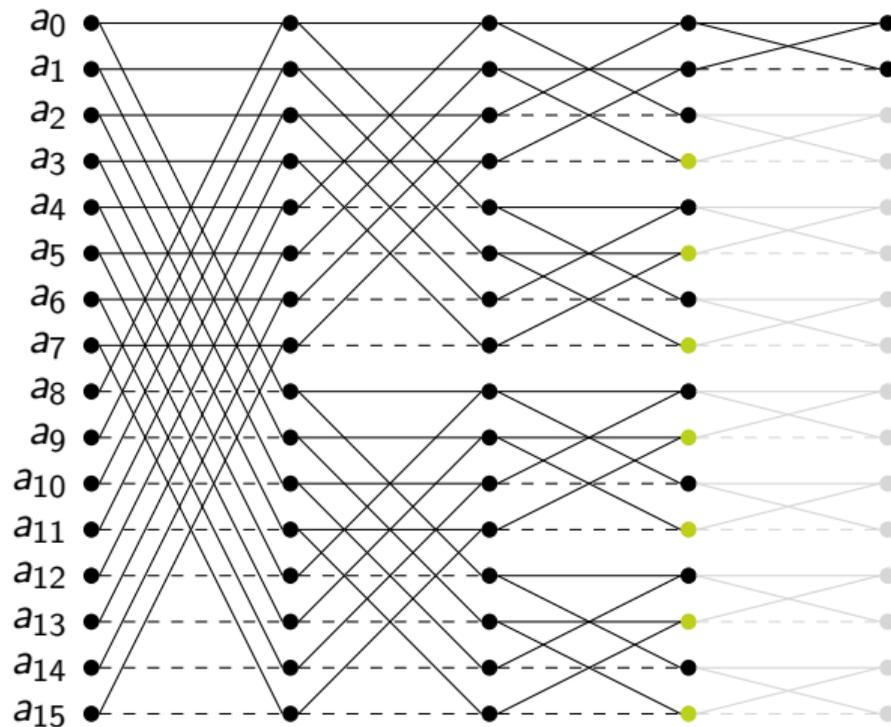


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

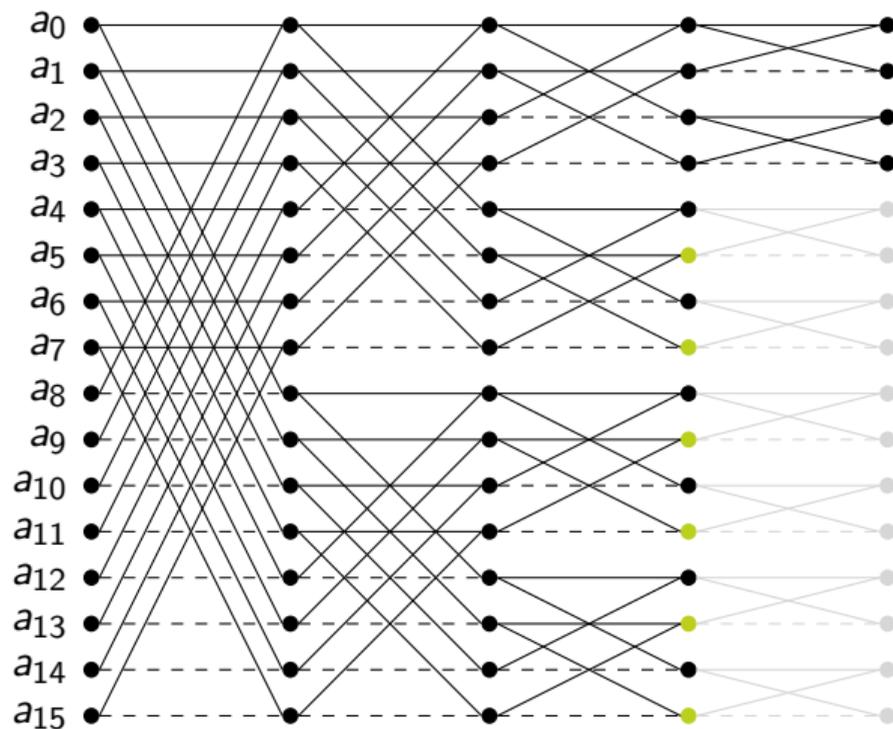


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

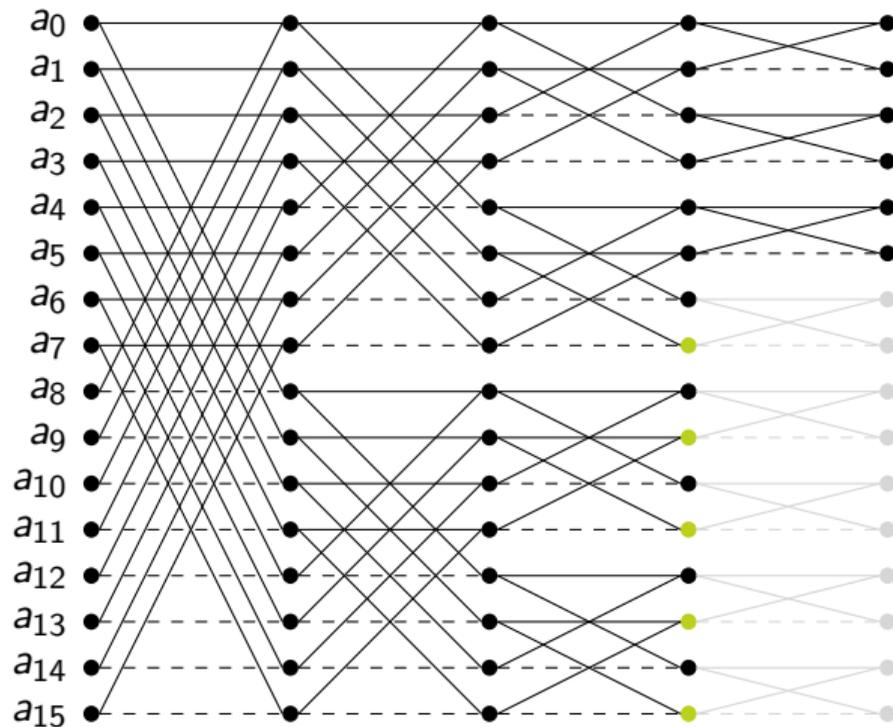


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

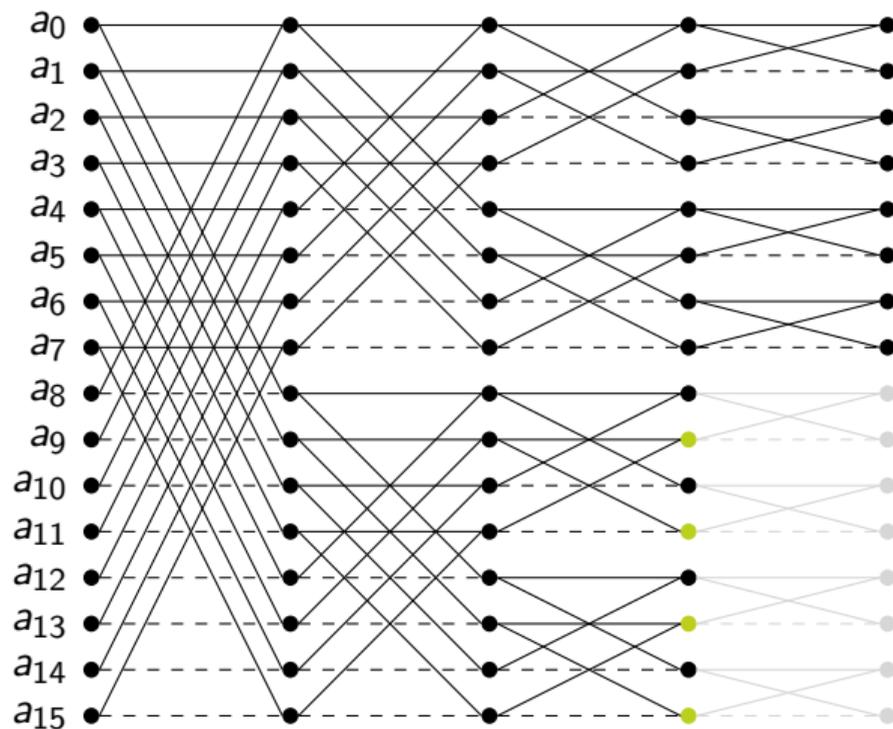


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

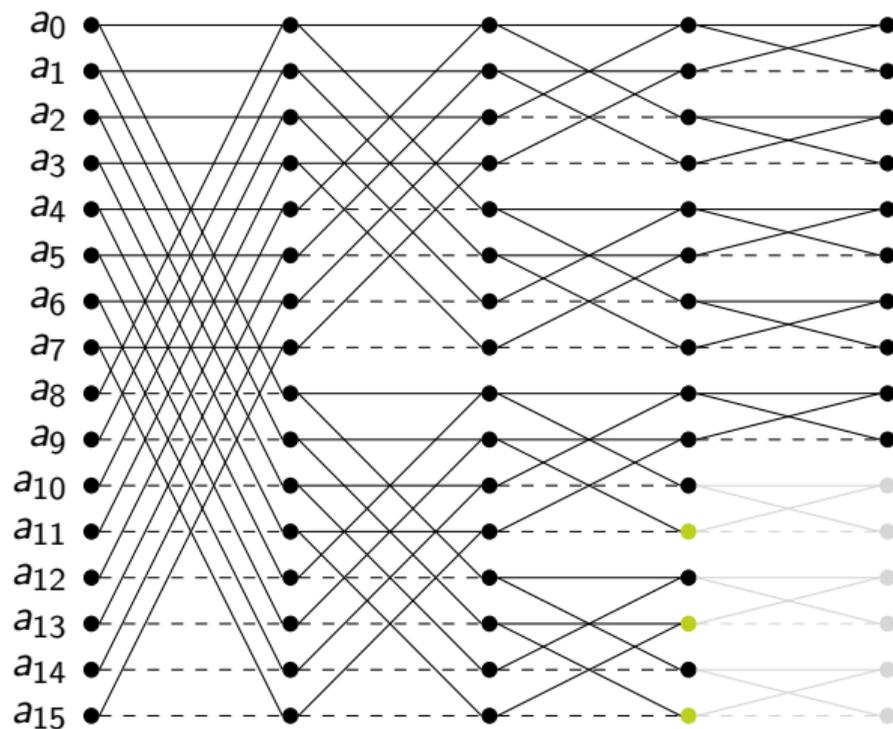


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

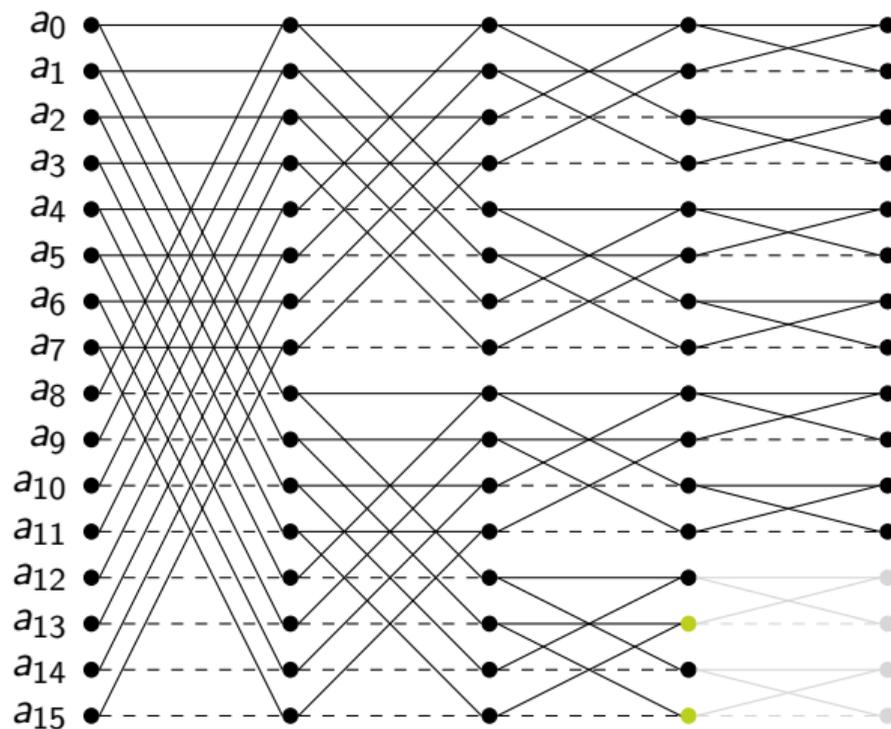


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

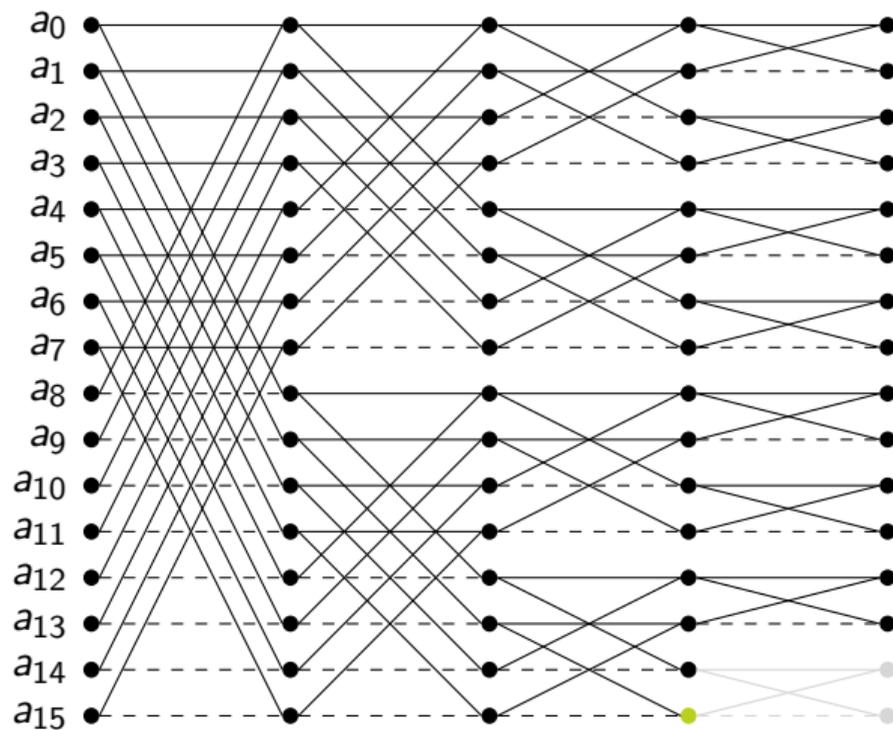


Figure: 4 layer radix 2 NTT signal flow

4 Layer Merge

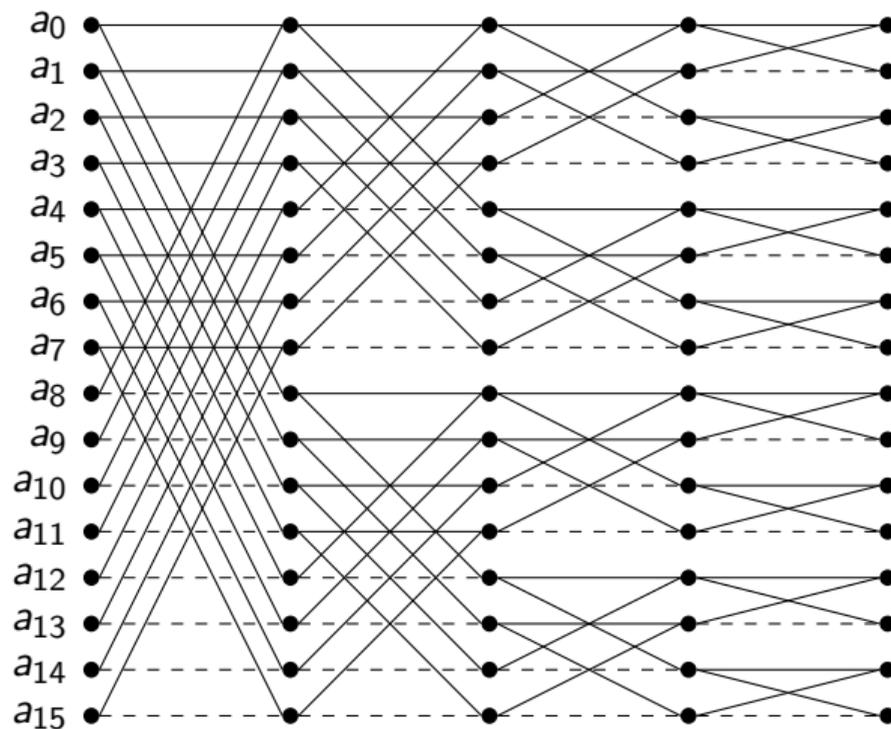


Figure: 4 layer radix 2 NTT signal flow

Algorithm: Kyber PKE key gen.

Output: public key: $pk = (\hat{\mathbf{t}}, \rho)$

Output: secret key: $sk = (\hat{\mathbf{s}})$

- 1 $\rho, \sigma \in \{0, 1\}^{256} \leftarrow \text{sampleUniform}()$
 - 2 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 3 $\mathbf{s}, \mathbf{e} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^m(\sigma)$
 - 4 $\hat{\mathbf{t}} \leftarrow \hat{\mathbf{A}} \circ \text{NTT}(\mathbf{s}) + \text{NTT}(\mathbf{e})$
 - 5 **return** (pk, sk)
-

Algorithm: Kyber PKE decryption

Input : secret key: $sk = (\hat{\mathbf{s}})$

Input : compressed ciphertext: (\mathbf{u}', v')

Output: message $m \in \mathcal{R}_q$

- 1 $\mathbf{u} \leftarrow \text{Decompress}(\mathbf{u}')$
 - 2 $v \leftarrow \text{Decompress}(v')$
 - 3 **return** $m \leftarrow v - \text{iNTT}(\hat{\mathbf{s}}^T \circ \text{NTT}(\mathbf{u}))$
-

Algorithm: Kyber PKE encryption

Input : public key: $pk = (\hat{\mathbf{t}}, \rho)$

Input : message: $m \in \mathcal{R}_q$

Input : random coins: $\mu \in \{0, 1\}^{256}$

Output: ciphertext (\mathbf{u}', v')

- 1 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 2 $\mathbf{r} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^m(\mu)$
 - 3 $\mathbf{e}_1 \in \mathcal{R}_q^{k \times 1}, \mathbf{e}_2 \in \mathcal{R}_q \leftarrow \text{sampleCBD}^{m_2}(\mu)$
 - 4 $\hat{\mathbf{r}} \leftarrow \text{NTT}(\mathbf{r})$
 - 5 $\mathbf{u} \leftarrow \text{iNTT}(\hat{\mathbf{A}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_1$
 - 6 $v \leftarrow \text{iNTT}(\hat{\mathbf{t}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_2 + m$
 - 7 **return** $(\text{Compress}(\mathbf{u}), \text{Compress}(v))$
-

Optimization: Barrett Reduction

Algorithm: Packed Barrett Reduction [BKS19]

Input : $a = (a_t \parallel a_b)$

Output: $c = (c_t \parallel c_b) \bmod^{\pm} q$

```
1 smulbb t0, a,  $\lfloor \frac{2^{26}}{q} \rfloor$ 
2 smulbt t1, a,  $\lfloor \frac{2^{26}}{q} \rfloor$ 
3 asr t0, t0, #26
4 asr t1, t1, #26
5 smulbb t0, t0, q
6 smulbb t1, t1, q
7 pkhbt t0, t0, t1, lsl #16
8 usub16 r, a, t0
```

Algorithm: Improved Packed Barrett Reduction

Input : $a = (a_t \parallel a_b)$

Output: $c = (c_t \parallel c_b) \bmod^{\pm} q$

```
1 smlawb t0,  $-\lfloor \frac{2^{32}}{q} \rfloor$ , a, 2^{15}
2 smlabt t0, q, t0, a
3 smlawt t1,  $-\lfloor \frac{2^{32}}{q} \rfloor$ , a, 2^{15}
4 smulbt t1, q, t1
5 add t1, a, t1, lsl #16
6 pkhbt c, t0, t1, lsl #16
```

► Note: Output range not in $[0, q)$ but $[-\frac{q-1}{2}, \frac{q-1}{2}]$ for odd q

Optimization: Matrix-vector and inner product

Optimization based on technique presented in [Bec+21]:

- ▶ Recall base multiplication for Kyber: Let $\hat{a} = \hat{\mathbf{A}}_{m,n}$, $\hat{s} = \hat{\mathbf{s}}_m$. For $\hat{c} = \hat{a} \circ \hat{s}$

$$\hat{c}_{2i} + \hat{c}_{2i+1}X = (\hat{a}_{2i} + \hat{a}_{2i+1}X)(\hat{s}_{2i} + \hat{s}_{2i+1}X) \bmod (X^2 - \zeta^{2br_7(i)+1}), \text{ with}$$

$$\hat{c}_{2i} = \hat{a}_{2i}\hat{s}_{2i} + \hat{a}_{2i+1}\hat{s}_{2i+1}\zeta^{2br_7(i)+1}$$

$$\hat{c}_{2i+1} = \hat{a}_{2i}\hat{s}_{2i+1} + \hat{s}_{2i}\hat{a}_{2i+1}$$

- ▶ Notice:

$$\begin{bmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix} \cdot \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix}$$

- ▶ Asymmetric multiplication: Cache/Reuse multiplication with twiddle factors during base multiplication.

Optimization: Matrix-vector and inner product

Optimization based on technique presented in [Bec+21]:

- ▶ Recall base multiplication for Kyber: Let $\hat{a} = \hat{\mathbf{A}}_{m,n}$, $\hat{s} = \hat{\mathbf{s}}_m$. For $\hat{c} = \hat{a} \circ \hat{s}$

$$\hat{c}_{2i} + \hat{c}_{2i+1}X = (\hat{a}_{2i} + \hat{a}_{2i+1}X)(\hat{s}_{2i} + \hat{s}_{2i+1}X) \bmod (X^2 - \zeta^{2br_7(i)+1}), \text{ with}$$

$$\hat{c}_{2i} = \hat{a}_{2i}\hat{s}_{2i} + \hat{a}_{2i+1}\hat{s}_{2i+1}\zeta^{2br_7(i)+1}$$

$$\hat{c}_{2i+1} = \hat{a}_{2i}\hat{s}_{2i+1} + \hat{s}_{2i}\hat{a}_{2i+1}$$

- ▶ Notice:

$$\begin{bmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix} \cdot \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix}$$

- ▶ Asymmetric multiplication: Cache/Reuse multiplication with twiddle factors during base multiplication.

Optimization: Matrix-vector and inner product

Optimization based on technique presented in [Bec+21]:

- Recall base multiplication for Kyber: Let $\hat{a} = \hat{\mathbf{A}}_{m,n}$, $\hat{s} = \hat{\mathbf{s}}_m$. For $\hat{c} = \hat{a} \circ \hat{s}$

$$\hat{c}_{2i} + \hat{c}_{2i+1}X = (\hat{a}_{2i} + \hat{a}_{2i+1}X)(\hat{s}_{2i} + \hat{s}_{2i+1}X) \bmod (X^2 - \zeta^{2br_7(i)+1}), \text{ with}$$

$$\hat{c}_{2i} = \hat{a}_{2i}\hat{s}_{2i} + \hat{a}_{2i+1}\hat{s}_{2i+1}\zeta^{2br_7(i)+1}$$

$$\hat{c}_{2i+1} = \hat{a}_{2i}\hat{s}_{2i+1} + \hat{s}_{2i}\hat{a}_{2i+1}$$

- Notice:

$$\begin{bmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix} \cdot \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix}$$

- Asymmetric multiplication: Cache/Reuse multiplication with twiddle factors during base multiplication.

Optimization: Matrix-vector and inner product

Optimization based on technique presented in [Bec+21]:

- Recall base multiplication for Kyber: Let $\hat{a} = \hat{\mathbf{A}}_{m,n}$, $\hat{s} = \hat{\mathbf{s}}_m$. For $\hat{c} = \hat{a} \circ \hat{s}$

$$\hat{c}_{2i} + \hat{c}_{2i+1}X = (\hat{a}_{2i} + \hat{a}_{2i+1}X)(\hat{s}_{2i} + \hat{s}_{2i+1}X) \bmod (X^2 - \zeta^{2br_7(i)+1}), \text{ with}$$

$$\hat{c}_{2i} = \hat{a}_{2i}\hat{s}_{2i} + \hat{a}_{2i+1}\hat{s}_{2i+1}\zeta^{2br_7(i)+1}$$

$$\hat{c}_{2i+1} = \hat{a}_{2i}\hat{s}_{2i+1} + \hat{s}_{2i}\hat{a}_{2i+1}$$

- Notice:

$$\begin{bmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{bmatrix} \cdot \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix} = \begin{bmatrix} \cdot \\ \cdot \\ \cdot \end{bmatrix}$$

- Asymmetric multiplication: Cache/Reuse multiplication with twiddle factors during base multiplication.

Algorithm: Kyber PKE key gen.

Output: public key: $pk = (\hat{\mathbf{t}}, \rho)$

Output: secret key: $sk = (\hat{\mathbf{s}})$

- 1 $\rho, \sigma \in \{0, 1\}^{256} \leftarrow \text{sampleUniform}()$
 - 2 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 3 $\mathbf{s}, \mathbf{e} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{\eta_1}(\sigma)$
 - 4 $\hat{\mathbf{t}} \leftarrow \hat{\mathbf{A}} \circ \text{NTT}(\mathbf{s}) + \text{NTT}(\mathbf{e})$
 - 5 **return** (pk, sk)
-

Algorithm: Kyber PKE encryption

Input : public key: $pk = (\hat{\mathbf{t}}, \rho)$

Input : message: $m \in \mathcal{R}_q$

Input : random coins: $\mu \in \{0, 1\}^{256}$

Output: ciphertext $(\mathbf{u}', \mathbf{v}')$

- 1 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 2 $\mathbf{r} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{\eta_1}(\mu)$
 - 3 $\mathbf{e}_1 \in \mathcal{R}_q^{k \times 1}, \mathbf{e}_2 \in \mathcal{R}_q \leftarrow \text{sampleCBD}^{\eta_2}(\mu)$
 - 4 $\hat{\mathbf{r}} \leftarrow \text{NTT}(\mathbf{r})$
 - 5 $\mathbf{u} \leftarrow \text{iNTT}(\hat{\mathbf{A}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_1$
 - 6 $\mathbf{v} \leftarrow \text{iNTT}(\hat{\mathbf{t}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_2 + m$
 - 7 **return** $(\text{Compress}(\mathbf{u}), \text{Compress}(\mathbf{v}))$
-

Algorithm: Kyber PKE key gen.

Output: public key: $pk = (\hat{\mathbf{t}}, \rho)$

Output: secret key: $sk = (\hat{\mathbf{s}})$

- 1 $\rho, \sigma \in \{0, 1\}^{256} \leftarrow \text{sampleUniform}()$
 - 2 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 3 $\mathbf{s}, \mathbf{e} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{\eta_1}(\sigma)$
 - 4 $\hat{\mathbf{t}} \leftarrow \hat{\mathbf{A}} \circ \text{NTT}(\mathbf{s}) + \text{NTT}(\mathbf{e})$
 - 5 **return** (pk, sk)
-

Algorithm: Kyber PKE encryption

Input : public key: $pk = (\hat{\mathbf{t}}, \rho)$

Input : message: $m \in \mathcal{R}_q$

Input : random coins: $\mu \in \{0, 1\}^{256}$

Output: ciphertext $(\mathbf{u}', \mathbf{v}')$

- 1 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times k} \leftarrow \text{sampleUniform}(\rho)$
 - 2 $\mathbf{r} \in \mathcal{R}_q^{k \times 1} \leftarrow \text{sampleCBD}^{\eta_1}(\mu)$
 - 3 $\mathbf{e}_1 \in \mathcal{R}_q^{k \times 1}, \mathbf{e}_2 \in \mathcal{R}_q \leftarrow \text{sampleCBD}^{\eta_2}(\mu)$
 - 4 $\hat{\mathbf{r}} \leftarrow \text{NTT}(\mathbf{r})$
 - 5 $\mathbf{u} \leftarrow \text{iNTT}(\hat{\mathbf{A}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_1$
 - 6 $\mathbf{v} \leftarrow \text{iNTT}(\hat{\mathbf{t}}^T \circ \hat{\mathbf{r}}) + \mathbf{e}_2 + m$
 - 7 **return** $(\text{Compress}(\mathbf{u}), \text{Compress}(\mathbf{v}))$
-

Better accumulation based on [Chu+21]:

- ▶ Kyber's small prime allows for accumulation without intermediate reductions in the matrix-vector product.

Section 3

Dilithium

- ▶ Three different parameter sets: Dilithium2, Dilithium3, Dilithium5

Table: Overview of Dilithium's parameter sets [Bai+20]

Scheme	NIST level	(k, l)	η	τ	γ_1	γ_2	#reps	pk	sig
Dilithium2	2	(4, 4)	2	39	2^{17}	$(q-1)/88$	4.25	1312 B	2420 B
Dilithium3	3	(6, 5)	4	49	2^{19}	$(q-1)/32$	5.1	1952 B	3293 B
Dilithium5	5	(8, 7)	2	60	2^{19}	$(q-1)/32$	3.85	2592 B	4595 B

- ▶ Same q, n for the three variants
 - ⇒ nice for optimizing
- ▶ In contrast to Kyber, $2n$ -th primitive root of unity exists
 - ⇒ Complete NTT

Algorithm: Dilithium key generation

Output: secret key

$$sk = (\rho, K, tr, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$$

Output: public key $pk = (\rho, \mathbf{t}_1)$

- 1 $\rho, \varsigma, K \in \{0, 1\}^{256} \leftarrow$
sampleUniform();
 - 2 $\mathbf{s}_1 \in [-\eta, \eta]^{l \times 1}, \mathbf{s}_2 \in [-\eta, \eta]^{k \times 1} \leftarrow$
sampleUniform(ς);
 - 3 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times l} \leftarrow$ ExpandA(ρ);
 - 4 $\mathbf{t} \leftarrow$ iNTT($\hat{\mathbf{A}} \circ$ NTT(\mathbf{s}_1)) + \mathbf{s}_2 ;
 - 5 $(\mathbf{t}_1, \mathbf{t}_0) \leftarrow$ Power2Round(\mathbf{t});
 - 6 $tr \in \{0, 1\}^{256} \leftarrow$ CRH($\rho \parallel \mathbf{t}_1$);
 - 7 **return** (pk, sk)
-

Algorithm: Dilithium verification

Input : public key $pk = (\rho, \mathbf{t}_1)$

Input : message: $M \in \{0, 1\}^*$

Input : signature $\sigma = (\mathbf{z}, \mathbf{h}, \tilde{c})$

Output: signature valid or signature
invalid

- 1 $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times l} \leftarrow$ ExpandA(ρ);
 - 2 $c \leftarrow$ SampleInBall(\tilde{c});
 - 3 $\mu \in \{0, 1\}^{512} \leftarrow$ CRH(CRH($\rho \parallel \mathbf{t}_1$) \parallel M);
 - 4 $\mathbf{w}'_1 \leftarrow$ UseHint($\mathbf{h},$ iNTT($\hat{\mathbf{A}} \circ$ NTT(\mathbf{z}) –
NTT(c) \circ NTT($2^d \cdot \mathbf{t}_1$)));
 - 5 **if** $\|\mathbf{z}\|_\infty < \gamma_1 - \beta$ and $\tilde{c} = \text{H}(\mu \parallel \mathbf{w}'_1)$
and # of 1's in $\mathbf{h} \leq \omega$ **then**
 - 6 | **return** signature valid;
 - 7 **else**
 - 8 | **return** signature invalid;
-

Algorithm: Dilithium signing

```

Input : secret key  $sk = (\rho, K, tr, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$ 
Input : message:  $M \in \{0, 1\}^*$ 
Output: signature  $\sigma = (\mathbf{z}, \mathbf{h}, \tilde{c})$ 
1  $\hat{\mathbf{A}} \in \mathcal{R}_q^{k \times l} \leftarrow \text{ExpandA}(\rho);$ 
2  $\mu \in \{0, 1\}^{512} \leftarrow \text{CRH}(tr \| M);$ 
3  $\kappa \leftarrow 0, (\mathbf{z}, \mathbf{h}) \leftarrow \perp;$ 
4  $\rho' \in \{0, 1\}^{512} \leftarrow \text{CRH}(K \| \mu);$ 
5  $\hat{\mathbf{s}}_1 \leftarrow \text{NTT}(\mathbf{s}_1), \hat{\mathbf{s}}_2 \leftarrow \text{NTT}(\mathbf{s}_2), \hat{\mathbf{t}}_0 := \text{NTT}(\mathbf{t}_0);$ 
6 while  $(\mathbf{z}, \mathbf{h}) = \perp$  do
7    $\mathbf{y} \in \mathcal{R}_q^{l \times 1} \leftarrow \text{ExpandMask}(\rho', \kappa);$ 
8    $\mathbf{w} \leftarrow \text{iNTT}(\hat{\mathbf{A}} \circ \text{NTT}(\mathbf{y}));$ 
9    $\mathbf{w}_1 \leftarrow \text{HighBits}(\mathbf{w}, 2\gamma_2);$ 
10   $\tilde{c} \leftarrow \text{H}(\mu \| \mathbf{w}_1);$ 
11   $c \leftarrow \text{SampleInBall}(\tilde{c});$ 
12   $\hat{c} \leftarrow \text{NTT}(c);$ 
13   $\mathbf{z} \leftarrow \mathbf{y} + \text{iNTT}(\hat{c} \circ \hat{\mathbf{s}}_1);$ 
14   $\mathbf{r}_0 \leftarrow \text{LowBits}(\mathbf{w} - \text{iNTT}(\hat{c} \circ \hat{\mathbf{s}}_2), 2\gamma_2);$ 
15  if  $\|\mathbf{z}\|_\infty \geq \gamma_1 - \beta$  or  $\|\mathbf{r}_0\|_\infty \geq \gamma_2 - \beta$  then
16     $(\mathbf{z}, \mathbf{h}) \leftarrow \perp$ 
17  else
18     $\mathbf{h} \leftarrow \text{MakeHint}(-\text{iNTT}(\hat{c} \circ \hat{\mathbf{t}}_0), \mathbf{w} -$ 
19       $\text{iNTT}(\hat{c} \circ \hat{\mathbf{s}}_2 + \text{iNTT}(\hat{c} \circ \hat{\mathbf{t}}_0)), 2\gamma_2);$ 
20    if  $\|\text{iNTT}(\hat{c} \circ \hat{\mathbf{t}}_0)\|_\infty \geq \gamma_2$  or # of 1's in
       $\mathbf{h} > \omega$  then
21       $(\mathbf{z}, \mathbf{h}) \leftarrow \perp$ 
22   $\kappa \leftarrow \kappa + l;$ 
23 return  $\sigma = (\mathbf{z}, \mathbf{h}, \tilde{c})$ 
    
```

Similar techniques as for Kyber:

- ▶ Better layer merging: Merge layers 7–5, 4–2, 1–0, instead of 7–6, 5–4, 3–2, 1–0.
- ▶ CT-Butterflies for i NTT

Similar techniques as for Kyber:

- ▶ Better layer merging: Merge layers 7–5, 4–2, 1–0, instead of 7–6, 5–4, 3–2, 1–0.
- ▶ CT-Butterflies for i NTT

Optimization: Small NTTs

- ▶ Recall: c consists of τ -many ± 1 s, $\mathbf{s}_1, \mathbf{s}_2$ have elements in $[-\eta, \eta]$
 - $\Rightarrow c\mathbf{s}_1$ and $c\mathbf{s}_2$ bounded by $\tau\eta$
 - \Rightarrow Regard computation as in $\mathbb{Z}_{q'}$ with $q' > 2\tau\eta$ [Chu+21]
- ▶ Some freedom for choosing q'

Table: Choosing q'

Scheme	η	τ	$2\tau\eta$	q'
Dilithium2	2	39	156	$F_3 = 257$
Dilithium3	4	49	392	769
Dilithium5	2	60	240	$F_3 = 257$

Optimization: Small NTTs

- ▶ Recall: c consists of τ -many ± 1 s, $\mathbf{s}_1, \mathbf{s}_2$ have elements in $[-\eta, \eta]$
 - ⇒ $c\mathbf{s}_1$ and $c\mathbf{s}_2$ bounded by $\tau\eta$
 - ⇒ Regard computation as in $\mathbb{Z}_{q'}$ with $q' > 2\tau\eta$ [Chu+21]
- ▶ Some freedom for choosing q'

Table: Choosing q'

Scheme	η	τ	$2\tau\eta$	q'
Dilithium2	2	39	156	$F_3 = 257$
Dilithium3	4	49	392	769
Dilithium5	2	60	240	$F_3 = 257$

Optimization: Small NTTs

- ▶ Recall: c consists of τ -many ± 1 s, $\mathbf{s}_1, \mathbf{s}_2$ have elements in $[-\eta, \eta]$
 - ⇒ $c\mathbf{s}_1$ and $c\mathbf{s}_2$ bounded by $\tau\eta$
 - ⇒ Regard computation as in $\mathbb{Z}_{q'}$ with $q' > 2\tau\eta$ [Chu+21]
- ▶ Some freedom for choosing q'

Table: Choosing q'

Scheme	η	τ	$2\tau\eta$	q'
Dilithium2	2	39	156	$F_3 = 257$
Dilithium3	4	49	392	769
Dilithium5	2	60	240	$F_3 = 257$

Optimization: Small NTTs

- ▶ Recall: c consists of τ -many ± 1 s, $\mathbf{s}_1, \mathbf{s}_2$ have elements in $[-\eta, \eta]$
 - ⇒ $c\mathbf{s}_1$ and $c\mathbf{s}_2$ bounded by $\tau\eta$
 - ⇒ Regard computation as in $\mathbb{Z}_{q'}$ with $q' > 2\tau\eta$ [Chu+21]
- ▶ Some freedom for choosing q'

Table: Choosing q'

Scheme	η	τ	$2\tau\eta$	q'
Dilithium2	2	39	156	$F_3 = 257$
Dilithium3	4	49	392	769
Dilithium5	2	60	240	$F_3 = 257$

FNT for Dilithium2 and Dilithium3

- ▶ CT-Butterfly: $(a, b) \mapsto (a + \omega b, a - \omega b)$ can be implemented with `m1a` and `m1s`
- ▶ First $t = 3$ layers have power of two twiddle factor
 - ⇒ Efficient implementation without loading using barrel shifter and $\log \omega$ as twiddle factor

Algorithm: `CT_FNT(a, b, logW)`.

Input : $(a, b) = (a, b)$

Output: $(a, b) = (a + 2^{\log W} b, a - 2^{\log W} b)$

- 1 `add a, a, b, lsl #logW;`
 - 2 `sub b, a, b, lsl #(logW+1);`
-

- ▶ Incompatible with FNT over F_3
- ▶ Kyber NTT/iNTT with $q' = 769$ and most reductions left out
- ▶ Experiments with $q' = F_4 = 65537$ yielding no speed-up over $q' = 769$
- ▶ $q' = 3329$ also possible for code re-use

Section 4

Results

Benchmarking Setup

- ▶ Based on pqm4
- ▶ Clock reduced to 24 MHz
- ▶ `arm-none-eabi-gcc` version 10.2.1 with `-O3`
- ▶ Keccak from pqm4
- ▶ Randomness from hardware RNG

Benchmarking Setup

- ▶ Based on pqm4
- ▶ Clock reduced to 24 MHz
- ▶ `arm-none-eabi-gcc` version 10.2.1 with `-O3`
- ▶ Keccak from pqm4
- ▶ Randomness from hardware RNG

Benchmarking Setup

- ▶ Based on pqm4
- ▶ Clock reduced to 24 MHz
- ▶ `arm-none-eabi-gcc` version 10.2.1 with `-O3`
- ▶ Keccak from pqm4
- ▶ Randomness from hardware RNG

Benchmarking Setup

- ▶ Based on pqm4
- ▶ Clock reduced to 24 MHz
- ▶ `arm-none-eabi-gcc` version 10.2.1 with `-O3`
- ▶ Keccak from pqm4
- ▶ Randomness from hardware RNG

Benchmarking Setup

- ▶ Based on pqm4
- ▶ Clock reduced to 24 MHz
- ▶ `arm-none-eabi-gcc` version 10.2.1 with `-O3`
- ▶ Keccak from pqm4
- ▶ Randomness from hardware RNG

NTT related functions

Table: Cycle counts for transformation operations of Kyber and Dilithium. NTT and iNTT correspond to the schemes default transformations, i.e., $q = 3329$ for Kyber and $q = 8380417$ for Dilithium. The NTT with $q = 257$ is deployed for Dilithium2 and Dilithium5, and the NTT with $q = 769$ is used used for Dilithium3.

	Prime	Implementation	NTT	iNTT	basemul
Kyber	$q = 3329$	[Alk+20]	6 852	6 979	2 317
		This work	5 992	5 491/6 282 ^a	1 613 ^b
Dilithium	$q = 8380417$	[GKS20]	8 540	8 923	1 955
		This work	8 093	8 415	1 955
	$q = 257$	This work	5 524	5 563	1 225
	$q = 769$	[Abd+21] (6-layer)	4 852	4 817	2 966
This work		5 200	5 537	1 740	

^a First value is for speed-optimization, second for stack-optimization.

^b Asymmetric basemul as used in the IP (enc). As the basemul in the MVP and IP consists of individual function calls, the cycle count is not straight forward to measure.

Kyber: Matrix-vector and inner product

Table: Cycle counts for matrix-vector and inner products used in Kyber.

implementation	variant	operation	Kyber-512	Kyber-768	Kyber-1024
pqm4		Matrix-Vector Product ^a	66 291	127 634	209 517
		Matrix-Vector Product ^b	226 580	484 077	840 498
		Inner Product (enc)	11 978	14 696	17 429
		Inner Product (dec)	29 888	41 910	53 792
This work	speed	Matrix-Vector Product ^a	55 746	106 380	172 152
		Matrix-Vector Product ^b	211 606	457 213	796 349
		Inner Product (enc)	8 762	10 331	11 898
		Inner Product (dec)	23 425	32 354	41 275
	stack	Matrix-Vector Product ^a	58 028	112 503	184 149
		Matrix-Vector Product ^b	214 053	463 590	808 206
		Inner Product (enc)	11 218	13 877	16 733
		Inner Product (dec)	24 722	34 167	43 619

^a Measurement excluding the hashing.

^b Measurement including the hashing.

Kyber: Scheme performance

Table: Cycle counts and stack usage for Kyber for the key generation, encapsulation, and decapsulation. Cycle counts are averaged over 100 executions.

implementation	variant	Kyber-512		Kyber-768		Kyber-1024		
		cc	stack [B]	cc	stack [B]	cc	stack [B]	
pqm4, [Alk+20]	K	458k	2 220	745k	3 100	1 188k	3 612	
	E	553k	2 308	899k	2 780	1 373k	3 292	
	D	513k	2 324	839k	2 804	1 294k	3 324	
This work	speed	K	443k	4 272	718k	5 312	1 138k	6 336
		E	536k	5 376	870k	6 416	1 324k	7 432
		D	487k	5 384	796k	6 432	1 227k	7 448
	stack	K	444k	2 220	724k	2 736	1 149k	3 256
		E	540k	2 308	879k	2 808	1 341k	3 328
		D	492k	2 324	807k	2 824	1 246k	3 352

Dilithium: Scheme performance

Table: Cycle counts and stack usage for Dilithium. K, S, and V correspond to the key generation, signature generation, and signature verification. Cycle counts are averaged over 10000 executions.

implementation	variant	Dilithium2		Dilithium3		Dilithium5		
		cc	stack [B]	cc	stack [B]	cc	stack [B]	
pqm4, [GKS20]	K	1 602k	38k	2 835k	61k	4 836k	98k	
	S	4 336k	49k	6 721k	74k	9 037k	115k	
	V	1 579k	36k	2 700k	58k	4 718k	93k	
This work	speed	K	1 596k	8 508	2 827k	9 540	4 829k	11 696
		S	4 093k	49k	6 623k	69k	8 803k	116k
		V	1 572k	36k	2 692k	58k	4 707k	93k

Faster Kyber and Dilithium on the Cortex-M4

Amin Abdulrahman^{1,2} Vincent Hwang^{3,4} Matthias J. Kannwischer³ Amber
Sprenkels⁵

¹Ruhr University Bochum, Germany

²Max Planck Institute for Security and Privacy, Bochum, Germany

³Academia Sinica, Taipei, Taiwan

⁴National Taiwan University, Taipei, Taiwan

⁵Digital Security Group, Radboud University, Nijmegen, The Netherlands

23rd June 2022

Bibliography I

- [Abd+21] Amin Abdulrahman et al. *Multi-moduli NTTs for Saber on Cortex-M3 and Cortex-M4*. Cryptology ePrint Archive, Report 2021/995. <https://ia.cr/2021/995>. 2021.
- [Alk+20] Erdem Alkim et al. 'Cortex-M4 Optimizations for $\{R,M\}$ LWE Schemes'. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems 2020.3* (June 2020), pp. 336–357. DOI: 10.13154/tches.v2020.i3.336–357. URL: <https://tches.iacr.org/index.php/TCHES/article/view/8593>.
- [Bai+20] Shi Bai et al. *CRYSTALS-Dilithium: Algorithm Specifications And Supporting Documentation (version 3.0)*. Submission to round 3 of the NIST post-quantum project [Nat]. Oct. 2020.
- [Bec+21] Hanno Becker et al. *Neon NTT: Faster Dilithium, Kyber, and Saber on Cortex-A72 and Apple M1*. Cryptology ePrint Archive, Report 2021/986. <https://ia.cr/2021/986>. 2021.

- [BKS19] Leon Botros, Matthias J. Kannwischer, and Peter Schwabe. ‘Memory-Efficient High-Speed Implementation of Kyber on Cortex-M4’. In: *Progress in Cryptology – AFRICACRYPT 2019*. Ed. by Johannes Buchmann, Abderrahmane Nitaj, and Tajjeeddine Rachidi. Cham: Springer International Publishing, 2019, pp. 209–228. ISBN: 978-3-030-23696-0.
- [Chu+21] Chi-Ming Marvin Chung et al. ‘NTT Multiplication for NTT-unfriendly Rings: New Speed Records for Saber and NTRU on Cortex-M4 and AVX2’. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems 2021.2* (Feb. 2021), pp. 159–188. DOI: 10.46586/tches.v2021.i2.159–188. URL: <https://tches.iacr.org/index.php/TCHES/article/view/8791>.

- [GKS20] Denisa O. C. Greconici, Matthias J. Kannwischer, and Amber Sprenkels. ‘Compact Dilithium Implementations on Cortex-M3 and Cortex-M4’. In: *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2021.1 (Dec. 2020), pp. 1–24. DOI: 10.46586/tches.v2021.i1.1-24. URL: <https://tches.iacr.org/index.php/TCHES/article/view/8725>.
- [Nat] National Institute of Standards and Technology. *Post-Quantum Cryptography Standardization Project*. Accessed: 2021-04-04. URL: <https://csrc.nist.gov/projects/post-quantum-cryptography>.

