



Towards ML-KEM & ML-DSA on OpenTitan

Amin Abdulrahman¹ Felix Oberhansl² Hoang Nguyen Hien Pham^{3,4}
Jade Philipoom^{5,6} Peter Schwabe^{1,7} Tobias Stelzer² Andreas Zankl^{2,8}

August 23, 2024

¹Max Planck Institute for Security and Privacy (MPI-SP)

²Fraunhofer Institute for Applied and Integrated Security (AISEC)

³BULL SAS ⁴Université Grenoble Alpes

⁵zeroRISC ⁶OpenTitan

⁷Radboud University ⁸Technical University of Munich (TUM)

Introduction

Asymmetric Cryptography at Risk



Figure 1: IBM Quantum System One¹

¹https://www.flickr.com/photos/ibm_research_zurich/51248690716/

- July 2016: NIST Post-Quantum Cryptography (PQC) project

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

Post-Quantum Cryptography

- July 2016: NIST Post-Quantum Cryptography (PQC) project
- July 2022: NIST announces schemes for standardization

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

Post-Quantum Cryptography

- July 2016: NIST Post-Quantum Cryptography (PQC) project
- July 2022: NIST announces schemes for standardization
 - **KEM:** KYBER²
 - **Digital Signatures:** DILITHIUM³, SPHINCS⁺⁴, FALCON⁵

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

Post-Quantum Cryptography

- July 2016: NIST Post-Quantum Cryptography (PQC) project
- July 2022: NIST announces schemes for standardization
 - **KEM:** *KYBER*²
 - **Digital Signatures:** *DILITHIUM*³, *SPHINCS+*⁴, *FALCON*⁵

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

Post-Quantum Cryptography

- July 2016: NIST Post-Quantum Cryptography (PQC) project
- July 2022: NIST announces schemes for standardization
 - **KEM:** *KYBER*²
 - **Digital Signatures:** *DILITHIUM*³, *SPHINCS+*⁴, *FALCON*⁵
- Draft standards: FIPS 203 (ML-KEM) and FIPS 204 (ML-DSA)

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

Post-Quantum Cryptography

- July 2016: NIST Post-Quantum Cryptography (PQC) project
- July 2022: NIST announces schemes for standardization
 - **KEM:** *KYBER*²
 - **Digital Signatures:** *DILITHIUM*³, *SPHINCS*⁺⁴, *FALCON*⁵
- Draft standards: FIPS 203 (*ML-KEM*) and FIPS 204 (*ML-DSA*)
- PQC is going real-world: Integration into TLS by Google, Cloudflare, and Mozilla

²[SAB+22]

³[LDK+22]

⁴[HBD+22]

⁵[PFH+22]

What's inside?

- Vastly more complex than traditional, asymmetric cryptography

What's inside?


- Vastly more complex than traditional, asymmetric cryptography
 - Common thread: Hashing and polynomial arithmetic
 - Heavy usage of SHAKE and SHA-3
 - Polynomial arithmetic over $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$, where $n = 256$, $q = 3329$ for ML-KEM, and $q = 8380417$ for ML-DSA
- Coefficients are “small” integers

What's inside?

- Vastly more complex than traditional, asymmetric cryptography
- Common thread: Hashing and polynomial arithmetic
 - Heavy usage of SHAKE and SHA-3
 - Polynomial arithmetic over $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$, where $n = 256$, $q = 3329$ for ML-KEM, and $q = 8380417$ for ML-DSA
 - Coefficients are “small” integers
- Efficient polynomial arithmetic achieved through NTT-based multiplication
 - Variant of the discrete Fourier transform over finite fields


NTT-based polynomial multiplication

- Fast NTT with Cooley-Tukey (CT) or Gentleman-Sande (GS) FFT algorithms in $\mathcal{O}(n \log n)$
 - Divide & Conquer approach
 - $\log n$ divide-steps with 128 parallel “butterfly” operations on coefficient pair




NTT-based polynomial multiplication

- Fast NTT with Cooley-Tukey (CT) or Gentleman-Sande (GS) FFT algorithms in $\mathcal{O}(n \log n)$
 - Divide & Conquer approach
 - $\log n$ divide-steps with 128 parallel “butterfly” operations on coefficient pair




→ Central operations: Addition, subtraction, (modular) multiplication with a constant

Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg


- Reduced RV32 instruction set

Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg

OpenTitan OTBN


- Reduced RV32 instruction set
- Big number instruction set

Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg

OpenTitan OTBN


- Reduced RV32 instruction set
- Big number instruction set
- 32 256-bit wide WDRs

Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg

OpenTitan OTBN

- Reduced RV32 instruction set
- Big number instruction set
- 32 256-bit wide WDRs
- `bn.addm`
`/bn.subm`

Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg

OpenTitan OTBN

- Reduced RV32 instruction set
- Big number instruction set
- 32 256-bit wide WDRs
- `bn.addm`
`/bn.subm`
- 64×64 -bit multiply-acc. unit




Figure 2: OTBN block architecture⁶

¹https://opentitan.org/book/hw/ip/otbn/doc/otbn_blockarch.svg

Implementing ML-KEM & ML-DSA on OTBN

Approach

Establish a baseline for ML-KEM and ML-DSA on OTBN

- Using state-of-the-art implementation techniques
- Optimization tailored to the architecture
 - Make use of WDRs whenever possible
 - Leverage `bn.addm/bn.subm`
- Only modification required: More memory

Profiling on OTBN

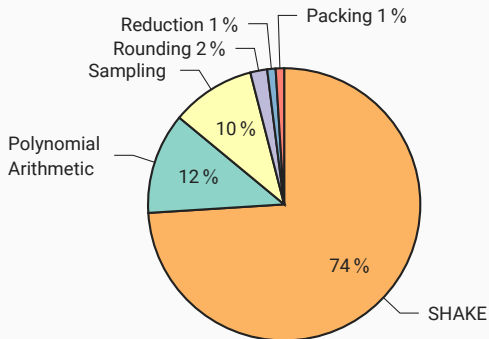


Figure 3: Profiling of ML-DSA-65 verification on OTBN.

Extensions to OTBN

Leveraging the KMAC Block

- High-speed, high-security Keccak accelerator available
 - 4 cycles/round
 - Masked
- Implementation in the OTBN Python simulator by Philipoom⁷
- Interfacing through WSRs
- We call OTBN with KMAC OTBN^{KMAC}

⁷<https://github.com/jadephilipoom/opentitan/commit/e86be3446204f439c41c142b077a4ca8b449b1c9>

Profiling on OTBN^{KMAC}




Figure 4: Profiling of ML-DSA-65 verification on OTBN^{KMAC}.

Challenges

```
1 bn.and coeffa, coeffs, consts >> 192 /* Mask out coeffs from buffer*/
2 bn.and coeffb, coeffs, consts >> 192
3
4 /* Plantard multiplication: Twiddle * coeffb */
5 bn.mulqacc.wo.z coeffb, coeffb.0, twiddle.0, 192 /* (coeffb*R) */
6 bn.add          coeffb, consts, coeffb >> 160 /* +1 */
7 bn.mulqacc.wo.z coeffb, coeffb.1, consts.2, 0 /* *q */
8 bn.rshi         wtmp, consts, coeffb >> 32 /* >> d */
9 /* Butterfly */
10 bn.subm coeffb, coeffa, wtmp
11 bn.addm coeffa, coeffa, wtmp
12
13 bn.rshi coeffs, coeffa, coeffs >> 32 /* Shift results to buffer */
14 bn.rshi coeffs, coeffb, coeffs >> 32
```

Listing 1: CT butterfly on OTBN.

Challenges

```
1 bn.and coeffa, coeffs, consts >> 192 /* Mask out coeffs from buffer */
2 bn.and coeffb, coeffs, consts >> 192
3
4 /* Plantard multiplication: Twiddle * coeffb */
5 bn.mulqacc.w0.z coeffb, coeffb.0, twiddle.0, 192 /* (coeffb*R) */
6 bn.add coeffb, coeffb, consts, coeffb >> 160 /* +1 */
7 bn.mulqacc.w0.z coeffb, coeffb.1, consts.2, 0 /* *q */
8 bn.rshi wtmp, coeffb, consts, coeffb >> 32 /* >> d */
9 /* Butterfly */
10 bn.subm coeffb, coeffa, wtmp
11 bn.addm coeffa, coeffa, wtmp
12
13 bn.rshi coeffs, coeffa, coeffs >> 32 /* Shift results to buffer */
14 bn.rshi coeffs, coeffb, coeffs >> 32
```

Listing 1: CT butterfly on OTBN.

- High data-movement overhead
- No SIMD capabilities
- Generally: ISA not made for arithmetic on small integers

Speeding-up polynomial arithmetic

Approach

Make better use of 256-bit wide registers.

→ Interpret WDRs as vectors of smaller elements, e.g.,
 16×16 -bit or 8×32 -bit.




Figure 5: Vectorized (modular) arithmetic: `bn.{addv, subv, mulv}{m}`




Figure 5: Vectorized (modular) arithmetic: `bn.{addv, subv, mulv}{m}`

Example

```
bn.mulvm.8S w0, w1, w2
```

bn.trn1, bn.trn2




Figure 6: Vector transpose (odd/even indices): `bn.trn1`, `bn.trn2`

bn.trn1, bn.trn2




Figure 6: Vector transpose (odd/even indices): `bn.trn1`, `bn.trn2`

Example

`bn.trn1.8S w0, w1, w2`




Figure 7: Vectorized {right,left} bit shift: `bn.shv`




Figure 7: Vectorized {right,left} bit shift: `bn.shv`

Example

```
bn.shv.8S w0, w1 >> 3
```

Impact of the Extensions

Evaluating OTBN^{KMAC}_{Ext.}

```
1 bn.mulvm.1.8S tmp, vec8, twiddles, 0
2 bn.subvm.8S   vec8, vec0, tmp
3 bn.addvm.8S   vec0, vec0, tmp
```

Listing 2: CT butterfly on OTBN^{KMAC}_{Ext.}.

Profiling on OTBN^{KMAC}




Figure 4: Profiling of ML-DSA-65 verification on OTBN^{KMAC}.

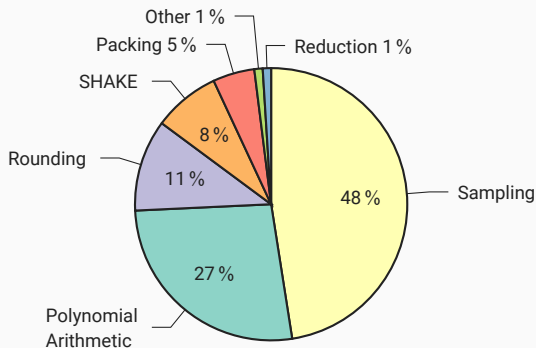


Figure 8: Profiling of ML-DSA-65 verification on OTBN^{KMAC}_{Ext.}.

Full Scheme Benchmarks

Table 1: ML-DSA-65 full scheme benchmarks. All numbers refer to cycles. Median result was selected, if given. 10 000 iterations for our measurements.

Platform	Key Gen.	Sign	Verify
OTBN	2 190 278 ($\times 8.39$)	4 490 766 ($\times 6.44$)	2 107 440 ($\times 8.22$)
OTBN ^{KMAC}	438 154 ($\times 1.68$)	1 842 696 ($\times 2.64$)	493 307 ($\times 1.92$)
OTBN ^{KMAC} _{Ext.}	261 000 ($\times 1.00$)	697 203 ($\times 1.00$)	256 327 ($\times 1.00$)
OpenTitan [SOSK23] ^{b,c}	— —	— —	1 488 526 ($\times 5.81$)
Skylake [LDK+22] ^a	154 308 ($\times 0.59$)	342 708 ($\times 0.49$)	154 622 ($\times 0.60$)
Cortex-M4 [HAZ+24] ^a	2 390 080 ($\times 9.16$)	4 878 759 ($\times 7.00$)	2 289 269 ($\times 8.93$)

^a Own benchmarks.

^b Including modified variant of OTBN, parts of the execution on Ibex Core.

^c Round 3 DILITHIUM.

Conclusion

Five new instruction classes:

- `bn.addv{m}{.8S, .16H}`
- `bn.subv{m}{.8S, .16H}`
- `bn.mulv{m}{.8S, .16H, .1}`
- `bn.trn1{.2Q, .4D, .8S, .16H}, bn.trn2{.2Q, .4D, .8S, .16H}`
- `bn.shv{.8S, .16H}`

Conclusion

Five new instruction classes:

- `bn.addv{m}{.8S, .16H}`
- `bn.subv{m}{.8S, .16H}`
- `bn.mulv{m}{.8S, .16H, .1}`
- `bn.trn1{.2Q, .4D, .8S, .16H}, bn.trn2{.2Q, .4D, .8S, .16H}`
- `bn.shv{.8S, .16H}`

Result

Longer critical path, but: We achieve speed-ups of a **factor of 6 to 9** with only **11% area overhead** for OTBN and not even **2% for Top-Earlgrey**.

- [HAZ+24] Junhao Huang et al. “Revisiting Keccak and Dilithium Implementations on ARMv7-M”. In: *IACR TCHES* 2024.2 (2024), pp. 1–24. DOI: [10.46586/tches.v2024.i2.1-24](https://doi.org/10.46586/tches.v2024.i2.1-24).
- [HBD+22] Andreas Hülsing et al. *SPHINCS⁺*. Tech. rep. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. National Institute of Standards and Technology, 2022.
- [LDK+22] Vadim Lyubashevsky et al. *CRYSTALS-DILITHIUM*. Tech. rep. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. National Institute of Standards and Technology, 2022.

- [PFH+22] Thomas Prest et al. *FALCON*. Tech. rep. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. National Institute of Standards and Technology, 2022.
- [SAB+22] Peter Schwabe et al. *CRYSTALS-KYBER*. Tech. rep. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. National Institute of Standards and Technology, 2022.

- [SOSK23] Tobias Stelzer et al. “Enabling Lattice-Based Post-Quantum Cryptography on the OpenTitan Platform”. In: *Proceedings of the 2023 Workshop on Attacks and Solutions in Hardware Security*. ASHES '23. <https://dl.acm.org/doi/10.1145/3605769.3623993>. New York, NY, USA: Association for Computing Machinery, Nov. 2023, pp. 51–60. ISBN: 9798400702624. DOI: [10.1145/3605769.3623993](https://doi.org/10.1145/3605769.3623993). (Visited on 11/30/2023).