High-Speed Hardware Implementations of Post-Quantum Cryptography Digital Signature Schemes

Luke Beckwith, Robert Wallace,
Duc T. Nguyen, Kamyar Mohajerani,
and Kris Gaj

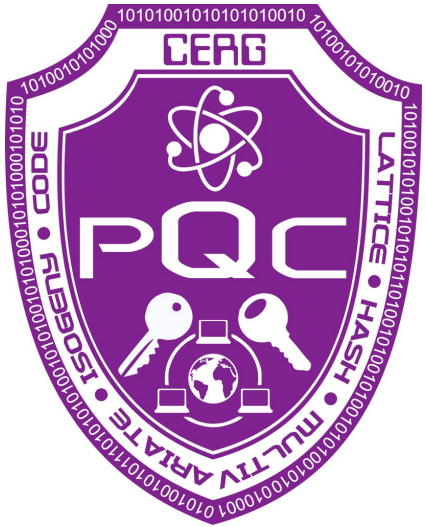

GEORGE MASON UNIVERSITY

Co-Authors

Luke Beckwith

Robert Wallace

Duc T. Nguyen

Kamyar Mohajerani

Post-Quantum Cryptography (PQC)

- Public-key cryptographic algorithms for which there are no known attacks using quantum computers
  - Capable of being implemented using any traditional methods, including software and hardware
  - Running efficiently on any modern computing platforms: PCs, tablets, smartphones, servers with FPGA accelerators, etc.
- Based entirely on traditional semiconductor VLSI technology!

The biggest revolution in cryptography since the invention of public-key cryptography in 1970s!!!

Types and Security Levels of PQC Schemes

Types

1. Public Key Encryption (PKE)
2. Key Encapsulation Mechanism (KEM)
3. Digital Signature (DS)

Security Levels

Level	Standard w/ Equivalent Security
1	AES-128
2	SHA-256
3	AES-192
4	SHA-384
5	AES-256

Only Level 5 schemes allowed in the U.S. National Security Systems

NIST PQC Standardization Process

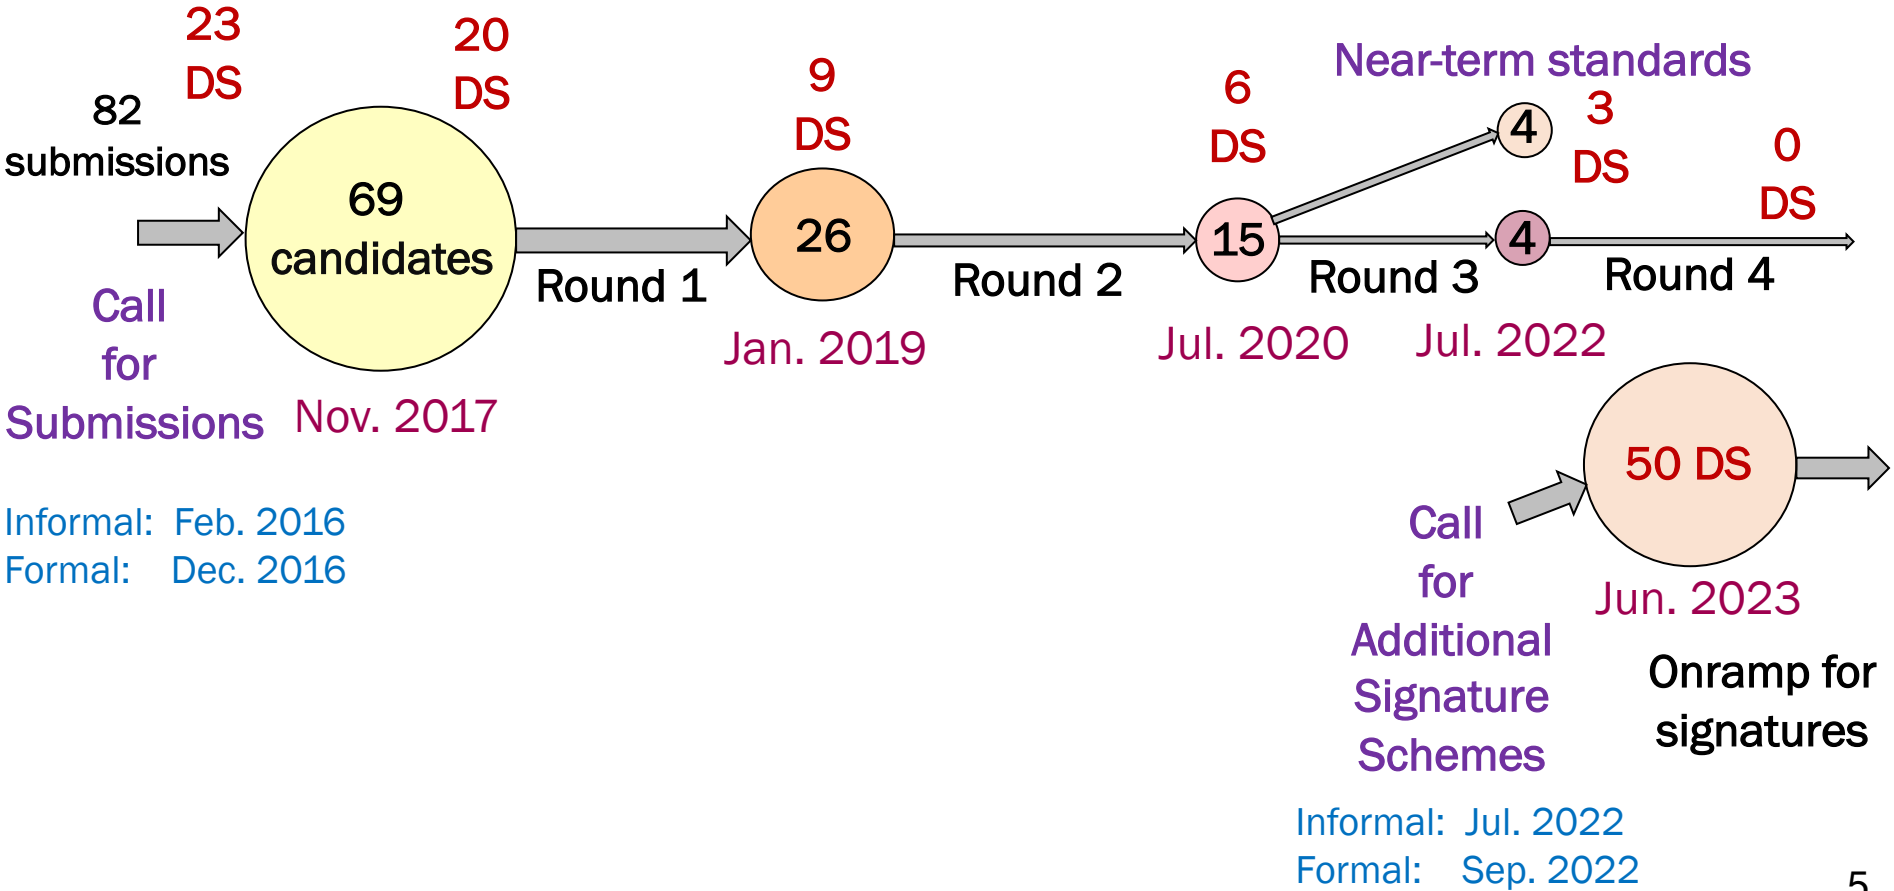

Focus of this Talk

Near-Term Digital Signature Standards

Lattice-based

CRYSTALS-DILITHIUM

FALCON

Symmetric-based (hash-based)

SPHINCS+

Candidate in a New Onramp Signature Process

Code-based

LESS:
Linear Equivalence Signature Scheme

<

https://www.less-project.com

CRYSTALS-Dilithium: High-Level Architecture

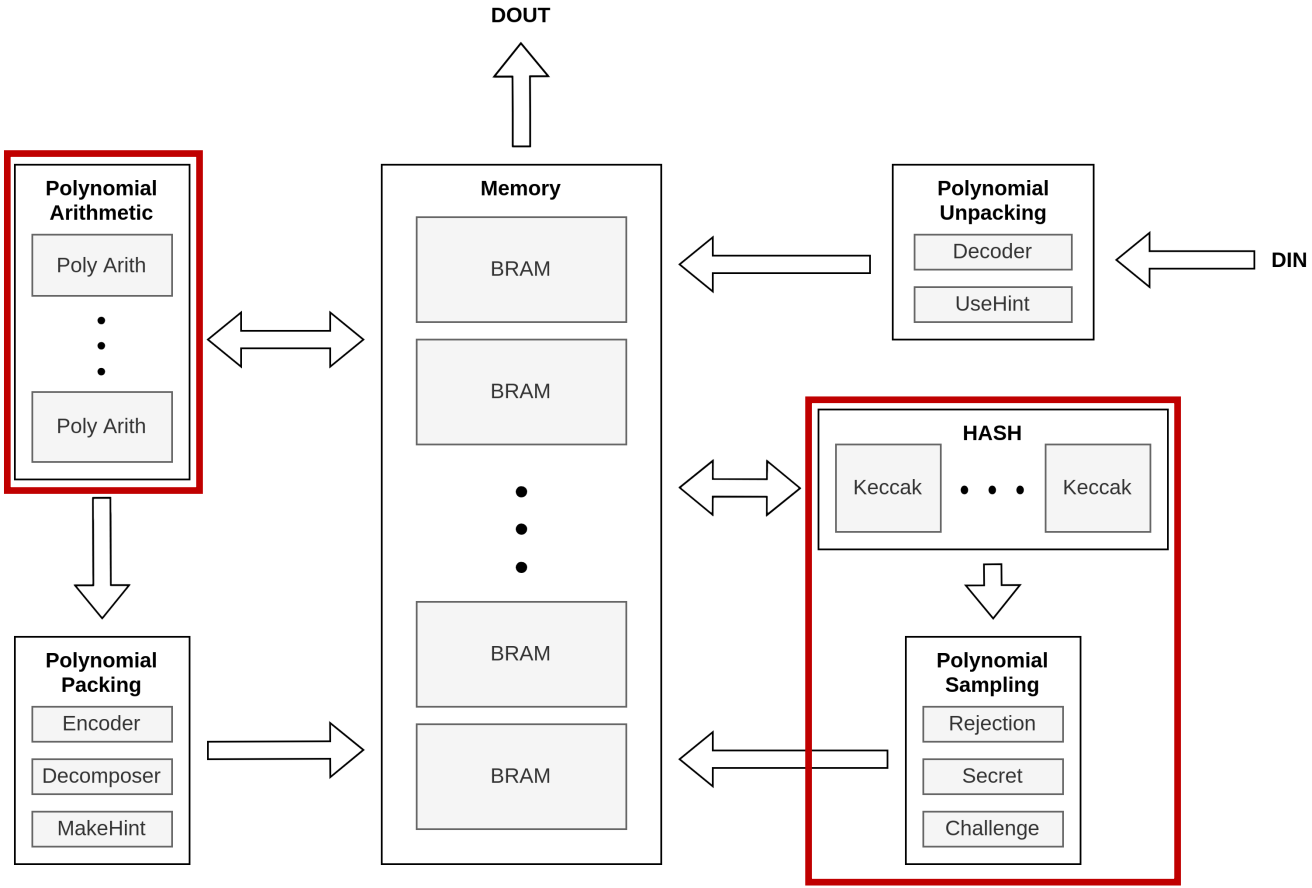

Polynomial Sampling – Bottleneck no. 1

- Core elements are matrices and vectors of polynomials
- Sampled uniformly from the output of the SHA-3 extendable-output function (XOF) SHAKE

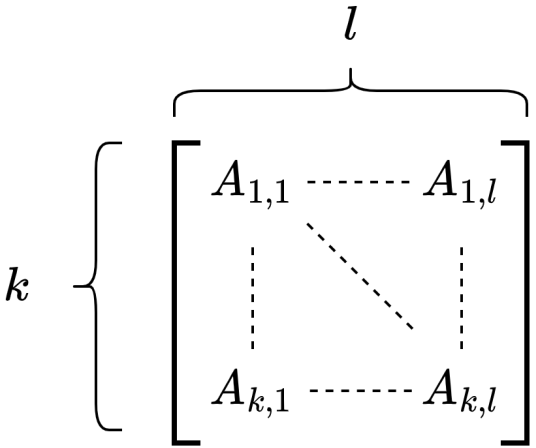

$$A_{i,j} \in R_q = a_{N-1}x^{N-1} + \ldots + a_1 x^1 + a_0$$

where R_q is the ring $Z_q[X]/(X^N + 1)$

Polynomial Sampling - Number of bytes per matrix

Parameter	Dilithium2	Dilithium3	Dilithium5
Dimensions (k, l)	(4, 4)	(6, 5)	(8, 7)
#Polynomial coefficients (N)	256		
Total number of coefficients $(k \cdot l \cdot N)$	4096	7680	14336
Coefficient Modulus (q)	8,380,417 = $2^{23} - 2^{13} + 1$		
Total number of bytes	12 kbytes	23 kbytes	43 kbytes

Solution – Using multiple (3-4) Keccak cores in High-Speed Implementations

Polynomial Arithmetic - Dilithium Bottleneck no. 2

- Performs polynomial multiplication using the Number Theoretic Transform (NTT)
- $C(x) = A(x) \cdot B(x) = NTT^{-1}(NTT(A) \circ NTT(B)) = NTT^{-1}(\hat{A} \circ NTT(B))$,
  where $\hat{A} = A$ in the NTT domain, $\circ$ is a point-wise multiplication

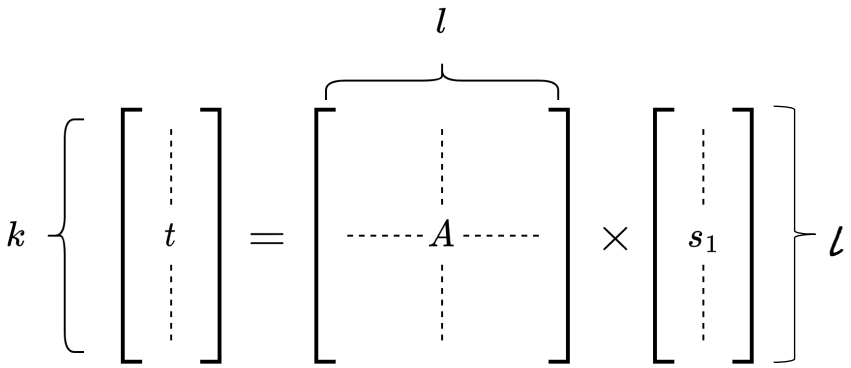

Matrix A is sampled in the NTT-domain

- Input vector s_1 must be converted using the NTT
- Result vector t must be converted back using the Inverse NTT

l	NTT transformations
k	inverse NTT transformations
$k \cdot l$	point-wise multiplications
$k \cdot (l-1)$	additions

Dilithium: Polynomial Arithmetic – NTT Unit

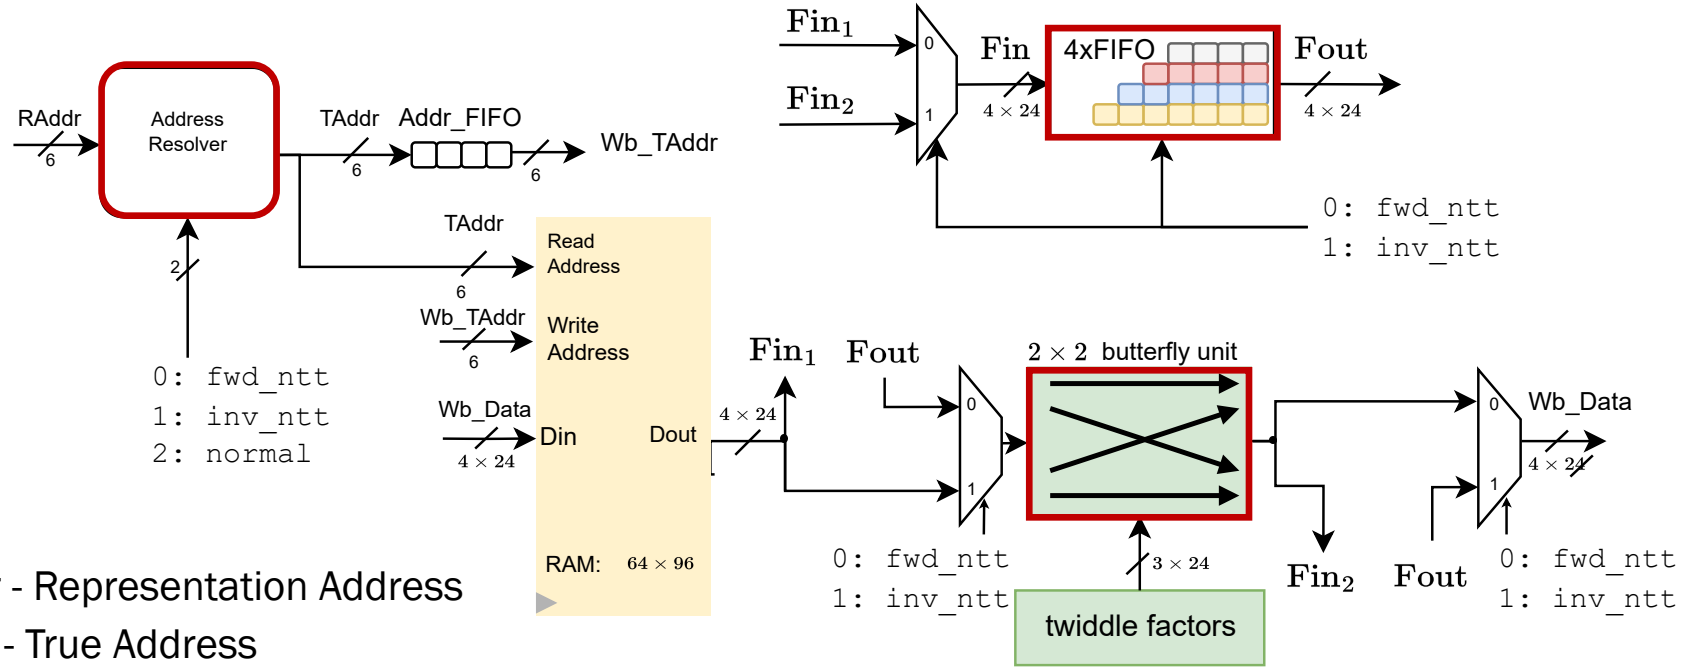

RAddr - Representation Address

TAddr - True Address

Number of clock cycles proportional to $\mathbf{O}(N\ log(N))$

Dilithium: NTT Unit– 2 x 2 Butterfly Unit

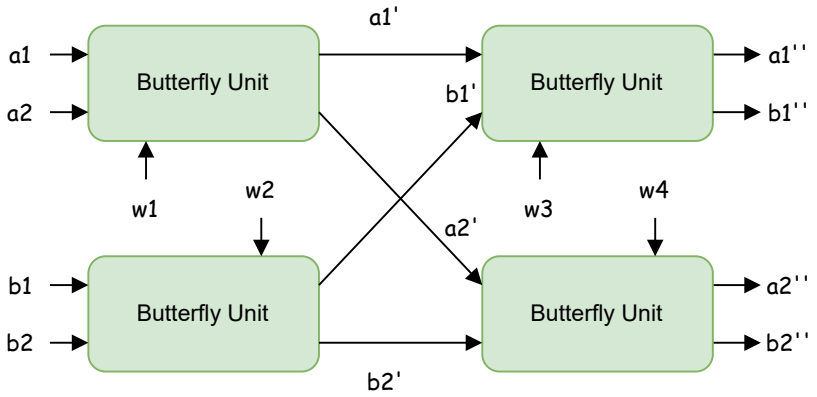

Operations of each Butterfly Unit

- Cooley-Tukey Butterfly (Forward NTT):
  $a1' = a1 + \omega \cdot a2 \; mod \; q$
  $a2' = a1 - \omega \cdot a2 \; mod \; q$
- Gentleman-Sande Butterfly (Inverse NTT):
  $a1' = a1 + a2 \; mod \; q,$
  $a2 = (a1 - a2) \cdot \omega \; mod \; q$
- Inverse also requires scaling by $\frac{1}{N} = \left(\frac{1}{2}\right)^{\log_2 N}$

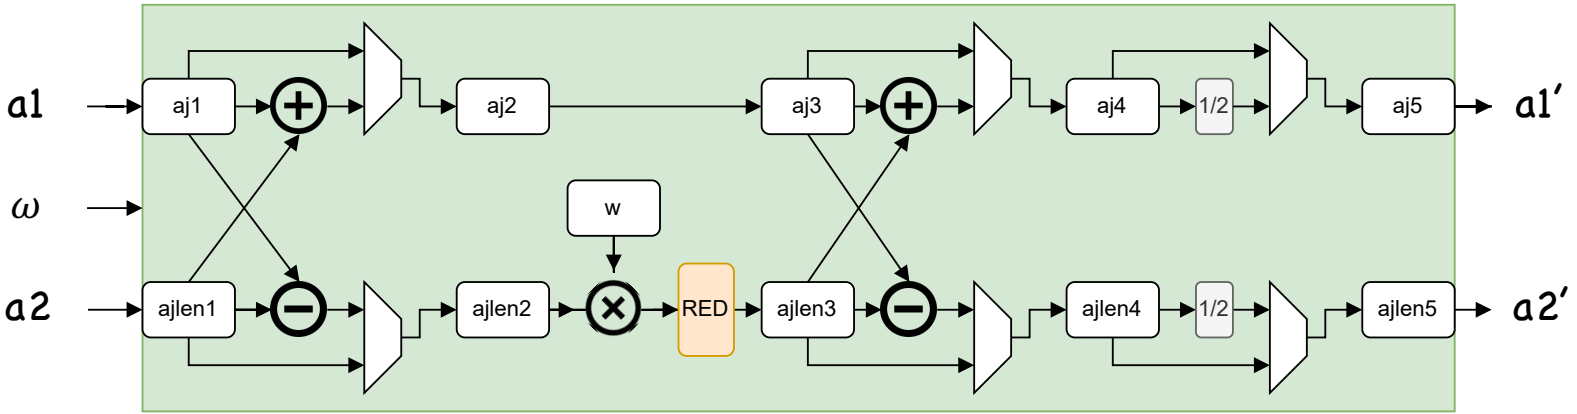

Dilithium: NTT Unit– 4xFIFO

Forward NTT

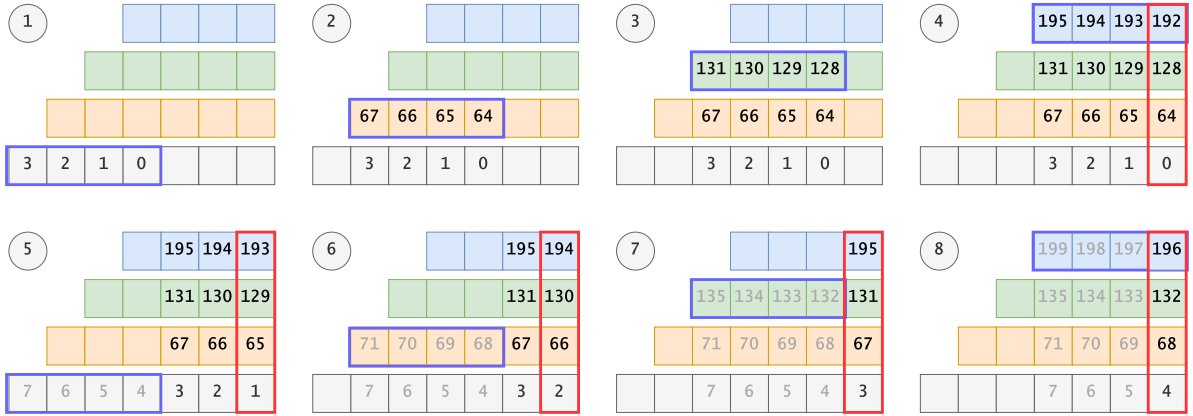

Inverse NTT

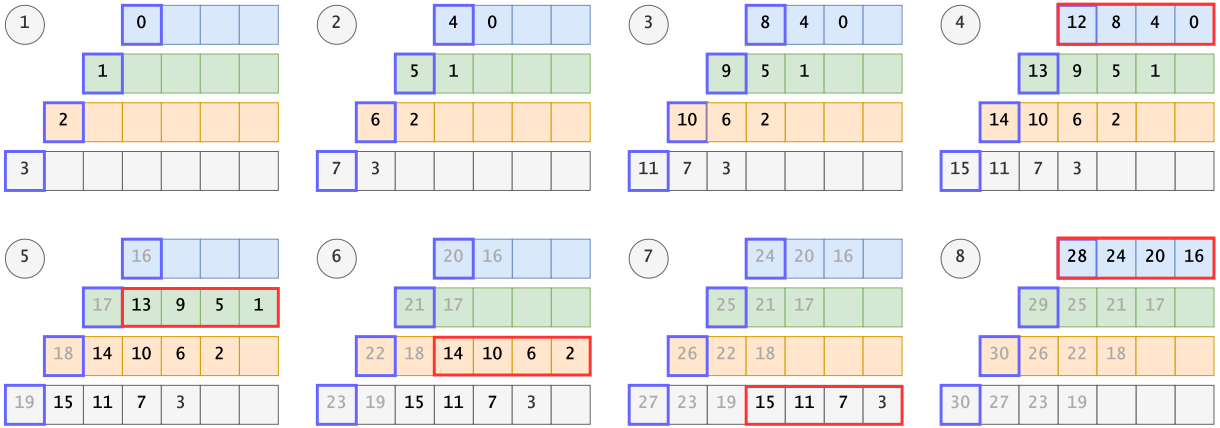

Dilithium: Example of Operation Scheduling - Verification

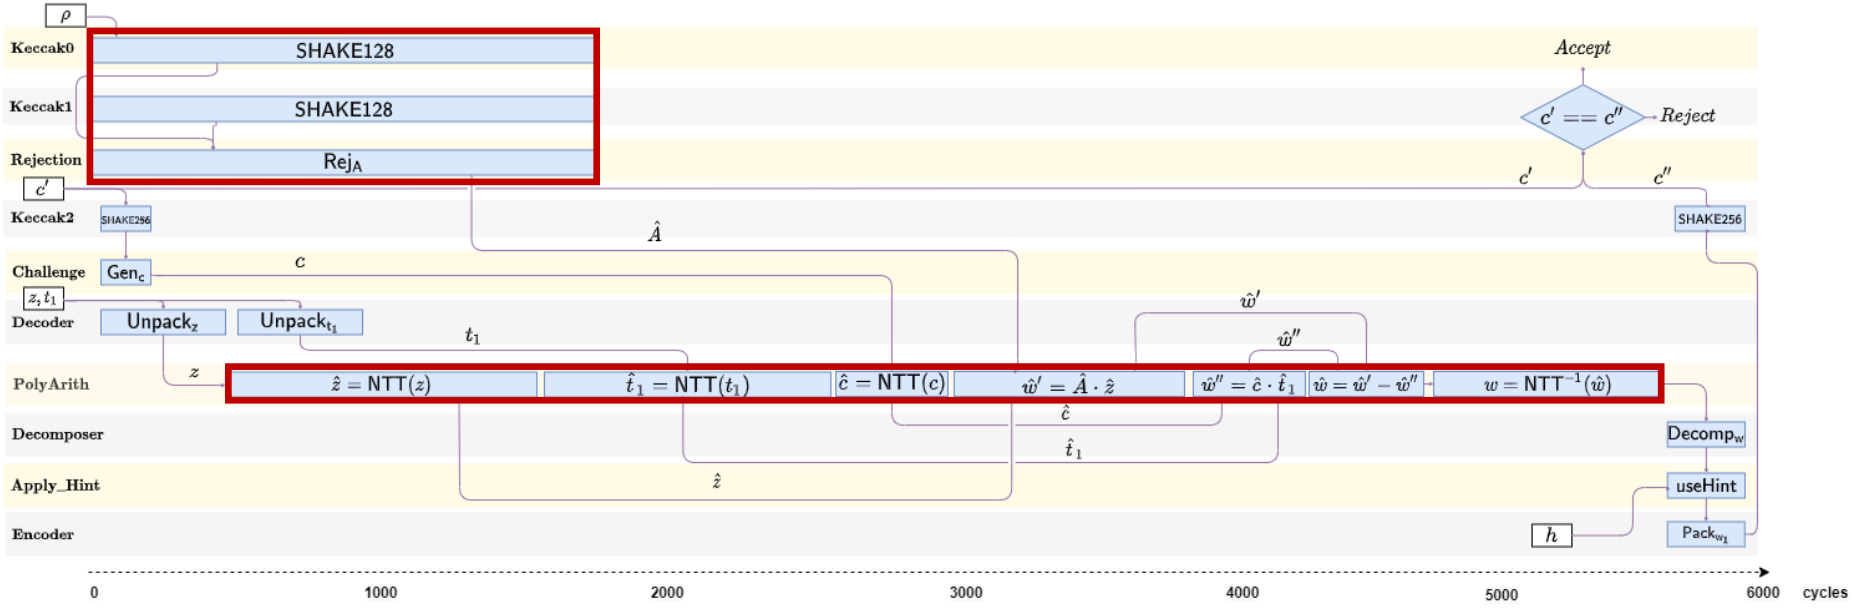

Bottlenecks: Sampling of A matrix + calculation of $Az - ct_1$

FALCON (Verification only): High-Level Architecture

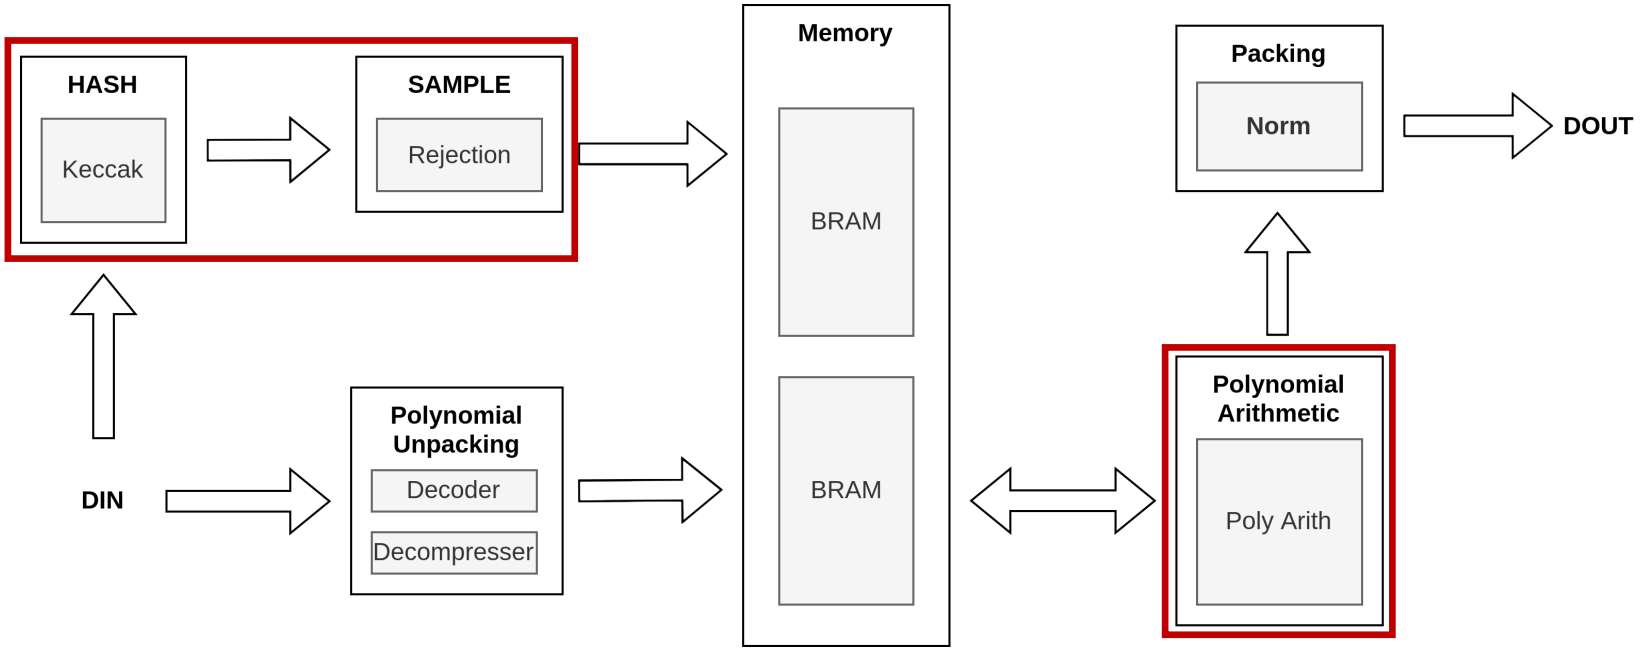

- KeyGen and Sign are much more complicated
- Require implementing floating point Fast Fourier Transform (FFT)
- Hardware implementations not reported yet

FALCON: Major Operations of Verification

Given message M, signature $(r \in \{0,1\}^{320}, s_2 \in Z_q[x])$, and the public key $(h \in Z_q[x])$, calculate:

$$c \leftarrow HashToPoint(r||M)$$

$$s_1 = c - h * s_2$$

If $\left|\left|s_1, s_2\right|\right|^2 = \sum_{i=0}^{n-1}(s_1^2[i] + s_2^2[i])$ is less than the parameter value β, accept

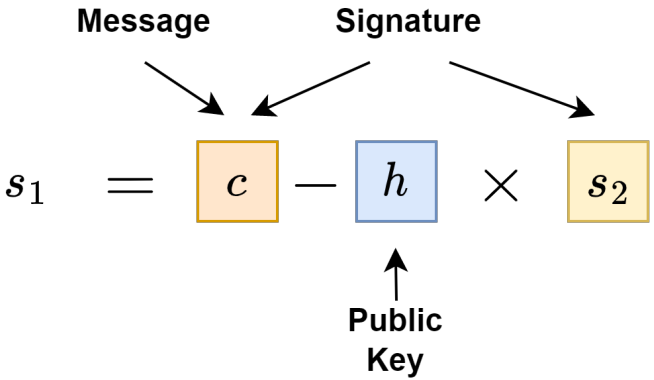

Parameter	FALCON512	FALCON1024
Coefficient Modulus (q)	$12{,}289 = 12 \cdot 2^{10} + 1$	
Number of Polynomial Coefficients (N)	512	1024

Latency of Major Operations in µs: Level 5 – Artix-7 FPGA

Lightweight

1 x NTT, 1 Keccak

High-Speed

1 x NTT, 1 Keccak

4 x NTT, 4 Keccak

4 x NTT, 4 Keccak

2 x NTT, 3 Keccak

1 x NTT, 1 Keccak

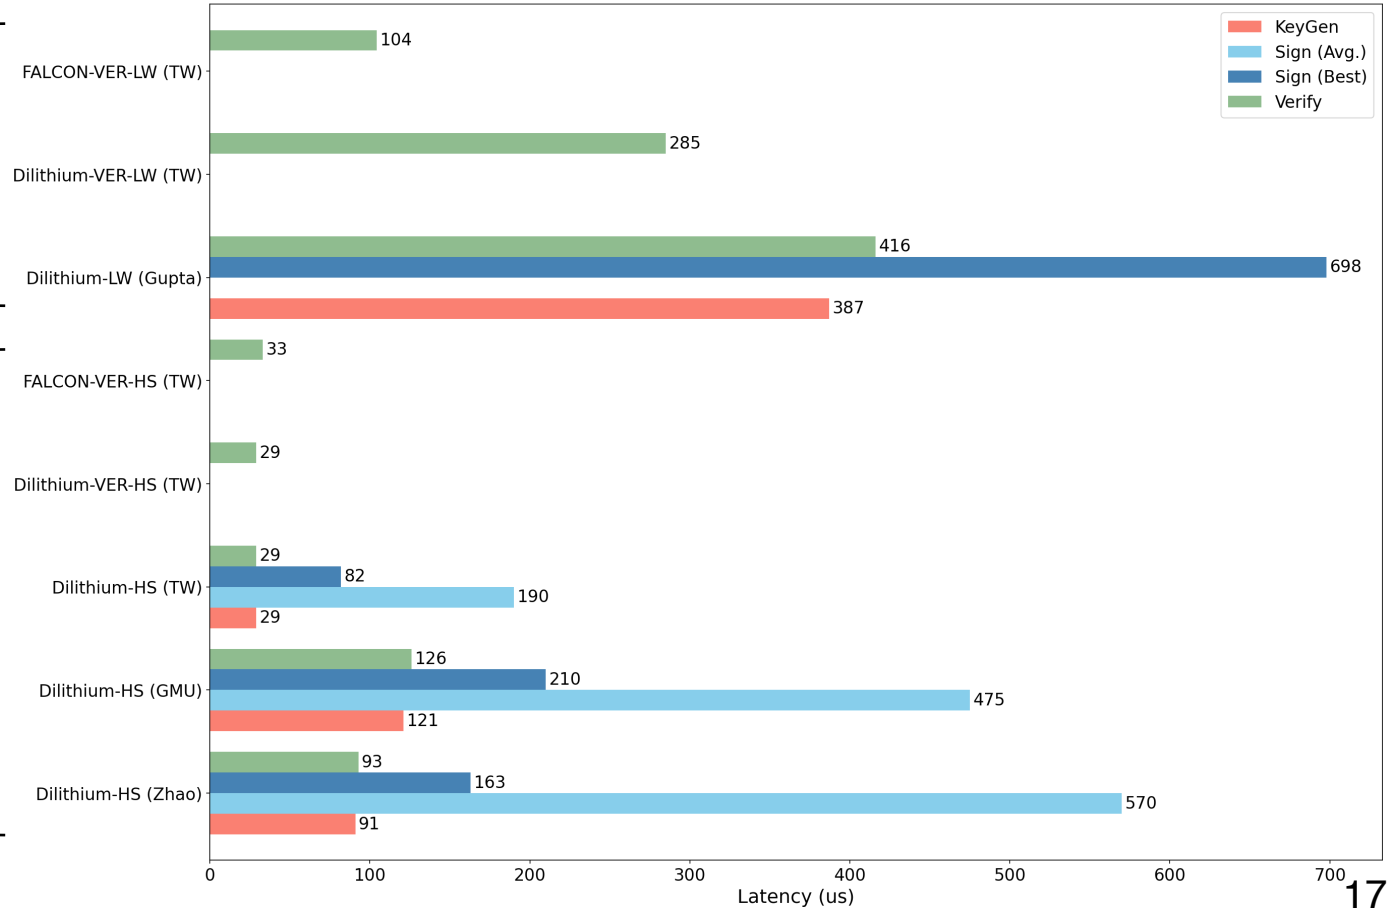

Resource Utilization: Level 5 – Artix-7 FPGA

Lightweight

1 x NTT, 1 Keccak

High-Speed

1 x NTT, 1 Keccak

4 x NTT, 4 Keccak

4 x NTT, 4 Keccak

2 x NTT, 3 Keccak

1 x NTT, 1 Keccak

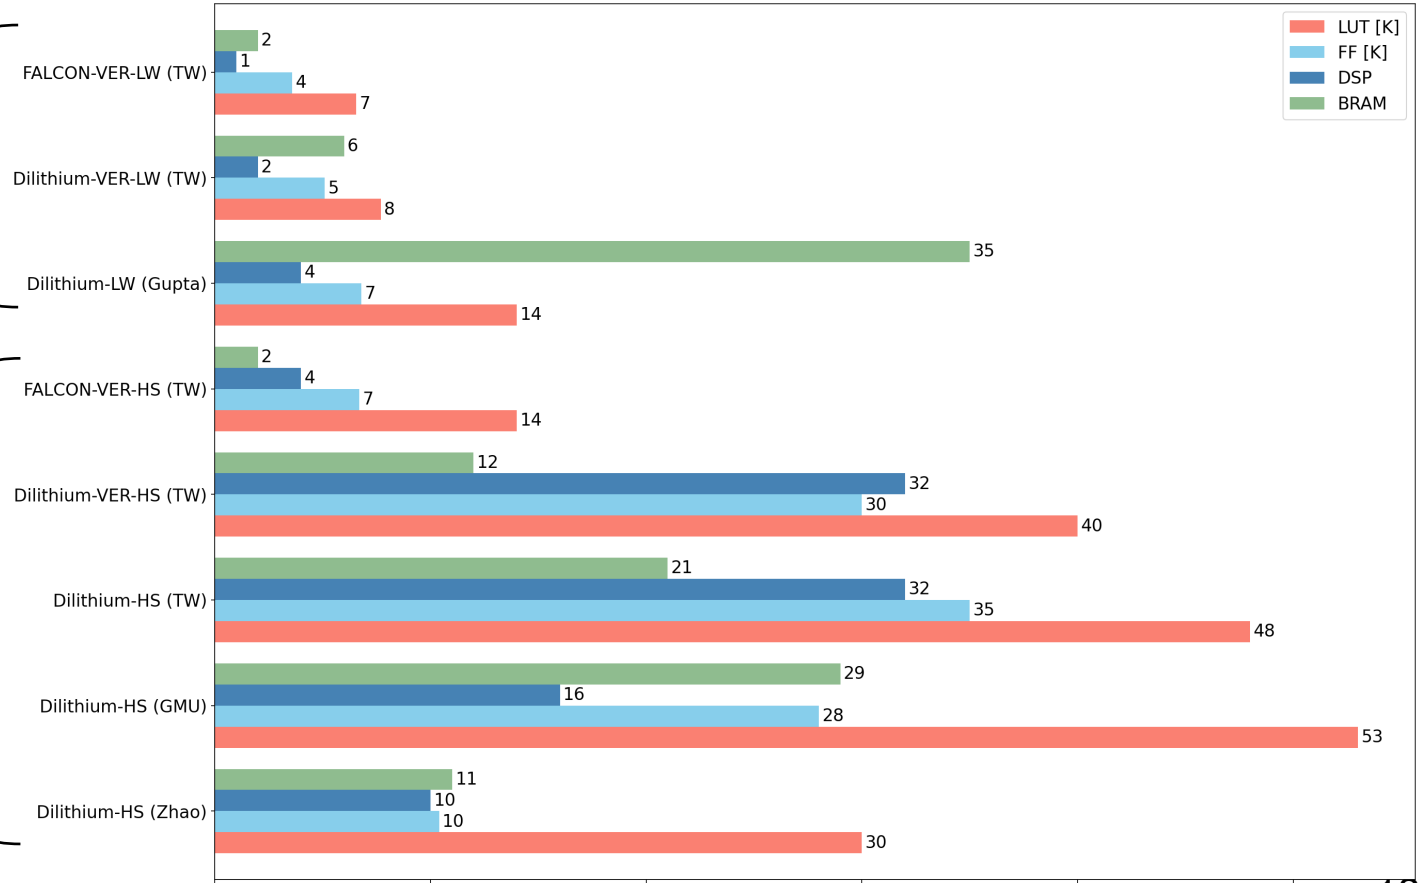

LESS: Conversion of a Matrix to the Reduced Row Echelon Form

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 6 & 0 & 4 \\ 0 & 1 & 5 & 0 & 3 & 1 & 2 \\ 0 & 0 & 0 & 1 & 6 & 3 & 4 \end{bmatrix}$$

Reduced Row Echelon Form (RREF):
(1) the leading entry of every nonzero row is to the right of the leading entry of every row above it
(2) the leading entry in every non-zero row is 1,
(3) each column containing a leading 1 has zeros in all of its other entries.

LESS: Conversion to RREF – Example – Part 1

n=7
k=3
q=7

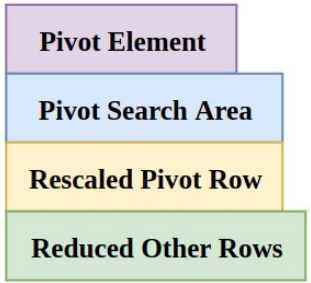

A.

2	2	3	3	1	4	3
3	3	1	5	1	4	3
5	3	1	2	2	2	6

B.

2	2	3	3	1	4	3
3	3	1	5	1	4	3
5	3	1	2	2	2	6

C.

2	2	3	3	1	4	3
3	3	1	5	1	4	3
5	3	1	2	2	2	6

D.

1	1	5	5	4	2	5
3	3	1	5	1	4	3
5	3	1	2	2	2	6

E.

1	1	5	5	4	2	5
0	0	0	4	3	5	2
0	5	4	5	3	6	2

LESS: Conversion to RREF – Example – Part 2

n=7
k=3
q=7

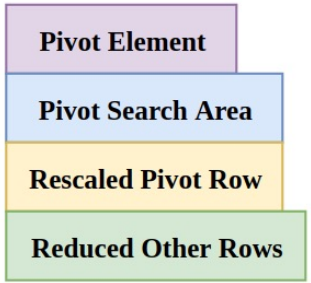

F.
$$\begin{bmatrix} 1 & 1 & 5 & 5 & 4 & 2 & 5 \\ 0 & 0 & 0 & 4 & 3 & 5 & 2 \\ 0 & 5 & 4 & 5 & 3 & 6 & 2 \end{bmatrix}$$

G.
$$\begin{bmatrix} 1 & 1 & 5 & 5 & 4 & 2 & 5 \\ 0 & 0 & 0 & 4 & 3 & 5 & 2 \\ 0 & 5 & 4 & 5 & 3 & 6 & 2 \end{bmatrix}$$

H.
$$\begin{bmatrix} 1 & 1 & 5 & 5 & 4 & 2 & 5 \\ 0 & 1 & 5 & 1 & 2 & 4 & 6 \\ 0 & 0 & 0 & 4 & 3 & 5 & 2 \end{bmatrix}$$

I.
$$\begin{bmatrix} 1 & 0 & 0 & 4 & 2 & 5 & 6 \\ 0 & 1 & 5 & 1 & 2 & 4 & 6 \\ 0 & 0 & 0 & 4 & 3 & 5 & 2 \end{bmatrix}$$

LESS: Conversion to RREF – Example – Part 2

n=7
k=3
q=7

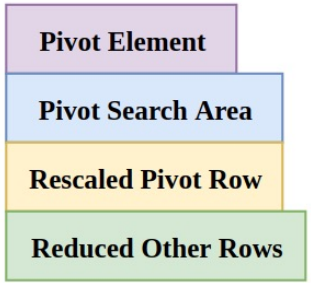

J.

1	0	0	4	2	5	6
0	1	5	1	2	4	6
0	0	0	4	3	5	2

K.

1	0	0	4	2	5	6
0	1	5	1	2	4	6
0	0	0	4	3	5	2

L.

1	0	0	4	2	5	6
0	1	5	1	2	4	6
0	0	0	1	6	3	4

M.

1	0	0	0	6	0	4
0	1	5	0	3	1	2
0	0	0	1	6	3	4

N.

1	0	0	0	6	0	4
0	1	5	0	3	1	2
0	0	0	1	6	3	4

Scheduling of Major Operations within RREF

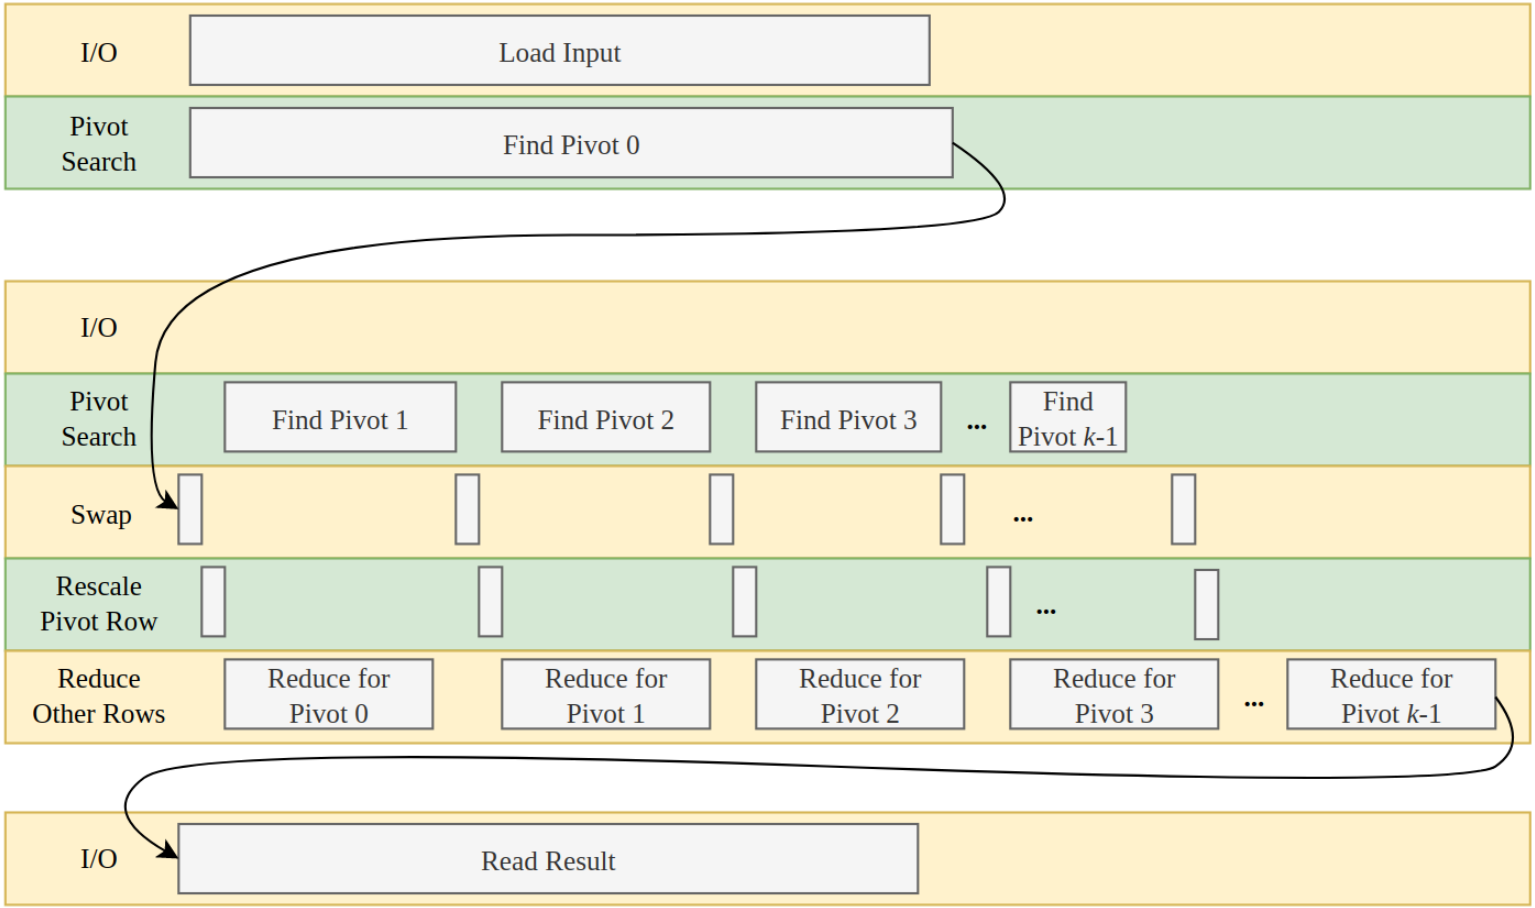

LESS: RREF Unit: Top-Level Block Diagram

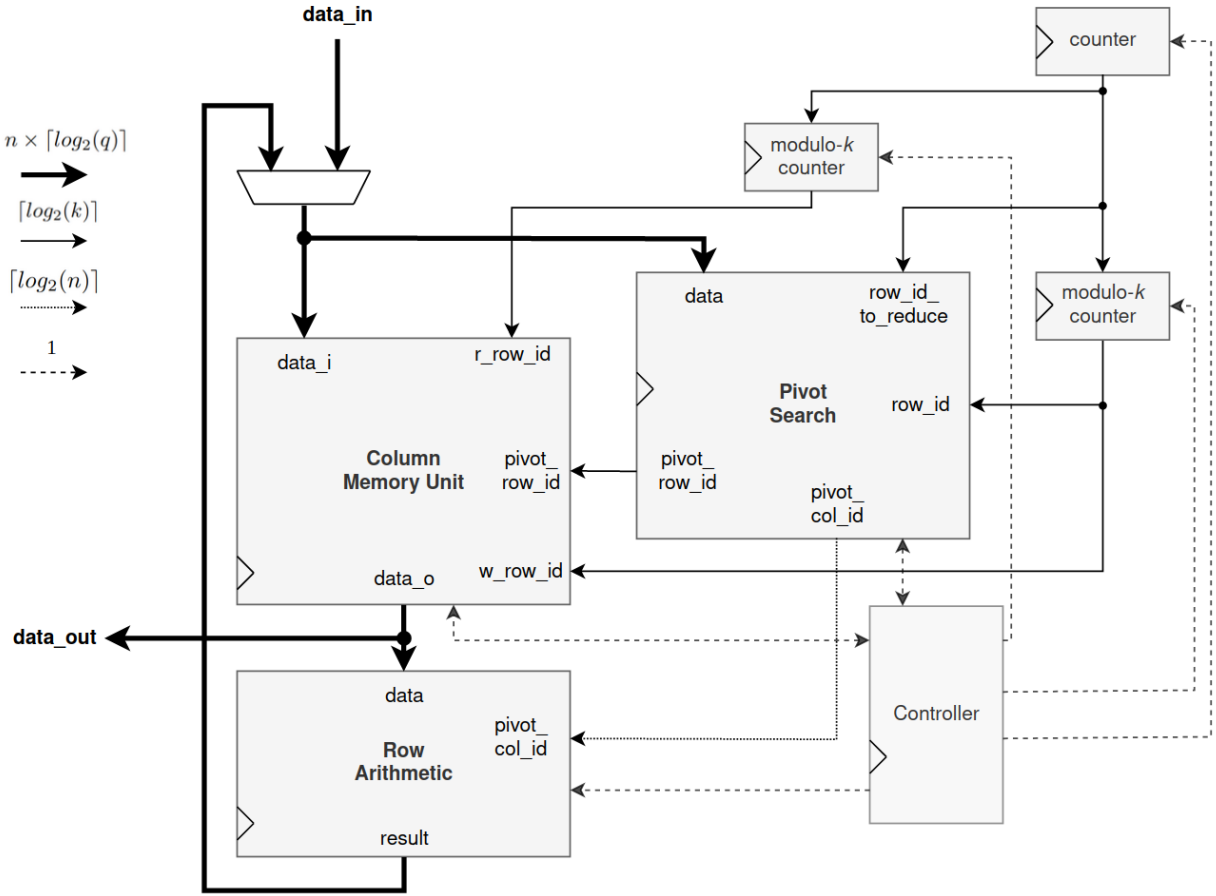

LESS: RREF Unit: Column Memory Unit

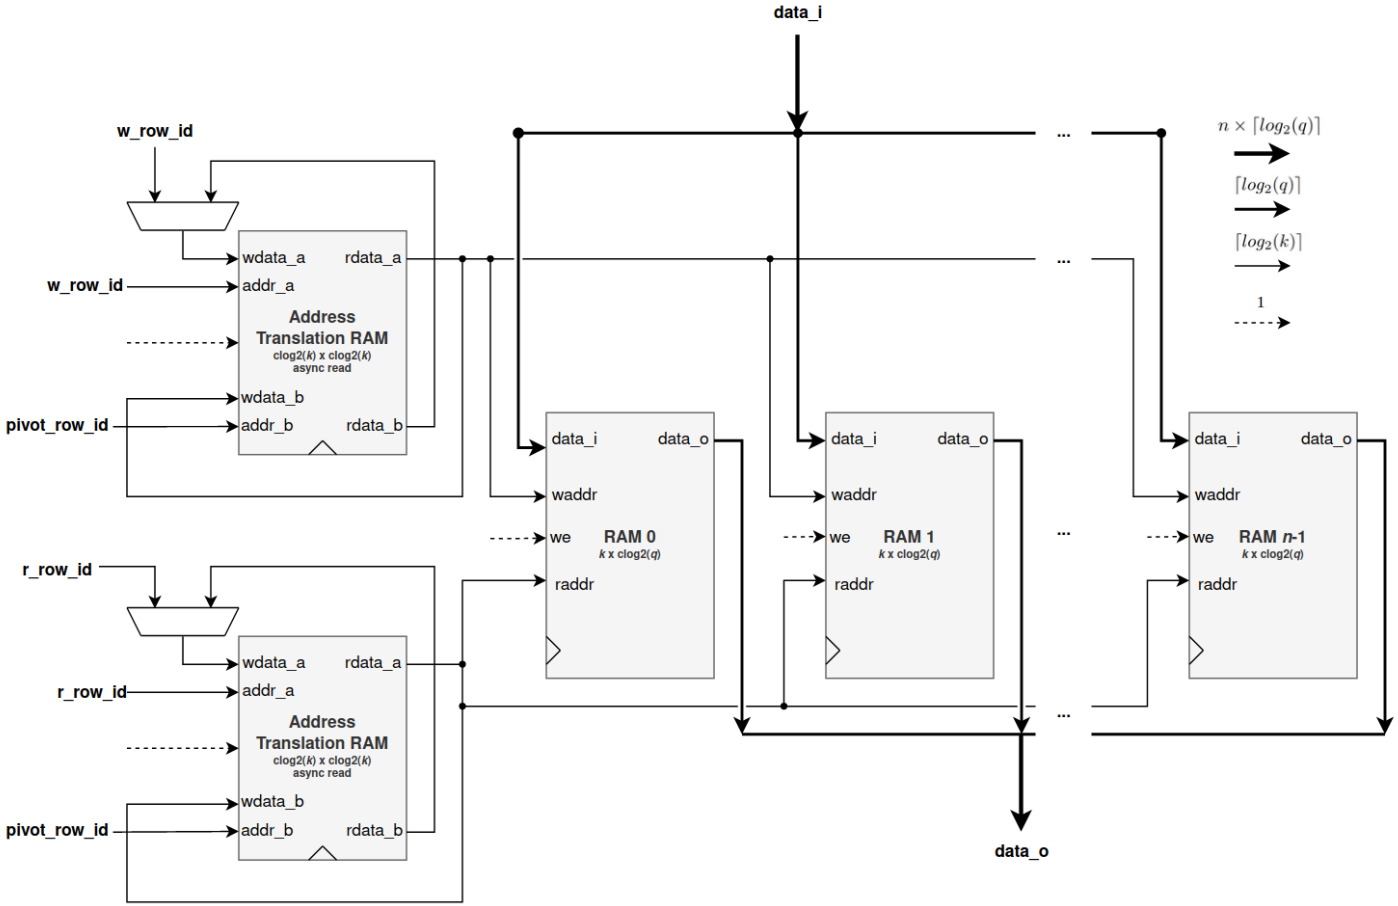

LESS: RREF Unit: Row Arithmetic Unit – Rescale Pivot Row

2	2	3	3	1	4	3
3	3	1	5	1	4	3
5	3	1	2	2	2	6

1	1	5	5	4	2	5
3	3	1	5	1	4	3
5	3	1	2	2	2	6

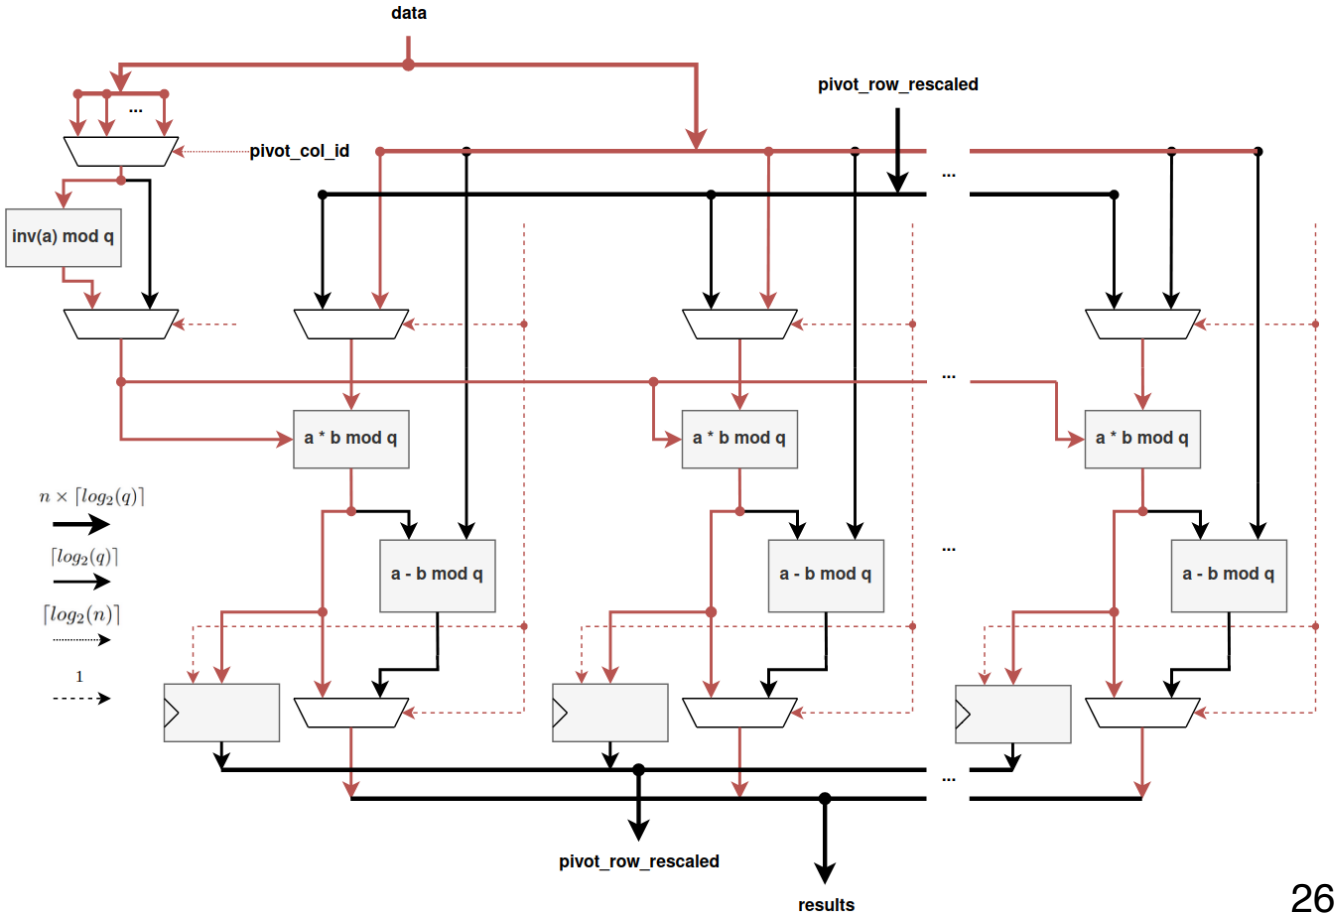

LESS: RREF Unit: Row Arithmetic Unit – Reduce Other Rows

1	1	5	5	4	2	5
3	3	1	5	1	4	3
5	3	1	2	2	2	6

1	1	5	5	4	2	5
0	0	0	4	3	5	2
0	5	4	5	3	6	2

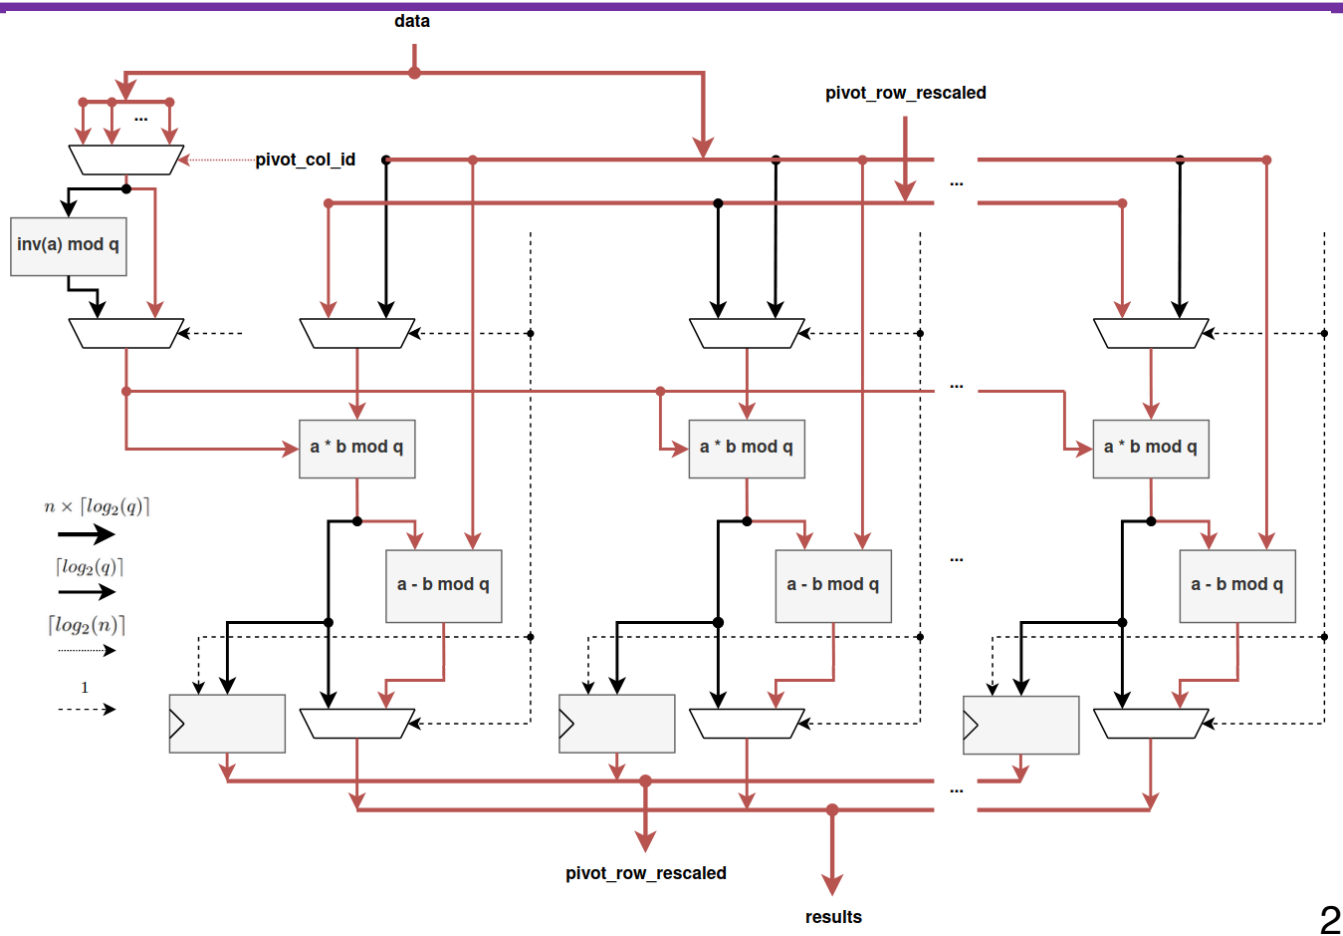

Public Key and Signature Sizes: LESS vs. Near-Term Standards

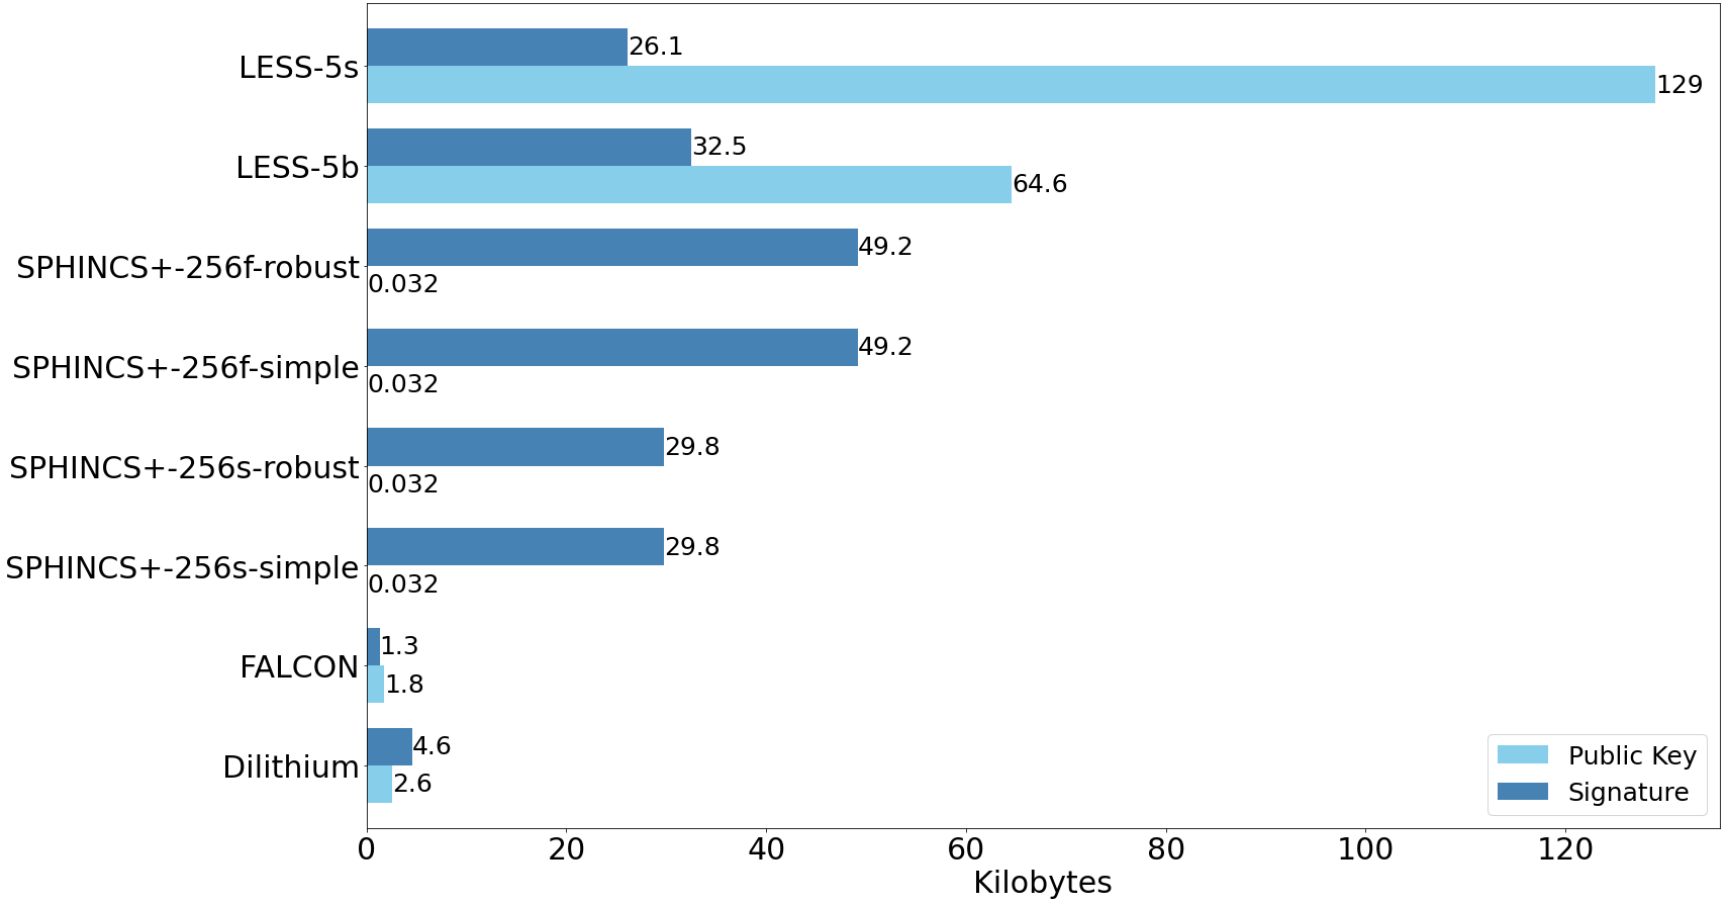

Latencies in µs on Artix-7: LESS vs. Near-Term Standards

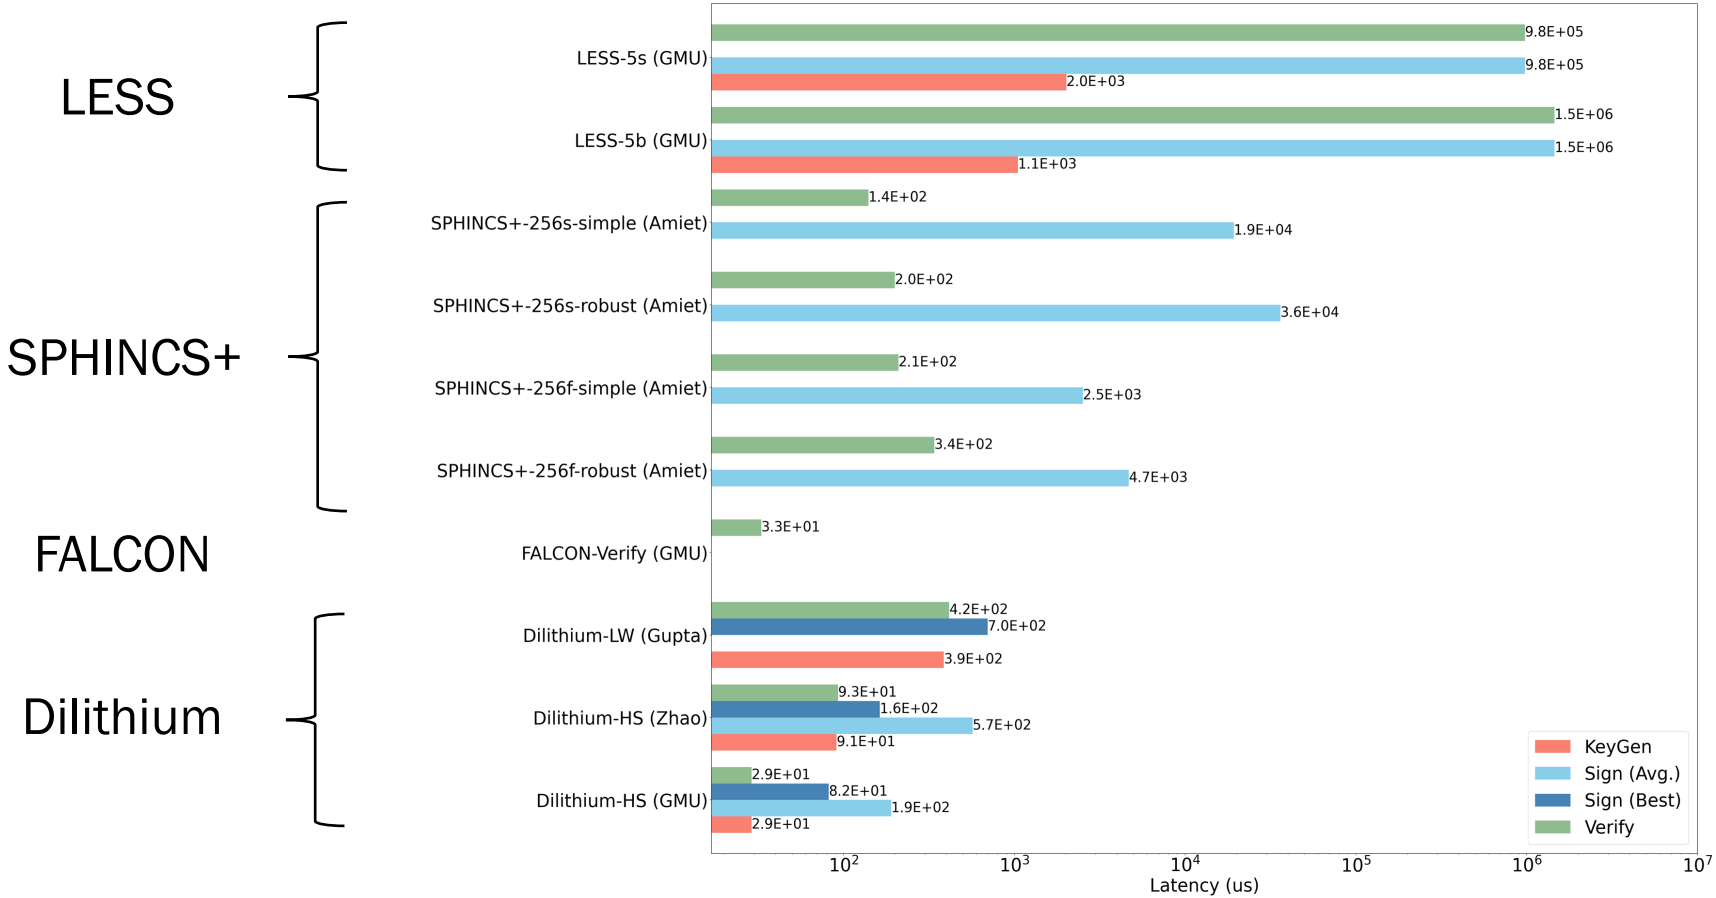

Conclusions

- **FALCON**
  - + Smallest sum signature size + public key size (certificate size)
  - + Very efficient verification
  - − Complex key generation and signing
- **Dilithium**
  - + Second smallest sum signature size + public key size (certificate size)
  - + Efficient key generation, signing, and verifying
- **LESS**
  - + Security dependent on a different problem than all other near-term NIST standards
  - +/− Signature sizes comparable to SPHINCS+ but larger than in FALCON & Dilithium
  - − Longer public-key sizes than in all other near-term NIST standards
  - − Orders of magnitude longer execution times for signing and verifying

Q&A

Cryptographic Engineering Research Group

Thank You!

Questions?

Comments?

CERG: http://cryptography.gmu.edu

ATHENa: http://cryptography.gmu.edu/athena
Menu Field: PQC

Cryptographic Engineering Research Group