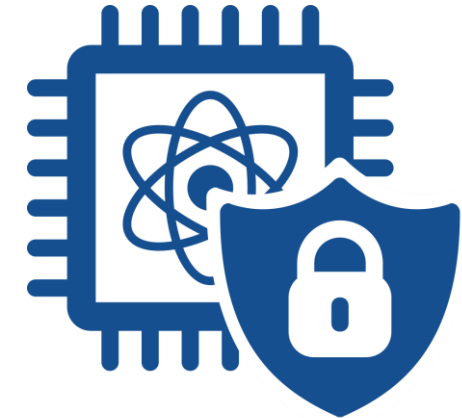


SDitH in Hardware



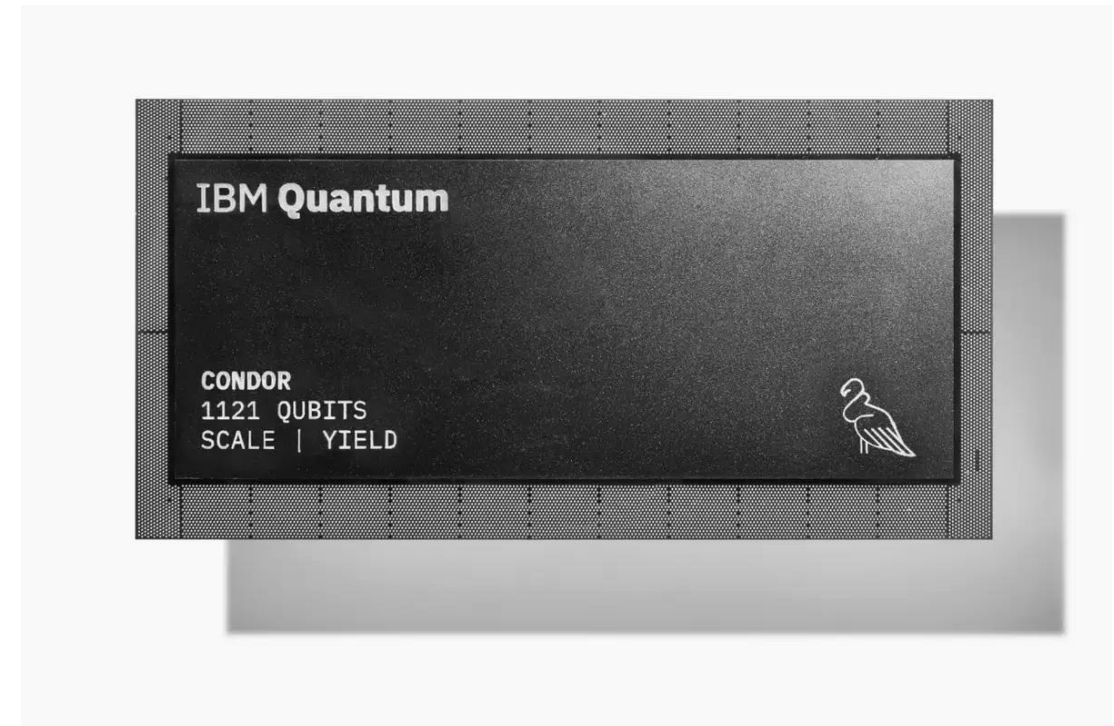
Sanjay Deshpande, James Howe, Jakub Szefer, and
Dongze (Steven) Yue

5th NIST PQC Standardization Conference
April 12, 2024

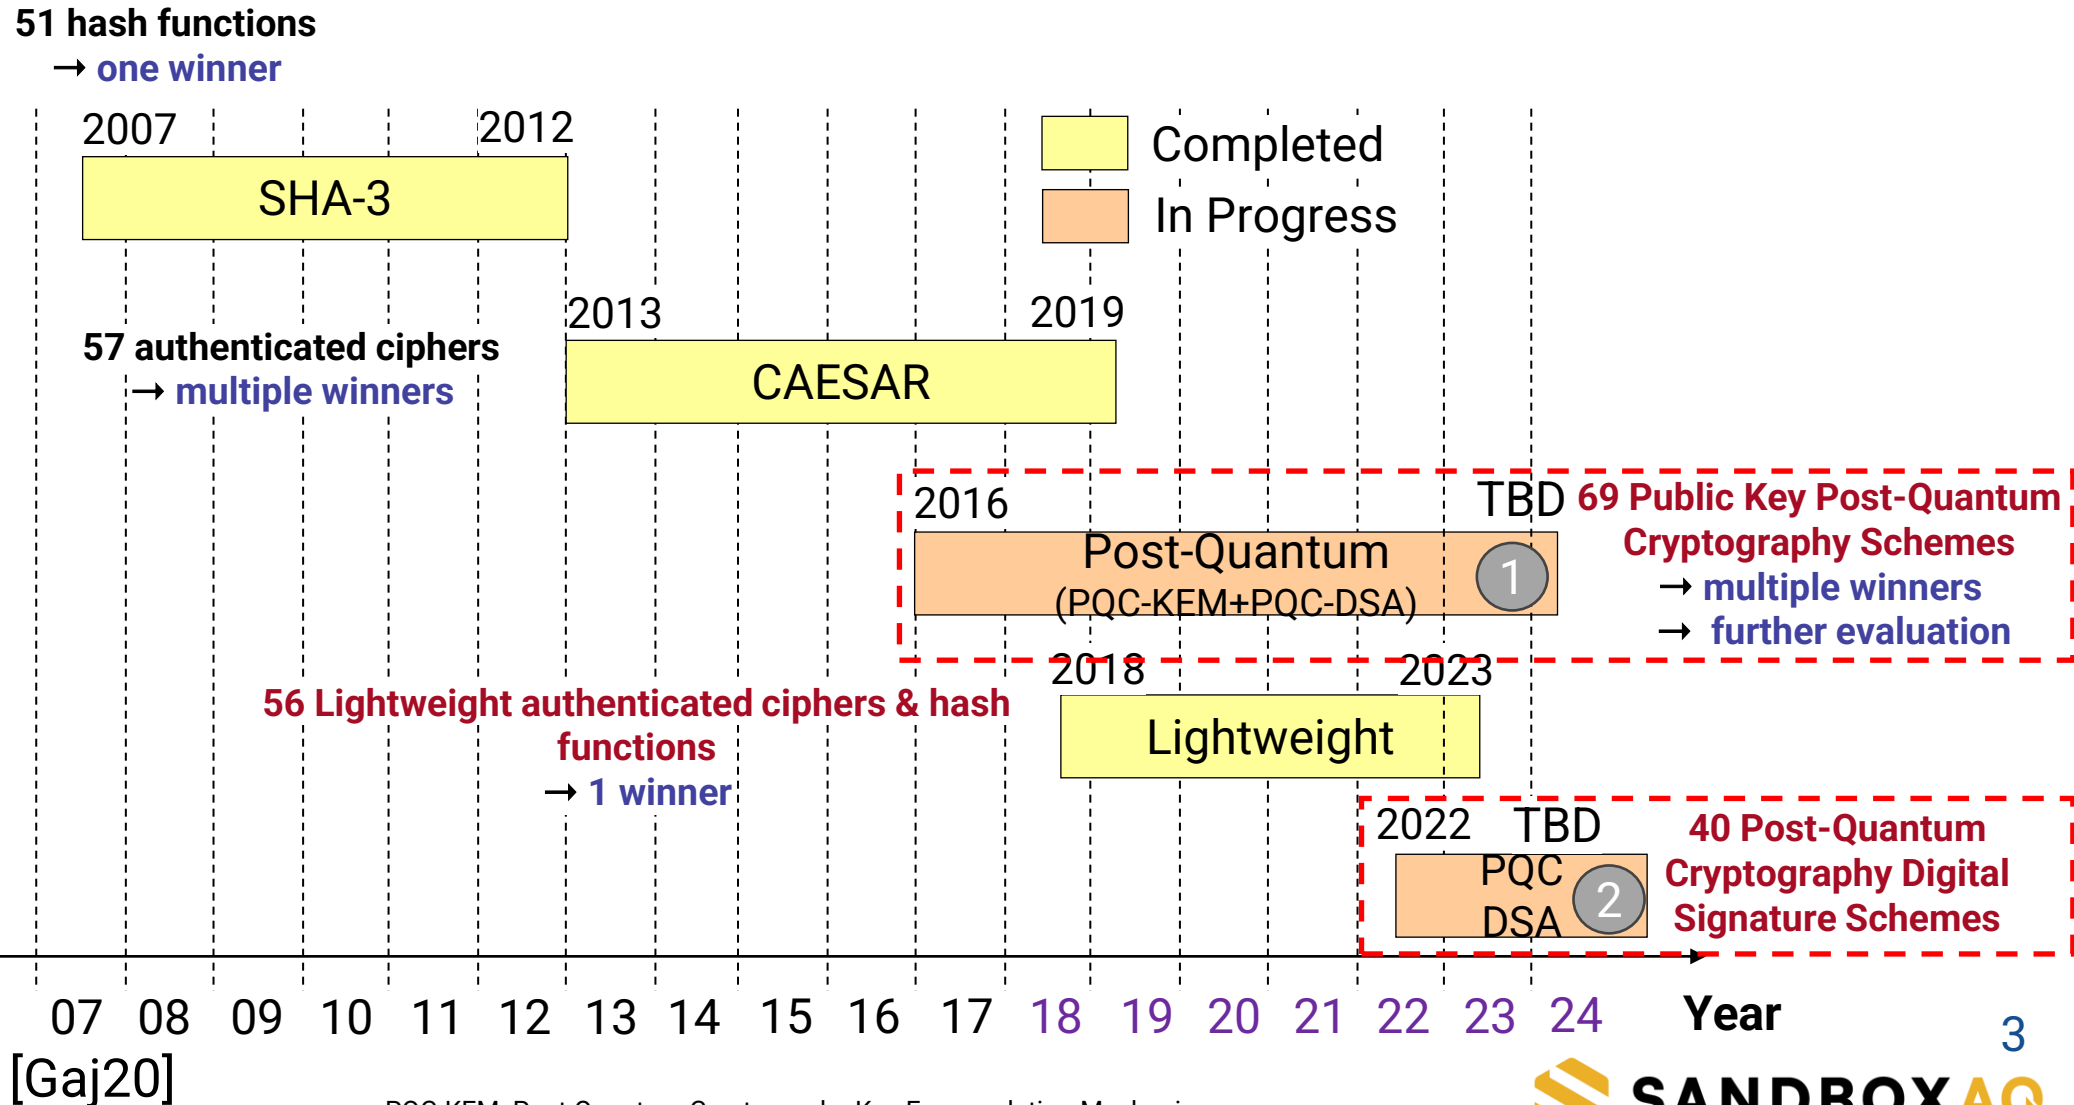


Motivation

- Quantum Computing holds tremendous potential that could solve complex problems that are out of reach for current high-performance computers
 - Life-saving pharmaceuticals
 - Green-battery technology
- However, they also pose significant cybersecurity risks
 - Can easily break existing standards of public key cryptography
 - Can jeopardize payment systems, encrypted chat, emails, etc.
- The [Quantum Insider's report](#) from 2022 forecasts the quantum security market worth \$10 billion by 2030
- Currently, we do not have large-scale quantum computers
 - In 2023, IBM [announced](#) the 1,121-qubit quantum processor "Condor"
- Hence, there is a need for Quantum-safe Cryptography!
 - Post quantum cryptography emerges as a beacon of hope



NIST Post Quantum Cryptography Standardization Effort



PQC-KEM: Post Quantum Cryptography-Key Encapsulation Mechanism
 PQC-DSA: Post Quantum Cryptography-Digital Signature Algorithm

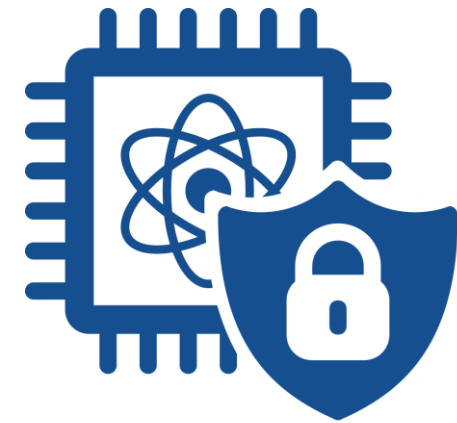


Outline

- Introduction
 - SDitH Signature Scheme
- Hardware Design and Challenges
- Comparison with Related and Relevant Work
- Conclusion and Future Work



Introduction



SDitH Parameter Sets

- Two Variants of the Algorithm
 - **Hypercube**
 - Threshold
- Three Security Levels
- Two Syndrome Decoding Fields
 - **GF256 and GF251**
- **d=2** splits for L3 and L5 Parameter Sets

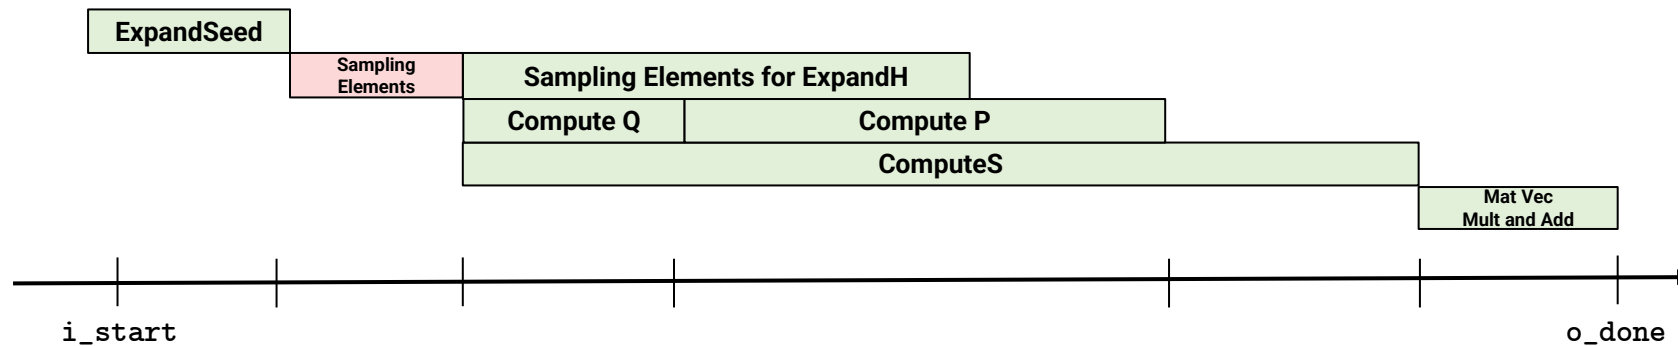
		SDitH Parameters	NIST Security Categories		
			L1	L3	L5
Signature Parameters	NIST Security Level		143	207	272
	λ (Security Target)		128	192	256
	N^D (# Secret Shares)		2^8	2^8	2^8
	τ (Repetition Rate)		17	26	34
Syndrome Decoding	q (SD Base Field Size)		251/256	251/256	251/256
	m (Code Length)		230	352	480
	k (Vector Dimension)		126	193	278
	w (Hamming Weight Bound)		79	120	150
	d (d Splitting Size)		1	2	2
Multi-Party Computation	t (# Random evaluation points)		3	3	4
	\mathbb{F}_q (SD base field)		\mathbb{F}_q	\mathbb{F}_q	\mathbb{F}_q
	η (Field extension size)		4	4	4
	$\mathbb{F}_{\text{points}}$ (Field extension of \mathbb{F}_q)		\mathbb{F}_{q^η}	\mathbb{F}_{q^η}	\mathbb{F}_{q^η}
	p (False positive probability)		$2^{-71.2}$	$2^{-72.4}$	$2^{-94.8}$
Output Sizes	pk Size (in Bytes)		120	183	234
	sk Size (in Bytes)		404	616	812
	Max Signature Size (in Bytes)		8 260	19 206	33 448



SDitH Key Generation

Algorithm 1 SDitH – Key Generation

- 1: $\text{seed}_{\text{root}} \leftarrow \{0, 1\}^\lambda$
 - 2: $(\text{seed}_{\text{wit}}, \text{seed}_H) \leftarrow \text{ExpandSeed}(\text{salt} := 0, \text{seed}_{\text{root}}, 2)$
 - 3: $(Q, S, P) \leftarrow \text{SampleWitness}(\text{seed}_{\text{wit}})$
 - 4: $s = \text{Serialize}(S)$
 - 5: $(s_A, s_B) = \text{Parse}(s, \mathbb{F}_q^k, \mathbb{F}_q^{m-k})$
 - 6: $H' \leftarrow \text{ExpandH}(\text{seed}_H)$
 - 7: $y = s_B + H' s_A$
 - 8: $Q' = \text{TruncateQ}(Q)$
 - 9: $\text{wit_plain} = \text{Serialize}(s_A, Q', P)$
 - 10: **return** $(pk = (\text{seed}_H, y), sk = (\text{seed}_H, y, \text{wit_plain}))$
-



- Variable time due to rejection sampling
- Constant time



SDitH Sign

- Only a little scope for parallelism at the algorithm level
- Processing Message (m) input happens much later in the algorithm
 - Hence, could be divided in to *Offline* and *Online* Parts

Algorithm 10 SD-in-the-Head – Hypercube Variant – Signature Algorithm

Input: a secret key $sk = (\text{seed}_H, y, \text{wit_plain})$ and a message $m \in \{0, 1\}^*$

- 1: $\text{salt} \leftarrow \{0, 1\}^{2\lambda}$
- 2: $\text{mseed} \leftarrow \{0, 1\}^\lambda$
- 3: $H' \leftarrow \text{ExpandH}(\text{seed}_H)$ $\triangleright H' \in \mathbb{F}_q^{(n-k) \times k}$
- 4: $\{\text{rseed}[e]\}_{e \in [1:\tau]} \leftarrow \text{ExpandSeed}(\text{salt}, \text{mseed}, \tau)$ $\triangleright \text{rseed}[e] \in \{0, 1\}^\lambda$
- 5: **for** $e \in [1:\tau]$ **do**
- 6: $(\text{seed}[e][i])_{i \in [1:2^D]} \leftarrow \text{TreePRG}(\text{salt}, \text{rseed}[e])$
- 7: $\text{acc} = 0$ $\triangleright \text{acc} \in \mathbb{F}_q^{k+2w+t(2d+1)\eta}$
- 8: $\text{input_mshare}[e][p] = 0$ for all $(e, p) \in [1:\tau] \times [1:D]$ $\triangleright \text{input_mshare}[e][i] \in \mathbb{F}_q^{k+2w+t(2d+1)\eta}$
- 9:
- 10: **for** $i \in [1:2^D]$ **do**
- 11: **if** $i \neq 2^D$ **then**
- 12: $\text{input_share}[e][i] \leftarrow \text{SampleFieldElements}(\text{salt}, \text{seed}[e][i], k + 2w + t(2d + 1)\eta)$ $\triangleright \text{input_share}[e][i] \in \mathbb{F}_q^{k+2w+t(2d+1)\eta}$
- 13:
- 14: $\text{acc} += \text{input_share}[e][i]$
- 15: $\text{state}[e][i] = \text{seed}[e][i]$
- 16: **for** $p \in [1:D]$: the p^{th} bit of $i - 1$ is zero, **do**
- 17: $\text{input_mshare}[e][p] += \text{input_share}[e][i]$
- 18: **else**
- 19: $\text{acc.wit}, \text{acc.beav.ab}, \text{acc.beav.c} = \text{acc}$
- 20: $\text{beav.ab.plain}[e] = \text{acc.beav.ab} + \text{SampleFieldElements}(\text{salt}, \text{seed}[e][i], 2dt\eta)$
- 21: $\text{beav.c.plain}[e] = \text{beav.c.plain} \leftarrow \text{InnerProducts}(\text{beav.ab.plain})$ $\triangleright a \cdot b = c$
- 22: $\text{aux}[e] = (\text{wit_plain} - \text{acc.wit}, \text{beav.c.plain}[e] - \text{acc.beav.c})$ $\triangleright \text{aux}[e] \in \mathbb{F}_q^{k+2w+t\eta}$
- 23: $\text{state}[e][i] = (\text{seed}[e][i], \text{aux}[e])$
- 24: $\text{com}[e][i] = \text{Commit}(\text{salt}, e, i, \text{state}[e][i])$
- 25: $h_1 = \text{Hash}_1(\text{seed}_H, y, \text{salt}, \text{com}[1][1], \dots, \text{com}[\tau][2^D])$
- 26: $(\text{chal}[e])_{e \in [1:\tau]} \leftarrow \text{ExpandMPCChallenge}(h_1, \tau)$
- 27: **for** $e \in [1:\tau]$ **do**
- 28: $\text{input_plain}[e] = (\text{wit_plain}, \text{beav.ab.plain}[e], \text{beav.c.plain}[e])$
- 29: $\text{broad_plain}[e] \leftarrow \text{ComputePlainBroadcast}(\text{input_plain}[e], \text{chal}[e], (H', y))$
- 30: **for** $p \in [1:D]$ **do**
- 31: $\text{broad_share}[e][p] = \text{PartyComputation}(\text{input_mshare}[e][p], \text{chal}[e],$
- 32: $(H', y), \text{broad_plain}[e], \text{False})$
- 33: $\triangleright \text{broad_share}[e][p] \in \mathbb{F}_q^{(2d+1)t\eta}$
- 34: $h_2 = \text{Hash}_2(m, \text{salt}, h_1, \{\text{broad_plain}[e], \{\text{broad_share}[e][p]\}_{p \in [1:D]}\}_{e \in [1:\tau]})$
- 35: $\{i^*[e]\}_{e \in [1:\tau]} \leftarrow \text{ExpandViewChallenge}(h_2, 1)$
- 36: **for** $e \in [1:\tau]$ **do**
- 37: $\text{path}[e] \leftarrow \text{GetSeedSiblingPath}(\text{rseed}[e], i^*[e])$
- 38: **if** $i^*[e] = 2^D$ **then**
- 39: $\text{view}[e] = \text{path}[e]$
- 40: **else**
- 41: $\text{view}[e] = (\text{path}[e], \text{aux}[e])$
- 42: $\sigma = (\text{salt} \parallel h_2 \parallel (\text{view}[e], \text{broad_plain}[e], \text{com}[e][i^*[e]])_{e \in [1:\tau]})$
- 43: **return** σ

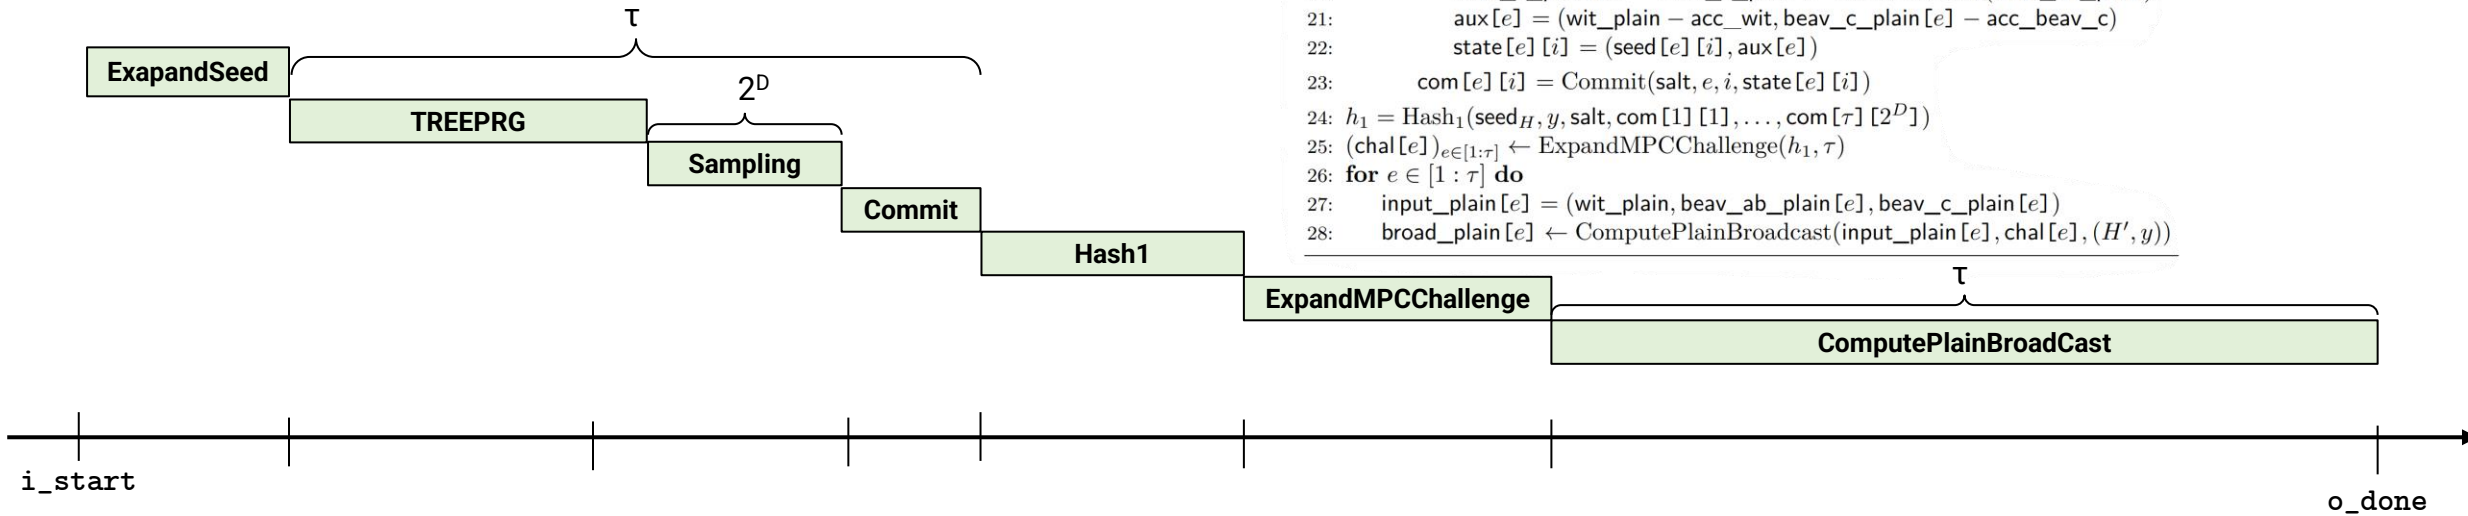


SDitH Sign - Offline

Algorithm 2a SDitH – Hypercube Variant – Signature Generation (Offline Part)

Input: a secret key $sk = (\text{seed}_H, y, \text{wit_plain})$ and a message $m \in \{0, 1\}^*$

- 1: $\text{salt} \leftarrow \{0, 1\}^{2\lambda}, \text{mseed} \leftarrow \{0, 1\}^\lambda$
- 2: $H' \leftarrow \text{ExpandH}(\text{seed}_H)$
- 3: $\{\text{rseed}[e]\}_{e \in [1:\tau]} \leftarrow \text{ExpandSeed}(\text{salt}, \text{mseed}, \tau)$
- 4: **for** $e \in [1:\tau]$ **do**
- 5: $(\text{seed}[e][i])_{i \in [1:2^D]} \leftarrow \text{TreePRG}(\text{salt}, \text{rseed}[e])$
- 6: $\text{acc} = 0$
- 7: $\text{input_mshare}[e][p] = 0$ for all $(e, p) \in [1:\tau] \times [1:D]$
- 8: **for** $i \in [1:2^D]$ **do**
- 9: **if** $i \neq 2^D$ **then**
- 10: $\text{input_share}[e][i] \leftarrow \text{SampleFieldElements}(\text{salt}, \text{seed}[e][i], k + 2w + t(2d + 1)\eta)$
- 11:
- 12: $\text{acc} += \text{input_share}[e][i]$
- 13: $\text{state}[e][i] = \text{seed}[e][i]$
- 14: **for** $p \in [1:D]$: the p^{th} bit of $i - 1$ is zero, **do**
- 15: $\text{input_mshare}[e][p] += \text{input_share}[e][i]$
- 16: **else**
- 17: $\text{acc_wit}, \text{acc_beav_ab}, \text{acc_beav_c} = \text{acc}$
- 18: $\text{beav_ab_plain}[e] = \text{acc_beav_ab} + \text{SampleFieldElements}(\text{salt}, \text{seed}[e][i], 2d\eta)$
- 19:
- 20: $\text{beav_c_plain}[e] = \text{beav_c_plain} \leftarrow \text{InnerProducts}(\text{beav_ab_plain})$
- 21: $\text{aux}[e] = (\text{wit_plain} - \text{acc_wit}, \text{beav_c_plain}[e] - \text{acc_beav_c})$
- 22: $\text{state}[e][i] = (\text{seed}[e][i], \text{aux}[e])$
- 23: $\text{com}[e][i] = \text{Commit}(\text{salt}, e, i, \text{state}[e][i])$
- 24: $h_1 = \text{Hash}_1(\text{seed}_H, y, \text{salt}, \text{com}[1][1], \dots, \text{com}[\tau][2^D])$
- 25: $(\text{chal}[e])_{e \in [1:\tau]} \leftarrow \text{ExpandMPCChallenge}(h_1, \tau)$
- 26: **for** $e \in [1:\tau]$ **do**
- 27: $\text{input_plain}[e] = (\text{wit_plain}, \text{beav_ab_plain}[e], \text{beav_c_plain}[e])$
- 28: $\text{broad_plain}[e] \leftarrow \text{ComputePlainBroadcast}(\text{input_plain}[e], \text{chal}[e], (H', y))$



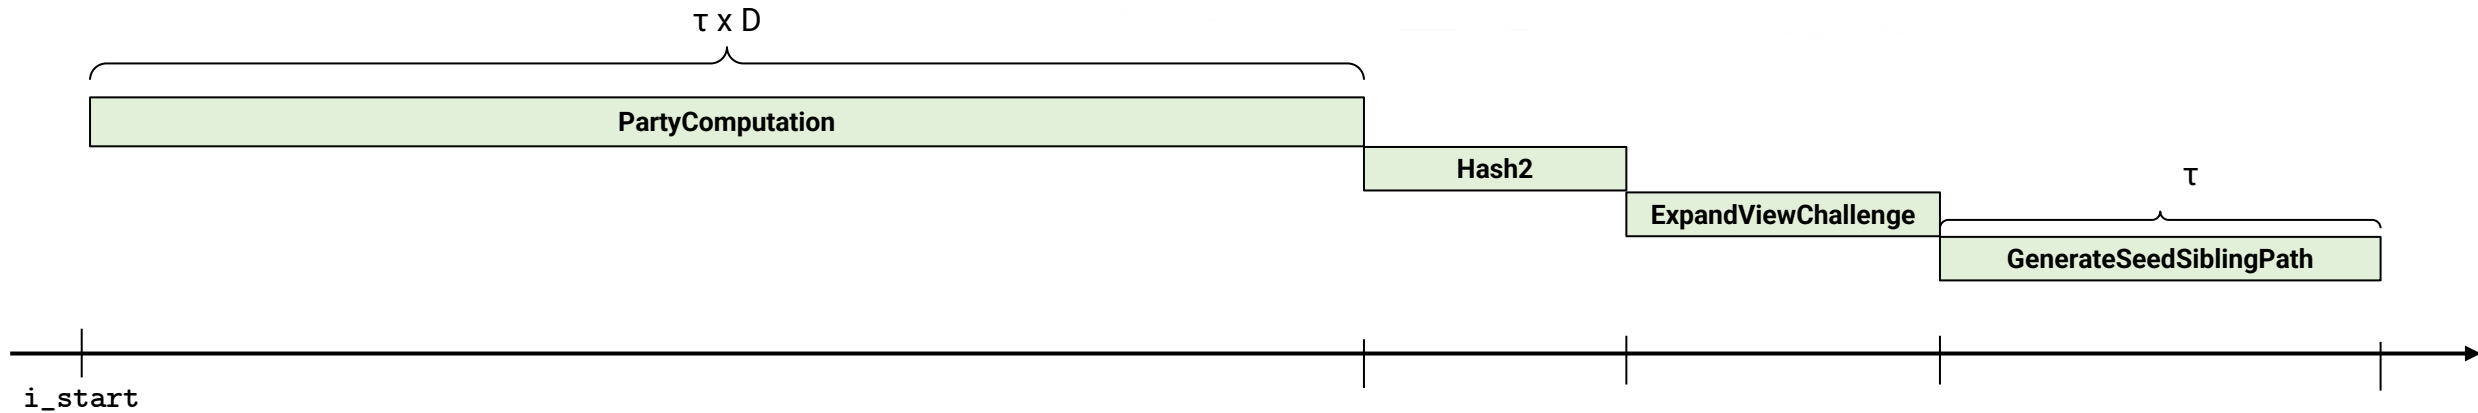
SDitH Sign - Online

Algorithm 2b SDitH – Hypercube Variant – Signature Generation (Online Part)

```

29: for  $e \in [1 : \tau]$  do
30:   for  $p \in [1 : D]$  do
31:      $\text{broad\_share}[e][p] = \text{PartyComputation}(\text{input\_mshare}[e][p], \text{chal}[e],$ 
32:                                              $(H', y), \text{broad\_plain}[e], \text{False})$ 
33:  $h_2 = \text{Hash}_2(m, \text{salt}, h_1, \{\text{broad\_plain}[e], \{\text{broad\_share}[e][p]\}_{p \in [1:D]}\}_{e \in [1:\tau]})$ 
34:  $\{i^*[e]\}_{e \in [1:\tau]} \leftarrow \text{ExpandViewChallenge}(h_2, 1)$ 
35: for  $e \in [1 : \tau]$  do
36:    $\text{path}[e] \leftarrow \text{GetSeedSiblingPath}(\text{rseed}[e], i^*[e])$ 
37:   if  $i^*[e] = 2^D$  then
38:      $\text{view}[e] = \text{path}[e]$ 
39:   else
40:      $\text{view}[e] = (\text{path}[e], \text{aux}[e])$ 
41: return  $\sigma = (\text{salt} \mid h_2 \mid (\text{view}[e], \text{broad\_plain}[e], \text{com}[e][i^*[e]])_{e \in [1:\tau]})$ 

```



o_done

10



SDitH Verify

- Similar to sign_offline and sign_online not so much scope for the parallelism at the algorithmic level
- Possibility of parallelism at the function/module level at cost of additional hardware
- Unrolling the for loops at the cost of duplicating modules

Algorithm 7 SDitH – Hypercube Variant – Verification Algorithm

Input: a public key $pk = (\text{seed}_H, y)$, a signature σ and a message $m \in \{0, 1\}^*$

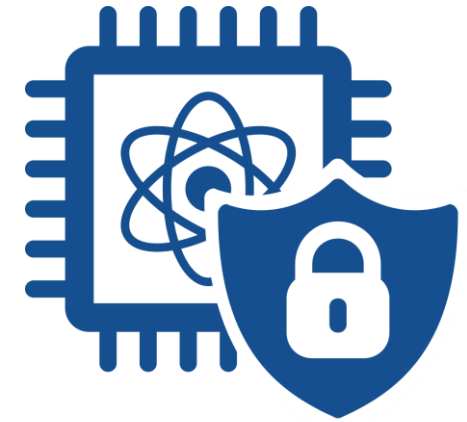
```

1: Parse  $\sigma$  as  $(\text{salt} \mid h_2 \mid (\text{view}[e], \text{broad\_plain}[e], \text{com}[e][i^*[e]])_{e \in [1:\tau]})$ 
2:  $H' \leftarrow \text{ExpandH}(\text{seed}_H)$ 
3:  $\{i^*[e]\}_{e \in [1:\tau]} \leftarrow \text{ExpandViewChallenge}(h_2, 1)$ 
4: for  $e \in [1:\tau]$  do
5:    $(\text{seed}[e][i])_{i \in [1:2^D \setminus i^*[e]]} \leftarrow \text{GetLeavesFromSiblingPath}(i^*[e], \text{salt}, \text{path}[e])$ 
6:   for  $i \in [1:2^D \setminus i^*[e]]$  do
7:     if  $i \neq 2^D$  then
8:        $\text{state}[e][i] = \text{seed}[e][i]$ 
9:     else
10:       $\text{state}[e][i] = (\text{seed}[e][i], \text{aux}[e])$ 
11:       $\text{com}[e][i] = \text{Commit}(\text{salt}, e, i, \text{state}[e][i])$ 
12:  $h_1 = \text{Hash}_1(\text{seed}_H, y, \text{salt}, \text{com}[1][1], \dots, \text{com}[\tau][2^D])$ 
13:  $\text{chal} \leftarrow \text{ExpandMPCChallenge}(h_1, \tau)$ 
14: for  $e \in [1:\tau]$  do
15:    $\text{input\_mshare}^*[e][p] = 0$  for all  $(e, p) \in [1:\tau] \times [1:D]$ 
16:   for  $i \in [1:2^D \setminus i^*[e]]$  do
17:     if  $i \neq 2^D$  then
18:        $\text{input\_share}[e][i] \leftarrow \text{SampleFieldElements}(\text{salt}, \text{seed}[e][i], k + 2w + t(2d + 1)\eta)$ 
19:
20:     else
21:        $\text{beav\_ab\_plain}[e][2^D] = \text{SampleFieldElements}(\text{salt}, \text{seed}[e][2^D], 2d\eta)$ 
22:
23:        $\text{input\_share}[e][2^D] = (\text{aux}[e] \mid \text{beav\_ab\_plain}[e][2^D])$ 
24:   for  $p \in [1:D]$ : the  $p^{\text{th}}$  bit of  $i - 1$  and  $i^*[e]$  are different do
25:      $\text{input\_mshare}'[e][p] += \text{input\_share}[e][i]$ 
26:   for  $p \in [1:D]$  do
27:     if the  $p^{\text{th}}$  bit of  $i^*[e]$  is 1 then
28:        $\text{broad\_share}[e][p] = \text{PartyComputation}(\text{input\_mshare}'[e][p], \text{chal},$ 
29:          $(H', y), \text{broad\_plain}, \text{False})$ 
30:     else
31:        $\text{broad\_share}[e][p] = \text{broad\_plain}[e] - \text{PartyComputation}(\text{input\_mshare}'[e][p], \text{chal},$ 
32:          $(H', y), \text{broad\_plain}, \text{True})$ 
33:
34:  $h'_2 = \text{Hash}_2(m, \text{salt}, h_1, \{\text{broad\_plain}[e], \{\text{broad\_share}[e][p]\}_{p \in [1:D]}\}_{e \in [1:\tau]})$ 
35: return  $h_2 \stackrel{?}{=} h'_2$ 

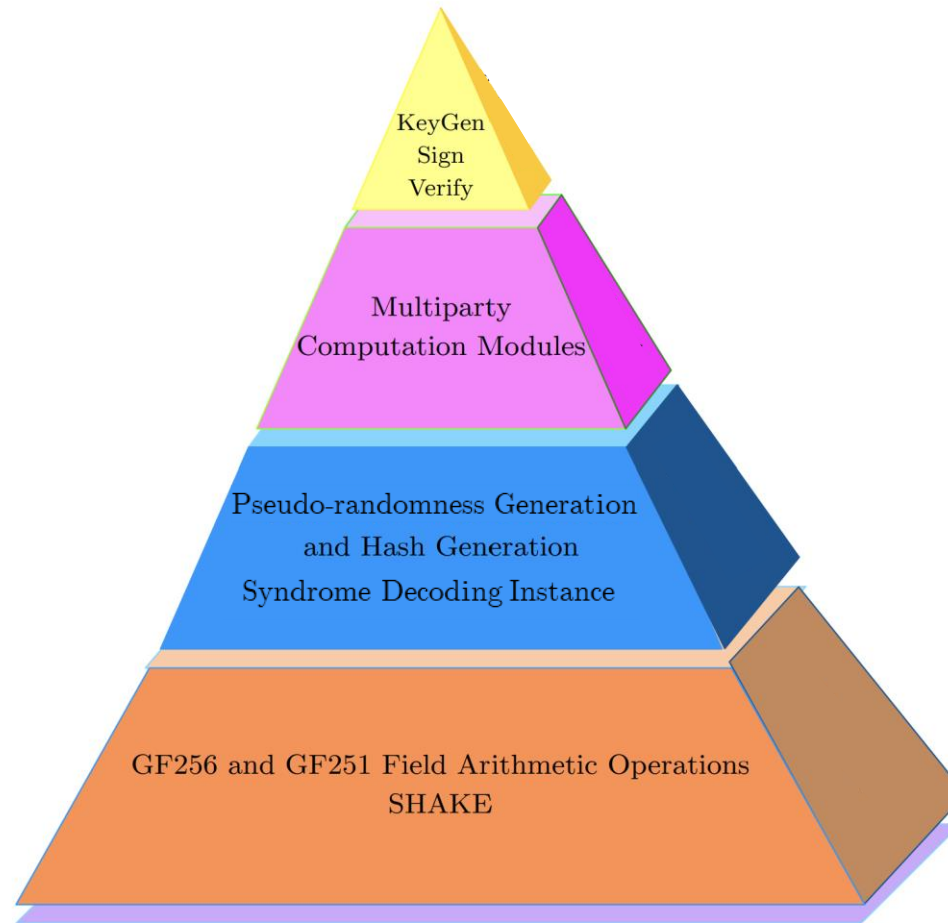
```



Hardware Design and Challenges



Hardware Design Architecture

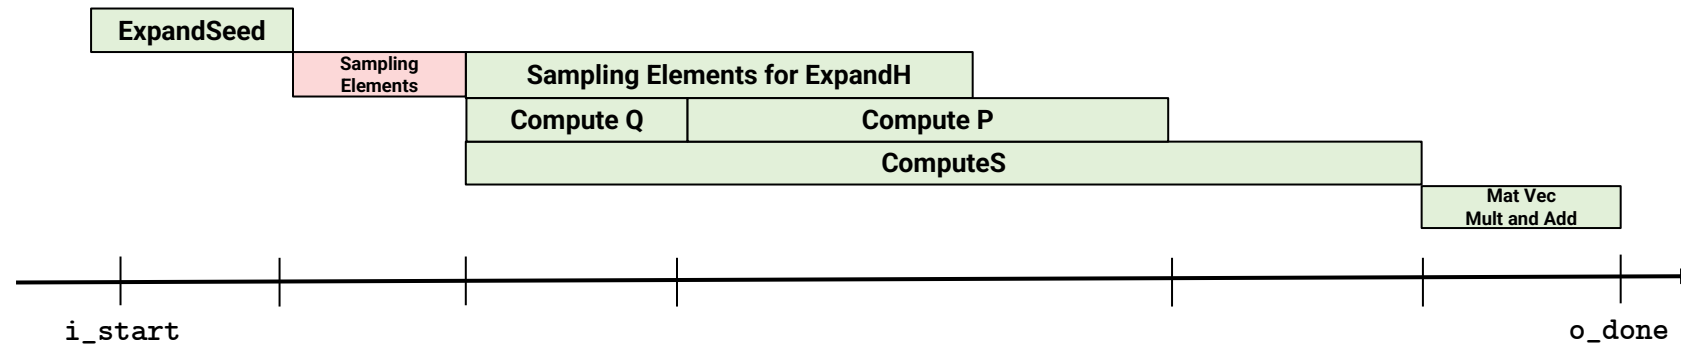


Our Contributions

- First parameterizable hardware realization of Hypercube Variant of SDitH Signature Scheme
- Two Variants of Syndrome Decoding Modules
 - Sample first, then multiply
 - Sample and multiply on the fly
- Split Hardware Implementation of Sign into *Offline* and *Online* phases
- Drastic Reduction in terms of Clock Cycles when compared to
 - Key Generation – Up to **250x**
 - Signature Generation – Up to **3.4x**
 - Signature Verification – Up to **2.2x**



Syndrome Computation Module – Key Generation



- $Y = s_B + H's_A$
- Syndrome Computation needs to be done after S (s_A, s_B) is computed by ComputeS module
- Hence, Sample First then Multiply approach (STFM)



Syndrome Computation Module – Sign and Verify

Algorithm 2a SDiH – Hypercube Variant – Signature Generation (Offline Part)

Input: a secret key $sk = (\text{seed}_H, y, \text{wit_plain})$ and a message $m \in \{0, 1\}^*$

1: $\text{salt} \leftarrow \{0, 1\}^{2\lambda}, \text{mseed} \leftarrow \{0, 1\}^\lambda$

2: $H' \leftarrow \text{ExpandH}(\text{seed}_H)$

Algorithm 4 ComputePlainBroadcast

Input: $\text{input_plain} := (\text{wit_plain}, \text{beav_ab_plain}, \text{beav_c_plain}), \text{chal}, (H', y)$

Output: broad_plain

1: $(s_A, Q', P) \leftarrow \text{Parse}(\text{wit_plain}, \mathbb{F}_q^k, (\mathbb{F}_q^{w/d})^d, (\mathbb{F}_q^{w/d})^d)$

2: $(a, b) \leftarrow \text{Parse}(\text{beav_ab_plain}, (\mathbb{F}_{q^n}^d)^t)$

3: $c \leftarrow \text{Parse}(\text{beav_c_plain}, \mathbb{F}_{q^n}^t)$

4: $(r, \varepsilon) \leftarrow \text{Parse}(\text{chal}, \mathbb{F}_{q^n}^t, (\mathbb{F}_{q^n}^d)^t)$

5: $s = (s_A \parallel y + H' s_A)$

6: $Q = \text{CompleteQ}(Q', 1)$

7: $S \leftarrow \text{Parse}(s, (\mathbb{F}_q^{m/d})^d)$

8: **for** $j \in [1 : t]$ **do**

9: **for** $\nu \in [1 : d]$ **do**

10: $\alpha[j][\nu] = \varepsilon[j][\nu] \otimes \text{Evaluate}(Q[\nu], r[j]) + a[j][\nu]$

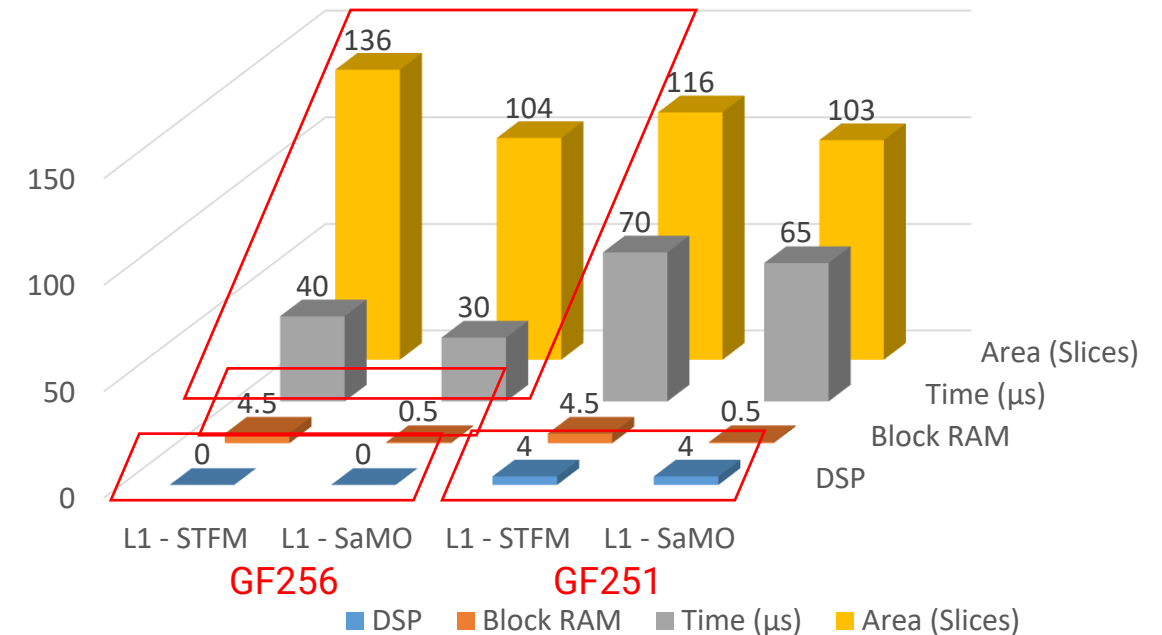
11: $\beta[j][\nu] = \text{Evaluate}(S[\nu], r[j]) + b[j][\nu]$

12: $\text{broad_plain} = \text{Serialize}(\alpha, \beta)$

13: **return** broad_plain

- y and S_a are inputs.
- Only need H' to compute the syndrome.
- “Sample and Multiply On the Fly” (SaMO) approach

Comparison of Syndrome Computation – STFM vs SaMO

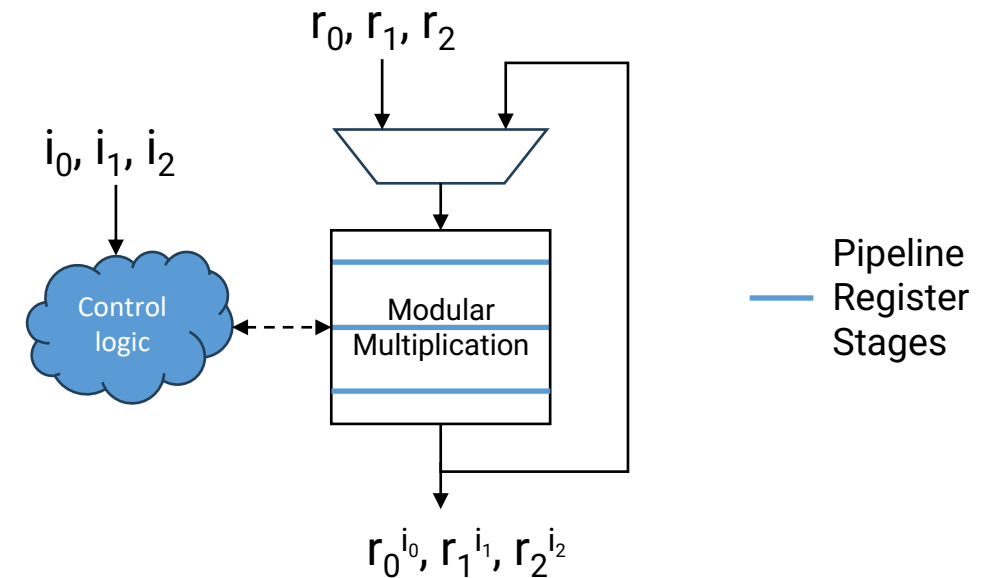


Evaluate Module

- Evaluate - takes as input an \mathbb{F}_q -vector Q representing the coefficients of a polynomial $\mathbb{F}_q[X]$ and a point $r \in \mathbb{F}_{\text{points}}$ and computes evaluation as follows:

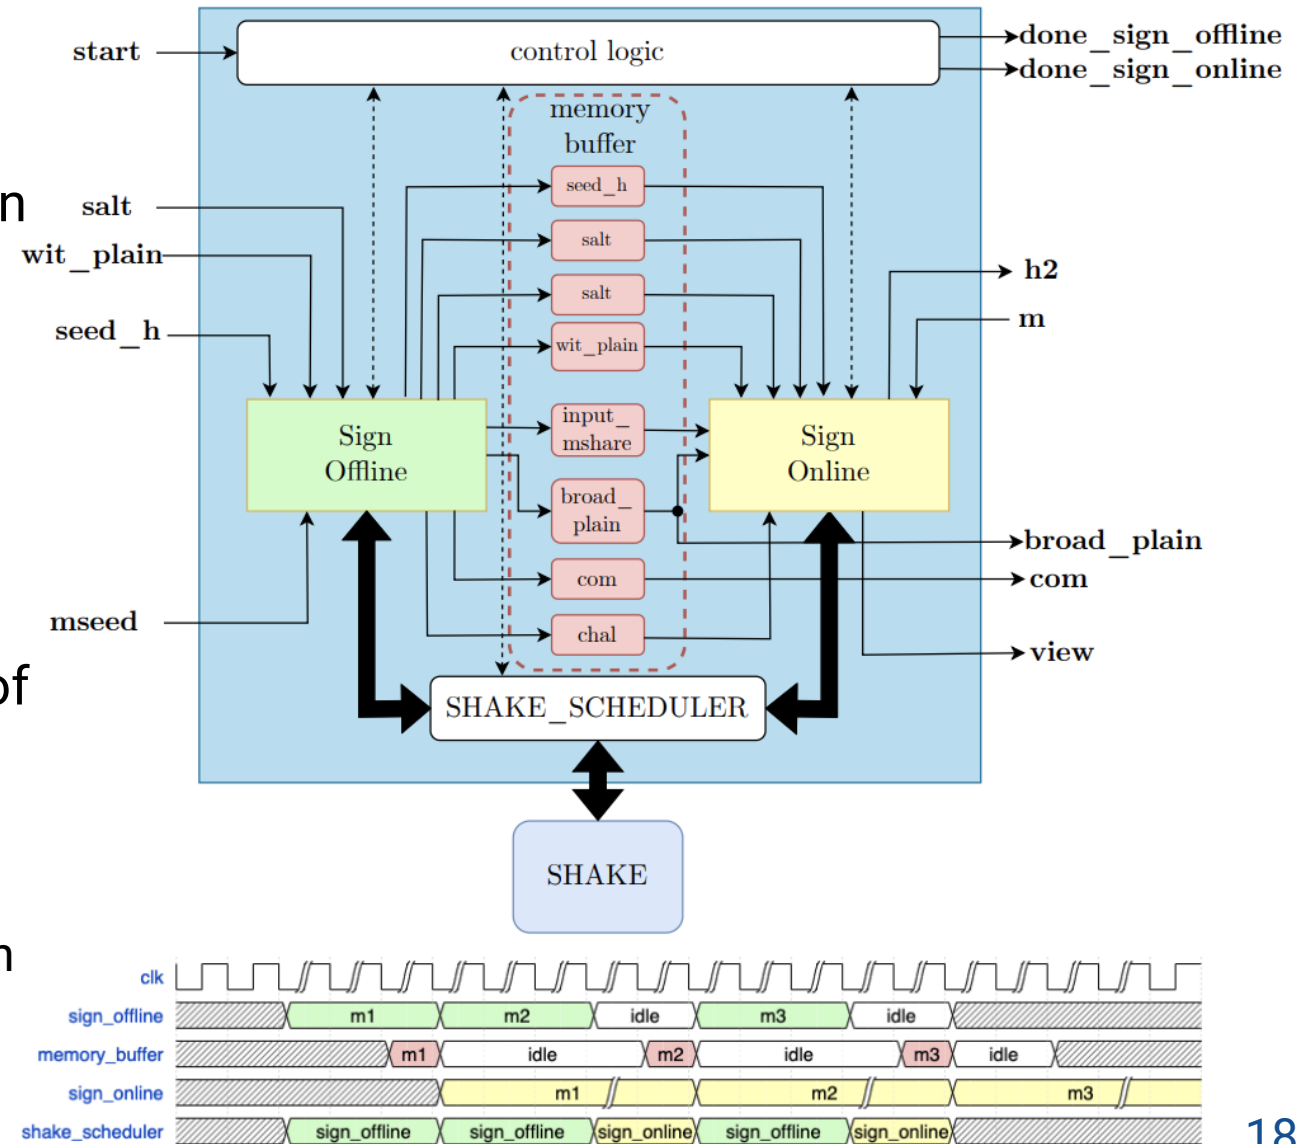
$$\text{Evaluate} : \begin{cases} \bigcup_{|Q|} (\mathbb{F}_q)^{|Q|} \times \mathbb{F}_{q^n} & \rightarrow \mathbb{F}_{q^n} \\ (Q, r) & \mapsto \sum_{i=1}^{|Q|} Q[i] \cdot r^{i-1} \end{cases} \quad \text{where } r^{i-1} = \underbrace{r \otimes r \otimes \dots \otimes r}_{i-1 \text{ times}} .$$

- Evaluate is used in both sign and verify operations. It contributes to:
 - 99% of cycles of the online part of the signing
 - 70%-90% of clock cycles in the verification based on the parameter set
- r^{i-1} is a 32-bit modular exponentiation; it is an expensive operation
 - Software implementation (target device Intel Xeon E-2378 CPU) accomplishes this by two large look-up-tables (370 KB to 1.5 MB for full design - based on parameter set)
 - Our lightweight target, Artix 7 FPGA, does not have these resources. Hence, we take an on-the-fly computation approach

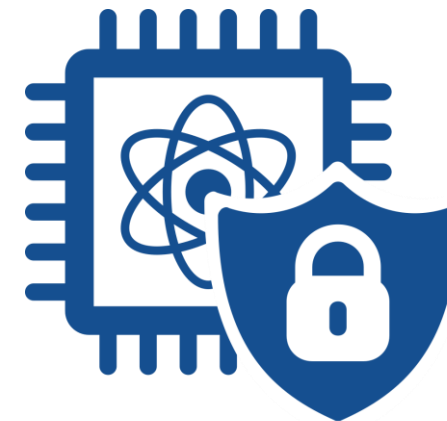


Signature Generation Module

- The block shown is for our area optimized implementation of SDitH signature generation scheme
- Signature generation is divided into two phases offline and online – they can run in parallel
- SHAKE256 is a hash function that is used in both the offline and online phases
- However, SHAKE256 is area expensive 31% of overall hardware design
- Hence, we design an optimized SHAKE scheduler such that
 - the same SHAKE module is switched between Offline and Online phases without wasting cycles and additional area

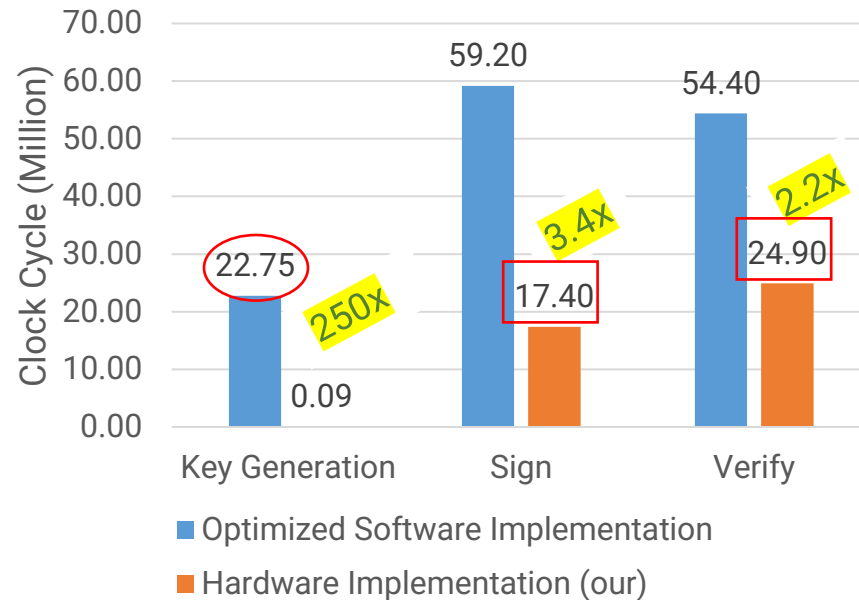


Comparison with Related and Relevant Work

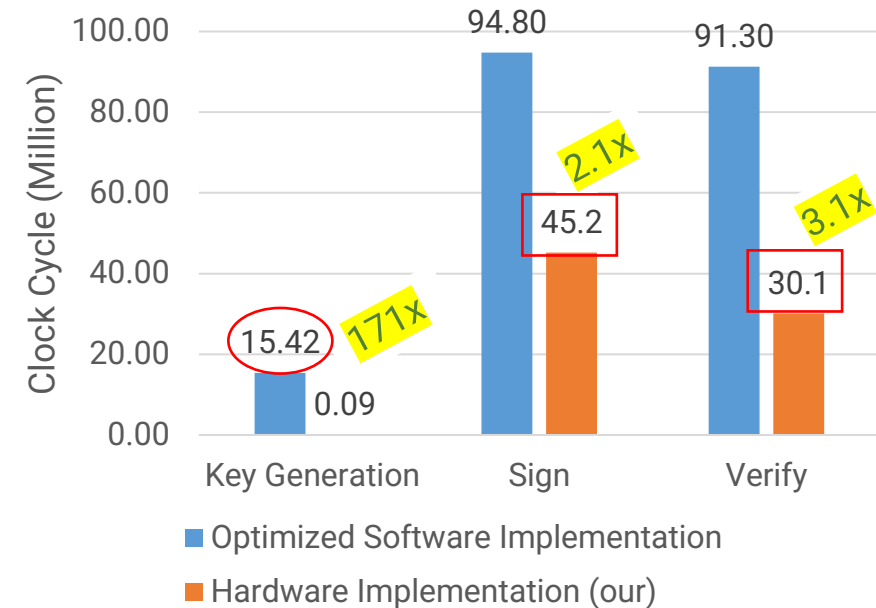


Clock Cycles Comparison – Optimized Software v/s Our Hardware Implementation – Hypercube Variant


Clock Cycle Comparison for Security Level 5 – GF256



Clock Cycle Comparison for Security Level 5 – GF251



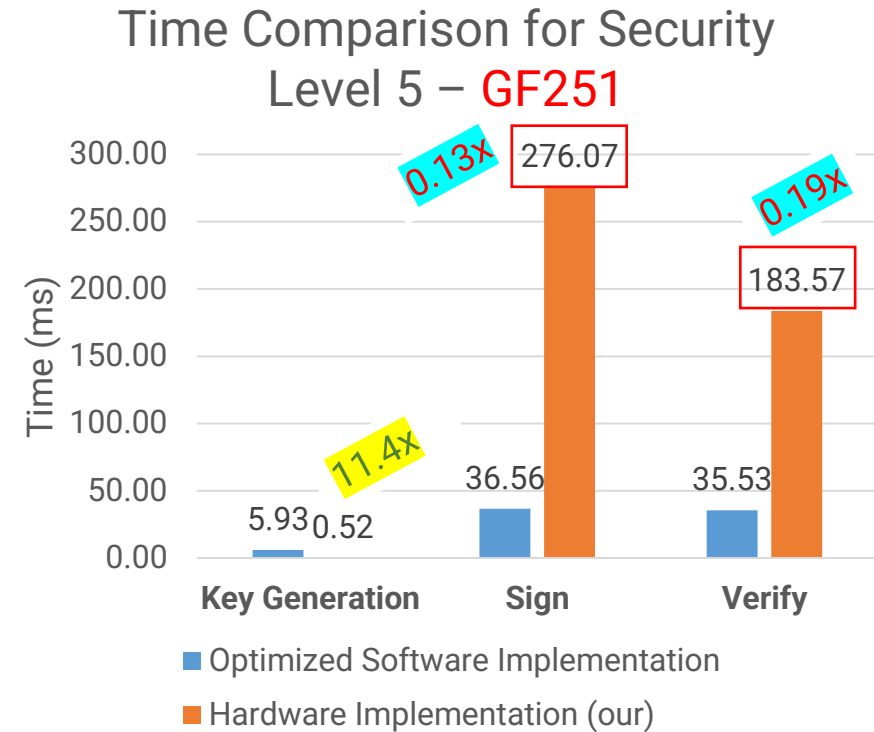
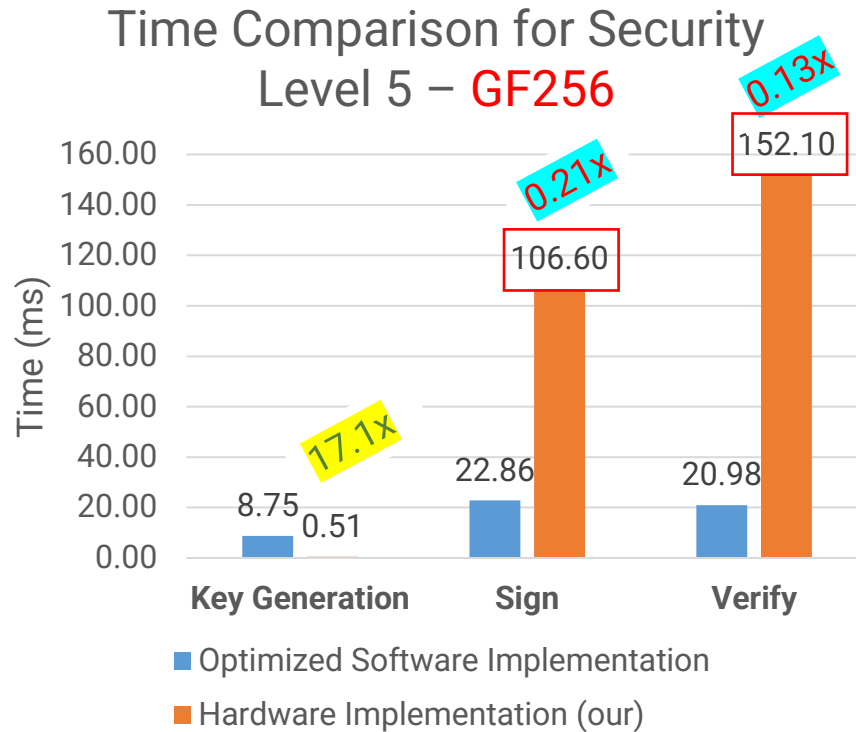
 Galois Field New Instructions from Intel are used

 ~70-99% of the clock cycles are taken in the 'sign_online' and 'verify' modules by the 'Evaluate' module

 Improvement
20

Time Comparison – Optimized Software v/s Our Hardware Implementation – Hypercube Variant

Operating Frequency:
 Intel Xeon Processor = 2.6 GHz
 Xilinx Artix 7 FPGA = 164 MHz

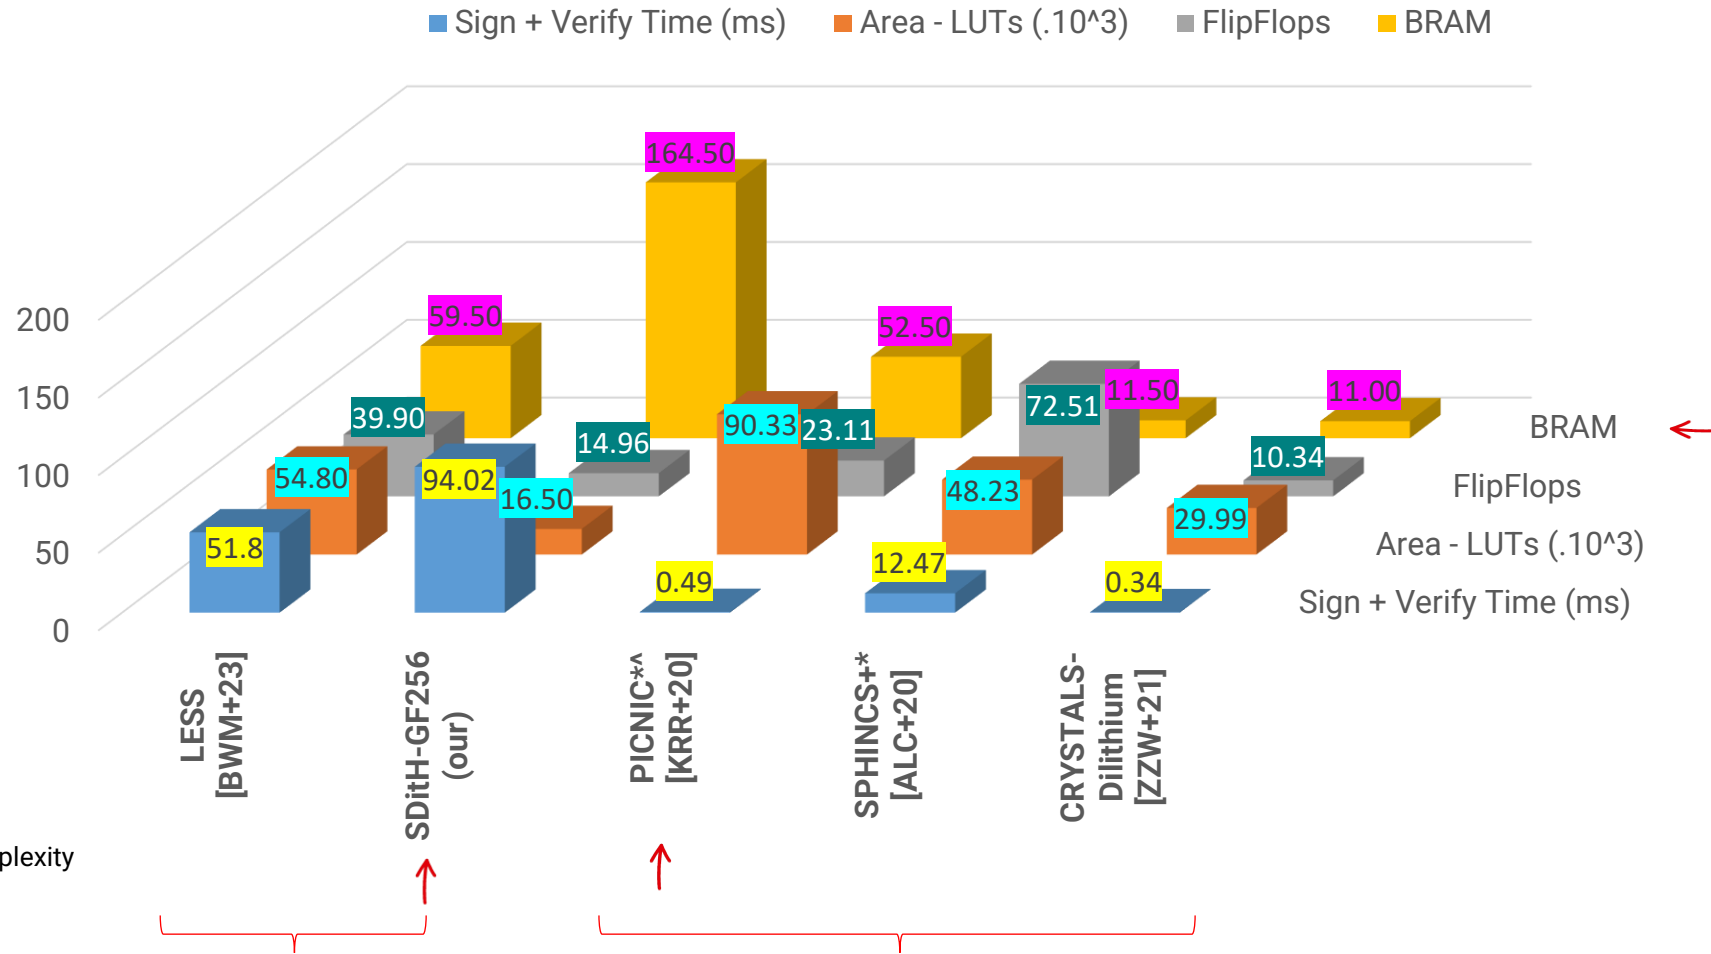


~70-99% of the clock cycles are taken in the 'sign_online' and 'verify' modules by the 'Evaluate' module

Improvement
 Decline



Comparison with other PQC-DSA candidates – Security Level 1

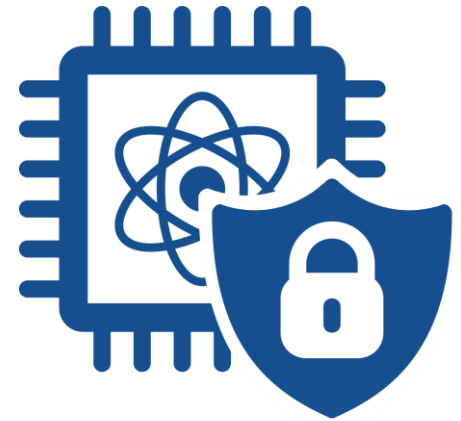


*No KeyGen
 ^Low Multiplication Complexity

Latest NIST Competition Candidates

Old NIST Competition Candidates

Conclusion and Future Work



Conclusion

- This work presents first hardware realization of SDitH Signature Scheme.
 - Parameterizable across three Security Levels and Two Arithmetic Fields.
- SDitH could be realized as a light-weight implementation. However, the memory consumption is higher.



Future Work

- The lower-level modules implemented as part of this work could be used to construct the 'threshold' variant of SDitH easily.
- Module level parallelism could be exploited to build a high-performance design which could speed-up the sign and verify operations.
- The MPC hardware modules' components could be reused outside SDitH.



References

[Gaj20] Kris Gaj, Implementation and Benchmarking of Round 2 Candidates in the NIST Post-Quantum Cryptography Standardization Process Using FPGAs, NIST Seminars, Oct 2020.

[BWM+23] Luke Beckwith, Robert Wallace, Kamyar Mohajerani, and Kris Gaj. A high-performance hardware implementation of the less digital signature scheme. In Thomas Johansson and Daniel Smith-Tone, editors, Post-Quantum Cryptography, pages 57–90, Cham, 2023. Springer Nature Switzerland.

[KRR+20] Daniel Kales, Sebastian Ramacher, Christian Rechberger, Roman Walch, and Mario Werner. Efficient FPGA implementations of lowmc and picnic. In Stanislaw Jarecki, editor, Topics in Cryptology – CT-RSA 2020, pages 417–441, Cham, 2020. Springer International Publishing.

[ZZW+23] Cankun Zhao, Neng Zhang, Hanning Wang, Bohan Yang, Wenping Zhu, Zhengdong Li, Min Zhu, Shouyi Yin, Shaojun Wei, and Leibo Liu. A compact and high-performance hardware architecture for CRYSTALS-Dilithium. IACR Transactions on Cryptographic Hardware and Embedded Systems, 2022(1):270–295, Nov. 2021.

[ALC+20] Dorian Amiet, Lukas Leuenberger, Andreas Curiger, and Paul Zbinden. Fpga-based sphincs+ implementations: Mind the glitch. In 2020 23rd Euromicro Conference on Digital System Design (DSD), pages 229–237, 2020.



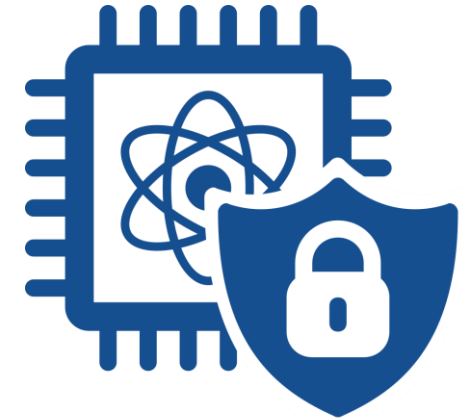
Sanjay Deshpande, James Howe, Jakub Szefer, and Dongze Yue, "SDitH in Hardware", in Transactions on Cryptographic Hardware and Embedded Systems (TCHES), September 2024.



<https://ia.cr/2024/069>

Thank you!

Sanjay Deshpande
email: sanjay.deshpande@yale.edu



Computer Architecture
and Security Lab (CASLAB)
<https://caslab.io>

