

BASALISC: Programmable Hardware Accelerator for BGV Fully Homomorphic Encryption

Robin Geelen¹, Michiel Van Beirendonck¹

H. V. L. Pereira¹, B. Huffman², T. McAuley³, B. Selfridge², D. Wagner², G. Dimou³,
I. Verbauwhede¹, F. Vercauteren¹, D. W. Archer²

¹COSIC KU Leuven, ²Galois, Inc., and ³Niobium Microsystems

CHES 2023, September 12

DARPA DPRIVE

- ▶ Dedicated ASIC acceleration of BGV
 - 150mm² 12nm chip
 - 10,000× faster than software reference
 - 128-bit security
 - Support *bootstrapping*



DARPA DPRIVE

- ▶ Dedicated ASIC acceleration of BGV
 - 150mm² 12nm chip
 - 10,000× faster than software reference
 - 128-bit security
 - Support *bootstrapping*
- ▶ Four teams of researchers
 - **Galois**, Duality, SRI, and Intel



| galois |

DARPA DPRIVE

- ▶ Dedicated ASIC acceleration of BGV
 - 150mm² 12nm chip
 - 10,000× faster than software reference
 - 128-bit security
 - Support *bootstrapping*
- ▶ Four teams of researchers
 - **Galois**, Duality, SRI, and Intel
- ▶ Several phases
 - **Phase 1**: design, implementation and verification of system architecture and IP blocks



| galois |

DARPA DPRIVE

- ▶ Dedicated ASIC acceleration of BGV
 - 150mm² 12nm chip
 - 10,000× faster than software reference
 - 128-bit security
 - Support *bootstrapping*
- ▶ Four teams of researchers
 - **Galois**, Duality, SRI, and Intel
- ▶ Several phases
 - **Phase 1**: design, implementation and verification of system architecture and IP blocks
 - **Phase 2**: software integration, hardware emulation, full system verification and validation, CKKS scheme



| galois |

Outline

- ① Hardware Design
- ② Bootstrapping
- ③ Evaluation
- ④ Ongoing Work

Outline

- ① Hardware Design
- ② Bootstrapping
- ③ Evaluation
- ④ Ongoing Work

BGV

- ▶ Ciphertexts are vectors of polynomials $c(X)$
 - Large polynomial multiplication through NTT: $C[k] = c(\omega^k)$
 - Large integer multiplication through RNS: $q = \prod q_i$

BGV

- ▶ Ciphertexts are vectors of polynomials $c(X)$
 - Large polynomial multiplication through NTT: $C[k] = c(\omega^k)$
 - Large integer multiplication through RNS: $q = \prod q_i$
 - 32-bit $q_i = c_i \cdot 2^k + 1$

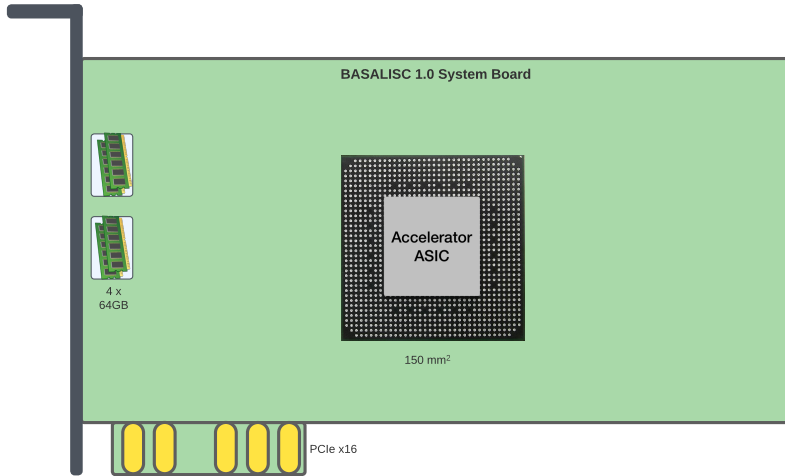
BGV

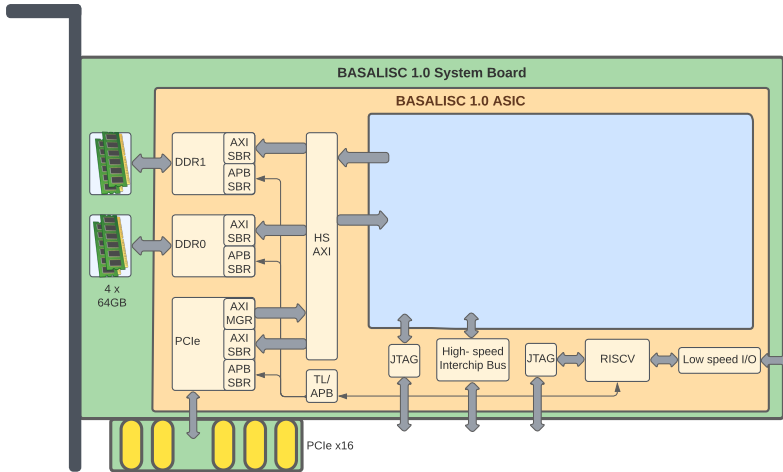
- ▶ Ciphertexts are vectors of polynomials $c(X)$
 - Large polynomial multiplication through NTT: $C[k] = c(\omega^k)$
 - Large integer multiplication through RNS: $q = \prod q_i$
 - 32-bit $q_i = c_i \cdot 2^k + 1$
- ▶ Basic Homomorphic Operations
 - Addition, multiplication, permutation
- ▶ Auxiliary Homomorphic Operations
 - Modulus switching, key switching
- ▶ Bootstrapping

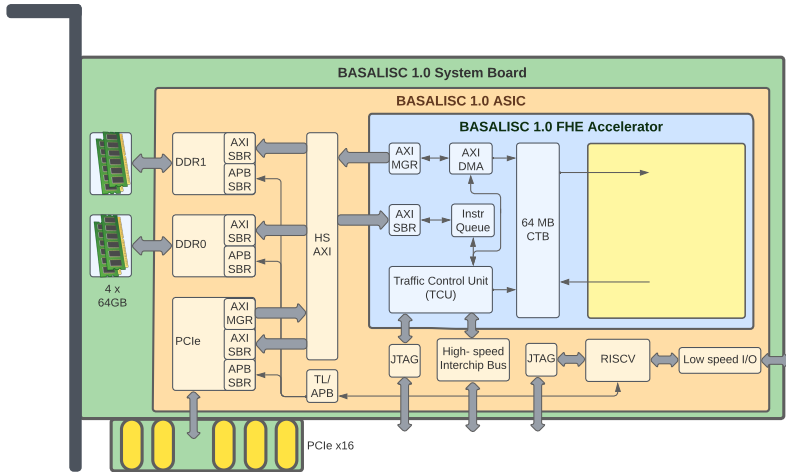
BGV 128-bit Security Parameters

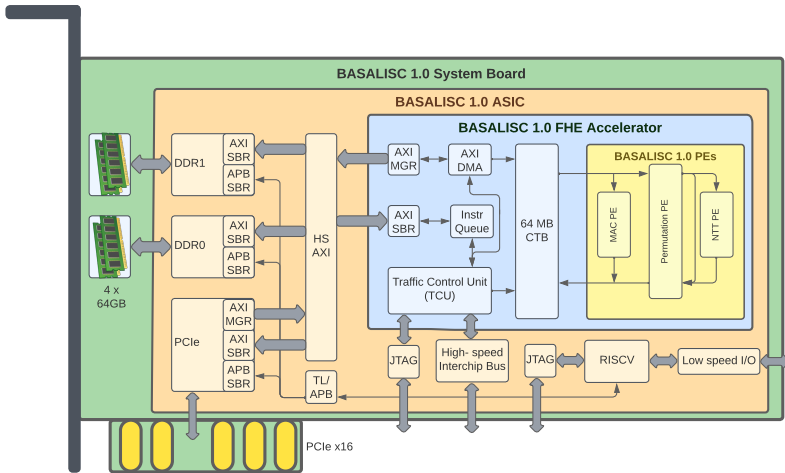
Parameter	Range	Example
Ring dimension N	$2^9 - 2^{16}$	2^{16}
Ciphertext packing ℓ	$2 - 2^{16}$	64 slots
Max $\log_2(QP)$ for key switching	20 - 1782	1782 bits
Max $\log_2(Q)$ for ciphertext	20 - 1782	1263 bits
Max multiplicative depth L	N/A	31

Ciphertext: 21 MB, Key-switch key: 112 MB










BASALISC High-Throughput NTT PE




Example radix- 2^2 butterfly

BASALISC High-Throughput NTT PE

1 Radix- 2^8 butterfly


- 2^{16} -point NTT with 2 passes
- 32 Tb/s NTT throughput



Example radix- 2^2 butterfly

BASALISC High-Throughput NTT PE


- 1 Radix- 2^8 butterfly
 - 2^{16} -point NTT with 2 passes
 - 32 Tb/s NTT throughput
- 2 Twiddle-factor factory unit



Example radix- 2^2 butterfly

BASALISC High-Throughput NTT PE


- 1 Radix- 2^8 butterfly
 - 2^{16} -point NTT with 2 passes
 - 32 Tb/s NTT throughput
- 2 Twiddle-factor factory unit
- 3 Optimized multipliers
 - $q_i = c_i \cdot 2^k + 1$



Example radix- 2^2 butterfly

BASALISC High-Throughput NTT PE

- 1 Radix- 2^8 butterfly
 - 2^{16} -point NTT with 2 passes
 - 32 Tb/s NTT throughput
- 2 Twiddle-factor factory unit
- 3 Optimized multipliers
 - $q_i = c_i \cdot 2^k + 1$
- 4 Conflict-free data layout



Example radix- 2^2 butterfly

BASALISC Conflict-free CTB Layout

NTT_1

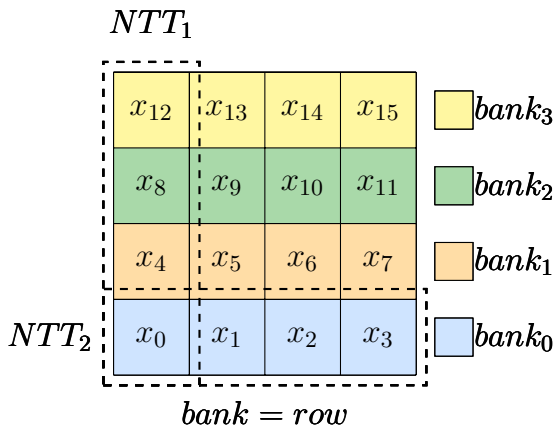
x_{12}	x_{13}	x_{14}	x_{15}
x_8	x_9	x_{10}	x_{11}
x_4	x_5	x_6	x_7
x_0	x_1	x_2	x_3

BASALISC Conflict-free CTB Layout

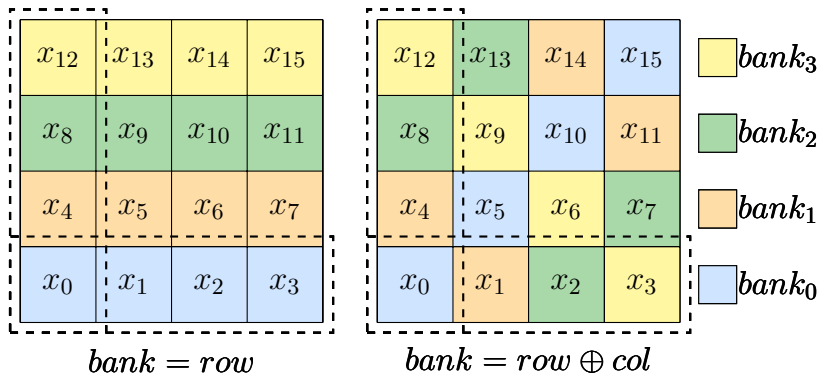
x_{12}	x_{13}	x_{14}	x_{15}
x_8	x_9	x_{10}	x_{11}
x_4	x_5	x_6	x_7
x_0	x_1	x_2	x_3

NTT_2

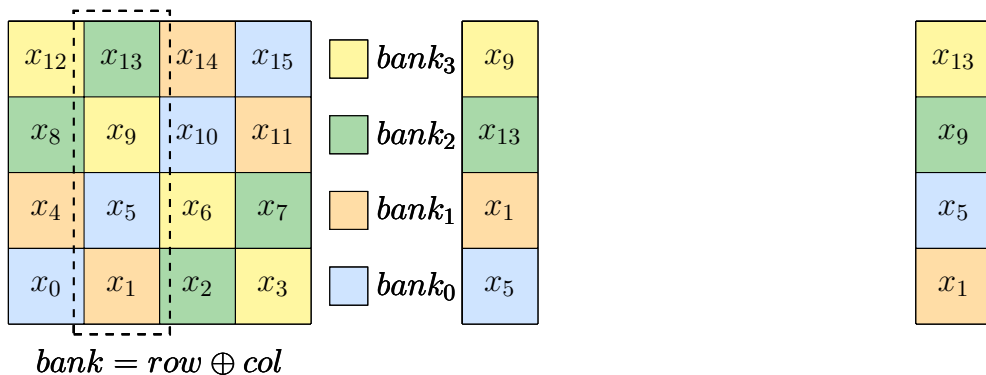
BASALISC Conflict-free CTB Layout



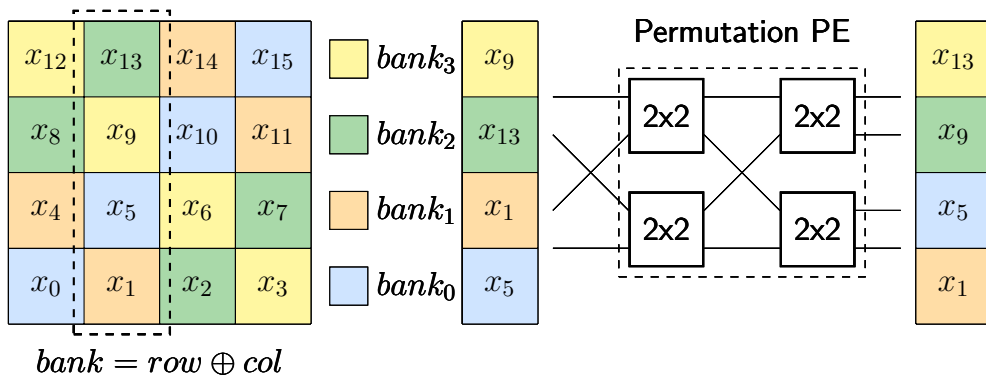
BASALISC Conflict-free CTB Layout



BASALISC Permutation PE




BASALISC Permutation PE




BASALISC Permutation PE

- ▶ Implements permutations: $i \mapsto i \oplus c$




BASALISC Permutation PE

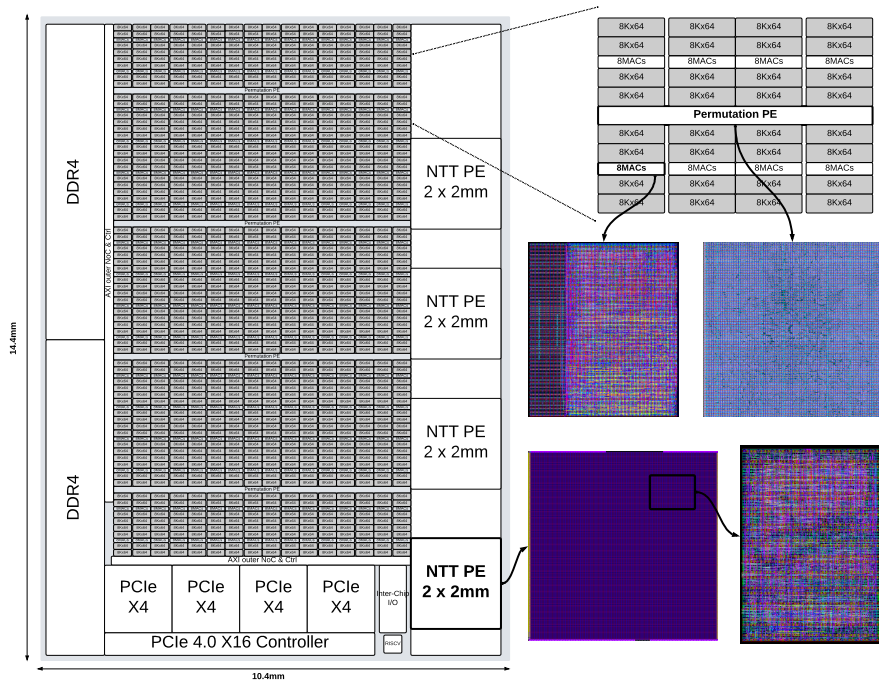
- ▶ Implements permutations: $i \mapsto i \oplus c$
- ▶ We slightly generalized: $i \mapsto (i \cdot a + b) \oplus c$



BASALISC Permutation PE

- ▶ Implements permutations: $i \mapsto i \oplus c$
- ▶ We slightly generalized: $i \mapsto (i \cdot a + b) \oplus c$
 - Supports BGV automorphisms: $\phi_k : c(X) \mapsto c(X^k)$ with the same hardware





Outline

- ① Hardware Design
- ② Bootstrapping
- ③ Evaluation
- ④ Ongoing Work

Bootstrapping Procedure

- ▶ **Homomorphic decryption** reduces noise

Bootstrapping Procedure

- ▶ **Homomorphic decryption** reduces noise
- ▶ Existing approaches (using odd prime p):
 - Ciphertext modulus $q = p^e + 1$
 - Exact division by p^k

Bootstrapping Procedure


- ▶ **Homomorphic decryption** reduces noise
- ▶ Existing approaches (using odd prime p):
 - Ciphertext modulus $q = p^e + 1$
 - Exact division by p^k
- ▶ **Problem:** Montgomery multipliers restricted to *NTT-friendly* primes
 - This means bits 16:1 are tied to 0 and bit 0 is tied to 1
 - But q and p^k are not NTT-friendly

NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k


NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k
- ▶ Simplified decryption function:



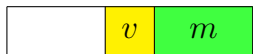
NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k
- ▶ Simplified decryption function:



NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k
- ▶ Simplified decryption function:



A horizontal bar divided into three segments: an empty white segment on the left, a yellow segment in the middle containing the variable v , and a green segment on the right containing the variable m .

$$(c_0, c_1)$$



A horizontal bar divided into four segments: an empty white segment on the left, a yellow segment containing v , a green segment containing m , and an empty white segment on the right.

$$c'_i \leftarrow [p^{e-r} c_i]_q$$





A horizontal bar divided into five segments: a red segment on the left containing r , an empty white segment, a yellow segment containing v , a green segment containing m , and an empty white segment on the right.

$$u \leftarrow c'_0 + c'_1 \cdot s$$





NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k
- ▶ Simplified decryption function:



NTT-friendly Bootstrapping

- ▶ Exclusive use of NTT-friendly primes
 - Ciphertext modulus $q = u \cdot p^e + 1$
 - No exact division by p^k
- ▶ Simplified decryption function:



NTT-friendly Bootstrapping

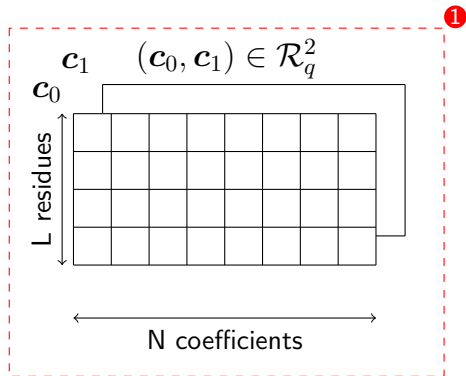
- ▶ In the paper:
 - Generalization to arbitrary q
 - Correction factor management
 - Instantiation in RNS

Outline

- ① Hardware Design
- ② Bootstrapping
- ③ Evaluation**
- ④ Ongoing Work

Software Toolchain


- ▶ Compiler for three-level instruction set:
 - 1 **Macro-instructions**: operations on ciphertexts and keys



Software Toolchain

► Compiler for three-level instruction set:


- 1 **Macro-instructions**: operations on ciphertexts and keys
- 2 **Mid-level instructions**: operations on residue polynomials



Software Toolchain

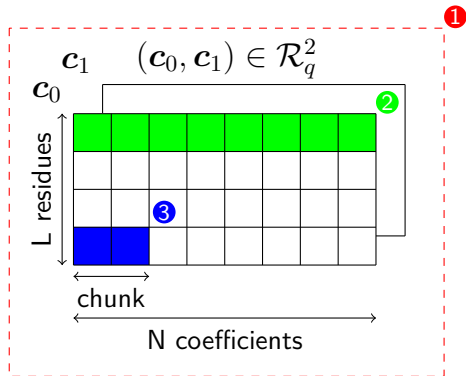
► Compiler for three-level instruction set:

- 1 **Macro-instructions**: operations on ciphertexts and keys
- 2 **Mid-level instructions**: operations on residue polynomials
- 3 **Micro-instructions**: operations on smaller chunks of data



Software Toolchain

- ▶ Compiler for three-level instruction set:
 - 1 **Macro-instructions**: operations on ciphertexts and keys
 - 2 **Mid-level instructions**: operations on residue polynomials
 - 3 **Micro-instructions**: operations on smaller chunks of data
- ▶ Performance and correctness simulator: works at micro-level



Performance

- ▶ HElib on single-core Intel Xeon @2.6GHz
- ▶ BASALISC @1GHz via simulation

	HElib	BASALISC	Speedup
NTT	27 ms	11 μ s	$2.5 \cdot 10^3 \times$
Add/Sub	4 ms	8 μ s	$5.0 \cdot 10^2 \times$
Multiply	58 ms	20 μ s	$2.9 \cdot 10^3 \times$
Permutation	12 ms	11 μ s	$1.1 \cdot 10^3 \times$
Key switching	580 ms	292 μ s	$2.0 \cdot 10^3 \times$

Performance

- ▶ HElib on single-core Intel Xeon @2.6GHz
- ▶ BASALISC @1GHz via simulation

	HElib	BASALISC	Speedup
NTT	27 ms	11 μ s	$2.5 \cdot 10^3 \times$
Add/Sub	4 ms	8 μ s	$5.0 \cdot 10^2 \times$
Multiply	58 ms	20 μ s	$2.9 \cdot 10^3 \times$
Permutation	12 ms	11 μ s	$1.1 \cdot 10^3 \times$
Key switching	580 ms	292 μ s	$2.0 \cdot 10^3 \times$
Bootstrapping	160 s	40 ms	$4.0 \cdot 10^3 \times$
DB lookup	2,325 s	267 ms	$8.7 \cdot 10^3 \times$
LR training	217,000 s	40,500 ms	$5.4 \cdot 10^3 \times$

Comparison with Different Accelerators

	BASALISC	F1	BTS	CraterLake
Scheme	BGV	BGV/CKKS	CKKS	CKKS
Area	150mm²	150mm ²	374mm ²	472mm ²
Technology	12nm	12/14nm	7nm	12/14nm
Power	115W	180W	163W	320W
Bootstrapping speedup	4,000×	1,830×/1,195×	2,237×	4,400×
LR speedup	5,400×	—/7,200×	1,306×	2,978×

Outline

- ① Hardware Design
- ② Bootstrapping
- ③ Evaluation
- ④ Ongoing Work

Changes Since End of DPRIVE Phase 1

- ▶ Support two schemes: CKKS and BGV
- ▶ Change from 32-bit to 64-bit q_i in RNS
 - Accommodate precision requirements in CKKS
 - Fewer NTT units due to increased area

Changes Since End of DPRIVE Phase 1

- ▶ Support two schemes: CKKS and BGV
- ▶ Change from 32-bit to 64-bit q_i in RNS
 - Accommodate precision requirements in CKKS
 - Fewer NTT units due to increased area
- ▶ On-chip PRNG eliminates 50% of key-switch memory bandwidth

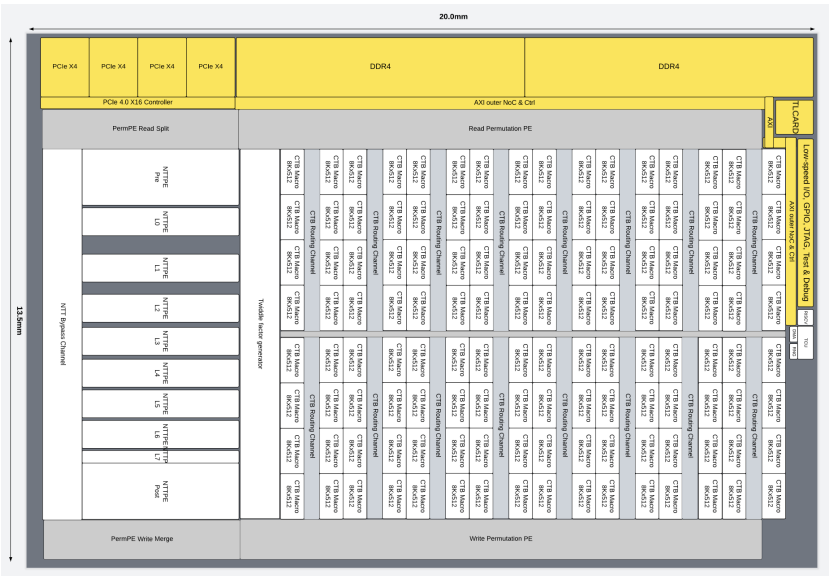
Changes Since End of DPRIVE Phase 1

- ▶ Support two schemes: CKKS and BGV
- ▶ Change from 32-bit to 64-bit q_i in RNS
 - Accommodate precision requirements in CKKS
 - Fewer NTT units due to increased area
- ▶ On-chip PRNG eliminates 50% of key-switch memory bandwidth
- ▶ Increased die size $270\text{mm}^2 = 20\text{mm} \times 13.5\text{mm}$

Current Status of BASALISC

- ▶ Floorplan and place and route complete
- ▶ On track for tape-out by November 30
- ▶ Verification tests in progress
- ▶ Sign-off flow on track

Latest BASALISC Floorplan



Conclusion

- ▶ New ASIC architecture for the BGV and CKKS schemes
- ▶ Combines efficiency with programmability
 - Novel processing elements
 - Bootstrapping enables optimized multipliers
 - Software toolchain
- ▶ Tape-out by the end of this year

Thank you for your attention!

robin.geelen@esat.kuleuven.be

michiel.vanbeirendonck@esat.kuleuven.be

BASALISC Die Plot

