A Flexible RNS-CKKS Processor for FHE-Based Privacy-Preserving Computing

Hyunhoon Lee*, Hyeokjun Kwon*, Youngjoo Lee
Pohang University of Science and Technology, Pohang, Korea

POSTECH
POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Abstract

Fully Homomorphic Encryption (FHE) has emerged as a crucial privacy-preserving solution for modern server systems handling sensitive data. Among FHE schemes, the CKKS approach based on Ring Learning with Error (RLWE) and the residue number system (RNS) is considered promising. However, efficient handling of FHE operations, particularly the bootstrapping step, remains a challenge due to significant computational costs. This paper proposes an integrated high-efficiency FHE processor tailored to meet the demands of RNS-CKKS schemes. The processor features novel design-level optimizations to reduce energy consumption and processing latency, including inter-/intra-set scheduling of residue polynomials and cost-reduced computing engines. Implemented in 28nm CMOS, the proposed processor demonstrates energy efficiencies outperforming recent works. The architecture includes dedicated computing engines for NTT/iNTT acceleration, base conversion, and arithmetic operations, managed by a top-level controller. The paper presents detailed designs for each computing engine, highlighting optimizations to support arbitrary input sizes and reduce on-chip memory requirements. Performance evaluation shows significant energy savings and latency improvements compared to existing architectures, making it a highly energy-efficient solution for RNS-CKKS-based FHE systems.

Proposed Design

❖ Proposed NTT/iNTT hardware engine

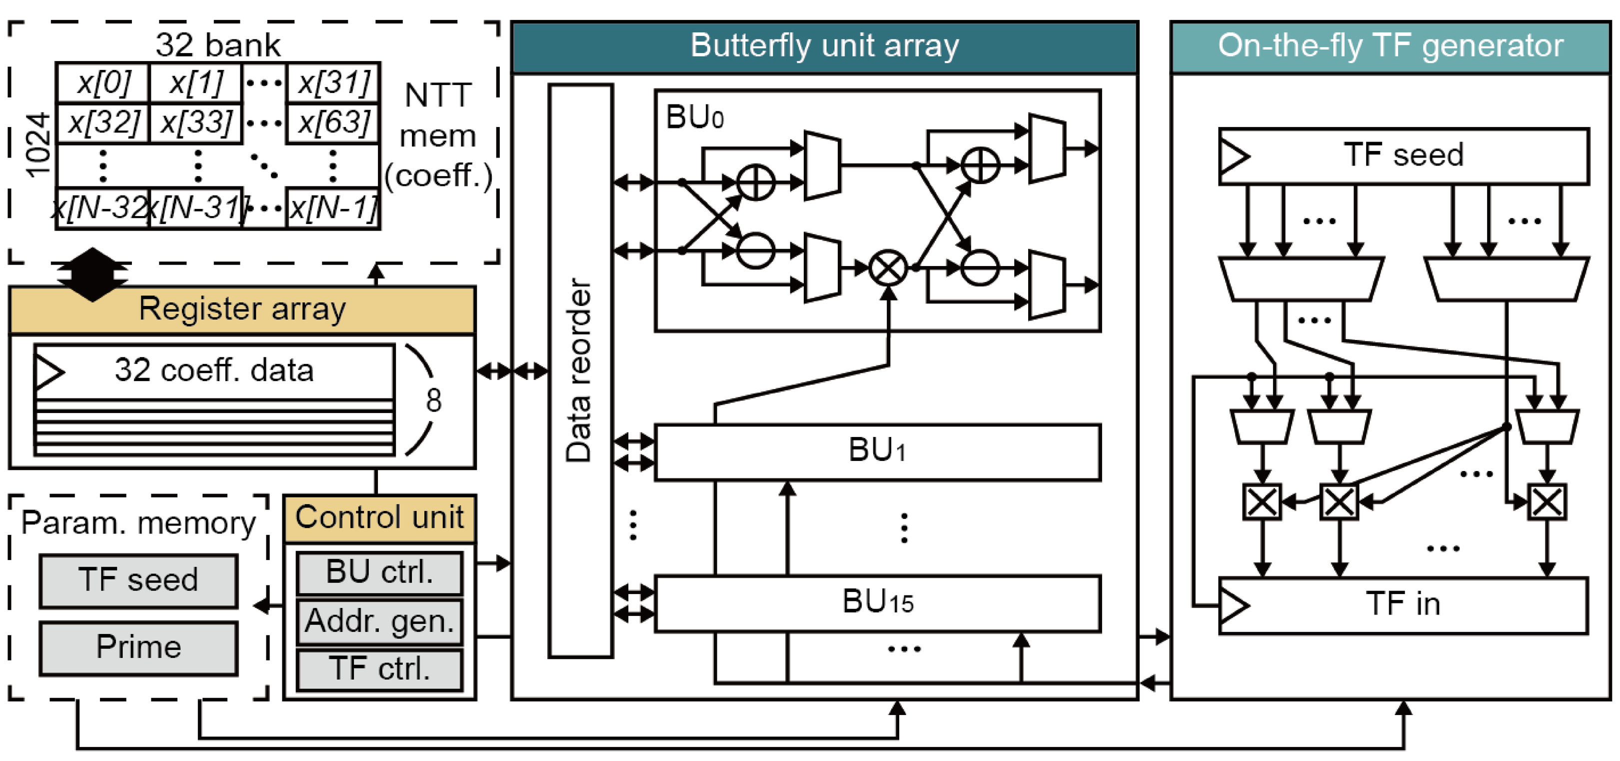

❖ Proposed base conversion (Bconv) engine

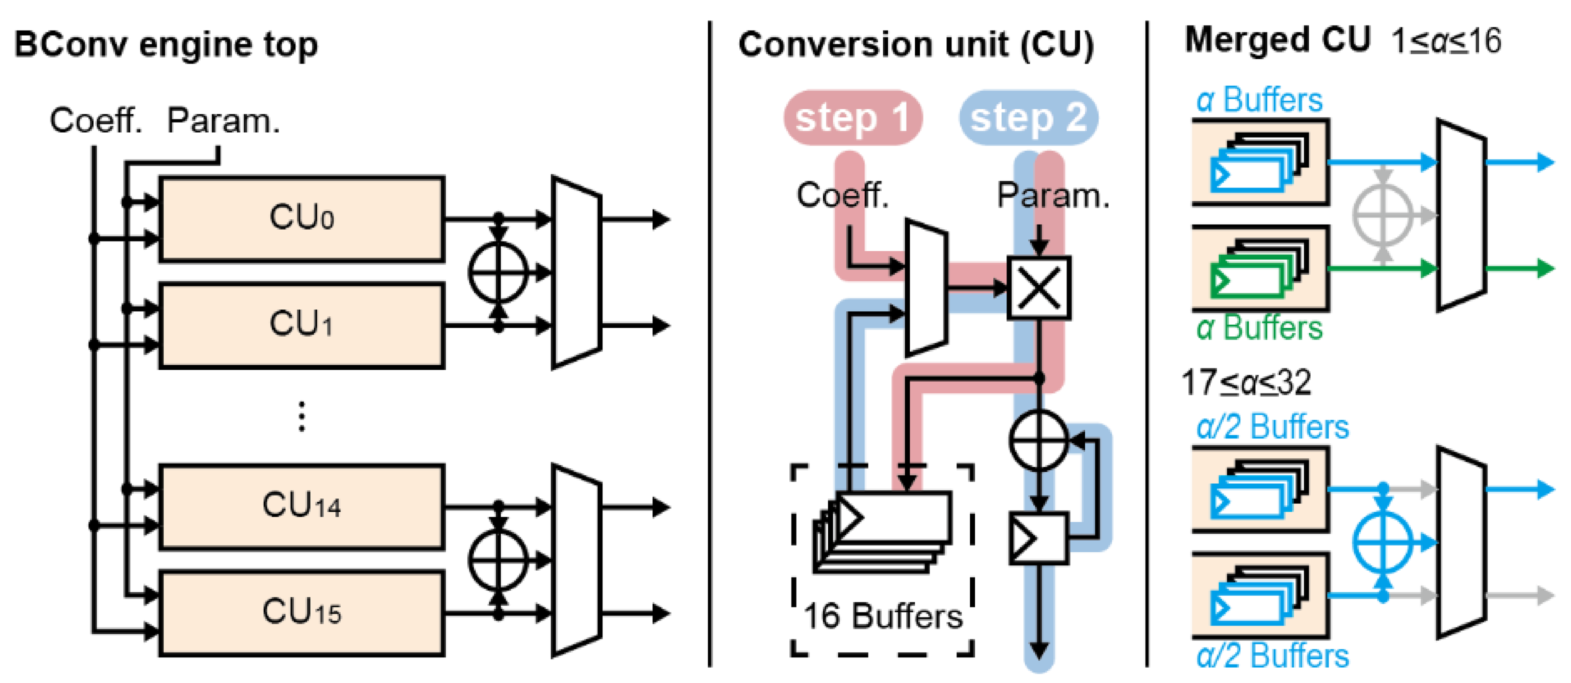

❖ Proposed modular arithmetic engine

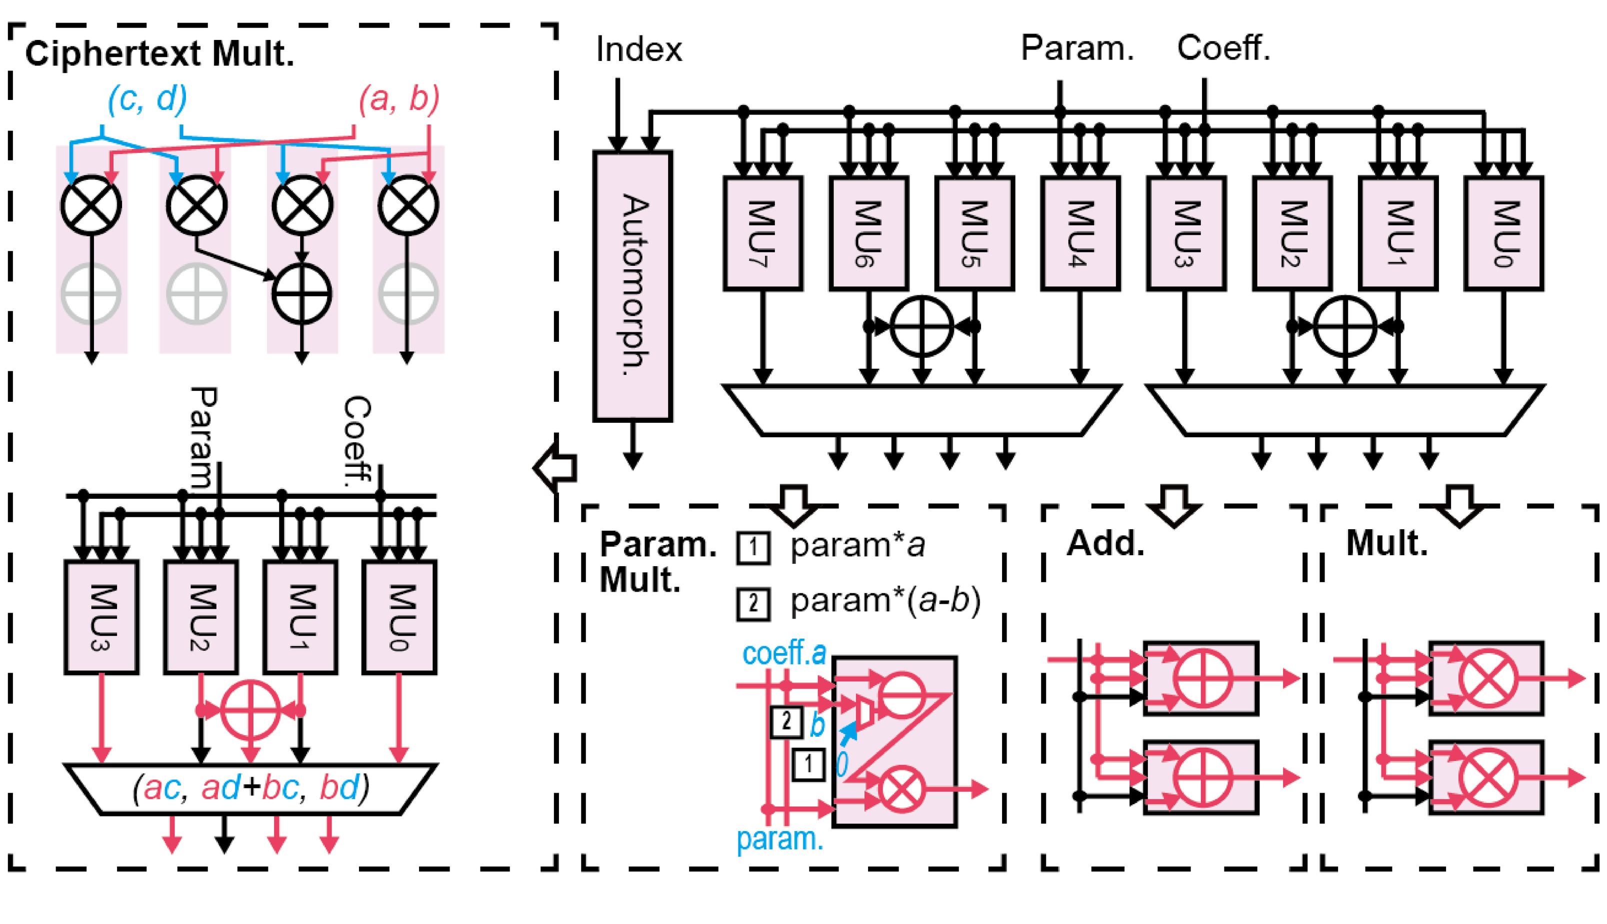

Energy-efficient scheduling system

❖ Intra- / Inter-set scheduling of residue polynomials

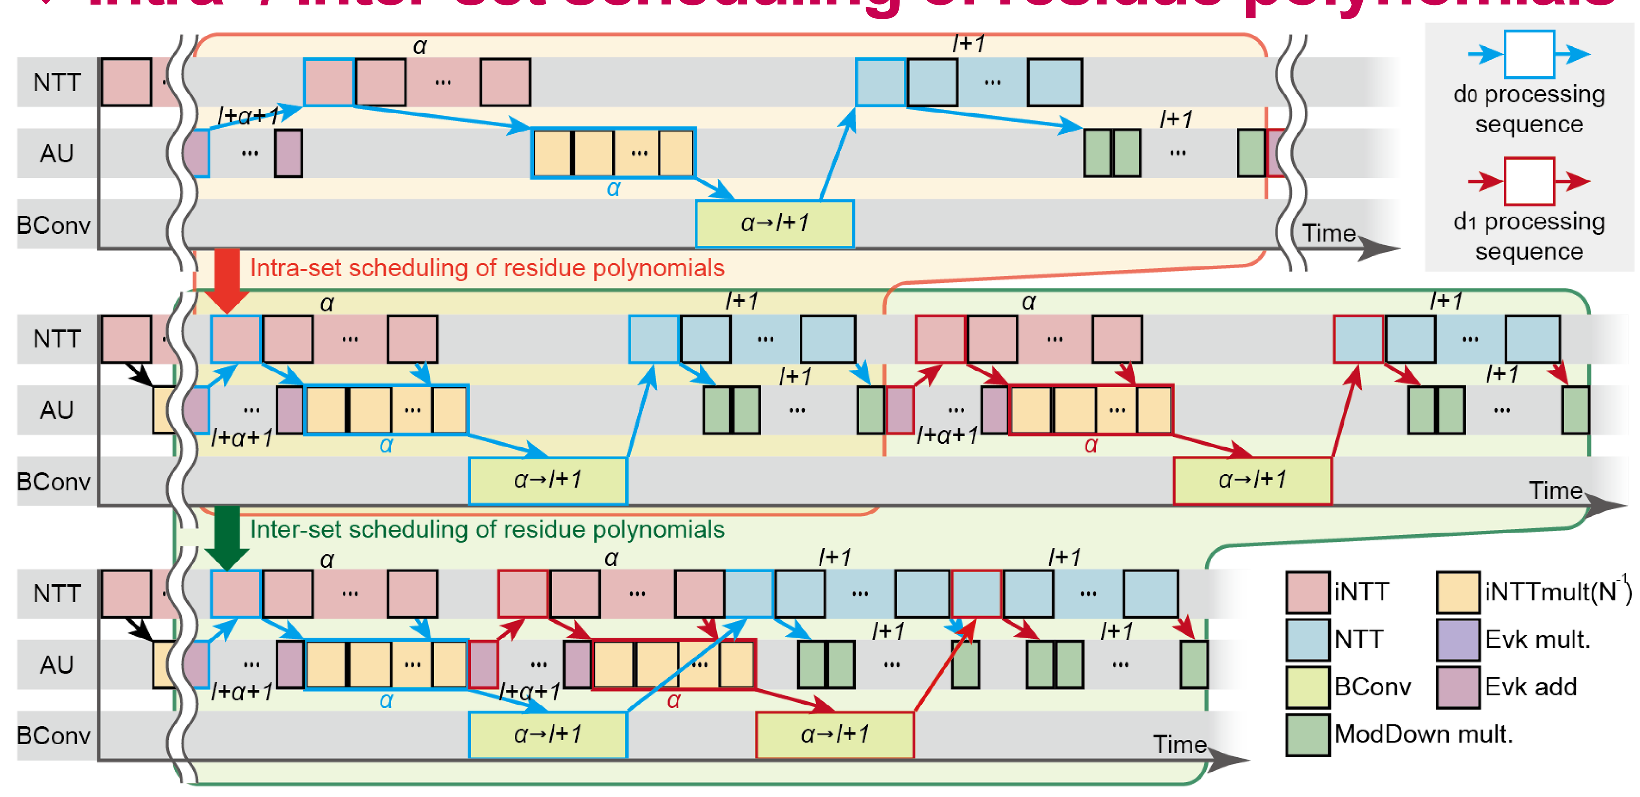

Implementation Results

❖ Processor layout

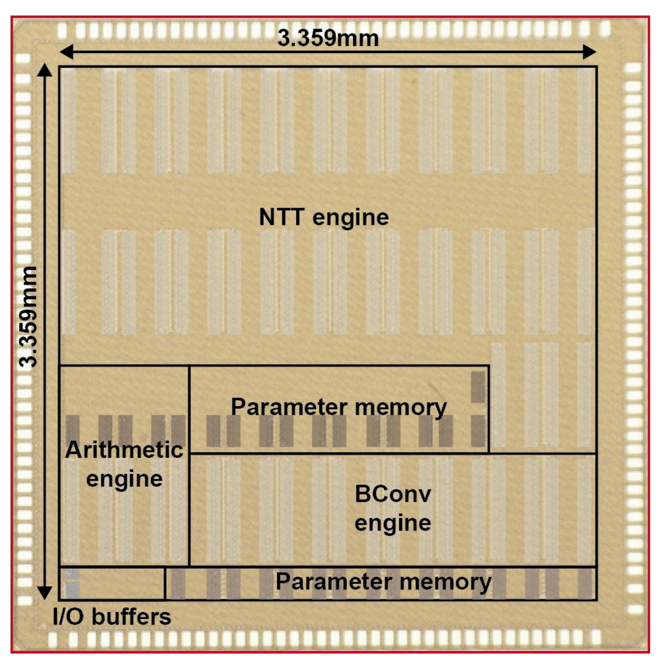

❖ Comparison to other state-of-the-art accelerators

		CICC'18 [1]	ISSCC'19 [2]	ISSCC'23 [3]	MICRO'21 [4]	ISCA'22 [5]	HPCA'23 [6]	This work		
Platform		ASIC	ASIC	ASIC	Architecture	Architecture	FPGA	ASIC		
Technology		40nm	40nm	28nm	12/14nm	12/14nm	16nm	28nm		
Frequency		300MHz	12-72MHz	500MHz	1GHz	1GHz	450MHz	333MHz		
Voltage		0.9V	0.68-1.1V	0.9V	-	-	-	1V		
Power		216.5mW	7~10mW	4W/12W	113W[a]	<320W	-	180mW		
Area		2.05mm²	0.28mm²	42.96mm²	54.56mm² [a]	240.5mm² [a]	-	11.28mm²		
Functionality	Application	PQC	PQC	HE	FHE	FHE	FHE	FHE		
	HE support	No	No	Paillier[b]	RNS-CKKS	RNS-CKKS	RNS-CKKS	RNS-CKKS		
	Supported HE operation	-	-	Partially (CAdd, PMult)	Fully (CAdd, CMult, Rot)	Fully (CAdd, CMult, Rot)	Fully (CAdd, CMult, Rot)	Fully (CAdd, CMult, Rot)		
	Flexible parameters	Bit-width, N	Bit-width, N	Bit-width[b]	Bit-width, N, l, α	Bit-width, N, l, α	Bit-width, N, l, α	Bit-width, N, l, α		
	$\log N$	6~11	6~11	-	~15	~17	~16	~17		
	Coefficient bit-width	<32 bit	<24 bit	-	<32 bit	<28 bit	<32 bit	<62 bit		
RNS-CKKS Bootstrapping	($\log N$, l_{max}, α)	-	-	-	(15,24,n/a)	(16,57,n/a)	(16,57,n/a)	(15,12,4)	(16,24,8)	(17,40,20)
	Security level	-	-	-	≈80	≈80	n/a	≈128	≈128	>128
	# of slots	-	-	-	1	32768	32768	16384	32768	65536
	Throughput	-	-	-	769.2boots/s	255.8boots/s	7.9boots/s	1.4boots/s	0.4boots/s	0.07boots/s
	Energy eff.	-	-	-	11.5mJ/boot[a]	775.7mJ/boot[a]	3267.8mJ/boot[a]	43.8mJ/boot	179.2mJ/boot	874mJ/boot
	Energy eff. per slot	-	-	-	11505μJ/boot/slot	23.7μJ/boot/slot	99.7μJ/boot/slot	2.7 μJ/boot/slot	5.5 μJ/boot/slot	13.3 μJ/boot/slot

[a] FHE accelerator part only
[b] Paillier does not belong to lattice-based cryptosystems; and therefore, it cannot be represented as a polynomial form.

[1] S. Song et al., "LEIA: A 2.05mm² 140mW Lattice Encryption Instruction Accelerator in 40nm CMOS," *IEEE CICC*, 2018.
[2] U. Banerjee et al., "An Energy-Efficient Configurable Lattice Cryptography Processor for the Quantum-Secure Internet of Things," *ISSCC*, pp. 46-48, 2019.
[3] G. Shi et al., "A 28nm 68MOPS 0.18μJ/Op Paillier Homomorphic Encryption Processor with Bit-Serial Sparse Ciphertext Computing," *ISSCC*, pp.242-243, 2023.
[4] N. Samardzic et al., "F1: A Fast and Programmable Accelerator for Fully Homomorphic Encryption," *IEEE/ACM MICRO,* pp.1295-1309, 2021.
[5] N. Samardzic et al., "CraterLake: A Hardware Accelerator for Efficient Unbounded Computation on Encrypted Data," *ACM/IEEE ISCA,* pp.173-187, 2022.
[6] Y. Yang et al., "Poseidon: Practical Homomorphic Encryption Accelerator," *IEEE HPCA*, pp.870-881, 2023.