

From Leuven with Love: An overview of arithmetic advances in the TFHE space



KU LEUVEN

T. de Ruijter¹, J.-P. D'Anvers^{1,2}, X. Pottier¹, M. Van Beirendonck² and I. Verbauwhede^{1,2}

¹COSIC, KU Leuven – ²Belfort Labs

Arithmetic advances in TFHE

- 2.06× faster bootstrapping or 50% noise reduction: Mean Compensation [1]
- 7.51× faster (Multi-)Scalar Multiplications through Pippenger style computations: SMOOTHIE [2]
- 2.75× faster Vector Addition through efficient Digit Extraction: Head Start [3]



COSIC



Belfort

Don't be mean: Reducing Approximation Noise through Mean Compensation [1]

- Compensate body based on expected error terms due to rounding
- $b_{\text{comp}} - \sum_{i=1}^n \hat{a}_i \cdot s_i = b - \mu_s \cdot \sum_{i=1}^n u_i - \sum_{i=1}^n \hat{a}_i \cdot s_i = m + e + \sum_{i=1}^n u_i \cdot (s_i - \mu_s)$
- Applies to modulus switching and gadget decompositions
- Allows parameters sets with up to 2.06× faster bootstrapping
- Already implemented and used in the TFHE-rs library [4]

	# rerand.	Variance compon. + body rounding		Mean compon.		Total noise	
		Measured	Rel.	Measured	Rel.	Measured	Rel.
Baseline [4]	/	$1.08 \cdot 10^{-6}$	100%	$1.07 \cdot 10^{-6}$	100%	$2.14 \cdot 10^{-6}$	100%
Rerand. [5]	15.3	$1.04 \cdot 10^{-6}$	96%	$2.72 \cdot 10^{-8}$	3%	$1.06 \cdot 10^{-6}$	49%
Mean comp.	/	$1.08 \cdot 10^{-6}$	100%	0	0%	$1.06 \cdot 10^{-6}$	50%
Hybrid	1.5	$1.05 \cdot 10^{-6}$	97%	0	0%	$1.05 \cdot 10^{-6}$	49%

References & Funding

- [1] T. de Ruijter, J.-P. D'Anvers, and I. Verbauwhede, "Don't be mean: Reducing approximation noise in TFHE through mean compensation," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2026, Jan. 2026.
- [2] X. Pottier, J.-P. D'Anvers, T. de Ruijter, and I. Verbauwhede, *SMOOTHIE: (multi-)scalar multiplication optimizations on TFHE*, Cryptology ePrint Archive, Paper 2025/1267, 2025. [Online]. Available: <https://eprint.iacr.org/2025/1267>
- [3] J.-P. D'Anvers, X. Pottier, T. de Ruijter, and I. Verbauwhede, *Head start: Digit extraction in TFHE from MSB to LSB*, Cryptology ePrint Archive, Paper 2025/2012, 2025. [Online]. Available: <https://eprint.iacr.org/2025/2012>
- [4] Zama, *TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data*, <https://github.com/zama-ai/tfhe-rs>, 2022.
- [5] O. Bernard, M. Joye, N. P. Smart, and M. Walter, "Drifting towards better error probabilities in fully homomorphic encryption schemes," in *Advances in Cryptology – EUROCRYPT 2025*, 2025.



ERC Adv. Grant No. 101020005; CSF No. VR20192203; FWO PhD SB No. 1S93126N

SMOOTHIE: (Multi-)Scalar Multiplication Optimizations On TFHE [2]

- TFHE adapted Pippenger method computing $\sum_{i=1}^k ct_i \cdot s_i$
- $ct \cdot s = ct \cdot \sum_{w=0}^{W-1} 2^{c \cdot w} s_w = \sum_{w=0}^{W-1} s_w \cdot ct_w$ and $ct_w = 2^{c \cdot w} ct$

$$ct \cdot s = \underbrace{\sum_{j=0}^{2^c-1} j \cdot B_j}_{\text{Bucket Aggregation}} \quad \text{and} \quad B_j = \underbrace{\sum_{w=0: s_w==j}^{W-1} ct_w}_{\text{Bucket Accumulation}}$$

- Reduces window count through sliding and negative windows
- Avoids full carry propagation through a global plaintext offset
- Replaces the multiplication in Bucket Aggregation by just two additions through Bucket Merging
- Large vector addition is the bottleneck: improved in Head Start

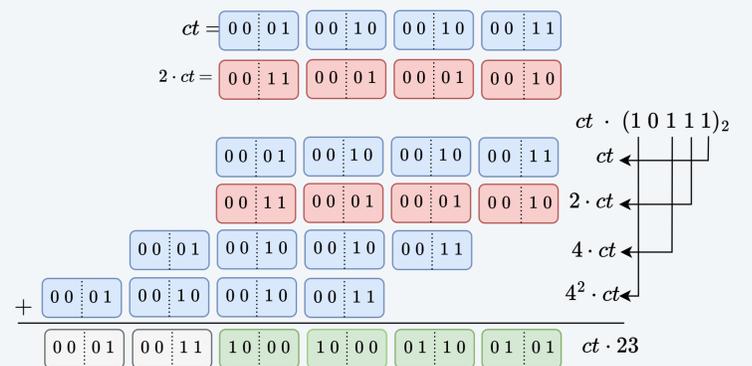
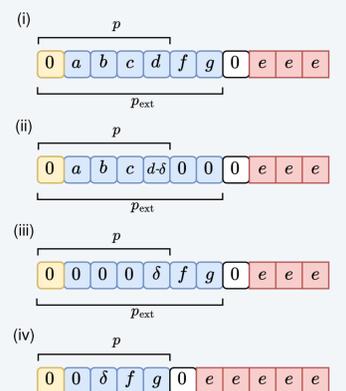


Figure 1: Scalar Multiplication as implemented in TFHE-rs

Head Start: Digit Extraction in TFHE from MSB to LSB [3]

- Faster vector addition through increased plaintext space
- Accurate extract using DirtyMSB: (i) Bootstrapping the MSBs returns (ii) an approximation of the MSBs with a possible error of δ . Subtracting (ii) from (i) yields (iii) a representation of the LSBs plus the error δ . After up-shifting (iv), this can be bootstrapped to correctly represent the input.



#Elements	TFHE-rs [4]	SMOOTHIE	Combined	Improv. over TFHE-rs
2400	111.3s	16.6s	8.1s	×13.73
4096	186.0s	26.9s	11.84s	×15.71
9216	421.6s	56.1s	25.19s	×16.74

Table 1: Computation time for a multi-scalar multiplication of 16-bit scalars and ciphertexts using 16 threads on an AMD EPYC 9174F 16-core processor.