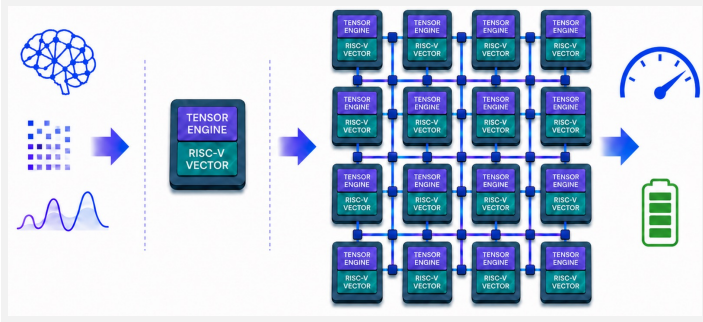


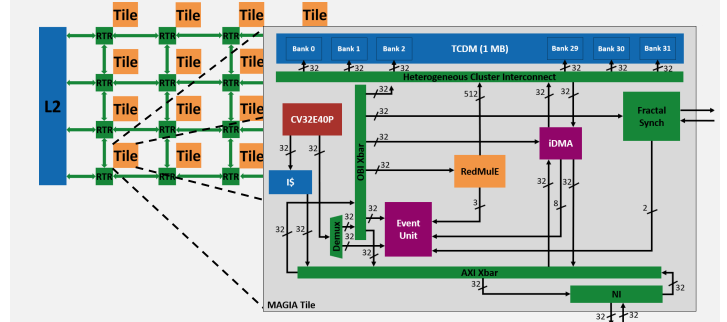
# MAGIA-V: A Heterogeneous Zve32d+GEMM Tile for Emerging Mesh-of-Tiles Accelerators

Luca Balboni<sup>1</sup>, Alessandro Nadalini<sup>2</sup>, Yvan Tortorella<sup>1</sup>, Riccardo Fiorani Gallotta<sup>1,2,3</sup>, Davide Rossi<sup>1,2</sup>, Francesco Conti<sup>1,2</sup>  
<sup>1</sup> Fondazione Chips-IT | <sup>2</sup> University of Bologna | <sup>3</sup> University of Pavia

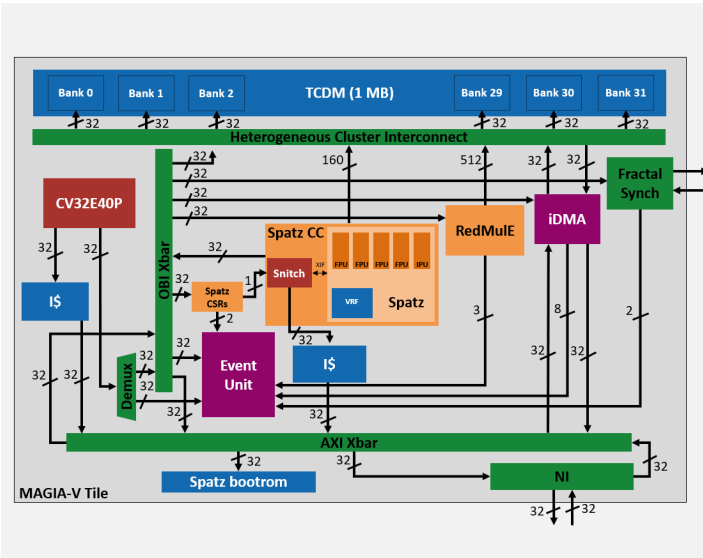
## MOTIVATION



## VANILLA MAGIA ARCHITECTURE



## MAGIA-V TILE


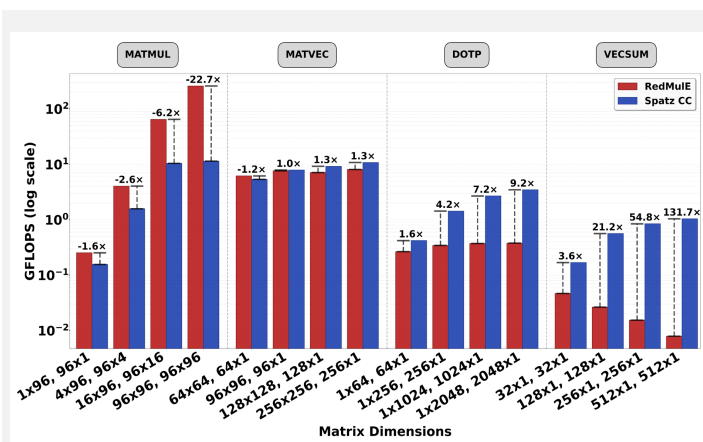


## PHYSICAL IMPLEMENTATION

**Technology:** GF 12nm  
**Target Freq:** 800 MHz  
**Max Freq:** 769 MHz

Component	Area (mm <sup>2</sup> )	Notes
MAGIA tile (baseline)	1.22	Total cell area
Memory macros	0.95	Dominant
Standard-cell logic	0.29	
<b>Spatz CC overhead</b>	<b>0.07</b>	<b>+5.8%</b>
P&R area (square)	1.69	76% Silicon Utilization

## BENCHMARKING:



## REFERENCES:

- [1] V. Isachi et al. "FractalSync: Lightweight Scalable Global Synchronization of Massive BSP AI Accelerators." CF '25, ACM, 2025, pp. 84–87.
- [2] M. Perotti et al. "Spatz: Clustering Compact RISC-V-Based Vector Units to Maximize Computing Efficiency." IEEE TCAD 44.7 (2025), pp. 2488–2502.
- [3] Y. Tortorella et al. "RedMule: A Mixed-Precision Matrix–Matrix Operation Engine for Flexible and Energy-Efficient On-Chip Linear Algebra and TinyML Training Acceleration." Future Gen. Comput. Syst. 149 (2023), pp. 122–135.