Reconfigurable vs. Specialized Architectures for Wireless Baseband Processing: A Spatial Array Perspective

Ali Rasteh, Ishaan Shivhare, Andrew Hennessee, Sundeep Rangan, Brandon Reagen

Motivation & Problem

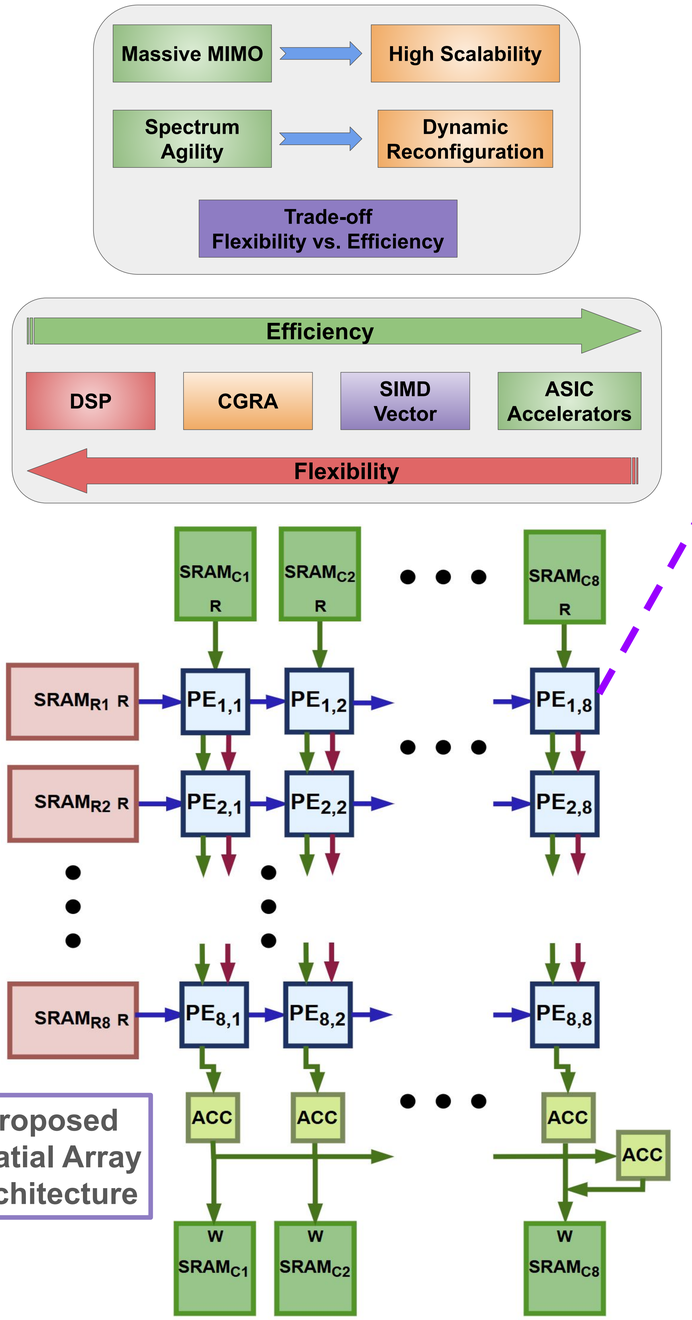

- Weights and data are injected from up and left side respectively.
- 2 operation modes: Accumulation & Element-wise.

Architecture & Results

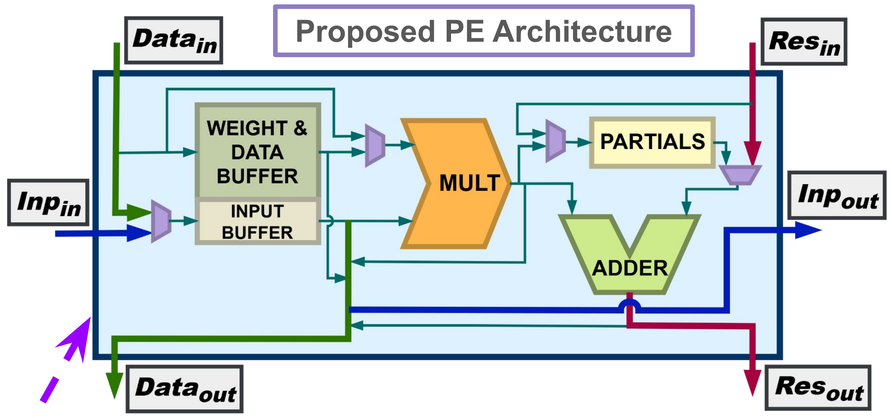

Kernel	Data Size	Weight Size	Data Type
FIR	1024x1	32x1	Complex
MatMat	1024x16	16x16	Real
MatVec	1024x16	16x1	Complex
OuterProduct	1024x32	1024x32	Complex
VectorMagSq	1024x1	-	Complex

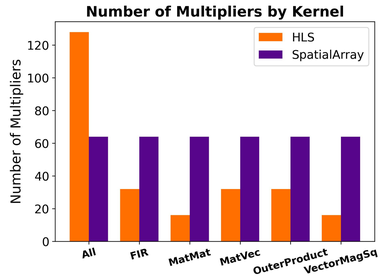

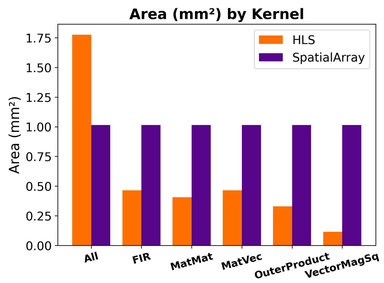

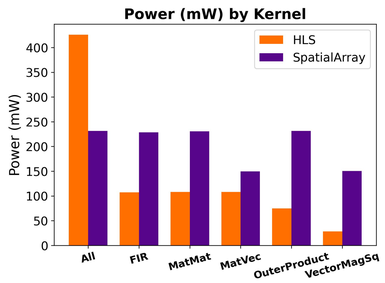

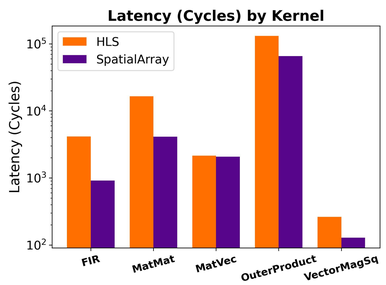

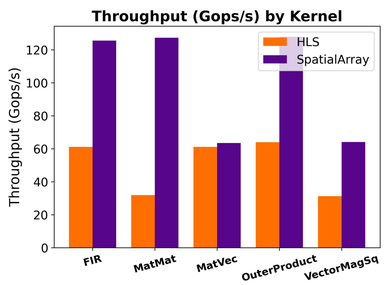

Technology node: SAED 32nm LVT

Comparisons & Takeaways

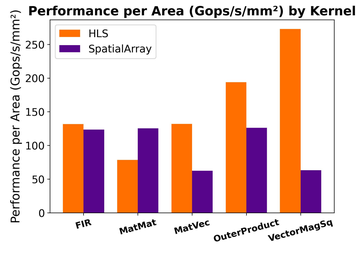

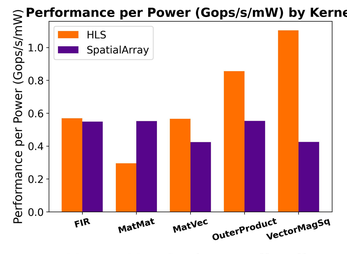

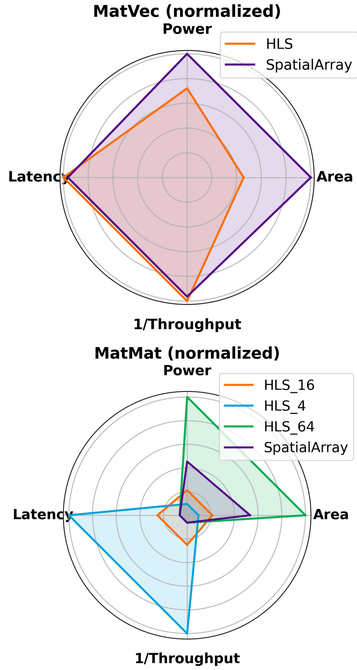

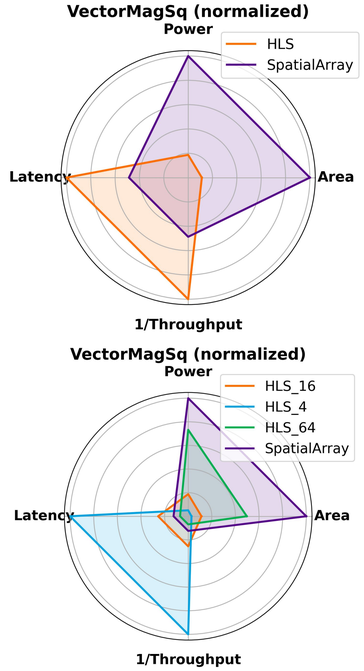

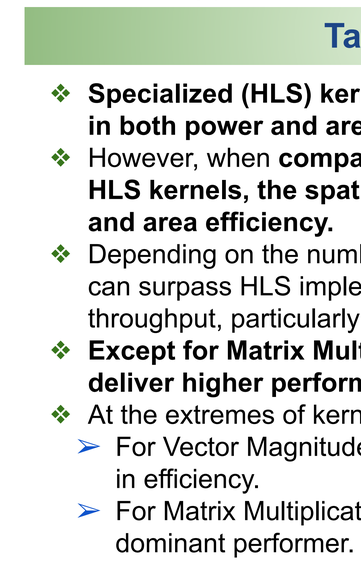

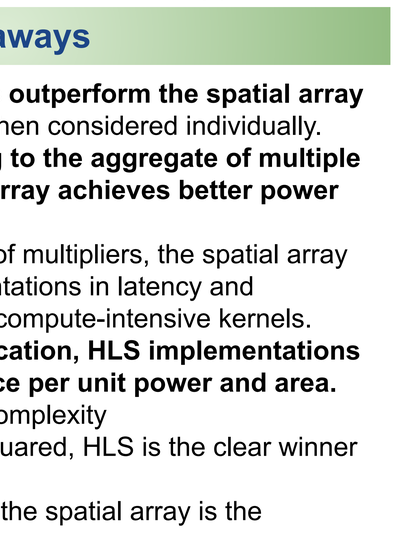

Takeaways

- **Specialized (HLS) kernels outperform the spatial array in both power and area** when considered individually.
- However, when **comparing to the aggregate of multiple HLS kernels, the spatial array achieves better power and area efficiency.**
- Depending on the number of multipliers, the spatial array can surpass HLS implementations in latency and throughput, particularly for compute-intensive kernels.
- **Except for Matrix Multiplication, HLS implementations deliver higher performance per unit power and area.**
- At the extremes of kernel complexity
  - For Vector Magnitude Squared, HLS is the clear winner in efficiency.
  - For Matrix Multiplication, the spatial array is the dominant performer.