

Vertex-Flow: A High-Throughput Spatial Accelerator for Synapse-Based Deep Learning on the Meridian Platform

Anonymous Author(s)

Affiliations withheld for double-blind review

Unofficial template. Delete this box before you submit. It is built on the IEEEtran class in the format the DAC 2027 submission page asks for, and it is not affiliated with or endorsed by DAC, IEEE or ACM. The results and references in this sample are fictional.

Abstract—Spatial accelerators dominate deep learning acceleration thanks to their parallelism and local memory hierarchies, but standard grid-based arrays fall short of the throughput and energy efficiency modern workloads need under tight power budgets. We introduce Vertex-Flow, a spatial accelerator for synapse-based deep learning workloads on the Meridian platform. A hybrid dataflow routes weights and activations dynamically to keep Processing Elements (PEs) busy, and localized weight buffering with adaptive accumulation routing cuts energy. Fabricated in a 12nm node, Vertex-Flow delivers $2.0\times$ the throughput and 46% lower energy per operation than state-of-the-art spatial arrays.

Index Terms—spatial accelerators, dataflow, deep learning hardware, energy efficiency

I. INTRODUCTION

Deep neural networks (DNNs) have revolutionized cognitive computing tasks, prompting significant interest in high-performance hardware design. General-purpose processors are ill-suited for the computational intensity of DNNs, leading to the rise of specialized accelerators such as Nexa’s TPU, Nexa-Core [1], and Synapse-V1 [2]. Spatial accelerators, consisting of two-dimensional grids of Processing Elements (PEs), offer high throughput by distributing computation across localized nodes and reducing the global memory access bottleneck.

However, spatial accelerators face challenges when executing diverse layers of DNNs (e.g., dense layers, sparse convolutions). Fixed dataflows, such as weight-stationary or output-stationary, limit the flexibility and utilization of the array. The Meridian platform [3] provides a robust emulation framework to evaluate alternative hardware architectures.

In this work, we propose **Vertex-Flow**, a spatial accelerator with an adaptive dataflow architecture designed to improve execution efficiency. Our main contributions are:

- We design a reconfigurable PE structure that supports both weight-stationary and hybrid activation-stationary modes.
- We introduce an on-chip network routing mechanism to dynamically manage accumulation paths.
- We evaluate Vertex-Flow on the Meridian platform and show substantial improvements in latency and power over baseline systems.

The rest of the paper is organized as follows: Section II reviews related work. Section III presents the proposed Vertex-

Flow architecture. Section IV details our experimental setup, and Section V discusses the results. Section VI provides a design space exploration, and Section VII concludes the paper.

II. RELATED WORK

Prior work on DNN accelerators has focused on optimizing memory architectures and scheduling algorithms. For instance, Nexa-Core [1] utilized a standard 2D systolic array with a weight-stationary dataflow, which suffers from low utilization during sparse convolution operations. To address this, Synapse-V1 [2] introduced row-stationary dataflow scheduling, mapping different filter rows to PE rows to reduce local buffer read energy.

Additionally, on-chip network design plays a crucial role in spatial accelerators. Nimbus-NoC [4] proposed a low-latency network-on-chip that uses virtual channels to reduce routing congestion in large-scale neuromorphic chips. Vertex-Alpha [5] scaled these accelerators to sub-10nm nodes but highlighted the severe thermal dissipation challenges associated with continuous high-frequency operations. Scheduling algorithms for these heterogeneous arrays have been proposed by Apex Research Lab [6], [7]. Dataflow optimizations for convolutional operations on these platforms were further optimized in [8]. Additionally, memory hierarchy design is a key design aspect in related architectures [9], [10]. Unlike these static or semi-static designs, Vertex-Flow dynamically adapts its accumulation pathways to match the tensor dimensions of each individual layer.

III. PROPOSED METHOD

The core architecture of Vertex-Flow comprises a hierarchical network of processing elements connected to a global weight buffer. The PEs are arranged in a 2D grid, where each node is capable of performing multiply-accumulate (MAC) operations and routing values to neighboring nodes.

A. PE Microarchitecture

Each PE contains a local register file for weights, an input register for activations, and an output accumulator. By using a local controller, the PE can choose to hold the weight locally while streaming activations (Weight-Stationary mode), or stream both weights and activations to accumulate results on the fly.




Fig. 1. Microarchitecture of the Vertex-Flow PE array with dynamic routing. Activation and weight streams are managed locally.

B. Dynamic Dataflow Routing

The primary innovation in Vertex-Flow is the dynamic configuration of data paths. During compilation, the layers are partitioned, and routing commands are generated. Let E_{tot} represent the energy required to execute a layer, which we model as:

$$E_{\text{tot}} = \sum_{l=1}^L \left(N_{\text{mac}}^{(l)} \cdot E_{\text{op}} + N_{\text{rd}}^{(l)} \cdot E_{\text{mem}} + N_{\text{rt}}^{(l)} \cdot E_{\text{noc}} \right) \quad (1)$$

where $N_{\text{mac}}^{(l)}$ is the number of MAC operations in layer l , E_{op} is the energy consumption of a single MAC operation, $N_{\text{rd}}^{(l)}$ is the count of buffer reads with memory energy E_{mem} , and $N_{\text{rt}}^{(l)}$ represents the number of NoC routing hops with routing energy E_{noc} . By optimizing the routing paths, we reduce $N_{\text{rt}}^{(l)}$ and minimize total power.

IV. EXPERIMENTAL SETUP

We evaluated Vertex-Flow by modeling the architecture in SystemC and simulating it on the Meridian platform [3]. The platform provides cycle-accurate emulation of spatial accelerators and computes power consumption based on cell library characterization.

Our baseline system is a conventional 2D systolic array with static weight-stationary execution. The benchmarks consist of representative workloads from the deep learning domain:

- **ResNet-50:** A widely used convolutional neural network for image classification.
- **BERT-Base:** A transformer-based model for natural language processing.
- **Fictional-Transformer:** An advanced model with dynamic token sparsity.

The simulation was carried out assuming a 12nm CMOS technology node with a clock frequency of 1.2 GHz. Memory energy was extracted using CACTI models for local SRAM buffers.

V. RESULTS AND ANALYSIS

The simulation results indicate substantial gains in both throughput and energy efficiency. Table I summarizes the performance comparison of Vertex-Flow against three baseline architectures.

Compared to Vertex-Alpha, our architecture doubles the throughput while operating under a similar technology node. This improvement is primarily driven by the high PE utilization rate, which reaches 92% on average across all ResNet-50 layers, compared to only 64% in Vertex-Alpha. Furthermore, our dynamic routing reduces the network-on-chip power by 35%, resulting in an overall energy efficiency of 0.78 pJ/OP.

VI. DISCUSSION

The trade-off in the Vertex-Flow design is the control logic overhead within each PE. The addition of the dynamic router increases the PE area by 8.4% compared to a static systolic array node. However, when considering the entire chip area, the area overhead is less than 2% since the SRAM buffers dominate the silicon area.

Furthermore, we investigated the scalability of Vertex-Flow by scaling the grid from 8×8 to 32×32 PEs. As shown in our scalability analysis, the routing delay scales sub-linearly with the array size due to the hierarchical structure of the NoC routing pathways.

VII. CONCLUSION

In this paper, we presented Vertex-Flow, a high-throughput spatial accelerator for synapse-based deep learning. By employing dynamic dataflow routing and localized weight buffering on the Meridian platform, Vertex-Flow resolves utilization bottlenecks inherent in traditional systolic arrays. Our experimental evaluation in 12nm technology demonstrates a throughput of 2048.8 GOPS and an energy efficiency of 0.78 pJ/OP, representing a significant advancement over state-of-the-art accelerators. Future work will investigate thermal-aware task scheduling and dynamic voltage scaling on the Vertex-Flow grid.

REFERENCES

- [1] S. Jenkins and D. Chen, "A 28nm spatial accelerator for edge deep learning," *Journal of Fictional Microarchitecture*, vol. 12, no. 3, pp. 45–56, 2024.
- [2] H. Tanaka and E. Rostova, "Synapse-v1: Reconfigurable dataflow architectures," in *Proceedings of the International Conference on Fictional Hardware*, 2023, pp. 112–120.
- [3] A. Vance and L. Thorne, "The meridian platform: A unified framework for hardware-in-the-loop emulation," *Systems Emulation Journal*, vol. 8, no. 2, pp. 130–142, 2021.
- [4] C. Dupont and R. Smith, "Nimbus-noc: Low-latency on-chip networks for spatial arrays," in *Symposium on Fictional Networks*, 2022, pp. 45–52.
- [5] M. Vance and S. Patel, "Vertex-alpha: Scaling spatial accelerators to sub-10nm nodes," *Transactions on Computer-Aided Systems*, vol. 34, no. 1, pp. 89–101, 2025.
- [6] A. Okafor and G. Bennett, "Energy-efficient scheduling algorithms for heterogeneous cores," in *Design Automation Workshop*, 2023, pp. 201–210.
- [7] E. Rostova and L. O'Connor, "Thermal-aware mapping of neural networks on apex cores," in *Proceedings of the Fictional Design Automation Conference*, 2024, pp. 312–320.
- [8] Y. Sato and C. Oswald, "Dataflow optimizations for dilated convolutions on synapse systems," *Journal of Systems Design*, vol. 15, no. 4, pp. 233–245, 2022.
- [9] M. Chang and S. Jenkins, "Heterogeneous memory systems for tile-based processing," *Fictional VLSI Letters*, vol. 9, no. 1, pp. 14–18, 2023.
- [10] T. Wright and S. Martinez, "Vertex-mesh: A tiled architecture for low-power neural inference," in *International Conference on Integrated Circuits*, 2024, pp. 67–75.

TABLE I
PERFORMANCE AND ENERGY EFFICIENCY COMPARISON ON MERIDIAN PLATFORM

Architecture	Node [nm]	Freq. [MHz]	Throughput [GOPS]	Energy [pJ/OP]
Nexa-Core [1]	28	500	128.4	4.56
Synapse-V1 [2]	16	800	512.0	2.10
Vertex-Alpha [5]	12	1000	1024.0	1.45
Vertex-Flow (Ours)	12	1200	2048.8	0.78