



Sparseloop: **An Analytical Approach to** **Sparse Tensor Accelerator Modeling**

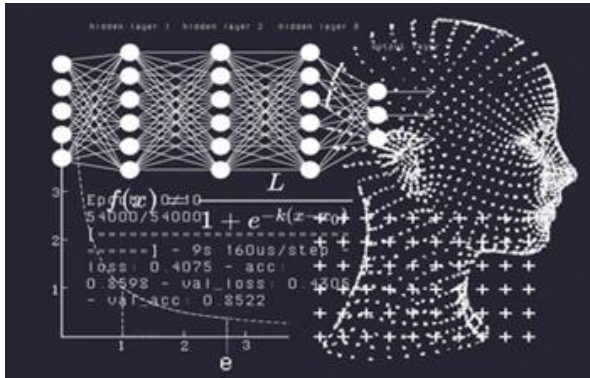
Yannan Nellie Wu¹, Po-An Tsai², Angshuman Parashar²,
Vivienne Sze¹, Joel Emer^{1,2}

¹MIT


²NVIDIA

<http://sparseloop.mit.edu/>

Many Applications Use Sparse Tensor Algebra



Sparse Neural Networks



Circuit Simulations

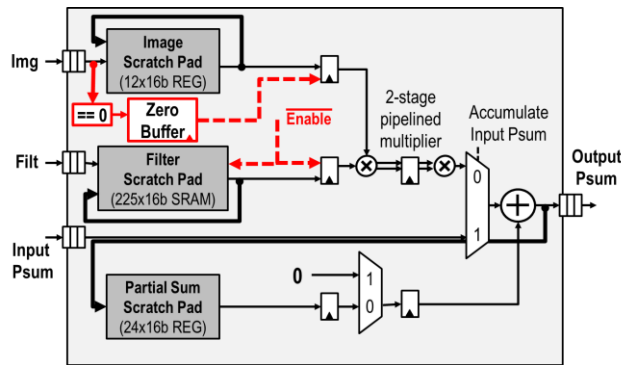


Data Science

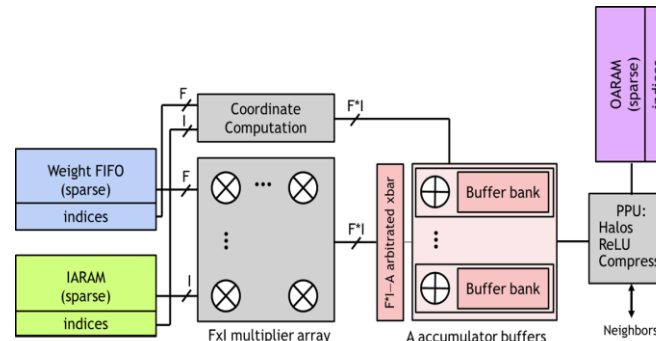
Inefficient Processing on General-Purpose Processors

An Explosion of Sparse Tensor Accelerators

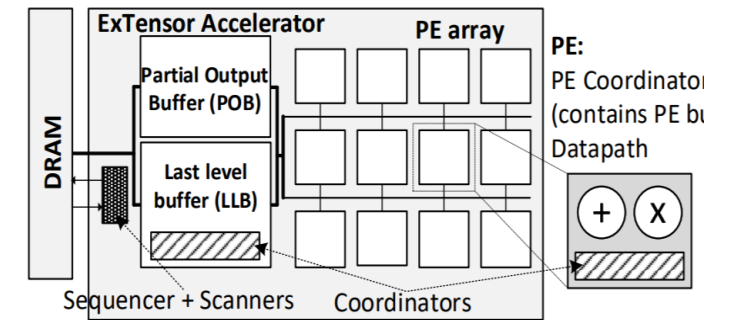
Large Design Space



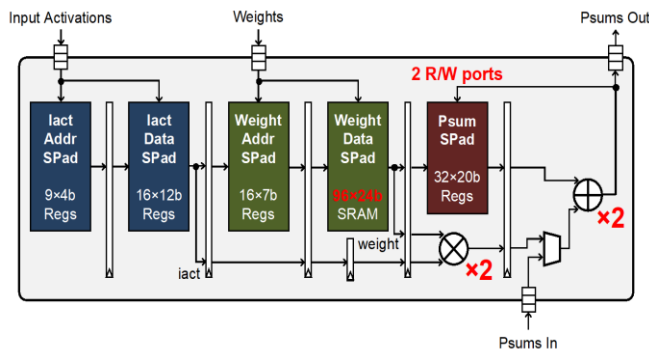
Eyeriss [JSSC2017]



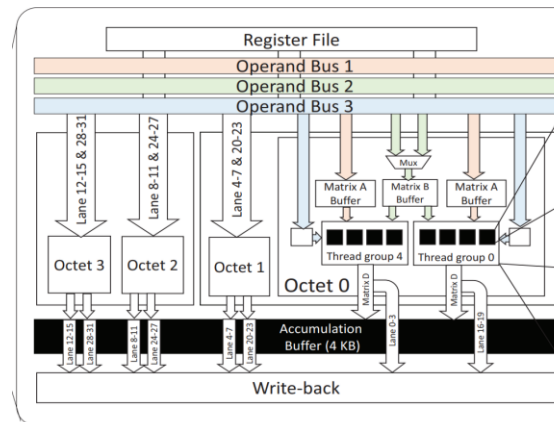
SCNN [ISCA2017]



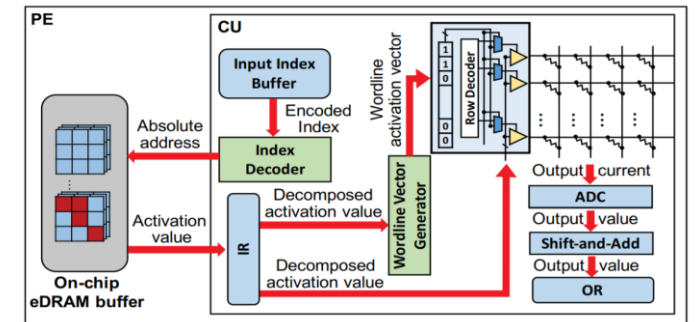
ExTensor [MICRO2019]



Eyeriss V2 [JETCAS2019]



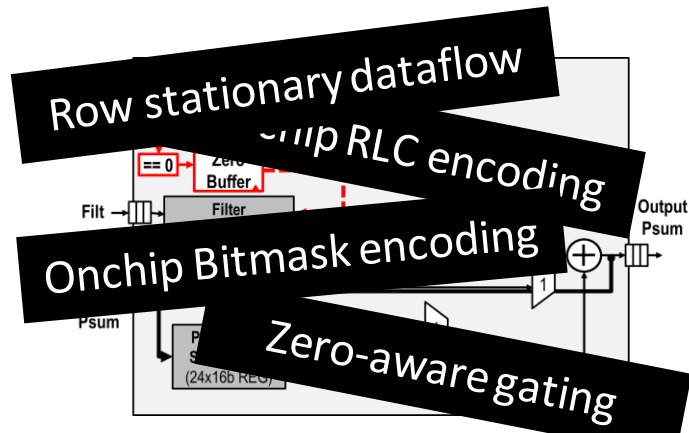
DSTC [ISCA2021]



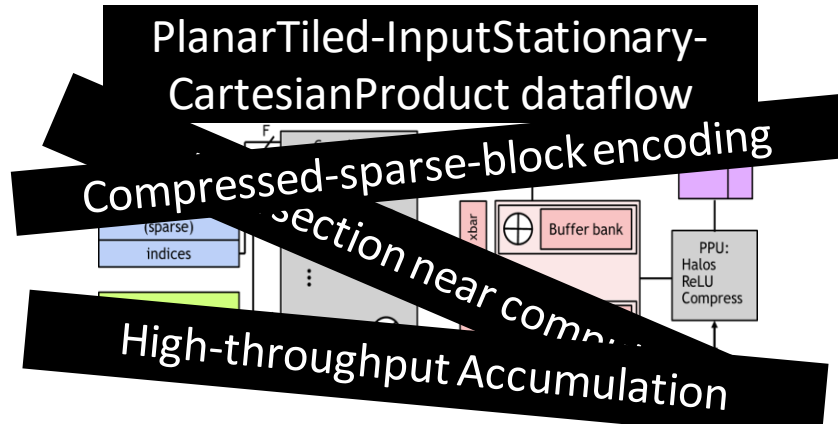
Sparse-ReRAM [ISCA2019]

Rely on Diverse Design-Specific Terminologies

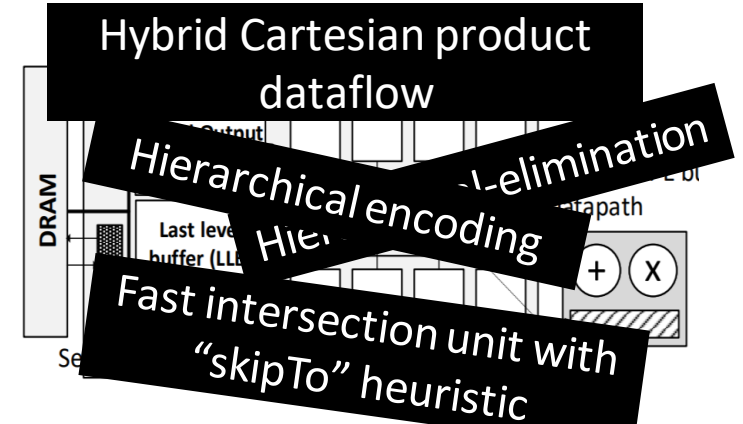
Large, Unstructured, and Confusing Design Space



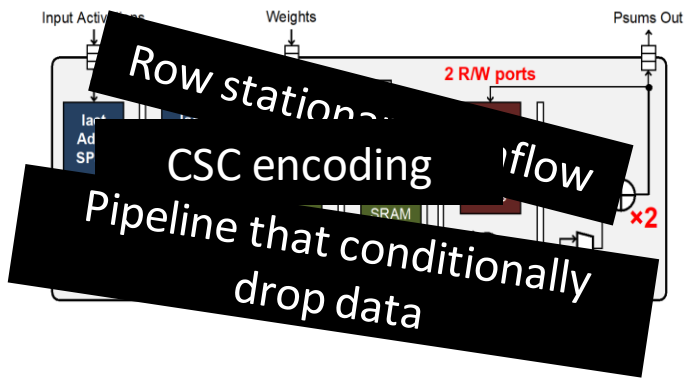
Eyeriss [JSSC2017]



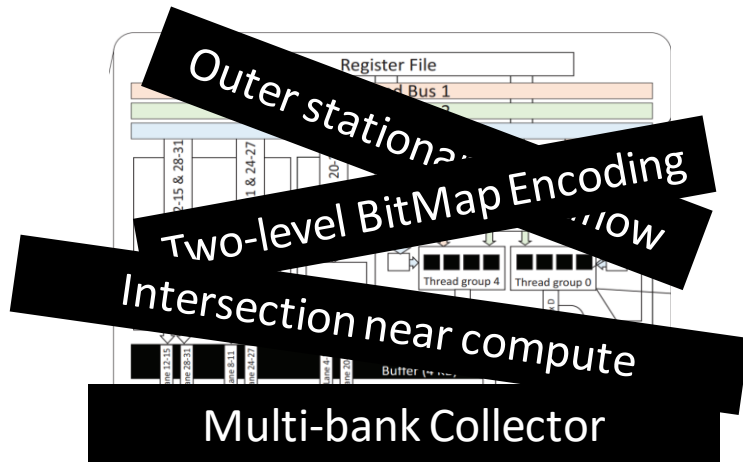
SCNN [ISCA2017]



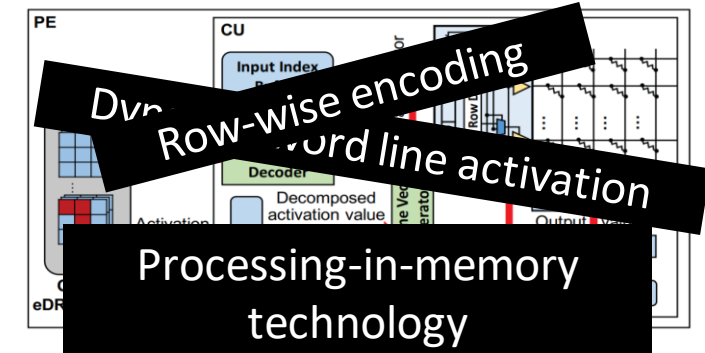
ExTensor [MICRO2019]



Eyeriss V2 [JETCAS2019]



DSTC [ISCA2021]




Sparse-ReRAM [ISCA2019]



**Important to
systematically understand and explore the design space**

**Requirements
A Modeling Framework**




Existing Modeling Frameworks are Insufficient

(Design-Specific) Cycle-Level Simulators

*SCNN[ISCA16], STONNE[CAL21],
MAGNET[ICCAD19], DNNBuilder[ICCAD18], etc.*

Slow



Inflexible



General Analytical Modeling Frameworks

*Timeloop[ISPASS19], MAESTRO[MICRO19],
Scale-Sim[ISPASS20], CoSA[ISCA21], etc.*

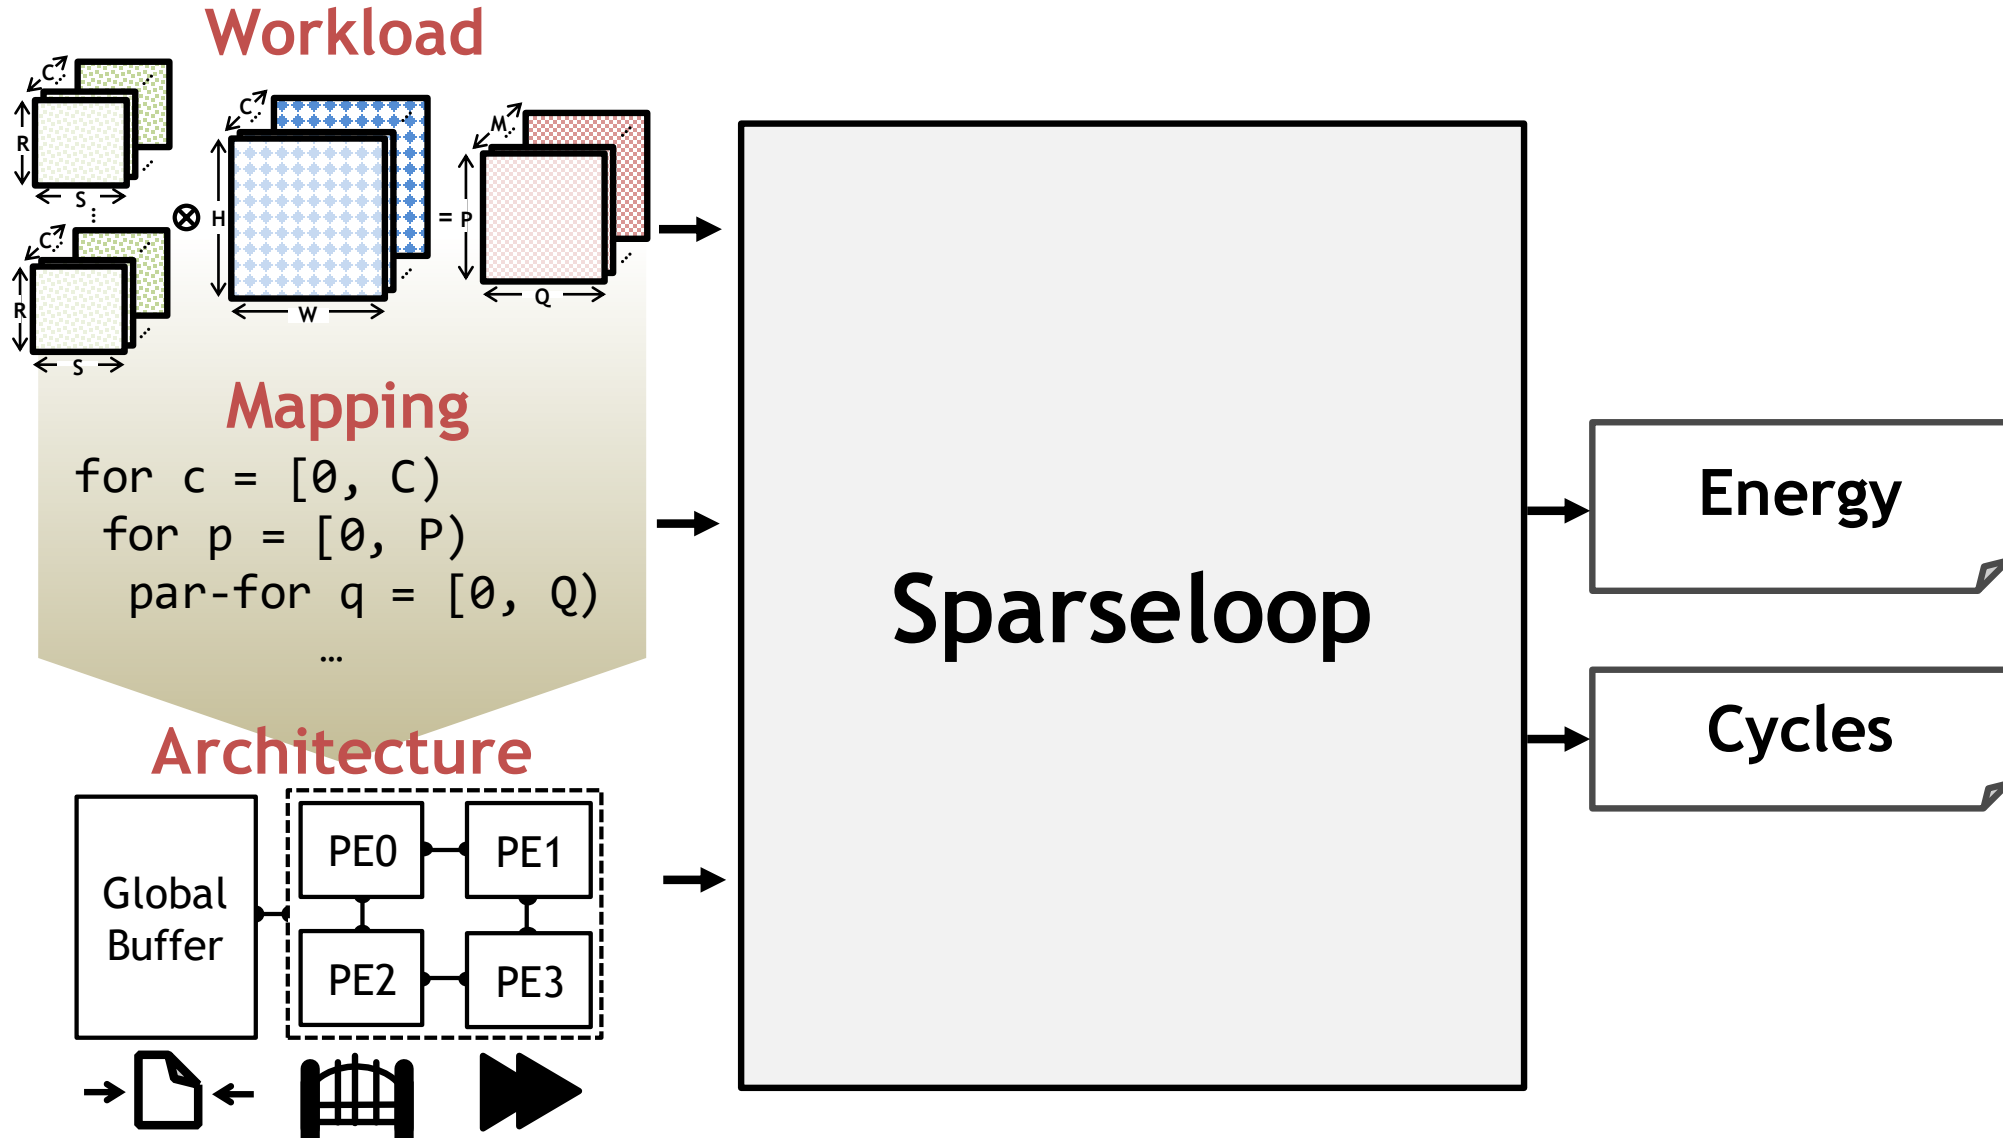
No Sparsity Support

0			0	0
0	0			
0				
0				0
0				

Solution

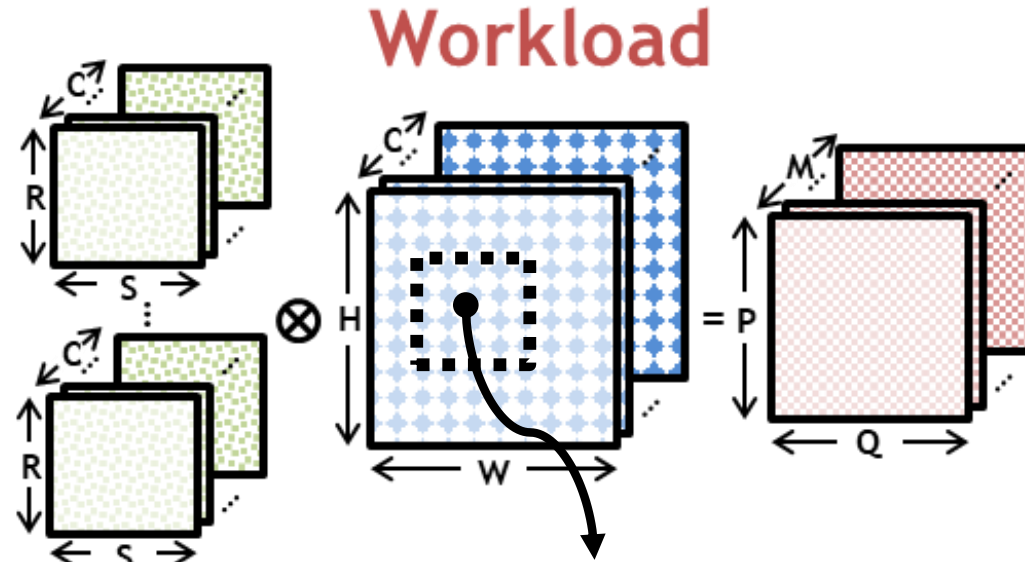
**Sparseloop: The First Analytical Modeling Framework
for Sparse Tensor Accelerators**

Sparseloop High-Level Framework



Challenge: Slow Workload Characterization

Accelerator Performance is Data-dependent

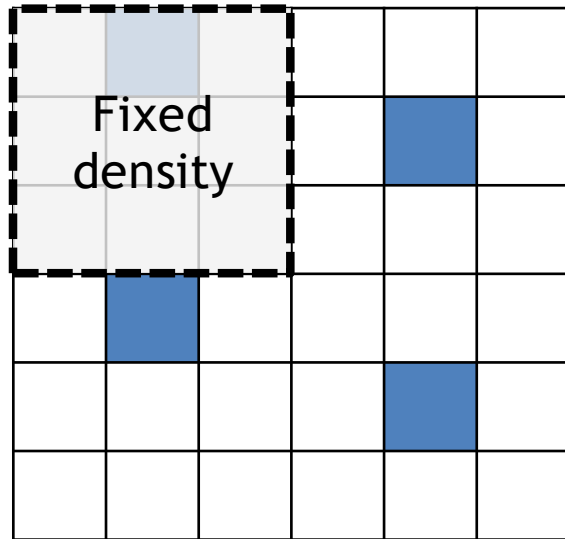


Nonzero values locations
in various subtensors

Traversing the exact values can be very slow

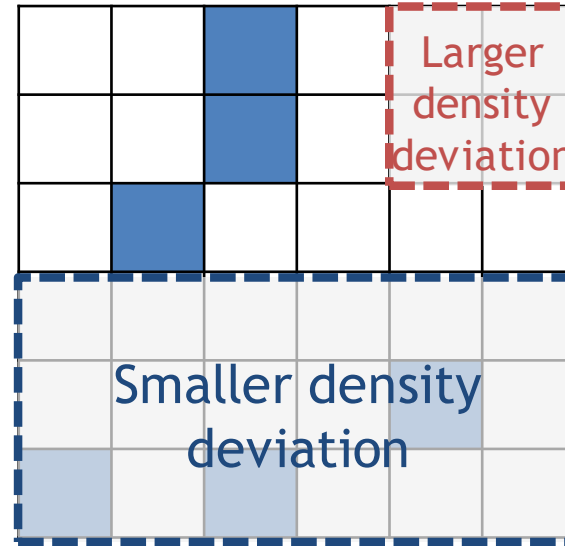
Sparseloop Solution: Statistical Characterization

Fixed-Structured



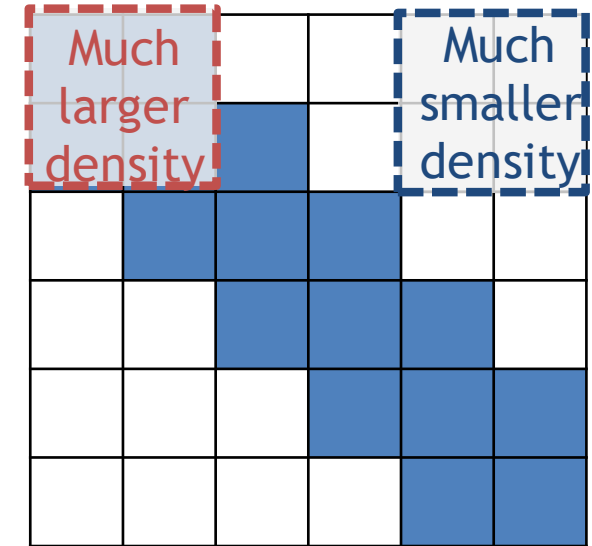
Structured Pruned
DNNs

Uniform Random



Unstructured Pruned
DNNs

Banded Distribution




Scientific
Simulations

Statistical Modeling Ensures Both Speed and Accuracy

Challenge: Unstructured Architecture Description

High-Level Opportunities



Zero Values

Can be Compressed Away

$$x \times 0 = 0$$

$$x + 0 = x$$

Ineffectual Operations

Can be Eliminated

Can be Exploited Differently at Different Architecture Levels

Sparseloop Solution: Sparse Acceleration Features



Format

Choice of tensor representations



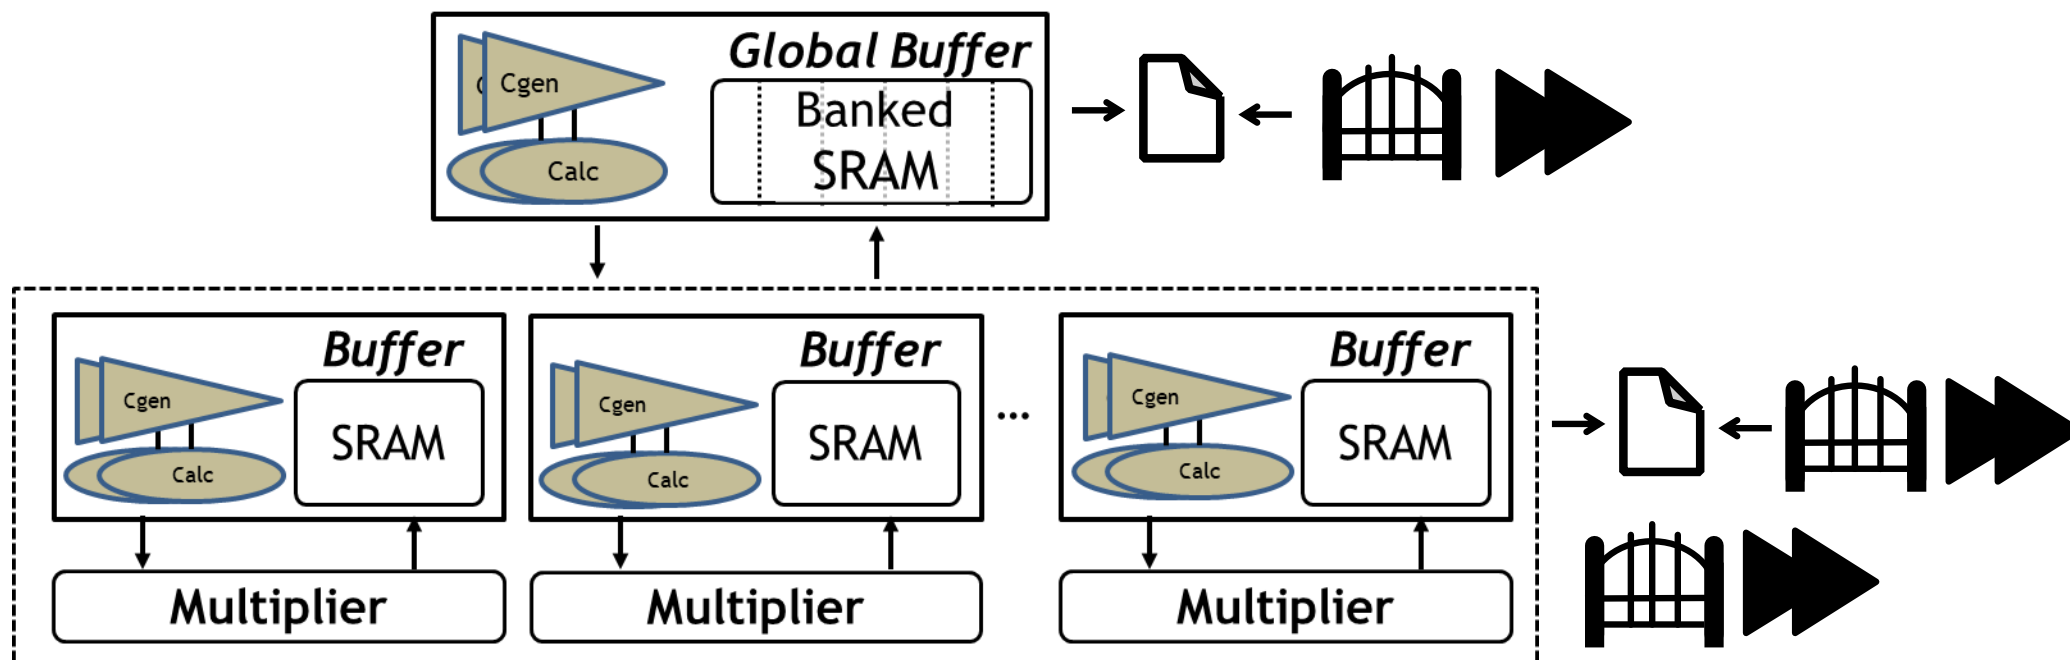
Gating

Explicitly let the hardware staying idle



Skipping

Explicitly fast forward to next effectual operation

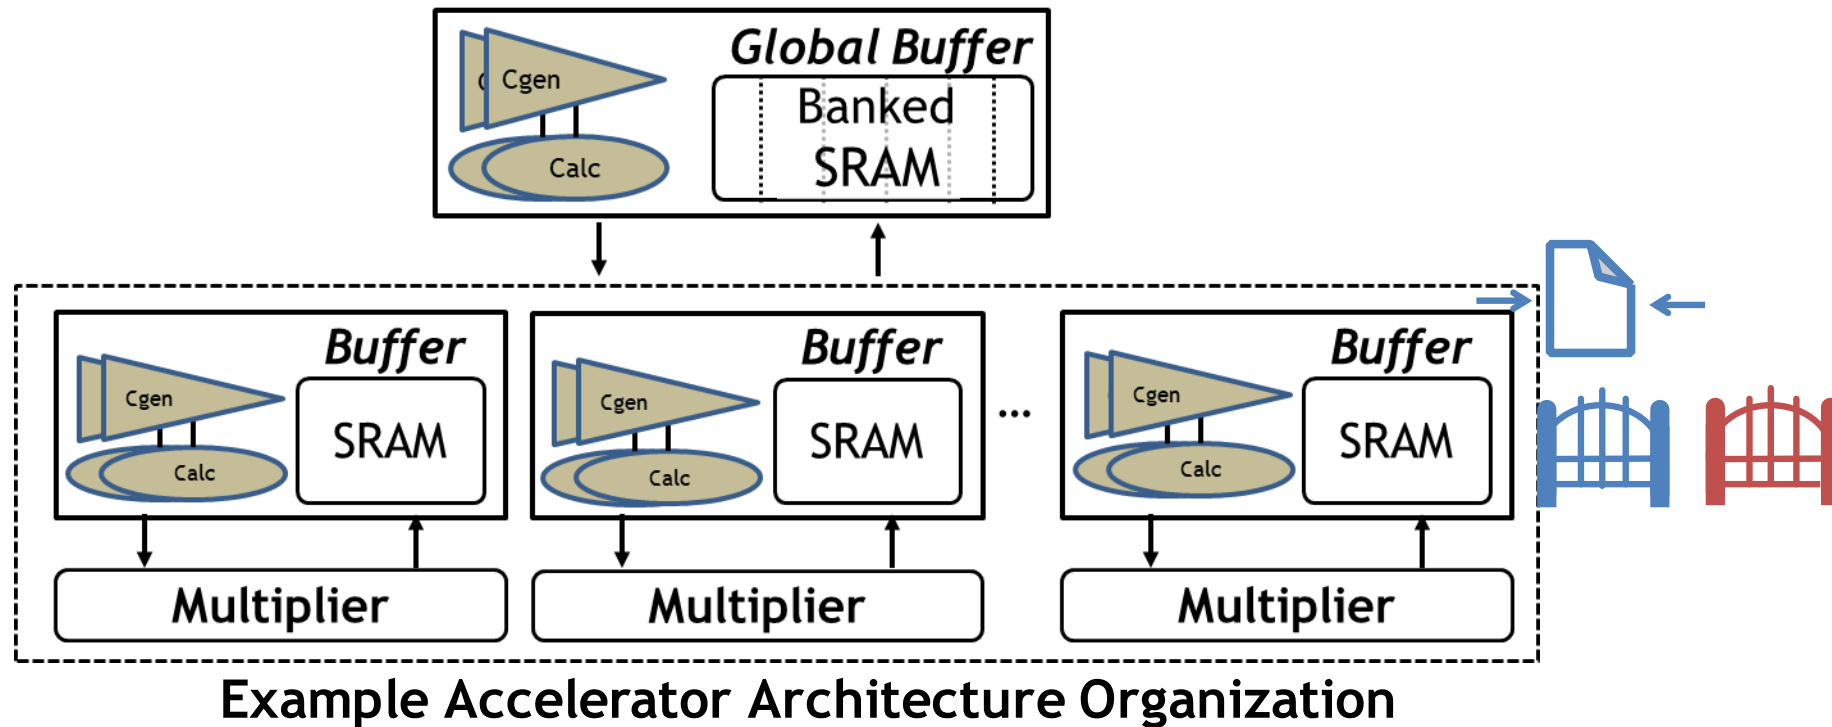


Example Accelerator Architecture Organization

Systematic Descriptions with Various SAF Combinations

Eyeriss-Style
[ISCA2016]

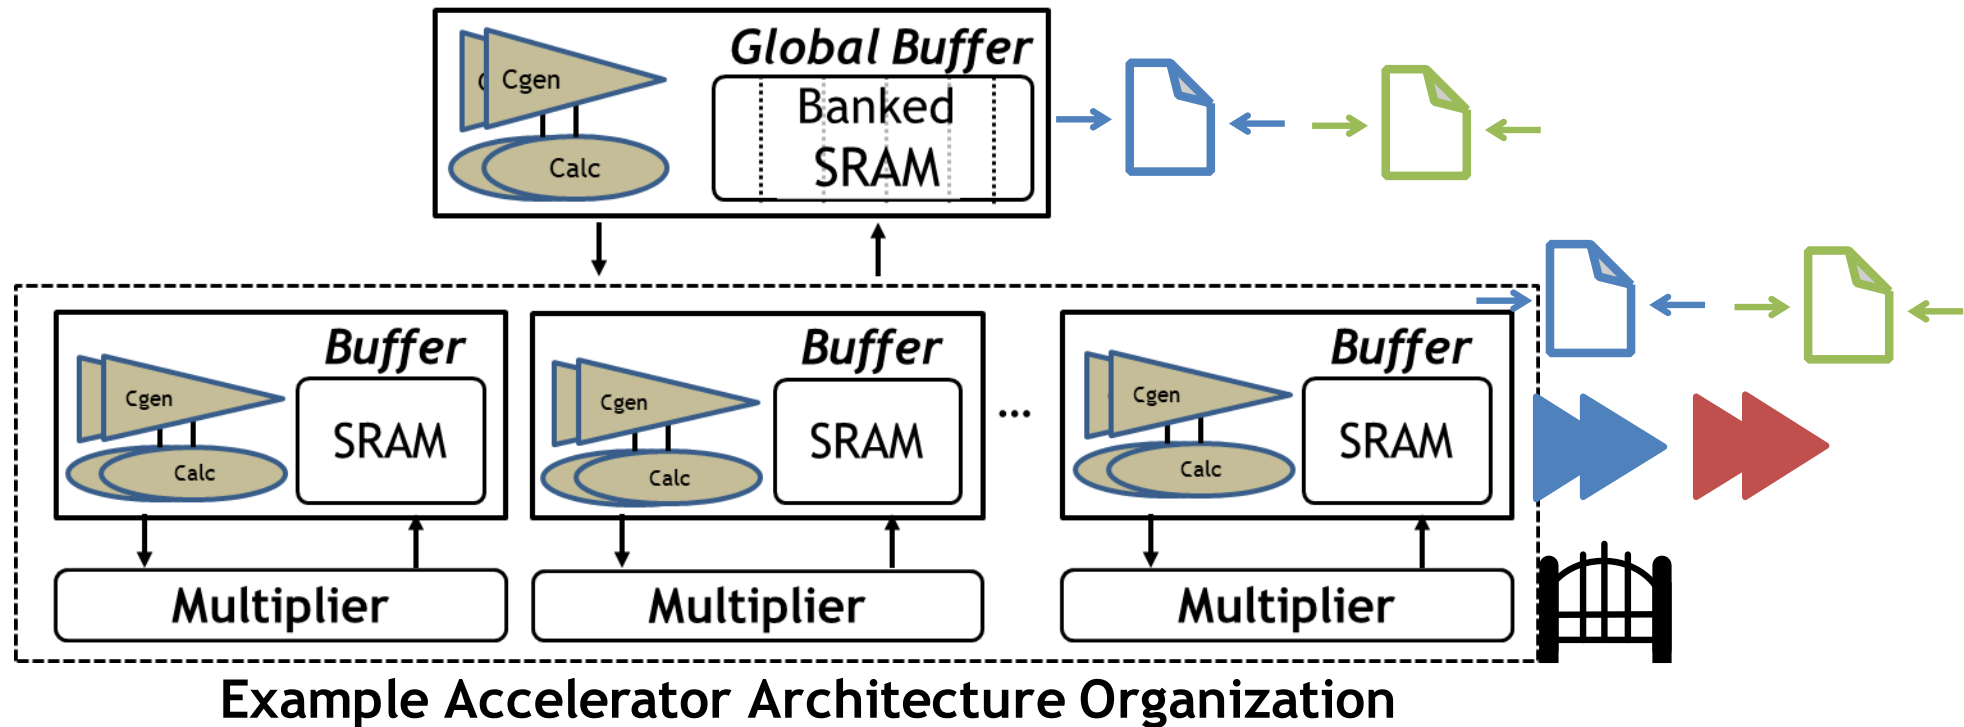

*Different color coding
represents features applied
to different operand*



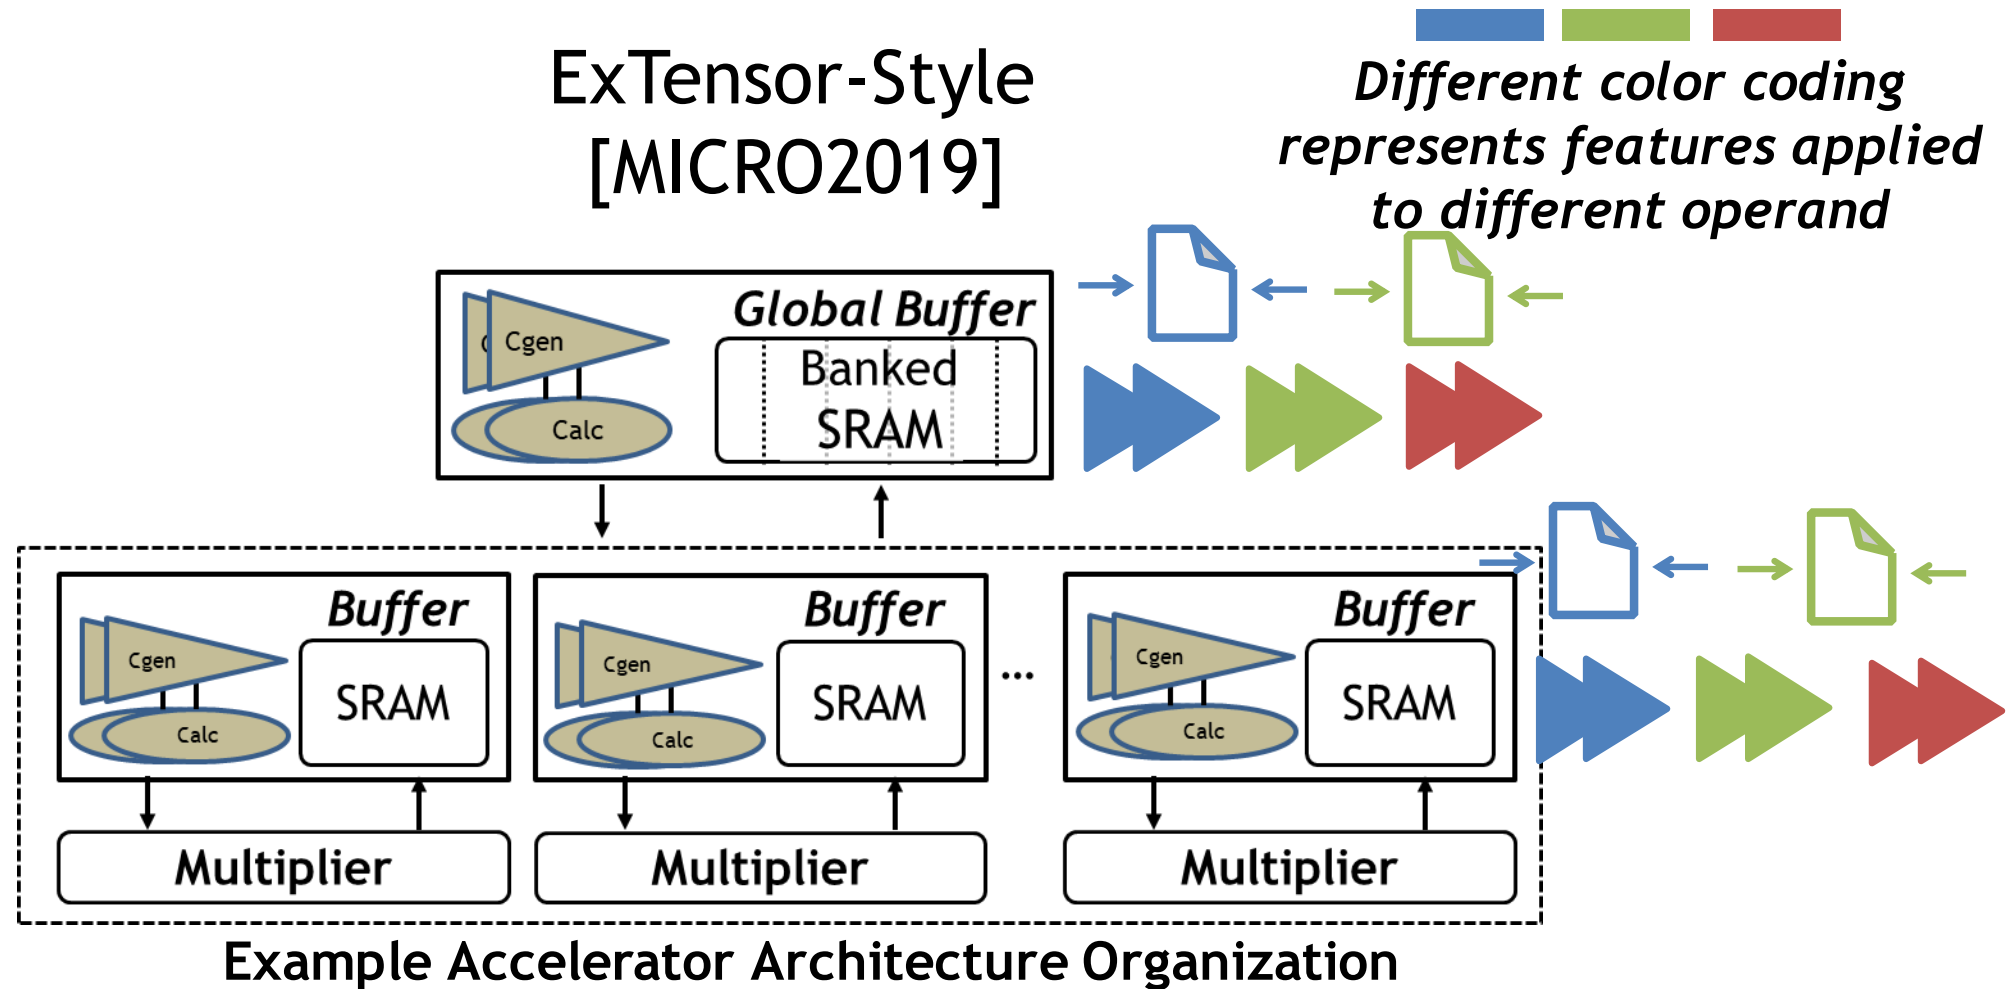
Systematic Descriptions with Various SAF Combinations

Eyeriss V2-Style
[JETCAS2019]


*Different color coding
represents features applied
to different operand*

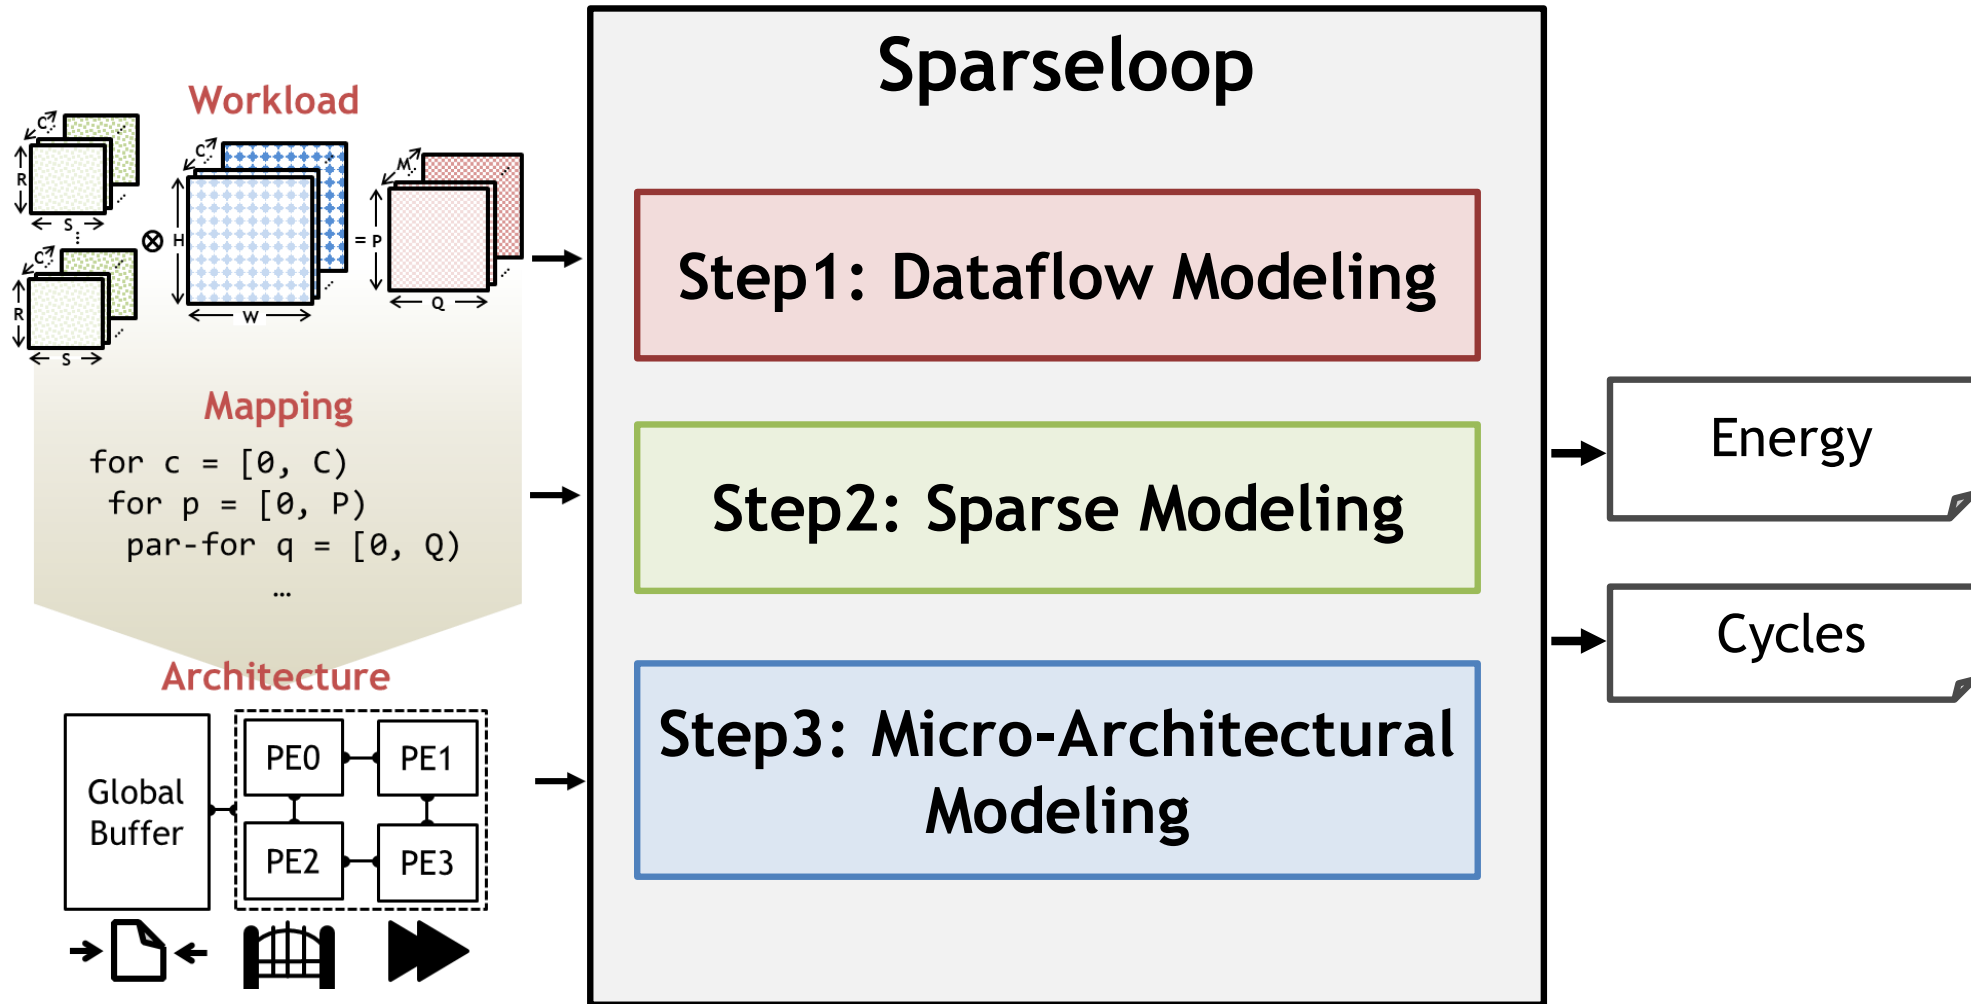


Systematic Descriptions with Various SAF Combinations



Challenge: Complex Interactions Lead Slow Modeling Speed

Sparseloop Solution: Decoupled Modeling



Keep Modeling Complexity Tractable

Modeling Speed and Accuracy

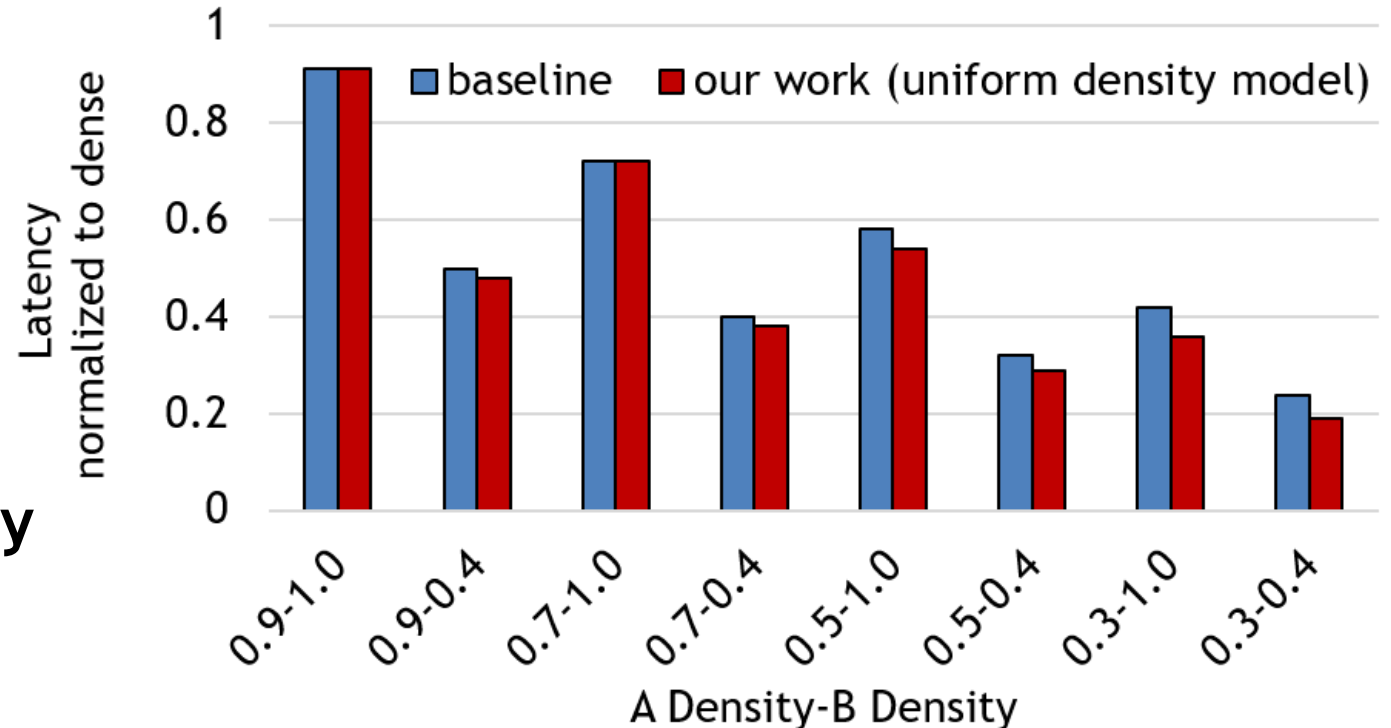
- Speed

- >2000x faster compared to cycle-level simulations
- Months -> Hours

- Accuracy

- Validated on well-known DNN accelerators
 - Maintains relative trends
 - Achieves 0.1% - 8% error in cycle counts and energy consumption




Example DSTC [ISCA21] Validation



More Details in Paper!

- How to build the next-generation sparse tensor core accelerator?
 - *short answer: explore support for different sparsity ratios*
- What happens when we use a sparse DNN accelerator to run much sparser HPC workloads? Or vice versa?
 - *short answer: sparse acceleration features become ineffective for inappropriate workloads*
- ...

Summary

- Sparseloop is a fast, accurate, and flexible analytical modeling framework that enables tensor accelerator design space exploration
 - **Fast:** achieves >2000x speedup compared to cycle-level simulations
 - **Accurate:** maintains relative trend and achieves 0.1% - 8% error on cycles counts and energy consumption
 - **Flexible:** helps designers understand the critical design trade-offs
- Resources
 - Artifact   
 - Tutorial at: <http://sparseloop.mit.edu/>

