

Duality Cache

for Data Parallel Acceleration

Daichi Fujiki

Scott Malke

Reetuparna Das

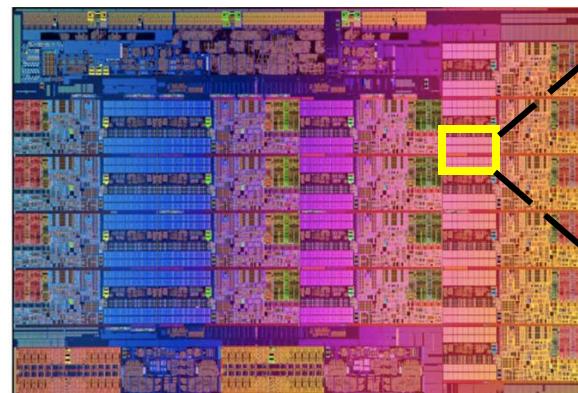


Duality = Storage + Compute

Why compute in-cache ?

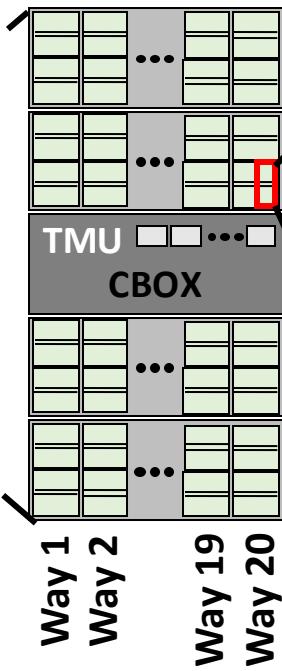
Transforming caches into massively parallel vector ALUs

18-core Xeon processor
45 MB LLC



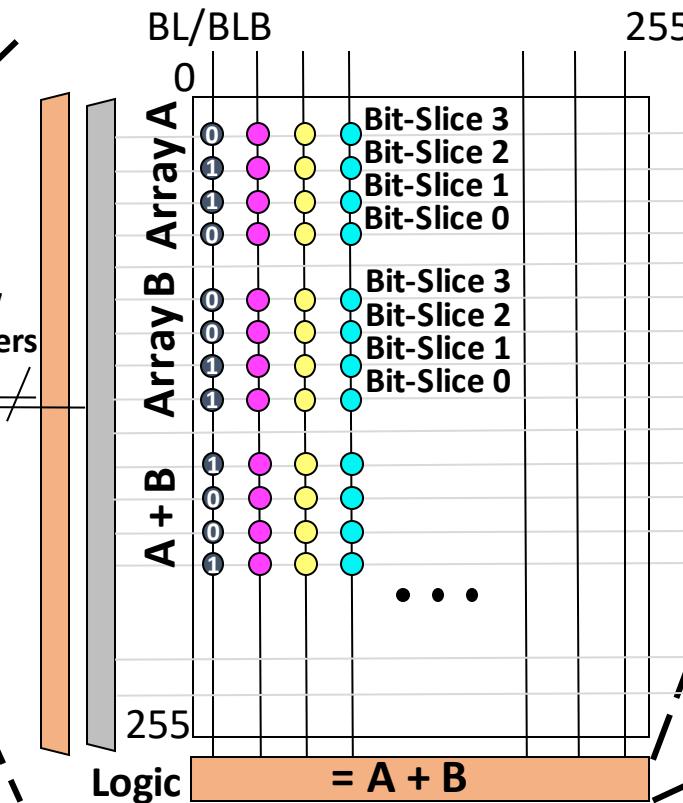
18 LLC slices

2.5MB LLC slice



360 ways

8kB SRAM array



Transforming caches into massively parallel vector ALUs

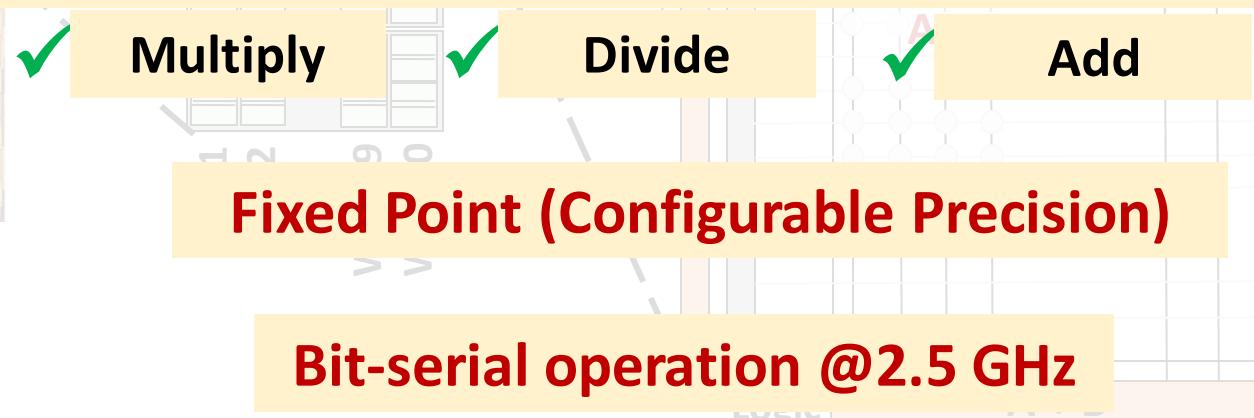
18-core Xeon processor
45 MB LLC



Bitline ALU



Passive Last Level Cache transformed into ~ 1 million bit-serial active ALUs



18 LLC slices

32kB
data
bank

8kB array

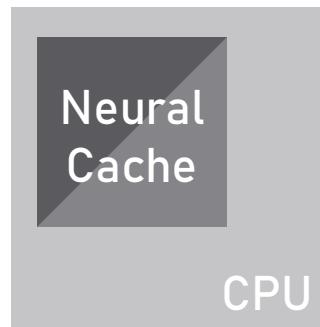
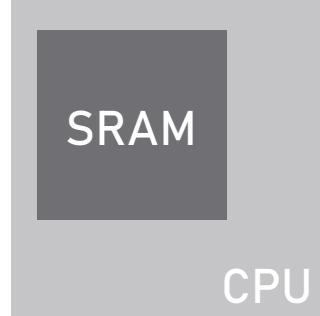
360 ways

5760 arrays

45 MB LLC →

1,474,560 ALUs

Neural Cache [ISCA '18]



35 MB LLC
Cache Mode



Accelerator Mode

Neural Network

1 million *bit-line* ALUs

Digital Bit-Serial Operations @ 2.5GHz

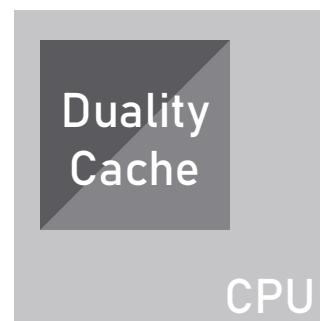
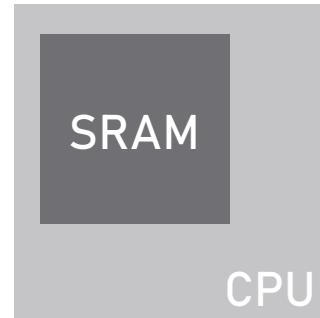
Supported Operations

- ✓ Integer Operations

Programming Model

- ✓ Manually Mapped CNN Kernels

Duality Cache



35 MB LLC

Cache Mode

1 million *bit-line* ALUs

</> Accelerator Mode

Data Parallel Applications

Digital Bit-Serial Operations @ 2.5GHz

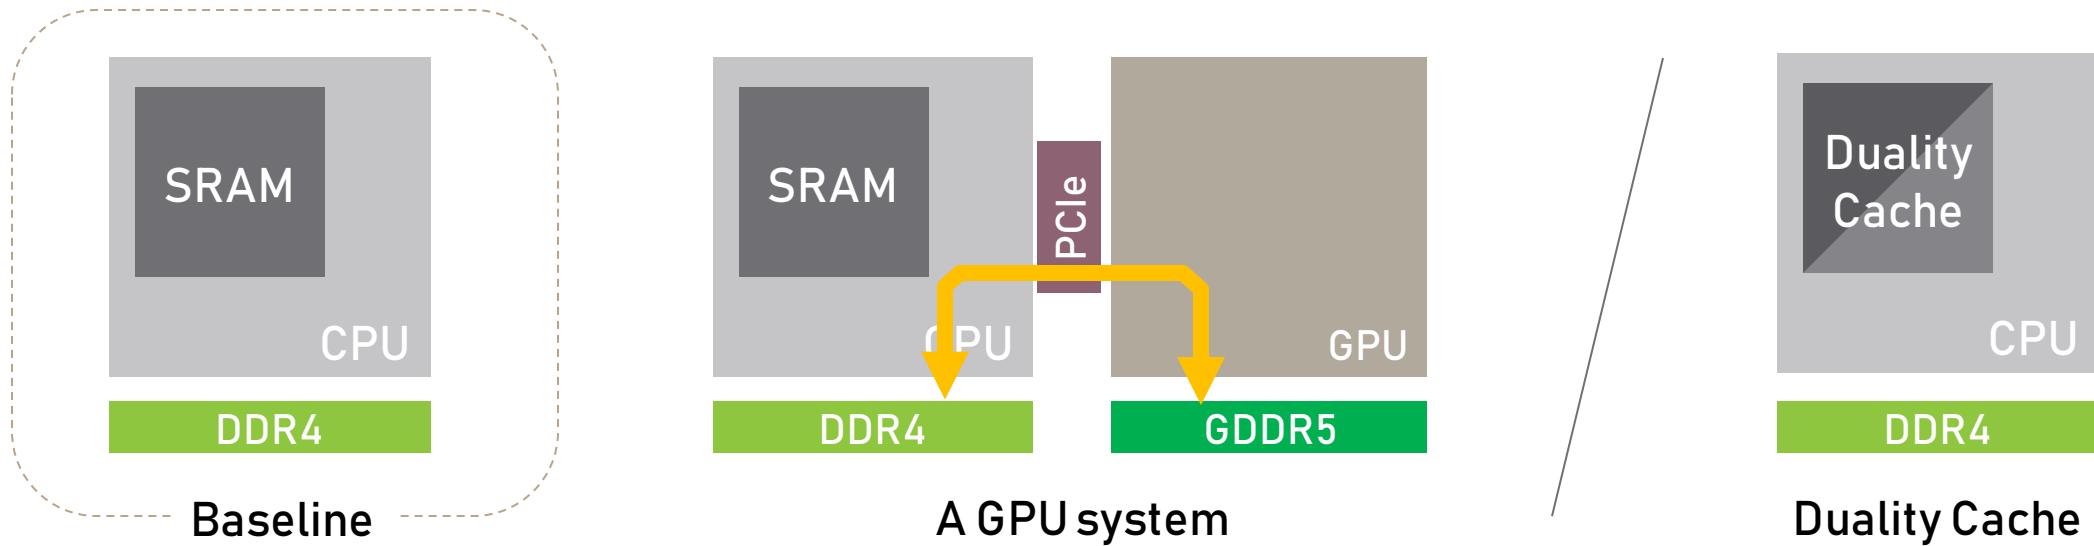
Supported Operations

- ✓ Integer Operations
- ✓ Floating Point Operations
- ✓ Transcendental Functions (sin, cos, log etc.)

Programming Model

- ✓ SIMD & VLIW execution
- ✓ CUDA / OpenACC programs

Duality Cache Benefits



1. Reduced data movement
 2. Cost
 3. On chip memory capacity
- ✓ No memcpy. Efficient serial/parallel interleaving.
 + 471 mm² +250W (GPU) vs. ✓ +30 mm² +6W = ~3.5% (DC)
 ✓ ~10x thread capacity. Flexible cache allocation.

Outline

Background

- ▶ Integer / Logical Operations
- Floating Point Operations
- Transcendental Functions

Execution Model

Programming Model

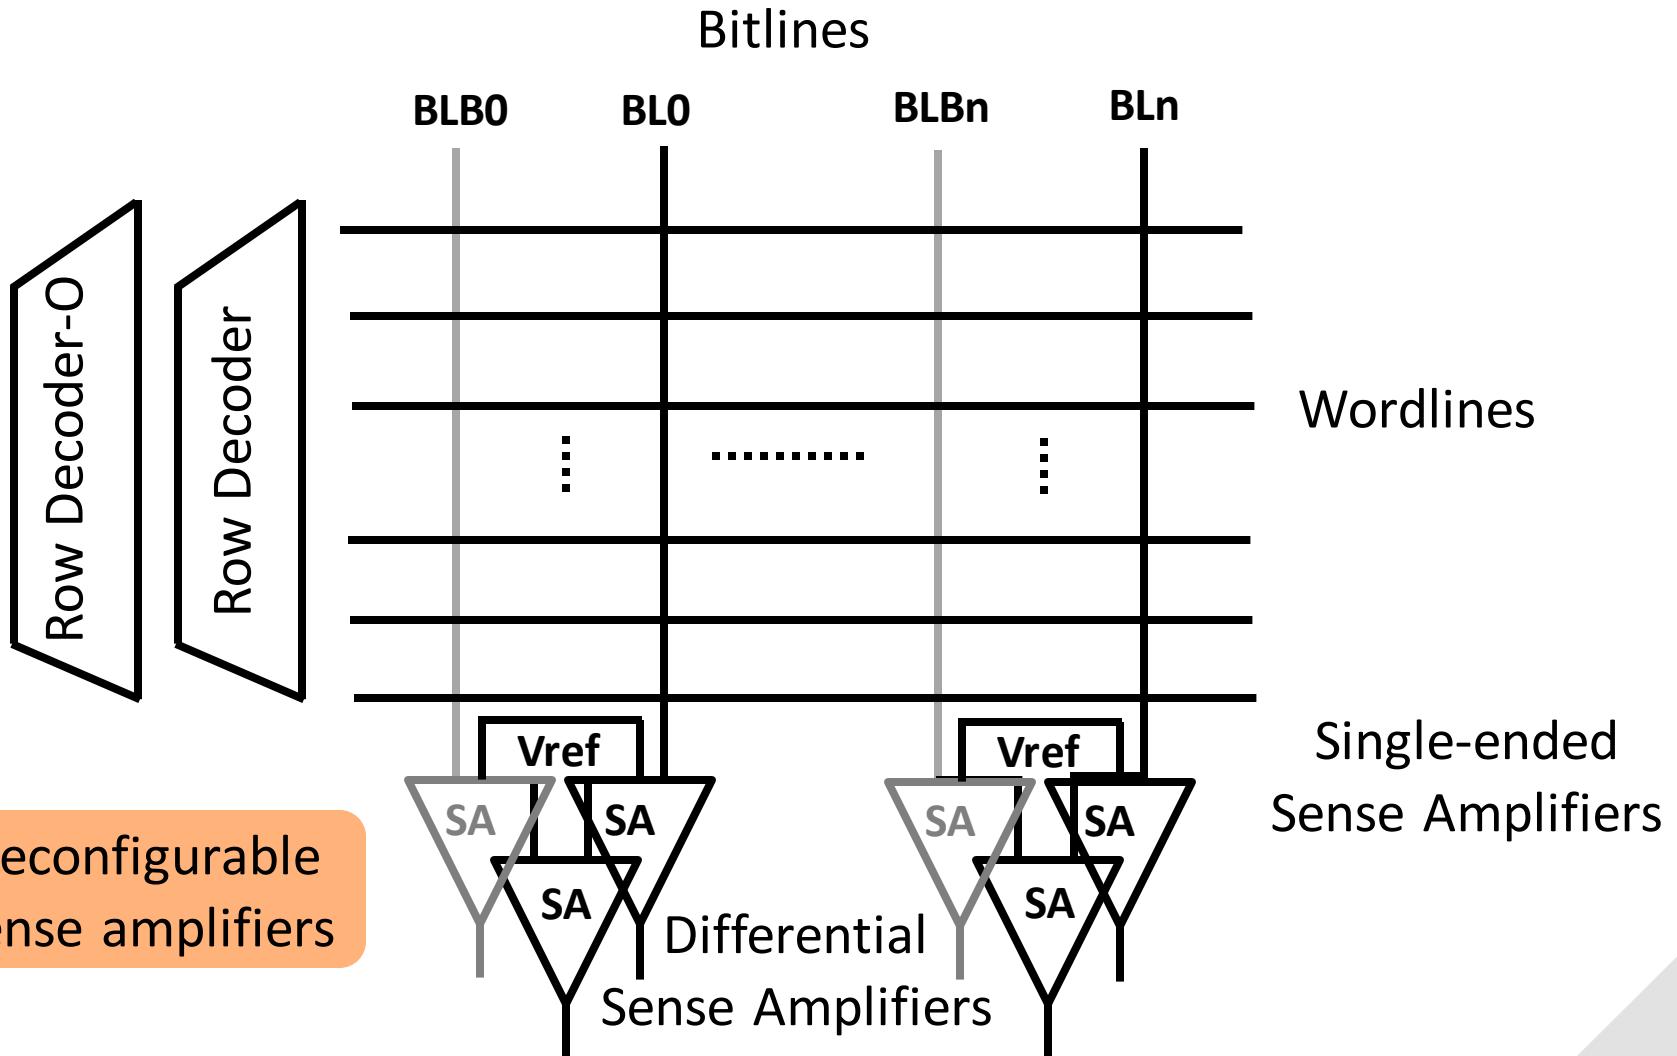
Compiler

Methodology / Results

Logical Operations In-SRAM

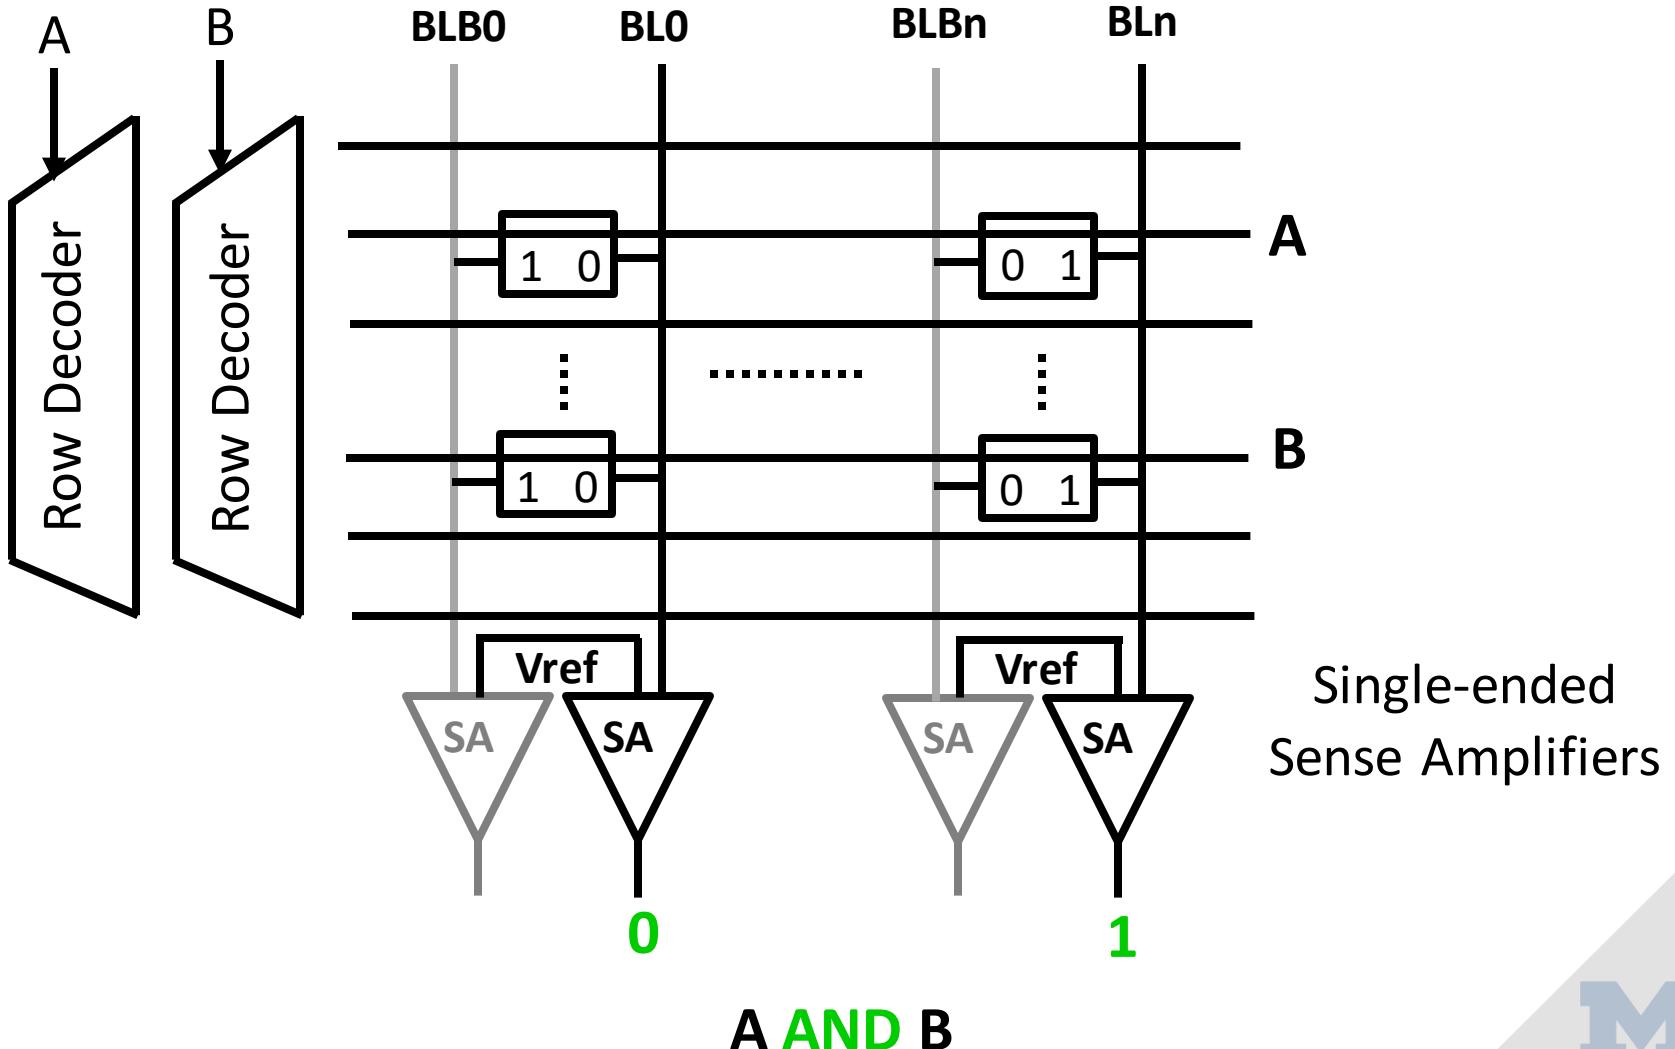
Changes

Additional
row decoder

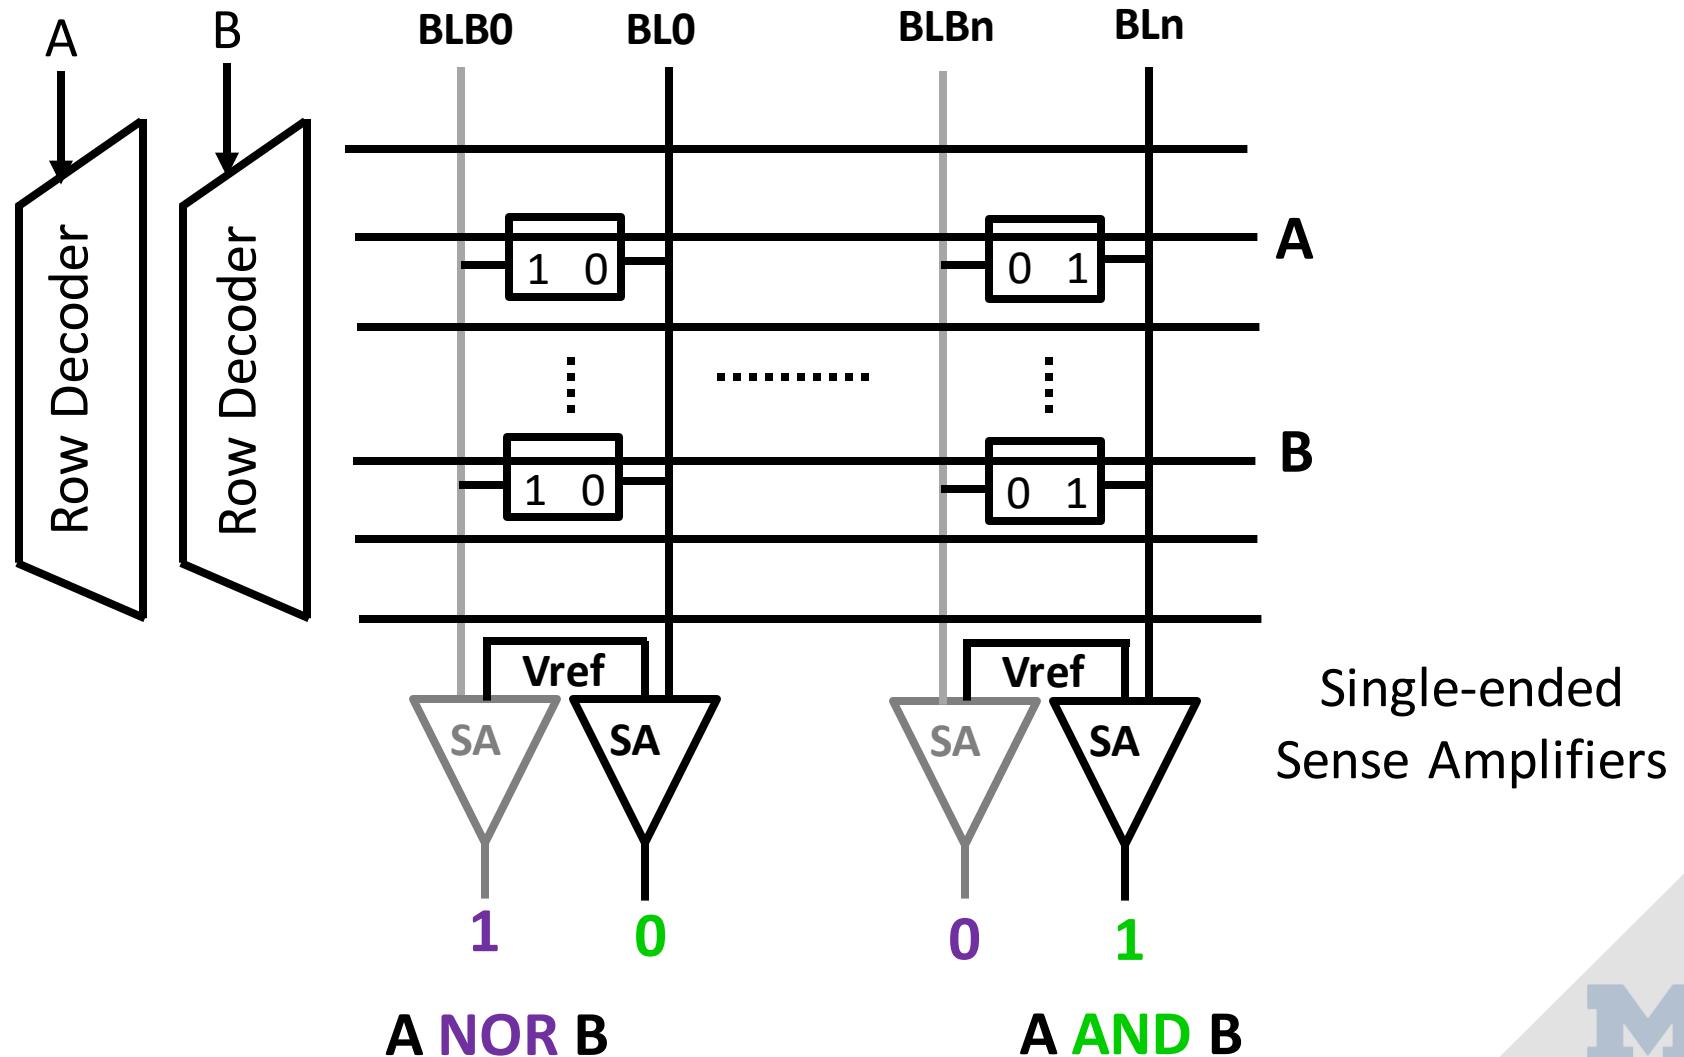


Logical Operations In-SRAM

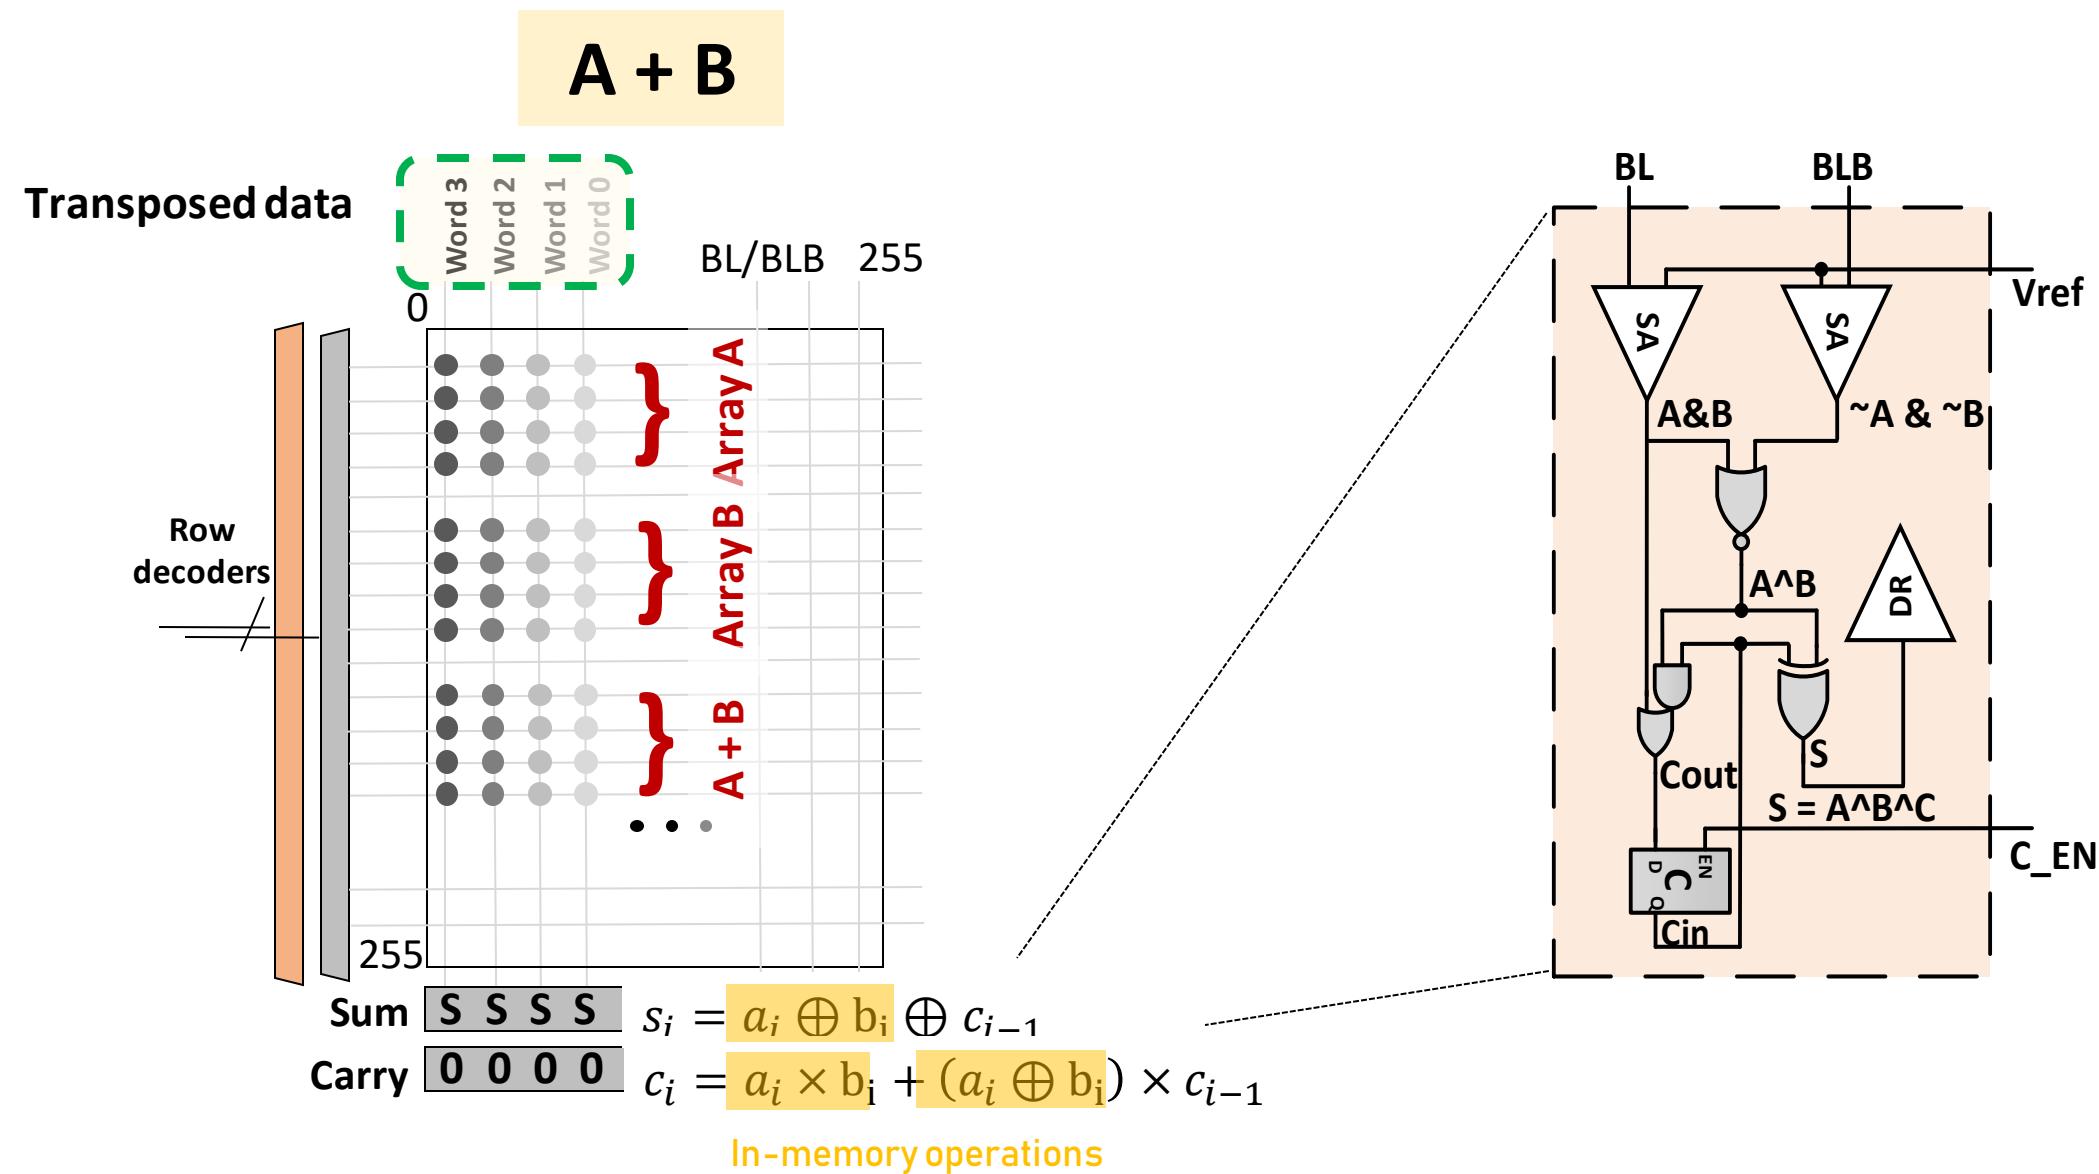
A AND B



Logical Operations In-SRAM

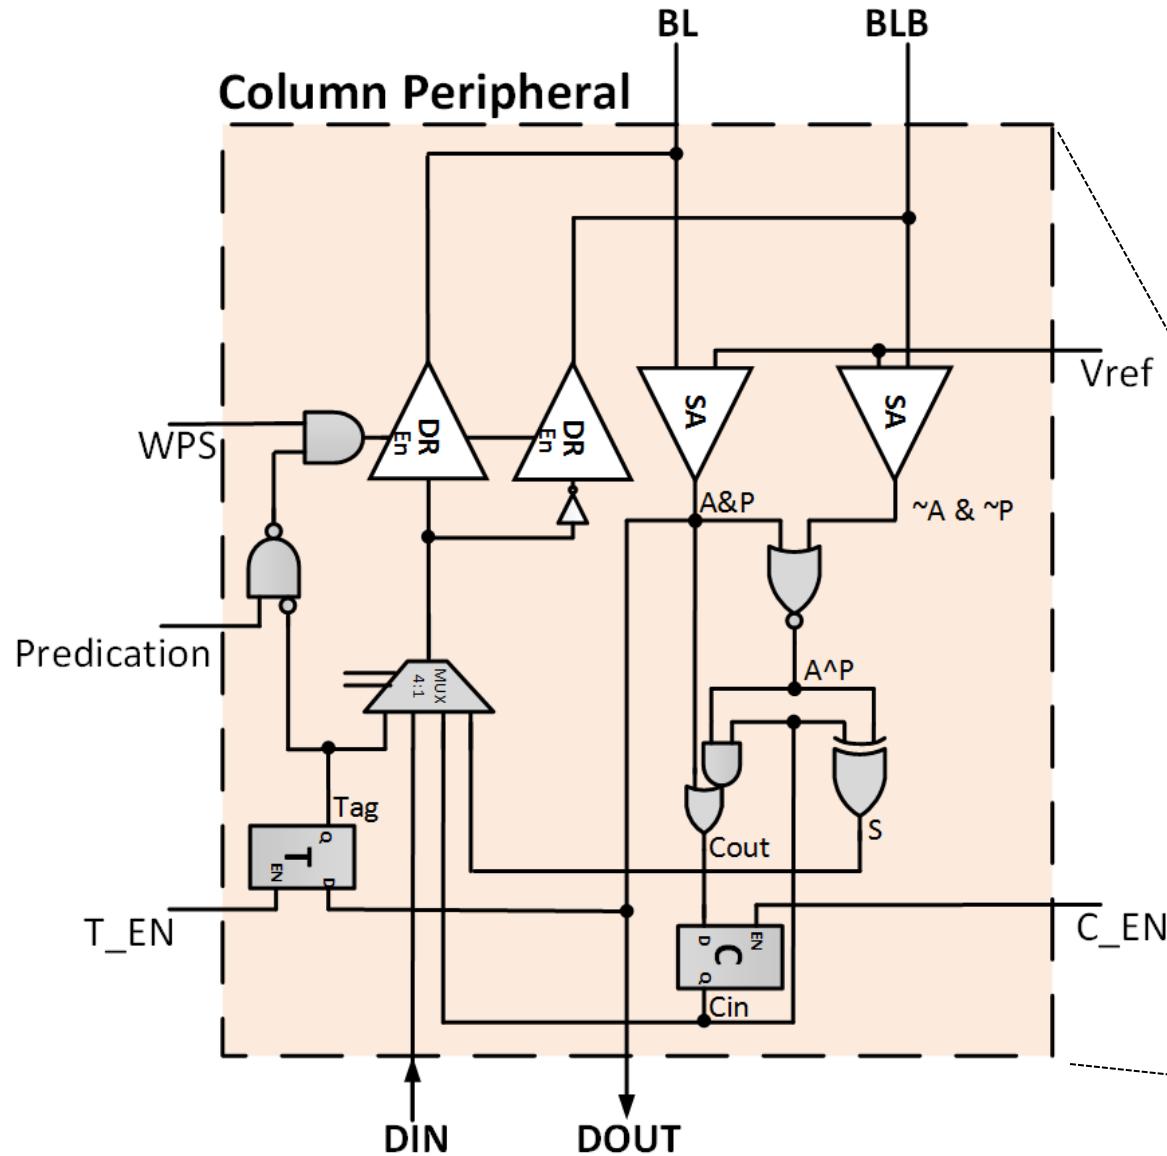


Bit-Serial Integer Operations

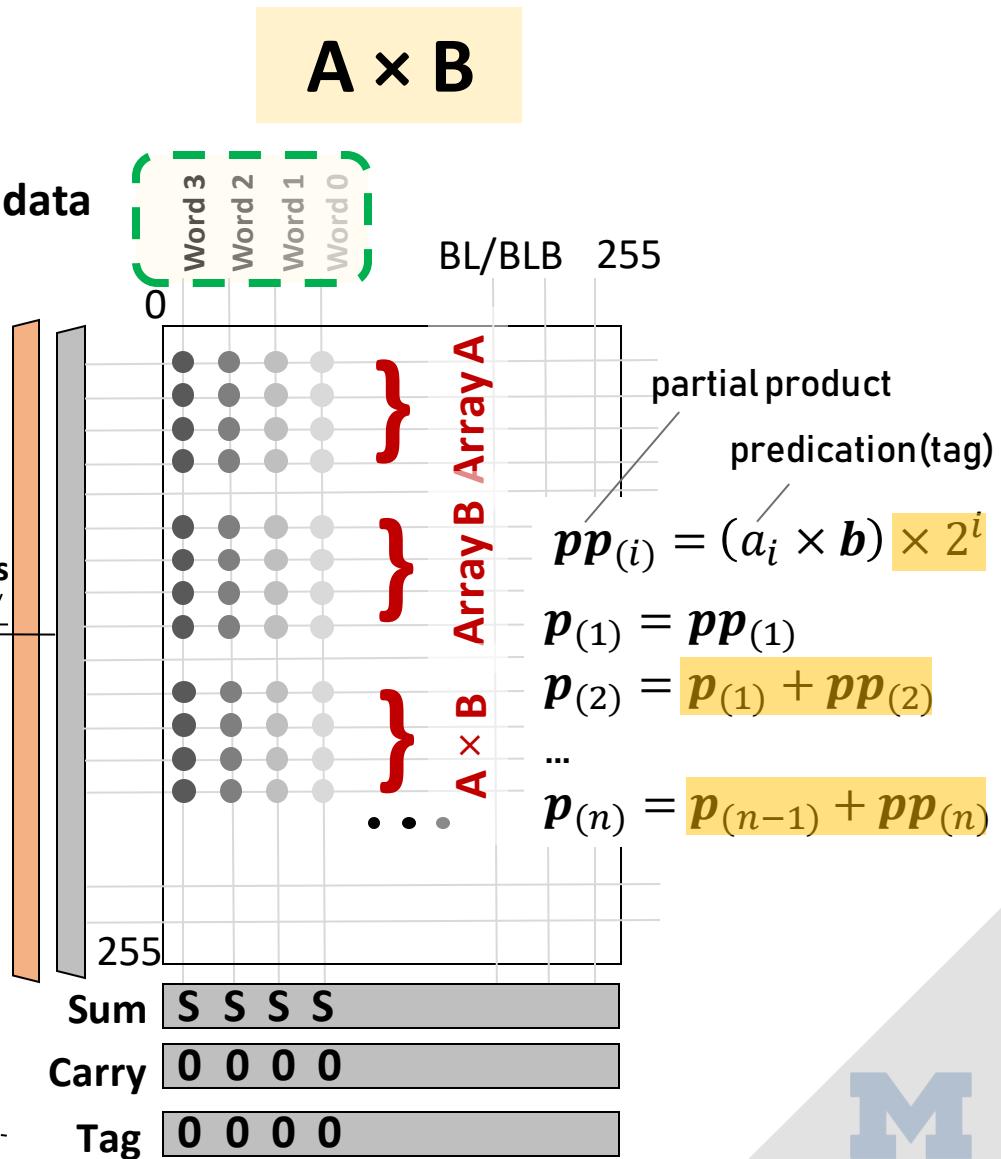


Bit-Serial Integer Operations

$a_i := i^{\text{th}}$ bit of bit vector a
 $a_{(i)} := \text{bit vector } a \text{ after } i^{\text{th}}$ iteration

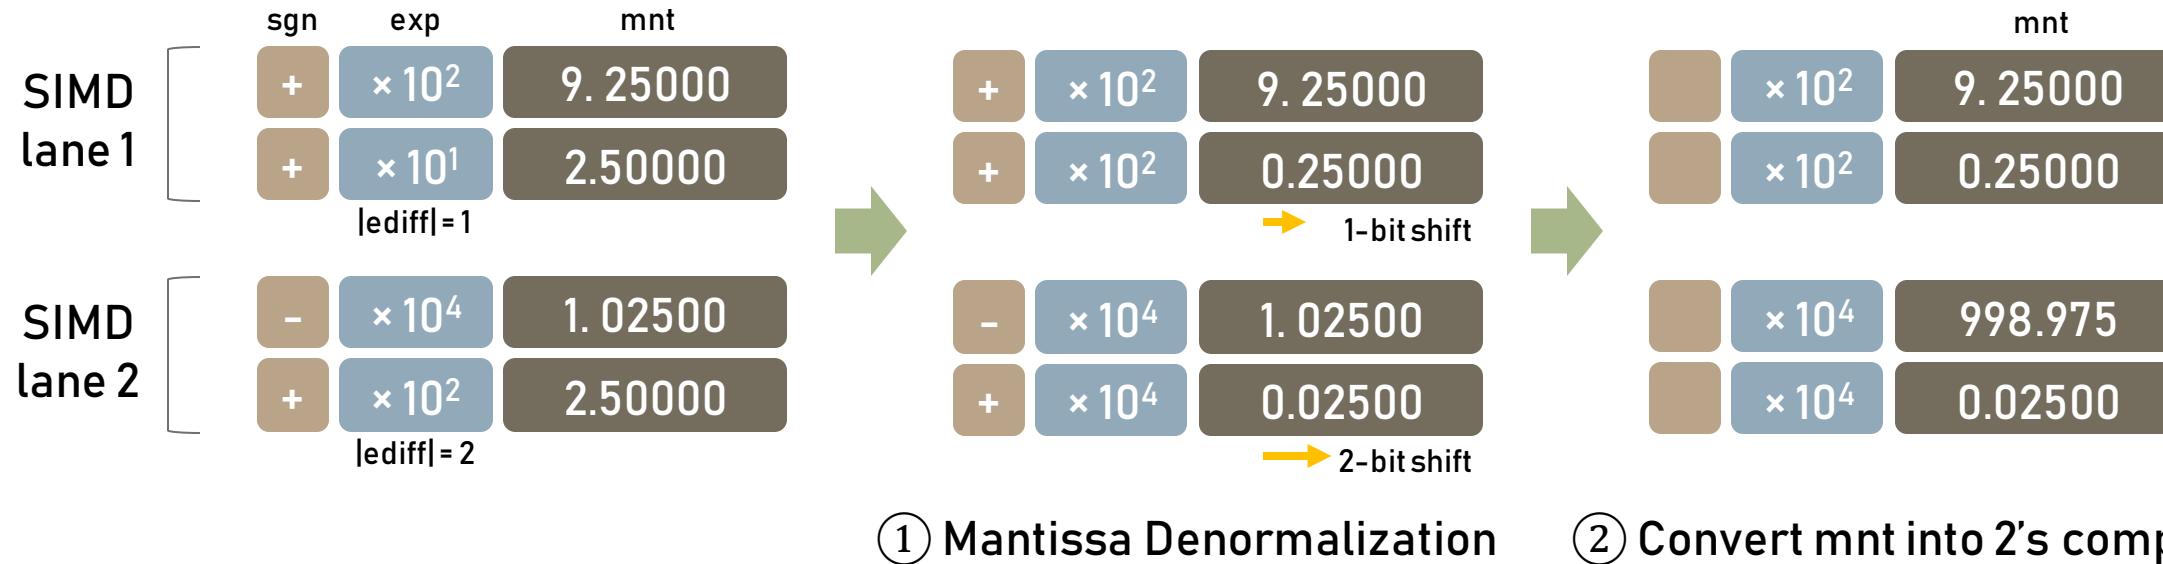


Transposed data



Bit-Serial Floating Point Operations

A + B



① Mantissa Denormalization

② Convert mnt into 2's complement format



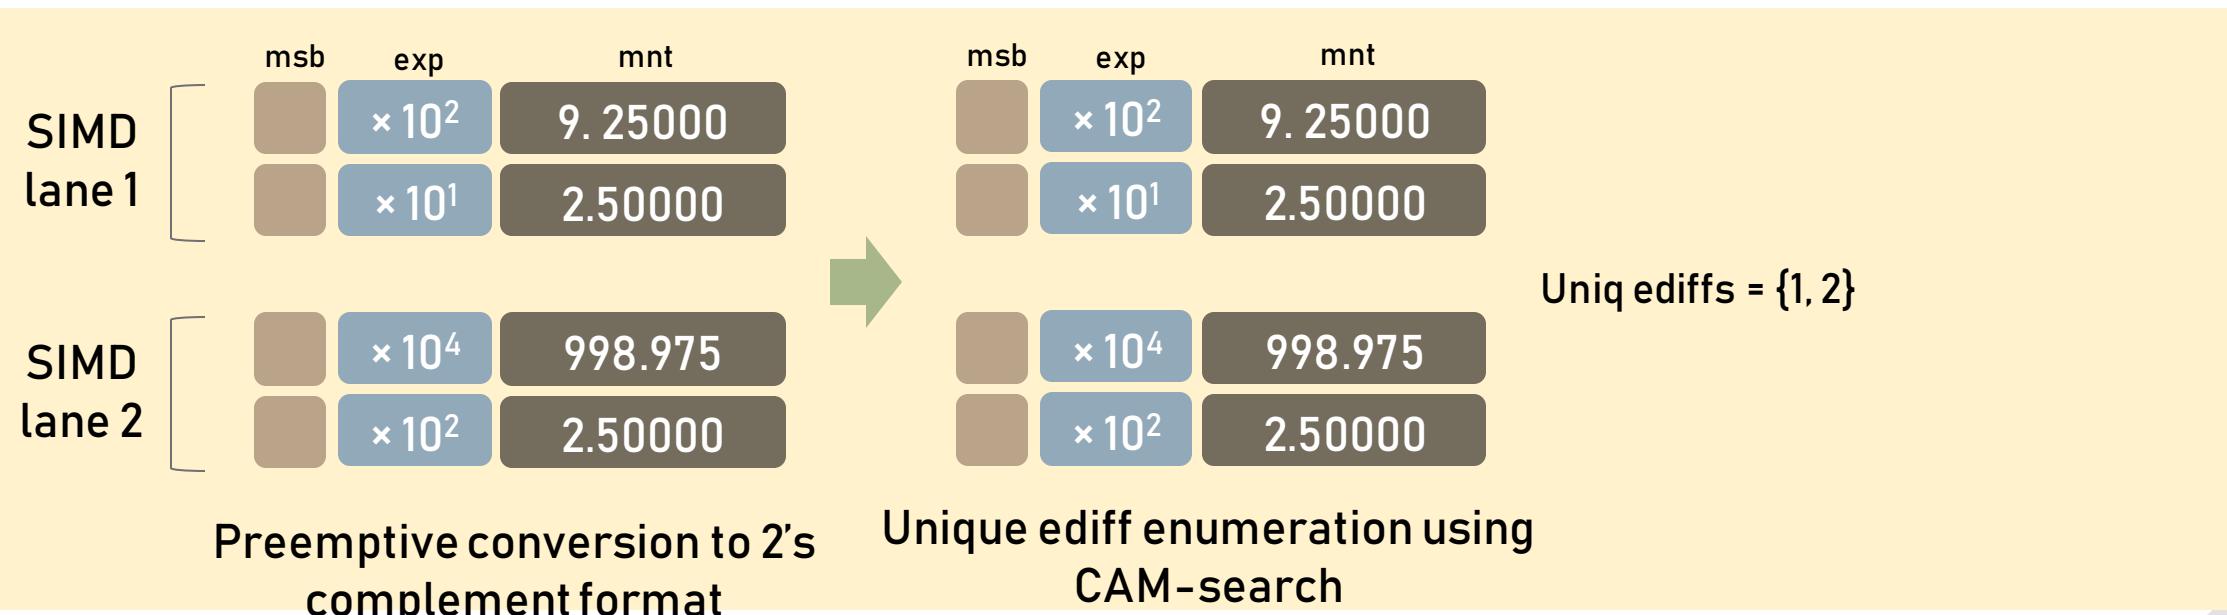
Up to $mnt_{bit} * \#lanes$ cycles for shift ops



Conversion cost (signed \leftrightarrow 2's complement)

Bit-Serial Floating Point Operations

A + B



Bit-Serial Floating Point Operations

A + B

	sgn	exp	mnt				
SIMD lane 1	[+]	$\times 10^2$	9.25000	[+]	$\times 10^1$	2.50000	$ ediff = 1$
SIMD lane 2	-	$\times 10^4$	1.02500	[+]	$\times 10^2$	2.50000	$ ediff = 2$
SIMD lane 3	-	$\times 10^3$	1.02500	[+]	$\times 10^1$	2.50000	$ ediff = 2$
SIMD lane 4	-	$\times 10^3$	1.02500	[+]	$\times 10^1$	2.50000	$ ediff = 2$
	:						:
SIMD lane 256	-	$\times 10^3$	1.00500	[+]	$\times 10^2$	2.50100	$ ediff = 1$

8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift
8 cycle ediff read + 23 cycle mnt shift

7,936 cycles (total)



Up to $mnt_{bit} * \#lanes$ cycles for shift ops



Unique ediff enumeration using CAM-search

Uniq ediffs = {1, 2}

$ ediff = 1$	23 cycle mnt shift
$ ediff = 2$	23 cycle mnt shift
	+ CAM-search cycle

Bit-Serial Floating Point Operations

A + B

SIMD lane 1

	exp	mnt
$\times 10^2$	9.25000	
$\times 10^2$	0.25000	

 $+ =$

	exp	mnt
$\times 10^2$	10.0000	

SIMD lane 2

	exp	mnt
$\times 10^4$	998.975	
$\times 10^4$	0.02500	

 $+ =$

	exp	mnt
$\times 10^4$	999.000	

② Convert mnt into complement format

③ Perform addition

sgn	exp	mnt
+	$\times 10^2$	10.0000



sgn	exp	mnt
-	$\times 10^4$	1.00000

④ Convert back to sign expression

sgn	exp	mnt
+	$\times 10^3$	1.0000



sgn	exp	mnt
-	$\times 10^4$	1.00000

⑤ Normalization

Bit-Serial Floating Point Operations

A + B

SIMD lane 1

	exp	mnt
$\times 10^2$	9.25000	
$\times 10^2$	0.25000	

 $+ =$

	exp	mnt
$\times 10^2$	10.0000	

SIMD lane 2

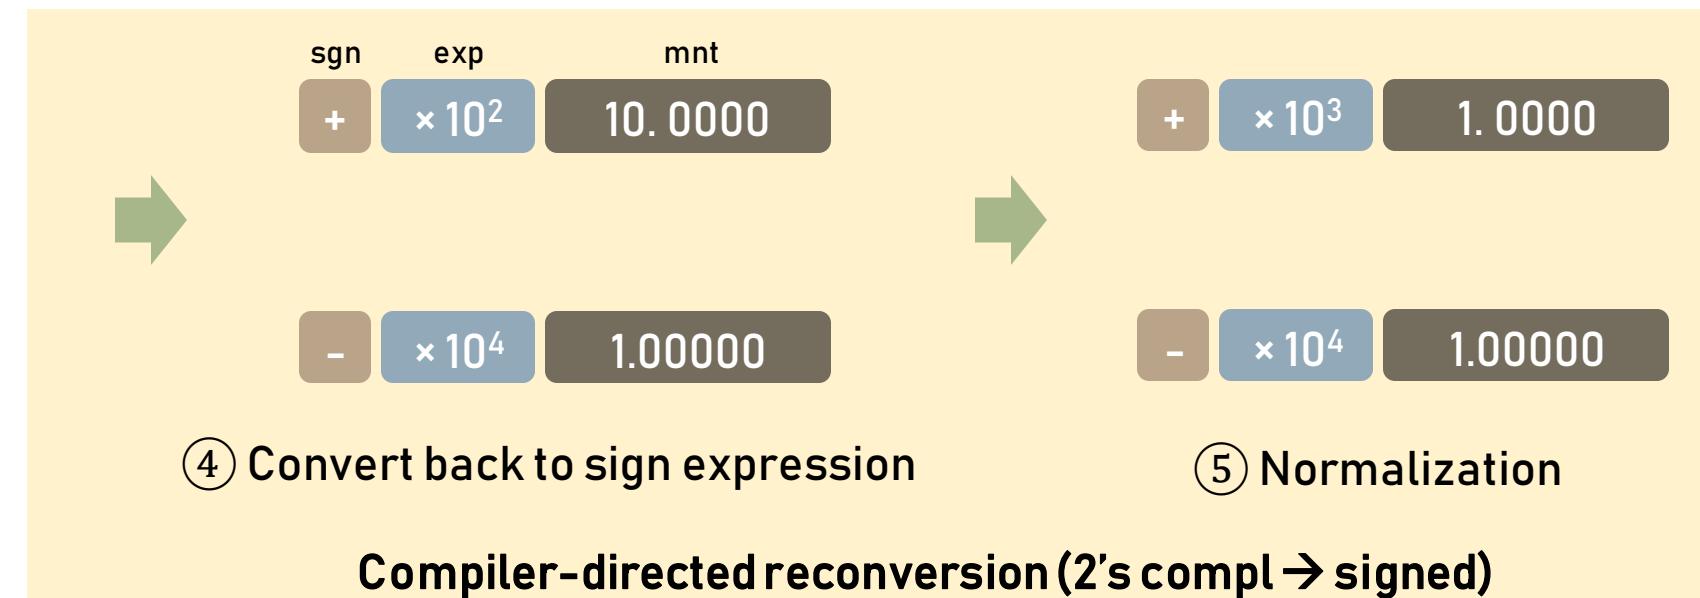
	exp	mnt
$\times 10^4$	998.975	
$\times 10^4$	0.02500	

 $+ =$

	exp	mnt
$\times 10^4$	999.000	

② Convert mnt into complement format

③ Perform addition



Bit-Serial Floating Point Operations

A + B

SIMD lane 1

	exp	mnt
$\times 10^2$	9.25000	
$\times 10^2$	0.25000	

 $+ =$

	exp	mnt
$\times 10^2$	10.0000	

SIMD lane 2

	exp	mnt
$\times 10^4$	998.975	
$\times 10^4$	0.02500	

 $+ =$

	exp	mnt
$\times 10^4$	999.000	

② Convert mnt into complement format

③ Perform addition

SIMD lane 1

	exp	mnt
$\times 10^2$	9.25000	
$\times 10^1$	2.50000	

 $+ =$

	exp	mnt
$\times 10^2$	10.0000	



	exp	mnt
$\times 10^3$	1.0000	

SIMD lane 2

	exp	mnt
$\times 10^4$	998.975	
$\times 10^2$	2.50000	

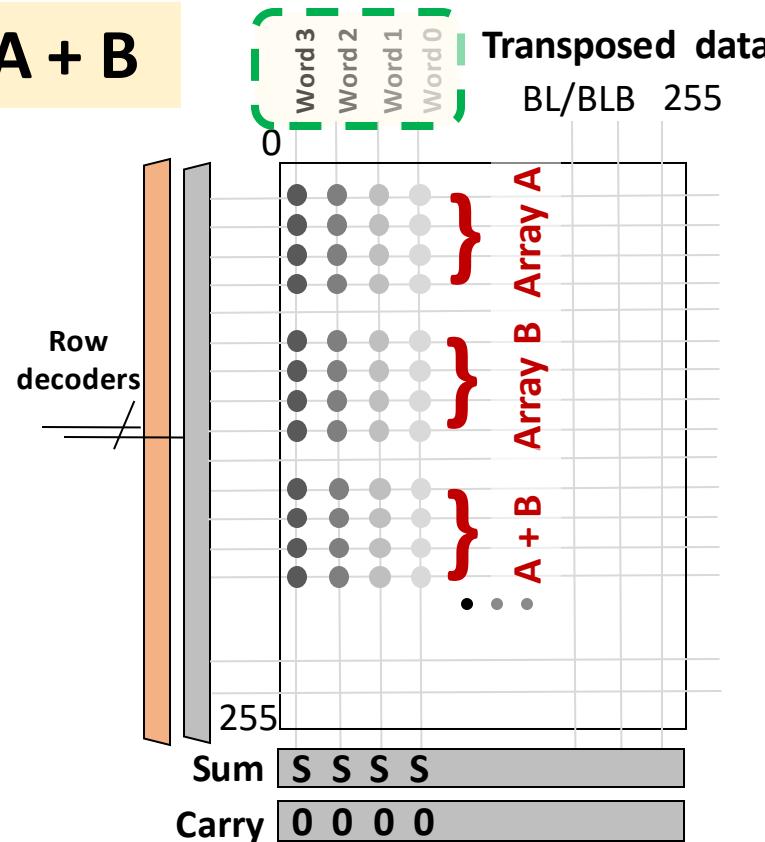
 $+ =$

	exp	mnt
$\times 10^4$	999.000	

③ Perform shift+add
 $s_i = a_i + b_{i+\text{ediff}}$

④ Partial normalization

A + B



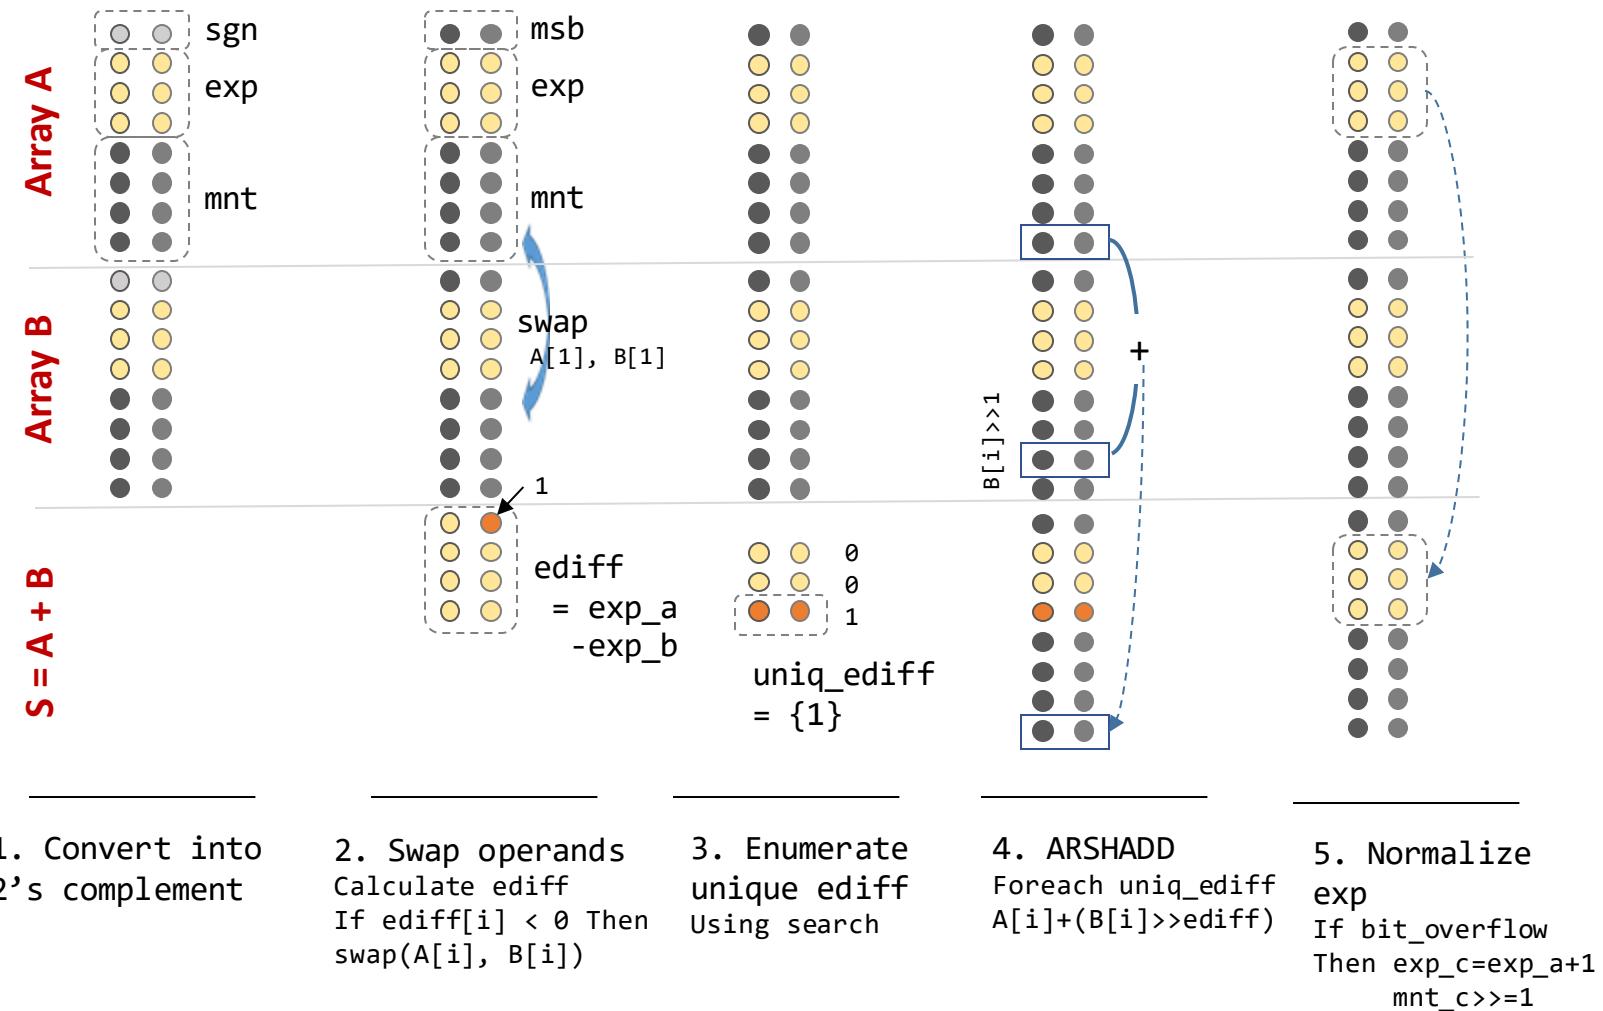
Transposed data

BL/BLB 255

Row
decoders

255

Sum **S S S S**
Carry **0 0 0 0**



Transposed data

BL/BLB 255

Array A

msb
exp
mnt

Swap
 $A[1], B[1]$

ediff
 $= \text{exp}_a - \text{exp}_b$

exp
mnt

uniq_ediff
 $= \{1\}$

exp
mnt

$B[i] \gg 1$

+

exp
mnt

+

exp
mnt

Latency Optimizations (Integer, FP)



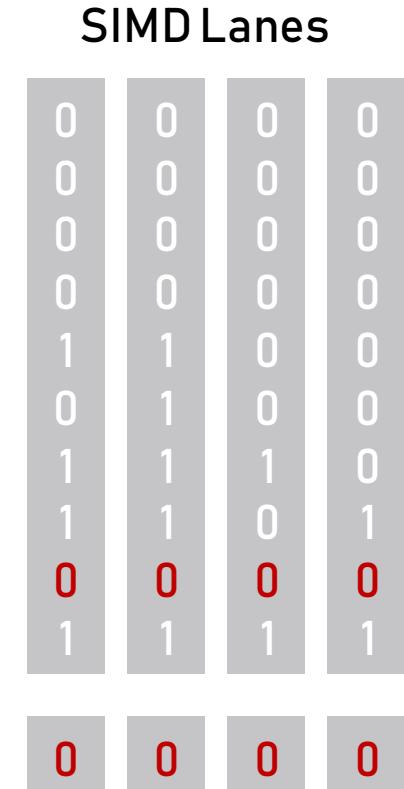
Search leading zeros.

Skip inner loop iterations & reduce inner loop cycles.

$$\begin{array}{r} 001010 \\ \times 000101 \\ \hline \end{array}$$

n_a n_b

e.g. Multiplication $O(n^2) \rightarrow O(n_a \times n_b)$



Check tag bits are all zero.

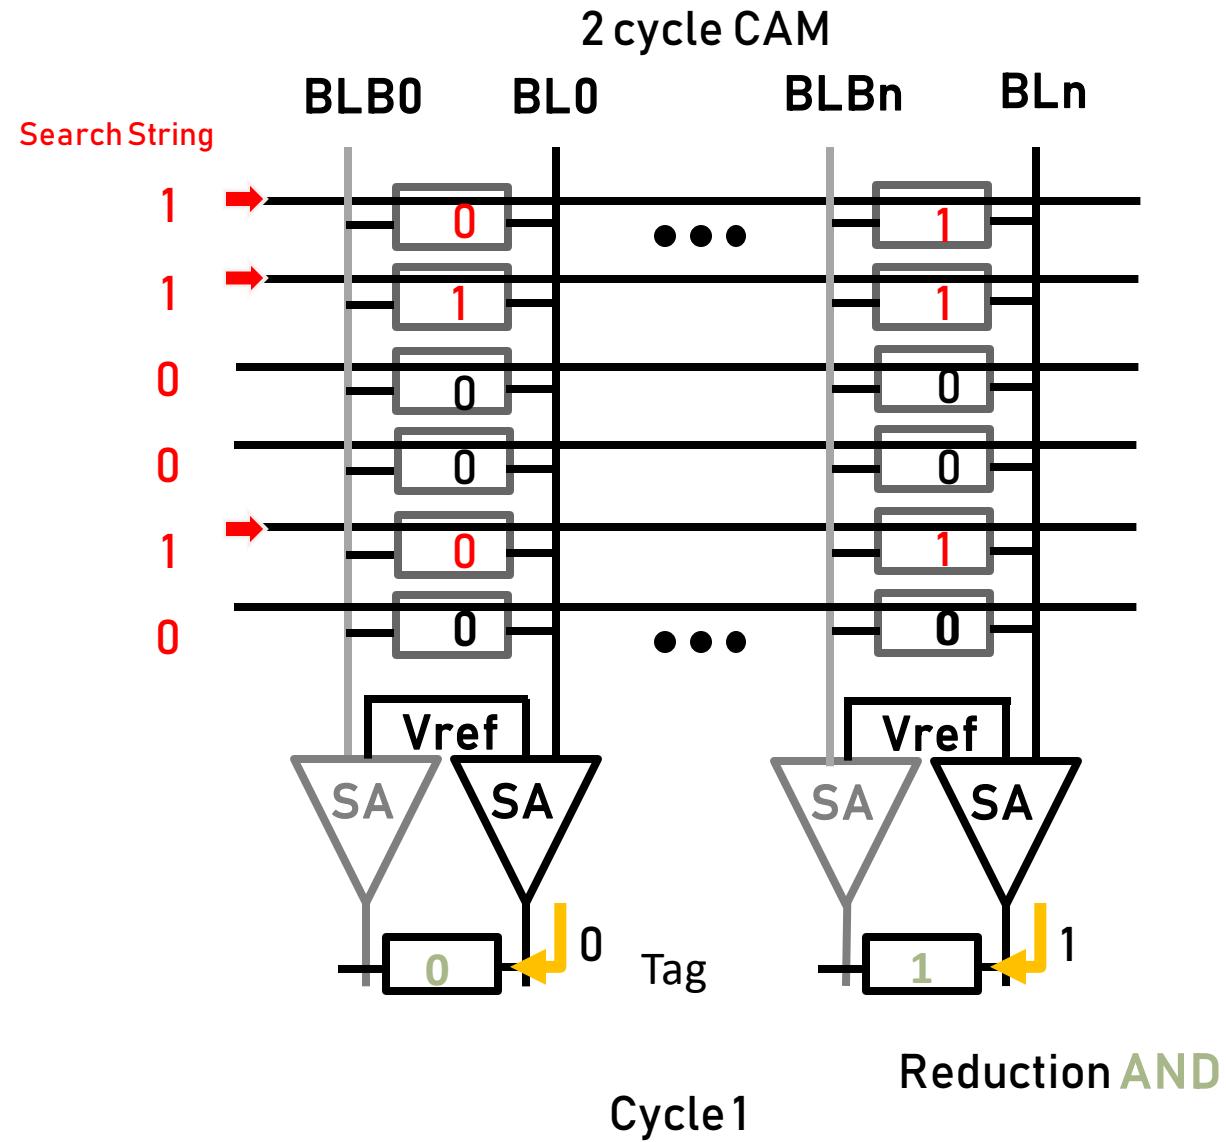
Skip the inner loop iteration.

Useful for (const) integer multiplications.

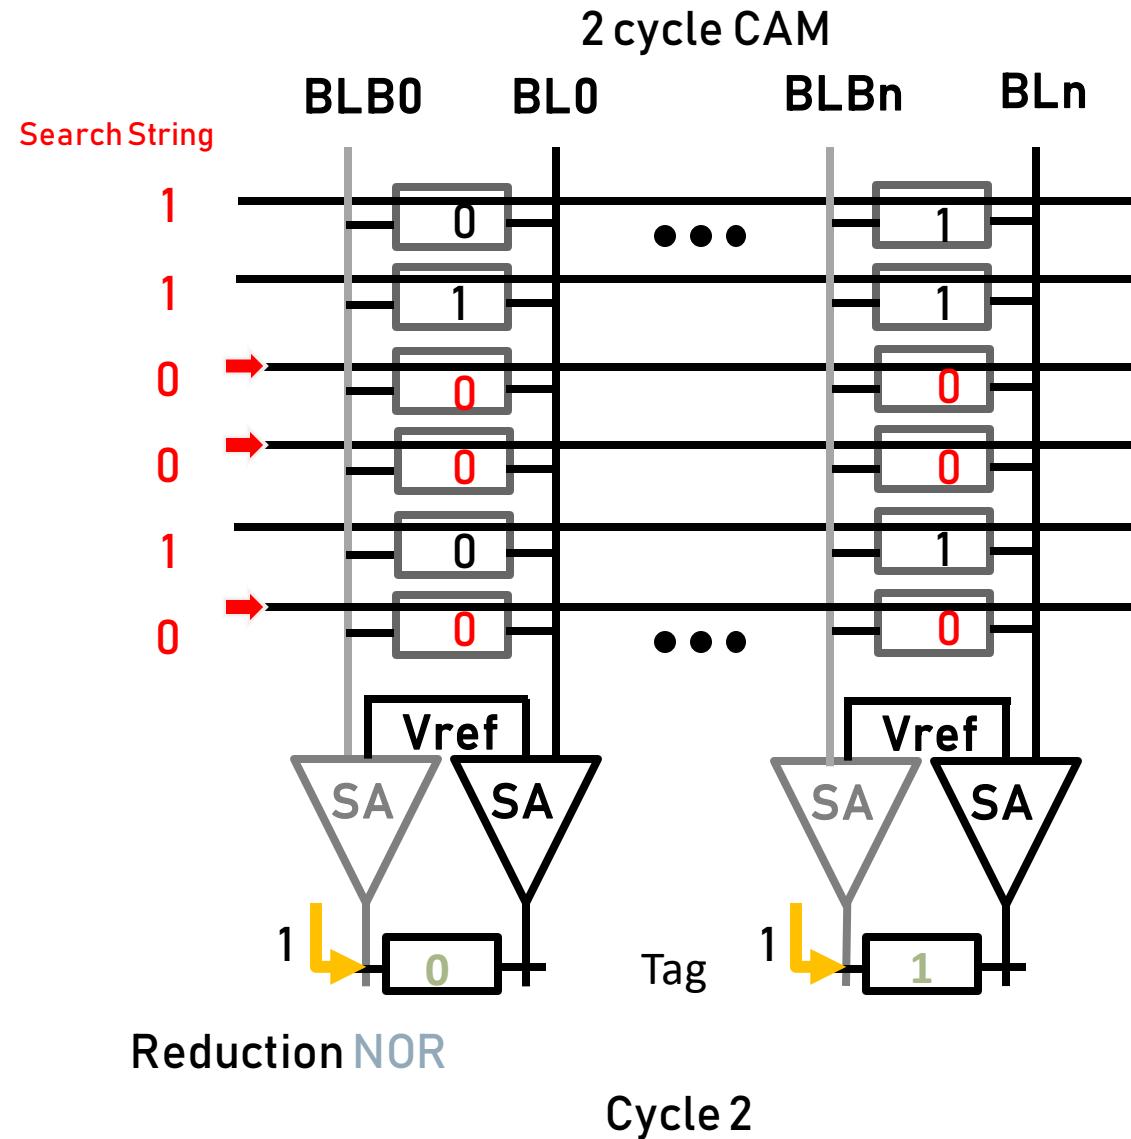
Leading Zero Search

Zero Tag Search

Optimizations with CAM Search



Optimizations with CAM Search



Ediff enumeration

1. Perform **leading zero search** to limit the search space for CAM search.
(e.g. $\forall \text{ediff} < 8 \Rightarrow 4\text{-bit CAM search}$)
2. Perform **CAM search** over ediff vector for values $0 \leq i \leq 2^n$.
3. If any hit (wired OR of tag), perform ARSHADD

Supported In-Cache Operations

Operation	Type	Algorithm	Latency	Optimizations
add, sub	uint, int	[2], Bit-serial	$O(n)$	-
mul	uint, int	[2], Bit-serial	$O(n^2)$	LZS (multiplicand), ZTS
div, rem				divisor), ZTS,
and, or, xor				
shl, shr				search
add, sub				search
mul				
div				
sin, cos				CORDIC
exp				CORDIC
log	fixed point	CORDIC	$O(nk)$	Parallel CORDIC
sqrt	fixed point	CORDIC	$O(nk)$	Parallel CORDIC
rsqrt	float	Fast Inv Sqrt	$O(n^2)$	

Algorithm
[1] Compute Caches (HPCA'17)
[2] Neural Cache (ISCA'18)

Latency
n: Data Bit Length (bit)
k: CORDIC iteration count

Optimizations
LZS: Leading Zero Search
ZTS: Zero Tag Search
ZRS: Zero Residue Search



Transcendental functions using CORDIC

Express an angle θ using add/sub of a series. ($0 \leq \theta < \pi/2$)

$$\theta = \sum \pm \alpha_i$$

Vector rotation requires multiplications with $\tan(\alpha_i)$.

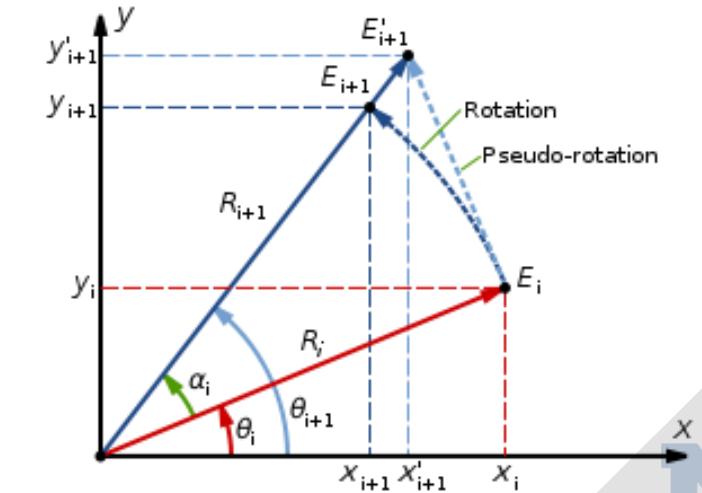
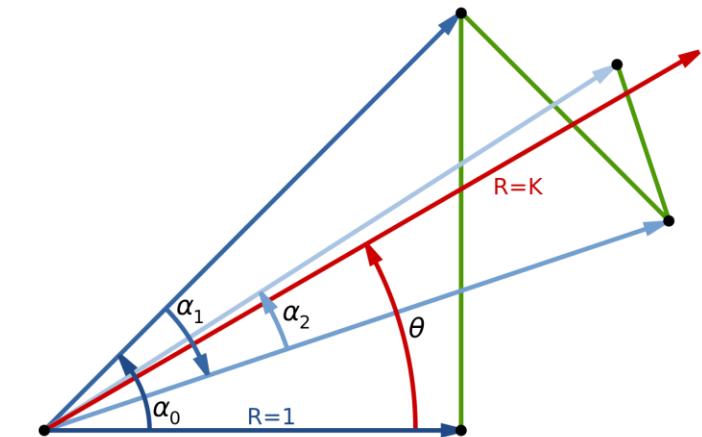
$$\begin{pmatrix} x_i \\ y_i \end{pmatrix} = \begin{pmatrix} \cos(\alpha_i) & -\sin(\alpha_i) \\ \sin(\alpha_i) & \cos(\alpha_i) \end{pmatrix} \begin{pmatrix} x_{i-1} \\ y_{i-1} \end{pmatrix} = \cos(\alpha_i) \begin{pmatrix} 1 & -\tan(\alpha_i) \\ \tan(\alpha_i) & 1 \end{pmatrix} \begin{pmatrix} x_{i-1} \\ y_{i-1} \end{pmatrix}$$

CORDIC performs pseudo rotation so that $\tan(\alpha_i) = \pm 2^{-i}$

$$\begin{pmatrix} x_i \\ y_i \end{pmatrix} = K \begin{pmatrix} x_{i-1} \mp y_{i-1} \times 2^{-i} \\ y_{i-1} \pm x_{i-1} \times 2^{-i} \end{pmatrix}, \quad \theta = \sum \pm \arctan 2^{-i}$$

Benefits

- No multiplications.
- Can be combined with Bit-serial algorithms.
- Constant static α_i series for all possible inputs. No LUT required.
- Highly parallelizable for large SIMD processor.
- Operation level parallelism.



Outline

Background

Integer / Logical Operations

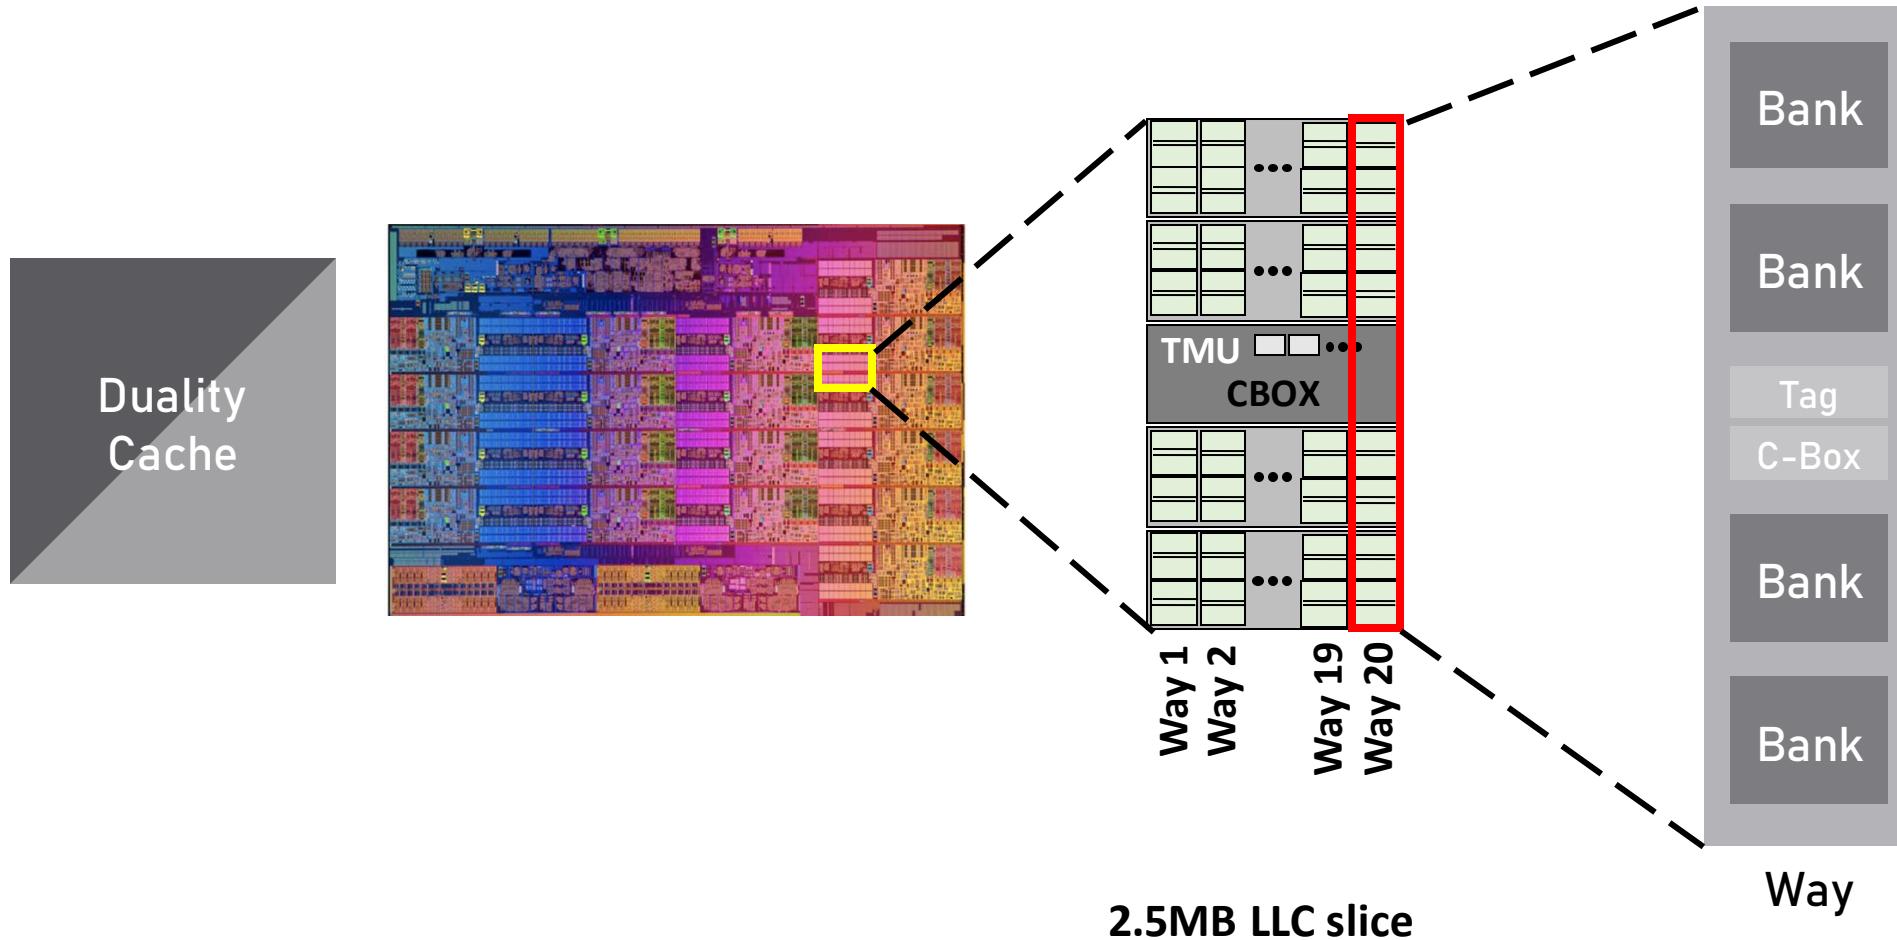
Floating Point Operations

Transcendental Functions

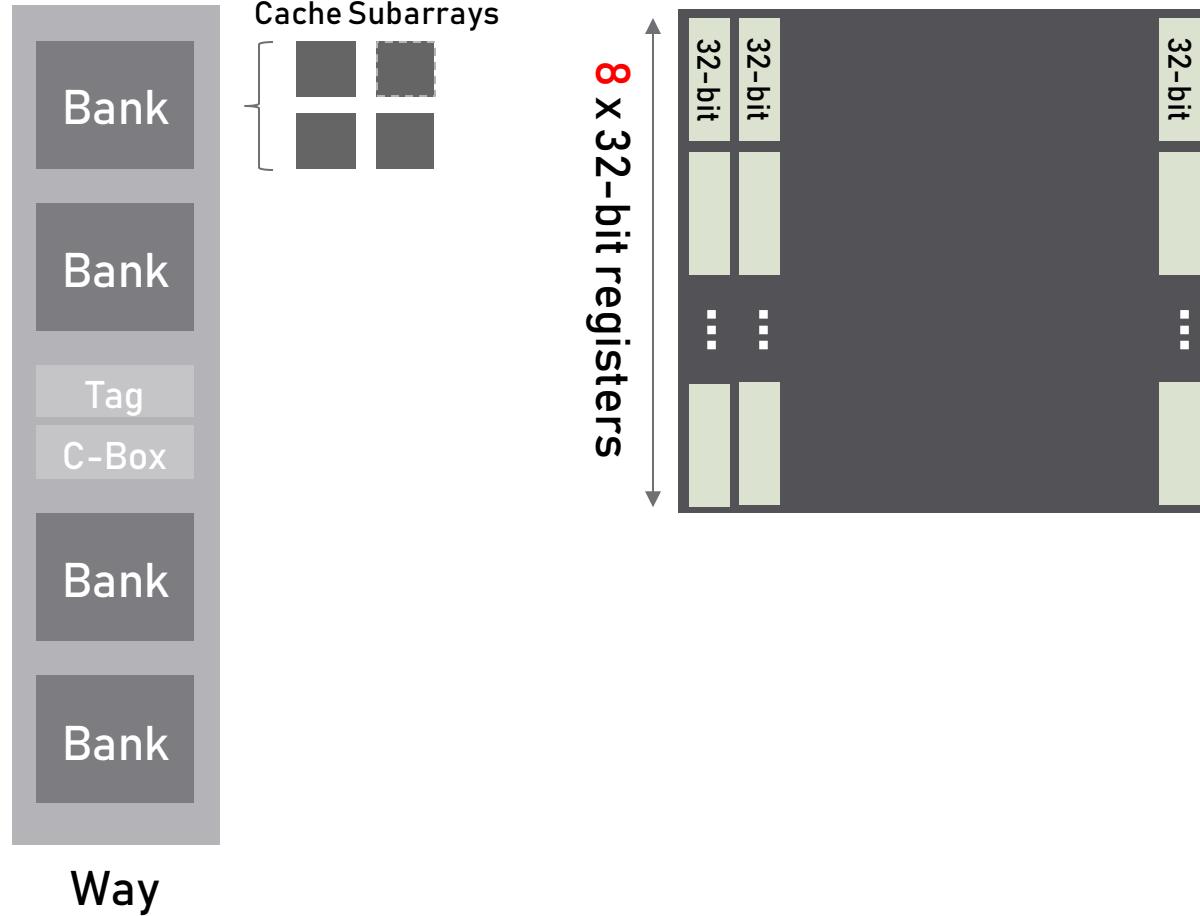
- ▶ Execution Model
- ▶ Programming Model
- ▶ Compiler

Methodology / Results

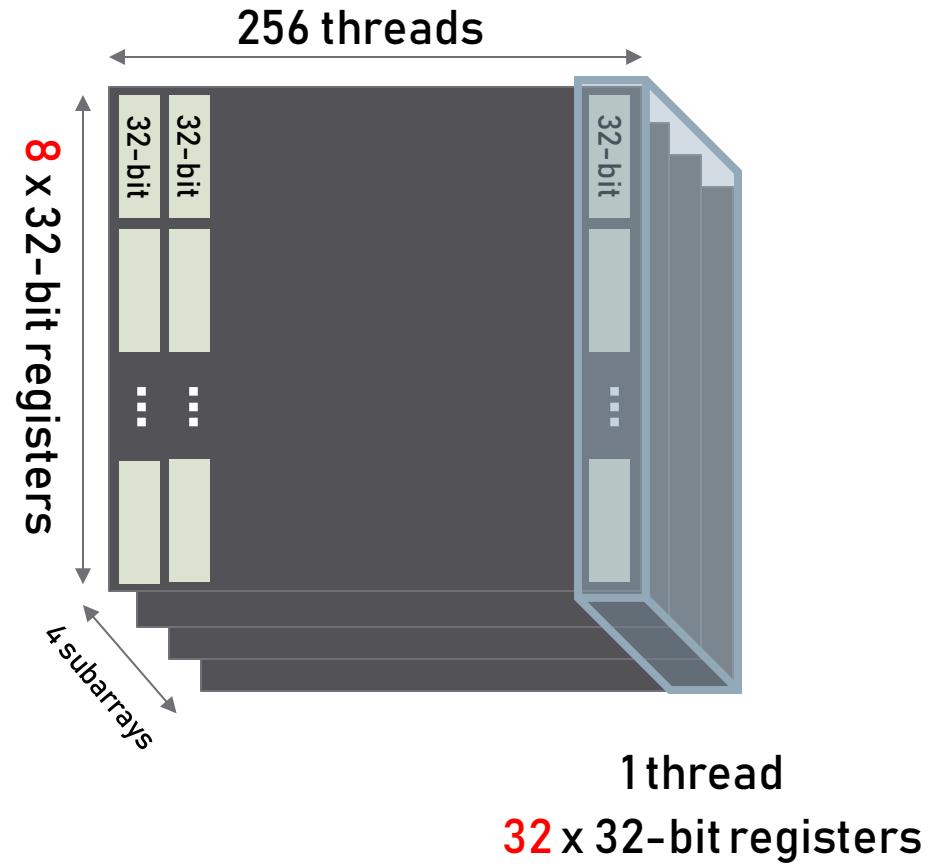
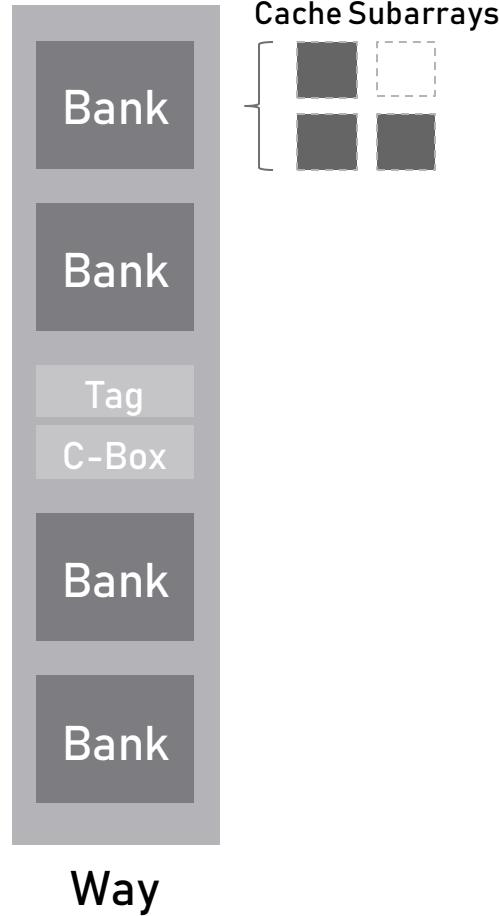
Execution Model



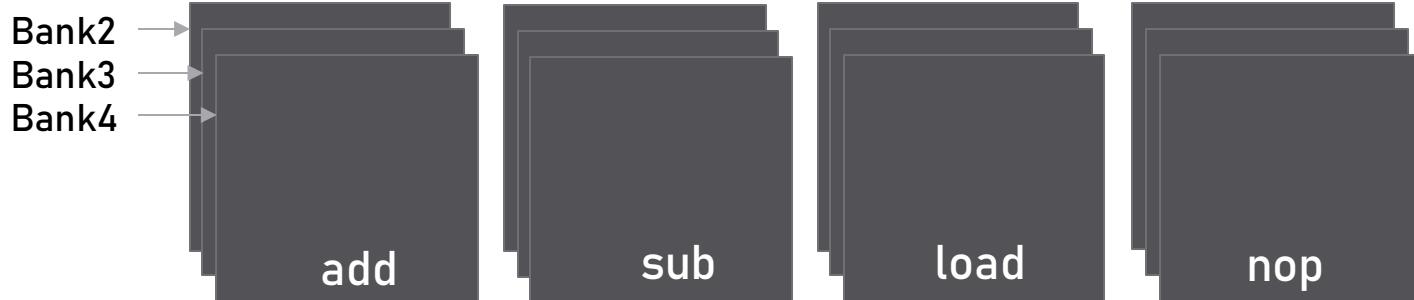
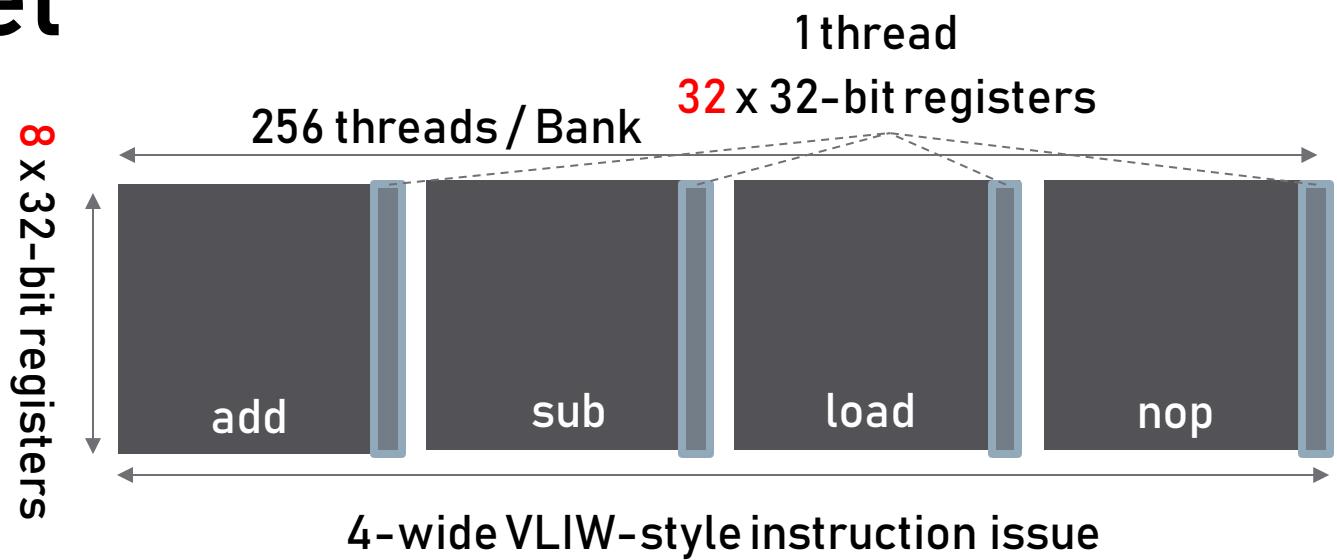
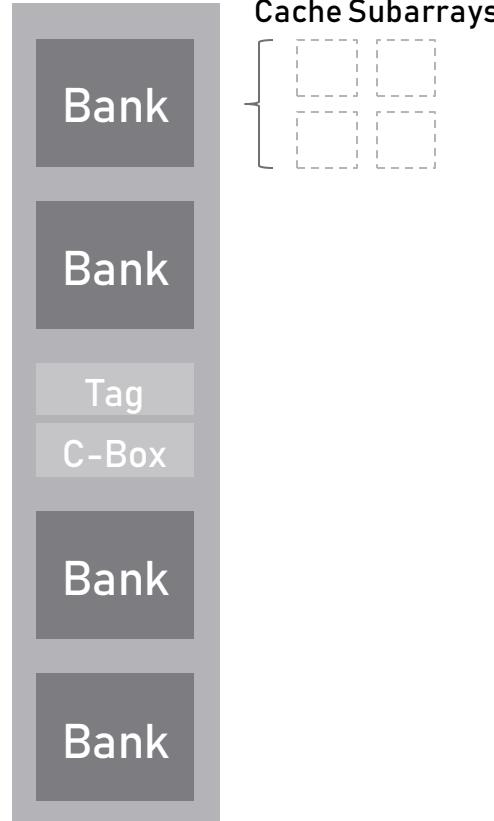
Execution Model



Execution Model

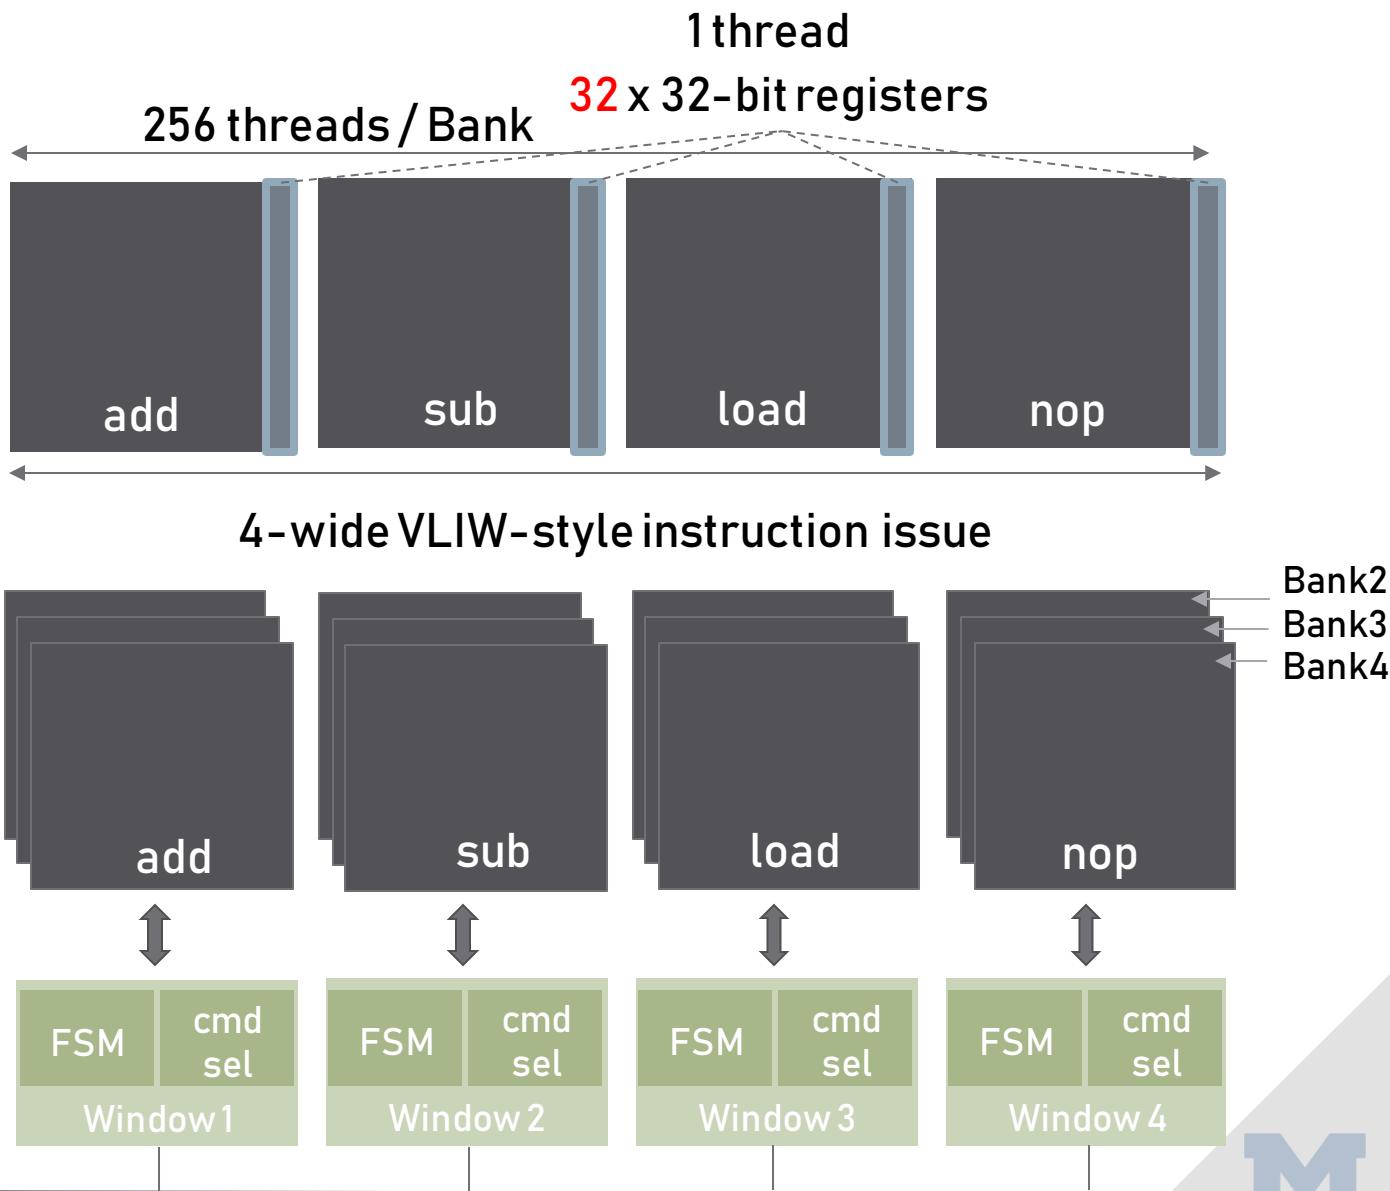
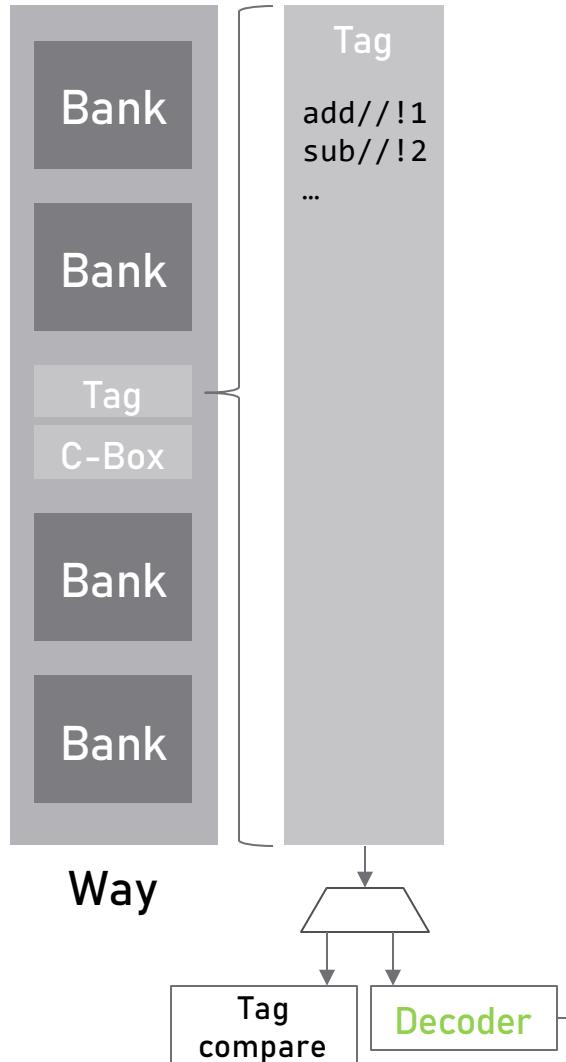


Execution Model

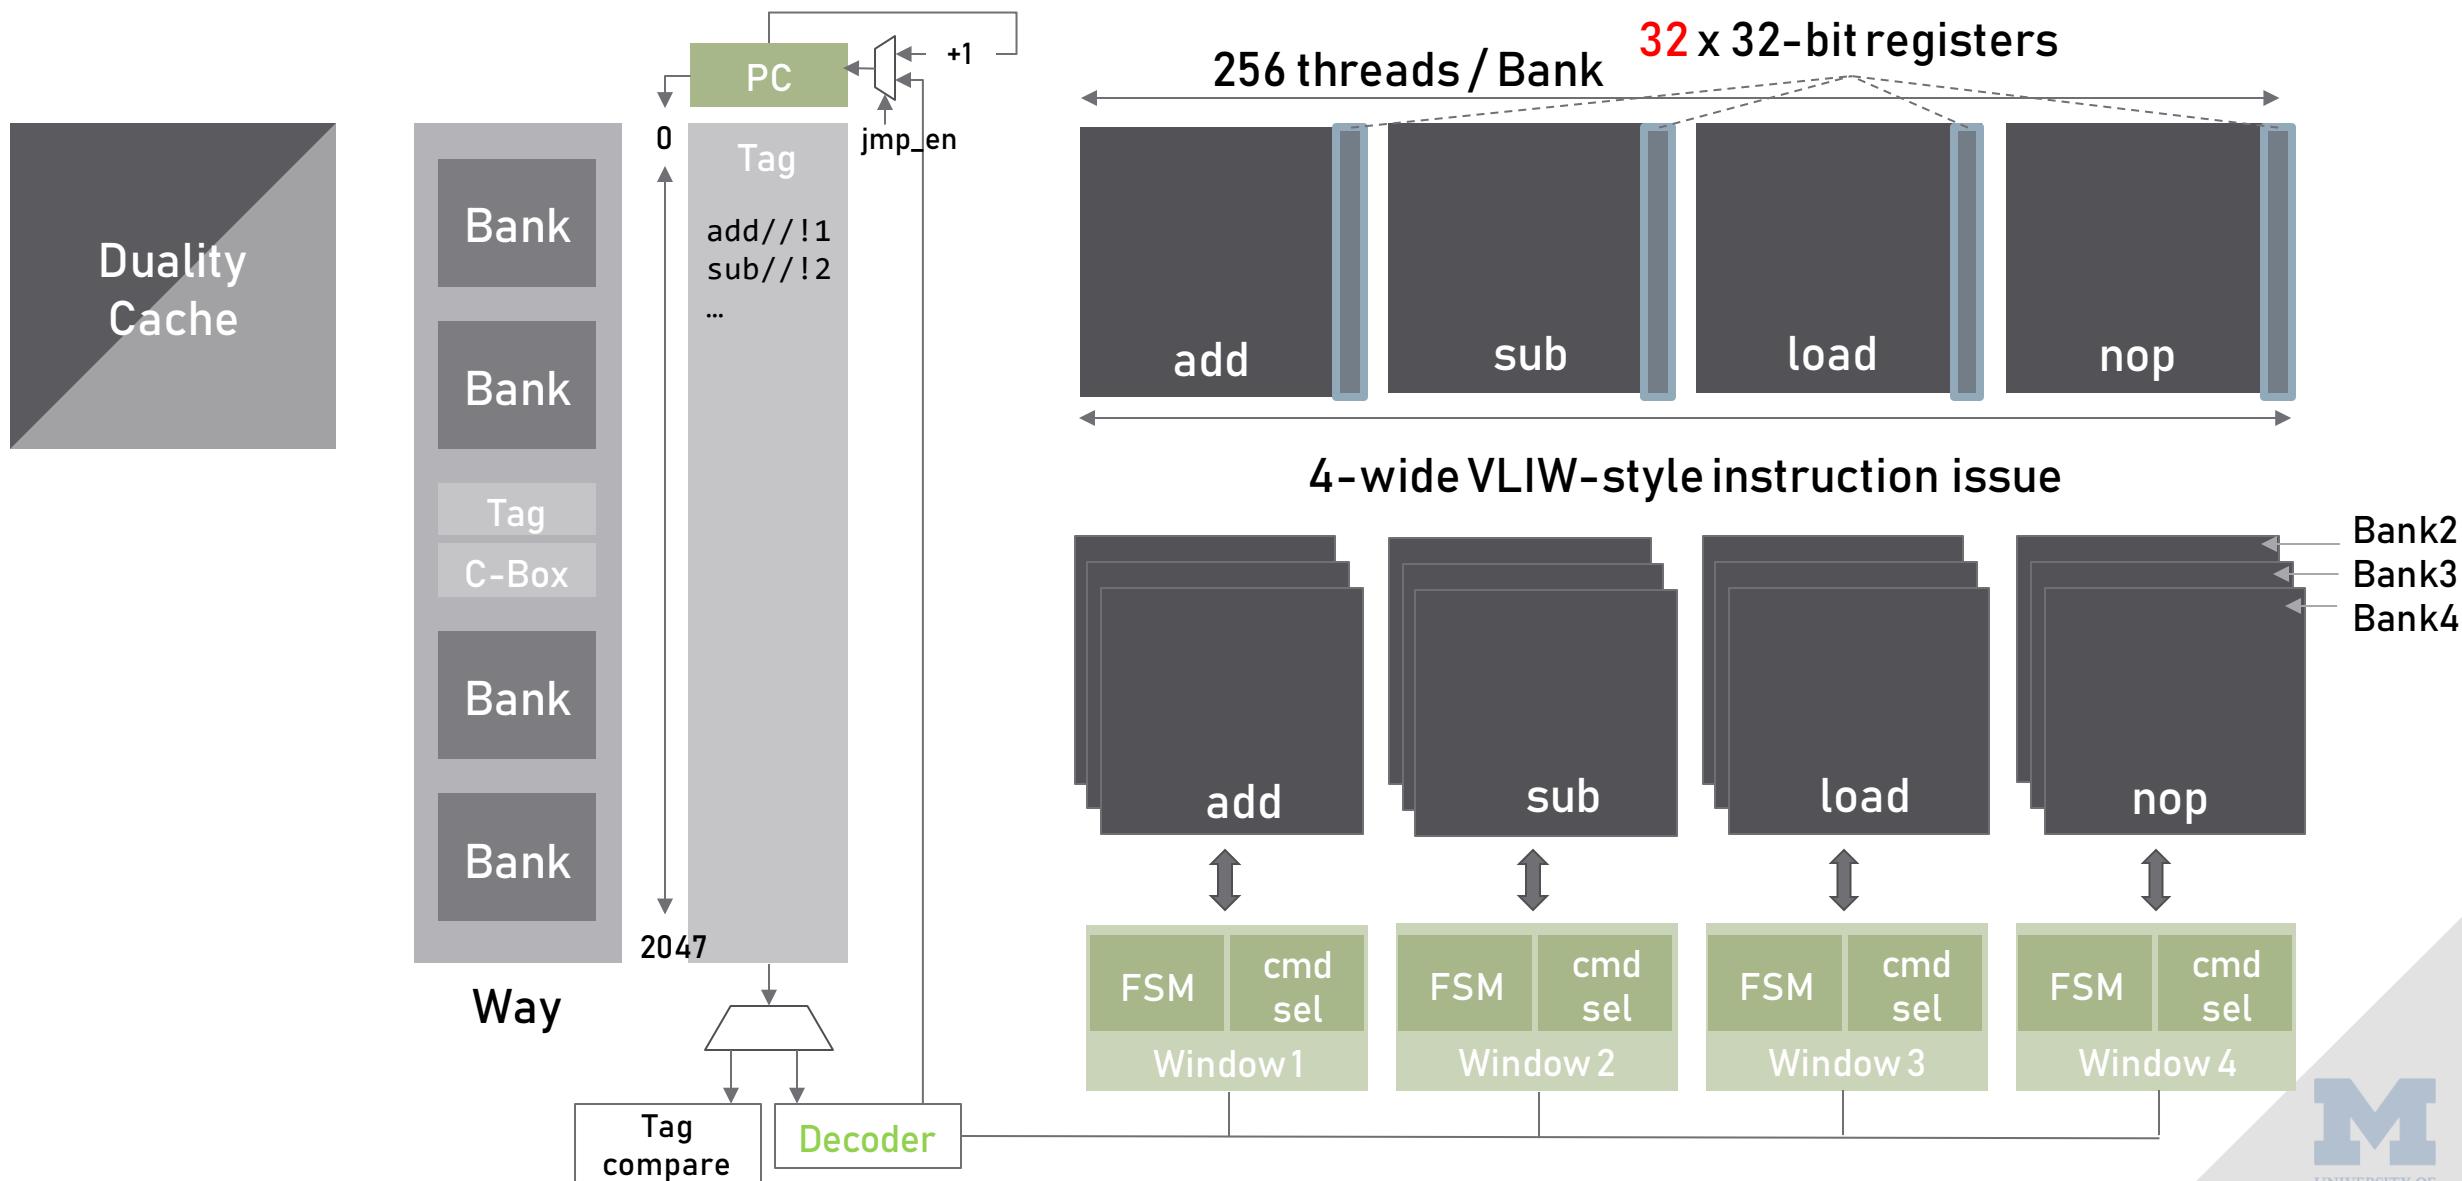


1024 thread 4-wide VLIW SIMD processor

Execution Model

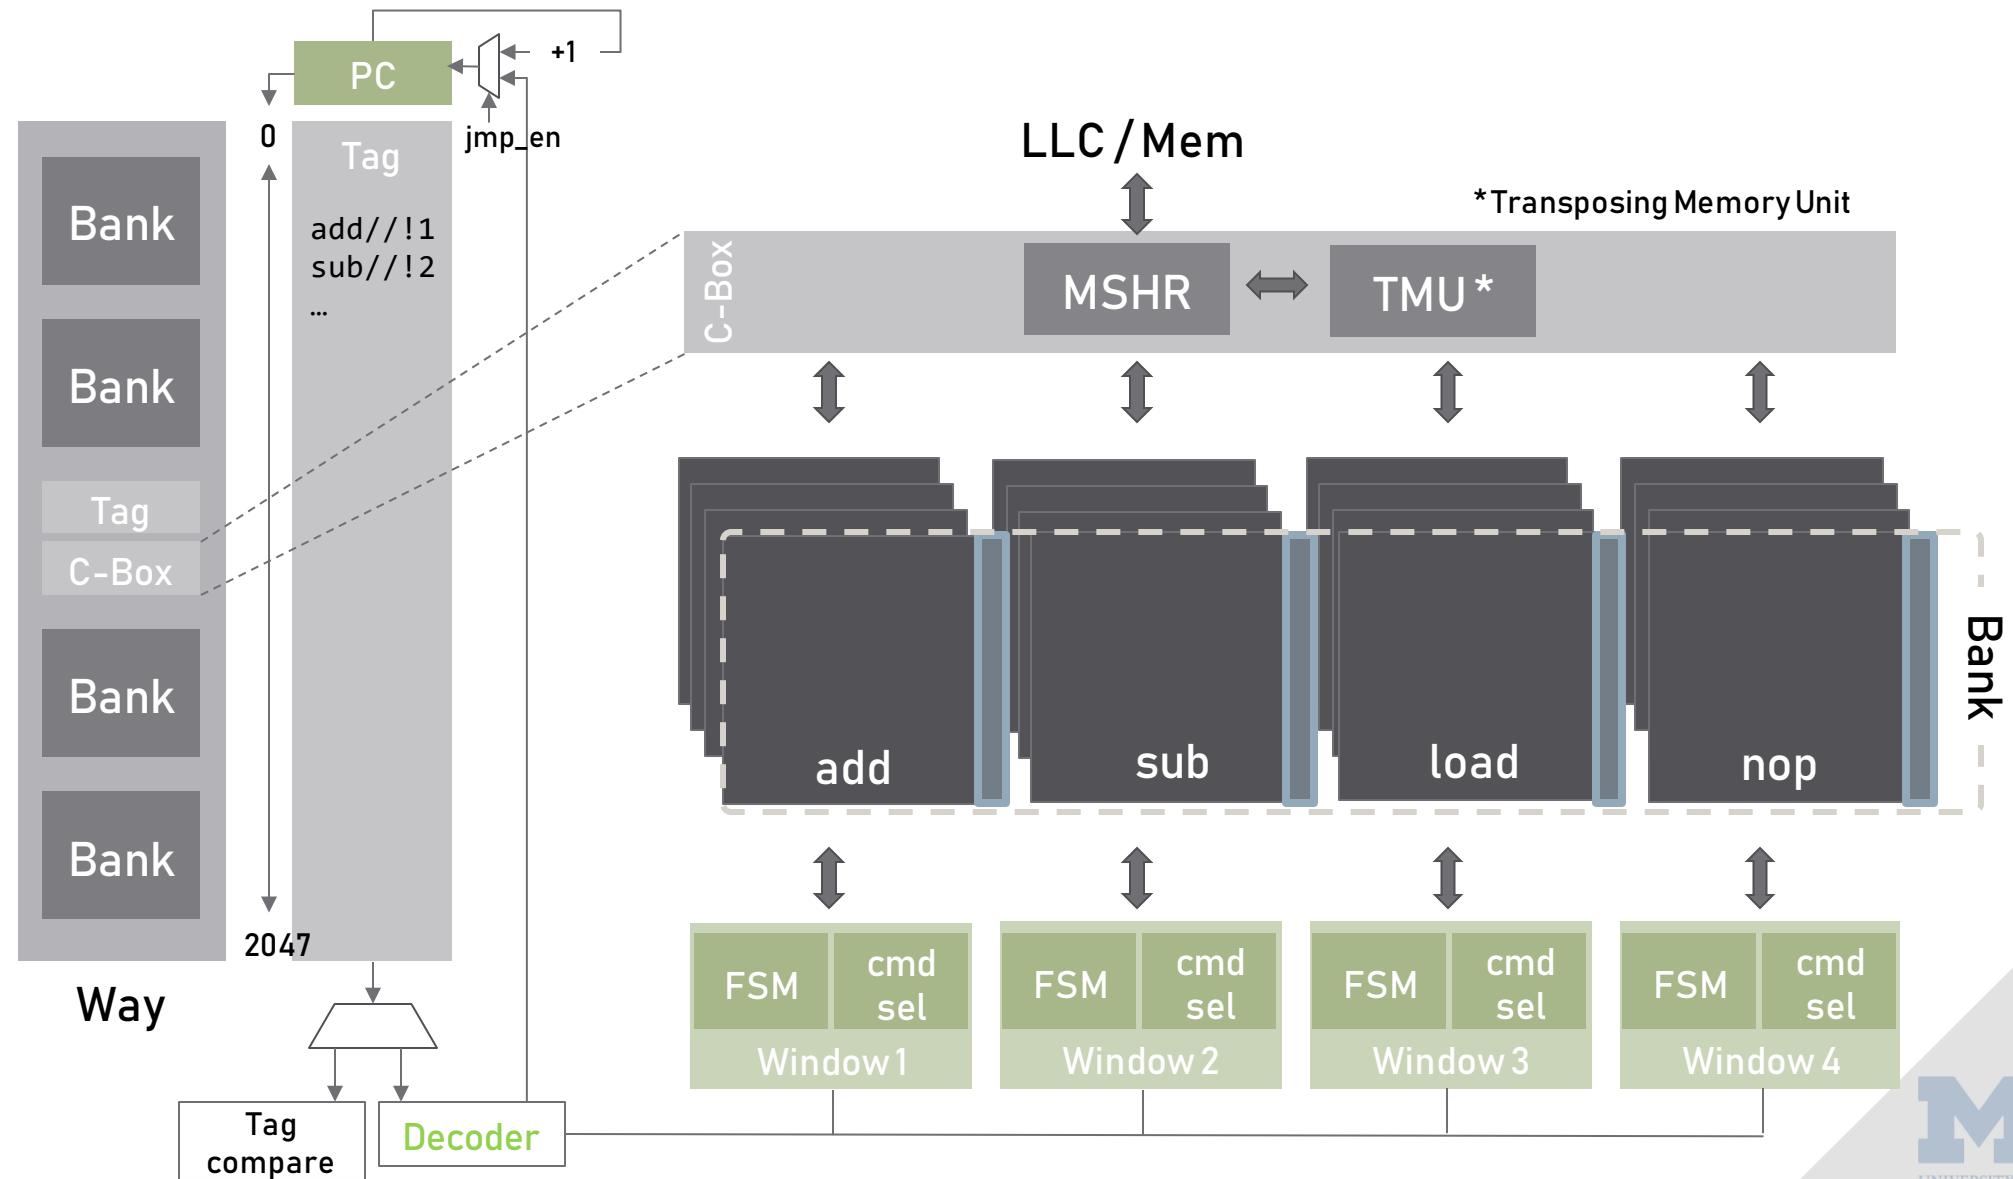


Execution Model



Execution Model

33

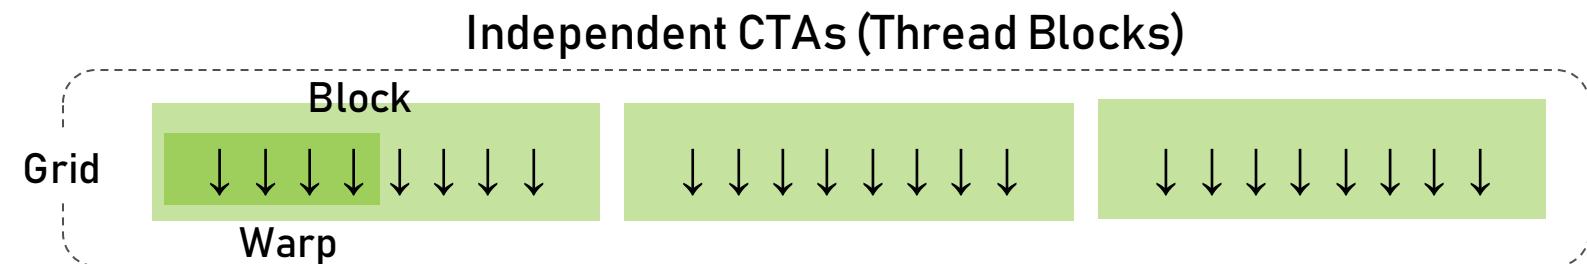


Programming Model



```
//kernel
__global__ void vecadd (const float* A, const float* B, float* C, int n_el) {
    int tid = blockDim.x * blockIdx.x + threadIdx.x;
    if (tid < n_el)
        C[tid] = A[tid] + B[tid];
}

//kernel invocation
vecadd <<<blocksPerGrid, threadsPerBlock>>> (A, B, C, n_el);
```



Programming Model

OpenACC

More Science, Less Programming

Data Movement
Management

Parallel Region

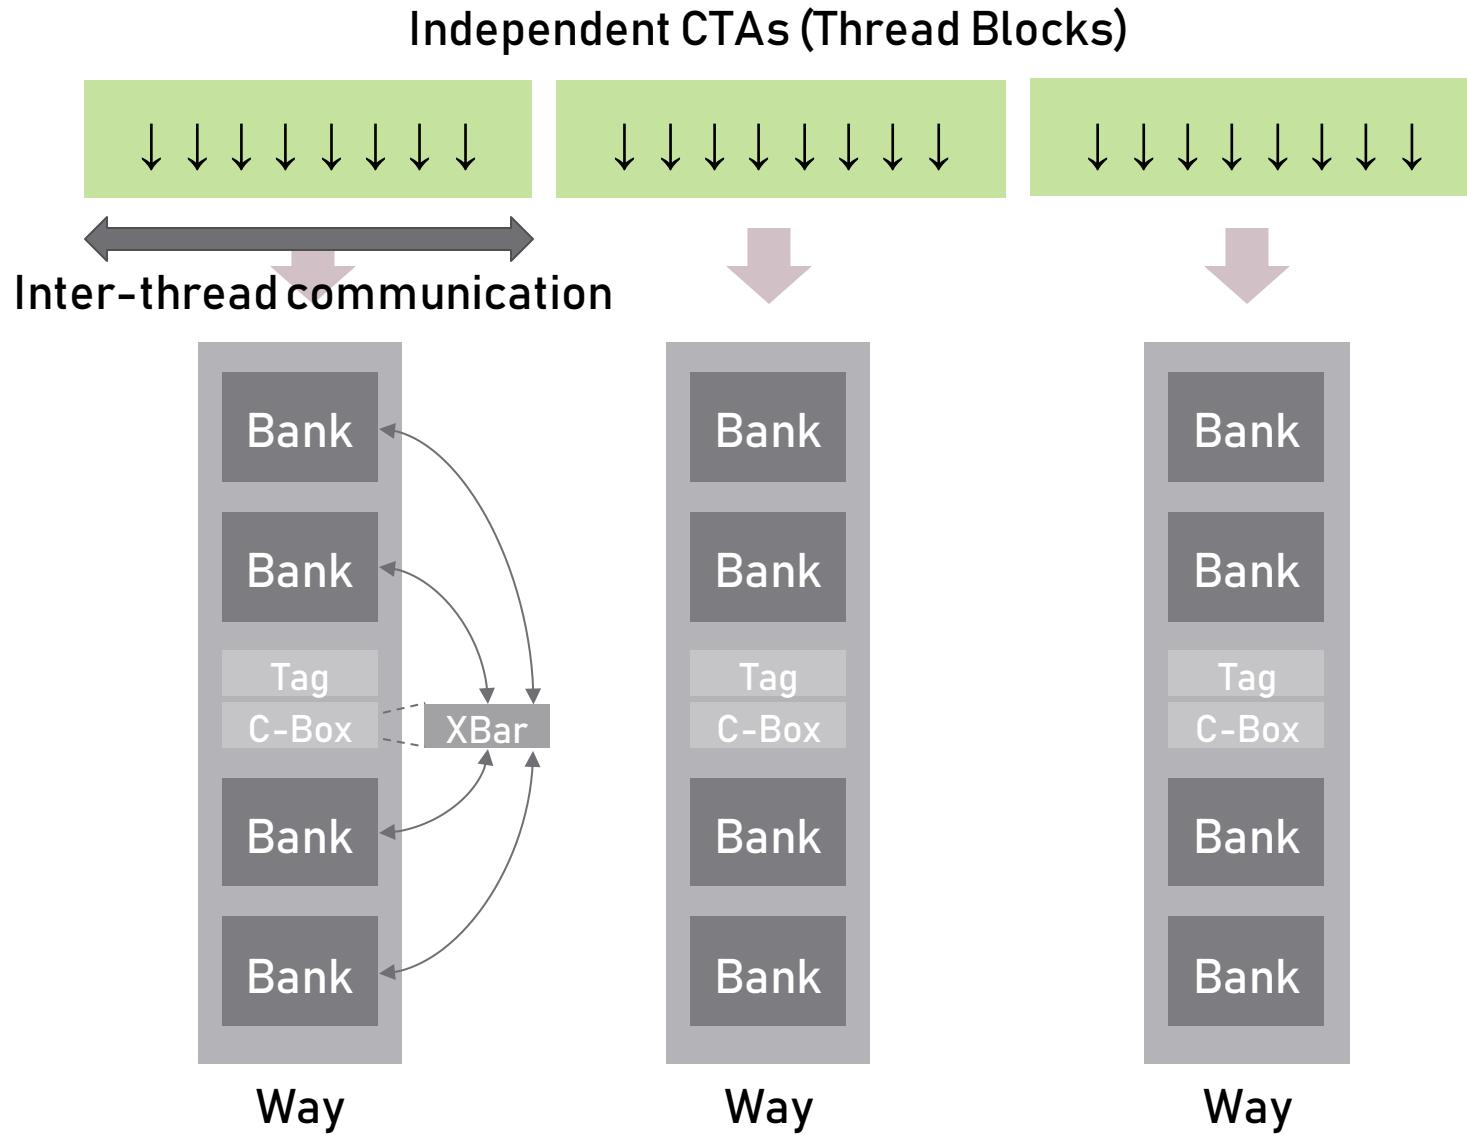
Loop Mapping
Optimization

```
#pragma acc data copyin(a,b) copyout(c)
{
    ...
    #pragma acc parallel
    {
        #pragma acc loop gang vector
        for (i = 0; i < n; ++i) {
            c[i] = a[i] + b[i];
            ...
        }
        ...
    }
}
```

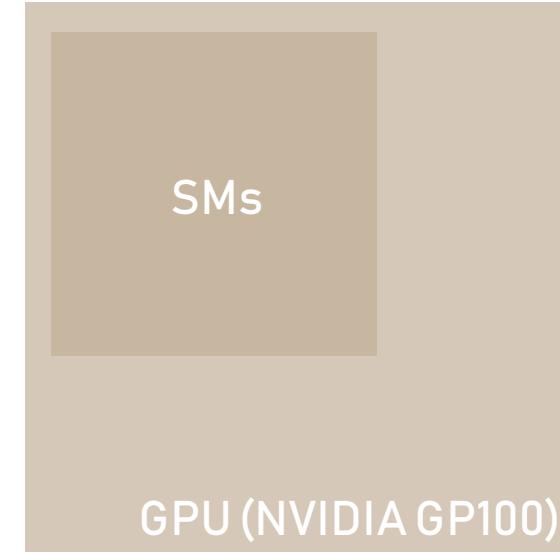
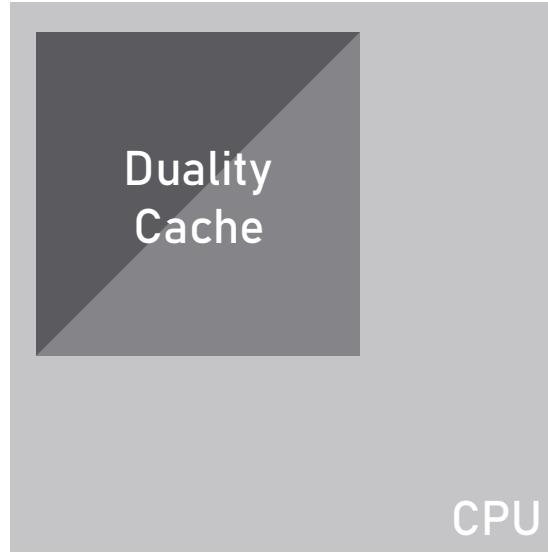
Simple and seamless directive-based parallel programming paradigm for heterogeneous architectures.

Target Architectures: Native x86, NVIDIA GPU, OpenCL
Compiler Support: PGI, gcc9.1, Riken Omni, OpenUH

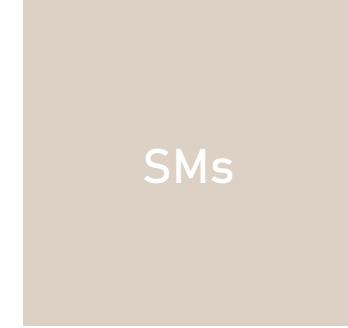
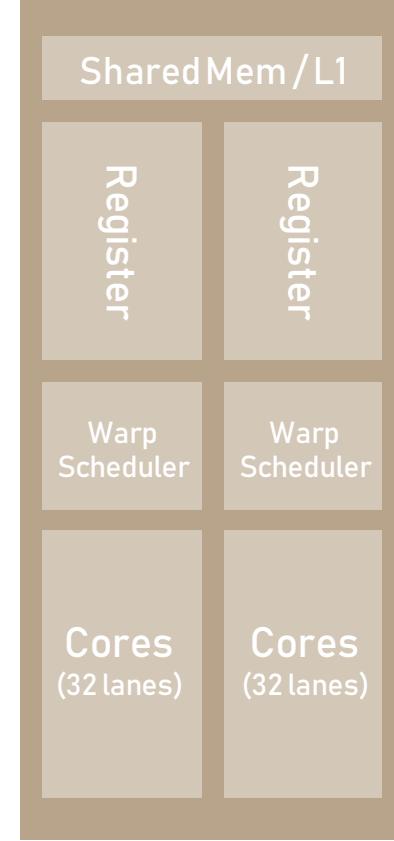
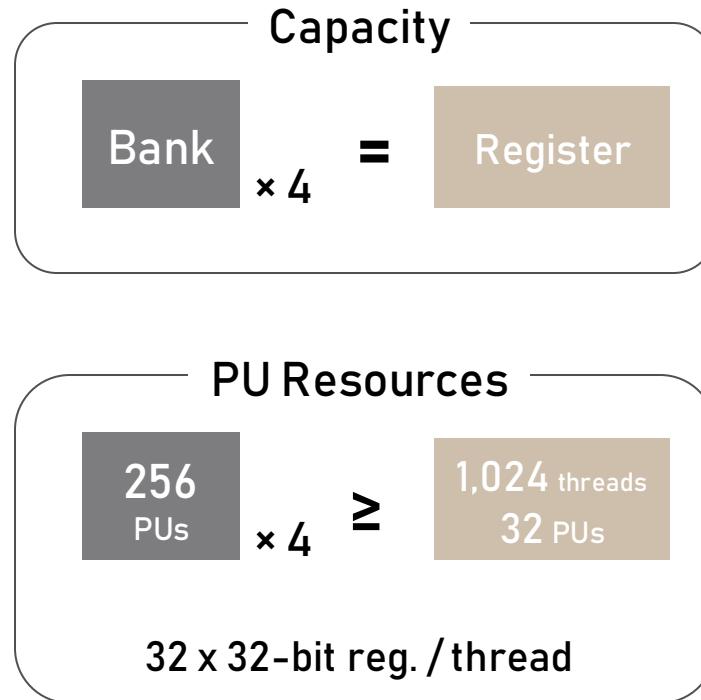
Programming Model



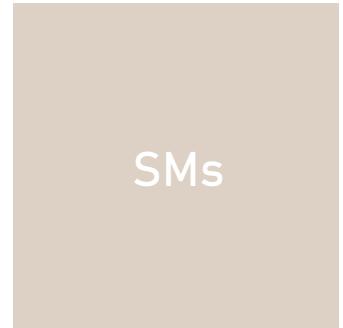
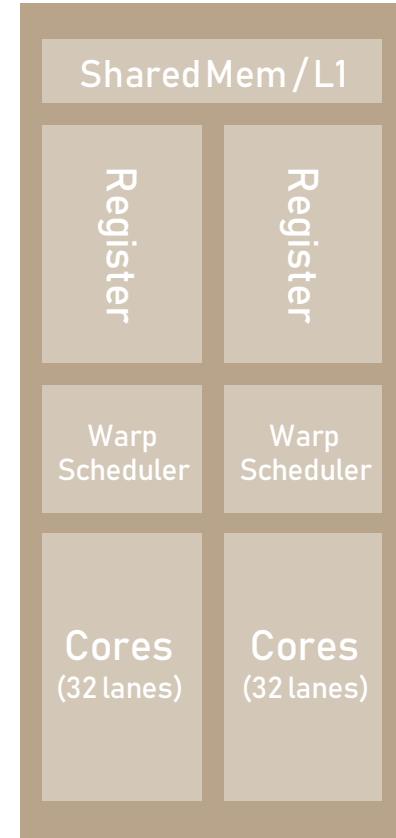
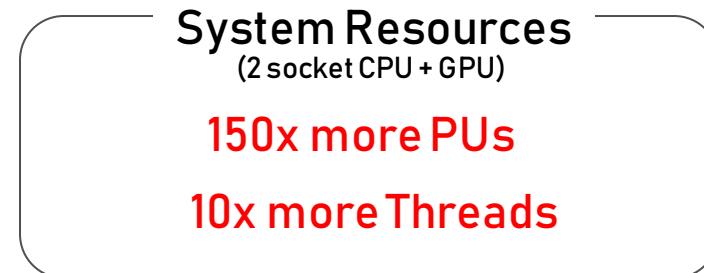
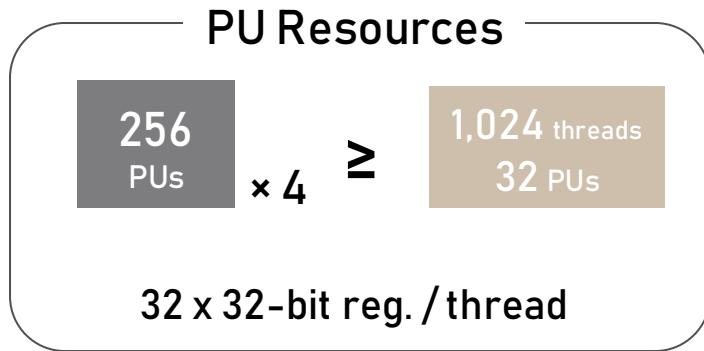
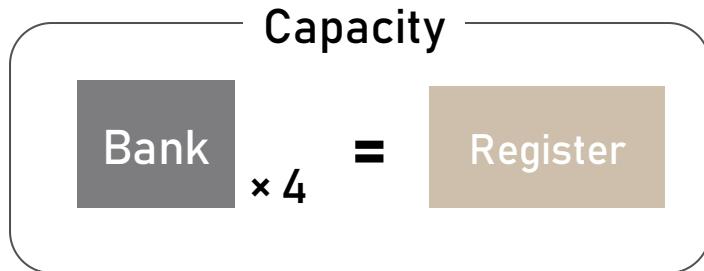
Resource Comparison



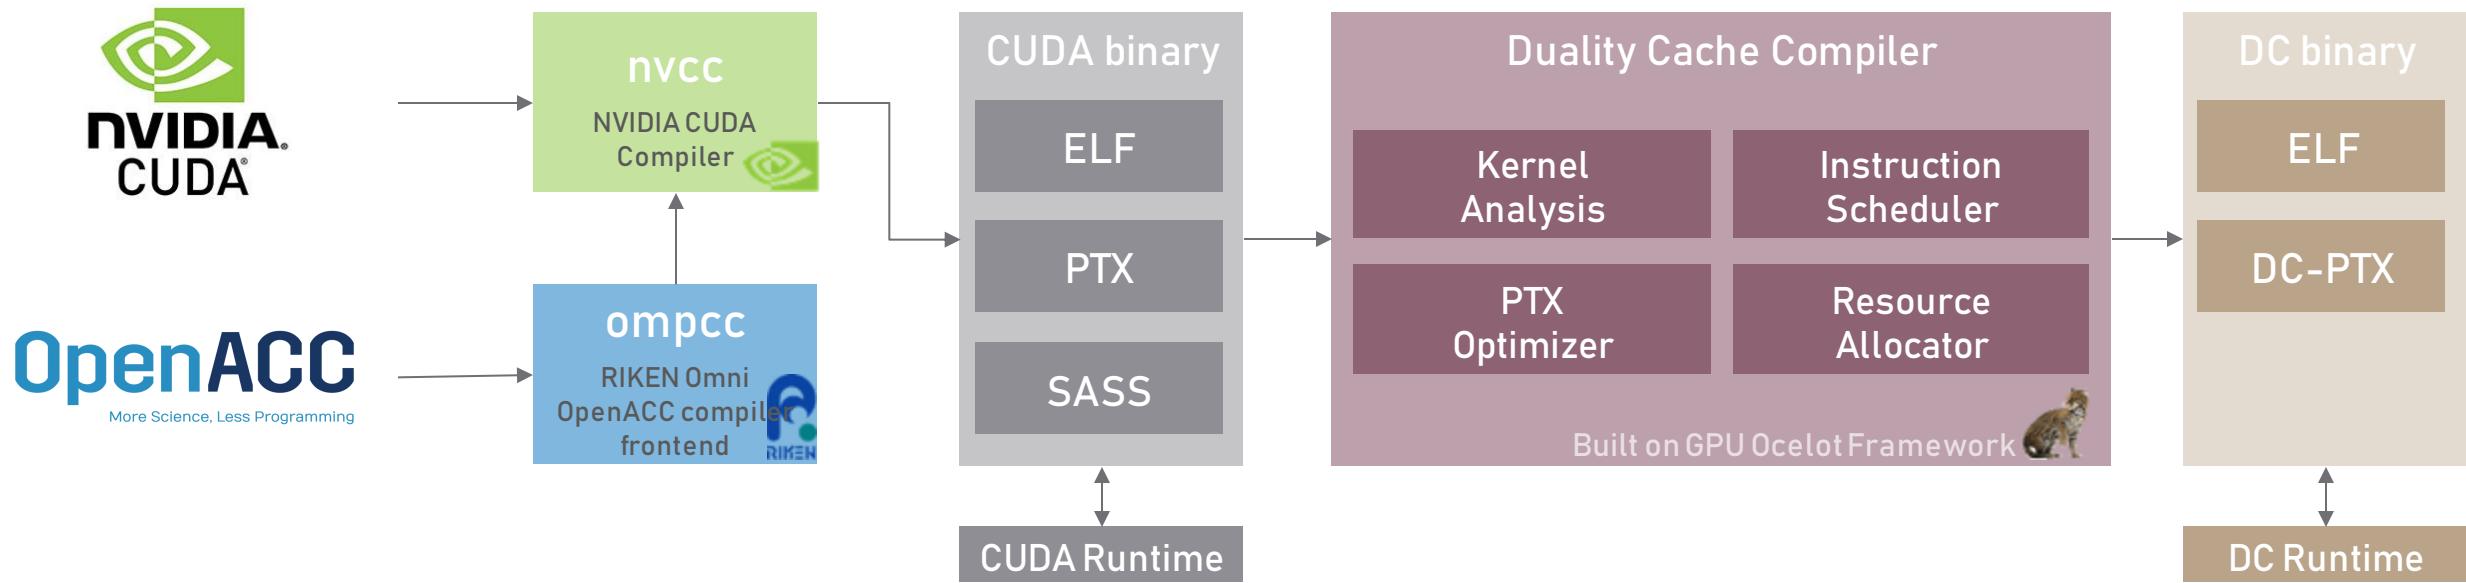
Resource Comparison



Resource Comparison

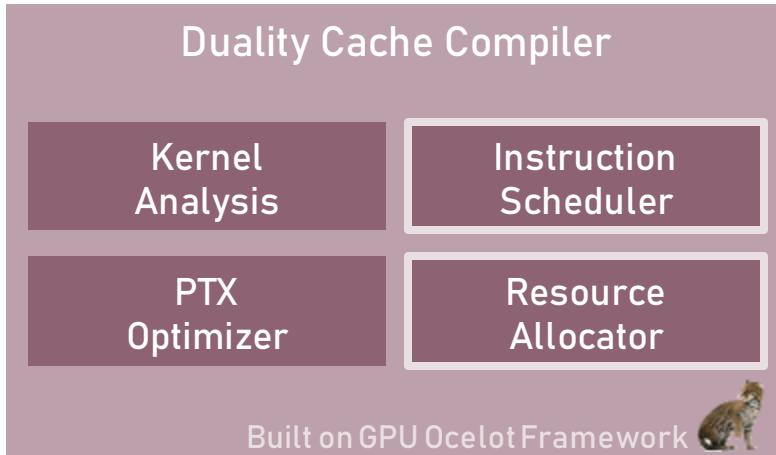


Compiler



Compiler Optimization

(1) Register Pressure Aware Instruction Scheduling



Instruction Scheduling First

- ✓ Maximize parallelism utilizing abundant shared register resources in VLIW architecture
- ! Result in frequent register spill

Resource Allocation First

- ✓ Optimized (small) resource usage and reuse distance
- ! Introduces many false dependencies (less parallelism)

Bottom-Up Greedy
(BUG) based

Instruction
Scheduler

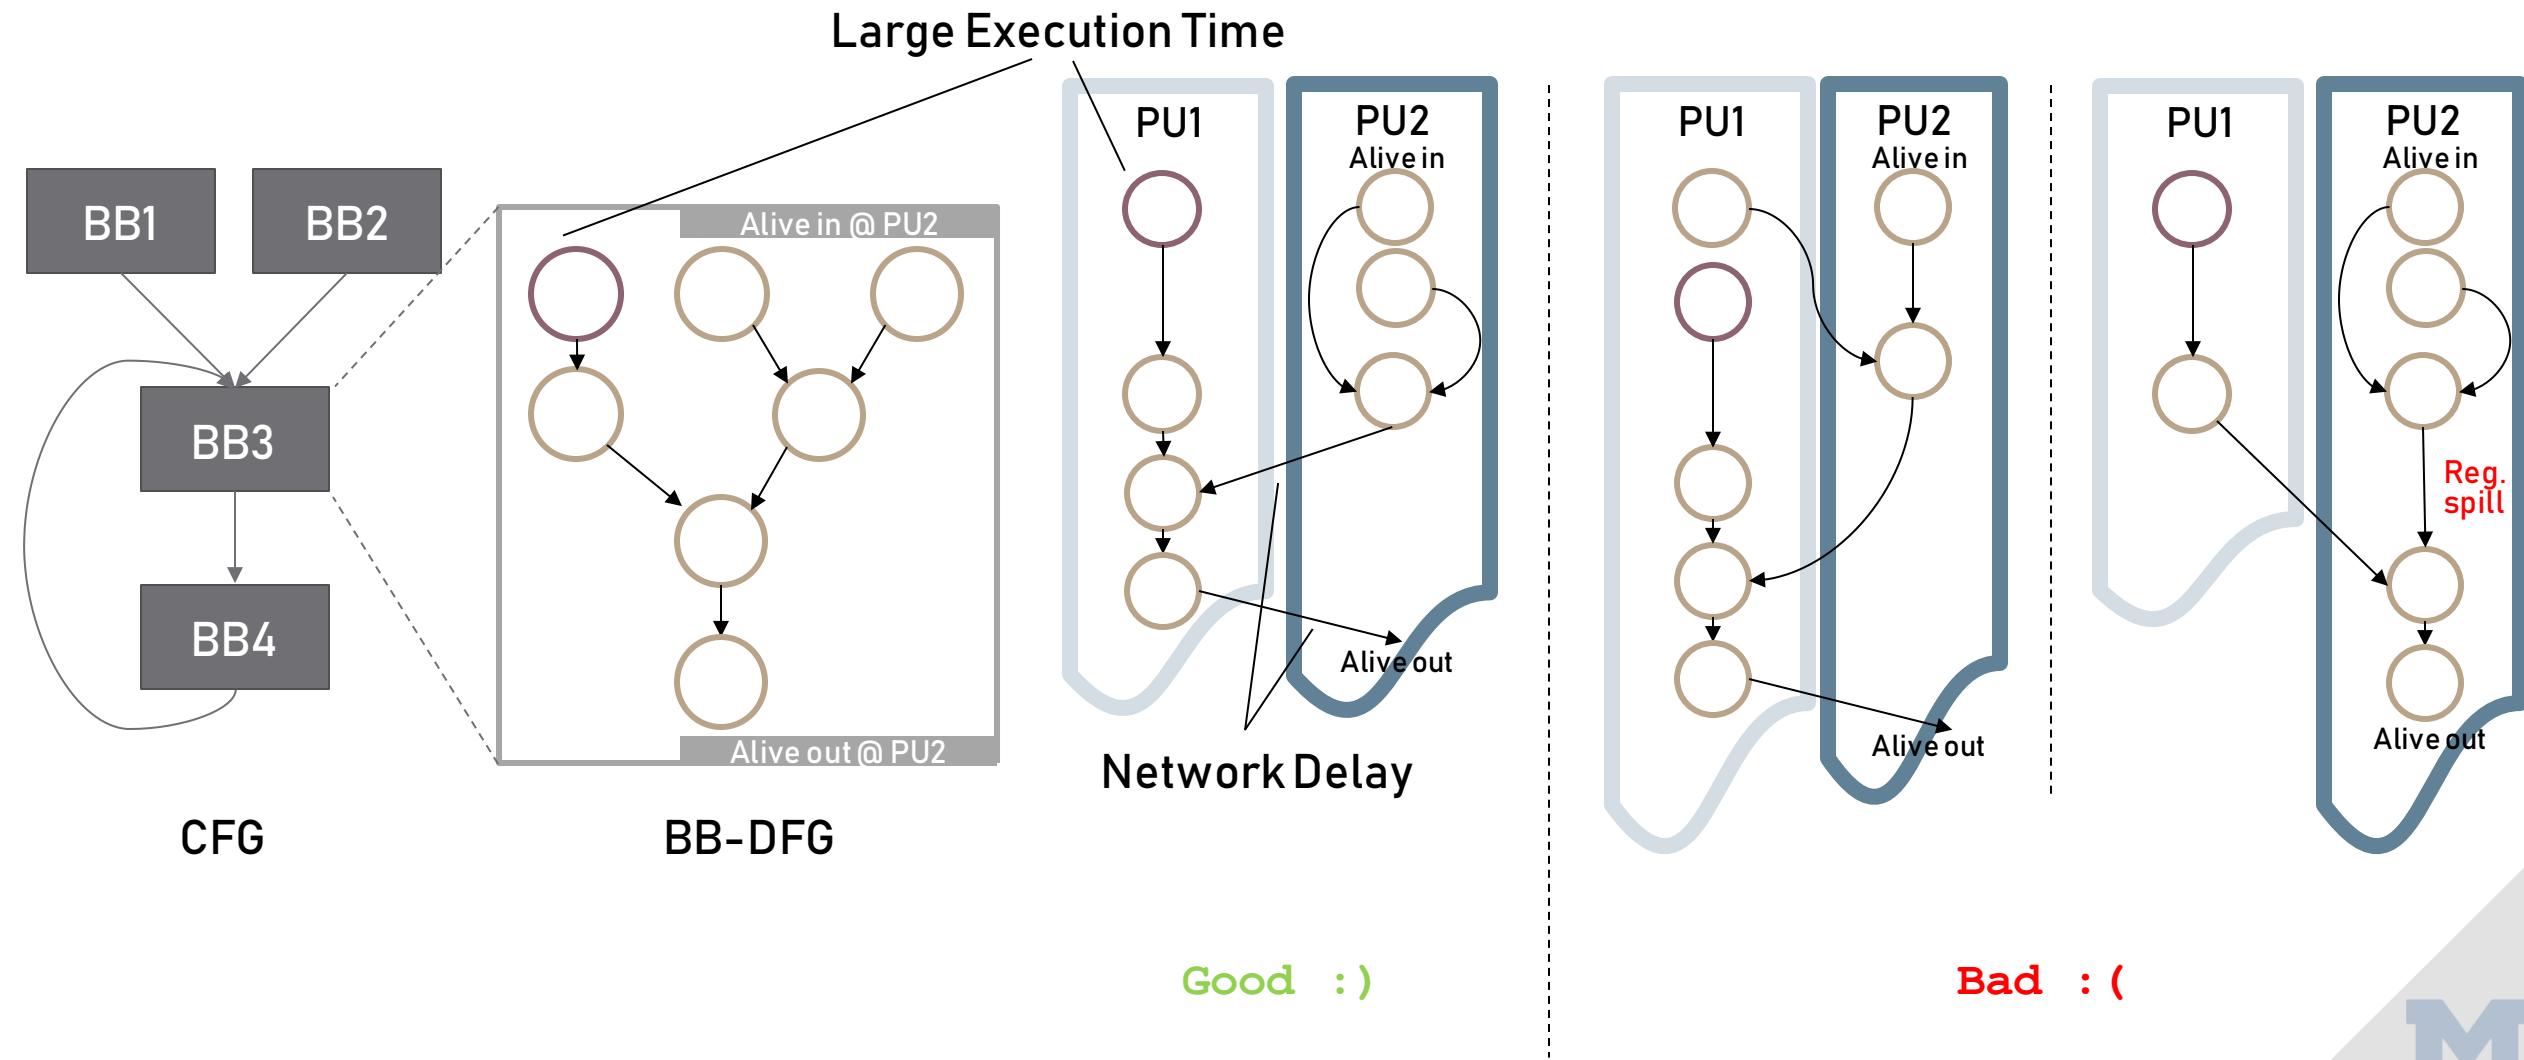
Resource
Allocator

Linear Scan Register
Allocation based

Interactively perform resource allocation and instruction scheduling.

Compiler Optimization

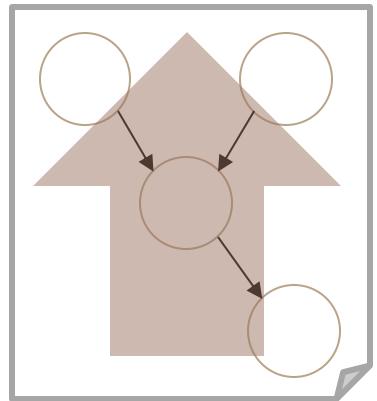
(1) Register Pressure Aware Instruction Scheduling



Compiler Optimization

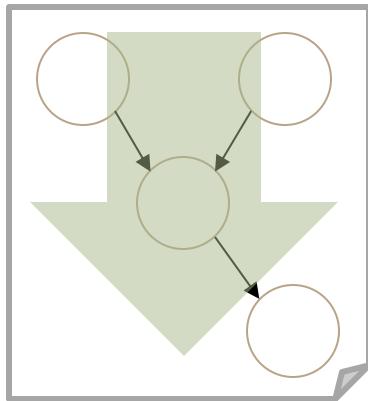
(1) Register Pressure Aware Instruction Scheduling

Bottom-Up Greedy

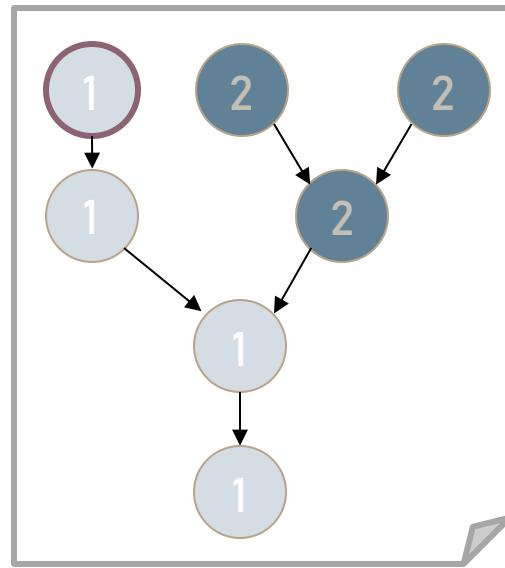


Collect candidate assignments

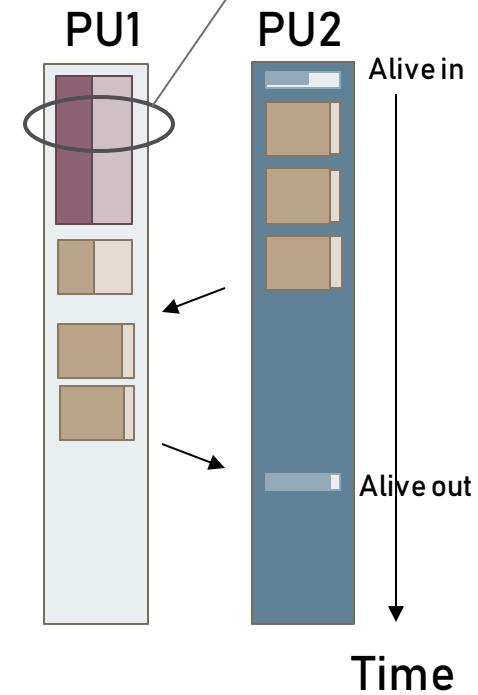
[Ellis 1986]



Make final assignments
Minimize data transfer latency by taking both operand & successor location into consideration

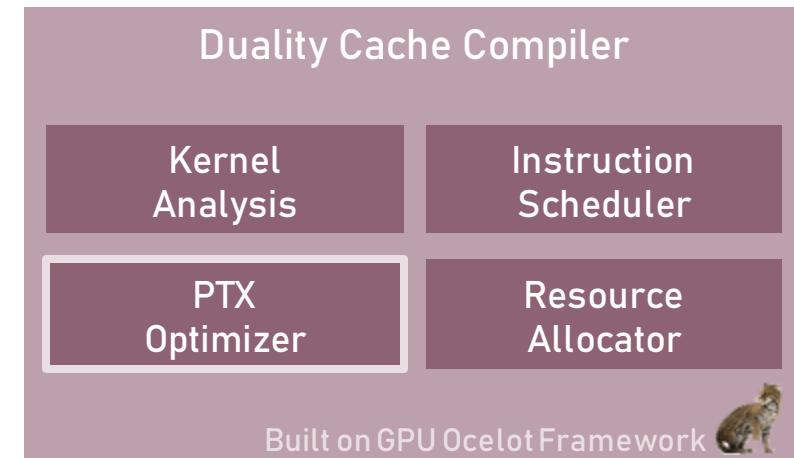


Register pressure

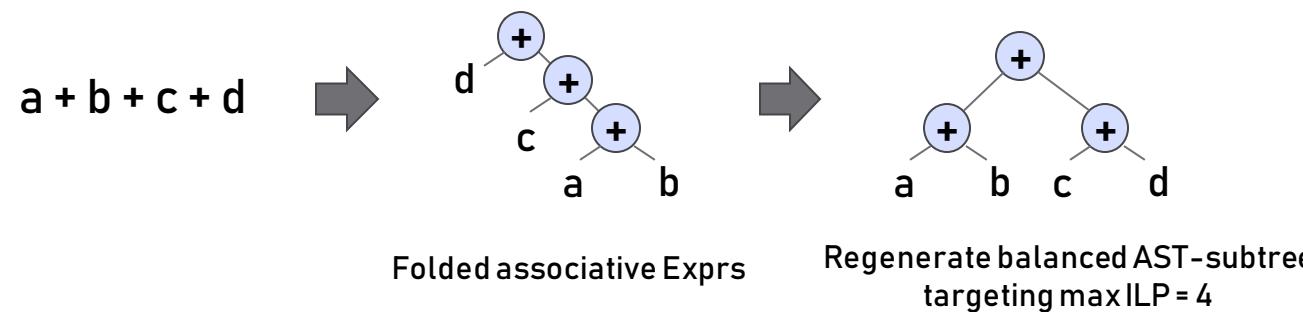


Compiler Optimization

(2) PTX Optimizations

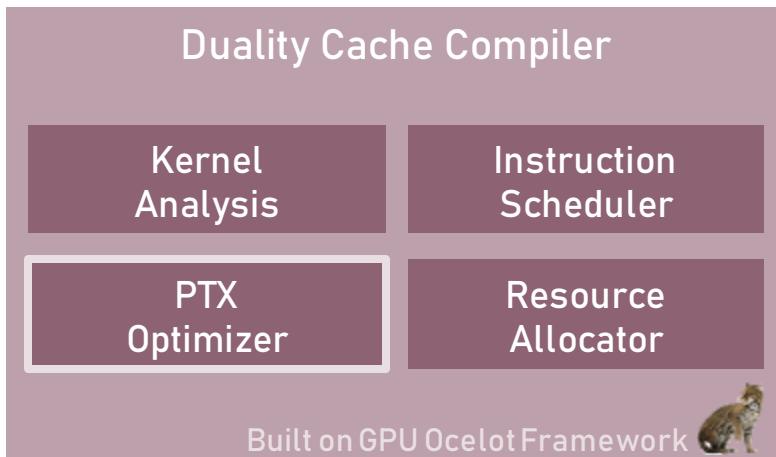


I. AST Balancing

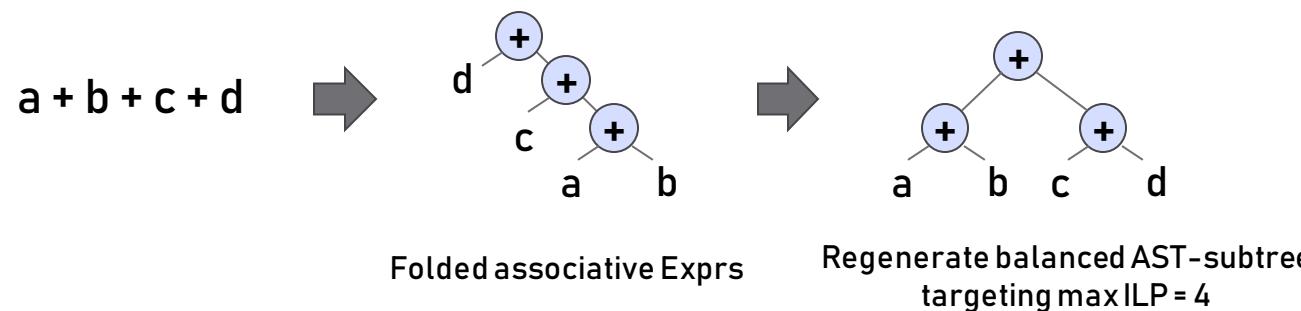


Compiler Optimization

(2) PTX Optimizations



I. AST Balancing



II. Thread Independent Variable Isolation

```
_device__kernel
{
    int bid = blockIdx.x;
    for (int i = 0; i < ITER_CNT; ++i)
    {
        //Do something
    }
}
```

bid and i are thread independent variable.

↓
Loop unroll / const fold.

Do not store them in thread private registers to reduce reg pressure.

Outline

Background

Integer / Logical Operations

Floating Point Operations

Transcendental Functions

Execution Model

Programming Model

Compiler

➡ Methodology/Results

Evaluation Methodology

◆ Benchmarks

- Rodinia
 - backprop
 - bfs
 - b+tree
 - dwt2d
 - hotspot
 - hotspot3D
 - hybridsort
 - nw
 - streamcluster
 - gaussian
 - heartwall
 - leukocyte
 - lud
 - nn
- PathScale OpenACC benchmark
 - divergence
 - gradient
 - lapsgrb
 - laplacian
 - tricubic
 - tricubic2
 - uxx1
 - vecadd
 - wave13pt
 - gameoflife
 - gaussblur
 - matvec
 - whispering

	CPU 2-sockets	GPU 1-card, PCIe v3	Duality Cache
Processor	Intel Xeon E5-2597 v3, 2.6GHz, 28 cores, 56 threads	NVIDIA GP100, 1.6GHz, 3840 cuda cores	35MB Duality Cache, 2.6 GHz
On-chip Memory	78.96 MB	9.14 MB	78.96 MB
Off-chip Memory	64 GB DDR4	12GB GDDR5 + 64GB DDR4 (host)	64 GB DDR4
Total System Area	912 mm ²	1383 mm ²	942 mm ²
TDP	290 W	640 W	296 W
Profiler / Simulator (Performance)	perf	NVPROF	GPU Ocelot + Ramulator
Profiler / Simulator (Energy)	Intel RAPL Interface	NVIDIA System Management Interface	Trace based simulation

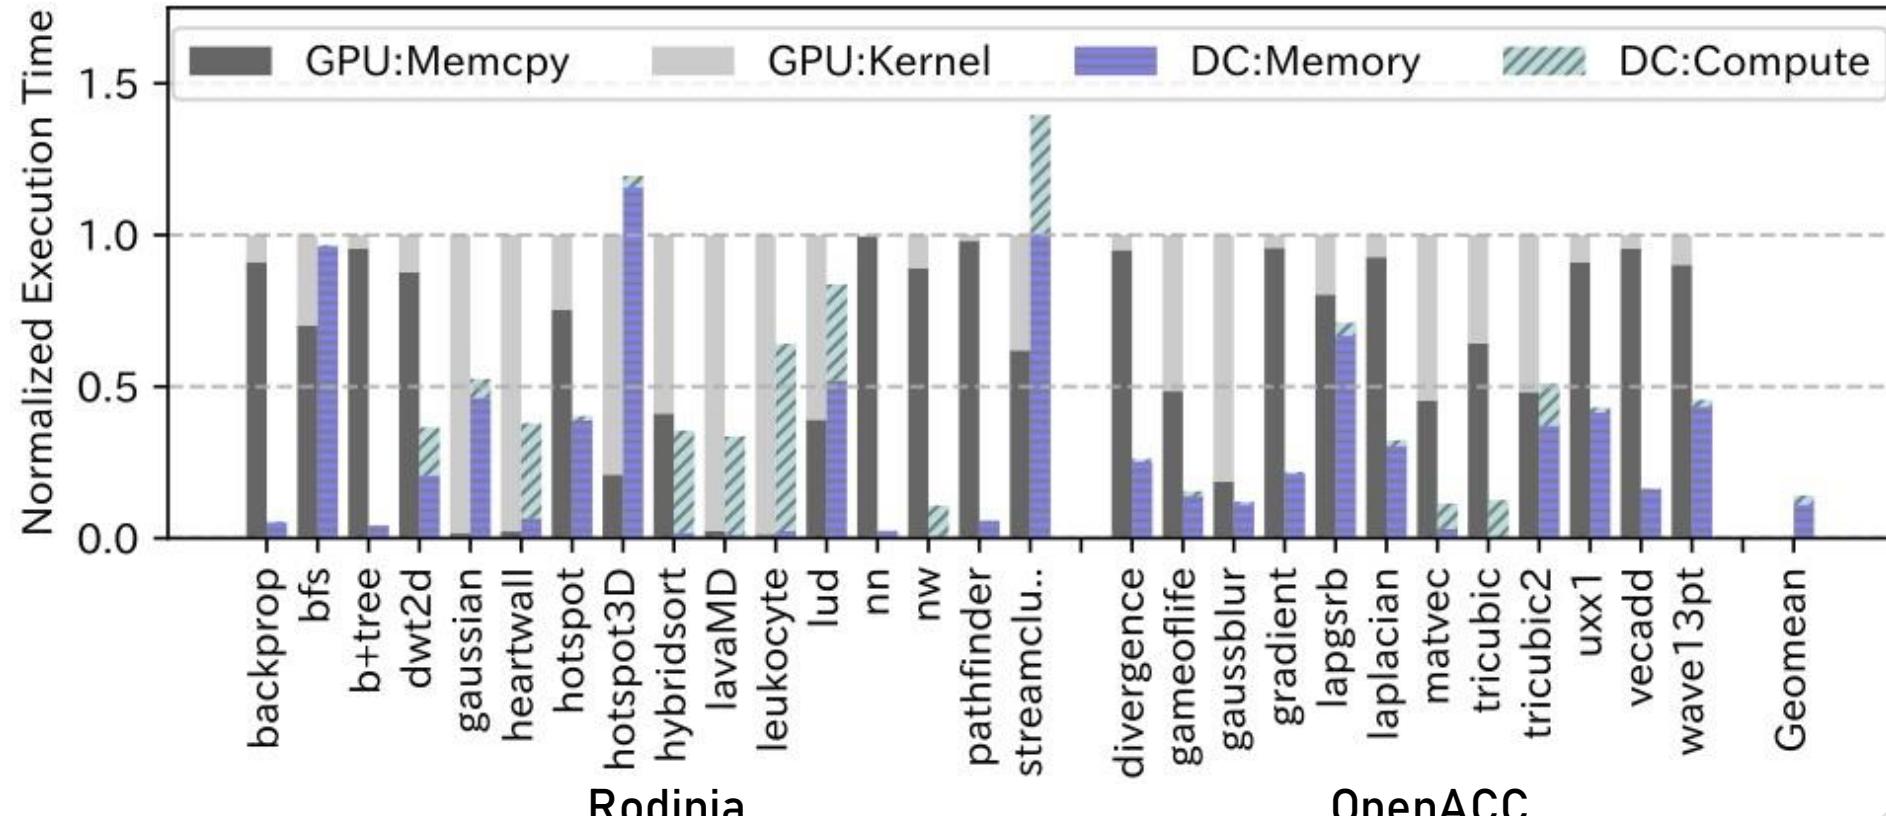
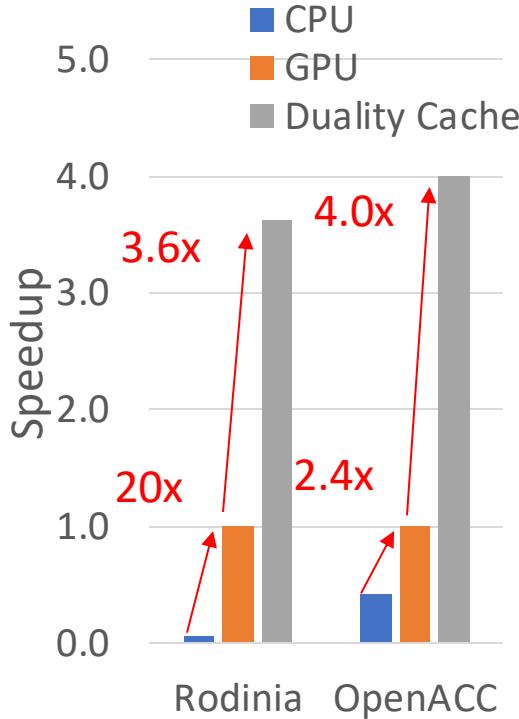
Area / Power Overhead

	Area (mm ²)	Power (W)	Area Overhead
CPU	456	145	-
Compute Cache Peripheral	3.15	2.96	0.69%
TMU	5.32	0.06	1.17%
Controller / FSM	6.16	0.33	1.35%
MSHR	0.86	0.05	0.19%
Local Crossbar	0.28	0.01	0.06%
Total	471.77	148.4	3.50%

Duality Cache has only 3.5% area overhead

System Speedup

CPU, DC: DDR4, GPU: GDDR5 + memcpy (DDR4 \leftrightarrow PCIe v3 \leftrightarrow GDDR5)



Key performance factors

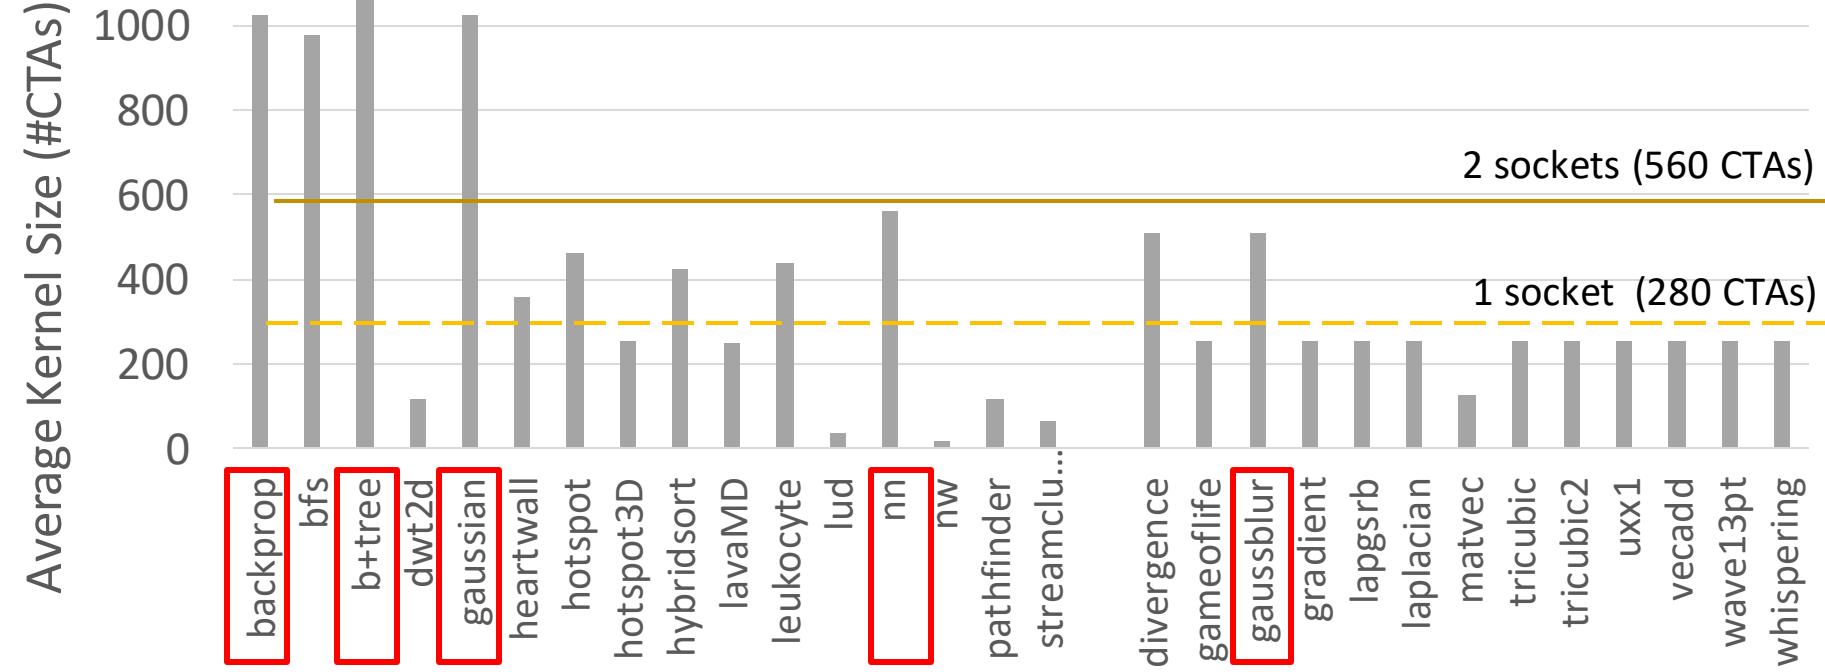
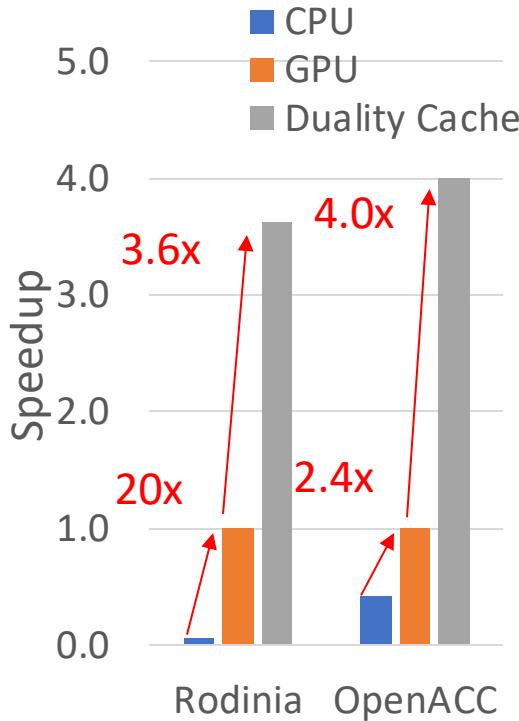
- Reduced Memcpy time
- Massively parallel execution
- Compute / memory access overlap
- Flexible cache allocation

Reduced data transfer (memcpy) cost via external bus (PCIe).

Small reuse factor / large working set make memcpy costly for GPU.

Massively Parallel Execution

PU Utilization



Key performance factors

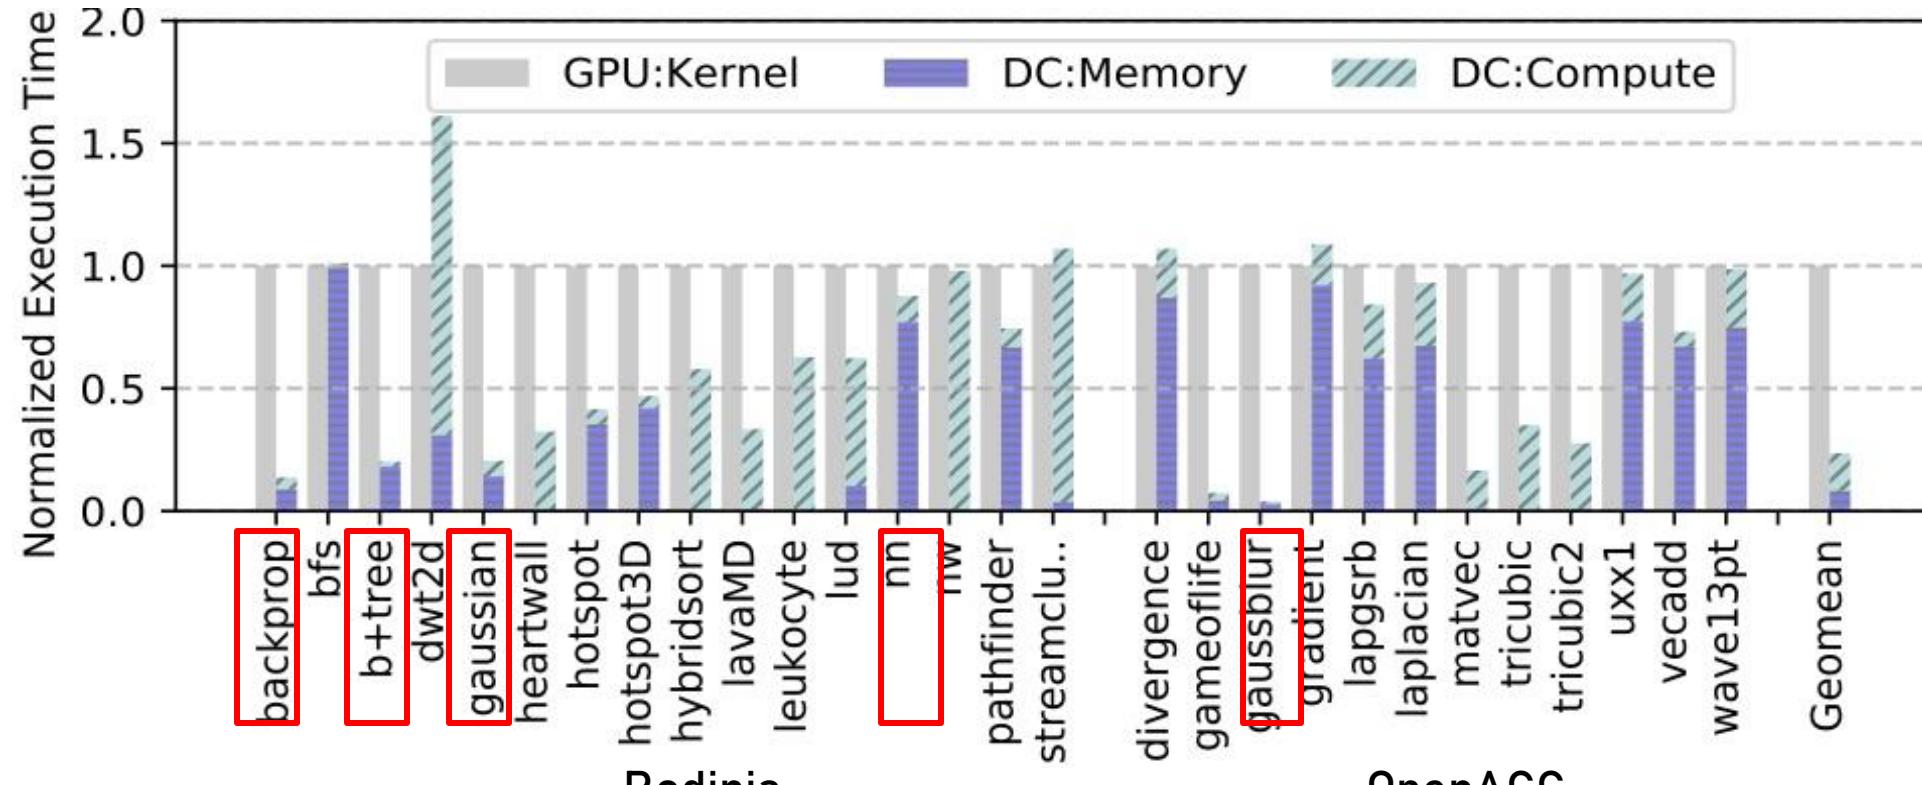
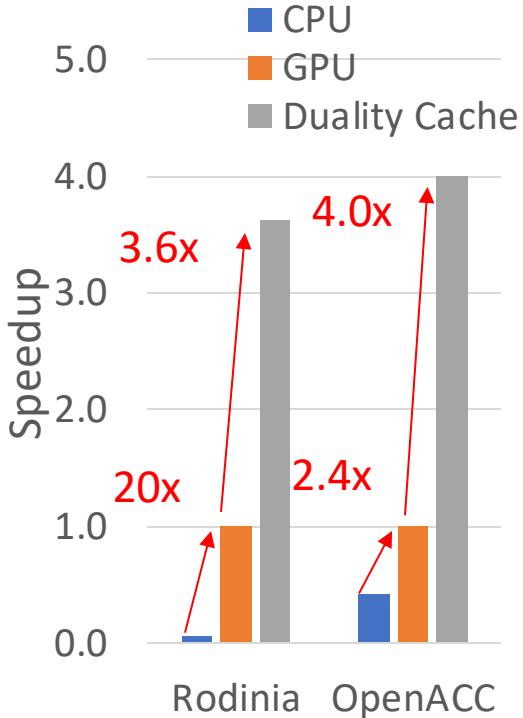
- Reduced Memcpy time
- **Massively parallel execution**
- Compute / memory access overlap
- Flexible cache allocation

Some kernels have high level of parallelism (e.g. backprop, b+tree, nn, gaussian, gausblur, etc.)

280 CTA/ socket (DC) vs. 60 CTA / card (GPU) (reg. capacity based)

Kernel Speedup

DC: **GDDR5**, GPU: GDDR5 + **No memcpy**



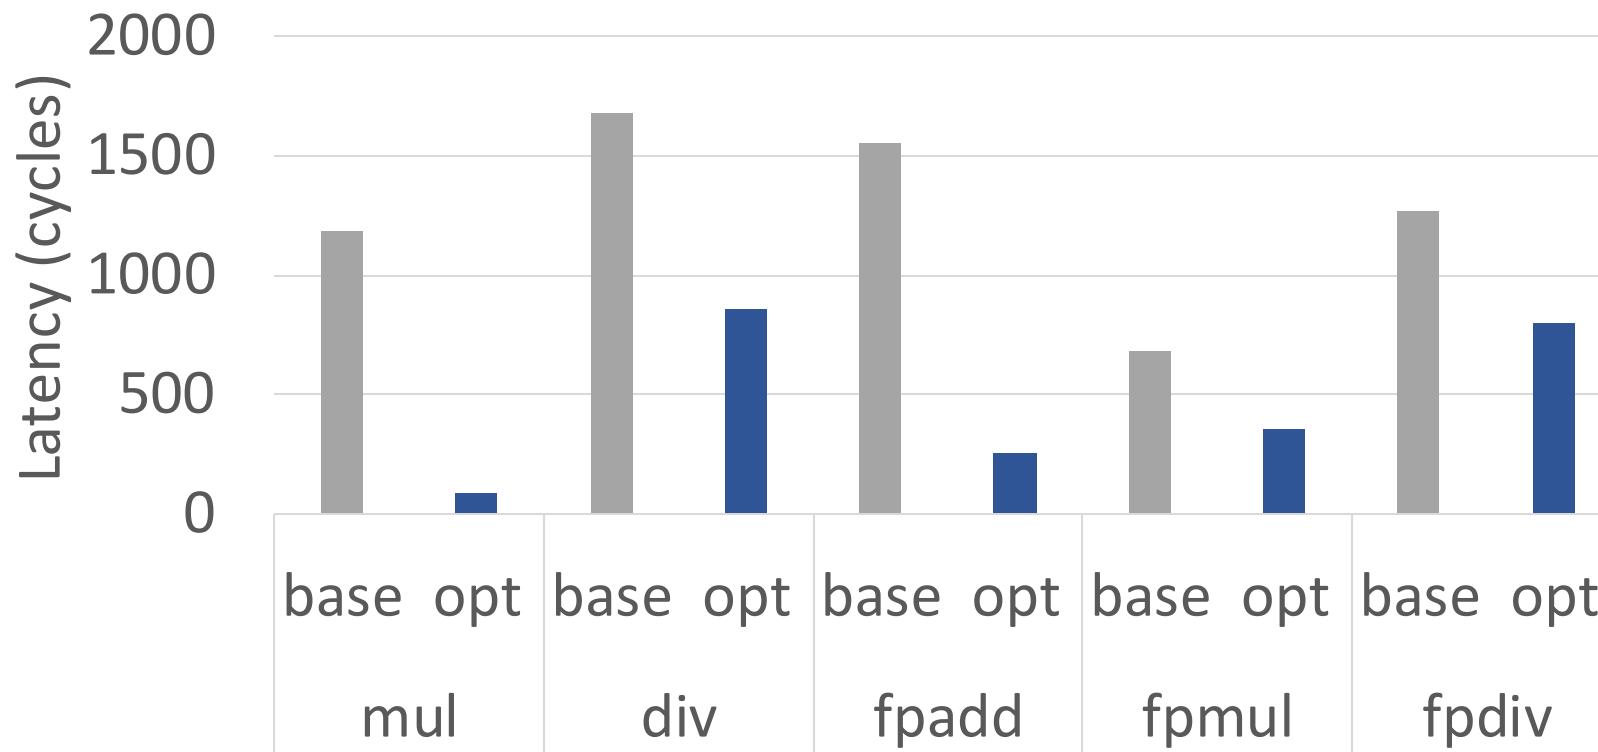
Key performance factors

- Reduced Memcpy time
- **Massively parallel execution**
- Compute / memory access overlap
- Flexible cache allocation

Some kernels with high level of parallelism (e.g. backprop, b+tree, nn, gaussian, gaussblur, etc.) **significantly reduce execution time**.

Memory bounded applications (hotspot3d, streamcluster) will benefit from GDDR5

Operation Latency



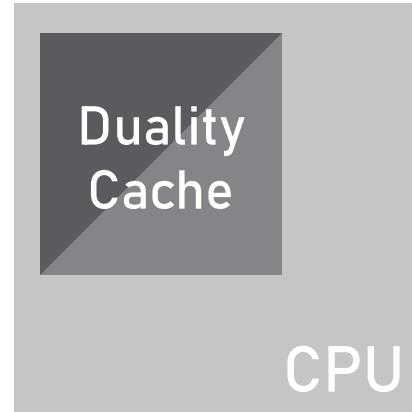
Integer multiplication is often used to calculate address or some variables based on induction variables as an operand

→ They contain many **leading zeros** which we can skip by our optimization. (13x faster)

Floating point addition has **small dynamic range**.

→ The number of **unique ediff** found usually has its peak at 1 in the distribution. (6.1x faster)

Conclusion



Contributions

In-cache computing framework for general purpose programming

- Used SIMD programming model for the programming frontend
- Developed a compiler for in-cache computing
- Enhanced computation primitives (all digital)

Results

Overall Efficiency



**1450x over CPU
52x over GPU**

Performance



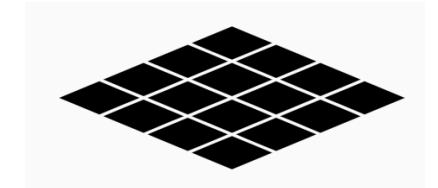
**72x over CPU
4x over GPU**

Energy



**20x over CPU
5.8x over GPU**

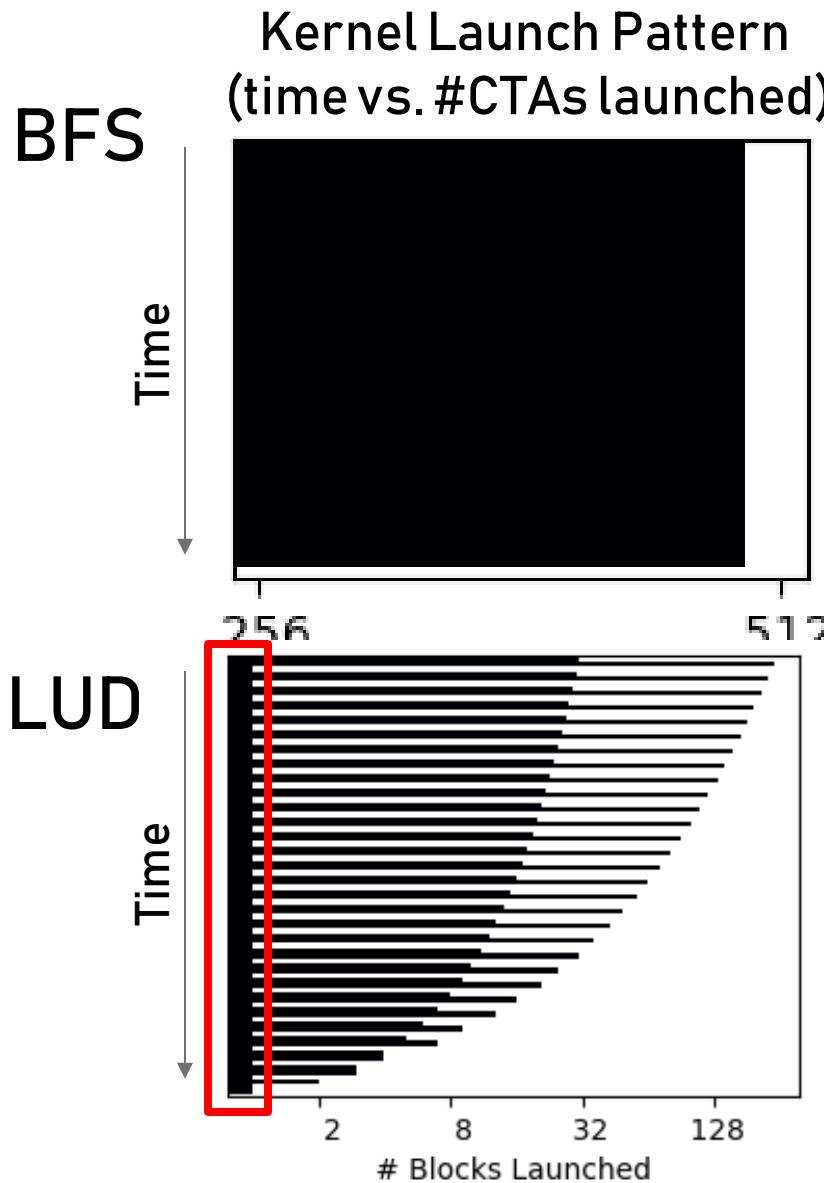
Area



**Duality Cache needs 15.7 mm²
== 3.5% over CPU
GPU = 471 mm²**

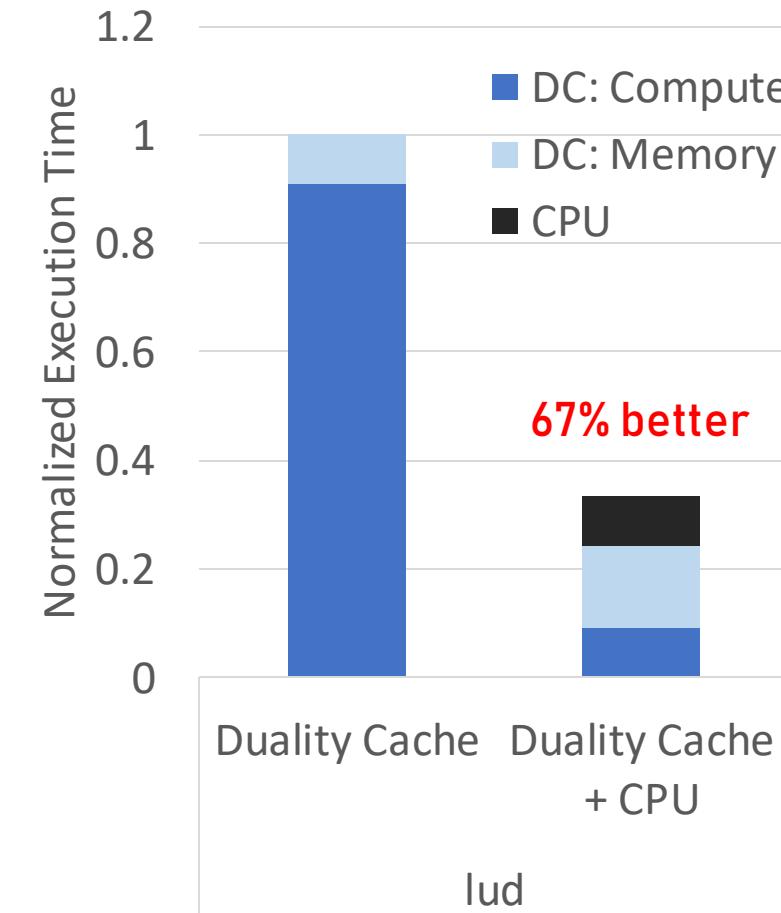
Backup Slides

DC + CPU Hybrid Execution

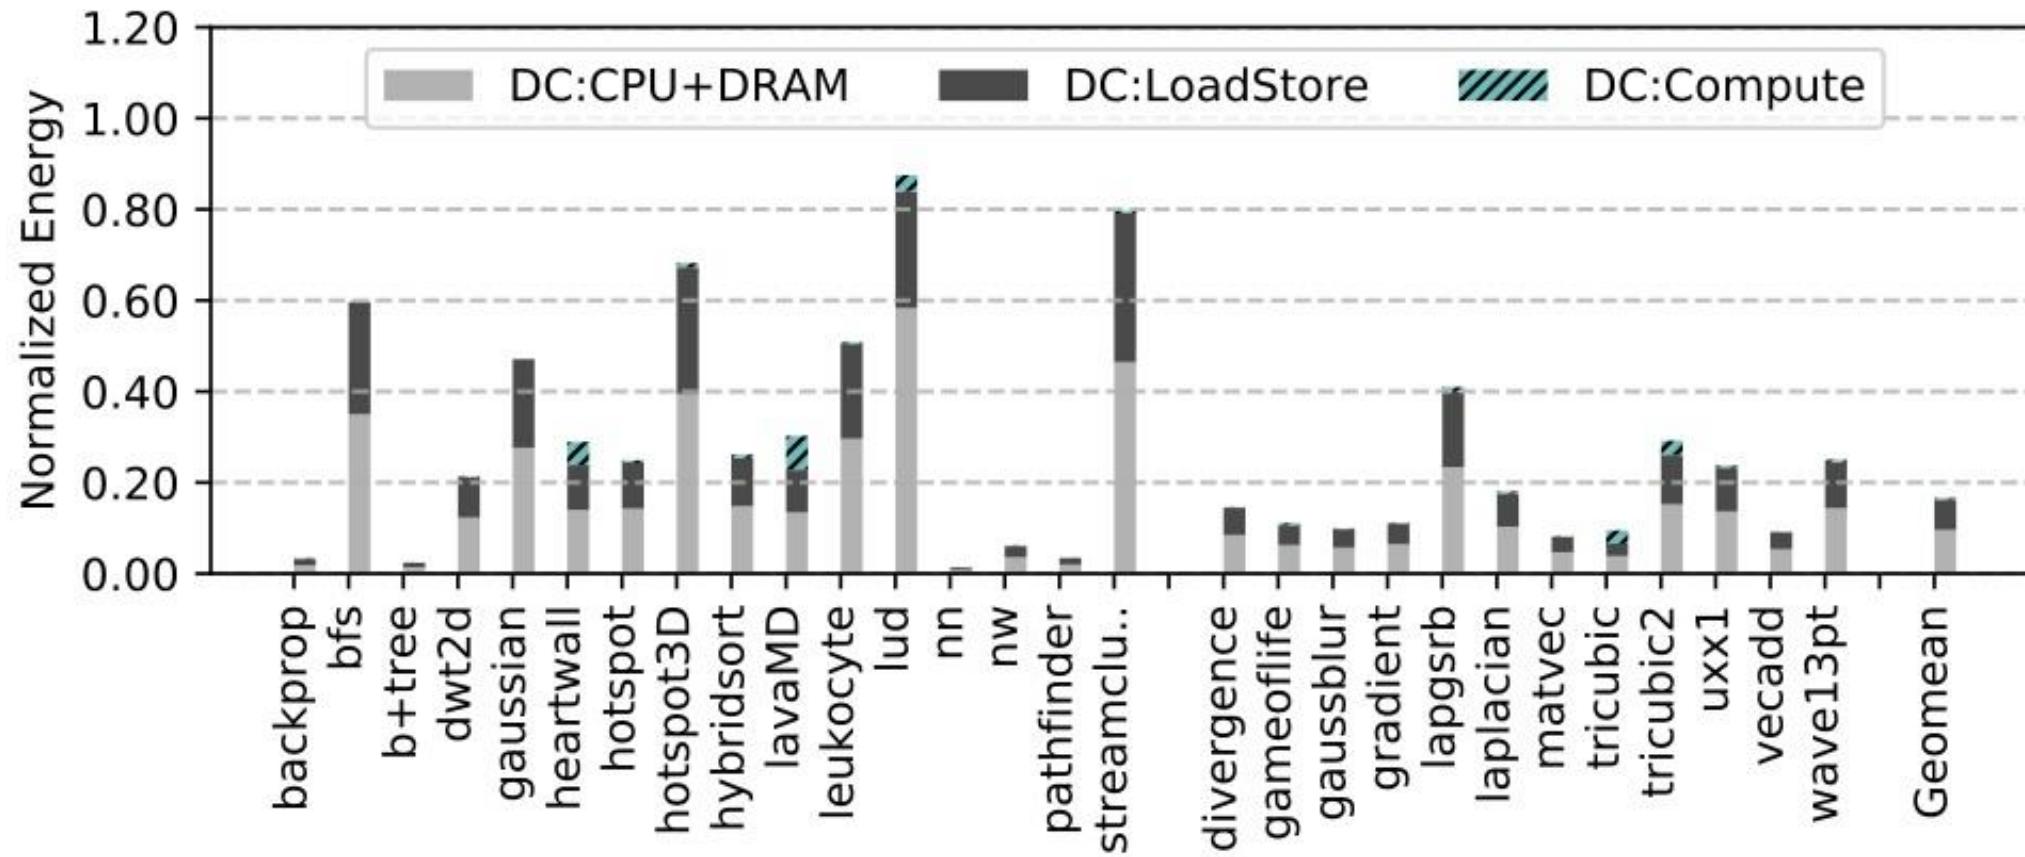


Enough parallelism to keep all PUs busy

Execute small kernels on CPU threads

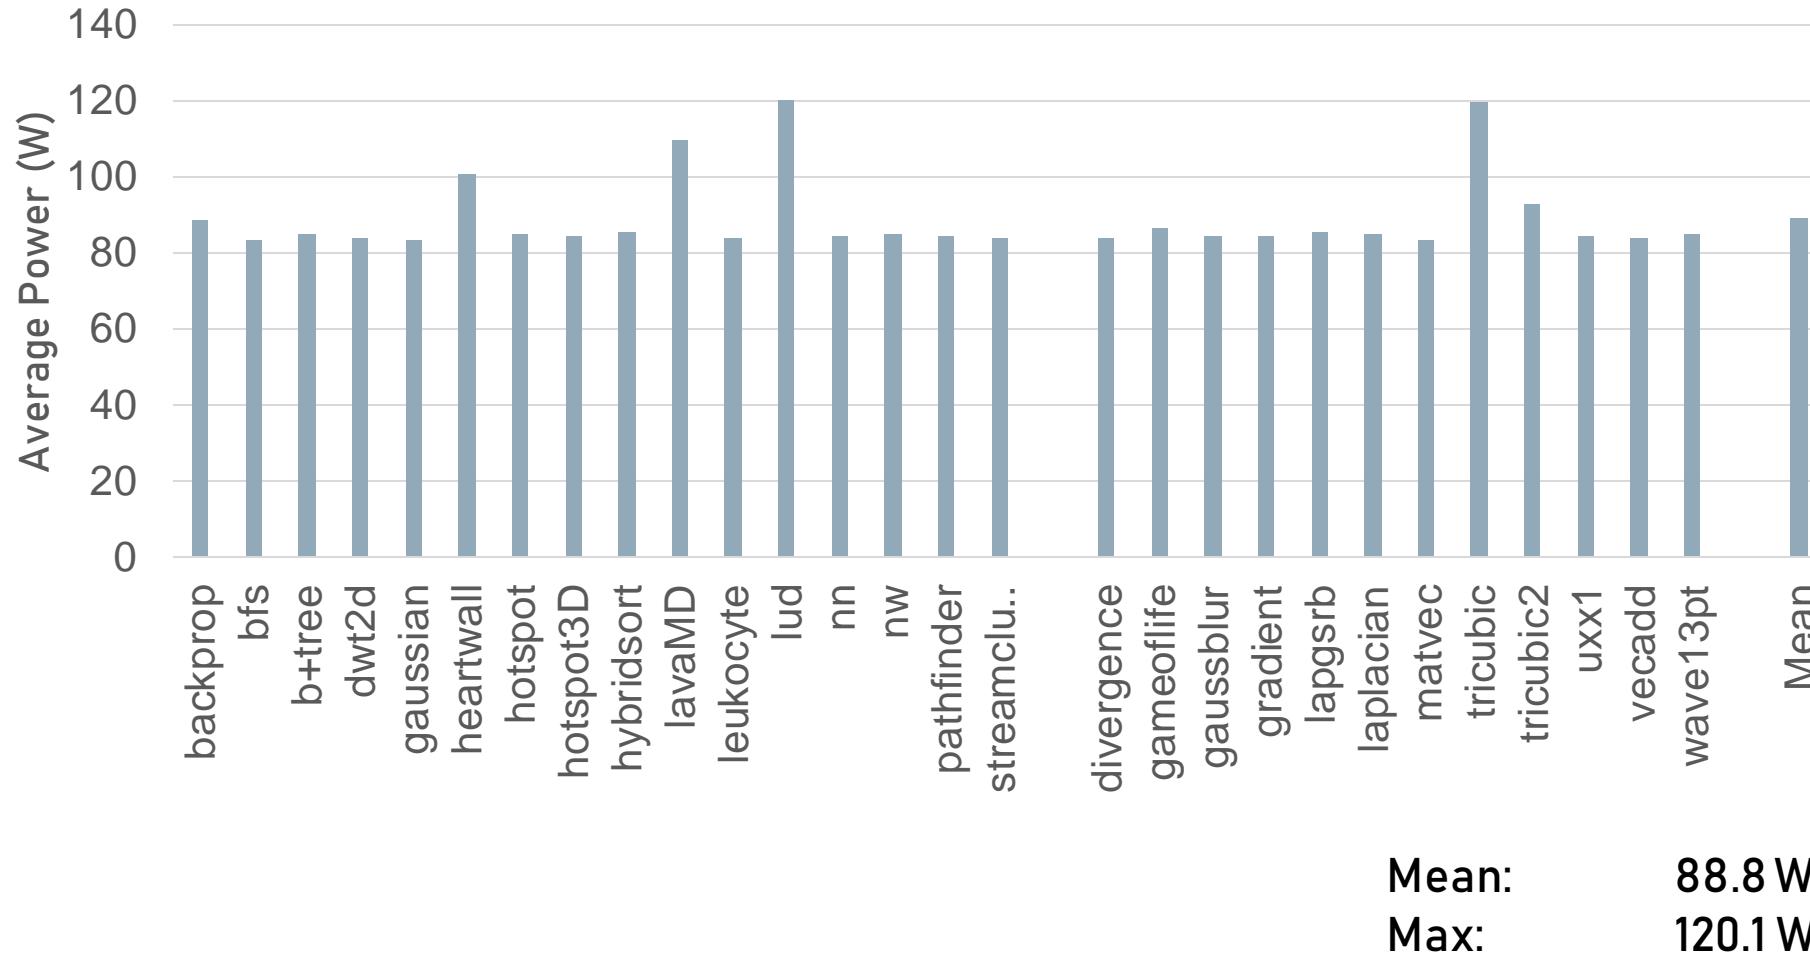


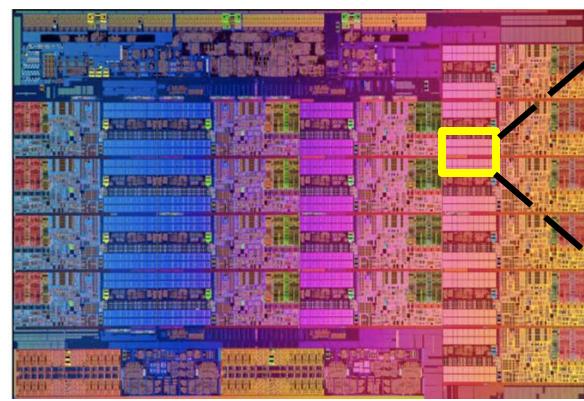
System Energy



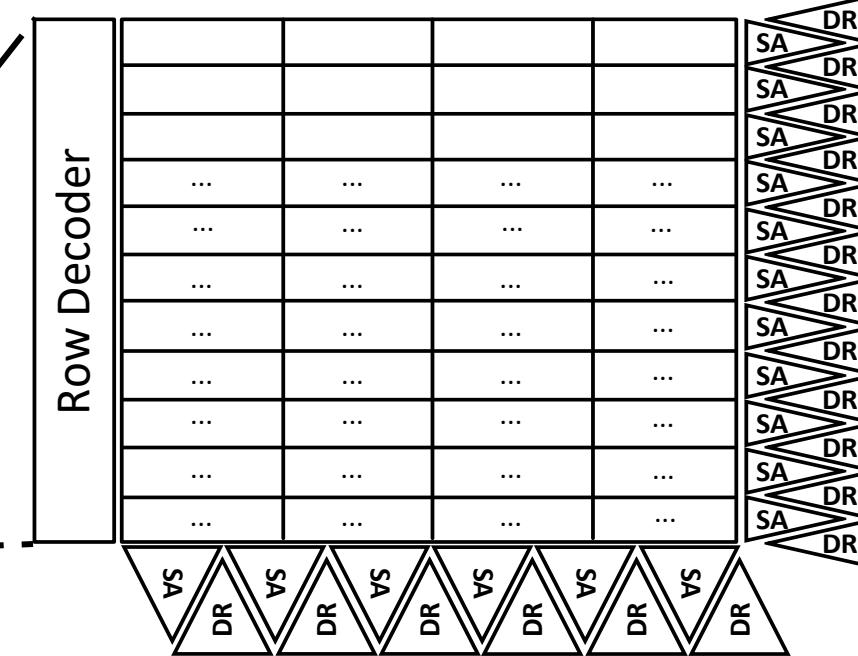
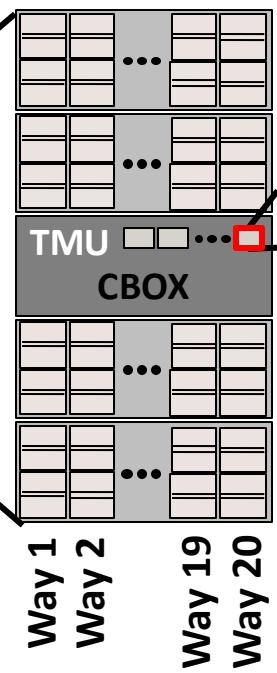
Because of the reduced execution time, we achieve 5.34× energy efficiency compared to GPU. One core is active during execution to serve instructions. This makes CPU and DRAM access dominant in energy consumption.

Average Power

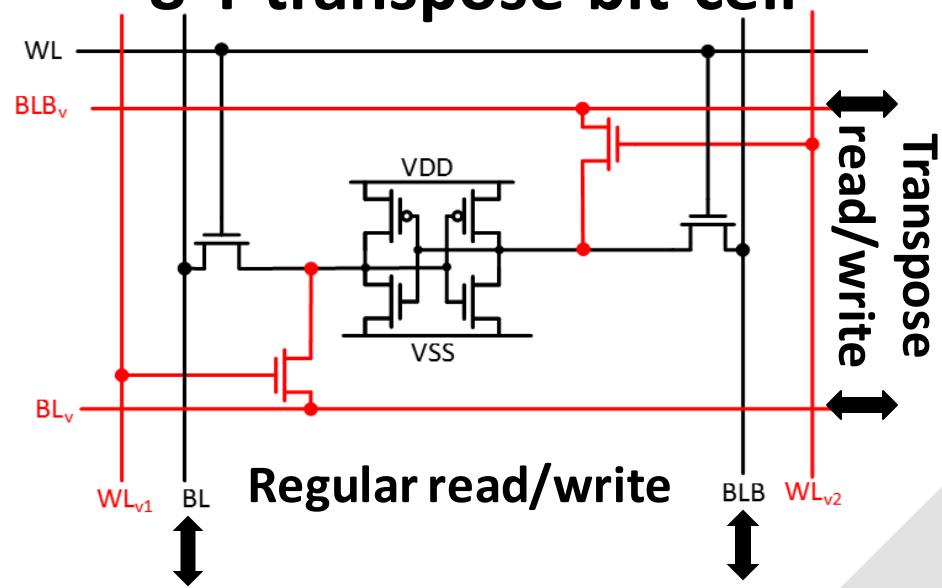




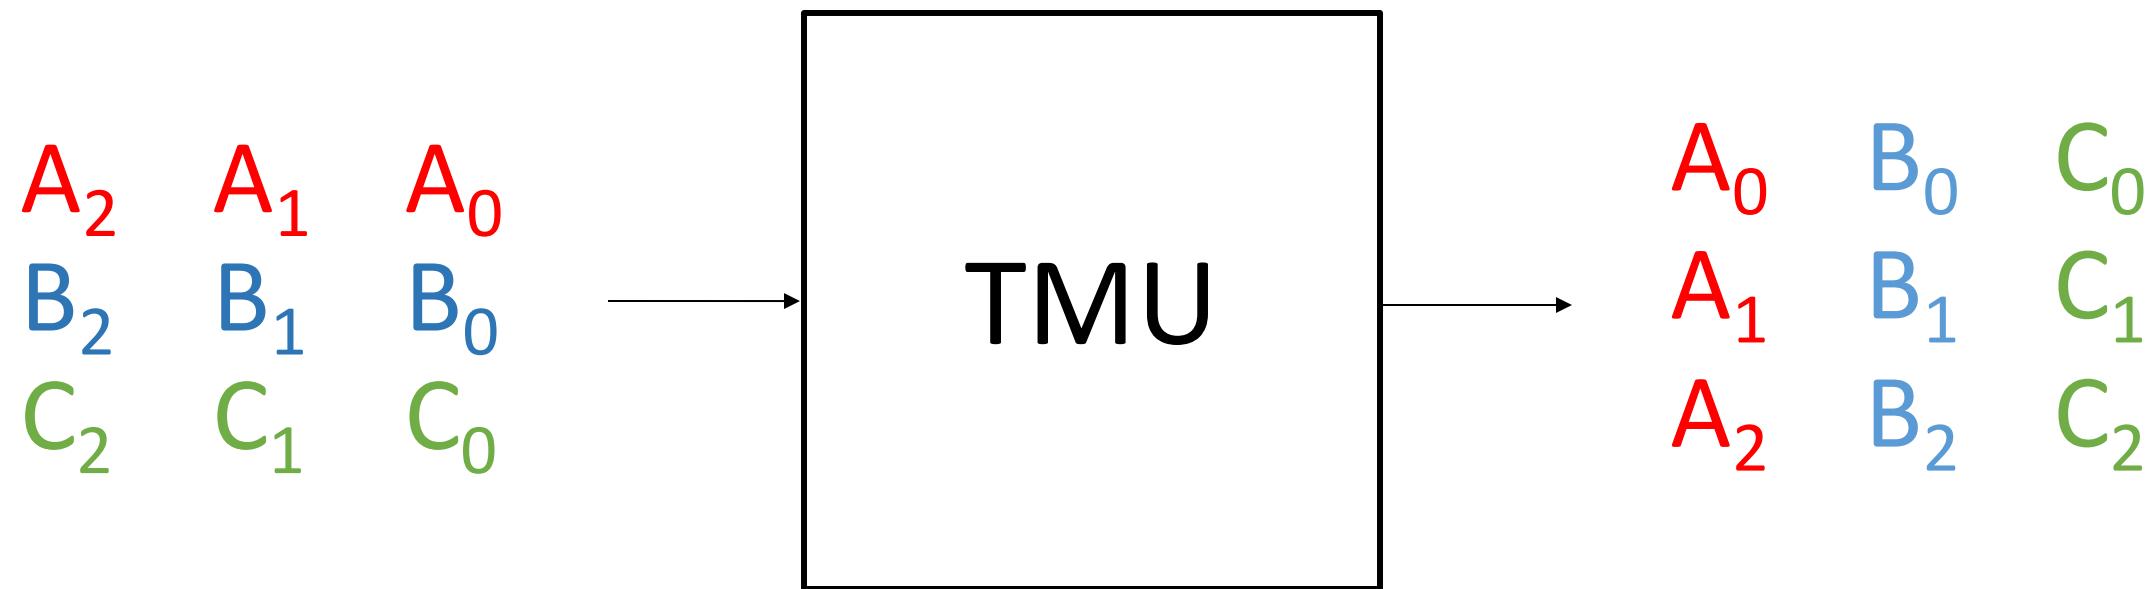
Transpose



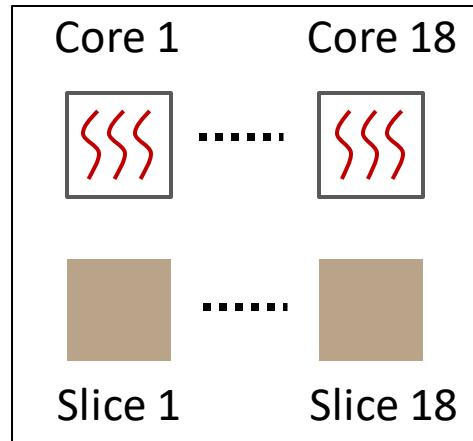
8-T transpose bit-cell



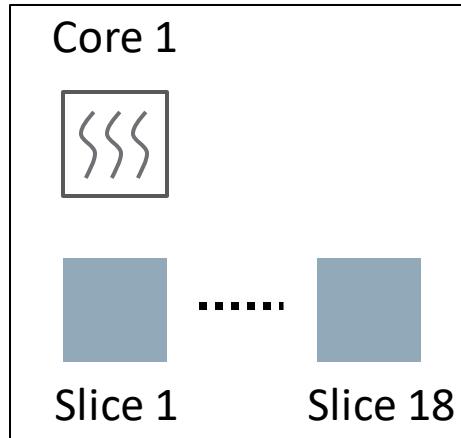
Transpose



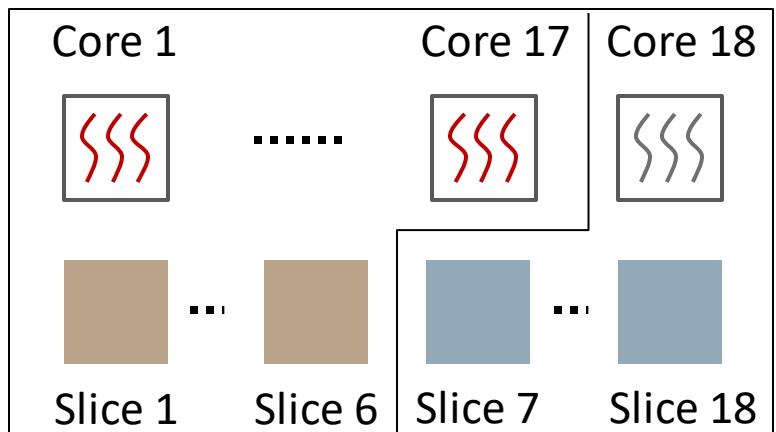
Duality Cache Operation Modes



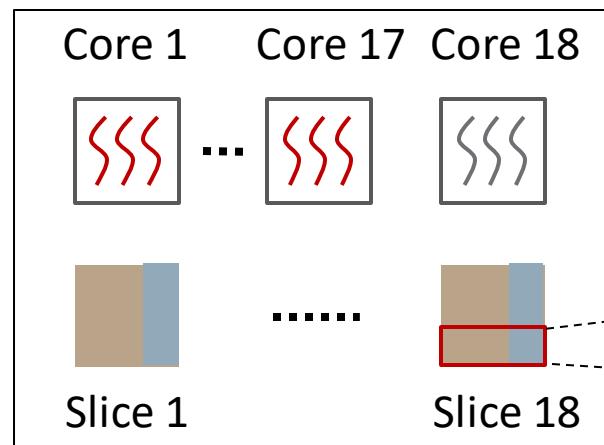
CPU only mode



Duality Cache only mode



Hybrid mode - Slice partitioned



Hybrid mode - Way partitioned



Cores running
non-DC threads



Cores running
DC-related threads



LLC used for
non-DC threads



LLC used for
NC-related threads