

P&S Heterogeneous Systems

GPU Performance Considerations

Dr. Juan Gómez Luna

Prof. Onur Mutlu

ETH Zürich

Fall 2022

31 October 2022

GPU Memories

Traditional Program Structure

- CPU threads and GPU kernels
 - Sequential or modestly parallel sections on CPU
 - Massively parallel sections on GPU

Serial Code (host)


Parallel Kernel (device)

```
KernelA<<< nBlk, nThr >>>(args);
```

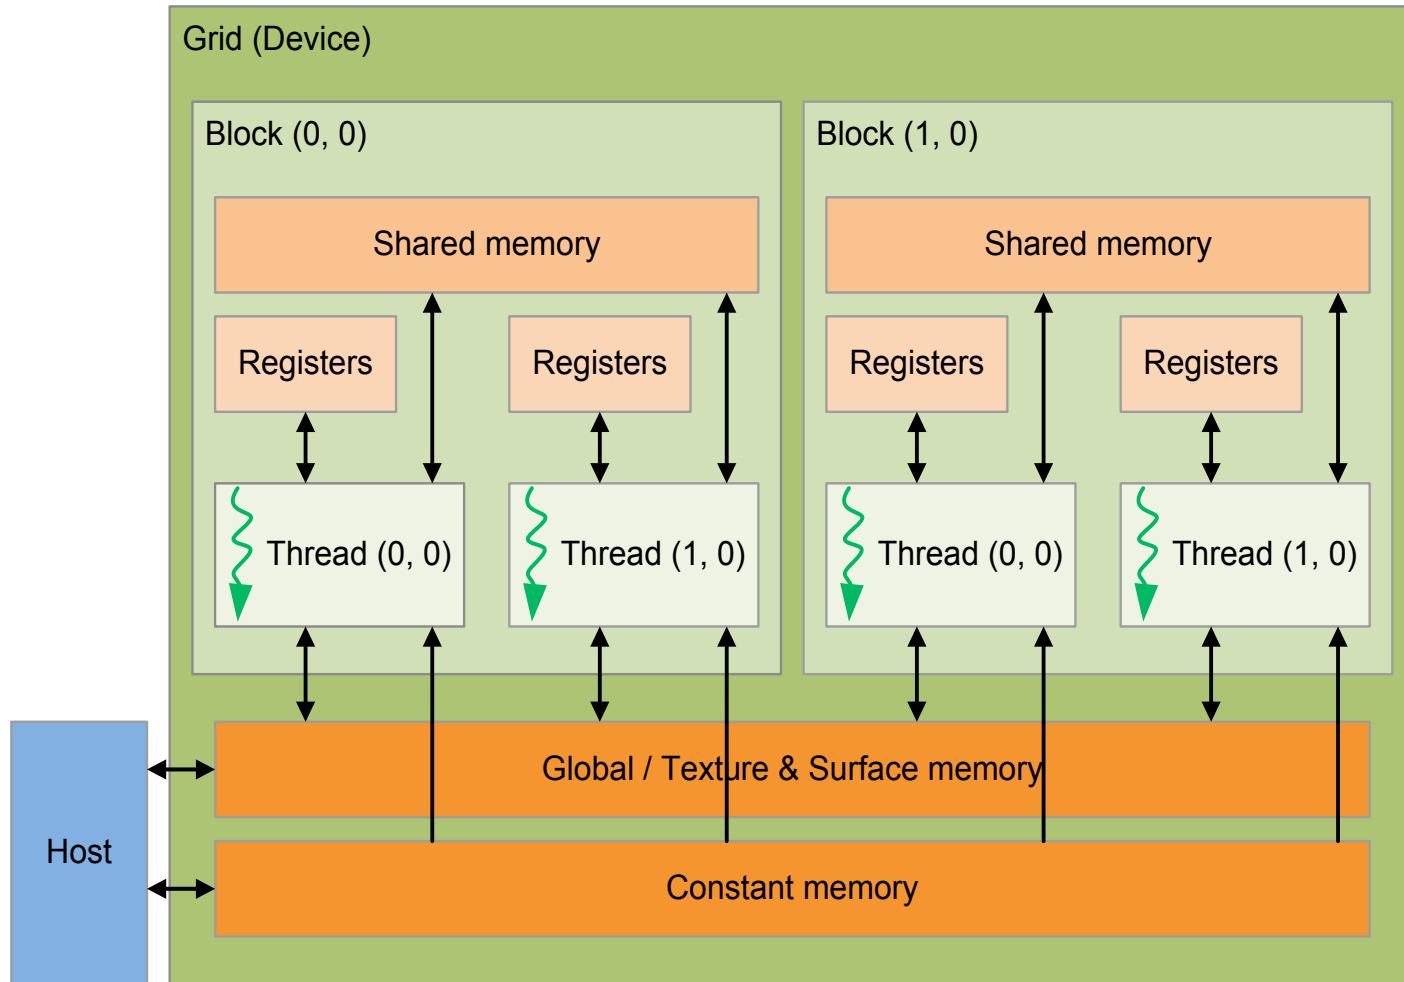
Serial Code (host)

Parallel Kernel (device)

```
KernelB<<< nBlk, nThr >>>(args);
```



Memory Hierarchy in CUDA Programs



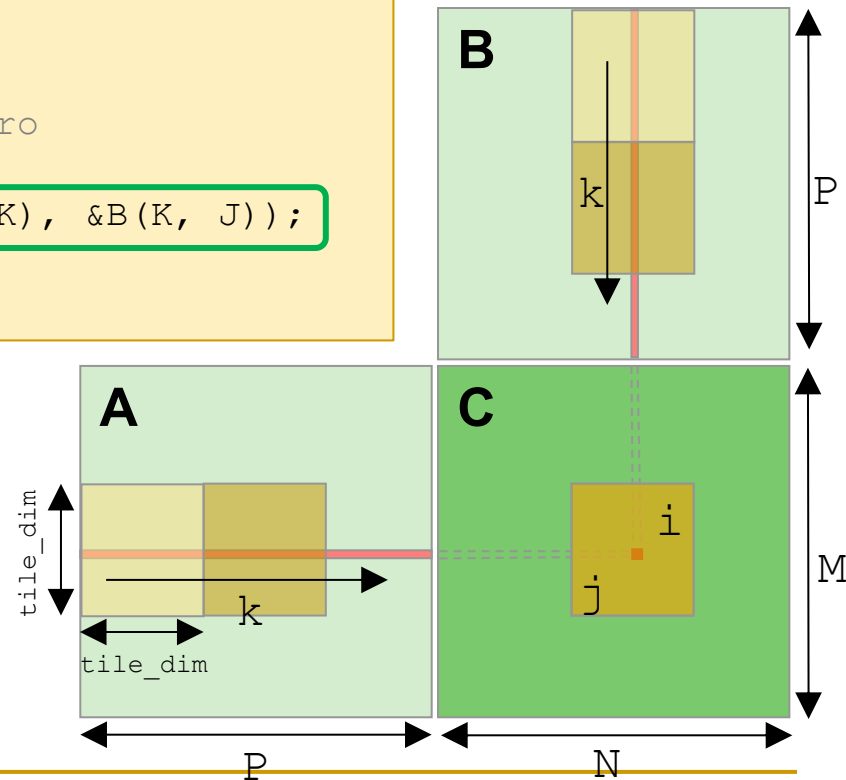
Tiled Matrix Multiplication (II)

- **Tiled implementation** operates on submatrices (tiles or blocks) that fit fast memories (cache, scratchpad, RF)

```
#define A(i,j) matrix_A[i * P + j]
#define B(i,j) matrix_B[i * N + j]
#define C(i,j) matrix_C[i * N + j]

for (I = 0; I < M; I += tile_dim){
    for (J = 0; J < N; J += tile_dim){
        Set_to_zero(&C(I, J)); // Set to zero
        for (K = 0; K < P; K += tile_dim)
            Multiply_tiles(&C(I, J), &A(I, K), &B(K, J));
    }
}
```

Multiply small submatrices (tiles or blocks) of size `tile_dim` x `tile_dim`



Tiled Matrix-Matrix Multiplication (V)

```
__shared__ float A_s[TILE_DIM][TILE_DIM];  
__shared__ float B_s[TILE_DIM][TILE_DIM];
```

 ———— Declare arrays in shared memory

```
unsigned int row = blockIdx.y*blockDim.y + threadIdx.y;  
unsigned int col = blockIdx.x*blockDim.x + threadIdx.x;
```

```
float sum = 0.0f;
```

```
for(unsigned int tile = 0; tile < N/TILE_DIM; ++tile) {
```

```
    // Load tile to shared memory
```

```
    A_s[threadIdx.y][threadIdx.x] = A[row*N + tile*TILE_DIM + threadIdx.x];  
    B_s[threadIdx.y][threadIdx.x] = B[(tile*TILE_DIM + threadIdx.y)*N + col];  
    __syncthreads();
```

Threads wait for each other to finish loading before computing

```
    // Compute with tile
```

```
    for(unsigned int i = 0; i < TILE_DIM; ++i) {  
        sum += A_s[threadIdx.y][i]*B_s[i][threadIdx.x];  
    }
```

```
    __syncthreads();
```

 Threads wait for each other to finish computing before loading

```
}
```

```
C[row*N + col] = sum;
```

Performance Considerations


Performance Considerations

- Main bottlenecks
 - ❑ CPU-GPU data transfers
 - ❑ Global memory access
- Memory access
 - ❑ Latency hiding
 - Occupancy
 - ❑ Memory coalescing
 - ❑ Data reuse
 - Shared memory usage
- SIMD (Warp) Utilization: Divergence
- Other considerations
 - ❑ Atomic operations: Serialization
 - ❑ Data transfers between CPU and GPU
 - Overlap of communication and computation

Memory Access


Latency Hiding via Warp-Level FGMT

- Warp: A set of threads that execute the same instruction (on different data elements)
- Fine-grained multithreading
 - One instruction per thread in pipeline at a time (No interlocking)
 - Interleave warp execution to hide latencies
- Register values of all threads stay in register file
- FGMT enables long latency tolerance
 - Millions of pixels



Latency Hiding and Occupancy


- FGMT can hide **long latency operations** (e.g., memory accesses)
- **Occupancy**: ratio of **active warps** to the maximum number of warps per GPU core



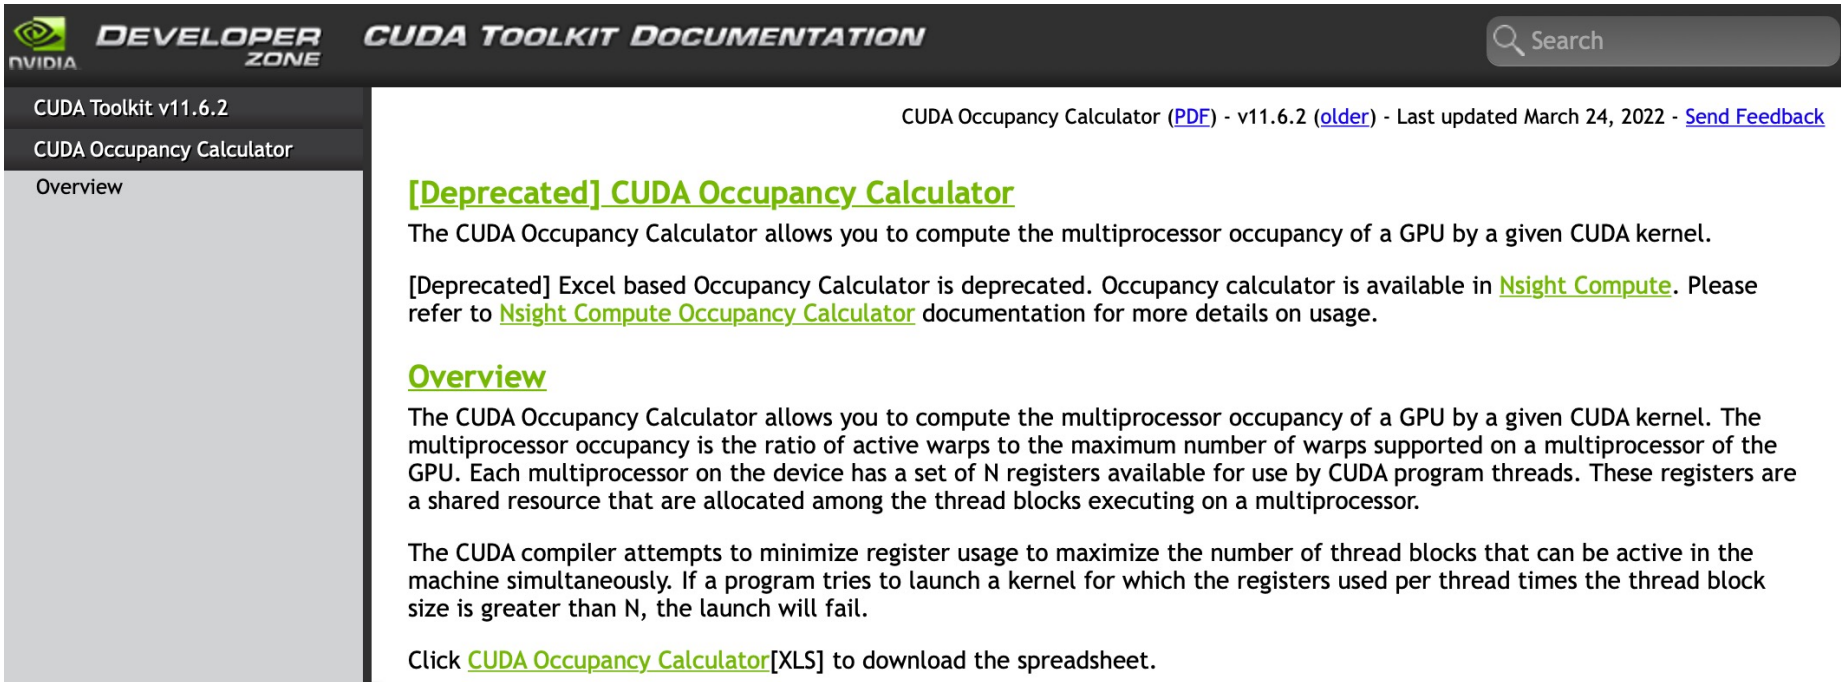
Occupancy

- GPU core, a.k.a. SM, resources (typical values)
 - ❑ Maximum number of warps per SM (64)
 - ❑ Maximum number of blocks per SM (32)
 - ❑ Register usage (256KB)
 - ❑ Shared memory usage (64KB)
- Occupancy calculation
 - ❑ Number of threads per block (defined by the programmer)
 - ❑ Registers per thread (known at compile time)
 - ❑ Shared memory per block (defined by the programmer)

CUDA Occupancy Calculator (I)



CUDA Occupancy Calculator (II)



The screenshot shows the NVIDIA Developer Zone documentation page for the CUDA Occupancy Calculator. The header includes the NVIDIA logo, 'DEVELOPER ZONE', and 'CUDA TOOLKIT DOCUMENTATION'. A search bar is located in the top right. The left sidebar contains a navigation menu with 'CUDA Toolkit v11.6.2', 'CUDA Occupancy Calculator', and 'Overview'. The main content area features a title 'CUDA Occupancy Calculator' with links to a PDF and older versions, and a 'Send Feedback' link. Below this is a section titled '[Deprecated] CUDA Occupancy Calculator' with a paragraph explaining its purpose. A paragraph follows stating that the Excel-based calculator is deprecated and pointing to 'Nsight Compute' documentation. An 'Overview' section explains how the calculator works by computing multiprocessor occupancy based on active warps and registers. A final paragraph notes that the CUDA compiler attempts to minimize register usage to maximize thread blocks. The page concludes with a link to download the spreadsheet.

NVIDIA DEVELOPER ZONE CUDA TOOLKIT DOCUMENTATION

Search

CUDA Toolkit v11.6.2

CUDA Occupancy Calculator

Overview

CUDA Occupancy Calculator ([PDF](#)) - v11.6.2 ([older](#)) - Last updated March 24, 2022 - [Send Feedback](#)

[Deprecated] CUDA Occupancy Calculator

The CUDA Occupancy Calculator allows you to compute the multiprocessor occupancy of a GPU by a given CUDA kernel.

[Deprecated] Excel based Occupancy Calculator is deprecated. Occupancy calculator is available in [Nsight Compute](#). Please refer to [Nsight Compute Occupancy Calculator](#) documentation for more details on usage.

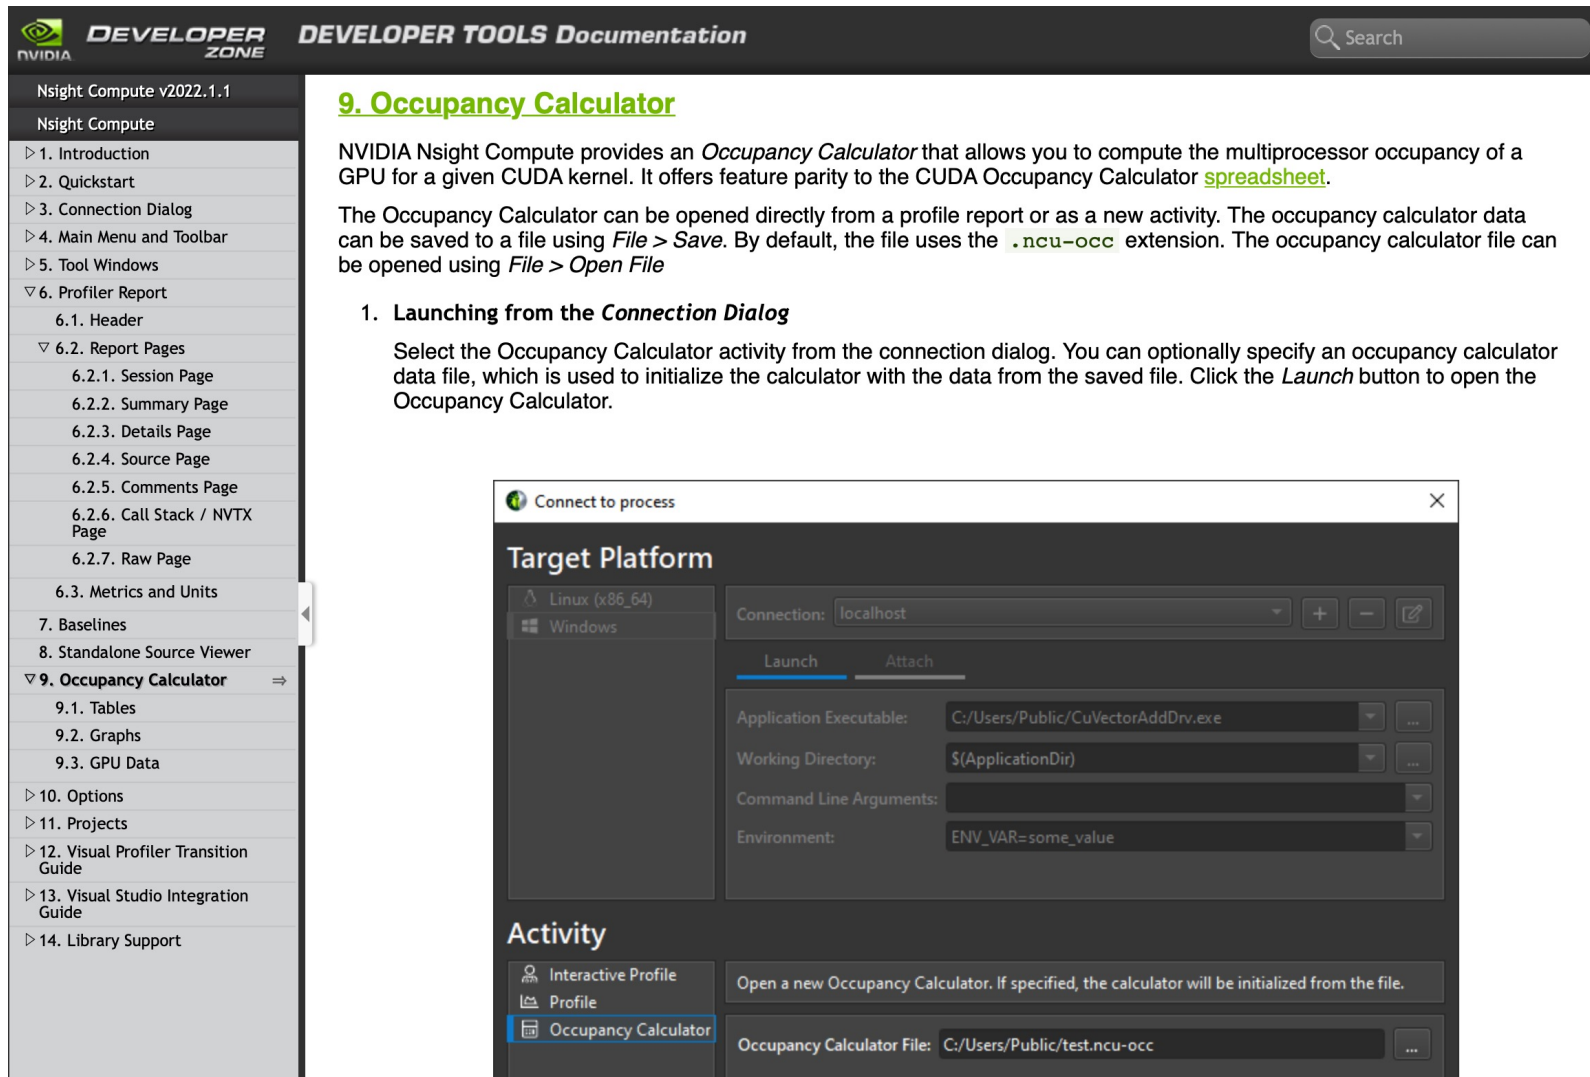
Overview

The CUDA Occupancy Calculator allows you to compute the multiprocessor occupancy of a GPU by a given CUDA kernel. The multiprocessor occupancy is the ratio of active warps to the maximum number of warps supported on a multiprocessor of the GPU. Each multiprocessor on the device has a set of N registers available for use by CUDA program threads. These registers are a shared resource that are allocated among the thread blocks executing on a multiprocessor.

The CUDA compiler attempts to minimize register usage to maximize the number of thread blocks that can be active in the machine simultaneously. If a program tries to launch a kernel for which the registers used per thread times the thread block size is greater than N, the launch will fail.

Click [CUDA Occupancy Calculator](#)[XLS] to download the spreadsheet.

CUDA Occupancy Calculator (III)



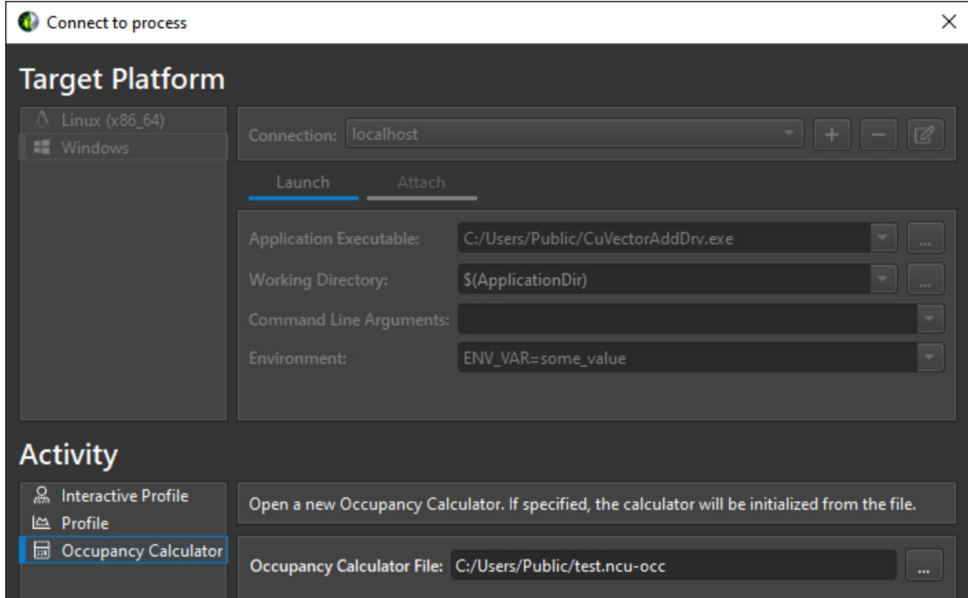
9. Occupancy Calculator

NVIDIA Nsight Compute provides an *Occupancy Calculator* that allows you to compute the multiprocessor occupancy of a GPU for a given CUDA kernel. It offers feature parity to the CUDA Occupancy Calculator [spreadsheet](#).

The Occupancy Calculator can be opened directly from a profile report or as a new activity. The occupancy calculator data can be saved to a file using *File > Save*. By default, the file uses the `.ncu-occ` extension. The occupancy calculator file can be opened using *File > Open File*.

- 1. Launching from the *Connection Dialog***

Select the Occupancy Calculator activity from the connection dialog. You can optionally specify an occupancy calculator data file, which is used to initialize the calculator with the data from the saved file. Click the *Launch* button to open the Occupancy Calculator.



Memory Layout of a Matrix in C

$M_{0,0}$	$M_{1,0}$	$M_{2,0}$	$M_{3,0}$
$M_{0,1}$	$M_{1,1}$	$M_{2,1}$	$M_{3,1}$
$M_{0,2}$	$M_{1,2}$	$M_{2,2}$	$M_{3,2}$
$M_{0,3}$	$M_{1,3}$	$M_{2,3}$	$M_{3,3}$

M
↓


$M_{0,0}$	$M_{1,0}$	$M_{2,0}$	$M_{3,0}$	$M_{0,1}$	$M_{1,1}$	$M_{2,1}$	$M_{3,1}$	$M_{0,2}$	$M_{1,2}$	$M_{2,2}$	$M_{3,2}$	$M_{0,3}$	$M_{1,3}$	$M_{2,3}$	$M_{3,3}$
-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------	-----------

The DRAM Subsystem


The Top-Down View

DRAM Subsystem Organization

- Channel
- DIMM
- Rank
- Chip
- Bank
- Row/Column



The DRAM Subsystem




Breaking down a DIMM (module)


DIMM (Dual in-line memory module)



Side view




Front of DIMM




Rank 0: collection of 8 chips

Back of DIMM

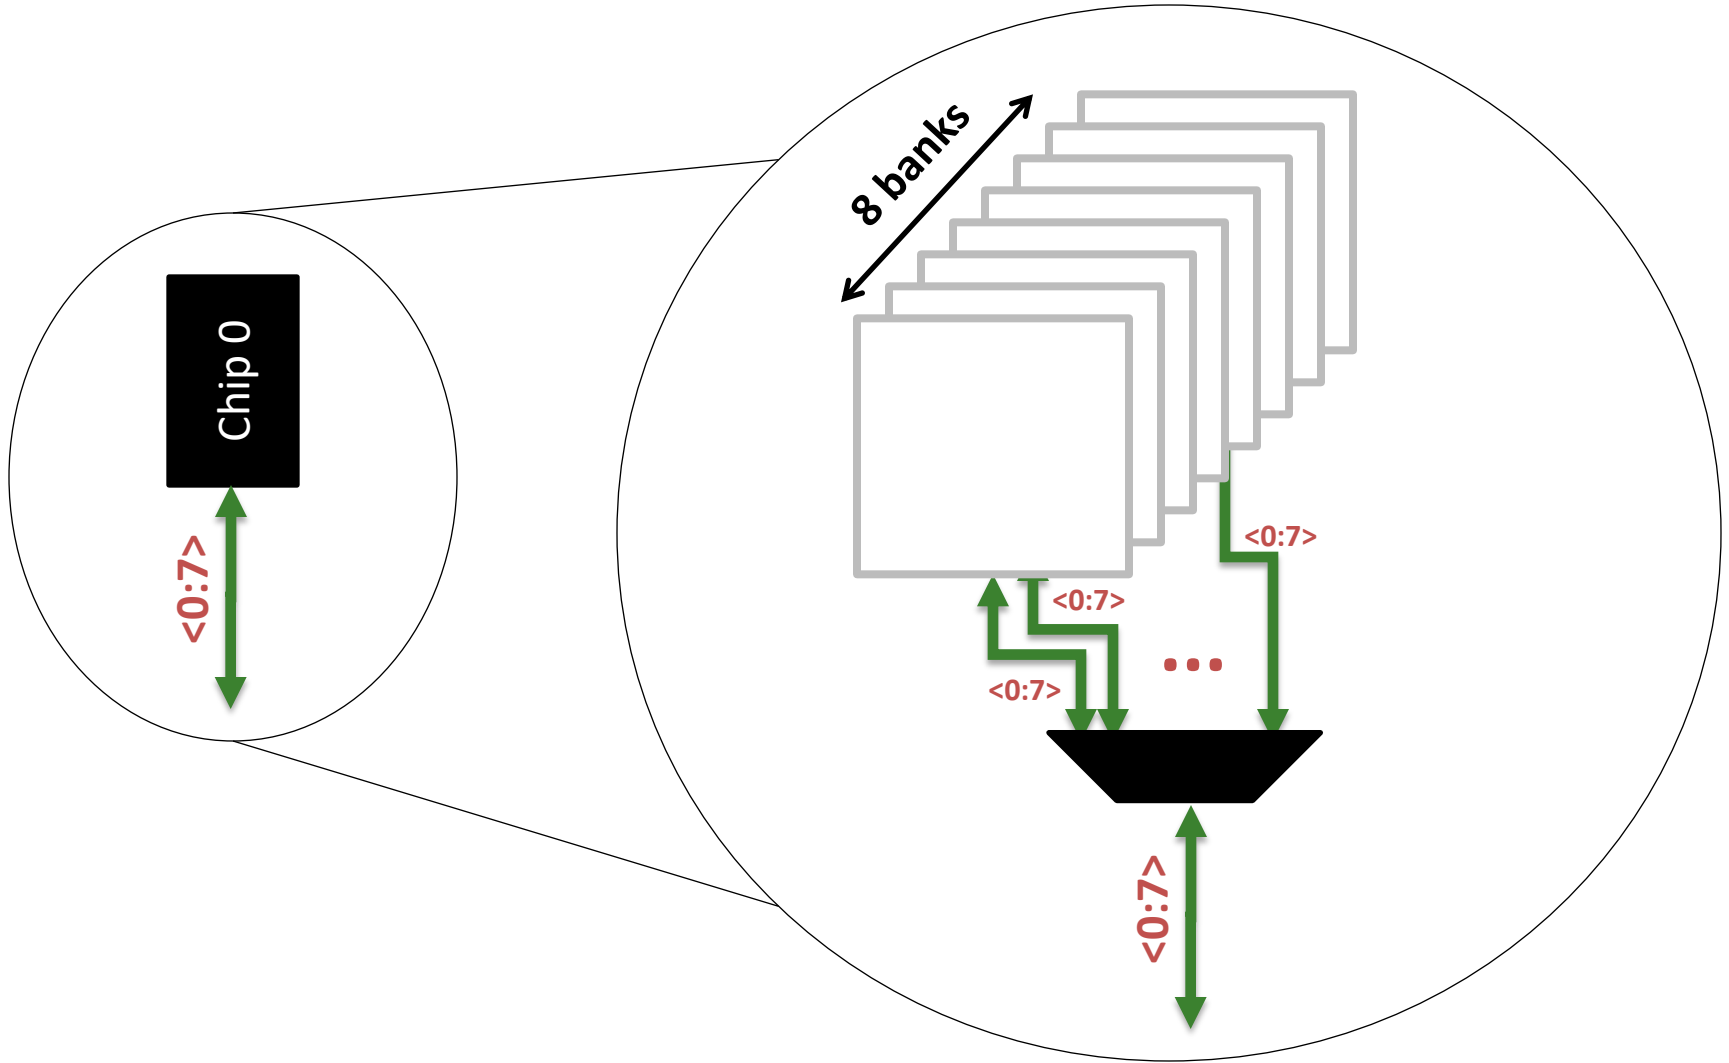


Rank 1


Breaking down a Rank




Breaking down a Chip



Inside a DRAM Chip



DRAM Cell Operation




1. ACTIVATE (ACT)

2. READ/WRITE

3. PRECHARGE (PRE)

DRAM Cell Operation - ACTIVATE




1. ACTIVATE (ACT)

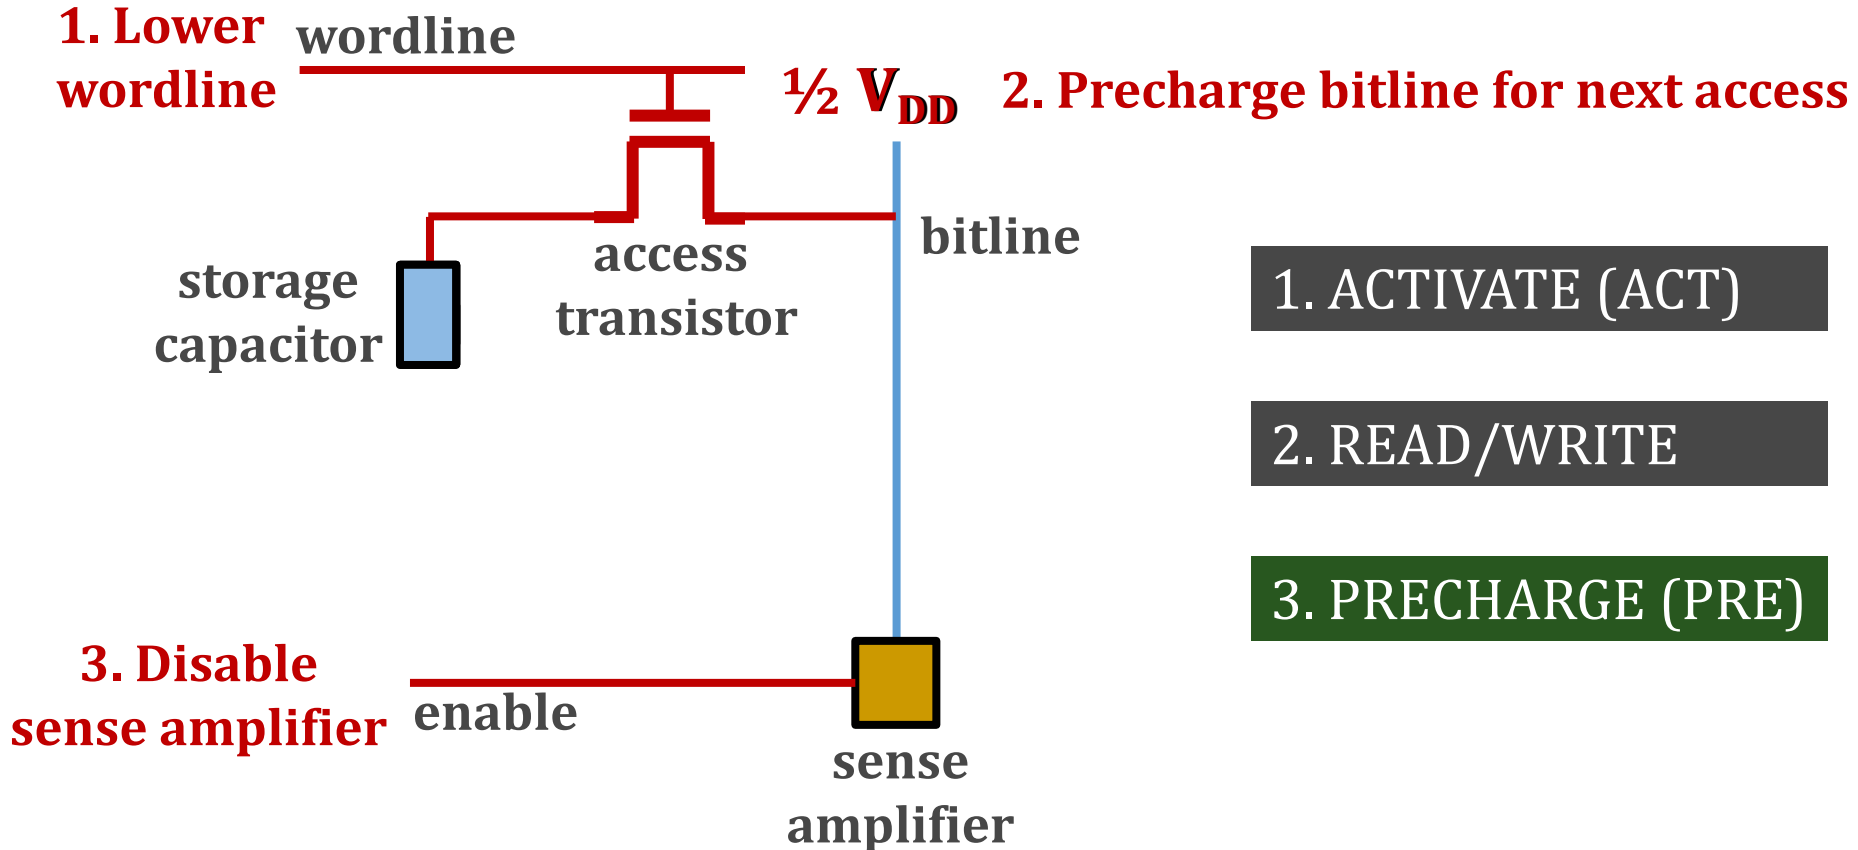
2. READ/WRITE

3. PRECHARGE (PRE)

DRAM Cell Operation – READ/WRITE




DRAM Cell Operation - PRECHARGE



DRAM Bank Operation

Access Address:
(Row 0, Column 0)
(Row 0, Column 1)
(Row 0, Column 85)
(Row 1, Column 0)



DRAM Burst

- Accessing data in different bursts (rows)

- Need to access the array again



- Accessing data in the same burst (row)


- No need to access the array again, just the multiplexer



- Accessing data in the same burst is faster than accessing data in different bursts

Recall: Memory Banking

- Memory is divided into **banks** that can be accessed independently; banks share address and data buses (to minimize pin cost)
- Can start and complete one bank access per cycle
- Can sustain N concurrent accesses if all N go to different banks



Multiple Banks (Interleaving) and Channels


- Multiple banks
 - Enable **concurrent DRAM accesses**
 - Bits in address determine which bank an address resides in
- Multiple independent channels serve the same purpose
 - But they are even better because they have **separate data buses**
 - **Increased bus bandwidth**
- Enabling more concurrency requires reducing
 - Bank conflicts
 - Channel conflicts
- How to select/randomize bank/channel indices in address?
 - Lower order bits have more entropy
 - Randomizing hash functions (XOR of different address bits)

Latency Hiding with Multiple Banks

- With one bank, time still wasted in between bursts



- Latency can be hidden by having multiple banks



- Need many threads to simultaneously access memory to keep all banks busy
 - Achieved with having high occupancy in GPU cores (SMs)
 - Similar idea to hiding pipeline latency in the core

Lecture on Memory Organization & Technology

Breaking down a Chip

The diagram shows a hierarchical view of a chip's memory organization. On the left, a small circle labeled 'Chip 0' contains a vertical double-headed arrow labeled '<0:7>'. A larger circle on the right provides a detailed view of the chip's internal structure, showing a stack of 8 memory banks. Below the banks, a black trapezoidal shape represents a memory array or controller, with multiple data paths labeled '<0:7>' connecting the banks to it. A vertical double-headed arrow labeled '<0:7>' is also shown at the bottom of this structure. A 'zoom' watermark is visible in the bottom right corner of the video frame.

Onur Mutlu

DEPARTMENT OF INFORMATION TECHNOLOGY AND ELECTRICAL ENGINEERING (D-ITET)

Digital Design & Comp. Arch. - Lecture 22: Memory Organization & Technology (ETH Zürich, Spring '21)

2,906 views • Streamed live on May 21, 2021

83 1 SHARE SAVE ...



Onur Mutlu Lectures

19.6K subscribers

Digital Design and Computer Architecture, ETH Zürich, Spring 2021 (

<https://safari.ethz.ch/digitaltechnik...>)

ANALYTICS


EDIT VIDEO

Memory Coalescing (I)

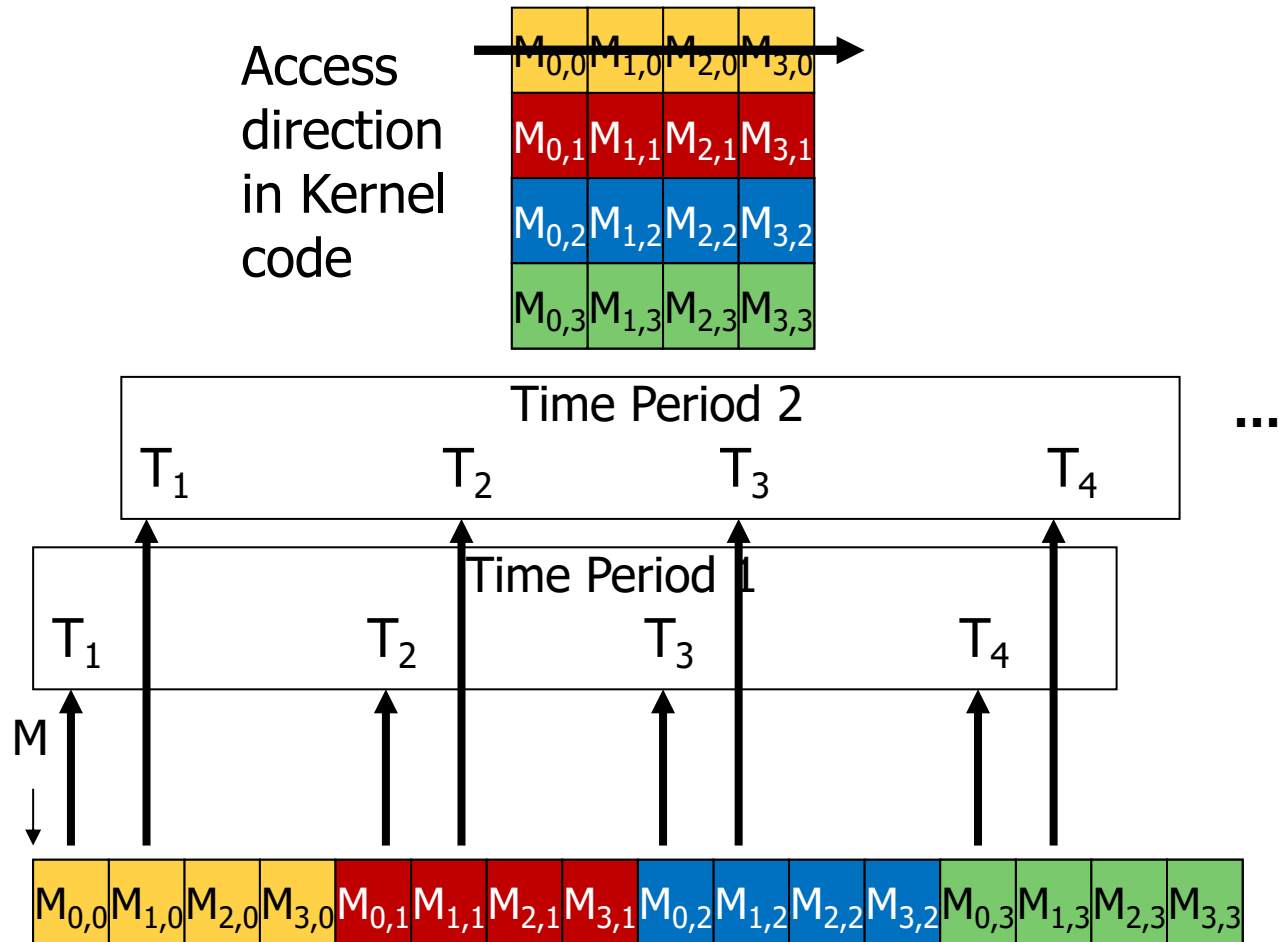
- When threads in the same warp access consecutive memory locations in the same burst, the accesses can be combined and served by one burst
 - One DRAM transaction is needed
 - Known as memory coalescing
- If threads in the same warp access locations not in the same burst, accesses cannot be combined
 - Multiple transactions are needed
 - Takes longer to service data to the warp
 - Sometimes called memory divergence

Memory Coalescing (II)

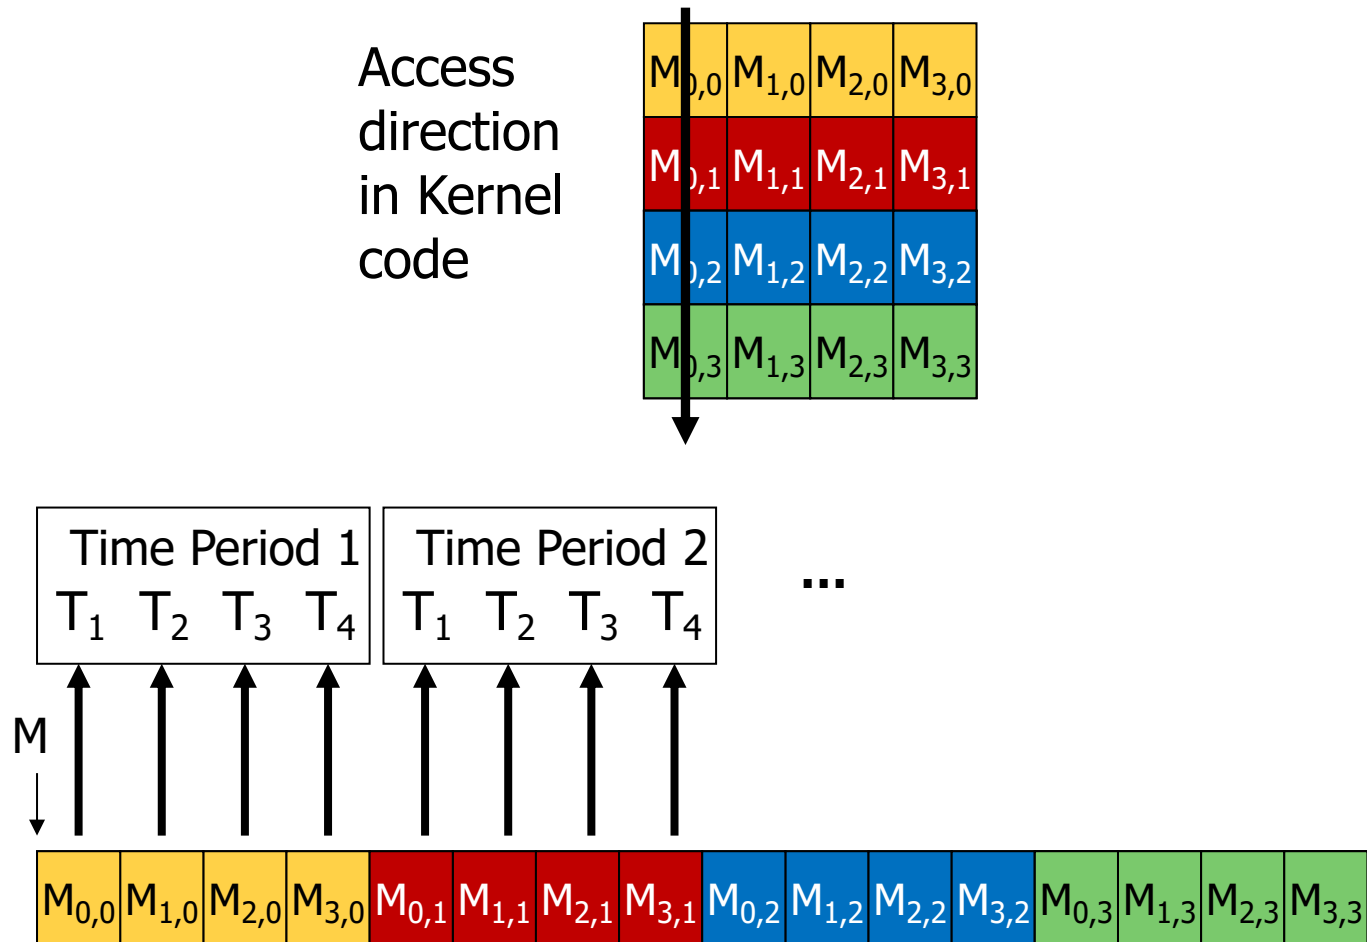
- When accessing global memory, we want to make sure that **concurrent threads access nearby memory locations**
- **Peak bandwidth** utilization occurs when all threads in a warp access **one cache line** (or several consecutive cache lines)



Uncoalesced Memory Accesses



Coalesced Memory Accesses




AoS vs. SoA

■ Array of Structures vs. Structure of Arrays


Structure of
Arrays
(SoA)

```
struct foo{  
    float a[8];  
    float b[8];  
    float c[8];  
    int d[8];  
} A;
```



Array of
Structures
(AoS)


```
struct foo{  
    float a;  
    float b;  
    float c;  
    int d;  
} A[8];
```




CPUs Prefer AoS, GPUs Prefer SoA

■ Linear and strided accesses

GPU



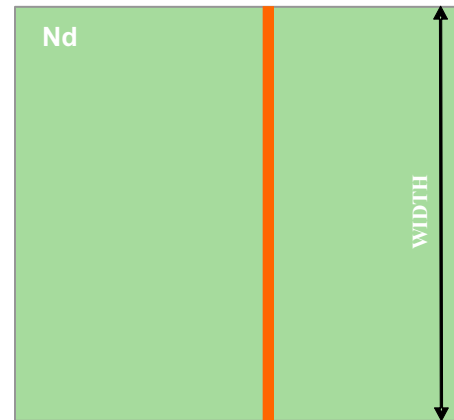
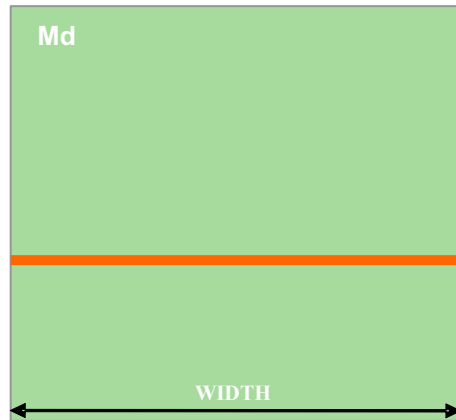
CPU





AMD Kaveri A10-7850K

Use Shared Memory to Improve Coalescing

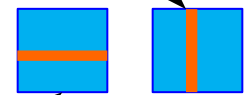
Original
Access
Pattern



Tiled
Access
Pattern




Copy into
scratchpad
memory



Perform
multiplication
with scratchpad
values

Data Reuse


- Same memory locations accessed by neighboring threads



```
for (int i = 0; i < 3; i++){  
    for (int j = 0; j < 3; j++){  
        sum += gauss[i][j] * Image[(i+row-1)*width + (j+col-1)];  
    }  
}
```

Data Reuse: Tiling

- To take advantage of data reuse, we divide the input into **tiles** that can be loaded into **shared memory**



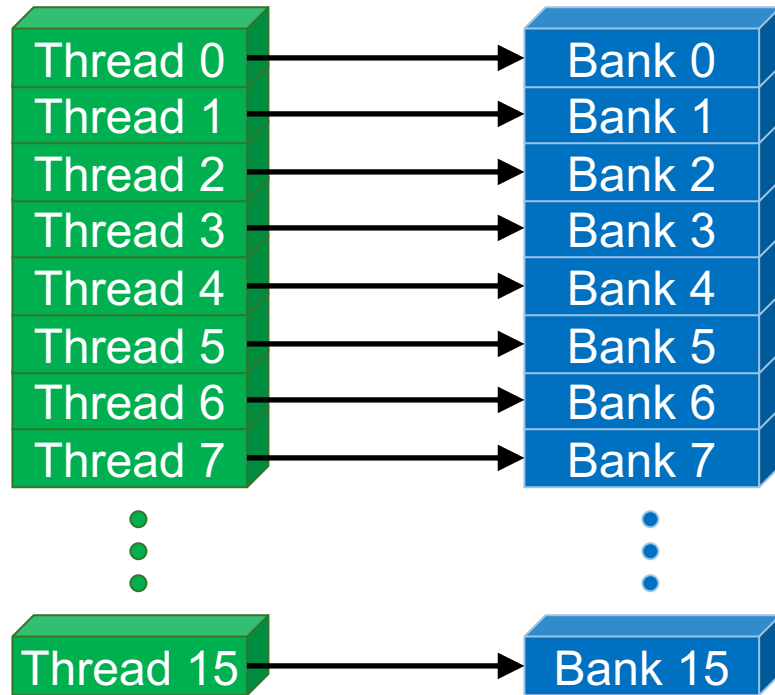
```
__shared__ int l_data[(L_SIZE+2)*(L_SIZE+2)];  
...  
Load tile into shared memory  
__syncthreads();  
for (int i = 0; i < 3; i++){  
    for (int j = 0; j < 3; j++){  
        sum += gauss[i][j] * l_data[(i+l_row-1)*(L_SIZE+2)+j+l_col-1];  
    }  
}
```

Shared Memory

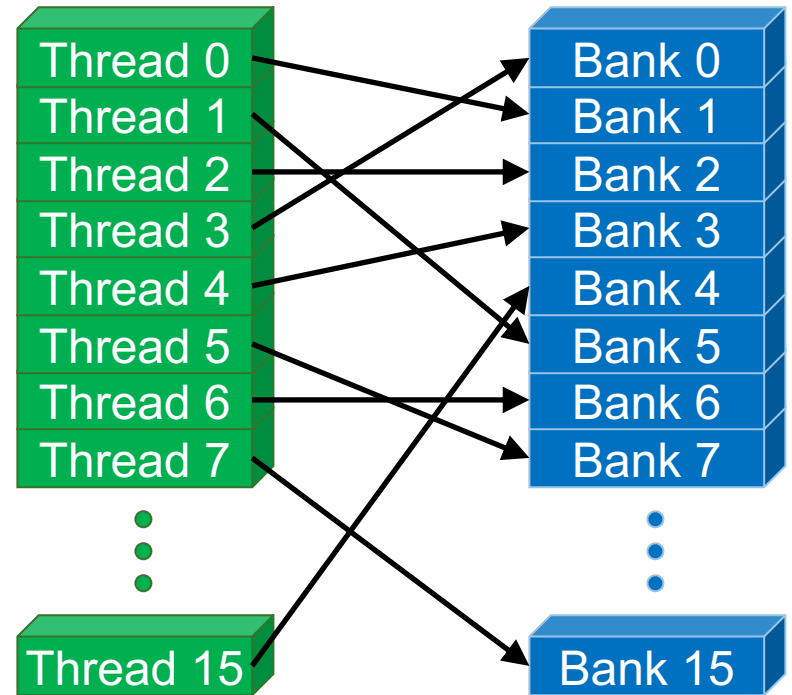
- Shared memory is an **interleaved (banked) memory**
 - Each bank can service one address per cycle
- Typically, 32 banks in NVIDIA GPUs
 - Successive 32-bit words are assigned to successive banks
 - **Bank = Address % 32**
- Bank conflicts are **only possible within a warp**
 - No bank conflicts between different warps

Shared Memory Bank Conflicts (I)

- Bank conflict free



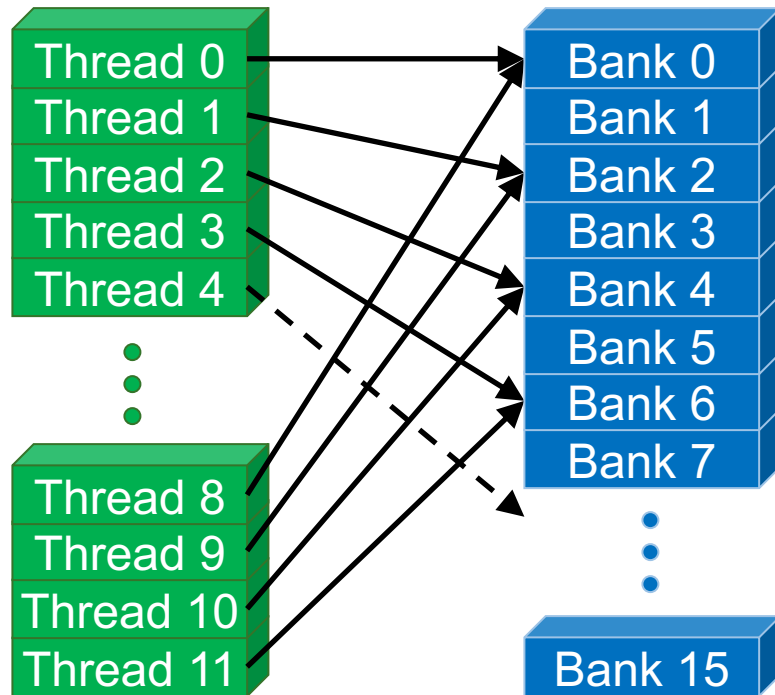
Linear addressing: stride = 1




Random addressing 1:1

Shared Memory Bank Conflicts (II)

■ N-way bank conflicts





2-way bank conflict: stride = 2





8-way bank conflict: stride = 8

Use Shared Memory to Improve Coalescing


Original
Access
Pattern



Tiled
Access
Pattern



Copy into
scratchpad
memory



Perform
multiplication
with scratchpad
values



Reducing Shared Memory Bank Conflicts

- Bank conflicts are only possible within a warp
 - No bank conflicts between different warps
- If strided accesses are needed, some optimization techniques can help
 - Padding
 - Randomized mapping
 - Rau, “Pseudo-randomly interleaved memory,” ISCA 1991
 - Hash functions
 - V.d.Braak+, “Configurable XOR Hash Functions for Banked Scratchpad Memories in GPUs,” IEEE TC, 2016

SIMD Utilization


Threads Can Take Different Paths in Warp-based SIMD

- Each thread can have **conditional control flow instructions**
- Threads can execute different control flow paths



Control Flow Problem in GPUs/SIMT

- A GPU uses a SIMD pipeline to save area on control logic
 - Groups scalar threads into warps
- **Branch divergence** occurs when threads inside warps branch to different execution paths

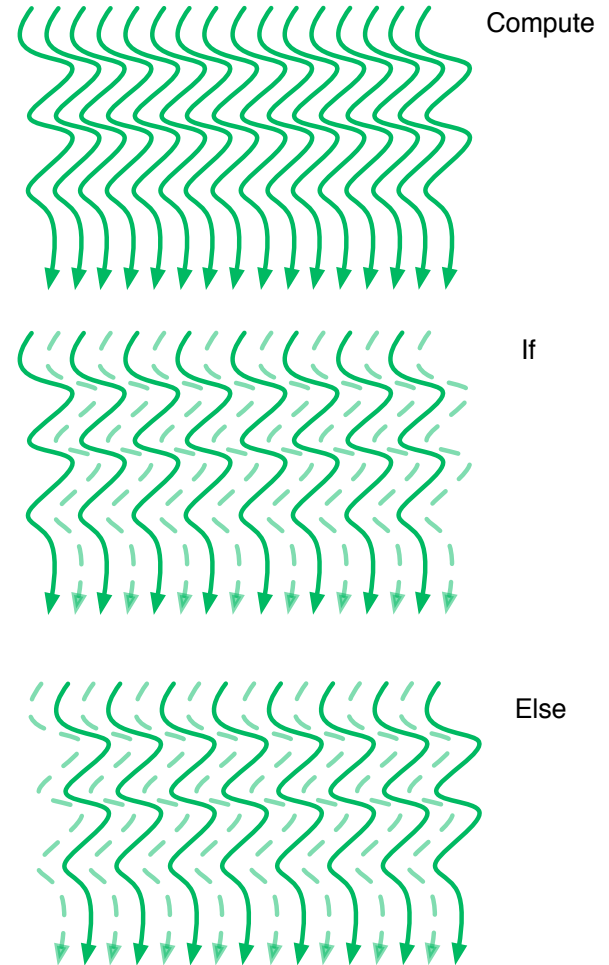


**This is the same as conditional/predicated/masked execution.
Recall the Vector Mask and Masked Vector Operations?**

SIMD Utilization

■ Intra-warp divergence

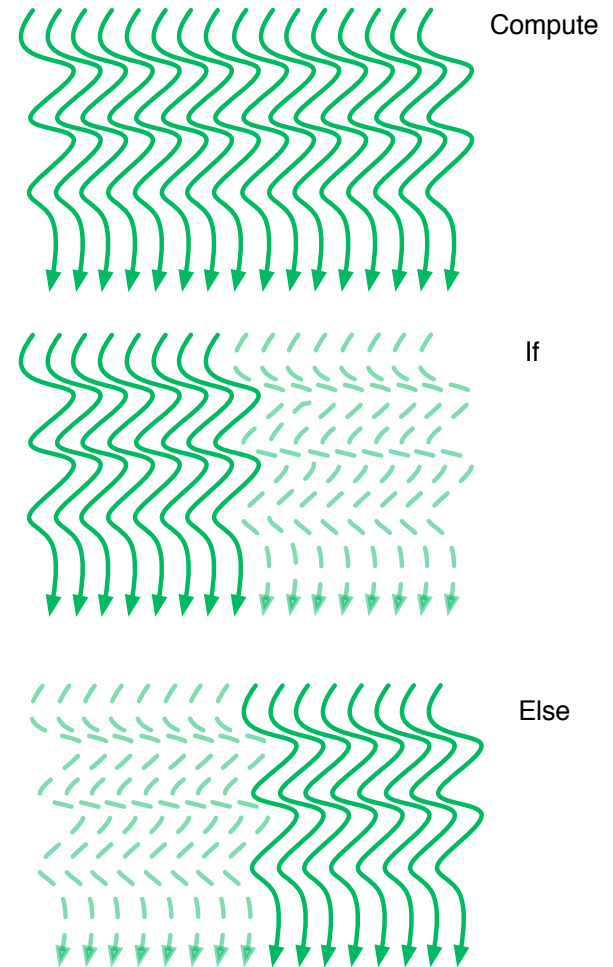
```
Compute(threadIdx.x);  
if (threadIdx.x % 2 == 0){  
    Do_this(threadIdx.x);  
}  
else{  
    Do_that(threadIdx.x);  
}
```




Increasing SIMD Utilization

■ Divergence-free execution

```
Compute(threadIdx.x);  
if (threadIdx.x < 32){  
    Do_this(threadIdx.x * 2);  
}  
else{  
    Do_that((threadIdx.x%32)*2+1);  
}
```



Vector Reduction: Naïve Mapping (I)



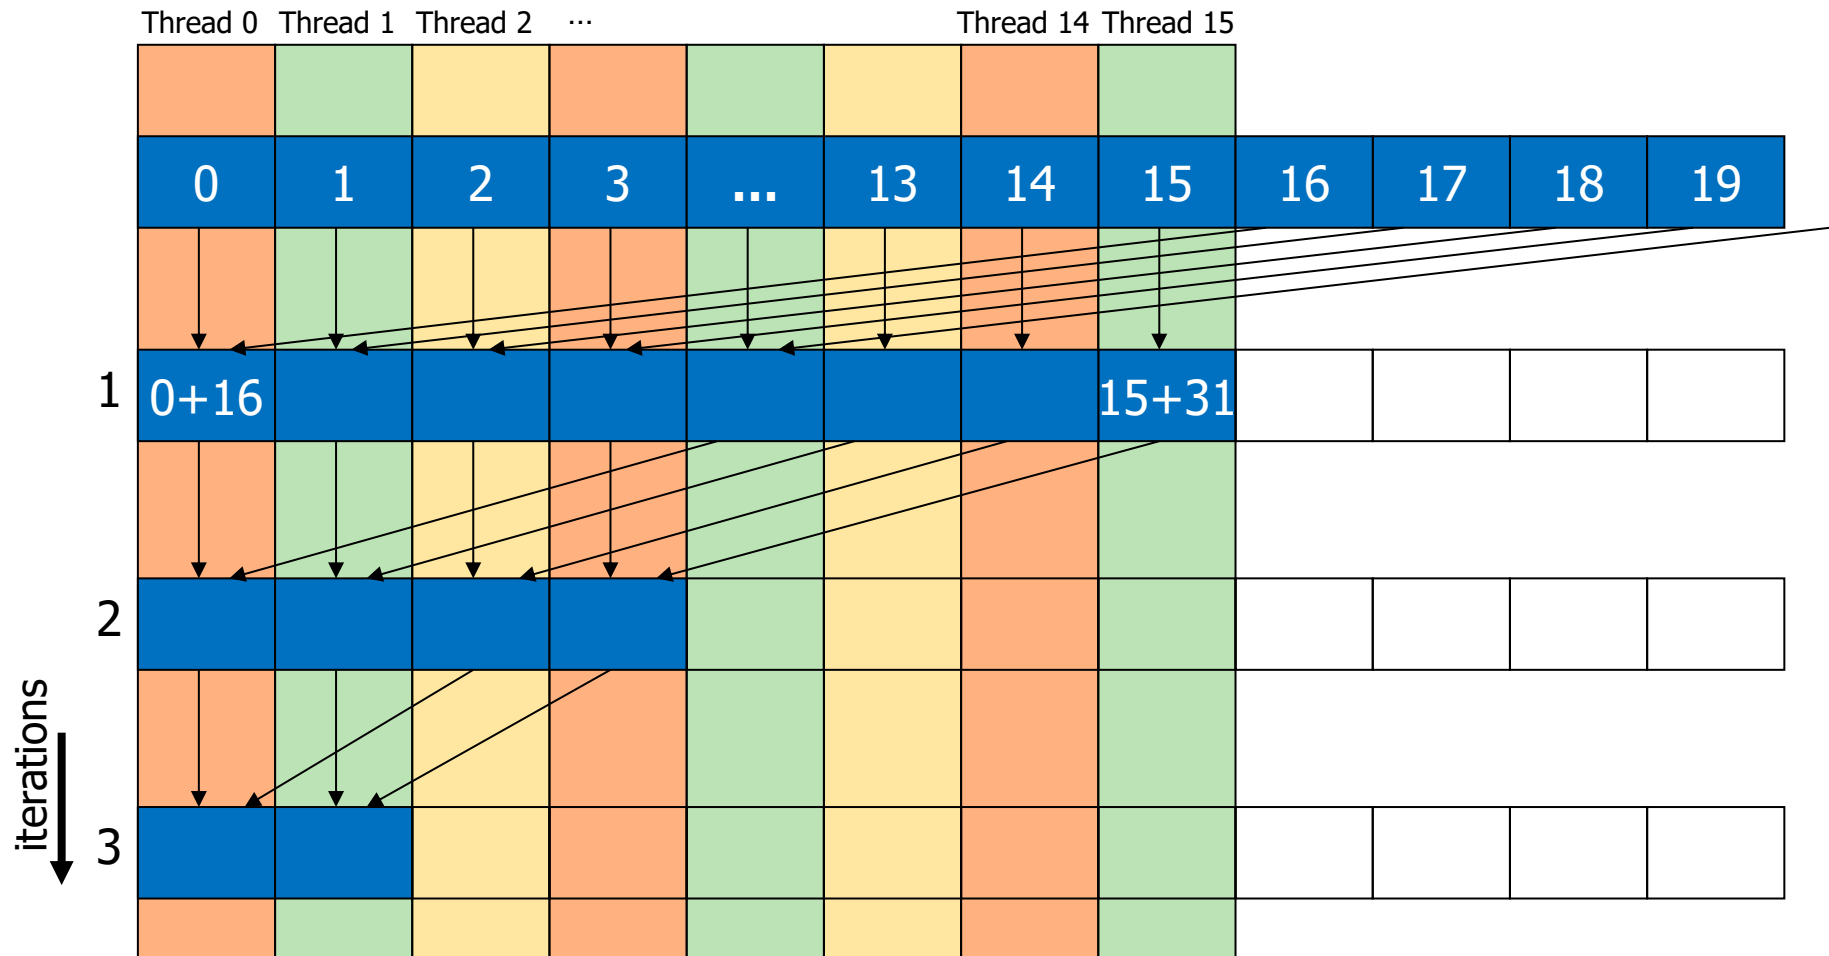
Vector Reduction: Naïve Mapping (II)

- Program with **low SIMD utilization**

```
__shared__ float partialSum[]  
  
unsigned int t = threadIdx.x;  
  
for (int stride = 1; stride < blockDim.x; stride *= 2) {  
  
    __syncthreads();  
  
    if (t % (2*stride) == 0)  
        partialSum[t] += partialSum[t + stride];  
  
}
```

Divergence-Free Mapping (I)

- All active threads belong to the same warp



Divergence-Free Mapping (II)

- Program with high SIMD utilization

```
__shared__ float partialSum[]  
  
unsigned int t = threadIdx.x;  
  
for (int stride = blockDim.x; stride > 0;  stride >> 1){  
  
    __syncthreads();  
  
    if (t < stride)  
        partialSum[t] += partialSum[t + stride];  
  
}
```


Atomic Operations

Atomic Operations (I)

- CUDA provides **atomic instructions** on shared memory and global memory
 - They perform **read-modify-write** operations atomically

- Arithmetic functions

- Add, sub, max, min, exch, inc, dec, CAS

`int atomicAdd(int*, int);`

Return value (old value)

Pointer to shared memory or global memory

Value to add

- Bitwise functions


- And, or, xor

- Datatypes: int, uint, ull, float (half, single, double)*

* Datatypes for different atomic operations in <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#atomic-functions>

Atomic Operations (II)

- Atomic operations serialize the execution if there are atomic conflicts



Uses of Atomic Operations

■ Computation

- Atomics on an array that will be the output of the kernel
- Example
 - Histogram, reduction

■ Synchronization

- Atomics on memory locations that are used for synchronization or coordination
- Example
 - Counters, locks, flags...


- Use them to prevent **data races** when more than one thread need to update the same memory location

Image Histogram

- Histograms are widely used in **image processing**
 - Some **computation before voting** in the histogram may be needed

```
For (each pixel i in image I){  
    Pixel = I[i]                // Read pixel  
    Pixel' = Computation(Pixel) // Optional computation  
    Histogram[Pixel']++         // Vote in histogram bin  
}
```

- Parallel threads frequently incur **atomic conflicts** in image histogram computation



Optimized Parallel Reduction

- 7 versions in CUDA samples: Tree-based reduction in shared memory
 - ❑ Version 0: No whole warps active
 - ❑ Version 1: Contiguous threads, but many bank conflicts
 - ❑ Version 2: No bank conflicts
 - ❑ Version 3: First level of reduction when reading from global memory
 - ❑ Version 4: Warp shuffle or unrolling of final warp
 - ❑ Version 5: Warp shuffle or complete unrolling
 - ❑ Version 6: Multiple elements per thread sequentially


Reduction with Atomic Operations

- 3 new versions of reduction based on 3 previous versions
 - Version 0: No whole warps active
 - Version 3: First level of reduction when reading from global memory
 - Version 6: Multiple elements per thread sequentially
- New versions 7, 8, and 9
 - Replace the `for` loop (tree-based reduction) with one shared memory atomic operation per thread

Asynchronous Data Transfers between CPU and GPU

CUDA Streams

- **CUDA streams** (command queues in OpenCL)
- Sequence of operations that are performed in order
 - 1. Data transfer CPU-GPU
 - 2. Kernel execution
 - D input data instances, B blocks
 - #Streams: $(D / \text{\#Streams})$ data instances, $(B / \text{\#Streams})$ blocks
 - 3. Data transfer GPU-CPU



Asynchronous Transfers between CPU & GPU


- Computation **divided into #Streams**

- D input data instances, B blocks

- #Streams

- $D/\#Streams$ data instances

- $B/\#Streams$ blocks



- Estimates

$$t_E + \frac{t_T}{\#Streams}$$

$t_E \geq t_T$ (dominant kernel)

$$t_T + \frac{t_E}{\#Streams}$$

$t_T > t_E$ (dominant transfers)

Overlap of Data Transfers and Kernel Execution

Code for devices that do not support concurrent data transfers

```
// Create streams
int number_of_streams = 32;
cudaStream_t stream[number_of_streams]; // Stream declaration
for(int i = 0; i < number_of_streams; ++i)
    cudaStreamCreate(&stream[i]); // Stream creation

// CPU-GPU data transfers
for (int i = 0; i < number_of_streams; ++i)
    cudaMemcpyAsync(inputDevPtr + i * size, hostPtr + i * size, size,
                    cudaMemcpyHostToDevice, stream[i]);

// Kernel launches
for (int i = 0; i < number_of_streams; ++i)
    MyKernel<<<num_blocks / number_of_streams, num_threads, 0, stream[i]>>>
        (outputDevPtr + i * size, inputDevPtr + i * size, size);

// GPU-CPU data transfers
for (int i = 0; i < number_of_streams; ++i)
    cudaMemcpyAsync(hostPtr + i * size, outputDevPtr + i * size, size,
                    cudaMemcpyDeviceToHost, stream[i]);

cudaDeviceSynchronize(); // Explicit synchronization

// Destroy streams
for (int i = 0; i < number_of_streams; ++i)
    cudaStreamDestroy(stream[i]); // Stream destruction
```


Check CUDA programming guide
<https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html#streams>

Use Case: Video Processing


- Applications with independent computation on different data instances can benefit from asynchronous transfers
- For instance, **video processing**

Non-streamed execution

A sequence of 6 frames is transferred to device





6 x b blocks compute on the sequence of frames







Streamed execution

A chunk of 2 frames is transferred to device



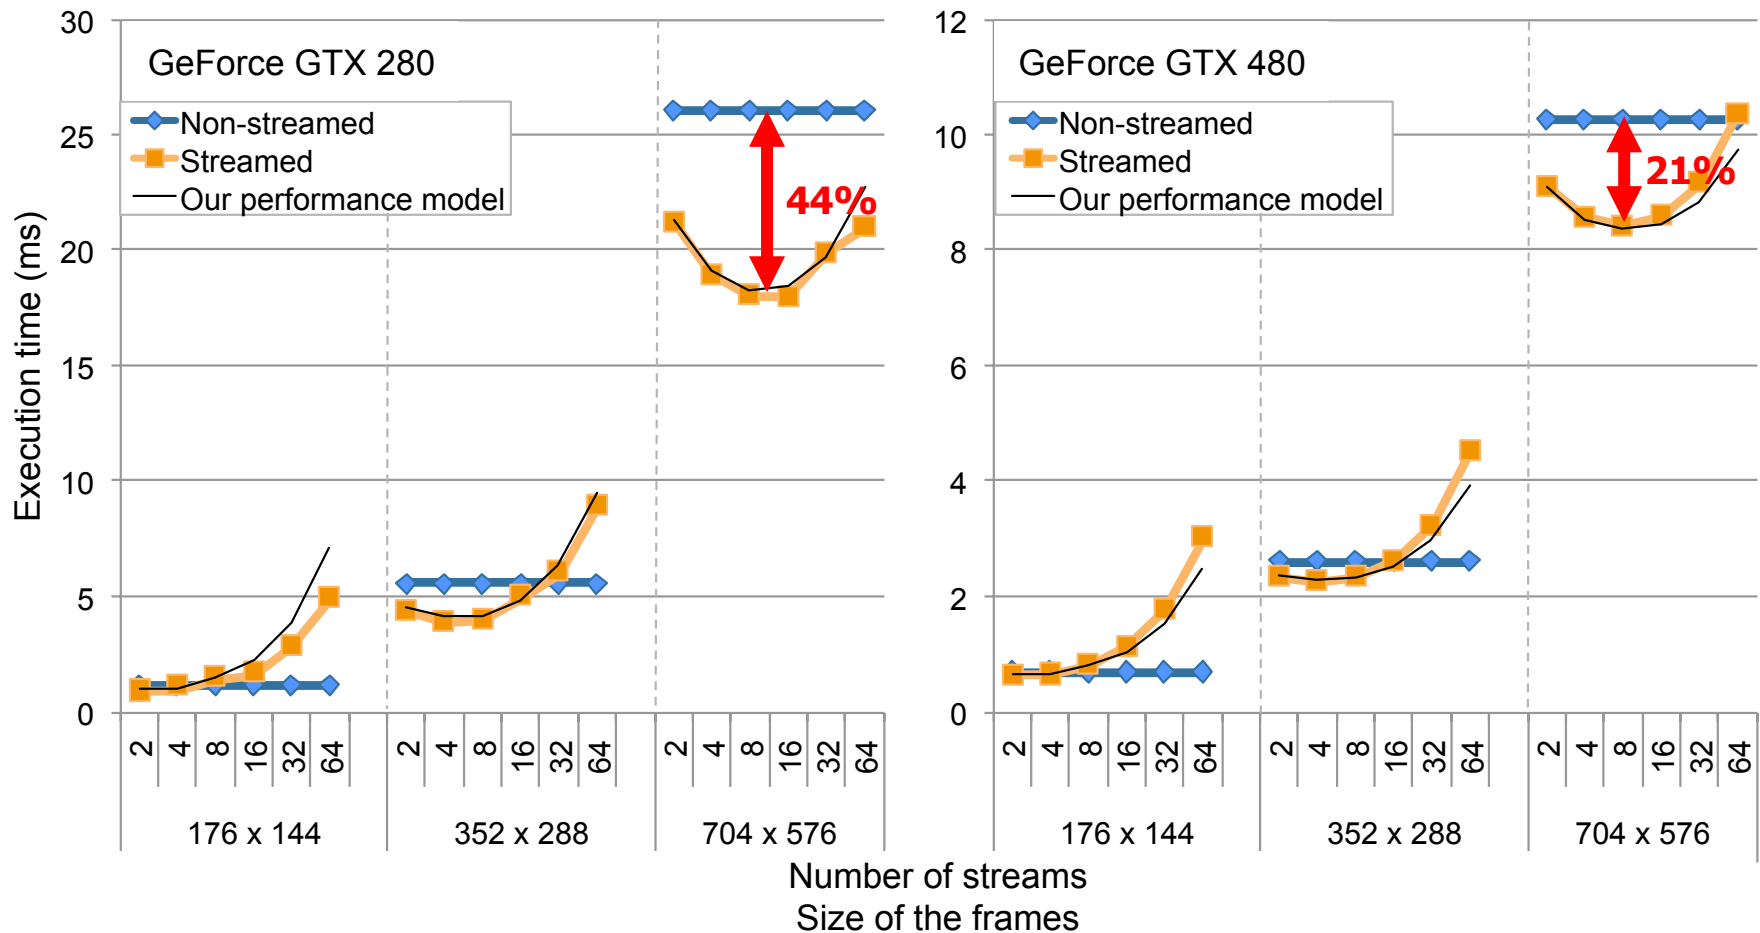
2 x b blocks compute on the chunk, while the second chunk is being transferred



Execution time saved thanks to streams


Video Processing: Performance Results (I)

■ 256-bin histogram calculation



Video Processing: Performance Results (II)

■ RGB-to-grayscale conversion

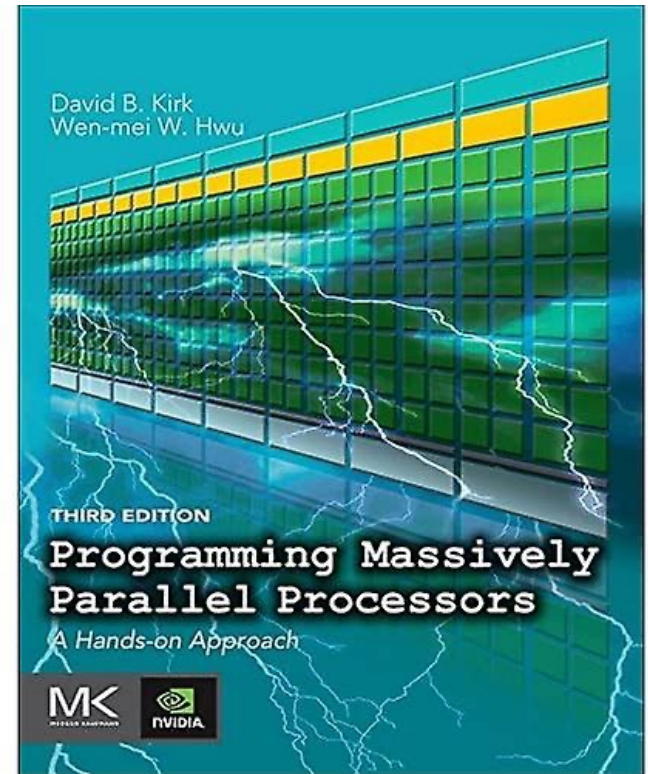


Performance Considerations

- Main bottlenecks
 - ❑ CPU-GPU data transfers
 - ❑ Global memory access
- Memory access
 - ❑ Latency hiding
 - Occupancy
 - ❑ Memory coalescing
 - ❑ Data reuse
 - Shared memory usage
- SIMD (Warp) Utilization: Divergence
- Other considerations
 - ❑ Atomic operations: Serialization
 - ❑ Data transfers between CPU and GPU
 - Overlap of communication and computation

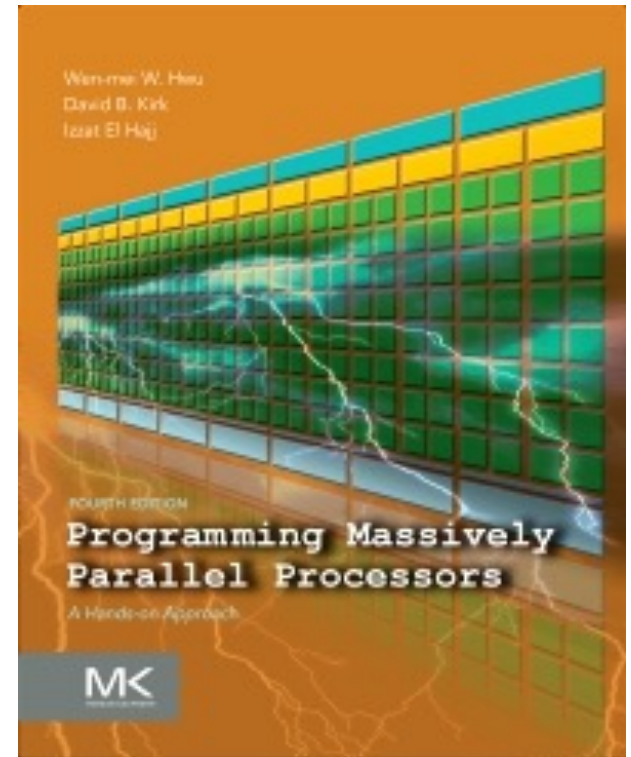
Recommended Readings (I)

- Hwu and Kirk, “**Programming Massively Parallel Processors,**” Third Edition, 2017
 - Chapter 5: Performance considerations
 - Chapter 18 - Programming a heterogeneous computing cluster, Section 18.5



Recommended Readings (II)

- Hwu and Kirk and El Hajj, “[Programming Massively Parallel Processors](#),” Fourth Edition, 2022
 - Chapter 6 - Performance considerations
 - Chapter 20 - Programming a heterogeneous computing cluster, Section 20.5



P&S Heterogeneous Systems

GPU Performance Considerations

Dr. Juan Gómez Luna

Prof. Onur Mutlu

ETH Zürich

Fall 2022

31 October 2022