高性能计算的新发展

基于图形处理器的并行计算及CUDA编程

Ying Liu, Associate Prof., Ph.D

Graduate University of Chinese Academy of Sciences
Research Center on Fictitious Economy and Data Science

Threads

- Threads
- Warp
- Resource Allocation
- Control Flow
- Memory Coalescing

Streaming Multiprocessor (SM)

- Streaming Multiprocessor (SM)
  - 8 Streaming Processors (SP)
  - 2 Special Function Units (SFU)
- Multi-threaded instruction dispatch
  - 1 to 768 threads active
- 16 KB shared memory per SM
- 86.4 GB/s memory bandwidth

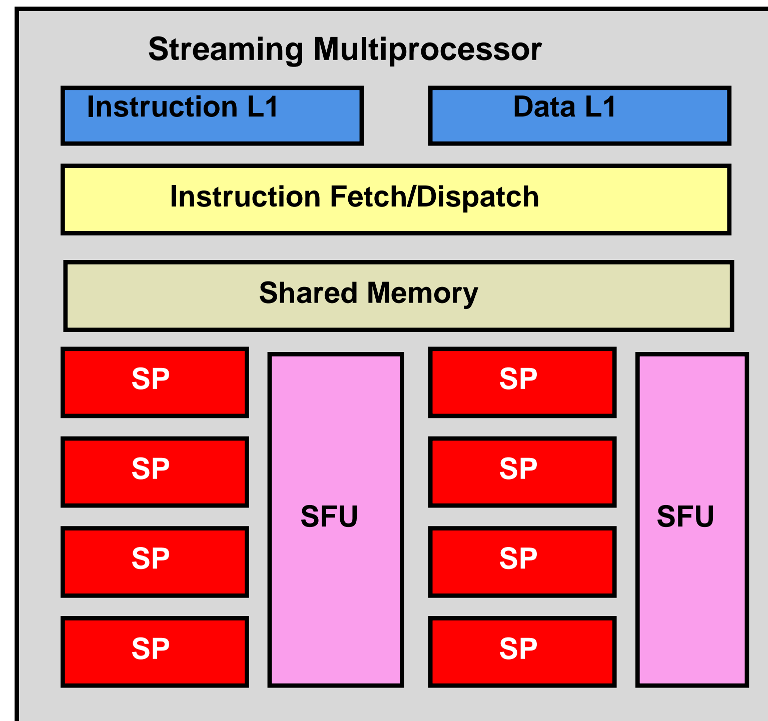

Block IDs and Thread IDs

- Each thread uses IDs to decide what data to work on
  - Block ID: 1D or 2D
  - Thread ID: 1D, 2D, or 3D
- Simplifies memory addressing when processing multidimensional data
  - Image processing
  - Solving PDEs on volumes
  - …

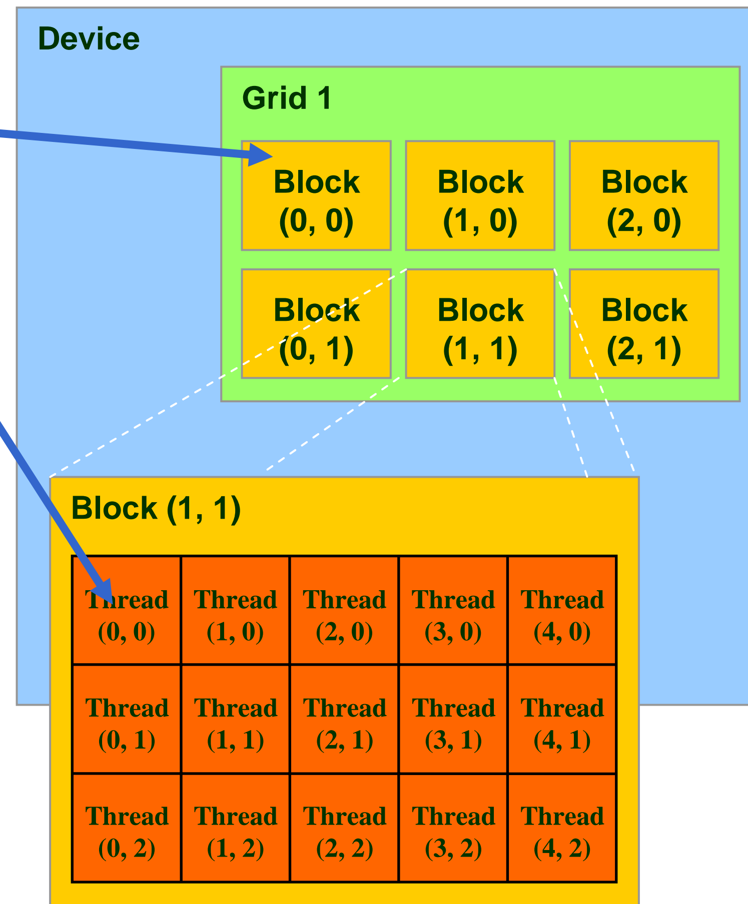

CUDA Thread Block

- Threads are grouped into thread blocks
- Kernel = a grid of thread blocks

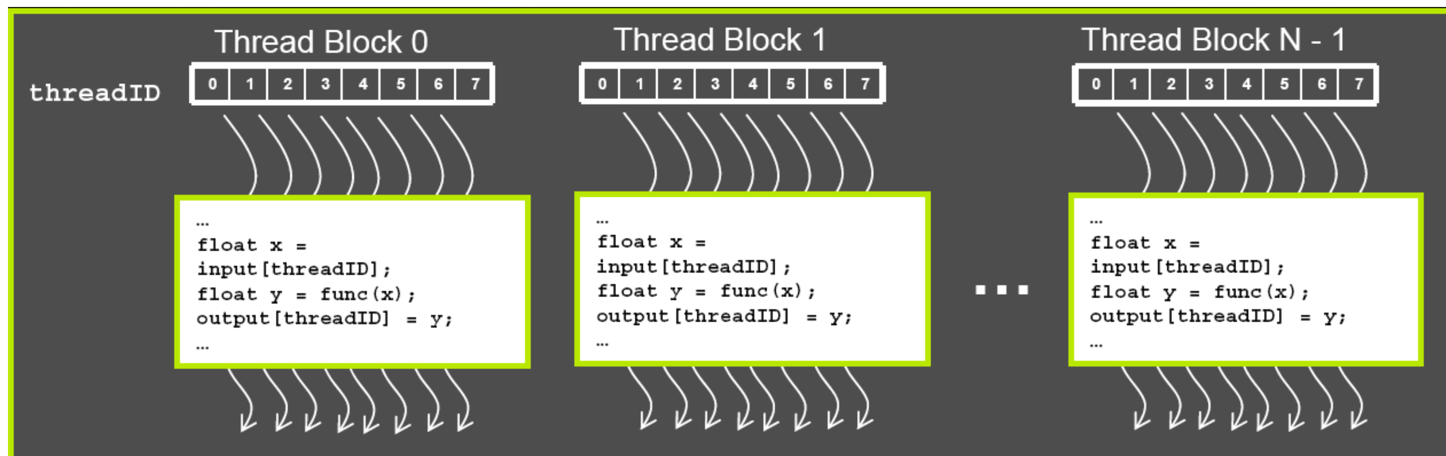

CUDA Thread Block (Cont.)

- Programmer declares block:
  - Block size 1 to 512 concurrent threads
  - Block shape 1D, 2D, or 3D
- Threads have thread id numbers within block
  - Thread program uses thread id to select work and address shared data

CUDA Thread Block

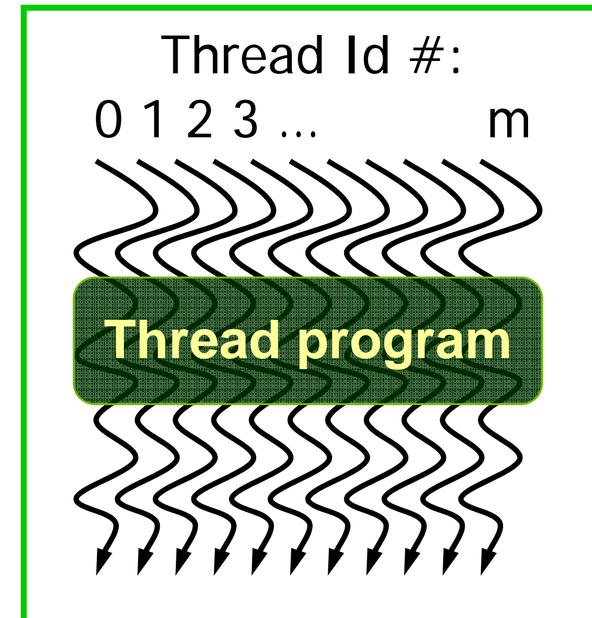

CUDA Thread Block (Cont.)

- All threads in a block execute the same kernel program (SPMD)
- Threads in the same block are able to share data and synchronize while doing their share of the work
- Threads in different blocks are not able to cooperate
  - Each block can execute in any order relative to other blocks

Transparent Scalability

- Hardware is free to assigns blocks to any stream multi-processor (SM) at any time
  - A kernel scales across any number of parallel processors

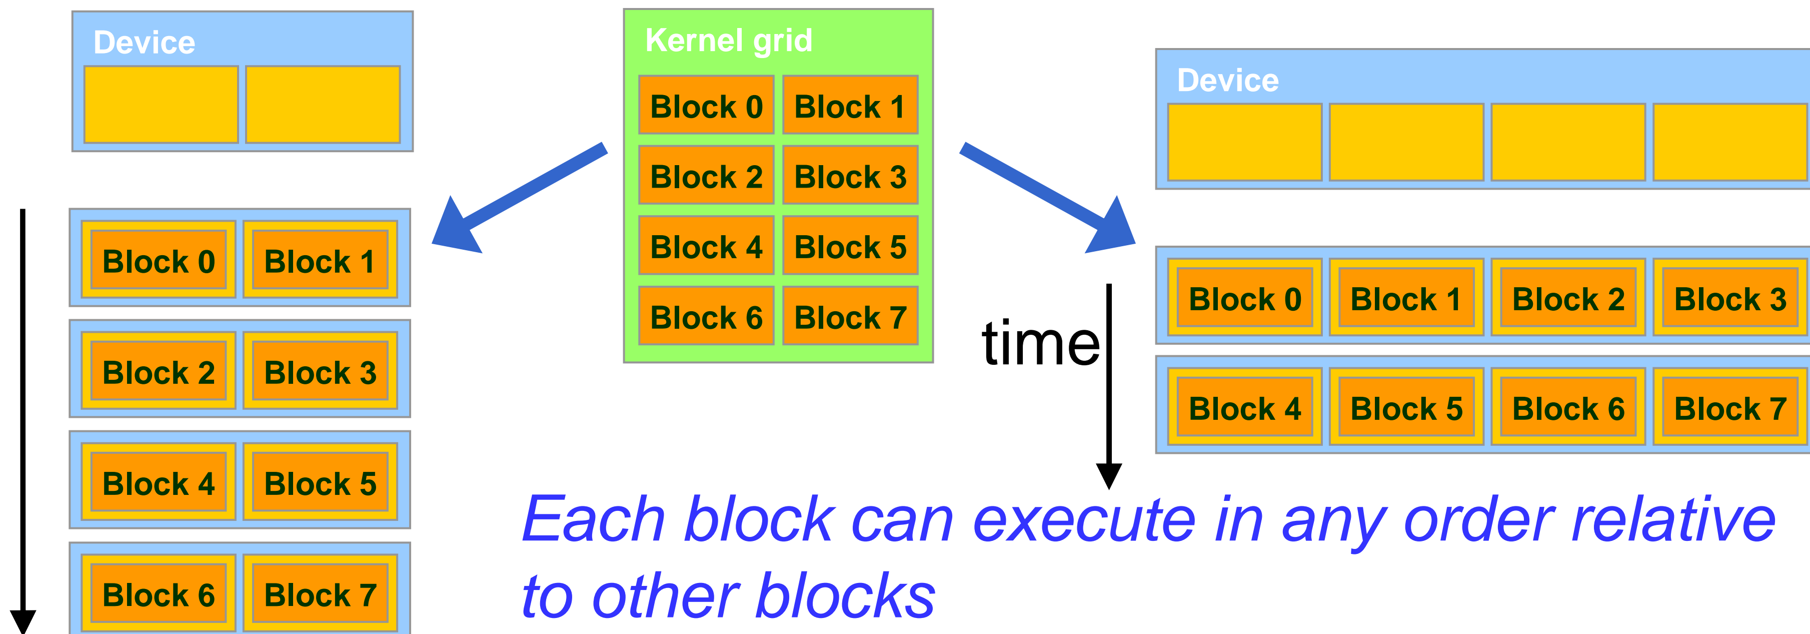

Each block can execute in any order relative to other blocks

Threads

- Threads
- Warp
- Resource Allocation
- Control Flow
- Memory Coalescing

Thread Execution

- A warp of 32 threads physically running on a SM
  - Sharing instructions
  - 4 cycles for 1 warp instruction
    - Why?
  - Dynamically scheduled by SM
    - Executed when operands ready

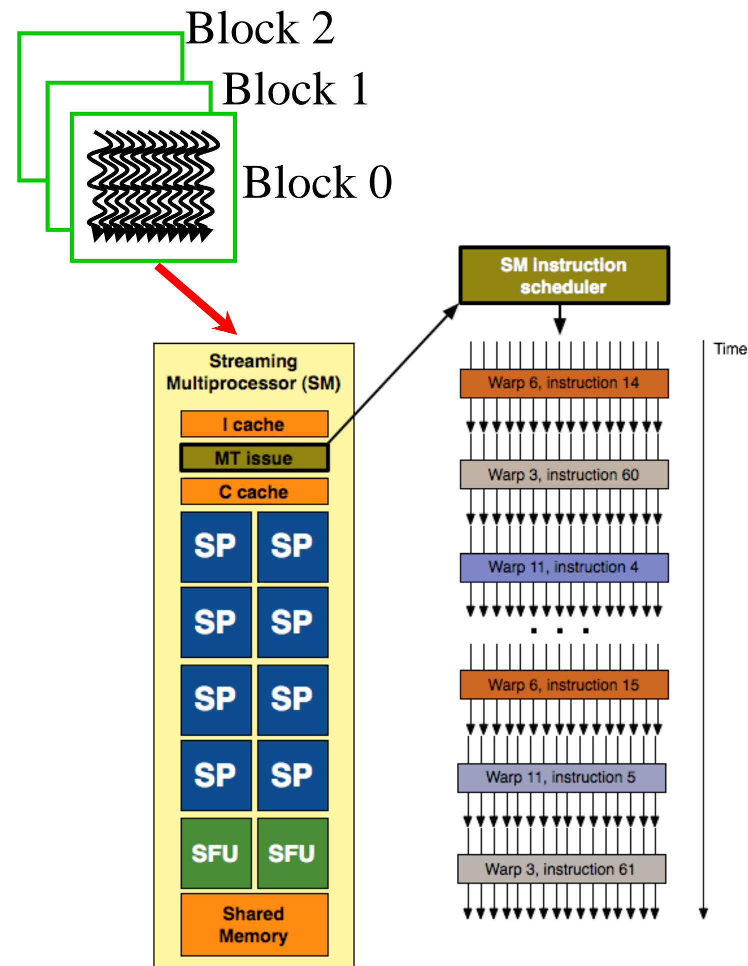

Streaming Multiprocessor (SM)

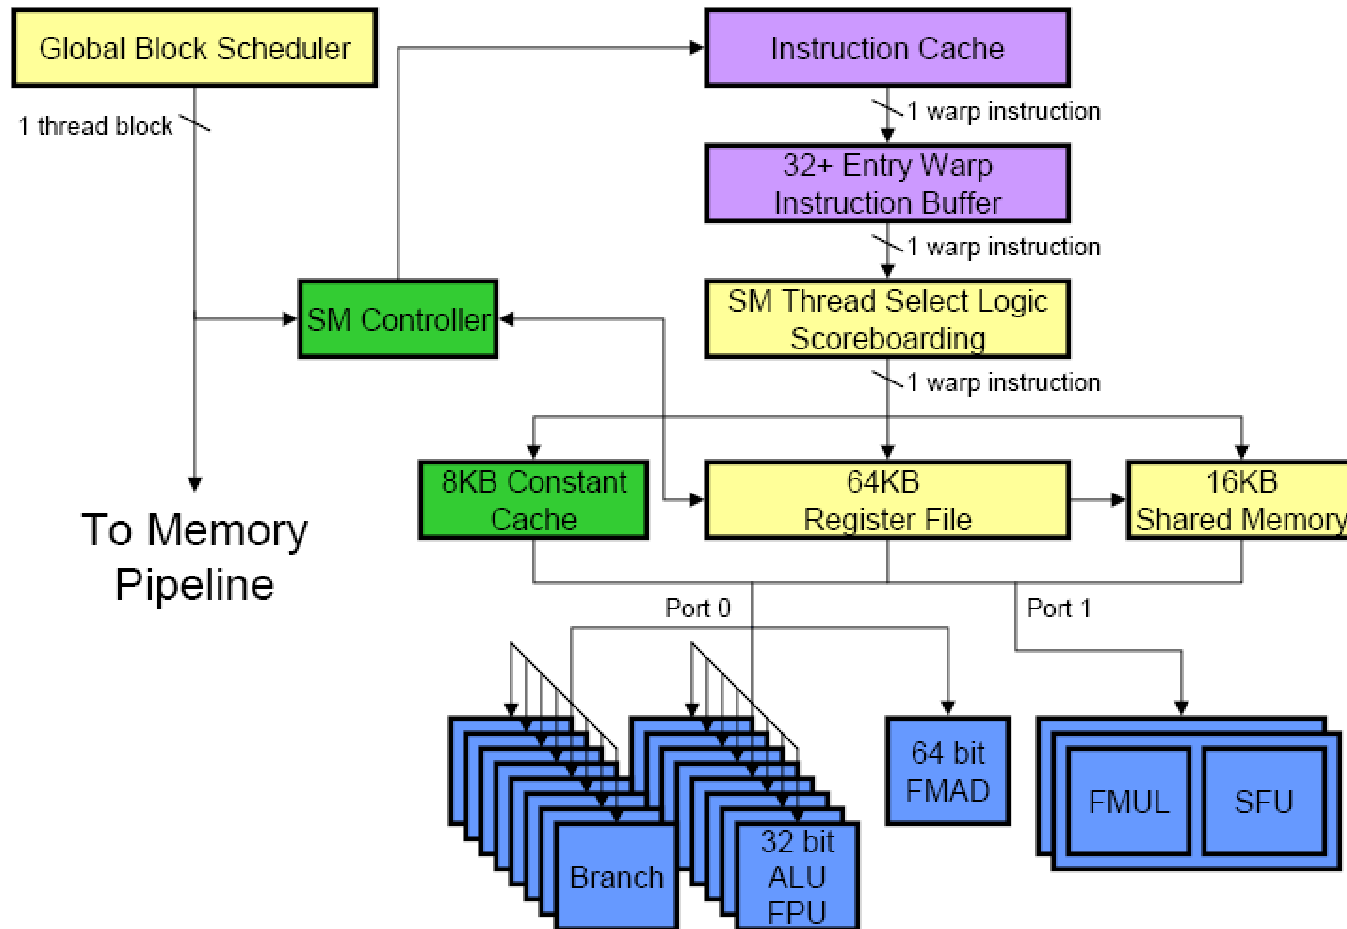

Warp Scheduling

- One instruction fetch per clock cycle
  - from L1 instruction cache
  - load into the instruction buffer
- One warp is selected to execute whose data is ready
  - from the instruction buffer
  - use scoreboard to prevent hazards
- All threads in a warp execute the same instruction
  - SM broadcasts the instruction to 32 threads in the warp

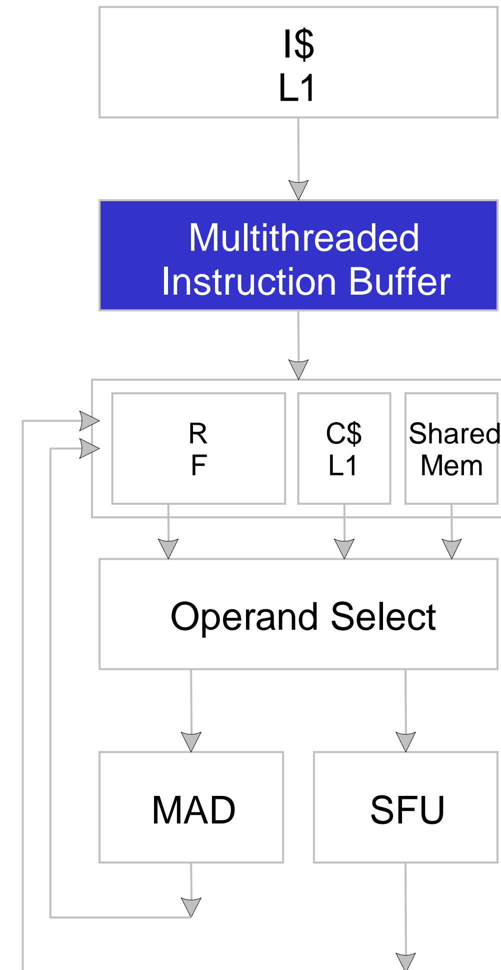

Scoreboarding

- Scoreboard monitors the operands of the instructions in the instruction buffer
  - Give a ‘ready’ mark when all the operands are ready
  - Prevent hazards
- Scoreboarding guarantees the execution of the warp with no interrupt
- Separate memory/processor pipeline

How Thread Blocks Are Partitioned

- Thread blocks are partitioned into warps
  - Thread IDs within a warp are consecutive and increasing
  - Warp 0 starts with Thread ID 0
- Partitioning is always the same
  - Thus you can use this knowledge in control flow
  - However, the exact size of warps may change from generation to generation
- **However, DO NOT rely on any ordering between warps**
  - If there are any dependencies between threads, you must __syncthreads() to get correct results

G80 Example: Executing Thread Blocks

- Threads are assigned to Streaming Multiprocessors in block granularity
  - Up to 8 blocks to each SM as resource allows
  - SM in G80 can take up to 768 threads
    - Could be 256 (threads/block) * 3 blocks
    - Or 128 (threads/block) * 6 blocks, etc.
  - If 3 blocks are assigned to an SM and each block has 256 threads, how many warps are there in an SM?
    - Each block is divided into 256/32 = 8 warps
    - There are 8 * 3 = 24 warps
    - At any point in time, only one of the 24 warps will be selected for instruction fetch and execution!
- Threads run concurrently

G80 Example: Thread Scheduling

- Each block is executed as 32-thread warps
  - An implementation decision, not part of the CUDA programming model
- Warps are scheduling units in SM
  - Warp is a SIMD instruction

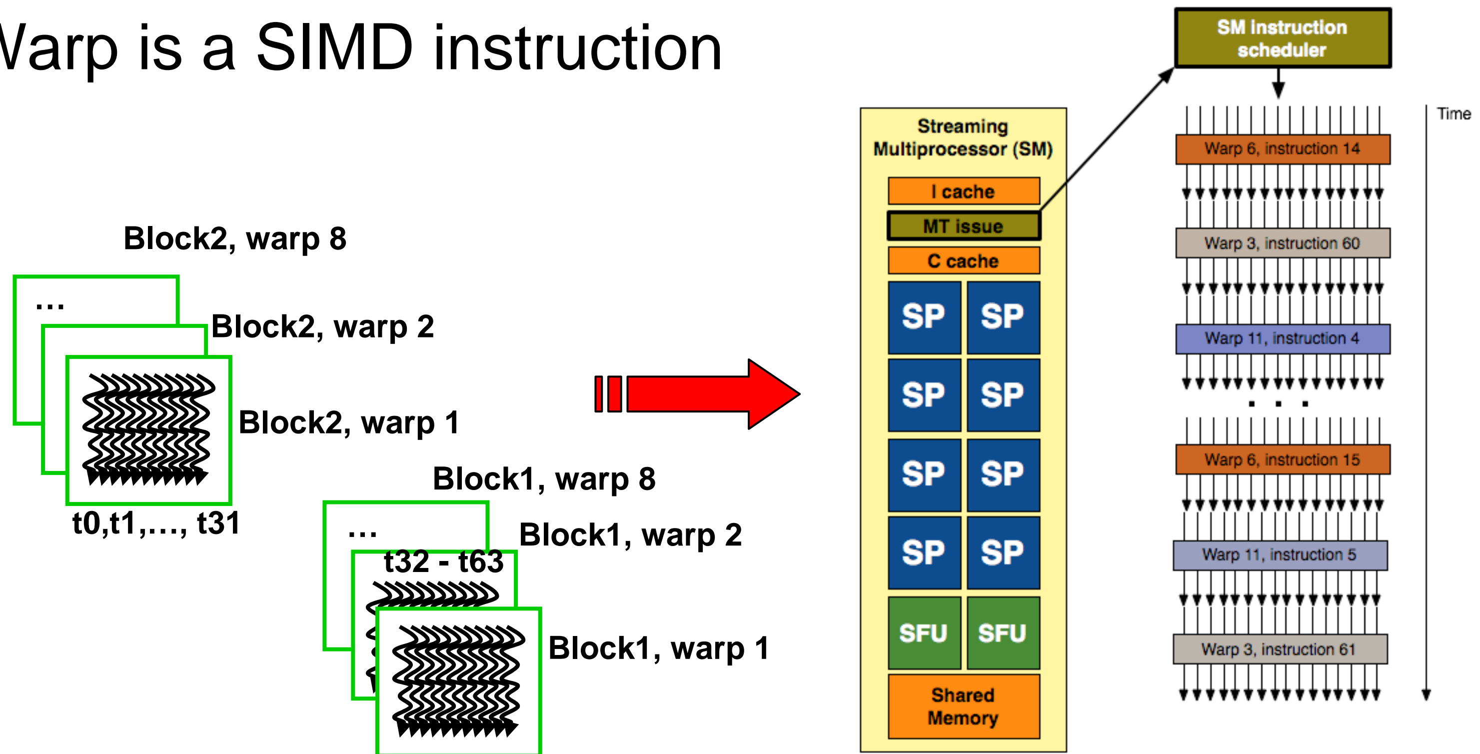

G80 Example: Thread Scheduling (Cont.)

- SM hardware implements **zero-overhead** warp scheduling
  - Warps whose next instruction has its operands ready for consumption are eligible for execution
  - Eligible warps are selected for execution on a prioritized scheduling policy
    - Round-Robin
  - All threads in a warp execute the same instruction when selected

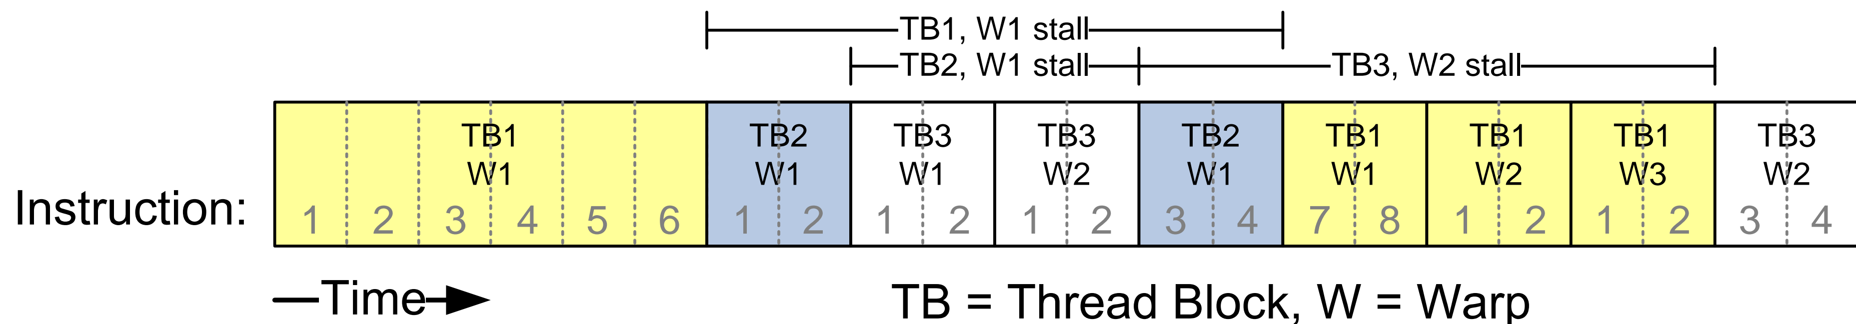

Hiding Memory Latency

- 4 clock cycles needed to dispatch a single instruction for all threads in a warp
  - If one global memory access is needed for every *n* instructions
  - A 400-cycle global memory latency
  - A minimum 400/4*n* warps are needed to tolerate the latency

Threads

- Threads
- Warp
- Resource Allocation
- Control Flow
- Memory Coalescing

Programmer View of Register File

- There are 8192 registers in each SM in G80
  - This is an implementation decision, not part of CUDA
  - Registers are dynamically partitioned across all blocks assigned to the SM
  - Once assigned to a block, the register is NOT accessible by threads in other blocks
  - Each thread in the same block only access registers assigned to itself

Matrix Multiplication Example

- If each block has 16X16 threads and each thread uses 10 registers, how many thread can run on each SM?
  - Each block requires 10*256 = 2560 registers
  - 8192 = **3** * 2560 + change
  - So, three blocks can run on an SM as far as registers are concerned
- How about if each thread increases the use of registers by 1?
  - Each block now requires 11*256 = 2816 registers
  - 8192 < 2816 *3
  - Only two blocks can run on an SM, **1/3 reduction of thread-level parallelism (TLP)**!!!

Dynamic Partitioning

- Dynamic partitioning of SM resources gives more flexibility to compilers/programmers
  - One can run a smaller number of threads that require many registers each or a large number of threads that require few registers each
    - This allows for finer grain threading than traditional CPU threading models
  - The compiler can tradeoff between instruction-level parallelism and thread level parallelism

ILP vs. TLP Example

- Assume that a kernel has 256-thread blocks, 4 independent instructions for each global memory load in the thread program, and each thread uses 10 registers, global loads have 200 cycles
  - 3 blocks can run on each SM
- If a compiler can use one more register to change the dependence pattern so that 8 independent instructions exist for each global memory load
  - Only two can run on each SM
  - However, one only needs 200/(8*4) = 7 warps to tolerate the memory latency
  - Two blocks have 16 warps. The performance can be actually higher!

G80 Tiling Granularity Considerations

- For Matrix Multiplication, should I use 8X8, 16X16 or 32X32 tiles?
  - For 8X8, we have 64 threads per block. Since each SM can take up to 768 threads, it can take up to 12 blocks. However, each SM can only take up to 8 blocks, only 512 threads will go into each SM!
  - For 16X16, we have 256 threads per block. Since each SM can take up to 768 threads, it can take up to 3 blocks and achieve full capacity unless other resource considerations overrule.
  - For 32X32, we have 1024 threads per block. Not even one can fit into an SM!

G80 Shared Memory and Threading

- Each SM in G80 has 16KB shared memory
  - Dynamic resource: Blocks under the same SM share the 16KB, but different blocks can’t read data of others
  - SM size is implementation dependent!
  - For TILE_WIDTH = 16, each thread block uses 2*256*4B = 2KB of shared memory
  - TILE_WIDTH 32 would lead to 2*32*32*4B= 8KB shared memory usage per thread block, allowing only up to two thread blocks active at the same time

Tiling Size Effects

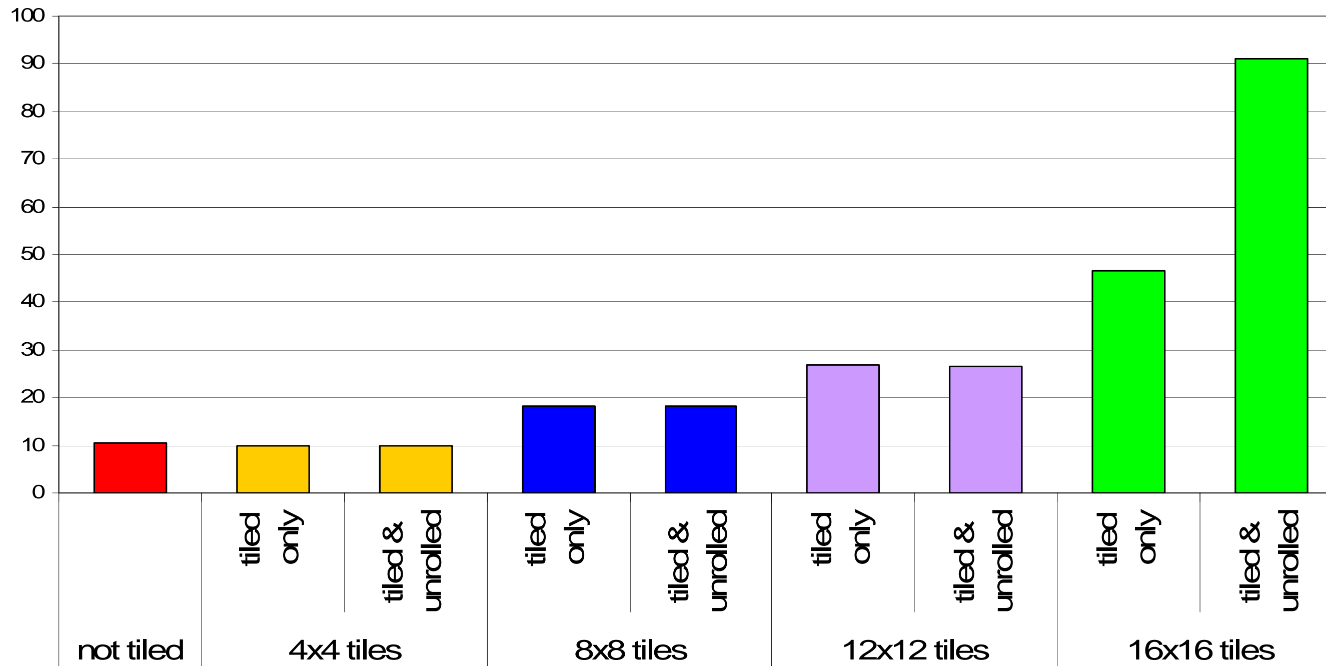

Threads

- Threads
- Warp
- Resource Allocation
- Control Flow
- Memory Coalescing

Control Flow Instructions

- Main performance concern with branching is divergence
  - Threads within a single warp take different paths
  - Different execution paths are serialized in G80
    - The control paths taken by the threads in a warp are traversed one at a time until there is no more

Control Flow Instructions

- A common case: avoid divergence when branch condition is a function of thread ID
  - Example with divergence:
    - If (threadIdx.x > 2) { }
    - This creates two different control paths for threads in a block
    - Branch granularity < warp size; threads 0, 1 and 2 follow different path than the rest of the threads in the first warp
  - Example without divergence:
    - If (threadIdx.x / WARP_SIZE > 2) { }
    - Also creates two different control paths for threads in a block
    - Branch granularity is a whole multiple of warp size
    - All threads in any given warp follow the same path

Parallel Reduction

- Given an array of values, "reduce" them to a single value in parallel
- Examples
  - sum reduction: sum of all values in the array
  - max/min reduction: maximum/minimum of all values in the array
- Typically parallel implementation
  - Recursively halve # threads, add two values per thread
  - Takes log(n) steps for n elements, requires n/2 threads

A Vector Reduction Example

- Use shared memory
  - The original vector is in device global memory
  - The shared memory used to hold a partial sum vector
  - Each iteration brings the partial sum vector closer to the final sum
  - The final solution will be in element 0

A simple implementation

- Assume we have already loaded array into

```
__shared__ float partialSum[blockDim.x]

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

Vector Reduction with Branch Divergence

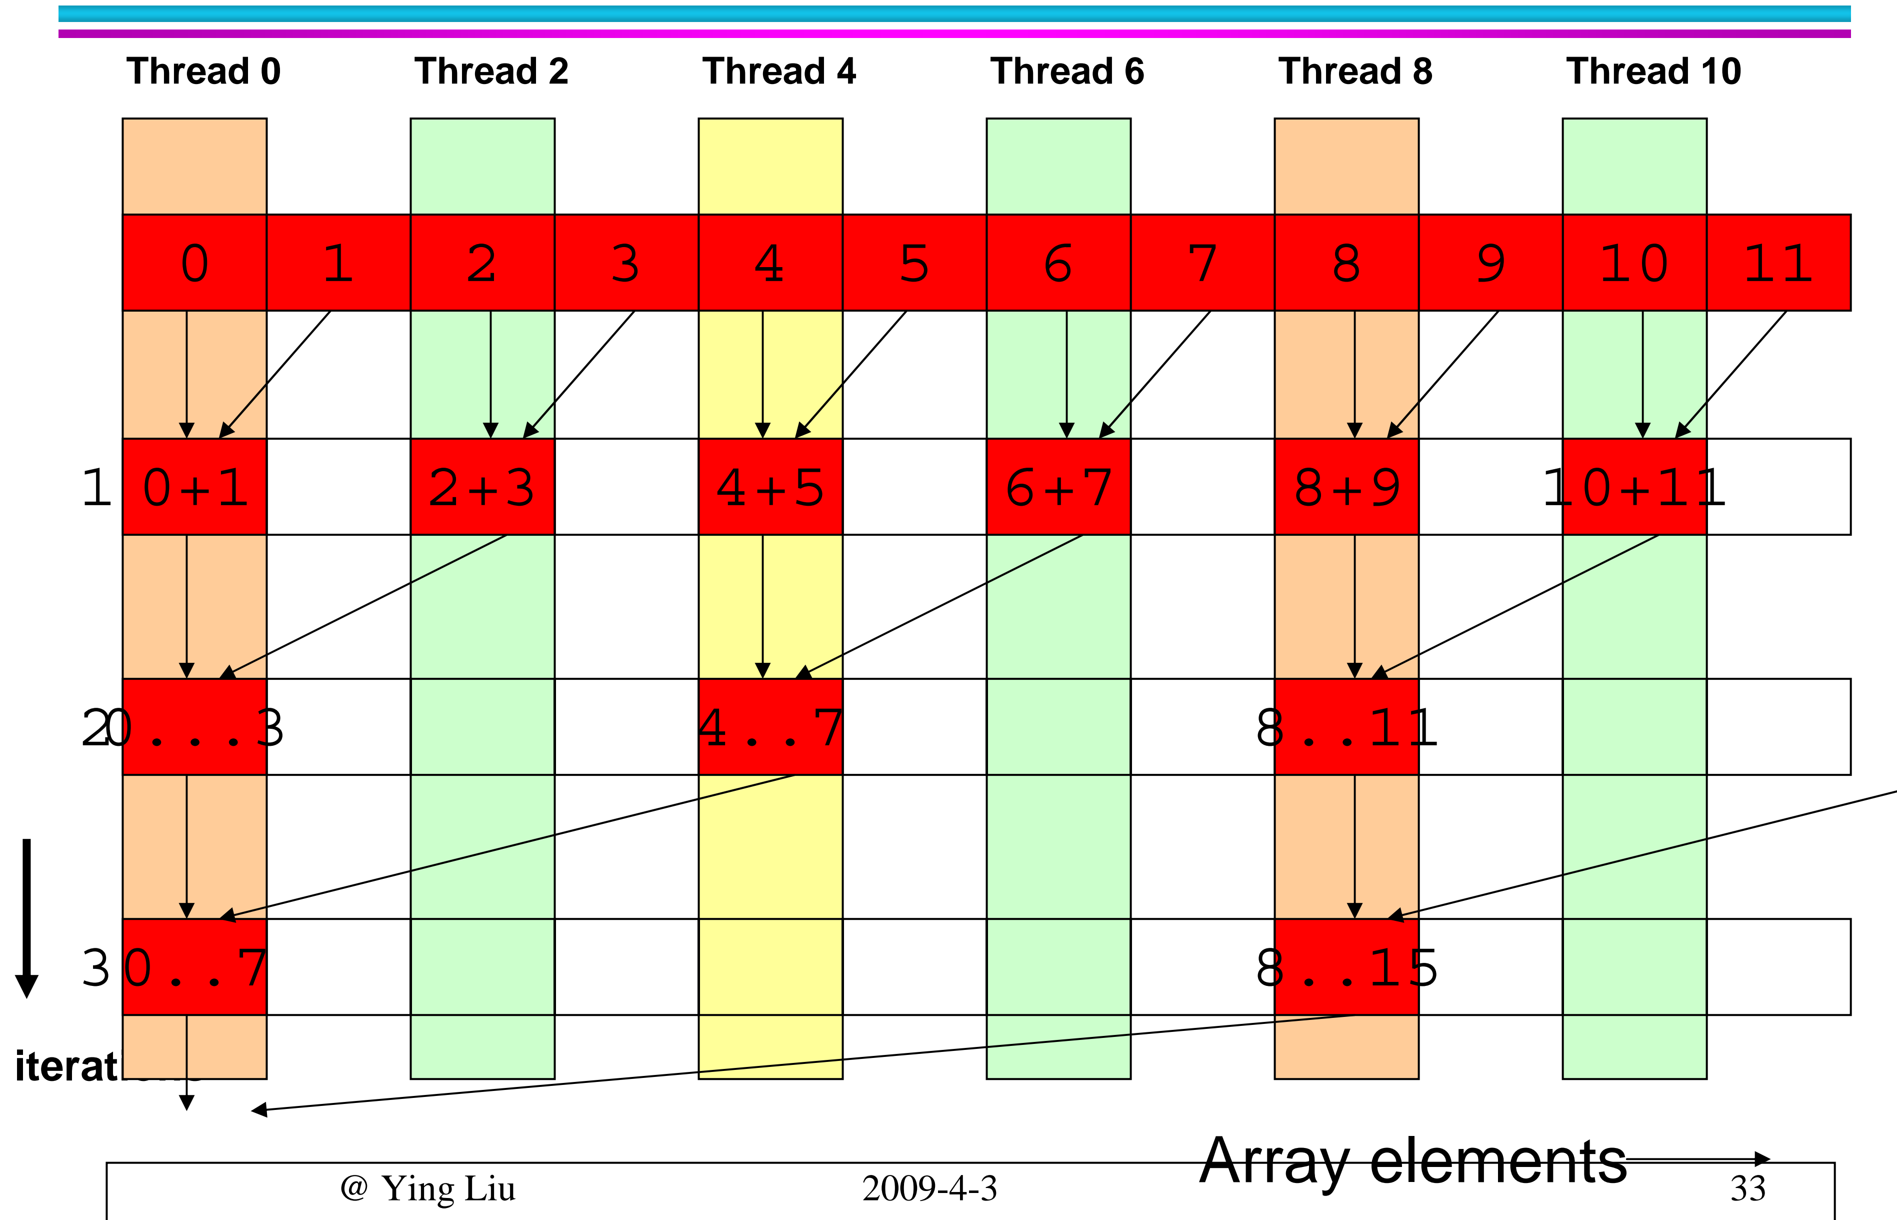

Some Observations

- In each iteration, two control flow paths will be sequentially traversed for each warp
  - Threads that perform addition and threads that do not
  - Threads that do not perform addition may cost extra cycles depending on the implementation of divergence
- No more than half of threads will be executing at any time
  - All odd index threads are disabled right from the beginning!
  - On average, less than ¼ of the threads will be activated for all warps over time
  - After the 5th iteration, entire warps in each block will be disabled, poor resource utilization but no divergence
    - This can go on for a while, up to 4 more iterations ($512/32=16=2^4$), where each iteration only has one thread activated until all warps retire

Shortcomings of the Implementation

- Assume we have already loaded array into
  __shared__ float partialSum[blockDim.x]

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

BAD: Divergence due to interleaved branch decisions

A Better Implementation

- Assume we have already loaded array into

```
__shared__ float partialSum[blockDim.x]

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

No Divergence Until < 16 Sub-sums

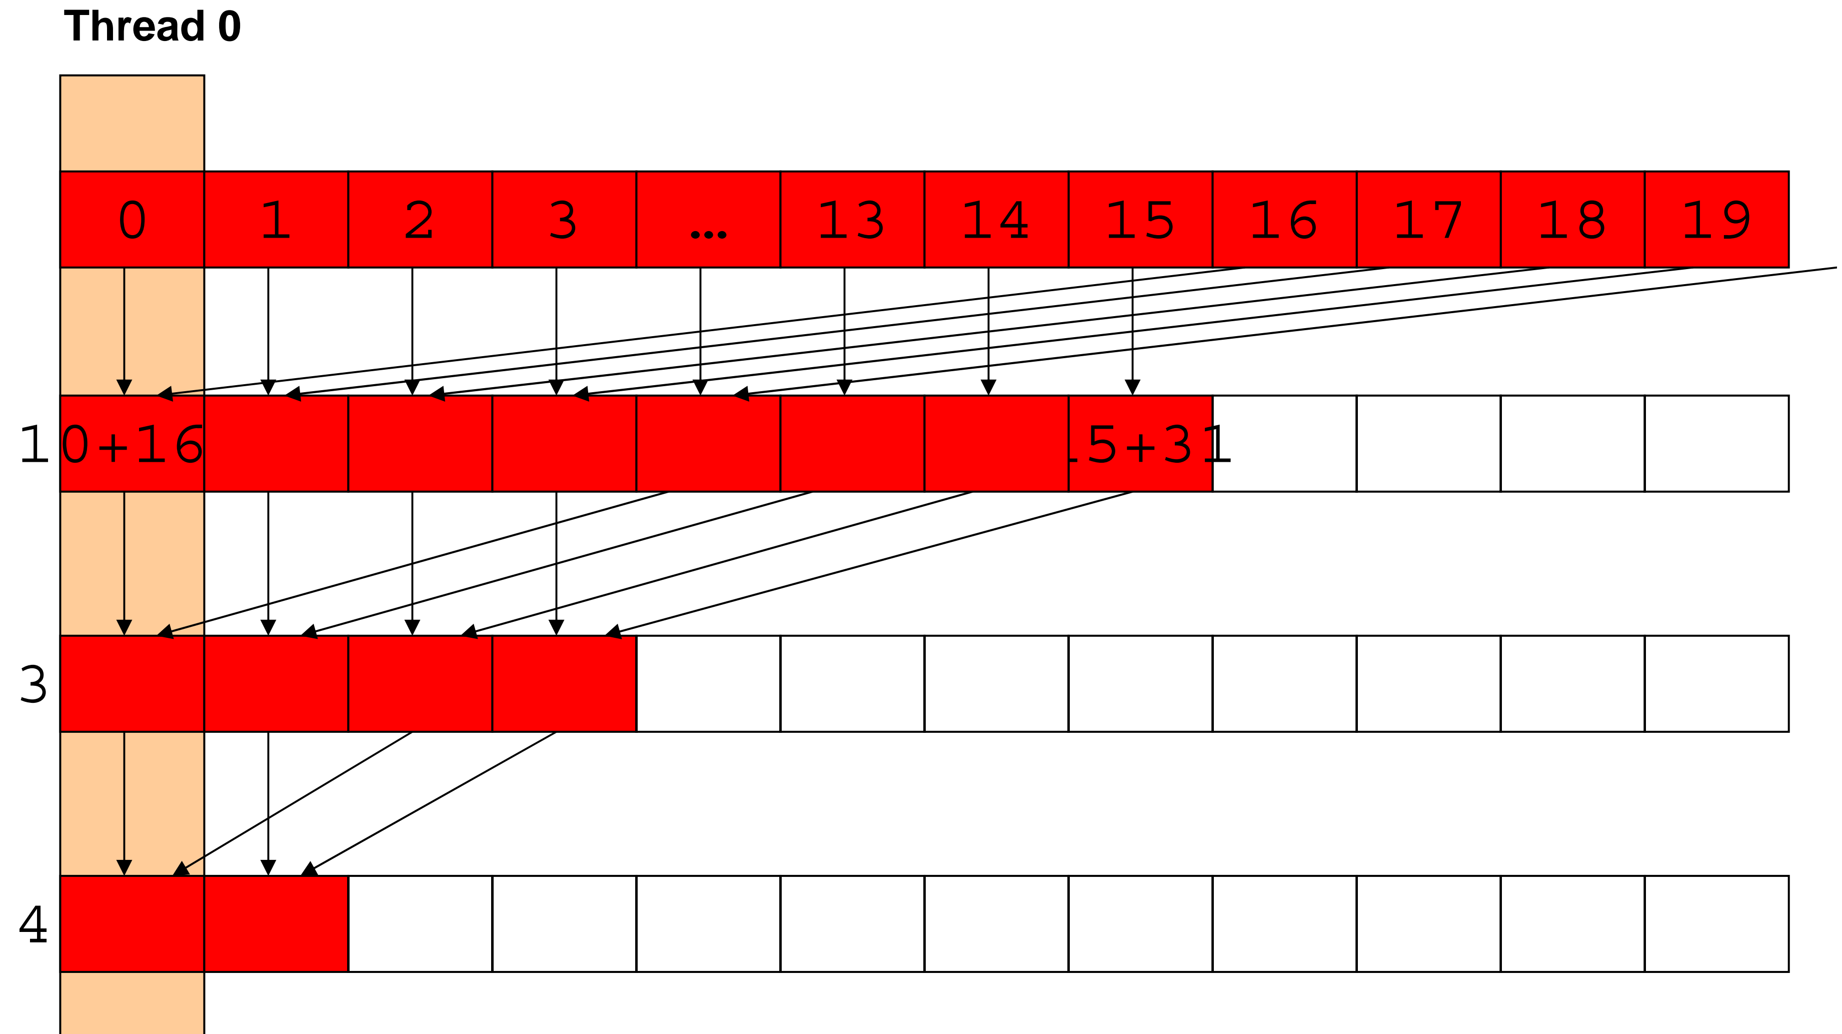

Threads

- Threads
- Warp
- Resource Allocation
- Control Flow
- Memory Coalescing

Memory Bandwidth

- Device memory should be minimized
  - Device memory is much higher latency and lower bandwidth than on-chip memory
- Global memory bandwidth is used most efficiently when the simultaneous memory accesses by threads in a half-warp (16 threads) can be coalesced into a single memory transaction
- A memory transaction can be 32 bytes (compute capability 1.2+), 64 bytes, 128 bytes

Memory Bandwidth

- Device reads 32-bit, 64-bit or 128-bit words from global memory into registers in a single transaction
  - Data type must be sizeof(type) = 4, 8, 16 bytes
  - Data type must be aligned
- The alignment specifiers _align_(8) or _align_(16)

Memory Bandwidth

```
Struct _align_ (8) {
    float a;
    float b;
}
Struct _align_ (16) {
    float a;
    float b;
    float c;
}
Struct _align_ (16) {
    float a;
    float b;
    float c;
    float d;
    float e;
}
```

Global Memory

Coalesced Access

Global Memory

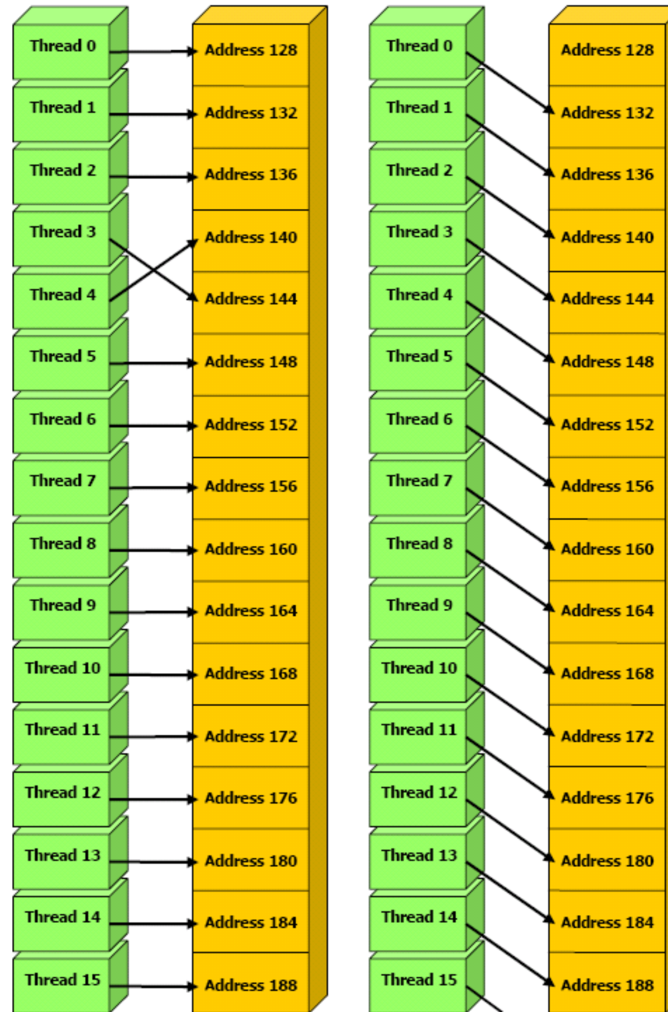

Non-coalesced Access

Global Memory

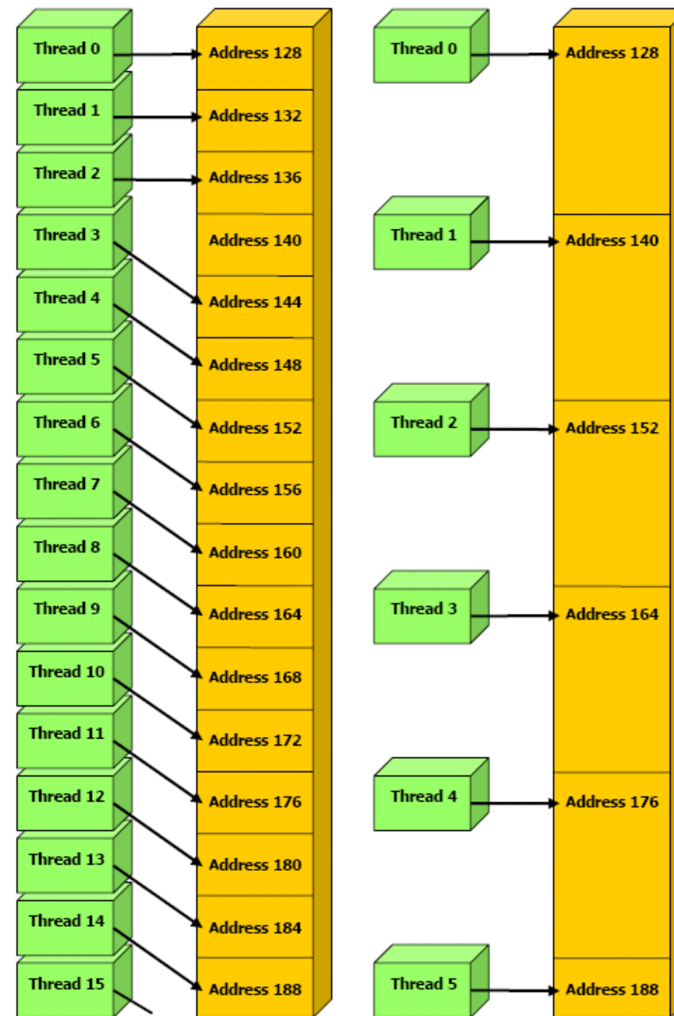

Non-coalesced Access

Memory Access Pattern

Original Access Pattern

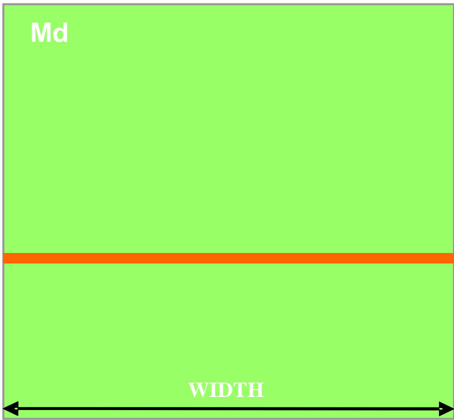

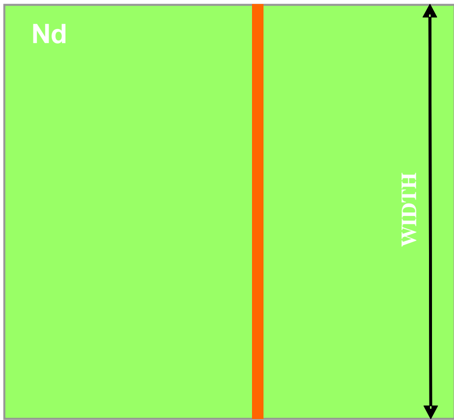

Tiled Access Pattern

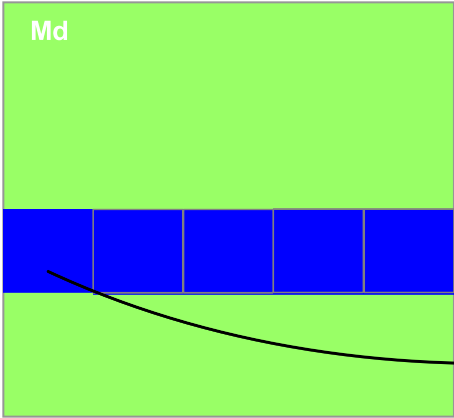

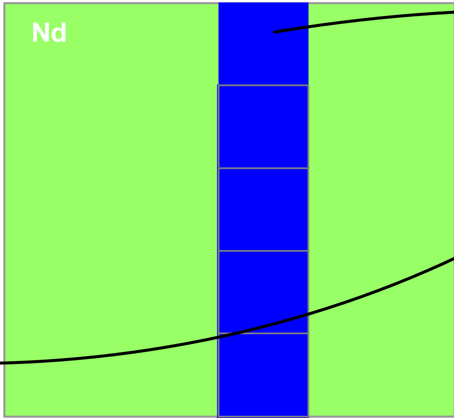

Copy into scratchpad memory

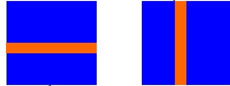

Perform multiplication with scratchpad values

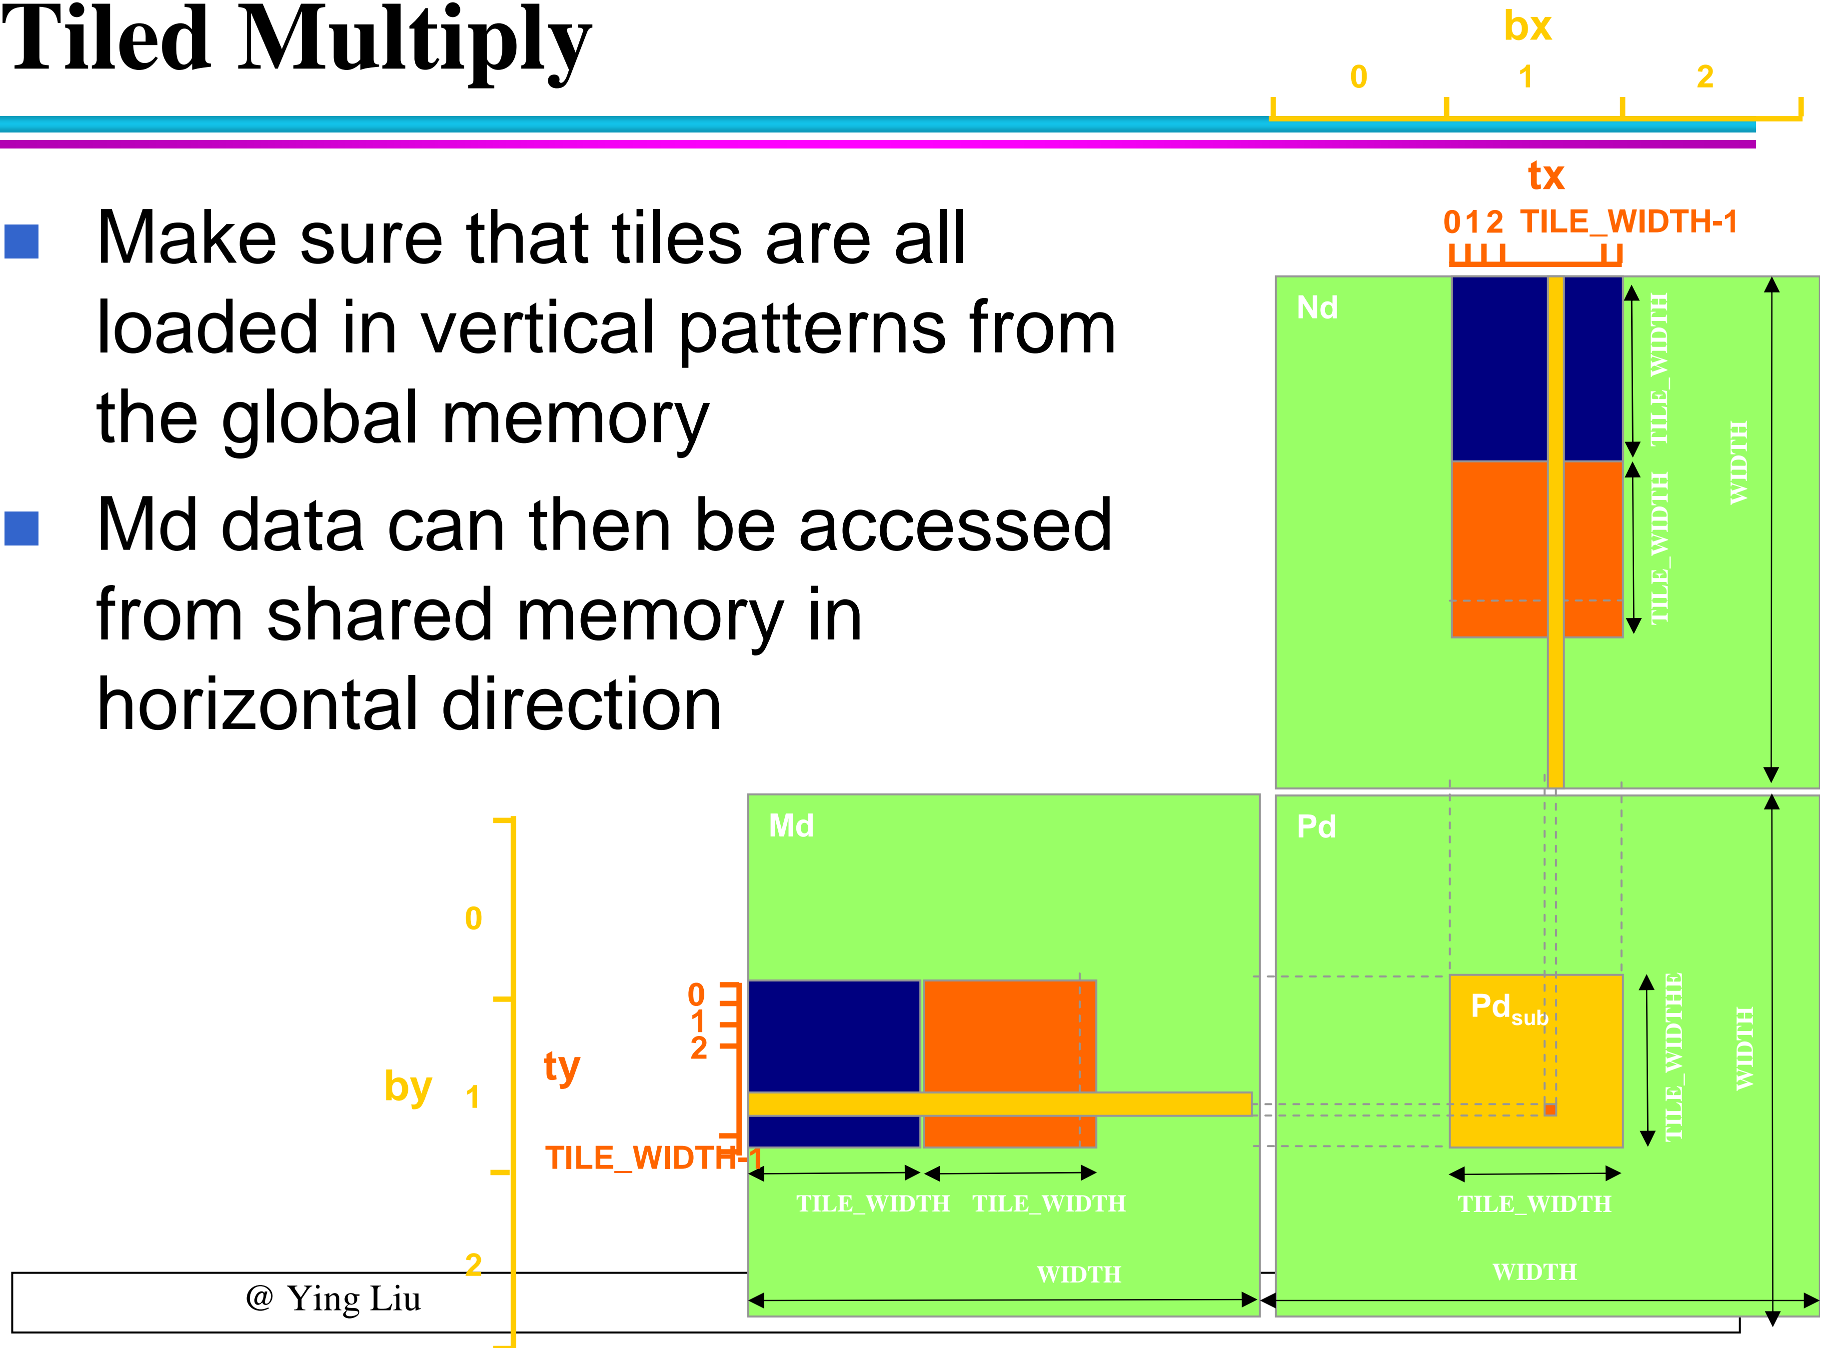
Tiled Multiply
Make sure that tiles are all loaded in vertical patterns from the global memory
Md data can then be accessed from shared memory in horizontal direction
bx
0
1
2
tx
0 1 2 TILE_WIDTH-1
Nd
TILE_WIDTH
TILE_WIDTH
WIDTH
Md
Pd
Pdsub
TILE_WIDTHE
WIDTH
by
0
1
2
ty
0
1
2
TILE_WIDTH-1
TILE_WIDTH
TILE_WIDTH
WIDTH
TILE_WIDTH
WIDTH
@ Ying Liu

Tiling Size Effects

- For good bandwidth utilization, accesses should be aligned and consist of 16 contiguous words
- Tile size 16X16 minimal required to achieve full coalescing
  - Both reduction of global memory accesses and more efficient execution of the accesses