



CUDA Performance

Patrick Cozzi
University of Pennsylvania
CIS 565 - Fall 2012

1

Announcements

- Project 2 due Friday 10/12
 - Be ready to present on Monday, 10/15
- Move class on Halloween to Tuesday, 10/30?

2

Acknowledgements

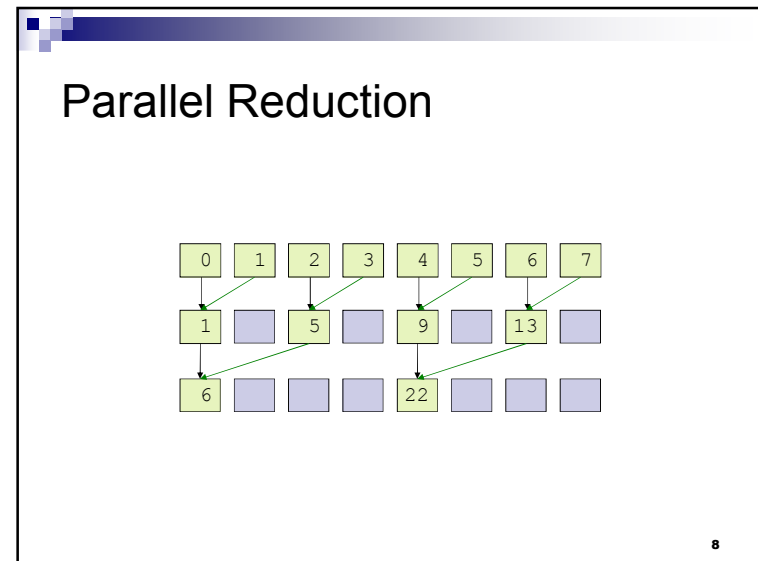
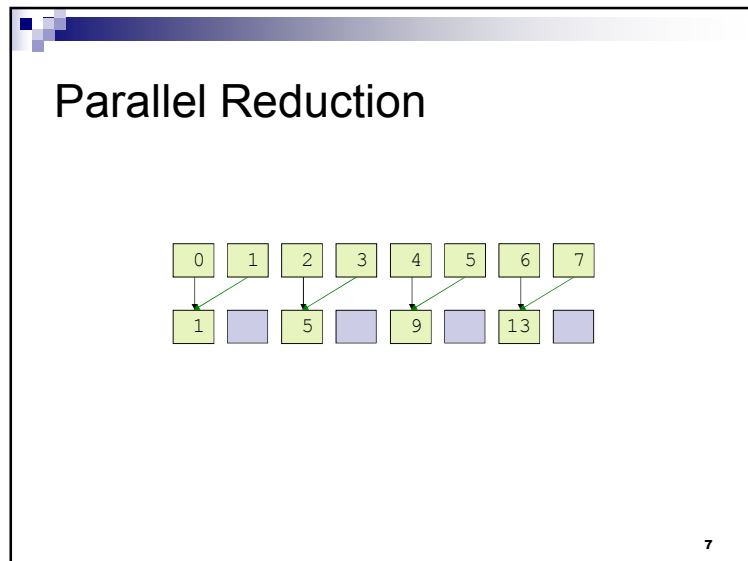
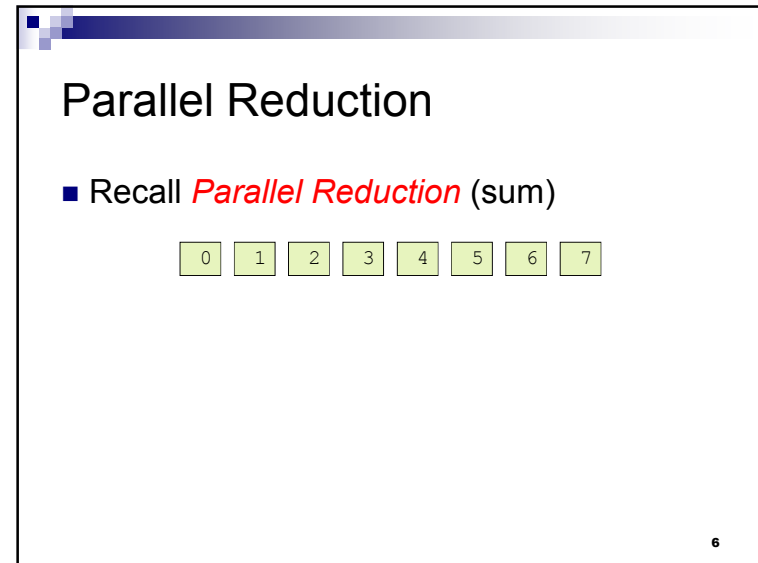
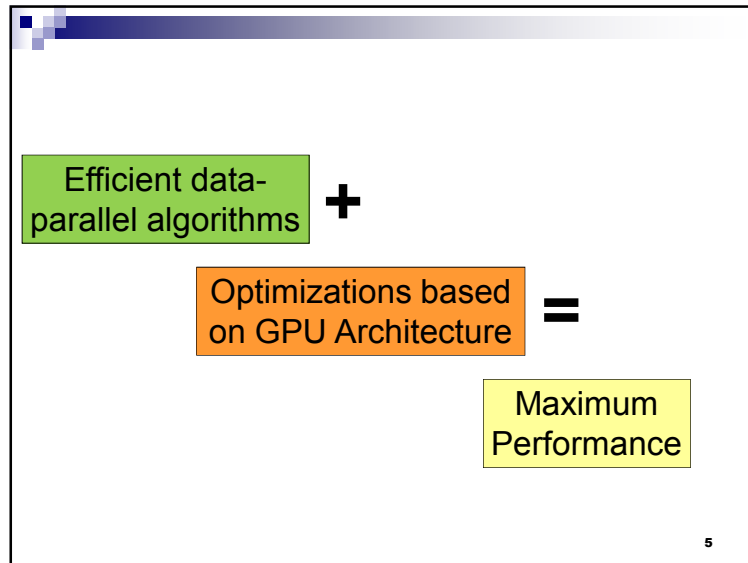
- Some slides from [Varun Sampath](#)

3


Agenda

- Parallel Reduction Revisited
- Warp Partitioning
- Memory Coalescing
- Bank Conflicts
- Dynamic Partitioning of SM Resources
- Data Prefetching
- Instruction Mix
- Loop Unrolling

4




Parallel Reduction



9

Parallel Reduction

- Similar to brackets for a basketball tournament
- $\log(n)$ passes for n elements
- How would you implement this in CUDA?



10

```

__shared__ float partialSum[];
// ... load into shared memory
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];
}
    
```

Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

11

```

__shared__ float partialSum[];
// ... load into shared memory
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];
}
    
```

Computing the sum for the elements in shared memory

Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

12

```

__shared__ float partialSum[];
// ... load into shared memory
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];
}

```

Stride:
1, 2, 4, ...

13
Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

```

__shared__ float partialSum[];
// ... load into shared memory
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];
}

```

Why?

14
Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

```

__shared__ float partialSum[];
// ... load into shared memory
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];
}

```

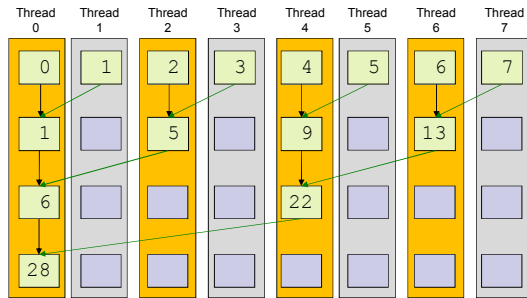
• Compute sum in same shared memory
• As stride increases, what do more threads do?

15
Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Parallel Reduction

16

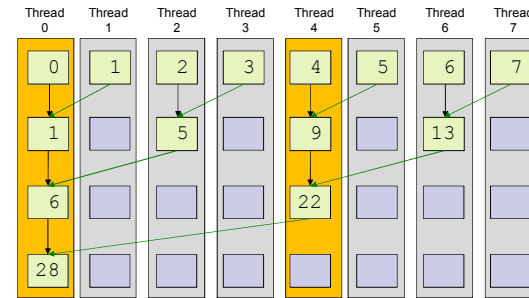
Parallel Reduction



- 1st pass: threads 1, 3, 5, and 7 don't do anything
- Really only need $n/2$ threads for n elements

17

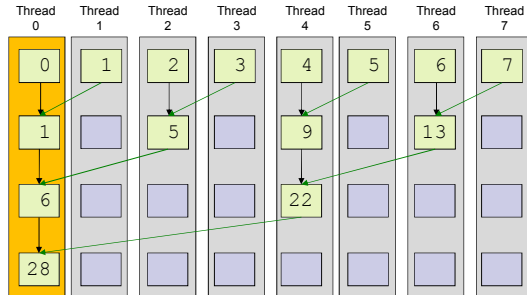
Parallel Reduction



- 2nd pass: threads 2 and 6 also don't do anything

18


Parallel Reduction



- 3rd pass: thread 4 also doesn't do anything

19

Parallel Reduction



- In general, number of required threads cuts in half after each pass


20

Parallel Reduction

- What if we *tweaked* the implementation?


21

Parallel Reduction




22

Parallel Reduction




23

Parallel Reduction



24


Parallel Reduction



25

```

__shared__ float partialSum[]
// ... load into shared memory
unsigned int t = threadIdx.x;
for(unsigned int stride = blockDim.x / 2;
    stride > 0;
    stride /= 2)
{
    __syncthreads();
    if (t < stride)
        partialSum[t] +=
            partialSum[t + stride];
}
    
```



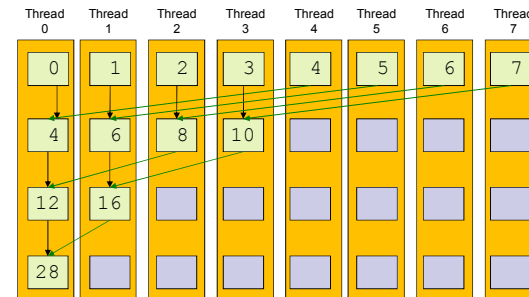
Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html> 26

```

__shared__ float partialSum[]
// ... load into shared memory
unsigned int t = threadIdx.x;
for(unsigned int stride = blockDim.x / 2;
    stride > 0;
    stride /= 2)
{
    __syncthreads();
    if (t < stride)
        partialSum[t] +=
            partialSum[t + stride];
}
    
```

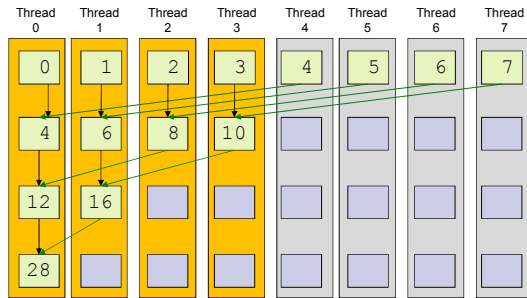
Code from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html> 27

Parallel Reduction



28

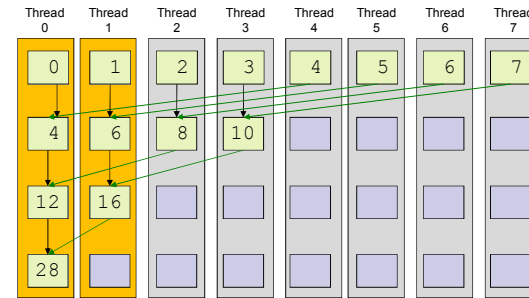
Parallel Reduction



- 1st pass: threads 4, 5, 6, and 7 don't do anything
- Really only need $n/2$ threads for n elements

29

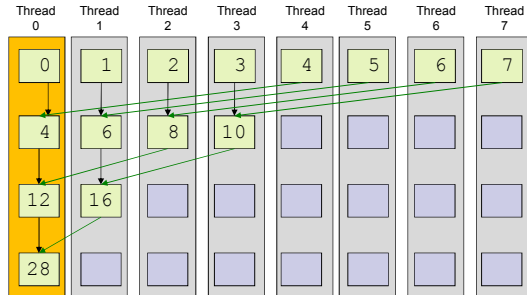
Parallel Reduction



- 2nd pass: threads 2 and 3 also don't do anything

30

Parallel Reduction




- 3rd pass: thread 1 also doesn't do anything

31

Parallel Reduction

- What is the difference?



32

Parallel Reduction

- What is the difference?

```
if (t % (2 * stride) == 0)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 1, 2, 4, ...

```
if (t < stride)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 4, 2, 1, ...

33

Warp Partitioning

- **Warp Partitioning**: how threads from a block are divided into warps
- Knowledge of warp partitioning can be used to:
 - Minimize divergent branches
 - Retire warps early

34


Warp Partitioning

- Partition based on **consecutive increasing threadIdx**

35

Warp Partitioning

- 1D Block
 - `threadIdx.x` between 0 and 512 (G80/GT200)
 - Warp n
 - Starts with thread $32n$
 - Ends with thread $32(n + 1) - 1$
 - Last warp is padded if block size is not a multiple of 32



36

Warp Partitioning

■ 2D Block

- Increasing `threadIdx` means
 - Increasing `threadIdx.x`
 - Starting with row `threadIdx.y == 0`

37

Warp Partitioning

■ 2D Block




Image from <http://courses.engr.illinois.edu/ece498/al/textbook/Chapter5-CudaPerformance.pdf> 38

Warp Partitioning

■ 3D Block

- Start with `threadIdx.z == 0`
- Partition as a 2D block
- Increase `threadIdx.z` and repeat

39

Warp Partitioning

Divergent branches are within a warp!

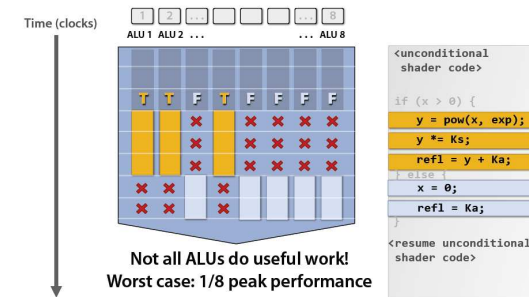


Image from: http://bps10.idav.ucdavis.edu/talks/03-fatahalian_gpuArchTeraflop_BPS_SIGGRAPH2010.pdf 40

Warp Partitioning

- For `warpSize == 32`, does any warp have a divergent branch with this code:

```
if (threadIdx.x > 15)
{
    // ...
}
```

41

Warp Partitioning

- For any `warpSize > 1`, does any warp have a divergent branch with this code:

```
if (threadIdx.x > warpSize - 1)
{
    // ...
}
```

42

Warp Partitioning

- Given knowledge of warp partitioning, which parallel reduction is better?

```
if (t % (2 * stride) == 0)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 1, 2, 4, ...


```
if (t < stride)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 4, 2, 1, ...


43

Warp Partitioning

- Pretend `warpSize == 2`



stride = 1, 2, 4, ...




stride = 4, 2, 1, ...

44

Warp Partitioning


1st Pass



45

Warp Partitioning


2nd Pass



46

Warp Partitioning

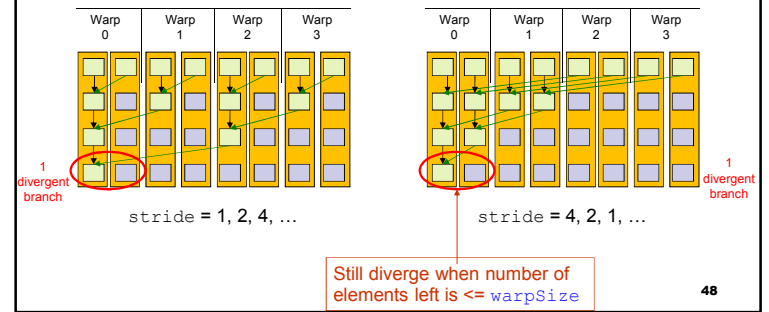
2nd Pass



47

Warp Partitioning

2nd Pass



48

Warp Partitioning

- Good partitioning also allows warps to be retired early.

- Better hardware utilization

```
if (t % (2 * stride) == 0)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 1, 2, 4, ...


```
if (t < stride)
    partialSum[t] +=
        partialSum[t + stride];
```

stride = 4, 2, 1, ...


49

Warp Partitioning

- Parallel Reduction



stride = 1, 2, 4, ...




stride = 4, 2, 1, ...


50

Warp Partitioning

- 1st Pass



stride = 1, 2, 4, ...




stride = 4, 2, 1, ...


51

Warp Partitioning

- 1st Pass



stride = 1, 2, 4, ...




stride = 4, 2, 1, ...

52

Warp Partitioning


2nd Pass



53

Warp Partitioning


2nd Pass



54

Memory Coalescing

- Given a matrix stored *row-major* in *global memory*, what is a *thread's* desirable access pattern?




55

Image from: http://bpsi10.idav.ucdavis.edu/talks/03-fatahalian_gpuArchTeraflop_BPS_SIGGRAPH2010.pdf

Memory Coalescing

- Given a matrix stored *row-major* in *global memory*, what is a *thread's* desirable access pattern?



56

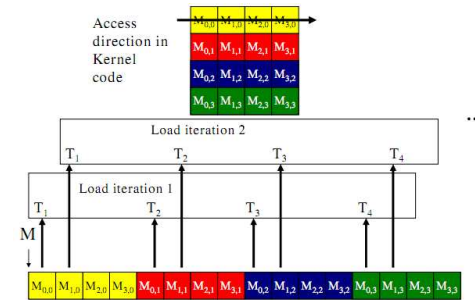
Image from: <http://courses.engr.illinois.edu/ece498/all/textbook/Chapter5-CudaPerformance.pdf>

Memory Coalescing

- Given a matrix stored *row-major* in *global memory*, what is a *thread's* desirable access pattern?
 - a) column after column
 - *Individual threads* read increasing, consecutive memory address
 - b) row after row
 - *Adjacent threads* read increasing, consecutive memory addresses

57

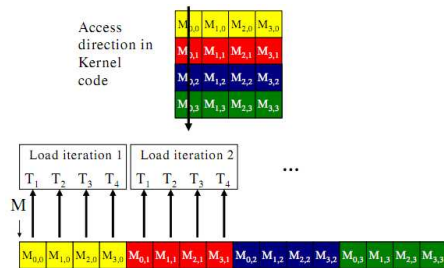
Memory Coalescing



a) column after column

Image from: <http://courses.engr.illinois.edu/ece498/al/textbook/Chapter5-CudaPerformance.pdf> 58


Memory Coalescing



b) row after row

Image from: <http://courses.engr.illinois.edu/ece498/al/textbook/Chapter5-CudaPerformance.pdf> 59

Memory Coalescing



Recall warp partitioning; if these threads are in the same warp, global memory addresses are increasing and consecutive across warps.

Image from: <http://courses.engr.illinois.edu/ece498/al/textbook/Chapter5-CudaPerformance.pdf> 60

Memory Coalescing

- Global memory bandwidth (DRAM)
 - G80 – 86.4 GB/s
 - GT200 – 150 GB/s
- Achieve peak bandwidth by requesting large, consecutive locations from DRAM
 - Accessing random location results in much lower bandwidth

61

Memory Coalescing

- *Memory coalescing* – rearrange access patterns to improve performance
- Useful today but will be less useful with large on-chip caches

62

Memory Coalescing


- The GPU coalesce consecutive reads in a *half-warp* into a single read
- *Strategy*: read global memory in a coalesce-able fashion into shared memory
 - Then access shared memory randomly at maximum bandwidth
 - Ignoring *bank conflicts*...

63

See Appendix G in the NVIDIA CUDA C Programming Guide for coalescing alignment requirements

Bank Conflicts

- Shared Memory
 - Sometimes called a *parallel data cache*
 - Multiple threads can access shared memory at the same time
 - Memory is divided into *banks*




64

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

■ Banks

- Each bank can service one address per two cycles
- Per-bank bandwidth: 32-bits per two (shader clock) cycles
- Successive 32-bit words are assigned to successive banks



65

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts


- **Bank Conflict:** Two simultaneous accesses to the same bank, but not the same address

- Serialized

- G80-GT200: 16 banks, with 8 SPs concurrently executing

- Fermi: 32 banks, with 16 SPs concurrently executing

- What does this mean for conflicts?




66

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts


■ Bank Conflicts?

- Linear addressing stride == 1



■ Bank Conflicts?

- Random 1:1 Permutation




67

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts


■ Bank Conflicts?

- Linear addressing stride == 2



■ Bank Conflicts?

- Linear addressing stride == 8




68

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

Fast Path 1 (G80)

- All threads in a half-warp access different banks




69

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

Fast Path 2 (G80)

- All threads in a half-warp access the same address




70

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

Slow Path (G80)

- Multiple threads in a half-warp access the same bank
- Access is serialized
- What is the cost?



71

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

```
__shared__ float shared[256];  
// ...  
float f = shared[index + s * threadIdx.x];
```


For what values of s is this conflict free?

- Hint: The G80 has 16 banks

72

Bank Conflicts

```
__shared__ float shared[256];  
// ...  
float f = shared[index + s * threadIdx.x];
```



73

Image from <http://courses.engr.illinois.edu/ece498/al/Syllabus.html>

Bank Conflicts

- Without using a profiler, how can we tell what kind of speedup we can expect by removing bank conflicts?
- What happens if more than one thread in a warp writes to the same shared memory address (non-atomic instruction)?

74

75

76



SM Resource Partitioning

- Recall a SM dynamically partitions resources:

79

SM Resource Partitioning


- Recall a SM dynamically partitions resources:

<u>G80 Limits</u>	
Thread block slots	8
Thread slots	768
Registers	8K registers / 32K memory
Shared memory	16K

80

SM Resource Partitioning


- We can have
 - 8 blocks of 96 threads
 - 4 blocks of 192 threads
 - But not 8 blocks of 192 threads



81

SM Resource Partitioning

- We can have (assuming 256 thread blocks)
 - 768 threads (3 blocks) using 10 registers each
 - 512 threads (2 blocks) using 11 registers each




82

SM Resource Partitioning

- We can have (assuming 256 thread blocks)
 - 768 threads (3 blocks) using 10 registers each
 - 512 threads (2 blocks) using 11 registers each

- More registers decreases thread-level parallelism
 - Can it ever increase performance?



83

SM Resource Partitioning

- **Performance Cliff:** Increasing resource usage leads to a dramatic reduction in parallelism
 - For example, increasing the number of registers, unless doing so hides latency of global memory access

84

SM Resource Partitioning

- CUDA Occupancy Calculator

- http://developer.download.nvidia.com/compute/cuda/CUDA_Occupancy_calculator.xls

85

Data Prefetching

- Independent instructions between a global memory read and its use can hide memory latency

```
float m = Md[i];  
float f = a * b + c * d;  
float f2 = m * f;
```

86

Data Prefetching

- Independent instructions between a global memory read and its use can hide memory latency

```
float m = Md[i];  
float f = a * b + c * d;  
float f2 = m * f;
```

← Read global memory

87

Data Prefetching

- Independent instructions between a global memory read and its use can hide memory latency

```
float m = Md[i];  
float f = a * b + c * d;  
float f2 = m * f;
```

↑
Execute instructions
that are not dependent
on memory read

88

Data Prefetching

- Independent instructions between a global memory read and its use can hide memory latency

```
float m = Md[i];  
float f = a * b + c * d;  
float f2 = m * f;
```

Use global memory after the above line from enough warps hide the memory latency

89

Data Prefetching

- Prefetching** data from global memory can effectively increase the number of independent instructions between global memory read and use

90

Data Prefetching

- Recall tiled matrix multiply:

```
for (/* ... */) {  
    // Load current tile into shared memory  
    __syncthreads();  
    // Accumulate dot product  
    __syncthreads();  
}
```

91

Data Prefetching

- Tiled matrix multiply with prefetch:

```
// Load first tile into registers  
for (/* ... */) {  
    // Deposit registers into shared memory  
    __syncthreads();  
    // Load next tile into registers  
    // Accumulate dot product  
    __syncthreads();  
}
```

92

Data Prefetching

- Tiled matrix multiply with prefetch:

```
// Load first tile into registers

for (/* ... */)
{
    // Deposit registers into shared memory
    __syncthreads();
    // Load next tile into registers
    // Accumulate dot product
    __syncthreads();
}
```

93

Data Prefetching

- Tiled matrix multiply with prefetch:

```
// Load first tile into registers

for (/* ... */)
{
    // Deposit registers into shared memory
    __syncthreads();
    // Load next tile into registers
    // Accumulate dot product
    __syncthreads();
}
```

Prefetch for next iteration of the loop

94

Data Prefetching

- Tiled matrix multiply with prefetch:

```
// Load first tile into registers

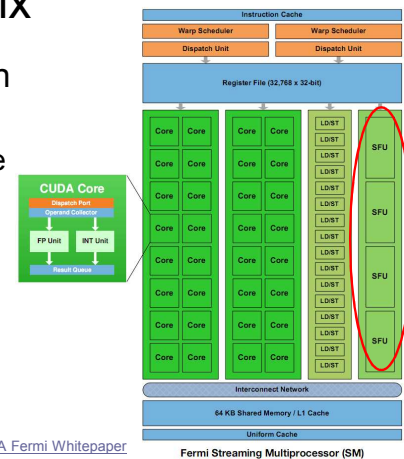
for (/* ... */)
{
    // Deposit registers into shared memory
    __syncthreads();
    // Load next tile into registers
    // Accumulate dot product
    __syncthreads();
}
```

These instructions executed by enough threads will hide the memory latency of the prefetch

95

Instruction Mix

- Special Function Units (SFUs)
 - Use to compute `__sinf()`, `__expf()`
 - Only 4, each can execute 1 instruction per clock



96

Image: NVIDIA Fermi Whitepaper

Fermi Streaming Multiprocessor (SM)

Loop Unrolling

```
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

■ Instructions per iteration

- One floating-point multiply
- One floating-point add
- What else?

97

Loop Unrolling

```
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

■ Other instructions per iteration

- Update loop counter

98

Loop Unrolling

```
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

■ Other instructions per iteration

- Update loop counter
- Branch

99

Loop Unrolling

```
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

■ Other instructions per iteration

- Update loop counter
- Branch
- Address arithmetic

100

Loop Unrolling

```
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

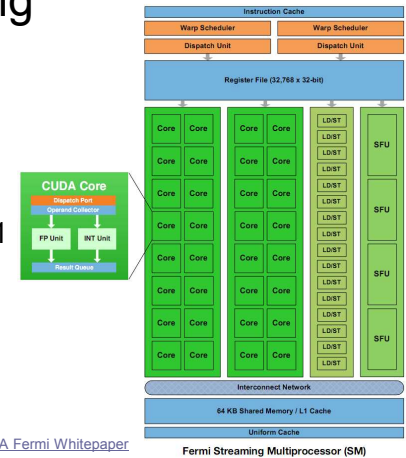
■ Instruction Mix

- 2 floating-point arithmetic instructions
- 1 loop branch instruction
- 2 address arithmetic instructions
- 1 loop counter increment instruction

101

Loop Unrolling

- Only 1/3 are floating-point calculations
 - But I want my full theoretical 1 TFLOP (Fermi)
 - Consider *loop unrolling*



102

Image: [NVIDIA Fermi Whitepaper](#)

Fermi Streaming Multiprocessor (SM)

Loop Unrolling

```
Pvalue +=
Ms[ty][0] * Ns[0][tx] +
Ms[ty][1] * Ns[1][tx] +
...
Ms[ty][15] * Ns[15][tx]; // BLOCK_SIZE = 16
```

- No more loop
 - No loop count update
 - No branch
 - Constant indices – no address arithmetic instructions

103

Loop Unrolling

■ Automatically:

```
#pragma unroll BLOCK_SIZE
for (int k = 0; k < BLOCK_SIZE; ++k)
{
    Pvalue += Ms[ty][k] * Ns[k][tx];
}
```

■ Disadvantages to unrolling?

104