

CUDA: Performance Considerations

Mark Greenstreet

CpSc 418 – November 23, 2018

- Occupancy
- Thread Divergence
- Instruction Mix



Unless otherwise noted or cited, these slides are copyright 2018 by Mark Greenstreet and are made available under the terms of the Creative Commons Attribution 4.0 International license <http://creativecommons.org/licenses/by/4.0/>

First a quick summary of the last lecture

- Floating point operation:
 - ▶ GPU single-precision floating point is much faster than double-precision
 - ▶ Use `floats` instead of `doubles` when you can
 - ★ I would give the opposite advice for programming an x86.
 - ★ Pay attention when using libraries.
 - ▶ Remember the `f` when writing floating point constants: e.g. `3.14159265f`.
- Memory
 - ▶ Estimate your CGMA – global accesses are probably your bottleneck
 - ▶ Coalesce global memory accesses:
 - ★ Warp locality matters.
 - ★ Temporal locality, not so much.
 - ▶ Avoid shared memory branch conflicts
 - ★ Think about how array indices connect to thread indices.

Measuring Time

- From piazza, but worth repeating to make sure everyone gets it.
- Adjust `ntrials` for `saxpy` or `m` for `recur1` and `recur2` to get a `total_time` between 10ms and 0.25 seconds.
 - ▶ 1 ms = 1 millisecond = $1.0e-03$ seconds.
 - ▶ 10 ms = $1.0e-02$ seconds.
 - ▶ 0.25 seconds = $2.5e-01$ seconds.
- If you have a `total_time` of 0 seconds, or $1.5e-05$ seconds, then you are seeing the granularity of the linux system clock and your have lots of measurement error.
 - ▶ To get accurate, repeatable measurements, aim for 100–1000 or more “OS clock” ticks.
- If you have a `total_time` of 1 or 2 seconds, you may be messing with the frame refresh for a user on the console.
 - ▶ Be a good citizen.

Occupancy

- We want lots of threads running on each SM. How do we get this?
- To keep every SM busy, the number of blocks should be greater than the number of SMs.
 - ▶ An SM can have 2048 threads, but a block can have at most 1024 threads.
 - ▶ We want at least two blocks per SM. More is better.
- Limits on occupancy
 - ▶ At most 8 blocks per SM.
 - ▶ 2^{16} registers/SM. Use
`nvcc -O3 -c --ptxas-options -v examples.cu`
 - ▶ 96Kbytes of shared memory per SM. At most 48Kbytes/block.
 - ▶ Check out the occupancy calculator:
[CUDA_Occupancy_calculator.xls](#)

Predicated Execution

- What about `if` (or `for`)?
- When a warp executes the `if`, all threads execute it?
 - ▶ What if some threads take the then-branch and others take the else-branch?
- Predicated execution.
 - ▶ Each thread keeps a bit saying whether it's a then-thread or an else-thread.
 - ▶ The GPU fetches instructions from the then and else clauses, and marks them as then-instructions or else-instructions
 - ▶ The then-threads execute then-instructions and ignore else-instructions.
 - ★ The else-instructions are still fetched and dispatched for the SPs that need them.
 - ★ They are treated as NOPs by the SPs for then-threads, but they take up execution time.
 - ▶ Likewise, the else-threads execute else-instructions and ignore then-instructions.
- Predication handles multiple levels of control flow nesting.

Thread Divergence and Performance

Thread Divergence and `__syncthreads()`

Instruction Mix

- We measure our program performance in terms of the critical, unavoidable operations
 - ▶ Typically “floating point operations” for matrix-multiplication or other scientific computing applications.
 - ▶ Often main memory accesses for sorting, or other data-intensive applications.
- But, the program does other operations as well
 - ▶ This is where you see me counting instructions on my fingers during lecture.
 - ▶ Optimizing performance can involve minimizing this overhead:
 - ★ Good algorithm design.
 - ★ Memory access optimization.
 - ★ Loop unrolling

Bigger Kernels

```
--global__myKernel (...) {  
    do something  
}
```

Unless *do something* is big, kernel launch takes most of the time.

- We can launch a big-grid
 - ▶ If we have a huge number of array elements than each need a small amount of work, this can be a good idea.
 - ▶ **BUT** we're likely to create a memory-bound problem.
- Or, we can make each thread do many somethings.

```
--global__myKernel(int m, ...) {  
    for(int i = 0; i < m; i++)  
        do something  
}
```

Loop Limitations

- It takes two or three instructions per loop iteration to manage the loop:
 - ▶ One to update the loop index
 - ▶ One or two to check the loop bounds and branch.
 - ▶ If *do something* is only three or four instructions, then 40-50% of the execution time is for loop management.
- If each iteration of *do something* depends on the previous one
 - ▶ Then the long latency of the SP pipelines can limit performance.
 - ▶ Even if we have 48 warps running.



Loop Unrolling

- Have each loop iteration perform multiple copies of the loop body

```
__global__ myKernel(int m, ...) {  
    for(int i = 0; i < m; i += 4) {  
        do something 1  
        do something 2  
        do something 3  
        do something 4  
    }  
}
```

- More “real work” for each time the loop management code is executed.
- Need to make sure that `m` is a multiple of four, or handle end-cases separately.
- Often, we need more registers.

Unrolling – the plots



This example is from [2015W2 HW3](#), Q1.

Preview

November 23: CUDA Performance, Part 2

November 26: CUDA Performance/Examples

November 28: CUDA (and other) Examples

November 30: Review

Dec 3: Final Exam review session (to be scheduled)

Dec 4: Final Exam – 8:38am (ouch!)