

CUDA optimization

Philip Blakely

Laboratory for Scientific Computing, Cambridge

Outline

- 1 Should we optimize
- 2 Memory layout
- 3 Streaming Multiprocessor usage
- 4 Code optimizations
- 5 Which optimizations
- 6 Larger CUDA programs

To consider first...

Is it worth optimizing it?

- In scientific computing, answer usually “yes” ...
- unless it takes a lot of effort to implement ...
- or you're only going to use it once.
- But - only optimize actual bottlenecks
- No point in optimizing part of the code that takes 1% of the run-time.
- Use NVIDIA Nsight `nsys-ui` to check where bottle-necks occur.
- (Slightly tricky to set up; need to match CUDA Driver with NSight exactly, and set some kernel permissions.)

Are you making progress?

- Optimization can be fun - and time-consuming.
- Include timing to make sure optimizations actually work!
- Some techniques suggested here may either not make a large difference or may worsen performance if they affect some other aspect of the algorithm
- You need a feel of how different techniques can affect performance
- The techniques presented here are in rough order of potential improvements, starting with those most likely to give good speed-up.
- Start with a working code, and check at all stages whether the results are still correct.

Change the algorithm

Outside the scope of this course...

- Best serial algorithm not necessarily the best parallel algorithm
- For example, there exist explicitly parallelised versions of sorting algorithms
- If you're following a one-element-per-cell approach, then this is *probably* the best approach.
- (Unless there is little work per thread, in which case do multiple cells per thread in a loop.)
- Do some reading around your field/algorithm, see what has been done before (even back to 1970s).
- Take care - a parallel algorithm suited to a small cluster may not scale well for massively parallel architectures.

Use more than one GPU

- Previously, a code could be sped up using multiple cores / CPUs
- This is still the case for GPUs
- See next lecture for combining MPI with GPUs
- Likely to give good speed-up

Outline

- 1 Should we optimize
- 2 Memory layout**
- 3 Streaming Multiprocessor usage
- 4 Code optimizations
- 5 Which optimizations
- 6 Larger CUDA programs

Device-Host memory transfer

- GPUs have their advantage in on-GPU memory-bandwidth, and in multi-processors
- Device↔Host memory transfers are relatively slow
- (slower than global memory to SM bandwidth by an order of magnitude)
- It may be an advantage to do more calculation on device rather than send data back to host and calculate there even if the extra calculation is not very well parallelisable

Global memory coalescence

- On-device memory-bandwidth is only attained if the access is coalesced
- Read the Programming Guide carefully
- Use NVIDIA Nsight
- Reconsider your data-layout
- Write small test-codes to optimize

Best not to access global memory too much:

- Make use of shared memory as much as possible
- Much faster access than global memory
- May be useful to do “in-place” operations (see `euler.cu`) to reduce shared-memory usage

Outline

- 1 Should we optimize
- 2 Memory layout
- 3 Streaming Multiprocessor usage**
- 4 Code optimizations
- 5 Which optimizations
- 6 Larger CUDA programs

Occupancy

- Recall that multiple thread blocks can reside on a multiprocessor and can be switched rapidly to avoid memory latency
- So multiple blocks can start while other blocks wait for memory reads
- So if many thread blocks can fit on a multiprocessor, we can hide latency
- The number of blocks that can be held on a multiprocessor is limited by
 - Shared memory use
 - Register use
- Reduce these to increase occupancy (but only up to about 25% is worth it)
- Low occupancy usually degrades performance

Calculating occupancy

Consider a typical NVIDIA GPU:

- Each multiprocessor has
 - 32,768 32-bit registers
 - 1536 simultaneous threads
 - 48KB shared memory
- so need < 22 registers per thread for 100% occupancy, i.e. all threads held in the SM simultaneously.
- and (for 512-thread blocks):
 $48 \times 1024 \times \frac{512}{1536} = 16384$ bytes of shared-memory per block

$$\text{Occupancy} = \frac{\begin{array}{l} \text{No. of active warps per multiprocessor,} \\ \text{given registers/shared memory} \end{array}}{\text{Max no. of warps per multiprocessor}}$$

Usually easier to use the occupancy spreadsheet

`/lsc/opt/cuda-11.3/tools/CUDA_Occupancy_Calculator.xls`

Check register and shared-memory usage

Compile using `-Xptxas -v` option:

```
nvcc euler.cu -o euler -O3 -Xptxas -v
```

```
...
```

```
ptxas info : Compiling entry function
```

```
'_Z15getSLICflux_GPUILi16ELi32ELi1EEv4Gridf' for 'sm_61'
```

```
ptxas info : Function properties for
```

```
_Z15getSLICflux_GPUILi16ELi32ELi1EEv4Gridf
```

```
0 bytes stack frame, 0 bytes spill stores, 0 bytes spill  
loads
```

```
ptxas info : Used 50 registers, 34816 bytes smem, 364  
bytes
```

```
cmem[0], 108 bytes cmem[2]
```

Translation:

Function `void getSLICflux_GPU<16,32,1>(Grid, float)` uses:

50 registers per thread,

34816 bytes user-allocated shared memory,

364 bytes of user-defined constant memory

108 bytes of compiler-generated constants in constant memory

where `c++filt` was used to turn the kernel name into something readable.

Occupancy spreadsheet example

- Recall the Euler example from earlier
- The shared memory approach uses 4140 bytes of shared memory for 8×8 blocks and 36 registers
- Which of these affects occupancy?
- According to the Occupancy spreadsheet, both do.
- In this case, not a substantial difference.

Reduce registers

- Use `--maxrregcount` to force fewer registers to be used.
- Register overflow goes into local memory, which has same latency as global memory.
- However, if reducing registers increases occupancy, it can be beneficial.
- Hand-tuning the code to reduce registers can work
- Approximately equivalent to reducing local variables in function

Outline

- 1 Should we optimize
- 2 Memory layout
- 3 Streaming Multiprocessor usage
- 4 Code optimizations**
- 5 Which optimizations
- 6 Larger CUDA programs

Make the compiler do the work

- The compiler is very good at optimizing if you give it a chance
- Make as many constants available at compile-time as possible
- Can use templates to do this (see `euler.cu`)

```
template<int coord> void __global__ getFlux();  
  
getFlux<0><<<gridDim, blockDim>>>();
```

- Branching and calculations dependent on `coord` can be optimized out.
- The CUDA compiler appears to optimize certain constructs fairly aggressively, such as constant-size loops
- Exact mathematical expressions such as `sin(9.8)` do not appear to be optimized out
(there could be accuracy issues with doing so anyway)

Loop counters

- In loops, use signed integers, rather than unsigned.
- May be the opposite to what you thought
- Overflow on signed integers is undefined (in C++) so compiler can get away without overflow checks for signed arithmetic

```
for (i = 0; i < n; i++) {  
    out[i] = in[offset + stride*i];  
}
```

`stride*i` could overflow, so it's up to you to ensure that it doesn't.

- Recall that threads in a warp will run in step where possible.
- If different threads take different branches, threads taking different branches execute sequentially.
- So, make sure all threads take the same branch
- Or reduce amount of branching as much as possible.

Diagnose this using the CUDA visual profiler.

Warning

The following suggestions may well affect the accuracy of your code. Use them only if you *know* this is not a problem. Test your code with and without these suggestions and compare results.

- Use single precision unless `double` is absolutely necessary.
- Factor of 8 slowdown on older cards, factor of at least 2 even on latest cards
- Use the `--use_fast_math` compiler option - faster versions of `exp`, `sin`, `cos` etc.
- Use the `--ftz`, `--prec-div`, and `--prec-sqrt` compiler options.

For those really concerned about IEEE 754 accuracy, see Appendix E of the Programming Guide.

Instruction counting

- When most other methods have been applied, instruction counting may be useful.
- The CUDA Programming Guide contains clock-cycle counts for various instructions (5.4.1)
- For example: (compute capability 6.0)
 - 32-bit floating-point add/multiply:
64 instructions/clock-cycle
 - 64-bit floating-point add/multiply:
32 instructions/clock-cycle
 - 32-bit floating-point reciprocal, reciprocal square-root, sine/cosine:
16 instructions/clock-cycle
- Type conversions:
16 instructions/clock-cycle

General rule: Reduce the number of expensive operations.
Use temporary variables to help the compiler if necessary.

- Can sometimes replace branching by special floating point or bit-wise operations
 - `copysign(x,y)` and `signbit(x)`
 - `sincos(x, sptr, cptr)`
- For integer operations: See `CUDA_Math_API.pdf`
Includes bit-counting functions such as `__popc`.
- Or search for “Bit Twiddling Hacks” (S.E. Anderson):

```
r = y ^ ((x ^ y) & -(x < y)); // min(x, y)
```


(Interesting read even if you don't need these techniques)

C++ can be used to make code more readable while allowing compiler to optimize as usual

Using expression templates:

```
a = b + 3.2 * c - d*e;
```

in CPU code might be expanded to

```
--global-- (float* a, float* b, float* c, float* d,  
            float* e, int N)  
{  
    const int i = ...;  
    a[i] = b[i] + 3.2f * c[i] - d[i]*e[i];  
}
```

by the compiler alone.

See “C++ Templates” - Vandevoorde & Josuttis for details

Outline

- 1 Should we optimize
- 2 Memory layout
- 3 Streaming Multiprocessor usage
- 4 Code optimizations
- 5 Which optimizations**
- 6 Larger CUDA programs

Optimization approaches

First - check whether algorithm is bandwidth or instruction bound

- Imagine two stripped-down versions of the code:
 - ① No calculation - just appropriate memory reads/writes
 - ② No memory reads/writes - just calculation (be careful the compiler can't optimize them away)
- Compare times for all three versions (including original)
- In the ideal case, the full version is much faster than sum of its parts
- Why? Executing instructions can hide memory latency
- See whether memory access or calculations take longer, and optimize accordingly
- In practice, Nsight System or Nsight Compute can give the same kind of thing (so long as you have a recent GPU...)

If your kernel is bandwidth bound:

- Check for global-memory read coalescence
- Use shared memory as programmer-designed cache
- Use constant memory if there are e.g. stencil coefficients used in all kernels
- Rethink data-structures - Struct-Of-Arrays versus Array-Of-Structs

If your kernel is compute bound (less likely):

- Strength-reduction

- ① Replace multiplication by shift operation or addition if possible
- ② Replace division by multiplication by reciprocal
- ③ Reduce number of operations - use algebraic identities if possible

$$a*x*x*x + b*x*x + c*x + d == d + x*(c + x*(b + a*x))$$

- ④ Reduce use of `pow()` if possible
 - ⑤ Avoid expensive recomputation - precompute values and store in temporary variables
- Reduce (divergent) branching

If sum of memory-access and compute versions is about the same as standard kernel, you're latency bound, i.e. the hardware cannot overlap the memory-access and computation.

- Move memory-reads to early in kernel - allows instructions that don't need data to execute and hide memory-read time
- Reduce number of `__syncthreads()` - while maintaining correctness
- Increase thread-block occupancy
 - Reduce shared-memory use
 - Reduce register-use
- (Allows threads to be swapped out while their data is fetched from global memory).
- Increase the amount of data processed per kernel launch

Outline

- 1 Should we optimize
- 2 Memory layout
- 3 Streaming Multiprocessor usage
- 4 Code optimizations
- 5 Which optimizations
- 6 Larger CUDA programs**

- NVIDIA cards all support full C++ (arbitrary pointer dereference)
- Unified address space for all variables and pointers
- Virtual functions, function pointers are allowed on device
- `new` and `delete` are supported (probably not efficiently)
- Ideally, you should probably keep code at this level on the CPU and use the GPU for pure computation.
- However, for more complex codes, may be necessary to use these features.

- If writing anything longer than a single algorithm, see the Software Design course.
- `nvcc` can compile all C++ code
- C++ control code should be kept separate from kernels and kernel calls
- Run-time API calls can be put in C++ code and compiled with `gcc` - suffix `.C`
- Kernels and the `<<<...>>>` calls should be put in `.cu` files and compiled with `nvcc`
- A `__device__` or `__global__` function does not have to be in the same file as the call to it (unless templating is involved).

Linking across object files

- When compiling an object file with device code, use:
`nvcc -dc a.cu -o a.o`
Similar to `-c` option for `gcc`
- When linking separate objects containing device code, use:
`nvcc -dlink a.o b.o -o myProg`
- However, the CUDA compiler now cannot inline device functions called between files, and must introduce a general function-call stack.
- This results in more register usage.
- This may significantly impact performance (e.g. reduce by a third).

Read `CUDA_Compiler_Driver_NVCC.pdf` for compiler options

- `-O<n>` - Optimize to level n . Typically use $n = 3$
- `-arch gpubarch` Assume given GPU architecture:
`compute_50`, `compute_52`, `compute_60`
when optimizing
- `-code code` Compile code for target GPU compute capability:
`sm_50`, `sm_52`, `sm_52`, `sm_60`, `sm_70`, `sm_80`

As with all programming, the best approach is not to write any code!
NVIDIA ships some libraries as part of the CUDA Toolkit:

- Thrust - replacement for C++ STL
- NVIDIA C++ Standard Library `libcudacxx`
<https://nvidia.github.io/libcudacxx/>
- CUBLAS - BLAS for CUDA
- CUSPARSE - Sparse matrix calculations
- CUSOLVER - Covers BLAS, Sparse, and Sparse refactorization
- CUFFT - FFT calculations
- CURAND - Random number generation
- nvGRAPH - Graph creation and manipulation

And some are available from NVIDIA's GitHub repository:

- GVDB-voxels - Sparse volume compute and rendering
- AMGX - Multigrid (elliptic) linear solvers
- NCCL - Inter-GPU communication algorithms
- CUTLASS - matrix-multiplication via C++ templates
- Kokkos - C++ abstractions for running CUDA code, e.g. array classes etc.
- ... and various Deep Learning algorithms.

<http://docs.nvidia.com/cuda/thrust/index.html>

C++ has STL for various algorithms on containers:

- Sorting a list of strings
- Summing over a list of integers
- Various operations on stacks, heaps, maps.

Thrust creates versions of these that run on a GPU.

```
// Taken from Thrust examples
#include <thrust/count.h>
#include <thrust/device_vector.h>
...
// put three 1s in a device_vector
thrust::device_vector<int> vec(5,0);
vec[1] = 1;
vec[3] = 1;
vec[4] = 1;

// count the 1s
int result = thrust::count(vec.begin(), vec.end(), 1);
// result == 3
```

- BLAS (Basic Linear Algebra Subprograms) is a library of functions for operating on vectors and matrices.
- Includes operations on symmetric, triangular, and Hermitian matrices.
- Well-defined standard and interface for FORTRAN, C, and C++.
- First published in 1979, many implementations exist
- cuBLAS implements some of these functions in a precompiled CUDA library - available with `<cublas.h>` header and `-lcublas` library.
- If you only need vector/matrix operations - use this!

cuBLAS example

```
  cublasAlloc(N, sizeof(float), (void **)&x);
  cublasAlloc(N, sizeof(float), (void **)&y);

  cublasSetVector(N, sizeof(float), xHost, 1, x, 1);
  cublasSetVector(N, sizeof(float), yHost, 1, y, 1);

  clock_t start = clock();
  float tot = 0;
  for(int i=0 ; i < M ; i++)
  {
    tot += cublasSdot(N, x, 1, y, 1);
  }

  clock_t end = clock();
```

- Fast Fourier Transforms are a common application.
- cuFFT implements these in CUDA - use `cufft.h` header and `-lcufft` library.
- Includes:
 - 1D, 2D, 3D FFTs - real and complex data
 - In-place and out-of-place transforms
 - Double-precision on compatible hardware

More information is available from:

- CUDA Best Practices Guide (similar material to but more depth than this lecture)
- Whitepapers in the CUDA SDK (usually well written)
- Research articles on CUDA (may contain useful code snippets)
- Source code for CUDA applications often made freely available by authors (harder to read).