

Introduction, CUDA Basics

Jiří Filipovič

Fall 2015

Language

I will speak English if there is any foreign student in the class

- expect Slovak :-)

But I understand that

- my English is not perfect
- your English may or may not be perfect
- so feel free to interrupt me and ask me, if you do not understand

If you do not feel comfortable to ask me in English

- ask me in Czech/Slovak
- (but English is preferred)

About the class

The class is focused on algorithm design and programming of *general purpose* computing applications on *vector processors*

About the class

The class is focused on algorithm design and programming of *general purpose* computing applications on *vector processors*

We will focus to CUDA GPU first:

- C for CUDA is good for teaching (easy API, lot of examples available, mature compilers and tools)
- GPUs are wide-spread and powerful
- restricted to NVIDIA GPUs and x86 CPUs

About the class

The class is focused on algorithm design and programming of *general purpose* computing applications on *vector processors*

We will focus to CUDA GPU first:

- C for CUDA is good for teaching (easy API, lot of examples available, mature compilers and tools)
- GPUs are wide-spread and powerful
- restricted to NVIDIA GPUs and x86 CPUs

After learning CUDA, we focus to OpenCL

- programming model very similar to CUDA, easy to learn when you already know CUDA
- can be used also with AMD GPUs, Intel MIC (Xeon Phi), Cell
- we will focus on code optimizations for x86, Intel MIC and AMD GPUs

The class is practically oriented – besides efficient parallelization, we will focus on writing efficient code.

What is offered

You will learn:

- architecture of NVIDIA and AMD GPUs, Xeon Phi
- architecture-aware design of data-parallel algorithms
- programming in C for CUDA and OpenCL
- performance tuning and profiling
- basic tools and libraries for CUDA GPUs
- use cases

What is expected from you

During the semester, you will work on a practically oriented project

- important part of your total score in the class
- the same task for everybody, we will compare speed of your implementation
- 50 + 20 points of total score
 - working code: 25 points
 - efficient implementation: 25 points
 - speed of your code relative to your class mates: 20 points (only to improve your final grading)

Exam (oral or written, depending on the number of students)

- 50 points

Grading

For those finishing by exam:

- A: 92–100
- B: 86–91
- C: 78–85
- D: 72–77
- E: 66–71
- F: 0–65 pts

For those finishing by colloquium:

- 50 pts

Materials – CUDA

CUDA documentation (installed as a part of CUDA Toolkit, downloadable from *developer.nvidia.com*)

- CUDA C Programming Guide (most important properties of CUDA)
- CUDA C Best Practices Guide (more detailed document focusing on optimizations)
- CUDA Reference Manual (complete description of C for CUDA API)
- other useful documents (nvcc guide, PTX language description, library manuals, ...)

CUDA article series, Supercomputing for the Masses

- <http://www.ddj.com/cpp/207200659>

Materials – OpenCL


- OpenCL 1.1 Specification
- AMD Accelerated Parallel Processing Programming Guide
- Intel OpenCL SDK Programming Guide
- Writing Optimal OpenCL Code with Intel OpenCL SDK

Materials – Parallel Programming

- Ben-Ari M., Principles of Concurrent and Distributed Programming, 2nd Ed. Addison-Wesley, 2006
- Timothy G. Mattson, Beverly A. Sanders, Berna L. Massingill, Patterns for Parallel Programming, Addison-Wesley, 2004


Motivation – GPU arithmetic performance

Theoretical GFLOP/s



Motivation – GPU memory bandwidth

Theoretical GB/s



Motivation – programming complexity

OK, GPUs are more powerful, but GPU programming is substantially more difficult, right?

- well, it is more difficult comparing to writing serial C/C++ code...
- but can we compare it to serial code?

Motivation – programming complexity

OK, GPUs are more powerful, but GPU programming is substantially more difficult, right?

- well, it is more difficult comparing to writing serial C/C++ code...
- but can we compare it to serial code?

Moore's Law

Number of transistors on a single chip **doubles** every 18 months

Motivation – programming complexity

OK, GPUs are more powerful, but GPU programming is substantially more difficult, right?

- well, it is more difficult comparing to writing serial C/C++ code...
- but can we compare it to serial code?

Moore's Law

Number of transistors on a single chip **doubles** every 18 months

Corresponding growth of performance comes from

- **in the past:** frequency increase, parallelism of instructions, of-of-order instruction processing, caches, etc.
- **today:** vector instructions, increase in number of cores

Motivation – paradigm change

Moore's Law consequences:

- **in the past:** changes were important for compiler developers; application developers didn't need to worry
- **today:** in order to utilize state-of-the-art processors, it is necessary to write parallel and vectorized code
 - it is necessary to find parallelism in the problem being solved, which is a task for a programmer, not for a compiler (at least for now)
 - writing efficient code for modern CPUs is similarly difficult as writing for GPUs

Motivation – Applications

Use of GPU for general computations is a dynamically developing field with broad applicability

Motivation – Applications

Use of GPU for general computations is a dynamically developing field with broad applicability

- high-performance scientific calculations
 - computational chemistry
 - physical simulations
 - image processing
 - medicine
 - and others. . .

Motivation – Applications

Use of GPU for general computations is a dynamically developing field with broad applicability

- high-performance scientific calculations
 - computational chemistry
 - physical simulations
 - image processing
 - medicine
 - and others. . .
- performance-hungry home and desktop applications
 - encoding/decoding of multimedia data
 - game physics
 - image editing, 3D rendering
 - etc.

Why are GPUs so powerful?

Types of Parallelism

- Task parallelism
 - decomposition of a task into the problems that may be processed in parallel
 - usually more complex tasks performing different actions
 - usually more frequent (and complex) synchronization
 - ideal for small number of high-performance processors
- Data parallelism
 - parallelism on the level of data structures
 - usually the same operations on many items of a data structure
 - finer-grained parallelism allows for simple construction of individual processors

Why are GPUs so powerful?


From programmer's perspective

- some problems are rather data-parallel, some task-parallel (graph traversal vs. edge processing)


From hardware perspective

- processors for data-parallel tasks may be **simpler**
- it is possible to achieve **higher arithmetic performance** with the same size of processor
- simpler memory access patterns allow for **high-throughput memory** designs

GPU Architecture



CPU



GPU

GPU Architecture

Main differences compared to CPU

- high parallelism: tens thousands threads needed to utilize high-end GPUs
- SIMT model: subsets of threads runs in lock-step mode
- distributed on-chip memory: subsets of threads shares their private memory
- restricted caching capabilities: small cache, often read-only

Algorithms usually need to be redesigned to be efficient on GPU.

GPU Architecture


Within the system:

- co-processor with dedicated memory (discrete GPU)
- asynchronous processing of instructions
- attached using PCI-E to the rest of the system (discrete GPU)

CUDA

CUDA (Compute Unified Device Architecture)

- architecture for parallel computations developed by NVIDIA
- provides a new programming model, allows efficient implementation of general GPU computations
- may be used in multiple programming languages



G80 Processor

G80


- first CUDA processor
- 16 multiprocessors
- each multiprocessor
 - 8 scalar processors
 - 2 units for special functions
 - up to 768 threads
 - HW for thread switching and scheduling
 - threads are grouped into warps by 32
 - SIMT
 - native synchronization within the multiprocessor

G80 Memory Model

Memory model

- 8192 registers shared among all threads of a multiprocessor
- 16 kB of shared memory
 - local within the multiprocessor
 - as fast as registry (under certain constraints)
- constant memory
 - cached, read-only
- texture memory
 - cached with 2D locality, read-only
- global memory
 - non cached, read-write
- data transfers between global memory and system memory through PCI-E

G80 Processor



C for CUDA

C for CUDA is extension of C for parallel computations


- explicit separation of host (CPU) and device (GPU) code
- thread hierarchy
- memory hierarchy
- synchronization mechanisms
- API

Thread Hierarchy

Thread hierarchy

- threads are organized into blocks
- blocks form a grid
- problem is decomposed into sub-problems that can be run independently in parallel (blocks)
- individual sub-problems are divided into small pieces that can be run cooperatively in parallel (threads)
- scales well

Thread Hierarchy

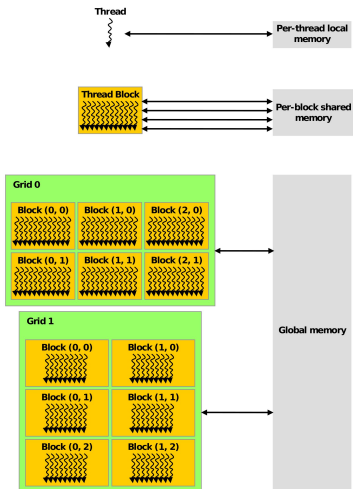


Memory Hierarchy

More memory types:

- different visibility
- different lifetime
- different speed and behavior
- brings good scalability

Memory Hierarchy



An Example – Sum of Vectors

We want to sum vectors a and b and store the result in vector c

An Example – Sum of Vectors

We want to sum vectors a and b and store the result in vector c
We need to find parallelism in the problem.

An Example – Sum of Vectors

We want to sum vectors a and b and store the result in vector c

We need to find parallelism in the problem.

Serial sum of vectors:

```
for (int i = 0; i < N; i++)  
    c[i] = a[i] + b[i];
```

An Example – Sum of Vectors

We want to sum vectors a and b and store the result in vector c

We need to find parallelism in the problem.

Serial sum of vectors:

```
for (int i = 0; i < N; i++)  
    c[i] = a[i] + b[i];
```

Individual iterations are independent – it is possible to parallelize, scales with the size of the vector.

An Example – Sum of Vectors

We want to sum vectors a and b and store the result in vector c

We need to find parallelism in the problem.

Serial sum of vectors:

```
for (int i = 0; i < N; i++)  
    c[i] = a[i] + b[i];
```


Individual iterations are independent – it is possible to parallelize,
scales with the size of the vector.

i -th thread sums i -th component of the vector:

```
c[i] = a[i] + b[i];
```

How do we find which thread we are?

Thread Hierarchy



Thread and Block Identification

C for CUDA has built-in variables:

- **threadIdx.**{x, y, z} tells position of a thread in a block
- **blockDim.**{x, y, z} tells size of the block
- **blockIdx.**{x, y, z} tells position of the block in grid (z always equals 1)
- **gridDim.**{x, y, z} tells grid size (z always equals 1)

An Example – Sum of Vectors

Thus we calculate the position of the thread (grid and block are one-dimensional):

An Example – Sum of Vectors

Thus we calculate the position of the thread (grid and block are one-dimensional):

```
int i = blockIdx.x*blockDim.x + threadIdx.x;
```

An Example – Sum of Vectors

Thus we calculate the position of the thread (grid and block are one-dimensional):

```
int i = blockIdx.x*blockDim.x + threadIdx.x;
```

Whole function for parallel summation of vectors:

```
__global__ void addvec(float *a, float *b, float *c){  
    int i = blockIdx.x*blockDim.x + threadIdx.x;  
    c[i] = a[i] + b[i];  
}
```

An Example – Sum of Vectors

Thus we calculate the position of the thread (grid and block are one-dimensional):

```
int i = blockIdx.x*blockDim.x + threadIdx.x;
```

Whole function for parallel summation of vectors:

```
__global__ void addvec(float *a, float *b, float *c){
    int i = blockIdx.x*blockDim.x + threadIdx.x;
    c[i] = a[i] + b[i];
}
```

The function defines so called kernel; we specify how many threads and what structure will be run when calling.

Function Type Quantifiers

C syntax enhanced by quantifiers defining where the code is run and from where it may be called:

- **__device__** function is run on device (GPU) only and may be called from the device code only
- **__global__** function is run on device (GPU) only and may be called from the host (CPU) code only
- **__host__** function is run on host only and may be called from the host only
- **__host__** and **__device__** may be combined – function is compiled for both then

The following steps are needed for the full computation:

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data
- *allocate memory on GPU*

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data
- *allocate memory on GPU*
- *copy vectors a and b to GPU*

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data
- *allocate memory on GPU*
- *copy vectors a a b to GPU*
- **compute the sum on GPU**

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data
- *allocate memory on GPU*
- *copy vectors a a b to GPU*
- **compute the sum on GPU**
- *store the result from GPU into c*

The following steps are needed for the full computation:

- allocate memory for vectors and fill it with data
- *allocate memory on GPU*
- *copy vectors a a b to GPU*
- **compute the sum on GPU**
- *store the result from GPU into c*
- use the result in *c* :-)

When managed memory is used (requires GPU with computing capability 3.0 and CUDA 6.0 or better), steps written in italics are not required.

An Example – Sum of Vectors

CPU code that fills a and b and computes c

```
#include <stdio.h>
#define N 64
int main(){
    float a[N], b[N], c[N];
    for (int i = 0; i < N; i++) {
        a[i] = i;
        b[i] = i*3;
    }

    // GPU code will be here

    for (int i = 0; i < N; i++)
        printf("%f, ", c[i]);
    return 0;
}
```

GPU Memory Management

Using managed memory, CUDA maintains memory transfers between CPU and GPU automatically.

- memory coherency is guaranteed
- GPU memory cannot be used when any GPU kernel is running

Memory operations can be programmed explicitly

```
cudaMalloc(void** devPtr, size_t count);  
cudaFree(void* devPtr);  
cudaMemcpy(void* dst, const void* src, size_t count,  
           enum cudaMemcpyKind kind);
```

An Example – Sum of Vectors

Running the kernel:

- kernel is called as a function; between the name and the arguments, there are triple angle brackets with specification of grid and block size
- we need to know block size and their count
- we will use 1D block and grid with fixed block size
- the size of the grid is determined in a way to compute the whole problem of vector sum

For vector size divisible by 32:

```
#define BLOCK 32
addvec<<<N/BLOCK, BLOCK>>>(d_a, d_b, d_c);
```

How to solve a general vector size?

An Example – Sum of Vectors

We will modify the kernel source:

```
__global__ void addvec(float *a, float *b, float *c, int n){  
    int i = blockIdx.x*blockDim.x + threadIdx.x;  
    if (i < n) c[i] = a[i] + b[i];  
}
```

And call the kernel with sufficient number of threads:

```
addvec<<<N/BLOCK + 1, BLOCK>>>(d_a, d_b, d_c, N);
```

An Example – Running It

Now we just need to compile it :-)

```
nvcc -o vecadd vecadd.cu
```

Where to work with CUDA?

- on a remote computer: airacuda.fi.muni.cz, accounts will be made
- your own machine: download and install CUDA toolkit and SDK from developer.nvidia.com
- source code used in lectures will be published as a part of course materials