

Hands-On GPU Programming with Python and CUDA

Explore high-performance parallel computing with CUDA



Dr. Brian Tuomanen

Packt

www.packt.com

Dr. Brian Tuomanen

Hands-On GPU Programming with Python and CUDA: Explore high-performance parallel computing with CUDA



[continue reading](#)

Build real-world applications by composing effective GPU code, CUDA kernels, and device functions with the most recent top features of Python 3.7, CUDA 9 and CUDA 10

Key Features

- Expand your background in GPU development—ll begin by learning how exactly to apply Amdahl's PyCUDA, scikit-cuda, and Nsight
- Effectively use CUDA libraries such as cuBLAS, cuFFT, and cuSolver
- Apply GPU programming to modern data science applications

Book Description

Hands-On GPU Programming with Python and CUDA hits the ground running: you'll learn the rules, use a code profiler to recognize bottlenecks in your Python code, and set up an appropriate GPU programming environment.

ESTABLISHING Your GPU Development Environment? ll then observe how to "query" ll explore some of the more well-known NVIDIA libraries, such as for example cuFFT and cuBLAS's features and duplicate arrays of data to and from the GPU. In the final chapter, you'll see some topics and applications related to GPU development that you may desire to pursue, including AI, images, and blockchain.

With a solid background in place, you will today apply your new-found knowledge to build up your own GPU-based deep neural network from scratch.

ll get to grips with profiling GPU code efficiently and fully ensure that you debug your code using Nsight IDE. Next, you'll explore the GPU's own memory. ll explore advanced topics, such as warp shuffling, dynamic parallelism, and PTX assembly. ll start code directly onto the GPU and write full-blown GPU kernels and gadget functions in CUDA C. By the finish of this book, you will be able to use GPU programming to complications related to data science and high-performance computing.

What you would learn

- Launch GPU code directly from Python
- Write effective and efficient GPU kernels and device functions
- Use libraries such as for example cuFFT, cuBLAS, and cuSolver
- Debug and profile your code with Nsight and Visual Profiler
- Apply GPU programming to datascience problems
- Build a GPU-based deep neural network from scratch
- Explore advanced GPU hardware features, such as for example warp shuffling

Who this book is for

Hands-On GPU Development with Python and CUDA is for developers and data scientists who want to learn the basics of effective GPU programming to improve performance using Python code. You should have an understanding of first-year university or university-level engineering mathematics and physics, and have some experience with Python as well as in any C-based program writing language such as C, C++, Go, or Java.

Table of Contents

- Why GPU Development? You'll Get Started with PyCUDA
- Streams, Events, Contexts, and Concurrency
- Debugging and Profiling Your CUDA Code
- Kernels, Threads, Blocks, and Grids
- Using the CUDA Libraries with Scikit-CUDA
- Draft complete
- The CUDA Gadget Function Libraries and Thrust
- Implementing a Deep Neural Network
- Working with Compiled GPU Code
- Efficiency Optimization in CUDA
- Where you can Go from Here



[continue reading](#)

download Hands-On GPU Programming with Python and CUDA: Explore high-performance parallel computing with CUDA ebook

download Hands-On GPU Programming with Python and CUDA: Explore high-performance parallel computing with CUDA ebook

[download free CRYPTO CURRENCY BASICS: BITCOIN FOR BEGINNERS \(mini series Book 1\) e-book](#)

[download free Bitcoin: Is This The Future? \(Art & Science of Digital Currency\) txt](#)

[download free Fintech Explained fb2](#)