

GPU Parallel Program Development Using CUDA

GPU Parallel Program Development using CUDA teaches GPU programming by showing the differences among different families of GPUs. This approach prepares the reader for the next generation and future generations of GPUs. The book emphasizes concepts that will remain relevant for a long time, rather than concepts that are platform-specific. At the same time, the book also provides platform-dependent explanations that are as valuable as generalized GPU concepts. The book consists of three separate parts; it starts by explaining parallelism using CPU multi-threading in Part I. A few simple programs are used to demonstrate the concept of dividing a large task into multiple parallel sub-tasks and mapping them to CPU threads. Multiple ways of parallelizing the same task are analyzed and their pros/cons are studied in terms of both core and memory operation. Part II of the book introduces GPU massive parallelism. The same programs are parallelized on multiple Nvidia GPU platforms and the same performance analysis is repeated. Because the core and memory structures of CPUs and GPUs are different, the results differ in interesting ways. The end goal is to make programmers aware of all the good ideas, as well as the bad ideas, so readers can apply the good ideas and avoid the bad ideas in their own programs. Part III of the book provides pointer for readers who want to expand their horizons. It provides a brief introduction to popular CUDA libraries (such as cuBLAS, cuFFT, NPP, and Thrust), the OpenCL programming language, an overview of GPU programming using other programming languages and API libraries (such as Python, OpenCV, OpenGL, and Apple's Swift and Metal,) and the deep learning library cuDNN.

This book teaches GPU programming by introducing CPU multi-threaded programming and bases GPU massively-parallel programming on this foundation. The differences among families of GPUs are also studied. The book also explores CUDA libraries, OpenCL, GPU programming with other languages and API libraries, and the deep learning library cuDNN.



71,80 €

71,80 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9780367572242

Medium: Buch

ISBN: 978-0-367-57224-2

Verlag: Chapman and Hall/CRC

Erscheinungstermin: 30.06.2020

Sprache(n): Englisch

Auflage: 1. Auflage 2020

Serie: Chapman & Hall/CRC

Computational Science

Produktform: Kartoniert

Gewicht: 891 g

Seiten: 476

Format (B x H): 178 x 254 mm

