


The book was found

CUDA Programming: A Developer's Guide To Parallel Computing With GPUs (Applications Of Gpu Computing)



Synopsis

If you need to learn CUDA but don't have experience with parallel computing, *CUDA Programming: A Developer's Introduction* offers a detailed guide to CUDA with a grounding in parallel fundamentals. It starts by introducing CUDA and bringing you up to speed on GPU parallelism and hardware, then delving into CUDA installation. Chapters on core concepts including threads, blocks, grids, and memory focus on both parallel and CUDA-specific issues. Later, the book demonstrates CUDA in practice for optimizing applications, adjusting to new hardware, and solving common problems. Comprehensive introduction to parallel programming with CUDA, for readers new to both. Detailed instructions help readers optimize the CUDA software development kit. Practical techniques illustrate working with memory, threads, algorithms, resources, and more. Covers CUDA on multiple hardware platforms: Mac, Linux and Windows with several NVIDIA chipsets. Each chapter includes exercises to test reader knowledge.

Book Information

Series: Applications of Gpu Computing

Paperback: 600 pages

Publisher: Morgan Kaufmann; 1 edition (November 27, 2012)

Language: English

ISBN-10: 0124159338

ISBN-13: 978-0124159334

Product Dimensions: 7.5 x 1.3 x 9.2 inches

Shipping Weight: 2.2 pounds (View shipping rates and policies)

Average Customer Review: 3.8 out of 5 stars Â Â See all reviews Â (11 customer reviews)

Best Sellers Rank: #575,558 in Books (See Top 100 in Books) #60 in Â Books > Computers & Technology > Programming > Parallel Programming #772 in Â Books > Textbooks > Computer Science > Software Design & Engineering #1626 in Â Books > Computers & Technology > Programming > Software Design, Testing & Engineering > Software Development

Customer Reviews

To put this review into context, I come from a background of formal training in computer science, and my day-to-day job is systems engineering. With that said, I'm very comfortable with both systems architecture and the C programming language. It seemed that this book was perfectly targeted for somebody like me, as it assumes an substantial existing knowledge base and doesn't slow down to cover much of the basics. However, I do find that the manner in which the author

covers things to be rather tedious. The writing style itself could use some work, and it frequently makes use of analogies that confuse rather than illuminate. The author will continuously repeat himself on several concepts, and although I agree repetition is critical to learning, it can be done so much at times that it is a hindrance. The book also doesn't necessarily follow a progression of increasing "detail". So while introducing a new concept, the author often goes down the rabbit hole of unnecessary detail and device specific optimizations, and other tangents that are not strictly critical at that time. Another slight complaint is that the CUDA spec has changed dramatically, and thus there are significant changes between CUDA 1/2/3. As a beginner, the complexity of learning something across multiple versions is also frustrating at times. I would've greatly preferred a focus on CUDA 3, and then added notes about backward compatibility in a later chapter, as opposed to covering all three simultaneously. On the plus side, I do like that the book doesn't spend much time teaching C or basic systems architecture. There is not a lot of introductory material to get through before you get to the meat of CUDA.

Do not buy this item! My background is Master level CS, some CUDA experience, but none in a professional capacity. The best part of these 600 pages of text is their promise: learning how to program with CUDA as an already experienced developer. But that's where the good parts end, sadly. What stands out first and foremost is the incoherent stream of arbitrary facts and figures related to CUDA programming, nVidia GPUs, parallel programming, and profiling metrics of the author's test programs. It's all over the place! Rarely (if ever) are basic or advanced concepts clearly or concisely explained, the index is insufficient at best (best just Google it or read StackOverflow threads), and what information can be distilled has to be extracted piece by piece by the reader. A term will often be used a fair number of times before the author gets around to any sort of definition or comprehensive description. And when he does, he resorts to car analogies, supermarket queues and whatnot like his audience is completely new to programming ('imagine computer memory like shelf storage, and you're the manager of the warehouse' could very well be a phrase used in this text). The author's idea of explaining complex concepts seems to be to just use a different analogy every time a particular aspect of that concept is relevant to the current subject. I could go on in some detail about other noticeably poor parts of this item, like its wordiness, lack of vision and structure, quality of its figures and tables ("Descriptions are for suckers!"), endless source code listings in huge font ("Got to get that page count up, or we won't be able to charge \$45"), lack of overview, and much, much more, but I'll leave that to others.

I have been looking over almost all of the books on GPGPU programming for three months now and IMHO this book is presently the best one to select for Nvidia hardware understanding and program development. "CUDA Programming" meets high standards for in depth hardware exploration and program approaches. There are almost 100 pages of code in this 550 page book that help the reader get acquainted with CUDA programming differences from standard C++. Although this is a fairly deep read, it delivers a host of understanding about GPU hardware architectures and how they create a demand for programming a certain way that supports the high throughput potential that the new GPU cards can deliver. There are new terms to learn in this arena: Grids, Blocks, Threads, Warps, and Kernels that are related to the parallel programming paradigm. The author does a good job of explaining what they are and the how and why of their proper use. I'm sure that there will be less erratic introductions of this material in the future, but for 2012, this is the best there is in print. No doubt, you will need to reread the material three times to get it all down correctly. That is why you will need many program examples to make it all clear. The book also covers the various programming libraries that surround the CUDA dialect of C++ and is helpful in understanding their usage. There is also an introduction to CUDA debugging tools and several program timing comparisons. The book devotes chapter 10 to the various support libraries and describes their area of support briefly. The current release of the CUDA SDK is version 5. This SDK contains a few dozen example programs and I believe that you will develop most of your understanding of CUDA from it.

[Download to continue reading...](#)

CUDA Programming: A Developer's Guide to Parallel Computing with GPUs (Applications of Gpu Computing) Programming Massively Parallel Processors: A Hands-on Approach (Applications of GPU Computing Series) CUDA Handbook: A Comprehensive Guide to GPU Programming, The CUDA by Example: An Introduction to General-Purpose GPU Programming CUDA for Engineers: An Introduction to High-Performance Parallel Computing Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers (2nd Edition) Parallel Scientific Computing in C++ and MPI: A Seamless Approach to Parallel Algorithms and their Implementation Introduction to Parallel Computing: Design and Analysis of Parallel Algorithms Parallel Programming with Intel Parallel Studio XE Short Stories in Spanish: New Penguin Parallel Text (New Penguin Parallel Texts) (Spanish and English Edition) Learn German: Parallel Text - Easy, Funny Stories (German - English) - Bilingual (Learning German with Parallel Text Book 1) Learn German III: Parallel Text - Easy Stories (German - English) Bilingual - Dual Language (Learning German with Parallel Text 3) (German Edition) Dependable Computing for Critical Applications 5 (Dependable

Computing and Fault-Tolerant Systems) Java: The Simple Guide to Learn Java Programming In No Time (Programming,Database, Java for dummies, coding books, java programming) (HTML,Javascript,Programming,Developers,Coding,CSS,PHP) (Volume 2) Deep Belief Nets in C++ and CUDA C: Volume III: Convolutional Nets (Volume 3) Deep Belief Nets in C++ and CUDA C: Volume 1: Restricted Boltzmann Machines and Supervised Feedforward Networks Programming Microsoft Office 365 (includes Current Book Service): Covers Microsoft Graph, Office 365 applications, SharePoint Add-ins, Office 365 Groups, and more (Developer Reference) Big CPU, Big Data: Solving the World's Toughest Computational Problems with Parallel Computing GPU Pro 7: Advanced Rendering Techniques GPU Pro 6: Advanced Rendering Techniques

[Dmca](#)