1. Record Nr. UNINA9910811688503321

Autore	Cook Shane
Titolo	CUDA programming : a developer's guide to parallel computing with GPUs / / Shane Cook
Pubbl/distr/stampa	Amsterdam ; ; Boston, : Elsevier/MK, c2013
ISBN	1-283-71645-3 0-12-415988-5
Edizione	[1st edition]
Descrizione fisica	1 online resource (591 p.)
Collana	Applications of GPU Computing Series
Disciplina	004.35 005.275
Soggetti	Parallel programming (Computer science)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Description based upon print version of record.
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Front Cover; CUDA Programming: A Developer's Guide to ParallelComputing with GPUs; Copyright; Contents; Preface; Chapter 1 - A Short History of Supercomputing; INTRODUCTION; VON NEUMANN ARCHITECTURE; CRAY; CONNECTION MACHINE; CELL PROCESSOR; MULTINODE COMPUTING; THE EARLY DAYS OF GPGPU CODING; THE DEATH OF THE SINGLE-CORE SOLUTION; NVIDIA AND CUDA; GPU HARDWARE; ALTERNATIVES TO CUDA; CONCLUSION; Chapter 2 - Understanding Parallelism with GPUs; INTRODUCTION; TRADITIONAL SERIAL CODE; SERIAL/PARALLEL PROBLEMS; CONCURRENCY; TYPES OF PARALLELISM; FLYNN'S TAXONOMY; SOME COMMON PARALLEL PATTERNS CONCLUSIONChapter 3 - CUDA Hardware Overview; PC ARCHITECTURE; GPU HARDWARE; CPUS AND GPUS; COMPUTE LEVELS; Chapter 4 - Setting Up CUDA; INTRODUCTION; INSTALLING THE SDK UNDER WINDOWS; VISUAL STUDIO; LINUX; MAC; INSTALLING A DEBUGGER; COMPILATION MODEL; ERROR HANDLING; CONCLUSION; Chapter 5 - Grids, Blocks, and Threads; WHAT IT ALL MEANS; THREADS; BLOCKS; GRIDS; WARPS; BLOCK SCHEDULING; A PRACTICAL EXAMPLE-HISTOGRAMS; CONCLUSION; Chapter 6 - Memory Handling with CUDA; INTRODUCTION; CACHES; REGISTER USAGE; SHARED MEMORY; CONSTANT MEMORY; GLOBAL MEMORY; TEXTURE MEMORY; CONCLUSION

Chapter 7 - Using CUDA in PracticeINTRODUCTION; SERIAL AND PARALLEL CODE; PROCESSING DATASETS; PROFILING; AN EXAMPLE USING AES; CONCLUSION; References; Chapter 8 - Multi-CPU and Multi-GPU Solutions; INTRODUCTION; LOCALITY; MULTI-CPU SYSTEMS; MULTI-GPU SYSTEMS; ALGORITHMS ON MULTIPLE GPUS; WHICH GPU?; SINGLE-NODE SYSTEMS; STREAMS; MULTIPLE-NODE SYSTEMS; CONCLUSION; Chapter 9 - Optimizing Your Application; STRATEGY 1: PARALLEL/SERIAL GPU/CPU PROBLEM BREAKDOWN; STRATEGY 2: MEMORY CONSIDERATIONS; STRATEGY 3: TRANSFERS; STRATEGY 4: THREAD USAGE, CALCULATIONS, AND DIVERGENCE; STRATEGY 5: ALGORITHMS
STRATEGY 6: RESOURCE CONTENTIONSSTRATEGY 7: SELF-TUNING APPLICATIONS; CONCLUSION; Chapter 10 - Libraries and SDK; INTRODUCTION; LIBRARIES; CUDA COMPUTING SDK; DIRECTIVE-BASED PROGRAMMING; WRITING YOUR OWN KERNELS; CONCLUSION; Chapter 11 - Designing GPU-Based Systems; INTRODUCTION; CPU PROCESSOR; GPU DEVICE; PCI-E BUS; GEFORCE CARDS; CPU MEMORY; AIR COOLING; LIQUID COOLING; DESKTOP CASES AND MOTHERBOARDS; MASS STORAGE; POWER CONSIDERATIONS; OPERATING SYSTEMS; CONCLUSION; Chapter 12 - Common Problems, Causes, and Solutions; INTRODUCTION; ERRORS WITH CUDA DIRECTIVES; PARALLEL PROGRAMMING ISSUES
ALGORITHMIC ISSUESFINDING AND AVOIDING ERRORS; DEVELOPING FOR FUTURE GPUS; FURTHER RESOURCES; CONCLUSION; References; Index

Sommario/riassunto

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. <u