

Intel Fpga Sdk For Opencil Programming Guide Altera

OpenCL Programming Guide *Heterogeneous Computing with OpenCL* [OpenCL Programming by Example](#) [OpenCL in Action](#) [WebGL Programming Guide](#) [OpenVX Programming Guide](#) *Programming Massively Parallel Processors* **CUDA by Example Heterogeneous Computing with OpenCL 2.0** [Vulkan Programming Guide](#) **CUDA Programming** [OpenGL ES 3.0 Programming Guide](#) **OpenGL Programming Guide Introduction to Concurrency in Programming Languages** **Introduction to High Performance Scientific Computing** **Parallel Programming with OpenACC** *Data Parallel C++ PThreads Programming Euro-Par 2012 Parallel Processing* **Patterns for Parallel Programming Heterogeneous Computing Architectures GPU Parallel Program Development Using CUDA** **An Introduction to Parallel Programming Intel Xeon Phi Coprocessor Architecture and Tools** **The CUDA Handbook Parallel Programming with MPI** *Hands-On GPU Programming with CUDA* **OpenGL Programming Guide** *CUDA Application Design and Development GPU Pro 6* [OpenGL Programming Guide Numerical Computations with GPUs](#) *Introduction to Parallel Computing* **Intel Xeon Phi Coprocessor High Performance Programming Multicore and GPU Programming GPU Pro 5** **OpenGL Programming Guide** [OpenGL Shading Language](#) **GPU Pro 360 Guide to Rendering Learning OpenGL ES for iOS**

This is likewise one of the factors by obtaining the soft documents of this **Intel Fpga Sdk For Opencil Programming Guide Altera** by online. You might not require more grow old to spend to go to the book introduction as without difficulty as search for them. In some cases, you likewise get not discover the proclamation Intel Fpga Sdk For Opencil Programming Guide Altera that you are looking for. It will extremely squander the time.

However below, later you visit this web page, it will be thus very easy to acquire as with ease as download lead Intel Fpga Sdk For Opencil Programming Guide Altera

It will not give a positive response many grow old as we run by before. You can do it even though play a role something else at house and even in your workplace. fittingly easy! So, are you question? Just exercise just what we meet the expense of under as capably as review **Intel Fpga Sdk For Opencil Programming Guide Altera** what you taking into consideration to read!

[OpenVX Programming Guide](#) Jul 21 2022 OpenVX is the computer vision API adopted by many high-performance processor vendors. It is quickly becoming the preferred way to write fast and power-efficient code on embedded systems. OpenVX Programming Guidebook presents definitive information on OpenVX 1.2 and 1.3, the Neural Network, and other extensions as well as the OpenVX Safety Critical standard. This book gives a high-level overview of the OpenVX standard, its design principles, and overall structure. It covers computer vision functions and the graph API, providing examples of usage for the majority of the functions. It is intended both for the first-time user of OpenVX and as a reference for experienced OpenVX developers. Get to grips with the OpenVX standard and gain insight why various options were chosen Start developing efficient OpenVX code instantly Understand design principles and use them to create robust code Develop consumer and industrial products that use computer vision to understand and interact with the real world [WebGL Programming Guide](#) Aug 22 2022 Using WebGL®, you can create sophisticated interactive 3D graphics inside web browsers, without plug-ins. WebGL makes it possible to build a new generation of 3D web games, user interfaces, and information visualization solutions that will run on any standard web browser, and on PCs, smartphones, tablets, game consoles, or other devices. WebGL Programming Guide will help you get started quickly with interactive WebGL 3D programming, even if you have no prior knowledge of HTML5, JavaScript, 3D graphics, mathematics, or OpenGL. You'll learn step-by-step, through realistic examples, building your skills as you move from simple to complex solutions for building visually appealing web pages and 3D applications with WebGL. Media, 3D graphics, and WebGL pioneers Dr. Kouichi Matsuda and Dr. Rodger Lea offer easy-to-understand tutorials on key aspects of WebGL, plus 100 downloadable sample programs, each demonstrating a specific WebGL topic. You'll move from basic techniques such as rendering, animating, and texturing triangles, all the way to advanced techniques such as fogging, shadowing, shader switching, and displaying 3D models generated by Blender or other authoring tools. This book won't just teach you WebGL best practices, it will give you a library of code to jumpstart your own projects. Coverage includes: • WebGL's origin, core concepts, features, advantages, and integration with other web standards • How and basic WebGL functions work together to deliver 3D graphics • Shader development with OpenGL ES Shading Language (GLSL ES) • 3D scene drawing: representing user views, controlling space volume, clipping, object creation, and perspective • Achieving greater realism through lighting and hierarchical objects • Advanced techniques: object manipulation, heads-up displays, alpha blending, shader switching, and more • Valuable reference appendixes covering key issues ranging from coordinate systems to matrices and shader loading to web browser settings This is the newest text in the OpenGL Technical Library, Addison-Wesley's definitive collection of programming guides an reference manuals for OpenGL and its related technologies. The Library enables programmers to gain a practical understanding of OpenGL and the other Khronos application-programming libraries including OpenGL ES and OpenCL. All of the technologies in the OpenGL Technical Library evolve under the auspices of the Khronos Group, the industry consortium guiding the evolution of modern, open-standards media APIs.

Introduction to Concurrency in Programming Languages Nov 13 2021 Exploring how concurrent programming can be assisted by language-level techniques, *Introduction to Concurrency in Programming Languages* presents high-level language techniques for dealing with concurrency in a general context. It provides an understanding of programming languages that offer concurrency features as part of the language definition. The book supplies a conceptual framework for different aspects of parallel algorithm design and implementation. It first addresses the limitations of traditional programming techniques and models when dealing with concurrency. The book then explores the current state of the art in concurrent programming and describes high-level language constructs for concurrency. It also discusses the historical evolution of hardware, corresponding high-level techniques that were developed, and the connection to modern systems, such as multicore and manycore processors. The remainder of the text focuses on common high-level programming techniques and their application to a range of algorithms. The authors offer case studies on genetic algorithms, fractal generation, cellular automata, game logic for solving Sudoku puzzles, pipelined algorithms, and more. Illustrating the effect of concurrency on programs written in familiar languages, this text focuses on novel language abstractions that truly bring concurrency into the language and aid analysis and compilation tools in generating efficient, correct programs. It also explains the complexity involved in taking advantage of concurrency with regard to program correctness and performance.

Hands-On GPU Programming with CUDA Sep 30 2020 Explore different GPU programming methods using libraries and directives, such as OpenACC, with extension to languages such as C, C++, and Python Key Features Learn parallel programming principles and practices and performance analysis in GPU computing Get to grips with distributed multi GPU programming and other approaches to GPU programming Understand how GPU acceleration in deep learning models can improve their performance Book Description Compute Unified Device Architecture (CUDA) is NVIDIA's GPU computing platform and application programming

interface. It's designed to work with programming languages such as C, C++, and Python. With CUDA, you can leverage a GPU's parallel computing power for a range of high-performance computing applications in the fields of science, healthcare, and deep learning. Learn CUDA Programming will help you learn GPU parallel programming and understand its modern applications. In this book, you'll discover CUDA programming approaches for modern GPU architectures. You'll not only be guided through GPU features, tools, and APIs, you'll also learn how to analyze performance with sample parallel programming algorithms. This book will help you optimize the performance of your apps by giving insights into CUDA programming platforms with various libraries, compiler directives (OpenACC), and other languages. As you progress, you'll learn how additional computing power can be generated using multiple GPUs in a box or in multiple boxes. Finally, you'll explore how CUDA accelerates deep learning algorithms, including convolutional neural networks (CNNs) and recurrent neural networks (RNNs). By the end of this CUDA book, you'll be equipped with the skills you need to integrate the power of GPU computing in your applications. What you will learn Understand general GPU operations and programming patterns in CUDA Uncover the difference between GPU programming and CPU programming Analyze GPU application performance and implement optimization strategies Explore GPU programming, profiling, and debugging tools Grasp parallel programming algorithms and how to implement them Scale GPU-accelerated applications with multi-GPU and multi-nodes Delve into GPU programming platforms with accelerated libraries, Python, and OpenACC Gain insights into deep learning accelerators in CNNs and RNNs using GPUs Who this book is for This beginner-level book is for programmers who want to delve into parallel computing, become part of the high-performance computing community and build modern applications. Basic C and C++ programming experience is assumed. For deep learning enthusiasts, this book covers Python InterOps, DL libraries, and practical examples on performance estimation.

OpenCL Programming Guide Dec 26 2022 Using the new OpenCL (Open Computing Language) standard, you can write applications that access all available programming resources: CPUs, GPUs, and other processors such as DSPs and the Cell/B.E. processor. Already implemented by Apple, AMD, Intel, IBM, NVIDIA, and other leaders, OpenCL has outstanding potential for PCs, servers, handheld/embedded devices, high performance computing, and even cloud systems. This is the first comprehensive, authoritative, and practical guide to OpenCL 1.1 specifically for working developers and software architects. Written by five leading OpenCL authorities, OpenCL Programming Guide covers the entire specification. It reviews key use cases, shows how OpenCL can express a wide range of parallel algorithms, and offers complete reference material on both the API and OpenCL C programming language. Through complete case studies and downloadable code examples, the authors show how to write complex parallel programs that decompose workloads across many different devices. They also present all the essentials of OpenCL software performance optimization, including probing and adapting to hardware. Coverage includes Understanding OpenCL's architecture, concepts, terminology, goals, and rationale Programming with OpenCL C and the runtime API Using buffers, sub-buffers, images, samplers, and events Sharing and synchronizing data with OpenGL and Microsoft's Direct3D Simplifying development with the C++ Wrapper API Using OpenCL Embedded Profiles to support devices ranging from cellphones to supercomputer nodes Case studies dealing with physics simulation; image and signal processing, such as image histograms, edge detection filters, Fast Fourier Transforms, and optical flow; math libraries, such as matrix multiplication and high-performance sparse matrix multiplication; and more Source code for this book is available at <https://code.google.com/p/opencl-book-samples/>

OpenCL in Action Sep 23 2022 Summary OpenCL in Action is a thorough, hands-on presentation of OpenCL, with an eye toward showing developers how to build high-performance applications of their own. It begins by presenting the core concepts behind OpenCL, including vector computing, parallel programming, and multi-threaded operations, and then guides you step-by-step from simple data structures to complex functions. About the Technology Whatever system you have, it probably has more raw processing power than you're using. OpenCL is a high-performance programming language that maximizes computational power by executing on CPUs, graphics processors, and other number-crunching devices. It's perfect for speed-sensitive tasks like vector computing, matrix operations, and graphics acceleration. About this Book OpenCL in Action blends the theory of parallel computing with the practical reality of building high-performance applications using OpenCL. It first guides you through the fundamental data structures in an intuitive manner. Then, it explains techniques for high-speed sorting, image processing, matrix operations, and fast Fourier transform. The book concludes with a deep look at the all-important subject of graphics acceleration. Numerous challenging examples give you different ways to experiment with working code. A background in C or C++ is helpful, but no prior exposure to OpenCL is needed. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book. What's Inside Learn OpenCL step by step Tons of annotated code Tested algorithms for maximum performance ***** Table of Contents PART 1 FOUNDATIONS OF OPENCL PROGRAMMING Introducing OpenCL Host programming: fundamental data structures Host programming: data transfer and partitioning Kernel programming: data types and device memory Kernel programming: operators and functions Image processing Events, profiling, and synchronization Development with C++ Development with Java and Python General coding principles PART 2 CODING PRACTICAL ALGORITHMS IN OPENCL Reduction and sorting Matrices and QR decomposition Sparse matrices Signal processing and the fast Fourier transform PART 3 ACCELERATING OPENGL WITH OPENCL Combining OpenCL and OpenGL Textures and renderbuffers

OpenGL Programming Guide Dec 14 2021 Includes Complete Coverage of the OpenGL® Shading Language! Today's OpenGL software interface enables programmers to produce extraordinarily high-quality computer-generated images and interactive applications using 2D and 3D objects, color images, and programmable shaders. OpenGL® Programming Guide: The Official Guide to Learning OpenGL®, Version 4.3, Eighth Edition, has been almost completely rewritten and provides definitive, comprehensive information on OpenGL and the OpenGL Shading Language. This edition of the best-selling "Red Book" describes the features through OpenGL version 4.3. It also includes updated information and techniques formerly covered in OpenGL® Shading Language (the "Orange Book"). For the first time, this guide completely integrates shader techniques, alongside classic, functioncentric techniques. Extensive new text and code are presented, demonstrating the latest in OpenGL programming techniques. OpenGL® Programming Guide, Eighth Edition, provides clear explanations of OpenGL functionality and techniques, including processing geometric objects with vertex, tessellation, and geometry shaders using geometric transformations and viewing matrices; working with pixels and texture maps through fragment shaders; and advanced data techniques using framebuffer objects and compute shaders. New OpenGL features covered in this edition include Best practices and sample code for taking full advantage of shaders and the entire shading pipeline (including geometry and tessellation shaders) Integration of general computation into the rendering pipeline via compute shaders Techniques for binding multiple shader programs at once during application execution Latest GLSL features for doing advanced shading techniques Additional new techniques for optimizing graphics program performance

OpenGL Shading Language Oct 20 2019 "As the 'Red Book' is known to be the gold standard for OpenGL, the 'Orange Book' is considered to be the gold standard for the OpenGL Shading Language. With Randi's extensive knowledge of OpenGL and GLSL, you can be assured you will be learning from a graphics industry veteran. Within the pages of the second edition you can find topics from beginning shader development to advanced topics such as the spherical harmonic lighting model and more." —David Tommeraasen, CEO/Programmer, Plasma Software "This will be the definitive guide for OpenGL shaders; no other book goes into this detail. Rost has done an excellent job at setting the stage for shader development, what the purpose is, how to do it, and how it all fits together. The book includes great examples and details, and good additional coverage of 2.0 changes!" —Jeffery Galinovsky, Director of Emerging Market Platform Development, Intel Corporation "The coverage in this new edition of the book is pitched just right to help many new shader-writers get started, but with enough deep information for the 'old hands.'" —Marc Olano, Assistant Professor, University of Maryland "This is a really great book on GLSL—well written and organized, very accessible, and with good real-world examples and sample code. The topics flow naturally and easily, explanatory code fragments are inserted in very logical places to illustrate concepts, and all in all, this book makes an excellent tutorial as well as a reference." —John Carey, Chief Technology Officer, C.O.R.E. Feature Animation OpenGL® Shading Language, Second Edition, extensively updated for OpenGL 2.0, is the experienced application

programmer's guide to writing shaders. Part reference, part tutorial, this book thoroughly explains the shift from fixed-functionality graphics hardware to the new era of programmable graphics hardware and the additions to the OpenGL API that support this programmability. With OpenGL and shaders written in the OpenGL Shading Language, applications can perform better, achieving stunning graphics effects by using the capabilities of both the visual processing unit and the central processing unit. In this book, you will find a detailed introduction to the OpenGL Shading Language (GLSL) and the new OpenGL function calls that support it. The text begins by describing the syntax and semantics of this high-level programming language. Once this foundation has been established, the book explores the creation and manipulation of shaders using new OpenGL function calls. OpenGL® Shading Language, Second Edition, includes updated descriptions for the language and all the GLSL entry points added to OpenGL 2.0; new chapters that discuss lighting, shadows, and surface characteristics; and an under-the-hood look at the implementation of RealWorldz, the most ambitious GLSL application to date. The second edition also features 18 extensive new examples of shaders and their underlying algorithms, including Image-based lighting Lighting with spherical harmonics Ambient occlusion Shadow mapping Volume shadows using deferred lighting Ward's BRDF model The color plate section illustrates the power and sophistication of the OpenGL Shading Language. The API Function Reference at the end of the book is an excellent guide to the API entry points that support the OpenGL Shading Language. Also included is a convenient Quick Reference Card to GLSL.

OpenGL ES 3.0 Programming Guide Jan 15 2022 OpenGL® ES™ is the industry's leading software interface and graphics library for rendering sophisticated 3D graphics on handheld and embedded devices. The newest version, OpenGL ES 3.0, makes it possible to create stunning visuals for new games and apps, without compromising device performance or battery life. In the OpenGL® ES™ 3.0 Programming Guide, Second Edition, the authors cover the entire API and Shading Language. They carefully introduce OpenGL ES 3.0 features such as shadow mapping, instancing, multiple render targets, uniform buffer objects, texture compression, program binaries, and transform feedback. Through detailed, downloadable C-based code examples, you'll learn how to set up and program every aspect of the graphics pipeline. Step by step, you'll move from introductory techniques all the way to advanced per-pixel lighting and particle systems. Throughout, you'll find cutting-edge tips for optimizing performance, maximizing efficiency with both the API and hardware, and fully leveraging OpenGL ES 3.0 in a wide spectrum of applications. All code has been built and tested on iOS 7, Android 4.3, Windows (OpenGL ES 3.0 Emulation), and Ubuntu Linux, and the authors demonstrate how to build OpenGL ES code for each platform. Coverage includes EGL API: communicating with the native windowing system, choosing configurations, and creating rendering contexts and surfaces Shaders: creating and attaching shader objects; compiling shaders; checking for compile errors; creating, linking, and querying program objects; and using source shaders and program binaries OpenGL ES Shading Language: variables, types, constructors, structures, arrays, attributes, uniform blocks, I/O variables, precision qualifiers, and invariance Geometry, vertices, and primitives: inputting geometry into the pipeline, and assembling it into primitives 2D/3D, Cubemap, Array texturing: creation, loading, and rendering; texture wrap modes, filtering, and formats; compressed textures, sampler objects, immutable textures, pixel unpack buffer objects, and mipmapping Fragment shaders: multitexturing, fog, alpha test, and user clip planes Fragment operations: scissor, stencil, and depth tests; multisampling, blending, and dithering Framebuffer objects: rendering to offscreen surfaces for advanced effects Advanced rendering: per-pixel lighting, environment mapping, particle systems, image post-processing, procedural textures, shadow mapping, terrain, and projective texturing Sync objects and fences: synchronizing within host application and GPU execution This edition of the book includes a color insert of the OpenGL ES 3.0 API and OpenGL ES Shading Language 3.0 Reference Cards created by Khronos. The reference cards contain a complete list of all of the functions in OpenGL ES 3.0 along with all of the types, operators, qualifiers, built-ins, and functions in the OpenGL ES Shading Language.

GPU Pro 6 Jun 27 2020 The latest edition of this bestselling game development reference offers proven tips and techniques for the real-time rendering of special effects and visualization data that are useful for beginners and seasoned game and graphics programmers alike. Exploring recent developments in the rapidly evolving field of real-time rendering, GPU Pro6: Advanced Rendering Techniques assembles a high-quality collection of cutting-edge techniques for advanced graphics processing unit (GPU) programming. It incorporates contributions from more than 45 experts who cover the latest developments in graphics programming for games and movies. The book covers advanced rendering techniques that run on the DirectX or OpenGL runtimes, as well as on any other runtime with any language available. It details the specific challenges involved in creating games across the most common consumer software platforms such as PCs, video consoles, and mobile devices. The book includes coverage of geometry manipulation; rendering techniques, handheld devices programming, effects in image space, shadows, 3D engine design, graphics-related tools, and environmental effects. It also includes a dedicated section on general purpose GPU programming that covers CUDA, DirectCompute, and OpenCL examples. In color throughout, GPU Pro6 presents ready-to-use ideas and procedures that can help solve many of your daily graphics programming challenges. Example programs with downloadable source code are also provided on the book's CRC Press web page.

An Introduction to Parallel Programming Feb 04 2021 An Introduction to Parallel Programming, Second Edition presents a tried-and-true tutorial approach that shows students how to develop effective parallel programs with MPI, Pthreads and OpenMP. As the first undergraduate text to directly address compiling and running parallel programs on multi-core and cluster architecture, this second edition carries forward its clear explanations for designing, debugging and evaluating the performance of distributed and shared-memory programs while adding coverage of accelerators via new content on GPU programming and heterogeneous programming. New and improved user-friendly exercises teach students how to compile, run and modify example programs. Takes a tutorial approach, starting with small programming examples and building progressively to more challenging examples Explains how to develop parallel programs using MPI, Pthreads and OpenMP programming models A robust package of online ancillaries for instructors and students includes lecture slides, solutions manual, downloadable source code, and an image bank New to this edition: New chapters on GPU programming and heterogeneous programming New examples and exercises related to parallel algorithms

Introduction to Parallel Computing Mar 25 2020 Advancements in microprocessor architecture, interconnection technology, and software development have fueled rapid growth in parallel and distributed computing. However, this development is only of practical benefit if it is accompanied by progress in the design, analysis and programming of parallel algorithms. This concise textbook provides, in one place, three mainstream parallelization approaches, Open MPP, MPI and OpenCL, for multicore computers, interconnected computers and graphical processing units. An overview of practical parallel computing and principles will enable the reader to design efficient parallel programs for solving various computational problems on state-of-the-art personal computers and computing clusters. Topics covered range from parallel algorithms, programming tools, OpenMP, MPI and OpenCL, followed by experimental measurements of parallel programs' run-times, and by engineering analysis of obtained results for improved parallel execution performances. Many examples and exercises support the exposition.

Multicore and GPU Programming Jan 23 2020 Multicore and GPU Programming offers broad coverage of the key parallel computing skillsets: multicore CPU programming and manycore "massively parallel" computing. Using threads, OpenMP, MPI, and CUDA, it teaches the design and development of software capable of taking advantage of today's computing platforms incorporating CPU and GPU hardware and explains how to transition from sequential programming to a parallel computing paradigm. Presenting material refined over more than a decade of teaching parallel computing, author Gerassimos Barlas minimizes the challenge with multiple examples, extensive case studies, and full source code. Using this book, you can develop programs that run over distributed memory machines using MPI, create multi-threaded applications with either libraries or directives, write optimized applications that balance the workload between available computing resources, and profile and debug programs targeting multicore machines. Comprehensive coverage of all major multicore programming tools, including threads, OpenMP, MPI, and CUDA Demonstrates parallel programming design patterns and examples of how different tools and paradigms can be integrated for superior performance Particular focus on the emerging area of divisible load theory and its impact on load balancing and distributed systems Download source code, examples, and instructor support materials on the book's

companion website

Patterns for Parallel Programming May 07 2021 The Parallel Programming Guide for Every Software Developer From grids and clusters to next-generation game consoles, parallel computing is going mainstream. Innovations such as Hyper-Threading Technology, HyperTransport Technology, and multicore microprocessors from IBM, Intel, and Sun are accelerating the movement's growth. Only one thing is missing: programmers with the skills to meet the soaring demand for parallel software. That's where Patterns for Parallel Programming comes in. It's the first parallel programming guide written specifically to serve working software developers, not just computer scientists. The authors introduce a complete, highly accessible pattern language that will help any experienced developer "think parallel"-and start writing effective parallel code almost immediately. Instead of formal theory, they deliver proven solutions to the challenges faced by parallel programmers, and pragmatic guidance for using today's parallel APIs in the real world. Coverage includes: Understanding the parallel computing landscape and the challenges faced by parallel developers Finding the concurrency in a software design problem and decomposing it into concurrent tasks Managing the use of data across tasks Creating an algorithm structure that effectively exploits the concurrency you've identified Connecting your algorithmic structures to the APIs needed to implement them Specific software constructs for implementing parallel programs Working with today's leading parallel programming environments: OpenMP, MPI, and Java Patterns have helped thousands of programmers master object-oriented development and other complex programming technologies. With this book, you will learn that they're the best way to master parallel programming too.

CUDA Programming Feb 16 2022 'CUDA Programming' 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.

Numerical Computations with GPUs Apr 25 2020 This book brings together research on numerical methods adapted for Graphics Processing Units (GPUs). It explains recent efforts to adapt classic numerical methods, including solution of linear equations and FFT, for massively parallel GPU architectures. This volume consolidates recent research and adaptations, covering widely used methods that are at the core of many scientific and engineering computations. Each chapter is written by authors working on a specific group of methods; these leading experts provide mathematical background, parallel algorithms and implementation details leading to reusable, adaptable and scalable code fragments. This book also serves as a GPU implementation manual for many numerical algorithms, sharing tips on GPUs that can increase application efficiency. The valuable insights into parallelization strategies for GPUs are supplemented by ready-to-use code fragments. Numerical Computations with GPUs targets professionals and researchers working in high performance computing and GPU programming. Advanced-level students focused on computer science and mathematics will also find this book useful as secondary text book or reference.

OpenGL Programming Guide May 27 2020 Please note that this title's color insert (referred to as "Plates" within the text) is not available for this digital product. OpenGL is a powerful software interface used to produce high-quality, computer-generated images and interactive applications using 2D and 3D objects, bitmaps, and color images. The OpenGL® Programming Guide, Seventh Edition, provides definitive and comprehensive information on OpenGL and the OpenGL Utility Library. The previous edition covered OpenGL through Version 2.1. This seventh edition of the best-selling "red book" describes the latest features of OpenGL Versions 3.0 and 3.1. You will find clear explanations of OpenGL functionality and many basic computer graphics techniques, such as building and rendering 3D models; interactively viewing objects from different perspective points; and using shading, lighting, and texturing effects for greater realism. In addition, this book provides in-depth coverage of advanced techniques, including texture mapping, antialiasing, fog and atmospheric effects, NURBS, image processing, and more. The text also explores other key topics such as enhancing performance, OpenGL extensions, and cross-platform techniques. This seventh edition has been updated to include the newest features of OpenGL Versions 3.0 and 3.1, including Using framebuffer objects for off-screen rendering and texture updates Examples of the various new buffer object types, including uniform-buffer objects, transform feedback buffers, and vertex array objects Using texture arrays to increase performance when using numerous textures Efficient rendering using primitive restart and conditional rendering Discussion of OpenGL's deprecation mechanism and how to verify your programs for future versions of OpenGL This edition continues the discussion of the OpenGL Shading Language (GLSL) and explains the mechanics of using this language to create complex graphics effects and boost the computational power of OpenGL. The OpenGL Technical Library provides tutorial and reference books for OpenGL. The Library enables programmers to gain a practical understanding of OpenGL and shows them how to unlock its full potential. Originally developed by SGI, the Library continues to evolve under the auspices of the Khronos OpenGL ARB Working Group, an industry consortium responsible for guiding the evolution of OpenGL and related technologies.

Parallel Programming with MPI Nov 01 2020 Mathematics of Computing -- Parallelism.

Heterogeneous Computing Architectures Apr 06 2021 Heterogeneous Computing Architectures: Challenges and Vision provides an updated vision of the state-of-the-art of heterogeneous computing systems, covering all the aspects related to their design: from the architecture and programming models to hardware/software integration and orchestration to real-time and security requirements. The transitions from multicore processors, GPU computing, and Cloud computing are not separate trends, but aspects of a single trend-mainstream; computers from desktop to smartphones are being permanently transformed into heterogeneous supercomputer clusters. The reader will get an organic perspective of modern heterogeneous systems and their future evolution.

Euro-Par 2012 Parallel Processing Jun 08 2021 This book constitutes the thoroughly refereed proceedings of the 18th International Conference, Euro-Par 2012, held in Rhodes Islands, Greece, in August 2012. The 75 revised full papers presented were carefully reviewed and selected from 228 submissions. The papers are organized in topical sections on support tools and environments; performance prediction and evaluation; scheduling and load balancing; high-performance architectures and compilers; parallel and distributed data management; grid, cluster and cloud computing; peer to peer computing; distributed systems and algorithms; parallel and distributed programming; parallel numerical algorithms; multicore and manycore programming; theory and algorithms for parallel computation; high performance network and communication; mobile and ubiquitous computing; high performance and scientific applications; GPU and accelerators computing.

Heterogeneous Computing with OpenCL 2.0 Apr 18 2022 Heterogeneous Computing with OpenCL 2.0 teaches OpenCL and parallel programming for complex systems that may include a variety of device architectures: multi-core CPUs, GPUs, and fully-integrated Accelerated Processing Units (APUs). This fully-revised edition includes the latest enhancements in OpenCL 2.0 including: . Shared virtual memory to increase programming flexibility and reduce data transfers that consume resources . Dynamic parallelism which reduces processor load and avoids bottlenecks . Improved imaging support and integration with OpenGL Designed to work on multiple platforms, OpenCL will help you more effectively program for a heterogeneous future. Written by leaders in the parallel computing and OpenCL communities, this book explores memory spaces, optimization techniques, extensions, debugging and profiling. Multiple case studies and examples illustrate high-performance algorithms, distributing work across heterogeneous systems, embedded domain-specific languages, and will give you hands-on OpenCL experience to address a range of fundamental parallel algorithms. Updated content to cover the latest developments in OpenCL 2.0, including improvements in memory handling, parallelism, and imaging support Explanations of principles and strategies to learn parallel programming with OpenCL, from understanding the abstraction models to thoroughly testing and debugging complete applications Example code covering image analytics, web plugins, particle simulations, video editing, performance optimization, and more

OpenGL Programming Guide Nov 20 2019 Explaining how graphics programs using Release 1.1, the latest release of OpenGL, this book presents the overall structure of OpenGL and discusses in detail every OpenGL feature including the new features introduced in Release 1.1. Numerous programming examples in C show how to use OpenGL functions. Also includes 16 pages of full-color examples.

OpenCL Programming by Example Oct 24 2022 This book follows an example-driven, simplified, and practical approach to using OpenCL for general purpose GPU programming. If you are a beginner in parallel

programming and would like to quickly accelerate your algorithms using OpenCL, this book is perfect for you! You will find the diverse topics and case studies in this book interesting and informative. You will only require a good knowledge of C programming for this book, and an understanding of parallel implementations will be useful, but not necessary.

Heterogeneous Computing with OpenCL Nov 25 2022 Heterogeneous Computing with OpenCL, Second Edition teaches OpenCL and parallel programming for complex systems that may include a variety of device architectures: multi-core CPUs, GPUs, and fully-integrated Accelerated Processing Units (APUs) such as AMD Fusion technology. It is the first textbook that presents OpenCL programming appropriate for the classroom and is intended to support a parallel programming course. Students will come away from this text with hands-on experience and significant knowledge of the syntax and use of OpenCL to address a range of fundamental parallel algorithms. Designed to work on multiple platforms and with wide industry support, OpenCL will help you more effectively program for a heterogeneous future. Written by leaders in the parallel computing and OpenCL communities, Heterogeneous Computing with OpenCL explores memory spaces, optimization techniques, graphics interoperability, extensions, and debugging and profiling. It includes detailed examples throughout, plus additional online exercises and other supporting materials that can be downloaded at http://www.heterogeneouscompute.org/?page_id=7 This book will appeal to software engineers, programmers, hardware engineers, and students/advanced students. Explains principles and strategies to learn parallel programming with OpenCL, from understanding the four abstraction models to thoroughly testing and debugging complete applications. Covers image processing, web plugins, particle simulations, video editing, performance optimization, and more. Shows how OpenCL maps to an example target architecture and explains some of the tradeoffs associated with mapping to various architectures Addresses a range of fundamental programming techniques, with multiple examples and case studies that demonstrate OpenCL extensions for a variety of hardware platforms

Parallel Programming with OpenACC Sep 11 2021 Parallel Programming with OpenACC is a modern, practical guide to implementing dependable computing systems. The book explains how anyone can use OpenACC to quickly ramp-up application performance using high-level code directives called pragmas. The OpenACC directive-based programming model is designed to provide a simple, yet powerful, approach to accelerators without significant programming effort. Author Rob Farber, working with a team of expert contributors, demonstrates how to turn existing applications into portable GPU accelerated programs that demonstrate immediate speedups. The book also helps users get the most from the latest NVIDIA and AMD GPU plus multicore CPU architectures (and soon for Intel® Xeon Phi™ as well). Downloadable example codes provide hands-on OpenACC experience for common problems in scientific, commercial, big-data, and real-time systems. Topics include writing reusable code, asynchronous capabilities, using libraries, multicore clusters, and much more. Each chapter explains how a specific aspect of OpenACC technology fits, how it works, and the pitfalls to avoid. Throughout, the book demonstrates how the use of simple working examples that can be adapted to solve application needs. Presents the simplest way to leverage GPUs to achieve application speedups Shows how OpenACC works, including working examples that can be adapted for application needs Allows readers to download source code and slides from the book's companion web page

PThreads Programming Jul 09 2021 With threads programming, multiple tasks run concurrently within the same program. They can share a single CPU as processes do or take advantage of multiple CPUs when available. They provide a clean way to divide the tasks of a program while sharing data.

Intel Xeon Phi Coprocessor Architecture and Tools Jan 03 2021 Intel® Xeon Phi™ Coprocessor Architecture and Tools: The Guide for Application Developers provides developers a comprehensive introduction and in-depth look at the Intel Xeon Phi coprocessor architecture and the corresponding parallel data structure tools and algorithms used in the various technical computing applications for which it is suitable. It also examines the source code-level optimizations that can be performed to exploit the powerful features of the processor. Xeon Phi is at the heart of world's fastest commercial supercomputer, which thanks to the massively parallel computing capabilities of Intel Xeon Phi processors coupled with Xeon Phi coprocessors attained 33.86 teraflops of benchmark performance in 2013. Extracting such stellar performance in real-world applications requires a sophisticated understanding of the complex interaction among hardware components, Xeon Phi cores, and the applications running on them. In this book, Rezaur Rahman, an Intel leader in the development of the Xeon Phi coprocessor and the optimization of its applications, presents and details all the features of Xeon Phi core design that are relevant to the practice of application developers, such as its vector units, hardware multithreading, cache hierarchy, and host-to-coprocessor communication channels. Building on this foundation, he shows developers how to solve real-world technical computing problems by selecting, deploying, and optimizing the available algorithms and data structure alternatives matching Xeon Phi's hardware characteristics. From Rahman's practical descriptions and extensive code examples, the reader will gain a working knowledge of the Xeon Phi vector instruction set and the Xeon Phi microarchitecture whereby cores execute 512-bit instruction streams in parallel.

OpenGL Programming Guide Aug 30 2020 Complete Coverage of OpenGL 4.5--the Latest Version (Includes 4.5, 4.4, SPIR-V, and Extensions) The latest version of today's leading worldwide standard for computer graphics, OpenGL 4.5 delivers significant improvements in application efficiency, flexibility, and performance. OpenGL 4.5 is an exceptionally mature and robust platform for programming high-quality computer-generated images and interactive applications using 2D and 3D objects, color images, and shaders. OpenGL Programming Guide, Ninth Edition, presents definitive, comprehensive information on OpenGL 4.5, 4.4, SPIR-V, OpenGL extensions, and the OpenGL Shading Language. It will serve you for as long as you write or maintain OpenGL code. This edition of the best-selling "Red Book" fully integrates shader techniques alongside classic, function-centric approaches, and contains extensive code examples that demonstrate modern techniques. Starting with the fundamentals, its wide-ranging coverage includes drawing, color, pixels, fragments, transformations, textures, framebuffers, light and shadow, and memory techniques for advanced rendering and nongraphical applications. It also offers discussions of all shader stages, including thorough explorations of tessellation, geometric, and compute shaders. New coverage in this edition includes Thorough coverage of OpenGL 4.5 Direct State Access (DSA), which overhauls the OpenGL programming model and how applications access objects Deeper discussions and more examples of shader functionality and GPU processing, reflecting industry trends to move functionality onto graphics processors Demonstrations and examples of key features based on community feedback and suggestions Updated appendixes covering the latest OpenGL libraries, related APIs, functions, variables, formats, and debugging and profiling techniques

CUDA Application Design and Development Jul 29 2020 The book then details the thought behind CUDA and teaches how to create, analyze, and debug CUDA applications. Throughout, the focus is on software engineering issues: how to use CUDA in the context of existing application code, with existing compilers, languages, software tools, and industry-standard API libraries."--Pub. desc.

Programming Massively Parallel Processors Jun 20 2022 Programming Massively Parallel Processors: A Hands-on Approach, Second Edition, teaches students how to program massively parallel processors. It offers a detailed discussion of various techniques for constructing parallel programs. Case studies are used to demonstrate the development process, which begins with computational thinking and ends with effective and efficient parallel programs. This guide shows both student and professional alike the basic concepts of parallel programming and GPU architecture. Topics of performance, floating-point format, parallel patterns, and dynamic parallelism are covered in depth. This revised edition contains more parallel programming examples, commonly-used libraries such as Thrust, and explanations of the latest tools. It also provides new coverage of CUDA 5.0, improved performance, enhanced development tools, increased hardware support, and more; increased coverage of related technology, OpenCL and new material on algorithm patterns, GPU clusters, host programming, and data parallelism; and two new case studies (on MRI reconstruction and molecular visualization) that explore the latest applications of CUDA and GPUs for scientific research and high-performance computing. This book should be a valuable resource for advanced students, software engineers, programmers, and hardware engineers. New coverage of CUDA 5.0, improved performance, enhanced development tools, increased hardware support, and more Increased coverage of related technology, OpenCL and new material on algorithm patterns, GPU clusters, host programming, and data parallelism Two new case studies (on MRI reconstruction and molecular visualization) explore the latest applications of CUDA and GPUs for scientific research and high-performance computing

CUDA by Example May 19 2022 CUDA is a computing architecture designed to facilitate the development of parallel programs. In conjunction with a comprehensive software platform, the CUDA Architecture enables programmers to draw on the immense power of graphics processing units (GPUs) when building high-performance applications. GPUs, of course, have long been available for demanding graphics and game applications. CUDA now brings this valuable resource to programmers working on applications in other domains, including science, engineering, and finance. No knowledge of graphics programming is required—just the ability to program in a modestly extended version of C. CUDA by Example, written by two senior members of the CUDA software platform team, shows programmers how to employ this new technology. The authors introduce each area of CUDA development through working examples. After a concise introduction to the CUDA platform and architecture, as well as a quick-start guide to CUDA C, the book details the techniques and trade-offs associated with each key CUDA feature. You'll discover when to use each CUDA C extension and how to write CUDA software that delivers truly outstanding performance. Major topics covered include Parallel programming Thread cooperation Constant memory and events Texture memory Graphics interoperability Atomics Streams CUDA C on multiple GPUs Advanced atomics Additional CUDA resources All the CUDA software tools you'll need are freely available for download from NVIDIA. <http://developer.nvidia.com/object/cuda-by-example.html>

Vulkan Programming Guide Mar 17 2022 The Definitive Vulkan™ Developer's Guide and Reference: Master the Next-Generation Specification for Cross-Platform Graphics The next generation of the OpenGL specification, Vulkan, has been redesigned from the ground up, giving applications direct control over GPU acceleration for unprecedented performance and predictability. Vulkan™ Programming Guide is the essential, authoritative reference to this new standard for experienced graphics programmers in all Vulkan environments. Vulkan API lead Graham Sellers (with contributions from language lead John Kessenich) presents example-rich introductions to the portable Vulkan API and the new SPIR-V shading language. The author introduces Vulkan, its goals, and the key concepts framing its API, and presents a complex rendering system that demonstrates both Vulkan's uniqueness and its exceptional power. You'll find authoritative coverage of topics ranging from drawing to memory, and threading to compute shaders. The author especially shows how to handle tasks such as synchronization, scheduling, and memory management that are now the developer's responsibility. Vulkan™ Programming Guide introduces powerful 3D development techniques for fields ranging from video games to medical imaging, and state-of-the-art approaches to solving challenging scientific compute problems. Whether you're upgrading from OpenGL or moving to open-standard graphics APIs for the first time, this guide will help you get the results and performance you're looking for. Coverage includes Extensively tested code examples to demonstrate Vulkan's capabilities and show how it differs from OpenGL Expert guidance on getting started and working with Vulkan's new memory system Thorough discussion of queues, commands, moving data, and presentation Full explanations of the SPIR-V binary shading language and compute/graphics pipelines Detailed discussions of drawing commands, geometry and fragment processing, synchronization primitives, and reading Vulkan data into applications A complete case study application: deferred rendering using complex multi-pass architecture and multiple processing queues Appendixes presenting Vulkan functions and SPIR-V opcodes, as well as a complete Vulkan glossary Example code can be found here: [Example code can be found here: https://github.com/vulkanprogrammingguide/examples](https://github.com/vulkanprogrammingguide/examples)

Intel Xeon Phi Coprocessor High Performance Programming Feb 22 2020 Authors Jim Jeffers and James Reinders spent two years helping educate customers about the prototype and pre-production hardware before Intel introduced the first Intel Xeon Phi coprocessor. They have distilled their own experiences coupled with insights from many expert customers, Intel Field Engineers, Application Engineers and Technical Consulting Engineers, to create this authoritative first book on the essentials of programming for this new architecture and these new products. This book is useful even before you ever touch a system with an Intel Xeon Phi coprocessor. To ensure that your applications run at maximum efficiency, the authors emphasize key techniques for programming any modern parallel computing system whether based on Intel Xeon processors, Intel Xeon Phi coprocessors, or other high performance microprocessors. Applying these techniques will generally increase your program performance on any system, and better prepare you for Intel Xeon Phi coprocessors and the Intel MIC architecture. A practical guide to the essentials of the Intel Xeon Phi coprocessor Presents best practices for portable, high-performance computing and a familiar and proven threaded, scalar-vector programming model Includes simple but informative code examples that explain the unique aspects of this new highly parallel and high performance computational product Covers wide vectors, many cores, many threads and high bandwidth cache/memory architecture

The CUDA Handbook Dec 02 2020 The CUDA Handbook begins where CUDA by Example (Addison-Wesley, 2011) leaves off, discussing CUDA hardware and software in greater detail and covering both CUDA 5.0 and Kepler. Every CUDA developer, from the casual to the most sophisticated, will find something here of interest and immediate usefulness. Newer CUDA developers will see how the hardware processes commands and how the driver checks progress; more experienced CUDA developers will appreciate the expert coverage of topics such as the driver API and context migration, as well as the guidance on how best to structure CPU/GPU data interchange and synchronization. The accompanying open source code—more than 25,000 lines of it, freely available at www.cudahandbook.com—is specifically intended to be reused and repurposed by developers. Designed to be both a comprehensive reference and a practical cookbook, the text is divided into the following three parts: Part I, Overview, gives high-level descriptions of the hardware and software that make CUDA possible. Part II, Details, provides thorough descriptions of every aspect of CUDA, including Memory Streams and events Models of execution, including the dynamic parallelism feature, new with CUDA 5.0 and SM 3.5 The streaming multiprocessors, including descriptions of all features through SM 3.5 Programming multiple GPUs Texturing The source code accompanying Part II is presented as reusable microbenchmarks and microdemos, designed to expose specific hardware characteristics or highlight specific use cases. Part III, Select Applications, details specific families of CUDA applications and key parallel algorithms, including Streaming workloads Reduction Parallel prefix sum (Scan) N-body Image Processing These algorithms cover the full range of potential CUDA applications.

Data Parallel C++ Aug 10 2021 Learn how to accelerate C++ programs using data parallelism. This open access book enables C++ programmers to be at the forefront of this exciting and important new development that is helping to push computing to new levels. It is full of practical advice, detailed explanations, and code examples to illustrate key topics. Data parallelism in C++ enables access to parallel resources in a modern heterogeneous system, freeing you from being locked into any particular computing device. Now a single C++ application can use any combination of devices—including GPUs, CPUs, FPGAs and AI ASICs—that are suitable to the problems at hand. This book begins by introducing data parallelism and foundational topics for effective use of the SYCL standard from the Khronos Group and Data Parallel C++ (DPC++), the open source compiler used in this book. Later chapters cover advanced topics including error handling, hardware-specific programming, communication and synchronization, and memory model considerations. Data Parallel C++ provides you with everything needed to use SYCL for programming heterogeneous systems. What You'll Learn Accelerate C++ programs using data-parallel programming Target multiple device types (e.g. CPU, GPU, FPGA) Use SYCL and SYCL compilers Connect with computing's heterogeneous future via Intel's oneAPI initiative Who This Book Is For Those new data-parallel programming and computer programmers interested in data-parallel programming using C++.

GPU Parallel Program Development Using CUDA Mar 05 2021 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.

GPU Pro 5 Dec 22 2019 In GPU Pro5: Advanced Rendering Techniques, section editors Wolfgang Engel, Christopher Oat, Carsten Dachsbacher, Michal Valient, Wessam Bahnassi, and Marius Bjorge have once again assembled a high-quality collection of cutting-edge techniques for advanced graphics processing unit (GPU) programming. Divided into six sections, the book covers rendering, lighting, effects in image space, mobile devices, 3D engine design, and compute. It explores rasterization of liquids, ray tracing of art assets that would otherwise be used in a rasterized engine, physically based area lights, volumetric light effects, screen-space grass, the usage of quaternions, and a quadtree implementation on the GPU. It also addresses the latest developments in deferred lighting on mobile devices, OpenCL optimizations for mobile devices, morph targets, and tiled deferred blending methods. In color throughout, GPU Pro5 is the only book that incorporates contributions from more than 50 experts who cover the latest developments in graphics programming for games and movies. It presents ready-to-use ideas and procedures that can help solve many of your daily graphics programming challenges. Example programs with source code are provided on the book's CRC Press web page.

GPU Pro 360 Guide to Rendering Sep 18 2019 Wolfgang Engel's GPU Pro 360 Guide to Rendering gathers all the cutting-edge information from his previous seven GPU Pro volumes into a convenient single source anthology that covers real-time rendering. This volume is complete with 32 articles by leading programmers that focus on the ability of graphics processing units to process and generate rendering in exciting ways. GPU Pro 360 Guide to Rendering is comprised of ready-to-use ideas and efficient procedures that can help solve many rendering programming challenges that may arise. Key Features: Presents tips and tricks on real-time rendering of special effects and visualization data on common consumer software platforms such as PCs, video consoles, and mobile devices Covers specific challenges involved in creating games on various platforms Explores the latest developments in the rapidly evolving field of real-time rendering Takes a practical approach that helps graphics programmers solve their daily challenges

Learning OpenGL ES for iOS Aug 18 2019 Get Started Fast with Modern OpenGL ES Graphics Programming for iPhone, iPod touch, and iPad OpenGL ES technology underlies the user interface and graphical capabilities of Apple's iPhone, iPod touch, and iPad—as well as devices ranging from video-game consoles and aircraft-cockpit displays to non-Apple smartphones. In this friendly, thorough introduction, Erik M. Buck shows how to make the most of Open GL ES in Apple's iOS environment. This highly anticipated title focuses on modern, efficient approaches that use the newest versions of OpenGL ES, helping you avoid the irrelevant, obsolete, and misleading techniques that litter the Internet. Buck embraces Objective-C and Cocoa Touch, showing how to leverage Apple's powerful, elegant GLKit framework to maximize your productivity, achieve tight platform integration, and deliver exceptionally polished apps. If you've written C or C++ code and know object-oriented programming basics, this title brings together everything you need to fully master OpenGL ES graphics for iOS—including downloadable examples specifically designed to jumpstart your own projects. Coverage includes • Understanding core OpenGL ES computer graphics concepts and iOS graphics architecture • Integrating Cocoa Touch with OpenGL ES to leverage the power of Apple's platform • Creating textures from start to finish: opacity, blending, multi-texturing, and compression • Simulating ambient, diffuse, and specular light • Using transformations to render 3D geometric objects from any point of view • Animating scenes by controlling time through application logic • Partitioning data to draw expansive outdoor scenes with rolling terrain • Detecting and handling user interaction with 3D geometry • Implementing special effects ranging from skyboxes to particles and billboards • Systematically optimizing graphics performance • Understanding the essential linear algebra concepts used in computer graphics • Designing and constructing a complete simulation that incorporates everything you've learned

Introduction to High Performance Scientific Computing Oct 12 2021 Based on a course developed by the author, Introduction to High Performance Scientific Computing introduces methods for adding parallelism to numerical methods for solving differential equations. It contains exercises and programming projects that facilitate learning as well as examples and discussions based on the C programming language, with additional comments for those already familiar with C++. The text provides an overview of concepts and algorithmic techniques for modern scientific computing and is divided into six self-contained parts that can be assembled in any order to create an introductory course using available computer hardware. Part I introduces the C programming language for those not already familiar with programming in a compiled language. Part II describes parallelism on shared memory architectures using OpenMP. Part III details parallelism on computer clusters using MPI for coordinating a computation. Part IV demonstrates the use of graphical programming units (GPUs) to solve problems using the CUDA language for NVIDIA graphics cards. Part V addresses programming on GPUs for non-NVIDIA graphics cards using the OpenCL framework. Finally, Part VI contains a brief discussion of numerical methods and applications, giving the reader an opportunity to test the methods on typical computing problems.