

Parallel Programming In C With Mpi And Openmp

If you have a working knowledge of Haskell, this hands-on book shows you how to use the language's many APIs and frameworks for writing both parallel and concurrent programs. You'll learn how parallelism exploits multicore processors to speed up computation-heavy programs, and how concurrency enables you to write programs with threads for multiple interactions. Author Simon Marlow walks you through the process with lots of code examples that you can run, experiment with, and extend. Divided into separate sections on Parallel and Concurrent Haskell, this book also includes exercises to help you become familiar with the concepts presented: Express parallelism in Haskell with the Eval monad and Evaluation Strategies Parallelize ordinary Haskell code with the Par monad Build parallel array-based computations, using the Repa library Use the Accelerate library to run computations directly on the GPU Work with basic interfaces for writing concurrent code Build trees of threads for larger and more complex programs Learn how to build high-speed concurrent network servers Write distributed programs that run on multiple machines in a network

This open access book is a modern guide for all C++ programmers to learn Threading Building Blocks (TBB). Written by TBB and parallel programming experts, this book reflects their collective decades of experience in developing and teaching parallel programming with TBB, offering their insights in an approachable manner. Throughout the book the authors present numerous examples and best practices to help you become an effective TBB programmer and leverage the power of parallel systems. Pro TBB starts with the basics, explaining parallel algorithms and C++'s built-in standard template library for parallelism. You'll learn the key concepts of managing memory, working with data structures and how to handle typical issues with synchronization. Later chapters apply these ideas to complex systems to explain performance tradeoffs, mapping common parallel patterns, controlling threads and overhead, and extending TBB to program heterogeneous systems or system-on-chips. What You'll Learn Use Threading Building Blocks to produce code that is portable, simple, scalable, and more understandable Review best practices for parallelizing computationally intensive tasks in your applications Integrate TBB with other threading packages Create scalable, high performance data-parallel programs Work with generic programming to write efficient algorithms Who This Book Is For C++ programmers learning to run applications on multicore systems, as well as C or C++ programmers without much experience with templates. No previous experience with parallel programming or multicore processors is required.

Parallel Programming: Concepts and Practice provides an upper level introduction to parallel programming. In addition to covering general parallelism concepts, this text teaches practical programming skills for both shared memory and distributed memory architectures. The authors' open-source system for automated code evaluation provides easy access to parallel computing resources, making the book particularly suitable for classroom settings. Covers parallel programming approaches for single computer nodes and HPC clusters: OpenMP, multithreading, SIMD vectorization, MPI, UPC++ Contains numerous practical parallel programming exercises Includes access to an automated code evaluation tool that enables students the opportunity to program in a web browser and receive immediate feedback on the result validity of their program Features an example-based teaching of concept to enhance learning outcomes

Foreword by Bjarne Stroustrup Software is generally acknowledged to be the single greatest obstacle preventing mainstream adoption of massively-parallel computing. While sequential applications are routinely ported to platforms ranging from PCs to mainframes, most parallel programs only ever run on one type of machine. One reason for this is that most parallel programming systems have failed to insulate their users from the architectures of the machines on which they have run. Those that have been platform-independent have usually also had poor performance. Many researchers now believe that object-oriented languages may offer a solution. By hiding the architecture-specific constructs required for high performance inside platform-independent abstractions, parallel object-oriented programming systems may be able to combine the speed of massively-parallel computing with the comfort of sequential programming. Parallel Programming Using C++ describes fifteen parallel programming systems based on C++, the most popular object-oriented language of today. These systems cover the whole spectrum of parallel programming paradigms, from data parallelism through dataflow and distributed shared memory to message-passing control parallelism. For the parallel programming community, a common parallel application is discussed in each chapter, as part of the description of the system itself. By comparing the implementations of the polygon overlay problem in each system, the reader can get a better sense of their expressiveness and functionality for a common problem. For the systems community, the chapters contain a discussion of the implementation of the various compilers and runtime systems. In addition to discussing the performance of polygon overlay, several of the contributors also discuss the performance of other, more substantial, applications. For the research community, the contributors discuss the motivations for and philosophy of their systems. As well, many of the chapters include critiques that complete the research arc by pointing out possible future research directions. Finally, for the object-oriented community, there are many examples of how encapsulation, inheritance, and polymorphism can be used to control the complexity of developing, debugging, and tuning parallel software.

Parallel Computing: Fundamentals, Applications and New Directions

Mastering DPC++ for Programming of Heterogeneous Systems using C++ and SYCL

Data Parallel C++

Introduction to Parallel Computing

Vector Models for Data-parallel Computing

Programming Massively Parallel Processors

A comprehensive overview of OpenMP, the standard application programming interface for shared memory parallel computing—a reference for students and professionals. "I hope that readers will learn to use the full expressibility and power of OpenMP. This book should provide an excellent introduction to beginners, and the performance section should help those with some experience who want to push OpenMP to its limits." —from the foreword by David J. Kuck, Intel Fellow, Software and Solutions Group, and Director, Parallel and Distributed Solutions, Intel Corporation OpenMP, a portable programming interface for shared memory parallel computers, was adopted as an informal standard in 1997 by computer scientists who wanted a unified model on which to base programs for shared memory systems. OpenMP is now used by many software developers; it offers significant advantages over both hand-threading and MPI. Using OpenMP offers a comprehensive introduction to parallel programming concepts and a detailed overview of OpenMP. Using OpenMP discusses hardware developments, describes where OpenMP is applicable, and compares OpenMP to other programming interfaces for shared and distributed memory parallel architectures. It introduces the individual features of OpenMP, provides many source code examples that demonstrate the use and functionality of the language constructs, and offers tips on writing an efficient OpenMP program. It describes how to use OpenMP in full-scale applications to achieve high performance on large-scale architectures, discussing several case studies in detail, and offers in-depth troubleshooting advice. It explains how OpenMP is translated into explicitly multithreaded code, providing a valuable behind-the-scenes account of OpenMP program performance. Finally, Using OpenMP considers trends likely to influence OpenMP development, offering a glimpse of the possibilities of a future OpenMP 3.0 from the vantage point of the current OpenMP 2.5. With multicore computer use increasing, the need for a comprehensive introduction and overview of the standard interface is clear. Using OpenMP provides an essential reference not only for students at both undergraduate and graduate levels but also for professionals who intend to parallelize existing codes or develop new parallel programs for shared memory computer architectures.

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. 0321228111B08232004 Enhance your enterprise application development skills by mastering parallel programming techniques in .NET and C# Key FeaturesWrite efficient, fine-grained, and scalable parallel code with C# and .NET CoreExperience how parallel programming works by building a powerful applicationLearn the fundamentals of multithreading by working with IIS and KestrelBook Description In today's world, every CPU has a multi-core processor. However, unless your application has implemented parallel programming, it will fail to utilize the hardware's full processing capacity. This book will show you how to write modern software on the optimized and high-performing .NET Core 3 framework using C# 8. Hands-On Parallel Programming with C# 8 and .NET Core 3 covers how to build multithreaded, concurrent, and optimized applications that harness the power of multi-core processors. Once you've understood the fundamentals of threading and concurrency, you'll gain insights into the data structure in .NET Core that supports parallelism. The book will then help you perform asynchronous programming in C# and diagnose and debug parallel code effectively. You'll also get to grips with the new Kestrel server and understand the difference between the IIS and Kestrel operating models. Finally, you'll learn best practices such as test-driven development, and run unit tests on your parallel code. By the end of the book, you'll have developed a deep understanding of the core concepts of concurrency and asynchrony to create responsive applications that are not CPU-intensive. What you will learnAnalyze and break down a problem statement for parallelismExplore the APM and EAP patterns and how to move legacy code to TaskApply reduction techniques to get aggregated resultsCreate PLINQ queries and study the factors that impact their performanceSolve concurrency problems caused by producer-consumer race conditionsDiscover the synchronization primitives available in .NET CoreUnderstand how the threading model works with IIS and KestrelFind out how you can make the most of server resourcesWho this book is for If you want to learn how task parallelism is used to build robust and scalable enterprise architecture, this book is for you. Whether you are a beginner to parallelism in C# or an experienced architect, you'll find this book useful to gain insights into the different threading models supported in .NET Standard and .NET Core. Prior knowledge of C# is required to understand the concepts covered in this book.

An overview of the most prominent contemporary parallel processing programming models, written in a unique tutorial style. With the coming of the parallel computing era, computer scientists have turned their attention to designing programming models that are suited for high-performance parallel computing and supercomputing systems. Programming parallel systems is complicated by the fact that multiple processing units are simultaneously computing and moving data. This book offers an overview of some of the most prominent parallel programming models used in high-performance computing and supercomputing systems today. The chapters describe the programming models in a unique tutorial style rather than using the formal approach taken in the research literature. The aim is to cover a wide range of parallel programming models, enabling the reader to understand what each has to offer. The book begins with a description of the Message Passing Interface (MPI), the most common parallel programming model for distributed memory computing. It goes on to cover one-sided communication models, ranging from low-level runtime libraries (GASNet, OpenSHMEM) to high-level programming models (UPC, GA, Chapel); task-oriented programming models (Charm++, ADLB, Scioto, Swift, CnC) that allow users to describe their computation and data units as tasks so that the runtime system can manage computation and data movement as necessary; and parallel programming models intended for on-node parallelism in the context of multicore architecture or attached accelerators (OpenMP, Cilk Plus, TBB, CUDA, OpenCL). The book will be a valuable resource for graduate students, researchers, and any scientist who works with data sets and large computations. Contributors Timothy Armstrong, Michael G. Burke, Ralph Butler, Bradford L. Chamberlain, Sunita Chandrasekaran, Barbara Chapman, Jeff Daily, James Dinan, Deepak Eachempati, Ian T. Foster, William D. Gropp, Paul Hargrove, Wen-mei Hwu, Nikhil Jain, Laxmikant Kale, David Kirk, Kath Knoke, Ariram Krishnamoorthy, Jeffery A. Kuehn, Alexey Kukanov, Charles E. Leiserson, Jonathan Lifflander, Ewing Lusk, Tim Mattson, Bruce Palmer, Steven C. Pieper, Stephen W. Poole, Arch D. Robison, Frank Schlimbach, Rajeev Thakur, Abhinav Vishnu, Justin M. Wozniak, Michael Wilde, Kathy Yelick, Yili Zheng

A Hands-on Approach

How to Write Parallel Programs

Parallel Programming in MPI and OpenMP

A practical guide with examples in C

Parallel and Distributed Programming Using C++

Parallel Programming in OpenMP

Optimize code for multi-core processors with Intel's Parallel Studio Parallel programming is rapidly becoming a "must-know" skill for developers. Yet, where to start? This teach-yourself tutorial is an ideal starting point for developers who already know Windows C and C++ and are eager to add parallelism to their code. With a focus on applying tools, techniques, and language extensions to implement parallelism, this essential resource teaches you how to write programs for multicore and leverage the power of multicore in your programs. Sharing hands-on case studies and real-world examples, the authors examine the challenges of each project and show you how to overcome them. Explores conversion of serial code to parallel Focuses on implementing Intel Parallel Studio Highlights the benefits of using parallel code Addresses error and performance optimization of code Includes real-world scenarios that illustrate the techniques of advanced parallel programming situations Parallel Programming with Intel Parallel Studio dispels any concerns of difficulty and gets you started creating faster code with Intel Parallel Studio.

The era of practical parallel programming has arrived, marked by the popularity of the MPI and OpenMP software standards and the emergence of commodity clusters as the hardware platform of choice for an increasing number of organizations. This exciting new book, Parallel Programming in C with MPI and OpenMP addresses the needs of students and professionals who want to learn how to design, analyze, implement, and benchmark parallel programs in C using MPI and/or OpenMP. It introduces a rock-solid design methodology with coverage of the most important MPI functions and OpenMP directives. It also demonstrates, through a wide range of examples, how to develop parallel programs that will execute efficiently on today's parallel platforms. If you are an instructor who has adopted the book and would like access to the additional resources, please contact your local sales rep. or Michelle Flomenhofs at: michelle.flomenhofs@mcgraw-hill.com.

About The Book: Your CPU meter shows a problem. One core is running at 100 percent, but all the other cores are idle. Your application is CPU-bound, but you are using only a fraction of the computing power of your multicore system. Is there a way to get better performance? The answer, in a nutshell, is parallel programming. Where you once would have written the kind of sequential code that is familiar to all programmers, you now find that this no longer meets your performance goals. To use your system's CPU resources efficiently, you need to split your application into pieces that can run at the same time. Of course, this is easier said than done. Parallel programming has a reputation for being the domain of experts and a minefield of subtle, hard-to-reproduce software defects. Everyone seems to have a favorite story about a parallel program that did not behave as expected because of a mysterious bug. These stories should inspire a healthy respect for the difficulty of the problems you will face in writing your own parallel programs. Fortunately, help has arrived. The Parallel Patterns Library (PPL) and the Asynchronous Agents Library introduce a new programming model for parallelism that significantly simplifies the job. Behind the scenes are sophisticated algorithms that dynamically distribute computations on multicore architectures. In addition, Microsoft's Visual Studio® 2010 development system includes debugging and analysis tools to support the new parallel programming model.

Parallel Programming in C with MPI and OpenMP Tata McGraw-Hill Education Parallel Programming in C with MPI and OpenMP McGraw-Hill Education

The PARC++ Library

Concepts and Practice

for Multicore and Cluster Systems

PARALLEL PROGRAMMING WITH MICROSOFT VISUAL C++, DESIGN PATTERNS FOR DECOMPOSITION AND COORDINATION O (With CD)

Techniques for Multicore and Multithreaded Programming

Parallel Programming in C with MPI and OpenMP

Data science has taken the world by storm. Every field of study and area of business has been affected as people increasingly realize the value of the incredible quantities of data being generated. But to extract value from those data, one needs to be tra

With the rise of multi-core architecture, parallel programming is an increasingly important topic for software engineers and computer system designers. Written by well-known researchers Larry Snyder and Calvin Lin, this highly anticipated first edition emphasizes the principles underlying parallel computation, explains the various phenomena, and clarifies why these phenomena represent opportunities or barriers to successful parallel programming. Ideal for an advanced upper-level undergraduate course, Principles of Parallel Programming supplies enduring knowledge that will outlive the current hardware and software, aiming to inspire future researchers to build tomorrow's solutions.

Programming is now parallel programming. Much as structured programming revolutionized traditional serial programming decades ago, a new kind of structured programming, based on patterns, is relevant to parallel programming today. Parallel computing experts and industry insiders Michael McCool, Arch Robison, and James Reinders describe how to design and implement maintainable and efficient parallel algorithms using a pattern-based approach. They present both theory and practice, and give detailed concrete examples using multiple programming models. Examples are primarily given using two of the most popular and cutting edge programming models for parallel programming: Threading Building Blocks, and Cilk Plus. These architecture-independent models enable easy integration into existing applications, preserve investments in existing code, and speed the development of parallel applications. Examples from realistic contexts illustrate patterns and themes in parallel algorithm design that are widely applicable regardless of implementation technology. The patterns-based approach offers structure and insight that developers can apply to a variety of parallel programming models Develops a composable, structured, scalable, and machine-independent approach to parallel computing Includes detailed examples in both Cilk Plus and the latest Threading Building Blocks, which support a wide variety of computers

This guide shows Visual C++ programmers how to effectively take advantage of the multicore capabilities of modern PCs using the Microsoft platform.

Structured Parallel Programming

Hands-On Parallel Programming with C# 8 and .NET Core 3

Parallel and Concurrent Programming in Haskell

Patterns for Efficient Computation

Portable Shared Memory Parallel Programming

Hands-On GPU-Accelerated Computer Vision with OpenCV and CUDA

Innovations in hardware architecture, like hyper-threading or multicore processors, mean that parallel computing resources are available for inexpensive desktop computers. In only a few years, many standard software products will be based on concepts of parallel programming implemented on such hardware, and the range of applications will be much broader than that of scientific computing, up to now the main application area for parallel computing. Rauber and Runger take up these recent developments in processor architecture by giving detailed descriptions of parallel programming techniques that are necessary for developing efficient programs for multicore processors as well as for parallel cluster systems and supercomputers. Their book is structured in three main parts, covering all areas of parallel computing: the architecture of parallel systems, parallel programming models and environments, and the implementation of efficient application algorithms. The emphasis lies on parallel programming techniques needed for different architectures. For this second edition, all chapters have been carefully revised. The chapter on architecture of parallel systems has been updated considerably, with a greater emphasis on the architecture of multicore systems and adding new material on the latest developments in computer architecture. Lastly, a completely new chapter on general-purpose GPUs and the corresponding programming techniques has been added. The main goal of the book is to present parallel programming techniques that can be used in many situations for a broad range of application areas and which enable the reader to develop correct and efficient parallel programs. Many examples and exercises are provided to show how to apply the techniques. The book can be used as both a textbook for students and a reference book for professionals. The material presented has been used for courses in parallel programming at different universities for many years.

This volume gives an overview of the state-of-the-art with respect to the development of all types of parallel computers and their application to a wide range of problem areas. The international conference on parallel computing ParCo97 (Parallel Computing 97) was held in Bonn, Germany from 19 to 22 September 1997. The first conference in this biannual series was held in 1983 in Berlin. Further conferences were held in Leiden (The Netherlands), London (UK), Grenoble (France) and Gent (Belgium). From the

outset the aim with the ParCo (Parallel Computing) conferences was to promote the application of parallel computers to solve real life problems. In the case of ParCo97 a new milestone was reached in that more than half of the papers and posters presented were concerned with application aspects. This fact reflects the coming of age of parallel computing. Some 200 papers were submitted to the Program Committee by authors from all over the world. The final programme consisted of four invited papers, 71 contributed scientific/industrial papers and 45 posters. In addition a panel discussion on Parallel Computing and the Evolution of Cyberspace was held. During and after the conference all final contributions were refereed. Only those papers and posters accepted during this final screening process are included in this volume. The practical emphasis of the conference was accentuated by an industrial exhibition where companies demonstrated the newest developments in parallel processing equipment and software. Speakers from participating companies presented papers in industrial sessions in which new developments in parallel computing were reported.

Expert guidance for those programming today's dual-core processors PCs As PC processors explode from one or two to now eight processors, there is an urgent need for programmers to master concurrent programming. This book dives deep into the latest technologies available to programmers for creating professional parallel applications using C#, .NET 4, and Visual Studio 2010. The book covers task-based programming, coordination data structures, PLINQ, thread pools, asynchronous programming model, and more. It also teaches other parallel programming techniques, such as SIMD and vectorization. Teaches programmers professional-level, task-based, parallel programming with C#, .NET 4, and Visual Studio 2010 Covers concurrent collections, coordinated data structures, PLINQ, thread pools, asynchronous programming model, Visual Studio 2010 debugging, and parallel testing and tuning Explores vectorization, SIMD instructions, and additional parallel libraries Master the tools and technology you need to develop thread-safe concurrent applications for multi-core systems, with Professional Parallel Programming with C#.

Parallel and High Performance Computing offers techniques guaranteed to boost your code's effectiveness. Summary Complex calculations, like training deep learning models or running large-scale simulations, can take an extremely long time. Efficient parallel programming can save hours—or even days—of computing time. Parallel and High Performance Computing shows you how to deliver faster run-times, greater scalability, and increased energy efficiency to your programs by mastering parallel techniques for multicore processor and GPU hardware. About the technology Write fast, powerful, energy efficient programs that scale to tackle huge volumes of data. Using parallel programming, your code spreads data processing tasks across multiple CPUs for radically better performance. With a little help, you can create software that maximizes both speed and efficiency. About the book Parallel and High Performance Computing offers techniques guaranteed to boost your code's effectiveness. You'll learn to evaluate hardware architectures and work with industry standard tools such as OpenMP and MPI. You'll master the data structures and algorithms best suited for high performance computing and learn techniques that save energy on handheld devices. You'll even run a massive tsunami simulation across a bank of GPUs. What's inside Planning a new parallel project Understanding differences in CPU and GPU architecture Addressing underperforming kernels and loops Managing applications with batch scheduling About the reader For experienced programmers proficient with a high-performance computing language like C, C++, or Fortran. About the author Robert Robey works at Los Alamos National Laboratory and has been active in the field of parallel computing for over 30 years. Yuliana Zamora is currently a PhD student and Siebel Scholar at the University of Chicago, and has lectured on programming modern hardware at numerous national conferences. Table of Contents PART 1 INTRODUCTION TO PARALLEL COMPUTING 1 Why parallel computing? 2 Planning for parallelization 3 Performance limits and profiling 4 Data design and performance models 5 Parallel algorithms and patterns PART 2 CPU: THE PARALLEL WORKHORSE 6 Vectorization: FLOPs for free 7 OpenMP that performs 8 MPI: The parallel backbone PART 3 GPUS: BUILT TO ACCELERATE 9 GPU architectures and concepts 10 GPU programming model 11 Directive-based GPU programming 12 GPU languages: Getting down to basics 13 GPU profiling and tools PART 4 HIGH PERFORMANCE COMPUTING ECOSYSTEMS 14 Affinity: Truce with the kernel 15 Batch schedulers: Bringing order to chaos 16 File operations for a parallel world 17 Tools and resources for better code

With Examples in R, C++ and CUDA

Parallel and High Performance Computing

An Introduction to Parallel Programming

Object-oriented Parallel Programming in C++

Using OpenMP

Programming Models for Parallel Computing

A complete source of information on almost all aspects of parallel computing from introduction, to architectures, to programming paradigms, to algorithms, to programming standards. It covers traditional Computer Science algorithms, scientific computing algorithms and data intensive algorithms.

In the last few years, courses on parallel computation have been developed and offered in many institutions in the UK, Europe and US as a recognition of the growing significance of this topic in mathematics and computer science. There is a clear need for texts that meet the needs of students and lecturers and this book, based on the author's lecture at ETH Zurich, is an ideal practical student guide to scientific computing on parallel computers working up from a hardware instruction level, to shared memory machines, and finally to distributed memory machines. Aimed at advanced undergraduate and graduate students in applied mathematics, computer science, and engineering, subjects covered include linear algebra, fast Fourier transform, and Monte-Carlo simulations, including examples in C and, in some cases, Fortran. This book is also ideal for practitioners and programmers.

Concurrent C is a superset of C that provides parallel programming facilities such as those for the declaring and creating processes, for process synchronization and interaction, and for process termination and abortion. Concurrent C was designed for the effective utilization of multiprocessors and multicomputers. Concurrent C, as a compile-time option, also works with C++, an object-oriented superset of C.

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++, 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++.

C++ Parallel Programming with Threading Building Blocks

R Programming for Data Science

Sequential and Parallel Programming with C and Fortran

Introduction to Parallel Programming

Build solid enterprise software using task parallelism and multithreading

Parallel Programming Using C++

This text takes complicated and almost unapproachable parallel programming techniques and presents them in a simple, understandable manner. It covers the fundamentals of programming for distributed environments like Internets and Intranets as well as the topic of Web Based Agents.

Parallel Computing for Data Science: With Examples in R, C++ and CUDA is one of the first parallel computing books to concentrate exclusively on parallel data structures, algorithms, software tools, and applications in data science. It includes examples not only from the classic "n observations, p variables" matrix format but also from time series, 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

Mathematics of Computing -- Parallelism.

Parallel Programming with MPI

Parallel Programming with Microsoft Visual C++

Professional Parallel Programming with C#

Pro TBB

The Concurrent C Programming Language

Discover how CUDA allows OpenCV to handle complex and rapidly growing image data processing in computer and machine vision by accessing the power of GPU Key FeaturesExplore examples to leverage the GPU processing power with OpenCV and CUDAEnhance the performance of algorithms on embedded hardware platformsDiscover C++ and Python libraries for GPU accelerationBook Description Computer vision has been revolutionizing a wide range of industries, and OpenCV is the most widely chosen tool for computer vision with its ability to work in multiple programming languages. Nowadays, in computer vision, there is a need to process large images in real time, which is difficult to handle for OpenCV on its own. This is where CUDA comes into the picture, allowing OpenCV to leverage powerful NVIDIA GPUs. This book provides a detailed overview of integrating OpenCV with CUDA for practical applications. To start with, you'll understand GPU programming with CUDA, an essential aspect for computer vision developers who have never worked with GPUs. You'll then move on to exploring OpenCV acceleration with GPUs and CUDA by walking through some practical examples. Once you have got to grips with the core concepts, you'll familiarize yourself with deploying OpenCV applications on NVIDIA Jetson TX1, which is popular for computer vision and deep learning applications. The last chapters of the book explain PyCUDA, a Python library that leverages the power of CUDA and GPUs for accelerations and can be used by computer vision developers who use OpenCV with Python. By the end of this book, you'll have enhanced computer vision applications with the help of this book's hands-on approach. What you will learnUnderstand how to access GPU device properties and capabilities from CUDA programsLearn how to accelerate searching and sorting algorithmsDetect shapes such as lines and circles in imagesExplore object tracking and detection with algorithmsProcess videos using different video analysis techniques in Jetson TX1Access GPU device properties from the PyCUDA programUnderstand how kernel execution worksWho this book is for This book is a go-to guide for you if you are a developer working with OpenCV and want to learn how to process more complex image data by exploiting GPU processing. A thorough understanding of computer vision concepts and programming languages such as C++ or Python is expected.

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

Break into the powerful world of parallel GPU programming with this down-to-earth, practical guide Designed for professionals across multiple industrial sectors, Professional CUDA C Programming presents CUDA -- a parallel computing platform and programming model designed to ease the development of GPU programming -- fundamentals in an easy-to-follow format, and teaches readers how to think in parallel and implement parallel algorithms on GPUs. Each chapter covers a specific topic, and includes workable examples that demonstrate the development process, allowing readers to explore both the "hard" and "soft" aspects of GPU programming. Computing architectures are experiencing a fundamental shift toward scalable parallel computing motivated by application requirements in industry and science. This book demonstrates the challenges of efficiently utilizing compute resources at peak performance, presents modern techniques for tackling these challenges, while increasing accessibility for professionals who are not necessarily parallel programming experts. The CUDA programming model and tools empower developers to write high-performance applications on a scalable, parallel computing platform: the GPU. However, CUDA itself can be difficult to learn without extensive programming experience. Recognized CUDA authorities John Cheng, Max Grossman, and Ty McKeercher guide readers through essential GPU programming skills and best practices in Professional CUDA C Programming, including: CUDA Programming Model GPU Execution Model GPU Memory model Streams, Event and Concurrency Multi-GPU Programming CUDA Domain-Specific Libraries Profiling and Performance Tuning The book makes complex CUDA concepts easy to understand for anyone with knowledge of basic software development with exercises designed to be both readable and high-performance. For the professional seeking entrance to parallel computing and the high-performance computing community, Professional CUDA C Programming is an invaluable resource, with the most current information available on the market.

Software -- Programming Techniques.

Effective techniques for processing complex image data in real time using GPUs

Is Parallel Programming Hard

Design Patterns for Decomposition and Coordination of Multicore Architectures

Parallel Computing for Data Science

Patterns for Parallel Programming

A First Course

In modern computer science, there exists no truly sequential computing system; and most advanced programming is parallel programming. This is particularly evident in modern application domains like scientific computation, data science, machine intelligence, etc. This lucid introductory textbook will be invaluable to students of computer science and technology, acting as a self-contained primer to parallel programming. It takes the reader from introduction to expertise, addressing a broad gamut of issues. It covers different parallel programming styles, describes parallel architecture, includes parallel programming frameworks and techniques, presents algorithmic and analysis techniques and discusses parallel design and performance issues. With its broad coverage, the book can be useful in a wide range of courses; and can also prove useful as a ready reckoner for professionals in the field.

Parallel Programming

Parallel Programming with Intel Parallel Studio XE

Master Parallel Extensions with .NET 4

Professional CUDA C Programming

Principles of Parallel Programming