
Online Library R High Performance Programming

This is likewise one of the factors by obtaining the soft documents of this **R High Performance Programming** by online. You might not require more become old to spend to go to the books establishment as well as search for them. In some cases, you likewise realize not discover the proclamation R High Performance Programming that you are looking for. It will unconditionally squander the time.

However below, as soon as you visit this web page, it will be consequently very easy to get as well as download lead R High Performance Programming

It will not take many times as we explain before. You can realize it even though accomplish something else at home and even in your workplace. consequently easy! So, are you question? Just exercise just what we have the funds for below as competently as review **R High Performance Programming** what you in the same way as to read!

KEY=HIGH - ASHTYN MCKENZIE

R HIGH PERFORMANCE PROGRAMMING

Packt Publishing Ltd This book is for programmers and developers who want to improve the performance of their R programs by making them run faster with large data sets or who are trying to solve a pesky performance problem.

INTEL XEON PHI COPROCESSOR HIGH PERFORMANCE PROGRAMMING

Newnes 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

ADVANCED R

CRC Press An Essential Reference for Intermediate and Advanced R Programmers Advanced R presents useful tools and techniques for attacking many types of R programming problems, helping you avoid mistakes and dead ends. With more than ten years of experience programming in R, the author illustrates the elegance, beauty, and flexibility at the heart of R. The book develops the necessary skills to produce quality code that can be used in a variety of circumstances. You will learn: The fundamentals of R, including standard data types and functions Functional programming as a useful framework for solving wide classes of problems The positives and negatives of metaprogramming How to write fast, memory-efficient code This book not only helps current R users become R programmers but also shows existing programmers what's special about R. Intermediate R programmers can dive deeper into R and learn new strategies for solving diverse problems while programmers from other languages can learn the details of R and understand why R works the way it does.

EFFICIENT R PROGRAMMING

A PRACTICAL GUIDE TO SMARTER PROGRAMMING

"O'Reilly Media, Inc." There are many excellent R resources for visualization, data science, and package development. Hundreds of scattered vignettes, web pages, and forums explain how to use R in particular domains. But little has been written on how to simply make R work effectively—until now. This hands-on book teaches novices and experienced R users how to write efficient R code. Drawing on

years of experience teaching R courses, authors Colin Gillespie and Robin Lovelace provide practical advice on a range of topics—from optimizing the set-up of RStudio to leveraging C++—that make this book a useful addition to any R user’s bookshelf. Academics, business users, and programmers from a wide range of backgrounds stand to benefit from the guidance in *Efficient R Programming*. Get advice for setting up an R programming environment Explore general programming concepts and R coding techniques Understand the ingredients of an efficient R workflow Learn how to efficiently read and write data in R Dive into data carpentry—the vital skill for cleaning raw data Optimize your code with profiling, standard tricks, and other methods Determine your hardware capabilities for handling R computation Maximize the benefits of collaborative R programming Accelerate your transition from R hacker to R programmer

MASTERING PARALLEL PROGRAMMING WITH R

Packt Publishing Ltd Master the robust features of R parallel programming to accelerate your data science computations About This Book Create R programs that exploit the computational capability of your cloud platforms and computers to the fullest Become an expert in writing the most efficient and highest performance parallel algorithms in R Get to grips with the concept of parallelism to accelerate your existing R programs Who This Book Is For This book is for R programmers who want to step beyond its inherent single-threaded and restricted memory limitations and learn how to implement highly accelerated and scalable algorithms that are a necessity for the performant processing of Big Data. No previous knowledge of parallelism is required. This book also provides for the more advanced technical programmer seeking to go beyond high level parallel frameworks. What You Will Learn Create and structure efficient load-balanced parallel computation in R, using R's built-in parallel package Deploy and utilize cloud-based parallel infrastructure from R, including launching a distributed computation on Hadoop running on Amazon Web Services (AWS) Get accustomed to parallel efficiency, and apply simple techniques to benchmark, measure speed and target improvement in your own code Develop complex parallel processing algorithms with the standard Message Passing Interface (MPI) using RMPI, pbdMPI, and SPRINT packages Build and extend a parallel R package (SPRINT) with your own MPI-based routines Implement accelerated numerical functions in R utilizing the vector processing capability of your Graphics Processing Unit (GPU) with OpenCL Understand parallel programming pitfalls, such as deadlock and numerical instability, and the approaches to handle and avoid them Build a task farm master-worker, spatial grid, and hybrid parallel R programs In Detail R is one of the most popular programming languages used in data science. Applying R to big data and complex analytic tasks requires the harnessing of scalable compute resources. Mastering Parallel Programming with R presents a comprehensive and practical treatise on how to build highly scalable and efficient algorithms in R. It will teach you a variety of parallelization techniques, from simple use of R's built-in parallel package versions of `lapply()`, to high-level AWS cloud-based Hadoop and Apache Spark frameworks. It will also teach you low level scalable parallel programming using RMPI and pbdMPI for message passing, applicable to clusters and supercomputers, and how to exploit thousand-fold simple processor GPUs through ROpenCL. By the end of the book, you will understand the factors that influence parallel efficiency, including assessing code performance and implementing load balancing; pitfalls to avoid, including deadlock and numerical instability issues; how to structure your code and data for the most appropriate type of parallelism for your problem domain; and how to extract the maximum performance from your R code running on a variety of computer systems. Style and approach This book leads you chapter by chapter from the easy to more complex forms of parallelism. The author's insights are presented through clear practical examples applied to a range of different problems, with comprehensive reference information for each of the R packages employed. The book can be read from start to finish, or by dipping in chapter by chapter, as each chapter describes a specific parallel approach and technology, so can be read as a standalone.

JULIA: HIGH PERFORMANCE PROGRAMMING

Packt Publishing Ltd Leverage the power of Julia to design and develop high performing programs About This Book Get to know the best techniques to create blazingly fast programs with Julia Stand out from the crowd by developing code that runs faster than your peers' code Complete an extensive data science project through the entire cycle from ETL to analytics and data visualization Who This Book Is For This learning path is for data scientists and for all those who work in technical and scientific computation projects. It will be great for Julia developers who are interested in high-performance technical computing. This learning path assumes that you already have some basic working knowledge of Julia's syntax and high-level dynamic languages such as MATLAB, R, Python, or Ruby. What You Will Learn Set up your Julia environment to achieve the highest productivity Solve your tasks in a high-level dynamic language and use types for your data only when needed Apply Julia to tackle problems concurrently and in a distributed environment Get a sense of the possibilities and limitations of Julia's performance Use Julia arrays to write high performance code Build a data science project through the entire cycle of ETL, analytics, and data visualization Display graphics and visualizations to carry out modeling and simulation in Julia Develop your own packages and contribute to the Julia Community In Detail In this learning path, you will learn to use an interesting and dynamic programming language—Julia! You will get a chance to tackle your numerical and data problems with Julia. You'll begin the journey by setting up a running Julia platform before exploring its various built-in types. We'll then move on to the various functions and constructs in Julia. We'll walk through the two important collection types—arrays and matrices in Julia. You will dive into how Julia uses type information to achieve its performance goals, and how to use multiple dispatch to help the compiler emit high performance machine code. You will see how Julia's design makes code fast, and you'll see its distributed computing capabilities. By the end of this learning path, you will see how data works using simple statistics and analytics, and you'll discover its high and dynamic performance—its real strength, which makes it particularly useful in highly intensive computing tasks. This learning path combines some of the best that Packt has to offer in one complete, curated package. It includes content from the following Packt products: *Getting Started with Julia* by Ivo Balvaert *Julia High Performance* by Avik Sengupta *Mastering Julia*

by Malcolm Sherrington Style and approach This hands-on manual will give you great explanations of the important concepts related to Julia programming.

PARALLEL AND HIGH PERFORMANCE COMPUTING

Simon and Schuster 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

LEARNING R PROGRAMMING

Packt Publishing Ltd Become an efficient data scientist with R About This Book Explore the R language from basic types and data structures to advanced topics Learn how to tackle programming problems and explore both functional and object-oriented programming techniques Learn how to address the core problems of programming in R and leverage the most popular packages for common tasks Who This Book Is For This is the perfect tutorial for anyone who is new to statistical programming and modeling. Anyone with basic programming and data processing skills can pick this book up to systematically learn the R programming language and crucial techniques. What You Will Learn Explore the basic functions in R and familiarize yourself with common data structures Work with data in R using basic functions of statistics, data mining, data visualization, root solving, and optimization Get acquainted with R's evaluation model with environments and meta-programming techniques with symbol, call, formula, and expression Get to grips with object-oriented programming in R: including the S3, S4, RC, and R6 systems Access relational databases such as SQLite and non-relational databases such as MongoDB and Redis Get to know high performance computing techniques such as parallel computing and Rcpp Use web scraping techniques to extract information Create RMarkdown, an interactive app with Shiny, DiagramR, interactive charts, ggvis, and more In Detail R is a high-level functional language and one of the must-know tools for data science and statistics. Powerful but complex, R can be challenging for beginners and those unfamiliar with its unique behaviors. Learning R Programming is the solution - an easy and practical way to learn R and develop a broad and consistent understanding of the language. Through hands-on examples you'll discover powerful R tools, and R best practices that will give you a deeper understanding of working with data. You'll get to grips with R's data structures and data processing techniques, as well as the most popular R packages to boost your productivity from the offset. Start with the basics of R, then dive deep into the programming techniques and paradigms to make your R code excel. Advance quickly to a deeper understanding of R's behavior as you learn common tasks including data analysis, databases, web scraping, high performance computing, and writing documents. By the end of the book, you'll be a confident R programmer adept at solving problems with the right techniques. Style and approach Developed to make learning easy and intuitive, this book comes packed with a wide variety of statistical and graphical techniques and a wealth of practical information for anyone looking to get started with this exciting and powerful language.

HIGH PERFORMANCE COMPUTING AND THE ART OF PARALLEL PROGRAMMING

AN INTRODUCTION FOR GEOGRAPHERS, SOCIAL SCIENTISTS AND ENGINEERS

Routledge This book provides a non-technical introduction to High Performance Computing applications together with advice about how beginners can start to write parallel programs. The authors show what HPC can offer geographers and social scientists and how it can be used in GIS. They provide examples of where it has already been used and suggestions for other areas of application in geography and the social sciences. Case studies drawn from geography explain the key principles and help to understand the logic and thought processes that lie behind the parallel programming.

JULIA HIGH PERFORMANCE

Packt Publishing Ltd Design and develop high performing programs with Julia About This Book Learn to code high reliability and high performance programs Stand out from the crowd by developing code that runs faster than your peers' codes This book is intended for developers who are interested in high performance technical programming. Who This Book Is For This book is for beginner and intermediate Julia programmers who are interested in high performance technical computing. You will have a basic familiarity with Julia syntax, and have written some small programs in the language. What You Will Learn Discover the secrets behind Julia's speed Get a sense of the possibilities and limitations of Julia's performance Analyze the performance of Julia programs Measure the time and memory taken by Julia programs Create fast machine code using Julia's type information Define and call functions without compromising Julia's performance Understand number types in Julia Use Julia arrays to write high performance code Get an overview of Julia's distributed computing capabilities In Detail Julia is a high performance, high-level dynamic language designed to address the requirements of high-level numerical and scientific computing. Julia brings solutions to the complexities faced by developers while developing elegant and high performing code. Julia High Performance will take you on a journey to understand the performance characteristics of your Julia programs, and enables you to utilize the promise of near C levels of performance in Julia. You will learn to analyze and measure the performance of Julia code, understand how to avoid bottlenecks, and design your program for the highest possible performance. In this book, you will also see how Julia uses type information to achieve its performance goals, and how to use multiple dispatch to help the compiler to emit high performance machine code. Numbers and their arrays are obviously the key structures in scientific computing – you will see how Julia's design makes them fast. The last chapter will give you a taste of Julia's distributed computing capabilities. Style and approach This is a hands-on manual that will give you good explanations about the important concepts related to Julia programming.

HIGH PERFORMANCE PYTHON

PRACTICAL PERFORMANT PROGRAMMING FOR HUMANS

"O'Reilly Media, Inc." Your Python code may run correctly, but you need it to run faster. Updated for Python 3, this expanded edition shows you how to locate performance bottlenecks and significantly speed up your code in high-data-volume programs. By exploring the fundamental theory behind design choices, High Performance Python helps you gain a deeper understanding of Python's implementation. How do you take advantage of multicore architectures or clusters? Or build a system that scales up and down without losing reliability? Experienced Python programmers will learn concrete solutions to many issues, along with war stories from companies that use high-performance Python for social media analytics, productionized machine learning, and more. Get a better grasp of NumPy, Cython, and profilers Learn how Python abstracts the underlying computer architecture Use profiling to find bottlenecks in CPU time and memory usage Write efficient programs by choosing appropriate data structures Speed up matrix and vector computations Use tools to compile Python down to machine code Manage multiple I/O and computational operations concurrently Convert multiprocessing code to run on local or remote clusters Deploy code faster using tools like Docker

MASTERING PARALLEL PROGRAMMING WITH R

Packt Publishing Ltd Master the robust features of R parallel programming to accelerate your data science computations About This Book Create R programs that exploit the computational capability of your cloud platforms and computers to the fullest Become an expert in writing the most efficient and highest performance parallel algorithms in R Get to grips with the concept of parallelism to accelerate your existing R programs Who This Book Is For This book is for R programmers who want to step beyond its inherent single-threaded and restricted memory limitations and learn how to implement highly accelerated and scalable algorithms that are a necessity for the performant processing of Big Data. No previous knowledge of parallelism is required. This book also provides for the more advanced technical programmer seeking to go beyond high level parallel frameworks. What You Will Learn Create and structure efficient load-balanced parallel computation in R, using R's built-in parallel package Deploy and utilize cloud-based parallel infrastructure from R, including launching a distributed computation on Hadoop running on Amazon Web Services (AWS) Get accustomed to parallel efficiency, and apply simple techniques to benchmark, measure speed and target improvement in your own code Develop complex parallel processing algorithms with the standard Message Passing Interface (MPI) using RMPI, pbdMPI, and SPRINT packages Build and extend a parallel R package (SPRINT) with your own MPI-based routines Implement accelerated numerical functions in R utilizing the vector processing capability of your Graphics Processing Unit (GPU) with OpenCL Understand parallel programming pitfalls, such as deadlock and numerical instability, and the approaches to handle and avoid them Build a task farm master-worker, spatial grid, and hybrid parallel R programs In Detail R is one of the most popular programming languages used in data science. Applying R to big data and complex analytic tasks requires the harnessing of scalable compute resources. Mastering Parallel Programming with R presents a comprehensive and practical treatise on how to build highly scalable and efficient algorithms in R. It will teach you a variety of parallelization techniques, from simple use of R's built-in parallel package versions of lapply(), to high-level AWS cloud-based Hadoop and Apache Spark frameworks. It will also teach you low level scalable parallel programming using RMPI and pbdMPI for message passing, applicable to clusters and supercomputers, and how to exploit thousand-fold simple processor GPUs through ROpenCL. By the end of the book, you will understand the factors that influence parallel efficiency, including assessing code performance and implementing load balancing; pitfalls to avoid, including

deadlock and numerical instability issues; how to structure your code and data for the most appropriate type of parallelism for your problem domain; and how to extract the maximum performance from your R code running on a variety of computer systems. Style and approach This book leads you chapter by chapter from the easy to more complex forms of parallelism. The author's insights are presented through clear practical examples applied to a range of different problems, with comprehensive reference information for each of the R packages employed. The book can be read from start to finish, or by dipping in chapter by chapter, as each chapter describes a specific parallel approach and technology, so can be read as a standalone.

HIGH PERFORMANCE PROGRAMMING FOR SOFT COMPUTING

CRC Press This book examines the present and future of soft computer techniques. It explains how to use the latest technological tools, such as multicore processors and graphics processing units, to implement highly efficient intelligent system methods using a general purpose computer.

INTEL XEON PHI PROCESSOR HIGH PERFORMANCE PROGRAMMING

KNIGHTS LANDING EDITION

Morgan Kaufmann This book is an all-in-one source of information for programming the Second-Generation Intel Xeon Phi product family also called Knights Landing. The authors provide detailed and timely Knights Landingspecific details, programming advice, and real-world examples. The authors distill their years of Xeon Phi programming experience coupled with insights from many expert customers — Intel Field Engineers, Application Engineers, and Technical Consulting Engineers — to create this authoritative book on the essentials of programming for Intel Xeon Phi products. Intel® Xeon Phi™ Processor High-Performance Programming is useful even before you ever program a system with an Intel Xeon Phi processor. To help 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 processors, or other high-performance microprocessors. Applying these techniques will generally increase your program performance on any system and prepare you better for Intel Xeon Phi processors. A practical guide to the essentials for programming Intel Xeon Phi processors Definitive coverage of the Knights Landing architecture Presents best practices for portable, high-performance computing and a familiar and proven threads and vectors programming model Includes real world code examples that highlight usages of the unique aspects of this new highly parallel and high-performance computational product Covers use of MCDRAM, AVX-512, Intel® Omni-Path fabric, many-cores (up to 72), and many threads (4 per core) Covers software developer tools, libraries and programming models Covers using Knights Landing as a processor and a coprocessor

HIGH-PERFORMANCE COMPUTING ON THE INTEL(R) XEON PHI

HOW TO FULLY EXPLOIT MIC ARCHITECTURES

Springer The aim of this book is to explain to high-performance computing (HPC) developers how to utilize the Intel(r) Xeon Phi series products efficiently. To that end, it introduces some computing grammar, programming technology and optimization methods for using many-integrated-core (MIC) platforms and also offers tips and tricks for actual use, based on the authors first-hand optimization experience. The material is organized in three sections. The first section, Basics of MIC, introduces the fundamentals of MIC architecture and programming, including the specific Intel MIC programming environment. Next, the section on Performance Optimization explains general MIC optimization techniques, which are then illustrated step-by-step using the classical parallel programming example of matrix multiplication. Finally, Project development presents a set of practical and experience-driven methods for using parallel computing in application projects, including how to determine if a serial or parallel CPU program is suitable for MIC and how to transplant a program onto MIC. This book appeals to two main audiences: First, software developers for HPC applications it will enable them to fully exploit the MIC architecture and thus achieve the extreme performance usually required in biological genetics, medical imaging, aerospace, meteorology and other areas of HPC. Second, students and researchers engaged in parallel and high-performance computing it will guide them on how to push the limits of system performance for HPC applications. "

SOFTWARE OPTIMIZATION FOR HIGH-PERFORMANCE COMPUTING

Prentice Hall Professional The hands-on guide to high-performance coding and algorithm optimization. This hands-on guide to software optimization introduces state-of-the-art solutions for every key aspect of software performance - both code-based and algorithm-based. Two leading HP software performance experts offer comparative optimization strategies for RISC and for the new Explicitly Parallel Instruction Computing (EPIC) design used in Intel IA-64 processors. Using many practical examples, they offer specific techniques for: Predicting and measuring performance - and identifying your best optimization opportunities Storage optimization: cache, system memory, virtual memory, and I/O Parallel processing: distributed-memory and shared-memory (SMP and ccNUMA) Compilers and loop optimization Enhancing parallelism: compiler directives, threads, and message passing Mathematical libraries and algorithms Whether you're a developer, ISV, or technical researcher, if you need to optimize high-performance software on today's leading processors, one book delivers the advanced techniques and code examples you need: Software Optimization for High Performance Computing.

R FOR PROGRAMMERS

MASTERING THE TOOLS

CRC Press Unlike other books about R, written from the perspective of statistics, R for Programmers: Mastering the Tools is written from the perspective of programmers, providing a channel for programmers with expertise in other programming languages to quickly understand R. The contents are divided into four sections: The first section consists of the basic

HASKELL HIGH PERFORMANCE PROGRAMMING

Packt Publishing Ltd Boost the performance of your Haskell applications using optimization, concurrency, and parallel programming About This Book Explore the benefits of lazy evaluation, compiler features, and tools and libraries designed for high performance Write fast programs at extremely high levels of abstraction Work through practical examples that will help you address the challenges of writing efficient code Who This Book Is For To get the most out of this book, you need to have a working knowledge of reading and writing basic Haskell. No knowledge of performance, optimization, or concurrency is required. What You Will Learn Program idiomatic Haskell that's also surprisingly efficient Improve performance of your code with data parallelism, inlining, and strictness annotations Profile your programs to identify space leaks and missed opportunities for optimization Find out how to choose the most efficient data and control structures Optimize the Glasgow Haskell Compiler and runtime system for specific programs See how to smoothly drop to lower abstractions wherever necessary Execute programming for the GPU with Accelerate Implement programming to easily scale to the cloud with Cloud Haskell In Detail Haskell, with its power to optimize the code and its high performance, is a natural candidate for high performance programming. It is especially well suited to stacking abstractions high with a relatively low performance cost. This book addresses the challenges of writing efficient code with lazy evaluation and techniques often used to optimize the performance of Haskell programs. We open with an in-depth look at the evaluation of Haskell expressions and discuss optimization and benchmarking. You will learn to use parallelism and we'll explore the concept of streaming. We'll demonstrate the benefits of running multithreaded and concurrent applications. Next we'll guide you through various profiling tools that will help you identify performance issues in your program. We'll end our journey by looking at GPGPU, Cloud and Functional Reactive Programming in Haskell. At the very end there is a catalogue of robust library recommendations with code samples. By the end of the book, you will be able to boost the performance of any app and prepare it to stand up to real-world punishment. Style and approach This easy-to-follow guide teaches new practices and techniques to optimize your code, and then moves towards more advanced ways to effectively write efficient Haskell code. Small and simple practical examples will help you test the concepts yourself, and you will be able to easily adapt them for any application.

LANGUAGES AND COMPILERS FOR HIGH PERFORMANCE COMPUTING

17TH INTERNATIONAL WORKSHOP, LCPC 2004, WEST LAFAYETTE, IN, USA, SEPTEMBER 22-24, 2004, REVISED SELECTED PAPERS

Springer Science & Business Media This book constitutes the thoroughly refereed post-proceedings of the 17th International Workshop on Languages and Compilers for High Performance Computing, LCPC 2004, held in West Lafayette, IN, USA in September 2004. The 33 revised full papers presented were carefully selected during two rounds of reviewing and improvement. The papers are organized in topical sections on compiler infrastructures; predicting and reducing memory access; locality, tiling, and partitioning; tools and techniques for parallelism and locality; Java for high-performance computing; high-level languages and optimizations; large-scale data sharing; performance studies; program analysis; and exploiting architectural features.

SEAMLESS R AND C++ INTEGRATION WITH RCPP

Springer Science & Business Media Rcpp is the glue that binds the power and versatility of R with the speed and efficiency of C++. With Rcpp, the transfer of data between R and C++ is nearly seamless, and high-performance statistical computing is finally accessible to most R users. Rcpp should be part of every statistician's toolbox. -- Michael Braun, MIT Sloan School of Management "Seamless R and C++ integration with Rcpp" is simply a wonderful book. For anyone who uses C/C++ and R, it is an indispensable resource. The writing is outstanding. A huge bonus is the section on applications. This section covers the matrix packages Armadillo and Eigen and the GNU Scientific Library as well as Rinside which enables you to use R inside C++. These applications are what most of us need to know to really do scientific programming with R and C++. I love this book. -- Robert McCulloch, University of Chicago Booth School of Business Rcpp is now considered an essential package for anybody doing serious computational research using R. Dirk's book is an excellent companion and takes the reader from a gentle introduction to more advanced applications via numerous examples and efficiency enhancing gems. The book is packed with all you might have ever wanted to know about Rcpp, its cousins (RcppArmadillo, RcppEigen .etc.), modules, package development and sugar. Overall, this book is a must-have on your shelf. -- Sanjog Misra, UCLA Anderson School of Management The Rcpp package represents a major leap forward for scientific computations with R. With very few lines of C++ code, one has R's data structures readily at hand for further computations in C++. Hence, high-level numerical programming can be made in C++ almost as easily as in R, but often with a substantial speed

gain. Dirk is a crucial person in these developments, and his book takes the reader from the first fragile steps on to using the full Rcpp machinery. A very recommended book! -- Søren Højsgaard, Department of Mathematical Sciences, Aalborg University, Denmark "Seamless R and C ++ Integration with Rcpp" provides the first comprehensive introduction to Rcpp. Rcpp has become the most widely-used language extension for R, and is deployed by over one-hundred different CRAN and BioConductor packages. Rcpp permits users to pass scalars, vectors, matrices, list or entire R objects back and forth between R and C++ with ease. This brings the depth of the R analysis framework together with the power, speed, and efficiency of C++. Dirk Eddelbuettel has been a contributor to CRAN for over a decade and maintains around twenty packages. He is the Debian/Ubuntu maintainer for R and other quantitative software, edits the CRAN Task Views for Finance and High-Performance Computing, is a co-founder of the annual R/Finance conference, and an editor of the Journal of Statistical Software. He holds a Ph.D. in Mathematical Economics from EHESS (Paris), and works in Chicago as a Senior Quantitative Analyst.

PARALLEL R

DATA ANALYSIS IN THE DISTRIBUTED WORLD

"O'Reilly Media, Inc." It's tough to argue with R as a high-quality, cross-platform, open source statistical software product—unless you're in the business of crunching Big Data. This concise book introduces you to several strategies for using R to analyze large datasets, including three chapters on using R and Hadoop together. You'll learn the basics of Snow, Multicore, Parallel, Segue, RHIPE, and Hadoop Streaming, including how to find them, how to use them, when they work well, and when they don't. With these packages, you can overcome R's single-threaded nature by spreading work across multiple CPUs, or offloading work to multiple machines to address R's memory barrier. Snow: works well in a traditional cluster environment Multicore: popular for multiprocessor and multicore computers Parallel: part of the upcoming R 2.14.0 release R+Hadoop: provides low-level access to a popular form of cluster computing RHIPE: uses Hadoop's power with R's language and interactive shell Segue: lets you use Elastic MapReduce as a backend for lapply-style operations

THE ART OF R PROGRAMMING

A TOUR OF STATISTICAL SOFTWARE DESIGN

No Starch Press R is the world's most popular language for developing statistical software: Archaeologists use it to track the spread of ancient civilizations, drug companies use it to discover which medications are safe and effective, and actuaries use it to assess financial risks and keep economies running smoothly. The Art of R Programming takes you on a guided tour of software development with R, from basic types and data structures to advanced topics like closures, recursion, and anonymous functions. No statistical knowledge is required, and your programming skills can range from hobbyist to pro. Along the way, you'll learn about functional and object-oriented programming, running mathematical simulations, and rearranging complex data into simpler, more useful formats. You'll also learn to:

- Create artful graphs to visualize complex data sets and functions
- Write more efficient code using parallel R and vectorization
- Interface R with C/C++ and Python for increased speed or functionality
- Find new R packages for text analysis, image manipulation, and more
- Squash annoying bugs with advanced debugging techniques

Whether you're designing aircraft, forecasting the weather, or you just need to tame your data, The Art of R Programming is your guide to harnessing the power of statistical computing.

HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS

FOUNDATION FOR AMERICA'S INFORMATION FUTURE : A REPORT

R PROGRAMMING FOR DATA SCIENCE

Lulu.com 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

HIGH PERFORMANCE BROWSER NETWORKING

WHAT EVERY WEB DEVELOPER SHOULD KNOW ABOUT NETWORKING AND WEB PERFORMANCE

"O'Reilly Media, Inc." How prepared are you to build fast and efficient web applications? This eloquent book provides what every web developer should know about the network, from fundamental

limitations that affect performance to major innovations for building even more powerful browser applications—including HTTP 2.0 and XHR improvements, Server-Sent Events (SSE), WebSocket, and WebRTC. Author Ilya Grigorik, a web performance engineer at Google, demonstrates performance optimization best practices for TCP, UDP, and TLS protocols, and explains unique wireless and mobile network optimization requirements. You'll then dive into performance characteristics of technologies such as HTTP 2.0, client-side network scripting with XHR, real-time streaming with SSE and WebSocket, and P2P communication with WebRTC. Deliver superlative TCP, UDP, and TLS performance Speed up network performance over 3G/4G mobile networks Develop fast and energy-efficient mobile applications Address bottlenecks in HTTP 1.x and other browser protocols Plan for and deliver the best HTTP 2.0 performance Enable efficient real-time streaming in the browser Create efficient peer-to-peer videoconferencing and low-latency applications with real-time WebRTC transports

R PROGRAMMING BY EXAMPLE

PRACTICAL, HANDS-ON PROJECTS TO HELP YOU GET STARTED WITH R

Packt Publishing Ltd This step-by-step guide demonstrates how to build simple-to-advanced applications through examples in R using modern tools. About This Book Get a firm hold on the fundamentals of R through practical hands-on examples Get started with good R programming fundamentals for data science Exploit the different libraries of R to build interesting applications in R Who This Book Is For This book is for aspiring data science professionals or statisticians who would like to learn about the R programming language in a practical manner. Basic programming knowledge is assumed. What You Will Learn Discover techniques to leverage R's features, and work with packages Perform a descriptive analysis and work with statistical models using R Work efficiently with objects without using loops Create diverse visualizations to gain better understanding of the data Understand ways to produce good visualizations and create reports for the results Read and write data from relational databases and REST APIs, both packaged and unpackaged Improve performance by writing better code, delegating that code to a more efficient programming language, or making it parallel In Detail R is a high-level statistical language and is widely used among statisticians and data miners to develop analytical applications. Often, data analysis people with great analytical skills lack solid programming knowledge and are unfamiliar with the correct ways to use R. Based on the version 3.4, this book will help you develop strong fundamentals when working with R by taking you through a series of full representative examples, giving you a holistic view of R. We begin with the basic installation and configuration of the R environment. As you progress through the exercises, you'll become thoroughly acquainted with R's features and its packages. With this book, you will learn about the basic concepts of R programming, work efficiently with graphs, create publication-ready and interactive 3D graphs, and gain a better understanding of the data at hand. The detailed step-by-step instructions will enable you to get a clean set of data, produce good visualizations, and create reports for the results. It also teaches you various methods to perform code profiling and performance enhancement with good programming practices, delegation, and parallelization. By the end of this book, you will know how to efficiently work with data, create quality visualizations and reports, and develop code that is modular, expressive, and maintainable. Style and Approach This is an easy-to-understand guide filled with real-world examples, giving you a holistic view of R and practical, hands-on experience.

CONQUERING BIG DATA WITH HIGH PERFORMANCE COMPUTING

Springer This book provides an overview of the resources and research projects that are bringing Big Data and High Performance Computing (HPC) on converging tracks. It demystifies Big Data and HPC for the reader by covering the primary resources, middleware, applications, and tools that enable the usage of HPC platforms for Big Data management and processing. Through interesting use-cases from traditional and non-traditional HPC domains, the book highlights the most critical challenges related to Big Data processing and management, and shows ways to mitigate them using HPC resources. Unlike most books on Big Data, it covers a variety of alternatives to Hadoop, and explains the differences between HPC platforms and Hadoop. Written by professionals and researchers in a range of departments and fields, this book is designed for anyone studying Big Data and its future directions. Those studying HPC will also find the content valuable.

HIGH PERFORMANCE COMPUTING

ISC HIGH PERFORMANCE 2016 INTERNATIONAL WORKSHOPS, EXACOMM, E-MUCOCOS, HPC-IODC, IXPUG, IWOPH, P³MA, VHPC, WOPSSS, FRANKFURT, GERMANY, JUNE 19-23, 2016, REVISED SELECTED PAPERS

Springer This book constitutes revised selected papers from 7 workshops that were held in conjunction with the ISC High Performance 2016 conference in Frankfurt, Germany, in June 2016. The 45 papers presented in this volume were carefully reviewed and selected for inclusion in this book. They stem from the following workshops: Workshop on Exascale Multi/Many Core Computing Systems, E-MuCoCoS; Second International Workshop on Communication Architectures at Extreme Scale, ExaComm; HPC I/O in the Data Center Workshop, HPC-IODC; International Workshop on OpenPOWER for HPC, IWOPH; Workshop on the Application Performance on Intel Xeon Phi - Being Prepared for KNL and Beyond, IXPUG; Workshop on Performance and Scalability of Storage Systems, WOPSSS; and International Workshop on Performance Portable Programming Models for Accelerators, P3MA.

TOOLS FOR HIGH PERFORMANCE COMPUTING 2016

PROCEEDINGS OF THE 10TH INTERNATIONAL WORKSHOP ON PARALLEL TOOLS FOR HIGH PERFORMANCE COMPUTING, OCTOBER 2016, STUTTGART, GERMANY

Springer This book presents the proceedings of the 10th International Parallel Tools Workshop, held October 4-5, 2016 in Stuttgart, Germany - a forum to discuss the latest advances in parallel tools. High-performance computing plays an increasingly important role for numerical simulation and modelling in academic and industrial research. At the same time, using large-scale parallel systems efficiently is becoming more difficult. A number of tools addressing parallel program development and analysis have emerged from the high-performance computing community over the last decade, and what may have started as collection of small helper script has now matured to production-grade frameworks. Powerful user interfaces and an extensive body of documentation allow easy usage by non-specialists.

INTRODUCTION TO HIGH PERFORMANCE SCIENTIFIC COMPUTING

Lulu.com

LEARNING .NET HIGH-PERFORMANCE PROGRAMMING

Packt Publishing Ltd This book will help you understand what "programming for performance" means, and use effective coding patterns and techniques to optimize your .NET applications. You will begin by understanding what "high performance coding" means, and the different performance concerns. You will see how CLR works and get an understanding of concepts such as memory management, garbage collection, and thread life cycles. You will proceed to learn about the theoretical and practical concepts of PLINQ programming. You will also see what Big Data is, and how to architect a Big Data solution to manipulate large datasets. Finally, you will learn how to launch and analyze a profile session and execute tests against a code block or application for performance analysis. By the end of this book, you will have a complete understanding of efficient programming using high-performance techniques, and will able to write highly optimized applications.

TESTING R CODE

CRC Press Learn how to write R code with fewer bugs. The problem with programming is that you are always one typo away from writing something silly. Likewise with data analysis, a small mistake in your model can lead to a big mistake in your results. Combining the two disciplines means that it is all too easy for a missed minus sign to generate a false prediction that you don't spot until it's too late. Testing is the only way to be sure that your code, and your results, are correct. Testing R Code teaches you how to perform development-time testing using the testthat package, allowing you to ensure that your code works as intended. The book also teaches run-time testing using the assertive package; enabling your users to correctly run your code. After beginning with an introduction to testing in R, the book explores more advanced cases such as integrating tests into R packages; testing code that accesses databases; testing C++ code with Rcpp; and testing graphics. Each topic is explained with real-world examples, and has accompanying exercises for readers to practise their skills — only a small amount of experience with R is needed to get started!

HIGH PERFORMANCE COMPUTING SYSTEMS. PERFORMANCE MODELING, BENCHMARKING, AND SIMULATION

8TH INTERNATIONAL WORKSHOP, PMBS 2017, DENVER, CO, USA, NOVEMBER 13, 2017, PROCEEDINGS

Springer This book constitutes the refereed proceedings papers from the 8th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems, PMBS 2017, held in Denver, Colorado, USA, in November 2017. The 10 full papers and 3 short papers included in this volume were carefully reviewed and selected from 36 submissions. They were organized in topical sections named: performance evaluation and analysis; performance modeling and simulation; and short papers.

R AND DATA MINING

EXAMPLES AND CASE STUDIES

Academic Press R and Data Mining introduces researchers, post-graduate students, and analysts to data mining using R, a free software environment for statistical computing and graphics. The book provides practical methods for using R in applications from academia to industry to extract knowledge from vast amounts of data. Readers will find this book a valuable guide to the use of R in tasks such as classification and prediction, clustering, outlier detection, association rules, sequence analysis, text mining, social network analysis, sentiment analysis, and more. Data mining techniques are growing in

popularity in a broad range of areas, from banking to insurance, retail, telecom, medicine, research, and government. This book focuses on the modeling phase of the data mining process, also addressing data exploration and model evaluation. With three in-depth case studies, a quick reference guide, bibliography, and links to a wealth of online resources, R and Data Mining is a valuable, practical guide to a powerful method of analysis. Presents an introduction into using R for data mining applications, covering most popular data mining techniques Provides code examples and data so that readers can easily learn the techniques Features case studies in real-world applications to help readers apply the techniques in their work

JULIA 1.0 PROGRAMMING

DYNAMIC AND HIGH-PERFORMANCE PROGRAMMING TO BUILD FAST SCIENTIFIC APPLICATIONS, 2ND EDITION

Packt Publishing Ltd Enter the exciting world of Julia, a high-performance language for technical computing Key Features Leverage Julia's high speed and efficiency for your applications Work with Julia in a multi-core, distributed, and networked environment Apply Julia to tackle problems concurrently and in a distributed environment Book Description The release of Julia 1.0 is now ready to change the technical world by combining the high productivity and ease of use of Python and R with the lightning-fast speed of C++. Julia 1.0 programming gives you a head start in tackling your numerical and data problems. You will begin by learning how to set up a running Julia platform, before exploring its various built-in types. With the help of practical examples, this book walks you through two important collection types: arrays and matrices. In addition to this, you will be taken through how type conversions and promotions work. In the course of the book, you will be introduced to the homo-iconicity and metaprogramming concepts in Julia. You will understand how Julia provides different ways to interact with an operating system, as well as other languages, and then you'll discover what macros are. Once you have grasped the basics, you'll study what makes Julia suitable for numerical and scientific computing, and learn about the features provided by Julia. By the end of this book, you will also have learned how to run external programs. This book covers all you need to know about Julia in order to leverage its high speed and efficiency for your applications. What you will learn Set up your Julia environment to achieve high productivity Create your own types to extend the built-in type system Visualize your data in Julia with plotting packages Explore the use of built-in macros for testing and debugging, among other uses Apply Julia to tackle problems concurrently Integrate Julia with other languages such as C, Python, and MATLAB Who this book is for Julia 1.0 Programming is for you if you are a statistician or data scientist who wants a crash course in the Julia programming language while building big data applications. A basic knowledge of mathematics is needed to understand the various methods that are used or created during the course of the book to exploit the capabilities that Julia is designed with.

HIGH PERFORMANCE PARALLELISM PEARLS VOLUME TWO

MULTICORE AND MANY-CORE PROGRAMMING APPROACHES

Morgan Kaufmann High Performance Parallelism Pearls Volume 2 offers another set of examples that demonstrate how to leverage parallelism. Similar to Volume 1, the techniques included here explain how to use processors and coprocessors with the same programming - illustrating the most effective ways to combine Xeon Phi coprocessors with Xeon and other multicore processors. The book includes examples of successful programming efforts, drawn from across industries and domains such as biomed, genetics, finance, manufacturing, imaging, and more. Each chapter in this edited work includes detailed explanations of the programming techniques used, while showing high performance results on both Intel Xeon Phi coprocessors and multicore processors. Learn from dozens of new examples and case studies illustrating "success stories" demonstrating not just the features of Xeon-powered systems, but also how to leverage parallelism across these heterogeneous systems. Promotes write-once, run-anywhere coding, showing how to code for high performance on multicore processors and Xeon Phi Examples from multiple vertical domains illustrating real-world use of Xeon Phi coprocessors Source code available for download to facilitate further exploration

HIGH PERFORMANCE COMPUTING: TECHNOLOGY, METHODS AND APPLICATIONS

Elsevier High Performance Computing is an integrated computing environment for solving large-scale computational demanding problems in science, engineering and business. Newly emerging areas of HPC applications include medical sciences, transportation, financial operations and advanced human-computer interface such as virtual reality. High performance computing includes computer hardware, software, algorithms, programming tools and environments, plus visualization. The book addresses several of these key components of high performance technology and contains descriptions of the state-of-the-art computer architectures, programming and software tools and innovative applications of parallel computers. In addition, the book includes papers on heterogeneous network-based computing systems and scalability of parallel systems. The reader will find information and data relative to the two main thrusts of high performance computing: the absolute computational performance and that of providing the most cost effective and affordable computing for science, industry and business. The book is recommended for technical as well as management oriented individuals.

BIG DATA ANALYTICS FOR CYBER-PHYSICAL SYSTEMS

MACHINE LEARNING FOR THE INTERNET OF THINGS

Elsevier Big Data Analytics in Cyber-Physical Systems: Machine Learning for the Internet of Things examines sensor signal processing, IoT gateways, optimization and decision-making, intelligent mobility, and implementation of machine learning algorithms in embedded systems. This book focuses on the interaction between IoT technology and the mathematical tools used to evaluate the extracted data of those systems. Each chapter provides the reader with a broad list of data analytics and machine learning methods for multiple IoT applications. Additionally, this volume addresses the educational transfer needed to incorporate these technologies into our society by examining new platforms for IoT in schools, new courses and concepts for universities and adult education on IoT and data science. . Bridges the gap between IoT, CPS, and mathematical modelling. Features numerous use cases that discuss how concepts are applied in different domains and applications. Provides "best practices", "winning stories" and "real-world examples" to complement innovation. Includes highlights of mathematical foundations of signal processing and machine learning in CPS and IoT.

JULIA 1.0 PROGRAMMING COMPLETE REFERENCE GUIDE

DISCOVER JULIA, A HIGH-PERFORMANCE LANGUAGE FOR TECHNICAL COMPUTING

Packt Publishing Ltd Learn dynamic programming with Julia to build apps for data analysis, visualization, machine learning, and the web Key FeaturesLeverage Julia's high speed and efficiency to build fast, efficient applicationsPerform supervised and unsupervised machine learning and time series analysisTackle problems concurrently and in a distributed environmentBook Description Julia offers the high productivity and ease of use of Python and R with the lightning-fast speed of C++. There's never been a better time to learn this language, thanks to its large-scale adoption across a wide range of domains, including fintech, biotech and artificial intelligence (AI). You will begin by learning how to set up a running Julia platform, before exploring its various built-in types. This Learning Path walks you through two important collection types: arrays and matrices. You'll be taken through how type conversions and promotions work, and in further chapters you'll study how Julia interacts with operating systems and other languages. You'll also learn about the use of macros, what makes Julia suitable for numerical and scientific computing, and how to run external programs. Once you have grasped the basics, this Learning Path goes on to how to analyze the Iris dataset using DataFrames. While building a web scraper and a web app, you'll explore the use of functions, methods, and multiple dispatches. In the final chapters, you'll delve into machine learning, where you'll build a book recommender system. By the end of this Learning Path, you'll be well versed with Julia and have the skills you need to leverage its high speed and efficiency for your applications. This Learning Path includes content from the following Packt products: Julia 1.0 Programming - Second Edition by Ivo BalbaertJulia Programming Projects by Adrian SalceanuWhat you will learnCreate your own types to extend the built-in type systemVisualize your data in Julia with plotting packagesExplore the use of built-in macros for testing and debuggingIntegrate Julia with other languages such as C, Python, and MATLABAnalyze and manipulate datasets using Julia and DataFramesDevelop and run a web app using Julia and the HTTP packageBuild a recommendation system using supervised machine learningWho this book is for If you are a statistician or data scientist who wants a quick course in the Julia programming language while building big data applications, this Learning Path is for you. Basic knowledge of mathematics and programming is a must.

INTRODUCTION TO HIGH PERFORMANCE COMPUTING FOR SCIENTISTS AND ENGINEERS

CRC Press Written by high performance computing (HPC) experts, Introduction to High Performance Computing for Scientists and Engineers provides a solid introduction to current mainstream computer architecture, dominant parallel programming models, and useful optimization strategies for scientific HPC. From working in a scientific computing center, the author