

Parallel Programming For Multicore And Cluster Systems

Thank you for downloading Parallel Programming For Multicore And Cluster Systems. Maybe you have knowledge that, people have search hundreds times for their favorite novels like this Parallel Programming For Multicore And Cluster Systems, but end up in malicious downloads.

Rather than reading a good book with a cup of tea in the afternoon, instead they cope with some malicious virus inside their laptop.

Parallel Programming For Multicore And Cluster Systems is available in our digital library an online access to it is set as public so you can download it instantly.

Our book servers spans in multiple locations, allowing you to get the most less latency time to download any of our books like this one.

Kindly say, the Parallel Programming For Multicore And Cluster Systems is universally compatible with any devices to read

High Performance Parallel Runtimes Michael Klemm 2021-02-08 This book focuses on the theoretical and practical aspects of parallel programming systems for today's high performance multi-core processors and discusses the efficient implementation of key algorithms needed to implement parallel programming models. Such implementations need to take into account the specific architectural aspects of the underlying computer architecture and the features offered by the execution environment. This book briefly reviews key concepts of modern computer architecture, focusing particularly on the performance of parallel codes as well as the relevant concepts in parallel programming models. The book then turns towards the fundamental algorithms used to implement the parallel programming models and discusses how they interact with modern processors. While the book will focus on the general mechanisms, we will mostly use the Intel processor architecture to exemplify the implementation concepts discussed but will present other processor architectures where appropriate. All algorithms and concepts are discussed in an easy to understand way with many illustrative examples, figures, and source code fragments. The target audience of the book is students in Computer Science who are studying compiler construction, parallel programming, or programming systems. Software developers who have an interest in the core algorithms used to implement a parallel runtime system, or who need to educate themselves for projects that require the algorithms and concepts discussed in this book will also benefit from reading it.

Embedded Computing Systems: Applications, Optimization, and Advanced Design Khalgui, Mohamed 2013-04-30 Embedded computing systems play an important and complex role in the functionality of electronic devices. With our daily routines becoming more reliant on electronics for personal and professional use, the understanding of these computing systems is crucial. Embedded Computing Systems: Applications, Optimization, and Advanced Design brings together theoretical and technical concepts of intelligent embedded control systems and their use in hardware and software architectures. By highlighting formal modeling, execution models, and optimal implementations, this reference source is essential for experts, researchers, and technical supporters in the industry and academia.

Parallel R Q. Ethan McCallum 2011-10-21 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, RHIFE, 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 RHIFE: uses Hadoop's power with R's language and interactive shell Segue: lets you use Elastic MapReduce as a backend for lapply-style operations

Parallel Programming Thomas Rauber 2010-03-10 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. The main goal of the book is to present parallel programming techniques that can be used in many situations for many 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 presented material has been used for courses in parallel programming at different universities for many years.

Programmverifikation Krzysztof R. Apt 2013-03-07

Programming Multicore and Many-core Computing Systems Sabri Pllana 2017-01-23 Programming multi-core and many-core computing systems Sabri Pllana, Linnaeus University, Sweden Fatos Xhafa, Technical University of Catalonia, Spain Provides state-of-the-art methods for programming multi-core and many-core systems The book comprises a selection of twenty two chapters covering: fundamental techniques and algorithms; programming approaches; methodologies and frameworks; scheduling and management; testing and evaluation methodologies; and case studies for programming multi-core and many-core systems. Program development for multi-core processors, especially for heterogeneous multi-core processors, is significantly more complex than for single-core processors. However, programmers have been traditionally trained for the development of sequential programs, and only a small percentage of them have experience with parallel programming. In the past, only a relatively small group of programmers interested in High Performance Computing (HPC) was concerned with the parallel programming issues, but the situation has changed dramatically with the appearance of multi-core processors on commonly used computing systems. It is expected that with the pervasiveness of multi-core processors, parallel programming will become mainstream. The pervasiveness of multi-core processors affects a large spectrum of systems, from embedded and general-purpose, to high-end computing systems. This book assists programmers in mastering the efficient programming of multi-core systems, which is of paramount importance for the software-intensive industry towards a more effective product-development cycle. Key features: Lessons, challenges, and roadmaps ahead. Contains real world examples and case studies. Helps programmers in mastering the efficient programming of multi-core and many-core systems. The book serves as a reference for a larger audience of practitioners, young researchers and graduate level students. A basic level of programming knowledge is required to use this book.

High Performance Computing and Applications Wu Zhang 2010-03-10 The Second International Conference on High-Performance Computing and Applications (HPCA 2009) was a follow-up event of the successful HPCA 2004. It was held in Shanghai, a beautiful, active, and modern city in China, August 10–12, 2009. It served as a forum to present current work by researchers and software developers from around the world as well as to highlight activities in the high-performance computing area. It aimed to bring together research scientists, application pioneers, and software developers to discuss problems and solutions and to identify new issues in this area. This conference emphasized the development and study of novel approaches for high-performance computing, the design and analysis of high-performance numerical algorithms, and their scientific, engineering, and industrial applications. It offered the conference participants a great opportunity to exchange the latest research results, heighten international collaboration, and discuss future research ideas in HPCA. In addition to 24 invited presentations, the conference received over 300 contributed submissions from over ten countries and regions worldwide, about 70 of which were accepted for presentation at HPCA 2009. The conference proceedings contain some of the invited presentations and contributed submissions, and cover such research areas of interest as numerical algorithms and solutions, high-performance and grid computing, novel approaches to high-performance computing, massive data storage and processing, hardware acceleration, and their wide applications.

Big Data in Engineering Applications Sanjiban Sekhar Roy 2018-05-02 This book presents the current trends, technologies, and challenges in Big Data in the diversified field of engineering and sciences. It covers the applications of Big Data ranging from conventional fields of mechanical engineering, civil engineering to electronics, electrical, and computer science to areas in pharmaceutical and biological sciences. This book consists of contributions from various authors from all sectors of academia and industries, demonstrating the imperative application of Big Data for the decision-making process in sectors where the volume, variety, and velocity of information keep increasing. The book is a useful reference for graduate students, researchers and scientists interested in exploring the potential of Big Data in the application of engineering areas.

Linux in a nutshell Ellen Siever 2005

Parallel Computing Barbara Chapman 2010-01-01 Parallel computing technologies have brought dramatic changes to mainstream computing; the majority of today's PC's, laptops and even notebooks incorporate multiprocessor chips with up to four processors. Standard components are increasingly combined with GPU's (Graphics Processing Unit), originally designed for high-speed graphics processing, and FPGA's (Free Programmable Gate Array) to build parallel computers with a wide spectrum of high-speed processing functions. The scale of this powerful hardware is limited only by factors such as energy consumption and thermal control However, in addition to hardware factors, the practical use of petascale and exascale machines is often hampered by the difficulty of

developing software which will run effectively and efficiently on such architecture This book includes selected and refereed papers, presented at the 2009 international Parallel Computing conference (ParCo2009), which set out to address these problems. It provides a snapshot of the state-of-the-art of parallel computing technologies in hardware, application and software development Areas covered include: numerical algorithms, grid and cloud computing, programming - including GPU and cell programming. The book also includes papers presented at the six mini-symposia held at the conference

Parallel Computing is Everywhere S. Bassini 2018-03-07 The most powerful computers work by harnessing the combined computational power of millions of processors, and exploiting the full potential of such large-scale systems is something which becomes more difficult with each succeeding generation of parallel computers. Alternative architectures and computer paradigms are increasingly being investigated in an attempt to address these difficulties. Added to this, the pervasive presence of heterogeneous and parallel devices in consumer products such as mobile phones, tablets, personal computers and servers also demands efficient programming environments and applications aimed at small-scale parallel systems as opposed to large-scale supercomputers. This book presents a selection of papers presented at the conference: Parallel Computing (ParCo2017), held in Bologna, Italy, on 12 to 15 September 2017. The conference included contributions about alternative approaches to achieving High Performance Computing (HPC) to potentially surpass exa- and zetascale performances, as well as papers on the application of quantum computers and FPGA processors. These developments are aimed at making available systems better capable of solving intensive computational scientific/engineering problems such as climate models, security applications and classic NP-problems, some of which cannot currently be managed by even the most powerful supercomputers available. New areas of application, such as robotics, AI and learning systems, data science, the Internet of Things (IoT), and in-car systems and autonomous vehicles were also covered. As always, ParCo2017 attracted a large number of notable contributions covering present and future developments in parallel computing, and the book will be of interest to all those working in the field.

Parallel Programming for Modern High Performance Computing Systems Pawel Czarnul 2018-03-05 In view of the growing presence and popularity of multicore and manycore processors, accelerators, and coprocessors, as well as clusters using such computing devices, the development of efficient parallel applications has become a key challenge to be able to exploit the performance of such systems. This book covers the scope of parallel programming for modern high performance computing systems. It first discusses selected and popular state-of-the-art computing devices and systems available today, These include multicore CPUs, manycore (co)processors, such as Intel Xeon Phi, accelerators, such as GPUs, and clusters, as well as programming models supported on these platforms. It next introduces parallelization through important programming paradigms, such as master-slave, geometric Single Program Multiple Data (SPMD) and divide-and-conquer. The practical and useful elements of the most popular and important APIs for programming parallel HPC systems are discussed, including MPI, OpenMP, Pthreads, CUDA, OpenCL, and OpenACC. It also demonstrates, through selected code listings, how selected APIs can be used to implement important programming paradigms. Furthermore, it shows how the codes can be compiled and executed in a Linux environment. The book also presents hybrid codes that integrate selected APIs for potentially multi-level parallelization and utilization of heterogeneous resources, and it shows how to use modern elements of these APIs. Selected optimization techniques are also included, such as overlapping communication and computations implemented using various APIs. Features: Discusses the popular and currently available computing devices and cluster systems Includes typical paradigms used in parallel programs Explores popular APIs for programming parallel applications Provides code templates that can be used for implementation of paradigms Provides hybrid code examples allowing multi-level parallelization Covers the optimization of parallel programs

Parallel Computing Architectures and APIs Vivek Kale 2019-12-02 This book addresses the challenge of provisioning of parallel hardware and software enabling high-performance parallel processing in the context of requirements dictated by big data systems. It gives description of parallel programming techniques that are necessary for efficient programs for multicore processors and parallel cluster systems.

Recent Advances in Parallel Virtual Machine and Message Passing Interface Matti Ropo 2009-08-28 This book constitutes the refereed proceedings of the 16th European PVM/MPI Users' Group Meeting on Recent Advances in Parallel Virtual Machine and Message Passing Interface, EuroPVM/MPI 2009, held in Espoo, Finland, September 7-10, 2009. The 27 papers presented were carefully reviewed and selected from 48 submissions. The volume also includes 6 invited talks, one tutorial, 5 poster abstracts and 4 papers from the special session on current trends in numerical simulation for parallel engineering environments. The main topics of the meeting were Message Passing Interface (MPI) performance issues in very large systems, MPI program verification and MPI on multi-core architectures.

Parallele Programmierung Thomas Rauber 2012-09-04 Multiprozessor-Desktoprechner, Cluster von PCs und Innovationen wie Hyperthreading oder Multicore-Prozessoren machen parallele Rechenressourcen allgegenwärtig. Die Ausnutzung dieser Rechenleistung ist jedoch nur durch parallele Programmieretechniken möglich. Das Buch stellt diese Techniken für herkömmliche Parallelrechner und für neuartige Plattformen umfassend dar. Neben den Grundlagen der parallelen Programmierung werden Programmierumgebungen wie Pthreads, Java-Threads,

OpenMP, MPI oder PVM sowie die zugehörigen Programmiermodelle behandelt.

Complex Systems Design & Management Marc Aiguier 2010-10-01 This book contains all refereed papers that were accepted to the "Complex Systems Design & Management" (CSDM 2010) international conference that took place in Paris (France), October 27 – 29, 2010 (Website: <http://www.csdm2010.csdm.fr>). These proceedings covers the most recent trends in the emerging field of complex systems sciences & practices from an industrial and academic perspective, including the main industrial domains (transport, defense & security, electronics, energy & environment, health, communications & media, e-services), scientific & technical topics (systems fundamentals, systems architecture & engineering, systems metrics & quality, systemic tools) and system types (transportation systems, embedded systems, software & information systems, systems of systems, artificial ecosystems). The CSDM 2010 conference is organized under the guidance of the CESAMES non profit organization (Website: <http://www.cesames.net>).

Parallel Programming Bertil Schmidt 2017-11-20 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

Algorithms Sequential & Parallel: A Unified Approach Russ Miller 2012-12-20 Equip yourself for success with a state-of-the-art approach to algorithms available only in Miller/Boxer's ALGORITHMS SEQUENTIAL AND PARALLEL: A UNIFIED APPROACH, 3E. This unique and functional text gives you an introduction to algorithms and paradigms for modern computing systems, integrating the study of parallel and sequential algorithms within a focused presentation. With a wide range of practical exercises and engaging examples drawn from fundamental application domains, this book prepares you to design, analyze, and implement algorithms for modern computing systems. Important Notice: Media content referenced within the product description or the product text may not be available in the ebook version.

Operating Systems for Supercomputers and High Performance Computing Balazs Gerofi 2019-10-15 Few works are as timely and critical to the advancement of high performance computing than is this new up-to-date treatise on leading-edge directions of operating systems. It is a first-hand product of many of the leaders in this rapidly evolving field and possibly the most comprehensive. This new and important book masterfully presents the major alternative concepts driving the future of operating system design for high performance computing. In particular, it describes the major advances of monolithic operating systems such as Linux and Unix that dominate the TOP500 list. It also presents the state of the art in lightweight kernels that exhibit high efficiency and scalability at the loss of generality. Finally, this work looks forward to possibly the most promising strategy of a hybrid structure combining full service functionality with lightweight kernel operation. With this, it is likely that this new work will find its way on the shelves of almost everyone who is in any way engaged in the multi-discipline of high performance computing. (From the foreword by Thomas Sterling)

Rechnerorganisation und Rechnerentwurf David Patterson 2011-03-09 Mit der deutschen Übersetzung zur vierten Auflage des amerikanischen Klassikers Computer Organization and Design. The Hardware/Software Interface ist das Standardwerk zur Rechnerorganisation wieder auf dem neusten Stand - David A. Patterson und John L. Hennessy gewähren die gewohnten Einblicke in das Zusammenwirken von Hard- und Software, Leistungseinschätzungen und zahlreicher Rechnerkonzepte in einer Tiefe, die zusammen mit klarer Didaktik und einer eher lockeren Sprache den Erfolg dieses weltweit anerkannten Standardwerks begründen. Patterson und Hennessy achten darauf, nicht nur auf das "Wie" der dargestellten Konzepte, sondern auch auf ihr "Warum" einzugehen und zeigen damit Gründe für Veränderungen und neue Entwicklungen auf. Jedes der Kapitel steht für einen deutlich umrissenen Teilbereich der Rechnerorganisation und ist jeweils gleich aufgebaut: Eine Einleitung, gefolgt von immer tiefgreifenderen Grundkonzepten mit steigender Komplexität. Darauf eine aktuelle Fallstudie, "Fallstricke und Fehlschlüsse", Zusammenfassung und Schlussbetrachtung, historische Perspektiven und Literaturhinweise sowie Aufgaben. Umfangreiches Zusatzmaterial (Werkzeuge mit Tutorien etc.) steht auf der beiliegenden CD-ROM zur Verfügung.

Parallel Computing Architectures and APIs Vivek Kale 2019-12-06 Parallel Computing Architectures and APIs: IoT Big Data Stream Processing commences from the point high-performance uniprocessors were becoming increasingly complex, expensive, and power-hungry. A basic trade-off exists between the use of one or a small number of such complex processors, at one extreme, and a moderate to very large number of simpler processors, at the other. When combined with a high-bandwidth, interprocessor communication facility leads to significant simplification of the design process. However, two major roadblocks prevent the widespread adoption of such moderately to massively parallel architectures: the interprocessor communication bottleneck, and the difficulty and high cost of algorithm/software development. One of the most important reasons for studying parallel computing

architectures is to learn how to extract the best performance from parallel systems. Specifically, you must understand its architectures so that you will be able to exploit those architectures during programming via the standardized APIs. This book would be useful for analysts, designers and developers of high-throughput computing systems essential for big data stream processing emanating from IoT-driven cyber-physical systems (CPS). This pragmatic book: Devolves uniprocessors in terms of a ladder of abstractions to ascertain (say) performance characteristics at a particular level of abstraction Explains limitations of uniprocessor high performance because of Moore's Law Introduces basics of processors, networks and distributed systems Explains characteristics of parallel systems, parallel computing models and parallel algorithms Explains the three primary categorical representatives of parallel computing architectures, namely, shared memory, message passing and stream processing Introduces the three primary categorical representatives of parallel programming APIs, namely, OpenMP, MPI and CUDA Provides an overview of Internet of Things (IoT), wireless sensor networks (WSN), sensor data processing, Big Data and stream processing Provides introduction to 5G communications, Edge and Fog computing Parallel Computing Architectures and APIs: IoT Big Data Stream Processing discusses stream processing that enables the gathering, processing and analysis of high-volume, heterogeneous, continuous Internet of Things (IoT) big data streams, to extract insights and actionable results in real time. Application domains requiring data stream management include military, homeland security, sensor networks, financial applications, network management, web site performance tracking, real-time credit card fraud detection, etc.

Cloud Computing and Big Data: Technologies, Applications and Security Mostapha Zbakh 2018-07-27 This book addresses topics related to cloud and Big Data technologies, architecture and applications including distributed computing and data centers, cloud infrastructure and security, and end-user services. The majority of the book is devoted to the security aspects of cloud computing and Big Data. Cloud computing, which can be seen as any subscription-based or pay-per-use service that extends the Internet's existing capabilities, has gained considerable attention from both academia and the IT industry as a new infrastructure requiring smaller investments in hardware platforms, staff training, or licensing software tools. It is a new paradigm that has ushered in a revolution in both data storage and computation. In parallel to this progress, Big Data technologies, which rely heavily on cloud computing platforms for both data storage and processing, have been developed and deployed at breathtaking speed. They are among the most frequently used technologies for developing applications and services in many fields, such as the web, health, and energy. Accordingly, cloud computing and Big Data technologies are two of the most central current and future research mainstreams. They involve and impact a host of fields, including business, scientific research, and public and private administration. Gathering extended versions of the best papers presented at the Third International Conference on Cloud Computing Technologies and Applications (CloudTech'17), this book offers a valuable resource for all Information System managers, researchers, students, developers, and policymakers involved in the technological and application aspects of cloud computing and Big Data.

Rechnerorganisation und Rechnerentwurf David Patterson 2016-05-24 Mit der deutschen Übersetzung zur fünfter Auflage des amerikanischen Klassikers Computer Organization and Design - The Hardware/Software Interface ist das Standardwerk zur Rechnerorganisation wieder auf dem neusten Stand - David A. Patterson und John L. Hennessy gewähren die gewohnten Einblicke in das Zusammenwirken von Hard- und Software, Leistungseinschätzungen und zahlreicher Rechnerkonzepte in einer Tiefe, die zusammen mit klarer Didaktik und einer eher lockeren Sprache den Erfolg dieses weltweit anerkannten Standardwerks begründen. Patterson und Hennessy achten darauf, nicht nur auf das "Wie" der dargestellten Konzepte, sondern auch auf ihr "Warum" einzugehen und zeigen damit Gründe für Veränderungen und neue Entwicklungen auf. Jedes der Kapitel steht für einen deutlich umrissenen Teilbereich der Rechnerorganisation und ist jeweils gleich aufgebaut: Eine Einleitung, gefolgt von immer tiefgreifenderen Grundkonzepten mit steigender Komplexität. Darauf eine aktuelle Fallstudie, "Fallstricke und Fehlschlüsse", Zusammenfassung und Schlussbetrachtung, historische Perspektiven und Literaturhinweise sowie Aufgaben. In der neuen Auflage sind die Inhalte in den Kapiteln 1-5 an vielen Stellen punktuell verbessert und aktualisiert, mit der Vorstellung neuerer Prozessoren worden, und der Kapitel 6... from Client to Cloud wurde stark überarbeitet Umfangreiches Zusatzmaterial (Werkzeuge mit Tutorien etc.) steht Online zur Verfügung.

Numerische Simulation in der Moleküldynamik Michael Griebel 2013-03-07 Das Buch behandelt Methoden des wissenschaftlichen Rechnens in der Moleküldynamik, einem Bereich, der in vielen Anwendungen der Chemie, der Biowissenschaften, der Materialwissenschaften, insbesondere der Nanotechnologie, sowie der Astrophysik eine wichtige Rolle spielt. Es führt in die wichtigsten Simulationstechniken zur numerischen Behandlung der Newtonschen Bewegungsgleichungen ein. Der Schwerpunkt liegt hierbei auf der schnellen Auswertung kurz- und langreichweitiger Kräfte mittels Linked Cell-, P³M-, Baum- und Multipol-Verfahren, sowie deren paralleler Implementierung und Lastbalancierung auf Rechensystemen mit verteiltem Speicher. Die einzelnen Kapitel beinhalten darüberhinaus detaillierte Hinweise, um die Verfahren Schritt für Schritt in ein Programmpaket umzusetzen. In zahlreichen farbigen Abbildungen werden Simulationsergebnisse für eine Reihe von Anwendungen präsentiert.

Applied Computing to Support Industry: Innovation and Technology Mohammed I. Khalaf 2020-01-07 This book constitutes the refereed proceedings of the First International Conference on Applied Computing to Support

Industry: Innovation and Technology, ACRIT 2019, held in Ramadi, Iraq, in September 2019. The 38 revised full papers and 1 short paper were carefully reviewed and selected from 159 submissions. The papers of this volume are organized in topical sections on theory, methods and tools to support computer science; computer security and cryptography; computer network and communication; real world application in information science and technology.

GeoComputation Robert J. Abraham 2014-06-23 A revision of Openshaw and Abraham's seminal work, GeoComputation, Second Edition retains influences of its originators while also providing updated, state-of-the-art information on changes in the computational environment. In keeping with the field's development, this new edition takes a broader view and provides comprehensive coverage across the

Introduction to Parallel Computing Zbigniew J. Czech 2017-01-11 The constantly increasing demand for more computing power can seem impossible to keep up with. However, multicore processors capable of performing computations in parallel allow computers to tackle ever larger problems in a wide variety of applications. This book provides a comprehensive introduction to parallel computing, discussing theoretical issues such as the fundamentals of concurrent processes, models of parallel and distributed computing, and metrics for evaluating and comparing parallel algorithms, as well as practical issues, including methods of designing and implementing shared- and distributed-memory programs, and standards for parallel program implementation, in particular MPI and OpenMP interfaces. Each chapter presents the basics in one place followed by advanced topics, allowing novices and experienced practitioners to quickly find what they need. A glossary and more than 80 exercises with selected solutions aid comprehension. The book is recommended as a text for advanced undergraduate or graduate students and as a reference for practitioners.

Delivery and Adoption of Cloud Computing Services in Contemporary Organizations Chang, Victor 2015-03-31 The ubiquity of technology has not only brought the need for computer knowledge to every aspect of the modern business world; it has also increased our need to safely store the data we are now creating at a rate never experienced before. Delivery and Adoption of Cloud Computing Services in Contemporary Organizations brings together the best practices for storing massive amounts of data. Highlighting ways cloud services can work effectively in production and in real time, this book is an essential reference source for professionals and academics of various disciplines, such as computer science, consulting, information technology, information and communication sciences, healthcare, and finance.

Multicore: Thomas Rauber 2007-10-27 Multicore-Prozessoren mit zwei oder mehreren Prozessorkernen erhöhen die Leistungsfähigkeit aller Computer immens. Doch nur spezielle Techniken gewährleisten die tatsächlich schnellere Programmbearbeitung und optimale Nutzung dieser Leistungsfähigkeit. Die Autoren zeigen hier erstmals, wie Software-Entwickler parallele Programme mittels Software-Threads zur schnellen Ausführung auf Multicore-Prozessoren erstellen. Umfassend erläutern sie alle Aspekte des Themas: parallele Programmiermodelle, Konzepte der Thread-Programmierung, die Programmierumgebungen Pthreads, Java-Threads und OpenMP sowie Sprachkonstrukte und neuere Programmieransätze.

Parallel Programming with Microsoft Visual C++ Colin Campbell 2011 This guide shows Visual C++ programmers how to effectively take advantage of the multicore capabilities of modern PCs using the Microsoft platform.

OpenMP in a New Era of Parallelism Rudi Eigenmann 2008-05-01 OpenMP is a widely accepted, standard application programming interface (API) for high-level shared-memory parallel programming in Fortran, C, and C++. Since its introduction in 1997, OpenMP has gained support from most high-performance compiler and hardware vendors. Under the direction of the OpenMP Architecture Review Board (ARB), the OpenMP specification has evolved, including the recent release of Specification 3.0. Active research in OpenMP compilers, runtime systems, tools, and environments drives its evolution, including new features such as tasking. The community of OpenMP researchers and developers in academia and industry is united under cOMPunity (www.compunity.org). This organization has held workshops on OpenMP around the world since 1999: the European Workshop on OpenMP (EWOMP), the North American Workshop on OpenMP Applications and Tools (WOMPAT), and the Asian Workshop on OpenMP Experiences and Implementation (WOMPEI) attracted annual audiences from academia and industry. The International Workshop on OpenMP (IWOMP) consolidated these three workshop series into a single annual international event that rotates across the previous workshop sites. The first IWOMP meeting was held in Eugene, Oregon, USA. IWOMP 2006 took place in Reims, France, and IWOMP 2007 in Beijing, China. Each workshop drew over 60 participants from research and industry throughout the world. IWOMP 2008 continued the series with technical papers, panels, tutorials, and OpenMP status reports. The first IWOMP workshop was organized under the auspices of cOMPunity.

Smart Multicore Embedded Systems Massimo Torquati 2013-11-09 This book provides a single-source reference to the state-of-the-art of high-level programming models and compilation tool-chains for embedded system platforms. The authors address challenges faced by programmers developing software to implement parallel applications in embedded systems, where very often they are forced to rewrite sequential programs into parallel software, taking into account all the low level features and peculiarities of the underlying platforms. Readers will benefit from these authors' approach, which takes into account both the application requirements and the platform specificities of various embedded systems from different industries. Parallel programming tool-chains are described that take as input parameters both the application and the platform model, then determine relevant transformations and mapping

decisions on the concrete platform, minimizing user intervention and hiding the difficulties related to the correct and efficient use of memory hierarchy and low level code generation.

The OpenCL Programming Book <http://www.fixstars.com/en/opencv/book/>

Parallelism in Matrix Computations Efstratios Gallopoulos 2015-07-25 This book is primarily intended as a research monograph that could also be used in graduate courses for the design of parallel algorithms in matrix computations. It assumes general but not extensive knowledge of numerical linear algebra, parallel architectures, and parallel programming paradigms. The book consists of four parts: (I) Basics; (II) Dense and Special Matrix Computations; (III) Sparse Matrix Computations; and (IV) Matrix functions and characteristics. Part I deals with parallel programming paradigms and fundamental kernels, including reordering schemes for sparse matrices. Part II is devoted to dense matrix computations such as parallel algorithms for solving linear systems, linear least squares, the symmetric algebraic eigenvalue problem, and the singular-value decomposition. It also deals with the development of parallel algorithms for special linear systems such as banded, Vandermonde, Toeplitz, and block Toeplitz systems. Part III addresses sparse matrix computations: (a) the development of parallel iterative linear system solvers with emphasis on scalable preconditioners, (b) parallel schemes for obtaining a few of the extreme eigenpairs or those contained in a given interval in the spectrum of a standard or generalized symmetric eigenvalue problem, and (c) parallel methods for computing a few of the extreme singular triplets. Part IV focuses on the development of parallel algorithms for matrix functions and special characteristics such as the matrix pseudospectrum and the determinant. The book also reviews the theoretical and practical background necessary when designing these algorithms and includes an extensive bibliography that will be useful to researchers and students alike. The book brings together many existing algorithms for the fundamental matrix computations that have a proven track record of efficient implementation in terms of data locality and data transfer on state-of-the-art systems, as well as several algorithms that are presented for the first time, focusing on the opportunities for parallelism and algorithm robustness.

Languages and Compilers for Parallel Computing Sanjay Rajopadhye 2013-01-18 This book constitutes the thoroughly refereed post-conference proceedings of the 24th International Workshop on Languages and Compilers for Parallel Computing, LCPC 2011, held in Fort Collins, CO, USA, in September 2011. The 19 revised full papers presented and 19 poster papers were carefully reviewed and selected from 52 submissions. The scope of the workshop spans the theoretical and practical aspects of parallel and high-performance computing, and targets parallel platforms including concurrent, multithreaded, multicore, accelerator, multiprocessor, and cluster systems.

A Framework for Productive, Efficient and Portable Parallel Computing Ekaterina Gonina 2013 Developing efficient parallel implementations and fully utilizing the available resources of parallel platforms is now required for software applications to scale to new generations of processors. Yet, parallel programming remains challenging to programmers due to the requisite low-level knowledge of the underlying hardware and parallel computing constructs. Developing applications that effectively utilize parallel hardware is restricted by poor programmer productivity, low-level implementation requirements, and limited portability of the application code. These restrictions in turn impede experimentation with various algorithmic approaches for a given application. Currently, the programming world is divided into two types of programmers: application writers who focus on designing and prototyping applications and algorithms, and efficiency programmers who focus on extracting performance for particular compute kernels. The gap between these two types of programmers is referred to as "the implementation gap". In this dissertation, we present a software environment that aims to bridge the implementation gap and enable application writers to productively utilize parallel hardware by reusing the work of efficiency programmers. Specifically, we present PyCASP, a Python-based software framework that automatically maps Python application code to a variety of parallel platforms. PyCASP is an application-domain-specific framework that uses a systematic, pattern-oriented approach to offer a single productive software development environment for application writers. PyCASP targets audio content analysis applications, but our methodology is designed to be applicable to any application domain. Using PyCASP, applications can be prototyped in Python code and our environment enables them to automatically scale their performance to modern parallel processors such as GPUs, multicore CPUs and compute clusters. We use the Selective Embedded JIT Specialization (SEJITS) mechanism to realize the pattern-based design of PyCASP in software. We use SEJITS to implement PyCASP's components and to enable automatic parallelization of specific audio content analysis application patterns on a variety of parallel hardware. By focusing on one application domain, we enable efficient composition of computations using three structural patterns: MapReduce, Iterator and Pipe-and-Filter. To illustrate our approach, we study a set of four example audio content analysis applications that are architected and implemented using PyCASP: a speaker verification system, a speaker diarization system, a music recommendation system and a video event detection system. We describe the detailed implementation of two computational components of PyCASP: a Gaussian Mixture Model (GMM) component and a Support Vector Machine (SVM) component and their use in implementing the example applications. We also analyze composition of computations using the three structural patterns and implement the available optimizations for composing computations in audio analysis applications. We evaluate our approach with results on productivity and performance using the two computational components and the four example applications. Our results illustrate that we can prototype the full-functioning applications in Python using 10 - 60x

less lines of code than equivalent implementations using low-level languages. Our PyCASP components and example applications achieve and often exceed the efficiency of comparable hand-tuned low-level implementations. In addition to specialization, adding the optimizations for composing components in these applications can give up to 30% performance improvement. We show that applications written using PyCASP can be run on multiple parallel hardware backends with little or no application code change. PyCASP also enables applications to scale from one desktop GPU to a cluster of GPUs with little programmer effort. Combining all of the specialization and composition techniques, our example applications are able to automatically achieve 50-1000x faster-than-real-time performance on both multi-core CPU and GPU platforms and 15.5x speedup on 16-node cluster of GPUs showing near-optimal scaling.

Languages and Compilers for Parallel Computing Keith Cooper 2011-02-24 This book constitutes the thoroughly refereed post-proceedings of the 23rd International Workshop on Languages and Compilers for Parallel Computing, LCPC 2010, held in Houston, TX, USA, in October 2010. The 18 revised full papers presented were carefully reviewed and selected from 47 submissions. The scope of the workshop spans foundational results and practical experience, and targets all classes of parallel platforms including concurrent, multithreaded, multicore, accelerated, multiprocessor, and cluster systems

Internet of Things, Smart Spaces, and Next Generation Networks and Systems Olga Galinina 2016-09-19 This book constitutes the joint refereed proceedings of the 16th International Conference on Next Generation Wired/Wireless Advanced Networks and Systems, NEW2AN 2016, and the 9th Conference on Internet of Things and Smart Spaces, ruSMART 2016, held in St. Petersburg, Russia, in September 2016. The 69 revised full papers were carefully reviewed and selected from 204 submissions. The 12 papers selected for ruSMART are organized in topical sections on new generation of smart services; smart services serving telecommunication networks; role of context for smart services; and smart services in automotive industry. The 57 papers from NEW2AN deal with the following topics: cooperative communications; wireless networks; wireless sensor networks; security issues; IoT and industrial IoT; NoC and positioning; ITS; network issues; SDN; satellite communications; signals and circuits; advanced materials and their properties; and economics and business.

Encyclopedia of Parallel Computing David Padua 2011-09-08 Containing over 300 entries in an A-Z format, the Encyclopedia of Parallel Computing provides easy, intuitive access to relevant information for professionals and researchers seeking access to any aspect within the broad field of parallel computing. Topics for this comprehensive reference were selected, written, and peer-reviewed by an international pool of distinguished researchers in the field. The Encyclopedia is broad in scope, covering machine organization, programming languages, algorithms, and applications. Within each area, concepts, designs, and specific implementations are presented. The highly-structured essays in this work comprise synonyms, a definition and discussion of the topic, bibliographies, and links to related literature. Extensive cross-references to other entries within the Encyclopedia support efficient, user-friendly searches for immediate access to useful information. Key concepts presented in the Encyclopedia of Parallel Computing include; laws and metrics; specific numerical and non-numerical algorithms; asynchronous algorithms; libraries of subroutines; benchmark suites; applications; sequential consistency and cache coherency; machine classes such as clusters, shared-memory multiprocessors, special-purpose machines and dataflow machines; specific machines such as Cray supercomputers, IBM's cell processor and Intel's multicore machines; race detection and auto parallelization; parallel programming languages, synchronization primitives, collective operations, message passing libraries, checkpointing, and operating systems. Topics covered: Speedup, Efficiency, Isoefficiency, Redundancy, Amdahls law, Computer Architecture Concepts, Parallel Machine Designs, Benchmarks, Parallel Programming concepts & design, Algorithms, Parallel applications. This authoritative reference will be published in two formats: print and online. The online edition features hyperlinks to cross-references and to additional significant research. Related Subjects: supercomputing, high-performance computing, distributed computing

Flexible Approaches in Data, Information and Knowledge Management Olivier Pivert 2013-09-12 This volume showcases contributions from internationally-known researchers in the field of information management. Most of the approaches presented here make use of fuzzy logic, introduced by L.A. Zadeh almost 50 years ago, which constitute a powerful tool to model and handle gradual concepts. What all of these contributions have in common is placing the user at the center of the information system, be it for helping him/her to query a data set, to handle imperfect information, or to discover useful knowledge from a massive collection of data. Researchers working in data and knowledge management will greatly benefit from this collection of up-to-date studies. This may be also an invaluable source of information for postgraduate students interested in advanced information management techniques.