# Bioinformatics High Performance Parallel Computer Architectures Embedded Multi Core Systems

High Performance Computing Systems and ApplicationsAdvances in ComputersCatalyzing Inquiry at the Interface of Computing and BiologyBeowulf Cluster Computing with LinuxDistributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted LivingDeveloping Bioinformatics Computer SkillsBioinformatics and Biomedical EngineeringReconfigurable LogicAdvances in Bioinformatics and Computational BiologyPattern Recognition in BioinformaticsHigh-Performance Computing Using FPGAsMassively Parallel Evolutionary Computation on GPGPUsParallel ComputingApplied Computing to Support Industry: Innovation and TechnologyR/parallelEuro-Par 2013: Parallel Processing WorkshopsComputational Modeling in Bioengineering and BioinformaticsHandbook of Bioinspired Algorithms and ApplicationsHigh Performance Computing in Power and Energy SystemsParallel ProgrammingHands-On GPU Computing with PythonEmerging Trends in Computational Biology, Bioinformatics, and Systems BiologyProgramming Massively Parallel ProcessorsGrid Computing for Bioinformatics and Computational BiologyApplications and Developments in Grid, Cloud, and High Performance ComputingParallel Computing for Bioinformatics and Computational

BiologySupercomputingEuro-Par 2014: Parallel Processing WorkshopsTechniques for Optimizing ApplicationsBioinformatics ComputingPetascale ComputingHigh-Performance Computational Solutions in Protein BioinformaticsAnt Colony Optimization and Swarm IntelligenceDeep Learning and Parallel Computing Environment for Bioengineering SystemsBioinformaticsMulticore ComputingAdvances in High Performance ComputingHigh-Performance Modelling and Simulation for Big Data ApplicationsParallel Computing TechnologiesVLSI

# High Performance Computing Systems and Applications

New sequencing technologies have broken many experimental barriers to genome scale sequencing, leading to the extraction of huge quantities of sequence data. This expansion of biological databases established the need for new ways to harness and apply the astounding amount of available genomic information and convert it into substantive biological understanding. A complilation of recent approaches from prominent researchers, Bioinformatics: High Performance Parallel Computer Architectures discusses how to take advantage of bioinformatics applications and algorithms on a variety of modern parallel architectures. Two factors continue to drive the increasing use of modern parallel computer architectures to address problems in computational biology and bioinformatics: high-throughput techniques for DNA sequencing and gene expression

analysis—which have led to an exponential growth in the amount of digital biological data—and the multi- and many-core revolution within computer architecture. Presenting key information about how to make optimal use of parallel architectures, this book: Describes algorithms and tools including pairwise sequence alignment, multiple sequence alignment, BLAST, motif finding, pattern matching, sequence assembly, hidden Markov models, proteomics, and evolutionary tree reconstruction Addresses GPGPU technology and the associated massively threaded CUDA programming model Reviews FPGA architecture and programming Presents several parallel algorithms for computing alignments on the Cell/BE architecture, including linear-space pairwise alignment, syntenic alignment, and spliced alignment Assesses underlying concepts and advances in orchestrating the phylogenetic likelihood function on parallel computer architectures (ranging from FPGAs upto the IBM BlueGene/L supercomputer) Covers several effective techniques to fully exploit the computing capability of many-core CUDA-enabled GPUs to accelerate protein sequence database searching, multiple sequence alignment, and motif finding Explains a parallel CUDA-based method for correcting sequencing base-pair errors in HTSR data Because the amount of publicly available sequence data is growing faster than single processor core performance speed, modern bioinformatics tools need to take advantage of parallel computer architectures. Now that the era of the many-core processor has begun, it is expected that future mainstream processors will be parallel systems. Beneficial to anyone actively involved in research and applications, this book helps you to get

the most out of these tools and create optimal HPC solutions for bioinformatics.

## Advances in Computers

Traditionally, parallel computing has been associated with special purpose applications designed to run in complex computing clusters, specifically set up with a software stack of dedicated libraries together with advanced administration tools to manage co Traditionally, parallel computing has been associated with special purpose applications designed to run in complex computing clusters, specifically set up with a software stack of dedicated libraries together with advanced administration tools to manage complex IT infrastructures. These High Performance Computing (HPC) solutions, although being the most efficient solutions in terms of performance and scalability, impose technical and practical barriers for most common scientists whom, with reduced IT knowledge, time and resources, are unable to embrace classical HPC solutions without considerable efforts. Moreover, two important technology advances are increasing the need for parallel computing. For example in the bioinformatics field, and similarly in other experimental science disciplines, new high throughput screening devices are generating huge amounts of data within very short time which requires their analysis in equally short time periods to avoid delaying experimental analysis. Another important technological change involves the design of new processor chips. To increase raw performance the current strategy is to increase the number

of processing units per chip, so to make use of the new processing capacities parallel applications are required. In both cases we find users that may need to update their current sequential applications and computing resources to achieve the increased processing capacities required for their particular needs. Since parallel computing is becoming a natural option for obtaining increased performance and it is required by new computer systems, solutions adapted for the mainstream should be developed for a seamless adoption. In order to enable the adoption of parallel computing, new methods and technologies are required to remove or mitigate the current barriers and obstacles that prevent many users from evolving their sequential running environments. A particular scenario that specially suffers from these problems and that is considered as a practical case in this work consists of bioinformaticians analyzing molecular data with methods written with the R language. In many cases, with long datasets, they have to wait for days and weeks for their data to be processed or perform the cumbersome task of manually splitting their data, look for available computers to run these subsets and collect back the previously scattered results. Most of these applications written in R are based on parallel loops. A loop is called a parallel loop if there is no data dependency among all its iterations, and therefore any iteration can be processed in any order or even simultaneously, so they are susceptible of being parallelized. Parallel loops are found in a large number of scientific applications. Previous contributions deal with partial aspects of the problems suffered by this kind of users, such as providing access to additional computing resources or enabling the

codification of parallel problems, but none takes proper care of providing complete solutions without considering advanced users with access to traditional HPC platforms. Our contribution consists in the design and evaluation of methods to enable the easy parallelization of applications based in parallel loops written in R using non-dedicated environments as a computing platform and considering users without proper experience in parallel computing or system management skills. As a proof of concept, and in order to evaluate the feasibility of our proposal, an extension of R, called R/parallel, has been developed to test our ideas in real environments with real bioinformatics problems. The results show that even in situations with a reduced level of information about the running environment and with a high degree of uncertainty about the quantity and quality of the available resources it is possible to provide a software layer to enable users without previous knowledge and skills adapt their applications with a minimal effort and perform concurrent computations using the available computers. Additionally of proving the feasibility of our proposal, a new self-scheduling scheme, suitable for parallel loops in dynamics environments has been contributed, the results of which show that it is possible to obtain improved performance levels compared to previous contributions in best-effort environments. The main conclusion is that, even in situations with limited information about the environment and the involved technologies, it is possible to provide the mechanisms that will allow users without proper knowledge and time restrictions to conveniently make use and take advantage of parallel computing technologies, so closing the gap between classical

HPC solutions and the mainstream of users of common applications, in our case, based in parallel loops with R. mplex IT infrastructures. These High Performance Computing (HPC) solutions, although being the most efficient solutions in terms of performance and scalability, impose technical and practical barriers for most common scientists whom, with reduced IT knowledge, time and resources, are unable to embrace classical HPC solutions without considerable efforts. Moreover, two important technology advances are increasing the need for parallel computing. For example in the bioinformatics field, and similarly in other experimental science disciplines, new high throughput screening devices are generating huge amounts of data within very short time which requires their analysis in equally short time periods to avoid delaying experimental analysis. Another important technological change involves the design of new processor chips. To increase raw performance the current strategy is to increase the number of processing units per chip, so to make use of the new processing capacities parallel applications are required. In both cases we find users that may need to update their current sequential applications and computing resources to achieve the increased processing capacities required for their particular needs. Since parallel computing is becoming a natural option for obtaining increased performance and it is required by new computer systems, solutions adapted for the mainstream should be developed for a seamless adoption. In order to enable the adoption of parallel computing, new methods and technologies are required to remove or mitigate the current barriers and obstacles that prevent many users from evolving their sequential running

environments. A particular scenario that specially suffers from these problems and that is considered as a practical case in this work consists of bioinformaticians analyzing molecular data with methods written with the R language. In many cases, with long datasets, they have to wait for days and weeks for their data to be processed or perform the cumbersome task of manually splitting their data, look for available computers to run these subsets and collect back the previously scattered results. Most of these applications written in R are based on parallel loops. A loop is called a parallel loop if there is no data dependency among all its iterations, and therefore any iteration can be processed in any order or even simultaneously, so they are susceptible of being parallelized. Parallel loops are found in a large number of scientific applications. Previous contributions deal with partial aspects of the problems suffered by this kind of users, such as providing access to additional computing resources or enabling the codification of parallel problems, but none takes proper care of providing complete solutions without considering advanced users with access to traditional HPC platforms. Our contribution consists in the design and evaluation of methods to enable the easy parallelization of applications based in parallel loops written in R using non-dedicated environments as a computing platform and considering users without proper experience in parallel computing or system management skills. As a proof of concept, and in order to evaluate the feasibility of our proposal, an extension of R, called R/parallel, has been developed to test our ideas in real environments with real bioinformatics problems. The results show that even in situations with a

reduced level of information about the running environment and with a high degree of uncertainty about the quantity and quality of the available resources it is possible to provide a software layer to enable users without previous knowledge and skills adapt their applications with a minimal effort and perform concurrent computations using the available computers. Additionally of proving the feasibility of our proposal, a new self-scheduling scheme, suitable for parallel loops in dynamics environments has been contributed, the results of which show that it is possible to obtain improved performance levels compared to previous contributions in best-effort environments. The main conclusion is that, even in situations with limited information about the environment and the involved technologies, it is possible to provide the mechanisms that will allow users without proper knowledge and time restrictions to conveniently make use and take advantage of parallel computing technologies, so closing the gap between classical HPC solutions and the mainstream of users of common applications, in our case, based in parallel loops with R.

## Catalyzing Inquiry at the Interface of Computing and Biology

This open access book was prepared as a Final Publication of the COST Action IC1406 "High-Performance Modelling and Simulation for Big Data Applications (cHiPSet)" project. Long considered important pillars of the scientific method, Modelling and Simulation have evolved from traditional discrete numerical

methods to complex data-intensive continuous analytical optimisations. Resolution, scale, and accuracy have become essential to predict and analyse natural and complex systems in science and engineering. When their level of abstraction raises to have a better discernment of the domain at hand, their representation gets increasingly demanding for computational and data resources. On the other hand, High Performance Computing typically entails the effective use of parallel and distributed processing units coupled with efficient storage, communication and visualisation systems to underpin complex data-intensive applications in distinct scientific and technical domains. It is then arguably required to have a seamless interaction of High Performance Computing with Modelling and Simulation in order to store, compute, analyse, and visualise large data sets in science and engineering. Funded by the European Commission, cHiPSet has provided a dynamic trans-European forum for their members and distinguished guests to openly discuss novel perspectives and topics of interests for these two communities. This cHiPSet compendium presents a set of selected case studies related to healthcare, biological data, computational advertising, multimedia, finance, bioinformatics, and telecommunications.

## Beowulf Cluster Computing with Linux

Computational Modeling in Bioengineering and Bioinformatics promotes complementary disciplines that hold great promise for the advancement of

research and development in complex medical and biological systems, and in the environment, public health, drug design, and so on. It provides a common platform by bridging these two very important and complementary disciplines into an interactive and attractive forum. Chapters cover biomechanics and bioimaging, biomedical decision support system, data mining, personalized diagnoses, bio-signal processing, protein structure prediction, tissue and cell engineering, biomedical image processing, analysis and visualization, high performance computing and sports bioengineering. The book's chapters are the result of many international projects in the area of bioengineering and bioinformatics done at the Research and Development Center for Bioengineering BioIRC and by the Faculty of Engineering at the University of Kragujevac, Serbia. Presents recent advances at the crossroads of biomedical engineering and bioinformatics, one of the hottest areas in biomedical and clinical research Discusses a wide range of leading-edge research topics, including biomechanics and bioimaging, biomedical decision support systems, data mining, personalized diagnoses, bio-signal processing, protein structure prediction, tissue and cell engineering, amongst others Includes coverage of biomechanical, bioengineering and computational methods of treatment and diagnosis

## Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living

The mystique of biologically inspired (or bioinspired) paradigms is their ability to describe and solve complex relationships from intrinsically very simple initial conditions and with little or no knowledge of the search space. Edited by two prominent, well-respected researchers, the Handbook of Bioinspired Algorithms and Applications reveals the

## Developing Bioinformatics Computer Skills

Recent developments in computer science enable algorithms previously perceived as too time-consuming to now be efficiently used for applications in bioinformatics and life sciences. This work focuses on proteins and their structures, protein structure similarity searching at main representation levels and various techniques that can be used to accelerate similarity searches. Divided into four parts, the first part provides a formal model of 3D protein structures for functional genomics, comparative bioinformatics and molecular modeling. The second part focuses on the use of multithreading for efficient approximate searching on protein secondary structures. The third and fourth parts concentrate on finding 3D protein structure similarities with the support of GPUs and cloud computing. Parts three and four both describe the acceleration of different methods. The text will be of interest to researchers and software developers working in the field of structural bioinformatics and biomedical databases.

# Bioinformatics and Biomedical Engineering

The twin challenge of meeting global energy demands in the face of growing economies and populations and restricting greenhouse gas emissions is one of the most daunting ones that humanity has ever faced. Smart electrical generation and distribution infrastructure will play a crucial role in meeting these challenges. We would need to develop capabilities to handle large volumes of data generated by the power system components like PMUs, DFRs and other data acquisition devices as well as by the capacity to process these data at high resolution via multi-scale and multi-period simulations, cascading and security analysis, interaction between hybrid systems (electric, transport, gas, oil, coal, etc.) and so on, to get meaningful information in real time to ensure a secure, reliable and stable power system grid. Advanced research on development and implementation of market-ready leading-edge high-speed enabling technologies and algorithms for solving real-time, dynamic, resource-critical problems will be required for dynamic security analysis targeted towards successful implementation of Smart Grid initiatives. This books aims to bring together some of the latest research developments as well as thoughts on the future research directions of the high performance computing applications in electric power systems planning, operations, security, markets, and grid integration of alternate sources of energy, etc.

# Reconfigurable Logic

Explore GPU-enabled programmable environment for machine learning, scientific applications, and gaming using PuCUDA, PyOpenGL, and Anaconda Accelerate Key Features Understand effective synchronization strategies for faster processing using GPUs Write parallel processing scripts with PyCuda and PyOpenCL Learn to use the CUDA libraries like CuDNN for deep learning on GPUs Book Description GPUs are proving to be excellent general purpose-parallel computing solutions for high performance tasks such as deep learning and scientific computing. This book will be your guide to getting started with GPU computing. It will start with introducing GPU computing and explain the architecture and programming models for GPUs. You will learn, by example, how to perform GPU programming with Python, and you'll look at using integrations such as PyCUDA, PyOpenCL, CuPy and Numba with Anaconda for various tasks such as machine learning and data mining. Going further, you will get to grips with GPU work flows, management, and deployment using modern containerization solutions. Toward the end of the book, you will get familiar with the principles of distributed computing for training machine learning models and enhancing efficiency and performance. By the end of this book, you will be able to set up a GPU ecosystem for running complex applications and data models that demand great processing capabilities, and be able to efficiently manage memory to compute your application effectively and quickly. What you will learn Utilize Python libraries and frameworks for GPU

acceleration Set up a GPU-enabled programmable machine learning environment on your system with Anaconda Deploy your machine learning system on cloud containers with illustrated examples Explore PyCUDA and PyOpenCL and compare them with platforms such as CUDA, OpenCL and ROCm. Perform data mining tasks with machine learning models on GPUs Extend your knowledge of GPU computing in scientific applications Who this book is for Data Scientist, Machine Learning enthusiasts and professionals who wants to get started with GPU computation and perform the complex tasks with low-latency. Intermediate knowledge of Python programming is assumed.

## Advances in Bioinformatics and Computational Biology

This book constitutes the refereed proceedings of the 7th Brazilian Symposium on Bioinformatics, BSB 2012, held in Campo Grande, Brazil, in August 2012. The 16 regular papers presented were carefully reviewed and selected for inclusion in this book. It also contains a joint paper from two of the guest speakers. The Brazilian Symposium on Bioinformatics covers all aspects of bioinformatics and computational biology, including sequence analysis; motifs, and pattern matching; biological databases, data management, data integration, and data mining; biomedical text mining; structural, comparative, and functional genomics; personal genomics; protein structure, modeling, and simulation; gene identification, regulation and expression analysis; gene and protein interaction and networks;

molecular docking; molecular evolution and phylogenetics; computational systems biology; computational proteomics; statistical analysis of molecular sequences; algorithms for problems in computational biology; applications in molecular biology, biochemistry, genetics, medicine, microbiology and associated subjects.

## Pattern Recognition in Bioinformatics

Recently the world celebrated the 60th anniversary of the invention of the first transistor. The first integrated circuit (IC) was built a decade later, with the first microprocessor designed in the early 1970s. Today, ICs are a part of nearly every aspect of our daily lives. They help us live longer and more comfortably, and do more, faster. All this is possible because of the relentless search for new materials, circuit designs, and ideas happening on a daily basis at industrial and academic institutions around the globe. Showcasing the latest advances in very-large-scale integrated (VLSI) circuits, VLSI: Circuits for Emerging Applications provides a balanced view of industrial and academic developments beyond silicon and complementary metal–oxide–semiconductor (CMOS) technology. From quantum-dot cellular automata (QCA) to chips for cochlear implants, this must-have resource: Investigates the trend of combining multiple cores in a single chip to boost performance of the overall system Describes a novel approach to enable physically unclonable functions (PUFs) using intrinsic features of a VLSI chip Examines the VLSI implementations of major symmetric and asymmetric key

cryptographic algorithms, hash functions, and digital signatures Discusses nonvolatile memories such as resistive random-access memory (Re-RAM), magneto-resistive RAM (MRAM), and floating-body RAM (FB-RAM) Explores organic transistors, soft errors, photonics, nanoelectromechanical (NEM) relays, reversible computation, bioinformatics, asynchronous logic, and more VLSI: Circuits for Emerging Applications presents cutting-edge research, design architectures, materials, and uses for VLSI circuits, offering valuable insight into the current state of the art of micro- and nanoelectronics.

## High-Performance Computing Using FPGAs

The two volume set LNCS 9043 and 9044 constitutes the refereed proceedings of the Third International Conference on Bioinformatics and Biomedical Engineering, IWBBIO 2015, held in Granada, Spain, in April 2015. The 135 papers presented were carefully reviewed and selected from 268 submissions. The scope of the conference spans the following areas: bioinformatics for healthcare and diseases, biomedical engineering, biomedical image analysis, biomedical signal analysis, computational genomics, computational proteomics, computational systems for modelling biological processes, e Health, next generation sequencing and sequence analysis, quantitative and systems pharmacology, Hidden Markov Model (HMM) for biological sequence modeling, advances in computational intelligence for bioinformatics and biomedicine, tools for next generation sequencing data

analysis, dynamics networks in system medicine, interdisciplinary puzzles of measurements in biological systems, biological networks, high performance computing in bioinformatics, computational biology and computational chemistry, advances in drug discovery and ambient intelligence for bio emotional computing.

## Massively Parallel Evolutionary Computation on GPGPUs

"This book provides insight into the current trends and emerging issues by investigating grid and cloud evolution, workflow management, and the impact new computing systems have on the education fields as well as the industries"--Provided by publisher.

## Parallel Computing

## Applied Computing to Support Industry: Innovation and Technology

Evolutionary algorithms (EAs) are metaheuristics that learn from natural collective behavior and are applied to solve optimization problems in domains such as scheduling, engineering, bioinformatics, and finance. Such applications demand

acceptable solutions with high-speed execution using finite computational resources. Therefore, there have been many attempts to develop platforms for running parallel EAs using multicore machines, massively parallel cluster machines, or grid computing environments. Recent advances in general-purpose computing on graphics processing units (GPGPU) have opened up this possibility for parallel EAs, and this is the first book dedicated to this exciting development. The three chapters of Part I are tutorials, representing a comprehensive introduction to the approach, explaining the characteristics of the hardware used, and presenting a representative project to develop a platform for automatic parallelization of evolutionary computing (EC) on GPGPUs. The 10 chapters in Part II focus on how to consider key EC approaches in the light of this advanced computational technique, in particular addressing generic local search, tabu search, genetic algorithms, differential evolution, swarm optimization, ant colony optimization, systolic genetic search, genetic programming, and multiobjective optimization. The 6 chapters in Part III present successful results from real-world problems in data mining, bioinformatics, drug discovery, crystallography, artificial chemistries, and sudoku. Although the parallelism of EAs is suited to the single-instruction multiple-data (SIMD)-based GPU, there are many issues to be resolved in design and implementation, and a key feature of the contributions is the practical engineering advice offered. This book will be of value to researchers, practitioners, and graduate students in the areas of evolutionary computation and scientific computing.

# R/parallel

Deep Learning and Parallel Computing Environment for Bioengineering Systems delivers a significant forum for the technical advancement of deep learning in parallel computing environment across bio-engineering diversified domains and its applications. Pursuing an interdisciplinary approach, it focuses on methods used to identify and acquire valid, potentially useful knowledge sources. Managing the gathered knowledge and applying it to multiple domains including health care, social networks, mining, recommendation systems, image processing, pattern recognition and predictions using deep learning paradigms is the major strength of this book. This book integrates the core ideas of deep learning and its applications in bio engineering application domains, to be accessible to all scholars and academicians. The proposed techniques and concepts in this book can be extended in future to accommodate changing business organizations' needs as well as practitioners' innovative ideas. Presents novel, in-depth research contributions from a methodological/application perspective in understanding the fusion of deep machine learning paradigms and their capabilities in solving a diverse range of problems Illustrates the state-of-the-art and recent developments in the new theories and applications of deep learning approaches applied to parallel computing environment in bioengineering systems Provides concepts and technologies that are successfully used in the implementation of today's intelligent data-centric critical systems and multi-media Cloud-Big data

# Euro-Par 2013: Parallel Processing Workshops

Although the highly anticipated petascale computers of the near future will perform at an order of magnitude faster than today's quickest supercomputer, the scaling up of algorithms and applications for this class of computers remains a tough challenge. From scalable algorithm design for massive concurrency toperformance analyses and scientific visualization, Petascale Computing: Algorithms and Applications captures the state of the art in high-performance computing algorithms and applications. Featuring contributions from the world's leading experts in computational science, this edited collection explores the use of petascale computers for solving the most difficult scientific and engineering problems of the current century. Covering a wide range of important topics, the book illustrates how petascale computing can be applied to space and Earth science missions, biological systems, weather prediction, climate science, disasters, black holes, and gamma ray bursts. It details the simulation of multiphysics, cosmological evolution, molecular dynamics, and biomolecules. The book also discusses computational aspects that include the Uintah framework, Enzo code, multithreaded algorithms, petaflops, performance analysis tools, multilevel finite element solvers, finite element code development, Charm++, and the Cactus framework. Supplying petascale tools, programming methodologies, and an eight-page color insert, this volume addresses the challenging problems of developing application codes that can take advantage of the architectural features

of the new petascale systems in advance of their first deployment.

## Computational Modeling in Bioengineering and Bioinformatics

This book constitutes thoroughly refereed post-conference proceedings of the workshops of the 19th International Conference on Parallel Computing, Euro-Par 2013, held in Aachen, Germany in August 2013. The 99 papers presented were carefully reviewed and selected from 145 submissions. The papers include seven workshops that have been co-located with Euro-Par in the previous years: - Big Data Cloud (Second Workshop on Big Data Management in Clouds) - Hetero Par (11th Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) - HiBB (Fourth Workshop on High Performance Bioinformatics and Biomedicine) - OMHI (Second Workshop on On-chip Memory Hierarchies and Interconnects) - PROPER (Sixth Workshop on Productivity and Performance) - Resilience (Sixth Workshop on Resiliency in High Performance Computing with Clusters, Clouds, and Grids) - UCHPC (Sixth Workshop on Un Conventional High Performance Computing) as well as six newcomers: - DIHC (First Workshop on Dependability and Interoperability in Heterogeneous Clouds) - Fed ICI (First Workshop on Federative and Interoperable Cloud Infrastructures) - LSDVE (First Workshop on Large Scale Distributed Virtual Environments on Clouds and P2P) - MHPC (Workshop on Middleware for HPC and Big Data Systems) -PADABS ( First Workshop on Parallel and Distributed Agent Based Simulations) - ROME (First

Workshop on Runtime and Operating Systems for the Many core Era) All these workshops focus on promotion and advancement of all aspects of parallel and distributed computing.

## Handbook of Bioinspired Algorithms and Applications

Every area of science and engineering today has to process voluminous data sets. Using exact, or even approximate, algorithms to solve intractable problems in critical areas, such as computational biology, takes time that is exponential in some of the underlying parameters. Parallel computing addresses this issue and has become affordable with the advent of multicore architectures. However, programming multicore machines is much more difficult due to oddities existing in the architectures. Offering insights into different facets of this area, Multicore Computing: Algorithms, Architectures, and Applications focuses on the architectures, algorithms, and applications of multicore computing. It will help readers understand the intricacies of these architectures and prepare them to design efficient multicore algorithms. Contributors at the forefront of the field cover the memory hierarchy for multicore and manycore processors, the caching strategy Flexible Set Balancing, the main features of the latest SPARC architecture specification, the Cilk and Cilk++ programming languages, the numerical software library Parallel Linear Algebra Software for Multicore Architectures (PLASMA), and the exact multipattern string matching algorithm of Aho-Corasick. They also

describe the architecture and programming model of the NVIDIA Tesla GPU, discuss scheduling directed acyclic graphs onto multi/manycore processors, and evaluate design trade-offs among Intel and AMD multicore processors, IBM Cell Broadband Engine, and NVIDIA GPUs. In addition, the book explains how to design algorithms for the Cell Broadband Engine and how to use the backprojection algorithm for generating images from synthetic aperture radar data.

## High Performance Computing in Power and Energy Systems

High Performance Computing Systems and Applications contains fully refereed papers from the 15th Annual Symposium on High Performance Computing. These papers cover both fundamental and applied topics in HPC: parallel algorithms, distributed systems and architectures, distributed memory and performance, high level applications, tools and solvers, numerical methods and simulation, advanced computing systems, and the emerging area of computational grids. High Performance Computing Systems and Applications is suitable as a secondary text for graduate level courses, and as a reference for researchers and practitioners in industry.

## Parallel Programming

Comprehensive and concise, this handbook has chapters on computing visualization, large database designs, advanced pattern matching and other key bioinformatics techniques. It is a practical guide to computing in the growing field of Bioinformatics--the study of how information is represented and transmitted in biological systems, starting at the molecular level.

## Hands-On GPU Computing with Python

Advances in computer science and technology and in biology over the last several years have opened up the possibility for computing to help answer fundamental questions in biology and for biology to help with new approaches to computing. Making the most of the research opportunities at the interface of computing and biology requires the active participation of people from both fields. While past attempts have been made in this direction, circumstances today appear to be much more favorable for progress. To help take advantage of these opportunities, this study was requested of the NRC by the National Science Foundation, the Department of Defense, the National Institutes of Health, and the Department of Energy. The report provides the basis for establishing cross-disciplinary collaboration between biology and computing including an analysis of potential impediments and strategies for overcoming them. The report also presents a wealth of examples that should encourage students in the biological sciences to look for ways to enable them to be more effective users of computing in their

studies.

## Emerging Trends in Computational Biology, Bioinformatics, and Systems Biology

Offers a structured approach to biological data and the computer tools needed to analyze it, covering UNIX, databases, computation, Perl, data mining, data visualization, and tailoring software to suit specific research needs.

## Programming Massively Parallel Processors

Enabling technologies - An overview of cluster computing / Thomas Sterling / - Node Hardware / Thomas Sterling / - Linux / Peter H. Beckman / - Network Hardware / Thomas Sterling / - Network Software / Thomas Sterling / - Setting Up clusters : installation and configuration - How fast is my beowulf? / David Bailey / - Parallel programming / - Parallel programming with MPI / William Gropp / - Advanced topics in MPI programming / William Gropp / - Parallel programming with PVM / Al Geist / - Fault-tolerant and adaptive programs with PVM / Al Geist / - Managing clusters / - Cluster workload management / James Patton Jones / - Condor : a distributed job scheduler / - Maui scheduler : A multifunction cluster scheduler / David B. Jackson / - PBS : portable batch system / James Patton Jones / - PVFS :

parallel virtual file system / Walt Ligon / - Chiba city : the Argonne scalable cluster.

## Grid Computing for Bioinformatics and Computational Biology

Software -- Programming Techniques.

## Applications and Developments in Grid, Cloud, and High Performance Computing

High-Performance Computing using FPGA covers the area of high performance reconfigurable computing (HPRC). This book provides an overview of architectures, tools and applications for High-Performance Reconfigurable Computing (HPRC). FPGAs offer very high I/O bandwidth and fine-grained, custom and flexible parallelism and with the ever-increasing computational needs coupled with the frequency/power wall, the increasing maturity and capabilities of FPGAs, and the advent of multicore processors which has caused the acceptance of parallel computational models. The Part on architectures will introduce different FPGA-based HPC platforms: attached co-processor HPRC architectures such as the CHREC's Novo-G and EPCC's Maxwell systems; tightly coupled HRPC architectures, e.g. the Convey hybrid-core computer; reconfigurably networked HPRC architectures, e.g. the QPACE system, and standalone HPRC architectures such as

EPFL's CONFETTI system. The Part on Tools will focus on high-level programming approaches for HPRC, with chapters on C-to-Gate tools (such as Impulse-C, AutoESL, Handel-C, MORA-C++); Graphical tools (MATLAB-Simulink, NI LabVIEW); Domain-specific languages, languages for heterogeneous computing(for example OpenCL, Microsoft's Kiwi and Alchemy projects). The part on Applications will present case from several application domains where HPRC has been used successfully, such as Bioinformatics and Computational Biology; Financial Computing; Stencil computations; Information retrieval; Lattice QCD; Astrophysics simulations; Weather and climate modeling.

# Parallel Computing for Bioinformatics and Computational Biology

This book constitutes the refereed proceedings of the 7th International Conference on Parallel Computing Technologies, PaCT 2003, held in Novosibirsk, Russia in September 2003. The 38 revised full papers presented together with 4 invited papers and 10 poster papers were carefully reviewed and selected from 78 submissions. The papers are organized in topical sections on theory, software, applications, and tools. A broad variety of parallel processing issues and distributed computing in general are addressed.

# Supercomputing

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.

# Euro-Par 2014: Parallel Processing Workshops

The only single, up-to-date source for Grid issues in bioinformatics and biology Bioinformatics is fast emerging as an important discipline for academic research and industrial applications, creating a need for the use of Grid computing techniques for large-scale distributed applications. This book successfully presents Grid algorithms and their real-world applications, provides details on modern and ongoing research, and explores software frameworks that integrate bioinformatics and computational biology. Additional coverage includes: * Bio-ontology and data mining * Data visualization * DNA assembly, clustering, and mapping * Molecular

evolution and phylogeny * Gene expression and micro-arrays * Molecular modeling and simulation * Sequence search and alignment * Protein structure prediction * Grid infrastructure, middleware, and tools for bio data Grid Computing for Bioinformatics and Computational Biology is an indispensable resource for professionals in several research and development communities including bioinformatics, computational biology, Grid computing, data mining, and more. It also serves as an ideal textbook for undergraduate- and graduate-level courses in bioinformatics and Grid computing.

## Techniques for Optimizing Applications

This book constitutes the refereed proceedings of the 10th International Work-Conference on Artificial Neural Networks, IWANN 2009, held in Salamanca, Spain in June 2009. The 167 revised full papers presented together with 3 invited lectures were carefully reviewed and selected from over 230 submissions. The papers are organized in thematic sections on theoretical foundations and models; learning and adaptation; self-organizing networks, methods and applications; fuzzy systems; evolutionary computation and genetic algoritms; pattern recognition; formal languages in linguistics; agents and multi-agent on intelligent systems; brain-computer interfaces (bci); multiobjetive optimization; robotics; bioinformatics; biomedical applications; ambient assisted living (aal) and ambient intelligence (ai); other applications.

**Bioinformatics Computing**

The field of bioinformatics and computational biology arose due to the need to apply techniques from computer science, statistics, informatics, and applied mathematics to solve biological problems. Scientists have been trying to study biology at a molecular level using techniques derived from biochemistry, biophysics, and genetics. Progress has greatly accelerated with the discovery of fast and inexpensive automated DNA sequencing techniques. As the genomes of more and more organisms are sequenced and assembled, scientists are discovering many useful facts by tracing the evolution of organisms by measuring changes in their DNA, rather than through physical characteristics alone. This has led to rapid growth in the related fields of phylogenetics, the study of evolutionary relatedness among various groups of organisms, and comparative genomics, the study of the correspondence between genes and other genomic features in different organisms. Comparing the genomes of organisms has allowed researchers to better understand the features and functions of DNA in individual organisms, as well as provide insights into how organisms evolve over time. The first four chapters of Advances in Computers focus on algorithms for comparing the genomes of different organisms. Possible concrete applications include identifying the basis for genetic diseases and tracking the development and spread of different forms of Avian flu. As researchers begin to better understand the function of DNA, attention has begun shifting towards the actual proteins produced by DNA.

The final two chapters explore proteomic techniques for analyzing proteins directly to identify their presence and understand their physical structure. Written by active PhD researchers in computational biology and bioinformatics

# Petascale Computing

Discover how to streamline complex bioinformatics applications withparallel computing This publication enables readers to handle more complexbioinformatics applications and larger and richer data sets. As theeditor clearly shows, using powerful parallel computing tools canlead to significant breakthroughs in deciphering genomes,understanding genetic disease, designing customized drug therapies,and understanding evolution. A broad range of bioinformatics applications is covered withdemonstrations on how each one can be parallelized to improveperformance and gain faster rates of computation. Current parallelcomputing techniques and technologies are examined, includingdistributed computing and grid computing. Readers are provided witha mixture of algorithms, experiments, and simulations that providenot only qualitative but also quantitative insights into thedynamic field of bioinformatics. Parallel Computing for Bioinformatics and Computational Biology isa contributed work that serves as a repository of case studies,collectively demonstrating how parallel computing streamlinesdifficult problems in bioinformatics and produces better results.Each of the chapters is authored by an established expert in thefield and carefully edited

to ensure a consistent approach and highstandard throughout the publication. The work is organized into five parts: * Algorithms and models * Sequence analysis and microarrays * Phylogenetics * Protein folding * Platforms and enabling technologies Researchers, educators, and students in the field of bioinformaticswill discover how high-performance computing can enable them tohandle more complex data sets, gain deeper insights, and make newdiscoveries.

## High-Performance Computational Solutions in Protein Bioinformatics

This book constitutes the refereed post-conference proceedings of the 5th Russian Supercomputing Days, RuSCDays 2019, held in Moscow, Russia, in September 2019. The 60 revised full papers presented were carefully reviewed and selected from 127 submissions. The papers are organized in the following topical sections: parallel algorithms; supercomputer simulation; HPC, BigData, AI: architectures, technologies, tools; and distributed and cloud computing.

## Ant Colony Optimization and Swarm Intelligence

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

## Deep Learning and Parallel Computing Environment for Bioengineering Systems

In the post-genomic era, a holistic understanding of biological systems and p-cesses,inalltheircomplexity,is criticalincomprehendingnature'schoreography of life. As a result, bioinformatics involving its two main disciplines, namely, the life sciences and the computational sciences, is fast becoming a very promising multidisciplinary research ?eld. With the ever-increasing application of lar-scalehigh-throughputtechnologies,suchasgeneorproteinmicroarraysandmass spectrometry methods, the enormous body of information is growing rapidly.

Bioinformaticians are posed with a large number of di?cult problems to solve, arising not only due to the complexities in acquiring the molecular infor- tion but also due to the size and nature of the generated data sets and/or the limitations of the algorithms required for analyzing these data. Although the ?eld of bioinformatics is still in its embryonic stage, the recent advancements in computational and information-theoretic techniques are enabling us to c-ductvariousinsilicotestingandscreeningofmanylab-basedexperimentsbefore these are actually performed in vitro or in vivo. These in silico investigations are providing new insights for interpretation and establishing a new direction for a deeper understanding. Among the various advanced computational methods currently being applied to such studies, the pattern recognition techniques are mostly found to be at the core of the whole discovery process for apprehending the underlying biological knowledge. Thus, we can safely surmise that the - going bioinformatics revolution may, in future, inevitably play a major role in many aspects of medical practice and/or the discipline of life sciences.

## Bioinformatics

Programming Massively Parallel Processors discusses the basic concepts of parallel programming and GPU architecture. Various techniques for constructing parallel programs are explored in detail. Case studies demonstrate the development process, which begins with computational thinking and ends with effective and

efficient parallel programs. This book describes computational thinking techniques that will enable students to think about problems in ways that are amenable to high-performance parallel computing. It utilizes CUDA (Compute Unified Device Architecture), NVIDIA's software development tool created specifically for massively parallel environments. Studies learn how to achieve both high-performance and high-reliability using the CUDA programming model as well as OpenCL. This book is recommended for advanced students, software engineers, programmers, and hardware engineers. Teaches computational thinking and problem-solving techniques that facilitate high-performance parallel computing. Utilizes CUDA (Compute Unified Device Architecture), NVIDIA's software development tool created specifically for massively parallel environments. Shows you how to achieve both high-performance and high-reliability using the CUDA programming model as well as OpenCL.

## Multicore Computing

The two volumes LNCS 8805 and 8806 constitute the thoroughly refereed post-conference proceedings of 18 workshops held at the 20th International Conference on Parallel Computing, Euro-Par 2014, in Porto, Portugal, in August 2014. The 100 revised full papers presented were carefully reviewed and selected from 173 submissions. The volumes include papers from the following workshops: APCI&E (First Workshop on Applications of Parallel Computation in Industry and

Engineering - BigDataCloud (Third Workshop on Big Data Management in Clouds) - DIHC (Second Workshop on Dependability and Interoperability in Heterogeneous Clouds) - FedICI (Second Workshop on Federative and Interoperable Cloud Infrastructures) - Hetero Par (12th International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms) - HiBB (5th Workshop on High Performance Bioinformatics and Biomedicine) - LSDVE (Second Workshop on Large Scale Distributed Virtual Environments on Clouds and P2P) - MuCoCoS (7th International Workshop on Multi-/Many-core Computing Systems) - OMHI (Third Workshop on On-chip Memory Hierarchies and Interconnects) - PADAPS (Second Workshop on Parallel and Distributed Agent-Based Simulations) - PROPER (7th Workshop on Productivity and Performance) - Resilience (7th Workshop on Resiliency in High Performance Computing with Clusters, Clouds, and Grids) - REPPAR (First International Workshop on Reproducibility in Parallel Computing) - ROME (Second Workshop on Runtime and Operating Systems for the Many Core Era) - SPPEXA (Workshop on Software for Exascale Computing) - TASUS (First Workshop on Techniques and Applications for Sustainable Ultrascale Computing Systems) - UCHPC (7th Workshop on Un Conventional High Performance Computing) and VHPC (9th Workshop on Virtualization in High-Performance Cloud Computing.

# **Advances in High Performance Computing**

Emerging Trends in Computational Biology, Bioinformatics, and Systems Biology discusses the latest developments in all aspects of computational biology, bioinformatics, and systems biology and the application of data-analytics and algorithms, mathematical modeling, and simu- lation techniques. • Discusses the development and application of data-analytical and theoretical methods, mathematical modeling, and computational simulation techniques to the study of biological and behavioral systems, including applications in cancer research, computational intelligence and drug design, high-performance computing, and biology, as well as cloud and grid computing for the storage and access of big data sets. • Presents a systematic approach for storing, retrieving, organizing, and analyzing biological data using software tools with applications to general principles of DNA/RNA structure, bioinformatics and applications, genomes, protein structure, and modeling and classification, as well as microarray analysis. • Provides a systems biology perspective, including general guidelines and techniques for obtaining, integrating, and analyzing complex data sets from multiple experimental sources using computational tools and software. Topics covered include phenomics, genomics, epigenomics/epigenetics, metabolomics, cell cycle and checkpoint control, and systems biology and vaccination research. • Explains how to effectively harness the power of Big Data tools when data sets are so large and complex that it is difficult to process them using conventional database management systems or traditional data processing applications. Discusses the development and application of data-analytical and theoretical

methods, mathematical modeling and computational simulation techniques to the study of biological and behavioral systems. Presents a systematic approach for storing, retrieving, organizing and analyzing biological data using software tools with applications. Provides a systems biology perspective including general guidelines and techniques for obtaining, integrating and analyzing complex data sets from multiple experimental sources using computational tools and software.

## High-Performance Modelling and Simulation for Big Data Applications

Most of the papers in this volume were presented at the NATO Advanced Research Workshop High Performance Computing: Technology and Application, held in Cetraro, Italy from 24 to 26 of June, 1996. The main purpose of the Workshop was to discuss some key scientific and technological developments in high performance computing, identify significant trends and defme desirable research objectives. The volume structure corresponds, in general, to the outline of the workshop technical agenda: general concepts and emerging systems, software technology, algorithms and applications. One of the Workshop innovations was an effort to extend slightly the scope of the meeting from scientific/engineering computing to enterprise-wide computing. The papers on performance and scalability of database servers, and Oracle DBMS reflect this attempt We hope that after reading this collection of

papers the readers will have a good idea about some important research and technological issues in high performance computing. We wish to give our thanks to the NATO Scientific and Environmental Affairs Division for being the principal sponsor for the Workshop. Also we are pleased to acknowledge other institutions and companies that supported the Workshop: European Union: European Commission DGIII-Industry, CNR: National Research Council of Italy, University of Calabria, Alenia Spazio, Centro Italiano Ricerche Aerospaziali, ENEA: Italian National Agency for New Technology, Energy and the Environment, Fujitsu, Hewlett Packard-Convex, Hitachi, NEC, Oracle, and Silicon Graphics-Cray Research. Editors January 1997 vii LIST OF CONTRIBUTORS Ecole Nonnale Su¢rieure de Lyon, 69364 Abarbanel. Robert M.

## Parallel Computing Technologies

This book constitutes the refereed proceedings of the 5th International Workshop on Ant Colony Optimization and Swarm Intelligence, ANTS 2006, held in Brussels, Belgium, in September 2006. The 27 revised full papers, 23 revised short papers, and 12 extended abstracts presented were carefully reviewed and selected from 115 submissions.

## VLSI

During the last three decades, reconfigurable logic has been growing steadily and can now be found in many different fields. Field programmable gate arrays (FPGAs) are one of the most famous architecture families of reconfigurable devices. FPGAs can be seen as arrays of logic units that can be reconfigured to realize any digital systems. Their high versatility has enabled designers to drastically reduce time to market, and made FPGAs suitable for prototyping or small production series in many branches of industrial products. In addition, and thanks to innovations at the architecture level, FPGAs are now conquering segments of mass markets such as mobile communications. Reconfigurable Logic: Architecture, Tools, and Applications offers a snapshot of the state of the art of reconfigurable logic systems. Covering a broad range of architectures, tools, and applications, this book: Explores classical FPGA architectures and their supporting tools Evaluates recent proposals related to FPGA architectures, including the use of network-on-chips (NoCs) Examines reconfigurable processors that merge concepts borrowed from the reconfigurable domain into processor design Exploits FPGAs for high-performance systems, efficient error correction codes, and high-bandwidth network routers with built-in security Expounds on emerging technologies to enhance FPGA architectures, improve routing structures, and create non-volatile configuration flip-flops Reconfigurable Logic: Architecture, Tools, and Applications reviews current trends in reconfigurable platforms, providing valuable insight into the future potential of reconfigurable systems.

ROMANCE  ACTION & ADVENTURE  MYSTERY & THRILLER  BIOGRAPHIES & HISTORY  CHILDREN'S  YOUNG ADULT  FANTASY  HISTORICAL FICTION  HORROR  LITERARY FICTION  NON-FICTION  SCIENCE FICTION