Hpca 18 Call For Papers

Our world is being revolutionized by data-driven methods: access to large amounts of data has generated new insights and opened exciting new opportunities in commerce, science, and computing applications. Processing the enormous quantities of data necessary for these advances requires large clusters, making distributed computing paradigms more crucial than ever. MapReduce is a programming model for expressing distributed computations on massive datasets and an execution framework for large-scale data processing on clusters of commodity servers. The programming model provides an easy-to-understand abstraction for designing scalable algorithms, while the execution framework transparently handles many system-level details, ranging from scheduling to synchronization to fault tolerance. This book focuses on MapReduce algorithm design, with an emphasis on text processing algorithms common in natural language processing, information retrieval, and machine learning. We introduce the notion of MapReduce design patterns, which represent general reusable solutions to commonly occurring problems across a variety of problem domains. This book not only intends to help the reader "think in MapReduce", but also discusses limitations of the programming model as well. This volume is a printed version of a work that appears in the Synthesis Digital Library of Engineering and Computer Science. Synthesis Lectures provide concise, original presentations of important research and development topics, published quickly, in digital and print formats. For more information visit www.morganclaypool.com

This text on parallel architectures and compilation techniques covers such topics as: loop transformations; shared memory design techniques; specialized multiprocessor systems; parallel programming languages; JAVA and multithreading processors; register allocation; and branch and value prediction."

Quantum mechanics, the subfield of physics that describes the behavior of very small (quantum) particles, provides the basis for a new paradigm of computing. First proposed in the 1980s as a way to improve computational modeling of quantum systems, the field of quantum computing has recently garnered significant attention due to progress in building small-scale devices. However, significant technical advances will be required before a large-scale, practical quantum computer can be achieved. Quantum Computing: Progress and Prospects provides an introduction to the field, including the unique characteristics and constraints of the technology, and assesses the feasibility and implications of creating a functional quantum computer capable of addressing real-world problems. This report considers hardware and software requirements, quantum algorithms, drivers of advances in quantum computing and quantum devices, benchmarks associated with relevant use cases, the time and resources required, and how to assess the probability of success.

This book constitutes the refereed proceedings of the 22nd International Conference on Architecture of Computing Systems, ARCS 2009, held in Delft, The Netherlands, in March 2009. The 21 revised full papers presented together with 3 keynote papers were carefully reviewed and selected from 57 submissions. This year's special focus is set on energy awareness. The papers are organized in topical sections on compilation technologies, reconfigurable hardware and applications, massive parallel architectures, organic computing, memory architectures, enery awareness, Java processing, and chip-level multiprocessing.

Is your memory hierarchy stopping your microprocessor from performing at the high level it should be? Memory Systems: Cache, DRAM, Disk shows you how to resolve this problem. The book tells you everything you need to know about the logical design and operation, physical design and operation, performance characteristics and resulting design trade-offs, and the energy consumption of modern memory hierarchies. You learn how to to tackle the challenging optimization problems that result from the side-effects that can appear at any point in the entire hierarchy. As a result you will be able to design and emulate the entire memory hierarchy. Understand all levels of the system hierarchy -Xcache, DRAM, and disk. Evaluate the system-level effects of all design choices. Model performance and energy consumption for each component in the memory hierarchy.

Computer Architecture

Many modern computer systems, including homogeneous and heterogeneous architectures, support shared memory in hardware. In a shared memory system, each of the processor cores may read and write to a single shared address space. For a shared memory machine, the memory consistency model defines the architecturally visible behavior of its memory system. Consistency definitions provide rules about loads and stores (or memory reads and writes) and how they act upon memory. As part of supporting a memory consistency model, many machines also provide cache coherence protocols that ensure that multiple cached copies of data are kept up-to-date. The goal of this primer is to provide readers with a basic understanding of consistency and coherence. This understanding includes both the issues that must be solved as well as a variety of solutions. We present both high-level concepts as well as specific, concrete examples from real-world systems. This second edition reflects a decade of advancements since the first edition and includes, among other more modest changes, two new chapters: one on consistency and coherence for non-CPU accelerators (with a focus on GPUs) and one that points to formal work and tools on consistency and coherence.

This book constitutes the refereed proceedings of the 36th International Conference on High Performance Computing, ISC High Performance 2021, held virtually in June/July 2021. The 24 full papers presented were carefully reviewed and selected from 74 submissions. The papers cover a broad range of topics such as architecture, networks, and storage; machine learning, AI, and emerging technologies; HPC algorithms and applications; performance modeling, evaluation, and analysis; and programming environments and systems software.

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.

Distributed and Cloud Computing: From Parallel Processing to the Internet of Things offers complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing. It is the first modern, up-to-date distributed systems textbook; it explains how to create high-performance, scalable, reliable systems, exposing the design principles, architecture, and innovative applications of parallel, distributed, and cloud computing systems. Topics covered by this book include: facilitating management, debugging, migration, and disaster recovery through virtualization; clustered systems for research or ecommerce applications; designing systems as web services; and social networking systems using peer-to-peer computing. The principles of cloud computing are discussed using examples from open-source and commercial applications, along with case studies from the leading distributed computing vendors such as Amazon, Microsoft, and Google. Each chapter includes exercises and further reading, with lecture slides and more available online. This book will be ideal for students taking a distributed systems or distributed computing class, as well as for professional system designers and engineers looking for a reference to the latest distributed technologies including cloud, P2P and grid computing. Complete coverage of modern distributed computing technology including clusters, the grid, service-oriented architecture, massively parallel processors, peer-to-peer networking, and cloud computing Includes case studies from the leading distributed computing vendors: Amazon, Microsoft, Google, and more Explains how to use virtualization to facilitate management, debugging, migration, and disaster recovery Designed for undergraduate or graduate students taking a distributed systems course—each chapter includes exercises and further reading, with lecture slides and more available online

HPCAThe Seventh International Symposium on High-Performance Computer Architecture : Proceedings, 19-24 January 2001, Monterrey, Nuevo Leon, MéxicoIEEE

A survey of architectural mechanisms and implementation techniques for exploiting fine- and coarse-grained parallelism within microprocessors. Beginning with a review of past techniques, the monograph provides a comprehensive account of state-of-the-art techniques used in microprocessors, covering both the concepts involved and implementations in sample processors. The whole is rounded off with a thorough review of the research techniques that will lead to future microprocessors. XXXXXXX Neuer Text This monograph surveys architectural mechanisms and implementation techniques for exploiting fine-grained and coarse-grained parallelism within microprocessors. It presents a comprehensive account of state-of-the-art techniques used in microprocessors that covers both the concepts involved and possible implementations. The authors also provide application-oriented methods and a thorough review of the research techniques that will lead to the development of future processors.

Since its first volume in 1960, Advances in Computers has presented detailed coverage of innovations in hardware and software and in computer theory, design, and applications. It has also provided contributors with a medium in which they can examine their subjects in greater depth and breadth than that allowed by standard journal articles. As a result, many articles have become standard references that continue to be of significant, lasting value despite the rapid growth taking place in the field. This volume is organized around engineering large scale software systems. It discusses which technologies are useful for building these systems, which are useful to incorporate in these systems, and which are useful to evaluate these systems.

at the distributed virtual Program Committee meeting. Each paper's review recomm- dationswere carefully checkedfor consistency; in many instances, the Vice Chairs read the papers themselves when the reviews did not seem suf?cient to make a decision. Throughout the reviewing process, I received a tremendous amount of help and advice from General Co-chair Manish Parashar, Steering Chair Viktor Prasanna, and last year's Program Chair Srinivas Aluru; I am very grateful to them. My thanks also go to the Publications Chair Sushil Prasad for his outstanding efforts in putting the proceedings together. Finally, I thank all the authors for their contributions to a hi- quality technical program. I wish all the attendees a very enjoyable and informative meeting. December 2008 P. Sadayappan Message from the General Co-chairs and the Vice General Co-chairs On behalf of the organizers of the 15th International Conference on High-Performance Computing(HiPC), it is our pleasureto present these proceedingsand we hopeyou will ?nd them exciting and rewarding. TheHiPCcallforpapers,onceagain,receivedanoverwhelmingresponse,attracting 317submissionsfrom27countries.P.Sadayappan,theProgramChair,andthe Program Committee worked with remarkablededication to put together an outstandingtechnical program consisting of the 46 papers that appear in these proceedings.

The computing world today is in the middle of a revolution: mobile clients and cloud computing have emerged as the dominant paradigms driving programming and hardware innovation today. The Fifth Edition of Computer Architecture focuses on this dramatic shift, exploring the ways in which software and technology in the cloud are accessed by cell phones, tablets, laptops, and other mobile computing devices. Each chapter includes two real-world examples, one mobile and one datacenter, to illustrate this revolutionary change. Updated to cover the mobile computing revolution Emphasizes the two most important topics in architecture today: memory hierarchy and parallelism in all its forms.

Develops common themes throughout each chapter: power, performance, cost, dependability, protection, programming models, and emerging trends ("What's Next") Includes three review appendices in the printed text. Additional reference appendices are available online. Includes updated Case Studies and completely new exercises.

Most emerging applications in imaging and machine learning must perform immense amounts of computation while holding to strict limits on energy and power. To meet these goals, architects are building increasingly specialized compute engines tailored for these specific tasks. The resulting computer systems are heterogeneous, containing multiple processing cores with wildly different execution models. Unfortunately, the cost of producing this specialized hardware—and the software to control it—is astronomical. Moreover, the task of porting algorithms to these heterogeneous machines typically requires that the algorithm be partitioned across the machine and rewritten for each specific architecture, which is time consuming and prone to error. Over the last several years, the authors have approached this problem using domain-specific languages (DSLs): high-level programming languages customized for specific domains, such as database manipulation, machine learning, or image processing. By giving up generality, these languages are able to provide high-level abstractions to the developer while producing high performance output. The purpose of this book is to spur the adoption and the creation of domain-specific languages, especially for the task of creating hardware designs. In the first chapter, a short historical journey explains the forces driving computer architecture today. Chapter 2 describes the various methods for producing designs for accelerators, outlining the push for more abstraction and the tools that enable designers to work at a higher conceptual level. From there, Chapter 3 provides a brief introduction to image processing algorithms and hardware design patterns for implementing them. Chapters 4 and 5 describe and compare Darkroom and Halide, two domain-specific languages created for image processing that produce high-performance designs for both FPGAs and CPUs from the same source code, enabling rapid design cycles and quick porting of algorithms. The final section describes how the DSL approach also simplifies the problem of interfacing between application code and the accelerator by generating the driver stack in addition to the accelerator configuration. This book should serve as a useful introduction to domain-specialized computing for computer architecture students and as a primer on domain-specific languages and image processing hardware for those with more experience in the field.

Topics covered in this text include: microarchitecture; memory architectures; multiprocessor systems; code generation techniques; energy and thermal management; prediction techniques; application-specific designs; performance modelling and analysis; and latency tolerance techniques.

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

This book constitutes the thoroughly refereed post-proceedings of the Second International Workshop on Power-Aware Computer Systems, PACS 2002, held in Cambridge, MA, USA, in February 2002. The 13 revised full papers presented were carefully selected for inclusion in the book during two rounds of reviewing and revision. The papers are organized in topical sections on power-aware architecture and microarchitecture, power-aware real-time systems, power modeling and monitoring, and power-aware operating systems and compilers.

The complexity of modern computer networks and systems, combined with the extremely dynamic environments in which they operate, is beginning to outpace our ability to manage them. Taking yet another page from the biomimetics playbook, the autonomic computing paradigm mimics the human autonomic nervous system to free system developers and administrators from performing and overseeing low-level tasks. Surveying the current path toward this paradigm, Autonomic Computing: Concepts, Infrastructure, and Applications offers a comprehensive overview of state-of-the-art research and implementations in this emerging area. This book begins by introducing the concepts and requirements of autonomic computing and exploring the architectures required to implement such a system. The focus then shifts to the approaches and infrastructures, including control-based and recipe-based concepts, followed by enabling systems, technologies, and services proposed for achieving a set of "self-*" properties, including self-configuration, self-healing, self-optimization, and self-protection. In the final section, examples of real-world implementations reflect the potential of emerging autonomic systems, such as dynamic server allocation and runtime reconfiguration and repair. Collecting cutting-edge work and perspectives from leading experts, Autonomic Computing: Concepts, Infrastructure, and Applications reveals the progress made and outlines the future challenges still facing this exciting and dynamic field.

This book constitutes the thoroughly refereed post conference papers of the First International Conference on Blockchain and Trustworthy Systems, Blocksys 2019, held in Guangzhou, China, in December 2019. The 50 regular papers and the 19 short papers were carefully reviewed and selected from 130 submissions. The papers are focus on Blockchain and trustworthy systems can be applied to many fields, such as financial services, social management and supply chain management.

This book constitutes the thoroughly refereed post-conference proceedings of the Second International Conference on High Performance Computing and Applications, HPCA 2009, held in Shangahi, China, in August 2009. The 71 revised papers presented together with 10 invited presentations were carefully selected from 324 submissions. The papers cover topics such as numerical algorithms and solutions; high performance and grid computing; novel approaches to high performance computing; massive data storage and processsing; and hardware acceleration.

This book constitutes the thoroughly refereed post-conference proceedings of the 13th International Conference on Information Security and Cryptology, Inscrypt 2017, held in Xi'an, China, in November 2017. The 27 revised full papers presented together with 5 keynote speeches were carefully reviewed and selected from 80 submissions. The papers are organized in the following topical sections: cryptographic protocols and algorithms; digital signatures; encryption; cryptanalysis and attack; and applications.

Computer Architecture: A Quantitative Approach, Sixth Edition has been considered essential reading by instructors, students and practitioners of computer design for over 20 years. The sixth edition of this classic textbook from Hennessy and Patterson, winners of the 2017 ACM A.M. Turing Award recognizing contributions of lasting and major technical importance to the computing field, is fully revised with the latest developments in processor and system architecture. The text now features examples from the RISC-V (RISC Five) instruction set architecture, a modern RISC instruction set developed and designed to

be a free and openly adoptable standard. It also includes a new chapter on domain-specific architectures and an updated chapter on warehouse-scale computing that features the first public information on Google's newest WSC. True to its original mission of demystifying computer architecture, this edition continues the longstanding tradition of focusing on areas where the most exciting computing innovation is happening, while always keeping an emphasis on good engineering design. Winner of a 2019 Textbook Excellence Award (Texty) from the Textbook and Academic Authors Association Includes a new chapter on domain-specific architectures, explaining how they are the only path forward for improved performance and energy efficiency given the end of Moore’s Law and Dennard scaling Features the first publication of several DSAs from industry Features extensive updates to the chapter on warehouse-scale computing, with the first public information on the newest Google WSC Offers updates to other chapters including new material dealing with the use of stacked DRAM; data on the performance of new NVIDIA Pascal GPU vs. new AVX-512 Intel Skylake CPU; and extensive additions to content covering multicore architecture and organization Includes "Putting It All Together" sections near the end of every chapter, providing real-world technology examples that demonstrate the principles covered in each chapter Includes review appendices in the printed text and additional reference appendices available online Includes updated and improved case studies and exercises ACM named John L. Hennessy and David A. Patterson, recipients of the 2017 ACM A.M. Turing Award for pioneering a systematic, quantitative approach to the design and evaluation of computer architectures with enduring impact on the microprocessor industry

Children's palliative care has developed rapidly as a discipline, as health care professionals recognise that the principles of adult palliative care may not always be applicable to children at the end of life. The unique needs of dying children are particularly evident across Africa, where the scale of the problem is overwhelming and the figures so enormous that they are barely comprehensible. Written by a group with wide experience of caring for dying childrenin Africa, this book provides practical, realistic guidance on improving access to, and delivery of, palliative care in this demanding setting. It looks at the themes common to palliative care - including communication, assessment, symptom management, psychosocial issues, ethical dilemmas, end of lifecare, and tips for the professional on compassion and conservation of energy - but always retains the focus on the particular needs of the health care professional in Africa. Whilst containing some theory, the emphasis is on practical action throughout. It will provide health care professonals working in Africa, and other resource-poor settings, with the confidence, knowledge, and capacity to improve care for the terminally ill child in constrained and demanding environments.

Computing and information and communications technology (ICT) has dramatically changed how we work and live, has had profound effects on nearly every sector of society, has transformed whole industries, and is a key component of U.S. global leadership. A fundamental driver of advances in computing and ICT has been the fact that the single-processor performance has, until recently, been steadily and dramatically increasing year over years, based on a combination of architectural techniques, semiconductor advances, and software improvements. Users, developers, and innovators were able to depend on those increases, translating that performance into numerous technological innovations and creating successive generations of ever more rich and diverse products, software services, and applications that had profound effects across all sectors of society. However, we can no longer depend on those extraordinary advances in single-processor performance continuing. This slowdown in the growth of single-processor computing performance has its roots in fundamental physics and engineering constraints--multiple technological barriers have converged to pose deep research challenges, and the consequences of this shift are deep and profound for computing and for the sectors of the economy that depend on and assume, implicitly or explicitly, ever-increasing performance. From a technology standpoint, these challenges have led to heterogeneous multicore chips and a shift to alternate innovation axes that include, but are not limited to, improving chip performance, mobile devices, and cloud services. As these technical shifts reshape the computing industry, with global consequences, the United States must be prepared to exploit new opportunities and to deal with technical challenges. The New Global Ecosystem in Advanced Computing: Implications for U.S. Competitiveness and National Security outlines the technical challenges, describe the global research landscape, and explore implications for competition and national security.

A vision of the future of education in which the classroom experience is distributed across space and time without compromising learning. What if there were a model for learning in which the classroom experience was distributed across space and time--and students could still have the benefits of the traditional classroom, even if they can't be present physically or learn synchronously? In this book, two experts in online learning envision a future in which education from kindergarten through graduate school need not be tethered to a single physical classroom. The distributed classroom would neither sacrifice students' social learning experience nor require massive development resources. It goes beyond hybrid learning, so ubiquitous during the COVID-19 pandemic, and MOOCs, so trendy a few years ago, to reimagine the classroom itself. David Joyner and Charles Isbell, both of Georgia Tech, explain how recent developments, including distance learning and learning management systems, have paved the way for the distributed classroom. They propose that we dispense with the dichotomy between online and traditional education, and the assumption that online learning is necessarily inferior. They describe the distributed classroom's various delivery modes for in-person students, remote synchronous students, and remote asynchronous students; the goal would be a symmetry of experiences, with both students and teachers able to move from one mode to another. With The Distributed Classroom, Joyner and Isbell offer an optimistic, learner-centric view of the future of education, in which every person on earth is turned into a potential learner as barriers of cost, geography, and synchronicity disappear.

This book constitutes the thoroughly refereed post-conference proceedings of the workshops held at the 37th International Symposium on Computer Architecture, ISCA 2010, in Saint-Malo, France, in June 2010. The 28 revised full papers presented were carefully reviewed and selected from the lectures given at 5 of these workshops. The papers address topics ranging from novel memory architectures to emerging application design and performance analysis and encompassed the following workshops: A4MMC, applications for multi- and many-cores; AMAS-BT, 3rd workshop on architectural and micro-architectural support for binary translation; EAMA, the 3rd Workshop for emerging

applications and many-core architectures; WEED, 2nd Workshop on energy efficient design, as well as WIOSCA, the annual workshop on the interaction between operating systems and computer architecture.

This book constitutes the refereed proceedings of the 19th International Conference on CParallel and Distributed Computing, Applications and Technologies, PDCAT 2018, held in Jeju Island, South Korea, in August 2018. The 35 revised full papers presented along with the 14 short papers and were carefully reviewed and selected from 150 submissions. The papers of this volume are organized in topical sections on wired and wireless communication systems, high dimensional data representation and processing, networks and information security, computing techniques for efficient networks design, electronic circuits for communication systems.

This book constitutes the refereed proceedings of the Second International Symposium on Benchmarking, Measuring, and Optimization, Bench 2019, held in Denver, CO, USA, in November 2019. The 20 full papers and 11 short papers presented were carefully reviewed and selected from 79 submissions. The papers are organized in topical sections named: Best Paper Session; AI Challenges on Cambircon using AIBenc; AI Challenges on RISC-V using AIBench; AI Challenges on X86 using AIBench; AI Challenges on 3D Face Recognition using AIBench; Benchmark; AI and Edge; Big Data; Datacenter; Performance Analysis; Scientific Computing.

Originally developed to support video games, graphics processor units (GPUs) are now increasingly used for general-purpose (non-graphics) applications ranging from machine learning to mining of cryptographic currencies. GPUs can achieve improved performance and efficiency versus central processing units (CPUs) by dedicating a larger fraction of hardware resources to computation. In addition, their general-purpose programmability makes contemporary GPUs appealing to software developers in comparison to domain-specific accelerators. This book provides an introduction to those interested in studying the architecture of GPUs that support general-purpose computing. It collects together information currently only found among a wide range of disparate sources. The authors led development of the GPGPU-Sim simulator widely used in academic research on GPU architectures. The first chapter of this book describes the basic hardware structure of GPUs and provides a brief overview of their history. Chapter 2 provides a summary of GPU programming models relevant to the rest of the book. Chapter 3 explores the architecture of GPU compute cores. Chapter 4 explores the architecture of the GPU memory system. After describing the architecture of existing systems, Chapters \ref{ch03} and \ref{ch04} provide an overview of related research. Chapter 5 summarizes cross-cutting research impacting both the compute core and memory system. This book should provide a valuable resource for those wishing to understand the architecture of graphics processor units (GPUs) used for acceleration of general-purpose applications and to those who want to obtain an introduction to the rapidly growing body of research exploring how to improve the architecture of these GPUs.

This book constitutes the refereed proceedings of the 35th International Conference on High Performance Computing, ISC High Performance 2020, held in Frankfurt/Main, Germany, in June 2020.* The 27 revised full papers presented were carefully reviewed and selected from 87 submissions. The papers cover a broad range of topics such as architectures, networks & infrastructure; artificial intelligence and machine learning; data, storage & visualization; emerging technologies; HPC algorithms; HPC applications; performance modeling & measurement; programming models & systems software. *The conference was held virtually due to the COVID-19 pandemic. Chapters "Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) Streaming-Aggregation Hardware Design and Evaluation", "Solving Acoustic Boundary Integral Equations Using High Performance Tile Low-Rank LU Factorization", "Scaling Genomics Data Processing with Memory-Driven Computing to Accelerate Computational Biology", "Footprint-Aware Power Capping for Hybrid Memory Based Systems", and "Pattern-Aware Staging for Hybrid Memory Systems" are available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

This book constitutes the proceedings of the 23rd International Conference on Parallel and Distributed Computing, Euro-Par 2017, held in Santiago de Compostela, Spain, in August/September 2017. The 50 revised full papers presented together with 2 abstract of invited talks and 1 invited paper were carefully reviewed and selected from 176 submissions. The papers are organized in the following topical sections: support tools and environments; performance and power modeling, prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; parallel and distributed data management and analytics; cluster and cloud computing; distributed systems and algorithms; parallel and distributed programming, interfaces and languages; multicore and manycore parallelism; theory and algorithms for parallel computation and networking; prallel numerical methods and applications; and accelerator computing.

The two-volume set LNCS 11944-11945 constitutes the proceedings of the 19th International Conference on Algorithms and Architectures for Parallel Processing, ICA3PP 2019, held in Melbourne, Australia, in December 2019. The 73 full and 29 short papers presented were carefully reviewed and selected from 251 submissions. The papers are organized in topical sections on: Parallel and Distributed Architectures, Software Systems and Programming Models, Distributed and Parallel and Network-based Computing, Big Data and its Applications, Distributed and Parallel Algorithms, Applications of Distributed and Parallel Computing, Service Dependability and Security, IoT and CPS Computing, Performance Modelling and Evaluation.

With growing interest in computer security and the protection of the code and data which execute on commodity computers, the amount of hardware security features in today's processors has increased significantly over the recent years. No longer of just academic interest, security features inside processors have been embraced by industry as well, with a number of commercial secure processor architectures available today. This book aims to give readers insights into the principles behind the design of academic and

commercial secure processor architectures. Secure processor architecture research is concerned with exploring and designing hardware features inside computer processors, features which can help protect confidentiality and integrity of the code and data executing on the processor. Unlike traditional processor architecture research that focuses on performance, efficiency, and energy as the first-order design objectives, secure processor architecture design has security as the first-order design objective (while still keeping the others as important design aspects that need to be considered). This book aims to present the different challenges of secure processor architecture design to graduate students interested in research on architecture and hardware security and computer architects working in industry interested in adding security features to their designs. It aims to educate readers about how the different challenges have been solved in the past and what are the best practices, i.e., the principles, for design of new secure processor architectures. Based on the careful review of past work by many computer architects and security researchers, readers also will come to know the five basic principles needed for secure processor architecture design. The book also presents existing research challenges and potential new research directions. Finally, this book presents numerous design suggestions, as well as discusses pitfalls and fallacies that designers should avoid.

This title gives students an integrated and rigorous picture of applied computer science, as it comes to play in the construction of a simple yet powerful computer system.

This collection of case studies contains contributions illustrating the application of formal methods to real-life problems with industrial relevance.