

HPCA 18 Call For Papers

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.

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.

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.

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.

Index of Conference Proceedings

Fundamental Approaches to Software Engineering

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

C++ Parallel Programming with Threading Building Blocks

A Holistic Design Exploration

Euro-Par 2007 Workshops: Parallel Processing

Emerging Memory Technologies

This book explores the design implications of emerging, non-volatile memory (NVM) technologies on future computer memory hierarchy architecture designs. Since NVM technologies combine the speed of SRAM, the density of DRAM, and the non-volatility of Flash memory, they are very attractive as the basis for future universal memories. This book provides a holistic perspective on the topic, covering modeling, design, architecture and applications. The practical information included in this book will enable designers to exploit emerging memory technologies to improve significantly the performance/power/reliability of future, mainstream integrated circuits.

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.

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. XXXXXX 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.

The implementation of object-oriented languages has been an active topic of research since the 1960s when the first Simula compiler was written. The topic received renewed interest in the early 1980s with the growing popularity of object-oriented programming languages such as c++ and Smalltalk, and got another boost with the advent of Java.

Polymorphic calls are at the heart of object-oriented languages, and even the first implementation of Simula-67 contained their classic implementation via virtual function tables. In fact, virtual function tables predate even Simula-for example, Ivan Sutherland's Sketchpad drawing editor employed very similar structures in 1960. Similarly, during the 1970s and 1980s the implementers of Smalltalk systems spent considerable efforts on implementing polymorphic calls for this dynamically typed language where virtual function tables could not be used. Given this long history of research into the implementation of polymorphic calls, and the relatively mature standing it achieved over time, why, one might ask, should there be a new book in this field? The answer is simple. Both software and hardware have changed considerably in recent years, to the point where many assumptions underlying the original work in this field are no longer true. In particular, virtual function tables are no longer sufficient to implement polymorphic calls even for statically typed languages; for example, Java's interface calls cannot be implemented this way. Furthermore, today's processors are deeply pipelined and can execute instructions out-of order, making it difficult to predict the execution time of even simple code sequences.

40th IFIP WG 6.1 International Conference, FORTE 2020, Held as Part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta, June 15 – 19, 2020, Proceedings

Pro TBB

13th International Conference, Inscrypt 2017, Xi'an, China, November 3 – 5, 2017, Revised Selected Papers

35th International Conference, ISC High Performance 2020, Frankfurt/Main, Germany, June 22 – 25, 2020, Proceedings

Processor Architecture

Third International Conference, BlockSys 2021, Guangzhou, China, August 5 – 6, 2021, Revised Selected Papers

Efficient Processing of Deep Neural Networks

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 processing; and hardware acceleration.

This book constitutes the strictly refereed post-workshop proceedings of the ACM SIGPLAN Workshop on Languages, Compilers, and Tools for Embedded Systems, LCTES '98, held in Montreal, Canada, in June 1998. The 19 revised papers presented were carefully reviewed and selected from a total of 54 submissions for inclusion in the book; also included are one full paper and an abstract of an invited contribution. The papers address all current aspects of research and development in the rapidly growing area of embedded systems and real-time computing.

From driving, flying, and swimming, to digging for unknown objects in space exploration, autonomous robots take on varied shapes and sizes. In part, autonomous robots are designed to perform tasks that are too dirty, dull, or dangerous for humans. With nontrivial autonomy and volition, they may soon claim their own place in human society. These robots will be our allies as we strive for understanding our natural and man-made environments and build positive synergies around us. Although we may never perfect replication of biological capabilities in robots, we must harness the inevitable emergence of robots that synchronizes with our own capacities to live, learn, and grow. This book is a snapshot of motivations and methodologies for our collective attempts to transform our lives and enable us to cohabit with robots that work with and for us. It reviews and guides the reader to seminal and continual developments that are the foundations for successful paradigms. It attempts to demystify the abilities and limitations of robots. It is a progress report on the continuing work that will fuel future endeavors. Table of Contents: Part I: Preliminaries/Agency, Motion, and Anatomy/Behaviors / Architectures / Affect/Sensors / Manipulators/Part II: Mobility/Potential Fields/Roadmaps / Reactive Navigation / Multi-Robot Mapping: Brick and Mortar Strategy / Part III: State of the Art / Multi-Robotics Phenomena / Human-Robot Interaction / Fuzzy Control / Decision Theory and Game Theory / Part IV: On the Horizon / Applications: Macro and Micro Robots / References / Author Biography / Discussion

This book constitutes the proceedings of the 8th International Conference on Future Data and Security Engineering, FDSE 2021, held in Ho Chi Minh City, Vietnam, in November 2021.* The 28 full papers and 8 short were carefully reviewed and selected from 168 submissions. The selected papers are organized into the following topical headings: big data analytics and distributed systems; security and privacy engineering; industry 4.0 and smart city: data analytics and security; blockchain and access control; data analytics and healthcare systems; and short papers: security and data engineering. * The conference was held virtually due to the COVID-19 pandemic.

A Primer on Memory Consistency and Cache Coherence, Second Edition

Efficient Polymorphic Calls

Algorithms and Architectures for Parallel Processing

Principles of Secure Processor Architecture Design

General-Purpose Graphics Processor Architectures

Mid-America

ISC High Performance 2016 International Workshops, ExaComm, E-MuCoCoS, HPC-IODC, IXPUG, IWOPH, P^3MA, VHPC, WOPSSS, Frankfurt, Germany, June 19–23, 2016, Revised Selected Papers

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 is a primer written for computer architects in the new and rapidly evolving field of deep learning. It reviews how machine learning has evolved since its inception in the 1960s and tracks the key developments leading up to the emergence of the powerful deep learning techniques that emerged in the last decade. Machine learning, and specifically deep learning, has been hugely disruptive in many fields of computer science. The success of deep learning techniques in solving notoriously difficult classification and regression problems has resulted in their rapid adoption in solving real-world problems. The emergence of deep learning is widely attributed to a virtuous cycle whereby fundamental advancements in training deeper models were enabled by the availability of massive datasets and high-performance computer hardware. It also reviews representative workloads, including the most commonly used datasets and seminal networks across a variety of domains. In addition to discussing the workloads themselves, it also details the most popular deep learning tools and show how aspiring practitioners can use the tools with the workloads to characterize and optimize DNNs. The remainder of the book is dedicated to the design and optimization of hardware and architectures for machine learning. As high-performance hardware was so instrumental in the success of machine learning becoming a practical solution, this chapter recounts a variety of optimizations proposed recently to further improve future designs. Finally, it presents a review of recent research published in the area as well as a taxonomy to help readers understand how various contributions fall in context.

Computer Architecture

This book constitutes the proceedings of the 17th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment, DIMVA 2020, held in Lisbon, Portugal, in June 2020. The 13 full papers presented in this volume were carefully reviewed and selected from 45 submissions. The contributions were organized in topical sections named: vulnerability discovery and analysis; attacks; web security; and detection and containment. *The conference was held virtually due to the COVID-19 pandemic.*

From Parallel Processing to the Internet of Things

Progress and Prospects

Data-intensive Text Processing with MapReduce

Formal Techniques for Distributed Objects, Components, and Systems

Second Edition

Fundamentals of Superscalar Processors

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 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.

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

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.

Modern Processor Design

Second International Conference, HPCA 2009, Shanghai, China, August 10–12, 2009, Revised Selected Papers

Architectures, Algorithms, and Applications

Introduction to Reconfigurable Computing

8th International Conference, FDSE 2021, Virtual Event, November 24–26, 2021, Proceedings

Quantum Computing

HPPC 2007, UNICORE Summit 2007, and VHPC 2007, Rennes, France, August 28–31, 2007, Revised Selected Papers

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

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 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.

This book provides a structured treatment of the key principles and techniques for enabling efficient processing of deep neural networks (DNNs). DNNs are currently widely used for many artificial intelligence (AI) applications, including computer vision, speech recognition, and robotics. While DNNs deliver state-of-the-art accuracy on many AI tasks, it comes at the cost of high computational complexity. Therefore, techniques that enable efficient processing of deep neural networks to improve metrics—such as energy-efficiency, throughput, and latency—without sacrificing accuracy or increasing hardware costs are critical to enabling the wide deployment of DNNs in AI systems. The book includes background on DNN processing; a description and taxonomy of hardware architectural approaches for designing DNN accelerators; key metrics for evaluating and comparing different designs; features of the DNN processing that are amenable to hardware/algorithm co-design to improve energy efficiency and throughput; and opportunities for applying new technologies. Readers will find a structured introduction to the field as well as a formalization and organization of key concepts from contemporary works that provides insights that may spark new ideas. Conceptual and precise, Modern Processor Design brings together numerous microarchitectural techniques in a clear, understandable framework that is easily accessible to both graduate and undergraduate students. Complex practices are distilled into foundational principles to reveal the authors insights and hands-on experience in the effective design of contemporary high-performance micro-processors for mobile, desktop, and server markets. Key theoretical and foundational principles are presented in a systematic way to ensure comprehension of important implementation issues. The text presents fundamental concepts and foundational techniques such as processor design, pipelined processors, memory and I/O systems, and especially superscalar organization and implementations. Two case studies and an extensive survey of actual commercial superscalar processors reveal real-world developments in processor design and performance. A thorough overview of advanced instruction flow techniques, including developments in advanced branch predictors, is incorporated. Each chapter concludes with homework problems that will institute the groundwork for emerging techniques in the field and an introduction to multiprocessor systems.

17th International Conference on Information Technology-New Generations (ITNG 2020)

On-Chip Networks

Annual cumulation

Deep Learning for Computer Architects

Languages, Compilers, and Tools for Embedded Systems

Design, Architecture, and Applications

2021 ACM IEEE 48th Annual International Symposium on Computer Architecture (ISCA)

[2], *The Cell Processor from Sony, Toshiba and IBM (STI)* [3], and *the Sun UltraSPARC T1 (formerly codenamed Niagara)* [4] signal the growing popularity of such systems. Furthermore, Intel's very recently announced 80-core TeraFLOP chip [5] exemplifies the irreversible march toward many-core systems with tens or even hundreds of processing elements. 1.2 *The Dawn of the Communication-Centric Revolution* The multi-core thrust has ushered the gradual displacement of the computat- centric design model by a more communication-centric approach [6]. The large, sophisticated monolithic modules are giving way to several smaller, simpler p- cessing elements working in tandem. This trend has led to a surge in the popularity of multi-core systems, which typically manifest themselves in two distinct incarnations: heterogeneous Multi-Processor Systems-on-Chip (MPSoC) and homogeneous Chip Multi-Processors (CMP). The SoC philosophy revolves around the technique of Platform-Based Design (PBD) [7], which advocates the reuse of Intellectual Property (IP) cores in flexible design templates that can be customized accordingly to satisfy the demands of particular implementations. The appeal of such a modular approach lies in the substantially reduced Time-To- Market (TTM) incubation period, which is a direct outcome of lower circuit complexity and reduced design effort. The whole system can now be viewed as a diverse collection of pre-existing IP components integrated on a single die.

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

This book constitutes the proceedings of the 40th IFIP WG 6.1 International Conference on Formal Techniques for Distributed Objects, Components, and Systems, FORTE 2020, held in Valletta, Malta, in June 2020, as part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020.*The 10 full papers and 1 short paper presented were carefully reviewed and selected from 25 submissions. The conference is dedicated to fundamental research on theory, models, tools, and applications for distributed systems. *The conference was held virtually due to the COVID-19 pandemic. Chapter 'Conformance-Based Doping Detection for Cyber-Physical Systems' is available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

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.

19th International Conference, ICA3PP 2019, Melbourne, VIC, Australia, December 9–11, 2019, Proceedings, Part 1

17th International Conference, DIMVA 2020, Lisbon, Portugal, June 24–26, 2020, Proceedings

22nd International Conference, FASE 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings

Future Data and Security Engineering, Big Data, Security and Privacy, Smart City and Industry 4.0 Applications

Computer Architecture

Memory Systems

Compiling Algorithms for Heterogeneous Systems

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.

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.

Parallel and distributed processing, although within the focus of computer science research for a long time, is gaining more and more importance in a wide spectrum of applications. These proceedings aim to demonstrate the use of parallel and distributed processing concepts in different application fields, and attempt to spark interest in novel research directions to advance the embracing model of high-performance computing research in general. The objective of these workshops is to specifically address researchers coming from university, industry and governmental research organizations and application-oriented companies, in order to close the gap between purely scientific research and the applicability of the research ideas to real-life problems. Euro-Par is an annual series of international conferences dedicated to the promotion and advancement of all aspects of parallel and distributed computing. The 2007 event was the 13th issue of the conference. Euro-Par has for a long time been eager to attract colocated events sharing the same goal of promoting the development of parallel and distributed computing, both as an industrial technique and an academic discipline, extending the frontier of both the state of the art and the state of the practice. Since 2006, Euro-Par offers researchers the chance to colocate advanced technical workshops back-to-back with the main conference. This is for a mutual benefit: the workshops can take advantage of all technical and social facilities which are set up for the conference, so that the organizational tasks are kept to a minimal level; the conference can rely on workshopsto experiment with specific areas of research which are not yet mature enough, or too specific, to lead to an official, full-fledged topic at the conference.

This book constitutes the thoroughly refereed post conference papers of the Third International Conference on Blockchain and Trustworthy Systems, Blocksys 2021, held in Guangzhou, China, in August 2021.*The 38 full papers and the 12 short papers were carefully reviewed and selected from 98 submissions. The papers are organized in topical sections: Contents Blockchain and Data Mining; Performance Optimization of Blockchain; Blockchain Security and Privacy; Theories and Algorithms for Blockchain; Blockchain and Internet of Things; Blockchain and Smart Contracts; Blockchain Services and Applications; Trustworthy System Development.*

Cache, DRAM, Disk

HPCA

Network-on-Chip Architectures

From Dataflow to Superscalar and Beyond

Information Security and Cryptology

A Quantitative Approach

The NAS Parallel Benchmarks

This book targets engineers and researchers familiar with basic computer architecture concepts who are interested in learning about on-chip networks. This work is designed to be a short synthesis of the most critical concepts in on-chip network design. It is a resource for both understanding on-chip network basics and for providing an overview of state-of-the-art research in on-chip networks. We believe that an overview that teaches both fundamental concepts and highlights state-of-the-art designs will be of great value to both graduate students and industry engineers. While not an exhaustive text, we hope to illuminate fundamental concepts for the reader as well as identify trends and gaps in on-chip network research. With the rapid advances in this field, we felt it was timely to update and review the state of the art in this second edition. We introduce two new chapters at the end of the book. We have updated the latest research of the past years throughout the book and also expanded our coverage of fundamental concepts to include several research ideas that have now made their way into products and, in our opinion, should be textbook concepts that all on-chip network practitioners should know. For example, these fundamental concepts include message passing, multicast routing, and bubble flow control schemes.

This work is a comprehensive study of the field. It provides an entry point to the novice willing to move in the research field reconfigurable computing, FPGA and system on programmable chip design. The book can also be used as teaching reference for a graduate course in computer engineering, or as reference to advance electrical and computer engineers. It provides a very strong theoretical and practical background to the field, from the early Estrin's machine to the very modern architecture such as embedded logic devices.

This book is Open Access under a CC BY licence. This book constitutes the proceedings of the 22nd International Conference on Fundamental Approaches to Software Engineering, FASE 2019, which took place in Prague, Czech Republic in April 2019, held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019.The 24 papers presented in this volume were carefully reviewed and selected from 94 submissions. The papers are organized in topical sections named: software verification; model-driven development and model transformation; software evolution and requirements engineering; specification, design, and implementation of particular classes of systems; and software testing.

This volume presents the 17th International Conference on Information Technology—New Generations (ITNG), and chronicles an annual event on state of the art technologies for digital information and communications. The application of advanced information technology to such domains as astronomy, biology, education, geosciences, security, and healthcare are among the themes explored by the ITNG proceedings. Visionary ideas, theoretical and experimental results, as well as prototypes, designs, and tools that help information flow to end users are of special interest. Specific topics include Machine Learning, Robotics, High Performance Computing, and Innovative Methods of Computing. The conference features keynote speakers; a best student contribution award, poster award, and service award; a technical open panel, and workshops/exhibits from industry, government, and academia.

High Performance Computing

The Engineering of Large Systems

Detection of Intrusions and Malware, and Vulnerability Assessment

An Historical Review
High Performance Computing and Applications
Distributed and Cloud Computing

ACM SIGPLAN Workshop LCTES '98, Montreal, Canada, June 19-20, 1998, Proceedings

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.

This book is the first technical overview of autonomous vehicles written for a general computing and engineering audience. The authors share their practical experiences of creating autonomous vehicle systems. These systems are complex, consisting of three major subsystems: (1) algorithms for localization, perception, and planning and control; (2) client systems, such as the robotics operating system and hardware platform; and (3) the cloud platform, which includes data storage, simulation, high-definition (HD) mapping, and deep learning model training. The algorithm subsystem extracts meaningful information from sensor raw data to understand its environment and make decisions about its actions. The client subsystem integrates these algorithms to meet real-time and reliability requirements. The cloud platform provides offline computing and storage capabilities for autonomous vehicles. Using the cloud platform, we are able to test new algorithms and update the HD map, train better recognition, tracking, and decision models. This book consists of nine chapters. Chapter 1 provides an overview of autonomous vehicle systems; Chapter 2 focuses on localization technologies; Chapter 3 discusses traditional techniques used for perception; Chapter 4 discusses deep learning based techniques for perception; Chapter 5 introduces the planning and control sub-system, especially prediction and routing technologies; Chapter 6 focuses on motion planning and feedback control of the planning and control subsystem; Chapter 7 introduces reinforcement learning-based planning and control; Chapter 8 delves into the details of client systems design; and Chapter 9 provides the details of cloud platforms for autonomous driving. This book should be useful to students, researchers, and practitioners alike. Whether you are an undergraduate or a graduate student interested in autonomous driving, you will find herein a comprehensive overview of the whole autonomous vehicle technology stack. If you are an autonomous driving practitioner, the many practical techniques introduced in this book will be of interest to you. Researchers will also find plenty of references for an effective, deeper exploration of the various technologies.

Blockchain and Trustworthy Systems

Shared-Memory Synchronization

Creating Autonomous Vehicle Systems