

Altera Sdk For Opencil

This book constitutes the refereed proceedings of the 9th International Conference on Data-Driven Innovation, ICT Innovations 2017, held in Skopje, Macedonia, in September 2017. The 26 full papers presented were carefully reviewed and selected from 90 submissions. They cover the following topics: big data analytics, cloud computing, data mining, digital signal processing, e-health, embedded systems, emerging mobile technologies, multimedia, Internet of Things (IoT), machine learning, software engineering, security and cryptography, coding theory, wearable technologies, wireless communication, and sensor networks.

Design of FPGA-Based Computing Systems with OpenCLSpringer

This book presents novel compiler techniques, which combine a rigorous mathematical framework, novel program analyses and digital hardware design to advance current high-level synthesis tools and extend their scope beyond the industrial 'state of the art'. Implementing computation on customised digital hardware plays an increasingly important role in the quest for energy-efficient high-performance computing. Field-programmable gate arrays (FPGAs) gain efficiency by encoding the computing task into the chip's physical circuitry and are gaining rapidly increasing importance in the processor market, especially after recent announcements of large-scale deployments in the data centre. This is driving, more than ever, the demand for higher design entry abstraction levels, such as the automatic circuit synthesis from high-level languages (high-level synthesis). The techniques in this book apply formal reasoning to high-level synthesis in the context of demonstrably practical applications. /pp

Heterogeneous Computing with OpenCL, Second Edition teaches OpenCL and parallel programming for complex systems that may include a variety of device architectures: multi-core CPUs, GPUs, and fully-integrated Accelerated Processing Units (APUs) such as AMD Fusion technology. It is the first textbook that presents OpenCL programming appropriate for the classroom and is intended to support a parallel programming course. Students will come away from this text with hands-on experience and significant knowledge of the syntax and use of OpenCL to address a range of fundamental parallel algorithms. Designed to work on multiple platforms and with wide industry support, OpenCL will help you more effectively program for a heterogeneous future. Written by leaders in the parallel computing and OpenCL communities, Heterogeneous Computing with OpenCL explores memory spaces, optimization techniques, graphics interoperability, extensions, and debugging and profiling. It includes detailed examples throughout, plus additional online exercises and other supporting materials that can be downloaded at http://www.heterogeneouscompute.org/?page_id=7 This book will appeal to software engineers, programmers, hardware engineers, and students/advanced students. Explains principles and strategies to learn parallel programming with OpenCL, from understanding the four abstraction models to thoroughly testing and debugging complete applications. Covers image processing, web plugins, particle simulations, video editing, performance optimization, and more. Shows how OpenCL maps to an example target architecture and explains some of the tradeoffs associated with mapping to various architectures Addresses a range of fundamental programming techniques, with multiple examples and case studies that demonstrate OpenCL extensions for a variety of hardware platforms

This book describes computational problems related to kernel density estimation (KDE) – one of the most important and widely used data smoothing techniques. A very detailed description of novel FFT-based algorithms for both KDE computations and bandwidth selection are presented. The theory of KDE appears to have matured and is now well developed and understood. However, there is not much progress observed in terms of performance improvements. This book is an attempt to remedy this. The book primarily addresses researchers and advanced graduate or postgraduate students who are interested in KDE and its computational aspects. The book contains both some background and much more sophisticated material, hence also more experienced researchers in the KDE area may find it interesting. The presented material is richly illustrated with many numerical examples using both artificial and real datasets. Also, a number of practical applications related to KDE are presented.

This book constitutes revised selected papers from of the workshops held at 24th International Conference on Parallel and Distributed Computing, Euro-Par 2018, which took place in Turin, Italy, in August 2018. The 64 full papers presented in this volume were carefully reviewed and selected from 109 submissions. Euro-Par is an annual, international conference in Europe, covering all aspects of parallel and distributed processing. These range from theory to practice, from small to the largest parallel and distributed systems and infrastructures, from fundamental computational problems to full-edged applications, from architecture, compiler, language and interface design and implementation to tools, support infrastructures, and application performance aspects.

It constitutes the refereed proceedings of the 4th Asian Supercomputing Conference, SCFA 2018, held in Singapore in March 2018. Supercomputing Frontiers will be rebranded as Supercomputing Frontiers Asia (SCFA), which serves as the technical programme for SCA18. The technical programme for SCA18 consists of four tracks: Application, Algorithms & Libraries Programming System Software Architecture, Network/Communications & Management Data, Storage & Visualisation The 20 papers presented in this volume were carefully reviewed and selected from 60 submissions.

This book presents the proceedings of two conferences, the 37th and 38th in the WoTUG series; Communicating Process Architectures (CPA) 2015, held in Canterbury, England, in August 2015, and CPA 2016, held in Copenhagen, Denmark, in August 2016. Fifteen papers were accepted for presentation at the 2015 conference. They cover a spectrum of concurrency concerns: mathematical theory, programming languages, design and support tools, verification, multicore infrastructure and applications ranging from

supercomputing to embedded. Three workshops and two evening fringe sessions also formed part of the conference, and the workshop position papers and fringe abstracts are included in this book. Fourteen papers covering the same broad spectrum of topics were presented at the 2016 conference, one of them in the form of a workshop. They are all included here, together with abstracts of the five fringe sessions from the conference.

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.

This book focuses on computing devices and their design at various levels to combat variability. The authors provide a review of key concepts with particular emphasis on timing errors caused by various variability sources. They discuss methods to predict and prevent, detect and correct, and finally conditions under which such errors can be accepted; they also consider their implications on cost, performance and quality. Coverage includes a comparative evaluation of methods for deployment across various layers of the system from circuits, architecture, to application software. These can be combined in various ways to achieve specific goals related to observability and controllability of the variability effects, providing means to achieve cross layer or hybrid resilience.

This book constitutes the refereed proceedings of the 11th Joint Conference on Knowledge-Based Software-Engineering, JCKBSE 2014, held in Volgograd, Russia, in September 2014. The 59 full and 3 short papers presented were carefully reviewed and selected from 197 submissions. The papers are organized in topical sections on methodology and tools for knowledge discovery and data mining; methods and tools for software engineering education; knowledge technologies for semantic web and ontology engineering; knowledge-based methods and tools for testing, verification and validation, maintenance and evolution; natural language processing, image analysis and recognition; knowledge-based methods and applications in information security, robotics and navigation; decision support methods for software engineering; architecture of knowledge-based systems, including intelligent agents and softbots; automating software design and synthesis; knowledge management for business processes, workflows and enterprise modeling; knowledge-based methods and applications in bioscience, medicine and justice; knowledge-based requirements engineering, domain analysis and modeling; intelligent user interfaces and human-machine interaction; lean software engineering; program understanding, programming knowledge, modeling programs and programmers.

From the Foreword: "Big Data Management and Processing is [a] state-of-the-art book that deals with a wide range of topical themes in the field of Big Data. The book, which probes many issues related to this exciting and rapidly growing field, covers processing, management, analytics, and applications... [It] is a very valuable addition to the literature. It will serve as a source of up-to-date research in this continuously developing area. The book also provides an opportunity for researchers to explore the use of advanced computing technologies and their impact on enhancing our capabilities to conduct more sophisticated studies." ---Sartaj Sahni, University of Florida, USA "Big Data Management and Processing covers the latest Big Data research results in processing, analytics, management and applications. Both fundamental insights and representative applications are provided. This book is a timely and valuable resource for students, researchers and seasoned practitioners in Big Data fields. --Hai Jin, Huazhong University of Science and Technology, China Big Data Management and Processing explores a range of big data related issues and their impact on the design of new computing systems. The twenty-one chapters were carefully selected and feature contributions from several outstanding researchers. The book endeavors to strike a balance between theoretical and practical coverage of innovative problem solving techniques for a range of platforms. It serves as a repository of paradigms, technologies, and applications that target different facets of big data computing systems. The first part of the book explores energy and resource management issues, as well as legal compliance and quality management for Big Data. It covers In-Memory computing and In-Memory data grids, as well as co-scheduling for high performance computing applications. The second part of the book includes

comprehensive coverage of Hadoop and Spark, along with security, privacy, and trust challenges and solutions. The latter part of the book covers mining and clustering in Big Data, and includes applications in genomics, hospital big data processing, and vehicular cloud computing. The book also analyzes funding for Big Data projects.

This book constitutes the proceedings of the 14th International Workshop on Open MP, IWOMP 2018, held in Barcelona, Spain, in September 2018. The 16 full papers presented in this volume were carefully reviewed and selected for inclusion in this book. The papers are organized in topical sections named: best paper; loops and OpenMP; OpenMP in heterogeneous systems; OpenMP improvements and innovations; OpenMP user experiences: applications and tools; and tasking evaluations.

Parallel computing has been the enabling technology of high-end machines for many years. Now, it has finally become the ubiquitous key to the efficient use of any kind of multi-processor computer architecture, from smart phones, tablets, embedded systems and cloud computing up to exascale computers. This book presents the proceedings of ParCo2013 – the latest edition of the biennial International Conference on Parallel Computing – held from 10 to 13 September 2013, in Garching, Germany. The conference focused on several key parallel computing areas. Themes included parallel programming models for multi- and manycore CPUs, GPUs, FPGAs and heterogeneous platforms, the performance engineering processes that must be adapted to efficiently use these new and innovative platforms, novel numerical algorithms and approaches to large-scale simulations of problems in science and engineering. The conference programme also included twelve mini-symposia (including an industry session and a special PhD Symposium), which comprehensively represented and intensified the discussion of current hot topics in high performance and parallel computing. These special sessions covered large-scale supercomputing, novel challenges arising from parallel architectures (multi-/manycore, heterogeneous platforms, FPGAs), multi-level algorithms as well as multi-scale, multi-physics and multi-dimensional problems. It is clear that parallel computing – including the processing of large data sets (“Big Data”) – will remain a persistent driver of research in all fields of innovative computing, which makes this book relevant to all those with an interest in this field.

This book makes powerful Field Programmable Gate Array (FPGA) and reconfigurable technology accessible to software engineers by covering different state-of-the-art high-level synthesis approaches (e.g., OpenCL and several C-to-gates compilers). It introduces FPGA technology, its programming model, and how various applications can be implemented on FPGAs without going through low-level hardware design phases. Readers will get a realistic sense for problems that are suited for FPGAs and how to implement them from a software designer’s point of view. The authors demonstrate that FPGAs and their programming model reflect the needs of stream processing problems much better than traditional CPU or GPU architectures, making them well-suited for a wide variety of systems, from embedded systems performing sensor processing to large setups for Big Data number crunching. This book serves as an invaluable tool for software designers and FPGA design engineers who are interested in high design productivity through behavioural synthesis, domain-specific compilation, and FPGA overlays. Introduces FPGA technology to software developers by giving an overview of FPGA programming models and design tools, as well as various application examples; Provides a holistic analysis of the topic and enables developers to tackle the architectural needs for Big Data processing with FPGAs; Explains the reasons for the energy efficiency and performance benefits of FPGA processing; Provides a user-oriented approach and a sense for where and how to apply FPGA technology.

This two-volume set constitutes the proceedings of the Third Conference on Creativity in Intellectual Technologies and Data Science, CIT&DS 2019, held in Volgograd, Russia, in September 2019. The 67 full papers, 1 short paper and 3 keynote papers presented were carefully reviewed and selected from 231 submissions. The papers are organized in topical sections in the two volumes. Part I: cyber-physical systems and Big Data-driven world. Part II: artificial intelligence and deep learning technologies for creative tasks; intelligent technologies in social engineering.

P2P, Grid, Cloud and Internet computing technologies have been very fast established as breakthrough paradigms for solving complex problems by enabling aggregation and sharing of an increasing variety of distributed computational resources at large scale. The aim of this volume is to provide latest research findings, innovative research results, methods and development techniques from both theoretical and practical perspectives related to P2P, Grid, Cloud and Internet computing as well as to reveal synergies among such large scale computing paradigms. This proceedings volume presents the results of the 11th International Conference on P2P, Parallel, Grid, Cloud And Internet Computing (3PGCIC-2016), held November 5-7, 2016, at Soonchunhyang University, Asan, Korea

The second of two volumes in the Electronic Design Automation for Integrated Circuits Handbook, Second Edition, Electronic Design Automation for IC Implementation, Circuit Design, and Process Technology thoroughly examines real-time logic (RTL) to GDSII (a file format used to transfer data of semiconductor physical layout) design flow, analog/mixed signal design, physical verification, and technology computer-aided design (TCAD). Chapters contributed by leading experts authoritatively discuss design for manufacturability (DFM) at the nanoscale, power supply network design and analysis, design modeling, and much more. New to This Edition: Major updates appearing in the initial phases of the design flow, where the level of abstraction keeps rising to support more functionality with lower non-recurring engineering (NRE) costs Significant revisions reflected in the final phases of the design flow, where the complexity due to smaller and smaller geometries is compounded by the slow progress of shorter wavelength lithography New coverage of cutting-edge applications and approaches realized in the decade since publication of the previous edition—these are illustrated by new chapters on 3D circuit integration and clock design Offering improved depth and modernity, Electronic Design Automation for IC Implementation, Circuit Design, and Process Technology provides a valuable, state-of-the-art reference for electronic design automation (EDA) students, researchers, and professionals.

The 36 full papers presented in this volume were carefully reviewed and selected from 142 submissions. They deal with parallel and distributed computing in general, focusing on support tools and environments; performance and power modeling, prediction and evaluation; scheduling and load balancing; high performance architectures and compilers; data management, analytics and deep learning; 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; parallel numerical methods and applications; accelerator computing; algorithms and systems for bioinformatics; and algorithms and systems for digital humanities.

Hardware acceleration in the form of customized datapath and control circuitry tuned to specific applications has gained popularity for its promise to utilize transistors more efficiently. Historically, the computer architecture community has focused on general-purpose processors, and extensive research infrastructure has been developed to support research efforts in this domain. Envisioning future computing systems with a diverse set of general-purpose cores and accelerators, computer architects must add accelerator-related research infrastructures to their toolboxes to explore future heterogeneous systems. This book serves as a primer for the field, as an overview of the vast literature on accelerator architectures and their design flows, and as a resource guidebook for researchers working in related areas.

The Complete Guide to OpenACC for Massively Parallel Programming Scientists and technical professionals can use OpenACC to leverage the immense power of modern GPUs without the complexity traditionally associated with programming them. OpenACC™ for Programmers is one of the first comprehensive and practical overviews of OpenACC for massively parallel programming. This book integrates contributions from 19 leading parallel-programming experts from academia, public research organizations, and industry. The authors and editors explain each key concept behind OpenACC, demonstrate how to use essential OpenACC development tools, and thoroughly explore each OpenACC feature set. Throughout, you'll find realistic examples, hands-on exercises, and case studies showcasing the efficient use of OpenACC language constructs. You'll discover how OpenACC's language constructs can be translated to maximize application performance, and how its standard interface can target multiple platforms via widely used programming languages. Each chapter builds on what you've already learned, helping you build practical mastery one step at a time, whether you're a GPU programmer, scientist, engineer, or student. All example code and exercise solutions are available for download at GitHub. Discover how OpenACC makes scalable parallel programming easier and more practical Walk through the OpenACC spec and learn how OpenACC directive syntax is structured Get productive with OpenACC code editors, compilers, debuggers, and performance analysis tools Build your first real-world OpenACC programs Exploit loop-level parallelism in OpenACC, understand the levels of parallelism available, and maximize accuracy or performance Learn how OpenACC programs are compiled Master OpenACC programming best practices Overcome common performance, portability, and interoperability challenges Efficiently distribute tasks across multiple processors Register your product at informit.com/register for convenient access to downloads, updates, and/or corrections as they become available.

The second edition of this accepted reference work has been updated to reflect the rapid developments in the field and now covers both 2D and 3D imaging. Written by expert practitioners from leading companies operating in machine vision, this one-stop handbook guides readers through all aspects of image acquisition and image processing, including optics, electronics and software. The authors approach the subject in terms of industrial applications, elucidating such topics as illumination and camera calibration. Initial chapters concentrate on the latest hardware aspects, ranging from lenses and camera systems to camera-computer interfaces, with the software necessary discussed to an equal depth in later sections. These include digital image basics as well as image analysis and image processing. The book concludes with extended coverage of industrial applications in optics and electronics, backed by case studies and design strategies for the conception of complete machine vision systems. As a result, readers are not only able to understand the latest systems, but also to plan and evaluate this technology. With more than 500 images and tables to illustrate relevant principles and steps.

This book follows an example-driven, simplified, and practical approach to using OpenCL for general purpose GPU programming. If you are a beginner in parallel programming and would like to quickly accelerate your algorithms using OpenCL, this book is perfect for you! You will find the diverse topics and case studies in this book interesting and informative. You will only require a good knowledge of C programming for this book, and an understanding of parallel implementations will be useful, but not necessary.

This book covers the latest approaches and results from reconfigurable computing architectures employed in the finance domain. So-called field-programmable gate arrays (FPGAs) have already shown to outperform standard CPU- and GPU-based computing architectures by far, saving up to 99% of energy depending on the compute tasks. Renowned authors from financial mathematics, computer architecture and finance business introduce the readers into today's challenges in finance IT, illustrate the most advanced approaches and use cases and present currently known methodologies for integrating FPGAs in finance systems together with latest results. The complete algorithm-to-hardware flow is covered holistically, so this book serves as a hands-on guide for IT managers, researchers and quants/programmers who think about integrating FPGAs into their current IT systems.

The book presents the proceedings of four conferences: The 26th International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'20), The 18th International Conference on Scientific Computing (CSC'20); The 17th International Conference on Modeling, Simulation and Visualization Methods (MSV'20); and The 16th International Conference on Grid, Cloud, and Cluster Computing (GCC'20). The conferences took place in Las Vegas, NV, USA, July 27-30, 2020. The conferences are part of the larger 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20), which features 20 major tracks. Authors include academics,

researchers, professionals, and students. Presents the proceedings of four conferences as part of the 2020 World Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE'20); Includes the research tracks Parallel and Distributed Processing, Scientific Computing, Modeling, Simulation and Visualization, and Grid, Cloud, and Cluster Computing; Features papers from PDPTA'20, CSC'20, MSV'20, and GCC'20.

An important working resource for engineers and researchers involved in the design, development, and implementation of signal processing systems The last decade has seen a rapid expansion of the use of field programmable gate arrays (FPGAs) for a wide range of applications beyond traditional digital signal processing (DSP) systems. Written by a team of experts working at the leading edge of FPGA research and development, this second edition of FPGA-based Implementation of Signal Processing Systems has been extensively updated and revised to reflect the latest iterations of FPGA theory, applications, and technology. Written from a system-level perspective, it features expert discussions of contemporary methods and tools used in the design, optimization and implementation of DSP systems using programmable FPGA hardware. And it provides a wealth of practical insights—along with illustrative case studies and timely real-world examples—of critical concern to engineers working in the design and development of DSP systems for radio, telecommunications, audio-visual, and security applications, as well as bioinformatics, Big Data applications, and more. Inside you will find up-to-date coverage of: FPGA solutions for Big Data Applications, especially as they apply to huge data sets The use of ARM processors in FPGAs and the transfer of FPGAs towards heterogeneous computing platforms The evolution of High Level Synthesis tools—including new sections on Xilinx's HLS Vivado tool flow and Altera's OpenCL approach Developments in Graphical Processing Units (GPUs), which are rapidly replacing more traditional DSP systems FPGA-based Implementation of Signal Processing Systems, 2nd Edition is an indispensable guide for engineers and researchers involved in the design and development of both traditional and cutting-edge data and signal processing systems. Senior-level electrical and computer engineering graduates studying signal processing or digital signal processing also will find this volume of great interest.

This book constitutes the proceedings of the 15th International Symposium on Applied Reconfigurable Computing, ARC 2019, held in Darmstadt, Germany, in April 2019. The 20 full papers and 7 short papers presented in this volume were carefully reviewed and selected from 52 submissions. In addition, the volume contains 1 invited paper. The papers were organized in topical sections named: Applications; partial reconfiguration and security; image/video processing; high-level synthesis; CGRAs and vector processing; architectures; design frameworks and methodology; convolutional neural networks.

As predicted by Gordon E. Moore in 1965, the performance of computer processors increased at an exponential rate. Nevertheless, the increases in computing speeds of single processor machines were eventually curtailed by physical constraints. This led to the development of parallel computing, and whilst progress has been made in this field, the complexities of parallel algorithm design, the deficiencies of the available software development tools and the complexity of scheduling tasks over thousands and even millions of processing nodes represent a major challenge to the construction and use of more powerful parallel systems. This book presents the proceedings of the biennial International Conference on Parallel Computing (ParCo2015), held in Edinburgh, Scotland, in September 2015. Topics covered include computer architecture and performance, programming models and methods, as well as applications. The book also includes two invited talks and a number of mini-symposia. Exascale computing holds enormous promise in terms of increasing scientific knowledge acquisition and thus contributing to the future well-being and prosperity of mankind. A number of innovative approaches to the development and use of future high-performance and high-throughput systems are to be found in this book, which will be of interest to all those whose work involves the handling and processing of large amounts of data.

This book gathers 14 of the most promising papers presented at the 18th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2017), which was held on June 26–28, 2017 in Kanazawa, Japan. The aim of this conference was to bring together researchers and scientists, businessmen and entrepreneurs, teachers, engineers, computer users, and students to discuss the various fields of computer science and to share their experiences and exchange new ideas and information in a meaningful way. The book presents research findings concerning all aspects (theory, applications and tools) of computer and information science, and discusses the practical challenges encountered along the way, as well as the solutions adopted to solve them.

This book provides wide knowledge about designing FPGA-based heterogeneous computing systems, using a high-level design environment based on OpenCL (Open Computing language), which is called OpenCL for FPGA. The OpenCL-based design methodology will be the key technology to exploit the potential of FPGAs in various applications such as low-power embedded applications and high-performance computing. By understanding the OpenCL-based design methodology, readers can design an entire FPGA-based computing system more easily compared to the conventional HDL-based design, because OpenCL for FPGA takes care of computation on a host, data transfer between a host and an FPGA, computation on an FPGA with a capable of accessing external DDR memories. In the step-by-step way, readers can understand followings: how to set up the design environment how to write better codes systematically considering architectural constraints how to design practical applications

This book constitutes the refereed proceedings of the Third International Conference on Smart Computing and Communications, SmartCom 2018, held in Tokyo, Japan, in December 2018. The 45 papers presented in this volume were carefully reviewed and selected from 305 submissions. They focus on topics from smart data to smart communications, as well as smart cloud computing to smart security.

This textbook for courses in Embedded Systems introduces students to necessary concepts, through a hands-on approach. It gives a great introduction to FPGA-based microprocessor system design using state-of-the-art boards, tools, and microprocessors from Altera/Intel® and Xilinx®. HDL-based designs (soft-core), parameterized cores (Nios II and MicroBlaze), and ARM Cortex-A9 design are discussed, compared and explored using many hand-on designs projects. Custom IP for HDMI coder, Floating-point operations, and FFT bit-swap are developed, implemented, tested and speed-up is measured.

Downloadable files include all design examples such as basic processor synthesizable code for Xilinx and Altera tools for PicoBlaze, MicroBlaze, Nios II and ARMv7 architectures in VHDL and Verilog code, as well as the custom IP projects. Each Chapter has a substantial number of short quiz questions, exercises, and challenging projects. Explains soft, parameterized, and hard core systems design tradeoffs; Demonstrates design of popular KCPSM6 8 Bit microprocessor step-by-step; Discusses the 32 Bit ARM Cortex-A9 and a basic processor is synthesized; Covers design flows for both FPGA Market leaders Nios II Altera/Intel and MicroBlaze Xilinx system; Describes Compiler-Compiler Tool development; Includes a substantial number of Homework's and FPGA exercises and design projects in each chapter.

This book brings together a selection of the best papers from the eighteenth edition of the Forum on Specification and Design Languages Conference (FDL), which took place on September 14-16, 2015, in Barcelona, Spain. FDL is a well-established international forum devoted to dissemination of research results, practical experiences and new ideas in the application of specification, design and verification languages to the design, modeling and verification of integrated circuits, complex hardware/software embedded systems, and mixed-technology systems.

[Copyright: fa98f7d5e0439e9df3ddead76e66b2b2](https://www.researchgate.net/publication/321111111)