

An Fpga Based Accelerator For Multiple Real Time Template

This is likewise one of the factors by obtaining the soft documents of this An Fpga Based Accelerator For Multiple Real Time Template by online. You might not require more time to spend to go to the books introduction as with ease as search for them. In some cases, you likewise realize not discover the pronouncement An Fpga Based Accelerator For Multiple Real Time Template that you are looking for. It will agreed squander the time.

However below, in the manner of you visit this web page, it will be consequently very simple to acquire as skillfully as download guide An Fpga Based Accelerator For Multiple Real Time Template

It will not acknowledge many period as we explain before. You can reach it though play a part something else at house and even in your workplace. consequently easy! So, are you question? Just exercise just what we provide below as with ease as review An Fpga Based Accelerator For Multiple Real Time Template what you when to read!

Context-Aware Systems and Applications Phan Cong Vinh 2021 This book constitutes the refereed post-conference proceedings of the International Conference on Context-Aware Systems and Applications,

held in October 2021. Due to COVID-19 pandemic the conference was held virtually. The 25 revised full papers presented were carefully selected from 52 submissions. The papers cover a wide spectrum of modern approaches and techniques for smart computing systems and their applications.

Robotic Computing on FPGAs Shaoshan Liu 2022-05-31 This book provides a thorough overview of the state-of-the-art field-programmable gate array (FPGA)-based robotic computing accelerator designs and summarizes their adopted optimized techniques. This book consists of ten chapters, delving into the details of how FPGAs have been utilized in robotic perception, localization, planning, and multi-robot collaboration tasks. In addition to individual robotic tasks, this book provides detailed descriptions of how FPGAs have been used in robotic products, including commercial autonomous vehicles and space exploration robots.

Parallel Computing: On the Road to Exascale G.R. Joubert 2016-04-28 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.

Design of FPGA-Based Computing Systems with OpenCL Hasitha Muthumala Waidyasooriya 2017-10-24 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

Design Space Exploration and Resource Management of Multi/Many-Core Systems Amit Kumar Singh 2021-05-10 The increasing demand of processing a higher number of applications and related data on computing platforms has resulted in reliance on multi-/many-core chips as they facilitate parallel processing. However, there is a desire for these platforms to be energy-efficient and reliable, and they need to perform secure computations for the interest of the whole community. This book provides perspectives on the aforementioned aspects from leading researchers in terms of state-of-the-art contributions and upcoming trends.

Applied Reconfigurable Computing. Architectures, Tools, and Applications Nikolaos Voros 2018-04-25 This book constitutes the proceedings of the 14th International Conference on Applied Reconfigurable Computing, ARC 2018, held in Santorini, Greece, in May 2018. The 29 full papers and 22 short presented in this volume were carefully reviewed and selected from 78 submissions. In addition, the volume contains 9 contributions from research projects. The papers were organized in topical sections

named: machine learning and neural networks; FPGA-based design and CGRA optimizations; applications and surveys; fault-tolerance, security and communication architectures; reconfigurable and adaptive architectures; design methods and fast prototyping; FPGA-based design and applications; and special session: research projects.

FPGA Based Accelerator for Haplotype Inference Naim Majeed Harb 2008 Hardware accelerators have been used to accelerate various bioinformatics applications without altering their accuracy. These accelerators were also seen to accelerate sophisticated algorithms where powerful computational techniques are used to accumulate, analyze, simulate and estimate biological data. These accelerators are hardware accelerators mostly made up of Field Programmable Gate Array (FPGA) hybrid systems or multiple FPGA systems. One bioinformatics application in need for acceleration is the haplotype inference application. This application is essential in producing maps used to identify complex diseases. It is also used in finding phylogenetic trees that provide relationships among different populations. The main objective of this thesis is to build an FPGA-hybrid system connected to a host PC that will accelerate PHASE (one important haplotype inference application) and enhance its processing time while maintaining the same accuracy and functionality.

FPGA Based Accelerators for Financial Applications Christian De Schryver 2015-07-30 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.

Software and Hardware Co-optimization for Deep Learning Algorithms on FPGA Chen Wu 2022

Over recent years, deep learning paradigms such as convolutional neural networks (CNNs) have shown great success in various families of tasks including object detection and autonomous driving, etc. To extend such success to non-euclidean data, graph convolutional networks (GCNs) have been introduced, and have quickly attracted industrial and academia attention as a popular solution to real-world problems. However, both CNNs and GCNs often have huge computation and memory complexity, which calls for specific hardware architectures to accelerate these algorithms. In this dissertation, we propose several architectures to accelerate CNNs and GCNs based on FPGA platforms. We start from the domain-specific FPGA-overlay processor (OPU) on commonly used CNNs, such as VGG, Inception, ResNet, and YoloV2. The data is first quantized to 8-bit fixed-point with little accuracy loss to reduce computation complexity and memory requirement. A fully-pipelined dataflow architecture is proposed to accelerate the typical layers (i.e., convolutional, pooling, residual, inception, and activation layers) in CNNs. Experimental results show that OPU is 9.6 faster than GPU Jetson TX2 on a cascaded of three CNNs, which are used for the curbside parking system. However, 8-bit fixed-point data representation always need re-training to maintain accuracy for deep CNNs. In this way, we propose a low precision (8-bit) floating-point (LPFP) quantization method for FPGA-based acceleration to overcome the above limitation. Without any re-training, LPFP finds an optimal 8-bit data representation with negligible top-1/top-5 accuracy loss (within 0.5%/0.3% in our experiments, respectively, and significantly better than existing methods for deep CNNs). Furthermore, we implement one 8-bit LPFP multiplication by one 4-bit multiply-adder (MAC) and one 3-bit adder. Therefore, we can implement four 8-bit LPFP multiplications using one DSP48E1 of Xilinx Kintex-7 family or one DSP48E2 of Xilinx Ultrascale/Ultrascale Plus family whereas one DSP can only implement two 8-bit fixed-point multiplications. Experiments on six typical CNNs for inference show that on average, we improve throughput by 1.5 over existing FPGA accelerators. Particularly for VGG16

and Yolo, compared with seven FPGA accelerators, we improve average throughput by 3.5 and 27.5 and average throughput per DSP by 4.1 and 5, respectively. CNNs quantized with mixed precision, on the other hand, benefits from low precision while maintaining accuracy. To better leverage the advantages of mixed precision, we propose a Mixed Precision FPGA-based Overlay Processor (MP-OPU) for both conventional and lightweight CNNs. The micro-architecture of MP-OPU considers sharing of computation core with mixed precision weights and activations to improve computation efficiency. In addition, run-time scheduling of external memory access and data arrangement are optimized to further leverage the advantages of mixed precision data representation. Our experimental results show that MP-OPU reaches 4.92 TOPS peak throughput when implemented on Xilinx VC709 FPGA (with all DSPs configured to support 2-bit multipliers). Moreover, MP-OPU achieves 12.9 latency reduction and 2.2 better throughput per DSP for conventional CNNs, while 7.6 latency reduction and 2.9 better throughput per DSP for lightweight CNNs, all on average compared with existing FPGA accelerators/processors, respectively.

Graph convolutional networks (GCNs) have been introduced to effectively process non-euclidean graph data. However, GCNs incur large amount of irregularity in computation and memory access, which prevents efficient use of previous CNN accelerators/processors. In this way, we propose a lightweight FPGA-based accelerator, named LW-GCN, to tackle irregularity in computation and memory access in GCN inference. We first decompose the main GCN operations into Sparse Matrix-Matrix Multiplication (SpMM) and Matrix-Matrix Multiplication (MM). Thereafter, we propose a novel compression format to balance work- load across PEs and prevent data hazards. In addition, we quantize the data into 16-bit fixed-point and apply workload tiling, and map both SpMM and MM onto a uniform archi- tecture on resource limited devices. Evaluations on GCN and GraphSAGE are performed on Xilinx Kintex-7 FPGA with three popular datasets. Compared with existing CPU, GPU and state-of-the-art FPGA-based accelerator, LW-GCN reduces latency by up to 60, 12 and 1.7 and increases power efficiency by up to 912, 511 and 3.87, respectively. Moreover, compared with Nvidia's latest edge GPU Jetson Xavier NX,

LW-GCN achieves speedup and energy savings of 32 and 84, respectively. At last, we extend our GCN inference accelerator to a GCN training accelerator, called SkeletonGCN. To better fit the properties of GCN training, we add more software-hardware co-optimizations. First, we simplify the non-linear operations in GCN training to better fit the FPGA computation, and identify reusable intermediate results to eliminate redundant computation. Second, we optimize the previous compression format to further reduce memory bandwidth while allowing efficient decompression on hardware. Finally, we propose a unified architecture to support SpMM, MM and MM with transpose, all on the same group of PEs to increase DSP utilization on FPGA. Evaluations are performed on Xilinx Alveo U200 board. Compared with existing FPGA-based accelerator on the same network architecture, SkeletonGCN can achieve up to 11.3 speedup while maintaining the same training accuracy with 16-bit fixed-point data representation. In addition, SkeletonGCN is 178 and 13.1 faster than state-of-the-art CPU and GPU implementation on popular datasets, respectively. To summarize, we have been working on FPGA-based acceleration for deep learning algorithms of CNNs and GCNs in both inference and training process. All the accelerators/processors were hand-coded and have been fully verified. In addition, the related tool chains for generating golden results and running instructions for the accelerators/processors have also been finished.

Cyber Physical Systems. Design, Modeling, and Evaluation Roger Chamberlain 2019-04-12 This book constitutes the proceedings of the 7th International Workshop on Design, Modeling, and Evaluation of Cyber Physical Systems, CyPhy2017, held in conjunction with ESWeek 2017, in Seoul, South Korea, in October 2017. The 10 papers presented together with 1 extended and 1 invited abstracts in this volume were carefully reviewed and selected from 16 submissions. The conference presents a wide range of domains including robotics; smart homes, vehicles, and buildings; medical implants; and future-generation sensor networks.

High Performance Computing - HiPC 2007 Srinivas Aluru 2008-01-22 This book constitutes the refereed

proceedings of the 14th International Conference on High-Performance Computing, HiPC 2007, held in Goa, India, in December 2007. The 53 revised full papers presented together with the abstracts of five keynote talks were carefully reviewed and selected from 253 submissions. The papers are organized in topical sections on a broad range of applications including I/O and FPGAs, and microarchitecture and multiprocessor architecture.

Advances in Computers Ali R. Hurson 2022-04-01 Advances in Computers, Volume 126 presents innovations in computer hardware, software, theory, design and applications, with this updated volume including new chapters on VLSI for Super-Computing: Creativity in R+D from Applications and Algorithms to Masks and Chips, Bulk Bitwise Execution Model in Memory: Mechanisms, Implementation, and Evaluation, Embracing the Laws of Physics: Three Reversible Models of Computation, WSNs in Environmental Monitoring: Data Acquisition and Dissemination Aspects, Energy efficient implementation of tensor operations using dataflow paradigm for machine learning, and A Run-Time Job Scheduling Algorithm for Cluster Architectures with DataFlow Accelerators. Contains novel subject matter that is relevant to computer science Includes the expertise of contributing authorsPresents an easy to comprehend writing style

Proceedings of International Scientific Conference on Telecommunications, Computing and Control Nikita Voinov 2021-04-28 This book provides a platform for academics and practitioners for sharing innovative results, approaches, developments, and research projects in computer science and information technology, focusing on the latest challenges in advanced computing and solutions introducing mathematical and engineering approaches. The book presents discussions in the area of advances and challenges of modern computer science, including telecommunications and signal processing, machine learning and artificial intelligence, intelligent control systems, modeling and simulation, data science and big data, data visualization and graphics systems, distributed, cloud and high-performance computing, and software engineering. The papers included are presented at

TELECCON 2019 organized by Peter the Great St. Petersburg University during November 18–19, 2019.

Hardware Accelerator Systems for Artificial Intelligence and Machine Learning 2021-03-28 Hardware Accelerator Systems for Artificial Intelligence and Machine Learning, Volume 122 delves into artificial intelligence and the growth it has seen with the advent of Deep Neural Networks (DNNs) and Machine Learning. Updates in this release include chapters on Hardware accelerator systems for artificial intelligence and machine learning, Introduction to Hardware Accelerator Systems for Artificial Intelligence and Machine Learning, Deep Learning with GPUs, Edge Computing Optimization of Deep Learning Models for Specialized Tensor Processing Architectures, Architecture of NPU for DNN, Hardware Architecture for Convolutional Neural Network for Image Processing, FPGA based Neural Network Accelerators, and much more. Updates on new information on the architecture of GPU, NPU and DNN Discusses In-memory computing, Machine intelligence and Quantum computing Includes sections on Hardware Accelerator Systems to improve processing efficiency and performance

FPGA-BASED Hardware Accelerators Ioulia Skliarova 2019-05-30 This book suggests and describes a number of fast parallel circuits for data/vector processing using FPGA-based hardware accelerators. Three primary areas are covered: searching, sorting, and counting in combinational and iterative networks. These include the application of traditional structures that rely on comparators/swappers as well as alternative networks with a variety of core elements such as adders, logical gates, and look-up tables. The iterative technique discussed in the book enables the sequential reuse of relatively large combinational blocks that execute many parallel operations with small propagation delays. For each type of network discussed, the main focus is on the step-by-step development of the architectures proposed from initial concepts to synthesizable hardware description language specifications. Each type of network is taken through several stages, including modeling the desired functionality in software, the retrieval and automatic conversion of key functions, leading to specifications for optimized hardware modules. The resulting specifications are then synthesized, implemented, and tested in FPGAs using commercial

design environments and prototyping boards. The methods proposed can be used in a range of data processing applications, including traditional sorting, the extraction of maximum and minimum subsets from large data sets, communication-time data processing, finding frequently occurring items in a set, and Hamming weight/distance counters/comparators. The book is intended to be a valuable support material for university and industrial engineering courses that involve FPGA-based circuit and system design.

Architecture of Computing Systems - ARCS 2012 Andreas Herkersdorf 2012-02-09 This book constitutes the refereed proceedings of the 25th International Conference on Architecture of Computing Systems, ARCS 2012, held in Munich, Germany, in February/March 2012. The 20 revised full papers presented in 7 technical sessions were carefully reviewed and selected from 65 submissions. The papers are organized in topical sections on robustness and fault tolerance, power-aware processing, parallel processing, processor cores, optimization, and communication and memory.

Synthesis and Optimization of FPGA-Based Systems Valery Sklyarov 2014-03-14 The book is composed of two parts. The first part introduces the concepts of the design of digital systems using contemporary field-programmable gate arrays (FPGAs). Various design techniques are discussed and illustrated by examples. The operation and effectiveness of these techniques is demonstrated through experiments that use relatively cheap prototyping boards that are widely available. The book begins with easily understandable introductory sections, continues with commonly used digital circuits, and then gradually extends to more advanced topics. The advanced topics include novel techniques where parallelism is applied extensively. These techniques involve not only core reconfigurable logical elements, but also use embedded blocks such as memories and digital signal processing slices and interactions with general-purpose and application-specific computing systems. Fully synthesizable specifications are provided in a hardware-description language (VHDL) and are ready to be tested and incorporated in engineering designs. A number of practical applications are discussed from areas such as data processing and vector-

based computations (e.g. Hamming weight counters/comparators). The second part of the book covers the more theoretical aspects of finite state machine synthesis with the main objective of reducing basic FPGA resources, minimizing delays and achieving greater optimization of circuits and systems.

3rd Kuala Lumpur International Conference on Biomedical Engineering 2006 F. Ibrahim 2007-04-28 The Kuala Lumpur International Conference on Biomedical Engineering (BioMed 2006) was held in December 2006 at the Palace of the Golden Horses, Kuala Lumpur, Malaysia. The papers presented at BioMed 2006, and published here, cover such topics as Artificial Intelligence, Biological effects of non-ionising electromagnetic fields, Biomaterials, Biomechanics, Biomedical Sensors, Biomedical Signal Analysis, Biotechnology, Clinical Engineering, Human performance engineering, Imaging, Medical Informatics, Medical Instruments and Devices, and many more.

Reconfigurable and Adaptive Computing Nadia Nedjah 2018-10-09 Reconfigurable computing techniques and adaptive systems are some of the most promising architectures for microprocessors. Reconfigurable and Adaptive Computing: Theory and Applications explores the latest research activities on hardware architecture for reconfigurable and adaptive computing systems. The first section of the book covers reconfigurable systems. The book presents a software and hardware codesign flow for coarse-grained systems-on-chip, a video watermarking algorithm for the H.264 standard, a solution for regular expressions matching systems, and a novel field programmable gate array (FPGA)-based acceleration solution with MapReduce framework on multiple hardware accelerators. The second section discusses network-on-chip, including an implementation of a multiprocessor system-on-chip platform with shared memory access, end-to-end quality-of-service metrics modeling based on a multi-application environment in network-on-chip, and a 3D ant colony routing (3D-ACR) for network-on-chip with three different 3D topologies. The final section addresses the methodology of system codesign. The book introduces a new software–hardware codesign flow for embedded systems that models both processors and intellectual property cores as services. It also proposes an efficient algorithm for dependent task

software–hardware codesign with the greedy partitioning and insert scheduling method (GPISM) by task graph.

Applied Reconfigurable Computing. Architectures, Tools, and Applications Fernando Rincón 2020-03-25

This book constitutes the proceedings of the 16th International Symposium on Applied Reconfigurable Computing, ARC 2020, held in Toledo, Spain, in April 2020. The 18 full papers and 11 poster presentations presented in this volume were carefully reviewed and selected from 40 submissions. The papers are organized in the following topical sections: design methods & tools; design space exploration & estimation techniques; high-level synthesis; architectures; applications.

Attaining High Performance Communications Ada Gavrilovska 2016-04-19 Technological Advances and

Problems of High Performance Communications An ecosystem of solutions along a stack of technology layers Cohesively collecting state-of-the-art contributions from leading researchers in industry, national laboratories, and academia, Attaining High Performance Communications: A Vertical Approach discusses various issues pertaining to high performance communications in a particular layer of a vertical stack. It explores efficient interconnection hardware, the architectural aspects of network adapters and their integration with processor cores, the design of scalable and robust high performance end-to-end communications services and protocols, and system services and tools for new multi-core environments. No single solution applied at one particular layer can help applications solve all performance-related issues with communication services. Instead, this book shows that a coordinated effort is needed among the layers. It covers many different types of technologies and layers across the stack, from the architectural features of the hardware, through the protocols and their implementation in operating system kernels, to the manner in which application services and middleware are using underlying platforms. The book also describes key developments in high-end platforms, high performance interconnection fabrics and communication libraries, and multi- and many-core systems. This volume addresses the challenges involved in emerging types of communications applications, platforms, and

services. Examining each layer in the vertical stack, it illustrates how to eliminate bottlenecks and provide optimization opportunities.

FPGA-based Implementation of Signal Processing Systems Roger Woods 2017-05 Revised edition of: FPGA-based implementation of signal processing systems / Roger Woods ... [et al.]. 2008.

Reconfigurable Embedded Control Systems: Applications for Flexibility and Agility Khalgui, Mohamed 2010-11-30 "This book addresses the development of reconfigurable embedded control systems and describes various problems in this important research area, which include static and dynamic (manual or automatic) reconfigurations, multi-agent architectures, modeling and verification, component-based approaches, architecture description languages, distributed reconfigurable architectures, real-time and low power scheduling, execution models, and the implementation of such systems"--

High Performance Computing for Big Data Chao Wang 2017-10-16 High-Performance Computing for Big Data: Methodologies and Applications explores emerging high-performance architectures for data-intensive applications, novel efficient analytical strategies to boost data processing, and cutting-edge applications in diverse fields, such as machine learning, life science, neural networks, and neuromorphic engineering. The book is organized into two main sections. The first section covers Big Data architectures, including cloud computing systems, and heterogeneous accelerators. It also covers emerging 3D IC design principles for memory architectures and devices. The second section of the book illustrates emerging and practical applications of Big Data across several domains, including bioinformatics, deep learning, and neuromorphic engineering. Features Covers a wide range of Big Data architectures, including distributed systems like Hadoop/Spark Includes accelerator-based approaches for big data applications such as GPU-based acceleration techniques, and hardware acceleration such as FPGA/CGRA/ASICs Presents emerging memory architectures and devices such as NVM, STT- RAM, 3D IC design principles Describes advanced algorithms for different big data application domains Illustrates novel analytics techniques for Big Data applications, scheduling, mapping, and partitioning

methodologies Featuring contributions from leading experts, this book presents state-of-the-art research on the methodologies and applications of high-performance computing for big data applications. About the Editor Dr. Chao Wang is an Associate Professor in the School of Computer Science at the University of Science and Technology of China. He is the Associate Editor of ACM Transactions on Design Automations for Electronics Systems (TODAES), Applied Soft Computing, Microprocessors and Microsystems, IET Computers & Digital Techniques, and International Journal of Electronics. Dr. Chao Wang was the recipient of Youth Innovation Promotion Association, CAS, ACM China Rising Star Honorable Mention (2016), and best IP nomination of DATE 2015. He is now on the CCF Technical Committee on Computer Architecture, CCF Task Force on Formal Methods. He is a Senior Member of IEEE, Senior Member of CCF, and a Senior Member of ACM.

Architecture Exploration of FPGA Based Accelerators for BioInformatics Applications B. Sharat Chandra Varma 2016-03-02 This book presents an evaluation methodology to design future FPGA fabrics incorporating hard embedded blocks (HEBs) to accelerate applications. This methodology will be useful for selection of blocks to be embedded into the fabric and for evaluating the performance gain that can be achieved by such an embedding. The authors illustrate the use of their methodology by studying the impact of HEBs on two important bioinformatics applications: protein docking and genome assembly. The book also explains how the respective HEBs are designed and how hardware implementation of the application is done using these HEBs. It shows that significant speedups can be achieved over pure software implementations by using such FPGA-based accelerators. The methodology presented in this book may also be used for designing HEBs for accelerating software implementations in other domains besides bioinformatics. This book will prove useful to students, researchers, and practicing engineers alike.

Computers as Components Wayne Hendrix Wolf 2005 This work unravels the complexity of embedded systems, e.g. cell phones, microwaves, and information appliances, and of the process, tools and

techniques necessary for designing them.

Advanced Parallel Processing Technologies Yong Dou 2017-09-13 This book constitutes the proceedings of the 12th International Symposium on Advanced Parallel Processing Technologies, APPT 2017, held in Santiago de Compostela, Spain, in August 2017. The 11 regular papers presented in this volume were carefully reviewed and selected from 18 submissions. They deal with the recent advances in big data processing; parallel architectures and systems; parallel software; parallel algorithms and artificial intelligence applications; and distributed and cloud computing.

Parallel Computing: Technology Trends I. Foster 2020-03-25 The year 2019 marked four decades of cluster computing, a history that began in 1979 when the first cluster systems using Components Off The Shelf (COTS) became operational. This achievement resulted in a rapidly growing interest in affordable parallel computing for solving compute intensive and large scale problems. It also directly led to the founding of the ParCo conference series. Starting in 1983, the International Conference on Parallel Computing, ParCo, has long been a leading venue for discussions of important developments, applications, and future trends in cluster computing, parallel computing, and high-performance computing. ParCo2019, held in Prague, Czech Republic, from 10 – 13 September 2019, was no exception. Its papers, invited talks, and specialized mini-symposia addressed cutting-edge topics in computer architectures, programming methods for specialized devices such as field programmable gate arrays (FPGAs) and graphical processing units (GPUs), innovative applications of parallel computers, approaches to reproducibility in parallel computations, and other relevant areas. This book presents the proceedings of ParCo2019, with the goal of making the many fascinating topics discussed at the meeting accessible to a broader audience. The proceedings contains 57 contributions in total, all of which have been peer-reviewed after their presentation. These papers give a wide ranging overview of the current status of research, developments, and applications in parallel computing.

Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design Nan Zheng

2019-10-18 Explains current co-design and co-optimization methodologies for building hardware neural networks and algorithms for machine learning applications This book focuses on how to build energy-efficient hardware for neural networks with learning capabilities—and provides co-design and co-optimization methodologies for building hardware neural networks that can learn. Presenting a complete picture from high-level algorithm to low-level implementation details, *Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design* also covers many fundamentals and essentials in neural networks (e.g., deep learning), as well as hardware implementation of neural networks. The book begins with an overview of neural networks. It then discusses algorithms for utilizing and training rate-based artificial neural networks. Next comes an introduction to various options for executing neural networks, ranging from general-purpose processors to specialized hardware, from digital accelerator to analog accelerator. A design example on building energy-efficient accelerator for adaptive dynamic programming with neural networks is also presented. An examination of fundamental concepts and popular learning algorithms for spiking neural networks follows that, along with a look at the hardware for spiking neural networks. Then comes a chapter offering readers three design examples (two of which are based on conventional CMOS, and one on emerging nanotechnology) to implement the learning algorithm found in the previous chapter. The book concludes with an outlook on the future of neural network hardware. Includes cross-layer survey of hardware accelerators for neuromorphic algorithms Covers the co-design of architecture and algorithms with emerging devices for much-improved computing efficiency Focuses on the co-design of algorithms and hardware, which is especially critical for using emerging devices, such as traditional memristors or diffusive memristors, for neuromorphic computing *Learning in Energy-Efficient Neuromorphic Computing: Algorithm and Architecture Co-Design* is an ideal resource for researchers, scientists, software engineers, and hardware engineers dealing with the ever-increasing requirement on power consumption and response time. It is also excellent for teaching and training undergraduate and graduate students about the latest generation neural networks

with powerful learning capabilities.

Signal and Information Processing, Networking and Computers Yue Wang 2020-12-17 This book collects selected papers from the 7th Conference on Signal and Information Processing, Networking and Computers held in Rizhao, China, on September, 2020. The 7th International Conference on Signal and Information Processing, Networking and Computers (ICSINC) was held in Rizhao, China, on September, 2020.

OpenMP in a Modern World: From Multi-device Support to Meta Programming Michael Klemm 2022-09-20 This book constitutes the proceedings of the 18th International Workshop on OpenMP, IWOMP 2022, held in Chattanooga, TN, USA, in September 2022. The 11 full papers presented in this volume were carefully reviewed and selected for inclusion in this book from the 13 submissions. The papers are organized in topical sections named: ?OpenMP and multiple nodes; exploring new and recent OpenMP extensions; effective use of advanced heterogeneous node architectures; OpenMP tool support; OpenMP and multiple translation units. Chapter "Improving Tool Support for Nested Parallel Regions with Introspection Consistency" is published Open Access and licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>).

High-Performance Computing Using FPGAs Wim Vanderbauwhede 2013-08-23 High-Performance Computing using FPGA covers the area of high performance reconfigurable computing (HPRC). This book provides an overview of architectures, tools and applications for High-Performance Reconfigurable Computing (HPRC). FPGAs offer very high I/O bandwidth and fine-grained, custom and flexible parallelism and with the ever-increasing computational needs coupled with the frequency/power wall, the increasing maturity and capabilities of FPGAs, and the advent of multicore processors which has caused the acceptance of parallel computational models. The Part on architectures will introduce different FPGA-based HPC platforms: attached co-processor HPRC architectures such as the CHREC's Novo-G and EPCC's Maxwell systems; tightly coupled HPRC architectures, e.g. the Convey hybrid-core computer;

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

Accessing an FPGA-based Hardware Accelerator in a Paravirtualized Environment Wei Wang 2013 In this thesis we present pvFPGA, the first system design solution for virtualizing an FPGA - based hardware accelerator on the x86 platform. The accelerator design on the FPGA can be used for accelerating various applications, regardless of the application computation latencies. Our design adopts the Xen virtual machine monitor (VMM) to build a paravirtualized environment, and a Xilinx Virtex - 6 as an FPGA accelerator. The accelerator communicates with the x86 server via PCI Express (PCIe). In comparison to the current GPU virtualization solutions, which primarily intercept and redirect API calls to the hosted or privileged domain's user space, pvFPGA virtualizes an FPGA accelerator directly at the lower device driver layer. This gives rise to higher efficiency and lower overhead. In pvFPGA, each unprivileged domain allocates a shared data pool for both user - kernel and inter-domain data transfer. In addition, we propose the coprovisor, a new component that enables multiple domains to simultaneously access an FPGA accelerator. The experimental results have shown that 1) pvFPGA achieves close-to-zero overhead compared to accessing the FPGA accelerator without the VMM layer, 2) the FPGA accelerator is successfully shared by multiple domains, 3) distributing different maximum data transfer bandwidths to different domains can be achieved by regulating the size of the shared data pool at the

split driver loading time, 4) request turnaround time is improved through DMA (Direct Memory Access) context switches implemented by the coprocessor.

2016 26th International Conference on Field Programmable Logic and Applications (FPL) IEEE Staff
2016-08-29 The International Conference on Field Programmable Logic and Applications (FPL) is the first and largest conference covering the rapidly growing area of field programmable logic. During the past 26 years, many of the advances achieved in reconfigurable system architectures, applications, embedded processors, design automation methods (EDA) and tools have been first published in the proceedings of the FPL conference series. FPL 2016 will offer the following five conference tracks: Architectures and Technology, Applications and Benchmarks, Design Methods and Tools, Self-aware and Adaptive Systems, Surveys, Trends and Education.

IP Cores Design from Specifications to Production Khaled Salah Mohamed 2015-08-27 This book describes the life cycle process of IP cores, from specification to production, including IP modeling, verification, optimization, and protection. Various trade-offs in the design process are discussed, including those associated with many of the most common memory cores, controller IPs and system-on-chip (SoC) buses. Readers will also benefit from the author's practical coverage of new verification methodologies, such as bug localization, UVM, and scan-chain. A SoC case study is presented to compare traditional verification with the new verification methodologies. Discusses the entire life cycle process of IP cores, from specification to production, including IP modeling, verification, optimization, and protection; Introduce a deep introduction for Verilog for both implementation and verification point of view. Demonstrates how to use IP in applications such as memory controllers and SoC buses. Describes a new verification methodology called bug localization; Presents a novel scan-chain methodology for RTL debugging; Enables readers to employ UVM methodology in straightforward, practical terms.

Hardware Accelerators in Data Centers Christoforos Kachris 2018-08-21 This book provides readers with an overview of the architectures, programming frameworks, and hardware accelerators for typical cloud

computing applications in data centers. The authors present the most recent and promising solutions, using hardware accelerators to provide high throughput, reduced latency and higher energy efficiency compared to current servers based on commodity processors. Readers will benefit from state-of-the-art information regarding application requirements in contemporary data centers, computational complexity of typical tasks in cloud computing, and a programming framework for the efficient utilization of the hardware accelerators.

VLSI and Hardware Implementations using Modern Machine Learning Methods Sandeep Saini 2022-01-19 Machine learning is a potential solution to resolve bottleneck issues in VLSI via optimizing tasks in the design process. This book aims to provide the latest machine-learning–based methods, algorithms, architectures, and frameworks designed for VLSI design. The focus is on digital, analog, and mixed-signal design techniques, device modeling, physical design, hardware implementation, testability, reconfigurable design, synthesis and verification, and related areas. Chapters include case studies as well as novel research ideas in the given field. Overall, the book provides practical implementations of VLSI design, IC design, and hardware realization using machine learning techniques. Features: Provides the details of state-of-the-art machine learning methods used in VLSI design Discusses hardware implementation and device modeling pertaining to machine learning algorithms Explores machine learning for various VLSI architectures and reconfigurable computing Illustrates the latest techniques for device size and feature optimization Highlights the latest case studies and reviews of the methods used for hardware implementation This book is aimed at researchers, professionals, and graduate students in VLSI, machine learning, electrical and electronic engineering, computer engineering, and hardware systems.

Bioinformatics Information Resources Management Association 2013-03-31 "Bioinformatics: Concepts, Methodologies, Tools, and Applications highlights the area of bioinformatics and its impact over the medical community with its innovations that change how we recognize and care for illnesses"--Provided

by publisher.

Correct Hardware Design and Verification Methods Dominique Borrione 2005-09-19 This book constitutes the refereed proceedings of the 13th IFIP WG 10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods, CHARME 2005, held in Saarbrücken, Germany, in October 2005. The 21 revised full papers and 18 short papers presented together with 2 invited talks and one tutorial were carefully reviewed and selected from 79 submissions. The papers are organized in topical sections on functional approaches to design description, game solving approaches, abstraction, algorithms and techniques for speeding (DD-based) verification, real time and LTL model checking, evaluation of SAT-based tools, model reduction, and verification of memory hierarchy mechanisms.

Functional Verification of Dynamically Reconfigurable FPGA-based Systems Lingkan Gong 2014-10-08 This book analyzes the challenges in verifying Dynamically Reconfigurable Systems (DRS) with respect to the user design and the physical implementation of such systems. The authors describe the use of a simulation-only layer to emulate the behavior of target FPGAs and accurately model the characteristic features of reconfiguration. Readers are enabled with this simulation-only layer to maintain verification productivity by abstracting away the physical details of the FPGA fabric. Two implementations of the simulation-only layer are included: Extended Re Channel is a System C library that can be used to check DRS designs at a high level; ReSim is a library to support RTL simulation of a DRS reconfiguring both its logic and state. Through a number of case studies, the authors demonstrate how their approach integrates seamlessly with existing, mainstream DRS design flows and with well-established verification methodologies such as top-down modeling and coverage-driven verification.