Editorial

Selected Papers from the 2011 International Conference on Reconfigurable Computing and FPGAs (ReConFig 2011)

René Cumplido,[1] Peter Athanas,[2] and Jürgen Becker[3]

[1] *Reconfigurable Computing Group, Computer Science Department, National Institute for Astrophysics, Optics and Electronics, 72840 Puebla, PUE, Mexico*

[2] *Department of Electrical and Computer Engineering, Virginia Tech Configurable Computing Laboratory, Blacksburg, VA 24061, USA*

[3] *Institute for Information Processing Technology (ITIV), Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany*

Correspondence should be addressed to René Cumplido; rcumplido@ccc.inaoep.mx

The seventh edition of the International Conference on Reconfigurable Computing and FPGAs (ReConFig 2011) was held in Cancun, Mexico, from November 30 to December 2, 2011. This special issue covers actual and future trends on reconfigurable computing and FPGA technology given by academic and industrial specialists from all over the world. All papers in this special issue are extended versions of selected papers presented at ReConFig 2011, for final publication they were peer-reviewed to ensure that they are presented with the breadth and depth expected from this high-quality journal.

There are a total of 11 papers in this issue. The following 4 papers correspond to the track titled *General sessions*. In "*Analysis of fast radix-10 digit recurrence algorithms for fixed-point and floating-point dividers on FPGAs*," M. Baesler and S. O. Voigt present five different radix-10 digit recurrence dividers for FPGA architectures. All five architectures apply a radix-10 digit recurrence algorithm but differ in the quotient digit selection (QDS) function. In "*Runtime scheduling, allocation and execution of real-time hardware tasks onto Xilinx FPGAs subject to fault occurrence*," Iturbe et al. present describes a novel way to exploit the computation capabilities delivered by modern field-programmable gate Arrays (FPGAs), not only towards a higher performance, but also towards an improved reliability. Computation-specific pieces of circuitry are dynamically scheduled and allocated to different resources on the chip based on a set of novel algorithms which are described in detail. In "*Object recognition and pose estimation on embedded hardware: SURF-based system designs accelerated by FPGA Logic,*" Schaeferling et al. describe two embedded systems for object detection and pose estimation using sophisticated point features. The feature detection step of the Speeded-Up Robust Features (SURF) algorithm is accelerated by a special IP core. The first system performs object detection and is completely implemented in a single medium-size Virtex-5 FPGA. The second system is an augmented reality platform, which consists of an ARM-based microcontroller and intelligent FPGA-based cameras which support the main system. In "*Adaptive multiclient network-on-chip memory core: hardware architecture, software abstraction layer and application exploration,*" D. Göhringer et al. present the hardware architecture and the software abstraction layer of an adaptive multiclient Network-on-Chip (NoC) memory core. The advantages of the novel memory core in terms of performance, flexibility, and user-friendliness are shown using a real-world image processing application.

One paper is within the area of *security and cryptography*. In "*A hardware-accelerated ECDLP with high-performance modular multiplication,*" Judge et al. demonstrate a successful attack on ECC over prime field using the Pollard rho algorithm implemented on a hardware-software cointegrated platform. They propose a high-performance architecture for multiplication over prime field using specialized DSP blocks in the FPGA.

One paper is within the area of *productivity environments and high-level languages*. In "*HwPMI: An extensible performance monitoring infrastructure for improving hardware*

design and productivity on FPGAs," A. G. Schmidt et al. present the hardware performance monitoring infrastructure (HwPMI), which includes a collection of software tools and hardware cores that can be used to profile the current design, recommend/insert performance monitors directly into the HDL or netlist, and retrieve the monitored data with minimal invasion to the design. Three applications are used to demonstrate and evaluate HwPMI's capabilities.

Two papers are within the area of *digital signal processing*. In "*high-level design space and flexibility exploration for adaptive, energy-efficient WCDMA channel estimation architectures,*" Z. E. Rákossy et al. conduct a case study for representatives of two complexity classes of WCDMA channel estimation algorithms and explore the effect of flexibility on energy efficiency using different implementation options. They also propose new design guidelines for both highly specialized architectures and highly flexible architectures using high-level synthesis, to enable the required performance and flexibility to support multiple applications. In "*Configurable transmitter and systolic channel estimator architectures for data-dependent superimposed training communications systems,*" E. Romero et al. present a configurable ST/DDST transmitter and architecture based on array processors (AP) for DDST channel estimation. The high performance and reduced hardware of the proposed architectures leads to the conclusion that the DDST concept can be applied in current communications standards.
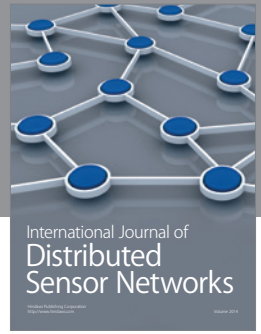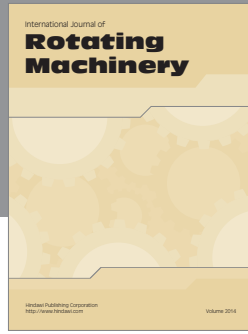
Two papers deal with comparison between FPGAs and GPUs. In "*Novel dynamic partial reconfiguration implementation of K-means clustering on FPGAs: comparative results with GPPs and GPUs,*" H. M. Hussain et al. describe a parameterized implementation of the K-means clustering algorithm in field programmable gate array (FPGA) is presented and compared with previous FPGA implementation as well as recent implementations on graphics processing units (GPUs) and GPPs. The proposed FPGA has higher performance in terms of speed-up over previous GPP and GPU implementations. In "*Multidimensional Costas arrays and their enumeration using GPUs and FPGAs,*" R. A. Arce-Nazario and J. Ortiz-Ubarri present the first proposed implementations for enumerating these multidimensional arrays in GPUs and FPGAs, as well as the first discussion of techniques to prune the search space and reduce enumeration run time. Both GPU and FPGA implementations rely on Costas array symmetries to reduce the search space and perform concurrent explorations over the remaining candidate solutions.

Finally, one paper is within the area of *reconfiguration techniques*. In "*Transparent runtime migration of loop-based traces of processor instructions to reconfigurable processing units,*" J. Bispo et al. focus on the mapping of loop-based instruction traces (called Megablocks) to RPUs. The proposed approach considers offline partitioning and mapping stages without ignoring their future runtime applicability. They present a toolchain that automatically extracts specific trace-based loops, called megablocks, from MicroBlaze instruction traces and generates an RPU for executing those loops. The hardware infrastructure is able to move loop execution from the microprocessor to the RPU transparently, at runtime, and without changing the executable binaries.