



Special Issue on Applied Reconfigurable Computing

Christian Hochberger¹ · Brent Nelson²

Received: 20 February 2021 / Revised: 20 February 2021 / Accepted: 20 February 2021 / Published online: 31 March 2021
© The Author(s) 2021

1 Editorial

Applied reconfigurable computing refers to the use of reconfigurable devices to create complex custom computing systems. Examples of the types of devices focused on by the papers in this issue include field-programmable gate arrays (FPGAs) and coarse-grained reconfigurable arrays (CGRAs). Both allow pre-fabricated silicon devices to be configured and reconfigured at run-time to execute high-performance computations. For many important computational areas the resulting systems can perform orders of magnitude faster than software solutions. Two of these computational areas include high-performance signal processing and control engineering as well as neural networks and machine learning. In addition, sophisticated design tools can aid in the creation of such reconfigurable systems and applications, and the development of such tools is an important research area in the field. The papers in this special issue focus on these three areas.

Papers for this special issue were invited from the 15th International Symposium on Applied Reconfigurable Computing (ARC 2019) held at Technische Universität Darmstadt in Darmstadt, Germany from April 9th to 11th 2019. A total of 14 research groups were invited to prepare extended versions of their papers from the symposium and, after a rigorous review process, 8 papers were selected for inclusion in this special issue.

This issue begins with two papers in the area of high-performance signal processing and control engineering. The paper *»UltraSynth: Insights of a CGRA Integration into a Control Engineering Environment«* explains how a CGRA

could be used in a control environment to execute complex control algorithms. Many facets are discussed like CGRA size and structure or tool runtimes. The paper *»Real-time FPGA implementation of parallel connected component labelling for a 4K video stream«* demonstrates that modern FPGAs can be used to process 4K/UHD video streams for complex analysis like connected component analysis, which is highly demanding.

The issue next turns to the area of convolutional neural networks. The paper *»FPGA-based Inter-layer Pipelined Accelerators for Filter-wise Weight-balanced Sparse Fully Convolutional Networks with Overlapped Tiling«* explains how convolutional neural networks (CNN) can be implemented efficiently by eliminating weights such that all filters have the same number of zero weighted coefficients. Skipping zero weighted coefficients is then realized by internal memories of the FPGAs. The paper *»A Parametrizable High-Level Synthesis Library for Accelerating Neural Networks on FPGAs«* presents a high-level-synthesis library which simplifies the implementation of arbitrary CNNs on FPGAs using deep pipelines to ensure high clock frequencies. Its usage is demonstrated in the HighFlipCV library which is an open source HLS library for image processing. The last paper in this section *»Efficient Design of Pruned Convolutional Neural Networks on FPGA«* proposes a method to prune CNNs and shows an architecture that efficiently executes pruned networks on an FPGA. Particularly, the interplay between pruning and batching is studied in detail.

The concluding set of papers focuses on tools and techniques for the creation of reconfigurable systems and applications. The first paper in this group is *»The TaPaSCo Open-Source Toolflow for the Automated Composition of Task-Based Parallel Reconfigurable Computing Systems«*. The TaPaSCo framework presented in the paper provides an automated toolflow for the construction of heterogeneous many-core architectures from custom processing elements, and a simple, uniform programming interface for using spatially distributed, parallel computation on FPGAs. The second tools paper is *»Generating high-performance FPGA accelerator designs for big data analytics with*

✉ Christian Hochberger
hochberger@rs.tu-darmstadt.de

Brent Nelson
brentnelson@byu.edu

¹ Technische Universität Darmstadt, Darmstadt, Germany

² Brigham Young University, Provo, UT, USA

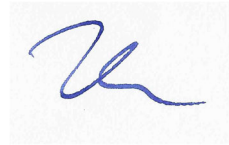
Fletcher and Apache Arrow«. This paper describes the open source tool Fletcher, built on top of Apache Arrow, and which provides a framework designed to deal with the challenges of integrating FPGA-based accelerators into big data analytics pipelines. The combined system accomplishes this by providing a common in-memory format for tabular data sets, a set of associated hardware components, and an associated open source tool to generate hardware interfaces for use by big data analytics accelerators. The final tools paper is »*Evaluation of Static Mapping for Dynamic Space-Shared Multi-Task Processing on FPGAs*«. It first analyzes the shortcomings of previous approaches for dynamically mapping multiple heterogenous tasks onto an executing FPGA and then proposes a static partitioning and mapping approach and tool for the problem, demonstrating more than a doubling of system throughput compared to previous methods.

These eight contributions encompass a wide range of research topics, thereby appealing to both the experts in the field and those who want a snapshot of the current breadth of research in the rapidly expanding field of applied reconfigurable computing.

We would like to thank the reviewers for their constructive comments and many suggestions they made to improve the articles. We also thank the authors for their contributions and work to revise their submissions, with the result being an excellent special issue. We hope that the

reader will benefit from this selection of research highlights from the field of reconfigurable computing.

Sincerely,



Christian Hochberger



Brent Nelson

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.