Introduction to the special section on FPL 2020

Mentens, N.; Sousa, L.; Trancoso, P.

Introduction to the Special Section on FPL 2020

NELE MENTENS, KU Leuven and Leiden University, Belgium and The Netherlands
LIONEL SOUSA, INESC-ID, Instituto Superior Tecnico, Universidade de Lisboa, Portugal
PEDRO TRANCOSO, Chalmers University of Technology, Sweden

ACM Reference format:
Nele Mentens, Lionel Sousa, and Pedro Trancoso. 2022. Introduction to the Special Section on FPL 2020. *ACM Trans. Reconfig. Technol. Syst.* 15, 4, Article 35 (December 2022), 2 pages.
https://doi.org/10.1145/3536336

1 INTRODUCTION

The International Conference on Field Programmable Logic and Applications (FPL) was the first and remains the largest conference in the important area of field-programmable logic and reconfigurable computing. The 30th edition of FPL was scheduled to be from August 31 to September 4, 2020, in the Chalmers Conference Center in Gothenburg, Sweden, but was moved to a virtual format due to the coronavirus disease (COVID-19). From 158 submissions, the program committee selected 24 full papers and 28 short papers to be presented in the conference. The FPL Program co-Chairs invited the authors of the best papers to submit an extended version of their FPL published work for composing a Special Issue of the ACM Transactions on Reconfigurable Technology and Systems.

Six extended articles that went through a completely new review process have been accepted to be published in this Special Issue. These articles bring new results of research efforts in reconfigurable computing, in the areas of placement and connection of nodes and hard-blocks, near-memory processing and HBM, NoCs, and aging in FPGAs.

We acknowledge the support of all reviewers, which are fundamental in the article selection process, also for giving valuable suggestions to the authors. Thanks also go to the authors who submitted articles, and to the ACM TRETS support team. We also thank Professor Deming Chen, Editor-in-Chief of ACM TRETS, for hosting this special issue.

The article *Exploiting HBM on FPGAs for Data Processing* focuses on the potential to exploit High Bandwidth Memory (HBM) for FPGA acceleration of data analytics workloads. The authors investigate different aspects of the computation as well as data partitioning and placement. For the evaluation of the FPGA+HBM setup, the authors integrate into an in-memory database system three relevant workloads: range selection, hash join, and stochastic gradient descent. The results show large performance benefits (6–18×) of the proposed approach when compared to traditional server systems used for the same workloads justifying the use of HBM for FPGA accelerators for these workloads.

The article *Detailed Placement for Dedicated LUT-level FPGA Interconnect* studies the impact of dedicated placement on FPGA architectures with direct connections between the Look-Up Tables (LUTs). The authors propose a novel algorithm that orchestrates different Linear Programs (LPs)

1936-7406/2022/12-ART35
https://doi.org/10.1145/3536336

and Integer Linear Programs (ILPs) for the placement and connection of LUTs in the FPGA. The proposed algorithm was evaluated with several examples and showed that the approach of using direct connections reduces the critical path delay considerably when compared to alternative architecture-oblivious placement algorithms.

The article *RapidLayout: Fast Hard Block Placement of FPGA-optimized Systolic Arrays Using Evolutionary Algorithms* proposes the use of evolutionary algorithms for the placement of hard-block intensive designs. The proposed method obviates the need of manual placement constraints by treating hard block placement as a multi-objective optimization problem. The resulting automated design flow is called RapidLayout and is based on the RapidWright framework of Xilinx. RapidLayout runs significantly faster than existing placement-and-routing flow while the achieved quality of results is similar.

The article *Accelerating Weather Prediction Using Near-memory Reconfigurable Fabric* proposes and evaluates near-memory acceleration of large-scale weather prediction simulations. The acceleration is achieved through reconfigurable logic and high-bandwidth memory. The resulting accelerator is designed using a high-level synthesis approach and is called NERO. The conducted experiments show a significant increase in performance and decrease in energy consumption compared to a 16-core POWER9 system.

The article *HopliteML: Evolving Application Customized FPGA NoCs with Adaptable Routers and Regulators* shows how to evolve hybrid FPGA NoC switch configurations and trace regulation to provide feasibility, worst-case latency, and cost improvements over homogeneous FPGA NoCs. By using fine-grained per-switch static configuration models, it was possible in practice to evolve NoC configurations that offer significant improvements in feasibility, and in worst-case latency over real-world applications.

The article *Inducing Non-uniform FPGA Aging Using Configuration-based Short Circuits* shows the interest of a novel method of accelerating FPGA aging by configuring these devices for high on-chip currents and temperatures through the implementation of thousands of short circuits. Patterns of ring oscillators are placed across the chip to characterize the operating frequency of the FPGA fabric. Over the course of several months of running the short circuits, with daily characterization of the FPGA performance, experiments give insight into the non-uniform nature of the aging caused by the short circuits.