

EXTRA: Towards an Efficient Open Platform for Reconfigurable High Performance Computing

Cătălin Bogdan Ciobanu*, Ana Lucia Varbanescu*, Dionisios Pnevmatikatos[†], George Charitopoulos[†], Xinyu Niu[‡], Wayne Luk[‡], Marco D. Santambrogio[§], Donatella Sciuto[§], Muhammed Al Kadi[¶], Michael Huebner[¶], Tobias Becker^{||}, Georgi Gaydadjiev^{||}, Andreas Brokalakis**, Antonis Nikitakis**, Alex J. W. Thom^{††}, Elias Vansteenkiste^{‡‡}, Dirk Stroobandt^{‡‡}

*UvA, the Netherlands, [†]Telecommunications Systems Institute, Greece

[‡]Imperial College London, UK, [§]Politecnico di Milano, Italy

[¶]Ruhr-Universität Bochum, Germany, ^{||}Maxeler, UK

**University of Cambridge, UK, ^{††}Synelixis, Greece, ^{‡‡}Ghent University, Belgium

Abstract—To handle the stringent performance requirements of future exascale-class applications, High Performance Computing (HPC) systems need ultra-efficient heterogeneous compute nodes. To reduce power and increase performance, such compute nodes will require hardware accelerators with a high degree of specialization. Ideally, dynamic reconfiguration will be an intrinsic feature, so that specific HPC application features can be optimally accelerated, even if they regularly change over time. In the EXTRA project, we create a new and flexible exploration platform for developing reconfigurable architectures, design tools and HPC applications with run-time reconfiguration built-in as a core fundamental feature instead of an add-on. EXTRA covers the entire stack from architecture up to the application, focusing on the fundamental building blocks for run-time reconfigurable exascale HPC systems: new chip architectures with very low reconfiguration overhead, new tools that truly take reconfiguration as a central design concept, and applications that are tuned to maximally benefit from the proposed run-time reconfiguration techniques. Ultimately, this open platform will improve Europe’s competitive advantage and leadership in the field.

I. INTRODUCTION

As power consumption of HPC systems skyrockets with ever more compute intensive tasks, each subtask should be handled with near-optimal power efficiency. This necessarily means that the system has to adapt itself optimally to the current needs of the application. As a result, exascale HPC systems need to be heterogeneous and employ ultra-efficient compute nodes. Some of these nodes will be high performance general purpose processors but other nodes will have to be customized for specific computational kernels and use application specific hardware in order to provide high performance and maximally exploit parallelism at all levels and improve the system energy efficiency. At the same time, these hardware specific nodes will need to be flexible enough to adapt to different applications and their specific requirements, as it is both technically and economically infeasible to include non-programmable acceleration nodes for all possible applications that run on a typical HPC system. The reconfigurability of such hardware nodes becomes therefore mandatory. Applications

that significantly change their behaviour during execution benefit the most from run-time reconfiguration support.

Energy efficiency is one of the main drivers to implement applications on dedicated or reconfigurable hardware rather than on a standard microprocessor system [1]. Higher energy efficiency translates into less heat dissipation. In modern data centers, as much as 50% or more of the overall power consumption is due to non-computing activities such as cooling [2]. Recent work on various HPC applications has shown that GPUs can deliver increased performance and power efficiency over CPUs [3]. However, the overall GPU power consumption is still very high, reaching 300W per card, limiting their deployment in large-scale HPC systems. Reconfigurable devices such as Field Programmable Gate Arrays (FPGAs) are a valid alternative that can provide very high computational performance by creating customised datapaths combined with high power and energy efficiency. The study in [3] shows systematic performance (up to 1.55X) and energy benefits (2.9X to 3.9X) for FPGA implementations for Barrier Option Pricing, Particle Filter, and Reverse-Time Migration when compared to GPUs.

Accelerators alone will not replace the traditional General Purpose processors, there will always be a part of the HPC system constituted by GPPs that will run the OS and control the accelerators. In essence this will lead to heterogeneous architectures. One possible approach is to combine General Purpose Processors (GPPs) with reconfigurable devices, e.g., FPGAs. One representative example of such systems is represented by the Maxeler systems. Very efficient implementations [4] have been obtained on Maxeler hardware while accelerating streaming applications. Micron (former Convey) HC-1 and HC-2 combine Intel Xeon CPUs with FPGAs [5]. Particle physics experiments such as at CERN deal with the high throughput requirements of real-time sensor data and rely almost exclusively on FPGAs for their speed, density, computational power, flexibility, and intrinsic radiation tolerance [6]. IBM has recently announced [7] its strategy for FPGA-enabled acceleration within its POWER8 and OpenPower initiatives.

Microsoft has also recently adopted the dataflow computing approach, programming their own FPGAs [8] to accelerate various BING search engine algorithms. Intel is also expected to produce data center chips combining Xeon CPUs and Altera FPGAs in the near future [9].

Several run-time reconfigurable systems have been proposed over the years. However, several obstacles prevent them from becoming mainstream:

- The tools required for programming such run-time reconfigurable systems still face substantial reconfiguration overheads, which prevent them from being used for large-scale deployment;
- The run-time reconfigurable systems have to use existing FPGA architectures, which are not specifically built with run-time reconfiguration in mind, and therefore lack in efficiency for maximally exploiting possible run-time reconfiguration benefits;
- For newly proposed reconfigurable architectures, the optimal granularity of the reconfiguration infrastructure is still undecided. A low-level reconfiguration infrastructure (such as in current FPGAs) has higher flexibility but larger reconfiguration time, compared to a coarser granularity;
- HPC applications are not optimized for exploiting the available reconfigurability. This is partly because current toolchains do not maximize programmability and designer productivity.

In the EXTRA (Exploiting eXascale Technology with Reconfigurable Architectures) project, we aim to develop an integrated environment for developing and programming reconfigurable architectures with built-in run-time reconfiguration. The idea of this new and flexible exploration platform is to enable the joint optimization of architecture, tools, applications, and reconfiguration technology in order to prepare for the necessary HPC hardware nodes of the future.

The remainder of this paper is organized as follows: the research objectives are described in Section II. Section III presents our approach, and Section IV describes the main challenges. Finally, the paper is concluded in Section V.

II. RESEARCH OBJECTIVES

The main objective of the EXTRA project is to develop an open source research platform for continued research on reconfiguration architectures and tools. The goal is to find architectures and tools that match the next-generation HPC application requirements within a virtual tool environment. Versatile Place and Route (VPR) has been a virtual environment to develop and test placement and routing tools for typical FPGA architectures. We want to provide a platform for run-time reconfiguration that will enable increased research efforts on run-time reconfiguration in Europe.

Because the exploitation of system reconfigurability is relatively new, more research is needed on the optimal HPC architectures that can maximally benefit from reconfiguration, on improvements in the tools to exploit reconfiguration while

designing high performance and power-efficient implementations, and on the application optimizations. Therefore, we identify three Key Objectives (KO) for the success of EXTRA.

KO1. We target the development and promotion of an open reconfigurable technology exploration platform that combines a reconfigurable architecture description with reconfigurable design tools and thus allows to evaluate and optimize reconfigurable applications.

KO2. We aim to make significant contributions to the development of reconfigurable architectures, reconfigurable tools, and the optimization of reconfigurable HPC applications.

KO3. We will validate both the platform and our proposed improvements using the EXTRA ecosystem to implement three HPC applications, with the aim to improve performance, area and power efficiency.

To achieve these key objectives, we identified six major technical objectives that must be achieved.

1: Enable a co-design approach for developing reconfigurable HPC architectures, tools and applications.

Our open research platform allows individual contributions to be tested using the complete chain from device up to the application. This is the first time that such a holistic approach is proposed. The HW/SW partitioning may be semi-automatic, when the tools guide the user by providing profiling data, or automatic. For the improvement and exploration of new design tools for FPGA implementation, a common research platform exists. VPR (now part of the VTR framework [10]) is a common platform for the exploration and improvement of design tools for FPGA implementation.

2: Include reconfigurability as an explicit design concept in future HPC systems design.

Although VPR offers the ability to describe a typical FPGA architecture and offers an open source tool framework for placement and routing tools, it does not include reconfiguration as a specific design option and the architecture descriptions are limited to classical FPGA architectures. We intend to build a conceptually similar infrastructure, but with a lot more flexibility in the architecture choices and a focus on reconfigurability as a design concept.

Both Xilinx and Altera have recently announced that they see a large market potential for their FPGAs in data centers. This enforces our belief that reconfigurable architectures will be essential to the success of future exascale systems. But as reconfigurable devices get larger and more complex, reconfiguring the entire device takes longer and requires more energy; so partial reconfigurability is becoming increasingly important. Thus, both Altera and Xilinx have support for designs with partial run-time reconfiguration in their current devices (Stratix-V [11] and Virtex-7 [12]). In this context, EXTRA will explore (partial) reconfigurability as a specific design feature in future HPC systems, aiming to enable it fully in new reconfigurable architectures, new design tools, and re-engineered applications, optimized for reconfigurability.

3: Speed up the reconfiguration process through novel reconfiguration approaches for processing, BRAMs, special blocks and interconnection in a coarse-grain reconfigura-

tion architecture.

The duration of the reconfiguration process is one of the important bottlenecks in the current reconfigurable systems. By reducing the time needed to reconfigure the hardware, more tasks can be accelerated, significantly reducing the overhead of programming new kernels into the hardware. In order to achieve this goal, EXTRA will consider all the relevant components: processing, BRAMs, specialized blocks and interconnect. Furthermore, the development of optimized hardware structures for HPC workloads allow EXTRA to reduce the need to reconfigure, as a certain level of adaptability is built-in at design-time.

4: Provide just-in-time synthesis methods for reconfiguration on the fly, based on application requirements.

For specific applications, on the fly configuration generation support is desirable, enhancing run-time flexibility. Thus, an adequate configuration scheme for the EXTRA hardware has to be developed to meet the reconfiguration performance (or frequency) requirements. In order to generate this bitstream, a toolset - including synthesis, mapping, placement and routing under hard timing constraint - must be developed.

5: Improve the HPC applications under consideration.

To determine the impact of reconfiguration on HPC applications, we will combine runtime reconfiguration with a domain specific compilation infrastructure (i.e., applications using numerical techniques such as the finite element method). Runtime reconfiguration will add a dynamic layer on top of the existing infrastructure to provide the flexibility to adjust to different data sets as well as differences in numerical properties that may be discovered at runtime. Two other applications will be considered for testing the impact of reconfiguration on HPC workloads: a highly parallel application for the analysis of 3D hrCT images and Quantum Monte Carlo (QMC) methods (both Variational and Diffusion). We note that although parallel and/or hybrid implementations (using CPUs and GPUs) exist, to fully exploit the power of reconfigurable architectures, a fresh implementation of the algorithms is necessary.

6: Suggest new reconfiguration features for future technologies.

We will investigate and suggest practical improvements and necessary features to improve the technology constraints in future reconfigurable systems. We will focus on possible improvements in reconfiguration infrastructure, achieving tighter coupling with the compute cores, and providing hardware support for monitoring and emergency situation management.

III. MAIN APPROACH

The main assumption in the EXTRA project is that system reconfigurability will be a key concept in future HPC systems. In order to develop reconfigurable hardware HPC systems, we need (i) to design completely new system architectures that are inherently reconfigurable, (ii) to develop new tools that enable efficient reconfiguration, and (iii) to identify the applications that can best exploit this novel concept of reconfigurability.

The EXTRA project will tackle all three issues and propose initial architectures, tools and applications that benefit from

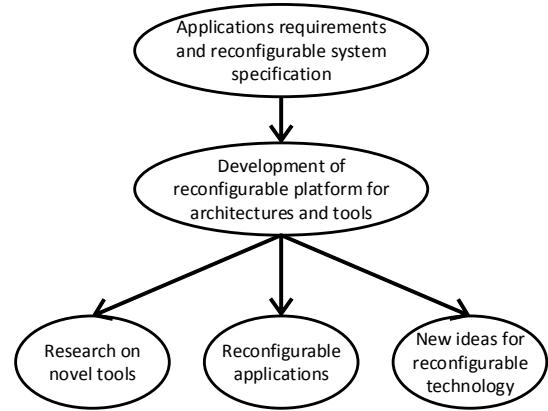


Fig. 1. Conceptual overview of the EXTRA project

reconfigurability. We will focus on building the necessary infrastructure for enabling continued research towards reconfigurable HPC systems for exascale applications while, at the same time, presenting initial solutions that prove that our reconfigurability concept enables more efficient systems and application implementations. It is important to note that reconfiguration can only bring the necessary power efficiency to HPC systems if the reconfiguration can be done while the application is running (run-time reconfiguration) and that the reconfiguration overhead should be significantly smaller than what current systems can offer. Hence, EXTRA will devote significant effort to minimizing the reconfiguration overhead.

The overall approach of the EXTRA project is visually explained in Figure 1. We will investigate how run-time reconfiguration can benefit exascale HPC applications. Based on the application requirements, we will specify the main system requirements for maximally exploiting the benefits of run-time reconfiguration. This will then be the basis for the further work in the project.

The main focus of the EXTRA project is the development of an open source exploration platform that allows the joint investigation of reconfigurable architectures, tools, and applications. The concept is that this open platform will enable many researchers (especially in Europe) to explore novel reconfigurable architectures independently from current commercial vendor solutions. At the same time, the platform provides several hooks within the tool flow to enable tool developers to investigate new tool metrics and propose new tools for designing HPC applications on chosen reconfigurable architectures. These tools will also inherently have reconfigurability included, which is not the case today. Finally, the combination of available reconfigurable architecture descriptions and tools to develop implementations on these architectures provides application developers with an easy to use platform for optimizing their applications. Again, run-time reconfiguration is available everywhere and allows application developers to optimize their applications for it and to evaluate the benefits for their applications using the platform. The open exploration platform for architectures, tools and applications

will also allow the EXTRA consortium partners to make significant contributions in just-in-time synthesis tools for reconfigurable architectures, to efficiently optimize applications for maximally exploiting reconfiguration and evaluating their performance, and to suggest novel reconfiguration technology concepts to improve the efficiency of the reconfiguration within the architectures (bottom part of Figure 1).

We will demonstrate our open source exploration platform and make it available to other researchers in order to create a strong momentum towards research in reconfigurable HPC systems, architectures, tools and applications. Also, the benefits of reconfiguration will be demonstrated by actual implementations of the three EXTRA applications in finite elements, medical imaging, and scientific computing applications on modern commercial reconfigurable devices.

IV. CHALLENGES

The overall EXTRA approach consists of three main parts.

- 1) A thorough investigation on run-time reconfiguration requirements in exascale HPC applications and the specification of system requirements for maximally exploiting the benefits of run-time reconfiguration. *The main research challenges* here are (1) to analyze and characterize the workloads of three HPC application domains, (2) to specify metrics and validation strategies, and finally, (3) to integrate and demonstrate the results, showing that the initial requirements are met.
- 2) The development of an open source reconfiguration platform that allows the joint investigation of reconfigurable architectures, tools, and applications. We focus on the employment of Field Programmable Gate Arrays (FPGAs) as custom hardware accelerators to speed up the hot portion of target applications. Productivity is guaranteed due to the adoption of broadly known programming languages such as C. The open platform development is done on two concurrent directions with tight interaction between them: the *reconfigurable architecture* itself and *the design tools*. *The main challenges for designing the platform itself* are (1) the definition of optimal granularity, based on the analysis of state-of-the-art reconfigurable architectures, (2) the design and development of the tools and (micro)architectural support for optimal interaction between the CPU and the FPGA accelerators, (3) the design space exploration for reconfigurable HPC applications, and (4) the design of the architectural and circuit models for evaluating the feasibility and potential of the proposed platform. The framework for the design tools that use the reconfigurability of the architecture for implementing applications must keep pace with the platform development as well. Finally, the actual integration of the reconfigurable processing units in an exascale system will pose significant challenges.
- 3) Significant contributions in just-in-time synthesis tools for reconfigurable architectures, in efficient optimization of HPC applications, and guidelines for future reconfiguration technology. *The challenges* here include (1)

the development of tools and methods to enable the just-in-time synthesis of configuration for the reconfigurable hardware, (2) systematic analysis, selection, and optimization of applications' functions and structures that can be optimized through reconfiguration, (3) the impact evaluation of the new optimization techniques. Finally, we aim to develop techniques and guidelines that improve the potential of future reconfigurable technology by learning from the past mistakes, i.e., based on all the feedback obtained along in the project. The main challenges here are (1) to collect this feedback coherently and comprehensively, and (2) to transform these issues into actionable points with potential impact on the future of reconfigurable HPC.

V. SUMMARY

In conclusion, this project focuses on the fundamental building blocks for run-time reconfigurable exascale HPC systems: new reconfigurable architectures with very low reconfiguration overhead, new tools that truly take reconfiguration as a design concept, and applications that are tuned to maximally exploit run-time reconfiguration techniques. The developed exploration platform ensures a smooth and efficient co-design of architecture, tools and applications.

ACKNOWLEDGMENTS

This project has received funding from the EU Horizon 2020 research and innovation programme under grant No 671653.

REFERENCES

- [1] J. Rabaey, *Low Power Design Essentials*. Springer, 2009.
- [2] J. Liu *et al.*, "Project Genome: Wireless Sensor Network for Data Center Cooling," *The Architecture Journal*, December 2008. [Online]. Available: research.microsoft.com/apps/pubs/default.aspx?id=78813
- [3] X. Niu *et al.*, "Automating elimination of idle functions by run-time reconfiguration," in *FCCM 2013*, April 2013, pp. 97–104.
- [4] C. Tomas *et al.*, "Acceleration of the Anisotropic PSPI Imaging Algorithm with Dataflow Engines," in *82nd Annual Meeting and International Exposition of the Society of Exploration Geophysics-SEG*, 2012. [Online]. Available: publications.crs4.it/pubdocs/2012/TCOPTSB12
- [5] T. Brewer, "Instruction Set Innovations for the Convey HC-1 Computer," *IEEE Micro*, vol. 30, no. 2, pp. 70–79, 2010.
- [6] L. Musa, "FPGAs in high energy physics experiments at CERN," in *FPL 2008*, Sept 2008, pp. 2–2.
- [7] T. P. Morgan, "IBM Forging Bigger Power8 Systems, Adding FPGA Acceleration," Jul. 2014. [Online]. Available: www.enterprisetech.com/2014/07/28/ibm-forging-bigger-power8-systems-adding-fpga-acceleration/
- [8] J. Clark, "Microsoft 'Catapults' geriatric Moore's Law from CERTAIN DEATH: FPGAs DOUBLE data center throughput despite puny power pump-up, we're told," Jun. 2014. [Online]. Available: www.theregister.co.uk/2014/06/16/microsoft_catapult_fpgas/
- [9] T. P. Morgan, "How Intel is Hedging on the Future of Compute with Altera Buy," Jun. 2014. [Online]. Available: www.theplatform.net/2015/06/01/how-intel-is-hedging-on-the-future-of-compute-with-altera-buy/
- [10] J. Luu *et al.*, "VTR 7.0: Next Generation Architecture and CAD System for FPGAs," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 7, no. 2, pp. 6:1–6:30, Jul. 2014.
- [11] "Increasing Design Functionality with Partial and Dynamic Reconfiguration in 28-nm FPGAs, Altera White Paper, WP-01137-1.0," Jul. 2010.
- [12] B. Przybus, "Xilinx Redefines Power, Performance, and Design Productivity with Three New 28 nm FPGA Families: Virtex-7, Kintex-7, and Artix-7 Devices, Xilinx White Paper WP373 (v1.0)," Jun. 2010.