

# Multi-FPGA Evaluation Platform for Disaggregated Computing

Dimitris Theodoropoulos  
dtheodor@ics.forth.gr

Nikolaos Alachiotis  
nalachio@ics.forth.gr

Dionisios Pnevmatikatos  
pnevmati@ics.forth.gr

Computer Architecture and VLSI Systems Laboratory  
Institute of Computer Science, Foundation for Research and Technology - Hellas (FORTH)

**Abstract**—We present a versatile FPGA-based evaluation platform for exploring alternative execution strategies on disaggregated environments for applications, considering different processing block types: compute cores, memory, and accelerators. Developers can interconnect different blocks types in order to create optimal configurations. A user-level software library allows quick mapping of applications on real hardware. We have implemented a fully working prototype using three ZC706 FPGA boards, and evaluated different software / hardware configurations of a matrix multiplication benchmark.

## I. INTRODUCTION

Our work provides a full prototype development platform for the thorough exploration of alternative execution strategies of applications in disaggregated computing environments. We propose a versatile architecture that employs three different types of processing nodes (blocks): compute cores, memory, and accelerators. A thin, user-level library implements the management API and facilitates block configuration, remote memory allocation, data transfers, and accelerator deployment, collectively minimizing the required porting effort. Recent research projects focus on datacenter architectures supported by reconfigurable modules [1] [2]. In the industrial domain, large companies also propose processing platforms, tightly coupled with FPGA-based reconfigurable coprocessors Micron [3] [4].

## II. DISAGGREGATED PLATFORM

Our architecture (Figure 1) relies on three different types of processing platforms (blocks), i.e. general-purpose computing (compute block - cBlock), host large memory resources (memory block - mBlock), and support hardware acceleration (accelerator block - aBlock), communicating over a low-latency interconnect. Within each block, the interconnect interfaces with DMA engines connected to the local PU for efficient data transfers. The supporting software stack exposes, at user-level, a set of functions that facilitate block-to-block operations, such as remote memory allocation, block-to-block memcpy operations, remote variable iteration, software / hardware task invocation / monitoring, and system debugging.

We use the matrix multiplication benchmark with different mappings on three ZC706 interconnected boards; (1) the PS allocates input L, R and output LR matrices in the cBlock DDR, (2) the PS-MEM allocates L, LR matrices in the cBlock DDR and R matrix in the mBlock DDR, and (3) PS-ACCEL allocates L, R and LR matrices in the aBlock DDR. Table I shows execution time for mappings of different matrix sizes. As expected, for small sizes, the PS mapping

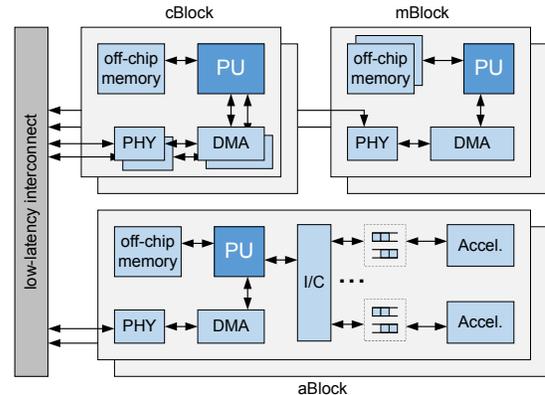


Fig. 1. High-level architecture of the proposed evaluation platform.

TABLE I  
EXECUTION TIME OF THE MATRIX MULTIPLICATION MAPPINGS

Matrix sizes	PS (sec)	PS-MEM (sec)	PS-ACCEL (sec)
256 × 256	1.18	2.06	1.71
512 × 512	9.47	14.55	4.95
1024 × 1024	75.79	110.00	38.30

performs calculations faster, since most data fit well into the PS7 caches. However, for larger sizes, the reconfigurable hardware processing efficiency of the PS-ACCEL leads up to 2x shorter execution time, compared with the baseline PS.

## III. CONCLUSIONS

We presented an evaluation platform for exploring application execution on disaggregated environments. The platform supports different types of resources, which developers can instantiate and interconnect as needed, as well as helps to detect optimal configurations that fit the application requirements.

**Acknowledgement:** This work has been supported in part by EU H2020 ICT project dRedBox, contract #687632.

## REFERENCES

- [1] I. Mavroidis, et al., “ECOSCALE: Reconfigurable computing and runtime system for future exascale systems,” in *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2016, pp. 696–701.
- [2] C. Kachris, et al., “The VINEYARD approach: Versatile, Integrated, Accelerator-based, Heterogeneous Data Centres,” in *International Symposium on Applied Reconfigurable Computing (ARC 2016)*, 2016, pp. 3–13.
- [3] Micron Technology, Inc., “Convey HC2,” <https://www.micron.com/about/about-the-convey-computer-acquisition/hc-series>, [Online; accessed 31-Oct-2016].
- [4] A. Caulfield, et al., “A cloud-scale acceleration architecture,” in *Proceedings of the 49th Annual IEEE/ACM International Symposium on Microarchitecture*, 2016.