

# A Reconfigurable Heterogeneous Microserver Architecture for Energy-efficient Computing

Martin Kaiser, René Griessl,  
Jens Hagemeyer, Dirk Jungewelter,  
Florian Porrhmann, Sarah Pilz and  
Mario Porrhmann  
CITEC, Bielefeld University  
Bielefeld, Germany

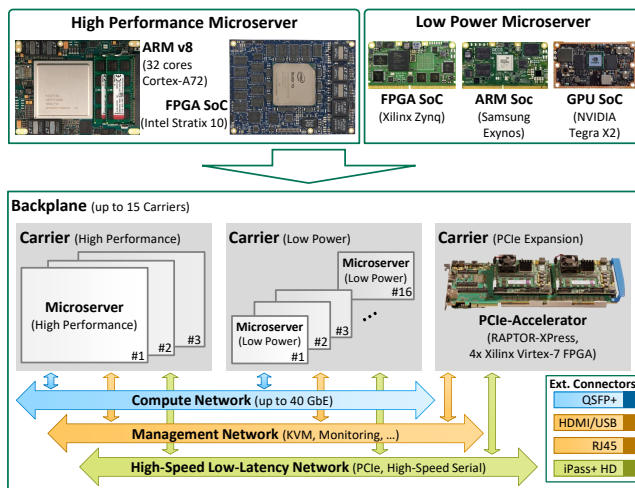
{mkaiser, rgriessl, jhagemeyer, djungewelter,  
fporrhmann, spilz, mporrhmann}@cit-ec.uni-bielefeld.de

Micha vor dem Berge, Stefan Krupop  
christmann informationstechnik + medien GmbH  
Ilsede, Germany  
{micha.vordemberge, stefan.krupop}  
@christmann.info

## Extended Abstract

Using microserver-based system architectures for scale-out applications in cloud computing and HPC could provide a significant advantage in energy-efficiency and TCO, as they offer performance, scalability and high integration density. Specialized server architectures and GPU-based hardware accelerators are increasingly used instead of traditional, CPU-based architectures to achieve higher performance and better energy efficiency. Recently, FPGA-based acceleration is being on the rise due to novel programming models. In this work we propose a scale-out solution for high-density, heterogeneous microservers integrating a high-speed, low-latency communication infrastructure.

has been integrated into our platform. These PCIe-based accelerators can be interconnected to each other to form a unique, densely coupled FPGA cluster. The server architecture permits sharing a PCIe device across multiple microservers or a dynamic assignment of PCIe devices to a dedicated microserver. In contrast to existing microserver platforms that support only homogeneous populations, the proposed modular microserver enables a seamless combination of all these technologies in a single enclosure. This allows for fine-tuning the platform towards a specific application, offering a densely coupled, highly integrated heterogeneous server architecture including a scalable, high-speed, low-latency communication infrastructure. The platform targets form factors on microserver-, baseboard-, backplane-, and chassis-level that match the requirements of today's data centers, enabling hot-swapping and hot-plugging of system components as well as easy integration into existing data center racks.[1]



**Figure 1: Heterogeneous, modular microserver architecture including GPU and FPGA acceleration**

The presented next-generation modular microserver incorporates a broad spectrum of heterogeneous target architectures, making it a versatile platform for a wide range of applications. Specifically, state-of-the-art x86 processors, 64-bit ARM SoCs and server processors, FPGAs, GPUs and others can be integrated. All major processing architectures (CPU, GPU and FPGA) are available in a high-performance as well as a low-power variant. In addition, a custom PCIe-based accelerator card with up to 4 FPGAs

One of the greatest challenges in using heterogeneous server architectures is to program the different kinds of target devices efficiently. As a common programming language, all integrated computation nodes support OpenCL. This includes our custom FPGA-platform RAPTOR-XPress, which comes with a comprehensive software environment and is fully integrated into the Xilinx<sup>®</sup> SDAccel<sup>™</sup> design flow.

Our current work is to extend this design flow to enable communication between multiple FPGAs in a tightly coupled cluster via high-speed serial transceivers. Several benchmark applications are used to assess the performance, scalability and energy efficiency of different applications, e.g. Locality Sensitive Hashing for DNA-processing in bio-computing appliances. These algorithms will be implemented using OpenCL on every compute platform and compared with optimized implementations based on, e.g., OpenMP, VHDL and CUDA. All FPGA implementations will use the high-speed communication infrastructure of the RAPTOR-XPress and will be partitioned across multiple FPGAs.

**Summary** The introduced novel server architecture provides heterogeneous computation nodes for an application-driven, configurable high-speed communication infrastructure, which enables easy usage of hardware accelerators, especially multiple, tightly coupled FPGAs via OpenCL.

## Acknowledgment

This research was supported by the EU Horizon 2020 funded project M2DC (Grant Agreement no. 688201, Webpage: [www.m2dc.eu](http://www.m2dc.eu)). This work was also funded as part of the Cluster of Excellence – Cognitive Interaction Technology *CITEC* (EXC 277), Bielefeld University and the European Fond for regional Development (Europäischer Fond für regionale Entwicklung (EFRE no. 0400079)).

## Keywords

HPC, Cloud Computing, Heterogeneous Microserver, Accelerator, OpenCL, High-Speed Low Latency Communication, Energy Efficiency, Scale-Out Server

## 1. REFERENCES

- [1] A. Oleksiak, M. Kierzynka, W. Piatek, G. Agosta, A. Barengi, M. Porrman, J. Hagemeyer, R. Griessl, J. Lachmair, M. Peykanu, L. Tigges, M. vor dem Berge, W. Christmann, S. Krupop, A. Carbon, L. Cudennec, T. Goubier, J.-M. Philippe, S. Rosinger, D. Schlitt, C. Pieper, C. Adeniyi-Jones, C. Brandolese, W. Fornaciari, G. Pelosi, M. Cecowski, R. Plestenjak, J. Cinkelj, J. Setoain, L. Ceva, and U. Janssen. M2DC – Modular Microserver DataCentre with heterogeneous hardware. *Microprocessors and Microsystems*, 52:117–130, 2017.