

The MSA story of the DEEP projects family

1) Motivation

General purpose systems

- + Highly flexible
- High energy consumption
- + Preferred by many applications

Highly scalable systems

- Few (highly parallelizable) codes can fully exploit them
- + Highly energy efficient

2) Can one combine the best of these two worlds into a single system? → Yes! Exploit heterogeneity!

Homogeneous cluster

- General purpose CPUs attached to a high-speed network
 - + Easy to use, very flexible
 - Power hungry

Traditional heterogeneous cluster

- Attach accelerators (e.g. GPUs) to each CPU
 - + Energy efficient, easy management
 - Static assignment of accelerators to CPUs

3) The basis for the MSA: The Cluster-Booster Concept

The MSA developed in DEEP-EST builds on the so-called Cluster-Booster architecture.

Cluster-Booster architecture

- + Energy efficient, high flexibility, dynamic resource assignment

Does this work?

The Cluster-Booster architecture was first conceptualized and proven with prototypes in the DEEP project. It is a combination of a standard HPC Cluster and a tightly connected HPC Booster built of many-core processors or accelerators. The second project DEEP-ER evolved this architecture to address two significant Exascale computing challenges: highly scalable and efficient parallel I/O and system resiliency. Co-Design was the key to tackle these challenges ? through thoroughly integrated development of new hardware and software components, fine-tuned with actual HPC applications in mind. Results of these two projects showed: **Yes, it works!**

4) Towards a modular supercomputing architecture - The theory

The idea of a MSA is a generalization to any number of specialized modules to address diverse application needs. An example could be a system arranged like this:

Multiple specialized modules will allow a wide range of different applications to efficiently use the system. Each application has its own way of using the MSA. One scenario where the workflows of several applications are distributed over the system could be: