ExaNoDe builds groundbreaking 3D prototype of compute element for exascale

The ExaNoDe project has pioneered 3DIC integration, global shared memory, and a complete software stack that paves the way to a European exascale compute node. Detailed modelling and simulation has shown that a realization using state-of-the-art 7 nm Arm-core based chiplets with silicon interposer and HBM2 would be a modular, cost-effective and energy efficient avenue to achieve multi-teraflops heterogeneous compute nodes.

Barcelona, 10 October 2019

The European ExaNoDe project has built a groundbreaking compute node prototype paving the way to exascale, combining 3DIC with multi-chip-module integration technologies, heterogeneous compute elements with Arm cores and FPGA acceleration and the UNIMEM memory system, all powered by a high-performance, high-productivity software stack.

Denis Dutoit, research engineer at CEA-Leti and the coordinator of ExaNoDe, notes: "Affordability and power consumption are the main hurdles for an exascale-class compute node. In the ExaNoDe project, we have built a complete prototype that integrates multiple core technologies: a 3D active interposer with chiplets, Arm cores with FPGA acceleration, a global address space, high-performance and productive programming environment, which will enable European technology to satisfy the requirements of exascale HPC."

The ExaNoDe protoype is part of the disruptive change required to provide the necessary compute density and power efficiency for an operational exascale machine. Taking as a basis an innovative interposer developed by CEA, ExaNoDe allows the combination of multiple system-on-chips (SoC) chiplets, forming a three-dimensional integrated circuit (3DIC). This delivers multiple advantages, such as:

- higher chip fabrication yields thanks to the smaller chip size
- reduced costs of customization, as the modular design allows combination of the most cutting-edge technology with lower-cost, more established technology as required
- the flexibility to slot in compute elements – such as CPUs and accelerators – in a single chip for different applications, resulting in greater performance at lower design costs
- reduced inter-chip communication distances, resulting in improved energy efficiency

The UNIMEM memory system, which was created in the EUROSERVER project and is being brought to scale in the EuroEXA project, allows the creation of shared memory among multiple compute nodes. The UNIMEM shared memory is accessible through a non-coherent global address space, and is made visible to the programmer via a native UNIMEM API, standard MPI-3.0 and GPI-2. Advances in OmpSs-2@Cluster and

OpenStream allow programmers to exploit the ExaNoDe architecture through a multi-node task-based programming model. In order to increase the resilience and improve the manageability of the compute node, the software stack also includes virtualization, with checkpointing and virtualization of the UNIMEM capabilities.

Finally, ExaNoDe's research activities also extend to applications. Several application areas have been selected to ensure broad coverage, including materials science and engineering. So-called 'mini applications' – self-contained and based on real-life applications – have been developed and ported to the architecture via the programming models and communication application programming interfaces (APIs). Initial work has been performed to accelerate the key kernels on the compute node's FPGA logic, and this expertise will be brought to future and ongoing projects such as EuroEXA. ETHZ developed the open source ExaConv convolutional neural network accelerator to accelerate neural network training as a demonstration of heterogeneous integration.

About ExaNoDe

Thirteen partners from six European countries were involved in the ExaNoDe project: CEA (France), Arm (UK), the University of Manchester (UK), ETH Zürich (Switzerland), CNRS (France), Kalray (France), FORTH (Greece), Virtual Open Systems (France), Fraunhofer ITWM (Germany), Barcelona Supercomputing Center (Spain), Forschungszentrum Jülich (Germany), Atos (France) and scapos (Germany).

CEA, Arm, ETH Zürich and Kalray deployed their knowledge of silicon-level power-management techniques, chiplet and nanotechnology design. CNRS supported CEA in the assembly and packaging of devices. FORTH contributed its expertise in device-to-PCB integration and in the implementation of the UNIMEM memory scheme, firmware and operating systems. Virtual Open Systems provided the virtualized checkpointing and UNIMEM virtualization. Fraunhofer, BSC and University of Manchester enabled the programming environment. Forschungszentrum Jülich and CEA brought their expertise in mini-applications. ATOS provided HPC end-user requirements ensuring that all the technology is appropriate for deployment and integration. scapos ensured that the consortium worked together and successfully delivered the vision of the ExaNoDe project.

This research project has been supported by the European Commission under the Horizon 2020 Framework Programme with grant number 671578.

Website: http://exanode.eu/

Further information

Madeleine Gray, Barcelona Supercomputing Center
Email: madeleine.gray@bsc.es
Tel: (+34) 93 415 739