Energy-efficient Heterogeneous Computing at EXASCALE

In order to sustain the ever-increasing demand of storing, transferring and processing data, HPC servers need to improve their capabilities. Scaling the number of cores alone is not a feasible solution any more due to the increasing utility costs and power consumption limitations. While current HPC systems can offer petaflop performance, their architecture limits their capabilities in terms of scalability and energy consumption. Extrapolating from the top HPC systems, such as China's Tianhe-2 Supercomputer, we estimate that sustaining exaflop performance requires an enormous 1GW power. Similar, albeit smaller, figures are obtained by extrapolating even the best system of the Green 500 list as an initial reference.

To address the energy efficiency and scalability requirements of exascale computing systems, ECOSCALE will analyse the characteristics and trends of current and future HPC applications in order to provide a hybrid and hierarchical many-core + shared distributed reconfigurable accelerator platform. This will be co-developed with a high-level OpenCL application development environment, a runtime system for managing the machine and applications, and a middleware controlling the reconfigurable accelerators. This co-designed ECOSCALE hardware and software components will be used to provide massive boosts in performance and energy efficiency on real world applications.

Industrial Applications

Two applications have been presented to validate the capabilities of the ECOSCALE platform in real industrial processes:

- Smart city application: Based on computer vision, the system is capable to detect patterns of different objects (e.g. pedestrians, vehicles) and feed other smart city subsystems, such as street lighting systems, public/commercial building energy management systems, traffic management systems, etc.

- Reservoir simulation: It is the state-of-the-art technology predicting oil field performance under several production schemes and it requires extreme processing power. The developed high-end reservoir simulation framework will exploit the unique characteristics of the end platform so as to achieve order(s) of magnitude faster and more power efficient simulation runs when compared with the current systems.

For further information, please visit ECOSCALE project website:
http://www.ecoscale.eu

Project coordinator:
Iakovos Mavroidis
Telecommunication Systems Institute (TSI)
Electronic and Computer Engineering Department
Technical University of Crete (TUC)
Chania, GR-73100, Greece
Tel.: +302810254899
Email: iakovosmavro@gmail.com

TSI

Queen's University Belfast

life.augmented

acciona

MANCHESTER 1824
The University of Manchester

POLITECNICO DI TORINO

CHALMERS
UNIVERSITY OF TECHNOLOGY

SYNELIXIS

HORIZON 2020

Description and partners

In order to develop the technology base for future low power large scale HPC computing platforms, ECOSCALE is bringing together a consortium of leading European institutes and industrial partners providing profound experience in fundamental research, product development and applications in the domains HPC, low energy computing, EDA and system and application programming.

Telecommunications Systems Institute (Greece)

Operating within the framework of the Technical University of Crete (T.U.C.) since 1995, has established internationally acknowledged excellence in Computer Architecture and Networks and in Signal and Data Processing with its participation in more than 80 European projects.

Queen's University of Belfast (United Kingdom)

Queen's University is the flagship academic institution of Northern Ireland, a member of the prestigious Russell Group and the 8th best University in the UK in research intensity.

STMicroelectronics (Italy)

STMicroelectronics is one of the world's largest semiconductor companies with net revenues of US$ 7.40 billion in 2014 and approximately 43,600 employees all around the world.

Acciona Infraestructuras (Spain)

Is a leading European construction company operating under the sustainability principles. It has an international presence in more than 30 countries, employing 15.800 people.

University of Manchester (United Kingdom)

Is the largest single-site university in the UK with 25 Nobel laureates among current and former staff and students and with a 50 year tradition in computer science, UoM is ranked 5th in the UK.

Politecnico di Torino (Italy)

Is a leading European technical university, with about 33,000 students, 1500 professors and staff, and a very strong international tradition on teaching and research.

Chalmers University of Technology (Sweden)

Chalmers is among the top technical universities in Sweden, founded in 1829. 65% of its 250 million EUR annual turnover is related to research.

Synelixis (Greece)

Synelixis Solutions Ltd is a high-tech SME, providing turnkey cloud, automation, security and HPC solutions; its technology superiority is a result of extensive R&D covering more than 35% of its turnover.

Technical description

Driven by the characteristics and trends of future HPC applications, ECOSCALE will co-design **a) novel HPC applications** and **b) the novel hierarchical UNIMEM+UNILOGIC architecture** in order to reach exascale performance while achieving exascale class energy-efficiency. The novel UNILOGIC (Unified Logic) architecture, introduced within ECOSCALE for the first time, is an extension of the UNIMEM architecture, proposed within the **EUROSERVER** project. UNIMEM provides shared partitioned global address space while UNILOGIC provides shared partitioned reconfigurable resources. The UNIMEM architecture gives the user the option to move tasks and processes close to data instead of moving data around and thus it reduces significantly the data traffic and related energy consumption and delays. The proposed UNILOGIC+UNIMEM architecture partitions the design into several Worker nodes that communicate through a fat-tree communication infrastructure, similar to the one shown in the figure below.

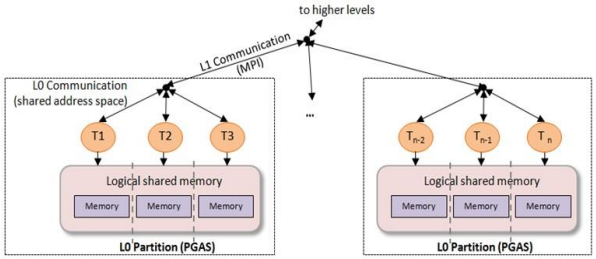

These Worker nodes correspond to the partitions of the HPC application. Each Worker node is an entire sub-system including processing units, memory, and storage. Within a PGAS domain (consisting of several Workers), the proposed architecture offers a shared partitioned global address space and shared partitioned reconfigurable resources that can be accessed via regular load and store instructions without using any global cache coherent mechanism.

In order to further tackle the scalability problems in an exascale machine, ECOSCALE targets **to decrease the number of required interconnected compute nodes** (a compute node is called Worker node in ECOSCALE), which is becoming a critical factor. This number is related both to the computing power and to the energy efficiency level provided by each node. The more computing power provided by a compute node, the fewer compute nodes are required in the HPC system.

On the other hand, the higher the energy efficiency is, the more nodes can be used without breaking the total budget. The approach taken by several architectures, such as the EUROSERVER architecture, is to integrate several low-power CPUs in a single compute node. However, technology advancements (3D-stacking, tri-gates, etc.) are not enough by themselves to provide an exascale solution. Novel architectural approaches are also required. **Using energy-efficient reconfigurable accelerators that can provide significant energy optimizations for data flow applications can further improve the energy efficiency of a compute node.**

While reconfigurable devices (such as FPGAs) have been proven to be significantly more energy efficient than other traditional multi-core architectures (both CPUs and GPUs), enabling up to 25X better performance/watt, their use is still limited, since the path to porting an application onto reconfigurable hardware is often prohibitively cumbersome. Things are even harder in the HPC domain where thousands or millions of reconfigurable accelerators have to be managed.

Therefore, ECOSCALE aims to facilitate this path by providing a novel methodology and architecture to automatically execute HPC applications onto an HPC platform that supports thousands or millions of reconfigurable hardware blocks, while taking into account the projected trends and characteristics of HPC applications. Within this context, the project aims at linking and extending various disconnected existing FPGA-based acceleration approaches and adapting them to work in an HPC environment providing a novel framework from the ground up. In order to efficiently do so, we follow a holistic approach providing solutions for all the aspects of an HPC environment, ranging from architecture and runtime optimizations to high level synthesis (HLS) and hardware virtualization.