textarossa

To achieve high performance and high energy efficiency on near-future exascale computing systems, three key technology gaps needs to be bridged. These gaps include: energy efficiency and thermal control; extreme computation efficiency via HW acceleration and new arithmetics; methods and tools for seamless integration of reconfigurable accelerators in heterogeneous HPC multi-node platforms. TEXTAROSSA aims at tackling this gap through a co-design approach to heterogeneous HPC solutions, supported by the integration and extension of HW and SW IPs, programming models and tools derived from European research.

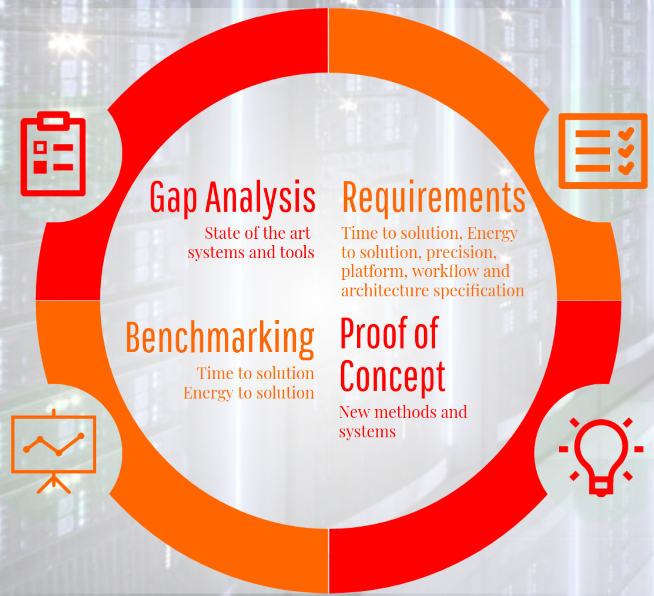

Technical goals

TEXTAROSSA aims at providing key technological advances across the HPC stack and validate them on new development platforms representative of future HPC systems, using a wide range of applications from different domains, both within traditional HPC and coming from emerging domains.

Energy efficiency and thermal control via innovative two-phase cooling technology at node and rack level, fully integrated in an optimized multi-level runtime resource management driven by power, energy, and thermal models fed by on-board sensor data;

Sustained application performance through efficient exploitation of highly concurrent accelerators (GPUs and FPGAs) by focusing on data/stream locality, efficient algorithms and programming models, tuned libraries and innovative IPs;

Seamless integration of reconfigurable accelerators by extending field-proven tools for the design and implementation such as Vitis and OmpSs@FPGA to support new IPs and methodologies such as mixed-precision computing and power monitoring and control;

Development of new IPs for mixed-precision AI computing, data compression, security, power monitoring and control, and scheduling;

Integrated Development Platforms by developing two architecturally different, heterogeneous Integrated Development Vehicles (IDVs), one as a dedicated testbed for two-phase cooling technology, and one supporting the wider range of project technical goals.

Co-design approach

From a methodology point of view TEXTAROSSA adopts a co-design process as key strategy for Fast Forward and Exascale computing, considering the entire system stack from underlying technologies to applications.

The co-design process concerns five layers covering the whole HPC stack:

User Application representing a wide range of scenarios, from mathematical libraries, to miniApps and flagship codes for numerical modelling with massive parallelism in HPC/HPDA/AI applications.

Runtime Services ensuring that application requirements are dynamically satisfied and mapped onto system resources, and including execution models with workload handling, fault tolerance and data management.

Programming Models underlying the applications, they define the toolchains and SW development tools able to implement applications in parallel architectures.

System Architecture including the processor core's microarchitecture, the arrangement of cores within a chip, memory hierarchy, system interconnect, and storage subsystems.

Hardware Platforms Hardware platform at node and rack level able to achieve performance requirements in terms of computing power and energy consumption.

textarossa

Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

The TEXTAROSSA project is funded by the EuroHPC Joint Undertaking under Grant Agreement no 956831

Starting date: April 1st, 2021
Duration: 36 months

Contact: info-textarossa@polimi.it
https://textarossa.eu

TEXTAROSSA Consortium

ENEA, Italy (Coordinator)
Frahnhofer, Germany
CINI, Italy
INRIA, France
ATOS, France
E4, Italy
BSC, Spain
PSNC, Poland
INFN, Italy
CNR, Italy
IN QUATTRO, Italy

Application Use Cases

To address the variety of application domains of future Exascale systems, TEXTAROSSA applications include:

- basic mathematical building blocks: *MathLib*;
- traditional HPC applications: *UrbanAir*, *TNM*, *HEP*, and *RTM*;
- applications from emerging domains: *RAIDER*, *DPSNN*, *Danger Detection*.

MathLib
CNR, FHG, INRIA, ENEA

High performance numerical methods for HPC, HPDA, HPC-AI, including linear algebra, and graph computation

UrbanAir
Air Pollution Model
PSNC

Modelling and forecasting of the concentration and dispersion of air pollutants at meso-scale and city-scale

RAIDER
INFN

Real-time data analytics on heterogeneous distributed systems, processing data streams through Deep Neural Networks

TNM Quantum Simulation
INFN

Tensor Network Method to study in and out of equilibrium properties of strongly correlated many-body quantum systems

Brain Simulation DPSNN
INFN

Distributed and Plastic Spiking Neural Network model of the brain cortex behavior

High Energy Physics
INFN

Optimization of high energy physics simulation and data analysis frameworks

Smart Cities Danger Detection
CINI

Smoke and fire detection in a smart city context, implemented through Convolutional Neural Networks on edge servers

Oil & Gas
Reverse Time Migration
FHG

High performance, energy efficient Reverse Time Migration for Oil & Gas and Geo-Services applications