This project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement No 671668

## MANGO Highlights

MANGO FETHPC-2014 project addresses **manycore architecture exploration** in HPC. It is building a system made of general-purpose nodes (Xeon+GPGPU) coupled with **Heterogeneous Nodes** (HNs). Such nodes include a large-scale cluster of high-capacity FPGAs as well as a robust, scalable interconnect for a multi-FPGA manycore system. The MANGO manycore compute units offer deeply customizable architectures, particularly GPU-like / vector cores that can be instantiated/configured in an application-driven fashion. The FPGA fabric is used as an acceleration platform, but it can also be used as an emulation infrastructure for designing accelerators to be subsequently mapped to different technologies. Importantly, to such a heterogeneous platform MANGO provides a **stable software ecosystem**, independent of the final underlying hardware technologies, encompassing the programming infrastructure and RunTime Manager. The manycore infrastructure provides native hardware-level isolation and partitioning mechanisms for QoS-aware capacity computing HPC applications. Furthermore, MANGO is exploring ultra energy-efficient cooling based on a new concept of **two-phase passive thermosyphon**, enabling PUE values as low as 1.02. The MANGO approach is being demonstrated with applications exhibiting stringent high-performance and QoS requirements, such as videotranscoding, medical imaging, DSP and real-time crypto-processing.

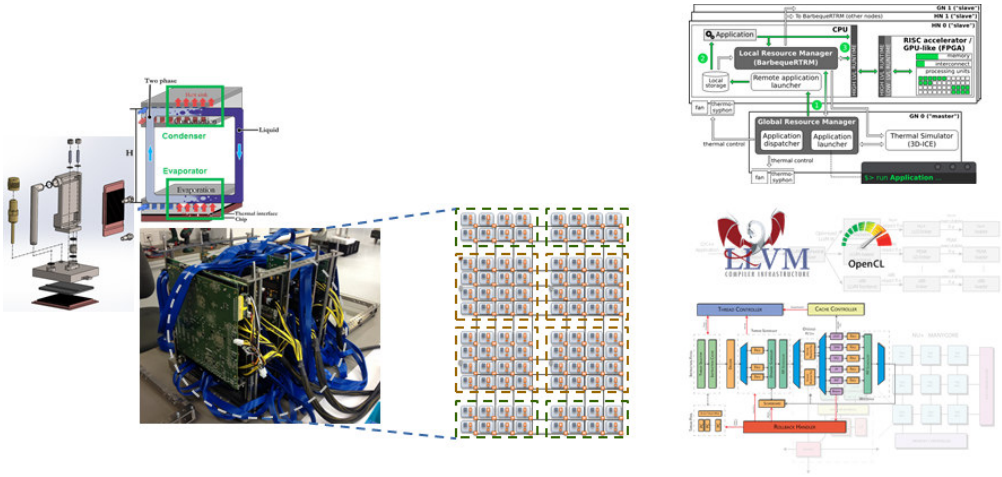


## Anticipated technologies suggested for inclusion in EsD projects

The MANGO key contributions to EsD will include:

- Customizable, **software-programmable GPU-like/vector accelerators** that can be readily configured according to the application requirements. Customization will involve floating-point precision values, number of vector lanes and hardware threads, organization and size of the non-coherent memory, and many additional architecture-level knobs. The software-programmable cores can be easily coupled with custom hardware for specific algorithm acceleration, possibly generated by means of commercial toolflows.
- An advanced infrastructure for interconnecting FPGAs and accelerators in a **large-scale manycore system**. The infrastructure includes a dedicated rack/board design as well as an advanced network with QoS and isolation mechanisms embracing clusters of FPGAs. This enables architecture-wide customization, memory partitioning, as well as some form of close-to-data computing, as key enablers for QoS- and power-aware HPC. As its ultimate goal, the MANGO infrastructure will make HPC ready for FPGA acceleration *at scale*.
- Innovative concept for two-phase **energy-efficient cooling**, based on a gravity-driven passive thermosyphon coupled with a two-level, thermal/energy-aware RunTime Manager. This concept has been demonstrated in MANGO and shown to enable values of **Power Usage Effectiveness (PUE) as**

**low as 1.02**, as opposed to standard air-based cooling systems, reaching a PUE of around 1.60 as well as liquid cooling, reaching a PUE of 1.10.

## Access to MANGO technologies and pre-requisites

- The MANGO customizable vector/GPU-like *nu+* core relies on an LLVM backend developed from scratch and already released in a first version. OpenCL support and related libraries are expected to be provided soon for full integration with standard heterogeneous programming flows. This will also be readily coupled with commercial OpenCL-based HLS flows, for those scenarios where developers need a full custom (non-software programmable) hardware block.
- For the multi-FPGA / manycore infrastructure, the features of the custom interconnect are mostly hidden to applications and high-level software. On the other hand, configuration knobs (for mapping, partitioning, isolation) are exposed to the RTMS. The current prototype supports non-proprietary memory and I/O interfaces, including PCIe, Gigabit Ethernet, and DDR3 controllers, for a straightforward integration with general-purpose nodes (already demonstrated in operation).
- The Global RunTime Management System (RTMS) is based on the SLURM open-source manager, widely used by research and industry (BSC, Bull, CEA, HP, NVIDIA, Intel…). Policies are implemented as plugins. The plugin-based approach requires no modifications to the SLURM core. The Local RTMS is based on the Barbeque open-source project.
- The innovative MANGO cooling system involves mechanical design at the board/rack level. MANGO developed a general methodology for cooling design that can be readily applied to next-generation HPC systems and fully integrated in the EsD technological roadmap.

## Maturity of MANGO technologies

Key innovations in MANGO have been already demonstrated. The MANGO intermediate Review Meeting was held in Naples on May 10th 2017, and all achievements described here were successfully shown in action, in line with the project roadmap and planned timing. According to the MANGO workplan, the key technologies described here are expected to reach a **Technology Readiness Level (TRL) of 6 to 7** by the end of the project in October 2018, providing stable and ready-to-integrate solutions for customizable compute units, multi-FPGA manycore infrastructure, and energy-efficient passive cooling as inputs to **pre EsD1-2 Phase A**.

## What's next

The timing and maturity of MANGO technologies fit the EsD roadmap. In the current stage, partners are porting relevant applications to the MANGO system, including videotranscoding, medical imaging, DSP and real-time crypto-processing. These applications will fully demonstrate the MANGO concept with real-world, complex, data-intensive use-cases, with stringent QoS, performance, and energy-efficiency requirements.
As its core contribution, MANGO will provide a few ***key missing pieces*** for EsD:

- Customizable compute units that can be specified in an application-driven fashion;
- Comprehensive, scalable, future-proof infrastructure support to hardware acceleration in HPC;
- Innovative passive cooling solutions enabling unprecedented values of PUE.

MANGO has some important ***complementarities with other FETHPC projects***. Potential synergies will involve the incorporation of solutions for standard CPUs/accelerators, storage and new memory technologies, and advanced programming models.
As a next step, on top of the current project roadmap, MANGO wants to actively promote factual interplay with other projects, to demonstrate successful integration of innovations coming from different sources. This pre-EsD integration will be essential for the effective exploitation of key research results from current projects. As a further action item, thus, we will soon solicit focused exchange actions with selected FETHPC projects aimed at demonstrating the integration of different, complementary technologies.