

## [OmpSs shown to accelerate applications in FPGA devices and exploit parallelism in clusters](#)

The Barcelona Supercomputing Center (BSC) [Parallel Programming Models team](#) has demonstrated that OmpSs@FPGA can be used to accelerate applications suitable for execution on FPGA devices. The demonstration has shown that with OmpSs@FPGA the performance of a matrix multiplication benchmark accelerated within an FPGA increases 3.5 times with respect to a UDOO board.

In a second demonstration, BSC's team ran the OmpSs matrix multiplication benchmark in a cluster of UDOO boards. Each node had an ARM Cortex-A9 quad core processor and ran Ubuntu Linux. They were connected with a fast ethernet connection. Two directives, target and task, were added to the matrix multiplication kernel in order to spawn the kernel on the different cores on the two boards. The team demonstrated that the same binary application can either be run on a single board or in a distributed way using a number of boards. In this specific demonstration, two boards were used.

For each experiment the benchmark reported the GFlops obtained; in all cases the performance increased. Moving from 1 to 4 cores improves performance almost fourfold. Moving to cluster, the benchmark delivers an additional 1.4 increase in performance.

Xavier Martorell, Parallel Programming Models Group Manager at BSC, states: "The team has been working on the cluster and FPGA environments over the last few years, and we think that getting both execution environments stable for running applications is critical to understanding their potential. We have been working hard on the cluster environment, and with AXIOM we are progressing in the area of FPGA support."

Both demonstrations were carried out in the framework of the [AXIOM project](#) (Agile, eXtensible, fast I/O Module for the cyber-physical era), whose goal is to design a European-manufactured single board computer. Such computers will be the heart of future smart applications from smart homes to smart video surveillance, distributed sensors, the Internet of Things and cybersecurity.

BSC is evaluating the benefits of using FPGAs in production environments. One of its research challenges is to merge its [OmpSs programming model](#) with FPGA engines. This will allow researchers to identify the benefits of both software and hardware approaches in order to get the best out of each, and support the adoption of OmpSs on users' own applications.

### **About the OmpSs Programming Model**

OmpSs integrates features from the StarSs programming model developed by BSC into a single programming model. In particular, it aims to extend OpenMP with new directives to support asynchronous

parallelism and heterogeneity (devices like GPUs, and FPGAs). However, it can also be understood as new directives extending other accelerator-based APIs like CUDA or OpenCL. The OmpSs environment is built on top of the BSC Mercurium compiler and Nanos++ runtime system.

Asynchronous parallelism is enabled in OmpSs by the use of data dependencies between the different tasks of the program. To support heterogeneity, a new construct is introduced: the target construct.

In OmpSs the task construct also allows the annotation of function declarations or definitions in addition to structured-blocks. When a function is annotated with the task construct each invocation of that function becomes a task creation point.

**OmpSs@cluster** allows the exploitation of shared memory parallelism on a cluster of nodes with accelerators.

**OmpSs@FPGA** allows kernels to be executed on FPGA devices.

[pm.bsc.es/ompss](http://pm.bsc.es/ompss)

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

---

**Source URL (retrieved on 16 Jul 2024 - 01:07):** <https://www.bsc.es/es/news/bsc-news/ompss-shown-accelerate-applications-fpga-devices-and-exploit-parallelism-clusters>