

Programming Models



Our team explores the **programming** approaches for **homogeneous and heterogeneous architectures**, including the use of accelerators and distributed clusters. We propose the **OmpSs programming model**, a task-based approach. OmpSs seeks to improve programmability of applications in such environments, without sacrificing performance. We actively participate in the **OpenMP standard** committee, and use OmpSs for rapid prototyping of new ideas that can be then considered for OpenMP. Our team also proposes the use of **Domain Specific Languages (DSLs)** to permit scientists with higher level languages, running on top of OmpSs, to better code their applications.

Objectives

The main goal of the team is to investigate new and current programming paradigms and the associated runtime system support to provide high performance to parallel applications. The target architectures range from multicore and SMT processors to shared- and distributed-memory systems. Our main platform for research and development is a multicore-based cluster with accelerators.

In terms of the programming model, we are proposing extensions to OpenMP to improve the expressiveness of the model. The Mercurium Compiler is able to exploit multiple levels of parallelism and generate work from multiple simultaneously executing threads. Once parallelism is spawned on a coarse level, new opportunities for parallelism on a finer level result in the generation of work for all or for restricted groups of processors. Although our primary focus is on OpenMP, we are also investigating ways to exploit the parallelism in environments with accelerators, and in distributed memory architectures.

The runtime system offers the basic services to spawn or join parallelism and synchronise across threads. The Nanos++ API is simple enough to be used for multithreaded programming on a variety of multiprocessor platforms. The goal is to have a platform that is easy to modify to incorporate new functionalities as needed by the new architectures. For instance, currently it is being used to experiment with accelerator support on GPUs, Xeon PHI and FPGAs.

Domain Specific Languages, executed on top of OmpSs, provide ease of use for scientists to code their applications.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 31 Mar 2025 - 10:33): <https://www.bsc.es/ca/discover-bsc/organisation/scientific-structure/programming-models>