



THE ANTAREX PROJECT

Autotuning and Adaptivity Approach for Energy Efficient Exascale HPC Systems

The ANTAREX research project, coordinated by prof. Cristina Silvano from Politecnico di Milano, has been granted in the H2020 Future and Emerging Technologies programme on High Performance Computing. The project involves CINECA, the Italian Tier-0 Supercomputing Centre and IT4Innovations, the Czech Tier-1 Supercomputing Center. The Consortium also includes three top-ranked academic partners (ETH Zurich, University of Porto, and INRIA). Industrial partners include one of the Italian leading biopharmaceutical companies (Dompé) and the top European navigation software company (Sygic). Being one of the nineteen research projects in FET-HPC-2014, ANTAREX brings its partners on the forefront of the European research in High Performance Computing. The project just started on September the 1st, 2015.

The main goal of the ANTAREX project is to provide a breakthrough approach to map, runtime manage and autotune applications for green and heterogeneous High Performance Computing systems up to the Exascale level. The approach will be based on the tool-flow shown in Fig. 1. One key innovation of the proposed approach consists of introducing a separation of concerns (where self-adaptivity and energy efficient strategies are specified aside to application functionalities) promoted by the definition of a Domain Specific Language (DSL) inspired by aspect-oriented programming concepts for heterogeneous systems. The new DSL will be introduced for expressing adaptivity/energy/performance strategies and to enforce at runtime application autotuning and resource and power management. The goal is to support the parallelism, scalability and adaptability of a dynamic workload by exploiting the full system capabilities (including energy management) for emerging large-scale and extreme-scale systems, while reducing the Total Cost of Ownership (TCO) for companies and public organizations.

The ANTAREX project is driven by two use cases chosen to address the self-adaptivity and scalability characteristics of two highly relevant HPC application scenarios. The two use cases are: (1) a biopharmaceutical HPC application for accelerating drug discovery deployed on the 1.21 PetaFlops heterogeneous NeXtScale Intel-based IBM system at CINECA; and (2) a self-adaptive navigation system to be used in smart cities deployed on the server-side on a heterogeneous Intel-based 1.46 PetaFlops class system provided by IT4Innovations Supercomputing Center. All the key ANTAREX software innovations will be designed and engineered since the beginning to be scaled-up to the Exascale level. Performance metrics extracted from the two use cases will be modelled to extrapolate these results towards Exascale systems. These use cases have been selected due to their significance in emerging application trends and thus by their direct economic exploitability and relevant social impact.

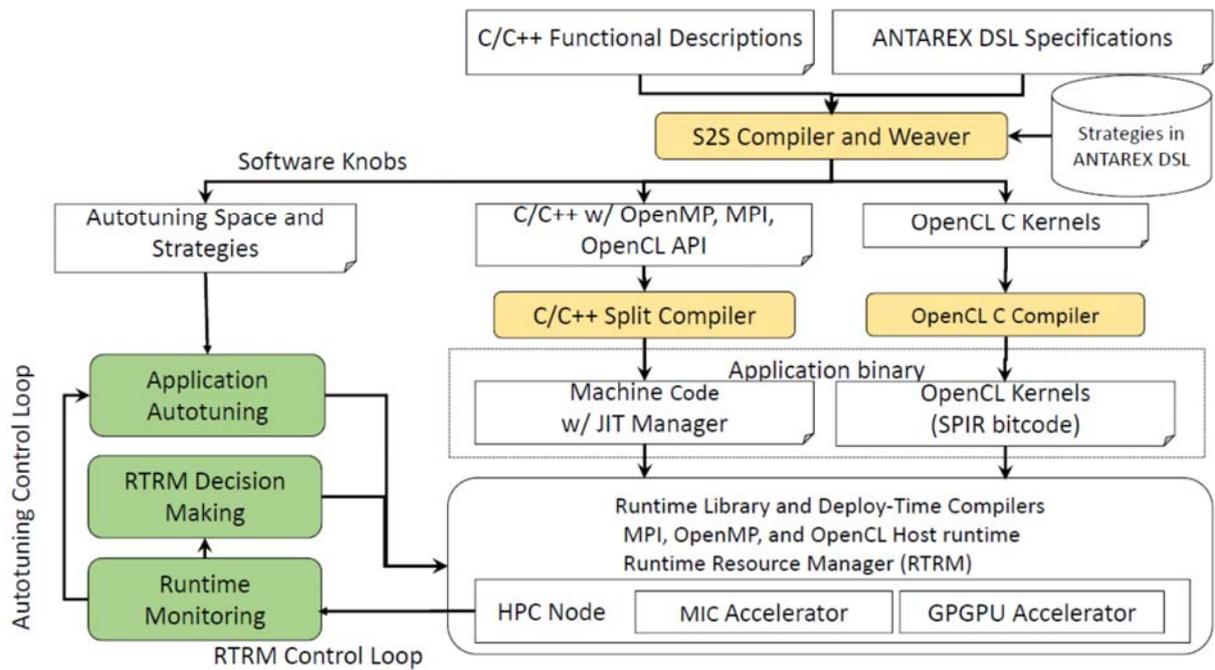


Figure 1: The ANTAREX tool-flow



Figure 2: The ANTAREX Group Photo at the Kick-off Meeting held in September 2015 at CINECA (Italy).