The computing world is changing rapidly. All devices – from mobile phones and personal computers to high-performance supercomputers – are based on multi-/manycore architectures that become increasingly complex, hierarchical and heterogeneous. The huge capacity of modern supercomputers allows complex problems, previously thought impossible, to be solved. Performance of the best supercomputers in the world is measured in Petaflops providing unprecedentedly powerful instruments for research. However, the efficient usage of all opportunities offered by modern computing systems represents a global challenge. Since applications far outlive any computer system, both scientific and commercial environments have to face up to this challenge, with all specific problems arising from the diversity of upcoming parallel architectures.

This Project addresses the urgent need for theoretical and practical technologies of an accurate and efficient design of scientific applications, highly parallel and portable methods, as well as extreme-scale parallel codes to be able to solve large problems, using the current and prospective generations of high performance systems.

The main vision of this Project is aimed at solving the issues of the efficient usage of emerging computing systems in the practice, not only in the field of computer sciences, but also in physics, chemistry, medicine and others disciplines where the computational simulations are necessary. The goal of this research is to investigate, improve and simplify the process of porting selected scientific applications for a wide range of modern computing systems with shared and distributed memory. The research hypothesis assumes that is possible to automate this process, and to provide the performance portability for a given application using both current and future architectures.

The subject of this proposal implies a profound and efficient exploration of emerging multi-/manycore architectures considering their complex, hierarchical and heterogeneous nature. Additionally, the planned HPC Benchmark Challenge will enable the performance comparison of a wide range of modern computing systems with shared- and distributed-memory models. The expected outcomes should also provide a collection of general knowledge for the process of porting parallel applications, as a guidance for optimal direction of development of scientific codes. The developed approaches can help in understanding the evolution of computing architectures, as well as providing basis for generalizing the solution for the performance portability issues. The Project results will help application developers to analyze and model the portable computing kernels, and outline them how to automate the process of porting parallel codes. The outcomes of the Project can also make suggestions for the creators of new scientific algorithms by illustrating them new programming abstractions.

The general conception of the Project realization is based on the development of methods and algorithms that allow us to provide the code and performance portability, as well as self-adaptable applications. In this work, we face the challenge of ensuring the performance portability for two complex scientific applications:

- Multidimensional Positive Definite Advection Transport Algorithm;
- parallel implementation of a numerical model of the dendritic solidification process.

It is expected that the potential of the Project results will help the developers of these applications to harness the power of the emerging platforms in the practical usage, giving an exciting opportunity for developing more complex simulations than ever before. Additionally, the expected outcomes should also provide a collection of general knowledge for the process of porting parallel applications, as a guidance for optimal direction of development of scientific codes.