DESCRIPTION FOR THE GENERAL PUBLIC (IN ENGLISH)

Effective parallel algorithms for solving partial differential equations

Accurate weather forecast, modern skyscrapers construction design, development of cheap LED lightbulbs or fuel efficient cars, and other technologically advanced products or services driven by the applied sciences — many of them rely on computationally intensive simulations of mathematical models based on partial differential equations (PDEs).

It turns out that in order to obtain accurate results one has to solve problems with huge numbers of unknowns — millions or more. The enormous size and complexity of such computation makes it necessary to use powerful supercomputers, which usually operate many processing nodes working in parallel for improved throughput and number crunching capabilities.

Solving PDEs on a parallel computer requires subtle interaction between hardware, software and specific properties of the very mathematical equation. Usually, to maintain reasonable cost of the solution (in terms of computer time and memory requirements), numerical scientists have to use fine tuned approximations, because standard methods just fail.

One of the most efficient ways to construct parallel solution algorithms for PDEs is the Domain Decomposition Method (DDM), based on the *divide and conquer* paradigm. An advantage of DDM is that concurrency is then directly built into the algorithm. Surprisingly, this simple idea requires quite a lot of mathematical thought to make it working (and it turns out, *it is working* very well in many cases). In particular, in order to make the method efficient and robust, one has to exploit subtle mathematical properties of the underlying PDE.

In this project, we will develop, analyze or apply Domain Decomposition methods to solve very large systems of equations derived from PDEs important in science and industry. These algorithms will make it possible to tackle more complex models on modern computer architectures.