A Highly Scalable Implementation of Balancing Domain Decomposition by Constraints

Javier Principe, Santiago Badia, and Alberto F. Martín

Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE),
Universitat Politècnica de Catalunya (UPC),
Parc Mediterrani de la Tecnologia, Esteve Terrades 5, 08860 Castelldefels, Spain.
e-mail: {sbadia,amartin,principe}@cimne.upc.edu

ABSTRACT

The numerical approximation of partial differential equations (PDEs) by the finite element (FE) method requires the solution of huge sparse linear systems which is only possible by efficiently exploiting exploiting current multicore-based distributed-memory machines. We do so using the Balancing Domain Decomposition by Constraints method (BDDC) for which optimal condition number bounds can be proved. It requires the solution of local problems and communication among neighboring subdomains only, which is embarrassingly parallel, but also the solution of a "small" coarse-grid problem that couples all the subdomains, providing a global mechanism for exchanging information. For large-scale problems, the coarse problem rapidly becomes the bottleneck of the algorithm and we devised three strategies to tackle this problem. **The first strategy**, developed in [1], consists in exploiting the orthogonality (with respect to energy inner product) of the coarse and fine correction spaces of the BDDC preconditioner to develop an algorithm in which the corresponding corrections are computed in parallel. Weak scalability up to 27K cores has been obtained in state-of-the-art multicore-based distributed-memory machines (HELIOS and CURIE). **The second strategy** is to use inexact/approximate solvers with linear complexity for the coarse problem. On top of the first strategy, in [2] we explore several combinations obtaining weak scalabity in the largest supercomputer in Europe (JUQUEEN) till 65K cores.

In this contribution we discuss **the third strategy**: to recursively apply the BDDC method to solve the coarse problem, resulting in a multilevel algorithm. A key aspect to efficiently implement this strategy on top of the first one, in a multilevel setting, is how to map the computations/communications at each level to MPI tasks in such a way that a high degree of overlapping is achieved. In order to reach such goal, a strategically defined hierarchy of MPI communicators is employed. Besides, in order to implement the transfer of coarse-grid problem related data among pairs of consecutive levels, say level k and $k+1$, there are as many groups as MPI tasks on level $k+1$ and each group has a root MPI task in level $k+1$ that collects (partial) contributions to the coarse problem from MPI tasks in level k. This way, coarse problem matrix and vectors are never centralized in a MPI task but always distributed among the MPI tasks of the next level, which permits to maximize multilevel parallelism. Preliminary results show scalability till 400Kcores.

REFERENCES

[1] S. Badia, A. F. Martín, J. Principe. A Highly Scalable Parallel Implementation of Balancing Domain Decomposition by Constraints. *SIAM Journal of Scientific Computing*, Vol. **36(2)**, C190–C218, 2014.

[2] S. Badia, A. F. Martín, J. Principe. On an overlapped coarse/fine implementation of balancing domain decomposition with inexact solvers. Submitted, 2014.