11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
July 20–25, 2014, Barcelona, Spain

A HIGHLY SCALABLE IMPLEMENTATION OF BALANCING DOMAIN DECOMPOSITION BY CONSTRAINTS

Javier Principe[1,2], Santiago Badia[1,2] and Alberto F. Martín[1,2]

[1] Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE), 08860 Castelldefels, Spain ({sbadia,amartin,principe}@cimne.upc.edu).
[2] Universitat Politècnica de Catalunya, 08034 Barcelona, Spain.

Key words: *parallel computing, scalability, balancing domain decomposition*

The numerical approximation of partial differential equations (PDEs) by the finite element (FE) method requires the solution of sparse linear systems with several hundreds and even thousands of millions of equations/unknowns, which is only possible by appropriately exploiting current multicore-based distributed-memory machines. A natural strategy to achieve this goal is to rely on domain decomposition (DD) methods, in which the divide and conquer principle is exploited through the solution of local problems and communication among neighboring subdomains only, resulting in highly parallel preconditioners. Two-level DD preconditioners also include a global correction in order to achieve quasi-optimal condition number bounds, i.e., independent of the number of subdomains and global problem size, for second-order coercive problems. The global correction involves the solution of a "small" coarse-grid problem that couples all the subdomains, providing a global mechanism for exchanging information. We focus on the Balancing DD by Constraints preconditioner [1] (BDDC) for which optimal condition number bounds can be proved.

However, how the scalability of optimal algorithms behaves in practice depends on a number of factors, the most important being the cost of the solution of the coarse problem, whose size increases (at best) linearly with respect to the number of subdomains. For large-scale problems, the coarse problem rapidly becomes the bottleneck of the algorithm [2]. In this contribution we discuss three strategies to tackle this problem.

The first strategy, developed in [3], 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. We propose a novel parallelization approach of BDDC based on overlapped fine-grid and coarse-grid duties in time and we present a discussion of how these novel techniques are exploited in order to reach maximum performance benefit. A weak scala-

bility study of the new parallelization approach for the 3D Poisson and linear elasticity problems up to 27K cores on structured meshes and up to 8K cores on unstructured meshes on a pair of state-of-the-art multicore-based distributed-memory machines (HELIOS and CURIE) confirms the success of this strategy.

The second strategy is to use inexact/approximate solvers, as in [5]. The BDDC preconditioner requires the solution of several local problems and the global problem and different strategies can be used for their (approximate) solution. On top of the first strategy, in [4] we explore several combinations, including a single application of an algebraic multigrid preconditioner (AMG) or even an AMG-preconditioned Krylov method with a coarse relative residual tolerance, for the approximate solution of local Dirichlet/Neumann problems and/or the global coarse-grid problem.

The third strategy is to recursively apply the BDDC method to solve the coarse problem, resulting in a multilevel algorithm. This idea has already been developed in [6] but we combine it with **the first strategy** above in a work currently under development. A key aspect to efficiently generalize this strategy 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.

REFERENCES

[1] C. R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM Journal on Scientific Computing*, Vol. **25 (1)**, 246–258, 2003.

[2] S. Badia, A. F. Martín, J. Principe. Implementation and scalability analysis of balancing domain decomposition methods. Archives of Computational Methods in Engineering, Vol. **20 (3)**, 239–262, 2013.

[3] S. Badia, A. F. Martín, J. Principe. A Highly Scalable Parallel Implementation of Balancing Domain Decomposition by Constraints. Submitted, 2013.

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

[5] O. Rheinbach, Parallel Iterative Substructuring in Structural Mechanics, Archives of Computational Methods in Engineering, Vol. **16 (4)**, 425–463, 2009.

[6] J. Mandel, B. Sousedík and C. R. Dohrmann, Multispace and multilevel BDDC, Computing, Vol. **83 (2)**, 55 – 85, 2008.