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

IMPLICIT MULTIPHYSICS SOLVERS

Santiago Badia[1,2]

[1] Centre Internacional de Mètodes Numèrics en Enginyeria (CIMNE), 08860 Castelldefels, Spain (sbadia@cimne.upc.edu; http://badia.rmee.upc.edu).
[2] Universitat Politècnica de Catalunya, 08034 Barcelona, Spain.

Key words: *Scalable solvers, Large Scale Computing, Multiphysics Problems, Domain Decomposition Methods, Preconditioning*

The simulation of many problems in computational mechanics and physics require the use of implicit solvers. Implicit solvers can deal with steady problems and can solve transient problems without the need to reduce the time step size in order to satisfy a stringent CFL condition. This last situation is particularly important when dealing with multiphysics problems involving very different time scales but one is only interested in slow time scales. Irremediably, implicit solvers require to solve linear systems. Increasingly more realistic simulations of complex physics and engineering three-dimensional phenomena end up with very large linear systems, which can easily reach 10^{10} unknowns. And some other applications, like full grid-based direct numerical simulations of plasma physics via kynetic models, are still out of the reach. The only way to deal with these large scale simulations is to make use of distributed-memory platforms in an efficient way. Current petascale supercomputers involve tens of thousands of interconnected processors. The design of robust linear solvers in these ranges and in the forthcoming exascale machines of the future pose a complicated problem in terms of mathematics and implementation issues.

In this work, we talk about the development of scalable domain decomposition preconditioners for Krylov iterative linear solvers for the solution of linear systems arising from finite element simulations of different physical problems. Domain decomposition techniques involve a partition of the global meshes into submeshes, e.g., as many submeshes (subdomains) as processors to be used. At every processor, we can solve local (Dirichlet and Neumann) problems, depending on the type of boundary conditions on the interface. However, preconditioners that only involve local (subdomain) corrections are not (weakly) scalable, i.e., as we increase the size of the global system and the number of processors in such a way that the load per processor is kept fixed, the number of iterations blow up. Scalable domain decomposition involves some type of coarse-grid correction that couples all the subdomains. This way, the method is wekly scalable but now it involves a coarse-grid correction that serializes all the subdomains and can harm practical scalability.

We consider balancing domain decomposition methods, in particular balancing domain decomposition by constraintes (BDDC) techniques. Our choice is based on the fact that BDDC preconditioners involve fine and coarse corrections that can be orthogonal in the energy norm, allowing for additional paralelization. In recent works [1], we have proposed an efficient implementation of BDDC with overlapped fine/coarse duties. Two types of processors are considered, namely fine processors associated to one subdomain and coarse processors that are used to solve the coarse problem. Hybrid MPI/OpenMP implementations that use OpenMP to extract inter-node parallelism for the coarse solver are considered. We have performed a comprehensive set of numerical experiments, using both direct and inexact (AMG-type) solvers for the fine and coarse corrections, and observed perfect weak scalability till 30,000 processors in many ranges of applicability, when dealing with coercive problems, e.g., Laplacian and linear elasticity problems. Finally, we have pushed this approach further, by using BDDC recursively, i.e. solving the coarse problem in parallel by using a BDDC preconditioner iterative solver. Based on numerical results and estimations, the overlapped/multilevel BDDC implementation (with only three-levels) can keep excellent weak scalability till hundreds of thousands of processors.

Many interesting problems are not governed by coercive (positive-definite) systems. E.g., incompressible fluid simulations are governed by indefinite (saddle-point) systems of equations. In this situations, we can apply the same preconditioners, but (in most situations) there is no math theory that proves scalability. Recent results [2] prove that the BDDC preconditioner is (at least) well-posed for this type of problems. An alternative to the monolithic approach is to use block-preconditioners (also called physics-based preconditioning in some disciplines). It is based on the definition of incomplete block-LU factorizations of the original problem that are spectrally equivalent to the original matrix but only involve one-physics (coercive) problems. The one-physics diagonal blocks can further be replaced by their BDDC preconditioners. This way, we end up with preconditioners that are both optimal and scalable. The key of this approach is the approximaton of the Schur complement matrices in the block-LU approximation such that the resulting preconditioner is spectrally equivalent. This type of preconditioners have already been developed, e.g., for the Stokes solvers.

Finally, we will use our BDDC overlapped/multilevel and block-preconditioning implementations in the simulation of realistic problems, e.g., in computational fluid dynamics and MHD simulations.

REFERENCES

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

[2] S. Badia, A. F. Martín. Balancing domain decomposition preconditioning for the discrete Stokes problem with continuous pressures. In preparation.