Martin J. Gander

Domain Decomposition Methods with Convergence Rates Faster than Multigrid

Dept. of Mathematics and Statistics
McGill University
805 Sherbrooke Street West
Montreal
QC
H3A 2K6
mgander@math.mcgill.ca

Most large scale simulations, from aircraft carriers to Jumbo-Jets, are only possible with parallel computers. Often codes are available for partial problems, like the wing or the engine of the plane. A natural paradigm to combine such codes to simulate the entire aircraft is to use domain decomposition techniques. But the performance of these techniques depends very much on the strength of the physical coupling between the pieces of the model.

The Schwarz methods are a class of domain decomposition methods. They are based on a theoretical tool to prove existence and uniqueness, invented by Schwarz in 1869, have been investigated in detail over the last two decades and are now well understood, see for example the survey articles by Chan and Mathew, Xu, or the book by Smith, Bjørstad, and Gropp, and the book by Quarteroni and Valli. Optimal convergence results exist for the Schwarz methods in the sense that the condition number of the preconditioned system is independent of (or only weakly dependent on) the mesh parameter and the number of subdomains. Thus asymptotically Schwarz methods have optimal scalability.

These optimality results contain however constants which remain unknown in the analysis. Thus they do not imply that the current Schwarz methods have optimal performance. They do not guarantee either that domain decomposition methods are competitive to other parallel methods. Thus the word "optimal" can be misleading. Indeed a comparison by Gaertel and Kessel in 1992 of an "optimal" domain decomposition method with a simple multi-grid algorithm implemented in parallel showed that, although the domain decomposition algorithm scales optimally, the parallelized multi-grid algorithm is more than an order of magnitude faster on 9 and 16 processors for an elliptic model problem. So why are the Schwarz methods so slow ? Our analysis reveals that the main reason are the transmission conditions employed at the artificial interfaces between the subdomains. The classical Schwarz algorithm uses Dirichlet transmission conditions to exchange information between subdomains, which is fatal for the performance. But fast convergence was not of interest to Schwarz, who used his tool only to prove existence and uniqueness of solutions, and not to actually compute them.

Instead of Dirichlet transmission conditions one should use transmission conditions which decouple the subdomain problems as much as possible. And precisely this is possible: there are well known boundary conditions to truncate infinite domains when the corresponding problem needs to be solved on a finite computer, namely the transparent or absorbing boundary conditions. They decouple as well as possible the exterior problem (which is then not even solved) from the interior, computational one. Fundamental contributions for hyperbolic problems can be found in an early paper by Enquist and Majda from 1977, and for the case of parabolic problems in the work by Halpern in 1986. Optimized Schwarz methods use, instead of the classical transmission conditions of Dirichlet type, absorbing or approximately absorbing transmission conditions. The basic algorithm stays the same, only the information which is exchanged between subdomains is replaced by physically more valuable information, which decouples the subdomain problems much more effectively. The impact of this small change on the performance is dramatic. If one uses absorbing boundary conditions to exchange information, optimal convergence results limited only by the physics of the problem can be achieved: convergence is reached in a finite number of steps related to the number of subdomains used, a result first derived for elliptic problems and a special decomposition by Nataf et al. in 1995. But simple local approximations to the absorbing boundary conditions suffice already to speed up the algorithms by orders of magnitudes.

Optimized Schwarz methods are designed to weaken the coupling between subdomain problems, even if the physical coupling is strong, and use transmission conditions which take the physics of the underlying problem into account. I will present in this talk an analysis and numerical experiments for a symmetric positive definite model problem to illustrate the dramatic change in the convergence rate of optimized Schwarz methods compared to classical ones. The optimized methods converge often with an order of magnitude less iterations than classical Schwarz methods at the same cost per iteration, and attain contraction rates which are comparable to those of multigrid methods.