Coarse Space Components of Domain Decomposition Algorithms

Olof Widlund
Courant Institute of Mathematical Sciences
New York University
http://www.cs.nyu.edu/cs/faculty/widlund/

Domain Decomposition Methods

Consider a finite element approximation of, e.g., the equations of linear elasticity or a self-adjoint scalar elliptic problem.

The domain Ω is subdivided into non-overlapping subdomains (substructures) Ω_i. In between the interface Γ. Each subdomain is the union of elements of the finite element triangulation.

Two families: the *iterative substructuring algorithms*, using solvers on the Ω_i, each with thousands of degrees of freedom, and the *overlapping Schwarz methods*, using solvers on an overlapping set of subdomains Ω_i', often obtained by adding layers of elements to the Ω_i.

The preconditioner of the finite element problem are built from these solvers and all preconditioners considered also include a coarse, global solver with a few degrees of freedom for each subdomain. A Krylov space method - conjugate gradients or GMRES - is always used to accelerate the convergence.

In the beginning, coarse spaces were not used. Only the continuous problems were considered and then it unclear what a coarse problem might be. The algorithms based on overlapping subdomains were considered by Schwarz (1870), Sobolev (1936), and Babuška (1957).

Algorithms with non-overlapping subdomains were considered in a Poincaré-Steklov framework by Agoshkov, Lebedev, and Quarteroni and others. An important early paper is by Dryja (1981), which provides an optimal preconditioner. The theory, as far as I know, was only for interfaces without cross points or edges. Therefore, the decompositions were essentially only into strips.

The introduction of a second level dates back to the mid-eighties. Four papers by Bramble, Pasciak, and Schatz (Math Comp 1986-1989) were crucial for the development of theory of iterative substructuring; these papers all existed in preprint form by the time of DD1 (January 1987).

In the first of these papers, on problems in two dimensions, the substructures are triangles and the coarse space the piece-wise linears on this coarse triangulation; cf. geometric multigrid. There is one local space for each of the edges of the interface. A

$$C(1 + \log(H/h))^2$$

bound was established for the condition number of the algorithm.

This result are obtained, subdomain by subdomain. C is therefore independent of the number of subdomains and the result is also valid uniformly for any scalar problem

$$-\,\mathrm{div}(a(x)\,\mathbf{grad})u(x) = f(x),$$

where $a(x) = a_i, x \in \Omega_i$ where the a_i are arbitrary, positive constants.

An important tool is a finite element Sobolev inequality, valid for plane domains,

$$\|u_h\|^2_{L^\infty(\Omega_i)} \leq C(1 + \log(H/h))\|u_h\|^2_{H^1(\Omega_i)}.$$

This bound cannot be improved (Brenner); it is a genuine finite element result.

Before discussing the fourth of the BPS papers, it makes sense to look at the geometry of the decomposition of a domain in three dimensions. There is an $interface\ \Gamma$ containing the finite element nodes which belong to the closure of at least two subdomains. The interface is decomposed into $faces$, $edges$, and $vertices$: the nodes on a face F^{ij} belong to a pair of subdomains Ω_i and Ω_j, edges and vertices make up the boundary of faces with edges typically common to at least three subdomains, and vertices are end points of edges.

Such decompositions make sense even for quite irregular subdomains, such as those delivered by mesh partitioners, and each of these objects is defined by an equivalence class of nodes with a common set of subdomain indices. For many iterative substructuring methods as well as for some methods based on overlapping decompositions, there are basis functions of coarse spaces directly associated with these geometric objects.

The union of the edges and vertices of the interface is the $wire\ basket$. Individual subdomains also have wire baskets.

BPS IV concerns wire basket algorithms. Instead of working with a conventional coarse space, for which to this day, strong results independent on the values of the a_i have not been derived, the unknowns of the coarse problem are the average values over the subdomain wire baskets. A full description of such an algorithm is quite complicated and is not given.

A version of the BPS IV algorithms was developed and analyzed in the 1990 PhD thesis of Barry Smith. Smith also implemented it (SISC 1992) on parallel computers prior to moving on to the development of PETSc. He also took the initiative to a joint project with Dryja and W., which led to the development and analysis of a large number of $primal\ iterative\ substructuring$ algorithms (SINUM 1994).

The analysis was carried out in an abstract Schwarz framework, which had its roots in the contribution of Pierre-Louis Lions to DD1. All bounds are of the form $C(1 + \log(H/h))^2$, with a few exceptions, and most are independent of coefficient jumps. Smith also wrote a fine book with Bjørstad and Gropp.

It deserves to be mentioned, that there were other relevant contributions at DD1, in particular work by Glowinski and Wheeler, with algorithms which resembles one-level FETI methods. The importance of this work has been overlooked; see, however, Chapter 1 of T. & W.

By the time of DD2, two-level $additive\ Schwarz$ methods had been developed and proven to be optimal and scalable (independent of the number of subdomains) for problems with constant coefficients; cf. Dryja's paper in DD2 proceedings. At first, generous overlap was assumed but the methods soon turned out to work best with modest overlap. This led to an analysis of the case of small overlap and the bound

$$\kappa(T_{as}) \leq C(1 + H/\delta),$$

shown to be best possible by Brenner; see (SISC 1994 & 2002). Already at the time of DD3, it was realized that these and the iterative substructuring algorithms could be analyzed in a common framework.

The two-level overlapping Schwarz methods require two communication steps per iteration. One of them can be eliminated resulting in $restricted\ additive\ Schwarz$ methods, invented by Cai and Sarkis (SISC 1999). These algorithms have been studied extensively and they also typically require fewer iteration steps.

Already at the time of DD1, it was realized that a coarse component, which provides at least a minimal amount of global transfer of information across the entire domain, is required to obtain bounds which are independent of the number of subdomains.

This comes as no surprise to any student of multigrid. What makes the two families different, is that only two levels are required for a domain decomposition method even for very large problems, which limits the number of communication steps. The great repertoire of coarse spaces adds to this promise.

While many of the domain decomposition methods have been developed using exact, Cholesky solvers, there is now an active interest in the development and analysis of methods based on inexact solvers, such as multigrid. In parallel, there has been important progress on domain decomposition methods for more than two levels, in particular by Tu and by Klawonn and Rheinbach. With the relative cost of communication increasing with the number of processors of massively parallel computing systems, the domain decomposition algorithms appear to be quite competitive

The extension of any domain decomposition developed for scalar elliptic problems to the equations of linear elasticity requires a modification of the coarse space to accomodate the larger null space for problems with natural boundary conditions. There are then six rigid body modes instead of a single constant. In many cases, this work is relatively routine, see Chapter 8 of T. & W. A successful approach is often to construct an interpolation operator which reproduces all rigid body modes and which also can be bounded uniformly or with a factor $C(1+\log(H/h))$. Examples of this type of activity is given in two papers by Pavarino and W., (SINUM 2000).

This requirement was formalized by Mandel in terms of the $null\ space\ condition$ (CMAME 1990) and it is also explained well in the book by Smith, Bjørstad, and Gropp.

Dual Iterative Substructuring Methods

These important domain decomposition algorithms date back at least to DD2 and a paper by Bourgat, Glowinski, Le Tallec, and Vidrascu. This development led to the development of *balancing Neumann-Neumann* methods with coarse space components.

An important role in the description and analysis of the Neumann-Neumann algorithms is played by a family of weighted counting functions δ_i, which are associated with the individual $\partial\Omega_i$. They are defined for $\gamma \in [1/2, \infty)$ by a sum of contributions from Ω_i and its relevant next neighbors,

$$\delta_i(x) = \frac{\sum_{j \in N_x} a_j^\gamma}{a_i^\gamma}, \quad x \in \partial\Omega_{i,h} \cap \Gamma_h,$$

where N_x is the set of indices j of the subregions such that $x \in \partial\Omega_{j,h}$.

Their pseudoinverses $\delta_i^\dagger$ are defined by

$$\delta_i^\dagger(x) = (\delta_i(x))^{-1}, \quad x \in \partial\Omega_{i,h} \cap \Gamma_h.$$

They provide a partition of unity:

$$\sum_i R_i^T \delta_i{}^\dagger(x) \equiv 1, \quad x \in \Gamma_h,$$

for any Ω_i such that $\partial\Omega_i \cap \partial\Omega_D = \emptyset$. These functions span the coarse space of the algorithm. Note that each of them can be written as the sum of face, edge, and vertex functions associated with the decomposition of the interface.

The local space V_i consists of discrete harmonic functions with non-zero interface values only on $\partial\Omega_i$. A scaled Neumann problem given by the bilinear form

$$\tilde{a}_i(u,v) = a_i \int_{\Omega_i} \nabla(\delta_i u) \cdot \nabla(\delta_i v) dx,$$

is used to define the local parts of a hybrid Schwarz method.
$C(1+\log(H/h))^2$ bounds have been established with C independent of the number of substructures and of jumps in the coefficients across the interface around (1995). These algorithms have proven very successful and have been used extensively. They have been modified for linear elasticity.

What is now called *one-level FETI* methods were introduced by Farhat and Roux around 1990 and first analyzed by Mandel and Brezina (Math Comp 1996). The analysis has also been further refined. Instead of describing these methods, the more recent FETI–DP and BDDC algorithms will now be considered.

FETI–DP and BDDC

These more recent methods only lead to positive definite subproblems. They are defined in terms a set of primal continuity constraints which are satisfied throughout the iteration. A pair of FETI–DP and BDDC preconditioned systems have essentially identical spectra if they employ the same primal constraints. Here is a figure:

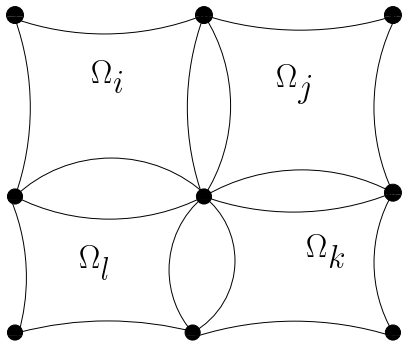

Figure 1: Decomposition of subdomains for a FETI–DP method.

The primal constraints in this case make the values at the subdomain vertices global. Note that we obtain multiple values at all other nodes on the interface. The partially subassembled stiffness matrix of this alternative finite element model is used to define the preconditioners.

In a FETI–DP algorithm, the continuity at the edge nodes is enforced by using Lagrange multipliers and the rate of convergence is enhanced by solving Dirichlet problems on each subdomain in each iteration. The conjugate gradient algorithm is used to find the correct values of the Lagrange multipliers.

In a BDDC algorithm, continuity is instead restored in each step by computing weighted averages across the interface. This leads to non-zero residuals at some nodes interior to the subdomains, and in each iteration, these residuals are eliminated by using subdomain Dirichlet solves.

We can thus think of the primal constraints as providing a global component of these preconditioners.

For problems in three dimensions, primal variables associated with point constraints alone do not lead to competitive algorithms; this is technically closely related to the issues raised in early studies of primal iterative substructuring methods. Instead, or in addition, averages (and moments) over faces or, preferably edges, should have common values across the interface.

The selection a small and effective set of primal constraints for elasticity problems with large jumps in the Lamé parameters has been very challenging, see Klawonn and W. (CPAM 2006). The resulting recipes have proven very successful for very difficult problems, see Klawonn and Rheinbach.

Additional Purposes for Coarse Spaces

In work on incompressible Stokes, almost incompressible elasticity, and Maxwell's equation, the choice of coarse spaces require additional care.

By the divergence theorem, a divergence-free extension of boundary data is only possible if there is a zero net flux across the boundary. If for a overlapping Schwarz method for almost incompressible elasticity a coarse component $\mathbf{u}_0$ of a given $\mathbf{u}$ can be chosen with same net fluxes across subdomain boundaries, then the remainder, $\mathbf{w} := \mathbf{u} - \mathbf{u}_0$, allow a divergence free extension and a successful decomposition of $\mathbf{w}$ into local components. These ideas have been explored repeatedly for balancing Neumann-Neumann, FETI-DP, and BDDC algorithms.

For Maxwell's equation, curl-free extension are desirable for very similar reasons. This is being explored currently in joint work by Dohrmann and W. for problems in two dimensions.