
The Development of Coarse Spaces for Domain Decomposition Algorithms

Olof B. Widlund¹

Courant Institute, 251 Mercer Street, New York, NY 10012, USA
widlund@cims.nyu.edu

Summary. The importance of using coarse components, and thus at least one additional level, in the design of domain decomposition methods has been understood for at least twenty years. For many problems of interest, such a device, which provides at least a minimal amount of global transfer of information in each step, is necessary in order to obtain convergence rates which are independent of the number of subdomains. An historical overview, colored by the scientific history of its author, is given of the development of such coarse components of the domain decomposition algorithms. These algorithms are all preconditioned conjugate gradient methods or they are accelerated by using some alternative Krylov space method. The preconditioners are built from solvers of the given problem restricted to subdomains and a coarse approximation which often can be quite exotic.

1 Introduction

We will consider finite element approximations of, e.g., a self-adjoint scalar elliptic problem or the equations of linear elasticity. The domain Ω of the partial differential equation is subdivided into non-overlapping subdomains (substructures) Ω_i ; there can be very many of them, in particular, when massively parallel computer systems are employed. Between the subdomains, we have the interface Γ ; Γ_h is its set of finite element nodes. Each subdomain is the union of elements of the finite element triangulation.

There are two main families of domain decomposition algorithms: the *iterative substructuring algorithms*, using solvers of the finite element problems restricted to the Ω_i , each often with tens of thousands degrees of freedom, and the *overlapping Schwarz methods*, using solvers on a set of overlapping subdomains Ω'_i , often obtained by adding layers of elements to the individual Ω_i 's. Exact solvers are often used to solve these local problems as in much of traditional finite element practice.

The preconditioner of the finite element problem 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.

Early on, coarse spaces were not used and only continuous problems were considered; in fact it is unclear what a coarse problem then might be. Algorithms based

on overlapping subdomains were considered by Schwarz [1890], Sobolev [1936], and Babuška [1957] and algorithms with non-overlapping subdomains, in a Poincaré-Steklov framework, by Agoshkov [1988], Agoshkov and Lebedev [1985] and Quarteroni and Valli [1999].

An important early paper by Dryja [1981] provides an optimal preconditioner for a case of two subdomains and an elliptic finite difference problem. It was followed, in the 1980'ies, by a number of other studies for problems where the interfaces are without cross points (or cross edges), i.e., all finite element nodes on the interface are common to the boundary of only two subdomains; the decompositions of the domains were into strips.

2 Early Two-level Domain Decomposition Methods

The successful introduction of a second, coarse level dates to the mid-eighties. In particular, the first and fourth paper in a series of four, Bramble et al. [1986, 1989] have been crucial for the development of the theory of iterative substructuring methods for more general decompositions; these papers existed at least in preprint form by the time of the first international conference on domain decomposition methods, DDI, held in Paris in January 1987. Already at this time, it was realized that a coarse component, to provide 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.

In the first of these papers, on problems in two dimensions, the substructures are triangles and the coarse space is spanned by continuous, piece-wise linear functions on this coarse triangulation in a set-up resembling that of geometric multigrid. There is one local space for each of the edges of the interface. A $C(1 + \log(H/h))^2$ bound is established for the condition number of the preconditioned operator; here H is a typical subdomain diameter and h that of the finite element triangles. These logarithmic factors arise when we divide the trace of finite element functions on the interface into a sum of functions with a nonzero trace only on one edge.

The main result is obtained in an analysis for one subdomain at a time. As a consequence, the constant C is independent of the number of subdomains and the result is valid uniformly for any scalar problem

$$-div(a(x)grad)u(x) = f(x),$$

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

An important tool, used in this work, is a finite element Sobolev inequality, valid for plane domains:

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

This is a genuine finite element result and Brenner and Sung [2000] have shown that this result cannot be improved.

Before discussing Bramble et al. [1989], which is regarded as the most important in the series, we consider the geometry of the decomposition of a domain in three dimensions. The interface Γ contains all the finite element nodes which belong to the closure of at least two subdomains. It 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 can be defined even for quite irregular subdomains, such as those delivered by mesh partitioners. Each of these geometric objects can be defined in terms of an equivalence class of nodes with a common set of subdomain indices. For many iterative substructuring methods, as well as for some more recent methods based on overlapping decompositions, there are basis functions of coarse spaces directly associated with these geometric objects: $\theta_{F^{ij}}$, $\theta_{E^{ik}}$, and $\theta_{V^{i\ell}}$. They are defined by the value 1 on the set of nodes in question and vanish at all other nodes on Γ and they are discrete harmonic, i.e., the values inside the subdomains give a minimal energy extension. Therefore, they form a partition of unity for any subdomain which is interior to Ω .

The union of the edges and vertices of the interface in three dimensions is known as the *wire basket* and individual subdomains also have wire baskets. Bramble et al. [1989] concerns wire basket algorithms. Instead of working with a conventional coarse space, for which, to this day, no strong results, independent of the values of the a_i , have been derived for three dimensions, the coarse space functions are given in terms of their values on the wire basket of the subdomains. The values on a face are then given in terms of the values on its boundary; this establishes continuity across Γ . A corresponding interpolation operator, into the coarse space, will reproduce constants. Technically, this coarse space is of large dimension, but this is compensated for by using a particular inexact solver in terms of one variable per subdomain, namely the average values over the subdomain wirebaskets. The values at the subdomain nodes are then computed locally. It is interesting to note that these first successful algorithms of this kind are also among the most complicated.

A version of these algorithms is developed and analyzed in the 1990 PhD thesis of Barry Smith. It was also implemented on parallel processors, see Smith [1993]. Smith then moved on to the development of PETSc. He also took the initiative to a joint project with Dryja and this author, which led to the development and analysis of a large number of *primal iterative substructuring* algorithms, see Dryja et al. [1994]. The analysis in that paper is carried out in an *abstract Schwarz* framework, which has its roots in a DD1 contribution of Pierre-Louis Lions [1988]. All bounds are, with a few exceptions, of the form $C(1 + \log(H/h))^2$ and most of them are independent of coefficient jumps. Smith also later wrote a pioneering book, see Smith et al. [1996].

Another important contribution at DD1, is a paper by Glowinski and Wheeler [1988]; their algorithms resembles one-level FETI methods. The importance of this work has been overlooked; see, however [Toselli and Widlund, 2005, Section 1.3.5].

By the time of DD2, the first two-level *additive Schwarz* methods had been developed and shown to be optimal and scalable, i.e., with convergence rates independent of the number of subdomains, for problems with moderately varying coefficients; cf. Dryja [1989]. These preconditioners are built from solvers on the set of overlapping subdomains and a conventional coarse spaces just as that of Bramble et al. [1986]. At first, a generous overlap was assumed but the methods work most efficiently 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 Dryja and Widlund [1994] and Brenner [2000]. Here δ measures the overlap between neighboring subdomains.

Already at the time of DD3, it was realized that these and the iterative substructuring algorithms could be analyzed in a common *abstract Schwarz framework*; i see [Toselli and Widlund, 2005, Chapter 2].

3 Additional Comments

It comes as no surprise to any student of multigrid that a global component of the preconditioner is very important. What makes the two families different, is that only two levels are required for a domain decomposition method even for very large problems. This limits the number of communication steps. 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 [1999]. These algorithms have been studied extensively and they also typically require fewer iterations.

The great repertoire of coarse spaces has made it possible to develop fast methods with convergence rates independent of even large jumps in the material properties across the interface. It is also easy to extend the overlapping Schwarz methods to more than two levels and progress has also been made recently on introducing additional levels for the iterative substructuring methods of Section 4; see in particular Tu [2007] and Klawonn and Rheinbach [2007]. This work is increasingly relevant for very large problems and massively parallel computing systems for which the coarse space will be of very large dimension and present a bottle neck.

The extension of any domain decomposition developed for scalar elliptic problems to the equations of linear elasticity requires a modification of the coarse spaces to accommodate the larger null space for these problems; in three dimensions, there are six rigid body modes of zero energy instead of a single constant. This *null space condition* for the coarse space was formalized in Mandel [1990] and it is also explained well in Smith et al. [1996]. In many cases, this work is relatively routine, see [Toselli and Widlund, 2005, Chapter 8]. A successful approach begins by constructing an interpolation operator, which reproduces all rigid body modes, and with an energy that can be bounded uniformly or with a factor $C(1 + \log(H/h))$.

4 Other Iterative Substructuring Methods

Other important domain decomposition algorithms date back at least to DD2, see Bourgat et al. [1989]. This development led to *balancing Neumann-Neumann* methods with coarse space components; cf. [Toselli and Widlund, 2005, Section 6.2]. An important role in the description and analysis of the Neumann-Neumann algorithms is played by a family of weighted counting functions δ_i^\dagger , which are associated with the individual $\partial\Omega_i$. They are defined, for $\gamma \in [1/2, \infty)$, by a sum of contributions from the coefficient in Ω_i and its relevant next neighbors;

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

Here N_x is the set of indices j of the subregions such that $x \in \partial\Omega_{j,h}$. A subscript h denotes the set of nodes on the set in question. These functions 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 $\partial\Omega_i \cap \partial\Omega_D = \emptyset$, and they span the coarse space of the algorithm. Here R_i^T provides an extension by zero to the nodes of $\Gamma \setminus \partial\Omega_i$. If the coefficients are constant in each subdomain, each of the δ_i^\dagger can be written as the linear combination of the face, edge, and vertex functions of the interface.

The local space V_i for the balancing methods have of 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 and $C(1 + \log(H/h))^2$ bounds can be 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, in a modified form, for problems of elasticity.

What is now called the *one-level FETI* methods were introduced in Farhat and Roux [1991] and first analyzed in Mandel and Brezina [1996]. Instead of describing these methods, we will now consider the more recent FETI-DP and BDDC algorithms.

5 FETI-DP and BDDC

The FETI-DP methods were introduced in Farhat et al. [2001, 2000] and the BDDC methods in Dohrmann [2003]. These more recent methods only require the solution of positive definite problems. They are defined in terms a set of *primal* continuity constraints which are satisfied throughout the iteration; see Figure 1. A pair of FETI-DP and BDDC preconditioned systems have essentially identical spectra if they employ the same primal constraints.




Fig. 1. Decomposition of subdomains for FETI-DP and BDDC methods.

The primal constraints in this case make the values at the subdomain vertices global, while 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. It is important to note that a linear system of equations of this kind will have a positive definite matrix and that it can be solved at a cost which is much smaller than that for a fully assembled matrix.

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

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.

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 Widlund [2006]. The resulting recipes have proven successful for very difficult problems, see Klawonn and Rheinbach [2006]. In spite of the seemingly different coarse components of these algorithms, the tools of analysis are essentially the same as for the older iterative substructuring methods.

6 Additional Roles for Coarse Spaces

In work on incompressible Stokes, almost incompressible elasticity, and Maxwell's equations, 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 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 interface values of the remainder, $\mathbf{w} := \mathbf{u} - \mathbf{u}_0$, will allow for 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; see e.g., Li and Widlund [2006] and more recently for overlapping Schwarz methods which use coarse spaces borrowed from primal substructuring methods. Taking account of the net flux across the subdomain boundaries is a necessity, for almost incompressible elasticity, since we have to make sure that a divergence free function can be partitioned into components in the same class; otherwise the energy of these local components would greatly exceed that of the given function. For Maxwell's equation, curl-free extension are desirable for very similar reasons.

References

- Valeri I. Agoshkov. Poincaré-Steklov operators and domain decomposition methods in finite dimensional spaces. In Roland Glowinski, Gene H. Golub, Gérard A. Meurant, and Jacques Périaux, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, Philadelphia, PA, 1988. SIAM.
- Valeri I. Agoshkov and Vyacheslav I. Lebedev. Poincaré-Steklov operators and the methods of partition of the domain in variational problems. In Guri I. Marchuk, editor, *Computational Processes and Systems*, pages 173–227, Moscow, 1985. Nauka. In Russian.
- Ivo Babuška. Über Schwarzsche Algorithmen in partiellen Differentialgleichungen der mathematischen Physik. *ZAMM*, 37(7/8):243–245, 1957.
- Jean-François Bourgat, Roland Glowinski, Patrick Le Tallec, and Marina Vidrascu. Variational formulation and algorithm for trace operator in domain decomposition calculations. In Tony Chan, Roland Glowinski, Jacques Périaux, and Olof Widlund, editors, *Domain Decomposition Methods. Second International Symposium on Domain Decomposition Methods*, pages 3–16, Philadelphia, PA, 1989. SIAM. Los Angeles, California, January 14–16, 1988.
- James H. Bramble, Joseph E. Pasciak, and Alfred H. Schatz. The construction of preconditioners for elliptic problems by substructuring, I. *Math. Comp.*, 47(175):103–134, 1986.
- James H. Bramble, Joseph E. Pasciak, and Alfred H. Schatz. The construction of preconditioners for elliptic problems by substructuring, IV. *Math. Comp.*, 53(187):1–24, 1989.
- Susanne C. Brenner. Lower bounds of two-level additive Schwarz preconditioners with small overlap. *SIAM J. Sci. Comput.*, 21(5):1657–1669, 2000.
- Susanne C. Brenner and Li-yeng Sung. Discrete Sobolev and Poincaré inequalities via Fourier series. *East-West J. of Numer. Math*, 8:83–92, 2000.
- Xiao-Chuan Cai and Marcus Sarkis. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM J. Sci. Comp.*, 21:239–247, 1999.
- Clark R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci Comput.*, 25(1):246–258, 2003.
- Maksymilian Dryja. An algorithm with a capacitance matrix for a variational-difference scheme. In Guri I. Marchuk, editor, *Variational-Difference Methods in Mathematical Physics*, pages 63–73, Novosibirsk, 1981. USSR Academy of Sciences.
- Maksymilian Dryja. An additive Schwarz algorithm for two- and three-dimensional finite element elliptic problems. In Tony Chan, Roland Glowinski, Jacques Périaux, and Olof Widlund, editors, *Domain Decomposition Methods. Second International Symposium on Domain Decomposition Methods*, pages 168–172, Philadelphia, PA, 1989. SIAM. Los Angeles, California, January 14–16, 1988.
- Maksymilian Dryja, Barry F. Smith, and Olof B. Widlund. Schwarz analysis of iterative substructuring algorithms for elliptic problems in three dimensions. *SIAM J. Numer. Anal.*, 31(6):1662–1694, December 1994.
- Maksymilian Dryja and Olof B. Widlund. Domain decomposition algorithms with small overlap. *SIAM J. Sci. Comput.*, 15(3):604–620, May 1994.
- Charbel Farhat, Michel Lesoinne, Patrick Le Tallec, Kendall Pierson, and Daniel Rixen. FETI-DP: A dual-primal unified FETI method – part I: A faster alter-

- native to the two-level FETI method. *Internat. J. Numer. Methods Engrg.*, 50:1523–1544, 2001.
- Charbel Farhat, Michel Lesoinne, and Kendall Pierson. A scalable dual-primal domain decomposition method. *Numer. Linear Algebra Appl.*, 7(7–8):687–714, 2000.
- Charbel Farhat and Francois-Xavier Roux. A Method of Finite Element Tearing and Interconnecting and its Parallel Solution Algorithm. *Int. J. Numer. Meth. Engrg.*, 32:1205–1227, 1991.
- Roland Glowinski and Mary F. Wheeler. Domain decomposition and mixed finite element methods for elliptic problems. In Roland Glowinski, Gene H. Golub, Gérard A. Meurant, and Jacques Périaux, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, pages 144–172, Philadelphia, PA, 1988. SIAM. Paris, France, January 7–9, 1987.
- Axel Klawonn and Oliver Rheinbach. A parallel implementation of Dual-Primal FETI methods for three dimensional linear elasticity using a transformation of basis. *SIAM J. Sci. Comput.*, 28(5):1886–1906, 2006.
- Axel Klawonn and Oliver Rheinbach. Inexact FETI-DP methods. *Internat. J. Numer. Methods Engrg.*, 69:284–307, 2007.
- Axel Klawonn and Olof B. Widlund. Dual-Primal FETI methods for linear elasticity. *Comm. Pure Appl. Math.*, 59(11):1523–1572, November 2006.
- Jing Li and Olof B. Widlund. BDDC algorithms for incompressible Stokes equations. *SIAM J. Numer. Anal.*, 44(6):2432–2455, 2006.
- Pierre-Louis Lions. On the Schwarz alternating method. I. In Roland Glowinski, Gene H. Golub, Gérard A. Meurant, and Jacques Périaux, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, pages 1–42, Philadelphia, PA, 1988. SIAM. Paris, France, January 7–9, 1987.
- Jan Mandel. Iterative solvers by substructuring for the p-version finite element method. *Comp. Methods Appl. Mech. Engrg.*, 80:117–128, 1990.
- Jan Mandel and Marian Brezina. Balancing domain decomposition for problems with large jumps in coefficients. *Math. Comp.*, 65(216):1387–1401, 1996.
- Alfio Quarteroni and Alberto Valli. *Domain Decomposition Methods for Partial Differential Equations*. Oxford Science Publications, 1999.
- Hermann A. Schwarz. *Gesammelte Mathematische Abhandlungen*, volume 2, pages 133–143. Springer, Berlin, 1890. First published in *Vierteljahrsschrift der Naturforschenden Gesellschaft in Zürich*, volume 15, 1870, pp. 272–286.
- Barry F. Smith. A parallel implementation of an iterative substructuring algorithm for problems in three dimensions. *SIAM J. Sci. Comput.*, 14(2):406–423, March 1993.
- Barry F. Smith, Petter E. Bjørstad, and William Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, 1996.
- Sergey L. Sobolev. L’Algorithme de Schwarz dans la Théorie de l’Elasticité. *Comptes Rendus (Doklady) de l’Académie des Sciences de l’URSS*, IV((XIII) 6):243–246, 1936.
- Andrea Toselli and Olof Widlund. *Domain Decomposition Methods - Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin Heidelberg New York, 2005.
- Xuemin Tu. Three-level BDDC in two dimensions. *Internat. J. Numer. Methods Engrg.*, 69:33–59, 2007.