MathematicS

MathS In A.

In Action

Victorita Dolean, Frédéric Hecht, Pierre Jolivet,
Frédéric Nataf & Pierre-Henri Tournier

Recent Advances in Adaptive Coarse Spaces and Availability in Open Source Libraries

Volume 11 (2022), p. 61-71.

https://doi.org/10.5802/msia.17

MathematicS In Action *est membre du*
Centre Mersenne pour l'édition scientifique ouverte
http://www.centre-mersenne.org/
e-ISSN : 2102-5754

MathematicS In Action
Vol. 11, 61-71 (2022)

Recent Advances in Adaptive Coarse Spaces and Availability in Open Source Libraries

Victorita Dolean [*]
Frédéric Hecht [**]
Pierre Jolivet [***]
Frédéric Nataf [†]
Pierre-Henri Tournier [‡]

[*] Université Côte d'Azur-CNRS- Labo Dieudonné, Nice, France
E-mail address: victorita.dolean@univ-cotedazur.fr
[**] Sorbonne Université, Université de Paris, Inria Equipe Alpines, Laboratoire Jacques-Louis Lions, F-75005 Paris, France
E-mail address: frederic.hecht@sorbonne-universite.fr
[***] IRIT, CNRS, F-31071 Toulouse, France
E-mail address: pierre.jolivet@enseeiht.fr
[†] Sorbonne Université, Université de Paris, Inria Equipe Alpines, Laboratoire Jacques-Louis Lions, F-75005 Paris, France
E-mail address: frederic.nataf@sorbonne-universite.fr
[‡] Sorbonne Université, Université de Paris, Inria Equipe Alpines, Laboratoire Jacques-Louis Lions, F-75005 Paris, France
E-mail address: pierre-henri.tournier@sorbonne-universite.fr.

Abstract

Scalability of parallel solvers for problems with high heterogeneities relies on adaptive coarse spaces built from generalized eigenvalue problems in the subdomains. The corresponding theory is powerful and flexible but the development of an efficient parallel implementation is challenging. We report here on recent advances in adaptive coarse spaces and on their open source implementations.

1. Introduction

All domains of physical sciences, engineering or medicine make use of scientific computing. For instance more and more imaging methods in earth science or in medicine rely on the direct simulation of the data acquisition process when solving the inverse problems. For earth science applications, the volume of data is naturally very large. In medical imaging, nearly real time methods are demanded. These requirements can only be met with parallel computers and algorithms designed for these new computing architectures.

As an example, we consider the modeling of a microwave imaging system, for the detection and monitoring of brain strokes. The prototype in Figure 1.1 was developed by the company EMTensor GmbH and studied in the framework of the ANR project MEDIMAX. The data acquired with this device are used as input for an inverse problem associated with the time-harmonic Maxwell's equations, which makes it possible to estimate the complex electric permittivity of the brain tissues of a patient affected by a stroke. Indeed, a stroke results in a variation of the complex electric permittivity inside a region of the brain, thus it can be detected and monitored by clinicians thanks to an image of the brain displaying the values of this property. Taking

This work was granted access to the HPC resources of OCCIGEN@CINES under the allocation 2020-067730 granted by GENCI..
Keywords: Example, Applied mathematics, Journal.
2020 *Mathematics Subject Classification:* 00X99.

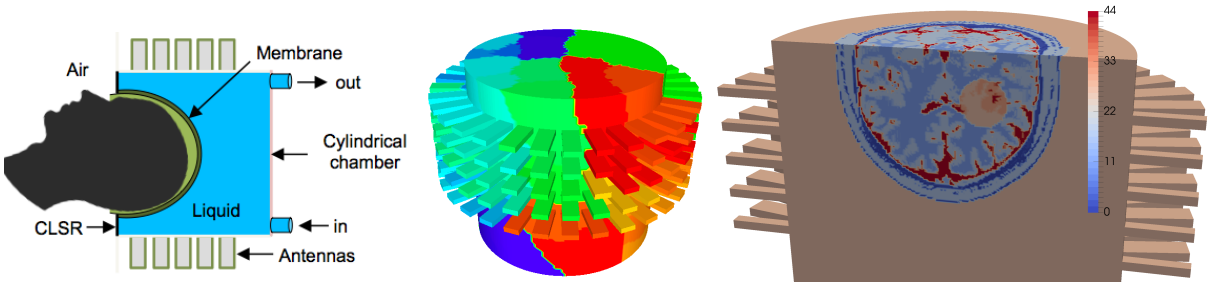

FIGURE 1.1. The microwave imaging system prototype developed by EMTensor GmbH, the decomposed computational domain, and the imaginary part of the relative complex permittivity of a virtual head model immersed in the imaging chamber, with a simulated ellipsoid-shaped stroke.

advantage of several levels of parallelism, in [42] we met the requirement of a nearly real time imaging with an elapsed time of two minutes when running on 2048 cores.

The paper is organized as follows. In Section 2, we recall the original one-level Schwarz method and in Section 3 the need for two-level methods in order to reach scalability. In Section 4, we introduce the GenEO method and point to several of its extensions. In Section 5, some large scale computations are displayed.

2. One level Schwarz methods

H.A. Schwarz was a German analyst of the 19th century. He was interested in proving the existence and uniqueness of the Poisson problem. At his time, there were no Sobolev spaces nor Lax–Milgram theorem. The only available tool was the Fourier transform, limited by its very nature to simple geometries. In order to consider more general situations, H.A. Schwarz devised an iterative algorithm for solving Poisson problems set on a union of simple geometries, see [38]. For a historical presentation of these kind of methods see [15].

Let the domain Ω be the union of a disk and a rectangle, see Figure 2.1. Consider the Poisson problem which consists in finding $u : \Omega \to \mathbb{R}$ such that:

$$\begin{aligned} -\Delta(u) &= f \quad \text{in } \Omega \\ u &= 0 \quad \text{on } \partial\Omega. \end{aligned} \tag{2.1}$$

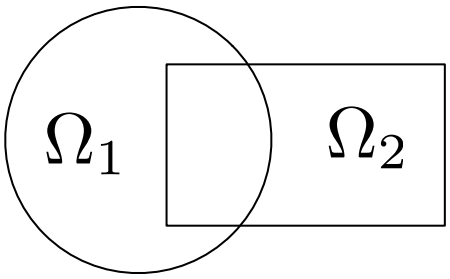

FIGURE 2.1. A complex domain made from the union of two simple geometries

Definition 2.1 (Original Schwarz algorithm)**.** The Schwarz algorithm is an iterative method based on solving alternatively sub-problems in domains Ω_1 and Ω_2. It updates $(u_1^n, u_2^n) \to (u_1^{n+1}, u_2^{n+1})$ by:

$$\begin{array}{rlcrl} -\Delta(u_1^{n+1}) = f & \text{in } \Omega_1 & & -\Delta(u_2^{n+1}) = f & \text{in } \Omega_2 \\ u_1^{n+1} = 0 & \text{on } \partial\Omega_1 \cap \partial\Omega & \text{then,} & u_2^{n+1} = 0 & \text{on } \partial\Omega_2 \cap \partial\Omega \\ u_1^{n+1} = u_2^n & \text{on } \partial\Omega_1 \cap \overline{\Omega_2}. & & u_2^{n+1} = u_1^{n+1} & \text{on } \partial\Omega_2 \cap \overline{\Omega_1}. \end{array} \tag{2.2}$$

H.A. Schwarz proved the convergence of the algorithm and thus the well-posedness of the Poisson problem in complex geometries.

With the advent of digital computers, this method also acquired a practical interest as an iterative linear solver. Subsequently, parallel computers became available and a small modification of the algorithm [26] makes it suited to these architectures. Its convergence can be proved using the maximum principle [25].

Our main focus is high performance computing which implies to have thousands of subdomains and also to be able to work with third-party linear solver libraries. Let N be the number of subdomains, the domain Ω is decomposed into N overlapping subdomains $(\Omega_i)_{1\le i\le N}$. Rather than iterating on the collection of local approximate solutions $(u_i^n)_{1\le i\le N}$, it is then more convenient to consider related algorithms where we iterate on a global approximate solution $u^n : \Omega \to \mathbb{R}$. In order to do this, we introduce a partition of unity. We proceed here in an informal way.

Definition 2.2 (Extension operators and partition of unity)**.** Let the extension operator E_i be such that $E_i(w_i) : \Omega \to \mathbb{R}$ is the *extension* of a function $w_i : \Omega_i \mapsto \mathbb{R}$, by zero outside Ω_i. We also define the *partition of unity functions* $\chi_i : \Omega_i \to \mathbb{R}$, $\chi_i \ge 0$ and $\chi_i(x) = 0$ for $x \in \partial\Omega_i \setminus \partial\Omega$ and such that:

$$w = \sum_{i=1}^{N} E_i(\chi_i \, w_{|\Omega_i}) \tag{2.3}$$

for any function $w : \Omega \mapsto \mathbb{R}$.

The Schwarz algorithm of (2.2) iterates on a pair of local functions which, until convergence, do not match in the overlap. As explained in [10, 12], its natural generalization to globally defined iterates is the RAS (Restricted Additive Schwarz) algorithm [7]:

For an approximate solution u^n to (2.1), we first solve in parallel subproblems

$$\begin{array}{rl} -\Delta(u_i^{n+1}) = f & \text{in } \Omega_i \\ u_i^{n+1} = 0 & \text{on } \partial\Omega_i \cap \partial\Omega \\ u_i^{n+1} = u^n & \text{on } \partial\Omega_i \setminus \partial\Omega. \end{array} \tag{2.4}$$

$$\text{followed by } u^{n+1} := \sum_{i=1}^{N} E_i(\chi_i \, u_i^{n+1}) \quad \text{in } \Omega\,.$$

Surprisingly enough, at least in the two-subdomain case, an equivalence between algorithms (2.2) and (2.4) holds independently of the choice of the partition of unity functions, see Lemma 1.5. in [10] at the continuous level.

3. Need for global transfer of information

When the number of subdomains is large, plateaus appear in the convergence of Schwarz domain decomposition methods. This is the case even for a simple model such as the Poisson problem (2.1). The problem of the one level method comes from the fact that in the Schwarz method there is a lack of a global exchange of information. Data are exchanged only from one

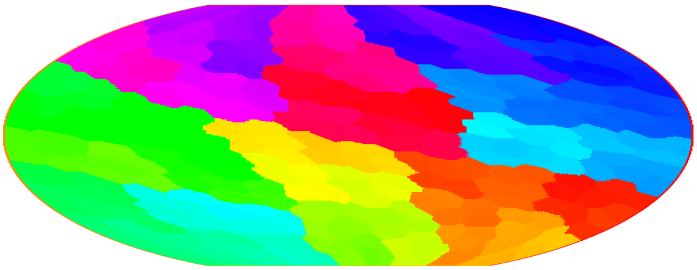

FIGURE 3.1. Decomposition into 18×18 subdomains. One color for each subdomain.

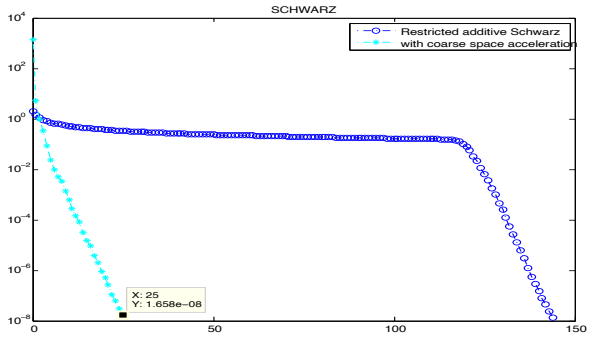

FIGURE 3.2. Convergence curves with and without a coarse space correction for a decomposition into 64 strips

subdomain to its direct neighbors. But the solution in each subdomain depends on the right-hand side in all subdomains. Let us denote by N_d the number of subdomains in one direction. Then, for instance, the leftmost domain of Figure 3.1 needs at least N_d iterations before being aware about the value of the right-hand side f in the rightmost subdomain. In practice, this long-distance effect cannot be neglected when the domain has a large inner radius or when it is elongated with Neumann boundary conditions on the long boundary. It is then observed that the length of the plateau is thus typically related to the number of subdomains in one direction as exemplified in Figure 3.2.

4. Coarse Space constructions

In order to analyze and fix this weakness of the one-level method, it is necessary to go beyond the above qualitative explanation relating the length of the delay in the convergence to the diameter of the graph of the connections between subdomains. In [32] an adaptive method was proposed at the continuous level. In order to generalize it, a first key tool for this is the Fictitious Space Lemma (see [33] for the original paper and [19] for a modern exposition) which can be seen as the *Lax–Milgram theorem* of domain decomposition methods. It is formulated as an abstract result of the theory of Hilbert spaces but actually it encompasses almost all (if not all) known domain decomposition methods: Schwarz methods, P.L. Lions algorithm, Balancing Neumann-Neuman and FETI methods. The second key tool is the deflation technique that comes from the linear algebra community, see e.g. [14, 34, 41] and references therein. This is connected as well to augmented or recycled Krylov space methods, see e.g. [8, 13, 35] or [37] and references therein.

For instance in [39] and in [10] and references therein, it is explained how to add a so-called coarse space correction to the original one-level method in order to get a scalable method with

a guaranteed condition number for the preconditioned system. More precisely, let equation (2.1) (or any symmetric positive definite system) be discretized by a finite element method so that the resulting linear system to be solved reads:

$$A\mathbf{U} = \mathbf{F} \tag{4.1}$$

where $\mathbf{U} \in \mathbb{R}^{\#\mathcal{N}}$ is the vector of degrees of freedom for a set of indices denoted by $\mathcal{N}$.

In order to define the domain decomposition method at the discrete level, we first decompose the global mesh $\mathcal{T}_h$ into overlapping sub meshes $(\mathcal{T}_{hi})_{1\leq i\leq N}$. This induces a decomposition of the global set of degrees of freedom $\mathcal{N}$ into N sub sets $(\mathcal{N}_i)_{1\leq i\leq N}$. Let $(R_i)_{1\leq i\leq N}$ be the restriction operator from the global set of indices $\mathcal{N}$ to the local one $\mathcal{N}_i$. The one-level ASM (additive Schwarz method) preconditioner which is a symmetrized version of the RAS method reads:

$$M_{ASM}^{-1} := \sum_{i=1}^{N} R_i^T \left(R_i\, A\, R_i^T\right)^{-1} R_i \tag{4.2}$$

Compared to RAS, its symmetry allows for the use of Hilbert space theory for its spectral analysis. As a result, a coarse space with a provable efficiency can be defined in several steps.

First, let D_i be diagonal square matrices of size $\#\mathcal{N}_i$ that define a partition of unity, that is for all $\mathbf{U} \in \mathbb{R}^{\#\mathcal{N}}$:

$$\mathbf{U} = \sum_{i=1}^{N} R_i^T\, D_i\, R_i\, \mathbf{U}\,. \tag{4.3}$$

A simple choice for D_i is that for each degree of freedom $k \in \mathcal{N}_i$, we set $(D_i)_{kk} := 1/\mu_k$ where μ_k denotes the number of subsets k belongs to. Note that other choices are possible and popular particularly the ones for which D_i is a Boolean matrix since it saves communications and avoids round-off errors.

Now, in each subdomain, let A_i^{Neu} be the matrix associated to the restriction of the variational formulation to the sub mesh $\mathcal{T}_{hi}$. We define a generalized eigenvalue problem

$$D_i R_i\, A\, R_i^T D_i \mathbf{V}_{ik} = \lambda_{ik} A_i^{Neu} \mathbf{V}_{ik} \tag{4.4}$$

Let τ be a user-defined threshold, we assume that the rectangular matrix Z_0 defined by the concatenation of the vectors $R_i^T D_i \mathbf{V}_{ik}$ for all $1 \leq i \leq N$ such that $\lambda_{ik} > 1/\tau$ is full rank. The coarse space is defined as the range of Z_0.

Following [28], a two-level preconditioner is defined as follows:

$$\boxed{M_{2,HSM}^{-1} := Z_0 (Z_0^T A Z_0)^{-1} Z_0^T + (I - P_0)\, M_{ASM}^{-1} (I - P_0^T)\,,} \tag{4.5}$$

where P_0 is the A orthogonal projection on the coarse space V_0:

$$P_0 := Z_0 (Z_0^T A Z_0)^{-1} Z_0^T A\,. \tag{4.6}$$

It is then possible to have a full control of the spectrum of the preconditioned operator, see Theorem 7.23 in [10]:

Theorem 4.1 (Hybrid Schwarz algorithm)**.** *Let τ be a user-defined parameter to build the GenEO coarse space as above.*

The eigenvalues of the hybrid Schwarz preconditioned system satisfy the following estimate

$$\boxed{\frac{1}{1 + k_1\, \tau} \leq \lambda(M_{2,HSM}^{-1}\, A) \leq k_0\,,} \tag{4.7}$$

where k_0 is the number of neighbors of a subdomain plus one, k_1 is the maximum multiplicity of the intersection between subdomains and $M_{2,HSM}^{-1}$ is defined by (4.5)*.*

We have considered here only the Additive Schwarz method but this kind of result has been extended to:

- FETI method [16, 40]
- P.L. Lions algorithm [20]
- Inexact Coarse solve [30]
- Saddle point problem [31]
- Boundary Element Methods [29]
- Multiscale Finite Element methods [27]
- Time dependent Maxwell system [5]
- Least Square problems [9]
- Purely algebraic settings [9, 17]

Note that for Helmholtz or frequency Maxwell type problems, efficient coarse spaces are more easily built when they are based on a coarse grid discretisation of the underlying variational form as it is classically done in multigrid methods, see [18] and references therein for a detailed mathematical analysis and [4] for extensive numerical tests of various approaches.

5. **Libraries for Large Scale Computations**

We have thus a strong set of mathematical results that enable to build efficient domain decomposition methods. But it comes at the expense of some difficulties since the implementation of these methods is not so easy due to the need to have access to the variational formulation to build the local matrices $(A_i^{Neu})_{1\leq i\leq N}$. Also, an efficient parallel implementation needs a careful organization of the code and the use of the message passing interface MPI on distributed memory machines.

This motivated the development of open source implementations of the GenEO coarse space:

- the C++/MPI library *hpddm* [22] either as an autonomous library or interfaced with FreeFem [21] or PETSc [2]. The library also provides access to the GCRODR method [35] useful when solving linear systems with multiple right hand sides.
- the FreeFem domain decomposition library *ffddm* [43]
- a Dune solver as described in [6]

It is worth noticing that the access to the variational formulation is made easy if one makes use of domain specific languages or libraries such as *FreeFem* [21], *Dune* [3] or *Firedrake* [36]. It facilitates the encapsulation of two-level methods. The resulting scripts are then quite compact and do not require to write MPI lines of code although they enable the parallel solving of all kinds of equations with the above mentioned methods.

We now give a few examples of computations performed using these libraries.

In [31], we solve three dimensional elasticity problems for steel-rubber ($\nu = 0,4999$) structures (see Figure 5.1) discretized by a finite element method with continuous pressure with up to a billion degrees of freedom on 16,800 cores.

Considering the brain imaging problem mentioned in the introduction, we solved the forward problem (at frequency 1 GHz) on the domain shown in Figure 1.1 (center), discretized with order 1 Nedelec's edge finite elements using 40 grid-points per wavelength, resulting in a linear system of size $n \approx 1.6 \times 10^7$. We take 729 subdomains (with a two-cell overlap) and for the two-level

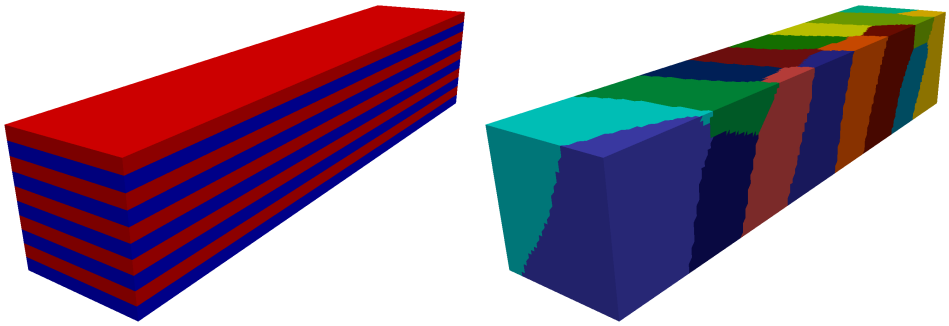

Figure 5.1. Heterogeneous beam composed of 10 alternating layers of rubber material $(E_1, \nu_1) = (1.\,10^7, 0.4999)$ and steel material $(E_2, \nu_2) = (2\,10^9, 0.35)$. Coefficient distribution (left) and mesh partitioning into 16 subdomains by the automatic graph partitioner *Metis* (right).

Table 5.1. Numerical performance for the microwave imaging system problem for three different material configurations inside the imaging chamber. In the Time columns we report the total time (the execution time for GMRES) in seconds.

	#2-level	Time	#1-level	Time
homogeneous liquid	28	63.4(8.6)	30	53.1(6.4)
head model	28	64.1(9.2)	32	53.4(6.9)
non-conductive cylinder	29	62.3(9.4)	125	83.5(38.2)

preconditioner a coarse problem of size $\approx 3.8 \times 10^4$. We tested a two-level preconditioner and the corresponding one-level version for three different material configurations: the imaging chamber filled just with a homogeneous matching liquid, the virtual head model of Figure 1.1, and a plastic-filled cylinder immersed in the matching liquid; the last test case is the most difficult because plastic is a non-conductive material. Then, waves are not damped in the chamber and it will delay the convergence of the domain decomposition method. In Table 5.1 we see that the performance of the one-level method deteriorates badly for the non-conductive cylinder, but the performance of the two-level method is uniform across all three cases, i.e. it appears robust with respect to the type of heterogeneity.

In [11], a two-level domain-decomposition preconditioner is used to solve large scale Helmholtz-type problems arising in the context of frequency-domain full-waveform inversion (FWI) for seismic imaging. The performance of the solver is assessed for cartesian meshes as well as unstructured meshes adapted to the size of the local wavelength (h-adaptivity) allowing to significantly reduce the number of unknowns in the discretization. The coarse problem in the two-level preconditioner is based on a coarse mesh discretization and is solved at each iteration in an approximate manner with an inner one-level method retaining the same spatial decomposition as the first level in order to minimize communications. Figure 5.2 shows a wavefield solution at frequency $f = 20$ Hz for the 3D acoustic SEG/EAGE Overthrust model computed using P3 finite elements on an unstructured mesh adapted to the local wavelength, with 5 points per wavelength. Table 5.2 reports simulation statistics for different frequencies with regular and adapted meshes for the Overthrust model. We can see for example that it takes 37 seconds and 30 GMRES iterations to solve the forward problem at frequency $f = 20$ Hz on 16960 cores, with 2.285 billion unknowns.

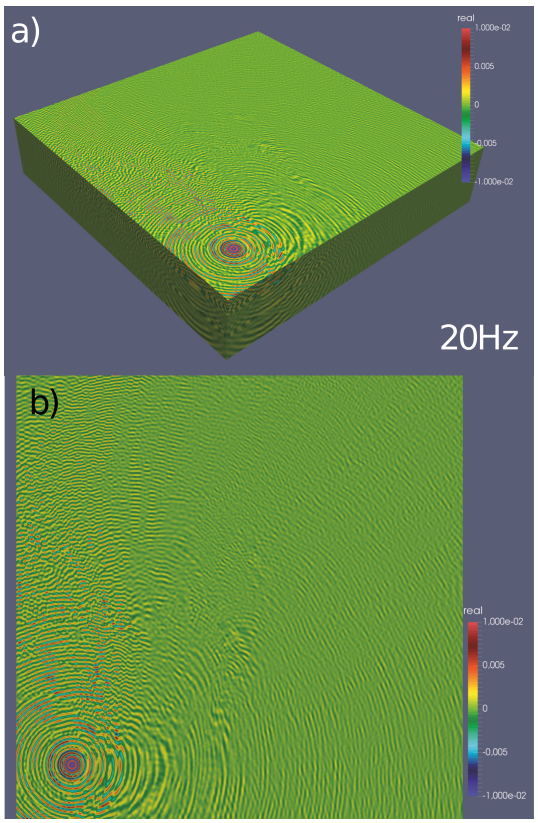

FIGURE 5.2. 20 Hz wavefield for the acoustic SEG/EAGE Overthrust model computed using P3 finite elements on an unstructured mesh adapted to the local wavelength, with 5 points per wavelength – FreeFem software

TABLE 5.2. Simulation statistics for the acoustic SEG/EAGE Overthrust model with regular and adaptive meshes. $Freq(Hz)$: frequency; #core: number of cores; #elts: number of elements; #dofs: number of degrees of freedom; #it: iteration count. Elapsed time (seconds) in GMRES.

Regular mesh					
Freq (Hz)	#core	#elts (M)	#dofs (M)	#it	GMRES
5	265	16	74	7	16s
10	2,120	131	575	15	33s
Adaptive mesh					
Freq (Hz)	#core	#elts (M)	#dofs (M)	#it	GMRES
10	2,120	63	286	14	15s
20	16,960	506	2,285	30	37s

In [23], the flexibility of PETSc [1, 2] was leveraged to try new applications of GenEO. First, the method was used on meshes with quadrilateral elements, see Figure 5.3, in comparison to previous FreeFEM results where only simplices (lines, triangles, and tetrahedra) are used. Then, GenEO was also used to precondition an eigensolver, LOBPCG [24], which requires the application of the preconditioned operator on a set of multiple vectors. This is particularly

efficient since HPDDM can leverage the high efficiency of subdomain solvers when dealing with blocks of right-hand sides. Eventually, the GenEO second-level correction was benchmarked in the context of a nonlinear problem, the Liouville–Bratu–Gelfand equation, solved with a Newton method. It was shown that it is possible to reuse the coarse operator to solve the linearized system without hindering the convergence of the outer solver.

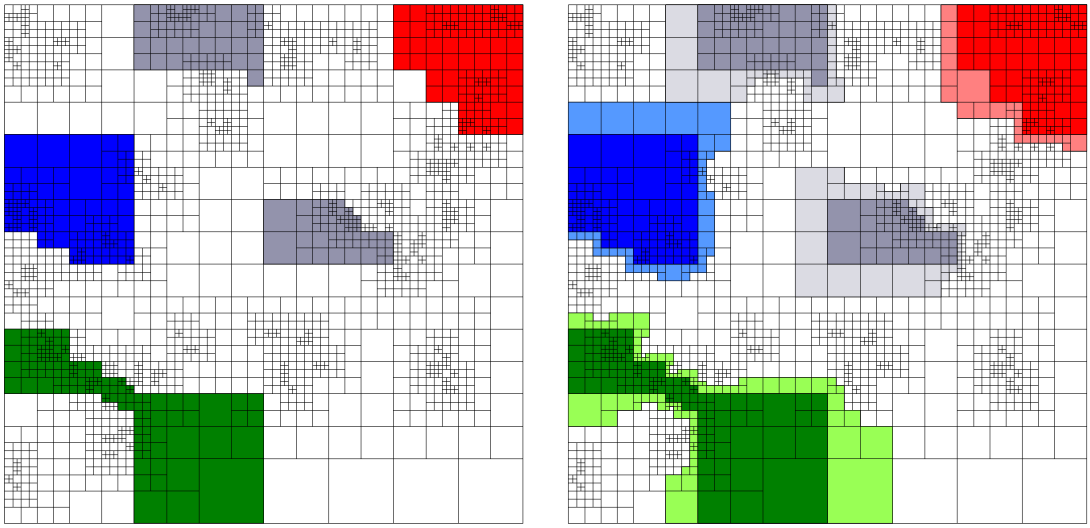

(A) Initial non-overlapping decomposition.

(B) Overlapping decomposition on which $(A_i^{Neu})_{1\leq i\leq N}$ are assembled

Figure 5.3. Automatic generation of overlapping subdomains and assembly of local Neumann operators needed by (4.4). For clarity, only four subdomains, one of which is disconnected, are colored. Figure taken from [23]

Also, in [6], the scalability of the GenEO solver of Dune was assessed on more than 15,000 cores of the UK national supercomputer Archer, solving an aerospace composite problem with over 200 million degrees of freedom in a few minutes.

References

[1] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, and Hong Zhang. PETSc Users Manual. Technical Report ANL-95/11 - Revision 3.5, Argonne National Laboratory, 2014.

[2] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser, 1997.

[3] Peter Bastian, Felix Heimann, and Sven Marnach. Generic implementation of finite element methods in the distributed and unified numerics environment (DUNE). *Kybernetika*, 46(2):294–315, 2010.

[4] Niall Bootland, Victorita Dolean, Pierre Jolivet, and Pierre-Henri Tournier. A comparison of coarse spaces for Helmholtz problems in the high frequency regime, 2021. `https://arxiv.org/abs/2012.02678`.

[5] Niall Bootland, Victorita Dolean, Frédéric Nataf, and Pierre-Henri Tournier. Two-level DDM preconditioners for positive Maxwell equations, 2020. `https://arxiv.org/abs/2012.02388`.

[6] Richard Butler, Tim Dodwell, Anne Reinarz, Anhad Sandhu, Robert Scheichl, and Linus Seelinger. High-performance Dune modules for solving large-scale, strongly anisotropic elliptic problems with applications to aerospace composites. *Comput. Phys. Commun.*, 249:106997, 2020.

[7] Xiao-Chuan Cai and Marcus Sarkis. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM J. Sci. Comput.*, 21:239–247, 1999.

[8] Andrew Chapman and Yousef Saad. Deflated and augmented Krylov subspace techniques. *Numer. Linear Algebra Appl.*, 4(1):43–66, 1997.

[9] Hussam Al Daas, Pierre Jolivet, and Jennifer Scott. A Robust Algebraic Domain Decomposition Preconditioner for Sparse Normal Equations. `https://arxiv.org/abs/2107.09006`, 2021.

[10] Victorita Dolean, Pierre Jolivet, and Frédéric Nataf. *An Introduction to Domain Decomposition Methods: algorithms, theory and parallel implementation*. Society for Industrial and Applied Mathematics, 2015.

[11] Victorita Dolean, Pierre-Henri Tournier, Pierre Jolivet, and Stéphane Operto. Large-scale frequency-domain seismic wave modeling on h-adaptive tetrahedral meshes with iterative solver and multi-level domain-decomposition preconditioners. In *SEG Technical Program Expanded Abstracts 2020*, pages 2683–2688. Society of Exploration Geophysicists, 2020.

[12] Evridiki Efstathiou and Martin J. Gander. Why restricted additive Schwarz converges faster than additive Schwarz. *BIT*, 43:945–959, 2003.

[13] Jocelyne Erhel and Frédéric Guyomarc'h. An augmented conjugate gradient method for solving consecutive symmetric positive definite linear systems. *SIAM J. Matrix Anal. Appl.*, 21(4):1279–1299, 2000.

[14] Yogi A. Erlangga and Reinhard Nabben. Deflation and balancing preconditioners for Krylov subspace methods applied to nonsymmetric matrices. *SIAM J. Matrix Anal. Appl.*, 30(2):684–699, 2008.

[15] Martin J. Gander. Schwarz methods over the course of time. *Electron. Trans. Numer. Anal.*, 31:228–255, 2008.

[16] Pierre Gosselet, Daniel Rixen, François-Xavier Roux, and Nicole Spillane. Simultaneous FETI and block FETI: Robust domain decomposition with multiple search directions. *Int. J. Numer. Meth. Engng.*, 104(10):905–927, 2015.

[17] Loïc Gouarin and Nicole Spillane. Fully algebraic domain decomposition preconditioners with adaptive spectral bounds. https://arxiv.org/abs/2106.10913, 2021.

[18] Ivan G. Graham, Euan A. Spence, and Jun Zou. Domain Decomposition with local impedance conditions for the Helmholtz equation with absorption. *SIAM J. Numer. Anal.*, 58(5):2515–2543, 2020.

[19] Michael Griebel and Peter Oswald. On the abstract theory of additive and multiplicative Schwarz algorithms. *Numer. Math.*, 70(2):163–180, 1995.

[20] Ryadh Haferssas, Pierre Jolivet, and Frédéric Nataf. An additive Schwarz method type theory for Lions's algorithm and a symmetrized optimized restricted additive Schwarz method. *SIAM J. Sci. Comput.*, 39(4):A1345–A1365, 2017.

[21] Frédéric Hecht. New development in FreeFem++. *J. Numer. Math.*, 20(3-4):251–265, 2012.

[22] Pierre Jolivet and Frédéric Nataf. HPDDM: High-Performance Unified framework for Domain Decomposition methods, MPI-C++ library, 2014. `https://github.com/hpddm/hpddm`.

[23] Pierre Jolivet, Jose E. Roman, and Stefano Zampini. KSPHPDDM and PCHPDDM: Extending PETSc with advanced Krylov methods and robust multilevel overlapping Schwarz preconditioners. *Comput. Math. Appl.*, 84:277–295, 2021.

[24] Andrew V. Knyazev, Merico E. Argentati, Ilya Lashuk, and Evgueni E. Ovtchinnikov. Block locally optimal preconditioned eigenvalue xolvers (BLOPEX) in HYPRE and PETSc. *SIAM J. Sci. Comput.*, 29(5):2224–2239, 2007.

[25] Pierre-Louis Lions. On the Schwarz alternating method. II. In Tony Chan, Roland Glowinski, Jacques Périaux, and Olof Widlund, editors, *Domain Decomposition Methods*, pages 47–70. Society for Industrial and Applied Mathematics, 1989.

[26] Pierre-Louis Lions. On the Schwarz alternating method. III: A variant for nonoverlapping subdomains. In Tony F. Chan, Roland Glowinski, Jacques Périaux, and Olof Widlund, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*. Society for Industrial and Applied Mathematics, 1990.

[27] Chupeng Ma, Robert Scheichl, and Tim Dodwell. Novel design and analysis of generalized FE methods based on locally optimal spectral approximations. `https://arxiv.org/abs/2103.09545`, 2021.

[28] Jan Mandel. Balancing domain decomposition. *Commun. Numer. Methods Eng.*, 9:233–241, 1992.

[29] Pierre Marchand, Xavier Claeys, Pierre Jolivet, Frédéric Nataf, and Pierre-Henri Tournier. Two-level preconditioning for the h-version boundary element approximation of hypersingular operator with GenEO. *Numer. Math.*, 146(3):597–628, 2020.

[30] Frédéric Nataf. Mathematical Analysis of Robustness of Two-Level Domain Decomposition Methods with respect to Inexact Coarse Solves. *Numer. Math.*, 2020.

[31] Frédéric Nataf and Pierre-Henri Tournier. A GenEO Domain Decomposition method for Saddle Point problems. `https://arxiv.org/abs/1911.01858`, 2021.

[32] Frédéric Nataf, Hua Xiang, Victorita Dolean, and Nicole Spillane. A coarse space construction based on local Dirichlet to Neumann maps. *SIAM J. Sci. Comput.*, 33(4):1623–1642, 2011.

[33] Sergey V. Nepomnyaschikh. Mesh theorems of traces, normalizations of function traces and their inversions. *Sov. J. Numer. Anal. Math. Model.*, 6:1–25, 1991.

[34] Roy A. Nicolaides. Deflation of conjugate gradients with applications to boundary value problems. *SIAM J. Numer. Anal.*, 24(2):355–365, 1987.

[35] Michael L. Parks, Eric de Sturler, Greg Mackey, Duane D. Johnson, and Spandan Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM J. Sci. Comput.*, 28(5):1651–1674, 2006.

[36] Florian Rathgeber, David A. Ham, Lawrence Mitchell, Michael Lange, Fabio Luporini, Andrew T. T. McRae, Gheorghe-Teodor Bercea, Graham R. Markall, and Paul H. J. Kelly. Firedrake: automating the finite element method by composing abstractions. *ACM Trans. Math. Softw.*, 43(3):1–27, 2016.

[37] Yousef Saad. Analysis of augmented Krylov subspace methods. *SIAM J. Matrix Anal. Appl.*, 18(2):435–449, 1997.

[38] Hermann A. Schwarz. Über einen Grenzübergang durch alternierendes Verfahren. *Vierteljahrsschrift der Naturforschenden Gesellschaft in Zürich*, 15:272–286, 1870.

[39] Nicole Spillane, Victorita Dolean, Patrice Hauret, Frédéric Nataf, Clemens Pechstein, and Robert Scheichl. Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps. *Numer. Math.*, 126(4):741–770, 2014.

[40] Nicole Spillane and Daniel Rixen. Automatic spectral coarse spaces for robust finite element tearing and interconnecting and balanced domain decomposition algorithms. *Int. J. Numer. Meth. Engng.*, 95(11):953–990, 2013.

[41] Jok M. Tang, Reinhard Nabben, Cornelis Vuik, and Yogi A. Erlangga. Comparison of two-level preconditioners derived from deflation, domain decomposition and multigrid methods. *J. Sci. Comput.*, 39(3):340–370, 2009.

[42] Pierre-Henri Tournier, Marcella Bonazzoli, Victorita Dolean, Francesca Rapetti, Frédéric Hecht, Frédéric Nataf, Iannis Aliferis, Ibtissam El Kanfoud, Claire Migliaccio, Maya De Buhan, et al. Numerical Modeling and High-Speed Parallel Computing: New Perspectives on Tomographic Microwave Imaging for Brain Stroke Detection and Monitoring. *IEEE Trans. Antennas Propag.*, 59(5):98–110, 2017.

[43] Pierre-Henri Tournier and Frédéric Nataf. FFDDM: FreeFem Domain Decomposition Methd, 2019. `https://doc.freefem.org/documentation/ffddm/index.html`.