CSMA 2026
17ème Colloque National en Calcul des Structures
18-22 Mai 2026, Presqu'île de Giens (Var)

Two-Level Schwarz Preconditioning for Large-Scale HHO Problems

B. Sidi Hida[1], N. Tardieu[2], P. Jolivet[3]

[1] *EDF (Électricité de France), Palaiseau, bahaa-eddine.sidi-hida@edf.fr*
[2] *EDF (Électricité de France), Palaiseau, nicolas.tardieu@edf.fr*
[3] *LIP6 (Laboratoire d'Informatique - CNRS - Sorbonne Université), Paris, pierre@joliv.et*

Résumé — This work presents a two-level Schwarz domain-decomposition preconditioning framework with a GenEO coarse space for finite element systems discretized with Hybrid High-Order (HHO) methods and assembled in code_aster. It combines a concise state-of-the-art review on SPD and saddle-point systems, an implementation contribution ensuring consistent index mapping between code_aster and PETSc for seamless coupling with iterative solvers and preconditioners, and a numerical study on linear elasticity benchmarks.
Mots clés — Domain decomposition methods, Two-level Schwarz, HHO methods.

1 Background (SPD) : overlapping DD preconditioners and coarse spaces

We consider a symmetric positive definite (SPD) system $A\mathbf{u} = \mathbf{b}$ arising from a conforming finite-element discretization of an elliptic problem. Let the mesh be partitioned into N_s subdomains with an overlap of thickness $\delta > 0$. Denote by R_i the Boolean restriction from the global space to the local unknowns on the overlapping subdomain i, and by D_i a diagonal "partition of unity" weight such that

$$\sum_{i=1}^{N_s} R_i^\top D_i R_i = I.$$

The basic one-level overlapping preconditioners are the Additive Schwarz Method (ASM) and its restricted variant (RAS) [1]. In particular, the (one-level) RAS reads

$$P_{\mathrm{RAS}}^{-1} = \sum_{i=1}^{N_s} R_i^\top D_i \left(R_i A R_i^\top\right)^{-1} R_i, \tag{1}$$

while ASM corresponds to dropping the weights D_i in (1). These methods are effective but *not scalable* : as N_s grows (subdomain diameter $H\downarrow$), the condition number $\kappa(P^{-1}A)$ deteriorates roughly like a positive power of H/δ, and the iteration count increases [1]. This motivates a second level (a coarse correction) to recover scalability [2].

Two-level additive Schwarz. In the scalable two-level variant, we enrich (1) with a coarse space $V_0 = \mathrm{range}(Z)$, $Z \in \mathbb{R}^{n\times m}$ $(m \ll n)$, yielding

$$M^{-1} = P_{\mathrm{RAS}}^{-1}\left(I - AZE^{-1}Z^\top\right) + ZE^{-1}Z^\top, \qquad E = Z^\top AZ.$$

This expression applies a single coarse solve per iteration and is preferred for parallel efficiency [2]. With a suitable V_0, the preconditioned operator $M^{-1}A$ admits condition-number bounds independent of N_s (and only weakly dependent on coefficient variations), restoring scalability [1].

Classical and spectral coarse spaces. For SPD scalar problems, early coarse spaces such as Nicolaides' (piecewise constants per subdomain) already yield two-level scalability for model Poisson-type operators. To handle strong heterogeneities and systems (e.g., elasticity), *spectral* coarse spaces built from localized generalized eigenproblems are preferred; the abstract two-level ASM theory (stable decomposition, strengthened Cauchy–Schwarz) and its consequences for coarse space design are detailed in [1].

GenEO coarse space (SC'13 formulation). Following [2], the deflation/coarse matrix is assembled as

$$Z \;=\; \begin{bmatrix} R_1^\top W_1 & R_2^\top W_2 & \cdots & R_{N_s}^\top W_{N_s} \end{bmatrix} \;\in\; \mathbb{R}^{n \times \sum_{i=1}^{N_s} \nu_i},$$

where, for each subdomain i, the block W_i collects the selected local eigenvectors post-weighted by the partition-of-unity :

$$W_i \;=\; \begin{bmatrix} D_i\Lambda_{i1} & D_i\Lambda_{i2} & \cdots & D_i\Lambda_{i\nu_i} \end{bmatrix} \;\in\; \mathbb{R}^{n_i \times \nu_i}.$$

The vectors $\{\Lambda_{ij}\}_{1 \le j \le \nu_i}$ are chosen among the eigenvectors associated with the smallest eigenvalues (in magnitude) of the following local generalized eigenproblem :

$$A_i^\delta \Lambda \;=\; \lambda D_i R_{i,0}^\top R_{i,0} A_i^\delta D_i \Lambda,$$

where A_i^δ is the (unassembled) local matrix yielded by the discretization of the bilinear form on the overlapped space V_i^δ (a "floating" subdomain when $\Omega_i^\delta \cap \partial\Omega = \varnothing$), $R_{i,0}$ is the restriction from V_i^δ to its overlap with neighbors, and D_i is a diagonal partition-of-unity weight on the overlap. A threshold criterion (e.g. keep all $\lambda \le \tau$) selects the ν_i columns of W_i. The coarse matrix is

$$E \;=\; Z^\top A Z.$$

This construction yields a two-level method whose condition number is independent of the number of subdomains and robust to large coefficient jumps [2].

Practical remarks (SPD). In practice, one-layer overlap is often sufficient; local operators $R_i A R_i^\top$ are assembled with appropriate boundary conditions (Dirichlet for ASM/RAS). Two-level RAS with a GenEO coarse space (as in HPDDM) delivers weak/strong scaling on large 2D/3D heterogeneous benchmarks, in line with the theory [2].

2 Background (saddle-point) : two-level DD with GenEO

We consider the saddle-point matrix

$$A \;=\; \begin{bmatrix} A_{00} & B_{10}^T \\ B_{10} & -A_{11} \end{bmatrix}, \qquad A_{00} \in \mathbb{R}^{n\times n} \text{ SPD},\; B_{10} \in \mathbb{R}^{m\times n} \text{ full row rank},\; A_{11} \succeq 0.$$

Block factorization shows that solving $A\,[u;\,p]^\top = [f;g]^\top$ reduces to (i) solves with A and (ii) a Schur complement system for the dual unknowns :

$$S \;=\; A_{11} \;+\; B_{10} A_{00}^{-1} B_{10}^\top.$$

We precondition A by a scalable two-level overlapping Schwarz method M_A^{-1} (e.g., RAS/ORAS with a GenEO coarse space, cf. [1]), and use it to define a first spectrally equivalent approximation of S :

$$P_S \;=\; A_{11} \;+\; B_{10} M_A^{-1} B_{10}^\top. \tag{2}$$

Domain-decomposition structure of P_S. Let R_i denote the Boolean restriction to overlapping subdomain i, $A_i = R_i A_{00} R_i^\top$, and R_0 the coarse restriction. Following the element/overlap-based splitting, introduce $\widetilde{R}_i$ as the restriction to the support of $\mathrm{Im}(BR_i^\top)$ and diagonal weights $\widetilde{D}_i$ with $\sum_i \widetilde{R}_i^\top \widetilde{D}_i \widetilde{R}_i = I_m$. Let us define local lifted constraint blocks $\widetilde{B}_i = \widetilde{R}_i B_{10} R_i^\top$ and local positive semidefinite contributions $\widetilde{A}_i$ such that

$$A_{11} \;\preceq\; \sum_{i=1}^{N_s} \widetilde{R}_i^\top \widetilde{A}_i \widetilde{R}_i \;\preceq\; \widetilde{k}_1 A_{11}.$$

With these ingredients, a spectrally equivalent splitting of P_S reads

$$S_0 \;=\; \widetilde{R}_0^\top \widetilde{B}_0 \,(R_0 A_{00} R_0^\top)^{-1} \widetilde{B}_0^\top \widetilde{R}_0, \tag{3}$$

$$S_1 \;=\; \sum_{i=1}^{N_s} \widetilde{R}_i^\top \left(\widetilde{A}_i \;+\; \widetilde{B}_i A_i^{-1} \widetilde{B}_i^\top \right) \widetilde{R}_i, \tag{4}$$

$$M_S \;=\; S_0 \;+\; S_1, \tag{5}$$

and M_S is spectrally equivalent to P_S (hence to S). The difficult part is S_1, for which we now build a two-level DD preconditioner.

One-level Neumann–Neumann for S_1 and the need for a coarse space. A one-level preconditioner for S_1 (Neumann–Neumann flavor) is

$$M_{S_1,\mathrm{one}}^{-1} \;=\; \sum_{i=1}^{N_s} \widetilde{R}_i^{\top}\, \widetilde{D}_i \left(\widetilde{A}_i \;+\; \widetilde{B}_i A_i^{-1} \widetilde{B}_i^{\top}\right)^{-1} \widetilde{D}_i \widetilde{R}_i,$$

but it is not scalable. As in the SPD case, scalability is restored by an adaptive coarse correction derived from a localized generalized eigenproblem. [3]

GenEO-type coarse space for S_1. For each subdomain i, define the local EVP

$$\widetilde{D}_i \widetilde{R}_i \Big(\sum_{j\in\mathcal{N}(i)} \widetilde{R}_j^{\top} \big(\widetilde{A}_j + \widetilde{B}_j A_j^{-1} \widetilde{B}_j^{\top}\big) \widetilde{R}_j \Big) \widetilde{R}_i^{\top}\, \widetilde{D}_i \phi \;=\; \lambda \left(\widetilde{A}_i + \widetilde{B}_i A_i^{-1} \widetilde{B}_i^{\top}\right) \phi, \tag{6}$$

retain the eigenvectors with $\lambda \le \tau_{S_1}$, stack them (optionally PoU-weighted) into local blocks $\widetilde{W}_i$, and assemble the coarse operator

$$Z_{S_1} \;=\; \big[\widetilde{R}_1^{\top} \widetilde{D}_1 \widetilde{W}_1 \;\cdots\; \widetilde{R}_{N_s}^{\top} \widetilde{D}_{N_s} \widetilde{W}_{N_s}\big], \qquad E_{S_1} \;=\; Z_{S_1}^{\top} S_1 Z_{S_1}.$$

This yields the scalable two-level preconditioner

$$M_{S_1}^{-1} \;=\; Z_{S_1} E_{S_1}^{-1} Z_{S_1}^{\top} \;+\; (I - P_{S_1}) \Big(\sum_{i=1}^{N_s} \widetilde{R}_i^{\top}\, \widetilde{D}_i \big(\widetilde{A}_i + \widetilde{B}_i A_i^{-1} \widetilde{B}_i^{\top}\big)^{-1} \widetilde{D}_i \widetilde{R}_i \Big) (I - P_{S_1})^{\top}, \tag{7}$$

where $P_{S_1} = Z_{S_1} E_{S_1}^{-1} Z_{S_1}^{\top} S_1$ is the S_1-orthogonal projector. In practice, kernels (if any) are injected into Z_{S_1} and local pseudo-inverses are used. [3]

Final Schur complement and outer Krylov. Define $N_S = S_0 + M_{S_1}$ and solve $S\,p = \widehat{g}$ by (flexible) GMRES preconditioned by N_S^{-1}. The overall method combines a two-level Schwarz for the primal block with a two-level Neumann–Neumann for the Schur complement, and is robust for nearly incompressible, heterogeneous elasticity (continuous pressure) with scalable iteration counts. [1, 3]

3 Index sets and coupling between `code_aster` and `PETSc`

Context and idea The assembly of Neumann matrices in `code_aster` is done correctly, as they coincide exactly with the local subdomain matrices. This is trivial for the user as he only needs to enable an option to obtain these matrices, while under the hood the assembly relies on ghost cells to recover them. The challenge is therefore not assembly, but *exposing* the block structure to the solver in order to : (1) separate the different physical fields, (2) attach a two–level domain–decomposition preconditioner to the physical blocks, (3) and keep the whole setup consistent across MPI processes.

Brief definitions **`FIELDSPLIT (PETSc).`** A block preconditioner that partitions the unknown vector into *fields* (e.g., physical vs. Lagrange) and applies a block solution scheme (additive, multiplicative, or Schur). The benefit is to isolate the physics (where effective preconditioners are known) while keeping the couplings under control.

`HPDDM (PETSc).` A two–level domain–decomposition preconditioner designed for large-scale parallelism : local subdomain solvers plus a coarse space (generic modes) to ensure robustness and scalability as the number of subdomains increases.

Practical problem The matrices are already parallelized, but the solver needs three clear pieces of information : (i) “which indices belong to each field ?” (ii) “how to provide HPDDM with a *contiguous* and global view of each block ?” (iii) “which local submatrix must be supplied to HPDDM’s auxiliary interface ?”

3.1 Setup outline

— **Separate the fields.** Starting from the global numbering in `code_aster`, extract two index lists : field 1 and field 2. On each rank, keep only *owned* indices (no ghosts) to avoid duplicates. This separation is sufficient to enable `FIELDSPLIT` (field 1 block / field 2 block).
— **Make each block globally contiguous.** HPDDM expects the field indices to be numbered from $0, \ldots, n_{\text{FIELD}} - 1$ at the global level. Gather the field's indices from all ranks, sort and unique them to obtain a global list ; each rank then projects onto this list to build its consistent local view. Result : all ranks refer to the same block, with no gaps or ambiguous overlaps.
— **Extract the field's local submatrix.** From the local matrix and the local→global map, retain only the owned rows/columns that belong to the field of interest. This local submatrix is the building block passed to HPDDM's auxiliary interface (used to construct the subdomain solvers and the coarse space).

3.2 Why this addresses our problem

— *Numerical clarity.* `FIELDSPLIT` cleanly separates the physics, enabling well-known block strategies while keeping couplings under control.
— *MPI consistency.* The "owned indices ∩ field tags" rule avoids ghost-related duplicates ; projection onto the global list ensures that all ranks share exactly the same view of each field block.

4 Numerical Test

4.1 3D Linear Elasticity on a Cube

We consider a three-dimensional linear elasticity problem on the unit cube $\Omega = [0,1]^3$, discretized with the Hybrid High-Order (HHO) method [6, 7]. The numerical experiments are performed on two large-scale configurations, involving approximately 6.5×10^6 degrees of freedom on 64 MPI processes and 2.2×10^7 degrees of freedom on 216 MPI processes after assembly in `code_aster`.

4.1.1 Mechanical model

The test case corresponds to a standard small-strain linear elasticity problem on the unit cube. The equations read

$$-\nabla \cdot \sigma(u) = 0, \qquad \sigma(u) = 2\mu\varepsilon(u) + \lambda \operatorname{tr}(\varepsilon(u))I,$$

with one face of the cube clamped (zero displacement) and a uniform pressure applied on the opposite face. All other faces are traction-free.

The system assembled by `code_aster` is therefore a symmetric positive definite matrix resulting from the HHO discretization of the elasticity operator. No Lagrange multipliers or additional linear constraints are used in this configuration.

4.1.2 Solver and preconditioner

The linear systems are solved inside `code_aster` using the PETSc infrastructure. We employ `FGMRES` together with two scalable preconditioning strategies based on a two-level overlapping Schwarz method (HPDDM) equipped with a GenEO coarse space. In the monolithic configuration, HPDDM is applied directly to the full matrix A. In the `FIELDSPLIT` configuration, a Schur-type decomposition is used and HPDDM is applied to both blocks of the system, including the Schur complement. The experiments target large-scale HHO elasticity problems involving between 6.5×10^6 and 2.2×10^7 degrees of freedom.

Interpretation. Figures 1 compare two preconditioning strategies that both rely on a two-level overlapping Schwarz method implemented through HPDDM and equipped with a GenEO coarse space. All numerical experiments were performed on the SELENA HPC cluster of EDF. In the monolithic configuration, HPDDM is applied directly to the full HHO stiffness matrix, leading to a nearly constant

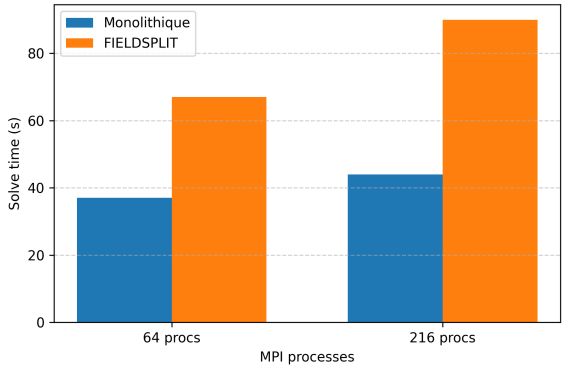

(a) Solve time vs. number of MPI processes

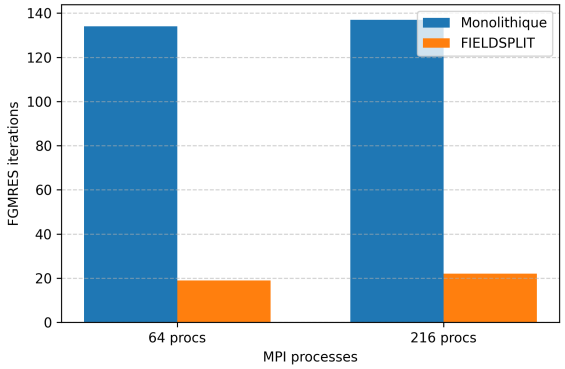

(b) FGMRES iteration count vs. number of MPI processes

FIGURE 1 – Comparison of monolithic and FIELDSPLIT preconditioning strategies for the 3D HHO elasticity test case

FGMRES iteration count when scaling from 6.5×10^6 degrees of freedom on 64 MPI processes (134 iterations, 37 s) to 2.2×10^7 degrees of freedom on 216 processes (137 iterations, 44 s). This behavior indicates that the HPDDM–GenEO preconditioner effectively preserves the spectral properties of the global operator at scale. In the `FIELDSPLIT` configuration, HPDDM is applied to both blocks of the Schur-type decomposition, including the Schur complement, which results in a drastic reduction of the Krylov iteration count to 19 iterations on 64 processes and 22 iterations on 216 processes. This improved convergence reflects a better-conditioned preconditioned system at the algorithmic level. However, the additional cost associated with the Schur complement construction and application increases the cost per iteration, leading to solver times of 67 s and 90 s, respectively. Overall, these results highlight a trade-off between iteration efficiency and time-to-solution, where the monolithic HPDDM–GenEO configuration remains more time-efficient for this SPD elasticity problem, while the FIELDSPLIT–HPDDM strategy offers enhanced robustness and flexibility for more complex large-scale systems.

4.2 Cylinder with multipoint constraints

The second test reproduces the cylinder with multipoint constraints benchmark introduced by Kruse et al. [8]. We keep exactly the same finite element discretization used in their study, while solving the resulting linear system with a PETSc configuration based on FIELDSPLIT and HPDDM.

The test was run using FGMRES as the global Krylov solver, combined with a FIELDSPLIT preconditioner of Schur type separating the physical degrees of freedom and the Lagrange multipliers. Both the A00 block and the Schur complement were preconditioned with HPDDM, using 40 selected eigenvalues (eps_nev = 40) per subdomain for A00 and 20 (eps_nev = 20) for the Schur complement. The simulation was executed on 192 processors of the SELENA cluster. For a relative tolerance of 1e-8, the global solver required 64 outer iterations and converged in 2772 seconds. These results demonstrate that such a configuration is viable and efficient on this benchmark, and they provide a solid basis for future comparisons with fully monolithic solution strategies.

5 Conclusion

This work establishes a practical and scalable workflow for solving large finite-element systems assembled in `code_aster` using PETSc and two-level Schwarz preconditioning. For the 3D HHO elasticity test case, we demonstrated that both a monolithic HPDDM–GenEO configuration and a FIELDSPLIT–HPDDM strategy exhibit scalable behavior on large problem sizes, ranging from 6.5×10^6 to 2.2×10^7 degrees of freedom. The monolithic approach achieves lower time-to-solution while maintaining a nearly constant Krylov iteration count, whereas the FIELDSPLIT configuration, with HPDDM applied to both blocks of a Schur-type decomposition, significantly reduces the number of Krylov iterations at the expense of higher per-iteration costs. In addition, we showed that a FIELDSPLIT–HPDDM configuration

successfully handles a benchmark with multipoint constraints on 192 processors. Altogether, these results demonstrate that both monolithic and block-structured strategies based on HPDDM–GenEO are viable at scale, and they provide a solid foundation for systematic comparisons and the development of robust FIELDSPLIT formulations for HHO and coupled large-scale systems.

Références

[1] V. Dolean, P. Jolivet, F. Nataf. *An Introduction to Domain Decomposition Methods : Algorithms, Theory, and Parallel Implementation*, SIAM, Philadelphia, 2015. DOI : 10.1137/1.9781611974063.

[2] P. Jolivet, F. Hecht, F. Nataf, C. Prud'homme. *Scalable Domain Decomposition Preconditioners for Heterogeneous Elliptic Problems*, SC '13 : Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, ACM, Article 80, 11 pages, 2013. DOI : 10.1145/2503210.2503282.

[3] F. Nataf, P.-H. Tournier. *GenEO coarse spaces for saddle-point problems : a two-level domain decomposition preconditioner*, Comptes Rendus Mécanique, 2023.

[4] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, K. Mei, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, J. Zhang. *PETSc Users Manual*, Argonne National Laboratory, ANL-95/11, 2024. Available at `https://petsc.org`.

[5] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, R. Scheichl. *Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps*, Numerische Mathematik, 126 :741–770, 2014. DOI : 10.1007/s00211-013-0576-y.

[6] Cicuttin, M., Ern, A., & Pignet, N. (2021). Hybrid high-order methods. A primer with application to solid mechanics. HAL. `https://hal.science/hal-03262027v2` doi :10.1007/978-3-030-81477-9

[7] Pignet, N. (2019). Hybrid High-Order methods for nonlinear solid mechanics (Thèse de doctorat, Université Paris-Est Marne la Vallée). HAL.`https://theses.hal.science/tel-02318157`

[8] C. Kruse, V. Darrigrand, N. Tardieu, M. Arioli, U. Rüde. *Application of an iterative Golub–Kahan algorithm to structural mechanics problems with multi-point constraints*, Advances in Modeling and Simulation in Engineering Sciences, 7 :45, 2020. DOI : 10.1186/s40323-020-00181-2.