

PROPOSITION DE SUJET DE THESE

Intitulé : Adaptive memory-aware hybrid linear solvers for compressible flows

Référence : **SNA-DAAA-2026-01**
(à rappeler dans toute correspondance)

Début de la thèse : 01/10/26

Date limite de candidature : 31/05/26

Mots clés : Hybrid Krylov solvers, memory constraints, restricted additive Schwarz preconditioner, LU factorization, block low-rank matrices, mixed precision algorithms, weighted graph partitioning

Profil et compétences recherchées

Master's degree or equivalent with a background in Applied Mathematics or Computer Science. A keen interest in numerical linear algebra and programming (C++, Fortran, Python) would be welcome.

Présentation du projet doctoral, contexte et objectif

Context

High-fidelity simulations of turbulent compressible flows in aerodynamics typically require the numerical analysis of three-dimensional flows around complex geometries. In optimization, linear analysis or data assimilation area, typical Computational Fluid Dynamics (CFD) workflows, such as fixed-point iterations or adjoint-state methods, require the solution of large, sparse, non-symmetric and ill-conditioned linear systems. An active area of research focuses on developing parallel, robust and efficient solvers capable of delivering solutions to such systems within a prescribed error tolerance. A key challenge, especially as problem sizes approach billions of unknowns, is the design of effective preconditioning operators. Although the solve phase still account for a substantial portion of the total CPU time in parallel simulations, the number of computing cores is usually dictated by the requirements of the CFD study itself, rather than by the needs of the linear solver.

Objectives

Recent studies [1-4] have demonstrated the strong performance of hybrid direct-iterative strategies. In these approaches, from the algebraic decomposition of the matrix, a flexible Krylov method is employed with a domain decomposition method as preconditioner [5] and an approximate direct method [6] as subdomain solver. The main objective is now to design a preconditioning operator that fully exploits the memory budget already allocated for the CFD simulation. The first MPI paradigm splits the study domain geometrically into well-balanced partitions. We aim to introduce a second MPI paradigm devoted to linear systems: enlarging subdomains for the local approximate direct solvers may significantly enhance the global numerical efficiency of the approach. To mitigate the computational and memory costs traditionally associated with direct solvers, we rely on the general-purpose multifrontal *MUMPS* solver [7] that exploits variable precision and possible low-rank property of matrices [8]. Initial numerical experiments using a fixed accuracy Block Low-Rank (BLR) multifrontal factorization [9, 10, 11] as the subdomain solver have already shown promising CPU-time reductions, together with a significant memory compression of the *L* and *U* factors compared with a classical full *LU* factorization. The new HPC capabilities available in *MUMPS*, combined with variable subdomain sizes, may open further opportunities for performance improvements. In addition, adapting the accuracy of the subdomain solver to the subdomain stiffness may offer further benefits [12, 13]. To preserve scalability under such heterogeneous configurations, we will also investigate load-balancing techniques with weighted graph partitioning.

Key steps

The ONERA CFD code *SoNICS* [14] will rely on the *MUMPS* library to define a mixed-precision BLR-LU solver within each subdomain, initially using a uniform accuracy. A second MPI parallelization strategy will then be developed through the *ParaDiGM* library [15], already used by *SoNICS*, to generate a new partitioning better suited to linear system solves. Subsequently, a non-uniform accuracy strategy will be explored, adjusting the subdomain accuracy in combination with load-balancing techniques. Numerical

experiments on challenging test cases will be performed at each stage to assess the potential and the limitations of the proposed approach.

Bibliography

- [1] D. N. Wakam, F. Pacull, [Memory efficient hybrid algebraic solvers for linear systems arising from compressible flows](#), Computers & Fluids, 80, 2013.
- [2] N. Guilbert, *Amélioration de l'inversion de grands systèmes creux pour la simulation numérique en mécanique des fluides*, Thèse, 2021.
- [3] M. Jadoui, C. Blondeau, E. Martin, F. Renac, and F.-X. Roux, [Comparative study of inner-outer Krylov solvers for linear systems in structured and high-order unstructured CFD problems](#), Computers & Fluids, 244, 2022.
- [4] Y. Jang, L. Grigori, E. Martin, C. Content, [Randomized flexible GMRES with deflated restarting](#), Numerical Algorithms, 98, 2024.
- [5] X.-C. Cai and M. Sarkis, [A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems](#), SIAM J. Sci. Comput. 21(2), pp 792–797, 1999.
- [6] M. Benzi, [Preconditioning Techniques for Large Linear Systems: A Survey](#), J. Comput. Phys. 182 (2), pp 418-477, 2002.
- [7] MUMPS group, [MUMPS overview and recent features](#), User Days 2023, Sorbonne University.
- [8] Zan Xu, L. Cambier, J.J. Alonso, E. Darve, [Towards a Scalable Hierarchical High-order CFD Solver](#), AIAA Scitech 2021 Forum, 2021.
- [9] P. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, and C. Weisbecker, [Improving multifrontal methods by means of block low-rank representations](#), SIAM J. Sci. Comput., 37(3):A1451–A1474, 2015.
- [10] P. Amestoy, A. Buttari, J.-Y. L'Excellent, T. Mary, [Performance and Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures](#), ACM Transactions on Mathematical Software, 45 (1), 2019.
- [11] P. Amestoy, O. Boiteau, A. Buttari, M. Gerest, F. Jézéquel, J.-Y. L'Excellent, T. Mary, [Mixed precision low-rank approximations and their application to block low-rank LU factorization](#), IMAJNA, 43 (4), 2023.
- [12] Z. Jorti, [Fast solution of sparse linear systems with adaptive choice of preconditioners](#), Thesis, Sorbonne Université, 2019.
- [13] M. Franco, [Efficient solvers for the implicit time integration of matrix-free high-order methods](#), PhD Thesis, 2022.
- [14] D. Gueyffier, S. Plot, M. Soismier, [SoNICS: a new generation CFD software for satisfying industrial users needs](#), OTAN/STO/Workshop AVT-366, 2022.
- [15] E. Quemerais et al., *Parallel Distributed General Mesh (ParaDiGM) library*, ONERA public GitLab repository under LGPL-3.0 license. <https://github.com/onera/paradigm>

Collaborations envisagées

Laboratory LIP6 at Sorbonne University.

Team PEQUAN (Performance and Quality of Numerical Algorithms).

Laboratoire d'accueil à l'ONERA

Département : Aérodynamique, Aéroélasticité, Acoustique
 Lieu (centre ONERA) : 29 avenue de la Division Leclerc, 92 322 Châtillon Cedex, France
Contact : Emeric MARTIN / Cédric CONTENT
 Tél. : 01 46 73 43 89 / 01 46 73 42 56
 Email : emeric.martin@onera.fr, cedric.content@onera.fr

Directeur de thèse

Nom : Pierre JOLIVET
 Laboratoire : Sorbonne Université - LIP6, office 26-00/311, 4 place Jussieu, 75252, Paris Cedex 05, France
 Tél. : 01 44 27 80 29
 Email : pierre.jolivet@lip6.fr

Pour plus d'informations : <https://www.onera.fr/rejoindre-onera/la-formation-par-la-recherche>