

# 2025 Activity Report

RESEARCH CENTRE: Inria Centre at the University of Bordeaux

IN PARTNERSHIP WITH: Airbus Central Research & Technology, Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique

  
Project-Team

## CONCACE

Numerical and Parallel Composability for High Performance Computing





## **Project-Team CONCACE**

*Creation of the Project-Team: 2022 September 01*

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

## Keywords

### Computer sciences and digital sciences

- A1.1.4. – High performance computing
- A1.1.5. – Exascale
- A1.1.9. – Fault tolerant systems
- A6.2.5. – Numerical Linear Algebra
- A6.2.7. – HPC for machine learning
- A6.3. – Computation-data interaction
- A7.1. – Algorithms
- A8.2. – Optimization
- A8.10. – Computer arithmetic
- A9.2. – Machine learning
- A9.7. – AI algorithmics
- A9.10. – Hybrid approaches for AI

### Other research topics and application domains

- B3.3.1. – Earth and subsoil
- B4.2.2. – Fusion
- B5.2.3. – Aviation
- B5.5. – Materials
- B9.5.1. – Computer science
- B9.5.2. – Mathematics
- B9.5.4. – Chemistry
- B9.5.6. – Data science

## Contents

<b>Project-Team CONCACE</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>5</b>
<b>2 Overall objectives</b>	<b>6</b>
<b>3 Research program</b>	<b>8</b>
3.1 Numerical algebra methodologies in model and data-driven scientific computing . . . . .	9
3.1.1 Scientific computing in large size linear algebra . . . . .	9
3.1.2 Scientific computing in large dimension multi-linear algebra . . . . .	9
3.1.3 Scientific continuum between large size and large dimension . . . . .	10
3.2 Composition of parallel numerical algorithms from a sequential expression . . . . .	10
<b>4 Application domains</b>	<b>11</b>
4.1 Aeroacoustics Simulation . . . . .	11
<b>5 Highlights of the year</b>	<b>12</b>
5.1 Awards . . . . .	12
<b>6 Latest software developments, platforms, open data</b>	<b>12</b>
6.1 Latest software developments . . . . .	12
6.1.1 composyx . . . . .	12
6.1.2 ScalFMM . . . . .	12
6.1.3 CPPDiodon . . . . .	13
6.1.4 FMR . . . . .	14
<b>7 New results</b>	<b>14</b>
7.1 Error Estimates for Sparse Tensor Products of B-spline Approximation Spaces . . . . .	14
7.2 On some orthogonalization schemes in Tensor Train format . . . . .	14
7.3 A note on TT-GMRES for the solution of parametric linear systems . . . . .	15
7.4 Solving eigenvalue problems in high dimensions using contour integration and Tensor Train format . . . . .	15
7.5 A Tensor Train solver for the Magnetic Moment Method . . . . .	15
7.6 Generalized Golub–Kahan Bidiagonalization for Nonsymmetric Saddle-Point Systems . . . . .	16
7.7 A note on the partial convergence management for the solution of symmetric linear systems with multiple right-hand sides . . . . .	16
7.8 A Scalable and Parameter-Robust Preconditioner for a Second Gradient of Dilation Regularization Applied to a Mechanics Problem . . . . .	16
7.9 Neural network preconditioning: a case study for the solution of the parametric Helmholtz equation . . . . .	16
7.10 Memory- and compute-optimized geometric multigrid GMGPolar for curvilinear coordinate representations -Applications to fusion plasma . . . . .	17
7.11 Complexity analysis and scalability of a matrix-free extrapolated geometric multigrid solver for curvilinear coordinates representations from fusion plasma applications . . . . .	17
7.12 Fault-tolerant numerical iterative algorithms at scale . . . . .	17
7.13 A filtered multilevel Monte Carlo method for estimating the expectation of cell-centered discretized random fields . . . . .	18
7.14 Multilevel Monte Carlo methods for ensemble variational data assimilation . . . . .	18
7.15 Convergence analysis of overlapping domain decomposition preconditioners for nonlinear problems . . . . .	18
7.16 Robustness and reliability of state-space, frame-based modeling for thermoacoustics . . . . .	19
7.17 Juxtaposing the fourth order vibrational operator perturbation theory CVPT(4) and the adaptive VCI (A-VCI): Accuracy, vibrational resonances and polyads of C <sub>2</sub> H <sub>4</sub> and C <sub>2</sub> D <sub>4</sub> . . . . .	19

7.18	Approximation Algorithms for Scheduling with/without Deadline Constraints where Rejection Costs are Proportional to Processing Times . . . . .	19
7.19	Guix-HPC Activity Report 2023-2024 . . . . .	20
<b>8</b>	<b>Partnerships and cooperations</b>	<b>20</b>
8.1	European initiatives . . . . .	20
8.1.1	H2020 projects . . . . .	20
8.2	National initiatives . . . . .	21
<b>9</b>	<b>Dissemination</b>	<b>24</b>
9.1	Promoting scientific activities . . . . .	24
9.1.1	Scientific events: organisation . . . . .	24
9.1.2	Scientific events: selection . . . . .	25
9.1.3	Journal . . . . .	25
9.1.4	Invited talks . . . . .	25
9.1.5	Scientific expertise . . . . .	25
9.1.6	Research administration . . . . .	25
9.2	Teaching - Supervision - Juries - Educational and pedagogical outreach . . . . .	26
9.2.1	Supervision . . . . .	26
9.2.2	Juries . . . . .	26
<b>10</b>	<b>Scientific production</b>	<b>27</b>
10.1	Major publications . . . . .	27
10.2	Publications of the year . . . . .	28
10.3	Cited publications . . . . .	32

# 1 Team members, visitors, external collaborators

## Research Scientists

- Luc Giraud [Team leader, INRIA, Senior Researcher, HDR]
- Carola Kruse [Team leader, CERFACS, Senior Researcher]
- Jean-Marie Couteyen [Team leader, AIRBUS, Senior Researcher, from Feb 2025]
- Emmanuel Agullo [INRIA, Researcher]
- Pierre Benjamin [AIRBUS, Senior Researcher]
- Olivier Coulaud [INRIA, Senior Researcher, HDR]
- Clement Guillet [INRIA, ISFP, from Oct 2025]
- Sofiane Haddad [AIRBUS, Senior Researcher]
- Paul Mycek [CERFACS, Senior Researcher]
- Guillaume Sylvand [AIRBUS, Senior Researcher]

## Post-Doctoral Fellows

- Joanna Bisch [INRIA, from Sep 2025, Post-Doctoral Fellow]
- Théo Boivin [CERFACS, from Apr 2025, Post-Doctoral Fellow]
- Augustin Leclerc [INRIA, Post-Doctoral Fellow]
- Stojche Nakov [INRIA, Post-Doctoral Fellow]

## PhD Students

- Theo Briquet [INRIA]
- Sai Aakash Dasari [CERFACS]
- Hugo Dodelin [INRIA]
- El Mehdi Ettaouchi [EDF, CIFRE]
- Antoine Gicquel [INRIA]
- Alexandre Malhene [INRIA]
- Clément Peaucelle [AIRBUS, CIFRE]
- Amine El Mehdi Zekri [TOULOUSE INP]

## Technical Staff

- Ludovic Courtes [INRIA, Engineer, (SED, 20 %)]
- Arthur Gouinguenet [INRIA, Engineer, from Oct 2025]
- Eragul Korkmaz [INRIA, Engineer]
- Gilles Marait [INRIA, Engineer, (SED, 90%)]
- Florent Pruvost [INRIA, Engineer, (SED, 40%)]

## Interns and Apprentices

- Mathis Foussac [INRIA, Apprentice, from Nov 2025]
- Mathis Foussac [INRIA, Intern, from Jun 2025 until Aug 2025]
- Vivien Nebout [INRIA, Intern, from Jun 2025 until Aug 2025]
- Juliette Petit [INRIA, Intern, from Jun 2025 until Aug 2025]

## Administrative Assistants

- Catherine Cattaert Megrat [INRIA]
- Marie-Melissandre Roy [INRIA]

## External Collaborators

- Luciano Drozda [CERFACS]
- Marek Felsoci [UVSQ, from Dec 2025]
- Jean-Rene Poirier [TOULOUSE INP, HDR]

## 2 Overall objectives

Over the past few decades, there have been innumerable science, engineering and societal breakthroughs enabled by the development of high performance computing (HPC) applications, algorithms and architectures. These powerful tools have enabled researchers to find computationally efficient solutions to some of the most challenging scientific questions and problems in medicine and biology, climate science, nanotechnology, energy, and environment – to name a few – in the field of **model-driven computing**. Meanwhile the advent of network capabilities and IoT, next generation sequencing, . . . tend to generate a huge amount of data that deserves to be processed to extract knowledge and possible forecasts. These calculations are often referred to as **data-driven calculations**. These two classes of challenges have a common ground in terms of numerical techniques that lies in the field of linear and multi-linear algebra. They do also share common bottlenecks related to the size of the mathematical objects that we have to represent and work on; those challenges retain a growing attention from the computational science community.

In this context, the purpose of the concace project, is **to contribute to the design of novel numerical tools for model-driven and data-driven calculations** arising from challenging academic and industrial applications. The solution of these challenging problems requires a multidisciplinary approach involving applied mathematics, computational and computer sciences. In applied mathematics, it essentially involves **advanced numerical schemes both in terms of numerical techniques and data representation of the mathematical objects** (e.g., compressed data, low-rank tensor [84, 91, 80] low-rank hierarchical matrices [82, 69]). In computational science, it involves large scale parallel heterogeneous computing and **the design of highly composable algorithms**. Through this approach, concace intends to contribute to all the steps that go from the design of new robust and accurate numerical schemes to the flexible implementations of the associated algorithms on large computers. To address these research challenges, researchers from Inria, Airbus Central R&T and Cerfacs have decided to combine their skills and research efforts to create the Inria concace project team, which will allow them to cover the entire spectrum, from fundamental methodological concerns to full validations on challenging industrial test cases. Such a joint project will enable a real synergy between basic and applied research with complementary benefits to all the partners. The main benefits for each partner are given below:

- Airbus Central R&T
  - Push our specific needs and use-cases towards the academic world to stimulate research in particular directions;

- Remain at the level of the scientific state of the art, this collaboration allows us to facilitate the feedback by exposing directly our challenges and industrial applications to eventually facilitate the transfer of research in our design tools;
  - The Inria research model will naturally be extended to Airbus, allowing for the multiplication of ambitious, very upstream and long-term research, while at the same time directly applying to the needs expressed by Airbus;
  - Benefit from the very high-level international network of the Inria team (e.g., Univ. of Tennessee Knoxville, Barcelona supercomputing center, Julich supercomputing center, Lawrence Berkeley National Lab, Sandia National Lab, etc.).
- Cerfacs
    - Join forces, in terms of skills and expertise, with Inria and Airbus to make faster and more effective progress on the research areas addressed by the team;
    - Bring scientific challenges from industrial applications through our privileged relationship with our industrial partners;
    - Reciprocally, promote the developed methodologies and the obtained results towards our industrial partners;
    - Naturally interact with the national and european HPC ecosystems, as a member of the EuroHPC national competence center on HPC, to promote the research activities and tools of the team and to meet novel scientific challenges where our methodologies or tools apply.
  - Inria
    - Reinforce the impact of our research through a direct contact and close interactions with real scientific and technical challenges;
    - Feed the virtuous feedback cycle between academic research and industrially-relevant applications enabling the emergence of new research avenues;
    - Create a privileged space for an open scientific dialogue enabling the fostering of existing synergies and to create new ones, in particular when one of the industrial partners is a large group whose spectrum of scientific problems is very broad.

In addition to the members of these entities, two other external collaborators will be strongly associated: Jean-René Poirier, from Laplace Laboratory at University of Toulouse) and Oguz Kaya, from LISN (Laboratoire Interdisciplinaire des Sciences du Numérique) at University of Saclay.

The scientific objectives described in Section 4 contain two main topics which cover numerical and computational methodologies. Each of the topic is composed of a methodological component and its validation counterpart to fully assess the relevance, robustness and effectiveness of the proposed solutions. First, we address numerical linear and multilinear algebra methodologies for model- and data-driven scientific computing. Second, because there is no universal single solution but rather a large panel of alternatives combining many of the various building boxes, we also consider research activities in the field of composition of parallel algorithms and data distributions to ease the investigation of this combinatorial problem toward the best algorithm for the targeted problem.

To illustrate on a single but representative example of **model-driven problems** that the joint team will address we can mention one encountered at Airbus that is related to large aero-acoustic calculations. The reduction of noise produced by aircraft during take-off and landing has a **direct societal and environmental impact on the populations (including citizen health)** located around airports. To comply with new noise regulation rules, novel developments must be undertaken to preserve the competitiveness of the European aerospace industry. In order to design and optimize new absorbing materials for acoustics and reduce the perceived sound, one must be able to simulate the propagation of an acoustic wave in an aerodynamic flow: The physical phenomenon at stake is **aero-acoustics**. The complex and chaotic nature of fluid mechanics requires simplifications in the models used. Today, we consider the flow as non-uniform only in a small part of the space (in the jet flow of the reactors mainly) which will be meshed in volume finite elements, and everywhere else the flow will be considered as uniform, and the acoustic propagation will be treated with

surface finite elements. This brings us back to the solution of a linear system with dense and sparse parts, an atypical form for which there is no "classical" solver available. We therefore have to work on the coupling of methods (direct or iterative, dense or sparse, compressed or not, etc.), and to **compose different algorithms** in order to be able to handle very large industrial cases. While there are effective techniques to solve each part independently from one another, there is no canonical, efficient solution for the coupled problem, which has been much less studied by the community. Among the possible improvements to tackle such a problem, **hybridizing simulation and learning** represents an alternative which allows one to reduce the complexity by avoiding as much as possible local refinements and therefore reduce the size of the problem.

Regarding **data-driven calculation**, climate data analysis is one of the application domains that generate huge amounts of data, either in the form of measurements or computation results. The ongoing effort between the climate modeling and weather forecasting community to mutualize digital environment, including codes and models, leads the climate community to use finer models and discretization generating an ever growing amount of data. The analysis of these data, mainly based on classical numerical tools with a strong involvement of linear algebra ingredients, is facing new scalability challenges due to this growing amount of data. Computed and measured data have **intrinsic structures** that could be naturally exploited by **low rank tensor representations** to best reveal the hidden structure of the data while addressing the scalability problem. The close link with the **CECI team at Cerfacs** will provide us with the opportunity to study novel numerical methodologies based on tensor calculation. Contributing to a better understanding of the mechanisms governing the climate change would obviously have significant **societal and economical impacts on the population**. This is just an illustration of a possible usage of our work, we could also have possibly mentioned an on-going collaboration where our tools will be used in the context of a steel company to reduce the data volume generated by IoT to be transferred on the cloud for the analysis. The methodological part described in Section 4 covers mostly two complementary topics: the first in the field of numerical scientific computing and the second in the core of computational sciences.

To sum-up, for each of the methodological contributions, we aim to find at least one dimensioning application, preferably from a societal challenge, which will allow us to validate these methods and their implementations at full-scale. The search for these applications will initially be carried out among those available at Airbus or Cerfacs, but the option of seeking them through collaborations outside the project will remain open. The ambition remains to develop generic tools whose implementations will be made accessible via their deposit in the public domain.

### 3 Research program

The methodological component of our proposal concerns the expertise for the design as well as the efficient and scalable implementation of highly parallel numerical algorithms. We intend to go from numerical methodology studies to design novel numerical schemes up to the full assessment at scale in real case academic and industrial applications thanks to advanced HPC implementations.

Our view of the research activity to be developed in **CONCACE** is to systematically assess the methodological and theoretical developments in real scale calculations mostly through applications under investigations by the industrial partners (namely Airbus Central R&T and Cerfacs).

We first consider in Section 4.1 topics concerning parallel linear and multi-linear algebra techniques that currently appear as promising approaches to tackle huge problems both in size and in dimension on large numbers of cores. We highlight the linear problems (linear systems or eigenproblems) because they are in many large scale applications the main bottleneck and the most computational intensive numerical kernels. The second research axis, presented in Section 4.2, is related to the challenge faced when advanced parallel numerical toolboxes need to be composed to easily find the best suited solution both from a numerical but also parallel performance point of view.

In short the research activity will rely on two scientific pillars, the first dedicated to the development of new mathematical methods for linear and multilinear algebra (both for model-driven and data-driven calculations). The second pillar will be on parallel computational methods enabling to easily compose in a parallel framework the packages associated with the methods developed as outcome of the first pillar. The mathematical methods from the first pillar can mathematically be composed, the challenge will be to do on large parallel computers thank to the outcome of the second pillar. We will still validate on real applications and at scale (problem and platform) in close collaborations with application experts.

### 3.1 Numerical algebra methodologies in model and data-driven scientific computing

At the core of many simulations, one has to solve a linear algebra problem that is defined in a vector space and that involves linear operators, vectors and scalars, the unknowns being usually vectors or scalars, *e.g.* for the solution of a linear system or an eigenvalue problem. For many years, in particular in model-driven simulations, the problems have been reformulated in classical matrix formalism possibly unfolding the spaces where the vectors naturally live (typically 3D PDEs) to end up with classical vectors in  $R^n$  or  $C^n$ . For some problems, defined in higher dimension (*e.g.*, time dependent 3D PDE), the other dimensions are dealt in a problem specific fashion as unfolding those dimensions would lead to too large matrices/vectors. The concace research program on numerical methodology intends to address the study of novel numerical algorithms to continue addressing the mainstream approaches relying on **classical matrix formalism** but also to investigate alternatives where the structure of the underlying problem is kept preserved and all dimensions are dealt with equally. This latter research activity mostly concerns **linear algebra in tensor spaces**. In terms of algorithmic principles, we will lay an emphasis on hierarchy as a unifying principle for the numerical algorithms, the data representation and processing (including the current hierarchy of arithmetic) and the parallel implementation towards scalability.

#### 3.1.1 Scientific computing in large size linear algebra

As an extension of our past and on-going research activities, we will continue our works on numerical linear algebra for model-driven applications that rely on classical vectorial spaces defined on  $R^n$  and  $C^n$ , where vectors and matrices are classical sparse or dense objects encountered in regular numerical linear algebra computations.

The main numerical algorithms we are interested in are:

- **Matrix decompositions** including classical ones such as the  $QR$  factorization that plays a central role in block Krylov solvers [65, 79], randomized range finder algorithms [67, 68], to name a few, as building orthonormal basis of subspaces guarantees numerical robustness. But also other factorizations, not used in classical linear algebra for model-driven calculation, such as non-negative factorization encountered in data-science for multi-variable analysis [78, 73].
- **Iterative solvers** both for linear system solutions and for eigenproblems. Regarding linear systems, we will pay a particular attention to advanced numerical techniques such as multi-level preconditioning, hybrid direct-iterative (both algebraic and PDE driven interface boundary conditions) and the solution of augmented systems (*e.g.*, Karush-Kuhn-Tucker or KKT) [85, 86]. We will investigate variants of nested subspace methods, possibly with subspace augmentation or deflation. In the multiple right-hand sides or left-hand sides cases, we will further study the possible orthogonalization variants and the trade-off between the associated parallel scalability and robustness. A particular attention will be paid to the communication hiding approaches and the investigation of their block extensions. For eigenproblem solutions, we will consider novel nested subspace techniques to further extend the numerical capabilities of the recently proposed AVCI [90, 87] technique as well as contour based integral equations (that intensively use linear systems techniques mentioned above).

In that context, we will consider the benefit of using hybridization between simulation and learning in order to reduce the complexity of classical approaches by diminishing the problem size or improving preconditioning techniques. In a longer term perspective, we will also conduct an active technological watch activity with respect to quantum computing to better understand how such a advanced computing technology can be synergized with classical scientific computing.

#### 3.1.2 Scientific computing in large dimension multi-linear algebra

This work will mostly address linear algebra problems defined in large dimensional spaces as they might appear either in model-driven simulations or data-driven calculations. In particular we will be interested in tensor vectorial spaces where the intrinsic mathematical structures of the objects have to be exploited to design efficient and effective numerical techniques.

The main numerical algorithms we are interested in are:

- Low-rank tensor decompositions for model- and data-driven, some of them rely on some numerical techniques considered in the previous section [75, 77];
- Extension of iterative numerical linear solvers (linear systems and eigensolvers) to tensor vectorial spaces to handle problems that were previously vectorized to be amenable to solution by classical linear algebra techniques;
- Study preconditioning and domain decomposition techniques suited for the solution of stochastic PDEs (encountered in some Uncertainty Quantification context) [95] leading to large dimension or preconditioning based on a low-rank approximation of the tensorization of the dense matrix in Boundary Element Method solver [63, 66, 92].

### 3.1.3 Scientific continuum between large size and large dimension

Novel techniques for large size and large dimension problems tend to reduce the memory footprint and CPU consumption through data compression such as low-rank approximations (hierarchical matrices for dense and sparse calculation, tensor decomposition [76, 93, 88]) or speed up the algorithm (fast multipole method, randomized algorithm [83, 89] [94, 68]) to reduce the time and energy to solution. Because of the compression, the genuine data are represented with lower accuracy possibly in a hierarchical manner. Understanding the impact of this lower precision data representation through the entire algorithm is an important issue for developing robust, “accurate” and efficient numerical schemes for current and emerging computing platforms from laptop commodity to supercomputers. Mastering the **trade-off between performance and accuracy** will be part of our research agenda [71, 74].

Because the low precision data representation can have diverse origins, this research activity will naturally cover the multi-precision arithmetic calculation in which the data perturbation comes entirely from the data encoding, representation and calculation in IEEE (or more exotic Nvidia GPU or Google TPU) floating point numbers. This will result in variable accuracy calculations. This general framework will also enable us to address soft error detection [62] and study possible mitigation schemes to design resilient algorithms.

## 3.2 Composition of parallel numerical algorithms from a sequential expression

A major breakthrough for exploiting multicore machine [70] is based on a data format and computational technique originally used in an out-of-core context [81]. This is itself a refinement of a broader class of numerical algorithms – namely, “updating techniques” – that were not originally developed with specific hardware considerations in mind. This historical anecdote perfectly illustrates the need to separate data representation, algorithmic and architectural concerns when developing numerical methodologies. In the recent past, we have contributed to the study of the sequential task flow (STF) programming paradigm, that enabled us to abstract the complexity of the underlying computer architecture [61, 60, 59]. In the *conca* project, we intend to go further by abstracting the numerical algorithms and their dedicated data structures. We strongly believe that combining these two abstractions will allow us to **easily compose toolbox algorithms and data representations in order to study combinatorial alternatives towards numerical and parallel computational efficiency**. We have demonstrated this potential on domain decomposition methods for solving sparse linear systems arising from the discretisation of PEDs, that has been implemented in the *maphys++* parallel package.

Regarding the abstraction of the target architecture in the design of numerical algorithms, the STF paradigm has been shown to significantly reduce the difficulty of programming these complex machines while ensuring high computational efficiency. However, some challenges remain. The first major difficulty is related to the scalability of the model at large scale where handling the full task graph associated with the STF model becomes a severe bottleneck. Another major difficulty is the inability (at a reasonable runtime cost) to efficiently handle fine-grained dynamic parallelism, such as numerical pivoting in the Gaussian elimination where the decision to be made depends on the outcome of the current calculation and cannot be known in advance or described in a task graph. These two challenges are the ones we intend to study first.

With respect to the second ingredient, namely the abstraction of the algorithms and data representation, we will also explore whether we can provide additional separation of concerns beyond that offered by a task-based design. As a seemingly simple example, we will investigate the possibility of abstracting the matrix-vector product, basic kernel at the core of many numerical linear algebra methods, to cover the case

of the fast multipole method (FMM, at the core of the ScalFMM library). FMM is mathematically a block matrix-vector product where some of the operations involving the extra-diagonal blocks with hierarchical structure would be compressed analytically. Such a methodological step forward will consequently allow the factorisation of a significant part of codes (so far completely independent because no bridge has been made upstream) including in particular the ones dealing with  $\mathcal{H}$ -matrices. The easy composition of these different algorithms will make it possible to explore the combinatorial nature of the possible options in order to best adapt them to the size of the problem to be treated and the characteristics of the target computer. \*Offering such a continuum of numerical methods rather than a discrete set of tools is part of the team's objectives\* It is a very demanding effort in terms of HPC software engineering expertise to coordinate the overall technical effort.

We intend to strengthen our engagement in reproducible and open science. Consequently, we will continue our joint effort to ensure consistent deployment of our parallel software; this will contribute to improve its impact on academic and industrial users. The software engineering challenge is related to the increasing number of software dependencies induced by the desired capability of combining the functionality of different numerical building boxes, e.g., a domain decomposition solver (such as `maphys++`) that requires advanced iterative schemes (such as those provided by `fabulous`) as well as state-of-the-art direct methods (such as `pastix`, `mumps`, or `qr_mumps`), deploying the resulting software stack can become tedious [64].

In that context, we will consider the benefit of using hybridization between simulation and learning in order to reduce the complexity of classical approaches by diminishing the problem size or improving preconditioning techniques. In a longer term perspective, we will also conduct an active technological watch activity with respect to quantum computing to better understand how such a advanced computing technology can be synergized with classical scientific computing.

## 4 Application domains

**Participants:** Emmanuel Agullo, Pierre Benjamin, Théo Briquet, Olivier Coulaud, Antoine Gicquel, Luc Giraud, Sofiane Haddad, Esragul Korkmaz, Carola Kruse, Paul Mycek, Gilles Marait, Clément Peaucelle, Guillaume Sylvand

We have a major application domain in acoustic simulations that is provided by Airbus CR & T and a few more through collaborations in the context of ongoing projects, that include: plasma simulation (ESA contract and ANR Maturation), Electric device design (ANR TensorVim) and nanoscale simulation platform (ANR Diwina).

### 4.1 Aeroacoustics Simulation

This domains is in the context of a long term collaboration with Airbus Research Centers. Wave propagation phenomena intervene in many different aspects of systems design at Airbus. They drive the level of acoustic vibrations that mechanical components have to sustain, a level that one may want to diminish for comfort reason (in the case of aircraft passengers, for instance) or for safety reason (to avoid damage in the case of a payload in a rocket fairing at take-off). Numerical simulations of these phenomena plays a central part in the upstream design phase of any such project [72]. Airbus Central R & T has developed over the last decades an in-depth knowledge in the field of Boundary Element Method (BEM) for the simulation of wave propagation in homogeneous media and in frequency domain. To tackle heterogeneous media (such as the jet engine flows, in the case of acoustic simulation), these BEM approaches are coupled with volumic finite elements (FEM). We end up with the need to solve large (several millions unknowns) linear systems of equations composed of a dense part (coming for the BEM domain) and a sparse part (coming from the FEM domain). Various parallel solution techniques are available today, mixing tools created by the academic world (such as the Mumps and Pastix sparse solvers) as well as parallel software tools developed in-house at Airbus (dense solver SPIDO, multipole solver,  $\mathcal{H}$ -matrix solver with an open sequential version available online). In the current state of knowledge and technologies, these methods do not permit to tackle the simulation of

aeroacoustics problems at the highest acoustic frequencies (between 5 and 20 kHz, upper limits of human audition) while considering the whole complexity of geometries and phenomena involved (higher acoustic frequency implies smaller mesh sizes that lead to larger unknowns number, a number that grows like  $f^2$  for BEM and  $f^3$  for FEM, where  $f$  is the studied frequency). The purpose of the study in this domain is to develop advanced solvers able to tackle this kind of mixed dense/sparse linear systems efficiently on parallel architectures.

## 5 Highlights of the year

### 5.1 Awards

- We are delighted to welcome a new member to the team. Clément Guillet has joined us as an ISFP for the 2025 campaign.
- A new release of **Composyx** has been released to the public.
- The **CONCACE** Steering Committee approved the continuation of the team, subject to the completion of the ongoing Inria scientific evaluation process.

## 6 Latest software developments, platforms, open data

### 6.1 Latest software developments

#### 6.1.1 composyx

**Name:** Numerical and parallel composability for high performance computing

**Keywords:** Numerical algorithm, Parallel computing, Linear algebra, Task-based algorithm, Dense matrix, Sparse matrix, Hierarchical matrix, FMM, C++

**Functional Description:** Composable numerical and parallel linear algebra library

**URL:** <https://gitlab.inria.fr/composyx/composyx>

**Contact:** Emmanuel Agullo

#### 6.1.2 ScalFMM

**Name:** Scalable Fast Multipole Method

**Keywords:** N-body, Fast multipole method, Parallelism, MPI, OpenMP

**Scientific Description:** ScalFMM is a software library to simulate N-body interactions using the Fast Multipole Method. The library offers two methods to compute interactions between bodies when the potential decays like  $1/r$ . The first method is the classical FMM based on spherical harmonic expansions and the second is the Black-Box method which is an independent kernel formulation (introduced by E. Darve @ Stanford). With this method, we can now easily add new non oscillatory kernels in our library. For the classical method, two approaches are used to decrease the complexity of the operators. We consider either matrix formulation that allows us to use BLAS routines or rotation matrix to speed up the M2L operator.

ScalFMM intends to offer all the functionalities needed to perform large parallel simulations while enabling an easy customization of the simulation components: kernels, particles and cells. It works in parallel in a shared/distributed memory model using OpenMP and MPI. The software architecture has been designed with two major objectives: being easy to maintain and easy to understand. There is two main parts: the management of the octree and the parallelization of the method the kernels. This new architecture allow us to easily add new FMM algorithm or kernels and new paradigm of parallelization.

The version 3.0 of the library is a partial rewriting of the version 2.0 in modern C++ ( C++17) to increase the genericity of the approach. This version is also the basic framework for studying numerical and parallel composability within Concace.

**Functional Description:** Compute N-body interactions using the Fast Multipole Method for large number of objects

**Release Contributions:** ScalFmm is a high performance library for solving n-body problems in astrophysics and electrostatics. It is based on the fast multipole method (FMM) and is highly parallel

**News of the Year:** Performance improvements in version 3.0. For the moment, this version only considers the interpolation approach. New features - the target particles can be different from the source particles - possibility to consider a non-mutual approach in the direct field - the low rank approximation of the transfer operator is taken into account.

**URL:** <https://gitlab.inria.fr/solverstack/ScalFMM>

**Publications:** [hal-00911856](#), [hal-01359458](#), [hal-00746089](#), [hal-01517153](#), [hal-01474556](#), [hal-01517153](#)

**Contact:** Olivier Coulaud

**Participants:** Olivier Coulaud, Pierre Est erie

### 6.1.3 CPPDiodon

**Name:** Parallel C++ library for Multivariate Data Analysis of large datasets.

**Keywords:** SVD, PCA, Classification

**Scientific Description:** Diodon provides executables and functions to compute multivariate data Analysis such as: Singular Value Decomposition (SVD), Principal Component Analysis (PCA) and variants (with different pre-treatments), Multidimensional Scaling (MDS), Correspondence Analysis (CoA), Canonical Correlation Analysis (CCA, future work), Multiple Correspondence Analysis (MCoA, future work). All these methods rely on a Singular Value Decomposition (SVD) of a 2D matrix. For small size matrices the SVD can be directly computed using a sequential or multi-threaded LAPACK solver such as OpenBlas or Intel MKL. For large matrices the SVD becomes time consuming and we use a Randomized Singular Value Decomposition method (rSVD) instead of the exact SVD which implementation is given by the FMR library. FMR can perform computations of the rSVD on parallel shared and distributed memory machines using adequate parallel dense linear algebra routines internally such as OpenBlas or Intel MKL on a shared memory node and Chameleon for distributed memory nodes (MPI).

**Functional Description:** Dimension reduction by multivariate data analysis. Diodon is a list of functions and drivers that implement in C++ and Python (i) pre-processing, SVD and post-processing with a wide variety of methods, (ii) random projection methods for SVD execution which allows to circumvent the time limitation in the calculation of the SVD, and (iii) a C++ implementation of the SVD with random projection to an imposed range or precision, connected to the MDS, PCA, CoA.

**Release Contributions:** Initial release of cppdiodon : a parallel C++ library for Multivariate Data Analysis of large datasets. Contains methods to compute Singular Value Decomposition (SVD), Randomized SVD, Principal Component Analysis (PCA), Multidimensional Scaling (MDS) and Correspondence Analysis (CoA). Handles text and hdf5 files. Parallel (mpi, threads, cuda) randomized SVD and EVD (for symmetric matrices) provided by FMR. Use multithreaded Lapack or Chameleon (distributed systems + GPUs).

**URL:** <https://gitlab.inria.fr/diodon/cppdiodon>

**Publications:** [tel-04468314](#), [hal-03773985](#)

**Contact:** Florent Pruvost

**Partner:** INRAE

### 6.1.4 FMR

**Name:** Fast Methods for Randomized numerical linear algebra

**Keyword:** SVD

**Scientific Description:** Fast Dense Standard and Randomized Numerical Linear Algebra is a library that allows to compute singular values or eigenvalues of large dense matrices by random linear algebra techniques. It is based on the random projection method (Gaussian or fast Hadamard/Fourier) or row/column selection (Nystrom method and variants). The library is developed in C++ and proposes a shared memory parallelization and a distributed approach with Chameleon (<https://gitlab.inria.fr/solverstack/chameleon>).

**Functional Description:** Fast Dense Standard and Randomized Numerical Linear Algebra is a library that allows to compute singular values or eigenvalues of large dense matrices by random linear algebra techniques. It is based on the random projection method (Gaussian or fast Hadamard/Fourier) or row/column selection (Nystrom method and variants). The library is developed in C++ and proposes a shared memory parallelization and a distributed approach with Chameleon (<https://gitlab.inria.fr/solverstack/chameleon>).

**URL:** <https://gitlab.inria.fr/compose/legacystack/fmr>

**Publications:** [hal-01228519](#), [hal-01334747](#), [hal-03773985](#)

**Contact:** Olivier Coulaud

**Participants:** Olivier Coulaud, Florent Pruvost

## 7 New results

**Participants:** All team members.

### 7.1 Error Estimates for Sparse Tensor Products of B-spline Approximation Spaces

This work introduces and analyzes B-spline approximation spaces defined on general geometric domains obtained through a mapping from a parameter domain. These spaces are constructed as sparse-grid tensor products of univariate spaces in the parameter domain and are mapped to the physical domain via a geometric parametrization. Both the univariate approximation spaces and the geometric mapping are built using maximally smooth B-splines. We construct two such spaces, employing either the sparse-grid combination technique or the hierarchical subspace decomposition of sparse-grid tensor products, and we prove their mathematical equivalence. Furthermore, we derive approximation error estimates and inverse inequalities that highlight the advantages of sparse-grid tensor products. Specifically, under suitable regularity assumptions on the solution, these spaces achieve the same approximation order as standard tensor product spaces while using significantly fewer degrees of freedom. Additionally, our estimates indicate that, in the case of non-tensor-product domains, stronger regularity assumptions on the solution—particularly concerning isotropic (non-mixed) derivatives—are required to achieve optimal convergence rates compared to sparse-grid methods defined on tensor-product domains.

For more details on this work we refer to [48].

### 7.2 On some orthogonalization schemes in Tensor Train format

In the framework of tensor spaces, we consider orthogonalization algorithms to generate an orthogonal basis of a tensor subspace from a set of linearly independent tensors. All variants, except for the Householder transformation, are straightforward extensions of well-known algorithms in matrix computation to tensors. In particular, we experimentally study the loss of orthogonality of six orthogonalization methods: Classical

and Modified Gram-Schmidt with (CGS2, MGS2) and without (CGS, MGS) re-orthogonalization, the Cholesky-QR, and the Householder transformation. To overcome the curse of dimensionality, we represent tensors with a low-rank approximation using the Tensor Train (TT) formalism. Additionally, we introduce recompression steps in the standard algorithm outline through the TT-round method at a prescribed accuracy. After describing the structure and properties of the algorithms, we illustrate their loss of orthogonality with numerical experiments. Although no formal proof exists at this time, we observe very clearly that the well-established properties verified over decades of research by the round error analysis community in matrix computation appear to extend to the case of low-rank tensors, with the unit round-off replaced by the TT-round accuracy. The computational analysis for each orthogonalization scheme in terms of memory requirements and computational complexity, measured as a function of the number of TT-round operations, which happens to be the most computationally expensive operation, completes the study.

For more details on this work we refer to [18].

### 7.3 A note on TT-GMRES for the solution of parametric linear systems

We study the solution of linear systems with tensor product structure using the Generalized Minimal RESidual (GMRES) algorithm. To manage the computational complexity of high dimensional problems our approach relies on low-rank tensor representation, focusing specifically on the Tensor Train format. We implement and experimentally study the TT-GMRES algorithm. Our analysis bridges the heuristic methods proposed for TT-GMRES by Dolgov in [Russian J. Numer. Anal. Math. Modelling, 28 (2013), pp. 149–172] and the theoretical framework of inexact GMRES by Simoncine and Szylid [SIAM J. Sci. Comput. 25 (2003), pp. 454–477]. This approach is particularly relevant in a scenario where a  $(d + 1)$ -dimensional problem arises from concatenating a sequence of  $d$ -dimensional problems, as in the case of a parametric linear operator or parametric right-hand-side formulation. Thus, we provide backward error bounds that link the accuracy of the computed  $(d + 1)$ -dimensional solution to the numerical quality of the extracted  $d$ -dimensional solutions. This facilitates the prescription of a convergence threshold ensuring that the  $d$ -dimensional solutions extracted from the  $(d + 1)$ -dimensional result have the desired accuracy once the solver converges. We illustrate these results with academic examples across varying dimensions and sizes. Our experiments indicate that the TT-GMRES retains the theoretical rounding error properties observed in matrix-based GMRES.

For more details on this work we refer to [17].

### 7.4 Solving eigenvalue problems in high dimensions using contour integration and Tensor Train format

In high-dimensional settings, solving eigenvalue problems is hindered by the curse of dimensionality, particularly when only a subset of eigenpairs within a prescribed spectral interval is sought. In this work, we investigate an adaptation of the FEAST algorithm, originally developed for symmetric eigenproblems based on contour integration, to computations where both operators and vectors are represented in the Tensor Train (TT) format. This representation drastically reduces memory and computational demands. We introduce an adaptive scheme for determining the projection subspace dimension by incorporating a rank-revealing Modified Gram-Schmidt procedure with pivoting tailored to TT-vectors. A perturbation-based analysis provides explicit bounds on the attainable residual accuracy, from which we derive a robust stopping criterion for the proposed TT-FEAST algorithm. Moreover, we design a continuation strategy that gradually refines convergence and rounding tolerances to effectively control memory growth during iterations. To demonstrate the effectiveness of TT-FEAST as a viable alternative to existing high-dimensional eigensolvers when a few eigenvalues are required, we present numerical experiments on problems up to twelve dimensions, including the Laplacian and a vibrational Hamiltonian operator.

For more details on this work we refer to [44].

### 7.5 A Tensor Train solver for the Magnetic Moment Method

In this work, we focus on enhancing the computational efficiency of the Magnetic Moment Method (MMM) using low-rank tensor representations, specifically the Tensor Train (TT) formats. By transforming the dense linear system generated by MMM into TT format, we achieve significant reductions in both computational cost and memory usage. Furthermore, when the problem structure allows—such as in cases with regular grids

and binary-compatible sizes—the Quantized Tensor Train (QTT) format is employed to exploit additional compression through quantization, leading to even larger performance gains. The proposed methods are tested on several problems involving a ferromagnetic part with a regular mesh. For simple cases, our approach demonstrates superior performance compared to traditional techniques. As the problem complexity increases, the TT- and QTT-based methods remain competitive, maintaining efficiency while addressing the added computational challenges.

This work is part of Amine Zekri’s PhD thesis and is carried out in the context of the TensoVim ANR project. For more details on this work we refer to [25].

## **7.6 Generalized Golub–Kahan Bidiagonalization for Nonsymmetric Saddle-Point Systems**

The generalized Golub–Kahan bidiagonalization has been used to solve saddle-point systems where the leading block is symmetric and positive definite. We extend this iterative method for the case where the symmetry condition no longer holds. We do so by relying on the known connection the algorithm has with the conjugate gradient method and following the line of reasoning that adapts the latter into the full orthogonalization method. We propose appropriate stopping criteria based on the residual and an estimate of the energy norm for the error associated with the primal variable. Numerical comparison with GMRES highlights the advantages of our proposed strategy regarding its low memory requirements and the associated implications.

For more details on this work we refer to [20].

## **7.7 A note on the partial convergence management for the solution of symmetric linear systems with multiple right-hand sides**

We consider the solution of large sparse symmetric linear systems with multiple right-hand sides available simultaneously. Based on the partial convergence detection and management, described in IB-BGMRES [Linear Algebra Appl., 419 (2006), pp. 265-285] and the breakdown-free idea discussed in [BIT Numer. Math., 57 (2017), pp. 379-403], the block conjugate residual and block conjugate gradient methods with partial convergence management are proposed. It enables to select the directions to use for extending the search space from one iteration to the next by choosing the directions that contribute the most to the residual norms. We illustrate the numerical and computational benefits of these two novel block conjugate direction variants on a set of simple academic examples enabling reproducible experiments.

For more details on this work we refer to [47].

## **7.8 A Scalable and Parameter-Robust Preconditioner for a Second Gradient of Dilation Regularization Applied to a Mechanics Problem**

We propose a rigorous analysis of the second gradient of dilation regularization as it is used in many geomechanical applications. In this method, a new primal unknown, modeling the trace of the displacement gradient, is introduced leading to a saddle-point system. The resulting balance equations depend on parameters that range on several orders of magnitude. We prove the well-posedness of the continuous and discrete problems using parameter dependent norms. This allows the definition of a robust block preconditioner for the discrete problem with respect to the parameter variations and to mesh refinement. Numerical results confirm our theoretical findings.

For more details on this work we refer to [49].

## **7.9 Neural network preconditioning: a case study for the solution of the parametric Helmholtz equation**

This work presents a hybrid numerical approach for solving linear systems arising from the discretization of the two-dimensional parametric Helmholtz equation. A convolutional neural network based on the U-Net architecture is trained in an unsupervised manner to approximate the inverse of the discretized Helmholtz operator, using a loss function involving the residual norm of the linear system. The trained network is used as

a nonlinear preconditioner within the Flexible GMRES (FGMRES) algorithm. Numerical experiments show that while the neural network is not accurate enough to act as a standalone solver, it significantly improves the convergence of FGMRES when employed as a preconditioner. The neural preconditioner demonstrates robust performance and generalization capabilities with respect to variations in the velocity field and the domain size. Comparisons with classical algebraic preconditioners based on sparsified LU factorizations indicate superior efficiency of the neural approach under equivalent conditions. We believe that the proposed method is not tied to a specific neural architecture and can be extended to other parametric PDEs.

For more details on this work we refer to [46].

### **7.10 Memory-and compute-optimized geometric multigrid GMGPolar for curvilinear coordinate representations -Applications to fusion plasma**

Tokamak fusion reactors are actively studied as a means of realizing energy production from plasma fusion. However, due to the substantial cost and time required to construct fusion reactors and run physical experiments, numerical experiments are indispensable for understanding plasma physics inside tokamaks, supporting the design and engineering phase, and optimizing future reactor designs. Geometric multigrid methods are optimal solvers for many problems that arise from the discretization of partial differential equations. It has been shown that the multigrid solver GMGPolar solves the 2D gyrokinetic Poisson equation in linear complexity and with only small memory requirements compared to other state-of-the-art solvers. In this paper, we present a completely refactored and object-oriented version of GMGPolar which offers two different matrix-free implementations. Among other things, we leverage the Sherman-Morrison formula to solve cyclic tridiagonal systems from circular line solvers without additional fill-in and we apply reordering to optimize cache access of circular and radial smoothing operations. With the Give approach, memory requirements are further reduced and speedups of four to seven are obtained for usual test cases. For the Take approach, speedups of 16 to 18 can be attained.

For more details on this work we refer to [50].

### **7.11 Complexity analysis and scalability of a matrix-free extrapolated geometric multigrid solver for curvilinear coordinates representations from fusion plasma applications**

Tokamak fusion reactors are promising alternatives for future energy production. Gyrokinetic simulations are important tools to understand physical processes inside tokamaks and to improve the design of future plants. In gyrokinetic codes such as Gysela, these simulations involve at each time step the solution of a gyrokinetic Poisson equation defined on disk-like cross sections. The authors of [14,15] proposed to discretize a simplified differential equation using symmetric finite differences derived from the resulting energy functional and to use an implicitly extrapolated geometric multigrid scheme tailored to problems in curvilinear coordinates. In this article, we extend the discretization to a more realistic partial differential equation and demonstrate the optimal linear complexity of the proposed solver, in terms of computation and memory. We provide a general framework to analyze floating point operations and memory usage of matrix-free approaches for stencil-based operators. Finally, we give an efficient matrix-free implementation for the considered solver exploiting a task-based multithreaded parallelism which takes advantage of the disk-shaped geometry of the problem. We demonstrate the parallel efficiency for the solution of problems of size up to 50 million unknowns.

For more details on this work we refer to [23].

### **7.12 Fault-tolerant numerical iterative algorithms at scale**

This work investigates how to protect numerical iterative algorithms from all types of errors that can strike at scale: fail-stop errors (a.k.a. failures) and silent errors, striking both as computation errors and memory bit-flips. We combine various techniques: detectors for computation errors, checksums for memory errors, and checkpoint/restart for failures. The objective is to minimize the expected time per iteration of the algorithm. We design a hierarchical pattern that combines and interleaves all these fault-tolerance mechanisms, and we determine the optimal periodic pattern that achieves this objective. We instantiate these

results for the performance analysis of the Preconditioned Conjugate Gradient (PCG) algorithm: we report several scenarios where the optimal pattern dramatically decreases the overhead due to error mitigation.

For more details on this work we refer to [24].

### **7.13 A filtered multilevel Monte Carlo method for estimating the expectation of cell-centered discretized random fields**

We investigate the use of multilevel Monte Carlo (MLMC) methods for estimating the expectation of discretized random fields. Specifically, we consider a setting in which the input and output vectors of numerical simulators have inconsistent dimensions across the multilevel hierarchy. This requires the introduction of grid transfer operators borrowed from multigrid methods. By adapting mathematical tools from multigrid methods, we perform a theoretical spectral analysis of the MLMC estimator of the expectation of discretized random fields, in the specific case of linear, symmetric and circulant simulators. We then propose filtered MLMC (F-MLMC) estimators based on a filtering mechanism similar to the smoothing process of multigrid methods, and we show that the filtering operators improve the estimation of both the small- and large-scale components of the variance, resulting in a reduction of the total variance of the estimator. Next, the conclusions of the spectral analysis are experimentally verified with a one-dimensional illustration. Finally, the proposed F-MLMC estimator is applied to the problem of estimating the discretized variance field of a diffusion-based covariance operator, which amounts to estimating the expectation of a discretized random field. The numerical experiments support the conclusions of the theoretical analysis even with non-linear simulators, and demonstrate the improvements brought by the F-MLMC estimator compared to both a crude MC and an unfiltered MLMC estimator.

For more details on this work we refer to [15].

### **7.14 Multilevel Monte Carlo methods for ensemble variational data assimilation**

Ensemble variational data assimilation relies on ensembles of forecasts to estimate the background error covariance matrix  $B$ . The ensemble can be provided by an Ensemble of Data Assimilations (EDA), which runs independent perturbed data assimilation and forecast steps. The accuracy of the ensemble estimator of  $B$  is strongly limited by the small ensemble size that is needed to keep the EDA computationally affordable. We investigate here the potential of the multilevel Monte Carlo (MLMC) method, a type of multifidelity Monte Carlo method, to improve the accuracy of the standard Monte-Carlo estimator of  $B$  while keeping the computational cost of ensemble generation comparable. MLMC exploits the availability of a range of discretization grids, thus shifting part of the computational work from the original assimilation grid to coarser ones. MLMC differs from the mere averaging of statistical estimators, as it ensures that no bias from the coarse resolution grids is introduced in the estimation. The implications for ensemble variational data assimilation systems based on EDAs are discussed. Numerical experiments with a quasi-geostrophic model demonstrate the potential of the approach, as MLMC yields more accurate background error covariances and reduced analysis error. The challenges involved in cycling a multilevel variational data assimilation system are identified and discussed.

For more details on this work we refer to [19].

### **7.15 Convergence analysis of overlapping domain decomposition preconditioners for nonlinear problems**

Numerical simulations of nonlinear partial differential equations often involve solving large nonlinear systems, for which Newton's method is widely employed due to its fast convergence near the solution. However, its performance can deteriorate in the presence of strong nonlinearities or poor initial guesses. Nonlinear overlapping domain decomposition methods, such as RASPEN and Substructured RASPEN (SRASPEN), have proven effective in addressing these challenges. Because SRASPEN reduces the problem size by restricting computations to a substructure, it does not update the solution outside the substructure, so that no natural initial guesses for the nonlinear local solution exists that might lead to additional inner subdomain nonlinear iterations or even prevent the local solvers to converge. In this study, we analyze the convergence of RASPEN. We show how domain decomposition improves the convergence rate of the Newton's method by highlighting the key role of the substructure on the global error contraction. Moreover, our analysis

provides insight into an inexpensive modification to SRASPEN that mitigates the lack of iterations outside the substructure. The proposed variant significantly reduces computational cost while improving overall efficiency compared to existing techniques in the literature. Numerical experiments confirm the computational performance and robustness of the improved SRASPEN, establishing it as a reliable approach for solving large-scale nonlinear systems.

This work is part of Ettaouchi El Mehdi's PhD thesis and is carried out in collaboration with Nicolas Tardieu (EDF). For more details on this work we refer to [21]

### **7.16 Robustness and reliability of state-space, frame-based modeling for thermoacoustics**

The Galerkin modal expansion is a well-known method used to develop reduced order models for thermoacoustics. A known issue is the appearance of Gibbs-type oscillations on velocity fluctuations at the interface between subdomains and at boundary conditions. Recent work of Laurent et al. (2019) and Laurent et al. (2021) have shown that it is possible to overcome this issue by using an over-completed frame, instead of a Galerkin modal basis. However, the low-order modeling based on this frame modal expansion may generate spurious modes. In this paper, the origin of these non-physical modes is identified and a method is proposed to automatically remove them from the outcome. By preventing any interaction between the physical and non-physical components, the proposed methodology drastically improves the robustness and reliability of the frame modal expansion modeling for thermoacoustics.

For more details on this work we refer to [16].

### **7.17 Juxtaposing the fourth order vibrational operator perturbation theory CVPT(4) and the adaptive VCI (A-VCI): Accuracy, vibrational resonances and polyads of C<sub>2</sub>H<sub>4</sub> and C<sub>2</sub>D<sub>4</sub>**

Ab initio prediction of anharmonic vibrational spectra produces an increasing computational overhead for larger molecules, requesting a balance between an accuracy and resources. Two complementary fundamental quantum mechanical approaches, the perturbative and variational, have various strong and weak features, depending on a specific target problem. The vibrational perturbation theory (VPT) treats weak couplings and strong resonances separately, relying on somewhat artificial criteria. In contrast, the more precise but computationally intense variational configuration interaction (VCI) method treats all couplings in universal manner. The active ongoing development of approaches to solving vibrational problems requires an update of comparative benchmarks, helping to choose the best theoretical tools for a particular target. In this work, the performance of two particular modern implementations of these methods was juxtaposed: the second and fourth order operator canonical perturbation theory CVPT(2,4) and a recently proposed adaptive vibrational configuration interaction method (A-VCI). Two practically important C<sub>2</sub>H<sub>4</sub> and C<sub>2</sub>D<sub>4</sub> molecules and an accurate CCSD(T)/cc-pVQZ four-body sextic normal mode PES were employed for benchmarking. The comprehensive picture of vibrational resonances and the polyad quantum number was revealed. A new quadratic resonance criterion is proposed and its efficiency in elucidating polyad structures is demonstrated. A striking observation was made that CVPT(2) often produces better predictions of fundamental frequencies, while CVPT(4) demonstrates an excellent level of correlation with A-VCI results for both fundamentals and two-quanta states.

For more details on this work we refer to [22]

### **7.18 Approximation Algorithms for Scheduling with/without Deadline Constraints where Rejection Costs are Proportional to Processing Times**

In this work, we address two offline job scheduling problems, where jobs can either be processed on a limited supply of energy-efficient machines on the edge, or offloaded to an unlimited supply of energy-inefficient machines on the cloud (called rejected in our context). The goal is to minimize the total energy consumed in processing all tasks. We consider a first scheduling problem with no due date (or deadline) constraints, and we formulate it as a scheduling problem with rejection, where the cost of rejecting a job is directly proportional to its processing time. We introduce a novel  $5/4(1+\varepsilon)$  approximation algorithm BEKP by

associating it with a Multiple Subset Sum problem for this version. Our algorithm is an improvement over the existing literature, which provides a  $(3/2 - 1/2m)$  approximation for scenarios with arbitrary rejection costs. In the second scheduling problem, jobs have due date (or deadline) constraints, and the goal is to minimize the weighted number of late jobs. In our context, if a job is late, it is offloaded (rejected) to an energy-inefficient machine on the cloud, which incurs a cost directly proportional to its processing time of the job. We position this problem in the literature, and introduce a novel  $(1 - (m - 1)^m/m^m)$ -approximation algorithm MDP for this version, where we got our inspiration from an algorithm for the interval selection problem with a  $(1 - m^m/(m + 1)^m)$  approximation ratio for arbitrary rejection costs. We evaluate and discuss the effectiveness of our approaches through a series of experiments, comparing them to existing algorithms.

For more details on this work we refer to [14]

## 7.19 Guix-HPC Activity Report 2023-2024

Guix-HPC is a collaborative effort to bring reproducible software deployment to scientific workflows and high-performance computing (HPC). Guix-HPC builds upon the GNU Guix software deployment tools and aims to make them useful for HPC practitioners and scientists concerned with dependency graph control and customization and, uniquely, reproducible research. This report—our seventh report!—highlights key achievements of Guix-HPC between our previous report a year ago and today, February 2025. This year was marked by exciting developments for HPC and reproducible workflows. Significant advances were made in integrating Guix into the complex software landscape of HPC, taking the roles of software manager, workflow execution engine, backend for generating container images, or provider for the complete operating system layer. Support for reproducing computations from the past was also much improved. And, as usual, we have been using Guix for research, and teaching other researchers how to get started.

For more details on this work we refer to [42].

## 8 Partnerships and cooperations

**Participants:** All permanent members.

### 8.1 European initiatives

#### 8.1.1 H2020 projects

##### EoCoE-3

**Title:** Energy oriented Centre of Excellence for computer applications

**Duration:** 2024-2026

**Coordinator:** CEA

**Inria coordinator:** Bruno Raffin

**CONCACE contact:** Emmanuel Agullo

##### Partners:

- Agenzia Nazionale per le Nuove Tecnologie, l’Energia e lo Sviluppo Economico Sostenibile (Italy)
- Barcelona Supercomputing Center - Centro Nacional de Supercomputacion (Spain)
- Centre Europeen de Recherche et de Formation Avancee en Calcul Scientifique (France)
- Centre National de la Recherche Scientifique CNRS (France)
- Commissariat a l’Energie Atomique et aux Energies Alternatives (France)

- Consiglio Nazionale delle Ricerche (Italy)
- Forschungszentrum Julich GmbH (Germany)
- Fraunhofer Gesellschaft zur Foerderung der Angewandten Forschung E.V. (Germany)
- Inria
- Max-Planck-Gesellschaft zur Forderung der Wissenschaften EV (Germany)
- Rheinisch-Westfaelische Technische Hochschule Aachen (Germany)
- Universita degli Studi di Roma Torvergata (Italy)
- Universita degli Studi di Trento (Italy)
- Universite Libre de Bruxelles (Belgium)
- Universite Paris-Sud (France)

**Inria contact:** Bruno Raffin (Datamove)

**Summary:** The Concace team (Inria, Cerfacs) participates in the Energy-oriented Centre of Excellence (EoCoE-III), starting in January 2024. The project applies cutting-edge exascale computational methods in its mission to accelerate the transition to the production, storage and management of clean, decarbonized energy. EoCoE-III is anchored in the High Performance Computing (HPC) community and targets research institutes and key commercial players who develop and enable energy-relevant numerical models to be run on exascale supercomputers, demonstrating their benefits for the net-zero energy transition. The project will draw on the experience of two successful previous projects EoCoE-I and -II, where a large set of diverse computer applications from four such energy domains achieved significant efficiency gains thanks to a multidisciplinary expertise in applied mathematics and supercomputing. EoCoE-III channels its efforts into 5 exascale lighthouse applications in the low-carbon sectors of Energy Materials, Water, Wind and Fusion. This multidisciplinary effort will harness innovations in computer science and mathematical algorithms within a tightly integrated co-design approach to overcome performance bottlenecks and to anticipate HPC hardware developments. A world-class consortium of 16 complementary partners forms a unique network of expertise in energy science, scientific computing and HPC, including 3 leading European supercomputing centres.

## 8.2 National initiatives

### MAMBO

**Duration:** 2018 – 2026

**CONCACE contact:** Guillaume Sylvand

**Funding:** DGAC

**Partners:**

- CEA
- Inria
- CNRS

**Summary:** MAMBO ("Méthodes Avancées pour la Modélisation du Bruit moteur et aviOn") is a project funded by the DGAC, bringing together 21 academic and industrial partners from France and Europe, including Airbus, Cerfacs, and Inria. For Inria, the key challenge of this project lies in addressing new research problems and developing innovative numerical methods by enhancing the capabilities of its software, particularly in the field of high-performance computing.

## SANTANA

**Duration:** 2024 - 2027

**CONCACE contact:** Carola Kruse

**Funding:** DGAC

**Partners:**

- Airbus
- Cerfacs
- DLR
- ONERA

**Summary:** Santana is a project funded by the DGAC that focuses on the development and enhancement of the CODA solver, which is mainly used for the computation of aerodynamic effects. Within this CFD software, the Newton-Krylov solver has been identified as the most computationally expensive component. The role of Cerfacs/Concace is to optimize the linear solution step in the Newton-Krylov solver, with the goal of achieving substantial reductions in computation time.

## PEPR Numpex

**Duration:** 2023 – 2028

**CONCACE contact:** Emmanuel Agullo, Luc Giraud

**Funding:** ANR

**Partners:**

- CEA
- Inria
- CNRS

**Summary:** NumPEX is a French program dedicated to Exascale: High-performance computing (HPC), high-performance data analytics (HPDA), and Artificial Intelligence (AI) pose significant challenges across scientific, societal, economic, and ethical realms. These technologies, including modeling and data analysis, are crucial decision support tools addressing societal issues and competitiveness in French research and development. Digital resources, essential across science and industry, demand high-performance hardware. HPC enables advanced modeling, while HPDA handles heterogeneous and massive data. The solution to exploding demand is the upcoming “exascale” computers, a new generation with extraordinary capabilities.

In this context, the French Exascale program NumPEX aims at designing and developing the software components that will equip future exascale machines. NumPEX will deliver Exascale-grade numerical methods, softwares, and training, allowing France to remain one of the leaders in the field. It will contribute to take bridging the gap between cutting-edge software development and application domains to prepare the major scientific and industrial application codes to fully exploit the capabilities of these machines. Application domains of the NumPEX program include, but are not limited to, weather forecasting and climate, aeronautics, automotive, astrophysics, high energy physics, material science, energy production and management, biology and health.

Numpex is organized in 7 scientific pillar projects, we are directly involved in two of them namely:

Exa-MA : Methods and Algorithms for Exascale;

Exa-SofT : HPC softwares and tools.

**TensorVIM****Duration:** 2023 – 2026**Coordinator:** LAPLACE**CONCACE contact:** Olivier Coulaud**Funding:** ANR**Partners:**

- Inria
- LAPLACE
- G2ELaB

**Summary:** The aim of this project is to develop high-performance computational tools for the rapid implementation of low-frequency electromagnetic simulations for electrical applications. We consider an approach based on volume integral methods using low-rank approximations. Instead of using classical compression techniques such as the fast multipole method or the hierarchical matrix approach, we propose to investigate the use of low-rank tensors to accelerate the computation of the solution of the linear system. The tools developed will be used for the modeling of various devices (PCB modeling, Electrical Machines) with the main goal of improving their energy performance.

**Maturation****Title:** MASSively parallel sparse grid PIC algorithms for low TemperatURE plAsmas SimulaTIONs**Duration:** 2023 – 2026**Coordinator:** Laurent Garrigues (Laplace)**CONCACE contact:** Luc Giraud**Funding:** ANR**Partners:**

- Laplace Lab
- IMT
- Inria

**Summary:** The simulation under real conditions of partially magnetized low temperature plasmas by Lagrangian approaches, though using powerful Particle-In-Cell (PIC) techniques supplemented with efficient high-performance computing methods, requires considerable computing resources for large plasma densities. This is explained by two main limitations. First, stability conditions that constrain the numerical parameters to resolve the small space and time scales. These numerical parameters are the mesh size of the grid used to compute the electric field and the time step between two consecutive computations. Second, PIC methods rely on a sampling of the distribution function by numerical particles whose motion is time integrated in the self-consistent electric field. The PIC algorithm remains close to physics and offers an incomparable efficiency with regard to Eulerian methods, discretizing the distribution function onto a mesh. It is widely and successfully operated for the discretization of kinetic plasma models for more than 40 years. Nonetheless, to spare the computational resources, the number of numerical particles is limited compared to that of the physical particles. Inherent to this “coarse” sampling, PIC algorithms produce numerical approximations prone to statistical fluctuations that vanish slowly with the mean number of particles per cell. The mesh accessible on typical high performance computing machines may  $10^9$  cells, which brings the mesh size close to the scale of the physics, but the mean number of numerical particles in each cell shall be limited, to mitigate the memory footprint as well as the computational time. A breakthrough is therefore necessary to reduce

the computational resources by orders of magnitude and make possible the use of explicit PIC method for large scale and/or densities for 3D computations.

This is the issue addressed within the MATURATION project aiming at introducing a new class of PIC algorithms with an unprecedented computational efficiency, by analyzing and improving, parallelizing and optimizing as well as benchmarking, in the demanding context of partially magnetized low temperature plasmas through 2D large scale and 3D computations, a method recently proposed in the literature, based on a combination of sparse grid techniques and PIC algorithm.

## Diwina

**Title:** Magnetic Digital Twins for Spintronics : nanoscale simulation platform

**Duration:** 2023 – 2026

**Coordinator:** Institut Neel

**CONCACE contact:** Olivier Coulaud

**Funding:** ANR

### Partners:

- CMAP
- Inria
- Institut Neel
- SPINTEC

**Summary:** The DiWiNa project aims at developing a unified open-access platform for spintronic numerical twins, ie, codes for micromagnetic/spintronic simulations with sufficiently-high reliability and speed so that they can be trusted and used as reality. The simulations will be bridged to the advanced microscopy techniques used by the community, through plugins to convert the statics or time-resolved 3D vector-fields into contrast maps for the various techniques, including their experimental transfer functions. To achieve this, we bring together experts from different disciplines to address the various challenges: spintronics for the core simulations, mathematics for trust, algorithmics for speed, experimentalists for the bridge with microscopy. Practical work consists of checking the time-integration stability of spintronic torque involved in the dynamics when implemented in the versatile finite-element framework, improve the calculation speed through advanced libraries, build the bridge with microscopies through rendering tools, and encapsulate these three key ingredients into a user-friendly Python ecosystem. Through open-access and versatile user-friendly encapsulation, we expect that this platform is suited to serve the needs of the entire physics and engineering community of spintronics. The platform will be unique in its features, ranging from simulation to the direct and practical comparison with experiments. It will contribute to reduce considerably the number of experimental screening for the faster development of new spintronic devices, which are expected to play a key role in energy saving.

## 9 Dissemination

All permanent members

### 9.1 Promoting scientific activities

#### 9.1.1 Scientific events: organisation

##### General chair, scientific chair

- E. Agullo: co-chair of [Compas 2025 - Bordeaux, June 2025](#)

### 9.1.2 Scientific events: selection

#### Member of the conference program committees

- PDSEC: Olivier Coulaud, Luc Giraud,
- SC25: Emmanuel Agullo
- Luc Giraud is member of the Gene Golub SIAM Summer School. The thirteen Gene Golub SIAM Summer School was entitled "Frontiers in multi-dimensional pattern formation", Concordia University - Montreal, Quebec, CA - August 4th to 15th, 2025

### 9.1.3 Journal

**Reviewer - reviewing activities** ACM TOMS, SIAM SISC, SIAM SIMAX, International Journal for Numerical Methods in Engineering.

### 9.1.4 Invited talks

- A journey through some numerical linear algebra algorithms with variable accuracy storage; Luc Giraud, Emmanuel Agullo, Olivier Coulaud, Martina Iannacito, Mohammad Issa, Gilles Marait, Miroslav Rozložnik CAS-ANLA 2025 - The Chinese Academy of Sciences Workshop on Approximate computing in Numerical Linear Algebra, Apr 2025, Beijing, China.

### 9.1.5 Scientific expertise

- Emmanuel Agullo is member of the Cerfacs Evaluation committee
- Luc Giraud is
  - member of the board on Modelization, Simulation and data analysis of the Competitiveness Cluster for Aeronautics, Space and Embedded Systems.
  - member of the scientific council of the ONERA Lab LMA2S (Laboratoire de Mathématiques Appliquées à l'Aéronautique et au Spatial).
  - member of the scientific council of GDR Calcul,
  - scientific advisor at Cerfacs.
- Carola Kruse referee for Icelandic Research Fund proposals.
- Guillaume Sylvand is
  - expert in Numerical Simulation and HPC at Airbus.
  - member of the scientific council of the ORAP.

### 9.1.6 Research administration

- Emmanuel Agullo is member of the Technological Development Commission (CDT) and Bureau du Comité des Projets (BCP) at the Inria Centre at the University of Bordeaux.
- Luc Giraud is
  - techniques pilot for the expert group for the evaluation of French research entities (UMRs and EAs) relatively to the protection of scientific and technological properties (PPST) on information and communication sciences and technologies (STIC),
  - the representative of Inria at the GENCI evaluation and ressource allocation committees,
  - scientific expert of the GENCI committee on scientific computing.

## 9.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

- Post graduate level/Master:
  - E. Agullo: Numerical algorithms 20h, advanced Numerical Numerical Algebra 8h, and Implementation of HPC dense linear algebra kernels 8h, at Bordeaux INP (ENSEIRB-MatMeca).
  - L. Giraud: Introduction to intensive computing and related programming tools 20h, INSA Toulouse; Advanced numerical linear algebra 10h, ENSEEIHT Toulouse.
  - C. Kruse: Iterative methods in linear algebra, 28h, ENSEEIHT Toulouse.
  - P. Mycek: Multifidelity methods 14h, ModIA (cursus en alternance, INSA/N7), Toulouse.

### 9.2.1 Supervision

- PhD in progress: Alexandre Malhene; Abstraction of subspace methods in numerical linear algebra; started October 2024, E. Agullo, L. Giraud.
- PhD in progress: Hugo Dodelin; Abstraction of parallel execution models; started October 2024, E. Agullo, O. Coulaud.
- PhD in progress: Théo Briquet; machine learning techniques for rank prediction of  $\mathcal{H}$ -matrices; started October 2023, L. Giraud, P. Mycek, G. Sylvand.
- PhD in progress: El Mehdi Ettaouchi; nonlinear domain decomposition techniques in geosciences; started March 2023, L. Giraud, C. Kruse, N. Tardieu (EDF).
- PhD in progress: Sai Aakash Dasari; Scalable multigrid methods for tokamak geometries; started Oct. 2024, C. Kruse, P. Mycek.
- PhD in progress: Antoine Gicquel; Acceleration of the matrix-vector product by the fast multipole method for heterogeneous machine clusters; started Nov. 2023, O. Coulaud, B. Bramas.
- PhD in progress: Andrea Lagardère; Méthode Quasi-Trefftz Couplée pour l'Aéroacoustique; started April 2024, G. Sylvand, S. Tordeux.
- PhD in progress: Clément Peaucelle; Composabilité en Algèbre Linéaire Haute Performance - Application à l'Aéroacoustique et à l'Électromagnétisme; started Jan. 2025, E. Agullo, G. Sylvand.
- PhD in progress: Amine Zekri ; Low-rank tensor solver for magnetostatic problems for electric power applications, started October 2023; O. Coulaud, J.R. Poirier
- PhD thesis defended on December 12, 2025: Atte Tori; Towards a fast task-based parallel tensor solver for high-dimensional problems, O. Coulaud, O. Kaya (LISN, Paris-Saclay) ; Jury: Grey Ballard (Reviewer), Julien Langou, Professor (Reviewer), Alfredo Buttari (Examiner), Thomas Hérault (Examiner), Mariya Ishteva (Examiner), Samuel Thibault (Examiner)

### 9.2.2 Juries

#### PhD defense

- El Hachimi Anas, "Tensor-Based Computational Methods: Algorithms, Theory, and Applications"; Spécialité : Mathématiques Appliquées, Université du Littoral Côte d'Opale and Université Mohammed VI Polytechnique, Maroc; referees: Nicolas Gillis, Université de Mons - Luc Giraud, Inria, Stefano Serra Capizzano - Université Insubria; members: Lahcen Maniar - Université Cadi Ayyad Marrakech, Hassane Sadok (president) - Université du Littoral Côte d'Opale, Françoise Tisseur - Université de Manchester, Khalide Jbilou - Université du Littoral Côte d'Opale, Ahmed Ratnani - Université Mohammed 6 Polytechnique; Jul. 22, 2025.

- Antoine Ronsain, "Modélisation robuste des politiques climatiques : Une analyse critique et stochastique des modèles DICE et RICE"; Spécialité : Mathématiques et Applications, ENAC-LAB - Laboratoire de Recherche ENAC; referees: Aude Pommeret - Université Savoie Mont Blanc, Christine Solnon - INSA Lyon (president), Cyril Allignol - École Nationale de l'Aviation Civile, Julien Lefevre - CIRED, Alexandre Gondran - École Nationale de l'Aviation Civile, Estelle Malavolti - École Nationale de l'Aviation Civile, Pierre Benjamin - Airbus; Dec. 9, 2025.
- Mouhssine Abdellatif, "Vector extrapolation methods with applications to geometric multigrid and nonlinear least-squares problems"; Spécialité : Mathématiques Appliquées, University of the Littoral Opal Coast and Mohammed VI Polytechnic University; referees: Kees Vuik - Delft University of Technology, Martin Gander - University of Geneva, Luc Giraud - Inria; members: Lahcen Maniar (president) - Cadi Ayyad University, Carole Rosier - University of the Littoral Opal Coast, Ahmed Ratnani - Mohammed VI Polytechnic University, Hassane Sadok - University of the Littoral Opal Coast; Dec. 29, 2025.

#### HDR defense

- Nicole Spilanne, "High Performance Krylov Subspace Solvers with Preconditioning and Deflation"; Habilitation à Diriger des Recherches Institut Polytechnique de Paris CMAP (CNRS, École polytechnique); referees: Grégoire Allaire, Professor at École polytechnique - Martin Gander, Professor at Université de Genève - Yousef Saad, Professor at University of Minnesota; members: Stéphanie Chaillat, CNRS Senior Researcher at EPFL - Marc Embree, Professor at Virginia Tech - Virginie Erlacher, Professor at École des Ponts - Luc Giraud, Senior Researcher at Inria (President) - Laura Grigori, Professor at EPFL - Axel Klawonn, Professor at University of Cologne; Oct. 24, 2025.

## 10 Scientific production

### 10.1 Major publications

- [1] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner and T. Takahashi. 'Task-Based FMM for Multicore Architectures'. In: *SIAM Journal on Scientific Computing* 36.1 (2014), pp. 66–93. DOI: [10.1137/130915662](https://doi.org/10.1137/130915662). URL: <https://hal.inria.fr/hal-00911856>.
- [2] E. Agullo, A. Buttari, A. Guermouche, J. Herrmann and A. Jégou. 'Task-based parallel programming for scalable matrix product algorithms'. In: *ACM Transactions on Mathematical Software* (2023). URL: <https://hal.science/hal-03936659>.
- [3] E. Agullo, L. Giraud and L. Poirel. 'Robust preconditioners via generalized eigenproblems for hybrid sparse linear solvers'. In: *SIAM Journal on Matrix Analysis and Applications* 40.2 (2019), pp. 417–439. DOI: [10.1137/17M1153765](https://doi.org/10.1137/17M1153765). URL: <https://hal.inria.fr/hal-02074474>.
- [4] E. Bourne, P. Leleux, K. Kormann, C. Kruse, V. Grandgirard, Y. Güçlü, M. J. Kühn, U. Rüde, E. Sonnendrücker and E. Zoni. *Solver comparison for Poisson-like equations on tokamak geometries*. 23rd Sept. 2022. URL: <https://cea.hal.science/cea-03786723>.
- [5] B. Bramas, O. Coulaud and G. Sylvand. 'Time-domain BEM for the wave equation on distributed-heterogeneous architectures: A blocking approach'. In: *Parallel Computing* 49 (July 2015), pp. 66–82. DOI: [10.1016/j.parco.2015.07.005](https://doi.org/10.1016/j.parco.2015.07.005). URL: <https://hal.inria.fr/hal-01200727>.
- [6] F. Casenave, A. Ern and G. Sylvand. 'Coupled BEM-FEM for the convected Helmholtz equation with non-uniform flow in a bounded domain'. In: *Journal of Computational Physics* 257.A (15th Jan. 2014), pp. 627–644. DOI: [10.1016/j.jcp.2013.10.016](https://doi.org/10.1016/j.jcp.2013.10.016). URL: <https://hal.science/hal-00805477>.
- [7] S. Cools, E. F. Yetkin, E. Agullo, L. Giraud and W. Vanroose. 'Analyzing the Effect of Local Rounding Error Propagation on the Maximal Attainable Accuracy of the Pipelined Conjugate Gradient Method'. In: *SIAM Journal on Matrix Analysis and Applications* 39.1 (13th Mar. 2018), pp. 426–450. DOI: [10.1137/17M1117872](https://doi.org/10.1137/17M1117872). URL: <https://hal.inria.fr/hal-01753411>.
- [8] V. Darrigrand, A. Dumitrescu, C. Kruse and U. Rüde. 'Inexact inner-outer Golub–Kahan bidiagonalization method: A relaxation strategy'. In: *Numerical Linear Algebra with Applications* (12th Dec. 2022). DOI: [10.1002/nla.2484](https://doi.org/10.1002/nla.2484). URL: <https://hal.inria.fr/hal-03960074>.

- [9] D. A. Di Pietro, P. Matalon, P. Mycek and U. Rde. ‘High-order multigrid strategies for HHO discretizations of elliptic equations’. In: *Numerical Linear Algebra with Applications* (22nd June 2022). DOI: [10.1002/nla.2456](https://doi.org/10.1002/nla.2456). URL: <https://hal.science/hal-03531293>.
- [10] A. Franc, P. Blanchard and O. Coulaud. ‘Nonlinear mapping and distance geometry’. In: *Optimization Letters* 14.2 (2020), pp. 453–467. DOI: [10.1007/s11590-019-01431-y](https://doi.org/10.1007/s11590-019-01431-y). URL: <https://hal.inria.fr/hal-02124882>.
- [11] L. Giraud, Y.-F. Jing and Y. Xiang. ‘A block minimum residual norm subspace solver with partial convergence management for sequences of linear systems’. In: *SIAM Journal on Matrix Analysis and Applications* 43.2 (2022), pp. 710–739. DOI: [10.1137/21m1401127](https://doi.org/10.1137/21m1401127). URL: <https://hal.inria.fr/hal-03546496>.
- [12] Y. Jang, J. Bonelle, C. Kruse, F. Hlsemann and U. Ruede. *Fast Linear Solvers for Incompressible CFD Simulations with Compatible Discrete Operator Schemes*. 27th Apr. 2023. URL: <https://hal.science/hal-04087358>.
- [13] M. Odunlami, V. Le Bris, D. Bgu, I. Baraille and O. Coulaud. ‘A-VCI: A flexible method to efficiently compute vibrational spectra’. In: *Journal of Chemical Physics* 146.21 (7th June 2017). DOI: [10.1063/1.4984266](https://doi.org/10.1063/1.4984266). URL: <https://hal.inria.fr/hal-01534134>.

## 10.2 Publications of the year

### International journals

- [14] O. Beaumont, R. Bouzel, L. Eyraud-Dubois, E. Korkmaz, L. L. Pilla and A. van Kempen. ‘Approximation Algorithms for Scheduling with/without Deadline Constraints where Rejection Costs are Proportional to Processing Times’. In: *IEEE Transactions on Parallel and Distributed Systems* 36.12 (Dec. 2025), pp. 2596–2608. DOI: [10.1109/TPDS.2025.3605674](https://doi.org/10.1109/TPDS.2025.3605674). URL: <https://hal.science/hal-04745701> (cit. on p. 20).
- [15] J. Briant, P. Mycek, M. Destouches, O. Goux, S. Gratton, S. Grol, E. Simon and A. T. Weaver. ‘A filtered multilevel Monte Carlo method for estimating the expectation of cell-centered discretized random fields’. In: *International Journal for Uncertainty Quantification* (2025). DOI: [10.1615/Int.J.UncertaintyQuantification.2025057754](https://doi.org/10.1615/Int.J.UncertaintyQuantification.2025057754). URL: <https://hal.science/hal-04337114>. In press (cit. on p. 18).
- [16] M. Cances, L. Giraud, M. Bauerheim, L. Y. Gicquel and F. Nicoud. ‘Robustness and reliability of state-space, frame-based modeling for thermoacoustics’. In: *Journal of Computational Physics* 520 (2025), p. 113472. DOI: [10.1016/j.jcp.2024.113472](https://doi.org/10.1016/j.jcp.2024.113472). URL: <https://hal.science/hal-04734953> (cit. on p. 19).
- [17] O. Coulaud, L. Giraud and M. Iannacito. ‘A note on TT-GMRES for the solution of parametric linear systems’. In: *Electronic Transactions on Numerical Analysis* 62 (27th Jan. 2025), pp. 163–187. DOI: [10.1553/etna\\_vol62s163](https://doi.org/10.1553/etna_vol62s163). URL: <https://inria.hal.science/hal-04986230> (cit. on p. 15).
- [18] O. Coulaud, L. Giraud and M. Iannacito. ‘On some orthogonalization schemes in Tensor Train format’. In: *BIT Numerical Mathematics* 65.45 (31st Oct. 2025). DOI: [10.1007/s10543-025-01086-5](https://doi.org/10.1007/s10543-025-01086-5). URL: <https://inria.hal.science/hal-05343109> (cit. on p. 15).
- [19] M. Destouches, P. Mycek, S. Grol, A. Weaver, S. Gratton and E. Simon. ‘Multilevel Monte Carlo methods for ensemble variational data assimilation’. In: *Nonlinear Processes in Geophysics* 32.2 (23rd June 2025), pp. 167–187. DOI: [10.5194/npg-32-167-2025](https://doi.org/10.5194/npg-32-167-2025). URL: <https://hal.science/hal-04869802> (cit. on p. 18).
- [20] A. Dumitrasc, C. Kruse and U. Rde. ‘Generalized Golub–Kahan Bidiagonalization for Nonsymmetric Saddle-Point Systems’. In: *SIAM Journal on Matrix Analysis and Applications* 46.1 (13th Feb. 2025), pp. 370–392. DOI: [10.1137/23M160760X](https://doi.org/10.1137/23M160760X). URL: <https://hal.science/hal-04963962> (cit. on p. 16).

- [21] E. M. Ettaouchi, L. Giraud, C. Kruse and N. Tardieu. ‘Convergence analysis of overlapping domain decomposition preconditioners for nonlinear problems’. In: *SIAM Journal on Scientific Computing* 47.6 (17th Nov. 2025), A3464–A3493. DOI: [10.1137/25M1735585](https://doi.org/10.1137/25M1735585). URL: <https://hal.science/hal-05412939> (cit. on p. 19).
- [22] V. Le Bris, S. Krasnoshchekov, E. Dobrolyubov, I. Efremov, I. Polyakov, O. Coulaud and D. Bégué. ‘Juxtaposing the fourth order vibrational operator perturbation theory CVPT(4) and the adaptive VCI (A-VCI): Accuracy, vibrational resonances and polyads of C<sub>2</sub>H<sub>4</sub> and C<sub>2</sub>D<sub>4</sub>’. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 346 (Dec. 2025), p. 109586. DOI: [10.1016/j.jqsrt.2025.109586](https://doi.org/10.1016/j.jqsrt.2025.109586). URL: <https://univ-pau.hal.science/hal-05283595> (cit. on p. 19).
- [23] P. Leleux, C. Schwarz, M. Kühn, C. Kruse and U. Rüdé. ‘Complexity analysis and scalability of a matrix-free extrapolated geometric multigrid solver for curvilinear coordinates representations from fusion plasma applications’. In: *Journal of Parallel and Distributed Computing* 205 (Nov. 2025), p. 105143. DOI: [10.1016/j.jpdc.2025.105143](https://doi.org/10.1016/j.jpdc.2025.105143). URL: <https://hal.science/hal-05164801> (cit. on p. 17).
- [24] A. Tremodeux, A. Benoit, E. Agullo, T. Herault, L. Giraud and Y. Robert. ‘Fault-tolerant numerical iterative algorithms at scale’. In: *International Journal of High Performance Computing Applications* (2025). URL: <https://inria.hal.science/hal-05234063> (cit. on p. 18).
- [25] A. Zekri, O. Coulaud, L. Giraud and J.-R. Poirier. ‘A Tensor Train solver for the Magnetic Moment Method’. In: *IEEE Transactions on Magnetics* (2025), pp. 1–1. DOI: [10.1109/TMAG.2025.3618391](https://doi.org/10.1109/TMAG.2025.3618391). URL: <https://hal.science/hal-05311300> (cit. on p. 16).

#### Invited conferences

- [26] L. Giraud, E. Agullo, O. Coulaud, M. Iannacito, M. Issa, G. Marait and M. Rozloznik. ‘A journey through some numerical linear algebra algorithms with variable accuracy storage’. In: *CAS-ANLA 2025 - The Chinese Academy of Sciences Workshop on Approximate computing in Numerical Linear Algebra*. Beijing, China, 23rd Apr. 2025. URL: <https://hal.science/hal-05049142>.

#### International peer-reviewed conferences

- [27] H. Godé, C. Kruse, R. Angersbach, H. Köstler, M. Bauerheim and U. Rüdé. ‘Comparison of Multigrid and Machine Learning-Based Poisson Solvers’. In: *Parallel Processing and Applied Mathematics 15th International Conference, PPAM 2024*. PPAM 2024 - Parallel Processing and Applied Mathematics. Vol. Lecture Notes in Computer Science. Lecture Notes in Computer Science LNCS-15581. Ostrava, Czech Republic: Springer Nature Switzerland, 1st Apr. 2025, pp. 174–189. DOI: [10.1007/978-3-031-85703-4\\_12](https://doi.org/10.1007/978-3-031-85703-4_12). URL: <https://hal.science/hal-05164864>.
- [28] A.-K. M. E. Maarouf, L. Giraud, A. Guermouche and T. Guignon. ‘Sparse Matrix Ordering for Fine Grain Parallel Triangular Solve Using SIMD’. In: *Parallel Processing and Applied Mathematics (PPAM 2024)*. PPAM 2024 - 15th International Conference on Parallel Processing & Applied Mathematics. Vol. LNCS. Lecture Notes in Computer Science 15579. Ostrava, Czech Republic: Springer Nature Switzerland, 1st Apr. 2025, pp. 51–64. DOI: [10.1007/978-3-031-85697-6\\_4](https://doi.org/10.1007/978-3-031-85697-6_4). URL: <https://hal.science/hal-05185404>.
- [29] A. Torri, P. Dominikowski, B. Pointal, O. Kaya, L. Lima Pilla and O. Coulaud. ‘Near-Optimal Contraction Strategies for the Scalar Product in the Tensor-Train Format’. In: *Euro-Par 2025: Parallel Processing*. Euro-Par 2025 - 31 International European Conference on Parallel and Distributed Computing. Vol. 15902. Lecture Notes in Computer Science. Dresden, Germany: Springer Nature Switzerland, 22nd Aug. 2025, pp. 63–77. DOI: [10.1007/978-3-031-99872-0\\_5](https://doi.org/10.1007/978-3-031-99872-0_5). URL: <https://hal.science/hal-05285400>.
- [30] Y. Xiang and L. Giraud. ‘Convolution neural operator preconditioning for the solution of some heterogeneous PDES’. In: *DTE - AICOMAS 2025 - Digital Twins in Engineering & Artificial Intelligence and Computational Methods in Applied Science*. Paris, France, 21st Feb. 2025. URL: <https://inria.hal.science/hal-04961884>.

### National peer-reviewed Conferences

- [31] H. Dodelin. ‘Peut-on exprimer un produit de matrices distribué haute performance avec un modèle Map-Reduce ?’ In: *COMPAS2025*. Conférence francophone d’informatique en Parallélisme, Architecture et Système (COMPAS 2025). COMPAS2025. Bordeaux, France, 11th May 2025. URL: <https://hal.science/hal-05063477>.
- [32] A. Gicquel. ‘A composable abstraction of hierarchical methods for matrix-vector product acceleration’. In: *COMPAS 2025 - Conférence francophone d’informatique en Parallélisme, Architecture et Système*. Bordeaux, France, 24th June 2025. URL: <https://inria.hal.science/hal-05185249>.
- [33] E. Korkmaz, E. Agullo and G. Sylvand. ‘High performance solvers for aeroacoustic simulations for Airbus’. In: *COMPAS 2025 - Conférence francophone d’informatique en Parallélisme, Architecture et Système*. Bordeaux, France, 24th June 2025. URL: <https://inria.hal.science/hal-05175741>.

### Conferences without proceedings

- [34] A. Gicquel. ‘A composable abstraction of hierarchical methods for matrix-vector product acceleration’. In: *Sparse Days 2025*. Toulouse, France, 2nd June 2025. URL: <https://inria.hal.science/hal-05184734>.
- [35] A. Gicquel, O. Coulaud, B. Bramas and E. Agullo. ‘A composable abstraction of hierarchical methods for matrix-vector product acceleration’. In: *Workshop on Approximate Computing in Numerical Linear Algebra 2025*. Paris, France, 4th Dec. 2025. URL: <https://inria.hal.science/hal-05397428>.
- [36] A. Malhéné, E. Agullo, L. Giraud and M. Rozloznic. ‘Practical stopping criteria for GMRES’. In: *Sparse Days meeting 2025*. Toulouse, France, 2nd June 2025. URL: <https://inria.hal.science/hal-05200599>.
- [37] S. Nakov, E. Agullo, A. Buttari, L. Giraud and G. Marait. ‘Scalable Domain Decomposition Methods for Large Sparse Linear Systems on Modern Architectures’. In: *Sparse Days 2025*. Toulouse, France, 2nd June 2025. URL: <https://inria.hal.science/hal-05248757>.
- [38] S. Nakov, E. Agullo, A. Buttari, L. Giraud and G. Marait. ‘Scalable Domain Decomposition Methods for Large Sparse Linear Systems on Modern Architectures’. In: *COMPAS 2025 - Conférence francophone d’informatique en Parallélisme, Architecture et Système*. Bordeaux, France, 24th June 2025. URL: <https://inria.hal.science/hal-05248772>.
- [39] C. Peaucelle, E. Agullo and G. Sylvand. ‘Using Floating Point Compression to increase H-Matrices performance’. In: *COMPAS 2025 - Conférence francophone d’informatique en Parallélisme, Architecture et Système*. Bordeaux, France, 24th June 2025. URL: <https://hal.science/hal-05186849>.
- [40] A. E. M. Zekri, O. Coulaud, L. Giraud and J.-R. Poirier. ‘Tensor Train-Based Approach for the Magnetic Moment Method’. In: *Compumag 2025 - 25th International Conference on the Computation of Electromagnetic Fields*. Naples, Italy, 22nd June 2025. URL: <https://hal.science/hal-05291496>.

### Reports & preprints

- [41] M. A. Abouabdallah, O. Coulaud, N. Peyrard and A. Franc. *Computing WSBM conditional marginals with Tensor-Train decomposition*. 30th May 2025. URL: <https://hal.inrae.fr/hal-04394024>.
- [42] C. Acary-Robert, E. Agullo, B. Arrondeau, L. Bilke, D. Bissuel, L. Courtès, C. J. Doering, R. Garbage, K. Hinsén, A. Isaac, E. Medernach, S. Camarasu-Pop, P. Prins, C. Santos, P. Swartvagher, S. Tournier and R. Wurmus. *Guix-HPC Activity Report 2023-2024*. Inria Bordeaux - Sud Ouest, 11th Feb. 2025, pp. 1–36. URL: <https://hal.science/hal-04942752> (cit. on p. 20).
- [43] E. Agullo, A. Buttari, A. Guermouche and A. Jégo. *Redundant computations in task-based parallelism with applications to communication-reducing algorithms*. 22nd July 2025. URL: <https://hal.science/hal-05176537>.
- [44] O. Coulaud, L. Giraud, M. Iannacito and M. Issa. *Solving eigenvalue problems in high dimensions using contour integration and Tensor Train format*. RR-9586. Inria, Apr. 2025, p. 19. URL: <https://inria.hal.science/hal-05017425> (cit. on p. 15).

- [45] E. M. Ettaouchi, L. Giraud, C. Kruse and N. Tardieu. *Convergence analysis of overlapping domain decomposition preconditioners for nonlinear problems*. RR-9575. INRIA, 20th Feb. 2025. URL: <https://hal.science/hal-04958735>.
- [46] L. Giraud, C. Kruse, P. Mycek, M. Shpakovych and Y. Xiang. *Neural network preconditioning: a case study for the solution of the parametric Helmholtz equation*. RR-9593. Inria Centre at the University of Bordeaux, France, 10th July 2025. URL: <https://hal.science/hal-05157038> (cit. on p. 17).
- [47] L. Giraud and Y. Xiang. *A note on the partial convergence management for the solution of symmetric linear systems with multiple right-hand sides*. RR-9574. Inria Centre at the University of Bordeaux, 30th Jan. 2025. URL: <https://inria.hal.science/hal-04922247> (cit. on p. 16).
- [48] C. Guillet. *Error Estimates for Sparse Tensor Products of B-spline Approximation Spaces*. RR-9598. Centre Inria de l'Université de Bordeaux, 22nd Oct. 2025. URL: <https://hal.science/hal-05326739> (cit. on p. 14).
- [49] C. Kruse, A. Clara Ordonez Egas and N. Tardieu. *A Scalable and Parameter-Robust Preconditioner for a Second Gradient of Dilation Regularization Applied to a Mechanics Problem*. Cerfacs; EDF R&D, 13th Aug. 2025. URL: <https://hal.science/hal-05209343> (cit. on p. 16).
- [50] J. Litz, P. Leleux, C. Kruse, J. Gedicke and M. J. Kühn. *Memory-and compute-optimized geometric multigrid GMGPolar for curvilinear coordinate representations -Applications to fusion plasma*. Cerfacs - Toulouse, France, 9th July 2025. URL: <https://hal.science/hal-05153605> (cit. on p. 17).
- [51] A. Tremodeux, E. Agullo, A. Benoit, L. Giraud, T. Herault and Y. Robert. *Fault-tolerant numerical iterative algorithms at scale*. RR-9567. Inria Lyon, Jan. 2025. URL: <https://inria.hal.science/hal-04872041>.
- [52] Y. Xiang. *Randomized LOBPCG algorithm with linear dimension reduction*. Inria Centre at the University of Bordeaux, France, 10th Feb. 2025, p. 33. URL: <https://inria.hal.science/hal-04937938>.
- [53] Y. Xiang. *Unsupervised convolution neural operator preconditioning for the solution of some heterogeneous fluid PDEs*. Inria Centre at the University of Bordeaux, 14th Jan. 2025. URL: <https://inria.hal.science/hal-04886933>.

#### Other scientific publications

- [54] P. Dominikowski, A. Torri, B. Pointal, O. Kaya, L. Lima Pilla and O. Coulaud. 'Exploring Near-Optimal Contraction Strategies for the Scalar Product in the Tensor-Train Format'. In: IPDPS 2025 - 39th IEEE International Parallel & Distributed Processing Symposium. Milan, Italy, 3rd June 2025, pp. 1274–1276. DOI: [10.1109/IPDPSW66978.2025.00210](https://doi.org/10.1109/IPDPSW66978.2025.00210). URL: <https://hal.science/hal-05304665>.
- [55] A. Gicquel, O. Coulaud and B. Bramas. 'Towards a composable abstraction of hierarchical methods for matrix-vector product acceleration'. In: COMPAS 2025 - Conférence francophone d'informatique en Parallélisme, Architecture et Système. Bordeaux, France, 24th June 2025. URL: <https://inria.hal.science/hal-05184758>.
- [56] E. Korkmaz. *ResourceMonitor: Coarse-grain resource monitoring tool for HPC applications*. 17th Nov. 2025. URL: <https://hal.science/hal-05368647>.

#### Software

- [57] [SW] O. Coulaud, A. Franc, J.-M. Frigerio, L. Leandro, R. Peressoni and F. Pruvost, *cpdiodon* version 0.4, 27th Feb. 2025. LIC: CeCILL-C Free Software License Agreement. HAL: [hal-04971369](https://hal.science/hal-04971369), URL: <https://hal.science/hal-04971369>, vcs: <https://gitlab.inria.fr/diodon/cppdiodon>, SWHID: [swh:1:dir:db17fbf0bb5faf4aed25a7edb9db89f9ba733ab1;origin=https://hal.archives-ouvertes.fr/hal-04971369;visit=swh:1:snp:39a10b280d848a8fcbd107216cbb3e3a1a7374aef;anchor=swh:1:rel:c73f2a7693dde1920e3ebb629413f85c84cca1c7;path=/](https://sw.hal.archives-ouvertes.fr/hal-04971369).

- [58] [SW] A. Franc, F. Pruvost, R. Peressoni, J.-M. Frigerio and O. Coulaud, *pydiodon* version 0.1.0, 4th Apr. 2025. LIC: GNU General Public License v3.0 or later. HAL: [hal-05021182](https://hal.inria.fr/hal-05021182), URL: <https://inria.hal.science/hal-05021182>, vcs: <https://gitlab.inria.fr/diodon/pydiodon>, SWHID: [swh:1:dir:5acef1b48334f9a614da1023e41b9ee0ceb5a214;origin=https://gitlab.inria.fr/diodon/pydiodon;visit=swh:1:snp:0ff620e0c9a270ff0ccb4afd0fe8cd aa70a0312a;anchor=swh:1:rev:d1be4783f88682bb77840a93b059bcdca81f0fa5](https://swh.cs.cmu.edu/swh-repository/urn:uuid:1:dir:5acef1b48334f9a614da1023e41b9ee0ceb5a214;origin=https://gitlab.inria.fr/diodon/pydiodon;visit=swh:1:snp:0ff620e0c9a270ff0ccb4afd0fe8cd aa70a0312a;anchor=swh:1:rev:d1be4783f88682bb77840a93b059bcdca81f0fa5).

### 10.3 Cited publications

- [59] E. Agullo, O. Aumage, B. Bramas, O. Coulaud and S. Pitoiset. ‘Bridging the gap between openMP and task-based runtime systems for the fast multipole method’. In: *IEEE Transactions on Parallel and Distributed Systems* 28.10 (2017). DOI: [10.1109/TPDS.2017.2697857](https://doi.org/10.1109/TPDS.2017.2697857) (cit. on p. 10).
- [60] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner and T. Takahashi. ‘Task-based FMM for heterogeneous architectures’. In: *Concurrency and Computation: Practice and Experience* 28.9 (June 2016), pp. 2608–2629. DOI: [10.1002/cpe.3723](https://doi.org/10.1002/cpe.3723). URL: <http://doi.wiley.com/10.1002/cpe.3723> (cit. on p. 10).
- [61] E. Agullo, B. Bramas, O. Coulaud, E. Darve, M. Messner and T. Takahashi. ‘Task-Based FMM for Multicore Architectures’. In: *SIAM Journal on Scientific Computing* 36.1 (2014), pp. 66–93. DOI: [10.1137/130915662](https://hal.inria.fr/hal-00911856). URL: <https://hal.inria.fr/hal-00911856> (cit. on p. 10).
- [62] E. Agullo, S. Cools, E. Fatih-Yetkin, L. Giraud, N. Schenkels and W. Vanroose. ‘On soft errors in the conjugate gradient method: sensitivity and robust numerical detection’. In: *SIAM Journal on Scientific Computing* 42.6 (Nov. 2020). DOI: [10.1137/18M122858X](https://doi.org/10.1137/18M122858X). URL: <https://hal.inria.fr/hal-03022845> (cit. on p. 10).
- [63] E. Agullo, E. Darve, L. Giraud and Y. Harness. ‘Low-Rank Factorizations in Data Sparse Hierarchical Algorithms for Preconditioning Symmetric Positive Definite Matrices’. In: *SIAM Journal on Matrix Analysis and Applications* 39.4 (Oct. 2018), pp. 1701–1725. URL: <https://hal.inria.fr/hal-01940053> (cit. on p. 10).
- [64] E. Agullo, M. Felšöci and G. Sylvand. *A comparison of selected solvers for coupled FEM/BEM linear systems arising from discretization of aeroacoustic problems: literate and reproducible environment*. Technical Report RT-0513. Inria Bordeaux Sud-Ouest, June 2021, p. 100. URL: <https://hal.inria.fr/hal-03263620> (cit. on p. 11).
- [65] E. Agullo, L. Giraud and Y.-F. Jing. ‘Block GMRES method with inexact breakdowns and deflated restarting’. In: *SIAM Journal on Matrix Analysis and Applications* 35.4 (2014), pp. 1625–1651 (cit. on p. 9).
- [66] E. Agullo, L. Giraud and L. Poirel. ‘Robust preconditioners via generalized eigenproblems for hybrid sparse linear solvers’. In: *SIAM Journal on Matrix Analysis and Applications* 40.2 (2019), pp. 417–439. DOI: [10.1137/17M1153765](https://doi.org/10.1137/17M1153765). URL: <https://hal.inria.fr/hal-02074474> (cit. on p. 10).
- [67] P. Blanchard. ‘Fast hierarchical algorithms for the low-rank approximation of matrices, with applications to materials physics, geostatistics and data analysis’. PhD thesis. Bordeaux, 2017. URL: <https://tel.archives-ouvertes.fr/tel-01534930> (cit. on p. 9).
- [68] P. Blanchard, O. Coulaud and E. Darve. *Fast hierarchical algorithms for generating Gaussian random fields*. Research Report 8811. Inria Bordeaux Sud-Ouest, Dec. 2015. URL: <https://hal.inria.fr/hal-01228519> (cit. on pp. 9, 10).
- [69] S. Börm, L. Grasedyck and W. Hackbusch. *Hierarchical Matrices*. Tech. rep. 2003, pp. 1–173 (cit. on p. 6).
- [70] A. Buttari, J. Langou, J. Kurzak and J. Dongarra. ‘Parallel tiled QR factorization for multicore architectures’. In: *Concurrency and Computation: Practice and Experience* 20.13 (2008), pp. 1573–1590 (cit. on p. 10).
- [71] E. Carson, N. J. Higham and S. Pranesh. ‘Three-Precision GMRES-Based Iterative Refinement for Least Squares Problems’. In: *SIAM Journal on Scientific Computing* 42.6 (Jan. 2020), A4063–A4083. DOI: [10.1137/20m1316822](https://doi.org/10.1137/20m1316822) (cit. on p. 10).

- [72] F. Casenave, A. Ern and G. Sylvand. ‘Coupled BEM-FEM for the convected Helmholtz equation with non-uniform flow in a bounded domain’. In: *Journal of Computational Physics* 257.A (Jan. 2014). 23 pages, 9 figures, pp. 627–644. DOI: [10.1016/j.jcp.2013.10.016](https://doi.org/10.1016/j.jcp.2013.10.016). URL: <https://hal.archives-ouvertes.fr/hal-00805477> (cit. on p. 11).
- [73] A. Cichocki, R. Zdunek, A. H. Phan and S.-i. Amari. *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-way Data Analysis and Blind Source Separation*. Wiley, 2009 (cit. on p. 9).
- [74] S. Cools, E. F. Yetkin, E. Agullo, L. Giraud and W. Vanroose. ‘Analyzing the Effect of Local Rounding Error Propagation on the Maximal Attainable Accuracy of the Pipelined Conjugate Gradient Method’. In: *SIAM Journal on Matrix Analysis and Applications* 39.1 (Mar. 2018), pp. 426–450. DOI: [10.1137/17M1117872](https://doi.org/10.1137/17M1117872). URL: <https://hal.inria.fr/hal-01753411> (cit. on p. 10).
- [75] O. Coulaud, A. A. Franc and M. Iannacito. *Extension of Correspondence Analysis to multiway data-sets through High Order SVD: a geometric framework*. Research Report RR-9429. Inria Bordeaux - Sud-Ouest ; Inrae, Nov. 2021. URL: <https://hal.inria.fr/hal-03418404> (cit. on p. 10).
- [76] A. Falco. ‘Comblent l’écart entre  $\mathcal{H}$ -Matrices et méthodes directes creuses pour la résolution de systèmes linéaires de grandes tailles’. Theses. Université de Bordeaux, June 2019. URL: <https://tel.archives-ouvertes.fr/tel-02183902> (cit. on p. 10).
- [77] A. A. Franc, P. Blanchard and O. Coulaud. ‘Nonlinear mapping and distance geometry’. In: *Optimization Letters* 14.2 (2020), pp. 453–467. DOI: [10.1007/s11590-019-01431-y](https://doi.org/10.1007/s11590-019-01431-y). URL: <https://hal.inria.fr/hal-02124882> (cit. on p. 10).
- [78] N. Gillis. *Nonnegative Matrix Factorization*. Society for Industrial and Applied Mathematics, Jan. 2020. DOI: [10.1137/1.9781611976410](https://doi.org/10.1137/1.9781611976410) (cit. on p. 9).
- [79] L. Giraud, Y.-F. Jing and Y. Xiang. *A block minimum residual norm subspace solver for sequences of multiple left and right-hand side linear systems*. Research Report RR-9393. Inria Bordeaux Sud-Ouest, Feb. 2021, p. 60. URL: <https://hal.inria.fr/hal-03146213> (cit. on p. 9).
- [80] L. Grasedyck, W. Hackbusch and B. Nr. ‘An Introduction to Hierarchical ( H - ) Rank and TT - Rank of Tensors with Examples’. In: *Computational Methods in Applied Mathematics* 11.329 (2011), pp. 291–304 (cit. on p. 6).
- [81] B. Gunter and R. Van De Geijn. ‘Parallel out-of-core computation and updating of the QR factorization’. In: *ACM Transactions on Mathematical Software (TOMS)* 31.1 (2005), pp. 60–78 (cit. on p. 10).
- [82] W. Hackbusch. *Hierarchical Matrices: Algorithms and Analysis*. 1st. Springer Publishing Company, Incorporated, 2015 (cit. on p. 6).
- [83] N. Halko, P.-G. G. Martinsson and J. A. Tropp. ‘Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions’. In: *SIAM Review* 53.2 (2011), pp. 217–288. DOI: [10.1137/090771806](https://doi.org/10.1137/090771806). URL: <http://arxiv.org/abs/0909.4061> (cit. on p. 10).
- [84] T. G. Kolda and B. W. Bader. ‘Tensor Decompositions and Applications’. In: *SIAM Review* 51.3 (Aug. 2009), pp. 455–500. DOI: [10.1137/07070111X](https://doi.org/10.1137/07070111X). URL: <http://epubs.siam.org/doi/abs/10.1137/07070111X> (cit. on p. 6).
- [85] C. Kruse, V. Darrigrand, N. Tardieu, M. Arioli and U. Rüde. ‘Application of an iterative Golub-Kahan algorithm to structural mechanics problems with multi-point constraints’. In: *Adv. Model. Simul. Eng. Sci.* 7.1 (2020), p. 45. DOI: [10.1186/s40323-020-00181-2](https://doi.org/10.1186/s40323-020-00181-2). URL: <https://doi.org/10.1186/s40323-020-00181-2> (cit. on p. 9).
- [86] C. Kruse, M. Sosonkina, M. Arioli, N. Tardieu and U. Rüde. ‘Parallel solution of saddle point systems with nested iterative solvers based on the Golub-Kahan Bidiagonalization’. In: *Concurr. Comput. Pract. Exp.* 33.11 (2021). DOI: [10.1002/cpe.5914](https://doi.org/10.1002/cpe.5914). URL: <https://doi.org/10.1002/cpe.5914> (cit. on p. 9).
- [87] V. Le Bris, M. Odunlami, D. Bégué, I. Baraille and O. Coulaud. ‘Using computed infrared intensities for the reduction of vibrational configuration interaction bases’. In: *Phys. Chem. Chem. Phys.* 22 (13 2020), pp. 7021–7030. DOI: [10.1039/D0CP00593B](https://doi.org/10.1039/D0CP00593B). URL: <http://dx.doi.org/10.1039/D0CP00593B> (cit. on p. 9).

- [88] B. Lizé. ‘Résolution directe rapide pour les éléments finis de frontière en électromagnétisme et acoustique :  $\mathcal{H}$ -Matrices. Parallélisme et applications industrielles’. Theses. Université Paris-Nord - Paris XIII, June 2014. URL: <https://tel.archives-ouvertes.fr/tel-01244260> (cit. on p. 10).
- [89] P.-G. Martinsson and J. Tropp. ‘Randomized Numerical Linear Algebra: Foundations & Algorithms’. In: (2020). eprint: 2002.01387. URL: <http://arxiv.org/abs/2002.01387> (cit. on p. 10).
- [90] M. Odunlami, V. Le Bris, D. Bégué, I. Baraille and O. Coulaud. ‘A-VCI: A flexible method to efficiently compute vibrational spectra’. In: *The Journal of Chemical Physics* 146.21 (June 2017), p. 214108. DOI: 10.1063/1.4984266. URL: <http://aip.scitation.org/doi/10.1063/1.4984266> (cit. on p. 9).
- [91] I. V. Oseledets. ‘Tensor-Train Decomposition’. In: *SIAM Journal on Scientific Computing* 33.5 (Jan. 2011), pp. 2295–2317. DOI: 10.1137/090752286. URL: <https://doi.org/10.1137/090752286> (cit. on p. 6).
- [92] L. Poirel. ‘Algebraic domain decomposition methods for hybrid (iterative/direct) solvers’. Theses. Université de Bordeaux, Nov. 2018. URL: <https://tel.archives-ouvertes.fr/tel-02070618> (cit. on p. 10).
- [93] J.-R. Poirier, O. Coulaud and O. Kaya. ‘Fast BEM Solution for 2-D Scattering Problems Using Quantized Tensor-Train Format’. In: *IEEE Transactions on Magnetics* 56.3 (Mar. 2020), pp. 1–4. DOI: 10.1109/TMAG.2019.2954584. URL: <https://hal.inria.fr/hal-03150956> (cit. on p. 10).
- [94] G. Sylvand. ‘La méthode multipôle rapide en électromagnétisme. Performances, parallélisation, applications’. Theses. Ecole des Ponts ParisTech, June 2002. URL: <https://pastel.archives-ouvertes.fr/tel-00005683> (cit. on p. 10).
- [95] N. Venkovic, P. Mycek, L. Giraud and O. Le Maitre. *Recycling Krylov subspace strategies for sequences of sampled stochastic elliptic equations*. Research Report RR-9425. Inria Bordeaux - Sud Ouest, Oct. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03366966> (cit. on p. 10).