

RESEARCH CENTRE

**Inria Center  
at the University of Bordeaux**

IN PARTNERSHIP WITH:

**Centre Européen de Recherche et de  
Formation Avancée en Calcul  
Scientifique, Airbus Research &  
Technology**

2022

ACTIVITY REPORT

Project-Team

CONCACE

**Numerical and Parallel Composability for  
High Performance Computing**

**DOMAIN**

**Networks, Systems and Services,  
Distributed Computing**

**THEME**

**Distributed and High Performance  
Computing**

*Inria*

# Contents

<b>Project-Team CONCACE</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>3</b>
<b>3 Research program</b>	<b>5</b>
3.1 Numerical algebra methodologies in model and data-driven scientific computing . . . . .	5
3.1.1 Scientific computing in large size linear algebra . . . . .	6
3.1.2 Scientific computing in large dimension multi-linear algebra . . . . .	6
3.1.3 Scientific continuum between large size and large dimension . . . . .	6
3.2 Composition of parallel numerical algorithms from a sequential expression . . . . .	7
<b>4 Application domains</b>	<b>8</b>
4.1 Material physics . . . . .	8
4.2 Co-design of algorithms in scientific applications . . . . .	9
4.2.1 Numerical and parallel scalable hybrid solvers in large scale calculations . . . . .	9
4.2.2 Aeroacoustics Simulation . . . . .	9
<b>5 Highlights of the year</b>	<b>9</b>
<b>6 New software and platforms</b>	<b>10</b>
6.1 New software . . . . .	10
6.1.1 compose . . . . .	10
6.1.2 ScalFMM . . . . .	10
6.1.3 CPPDiodon . . . . .	11
6.1.4 FMR . . . . .	11
<b>7 New results</b>	<b>12</b>
7.1 A block minimum residual norm subspace solver with partial convergence management for sequences of linear systems . . . . .	12
7.2 Direct solution of larger coupled sparse/dense linear systems using low-rank compression on single-node multi-core machines in an industrial context . . . . .	12
7.3 Study of the processor and memory power consumption of coupled sparse/dense solvers .	13
7.4 Task-based randomized singular value decomposition and multidimensional scaling . . . .	13
7.5 Guix-HPC Activity Report 2020-2021 . . . . .	13
7.6 Decentralized in-order execution of a sequential task-based code for shared-memory archi- tectures . . . . .	14
7.7 Combining reduction with synchronization barrier on multi-core processors . . . . .	14
7.8 The backward stable variants of GMRES in variable accuracy . . . . .	14
7.9 A robust GMRES algorithm in Tensor Train format . . . . .	15
7.10 On some orthogonalization schemes in Tensor Train format . . . . .	15
7.11 High-order multigrid strategies for HHO discretizations of elliptic equations . . . . .	15
<b>8 Bilateral contracts and grants with industry</b>	<b>15</b>
8.1 Bilateral Grants with Industry . . . . .	16
<b>9 Partnerships and cooperations</b>	<b>16</b>
9.1 European initiatives . . . . .	16
9.1.1 H2020 projects . . . . .	16
9.1.2 Other european programs/initiatives . . . . .	18
9.2 National initiatives . . . . .	19

<b>10 Dissemination</b>	<b>20</b>
10.1 Promoting scientific activities	20
10.1.1 Scientific events: organisation	20
10.1.2 Scientific events: selection	20
10.1.3 Journal	20
10.1.4 Scientific expertise	20
10.1.5 Research administration	21
10.2 Teaching - Supervision - Juries	21
10.2.1 Teaching	21
10.2.2 Supervision	21
10.2.3 Juries	22
10.3 Popularization	22
10.3.1 Interventions	22
<b>11 Scientific production</b>	<b>22</b>
11.1 Major publications	22
11.2 Publications of the year	23
11.3 Cited publications	24

## Project-Team CONCACE

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

### 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. – High performance computing
- A7.1. – Algorithms
- A8.1. – Discrete mathematics, combinatorics
- A8.2. – Optimization
- 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
- B3.6. – Ecology
- B3.6.1. – Biodiversity
- 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

## **1 Team members, visitors, external collaborators**

### **Research Scientists**

- Luc Giraud [Team leader, INRIA, Senior Researcher, HDR]
- Guillaume Sylvand [Team leader, Airbus Central R & T, Senior Researcher]
- Carola Kruse [Team leader, Cerfacs, Senior Researcher]
- Emmanuel Agullo [INRIA, Researcher]
- Pierre Benjamin [Airbus Central R & T, Senior Researcher]
- Olivier Coulaud [INRIA, Senior Researcher, HDR]
- Sofiane Haddad [Airbus Central R & T, Senior Researcher]
- Paul Mycek [Cerfacs, Senior Researcher]

### **Post-Doctoral Fellows**

- Marvin Lasserre [INRIA, from Oct 2022]
- Maksym Shpakovych [INRIA, from Oct 2022]

### **PhD Students**

- Marek Felsoci [INRIA]
- Martina Iannacito [INRIA]
- Romain Peressoni [UNIV BORDEAUX]
- Yanfei Xiang [INRIA]

### **Technical Staff**

- Pierre Estérie [INRIA, Engineer]

### **Interns and Apprentices**

- Yuxuan Wang [INRIA, from Mar 2022 until Sep 2022]

### **Administrative Assistant**

- Flavie Blondel [INRIA, from Oct 2022]

### **External Collaborators**

- Jean-Rene Poirier [TOULOUSE INP, HDR]
- Ulrich Råde [Friedrich-Alexander-Universität Erlangen & Cerfacs, 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 [54, 62, 50] low-rank hierarchical matrices [52, 36]). 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 on-going 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 [32, 48], randomized range finder algorithms [34, 35], 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 [47, 41].
- **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) [55, 56]. 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 [61, 57] 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 [43, 46];
- 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) [66] leading to large dimension or preconditioning based on a low-rank approximation of the tensorization of the dense matrix in Boundary Element Method solver [27, 33, 63].

### 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 [45, 64, 58]) or speed up the algorithm (fast multipole

method, randomized algorithm [53, 59] [65, 35] 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 [38, 42].

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 [26] 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 [37] is based on a data format and computational technique originally used in an out-of-core context [51]. 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 [25, 24, 23]. In the concace 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 [29].

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 an advanced computing technology can be synergized with classical scientific computing.

## 4 Application domains

### 4.1 Material physics

**Participants:** Olivier Coulaud, Pierre Esterie.

Due to the increase of available computer power, new applications in nano science and physics appear such as study of properties of new materials (photovoltaic materials, bio- and environmental sensors, ...), failure in materials, nano-indentation. Chemists, physicists now commonly perform simulations in these fields. These computations simulate systems up to billion of atoms in materials, for large time scales up to several nanoseconds. The larger the simulation, the smaller the computational cost of the potential driving the phenomena, resulting in low precision results. So, if we need to increase the precision, there are two ways to decrease the computational cost. In the first approach, we improve algorithms and their parallelization and in the second way, we will consider a multiscale approach.

A domain of interest is the material aging for the nuclear industry. The materials are exposed to complex conditions due to the combination of thermo-mechanical loading, the effects of irradiation and the harsh operating environment. This operating regime makes experimentation extremely difficult and we must rely on multi-physics and multi-scale modeling for our understanding of how these materials behave in service. This fundamental understanding helps not only to ensure the longevity of existing nuclear reactors, but also to guide the development of new materials for 4th generation reactor programs and dedicated fusion reactors. For the study of crystalline materials, an important tool is dislocation dynamics (DD) modeling. This multiscale simulation method predicts the plastic response of a material from the underlying physics of dislocation motion. DD serves as a crucial link between the scale of molecular dynamics and macroscopic methods based on finite elements; it can be used to accurately describe the interactions of a small handful of dislocations, or equally well to investigate the global behavior of a massive collection of interacting defects.

To explore i.e. to simulate these new areas, we need to develop and/or to improve significantly models, schemes and solvers used in the classical codes. In the project, we want to accelerate algorithms arising in those fields.

We will focus on the following topics

- The interaction between dislocations is long ranged ( $O(1/r)$ ) and anisotropic, leading to severe computational challenges for large-scale simulations. In dislocation codes, the computation of interaction forces between dislocations is still the most CPU time consuming and has to be improved to obtain faster and more accurate simulations.
- In such simulations, the number of dislocations grows while the phenomenon occurs and these dislocations are not uniformly distributed in the domain. This means that strategies to dynamically construct a good load balancing are crucial to achieve high performance.
- From a physical and a simulation point of view, it will be interesting to couple a molecular dynamics model (atomistic model) with a dislocation one (mesoscale model). In such three-dimensional coupling, the main difficulties are firstly to find and characterize a dislocation in the atomistic

region, secondly to understand how we can transmit with consistency the information between the two micro and meso scales.

## 4.2 Co-design of algorithms in scientific applications

**Participants:** Emmanuel Agullo, Carola Kruse, Paul Mycek, Pierre Benjamin, Marek Felsoci, Luc Giraud, Gilles Marait, Guillaume Sylvand.

### 4.2.1 Numerical and parallel scalable hybrid solvers in large scale calculations

Parallel and numerically scalable hybrid solvers based on a fully algebraic coarse space correction have been theoretically studied and various advanced parallel implementations have been designed. Their parallel scalability has been initially investigated on large scale problems within the **EoCoE** project thanks to a close collaboration with the BSC and the integration of **MaPhyS** within the Alya software. This activity will further develop in the **EoCoE2** project. The performance has also been assessed on PRACE Tier-0 machine within a PRACE Project Access through a collaboration with CERFACS and Laboratoire de Physique des Plasmas at Ecole Polytechnique for the calculation of plasma propulsion. A comparative parallel scalability study with the Algebraic MultiGrid from Petsc has been conducted in that framework.

### 4.2.2 Aeroacoustics Simulation

This domain 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 play a central part in the upstream design phase of any such project [39]. 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 from 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

The team creation is the main highlight as **CONCACE** is the first Inria joint team with industry that involves two industrial partners, that gathers a diversity of scientific profiles and professional experiences. This feature is a real asset for future research and innovation.

## 6 New software and platforms

### 6.1 New software

#### 6.1.1 compose

**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/compose/compose>

**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 Conccace.

**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 Estérie

### 6.1.3 CPPDiodon

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

**Keywords:** SVD, PCA

**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).

**News of the Year:** Research report about task-based MDS: <https://hal.inria.fr/hal-03773985>. Speedup by a factor of about 3 of the HDF5 files reading thanks to the creation of several processes for this task (not handled internally by HDF5). Improved performances thanks to the new GEMM/SYMM A-stationnary of Chameleon. Improved performances in parallel MPI thanks to new data distributions SBC and TBC in the symmetric matrix case. New possibility for MDS solving using an eigen value decomposition (EVD) or by a randomized approach (rEVD).

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

**Publication:** [hal-03773985](https://hal.inria.fr/hal-03773985)

**Authors:** Olivier Coulaud, Florent Pruvost

**Contact:** Olivier Coulaud

**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/ch>)

**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>).

**News of the Year:** Research report about distributed task-based MDS algorithm using FMR: <https://hal.inria.fr/hal-03773985>. Speedup by a factor of about 3 of the HDF5 files reading thanks to the creation of several processes for this task (not handled internally by HDF5). Improved performances thanks to the new GEMM/SYMM A-stationnary of Chameleon. Improved performances in parallel MPI thanks to new data distributions SBC and TBC in the symmetric matrix case. We have implemented the computation of the eigenvalues by a random projection approach (REVD).

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

**Publications:** [hal-01228519](https://hal.inria.fr/hal-01228519), [hal-01334747](https://hal.inria.fr/hal-01334747), [hal-03773985](https://hal.inria.fr/hal-03773985)

**Contact:** Olivier Coulaud

**Participants:** Olivier Coulaud, Florent Pruvost, Romain Peressoni

## 7 New results

**Participants:** All team members.

### 7.1 A block minimum residual norm subspace solver with partial convergence management for sequences of linear systems

We are concerned with the iterative solution of linear systems with multiple right-hand sides available one group after another with possibly slowly-varying left-hand sides. For such sequences of linear systems, we first develop a new block minimum norm residual approach that combines two main ingredients. The first component exploits ideas from GCRO-DR [SIAM J. Sci. Comput., 28(5) (2006), pp. 1651–1674], enabling to recycle information from one solve to the next. The second component is the numerical mechanism to manage the partial convergence of the right-hand sides, referred to as inexact breakdown detection in IB-BGMRES [Linear Algebra Appl., 419 (2006), pp. 265–285], that enables the monitoring of the rank deficiency in the residual space basis expanded block-wise. Secondly, for the class of block minimum norm residual approaches, that relies on a block Arnoldi-like equality between the search space and the residual space (e.g., any block GMRES or block GCRO variants), we introduce new search space expansion policies defined on novel criteria to detect the partial convergence. These novel detection criteria are tuned to the selected stopping criterion and targeted convergence threshold to best cope with the selected normwise backward error stopping criterion, enabling to monitor the computational effort while ensuring the final accuracy of each individual solution. Numerical experiments are reported to illustrate the numerical and computational features of both the new block Krylov solvers and the new search space block expansion policies.

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

### 7.2 Direct solution of larger coupled sparse/dense linear systems using low-rank compression on single-node multi-core machines in an industrial context

While hierarchically low-rank compression methods are now commonly available in both dense and sparse direct solvers, their usage for the direct solution of coupled sparse/dense linear systems has been little investigated. The solution of such systems is though central for the simulation of many important physics problems such as the simulation of the propagation of acoustic waves around aircrafts. Indeed, the heterogeneity of the jet flow created by reactors often requires a Finite Element Method (FEM)

discretization, leading to a sparse linear system, while it may be reasonable to assume as homogeneous the rest of the space and hence model it with a Boundary Element Method (BEM) discretization, leading to a dense system. In an industrial context, these simulations are often operated on modern multicore workstations with fully-featured linear solvers. Exploiting their low-rank compression techniques is thus very appealing for solving larger coupled sparse/dense systems (hence ensuring a finer solution) on a given multicore workstation, and – of course – possibly do it fast. The standard method performing an efficient coupling of sparse and dense direct solvers is to rely on the Schur complement functionality of the sparse direct solver. However, to the best of our knowledge, modern fully- featured sparse direct solvers offering this functionality return the Schur complement as a non compressed matrix. In this paper, we study the opportunity to process larger systems in spite of this constraint. For that we propose two classes of algorithms, namely multi-solve and multi-factorization, consisting in composing existing parallel sparse and dense methods on well chosen submatrices. An experimental study conducted on a 24 cores machine equipped with 128 GiB of RAM shows that these algorithms, implemented on top of state-of-the-art sparse and dense direct solvers, together with proper low-rank assembly schemes, can respectively process systems of 9 million and 2.5 million total unknowns instead of 1.3 million unknowns with a standard coupling of compressed sparse and dense solvers.

For more details on this work we refer to [30, 31].

### **7.3 Study of the processor and memory power consumption of coupled sparse/dense solvers**

In the aeronautical industry, aeroacoustics is used to model the propagation of acoustic waves in air flows enveloping an aircraft in flight. This for instance allows one to simulate the noise produced at ground level by an aircraft during the takeoff and landing phases, in order to validate that the regulatory environmental standards are met. Unlike most other complex physics simulations, the method resorts to solving coupled sparse/dense systems. In a previous study, we proposed two classes of algorithms for solving such large systems on a relatively small workstation (one or a few multicore nodes) based on compression techniques. The objective of this study is to assess whether the positive impact of the proposed algorithms on time to solution and memory usage translates to the energy consumption as well. Because of the nature of the problem, coupling dense and sparse matrices, and the underlying solutions methods, including dense, sparse direct and compression steps, this yields an interesting processor and memory power profile which we aim to analyze in details.

For more details on this work we refer to [28].

### **7.4 Task-based randomized singular value decomposition and multidimensional scaling**

The multidimensional scaling (MDS) is an important and robust algorithm for representing individual cases of a dataset out of their respective dissimilarities. However, heuristics, possibly trading-off with robustness, are often preferred in practice due to the potentially prohibitive memory and computational costs of the MDS. The recent introduction of random projection techniques within the MDS allowed it to become competitive on larger test cases. The goal of this manuscript is to propose a high-performance distributed-memory MDS based on random projection for processing data sets of even larger size (up to one million items). We propose a task-based design of the whole algorithm and we implement it within an efficient software stack including state-of-the-art numerical solvers, runtime systems and communication layers. The outcome is the ability to efficiently apply robust MDS to large data sets on modern supercomputers. We assess the resulting algorithm and software stack to the point cloud visualization for analyzing distances between sequences in metabarcoding.

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

### **7.5 Guix-HPC Activity Report 2020-2021**

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 tool and aims to make it a better tool for HPC practitioners and scientists concerned with reproducible



research. This report highlights key achievements of Guix-HPC between our previous report a year ago and today, February 2022. This report highlights developments on GNU Guix proper, but also downstream on Guix-Jupyter, the Guix Workflow Language, upstream with Software Heritage integration, as well as experience reports on end-to-end reproducible research article authoring pipelines.

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

## 7.6 Decentralized in-order execution of a sequential task-based code for shared-memory architectures

The hardware complexity of modern machines makes the design of adequate programming models crucial for jointly ensuring performance, portability, and productivity in high-performance computing (HPC). Sequential task-based programming models paired with advanced runtime systems allow the programmer to write a sequential algorithm independently of the hardware architecture in a productive and portable manner, and let a third party software layer—the runtime system—deal with the burden of scheduling a correct, parallel execution of that algorithm to ensure performance. Many HPC algorithms have successfully been implemented following this paradigm, as a testimony of its effectiveness. Developing algorithms that specifically require fine-grained tasks along this model is still considered prohibitive, however, due to per-task management overhead, forcing the programmer to resort to a less abstract, and hence more complex “task+X” model. We thus investigate the possibility to offer a tailored execution model, trading dynamic mapping for efficiency by using a decentralized, conservative in-order execution of the task flow, while preserving the benefits of relying on the sequential task-based programming model. We propose a formal specification of the execution model as well as a prototype implementation, which we assess on a shared-memory multicore architecture with several synthetic workloads. The results show that under the condition of a proper task mapping supplied by the programmer, the pressure on the runtime system is significantly reduced and the execution of fine-grained task flows is much more efficient.

For more details on this work we refer to [40].

## 7.7 Combining reduction with synchronization barrier on multi-core processors

With the rise of multi-core processors with a large number of cores the need of shared memory reduction that perform efficiently on a large number of core is more pressing. Efficient shared memory reduction on these multi-core processors will help share memory programs being more efficient on these one. In this paper, we propose a reduction combined with barrier method that uses SIMD instructions to combine barriers signaling and reduction value read/write to minimize memory/cache traffic between cores thus, reducing barrier latency. We compare different barriers and reduction methods on three multi-core processors and show that proposed combining barrier/reduction method are 4 and 3.5 times faster than respectively GCC 11.1 and Intel 21.2 OpenMP 4.5 reduction.

For more details on this work we refer to [60].

## 7.8 The backward stable variants of GMRES in variable accuracy

In the context where the representation of the data is decoupled from the arithmetic used to process them, we investigate the backward stability of two backward-stable implementations of the GMRES method, namely the so-called Modified Gram-Schmidt (MGS) and the Householder variants. Considering data may be compressed to alleviate the memory footprint, we are interested in the situation where the leading part of the rounding error is related to the data representation. When the data representation of vectors introduces componentwise perturbations, we show that the existing backward stability analyses of MGS-GMRES and Householder-GMRES still apply. We illustrate this backward stability property in a practical context where an agnostic lossy compressor is employed and enables the reduction of the memory requirement to store the orthonormal Arnoldi basis or the Householder reflectors. Although technical arguments of the theoretical backward stability proofs do not readily apply to the situation where only the normwise relative perturbations of the vector storage can be controlled, we show experimentally that the backward stability is maintained; that is, the attainable normwise backward error is of the same order as the normwise perturbations induced by the data storage. We illustrate it with numerical experiments

in two practical different contexts. The first one corresponds to the use of an agnostic compressor where vector compression is controlled normwise. The second one arises in the solution of tensor linear systems, where low-rank tensor approximations based on Tensor-Train is considered to tackle the curse of dimensionality.

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

## 7.9 A robust GMRES algorithm in Tensor Train format

We consider the solution of linear systems with tensor product structure using a GMRES algorithm. To cope with the computational complexity in large dimension both in terms of floating point operations and memory requirement, our algorithm is based on low-rank tensor representation, namely the Tensor Train format. In a backward error analysis framework, we show how the tensor approximation affects the accuracy of the computed solution. With the backward perspective, we investigate the situations where the  $(d + 1)$ -dimensional problem to be solved results from the concatenation of a sequence of  $d$ -dimensional problems (like parametric linear operator or parametric right-hand side problems), we provide backward error bounds to relate the accuracy of the  $(d + 1)$ -dimensional computed solution with the numerical quality of the sequence of  $d$ -dimensional solutions that can be extracted from it. This enables to prescribe convergence threshold when solving the  $(d + 1)$ -dimensional problem that ensures the numerical quality of the  $d$ -dimensional solutions that will be extracted from the  $(d + 1)$ -dimensional computed solution once the solver has converged. The above mentioned features are illustrated on a set of academic examples of varying dimensions and sizes. For more details on this work we refer to [21].

## 7.10 On some orthogonalization schemes in Tensor Train format

In the framework of tensor spaces, we consider orthogonalization kernels to generate an orthogonal basis of a tensor subspace from a set of linearly independent tensors. In particular, we investigate numerically the loss of orthogonality of six orthogonalization methods, namely Classical and Modified Gram-Schmidt with (CGS2, MGS2) and without (CGS, MGS) re-orthogonalization, the Gram approach, and the Householder transformation. To tackle the curse of dimensionality, we represent tensor with low rank approximation using the Tensor Train (TT) formalism, and we introduce recompression steps in the standard algorithm outline through the TT-rounding method at a prescribed accuracy. After describing the algorithm structure and properties, we illustrate numerically that the theoretical bounds for the loss of orthogonality in the classical matrix computation round-off analysis results are maintained, with the unit round-off replaced by the TT-rounding accuracy. The computational analysis for each orthogonalization kernel in terms of the memory requirement and the computational complexity measured as a function of the number of TT-rounding, which happens to be the computational most expensive operation, completes the study.

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

## 7.11 High-order multigrid strategies for HHO discretizations of elliptic equations

This study compares various multigrid strategies for the fast solution of elliptic equations discretized by the Hybrid High-Order method. Combinations of h-, p- and hp-coarsening strategies are considered, combined with diverse intergrid transfer operators. Comparisons are made experimentally on 2D and 3D test cases, with structured and unstructured meshes, and with nested and non-nested hierarchies. Advantages and drawbacks of each strategy are discussed for each case to establish simplified guidelines for the optimization of the time to solution.

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

# 8 Bilateral contracts and grants with industry

**Participants:** Emmanuel Agullo, Luc Giraud, Guillaume Sylvand.

## 8.1 Bilateral Grants with Industry

Some on the ongoing PhD thesis are developed within bilareal contract with industry for PhD advisory such as

- Airbus CR&T for the PhD thesis of Marek Felsoci.
- IFPEN for the PhD of Aboul-Karim Mohamed El Maarouf,

In addition two post-docs, namely Maksym Shpakovych and Marvin Lasserre, are funded by the "plan de relance"

## 9 Partnerships and cooperations

**Participants:** Emmanuel Agullo, Olivier Coulaud, Luc Giraud, Guillaume Sylvand.

### 9.1 European initiatives

#### 9.1.1 H2020 projects

**EoCoE-2**

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

**Duration:** 2018-2022

**Coordinator:** CEA

**Inria coordinator:** Bruno Raffin

**CONCACE contact:** Luc Giraud

**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)
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV (Germany)
- RHEINISCH-WESTFAELISCHE TECHNISCHE HOCHSCHULE AACHEN (Germany)
- THE CYPRUS INSTITUTE (Cyprus)
- UNIVERSITA DEGLI STUDI DI ROMA TORVERGATA (Italy)
- UNIVERSITA DEGLI STUDI DI TRENTO (Italy)
- UNIVERSITE LIBRE DE BRUXELLES (Belgium)

- UNIVERSITE PARIS-SUD (France)
- UNIVERSITY OF BATH (UK)

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

**Summary:** The aim of the present proposal is to establish an Energy Oriented Centre of Excellence for computing applications, (EoCoE). EoCoE (pronounce “Echo”) will use the prodigious potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply. To achieve this goal, we believe that the present revolution in hardware technology calls for a similar paradigm change in the way application codes are designed. EoCoE will assist the energy transition via targeted support to four renewable energy pillars: Meteo, Materials, Water and Fusion, each with a heavy reliance on numerical modelling. These four pillars will be anchored within a strong transversal multidisciplinary basis providing high-end expertise in applied mathematics and HPC. EoCoE is structured around a central Franco-German hub coordinating a pan-European network, gathering a total of 8 countries and 23 teams. Its partners are strongly engaged in both the HPC and energy fields; a prerequisite for the long-term sustainability of EoCoE and also ensuring that it is deeply integrated in the overall European strategy for HPC. The primary goal of EoCoE is to create a new, long lasting and sustainable community around computational energy science. At the same time, EoCoE is committed to deliver high-impact results within the first three years. It will resolve current bottlenecks in application codes, leading to new modelling capabilities and scientific advances among the four user communities; it will develop cutting-edge mathematical and numerical methods, and tools to foster the usage of Exascale computing. Dedicated services for laboratories and industries will be established to leverage this expertise and to foster an ecosystem around HPC for energy. EoCoE will give birth to new collaborations and working methods and will encourage widely spread best practices.

## PRACE 6IP

**Title:** PRACE Sixth Implementation Phase

**Duration:** 2019-2022

**Partners:** see [list](#)

**Inria contact:** Luc Giraud

**Summary:** The mission of PRACE (Partnership for Advanced Computing in Europe) is to enable high-impact scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society. PRACE seeks to realise this mission by offering world class computing and data management resources and services through a peer review process. PRACE also seeks to strengthen the European users of HPC in industry through various initiatives. PRACE has a strong interest in improving energy efficiency of computing systems and reducing their environmental impact. The objectives of PRACE-6IP are to build on and seamlessly continue the successes of PRACE and start new innovative and collaborative activities proposed by the consortium. These include: assisting the development of PRACE 2; strengthening the internationally recognised PRACE brand; continuing and extend advanced training which so far provided more than 36 400 person-training days; preparing strategies and best practices towards Exascale computing, work on forward-looking SW solutions; coordinating and enhancing the operation of the multi-tier HPC systems and services; and supporting users to exploit massively parallel systems and novel architectures. The activities are designed to increase Europe’s research and innovation potential especially through: seamless and efficient Tier-0 services and a pan-European HPC ecosystem including national capabilities; promoting take-up by industry and new communities and special offers to SMEs; assistance to PRACE 2 development; proposing strategies for deployment of leadership systems; collaborating with the ETP4HPC, CoEs and other European and international organisations on future architectures,

## RISC2

**Title:** A network for supporting the coordination of High-Performance Computing research between Europe and Latin America

**Type:** H2020 (Coordinated Support Action)

**URL:** See also: [list](#)

**Duration:** 2021 - 2023

**Coordinator:** Barcelona Supercomputing Center (Spain)

**Inria coordinator:** Stéphane Lanteri

**CONCACE contact:** Luc Giraud

**Partners:** • Forschungszentrum Julich GMBH (Germany)

- Inria (France)
- Bull SAS (France)
- INESC TEC (Portugal)
- Universidade de Coimbra (Portugal)
- CIEMAT (Spain)
- CINECA (Italy)
- Universidad de Buenos Aires (Argentina)
- Universidad Industrial de Santander (Columbia)
- Universidad de le Republica (Uruguay)
- Laboratorio Nacional de Computacao Cientifica (Brazil)
- Centro de Investigacion y de Estudios Avanzados del Instituto Politecnico Nacional (Mexico)
- Universidad de Chile (Chile)
- Fundacao Coordenacao de Projetos Pesquisas e Estudos Tecnologicos COPPETEC (Brazil)
- Fundacion Centro de Alta Tecnologia (Costa Rica)

**Summary** Recent advances in AI and the Internet of things allow high performance computing (HPC) to surpass its limited use in science and defence and extend its benefits to industry, healthcare and the economy. Since all regions intensely invest in HPC, coordination and capacity sharing are needed. The EU-funded RISC2 project connects eight important European HPC actors with the main HPC actors from Argentina, Brazil, Chile, Colombia, Costa Rica, Mexico and Uruguay to enhance cooperation between their research and industrial communities on HPC application and infrastructure development. The project will deliver a cooperation roadmap addressing policy-makers and the scientific and industrial communities to identify central application areas, HPC infrastructure and policy needs.

### 9.1.2 Other european programs/initiatives

#### High Performance Spacecraft Plasma Interaction Software

**Duration:** 2022 - 2024

**Funding:** ESA

**Coordinator:** Sébastien Hess (ONERA)

**CONCACE contact:** Olivier Couland and Luc Giraud

**Summary:** Controlling the plasma environment of satellites is a key issue for the nation in terms of satellite design and propulsion. Three-dimensional numerical modelling is thus a key element, particularly in the preparation of future space missions. The SPIS code is today the reference in Europe for the simulation of these phenomena. The methods used to describe the physics of these plasmas are based on the representation of the plasma by a system of particles moving in a mesh (here unstructured) under the effect of the electric field which satisfies the Poisson equation. ESA has recently shown an interest in applications requiring complex 3D calculations, which may involve several tens of millions of cells and several tens of billions of particles, and therefore in a highly parallel and scalable version of the SPIS code.

## 9.2 National initiatives

### OPERA (Adaptive planar optics - ANR ASTRID Maturation)

**Duration:** 2019 – 2022

**Coordinator:** Stéphane Lanteri (Atlantis - SAM)

**CONCACE contact:** Luc Giraud

**Summary:** In the OPERA project, we are investigating and optimizing the properties of planar photonic devices based on metasurfaces using numerical modelling. The scientific and technical activities that constitute the project work programme are organized around 4 main workpackages. The numerical characterization of the optical properties of planar devices based on metasurfaces, as well as their optimization are at the heart of the activities and objectives of two horizontal (transversal) workpackages. These numerical methodologies will be integrated into the DIOGENeS software framework that will eventually integrate (1) discontinuous Galerkin-type methods that have been tested over the past 10 years for the discretization of Maxwell equations in time and frequency regimes, mainly for applications in the microwave band, (2) parallel resolution algorithms for sparse linear systems based on the latest developments in numerical linear algebra, (3) modern optimization techniques based on learning and metamodeling methods and (4) software components adapted to modern high performance computing architectures. Two vertical workpackages complete this program. One of them aims to demonstrate the contributions of methodological developments and numerical tools resulting from transversal workpackages through their application to diffusion/radiation control by passive planar devices. The other, more prospective, concerns the study of basic building blocks for the realization of adaptive planar devices.

### SOLHARIS: SOLvers for Heterogeneous Architectures over Runtime systems, Investigating Scalability

**Duration:** 2018 – 2022

**Coordinator:** Alfredo Buttari (IRIT)

**CONCACE contact:** Emmanuel Agullo

#### Partners:

- IRIT Institut de Recherche en Informatique de Toulouse
- Inria Bordeaux - Sud-Ouest and Lyon
- Airbus Central R&T
- CEA Commissariat à l'énergie atomique et aux énergies alternatives

**Summary:** The SOLHARIS project aims at addressing the issues related to the development of fast and scalable linear solvers for large-scale, heterogeneous supercomputers. Because of the complexity and heterogeneity of the targeted algorithms and platforms, this project intends to rely on modern

runtime systems to achieve high performance, programmability and portability. By gathering experts in computational linear algebra, scheduling algorithms and runtimes, **SOLHARIS** intends to tackle these issues through a considerable research effort for the development of numerical algorithms and scheduling methods that are better suited to the characteristics of large scale, heterogeneous systems and for the improvement and extension of runtime systems with novel features that more accurately fulfill the requirements of these methods. This is expected to lead to fundamental research results and software of great interest for researchers of the scientific computing community.

## 10 Dissemination

**Participants:** Emmanuel Agullo, Olivier Coulaud, Luc Giraud, Carola Kruse, Paul Mycek, Guillaume Sylvand.

### 10.1 Promoting scientific activities

#### 10.1.1 Scientific events: organisation

##### Member of the organizing committees

- Luc Giraud is member of **the Gene Golub SIAM Summer School**. The twelfth Gene Golub SIAM Summer School was entitled “Financial Analytics: Networks, Learning, and High-Performance Computing”.
- Carola Kruse and Paul Mycek are members of the organising committee of the **“Sparse Days 2022”**

#### 10.1.2 Scientific events: selection

**Chair of conference program committees** Emmanuel Agullo COMPAS steering committee chair on parallel computing.

**Member of the conference program committees** ICPP’22 (E. Agullo), IPDPS’22 (E. Agullo, O. Coulaud), PDSEC’22 (O. Coulaud, L. Giraud).

**Reviewer** ICPP’22 (E. Agullo), IPDPS’22 (E. Agullo, O. Coulaud), ISC’22 (C. Kruse), PDSEC’22 (O. Coulaud, L. Giraud), PPAM’22 (C. Kruse).

#### 10.1.3 Journal

##### Member of the editorial boards

- L. Giraud is member of the editorial board of the SIAM Journal on Scientific Computing (**SISC**).

**Reviewer - reviewing activities** Computer and Fluids, Computer Methods in Applied Mechanics and Engineering, SIAM J. Matrix Analysis and Applications, SIAM J. Scientific Computing, Journal of Computational Science, Journal of Computational Physics, IEEE Transactions on Parallel and Distributed Systems (TPDS)

#### 10.1.4 Scientific expertise

- 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 member of the scientific council of **GDR Calcul**.
- Guillaume Sylvand is
  - expert in Numerical Simulation and HPC at Airbus.
  - member of the scientific council of the **ORAP**.

### 10.1.5 Research administration

- Emmanuel Agullo is member of the CDT (Technological Development Commission) at inria Centre at the Bordeaux University.
- 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).

## 10.2 Teaching - Supervision - Juries

### 10.2.1 Teaching

- Post graduate level/Master:
  - E. Agullo: Operating systems 24h at Bordeaux University ; Dense linear algebra kernels 8h, Numerical algorithms 30h at Bordeaux INP (ENSEIRB-MatMeca).
  - O. Coulaud: Paradigms for parallel computing 8h, Introduction to Tensor methods 6 h at Bordeaux INP (ENSEIRB-MatMeca).
  - L. Giraud: Introduction to intensive computing and related programming tools 30h, INSA Toulouse; Advanced numerical linear algebra 10h, ENSEEIHT Toulouse.
  - C. Kruse: Adavanced topic in numerical linear algebra, 23h, FAU Erlangen.
  - P. Mycek: Multifidelity methods 25h, INSA Toulouse.

### 10.2.2 Supervision

- PhD in progress: Mohamed Anwar Abouabdallah ; Tensor-Train approach for inference in stochastic block models, application to biodiversity characterization ; started Oct 2019; O. Coulaud, A. Franc (**PLEIADE**), N. Peyrard (Inrae)
- PhD in progress: Marek Felsoci; Fast solvers for high-frequency aeroacoustics; started Oct. 2019; G. Sylvand, E Agullo.
- PhD completed: Martina Iannacito; Linear solvers in tensorial format for high dimensional problems; started Oct 2019; O. Coulaud, L. Giraud (defended on December 9, 2022)
- PhD in progress: Romain Peressoni; Fast multidimensional scaling method for the study of biodiversity; started Oct 2019; E. Agullo, O. Coulaud, A. Franc (**PLEIADE**)
- PhD in progress: Aboul-Karim Mohamed El Maarouf; Parallel fine grain incomplete LU factorization for the solution of sparse linear systems; started: Dec. 2019; L. Giraud, A. Guermouche (**HIEPACS**).
- PhD completed: Ana Clara Ordonez Egas; Solveur linéaire haute-performance pour la thermo-hydro-mécanique avec régularisation par second gradient de dilatation, started Nov. 2019; C. Kruse, N. Tardieu (EDF) (defended 25/11/2022).
- PhD completed: Yanfei Xiang; Solution of large linear systems with massive numbers of right-hand sides. Started Nov. 2019; L. Giraud, P. Mycek (defended on December 7, 2022).



### 10.2.3 Juries

- Committee selection: Luc Giraud was member of a jury for the hiring of an associate professor in Applied math for the Université du Littoral Côte d'Opale
- PhD defense
  - Emily Bourne, "Non-Uniform Numerical Schemes for the Modelling of Turbulence in the 5D GYSELA Code"; referees: Bruno Després, Virginie Ehlacher; members: Philippe Helluy, Carola Kruse, Claudia Negulescu, Eric Sonnendrücker, Michel Mehrenberger, Virginie Grandgirard; Aix-Marseille Université, 2 Dec. 2022,
  - Yishu Du, "Fault-tolerant algorithms for iterative applications and batch schedulers"; referees: Marc Casas, Luc Giraud; members: Fanny Duffossé, Francieli Zanon Boito, Yves Robert, Lois Marchal; ENS Lyon, 1 Dec. 2022.
  - Martina Iannacito, "Numerical linear algebra and data analysis in large dimensions using tensor format"; referees: Daniel Kressner, Karl Meerbergen; members: Olivier Coulaud, Luc Giraud, Alain Franc, Anthony Nouy, Valeria Simoncini, Nick Vannieuwenhoven; Université de Bordeaux, 9 Dec. 2022
  - Romain Lion, "Réplication de données pour la tolérance aux pannes dans un support d'exécution distribué à base de tâches"; referees: Franck Cappello, Cédric Bastoul; members: Camille Coti, Amina Guermouche, Leonardo Bautista Gomez, Luc Giraud, Samuel Thibault, Université de Bordeaux, 12 Dec. 2022
  - Margot Sirdey, "Méthode itérative de Trefftz pour la simulation d'ondes électromagnétiques en trois dimensions"; referees: Bruno Després, Stéphane Lanteri; members: Hélène Barucq, Luc Giraud, Lise-Marie Imbert-Gérard, Sébastien Pernet, Sébastien Tordeux; Université de Pau et des Pays de l'Adour, 20 Dec. 2022.
  - Bastien Vieuble, "Raffinement itératif en précision mixte pour la résolution de systèmes linéaires creux de grande taille"; referees: Julien Langou, Sherry X. Li; members: Emmanuel Agullo, Marc Baboulin, Alfredo Buttari, Erin carson, Nick Higham, Serge Gratton, Théo Mary; Toulouse INP, 30 Nov. 2022
  - Yanfei Xiang, "Solution of large linear systems with a massive number of right-hand sides and machine learning"; referees: Eric De Sturler, Andreas Frommer; members: Michael Bauerheil, Luc Giraud, Paul Mycek, Carola Kruse, Jayant Sengupta, Stéphane Lanteri; Université de Bordeaux, 7 Dec. 2022

## 10.3 Popularization

### 10.3.1 Interventions

In the context of **Competitiveness Cluster for Aeronautics, Space and Embedded Systems** Luc Giraud organized on November 21, 2022, a webinar on HPC-HPDA with invited speakers C. Lapeyre (Cerfacs), F. Ravache (Sorrac), S. Requena (GENCI).

## 11 Scientific production

### 11.1 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. Jégo. '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] 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>.
- [5] 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>.
- [6] 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>.
- [7] D. A. Di Pietro, P. Matalon, P. Mycek and U. Rude. ‘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>.
- [8] 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>.
- [9] 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>.
- [10] M. Odunlami, V. Le Bris, D. Begue, 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>.

## 11.2 Publications of the year

### International journals

- [11] V. Darrigrand, A. Dumitrasc, C. Kruse and U. Rude. ‘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>.
- [12] A.-k. Mohamed El Maarouf, L. Giraud, A. Guermouche and T. Guignon. ‘Combining reduction with synchronization barrier on multi-core processors’. In: *Concurrency and Computation: Practice and Experience* 35.1 (1st Dec. 2022). DOI: [10.1002/cpe.7402](https://doi.org/10.1002/cpe.7402). URL: <https://hal.inria.fr/hal-03948901>.

### International peer-reviewed conferences

- [13] C. Castes, E. Agullo, O. Aumage and E. Saillard. ‘Decentralized in-order execution of a sequential task-based code for shared-memory architectures’. In: *IPDPSW 2022 - IEEE International Parallel and Distributed Processing Symposium Workshops*. Lyon, France: IEEE, May 2022, pp. 552–561. DOI: [10.1109/IPDPSW55747.2022.00095](https://doi.org/10.1109/IPDPSW55747.2022.00095). URL: <https://hal.inria.fr/hal-03896030>.

### Conferences without proceedings

- [14] M. Iannacito, E. Agullo, O. Coulaud, L. Giraud, G. Marait and N. Schenkels. ‘GMRES in variable accuracy: a case study in low rank tensor linear systems’. In: *GAMM - Workshop on Applied and Numerical Linear Algebra 2022*. Prague, Czech Republic, 22nd Sept. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03826879>.

- [15] M. Iannacito, O. Coulaud and A. Franc. ‘Extension of Correspondence Analysis to multiway datasets through HOSVD: a geometric framework’. In: MDS 2022 - SIAM Conference on Mathematics of Data Science. San Diego / Hybrid, United States, 26th Sept. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03826894>.

#### Doctoral dissertations and habilitation theses

- [16] A. C. O. Egas. ‘Scalable linear solver for thermo-hydro-mechanics with a second gradient of dilation regularization problems’. Ecole Doctorale Mathématiques, Informatique et Télécommunications de Toulouse, 2022.
- [17] M. Iannacito. ‘Numerical linear algebra and data analysis in large dimensions using tensor format’. Université de Bordeaux, 9th Dec. 2022. URL: <https://theses.hal.science/tel-03952711>.
- [18] Y.-F. Xiang. ‘Solution of large linear systems with a massive number of right-hand sides and machine learning’. Université de Bordeaux, 7th Dec. 2022. URL: <https://theses.hal.science/tel-03967557>.

#### Reports & preprints

- [19] E. Agullo, O. Coulaud, A. Denis, M. Faverge, A. Franc, J.-M. Frigerio, N. Furmento, A. Guilbaud, E. Jeannot, R. Peressoni, F. Pruvost and S. Thibault. *Task-based randomized singular value decomposition and multidimensional scaling*. RR-9482. Inria Bordeaux - Sud Ouest; Inrae - BioGeCo, 9th Sept. 2022, p. 37. URL: <https://hal.inria.fr/hal-03773985>.
- [20] E. Agullo, O. Coulaud, L. Giraud, M. Iannacito, G. Marait and N. Schenkels. *The backward stable variants of GMRES in variable accuracy*. RR-9483. Inria, Sept. 2022, pp. 1–77. URL: <https://hal.inria.fr/hal-03776837>.
- [21] O. Coulaud, L. Giraud and M. Iannacito. *A robust GMRES algorithm in Tensor Train format*. RR-9484. Inria, 13th Sept. 2022, pp. 1–48. URL: <https://hal.inria.fr/hal-03776529>.
- [22] O. Coulaud, L. Giraud and M. Iannacito. *On some orthogonalization schemes in Tensor Train format*. RR-9491. Inria Bordeaux - Sud-Ouest, 13th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03850387>.

### 11.3 Cited publications

- [23] 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).
- [24] 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>.
- [25] 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>.
- [26] 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>.
- [27] 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>.

- [28] E. Agullo, M. Felšöci, A. Guermouche, H. Mathieu, G. Sylvand and B. Tagliaro. *Study of the processor and memory power consumption of coupled sparse/dense solvers*. Research Report RR-9463. Inria Bordeaux Sud-Ouest, Feb. 2022, p. 17. URL: <https://hal.inria.fr/hal-03589695>.
- [29] 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>.
- [30] E. Agullo, M. Felšöci and G. Sylvand. ‘Direct solution of larger coupled sparse/dense linear systems using low-rank compression on single-node multi-core machines in an industrial context’. In: *IPDPS 2022 - 36th IEEE International Parallel and Distributed Processing Symposium*. Lyon, France: IEEE, May 2022, p. 11. DOI: [10.1109/IPDPS53621.2022.00012](https://doi.org/10.1109/IPDPS53621.2022.00012). URL: <https://hal.inria.fr/hal-03774145>.
- [31] E. Agullo, M. Felšöci and G. Sylvand. *Direct solution of larger coupled sparse/dense linear systems using low-rank compression on single-node multi-core machines in an industrial context*. Research Report RR-9453. Inria Bordeaux Sud-Ouest, Feb. 2022, p. 25. URL: <https://hal.inria.fr/hal-03557692>.
- [32] 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.
- [33] 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>.
- [34] 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>.
- [35] 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>.
- [36] S. Börm, L. Grasedyck and W. Hackbusch. *Hierarchical Matrices*. Tech. rep. 2003, pp. 1–173.
- [37] 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.
- [38] 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).
- [39] 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>.
- [40] C. Castes, E. Agullo, O. Aumage and E. Saillard. *Decentralized in-order execution of a sequential task-based code for shared-memory architectures*. Research Report RR-9450. Inria Bordeaux - Sud Ouest, 2022, p. 30. URL: <https://hal.inria.fr/hal-03547334>.
- [41] 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.
- [42] 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>.
- [43] O. Coulaud, A. A. Franc and M. Iannacito. *Extension of Correspondence Analysis to multiway datasets 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>.

- [44] 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* (June 2022). DOI: [10.1002/nla.2456](https://doi.org/10.1002/nla.2456). URL: <https://hal.science/hal-03531293>.
- [45] A. Falco. ‘Combler l’cart entre  $\mathcal{H}$ -Matrices et mthodes directes creuses pour la rsolution de systmes linaires de grandes tailles’. Theses. Universit de Bordeaux, June 2019. URL: <https://tel.archives-ouvertes.fr/tel-02183902>.
- [46] 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>.
- [47] 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).
- [48] 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>.
- [49] 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>.
- [50] 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.
- [51] 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.
- [52] W. Hackbusch. *Hierarchical Matrices: Algorithms and Analysis*. 1st. Springer Publishing Company, Incorporated, 2015.
- [53] 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>.
- [54] 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>.
- [55] C. Kruse, V. Darrigrand, N. Tardieu, M. Arioli and U. Rde. ‘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>.
- [56] C. Kruse, M. Sosonkina, M. Arioli, N. Tardieu and U. Rde. ‘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>.
- [57] V. Le Bris, M. Odunlami, D. Bgu, 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>.
- [58] B. Liz. ‘Rsolution directe rapide pour les lments finis de frontire en lectromagntisme et acoustique :  $\mathcal{H}$ -Matrices. Paralllisme et applications industrielles’. Theses. Universit Paris-Nord - Paris XIII, June 2014. URL: <https://tel.archives-ouvertes.fr/tel-01244260>.
- [59] P.-G. Martinsson and J. Tropp. ‘Randomized Numerical Linear Algebra: Foundations & Algorithms’. In: (2020). eprint: [2002.01387](https://arxiv.org/abs/2002.01387). URL: <http://arxiv.org/abs/2002.01387>.
- [60] A.-K. Mohamed El Maarouf, L. Giraud, A. Guermouche and T. Guignon. ‘Combining reduction with synchronization barrier on multi-core processors’. working paper or preprint. Feb. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03577306>.

- [61] 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](https://doi.org/10.1063/1.4984266). URL: <http://aip.scitation.org/doi/10.1063/1.4984266>.
- [62] I. V. Oseledets. ‘Tensor-Train Decomposition’. In: *SIAM Journal on Scientific Computing* 33.5 (Jan. 2011), pp. 2295–2317. DOI: [10.1137/090752286](https://doi.org/10.1137/090752286). URL: <https://doi.org/10.1137/090752286>.
- [63] 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>.
- [64] 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](https://doi.org/10.1109/TMAG.2019.2954584). URL: <https://hal.inria.fr/hal-03150956>.
- [65] 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>.
- [66] 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>.