Research school on Iterative Methods for Partial Differential Equations (IMPDE2023)

Book of abstracts

Editors: Bastien Chaudet-Dumas, Liu-Di Lu, Lucas Perrin

Website of the event: https://impde2023.sciencesconf.org

Introduction

Due to the rise of supercomputers with more and more processors, the study of parallel algorithms has become a fast-growing research field in the last decades. In the context of the numerical solution of partial differential equations (PDEs), the complexity of the underlying phenomena (multiphysics formulation, high frequency wave propagation, long time scale, etc) often requires extremely fine space and/or time discretizations, which leads to very large systems. Iterative solvers are among the most efficient methods to tackle this type of problems.

This research school, which is exclusively dedicated to young researchers (PhD students, post-docs...), focuses on this research topic. The goal is to bring together young people working on iterative methods for PDEs, in order to offer them an opportunity to meet each other, to discuss their research works and also to initiate new collaborations.

The scientific program of this two-day event is divided into four thematic blocks: **Krylov Methods, Domain Decomposition Methods, Multigrid Methods** and **Parallelism in Time Methods**. Each block lasts half a day and is composed of an introductory mini-course on the topic, followed by two advanced presentations.

Krylov Methods

Krylov subspace methods

Hussam Al Daas*[1]

[1]STFC, Rutherford Appleton Laboratory – United Kingdom

Abstract

Krylov subspace methods are widely used to solve large-scale linear systems, eigenvalue problems, and other matrix computations. In this mini-course, we will provide a brief overview of Krylov subspace methods focusing on the two widely used methods, the conjugate gradient for Hermitian positive definite matrices and the generalised minimal residual for general matrices. We will show how both methods are derived and discuss their underlying algorithms, numerical behaviour, and convergence properties. We will also highlight recent advances in this amazing research field.

*Speaker

Preconditioning the stage equations of implicit Runge Kutta methods for parabolic PDEs

Michal Outrata*[1] and Martin Gander[2]

[1]Virginia Tech – United States
[2]University of Geneva – Switzerland

Abstract

When using Solving parabolic PDEs, we need to discretize in both space and time and for the time domain and implicit Runge-Kutta methods are a common choice for the latter. There, solving the stage equations is often the computational bottleneck, as the dimension of the stage equations Mk = b for an s-stage Runge-Kutta method becomes sn where the spatial discretization dimension n can be very large. Hence the solution process often requires the use of iterative solvers, whose convergence can be less than satisfactory. Moreover, due to the structure of the stage equations, the matrix M does not necessarily inherit any of the preferable properties of the spatial operator, making GMRES the go-to solver and hence there is a need for a preconditioner. Recently in (1) and also (2,3) a new block preconditioner was proposed and numerically tested with promising results.

Using spectral analysis and the particular structure of M, we study the properties of this class of preconditioners, focusing on the eigen properties of the preconditioned system, and we obtain interesting results for the eigenvalues of the preconditioned system for a general Butcher matrix. In particular, for low number of stages, i.e., s = 2,3, we obtain explicit formulas for the eigen properties of the preconditioned system and for general s we can explain and predict the characteristic features of the spectrum of the preconditioned system observed in (2). As the eigenvalues alone are known to **not** be sufficient to predict the GMRES convergence behavior in general, we also focus on the eigenvectors, which altogether allows us to give descriptive bounds of the GMRES convergence behavior for the preconditioned system.

We then numerically optimize the Butcher tableau for the performance of the entire solution process, rather than only the order of convergence of the Runge-Kutta method. To do so requires careful balancing of the numerical stability of the Runge-Kutta method, its order of convergence and the convergence of the iterative solver for the stage equations. We illustrate our result on test problem and then outline possible generalizations.

References

(1) M. Neytcheva, O. Axelsson. Numerical solution methods for implicit Runge-Kutta methods of arbitrarily high order. Proceedings of ALGORITHMY 2020, ISBN : 978-80-227-5032-5, 2020.
(2) M. M. Rana, V. E. Howle, K. Long, A. Meek, W. Milestone. A New Block Preconditioner

*Speaker

for Implicit Runge-Kutta Methods for Parabolic PDE Problems. SIAM Journal on Scientific Computing, vol (43): S475–S495, 2021.
(3) M. R. Clines, V. E. Howle, K. R. Long. Efficient order-optimal preconditioners for implicit Runge-Kutta and Runge-Kutta-Nystrom methods applicable to a large class of parabolic and hyperbolic PDEs. arXiv: https://arxiv.org/abs/2206.08991, 2022.

Randomization to speed-up Krylov subspace methods

Edouard Timsit*[1,2]

[1]Inria de Paris – Institut National de Recherche en Informatique et en Automatique – France
[2]Laboratoire Jacques-Louis Lions – Sorbonne Universite, Centre National de la Recherche Scientifique, Université Paris Cité – France

Abstract

The natural basis of the Krylov subspace is expected to be ill-conditioned, as the vectors grow colinear with the dominant eigenvector. It numerically infeasible to express vectors of the Krylov subspace in such a basis. For this reason, Krylov solvers usually perform a subsequent orthonormalization of the Krylov basis, which cost often defines that of the solver. For some Krylov solvers, randomization can speed up this orthonormalization process, while preserving performance and stability. In particular, we show that randomized orthogonal projections over the Krylov subspace produce quasi-optimal results.

*Speaker

Domain Decomposition Methods

Introduction to domain decomposition methods

Tommaso Vanzan*[1]

[1]Ecole Polytecnique Federale de Lausanne – Switzerland

Abstract

Domain decomposition methods are effective divide and conquer algorithms to solve large scale partial differential equations on modern high performance clusters. They are based on the idea of decomposing the computational domain into several subdomains, where the partial differential equation can be solved more easily and in parallel. The local solutions are then recombined together, and the process is iterated until convergence is reached. The first part of the mini-course will introduce the most classical one-level overlapping and nonoverlapping methods, recalling convergence results, discrete and continuous analogs and implementation details. In the second part, we will briefly touch more advanced topics such as multi-level methods and applications to multi-physics problems.

*Speaker

Parallel-friendly non local Optimized Schwarz Method (OSM) for electromagnetism

Roxane Atchekzai*[1], Xavier Claeys[2], and Matthieu Lecouvez[3]

[1]Laboratories – Laboratoire Jacques-Louis Lions (LJLL), CEA-DAM CESTA – France
[2]Laboratory – Laboratoire Jacques-Louis Lions (LJLL) – France
[3]Institution – CEA-DAM CESTA – France

Abstract

We consider the propagation of harmonic electromagnetic waves and we are aiming at an efficient domain decomposition strategy without overlap. In (1), it has been proved for general case of substructuring domain decomposition for harmonic wave problems that a non-local treatment on the transmission conditions accross interfaces allows to deal with cross points while maintaining a convergence speed that is robust with respect to the discretization parameter. In this non-local OSM method, the transmission conditions are implemented with a non-local exchange operator defined on the skeleton of the partition which is the union of the interfaces between subdomains. However, introducing non-local operations may affect the parallelism of the domain decomposition. That is why in this presentation, we will talk about different ways of reaching scalability and parallel efficiency despite the possible non-locality of the exchange operator in particular recycling and preconditionning techniques.

References

(1) X. Claeys. Non-Local Variant of the Optimised Schwarz Method for Arbitrary Non-Overlapping Subdomain Partitions. *ESAIM: Mathematical Modelling and Numerical Analysis* 55, no 2 (1st March 2021): 429-48.

*Speaker

Two-level domain decomposition methods on perforated domains

Miranda Boutilier*[1], Victorita Dolean[1,2], and Konstantin Brenner[1]

[1]Université Côte d'Azur – Laboratoire Jean-Alexandre Dieudonné – France
[2]University of Strathclyde – United Kingdom

Abstract

We propose a new coarse space for the two-level Restricted Additive Schwarz (RAS) preconditioner to efficiently solve elliptic equations on perforated domains. The numerous polygonal perforations in the model domain represent realistic structures (buildings, walls, etc.) in given areas of Nice, France. With the eventual goal of modeling urban floods via the nonlinear Diffusive Wave equation, we focus on the solution of linear problems on these domains. Similar to other multi-scale numerical methods, this coarse space is spanned by basis functions that are locally discrete harmonic inside of each subdomain. It is based on nodal degrees of freedom that occur at the intersection between the perforations and the subdomain boundaries. Numerical experiments are shown using the two-level RAS method as both a multiplicative iterative method and a Krylov preconditioner. These numerical results show that the new coarse space is very robust and accelerates iteration counts in both methods, independent of the complexity of the data. Additionally, we see that the coarse space has approximation properties that make it a strong coarse approximation in its own right.

*Speaker

Multigrid Methods

A mini-course on multigrid methods

Hardik Kothari*[1]

[1]Universitá della Svizzera Italiana – Switzerland

Abstract

Multigrid methods are some of the most powerful techniques for solving linear systems of equations arising from the discretization of partial differential equations (PDEs). In this mini-course, we will introduce fundamental concepts of multigrid methods, including their historical developments, mathematical foundations, and implementation. We will discuss in detail the multigrid algorithm and the key components of the method, namely: coarse grids, transfer operators, and the concepts of smoothers. In the later part of the course, we will extend the concepts from linear multigrid methods to their nonlinear counterparts and discuss the necessary modifications.

*Speaker

A multi-level spectral deferred correction method

Ikrom Akramov*[1], Robert Speck , Daniel Ruprecht[1], Matthew Emmet , Michael Minion , Matthias Bolten , and Rolf Krause

[1]Technical University Hamburg Harburg [Hamburg] – Germany

Abstract

Spectral deferred correction methods (SDC) are iterative schemes for computing the numerical solution of initial value problems. SDC can be understood as applying an appropriate preconditioner to a Picard iteration to achieve faster and more robust convergence to a collocation solution. It has been shown that SDC can achieve arbitrary order of accuracy and possesses good stability properties. In this talk, we will examine the combination of multigrid and SDC, known as the multi-level spectral deferred correction (MLSDC) method, where sweeps are performed on a hierarchy of levels, and the FAS correction term couples solutions across different levels. We will examine various strategies to reduce the computational cost of correction sweeps on the coarser levels. I will present numerical examples to demonstrate the impact of multi-level coarsening on the convergence and cost of SDC integration.

*Speaker

Multigrid preconditioners for the cardiac bidomain model: a performance analysis on HPC architectures

Edoardo Centofanti*[1]

[1]Dipartimento di Matematica, Università degli Studi di Pavia – Italy

Abstract

PETSc (Portable Extensible Toolkit for Scientific Computation) (1, 2), is a popular suite of data structures and routines for the numerical solution of partial differential equations. A feature of interest is the implementation of algebraic multigrid methods (3) as solvers and preconditioners, in order to provide fast, efficient, and scalable solutions to linear systems. Besides its own multigrid implementation, PETSc includes routines from HYPRE (4), a software library designed to solve large sparse linear equations on parallel computers and to run on the most modern architectures. This library has its cornerstone in the parallel implementation of the algebraic multigrid method known as BoomerAMG (Boomer Algebraic Multigrid (5)), suitable also for GPU computing. In this talk, we will explore the performances of these algorithms in preconditioning linear systems generated from the time and space finite element discretizations of the Bidomain cardiac model (6), a reaction-diffusion system of ordinary and partial differential equations describing the space-time evolution of cardiac potentials and ionic currents (7, 8). We will consider different settings for our Bidomain solvers and present scalability tests performed on structured and unstructured meshes, on both CPU and GPU parallel architectures.

References

(1) Satish Balay et al., PETSc Web page https://petsc.org/
(2) Richard Tran Mills et al., Toward performance-portable PETSc for GPU-based exascale systems. Parallel Computing, 108:102831, 2021.
(3) J. W. Ruge and K. St uben, Algebraic Multigrid in Multigrid Methods, chapter 4, pages 73-130, 1987. SIAM.
(4) Robert D. Falgout, Jim E. Jones and Ulrike Meier Yang, The Design and Implementation of hypre, a Library of Parallel High Performance Preconditioners. Numerical Solution of Partial Differential Equations on Parallel Computers, pages 267–294, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
(5) Van Emden Henson and Ulrike Meier Yang, BoomerAMG: A parallel algebraic multigrid solver and preconditioner. Applied Numerical Mathematics 41(1):155–177, 2002.
(6) P. Colli Franzone, L.F. Pavarino, S. Scacchi, Mathematical Cardiac Electrophysiology. 2014. Springer, MS&A. Modeling, Simulation and Applications 13.
(7) G. Plank, M. Liebmann, R.W. dos Santos, E.J. Vigmond, G. Haase Algebraic multigrid preconditioner for the cardiac bidomain model. 2007. IEEE Transactions on Biomedical Engineering, 54(4), 585-596.
(8) M. Pennacchio , V. Simoncini Algebraic multigrid preconditioners for the bidomain reaction diffusion system. 2009. Applied numerical mathematics, 59(12), 3033-3050.

*Speaker

Parallel in Time Methods

An introduction to Parallel-in-Time methods

Ausra Pogozelskyte*[1]

[1]Université de Genève = University of Geneva – Switzerland

Abstract

With the increase in the number of processors in supercomputers, additional spatial parallelism has been possible allowing the usage of finer models. However, in order to preserve the stability of numerical schemes, with a refinement of the spatial grid often comes a refinement of the time grid. This causes increased computational times.

In this mini-course, we will introduce some of the main parallel-in-time methods: Parareal, Multigrid Reduction-in-Time (MGRIT), Space-Time Multigrid (STMG), Schwarz Waveform Relaxation (SWR), ParaExp. Each of these methods will be briefly presented and a selection of their properties will be discussed.

*Speaker

Low-rank Parareal: a low-rank parallel-in-time integrator

Benjamin Carrel*[1], Martin J. Gander[2], and Bart Vandereycken[2]

[1]Department of Mathematics, University of Geneva – Switzerland
[2]University of Geneva – Switzerland

Abstract

The Parareal algorithm of Lions, Maday, and Turinici is a well-known time parallel algorithm for evolution problems. It is based on a Newton-like iteration, with cheap coarse corrections performed sequentially, and expensive fine solves performed in parallel.
On the other hand, the dynamical low-rank approximation (DLRA), proposed by Koch and Lubich, is a recent technique that allows for solving large-scale problems on the manifold of low-rank matrices. In particular, the method has a low memory footprint and is much faster than dense techniques. The cost and accuracy are mostly governed by the rank chosen for the approximation. Due to its large-scale purpose, having a time parallel algorithm would be particularly interesting.
After introducing the DLRA, we will see how to exploit its properties and get a low-rank Parareal algorithm. The new technique has both linear and superlinear convergence bounds, which are verified numerically. If time allows, we'll introduce a Krylov-based method for solving the DLRA, which is particularly useful for parabolic problems.

*Speaker

A unified analysis framework for iterative Parallel-in-Time algorithms

Thibaut Lunet*[1]

[1]Technische Universität Hamburg-Harburg – Germany

Abstract

Recent advances in supercomputing architectures in the last decades have pushed the scientific community to develop more efficient algorithms for high performance computing (HPC).
In particular, the switch from improving processor speed to increasing concurrency in HPC have motivated the development of new parallelization algorithms that can harness the computing power of massively parallel HPC architectures.
For two decades, there have been research efforts in developing parallel computing capabilities for time integration methods used the simulation of time dependent problems (Numerical Weather Prediction, Computational Fluid Dynamics, Fluid-Structure interactions, ...).
However, to develop such "Parallel-in-Time" algorithms (or PinT), one needs to deal with the sequential nature of time: it is necessary to know the past and the present before computing the future.
Hence, computing both past, present and future in parallel needs the development of new algorithms with different computation paradigms as the classical ones used for time-dependent problems.

Over the last decades, many research efforts have studied iterative Parallel-in-Time (PinT) algorithms, in particular Parareal (Lions, Maday, Turinici), PFASST (Minion and Emmett), MGRIT (Falgout, Friedhoff, Kolev, MacLachlan, Schroder) and a specific form of Space-Time Multigrid (Gander and Neumueller).
While various convergence analyses exist for each algorithm separately, it is difficult to connect them and compare convergence of these iterative PinT methods when applied to various model problems, and in applications.
In this talk, I will present a new approach that lets us analyze the convergence of these four iterative algorithms in a single framework. Following an idea of Gander and Hairer already used to analyze Parareal convergence, this framework is based on an abstract view of each iterative PinT algorithm and provides understanding of their different convergence mechanisms. We use it to show some key similarities and differences between all those iterative methods while focusing on the Dahlquist equation, the fundamental time-dependent test problem.

*Speaker

Posters

Discrete time analysis for domain decomposition

Arthur Arnoult*[1]

[1]Université Sorbonne Paris Nord – Université Paris XIII - Paris Nord – France

Abstract

Optimized Schwarz waveform relaxation (OSWR) is a domain decomposition algorithm for solving partial differential equations on small subdomains in order to accelerate numerical resolution. This poster shows a new approach, that provides new results in the convergence analysis of OSWR iterations for parabolic problems.
This new approach relies on the time discretization of the domain decomposition equations with backward Euler, in order to obtain a system of differential equations that can be analytically solved. Contrary to the classical method that choses the Robin parameter that minimizes the contraction ratio of the Fourier transform of the continuous in time solution, this method minimizes the contraction matrix norm of the discrete time solution.

This method allows to define efficient optimized Robin parameters that depend on the targeted iteration count, a property that is shared by the actual observed optimal parameters, while traditional Fourier analysis in the time direction leads to iteration independent parameters.
Numerical results show that this parameter is an accurate estimation of the optimal Robin parameter, which allows to perform the smallest number of iterations possible.

*Speaker

Partial and full improvement of the convergence of the Bl-BiCG and Bl-BiCGStab methods

Fatima Bouyghf*[1]

[1]Laboratoire de Mathématiques Nicolas Oresme – Université de Caen Normandie, Centre National de la Recherche Scientifique – France

Abstract

In this poster, we present a technique to improve the convergence of the block version of some Krylov methods for solving non-symmetric linear systems of equations with multiple right-hand sides. Especially, short-recurrence type methods. This technique is inspired from the famous block GMRES method (Bl-GMRES). Using orthogonal projectors to minimize the norm of the residual matrix computed in each method. The considered methods are block BiCG (Bl-BiCG) and block BiCGStab (Bl-BiCGStab) methods. To show the performance of our enhanced algorithms, we compare them with the Bl-GMRES method as the most optimal method in terms of precision.

*Speaker

One-shot domain decomposition methods for component-based model order reduction

Giulia Sambataro*[1], Angelo Iollo[2], and Tommaso Taddei[2]

[1]CERMICS – Ecole des Ponts ParisTech, Inria Paris – France
[2]Inria Bordeaux South-West – INRIA Bordeaux South-West, Université de Bordeaux (Bordeaux, France) – France

Abstract

We propose a component-based (CB) parametric model order reduction (pMOR) procedure for a class of problems in nonlinear mechanics with internal variables. The work is is motivated by applications to thermo-hydro-mechanical (THM) systems for radioactive waste disposal. The solution to this coupled system depends on several parameters, which might be related to the geometric configuration (e.g. the number of repositories, their distance or their size) or the material properties of the medium. The CB-pMOR formulation is based on overlapping subdomains. We first devise a constrained optimization procedure that minimizes jumps at components' interfaces subject to the (approximate) satisfaction of the PDE in each component. Then, we introduce suitable low-dimensional control variables to recast the optimization problem into an unconstrained nonlinear least-square problem that can be effectively solved using the Gauss-Newton method.

*Speaker

Convergence analysis of multi-step one-shot methods for linear inverse problems

Marcella Bonazzoli* , Houssem Haddar* , and Tuan Anh Vu*[1]

[1]École Nationale Supérieure de Techniques Avancées – MESRI – France

Abstract

In this work we are interested in general linear inverse problems where the corresponding forward problem is solved iteratively using fixed point methods. Then one-shot methods, which iterate at the same time on the forward problem solution and on the inverse problem un- known, can be applied. We analyze two variants of the so-called multi-step one-shot methods and establish sufficient conditions on the descent step for their convergence, by studying the eigenvalues of the block matrix of the coupled iterations. Several numerical experiments are provided to illustrate the convergence of these methods in comparison with the classical usual and shifted gradient descent. In particular, we observe that very few inner iterations on the forward problem are enough to guarantee good convergence of the inversion algorithm.

*Speaker