MS-10: Numerical linear algebra for PDEs

Organiser: Niall Madden (NUI Galway)

Theme: This mini-symposium will feature takes on varied topics broadly related to linear and nonlinear solvers for problems arising from the discretization of PDEs. As such, it will include elements of both theoretical and applied numerical linear algebra.

A boundary-layer preconditioner for singularly perturbed convection diffusion problems

Niall Madden

National University of Ireland, Galway

The numerical analysis of discretizations of singularly perturbed differential equations is an established sub-discipline within the study of the numerical approximation of solutions to differential equations. The motivatation stems the wide range of real-world problems whose solutions exhibit boundary and interior layers, and the challenges posed when trying to solve these problems numerically.

Consequently, much is known about how to accurately and stably discretize such equations in order to properly resolve the layer structure present in their continuum solutions. However, despite being a key step in the numerical simulation process, the study of efficient and accurate solution of the associated linear systems is somewhat neglected (though not entirely, see, e.g., [1, 2, 4]).

In this talk, we discuss problems associated with the application of direct solvers to these discretizations. We then propose a preconditioning strategy that is tuned to the matrix structure induced by using layer-adapted meshes for convection-diffusion equations, proving a strong condition-number bound on the preconditioned system in one spatial dimension, and a weaker bound in two spatial dimensions. Numerical results confirm the efficiency of the resulting preconditioners in one and two dimensions, with time-to-solution of less than one second for representative problems on 1024×1024 meshes and up to $40\times$ speedup over standard sparse direct solvers.

This talk is based on [3]; see also `https://arxiv.org/abs/2108.13468`.

This is joint work with Scott P. MacLachlan (Memorial University) and Thái Anh Nhan (Holy Names University).

Bibliography

[1] Carlos Echeverría, Jörg Liesen, Daniel B. Szyld, and Petr Tichý. Convergence of the multiplicative Schwarz method for singularly perturbed convection-diffusion problems discretized on a Shishkin mesh. *Electron. Trans. Numer. Anal.*, 48:40–62, 2018.

[2] S. MacLachlan and N. Madden. Robust solution of singularly perturbed problems using multigrid methods. *SIAM J. Sci. Comput.*, 35:A2225–A2254, 2013.

[3] Scott P. MacLachlan, Niall Madden, and Thái Anh Nhan. A boundary-layer preconditioner for singularly perturbed convection diffusion. *SIAM Journal on Matrix Analysis and Applications*, 43(2):561–583, 2022.

[4] T.A. Nhan and N. Madden. Cholesky factorisation of linear systems coming from finite difference approximations of singularly perturbed problems. In *BAIL 2014–Boundary and Interior Layers, Computational and Asymptotic Methods*, Lect. Notes Comput. Sci. Eng., pages 209–220. Springer International Publishing, 2015.

A scalable and robust vertex-star relaxation for high-order FEM

Patrick E. Farrell

University of Oxford

High-order finite element methods (FEM) offer numerous advantages. They are especially attractive on modern supercomputers, due to their arithmetic intensity and rapid convergence for smooth solutions. However, all aspects of a code must change at high-order, from the choice of basis functions to matrix-free assembly strategies and on to postprocessing and visualisation. A particularly important challenge is to develop preconditioners that operate matrix-free, without ever accessing even the local tensor for a single cell.

One promising strategy for preconditioners for high-order FEM is the use of p-multigrid. Pavarino proved that the two-level method with vertex patch relaxation for the high-degree problem and a low-order coarse space gives a solver that is robust in polynomial degree for symmetric and coercive problems [1]. However, for very high polynomial degree it is not feasible to assemble or factorize the matrices for each vertex patch, since they are dense.

In this work we introduce a direct solver for separable vertex patch problems that scales to very high polynomial degree on tensor product cells. The solver constructs a carefully-chosen tensor product basis that diagonalizes the blocks in the stiffness matrix for the internal degrees of freedom of each individual cell. As a result, the non-zero structure of the cell matrices is that of the graph connecting internal degrees of freedom to their projection onto the facets. In the new basis, the patch problem is as sparse as a low-order finite difference discretization, while having a sparser Cholesky factorization. We can thus afford to assemble and factorize the matrices for the vertex-patch problems, even for very high polynomial degree. In turn, this enables the use of fast p-multigrid solvers. In the non-separable case, the method can be applied as a preconditioner by approximating the problem with a separable surrogate.

We demonstrate the approach by solving the Poisson equation and a $H(\mathrm{div})$-conforming interior penalty discretization of linear elasticity in two dimensions at polynomial degree $p = 31$ and in three dimensions at $p = 15$.

This is joint work with Pablo D. Brubeck (Oxford). Supported by a Mathematical Institute departmental scholarship, and the Engineering and Physical Sciences Research Council, grants EP/R029423/1 and EP/W026163/1.

Bibliography

[1] L. F. Pavarino. Additive Schwarz methods for the p-version finite element method. *Numer. Math.* 66:493–515, (1993).

Flexible infinite GMRES for parameterized linear systems

Siobhán Correnty

KTH Royal Institute of Technology

We seek the numerical solution to the large sparse linear system

$$A(\mu)x(\mu) = b, \tag{2}$$

where $\mu \in \mathbb{C}$, $A(\mu) \in \mathbb{C}^{n\times n}$ nonsingular, analytic and nonlinear in μ, and $b \in \mathbb{C}^n$. Under these assumptions, the matrix $A(\mu)$ can be expressed locally by an infinite Taylor series expansion centered around origin, i.e.,

$$A(\mu) = \sum_{\ell=0}^{\infty} A_\ell \mu^\ell, \qquad A_\ell := A^{(\ell)}(0)\frac{1}{\ell!} \in \mathbb{C}^{n\times n}. \tag{3}$$

In our setting, we assume further that the Taylor coefficients in (3) do not vanish after a certain degree, and many of the derivatives of $A(\mu)$ are computationally available. The method proposed here efficiently approximates the solution to (2) for many values of the parameter μ simultaneously. This novel approach offers a significant reduction in complexity over the prior work [1].

The nonlinear dependence on the parameter μ in (2) is addressed with a technique called companion linearization, commonly used in the study of polynomial eigenvalue problems. The arising system, linear in the parameter μ, is approximated within a flexible right-preconditioned GMRES framework. The basis matrix for the Krylov subspace is built just once using the infinite Arnoldi method [2], a process independent of the truncation parameter m. As this process can be carried out in a finite number of operations, we, in theory, take $m \to \infty$ while constructing the basis matrix.

The preconditioner is applied almost exactly when the residual of the outer method is large, and with decreasing accuracy as the residual is reduced, as proposed in [3]. In practice, the level of accuracy can be relaxed dramatically without degrading convergence.

We analyze our method in a way which is analogous to the standard convergence theory for the method GMRES for linear systems. The competitiveness of our method is illustrated with large-scale problems arising from a finite element discretization of a Helmholtz equation with parameterized material coefficient.

This is joint work with Elias Jarlebring (KTH Royal Institute of Technology) and Kirk M. Soodhalter (Trinity College Dublin). This work was funded by The Swedish research council (VR).

Bibliography

[1] E. Jarlebring, S. Correnty. Infinite GMRES for parameterized linear systems. *Technical report*, arXiv:2102.04082v2, 2021

[2] E. Jarlebring, W. Michiels, and K. Meerbergen. A linear eigenvalue algorithm for the nonlinear eigenvalue problem. *Numer. Math.*, 122(1):169–195, 2012

[3] V. Simoncini and D. B. Szyld. Theory of inexact Krylov subspace methods and applications to scientific computing. *SIAM J. Sci. Comput.*, 25(2):454–477, 2002

Analysis of block GMRES using a *-algebra-based approach

Kirk M. Soodhalter

Trinity College Dublin

We discuss the challenges of extending convergence results of classical Krylov subspace methods to their block counterparts and propose a new approach to this analysis. Block KSMs such as block GMRES are generalizations of classical KSMs, and are meant to iteratively solve linear systems with multiple right-hand sides (a.k.a. a block right-hand side) all-at-once rather than individually. However, this all-at-once approach has made analysis of these methods more difficult than for classical KSMs because of the interaction of the different right-hand sides. We have proposed an approach built on interpreting the coefficient matrix and block right-hand side as being a matrix and vector over a *-algebra of square matrices. This allows us to sequester the interactions between the right-hand sides into the elements of the *-algebra and (in the case of GMRES) extend some classical GMRES convergence results to the block setting. We then discuss some challenges which remain and some ideas for how to proceed.

This is joint work with Marie Kubiínová from Czech Academy of Sciences, Institute of Geonics, Ostrava, Czech Republic (formerly)

Bibliography

[1] Marie Kubiínová and Kirk M. Soodhalter. Admissible and attainable convergence behavior of block Arnoldi and GMRES. *SIAM Journal on Matrix Analysis and Applications*, 41 (2), pp. 464-486, 2020.

Preconditioned iterative methods for multiple saddle-point systems arising from PDE-constrained optimization

John W. Pearson

University of Edinburgh

Optimization problems subject to PDE constraints form a mathematical tool that can be applied to a wide range of scientific processes, including fluid flow control, medical imaging, biological and chemical processes, and many others. These problems involve minimizing some function arising from a physical objective, while obeying a system of PDEs which describe the process. Of key interest is the numerical solution of the discretized linear systems arising from such problems, and in this talk we focus on preconditioned iterative methods for these systems.

In particular, we describe recent advances in the preconditioning of multiple saddle-point systems, specifically positive definite preconditioners which can be applied within MINRES, which may find considerable utility for solving these optimization problems as well as other applications. We discuss an inexact active-set method for large-scale nonlinear PDE-constrained optimization problems, coupled with block diagonal and block triangular preconditioners for multiple saddle-point systems which utilize suitable approximations for the relevant Schur complements.

Further, we discuss an alternative structure of a preconditioner for multiple saddle-point systems, which may be applied within the MINRES algorithm and lead to a guaranteed convergence rate, and often demonstrates superior convergence as opposed to widely-used block diagonal preconditioners.

This is joint work with Andreas Potschka (TU Clausthal), with associated papers available at [1, 2].

Bibliography

[1] John W. Pearson and Andreas Potschka. A Preconditioned Inexact Active-Set Method for Large-Scale Nonlinear Optimal Control Problems. arXiv preprint `arXiv:2112.05020`, 2021.

[2] John W. Pearson and Andreas Potschka. On Symmetric Positive Definite Preconditioners for Multiple Saddle-Point Systems. arXiv preprint `arXiv:2106.12433`, 2022.

A recycling preconditioning method for crack propagation problems

Xiao-Chuan Cai

University of Macau

In this talk, we discuss a recycling preconditioning method with auxiliary tip subspace for solving a sequence of highly ill-conditioned linear systems of equations of different sizes arising from elastic crack propagation problems discretized by an extended finite element method. The preconditioned linear systems are solved by a Krylov subspace method using a non-trivial initial guess constructed with a modification of an approximate solution around the crack tips. The strategy accelerates the convergence remarkably. Numerical experiments demonstrate the efficiency of the proposed algorithm applied to problems with several types of cracks.

This is a joint work with X. Chen.

Preconditioning the Stage Equations of Implicit Runge Kutta Methods

Michal Outrata

University of Geneva

When using implicit Runge-Kutta methods for solving parabolic PDEs, solving the stage equations is often the computational bottleneck, as the dimension of the stage equations

$$M\mathbf{k} = \mathbf{b}$$

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 [2] and also [1] 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 [1]. 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 to carefully balance the numerical stability of the Runge-Kutta method, its order of convergence, and also the convergence of the iterative solver for the stage equations. We illustrate our result on test problems with an advection-diffusion spatial operator and then outline possible generalizations.

Bibliography

[1] M. M. Rana, V. E. Howle, K. Long, A. Meek, W. Milestone. A New Block Preconditioner for Implicit Runge-Kutta Methods for Parabolic PDE Problems. *SIAM Journal on Scientific Computing*, volume (43): S475–S495, 2021

[2] 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

Provable convergence rate for asynchronous methods via randomize linear algebra

Daniel B. Szyld

Temple University

Asynchronous methods refer to parallel iterative procedures where each process performs its task without waiting for other processes to be completed, i.e., with whatever information it has locally available and with no synchronizations with other processes. For the numerical solution of a general partial differential equation on a domain, Schwarz iterative methods use a decomposition of the domain into two or more (usually overlapping) subdomains. In essence one is introducing new artificial boundary conditions Thus each process corresponds to a local solve with boundary conditions from the values in the neighboring subdomains.

Using this method as a solver, avoids the pitfall of synchronization required by the inner products in Krylov subspace methods. A scalable method results with either optimized Schwarz or when a coarse grid is added. Numerical results are presented on large three-dimensional problems illustrating the efficiency of asynchronous parallel implementations.

Most theorems show convergence of the asynchronous methods, but not a rate of convergence. In this talk, using the concepts of randomized linear algebra, we present provable convergence rate for the methods for a class of nonsymmetric linear systems.

Matrix equation techniques for certain evolutionary partial differential equations

Davide Palitta

Dipartimento di Matematica and AM^2 , Alma Mater Studiorum - Universita di Bologna, ‘ Piazza di Porta S. Donato, 5, I-40127 Bologna, Italy

In this talk we show how the linear system stemming from the all-at-once approach for certain evolutionary partial differential equations (PDEs) can be recast in terms of a Sylvester matrix equation which naturally encodes the separability of the time and space derivatives.

Combining appropriate projection techniques for the space operator together with a full exploitation of the structure of the discrete time derivative, we are able to efficiently solve problems with a tremendous number of degrees of freedom while maintaining a low storage demand in the solution process.

Such a scheme can be easily adapted to solve many different time-dependent PDEs and several numerical results are shown to illustrate the potential of our novel approach.

Bibliography

[1] Davide Palitta. Matrix Equation Techniques for certain Evolutionary Partial Differential Equations. *J Sci Comput* 87, 99 (2021).

Extrapolation methods as nonlinear Krylov methods

Conor McCoid

University of Geneva

Krylov methods are commonplace for solving of linear problems. Their use for nonlinear problems requires generalizing them. In linear examples some extrapolation methods have been shown to be equivalent to Krylov subspace methods. Since extrapolation methods can be applied to nonlinear problems, we can view these methods as nonlinear Krylov methods. To show the broad class of equivalences between these methods and others, we build each from their ancestral root-finding method, the multisecant equations, which are an extension of the secant equations to higher dimensions.

This work was completed under the supervision of Prof. Martin J. Gander (Geneva). Supported by the Swiss National Science Foundation.

A Domain Decomposition based preconditioner for Discretised Integral equations in two dimensions

V A Kandappan

Indian Institute of Technology, Madras

In this talk, we present a new preconditioner for dense linear systems arising from discretised integral equations in two dimensions. The developed preconditioner combines the traditional domain decomposition technique with hierarchical matrix representations, in particular the HODLR2D [1]. We apply this preconditioner to improve the conditioning of the system and thereby accelerate the convergence of the iterative solver. We present the preconditioner's performance through numerical experiments on dense linear systems from discretised integral equations in two dimensions. We also compare the performance of the developed new preconditioner with a block diagonal preconditioner.

This is joint work with Sivaram Ambikasaran (Indian Institute of Technology Madras)

Bibliography

[1] Kandappan, V. A., Vaishnavi Gujjula, and Sivaram Ambikasaran. HODLR2D: A new class of Hierarchical matrices. *arXiv preprint* arXiv:2204.05536, (2022).