University of Strathclyde

28th Biennial Conference on Numerical Analysis

25 – 28 June, 2019

Contents

Nothing in here

Introduction

Dear Participant,

On behalf of the Strathclyde Numerical Analysis and Scientific Computing Group, it is our pleasure to welcome you to the 28th Biennial Numerical Analysis conference. This is the sixth time the meeting has been held at Strathclyde, continuing the long series of conferences originally hosted in Dundee. We are delighted to have over 180 registered participants from all over the world, some who are familiar faces at this series of meetings, but also many who are attending for the first time.

The conference is rather unusual in the sense that it seeks to encompass all areas of numerical analysis, and the list of invited speakers reflects this aim. We have once again been extremely fortunate in securing a top line-up of plenary speakers, and we very much hope that you enjoy sampling the wide range of interesting topics which their presentations will cover. The specialised minisymposia and sessions of contributed talks scheduled cover a wealth of additional areas, confirming the continued strength of numerical analysis as an active field of research almost 55 years since the first meeting in the series.

The meeting is funded almost entirely from the registration fees of the participants. However, we are very pleased to be able to subsidise the student registration fee and provide additional financial support for some overseas participants thanks to the *Dundee Numerical Analysis Fund*, started by Professor Gene Golub from Stanford University in 2007. We are also thankful to have again received sponsorship from the UKIE section of the *Society for Industrial and Applied Mathematics* for three prizes for the best student talks which, if past meetings are anything to go by, will be a very hard-fought competition!

Outwith the scientific content, we hope you will also enjoy meeting new people and socialising with the other participants. We are indebted to the *City of Glasgow* for once again generously sponsoring a wine reception in the spectacular City Chambers building: this will take place on Tuesday evening. Another social highlight will be our conference dinner on Thursday evening, which will be held in the Trades Hall, another of Glasgow's impressive historic buildings.

Welcome to Glasgow, and enjoy the meeting!

Philip Knight
John Mackenzie
Alison Ramage

Conference Organising Committee

Information for participants

- **General**. There will be a registration desk in the foyer of the John Anderson building. The organisers can be contacted there during tea and coffee breaks.
- **Accommodation**. All rooms are in the Campus Village. Check-out time is 10:00 on day of departure. On Friday morning, luggage may be left in room JA327.
- **Meals**. Breakfast (Tue-Fri) is available from 07:30 until 09:00 in the Aroma Dining Room in the Lord Todd building. The times of lunches and dinners are as indicated in the conference programme. Dinner (Tue, Wed) will also be served in Aroma. Buffet lunches (Tue-Thu) will be served in the Urban Bean Java Cafe and (Fri) in the foyer outside JA325. Coffee and tea will be provided at the advertised times in the foyer outside JA325.
- **Lecture rooms**. These are all in the John Anderson building. The main auditorium (JA325) is down one floor from the main entrance, along with rooms JA314 and JA317. The additional rooms for parallel sessions are JA505, JA506 and JA507 (on level 5 of the John Anderson building).
- **Chairing sessions**. It is hoped that if you are listed as chairing a session, you will be willing to help in this way. Minisymposium organisers should organise chairpeople for their own sessions (including any contributed talks which follow) as appropriate. A break of 5 minutes has been allowed for moving between rooms. **Please keep speakers to the timetable!**
- **UKIE SIAM student prizes**. If you are chairing a session, please rate any student talks using the forms provided in the lecture rooms.
- **Book displays**. There will be books on display for the duration of the conference in room JA326. Room JA327 is also available for conference participants.
- **Reception**. A reception for all participants hosted by Glasgow City Council will be held in the City Chambers on Tuesday 23rd June from 20.00 to 21.00. Entry to the City Chambers is from George Square.
- **Conference dinner**. The conference dinner will be held in the Trades Hall of Glasgow on Thursday 29th June at 19:30 (for 20:00 dinner). The venue is located at 84 Glassford Street, Glasgow G1 1UH, which is 10 minutes walk from the conference venue. The guest speaker will be Professor Jeremy Levesley, University of Leicester.
- **Internet Access**. Complimentary WiFi is available throughout the campus from the 'WifiGuest' network, which should appear on the list of available networks on your portable device. This network uses the same authentication system as 'The Cloud' network found in public places across the UK. You can log in to 'WifiGuest' using existing 'The Cloud' credentials, or set up a new account which you can then use wherever 'The Cloud' is available. More comprehensive internet access is available to eduroam users. Computer terminals will be available from 09:00-17:00 Tuesday-Thursday in room JA512 of the John Anderson Building. If you require access to a fixed terminal, please contact the organisers to obtain a username/password.
- **Bar**. There is a bar in the Lord Todd building next to the Aroma Dining Room.
- **Sports facilities**. Delegates can access facilities at the University Sports Centre for a fee of £5 per visit.

Invited Speakers

Raymond Chan	City University of Hong Kong	rchan.sci@cityu.edu.hk
Paul Constantine	University of Colorado	paul.constantine@colorado.edu
Alistair Forbes	National Physical Laboratory	alistair.forbes@npl.co.uk
Vivette Girault	Sorbonne Université	girault@ann.jussieu.fr
Des Higham	University of Edinburgh	d.j.higham@ed.ac.uk
Natalia Kopteva	University of Limerick	natalia.kopteva@ul.ie
Gunilla Kreiss	Uppsala University	gunilla.kreiss@it.uu.se
Frances Kuo	University of NSW	f.kuo@unsw.edu.au
Ulrich Rüde	FAU Erlangen-Nürnberg	ulrich.ruede@fau.de
Carola-Bibiane Schönlieb	Univ of Cambridge	cbs31@cam.ac.uk
Holger Wendland	Universität Bayreuth	holger.wendland@uni-bayreuth.de
Margaret Wright	Courant Institute	mhw@cs.nyu.edu

Abstracts of Invited Talks

Flexible methodology for image segmentation

Raymond H. Chan *(City University of Hong Kong)*

In this talk, we introduce a SaT (Smoothing and Thresholding) method for multiphase segmentation of images corrupted with different degradations: noise, information loss and blur. At the first stage, a convex variant of the Mumford-Shah model is applied to obtain a smooth image. We show that the model has unique solution under different degradations. In the second stage, we apply clustering and thresholding techniques to find the segmentation. The number of phases is only required in the last stage, so users can modify it without the need of repeating the first stage again. The methodology can be applied to various kinds of segmentation problems, including color image segmentation, hyper-spectral image classification, and point cloud segmentation. Experiments demonstrate that our SaT method gives excellent results in terms of segmentation quality and CPU time in comparison with other state-of-the-art methods.

Joint work with: X.H. Cai (UCL) and T.Y. Zeng (CUHK)

The uncertainty contribution of numerical software

Alistair Forbes *(National Physical Laboratory)*

The National Physical Laboratory (NPL) is the UK's National Metrology Institute, responsible for defining and dissemination standard units of measurement, the metre, the kilogram, etc. Two key concepts in metrology are those of traceability and uncertainty. Measurements results need to be traceable back to relevant standard units though a chain of calibrations. Each measurement result must be accompanied by an uncertainty statement that summarises the factors that influence the accuracy of the measurement and quantifies their uncertainty contribution. Almost all measurement results involve computation, often significantly, and it is therefore necessary to assess the uncertainty contribution of the numerical software performing the computation. There are many ways in which numerical software can make a significant uncertainty contribution. There are issues relating to finite precision, numerical stability of algorithms and ill-conditioning. But there are also contributions related to model approximations and discretisation, algorithms that address a different computational aim, iterative schemes involving convergence tolerances and tuning parameters. Finally, the software implementing the algorithms might not have been coded correctly.

We can regard numerical software as fit for purpose if its uncertainty contribution is small compared to other uncertainty sources such as those associated with the input data. A dependable uncertainty quantification scheme also opens up the possibility of implementing simpler, quicker algorithms, perhaps in reduced precision, that still meet the requirements of the user in terms of accuracy.

In this talk, I will summarise some of the work we have undertaken at NPL relating to assessing the uncertainty contribution of numerical software with applications to the analysis of metrology data. I will also take the opportunity to discuss the redefinition of the SI units in terms of fundamental physical constants.

Discretization of some elastic models with implicit non linear constitutive relations

V. Girault *(Sorbonne Université)*

There are situations where linearization of classical Cauchy elasticity is not applicable. This occurs for example in the vicinity of crack tips and notches in a brittle material. In this case, the generalization of Cauchy elasticity proposed by K.R. Rajagopal, with implicit non linear constitutive relation, leads to a good description of the phenomenon. One of its advantages is that it permits the strain (i.e., the symmetric gradient of the displacement) to stay bounded, even when the stress is very large. This approach has not been much explored and is an interesting new line of research.

With this motivation, I shall discuss some finite element discretizations of a family of such models, where the strain is a non linear monotone function of the stress, but without coercivity. The problem consists of a system of two equations, the balance of linear momentum and the constitutive equation, in two unknowns, the displacement and the stress. The theoretical analysis is done by stabilizing the constitutive equation by the adjunction of a small stabilizing coercive term, and by proving convergence of the stabilized solution. Therefore, I shall discretize the stabilized form and prove convergence of its finite element solution. Also, as the resulting non linear system is large, I shall use an alternating-direction algorithm to

decouple the computation of the two equations, and present a numerical experiment. The monotone character of the operator plays a fundamental part in the numerical (and theoretical) analysis, as well as in the algorithm's convergence.

If there is time, I shall briefly present other applications to similar models.

This is joint work with A. Bonito, D. Guignard, K. Rajagopal, Department of Mathematics, TAMU, E. Süli, Mathematical Institute, University of Oxford

Walk this way

Desmond J. Higham *(University of Edinburgh)*

The notion of a walk around a graph is both natural and useful. The walker may follow any available route, with nodes and edges being revisited at any stage. However, in some settings, walks that *backtrack*—leave a node and then return to it on the next step—are best avoided. The concept of restricting attention to *nonbacktracking walks* has arisen, essentially independently, in a wide range of seemingly unrelated fields, including spectral graph theory, number theory, discrete mathematics, quantum chaos, random matrix theory, stochastic analysis, applied linear algebra and computer science. Nonbacktracking has more recently been considered in the context of matrix computation, where it has been shown to form the basis of effective algorithms in network science for community detection, centrality measurement and graph alignment. I will discuss some recent theoretical and practical advances that combine ideas from matrix polynomial theory, spectral analysis, mathematical physics and what Herbert S. Wilf called "generatingfunctionology." Topics will include

- localization effects,
- graph pruning for efficiency,
- alternative interpretations of nonbacktracking algorithms,
- a nonbacktracking version of Google's PageRank,
- extensions to nontriangulating, nonsquaring and, generally, non-k-cycling walks.

Results will be illustrated on some large scale networks arising in the social sciences. My nonbacktracking collaborators are Francesca Arrigo (Strathclyde), Peter Grindrod (Oxford) and Vanni Noferini (Aalto).

A posteriori error estimation on anisotropic meshes

Natalia Kopteva *(University of Limerick)*

Solutions of partial differential equations frequently exhibit corner singularities and/or sharp boundary and interior layers. To obtain reliable numerical approximations of such solutions in an efficient way, one may want to use meshes that are adapted to solution singularities. Such meshes can be constructed using a priori information on the exact solutions, however it is rarely available in real-life applications. Therefore the automated mesh construction by adaptive techniques is frequently employed. This approach requires no initial asymptotic understanding of the nature of the solutions and the solution singularity locations. Note that reliable adaptive algorithms are based on a posteriori error estimates, i.e. estimates of the error in terms of values obtained in the computation process: computed solution and current mesh. Such a posteriori error estimates for elliptic partial differential equations will be the subject of this talk.

Our main goal in this talk is to present residual-type a posteriori error estimates in the maximum norm, as well as in the energy norm, on anisotropic meshes, i.e. we allow mesh elements to have extremely high aspect ratios. The error constants in these estimates are independent of the diameters and the aspect ratios of mesh elements. Note also that, in contrast to some a posteriori error estimates on anisotropic meshes in the literature, our error constants do not involve so-called matching functions (that depend on the unknown error and, in general, may be as large as mesh aspect ratios).

To deal with anisotropic elements, a number of technical issues have been addressed. For example, an inspection of standard proofs for shape-regular meshes reveals that one obstacle in extending them to anisotropic meshes lies in the application of a scaled traced theorem when estimating the jump residual terms (this causes the mesh aspect ratios to appear in the estimator). For the estimation in the energy norm, a special quasi-interpolation operator is constructed on anisotropic meshes, which may be of independent interest.

In the second part of the talk, the focus will shift to the efficiency of energy-norm error estimators on anisotropic meshes.

A simple numerical example will be given that clearly demonstrates that the standard bubble function approach does not yield sharp lower error bounds. To

be more precise, short-edge jump residual terms in such bounds are not sharp. To remedy this, we shall present a new upper bound for the short-edge jump residual terms of the estimator.

We shall also touch on that certain perceptions need to be adjusted for the case of anisotropic meshes. In particular, it is not always the case that the computed-solution error in the maximum norm is closely related to the corresponding interpolation error.

Stability and accuracy for initial-boundary value problems revisited

Gunilla Kreiss *(Uppsala University)*

Stability and accuracy for a numerical method approximating an initial boundary value problem are inherently linked together. Stability means that perturbations of data have a bounded effect on the discrete solution, and is usually characterized by a precise estimate, which links the perturbations of data to perturbations of the solution. Such an estimate can be directly used to quantify the accuracy of the method. A very convenient and common way to investigate stability, and hence accuracy, is to use the energy method. If this approach fails one may instead attempt to get results by Laplace transforming in time. Such analysis is usually more involved, but sharper results may follow. In this talk we will discuss the two basic approaches, both in continuous and discrete settings. We will also show two examples where, even though the energy method is applicable, it is rewarding to consider the problem in the Laplace domain. In the first case we get sharper accuracy results, and in the second case we get sharper temporal bounds.

High dimensional integration and approximation: the Quasi-Monte Carlo (QMC) way

Frances Y. Kuo *(UNSW Sydney)*

High dimensional computation – that is, numerical computation in which there are very many or even infinitely many continuous variables – is a new frontier in scientific computing. Often the high dimensionality comes from uncertainty or randomness in the model or data (e.g., in groundwater flow it can arise from modeling the permeability field that is rapidly varying and uncertain). High dimensional problems present major challenges to computational resources, and require serious theoretical numerical analysis in devising new and effective methods.

This talk will begin with a contemporary review of Quasi-Monte Carlo (QMC) methods, which offer tailored point constructions for solving high dimensional integration and approximation problems by sampling. By exploiting the smoothness properties of the underlying mathematical functions, QMC methods are proven to achieve higher order convergence rates, beating standard Monte Carlo sampling. Moreover, QMC error bounds can be independent of the dimension under appropriate theoretical function space settings.

In recent years the modern QMC theory has been successfully applied to a number of applications in uncertainty quantification. This talk will showcase some ongoing works where we take QMC methods to new territories including neutron transport as a high dimensional PDE eigenvalue problem, high frequency wave scattering in random media, optimal control constrained by PDE with random coefficients, and a revolutionary approach to model random fields using periodic random variables. The common and essential theme among these collaborations with various international teams is that we are not just applying an off-the-shelf method; rather, we provide rigorous error and cost analysis to design QMC methods tailored to the features of the underlying mathematical functions in these applications.

Parallel multigrid for systems with a trillion unknowns and beyond

Ulrich Rüde *(Friedrich-Alexander-Universität Erlangen-Nürnberg and CERFACS, Toulouse)*

For discretized elliptic PDE, multigrid methods can provide asymptotically optimal solvers so that the number of operations grows only proportional to the number of unknowns. Such algorithms can then be used to develop *scalable* parallel solvers, i.e. concurrent algorithms that can solve twice the size of a system in the same time when it is given twice the number of processors. This is known as *weak scaling*. We will also discuss *strong scaling*, i.e. solving the same system in half the time when given twice as many processors, since this is significantly harder to achieve. However, multigrid methods do not only exhibit favorable asymptotic behavior. They are also efficient in an absolute sense, since they often beat other methods in terms of total number of flops and time spent for computing a sufficiently accurate solution for a given system size. To reach fast compute times, it is also essential that the methods are implemented with care and being aware of the architecture of modern CPUs, exploiting multicore architectures, instruction

level parallelism, and vector units. We will also discuss how parallel multigrid methods can be excellent companions for Multilevel Monte Carlo methods, where many problems of varying size must be solved.

Finally, we will present techniques that are necessary to realize multigrid methods that can exploit supercomputers with up to a million parallel processes. The largest example problem will be a Stokes system that is discretized with finite elements and that results in an indefinite linear system with more than 10^{13} unknowns. This corresponds to a solution vector that requires 80 TByte of memory. For solving these problems of this size, matrix free techniques have been developed that do not need to store the stiffness or mass matrices. The talk will conclude with examples from Earth mantle convection that are grand challenge problems in geophysics.

This is joint work with Dominik Bartuschat, Simon Bauer, Hans-Peter Bunge, Daniel Drzisga, Björn Gmeiner, Markus Huber, Lorenz John, Nils Kohl, Philippe Leleux, Marcus Mohr, Holger Stengel, Dominik Thönnes, Christian Waluga, Jens Weismüller, and Barbara Wohlmuth in the project TerraNeo within the DFG priority program *Software for Exascale Computing*.

Variational models and partial differential equations for mathematical imaging

Carola-Bibiane Schönlieb *(Cambridge University)*

Images are a rich source of beautiful mathematical formalism and analysis. Associated mathematical problems arise in functional and non-smooth analysis, the theory and numerical analysis of partial differential equations, harmonic, stochastic and statistical analysis, and optimisation. Starting with a discussion on the intrinsic structure of images and their mathematical representation, in this talk we will learn about variational models for image analysis and their connection to partial differential equations, and go all the way to the challenges of their mathematical analysis as well as the hurdles for solving these - typically non-smooth - models computationally. The talk is furnished with applications of the introduced models to digital art restoration, cancer imaging and forest conservation. If time allows I might say a few words about using machine learning to learn better variational models for imaging.

Kernel-based reconstructions for uncertainty quantification

Holger Wendland *(University of Bayreuth)*

In uncertainty quantification, an unknown quantity has to be reconstructed which depends typically on the solution of a partial differential equation. This partial differential equation itself depends on parameters, some of them are deterministic and some are random. To approximate the unknown quantity one thus has to solve the partial differential equation (usually numerically) for several instances of the parameters and then reconstruct the quantity from these simulations. As the number of parameters may be large, this becomes a high-dimensional reconstruction problem. In this talk, I will address the topic of reconstructing such unknown quantities using kernel-based reconstruction methods on sparse grids. I will give an introduction to the topic. After that, I will explain different reconstruction processes using kernel-based methods such as support vector machines and multiscale radial basis functions and the so-called Smolyak algorithm. I will discuss techniques for deriving error estimates and show convergence results. If time permits, I will also give numerical examples.

This talk is based upon joint work with R. Kempf (Bayreuth) and C. Rieger (Aachen/Bonn).

Optimization and machine learning: what are the connections?

Margaret Wright *(Courant Institute)*

To answer the question in the title, one might turn to the hundreds of millions of relevant websites produced by a favorite search engine. Unfortunately, the answers are highly inconsistent, nor are they necessarily trustworthy. (Some offer to charge a large fee so that you can become a data scientist in 6 weeks.) There are experts who claim that there is very little difference between machine learning and optimization; others assert that machine learning is merely a consumer of optimization; and still others state emphatically that optimization is a strict subset of machine learning. The speaker (an optimizer) will discuss the increasing number of interactions between machine learning and "classical" optimization, emphasizing recent and continuing substantive changes.

Minisymposia abstracts

Minisymposium M1

Fractional-derivative problems
Organisers
Martin Stynes and Yubin Yan

Error analysis of the Grünwald–Letnikov scheme for fractional initial-value problems with weakly singular solutions

Hu Chen & Finbarr Holland & Martin Stynes *(Beijing Computational Science Research Center)*

A convergence analysis is given for the Grünwald-Letnikov discretisation of a Riemann-Liouville fractional initial-value problem on a uniform mesh $t_m = m\tau$ with $m = 0, 1, \ldots, M$. For given smooth data, the unknown solution of the problem will usually have a weak singularity at the initial time $t = 0$. Our analysis is the first to prove a convergence result for this method while assuming such non-smooth behaviour in the unknown solution. In part our study imitates previous analyses of the L1 discretisation of such problems, but the introduction of some additional ideas enables exact formulas for the stability multipliers in the Grünwald-Letnikov analysis to be obtained (the earlier L1 analyses yielded only estimates of their stability multipliers). Armed with this information, it is shown that the solution computed by the Grünwald-Letnikov scheme is $O(\tau t_m^{\alpha-1})$ at each mesh point t_m; hence the scheme is globally only $O(\tau^\alpha)$ accurate, but it is $O(\tau)$ accurate for mesh points t_m that are bounded away from $t = 0$. Numerical results for a test example show that these theoretical results are sharp.

Mittag-Leffler Euler integrator for solving semilinear time fractional stochastic partial differential equations driven by fractionally integrated additive noise

Ye Hu *(LvLiang University, University of Chester)* & Changpin Li *(Shanghai University)* & Yubin Yan *(University of Chester)*

In this paper, we introduce a spectral element method in space and an exponential time discretization scheme in time to approximate a semilinear stochastic partial differential equation driven by fractionally integrated additive noise, which is the continuation of the work in Jin et al. where the numerical methods for the linear stochastic partial differential equations driven by fractionally integrated additive noise are considered. The solution of the problem involves three different Mittag-Leffler functions and the regularities of the solution are studied based on the properties of the Mittag-Leffler functions. The optimal convergence error estimates are obtained which precisely reflect how the fractional orders α, γ affect the convergence rates. Numerical experiments show that the numerical results are consistent with the theoretical findings.

Sharp spatial H^1-norm analysis of a finite element method for a time-fractional initial-boundary value problem

Chaobao Huang & Martin Stynes *(Beijing Computational Science Research Center)*

A time-fractional initial-boundary value problem $D_t^\alpha u - \Delta u = f$, where D_t^α is a Caputo fractional derivative of order $\alpha \in (0, 1)$, is considered on the space-time domain $\Omega \times [0, T]$, where $\Omega \subset \mathbb{R}^d$ $(d \geq 1)$ is a bounded Lipschitz domain. Typical solutions $u(x, t)$ of such problems have components that behave like a multiple of t^α as $t \to 0^+$, so the integer-order temporal derivatives of u blow up at $t = 0$. The numerical method of the paper uses a standard finite element method in space on a quasiuniform mesh and considers both the L1 discretisation and Alikhanov's L2-1_σ discretisation of the Caputo derivative on suitably graded temporal meshes. *Optimal error bounds in* $H^1(\Omega)$ are proved; no previous analysis of a discretisation of this problem using finite elements in space has established such a bound. Furthermore, the optimal grading of the temporal mesh can be deduced from our analysis. Numerical experiments show that our error bounds are sharp.

References

[1] Chaobao Huang and Martin Stynes, Optimal spatial H^1-norm analysis of a finite element method for a time-fractional diffusion equation. (Submitted for publication).

Mittag-Leffler Euler integrator for a stochastic fractional order equation with additive noise

Mihály Kovács & Stig Larsson & Fardin Saedpanah *(PPKE ITK)*

Motivated by fractional derivative models in viscoelasticity, a class of semilinear stochastic Volterra integro-

differential equations, and their deterministic counterpart, are considered. A generalized exponential Euler method, named here as the Mittag-Leffler Euler integrator, is used for the temporal discretization while the spatial discretization is performed by the spectral Galerkin method. The temporal rate of strong convergence is found to be (almost) twice compared to when the backward Euler method is used together with a convolution quadrature for time discretization.

Subdiffusion with a time-dependent coefficient: analysis and numerical solution

Bangti Jin & **Buyang Li** & Zhi Zhou *(The Polytechnic University of Hong Kong)*

In this work, a complete error analysis is presented for fully discrete solutions of the subdiffusion equation with a time-dependent diffusion coefficient, obtained by the Galerkin finite element method with conforming piecewise linear finite elements in space and backward Euler convolution quadrature in time. The regularity of the solutions of the subdiffusion model is proved for both nonsmooth initial data and incompatible source term. Optimal-order convergence of the numerical solutions is established using the proven solution regularity and a novel perturbation argument via freezing the diffusion coefficient at a fixed time. The analysis is supported by numerical experiments.

Error analysis for a fractional-derivative parabolic problem on quasi-graded meshes using barrier functions

Natalia Kopteva *(University of Limerick)* & **Xiangyun Meng** *(Beijing Computational Science Research Center)*

An initial-boundary value problem with a Caputo time derivative of fractional order $\alpha \in (0, 1)$ is considered, solutions of which typically exhibit a singular behaviour at an initial time. For this problem, we give a simple and general numerical-stability analysis using barrier functions, which yields sharp pointwise-in-time error bounds on quasi-graded temporal meshes with arbitrary degree of grading. L1-type and Alikhanov-type discretization in time are considered. In particular, those results imply that milder (compared to the optimal) grading yields optimal convergence rates in positive time. Semi-discretizations in time and full discretizations are addressed. The theoretical findings are illustrated by numerical experiments.

Blow-up in time-fractional initial-boundary value problems

Martin Stynes & Hu Chen *(Beijing Computational Science Research Center)*

Time-fractional initial-boundary value problems of the form $D_t^\alpha u - p\partial^2 u/\partial x^2 + cu = f$ are considered, where $D_t^\alpha u$ is a Caputo fractional derivative of order $\alpha \in (0, 1)$. As $\alpha \to 1^-$, we prove that the solution u converges, uniformly on the space-time domain, to the solution of the classical parabolic initial-boundary value problem where $D_t^\alpha u$ is replaced by $\partial u/\partial t$. Nevertheless, most of the rigorous analyses of numerical methods for this time-fractional problem have error bounds that blow up as $\alpha \to 1^-$, as we demonstrate. We show that in some cases these analyses can be modified to obtain robust error bounds that do not blow up as $\alpha \to 1^-$.

Convergence of $L1$-Galerkin FEMs for nonlinear time-fractional Schrödinger equations

Jilu Wang (Beijing Computational Science Research Center) & Dongfang Li *(Huazhong University of Science and Technology)* & Jiwei Zhang *(Wuhan University)*

In this work, a linearized $L1$-Galerkin finite element method is proposed to solve the multi-dimensional nonlinear time-fractional Schrödinger equation. In terms of a temporal-spatial error splitting argument, we prove that the finite element approximations in L^2-norm and L^∞-norm are bounded without any time stepsize conditions. More importantly, by using a discrete fractional Gronwall type inequality, optimal error estimates of the numerical schemes are obtained unconditionally, while the classical analysis for multi-dimensional nonlinear fractional problems always required certain time-step restrictions dependent on the spatial mesh size. Numerical examples are given to illustrate our theoretical results.

Laplace transform method for solving the fractional cable equation with nonsmooth data

Yanyong Wang & **Yubin Yan** & Amiya K. Pani *(Lvliang University, China; University of Chester; Indian Institute of Technology Bombay)*

We introduce two time discretization schemes for solving the time fractional cable equation. The time derivative is approximated by using the backward Euler method and the second order backward difference formula, respectively. The Riemann-Liouville fractional derivatives are approximated by using the backward Euler convolution quadrature method and the second order backward difference convolution quadrature method, respectively. The nonsmooth data error estimates with the convergence orders $O(k)$ and $O(k^2)$ are proved in detail. Instead of using the discretized operational calculus approach to prove the nonsmooth data error estimates of the time discretization schemes for solving time fractional cable equation as used in literature, we directly bound the kernel function in the resolvent and obtain the nonsmooth data error estimates. To the best of our knowledge, this is the first work to consider the nonsmooth data error estimates for solving the time fractional cable equation by directly bounding the kernel function in the resolvent. This argument may be applied to consider the nonsmooth data error estimates for solving the time fractional cable equation where the Riemann-Liouville fractional derivatives are approximated by using other schemes, for example, L1 scheme.

Minisymposium M2

Recent Advances in the Numerical Solution of Inverse Problems in Imaging

Organiser
Silvia Gazzola

Adaptive regularization parameter choice rules for large-scale problems

Malena Sabaté Landman & Silvia Gazzola *(University of Bath)*

This talk introduces a new class of adaptive regularization parameter choice strategies that can be efficiently applied when regularizing large-scale linear inverse problems using a combination of projection onto Krylov subspaces and Tikhonov regularization, and that can be regarded as special instances of bilevel optimization methods. The links between Gauss quadrature and Golub-Kahan bidiagonalization are exploited to prove convergence results for some of the considered approaches, and numerical tests are shown to give insight.

Learning a sampling pattern for MRI

Ferdia Sherry & Matthias J. Ehrhardt & Carola-Bibiane Schönlieb *(University of Cambridge)*

The discovery of the theory of compressed sensing brought the realisation that many inverse problems can be solved satisfactorily even when measurements are incomplete. This is particularly interesting in an imaging modality such as magnetic resonance imaging (MRI), where long acquisition times are problematic. The measurements taken in MRI are often modelled as samples of the Fourier transform of the signal to be recovered. In this work, we consider the problem of learning a sparse sampling pattern that can be used to obtain high quality reconstructions for images that are similar to the images used in training, while only having to sample along the learned sampling pattern. The approach that is taken is that of bilevel learning: a minimisation problem is solved, which has constraints that are formulated as the solution to variational regularisation problems (the problems that are solved to reconstruct a signal from measurements). The framework that we set up for learning a sampling pattern allows a sampling pattern of scattered points or of Cartesian lines in k-space to be learned and can easily be extended to other parametrisations of the sampling pattern.

Some augmented methods for treatment of ill-posed problems

Kirk M Soodhalter *(Trinity College Dublin)*

In this talk, we discuss some strategies for treating *discretize* ill-posed problems of the form,

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad \text{with} \quad \mathbf{A} \in \mathbb{R}^{n\times n} \quad \text{and} \quad \mathbf{b} \in \mathbb{R}^n,$$

which inherits a discrete version of ill conditioning from the continuum problem. We propose methods built on top of the Arnoldi-Tikhonov method of Lewis and Reichel, whereby one builds the Krylov subspace

$$\mathcal{K}_j(\mathbf{A}, \mathbf{w}) = \text{span}\left\{\mathbf{w}, \mathbf{A}\mathbf{w}, \mathbf{A}^2\mathbf{w}, \ldots, \mathbf{A}^{j-1}\mathbf{w}\right\}$$

where

$$\mathbf{w} \in \{\mathbf{b}, \mathbf{A}\mathbf{b}\}$$

and solves the discretized Tikhonov minimization problem projected onto that subspace. We propose to extend this strategy to setting of augmented Krylov subspace methods. Thus we project onto a sum of subspaces of the form $\mathcal{U} + \mathcal{K}_j$ where $\mathcal{U}$ is a fixed subspace

and $\mathcal{K}_j$ is a Krylov subspace. There are a number of approaches one can take to combine these strategies. We discuss them and demonstrate how they work with some example problems.

Variational models for image registration problems

Anis Theljani & Ke Chen *(University of Liverpool)*

Image registration is a core problem in Medical Imaging. It aims to align two or more images so that information can be compared and/or fused to highlight differences and complement information. Mathematically the task is to find a suitable transform that maps patterns and shapes of one image to those of another. The different imaging modalities lead to images with different properties and often complementary information which makes the task particularly useful. Although most methods work very well on mono-modal images, there is a high demand for improved models for registration of multi-modality images. In this talk, we present some new works on improved variational models for multi-modal images. Experiments and comparisons with competitive models show the good performance of the proposed models in generating an accurate diffeomorphic transformation. We also show how these models can be in used in a deep learning framework to order to get an accurate solution in fast speed for real time applications

Blind image deconvolution using a non-separable point spread function

Joab Winkler *(University of Sheffield)*

This paper considers the problem of the removal of blur from an image that is degraded by a non-separable point spread function (PSF), when information on the PSF is not known. The non-separable nature of the PSF implies that two blurred images, taken by the same system such that the PSF can be assumed to be the same for both images, are required to determine the PSF. The most difficult part of the computation is the determination of the size of the PSF because this problem reduces to the determination of the rank of two matrices (one matrix for the horizontal component of the PSF and one matrix for the vertical component of the PSF). It is shown that this computation requires the determination of the greatest common divisor of two polynomials, after they have been transformed to the Fourier domain. The Sylvester resultant matrix and its subresultant matrices are used for this computation. A structure-preserving matrix method is used to perform each deconvolution, and thereby compute deblurred forms of the given blurred images because this method preserves the Tœplitz structure of the coefficient matrix in the linear algebraic equation.

The presentation includes examples that demonstrate the method.

Expectation propagation for Poisson data

Chen Zhang & Simon Arridge & Bangti Jin *(University College London)*

The Poisson distribution arises naturally when dealing with data involving counts, and it has found many applications in inverse problems and imaging. In this work, we develop an approximate Bayesian inference technique based on expectation propagation for approximating the posterior distribution formed from the Poisson likelihood function and a Laplace type prior distribution, e.g., the anisotropic total variation prior. The approach iteratively yields a Gaussian approximation, and at each iteration, it updates the Gaussian approximation to one factor of the posterior distribution by moment matching. We derive explicit update formulas in terms of one-dimensional integrals, and also discuss stable and efficient quadrature rules for evaluating these integrals. The method is showcased on two-dimensional PET images.

Minisymposium M3

Trends in numerical methods for partial differential equations

Organisers

A M Portillo, M J Moreta and N Reguera

Transparency: boundary conditions that imitate the Cauchy problem for differential and finite-difference equations and systems

Vladimir A. Gordin & Alexander A. Shemendyuk *(National Research University "Higher School of Economics") & Hydrometeorological Centre of Russia)*

In many practical problems of mathematical physics the necessary complete set of physically adequate boundary conditions is absent. For example, in the problem of weather forecasting for a limited area V,

we can set Dirichlet boundary conditions at the border ∂V, where the right hand side of the conditions is taken from a larger scale (global) forecasting model. We can consider the difference between meteorological fields in these two numerical models. The dynamics of this difference show that waves coming out of the computational area are reflected from the boundary ∂V. It is not a physical effect, which worsens the regional model's forecasts.

We construct the special kind of boundary conditions (BC) that do not reflect outgoing waves both for differential and finite-difference problems. The mixed initial-boundary problem under such BC give the solution, which is identical to the solution of the corresponding Cauchy problem, i.e. the BC imitate the Cauchy problem.

The ICP BC may be constructed for a wide class of PDE. As a rule, such BC are not local: $Au = B(d_x u)$ for differential problems and $Au(x = 0) = Bu(x = h)$ for finite-difference ones. Here A, B are convolution operators (continuous or discrete) with respect to the variables that tangent to ∂V, and the variable x is directed along the normal to it. The coefficients of the convolution may be interpreted as the coefficients of the Pade – Hermite rational approximation. The bigger number of the ICP BC are necessary for higher order (with respect to x) of PDE (e.g. for the equation of rod transverse vibrations). We use here a more sophisticated kind of rational approximation. We confirm the ICP-property by numerical experiments for various difference schemes for various PDE.

The work was prepared within the framework of the Academic Fund Program at the National Research University Higher School of Economics (HSE) in 2018 – 2019 (grant No 18-05-0011) and by the Russian Academic Excellence Project 5-100.

References

[1] V.A.Gordin. "About Mixed Boundary Problem that Imitates Cauchy's Problem," Uspekhi Matematicheskikh Nauk, vol. 33, no. 5, pp. 181-182, 1978.

[2] V.A.Gordin. Mathematical Problems and Methods in Hydrodynamical Weather Forecasting. Gordon & Breach, 2000.

[3] V.A.Gordin. Mathematics, Computer, Weather Forecast and Other Scenaria of Mathematical Physics, Moscow: FIZMATLIT, 2010, 2013 (in Russian).

[4] V.A. Gordin, A.A.Shemendyuk. "Transparent" Boundary Conditions for the Equation of Rod Transverse Vibrations. Preprint https://arxiv.org/submit/2447606/view 2018.

A Schwarz method for a Rayleigh-Bénard problem

Henar Herrero & Francisco Pla *(Universidad de Castilla-La Mancha)*

A study of a Schwarz domain decomposition numerical method for a Rayleigh-Bénard convection problem is presented. The model equations are incompressible Navier-Stokes coupled with a heat equation under Boussinesq approximation. The problem is defined in a rectangular domain. The nonlinear stationary problem is dealt with by a Newton method. Each step in the Newton method is solved with a Schwarz domain decomposition method with the domain partitioned into two subdomains with appropriate interface conditions. The convergence properties of a linear case are studied theoretically in a simplified domain including two artificial parameters in the equations. The numerical resolution of the problem confirms the theoretical results. The convergence rate is less than one when overlap is considered. Convergence is optimal for some values of the parameters.

Solving efficiently time dependent singularly perturbed convection-diffusion systems

J.C. Jorge & C. Clavero *(Public University of Navarra)*

In this talk we deal with the efficient numerical resolution of one dimensional parabolic singularly perturbed systems of convection-diffusion-reaction type which are coupled by the reaction terms. We assume that the convection matrix is diagonal while we permit that the diffusion parameters can be different at each equation and even they can have different orders of magnitude. The numerical method used to solve the continuous problem combines a standard upwind finite difference scheme, defined on a special mesh of Shishkin type, to discretize in space and the fractional implicit Euler method together with an appropriate splitting by components to integrate in time. It is proven that the resulting numerical algorithm is uniformly convergent of first order in time and of almost first order in space. Besides, the splitting-by-components technique used here makes that only small tridiagonal linear systems must be solved to advance in time; because of this, the proposed algorithm is more efficient than other classical implicit methods used to solve the same type of problems. Some nu-

merical experiments are shown which corroborate in practice both the uniform convergence and the efficiency of the algorithm.

Exponential quadrature rules without order reduction for integrating linear initial boundary value problems

M. J. Moreta & B. Cano *(Universidad Complutense de Madrid)*

In this talk a technique is suggested to completely avoid the order reduction that it is observed when exponential quadrature rules are used to integrate linear initial boundary value problems as

$$\begin{aligned} u'(t) &= Au(t) + f(t), \quad 0 \le t \le T, \\ u(0) &= u_0, \\ \partial u(t) &= g(t), \end{aligned}$$

When the s nodes of the quadrature rule are chosen in a suitable way, we can obtain order $2s$ in time, while with the classical approach the maximum order that is reached is s.

A thorough error analysis is given for both the classical approach of integrating the problem firstly in space and then in time and of doing it in the reverse order in a suitable manner in order to avoid the order reduction. Time-dependent boundary conditions are considered with both approaches and full discretization formulas are given to implement the methods once the quadrature nodes have been chosen for the time integration and a particular (although very general) scheme is selected for the space discretization. Numerical experiments are shown which corroborate that, for example, with the suggested technique, order $2s$ is obtained when choosing the s nodes of the Gaussian quadrature rule.

References

[1] B. Cano, M. J. Moreta, *Exponential quadrature rules without order reduction for integrating linear initial boundary value problems*, SIAM J. NUMER. ANAL. Vol. 56, No. 3 (2018), pp. 1187–1209.

This work was funded by Ministerio de Economía y Competitividad, Junta de Castilla y León and FEDER through projects MTM 2015- 66837-P and VA024P17.

High-order full discretization for two-dimensional coupled seismic wave equations

A.M. Portillo *(University of Valladolid)*

Two-dimensional coupled seismic waves on a rectangular domain with initial–periodic boundary conditions are considered. Second order spatial derivatives, in the x direction, in the y direction and mixed derivative, are approximated by second and fourth order finite differences on a uniform grid. If the initial conditions have compact support contained in an open set included in the domain, the discrete energies considered are approximations of the continuous energy of second and fourth order respectively. The ordinary second order in time system obtained is transformed into a first order in time system in the displacement and velocity vectors. For the time integration of this system, second order and fourth order exponential splitting methods, which are geometric integrators, are proposed. These explicit splitting methods are not unconditionally stable and the stability condition for time step and space step ratio is deduced. In the numerical experiments are displayed energy errors for the second and fourth order finite differences at time $t = 0$ as well as the good behavior in the long time integration of the splitting methods.

Avoiding order reduction when integrating nonlinear problems with Strang splitting

N. Reguera & I. Alonso-Mallo & B. Cano *(Universidad de Burgos)*

In this work we focus on the problem of avoiding order reduction when using Strang splitting to integrate nonlinear initial boundary value problems with time-dependent boundary conditions. In the literature there are several references dealing with this problem [1], [3].

We are going to present the technique proposed in [1], along with the results obtained for the local and global errors. On the other hand, we will show a numerical comparison in terms of computational efficiency [2] between the techniques [1] and [3]. We will see that it is important to consider an exponential method, that avoids also order reduction, for the integration of the linear nonhomogeneous and stiff part in [3] so that it is similar in efficiency to the technique proposed in [1].

References

[1] I. Alonso-Mallo, B. Cano and N. Reguera, *Avoiding order reduction when integrating reaction-diffusion boundary value problems with exponential splitting methods*, J. Computational and Applied Math-

ematics **357** (2019) 228-250.

[2] I. Alonso-Mallo, B. Cano and N. Reguera, *Comparison of efficiency among different techniques to avoid order reduction with Strang splitting*, submitted for publication.

[3] L. Einkemmer and A. Ostermann, *Overcoming order reduction in diffusion-reaction splitting. Part 2: Oblique boundary conditions*, SIAM J. Sci. Comput. **38** (2016) A3741-A3757.

Minisymposium M4

Preconditioning and iterative methods for differential equations
Organisers
John Pearson and Jennifer Pestana

On least-squares commutator (LSC) preconditioning for incompressible two-phase flow

Niall Bootland & Andy Wathen & Chris Kees *(University of Strathclyde)*

Modelling two-phase incompressible flow gives rise to variable coefficient Stokes or Navier–Stokes equations that can be challenging to solve computationally. In particular, large jumps in the coefficients between the two phases can cause solver performance to degrade significantly. The least-square commutator (LSC) method is a largely automatic way to build a block preconditioner for incompressible flow problems. However, for two-phase flow we will see that weighting with the variable coefficients must be included in order to gain a more robust solver. Further, it is known that a weighting is also required on Dirichlet boundaries if one is to obtain mesh-independence. We show empirically that, unlike in the smooth coefficient case, for two-phase flow there is an interplay between such weightings which can be detrimental to convergence. In this talk we introduce and explore weighted LSC preconditioners and discuss practical aspects related to two-phase flow.

On Preconditioners and higher-order time discretizations for PDE-constrained optimization problems

Santolo Leveque & John Pearson *(University of Edinburgh)*

PDE-constrained optimization problems form a fertile research area, for which preconditioned iterative methods may be usefully applied. In fact, Krylov methods for symmetric indefinite matrices are often suitable for the saddle point systems arising from the discretization of such problems. However, the performance of these methods is related to the spectral properties of the matrices involved, thus advanced techniques for finding a robust preconditioner are required to accelerate the convergence of the iterative scheme.

In this talk, we consider preconditioned iterative methods for time-dependent PDE-constrained optimization problems (commencing with the optimal control of the heat equation), coupled with suitable discretizations in the time variable. The state-of-the-art in this field, whether finite element or finite difference discretizations are used for the spatial variable, is to apply a backward Euler method in time, as this has been found to work well with existing preconditioners. However, this discretization strategy results in methods that are of second order in space and only first order in time, thus in order to obtain an accurate solution it is reasonable to choose $\Delta t = \mathcal{O}(h^2)$, where Δt is the time step and h is the mesh-size in space. This results in a matrix system of huge dimension, and hence a very high CPU time is required for its solution.

In this talk we present a new method for solving certain time-dependent PDE-constrained optimization problems, that are of second order both in time and space. This involves a Crank–Nicolson discretization of the time derivative within the PDE, and a new preconditioner for the resulting saddle point system (which now has a more complex structure). Further, we show through numerical experiments that this approach can obtain a more accurate solution (in terms of discretization error) than the original (backward Euler) method, in less CPU time.

Preconditioners for boundary control of elliptic PDE

Daniel Loghin *(University of Birmingham)*

We are interested in developing efficient preconditioners for the solution of large sparse linear systems with coefficient matrix having the following block form

$$K = \begin{matrix} M & O & A^T \\ O & \alpha Q & -B^T \\ A & -B & O \end{matrix},$$

where $M, A \in \mathbb{R}^{n\times n}, Q \in \mathbb{R}^{m\times m}, B \in \mathbb{R}^{n\times m}$. The matrix K arises in the discrete form of the weak formulation.

Find $(y^*, u^*, p^*) \in \mathcal{Y} \times \mathcal{U} \times \mathcal{Y}$ such that $\forall (z, v, r) \in \mathcal{Y} \times \mathcal{U} \times \mathcal{Y}$

$$\begin{cases} m(y^*, z) + a(z, p^*) & = m(y_d, z), \\ \alpha q(u^*, v) - b(v, p^*) & = 0, \\ a(y^*, r) - b(u^*, r) & = \ell_f(r), \end{cases}$$

where (y^*, u^*, p^*) represents the global minimizer of a linear-quadratic programming problem expressed in weak form as follows

$$\begin{cases} \text{minimize} & \frac{1}{2} m(y - y_d, y - y_d) + \frac{\alpha}{2} q(u, u) \\ \text{subject to} & a(y, v) - b(u, v) = \ell_f(v) \quad \forall v \in \mathcal{Y}. \end{cases}$$

Solving linear systems with coefficient matrix K is often challenging: while sparse, the matrix can be large in practice and thus the approach often considered is iterative or hybrid (iterative-direct) accelerated by a good preconditioner. However, all the blocks, except for Q can be rank-deficient, depending on the underlying problem. This poses additional difficulties when designing block preconditioners based on the matrix structure.

In this talk, we consider a class of boundary preconditioners for solving the block systems arising from the finite element discretization of linear-quadratic elliptic control problems with Neumann boundary control. We show that an optimal approach can be obtained by solving an interior problem and a coupled boundary problem involving certain Steklov-Poincaré operators. The latter boundary problem can be preconditioned by a block matrix with rational entries involving a certain representation of a discrete fractional Sobolev norm associated with the Neumann part of the boundary.

Multigrid preconditioners for anisotropic space-fractional diffusion equations

Mariarosa Mazza[†] & Marco Donatelli[†] & Rolf Krause[‡] & Ken Trotti[†] *([†] University of Insubria, [‡] University of Italian Switzerland)*

We focus on a two-dimensional time-space diffusion equation with fractional derivatives in space. The use of Crank-Nicolson in time and finite differences in space leads to dense Toeplitz-like linear systems. Multigrid strategies that exploit such structure are particularly effective when the fractional orders are both close to 2. We seek to investigate how structure-based multigrid approaches can be efficiently extended to the case where only one of the two fractional orders is close to 2, i.e., when the fractional equation shows an intrinsic anisotropy. Precisely, we build a multigrid (block-banded-banded-block) preconditioner that has relaxed Jacobi as smoother and whose grid transfer operator is obtained with a semicoarsening technique. A further improvement in the robustness of the proposed multigrid method is attained using the V-cycle with semicoarsening as smoother inside an outer full-coarsening. Several numerical results confirm that the resulting multigrid preconditioner is computationally effective and outperforms current state of the art techniques.

Fast interior point solvers and preconditioning for PDE-constrained optimization

John Pearson *(University of Edinburgh)* & Jacek Gondzio & Margherita Porcelli & Martin Stoll

In this talk we consider the effective numerical solution of PDE-constrained optimization problems with additional box constraints on the state and control variables. Upon discretization, these may give rise to problems of quadratic or nonlinear programming form: a sensible solution strategy is to apply an interior point method, provided one can solve the large and structured matrix systems that arise at each Newton step. We therefore consider fast and robust preconditioned iterative methods for these systems, examining two cases: (i) where L^2 norms measure the misfit between state and desired state, as well as the control; (ii) with an additional L^1 norm term promoting sparsity in the control. Having motivated and derived our recommended preconditioners, and shown some theoretical results on saddle point systems, we present numerical results demonstrating the potency of our solvers.

Preconditioning for the 4D-Var method in data assimilation

Alison Ramage *(University of Strathclyde)*

The 4D-Var method is frequently used for variational data assimilation problems in applications like numerical weather prediction and oceanographic modelling. One key challenge is that the state vectors used in realistic applications contain a very large number of unknowns so it can be impossible to assemble, store or manipulate the matrices involved explicitly. We present a multilevel limited-memory approximation to the inverse Hessian and illustrate its effectiveness as a

preconditioner within a Gauss-Newton iteration.

Preconditioning for partial integro-differential equations

Ekkehard Sachs & Lukas Zimmer *(Trier University, Germany)*

As a motivating example we consider a system of partial integro-differential equations that occurs in a model for cell adhesion. We also take a look at the resulting optimization problem and give some theoretical results on necessary optimality conditions using semigroup theory for the analysis.

In the case of a numerical solution for the system equations we need to implement iterative solvers. However, the question of preconditioners poses some interesting challenges. Integral operators are compact operators in the appropriate function spaces and therefore their finite dimensional counterparts need no preconditioning. On the other hand some of the integral kernels used lead to the property that the action of the integral operator can also be viewed as an approximation of the Laplacian, hence preconditioning for discretizations would be necessary.

We use results from the analysis of Toeplitz matrices to give some insight into our findings.

Optimal operator preconditioning for pseudodifferential boundary problems

Jakub Stocek & Heiko Gimperlein & Carolina Urzúa Torres *(Heriot–Watt University)*

This work considers the Dirichlet problem for an elliptic pseudodifferential operator A in a bounded Lipschitz domain Ω, where Ω is either a subset of $\mathbb{R}^n$, or, more generally, in a Riemannian manifold Γ:

$$\begin{aligned} Au &= f \qquad \text{in } \Omega, \\ u &= 0 \qquad \text{in } \Gamma \setminus \overline{\Omega}. \end{aligned}$$

Such pseudodifferential boundary problems are of interest in applications. For instance, the integral fractional Laplacian $A = (-\Delta)^s$ and its variants $A = \operatorname{div}(c(x)\nabla^{2s-1}u)$ in a domain $\Omega \subset \mathbb{R}^n$ arise in the pricing of stock options, image processing, continuum mechanics, and in the movement of biological organisms or swarm robotic systems. Boundary integral formulations of the first kind for an elliptic boundary problem lead to equations for the weakly singular ($A = \mathsf{V}$) or hypersingular ($A = \mathsf{W}$) operators on a curve segment or open surface.

On the one hand, the bilinear form associated to A is nonlocal, and its Galerkin discretization results in dense matrices. On the other hand, the condition number of these Galerkin matrices is of order $\mathcal{O}(h^{-2|s|})$, where h is the size of the smallest cell of the mesh and $2s$ is the order of A. Therefore, the solution of the resulting linear system via iterative solvers becomes prohibitively slow on fine meshes.

We propose an operator preconditioner for general elliptic pseudodifferential equations in a domain in either $\mathbb{R}^n$ or a Riemannian manifold. For linear systems of equations arising from low-order Galerkin discretizations, we prove that the condition number is independent of the mesh size and of the choice of bases for test and trial functions.

The basic ingredient is a classical formula by Boggio for the fractional Laplacian, extended analytically. In the special case of the weakly and hypersingular operators on a line segment or a screen, our approach gives a unified, independent proof for a series of recent results by Hiptmair, Jerez-Hanckes, Nédélec and Urzúa-Torres.

Numerical examples illustrate the performance of the proposed preconditioner on quasi-uniform, graded and adaptively generated meshes. They show the increasing relevance of the regularity assumptions on the mesh with the order of the operator.

Parallel preconditioning for time-dependent PDE problems

Andy Wathen *(Oxford University)*

Monolithic (or all-at-once) discretizations of evolutionary problems most often give rise to nonsymmetric linear(ised) systems of equations which can be of very large dimension for PDE problems. In this talk we will describe preconditioners for such systems with guaranteed fast convergence via use of MINRES (not LSQR or CGNE) or GMRES. These results apply with standard time-stepping schemes and relate to Toeplitz matrix technology and preconditioning via circulants using the FFT. Simple parallel computational results will be shown for the heat equation and the wave equation.

This is joint work with Elle McDonald (CSIRO, Australia), Jennifer Pestana (Strathclyde University, UK) and Anthony Goddard (Durham University, UK).

Minisymposium M5

Finite element methods for Navier-Stokes equations: theory and algorithms
Organisers
Gabriel Barrenechea and Julia Novo

A posteriori error analysis for stabilized methods: a nonlinear Boussinesq problem, and a Stokes model under singular sources

Alejandro Allendes[1] & César Naranjo[2] & Enrique Otárola[3] & Abner J. Salgado[4]
[1,2] (Universidad Técnica Federico Santa María, Chile)
[3] (Universidad Técnica Federico Santa María, Chile)
[4] (University of Tennessee)

In this talk, we will present the design and analysis of a posteriori error estimators for low-order stabilized finite element approximations of a generalized Boussinesq problem, and also for a Stokes model with singular sources in two and three dimensional Lipschitz, but not necessarily convex, polytopal domains. The designed error estimators, which are mainly based on a Ritz projection of the errors, are proven to be reliable and locally efficient. The standard smallness assumption on the solution for the nonlinear problem is required. On the basis of these estimators we design a simple adaptive strategy that yields optimal rates of convergence for the numerical examples that we perform.

An a posteriori error estimator for the MHM method applied to Stokes and Brinkman problems

Rodolfo Araya & Ramiro Rebolledo & Frédéric Valentin *(Universidad de Concepción, Chile)*

In this work, we propose and analyze a residual type a posteriori error estimator for the Multiscale Hybrid-Mixed (MHM) method applied to Stokes and Brinkman equations. This error estimator relies on the multi-level structure of the MHM method, giving valuable information about the error of approximation. As a result, the error estimator accounts for a first-level global estimator defined on the skeleton of the partition and second-level contributions from element-wise approximations. The analysis establishes local efficiency and reliability of the complete estimator. Also, it yields a new face-adaptive strategy on the mesh's skeleton which avoids changing the topology of the global mesh. Specially designed to work on multiscale problems, the present estimator can leverage parallel computers since local error estimators are independent of each other.

Low-order divergence-free finite element methods in fluid mechanics

Alejandro Allendes*, **Gabriel R. Barrenechea**†, César Naranjo* & Julia Novo‡ * *(Universidad Técnica Federico Santa María, Chile)* † *(University of Strathclyde)* ‡ *(Universidad Autónoma de Madrid)*

It is a well-known fact that the finite element approximation of equations in fluid mechanics (e.g., the Navier-Stokes equations) is simpler, and more accurate, if the finite element method delivers an exactly divergence-free velocity. In fact, in such a case the convective term remains antisymmetric in the discrete setting, without the need to rewrite it, and the stability analysis can be greatly simplified. Now, when the finite element spaces used are the conforming $\mathbf{P}_1$ for velocity and $\mathbf{P}_0$ for pressure, then, in addition to add appropriate terms to stabilise the pressure, some extra work needs to be done in order to compensate for the lack of incompressibility of the discrete velocity.

In this talk I will present two recent applications of a technique developed previously in [2] to overcome the limitations described in the previous paragraph. The first example presented is the application of this idea to the steady-state Boussinesq equation [1]. Then, the transient Navier-Stokes equations will be analysed, where some error estimates independent of the Reynolds number will be shown.

References

[1] A. Allendes, G.R. Barrenechea, C. Naranjo : A divergence-free low-order stabilized finite element method for the steady state Boussinesq problem. *Computer Methods in Applied Mechanics and Engineering*, 340, 90–120, (2018).

[2] G.R. Barrenechea and F. Valentin : Consistent local projection stabilized finite element methods. SIAM *Journal on Numerical Analysis*, 48(5), 1801-1825, (2010).

Local pressure correction for the Stokes equations

Malte Braack & Utku Kaya *(University of Kiel)*

Pressure correction methods constitute the most widely used solvers for the time-dependent Navier-Stokes equations. There are several different pressure correction methods, where each time step usually consists in a predictor step for a non-divergence-free velocity, followed by a Poisson problem for the pressure (or pressure update), and a final velocity correction to obtain a divergence-free vector field. In some situations, the equations for the velocities are solved explicitly, so that the numerical most expensive step is the elliptic pressure problem. We here propose to solve this Poisson problem by a domain decomposition method which does not need any communication between the sub-regions. Hence, this system is perfectly adapted for parallel computation. We show under certain assumptions that this new scheme has the same order of convergence as the original pressure correction scheme (with global projection). Numerical examples for the Stokes system show the effectivity of this new pressure correction method. The convergence order $\mathcal{O}(k^2)$ for resulting velocity fields can be observed in the norm $l^2(0,T;L^2(\Omega))$.

Conforming mixed finite element approximation of the Navier-Stokes/Darcy-Forchheimer problem

Marco Discacciati & Sergio Caucao & Gabriel N. Gatica & Ricardo Oyarzúa *(Loughborough University)*

We present a mixed finite element approximation of the coupled system formed by the Navier-Stokes equations and either Darcy or Darcy-Forchheimer equations to model the filtration of an incompressible fluid through a porous medium. We consider the standard mixed formulation in the Navier-Stokes domain and the dual-mixed one in the porous-medium region, which yield the introduction of the trace of the pressure in the porous medium as a suitable Lagrange multiplier. We propose a stable Galerkin approximation that employs Bernardi-Raugel and Raviart-Thomas elements for the velocities, piecewise constants for the pressures, and continuous piecewise linear elements for the Lagrange multiplier. Numerical results illustrate the performance of the method and confirm the theoretical convergence rates. Finally, we discuss possible solution techniques based on domain decomposition methods that exploit the multi-physics nature of the problem to characterize optimal preconditioners.

Mixed finite elements applied to continuous data assimilation for the Navier-Stokes equations

Bosco García-Archilla & Julia Novo & Edriss S. Titi *(Universidad de Sevilla)*

We summarize the results presented in [2] and [3], where we introduce and analyze a finite element method applied to a continuous downscaling data assimilation algorithm for the numerical approximation of the two and three dimensional Navier-Stokes equations corresponding to given measurements on a coarse spatial scale, recently introduced in [1]. For representing the coarse mesh measurements we consider different types of interpolation operators including a Lagrange interpolant. We obtain uniform-in-time estimates for the error between a finite element approximation and the reference solution corresponding to the coarse mesh measurements. We consider both the case of a standard Galerkin method and a Galerkin method with grad-div stabilization. For the stabilized method we prove error bounds in which the constants do not depend on inverse powers of the viscosity. Some numerical experiments illustrate the theoretical results.

References

[1] A. Azouani, E. Olson, and E. S. Titi. Continuous data assimilation using general interpolant observables. *J. Nonlinear Sci.*, 24(2):277–304, 2014.

[2] B. García-Archilla, J. Novo, and E. S. Titi. Uniform in time error estimates for a finite element method applied to a downscaling data assimilation algorithm for the Navier-Stokes equations. *arXiv e-prints*, page arXiv:1807.08735, Mar. 2018.

[3] B. García-Archilla and J. Novo. Error analysis of fully discrete mixed finite element data assimilation schemes for the Navier-Stokes equations. *arXiv e-prints*, page arXiv:arXiv:1904.06113, Apr. 2019.

hp-Version discontinuous Galerkin methods on essentially arbitrarily-shaped elements

Emmanuil Georgoulis *(University of Leicester)*

We extend the applicability of the popular interior-penalty discontinuous Galerkin (dG) method discretizing advection-diffusion-reaction problems to meshes comprising of extremely general, essentially arbitrarily-shaped element shapes. In particular, our analysis allows for curved element shapes, arising, *without* the use of (iso-)parametric elemental maps. The feasibil-

ity of the method relies on the definition of a suitable choice of the discontinuity-penalization parameter, which turns out to be essentially independent on the particular element shape. A priori error bounds for the resulting method are given under very mild structural assumptions restricting the magnitude of the local curvature of element boundaries. Numerical experiments are also presented, indicating the practicality of the proposed approach. This work generalizes our earlier work detailed in the monograph [A.Cangiani, Z.Dong, E.H. Georgoulis and P.Houston, *hp*-version discontinuous Galerkin methods on polygonal and polyhedral meshes, SpringerBriefs in Mathematics, Springer, Cham, 2017] from polygonal/polyhedral meshes to essentially arbitrary element shapes involving curved faces without imposing any additional mesh conditions.

The talk is based on joint work with A. Cangiani (Nottingham) and Z. Dong (IACM-FORTH, Crete).

Adaptive augmented mixed methods for the Oseen problem

M. González & T.P. Barrios & J.M. Cascón *(Universidade da Coruña)*

The problem of computing the flow of a viscous and incompressible fluid at small Reynolds numbers is described by the Oseen equations. In the recent paper [1], we introduced a new augmented variational formulation for this problem in the pseudostress-velocity variables under homogeneous Dirichlet boundary conditions for the velocity, and we developed a simple a posteriori error analysis.

In this talk we will discuss a related method for the Oseen problem when non-homogeneous mixed boundary conditions are imposed. We will describe a new augmented dual-mixed variational formulation of the problem. Then, we will analyze the corresponding Galerkin scheme, and provide the rate of convergence when each row of the pseudostress is approximated by Raviart-Thomas elements and the velocity is approximated by continuous piecewise polynomials. Moreover, we will derive an a posteriori error indicator which is reliable and locally efficient, and show the performance of the corresponding adaptive algorithm through some numerical examples.

References

[1] T.P. Barrios, J.M. Cascón and M. González, Augmented mixed finite element method for the Oseen problem: A priori and a posteriori error analyses. Comput. Methods Appl. Mech. Engrg. 313 (2017), 216–238.

Higher order variational time discretisations for the incompressible flow problems

Gunar Matthies & Simon Becher *(Technische Universität Dresden)*

We present variational time discretisation schemes of higher order applied to incompressible flow problems that are described by Stokes, Oseen, or Navier–Stokes equations. As spatial discretisation, we will consider both inf-sup stable and equal-order pairs of finite element spaces for approximating velocity and pressure.

In cases of dominant convection, a spatial stabilization is needed. We will concentrate on local projection stabilization methods which allow to stabilise the streamline derivative, the divergence constraint and, if needed for equal-order pairs, the pressure gradient separately.

Starting from the well-known discontinuous Galerkin (dG) and continuous Galerkin-Petrov (cGP) methods we will present a two-parametric family of time discretisation schemes which combine variational and collocation conditions. The first parameter corresponds to the ansatz order while the second parameter is related to the global smoothness of the numerical solution. Hence, higher order schemes with higher order regularity in time can be obtained by adjusting the family parameters in the right way.

All members of the considered family show the same stability properties as either dG or cGP. Furthermore, the considered schemes provide a cheap post-processing to achieve better convergence orders in integral-based norms. In addition, the post-processing could be used for adaptive time-step control.

Two-grid based postprocessing of mixed finite element approximations to the Navier-Stokes equations

Julia Novo & Francisco Durango *(Universidad Autónoma de Madrid)*

In this talk we revise some postprocessing techniques that have been applied to mixed finite element approximations to the Navier-Stokes equations and introduce a generalized postprocessed method that depends on several parameters. The idea of postprocessing is the following: once a standard Galerkin

mixed finite element approximation is computed over a coarse mesh of size H at a fixed time T, the approximation is postprocessed over a finer mesh of size $h < H$. In the literature one can find different ways to postprocess the Galerkin approximation depending on the problem that is solved over the fine mesh at the final time T. Basically, one can solve a Stokes problem, an Oseen type problem or a Newton type problem. In the present talk we present a method based on several parameters that generalize the above postprocessing techniques. Depending on the values chosen for the different parameters one can recover one of the old (known) postprocessing procedures but also produce a new (different) method. We get error bounds for the generalized method valid for any of the values of the different parameters. In all the cases, the postprocessed method has a rate of convergence one unit bigger than the rate of convergence of the plain Galerkin method in terms of the coarse mesh H and optimal in terms of the fine mesh h. The computational added cost of postprocessing is however negligible compared to the cost of solving the Navier-Stokes equations in the time interval $[0, T]$ by means of the Gakerkin method. For the error analysis we do not assume nonlocal compatibility conditions for the true solutions of the Navier-Stokes equations.

Numerical comparisons of finite element stabilized methods for a 2D vortex dynamics simulation at high Reynolds number

Samuele Rubino & Naveed Ahmed *(University of Seville)*

In this work, we consider up-to-date and classical Finite Element (FE) stabilized methods for time-dependent incompressible flows. All studied methods belong to the Variational MultiScale (VMS) framework. So, different realizations of stabilized FE-VMS methods are compared using a high Reynolds number vortex dynamics simulation. In particular, a fully Residual-Based (RB)-VMS method is compared with the classical Streamline-Upwind Petrov–Galerkin (SUPG) method together with grad-div stabilization, a standard one-level Local Projection Stabilization (LPS) method, and a recently proposed LPS method by interpolation. These procedures do not make use of the statistical theory of equilibrium turbulence, and no ad-hoc eddy viscosity modeling is required for all methods. Applications to the simulation of a high Reynolds number flow with vortical structures on relatively coarse grids are showcased, by focusing on a two-dimensional plane mixing-layer flow. Both Inf-Sup Stable (ISS) and Equal Order (EO) **H1**-conforming FE pairs are explored using a second-order semi-implicit Backward Differentiation Formula (BDF2) in time. Based on the numerical studies conducted, it is concluded that the SUPG method using EO-FE pairs performs best among all methods. Furthermore, there seems to be no reason to extend the SUPG method by the higher order terms of the RB-VMS method.

IFISS: A computational laboratory for investigating incompressible flow problems

David Silvester *(The University of Manchester)*

The IFISS package is a computational laboratory for the interactive numerical study of incompressible flow problems. It includes algorithms for discretization by mixed finite element methods and a posteriori error estimation of the computed solutions, together with state-of-the-art preconditioned iterative solvers for the resulting discrete linear equation systems. The utility of the software is illustrated using two case studies. First, by computing a precise on-the-fly estimate of the discrete inf–sup constant when solving a Stokes flow problem. Second, by assessing the issues that arise when conventional mixed methods are used to approximate Beltrami flows or potential flows using a standard Navier–Stokes solver. The goal is to demonstrate that IFISS can be a valuable tool in both teaching and research.

References

[1] H. Elman, A. Ramage and D. Silvester. IFISS: A computational laboratory for investigating incompressible flow problems, SIAM Review, 56, 261273, 2014. https://doi.org/10.1137/120891393

[2] V. John, A. Linke, C. Merdon, M. Neilan and L. Rebholz. On the divergence constraint in mixed finite element methods for incompressible flows, SIAM Review, 593, 492–544, 2017. https://doi.org/10.1137/15M1047696

Minisymposium M6

Computational methods for model-driven optimization and control under uncertainty

Organisers
Dante Kalise and Mattia Zanella

Kinetic models for optimal control of wealth inequalities

Bertram Düring *(University of Sussex)*

We discuss optimal control strategies for kinetic models for wealth distribution in a simple market economy, which are designed to reduce the variance of the wealth density among the population. Our analysis is based on a finite time horizon approximation, or model predictive control, of the corresponding control problem for the microscopic agents' dynamic and results in an alternative theoretical approach to the taxation and redistribution policy at a global level.

Hybrid differential games and their application to a match race problem

Adriano Festa *(L'Aquila University) & Simone Cacace & Roberto Ferretti* (Roma tre University)

We discuss the general setting of a two-player hybrid differential game where the dynamics of the players follow some stochastic differential equations with the possibility to switch, paying a cost, from one to another. We analyse the abstract problem and apply it to the modelling of a "match race" competition between sailing boats. In a simplified version of this competition, the goal of both players is to proceed on the windward direction while trying to slow down the other player. We model the problem using the structure of a stochastic, zero-sum differential game in presence of hybrid dynamics, show that the game has a value, provide a convergent approximation scheme and validate the approach on some typical racing scenarios.

Controlling fluctuations in gas networks using stochastic Galerkin formulations

Stephan Gerster & Michael Herty *(RWTH Aachen University)*

Since gas system operators assume nearly constant consumption, there is a need to assess the risk of stochastic fluctuations. Typically, simulating stochastic processes by a Monte-Carlo method is computationally not feasible.
The idea to represent stochastic processes by orthogonal polynomials has been employed in uncertainty quantification and inverse problems. This approach is known as stochastic Galerkin formulation with a generalized polynomial chaos (gPC) expansion. The gPC expansions of the stochastic input are substituted into the governing equations. Then, they are projected by a Galerkin method to obtain deterministic evolution equations for the gPC coefficients.

So far, results for general hyperbolic systems are not available. A problem is posed by the fact that the deterministic Jacobian of the projected system differs from the random Jacobian of the original system and hence hyperbolicity is not guaranteed. Applications to hyperbolic conservation laws are in general limited to linear and scalar hyperbolic equations. We present a stochastic Galerkin formulation for isothermal Euler equations. Furthermore, we discuss numerically control policies to damp fluctuations in the network.

References

S. Gerster, M. Herty, *Entropies and Symmetrization of Hyperbolic Stochastic Galerkin Formulations*, Preprint-No. 488 (2019), RWTH Aachen University.

S. Gerster, M. Herty, A. Sikstel, *Hyperbolic Stochastic Galerkin Formulation for the p-System*, Preprint-No. 477 (2018), RWTH Aachen University.

S. Gerster, M. Herty, *Discretized Feedback Control for Systems of Linearized Hyperbolic Balance Laws*, MCRF, Vol. 9, No 3, 2019.

Parameter estimation for macroscopic pedestrian dynamics models

Susana N. Gomes *(Warwick)* & Andrew M. Stuart *(Caltech)* & Marie-Therese Wolfram *(Warwick)*

In this talk, we present a framework for estimating parameters in macroscopic models for crowd dynamics using data from individual trajectories. We consider a model for the unidirectional flow of pedestrians in a corridor which consists of a coupling between a density-dependent stochastic differential equation and a nonlinear partial differential equation for the density. In the stochastic differential equation for the trajectories, the velocity of a pedestrian decreases with the density according to the fundamental diagram. Although there is a general agreement on the basic shape of this dependence, its parametrization depends strongly on the measurement and averaging techniques used as well as the experimental setup considered. We will discuss identifiability of the parameters appearing in the fundamental diagram, introduce optimisation and Bayesian methods to perform the identification, and analyse the performance of the proposed methodology in various realistic situations. Finally, we discuss possible generalisations, including the effect of the form of the fundamental diagram and the use of

experimental data.

Polynomial approximation of Isaacs' equation and robust control of parabolic PDEs

Dante Kalise & Sudeep Kundu & Karl Kunisch *(University of Nottingham)*

We propose a numerical scheme for the approximation of high-dimensional, nonlinear Isaacs PDEs arising in robust optimal feedback control of nonlinear dynamics. The numerical method consists of a global polynomial ansatz together with separability assumptions for the calculation of high-dimensional integrals. The resulting Galerkin residual equation is solved by means of an alternating Newton-type/policy iteration method for differential games. We present numerical experiments illustrating the applicability of our approach in robust optimal control of nonlinear parabolic PDEs.

Kinetic models of traffic flow control via driver-assist vehicles

Andrea Tosin & Mattia Zanella *(Politecnico di Torino)*

In this talk, we present a hierarchical description of control problems for vehicular traffic, which aim to mitigate speed-dependent risk factors and to dampen structural uncertainties responsible for scattered, thus hardly predictable, aggregate trends. In particular, we implement mathematically the idea that a few automated cars can be controlled so as to align the speeds in the traffic stream either to each other or to some recommended optimal speed.

First, we discuss the modelling of stochastic microscopic binary interactions among the vehicles, including a probabilistic description of the penetration rate of automated vehicles. When the interactions involve one of such vehicles, they are further subject to a binary control problem, which aims to reduce the speed gap either with the leading vehicle or with a prescribed congestion-dependent speed. Then, we upscale the interaction rules to the global flow by means of a kinetic Boltzmann-type equation. Finally, we use the kinetic equation to investigate the impact of the microscopic control on the macroscopic flow. In particular, by means of suitable hydrodynamic limits, we show how to obtain consistent macroscopic traffic models, which incorporate the action of the control from the microscopic scale.

Ours is a markedly multiscale approach, particularly consistent with the idea that multi-agent systems, like vehicular traffic, need to be controlled by means of bottom-up strategies rather than assuming that it is possible to control directly their aggregate behaviour.

This talk is mainly based on the contents of the following works:

References

[1.] A. Tosin, M. Zanella. Kinetic-controlled hydrodynamics for traffic models with driver-assist vehicles, *Multiscale Model. Simul.*, 2019. doi:10.1137/18M1203766.

[2.] A. Tosin, M. Zanella. Control strategies for road risk mitigation in kinetic traffic modelling, *IFAC-PapersOnLine*, 51(9):67-72, 2018.

[3.] A. Tosin, M. Zanella. Uncertainty damping in kinetic traffic models by driver-assist controls, 2019 (preprint). doi:10.13140/RG.2.2.35871.41124.

Uncertainty quantification for moving boundary problems with applications in composites manufacturing

M.V. Tretyakov *(University of Nottingham)*

The considered UQ problems are motivated by modelling of one of the main manufacturing processes for producing advanced composites – resin transfer moulding (RTM). We consider models of the stochastic RTM process, which are formulated as random moving boundary problems. We study their properties, analytically in the one-dimensional case and numerically in the two-dimensional case. We will discuss Bayesian inversion algorithm for random moving boundary problems and a stochastic control setting. The talk is based on joint works with Marco Iglesias, Mikhail Matveev, Andreas Endruweit, Andy Long, Minho Park.

Spectral methods for optimal control problems for equations of Fokker–Planck type

Urbain Vaes & Dante Kalise & Grigorios Pavliotis *(Imperial College London)*

In this talk we present a novel methodology for controlling the Fokker–Planck equation in periodic and unbounded domains. Our method is based on the derivation of first-order optimality conditions, leading a forward-backward coupled system which is approxi-

mated by an expansion in either Fourier series or appropriately rescaled Hermite functions, depending on the domain. We propose a reduced gradient method for the synthesis of the optimal control signal, and we study the properties of the proposed numerical scheme both theoretically and by means of numerical experiments.

Structure preserving gPC methods for kinetic equations with uncertainties

Mattia Zanella *(Politecnico di Torino)*

We introduce and discuss numerical schemes for the approximation of kinetic equations that incorporate uncertain quantities. In contrast to a direct application of stochastic Galerkin generalized polynomial chaos (SG-gPC) methods which are widely considered for uncertainty quantification of differential equations, this class of schemes make use of a particle approach in the phase space coupled with a stochastic Galerkin expansion in the random space. The proposed methods naturally preserve the positivity of the statistical moments of the solution and are capable of achieving spectral accuracy in the random space. Several tests on kinetic models for collective phenomena validate the proposed methods both in the homogeneous and inhomogeneous setting, shading light on the influence of uncertainties in phase transition phenomena driven by noise such as their smoothing and confidence bands.

References

[1.] G. Dimarco, L. Pareschi, M. Zanella. Uncertainty quantification for kinetic models in socio-economic and life sciences. In *Uncertainty Quantification for Hyperbolic and Kinetic Equations*, Editors S. Jin, and L. Pareschi, SEMA SIMAI Springer Series, vol. 14, pp 151-191, 2017.

[2.] J. A. Carrillo, M. Zanella. Monte Carlo gPC methods for diffusive kinetic flocking models with uncertainties. Preprint `arXiv:1902.04518`, 2019.

[3.] J. A. Carrillo, L. Pareschi, M. Zanella. Particle based gPC methods for mean-field models of swarming with uncertainty. *Communications in Computational Physics*, 25(2): 508-531, 2019.

Minisymposium M7

Probabilistic numerical computation and high-dimensional data analysis

Organiser
Ivan Tyukin

A measure based construction of the finite element method

Mark A. Girolami *(University of Cambridge)*

The finite element method provides the means to numerically approximate the solution to linear and nonlinear systems of partial differential equations. It is without doubt one of the triumphs of applied mathematics in enabling the modelling and study of complex physical and natural systems and processes. However the uncertainty inherent in such models is restricted to that introduced by poorly defined values of physical properties such that empirical measurements in the form of data are used in the inverse-problem setting to infer plausible ranges of such values. However the uncertainty induced by model misspecification is ignored. By defining probability measures over the differential operators and the PDEs themselves along with a statistical representation of the model-reality mismatch a set of conditional probability measures emerge which provide the means to update the finite element solution in the light of empirical evidence and thus more informatively characterise uncertainty in the overall system. This talk will present the construction of such a statistical measure based finite element methodology and explore its properties theoretically and with applications.

Approximation with Gaussians of fixed and varying scales

Jeremy Levesley & Stephen Walsh *(University of Leicester)*

Approximation with Gaussians has advantages and disadvantages. One of the main disadvantages is that the analytic nature of the function means that approximation problems with Gaussians are often ill-conditioned. On the other hand, since Gaussians are smooth, they carry the possibility of optimal convergence orders when the smoothness of the underlying function is unknown (as it usually is in practice). Also, in high dimensions, most functions start to look like Gaussians.

In standard radial basis function applications Gaussians of a fixed width are used, and convergence is obtained by reducing the spacing between data points while leaving the shape of the Gaussian fixed and constant. We would like to explore a different scenario, which is to use a multiscale method, where the shape depends on the scale. We will explore this for a fixed shape at each scale and also a varying shape at each scale. The shape is changed related to end-point singularities in the function in order to mimic $h-p$ adaptivity in finite elements.

We will review results in the multilevel fixed shape paradigm. Collaborators in this work are Manolis Georgoulis, Peter Dong, Fuat Usta, Fazli Subhan (all at Leicester), and Simon Hubbert at Birkbeck. We will also show that the interpolation matrices in the multiscale, multishape scenario in one dimension are invertible, a new result.

Kernel stochastic separation theorems and mathematics for making data-driven artificial intelligence systems better

I. Tyukin & A.N. Gorban *(University of Leicester)*

Large and growing streams of data are ubiquitous in modern society and form the backbone of modern healthcare, public safety, services, and science. Sustained functioning and progress in these essential areas depend on the ability to extract and process information from large and growing data. Since processing overwhelmingly large volumes of data can no longer be accomplished by humans alone, we must rely on Artificial Intelligence (AI) systems built on state-of-the-art machine learning and data analytics technologies.

With the explosive pace of progress in computing, current AI systems are now capable of spotting minute patterns in large data sets and outperform humans in highly complicated cognitive tasks like chess, Go, and medical diagnosis. However, the super-human capability of modern AIs to learn from massive volumes of data make their conclusions vulnerable to data inconsistencies, poor data quality, and fundamental uncertainty inherent to any data. This uncertainty together with engineering constraints on AI's implementation lead to inevitable errors in data-driven AIs: incorrect face recognition matchings, incorrect cancer diagnosis, and Tesla and Uber crashes are just few examples.

In this talk we provide a mathematical formulation of the challenge of handling of AI errors and present a mechanism to address the challenge: stochastic separation theorems. We then further extend stochastic separation theorems to infinite-dimensional settings. A general separability result for two random sets and relaxed independence assumption is also established. We show that despite feature maps corresponding to a given kernel function may be infinite-dimensional, kernel separability characterizations can be expressed in terms of finite-dimensional volume integrals. These integrals allow to determine and quantify separability properties of an arbitrary kernel function. The theory is illustrated with numerical examples.

Minisymposium M8

Matrix methods for Networks
Organisers
Francesca Arrigo and Francesco Tudisco

Eigenvector-based centrality measures in multilayer networks

Francesca Arrigo & Francesco Tudisco *(University of Strathclyde)*

Complex networks are used to model several different types of interactions among various entities. Being so versatile, they provide a useful tool to understand today's world and answer questions about the underlying phenomenon or structure. Often in practice the question that needs answered is: which is the most important entity in the network? Several measures of importance, or centrality, have been introduced over the years. Among these, one of the most popular ones builds on the idea that the importance of a node stems from that of its neighbours. Mathematically this results in considering the Perron eigenvector of some suitable graph matrices. In this talk, we describe how to generalize this simple but effective idea to the more complicated setting of multi-layer networks, where entities interact on different levels and the higher-order structure is captured by tensors rather than matrices.

Spectral methods for certain inverse problems on graphs

Mihai Cucuringu *(University of Oxford & The Alan Turing Institute)*

We study problems that share an important common feature: they can all be solved by exploiting the spectrum of their corresponding graph Laplacian matrix.

We first consider the classic problem of establishing a statistical ranking of a set of items given a set of inconsistent and incomplete pairwise comparisons between such items. Instantiations of this problem occur in numerous applications in data analysis (e.g., ranking teams in sports data), computer vision, and machine learning. We formulate the above problem of ranking with incomplete noisy information as an instance of the group synchronization problem over the group SO(2) of planar rotations, whose least squares solution can be approximated by either a spectral or a semidefinite programming relaxation. We also present a simple spectral approach to the well-studied constrained clustering problem. It captures constrained clustering as a generalized eigenvalue problem with graph Laplacians. The proposed algorithm works in nearly-linear time, and consistently outperforms existing spectral approaches both in speed and quality. Furthermore, we propose an algorithm for clustering signed networks, where the edge weights between the nodes of the graph may take either positive or negative edges, and analyze it under a signed stochastic block model.

Metaplex networks: influence of the exo-endo structure of complex systems on diffusion

Gissell Estrada-Rodriguez & Ernesto Estrada & Heiko Gimperlein *(Maxwell Institute for Mathematical Sciences and Heriot-Watt University)*

In this talk I will introduce the concept of metaplexes and the dynamical systems on them. A metaplex combines the internal structure of the entities of a complex system with the discrete interconnectivity of these entities in a global topology. We focus here on the study of diffusive processes on metaplexes, both model and real-world examples. We provide theoretical and computational evidence pointing out the role of the endo- and exo-structure of the metaplexes in their global dynamics, including the role played by the size of the nodes, the location of sinks and sources, and the strength and range of the coupling between nodes. We show that the internal structure of brain regions (corresponding to the nodes of the network) in the macaque visual cortex metaplex dominates almost completely the global dynamics. On the other hand, in the linear metaplex chain and the landscape metaplex the diffusion dynamics display a combination of the endo- and exo-dynamics, which we explain analytically and in numerical results. Metaplexes are expected to facilitate the understanding of complex systems in an integrative way, which combine dynamical processes inside the nodes and between them.

The random walk centrality, revised

Dario Fasino *(University of Udine)*

Let H be the matrix containing the hitting times of the random walk on a connected undirected graph on n nodes. It is well known that the matrix $H + H^T$ is a multiple of the resistance matrix. On the other hand, the matrix $H - H^T$ has rank 2. In fact, there exist positive numbers $c_1, \ldots, c_n$ such that

$$H_{ij} - H_{ji} = \frac{1}{c_j} - \frac{1}{c_i}.$$

In a renowned paper by J. Noh and H. Rieger, c_i is defined in terms of an infinite series and is called the *random walk centrality* of node i, since it quantifies how central node i is regarding its potential to receive information randomly diffusing over the network. In practice, the value of c_i is quite good at recognizing node i as being a "core" node or a 'peripheral" node in the network. In this talk I propose a novel interpretation of c_i and new formulas to compute it from the pseudo-inverse of the Laplacian matrix of the network. Extensive numerical experiments show that c_i has a striking correlation with the so-called *random walk closeness centrality*, $C(i) = n/\sum_{j=1}^{n}(H_{ij} + H_{ji})$. The reason for that is, at the time of writing this abstract, not completely clear.

Bipartivity measures and methods

Philip Knight & Azhar Aleidan *(University of Strathclyde)*

Bipartivity is a well-defined concept in graph theory, namely that a graph's nodes can be divided into two sets in which there are no intra-set links. Often one finds that real-world networks have a close-to-bipartite structure, and there are several applications where it is helpful to both quantify and identify this structure. We compare measures based on links, walks and spectral information and show that these can give significantly different quantifications. Each measure leads to an algorithm for finding a nearby bipartite graph to any network. Finding the closest such graph is an NP hard problem, and we give some examples to show how close some simple matrix-based algorithms can come to finding an optimal solution in reasonable time.

Networks core–periphery detection with nonlinear Perron eigenvectors

Francesco Tudisco *(University of Strathclyde)* & Desmond J. Higham *(University of Edinburgh)*

Core–periphery detection is a highly relevant task in exploratory network analysis. Given a network of nodes and edges, one is interested in revealing the presence and measuring the consistency of a *core–periphery structure* using only the network topology. This mesoscale network structure consists of two sets: the core, a set of nodes that is highly connected across the whole network, and the periphery, a set of nodes that is well connected only to the nodes that are in the core. Networks with such a core–periphery structure have been observed in several applications, including economic, social, communication and citation networks.

In this work we propose a new core–periphery detection model based on the optimization of the core–periphery quality function

$$f(x) = \sum_{ij=1}^{n} A_{ij} \max\{x_i, x_j\} \tag{1}$$

on the unit sphere $\|x\| = 1$. While the quality measure f is highly nonconvex in general and thus hardly treatable, we show that the global maximum of f coincides with the *nonlinear* Perron eigenvector of a suitably defined parameter dependent matrix $M(x)$, i.e. the positive solution to the nonlinear eigenvector problem $M(x)\,x = \lambda\,x$.

Using recent advances in nonlinear Perron–Frobenius theory, we show that (1) has a unique solution and we propose a nonlinear power-method type scheme that (a) allows us to solve (1) with global convergence guarantees and (b) effectively scales to very large and sparse networks. Finally, we present several numerical experiments showing that the new method largely outperforms state-of-the-art techniques for core-periphery detection.

Minisymposium M9

Recent Advances in Continuous Optimisation

Organisers
Nick Gould and Lindon Roberts

An analysis of biased stochastic gradient descent methods

Derek Driggs & Jingwei Liang *(University of Cambridge)* & Carola Schönlieb *(University of Cambridge)*

Although stochastic gradient methods have received much attention in the machine learning and optimisation communities, almost all existing research considers unbiased stochastic gradient estimators. We discuss a framework for analysing several families of biased stochastic gradient methods that include SAGA, SAG, SVRG, and SARAH as special cases. We use this framework to develop a new biased stochastic gradient algorithm, Stochastic Average Recursive GradiEnt (SARGE), that combines SARAH's recursive gradient estimate with the SAGA gradient estimator. We show that biased stochastic gradient estimators including SARGE often outperform their unbiased counterparts on a variety of tasks from machine learning and imaging processing.

A block-coordinate Gauss-Newton method for nonlinear least squares

Jaroslav Fowkes & Coralia Cartis *(University of Oxford)*

We propose a block-coordinate Gauss-Newton method for nonlinear (nonconvex) least squares problems that computes Jacobians only on a subset of the optimization variables at a time. We investigate globalising this approach using either a regularization term or a trust-region model and show global complexity results as well as extensive computational results on CUTEst test problems. Furthermore, as our approach exhibits very slow rates of convergence on certain nonconvex problems, we design adaptive block size variants of our methods that can overcome these difficulties.

Solution of quadratic programming problems for fast approximate solution of linear programming problems

Ivet Galabova & Julian Hall *(University of Edinburgh)*

Linear programming problems (LP) are widely solved in practice and reducing the solution time is essential for many applications. This talk will explore solving a sequence of unconstrained quadratic programming (QP) problems to derive an approximate solution and bound on the optimal objective value of an LP. Techniques for solving these QP problems fast will be dis-

cussed.

An exact line search algorithm for piecewise smooth functions

Jonathan Grant-Peters *(University of Oxford)*

The problem of finding the local optimum of a 1D smooth function is one that has been solved for many years. Brent's method solves this problem with superlinear convergence while the Golden Section algorithm solves this problem robustly in linear time. However, when considering a piecewise smooth function, it is not difficult to construct an example for which Brent's method fails to converge faster than the Golden Section. We present an algorithm designed to converge quickly for a class of piecewise smooth functions, while being comparable to Brent's method for smooth functions.

HiGHS: a high-performance linear optimizer

Julian Hall *(University of Edinburgh)*

This talk will present `HiGHS`, a growing open-source repository of high-performance software for linear optimization based on award-winning computational techniques for the dual simplex method. The talk will give an insight into the work which has led to the creation of `HiGHS` and then set out the features which allows it to be used in a wide range of applications. Plans to extend the class of problems which can be solved using `HiGHS` will be set out.

On barrier and modified barrier multigrid methods for 3d topology optimization

Michal Kočvara & Alexander Brune *(University of Birmingham)*

Topology optimization of mechanical structures is a challenging problem due to its very large dimension easily reaching millions of variables. We propose to solve it by the Penalty-Barrier Multiplier (PBM) method, introduced by R. Polyak and later studied by Ben-Tal and Zibulevsky and others. The most computationally expensive part of the algorithm is the solution of linear systems arising from the Newton method used to minimize a generalized augmented Lagrangian. We use a special structure of the Hessian of this Lagrangian to reduce the size of the linear system and to convert it to a form suitable for a standard multigrid method. This converted system is solved approximately by a multigrid preconditioned MINRES method. The proposed PBM algorithm is compared with the optimality criteria (OC) method and an interior point (IP) method, both using a similar iterative solver setup. We apply all three methods to different loading scenarios. In our experiments, the PBM method clearly outperforms the other methods in terms of computation time required to achieve a certain degree of accuracy.

A deterministic approach to avoid saddle points

Lisa Maria Kreusser *(University of Cambridge)* & Stanley J. Osher *(UCLA)* & Bao Wang *(UCLA)*

Loss functions with a large number of saddle points are one of the main obstacles to training many modern machine learning models. For certain initial values, gradient descent (GD) converges to a saddle point. We call the region formed by these initial values the 'attraction region'. For quadratic functions, GD converges to a saddle point if the initial data is in a subspace of up to $n-1$ dimensions. In this talk, we prove that a small modification of the recently proposed Laplacian smoothing gradient descent (LSGD) [Osher, et al., arXiv:1806.06317] contributes to avoiding saddle points without sacrificing the convergence rate of GD. In particular, we show for a class of quadratic functions that the dimension of the LSGD's attraction region is $\lfloor (n-1)/2 \rfloor$ which is significantly smaller than GD's (n-1)-dimensional attraction region.

Dimensionality reduction techniques for global optimization

Adilet Otemissov & Coralia Cartis *(The Alan Turing Institute and The University of Oxford)*

We show that the scalability issues of Global Optimization algorithms can be alleviated for functions with low effective dimensionality, which are constant along certain linear subspaces. Such functions can be found in various applications, for example, in neural networks and complex engineering simulations. We propose the use of random subspace embeddings within any global minimization algorithm, extending Wang et al.'s (2013) approach. Using tools from random matrix theory and conic integral geometry, we investigate the efficacy and convergence of the method in its static and adaptive formulations, respectively. Us-

ing popular global solvers, we illustrate our algorithmic proposals/theoretical findings numerically. Joint work with Coralia Cartis.

A geometric integration approach to nonsmooth, nonconvex optimisation

Erlend S. Riis & Matthias M. Ehrhardt & Reinout Quispel & Carola-B. Schönlieb *(University of Cambridge)*

Discrete gradient methods are popular numerical methods from geometric integration for solving systems of ODEs. They are well-known for preserving structures of the continuous system such as energy dissipation/conservation. The preservation of dissipation makes discrete gradient methods interesting for optimisation problems. In this talk, we consider a derivative-free discrete gradient applied to dissipative ODEs such as gradient flow, thereby obtaining optimisation schemes that simultaneously are implementable in a black-box setting and retain favourable properties of gradient flow. We give a theoretical analysis of these schemes in the nonsmooth, nonconvex setting, and conclude with numerical results for the bilevel optimisation of regularisation parameters in image processing.

Improving the scalability of derivative-free optimization for nonlinear least-squares problems

Lindon Roberts & Coralia Cartis & Jan Fiala & Benjamin Marteau *(University of Oxford)*

In existing techniques for model-based derivative-free optimization, the computational cost of constructing local models and Lagrange polynomials can be high. As a result, these algorithms are not as suitable for large-scale problems as derivative-based methods. In this talk, I will introduce a derivative-free method based on exploration of random subspaces, suitable for nonlinear least-squares problems. This method has a substantially reduced computational cost (in terms of linear algebra), while still making progress using few objective evaluations.

Sketching for sparse linear least squares

Zhen Shao[1] & Coralia Cartis[1] & Jan Fiala[2]
([1] University of Oxford, [2] NAG Ltd)

We discuss sketching techniques for sparse Linear Least Squares (LLS) problems, that perform a randomised dimensionality reduction for more efficient and scalable solutions. We give theoretical bounds for the accuracy of the sketched solution/residual when hashing matrices are used for sketching, quantifying carefully the trade-off between the coherence of the original, un-sketched matrix and the sparsity of the hashing matrix. We then use these bounds to quantify the success of our algorithm that employs a sparse factorisation of the sketched matrix as a preconditioner for the original LLS, before applying LSQR. We extensively compare our algorithm to state-of-the-art direct and iterative solvers for large-scale and sparse LLS, with encouraging results.

Model based optimisation applied to black-box attacks in deep learning

Giuseppe Ughi & Prof. J. Tanner *(University of Oxford)*

Neural Network algorithms have achieved unprecedented performance in image recognition over the past decade. However, their application in real world use-cases, such as self driving cars, raises the question of whether it is safe to rely on them.

We generally associate the robustness of these algorithms with how easy it is to generate an adversarial example: a tiny perturbation of an image which leads it to be misclassified by the Neural Net (which classifies the original image correctly). Neural Nets are strongly susceptible to such adversarial examples, but when the architecture of the target neural net is unknown to the attacker it becomes more difficult to generate these examples efficiently.

In this Black-Box setting, we frame the generation of an adversarial example as an optimisation problem solvable via derivative free optimisation methods. Thus, we introduce an algorithm based on the BOBYQA model-based method and compare this to the current state of the art algorithm.

Minisymposium M10

Spectral Methods
Organisers
Sheehan Olver, Alex Townsend and Marcus Webb

Spectral problems and new resolvent based

methods

Matthew Colbrook *(University of Cambridge)*

I will present new results on spectral computations for linear operators on separable Hilbert spaces. This problem has a rich history, leading to the Solvability Complexity Index hierarchy - a measure of the difficulty of a computational problem. I will discuss classifications of spectral problems in this hierarchy. Some long-standing problems are solved (or shown to be impossible) revealing potential surprises. For example, the problem of computing spectra of compact operators, for which the method has been known for decades, is strictly harder than the problem of computing spectra of Schrödinger operators with bounded potentials, which has been open for more than half a century. The latter type of problem can be solved with a precise form of *error control* and without spectral pollution by using computational estimates of resolvent norms. Another resolvent based method I will introduce is the first set of general algorithms for computing *spectral measures* of self-adjoint and unitary operators. Numerical examples are given throughout, demonstrating not only the theoretical convergence guarantees of the algorithms, but also that they are competitive with "state-of-the-art" methods (which do not converge in general). These algorithms are members of a growing family of resolvent based methods, moving away from the "discretise-then-solve" paradigm. They are parallelisable and have natural links with, as well as making use of, spectral methods for solutions of PDEs.

The ultraspherical spectral element method

Daniel Fortunato *(Harvard University)* & Nicholas Hale *(Stellenbosch University)* & Alex Townsend *(Cornell University)*

We introduce a novel spectral element method based on the ultraspherical spectral method and the hierarchical Poincaré–Steklov scheme for solving general partial differential equations on polygonal unstructured meshes. Whereas traditional finite element methods have a computational complexity that scales as $\mathcal{O}(p^6)$ with polynomial degree p, our method scales as $\mathcal{O}(p^4)$, allowing for hp-adaptivity to be based on physical considerations instead of computational ones. Properties of the ultraspherical spectral method lead to almost banded well-conditioned linear systems, allowing for the element method to be competitive in the ultra-high-polynomial regime and to be robust to meshes that contain elements with small aspect ratios. The hierarchical Poincaré–Steklov scheme enables precomputed solution operators to be reused, allowing for fast elliptic solves in implicit and semi-implicit time-steppers. We develop an open-source software system to demonstrate the flexibility and robustness of the method.

A fast, robust solver for the Lippmann-Schwinger equation

Abinand Gopal & Per-Gunnar Martinsson *(University of Oxford)*

An incredibly effective strategy for solving certain partial differential equations is to represent the solution as an integral of a known kernel times an unknown density function and to then solve for the density by discretizing an integral equation. The resulting linear systems are typically dense and, after quadrature corrections necessary to achieve higher-order convergence, can be ill-conditioned and therefore challenging for iterative solvers. We present a new high-order solver for the Lippmann-Schwinger equation, an integral equation reformulation of the Helmholtz equation for inhomogeneous scatterers. The key idea is to utilize rank structure in a fast, robust way.

A sparse spectral method for Volterra integral equations on the triangle

Timon Gutleb & Sheehan Olver *(Imperial College London)*

We present a sparse spectral method to solve Volterra integral equations

$$\mathrm{V}_K u = f \qquad \text{or} \qquad (\lambda I + \mathrm{V}_K)u = f$$

with general kernels $K(x, y)$ using bivariate orthogonal polynomials on a triangle domain. The Volterra integral operator

$$(\mathrm{V}_K u)(x) := \int_0^x K(x, y)u(y)\mathrm{d}y$$

is shown to be sparse on a weighted Jacobi polynomial basis which allows us to achieve high efficiency and exponential convergence without being restricted to convolution kernel cases. We close by discussing convergence proofs and showing solutions to sample Volterra integral equations of first and second kind.

Resolvent techniques for computing the spectrum of differential operators

Andrew Horning & Alex Townsend *(Cornell University)*

The spectral decomposition of a differential operator $\mathcal{L}$ is fundamental to the study of differential equations: characterizing the spectrum and invariant subspaces of $\mathcal{L}$ allows one to construct simple solutions, analyze stability, and determine the asymptotic behavior of solutions. In classical and quantum physics, this decomposition often manifests in the familiar form of eigenvalues and eigenfunctions. However, more sophisticated radiative and dynamic models also admit resonant states, which are associated with the continuous spectrum and spectral measure of $\mathcal{L}$. We present a suite of new methods for computing the spectrum and spectral measure directly, providing access to measurable quantities of interest and an intuitive visualization of the operator $\mathcal{L}$. Our technique is rooted in the complex-analytic ideas of the resolvent-formalism, which describes the spectrum of $\mathcal{L}$ using the singularities of the resolvent $\mathcal{R}(z, \mathcal{L}) = (z\mathcal{I} - \mathcal{L})^{-1}$. We illustrate its advantages through a series of physically motivated examples.

Orthogonal polynomials on algebraic curves and surfaces

Sheehan Olver *(Imperial College)*

Multivariate orthogonal polynomials on the circle and sphere are classical: they are precisely Fourier series and spherical harmonics. The circle and sphere are special cases of algebraic curves/surfaces, that is, the zero set of a polynomial. This raises the following question: what are the properties of orthogonal polynomials on other algebraic curves? Can they be used for numerically solving surface PDEs? We will investigate this question on simple quadratic curves and surfaces (e.g. paraboloids and hyperboloids). We will also discuss examples where the boundary of the domain is specified by an algebraic curve, in particular, we will show how OPs can be used to solve PDEs in a half-disk.

Joint work with Ben Snowball (Imperial) and Yuan Xu (U. Oregon).

Continuous analogues of Krylov methods for differential operators

Marc Gilles & **Alex Townsend** *(Cornell University)*

Krylov subspace methods are a popular class of iterative algorithms for solving $Ax = b$ via matrix-vector products, with their convergence rates often depending on the condition number of A. In this talk we develop analogues of the conjugate gradient method, MINRES, and GMRES for solving $\mathcal{L}u = f$ via operator-function products, where $\mathcal{L}$ is an unbounded elliptic differential operator. Our Krylov methods employ operator preconditioning, have spectral convergence properties, and preserve the bilinear form associated to $\mathcal{L}$. They are practical algorithms that are computationally competitive for spectral discretizations.

A new Laplace solver for regions with corners

Nick Trefethen & Abi Gopal *(University of Oxford)*

Consider the Laplace equation with Dirichlet boundary data in a polygon or other domain with corners. We describe a new method for such problems of the flavour of the Method of Fundamental Solutions, but with probably root-exponential convergence, i.e., error $= O(\exp(-C\sqrt{N}))$ with $C > 1$. The idea is to represent the function as the real part of a rational function with poles exponentially clustered near each corner, leveraging a celebrated theorem in approximation theory by Donald Newman in 1964. A short Matlab implementation solves problems amazingly fast, as we will demonstrate.

Structured matrices and other desiderata for spectral methods

Marcus Webb *(KU Leuven)*

We discuss recent progress in designing the basis of a spectral method so as to give rise to highly structured linear systems, preserve self-adjointness of self-adjoint differential equations, and whose coefficients can be computed by a fast algorithm (for example, utilising the Fast Fourier Transform). The advantages of these properties are more pronounced in time-dependent problems and eigenvalue problems than in boundary value problems. We conclude with open questions about the feasibility of bases with these desiderata. This is joint work with Arieh Iserles (Cambridge).

Minisymposium M11

Numerical Methods for Partial Differential Equations in Unbounded Domains

Organisers

Buyang Li and Chunxiong Zheng

Fast and memory efficient solution of boundary integral formulations of the Schrödinger equation

Lehel Banjai & Mária López Fernández *(Heriot-Watt University)*

Fast and oblivious quadrature was introduced by López-Fernández, Lubich and Schädle for convolutions with a kernel whose Laplace transform is a sectorial operator. The algorithm can compute N steps of a convolution quadrature approximation of the convolution while using only $O(\log N)$ active memory and with $O(N \log N)$ computational complexity.

In this talk we describe how oblivious quadrature can be extended to some non-sectorial operators. In particular we present an application to the integral formulation of a non-linear Schrödinger equation describing the suppression of quantum beating.

A boundary integral equation method for linear elastodynamics problems on unbounded domains

Silvia Falletta & Giovanni Monegato & Letizia Scuderi *(Politecnico di Torino)*

We consider transient 3D elastic wave propagation problems in unbounded isotropic homogeneous media, which can be reduced to corresponding 2D ones. This is the case, for example, of problems defined on the exterior of a bounded rigid domain, which are invariant in one of the Cartesian directions. Thus, the linear elastodynamics problem that characterizes small variations of a displacement field $\mathbf{u}$ in a medium Ω^e having the above properties, caused by a body force $\mathbf{f}$, initial conditions $\mathbf{u_0}\mathbf{v_0}$ locally supported and a Dirichlet datum $\mathbf{g}$, is defined by the following system:
In [1] the two scalar equations are discretized by combining a Lubich time convolution quadrature formula with a classical space collocation method (see [3] and [2]). Several numerical results, including comparisons with the vector space-time boundary integral formulation approach are presented and discussed.

References

[1] S. Falletta and G. Monegato L. Scuderi, Two boundary integral equation methods for linear elastodynamics problems on unbounded domains, *submitted*.

[2] S. Falletta and G. Monegato L. Scuderi, A space-time BIE methods for nonhomogeneous exterior wave equation problems. The Dirichlet case, *IMA J. Numer. Anal.*, 32(1) (2012), pp. 202–226.

[3] C. Lubich, On the multistep time discretization of linear initial-boundary value problems and their boundary integral equations, *Numer. Math.*, 67(3) (1994), pp. 365–389.

Radiation boundary conditions for waves: extensions and applications

Thomas Hagstrom *(Southern Methodist University)*

The radiation of energy to the far field is a central feature of wave physics. As such efficient, convergent domain truncation algorithms are a necessary component of any wave simulation software. For the scalar wave equation and equivalent systems such as Maxwell's equations in a uniform dielectric medium, Complete Radiation Boundary Conditions (CRBC), which are optimized local radiation boundary condition sequences, provide a satisfactory solution: spectral convergence, rapid parameter selection based on sharp *a priori* error estimates, and effective corner/edge closures. Issues that arise in extending the method to more general problems include the treatment of so-called reverse modes, waves whose group and phase velocities are misaligned relative to the normal direction at the radiation boundary, as well as problems with inhomogeneities and nonlinearities. Here we will explore a number of ideas for constructing reliable and efficient domain truncation algorithms in these more difficult settings:

- Modifications of the CRBC recursions to handle reverse modes - this approach is successful for problems with single wave families;
- Rigorous *a priori* error analysis and optimization of *ad hoc* damping/stretching layers;
- Application of kernel compression/reduced order modeling to nonlocal formulations.

Specific applications to dispersive models of electromagnetic waves, the elastic wave equation, as well as waves in inhomogeneous media will be given.

Analyzing wave scattering problems in layered

media by using perfectly matched layers

Wangtao Lu *(Zhejiang University)*

In this talk, I will report some recent progress on using PMLs (Perfectly Matched Layers) for solving wave scattering problems in layered media.

The half-space matching method and its application to wave scattering in elastic plates

Yohanes Tjandrawidjaja & Vahan Baronian & Anne-Sophie Bonnet-Ben Dhia & Sonia Fliss
POEMS (CNRS - INRIA - ENSTA Paristech)

We have developed a new method for scattering problems in anisotropic media, where usual numerical methods are either too expensive or even not applicable. This so-called Half-Space Matching method has been first developed and validated in 2D. It consists in coupling several plane-wave representations of the solution in half-spaces surrounding the scatterer with a Finite Element computation of the solution around the scatterer. To ensure that all these representations match in the infinite intersections of the half-spaces, the traces of the solution on the edges of the half-planes are linked by Fourier-integral equations. In the case of a dissipative medium, the continuous problem is proved to be coercive plus compact, and the convergence of the discretization is ensured.

The method has been extended to the 3D case, for an application to non-destructive testing. The objective is to simulate the interaction of Lamb waves with a defect in an anisotropic elastic plate. The additional complexity compared to the 2D case lies on the representations which are obtained semi-analytically by decomposition on Lamb modes. In addition, the system of equations couples the FE representation in the bounded perturbed domain with not only the displacement, but also the normal stress of the solution on the infinite bands limiting the half-plates. The first preliminary result has been obtained in the isotropic case.

Coupling methods for elliptic problems in the unbounded domain

Liwei Xu *(University of Electronic Science and Technology of China)*

It is known that the coupling of finite element methods and boundary integral equation methods is an efficient method for the numerical solution of elliptic problems in the unbounded domain. In this talk, we first give a literature review on the coupling method of finite element methods and boundary element methods. Then we present a new analysis result on one kind of coupling, the so-called DtN-FEM method, and present a new coupling technique of finite element methods and boundary integral equation methods, the so-called DtD-FEM method. Numerical results will be presented to demonstrate the accuracy of the methods.

Optimal control in a bounded domain for wave propagating in the whole space

Wei Gong *(Chinese Academy of Sciences)* & Buyang Li *(The Hong Kong Polytechnic University)* & **Huanhuan Yang** *(Shantou University)*

We consider an optimal control problem in a bounded-domain Ω_0 under the constraint of a wave equation in the whole space. The problem is regularized and then reformulated as an initial-boundary value problem of the wave equation in a bounded domain $\Omega \supset \overline{\Omega}_0$ coupled with a set of boundary integral equations on $\partial\Omega$ taking account of wave propagation through the boundary. The well-posedness and stability of the reformulated problem are proved. A fully discrete finite element method is proposed for solving the reformulated problem. In particular, the wave equation in the bounded domain is discretized by an averaged central difference method in time, and the boundary integral equations are discretized in time by using the convolution quadrature generated by the second-order backward difference formula. The finite and boundary element methods are used for spatial discretization of the wave equation and the boundary integral equations, respectively. We prove the stability and convergence of the numerical method, finally validate the spatial and temporal convergence rates numerically in 2D.

Rational approximation of the square root function in the complex plane

Chunxiong Zheng *(University College London)*

The fast algorithm for the diffusion equation and the Schrödinger equation involves the efficient rational approximation of the square root function in the complex plane. Many methods have been proposed to solve this issue, and among them, Heron's algorithm seems to be very promising. In this talk, we make an error analysis for Heron's algorithm in suitable subdomains of

the complex plane. As an application, we consider the fast algorithm for the artificial boundary condition of the one-dimensional convection diffusion equation in unbounded domains.

Minisymposium M12

Rational Approximation
Organisers
Evan Gawlik and Yuji Nakatsukasa

Rational minimax iterations for computing fractional powers of matrices

Evan S. Gawlik *(University of Hawaii at Manoa)*

We construct a family of iterations for computing the principal square root of a square matrix A using Zolotarev's rational minimax approximants of the square root function. We show that these rational functions obey a recursion, allowing one to iteratively generate optimal rational approximants of $\sqrt{z}$ of high degree using compositions and products of low-degree rational functions. The corresponding iterations for the matrix square root converge to $A^{1/2}$ for any input matrix A having no nonpositive real eigenvalues. In special limiting cases, these iterations reduce to known iterations for the matrix square root: the lowest-order version is an optimally scaled Newton iteration, and for certain parameter choices, the principal family of Padé iterations is recovered. Theoretical results and numerical experiments indicate that the iterations perform especially well on matrices having eigenvalues with widely varying magnitudes.

After deriving iterations for the matrix square root, we generalize the above construction by deriving rational minimax iterations for the matrix p^{th} root, where $p > 2$ is an integer. The analysis of these iterations is considerably different from the case $p = 2$, owing to the fact that when $p > 2$, rational minimax approximants of the function $z^{1/p}$ do not obey a recursion. Nevertheless, we show that several of the salient features of the Zolotarev iterations for the matrix square root, including equioscillatory error, order of convergence, and stability, carry over to case $p > 2$. A key role in the analysis is played by the asymptotic behavior of rational minimax approximants on short intervals.

Approximating the pth root by composite rational functions

Evan S. Gawlik & **Yuji Nakatsukasa** *(Oxford University)*

A landmark result from rational approximation theory states that $x^{1/p}$ on $[0, 1]$ can be approximated by a type-(n, n) rational function with root-exponential accuracy. Motivated by the recursive optimality property of Zolotarev functions (for the square root and sign functions), we investigate approximating $x^{1/p}$ by composite rational functions of the form $r_k(x, r_{k-1}(x, r_{k-2}(\cdots(x, r_1(x, 1)))))$. While this class of rational functions ceases to contain the minimax (best) approximant for $p \geq 3$, we show that it roughly achieves pth-root exponential convergence with respect to the degree. Moreover, crucially, the convergence is *double exponential* with respect to the number of degrees of freedom, suggesting that composite rational functions are able to approximate $x^{1/p}$ and related functions (such as $|x|$ and the sector function) with exceptional efficiency.

Data-driven model order reduction for Rayleigh-damped second-order systems

Igor Pontes Duff & Pawan Goyal & Peter Benner *(Max Planck Institute for Dynamics of Complex Technical Systems)*

In this talk, we present a data-driven approach to identify second-order systems of the form

$$\begin{aligned} \mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{D}\dot{\mathbf{x}}(t) + \mathbf{K}\mathbf{x}(t) &= \mathbf{B}\mathbf{u}(t), \\ \mathbf{x}(0) = 0,\ \dot{\mathbf{x}}(0) = 0, & \\ \mathbf{y}(t) &= \mathbf{C}\mathbf{x}(t), \end{aligned}$$

with internal Rayleigh damping. This mean that the damping matrix is given as a linear combination of the mass and stiffness matrices, i.e., $\mathbf{D} = \alpha\mathbf{M} + \beta\mathbf{K}$. These systems typically appear in order to perform various engineering studies, e.g., vibrational analysis. The frequency response of the system can be given by the following rational structured function:

$$\mathbf{H}(s) = \mathbf{C}\left(s^2\mathbf{M} + s\mathbf{D} + \mathbf{K}\right)^{-1}\mathbf{B},$$

which is also known as the transfer function. In an experimental setup, the frequency response of a system can be measured via various approaches. As a consequence, given frequency samples, the identification of the underlying system relies on rational structured approximation. To that aim, we propose an interpolation-based approach, extending the Loewner framework for this class of systems. The efficiency of the proposed method is demonstrated by means of

various numerical benchmarks.

Minisymposium M13

Hierarchical methods in stochastic numerical approximations
Organiser
Abdul-Lateef Haji-Ali (Heriot-Watt)

Multilevel methods for fast Bayesian optimal experimental design

Joakim Beck & Ben Mansour Dia & Luis Espath & Raúl Tempone *King Abdullah University of Science and Technology (KAUST)*

We consider the problem of efficiently estimating the expected Shannon information gain between prior and posterior probability distributions, which is a nested integral and appears as a design criterion in Bayesian optimal experimental design for nonlinear models. In the context of partial differential equations with random input data, multilevel methods substantially reduce the computational work complexity of their single-level counterparts. In this talk, we first present a multilevel double loop Monte Carlo (MLDLMC) estimator and the associated optimal work complexity. The optimal values for the parameters of the MLDLMC estimator are determined by minimizing the average computational work subject to some specified error tolerance that should be satisfied with some desired probability. Then, we present a multilevel double loop stochastic collocation (MLDLSC) method for high-dimensional integration by deterministic quadrature on sparse grids. The methods are demonstrated numerically for the design of electrical impedance tomography experiments with the experimental goal being to infer fiber orientations in composite laminate materials.

Multilevel Monte Carlo using approximate distributions

Mike Giles & Oliver Sheridan-Methven *(University of Oxford)*

This talk will begin with a recap of the ideas and analysis behind the multilevel Monte Carlo (MLMC) method [1, 2]. It will then discuss a new application of the method, building on prior ideas developed with Hefter, Mayer and Ritter [3, 4, 5] to use different approximations of the inverse Normal CDF to generate increments of a Brownian motion.

The advantage of this is that the new approximations can be computed very efficiently, and the errors they introduce can be eliminated through a nested use of MLMC. Numerical results will show the effectiveness of the approach, and also the additional benefits which come from using reduced precision computer arithmetic. The supporting numerical analysis will be outlined briefly.

References

[1] M.B. Giles. Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617, 2008.

[2] M.B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.

[3] M.B. Giles, M. Hefter, L. Mayer, and K. Ritter. Random bit quadrature and approximation of distributions on Hilbert spaces. *Foundations of Computational Mathematics*, 19(1):205–238, 2019.

[4] M.B. Giles, M. Hefter, L. Mayer, and K. Ritter. Random bit multilevel algorithms for stochastic differential equations. *Journal of Complexity*, to appear, 2019.

[5] M.B. Giles, M. Hefter, L. Mayer, and K. Ritter. An adaptive random bit multilevel algorithm for SDEs. In *Multivariate Algorithms and Information-Based Complexity*. de Gruyter, 2019.

Inference with multilevel Monte Carlo

Kody J. H. Law & Ajay Jasra *(University of Manchester)*

Bayesian inference provides a principled and well-defined approach to the integration of data into an a priori known distribution, resulting in a posterior distribution. This talk will concern the problem of inference when the posterior involves continuous models which require approximation before inference can be performed. Typically one cannot sample from the posterior distribution directly, but can at best only evaluate it, up to a normalizing constant. Therefore one must resort to computationally-intensive inference algorithms in order to construct estimators. These algorithms are typically of Monte Carlo type, and include for example Markov chain Monte Carlo, importance samplers, and sequential Monte Carlo samplers. The multilevel Monte Carlo method provides a way of op-

timally balancing discretization and sampling error on a hierarchy of approximation levels, such that cost is optimized. Recently this method has been applied to computationally intensive inference. This non-trivial task can be achieved in a variety of ways. This talk will review 3 primary strategies which have been successfully employed to achieve optimal (or canonical) convergence rates ? in other words faster convergence than i.i.d. sampling at the finest discretization level. Some of the specific resulting algorithms, and applications, will also be presented.

Abstracts of Contributed Talks

Numerical methods for PDE-constrained optimization problems in particle dynamics

Mildred Aduamoah & John Pearson & Ben Goddard *(University of Edinburgh)*

A vast number of important applications in mathematics and engineering are governed by mathematical optimization problems. One crucial class of these problems that have gained significant recent attention are PDE-constrained optimization problems, which can be used to model problems from fluid flow, industrial processes, and biological mechanisms. However, to date, relatively little has been done by way of devising a systematic approach to tackle multiscale PDE-constrained optimization problems arising from particle dynamics. In this talk, we will introduce particle dynamics optimization models, discuss the real life applications they arise from, and explain the new numerical methods that we have developed to solve them.

Multilevel numerical algorithms for thin film flow

Mashael Aljohani & Dr Mark Walkley & Professor Peter Jimack *(University of Leeds)*

The thin film flow equations that we consider in this research are based upon a long-wave approximation to the Stokes equations. This results in a nonlinear parabolic system of two partial differential equations (PDEs) for two dependent variables (film thickness and pressure). When discretised implicitly in time, and using a standard finite difference stencil in space, this results in a large system of nonlinear algebraic equations to be solved at each time step.

For suitable linear equation systems the use of multilevel algorithms is well-established as a means of obtaining optimal (i.e.linear) time complexity. For some nonlinear equation systems multilevel algorithms may also be employed, and can also deliver optimal time complexity (or very close to optimal). However, there appears to be little research that seeks to systematically investigate which multilevel approach is the most effective for practical problems. In this work we consider three nonlinear multilevel algorithms to solve the discretised thin film flow PDEs: the Full Approximation Scheme (FAS), Newton-Multigrid (Newton-MG) and a new preconditioned Newton-Krylov variant.

Each of the nonlinear multilevel solvers will be described, and will be demonstrated to show linear (or close to linear) time complexity. For the Newton-Krylov algorithm, we will introduce our novel block-preconditioner based on the use of a multilevel algebraic multigrid (AMG) package. We will then present systematic numerical results to compare the three multilevel algorithms to show that, while all three are optimal, there are differences in performance that demonstrate the superiority of the new preconditioned Newton-Krylov approach.

Computational paradoxes in deep learning

Alexander Bastounis & Anders Hansen & Verner Vlacic *(University of Cambridge)*

One of the biggest problems in modern AI is that neural networks based on deep learning become highly unstable, and currently there does not exist any cure for this problem. The instabilities are typically so severe that images that look identical to the human eye may be classified with wildly different labels like a cat and a fire truck. In this talk we discuss theoretical paradoxical results in deep learning that shed light on the mysterious behaviour of modern AI.

1. Any given algorithm for computing the optimal neural network given a training set will fail to find the best neural network with high probability on uncountably many training sets, even if such an algorithm is given access to an oracle that allows the algorithm to avoid local but not global minima.

2. For a given arbitrarily small distance $\epsilon > 0$, there exists uncountably many classification problems (all of which are stable under ϵ perturbations) for which one can construct a neural network which correctly identifies an (arbitrarily large) training set and an (arbitrarily large) test set but such a neural network must be unstable under perturbations of size ϵ.

3. For any given $\epsilon > 0$, one can construct problems for which solving a neural network optimisation problem will yield an incorrect classification on a member of the training set, however, if one perturbs the training set by at most ϵ, solving the optimisation problem will yield a neural network that succeeds on the new training set.

The first of these paradoxes highlights an issue with the current understanding of neural networks – namely, the universal approximation theorem guarantees the

existence of a neural network for a given classification problem, but gives no method to find it and in fact our result shows that in general this is not possible.

The second paradox provides a theoretical demonstration of adversarial examples: that is, examples where a small change in the data yields to a wildly different output of the neural network. In this case, the training process has forced the choice of an unstable neural network.

Finally, the third paradox shows that quantifying the success of neural networks is a difficult task – whether or not a neural network will be successful can change drastically depending on arbitrarily small perturbations of the training set.

Convergence analysis of adaptive stochastic Galerkin FEM for elliptic PDEs with parametric uncertainty

Alex Bespalov [1] & Dirk Praetorius [2] & Leonardo Rocchi [1] & Michele Ruggeri [2] [1] *(University of Birmingham)* [2] *(Vienna University of Technology)*

We study adaptive stochastic Galerkin finite element approximations for a class of elliptic partial differential equations with parametric or uncertain coefficients. The proposed adaptive algorithm is steered by a novel reliable and efficient a posteriori error estimator that combines a two-level spatial estimator and a hierarchical parametric estimator, see [1]. The structure of the estimator is exploited in the algorithm to perform a balanced adaptive refinement of the spatial and parametric components of Galerkin approximations. In this talk, we present recent convergence results for this adaptive refinement procedure [2]. In particular, we show that the proposed algorithm drives the underlying energy error estimate to zero. We also report the results of numerical experiments that focus on comparison of the computational cost associated with employing different marking criteria for adaptive refinement.

References

[1] A. Bespalov, D. Praetorius, L. Rocchi, and M. Ruggeri, Goal-oriented error estimation and adaptivity for elliptic PDEs with parametric or uncertain inputs, *Comput. Methods Appl. Mech. Engrg.*, 345, pp. 951–982, 2019.

[2] A. Bespalov, D. Praetorius, L. Rocchi, and M. Ruggeri, Convergence of adaptive stochastic Galerkin FEM, *Preprint*, arXiv:1811.09462, 2018.

High order compact finite difference schemes on nonuniform grids for nonstationary PDEs

Raimondas Čiegis *(Vilnius Gediminas Technical University)*

In this talk we consider high-order compact finite difference schemes constructed on 1D non-uniform grids. We apply them to the parabolic equation:

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} - q(x)u + f(x,t),$$

and the Schrödinger equation:

$$\frac{\partial u}{\partial t} - i\Big(\frac{\partial^2 u}{\partial x^2} - q(x)u\Big) = 0,$$

where t and x are time and space variables, $u(x,t)$ is a real valued function for the parabolic equation and a complex valued function for the Schrödinger equation.

We consider a high-order compact semi-discrete finite difference scheme on a non-uniform grid defined on the three-point template for the parabolic problem

$$\widetilde{B}^h\big(U_{\bar{t}}^n + qU^{n-1/2} - f^{n-1/2}\big) = A_D^h U^{n-1/2},$$

and for the Schrödinger problem

$$\widetilde{B}^h\big(U_{\bar{t}}^n + iqU^{n-1/2}\big) = iA_D^h U^{n-1/2}.$$

Here the operator $\widetilde{B}_h$ is defined as

$$\widetilde{B}^h U_j := \alpha_j U_{j-1} + (1-\alpha_j-\beta_j)U_j + \beta_j U_{j+1}.$$

Stability of these schemes is investigated by using the spectral method. Computer experiments are applied in order to find critical grids for which the stability condition is violated.

We have proved that for the Schrödinger problem the proposed scheme is not A-stable. Examples of critical non-uniform grids are obtained. The analysis of ρ-stability is presented, it shows that ρ-stability can be guaranteed only if the time step is not too small $\tau \geq k_0(h)$. This nonclassical inverse stability condition regularizes the discrete scheme, but reduces the convergence accuracy to the second order.

A similar stability analysis is done also for the parabolic problem. It is interesting to note that in this case the large scale computational experiments have not produced any critical non-uniform grid when the stability condition $\operatorname{Re}\lambda_j < 0$ is violated for some eigenvalue.

In addition to FDS a compact finite element scheme with quadratic elements is constructed. This scheme can be seen as a high-order finite difference scheme, if values of the solution are considered only on the grid points. Basic stability and energy conservation properties of the obtained discrete scheme are obtained.

Numerical examples supporting our theoretical analysis are provided and discussed.

The MRE inverse problem for the elastic shear modulus

Penny J Davies & Eric Barnhill & Ingolf Sack *(University of Strathclyde)*

Magnetic resonance elastography (MRE) is a powerful technique for noninvasive determination of the biomechanical properties of tissue, with important applications in disease diagnosis. A typical experimental scenario is to induce waves in the tissue by time-harmonic external mechanical oscillation and then measure the tissue's displacement at fixed spatial positions 8 times during a complete time-period, extracting the dominant frequency signal from the discrete Fourier transform in time. Accurate reconstruction of the tissue's elastic moduli from MRE data is a challenging inverse problem, and I will describe a new approach based on combining approximations at different frequencies into a single overdetermined system.

Some highlights from NLAFET WP3 on parallel sparse direct solution

Iain Duff & Sébastien Cayrols & Florent Lopez & Stojce Nakov *(STFC Rutherford Appleton Laboratory and Cerfacs)*

The FET-HPC EU Horizon 2020 Project NLAFET just ended at the end of April. STFC-RAL was the lead partner for Workpackage 3 on sparse direct methods. We discuss some highlights from this work.

While it is recognized that direct methods are very powerful over a wide range of applications, they sometimes have issues on extremely large problems, particularly in 3 dimensions. Our approaches overcome this limitation by subdividing the problem through partitioning the matrix. The direct solver is then used on a submatrix of the original problem.

We illustrate this in two quite different ways. One is by preordering to singly bordered block diagonal form and the other is using an approach based on the block Cimmino algorithm.

The added bonus is that such a subdivision also introduces another level of parallelism and we are then able to combine distributed memory parallelism using MPI with shared memory parallelism using OpenMP or other runtime systems. We illustrate this with runs of our codes on industrial strength large test problems on a heterogeneous platform where we show that our codes, which are available on github, outperform other state-of-the-art codes.

Approximation of time domain boundary integral equations

Dugald B Duncan *(Heriot-Watt University)*

Time domain boundary integral equations (TDBIEs) are used to model time dependent wave scattering from surfaces in an infinite, homogeneous medium. We concentrate on the acoustic scattering case. The solution can be reconstructed anywhere in space and time in terms of a potential which is computed only on the surface. The approximation of the time dependent surface potential by numerical methods is expensive and sensitive to stability issues, but it can be cheaper than solving the wave equation in a very large space domain. We describe a practical implementation of a full space-time Galerkin TDBIE method derived and analysed by Ha Duong, and its connection with a simpler method, derived by Davies and Duncan, using "backward-in-time" collocation coupled with Galerkin in space.

Preconditioning techniques for the Darcy-Forchheimer model using Block-Centered finite difference method

Faisal A. Fairag & Mohsen G. Alshahrani *(King Fahd University of Petroleum and Minerals)*

We consider the Darcy-Forchheimer problem in two dimensions describing the fluid flow such as oil and gas in porous media

$$\begin{aligned} &\mu k^{-1}u + \beta\rho|u|u + \nabla p = 0, \quad \text{in } \Omega, \\ &\nabla \cdot u = f, \quad \text{in } \Omega, \quad u \cdot n = 0, \quad \text{on } \partial\Omega, \end{aligned} \tag{1}$$

with the compatibility condition $\int_\Omega f\, dxdy = 0$. The unknowns are the pressure p and the velocity $u = (u_1, u_2)$. Here μ, k and β denote the viscosity coefficient, the permeability tensor and the Forchheimer

number, respectively. n denotes the outward normal unit vector to $\partial\Omega$ and $|.|$ denotes the Euclidean norm. The function $f \in L^2(\Omega)$ and $\Omega \in R^2$. If there exist positive constants C_1 and C_2 such that: $C_1 \leq \frac{\mu}{k} \leq C_2$, $C_1 \leq \beta\rho \leq C_2$, then (1) is well-posed. If $\beta = 0$, then (1) becomes the Darcy's equation which describes a linear relation between the velocity and the pressure gradient.

Block-Centered Finite Difference discretization (BCFDM) and Newton-Raphson method yield a sparse and indefinite linear system of the form

$$\left[\begin{array}{cc|c} D_1 & M_3^T & I \otimes E^T \\ M_3 & D_2 & E^T \otimes I \\ \hline I \otimes E & E \otimes I & 0 \end{array}\right] \left[\begin{array}{c} U_1 \\ U_2 \\ P \end{array}\right] = \left[\begin{array}{c} b_1 \\ b_2 \\ b_3 \end{array}\right],$$

$$E = \begin{bmatrix} -1 & & \\ 1 & \ddots & \\ & \ddots & -1 \\ & & 1 \end{bmatrix} \tag{2}$$

where D_1 and D_2 are diagonal matrices and $\otimes$ denotes the Kronecker product. The focus of this study is to study the efficient solution of the linear system (2). For large problems with a sparse coefficient matrix, iterative methods are preferable. In this talk, we will investigate the performance of a preconditioned MINRES method which utilises the structure of the coefficient matrix (2).

Acknowledgement:

The authors acknowledge the support provided by King Abdulaziz City for Science and Technology (KACST) (project no. 1-17-003-0002) and KFUPM.

Two-step implicit higher order numerical integrator for stiff systems of initial – boundary value problems of ordinary differential equations

Johnson O. Fatokun* & Samuel I. Okoro *(Anchor University)*

In this paper, we consider the derivation of a Two-Step Implicit Higher Order Numerical Integrator for Stiff Systems of Ordinary Differential Equations. The exponentially fitted numerical method is developed using the method of trapezoidal interpolant for the numerical integration for stiff system. This method preserves the A-stability property of the numerical scheme and is also L-stable. Two theorems and one Lemma were proposed and proved which establish the A-stability and L-stability properties of the derived method. The local truncation error of the method is estimated from the continuous form of the method derived and presented. The analysis and numerical experiments show clearly that the method competes favourably with other known methods when applied to stiff systems of initial value problems of ordinary differential equations. The method is of at least order four (Oh4) in accuracy. Key words: Stiff systems, trapezoidal integrators, A-Stability, L-Stability, Implicit two-step method.

AMS Classification: 65L05

An adaptive primal-dual framework for nonsmooth convex minimization

Olivier Fercoq & Quoc Tran-Dinh & Ahmet Alacaoglu & Volkan Cevher *(LTCI, Télécom ParisTech, Université Paris-Saclay)*

We propose a new self-adaptive, double-loop smoothing algorithm to solve composite, nonsmooth, and constrained convex optimization problems.

Our algorithm is based on Nesterov's smoothing technique via general Bregman distance functions. It self-adaptively selects the number of iterations in the inner loop to achieve a desired complexity bound without requiring the accuracy a priori as in variants of Augmented Lagrangian methods (ALM). We prove $\mathcal{O}(\frac{1}{k})$-convergence rate on the last iterate of the outer sequence for both unconstrained and constrained settings in contrast to ergodic rates which are common in ALM as well as alternating direction method-of-multipliers literature.

Compared to existing inexact ALM or quadratic penalty methods, our analysis does not rely on the worst-case bounds of the subproblem solved by the inner loop. Therefore, our algorithm can be viewed as a restarting technique applied to the ASGARD method (Tran-Dinh et al, 2018) but with rigorous theoretical guarantees or as an inexact ALM with explicit inner loop termination rules and adaptive parameters. Our algorithm only requires to initialize the parameters once, and automatically updates them during the iteration process without tuning.

We illustrate the superiority of our methods via several examples as compared to the state-of-the-art.

Wilkinson test matrices

Carla Ferreira & Beresford Parlett *(University of Minho)*

In the 1950', J. H. Wilkinson introduced two families of symmetric tridiagonal integer matrices $W^{\pm}_{2m+1}$, for each positive integer m. All the next-to-diagonal entries are one. The diagonal entries are given by

$$\begin{aligned} & diag(W^{\pm}_{2m+1}) \\ = \; & diag(m, m-1, \ldots, 1, 0, \pm 1, \pm 2, \ldots, \pm m). \end{aligned}$$

Each matrix illustrated subtle difficulties in the automatic computation of eigenvectors and eigenvalues.

Most of the eigenvalues of $W^{\pm}_{2m+1}$ are close to diagonal entries. We developed the structure of their eigenvectors in a natural way which reveals that the envelopes of these eigenvectors all look the same to the naked eye. The shape is a badly dented bell curve. We also analyzed the eigenvectors of the remaining non-integer eigenvalues.

Variational time integration for evolutionary systems

S. Franz *(TU Dresden)*

Evolutionary systems are time dependent systems of first order partial differential equations. For a Hilbert space H and a weighted space-time Hilbert space given by

$$H_\rho(\mathbb{R}; H) : \Big\{ f : \mathbb{R} \to H \; : \; f \text{ meas.}, \int_R, \langle f(t), f(t) \rangle_H \mathrm{e}^{-2\rho t} \, \mathrm{d}t < \infty \Big\}$$

with an associated inner product

$$\langle f, g \rangle_\rho : \int_{\mathbb{R}} \langle f(t), g(t) \rangle_H \mathrm{e}^{-2\rho t} \, \mathrm{d}t$$

their simplest form is

$$(\partial_t M_0 + M_1 + A)U = F,$$

where

- M_0, M_1 are bounded, linear, self-adjoint operators on H,
- A is an unbounded, skew-selfadjoint operator on H,
- $\exists \rho_0 > 0$, $\gamma > 0$ such that for all $\rho \geq \rho_0$, $x \in H$ it holds $\langle (\rho M_0 + M_1)x, x \rangle_\rho \geq \gamma \langle x, x \rangle_\rho$.

Then the system has for each $\rho \geq \rho_0$ and $F \in H_\rho(\mathbb{R}, H)$ a unique solution $U \in H_\rho(\mathbb{R}, H)$, see [1], and it holds

$$\|U\|_\rho \leq \frac{1}{\gamma} \|F\|_\rho.$$

Examples of such systems are wave-equations, heat-equations, Maxwell-equations, linear elasticity, acoustic equations, combinations thereof and many many more.

In this talk we investigate the numerical approximation of these systems by using variational time integration methods like discontinuous Galerkin and continuous-Petrov-Galerkin. As examples we consider problems, where the type of the partial differential equation changes in the domain, see [2], or the homogenisation of rapidly changing coefficients, see [3, 4].

References

[1] R. Picard, A structural observation for linear material laws in classical mathematical physics, *Math. Methods Appl. Sci.*, **32**, 1768–1803, 2009.

[2] S. Franz, S. Trostorff, and M. Waurick, Numerical methods for changing type systems, *IMAJNA*, 2018, doi.org/10.1093/imanum/dry007.

[3] S. Franz and M. Waurick, Resolvent estimates and numerical implementation for the homogenisation of one-dimensional periodic mixed type problems, *ZAMM*, 98(7):1284–1294, 2018.

[4] S. Franz and M. Waurick, Homogenisation of parabolic/hyperbolic media, accepted for publication in BAIL2018 Proceedings, 2018.

A new two-step stable high accuracy implicit method for general second order nonlinear initial-value problems on a graded mesh

Bishnu Pada Ghosh *(South Asian University, India)*

A new two-step method of order three on a graded mesh for the numerical solution of general second order nonlinear initial-value problems $u'' = f(t, u, u')$, $u(t_0) = \gamma_0, u'(t_0) = \gamma_1$ is discussed. In practice, only a monotonically decreasing mesh will be employed. When applied to a test equation $u'' + 2\alpha u' + \beta^2 u = g(t), \alpha > \beta \geq 0 \geq$, the derived method is absolutely stable and superstable for a uniform mesh. The proposed method is applicable to both singular problems and problems containing boundary layers that arise

close to the point where one wishes to compute. Computational results are presented to demonstrate the validity of the proposed method. The numerical results clearly indicate that the proposed method produces better results in comparison with the existing methods.

Domain decomposition preconditioners for heterogeneous Helmholtz problems

Shihua Gong & Ivan G. Graham & Euan A. Spence *(University of Bath)*

We consider one-level additive Schwarz preconditioners for a family of Helmholtz problems of increasing difficulty characterized by wave number k. As k increases, the solution becomes more oscillatory and thus meshes need to be increasingly refined, leading to huge linear systems. Moreover, these linear systems become more indefinite and the minors of the linear systems may be even singular, which prevents the usage of many standard preconditioning techniques. We discuss additive Schwarz preconditioners, of which the action requires solution of independent subproblems with absorbing boundary conditions and (possibly) absorption added in the domain. Then the local solutions are glued together using prolongation operators defined by a partition of unity and nodal interpolation. Supporting theory for this preconditioner in the case of homogeneous Helmholtz problems is given by I. G. Graham, E. A. Spence and J. Zou in [Preprint arXiv:1806.03731.]. In this talk, we will present preliminary theoretical results and numerical experiments for this method applied to heterogeneous Helmholtz problems. We also show that the method is independent of the underlying finite element discretizations.

Analysis of circulant embedding methods for sampling stationary random fields

Ivan Graham *(University of Bath)*

This talk contains joint work with Markus Bachmayr, Frances Kuo, Van Kien Nguyen, Dirk Nuyens, Rob Scheichl and Ian Sloan ([1],[2]).

A standard problem in uncertainty quantification and in computational statistics is the rapid sam- pling of stationary Gaussian random fields with given covariance in a d-dimensional (physical) domain. In many applications it is sufficient to perform the sampling on a regular grid on a cube enclosing the physical domain, in which case the corresponding covariance matrix is nested block Toeplitz. After extension to a nested block circulant matrix, this can be diagonalised by FFT and, provided the circulant is positive definite, this provides a finite expansion of the field in terms of a deterministic basis, with coefficients given by i.i.d. standard normals. This is a discrete counter- part of the well-used Karhunen-Loève expansion, but (because of the FFT) it can be computed in log-optimal time.

Questions at the interface of analysis and numerical linear algebra which arise from this process are: (i) What is the most efficient way to extend the block Toeplitz matrix to a block circulant, so that positive definiteness is maintained and (ii) What is the decay rate of the eigenvalues of the block circulant (since this determines the efficiency of several UQ algorithms which use this type of expansion).

We give an overview of the answers to these questions, both for classical (non-smooth) circulant embedding [1] and for circulant embedding using smooth cut-off functions [2].

References

[1] I.G. Graham, F.Y. Kuo, D. Nuyens, R. Scheichl, and I.H. Sloan, Analysis of circulant embed- ding methods for sampling stationary random fields, SIAM J. Numer. Anal. 56(3), 1871 1895, 2018.

[2] M. Bachmayr, I.G. Graham, V.K. Nguyen and R. Scheichl, Unified analysis of periodisation- based sampling methods for Matérn covariances, in preparation.

Bayesian network PDEs for multiscale representations of porous materials

Eric Hall *(RWTH Aachen University)* & Kimoon Um *(Stanford University)* & Markos Katsoulakis *(University of Massachusetts Amherst)* & Daniel Tartakovsky *(Stanford University)*

Microscopic (pore-scale) properties of porous media affect and often determine their macroscopic (continuum- or Darcy-scale) counterparts. Understanding the relationship between processes on these two scales is essential to both the derivation of macroscopic models of, e.g., transport phenomena in natural porous media, and the design of novel materials, e.g., for energy storage. Microscopic properties exhibit complex statistical correlations and geometric constraints that present challenges for the estimation of macroscopic quantities of interest (QoIs), e.g., in the context of

global sensitivity analysis (GSA) of macroscopic QoIs with respect to microscopic material properties. We present a systematic way of building correlations into stochastic multiscale models through Bayesian Networks. The proposed framework allows us to construct the joint probability density function (PDF) of model parameters through causal relationships that are informed by domain knowledge and emulate engineering processes, e.g., the design of hierarchical nanoporous materials. These PDFs also serve as input for the forward propagation of parametric uncertainty thereby yielding Bayesian Network PDE. To assess the impact of causal relationships and micro-scale correlations on macroscopic material properties, we propose a moment-independent GSA and corresponding effect rankings for Bayesian Network PDE, based on the differential Mutual Information, that leverage the structure of Bayesian Networks and account for both correlated inputs and complex non-Gaussian QoIs. Our findings from numerical experiments, which feature a non-intrusive uncertainty quantification workflow, indicate two practical outcomes. Firstly, the inclusion of correlations through structured priors based on causal relationships informed by domain knowledge impacts predictions of QoIs which has important implications for engineering design. Secondly, we observe the inclusion of structured priors with non-trivial correlations yields different effect rankings than independent priors and moreover these rankings are more consistent with the anticipated physics of a model problem.

Numerical solution of the Hemker problem

Alan F. Hegarty *(University of Limerick)* & Eugene O'Riordan *(Dublin City University)*

In [1] Hemker proposed a model test problem in two space dimensions, comprising the simple constant coefficient linear singularly perturbed convection-diffusion equation

$$-\varepsilon\triangle u + u_x = 0, \quad \text{for} \quad x^2 + y^2 > 1;$$

on the unbounded domain $\mathbb{R}^2 \setminus \{x^2 + y^2 \leq 1\}$, exterior to the unit circle, with the boundary conditions $u(x, y) = 1$, if $x^2 + y^2 = 1$ and $u(x, y) \to 0$ as $x^2 + y^2 \to \infty$.

[1] notes that the domain is sufficiently complex to allow an exponential boundary layer as well as parabolic interior layers, and that to find a proper mesh is "not trivial at all".

In neighbourhoods of the points on the circle where the characteristics of the reduced problem $\tilde{v}_x = 0$ are tangential to the circle, there are transition regions, where the steep gradients in the solution migrate from the exponential boundary layers (located on the left side of the unit disk) to the parabolic internal layers emerging from the characteristic points $(0, \pm 1)$.

The challenge proposed in [1] is to design a numerical method, which produces stable and accurate approximations over the entire domain, for arbitrary small values of the perturbation parameter ε. We propose here such a method, based on appropriately chosen Shishkin meshes in different coordinate systems.

References

[1] P. W. Hemker, A singularly perturbed model problem for numerical computation, *J. Comp. Appl. Math*, **76**, 1996, 277–285.

On numerical simulations and a posteriori analysis for algebraic flux correction schemes

Abhinav Jha & Volker John *(Freie Universität Berlin, Weierstrass Institute, Berlin)*

Non-linear discretizations are necessary for convection-diffusion-reaction equations for obtaining accurate solutions that satisfy the discrete maximum principle (DMP). Algebraic stabilizations, also known as Algebraic Flux Correction (AFC) schemes, belong to the very few finite element discretizations that satisfy this property. The non-linearity of the system arises because of the presence of solution dependent limiters. Because of the non-linear nature of the problem, a new issue arises, which is the efficient solution of the system of equations.

The first part of the talk will address this issue and several methods will be discussed for solving the non-linear problem with a major focus on fixed point iterations and Newton's method. Different algorithmic components such as Anderson acceleration and dynamic damping will be discussed as well. The methods are compared on different parameters such as the number of iterations and computing time needed. Numerical examples are presented in 2d as well as 3d which will assess the different solvers.

The second part of the talk is devoted to the proposal of a new a posteriori error estimator for the AFC schemes. With a mild assumption on the interpolation error, we find a global upper bound in the energy norm of the system which is independent of the choice of limiters. Numerical results are presented in

2d with two different types of limiters.

Iterative methods for elliptic PDEs based on operator preconditioning

János Karátson *(ELTE University)*

A desirable property of an iterative method for a discretized partial differential equation is mesh independence. Then, if the auxiliary problems can be solved with optimal order of arithmetic operations (i.e. the cost $O(N)$ is proportional to the degrees of freedom), then the overall iteration also has this advantage. Such efficient preconditioners for discretized elliptic problems can be obtained via equivalent operator preconditioning, that is, the preconditioner shall be chosen as the discretization of a suitable auxiliary operator that is equivalent to the original one. This concept was elaborated by T. Manteuffel et al. [4] and then applied in several situations, see, e.g., [6, 7, 8]. The author's work is mainly based on joint research with O. Axelsson, e.g., [1, 2, 3], see also [5] for nonlinear problems.

This talk first gives a brief theoretical summary, including both linear and superlinear mesh independent convergence. Then various applications are shown, such as parallel preconditioning of transport type systems, shifted Laplace preconditioners for Helmholtz equations, and streamline diffusion preconditioning both for linear convection-diffusion equations and for semilinear parabolic systems.

References

[1] Axelsson, O., Karátson, J.,. Equivalent operator preconditioning for elliptic problems, *Numer. Algor.*, 50:297–380 (2009).

[2] Axelsson, O., Karátson, J., Robust preconditioning estimates for convection-dominated elliptic problems via a streamline Poincaré-Friedrichs inequality, *SIAM J. Numer. Anal.*, 52 (2014), No. 6, pp. 2957-2976.

[3] Axelsson, O., Karátson, J., Superlinear convergence of the GMRES for PDE-constrained optimization problems, *Numer. Funct. Anal. Optim.* 39 (2018), no. 9, 921–936.

[4] Faber, V., Manteuffel, T., Parter, S.V., On the theory of equivalent operators and applications to the numerical solution of uniformly elliptic partial differential equations, *Adv. in Appl. Math.,* 11 (1990), 109-163.

[5] Faragó I., Karátson J., *Numerical Solution of Nonlinear Elliptic Problems via Preconditioning Operators: Theory and Application.* Advances in Computation, Volume 11, NOVA Science Publishers, New York, 2002.

[6] Kirby R. C., From Functional Analysis to Iterative Methods, *SIAM Review*, 52(2) 269-293 (2010).

[7] Loghin, D., Wathen, A. J., Analysis of preconditioners for saddle-point problems, *SIAM J. Sci. Comput.* 25 (2004), no. 6, 2029–2049.

[8] Mardal, K-A.,. Winther, R., Preconditioning discretizations of systems of partial differential equations, *Numer. Linear Algebra Appl.* 18 (2011), no. 1, 1–40.

Numerical methods for Cordial Volterra integral equation with vanishing delays

Melusi Khumalo *(University of South Africa)*

In this talk, we consider some numerical approaches to solving Cordial Volterra integral equations (CVIEs) with vanishing delays. A number of authors have established an inextricable link between the compactness of a Cordial Volterra integral operator and the existence and uniqueness of the associated CVIE. The validity of such results are investigated in the case of CVIEs incorporating vanishing delays.

Collocation methods with uniform mesh as well as graded mesh are then implemented on some strategic examples. Convergence rates and existence and uniqueness of the collocation solution are considered.

Structured Sylvester and T-Sylvester equations

Ivana Kuzmanović Ivičić & Ninoslav Truhar *(University of Osijek)*

Sylvester and T-Sylvester equations are matrix equations of the form

$$AX + XB = E$$

and

$$AX + X^T B = E,$$

where A, B and E are given and X is an unknown matrix. Sylvester equations appear frequently in many areas of applied mathematics. For example, Sylvester equations play vital roles in matrix eigen-decom-

positions, control theory, model reduction, numerical solution of matrix differential Riccati equations and algebraic Riccati equations, image processing, and many more. On the other hand, T-Sylvester matrix equations have recently attracted the attention of researchers because of their relationship with palindromic eigenvalue problems.

We will present results about structured Sylvester and T-Sylvester equations, especially for structured problems with the system matrices of the form $A = A_0 + U_1V_1$ and $B = B_0 + U_2V_2$ where U_1, U_2, V_1, V_2 are small rank update matrices. We will give the Sherman-Morrison-Woodbury-type formula for solutions of this type of equation. The obtained formulas are used for construction of algorithms which solve the equations of the above form much more efficiently than the standard algorithms. Application of obtained algorithms will be illustrated in several examples.

This is joint work with Ninoslav Truhar.

Robust solvers for highly heterogeneous Maxwell's equations

Alexandros Kyriakis *(University of Strathclyde)*

Maxwell's equations are the fundamental equations which govern and explain the behaviour of electromagnetic waves. A key topic of interest is the investigation of highly heterogeneous problems with differing conductivities, which vary highly in magnitude along the domain. An analogue to this is Darcy's problem with a highly heterogeneous diffusion coefficient. Our interest lies in the stability and convergence to the solution when building an appropriate coarse space. Emphasis will be given to exposing some of the investigated cases. A few practical test cases will be presented, starting with Darcy's problem, using a two-level domain decomposition method. For each case we use the GenEO coarse space and investigate numerically using the specialised software FreeFem++.

Improving the condition number of sample covariance matrices for use in variational data assimilation

Amos S. Lawless, Jemima M. Tabeart, Sarah L. Dance, Nancy K. Nichols & Joanne A. Waller *(University of Reading & National Centre for Earth Observation)*

Variational data assimilation is used on a daily basis within numerical weather and ocean prediction, combining the latest observational data with numerical models to estimate the current state of the system. It formulates the problem as the minimisation of a weighted nonlinear least-squares objective function, in which the fit to the observations is balanced against the fit to a previous model forecast, with covariance matrices weighting the different sources of information. Until recently most operational forecasting centres have assumed that the errors in the observations are uncorrelated. However, this is not always true, especially when considering multichannel satellite data, and it is becoming more important to specify observation error correlations.

As observation error covariance matrices must usually be obtained by sampling methods, estimates are often degenerate or ill-conditioned, making it impossible to invert an observation error covariance matrix without the use of techniques to reduce its condition number. In this work we present new theory for two existing methods that can be used to 'recondition' any covariance matrix: ridge regression, and the minimum eigenvalue method. We compare these methods with multiplicative variance inflation, which cannot alter the condition number of a matrix, but is often used to account for neglected correlation information. We investigate the impact of reconditioning on variances and correlations of a general covariance matrix in both a theoretical and practical setting. The new theory shows that, for the same target condition number, both methods increase variances compared to the original matrix, with larger increases for ridge regression than the minimum eigenvalue method. We prove that the ridge regression method strictly decreases the absolute value of off-diagonal correlations. Theoretical comparison of the impact of reconditioning and multiplicative variance inflation on the data assimilation objective function shows that variance inflation alters information across all scales uniformly, whereas reconditioning has a larger effect on scales corresponding to smaller eigenvalues.

We then consider two numerical examples: a spatial correlation function, and an observation error covariance matrix arising from satellite interchannel correlations. The minimum eigenvalue method results in smaller overall changes to the correlation matrix than ridge regression, but can increase off-diagonal correlations. We find cases where standard deviations double in size, which will have implications for the solution of a data assimilation problem.

Doubly-stochastic scaling of adjacency matrices for community detection

Luce le Gorrec & Sandrine Mouysset & Daniel Ruiz

(Université de Toulouse)

Community detection in networks consists of finding groups of individuals such that a same group contains elements that are similar (behave similarly, share common features, ...), whereas elements of two different groups are different (have different features, ...). Various modularity measures have been designed to evaluate the correctness of a community structure for a given network. Most of these measures have been designed on simple unweighted networks, and some of them have been extended to be applied on weighted graphs. But it may happen that this generalisation presents a lack of connection with the native philosophy of the measure, making hard to know in what extend it still evaluates the correctness of community structures. We propose to revisit some similarity measures generalisation, or validate an existing one.

There exists a matrix transformation that has properties that may highlight the community structure of a network if applied on the network adjacency matrix: the doubly-stochastic scaling [3]. This scaling is interesting for community detection because:

- It leverages small and big communities by giving to the first higher edge weights.
- If a network has a flow—*i.e.* an imbalance between edges that enter and leave a community—the doubly-stochastic scaling erases the imbalance and highlights the community.
- The scaled graph is 1-regular—all its node degrees are 1.

In our study, we hence use the doubly-stochastic scaling of the adjacency matrix as a preprocessing for the Louvain's algorithm [1]. This algorithm has been proved to be one of the best community detection algorithms in the recent study [6]. A generic version of this algorithm proposed in [2] allows the user to maximise some other modularities. We use this version on some well-known benchmark test cases [5, 4] in order to first validate our modularity measures generalisations, then to compare the efficiency of these measures, and finally to show the added value of the doubly-stochastic scaling

References

[1] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, Fast unfolding of communities in large networks, Journal of Statistical Mechanics: Theory and Experiment, 2008 (2008), p. P10008.

[2] R. Campigotto, P. Conde-Cespedes, and J. Guillaume, A generalized and adaptive method for community detection, CoRR, abs/1406.2518 (2014).

[3] P. A. Knight, D. Ruiz, and B. Uçar, A symmetry preserving algorithm for matrix scaling, SIAM Journal on Matrix Analysis and Applications, 35 (2014), pp. 931955.

[4] A. Lancichinetti and S. Fortunato, Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities, Phys. Rev. E, 80 (2009), p. 016118.

[5] A. Lancichinetti, S. Fortunato, and F. Radicchi, Benchmark graphs for testing community detection algorithms, Physical Review E, 78 (2008), p. 046110.

[6] Z. Yang, R. Algesheimer, and C. Tessone, A comparative analysis of community detection algorithms on artificial networks, Scientific Reports, 6 (2016).

Multigrid-based augmented block-Cimmino method

Philippe Leleux *(CERFACS)* & Dr. Daniel Ruiz *(IRIT)* & Pr.Ulrich Rüde *(FAU, CERFACS)*

The Augmented Block Cimmino Distributed Solver (ABCD Solver[1]) is a hybrid method designed to solve large sparse unsymmetric linear systems of the form: $Ax = b$, where A is a full row rank $m \times n$ sparse matrix, x is a vector of size n and b is a vector of size m. The approach is based on the *block Cimmino row projection method* (**BC**) [2]. BC is applied on the system which is partitioned in row blocks. Convergence rate of BC is known to be slow and we rather solve the symmetric semi-positive definite and consistent system $Hx = k$ obtained when considering the fix point of the iterations. To accelerate the convergence of the block Cimmino method, we solve instead this system using a stabilized block Conjugate Gradient (BCG) [3].

The solver also offers the possibility to construct a larger system with an enlarged matrix $\begin{bmatrix} A & C \\ B & S \end{bmatrix}$ where the numerical orthogonality between partitions is enforced. As a result, the augmented block Cimmino method converges in one single iteration. This results in a pseudo-direct method (**ABCD**) [1] with the solution depending on projections, as in BC, and on the direct solution of a condensed system involving matrix S. Implementation of both ABCD and BC are available in the ABCD Solver package.

[1] http://abcd.enseeiht.fr

For large PDE problems, we investigate extensions of this augmentation approach, in which we relax the strict orthogonality between blocks so as to reduce the size of the augmentation matrices C, B, and S. The purpose is to control better the memory needs and computations. To do so, we exploit ideas from the multigrid framework. Assuming that we have at hand several levels of grids for the system, the original system is then augmented very similarly to the ABCD method, by enforcing the orthogonality between partitions on a coarse level. The result is an augmented system with approximated orthogonality on the fine grid level discretized operator. While in the exact ABCD approach the size of the augmentation can be large for highly connected subdomains, the new approach gives a way to control explicitly this size by choosing coarser levels of grids. This system can be solved with the classical block-Cimmino with fast linear convergence for a wide range of systems coming from PDE problems. We demonstrate the efficiency of the method on Helmholtz and Convection-Diffusion 2D problems.

References

[1] Iain S Duff, Ronan Guivarch, Daniel Ruiz, and Mohamed Zenadi. The augmented block Cimmino distributed method. SIAM Journal on Scientific Computing, 37(3):A1248A1269, 2015.

[2] Tommy Elfving. Block-iterative methods for consistent and inconsistent linear equations. Numerische Mathematik, 35(1):112, 1980.

[3] Daniel Ruiz. Solution of large sparse unsymmetric linear systems with a block iterative method in a multiprocessor environment. CERFACS TH/PA/9, 6, 1992.

Modeling round-off error in the fast gradient method for predictive control

Ian McInerney & Eric C. Kerrigan & George A. Constantinides *(Imperial College London)*

We present a method to determine the smallest data type for the Fast Gradient Method (FGM) to converge when solving a linear Model Predictive Control (MPC) problem in fixed-point arithmetic. We derive two models for the round-off error present in fixed-point arithmetic: a generic model and a structure-exploiting parametric model. We also propose a metric for measuring the amount of round-off error the FGM can tolerate before diverging. Using this metric and the round-off error models, we compute the minimum number of fractional bits needed for the fixed-point data type. We show that the structure-exploiting parametric model nearly halves the number of fractional bits needed to implement an example problem, significantly decreasing the resource usage for an implementation on a Field Programmable Gate Array (FPGA).

Recently, MPC has grown in popularity due to its ability to incorporate operating constraints in the computation of an optimal control action. At its core, MPC solves an optimization problem to determine the next control action, which for the Constrained Linear Quadratic Regulator (CLQR) is a Quadratic Program (QP). This popularity has led to MPC being implemented on smaller processors and FPGAs that do not contain hardware to accelerate floating-point computations.

The lack of floating-point acceleration necessitates the use of fixed-point representation, where the number is stored using a fixed number of bits separated into two segments: integer bits and fraction bits. When implementing the QP solver in fixed-point, the number of fractional bits must be large enough to ensure that the round-off error does not cause the FGM iterations to diverge. Prior work has shown that the convergence of the FGM in fixed-point is dependent upon the eigenvalues of the fixed-point representation of the QP's Hessian lying in the range $(0, 1)$.

In this work we present a new metric, called the rounding stability margin, that uses the pseudospectrum of the Hessian to compute how much round-off error can be introduced before the FGM iterations diverge. Since the Hessian H of the QP is normal, the pseudospectrum can provide direct information about the eigenvalue spectrum. We compute the ϵ-pseudospectrum of H at the points $\lambda \in \{0, 1\}$ to determine how large a perturbation H must undergo to shift its spectrum to include those points. This is then used to bound the spectrum of an additive perturbation matrix on H that represents the round-off errors.

We present two models for the round-off error introduced by moving the Hessian into a fixed-point representation. The first is a generic model that finds the worst-case round-off error for any CLQR problem by using the worst-case perturbation for every element of H. The second is a parametric model that exploits the Toeplitz structure of the Hessian for Schur-stable systems to find the cutoff diagonal beyond which all elements of H round to zero. The round-off error is then modeled as the worst-case for all diagonals before the cutoff, and as the actual element value for all diagonals after the cutoff.

These round-off error models are then used to compute the spectrum of the additive perturbation matrix representing the transformation to fixed-point. Combining these models with the rounding stability margin allows for the computation of the number of fractional bits required for the FGM to be convergent. We demonstrate that exploiting the structure of the MPC problem can reduce the number of fractional bits needed by 30–45%, and allow for a reduction in hardware usage and solution time by up to 77% and 25%, respectively, for an implementation of FGM on a FPGA.

A first-order system method for boundary layer problems

Niall Madden & J.H. Adler & S. MacLachlan & L. Zikatanov *(National University of Ireland Galway)*

We are interested in finite-element solutions of singularly perturbed problems, whose solutions feature boundary layers. There has long been concern that finite-element analyses are not appropriate for such problems, since it is often the case that the energy norms naturally associated with standard Galerkin methods are not strong enough to represent the layers present in the solution. Lin and Stynes (SINUM, 2012) were among the first to fully diagnose the key issue that, for a reaction-diffusion problem, the energy norm is not correctly "balanced". They also proposed a remedy, in the form of a mixed finite element method, for which the associated norm is balanced.

There have been various other attempts to either prove the convergence of solutions obtained with standard methods, but in non-induced norms, or to devise new schemes; see, e.g., Roos (Proc. BAIL 2016) for an overview. Our interest is in the latter approach and, in particular, the method proposed by Adler et al. (IMA J. Numer. Anal., 2016). This has advantages over the Lin and Stynes approach, including that the associated finite-element space is simplified, and that one can apply fast solvers.

In this talk we will examine how that framework can be enhanced so that it may provide further advantages, including being extendable to other singularly perturbed problems, and allowing natural adaptive refinement.

References

[1] Runchang Lin and Martin Stynes. A balanced finite element method for singularly perturbed reaction-diffusion problems. *SIAM J. Numer. Anal.*, 50(5): 2729–2743, 2012.

[2] Roos, Hans-G. Error estimates in balanced norms of finite element methods on layer-adapted meshes for second order reaction-diffusion problems. Proc. BAIL 2016, Lect. Notes Comput. Sci. Eng., Springer, 2017.

[3] James Adler, Scott MacLachlan, and Niall Madden. A first-order system Petrov-Galerkin discretisation for a reaction-diffusion problem on a fitted mesh. *IMA J. Numer. Anal.*, 36 (3):1281-1309, 2016.

Computing the Jordan structure of totally nonnegative matrices

Nicola Mastronardi & Teresa Laudadio & Paul Van Dooren *(Istituto per le Applicazioni del Calcolo "M. Picone")*

Given the factorization of a totally nonnegative matrix A of order n, into the product

$$A = B_1 B_2 \cdots B_{n-2} B_{n-1} D C_{n-1} C_{n-2} \cdots C_2 C_1,$$

with B_i, C_i^T lower bidiagonal totally nonnegative matrices and D diagonal one [1, 2], an algorithm for computing the size of the Jordan block associated to the zero eigenvalue was proposed in [3] with high relative accuracy in floating point arithmetic and $O(n^4)$ computational complexity.

In this talk we propose a modification of the latter algorithm that computes the Jordan structure [4] of A with high relative accuracy in $O(n^3)$ computational complexity.

References

[1] S. M. Fallat and C. R. Johnson, *Totally nonnegative matrices*, Princeton University Press, Princeton, NJ, (2011).

[2] P. Koev, *Accurate computations with totally nonnegative matrices*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 731–751.

[3] P. Koev, *Accurate Eigenvalues and Exact Zero Jordan Blocks of Totally Nonnegative Matrices*, Numer. Math., 141, (2019), pp. 693–713.

[4] N. Mastronardi, P. Van Dooren, *Computing the Jordan Structure of an Eigenvalue*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 949–966.

Path-following methods for the matrix numerical range, waves on a shear flow, and other

parametrised eigenvalue problems

Peter Maxwell *(Norwegian University of Science and Technology)*

An initial review of the path-following method in [1] for calculating the field of values boundary of a matrix is given. Thereafter, an adaptation of that algorithm to instead calculate the dispersion relation for linear surface waves on a shear flow is presented. Possible generalisations to other problem types are described.

The Johnson algorithm for calculating the field of values boundary of a matrix, A, is a parametrisation of the eigenproblem formed by projection of A onto the real axis after rotation by angle θ. By differentiating the (Hermitian) eigenproblem with respect to the real scalar θ and including an additional constraint, a system of ordinary differential equations is obtained. After an initial eigenproblem solve for the dominant eigenpair, Runge–Kutta numerical integration can then integrate along the boundary curve to calculate control points. Piecewise Hermite interpolation yields dense output. The numerical integration represents a nominal setup cost after which very many points on the curve can be calculated accurately at negligible cost.

For a given background velocity flow $U(z)$ in the horizontal direction at depth z and wave number $k \in [k_a, k_b]$, the dispersion relation is the curve described by the solution phase velocity $c(k)$ to the eigenproblem formed by the Rayleigh equation with linearised free-surface boundary condition, parametrised by k. Usually, the problem amounts to finding sufficiently many points on the curve. Using a collocation method with two point boundary row-replacement strategy to discretise the problem, a parametrised quadratic eigenvalue problem is obtained. A similar approach as before –differentiating with respect to k and forming a system of differential equations– permits integration along the dispersion relation curve. This has the benefit of swapping many QZ decompositions on a size $2N$ companion matrix for one QZ decomposition and some linear solves on a size N matrix. Again, dense output is obtained by piecewise Hermite interpolation. The resulting path-following algorithm preserves the accuracy in the initial collocation calculation and is asymptotically two orders of magnitude faster than other known numerical algorithms for this problem.

Finally, possible extensions of the method and uses in solving other problems are briefly mentioned.

Reference

[1] Loisel, S. and Maxwell, P. (2018), Path-Following Method to Determine the Field of Values of a Matrix with High Accuracy, *SIAM J. Matrix Anal. Appl.* 39-4 (2018), pp. 1726-1749, doi: 10.1137/17M1148608.

On divergence-free methods for double-diffusion equations in porous media

Paul E. Méndez & Raimund Bürger *(Universidad de Concepción)*, & Ricardo Ruiz-Baier *(Oxford University)*

A stationary Navier-Stokes-Brinkman model coupled to a system of advection-diffusion equations serves as a model for so-called double-diffusive viscous flow in porous media in which both heat and a solute within the fluid phase are subject to transport and diffusion. The solvability analysis of these governing equations results as a combination of compactness arguments and fixed-point theory. In addition an $\boldsymbol{H}(\mathrm{div})$-conforming discretisation is formulated by a modification of existing methods for Brinkman flows. The well-posedness of the discrete Galerkin formulation is also discussed, and convergence properties are derived rigorously. Computational tests confirm the predicted rates of error decay and illustrate the applicability of the methods for the simulation of bacterial bioconvection and thermohaline circulation problems.

Keywords: Viscous flow in porous media; doubly-diffusive problems; cross-diffusion; fixed-point theory; mixed finite element methods; a priori error estimation.

Mathematics Subject Classifications (2010): 65N30; 76S05; 76R50.

Relative perturbation bounds for regular quadratic eigenvalue problem

Suzana Miodragović & Ninoslav Truhar & Xin Liang & Peter Benner *(University of Osijek)*

We present new relative perturbation bounds for the eigensubspaces for the quadratic eigenvalue problem $\lambda^2 Mx + \lambda Cx + Kx = 0$, where M and K are nonsingular Hermitian and C is any Hermitian matrix. First, we derive the $\sin\Theta$ type theorems for the eigensubspaces of the regular matrix pairs (A, B), where both A and B are Hermitian matrices. Using a proper linearization and new relative perturbation bounds for regular matrix pairs (A, B), we develop correspond-

ing sin Θ type theorems for the eigensubspaces for the considered regular quadratic eigenvalue problem. Our bound can be applied to the gyroscopic systems which will also be shown. The obtained bounds will be illustrated by numerical examples.

An image restoration by a stochastic model

F.Z. Nouri & M. Benseghir *(Badji Mokhtar University)*

In this work, we tackle an image restoration problem by an approach based on a stochastic differential equation (SDE) with reflection. The drift and diffusion terms are rigorously formulated to express the concept of the stochastic anisotropic diffusion. For the numerical approximation, we consider the modified Euler scheme with a random stop time and a diffusion parameter depending on the image geometry structure. The numerical results obtained with the proposed algorithm, which clearly show the improvement of the edges, the smoothing of the noise and the connection of the broken lines, are very encouraging and they demonstrate its competitiveness compared to other methods.

Matrix equations techniques for time-dependent partial differential equations

Davide Palitta *(Max Planck Institute for Dynamics of Complex Technical Systems)*

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

Combining timely 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 different time-dependent PDEs and several numerical results are reported to illustrate the potential of our novel approach.

Parallel cross interpolation for high–precision computation of high–dimensional integrals

Dmitry Savostyanov & Sergey Dolgov *(University of Brighton)*

High–dimensional integrals appear in stochastic calculus, mathematical finance, quantum physics, etc. Textbook quadratures applied to d–dimensional integrals require n^d function evaluations, which is unfeasible for $d > 10$ (a notoriously known curse of dimensionality). Typically such integrals are treated using Monte Carlo (random samples) or quasi-MC algorithms (optimal lattice), but due to their relatively slow convergence they struggle to deliver very accurate results.

An alternative approach is offered by tensor product algorithms, which separate variables and hence lift the curse of dimensionality. The multivariate is reconstructed from a small adaptively chosen set of its univariate fibers, representing its behaviour along all dimensions. The resulting cross interpolation algorithm converges faster than MC and qMC and enables high–precision integration. A parallel version of tensor interpolation algorithm is based on independent exploration of all dimensions of the given problem.

The efficiency of the proposed algorithm is demonstrated for Ising integrals related to magnetic susceptibility of two–dimensional spin lattices. This application encourages evaluation of integrals in dimensions up to 1000 with very high precision, eventually allowing physicists to conjecture analytic formulae. Monte Carlo methods are not up to the challenge, while the tensor interpolation algorithm computes integrals to 100 decimal digits.

The success of tensor cross interpolation for calculation of Ising integrals encourages us to try it for a wider class of applications, where high precision is required, e.g. multivariate probability, stochastics and data science. We are hopeful that the proposed tensor cross interpolation algorithm will demonstrate fast convergence in these applications and eventually become a method of choice for high–dimensional integration.

What kind of tensors are compressible?

Tianyi Shi & Alex Townsend *(Cornell University)*

Tensors often have too many entries to be stored explicitly so it is essential to compress them into data sparse formats. I will identify three methodologies

that can be used to explain when a tensor is compressible. Each methodology leads to bounds on the compressibility of certain tensors, partially explaining the abundance of low-rank tensors in applied mathematics. In particular, I will focus on tensors with a so-called displacement structure, showing that solutions to Poisson equations on tensor-product geometries are highly compressible. As the rank bounds are constructive, I will develop an optimal-complexity spectrally-accurate 3D Poisson solver with $\mathcal{O}(n(\log n)^2(\log 1/\epsilon)^2)$ complexity for a smooth righthand side, where $n \times n \times n$ is the tensor discretization of the solution.

Minimax approximation by ridge functions, and neural networks

David E. Stewart & Palle Jorgensen *(University of Iowa)*

Neural networks with a single hidden layer approximate unknown functions given implicitly by a large set of labelled data. The problem that is considered here is the question of approximating continuous functions by sums of ridge functions with weight vectors in a given fixed set W. More specifically, we aim to determine how large is the largest unapproximable function by ridge functions of this type with Lipschitz constant one. This should give us a better understanding of how well sums of ridge functions with fixed weight vectors can approximate general functions. This information is particularly relevant to Extreme Learning Machines which do not train all weights.

Invertible piecewise rational interpolation

Jan Van lent *(UWE Bristol)*

This talk presents a simple method to interpolate strictly monotone data using a monotone rational spline that is continuous and has continuous derivatives. Each piece of the spline is a rational function of degree one. The inverse of the approximation is a spline of the same type and is obtained at the same time. This class of monotone and easily invertible functions can be useful for approximating cumulative density functions, for example. Previous methods for constructing splines of this type were based on solving nonlinear systems of equations. In this talk, a method will be presented that only requires solving a bidiagonal or tridiagonal linear system. The cost of constructing the spline and the cost of evaluating the spline are essentially the same as for a standard quadratic or cubic spline. However, such standard splines are not necessarily monotone functions for monotone data. Insuring monotonicity requires extra work and even if the spline function is monotone, evaluating the inverse function is not straightforward and requires extra computation. In summary, the method presented requires similar calculations to standard splines, but in addition insures monotonicity for monotone data and obtains the inverse spline function at no extra cost.

Compression properties in rank-structured solvers for Toeplitz linear systems

Heather Wilber & Bernhard Beckermann *(Université de Lille)* & Daniel Kressner *(EPFL)*

Toeplitz matrices are abundant in computational mathematics, and since they are also data sparse, it is natural to consider fast and superfast methods for solving linear systems involving them. If a matrix T is circulant in addition to being Toeplitz, then T is diagonalized by a unitary discrete Fourier transform. If T is not circulant, then applying a discrete Fourier transform results in a matrix with off-diagonal blocks that are well-approximated by low rank matrices. Surprisingly little is known about the compressibility of these highly structured matrices in a theoretically rigorous sense, even though this compressibility is relied upon in practice in a number of solvers. In this talk, we show that the compression properties of these matrices can be thoroughly explained as a consequence of their displacement structure. Our results lead to extremely efficient displacement-based compression strategies that can be used to formulate fully adaptive superfast rank-structured Toeplitz solvers. Our ideas also lead to superfast solvers for other linear systems involving matrices with special displacement structures.

Backward errors incurred by the CORK linearizations

Kuan Xu *(University of Science and Technology of China)* & Hongjia Chen *(Nanchang University)*

In this presentation, we show a unified treatment for the compact rational Krylov (CORK) linearization expressed in various commonly used bases, including Taylor, Newton, Lagrange, orthogonal, and rational basis functions. We construct one-sided factorizations that relate the eigenpairs of the CORK linearization and those of the original polynomial or rational eigenvalue problem. With these factorizations, we establish upper bounds in terms of norms of the coefficient

matrices and those of the basis functions for the backward error of an approximate eigenpair of the original eigenvalue problem relative to the backward error of an approximate eigenpair of the CORK linearization. These bounds suggest a scaling strategy to reduce the backward error of the computed eigenpairs. Numerical experiments show that the actual backward errors are successfully reduced by scaling and the errors, before and after scaling, are both well bounded by the bounds.

Optimal transport based reflection inversion: uncover model below reflectors

Yunan Yang & Björn Engquist *(New York University)*

The least-squares norm (L^2) as the objective function in reflection-dominated full waveform inversion (FWI) typically results in migration like features in the initial model updates. Further iterations only marginally improve the velocity model.

We will show that the failure is most likely that the L^2 norm is relatively insensitive to the lower frequency component of the data.

Inversions using the L^2 norm update the high-wavenumber components in the model first, which is good for migration but problematic for inversion.

On the other hand, with an intrinsic "transport" property, the optimal transport based misfit functions naturally exploit the lower frequency component in the data and generates low-wavenumber model update which is ideal for recovering the model kinematics.

Our focus is on recovering velocity even below the reflecting interfaces by reflections only. Successful recovery of several layered models, as well as a realistic benchmark with salt inclusion, demonstrate that the optimal-transport FWI is able first to reconstruct the low-wavenumber structure of the model and consequently accelerate the model convergence.

Fast iterative solver for stochastic Galerkin finite element approximations

Rawin Youngnoi & Daniel Loghin & Alex Bespalov *(University of Birmingham)*

Stochastic Galerkin finite element method (SGFEM) provides an efficient alternative to sampling methods for the numerical solution of linear elliptic PDE problems with parametric or random inputs. In order to compute the stochastic Galerkin solution for a given problem, one needs to solve a large coupled system of linear equations. Thus, an efficient iterative solver is a key ingredient of any SGFEM implementation. In this talk, we discuss a new preconditioning technique for SGFEM. We will present the available analysis and the results of numerical experiments for the model diffusion problems with affine and non-affine parametric representations of the coefficient. Specifically, we look at the efficiency of the solver (in terms of iteration counts for solving the underlying linear system) and compare the new preconditioner with other existing preconditioners for stochastic Galerkin matrices, such as the mean-based preconditioner from [C. E. Powell and H. Elman, IMA J. Numer. Anal., 29: 350–375, 2009] and the Kronecker product preconditioner from [E. Ullmann, SIAM J. Sci. Comput., 32(2): 923-946, 2010].