# Report on BIRS Workshop 09w5101: Advances and Perspectives on Numerical Methods for Saddle Point Problems

Howard Elman (University of Maryland)
Chen Greif (University of British Columbia)
Dominik Schötzau (University of British Columbia)

April 12, 2009 – April 17, 2009

Saddle point problems arise in a wide variety of mathematical fields, including models of incompressible flow, optimization, and control. The aim of this workshop was to bring together scientists who study numerical methods for solving saddle point problems to exchange ideas and examine the influences of different applications on development of algorithms. There were 39 participants of whom 7 were from Canada, 14 from the United States, and 18 from Europe and Asia. A total of 29 talks of length 30 minutes were presented during the five-day workshop.

## 1 Overview of the Field

In recent years, saddle point problems have taken an increasingly prominent role in mathematical modeling and applied science. Examples of settings where they appear include

1. fluid dynamics and magnetohydrodynamics, where models of incompressible fluids and interactions of electromagnetic fields and incompressible fluids are obtained from the numerical solution of saddle point problems;

2. general methods of optimization, which entails the minimization of cost functions subject to constraints consisting of bounds on solution values;

3. optimal control and PDE-constrained optimization, in which parameters associated with physical devices and systems are determined subject to constraints associated with partial differential equations.

Improvements in algorithm development offer the potential for increasingly accurate models to be solved. In addition, such developments have facilitated the use of modeling methods in many new arenas, such as liquid crystal devices, image processing, and computer graphics.

A common feature of the field is that the most effective solution algorithms take advantage of the specific structure of the problem, which has the generic form

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} f \\ g \end{bmatrix}.$$

Efficient algorithms typically are built on efficient approximations to the operator $A$ and the Schur complement $BA^{-1}B^T$. Different settings require attention to specific components of the problem. For example, in some settings, such as interior point methods for optimization, the matrix $A$ may be poorly understood, whereas in others, such as fluid dynamics, there is a large literature of techniques such as multilevel methods that can be used to handle it. Thus, different fields have a commonality of approach but they require attention to specific aspects of the problem domain they come from. A primary goal of this meeting was to facilitate the interaction of mathematical sciences working on the various aspects of saddle point projects, and to help foster cross-pollination of ideas developed within the different subfields.

## 2 Recent Developments

The size of the systems of interest can be very large for accurate representation of models. For example, accurate three-dimensional models of fluid flow discretized by finite element methods require thousands and even millions of degrees of freedom. Thus, for such models to become truly practical, there is a continued need for improvements in efficiency of numerical solution algorithms.

In recent years, there have been very significant improvements in all aspects of algorithms for saddle point problems [1], including specific approaches designed for fluid dynamics [3], general optimization [4], and control [2]. For example, new algorithms based on effective approximation to the Schur complement have facilitated the development of models that avoid costly integration of time-dependent problems to steady-state, and enable the use of implicit time-stepping methods.

Saddle point systems that arise from interior point methods in optimization pose two major challenges. First, not much is known about the operator $A$: it changes throughout the optimization iteration, which makes it difficult or impossible to identify specific properties of the operator that could be exploited in connection to the underlying application. A second difficulty is that the system becomes increasingly ill-conditioned as the solution is approached. As a result, inversion of $A$ is harder than in other settings. An example of an effective and widely used solution methodology is one where the focus is on preserving the constraints throughout the iteration, while approximating the operator $A$ by a simpler operator.

PDE-constrained optimization entails the solution of inverse problems, where the goal is to optimize the values of parameters associated with a simulation for which the underlying physics is modeled by a partial differential equation. Examples of applications include geodynamics, environmental modeling, and chemical processes. Improvements in algorithm technology for PDEs (including techniques for constrained problems considered in this workshop) now make it feasible to solve such optimization problems numerically. This entails many solutions of the equations associated with the constraints (for example, a model of incompressible flow) throughout the simulation.

## 3 Talks

The meeting was organized loosely according to the themes described above, with talks organized in sessions by topic. A summary of the abstracts of talks presented, grouped by topic, is given below.

**1. Problems in fluid dynamics and magnetohydrodynamics.**

**Constantin Bacuta** (University of Delaware)
*Multilevel Discretization of Saddle point Problems without the Discrete LBB Condition*
Using spectral results for Schur complement operators we prove a convergence result for the inexact Uzawa algorithm on general Hilbert spaces. We prove that for any symmetric and coercive saddle point problem, the inexact Uzawa algorithm converges, provided that the inexact process for inverting the residual at each step has the relative error smaller than a computable fixed threshold. As a consequence, we provide a new type of algorithms for discretizing saddle point problems, which implement the inexact Uzawa algorithm at the continuous level as a multilevel or adaptive algorithm. The discrete stability Ladyshenskaya-Babuška-Brezzi (LBB) condition might not be satisfied and the adaptivity is required only for solving symmetric and positive definite problems. The convergence result for the algorithm at the continuous level, combined with standard techniques of discretization and a posteriori error estimates leads to new and efficient algorithms for solving saddle point systems. Numerical results supporting the efficiency of the algorithm are presented for the Stokes Equations and the div?curl systems.

**Michele Benzi** (Emory University)
*Smoothers and Preconditioners for Saddle Point Problems Arising from the Incompressible Navier-Stokes Equations*
I will present recent work on iterative solvers for various discretizations of the incompressible Navier- Stokes equations, for both steady and unsteady flow cases. Although Picard linearization is used, many of the techniques and results are applicable to Newton linearization as well. The talk will focus on two classes of methods:

1. Coupled multigrid methods, with a focus on new smoothers;
2. Block triangular preconditioners for Krylov subspace methods.

More specifically, I will discuss the use of the Hermitian and skew-Hermitian splitting iteration as a smoother in a coupled multigrid method, as well as new preconditioners and a multigrid method for augmented Lagrangian formulations of the linearized Navier-Stokes equations. I will examine the performance of various solvers as the mesh size, Reynolds number, time step, and other problem parameters vary. Local Fourier analysis and extensive numerical tests indicate that fast convergence is achieved in many cases, with weak or no dependence on problem parameters.

**Dan Li and Xiaoxi Wei** (University of British Columbia)
*A Mixed Finite Element Method for Incompressible Magnetohydrodynamics*
Incompressible magnetohydrodynamics (MHD) models the interaction of electromagnetic fields with viscous, electrically conducting incompressible fluids. The numerical simulation of these problems requires discretizing a system of non-linear partial differential equations that couple the incompressible Navier-Stokes equations with Maxwells equations. In the talk, we will present a mixed finite element method for two-dimensional non-linear incompressible MHD. The fluid velocity is approximated using the divergence-conforming Brezzi-Douglas-Marini elements of degree k, while the pressure is discretized using discontinuous piecewise polynomials of total degree at most $k?1$. For the magnetic equations, we introduce a Lagrange multiplier related to the divergence constraint of the magnetic filed, and formulate a mixed variational setting. The magnetic field is then simulated using the first family of curl-conforming Nédélec elements of degree k, whereas the Lagrange multiplier belongs to the space of continuous piecewise polynomials of total degree at most k. Our method is proved inf-sup stable; it also provides exactly divergence-free velocity approximations. The use of Nédélec elements for the simulation of the magnetic field overcomes a drawback inherent in $H^1$-conforming nodal elements, namely that they cannot correctly capture the strongest singularity in non-convex domains. We carry out a complete a-priori error analysis for the proposed mixed method, and show optimal convergence in the energy norm for smooth solutions. For problems with re-entrant corners, we prove $\delta$-optimal convergence rates, i.e., the convergence is optimal up to a factor of $h^\delta$ for any $\delta > 0$. Finally, a series of numerical convergence tests on MHD benchmark problems are presented to highlight the practical performance of our method.

**Daniel Loghin** (University of Birmingham)
*Discrete Interpolation Norms with Applications*
The work I will describe concerns matrix representations of norms for Hilbert scales and their efficient implementation in a range of applications. These norms are products of integer and noninteger powers of Grammian matrices associated with the generating pair of spaces for the interpolation space. The case of interest is fractional Sobolev spaces both for positive and negative indices with applications arising in preconditioning techniques. Numerical examples will include boundary Schur complement preconditioning for domain decomposition of elliptic problems (both scalar equations and systems).

**Maxim Olshanskii** (Moscow Lomonosov State University)
*A Solver for the Stokes Type Problem with Variable Viscosity and Some Applications*
The talk concerns with an iterative technique for solving discretized Stokes type equations with varying viscosity coefficient. The subject is motivated by numerical solution of incompressible non-Newtonian fluid equations, but also important for several geophysical problems, such as magma migration and mantle convection. We build a special block preconditioner for the discrete system of equations and perform an analysis revealing its properties. In particular, the general analysis is applied to the linearized equations of the regularized Bingham model of viscoplastic fluid. Both theoretical analysis and numerical experiments show that the preconditioner leads to a significant improvement of an iterative method convergence and results almost insensible to viscosity variation.

**Catherine Powell** (University of Manchester)
*Solving Saddle Point Problems Arising in Stochastic Mixed Finite Element Problems*
In the last few years, interest in so-called stochastic finite element methods (SFEMs), which facilitate the approximation of statistics of solutions to PDEs with random data, has risen sharply. SFEMs based on Galerkin approximation satisfy an optimality condition, but unlike Monte Carlo methods, require the solution of a single (but very large) linear system of equations. The solution of this linear system of equations, which couples deterministic and stochastic degrees of freedom, is regarded as a serious bottleneck in computations. This difficulty is even more pronounced when we attempt to solve systems of PDEs with random data via stochastic mixed FEMs based on Galerkin approximation. To illustrate these challenges, we focus on the Darcy flow problem with random permeability coefficients, written as a first-order system. We derive the saddle-point systems that result from stochastic Galerkin approximation based on finite element spatial discretisations. These are orders of magnitude larger than the saddle-point systems arising for the deterministic problem. We report on fast and robust solvers and preconditioners based on multigrid methods which have proved successful for stochastically linear problems, and point to some challenges still to be met for stochastically nonlinear problems.

**Arnold Reusken** (RWTH Aachen)
*Properties of a New Finite Element Pair for Incompressible Two-Phase Flow Simulations*
We consider a domain $\Omega \subset \mathbb{R}3$ which contains two different immiscible incompressible newtonian phases (fluid-fluid or fluid-gas). The time-dependent domains which contain the phases are denoted by $\Omega_1 = \Omega_1(t)$ and $\Omega_2 = \Omega_2(t)$ with $\overline{\Omega}_1 \cup \overline{\Omega}_2 = \overline{\Omega}$. The interface between the two phases is denoted by $\Gamma = \Gamma(t)$. To model the forces at the interface we make the standard assumption that the surface tension balances the jump of the normal stress on the interface, i.e.

$$[\boldsymbol{\sigma}\mathbf{n}]_\Gamma = \tau\kappa\mathbf{n} \,,$$

with $\mathbf{n} = \mathbf{n}_\Gamma$ the unit normal at the interface, $\tau$ the surface tension coefficient (material parameter), $\kappa$ the curvature of $\Gamma$ and $\boldsymbol{\sigma}$ the stress tensor

$$\boldsymbol{\sigma} = -p\mathbf{I} + \mu\mathbf{D}(\mathbf{u}), \quad \mathbf{D}(\mathbf{u}) = \nabla\mathbf{u} + (\nabla\mathbf{u})^T.$$

We assume continuity of the velocity across the interface. In combination with the conservation laws of mass and momentum this yields the following standard model:

$$\begin{cases} \rho_i\big(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u}\cdot\nabla)\mathbf{u}\big) = -\nabla p + \rho_i\mathbf{g} + \operatorname{div}(\mu_i\mathbf{D}(\mathbf{u})) & \text{in } \Omega_i \\ \operatorname{div}\mathbf{u} = 0 & \text{in } \Omega_i \end{cases} \quad \text{for } i = 1,2$$
$$[\boldsymbol{\sigma}\mathbf{n}]_\Gamma = \tau\kappa\mathbf{n}, \quad [\mathbf{u}\cdot\mathbf{n}]_\Gamma = 0\,.$$

The vector $\mathbf{g}$ is a known external force (gravity). In addition we need initial conditions for $\mathbf{u}(x,0)$ and boundary conditions at $\partial\Omega$. For simplicity we assume homogeneous Dirichlet boundary conditions.
If for interface capturing a level set (or VOF) method is applied then the interface, which is implicitly given by the zero level of the level set function, is in general not aligned with the triangulation that is used in the discretization of the flow problem. This non-alignment causes severe difficulties w.r.t. the discretization of the localized surface tension force and the discretization of the flow variables. In cases with large surface tension forces the pressure has a large jump across the interface. In standard finite element spaces, due to the non-alignment, the functions are continuous across the interface and thus not appropriate for the approximation of the discontinuous pressure. In many simulations these effects cause large oscillations of the velocity close to the interface, so-called spurious velocities.
In this talk it is shown that an extended finite element space (XFEM) is much better suited for the discretization of the pressure variable. For the discretization of the velocity we use standard piecewise quadratics. We present (optimal) approximation error bounds and prove that the diagonally scaled mass matrix has a uniformly bounded spectral condition number. We address the issue of LBB-stability of the XFEM-$P_2$ pair. Results of numerical experiments are presented that illustrate properties of the XFEM space.

**Joachim Schöberl** (University of Aachen)
*Preconditioning for Divergence-free Hybrid-DG Finite Elements for the Navier Stokes Equations*
We consider a hybrid-DG version of the method by Cockburn, Kanschat and Schötzau for the discretization of the Navier Stokes equations. It uses H(div) conforming BDM finite elements, and allows to pose the incompressibility condition in a pointwise sense. We show how to apply kernel preserving multigrid methods for the efficient solution of the linear system.

**David Silvester** (University of Manchester)
*Fast Solvers for Unsteady Incompressible Flow*
Simulation of the motion of an incompressible fluid remains an important but very challenging problem. The resources required for accurate three-dimensional simulation of practical flows test even the most advanced computer hardware. The necessity for reliable and efficient solvers is widely recognised. This talk will focus on two components of such a solver: the error control used for self-adaptive time stepping; and the linear solver used at each time level. Conventional codes typically use semi-implicit time integration leading to a Poisson or Stokes-type system at every time step, but with a stability restriction on the time step. Our alternative approach is a stable version of the TRAB2 smart integrator originally developed by Gresho in the 1980s. Such fully-implicit time integration methods have no restriction on the time step, but have only become feasible in the last five years because of developments in solution techniques for the linear (or linearized) Oseen systems that arise at each time level. To this end, the preconditioning framework that we propose offers the possibility of uniformly fast convergence independent of the problem parameters (namely; the mesh size, the time step and the Reynolds number). In contrast, conventional multigrid solvers for this problem tend to work well if the time step is sufficiently small, but efficiency deteriorates rapidly if either the time step or the Reynolds number is increased.

**Alastair Spence** (University of Bath)
*The Calculation of Pure Imaginary Eigenvalues of Large Sparse Matrix Pencils with Application to Detecting Hopf Bifurcation in Large Systems*
The detection of a Hopf bifurcation in a large scale dynamical system that depends on a physical parameter often consists of computing the right-most eigenvalues of a sequence of large sparse eigenvalue problems. This is not only an expensive operation, but the computation of right-most eigenvalues is often unreliable for the commonly used methods for large sparse matrices. However, if good starting guesses are available for the parameter and the purely imaginary eigenvalue at the Hopf point, then efficient algorithms to accurately compute the bifurcation are available. In this talk, we propose a method for obtaining such good starting guesses, based on finding purely imaginary eigenvalues of a two-parameter eigenvalue problem (possibly arising after a linearisation process). The method utilises the Kronecker product and involves (in theory but not in practice) the solution of matrices of squared dimension. The problem is reformulated as an inexact inverse iteration method

that requires the solution of a sequence of Lyapunov equations with low rank right hand sides. It is this last fact that makes the method feasible for large systems. The power of the method is tested on numerical examples, one of which is a discretised PDE with two space dimensions.

**Walter Zulehner** (Johannes Kepler University Linz)
*Nonstandard Norms and Robust Estimates for Saddle Point Problems*
An important issue for constructing and analyzing efficient solvers for a saddle point problem is to understand the right mapping properties of the problem. We will concentrate on saddle point problems which result from the discretization of a system of partial differential equations. The mapping properties of the involved differential operators usually suggest the right norms for the discrete problems leading to mesh-independent estimates. These norms are quite often (discrete versions of) standard norms in Lebesgue or Sobolev spaces. If the saddle point problem contains critical parameters (like regulation parameters in optimal control problems) one would like to use norms leading to mesh-independent estimates which are also robust with respect to these critical parameters. Here standard norms usually do not the job. In this talk we will discuss the construction of norms for saddle point problems which lead to robust estimates: Firstly, on a purely algebraic level, a characterization of such norms is given for a general class of symmetric saddle point problems. Then we will apply these results to a family of optimal control problems.

## 2. Saddle point problems for general optimization.

**Jacek Gondzio** (University of Edinburgh)
*Inexact Constraint Preconditioners for Linear Systems Arising in Interior Point Methods*
Issues of indefinite preconditioning of reduced Newton systems (saddle point problems) arising in optimization with interior point methods will be addressed. Constraint preconditioners have shown much promise in this context. However, there are situations in which an unfavourable sparsity pattern of Jacobian matrix may adversely affect the preconditioner and make its inverse representation unacceptably dense hence too expensive to be used in practice. A remedy to such situations is proposed. An approximate constraint preconditioner is considered in which sparse approximation of the Jacobian is used instead of the complete matrix. Spectral analysis of the preconditioned matrix is performed and bounds on its non-unit eigenvalues are provided. Preliminary computational results are encouraging.

**Michael Overton** (New York University)
*Preconditioners for Semidefinite Programming*
Semidefinite programming has become a valuable paradigm whose practical impact is mainly limited by the large dense ill-conditioned systems of linear equations that arise when implementing interior-point methods that follow the central path. We investigate the behavior of iterative methods with preconditioners that depend on computing or approximating the smallest eigenvalues of the dual slack matrix, as well as the largest eigenvalues of the primal matrix variable when a primal-dual method is used. We also consider alternative preconditioners based on exploiting the Cholesky factorizations of these matrices. We consider both the Schur complement positive definite system of equations (using a preconditioned conjugate gradient method) and the underlying indefinite augmented system (using a preconditioned MINRES iteration). For the eigenvalue-based preconditioners, we present results characterizing the eigenvalues of the preconditioned linear systems on the central path in terms of ?, the variable that parameterizes the central path, making nondegeneracy assumptions. The cost of computing preconditioners can be amortized by reusing them to solve related linear systems corresponding to smaller values of ?. We also present a novel technique to improve the preconditioner for the Schur complement system by explicitly updating it during the conjugate gradient iteration.

**Michael Saunders** (Stanford University)
*An Active-Set Convex QP Solver Based on Regularized KKT Systems*
Implementations of the simplex method depend on basis repair to steer around near-singular basis matrices, and KKT-based QP solvers must deal with near-singular KKT systems. However, few sparse-matrix packages have the required rank-revealing features (we know of LUSOL, MA48, MA57, and HSL MA77).
For convex QP, we explore the idea of avoiding singular KKT systems by applying primal and dual regularization to the QP problem. A simplified single-phase active-set algorithm can then be developed. Warm starts are straightforward from any given active set, and the range of applicable KKT solvers expands.
QPBLUR is a prototype QP solver that makes use of the block-LU KKT updates in QPBLU (Hanh Huynhs PhD dissertation, 2008) but employs regularization and the simplified active-set algorithm. The aim is to provide a new QP subproblem solver for SNOPT for problems with many degrees of freedom. Numerical results confirm the robustness of the single-phase regularized QP approach.

## 3. PDE-constrained optimization.

**Pavel Bochev** (Sandia National Laboratories)
*Robust and Efficient Solvers via Optimization-Based Reformulation*
I will present a new optimizationbased approach for robust and efficient solution of PDE problems consisting of multiple physics operators with fundamentally different mathematical properties. This approach relies on three essential steps: decomposition of the original problem into subproblems for which robust solution algorithms are available; integration of the subproblems into an equivalent PDEconstrained optimization problem; and solution of the resulting optimization problem either directly as a fully coupled saddle-point algebraic system, or in the null space of the PDE constraints. This strategy gives rise to a general approach for synthesizing robust solvers for complex coupled problems from solvers for their simpler physics components.

**Sue Dollar** (Rutherford Appleton Laboratory)
*Projected Krylov Methods and PDE-Constrained Optimization*
Over recent years, constraint preconditioners have formed a popular class of preconditioners for saddlepoint problems. Additionally, such a preconditioner allows the user to apply a projected conjugate gradient method to solve the saddle-point problem. For simple PDE-constrained problems, the use of constraint preconditioners is very attractive and have been shown to be very effect. However, if the PDE in the constraints is, say, the Stokes equation, then it is not so obvious how a constraint preconditioner in conjunction with the projected CG method can be used. We will discuss this problem, consider a preconditioner which only replicates a subset of the constraints and formulate a new projected Krylov method that can be used within this framework.

**Tyrone Rees** (University of Oxford)
*Optimal Preconditioners for Problems in PDE-Constrained Optimization*
The problem of minimizing a cost functional subject to a constraint that is a partial differential equation arises widely in many areas of science and engineering. Such problems, when discretized, can be written in saddle-point form. The resulting system is generally of very large dimension and so matrix factorizations of any type must be avoided. We will show how to exploit the structure of the problem to give an optimal block diagonal preconditioner - based on standard multigrid cycles - to solve the problem all-at-once using MINRES. We look at the theory with respect to two different PDEs: Poisson's equation and the Stokes equations, and we give computational results in both cases.

**Olaf Schenk** (University of Basel)
*Scalable Inexact Interior-Point Algorithms for Large-scale Nonlinear Optimization based on Spike-Pardiso*
The availability of large-scale computing platforms comprised of thousands of multicore processors motivates the need for highly scalable sparse linear system solvers for symmetric indefinite matrices. These solvers must optimize parallel performance, processor (serial) performance, as well as memory requirements, while being robust across broad classes of applications and systems. We will present a new parallel solver that combines the desirable characteristics of direct methods (robustness) and iterative solvers (computational cost), while alleviating their drawbacks (memory requirements, lack of robustness). Our proposed hybrid solver is based on the general sparse solver Pardiso, and the Spike family of solvers. The resulting algorithm, called Spike-Pardiso, is as robust as a direct solver, much more reliable than classical preconditioned Krylov subspace methods, and much more scalable than direct sparse solvers.
We support our performance and parallel scalability claims using detailed experimental studies in large-scale nonlinear and nonconvex PDE-constrained optimization. The goal is to design a constrained optimization algorithm that emulates an efficient nonlinear programming approach. The optimization algorithm may utilize matrix-vector products with the constraint Jacobian, its transpose, and the Hessian of the Lagrangian together with appropriate preconditioners  quantities that are computable for many large-scale applications of interest  but must overcome the fact that exact factorizations of derivative matrices are impractical to obtain. Iterative linear system solvers present a viable alternative to direct factorization methods, but the benefits of these techniques are only realized if inexact step computations are controlled appropriately in order to guarantee global convergence of the algorithm for PDE problems.
The method is used in a PDE-constrained optimization algorithms which is designed for parallel scalability on distributed-memory architectures with thousands of cores. The optimization method is based on a line-search interior-point algorithm for large-scale continuous optimization, it is matrix-free in that it does not require the factorization of derivative matrices. Instead, it uses the new preconditioning method in Spike-Pardiso.
The preconditioning method is based on a Spike factorization, symmetric weighted matchings and eigenvalue spectral multilevel orderings for symmetric indefinite saddle-point matrices. We will show almost linear parallel scalability results on up to 1000 cores for the complete optimization problem, which is a new emerging important biomedical application and is related to antenna identification in hyperthermia cancer treatment planning.

**Martin Stoll** (University of Oxford)
*Block Triangular Preconditioners for PDE-Constraint Optimization with Application to Box Constraints*
We investigate the possibility of using block triangular preconditioner for saddle point problems arising in PDE-constrained optimization. In particular we focus on a cg-type method introduced by Bramble and Pasciak. The method is a well-known tool in the PDE community but has the drawback of requiring an appropriate scaling of the preconditioners to be applicable. This involves solving an eigenvalue estimation problem which is in general expensive. For examples coming from PDE-constrained optimization, we show that when the Chebyshev semi-iteration is used as a preconditioner the drawback of the Bramble-Pasciak method is easily overcome. We present an eigenvalue analysis for block triangular preconditioners and illustrate their competitiveness on some examples.

## 4. Linear algebra solution methods and analysis.

**Sabine Le Borne** (Tennessee Technological University)
*Algebraic Hierarchical Matrix Preconditioners for Saddle Point Problems*
Hierarchical (H-) matrices provide a powerful technique to compute and store approximations to dense matrices in a data-sparse format. In the past, H-matrix techniques have been exploited for the solution of dense subproblems of saddle point preconditioners, but they also led to the development of completely new preconditioners. While these preconditioners proved to be quite robust and efficient, one drawback of H-matrix-based preconditioners is their need for geometric information associated with the discretization that leads to the saddle point matrix. In this talk, we will introduce a blackbox technique that replaces the geometry-based construction of an H-matrix by an algebraic approach that only requires the (sparse) system matrix itself as input. Both the geometric bisection and the more efficient geometric domain-decomposition clustering techniques are replaced by their algebraic counterparts which are based on matrix graphs. As a result, we can now construct various H-matrix preconditioners from the system matrix directly. We will review some promising preconditioners from the literature and introduce some new preconditioning approaches for saddle point problems which benefit from these blackbox H-matrices. We will conclude with numerical results for the Oseen problem in two as well as three spatial dimensions in which we compare the setup times, storage requirements, and convergence properties of several preconditioners for a variety of test problems.

**Jörg Liesen** (Technical University of Berlin)
*On Nonsymmetric Saddle Point Matrices that Allow Conjugate Gradient Iterations*
Linear systems in saddle point form are usually highly indefinite, and this often slows down iterative solvers such as Krylov subspace methods. It has been noted by several authors that negating the second block row of a symmetric indefinite saddle point matrix leads to a nonsymmetric matrix whose spectrum is entirely contained in the right half plane. Is this alternative formulation of a saddle point problem useful? It seems so when the eigenvalues of the nonsymmetric matrix are still real, since then the system with the nonsymmetric matrix can be solved by a conjugate gradient method based on a non-standard inner product. In this talk I will discuss such a method, analyze its mathematical properties, and study its numerical behavior. I will discuss what happens when the conditions that guarantee real eigenvalues fail. This leads to the analysis of Krylov subspace methods based on indefinite inner products, a challenging research topic with many open questions.

**Miro Rozložník**
*Numerical Behavior of Saddle Point Solvers*
For largescale saddle point problems, the application of exact iterative schemes and preconditioners may be computationally expensive. In practical situations, only approximations to the inverses of the diagonal block or the related cross-product matrices are considered, giving rise to inexact versions of various solvers. Therefore, the approximation effects must be carefully studied. In this talk we study numerical behavior of several iterative Krylov subspace solvers applied to the solution of large-scale saddle point problems. Two main representatives of the segregated solution approach are analyzed: the Schur complement reduction method, based on an (iterative) elimination of primary variables and the null-space projection method which relies on a basis for the null-space for the constraints. We concentrate on the question what is the best accuracy we can get from inexact schemes solving either Schur complement system or the null-space projected system when implemented in finite precision arithmetic. The fact that the inner solution tolerance strongly influences the accuracy of computed iterates is known and was studied in several contexts.
In particular, for several mathematically equivalent implementations we study the influence of inexact solving the inner systems and estimate their maximum attainable accuracy. When considering the outer iteration process our rounding error analysis leads to results similar to ones which can be obtained assuming exact arithmetic. The situation is different when we look at the residuals in the original saddle point system. We can show that some implementations lead ultimately to residuals on the the roundoff unit level independently of the fact that the inner systems were solved inexactly on a much higher level than their level of limiting accuracy. Indeed, our results confirm that the generic and actually the cheapest implementations deliver the approximate solutions which satisfy either the second or the first block equation to the working accuracy. In addition, the schemes with a corrected direct substitution are also very attractive. We give a theoretical explanation for the behavior which

was probably observed or it is already tacitly known. The implementations that we pointed out as optimal are actually those which are widely used and suggested in applications.

**Valeria Simoncini** (University of Bologna)
*Spectral Analysis of Saddle Point Matrices with Indefinite Leading Blocks*
After a brief introduction to the spectral analysis of saddle point matrices, in this talk we present new estimates for the eigenvalue intervals for symmetric saddle-point and regularised saddle-point matrices in the case where the (1,1) block may be indefinite. These generalise known results for the definite (1,1) case. We also discuss spectral properties of the equivalent augmented formulation, which is an alternative to explicitly dealing with the indefinite (1,1) block.

**Fei Xue** (University of Maryland)
*Numerical Solutions of Eigenvalue Problems with Spectral Transformations*
This work concerns computational methods for solving the generalized eigenvalue problems used for linear stability analysis of models of incompressible flows. Our emphasis is on efficient iterative solution of the sequence of linear systems that arise when subspace iteration and implicitly restarted Arnoldi (IRA) method are applied to detect a few interior or rightmost eigenvalues. We first study some strategies to speed up the iterative solution of the linear systems with multiple right hand sides to efficiently implement shift-invert matrix vector products in subspace iteration steps. We provide new insights into the tuning of preconditioning proposed by Spence in this context, and propose a simplified approach to fully utilize the virtue of tuning. We show the deflation of converged Schur vectors leads to significant savings of computational cost. We also show by numerical experiments that the solutions to the linear systems in previous subspace iteration steps can be used as good initial guess for the linear system in the current step. We then explore the inexact shift-invert IRA method. Our tuning of preconditioning, based on the use of solved linear systems in the present and previous IRA cycles, proves to be quite effective for iterative solution of the linear system in the current IRA step. We provide some comments on the different motivations of our tuning and that by Freitag and Spence. Some refined heuristic estimation of relaxation of IRA is discussed. Finally we show that solvers with recycled subspace for the solution of a sequence of slowly changing linear systems can be used to further reduce the computational cost.

**5. Miscellaneous uses of saddle point problems.**

**Robert Bridson** (University of British Columbia)
*Making Cholesky Work for Saddle Point Systems*
We look at ways to transform saddle point systems, under certain assumptions, to be solved with tools for symmetric positive definite systems. Ill present ongoing work on ordering algorithms that provide for a guaranteed LDLT factorization (with no numerical pivoting) that often preserve sparsity, and then look at a new transformation of constrained regularized least squares problems to plain least-squares problems amenable to solution or preconditioning with QR. Ill draw my examples from a number of applications in computational physics and computer graphics, such as unsteady Stokes with free surfaces (leading to viscous jet buckling), isometrically deforming shells, and solid-fluid coupling.

**Abderrahman El Maliki** (University of Laval)
*Mixed Preconditioned Conjugate Projected Gradient Algorithm for Unilateral Contact Problems*
A new Mixed Preconditioned Conjugate Projected Gradient (MPCPG) algorithm is developed to solve minimization problems with equality or inequality constraints arising from unilateral contact problems in structural mechanics. The solution of the displacement unknowns and Lagrange multipliers are solved simultaneously. The success of the MPCPG is garanteed by 2?2 block symmetric indefinite preconditioner. The preconditioner involves the solution of two subsystems associated respectively with the displacement and the Lagrange multipliers. Both of these two subsystems have to be solved efficiently. We present some numerical results to illustrate the potential of the proposed method.

**Eldad Haber** (Emory University)
*Constraints in Image Registration*
In this talk we present a new and general framework for image registration when having additional constraints on the transformation. We demonstrate that registration without constraints leads to arbitrary results depending on the regularization, and in particular produces non-physical deformations. Having additional constraints based on the images introduces prior knowledge and contributes to reliability and uniqueness of the registration. In particular we consider recently proposed locally rigid transformations as an example. We propose a constrained optimization framework and discuss the different linear systems that arise for different formulations of this problem.

**Rich Lehoucq** (Sandia National Laboratories)
*Preconditioning Constrained Eigenvalue Problems*
My presentation introduces a robust preconditioning scheme for the numerical solution of the leftmost eigenvalues and corresponding eigenvectors of a constrained eigenvalue problem. This constrained eigenvalue problem is congruent to a nonsym-

metric eigenvalue problem with nontrivial Jordan blocks associated with infinite eigenvalues. The proposed preconditioning scheme is relevant to the application of Krylov subspace methods and preconditioned eigensolvers. The two key results are a semi-orthogonal decomposition and a transformation process that implicitly combines a preconditioning step followed by abstract projection onto the subspace associated with the finite eigenvalues. Numerical results demonstrate the effectiveness of the preconditioning scheme.

**Alison Ramage** (University of Strathclyde)
*Saddle Point Problems in Liquid Crystal Modelling*
Saddle-point problems occur frequently in liquid crystal modelling. For example, they arise whenever Lagrange multipliers are used for the pointwise-unit-vector constraints in director modelling, or in both general director and order tensor models when an electric field is present that stems from a constant voltage. Furthermore, in a director model with associated constraints and Lagrange multipliers, together with a coupled electric-field interaction, a particular double saddle-point structure arises. This talk will focus on a simple example of this type and discuss appropriate numerical solution schemes.

# 4 Outcome of the Meeting

The presentations at the meeting showed the state of the art in solution algorithms for saddle point problems across the spectrum of problem domains where they arise. Major contemporary issues that were considered include the wide variety of multilevel algorithms of use in the solution of problems in models of incompressible flow; the development of efficient solvers for problems in which randomness and uncertainty play a role; the use of nonstandard norms in development and analysis of saddle point solvers; the development of efficient preconditioning strategies for interior point methods in optimization; and the use of efficient algorithms for saddle point problems arising in new settings, including image registration, computer graphics of use in computer animation, and saddle point problems with indefinite and singular components.

We received many positive comments about the quality of the meeting and the opportunities to interact. We have also heard of several new collaborations initiated through the meeting, and we believe the prospects of cross-fertilization within the field have been significantly enhanced by it.

# References

[1] M. Benzi, G. Golub and J. Liesen, Numerical solution of saddle point problems, *Acta Numerica*, **14** (2005), 1–137.

[2] L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. v. B. Waanders (Editors), *Large-Scale PDE-Constrained Optimization*, Springer-Verlag, Berlin, 2003.

[3] H. C. Elman, D. J. Silvester and A. J. Wathen, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, Oxford, 2005.

[4] A. Forsgren, P. E. Gill and M. H. Wright, Interior Methods for Nonlinear Optimization, *SIAM REVIEW* **44** (2002), 525–597.