Advances in Linear Algebra and Huge-Scale Optimization

Abstracts

Luca Bergamaschi (Università di Padova)

Efficient preconditioning of the normal equations for large LPs

Interior point method for constrained linear programming problems require the repeated solution of symmetric saddle-point (KKT) type linear systems which can be reduced to symmetric positive definite systems (normal equations) by simple linear algebra. For large scale problems, the Cholesky factorization of the normal equations matrix can be excessively dense to be kept in memory and this prevents the use direct methods for the solution of the related linear system. In this presentation we will present a dynamic sparsification of the matrix of the constraint. This gives rise to an approximate matrix of the normal equation whose exact Cholesky factorization is used as the preconditioner. Sparsity of the proposed preconditioner is controlled by a user defined parameter. We will provide an estimate of the eigenvalue distribution of the preconditioned matrix as well as preliminary numerical results on very large constrained linear problems.

Jordi Castro Universitat Politècnica de Catalunya

A specialized interior-point algorithm for huge minimum convex cost flows in bipartite networks

Computing the Newton direction is the most computational expensive step of interior-point methods. The specialized interior-point algorithm implemented in the BlockIP package computes that direction by a combination of Cholesky factorizations and conjugate gradients. In this work we apply this algorithmic approach to solve very large instances of minimum cost flows problems in bipartite networks, for convex objective functions with diagonal Hessians (i.e., either linear, quadratic or separable nonlinear objectives). After analyzing the theoretical properties of the interior-point method for this kind of problems, we provide extensive computational experiments with linear and quadratic instances of up to one billion arcs and 200 and five million nodes in each subset of the node partition. For linear and quadratic instances our approach is compared with the barriers algorithms of CPLEX (both standard path-following and homogeneous-self-dual); for linear instances it is also compared with the different algorithms of the state-of-the-art network flow solver LEMON (namely: network simplex, capacity scaling, cost scaling and cycle cancelling). The specialized interior-point approach significantly outperformed the other approaches in most of the linear and quadratic transportation instances tested. In particular, it always provided a solution within the time limit and it never exhausted the 192 Gigabytes of memory of the server used for the runs. For assignment problems the network algorithms in LEMON were the most efficient option.

Daniela di Serafino Università della Campania "Luigi Vanvitelli")

On constraint-preconditioned Krylov solvers for saddle-point linear systems

Constraint preconditioners (CPs) have widely demonstrated their effectiveness on standard and regularized saddle-point systems arising in optimization methods, where the leading block of the system matrix is usually symmetric and positive definite in an appropriate subspace [1,2]. In this case, CPs allow the use of the conjugate gradient method, although both the saddle-point matrix and the preconditioner are indefinite. Here we discuss the design of constraint-preconditioned variants of other Krylov methods, starting from the work in [3] about standard saddle-point systems and extending it to regularized systems. In particular, by focusing on the underlying Lanczos and Arnoldi basis-generation processes, we develop constraint-preconditioned variants of various Krylov solvers, including the Lanczos version of CG, MINRES, SYMMLQ, GMRES(m) and DQGMRES, for regularized saddle-point systems. We illustrate the numerical behavior of these constraint-preconditioned Krylov solvers using symmetric and nonsymmetric systems arising from optimization and fluid-flow simulation. This is joint work with Dominique Orban (GERAD and École Polytechnique, Montréal, QC, Canada).

References

[1] M. Benzi, G.H. Golub, and J. Liesen, Numerical solution of saddle point problems, Acta Numerica, 14:1-137, 2005.
[2] M. D'Apuzzo, V. De Simone, and D. di Serafino, On mutual impact of numerical linear algebra and large-scale optimization with focus on interior point methods, Computational Optimization and Applications, 45:283-310, 2010.
[3] N.I.M. Gould, D. Orban, and T. Rees, Projected Krylov methods for saddle-point systems, SIAM Journal on Matrix Analysis and Applications, 35:1329-1343, 2014.

Anders Forsgren, KTH Royal Institute of Technology

On solving symmetric systems of linear equations arising in optimization

An important part of many optimization methods intended for smooth optimization problems is solving systems of linear equations where the matrix is symmetric. It could be methods based on linear equations as such, for example active-set methods for linear and convex quadratic programming, or methods based on solving nonlinear equations by solving a sequence of linear equations, for example interior methods.

It is not always so clear which equations to solve. Many methods are Newton-based, but it may be desirable to use quasi-Newton approximations of the Hessian. For interior methods, the aim is to compute solutions of a sequence of linear equations such that the solution converges to the solution of the original problem.

We discuss particular applications of solving augmented systems arising in interior methods, where the indefinite augmented system may be identified by a doubly augmented system, which is positive definite.

We also discuss quasi-Newton methods applied to solving a symmetric equation where the matrix is positive definite. We discuss relationship to the method of conjugate gradients and give a family of limited-memory methods that behave as the method of conjugate gradients in exact arithmetic. We illustrate the behavior on quadratic problems and demonstrate that there is a potential to gain overall efficiency by a limited-memory quasi-Newton method.

Finally, we will discuss an application of interior methods for optimization problems that arise in intensity-modulated radiation therapy. We show how higher-order interior methods may be highly useful for solving linear programs that arise.

The talk is based on joint work with David Ek, Lovisa Engberg, Kjell Eriksson, Philip E. Gill, Joshua Griffin, Björn Hårdemark and Tove Odland

Daniel Loghin University of Birmingham

Boundary preconditioning

Large scale problems are commonly solved using iterative methods. These methods are usually combined with preconditioning techniques aimed at rendering the solver performance optimal, i.e., independent of problem size, possibly also independent of other problem parameters. In the case of problems arising from the discretisation of PDE, the design of an efficient preconditioner is linked to the choice of partial differential operator. In particular, a suitable inclusion of the boundary operator in the preconditioning technique is essential. While this is well understood for simple (scalar) PDE, for complex applications this is not always a straightforward task. This is the case of, for example, boundary control problems or problems coupled at an interface.

In this talk I will discuss some classes of problems where a suitable choice of boundary preconditioner ensures the optimal performance of the iterative solver. Analysis will be presented, together with validating numerical experiments.

Bendetta Morini Università di Firenze

Subsampled trust-region methods for finite-sum minimization

Convex and non-convex finite-sum minimization arises in many scientific computing and machine learning applications. Recently, second-order methods have received great attention due to their distinguishing features with respect to first-order methods: resilience to problem ill-conditioning and low sensitivity to parameter tuning.

We propose a new trust-region method which is more efficient than the standard scheme in terms of cost-per iteration as it can employ suitable approximations of the objective function, gradient and Hessian matrix built via random subsampling techniques. We discuss the local and global properties for finding approximate first and second-order optimal points and show results from the numerical experience.

This is a joint work with Stefania Bellavia (University of Firenze) and Natasa Krejic (University of Novi Sad).

Margherita Porcelli Università di Bologna

A new interior point approach for low-rank semidefinite programs

We address the solution of semidefinite programming (SDP) problems in which the primal variable X is expected to be low-rank at optimality, a common situation in relaxations of combinatorial optimization (maximum cut) or in matrix completion. SDPs are solved efficiently using interior-point methods (IPMs), but such algorithms typically converge to a maximum-rank solution. We propose a new IPM approach which works with a low-rank X and gradually reveals the optimal (minimum) rank. Preliminary results show that using alternating directions improves the efficiency of the linear algebra.

This is a joint work with S. Bellavia and J. Gondzio.

Spyros Pougkakiotis University of Edinburgh

Fast solution methods for convex fractional differential equation optimization

Fractional differential equations (FDEs) appear in various fields of applied mathematics. Systems of equations produced by discretizing FDEs can be rather challenging to solve, mainly due to their size. Nevertheless, the discretized matrices enjoy a multilevel Toeplitz structure. Such matrices can be applied to a vector expeditiously, by exploiting the fast Fourier transform. In light of that, numerical FDEs are usually solved using some iterative method, alongside a suitable (usually multilevel circulant) preconditioner.

In this talk, we present some solution methods for convex FDE-optimization. In particular, we demonstrate that iterative methods with multilevel circulant preconditioners can efficiently solve equality constrained FDE optimization problems. Then, we propose a solver for general inequality optimization problems with FDE constraints, using preconditioned iterative methods as sub-solvers. Preliminary numerical results are reported.

Miroslav Rozloznik Czech Academy of Sciences

Numerical behavior of saddle-point solvers

A large amount of work has been devoted to solution techniques for saddle-point problems varying from fully direct approach, through the use of iterative stationary and Krylov subspace methods up to the combination of direct and iterative techniques including preconditioning. Significantly less attention has been paid to the numerical behavior of saddle-point solvers. In this contribution we study the behavior of stationary one-step or two-step splitting iteration methods that are often used also for solving large saddle-point problems. We show that inexact solutions of inner linear systems associated with matrix splittings may considerable influence the accuracy of approximate solutions computed in finite precision arithmetic. Then we extend these results also to the inexact saddle-point schemes and implementations that are based on the Schur-complement and null-space approaches.

David Silvester University of Manchester

Bespoke linear algebra for PDEs with random coefficients

We discuss the key role that bespoke linear algebra plays in modelling PDEs with random coefficients using stochastic Galerkin approximation methods. As a specific example, we consider nearly incompressible linear elasticity problems with an uncertain spatially varying Young's modulus. The uncertainty is modelled with a finite set of parameters with prescribed probability distribution. We introduce a novel three-field mixed variational formulation of the PDE model and focus on the efficient solution of the associated high-dimensional indefinite linear system of equations. Eigenvalue bounds for the preconditioned system are established and shown to be independent of the discretisation parameters and the Poisson ratio. This is joint work with Arbaz Khan and Catherine Powell.

Jemima Tabeart University of Reading and Imperial College London *Accounting for correlated errors in data assimilation: using linear algebra to improve computational efficiency*

In numerical weather prediction, observations of the atmosphere are combined with a previous forecast (or background) in order to produce the best initial state for the updated forecast. This optimisation is done using data assimilation (DA) algorithms, which weight the contribution of the observations and background according to their respective uncertainty. A good initial condition is crucial for good forecasts, but as the time allocated to the DA algorithm is very short, computational efficiency is also important.

In recent years, there has been a huge growth in the use of correlated errors between different observations in DA systems. Correlated observation errors allow users to make better use of existing observations, but estimated covariance matrices are often ill-conditioned. In this talk we will present analysis of the variational DA problem to understand the computational costs associated with introducing correlated observation error matrices. Using numerical linear algebra techniques, we find that the minimum eigenvalue of the observation error covariance matrix is important for the conditioning of the DA problem. This motivates investigation into 'reconditioning methods' which allow users to include correlation information in a more computationally efficient way. We introduce the first theoretical study of two reconditioning methods that are used at NWP centres. Finally, we discuss a case study using the operational Met Office 1D-Var system. We find that the qualitative conclusions of the linear theory hold in the non-linear experimental framework. We show that reconditioning methods improve convergence of the iterative DA system, but alter other aspects of the system such as quality control