Electronic Transactions on Numerical Analysis
Volume 28, 2007–2008

Contents

Abstract.
A regression problem is separable if the model can be represented as a linear combination of functions which have a nonlinear parametric dependence. The Gauss-Newton algorithm is a method for minimizing the residual sum of squares in such problems. It is known to be effective both when residuals are small, and when measurement errors are additive and the data set is large. The large data set result that the iteration asymptotes to a second order rate as the data set size becomes unbounded is sketched here. Variable projection is a technique introduced by Golub and Pereyra for reducing the separable estimation problem to one of minimizing a sum of squares in the nonlinear parameters only. The application of Gauss-Newton to minimize this sum of squares (the RGN algorithm) is known to be effective in small residual problems. The main result presented is that the RGN algorithm shares the good convergence rate behaviour of the Gauss-Newton algorithm on large data sets even though the errors are no longer additive. A modification of the RGN algorithm due to Kaufman, which aims to reduce its computational cost, is shown to produce iterates which are almost identical to those of the Gauss-Newton algorithm on the original problem. Aspects of the question of which algorithm is preferable are discussed briefly, and an example is used to illustrate the importance of the large data set behaviour.

Key Words.
nonlinear least squares, scoring, Newton's method, expected Hessian, Kaufman's modification, rate of convergence, random errors, law of large numbers, consistency, large data sets, maximum likelihood

AMS(MOS) Subject Classifications.
62-07, 65K99, 90-08

Abstract.
We establish decay bounds for the entries of $f(A)$, where A is a sparse (in particular, banded) $n \times n$ diagonalizable matrix and f is smooth on a subset of the complex plane containing the spectrum of A. Combined with techniques from approximation theory, the bounds are used to compute sparse (or banded) approximations to $f(A)$, resulting in algorithms that under appropriate conditions have linear complexity in the matrix dimension. Applications to various types of problems are discussed and illustrated by numerical examples.

Key Words.
matrix functions, sparse and banded matrices, decay rates, linear time algorithms, Chebyshev polynomials, Faber polynomials, density matrix, trace, determinant

AMS(MOS) Subject Classifications.
Primary 65F10, 65F30. Secondary 15A.

40 Convergence theory for inexact inverse iteration applied to the generalised nonsymmetric eigenproblem. *Melina A. Freitag and Alastair Spence*.

Abstract.
In this paper we consider the computation of a finite eigenvalue and corresponding right eigenvector of a large sparse generalised eigenproblem $\mathbf{Ax} = \lambda\mathbf{Mx}$ using inexact inverse iteration. Our convergence theory is quite general and requires few assumptions on $\mathbf{A}$ and $\mathbf{M}$. In particular, there is no need for $\mathbf{M}$ to be symmetric positive definite or even nonsingular. The theory includes both fixed and variable shift strategies, and the bounds obtained are improvements on those currently in the literature. In addition, the analysis developed here is used to provide a convergence theory for a version of inexact simplified Jacobi-Davidson. Several numerical examples are presented to illustrate the theory: including applications in nuclear reactor stability, with $\mathbf{M}$ singular and nonsymmetric, the linearised Navier-Stokes equations and the bounded finline dielectric waveguide.

Key Words.
inexact inverse iteration, nonsymmetric generalised eigenproblem

AMS(MOS) Subject Classifications.
Primary 65F15. Secondary 15A18, 65F50.

65 The structured distance to normality of an irreducible real tridiagonal matrix. *S. Noschese, L. Pasquini, and L. Reichel*.

Abstract.
The problem of computing the distance in the Frobenius norm of a given real irreducible tridiagonal matrix T to the algebraic variety of real normal irreducible tridiagonal matrices is solved. Simple formulas for computing the distance and a normal tridiagonal matrix at this distance are presented. The special case of tridiagonal Toeplitz matrices also is considered.

Key Words.
matrix nearness problem, distance to normality, real tridiagonal matrix, eigenvalue conditioning, Toeplitz matrix

AMS(MOS) Subject Classifications.
65F30, 65F50, 15A57, 65F35

78 Minimization of the spectral norm of the SOR operator in a mixed case. *A. Hadjidimos and P. Stratis*.

Abstract.
In this work we solve the problem of the minimization of the spectral norm of the SOR operator associated with a block two-cyclic consistently ordered matrix $A \in \mathbb{C}^{n,n}$, assuming that the corresponding Jacobi matrix has eigenvalues $\mu \in [-\beta, \beta] \cup$

$[-\imath\alpha, \imath\alpha]$, with $\beta \in [0, 1)$, $\alpha \in [0, +\infty)$ and $\imath = \sqrt{-1}$. Previous results obtained by other researchers are extended.

Key Words.
Jacobi and SOR iteration matrices, block two-cyclic consistently ordered matrix, spectral matrix norm

AMS(MOS) Subject Classifications.
65F10

95 A refined unsymmetric Lanczos eigensolver for computing accurate eigentriplets of a real unsymmetric matrix. *Jean Christophe Tremblay and Tucker Carrington Jr.*

Abstract.
For most unsymmetric matrices it is difficult to compute many accurate eigenvalues using the primitive form of the unsymmetric Lanczos algorithm (ULA). In this paper we propose a modification of the ULA. It is related to ideas used in [J. Chem. Phys. 122 (2005), 244107 (11 pages)] to compute resonance lifetimes. Using the refined ULA we suggest, the calculation of accurate extremal and interior eigenvalues is feasible. The refinement is simple: approximate right and left eigenvectors computed using the ULA are used to form a small projected matrix whose eigenvalues and eigenvectors are easily computed. There is no re-biorthogonalization of the Lanczos vectors and no need to store large numbers of vectors in memory. The method can therefore be used to compute eigenvalues of very large matrices. The idea is tested on several matrices.

Key Words.
eigenproblem, unsymmetric matrices, Lanczos algorithm

AMS(MOS) Subject Classifications.
15A18, 65F15, 65F50

115 Derivation of high-order spectral methods for time-dependent PDE using modified moments. *James V. Lambers.*

Abstract.
This paper presents a reformulation of Krylov Subspace Spectral (KSS) Methods, which build on Gene Golub's many contributions pertaining to moments and Gaussian quadrature, to produce high-order accurate approximate solutions to variable-coefficient time-dependent PDE. This reformulation serves two useful purposes. First, it more clearly illustrates the distinction between KSS methods and existing Krylov subspace methods for solving stiff systems of ODE arising from discretizions of PDE. KSS methods rely on perturbations of Krylov subspaces in the direction of the data, and therefore rely on directional derivatives of nodes and weights of Gaussian quadrature rules. Second, because these directional derivatives allow KSS methods to be described in terms of operator splittings, they facilitate stability analysis. It will be shown that under reasonable assumptions on the coefficients of the problem, certain KSS methods are unconditionally stable. This paper also discusses preconditioning similarity transformations that allow more general problems to benefit from this property.

Key Words.
spectral methods, Gaussian quadrature, variable-coefficient, Lanczos method, stability, heat equation, wave equation

AMS(MOS) Subject Classifications.
65M12, 65M70, 65D32

136 An application of the finite volume method to the bio-heat-transfer-equation in premature infants. *Martin Ludwig, Jochim Koch, and Bernd Fischer*.

Abstract.
In this report the development of a finite volume method for the time-accurate simulation of the temperature distribution in a premature infant inside an incubator or in an open radiant warmer is described. The real geometry of a premature infant is obtained from MRT-images. The infants thermoregulation is modelled by the so-called bio-heat-transfer-equation incorporating source terms and Neumann boundary conditions. The source terms describe the metabolic heat production, the blood flow and the respiratorical water loss whereas the Neumann boundary conditions model the heat transfer by transepidermal water loss, radiation, convection and conduction. The numerical solution is carried out by the developed finite volume method whose spatial discretization is done by a 3D-mesh-generator from CFD. For the time integration a semi-implicit multistep method is used. The arising large, sparse linear systems are efficiently solved with a Krylov subspace method. Some successful test runs using real life data are presented.

Key Words.
premature infant, bio-heat-transfer-equation, finite volume method, BiCGStab

AMS(MOS) Subject Classifications.
65M, 92C

149 A weighted-GCV method for Lanczos-hybrid regularization. *Julianne Chung, James G. Nagy, and Dianne P. O'Leary*.

Abstract.
Lanczos-hybrid regularization methods have been proposed as effective approaches for solving large-scale ill-posed inverse problems. Lanczos methods restrict the solution to lie in a Krylov subspace, but they are hindered by semi-convergence behavior, in that the quality of the solution first increases and then decreases. Hybrid methods apply a standard regularization technique, such as Tikhonov regularization, to the projected problem at each iteration. Thus, regularization in hybrid methods is achieved both by Krylov filtering and by appropriate choice of a regularization parameter at each iteration. In this paper we describe a weighted generalized cross validation (W-GCV) method for choosing the parameter. Using this method we demonstrate that the semi-convergence behavior of the Lanczos method can be overcome, making the solution less sensitive to the number of iterations.

Key Words.
generalized cross validation, ill-posed problems, iterative methods, Lanczos bidiagonalization, LSQR, regularization, Tikhonov

AMS(MOS) Subject Classifications.
65F20, 65F30

168 Variable-precision arithmetic considered perilous — a detective story. *Dirk Laurie*.

The scene of the crime.
In 2002, I was interested in computations to very high precision, having been challenged by Nick Trefethen to compute a number he called τ to 10 000 digits. This

number is the sum of ten other numbers, each of which was defined by a difficult computational problem. The story is told in [F. Bornemann, D. Laurie et al.: *The SIAM 100-Digit Challenge: A Study in High-Accuracy Numerical Computing*, SIAM, PA, 2004].

174 Implementing an interior point method for linear programs on a CPU-GPU system. *Jin Hyuk Jung and Dianne P. O'Leary*.

Abstract.
Graphics processing units (GPUs), present in every laptop and desktop computer, are potentially powerful computational engines for solving numerical problems. We present a mixed precision CPU-GPU algorithm for solving linear programming problems using interior point methods. This algorithm, based on the rectangular-packed matrix storage scheme of Gunnels and Gustavson, uses the GPU for computationally intensive tasks such as matrix assembly, Cholesky factorization, and forward and back substitution. Comparisons with a CPU implementation demonstrate that we can improve performance by using the GPU for sufficiently large problems. Since GPU architectures and programming languages are rapidly evolving, we expect that GPUs will be an increasingly attractive tool for matrix computation in the future.

Key Words.
GPGPU, Cholesky factorization, matrix decomposition, forward and back substitution, linear programming, interior point method, rectangular packed format

AMS(MOS) Subject Classifications.
90C05, 90C51, 15A23, 68W10

190 Quantum dynamical entropy and an algorithm by Gene Golub. *Giorgio Mantica*.

Abstract.
The problem of computing the quantum dynamical entropy introduced by Alicki and Fannes requires the trace of the operator function $F(\Omega) = -\Omega \log \Omega$, where Ω is a non-negative, Hermitean operator. Physical significance demands that this operator be a matrix of large order. We study its properties and we derive efficient algorithms to solve this problem, also implementable on parallel machines with distributed memory. We rely on a Lanczos technique for large matrix computations developed by Gene Golub.

Key Words.
quantum dynamical entropy, large matrices, Lanczos method, Montecarlo techniques

AMS(MOS) Subject Classifications.
65F10, 37M25, 81Q50

206 A generalization of the steepest descent method for matrix functions. *M. Afanasjew, M. Eiermann, O. G. Ernst, and S. Güttel*.

Abstract.
We consider the special case of the restarted Arnoldi method for approximating the product of a function of a Hermitian matrix with a vector which results when the restart length is set to one. When applied to the solution of a linear system of equations, this approach coincides with the method of steepest descent. We show

that the method is equivalent to an interpolation process in which the node sequence has at most two points of accumulation. This knowledge is used to quantify the asymptotic convergence rate.

Key Words.
matrix function, Krylov subspace approximation, restarted Krylov subspace method, restarted Arnoldi/Lanczos method, linear system of equations, steepest descent, polynomial interpolation

AMS(MOS) Subject Classifications.
65F10, 65F99, 65M20