Electronic Transactions on Numerical Analysis
Volume 54, 2021

Contents

AMS Subject Classifications.
15A60, 65F10, 65F35

51 Hypergraph edge elimination—A symbolic phase for Hermitian eigensolvers based on rank-1 modifications.
Karsten Kahl and Bruno Lang.

Abstract.
It is customary to identify sparse matrices with the corresponding adjacency or incidence graphs. For the solution of a linear system of equations using Gaussian elimination, the representation by its adjacency graph allows a symbolic factorization that can be used to predict memory footprints and enables the determination of near-optimal elimination orderings based on heuristics. The Hermitian eigenvalue problem on the other hand seems to evade such treatment at first glance due to its inherent iterative nature. In this paper we prove this assertion wrong by revealing a tight connection of Hermitian eigensolvers based on rank-1 modifications with a symbolic edge elimination procedure. A symbolic calculation based on the incidence graph of the matrix can be used in analogy to the symbolic phase of Gaussian elimination to develop heuristics which reduce memory footprint and computations. Yet, we also show that the question of an optimal elimination strategy remains NP-complete, in analogy to the linear systems case.

Key Words.
symmetric eigenvalue problem, hypergraphs, Gaussian elimination

AMS Subject Classifications.
65F15, 05C50, 05C65, 68R10

68 On the solution of the nonsymmetric T-Riccati equation.
Peter Benner and Davide Palitta.

Abstract.
The nonsymmetric T-Riccati equation is a quadratic matrix equation where the linear part corresponds to the so-called T-Sylvester or T-Lyapunov operator that has previously been studied in the literature. It has applications in macroeconomics and policy dynamics. So far, it presents an unexplored problem in numerical analysis, and both theoretical results and computational methods are lacking in the literature. In this paper we provide some sufficient conditions for the existence and uniqueness of a nonnegative minimal solution, namely the solution with component-wise minimal entries. Moreover, the efficient computation of such a solution is analyzed. Both the small-scale and large-scale settings are addressed, and Newton-Kleinman-like methods are derived. The convergence of these procedures to the minimal solution is proven, and several numerical results illustrate the computational efficiency of the proposed methods.

Key Words.
T-Riccati equation, M-matrices, minimal nonnegative solution, Newton-Kleinman method

AMS Subject Classifications.
65F30, 15A24, 49M15, 39B42, 40C05

Abstract.
We present robust and highly parallel multilevel non-overlapping Schwarz preconditioners to solve an interior penalty discontinuous Galerkin finite element discretization of a system of steady-state, singularly perturbed reaction-diffusion equations with a singular reaction operator using a GMRES solver. We provide proofs of convergence for the two-level setting and the multigrid V-cycle as well as numerical results for a wide range of regimes.

Key Words.
multilevel, Schwarz, preconditioner, multigrid, reaction, diffusion, discontinuous, Galerkin

AMS Subject Classifications.
65N55, 65N30, 65J10, 65F08

Abstract.
Image denoising using mean curvature leads to the problem of solving a nonlinear fourth-order integro-differential equation. The nonlinear fourth-order term comes from the mean curvature regularization functional. In this paper, we treat this high-order nonlinearity by reducing the nonlinear fourth-order integro-differential equation to a system of first-order equations. Then a cell-centered finite difference scheme is applied to this system. With a lexicographical ordering of the unknowns, the discretization of the mean curvature functional leads to a block pentadiagonal matrix. Our contributions are fourfold: (i) we give a new method for treating the high-order nonlinearity term; (ii) we express the discretization of this term in terms of simple matrices; (iii) we give an analysis for this new method and establish that the error is of first order; and (iv) we verify this theoretical result by illustrating the convergence rates in numerical experiments.

Key Words.
image denoising, mean curvature, cell-centered finite difference method, numerical analysis

AMS Subject Classifications.
68U10, 94A08, 65N06, 65N12

Abstract.
In this paper, we present the concept of perturbation bounds for the right eigenvalues of a quaternionic matrix. In particular, a Bauer-Fike-type theorem for the right eigenvalues of a diagonalizable quaternionic matrix is derived. In addition, perturbations of a quaternionic matrix are discussed via a block-diagonal decomposition and the Jordan canonical form of a quaternionic matrix. The location of the standard right eigenvalues of a quaternionic matrix and a sufficient condition for the stability

of a perturbed quaternionic matrix are given. As an application, perturbation bounds for the zeros of quaternionic polynomials are derived. Finally, we give numerical examples to illustrate our results.

Key Words.
quaternionic matrices, left eigenvalues, right eigenvalues, quaternionic polynomials, Bauer-Fike theorem, quaternionic companion matrices, quaternionic matrix norms

AMS Subject Classifications.
15A18, 15A66

150 New fractional pseudospectral methods with accurate convergence rates for fractional differential equations.
Shervan Erfani, Esmail Babolian, and Shahnam Javadi.

Abstract.
The main purpose of this paper is to introduce generalized fractional pseudospectral integration and differentiation matrices using a family of fractional interpolants, called fractional Lagrange interpolants. We develop novel approaches to the numerical solution of fractional differential equations with a singular behavior at an end-point. To achieve this goal, we present efficient and stable methods based on three-term recurrence relations, generalized barycentric representations, and Jacobi-Gauss quadrature rules to evaluate the corresponding matrices. In a special case, we prove the equivalence of the proposed fractional pseudospectral methods using a suitable fractional Birkhoff interpolation problem. In fact, the fractional integration matrix yields the stable inverse of the fractional differentiation matrix, and the resulting system is well-conditioned. We develop efficient implementation procedures for providing optimal error estimates with accurate convergence rates for the interpolation operators and the proposed schemes in the L^2-norm. Some numerical results are given to illustrate the accuracy and performance of the algorithms and the convergence rates.

Key Words.
convergence rate, fractional differential equations, fractional Birkhoff interpolation, fractional pseudospectral methods, fractional Lagrange interpolants, singularity

AMS Subject Classifications.
26A33, 41A05, 65M06, 65M12, 65L60

176 Pseudo-linear convergence of an additive Schwarz method for dual total variation minimization.
Jongho Park.

Abstract.
In this paper, we propose an overlapping additive Schwarz method for total variation minimization based on a dual formulation. The $O(1/n)$-energy convergence of the proposed method is proven, where n is the number of iterations. In addition, we introduce an interesting convergence property of the proposed method called *pseudo-linear convergence*; the energy decreases as fast as for linearly convergent algorithms until it reaches a particular value. It is shown that this particular value depends on the overlapping width δ, and the proposed method becomes as efficient as linearly convergent algorithms if δ is large. As the latest domain decomposition methods for total variation minimization are sublinearly convergent, the proposed

method outperforms them in the sense of the energy decay. Numerical experiments which support our theoretical results are provided.

Key Words.
domain decomposition method, additive Schwarz method, total variation minimization, Rudin–Osher–Fatemi model, convergence rate

AMS Subject Classifications.
65N55, 65Y05, 65K15, 68U10

198 Continuous time integration for changing type systems.
Sebastian Franz.

Abstract.
We consider variational time integration using continuous Galerkin-Petrov methods applied to evolutionary systems of changing type. We prove optimal-order convergence of the error in a cGP-like norm and conclude the paper with some numerical examples and conclusions.

Key Words.
evolutionary equations, changing type system, continuous Galerkin-Petrov, space-time approach

AMS Subject Classifications.
65J08, 65J10, 65M12, 65M60

210 Multigrid reduction in time with Richardson extrapolation.
R. D. Falgout, T. A. Manteuffel, B. O'Neill, and J. B. Schroder.

Abstract.
The advent of exascale computing will leave many users with access to more computational resources than they can simultaneously use, e.g., billion-way parallelism. In particular, this is true for time-dependent simulations that limit parallelism to the spatial domain. One method to add parallelism in time to existing simulation codes and thus take advantage of ever larger compute resources is Multigrid Reduction in Time (MGRIT). The goal is to achieve a smaller time-to-solution through parallelism in time. In this paper, MGRIT is enhanced with Richardson extrapolation in a cost-efficient way to produce a parallel-in-time method with improved accuracy. Overall, this leads to a large improvement in the accuracy per computational cost of MGRIT.

Key Words.
parallel time integration, high-performance-computing, multigrid-reduction-in-time, extrapolation-methods

AMS Subject Classifications.
65M20, 65M55, 65F08, 65F10, 65Y05

234 Additive Schwarz preconditioners for a localized orthogonal decomposition method.
Susanne C. Brenner, José C. Garay, and Li-yeng Sung.

Abstract.
We investigate a variant of the localized orthogonal decomposition method (Henning and Peterseim, [Multiscale Model. Simul., 11 (2013), pp. 1149–1175] and Målqvist

and Peterseim, [Math. Comp., 83 (2014), pp. 2583–2603]) for elliptic problems with rough coefficients. The construction of the basis of the multiscale finite element space is based on domain decomposition techniques, which is motivated by the recent work of Kornhuber, Peterseim, and Yserentant [Math. Comp., 87 (2018), pp. 2765–2774]. We also design and analyze additive Schwarz domain decomposition preconditioners for the resulting discrete problems.

Key Words.
multiscale, localized orthogonal decomposition, domain decomposition, additive Schwarz

AMS Subject Classifications.
65N12, 65N30, 65N55

Abstract.
Atmospheric tomography, i.e., the reconstruction of the turbulence profile in the atmosphere, is a challenging task for adaptive optics (AO) systems for the next generation of extremely large telescopes. Within the AO community, the solver of first choice is the so-called Matrix Vector Multiplication (MVM) method, which directly applies the (regularized) generalized inverse of the system operator to the data. For small telescopes this approach is feasible, however, for larger systems such as the European Extremely Large Telescope (ELT), the atmospheric tomography problem is considerably more complex, and the computational efficiency becomes an issue. Iterative methods such as the Finite Element Wavelet Hybrid Algorithm (FEWHA) are a promising alternative. FEWHA is a wavelet-based reconstructor that uses the well-known iterative preconditioned conjugate gradient (PCG) method as a solver. The number of floating point operations and the memory usage are decreased significantly by using a matrix-free representation of the forward operator. A crucial indicator for the real-time performance are the number of PCG iterations. In this paper, we propose an augmented version of FEWHA, where the number of iterations is decreased by 50% using an augmented Krylov subspace method. We demonstrate that a parallel implementation of augmented FEWHA allows the fulfilment of the real-time requirements of the ELT.

Key Words.
adaptive optics, atmospheric tomography, inverse problems, augmented Krylov subspace methods

AMS Subject Classifications.
65R32, 65Y05, 65Y20, 65B99, 85-08, 85-10

Abstract.
Filon-Clenshaw-Curtis (FCC) rules rank among the rapid and accurate quadrature rules for computing oscillatory integrals. In the implementation of the FCC rules, when the oscillator of the integral is nonlinear, its inverse has to be evaluated at several points. In this paper we suggest an approach based on interpolation, which

leads to a class of modifications of the original FCC rules in such a way that the modified rules do not involve the inverse of the oscillator function. In the absence of stationary points, two reliable and efficient algorithms based on the modified FCC (MFCC) rules are introduced. For each algorithm, an error estimate is verified theoretically and then illustrated by some numerical experiments. Also, some numerical experiments are carried out in order to compare the convergence speed of the two algorithms. In the presence of stationary points, an algorithm based on composite MFCC rules on graded meshes is developed. An error estimate is derived and illustrated by some numerical experiments.

Key Words.
Filon-Clenshaw-Curtis rule, oscillatory integral, nonlinear oscillator, stationary point, graded mesh

AMS Subject Classifications.
MSC 65D30, MSC 65T40

Abstract.
In this paper we introduce a finite element method for the Stokes equations with a massless immersed membrane. This membrane applies normal and tangential forces affecting the velocity and pressure of the fluid. Additionally, the points representing this membrane move with the local fluid velocity. We design and implement a high-accuracy cut finite element method (CutFEM) which enables the use of a structured mesh that is not aligned with the immersed membrane, and we formulate a time discretization that yields an unconditionally energy stable scheme. We prove that the stability is not restricted by the parameter choices that constrained previous finite element immersed boundary methods and illustrate the theoretical results with numerical simulations.

Key Words.
immersed boundary method, finite element method, numerical stability, CutFEM, unfitted methods

AMS Subject Classifications.
65N12, 65N30, 74F10

Abstract.
We study the numerical evaluation of the Legendre polynomials P_n on the interval $[-1, 1]$ via a three-term recurrence. We prove that in a neighborhood of an endpoint, the computed approximation exactly agrees with the line tangent to P_n at this endpoint. As a consequence, we obtain sharp error bounds for the recurrence.

Key Words.
Legendre polynomials, three-term recurrence, floating-point arithmetic

AMS Subject Classifications.
65D20, 65Q30, 33F05

Abstract.
The aim of this paper is to present a Nyström-type method for the numerical approximation of the solution of Volterra integral equations of the second kind having highly oscillatory kernels. The method is based on a mixed quadrature scheme which combines the classical product rule with a dilation quadrature formula. The convergence and the stability of the method are investigated and the accuracy of the presented approach is assessed by some numerical tests. The proposed procedure is also applied to the computation of initial scattering data related to the initial value problem associated to the Korteweg-de Vries equation.

Key Words.
Volterra integral equation, highly oscillatory kernels, Nyström method, mixed quadrature scheme, Korteweg-de Vries equation

AMS Subject Classifications.
65R20, 41A05, 45D05.

Abstract.
The well-known Reynolds equation is typically used to compute the pressure distribution for elasto-hydrodynamic contacts of parts, as, for instance, in radial slider bearings. In order to resolve local pressure phenomena like edge loading, a higher spatial resolution is needed. This causes problems for stationary solvers, like Gauss-Seidel iteration, which are well suited for the occurring nonlinearities. These problems can be overcome by applying multigrid methods. Since the Reynolds equation is nonlinear, expensive nonlinear multigrid methods are expected to be required. This paper introduces an approach to combine a linear multigrid method with a Gauss-Seidel solver on the finest level, which yields a similar convergence behavior as a nonlinear multigrid method but at much lower computational cost. The formulations are general so that analogous applications of the Reynolds equation, as, for instance, for axial slider bearings or hydrodynamic piston-liner contacts, are straightforward.

Key Words.
radial slider bearings, Reynolds equation, multigrid

AMS Subject Classifications.
76A20, 65N55, 76M12

Abstract.
We propose a two-level iterative scheme for solving general sparse linear systems. The proposed scheme consists of a sparse preconditioner that increases the norm of

the skew-symmetric part relative to the rest and makes the main diagonal of the coefficient matrix as close to the identity as possible so that the preconditioned system is as close to a shifted skew-symmetric matrix as possible. The preconditioned system is then solved via a particular Minimal Residual Method for Shifted Skew-Symmetric Systems (MRS). This leads to a two-level (inner and outer) iterative scheme where the MRS has short-term recurrences and satisfies an optimality condition. A preconditioner for the inner system is designed via a skew-symmetry-preserving deflation strategy based on the skew-Lanczos process. We demonstrate the robustness of the proposed scheme on sparse matrices from various applications.

Key Words.
symmetrizer, skew-symmetrizer, Krylov subspace method, shifted skew-symmetric system, skew-Lanczos method

AMS Subject Classifications.
65F08, 65F10, 65F50

Abstract.
Differential equations on metric graphs can describe many phenomena in the physical world but also the spread of information on social media. To efficiently compute the optimal setup of the differential equation for a given desired state is a challenging numerical analysis task. In this work, we focus on the task of solving an optimization problem subject to a linear differential equation on a metric graph with the control defined on a small set of Dirichlet nodes. We discuss the discretization by finite elements and provide rigorous error bounds as well as an efficient preconditioning strategy to deal with the large-scale case. We show in various examples that the method performs very robustly.

Key Words.
complex networks, optimal Dirichlet control, preconditioning, saddle point systems, error estimation

AMS Subject Classifications.
65F08, 65F50, 65M60, 65N15,

Abstract.
A standard approach to calculate the roots of a univariate polynomial is to compute the eigenvalues of an associated *confederate* matrix instead, such as, for instance, the companion or comrade matrix. The eigenvalues of the confederate matrix can be computed by Francis's QR algorithm. Unfortunately, even though the QR algorithm is provably backward stable, mapping the errors back to the original polynomial coefficients can still lead to huge errors. However, the latter statement assumes the use of a non-structure-exploiting QR algorithm. In [J. L. Aurentz et al., *Fast and backward stable computation of roots of polynomials*, SIAM J. Matrix Anal. Appl., 36 (2015), pp. 942–973] it was shown that a structure-exploiting QR algorithm for companion matrices leads to a structured backward error in the companion matrix.

The proof relied on decomposing the error into two parts: a part related to the recurrence coefficients of the basis (a monomial basis in that case) and a part linked to the coefficients of the original polynomial. In this article we prove that the analysis can be extended to other classes of comrade matrices. We first provide an alternative backward stability proof in the monomial basis using structured QR algorithms; our new point of view shows more explicitly how a structured, decoupled error in the confederate matrix gets mapped to the associated polynomial coefficients. This insight reveals which properties have to be preserved by a structure-exploiting QR algorithm to end up with a backward stable algorithm. We will show that the previously formulated companion analysis fits into this framework, and we analyze in more detail Jacobi polynomials (comrade matrices) and Chebyshev polynomials (colleague matrices).

Key Words.
backward error, structured QR, linearization, comrade matrix, colleague matrix, companion matrix

AMS Subject Classifications.
65H04, 65F15, 65G50

Abstract.
In this paper we propose a suitable combination of two collocation methods based on the zeros of Jacobi polynomials in order to approximate the solution of Fredholm integral equations on $[-1,1]$. One of the main interesting aspects of this procedure is that our approach is cheaper than the usual collocation method based on standard Lagrange interpolation using Jacobi zeros. Moreover, we can successfully manage functions with algebraic singularities at the endpoints. The error of the method is comparable with the error of the best polynomial approximation in suitable spaces of functions, equipped with the weighted uniform norm. The convergence and the stability of the method is proved, and some numerical tests, which confirm the theoretical estimates, are provided.

Key Words.
Fredholm integral equations, collocation method, polynomial approximation, orthogonal polynomials, modified moments

AMS Subject Classifications.
65R20, 45B05, 65D05, 65D32

Abstract.
We investigate moduli of planar circular quadrilaterals that are symmetric with respect to both coordinate axes. First we develop an analytic approach that reduces this problem to ODEs and then devise a numerical method to find out the accessory parameters. This method uses the Schwarz equation to determine a conformal

mapping of the unit disk onto a given circular quadrilateral. We also give an example of a circular quadrilateral for which the value of the conformal modulus can be found in analytic form. This example is used to validate the numeric calculations. We also apply another method, the so called hpFEM, for the numerical calculation of the moduli. These two different approaches provide results agreeing with high accuracy.

Key Words.
conformal capacity, conformal modulus, quadrilateral modulus, hp-FEM, numerical conformal mapping

AMS Subject Classifications.
65E05, 31A15, 30C85

Abstract.
We analyze three iterative methods that have been proposed in the computational physics community for the reconstruction of memory kernels in a stochastic delay differential equation known as the generalized Langevin equation. These methods use the autocorrelation function of the solution of this equation as input data. Although they have been demonstrated to be useful, a straightforward Laplace analysis does not support their conjectured convergence. We provide more detailed arguments to explain the good performance of these methods in practice. In the second part of this paper we investigate the solution of the generalized Langevin equation with a perturbed memory kernel. We establish sufficient conditions including error bounds such that the stochastic process corresponding to the perturbed problem converges to the unperturbed process in the mean square sense.

Key Words.
generalized Langevin equation, Laplace transform, strong convergence

AMS Subject Classifications.
60G15, 65R32, 65C20

Abstract.
In this work, we examine the numerical solution of a 1D distributed-order space-fractional diffusion equation. Discretizing the given problem by means of an implicit finite difference scheme based on the shifted Grünwald-Letnikov formula, the resulting linear systems show a Toeplitz structure. Then, by using well-known spectral tools for Toeplitz sequences, we determine the corresponding symbol describing its asymptotic eigenvalue distribution as the matrix size diverges. The spectral analysis is performed under different assumptions with the aim of estimating the intrinsic asymptotic ill-conditioning of the involved matrices. The obtained results suggest to precondition the involved linear systems with either a Laplacian-like preconditioner or with more general τ-preconditioners. Due to the symmetric positive definite nature of the coefficient matrices, we opt for the preconditioned conjugate gradient

method, and we compare the performances of our proposal with a Strang circulant alternative given in the literature.

Key Words.
fractional diffusion equations, Toeplitz matrices, spectral distribution, preconditioning

AMS Subject Classifications.
35R11, 15B05, 15A18, 65F08

Abstract.
Multigrid methods have proven to be an invaluable tool to efficiently solve large sparse linear systems arising in the discretization of Partial Differential Equations (PDEs). Algebraic multigrid methods and in particular adaptive algebraic multigrid approaches have shown that multigrid efficiency can be obtained without having to resort to properties of the PDE. Yet the required setup of these methods poses a not negligible overhead cost. Methods from machine learning have attracted attention to streamline processes based on statistical models being trained on the available data. Interpreting algebraically smooth error as an instance of a Gaussian process, we develop a new, data driven approach to construct adaptive algebraic multigrid methods. Based on Gaussian a priori distributions, kriging interpolation minimizes the mean squared error of the a posteriori distribution, given the data on the coarse grid. Going one step further, we exploit the quantification of uncertainty in the Gaussian process model in order to construct efficient variable splittings. Using a semivariogram fit of a suitable covariance model we demonstrate that our approach yields efficient methods using a single algebraically smooth vector.

Key Words.
multigrid, adaptivity, Gaussian processes

AMS Subject Classifications.
60G15, 65F08, 65F10, 65N22, 65N55

Abstract.
We introduce a new polynomial preconditioner for solving the discretized Helmholtz equation preconditioned with the complex shifted Laplace (CSL) operator. We exploit the localization of the spectrum of the CSL-preconditioned system to approximately enclose the eigenvalues by a non-convex 'bratwurst' set. On this set, we expand the function $1/z$ into a Faber series. Truncating the series gives a polynomial, which we apply to the Helmholtz matrix preconditioned by the shifted Laplacian to obtain a new preconditioner, the Faber preconditioner. We prove that the Faber preconditioner is nonsingular for degrees one and two of the truncated series. Our numerical experiments (for problems with constant and varying wavenumber) show that the Faber preconditioner reduces the number of GMRES iterations.

Key Words.
Helmholtz equation, shifted Laplace preconditioner, iterative methods, GMRES, preconditioning, Faber polynomials, 'bratwurst' sets

AMS Subject Classifications.
65F08, 65F10, 30C10, 30C20

558 Computing the matrix fractional power with the double exponential formula.
Fuminori Tatsuoka, Tomohiro Sogabe, Yuto Miyatake, Tomoya Kemmochi, and Shao-Liang Zhang.

Abstract.
Two quadrature-based algorithms for computing the matrix fractional power A^α are presented in this paper. These algorithms are based on the double exponential (DE) formula, which is well-known for its effectiveness in computing improper integrals as well as in treating nearly arbitrary endpoint singularities. The DE formula transforms a given integral into another integral that is suited for the trapezoidal rule; in this process, the integral interval is transformed into an infinite interval. Therefore, it is necessary to truncate the infinite interval to an appropriate finite interval. In this paper, a truncation method, which is based on a truncation error analysis specialized to the computation of A^α, is proposed. Then, two algorithms are presented—one where A^α is computed with a fixed number of abscissa points and one with A^α computed adaptively. Subsequently, the convergence rate of the DE formula for Hermitian positive definite matrices is analyzed. The convergence rate analysis shows that the DE formula converges faster than Gaussian quadrature when A is ill-conditioned and α is a non-unit fraction. Numerical results show that our algorithms achieve the required accuracy and are faster than other algorithms in several situations.

Key Words.
matrix function, matrix fractional power, numerical quadrature, double exponential formula

AMS Subject Classifications.
65F60, 65D30

581 A stable matrix version of the fast multipole method: stabilization strategies and examples.
Difeng Cai and Jianlin Xia.

Abstract.
The fast multipole method (FMM) is an efficient method for evaluating matrix-vector products related to certain discretized kernel functions. The method involves an underlying FMM matrix given by a sequence of smaller matrices (called generators for convenience). Although there has been extensive work in designing and applying FMM techniques, the stability of the FMM and the stable FMM matrix factorization have rarely been studied. In this work, we propose techniques that lead to stable operations with FMM matrices. One key objective is to give stabilization strategies that can be used to provide low-rank approximations and translation relations in the FMM satisfying some stability requirements. The standard Taylor expansions used in FMMs yield basis generators susceptible to instability. Here, we introduce some scaling factors to control the relevant norms of the generators and give a rigorous analysis of the bounds of the entrywise magnitudes. The second objective is to use the one-dimensional case as an example to provide an intuitive construction of FMM matrices satisfying some stability conditions and then convert an FMM matrix into a hierarchically semiseparable (HSS) form that admits stable

ULV-type factorizations. This bridges the gap between the FMM and stable FMM matrix factorizations. The HSS construction is done analytically and does not require expensive algebraic compression. Relevant stability studies are given, which show that the resulting matrix forms are suitable for stable operations. Note that the essential stabilization ideas are also applicable to higher dimensions. Extensive numerical tests are given to illustrate the reliability and accuracy.

Key Words.
numerical stability, fast multipole method, FMM matrix, scaling factor, low-rank approximation, HSS matrix

AMS Subject Classifications.
65F30, 65F35, 15A23, 15A60

610 On the regularization effect of stochastic gradient descent applied to least-squares.
Stefan Steinerberger.

Abstract.
We study the behavior of the stochastic gradient descent method applied to $\|Ax - b\|_2^2 \to \min$ for invertible matrices $A \in \mathbb{R}^{n\times n}$. We show that there is an explicit constant c_A depending (mildly) on A such that

$$\mathbb{E}\,\|Ax_{k+1} - b\|_2^2 \leq \left(1 + \frac{c_A}{\|A\|_F^2}\right)\|Ax_k - b\|_2^2 - \frac{2}{\|A\|_F^2}\left\|A^T A(x_k - x)\right\|_2^2.$$

This is a curious inequality since the last term involves one additional matrix multiplication applied to the error $x_k - x$ compared to the remaining terms: if the projection of $x_k - x$ onto the subspace of singular vectors corresponding to large singular values is large, then the stochastic gradient descent method leads to a fast regularization. For symmetric matrices, this inequality has an extension to higher-order Sobolev spaces. This explains a (known) regularization phenomenon: an energy cascade from large singular values to small singular values acts as a regularizer.

Key Words.
stochastic gradient descent, Kaczmarz method, least-squares, regularization

AMS Subject Classifications.
65F10, 65K10, 65K15, 90C06, 93E24