

Statement on Research

Gunnar Martinsson, August 28, 2009.

Overview of research interests: My research is primarily concerned with the development and analysis of computational techniques for linear algebraic calculations, and for calculating approximate solutions to linear Partial Differential Equations (PDEs), in particular elliptic PDEs.

Work on fast solvers for linear PDEs: The discretization of linear PDEs such as Laplace's equation and Helmholtz' equation results in linear systems of algebraic equations that can be very large. Matrices of size $10^6 \times 10^6$ are becoming fairly standard, and problem sizes two or three orders of magnitude larger still have been handled. Existing fast techniques for solving such large systems rely with few exceptions on so called *iterative methods* (such as GMRES, multigrid, *etc*) in which a sequence of successive improvements to an initial guess are generated by applying the system matrix to a sequence of vectors generated by the algorithm. In contrast, I seek to construct *direct* methods that solve the equation in one shot. Such methods are capable of solving technologically important equations that are not amenable to existing techniques (such as scattering problems close to resonant frequencies). The direct methods in addition tend to be more robust, and are much faster than iterative methods in environments where many equations involving the same coefficient matrix are to be solved.

That "direct" solvers of the type that I am interested in can in principle be constructed was demonstrated by the work of Beylkin, Coifman, and Rokhlin in the early 1990s, and actual algorithms were developed by Hackbusch a few years later. My research contribution has been to develop (in collaboration with Vladimir Rokhlin and Mark Tygert) and implement techniques that in practice beat the performance of state-of-the-art iterative methods in important environments. Specifically, such direct solvers have been developed for linear boundary value problems (BVPs) in two dimensions, for certain volume problems in two dimensions, and for BVPs on axi-symmetric geometries in three dimensions.

Work on numerical linear algebra: In numerical analysis, one commonly faces the task of determining an approximation to a given linear operator or matrix in the form of a low rank matrix. In 2006 Vladimir Rokhlin, Mark Tygert, and I developed an algorithm for this task that relies on randomized sampling to reduce the effective dimension of the range of the matrix. That algorithm, and various refined versions developed by myself and others, have proven capable of outperforming well-established deterministic methods (such as Krylov-subspace methods or rank-revealing QR factorizations) in almost every environment of interest. In some situations, they dramatically outperform classical algorithms. To mention two examples: (1) The randomized algorithms are capable of computing an accurate low-rank approximation to a matrix with a *single pass* over the data. (2) For a general $m \times n$ matrix of numerical rank k , classical algorithms require $O(mnk)$ operations, while randomized ones require $O(mn \log(k))$ operations, with a small constant of proportionality.

The randomized methods for computing low rank approximations to matrices have generated much interest in the last two years. It appears entirely possible that they might become industry standard for computing low-rank approximations, in particular in situations involving very large matrices that have to be processed out-of-core and on parallel machines.

The big picture: The overarching goal of my research is to develop numerical methods for elliptic PDEs that perform *far* better than existing techniques. In many environments, speed-ups of several orders of magnitude and computational accuracy to ten digits or more can be achieved. That such performance is in principle attainable by using integral equation formulations, high-order quadrature rules, and “fast” linear solvers (such as, *e.g.*, those based on the *Fast Multipole Method*) has been known since the work of Vladimir Rokhlin, Leslie Greengard, and others, in the 1980s. However, such techniques have not been widely adopted. In my view, the primary reasons are:

- (1) Existing fast solvers based on iterative methods are not robust enough for industrial strength implementations and are not fast enough for real-time simulations.
- (2) Methods for discretizing integral equations on surfaces in \mathbb{R}^3 are sorely lacking, in particular for non-smooth surfaces involving corners and edges.
- (3) Coding these methods from scratch requires much labor (and skill), and there is very little finished software available.

My research on direct solvers is aimed at overcoming (1). I would argue that it in fact has done so for several key problems in two dimensions. My current work is aimed at extending this work to problems in three dimensions, and to developed the discretization techniques mentioned in point (2).

Addressing point (3) will require a multi-year effort of a sizeable research group that includes numerical analysts and software development experts. However, a successful effort in this direction would have a massive impact on our scientific computing capabilities since it would accelerate several ubiquitously occurring computations (such as constructing approximate solutions to Laplace’s, Helmholtz’, and Maxwell’s equations). In certain important environments such as scattering problems, imaging, and engineering design, speed-ups by two orders of magnitude or more are possible. Much of my work is aimed at convincing decision makers of the value of such a program and at showing that its objectives are achievable by developing the core technologies and by demonstrating their power by solving model problems outside the reach of current methodologies. The long term goal is to secure the resources required for this program, and to build a new generation of software for linear boundary value problems.

Recent vita high-lights: A comprehensive list of journal articles and research presentation is provided in the attached vita. This list includes:

- Invited to give a plenary tutorial lecture at the *Neural Information Processing Systems (NIPS) Conference* in Vancouver in December 2009.
- Invited to submit a review article on randomized techniques in numerical linear algebra to *SIAM Review*. (Joint work with J. Tropp and my student N. Halko.)
- Review article on fast direct solvers for linear PDEs in *Acta Numerica 2009*. (Joint work with L. Greengard, D. Gueyffier, and V. Rokhlin.)
- Three funded NSF awards (totalling \$1.1M), including one Career award.
- Advisor for 4 doctoral students.

Textbook: I have been solicited by several publishers to write a book describing recent advances in my field of research (so called *Fast Multipole Methods*). I have signed a contract with Cambridge University Press and am currently writing the book.

Appendix: Specific research achievements:

This section summarizes my research contributions in different areas.

- (1) ***Fast direct solvers for elliptic PDEs:*** Let $A_N x = f$ denote a linear system of equations arising upon the discretization of a linear PDE or integral equation. A persisting theme of my research is to construct approximations to the solution operator $B = A_N^{-1}$ in different environments. The algorithms I am interested in are highly accurate (typically $\|B - A_N^{-1}\| < 10^{-10}$ or better), and have computational complexity that scales linearly or close to linearly with problem size N , with a small constant of proportionality. At this point, I have (sometimes with V. Rokhlin) developed such algorithms in the following environments:
 - (a) ***Integral operators in 1D:*** If A_N results from the discretization of almost any of the standard integral operators on the line, then A_N^{-1} can be constructed in $O(N)$ operations to high accuracy, very rapidly. (Inversion of a $100\,000 \times 100\,000$ matrix takes a few seconds on a standard desktop PC.) Similar results hold for Toeplitz matrices. See [15, 17].
 - (b) ***Boundary Integral Equation (BIE) in 2D:*** When A_N results from the discretization of any of the standard boundary integral operators associated with Laplace's equation on a finite domain in two dimensions, very fast inversion is possible (under weak constraints on the boundary of the domain) [14]. Similar results hold for Helmholtz' equation provided that the domain is not larger than a few hundred wave-lengths across. These results were implemented to very rapidly evaluate electro-static interactions in multi-phase dielectric media [6].
 - (c) ***Very large scatterers in two dimensions:*** For certain geometries (essentially scatterers that are very elongated), direct construction of the scattering matrix is possible even for scatterers that are thousands of wave-lengths long [16].
 - (d) ***Finite element discretization in 2D:*** Suppose that A_N results upon the finite element discretization of an equation of the form $-\Delta u + a \cdot \nabla u + b u = f$ on a domain in two dimensions. The coefficient functions a and b may be non-constant. Then A_N^{-1} can often be constructed in $O(N(\log N)^2)$ operations, with a low constant or proportionality [8]. Inversion of a matrix with $N = 10^6$ was executed in a few minutes on a desktop PC. The prototypical scheme of [8] has been significantly improved by my student Adrianna Gillman and we expect to publish an $O(N)$ version shortly.
 - (e) ***BIEs on rotationally symmetric surfaces in 3D:*** My student Patrick Young has constructed a fast direct solver for the BIEs associated with Laplace's equation on axi-symmetric geometries in \mathbb{R}^3 . Non axi-symmetric loads are handled using the FFT. The scheme is capable of inverting a $320\,000 \times 320\,000$ matrix to ten digits of accuracy in three minutes on a standard desktop PC. Application of the inverse takes 0.4s [20].
 - (f) ***General Boundary Integral Equation (BIE) in 3D:*** Preliminary work on inverting the BIEs associated with surface integral equations in \mathbb{R}^3 (for general surfaces) is reported in [3].

- (2) **Randomized methods for approximating matrices:** Let A be an $m \times n$ matrix and let ε denote a computational precision. Suppose that A admits an approximate factorization $\|A - BC\| \leq \varepsilon$ where B has k rows, and C has k columns. How do you find B and C ? How do you find the minimal integer k for which precision ε can be attained? Along with Vladimir Rokhlin, Mark Tygert, and Joel Tropp, I have recently demonstrated that *randomized* methods can dramatically outperform previously existing deterministic methods in most environments. Specific results include:
- (a) **The matrix A can rapidly be applied to vectors:** When $x \mapsto Ax$ can be evaluated rapidly, the textbook approach is to construct B and C via a Krylov subspace method such as Lanczos or Arnoldi. Methods based on randomized sampling methods match the CPU count of Krylov methods, while improving robustness, and dramatically loosening communication constraints [19, 5]. The latter point is important as parallelization is becoming imperative in many areas of scientific computing.
 - (b) **À posteriori error estimation:** The early versions of the randomized scheme either required the numerical rank of A to be known in advance, or that a separate error estimator be implemented. More recent versions integrate à posteriori error estimation into the randomized scheme at essentially no additional cost [4].
 - (c) **Improved error bounds:** [4] provides an error analysis of the randomized methods that significantly improves upon previously existing results. In particular, it is demonstrated that for a fixed rank k , it is possible to construct a factorization whose precision is arbitrarily close to the theoretically optimal one.
 - (d) **The matrix A has additional structure:** Suppose that A does not itself have low-rank, but can be hierarchically tessellated into submatrices that do (this is the characteristic property of the matrices A_N discussed in point (1) above). That randomized sampling can be used to approximate such “structured” matrices is demonstrated in [7].
- (3) **Fast numerical methods for solving equations on lattices:** When modeling atomic crystals and periodic composite materials, it is common that the mathematical equations are formulated on a finite integer lattice. There currently exist FFT-based extremely fast methods for solving the lattice equations, but these methods fundamentally work only on very simple geometries with periodic boundary conditions. I have developed methods that work on almost any geometry and with any boundary conditions. While the FFT scales (almost) linearly with the number of lattice nodes, my methods scale linearly with the number of *loaded* nodes, and with the number of nodes on the boundary.
- (a) **Design of lattices with specified bandgaps:** A simple technique for designing lattices with prescribed photonic and phononic bandgaps is described in [11]. (To my bemusement, this is my most highly cited article.)
 - (b) **Construction and analysis of a lattice fundamental solution:** The fundamental solution of a simple square or cubic lattice is quite well studied in the literature. Improvements upon some of these earlier results, and a

generalization to a broad class of linear equations on mono-atomic and multi-atomic lattices is given in [10]. A detailed asymptotic analysis is presented in [12].

- (c) **Homogenization via the fundamental solution:** The lattice fundamental solution can be analyzed to yield homogenized equations to any order of accuracy. The analysis gets particularly interesting for mechanical lattices since certain pathological geometries may have stiffness tensors with zero eigenvalues (for instance, the simple square lattice in 2D becomes degenerate when the nodes are pin-jointed), see [9].
- (d) **Lattice potential theory:** A lattice equivalent of the boundary integral equations of classical potential theory is derived in [13]. This paper also demonstrates that the resulting equations behave like Fredholm equations of the second kind, and are thus well suited for numerical work.
- (e) **Fast numerical methods:** With my student Adrianna Gillman, I have developed fast numerical methods that use the lattice boundary equations derived in [13] to very rapidly solve lattice equations. Due to the theoretically optimal scaling of the computational complexity, a lattice with 10^{12} nodes involving 10^6 loaded nodes can be handled on a desktop PC.

(4) **Miscellaneous:**

- (a) **Existence theory for solutions to the Eikonal equation:** The Eikonal equation for a function $u = u(x, y)$ reads $u_x^2 + u_y^2 = f(x, y)$. [2] provides a detailed existence and (non-) uniqueness results for solutions near points where $f(x, y) = 0$.
- (b) **The “interpolatory decomposition” of a matrix:** Given a matrix A of (approximate) low rank, it is often useful to find a set of columns that form a well-conditioned (approximate) basis for its range, and a set of rows that form a well-conditioned (approximate) basis for the range of A^* . [1] describes and analyzes algorithms for doing so. [18] provides analogous results for interpolation of finite-dimensional function spaces .

REFERENCES

- [1] H. Cheng, Z. Gimbutas, P. G. Martinsson, and V. Rokhlin, *On the compression of low rank matrices*, SIAM J. Sci. Comput. **26** (2005), no. 4, 1389–1404 (electronic). MR MR2143489 (2006a:65067)
- [2] E. Cornea, R. Howard, and P.G. Martinsson, *Solutions near singular points to the eikonal and related first-order nonlinear partial differential equations in two dimensions*, Differential and Integral Equations **14** (2001), 1441–1468.
- [3] L. Greengard, D. Gueyffier, P.G. Martinsson, and V. Rokhlin, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numer. **18** (2009), 243–275. MR MR2506042
- [4] N. Halko, P.G. Martinsson, and J. Tropp, *Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions*, 2009, Invited article for SIAM Review. Currently in review.
- [5] E. Liberty, F. Woolfe, P.G. Martinsson, V. Rokhlin, and M. Tygert, *Randomized algorithms for the low-rank approximation of matrices*, Proc. Natl. Acad. Sci. USA **104** (2007), no. 51, 20167–20172. MR MR2366406
- [6] P.G. Martinsson, *Fast evaluation of electro-static interactions in multi-phase dielectric media*, J. Comput. Phys. **211** (2006), no. 1, 289–299.
- [7] ———, *Approximation of structured matrices via randomized sampling*, Tech. Report 0806.2339, arXiv.org, 2008.
- [8] ———, *A fast direct solver for a class of elliptic partial differential equations*, J. Sci. Comput. **38** (2009), no. 3, 316–330. MR MR2475654
- [9] P.G. Martinsson and I. Babuška, *Homogenization of materials with periodic truss or frame micro-structures*, Math. Models Methods Appl. Sci. **17** (2007), no. 5, 805–832. MR MR2325840 (2008e:74066)
- [10] ———, *Mechanics of materials with periodic truss or frame micro-structures*, Arch. Ration. Mech. Anal. **185** (2007), no. 2, 201–234. MR MR2317789 (2008g:74070)
- [11] P.G. Martinsson and A.B. Movchan, *Vibrations of lattice structures and phononic bandgaps*, The Quarterly Journal of Mechanics and Applied Mathematics **56** (2003), 45–64.
- [12] P.G. Martinsson and G.J. Rodin, *Asymptotic expansions of lattice green's functions*, Proceedings of the Royal Society A (2002), no. 458, 2609–2622.
- [13] ———, *Boundary algebraic equations for lattice problems*, Proc. R. Soc. A **465** (2009), no. 2108, 2489–2503.
- [14] P.G. Martinsson and V. Rokhlin, *A fast direct solver for boundary integral equations in two dimensions*, J. Comp. Phys. **205** (2005), no. 1, 1–23.
- [15] ———, *An accelerated kernel independent fast multipole method in one dimension*, SIAM Journal of Scientific Computing **29** (2007), no. 3, 1160–11178.
- [16] ———, *A fast direct solver for scattering problems involving elongated structures*, Journal of Computational Physics **221** (2007), 288–302.
- [17] P.G. Martinsson, V. Rokhlin, and M. Tygert, *A fast algorithm for the inversion of general toeplitz matrices*, Computers and Mathematics with Applications **50** (2005), no. 5, 741 – 752.
- [18] ———, *On interpolation and integration in finite-dimensional spaces of bounded functions*, Tech. Report YALEU/DCS/RR-1317, Dept. of Computer Science, Yale University, 2005.
- [19] ———, *A randomized algorithm for the approximation of matrices*, Tech. Report 1361, Dept. of Computer Science, Yale University, 2006.
- [20] P. Young and P.G. Martinsson, *A direct solver for the rapid solution of boundary integral equations on axisymmetric surfaces in three dimensions*, 2009, In review.