STANFORD 50: State of the Art & Future Directions of Computational Mathematics & Numerical Computing

A conference celebrating the 50th anniversary of
George Forsythe's arrival at Stanford
and Gene Golub's 75th birthday


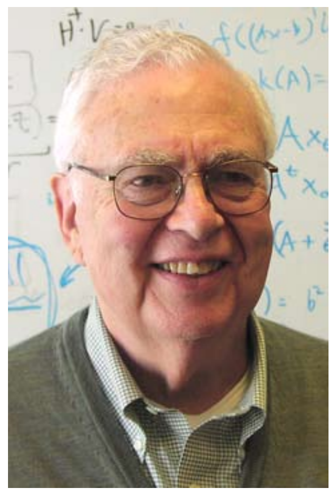

Incorporating the

Eighth Bay Area Scientific Computing Day (BASCD)

Stanford University
March 29–31, 2007

Contents

Welcome from the Organizers!

Stanford 50:
State of the Art and Future Directions of
Computational Mathematics and Numerical Computing

Stanford University, March 29–31, 2007
http://compmath50.stanford.edu

A scientific conference, incorporating the
Eighth Bay Area Scientific Computing Day (BASCD)

2007 marks the 50th anniversary of the arrival of George Forsythe at Stanford University. George ushered in a new era of computational mathematics both at Stanford and elsewhere. Over the past 50 years, Stanford has produced a continuous stream of outstanding scholars in computational mathematics. This progeny now inhabits the higher reaches of a great number of universities and has contributed much to science, industry, and commerce.

2007 also marks the 75th birthday of Gene Golub, who can rightfully claim to have carried the mantle after Forsythe's death in 1972. Gene is universally recognized as ambassador at large for scientific computation. The conference is to celebrate these milestones and to explore the rich future of this important field.

Talks are by invitation, with no parallel sessions. Everyone is invited to attend the meeting. Graduate students and junior scientists (PhD completed within the ten years prior to the meeting) have been invited to contribute to a poster session. Judges and attendees will select the best posters in both categories and the authors will be invited to talk on the final day.

2007 is also the 50th birthday of Fortran, and 60 years since the birth of numerical analysis with John von Neumann and Herman Goldstine's paper on numerical stability and backward error. For more numerical analysis history, we refer participants to *Milestones in Matrix Computation: The Selected Works of Gene H. Golub With Commentaries*, by Raymond Chan, Chen Greif, and Dianne O'Leary, published by Oxford University Press, 2007 (ISBN 978-0-19-920681-0). (Order forms are available at the registration desk.)

Welcome everyone to spring at Stanford! Let us celebrate the groundwork laid by George and Gene's NA Group and by Gene's SCCM Program as iCME takes over the lead in scientific computing at Stanford.

Thank you for joining us on this golden and double-diamond anniversary, and *Happy Birthday Gene* for the moment four weeks ago (28 February 2007 $24{:}00^{-}$).

Conference Schedule

Thursday, March 29, 2007

Time	Event	Location/Page
Registration & Opening		**Hewlett 200 Auditorium**
7:30am	*Registration desk opens*	
8:45– 9:00am	Gene Golub *Welcome and opening remarks*	
First Session	**Chair: James Varah**	**Hewlett 200 Auditorium**
9:00– 9:25am	Cleve Moler, The MathWorks, Inc. *Recollections of a Stanford NA groupie*	p15
9:25– 9:50am	Beresford Parlett, Univ of California, Berkeley *Stanford from 1958 to 1961*	p16
9:50–10:15am	Richard Brent, Australian National University *George Forsythe's last paper*	p9
10:15–10:40am	Paul Saylor, Univ of Illinois, Urbana-Champaign *Stanford's Foresight and Forsythe's Stanford*	p16
Second Session	**Chair: Richard Bartels**	**Hewlett 200 Auditorium**
11:00–11:25am	Pete Stewart, Univ of Maryland *A residual inverse power method*	p17
11:25–11:50am	Bill Gear, Princeton University *Future directions in petascale computing: Explicit methods for implicit problems*	p12
11:50–12:15pm	Paul Van Dooren, Université Catholique de Louvain *Optimizing PageRank by choosing outlinks*	p17
12:15–12:40pm	Sabine Van Huffel, Katholieke Universiteit Leuven *The impact of numerical linear algebra in computational biomedical signal processing*	p18
Third Session	**Chair: Victor Pereyra**	**Hewlett 200 Auditorium**
2:00– 2:25pm	Bertil Gustafsson, Stanford University *High order one-step difference methods for wave propagation*	p14
2:25– 2:50pm	Chris Paige, McGill University *Accuracy of Ritz values from a given subspace*	p15
2:50– 3:15pm	Bart De Moor, Katholieke Universiteit Leuven *Numerical linear algebra in subspace system identification*	p10
3:15– 3:40pm	Linda Petzold, Univ of California, Santa Barbara *Future directions in computational systems biology*	p16

Thursday, March 29, 2007 continued

Time	Event	Location/Page
Fourth Session	**Chair: Haesun Park**	**Hewlett 200 Auditorium**
4:00– 4:25pm	Philip Gill, Univ of California, San Diego *Iterative methods for generalized saddle-point problems*	p13
4:25– 4:50pm	Stephen Wright, Univ of Wisconsin, Madison *Finding sparse solutions of underdetermined systems: Gradient projection approaches*	p19
4:50– 5:15pm	Michael Ferris, Univ of Wisconsin, Madison *Optimization modeling: Recent enhancements and future extensions*	p11
5:15– 5:40pm	Michael Foster, Lenore Mullin, and Eun Park, CISE CCF, NSF *Grand challenges in computational mathematics* *and numerical/symbolic computing: An NSF view*	p15
Welcome Reception and Poster Session		**Packard**
6:00– 8:00pm	Graduate student posters	p20
	Junior scientist posters	p21
	Judges: Petter Bjørstad, Univ of Bergen Howard Elman, Univ of Maryland Michael Heath, Univ of Illinois, Urbana-Champaign James Nagy, Emory University Andy Wathen, Univ of Oxford	p21

Friday, March 30, 2007

Time	Event	Location/Page
First Session	**Chair: Gerard Meurant**	**Hewlett 200 Auditorium**
9:00– 9:25am	Andy Wathen, Univ of Oxford *Matrix iterations and saddle-point systems:* *From optimization to Navier-Stokes and back*	p19
9:25– 9:50am	Marcus Grote, Universität Basel *Computational wave propagation in bounded and unbounded domains*	p13
9:50–10:15am	Liliana Borcea, Rice University *Model reduction and electrical impedance tomography*	p9
10:15–10:40am	Martin Gander, Université de Genève *A best approximation problem with application to parallel computing*	p12
Second Session	**Chair: Walter Gander**	**Hewlett 200 Auditorium**
11:00–11:25am	Chen Greif, Univ of British Columbia *Block preconditioners for saddle point systems: a junction of linear algebra, constrained optimization, and PDEs*	p13
11:25–11:50am	Roland Freund, Univ of California, Davis *Krylov subspace-based dimension reduction of large-scale linear dynamical systems*	p12
11:50–12:15pm	Bernd Fischer, Universität zu Lübeck *Mathematics meets medicine*	p11
12:15–12:40pm	Steve Vavasis, Univ of Waterloo *An SVD-based approach to nonnegative matrix factorization*	p18
Third Session	**Chair: Åke Björck**	**Hewlett 200 Auditorium**
2:00– 2:25pm	Iain Duff, Rutherford Appleton Laboratory *Combining direct and iterative methods for the solution of large systems in different application areas*	p11
2:25– 2:50pm	Bo Kågström, Umeå Universitet *Product eigenvalue problems: Computing periodic deflating subspaces associated with a specified set of eigenvalues*	p14
2:50– 3:15pm	Ahmed Sameh, Purdue University *A parallel banded system solver*	p16
3:15– 3:40pm	Howard Elman, Univ of Maryland *The stochastic finite element method: Recent results and future directions*	p11

Friday, March 30, 2007 continued

Time	**Event**	**Location/Page**
Fourth Session	**Chair: Michele Benzi**	**Hewlett 200 Auditorium**
4:00– 4:25pm	Andrew Stuart, Univ of Warwick *MCMC in infinite dimensions*	p17
4:25– 4:50pm	Dianne O'Leary, Univ of Maryland *Parallel matrix computation: From the ILLIAC to quantum computing*	p15
4:50– 5:15pm	Grace Wahba, Univ of Wisconsin, Madison *A statistician's debt to numerical analysts*	p18
5:15– 5:40pm	William Kahan, Univ of California, Berkeley *Why I can debug some numerical programs and you can't*	p15
Banquet: Celebrating Gene Golub's 75th Birthday		**Stanford Faculty Club**
6:00–10:00pm	**After-Dinner Speaker:** Charles Van Loan, Cornell University *DoubleDeepSudden Impact* plus toasts from many friends	p19

Saturday, March 31, 2007

Time	Event	Location/Page
First Session	**Chair: Petter Bjørstad**	**Hewlett 200 Auditorium**
9:00– 9:25am	Rob Schreiber, Hewlett-Packard *Manycores in the future*	p16
9:25– 9:50am	Jack Dongarra, Univ of Tennessee *The challenge of multicore and specialized accelerators for mathematical software*	p10
9:50–10:15am	Winner of poster competition	
10:15–10:40am	Winner of poster competition	
Second Session	**Chair: Michael Heath**	**Hewlett 200 Auditorium**
11:00–11:25am	Carl de Boor, Univ of Wisconsin, Madison *Issues in multivariate polynomial interpolation*	p9
11:25–11:50am	Tony Chan, Univ of California, Los Angeles *Duality methods for nonlinear image processing*	p9
11:50–12:15pm	Winner of poster competition	
12:15–12:40pm	Winner of poster competition	
Third Session	**Chair: Nancy Nichols**	**Hewlett 200 Auditorium**
2:00– 2:25pm	Jim Demmel, Univ of California, Berkeley *Suggested extra credit questions for a future edition of Golub & Van Loan*	p10
2:25– 2:50pm	Nick Trefethen, Univ of Oxford *Beating Gauss quadrature*	p17
2:50– 3:15pm	Winner of poster competition	
3:15– 3:40pm	Winner of poster competition	
Panel Discussion: The Next 50 Years		**Hewlett 200 Auditorium**
4:00– 5:45pm	**Moderator: Bill Coughran**, Google Inc. Zhaojun Bai, Univ of California, Davis Margot Gerritsen, Stanford University Tammy Kolda, Sandia National Laboratories Paul Tupper, McGill University	p20

Abstracts

Model reduction and electrical impedance tomography

Liliana Borcea, Rice University

We present a novel inversion algorithm for electrical impedance tomography in two dimensions, based on a model reduction approach. The reduced models are resistor networks that arise in five-point stencil discretizations of the elliptic partial differential equation satisfied by the electric potential, on adaptive grids that are computed as part of the problem. We prove the unique solvability of the model reduction problem for a broad class of measurements of the Dirichlet to Neumann map. The size of the networks (reduced models) is limited by the precision of the measurements. The resulting grids are naturally refined near the boundary, where we make the measurements and where we expect better resolution of the images. To determine the unknown conductivity, we use the resistor networks to define a nonlinear mapping of the data that behaves as an approximate inverse of the forward map. Then we propose an efficient Newton-type iteration for finding the conductivity, using this map. We also show how to incorporate a priori information about the conductivity in the inversion scheme.

George Forsythe's last paper

Richard P. Brent
Australian National University, Australia

In 1949 George Forsythe attended some lectures at UCLA by John von Neumann on the topic of random number generation. Shortly before he died, Forsythe wrote a Stanford report STAN-CS-72-254 inspired by von Neumann's lectures. It was intended that this would form the basis of a joint paper with J. H. Ahrens and U. Dieter, who had discovered related results independently. However, after Forsythe died in April 1972, Don Knuth submitted the Stanford report to *Mathematics of Computation* and it was published with only minor changes as "Von Neumann's comparison method for random sampling from the normal and other distributions" (*Math. Comp.* 26, 1972, 817–826). This was Forsythe's last published paper, with the possible exception of a paper "Variational study of nonlinear spline curves" by E. H. Lee and Forsythe in *SIAM Review* (submitted before Forsythe's death but not published until 1973).

Ahrens and Dieter published a follow-up paper "Extensions of Forsythe's method for random sampling from the normal distribution" (*Math. Comp.* 27, 1973, 927–937), and I published an implementation "Algorithm 488: a Gaussian pseudo-random number generator" (*Comm. ACM* 17, 1974, 704–706).

In this talk I will describe von Neumann's elegant idea for sampling from the exponential distribution, Forsythe's generalization for sampling from a probability distribution whose density has the form $\exp(-G(x))$, where $G(x)$ is easy to compute (e.g., a polynomial), and my refinement of these ideas to give an efficient algorithm for generating pseudo-random numbers with a normal distribution. I will also (very briefly) mention some later developments.

Duality methods for nonlinear image processing

Tony Chan, University of California, Los Angeles

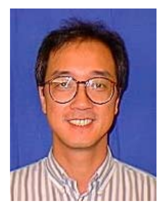

I review how a primal-dual algorithm for minimizing the total variation norm that Gene and I developed with Pep Mulet a decade ago is related to recent work by A. Chambolle on dual projection algorithms and by W. Yin and D. Goldfarb on second-order cone algorithms for the same problem.

Issues in multivariate polynomial interpolation

Carl de Boor, University of Wisconsin, Madison

A quick overview of some basic questions of current interest in multivariate polynomial interpolation.

Suggested extra credit questions for a future edition of Golub & Van Loan

James Demmel, University of California, Berkeley

A generation of researchers and students has benefitted immensely from the textbook "Matrix Computations" by Gene Golub and Charlie Van Loan. In this talk we imagine what a set of "extra credit" questions for a future edition might look like. Here are two examples for section 2.4:

Q1: You have an array of n double precision (64-bit) floating point numbers, and are allowed to do $n-1$ double extended (80-bit) floating point additions to compute their sum (and no other arithmetic operations or "bit-fiddling"). How small can you make the error bound?

Q2: We know that you can multiply matrices in $O(n^\omega)$ operations if and only if you can invert matrices in $O(n^\omega)$ operations.
True or false: You can multiply matrices in $O(n^{(\omega+\epsilon)})$ operations for any $\epsilon > 0$ if and only if you can invert matrices "stably" in $O(n^{(\omega+\epsilon)})$ operations for any $\epsilon > 0$.

Numerical linear algebra in subspace system identification

Bart De Moor, ESAT-SCD K. U. Leuven, Belgium

Subspace identification algorithms calculate a state-space model from input-output measurements of a linear system of the form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + w_k, \\ y_k &= Cx_k + Du_k + v_k, \end{aligned}$$

where $u_k \in \mathbf{R}^m$ and $y_k \in \mathbf{R}^l$, $k = 0, 1, \ldots, N$ are the given measured input and output sequences of the multivariable system with m inputs and l outputs. The consecutive states $x_k \in \mathbf{R}^n$, $k = 0, 1, \ldots$ are unknown, as are the (real) system matrices A, B, C and D of appropriate dimensions. The sequences $v_k \in \mathbf{R}^l$ and $w_k \in \mathbf{R}^n$ represent so-called measurement and process noises, which are supposed to be white, normally distributed zero mean with unknown covariance matrices

$$\mathbf{E}\left[\begin{pmatrix} w(k) \\ v(k) \end{pmatrix} \begin{pmatrix} w(t)^T & v(t)^T \end{pmatrix}\right] = \begin{pmatrix} Q & S \\ S^T & R \end{pmatrix} \delta_{kt} \ ,$$

in which δ_{kt} is the Kronecker delta.

Formidable as it may seem, subspace algorithms manage to identify the order of the system n (the number of difference equations needed to model the data appropriately) and to calculate the matrices A, B, C, D, Q, R and S. They start from building large block Hankel matrices with the input-output data, which are divided up in two parts called the 'past' and the 'future'. Conceptually, it can be shown that the state sequence can be obtained by calculating the intersection between certain 'past' and 'future' vector spaces (in the purely deterministic case, where there is no process and measurement noise), by calculating the principal angles and directions between 'past' and 'future' in the case of stochastic systems (no deterministic input u_k), or by calculating certain oblique projections in the case of the full state-space model given above.

We show that for all these calculations, tools from numerical linear algebra, such as the QR-decomposition, the SVD and eigenproblem solvers, are indispensable.

We also explain that these numerical linear algebra tools lead to new developments in control system design, where optimal control strategies can be computed directly from the data (so-called model free LQR and LQG).

The challenges of multicore and specialized accelerators for mathematical software

Jack Dongarra, University of Tennessee and Oak Ridge National Laboratory

Recent versions of microprocessors exhibit performance characteristics for 32-bit floating-point arithmetic (single precision) that are substantially higher than for 64-bit floating-point (double precision). Examples include Intel's Pentium IV and M processors, AMD's Opteron architectures, IBM's Cell processor, and various GPUs. Single precision operations can be performed up to two times faster on the Pentium and up to ten times faster on the Cell compared to double precision.

Our motivation is to exploit single precision whenever possible and resort to double precision at critical stages while attempting to provide full double precision results. The results described are fairly general and can be applied to various problems in linear algebra, such as solving large sparse systems using direct or iterative methods, and some eigenvalue problems. There are limitations, such as when the problem condition exceeds the reciprocal of single precision accuracy. In that case the double precision algorithm should be used.

Combining direct and iterative methods for the solution of large systems in different application areas

Iain S Duff
CCLRC Rutherford Appleton Laboratory and CERFACS

We are concerned with the solution of sets of linear equations where the matrices are of very high order. We first discuss sparse direct methods and consider the size of problems that they can currently solve. We then discuss the limitations of such methods, where current research is going in moving these limitations, and how far we might expect to go with direct solvers in the near future.

This leads us to the conclusion that very large systems, by which we mean three dimensional problems in more than a million degrees of freedom, require the assistance of iterative methods in their solution. However, even the strongest advocates and developers of iterative methods recognize their limitations when solving difficult problems, that is problems that are poorly conditioned and/or very unstructured. It is now universally accepted that sophisticated preconditioners must be used in such instances.

A very standard and sometimes successful class of preconditioners are based on incomplete factorizations or sparse approximate inverses, but we very much want to exploit the powerful software that we have developed for sparse direct methods over a period of more than thirty years. We thus discuss various ways in which a symbiotic relationship can be developed between direct and iterative methods in order to solve problems that would be intractable for one class of methods alone. In these approaches, we will use a direct factorization on a "nearby" problem or on a subproblem.

We then look at examples using this paradigm in four quite different application areas; the first solves a subproblem and the others a nearby problem using a direct method.

The stochastic finite element method: Recent results and future directions

Howard C. Elman, University of Maryland

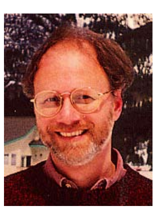

Traditional methods of mathematical modeling depend on the assumption that components of models such as diffusion coefficients or boundary conditions are known. In practice, however, such quantities may not be known with certainty and instead they may be represented as random functions; that is, a random variable for each point in the physical domain.

An approach for performing computational studies of models of this type is the stochastic finite element method, which is a generalization of finite element discretization for deterministic problems designed to handle problems posed with uncertainty. We discuss the use of this methodology to model elliptic partial differential equations when some terms in the problem are not known with certainty, and we explore efficient solution algorithms based on multigrid to solve the large algebraic systems that arise from it.

In addition, we discuss computational issues that will affect the capability of this methodology to generate useful information about uncertain models.

Optimization modeling: Recent enhancements and future extensions

Michael C. Ferris, University of Wisconsin, Madison

Modeling systems are an efficient way to develop the constraints and objectives for nonlinear programming problems. We outline several recent enhancements of such systems that facilitate grid solution techniques, complementarity or equilibrium constraints within optimization problems, model embedding, and explicit formulation of extended nonlinear programming problems. Further extensions of these systems to ease the modeling burden in specific contexts will also be proposed.

Mathematics meets medicine

Bernd Fischer, Universität zu Lübeck, Germany

Computational simulations of real-life phenomena often give rise to large systems and demand clever computational mathematics routines. In this talk we report on three projects along these lines, all of which arise in a medical environment.

The first is concerned with the time-accurate 3D simulation of the temperature distribution of premature infants. The simulation tool is used for hyperthermia planning and for the improvement of warming therapy devices. Its numerical challenge is the solution of the so-called bio-heat-equation equipped with complicated boundary conditions.

The second application deals with image registration. Here, one is looking for a transformation that aligns one image to another. Typical examples include the treatment verification of pre- and post-intervention images, study of temporal series of images, and the monitoring of time evolution of an agent injection subject to a patient-motion. A sound mathematical formulation leads to large-scale optimization problems.

Finally, we report on some activities in the context of nuclear medicine imaging. Because of the long imaging times, patient motion is inevitable and constitutes a serious problem for any reconstruction algorithm. The measured inconsistent projection data lead to reconstruction artifacts that can significantly affect the diagnostic accuracy. We briefly present a new reconstruction scheme that is capable of correcting for patient movement. Again, the mathematical treatment involves the solution of a large-scale numerical computing problem.

Krylov subspace-based dimension reduction of large-scale linear dynamical systems

Roland Freund, University of California, Davis

In recent years, Krylov subspace methods have become widely used tools for dimension reduction of large-scale linear dynamical systems. In this talk, we describe some basic properties of these methods, discuss a few applications, and mention some open problems.

A best approximation problem with application to parallel computing

Martin J. Gander, University of Geneva, Switzerland

The classical best approximation problem is the following: given a real-valued continuous function on a compact interval and a class of functions defined on the same interval, find an element in the class that realizes the distance of the function to the class. If the class is the linear space of polynomials of degree less than or equal to n, and the distance is measured in the L^∞ norm, then the approximation problem is called a Chebyshev best approximation problem.

We are interested in a best approximation problem in a more general setting: we search for a given function $f : C \to C$ the polynomial s_n^* of degree less than or equal to n that minimizes over all s of degree less than or equal to n the quantity

$$\sup_{z\in K}\left|\frac{s(z)-f(z)}{s(z)+f(z)}e^{-lf(z)}\right|,$$

where K is a compact set in C, and l is a non-negative real parameter. The solution of this best approximation problem is important in parallel computing: it leads to the fastest iterative domain decomposition methods.

Future directions in petascale computing: Explicit methods for implicit problems

Bill Gear, Princeton University
(with I. G. Kevrekidis, Princeton University and
Steven L. Lee, LLNL)

A combination of circumstances is causing a renewed interest in explicit methods for what are traditionally viewed as implicit problems when those problems become sufficiently large that massive parallelism is the only realistic computational approach. The difficulty with problems that exhibit diffusion or similar phenomena that lead to stiffness is that conventional methods for handling stiffness with large time steps require the implicit solution of a system of nonlinear equations at each time step (although typically one solution of a linear system is sufficient to get the required accuracy in a Newton-like step). However, the heavy load of inter-processor communication of direct methods in most cases is a significant factor, so iterative methods must be used. Unless there are suitable fast preconditioners to reduce the number of iterations, these may also be sufficiently time-consuming that other methods become more attractive.

While implicit methods *have* to be used if a problem is arbitrarily stiff, if we have some knowledge of the location of the eigenvalues, there are explicit methods that can be competitive. The first work in this area that led to codes was probably the Runge-Kutta Chebyshev methods, although related ideas have been around for some time. In these methods, high-stage RK methods are used, not to get a high order of accuracy (since second order often suffices for many PDEs), but to get extended regions of stability. Recently we have been studying a related class of methods—telescopic projective methods—that can achieve similar goals and also place stability regions in desired locations. These are methods that have the potential to be adaptive and for which second order can be obtained.

A further advantage of these methods is that they can be "wrapped around" single-step legacy codes or microscopic simulators for which we want to explore macroscopic phenomena.

Iterative methods for generalized saddle-point problems

Philip E. Gill, University of California, San Diego

We consider iterative methods for generalized saddle-point problems that arise in interior methods for general nonlinear optimization. Interior methods define a sequence of *KKT systems* that represent the symmetrized (but indefinite) equations associated with Newton's method for satisfying the perturbed optimality conditions. These equations involve both the primal and dual variables and become increasingly ill-conditioned as the optimization proceeds. In this context, an iterative linear solver must not only handle the ill-conditioning but also detect KKT matrices with incorrect inertia.

We focus on the application of the conjugate-gradient method to a certain "doubly-augmented system" that is positive definite with respect to both the primal and the dual variables. This property means that a standard preconditioned CG method involving both primal and dual variables will either terminate successfully or detect if the KKT matrix has wrong inertia.

Constraint preconditioning is a well-known technique for preconditioning the CG method on saddle-point problems. A family of constraint preconditioners is proposed that provably eliminates the inherent ill-conditioning. A considerable benefit of combining constraint preconditioning with the doubly-augmented system is that the preconditioner need not be applied exactly.

The talk is based on joint work with Anders Forsgren and Joshua Griffin.

Block preconditioners for saddle point systems: a junction of linear algebra, constrained optimization, and PDEs

Chen Greif, University of British Columbia, Canada

Saddle point linear systems are ubiquitous in science and engineering applications. The matrices associated with such systems are symmetric and indefinite, and have a 2x2 block structure with a zero block. These systems arise in constrained optimization, in variational formulation of PDEs, and in many other situations. In a large-scale setting it is desirable to take advantage of the block structure, and doing this requires knowing something about the underlying continuous problem and about the spectral structure of the operators involved.

In this talk we discuss solution techniques, addressing the question of which preconditioners should be used. We focus on an augmentation preconditioning technique in which the preconditioners are block diagonal with symmetric positive definite blocks and are based on augmented Lagrangian techniques. Interestingly, it is possible to show analytically that the more rank-deficient the (1,1) block of the original matrix is, the faster a preconditioned iterative scheme converges. Saddle point systems that arise in the time-harmonic Maxwell equations and interior-point methods in optimization are just two examples of situations where this feature of the preconditioner may come in handy. We discuss algebraic connections with other preconditioning approaches, and provide a few numerical examples.

Computational wave propagation in bounded and unbounded domains

Marcus Grote, University of Basel, Switzerland

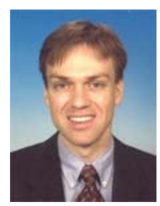

The accurate and reliable simulation of wave phenomena is of fundamental importance in a wide range of engineering applications such as fiber optics, wireless communication, sonar and radar technology, non-invasive testing, ultra-sound imaging, and optical microscopy. To address the wide range of difficulties involved, we consider symmetric interior penalty discontinuous Galerkin (IP-DG) methods, which easily handle elements of various types and shapes, irregular non-matching grids, and even locally varying polynomial order. Moreover, in contrast to standard (conforming) finite element methods, IP-DG methods yield an essentially diagonal mass matrix; hence, when coupled with explicit time integration, the overall numerical scheme remains truly explicit in time. To circumvent the stability (CFL) condition imposed on the time step by the smallest elements in the underlying mesh, we further propose energy conserving explicit local time-stepping schemes.

For problems set in an unbounded domain, an artificial boundary is required to confine the region of interest to a finite computational domain. Then, a nonreflecting boundary condition is required at the artificial boundary, which avoids spurious reflections from it. When a scatterer consists of several components, the use of a single artificial boundary to enclose the entire region of interest becomes too expensive. Instead, it is preferable to embed each component of the scatterer in a separate sub-domain. As waves may bounce back and forth between domains, they are no longer purely outgoing outside the computational domain, so that most standard approaches cannot be used. To overcome this difficulty, we show how to devise exact nonreflecting boundary conditions for multiple scattering problems, which avoid spurious reflections from the artificial boundary.

High order one-step difference methods for wave propagation

Bertil Gustafsson
Uppsala University and Stanford University

We have earlier constructed high order explicit one-step difference methods for linear wave propagation problems with variable coefficients. They use staggered grids, and are norm conserving without any restriction on the coefficients other than boundedness. In particular, they can be used for wave propagation in discontinuous media, without any special treatment of the interior boundaries.

A special advantage is the effective implementation. Once the coefficients of the problem are defined at all grid points, the difference scheme is applied everywhere in the interior without modification. In recent work with B. Engquist, A-K. Tornberg and P. Wahlund, we have applied the same principle when treating real boundaries, like solid walls. The coefficients of the PDE system are given extreme values on one side of the boundary, and in this way the domain of interest can be embedded in a regular domain, keeping the effective implementation of the algorithm. The accuracy is formally brought down to first order because of the boundary treatment. This error is independent of time, and is in most cases dominated by the formally higher order phase error, which grows with time. However, we will show that one can modify the algorithm, for both interior and exterior boundaries, such that second order accuracy is obtained. This is done by a modification of the coefficients near the boundary, which means that the effective implementation is not destroyed.

Product eigenvalue problems: Computing periodic deflating subspaces associated with a specified set of eigenvalues

Bo Kågström, Umeå University, Sweden

Let us consider a *linear discrete-time descriptor system* of the form

$$\begin{aligned} E_k x_{k+1} &= A_k x_k + B_k u_k \\ y_k &= C_k x_k + D_k u_k \end{aligned}$$

with state, input and output vectors x_k, u_k, and y_k, respectively. The matrices A_k, B_k, C_k, D_k, E_k are supposed to be of matching dimensions. In addition, we assume that the system is *periodic* for some period $p \geq 1$, i.e., $A_{k+p} = A_k$ for all integers k, and similarly for B_k, C_k, D_k, E_k. Computational tasks for such systems, which arise naturally from processes that exhibit seasonal or periodic behavior, can often be addressed by solving *product (or periodic) eigenvalue problems.*

For example, if all E_k are square and invertible then the system is *asymptotically stable* if and only if all eigenvalues of the monodromy matrix $\Pi = E_p^{-1} A_p E_{p-1}^{-1} A_{p-1} \ldots E_1^{-1} A_1$ lie strictly inside the unit disk.

Forming the product Π explicitly and applying a standard eigensolver may lead to disastrous numerical results. The most viable way to solve a general product eigenvalue problem is to compute a *generalized periodic real Schur form* (GPRSF), using the *periodic QZ algorithm* (Bojanczyk, Golub and Van Dooren 1992). The algorithm orthogonally transforms the matrix sequences E_k and A_k into upper (quasi-)triangular form at a cost only linear in p. In many applications, it is necessary to have the eigenvalues along the diagonal of the GPRSF in a certain order.

We present a direct method for reordering eigenvalues in the GPRSF of a regular K-cylic matrix pair sequence (A_k, E_k). Following and generalizing existing approaches, reordering consists of consecutively computing the solution to an associated periodic Sylvester-like matrix equation and constructing K pairs of orthogonal matrices. These pairs define an orthogonal K-cyclic equivalence transformation that swaps adjacent diagonal blocks in the GPRSF. An error analysis of this swapping procedure is presented, which extends existing results for reordering eigenvalues in the generalized real Schur form of a regular pair (A, E). Our direct reordering method is used to compute periodic deflating subspace pairs corresponding to a specified set of eigenvalues. This computational task arises in various applications, e.g., solving discrete-time periodic Riccati equations. We present computational experiments that confirm the stability and reliability of the eigenvalue reordering method. (Joint work with Robert Granat and Daniel Kressner at Umeå University.)

Why I can debug some numerical programs and you can't

William Kahan, University of California, Berkeley

The future promises teraflops in your laptop, petaflops in your supercomputer, and the inability to debug numerical programs on either. We discuss the reasons for this situation, and possible solutions.

Recollections of a Stanford NA groupie

Cleve Moler, The MathWorks, Inc.

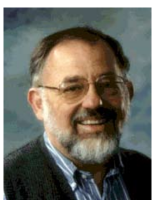

GEF, GHG, Encina, Polya, Serra, 9.63972, 576.

Grand challenges in computational mathematics and numerical/symbolic computing: An NSF view

Lenore Mullin, CISE CCF, NSF
(with Michael Foster and Eun Park)

Optimizing software to keep up with Moore's Law requires Grand Challenges for algorithm, language, and library developers. Is it possible to identify algorithms and data structures pervasive across scientific disciplines with deterministic properties? Can we design and build algebraically closed numeric and symbolic programming languages such that optimal designs can be verified both semantically *and* operationally? Cyber-enabled Discovery and Innovation (CDI), a multi-million dollar initiative at NSF, aims to explore radically new concepts, theories, and tools at the intersection of computational and physical worlds to address these issues. This talk will ask questions and pose answers to the community that will create Grand Challenges for Computational Mathematics.

Parallel matrix computation: from the ILLIAC to quantum computing

Dianne P. O'Leary, University of Maryland

The basic ideas behind parallel matrix computation were developed in the 1960s, 1970s, and 1980s. The single-instruction-multiple-data (SIMD) model was among the first ideas, implemented in machines such as the ILLIAC III and IV. Some later parallel machines implemented dataflow computing ideas.

Today, algorithms developed for these early machines are being revised and reused. For example, graphical processing units (GPUs) are cost-effective and widely-available SIMD parallel processors. An efficient implementation of an interior point algorithm for solving linear programming problems on GPUs, devised in collaboration with Jin Hyuk Jung, will be discussed.

In a second current application, algorithms for parallel matrix computation are not actually executed but instead used to design efficient machines. Specifically, efficient dataflow algorithms for the QR decomposition yield efficient designs for quantum computers, and the talk will focus on this rather surprising application (joint work with Gavin Brennen and Stephen Bullock).

Accuracy of Ritz values from a given subspace

Chris Paige, McGill University, Canada
(with M. E. Argentati, A. V. Knyazev, and I. Panayotov)

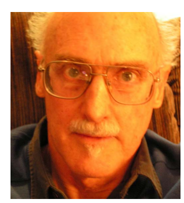

The rate of convergence of iterative methods has always been one of the interests of the Numerical Analysis group at Stanford. For example I read "On the asymptotic directions of the s-dimensional optimum gradient method" by George Forsythe (1968) while refereeing a paper for Gene Golub when I visited the group in 1972 as an early guest of Gene.

In this talk we generalize a well-known eigenvalue result. If $x, y \in \mathrm{C}^n$ are unit-length vectors ($x^H x = y^H y = 1$), where y is an approximation to an eigenvector x of $A = A^H \in \mathrm{C}^{n\times n}$ with $Ax = x\lambda$, $\lambda = x^H A x \in \mathrm{R}$, then the Rayleigh quotient $y^H A y$ satisfies

$$|\lambda - y^H A y| \leq \sin^2 \theta(x, y).\mathrm{spread}(A). \qquad (1)$$

Here, if $\lambda_1(A) \geq \cdots \geq \lambda_n(A)$ are the eigenvalues of A in descending order then $\text{spread}(A) \equiv \lambda_1(A) - \lambda_n(A)$, and $\theta(x, y) \equiv \cos^{-1}|x^H y| \in [0, \pi/2]$ is the acute angle between x and y.

We generalize this result to a higher-dimensional subspace $\mathcal{Y}$ approximating an invariant subspace $\mathcal{X}$ of A. Let $X, Y \in \mathbb{C}^{n\times k}$ be such that $X^H X = Y^H Y = I_k$, where $\mathcal{Y} \equiv \text{range}(Y)$ is an approximation to the invariant subspace $\mathcal{X} \equiv \text{range}(X)$ of A, so that $AX = X.X^H AX$. Let $\lambda(X^H AX)$ and $\lambda(Y^H AY) \in \mathbb{R}^k$ be the vectors of eigenvalues in descending order of $X^H AX$ and $Y^H AY$ respectively. The elements of $\lambda(Y^H AY)$ are called Ritz values in the Rayleigh-Ritz method for approximating the eigenvalues $\lambda(X^H AX)$ of A. Such approximations can be computed for example via the Lanczos or block-Lanczos methods for the Hermitian eigenproblem. Here we obtain new bounds on $\lambda(X^H AX) - \lambda(Y^H AY)$ of a form paralleling (1). We then ask whether such results might contribute to useful "rate of convergence" analyses for iterative eigenproblem methods for large sparse matrices.

Stanford from 1958 to 1961

Beresford Parlett, University of California, Berkeley

I will describe what Forsythe did with his graduate students and in his courses during this period when he was (a) promoting the work of J. H. Wilkinson and (b) acting as midwife for delivery of a Computer Science Department in the School of Humanities and Sciences.

Future directions in computational systems biology

Linda Petzold, University of California, Santa Barbara

As the biological sciences make their way through the 21st century, there will be an enormous need for systems-level analysis and quantitative methods that are well-integrated with the specific structure of the problems and the data. According to the recent NRC Report *Mathematics and 21st Century Biology*, "The exponentially increasing amounts of biological data at all scales of biological organization, along with comparable advances in computing power, create the potential for scientists to construct quantitative, predictive models of biological systems. Broad success would transform basic biology, medicine, agriculture, and environmental science." We illustrate some of the computational challenges in data analysis, model development and simulation via several biological problems.

A parallel banded system solver

Ahmed H. Sameh, Purdue University

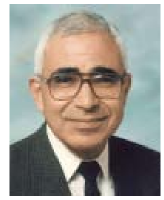

A hybrid parallel algorithm "SPIKE" is proposed for solving banded linear systems that are either dense or sparse within the band. Different versions of the algorithm may be chosen for achieving high performance depending on the parallel architecture and properties of the linear system under consideration. Numerical experiments are presented to demonstrate the effectiveness of the algorithm.

Support partially provided by NSF, DARPA, and Intel.

Stanford's Foresight and Forsythe's Stanford

Paul Saylor, Univ of Illinois, Urbana-Champaign

What Stanford Was Like
What the Time Was Like
Over a Four Year Period
Starting with the Arrival of This New Man
Professor George Forsythe, in 1957
Plus a Bonus Look-Ahead to the Future

Manycores in the future

Rob Schreiber, HP Labs

I'll survey some recent developments in processor chip architecture and the directions in which the field is headed, and consider their implications for parallel programming and scientific computing.

A residual inverse power method

G. W. Stewart, University of Maryland

The inverse power method involves solving shifted equations of the form $(A-\sigma I)v = u$. This talk describes a variant method in which shifted equations may be solved to a fixed reduced accuracy without affecting convergence. The idea is to alter the right-hand side to produce a correction step to be added to the current approximations. The digits of this step divide into two parts: leading digits that correct the solution and trailing garbage. Hence the step can be be evaluated to a reduced accuracy corresponding to the correcting digits. The cost is an additional multiplication by A at each step to generate the right-hand side. Analysis and experiments show that the method is suitable for normal and mildly nonnormal problems.

MCMC in infinite dimensions

Andrew Stuart, Warwick University, UK

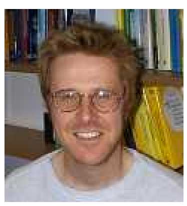

In many application areas it is of interest to sample a probability measure on a space of functions: an infinite dimensional sampling problem. Applications include molecular dynamics, signal processing, econometrics and data assimilation. For this reason it is important to be able to develop efficient algorithms to perform sampling for such problems. Markov Chain Monte Carlo (MCMC) has proved an effective tool in a wide variety of applications and it is natural to develop an understand of its computational complexity in the context of sampling function space.

In this talk I will illustrate the applications of interest; describe their common mathematical structure; and overview the theoretical understanding that has been developed for the sampling of problems with this mathematical structure.

Beating Gauss quadrature

Lloyd N. Trefethen, Oxford University, UK
(with Nicholas Hale)

We all know that Gauss quadrature points are in some sense optimal, and that they can be computed by the marvelous algorithm of Golub and Welsch. But as so often happens in mathematics, the optimality theorem conceals an assumption that may not always be reasonable—in this case, that the quality of a quadrature formula is determined by how high a degree of polynomial it can integrate exactly. If you drop this assumption, you find that alternative quadrature formulas can outperform Gauss for many integrands by a factor of about $\pi/2$. The new formulas involve nearly uniformly spaced nodes, without the usual clustering at endpoints, which can be a big advantage in PDE simulations by spectral methods. We show how to derive such formulas by conformal mapping and point out connections with previous work by Kosloff and Tal-Ezer, Alpert, and others. Fortunately, the Golub-Welsch algorithm is still applicable.

Optimizing PageRank by choosing outlinks

Paul Van Dooren, CESAME, Université Catholique de Louvain, Belgium
(with Cristobald de Kerchove and Laure Ninove)

Google has established its well-known PageRank that classifies the pages of the World Wide Web by scoring each of them. The PageRank of a page represents the probability of presence of a random surfer on that page. This surfer goes with probability c from one page to another page following the hyperlinks, and with probability $1 - c$ from one page to any page on the web with a prescribed probability. The PageRank vector can be seen as the normalized Perron vector of a positive matrix: the Google matrix, taking into account the random surfer motion described above.

If one wishes now to maximize one's own PageRank, one can only control one's own outlinks to other pages. The goal is to increase one element of the Perron vector by changing some elements of the Google matrix. We decribe an optimal strategy for selecting one's outlinks when they can all be chosen arbitrarily, as well as when some of the outinks are imposed in advance. We also address the same problem for a group of people who want to optimise their PageRank sum.

The impact of numerical linear algebra in computational biomedical signal processing

Sabine Van Huffel, Katholieke Universiteit Leuven, Belgium

In biomedical signal processing, the aim is to extract clinically, biochemically or pharmaceutically relevant information (e.g., metabolite concentrations in the brain) in terms of parameters out of low-quality measurements in order to enable an improved medical diagnosis. Typically, biomedical data are affected by large measurement errors, largely due to the non-invasive nature of the measurement process or the severe constraints to keep the input signal as low as possible for safety and bioethical reasons. Accurate and automated quantification of this information requires an ingenious combination of the following issues:

- an adequate pretreatment of the data,
- the design of an appropriate model and model validation,
- a fast and numerically robust model parameter quantification method,
- an extensive evaluation and performance study, using in-vivo and patient data, up to the embedding of the advanced tools into user-friendly interfaces to be used by clinicians.

The underlying computational signal processing problems can be solved by making use of linear algebra, signal processing, system theory and optimisation. In particular, it is shown how computational linear algebra kernels, such as the Singular Value Decomposition (SVD), Principal Component Analysis (PCA), Canonical Correlation Analysis (CCA), Least Squares, Total Least Squares, Independent Component Analysis (ICA), ..., can be used as building blocks for higher-level signal processing algorithms. In addition, the application of these algorithms and their benefits will be briefly illustrated in a variety of case studies, including Magnetic Resonance Spectroscopic Imaging and epileptic seizure detection.

An SVD-based approach to nonnegative matrix factorization

Stephen A. Vavasis, University of Waterloo, Canada

Nonnegative matrix factorization (NNMF) was introduced as a tool for datamining by Lee and Seung in 1999. NNMF attempts to approximate a matrix with nonnegative entries by a product of two low-rank matrices, also with nonnegative entries. We propose an approach for computing a NNMF that is based on an algorithm for singular value decomposition. Preliminary computational tests indicate that this method is able to identify features successfully in realistic datasets.

Parts of this talk represent joint work with Ali Ghodsi of University of Waterloo.

A statistician's debt to numerical analysts

Grace Wahba, University of Wisconsin, Madison

Statisticians, including this one, owe a huge debt to numerical analysts. Where would we be without the Singular Value Decomposition, Spline Algorithms, Matrix Computations (Golub and Van Loan)?

After briefly noting my collaboration with Gene and Michael Heath on Generalized Cross Validation (1979), which laid the foundation for much later work, I will describe some more recent work of my own and collaborators that relies on mathematical programming and convex cone algorithms for the numerical solution of large problems. These include Regularized Kernel Estimation for data sets with dissimilarity data rather than attribute data, and the LASSO-Patternsearch algorithm for finding patterns of high order interactions in risk factor models with large and extremely large attribute vectors.

Matrix iterations and saddle-point systems: From optimization to Navier-Stokes and back

Andy Wathen, Oxford University, UK

The Numerical Analysis group at Stanford and in particular Gene Golub have been deeply involved in iterative linear algebra since the late 1950s.

In this talk we discuss preconditioning and the iterative solution of saddle-point systems and draw together applications in Optimization and the PDEs of incompressible fluid flow, looking forward to the next challenging problems for these methodologies.

Finding sparse solutions of underdetermined systems: Gradient projection approaches

Stephen Wright, University of Wisconsin, Madison

We discuss optimization problems in which the objective consists of a linear least squares term (usually derived from an underdetermined linear system) added to a weighted ℓ-1 norm of the variables. Such problems arise in wavelet-based deconvolution, compressed sensing, and other applications. They have been the subject of intense research in recent years from both a theoretical and an algorithmic perspective. We give an overview of the various approaches, then focus on algorithms of gradient projection type. Some computational results are presented. (Joint work with Rob Nowak and Mario Figuerido.)

Banquet: Celebrating Gene Golub's 75th Birthday
DoubleDeepSudden Impact

After-Dinner Speaker: Charles Van Loan, Cornell University

plus toasts from many friends

Panel Discussion: The Next 50 Years

Moderator: Bill Coughran, Google Inc.

Panelists:
Zhaojun Bai, UC Davis
Margot Gerritsen, Stanford University
Tammy Kolda, Sandia National Laboratories
Paul Tupper, McGill University

   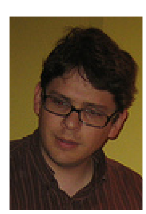

Posters

Graduate Student Posters

Topic Area Presenter(s)	*Title*
Dynamical Systems	
Sotiria Lampoudi UC Santa Barbara	*A computational algorithm for exploring the effect of excluded volume in chemical kinetics*
Stephanie Taylor UC Santa Barbara	*Analyzing phase dynamics of limit cycle systems with application to the circadian clock*
Image Processing	
David Gleich & Chris Maes, Stanford University	*Block SOR for colorizing images: Classical solutions for modern problems*
Dana Paquin Stanford University	*Multiscale deformable registration of medical images*
Image Transforms	
Boris Efros Ben-Gurion University	*DSHAS – Algorithm for real-time calculation of discrete X-ray transform over a sliding window*
Leif Christian Larsen Norwegian University of Science and Technology	*Speeding up transform algorithms for image compression using GPUs*
Least Squares Computations	
Yoza Hida & Jason Riedy, UC Berkeley	*Precise solutions for overdetermined linear least squares problems*
Kourosh Modarresi Stanford University	*Multi-level approach to Tikhonov regularization method*
Matrix Canonical Computations	
Lars Karlsson Umeå Universitet	*GUPTRI3: The next-generation staircase algorithms and library software for canonical structure information*
Matrix Computations	
Robert Granat Umeå Universitet	*Parallel algorithms and library software for Sylvester-type matrix equations*
Jonathan Moussa UC Berkeley	$O(N \log N)$ *tridiagonal eigensolver without the Fast Multipole Method*
Optimization	
Andrew Bradley Stanford University	*Initialization of the limited-memory quasi-Newton Hessian for discretized continuous optimization problems*
Nir Naor Ben-Gurion University	*Finding the best featured phases – Problem simplification*
PDE/CFD	
David Ketcheson University of Washington	*WENOCLAW: A higher-order wave propagation method for general hyperbolic systems*
Sarah Williams UC Davis	*A multiscale hybrid method for compressible fluids, with fluctuations*
Stochastic Methods	
Christian Perret ETH Zurich	*Improving the stability of numerical simulations of quantum dynamical systems using stochastic techniques*
Tiago Requeijo Stanford University	*Group dynamics of phototaxis: Interacting stochastic many-particle systems and their continuum limit*
Support Vector Machines	
Jin Hyuk Jung University of Maryland	*Adaptive constraint reduction for convex quadratic programs, with an application to support vector machines*
Di Zhao Louisiana Tech University	*Non-negative matrix factorization to speed up interior point method of SVM training*

Junior Scientist Posters

Topic Area Presenter(s)	*Title*
Fast Computation	
Laurent Demanet Stanford University	*Fast computation of Fourier integral operators*
Zhengji Zhao Lawrence Berkeley National Laboratory	*The linear scaling 3-dimensional fragment method for petascale nanoscience simulations*
Krylov Methods	
Sou-Cheng Choi Stanford University	*MINRES-QLP: A Krylov subspace method for singular symmetric linear equations and least-squares problems*
James Lambers Stanford University	*The evolution of Krylov subspace spectral methods*
Matrix Computation	
Pedher Johansson Umeå Universitet	*StratiGraph – software tools for matrix canonical computations*
Hyunsoo Kim Georgia Institute of Technology	*A framework of non-negative matrix factorizations via alternating non-negative least squares*
Optimization	
Holly Jin Cardinal Optimization Inc.	*Localization algorithms for ad hoc wireless networks*
Ofer Levi & Yisrael Parmet, Ben-Gurion University of the Negev, Israel	*Matching Pursuit under the gown of Linear Regression*
Uday Shanbhag University of Illinois at Urbana-Champaign	*Equilibrium programming under uncertainty*
PDEs	
Aboubacar Bagayogo University College of Saint-Boniface	*Hybrid grid generation – A symbolic programming approach*
Erik Boman Sandia National Laboratory	*Combinatorial scientific computing*
Vani Cheruvu National Center for Atmospheric Research	*A spectral finite volume/flux corrected transport method for shallow water equations*
Henrik Loef Stanford University	*Multigrid smoothers revisited: Parallelization on multi-core processors*

Judges


Petter Bjørstad
Univ of Bergen


Howard Elman
Univ of Maryland


Michael Heath
Univ of Illinois,
Urbana-Champaign


James Nagy
Emory Univ


Andy Wathen
Univ of Oxford

Participants

A-B
Bedros Afeyan
Benjamin Armbruster
Steven Ashby
Cleve Ashcraft
Aboubacar Bagayogo
Zhaojun Bai
Richard Bartels
Michele Benzi
Dan Berkenstock
Jean-Paul Berrut
Petter Bjørstad
Åke Björck
John Bodley
Daniel Boley
John Bolstad
Erik Boman
Liliana Borcea
Andrew Bradley
Richard Brent
Alex Brik
David Brown
Roland Bulirsch
Aydin Buluc
Jim Bunch
Dave Burgess

C-D
Yongyan Cai
Dongwei Cao
Ryan Cassidy
Simla Ceyhan
Anwei Chai
Raymond Chan
Tony Chan
Jiuping Chen
Vani Cheruvu
Sou-Cheng Choi
Gaurav Chopra
Paul Concus
Paul Constantine
Bill Coughran
Luis Crivelli
Jason Cui
Jose Cuminato
Biswa Nath Datta
Carl de Boor
Bart De Moor
John de Pillis
Laurent Demanet
Jim Demmel
Aaron Diaz
Chris Ding
Jack Dongarra
Tony Drummond
Iain Duff
John Dunec

E-F
Boris Efros
Lars Eldén
Howard Elman
Anne Elster
Koffi Enakoutsa
Oliver Ernst
Fariba Fahroo
Ying Wai Fan
Charbel Farhat
Carl Farrington
Michael Ferris
Bernd Fischer
Michael Foster
Roland Freund
Benjamin Friedlander
Michael Friedlander
Lawrence Friedman

G-H
Martin Gander
Walter Gander
Walter Gautschi
Bill Gear
Margot Gerritsen
Eldar Giladi
Philip Gill
David Gleich
Peter Glynn
Gene Golub
Bill Gragg
Robert Granat
Joseph Grcar
John Greenstadt
Chen Greif
Roger Grimes
Marcus Grote
Ming Gu
Bertil Gustafsson
Eldad Haber
Michael Heath
John Hench
Greg Henry
Matthew Herman
Yoza Hida
Robert Higdon
Mark Hoemmen
Maya Honda
Roger Horn
Jeffery Housman
Thomas Hunt

I-K
Jesus Izaguirre
Kathy Jensen
Holly Jin
Pedher Johansson
Jin Hyuk Jung
Bo Kågström
William Kahan
Thomas Kailath
Craig Kapfer
Lars Karlsson
Linda Kaufman
Kaustuv
Herbert Keller
David Ketcheson
David Keyes
Hyunsoo Kim
Plamen Koev
Tammy Kolda
Tzanio Kolev
Roland Krause
Felix Kwok

L-M
James Lambers
Sotiria Lampoudi
Leif Christian Larsen
Alan Laub
Steven Lee
Tin Lee
Steven Leon
Ofer Levi
Dan Li
Xiaoye Li
Yaokang Li
Yung-Ta Li
Ben-Shan Liao
Lek-Heng Lim
Svante Littmarck
Yifan Liu
Henrik Loef
Rashant Loyalka
Frank Luk
Wenxiu Ma
Chris Maes
Michael Mahoney
Arian Maleki
Osni Marques
Nicola Mastronardi
Omkar Mate
Aaron Melman
Emre Mengi
Gerard Meurant
Rita Meyer-Spasche
Peyman Milanfar
Kourosh Modarresi
Cleve Moler
Kam Morrella
Jonathan Moussa
Lenore Mullin
Walter Murray

N-Q
James Nagy
Nir Naor
Stephen Nash
Marian Nemec
Esmond Ng
Nhat Nguyen
Nancy Nichols
Silvia Noschese
Bradley Null
Dianne O'Leary
Julia Olkin
Michael Overton
Peter Pacheco
Chris Paige
John Panzer
Dana Paquin
Eun Park
Haesun Park
Beresford Parlett
Yisrael Parmet
Emanuel Parzen
Lionello Pasquini
Anssi Pennanen
Victor Pereyra
Christian Perret
Wesley Petersen
Linda Petzold
Robert Plummer
James Pool

R-T
Arthur Rallu
John Reid
Tiago Requeijo
Jason Riedy
Maxine Rockoff
Nab Raj Roshyara
Ahmed Sameh
David Saunders
Michael Saunders
Thomas Savarino
Paul Saylor
Sam Schechter
Inga Schierle
Rob Schreiber
Craig Schroeder
Jennifer Scott
Radu Serban
Stefano S. Capizzano
Uday Shanbhag
Bern Shen
Tamar Shinar
Horst Simon
Vadim Sokolov
Knut Solna
Philip Sternberg
Pete Stewart
Gilbert Strang
Thomas Strohmer
Andrew Stuart
Zheng Su
Kunio Tanabe
Peter Tang
Stefan Tang
Stephanie Taylor
Dilys Thomas
Peter Tsai

U-Z
Paul Van Dooren
Sabine Van Huffel
Charles Van Loan
Raf Vandebril
James Varah
Steve Vavasis
Christof Voemel
Grace Wahba
Guanyuan Wang
Qiqi Wang
Andy Wathen
Khela Weiler
Dahlia Weiss
John Welsch
Nick West
Sarah Williams
Joab Winkler
Stephen Wright
Leslie Wu
Jianlin Xia
Ichitaro Yamazaki
Chao Yang
Benjamin Yolken
Jinyun Yuan
Michael Zerzan
Hongyuan Zha
Di Zhao
Zhengji Zhao
Mingqiang Zhu
Zhisu Zhu

FOR_SYTHEtation

Some of George and Alexandra Forsythe's Books

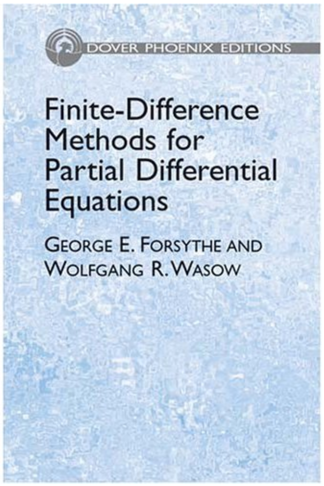

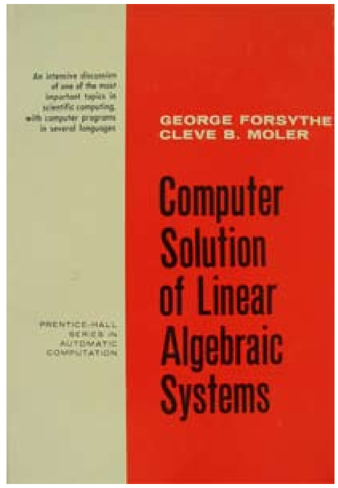

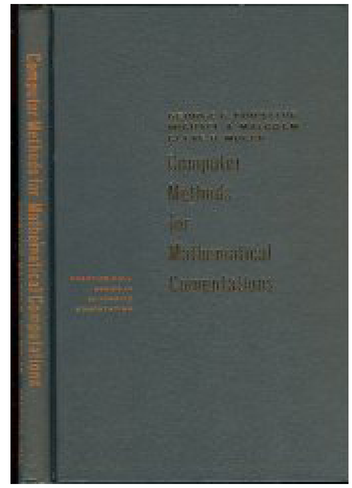

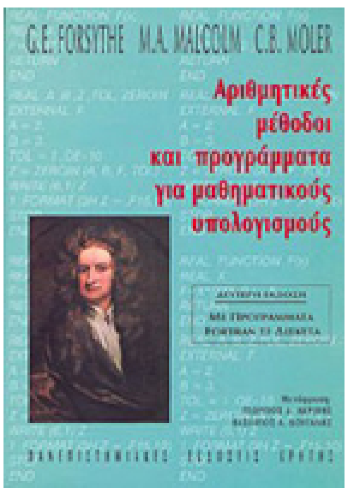

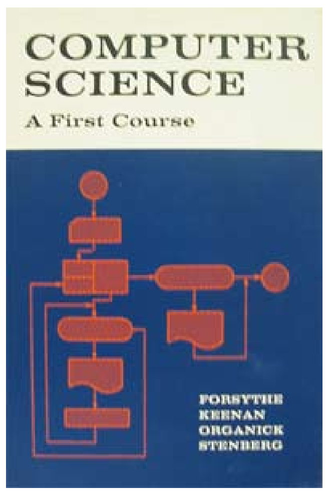

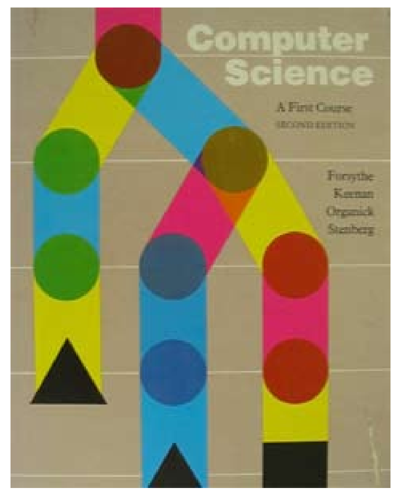

Some of Gene Golub's Books

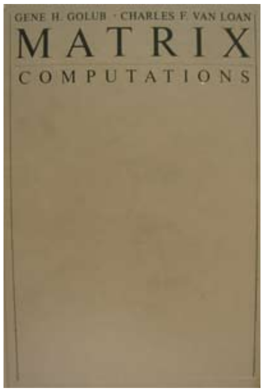

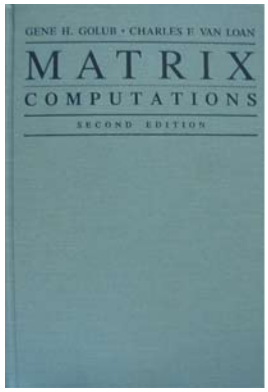

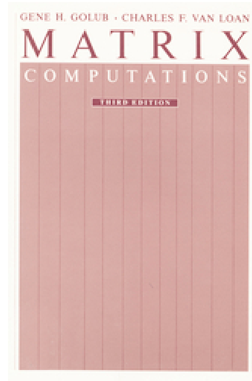

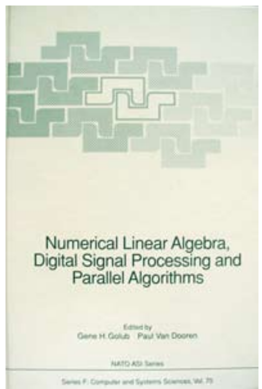

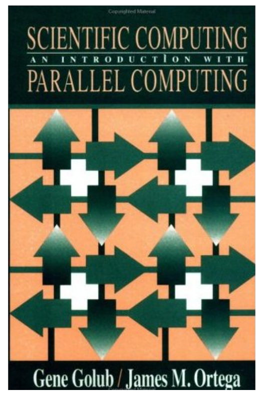

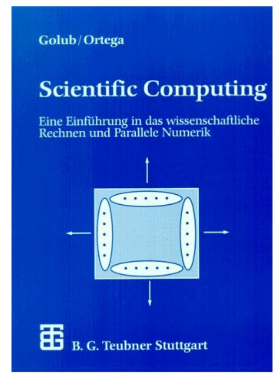

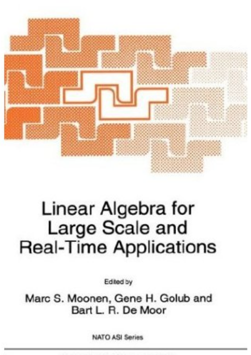

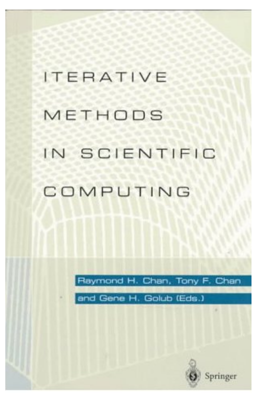

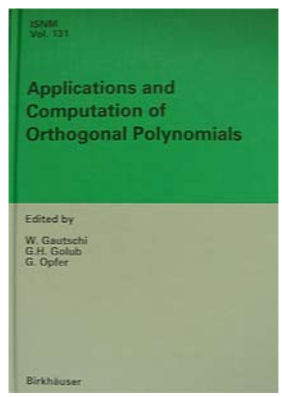

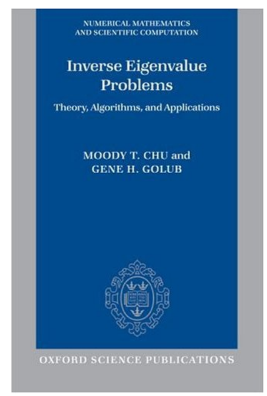

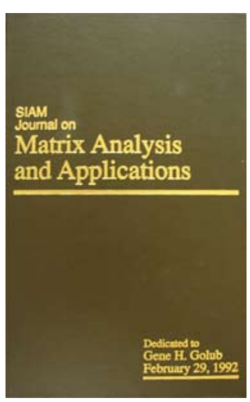

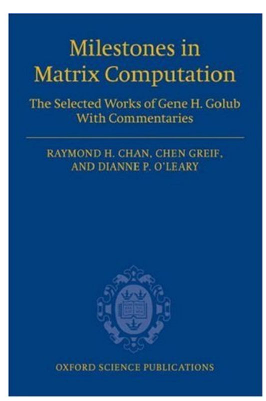

GENEalogy

George Forsythe's PhD Students

George Forsythe has $\geq$ 17 PhD students graduated from Stanford and $\geq$ 206 descendants all over the world.

Eldon Hansen
1960
$\geq$ 1 descendant

Beresford Parlett
1962
Univ of California, Berkeley
$\geq$ 44 descendants

James Ortega
1962
Univ of Virginia
$\geq$ 22 descendants

Betty Stone
1962
Lucent Technologies

Donald Fisher
1962

Ramon Moore
1963
Univ of Texas, El Paso
$\geq$ 7 descendants

Donald Grace
1964

Robert Causey
1964

Cleve Moler
1965
The MathWorks, Inc.
$\geq$ 41 descendants

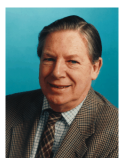

Roger Hockney
1966
Reading Univ
$\geq$ 3 descendants

William McKeeman
1966
The MathWorks Inc.
$\geq$ 1 descendant

James Varah
1967
Univ of British Columbia
$\geq$ 6 descendants

Paul Richman
1968
Bell Labs

Richard Brent
1971
Australian National Univ
$\geq$ 19 descendants

Alan George
1971
Univ of Waterloo
$\geq$ 17 descendants

David Stoutemyer
1972

Michael Malcolm
1973
Univ of Waterloo
$\geq$ 29 descendants

Some of Gene Golub's Postdocs

Iain Duff
Rutherford
Appleton
Laboratory, UK

Per Christian Hansen
Technical Univ of
Denmark
$\geq$ 16 descendants

Paul Van Dooren
Catholic University
of Louvain
Belgium

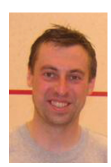

David Burgess
Yahoo!

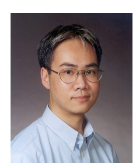

Wing Lok Justin Wan
Forsythe Fellow
Univ of Waterloo
Canada

Marko Huhtanen
Helsinki Univ
of Technology
Finland

Rasmus Larsen
L-3 Communications
Corp

Chen Greif
Univ of British
Columbia, Canada

Oren Livne
Univ of Utah
Salt Lake City

SungEun Jo
Samsung
Electro-Mechanics
South Korea

Gene's PhD Students

Gene Golub has ≥ 30 PhD students graduated from Stanford and ≥ 141 descendants all over the world.

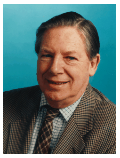
Roger Hockney
1966
Reading Univ
≥ 3 descendants
Deceased


Richard Bartels
1968
Univ of Waterloo
≥ 10 descendants


Michael Jenkins
1969
Queens Univ
≥ 2 descendants


Lyle Smith
1969
Lucent Technologies

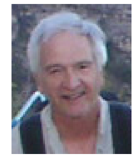
George Ramos
1970

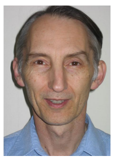
Richard Brent
1971
Australian National Univ
≥ 19 descendants


Michael Saunders
1972
Stanford Univ
≥ 5 descendants

John Palmer
1974


Richard Underwood
1975
Deceased


Dianne O'Leary
1976
Univ of Maryland
≥ 15 descendants


John Lewis
1976
Cray Inc.


Margaret Wright
1976
New York Univ


Michael Heath
1978
Univ of Illinois,
Urbana-Champaign
≥ 9 descendants


Franklin Luk
1978
Rensselaer Polytechnic Institute
≥ 23 descendants


Michael Overton
1979
New York Univ
≥ 4 descendants


Petter Bjørstad
1980
Univ of Bergen,
Norway
≥ 9 descendants


Daniel Boley
1981
Univ of Minnesota
≥ 4 descendants


Eric Grosse
1981
Google Inc.


Stephen Nash
1982
George Mason Univ
≥ 3 descendants


Mark Kent
1989


Ray Tuminaro
1989
Sandia National
Laboratories


Hongyuan Zha
1993
Georgia Tech
≥ 5 descendants


Oliver Ernst
1995
TU Bergakädemie
Freiberg

Xiaowei Zhan
1997


Tong Zhang
1999
Yahoo!


Nhat Nguyen
2000
Lockheed Martin Space
Systems Company


Urmi Holz
2002
National Security
Agency


James Lambers
2003
Stanford Univ


Yong Sun
2003
Accelet
Corporation


Sou-Cheng Choi
2006
Oracle USA, Inc.

Fondly Remembered NA Graduates and Faculty

Richard Underwood
1945–1985
with family Mary, Jennifer, and Christina

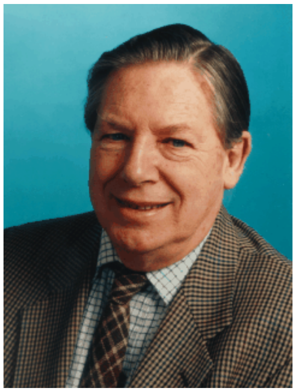

Roger Hockney
19xx–1999

George Forsythe
1917–1972

Jack Herriot
1916–2003

Joseph Oliger
1941–2005

James Wilkinson
1919–1986

Peter Henrici
1923–1987

Germund Dahlquist
1925–2005

Acknowledgments

Sponsors

Stanford 50 thanks the following institutional sponsors for their generous support:

- iCME: the Institute for Computational and Mathematical Engineering, Stanford University
- Faculty of Mathematics, University of Waterloo
- Stanford Computer Forum
- AFOSR: the Air Force Office of Scientific Research
- NSF: the National Science Foundation
- LLNL ISCR: the Institute for Scientific Computing Research, Computing Applications and Research Department, Lawrence Livermore National Laboratory
- COMSOL, Inc.
- The MathWorks, Inc.
- Yahoo! Inc.

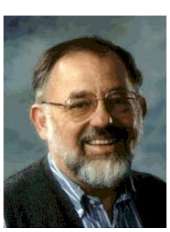

The mugs are a gift from Cleve Moler,
Chairman and Chief Scientist, The MathWorks, Inc.

They were designed by Jill Wright

Session Chairs


Richard Bartels
Univ of Waterloo


Victor Pereyra
Weidlinger
Associates


James Varah
Univ of
British Columbia


Haesun Park
Georgia Tech


Gerard Meurant
CEA, France


Walter Gander
ETH Zurich
Switzerland


Åke Björck
Linköping Univ
Sweden

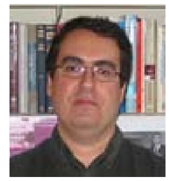
Michele Benzi
Emory Univ


Petter Bjørstad
Univ of Bergen
Norway

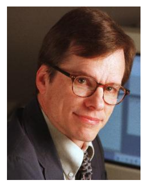
Michael Heath
Univ of Illinois,
Urbana-Champaign


Nancy Nichols
Univ of Reading
UK

Organizers

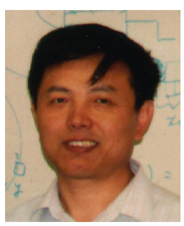
Ding-Zhu Du
Univ of Texas, Dallas

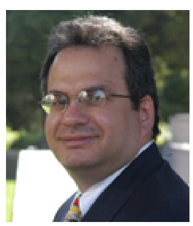
Charbel Farhat
Stanford Univ


Walter Murray
Stanford Univ


Michael Overton
New York Univ


Haesun Park
Georgia Tech

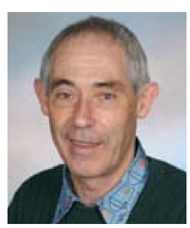
Michael Saunders
Stanford Univ


James Varah
Univ of
British Columbia

Events and meeting planning: Kam Morrella (Stanford Conference Services)
Schedule planning: Michael Overton
Webmaster: John Bodley
Program: Sou-Cheng Choi and Michael Saunders, using LaTeX 2_ε