Final Report for Period: 10/2006 - 09/2007 **Submitted on:** 01/15/2008
Principal Investigator: Park, Haesun . **Award ID:** 0633793
Organization: GA Tech Res Corp - GIT
Title:
Special Meeting: Workshop on Future Direction in Numerical Algorithms and Optimization

Project Participants

Senior Personnel

Name: Park, Haesun
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Golub, Gene
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Wu, Weili
Worked for more than 160 Hours: Yes
Contribution to Project:

Name: Du, Ding-Zhu
Worked for more than 160 Hours: Yes
Contribution to Project:
Dr. Du was one of the co-PI's but due to a COI (he was a program director at the NSF) with the NSF, Dr. Wu served as his substitute PI.

Post-doc

Graduate Student

Undergraduate Student

Technician, Programmer

Other Participant

Research Experience for Undergraduates

Organizational Partners

Other Collaborators or Contacts

The additional people who worked as the meeting organizers are:
Charbel Farhat,Mechanical Engineering,Stanford University
Walter Murray,Management Science and Engineering,Stanford University
Michael Overton,Courant Institute of Mathematical Sciences,New York University
Michael Saunders,Management Science and Engineering,Stanford University
James Varah,Computer Science,University of British Columbia

Activities and Findings

Research and Education Activities: (See PDF version submitted by PI at the end of the report)
The award was for supporting a workshop and the meeting program is attached.

Findings: (See PDF version submitted by PI at the end of the report)
The award was for supporting a workshop and the meeting report is attached.

Training and Development:

Outreach Activities:

Journal Publications

Books or Other One-time Publications

Web/Internet Site

Other Specific Products

Contributions

Categories for which nothing is reported:

Organizational Partners
Activities and Findings: Any Training and Development
Activities and Findings: Any Outreach Activities
Any Journal
Any Book
Any Web/Internet Site
Any Product
Any Contribution

NSF Workshop Report on FUTURE DIRECTIONS in NUMERICAL ALGORITHMS and OPTIMIZATION

Haesun Park[1], Ding-Zhu Du[2], and Gene Golub[3]

1 Executive Summary

The NSF workshop on future directions in numerical computing and optimization was held at Stanford University, Stanford, CA, on March 29-31, 2007. The stated goals of the workshop were to bring together researchers in numerical computing and optimization and to provide an opportunity for sharing ideas and information in order to identify the future research directions and initiatives by exploring critical and challenging research problems and topics, soliciting input from the community. The potential impact of such a workshop is that the result will provide value to the advancement in science within the communities, broader science and engineering disciplines, and eventually society by providing a better funding environment for developing the theory, models, tools, artifacts, and software.

The workshop on the future directions was strategically held in conjunction with the meeting on the history of numerical computing at Stanford University. The combined meeting provided a historically important forum for recognizing the accomplishments within the numerical computing and optimization communities over the past decades, and concurrently identify the current state-of-the-art and the challenging future directions. The meeting brought a large group of researchers of over two hundred fifty participants from the U.S., as well as from abroad, together in various stages of their careers from the key leaders in the field to graduate students.

A summary of the recommendations is as follows:

- Support for Fundamental Research in Numerical and Optimization Theory and Algorithm Development by the NSF: Continued support by the NSF, especially by the CISE for the research and education in fundamental theory and algorithm development in numerical computing and optimization is essential. There is a strong proven historical record that this fundamental research, from a computer science point of view, has contributed significantly to providing a basis for important application areas in science and engineering.

- New NSF Special Focus Program on Foundational Research in Massive Scale and High Dimensional Data Analysis for Knowledge Extraction: A special program by the NSF

[1]Division of Computational Science and Engineering, College of Computing, Georgia Institute of Technology, Atlanta, GA, 30332-0280, hpark@cc.gatech.edu

[2]Department of Computer Science and Engineering, University of Texas at Dallas, Richardson, TA 75083, dzdu@utdallas.edu

[3]Department of Computer Science, Stanford University, Stanford, CA 94305-9025, golub@stanford.edu

to support research in fundamental matrix decompositions and computations and multilinear algebra and computations is needed with regard to the dramatic development in massive scale, high dimensional, heterogeneous data analysis which has occurred in many ways that directly impact everyday life. This includes new algorithms and engines for efficiently and effectively searching databases and world wide web and network security, which can potentially have broader impacts in social computing as well.

- NSF Support on Applications of Numerical Methods and Optimization Algorithms to Areas of the Biological Sciences such as Computational biology/bioinformatics/systems biology in collaboration with the NIH: This type of program would have direct and enormous societal impact e.g., for early disease diagnosis and targeted treatment in the health sciences. Data analysis, machine learning, modeling and simulation, and the related numerical and optimization algorithms will be key in the development of tools for cancer diagnosis and treatment and tracking pandemic disease outbreaks. Also, discoveries from these areas will have concomitant impacts in many science and engineering application problems such as computational chemistry, quantum computing, and even systems performance analysis, to name a few.

- A Large Scale, Multi-divisional Support by the NSF in Numerical Modeling and Simulation for Inter-disciplinary Research in Science and Engineering: A large scale program by the NSF to promote research in multiscale problem formulation and solution, and model driven numerical computing and optimization will be essential as computation has been established as one of the three key paradigms for discovery in science and engineering.

- NSF Support for Research in Numerical Modeling and Simulation, Optimization, Analysis and Algorithm Tool Design: In addition to basic research support on theory and algorithm development, NSF program supporting practical computational tool design and development research will be essential for enabling the results developed in the numerical computing and optimization community to directly impact the application communities with the state-of-the-art techniques.

- NSF Program to Support for Algorithm, Tools, and Software Development Research for peta-scale and multicore architectures: Sustained NSF support for the development of algorithm and tools design for peta-scale and multicore architectures to drive the rapid discovery of improved techniques that will benefit not only scientific computing but will have broader impact in the general scientific community.

- NSF wide program to support Interdisciplinary and Cross-disciplinary Collaboration: Funding of collaborative research across disciplinary boundaries is highly recommended in order to bridge the gap between fundamental algorithm research and algorithm implementation that pushes the frontiers of research in science and engineering.

- An NSF Program to Support Education in Disciplines such as Computational Science and Engineering: One of the important topics in education in numerical computing

and optimization is how the community should position itself to take advantage of the trend to continue to make important contributions to interdisciplinary projects and use these experiences to identify new directions for fundamental optimization research. We need to give students a broad education in algorithmic techniques, so that they have a large "toolbox" to draw on when presented with an application. We need to give them experience in modeling a variety of applications in different contexts. They should gain experience in working directly with applications specialists and in learning how to communicate with them. Experienced researchers should also be ready to grasp the opportunities that interdisciplinary projects provide - opportunities to expand their own horizons, expand the reach of optimization in application areas, and open up new paradigms for fundamental research.

2 Introduction

Numerical Computing and Optimization involve the study of algorithms for continuous problems and have played fundamental roles in many Science and Engineering disciplines. Numerical computing, which originally sprang from mathematics, has been recognized as the area that spawned the new discipline of Computer Science. Currently, two generations after the Computer Science field has been established as an academic discipline, the role of numerical computing and optimization has become increasingly important. The numerous areas of research that have become modern disciplines demonstrate the crucial role that numerical algorithms and optimization play in a wide array of applications.

This community has been and will continue to play an important role in many interdisciplinary and cross-foundational programs supported by the National Science Foundation. Therefore, the result of this workshop is expected to provide an important guideline for identifying the strategic directions in many programs at the National Science Foundation that involve the areas of Numerical Computing and Optimization including the Division of Computing and Communications Foundation, the Computer and Information Science and Engineering Directorate.

From the workshop, it is clear that numerical computing and optimization are vital to progress in the application areas mentioned above and beyond. However, there are still numerous and important fundamental research problems to be tackled in these areas as well. Also very important is that the researchers in these areas proactively reach out to application communities to effectively contribute to important science and engineering problems. An immediate example of such efforts include massive scale high dimension data analysis for which many traditional algorithms have been discovered not to model the problems well. In addition, the recent paradigm change in high performance computing such as hyperthreading and multicore has resulted in ever higher demand for redesign of the fundamental scientific software for effective utilization of the novel architectural features. As a mature area now, the challenge is to produce results that are relevant to real applications and make the results readily accessible to researchers in other fields. These examples and more illustrate the numerous and fruitful collaborations that should exist between scientific computing researchers

and those in application areas.

3 Current Status of Numerical Computing and Optimizations

As agreed upon universally by now, computation has been established as one of the three paradigms for conducting research in science in addition to theory development and experimentation. Computational models are now playing an indispensable role in virtually all fields of science and engineering to better understand systems and phenomena at many different scales. Computation is essential in critical applications with great societal impact. This has been especially evident recently as computation has played a key role in addressing many security issues, e.g. face recognition, fingerprint classification, early warning systems, to name a few. The economic benefits and protection of human life of these new computational technologies are incalculable. This importance will only increase in the foreseeable future. Numerical computing and optimization have played significant roles in advancing many areas of science and engineering as the key area of computation dealing with continuous mathematical problems in contrast to discrete problems. The researchers in numerical computing and optimization have been engaged in an important mission leveraging the recent advancements in physical and biological science and engineering applications to focus and lead the development of modern research and education by developing effective and efficient algorithms and mathematical tools which would impact various applied disciplines. They have also been engaged in designing, improving, and analyzing fundamental algorithms that are powerful and indispensable across various application domains.

In the last few decades, there has been a tremendous amount of research activity and progress in large-scale numerical computing and optimization, especially in the area of large-scale matrix computations. As a result, we now have powerful algorithms and software for many basic large-scale computational tasks, such as the solution of large sparse systems of linear equations, the computation of selected eigenvalues of large sparse matrices, the solution of large-scale least-squares problems, or dimension reduction of large-scale linear dynamical systems. However, these basic computational tasks are hardly ever of interest by themselves. Instead, they usually arise as subproblems in numerical simulations of increasingly more complex systems and processes. For example, much of the recent work on dimension reduction of large-scale linear dynamical systems was driven by the need for such reduction techniques for very large-scale interconnect models of state-of-the-art VLSI circuits. However, the reduction of these interconnect models is only one of many challenging computational tasks that are required to simulate the behavior of a complete VLSI circuit.

In addition, the numerical computing community has made significant contributions throughout all stages of architectural development in creating reliable and fast general purpose computational software and tools. After decades of exponential increases, chip clock speeds have reached a plateau. It is difficult to estimate the magnitude of the discontinuity that the high performance computing (HPC) community is about to experience because of

the emergence of the next generation of multi-core and heterogeneous processor designs. For at least two decades, HPC programmers have taken it for granted that each successive generation of microprocessors would, either immediately or after minor adjustments, make their old software run substantially faster. But factors such as intractable physical barriers including too much heat, too much power consumption, and too much leaking voltage, to further increases in clock speeds, and physical limits on the number and bandwidth of pins on a single chip are converging to bring this free ride to an end. Despite the rapidly approaching obsolescence of familiar programming paradigms, there is currently no well understood alternative in whose viability the community can be confident. Though it is clear that the impact of the multi-core revolution will be ubiquitous, we believe that, in developing a programming model for this radically different environment, there are clear advantages to focusing on Linear Algebra (LA) in general and Dense Linear Algebra (DLA) in particular. For one thing, DLA libraries are critically important to Computational Science across an enormous spectrum of disciplines and applications, so a programming framework of the type based on DLA will certainly be indispensable and needs to be achieved as quickly as possible. But DLA also has strategic advantages as a research vehicle, because the methods and algorithms that underlie it have been so thoroughly studied and are so well understood. This background understanding will allow us to devise techniques that maximally exploit the resources of the microprocessor platforms under study.

An additional factor that contributes significantly to the need of new numerical algorithms and optimization techniques is the over-abundance of data. Due to today's exponential growth of the internet, world-wide-web, sensor networks, widespread availability of low-cost storage and computing power, a data analysis system is expected to effectively handle a massive amount of data in many different formats from many different sources, and users demand more efficient techniques to obtain useful information from the flood of data. These phenomena have been observed across numerous important areas such as biomedicine, surveillance, digital health, astronomy, business transactions, and many others. Recent research activities show that even system performance analysis techniques and network security analysis can be conducted using sophisticated data analysis techniques, and more research areas are demanding increased utilization of data mining and analysis, information retrieval, machine learning, and visualization to extract meaningful information from these massive data sets.

The databases storing the data continue to burgeon as new data is constantly being acquired at a faster pace, and sophisticated algorithms and software research will provide a means of extracting knowledge from this data that can be used in a highly charged competitive environment. There is a great need to serve users as efficiently and thoroughly as possible; knowledge is the means and data analysis is the key to achieve this. Efficient techniques for dimension reduction, data reduction, classification, clustering, and information retrieval have been developed, which heavily rely on numerical linear algebra and optimization techniques. The key examples include principal component analysis, linear and nonlinear discriminant analysis, support vector machines, k-means clustering, and latent semantic indexing. It is expected that numerical linear algebra and optimization will continue to play a

pivotal role in providing the mathematical modeling, algorithm development, and computational engine necessary to realize these goals of information extraction. The major research efforts should consist of both theoretical and experimental investigation since it is crucial to study these aspects in parallel: theoretical results provide a framework for developing algorithms and predicting their behavior, while through experimentation, theoretical results come to have practical value in information mining. Numerical computing and optimization will clearly play fundamental roles in new research advances required before data mining and analysis techniques can achieve their full potential.

4 Future Directions in Numerical Computing and Optimization and Applications in Science and Engineering

In this section, we present the future directions in numerical computing and optimization and the strategic application areas in science and engineering which were identified at the workshop. These include massive scale data analysis, computational systems biology, and numerical modeling and simulation that appear in numerous application areas such as computational fluid dynamics, computational electromagnetics, and demand new methods in PDE, nonlinear systems solvers, and nonlinear programming. In addition, with the advent of PetaScale computing there is a motivation to design new computational algorithms that exploit new hardware architectures. In general, it is clear that the exciting position of importance of the field of numerical computing and optimization in science and engineering areas has exploded. Never has the field been so important to so many different disciplines and it is expected that numerical computing will play a central role into the foreseeable future. Insights from computations will inform, influence, and direct novel science and engineering discovery. In order to contribute at the frontier of new directions, the community needs to be proactively involved in multi- and cross-disciplinary research activities, and produce reliable and practical computational tools. At the same time, fundamental research must continue since important new fundamental problems will be identified through these cross-disciplinary research and educational activities. Some of the important open research problems regarding fundamental theory and algorithm development in matrix computations and algorithms include exploiting matrix structure (displacement structure, sparsity, multi-level hierarchy) to save computational complexity and storage, developing efficient algorithms for solving conic convex optimization problems with general cones, developing algorithms to exploit new architectures, such as GPUs and massively parallel machines (PetaScale and Multicore), developing algorithms that are efficient for multiscale problem formulations, answering basic open questions about well-known algorithms, e.g., answering practical convergence issues for Krylov subspace methods.

With research in numerical computing and optimization taking place more and more in an interdisciplinary context, researchers in a vast range of application areas are turning to

numerical computing and optimization techniques for help in solving problems that lie on the critical path of their research agenda. Often, the resulting collaborative projects produce research of interest to both the numerical computing and optimization AND application communities. Such efforts are adding a great deal of vitality to modern numerical computing and optimization research.

The interesting applications for numerical computing and optimization researchers are, of course, those that require more than a straightforward application of an existing algorithm or code. Even when the application problem fits a known paradigm for which code exists, the process of modeling - expressing the application as an optimization problem - may be nontrivial. There are often a variety of ways to model the application, but some models are much easier to solve than others. (Such is frequently the case with models involving integer variables or discrete decisions, for instance.) Experience with formulation is needed, as well as interaction between numerical computing and optimization researchers and applications researchers to identify a level of model "resolution" that yields valuable insight yet can be solved in a reasonable amount of computer time.

In other applications, the problem may simply be too large to be solved with a general-purpose approach, or it may contain features (e.g. a type of uncertainty in the data) that have not yet been addressed in the numerical computing and optimization literature. In these cases, solution techniques need to be assembled or "engineered" from several algorithmic tools from optimization and numerical computing. The context in which the problem must be solved is also important in determining the appropriate approach. For example, it may only be necessary to find an approximate solution, or to solve many similar problems.

The prospects are excellent for continued expansion in this kind of interdisciplinary research involving numerical computing and optimization. Such developments are extremely healthy for the field. Exposure to a wider range of applications will help us to identify the key issues and to find new paradigms to be investigated at a more abstract (algorithmic and theoretical) level.

4.1 Massive Scale Data Analysis

The future of numerical analysis is not just conducting research in order to handle ever-larger problems but also includes a new focus on scientific research on data mining and related mathematical methods such as multilinear algebra and tensor decompositions.

More and more, practitioners in science and engineering are struggling under massive amounts of data from a combination of high through-put experimental devices and large-scale high-resolution simulations. In order to continue to advance scientific understanding, new numerical methods for analyzing large-scale (peta- and exa-scale) data are required. Currently, the methods that scientists and engineers are using are inaccurate due to their reliance on home-grown heuristics and do not scale even to current problem sizes. However, major advancements are possible by using novel advanced numerical algorithms. To the best of our knowledge, the funding for scientific data mining has not kept pace with the need. The only exception is biology, but the types of data mining problems in biology are more

akin to classical data mining problems in text retrieval, clustering, and so on.

This leads naturally to the next major theme: advanced numerical methods for data mining. Google's PageRank algorithm made linear algebra cool again, especially to students. In fact, linear algebra techniques play a major role in data mining and the need for matrix decompositions will continue to grow at a pace faster than ever due to the exponential growth in data size. And there will be a focus on inventive matrix decompositions such as nonnegative factorizations, sparse factorizations, independent components analysis, sample-based factorizations (like CUR), cluster-revealing factorization, and so on.

Finally, neither the scientific nor data mining world is necessarily two-dimensional and linear. The field of tensor decompositions for n-way arrays is similar to the field of linear algebra circa 1960. The next 30-40 years will see major advancements theoretically and computationally. These methods are proving useful in image analysis and numerous applications in data mining, particularly when looking at temporal data. Tensor methods are limited to multi-linear data, but we expect that more fundamental work in nonlinear multidimensional models will follow advancements in this area.

4.2 Computational Systems Biology

In spite of close to a century of dynamic modeling in a few areas of biology, notably genetics, epidemiology and ecology, a systematic approach to analyzing complexity in biological networks has only recently become feasible. Previously, progress was limited by the lack of suitable measurements and by conceptual and computational problems in simulating complex mathematical models. Advances in molecular biology over the past decade have made it possible to probe experimentally the causal relationships between processes initiated by individual molecules within a cell and their macroscopic phenotypic effects on cells and organisms. These studies provide increasingly detailed insights into the underlying networks, circuits, and pathways responsible for the basic functionality and robustness of biological systems, and create new and exciting opportunities for the development of quantitative and predictive modeling and simulation tools.

What are the challenges? Model development typically begins with potentially large amounts of noisy data. Models are highly multiscale: actions at the level of genes affect outcomes at the level of organisms and ecosystems; biochemical networks include both fast and slow reactions; the behaviors and actions of individual organisms must play a role in regulating populations and influencing the fluxes of elemental matter in ecosystems. When populations of chemical species or organisms are small, stochasticity can play a major role. The multiscale problems are often heterogeneous: individual-based models or discrete stochastic models at small time or spatial scales need to interact seamlessly with deterministic, differential equation-based models at larger scales. The stochastic simulations can take enormous amounts of computer time, particularly when spatial variation must be taken into account. There is a great need to develop both the mathematics and the computational software to accelerate these computations while maintaining confidence in their accuracy.

One of the hallmarks of systems biology is the interplay of experimentation and modeling.

One can begin with simple mathematical models and use these to guide experiments that refine the critical components of the model. Tools and methodologies will be required that go well beyond simulation, addressing the need for model development, model invalidation and experimental design in an environment where large amounts of uncertain data abound and it is often difficult or impossible to experimentally obtain the data one would ideally want to construct a model.

4.3 Numerical Modeling and Simulation

In order to facilitate the modeling and simulation of large-scale systems and processes, future research in large-scale numerical computing and optimization need to move beyond the solution of isolated basic computational tasks. In particular, techniques are needed that allow the use of existing algorithms for the solution of basic computational tasks as building blocks for the simulation of complete large-scale systems or processes. Issues that arise and need to be resolved in this context include the control of the accuracy to which the basic subproblems need to be solved, understanding the crucial properties of the algorithms for solving the basic subproblems that guarantee stability and sufficient accuracy of the overall simulation, and the ability to handle uncertainty in the data of the basic subproblems. Some subareas of numerical computing and optimization which have significant components in advancement in modeling and simulation and their future directions are identified in this section.

4.3.1 Fast Nonlinear System Solvers

In many computational science and engineering areas such as computational mechanics (fluids, structures, and fluid-structure interaction), computational electromagnetics, and computational nanoelectronics, just to cite a few examples, one is faced with the computational challenge of solving numerous very large nonlinear algebraic equations. Often, these nonlinear systems are solved using variants of Newtons method. Such an approach, in turn, requires the solution of very large sparse linear systems. In spite of great progress in the development of direct sparse solvers, current direct solvers, while reliable, cannot effectively handle very large systems. Moreover, current Krylov subspace methods with black-box preconditioners proved to be unreliable solvers. By black-box preconditioners we refer to the commonly used approximate factorization methods or approximate inverse preconditioners, i.e. preconditioners that consider only the underlying matrix of coefficients. For very large systems in which the matrix of coefficients is not generated explicitly, the reliability of available techniques is much worse.

The first challenge is regarding the development of reliable hybrid solvers within an adaptive library setting, which at one extreme can be used as pure direct solvers and at the other extreme as preconditioned iterative schemes. The design philosophy of such algorithms, however, will be different from the currently available solvers direct or iterative.

In addition to the above algorithmic factors, there are emerging architectural features that require novel approaches if high performance is to be realized. Considerations of power and

speed have motivated virtually every major microprocessor vendor to develop chip multiprocessor platforms (hyperthreading, multicore, Cell, etc). These platforms offer the potential of significantly higher computational speeds that do not correspondingly scale for the memory system, the Memory Wall phenomenon. For this reason, we often observe that algorithms with higher operation counts but lower memory reference counts perform much better in terms of wall-clock time.

An additional challenge is that the design and development of a new family of solvers that are optimized for memory access, as opposed to FLOP counts, which has been the traditional figure of merit.

This, in turn, leads to the next challenge that involves the design of an adaptive library of hybrid sparse system solvers that take into account both the characteristics of the sparse linear system at hand as well as the underlying architecture of the parallel computing platform to simultaneously achieve reliability, high performance, and parallel scalability.

4.3.2 Fast Linear System Solvers

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 generalizes deterministic finite element discretization strategies to handle problems posed with uncertainty. This methodology uses a discrete weak form in which the stochastic aspect of the problem is represented by a finite term series approximation, which is discretized in a manner analogous to finite elements. An advantage of this approach, in comparison with Monte Carlo methods, is that once the algebraic systems associated with such a discretization are solved, it is straightforward to obtain statistical properties of solutions such as moments and probability distributions. To achieve this advantage, it is critical to develop fast solvers for the associated systems, which are much larger than the problems that arise from deterministic models.

Many models of incompressible flow in computational fluid dynamics also give rise to algebraic systems of equations having the structure of saddle-point problems. These include the Navier-Stokes and Stokes equations, which model pure fluid flow, as well as more complex systems arising from linear stability analysis of flows and in models of flow coupled with thermal or chemical effects. The development of efficient iterative solution algorithms for these problems depends on the construction of effective preconditioners for them. It has been found that this can be done by focusing on the Schur complement associated with the saddle-point operators. Preconditioning methods are needed for a variety of models arising in CFD that derive from this observation. They have a common block-triangular structure in which the component blocks are tailored to handle specific problem classes.

4.3.3 Numerical PDE

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 solutions of a system of non-linear equations at each time step (although typically one solution of a linear system is sufficient to get the required accuracy in 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 so 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. Recently a class of methods called telescopic projective methods has been developed. Methods of this sort have the potential to be adaptive not only to control the error but to modify the method coefficients to adjust to the eigenvalues of the problem. So one direction that is potentially important is the development of automatic explicit integration methods for massively parallel systems. A further advantage of these methods is that they can be "wrapped around" single step legacy codes or microscopic simulators. This ties into an expanding area of interest in the analysis of systems via numerical tools. The determination of the macroscopic, low-dimensional behavior of systems that can only be described via very high-dimensional microscopic models because nobody has been able to analytically determine the equations for the closure of the model.

4.3.4 Multivariate Interpolation

Ever since its development in the 17th century, univariate polynomial interpolation has been one of the fundamental processes of Numerical Analysis (aka Scientific Computing). Its multivariate counterpart has not yet achieved the same standing, for various reasons. A chief reason is that, in $d > 1$ variables, there is no clear-cut choice of a polynomial subspace from which to interpolate given values at a given set of sites. Such a choice is made particularly difficult by the fact that, for any $n > 1$, there is no d-variable polynomial subspace of dimension n containing unique interpolants to values given at every n-set of data sites. It is an open and seemingly very difficult question to determine what is the smallest dimension, as a function of d and n, of a polynomial subspace that contains an interpolant to arbitrary values at every n-set of sites. In the late 80'ies, de Boor and Ron proposed a data-site dependent choice called least interpolation with many very nice properties. However, as a function of the data sites, this choice is only piecewise continuous. Hence one rather elusive goal is a continuous choice of an otherwise still reasonable interpolation space.

Assuming one has settled on a specific scheme for choosing the interpolation space for any given set of sites, there is the follow-up question of what happens to the resulting

'Lagrange interpolants' as interpolation sites coalesce. Here, too, the multivariate situation is much richer than the univariate one, with even some basic questions not yet understood. The choice of interpolation space also materially influences the shape of formulas for the interpolation error. However, except for some very specific and nice situations, we have as yet no handle on good error formulas, not even what form one might hope for them.

Most of the existing literature on multivariate polynomial interpolation does not deal with any of the above problems but concentrates on interpolation as a means of constructing approximations from the space of all polynomials of degree $< k$ for some k. (For this setting, Sauer and Xu have developed an error formula, but one might hope for something simpler.) Starting with the interpolation space severely restricts the number of interpolation sites, and somewhat restricts their location. More positively, it raises the question of how to choose 'good' interpolation sites so that the resulting interpolants are 'near best' with respect to the function norm of interest. Only very recently and essentially only in two variables so far, have there been any results in this direction, namely the 'Padua points'.

4.3.5 Nonlinear Programming in Modeling

Optimization models have traditionally been of the form to minimize or maximize a nonlinear function subject to nonlinear inequality and equality constraints on the variables. Specialized codes have allowed certain problem structures to be exploited algorithmically, for example simple bounds on variables. In a series of papers, Rockafellar and colleagues have introduced the notion of extended nonlinear programming, where the(primal) problem has the form: $\min_{x\in X} f_0(x)+\theta(f_1(x),\ldots,f_m(x))$ The function θ can be thought of as a generalized penalty function. However, for the following form it lends itself to a nice duality structure $\theta(u) = \sup_{y\in Y}\{y'u - k(y)\}$ This form can take on the value of ∞ and may well be nonsmooth. However it is guaranteed to be convex (proper and lower semicontinuous). An elegant duality framework can be derived from the Lagrangian: $\mathcal{L}(x,y) = f_0(x) + \sum_{i=1}^m y_i f_i(x) - k(y)$ where $x \in X, y \in Y$. Research is needed in making this problem format available to users within a modern modeling language. As special cases, it needs to be shown how to model piecewise linear penalties, least squares and L_1 approximation problems, as well as the notion of soft and hard constraints. Modelers need to be allowed to utilize cone constraints and pass on the underlying geometric structure to solvers. Particular examples show enormous promise both from a modeling and solution perspective.

4.3.6 Molecular Simulation

In the traditional paradigm of numerical analysis, the goal of a mathematical investigation of an algorithm is to show that as some parameter (step size, iteration number, order of approximation) goes to a limit, the numerical solution converges to the true solution to a problem. Frequently, this involves first proving that the original problem is well-posed. The algorithm is then analyzed as a perturbation to the problem. Even when a problem in its full complexity is not analyzable, usually a special case can be rigorously understood.

The problems attracting people's interest in molecular simulation have the feature that there are no analytically tractable versions of the problems available. Consider self-diffusion: the diffusion of one particle of a liquid among a bath of identical particles. For a given temperature, density, and particle interaction force, the rate at which a single particle diffuses through its neighbors is called the diffusion coefficient. It has long been well-understood how to estimate diffusion coefficients through numerical simulation, and this estimate suffices for practical purposes. For a mathematician trying to understand if these numerical simulations are reliable, a natural first move is to look at a case where we know that the diffusion coefficient is analytically available. Unfortunately, there is no such case. Even the fact that there is a finite rate of diffusion is not rigorously established for any models of this type.

A rigorous mathematical theory of self-diffusion, or of very many other physical problems, is not likely to emerge any time soon. So it is necessary for a mathematician interested in the numerics of these problems to broaden their view of rigor, and consider embracing some of the techniques and viewpoints of theoretical physics. Likewise, many physicists are realizing that algorithms are important enough to warrant study in their own right. They are coming to appreciate that the mathematical perspective brings a concreteness and clarity to scientific computing. It is predicted that the boundary between mathematicians and physicists working on scientific computing in this area will fade, to the benefit of all.

4.3.7 Reservoir Simulation, Coastal Ocean Modeling, and Design Optimization

The common challenges in the areas of continuous modeling and simulation, and design optimization of physical phenomena are due to multiple aspects including multi-scale character flows, strong nonlinearities, large uncertainties in parameters, and, because of the strong nonlinearities, strong sensitivities to parameters, in addition to limited data for validation and partially unknown physics. On top of these challenges, when it comes to computational tools, more difficulties are involved due to the necessary porting of codes to multi-core processor designs and there is a dire need for a very strong push for robust uncertainty assessments. These issues present huge challenges to the numerical computation and optimization community; the benefits of solving them are enormous, both economically and socially.

Some common concerns are: too strong a reliance of engineers (and scientists) on computer codes; codes are not always predictive because of nonlinearities and uncertainties; users of codes, and designers of codes, are not always aware of limitations of the codes, nor of uncertainty assessment strategies, leading to wrong predictions. This is quite common in, for example, large scale computations in atmospheric and ocean modeling and climate modeling. High density calculations (high grid density, multi-scale) often exposes weaknesses of empirical models for physics. There is a strong need for more lab experimentation for validation of computational models.

Unfortunately, there is a tendency for researchers to just believe that the more grid points are used, the better the model will predict the physics. However, predicting the physics cannot be done easily due to the physical assumptions of the models producing biased results. For example: in thermal simulation of enhanced oil recovery processes, it is generally assumed

that the kinetics at the reservoir scale satisfy Arrhenius equations. To find the reaction rates and activation energies, the models are calibrated using lab experimentation. However, it can easily be shown that at other scales, the results are unreliable. Arguably, the same is true for climate models.

4.4 PetaScale and Multi-Core Programming

Among the various factors that are driving the momentous changes now occurring in the design of microprocessors and high end systems, three stand out as especially notable: 1) the number of transistors on the chip will continue to double roughly every 18 months, but the speed of processor clocks will not continue to increase; 2) the number and bandwidth of pins on CPUs are reaching their limits and 3) there will be a strong drift toward hybrid systems for petascale (and larger) systems. Each of these factors has a somewhat different effect on the design space for future programming:
1) More transistors and slower clock speeds means multi-core designs and more parallelism required. The modus operandi of traditional processor design will increase the transistor density, speed up the clock rate, raise the voltage, which has now been blocked by a stubborn set of physical barriers, i.e., too much heat produced, too much power consumed, too much voltage leaked. Multi-core designs are a natural response to this situation. By putting multiple processor cores on a single die, architects can continue to increase the number of gates on the chip without increasing the power densities. Despite obvious similarities, multi-core processors are not equivalent to multiple-CPUs or to SMPs. Multiple cores on the same chip can share various caches (including TLB!) and they certainly share the bus. Extracting performance from this configuration of resources means that programmers must exploit increased thread-level parallelism (TLP) and efficient mechanisms for inter-processor communication and synchronization to manage resources effectively. The complexity of parallel processing will no longer be hidden in hardware by a combination of increased instruction level parallelism (ILP) and deep-and-narrow pipeline techniques, as it was with superscalar designs. It will have to be addressed in software.
2) Thicker memory walls means that communication efficiency will be even more essential. The processor to memory performance gap, which is already approaching a thousand cycles, is expected to grow, by 503) Limitations of commodity processors will further increase heterogeneity and system complexity: Experience has shown that tera- and petascale systems must use commodity off-the-shelf (COTS) processors as their foundation. Unfortunately, the trade-offs that are being and will continue to be made in the architecture of these general purpose multi-core processors are unlikely to deliver the capabilities that leading edge research applications require, even if the software is suitably modified. Consequently, in addition to all the different kinds of multithreading that multi-core systems may utilize - at the core-level, socket-level, board-level, and distributed - they are also likely to incorporate some constellation of special purpose processing elements such as hardware accelerators and GPUs.

For numerical software development, the essence of the problem is the dramatic increase

in complexity that software developers will have to confront. Dual-core machines are already common, and the number of cores is expected to roughly double with each processor generation. But contrary to the assumptions of the old model, programmers will not be able to consider these cores independently because they share on-chip resources in ways that separate processors do not. The proliferation of widely divergent design ideas shows that the question of how to best combine all these new resources and components is largely unsettled. When combined, these changes produce a picture of a future in which programmers must overcome software design problems that are vastly more complex and challenging than in the past in order to take advantage of the much higher degrees of concurrency and greater computing power that new architectures will offer.

We believe that these major trends will define, in large part at least, the design space for scientific software in the coming decade. But while it may be important for planning purposes to describe them in the abstract, to appreciate what they mean in practice, and therefore what their strategic significance may be for the development of new programming models, one has to look at how their effects play out in concrete cases. Early experience with these new architectures has taught us how they render traditional cornerstone numerical libraries obsolete, and how innovative techniques can exploit their parallelism and heterogeneity to address these problems.

One good way to appreciate the impact and significance of the multi-core revolution is to examine its effect on software packages that are widely familiar. The LAPACK/ScaLAPACK libraries fit that description. These libraries, which embody much of the research in the adaptation of block partitioned algorithms to parallel linear algebra software design, have served the HPC and Computational Science community remarkably well for twenty years. Both LAPACK and ScaLAPACK apply the idea of blocking in a consistent way to a wide range of algorithms in linear algebra (LA), including linear systems, least square problems, singular value decomposition, eigenvalue decomposition, etc., for problems with dense and banded coefficient matrices. ScaLAPACK also addresses the much harder problem of implementing these routines on top of distributed memory architectures. Yet it manages to keep close correspondence to LAPACK in the way the code is structured or organized. The design of these packages has had a major impact on how mathematical software has been written and used successfully during that time. Yet when you look at how these foundational libraries can be expected to fair on large-scale multi-core systems, it becomes clear that we are on the verge of a transformation in software design at least as potent as the change engendered a decade ago by message passing architectures, when the community had to rethink and rewrite many of its algorithms, libraries, and applications.

5 Recommendations

The recommendations to the NSF are itemized Section 1 Executive Summary.

6 List of Participants, Organizers, and Presenters

6.1 List of Participants

For the list of over two hundred fifty participants, see the attached pdf file which contains complete meeting program.

6.2 Organizers

- Ding-Zhu Du, University of Texas, Dallas, TX
- Charbel Farhat, Stanford University, Stanford, CA
- Walter Murray, Stanford University, Stanford, CA
- Michael Overton, New York University, New York, NY
- Haesun Park, Georgia Institute of Technology, GA
- Michael Saunders, Stanford University, CA
- James Varah, University of British Columbia, Vancouver, BC, Canada

6.3 Presenters and Panelists

The workshop occurred as a part of a combined meeting that covers historical, current state-of-the-art, and future directions in computational mathematics and numerical computing. It celebrated the milestones in the area of computational mathematics and numerical computing by recognizing the 50th anniversary of the arrival of George Forsythe at Stanford University which ushered in a new era of computational mathematics both at Stanford and elsewhere and explored the future of this important field. As a part of the meeting, the 75th birthday of Gene Golub was also celebrated who can rightfully claim to have carried the mantle after Forsythe's death in 1972 as a universally recognized ambassador at large for scientific computing. The speakers were invited from both the U.S. and from other countries as well. Here, we list only those US based invited speakers who were specifically asked to address the future directions in numerical computing and optimization in their research areas.

- Zhaojun Bai, University of California, Davis, Panelist
- Liliana Borcea, Rice University, Speaker
- Carl de Boor, University of Wisconsin, Speaker
- James Demmel, University of California, Berkeley, Speaker
- Jack Dongarra, University of Tennessee, Speaker

- Howard Elman, University of Maryland, Speaker
- Michael Ferris, University of Wisconsin, Speaker
- Roland Freund, University of California, Davis, Speaker
- William Gear, Princeton University, Speaker
- Margot Gerritsen, Stanford University, Speaker
- Philip Gill, University of California, San Diego, Speaker
- William Kahan, University of California, Berkeley, Speaker
- Tamara Kolda, Sandia National Laboratories, Panelist
- Dianne O'Leary, University of Maryland, Speaker
- Linda Petzold, University of California, Santa Barbara, Speaker
- Ahmed Sameh, Purdue University, Speaker
- G. W. Stewart, University of Maryland, Speaker
- Grade Wahba, University of Wisconsin, Speaker
- Stephen Wright, University of Wisconsin, Speaker

In addition, the following people from the National Science Foundation attended the meeting. A co-authored talk by Michael Foster, Lenore Mullin and Eun Park, was presented by Lenore Mullin.

- Michael Foster, CISE/National Science Foundation, Sponsor
- Lenore Mullin, CISE/National Science Foundation, Sponsor/Speaker
- Eun Park, CISE/National Science Foundation, Sponsor
- Tony Chan, MPS/National Science Foundation, Attendee

Acknowledgement

We would like to thank the following people for their contributions to this report: Carl de Boor, Jack Dongarra, Howard Elman, Michael Ferris, Roland Freund, William Gear, Margot Gerritsen, Tamara Kolda, Dianne O'Leary, Linda Petzold, Ahmed Sameh, Paul Tupper, and Stephen Wright.

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! **3**

Conference Schedule **4**

Thursday, March 29, 2007 4

Friday, March 30, 2007 6

Saturday, March 31, 2007 8

Abstracts **9**

Posters **20**

Graduate Student Posters 20

Junior Scientist Posters 21

Judges 21

Participants **22**

FOR_SYTHEtation **23**

Some of George and Alexandra Forsythe's Books 23

Some of Gene Golub's Books 24

GENEalogy **25**

George Forsythe's PhD Students 25

Some of Gene Golub's Postdocs 25

Gene's PhD Students 26

Fondly Remembered NA Graduates and Faculty **27**

Acknowledgments **28**

Sponsors 28

Session Chairs 29

Organizers 29

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

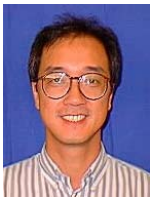

I review how a primal-dual algorithm for minimizing the total variation norm that Gene and I developed with Pep Mulet a decade ago is related to recent work by A. Chambolle on dual projection algorithms and by W. Yin and D. Goldfarb on second-order cone algorithms for the same problem.

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 sev-

eral 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 $\Pi = E_p^{-1} A_p E_{p-1}^{-1} A_{p-1} \dots E_1^{-1} A_1$ 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)

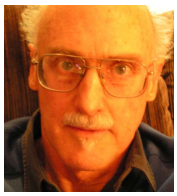

The rate of convergence of iterative methods has always been one of the interests of the Numerical Analysis group at Stanford. For example I read "On the asymptotic directions of the s-dimensional optimum gradient method" by George Forsythe (1968) while refereeing a paper for Gene Golub when I visited 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 $Ax = x\lambda$, $\lambda = x^H Ax \in \mathrm{R}$, then the Rayleigh quotient $y^H Ay$ satisfies

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

Here, if $\lambda_1(A) \geq \cdots \geq \lambda_n(A)$ are the eigenvalues of A in descending order then 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 $X, Y \in \mathrm{C}^{n \times k}$ be such that $X^H X = Y^H Y = I_k$, where $\mathcal{Y} \equiv$ range(Y) is an approximation to the invariant subspace $\mathcal{X} \equiv$ range(X) of A, so that $AX = X.X^H AX$. Let $\lambda(X^H AX)$ and $\lambda(Y^H AY) \in \mathrm{R}^k$ 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

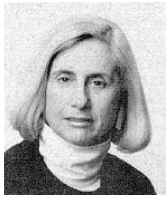

Statisticians, including this one, owe a huge debt to numerical analysts. Where would we be without the Singular Value Decomposition, Spline Algorithms, Matrix Computations (Golub and Van Loan)?

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

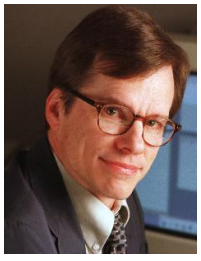

Michael Heath
Univ of Illinois,
Urbana-Champaign

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

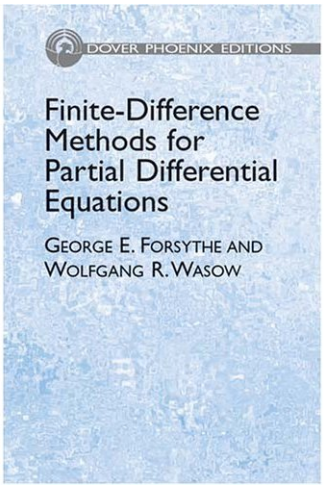

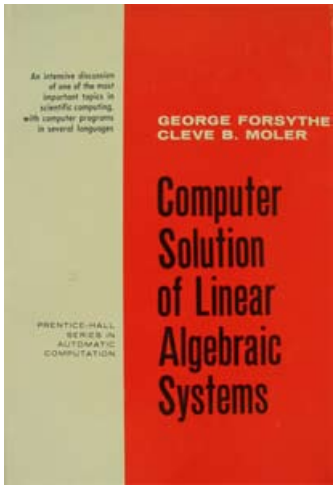

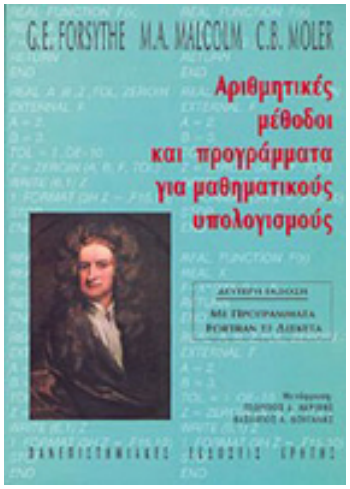

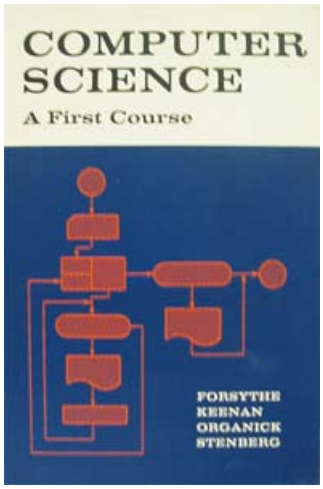

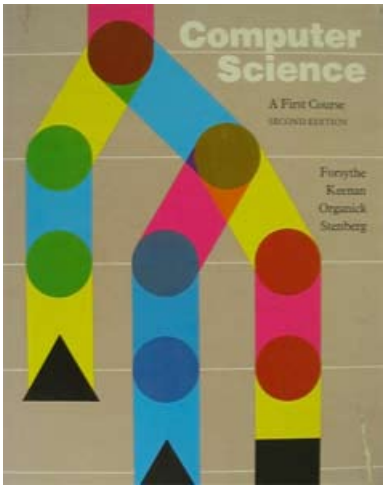

Some of Gene Golub's Books

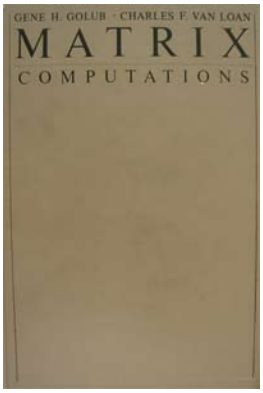

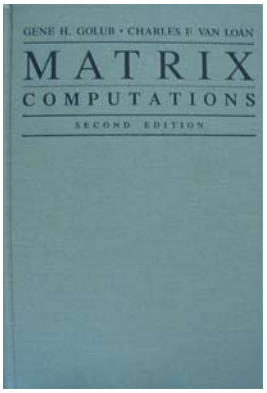

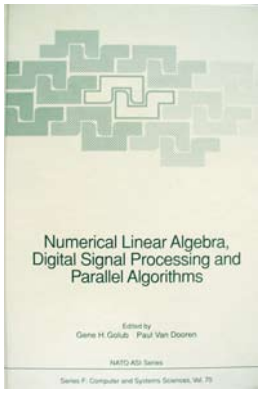

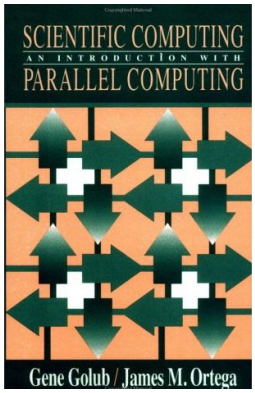

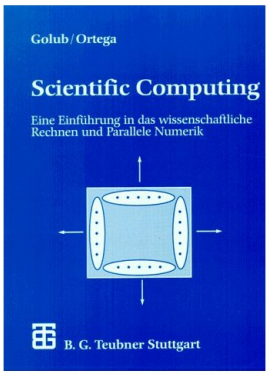

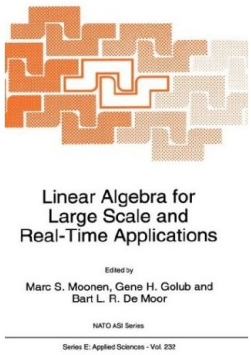

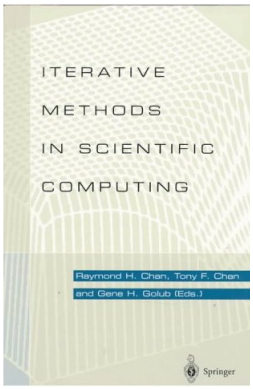

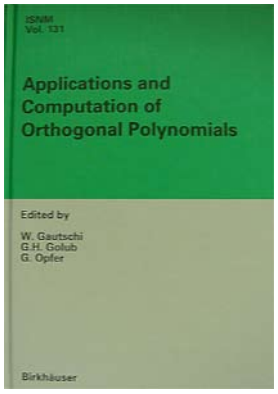

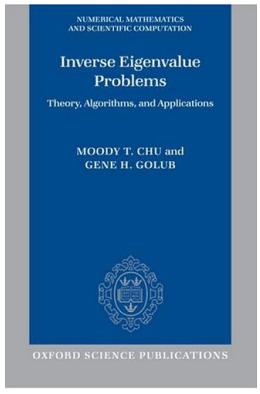

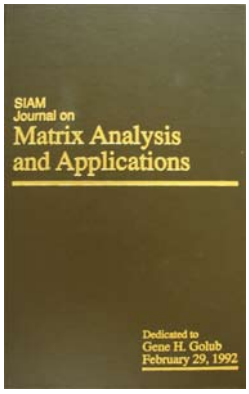

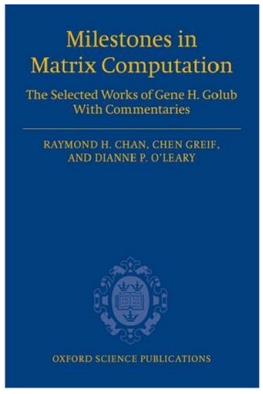

GENEalogy

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!

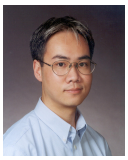

Wing Lok Justin Wan
Forsythe Fellow
Univ of Waterloo
Canada

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 Poly-
technic 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

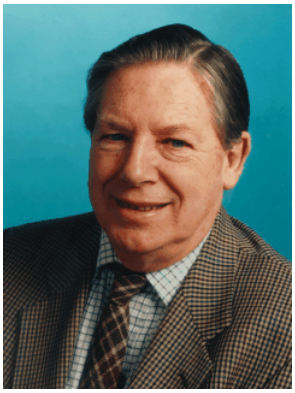

Roger Hockney
19xx–1999

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

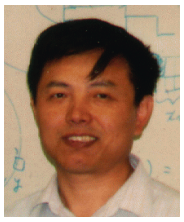
Ding-Zhu Du
Univ of Texas, Dallas


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