An Interview with Gene Golub

Higham, Nicholas J.

2008

MIMS EPrint: **2008.8**

ISSN 1749-9097

An Interview with Gene Golub[1]
by Nicholas J. Higham[2]

On July 3, 2005 I interviewed Gene Golub (1932–2007) during a visit he made to The University of Manchester to attend a workshop. This document provides an edited transcript of the interview.

The bibliography contains some books and papers mentioned directly or indirectly in the interview. For more on Golub and his work, I highly recommend *Milestones in Matrix Computation: The Selected Works of Gene H. Golub, with Commentaries* [3].

I am grateful to Louise Stait for transcribing the interview, and to Gail Corbett and Sven Hammarling for help with the editing.

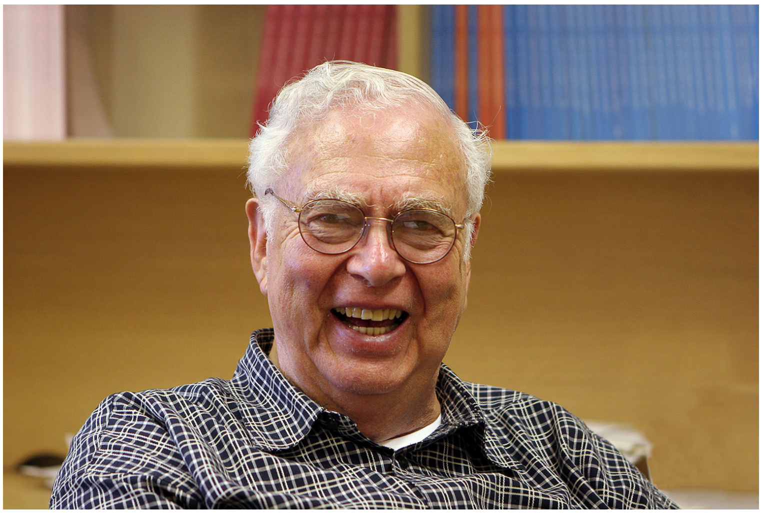

Gene Golub, July 2005. Photograph by N. J. Higham.

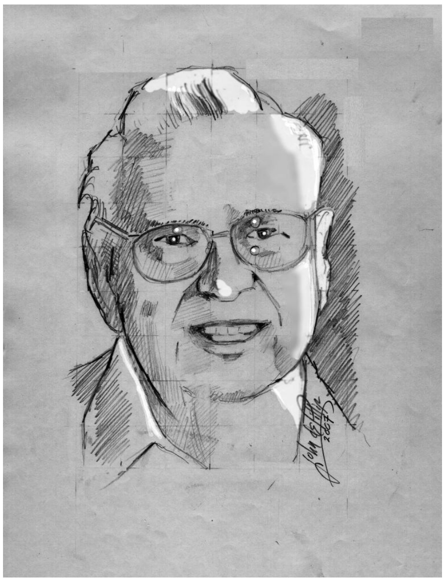

Gene Golub by John de Pillis (http://math.ucr.edu/~jdp), 2007.

[1]Document dated February 5, 2008.

[2]School of Mathematics, The University of Manchester, Manchester, M13 9PL, UK (higham@ma.man.ac.uk, http://www.ma.man.ac.uk/~higham).

NJH: How did you get interested in mathematics?

GHG: Originally, when I went to college, I wanted to become a chemist and thought of other things too, even becoming a political scientist, but I found that the subject that I liked to do the most and was best at was mathematics. I enjoyed the manipulations a lot and I still do.

NJH: Did you have any particular mentors at any stage of your career, up to, say, your Ph.D.?

GHG: Well, not very strong, but some people were quite influential. As a senior at the University of Illinois I took a statistics course from a man named K. A. Bush. He called me aside and said, "Would you like to major in statistics? I think you would do very well there." So he certainly played a strong role there. Perhaps I should explain my whole education. The first two years I went to a junior college in Chicago. It was called Wright Junior College and although none of the teachers were engaged in research, as far as I knew, they were excellent teachers: they were prepared and they were careful. And then from there I decided to go to the University of Chicago, which is such a tremendous change, and one of my teachers warned me that it probably was too dramatic a change. Nevertheless, I went to Chicago. The University of Chicago was a great distance from where I lived; it took an hour and a half to get there from where I was living at the time. So after a year's time I decided I really didn't like it and that I'd like to go away to university. But at the time I was growing up, going away to university was really done by people who had some money behind them. I didn't have any; my mother was a widow. So I had worked and I thought this would be my big splurge, going away to the University of Illinois for my final year. I came down to the University of Illinois in 1952, and I took several courses because there was a necessity to complete certain qualifications. I also took a programming course in January 1953 from J. P. Nash, who talked about programming for the Illiac. The Illiac had just become available and Nash was very influential in my life; he taught me how to program. At that time programming and elementary numerical analysis were included together. At the end of the term I was graduating from the University of Illinois. I had no idea what I was going to do. I had applied for various positions, but Nash called me aside and he said, "Would you like to be an assistant in the computing laboratory?" I knew I liked going to school, and although I applied elsewhere for real jobs I thought it was fun to go to school. I liked computing a lot even then; I wasn't particularly good at it but I liked it. That summer I was already an assistant and I guess I was making $95 a month, so that was good money for me then. I had a task to do. I was supposed to program Milne's method for solving differential equations. Well, I didn't really think it out very well. I think it was too big of a project for me, and of course it was later shown that Milne's method is a weakly stable method. But I worked on it and it kept me busy during the summer. I became a full-time student during the autumn of that year, 1953. But my main interest was in statistics. I was really going to do computing and statistics. I took a statistics course from a man called C. R. Rao. C. R. Rao is a great statistician and as it turns out he was visiting Illinois just that year. The course was probably beyond me, but he did lots of matrix manipulations: for instance, he didn't call it this, but he developed block Gaussian elimination and so forth. So that really expanded my vision. Also at Illinois there was a large group of people who did psychometrics and there was really a well established programme. And of course they had the Illiac, so it was one of the first psychometrics groups to have a computer available for it. In particular, there was a man there named Charles Wrigley, who was an assistant professor. Wrigley was from New Zealand originally and he had gotten his Ph.D. in England under Sir Cyril Burt, the man who had so much to do with the 11+ examination and may have faked the data in fact. And Wrigley was a lovely man. I don't know how to describe him, but he was a very kindly gentleman and he said, "Oh Gene, you are doing such interesting work and these matrix computations are so important." The thing we were able to do was compute the eigenvalues and eigenvectors of a 23 by 23 matrix in 15 minutes. Well, that for then was a big breakthrough because there were some classical cases in psychometrics that had never been done so extensively. I learnt a lot about factor analysis and in a way it's just one step away from the singular value decomposition. Rao devised a method called maximum likelihood factor analysis, or canonical factor analysis, which was a new form of factor analysis. The big question in factor analysis at that time was, "How many factors should there be?", and it's equivalent to saying, "How many nonzero singular values are there?" Rao had devised a statistical test for doing this. Much of this work was actually contained in the work by Lawley on factor analysis. So I was just very fortunate: J. P. Nash brought me to the computing laboratory, through that laboratory I met Charles Wrigley, I met C. R. Rao, and I met many other people. My first lecture was at the International Congress of Psychologists. There I met Louis Guttman, who is considered a great psychometrician, and Hans Eysenck from Britain, who was a very famous man. I never met Sir Cyril Burt. I met a lot of people through this connection and, as you can tell, it influenced the sort of things I liked to do.

NJH: Who was your Ph.D. advisor?

GHG: I was going to write a thesis in statistics, but the man who I was going to work under left abruptly. Then one of the leading professors at the computing laboratory, a man named A. H. Taub, agreed to be my advisor. Taub was quite a significant person at that time; he had worked with John von Neumann, and he had close connections with the whole Los Alamos community. He came to Illinois from the University of Washington. Taub was also the advisor of Bill Gear, Bob Gregory and other people too. So he had some very good students. He was not an easy man to work with. Taub took me on and he had a paper or a draft of a paper that von Neumann had written on using Chebyshev polynomials, and my thesis grew from there. While I was working on my thesis Taub invited Richard Varga to come by, just as I was finishing. Taub wasn't so interested in solving linear equations, but Varga came by and gave a talk and afterwards I told him what I had done and he said, "Oh, I've done something similar, let's write a paper together." Then my advisor said to me, "I understand Mr. Varga has done something similar. If he publishes before you, you don't get a degree because the thesis is supposed to be original." So I was a bit worried and I didn't really know how people behaved in the academic world, but I was pretty confident. Then the next year while I was in England on my NSF fellowship, Varga and I worked on our paper together. It was wonderful working with him, but as you may know when you are working on your thesis you really know the subject area that you are working on, so I knew all the details and various little tricks and so forth. So the paper was really a joint collaboration. I learnt a lot in working on it with him but I think I also helped substantially in the writing of the paper.

NJH: So that time in Cambridge was during the Ph.D.?

GHG: No, just after my Ph.D. I had an NSF fellowship and I was delayed to some extent; I didn't finish as I had hoped on time, so it took me a year before I went. Illinois was a wonderful place. The whole computing laboratory had a great sense of enthusiasm and there were a lot of young Ph.D.s who were beginning their careers. One was D. E. Muller of Muller's method; he's the son of a Nobel laureate. He's a wonderful person; he's still alive fortunately, and he developed Muller's Method and did other things. Bill Gear was a graduate student. He came a couple of years after I did; he was pretty spectacular as a student. The whole atmosphere was really very warm and we were often invited to faculty's homes, so I in my small way try to keep this Illinois tradition of inviting people to my home and trying to be friendly. I was treated to that and I loved it, and it meant a lot to me. I should just say one other word about my own family background. Both my parents were poor immigrants to the United States. They came in 1923 and they had really inferior jobs: my father delivered bread and my mother sewed in a shop. They had no idea what the academic world was like, but they knew it was good to go to college and of course that was impressed on us. So I continued always as a student, and of course after a while, because of education, you are separated in a way from your parents too. My father died when I was 16, but my mother was supportive throughout my career. She was very happy to see me carry on and continue.

NJH: Where had they emigrated from?

GHG: My father came from Ukraine and just this last week I went to the Ukraine to see the town where he lived. It was sort of like family folklore and my mother came from Latvia; they met in Chicago. I've been able to find on the Ellis Island website exactly when they left and when they arrived and where they went to in Chicago. So that was the life experience for me, and the other people I went to school with, they had similar experiences. 1923 was an important year in immigration in the United States. The racism that we saw in Europe eventually ended up in the holocaust. It was true everywhere, and in the United States organizations like the Klu Klux Klan pressured for immigration controls and one of the things that came about was that there were quotas for different nations. The quotas were really established based on the nineteenth-century United States, so if you were English or Swedish or French you came to the United States with relative ease. People from eastern Europe after 1923 were pretty well frozen out so it was a bad time throughout the world, and racism prevailed everywhere I think.

NJH: Going back to your year at Cambridge who were some of the people that you met there?

GHG: Well, I had met Velvel Kahan before and we knew each other from a couple of visits at Illinois. Illinois and Toronto had wanted to build a second computer together and Velvel was their man, so to speak, so I knew him and we shared an office with two other people: one was Colin Cryer, who was a research student, and another person was Nick Capon from Australia (I think he is retired now). So the four of us were in this tiny office. Velvel and I were the two postdocs there, and I knew David Wheeler because Wheeler had visited Illinois. He was sort of my teacher in numerical analysis. I'd never taken a numerical analysis course, but he designed a library for the University of Illinois Illiac, and I would read his codes and I learnt a lot by reading those codes. For instance, Gaussian elimination with partial pivoting: that's what he had implemented and that's how I learned how to do Gaussian elimination. And you know it's very interesting; as I've stated before, what goes around comes around. Gaussian elimination with partial pivoting was very popular because the storage requirement was relatively small and then it sort of fell by the wayside when people were doing direct LU factorization. Then of course it was brought back into play again after a few years during the initial interest in vector and parallel computing. So I learnt a lot. Another person who was at Illinois was Stanley Gill; if you recall, one of the earliest books on programming was by Wilkes, Wheeler, and Gill. Gill had devised a version of Runge-Kutta that reduced the amount of storage, and so he spent a year at Illinois and I had several conversations with him. He was a very, very nice man. He was from England—he became professor of computing science at Imperial College—and then, unfortunately, he died early, and now of course we are saddened to hear of Wheeler's death. Wheeler was a very creative man, but very modest and a very pleasant person; he had wonderful ideas. Numerical analysis was not his primary interest but in the early days that was a major component of computer science.

NJH: I'd like to ask you about the background to just a few of your papers. One I picked out was the 1965 paper "Numerical Methods for Solving Linear Least Squares Problems" in Numerische Mathematik, where you introduced the Householder QR factorization. What is the background to that?

GHG: After I had got my Ph.D. and spent some time in England, I then went to Berkeley and I was working with a physicist there for about five months and I really didn't enjoy it. I wasn't doing any numerical analysis; I loved doing numerical analysis, and matrix computations. So I took a job in Southern California at Thompson Ramo Wooldridge, and that was a very enjoyable experience. My boss was a man called Sam Conte. I had spent a summer there before. I decided least squares was a nice subject to work on, and I saw Householder's work on using orthogonal transformations for solving least squares problems. So I worked it out in further detail, but really Householder did much of the work for least squares using what we call Householder transformations. But it really goes further back if you look at the book by Turnbull and Aitken. Aitken had the idea much earlier to do that, but my paper introduces ideas of updating and column selection and so forth, so it was a much more detailed analysis of how to use orthogonal transformation. But I also made a big gaffe because I thought you could do iterative refinement in a straightforward manner, but we learned that only in certain circumstances can you get a good approximation.

NJH: A little later, in 1969, you have a paper "Matrix Decompositions and Statistical Calculations". Was this motivated by your earlier statistical interests?

GHG: Yes, and the fact that I worked on all these orthogonal decompositions and so forth. So I tried to get the statisticians interested in doing numerical computations. A few people were interested in it, but I don't think it had a heavy influence in statistics. I doubt if today people really use decomposition methods rather than normal equations. Statisticians are fairly fixed in their ways. That to me is rather sad for the following reason. In the very beginning of statistics the statisticians were very much aware of computing and their limitations. So they developed ideas on this, and R. A. Fischer had ideas and some of the statistical techniques were influenced by numerical computations. For instance, there are what are called Latin square designs, and basically those lead to diagonal normal equations. So there may be good statistical reasons for this too, but basically the design of the statistical experiment is in some way an influence on numerical computation. I think the early statisticians were very knowledgeable about the fact that computing was necessary. Of course we find that Hotelling dismissed Gaussian elimination because he thought that would not be a good way to proceed, and then he reinvented the Newton-Schulz method for inverting the matrix. So it's quite interesting: there was a big gap. In modern times, within the last fifteen years I'd say, there's really been more interest in numerical computing. There's a book by John Chambers, who has done work in numerical computing.

I'll digress here a bit. I spent a year in Zurich and I had some connection with Peter Huber. Huber is really a very fine statistician and he worked with robustness in statistics and various other things. He trained some students in Zurich and he was very much enamoured by computing. One of his students was a fellow by the name of Werner Stuetzle, and Stuetzle came to Stanford as a postdoc or assistant professor for a few years. And then he was involved with several people. I don't know if you are familiar with these names. Jerry Friedman, who came to Stanford; he's a physicist, then he got involved in statistical work because of Tukey. Then two young men who worked with Stuetzle, Rob Tibshirani and Trevor Hastie, and Werner was really influential with those people. Tibshirani and Trevor Hastie were younger, but they were all very good friends and later Werner left and he became the head of statistics at the University of Washington. These three guys have gone on and had an enormous effect in statistical computing at Stanford. They have a very well cited book, and then of course later at Stanford they hired David Donoho. Statistics is often regarded as a model for numerical analysis. It has some of the same kind of strengths and weaknesses that we do. When I was a graduate student they were going through their theory mode and they were trying to prove things about optimality and so forth. At least in the US people had lost their way to some extent, and they were trying to prove theoretical results. In the meantime, here in Britain, I think people have always had a very good eye for doing what you might call real statistics. In the US people admired the way British statisticians worked, but somehow the theoretical was emphasized, and it's only within the last fifteen years or so, as I see it, there has been a re-emphasis on re-thinking using computational methods in statistics, working in that direction of interesting statistical computations. I think it's heavily influenced the way statistics is done today.

NJH: Another very influential paper of yours is "Some Modified Matrix Eigenvalue Problems" in SIAM Review in 1973. That was one of your earlier papers on this topic and has led to a lot of other work, hasn't it?

GHG: Well I'm pleased that you say that because the origins are as follows. I was on sabbatical one year and I had been collecting results. For instance, I worked on eigenvalue problems with homogeneous constraints; that was at the suggestion of Henrici, as a matter of fact, and I devised a very simple algorithm for doing a computation, and then I discovered that this problem, minimizing quadratic forms subject to linear constraints, comes up in very many different places. So I had been working on a lot of small results, and then the time came I had to show I had done something on my sabbatical. So I wrote down this review paper. I love that paper, in part because it has a lot of nice useful tools for solving matrix problems.

NJH: You must have fond memories of Serra House, which was the Computer Science Department's home on the Stanford campus. Is there anything in particular that stands out in your memory about it?

GHG: Well, it was a terrific period, and originally everybody in the Computer Science Department and Computation Center was housed in Polya Hall. Polya was a professor at Stanford, and he was in fact even alive when Polya Hall was built and he was very much loved by everybody. Our department and the center kept on growing, and the center wanted of course more and more space. I don't know if you recall the large computation centres; they had rooms full of computers and people and everything else. So we outgrew our building and then there is always a scramble to find more space. The numerical analysts found themselves in Serra House. We had the ground floor of the house and people who worked in artificial intelligence had the upper floor. Forsythe's office was still in Polya Hall when he was head of the department. When he died he was still head of the department, and Jack Herriot, who worked with George on administering the department, was also at Polya Hall. But the younger numerical analysts and myself were all in that building, and as we had visitors like Wilkinson, Dahlquist, and Gautschi and others they all sat in Serra House. It had a nice atmosphere. We had a very wide hallway, and I collected reprints and they were all in this hallway. Even our secretary sat in the hallway as space got tighter and tighter. We had a kitchen, but the kitchen was turned into a terminal room. It was a very nice period. You may have seen the directory for Serra House at one time, with people like Dahlquist and Wilkinson and then the array of students that we had like Nick Trefethen, Randy LeVeque, Marsha Berger, Petter Bjorstad; there was just a great group of students that we had at the time.

NJH: You mentioned Wilkinson, who you obviously knew very well for many years. Do you have any particular memories of him?

GHG: Well I can't sum it up in a very simple statement. He was a wonderful man. I think not only was he brilliant, but he was full of life and he knew so much. He knew about music and literature, he loved a drink, he loved life. It was just good to be in his presence. He was a very charming, helpful person. Everybody felt happy being around him and loved to talk to him. He had a nice sense of humour, and he was altogether a very good person. He was special and I miss him greatly.

NJH: Over the years you've been a strong advocate of the Lanczos method. Did you ever meet Lanczos?

GHG: On two occasions. I met him at NPL one time. It was very auspicious because I went down with Velvel Kahan, who was at Cambridge too, and Lanczos talked about the singular value decomposition. You know, one of the nicest descriptions of it is in his book *Linear Differential Operators*. Afterwards I tried to speak to him. I said "P-p-r-o-ff-essor L-l-anczos...", and I was wanting to tell him about my work on Chebyshev polynomials. Then two powerful arms came around my shoulders and moved me aside. It was Velvel Kahan; he felt he had first priority. Lanczos later spoke at a meeting in Ireland that John Miller had organized. He gave a talk about computing instruments, computing methods, and he went slowly through the history until about 1940 and then he skipped, and the next thing he showed us an HP hand calculator. So he really hadn't come to terms with modern computing. I tried to talk to him afterwards about the SVD, for which I had by that time developed an algorithm, but he didn't seem very interested, and he died a couple of years after that. He was of course a great influence on all of us. I wish we could talk to him a little bit more because I've often felt the Lanczos method is really the Stieltjes algorithm with a different inner product, and he must have known that.

NJH: Another paper I'd like to ask you about the background to is your 1969 paper "Calculation of Gauss Quadrature Rules" with Welsch in Mathematics of Computation. How did you get interested in that particular topic?

GHG: Well, I'm not sure how I became particularly interested in quadrature. Paul Concus was a dear friend of mine and asked me some things about quadrature, and then I realised that all that's really necessary is to compute eigenvalues of matrices. I was seated in Peter Lax's office. I was on leave at the Courant Institute and I was sort of bored by what was taking place, so I happened to see a book by Wilf. There he had some of the relationships about computing the quadrature weights, and then I worked out the fact that all you need is the square of the first element of the eigenvector. Actually, other people had figured this out too and known this before. What I was able to do is to recognize that when you mix that with the QR method you can organize the QR method just to compute the first component of an eigenvector. So I was really pleased when that all fell into place. Later on I saw Rutishauser, and I told him that I had this algorithm for computing the quadrature rules and how I was doing it. The interesting thing is that he denied it could be done that way. Ordinarily I'm not a very confident person, but I knew there was a program that was working and that actually calculations could be done that way. What pleased me of course was that it eliminated all the use of all these large tables, and the knowledge that you can compute a quadrature rule in this way, and with the weights, has really been very useful in many different studies.

NJH: Many of your papers are joint. Do you have any particular way of working with your collaborators?

GHG: No. I think in current times I really talk to a lot of young people and I sort of suggest a way of doing something or other and then they go off and do it, and then there is a collaboration that way. I like working with people. I don't have a high ability of writing on my own, so I like to work with people who are willing to collaborate in that fashion.

NJH: Have you ever felt that a result, paper or book of yours has not received the attention it deserved, and, if so, which and why?

GHG: The work I've done on matrices, moments and quadrature—I don't think that is widely known, and yet I think it's very rich and I'm hoping there will be a book by Gerard Meurant and myself on that topic. On the whole, sometimes I'm astounded by the fact that people know of a paper that I have written. Actually, the books with Ortega are not so well known, and they contain a lot of good things.

NJH: Do you have any comments on the relationship between core linear algebra and numerical linear algebra, and how it has changed with time?

GHG: I feel most people who are in application areas learn linear algebra because they really want to do matrix analysis and I think the subject has been hijacked by the mathematicians who understand about applications; but when they teach their courses they don't necessarily emphasize the aspects we're familiar with. My own feeling is that the deeper you understand the matrix computation aspect, the better you understand what linear algebra is about. For instance, rank is a well known concept in mathematics. In truth, rank is a discrete topic, but if you look at it from a numerical point of view it almost has a continuous flavour to it. So I feel it's unfortunate, but the mathematical world has hijacked the linear algebra topics when it should really be matrix algebra, because people I think are interested in matrix analysis. If you recall, the journal SIMAX[3] has the phrase "Matrix Analysis and Applications" in the title. I did that very deliberately. I didn't want to see another linear algebra journal. My motto for this is the book by Bellman on matrix analysis and applications, so in a sense the journal is a tribute to him.

NJH: You actually founded two SIAM journals. Before SIMAX was SISC[4] in 1980,

[3]The SIAM Journal on Matrix Analysis and Applications, founded in 1988.

[4]The SIAM Journal on Scientific and Statistical Computing, renamed in 1993 the SIAM Journal on Scientific Computing.

so what was the motivation for SISC?

GHG: Well, Jim Ortega and I were working together on things. Ortega was really a great friend and colleague. We recognized that such journals as SINUM[5] were going so heavily into the theoretical aspects of computing, of numerical analysis. So we really wanted to redress the subject, and it's worked partially. But of course there is always a drifting that takes place.

NJH: I think you were involved in setting up NA-NET. Is that right, and what's the background of that?

GHG: Jim Wilkinson was at Stanford and he just said, "What's the address of so and so?", "What's the address of so and so?" So then I thought we should have an alias for everybody. Of course it wasn't so easy to find people's addresses then, so I made up a little list. I had seven or eight names and then we had somebody working in the computing centre who helped me out, and so we devised the beginning of NA-NET and the alias format. Originally, as messages came in we just circulated them, and then we decided to go into the digest format. I only did a couple of digests before we had the current format and that's really bloomed under Cleve's hand.

NJH: I must ask you about your book *Matrix Computations* with Charlie Van Loan, originally published in 1983, now in its third edition (1996). When did you actually start writing the book?

GHG: I gave some lectures at Johns Hopkins University. What happened is, they formed a Math Sciences Department at Johns Hopkins and the head was Roger Horn. I had known Roger Horn because he was in the first course of advanced numerical analysis I taught. I taught three quarters of the course; the first two quarters were out of Varga's book. Roger Horn is a wonderful fellow, and I hadn't realised what impact that had on him because, for instance, he told me he learned about the Perron–Frobenius theory from me and also the singular value decomposition, and those have both been of some importance to him. So Johns Hopkins had a series of lectures, and he invited me to give lectures. I gave ten lectures in two weeks or so. The class consisted of a whole bunch of people; I don't know specifically who now, but Charlie Van Loan was there and also Richard Bartels. Originally the idea was for the three of us to write a book, but then Richard's interests were developing otherwise; it was going to include a large component of math programming. So then it was just Charlie and I. He did most of the writing, as you probably know, but it was really a great stimulant for both us. For instance, although I had done some work on total least squares, we developed a paper on total least squares with a lot of analysis and showing how the computations will go, and so forth. So that was really a wonderful effort, and he would send me some section and I would say, "Oh, well, we can do better like this." There is the Sylvester equation that we did work on; again I looked at it and realised we could do one eigenvalue computation: we did one upper Hessenberg form and one Schur decomposition. So that was really a great period for both of us because we were incorporating new material in the book and we were also writing the book.

NJH: Were you surprised at how successful the book has been?

GHG: No, not at the beginning at any rate, because the publisher kept on saying, "When are you going to get that book out?", and I kept on saying, "Don't worry, it's going to be a good book. This will really have a market"—and it did. It's now sold over 50,000 copies, the three editions. If you Google the book you'll find that there are over 10,000 references, so Charlie said, "I think every 18 hours a paper is being

[5] The SIAM Journal on Numerical Analysis.

written which has reference to our book." I'm pleased that it's been a help to people. Electrical engineers I do know like it, and people in various specialties have found it of great use. We are planning in about a year, when Charlie is no longer chairman and will be on sabbatical, that I'll take a sabbatical and we will revise the book. So there is a question, how to limit a book: should you have two volumes, and if you do have two volumes what should they include, and so forth. I think it's been a great windfall for Johns Hopkins. You know, most of their editions probably sell in the hundreds, or a few thousand, so they have probably done well, and I am pleased for that. I think the commercial presses have acted in very poor ways, and I think it's good that academic institutions and professional organizations get the benefits of the books.

NJH: You've always travelled extensively. Are there any countries you haven't visited academically that you would like to visit?

GHG: I'd like to go to South Africa. I've had invitations, I've accepted them, and then somehow I couldn't quite make the long trip. There are some excellent people there.

NJH: One of your current research topics is the linear algebra of search engines. Are you surprised by this new application of this subject?

GHG: Not completely; I'm just so pleased. The way it happened is a student asked me about finding the stationary probability distribution of a large matrix. I don't think he knew exactly what he was asking about, so I spent a few hours. Of course, everything I suggested could have been done fifty years ago and it's really an acceleration of the power method. Now I have another paper with Chen Greif where we use the Arnoldi method for accelerating convergence, and that seems to work quite well too. I taught a course about linear algebra and computer science and one of the students was from Google. He said, "Well, that's only the stuff you talk about; how we order things is just a small part of search engines." Secondly, he said, "We use the singular value decomposition all the time", but he didn't go on any further to tell me how it's being used. So we're lucky at Stanford. Even I have to admit I have had a successful career. I think two great components are, first, Stanford has served as a magnet. People have come through, and I think I have greatly benefited by being in such a strong computer science department. Secondly, the Bay Area. It's not quite as true now as it was, but we have these various laboratories in the area: Lawrence Livermore Laboratory, Lawrence Berkeley Laboratory, NERSC, IBM had some activity in the past, and Intel. So the area within a fifty mile radius is honeycombed with organizations, and some are interested in numerical mathematics—well, some *should* be interested in numerical mathematics—so it's really fortunate being here.

NJH: What about future research plans?

GHG: Well, I've never really had a fixed programme in research, and I just like to do what comes up. I'm concerned that maybe I've over-farmed what I know. On the other hand, there always seems to be another application of Householder transformations and the singular value decomposition, but I reckon I'm coming to the end of that line of research. It's a complicated situation. The younger people, of course, they go in different directions. There are still problems of the sort that I mentioned, and maybe there aren't so many people who know all the intricate details that I know, although some people know them better.

NJH: Is there anything I should have asked that I didn't—that you think you would have wanted me to ask?

GHG: There are some people I would like to talk about. One is George Forsythe. He established what we have at Stanford, although in retrospect I have been there three times as long as he was there. He came in 1957, so it's close to fifty years, and in fact I am hoping to have some sort of celebration of that[6]. He had some wonderful students, like Moler, Varah, Ortega and Parlett. He really had first rate people, so really established a tone in numerical analysis that enabled us to build up something which was really good. Forsythe has been dead for over 30 years, but he was an important person, not because he was very original, because there aren't so many things you can point to as contributions from Forsythe, but he just pointed people in the right direction. In fact, one reason I worked on the SVD is because at the end of a lecture by Ben Rosen, Forsythe said, "Would somebody please figure out how to compute the pseudo-inverse of a matrix." He was very supportive of students. This came out again the other day: that he wanted to bring Moler to an early Householder meeting, and people actually opposed that, and Forsythe said, "Well, if you don't want Moler then I'm not going to attend", so just think now what an important component of a Householder meeting is having students there. Another man I greatly admire—I'll give you my icons in the field, partially—is David Young. I've read a lot of his thesis and his early work, and his thesis is genuinely a wonderful piece of work and I'm afraid people won't know it so well; there is that parameter that's so hard to compute, it seems. But the analysis there is just breathtaking in terms of the originality and showing what you can do with an algorithm by choosing one little number properly. Again he led the way. Later on people wrote about block SOR and various other aspects of the problem, but his thesis was really a fantastic piece of work, and I think it's something that we can all greatly admire. Dahlquist is one of my personal heros, not only because of his brilliance as a scientist, but as a human being he just was very concerned about other people, and he had many very good ideas. He was a very original man and a great scientist. Another person I admire greatly is Walter Gautschi, because he is such a serious scientist, such a serious scholar; he has done a lot of wonderful work. Those are just a few of the names off the top of my head; there are all kinds of other names.

NJH: One other thing that just comes to mind, because I was reading a book about this last week from SIAM, is the Mathematics Research Center at Madison. Did you have much contact with that?

GHG: Yes, in fact they even offered me a position at one time. I went to Chicago to see my mother and then I was going to go on to MRC for a few weeks. When I was in Chicago I had some spare time and I went to see a man called Victor Barcilon. He was telling me about his work on inverse eigenvalue problems, and he said to me he wanted to solve not the tridiagonal case but the five diagonal case. After about a week in Chicago I went up to Madison and there I started working with Carl de Boor. We worked on the tridiagonal case and while we were working on it I recognised the use of the Lanczos algorithm for inverse eigenvalue problems. We wrote a paper which sort of got things right in the inverse eigenvalue problem for tridiagonal matrices. Then I went back to Stanford and I started working on the five diagonal and the general case. Ironically, I sent Victor Barcilon a copy of the paper, but, as it turns out, although you more or less get a unique solution for the tridiagonal case when choosing the sets of eigenvalues, you don't get a unique solution if you have three sets of eigenvalues and want to build up the five diagonal case. Indeed, Gil Strang at a meeting in Moscow recently was talking about the problems arising when you have $n(n+1)/2$ eigenvalues—

[6] http://compmath50.stanford.edu/

the eigenvalues of all the leading principal submatrices; you are far from uniqueness in any of these cases. So Barcilon said, "Well, I'm not interested and if you can't give me a unique solution I don't want the solution", and he wrote some other paper. By the way, I now have a book with Moody Chu that just came out on inverse eigenvalue problems.

So I love that subject. I did work on Gauss quadrature and then I did some work on QR-like algorithms for tridiagonal matrices and stuff on moments, and then suddenly one day in the Lanczos algorithm it all came together and it was just wonderful. It was a big thrill to see that they all related to one another.

The other thing I have been so proud of I guess are the students that we've had at Stanford. Some have been my students and some have not been my students. We've just been very fortunate in the kinds of people that have come to Stanford. It's really a great period that we've had. My personal life hasn't been as happy as I would have liked, but when I look back at this aspect of life it's something that's going to give me great pleasure.

So when I think about the influences of my life, I think being a student at the University of Illinois, that really had an enormous impact on my life. Being at Cambridge for a year—although I didn't do very much—in terms of friendships I still have many friends even though it's over 45 years since I first came here, and not all of them are connected with the profession, but they all had a great impact on my life. And also being at Stanford of course; I've been there since 1960, so 45 years. The way we meet people is serendipitous: we bump into people, we talk to people, so that's really been a great satisfaction.

Bibliography

[1] Richard Bellman. *Introduction to Matrix Analysis.* McGraw-Hill, New York, second edition, 1970. xxiii+403 pp. Reprinted by Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1997. ISBN 0-89871-399-4.

[2] John M. Chambers. *Computational Methods for Data Analysis.* Wiley, New York, 1977. ISBN 0-471-02772-3. xiii+268 pp.

[3] Raymond H. Chan, Chen Greif, and Dianne P. O'Leary, editors. *Milestones in Matrix Computation: The Selected Works of Gene H. Golub, with Commentaries.* Oxford University Press, 2007. ISBN 978-0-19-920681-0. xi+565 pp.

[4] Jagdish Chandra and Stephen M. Robinson. *An Uneasy Alliance: The Mathematics Research Center at the University of Wisconsin,* 1956–1987. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2005. ISBN 0-89871-535-0. xi+89 pp.

[5] Moody T. Chu and Gene H. Golub. *Inverse Eigenvalue Problems: Theory, Algorithms, and Applications.* Oxford University Press, 2005. ISBN 0-19-856664-6. xvii+387 pp.

[6] Jack J. Dongarra, Gene H. Golub, Cleve Moler, and Keith Moore. Netlib and NA-Net: Building a scientific computing community. MIMS EPrint 2007.123, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, 2007. 18 pp. To appear in IEEE Annals of the History of Computing.

[7] G. H. Golub. Numerical methods for solving linear least squares problems. *Numer. Math.*, 7:206–216, 1965.

[8] G. H. Golub and J. H. Wilkinson. Note on the iterative refinement of least squares solution. *Numer. Math.*, 9:139–148, 1966.

[9] Gene H. Golub. Matrix decompositions and statistical calculations. In Roy C.

Milton and John A. Nelder, editors, *Statistical Computation*, pages 365–397. Academic Press, New York, 1969.

[10] Gene H. Golub. Some modified matrix eigenvalue problems. *SIAM Rev.*, 15(2): 318–334, 1973.

[11] Gene H. Golub and Gerard Meurant. Matrices, moments and quadrature. In D. F. Griffiths and G. A. Watson, editors, *Numerical Analysis* 1993*, Proceedings of the* 15*th Dundee Conference*, volume 303 of *Pitman Research Notes in Mathematics*, pages 105–156. Addison Wesley Longman, Harlow, Essex, UK, 1994. Reprinted in [3].

[12] Gene H. Golub and James M. Ortega. *Scientific Computing and Differential Equations: An Introduction to Numerical Methods*. Academic Press, New York, 1992. ISBN 0-12-289255-0. viii+337 pp.

[13] Gene H. Golub and James M. Ortega. *Scientific Computing: An Introduction with Parallel Computing*. Academic Press, New York, 1993. ISBN 0-12-289253-4. viii+442 pp.

[14] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, USA, third edition, 1996. ISBN 0-8018-5413-X (hardback), 0-8018-5414-8 (paperback). xxvii+694 pp.

[15] Gene H. Golub and John H. Welsch. Calculation of Gauss quadrature rules. *Math. Comp.*, 23(106):221–230, 1969.

[16] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer-Verlag, New York, 2001. ISBN 0-387-95284-5. xvi+533 pp.

[17] Cornelius Lanczos. *Linear Differential Operators*. Van Nostrand, New York, 1961. ISBN 0-486-68035-5. xvi+564 pp. Reprinted by Dover, New York, 1997.

[18] N. J. Mackintosh, editor. *Cyril Burt: Fraud or Framed?* Oxford University Press, 1995. ISBN 0-19-852336-X. vii+156 pp.

[19] H. W. Turnbull and A. C. Aitken. *An Introduction to the Theory of Canonical Matrices*. Blackie, London and Glasgow, 1932. xiii+200 pp. Reprinted with appendix, 1952.

[20] Maurice V. Wilkes, David J. Wheeler, and Stanley Gill. *The Preparation of Programs for an Electronic Digital Computer. With Special Reference to the EDSAC and the Use of a Library of Subroutines*. Addison-Wesley, Cambridge, Massachusetts, 1951. Reprinted by Tomash Publishers, Los Angeles, 1982, Volume 1 in the Charles Babbage Institute Reprint Series for the History of Computing, with a new Introduction by Martin Campbell-Kelly.