


CERFACS, 42 Avenue Gustave Coriolis, 31057 Toulouse Cedex, France



ACTIVITY REPORT  
of the  
PARALLEL ALGORITHMS PROJECT  
at  
CERFACS

JANUARY 2001 - DECEMBER 2001

TR/PA/01/105



# Contents

---

<b>1</b>	<b>Introduction (I. S. Duff)</b>	<b>5</b>
<b>2</b>	<b>Dense and Sparse Matrix Computations</b>	<b>8</b>
2.1	Updates to the Rutherford-Boeing sparse matrix collection . . . . .	8
2.2	Rank-revealing and incremental norm estimation . . . . .	8
2.3	Development of kernels for sparse numerical linear algebra . . . . .	9
2.4	Candidate-based dynamic scheduling for a distributed direct linear solver . . . . .	9
2.5	An analysis of MPI send/receive in the context of MUMPS . . . . .	10
2.6	Use of orderings for large entries on the diagonal . . . . .	11
2.7	The sparse BLAS . . . . .	11
2.8	Activities of a senior visitor . . . . .	12
<b>3</b>	<b>Iterative Methods and Preconditioning</b>	<b>14</b>
3.1	When modified Gram-Schmidt generates a well conditioned set of vectors . . . . .	14
3.2	Real/complex seed GMRES and block GMRES packages for sequential and parallel computation . . . . .	14
3.3	Sparse symmetric preconditioners for dense linear systems in electromagnetism . . . . .	15
3.4	Combining fast multipole techniques and approximate inverse preconditioners for large calculations in electromagnetism . . . . .	15
3.5	Grid transfer operators for highly variable coefficient problems in two-level non-overlapping domain decomposition methods . . . . .	16
3.6	Domain decomposition methods in semiconductor device modelling . . . . .	17
3.7	Spectral two-level preconditioners . . . . .	18
3.8	Combining OpenMP and MPI on clusters of symmetric multiprocessors: some basic promising experiments . . . . .	18
<b>4</b>	<b>Qualitative Computing</b>	<b>20</b>
4.1	Inner-Outer iterations . . . . .	20
4.2	Arnoldi method and the happy breakdown . . . . .	21
4.3	Inexact computing . . . . .	22
4.4	Geometric aspects of computing . . . . .	23
<b>5</b>	<b>Nonlinear Systems and Optimization</b>	<b>24</b>
5.1	SifDec: a lonesome SIF decoder . . . . .	24
5.2	CUTEr: a Constrained and Unconstrained Testing Environment revisited . . . . .	24
5.3	Componentwise fast convergence in the solution of full-rank systems of nonlinear equations . . . . .	25
5.4	Optimization techniques for the regularization of large-scale discrete forms of ill-posed problems . . . . .	26
<b>6</b>	<b>Conferences and Seminars</b>	<b>28</b>
6.1	Conferences and seminars attended by members of the Parallel Algorithms Project . . . . .	28
6.2	Conferences and seminars organized by the Parallel Algorithms Project . . . . .	30
6.3	Internal seminars organized within the Parallel Algorithms Project . . . . .	31

<b>7 Publications</b>	<b>32</b>
7.1 Journal Publications . . . . .	32
7.2 Conference Proceedings and Book Chapters . . . . .	32
7.3 Theses . . . . .	33
7.4 Technical Reports . . . . .	33

---

# 1 Introduction

---

The research programme conducted by the Parallel Algorithms Project combines the excitement of basic research discoveries with their use in the solution of large-scale problems in science and engineering in academic research, commerce, and industry. We are concerned both with underlying mathematical and computational science research, the development of new techniques and algorithms, and their implementation on a range of high performance computing platforms.

The description of our activities is presented in several subsections, but this is only to give a structure to the report rather than to indicate any compartmentalization in the work of the Project. Indeed one of the strengths of the Parallel Algorithms Project is that members of the Team work very much in consultation with each other so that there is considerable overlap and cross-fertilization between the areas demarcated in the subsequent pages. This cross-fertilization extends to formal and informal collaboration with other teams at CERFACS, the shareholders of CERFACS, and research groups and end users elsewhere. In fact, it is very interesting to me how much the research directions of the Project are increasingly influenced by problems from the partners.

Members of the Team very much play their full part in the wider academic and research community. They are involved in Programme Committees for major conferences, are editors and referees for frontline journals, and are involved in research and evaluation committees. These activities both help CERFACS to contribute to the scientific life of France, Europe and the world while at the same time maintaining the visibility of CERFACS within these communities. Some measure of the visibility of CERFACS can be found from the statistics of accesses to the Algo Web pages at CERFACS where we have recorded more than 43,000 hits during the reporting year. On our publication pages alone, people have requested no fewer than 223 different reports.

Our main approach in the direct solution of sparse equations continues to be the multifrontal technique originally pioneered at Harwell in the early 1980s. During this last period we have further developed the MUMPS package in conjunction with our colleagues at ENSEEIHT and elsewhere. The code is now very robust and has been downloaded by over 150 researchers from the MUMPS website. As I write this, preparations are well underway for a significant new release that will include better integrated ordering options and more sophisticated scheduling strategies. Recent work has included an in-depth study of the MPI send/receive mechanism and investigation of techniques to make the MUMPS code robust in the presence of different implementations of MPI. We have also explored the dynamic scheduling strategy to avoid a severe overestimate of storage requirements.

At the level of international efforts for standards in numerical linear algebra, we have been very involved in the development of a new standard for the Basic Linear Algebra Subprograms (or BLAS) coordinated through the BLAST Technical Forum. During this year we have developed the Fortran 95 implementation of the sparse BLAS and have submitted papers on this and on the design of the sparse BLAS to ACM Trans Math Softw. The codes are available on the web and we have had over a 130 downloads.

Although iterative methods remove many of the bottlenecks of direct approaches, particularly regarding memory, it is now well established that they can only be used in the solution of really challenging problems if the system is preconditioned to create a new system more amenable

to the iterative solver. During this last year we have continued our work on developing such preconditioners, including two-level schemes that have been used successfully in both dense electromagnetics applications and in domain decomposition methods for solving partial differential equations in device modelling. Work has continued on the use of the MUMPS direct solver as a preconditioner within a domain decomposition scheme and the resulting algorithm and code have been used with success in the solution of problems from drift diffusion in semiconductor device modelling in a joint collaboration with INRIA. On a more theoretical but very practical track, the loss of orthogonality when performing a QR factorization using modified Gram Schmidt has been extensively analysed. The GMRES and FGMRES codes that were discussed in a previous activity report are available through the CERFACS web and have attracted over 600 downloads, some from important establishments including partners of CERFACS. Our use of sparse preconditioning techniques within a fast multipole code from our partner EADS has resulted in the efficient solution of dense problems in over 1 million unknowns from highly challenging models of an Airbus aircraft.

The main area of interest for the Qualitative Computing Group concerns a deep understanding of the influence of finite-precision computation on complex scientific numerical applications. One of the major areas studied has been that of inner-outer iteration, an important application being the use of Krylov solvers (inner iteration) within an eigensystem calculation (outer iteration). It was shown experimentally that almost counter-intuitively, it is possible to relax the accuracy of the operator or inner iteration in a controlled way without affecting overall convergence. A project with our partner EDF has involved moving this work from a research environment to one of their major industrial codes. The second area of important research has been on Inexact Computing, which is based on the theory of homotopic perturbations. A global theory has been developed which allows one to understand why numerical artifacts occur when pseudospectra are computed in finite precision. Thirdly, on a more esoteric level, algorithms and theory have been developed for the use of hypercomplex numbers which can be a more natural way of representing actual physical processes.

We have been very active in the optimization area, with significant work on the organization of a large scale testing environment and an extension of theory to include componentwise convergent methods and their analysis. The work on solving ill-posed problems using regularization techniques has been much advanced and extended to several different application areas.

The Parallel Algorithms Project is heavily involved in the Advanced Training aspects of CERFACS' mission. We ran an internal training course for new recruits to all Projects at CERFACS to give them a basic understanding of high performance computing and numerical libraries. This course was also open to the shareholders of CERFACS. We are also involved in training through the "stagiaire" system and feel that this is extremely useful to young scientists and engineers in both their training and their career choice. Stagiaires who successfully completed their training in the Team during 2001 included Caroline Bousquet, Christophe Daniel, and Claire Mandry from UPS, and Gregoire Richard and Pierre Valentin from ENSEEIHT. Members of the Team have assisted in many lecture courses at other centres, including ENSICA, INPT, Météo, Toulouse 1 and UPS.

We run a regular series of "internal seminars" that are primarily for Team members to learn about each other's work and are also a good forum for young researchers to hone their presentational skills. We run a series of seminars from external people who are visiting the Team that attract a wider audience and this has benefited greatly this year through an extensive visitor programme. Indeed, our list of visitors is a veritable who's who of numerical analysts, including many distinguished scientists from Europe and the United States. We have included a list of the visitors at the end of this introduction. John Gilbert from Xerox Parc in California was one of our longer staying visitors this year and was involved in many interactions as witnessed in Section 2.8 of this report. As always, it was a pleasure to host Gene Golub for a week when he came for the Anniversary meeting. He is particularly helpful and inspiring to our younger researchers with his enthusiasm and

interest. During this past year, Dominique Orban successfully defended his thesis on “Interior-point methods for nonlinear programming”, a cotutelle thesis between INPT and the University of Namur in Belgium. He has now gone to Northwestern University to work as a post-doc with Jorge Nocedal.

The most significant meeting hosted by the Parallel Algorithm Project was the Sparse Day meeting held in June. We were able to take advantage of our visitors to design a very attractive programme that attracted 38 participants from not only outside Toulouse but also outside France. Some researchers from the partners of CERFACS also participated. The speakers and programme are described further in Section 6.2 of this report. We were also strongly involved in the CERFACS Anniversary meeting in October. Gene Golub from Stanford University and Gérard Meurant from CEA gave invited talks and two PhD students, Bruno Carpentieri and Jean-Christophe Rioual had the chance to present their work to a distinguished audience.

I am very pleased to report that, over the past year, we have continued our involvement in joint research projects with shareholders and with other teams at CERFACS. We are involved in research committees such as the CCT of CNES. We have two projects with EADS on preconditioning techniques in electromagnetics and a sponsored PhD on the study of iterative solution techniques for multiple right-hand sides. We still assist in the training programme for the Mastere, organized by ENM. We are involved in a contract with EDF to utilize our finding on inner-outer iteration within their Aster code. We help the other Projects at CERFACS on almost a daily basis, for example in giving advice to CFD and TSI on OpenMP, but sometimes collaborate on a more substantial level as for instance on the optimal use of our public domain linear solvers with EMC. The work on the performance evaluation on clusters was also of great interest for both other CERFACS Projects and shareholders (in particular CNES and MTO).

### Visitors to Parallel Algorithm Project in 2001

In alphabetical order, our visitors in the year 2001 included: Neculai Andrei (Research Institute for Informatics, Roumania), Mario Arioli (RAL, UK), Matthias Bollhofer (Chemnitz University, Germany), John Gilbert (Xerox Parc, California, USA), Gene Golub (Stanford University, California, USA), Laura Grigori (INRIA-LORIA, Nancy), Gundolf Haase (Institute for Computational Mathematics, Austria), Eldad Haber (University of British Columbia, Canada), Per Christian Hansen (Technical University of Denmark, Denmark), Gary Howell (Florida Insititute of Technology, USA), Jacko Koster (Parallab, Bergen, Norway), Jean-Yves L’Excellent (INRIA-ENSL, Lyon), Martin van Gijzen (NWO, The Netherlands), Erricos Kontoghiorghes (University of Neuchatel, Switzerland), Daniel Loghin (University of Oxford, UK), Mago Magolu monga Made (Université Libre de Bruxelles, Belgium), Luiz Carvalho (IME-UERJ, Rio de Janeiro, Brazil), Osni Marques (NERSC, California, USA), Emeric Martin (IRISA, Rennes), Gérard Meurant (CEA, Paris), John Miller (Dublin University, Dublin), Shane Mulligan (Dublin University, Dublin), Beresford Parlett (Berkeley University, USA), Bernard Philippe (IRISA, Rennes), Miro Rozložník (Academy of Sciences of the Czech Republic, Prague), Yousef Saad (University of Minnesota, USA), Annick Sartenaer (University of Namur, Belgium), Jennifer Scott (RAL, UK), Dan Sorensen (Rice University, Houston, USA), Masha Sosonkina (University of Minnesota, USA), Trond Steihaug (University of Bergen, Norway), Guillaume Sylvand (INRIA-CERMICS, Nice), Philippe Toint (University of Namur, Belgium), and Jean Tshimanga (University of Namur, Belgium).

Visitors from our partners have included: Guillaume Alléon (EADS), Serge Gratton (CNES), Pascal Richard (CNES), and Jean-Louis Vaudescal (EDF).

Iain S. Duff.

---

## 2 Dense and Sparse Matrix Computations

---

### 2.1 Updates to the Rutherford-Boeing sparse matrix collection

**I. S. Duff:** CERFACS, *France and Rutherford Appleton Laboratory, England;*  
**D. Orban:** CERFACS, *France.*

The Rutherford-Boeing sparse matrix collection [1] extends its predecessor, the Harwell-Boeing sparse matrix collection in a number of respects, including the representation of the problems in the Rutherford-Boeing sparse matrix format.

With assistance from a stagiaire, Pierre Valentin, we have converted all the data from the Harwell-Boeing sparse matrix collection into Rutherford-Boeing format and are documenting this data and incorporating other relevant data. Ultimately, the collection will consist of a comprehensive archive of test matrices, stored in Rutherford-Boeing format, along with the relevant related data, such as right-hand sides, solution estimates, orderings, etc.

In order to leave as much latitude as possible and to allow older solvers to solve problems from the Rutherford-Boeing collection without modification, a number of conversion tools, to and from the Matrix Market format, the Matlab format and the Harwell-Boeing format, will also be provided. Tools to plot or view on the screen the sparsity pattern of the matrices have been designed.

- 
- [1] I. S. Duff, R. Grimes, and J. G. Lewis. The Rutherford-Boeing Sparse Matrix Collection. Technical Report RAL-TR-97-031, Rutherford Appleton Laboratory, Oxon OX11 0QX, UK, 1997.

### 2.2 Rank-revealing and incremental norm estimation

**I. S. Duff:** CERFACS *and Rutherford Appleton Laboratory, England;* **C. Vömel:** CERFACS, *France*

We have developed an incremental approach to 2-norm estimation for triangular matrices which is important for the detection of ill-conditioning, one of the basic problems arising in the numerical solution of linear systems. Applications of our scheme include the calculation of forward error bounds based on the condition number, robust pivot selection criteria and rank-revealing factorizations, in particular, when *inverse* factors arise in the factorization. In [2], we have introduced such a scheme applicable for both dense and sparse matrices which can arise for example from a QR, a Cholesky or a LU factorization. If the explicit inverse of a triangular factor is available, as in the case of an implicit version of the LU factorization, we can relate our results to incremental condition estimation (ICE) presented in [1]. Incremental norm estimation (INE) extends directly from the dense to the sparse case without needing the modifications that are necessary for the sparse version of ICE. INE can be applied to complement ICE, since the product of the two estimates gives an estimate for the matrix condition number. Furthermore, when applied to matrix inverses, INE can be used as the basis of a rank-revealing factorization. The quality of our results on standard test cases is consistently high and demonstrates the general reliability of our scheme. A revised version of [2] which contains also a theoretical analysis of our scheme will appear in BIT.

- 
- [1] C. H. Bischof. Incremental condition estimation. *SIAM J. Matrix Analysis and Applications*, 11:312–322, 1990.
  - [2] I. S. Duff and C. Vömel. Incremental Norm Estimation for Dense and Sparse Matrices. Technical Report TR/PA/00/83, CERFACS, Toulouse, France, 2000. To appear in BIT.

## 2.3 Development of kernels for sparse numerical linear algebra

**I. S. Duff:** CERFACS and Rutherford Appleton Laboratory, England; **C. Vömel:** CERFACS, France

The Basic Linear Algebra Subprograms for sparse matrices, the Sparse BLAS, are defined by the BLAS Technical Forum and consist of a set of kernels providing basic operations for sparse matrices and vectors. A principal goal of the Sparse BLAS standard is to aid in the development of iterative solvers for large sparse linear systems by specifying, on the one hand, interfaces for a high-level description of vector and matrix operations for the algorithm developer and, on the other hand, leaving enough freedom for vendors to provide the most efficient implementation of the underlying algorithms for their specific architectures.

The Sparse BLAS standard defines interfaces and bindings for the three target languages: C, Fortran 77 and Fortran 95. Our Fortran 95 implementation is intended as a reference model for the Sparse BLAS. The design is based on the idea of matrix handles so that the user need not be concerned with the details of the underlying storage schemes. It is envisaged that these kernels will be widely used in the solution of sparse equations by iterative methods. Our work is documented in [1], [2] and [3]. The software together with its description [2] has been submitted for publication as algorithm to ACM Trans. Math. Softw.

- 
- [1] I. S. Duff and C. Vömel. Level 2 and Level 3 Basic Linear Algebra Subprograms for Sparse Matrices: A Fortran 95 instantiation. Technical Report TR/PA/00/18, CERFACS, Toulouse, France, 2000. [Compressed PS](#).
  - [2] I. S. Duff and C. Vömel. The Implementation of the Sparse BLAS in Fortran 95. Technical Report TR/PA/01/27, CERFACS, Toulouse, France, 2001. Submitted to *ACM Trans. Math. Softw.* [Compressed PS](#), [PDF](#).
  - [3] I. S. Duff, C. Vömel, and M. Youan. Implementing the Sparse BLAS in Fortran 95. Technical Report TR/PA/00/82, CERFACS, Toulouse, France, 2000. [Compressed PS](#).

## 2.4 Candidate-based dynamic scheduling for a distributed direct linear solver

**P. R. Amestoy:** ENSEEIHT, France; **I. S. Duff:** CERFACS and Rutherford Appleton Laboratory, England; **C. Vömel:** CERFACS, France

The asynchronous distributed memory multifrontal solver MUMPS ([1, 2]) exploits three types of parallelism when a sparse matrix is factorized. A first natural source of parallelism is established by independent branches of the assembly tree. Furthermore, tree nodes with a large enough contribution block can be updated in parallel by splitting the update between several slaves of the master that is factorizing the block of fully summed variables. Finally, the root node can be treated in parallel if it is big enough.

While the *master* processor of each node in the tree (i.e. the one that is responsible for the factorization of the block of fully summed variables) is chosen during the analysis phase, the *slaves* for the parallel update of large contribution blocks are only chosen during the factorization phase. This dynamic task scheduling takes place in order to balance the work load of the processors at run-time. Problems arise from offering too much freedom to the dynamic scheduling. If every processor

is a candidate for a slave then, on each processor, enough workspace has to be reserved during the analysis phase for the corresponding computational tasks. However, during the factorization, typically not all processors are actually needed as slaves (and, on a large number of processors, only a very few are needed), so the prediction of the required workspace will be overestimated. Thus the size of the problems that can be solved is reduced unnecessarily because of this difference between the prediction and allocation of memory in the analysis phase and the memory actually used during the factorization.

With the concept of *candidate processors* it is possible to address this issue. The concept originates in an algorithm presented in [3] and extends efficiently to MUMPS. For each node that requires slaves to be chosen dynamically during the factorization because of the size of its contribution block, we introduce a limited set of processors from which the slaves can be selected. While the master previously chose slaves from among all less loaded processors, the freedom of the dynamic scheduling can be reduced so that the slaves are only chosen from the candidates. This effectively allows us to exclude all non-candidates from the estimation of workspace during the analysis phase and leads to a more realistic prediction of workspace needed. Furthermore, the candidate concept allows us to better structure the computation since we can explicitly restrict the choice of the slaves to a certain group of processors and enforce a ‘subtree-to-subcube’ mapping principle.

Preliminary tests with a prototype version have shown the benefits of the concept that is currently being integrated into a compact scheduling module for MUMPS. It unifies static and dynamic mapping while at the same time taking account of tree modifications by node amalgamation and splitting.

- 
- [1] P. R. Amestoy, I. S. Duff, J.-Y. L’Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Analysis and Applications*, 23(1):15–41, 2001.
  - [2] P. R. Amestoy, I. S. Duff, J.-Y. L’Excellent, and X. S. Li. Analysis, tuning and comparison of two general sparse solvers for distributed memory computers. Technical Report TR/PA/00/72, CERFACS, 2000. To appear in *ACM Trans. Math. Softw.*
  - [3] A. Pothén and C. Sun. A Mapping Algorithm for Parallel Sparse Cholesky Factorization. 14(5):1253–1257, 1993.

## 2.5 An analysis of MPI send/receive in the context of MUMPS

**P. R. Amestoy:** ENSEEIHT, *France*; **I. S. Duff:** CERFACS and Rutherford Appleton Laboratory, *England*; **J. Y. L’Excellent:** INRIA-ENSL, *Lyon, France*; **X. S. Li:** NERSC-LBNL, *Berkeley, California, USA*

This work was developed from the research performed as part of the France-Berkeley project [2] and is intimately associated with the tuning of the MUMPS and SuperLU sparse direct solvers on distributed memory computers using MPI for message passing.

We examined the send and receive mechanisms of MPI in detail and considered how to implement message passing robustly so that performance is not significantly affected by changes to the MPI system. We discussed this within the context of two different parallel algorithms for sparse Gaussian elimination: a multifrontal solver (MUMPS), and a supernodal one (SuperLU). The performance of our initial strategies based on simple MPI point-to-point communication primitives is very sensitive to the MPI system, particularly the way MPI buffers are used. Using more sophisticated non-blocking communication primitives improves the performance robustness and scalability, but at the cost of increased code complexity.

We have submitted our report [1] to the journal *Parallel Computing*.

- 
- [1] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and X. S. Li. Impact of the implementation of MPI point-to-point communications on the performance of two general sparse solvers. Technical Report RT/APO/01/4, ENSEEIHT-IRIT, Oct 2001.
  - [2] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and Xiaoye S. Li. Analysis, tuning and comparison of two general sparse solvers for distributed memory computers. Technical Report TR/PA/00/72, CERFACS, Toulouse, France, October 2000. [Compressed PS](#).

## 2.6 Use of orderings for large entries on the diagonal

**I. S. Duff:** CERFACS and Rutherford Appleton Laboratory, England; **J. Koster:** PARALLAB, Bergen, Norway

Although our original work on this set of algorithms for permuting large entries to the diagonal was completed some time ago and the resulting code **MC64** was placed in the HSL mathematical software library in 2000, the most recent journal paper on this work appeared this year [3] and the algorithm and code has been used extensively over this past year by many people in the solution of large sparse systems of equations and in preconditioning techniques for sparse matrices.

In many cases, the routine can be very helpful when factorizing very unsymmetric systems as has been our own experience with **MA41** and **MUMPS** [1]. It has also proven almost a necessary preprocessor for algorithms that use static pivoting strategies as in **SuperLU** [5]. More recently, Gupta [4] has reaffirmed the importance of such a permutation in his work on sparse direct solvers and Benzi and his colleagues [2] have found that it is absolutely necessary to use **MC64** if preconditioning techniques are to be successful on highly indefinite and nonsymmetric matrices.

- 
- [1] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Analysis and Applications*, 23(1):15–41, 2001.
  - [2] M. Benzi, J. C. Haws, and M. Tũma. Preconditioning highly indefinite and nonsymmetric matrices. *SIAM J. Scientific Computing*, 22(4):1333–1353, 2000.
  - [3] I. S. Duff and J. Koster. On algorithms for permuting large entries to the diagonal of a sparse matrix. *SIAM J. Matrix Analysis and Applications*, 22(4):973–996, 2001.
  - [4] A. Gupta. Improved symbolic and numerical factorization algorithms for unsymmetric sparse matrices. Technical Report RC 22137 (99131), IBM T.J. Watson Research Center, Yorktown Heights, NY, August 2001.
  - [5] X. S. Li and J. W. Demmel. A scalable sparse direct solver using static pivoting. In *Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing*, San Antonio, Texas, March 22–24 1999.

## 2.7 The sparse BLAS

**I. S. Duff:** CERFACS and Rutherford Appleton Laboratory, England; **M. A. Heroux:** Sandia National Laboratories, Albuquerque, New Mexico, USA; **R. Pozo:** NIST, Gaithersburg, Maryland, USA

The effort by the BLAS Technical Forum to produce an updated version of the Basic Linear Algebra Subprograms has continued for many years and it was with some relief that everything came to a conclusion in 2001. The standard itself [1] should be appearing as two parts in the journal *International Journal of High Performance Computing Applications* and an article on the standard [2] has been submitted to *ACM Transactions on Mathematical Software*. Major extensions to the earlier BLAS include added functionality, mixed precision BLAS, and sparse BLAS. It is this latter area with which we have been primarily involved.

We discuss the interface design for the Sparse Basic Linear Algebra Subprograms (BLAS), the kernels in the recent standard from the BLAS Technical Forum that are concerned with unstructured sparse matrices. The motivation for such a standard is to encourage portable programming while allowing for library-specific optimizations. In particular, we show how this interface can shield one from concern over the specific storage scheme for the sparse matrix. This design makes it easy to add further functionality to the sparse BLAS in the future.

We illustrate, in our report [3], the use of the Sparse BLAS with examples in the three supported programming languages, Fortran 95, Fortran 77, and C. A Fortran 95 implementation of the sparse BLAS is described in the report [4] and discussed in Section 2.3.

- 
- [1] BLAS Technical Forum Standard. *The International Journal of High Performance Computing Applications*, 15(3–4), 2001.
  - [2] S. Blackford, J. Demmel, J. Dongarra, I. Duff, S. Hammarling, G. Henry, M. Heroux, L. Kaufman, A. Lumsdaine, A. Petitet, R. Pozo, K. Remington, and R. C. Whaley. An updated set of basic linear algebra subprograms (BLAS). *Submitted to ACM Trans. Math. Softw.*, 2002.
  - [3] I. S. Duff, M. A. Heroux, and R. Pozo. The Sparse BLAS. Technical Report TR/PA/01/24, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
  - [4] I. S. Duff and C. Vömel. The Implementation of the Sparse BLAS in Fortran 95. Technical Report TR/PA/01/27, CERFACS, Toulouse, France, 2001. Submitted to *ACM Trans. Math. Softw.* [Compressed PS](#), [PDF](#).

## 2.8 Activities of a senior visitor

**J. R. Gilbert:** XEROX PARC, *California, USA*

### Structure prediction and symbolic factorization for partial pivoting

The answers to several theoretical questions about nonzero structure and column dependencies in sparse  $LU$  factorization with partial pivoting could have consequences in the analysis phase of  $LU$  factorization codes. I have been working on some of them with Laura Grigori, a graduate student at LORIA in Nancy.

Laura has a symbolic all-at-once lower bound for the structure of the column elimination tree that we have recently described in a technical report [1] that we have submitted for publication. This lower bound applies to the case of so-called strong Hall matrices.

We have also proved a symbolic lower bound result for the structure of  $L$  and  $U$  under partial pivoting for all nonsingular matrices. This result was already known for the case of strong Hall matrices. However, the general case is of interest because many nonsymmetric  $LU$  codes do not preprocess the matrix to block triangular form. I hope that it will be possible to extend this work to get an exact lower bound; if we can push it through and also get an efficient way to implement it, we plan to experiment with using it as an improvement to the current structure prediction in SuperLU. We would also like to extend the column etree result from a symbolic bound to an exact bound.

During my time at CERFACS, I made a visit to LORIA and Laura made a visit to CERFACS to collaborate on this work. In December 2001, I served as a member of the jury during Laura's successful defence of her Ph.D. thesis at Nancy [2]; she is now a postdoc at NERSC in Berkeley, California.

### Weighted matching for symmetric indefinite factorization

The objective here is to use weighted bipartite matching (as in Duff and Koster's MC64 code) to identify good static (or at least initial) pivoting orders for symmetric indefinite factorization. Such orderings might also be useful for incomplete symmetric indefinite factorization or for the iterative solution of symmetric indefinite systems.

A maximum weight matching may not be symmetric, even for a symmetric matrix. One may ask to what extent one can capture the elements of such a matching in small symmetric block pivots by permuting the matrix symmetrically. The answer turns out to depend on the cycle structure of the matching permutation.

Iain Duff and I have started to experiment with the idea of first permuting the matrix symmetrically to have small diagonal blocks with large elements, then following this with another symmetric permutation that is chosen for sparsity, then using this as input to MA57, a symmetric indefinite multifrontal factorization. The sparsity permutation, approximate minimum degree, is computed on the quotient graph by the diagonal blocks, in order not to break up the block pivots. We plan to report on the results of these experiments during 2002.

### Support theory for preconditioning

Support graphs and support theory are a combinatorial approach to preconditioning sparse linear systems that date from the early 1990s but have recently seen some promising new developments. I gave a survey lecture on the topic, covering both our work and others', at the Sparse Day at CERFACS in June, with the intention of getting more researchers interested in working on it. Trond Steihaug of the University of Bergen pointed out a connection to a spanning tree preconditioner for network problems that has been studied by Resende and others in the optimization community; we hope to explore this in more detail. The Sparse Day talk also led to a visit to Jean Roman's group at LABRI in Bordeaux: we plan to collaborate on these and related ideas in the context of the block incomplete preconditioners that Jean's group is designing.

I hope to organize a workshop on combinatorial preconditioning sometime during 2002, perhaps at CERFACS, during which the sparse linear systems community can explore these and related problems.

### Matlab toolboxes

I updated the Matlab mesh partitioning toolbox, which Shang-Hua Teng and I wrote several years ago and which Tim Davis had most recently updated to run under version 5 of Matlab. The newly updated toolbox includes an interface to METIS and runs under Matlab 6. The updated toolbox is now available at <http://www.cerfacs.fr/algor/Softs/MESHPART/index.html>.

Gary Howell (of the Florida Institute of Technology, visiting CERFACS) and I experimented successfully with using f2c as an intermediary in connecting Fortran sparse matrix codes to Matlab. I also wrote a routine to read sparse matrices in Rutherford-Boeing format into Matlab. The routine is not yet comprehensive (it currently handles six of the R-B matrix types), but it demonstrates that using the new format with Matlab will be considerably simpler than using the older Harwell-Boeing format. Student Pierre Valentin helped a good deal with sample matrices from his summer work on the Rutherford-Boeing collection.

- 
- [1] J. R. Gilbert and L. Grigori. A note on the column elimination tree. Technical Report TR/PA/01/104, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
  - [2] L. Grigori. *Prédiction de structure et algorithmique parallèle pour la factorisation LU des matrices creuses*. PhD thesis, Université Henri Poincaré, 2001.

## 3 Iterative Methods and Preconditioning

### 3.1 When modified Gram-Schmidt generates a well conditioned set of vectors

L. Giraud: CERFACS, France; J. Langou: CERFACS, France

The modified Gram-Schmidt (MGS) algorithm is a well known orthogonalization method used in many linear algebra algorithms. For example, It is a fast way to compute a **QR** factorization. Although this method is widely used, some theoretical questions on its behaviour in finite precision arithmetic are still open.

It is well known that the set of computed vectors,  $\bar{\mathbf{Q}}$ , may lose orthogonality; however, it is experimentally observed that this set has full rank. In [1], we propose a theorem that gives a theoretical explanation of this phenomenon. This theorem applies to a large set of matrices. This is the set of not “too ill-conditioned” matrices defined by the relation  $c\mu\kappa < 0.1$  where  $c$  is a constant depending on  $m$  and  $n$ ,  $\kappa$  the condition number of the initial matrix and  $u$  the machine precision.

To complete this theorem, we exhibit a  $3 \times 3$  matrix, called *CERFACS*, that does not belong to this set but is indeed very close  $c\mu\kappa \sim 0.1$ . On that matrix, MGS generates a numerically rank deficient set of vectors  $\bar{\mathbf{Q}}$ . This justifies the sharpness of the domain of validity of our theorem. Finally, if we combine the result of our work with a well known result from Björck, we can show that two sweeps of MGS are indeed enough to get a matrix whose columns are orthogonal up to machine precision. We recover the famous sentence of W. Kahan [2] : “*twice is enough.*”

- 
- [1] L. Giraud and J. Langou. When modified Gram-Schmidt generates a well conditioned set of vectors. Technical Report TR/PA/01/17, CERFACS, Toulouse, France, 2001. To appear in IMAJNA. [Compressed PS](#), [PDF](#).
  - [2] B. Parlett. *The symmetric eigenvalue problem*. Prentice Hall, 1980.

### 3.2 Real/complex seed GMRES and block GMRES packages for sequential and parallel computation

L. Giraud: CERFACS, France; J. Langou: CERFACS, France

The solution of dense complex non-Hermitian linear systems arises in computational electromagnetics. In this context, iterative methods are appealing but with several right-hand sides iterative methods may become less attractive than a single factorization (direct method). A solution is to adapt iterative methods for multiple right-hand sides.

We have developed and interfaced with the EADS code a block GMRES solver and a seed GMRES solver. To extend the range of applicability, they are based on reverse communication and are available in real and complex, single and double arithmetics.

We have already some promising preliminary experiments observed on the test case CETAF with 5391 degrees of freedom and nine right-hand sides. Each of the right hand sides correspond to a different angle for the wave illuminating the object. In that experiment, the difference between two successive angles is one degree. The seed GMRES method converges in 226 iterations against 718 accumulated iterations for GMRES applied successively to each of the nine right-hand sides.

### 3.3 Sparse symmetric preconditioners for dense linear systems in electromagnetism

**B. Carpentieri:** CERFACS, *France*; **I. S. Duff:** CERFACS, *France*; **L. Giraud:** CERFACS, *France*;  
**M. Magolu monga Made:** *Université Libre de Bruxelles, Belgium*

In order to further develop the study presented in [1], we consider symmetric preconditioning strategies for the iterative solution of dense complex symmetric non-Hermitian systems arising in computational electromagnetics. In particular, we report on the numerical behaviour of the classical Incomplete Cholesky factorization as well as some of its recent variants and consider also well known factorized approximate inverses. We illustrate the difficulties that these techniques encounter on the linear systems under consideration and give some clues to explain their disappointing behaviour. We propose two symmetric preconditioners based on Frobenius-norm minimization that use a prescribed sparsity pattern. The numerical and computational efficiency of the proposed preconditioners are illustrated on a set of model problems arising both from academic and from industrial applications. More details on this work are available in [2].

- 
- [1] B. Carpentieri, I. S. Duff, and L. Giraud. Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism. *Numerical Linear Algebra with Applications*, 7(7-8):667–685, 2000.
  - [2] B. Carpentieri, I.S. Duff, L. Giraud, and M. Magolu monga Made. Sparse symmetric preconditioners for dense linear systems in electromagnetism. Technical Report TR/PA/01/35, CERFACS, Toulouse, France, 2001. Submitted. [Compressed PS](#), [PDF](#).

### 3.4 Combining fast multipole techniques and approximate inverse preconditioners for large calculations in electromagnetism

**B. Carpentieri:** CERFACS, *France*; **I. S. Duff:** CERFACS, *France*; **L. Giraud:** CERFACS, *France*;  
**G. Sylvand:** CERMICS-INRIA, *Sophia Antipolis, France*

For large electromagnetic calculations that involve several tens of thousands or a few million unknowns, the use of fast multipole techniques is mandatory to evaluate the matrix-vector product. For such simulations, we first investigate the numerical scalability of the approximate inverse preconditioner [1] implemented in a parallel distributed code [2]. Because the preconditioner naturally becomes more local when the size of the problem is increased, even though the Green functions decay rapidly, we observe that the convergence rate deteriorates when the size of the linear system increases. To overcome this drawback and improve the numerical robustness of the linear solver we propose an embedded scheme; it consists of a FGMRES Krylov solver for the outer iterations and a preconditioned GMRES inner scheme. For this outer solver, we use an accurate fast multipole calculation for the matrix-vector evaluation, preconditioned by a few inner preconditioned GMRES iterations. For the inner GMRES scheme, we use a less accurate fast multipole calculation for the matrix-vector computation and our Frobenius norm minimization approach as preconditioner. The efficiency of this numerical scheme is demonstrated on large test problems. The benefit of the new scheme is highlighted in Table 3, where we display the number of outer and inner fast multipole matrix-vector products as well as the elapsed time to solve a problem with around one million unknowns arising from a simulation on an Airbus aircraft. We consider GMRES(30) and FGMRES(5)/GMRES(20) because both use the same amount of memory. The target computer is a Compaq SC Alpha server.

Sphere with 1.0 M degree of freedoms				
GMRES(30)		FGMRES(5) + GMRES(20)		
# accurate FMM	Elapsed time	# accurate FMM	# less accurate FMM	Elapsed time
1196	11 hours	17	260	1 hour 30 mn
Airbus aircraft with 1.1 M degree of freedoms				
GMRES(30)		FGMRES(5) + GMRES(20)		
# accurate FMM	Elapsed time	# accurate FMM	# less accurate FMM	Elapsed time
no convergence	–	19	300	4 hour 20 mn

Table 3 : Numerical behaviour observed on a 16 processor Alpha Compaq Server.

- 
- [1] B. Carpentieri, I. S. Duff, and L. Giraud. Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism. *Numerical Linear Algebra with Applications*, 7(7-8):667–685, 2000.
- [2] G. Sylvand. *Résolution Itérative de Formulation Intégrale pour Helmholtz 3D : Applications de la Méthode Multipôle à des Problèmes de Grande Taille*. PhD thesis, Ecole Nationale des Ponts et Chaussées, 2002.

### 3.5 Grid transfer operators for highly variable coefficient problems in two-level non-overlapping domain decomposition methods

**L. Giraud:** CERFACS, France; **F. Guevara Vasquez:** Rice University, Houston, USA;  
**R. S. Tuminaro:** Sandia National Laboratories, Albuquerque, New Mexico, USA

We propose a robust interpolation scheme for non-overlapping two-level domain decomposition methods applied to two-dimensional elliptic problems with discontinuous coefficients. This interpolation is used to design a preconditioner closely related to the BPS scheme proposed in [1]. The definition of this interpolation is natural on structured meshes with uniform rectangular subdomains and we propose a generalization to unstructured meshes. This generalization preserves the constant function while taking into account possible discontinuities. The unstructured grid interpolation was inspired by AMGe [2] and also by the fact that it reduces to the original operator dependent interpolation on uniform meshes. Through numerical experiments, we show on structured and unstructured finite-element problems that the new preconditioning scheme reduces to the BPS method on smooth problems but outperforms it on problems with discontinuous coefficients. In particular it maintains good scalable convergence behaviour even when the jumps in the coefficients are not aligned with subdomain interfaces. This work has been presented during the Preconditioning 2001 Conference, and was completed during a visit to Sandia National Laboratory, Livermore after the conference. Results related to this work are reported in [3]

- 
- [1] J. H. Bramble, J. E. Pasciak, and A. H. Schatz. The construction of preconditioners for elliptic problems by substructuring I. *Math. Comp.*, 47(175):103–134, 1986.
- [2] M. Brezina, A. Cleary, R. Falgout, V. Henson, J. Jones, T. Manteuffel, S. McCormick, and J. Ruge. Algebraic multigrid based on element interpolation (AMGe). *SIAM J. Sci. Comput.*, 22(5):1570–1592, 2000.
- [3] L. Giraud, F. Guevara Vasquez, and R. S. Tuminaro. Grid transfer operators for highly variable coefficient problems in two-level non-overlapping domain decomposition methods. Tech. Rep. TR/PA/01/03, CERFACS, Toulouse, France, 2001. Submitted. [Compressed PS](#), [PDF](#).

## 3.6 Domain decomposition methods in semiconductor device modelling

**L. Giraud:** CERFACS, *France*; **J. Koster:** Parallab, *Norway*; **J.-C. Rioual:** CERFACS, *France*;  
**A. Marrocco:** INRIA, *France*

In the framework of a joint research project between INRIA (A. Marrocco), Parallab (J. Koster) and CERFACS (L. Giraud, J. C. Rioual) we are developing a parallel domain decomposition code for the solution of the drift diffusion equation involved in semiconductor device modelling. A nonlinear time dependent problem has to be solved where each nonlinear iteration requires the solution of a linear mixed finite-element problem resulting in a large sparse linear system.

In order to solve these linear systems in a parallel distributed memory environment using message passing, we investigate direct and iterative substructuring techniques. These techniques imply the computation of the local Schur complement matrices associated with each subdomain. If we use a classical sparse direct solver, the computation of the local Schur complement matrices is usually costly in computation time. We used some functionalities of the multifrontal parallel solver MUMPS [1] to compute them efficiently. Having an explicit distributed formulation of the Schur complement we can implement substructuring methods, direct or iterative. Direct substructuring is efficient and numerically stable. Iterative substructuring requires a good preconditioner for the Schur complement system. In collaboration with Parallab from Bergen University, we have tested the Parallab implementation [2] of a Balanced Neumann-Neumann preconditioner [2, 6] and compared it with another two-level preconditioner designed at CERFACS [3, 4]. We also investigate various scaling techniques on the Schur complement to improve the numerical stability of the method. The joint work with Parallel was partially funded by an Egide Aurora grant, that enabled J.-C. Rioual to visit Bergen and reciprocally J. Koster to come to at CERFACS. This work was presented during the 13<sup>th</sup> conference on Domain Decomposition Methods in Scientific Computing and more details are available in [5, 7].

- 
- [1] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Analysis and Applications*, 23(1):15–41, 2001.
  - [2] P. E. Bjørstad, J. Koster, and P. Krzyżanowski. Domain decomposition solvers for large scale industrial finite element problems. In *PARA2000 Workshop on Applied Parallel Computing*, volume 1947, pages 373 – 383. Lecture Notes in Computer Science, Springer-Verlag, 2000.
  - [3] L. M. Carvalho, L. Giraud, and G. Meurant. Local preconditioners for two-level non-overlapping domain decomposition methods. *Numerical Linear Algebra with Applications*, 8(4):207–227, 2001.
  - [4] L. M. Carvalho, L. Giraud, and P. Le Tallec. Algebraic two-level preconditioners for the Schur complement method. *SIAM J. Scientific Computing*, 22(6):1987 – 2005, 2001.
  - [5] L. Giraud, J. Koster, A. Marrocco, and J.-C. Rioual. Domain decomposition methods in semiconductor device modeling. Technical Report TR/PA/01/51, CERFACS, Toulouse, France, 2001. To appear in the proceedings of the 13<sup>th</sup> conference on Domain Decomposition Methods in Scientific Computing, 2000. [Compressed PS](#), [PDF](#).
  - [6] J. Mandel. Balancing domain decomposition. *Comm. Numer. Meth. Engrg.*, 9:233–241, 1993.
  - [7] J.-C. Rioual. *Solving linear systems in semiconductor device modeling on parallel distributed computers*. PhD thesis, CERFACS, Toulouse, France, 2002.

### 3.7 Spectral two-level preconditioners

**B. Carpentieri:** CERFACS, *France*; **I. S. Duff:** CERFACS, *France*; **L. Giraud:** CERFACS, *France*;  
**J.-C. Rioual:** CERFACS, *France*

When solving the left preconditioned linear system  $M_1Ax = M_1b$  with a Krylov method, the smallest eigenvalues of  $M_1A$  often slow down the convergence. In the symmetric positive definite case this situation is well-understood and arguments exist for unsymmetric systems to explain the bad effect of the smallest eigenvalues on the rate of convergence of the unsymmetric Krylov solver. We propose a class of spectral two-level preconditioners based on a low rank update that aims at shifting close to one these smallest eigenvalues of  $M_1A$ . Consequently the resulting two-level preconditioner does not suffer anymore from the effect of those small eigenvalues. Our technique requires the explicit computation of a few eigenvalues that makes it independent from the Krylov solver being used. Symmetric and symmetric positive definite variants can be derived for symmetric and symmetric positive definite linear systems. In that latter situation, the resulting preconditioner is similar to those proposed in [2] in the framework of domain decomposition for elliptic equations where the shape of the smallest eigenvectors might be *a priori* approximated. The effectiveness of the new preconditioners is demonstrated on symmetric non-Hermitian problems arising from electromagnetism [1] (similar problems to those considered in Section 3.3) and on symmetric positive definite and unsymmetric linear systems arising in domain decomposition for the simulation of semiconductor devices [3] (similar problems to those considered in Section 3.6). Although described for left preconditioners in this short section, this spectral two-level technique applies also for right preconditioners.

- 
- [1] B. Carpentieri. *Sparse preconditioners for dense linear systems in electromagnetic applications*. PhD thesis, CERFACS, Toulouse, France, 2002.
  - [2] L. M. Carvalho, L. Giraud, and P. Le Tallec. Algebraic two-level preconditioners for the Schur complement method. *SIAM J. Scientific Computing*, 22(6):1987 – 2005, 2001.
  - [3] J.-C. Rioual. *Solving linear systems in semiconductor device modeling on parallel distributed computers*. PhD thesis, CERFACS, Toulouse, France, 2002.

### 3.8 Combining OpenMP and MPI on clusters of symmetric multiprocessors: some basic promising experiments

**L. Giraud:** CERFACS, *France*

We present some experiments on different clusters of SMPs, where both distributed and shared memory parallel programming paradigms can be naturally combined. Although the platforms exhibit the same macroscopic memory organization, it appears that their individual overall performance is closely dependent on the ability of their hardware to efficiently exploit the local shared memory within the nodes. In that context, the cache blocking strategy appears to be very important not only to get good performance out of each individual processor but mainly to get good performance out of the overall computing node since sharing memory locally might become a severe bottleneck. On a simple benchmark, representative of many large simulation codes, we show through numerical experiments that mixing the two programming models enables to get attractive speed ups that compete with a pure distributed memory approach. There are many possible fields of application for this embedded parallelism including the development of numerical libraries [1]. This can be also of central interest for large industrial codes that have been developed for years on vector machines and then ported, with a significant manpower effort, on parallel distributed vector computers using message passing. In many cases those codes are efficient only on a moderate number of processors as they were not initially designed for parallel computers. Using OpenMP to parallelize most of the vectorial loops, in combination with MPI, might be a viable opportunity for smoothly moving these codes onto the emerging and promising platforms for intensive scientific computing that are the clusters of SMPs. Another illustration is from numerical linear algebra where we can mention

domain decomposition or more generally block preconditioning techniques for the solution of large sparse linear systems. For those techniques, the numerical scalability is often related to the number of blocks or subdomains. In classical parallel distributed implementations, one assigns one block per processor. Consequently increasing the number of processors for solving a given problem results in a less efficient numerical solver and then does not fully take advantage of the computing power of all the processors. Combining the two programming models enables us to efficiently exploit some parallelism at a block level, through the use of parallel direct solvers for shared memory for instance. In that context, the number of processors used to perform a given simulation can be increased without deteriorating the numerical property of the numerical algorithm since the number of blocks does not need to be increased. For more details on that work we refer to [2].

- 
- [1] J. J. Dongarra, S. Moore, and A. Trefethen. Numerical libraries and tools for scalable parallel cluster computing. *Int J. of High Performance Computing Applications*, 15(2):175–180, 2001.
  - [2] L. Giraud. Combining shared and distributed memory programming models on clusters of symmetric multiprocessors: Some basic promising experiments. Working Note WN/PA/01/19, CERFACS, Toulouse, France, 2001.

---

## 4 Qualitative Computing

---

*Group members: Françoise Chaitin-Chatelin, Claire Mandry, Tadas Meškauskas, Laurent Plantié and Elisabeth Travesias.*

The work of the Qualitative Computing Group is a collaborative effort to assess the validity of computer simulations. The central question is to give meaning to computer results which are seemingly wrong such as in chaotic computations. This goal can be reached by uncovering the laws of computation which govern finite precision computations in the neighbourhood of singularities. Some of these laws are now well understood. For example, one can cite i) the role of the normwise backward error to assess the reliability of numerical software in finite precision, ii) the role of nonnormality which makes approximated singularities appear much closer than they are in exact arithmetic.

A number of new laws have emerged more recently, which have been the focus of the Group's attention in recent years. The new laws concern in particular:

- a) *inexact computing* and the associated homotopic pseudospectrum and backward error as a fruitful framework to understand approximate numerical methods, in exact arithmetic,
- b) the unreasonable robustness of Krylov-type methods to perturbations in the data,
- c) the (underestimated) role of Geometry in Scientific Computing.

Our research and understanding is vitally nourished by work on practical numerical software applications in Physics and Technology, which come from CERFACS partners. We review below the work accomplished in 2001 in 4 related areas.

### 4.1 Inner-Outer iterations

It is well known that *asymptotic* methods, such as Newton-like methods for systems and the power method for eigenvalues, need to be performed with more and more accuracy as one gets closer to the solution.

On the contrary, *direct* methods such as Krylov-type methods exhibit a remarkable robustness to perturbations in the data. The only requirement is that, when the incomplete Krylov method is used *iteratively*, as in the *restarted* versions, the first steps in each new iteration loop are computed to full working accuracy, which can be later relaxed as the convergence proceeds inside the loop [1]. This amazing property results in substantial savings in the overall cost of running a 2 level-solver, when the *outer* iteration is of Krylov type.

Although a theoretical proof is still lacking, we have more and more empirical evidence, confirmed by similar results from researchers worldwide.

T. Meškauskas has worked on inner-outer iterations for mode solvers in structural mechanics in research supported by CERFACS industrial partner EDF (Electricité de France). He investigated an implicitly restarted Arnoldi method with shifts, coupled with a preconditioned conjugate gradient linear solver. This algorithm is applied to the generalized eigenvalue problem and fits into the general framework of inner-outer iterations when one iterative solver, referred to as the inner solver, is embedded into another one, referred to as the outer solver. In this case, each outer step of the eigensolver (an implicitly restarted Arnoldi method with shifts) requires the solution of a linear

system, provided by an iterative linear solver (preconditioned conjugate gradient), which is the inner iteration. As expected, the accuracy of the inner iteration can be *relaxed* when the outer process comes closer to the solution [2]. The *relaxation strategy* was implemented in the `Code_Aster` code and tested on large industrial problems.

- 
- [1] A. Bouras, F. Chaitin-Chatelin, and V. Frayssé. Solveurs itératifs imbriqués. Contract Report FR/PA/00/19, CERFACS, Toulouse, France, 2000.
  - [2] F. Chaitin-Chatelin and T. Meškauskas. Inner-outer iterations for mode solvers in structural mechanics: application to the Code\_Aster. Contract Report CR/PA/01/85, CERFACS, 2001.

## 4.2 Arnoldi method and the happy breakdown

A convincing theory of the convergence of an iterative Arnoldi method is still a tantalizing goal for numerical analysts in Linear Algebra. We approach this question by considering the incomplete Arnoldi method on a matrix of order  $n$  as a *finite*, or *direct* (as opposed to iterative or asymptotic) method. In this approach, it is of prime importance to understand the relation between the starting vector  $v_1$  and the happy breakdown at step  $k_0 \ll n$ . This requires, in particular, eliminating the conventional assumption that the matrix is non-derogatory. Because, if the matrix is non-derogatory, no early happy breakdown can occur unless the starting vector lies exactly in an invariant subspace in exact arithmetic.

In finite precision computations, it is well known that detecting the exact happy breakdown is difficult. Using the relation between  $v_1$  and  $k_0$ , a heuristic way to detect it is proposed in [6]. At the first SMAI Congress, this idea was presented in a poster, which was a finalist in best poster competition.

The quality of the eigenvalues computed by the Arnoldi method has been a topic of study for more than four years in the Qualitative Computing Group [1, 3, 6, 8]. These studies show that, for some initial vectors, it is not desirable to stop the Arnoldi algorithm in finite precision at the same time as in exact arithmetic. These particular initial vectors have certain components in the eigenvector basis which are too small to be considered in finite precision; only the components which are sufficiently large with respect to machine precision can be considered. So the notion of a neighbourhood of a happy breakdown is introduced.

A comparison between results from homotopic perturbations and results using the toolbox PRECISE [2, 7] and [5] has been performed.

At the beginning of 2001, with S. Gratton from CNES, a condition number formula was established, which expresses the sensitivity to the initial vector for the Arnoldi process [4]. In collaboration with Prof. B. N. Parlett from Berkeley University, California (USA), we try to apply it to determine the number of distinct eigenvalues of a symmetric matrix using the Lanczos process. Such a problem was posed by Prof. Lax (Courant Institute, New York).

- 
- [1] F. Chaitin-Chatelin. Comprendre les méthodes de Krylov en précision finie : le programme du Groupe Qualitative Computing au CERFACS. Technical Report TR/PA/00/11, CERFACS, 2000. [Compressed PS](#).
  - [2] F. Chaitin-Chatelin and E. Traviasas. PRECISE, a toolbox for assessing the quality of numerical methods and software. Technical Report TR/PA/00/12, CERFACS, Toulouse, France, 2000. Proceedings of the 16th IMACS World Congress, Lausanne, Switzerland. [Compressed PS](#).
  - [3] F. Chaitin-Chatelin, E. Traviasas, and L. Plantié. Understanding Krylov methods in finite precision. Technical Report TR/PA/00/40, 2000. Also in Proceedings of the Second Conference on Numerical Analysis and Applications, Rousse, Bulgaria, Springer, pp. 187–197. [Compressed PS](#).
  - [4] E. Traviasas F. Chaitin-Chatelin and S. Gratton. Sensitivity from the initial vector for Krylov based methods, 2001. Work in progress.
  - [5] A. Ilahi. *Validation du calcul sur ordinateur: application de la théorie des singularités algébriques*. Ph.D. dissertation, Université Toulouse I, June 1998. TH/PA/98/31.
  - [6] C. Mandry and E. Traviasas. Convergence de la méthode d’Arnoldi en précision finie en fonction du vecteur initial. Working Note WN/PA/01/36, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
  - [7] R. A. McCoy and V. Toumazou. PRECISE User’s Guide - Version 1.0. Technical Report TR/PA/97/38, CERFACS, Toulouse, France, 1997. [Compressed PS](#).
  - [8] E. Traviasas. *Sur le déploiement du champ spectral d’une matrice*. Ph.D. dissertation, Université Toulouse I, May 2000. TH/PA/00/30. [Compressed PS](#), [PDF](#).

### 4.3 Inexact computing

Inexact Computing can be described as follows: we are given two matrices  $A$  and  $E$  such that  $B = A + E$ , and we want to solve  $(B - zI)u = y$ . The rule of the game is that we are allowed only to use  $E$  and the resolvent field  $z \rightarrow (A - zI)^{-1}$  to solve the system.

To play such a game, it is useful to introduce the homotopic family  $A(t) = A + tE$ ,  $t \in \mathbb{C}$ , such that  $A(0) = A$  and  $A(1) = B$ . This describes in its simplest form the homotopic perturbation theory which has been studied and put to use by F. Chaitin-Chatelin in various guises for more than 30 years [1, 2, 3].

We interpret now the homotopic perturbation theory from a purely *information theoretic* point of view. This allows us to:

- i) interpret the computing activity in terms of knowledge acquisition,
- ii) structure the field of singularities of  $(A(t) - zI)^{-1}$  with two families of curves: the singular rays and the singular orbits,
- iii) give a *global geometric view* of the resolvent field for the family  $A(t) = A + tE$ , when there is **no restriction** on  $E$ . In particular  $\|E\|$  can be arbitrarily large. This is a big step forward since all previously known perturbation theories were local [1].

Such an analysis provides an efficient conceptual tool to understand how well separated eigenvalues can be grouped by the deviation matrix  $E$  induced by a numerical method, in exact arithmetic.

It also puts into fuller light the difference between  $E$  of rank 1 and  $E$  of rank greater than 1 [5]. In the second case, eigenvalues can go out of phase in a striking manner [4].

- 
- [1] F. Chaitin-Chatelin. *Spectral approximation of linear operators*. Academic Press, New York, 1983.
  - [2] F. Chaitin-Chatelin. *Valeurs propres de matrices*. Masson, Paris, 1988.
  - [3] F. Chaitin-Chatelin and V. Frayssé. *Lectures on Finite Precision Computations*. SIAM, Philadelphia, 1996.
  - [4] F. Chaitin-Chatelin and E. Traviasas. Homotopic perturbation - Unfolding the field of singularities of a matrix by a complex parameter: a global geometric approach. Technical Report TR/PA/01/84, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
  - [5] E. Traviasas. *Sur le déploiement du champ spectral d'une matrice*. Ph.D. dissertation, Université Toulouse I, May 2000. TH/PA/00/30. [Compressed PS](#), [PDF](#).

## 4.4 Geometric aspects of computing

In conventional scientific computing the four arithmetic operations (+, −, ×, /) are usually defined on scalars or real numbers, that is 1D vectors. Linear Algebra is performed on matrices and vectors of potentially very large dimension, but the numbers themselves are either real scalars (1 dimension) or complex scalars (2 dimensions). Hypercomplex scalars are usually not used in Numerical Linear Algebra. However, it is known since the days of Hamilton and Graves that one can multiply and divide real vectors of dimension  $2^k$ ,  $k = 0, 1, 2, \dots$ , which define algebras of hypercomplex numbers. This requires introducing an additional geometric operation, the conjugation. These 5 operations define the arithmetic-geometric core of Nature's Computation. For example, the laws of Classical Mechanics and of Maxwell's electromagnetism can be written most efficiently by using products of quaternions, which are 4D vectors. The real component can be interpreted as Time and the 3 imaginary components as Space.

This suggests using a quaternion parameter  $t$  with 4 dimensions, rather than a complex one with only 2 dimensions, to extract meaning from the singularities of  $A(t) = A + tE$ . Such a line of thought, which complements the point of view developed in Section 4.3, is currently under development and is expected to be completed in 2002.

A similar approach can be taken with 8D-octonions. Preliminary discussions have taken place with the TTN Group, in order to use the potential of the Virtual Reality machine to visualize the dynamics of knowledge acquisition, when computation is analysed by means of a multidimensional parameter  $t$  with 4 or 8 D in which one dimension, the first, is "felt" but not seen, the next 3 dimensions are spatial, hence visible, and the remaining 4 are non visible. F. Chaitin-Chatelin has presented several invited talks on this topic, which are listed below [1, 2, 3, 4, 5, 6].

- 
- [1] F. Chaitin-Chatelin. Calcul qualitatif et sens de la vie. Ecole d'Ingénieurs CPE, Lyon, 10 May 2001, talk.
  - [2] F. Chaitin-Chatelin. Computing thoughts. Workshop on Numerical methods for evolutionary problems, Peschici, Italy, 17-21 Sept. 2001, talk.
  - [3] F. Chaitin-Chatelin. Hypercomplex computation. IBM, New York, USA, 27 Dec. 2001, talk.
  - [4] F. Chaitin-Chatelin. Life computation. The Institute of Ecotechnics Conference on Time: Metapatterns, The Present Moment and Evolution. Aix-en-Provence, 26-29 Oct. 2001, talk.
  - [5] F. Chaitin-Chatelin. Nature's computation: a theory of qualitative computing. CERFACS, Toulouse, France, 30 March 2001, talk.
  - [6] F. Chaitin-Chatelin. Qualitative computing. IFIP WG 2.5 Meeting, Amsterdam, the Netherlands, 26-27 May 2001, talk.

## 5 Nonlinear Systems and Optimization

### 5.1 SifDec: a lonesome SIF decoder

**N. I. M. Gould:** *Rutherford Appleton Laboratory, England*; **D. Orban:** CERFACS, *France*;  
**Ph. L. Toint:** *Facultés Universitaires Notre-Dame de la Paix, Belgium*.

The Constrained and Unconstrained Testing Environment (CUTE) [1], briefly described in Section 5.2, strongly relies on a device used to convert SIF-encoded optimization problems into a set of Fortran 77 subroutines, which can in turn be used by solvers to evaluate the objective function values, constraint functions values, compute their derivatives and so forth. This device—the SIF decoder—has been isolated and ported to a variety of popular platforms, for a variety of popular Fortran 77 and Fortran 90/95 compilers. The resulting package has been named SifDec.

Although the main usage of SifDec is presently in conjunction with CUTEr (see Section 5.2), the fact that it is now isolated opens new doors, such as the implementation of a converter from SIF input format to Ampl [2] input format.

The general SifDec documentation may be found in the report [4], the software is formally described in the report [3] and the official SifDec website is located at <http://cuter.rl.ac.uk/cuter-www/sifdec>.

- 
- [1] I. Bongartz, A. R. Conn, N. I. M. Gould, and Ph. L. Toint. CUTE: Constrained and Unconstrained Testing Environment. *Transactions of the ACM on Mathematical Software*, 21(1):123–160, 1995.
  - [2] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Scientific Press, 1993.
  - [3] N. I. M. Gould, D. Orban, and Ph. L. Toint. CUTEr, a constrained and unconstrained testing environment, revisited. Technical Report TR/PA/01/04, CERFACS, Toulouse, France, 2001.
  - [4] N. I. M. Gould, D. Orban, and Ph. L. Toint. General SifDec documentation. Technical Report TR/PA/02/14, CERFACS, Toulouse, France, 2002.

### 5.2 CUTEr: a Constrained and Unconstrained Testing Environment revisited

**N. I. M. Gould:** *Rutherford Appleton Laboratory, England*; **D. Orban:** CERFACS, *France*;  
**Ph. L. Toint:** *Facultés Universitaires Notre-Dame de la Paix, Belgium*.

The Constrained and Unconstrained Testing Environment (CUTE) [1] is a versatile environment for testing small to large-scale nonlinear programming problems arising from both real practical applications and from academic circles. It provides Fortran tools for computing function values, gradients, Hessians, matrix-vector products and handles both dense and sparse problems. It has been designed with multi-platform environments in mind and the test problems are written using the SIF (Standard Input Format) description language, formerly used by LANCELOT [2]. CUTE also provides tools to help the users build their own interface to their optimization package, as well as ready-to-use interfaces to famous existing packages like MINOS and OSL.

The purpose of this research is to polish CUTEr, the new version of CUTE. CUTEr is available on a variety of popular platforms, including the quickly growing Linux platforms, for a variety of popular Fortran 77 and Fortran 90/95 compilers and is suited to heterogeneous local networks. Its installation phase, driven by *Imakefiles* make it portable across platforms, is faster, easier, more

efficient and makes better use of disk space and memory as architecture-dependent parts have been carefully isolated. CUTer provides more tools, with enhanced capabilities that implement recent developments in dense and sparse linear algebra and interfaces to more recent optimization packages like KNITRO and filterSQP. CUTer relies on the SIF decoder, described in Section 5.1.

The general CUTer documentation may be found in the report [4], the software is formally described in the report [3] and the official CUTer website is located at <http://cuter.rl.ac.uk/cuter-www/>.

- 
- [1] I. Bongartz, A. R. Conn, N. I. M. Gould, and Ph. L. Toint. CUTE: Constrained and Unconstrained Testing Environment. *Transactions of the ACM on Mathematical Software*, 21(1):123–160, 1995.
  - [2] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *LANCELOT: a Fortran package for large-scale nonlinear optimization (Release A)*. Springer Series in Computational Mathematics. Springer Verlag, Heidelberg, Berlin, New York, 1992.
  - [3] N. I. M. Gould, D. Orban, and Ph. L. Toint. CUTer, a constrained and unconstrained testing environment, revisited. Technical Report TR/PA/01/04, CERFACS, Toulouse, France, 2001.
  - [4] N. I. M. Gould, D. Orban, and Ph. L. Toint. General CUTer documentation. Technical Report TR/PA/02/13, CERFACS, Toulouse, France, 2002.

### 5.3 Componentwise fast convergence in the solution of full-rank systems of nonlinear equations

**N. I. M. Gould:** *Rutherford Appleton Laboratory, England*; **D. Orban:** *CERFACS, France*;  
**A. Sartenaer:** *Facultés Universitaires Notre-Dame de la Paix, Belgium*; **Ph. L. Toint:** *Facultés Universitaires Notre-Dame de la Paix, Belgium*.

Following the good local convergence properties of primal-dual interior point methods for nonlinear programming [2], the asymptotic convergence of parameterized variants of Newton’s method for the solution of nonlinear systems of equations is considered. The original system is perturbed by a term involving the variables and a scalar parameter which is driven to zero as the iteration proceeds. The exact local solutions to the perturbed systems then form a differentiable path leading to a solution of the original system, the scalar parameter determining the progress along the path. A homotopy-type algorithm, which involves an inner iteration in which the perturbed systems are approximately solved, is outlined. It is shown that asymptotically a single linear system is solved per update of the scalar parameter.

It turns out that a *componentwise*  $Q$ -superlinear rate in the sequence of iterates may be attained under standard assumptions, and that this rate may be made arbitrarily close to quadratic. The theoretical results presented in the original report have been extended to show that the exact same convergence rate, which still applies componentwise, also occurs in the residuals. Numerical experiments have been added to illustrate the new results and we discuss the relationships that this method shares with interior methods in constrained optimization. This work resulted in report [1], which was reviewed and revised in 2001 and accepted for publication in a special issue of *Mathematical Programming*.

- 
- [1] N. I. M. Gould, D. Orban, A. Sartenaer, and Ph. L. Toint. Componentwise fast convergence in the solution of full-rank systems of nonlinear equations. Technical Report TR/PA/00/56, CERFACS, Toulouse, France, 2000. To appear in *Mathematical Programming Serie B*. [Compressed PS](#), [PDF](#).
  - [2] N. I. M. Gould, D. Orban, A. Sartenaer, and Ph. L. Toint. Superlinear convergence of primal-dual interior point algorithms for nonlinear programming. Technical Report TR/PA/00/20, CERFACS, Toulouse, France, 2000. Preliminary version of the article published in SIOPT, vol. 11, nber 4, pp 974-1002. [Compressed PS](#).

## 5.4 Optimization techniques for the regularization of large-scale discrete forms of ill-posed problems

**M. Rojas:** CERFACS, *France*; **D. C. Sorensen:** *Rice Univeristy, Houston, Texas, USA*; **T. Steihaug:** *University of Bergen, Bergen, Norway*; **D. Noll:** *Université Paul Sabatier, Toulouse, France*; **G. Tanoh:** *Université Paul Sabatier, Toulouse, France*

This project consists of using optimization techniques for the numerical treatment of ill-posed problems from inverse problems in different application areas. Two of the main features of this type of problem are the ill-posedness of the operators involved and the presence of noise in the data. Regularization methods try to recover useful information about the solution of these problems by solving a related problem with better conditioning where the effect of the noise in the data is minimized.

The regularization problem can be formulated in different ways. One of the most popular regularization approaches is the classical Tikhonov regularization which is equivalent to the problem of minimizing a quadratic subject to a quadratic constraint. The latter is known in optimization as the trust-region subproblem.

We consider methods for both the linear and nonlinear formulations of ill-posed problems as well as the application of those techniques to the solution of problems with field data. At the core of this project is the method LSTRS for the large-scale trust-region subproblem presented in [1]. LSTRS is an iterative procedure that requires the solution of a large eigenvalue problem at each step. This fact has given rise to new problems in the area of large-scale eigenvalue computation, as we describe later.

**Linear Regularization.** We are interested in the solution of large-scale linear systems and linear least-squares problems where the coefficient matrix is highly ill-conditioned and the right-hand side is contaminated with noise. These problems arise in important areas such as seismic inversion, medical imaging and image restoration. Our work in this area includes:

- With Dan Sorensen from Rice University, Houston: solving seismic inversion problems such as those presented in [2].
- With Trond Steihaug from the University of Bergen, Norway: developing a method for the problem of minimizing a quadratic subject to quadratic and nonnegativity constraints in the large-scale and ill-posed setting [3]. This kind of problem arises when solving inverse problems in image restoration. Our technique is one of the very few methods available for the efficient computation of nonnegative restorations. We presented some preliminary results of this research effort at the Applied Inverse Problems Conference in Montecatini Terme, Italy, June 2001.
- We have done some work on image restoration problems without nonnegativity constraints, which we presented at the 17th Symposium on Mathematical Programming in Atlanta, Georgia, USA, August 2000.
- We have also used our techniques for linear regularization on applications from magnetic field reconstruction.

**Nonlinear Regularization.** There have been few attempts to treat ill-posed problems as nonlinear problems, mostly limited to the nonlinear least squares formulation. This is mainly due to the fact that until now there were no efficient methods that could handle the high-degree singularities present in these problems. Very recently we have presented LSTRS, a method for the large-scale trust-region subproblem that is able to handle such singularities and has proven to be very efficient in practice [1]. This method can be used at the heart of an optimization method relying on the trust-region globalization strategy, and such an optimization method can be used in turn to treat nonlinear ill-posed problems.

Our work in this area consists of the development of a trust-region interior-point method using LSTRS [1] to solve the subproblems. Applications will be in the area of medical tomography. This work is in collaboration with Germain Tanoh and Dominique Noll from Université Paul Sabatier, Toulouse [4].

**Preconditioning Large-Scale Eigenvalue Problems.** As we mentioned before, the main computation in LSTRS is a sequence of large-scale eigenvalue problems. Thus, the success of the trust-region method depends on how efficiently and accurately we can solve the eigenvalue problems. This is especially relevant in the ill-posed case where the eigenvalues of interest tend to be very clustered and therefore, eigensolvers based on the Lanczos method, such as ARPACK, will exhibit slow convergence. One way to overcome this problem is to precondition the eigenvalue problems. With Dan Sorensen from Rice University in Houston, we are working on the design of efficient preconditioners for the eigenvalue problems arising in LSTRS. By taking advantage of the special structure of these problems we have been able to design efficient preconditioners, and we have obtained some promising preliminary results for these new techniques.

- 
- [1] M. Rojas, S.A. Santos, and D.C. Sorensen. A new matrix-free algorithm for the large-scale trust-region subproblem. *SIAM J. Optimization*, 11(3):611–646, 2000.
  - [2] M. Rojas and D. C. Sorensen. A trust-region approach to the regularization of large-scale discrete ill-posed problems. Technical Report TR/PA/00/57, CERFACS, Toulouse, France, 2000. Revised September 25, 2001, Accepted for publication in SISC. [Compressed PS](#), [PDF](#).
  - [3] M. Rojas and T. Steihaug. An interior-point trust-region-based method for large-scale nonnegative regularization. Technical Report TR/PA/01/11, CERFACS, Toulouse, France, 2001. July 6, 2001. Revised December 19, 2001. Submitted. [Compressed PS](#), [PDF](#).
  - [4] G. Tanoh, D. Noll, and M. Rojas. A trust-region interior-point method for large-scale nonlinear nonconvex programming (working title). Technical report.

---

## 6 Conferences and Seminars

---

### 6.1 Conferences and seminars attended by members of the Parallel Algorithms Project

#### January

Seventh US-Mexico Workshop in Numerical Analysis, Merida, Yucatan, Mexico, January 7-12, I. S. DUFF, *MA57 - A new code for the solution of sparse symmetric indefinite systems*, invited talk.

#### March

Large Linear Systems Workshop, Toulouse, France, March 1-2. I. S. DUFF, *A new multifrontal code for the solution of sparse symmetric indefinite systems*, invited speaker. B. CARPENTIERI, L. GIRAUD, J. LANGOU, T. MEŠKAUSKAS, L. PLANTIÉ, M. ROJAS, E. TRAVIESAS, attendee.

10th SIAM Conference on Parallel Processing for Scientific Computing Portsmouth, VA, March 12-14. I. S. DUFF, *Performance and tuning of two distributed memory sparse solvers*, joint work with P. R. Amestoy, J.-Y. L'Excellent and X. S. Li, invited minisymposium talk. C. VÖMEL, attendee.

#### April

2001 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Industrial Applications, Tahoe City, California, April 29 - May 1. L. GIRAUD, *Domain Decomposition Technique for Highly Variable Coefficient Problems*, joint work with F. Guevara Vasquez and R. S. Tuminaro, contributed talk. B. CARPENTIERI, *Preconditioning design in electromagnetism*, joint work with I. S. Duff and L. Giraud, contributed talk. I. S. DUFF, conference committee, guest co-editor of proceedings.

#### May

Seminar at NERSC, Lawrence Berkeley National Lab., May 8. L. GIRAUD, *A relaxation strategy for inner-outer linear solvers in domain decomposition methods*, joint work with A. Bouras and V. Frayssé, seminar.

Seminar at Sandia National Laboratories, Livermore, May 9. L. GIRAUD, *A relaxation strategy for an inner-outer linear solver within a Schur complement method*, joint work with A. Bouras and V. Frayssé, seminar.

Ecole d'Ingénieurs CPE, Lyon, France, May 10. F. CHAITIN-CHATELIN, *Calcul Qualitatif et Sens de la vie*, talk.

Parallel CFD 2001 - Industrial day, Egmond aan Zee, The Netherlands, May 21-23. L. GIRAUD, *Parallel computing and fast 3D Poisson solvers in Meso-scale atmospheric simulations*, joint work with R. Guivarch, P. Jabouille and J. Stein. invited talk.

IFIP WG 2.5 Meeting, Amsterdam, the Netherlands, May 26-27. F. CHAITIN-CHATELIN, *Qualitative Computing*, talk.

5th IMACS Conference on Iterative Methods in Scientific Computing, Heraklion, Crete, May 28-31. J. LANGOU, *A note on the Modified Gram-Schmidt Algorithm*, joint work with L. Giraud, talk.

Algèbre Linéaire et Arithmétique Calcul Numérique et Parallèle, Rabat, Morocco, May 28-31. I. S. DUFF, *Direct methods for the solution of large sparse systems*, invited talk.

2001 Congrès national de mathématiques appliquées et industrielles, SMAI'01, Pompadour, France, May 28 - June 01. E. TRAVIESAS, *Décomposition de Hessenberg par la méthode d'Arnoldi : Etude théorique*, joint work with F. Chaitin-Chatelin, L. Plantié and S. Gratton, poster session. E. TRAVIESAS, *Décomposition de Hessenberg par la méthode d'Arnoldi : Implantation logicielle*, joint work with F. Chaitin-Chatelin, L. Plantié and C. Mandry, poster presentation. L. PLANTIÉ, *Le problème canonique des couches limites de type triple couche*, poster presentation. J.-C. RIOUAL, *Domain Decomposition Methods for Semiconductor Device Simulations*, joint work with L. Giraud and J. Koster and A. Marrocco, poster presentation.

## June

Biennial Conferences on Numerical Analysis University of Dundee, Scotland, UK, June 26-29. I. S. DUFF, *The use of a new sparse indefinite code in optimization packages*, talk. C. VÖMEL, *Incremental norm estimation for dense and sparse matrices*, talk.

## July

High Performance Algorithms and Software for Nonlinear Optimization, Erice, Sicily, June 30-July 8. D. ORBAN, *Full-rank systems of nonlinear equations*, talk.

Recent Developments in Large Scale Scientific Computing. Porto, Portugal, July 3-6. I. S. DUFF, *Introduction to sparsity and direct methods and multifrontal methods and their parallel implementation.*, invited tutorial talks.

SIAM Meeting, San Diego, USA, July 9-13. E. TRAVIESAS, *Finite Precision computations and the toolbox PRECISE*, talk.

40th Session of Scientific Seminar on Modern Methods in Scientific Computing and Applications, Montreal, Canada, July 9-20. T. MEŠKAUSKAS, attendee.

## August

Summer School on Sparse Matrix Computations, Technical University of Denmark, August 13-18. I. S. DUFF, invited lecturer.

Seminar at DTU, Lyngby, Denmark. August 30. I. S. DUFF, *The use of a new sparse indefinite code in optimization packages*.

## September

Workshop on Numerical Methods for Evolutionary Problems, Peschici, Italy, September 17-21. B. CARPENTIERI, *Iterative solution and preconditioning of boundary integral equations in electromagnetism*, joint work with I. S. Duff and L. Giraud, contributed talk. F. CHAITIN-CHATELIN, *Computing Thoughts*, talk.

Seminar cluster - CCT CNES, CNES, September 19. L. GIRAUD, *Mixing MPI and OpenMP on the CERFACS clusters: Some preliminary experiments*, invited talk.

## October

Mini-lecture course at OUCL on Direct Methods for Sparse Matrices. I. S. DUFF, invited lecturer. Course given jointly with J. Scott.

The Institute of Ecotechnics Conference on Time: Metapatterns, The Present Moment and Evolution, Aix-en-Provence, France, October 26-29. F. CHAITIN-CHATELIN, *Life Computation*, talk.

## November

SC 2001, Denver, Colorado, November 10-15. I. S. DUFF, *Sparse direct methods and software for systems of equations*, invited tutorial talk. (In workshop with J. Dongarra, V. Eijkhout, and D. Sorensen).

Workshop PARAMAS, Paris - Orsay, November 29-30. L. GIRAUD, *Solving large linear systems on parallel computers: the CERFACS experience*, invited talk.

## December

IBM, New York, USA, December 27. F. CHAITIN-CHATELIN, *Hypercomplex Computation*, talk.

## 6.2 Conferences and seminars organized by the Parallel Algorithms Project

### February

*A stopping criterion for the conjugate gradient algorithm in a finite element method framework.* February 15. M. Arioli (Rutherford Appleton Laboratory).

### June

Sparse Day Meeting, June 13.

*Sparse bidiagonalization by Householder transformations*, G. Howell (Florida). *On support-graph preconditioning*. J. Gilbert (Xerox Parc, Palo Alto). *Least-squares polynomial filters for ill-conditioned linear systems in image processing*, Y. Saad (Minneapolis). *An ILU that incorporates the*

*growth of the inverse triangular factors*, M. Bollhoefer (Berlin University of Technology). *A multilevel algorithm for wavefront reduction*, J. Scott (Rutherford Appleton Laboratory). *A preconditioning approach for sparse linear systems arising in interior-point methods*, M. Sosonkina (University of Minnesota, Duluth). *Analysis of preconditioners for saddle-point problems*, D. Loghin (Oxford University Computing Laboratory).

## September

*Residual smoothing techniques: Do they improve the limiting accuracy of iterative solvers?* September 13. M. Rozložník ( Academy of Sciences of the Czech Republic ).

## 6.3 Internal seminars organized within the Parallel Algorithms Project

### March

*Incremental norm Estimation for dense and sparse matrices* (joint work with I. S. Duff). March 23. C. Vömel.

### June

*Notes on the modified Gram-Schmidt Algorithm*. June 14. J. Langou.

### September

*Mixing MPI and OpenMP on the CERFACS clusters (PCs and Compaq): some preliminary experiments*. September 6. L. Giraud. Reference: WN/PA/01/19. This talk is available on the Web.

---

## 7 Publications

---

### 7.1 Journal Publications

- [PUB1] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Analysis and Applications*, 23(1):15–41, 2001.
- [PUB2] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and X. S. Li. Analysis and comparison of two general sparse solvers for distributed memory computers. *ACM Trans. Math. Softw.*, 27(4):xx–yy, December 2001.
- [PUB3] Z.-Z. Bai, I. S. Duff, and A. J. Wathen. A class of incomplete orthogonal factorization methods. I: Methods and theories. *BIT*, 41(1):53–70, 2001.
- [PUB4] L. M. Carvalho, L. Giraud, and G. Meurant. Local preconditioners for two-level non-overlapping domain decomposition methods. *Numerical Linear Algebra with Applications*, 8(4):207–227, 2001.
- [PUB5] L. M. Carvalho, L. Giraud, and P. Le Tallec. Algebraic two-level preconditioners for the Schur complement method. *SIAM J. Scientific Computing*, 22(6):1987 – 2005, 2001.
- [PUB6] I. S. Duff and J. Koster. On algorithms for permuting large entries to the diagonal of a sparse matrix. *SIAM J. Matrix Analysis and Applications*, 22(4):973–996, 2001.
- [PUB7] L. Giraud, R. Guivarch, and J. Stein. Parallel distributed fast 3D Poisson solver for meso-scale atmospheric simulations. *Int. J. of High Performance Computing Applications*, 15(1):36–46, 2001.
- [PUB8] N. I. M. Gould, D. Orban, A. Sartenaer, and Ph. L. Toint. Superlinear convergence of primal-dual interior point algorithms for nonlinear programming. *SIAM J. Optimization*, 11(4):974–1002, 2001.
- [PUB9] A. Meister and C. Vömel. Efficient preconditioning of linear systems arising from the discretization of hyperbolic conservation laws. *Advances in Computational Mathematics*, 14:49–73, 2001.

### 7.2 Conference Proceedings and Book Chapters

- [PRO1] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and X. S. Li. Performance and tuning of two distributed memory sparse solvers. Technical report, December 2000. Accepted for presentation at the Tenth SIAM Conference on Parallel Processing for Scientific Computing that will be held in Portsmouth, Virginia from March 12th-14th, 2001.
- [PRO2] F. Chaitin-Chatelin and T. Meškauskas. Computation with hypercomplex numbers. In *Proceedings of 3rd World Congress of Nonlinear Analysts (WCNA-2000), Catania, Italy*, volume 47, pages 3391–3400. Nonlinear Analysis, Elsevier, 2001.

## 7.3 Theses

- [THS1] D. Orban. *Méthodes de points intérieurs pour l'optimisation non-linéaire*. Ph.D. Thesis, FUNDP/INPT, May 2001. TH/PA/01/09.

## 7.4 Technical Reports

- [TRP1] B. Carpentieri, I.S. Duff, L. Giraud, and M. Magolu monga Made. Sparse symmetric preconditioners for dense linear systems in electromagnetism. Technical Report TR/PA/01/35, CERFACS, Toulouse, France, 2001. Submitted. [Compressed PS](#), [PDF](#).
- [TRP2] F. Chaitin-Chatelin and T. Meškauskas. Inner-outer iterations for mode solvers in structural mechanics: application to the Code\_Aster. Contract Report CR/PA/01/85, CERFACS, 2001.
- [TRP3] F. Chaitin-Chatelin, E. Traviesas, and A. Ilahi. Résolution géométrique d'une équation de degré 4 : Partie 1. Working Notes WN/PA/01/01, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP4] F. Chatin-Chatelin and E. Traviesas. Homotopic perturbation - Unfolding the field of singularities of a matrix by a complex parameter: a global geometric approach. Technical Report TR/PA/01/84, CERFACS, Toulouse, France, 2001. Preliminary version. [Compressed PS](#), [PDF](#).
- [TRP5] I. S. Duff, M. A. Heroux, and R. Pozo. The Sparse BLAS. Technical Report TR/PA/01/24, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP6] I. S. Duff and C. Vömel. The Implementation of the Sparse BLAS in Fortran 95. Technical Report TR/PA/01/27, CERFACS, Toulouse, France, 2001. Submitted to ACM Trans. Math. Softw. [Compressed PS](#), [PDF](#).
- [TRP7] J. R. Gilbert and L. Grigori. A note on the column elimination tree. Technical Report TR/PA/01/104, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP8] L. Giraud. Combining shared and distributed memory programming models on clusters of symmetric multiprocessors: Some basic promising experiments. Working Note WN/PA/01/19, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP9] L. Giraud, J. Koster, A. Marrocco, and J.-C. Rioual. Domain decomposition methods in semiconductor device modeling. Technical Report TR/PA/01/51, CERFACS, Toulouse, France, 2001. To appear in the proceedings of the 13<sup>th</sup> conference on Domain Decomposition Methods in Scientific Computing, 2000. [Compressed PS](#), [PDF](#).
- [TRP10] L. Giraud and J. Langou. When modified Gram-Schmidt generates a well-conditioned set of vectors. Technical Report TR/PA/01/17, CERFACS, Toulouse, France, 2001. To appear in IMAJNA. [Compressed PS](#), [PDF](#).
- [TRP11] L. Giraud, F. Guevara Vasquez, and R. S. Tuminaro. Grid transfer operators for highly variable coefficient problems in two-level non-overlapping domain decomposition methods. Tech. Rep. TR/PA/01/03, CERFACS, Toulouse, France, 2001. Submitted. [Compressed PS](#), [PDF](#).
- [TRP12] N. I. M. Gould, D. Orban, and Ph. L. Toint. CUTer, a constrained and unconstrained testing environment, revisited. Technical Report TR/PA/01/04, CERFACS, Toulouse, France, 2001.
- [TRP13] C. Mandry and E. Traviesas. Convergence de la méthode d'Arnoldi en précision finie en fonction du vecteur initial. Working Note WN/PA/01/36, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).

- [TRP14] The Parallel Algorithms Project. Scientific Report for 2000. Technical Report TR/PA/01/23, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP15] The Parallel Algorithms Project. Scientific Report for 2001. Technical Report TR/PA/01/105, CERFACS, Toulouse, France, 2001. [Compressed PS](#), [PDF](#).
- [TRP16] M. Rojas and T. Steihaug. An interior-point trust-region-based method for large-scale nonnegative regularization. Technical Report TR/PA/01/11, CERFACS, Toulouse, France, 2001. July 6, 2001. Revised December 19, 2001. Submitted. [Compressed PS](#), [PDF](#).