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

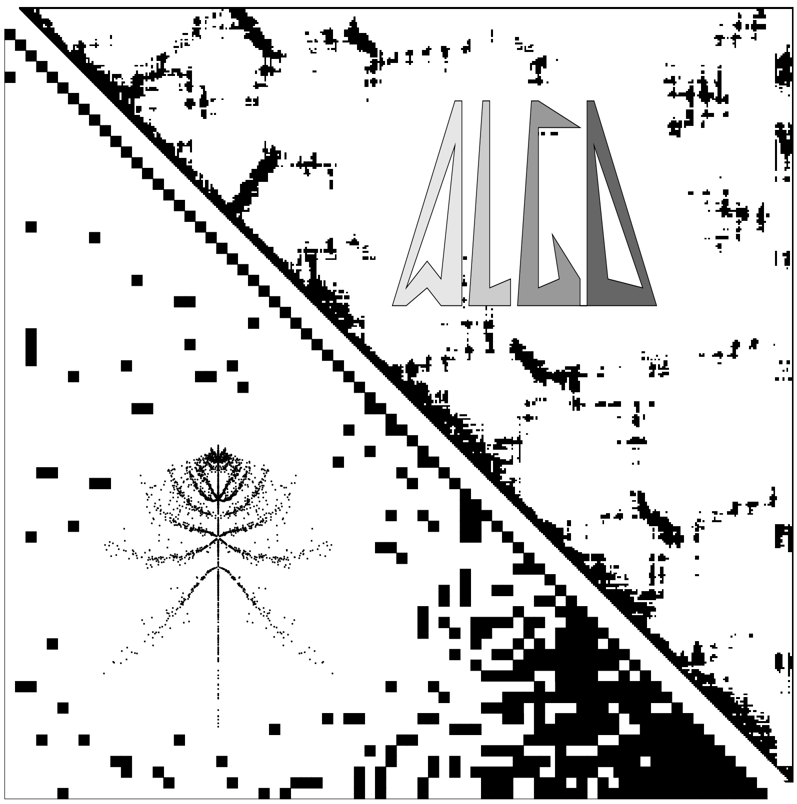

ACTIVITY REPORT
of the
PARALLEL ALGORITHMS PROJECT
at
CERFACS

JANUARY 1995 - DECEMBER 1996

SR/PA/97/03

Contents

1 Introduction

The research programme conducted by the Parallel Algorithm Project continues to combine the excitement of basic research discoveries with their utility in the solution of large-scale problems in science and engineering. We are thus concerned both with underlying mathematical and computational science research, the development of new techniques and algorithms, and their implementation on a range of computing platforms.

The description of our activities is divided into 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 Algorithm 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. Indeed this cross-fertilization extends to formal and informal collaboration with other teams at CERFACS, the partners of CERFACS, and collaborators and clients from outside.

One of the principal activities of the Project continues to be the solution of full and sparse systems of linear equations on high performance computers. This activity is basic to computational science and engineering and underpins other work by ourselves and others. Much of our work in this area is based on the platform provided by the Basic Linear Algebra Subprograms (BLAS) which we have developed for both parallel vector supercomputers (see previous Scientific Council Reports) and more recently for RISC processors that are used both in workstations and many high performance computers. A major tool in the design of a portable platform for these kernels has been our observation that the entire Level 3 BLAS suite can be constructed from the GEMM kernel for matrix-matrix multiply and the TRSM kernel for solving triangular systems. In the area of standard kernels and tools for supporting work on sparse matrices we have also continued our work on sparse Level 3 BLAS and on the design, contents, distribution and documentation of a new sparse matrix test collection.

We have developed Library quality software: for our now well recognized state of the art ordering scheme, the approximate minimum degree ordering (AMD); for the parallel solution of sparse symmetrically structured systems; and for the solution of sparse unsymmetric systems. Both these solvers use Level 3 BLAS kernels and the AMD ordering scheme and have been included in Release 12 of the Harwell Subroutine Library. Other work in sparse linear equation solution has involved rewriting and redesigning a general sparse code for distributed memory computers. Some preliminary experiments have been performed on the Meiko CS-2 and an IBM SP-2. We have also further developed the work on a parallel implementation of a multifrontal method for the QR factorization of a sparse matrix and have used this in new software for the solution of sparse least-squares problems. For this same problem, we have developed an approach which can use a standard LU factorization of the rectangular coefficient matrix and is particularly efficient when the matrix is nearly square. We are involved in a major EU LTR Project on parallel solvers called PARASOL. CERFACS role in the Project is to develop sparse direct solvers in cooperation with RAL. The other partners are PALLAS, GENIAS, and GMD (Germany), RAL (England), Parallab (Norway), and ONERA (France) and the industrial partners are DNV (Norway), MSC and INPRO (Germany), APEX (France), and Polyflow (Belgium). The Project is planned to continue until December 1998.

The use of preconditioning and partitioning techniques represent one of the most important and effective ways of exploiting parallelism, and we have several projects for solving large-scale systems arising from partial differential equations based on using domain decomposition methods and multigrid techniques implemented on both homogeneous and heterogeneous distributed memory computers and networks of workstations. We have also developed dynamic and static schedulers and used these in the parallel implementation of our block iterative methods on the Meiko CS-2, an IBM SP-2, and on a heterogeneous network of workstations. Some of this work on iterative methods has been concerned with the HPCN3S Project within the EUROPORT-1 Programme.

The Parallel Algorithms Project is involved in the European ESPRIT IV Project N. 20116 whose acronym is ODESIM. The other partners of ODESIM are CASA (Spain), CEIT (Spain), CR-FIAT (Italy), MATRA

DATAVISION (France) and Siemens (Germany). The aim of the project is the development of a set of software tools for kinematic and dynamic optimum design, designer-driven optimization and fully automatic optimization. The role of the Parallel Algorithms Project is to select the numerical optimization methods suitable for addressing the problems and to implement a parallel optimization process on a network of heterogeneous computers.
The algebraic eigenproblem has been a primary research area in linear algebra for many decades although it is only now that we are developing the theory and tools to handle the really large systems that occur in practical industrial applications. We have developed new algorithms and techniques for both symmetric and unsymmetric eigenproblems, including the case of complex non-Hermitian matrices, and have prototype codes that implement these ideas. The BLZPACK code has been used in problems in computational chemistry, including studies of the internal motions of proteins and the ground states of atoms and clusters. The code has also been used in shape optimization and plate bending problems with groups in Swansea and Brazil.
The main area of interest for the Qualitative Computing Group concerns the reliability of numerical software in finite precision arithmetic. The major aspects of this that have been studied during these past two years have been

- a study of instabilities in the neighbourhood of singularities,
- the development of a code (ARNCHEB) for computing part of the eigensystem of very non-normal matrices using the Arnoldi Tchebycheff method,
- the study of the influence of non-normality of iterative methods,
- the classification of scientific problems giving rise to non-normality, and
- the development of software tools for assessing the reliability and robustness of numerical methods (PRECISE).

The Qualitative Computing Group played a major role in the ILAY Workshop on eigenvalues in September and is involved in the European Project PINEAPL (HPCN, Fourth Framework Project:20018), which is a coordinated effort to produce a general purpose library of parallel numerical software. The role of CERFACS in this project is to analyse the numerical properties of the resulting library and the industrial applications. NAG is the leader of the consortium which includes British Aerospace, CERFACS, Thomson-CSF, CPS (Napoli), the Danish Hydraulic Institute, IBM SEMEA, the University of Manchester, Math-Tech, and PIAGGIO.
Most, if not all, of the techniques developed for the solution of linear systems have a direct influence on the performance of optimization software. We discuss some specific aspects of optimization in a separate section of this report.
Within the Parallel Algorithm Project, a small sub-team was working during the last two years in computational chemistry. Their primary research interests included Molecular Dynamics and Density Functional Theory. The Computational Chemistry Group was involved in collaborative projects with the CNRS GdR SPARCH, the French-German exchange programme PROCOPE, and with Grenoble and Erlangen under the EU TMR scheme.
The Parallel Algorithm Team is heavily involved in the Advanced Training aspects of CERFACS' mission. Not only have we given courses at CERFACS for partners and others but members of the Team have assisted in many lecture courses at other centres, including INPT in Toulouse. We have also been the primary host Team for the RECITE project and have welcomed many visitors to CERFACS for both short and long periods. We have been involved in several industrial contracts and have some at the planning stage. The interaction with the industrial scientists through these projects is as much an exercise in advanced training as in the research behind techniques used in the solution of the problem under consideration. During this reporting period, two more PhD students have joined the roll of honour of successful doctoral candidates. Tony Drummond and Vincent Toumazou both successfully defended their theses. Tony's thesis was on parallel implementations of block iterative schemes on distributed heterogeneous computing environments and Vincent's was on "Portraits spectraux de matrices : un outil d'analyse de la stabilité".
During the past two years, a major preoccupation of the Team has been the International Linear Algebra Year. We hosted four highly successful workshops; two in 1995 and two in 1996. Each attracted an attendance of between ninety and 120 people from all corners of the globe, with a good involvement of

engineers from industry and young research students. In 1995, both workshops were held at the UNESCO Centre in Toulouse. The first workshop (in September) was on Direct Methods for Sparse Matrices, the second (in October) was on Eigensystems. In 1996, we held a workshop on Optimization at Albi in April and completed the year with one on Iterative Methods at the UNESCO Centre in July. We also hosted a number of visits of both distinguished and junior scientists under the auspices of ILAY.
Information on ILAY and other aspects of the Project can be obtained through our World Wide Web pages that are accessible through the CERFACS home page (www.cerfacs.fr).

2 Dense & Sparse Matrix Computations

2.1 Computational kernels on vector and parallel machines and on RISC architectures (M.J. Daydé and I.S. Duff)

The Basic Linear Algebra Subprograms (BLAS), particularly the Level 3 BLAS, are one of the main platforms for the efficient implementation of a wide range of algorithms on modern high performance computers. The rapid development of new machines means that it is important to maintain a watching brief over new architectures and to develop and tune the BLAS in cases when the manufacturer fails to do so.

It is an impossible task to so redesign all of the BLAS (even just the Level 3 BLAS) for each machine so we have developed BLAS based on the GEMM, or matrix-matrix multiply, kernel so that only this one routine need be tuned for the machine. The GEMM kernel is usually the first to be optimized by the vendor and it is relatively easy to tune code for this operation. This work is described in detail in [2] and [3].

Our recent work, which has appeared appear in the report [2], shows that, even using a standard version for GEMM, our Level 3 BLAS can outperform some of the vendor-supplied BLAS on most machines including the Silicon Graphics Power Challenge and the DEC 8400 5/300 Turbo laser. Furthermore, if our kernels use the vendor-supplied version of GEMM, they can be made to run even faster.

We make extensive use of these new building blocks in linear algebra software and demonstrate that the high speed of these kernel operations can be reflected in the speed of the linear equation solvers [1].

[1] P. R. Amestoy, M. J. Daydé, I. S. Duff, and P. Morère, (1995), Linear Algebra Calculations on a Virtual Shared Memory Computer, *International Journal of High Speed Computing*, **7**, 21–43.

[2] M. J. Daydé and I. S. Duff, (1996), A Block Implementation of Level 3 BLAS for RISC Processors, Tech. Rep., CERFACS. TR/PA/96/06.

[3] M. J. Daydé, I. S. Duff, and A. Petitet, (1994), A Parallel Block Implementation of Level 3 BLAS Kernels for MIMD Vector Processors, *ACM Transactions on Mathematical Software*, **20**, 178–193.

2.2 Sparse BLAS (I.S. Duff, M. Marrone, G. Radicati and C. Vittoli)

The Basic Linear Algebra Subprograms (BLAS) have had a profound influence on the development of algorithms and software for the solution of systems of linear equations with full coefficient matrices. Although a set of sparse BLAS was developed [1], their use and manufacturer-supported availability are not very widespread. This is in part due to the fact that they are only Level 1 BLAS and suffer from the same shortcomings of inefficiency on modern computer architectures as in the full case. Another reason why they are not widely accepted is that, at least for direct sparse solution, it is more efficient to code the kernels using full linear algebra so that the full matrix BLAS can be used (see, for example, [4]).

However, in the solution of sparse equations using iterative methods, there is an urgent need for a standard interface for a sparse matrix by matrix multiplication and a sparse triangular solution routine. We have therefore designed routines for these cases as Level 3 BLAS so that the Level 2 equivalents are available as a particular case. In the sparse case, there is the problem of how the matrix is stored, and indeed storage schemes differ very widely and are usually applications dependent. We have therefore also designed routines for data conversion with a standard interface to include all the data structures we

are familiar with. Finally we also include two permutation routines to facilitate the efficient use of the kernels in iterative solvers.
In response to comments on our earlier report [2], we have made many changes to that proposal and have developed a revised proposal to address these points including designing a harness in Fortran 90 [3]. A prototype code for the sparse BLAS is available from the anonymous ftp server at RAL in file spblas.f in directory pub on seamus.cc.rl.ac.uk.

[1] D. Dodson, R. Grimes, and J. Lewis, (1991), Sparse extensions to the Fortran Basic Linear Algebra Subprograms, *ACM Transactions on Mathematical Software*, **17**, 253–263.

[2] I. S. Duff, M. Marrone, and G. Radicati, (1992), A proposal for user level sparse BLAS, Tech. Rep. RAL 92-087, Rutherford Appleton Laboratory.

[3] I. S. Duff, M. Marrone, G. Radicati, and C. Vittoli, (1995), A set of Level 3 Basic Linear Algebra Subprograms for sparse matrices, Tech. Rep. RAL-TR-95-049, Rutherford Appleton Laboratory.

[4] I. S. Duff, (1981), Full matrix techniques in sparse Gaussian elimination, In *Numerical Analysis Proceedings, Dundee 1981*, G. Watson, ed., Lecture Notes in Mathematics 912, Berlin, Springer-Verlag, 71–84.

2.3 The effect of putting large entries on the diagonal of sparse matrices (I.S. Duff and J. Koster)

We have developed software for permuting a sparse matrix so that the diagonal of the permuted matrix has entries of large absolute value. We have modified and extended Harwell Subroutine Library routine MC21, which finds a maximum transversal of any unsymmetric sparse matrix. A maximum transversal is a set of entries of maximum cardinality where no two members of the set are in the same row or column. For a structurally nonsingular matrix, this cardinality will be equal to the matrix order. Our extension consists of taking into account the values of the matrix entries and finding, among all maximum transversals, one that maximizes the minimum entry on the diagonal. We call such a transversal a bottleneck transversal. We have explored several ways of doing this and have designed one that is in practice similar in execution time to only a few calls to MC21 itself.
The permutation of a matrix to put large entries on the diagonal can be useful in several ways, and we are investigating several of these. For example, many current algorithms for the solution of sparse sets of linear equations (for example MA37 and MA41) perform an analysis phase that assumes that diagonal entries will be suitable as pivots during a subsequent numerical factorization phase.
We are also examining the use of bottleneck transversals in the context of a block iterative scheme like that developed in [1]. In such a scheme, a direct solution scheme is used on blocks of the matrix and the overall solution obtained by iterating over the blocks. By placing large entries on to the diagonal, we can improve the robustness of our direct solver, while at the same time helping the iterative part of the solver by reducing the influence of the overlapping part of the blocks.
Finally, we are examining the use of bottleneck transversals in techniques for obtaining preconditioning matrices for iterative solution. Clearly, it would appear that a diagonal preconditioner consisting of the bottleneck transversal might perform better than an arbitrary one and this effect might also influence more sophisticated preconditioners.
A fuller discussion of this work will shortly be available [2].

[1] M. Arioli, I.S. Duff, J. Noailles, and D. Ruiz, (1992), A block projection method for sparse matrices, *SIAM J. Sci. Stat. Comput.* 13, 47–70.

[2] I.S. Duff and J. Koster, (1997). The design and use of algorithms for permuting large entries to the diagonal of sparse matrices. To appear.

2.4 Threshold ordering of sparse matrices (M. Benzi, H. Choi and D. B. Szyld)

The effect of a threshold variant TPABLO of the permutation and partitioning algorithm PABLO [3] on the performance of preconditioned Krylov subspace methods was investigated. The goal of these (symmetric) permutations is to produce matrices with dense diagonal blocks, and in the threshold variant, with large entries in the diagonal blocks. TPABLO has already been used on matrices from Markov

chain models [2] in connection with stationary block iterations and iterative aggregation/disaggregation algorithms. In the present study [1], TPABLO is applied as a preprocessor for matrices arising from the discretization of non-selfadjoint elliptic partial differential equations.
The solvers tested where GMRES, BiCGStab, QMR and CGNR with various incomplete factorization preconditioners (ILU(0), ILUT, IQR). It was found that preprocessing the matrices with TPABLO has a positive effect on the overall performance, resulting in greater robustness and better convergence rates for highly nonsymmetric problems.
The experimental results were interpreted as follows. It is known that ILU preconditioning can fail for highly nonsymmetric and indefinite problems due to instabilities in the incomplete factorization. In this context, the instability refers to the fact that the norm of M^{-1} (where $M = LU$ is the incomplete LU preconditioner) can be extremely large, although the coefficient matrix A is actually well conditioned. This ill conditioning is caused by the long recurrences in the forward and backward triangular solves when the preconditioning is applied. In experiments with matrices arising from discretizations of PDEs, it was found that the incomplete factors tend to be much better conditioned with the TPABLO ordering than with the natural ordering. In other words, TPABLO has a stabilizing effect on the incomplete factorization. As a consequence, the preconditioned Krylov solvers were found to be more robust and more efficient when the TPABLO preprocessing was applied, especially for highly nonsymmetric problems.
This is a collaboration between M. Benzi (CERFACS) and D. B. Szyld and H. Choi from the Mathematics Department at Temple University in Philadelphia, USA.

[1] M. Benzi, H. Choi and D. B. Szyld (1997), Threshold ordering for preconditioning nonsymmetric problems, submitted to the *Hong Kong Workshop on Scientific Computing 1997*.

[2] H. Choi and D. B. Szyld (1996), Application of threshold partitioning of sparse matrices to Markov chains, *Proceedings of the IEEE International Computer Performance and Dependability Symposium, IPDS'96*, IEEE Computer Society Press, Los Alamitos, CA, 158–165.

[3] J. O'Neil and D. B. Szyld (1990), A block ordering method for sparse matrices, *SIAM J. Sci. Stat. Comput.* 11, 811–823.

2.5 An approximate minimum degree ordering (P.R. Amestoy, I.S. Duff and T.A. Davis)

The minimum degree reordering scheme for symmetric matrices is one of the oldest known reordering schemes for sparse matrices [3]. At each stage of the elimination, the next pivot is chosen as a diagonal entry in the reduced matrix with least number of other entries in its row. This very simple algorithm is easy to program and is a highly successful ordering scheme for reducing work and fill-in when performing Gaussian elimination. However, on larger problems, as commonly occur today, it can be quite slow. Most of the time is spent on the degree update, that is updating the row counts after a pivoting step. This cost is even more noticeable because of the recent increase in the speed of the numerical factorization algorithms.
The purpose of our work, which is described in more detail in [1], is to design an ordering scheme that closely approximates the minimum degree ordering, and so retains its good properties, while at the same time executes much more quickly than the best available implementation of minimum degree.
We use the same quotient graph model as most current implementations of minimum degree to represent the symbolic factorization but replace the costly degree update by an approximate update which in many cases will be exact. The resulting algorithm is typically much faster than previous minimum degree ordering algorithms and produces results that are comparable in quality with the best orderings from other minimum degree algorithms.
We show a comparison in Table 2.1 of our new algorithm with "MMD", the multiple minimum degree algorithm of Liu [2], and the minimum degree ordering from the Harwell Subroutine Library code "MA27". This clearly illustrates the power of our new algorithm. It has been incorporated in the Harwell Subroutine Library at Release 12 as subroutine "MC47".

[1] P. R. Amestoy, T. A. Davis, and I. S. Duff, (1996), An approximate minimum degree ordering algorithm, *SIAM J. Matrix Anal. Applic.*, **17**, 886–905.

[2] J. H. W. Liu, (1985), Modification of the minimum degree algorithm by multiple elimination, *ACM Transactions on Mathematical Software*, **11**, 141–153.

Matrix	Order	Entries	Number of entries in **L** (in thousands)			Time (seconds on SUN SPARC 10)		
			MC47	MMD	MA27	MC47	MMD	MA27
RAEFSKY3	21200	733784	4709	4779	5041	1.05	2.79	1.23
BCSSTK31	35588	572914	5115	5231	6056	4.55	11.60	7.92
FINAN512	74752	261120	4778	8180	8159	15.03	895.23	40.31
BBMAT	38744	1274141	19673	19876	21139	27.80	200.86	134.58
ORANI678	2529	85426	147	147	147	5.49	124.99	124.66
PSMIGR1	3140	410781	3020	2974	2966	10.61	186.07	229.51

Table 2.1: Comparison of MC47 with MMD and MA27 ordering.

[3] W. F. Tinney and J. W. Walker, (1967), Direct solutions of sparse network equations by optimally ordered triangular factorization, In *Proceeedings of the IEEE*, vol. 55, 1801–1809.

2.6 A combined unifrontal/multifrontal method (T.A. Davis and I.S. Duff)

Frontal methods are particularly powerful inasmuch as they make considerable use of the Level 3 BLAS kernels while keeping the amount of integer overhead and data movement low. However, they usually require a far greater number of floating-point operations than methods using an ordering based on Markowitz or minimum degree. Multifrontal methods attempt to capture the benefits of both the full arithmetic of frontal schemes and the orderings designed to reduce fill-in and arithmetic. They are quite successful at this but still involve more data movement and indirect addressing than a frontal method. The object of this work is to combine elements of a frontal solver into a multifrontal solver in order to obtain a factorization that is efficient in arithmetic, number of operations, and overhead operations.

In the multifrontal method, several frontal matrices are used. Each is used for one or more pivot steps, and the resulting Schur complement is summed with other Schur complements to generate another frontal matrix. Although this means that arbitrary sparsity patterns can be handled efficiently, extra work is required to add the Schur complements together and can be costly because indirect addressing is required. The (uni-)frontal method avoids this extra work by factorizing the matrix with a single frontal matrix. Rows and columns are added to the frontal matrix, and pivot rows and columns are removed. Data movement is simpler, but higher fill-in can result if the matrix cannot be permuted into a variable-band form with small profile. We consider a combined unifrontal/multifrontal algorithm to enable general fill-in reducing orderings to be applied without the data movement of previous multifrontal approaches.

Suppose we attempt to factorize a matrix with the frontal method, but we give it a smaller working array than it requires. The factorization proceeds until the point at which the front size would grow larger than the working array. At this point the frontal method halts, unable to proceed. However, the remaining terms in the current frontal matrix form a Schur complement (which we refer to as a contribution block). We store the contribution block and deallocate the working array. A new pivot is then selected based on a fill-in reducing heuristic on the whole reduced matrix, and a new frontal matrix started with this starting pivot, in a newly allocated working array. The contribution block from the first frontal matrix is eventually assembled into a subsequent frontal matrix. This is the combined unifrontal/multifrontal method [1].

We have incorporated this algorithm in the unsymmetric multifrontal code in the Harwell Subroutine Library, "MA38", and show, in Table 2.2, a comparison of this new code with the frontal code, "MA42", and an earlier version of the "MA38" code that did not include this device. The test matrices, from the Harwell-Boeing Collection and some large examples collected by Tim Davis, show quite clearly the advantage of using our technique.

[1] T.A. Davis and I.S. Duff. A combined unifrontal/multifrontal method for unsymmetric sparse matrices. Technical Report TR-95-020, Computer and Information Science Department, University of Florida.

Matrix	Order	Entries	Number of entries in **L** (10^6)			Time (seconds on CRAY C-98)		
			MA42	MA38 (old)	MA38	MA42	MA38 (old)	MA38
GRE1107	1107	5664	0.1	0.3	0.2	0.21	0.30	0.30
GEMAT11	4929	33185	2.2	0.4	0.3	15.95	0.18	0.20
ORANI678	2529	90158	4.1	1.1	1.0	9.62	2.07	1.36
LHR04	4101	82682	0.7	1.5	0.8	4.21	2.51	1.79
HYDR1	5308	23752	1.1	0.6	0.4	8.01	1.05	1.10

Table 2.2: Comparison of MA38 with earlier version and MA42.

2.7 MUPS: a MUltifrontal Parallel Solver for sparse unsymmetric sets of linear equations (P.R. Amestoy and I.S. Duff)

We have developed a new code, MA41, for the solution of sets of sparse unsymmetric equations using a multifrontal method. MA41 is similar to MA37 in the sense that the initial analysis is performed on a symmetrized pattern of the original matrix and assumes that pivots can be chosen from the diagonal in any order. It is thus best suited to systems which are structurally symmetric or nearly so and which are diagonally dominant. It differs from MA37 in several respects.

We have reorganized the structure of the code along the lines described in [4] and [3] so that it is no longer assumed that the code is being executed on a single processor. This makes the internal data structures a little more complicated but the overhead for uniprocessor running is almost negligible. We also require a more complex data management policy so that garbage collections do not interfere with parallel efficiency, see [2]. Another main change from MA37 is that we make extensive use of the higher Level BLAS which increases computational efficiency on all machines, particularly those with caches or vector chips [1].

The uniprocessor version is included in the standard release of the Harwell Subroutine Library (Release 12), but versions for shared memory parallel computers are also available on request. The code is designed so that it is relatively simple to develop a version for any shared memory machine since the only primitives used by the code involve task starting and locks. At present, we have fully tested versions for the CRAY YMP/C90 range of machines and for an Alliant FX/80 and are currently writing a version for multiprocessor Sun workstations. Versions of the code have also been designed and run on some virtual shared memory machines, for example on a BBN TC 2000 and a KSR-1.

We show, in Table 2.3 below, a comparison of the performance of the uniprocessor version of MA41 with MA37 on a range of our standard test problems from the Harwell-Boeing Collection and some matrices collected by Tim Davis.

Matrix	Order	Entries	Analyse time		Factorization time		Solve time	
			MA41	MA37	MA41	MA37	MA41	MA37
ORANI678	2529	90158	4.15	28.41	46.02	251.18	0.29	0.36
BCSSTK15	3948	60882	0.61	0.96	11.79	104.80	0.18	0.34
LNS 3937	3937	25407	0.32	0.34	1.79	13.22	0.07	0.12
SHERMAN3	5005	20033	0.61	0.31	1.07	3.23	0.05	0.06
GOODWIN	7320	324772	1.09	1.81	8.18	510.65	0.23	0.77
WANG3	26064	177168	8.63	5.09	299.50	1591.21	2.17	2.70

Table 2.3: Comparison of MA41 with MA37. Times in seconds on an HP 715/64 workstation.

[1] P. R. Amestoy and I. S. Duff, (1989), Vectorization of a multiprocessor multifrontal code, *Int. J. of Supercomputer Applics.*, **3**, 41–59.

[2] P. R. Amestoy and I. S. Duff, (1993), Memory management issues in sparse multifrontal methods on multiprocessors, *Int. J. of Supercomputer Applics.*, **7**, 64–82.

[3] I. S. Duff, (1986), Parallel implementation of multifrontal schemes, *Parallel Computing*, **3**, 193–204.

[4] I. S. Duff, (1989), Multiprocessing a sparse matrix code on the Alliant FX/8, *J. Comput. Appl. Math.*, **27**, 229–239.

2.8 Parallel Multifrontal Solvers (J.-Y. L'Excellent and I. S. Duff)

CERFACS is involved in the development of direct solvers within PARASOL, an EU ESPRIT IV LTR Project for "An Integrated Environment for Parallel Sparse Matrix Solvers". The work is based on the shared memory parallel multifrontal code MA41 developed by Amestoy and Duff [1], that was extended by Amestoy and Espirat [2] to a message passing version based on PVM.
Work at CERFACS has included developing an MPI version of this code. This version is being developed for various platforms such as the SP2, SGI Power Challenge, and the Cray T3D. We have also ported and experimented with the shared memory code MA41 on our SGI Power Challenge and have done a performance study on that machine. Finally we have been studying the test problems from the industrial partners of the PARASOL project and have performed an analysis to define the requirements of the code. We conclude that, in collaboration with RAL and ENSEEIHT, a distributed multifrontal solver for both symmetric and unsymmetric matrices should be developed under MPI, which will include out-of-core facilities, an element-entry and satisfy the PARASOL interface defined in [3].

[1] P. R. Amestoy and I. S. Duff, Memory management issues in sparse multifrontal methods on multiprocessors, *Int. J. Supercomputer Applics.* 7:64–82. 1993.

[2] V. Espirat, Développement d'une approche multifrontale pour machines a mémoire distribuée et réseau hétérogène de stations de travail, Rapport de stage de 3eme année, ENSEEIHT-IRIT, 1996.

[3] A. Supalov (editor), PARASOL Interface Specification, GMD SCAI, 1996.

2.9 A parallel code for the multifrontal QR factorization (P.R. Amestoy, I.S. Duff and C. Puglisi)

We have designed and implemented a parallel multifrontal **QR** decomposition algorithm of large sparse matrices for shared memory multiprocessor machines [2]. The most common use of such factorization is to solve linear least-squares problems, but several other uses can be listed such as eigenvalue calculations, the projection of a point into a manifold, etc. The algorithm is based on the multifrontal approach [3] and makes use of Householder transformations. Our approach is unusual inasmuch we keep the orthogonal factors **Q** which may be required in some applications (for example if an orthogonal basis for the column space of **A** is wanted) and may be crucial in the solution of ill-conditioned or weighted least-squares problems. Another novel aspect of this work is that we use relaxation of the sparsity pattern of both the original matrix and the frontal matrices to improve the performance by enabling more efficient use of the Level 3 BLAS.
The numerical factorization is preceded by a pre-processing phase, where the pattern of the matrix or the normal equation matrix is considered. In this phase, the matrix can optionally be permuted in block upper triangular form [4] and a column ordering can optionally be computed using a selection criterion to preserve sparsity (Approximate Minimum Degree algorithm [1]). The assembly tree and subsidiary information are subsequently constructed for actual factorization. Driven by the assembly tree, the numerical factorization computes the **Q** and **R** factors by partially factorizing dense matrices associated with each node of the assembly tree. Level 3 BLAS operations can optionally be used to speed the computation of the full **QR** factorization [5].
The principal feature of the code is that it can solve large least-squares problems using the corrected semi-normal equations, that is solving $\mathbf{R}^T\mathbf{R}\mathbf{x} = \mathbf{A}^T\mathbf{b}$, or by solving $\mathbf{Q}\mathbf{R}\mathbf{x} = \mathbf{b}$. In both cases, iterative refinement can be used and the error estimate is available.
We have considered the parallelization of the factorization, of the matrix-vector multiplication $\mathbf{Q}^T\mathbf{b}$ and the triangular solves $\mathbf{R}\mathbf{x} = \mathbf{b}$ and $\mathbf{R}^T\mathbf{x} = \mathbf{b}$. The tasks are distributed among processors according to the assembly tree for all four phases. We show, in Table 2.4, that the four phases of the solution of least squares problems parallelize well. The runs for this table were performed on an Alliant FX/80.

Matrices	Nb of Proc.	$\mathbf{A = QR}$ Time	Speedup	$\mathbf{Q^T b}$ Time	Speedup	$\mathbf{Rx = b}$ Time	Speedup	$\mathbf{R^T x = b}$ Time	Speedup
LARGE	1	198		2.59		2.69		2.44	
(28254×17264)	8	32	6.2	0.62	4.18	0.42	6.40	0.43	5.66
MEDIUM2	1	139		2.84		1.79		1.66	
(18794×12238)	8	29	4.8	0.77	3.69	0.32	5.59	0.37	4.52
LARGE2	1	992		10.98		5.49		3.10	
(56508×34528)	8	219	4.5	3.53	3.11	1.11	4.95	1.19	2.59
NFAC80	1	58		1.80		1.02		0.97	
(24964×6400)	8	10	5.8	0.34	5.29	0.18	5.67	0.18	5.25
NFAC100	1	112		2.98		1.65		1.54	
(39204×10000)	8	20	5.6	0.70	4.26	0.27	6.11	0.31	4.98

Table 2.4: Time (in seconds) and speedups for the multifrontal **QR** factorization and solve phases on the 8 processor Alliant FX/80.

We are modifing the code following the standards of the Harwell Subroutine Library in order to integrate it in the next release of this library.

[1] P. Amestoy, T. Davis, and I. Duff, (1996), An approximate minimum degree ordering algorithm, *SIAM J. Matrix Analysis and Applications* 17(4), 886–905..

[2] P. R. Amestoy, I. S. Duff, and C. Puglisi, (1996), Multifrontal **QR** factorization in a multiprocessor environment, *Numerical Linear Algebra with Applications* 3(4), 275–300.

[3] I. S. Duff and J. K. Reid, (1983), The multifrontal solution of indefinite sparse symmetric linear systems, *ACM Transactions on Mathematical Software* **9**, 302–325.

[4] A. Pothen and C. J. Fan, (1990), Computing the block triangular form of a sparse matrix, *ACM Transactions on Mathematical Software* **16**, 303–324.

[5] C. Puglisi, (1992), Modification of the Householder method based on the compact WY representation, *SIAM Journal on Scientific and Statistical Computing* **13**, 723–726.

2.10 Parallel sparse LU factorization (R.H. Bisseling and J. Koster)

We have developed an algorithm for the LU decomposition of large sparse general matrices on distributed-memory platforms.

The algorithms, originally intended for architectures with a rectangular mesh communication network, were adapted to more modern architectures where in general communication latency is high [2]. For example, the new code strives for larger sets of pivots that can be handled simultaneously. This way, fewer messages (with a higher volume) need to be sent, thereby amortising latency costs, and we observed a significant improvement of the performance on parallel machines like the MEIKO CS2-HA and the IBM SP2. Besides further reduction of the communication costs, we are also examining the design and use of more efficient local data storage schemes and the use of dense matrix kernels which we expect will considerably speed up the single node performance of the code.

This is a collaboration with Rob Bisseling of the Computational Science working group of the Department of Mathematics of Utrecht University, the Netherlands.

[1] J. Koster and R. H. Bisseling, An Improved Algorithm for Parallel Sparse LU Decomposition on a Distributed-Memory Multiprocessor, *Proceedings of the Fifth SIAM Conference on Applied Linear Algebra*, J. G. Lewis, ed., SIAM, 1994, pp. 397–401.

[2] J. Koster and R. H. Bisseling, Parallel Sparse LU Decomposition on a Distributed-Memory Multiprocessor, submitted.

2.11 The solution of linear least-squares problems (J. Cardenal, I.S. Duff and J. Jimenez)

We have designed and developed a general method for the linear least-squares solution of overdetermined and underdetermined systems. The method is particularly efficient when the coefficient matrix is quasi-square, that is when the number of rows and number of columns is almost the same. The numerical methods proposed in the literature for linear least-squares problems and minimum-norm solutions do not generally take account of this special characteristic. The proposed method is based on an LU factorization of the original quasi-square matrix **A**, assuming that **A** has full rank.
We have developed a framework based on the augmented system

$$\begin{bmatrix} \mathbf{I} & \mathbf{A} \\ \mathbf{A}^T & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{r} \\ \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{O} \end{bmatrix},$$

for presenting our algorithm and give a unified approach to solving both least-squares problems and minimum-norm problems.
In the overdetermined case, the **LU** factors are used to compute a basis for the null space of $\mathbf{A}^T$. The right-hand side vector **b** is then projected onto this subspace and the least-squares solution is obtained from the solution of this reduced problem. In the case of underdetermined systems, the desired solution is again obtained through the solution of a reduced system. The use of this method may lead to important savings in computational time for both dense and sparse matrices.
We have studied the performance of our algorithm on practical test problems arising in the solution of problems from the computer simulation of the kinematic behaviour of multibody systems. Some results from this study are shown in Table 2.5.

	Code	Test 1	Test 2	Test 3	Test 4	Test 5
Rows×Columns		117x105	170x161	181x180	385x361	574x526
Entries		615	1069	1200	1661	3365
Normal equations	LAPACK	57.41	203.26	238.61	1914.05	5990.06
	MA27	12.13	32.31	36.99	100.25	288.55
Augmented system	MA27	7.31	15.59	13.13	43.44	120.70
Proposed Method	LAPACK	29.83	47.71	56.36	391.71	2124.20
	MA48	4.53	7.25	4.09	21.07	67.65

Table 2.5: Solution of sparse least-squares problems. Times in milliseconds on Silicon Graphics Onyx workstation (150 MHz).

The test cases all come from multibody simulation problems and, although small by current sparse matrix standards, they are typical of realistic problems arising in this application, indeed the three smaller cases are more normal. It is thus quite surprising that the dense codes from LAPACK do so badly relatively to the sparse codes. From these results, we see that our proposed method does very well, especially if a good sparse code is used for the initial **LU** factorization of **A**. A full report on this work will shortly be available [1].

[1] J.M. Jimenez, J. Cardenal and I.S. Duff. A projection method for the solution of linear least-squares problems. Technical Report RAL-TR-96-013, Rutherford Appleton Laboratory.

2.12 Sparse matrix test problems (I.S. Duff)

Release I of the Harwell-Boeing Collection of sparse matrices consists of sparse matrices in a standard format from a wide range of application areas. The format used is basically a column-pointer row-index format and is described in [1]. There are both symmetric and unsymmetric systems, some rectangular matrices, and a few in unassembled element format. Sometimes only the pattern of the sparse matrix is given, and sometimes a right-hand side vector is supplied in addition to the reals.

We are just completing the design of a new release that will be called the Rutherford-Boeing Sparse Matrix Collection. A draft specification for this has already been circulated [2]. This redesign has proved to be a significant task since we wish to enhance the value and scope of the Collection while effectively maintaining upward compatibility with the present release. In particular, we have extended the format to include more information on matrices and on systems being solved. These include starting guesses, exact solutions, eigenvalues and vectors, singular values and vectors, permutations and geometric data. We have also made it a requirement that the format of the matrices enables them to be read easily from C programs. Although we may include matrix generation programs in the repository, we will not include generated matrices in the Collection unless they are explicitly submitted and accepted for the Collection. We propose this because we feel that one of the main strengths of the Collection is its reproducibility and do not want to jeopardize this aspect.

We have written complete Fortran subroutines both to read matrices in the Collection and to write a matrix in Harwell-Boeing format. These routines are included in the Harwell Subroutine Library as "MC36" and "MC37" and are distributed with the test Collection. Similar codes are currently being designed for the new Rutherford-Boeing formats.

Release I is available by anonymous ftp from both RAL (matisa.cc.rl.ac.uk) and CERFACS (ftp.cerfacs.fr), in both cases in directory pub/harwell_boeing. More recently, some researchers at NIST in Washington DC have developed an efficient World Wide Web interface to the Collection that will be used as the primary vehicle for accessing the Rutherford-Boeing Collection (`http://math.nist.gov/MatrixMarket`).

[1] I. S. Duff, R. G. Grimes, and J. G. Lewis, (1992), Users' Guide for the Harwell-Boeing Sparse Matrix Collection (Release I), Tech. Rep. RAL 92-086, Rutherford Appleton Laboratory.

[2] I. S. Duff, R. G. Grimes, and J. G. Lewis, (1996), The Rutherford-Boeing Sparse Matrix Collection, tech. rep., Rutherford Appleton Laboratory. To appear.

$\boxed{3}$ Iterative Methods and Preconditioning

3.1 Linear algebra kernels for parallel domain decomposition methods (L.M. Carvalho, I.S. Duff and L. Giraud)

The modelling of many applications in physics and engineering gives rise to systems of nonlinear partial differential equations (PDEs), which can be solved by nonlinear iterative procedures that linearize and discretize the equations. Due to memory limitations, the huge sparse linear systems involved in this approach can seldom be solved using only direct methods. Equally, the use of standard iterative techniques is not practical because of the slow convergence. We thus combine iterative techniques with direct methods for solving subproblems to accelerate the convergence. The used method is additionally suitable for parallel computing on distributed memory machines.

We study the parallel implementation of a domain decomposition method based on a Schur complement approach. We divide the problem into separate domains and, within each domain, we partition the variables into internal and interface variables. We can then solve for the interior variables of each domain using a direct method. The resulting problem in the interface variables has the Schur complement matrix as coefficient matrix. We solve this using a preconditioned iterative method.

The preconditioner that we use is a block Jacobi type, where the blocks correspond to the contribution to the diagonal blocks of the Schur complement matrix from each subdomain. We approximate the inverse of each of these blocks by tridiagonal matrices using the probing technique of Chan [2]. This approximation is then factorized and used as a preconditioner to accelerate the conjugate gradient iteration on the interface variables.

We have found that the performance of the method is strongly influenced by the direct method used to solve the subproblems. We choose to solve the subproblems with MA27 and have experimented with this solution scheme on various parallel environments including a Meiko CS-2HA, a cluster of workstations, an SP2 and a T3D. We have also studied the behaviour and the scalability of the code when using different message-passing libraries, PVM and MPI, and virtual shared memory paradigm on the Meiko (Elan) and on the T3D (shmem).

More details of these results can be found in [1].

[1] L. Carvalho, I.S. Duff, and L. Giraud, (1996). Linear algebra kernels for parallel domain decomposition methods, in *Advanced Computational Methods in Structural Mechanics*, M. Papadrakakis and G. Bugeda (eds), CIMNE, Barcelona, Spain, 1-17.

[2] T.F. Chan, and T.P. Mathew, (1992). The interface probing technique in domain decomposition, *SIAM J. Matrix Anal. Applic.* 13, 212-238.

3.2 Additive Schwarz for the Schur complement method (L.M. Carvalho and L. Giraud)

Domain decomposition methods for solving elliptic boundary problems have been receiving increased attention for the last two decades. To a large extent this is due to their potential application on new parallel computers.

We present a new preconditioner for nonoverlapping domain methods when solving self-adjoint elliptic equations in two dimensions and we report on some experimental results. This alternative can be viewed as

a generalization of a block diagonal preconditioner for the Schur complement matrix. This preconditioner introduces an overlap between the blocks by including the cross-points and their couplings in the diagonal blocks of the block Jacobi preconditioner. We will refer to this new preconditioner as Algebraic Additive Schwarz (AAS) since we sum the contributions of each block on the overlap.
The nonzero sub-blocks of the Schur complement are dense matrices, we consider approximations which are inexpensive to construct and invert. The algebraic approximation of the interface operator is constructed by using the probing technique studied in [2]. The proposed preconditioner belongs to the one-level type preconditioners. Therefore, it does not implement any coarse grid component. Currently, most preconditioners include a coarse correction to propagate the error globally and this is the next step of our work. We should stress that the AAS approach improves the convergence rate of the well-known block Jacobi method (see section 3.1) on some model problems while retaining the same computational complexity.

[1] L. Carvalho and L. Giraud, (1996), Additive Schwarz for the Schur complement method. TR/PA/96/51, CERFACS. Submitted to the proceedings of the Ninth International Conference on Domain Decomposition methods.

[2] T.F. Chan, and T.P. Mathew, (1992). The interface probing technique in domain decomposition, *SIAM J. Matrix Anal. Applic.* 13, 212-238.

3.3 Parallel domain decomposition for device modeling simulation (L.M. Carvalho and L. Giraud)

We developed a parallel implementation of a domain decomposition method without overlap for the numerical simulation of semiconductor devices in two dimensions. The approximate solution is achieved by a nonlinear iterative procedure, the main kernel of which consists in solving very large linear systems. The associated matrices are sparse, symmetric and positive definite (SPD) and the linear systems are solved by using a preconditioned conjugate gradient method applied to the Schur complement. This work is the development of a previous study by L. Giraud and R. Tuminaro [4] extended for parallel computation. In that paper, a new Schur complement domain decomposition preconditioner was proposed and promising experimental results for the sequential implementation were given. This work focusses upon the issues that arise in an efficient implementation of the method on parallel distributed platforms; mainly, the following topics are addressed:

- preconditioners for the Schur complement method,
- linear solvers on the subdomains,
- message passing versus virtual shared memory paradigm for the implementation of the communications,
- stopping criterion for the Schur complement method.

We consider the following 2^{nd} order self-adjoint elliptic problem on an open polygonal domain Ω included in $I\!R^2$:

$$\left\{ \begin{array}{rcll} -\nabla(a(x,y)\nabla u) & = & f(x,y) & \text{in} \quad \Omega \\ u & = & 0 & \text{on} \quad \partial\Omega, \quad \text{the boundary of } \Omega \;, \end{array} \right. \tag{3.1}$$

where $a(x,y) \in I\!R^2$ is a positive function on Ω. We assume that the domain Ω is partitioned into N non-overlapping subdomains Ω_1,..., Ω_N. We discretize (3.1) which results in the SPD linear system $Au = f$. Grouping the points corresponding to the interface among the subdomains (B) in the vector u_B and the ones corresponding to the interior (I) of the subdomains in u_I, we get to the reordered problem:

$$\begin{pmatrix} A_{II} & A_{IB} \\ A_{IB}^T & A_{BB} \end{pmatrix} \begin{pmatrix} u_I \\ u_B \end{pmatrix} = \begin{pmatrix} f_I \\ f_B \end{pmatrix} \;. \tag{3.2}$$

Eliminating u_I from the second block row of (3.2) leads to the following reduced equation for u_B:

$$Su_B = f_B - A_{IB}^T A_{II}^{-1} f_I \;, \qquad \text{where} \qquad S = A_{BB} - A_{IB}^T A_{II}^{-1} A_{IB} \tag{3.3}$$

is referred to as the Schur complement matrix.
We have implemented two preconditioners for the Schur complement system. The first is an approximate block Jacobi preconditioner constructed using the algebraic probing technique introduced by T. Chan [2]. Because of our ordering of the nodes on the interfaces, each diagonal block of the Schur complement describes the coupling of the nodes within the same edge of a subdomain. A banded approximation of these blocks is built with matrix-vector products between S and a few carefully chosen vectors. The resulting narrow banded matrices are good, and easily inverted, approximations to the original diagonal blocks. Much information has to be exchanged among neighbouring subdomains at this points for the computation of all the matrix-vector products required by the construction of the preconditioner; this can represent a significant part of the time required by the complete linear solver. However, we are able, by providing some extra storage, to organize the computation in such a way that each subdomain has only to perform two exchanges with each of its neighbours. Unfortunately, with this approximate block Jacobi preconditioner there is no coupling between two horizontal and vertical neighbour interfaces. The second preconditioner, proposed in [4], tries to overcome this lack of coupling by adding the cross-points that couple the interfaces. This new preconditioner corresponds to an algebraic additive Schwarz method, where the overlap appears only at the cross points. From a parallel implementation point of view, this preconditioner does not require more communications than the first one and has an equivalent float-point operations cost.

We study the performance of the code using either PVM or MPI as well as the shared memory paradigm available on Meiko CS2-HA (CERFACS) and Cray-T3D (CEA-Saclay). We measure the parallel performances for the simulation of a mosfet device on different platforms ranging from cluster of workstations to the Meiko, see Table 3.1 for preliminary results, and the Cray T3D.

	# domains		
MA27	2×1	2×2	2×4
size of S	128	255	508
# iterations	19	23	33
time per iteration (s)	0.074	0.038	0.021
global time (s)	3.41	1.88	1.15
speedup	1	1.81	2.97

Table 3.1: Performance of the parallel preconditioned solver on a 128×128 grid on an eight node Meiko-CS2

We study a stopping criterion based on the normwise backward error $\frac{||Su_b - f_b||}{||S|| ||u_b|| + ||f_b||}$ for the preconditioned conjugate gradient method in the solution of (3.3). This criterion uses the norm of S which can be easily bounded during the probe construction of the preconditioners. We also study, through the backward error, the influence of the accuracy of the local iterative solvers (the stopping criterion of the PCG is also based on the normwise backward error) on the accuracy of the global problem (3.2).

[1] L.M. Carvalho, I.S. Duff and L. Giraud , Linear Algebra Kernels for Parallel Domain Decomposition Methods, in *Advanced Computational methods in Structural Mechanics*, M. Papadrakakis and G. Bugeda (Eds.), CIMNE Publications, Barcelona, 1–17 (1996).

[2] T.F. Chan and T.P. Mathew, The interface probing technique in domain decomposition, *SIAM J. Matrix Analysis and Applications* 13: 212–238 (1992).

[3] I.S. Duff and J.K. Reid, *MA27 – A set of Fortran subroutines for solving sparse symmetric sets of linear equations*, Tech. Rep. AERE R10533, Rutherford Appleton Laboratory, London, 1982.

[4] L. Giraud and R. Tuminaro, Domain Decomposition Algorithms for the Drift-Diffusion Equations, in *Sixth SIAM Conference on Parallel Processing for Scientific Computing*, R. Syncovec, D. Keyes, M. Leuze, L. Petzold, and D. Reed (eds.), Philadelphia, 1993, SIAM, pp. 719–726.

3.4 Efficient linear solvers for large sparse complex symmetric non-Hermitian matrices (V. Frayssé, L. Giraud and S. Gratton)

We have implemented the classical restarted GMRES(m) method introduced by Y. Saad and M. Schultz [2] and the flexible variant FGMRES(m) proposed by Y. Saad [3]. The codes are available in single and double precision for real or complex arithmetic. They are based on a reverse communication scheme and the recommended stopping criterion is the backward error.

The development of these codes has been realized in the framework of an industrial contract between CNES and CERFACS. CNES provided two matrices arising in electromagnetism problems: both matrices turned out to be complex, dense and non-Hermitian. Their size was about 3000×3000. On these examples, neither GMRES nor FGMRES was competitive with the classical dense direct solvers available in LAPACK. Results are summarized in [1].

Further work is continuing on these codes in order to improve its performance on sequential computers and to allow to run it on a distributed memory environment. One of the target applications is an astrophysical problem (see Chapter 6) which results in a generalized eigenproblem. The shift-and-invert technique used for the eigensolver requires the solution of large and sparse linear systems. Preliminary tests show that GMRES seems to be an appropriate alternative to the direct band linear solver previously used.

[1] V. Frayssé, L. Giraud, and S. Gratton, (1996), Solveurs linéaires itératifs pour la résolution de systèmes complexes non hermitiens creux de grand taille, Contract Report FR/PA/96/34, CERFACS.

[2] Y. Saad and M. Schultz, (1986), GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.*, **7**, 856–869.

[3] Y. Saad, (1993), A flexible inner-outer preconditioned GMRES algorithm, *SIAM J. Sci. Comput.*, **14**, 461–469.

3.5 Solution of general linear systems of equations using block Krylov based iterative methods on distributed computing environments (M. Arioli, A. Drummond, I. S. Duff and D. Ruiz)

We studied the implementation of block Krylov based iterative methods on distributed computing environments for the solution of general linear systems of equations. First, we studied potential implementations of the classical conjugate gradient (CG) method on parallel environments. From the family of conjugate direction methods, we studied the Block Conjugate Gradient (Block-CG) method which is based on the classical CG method. The Block-CG works on a block of s right-hand side vectors instead of a single one as is the case of the Classical CG, and we studied the implementation of the Block-CG on distributed environments.

The complexity of the Block-CG method is higher than the complexity of the classical CG in terms of computations and memory requirements. We showed that the fact that an iteration of the Block-CG requires more computations than the classical CG makes the Block-CG more suitable for vector and parallel environments. Additionally, the increase in memory requirements is only a multiple of s which generally is much smaller than the size of the linear system being solved.

The Classical CG and Block-CG are suitable for the solution of symmetric and positive definite systems of equations. Furthermore, both methods guarantee, in exact arithmetic, to find a solution to positive definite systems in a finite number of iterations.

For non symmetric linear systems, we studied block row projection iterative methods for solving general linear systems of equations, and we were particularly interested in showing that the rate of convergence of some row projection methods can be accelerated with the Block-CG method. We review two block projection methods, the block Cimmino and the block Kaczmarz method.

Afterwards, we studied the implementation of an iterative procedure based on the block Cimmino method using the Block-CG method to accelerate its rate of convergence on distributed computing environments.

The complexity of the new iterative procedure is higher than the one of the Block-CG method in terms of computations and memory requirements. In this case, the main advantage is the extension of the application of the CG based methods to general linear systems of equations. We have developed a parallel implementation of the block Cimmino method with Block-CG acceleration that performs well on distributed computing environments.
This last parallel implementation opens a study of potential scheduling strategies for distributing tasks to a set of computing elements. We designed a scheduler for heterogeneous computing environments which is currently implemented inside the block Cimmino based solver and can be reused inside parallel implementations of several other iterative methods.
Lastly, we combine all of the above efforts for parallelizing iterative methods with preprocessing strategies to improve the numerical behavior of the block Cimmino based iterative solver.
We made a detailed study of the block conjugate gradient algorithm using various strategies for distributing the computation. Each strategy has its advantages and disadvantages but we found that, on most platforms and on most problems, it was best to use a master-slave distributed implementation in which the master performs an initial scheduling and distributes the tasks to the slaves. Thereafter the master is only used for global operations like computing inner products and testing for termination. The alternatives that we compared this strategy with were only to distribute matrix-vector multiplies and one where there is no master process but more redundant computation. We used the master-slave distributed implementation of the block conjugate gradient algorithm as the preconditioner for a distributed block Cimmino algorithm.
We have found it very beneficial to allow the scheduler to group closely related tasks on the same processor and have also experimented with different strategies for the initial subdivision of the equations for the block Cimmino algorithm. Choosing the best strategies in each case, we obtain impressive speedups, as is illustrated in Table 3.2 below.

Number of processors	1	2	4	8	10
Time	8.3	4.3	2.5	1.4	1.0
Speedup	-	1.9	3.3	5.9	8.3
Efficiency	-	0.95	0.83	0.74	0.83

Table 3.2: Time in seconds on IBM SP2 for solution of problem SHERMAN4 from the Harwell-Boeing Collection.

[1] M. Arioli, A. Drummond, I.S. Duff, and D. Ruiz, (1995). A parallel scheduler for block iterative solvers in heterogeneous computing environments, in *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing*, (Ed. by David H. Bailey et al.). SIAM, Philadelphia, 460–465.

[2] M. Arioli, A. Drummond, I.S. Duff, and D. Ruiz, (1995). Parallel block iterative solvers for heterogeneous computing environments, in *Algorithms and Parallel VLSI Architectures III*, (Ed. by M. Moonen and F. Catthoor). Elsevier, Amsterdam, 97–108.

[3] M. Arioli, I.S. Duff, J. Noailles, and D. Ruiz, (1992). A block projection method for sparse matrices, *SIAM J. Sci. Stat. Comput.* 13, 47–70.

[4] M. Arioli, I.S. Duff, D. Ruiz, and M. Sadkane, (1995). Block Lanczos techniques for accelerating the block Cimmino method, *SIAM J. Sci. Comput.* 16, 1478–1511.

[5] L.A. Drummond, (1995). Solution of general linear systems of equations using block Krylov based iterative methods on distributed computing environments, PhD Thesis. Report TH/PA/95/40, CERFACS.

[6] I.S. Duff, (Editor). (1994). Numerical Analysis Group. Progress Report. January 1991 - December 1993. Report RAL 94-062, RAL.

3.6 Sparse approximate inverse preconditioners (M. Benzi, M. Tůma and L. Giraud)

It is generally agreed that the problem of parallel preconditioning for solving large linear systems of equations is one of the main bottlenecks faced by large-scale modeling and simulation projects. State-of-the-art iterative methods, like Bi-CGStab and GMRES, which are frequently used in such simulations, must be preconditioned in order to be effective. Furthermore, parallel supercomputers now provide unprecedented computational power and offer the possibility of performing new and extremely large-scale simulations. Unfortunately, some of the best general-purpose preconditioners, like incomplete factorization (ILU) methods, are inherently sequential in nature, and are difficult to implement efficiently on current vector and parallel architectures. As a result, there is currently a growing interest in alternative solutions to this problem. At present, one of the most important alternatives to ILU is explicit preconditioning with sparse approximate inverses. The parallel implementation of these techniques is conceptually straightforward, and there is abundant evidence that approximate inverse preconditioners are able to deliver convergence rates close to those achieved by the best serial preconditioners for many problems; see, e.g., [1], [2], [4], [5], [6]).

With this approach, a sparse matrix $M \approx A^{-1}$ is explicitly constructed and used as a preconditioner for $Ax = b$. Applying an explicit preconditioner M within a step of the conjugate gradient method reduces to computing a matrix-vector product with M, an operation which offers considerably more scope for vectorization and parallelism than the triangular solves involved in ILU-type methods. It should be emphasized that there exist several distinct approaches to constructing an approximate inverse. One of the most popular methods consists in finding a sparse matrix M such that $||I - AM||_F$ is small in some sense; see, e.g., [5] and [6]. The construction of the approximate inverse can be done in parallel, since the computation of each column of M is independent from the others. This approach, commonly referred to as SPAI, has been quite successful in practice. However, there are several drawbacks to it. The construction of the preconditioner can be very time-consuming, and a large number of nonzeros in the approximate inverse is sometimes required in order to obtain a good rate of convergence. In addition, this preconditioner is not suitable for use with the conjugate gradient method, since in general M will not be symmetric and positive definite even if A is, and even ensuring that M is nonsingular is not at all easy.

An alternative approach is to construct an approximate inverse in factorized form, $M = M_U M_L$, where M_U and M_L are, respectively, upper and lower triangular. In other terms, $M_U \approx U^{-1}$ and $M_L \approx L^{-1}$ where L and U are the triangular factors of A. Clearly, it is trivial to ensure that M is nonsingular. If A is symmetric positive definite, then $M_U = M_L^T$ and the preconditioner is automatically symmetric and positive definite. Furthermore, a factorized approximate inverse contains more information than a standard approximate inverse with the same number of nonzeros, and this often translates in faster rates of convergence for the corresponding preconditioned iteration. It is important to realize that the approximate inverse factors can be computed directly from A, without any knowledge of the triangular factors L and U. An algorithm to accomplish this was proposed and developed in [1] for symmetric and positive definite matrices and in [2] for general matrices. Unlike the method in [7], this approach is well suited for matrices with arbitrary sparsity pattern. This method gives good convergence rates, often better than those obtained with SPAI, while requiring less storage for the preconditioner. On the other hand, while the preconditioner can be applied in parallel, its construction is rather sequential in nature. However, this computation was found to be typically one or two orders of magnitude faster than SPAI on a scalar or vector computer.

Results for an implementation on the vector computer CRAY C90 at Meteo France are reported in [3]. These results show that the sparse approximate inverse preconditioner results in faster execution than standard incomplete Cholesky (IC) preconditioning on vector machines even for problems for which the rate of convergence is inferior to that of IC.

Current work (with M. Tůma of the Czech Academy of Sciences, Prague) is focused on a systematic comparison of a number of different sparse approximate inverse preconditioners on a wide range of problems. The comparison will also include polynomial and incomplete factorization preconditioners.

In addition, at the end of 1996 a collaboration was started between AEROSPATIALE and M. Benzi and L. Giraud from CERFACS concerning the use of sparse approximate inverse preconditioners for large, dense, complex symmetric linear systems arising from boundary integral formulations of electromagnetism problems. For such problem, it was found that good results can be obtained with a SPAI-like approach in which the nonzero pattern of the approximate inverse can be determined a priori.

[1] M. Benzi, C. D. Meyer and M. Tůma, (1996). A sparse approximate inverse preconditioner for the conjugate gradient method, *SIAM J. Sci. Comput.* 17, 1135–1149.

[2] M. Benzi and M. Tůma, (1996). A sparse approximate inverse preconditioner for nonsymmetric linear systems, CERFACS Technical Report TR/PA/96/15. To appear in *SIAM J. Sci. Comput.*

[3] M. Benzi and M. Tůma, (1996). Approximate inverse preconditioning for the conjugate gradient method on a vector computer, in *Proceedings of the Prague Mathematical Conference 1996*, 29–34.

[4] E. Chow and Y. Saad, (1996). Approximate inverse preconditioners for general sparse matrices, University of Minnesota Supercomputer Institute Research Report UMSI 94/101. To appear in *SIAM J. Sci. Comput.*

[5] N. I. M. Gould and J. A. Scott, (1995). On approximate-inverse preconditioners, Rutherford Appleton Laboratory Technical Report RAL-95-026. To appear in *SIAM J. Sci. Comput.*

[6] M. Grote and T. Huckle, (1994). Parallel preconditioning with sparse approximate inverses. Stanford University Manuscript SCCM-94-03. To appear in *SIAM J. Sci. Comput.*

[7] L. Yu. Kolotilina and A. Yu. Yeremin (1993). Factorized sparse approximate inverse preconditioning I. Theory, *SIAM J. Matrix Anal. Applic.* 14, 45–58.

3.7 An investigation into the exploitation of parallel computer architectures in shape optimization (P. Hubsch)

The application of parallel processing to the problem of shape optimization is investigated. The objective function considered is the weight (or the volume) of the structure and the constraints taken into account are limit values for the effective stress as well as geometric constraints designed to safeguard the functionality of the analyzed structure. The finite element method (FEM) is used to compute the stress field throughout the structure in order to evaluate the stress constraints.
Because most of the computational cost is incurred through the analysis step, parallelization of the FE procedure should lead to a substantial reduction of the total CPU time needed. Within the FE procedure, all calculations which can be performed on an element level (e.g. stiffness calculation, strain and stress calculation) parallelize naturally. However, most of the CPU time in a typical FE procedure is spent with the solution of the equation.
Direct equations solvers do not parallelize very well in general and therefore an iterative solution scheme, the conjugate gradient method, was used. Unlike in most parallel FE procedures, the analyzed domain was not decomposed along element boundaries. The calculation of the stiffness matrix was shared between the processors according to element number, but the resulting stiffness terms were stored and processed according to equation number. This implies that an equation is not necessarily solved on the same processor it is generated on, a process which generates some additional message passing. However, this traffic is local and therefore not particularly costly. The advantage of the procedure is that the communication between neighboring processors of partial results along the interface of sub-domains typically associated with parallel iterative FE solvers is saved.
Some encouraging results on benchmark examples were achieved, with the parallel procedure breaking even with a sequential procedure based on a state of the art direct solver at about 12000 degrees of freedom.

3.8 Iterative solution of a sparse overdetermined linear system of equations (L. K. Lundin)

We have studied how to solve

$$\min_x ||Ax - b||_2 \tag{3.1}$$

by use of the Preconditioned Conjugate Gradient method for Least Squares (PCGLS) on parallel distributed computers. Problem (3.1) is to be solved a number of times (with varying b and constant A) as a part of the solution procedure for a certain 3D CFD problem. It is therefore important that a solution can be quickly computed. The following summarizes what we have done to achieve this goal:

Discovery of two levels of structure in A By experiments with different reorderings of A, it was discovered that A can be written as a sum of four Kronecker products

$$A = \sum_{i=1}^{4} E_i \otimes A_i \tag{3.2}$$

It was furthermore established that each of the four submatrices A_i has a blocked structure, where each block is either a (short) column vector or a circulant matrix.

Exploitation of the structure of A Since A appears in PCGLS solely in matrix-vector products, a matrix-vector product exploiting the structure of A was introduced in a previously written FORTRAN implementation of PCGLS. This led to a drastic reduction of the storage requirements for PCGLS. If N denotes the spatial resolution in each of the three dimensions of the underlying CFD problem, then the number of nonzeros in A is $O(N^3)$. The nested structure of A was exploited so that only the distinct values of A need to be stored and the storage requirement was reduced to $O(N)$. For $N = 70$ (corresponding to 10^6 columns in A) the storage was reduced from more than 190Mb to less than 10kb. The new matrix-vector product has furthermore proved to be 6-8 times faster (depending on the computer architecture) than the general matrix-vector product that it replaced.

Preconditioning A number of preconditioners have been tested. Some of these preconditioners were generated by software written by Benzi and Tůma [1]. The preconditioners were variants of incomplete orthogonalization schemes, which are often used with CGLS. Unfortunately these preconditioners were so dense that the total amount of computation was larger than if a diagonal preconditioner was used. The reason was mainly that the effectiveness of these preconditioners relies on a minimum degree reordering of the columns of A, something that in our case would destroy the structure of A. The most successful preconditioner so far has been an incomplete factorization of the block diagonal part of $A^T A$, tested in MATLAB. Depending on the geometry of the CFD problem the reduction in total amount of computation was from 2% to 15% relative to a diagonal preconditioner, provided that a proper drop tolerance was used. Since the reduction was very sensitive to the choice of drop tolerance, and since no heuristic for the optimal choice of drop tolerance could be found, this preconditioner was not implemented. At this point PCGLS is therefore used with a diagonal preconditioner that requires 15% less computations than plain CGLS overall. It is however still our ambition to find a more effective preconditioner.

Parallel implementation Our existing PCGLS implementation was parallelized using explicit communication (MPI), and tested successfully on 6 different parallel distributed computers (the same code has furthermore been tested successfully on 5 different uniprocessors).

Reduction of communication cost The standard formulation of PCGLS has (like the standard formulation of the Conjugate Gradient method) two inner products per iteration. The computation of these inner products on a parallel distributed computer require globally synchronizing communication. In [2] D'Azevedo, Eijkhout and Romine derive a formulation of CG with just one synchronization point; and they describe numerical experiments where the solution time is reduced by 5% to 13%. We applied the same approach to PCGLS, but found that the reduction in communication cost could not compensate for the computational overhead in our alternative formulation of PCGLS.

Performance We conducted various tests on 6 parallel distributed computers. Due to the low storage requirement for A some cachebased architectures (especially an SGI Power Challenge equipped with 16 R8000 processors) exhibited superlinear speed-up, even for systems with more than $1.5 \cdot 10^6$ unknowns. Due to this a wide range of speed-ups have been observed (see graphs). The time to solve (3.1) with $1.5 \cdot 10^6$ unknowns (4 correct decimal digits in the computed solution) was 250 seconds on 64 Cray T3D processors. On all 6 machines the time to perform 100 iterations on the same system was measured. From these measurements the time to solve for 4 correct digits on the 5 other machines were extrapolated. This test was also conducted on two uniprocessor vector architectures, a Cray Y-MP and a Cray C90. The table below summarizes the relative solution time for the various machines.

Machine	No. of processors	Relative solution time
SGI PC R8000	16	0.87
Cray T3D	64	1.00
IBM SP-2 Thin	32	1.30
IBM SP-2 Wide	16	1.32
Cray C90	1	3.44
SGI PC R10000	6	3.55
Meiko CS/2	14	3.70
Cray Y-MP	1	9.31

Parallel Speed-up of dPCGLS

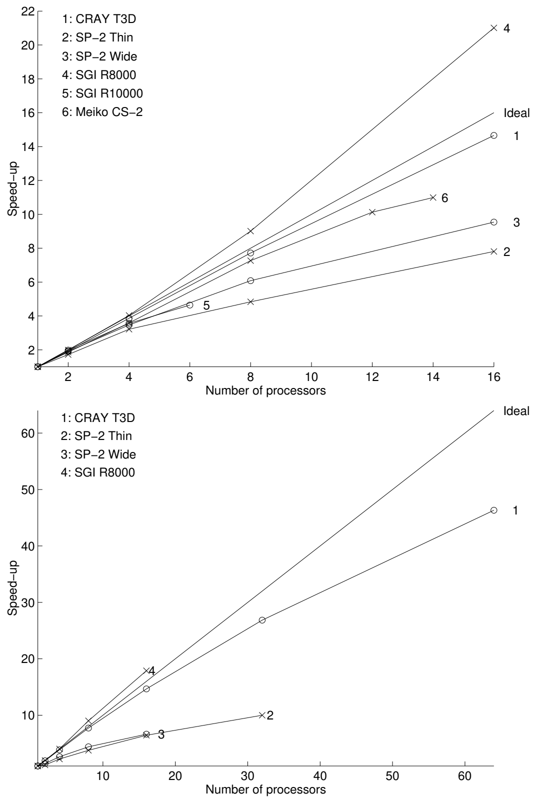

[1] M. Benzi and M. Tůma, A comparison of some preconditioning techniques for general sparse matrices, in S. D. Margenov and P. S. Vassilevski (Eds.), *Iterative Methods in Linear Algebra, II,* IMACS Series in Computational and Applied Mathematics Vol. 3, pp. 191–203, 1996.

[2] E. D'Azevedo, V. Eijkhout and C. Romine, Reducing Communication Costs in the Conjugate Gradient Algorithm on Distributed Memory Multiprocessors, in LAPACK Working Note 56, Tech. report, Computer Science Department, University of Tennessee, Knoxville, TN, 1993.

[3] L. K. Lundin, How to Compute a Rotating Velocity Field, to appear as a techn. report, CERFACS, 1997.

4 Multigrid Methods

Multigrid methods were considered from two points of view. One was as an iterative method for solving linear systems of equations. The other was in the context of numerically simulating flames with and without detailed, finite rate chemistry.
Of interest are the following papers and reports: [6], [8], [7], [9], [1], [2], [3], [4], [5], and [10].

4.1 Cache based multigrid methods (C.C. Douglas)

We investigate the effect computer caches can have on tailored multigrid algorithms. This topic is timely since larger caches with many cache lines are becoming quite common on computer processors.
By the term *cache*, we mean a fast memory unit closely coupled to the processor. In the interesting cases, the cache is further divided into *cache lines*, which hold copies of contiguous locations of main memory. The cache lines may hold data from quite separate parts of main memory.
For example, a cache with $.25 - 4$Mb of cache memory with .25Kb cache lines allows for many interesting possibilities, whereas a processor with 8Kb of cache memory with 2Kb cache lines does not. This is because in the first case many separate lines (or short vectors) can be held in cache simultaneously, whereas in the second case only 4 cache lines are available, which is insufficient in many situations.
Multigrid algorithms combine a number of operations in order to work. These include iterative methods (typically relaxation methods), residual computation, projection of residuals onto a coarser grid, and interpolation of corrections onto a finer grid. These are typically programmed as separate routines, which makes the components easy to replace and modify.
However, a number of components re-use data in a manner which is suitable for algorithms which are *cache aware*. We investigated which of these are really worth combining and which are not. This depends on how well or badly the cache system on a particular machine is designed.

4.2 MPI based multigrid solvers (C.C. Douglas)

At the Copper Mountain multigrid meeting, I talked briefly one evening about what I thought was new and interesting in parallel multigrid. I mentioned MPI, a message passing standard developed jointly by academics and computer companies. It appears to be catching on with the computer companies and there is a free version available from Argonne (in pub/mpi or pub/mpi/misc on the machine info.mcs.anl.gov) that runs on almost everything.
I put my two dimensional MPI based multigrid solvers on MGNet, which resides at Yale University and CERFACS, in the directory mgnet/Codes/douglas. This is freely available on the Internet, either by anonymous ftp or through a web browser. These solvers solve constant coefficient and variable coefficient elliptic partial differential equations on simple domains. They can be imbedded in time dependent solvers. Work is under progress with Luc Giraud to extend these solvers to three dimensional problems that are more complicated.

4.3 Flame simulation (C.C. Douglas, A. Ern and M.D. Smooke)

Detailed computer modeling of chemically reacting systems has sparked extensive interest among combustion scientists thanks to its ability to mimic experiments in several flame configurations and to increase our fundamental understanding of the underlying physical processes. While the ultimate goal in combustion modeling is the solution of unsteady, high-speed, three-dimensional, turbulent reacting systems, gas-jet laminar diffusion flames constitute a problem of practical and fundamental importance that can be analyzed in detail with current computer resources. The gas-jet diffusion flame is the basic element

of many combustion systems, such as diesel engines, ram jets, gas turbines, and industrial furnaces. Furthermore, the fundamental understanding of laminar diffusion flames plays a central role in modeling turbulent diffusion flames through the concept of laminar flamelets and in understanding the processes by which pollutants are formed.
The equations governing laminar diffusion flames are extremely difficult to analyze. Indeed, these equations are strongly coupled due to the interaction between fluid dynamics, transport, and chemistry processes. Moreover, as a result of the exponential dependence of the chemistry source terms on the temperature, these equations are highly nonlinear. Furthermore, high spatial activity regions (characterized by steep fronts and sharp peaks in the components of the computed solution) require disparate length scales to be resolved numerically. Finally, and most importantly, the large number of chemical species that must be included among the dependent unknowns for detailed investigation of the underlying chemical processes renders the computational cost of most combustion problems extremely high.
It is the purpose of the present study to present a numerical model of a laminar diffusion flame that yields accurate numerical solutions at a reasonable computational cost. The physical configuration is based on an inner cylindrical fuel stream surrounded by a coflowing oxidizer jet. The inlet velocities are high enough to produce a lifted flame with a triple flame ring structure at its base. Computationally, we solve the total mass, momentum, energy, and species conservation equations coupled with complex transport and finite rate chemistry models. We describe the vorticity–velocity formulation which casts the governing equations in an efficient form to be analyzed numerically. The numerical solution procedure involves damped Newton iterations, Krylov-type linear system solvers, and adaptive grid refinement.
The serial and parallel solution methods are described and the numerical results are presented using an IBM SP2. In particular, the parallel algorithms scale very well as a function of the number of processors.

Machine	Processors	Minutes
RS6000-580	1	602.6
SP2	8	40.2
SP2	16	20.9

[1] C. C. Douglas and J. Douglas, (1993), A unified convergence theory for abstract multigrid or multilevel algorithms, serial and parallel, *SIAM J. Numer. Anal.*, **30**, 136–158.

[2] C. C. Douglas, J. Douglas, and D. E. Fyfe, (1994), A unified multigrid theory for non-nested grids and/or quadrature, *East-West J. Numer. Math.*, **2**, 285–294.

[3] C. C. Douglas, A. Ern, and M. D. Smooke, (1994), Numerical simulation of laminar diffusion flames, *SIAM News*, **27**, 1, 12–13, 17.

[4] C. C. Douglas, A. Ern, and M. D. Smooke, (1995), High performance computing and numerical simulation of flames, Tech. Rep. TR/PA/95/34, CERFACS, Toulouse.

[5] C. C. Douglas, A. Ern, and M. D. Smooke, (1996), Multigrid solution of flame sheet problems on serial and parallel computers, *Paral. Algor. Appl.* To appear.

[6] C. C. Douglas, (1994), A generalized multigrid theory in the style of standard iterative methods, In *Multigrid Methods IV*, P. W. Hemker and P. Wesseling, eds., Basel, Birkäuser, 19–34.

[7] C. C. Douglas, (1995a), Caching in with multigrid algorithms: problems in two dimensions, Tech. Rep. TR/PA/95/15, CERFACS, Toulouse. To appear in Parallel Algorithms and Applications, 1996.

[8] C. C. Douglas, (1995b), Implementing abstract multigrid or multilevel methods, *Comput. Appl. Math*, **14**, 1–19.

[9] C. C. Douglas, (1995c), A sparse matrix approach to abstract multilevel solvers on serial and parallel computers, Tech. Rep. TR/PA/95/35, CERFACS, Toulouse.

[10] A. Ern, C. C. Douglas, and M. D. Smooke, (1995), Detailed chemistry modeling of laminar diffusion flames on parallel computers, *Inter. J. Supercomputer Applic. High Perfor. Comput.*, **9**, 167–186.

$\boxed{5}$ Eigenvalue Computations

5.1 Symmetric Eigenproblems: Tools (O.A. Marques)

The solutions of

$$Ax = \lambda Bx \tag{5.1}$$

where A and B are $n \times n$ matrices, x is a non null vector and λ is a scalar, are of great importance in a variety of disciplines. The above relation defines a *generalized eigenproblem*, a *standard eigenproblem* is obtained when B is equal to the identity matrix. In many cases, eigenproblems are associated with fundamental characteristics of differential and integral operators describing a physical phenomenon. The dimensions of A and B can thus reach dimensions of tens of thousands due to the complexity, the level of the discretization of a continuous problem or the precision required for the results. Real applications commonly require the computation of only a small subset of the n *eigenpairs* (λ, x): the *eigenvalues* λ of interest are those lying either in one of the extremities of the spectrum or in an interval $[\xi_1, \xi_2]$, together with the corresponding *eigenvectors* x. However, even the determination of a few solutions is a time consuming task, justifying a search for efficient algorithms.
The objective of the present activity is to improve and study efficient techniques to solve eigenproblems associated with symmetric operators. In particular, a public domain package, BLZPACK, has been developed [4]. The package is a Fortran 77 implementation of the block Lanczos method, in combination with a modified partial reorthogonalization and a selective orthogonalization strategies, for the computation of eigenvalues λ and eigenvectors x of the standard problem $Ax = \lambda x$ or the generalized problem $Ax = \lambda Bx$, where A and B are real symmetric matrices (in this case, a positive definite linear combination of A and B must exist). In the general case, the Lanczos algorithm requires calculations involving the matrices A, B and sets of vectors until convergence for the required solutions is reached. However, in the BLZPACK implementation, each time such calculations have to be performed the control is returned to the user, which means that A and B do not have to be passed as arguments. Since the matrices A and B are kept outside the code, the user is free to handle with them. Actually, the package is tailored to a class of applications for which at least one "inversion" of the operator $A_\sigma = A - \sigma B$ is feasible, where σ is a real scalar. With such an inverted operator the eigenvectors are preserved while the eigenvalues are remapped. Eigenvalues lying in a range of interest can be then set apart from the remaining eigenvalues, leading to better converge rates for the wanted solutions. In practical problems the inversion is replaced by a factorization $A_\sigma = LDL^T$, where L is a lower unit triangular matrix and D is a direct sum of 1×1 and 2×2 pivot blocks, leading to the solutions of systems of linear equations.

5.2 Symmetric Eigenproblems: Applications

This section describes a number of applications making use of the aforementioned Lanczos package.

The study of internal motions of proteins. This work is a collaboration with Yves-Henri Sanejouand (*Laboratoire de Physique Quantique, IRSAMC, Université Paul Sabatier*, Toulouse). The study of internal motions of molecules provides useful insights into the large amplitude conformational changes molecules experience during chemical reactions. In particular, theoretical *normal mode analyses* of proteins taking into account the lowest frequencies help to predict the nature of their conformational changes. As shown in [2] for *citrate synthase*, for example, theoretical results match those obtained with X-Ray crystallography. The normal mode method consists in obtaining an approximation for the motion history through the superposition of collective variables, namely the *normal mode coordinates*. These coordinates correspond to the solutions of a real symmetric eigenproblem whose dimension is equal to three times the number of atoms. A set of cases studied is given in [3]; the ribbon representation of *citrate synthase*, for instance, is plotted in Figure 5.1.a.

The ground state of atoms and clusters on a 2D grid. This work is part of the Computational Chemistry project at CERFACS (see Chapter 8 of this report for details) and is pursued since fall 1995 by C. A. Ullrich (CERFACS) and E. Suraud (*Laboratoire de Physique Quantique, IRSAMC, Université Paul Sabatier*, Toulouse). They developed a program package to compute the electronic ground state for atoms and simple metal clusters on a two-dimensional grid. For these N-electron systems, one has to solve (in a self-consistent manner) the so-called static Kohn-Sham equations

$$\left(-\frac{\nabla^2}{2} + v_s[\rho](\vec{r})\right) \varphi_j(\vec{r}) = \varepsilon_j \varphi_j(\vec{r}) \quad , \qquad j = 1, \ldots, N \tag{5.2}$$

featuring an effective local potential $v_s[\rho](\vec{r})$ depending on the density $\rho(\vec{r}) = \sum_{j=1}^{N} |\varphi_j(\vec{r})|^2$. The eigenproblem (5.2) is discretized on a finite rectangular grid using a finite-difference scheme in cylindrical coordinates. The resulting eigenvectors $\varphi_j(\vec{r})$, i.e. the electronic ground-state orbitals, serve as initial states for the investigations of dynamical (time-dependent) properties of the system considered. Figure 5.1.b shows how the orbital energy eigenvalues ε_j of a Na_8-cluster in the spherical jellium model approach their "exact" values for decreasing grid spacings $\Delta \to 0$. The corresponding total number of gridpoints (i.e., the dimension of the eigenproblem) increases from 820 ($\Delta = 1.0$) to 33930 ($\Delta = 0.15$).

A pseudoconsistent mass matrix for a triangle finite element. This work is a collaboration with R. M. Sydenstricker, A. L. G. A. Coutinho and L. Landau (Department of Civil Engineering, COPPE/UFRJ, Rio de Janeiro, Brazil). Many finite elements based on Reissner-Mindlin plate bending theory with discrete constraints have been developed in the past few years. This approach avoids the C continuity required by the classical Kirchhoff plate theory. However, in those elements the shape functions for deflections (or their derivatives) are not integrated. Particularly in DKT (discrete Kirchhoff triangle) and in DST (discrete shear triangle) elements, the lateral deflections are defined only on the boundaries. In [6] the transverse displacement for these elements is briefly discussed and a simple approach to define a pseudoconsistent load vector and mass matrix is presented.

Shape optimization and free vibration problems. This work is a collaboration with S. Bastos (Dept. of Civil Engineering, University of Wales, Swansea), who visited CERFACS in August, 1995. S. Bastos works with shape optimisation and her Ph.D. dissertation was about "Shape Optimisation of Shells Under Static and Free Vibration Conditions". A set of experiments have been performed with her code and BLZPACK. Real applications have been planned to be studied at Swansea.

5.3 Numerical Experiments with the HR and XHR Algorithms

This work was a collaboration with Prof. B. N. Parlett (Dept. of Mathematics, University of California, Berkeley) and Petter Wiberg (Department of Computing Science, Umeå University, Sweden), who visited CERFACS in 1996. The elegant structure of an unsymmetric tridiagonal matrix is usually destroyed when its eigenvalues and eigenvectors are computed by the QR algorithm. The HR algorithm, which was motivated by QR, is an attempt to exploit and respect the tridiagonal structure. Nevertheless, it can break down. Using a generalized Cholesky decomposition for a pseudo-symmetric matrix, an extended HR algorithm (XHR) can be defined [1]. Such an approach never breaks down. The objective here was to investigate the numerical behaviour of the XHR algorithm in the computation of eigenvalues of nonsymmetric tridiagonal matrices derived from real applications. The results of this study are reported in [7].

[1] Z. S. Liu, (1992), On the Extended HR Algorithm, Tech. Rep. PAM-564, Center for Pure and Applied Mathematics, University of California, Berkeley, USA.

[2] O. A. Marques and Y.-H. Sanejouand, (1995a), Hinge-Bending Motion in Citrate Syntase Arising from Normal Modes Calculations, *Proteins*, **23**, 557–560. Also CERFACS Report TR/PA/95/31, Toulouse, France.

[3] O. A. Marques and Y.-H. Sanejouand, (1995b), Protein Motions through Eigenanalyses: A Set of Study Cases, Tech. Rep. TR/PA/95/32, CERFACS, Toulouse, France.

[4] O. A. Marques, (1995), BLZPACK: Description and User's Guide, Tech. Rep. TR/PA/95/30, CERFACS, Toulouse, France.

[5] O. Marques, (1996), Eigensolvers and applications in finite elemnt analyses, In *Advanced Computational Methods in Structural Mechanics*, M. Papadrakakis and G. Bugeda, eds., International Centre for Numerical Methods in Engineering (CIMNE), Barcelona, Spain, 66–79.

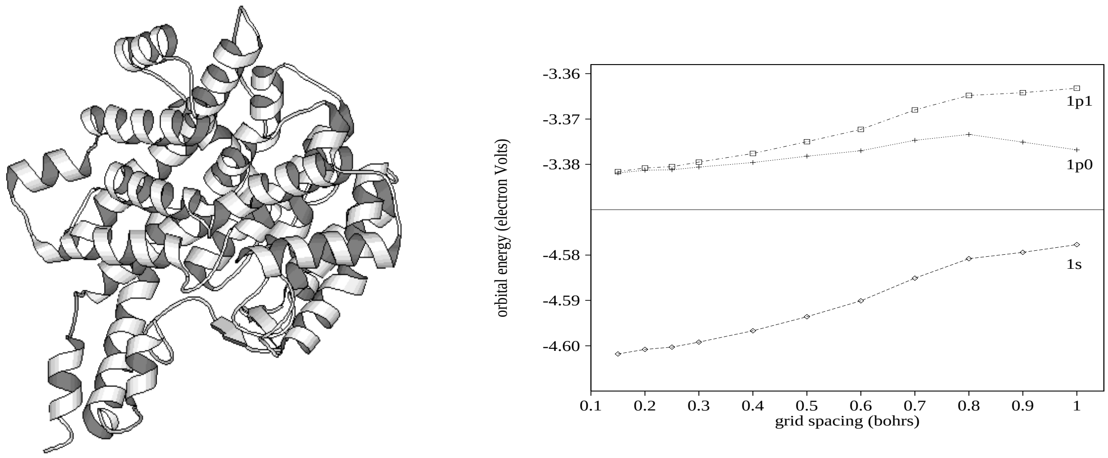

(a) (b)

Figure 5.1: a) ribbon representation for *citrate synthase* (8528 atoms), b) Na_8 orbital energies for different grid spacings.

[6] R. M. Sydenstricker, A. L. G. A. Coutinho, L. Landau, and O. A. Marques, (1995), Pseudoconsistent Load Vector and Mass Matrix for the Discrete Kirchhoff Triangle and the Discrete Shear Triangle Elements, *Comm. in Numer. Methods in Eng.*, **11**, 317–330. Also CERFACS Report TR/PA/95/44, Toulouse, France.

[7] P. Wiberg, (1996), A Study of the HR and Extended HR Methods for the Standard Eigenvalue Problem (Master's Thesis), Tech. Rep. UMINF 96.16, Department of Computing Science, Ume University.

6 Qualitative Computing

Group members: Françoise Chaitin-Chatelin, Valérie Frayssé, Serge Gratton, Arnaud Guermeur, Ali Harrabi, Abderrazak Ilahi, Alan McCoy, Jean-Christophe Rioual, Vincent Toumazou and Elisabeth Traviesas. Collaboration with Thierry Braconnier.

The central question studied in the Qualitative Computing Group concerns the reliability of numerical software and methods in finite precision arithmetic. This problem has been addressed through different aspects including a theory of computability in finite precision, convergence of iterative methods in finite precision, perturbation theory (conditioning and backward error analysis) in linear algebra, nonnormal eigenproblems (software for large nonhermitian matrices and tools for the investigation of spectral sensitivity). Our research is encouraged and enriched by the applications on physical problems coming from both the academic and industrial worlds.

It is well-known that a computation becomes unstable (and in particular computer results differ significantly from the **exact** ones) in the neighbourhood of a singularity. And a singularity is not generic : perturbations generated by the finite-precision arithmetic are such that one always solves a regular problem. However the properties of computations are ruled by the underlying singularity, and the distance to the mathematical singularity, *as viewed by the computer*, is an essential parameter to interpret the computation. The role of singularities should not be looked upon as a mathematical oddity, or a modelling clumsiness : they indeed play an essential role in the area of theoretical physics concerned with reduction of theories, and underlie some of the most difficult and intensively-studied problems in physics today. On the other hand, singularities of the same mathematical nature may have more or less influence on the computation, depending on their condition number, themselves under the influence of some parameters of the problem, such as the departure from normality for a matrix (related to a physical parameter such as the Reynolds number in CFD, the resistivity in MHD, etc). This influence, which is at first local, can gradually become global as the relevant parameters evolve. This results in the emergence of new collective phenomena when some singularities which are separated in exact arithmetic, are nevertheless viewed as connected in finite-precision computation - as well as in the physical reality. This is the essence of the new trend in the interpretation (based on **pseudo** rather than exact spectra) of some divergence phenomena in CFD, for example. The standard tool for investigating the stability properties of a solution process is linear perturbation theory, but its domain of validity can be below the natural threshold imposed by machine precision ($\sim 10^{-15}$ in double precision) in difficult cases. There is presently no available theory which describes the effects of large perturbations (beyond linear asymptotics) on the solution of nonlinear equations. One possibility to circumvent this theoretical gap is to use the computer to experiment with the effects of instability by working on samples of randomly perturbed problems. From a statistical analysis of the samples of data and solutions computed by means of a reliable algorithm, one can retrieve essential characteristics of the singularities as they are seen by the computer.

Qualitative Computing is proposed as a novel approach to scientific computation with finite-precision arithmetic. Some of its salient features are :

1. study of the stability and the reliability of numerical methods, specifically in the neighbourhood of singularities,
2. design of robust codes for computing some prescribed eigenelements of very large and nonnormal matrices (Arnoldi-Tchebycheff method),
3. influence of high nonnormality on the reliability of iterative methods in linear algebra,
4. identification of physical applications giving birth to highly nonnormal matrices, in CFD in particular,

5. software tools for assessing the quality (reliability, robustness) of numerical methods and software.

PRECISE (**PR**ecision **E**stimation and **C**ontrol **I**n **S**cientific and **E**ngineering computing) is a toolbox under MATLAB which offers a variety of software tools to assess the quality of numerical methods and software. It includes conditioning, backward errors, arithmetic / numeric quality index, distance to singularity and dangerous border, perturbed spectra, spectral portraits, very large eigenvalue computations. The heart of our software production is eigenvalues.

- ARNCHEB is a Fortran 77 code for the computation of some eigenvalues of large, sparse and nonnormal matrices. It implements the Arnoldi-Techbycheff method. Two versions are available, one for real matrices and one for complex matrices.
- A parallel code has been designed for the computation of spectral portraits for distributed memory machines using PVM. The spectral portrait is the tool of choice for getting an insight on the spectral sensitivity of highly nonnormal matrices.

The theoretical core of our research is described in the book *Lectures on Finite Precision Computations*, published by SIAM in April 1996 [13]. The question of the reliability of software in finite precision attracts a sustained interest in the community of applied numerical analysis. This is witnessed by the success of the minisymposium entitled "Influence of high-nonnormality on the reliability of iterative methods in computational linear algebra", that we organized for the IMACS'95 Conference on iterative methods (Blagoevgrad, Bulgaria, June 17-20 1995). We were also deeply involved in the organization of the second workshop on "Eigenvalues and beyond" of the International Linear Algebra Year where the problem of the spectral stability has been often addressed. These seminars ended with more opened questions than solutions, specifically on the role of non-normality in numerical computations. This makes our field even more challenging.
From the industrial point of view, our work has resulted in our participation in the HPCN EU PINEAPL project. In collaboration with the prime partner NAG, we will bring parts of the software PRECISE into an industrial-quality software. This 3-year project started effectively on January 1996.

We detail now some points that have been particularly investigated during the years 1995 and 1996.

6.1 Convergence of iterative methods in finite precision

In the literature of Linear Algebra, a lot of work is devoted to the convergence analysis of iterative methods. In many cases, one manages to prove the theoretical convergence of a method, but the behaviour of such methods on a computer with finite precision arithmetic remains in many cases an open question. We decided to start with the analysis of the behaviour, in finite precision arithmetic, of the simplest linear iterative scheme one can think of, that is the successive iteration method (SI) $x_0,\ x_{k+1} = Ax_k + b,\ k \geq 0$ where A is a real or complex matrix of order n and x is a real or complex vector of size n.
In exact arithmetic, the behaviour of (SI) is completely understood ; there is convergence for any x_0 if and only if $\rho(A) < 1$ where $\rho(A)$ is the spectral radius of A.
When (SI) is run on a computer with finite precision arithmetic, then for certain matrices A, the convergence is not guaranteed in practice when $\rho(A) < 1$ is true in exact arithmetic. It is clear that the phenomenon should be attributed to the conjunction of two factors : *i)* the nonnormality of A and *ii)* the finite precision of the computer arithmetic. We perform a straightforward analysis of the convergence of (SI) in finite precision from which we try to understand the subtle interplay between factors *i*) and *ii*) which takes place inside the computer, when the iteration matrix A has a high nonnormality. Why should nonnormality be an issue in finite precision ? Because only nonnormal matrices can display a significant amount of spectral instability. Therefore a small perturbation ΔA on A can result in a large perturbation of the spectrum. When the spectral instability of A is high, it appears that the convergence condition $\rho(A) < 1$ in exact arithmetic may not be robust enough for convergence in finite precision computations.
In [15], we present a detailed analysis of the convergence of successive iterations in finite precision, and we have contrasted it with exact convergence. This has been realized by considering two sequences of iterates, x_k defined with the matrix A, and y_k defined with the Schur form of A. The latter sequence y_k is easier to analyze because the Schur form S is invariant under the finite precision computation. Therefore, the perturbation created on S by the computer arithmetic is exactly triangular. The sequence y_k provides useful guidelines to analyze the general sequence x_k. This enables us to shed some light on the interplay

between nonnormality and finite precision in successive iterations. These results were presented on the occasion of the workshop on *Eigenvalues and Beyond* of the ILAY and published in the corresponding special issue of BIT [15] (Guest Editor F. Chaitin-Chatelin).

A similar study on the convergence of the power method has led to experimental evidence of chaos induced by round-off for large enough nonnormality [8]. This will be presented at the Conference on Computational Science for the 21st Century next May [9].

A very detailed study of the departure from normality which leads to significant improvements of results by Henrici [32] and Ostrowski [39] is presented in [30].

6.2 The Arnoldi-Tchebycheff algorithm for solving large nonhermitian eigenproblems

The code ARNCHEB is the implementation of the iterative Arnoldi-Tchebycheff algorithm for solving large sparse real nonsymmetric [4] and complex nonhermitian [5] eigenproblems in standard or generalized form $Ax = \lambda Bx$. It is a reliable and robust software designed to handle very large and highly nonnormal matrices [10]. It can solve

- the standard problem ($B = I$),
- the generalized problem: the inner linear systems are solved with a direct method (LU) when the matrix size is moderate, or with an iterative method (conjugate gradient when B is complex hermitian positive definite or GMRES otherwise).

Setting $A = I$, then solving $x = \lambda Bx$ allows to compute the eigenvalues which lie in the interior of the spectrum of B. Emphasis is put on the reliability of the code: the Krylov basis is reorthogonalised by the iterative modified Gram-Schmidt algorithm [6, 7]. The stopping criterion is based on the backward error to ensure robustness to high nonnormality [1, 2]. Performance and portability are enhanced by the systematic use of BLAS and LAPACK kernels [3]. An upgraded version of ARNCHEB based on a reverse-communication scheme is now available. Performance tests on distributed memory machine are in progress.
This code has been applied in Astrophysics by M. Rieutord (Observatoire Midi-Pyrénées) and L. Valdettaro (Politecnico di Milano, Italy) [40, 41]. The aim is to compute the oscillation (inertial) modes of an incompressible viscous fluid located between two concentric spheres in rotation. This serves as a model for stellar and planetary interiors. The discretization of the linearised Navier-Stokes equations leads to a generalized eigenproblem $Ax = \lambda Bx$. The matrices A and B are nonhermitian and B is possibly singular. The matrix B is real; the matrix A is real when the problem is axisymmetric, and complex otherwise. All eigenvalues have imaginary parts between $-i$ and i and have negative real parts (the equilibrium is stable). The size of the matrices depends on the resolution of the spectral decomposition i.e. of the product of the number of Chebyshev polynomials and the number of spherical harmonics used to discretize the problem. Both matrices are sparse and banded. The astrophysical case is concerned with very small values of a nondimensionnal parameter, the Ekman number, which measures the importance of the viscous force relative to the Coriolis force. As this parameter decreases, the spatial resolution has to be increased to handle all the scales, which results in large eigenproblems. Eigenvalues of interest are the closest to the imaginary axis. They are obtained via a shift-and-invert technique with several imaginary shifts in the interval $[-i, i]$. Using ARNCHEB, M. Rieutord and L. Valdettaro were able to solve their problem (up to a size of 7×10^4) for an Ekman number of the order of 10^{-8}. Improvements are being done (in particular on storage requirements) in order to be able to achieve Ekman numbers close to 10^{-12} which would be more realistic for the earth core modelling. This would mean a matrix of size 10^6.

6.3 Qualitative determination of Jordan forms and assessment of PRECISE

The computation of a Jordan form of a matrix is mathematically ill-conditioned, and many numerical methods for computing it are facing important numerical difficulties. However, a qualitative approach based on perturbed spectra was shown to be efficient in predicting the Jordan structure on many examples [17]. Important results to justify this approach theoretically have been obtained thanks to the Lidskii perturbation theory and resulted in an improvement of the method. Let A be a given matrix and $A(\varepsilon) = A + \varepsilon P$ where P is a perturbation matrix. The spectrum of $A(\varepsilon)$ can be studied in detail for ε small enough, but usually, the theorems do not specify the threshold value for ε. Thanks to the Lidskii perturbation theorem, we were able to determine this threshold and study its dependency with respect to the matrix structure and the departure from normality [35]. The assessment of PRECISE is developed in [34].

6.4 Spectral portraits

Spectral computations are crucial for the study of stability of physical phenomena. However, when highly nonnormal operators are involved, the computation of the eigenvalues is no longer accurate because of the high spectral instability. The trend today is to study the topological neighbourhood of the eigenvalues in order to get information about the sensitivity of the spectrum under perturbations. Classical tools such as the condition number are only valid up to the first order, thus are not informative enough. We want to consider the pseudo-spectrum of A, i.e., the set of all eigenvalues of the perturbed operators $A + \Delta A$, where $\|\Delta A\|$ is bounded. This amounts to plot the spectral portrait of the matrix A, i.e., to plot the map $z \longmapsto \mathrm{spp}(z) = \|A\| \, \|(A - zI)^{-1}\|$ for z in a given region of the complex plane. For large matrices, the computation of

$$\|(A - zI)^{-1}\|_2 = \frac{1}{\sigma_{\min}(A - zI)} = \frac{1}{\lambda_{\min}^{1/2}(A - zI)^*(A - zI)}$$

cannot be based on the singular value decomposition (SVD) which is too expensive. Two methods for computing $\|(A - zI)^{-1}\|$ via solving a hermitian eigenproblem were implemented [43]. The code is tuned for sparse matrices (Compressed Sparse Row format) and uses routines from BLAS, LAPACK, SPARSKIT and the Harwell subroutine library. The computation of $\mathrm{spp}(z)$ over a prescribed region of the complex plane gives rise to embarrassing parallelism. The evaluation of $\mathrm{spp}(z)$ can be done independently on any point of the grid but the time for evaluating $\mathrm{spp}(z)$ may vary a lot from one point to another. Therefore, we have designed a parallel code targeted for distributed memory machines that implements a load balancing strategy. The grid is divided into blocks which are dynamically allocated to the first processor available. The only required communications are the initial broadcast of the matrix to all slaves, each allocation of a block (master to slave) and the set of values $\mathrm{spp}(z)$ on the block (slave to master). There is neither synchronization nor communication between the slaves.

# proc / size of blocks	2	4	8	12	16	24
Maximal	1.95	3.66	6.73	9.6	11.99	16.32
1	1.64	3.85	7.48	11.07	14.47	20.68

Table 6.1: Speed-up for different number of nodes (with 2 processes per node)

Performance observed on the Meiko CS2 machine located at CERFACS for the code implemented with PVM are reported on Tables 6.1 and 6.2 on the Electromagnetism matrix (size 288×288) and a mesh size of 64×64 [20]. This matrix comes from an electromagnetism problem (submitted by F. Collinot and M. Farès from the Electromagnetism project at CERFACS) which deals with the diffraction of a transverse-magnetic wave by a periodic $2D$ structure. The solution x is computed using an iterative scheme $x_{k+1} = Ax_k + b_k$ where b_k tends to 0. This scheme converges provided the spectral radius of A is strictly smaller than 1. However, Figure 6.1 shows that the spectrum spreads out of the unit disc for relatively small matrix perturbations. This behaviour contributes to explain the reported divergence of this iterative scheme.

# proc / size of blocks	2	4	8	12
1	1.99	3.96	7.81	11.50
32	1.99	3.91	7.47	10.69
128	1.99	3.82	6.99	9.46
Maximal	2	3.75	6.91	9.82

Table 6.2: Speed-up for different number of nodes (with 1 process per node)

Recent work on spectral portraits was presented at the minisymposium on nonnormality of the IMACS'95 Conference on iterative methods. Parallelism issues were emphasized on the occasion of the workshop on Eigenvalues and Beyond of the International Linear Algebra Year.

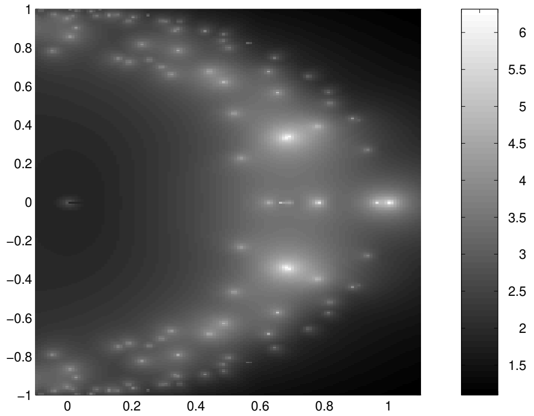

Figure 6.1: Spectral Portrait of the matrix from Electromagnetism.

The notions of pseudospectrum and spectral portrait can be extended to a matrix pair. This enables to analyse the generalized eigenproblem. This extension was motivated by a joint work with J.-R. Angilella and F. Nicoud from the CFD project at CERFACS, concerning the linear stability of a generalized Couette flow with injection and a pressure gradient. The discretization of the mathematical model results in a generalized eigenproblem in which the matrix pair $(A_{n,p}, B_{n,p})$ depend on the number of grid points n and on the order of the Taylor expansion p. The aim of the work was to determine an optimal pair (n, p) realizing the best compromise between the computational cost (increasing with n and p) and the quality of the approximation. The key point lies in using a reference eigenvalue $\lambda_{\rm ref}$ given from an accurate method (here a spectral method developed by Orszag [38]). For different pais (n, p), we determine the ε-pseudospectrum of the pair $(A_{n,p}, B_{n,p})$ where ε is the optimal backward error associated with $\lambda_{\rm ref}$. The smaller the pseudospectrum, the better the approximation scheme. Figures 6.2 and 6.3 respectively show the cases $(n, p) = (1000, 2)$ and $(300, 6)$ and the level curve correspond to the boundary of the ε-pseudospectrum for $\varepsilon = 10^{-11.2}$. Our experiments tend to show that a high order in the Taylor expansion is more crucial than a fine grid point [23]. This conclusion is in accordance with the expectations of the physicists. This application generated research on the perturbation theory for the generalized eigenproblem as described in the next section.

Since the notion of pseudospectrum is extremely useful in Linear Algebra, we expect its extension to operators to be powerful to analyse the convergence of numerical methods which require space and time coupling. It is a difficult subject where little is known so far. Several results have been obtained in this respect [31] for bounded operators in a Banach space. The continuity of the ε-pseudospectrum has been proved under a condition of uniform convergence. Only weaker results can be obtained for the collectively

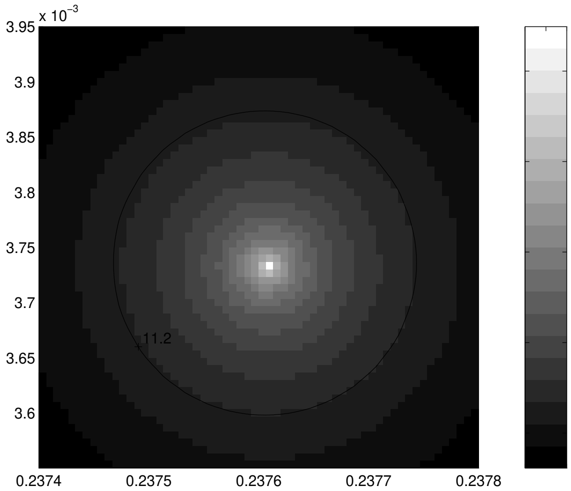

Figure 6.2: Spectral portrait of $(A_{1000,2}, B_{1000,2})$

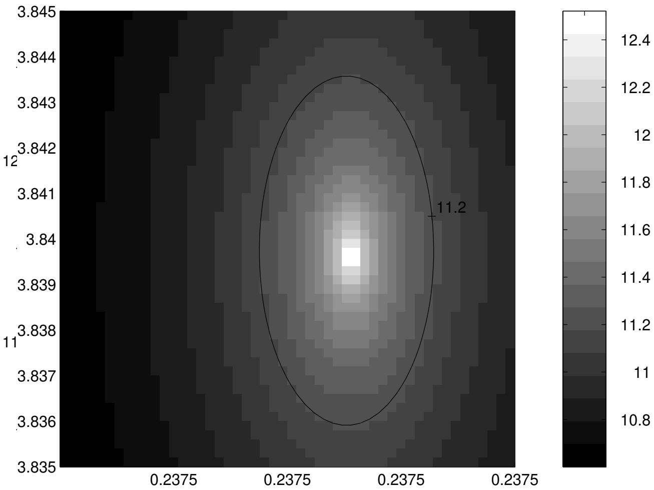

Figure 6.3: Spectral portrait of $(A_{300,6}, B_{300,6})$

compact convergence.

6.5 Backward error analysis in Linear Algebra

In order to assess the numerical reliability of the codes for computing spectral portraits, we were led to study the solution of linear systems of the kind $A^*Ax = b$, where A is a square matrix. These systems arise in the Lanczos process at the heart of the computation of the resolvent norm $\|(A - zI)^{-1}\|_2$. Considering normwise perturbations of A, we proved that the relative structured condition number of the linear system $A^*Ax = b$ is

$$C_{A^*A} = \|x^* \otimes A^{-1} + \left[(Ax)^* \otimes (A^*A)^{-1}\right] L_*\|_2 \frac{\|A\|_F}{\|x\|_2}.$$

See [22] for notations. Depending on x (that is, on the right-hand side b), C_{A^*A} can be of the order of $K_2(A)$ or $K_2^2(A)$ where $K(A) = \|A\|_2\|A^{-1}\|_2$ is the classical condition number of A. Upper and lower bounds for the corresponding backward error have also been derived [22].

The extension of the notion of pseudo-spectrum for a matrix to that for a pair of matrices motivated our work on the normwise perturbation theory for the regular generalized eigenproblem [25, 24]. The analysis follows immediately from that of the standard eigenproblem. When λ is a simple and finite eigenvalue of the matrix pair (A, B), its normwise condition number is

$$K = \frac{(\alpha + |\lambda|\beta)\|x\|\|x_*\|_*}{|\lambda||x_*^* Bx|}.$$

The backward error associated with the approximate eigenpair $(\tilde{\lambda}, \tilde{x})$ is

$$\eta(\tilde{\lambda}, \tilde{x}) = \frac{\|A\tilde{x} - \tilde{\lambda}B\tilde{x}\|}{(\alpha + |\tilde{\lambda}|\beta)\|\tilde{x}\|}.$$

However, in many situations, one seeks eigenvalues only, and the corresponding eigenvectors are not computed: the previous definition of the backward error no longer helps to check the numerical quality of the computed eigenvalue. The notion of *optimal backward error* [43] permits to judge the numerical quality of an eigenvalue without knowing any associated vector. The normwise backward error η associated with $\tilde{\lambda}$ only is optimal since it achieves $\min_{\tilde{x}} \eta(\tilde{\lambda}, \tilde{x})$. It is defined by $\eta_{\rm opt}(\tilde{\lambda}) = \min\{\, \varepsilon > 0; \exists u \neq 0,\ (A+\Delta A)u = \tilde{\lambda}(B + \Delta B)u$ such that $\|\Delta A\| \leq \varepsilon\alpha$ and $\|\Delta B\| \leq \varepsilon\beta \,\}$, and can be computed as

$$\eta_{\rm opt}(\tilde{\lambda}) = \frac{1}{(\alpha + |\tilde{\lambda}|\beta)\|(A - \tilde{\lambda}B)^{-1}\|}.$$

Clearly the optimal backward error is expensive to compute and is interesting mainly from a conceptual point of view. However, it permits to shed an interesting light on the convergence behaviour of iterative solvers for computing eigenvalues [44] or can be the only way to analyse the stability of some CFD problems. The last point is illustrated in a joint work with the CFD Project at CERFACS [23].

More work concerning the conditioning of basic factorisations in Linear Algebra is reported in the next section.

6.6 Reliability of an orbitography software – Least-squares

n important application of least-squares problems in space industry is orbitography. It consists in determining several parameters associated with data about the trajectory of a satellite. Given a model for the motion of the satellite, and a collection of physical measurements, one is led to solve an unconstrained nonlinear optimization problem to compute the parameters of the trajectory. Usually, this is done iteratively by a Gauss-Newton algorithm, and each step amounts to a linear least-squares problem. In the framework of its activities on the numerical quality of scientific software, CNES (French National Center for Space Research) undertook an audit with CERFACS to evaluate the reliability and accuracy of a software used in orbitography.

The first part of the audit performed by CERFACS was concerned only with the solution of the linearized least-squares problem. The method used to evaluate the numerical reliability of the algorithm is based on the Wilkinson backward error analysis. As a first result of this study, we propose alternative algorithms and make a comparative study on actual data. The study shows that it is possible to achieve backward stability up to machine precision and improve simultaneously the accuracy in the estimated parameters. As a result, we recommend the QR factorization to solve the linear least-squares problem occurring in this orbitography problem. Furthermore it has been shown that the cases of non-convergence for the Gauss-Newton algorithm that occur sometimes are a result of a weakness in the model. Two reports [26] and [21] summarize this activity which has also been presented at the *Conference on Least Squares Methods : theory, algorithms and applications*, on the occasion of Åke Björck's 60th birthday, Linköping University, January 9–10, 1995.

This study also proved to be fruitful from a more theoretical point of view. It generated research on perturbation theory and more particularly in finding condition numbers for various problems from Linear Algebra. The formula for the exact normwise condition number for the linear least-square problems, where perturbations on the data (resp., on the solution) are measured with the Frobenius norm (resp. the spectral norm), was established [29]. A similar study has been performed on three basic factorizations in Linear Algebra: LU, QR and the polar factorization VH. Formulae for the normwise conditioning of these six factors were also proved and compared with the existing bounds in the literature. An interesting result was obtained in particular on the R factor of QR: the exact condition number can be much smaller than the upper bound proposed in [42]. These results have been presented at the second workshop of the International Linear Algebra Year at CERFACS, entitled "Eigenvalues and Beyond", Toulouse, October 17-20, 1995.

The second part of the audit was initiated in 1996 and consisted in a bibliographical study of nonlinear optimisation problems developed in [28]. It appeared that, in the litterature on optimisation, the notion of conditioning is seldom used. A few results, related with sensitivity analysis, can be interpreted in terms of condition number. However, to our knowledge, there is no mention of backward error analysis for optimisation problems although this concept is fundamental for the assessment of reliability of algorithms. The collaboration with CNES continues in 1997 in order to develop more theory in this field and to perform the stability analysis of non linear least-squares problems with the help of the software tool PRECISE.

[1] M. Bennani and T. Braconnier. Comparative behaviour of eigensolvers on highly nonnormal matrices. Tech. Rep. TR/PA/94/23, CERFACS, 1994.

[2] M. Bennani and T. Braconnier. Stopping criteria for eigensolvers. Tech. Rep. TR/PA/94/22, CERFACS, 1994.

[3] M. Bennani, T. Braconnier, and J.-C. Dunyach. Solving large-scale nonnormal eigenproblems in the aeronautical industry using parallel BLAS. In W. Gentzsch and U. Harms, editors, *High-Performance Computing and Networking*, volume 796, pages 72–77. Springer-Verlag, 1994.

[4] T. Braconnier. The Arnoldi-Tchebycheff algorithm for solving large nonsymmetric eigenproblems. Tech. Rep. TR/PA/93/25, CERFACS, 1993.

[5] T. Braconnier. The Arnoldi-Tchebycheff algorithm for solving large complex non hermitian generalized eigenproblems. Tech. Rep. TR/PA/94/08, CERFACS, 1994.

[6] T. Braconnier. The role of the orthogonalization scheme used for eigensolvers applied to nonnormal matrices. Tech. Rep. TR/PA/94/20, CERFACS, 1994.

[7] T. Braconnier. Influence of orthogonality on the backward error and the stopping criterion for Krylov methods. Numerical Analysis Report 281, Dept. of Mathematics, University of Manchester, 1995. Submitted to *BIT*.

[8] T. Braconnier and F. Chaitin-Chatelin. Chaotic behavior for eigensolvers applied on highly nonnormal matrices in finite precision. Technical Report TR/PA/96/50 and Numerical Analysis Report 303, CERFACS and Manchester Centre for Computational Mathematics, 1996.

[9] T. Braconnier and F. Chaitin-Chatelin. Roundoff induces a chaotic behavior for eigensolvers applied on highly nonnormal matrices. In J. Periaux et al., editor, *Computational Science for the 21st Century*. Wiley, 1997. To appear.

[10] T. Braconnier, F. Chatelin, and J.-C. Dunyach. Highly nonnormal eigenvalue problems in the aeronautical industry. *Japan J. Ind. Appl. Math.*, 12:123–136, 1995.

[11] F. Chaitin-Chatelin. Le calcul qualitatif. Comment donner un sens à des résultats faux ? Tech. Rep. TR/PA/95/10, CERFACS, Mars 1995. Support de cours pour le DEA intensif Problèmes inverses en Astrophysique. Concepts mathématiques et méthodes de résolution, CIRM, Luminy, 20–24 Mars 95.

[12] F. Chaitin-Chatelin. Is nonnormality a serious computational difficulty in practice ?, 1996. Proceedings of the IFIP WoCo 7 **"The Quality of Numerical Software"**, Oxford, 8–12 July, Chapman and Hall.

[13] F. Chaitin-Chatelin and V. Frayssé. *Lectures on Finite Precision Computations*. SIAM, Philadelphia, 1996.

[14] F. Chaitin-Chatelin and V. Frayssé. Qualitative computing and nonnormality. In S. D. Margenov and P. S. Vassilevski, editors, *Iterative Methods in Linear Algebra II*, pages 34–44. IMACS Series in Computational and Applied Mathematics, 1996.

[15] F. Chaitin-Chatelin and S. Gratton. Convergence of successive iteration methods in finite precision under high nonnormality. *BIT*, 36:455–469, 1996.

[16] F. Chatelin. *Valeurs propres de matrices*. Masson, Paris, 1988.

[17] F. Chatelin and T. Braconnier. About the qualitative computation of Jordan forms. *ZAMM*, 74:105–113, 1994.

[18] F. Chatelin, V. Frayssé, and T. Braconnier. Computations in the neighbourhood of algebraic singularities. *Num. Funct. Anal. Opt.*, 16:287–302, 1995.

[19] V. Frayssé, L. Giraud, and S. Gratton. Solveurs linéaires itératifs pour la résolution de systèmes complexes non hermitiens creux de grand taille. Contract Report FR/PA/96/34, CERFACS, 1996.

[20] V. Frayssé, L. Giraud, and V. Toumazou. Parallel computation of spectral portraits on the Meiko CS2. In H. Liddell, A. Colbrook, B. Hertzberger, and P. Sloot, editors, *High-Performance Computing and Networking*, volume 1067, pages 312–318. Springer-Verlag, 1996.

[21] V. Frayssé and S. Gratton. Moindres carrés pour l'orbitographie - Etude de stabilité - Partie II. Contract Report FR/PA/95/28, CERFACS, 1995.

[22] V. Frayssé, S. Gratton, and V. Toumazou. Note on the conditioning of linear systems of the kind $A^*Ax = b$. Technical Report TR/PA/96/17, CERFACS, 1996.

[23] V. Frayssé, M. Gueury, F. Nicoud, and V. Toumazou. Spectral portraits for matrix pencils: a physical application. Technical Report TR/PA/96/19, CERFACS, 1996. Submitted to *IMA J. Numer. Anal.*

[24] V. Frayssé and V. Toumazou. About the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$, 1996. Submitted to *C.R.A.S.*

[25] V. Frayssé and V. Toumazou. A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$. Technical Report TR/PA/96/18, CERFACS, 1996. Submitted to *Numer. Linear Algebra Appl.*

[26] S. Gratton. Moindres carrés pour l'orbitographie - Etude de stabilité - Partie I. Contract Report IR/PA/95/20, CERFACS, 1995.

[27] S. Gratton. Utilisation du produit de Kronecker pour le calcul du conditionnement en norme de Frobenius. Tech. Rep. Annexe du rapport FR/PA/95/28, CERFACS, 1995.

[28] S. Gratton. Conditionnement en Optimisation. Etude bibliographique. Contract Report FR/PA/96/36, CERFACS, 1996.

[29] S. Gratton. On the condition number of linear least squares problems in Frobenius norm. *BIT*, 36:523–530, 1996.

[30] A. Harrabi. Défaut de normalité. Manuscript, CERFACS, 1996.

[31] A. Harrabi. Spectres et pseudospectres d'une famille d'opérateurs non bornés sous la convergence uniforme ou collectivement compacte. Manuscript, CERFACS, 1996.

[32] P. Henrici. Bounds for iterates, inverses, spectral variation and field of values of nonnormal matrices. *Numer. Math.*, 4:24–40, 1962.

[33] N. J. Higham and P. A. Knight. Matrix power in finite precision arithmetic. *SIAM J. Matrix Anal. Appl.*, 16:343–358, 1995.

[34] A. Ilahi. Analyse théorique de la méthode de perturbation PRECISE. Manuscript, CERFACS, 1996.

[35] A. Ilahi. Application pratique du diagramme de Newton et perturbations de Lidskii. Manuscript, CERFACS, 1996.

[36] O. Marques and V. Toumazou. Spectral portraits computation by a Lanczos method, augmented matrix version. Tech. Rep. TR/PA/95/05, CERFACS, 1995.

[37] O. Marques and V. Toumazou. Spectral portraits computation by a Lanczos method, normal equation version. Tech. Rep. TR/PA/95/02, CERFACS, 1995.

[38] S. A. Orszag. Accurate solution of the Orr-Sommerfeld stability equation. *J. Fluid Mech.*, 50:689–703, 1971.

[39] A. Ostrowski. Über die Stetigkeit von charakteristischen Wurzeln in Abhängigkeit von den Matrizenelementen. *Jahresber. Dtsch. Math. Verein.*, 60:40–42, 1957.

[40] M. Rieutord. Inertial modes in the liquid core of the earth. *Phys. Earth Plan. Int.*, 90:41–46, 1995.

[41] M. Rieutord and L. Valdettaro. Inertial waves in a rotating spherical shell, 1996. Submitted to *J. Fluid Mech.*

[42] G. W. Stewart. On the perturbation of LU, Cholesky and QR factorizations. *SIAM J. Matrix Anal. Appl.*, 14:1141–1145, 1993.

[43] V. Toumazou. *Portraits spectraux de matrices: un outil d'analyse de la stabilité*. Ph.D. dissertation, Université H. Poincaré, Nancy, 1996.

[44] E. Traviesas. Conditionnement de la factorisation de Cholesky – Analyse inverse des erreurs pour des problèmes de valeurs propres, 1996. Rapport de stage de Troisième Année I.U.P. M.I.C.S.S. CERFACS et Université Paul Sabatier, Toulouse, France.

$\boxed{7}$ Nonlinear Systems & Optimization

7.1 Some properties of Newton's method for smooth and nonsmooth equations (O. Burdakov)

Variational inequalities, nonlinear programming, complementarity problems and other problems can be reduced to nonsmooth equations, for which some generalizations of Newton's method are known. The Newton path, as a natural generalization of the Newton direction, was suggested by D. Ralph for enlarging the convergence region (globalization) of Newton-Robinson's method in the nonsmooth case. We investigated some properties of both the Newton direction and the Newton path, which seem to be basic for various globalization strategies. In particular, a simple formula for the derivative of an arbitrary norm of residuals along the Newton direction, derived earlier by the author for the smooth equations, was generalized for the derivative along the Newton path.

[1] O. Burdakov, Some properties of Newton's method for smooth and nonsmooth equations, in: "World Scientific Series in Applicable Analysis", Vol. 5, *Recent Trends in Optimization Theory and Applications*, World Scientific, pp. 17–24.

[2] O. Burdakov, On properties of Newton's method for smooth and nonsmooth equations. Tech. Rep. TR/PA/95/14, CERFACS, 1995.

7.2 On using the minimum spanning tree algorithm for optimal secant approximation of derivatives (O. Burdakov)

The following problem was considered. Given $m+1$ points $\{x_i\}_0^m$ in R^n which generate an m-dimensional linear manifold, construct for this manifold a maximally linearly independent basis that consists of vectors of the form $x_i - x_j$. This problem is present in, e.g., stable variants of the secant method, where it is required to approximate the Jacobian matrix f' of a nonlinear mapping f by using values of f computed at $m+1$ points. In this case, it is also desirable to have a combination of finite differences with maximal linear independence. As a natural measure of linear independence, we suggested a functional which is maximized to find an optimal combination of m pairs $\{x_i, x_j\}$. We have shown that the problem is not NP–hard but can be reduced to the minimum spanning tree (MST) problem, which is solved by an MST–type algorithm in $O(m^2 n)$ time.

[1] O. Burdakov, An MST-type algorithm for the optimal basis problem. Tech. Rep. TR/PA/95/22, CERFACS, 1995. Submitted to *BIT*.

7.3 Global optimization (B. Hamma)

Many problems in engineering can be seen as an optimization problem, i.e.:

$$\text{minimize } f(x) \text{ where } x \in I\!R^n$$

subject to the general:

(possibily non linear) inequality constraints : $C_j(x) \leq 0 \qquad j = 1, ..., q$
(possibily non linear) equality constraints : $C_j(x) = 0 \qquad j = q+1, ...m$

and simple bounds $l_i \leq x_i \leq u_i$,
where l_i and u_i could take the values $l_i = -\infty$ or/and $u_i = +\infty$.

Thus if we assume that there exists a subset $\mathbf{C}$ of $I\!R^n$ for which all the above constraints are fullfiled, solving an optimization problem means trying to find, for a given objective (real valued) function f, over a set of decision variables $x \in \mathbf{C}$ (a subset of $I\!R^n$), the lowest function value f_* and the corresponding x_*.

Usually, one just has to find a value f_{local} given by x_{local} such that f_{local} is the the lowest value "just" in a neighborhood of x_{local}.
This is called a *local optimization*. Many methods have been developed for solving local optimization problems. Most of them use the properties of the objective function (gradients , hessians) or/and properties of the set of the decision variables. Good routines and packages are available from NAG, HARWELL, MINOS, LANCELOT, MODULOPT, ...

Unfortunately, most of practical engineering applications can be formulated as optimization problems in wich the user could find many local minimizers with different function values. Then, **how to decide**? **How being sure** that the best of all possible minima is reached?
This is the aim of global optimization: **finding the best of all possible minima**, not determining just *a local minimum*.

In our work new algorithms based on Simulated annealing algorithms as well as multistart algorithms are introduced.

In collaboration with A. Torn and S. Viitanen of Abo Akademi, Finland, we have developed MTASS, a parallel simulated annealing code for global optimization.

[1] C. Bonnemoy and B. Hamma, La méthode du Recuit Simule: optimisation globale dans R^n, *RAIRO Aut. Prod. Inf. Ind.* **25**, Sept. 1991, pp.477-496.

[2] B. Hamma, Parallel Optimization of Interplanetary Trajectories, in Jack Dongarra et al. (Eds), *Applied Parallel Computing in Industrial Problems and Optimization, PARA'96*, Lyngby, Denmark, Lecture Notes in Computer Science **N.1184**, Springer-Verlag, 1996, pp. 369-378.

[3] B. Hamma, A. Torn and S. Viitanen, Parallel Simulated Annealing for continuous global optimization, CERFACS Technical Report TR/PA/96/44. Submitted to *Optimization Methods and Software.*

7.4 Degeneracy in linear programming (B. Hamma, P. Armand)

This is a continuation of joint work with Paul Armand, LACO, Laboratoire d'Arithmétique de Calcul Formel et d'Optimization, Université de Limoges. We are concerned with aspects of degeneracy in Linear Programming. Different rules allowing the enumeration of all vertexes which are neighbours of a degenerated one are considered, see [1], [2]. We give some geometric approach of the TNP rule in [3]).

[1] P. Armand, Programmation linéaire multicritère et problèmes de la dégenerescence, Ph. D. Thesis, Université de Limoges, 1991.

[2] T. Gal and F. Gue, A new pivoting rule for solving various degeneracy problems, *Operations Research Letters* **11**, 1992 pp. 23-32.

[3] P. Armand and B. Hamma, Interprétation géométrique d'une nouvelle règle de Pivot en Programmation Linéaire, in *Proceedings of CARI'96, Colloque Africain sur la Recherche en Informatique*, INRIA-ORSTOM (eds.), Libreville, Gabon, Oct. 9-15, 1996.

8 Computational Chemistry

The Computational Chemistry activity was started at CERFACS in Fall 1993 in order to develop applications of intensive computing to specific problems encountered in fundamental physics and chemistry. It is part of the Parallel Algorithms Project.

Computational chemistry is an intensive modeling and simulation process that allows one to explore, test, and sort hundreds of alternatives for new molecules on a computer before or instead of fabricating them in the laboratory. It is used by academic researchers to explore underlying properties of materials and to advance our knowledge of molecular structures and, more recently, of their dynamical properties. It is used by leading pharmaceutical, materials, semiconductor, and chemical companies to shorten the time to market, to reduce costs, and to scan molecular properties of target molecules before beginning commercial development.

Among all the different domains of computational chemistry two approaches of quite general interest are pursued at CERFACS: Molecular Dynamics (MD) and Density Functional Theory (DFT). In both cases we benefit from our experience gained in other fields of computational physics and chemistry: experience in simulating transport processes with Vlasov and Boltzmann equations for MD, experience in DFT from our former activity in the investigation of nuclear collisions and atoms in strong laser fields. These two lines of research (MD and DFT) have been further developed in the framework of various national and supranational structures: the CNRS's Groupement de Recherche *S.P.A.R.CH.* on transport properties of charged particles, a French-German project PROCOPE for 1995 on the dynamics of electrons in metallic clusters, and a project of European Network on electron dynamics presently in discussion.

8.1 MD and DFT in Physical Chemistry

Before summarizing the results we obtained last year we briefly mention some basic concepts in Computational Chemistry. Most forms of "organized" matter (molecules, clusters, solids,...) may be represented as collections of atoms sharing a small number of electrons (typically one or two "valence" electron per atom) to ensure cohesion of the system. Electrons are rapid, quantal particles while atoms (or, more precisely, ions) may in most cases be considered as slow, classical objects. Electrons and ions interact via the Coulomb force but develop qualitatively different dynamics (quantal/classical) at very different time scales. For most practical purposes, the electronic and ionic motions can thus be decoupled to within a very good approximation.

In recent years, DFT has become one of the leading methods of Quantum Chemistry. The basic idea is to describe a static many-electron system completely and exclusively in terms of its ground-state particle density. Owing to its numerical simplicity compared to more traditional quantum chemical approaches, DFT nowadays is the method of choice for most computations involving large molecules. On the other hand, the extension of DFT to time-dependent systems has proven to be a successful means for describing electron dynamics in atomic and molecular systems.

For practical reasons, approaches treating ionic motion usually take into account the effect of the electrons only via simple ansatzes. The resulting model is a *Molecular Dynamics* description of the ions, with a more or less sophisticated treatment of the electrons. The MD method by its mathematical simplicity permits a very wide range of possible applications as soon as classical motion is dominant compared to quantum effects.

In the last three years we have settled our MD and DFT projects in the field of *cluster physics*, for scientific as well as strategical reasons. Clusters interpolate between the microscopic molecules of quantum chemistry and the more macroscopic samples of solid state physics. Cluster physics represents a quite new but rapidly developing branch dealing with the so-called "new form" of mesoscopic matter. Potential applications are numerous, ranging from catalysis (large surface-to-volume ratio) to molecular electronics (cluster wires) or high-T_c superconductors based on fullerenes.

8.2 The MD project

The MD project was started by R. Glowinski in 1993 and driven by F. S. Zhang (CERFACS) in the framework of the Qualitative Computing Group (F. Chaitin-Chatelin, V. Frayssé) and in collaboration with IRSAMC at Université Paul Sabatier (F. Spiegelmann, E. Suraud). The goal of this project was to develop a high level version of MD based on an effective potential fitted to some *ab initio* calculations. This amounts to describing ionic motion in a *classical* way but on a *quantal* potential energy surface integrating electronic effects. As a first step we have performed systematic studies of the dynamical and thermodynamical properties of sodium clusters.

From the computational point of view the ionic propagation is based on a standard MD algorithm of Leap Frog type. But we aim at calculating on very long times in order to attain thermodynamical properties from the ergodic hypothesis. The demanding part of the project, still under consideration, lies in the evaluation of the quantal potential energy surface. It requires a diagonalization of the hamiltonian. The corresponding matrix is pretty small, of the size n of the cluster itself ($n \sim 10$ to 100) but diagonalization has to be performed at each time step, which requires a heavy numerical effort. The first attempts at parallelizing this part of the code (which consumes most of the CPU) are encouraging (V. Frayssé) and should be implemented for realistic cases in the near future.

The results obtained in this project concern both the dynamics [1, 2] and thermodynamics (in particular in connection with the study of phase transitions as a function of temperature) [3] of sodium clusters. In a field where numerical simulations are still scarce and often based on very schematic hamiltonians, our high level MD simulations have allowed us to give some answers to some open questions such as the formation of transient dimer states in cluster collisions [1, 2]. Our recent study of the temperature dependence of vibrational spectra will also provide one of the first systematic studies with a realistic hamiltonian [3].

8.3 The DF project

Following the recommendations of the 1994 scientific Audit of CERFACS, the DF project was started in fall 1994 by E. Suraud (Université Paul Sabatier), F. Calvayrac (ENSL Lyon and Université Paul Sabatier), P. G. Reinhard (Erlangen University, FRG) and L. Giraud and in relation with F. S. Zhang; it has been further developed by C. A. Ullrich since November 1995. The aim of this project is to develop a parallel version of a DF code for static and dynamical studies in computational chemistry. At the beginning, the calculations have been oriented on cluster dynamics in strongly out of equilibrium situations such as in the fragmentation experiments performed at CEA in Grenoble or under the influence of strong femtosecond laser pulses.

As a first phase of the project we have adapted to our problem two versions of a two dimensional, dynamical DF code originally developed for problems in nuclear resp. atomic physics in order to scan the difficulties and potential research axes. In a second phase, an analysis of the 2-D algorithm has allowed us to effectively start the writing of a powerful, parallel 3-D version. We plan to develop a modular and adaptive (to various machines and physical/chemical problems) version of this code. A first version of this 3-D code is already working on local workstations. Implementation on the T3D of IDRIS has been started in 1996, in the framework of a joined CERFACS-UPS program at IDRIS.

From the computational side, the quantal propagation of the electrons in a DF model reduces to a matrix propagation of the vectors representing the wave functions of each electron. This requires the evaluation of either the exponential of a huge matrix (on the 3-D grid) or its (approximate) inversion. Another difficulty lies in a fast evaluation of the Coulomb field by solving the Poisson equation. The latter aspect might not be ideally suited to parallel computations but benefits in turn from a large amount of experience which should be used here (L. Giraud). On the other hand, the matrix propagation should be more easily adapted to massive parallelization.

First results in this project (see Figure 1) have been obtained with a sequential, vectorized 2-D version of the code which will serve as benchmark and exploratory tool [4]–[9]. They show the feasibility of the project and its scientific interests, as well as the level of cluster physics, as in terms of developing a DF solver or for more fundamental questions on the onset of quantum chaos.

A brief note on the computational performance of the codes: In 2-D, the typical grid size is 40×80 points on a uniform grid assuming cylindrical symmetry. The propagation of the orbitals (described by partial differential equations of the time-dependent Kohn-Sham type) is performed via the Crank-

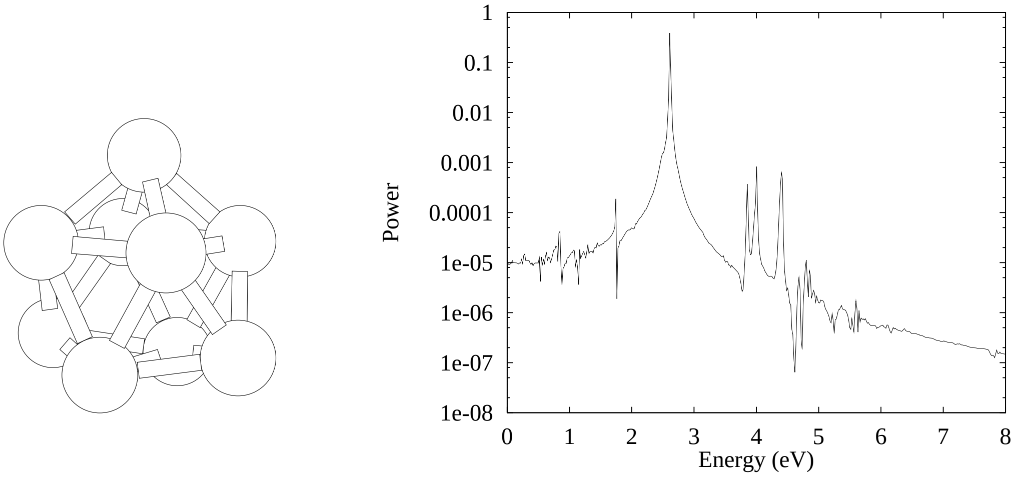

Figure 8.1: Left: Structure of a Na_9^+ Cluster. Right: Power spectrum of the linear dipole response of Na_9^+ within the spherical jellium model.

Nicholson method with or without the Peaceman-Rachford splitting of the evolution operator. The physical situation considered implies the computation of about 10000 time steps, each one taking about 10 s on a HP 9000/816 workstation (for three electronic orbitals). In our 3-D code, on the other hand, we have performed a number of optimizations using LAPACK libraries for the Crank-Nicholson scheme or, alternatively, implemented a method using alternating-direction time evolution in the direct or Fourier space where the partial derivatives transform into products. As a result, the execution time in 3-D is about the same as for the 2-D code, with typical grid sizes now in the order of $36 \times 36 \times 36$. The memory requirements in 3-D, however, are considerable (at least 32 Mo in the smaller cases of interest).

Furthermore, our quantal calculations have been complemented by a side project (E. Suraud, L. Feret) where a Vlasov approach is developed in place of the full quantal description. This leads to a huge simplification of the complexity of the code, even in 3-D, for the price of a loss of some quantal aspects. This may bring complementary information during the development phase of the quantal code. Parallelization of this code is also under progress in collaboration with L. Giraud, on the IDRIS' Cray T3D.

8.4 Perspectives

Most of the potential developments have been mentioned in the previous sections. Let us nevertheless briefly summarize the perspectives of our MD and DF projects. Both have provided very promising results which deserve well defined extensions. The MD project could integrate the development of Nosé-Hoover Generalized Molecular Dynamics algorithms, which are now used in many fields of physics and chemistry in order to consider systems at finite temperature in contact with a heat bath. The urgent goal of the DF project is of course the optimization of the first version of the 3-D code, with the perspective of evolving towards a powerful massively parallel 3-D time dependent DF solver. Both projects (MD and DF) should also evolve towards a wider spectrum of potential applications.

In both quantal and semi-classical calculations involved in the DF project the major bottleneck of massive parallelization lies in the Coulomb (Poisson) solver, which presently consumes typically 50% of CPU time. The importance of this problem is well known both in physics and chemistry, as well as in many other fields. We have settled a small group working on the building of an efficient parallel Coulomb solver. This group gathers F. Calvayrac, L. Giraud, M. Gross (who has recently developed a parallel version [in 3-D + symmetry] of the FFT Coulomb solver we presently use), E. Suraud and C. A. Ullrich. A first parallel version of the Coulomb solver has now been implemented, further tests and optimizations are currently under study.

As a last remark we would like to point out the original and exceptional opportunity offered by the Computational Chemistry project at CERFACS, due to the association of the expertise in high level computing developed in the Algo team with the scientific developments offered by a rapidly evolving field of physical chemistry. One may hence reasonably hope to settle quite rapidly a recognized and appreciated activity in CERFACS on this particular topic, with a relatively low-level investment.

8.5 Collaborations

- *Member of CNRS's Groupement de Recherche (GdR) S.P.A.R.CH.*
  Participation in conferences and training programs organized by the GdR on simulations of transport of charged particles in various fields of physics, chemistry and engineering studies.
- *French-German PROCOPE exchange program (1995 -).*
  This program is devoted to a study of electron dynamics with help of DF methods and is run by Prof. P. G. Reinhard of Erlangen University in FRG and by E. Suraud. The DF project, as presently developed in CERFACS constitutes the scientific project of this exchange program.
- *European Training and Mobility of Researchers program in preparation for 1995 to 1999.*
  This program will be devoted to the methods developed in the study of electron dynamics. A large component of the program will be based on the CERFACS DF project. The "hard" core of this project gathers the Grenoble CEA center (Prof. C. Guet), the Erlangen University (Prof. P. G. Reinhard) and Toulouse (E. Suraud).

[1] F. S. Zhang, F. Spiegelmann, E. Suraud, V. Frayssé, R. Poteau, R. Glowinski, and F. Chatelin, On the Formation of Transient $(Na_{19})_2$ and $(Na_{20})_2$ Cluster Dimers from Molecular Dynamics Simulations, *Phys. Lett. A* **193**, 75 (1994).

[2] F. S. Zhang, E. Suraud, F. Spiegelmann, V. Frayssé, F. Chatelin, R. Glowinski, Systematic Study of $(Na_n)_2$ Dimer Transient State in $Na_n + Na_n$ Collisions ($n = 8, 9, 19, and\ 20$), *Zeit. Phys. D* **35**, 131 (1995).

[3] F. S. Zhang, F. Spiegelmann, and E. Suraud, Temperature Dependent Structural and Vibrational Properties of Microclusters: Distance Dependent Tight-Binding Molecular Dynamics, *Zeit. Phys. D* (1996).

[4] F. Calvayrac, P. G. Reinhard, and E. Suraud, Nonlinear plasmon response in highly excited metallic clusters, *Phys. Rev. B* **52**, R17056 (1995).

[5] L. Mornas, F. Calvayrac, P. G. Reinhard, and E. Suraud, *Zeit. Phys. D* **38**, 71 (1996).

[6] L. Feret, E. Suraud, F. Calvayrac, and P. G. Reinhard, submitted to *J. Phys. B* (1996).

[7] C. A. Ullrich, A. Domps, F. Calvayrac, E. Suraud, and P.-G. Reinhard, Electron response of metallic clusters to strong laser pulses and energetic ion collisions, to appear in *Z. Phys. D* (also Tech. Rep. TR/CC/96/21, CERFACS, Toulouse, France).

[8] C. A. Ullrich, P.-G. Reinhard, and E. Suraud, Electron emission from strongly excited metal clusters, submitted to *Z. Phys. D* (also Tech. Rep. TR/CC/96/48, CERFACS, Toulouse, France).

[9] C. A. Ullrich, P.-G. Reinhard, and E. Suraud, Metallic clusters in strong femtosecond laser pulses, in preparation.

9 European Projects

9.1 Parallelization of a fluid dynamic code in the framework of a HPCN european contract (L.M. Carvalho and L. Giraud)

The Parallel Algorithm team is involved in Europort-1 (Esprit project EP 8421) project no 8336, whose acronym is HPCN3S. The partners of the HPCN3S consortium are CERFACS, EDF (France), CISE (Italy), IFP (France), SIMULOG (France) and the VKI (Belgium). The goal of this project is to develop on distributed memory platforms a parallel portable version of the existing code called N3S. This code solves the 3D Navier Stokes equation using unstructured meshes and implicit schemes. The N3S package is composed of two separate solvers, one aimed at simulating incompressible flows, the other at simulating compressible flows :

1. the incompressible solver is being developed by EDF-DER [1],
2. the compressible solver is jointly being developed by INRIA-SIMULOG.

In the first phase of this project which started on January 1^{st} 1994, the parallel algorithms project has been involved in the definition of the global parallelization strategy of the codes. In that respect a domain decomposition approach has been adopted. Another significant contribution in this project is the development of parallel iterative solvers for the solution of the linear systems that arise in the Stokes problem for the incompressible computation:

$\mathbf{Av} = \mathbf{f}$ for the velocity problem,

$\mathbf{BB}^{\mathrm{T}}\mathbf{p} = \mathbf{g}$ for the pressure problem.

In this respect, two parallel versions of the diagonally preconditioned conjugate gradient method have been implemented [2]. Results have been obtained on a Meiko CS-2 and an IBM SP2; both PVM and PARMACS have been tested as message passing library.

[1] J. Chabard, B. Metivet, G. Pot, and B. Thomas, (1992), An efficient finite element method for the computation of 3D turbulent incompressible flows, *Finite Elements in Fluids*, 869–895.

[2] E. D'Azevedo, V. Eijkhout, and C. Romine, (1992), Reducing Communication Costs in the Conjugate Gradient Algorithm on Distributed Memory Multiprocessors, Tech. Rep. Lapack Working Note 56, Oak Ridge National Laboratory, Oak Ridge.

[3] G. Degrez, L. Giraud, M. Loriot, A. Micelotta, B. Nitrosso, and A. Stoessel, (1995), Parallel Industrial CFD calculations with N3S, *Lecture Notes in Computer Science, High-Performance Computing and Networking, Springer-Verlag*, **919**, 820–825.

9.2 Parallel multibody simulation software in the framework of a HPCN european contract (L. Hamel, L. Giraud and C. Puglisi)

The Parallel Algorithms team has been involved in the development of a software tool in the framework of the CAPRI's subproject HIPERCOMBATS: *High PERformance COMputing in Multibody Analysis for Two-wheeler Suspension design* (Esprit - HPCN contrat no EP-9602 CAPRI).

The main objective of the project was to develop an interactive parallel software tool to improve the suspension design of two-wheelers. This project has been carried out by an international consortium of three partners with a high degree of complementary expertise: CEIT, a Spanish research centre whose main expertise is in multibody simulation, CERFACS, and Piaggio, one of the first European constructors of two-wheelers.
Following the directives of Piaggio, CEIT has developed a code simulating the two-wheeler suspension, pre-processing the experimental data given by the industrial partner, and post-processing the large amount of results produced by the simulations. CERFACS was responsible of the module that manages the concurent simulations on a network of non-dedicated workstations.
The project has started on January 1^{st} 1995 and lasted 18 months. In the first phase of the project, the Parallel Algorithms project has defined the global parallelization strategy of the new code and the tools required for the parallel implementation. The parallel platform targeted by the project was a heterogeneous network of workstations, that is the parallel platform available at Piaggio.
In the second phase, we have developed a C++ fault tolerant parallel manager suitable for intensive computations on such parallel platform. This module implements both fault recovering features and a load balancing strategy, that are key components for reliable and efficient computations on networks of heterogeneous non-dedicated workstations.
The parallel manager (developed by CERFACS) and the multibody simulation modules (developed by CEIT), have been integrated within an interactive Motif based pre-processor that has been delivered to the end user for the validation phase of the complete software.
One of the main results of this project has been to demonstrate that high performance computing can sometimes be easily introduced in industrial companies without requiring any equipment investments. The exploited parallelism can be considered as straightforward from a research point of view but is nevertheless a very powerful tool for industry. The code has been designed to run efficiently on a non-dedicated network of heterogeneous computers, then it is cost effective and could be used in any enterprise performing parametric studies ranging from SME to a large company like PIAGGIO V.E.

[1] S. Baldini, L. Giraud, L. Hamel, J. M. Jimenez, and L. M. Matey. Hipercombats : a parallel industrial tool for two-wheeler suspensions design. In *Proceedings of the International Conference on High-Performance Computing and Networking 1997 (to appear)*. Springer-Verlag, 1997.

[2] J. M.Jimenez, L. M. Matey, J. García, A. Avello, S. Baldini, L. Giraud, and L. Hamel. On the use of high performance computing in multi-body analysis for two-wheeled suspension design. In *Proceedings of the Symposium of Multi-Body dynamics (to appear), Bradford*, 1997.

9.3 European Project ODESIM, Optimum DESIgn for Multibody Systems in High Performance Computing and Networking (O. Burdakov, L. Giraud, B. Hamma)

The Parallel Algorithms Project is involved in the European ESPRIT IV Project N. 20116, whose acronym is ODESIM. The partners of the ODESIM consortium are CASA (Spain), CEIT (Spain), CERFACS, CR-FIAT (Italy), MATRA DATAVISION (France) and Siemens (Germany). The aim of the project is the development of a set of software tools for the kinematic and dynamic optimum design, designer-driven optimization and fully automatic optimization. ODESIM attempts to bring the world of multibody design closer to CAD environments through a tight integration between CAD programs and multibody solvers. This will allow CAD designers to create and optimize multyibody designs following the parametric or variational approach they currently use. ODESIM will use High Performance Computing and Networking techniques to perform parallel computations in the most CPU-intensive optimization problems.
In the fisrt part of the project that started Jannuary 1st, 1996, the Parallel Algorithms Project has been responsible of the definition of the Platforms, the parallel environments and the numerical optimization to be used. In this order we have retained the solution of heterogeneous computation on networks of worksations in PVM environment. Numerical Local optimization is adressed; however some kind of global optimization (by the use of multiple starts with different datas) is proposed for improving a given local minimization result. A parallelization strategy is being carried out to reduce the most CPU-intensive optimization problems.

9.4 The European project PINEAPL (A. McCoy, V. Toumazou and V. Frayssé)

PINEAPL (HPCN, Fourth Framework Project:20018) is a coordinated effort to produce a general purpose library of parallel numerical software. The role of CERFACS in this project is to analyse the numerical properties of the resulting library and the industrial applications.

A. McCoy (since June 1996) and V. Toumazou (since December 1996) are involved in this project, with V. Frayssé providing administrative and scientific support.

In this framework, the toolbox PRECISE, developed by the Qualitative Computing Group, shall be used for the analysis. For this purpose, the work performed has concerned producing Fortran versions of the PRECISE modules. This includes backward error analysis, spectral portraits, perturbed spectra and portable versions for parallel machines using PVM.

The development of PRECISE will result in a commercial quality product and therefore close collaboration with the other project partners has been initiated. To appreciate the context of PRECISE, A. McCoy has attended the IFIP conference (Oxford, July 1996) on Quality of Numerical Software and the PINEAPL technical meeting in Napoli (October 1996).

10 Conferences & Seminars

10.1 Conferences and seminars attended by members of the Parallel Algorithms Project

January 1995

Third ACME Conference on Computational Mechanics in the UK, Oxford, UK, January 4–5. I. S. Duff, *A review of frontal methods for solving linear systems,* invited talk.

Conference on Least Squares Methods : Theory, Algorithms and Applications, on the occasion of Åke Björck's 60th birthday. Linköping University, Sweden, January 9–10. I. S. Duff, *The use of augmented systems in the solution of linear least-squares problems,* invited talk. S. Gratton, J.C. Bergès and V. Frayssé, *Least squares problems in aerospace industry,* talk. C. Puglisi, attendee.

February 1995

Seminaire sur le Techniques Nouvelles de Traitement des Matrices Creuses pour les Problemes Industriels, Grenoble, February 13–14. I. S. Duff, *Solution of sparse problems using direct methods,* invited talk. L. Giraud, J. Koster, M.-M. Magolu and J. Patricio, attendees.

Seventh SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, USA, February 14–17. A. Drummond gave a talk and received a SIAM student award.

March 1995

GMD, EUROPORT 2 meeting, March 17. I. S. Duff, guest speaker.

EPFL, Lausanne, Switzerland, March 22. A. Drummond taught a PVM course.

Ecole Normale Supérieure de Lyon, Journées du PSMN, March 30–31. V. Toumazou and O. Marques, *Calcul de portraits spectraux,* talk. T. Braconnier, V. Frayssé and O. Marques, *Calculs de valeurs propres au CERFACS - Applications industrielles,* talk.

April 1995

Seventh Copper Mountain Conference on Multigrid Methods, Copper Mountain, Colorado (USA). C. Douglas, *A unified multigrid theory for non-nested grids and/or quadrature,* talk; *Cache based multigrid for serial and parallel computers,* talk.

HCM Workshop, Barcelona, Spain. L. Carvalho, L. Giraud, and M.-M. Magolu, *Parallel iterative schemes for finite element solutions,* talk.

PARASOL meeting, Frankfurt, Germany, April 24. I. S. Duff, representative for CERFACS and RAL.

May 1995

High Performance Computing and Networking 1995 (HPCN'95), Milan, Italy, May 1–3. I. S. Duff gave a tutorial on *Numerical Linear Algebra on High Performance Computers* (with J. Dongarra). L. Giraud, attendee.

June 1995

IBM T. J. Watson Research Center, Yorktown Heights, New York (USA). C. Douglas, *Burn, baby, burn: How to numerically simulate flames with finite rate chemistry while enjoying a nice high tea,* talk.

NSF-CBMS Meeting, San Francisco, California (USA), June 10–15. I. S. Duff spoke on frontal and multifrontal methods for sparse matrices.

Universitè Paul Sabatier, Toulouse, June 15. O. Burdakov, *On using the minimum spanning tree algorithm for optimal secant approximation of derivatives,* talk.

IMACS'95 Conference on iterative methods, Blagoevgrad, Bulgaria, June 17–20. F. Chaitin-Chatelin and V. Frayssé, *Qualitative computing and the robustness of numerical methods to high nonnormality,* talk. V. Toumazou and O. Marques, *Spectral portraits computations,* talk.

INFORMS Conference on Operations Research and Management Sciences, Singapore, June 25–28. O. Burdakov, *Multipoint secant approximation for nonlinear optimization,* invited talk.

Fox Prize Meeting, Dundee, UK, June 26. I. S. Duff, member of Adjudicating Committee to select prize winner.

16th Biennial Conference on Numerical Analysis, Dundee, UK, June 27–30. I. S. Duff and O. Marques, attendees.

July 1995

ICIAM'95, Hamburg, Germany, July 3–7. I. S. Duff, Chairman of CICIAM Steering Committee, invited minisymposium organizer, and invited speaker in minisymposium. C. Douglas, *Sparse methods for abstract multilevel solvers on serial and parallel computers,* talk; *Numerical solution of flame problems: formulations, flame sheets, and diffusion laminar flames with finite rate chemistry,* talk. O. Burdakov, *A generalization of the minimum spanning tree problem and its application to the optimal secant approximation of derivatives,* talk. F. Chaitin-Chatelin, *Qualitative computing and the robustness of numerical methods to high nonnormality,* talk (presented by S. Gratton). O. Marques and L. M. Carvalho, attendees.

Linear Algebra and its Applications, Manchester, UK, July 10–12. I. S. Duff, member of Programme Organizing Committee.

Black Sea University Summer School on "Chaitin Complexity and Applications", Mangalia, Romania, 27 June –6 July. F. Chaitin-Chatelin, *Why a theory of computability in finite precision can be useful for physics,* invited talk.

August 1995

Fourth Brazilian School on High Performance Scientific Computing, Rio de Janeiro, Brazil. C. Douglas, *Numerical simulation of flames,* talk; *Caching in with multigrid algorithms: problems in two dimensions,* talk.

September 1995

Purdue University, West Lafayette, Indiana (USA). C. Douglas, *Caching in with multigrid algorithms: problems in two dimensions,* talk.

IBM T. J. Watson Research Center, New York (USA). C. Douglas, *Caching in with multigrid algorithms: problems in two dimensions,* talk.

Summer School at Salsky Dvur, Czech Republic, September 12–14. I. S. Duff spoke on frontal and multifrontal methods for sparse matrices.

ILAY Workshop on Direct Methods, Toulouse, September 26–29. P. R. Amestoy, I. S. Duff and C. Puglisi, *Multifrontal parallel LU and QR factorizations,* invited talk. Parallel Algorithms Project, attendee.

October 1995

CRIHAN Board Meeting, Rouen, October 3. I. S. Duff, member of Scientific Council.

Hong Kong Institute of Science and Technology, Hong Kong. C. Douglas, *A unified multigrid theory for non-nested grids and/or quadrature,* talk.

International Conference on Parallel Algorithms (ICPA'95), Wuhan University, Wuhan, China. C. Douglas, *Cache based multigrid methods,* talk.

Xi'an Jiaotong University, Xi'an, China. C. Douglas, *Numerical simulation of flames with detailed, finite rate chemistry,* talk.

ILAY Workshop on Eigenvalues and Beyond, Toulouse, October 17–20. S. Gratton, *More condition numbers in linear algebra,* talk (presented by F. Chaitin-Chatelin). V. Toumazou, *Parallel spectral portraits,* talk. I. S. Duff, attendee, chairman, ILAY scientific chairman. Parallel Algorithm Project, attendee.

November 1995

BLAS Technical Workshop, Center for Research on Parallel Computation, University of Tennessee (USA), November 12–14. C. Douglas, *OBLAS: Objective Basic Linear Algebra Subprograms (One call for all BLAS),* talk. I. S. Duff, invited speaker and co-organizer.

Supercomputing '95, San Diego, California (USA), November 29–December 3. I. S. Duff gave a tutorial on Numerical Linear Algebra on High Performance Computers (joint with J. Dongarra, N. Nachtigal and D. Sorensen).

University of Strathclyde, Glasgow, Scotland. C. Douglas, *Numerical simulation of flames with detailed, finite rate chemistry,* talk.

December 1995

University of California, Berkeley (USA), December 7. I. S. Duff gave an invited seminar.

January 1996

IRSAMC, Université Paul Sabatier, Toulouse, January 25. C. Ullrich, *Density functional theory: a universal approach to the static and dynamic many-electron problem,* invited seminar.

March 1996

Conference on Algorithms and Computational Tools of Complex Systems organized by I. Prigogine at the Solvay Institute, Brussels, March 1–2. F. Chaitin-Chatelin, invited attendee.

April 1996

High-Performance Computing and Networking, Brussels, Belgium, April 15–19. V. Toumazou, *Parallel computation of spectral portraits on the Meiko CS2,* contributed talk. L. M. Carvalho gave a poster presentation.

State of the Art in Numerical Analysis, York, England. I. S. Duff was on the Program Committee and gave an invited talk. O. Marques, attendee.

ILAY Workshop on Linear Algebra in Optimization, Albi, April 22–25. I. S. Duff, attendee, chairman, ILAY scientific chairman. O. Burdakov, *An MST–type algorithm for the optimal basis problem*, contributed talk. Parallel Algorithms Project, attendee.

Copper Mountain Conference on Iterative Methods, Colorado (USA). I. S. Duff was session chairman and co-chair of discussion group.

May 1996

Worskhop on Optimum Design of Multibody Sytems, Toulouse. O. Burdakov gave three lectures: *Overview of multicriteria optimization problems, General formulation of optimization problems, Overview of automatic differentiation techniques*; L. Giraud, *General introduction to heteregenous computing*; B. Hamma, *Overview of global optimization techniques.*

5th SIAM Conference on Optimization, Victoria, Canada. O. Burdakov, *Interpolation methods for optimization and nonlinear equations,* contributed talk.

Department of Computer Science, Bilkent University, Ankara, Turkey, May 27. M. Benzi, *Preconditioning with sparse approximate inverses*, invited seminar.

CECAM Workshop on Non-Perturbative Methods for Solving Time Dependent Problems of Atoms and Molecules; LIC-C '96 Conference on The Two and More Electron Problem in Laser-Atom/Molecule Interactions, Crete, Greece, May 27-31. C. Ullrich, attendee.

NASA ICASE, Langley, Virginia (USA). O. Burdakov, *Interpolation methods for optimization and nonlinear equations,* invited seminar.

June 1996

Utah State University, Logan, Utah (USA). O. Burdakov, *Interpolation methods for optimization and nonlinear equations,* invited seminar.

Georgia Institute of Technology, Atlanta, Georgia (USA). O. Burdakov, *TV–norms and their applications*, invited seminar.

Ninth International Conference on Domain Decomposition Methods, Bergen, Norway, June 3–8. L. M. Carvalho gave a contributed talk.

ILAY Workshop on Iterative Methods, Toulouse, June 10–13. I. S. Duff, attendee, chairman, ILAY scientific chairman. Parallel Algorithm Project, attendee.

XIII Householder Symposium on Numerical Algebra, Pontresina, Switzerland, June 17–21 (attendance by invitation only). I. S. Duff, *Combinatorial aspects in the preconditioning and solution of sparse linear equations*, talk.

M. Benzi, *Preconditioning with sparse approximate inverses*, talk.

NATO Advanced Research Workshop on High Performance Computing, Cetraro, Italy. O. Burdakov, *Parallelization of interpolation methods for optimization and nonlinear equations*, talk.

July 1996

Numerical Methods and Computational Mechanics in Science and Engineering, Miskolc, Hungary. I. S. Duff gave an invited talk.

IFIP Working Conference on The Quality of Numerical Software, Oxford, England, July 8–12. F. Chaitin-Chatelin, *Is nonnormality a serious computational difficulty in practice?*, invited talk. A. Mc Coy, attendee.

Prague Mathematical Conference 1996, Prague, Czech Republic, July 8–12. M. Benzi, attendee.

Powellfest (Conference on numerical mathematics in honour of Mike Powell), Cambridge, England. I. S. Duff, attendee. O. Burdakov, *TV–norms and their applications*, contributed talk.

2nd World Congress of Nonlinear Analysts, Athens, Greece. O. Burdakov, *Interpolation methods for optimization and nonlinear equations*, talk.

August 1996

PARA96, Lingby, Denmark. I. S. Duff gave an invited talk. B. Hamma, *Parallel Optimization of Interplanetary Trajectories*, contributed talk.

EuroPar'96, Lyon, France. I. S. Duff gave an invited talk.

September 1996

ICTP Summer School on the Numerical Solution of PDEs, Trieste, Italy. I. S. Duff, invited lecturer.

Second ECCOMAS Symposium on Numerical Methods in Engineering, Paris, September 9–13. The Parallel Algorithms Project organized a minisymposium titled "Some CERFACS Activities in Scientific Computing for Engineering Applications" . The following talks were given by ALGO team members: L. M. Carvalho, *Parallel domain decomposition methods*; M. Benzi, *Preconditioning with sparse approximate inverses*; F. Chaitin-Chatelin, *Robust eigensolvers*; O. Burdakov, *TV–norms and their applications.*

CECAM Workshop on Going Beyond the Local Density Approximation in Physics and Chemistry, Lyon, September 9–13. C. Ullrich, *Time-dependent density functional theory: general formalism and non-perturbative applications to atoms and clusters,* contributed talk.

Scottish Computational Mathematics Symposium 1996, Strathclyde, Glasgow. I. S. Duff, attendee.

VECPAR II, Porto, Portugal. I. S. Duff gave an invited talk.

October 1996

Department of Computer Science, University of Tennessee, Knoxville, Tennessee (USA), October 4. M. Benzi, *Sparse approximate inverse preconditioning on vector computers*, invited seminar.

Scientific Computing Group, Los Alamos National Laboratory, Los Alamos, New Mexico (USA), October 7. M. Benzi, *Sparse approximate inverse preconditioning on vector computers*, invited seminar.

Second SIAM Conference on Sparse Matrices, Coeur d' Alene, Idaho (USA), October 9–11. I. S. Duff was on the Program Committee, was a panel member, and gave a contributed talk. M. Benzi, attendee.

CARI'96, Colloque Africain sur la Recherche en Informatique, Libreville, Gabon. B. Hamma, *Interprétation géométrique d'une nouvelle règle de Pivot en Programmation Linéaire*, poster.

November 1996

Réunion du GDR agrégat CNRS/CEA/Université Paul Sabatier, Toulouse, November 6–8. C. Ullrich, *Réponse électronique d'agrégats métalliques soumis à des impulsions laser intenses,* invited seminar.

Laboratoire de chimie théorique, Université P. et M. Curie, Paris, November 25. C. Ullrich, *Étude de la dynamique électronique des atomes et agrégats métalliques dans des faisceaux laser très intenses,* invited seminar.

10.2 Conferences and seminars organized by the Parallel Algorithms Project

February 1995

Nonlinear Electrodynamics: towards a 3D simulation, E. Suraud, Universitè Paul Sabatier, Toulouse, February 7.

September 1995

Continuous approaches to discrete optimization problems, P. Pardalos, University of Florida, September 8.

ILAY Workshop on Direct Methods, Toulouse, September 26–29.

October 1995

Lanczos for total least-squares problems, Å. Björck, Linköping University, Sweden, October 5.

Finding minimum height elimination trees for interval graphs in polynomial time, P. Heggernes, University of Bergen, Norway; *Storage schemes for the solution of the symmetric eigenvalue problem by Jacobi methods on multicomputers,* D. C. Gimenez, Universidad de Murcia, Spain; *Parallel LU factorisation on the VPP500,* D. J. Miron, Australian National University, October 6.

ILAY Workshop on Eigenvalues and Beyond, Toulouse, October 17–20.

Presentation of the Center for Scientific Computing in Finland, J. Rahola, CSC, Finland, October 25.

Dense linear algebra in electromagnetic scattering, J. Rahola, Center for Scientific Computing, Finland; *Computation of the principal singular vectors of Toeplitz matrices,* E. Sjöstrom, Linköping University, Sweden, October 26.

December 1995

Incomplete multilevel LU factorizations, J. C. Diaz, Center for Parallel and Scientific Computing, Tulsa, USA, December 13.

February 1996

An improved Lanczos algorithm for symmetric eigenproblems, M. Clint and K. Weston, Queen's University, Belfast and University of Ulster, Coleraine, February 15.

April 1996

Resolving degeneracy in LP and QP, R. Fletcher, University of Dundee, UK, April 16.

Dense factors of sparse matrices, R. Fletcher, University of Dundee, UK, April 19.

ILAY Workshop on Linear Algebra in Optimization, Albi, April 22–25.

Parallel fast Poisson solvers, M. Gross, Université Paul Sabatier, Toulouse, April 29.

May 1996

Dual iterative methods, A. Dax, Hydrological Service, Israel, May 10.

Workshop on Optimum Design of Multibody Systems, Toulouse, May 13–14.

June 1996

The potential fluid flow problem and the convergence rate of the minimal residual method, M. Tůma, Czech Academy of Sciences, Prague, June 4.

ILAY Workshop on Iterative Methods, Toulouse, June 10–13.

A breakdown of the block CG method, C. Broyden, Università di Bologna, June 18.

Iterative methods and analysis of linear systems arising in fluid dynamics, H. Elman, University of Maryland, June 24.

Theoretical and numerical results for a well-reservoir coupling problem, F. Dickstein, IM-UFRJ, Brazil, June 25.

Temporal characterization of data movement demands to predict performance of parallel programs, G. Alaghband, University of Colorado at Denver, June 26.

Large scale nonsymmetric eigenvalue problems and their numerical solution, Z. Bai, University of Kentucky, June 26.

July 1996

Dynamics of the power method, P. Knight, University of Strathclyde, July 3.

General truncation strategies for optimal Krylov methods, E. de Sturler, Swiss Center for Scientific Computing, July 22.

August 1996

The Matrix Market: a visual web resource for numerical matrix data, K. Remington, National Institute of Standards and Technology, August 29.

The PMESC programming library for distributed memory computers, S. Crivelli, University of Colorado at Boulder, August 30.

September 1996

A comparison of the HR method and the XHR method for the eigenvalue problem, P. Wiberg, University of Umeå, Sweden, September 24.

October 1996

Spectral transformations on small-signal stability assessment of large scale power systems, L. Bezerra, IM-UFSC, October 8.

Parallel computation of rotating flows, L. Lundin, Danish Technical University, October 23.

November 1996

Iterative methods for saddle-point problems with elliptic operators, Yu. A. Kuznetsov, Russian Academy of Sciences, Moscow, November 26..

Iterative methods for elliptic finite element problems with nonmatching grids, Yu. A. Kuznetsov, Russian Academy of Sciences, Moscow, Novemeber 27.

11 Publications & Reports

11.1 Book

[1] F. Chaitin-Chatelin and V. Frayssé, (1996), *Lectures on Finite Precision Computations*, SIAM, Philadelphia.

11.2 Journal publications

[1] P. R. Amestoy, T. A. Davis, and I. S. Duff, (1996), An approximate minimum degree ordering algorithm, *SIAM J. Matrix Anal. Applic.*, **17**, 886–905.

[2] P. R. Amestoy, M. Daydé, I. S. Duff, and P. Morére, (1995), Linear algebra calculations on a virtual shared memory computer, *Intern. J. High Speed Comput.*, **7**, 21–43.

[3] P. R. Amestoy, I. S. Duff, and C. Puglisi, (1996), Multifrontal QR factorization in a multiprocessor environment, *Num. Lin. Alg. Applic.*, **3**, 275–300.

[4] M. Arioli, I. S. Duff, D. Ruiz, and M. Sadkane, (1995), Block Lanczos techniques for accelerating the block Cimmino method, *SIAM J. Scientific Comput.*, **16**, 1478–1511.

[5] M. Bennani, F. Chatelin, and J.-C. Dunyach, (1995), Highly nonnormal eigenvalue problems in the aeronautical industry, *Japan J. Ind. Appl. Math.*, **12**, 123–136.

[6] M. Benzi, C. D. Meyer, and M. Tůma, (1996), A sparse approximate inverse preconditioner for the conjugate gradient method, *SIAM J. Scientific Comput.*, **17**, 1135–1149.

[7] O. Burdakov, (1996), On using the minimum spanning tree algorithm for optimal secant approximation of derivatives, *Z. angew. Math. und Mech.*, **76**, 389–390.

[8] F. Chaitin-Chatelin and S. Gratton, (1996), Convergence of successive iteration methods in finite precision under high nonnormality, *BIT*, **36**, 455–469.

[9] F. Chatelin, V. Frayssé, and T. Braconnier, (1995), Computations in the neighborhood of algebraic singularities, *Num. Funct. Anal. Opt.*, **16**, 287–302.

[10] C. C. Douglas, (1995), Implementing abstract multigrid or multilevel methods, *Comput. Appl. Math*, **14**, 1–19.

[11] I. S. Duff and J. K. Reid, (1996a), The design of MA48, a code for the direct solution of sparse unsymmetric linear systems of equations, *ACM Transactions on Mathematical Software*, **22**, 187–226.

[12] I. S. Duff and J. K. Reid, (1996b), Exploiting zeros on the diagonal in the direct solution of indefinite sparse symmetric linear systems, *ACM Transactions on Mathematical Software*, **22**, 227–257.

[13] I. S. Duff and J. A. Scott, (1996), The design of a new frontal code for solving sparse unsymmetric systems, *ACM Transactions on Mathematical Software*, **22**, 30–45.

[14] I. S. Duff, (1996), A review of frontal methods for solving linear systems, *Computer Physics Communications*, **97**, 45–52.

[15] A. Ern, C. C. Douglas, and M. D. Smooke, (1995), Detailed chemistry modeling of laminar diffusion flames on parallel computers, *Inter. J. Supercomputer Applic. High Perfor. Comput.*, **9**, 167–186.

[16] L. Giraud and R. S. Tuminaro, (1995), Time dependent solvers on distributed memory computers, *Calculateurs parallèles*, **7**, 255–269.

[17] L. Giraud, (1995), Block preconditioned conjugate gradient methods on a distributed virtual shared memory multiprocessor, *Intern. J. High Speed Comput.*, **7**, 161–190.

[18] S. Gratton, (1996), On the condition number of linear least squares problems in Frobenius norm, *BIT*, **36**, 523–530.

[19] O. A. Marques and Y.-H. Sanejouand, (1995), Hinge-Bending Motion in Citrate Syntase Arising from Normal Modes Calculations, *Proteins*, **23**, 557–560.

[20] R. M. Sydenstricker, A. L. G. A. Coutinho, L. Landau, and O. A. Marques, (1995), Pseudoconsistent Load Vector and Mass Matrix for the Discrete Kirchhoff Triangle and the Discrete Shear Triangle Elements, *Comm. in Numer. Methods in Eng.*, **11**, 317–330.

11.3 Conference proceedings and book chapters

[1] M. Arioli, A. Drummond, I. S. Duff, and D. Ruiz, (1995), A parallel scheduler for block iterative solvers in heterogeneous computing environments, In *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing*, Philadelphia, PA, SIAM, 460–465.

[2] P. Armand and B. Hamma, (1996), Interprétation géométrique d'une nouvelle règle de Pivot en Programmation Linéaire, In *Proceedings of CARI'96, Colloque Africain sur la Recherche en Informatique*, P. Moukeli, ed., INRIA-ORSTOM.

[3] M. Benzi and M. Tůma, (1996a), Approximate inverse preconditioning for the conjugate gradient method on a vector computer, In *Proceedings of the Prague Mathematical Conference 1996*, K. Segeth, ed., 29–34.

[4] M. Benzi and M. Tůma, (1996b), Approximate inverse preconditioning of iterative methods for nonsymmetric linear systems, In *Proceedings of the Copper Mountain Conference on Iterative Methods.*

[5] T. Braconnier and F. Chaitin-Chatelin, (1997), Roundoff Induces a Chaotic Behavior for Eigensolvers Applied on Highly Nonnormal Matrices, In *Computational Science for the 21st Century*, J. P. et al., ed., Wiley. To appear.

[6] O. P. Burdakov, (1995), Some properties of Newton's method for smooth and nonsmooth equations, In *Recent Trends in Optimization Theory and Applications*, R. P. Agarwal, ed., Philadelphia, PA, World Scientific, 17–24.

[7] L. Carvalho and L. Giraud, (1996), Parallel Domain Decomposition Experiments on the Meiko CS2-HA, In *High-Performance Computing and Networking (Lectures Notes in Computer Science; Vol. 1067)*, H. Liddell, A. Colbrook, B. Hertzberger, and P. Sloot, eds., Berlin, Germany, Springer-Verlag, 948–949.

[8] L. Carvalho, I. S. Duff, and L. Giraud, (1996), Linear algebra kernels for parallel domain decomposition methods, In *Advanced Computational Methods in Structural Mechanics*, M. Papadrakakis and G. Bugeda, eds., International Centre for Numerical Methods in Engineering (CIMNE), Barcelona, Spain, 1–17.

[9] F. Chaitin-Chatelin and V. Frayssé, (1996), Qualitative computing and nonnormality, In *Iterative Methods in Linear Algebra II*, S. D. Margenov and P. S. Vassilevski, eds., IMACS Series in Computational and Applied Mathematics, 34–44.

[10] F. Chaitin-Chatelin, (1997), Is nonnormality a serious computational difficulty in practice ?, In *The quality of Numerical Software*, R. Boisvert, ed., Chapmann and Hall, 300–314.

[11] G. Degrez, L. Giraud, M. Loriot, A. Micelotta, B. Nitrosso, and A. Stoessel, (1995), Parallel Industrial CFD calculations with N3S, In *Lecture Notes in Computer Science, High-Performance Computing and Networking, Springer-Verlag*, vol. 919, 820–825.

[12] I. S. Duff and J. A. Scott, (1996), Frontal software for the solution of sparse linear equations, In *Applied Parallel Computing in Industrial Problems and Optimization*, J. Wasniewski, J. Dongarra, K. Madsen, and D. Olesen, eds., Lecture Notes in Computer Science 1184, Berlin, Springer-Verlag, 227–238.

[13] V. Frayssé, L. Giraud, and V. Toumazou, (1996), Parallel computation of spectral portraits on the Meiko CS2, In *High-Performance Computing and Networking*, H. Liddell, A. Colbrook, B. Hertzberger, and P. Sloot, eds., vol. 1067, Springer-Verlag, 312–318.

[14] L. Giraud and R. Tuminaro, (1995), Domain decomposition algorithms for PDE problems with large scale variations, In *Domain Decomposition Methods in Scientific and Engineering Computing*, Providence, RI, AMS, 205–210.

[15] L. Giraud, M. Henry, L. Lefebvre, L. Perret, C. Puglisi, and D. Vaucher, (1995), Parallelization of ASTRYD, a software based on a time domain method for solving problems in vibraoacoustics, vol. 2, CETIM, 533–538.

[16] B. Hamma, (1996), Parallel Optimization of Interplanetary Trajectories, In *Applied Parallel Computing in Industrial Problems and Optimization*, J. D. et al, ed., Springer Verlag.

[17] O. Marques, (1996), Eigensolvers and applications in finite elemnt analyses, In *Advanced Computational Methods in Structural Mechanics*, M. Papadrakakis and G. Bugeda, eds., International Centre for Numerical Methods in Engineering (CIMNE), Barcelona, Spain, 66–79.

11.4 Theses

[1] A. L. Drummond. *Solution of general linear systems of equations using block Krylov based iterative methods on distributed computing environments*. PhD thesis, CERFACS, 1995. TH/PA/95/40.

[2] V. Toumazou. *Portraits spectraux de matrices: un outil d'analyse de la stabilité*. Ph.D. dissertation, Université H. Poincaré, Nancy, 1996.

11.5 Technical reports

[1] Parallel Algorithms Project. Rapport de fin de contrat Thomson: Resolution des equations de Maxwell tridimensionnelles dans le domaine frequentiel sur reseaux heterogenes de calculateurs. Technical Report FR/EL-PA/95/25, CERFACS, 1995.

[2] Parallel Algorithms Project. Activity report of the parallel algorithms project, october 1993 - december 1994. Technical Report TR/PA/95/37, CERFACS, 1995.

[3] M. Benzi and D. B. Szyld. Existence and uniqueness of splittings for stationary iterative methods with applications to alternating methods. Technical Report TR/PA/96/16, CERFACS, 1996. To appear in *Num. Math.*

[4] M. Benzi and M. Tůma. A sparse approximate inverse preconditioner for nonsymmetric linear systems. Technical Report TR/PA/96/15, CERFACS, 1996. To appear in *SIAM J. Sci. Comput.*

[5] M. Benzi. Remarks on the numerical solution of certain linear complementarity problems. Technical Report TR/PA/96/42, CERFACS, 1996. Submitted to *J. Comp. Appl. Math.*

[6] T. Braconnier and F. Chaitin-Chatelin. Chaotic behavior for eigensolvers applied on highly nonnormal matrices in finite precision. Technical Report TR/PA/96/50 and Numerical Analysis Report 303, CERFACS and Manchester Centre for Computational Mathematics, 1996.

[7] O. Burdakov. An MST-type algorithm for the optimal basis problem. Technical Report TR/PA/95/22, CERFACS, 1995. Submitted to *BIT*.

[8] O. Burdakov. On properties of Newton's method for smooth and nonsmooth equations. Technical Report TR/PA/95/14, CERFACS, 1995.

[9] Jesús Cardenal, Iain S. Duff, and José M. Jiménez. A projection method for the solution of rectangular systems. Technical Report RAL-TR-96-013, RAL, 1996. Also CERFACS Report TR/PA/96/05.

[10] L. M. Carvalho and L. Giraud. Additive schwarz for the schur complement method. Technical Report TR/PA/96/51, CERFACS, Toulouse, France, September 1996. Submitted to the proceedings of the Ninth Conference on Domain Decomposition Methods.

[11] F. Chaitin-Chatelin. La calcul qualitatif. Comment donner un sens à des résultats faux? Technical Report TR/PA/95/10, CERFACS, 1995.

[12] K. Andrew Cliffe, Iain S. Duff, and Jennifer A. Scott. Performance issues for frontal schemes on a cache-based high performance computer. Technical Report RAL-TR-97-001, RAL, 1997.

[13] Michel J. Daydé and Iain S. Duff. A blocked implemenation of Level 3 BLAS for RISC processors. Technical Report RAL-TR-96-014, RAL, 1996. Also ENSEEIHT-IRIT Technical Report RT/APO/96/1 and CERFACS Report TR/PA/96/06.

[14] Michel J. Daydé and Iain S. Duff. Use of computational kernels in full and sparse linear solvers, efficient code design on high-performance risc processors. Technical Report TR-PA-96-47, CERFACS, 1996. Also ENSEEIHT-IRIT Technical Report RT/APO/96/4.

[15] C. Douglas, A. Ern, and M. Smooke. High performance computing and numerical simulation of flames. Technical Report TR/PA/95/34, CERFACS, 1995.

[16] C. Douglas. Caching in with multigrid algorithms : Problems in two dimensions. Technical Report TR/PA/95/15, CERFACS, 1995.

[17] C. Douglas. A sparse matrix approach to abstract multilevel solvers on serial parallel computers. Technical Report TR/PA/95/35, CERFACS, 1995.

[18] C. Douglas. A review of numerous parallel multigrid methods. Technical Report TR/PA/95/36, CERFACS, 1995.

[19] Iain S. Duff. Sparse numerical linear algebra: direct methods and preconditioning. Technical Report RAL 96-047, RAL, 1996. Also CERFACS Report TR-PA-96-22.

[20] Iain S. Duff and Jennifer A. Scott. A comparison of frontal software with other sparse direct solvers. Technical Report RAL-TR-96-102, RAL, 1996.

[21] V. Frayssé, L. Giraud, and S. Gratton. Solveurs linéaires itératifs pour la résolution de systèmes complexes non hermitiens creux de grand taille. Contract Report FR/PA/96/34, CERFACS, 1996.

[22] V. Frayssé and S. Gratton. Moindres carres pour l'orbitographie. Etude de stabilite. Technical Report FR/PA/95/28, CERFACS, 1995.

[23] V. Frayssé, S. Gratton, and V. Toumazou. Note on the conditioning of linear systems of the kind $A^*Ax = b$. Technical Report TR/PA/96/17, CERFACS, 1996.

[24] V. Frayssé, M. Gueury, F. Nicoud, and V. Toumazou. Spectral portraits for matrix pencils: a physical application. Technical Report TR/PA/96/19, CERFACS, 1996. Submitted to *IMA J. Numer. Anal.*

[25] V. Frayssé and V. Toumazou. A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$. Technical Report TR/PA/96/18, CERFACS, 1996. Submitted to *Numer. Linear Algebra Appl.*

[26] L. Giraud and S. Gratton. Solveurs lineaires performants pour la resolution de systemes complexes non hermitiens creux de grande taille. Technical Report TR/PA/95/24, CERFACS, 1995.

[27] L. Giraud, P. Noyret, E.Sevault, and V. Van Kemenade. IPM - User's guide and reference manual. Tech. Rep. TR/PA/95/01, CERFACS, 1995.

[28] S. Gratton. Moindres carres pour l'orbitographie. Etude de stabilite. Technical Report IR/PA/95/20, CERFACS, 1995.

[29] S. Gratton. Conditionnement en Optimisation. Etude bibliographique. Contract Report FR/PA/96/36, CERFACS, 1996.

[30] B. Hamma. Parallel simulated annealing for continuous global optimization. Tech. Rep. TR/PA/96-44, CERFACS, 1996.

[31] O. Marques. BLZPACK: Description and User's guide. Technical Report TR/PA/95/30, CERFACS, 1995.

[32] O. A. Marques and V. Toumazou. Spectral portrait computation by a Lanczos method (normal equation version). Tech. Rep. TR/PA/95/02, CERFACS, 1995.

[33] O. A. Marques and V. Toumazou. Spectral portrait computation by a Lanczos method (augmented matrix version). Tech. Rep. TR/PA/95/05, CERFACS, 1995.

[34] E. Traviesas. Conditionnement de la factorisation de Cholesky – Analyse inverse des erreurs pour des problèmes de valeurs propres, 1996. Rapport de stage de Troisième Année I.U.P. M.I.C.S.S. CERFACS et Université Paul Sabatier, Toulouse, France.

[35] C. A. Ullrich and E. K. U. Gross. Many-electron atoms in strong femto-second laser pulses: a density functional study. Technical Report TR/PA/96/21, CERFACS, 1996. to appear in *Comments At. Mol. Phys.*

[36] C. A. Ullrich, A. Domps, F. Calvayrac, E. Suraud, and P.-G. Reinhard. Electron response of metallic clusters to strong laser pulses and energetic ion collisions. Technical Report TR/PA/96/48, CERFACS, 1996. to appear in *Z. Phys. D.*

[37] P. Wiberg. A Study of the HR and Extended HR Methods for the Standard Eigenvalue Problem (Master's Thesis). Technical Report UMINF 96.16, Department of Computing Science, Ume University, 1996.