N° Ordre : 151

Habilitation à Diriger des Recherches

Spécialité : Informatique

The Power of Backward Error Analysis

Puissance de l'Analyse Inverse des Erreurs

soutenue le 13 Septembre 2000 à

l'Institut National Polytechnique de Toulouse

par

Valérie FRAYSSÉ

CERFACS

devant le Jury composé de :

J.-C. Bergès	CNES, Toulouse	
F. Chaitin-Chatelin	Professeur à l'Université Toulouse I Group Leader at CERFACS	
I. S. Duff	Project Leader at CERFACS Group Leader at Rutherford Appleton Laboratory Visiting Professor at Strathclyde University	Rapporteur
G. H. Golub	Professor at Stanford University	
S. Hammarling	Principal Consultant at Nag Ltd Visiting Professor at Cranfield University	Rapporteur
B. Philippe	Directeur de Recherches à l'INRIA/IRISA	Rapporteur
P. Spitéri	Professeur à l'ENSEEIHT	
J.-L. Vaudescal	EDF Division Recherche et Développement, Clamart	

CERFACS Report TH/PA/00/65

Remerciements / Acknowledgements

Cette Habilitation à Diriger des Recherches est le point d'orgue d'une collaboration très riche, tant sur le plan scientifique que sur le plan humain, avec vous, Françoise Chaitin-Chatelin. Mon admiration pour votre maîtrise de la maïeutique et pour l'étendue de votre culture mathématique et philosophique est immense. Je vous dois beaucoup trop pour pouvoir jamais vous exprimer assez mes remerciements.

Je remercie Monsieur Jean-Claude Bergès du CNES d'avoir participé à mon jury d'Habilitation. Les collaborations industrielles que nous avons menées ensemble ont eu un rôle très important dans ma carrière scientifique.

Iain Duff, I will never forget that you have provided me, for more than ten years, with the professional environment necessary to accomplish this Habilitation. As the expert of numerical algorithms for sparse matrices, your opinion on my work was particularly valuable to me, and I thank you also for being a referee for my Habilitation.

My sincere thanks go to Professor Gene Golub from Stanford University for taking part in my jury and expressing support and interest for my work. It was for me a great honour and an immense pleasure.

Ever since he was my mentor at the XII Householder Symposium in California in 1993, Sven Hammarling has offered me continuous support, in particular through our collaboration in the PINEAPL European Project. An heir and promoter of Wilkinson's backward error analysis, Sven Hammarling was an obvious choice to act as a referee for this Habilitation. I am deeply grateful that he accepted this task.

Je remercie sincèrement Monsieur le Professeur Bernard Philippe, Directeur de Recherche à l'INRIA/IRISA, d'avoir accepté d'être l'un des rapporteurs de mon Habilitation à Diriger des Recherches.

Monsieur Pierre Spitéri, vous avez suivi et encouragé ma carrière depuis mes débuts comme étudiante à l'ENSEEIHT. Je suis très heureuse que vous soyez aujourd'hui mon correspondant à l'Institut National Polytechnique de Toulouse pour cette Habilitation et

je vous remercie très chaleureusement pour votre constant soutien.

Enfin, je remercie Monsieur Jean-Louis Vaudescal de la Division Recherche et Développement d'EDF pour l'intérêt qu'il a porté à mes récents travaux sur les itérations emboîtées et pour sa participation à ce jury.

Merci au CERFACS, à tout son personnel informatique et administratif si compétent et si disponible, à tous les membres du Projet Algorithmes Parallèles et à mes compagnons du Groupe Calcul Qualitatif. Après dix ans de carrière dans cet environnement très international et très évolutif, j'ai été amenée à rencontrer, collaborer, apprécier de nombreuses personnes au CERFACS. Beaucoup ont joué un rôle important dans ma vie professionnelle, mais aussi dans ma vie personnelle. Il m'est impossible des les citer tous ici et j'espère qu'ils me le pardonneront. Néanmoins je veux dire à mon alter ego Luc Giraud et à mes compères Serge Gratton et Vincent Toumazou combien notre amitié m'est chère. Et pour finir, je remercie profondément toute ma famille et tous les amis qui m'ont accompagnée.

The power of backward error analysis

Abstract
As the numerical simulations tackle physical problems of increasing difficulty and numerical models of increasing complexity, the coupling between the numerical approximation and the finite precision plays a crucial role to assess the domain of computability of the solutions. Backward error analysis appears today as one of the key to understand this coupling. Although this concept was primarily designed as a tool to understand the behaviour of numerical algorithms in finite precisions, its scope is much wider. The finite precision computations fit into the large scenery of the methods of approximations, where the backward error analysis appears to be the tool of choice to deal with any kind of uncertain computations.
After explaining this conceptual level, from which one can fully embrace the power of the backward error analysis, we propose a tour of the most recent techniques used to assess the reliability of numerical software. The current trend is to develop more and more refined models which are better and better suited to describe the perturbations generated by the finite precision arithmetic.
We then review three examples of approximation in exact arithmetic that have been enriched by the concept of backward error analysis: the homotopic perturbations, the pseudospectra of matrices and operators, and the convergence of iterative methods in Linear Algebra. This work concludes on the open question of the extraordinary robustness of Krylov methods with respect to inexact computation.

Keywords: finite precision computations, computability, backward error, conditioning, iterative methods, Krylov methods.

Puissance de l'analyse inverse des erreurs

Résumé
Alors que la simulation numérique s'emploie à résoudre des problèmes physiques de difficulté croissante, le couplage entre l'approximation numérique et le calcul à précision finie joue un rôle essentiel dans la calculabilité des solutions. L'analyse inverse des erreurs apparaît aujourd'hui comme l'une des clefs de la compréhension de ce couplage, en permettant de traiter le calcul incertain.
Après avoir exposé cette approche conceptuelle qui, seule, permet d'apprécier réellement la puissance de l'analyse inverse des erreurs, nous passons en revue les techniques les plus récentes pour l'évaluation de la fiabilité du logiciel numérique, et nous présentons trois domaines d'application où l'analyse inverse des erreurs permet l'étude de méthodes d'approximation. Nous terminons par une réflexion sur la question, toujours ouverte, de l'extraordinaire robustesse des méthodes de Krylov vis-à-vis au calcul incertain.

Mots-clés : calcul en précision finie, calculabilité, erreur inverse, conditionnement, méthodes itératives, méthodes de Krylov.

Contents

Foreword

This document presents a synthesis of my eleven years of research at CERFACS, where I started in September 1989 as a graduate student. The unifying theme of my research activities has been the use and development of backward error analysis. The introduction recalls the fundamentals of the theory of computability in finite precision, that we developed in the book *Lectures on Finite Precision Computations* co-authored with F. Chaitin-Chatelin and published by SIAM in 1996. Our approach consists in treating the finite precision approximation in the same way as any numerical approximation. The resulting framework proved to be extremely useful in understanding the subtleties of the coupling between the numerical and arithmetic parameters which govern the reliability of finite precision computations. The work I have developed afterwards is more concerned with the techniques of backward error analysis for numerical software and methods. This is the part that I have chosen to develop in more detail here. In this document, I would like to survey the evolution of the backward error analysis since its conceptualisation by J. H. Wilkinson in the sixties. I hope the reader will appreciate the power of this notion whose ever-growing impact goes now far beyond finite precision computations for which it was originally intended. I will explain how my own work and that of my co-workers in the Qualitative Computing Group led by F. Chaitin-Chatelin at CERFACS fit within the research efforts of the international community in the further development and use of backward error analysis.

Chapters 1 to 4 and the Bibliography, which consitute the core of this document, are devoted to the presentation of backward error analysis. Because I have chosen to paint a panorama of backward error analysis at the dawn of the 21^{st} century instead of just describing my own work, some readers may feel that my own contributions to the topic are not highlighted enough. Let me summarise them presently. Most of my research work performed before 1995 is included in the above mentioned book *Lectures on Finite Precision Computations*. This research monograph combines techniques from engineering and mathematics to describe the rigorous and novel theory of computability in finite precision. The theoretical analysis is supplemented by a wide range of numerical experiments and the software tool PRECISE is used extensively to explore the stability on the computer. This software tool was further enhanced to be part of the parallel library coordinated by Nag in the framework of the European Esprit project PINEAPL (1996–1998). The topic of nearness to singularity is addressed in a joint paper with F. Chatelin and T. Braconnier entitled "Computations in the neighbourhood of algebraic singularities" and published in

Num. Funct. Anal. Opt. in 1995. My contribution to backward error analysis in linear algebra is reflected in two papers, one with V. Toumazou entitled "A note on the normwise perturbation theory for the regular generalised eigenproblem $Ax = \lambda Bx$ published in *J. Numer. Linear Algebra Appl.* in 1998, and one with S. Gratton and V. Toumazou dealing with "Structured backward error and condition number for linear systems of the type $A^*Ax = b$" and published in BIT in 2000. My work related with high nonnormality, iterative methods and pseudospectra has appeared in various proceedings of international conferences. My recent studies on inner-outer schemes are currently under review.
As already mentioned, my research is intimately linked with software production. Apart from the PRECISE toolbox already mentioned, I contributed to packaged iterative solvers for linear systems and eigensolvers. Most of these routines are publicly available and are regularly downloaded worldwide by many researchers working in various areas ranging from geophysics or ocean modelling to theoretical physics. As a senior researcher in the Parallel Algorithms Project, I have been supervising a number of post-doctoral researchers, Ph.D. Student and trainees. I have been the main supervisor of the Ph.D. thesis of A. Bouras on inner-outer iterations, to be defended in September 2000. My teaching activities take place at three levels: regular internal training at CERFACS for the newcomers, lectures in academia on eigenvalues and on the numerical quality of software in finite precision, and specific training for the engineers from CERFACS shareholders.
Additionally, from its nature, CERFACS depends strongly on industrial funding and this is why I have been involved in a number of industrial collaborations mainly with CNES, the French space agency and Aerospatiale. A complete list of my publications followed by a curriculum vitae (in French) where my administrative and contractual activities are listed are available in this document. Finally, I include as an extension three papers and a technical report which develop some aspects of my work that are mentioned in the core of the document.

Chapter 1

Introduction

For a scientist or an engineer, it is of crucial importance to have confidence in the results of large codes and in the robustness of the models used. Indeed the emergence of supercomputers has allowed the intensive use of numerical simulation to replace physical experiments, even for *problems at the frontiers of instability*. Computed results may have very few or even no correct digits. It is really a challenging problem to design tools which offer information on the quality and validity of computer results, and which provide engineers and scientists with both quantitative and qualitative analyses, in order to help them extract the appropriate information from results that are seemingly wrong. The research I have been conducting in the last ten years fits in this context.

The theory of computability described in the first part of this chapter has been designed to give a theoretical framework to finite precision computations, which attempts to unify the treatment of numerical and arithmetic approximations. In the second part, we set out the main issues related to the assessment of the quality of reliable software and numerical methods.

1.1 Computability in finite precision

1.1.1 Numerical approximation

We consider the mathematical problem (P) which consists in solving the equation

$$(P) \qquad F(x) = y. \tag{1.1}$$

We set $H = F^{-1}$. We assume that there is a solution $x = H(y)$ which is locally unique. We suppose that F (resp., F^{-1}) is continuous in the neighbourhood of x (resp., y): the problem (P) is then well-posed in the sense of Hadamard.

Scientific computing mostly deals with problems (P) that cannot be solved as such (that is, by use of a simple inverse operation). Instead, they are replaced by a family of

approximations of well-posed problems

$$(P_\theta) \qquad F_\theta(x_\theta) = y_\theta, \tag{1.2}$$

where the parameter θ is meant to tend to zero. We set similarly $H_\theta = F_\theta^{-1}$. The central question is then to prove that the approximation x_θ tends towards x when θ tends to zero. This is the core of theoretical numerical analysis. A typical example is when (P) is a partial differential equation and (P_θ) is a discretization of it. This kind of approximation is governed by the well-known *Lax principle* which states

$$\text{consistency} \; + \; \text{stability} \implies \text{convergence}.$$

Consistency is the fact that

$$\begin{cases} \lim_{\theta\to 0} y_\theta = y, \\ \lim_{\theta\to 0} F_\theta(z) = F(z), \forall z \text{ in the neighbourhood of } x. \end{cases}$$

The stability is mathematically the equicontinuity[1] of the family of approximations $\{F_\theta\}$. The conjunction of both properties is a sufficient condition for the convergence of x_θ to x when $\theta \to 0$. In theoretical numerical analysis (that is, numerical analysis in exact arithmetic), F_θ and/or H_θ are known. The role of the numerical analyst is to establish the proofs of convergence.

The difference $|F_\theta(z) - F(z)|$ is called the *residual consistency error*. The Lax principle holds when replacing the residual consistency error by the *direct consistency error* $|H_\theta(z) - H(z)|$ which is more convenient when F_θ and/or F are not available (see Chapter 1, Theorem 2.2 of [19]). It is traditional in theoretical numerical analysis and functional analysis to prove convergence by inspecting the **direct** error $|x - x_\theta|$ [3, 30]. On the contrary, in applied numerical analysis, or numerical software (that is numerical analysis with finite precision), it is often the **residual** error $|y - F(x_\theta)|$ that plays the fundamental role [115]. This difference in attitude reveals much more than a matter of taste or tradition. If both attitudes are mathematically equivalent in exact arithmetic, they are not anymore equivalent *in finite precision* as we will see later.

There is a kind of approximation that one would not list naturally as being of the type defined by (1.2), but which however fits perfectly into this framework: it is finite precision approximation. When we started to think more conceptually about the principles that underly the finite precision computation, we realised that it is bound to obey also the Lax principle of approximation. This led us to define a theory of computability which is discussed in detail in the book *Lectures on Finite Precision Computations* co-authored with F. Chaitin-Chatelin [19]. Let us now recall the main ideas that underpin our understanding of finite precision computations.

[1]The family $\{F_\theta,\ 0 \le \theta < 1\}$ is *equicontinuous* at x if $\forall \varepsilon > 0$, $\exists \delta > 0$ independent of θ such that $\|x' - x\| \le \delta$ implies $\|F_\theta(x') - F_\theta(x)\| \le \varepsilon$ for all θ, $0 \le \theta < 1$. (See Chapter 1, Definition 2.1 of [19]).

1.1.2 Finite precision approximation

In this paragraph, we focus on the basic situation where $x = H(y)$ is exactly computable in a finite number of steps in exact arithmetic. This restriction is done to facilitate the presentation but iterative algorithms can also be fitted into the same framework [19]. We also assume that x is a **regular** point, i.e. that H has continuous first derivatives with respect to prescribed data at y (otherwise x would be a **singular** point). We look at the asymptotic properties of computations on an ideal computer whose precision increases when some parameter ε tends to zero (this parameter can be thought of as machine precision for instance). Let

$$(P_\varepsilon) \qquad F_\varepsilon(x_\varepsilon) = y_\varepsilon \tag{1.3}$$

be the problem solved with a computer of precision ε. We set $x_\varepsilon = H_\varepsilon(y_\varepsilon)$.

As in the case of numerical approximation, the central question is the convergence of $x - x_\varepsilon$ when $\varepsilon \to 0$. We say that x is *computable in finite precision* if and only if x_ε tends towards x when $\varepsilon \to 0$. If x is proved computable in finite precision, then its computer approximation in double precision will be more accurate than that in single precision: this might not be true otherwise.

Note that this framework is also suitable for any kind of *uncertain* computations, that is with data that are only known to some prescribed tolerance (such as in physical measurements).

The forward analysis carried out in Chapter 1 of [19] leads us to the following decomposition of the forward error

$$x - x_\varepsilon = E_1 + E_2 \tag{1.4}$$

where

- $E_1 = DH(y)(y - y_\varepsilon) + O(\varepsilon^2)$ expresses the stability (in the sense of sensitivity) of the original problem $x = F^{-1}(y)$. Here $DH(y)$ denotes the Jacobian of F^{-1} at y,
- $E_2 = H_\varepsilon(y_\varepsilon) - H(y_\varepsilon)$ represents the direct consistency error at y_ε.

We see that the forward error $|x - x_\varepsilon|$ is subjected to two distinct influences:

- the stability of the mathematical problem through A,
- the finite precision approximation which is, at best, of order 1 in ε since one can show that it always contains a term of the order of ε.

In this context, applying the Lax principle to finite precision computations leads to the following sufficient condition:

$$\text{arithmetic reliability} \implies \text{computability}.$$

By *arithmetic reliability* we mean the set of the two following properties

1. the family of approximations $\{H_\varepsilon\}$ is equicontinuous,

2. the direct consistency error at y is of order 1 in ε, i.e. $|H_\varepsilon(z) - H(z)| \leq C(y, F)\varepsilon$ for any z in the neighbourhood of y.

The constant $C(y, F)$ depends only on the data and not on the parameter ε. This property is often found in the literature under the name "numerical stability with respect to round-off". We find this term misleading in the sense that it omits the consistency property to retain only the idea of stability (equicontinuity).

An interesting example where the computability in finite precision can be easily proved is Gaussian Elimination. It is discussed in detail in [19]. The growth factor (that is the maximal element growth during the factorisation) plays the role of the constant $C(y, F)$.

In this paragraph, we have taken into account only one level of approximation ε. In practice we need to solve problems that mix several kinds of approximation.

1.1.3 Numerical and arithmetic coupling

In theoretical numerical analysis, it is often the case that different kinds of approximations are used: for instance the space and time discretization of a partial differential equation. According to the numerical method chosen to approximate the equation, a condition linking the different approximation parameters must be fulfilled to guarantee the convergence of the approximation. This is called *conditional stability*.

An interesting example of conditional stability is a fully discrete spectral method for a first-order hyperbolic mixed initial-boundary value problem treated by Trefethen and Trummer [112], where the notion of ε-pseudospectrum[2](in exact arithmetic) is used to express the coupling between the space and time discretizations (see also [19] for a description of this example).

It is inevitable that the last step of the approximation, that is the computation in finite precision, must also be coupled with the other numerical approximation parameters. A typical example is the approximation of the second derivative by finite differences: in finite precision, the discretization parameter has to be chosen larger than a threshold depending on the arithmetic precision ε.

Therefore, numerical stability in exact arithmetic cannot always ensure computability in finite precision. Instead, one has to consider a uniform stability with respect to *all* approximation parameters, i.e. numerical as well as arithmetic parameters. On a computer, a numerically stable method can become conditionally stable with respect to the arithmetic of the computer.

[2]The ε-pseudospectrum of a matrix A is the set of eigenvalues of all the matrices $A + \Delta A$, where some norm of ΔA is bounded by ε. For more details, see Section 3.2 in Chapter 3.

1.1.4 Exact versus finite precision computation

If the arithmetic approximation can be treated in the same framework as any numerical computation, there remains, however, fundamental differences between exact and inexact arithmetic. Indeed the analysis of the numerical approximation is completed when the convergence of x_θ towards x is proved. However, in finite precision, it is important to know how close x_ε is to x, and a small residual might not be enough to guarantee the proximity.
Let us first mention here the convergence of iterative stationary methods to solve linear systems $(I - A)x = b$. The necessary and sufficient condition of convergence $\rho(A) < 1$ (where $\rho(A)$ is the spectral radius of A, i.e. its largest eigenvalue in modulus) in exact arithmetic may not be sufficient in finite precision. When the matrix A is normal, the necessary and sufficient condition holds in finite precision, but when A is far from normal, this may not be the case any more (see the example of successive iterations [19, 22] where the residual in finite precision cannot be reduced when the nonnormality increases). In finite precision, one deals with a family of matrices $\{A_\varepsilon\}$ for which the condition should hold uniformly with respect to ε. Therefore mathematical conditions have to be modified so that they become robust in finite precision. This is one reason why the pseudo-spectrum is a useful concept: convergence conditions based on pseudo-spectra can be more robust to finite precision than those based on the spectrum alone.

Because one works in practice with a fixed and finite ε, finite precision determines a threshold below which it is impossible to reduce the consistency error. In return, this not only impacts back on the forward error, but affects already the computability of a solution. One has to revisit the mathematical conditions to obtain a solution by taking into account the finite precision parameter: this is the way to get new conditions that are robust with respect to the arithmetic.

Moreover, for numerical approximations, the difference $F - F_\theta$ can be known. For instance, in functional analysis, an operator T may be approximated by a family of operators $\{T_n\}$ for which the difference $T - T_n$ is available in some norm. In linear algebra, applying a Krylov method on a matrix A amounts to working at each step on a matrix A' so that $A - A'$ is exactly a homotopic perturbation of rank one (see Section 3.1 of Chapter 3). However, in finite precision, the problem actually solved by the computer $(F_\varepsilon(x_\varepsilon) = y_\varepsilon)$ is not known and is, by no means, uniquely defined. This is why one is bound to design models to represent the difference $F - F_\varepsilon$. One of the reasons for the success of the normwise model is that it deals uniquely with the norm $\|F - F_\varepsilon\|$ without any knowledge of the structure of this difference; but this model may fail to capture phenomena in finite precision which are strongly structured. How faithful a given model represents the reality of finite precision computations is the central question of backward error analysis. This will be addressed in Chapter 2.

Finally, we would like to stress the importance of using relative information in finite precision / uncertain computations. Indeed, for proving the convergence of x_θ towards x

in exact arithmetic, one can consider only the absolute norm $\|x_\theta - x\|$. In finite precision, because ε is fixed, we need criteria to decide whether the residual consistency error is small enough: this is done by comparing against external gauges such as the machine precision or the data accuracy. Since these gauges are intrinsically relative, the metric to express the forward and backward errors will naturally be chosen in a relative sense.

1.1.5 Computation in the neighbourhood of a singularity

The forward sensitivity analysis which allows us to reveal the contribution of the Jacobian of H (i.e. the condition number) to the forward error (see Equation 1.4) is valid under the assumption that H has continuous first derivatives at y. This assumption is valid as long as x is not a singular point. However, despite the fact that singularities are rare, their computational influence can be fatal. In exact arithmetic, the main borderline for computability of nonlinear problems is that of the class of *well-posed* problems. Since x is a continuous function of y, a small perturbation Δy on y (such as that induced by numerical approximation) induces a small perturbation Δx on the solution. Therefore, it is reasonable to look for numerical methods which converge in exact arithmetic ($\Delta x \to 0$ as $\Delta y \to 0$). But in finite precision arithmetic, the borderline is drawn closer to enclose only the subset of the well-posed problems consisting of those problems which are *regular* enough compared to machine precision. As a result, the borderline is fuzzy and depends on the machine precision of the available computer. This leads us to define the concept of the distance to singularity viewed by the computer [19, 34].

We are particularly interested in algebraic singularities such as multiple roots of polynomials or multiple eigenvalues of matrices. These cases offer a *hierarchy* of increasing singularities: simple / double / triple / ... , which is a hierarchy of decreasing regularities. Indeed a double root is singular when compared to simple roots (more regular points) but regular when compared to a triple root (more singular point). The notions regular / singular are context-dependent: they depend on the point of view chosen. For each level, one can define some stability through the notion of the Hölder continuity which allows us to relate the forward error $|\Delta x|$ to the backward error $|\Delta y|^{1/\gamma}$, where γ is the order of the Hölderian singularity.

1.1.6 Where does the backward error analysis à la Wilkinson stand in this picture ?

Let us go back to the essential condition that ensures computability in finite precision: the direct consistency error in the neighbourhood of y has to be of the order of 1 in ε. As mentioned in Paragraph 1.1.1, this property can be rewritten from the viewpoint of the residual consistency error, which gives

$$|F_\varepsilon(t) - F(t)| \leq C(y, F)\varepsilon \tag{1.5}$$

for any t in the neighbourhood of x. The left-hand side of Inequality (1.5) expresses the distance between the arithmetic approximation F_ε and the original problem F: it is a

backward error in the sense of Wilkinson.

When one wishes to assess the reliability of numerical software, one should first establish that the backward error is indeed of order 1 in ε and exhibit a bound for the backward error which reveals the constant $C(y, F)$. The analysis of this constant $C(y, F)$ may be useful in appreciating when the computability in finite precision is severely degraded or even impossible. It is the aim of *a priori* backward error analysis, as we shall see later.

Usually, one seeks for cases where the constant $C(y, F)$ is of the order of a few units so that the backward error is of the order of machine precision. This is the optimal reliability which is often referred to as "backward stability" in the field of numerical software.

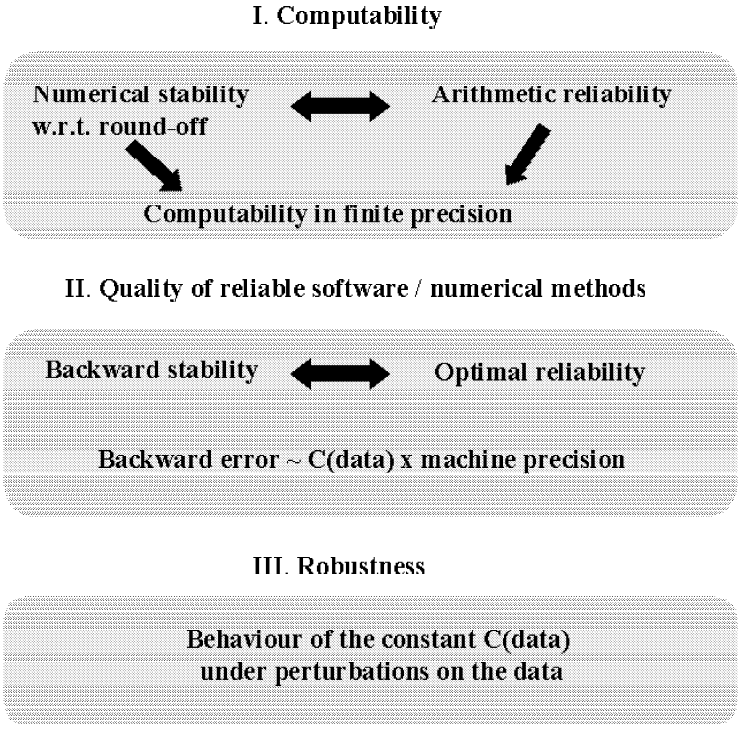

Figure 1.1: Finite precision computations at a glance

1.1.7 Conclusion

In this first part, we have presented the essence of our concept of finite precision computation. The picture given in Figure 1.1 adapted from [19] summarises the key notions that are indispensable for a complete treatment of finite precision. Amongst these notions, the most essential is the application of the Lax principle, for any numerical method as well as for the arithmetic approximation. As a consequence, all the approximation parameters, including the arithmetic one have to be coupled to determine a robust domain of com-

putability.

As already mentioned, this work is fully developed in the book *Lectures on Finite Precision Computations* with F. Chaitin-Chatelin [19]. The theory of computability is explained therein using a wide range of numerical examples. This theory is the framework of the work realised within the Qualitative Computing Group led by F. Chaitin-Chatelin at CERFACS. The book was published in 1996.
Since then, my research has been focused on backward error analysis, either for numerical software in finite precision or for approximation methods in exact arithmetic. This corresponds to level II in the diagram of Figure 1.1. The backward error analysis à la Wilkinson mentioned earlier is a complete field of research in itself. We will zoom into it in the rest of this document. The next section is devoted to explaining what are the main difficulties encountered when performing backward error analysis, mainly in the field of linear algebra.

1.2 Backward error analysis à la Wilkinson

Backward error analysis is now a well-established tool in numerical linear algebra. The idea of backward error analysis dates back to Givens [55] in the fifties and was fully developed by Wilkinson in the sixties. It was born in the context of finite precision computations. Assessing the quality of a solution computed in finite precision arithmetic has always been a major concern for software developers. In the early days of computing, it was even believed that rounding errors would render unfeasable the computer solution of large linear systems [75]; and at that time a system involving 20 unknowns was considered as large ... When practice made it clear that solving linear equations on computers was indeed feasible [44], it became necessary to design appropriate criteria to judge on the quality of the computed solutions. Equality (1.4) reveals the potential sources of inaccuracy. In particular, if a computed solution has a poor accuracy, one should be able to tell whether the inaccuracy is due either

i) either to the use of an algorithm which behaves badly in presence of rounding errors,

or

ii) to the mathematical problem to be solved, because it is highly sensitive to perturbations.

Of course, both difficulties can be present simultaneously. For a computer user, it is crucial to be able to distinguish between the two cases, because each of them calls for a different remedy. If the algorithm is responsible, one can try a better implementation or a different solution technique. It the mathematical problem turns out to be responsible for the poor accuracy, one can try to reformulate it into a less sensitive problem: this often means one must think at a higher level from which the problem is derived (the numerical discretization or even the physical problem for instance). Moreover, extended precision might help in the second case while it might be useless in the first case.

Backward error analysis was designed to allow the discrimination between both responsibilities. The essence of backward error analysis is to set the exact and the finite precision computation in a common framework by means of the following trick, which we call the Wilkinson principle [19]:

Consider the computed solution $\tilde{x}$ as the exact solution of a nearby problem.

This almost trivial idea turns out to be much more powerful than it looks at first sight.

i) It allows us to avoid the details of the computer arithmetic: the errors made during the course of the computation are interpreted in terms of equivalent perturbations in the given problem, and the computed quantities are *exact* for the perturbed problem.

ii) One advantage is that rounding errors are put on the same footing as errors in the original data. And the effect of uncertainty in data has usually to be considered in any case.

iii) It enables us to draw on powerful tools such as derivatives and perturbation theory.

iv) It allows us to factor out in the error bound, the contribution of the algorithm from the contribution of the problem.

v) Finally, it allows a great flexibility in the sensitivity analysis by providing a large choice of perturbations on the data of the problem to specify what "nearby" means.

Such an error analysis is referred to as backward error analysis because the errors are *reflected back* into the original problem. The backward error measures the minimal distance of the problem to be solved and the set of admissible perturbed problems, which is essentially done by specifying, in the original problem, the data allowed to vary and some measure for these variations.

The condition number is the favourite companion of the backward error. The condition number of a mapping measures the amplification of a perturbation on the input data. A mapping for which a small perturbation of the input data results in a large perturbation of the output data is said to be ill-conditioned. The name *condition number* is attributed to Turing although the notion of ill-conditioning was used much earlier. The first general theory of conditioning was developed by Rice [92]. As for the backward error, a condition number results from the preliminary choice of the input data allowed to vary and of some measures for these variations, together with the choice of the output data under consideration with a way to measure their variations too.

The backward error and the condition number complement each other beautifully, thanks to perturbation theory. For *regular* problems (that is problems which are not singular), they combine to give an estimation of the error on the computed solution via

the first-order bound

$$\text{forward error} \leq \text{condition number} \times \text{backward error}. \tag{1.6}$$

Indeed the backward error tells the amount of deviation from the original problem, i.e. by how much the input data have been perturbed. Then the condition number expresses how this input perturbation affects the solution. The product of the backward error and the condition number bounds the accuracy of the computed solution.

Some people restrict the term *backward error analysis* to the study of the backward error alone. This is not our choice. By backward error analysis, we will refer to both the backward error and the condition number, because both notions equally contribute to the assessment of the quality of an approximate solution. To perform a complete backward error analysis of some given problem, one should first define judiciously the mapping $d \overset{G}{\longmapsto} x$: the selection of the input data d is crucial as will be seen later. The mapping G and its datum d define the model which is meant to extract the essence of the computational process used to solve the equation $F(x) = y$. Then four steps are in order:

i) to define a metric on the input space $\mathcal{D}$ which the chosen data d belong to,

ii) to define a metric on the output space $\mathcal{X}$ where the solution x lies,

iii) to determine the backward error: given an approximate solution $\tilde{x}$, compute the distance from y to the set $\left\{\tilde{d};\ G(\tilde{d}) = \tilde{x}\right\}$ in the sense of the metric defined by *i)*,

iv) to determine the condition number: compute the norm of the Fréchet derivative of F with the metrics defined by *i)* and *ii)*.

The arithmetic quality of an algorithm is related to the size of the backward error, which should be as small as possible with respect to machine precision Ψ. The best one can do by running an algorithm on a computer is to introduce no more uncertainty than the unavoidable one that results from introducing the data onto the computer. This property is referred to as *backward stability* after Wilkinson [115]. Strictly speaking, this notion is closer to the consistency (used in theoretical numerical analysis) than to a stability notion (as a measure of sensitivity), as explained in Paragraph 1.1.2.

When the data are known with an accuracy of the order of η, then the backward error should be compared with η. This is important when η is significantly larger than machine precision and has a practical impact on iterative methods whose convergence can be monitored via the backward error. Such a situation is common in most applications outside mathematics: for instance the uncertainty in physical measurements leads to inaccurate data. Of course, the accuracy of the computed solution may decrease accordingly but, in many cases, what really matters is that the computed solution exactly solves a perturbed problem, indistinguishable from the original one at the level of the data uncertainty.

Since its introduction and development by Wilkinson in the sixties, backward error analysis has proved to be the tool of choice for the analysis of the numerical behaviour of algorithms in finite precision, and is now fully incorporated into high performance numerical libraries, at least as far as linear algebra is concerned. The package LAPACK [1], which aims at solving most problems in linear algebra, contains mainly algorithms which are guaranteed to be backward stable. The user is often provided with backward errors and condition number estimates (with different metrics) which make him or her able to have a critical view of the delivered numbers. Using a software package as LAPACK not only ensures getting the best performance on modern architectures, but also draws the best from the finite precision arithmetic, because the software is built upon two key principles: an efficient cache management exploiting the memory hierarchy of the computer *and* backward error analysis. The rest of this document, in illustrating various subtleties of backward error analysis, should hopefully convince the reader that the development of numerical software should be left to the best experts, and should make him or her able to appreciate the fine art put into it.

The scope of backward error analysis, and its related techniques, has greatly evolved in the last forty years, and yet this topic remains a very active field of numerical analysis and even more in linear algebra. In this document, we wish to review the recent advances in backward error analysis, that we structure into two main directions.

First, we will address in Chapter 2 the progress made in the use of backward error analysis for finite precision computations, mainly in the field of numerical linear algebra. As we mentioned earlier, backward error analysis is very much dependent on the initial specifications such as the input and output data and their corresponding metrics. A major difficulty is to make the appropriate choices so that the derived condition number and backward error produce a good error prediction. The choice of the model for the perturbations (data, metrics, structure) chosen to derive the condition number and the backward error is crucial. Indeed the accuracy of the error prediction given by the first-order inequality (1.6) depends on the ability of this model to represent faithfully the perturbations actually created by the algorithm when run in finite precision. The last twenty years have seen a tremendous refinement in tuning the models in order to better capture the behaviour of algorithms in finite precision arithmetic, leading to more and more complex and accurate models. We will give a series of examples where these refined models succeed in providing tighter error predictions. At the same time, the more complex the model, the more difficult the backward error analysis. The techniques for establishing the formulae for the backward error and the condition number have also evolved and we will show examples of systematic approaches that can help in deriving these formulae.

At this point in our writing, we have been intentionally ambiguous and used the terms "computed solution" or "approximate solution" indiscriminately. If backward error analysis was primarily designed for a better understanding of finite precision computation, nothing in its essence restricts it to finite precision. It can serve to assess the quality of *any* approximate solution, wherever the approximation comes from. For instance, back-

ward error analysis can be used to analyse various kinds of perturbations such as physical perturbations. It is also the backbone of eigenvalue and pseudo-eigenvalue analysis, an area which has been blossoming for the last fifteen years. It can also be used to study the convergence, in exact arithmetic, of approximation methods. In Chapter 3, we will illustrate the ever-growing role of backward error analysis outside finite precision.

We hope that this document gives the reader a flavour of the power and beauty of backward error analysis.

Chapter 2

Backward error analysis for finite precision computations

We mentioned in the introductory chapter that a major issue in analysing finite precision computations is that one never knows explicitly the problem which is actually solved by the computer. A way to circumvent this difficulty is to work with a model which is supposed to represent as faithfully as possible the behaviour of an algorithm in finite precision. The difficulty in designing a good model is a recurrent question in this chapter. The first four sections show how the models used in linear algebra are being refined to increase the performance of backward error analysis in numerical software. In the fifth section, we address the particular case of singular problems which require a specific treatment for the forward error is not anymore proportional to the backward error. We then present the notion of pseudosolution which stems directly from the backward error, and which has proved to be extremely useful to address the problems of conditional stability and robust convergence conditions evoked in the previous chapter. We end this chapter by an overview of the toolbox PRECISE.

2.1 Models

In order to go further in the subtleties of backward error analysis, we need to introduce some notations and definitions. Let G be a regular mapping such that

$$\begin{array}{rcl} G : \mathcal{D} & \rightarrow & \mathcal{X} \\ d & \longmapsto & x. \end{array}$$

We denote by $\|.\|_{\mathcal{D}}$ (resp., $\|.\|_{\mathcal{X}}$) a metric on the data (resp., solution) normed linear space $\mathcal{D}$ (resp., $\mathcal{X}$). The subordinate norm $\|L\|_{\mathcal{D},\mathcal{X}}$ of a linear operator L is defined by

$$\sup_{d \neq 0 \in \mathcal{D}} \frac{\|Ld\|_{\mathcal{X}}}{\|d\|_{\mathcal{D}}}.$$

Definition 2.1.1 *The condition number of G at d is*

$$K = \|G'(d)\|_{\mathcal{D},\mathcal{X}}$$

where $G'(d)$ is the Fréchet derivative of G at d.

Definition 2.1.2 *Let $\tilde{x}$ be an approximate solution of the equation $G(d) = x$. Then the backward error associated with $\tilde{x}$ is defined by*

$$\eta(\tilde{x}) = \min_{\Delta d \in \mathcal{D}} \{\varepsilon \geq 0;\ G(d + \Delta d) = \tilde{x};\ \|\Delta d\|_{\mathcal{D}} \leq \varepsilon\}\,.$$

Some remarks are now in order.

1. The characterisation of the condition number as the norm of the Fréchet derivative is due to Rice [92]. The condition number measures the stability of the mapping G at the point d.

2. The definition of the mapping $d \overset{G}{\longmapsto} x$ relies upon implicit choices which are of major importance. Let us take an example from the solution of linear systems. If A is a square $n \times n$ matrix with no particular properties (apart from being invertible), the linear system $Ax = b$ where x and b are two vectors of length n can be solved by a generic algorithm such as Gaussian Elimination with partial pivoting (GEPP). In such a context, choosing the mapping $A \longmapsto x = A^{-1}b$ seems a reasonable way to get useful information on the sensitivity of the linear system and on the quality of its computed solution. However, if A possesses strong structural properties, say A is a Vandermonde matrix defined solely by its generic coefficients $a = \{a_1, \ldots, a_n\}$, then one should use the appropriate Björck-Pereyra algorithm [7] which works directly on the coefficient vector a rather than on the entries of the matrix $A_{ij} = a_i^{j-1}$ [19, 71]. Clearly, the mapping $a \longmapsto x = A^{-1}b$ should expose the relevant information. In particular, the corresponding backward error will indicate whether an approximate solution solves a nearby Vandermonde system, as opposed to a nearby arbitrary linear system.

3. The definition of the backward error as a min holds because of the regularity properties of G. By convention, the backward error is infinite whenever the set $\left\{\tilde{d} \in \mathcal{D};\ G(\tilde{d}) = \tilde{x}\right\}$ of admissible perturbations is empty.

4. The condition number defined above is absolute (resp. relative) if the metrics defined on $\mathcal{X}$ and $\mathcal{D}$ are absolute (resp. relative). The same remark holds for the definition of the backward error.

For a choice of a mapping and for a choice of the metrics, one is able to perform a backward error analysis which amounts to deriving the condition number and the backward error according to Definitions 2.1.1 and 2.1.2 respectively; stated in this way, it is a purely mathematical game. The expertise of the numerical analyst comes when choosing the right mapping and the right metrics to serve a particular purpose which can be, for instance, to understand the numerical solution of some mathematical problem with some algorithm in finite precision. By "understand the numerical solution", we mean at least

- anticipate the presence of ill-conditioning and determine its causes if possible,
- assess the backward stability of the algorithm,
- predict the forward error.

Great care must be taken in choosing the model of perturbations, that is the data and the metrics, so that the model faithfully represents the behaviour of the algorithm in finite precision arithmetic. Even for linear systems, which are the most studied problem of linear algebra, the design of an efficient model is not a trivial task. The first model used to perform a backward error analysis on a square linear system is the *normwise* model on the mapping $(A, b) \longmapsto x = A^{-1}b$. The normwise metric on the input space $\mathcal{D} = \mathbb{C}^{n\times n} \times \mathbb{C}^n$ is

$$\|(\Delta A, \Delta b)\|_{\mathcal{D}} = \max \left\{ \frac{\|\Delta A\|}{\|A\|}, \frac{\|\Delta b\|}{\|b\|} \right\}$$

and the normwise metric on the output space $\mathcal{X} = \mathbb{C}^n$ is

$$\|\Delta x\|_{\mathcal{X}} = \frac{\|\Delta x\|}{\|x\|}$$

where $\|\ .\ \|$ is a classical subordinate norm in $\mathbb{C}^n$. The relative normwise condition number turns out to be

$$K_N = \frac{\left\|A^{-1}\right\| (\|A\|\,\|x\| + \|b\|)}{\|x\|} \tag{2.1}$$

and the normwise backward error associated with an approximate solution $\tilde{x}$ [93] is

$$\eta_N(\tilde{x}) = \frac{\|A\tilde{x} - b\|}{\|A\|\,\|\tilde{x}\| + \|b\|}. \tag{2.2}$$

For a long time, this model due to Wilkinson has been the only one available to analyse the finite precision solution of linear systems. However, its limitations were soon appreciated. In particular, K_N depends upon a diagonal scaling of the data A and b: it is possible to scale a linear system so that its normwise condition number becomes arbitrarily large, whereas the accuracy of its numerical solution may not affected if an appropriate linear solver is used. Moreover, when the first sparse solvers were developed, the researchers started wondering if a sparse solver in finite precision would solve a nearby sparse problem with the same sparsity pattern as the original one. More than fifteen years after Wilkinson's work, the tools to answer these questions were given by Skeel [99] with the design of a *componentwise* model inspired by the earlier work of Bauer [4]. In the componentwise model, the metric on the input space is

$$\|(\Delta A, \Delta b)\|_{\mathcal{D}} = \min_{\varepsilon>0} \{|\Delta A| \leq \varepsilon\,|A|\,,\ |\Delta b| \leq \varepsilon\,|b|\}$$

where the inequality $|\Delta A| \leq \varepsilon\,|A|$ means $|\Delta A_{ij}| \leq \varepsilon\,|A_{ij}|$ for any i and j. The metric on the output space is

$$\|\Delta x\|_{\mathcal{X}} = \frac{\|\Delta x\|_\infty}{\|x\|_\infty}.$$

The resulting condition number is

$$K_C = \frac{\left\| \left|A^{-1}\right| \left(|A|\,|x| + |b|\right) \right\|_\infty}{\|x\|_\infty}, \tag{2.3}$$

and the backward error associated with the approximate solution $\tilde{x}$ [87] is

$$\eta_C(\tilde{x}) = \max_{i=1,\ldots,n} \frac{|A\tilde{x} - b|_i}{(|A|\,|\tilde{x}| + |b|)_i}. \tag{2.4}$$

This backward error analysis is independent of any diagonal row scaling of the linear system. Moreover, the admissible perturbations have the same sparsity pattern as A, otherwise the backward error is infinite. The componentwise model meets its objectives and has been extremely successful in the numerical study of direct solvers (either dense or sparse) based on GEPP. The backward error returned by the general purpose linear solver in LAPACK is the componentwise backward error. Arioli, Demmel and Duff were the first to use it for sparse systems.

If the componentwise model prevails for linear systems solved by the Cholesky or the GEPP algorithm, the normwise model still performs well on other kinds of linear solvers such as QR factorisation, or iterative linear solvers such as Krylov-type methods [2]. In the latter case of iterative solvers the backward error is, in addition, an excellent stopping criterion; convergence can be declared when the backward error has reached machine precision (at best) or becomes smaller than the data uncertainty (if any). A residual alone would not serve as a stopping criterion because its size cannot be checked against any reference. For a linear system, the backward error is only a normalised residual (see Equalities 2.2 and 2.4), but the normalisation is the key!

In joint work with CNES [45], we have developed a GMRES implementation for complex matrices. This code was then further improved to deal with any type of arithmetic and to satisfy some software quality requirements enabling its efficient use on sequential and shared or distributed memory parallel computers. In particular, convergence is monitored though the normwise backward error, that the user can scale according to his or her needs. When preconditioning is used, the normwise backward error on the preconditioned system is also made available. The resulting packages [46, 47] have been put in the public domain, `www.cerfacs.fr/algor/Softs/`, with a non-commercial licence agreement. It is regularly downloaded by many researchers working in various areas ranging from geophysics or ocean modelling to theoretical physics. In particular, the complex version has been recently integrated into a public domain circuit simulator developed at Bell Labs.

In the light of these comments, one can see how delicate is the task of finding the appropriate model of perturbation in order understand the behaviour of an algorithm in finite precision. Given an algorithm, two situations arise:

1. one may want to assess the backward stability of the algorithm under study and to identify, whenever possible, the cases where the algorithm may fail to be backward

stable. This is usually the job of software developers. This approach can be called *a priori* backward error analysis, or

2. given a computed result, one may want to check whether it is the solution of a nearby problem. A formula for the backward error is needed. Any user of numerical algorithms, and in particular engineers, are likely to be in this situation. Deriving formulae for backward errors and condition numbers can be called *a posteriori* backward error analysis.

In the next two sections, we describe in detail these two different aspects of the backward error analysis.

2.2 *A priori* backward error analysis

The software designer wants to assess the backward stability of a newly designed algorithm. In particular, he or she may want to anticipate in which cases (that is often for which class of data) an algorithm may fail to deliver a backward stable answer. For such a purpose, the software designer has to cope with the propagation of the round-off errors, in order to appreciate how every floating-point operation reflects back as a perturbation onto the data: the details of the arithmetic are not hidden any more but this is the price to pay to obtain general results on the backward stability.

For instance, one of the first results achieved with this approach for linear systems is due to Wilkinson. We set $\gamma_n = nu/(1-nu)$ where u is the unit round-off. Let $\tilde{x}$ be the computed solution of the linear system $Ax = b$ with GEPP. Then $\tilde{x}$ solves exactly the perturbed linear system $(A + \Delta A)\tilde{x} = b$ with $\|\Delta A\|_\infty \leq 2n^2\gamma_n\rho_n \|A\|_\infty$, where ρ_n is the growth factor (ratio between the largest element computed during the factorisation and the largest entry of A, in modulus). This result ensures that, provided that the growth factor is reasonable, GEPP is normwise backward stable. Further work has been carried out after Wilkinson's result to identify the classes of matrices with large growth factors. Apart from contrived examples, very few cases of practical origin showing large growth factors have been reported [72].

A priori backward error analysis requires an intimate knowledge of algorithms and finite precision. Wilkinson performed the *a priori* backward error analysis of many fundamental algorithms in numerical analysis, including linear systems, eigenproblems or polynomials [116, 117]. Higham took over from him and has revisited backward error analysis in a really enjoyable recent book [72] where *a priori* backward error analyses are given for most problems and major algorithms in linear algebra, except for eigenproblems. The clarity of the text and the demonstrations stems from a hierarchical use of backward analysis results. The starting point is the backward analysis of an elementary floating-point operation. With the IEEE standard, the computed result $\mathrm{fl}(x \text{ op } y)$ satisfies

$$\mathrm{fl}(x \text{ op } y) = (x \text{ op } y)(1+\delta)$$

where $|\delta| \leq u$ and op is any of the four basic operations $+$, $-$, $\times$ and $/$. Building upon this property, one is able to establish the componentwise backward stability of inner products, compensated summation, solution of triangular systems with backward or forward solves. When tackling more complex problems, one should try to make as much use as possible of the above formalism and the results previously obtained. Higham's book has also proved to be very helpful in teaching because it provides the reader with many results and references.

2.3 *A posteriori* backward error analysis

A posteriori error analysis aims at deriving formulae for the backward error and the condition number once the model of perturbation has been defined. Recent advances in the techniques used for establishing these formulae have been achieved. We review them now.

The condition number of the mapping $d \overset{G}{\longmapsto} x$ is the norm of the Fréchet derivative. When the model of perturbations (that is the metrics) and the mapping are simple enough, the condition number can be obtained in two steps:

1. derive a first-order Taylor expansion of the mapping G at d and bound it with the given metrics,
2. exhibit a particular perturbation of the data which achieves the bound.

This approach is particularly successful in many problems of linear algebra (linear systems, matrix inversion, polynomials, eigenvalues) with the normwise and componentwise models.

However, for more complex problems and/or more complex metrics, the second step of this approach may not be practical. This happens in particular when dealing with very structured perturbations.

A step forward has been made in the recent work of Gratton [63] following [31], making use of the Kronecker product. Consider the mapping

$$\begin{array}{rcl} G : \mathbb{R}^{m\times n} & \to & \mathbb{R}^{r\times s} \\ d & \longmapsto & x, \end{array}$$

where the input space $\mathbb{R}^{m\times n}$ and the output space $\mathbb{R}^{r\times s}$ are both equipped with the Frobenius norm. Let $M(G,d)$ be the matrix of the linear application which represents the derivative $G'(d)$ in orthonormal bases of $\mathbb{R}^{m\times n}$ and $\mathbb{R}^{r\times s}$. Then the condition number of G is $\|M(G,d)\|_2$. The Kronecker product arises naturally when one looks for an expression of $M(G,d)$ in vector form so that the 2-norm is easily computable. Using this approach, Gratton was able to derive the condition number of several mathematical problems, such as

- the matrix exponential [81, 63],
- linear least-squares [62],

- the L and U factors of the LU factorisation of a square matrix [63],
- the Q and R factors of the QR factorisation of a square matrix [29, 63],
- the factors of the polar factorisation of a rectangular matrix [24].

Using the same approach, we have performed an analysis of the factors of the $L^T DL$ factorisation, in the context of a collaboration with CNES [17]. The accurate computation of the trajectory of a satellite using GPS measurements can be done using the LAMBDA method which basically solves an integer least-squares problem. The core of the LAMBDA method is an $L^T DL$ factorisation. We were able to derive the condition number for the L and D factors and use them to check the accuracy of this key step. Using real data, we were able to confirm the validity of their results to CNES.

The derivation of a formula for the backward error is usually more difficult than for the condition number, because in general it is a nonlinear problem. If the backward error reduces to a normalised residual in various problems of linear algebra (such as linear systems or eigenproblems), there are many cases where its formulation is very complex [50, 62]. The knowledge of a formula for the backward error and the condition number is also very useful for software designers. Sometimes, the formulae themselves reveal the source of potential numerical dangers. For instance, for the normwise relative condition number of the eigenvalue problem $Ax = \lambda x$, which is

$$K(\lambda) = \frac{\|A\|_2 \, \|x_*\|_2 \, \|x\|_2}{|\lambda| \, |x_*^* x|},$$

one sees that ill-conditioned problems are those where the left eigenvector x_* becomes orthogonal to the right eigenvector x, that is when the eigenvalue λ is close to being defective.
When the formulae are more obscure, it might be interesting to turn to an approach which consists in numerically finding data which maximise the condition number or the backward error. This can be done for instance through a direct search algorithm [73]. We found it particularly useful in practice and easy to do with MATLAB which offers well-suited optimisation routines.

2.4 Two examples of backward error analysis

It this section, we present two examples of our work in backward error analysis for linear algebra. The first one concerns the generalised eigenvalue problem. The second one deals with linear systems of the type $A^*Ax = b$ that arise for instance in Krylov methods for shifted and inverted eigenproblems. Both cases illustrate the improvement that can be gained in the understanding of the algorithmic behaviour in finite precision by appropriately choosing the mapping to be studied.

We give a quick overview of each of them. For the complete treatment, the reader is referred to the articles

- A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$, with V. Toumazou, published in the *Journal of Numerical Linear Algebra with Applications* [52], and
- Structured backward error and condition number for linear systems of the type $A^*Ax = b$, with S. Gratton and V. Toumazou, published in *BIT* [50].

Both articles are given in the Annexes.

2.4.1 Backward error analysis for the generalised eigenproblem

The generalised eigenproblem (P) : $Ax = \lambda Bx$, where A and B are two square matrices, often appears at the heart of physical applications. The problem (P), or equivalently the matrix pair (A, B), is said to be regular if there exists at least one complex number z for which $\det(A - zB)$ is nonzero. Otherwise it is called singular. Computing eigenvalues of singular matrix pairs belongs to the class of ill-posed problems [40, 41], i.e. their eigenvalues can change discontinuously as functions of A and B. In this work [52], we restrict our study to eigenproblems for regular matrix pairs which are well-posed problems. Regular matrix pairs (A, B) for which B is nonsingular have the same spectrum as the matrix $B^{-1}A$. When B is singular, the matrix pair admits at least one infinite eigenvalue.

In [101], Stewart and Sun establish the condition number of the regular generalised eigenproblem for the choice of a specific norm, namely the *chordal norm*. The main motivation for using such a norm is to treat finite and infinite eigenvalues within the same formalism. In such a case, it is useful to represent the eigenvalue as a pair of numbers and the chordal norm arises naturally to cope with this situation. The price to pay to handle infinite eigenvalues is to deal with this very special norm, the chordal norm, which is difficult to interpret in terms of physical perturbations.

However many situations arise in practice where one seeks a finite eigenvalue only, or where at least one of the two matrices A or B is nonsingular. It is then possible to use more *natural* norms to represent perturbations on the data A and B and on the solution λ, that is norms appropriate for easily representing physical perturbations. Our formulation relies upon a measure which has the property to be flexible because it can be absolute or relative with respect to the matrix A and/or B according to the user's prescription. The distance between the two matrix pairs (A, B) and (A', B') is defined by

$$\delta = \min\left\{\omega > 0; \left\|A - A'\right\| \leq \omega\alpha \text{ and } \left\|B - B'\right\| \leq \omega\beta\right\}.$$

In the context of perturbation analysis, keeping A (resp., B) unchanged amounts to set α (resp., β) to zero. Setting $\alpha = \|A\|$ and $\beta = \|B\|$ gives the classical normwise measure. Using this metric, we study the mapping $(A, B) \longmapsto (\lambda, x)$ for which we formulate the backward error and the condition number. This approach leads to two improvements in comparison with what existed in the literature:

- it performs better than the chordal metric in the prediction of the forward error,

- it is more appropriate for the analysis of the QZ algorithm than to apply the standard model ($C \longmapsto (\lambda, x)$) on the eigenproblem $Cx = \lambda x$ with $C = B^{-1}A$. In particular, it allows us to check whether an appproximate solution is the exact one of a nearby *generalized* eigenproblem.

This approach has been further developed for componentwise analysis [70] and extended to polynomial eigenproblems [102].

2.4.2 Linear systems of the type $A^*Ax = b$

We consider linear systems of the type

$$A^*Ax = b \tag{2.5}$$

where A is an $m \times n$ *rectangular* and complex matrix (with $m \geq n$) and x and b are two complex n-vectors. We assume that A is of full rank n and we denote its pseudo-inverse by $A^\dagger$. Such linear systems arise in a number of applications which require the computation of the smallest singular value $\sigma_{\min}$ of a large matrix A. A possible way to do it is to compute the smallest eigenvalue $\lambda_{\min}$ of A^*A with a Lanczos method with invert (then $\sigma_{\min} = \sqrt{\lambda_{\min}}$). Therefore, the computation of the Krylov basis requires us to solve repeatedly linear systems such as (2.5). A particular case of interest is the computation of the two-norm of resolvent matrices $(A - zI)^{-1}$. For instance, the computation of spectral portraits consists in evaluating $\left\|(A - zI)^{-1}\right\|_2 = \sigma_{\min}(A - zI)$ for z in a prescribed region of the complex plane. Systems of the form $A^*Ax = b$ occur also in other contexts, such as constrained linear least-squares problems [6] or bound constrained quadratic programming problems [85, 86].

Our purpose is to be able to predict reliably the computing error made in solving (2.5) in finite precision. This can be done classically by means of backward error analysis [116], provided that one has access to a condition number and the associated backward error for (2.5).

Classical sensitivity analysis of (2.5) leads to the conclusion that the condition number K is of the order of the square of the condition number of A, where the condition number of A is defined by $K_2(A) = \|A\|_2 \left\|A^\dagger\right\|_2$. Indeed, $K_2(A^*A) = \|A^*A\|_2 \left\|(A^*A)^{-1}\right\|_2 = \left(\|A\|_2 \left\|A^\dagger\right\|_2\right)^2 = K_2^2(A)$. Behind this result, there is the assumption that the linear system (2.5) is subject to normwise perturbations on the matrix $B = A^*A$. This condition number is relevant to predict a computing error when the algorithm used to solve (2.5) implicitly generates such global perturbations in finite precision. We can expect that this is the case if we form the product $B = A^*A$ explicitly and then solve (2.5) by means of a Cholesky factorisation.
However, most algorithms used to solve (2.5) never build the product A^*A. They rather involve either a factorisation of A (in the case of direct methods) or matrix-vector multiplications A^*x and Ax (in the case of iterative methods). Therefore, a condition number useful for predicting the error in the numerical solution of (2.5) by such algorithms should

take into account structured perturbations such as perturbations in the matrix A only, and the mapping under study should be $A \longmapsto x$ rather than $B \longmapsto x$.

Using the technique of the Kronecker product mentioned in Section 2.3, we derive the condition number and the backward error of the mapping $A \longmapsto x$ such that $A^*Ax = b$. The structured condition number of (2.5) can be as large as a quantity of the order of $K_2^2(A)$ when the solution x or the right-hand side b are in the direction of the right singular vector of A associated with the largest singular value. But choosing a solution or a right-hand side in the direction of the right singular vector associated with the smallest singular value of A permits to keep the structured condition number of (2.5) of the same order as the condition number of A. We illustrate that our model is appropriate on algorithms that do not explicitly build the matrix A^*A. Numerical experiments that show the relevance of this condition number in the prediction of the computing error are provided in [50].

2.5 Singular problems

2.5.1 Backward error analysis for well-posed singular problems

Well-posed problems which are so smooth that the solution is at least a C^1 function of the data, have been defined as **regular** problems. Their solutions computed in finite precision arithmetic with a reliable algorithm are generally good approximations of the exact ones, provided that the problems are not too ill-conditioned.

Any problem which is not regular is called **singular**. It is in the neighbourhood of such singularities that computational difficulties arise. Singular problems are not generic: under perturbations (such as those induced by finite precision arithmetic), they are usually transformed into regular problems which then appear as classically ill-conditioned (their condition number tends to infinity).

If one wants to compute a multiple root of multiplicity m, it is perturbed by finite precision computation into a set of m (often distinct) approximate roots. Each individual root is ill-conditioned: its condition number approaches infinity. Any method which is stable at a regular point becomes unstable with respect to a *linear* analysis at a singular point. In order to recover some stability, one has to consider the Hölder-continuity of order $\gamma = 1/m$: the error $\|\Delta x\|$ on the solution is proportional to $\|\Delta z\|^\gamma$, the γ^{th} power of the size of the perturbation on the data.

If one is not primarily interested in each individual approximation, one can recover stability by grouping the m approximate roots, that is, by taking the computed arithmetic mean as an approximation to the multiple root.

The linear condition number, which corresponds to some measure of the Fréchet derivative (see Definition 2.1.1) is a measure of stability for a regular problem. It tends to infinity whenever the problem approaches a singularity. For Hölder singularities of well-posed problems, we can likewise define normwise Hölderian condition numbers. It

measures the intrinsic difficulty of computing a multiple root, knowing its order γ, and taking into account the fact that the error $\|\Delta x\|$ is now proportional to the γ^{th} power of the size of the perturbation on the data.

Condition number

In linear algebra, the most important singular problems of Hölder type are multiple roots of polynomials and defective eigenvalues.
For roots of polynomials, the Hölder condition number was established by Wilkinson [116] for a componentwise metric. Let ξ be a root of multiplicity m of the polynomial $p(x) = \sum_{k=0}^{n} a_k x^k$. Then its Hölder condition number is

$$K^C_{1/m} = \frac{1}{|\xi|} \frac{m!}{\left|p^{(m)}(\xi)\right|} \left(\sum_{k=0}^{n} \left| a_k \xi^k \right| \right)^{1/m}.$$

A similar formula for the normwise condition number has been established by Gratton [63].

The Hölder condition number for a defective eigenvalue of index l is

$$K_{1/l}(\lambda) = \left\| D^{l-1} \right\|_2^{1/l},$$

where $D = (A - \lambda I)P$ is a nilpotent matrix and P is the spectral projection. Such a result can be derived directly from the Laurent expansion of the resolvent norm near λ [69, 76]. It can also be derived from the work of Lidskii [80] about the Puiseux series expansions of defective eigenvalues. This alternate derivation has been done by Moro, Burke and Overton [83]. The relationship between the Hölder condition number and the stratification diagram for defective eigenvalues is examined in [26].

Backward error

The distance to singularity is a backward error by definition: it is indeed the size of the perturbation to the data to render the problem singular. Note that the classical definition of the backward error is not affected by the singularity: the formula for the backward error associated with an approximate eigenvalue $\tilde{\lambda}$ or an approximate polynomial root $\tilde{x}$ hold even if one tries to compute a multiple root or a defective eigenvalue. It is because the backward error answers the question: what is the distance to the nearest polynomial (resp., eigenproblem) for which $\tilde{x}$ (resp., $\tilde{\lambda}$) is an exact root (resp., eigenvalue) ?
It would be different, of course, if one was looking for the nearest polynomial (resp., eigenproblem) having $\tilde{x}$ (resp., $\tilde{\lambda}$) as a multiple root (resp., defective eigenvalue).
This subtle difference has had some consequence on the definition of the distance to singularity, that we address in the next paragraph.

2.5.2 Condition number and distance to singularity

Regular problems which are ill-conditioned are close to singularity. But which singularity ? And how close ? This question of the distance to singularity has been central in numerical linear algebra for the past few years. There have been many efforts to relate the distance to singularity with the condition number. And in many cases, the question remains open. Several difficulties arise. First, the metrics chosen have a great impact on the complexity of the answers. Second, the definition of the singular problem is not always unique.

Let's start by an (apparently) simple example: the matrix inversion $A \longmapsto A^{-1}$. It is easy to define the set of singular problems: it is the set of singular matrices. With a normwise measure, the distance to singularity of a regular matrix is $1/(\|A\| \left\|A^{-1}\right\|)$. This theorem, often attributed to Turing, seems to have already been proved by Banach in the 20s. It was rediscovered by Eckart and Young in 1939 with the 2-norm, and proved by Gastinel for an arbitrary norm. It shows that the normwise condition number of a matrix is the reciprocal of the distance to singularity. The higher the condition number, the closer the matrix to singularity. In view of this simple and nice result, many efforts have been made to relate the componentwise condition number with the componentwise distance to singularity. The componentwise distance to singularity

$$\delta_C(A) = \min_{\varepsilon} \{ |\Delta A| \leq \varepsilon \, |A| \, ; \ (A + \Delta A) \text{ is singular } \} = \eta_C(0)$$

has been characterised by Rohn [94]. The evaluation of this formula has been proved to be of NP complexity [88]. The term $\eta_C(0)$ is the componentwise backward error associated with the approximate solution $z = 0$ for the problem consisting in finding the smallest eigenvalue (in modulus) of A: indeed, a singular matrix has a zero eigenvalue. It can be easily shown that the componentwise distance to singularity is always larger than the inverse of the componentwise condition number $K_C(A)$, which means that a well-conditioned matrix is far away from a singular matrix, in a componentwise sense. But what about the reciprocal statement ? Demmel [39] shows by complexity arguments that no simple relationship can hold between $\delta_C(A)$ and $K_C(A)$. Higham and he also conjecture that the ratio between $\delta_C(A)$ and the minimum condition number achievable by a diagonal scaling of A cannot be larger than a factor of the problem size n. This is what we also observed on our computer experiments using the toolbox PRECISE and reported in [19]. This conjecture has been proved by Rump [95, 96] who establishes that a matrix which is ill-conditioned in a componentwise sense is also close, in a componentwise sense, to a singular matrix. However, he also shows that, if the matrix is symmetric, and if the componentwise perturbations are constrained to preserve the symmetry, then an ill-conditioned matrix can be arbitrarily far from the nearest singular matrix, in this structured componentwise sense [97].

Matrix inversion is a typical example where the singular problem is uniquely defined and the complexity arises from the metric. Let's turn to examples where the nearest singular problems have been defined in various ways: this is the case in the literature of polynomials and eigenproblems. As suggested in the previous section, there are two ways

of defining the distance to singularity for a polynomial p having a root x:

$$\begin{aligned} \Delta &= \min\{\|\Delta p\|\,;\ (p+\Delta p) \text{ has a multiple root at } x\}\,, \\ \delta &= \min\{\|\Delta p\|\,;\ (p+\Delta p) \text{ has a multiple root } y \text{ near } x\}\,. \end{aligned}$$

Both definitions produce a different relationship with the condition number. Because of the constraint on p in Δ (x is exactly a multiple root of $(p+\Delta p)$), Δ is proportional to the reciprocal of the condition number. On the contrary, δ is proportional to some power of the reciprocal of the condition number. In the paper [34], co-authored with F. Chatelin and T. Braconnier and published in the *Journal of Numerical Functional Analysis and Optimization*, we investigate and illustrate this difference and we give a conjecture that links δ with the condition number. This paper is given in the Annex. Using a simple 3 by 3 matrix whose Jordan structure is parameterized, we give arguments of why we believe that the definition of δ is most useful *in the context of finite precision computations.* In finite precision computations, an essential parameter is what we call the "distance to singularity viewed by the computer" [19], that is the maximal size of perturbations which a regular problem can be subjected to, while remaining regular for the computer. Such a distance is often in practice a good approximation to δ: this reflects the fact that the actual computation at a regular point x is influenced by the presence of the closest singularity y in the neighbourhood of x. Gratton [63] has checked our conjecture in the context of polynomials in one variable by giving a formulation of the distance to singularity after refining the definition of δ, where the constraint "y near x" has to be specified mathematically.

2.6 Backward error and pseudosolutions

When we compared exact and finite precision computations in Section 1.1.4 of the Introduction, we emphasised the fact that some mathematical conditions (such as a necessary and sufficient conditions for convergence), established in exact arithmetic, may not be robust in finite precision. The notion of pseudosolution that we present now is one way to design more robust conditions.

Let $F(x) = y$ be the equation to be solved and $G(d) = x$ be the mapping under consideration to obtain the solution x. Let $\tilde{x}$ be an approximation of the solution x. The backward error $\eta(\tilde{x})$ associated with $\tilde{x}$ has been defined (see Definition 2.1.2) as the size of the smallest perturbation $\Delta d = \tilde{d} - d$ of the data d such that $\tilde{x}$ is the exact solution of the perturbed problem $G(\tilde{d}) = \tilde{x}$. The explicit formula for $\eta(\tilde{x})$ depends on the metrics.

Definition 2.6.1 *The set of the pseudosolutions for the mapping $G(d) = x$ is defined by*

$$\Sigma_\varepsilon(x) = \{z;\ \eta(z) \le \varepsilon\}$$

for any $\varepsilon > 0$.

The set of ε-pseudosolutions represents all the points in the solution space $\mathcal{X}$ that can be considered as an approximate solution of the equation with a backward error less than

ε. In particular, if the computed solution $\tilde{x}$ is such that $\eta(\tilde{x}) = \varepsilon$, then $\Sigma_\varepsilon(x)$ represents the set of points z in the neighbourhood of x which are indistinguishable from x by the computation. The notion is related to that of *domain of uncertainty* for the root x [113]. The larger the set $\Sigma_\varepsilon(x)$ around x, the more unstable the mapping G at d. Knowledge of x is not required; it is sufficient that the set $\Sigma_\varepsilon(x)$ be large.
It may be interesting to look at the border $\partial\Sigma_\varepsilon(x)$ of $\Sigma_\varepsilon(x)$ and determine whether it encloses a large region. When the border $\partial\Sigma_\varepsilon(x)$ is not computable because, for example, one has no explicit formula for the backward error or because each component of $\eta(\tilde{x})$ depends on the whole vector $\tilde{x}$, one can alternatively plot the sample of perturbed computed solutions obtained by random perturbations of the data of amplitude less than ε.

The graphical display of $\partial\Sigma_\varepsilon$ provides an informative description of ill-conditioning. Indeed, the first-order bound (1.6) for a given backward error would define a circle around the exact solution. The shape of the set of ε-pseudosolutions gives a more global and structured view of the effects of ill-conditioning by expressing how it diffuses in the solution space.

There are two problems which lend themselves very naturally to a graphical display of $\partial\Sigma_\varepsilon$. They consist of the computation of the set $Z = \{x_i\}_{i=1,\ldots,n}$ of roots of polynomials, and of the spectrum $\sigma = \{\lambda_i\}_{i=1,\ldots,n}$ of matrices, because both sets are in $\mathbb{C}$. The first one has been addressed by Mosier [84] under the name of *root loci*, used in linear control theory; see also Hinrichsen and Kelb [74] and Toh and Trefethen [104]. In [19], we compare the set of pseudo-zeros of polynomials for normwise and componentwise perturbations, and examine the relationship with pseudoeigenvalues of the associated companion matrices.
The second problem, that is the set of pseudoeigenvalues, has been a very active area of numerical analysis in the last ten years. We address it now.

The *ε-pseudospectrum* Σ_ε of a matrix A is the set defined by

$$\Sigma_\varepsilon = \{z \in \mathbb{C};\ \eta(z) \leq \varepsilon\}$$

where $\eta(z)$ is the backward error associated with the approximate eigenvalue z. The *spectral portrait* of a matrix [19, 57] consists of a graphical display of the map $z \longmapsto \eta(z)$ in a prescribed region of the complex plane: its level curve at the value ε is the border of Σ_ε. The definition of Σ_ε is metric dependent. If the normwise pseudospectrum has been so far the most often used, componentwise pseudospectrum or any more structured pseudospectra may play a role in specific areas (see [27] for pseudospectra with homotopic perturbations and [103] for an example of structured perturbations in control theory).

A feature that has greatly helped in the success of the normwise pseudospectra is that it is easily computable. Indeed $\eta_N(z) = 1/(\left\|(A - zI)^{-1}\right\|_2 \|A\|_2)$ is the scaled reciprocal of the norm of the resolvent. This quantity can also be interpreted as the normwise backward error of the scalar z as an approximate eigenvalue of A. It is also the relative distance of $(A - zI)$ to singularity. A large amount of literature has been devoted to computation methods to evaluate $\eta_N(z)$ and the spectral portrait. We have developed

parallel software tuned for distributed memory machines which computes the pseudospectrum in a predefined region of the complex plane by evaluating the resolvent norm with a Lanczos method [48, 105]. Incidentally, the work on linear systems of the type $A^*Ax = b$ described in Section 2.4.2 arose in this context from the necessity of validating the computational process to evaluate the resolvent norm. This software is part of the PRECISE toolbox [19] which is described in next section. It has been used to provide spectral portraits to MatrixMarket, the well-known repository of matrices maintained by NIST (`http://math.nist.gov/MatrixMarket/`). See also [110] and the reference therein for a survey on the computation of pseudospectra.

On the contrary, the componentwise pseudospectra is not computable in a closed form. Indeed its definition is based on the componentwise backward error associated with the approximate eigenvalue z which can be written as

$$\eta_C(z) = \min_{\varepsilon>0} \{ |\Delta A| \leq \varepsilon \, |A| \, ; \ (A + \Delta A - zI) \text{ is singular} \} .$$

Clearly, $\eta_C(z)$ is also the componentwise distance to singularity, and, as we mentioned in Section 2.5.2, it cannot be simply related to the componentwise condition number and its computation is NP-hard. When the pseudospectra are not easily computable, one can turn to the perturbed spectra which consist of the plot of the eigenvalues of $(A+\Delta A)$ where ΔA satisfies the chosen model of perturbations and varies in size. Many such plots are given in [19] and the corresponding software is also part of the toolbox PRECISE. These plots reflect the most influential Jordan structure in the vicinity of the matrix A. The notion of the "most influential Jordan structure" has been studied by Ilahi [76] in his Ph.D. thesis.

One of the first uses of pseudospectra in numerical analysis is attributed to Varah [114]. This idea is closely related to the earlier notion of the spectrum of a family of matrices [58]. In [19], we have chosen to define the pseudospectrum from the backward error, which allows flexibility according to the choice of the underlying model of perturbations. Its extension to the generalised eigenproblem was done by Toumazou [105]. Recently, Tisseur and Higham [103] have taken the same formalism to extend the notion of pseudospectra to polynomial eigenproblems (with normwise or structured perturbations). Because he is concerned with applications in numerical approximation (rather than finite precision computations), Trefethen [108] considers alternatively an absolute norm and defines absolute ε-pseudospectra, as we will see in Chapter 3. He and his co-workers have popularised plots of perturbed spectra (with complex normwise perturbations) in many instances (see for example [107, 108]).

We end by an example illustrating the need for robust conditions of convergence in finite precision. When an eigenvalue $\tilde{\lambda}$ is computed in finite precision with a backward error of ε, the ε-pseudospectrum encloses all the approximate eigenvalues having at least the same level of quality as $\tilde{\lambda}$. The larger the ε-pseudospectrum, the more unstable the eigenproblem. Matrices that are good candidates for having a large ε-pseudospectrum are those with a high departure from normality. The influence of high nonnormality on

eigenvalue computations has been one of our major concerns and has been addressed many times in the literature [19, 15, 20, 22, 33, 66, 105]. An example from electromagnetism (provided by the Electromagnetism Project at CERFACS) where the pseudospectra can help in understanding finite precision computations can be found in [105]: a stationary iterative scheme for solving a linear system diverges whereas the eigenvalues of the iteration matrix have been proved to be strictly smaller than 1 in modulus. The pseudospectrum reveals that a perturbation ΔA of size of the order of machine precision suffices to make $(A + \Delta A)$ have eigenvalues outside the unit circle. The condition that "the spectral radius of the iteration matrix is less than 1" is necessary and sufficient in exact arithmetic, but is not sufficient in finite precision. See also [19] for a similar academic example.

2.7 The toolbox PRECISE

PRECISE is a set of tools provided to help the user set up computer experiments to explore the impact, on the quality of convergence of numerical methods, of finite precision as well as other types of prescribed perturbations of the data.
Because *stability* is at the heart of the phenomena under study – mathematical as well as numerical stabilities –, PRECISE allows one to experiment about stability by a straightforward *randomisation* of selected data, then gets the computer to produce a sample of perturbed solutions and associated residuals, or a sample of perturbed spectra.

The idea of using random perturbations on a selection of data, or parameters, to get information on the stability of dynamical processes is very natural and very old. It has been used extensively in physics and technology, but it has not gained popularity in numerical analysis, nor in numerical software. However, the idea has often been recommended by the best specialists, as illustrated by the following quotation taken from [36]:
"In larger calculational problems, the relations between input data and output data are so complicated that it is difficult to directly apply the general formulas for the propagation of error. One should then investigate the sensitivity of the output data for errors in the input data by means of an *experimental perturbational calculation:* one performs the calculations many times with perturbed input data and studies the relation between the changes (perturbations) in the input data and the changes in the output data."

This quotation serves as an excellent introduction to PRECISE, which provides an experimental environment for the engineer or the software developer to test the robustness of a numerical method or of an algorithm with respect to finite precision and data uncertainty.
It allows one to perform a complete statistical backward error analysis on a numerical method or an algorithm to solve a general nonlinear problem of the form $F(x) = y$ (matrix or polynomial equation), at regular points, and in the neighbourhood of algebraic singularities. It provides an estimate of the distance to the nearest singularity viewed by the computer, as well as the order of this singularity. It can also help to perform a sensitivity analysis by means of graphical displays of samples of perturbed solutions.

PRECISE offers the following facilities:

1. a module for statistical backward error analysis: it provides a statistical estimation for:
   - condition numbers at regular and singular points, for the algorithm/method and the problem,
   - backward errors,
   - reliability and quality indexes,
   - distances to singularity, or dangerous borders,
   - order of Hölder-singularities.
2. a module for sensitivity analysis: it provides graphical displays of:
   - perturbed spectra,
   - spectral portraits and pseudospectra for matrices,
   - sensitivity portraits and sets of pseudozeros for polynomials,
   - divergence portraits for iterations depending on a parameter.

PRECISE has been intensively used, since 1988, in several industrial environments (IBM-France, Thomson-CSF and CERFACS) to test various *laws of computation* that emerge from invariant patterns of behaviour for computations in finite precision (see [19]) It has also been used, more classically, to assess the numerical quality of computations in industrial problems such as

- the flutter phenomenon for Aerospatiale [14, 56],
- an aeroelasticity problem for ONERA (Division Hélicoptères) [5, 13],
- electromagnetic guided waves for Thomson-CSF [13],
- the reliability of an orbitography software for CNES [21, 23, 49, 62, 63],
- fluid dynamics and electromagnetism at CERFACS [48, 51, 105],
- ambiguity resolution in GPS (Jason project) for CNES [18],
- astrophysics for Observatoire Midi-Pyrénées and Politechnico Milano.

The PRECISE code was a vital part of the HPCN (High Performance Computing and Networking) European project PINEAPL (1996-98) to produce a general purpose library of parallel numerical software suitable for a wide range of computationally intensive industrial applications and to port several application codes which use this library to parallel computers. The industrial consortium led by NAG included British Aerospace, CERFACS, LCR Thomson-CSF, CPS (Napoli, Italy), the Danish Hydraulic Institute (Denmark), IBM SEMEA, the University of Manchester (UK), Math-Tech, and Piaggio (Italy).

PRECISE was translated from MATLAB into Fortran to allow large and realistic problems to be handled and was used to test each item of numerical software produced during the project. The toolbox is now available as freeware from the CERFACS Web pages `http://www.cerfacs.fr/algor/Softs/PRECISE/index.html` [82].

Before closing, it is worth stating that PRECISE has proved an extremely versatile tool to test algorithmic behaviour on the computer [76, 106], as illustrated by the following example. The quantity $\|A^{-1}\|_2$, where $\|.\|_2$ is the spectral norm of a very large matrix A can be computed using the Lanczos algorithm on A^*A, with invert. Such an algorithm requires the solution of linear systems of the kind $A^*Ax = b$. Extensive experimentation with PRECISE gave us the necessary intuition for the discovery of the theoretical formulae for the condition number and the backward error for such linear systems where only A is perturbed, as described in Section 2.4.2. This study is of importance to design a reliable code for computing the spectral portrait of a matrix [50, 82]. So PRECISE helped design itself.

As a conclusion, PRECISE is not intended to be yet another piece of software for automatic control of round-off error propagation. It is as much a matter of personal taste as of performance that should guide the user amongst the available methods and software for automatic control of accuracy. We view PRECISE as a help to investigate difficult cases, such as computations in the neighbourhood of a singularity, or computations in the presence of high nonnormality, to get better insight on the underlying mathematical instability. The better understanding of the problem provided by PRECISE allows in turn a better use of current software for error control.

Chapter 3

Backward error analysis for approximation methods in exact arithmetic

In this chapter, we would like to widen the scope of backward error analysis to applications that are not primarily concerned with finite precision computations. We review three domains that have been enriched by the concept of backward error analysis. We start by showing that the model of homotopic perturbations defines a backward error analysis which is helpful in understanding the principle of approximation of Krylov methods. Then we come back to the notion of pseudosolutions which stems directly from the definition of the backward error, and we examine its applications in exact arithmetic. Finally we see how the concept of backward error analysis can be incorporated into the study of the mathematical convergence of iterative methods. In particular we show that backward error analysis can be used as a useful framework for the analysis of the convergence of embedded iterative solvers.

3.1 Homotopic perturbations

The work realised by Chaitin-Chatelin [31, 32] and her co-workers on homotopic perturbations is an original contribution for the use of backward error analysis in approximation methods. An homotopic perturbation of a matrix A is of the form $\Delta A = tE$ where E is a prescribed deviation matrix and t is a scalar (real or complex). Homotopic unfolding is a tool of choice for an asymptotic analysis of the Jordan structure of a matrix, as shown by Lidskii's theorem [80]. This aspect has been central to the thesis of Ilahi [76].

Using the homotopic framework for the eigenvalue problem (either standard or generalised), Chaitin-Chatelin, Toumazou and Traviesas have derived the backward errors and associated pseudospectra [27]. They have also shown how the Krylov approximation can be analysed in exact arithmetic in terms of homotopic perturbations. In the Arnoldi method, at iteration k, the eigenvalues of the Hessenberg matrix H_k are supposed to approximate

at most k eigenvalues of A. The matrix A, the Hessenberg matrix H_k and the Krylov orthonormal basis $V_k = [v_1, \ldots, v_k]$ satisfy the identity $AV_k = V_k H_k + h_{k+1,k} v_{k+1} e_k^T$, which can be reformulated as $(A + h_{k+1,k} E_k) V_k = V_k H_k$ where $E_k = -v_{k+1} e_k^T$. Therefore the k^{th} step gives an exact orthonormal basis for an invariant subspace of $A + h_{k+1,k} E_k$ which is a rank one homotopic perturbation of A: all the eigenvalues of H_k have the same homotopic backward error $|h_{k+1,k}|$ with respect to A in exact arithmetic. The quantity $h_{k+1,k}$ can be viewed as the method error for the Arnoldi method. For more on the implications for finite precision behaviour, see [16, 28, 106].

3.2 Pseudospectra

3.2.1 Pseudospectra of matrices and nonnormality

Pseudospectra of matrices have been defined in Section 2.6 of Chapter 2, where their role in helping define robust conditions for convergence in finite precision has been illustrated. However, pseudospectra play a role beyond finite precision computation. Their use has been particularly emphasised in computational physics, thanks in particular to the many works of Trefethen. As an example of such a use of pseudospectra, we have had a collaboration with the CFD Project at CERFACS, in which the pseudospectra help to find the best trade-off between two discretization levels: increasing the order of the Taylor expansion or refining the mesh. This application has been one of the first to use the notion of pseudospectra of a pencil of matrices (instead of a matrix alone) and is reported in [51]. The matrices that occur in practice often depend, implicitly or explicitly, on one or several parameters which can be the order n of the matrix A itself, or a physical parameter such as the Reynolds number or the Péclet number. Whenever the spectral instability of the family of matrices under consideration is such that the index of at least one eigenvalue is unbounded and/or the condition number of the Jordan basis is unbounded under the parameter variation, we shall say – in a somewhat loose sense– that this family of matrices is **highly nonnormal**.
Many examples show the physical origin of high nonnormality: see for instance [14, 78, 79, 91, 111]. The coupling between physical phenomena is often transferred to the numerical approximation of evolution equations as a requirement for a coupling between parameters such as time steps and mesh sizes to ensure numerical stability. Without a proper restriction on the discretization sizes, the numerical method can be unstable. So we expect that nonnormality in Physics may have an impact on the numerical stability of the approximation methods, whenever this stability is dependent on the parameters of the method.

This is indeed the case. For example, it has been well known for a long time that for fully discrete evolution equations, the condition which requires that the spectrum of the spatial discretization matrix lies in the stability region for the time-stepping formula is only a *necessary* condition for stability whenever the discontinuous operator is nonnormal. Recently, Trefethen and Reddy [90] have proposed a *necessary and sufficient* condition by means of the ε-pseudospectra.

More generally, the realization that, when there is spectral discontinuity with respect to a parameter, there can be nevertheless continuity for the pseudospectra, led Trefethen [108] to propose what can be called the following principle: condition for convergence of numerical methods should be based on **pseudospectra** rather than on exact spectra. Trefethen applies his principle to the analysis of the convergence, in *exact arithmetic*, of numerical methods for various PDEs, mainly in computational fluid dynamics (CFD) [111]. Trefethen's principle is also well popularised in the SIAM News article "Are eigenvalues overvalued ?" [35].

An area of computational physics where nonnormality is encountered not infrequently is CFD. For example, it occurs in the study of parallel shear flows in fluid dynamics, whose behaviour is governed by the Orr-Sommerfeld operator (see [42]). The discretization of this nonnormal operator leads to a nonnormal matrix whose departure from nonnormality increases with the Reynolds number of the flow [89].

There are interesting consequences of nonnormality for hydrodynamic stability [111]. In order to explain the discrepancy between the computational predictions of the eigenvalue analysis and laboratory experiments in certain flows, called "subcritical transition to turbulence", one traditionally explains it by a failure of the linearisation about the laminar solution, recommending that one looks closer at the nonlinear terms, or linearises about another solution ("secondary instability"). Recently the complementary view has emerged that the cause might sometimes be found in the high nonnormality of the linearised problem. Schmid and Henningson [98] and Gustavsson [64] proved that the operators that arise in Poiseuille and Couette flows are in a sense *exponentially* far from normal. Hence the stability analysis cannot be based on the exact spectrum, but should be based on the pseudospectra [111]. This seems to be a very promising new direction in computational turbulence.

When the nonnormality is not in the mathematical equation but in the approximation method, one should also be cautious as is illustrated by the fully discrete spectral approximation of a first order mixed initial boundary value problem devised by Trefethen and Trummer [112]. With the Chebyshev collocation method in space, they show that the numerical stability of the discretization in space coupled with the discretization in time is at risk because the differentiation matrix is nonnormal: it is a low rank perturbation of a fully defective Jordan block. Numerical stability in exact arithmetic is guaranteed only with an appropriate coupling between time and space steps.

Since the question of how the physical instability reflects in the equations of the model is central, the extension of the notion of pseudospectra of matrices to operators arises naturally. We address it now.

3.2.2 Pseudospectra of operators

Let T be a closed linear operator on a Banach space. For $\varepsilon > 0$, the natural definition of the ε-pseudospectrum of the operator T is

$$\sigma_\varepsilon = \left\{ z \in \rho(T);\ \left\| (zI - T)^{-1} \right\| \geq 1/\varepsilon \right\} \ \cup\ \sigma(T), \tag{3.1}$$

where $\rho(T)$ is the resolvent set of T and $\sigma(T)$ is the spectrum of T (see [30] for detailed definitions). This definition uses an absolute norm because it is set in the context of exact arithmetic. Based on this definition, Harrabi [65] establishes the continuity of the ε-pseudospectrum (for $\varepsilon > 0$). If T and $(T_n)_{n\in\mathbb{N}}$ are bounded operators in a Banach space such that $\|T - T_n\| \to 0$ (uniform convergence), then the spectrum of the operator family $(T_n)_{n\in\mathbb{N}}$ is equal to the spectrum of T. Harrabi has also shown weaker results that hold for collectively compact convergence [30].

The definition of the pseudospectrum given above is the conventional definition with non-strict inequality in (3.1) [109]. It makes the pseudospectra a closed set and represents the natural extension of the definition of pseudospectra for matrices (finite dimension) to operators (infinite dimension). An alternate definition is derived by taking a strict inequality instead of $\geq$ in 3.1 [25] and the corresponding set $\omega_\varepsilon(T)$ is called *spectral value set* in [37, 53]. Under a certain hypothesis, $\sigma_\varepsilon(T)$ is the closure of $\omega_\varepsilon(T)$. This hypothesis assumes that the norm of the resolvent of T is not constant on an open set in $\rho(T)$. This hypothesis is always satisfied for T in a Hilbert space or in a Banach space of finite dimension [8, 69]. However, for an operator in an infinite dimensional Banach space, the validity of the hypothesis is still an open question in functional analysis.

The relationship between the pseudospectrum of an operator and that of its discretizations is an interesting topic that keeps developing. See in particular [109] and the references therein, and [67, 68].

3.3 Convergence of iterative methods in Linear Algebra

Backward error analysis is now more and more recognised as a tool of choice to analyse the convergence of numerical methods in *exact arithmetic*. We give two examples of such a use.

3.3.1 Convergence of the Power method

As a first example, the theoretical convergence of the Power method for the computation of the dominant eigenvalue of a matrix has recently been revisited by Ilahi, Bouras and Chaitin-Chatelin [77]. Until now, most of the convergence results for the power method available in the literature assumed that the dominant eigenvalue is unique and simple or semi-simple. The work of Ilahi et al. considers the case where this dominant eigenvalue is defective and the single assumption remains that there are no two distinct eigenvalues of

largest modulus. In addition, the convergence of the method is analysed in terms of the backward error associated with the eigenpair $(\lambda^{(k)}, u^{(k)})$ at the k^{th} iteration, that is

$$\eta(\lambda^{(k)}, u^{(k)}) = \frac{\left\| Au^{(k)} - \lambda^{(k)} u^{(k)} \right\|_2}{\|A\|_2},$$

and also through the backward error associated with the eigenvalue alone

$$\eta(\lambda^{(k)}) = \frac{1}{\|A\|_2 \left\| (A - \lambda^{(k)} I)^{-1} \right\|_2}.$$

It is proved in particular that, if the dominant eigenvalue is defective with an index $l \geq 2$, the backward error $\eta(\lambda^{(k)}, u^{(k)})$ converges as $1/k^2$ whereas $\eta(\lambda^{(k)})$ converges as $(1/k^l) \leq (1/k^2)$. However, the forward errors on the eigenvalue or on the eigenvector both converge as $1/k$. Because it illustrates the subtle convergence properties expressed as forward / backward errors, this example is very illustrative of the insight that can be gained by making full use of backward error analysis.

Note that the backward errors given above are defined in a relative sense (they are scaled with the norm of A). Strictly speaking, this is useless for the analysis in exact arithmetic but this was done on purpose in view of future applications in finite precision.

3.3.2 Embedded iterative solvers

Finally, we would like to discuss the convergence of embedded iterative solvers, an area which, in our opinion, will attract growing interest in the near future.

Iterative processes are widely used in scientific computing and in particular in linear algebra. We are in interested in the specific case when two iterative solvers are embedded. The crucial question arises then: *what is the best strategy for stopping the inner iterations for ensuring the convergence of the outer iterations while minimising the global computational cost ?* In the context of matrix computation on large-scale problems, saving inner iterations usually means making less matrix-vector products, and that represents a potential reduction of the computational cost. If there are various studies for the behaviour of inexact Newton-like methods, the interest for the study of embedded linear solvers is more recent, and stems in part from the growing success of iterative Krylov methods. This question has been partially addressed by numerical experts since the eighties, in the context of Newton-like methods [38, 43, 54, 59, 60, 100]. It is generally concluded, as one could expect, that the accuracy of the inner iteration needs to be *increased* when the outer process comes closer to the solution. The proposed strategies for monitoring the inner iterations have been so far very problem- and method-dependent. More recently in the late nineties, it has been emphasised that the behaviour of embedded solvers involving a Krylov outer process was very different than that of inexact Newton-like methods [60, 61]. A strategy for monitoring the accuracy of inner iterations is proposed in the framework of symmetric eigenvalue problems with homogeneous linear constraints in [61]. Why do

Krylov methods behave so differently ?

Our recent results on the control of embedded iterations, in the case where the outer solver is a Krylov method, have been presented in the report [10] entitled "A relaxation strategy for inexact matrix-vector products for Krylov methods". This report, submitted to the Journal of Numerical Linear Algebra with Applications, is given in the Annex. Less studied in the literature because they are more recent, these methods seems to have very strong robustness properties: as outer processes, they accept larger and larger inaccuracies from the inner process, provided that these inaccuracies are properly monitored. Our analysis is set in the framework of backward error analysis, which is the most powerful tool for handling inexact data and which naturally applies to embedded iterative processes. In order to better understand the behaviour of outer Krylov processes with inner processes whose accuracy can vary and is controllable (such as another iterative method), we first study the linear Krylov method GMRES with inexact matrix-vector products. The matrix-vector product, taken as the inner process, is used each time a vector is added to the Krylov basis. By perturbing each matrix-vector product, we are able to control its accuracy and in this way check the impact of various strategies for monitoring the inner accuracy on the outer convergence. We have shown in this way that an interesting and efficient strategy consists in linking the accuracy of the matrix-vector product to the reciprocal of the outer residual. Therefore, the matrix-vector products become less and less accurate as the outer convergence proceeds. This approach is called the *relaxation strategy*. In practice, we observe that only the first Krylov vectors are computed with the targeted outer accuracy, while the subsequent vectors are soon subjected to large perturbations: nevertheless, the global outer convergence is not dramatically changed.

However, it is important to observe that the relaxation strategy that we have defined suffers from an important drawback: it is not scaling independent, while the convergence of GMRES is. In our first tests reported in [10], the matrices are very well suited to our strategy. But it is easy to design test cases for which our relaxation strategy will fail to meet its objective. It is our intention to work on this problem in the near future.

The robustness of Krylov methods to perturbations of the Krylov basis is a remarkable fact which has many important implications. Indeed, inexact matrix-vector products arise naturally in multipole methods which have become popular in the solution of electromagnetism problems. Moreover, most of the embedded iterative solvers with an outer Krylov scheme can benefit from the results observed in [10]. We have started to explore two applications that we now briefly review.

Inner-outer iterations in eigenvalue computations

The Arnoldi method and its variants are one of the most popular techniques for solving large-scale nonsymmetric eigenproblems. This Krylov method allows us to approximate the periphery of the spectrum of a matrix A. The matrix A is only used through matrix-vector products for the construction of the Krylov basis onto which it is projected. When

one seeks internal eigenvalues (for example those close to some prescribed value σ), one can apply a shift and invert technique which results in computing the dominant eigenvalues of the matrix $(A - \sigma I)^{-1}$. The construction of a Krylov basis of size m requires the solution of m linear systems of the type $(A-\sigma I)z_{k+1} = z_k$. When the matrix A is so large that even a sparse direct solver cannot be applied, these systems have to be solved with an iterative method which becomes the inner process.

In the report [11], we first study the Arnoldi method with inexact matrix-vector products and define a similar relaxation strategy as for GMRES in [10]. A comparison with the inexact power method underlines the fundamental differences between Krylov-type methods and Newton-type methods. We then propose two relaxation strategies for the control of the inner accuracy of the shifted and inverted Arnoldi method. These two strategies reflect the duality inherent in inversion methods. Indeed, one can consider either the computation of the smallest eigenvalue λ of A (final goal) or the computation of the largest eigenvalue $\mu = 1/\lambda$ of $B = A^{-1}$ (means). Consequently, two backward errors can be considered: the one associated with the computed eigenvalue $\hat{\lambda}$ of A and the one associated with the computed eigenvalue $\hat{\mu}$ of B. Therefore, there are two ways of assessing the quality of the convergence, and it is interesting to study which one is the best suited for a Krylov solver in finite precision. We cannot yet propose a definitive choice between both viewpoints, and this important question remains open for future work [9].

Domain decomposition

Domain decomposition is a frequently used technique for the solution of partial differential equations on parallel computers. A finite-element discretization followed by an appropriate renumbering of the unknowns leads to a structured linear system, whose diagonal blocks A_{ii} (local matrices) correspond to the discretization of the operator restricted to the interior nodes for each subdomain. A subdomain is connected (via the coupling blocks $A_{\Gamma i}$ et $A_{i\Gamma}$) to the last diagonal block A_Γ associated with the interface nodes. The elimination of the unknowns associated with the interior nodes leads to the condensed system $Sx = f$, where $S = A_\Gamma - \Sigma_i(A_{i\Gamma}A_{ii}^{-1}A_{\Gamma i})$ is the Schur complement. It is not reasonable to assemble the Schur complement S explicitly because the matrices A_{ii}^{-1} operate on much larger spaces (depending on the discretization size). Additionally, even if the Schur complement operates on a smaller space (that of the interfaces), it may still be quite large and is much denser than the matrices it involves. Therefore, the condensed system is often solved using iterative methods such as the conjugate gradient method (when the Schur complement is positive definite): at each outer iteration, the matrix-vector product Sx requires the solution of the local linear systems with coefficient matrices A_{ii}^{-1}. The latter can be solved in turn using an iterative method (inner process). Here again, embedded iterations come in.

In [12], we describe the direct application of the relaxation strategy described in [10]. The control on the accuracy of the matrix-vector product Sx must be transferred in fact to the control on the accuracy for the solution of each local system: this is done using backward error analysis arguments. The relaxation scheme is applied to a heterogeneous

anisotropic problem representative of those arising in semi-conductor modelling. The numerical experiments show that a significant number of matrix-vector products can be saved using the relaxation strategy.

Chapter 4

Conclusion

Backward error analysis was primarily designed as a tool to understand the behaviour of numerical algorithms in finite precision; however its scope is much wider. As explained in the introductory part, finite precision computations fit into the wider scenery of methods of approximations, where the backward error analysis appears to be the tool of choice to deal with any kind of uncertain computations. At this conceptual level, one can fully embrace the power of backward error analysis.

The apparent simplicity of its principle hides great complexity, visible as soon as one gets into the subtleties of finite precision computations. A tour of the most recent techniques used to assess the reliability of numerical software has been proposed in the second chapter. The current trend is to develop more and more refined models which are better and better suited to describe the perturbations generated by finite precision arithmetic.

As numerical simulations tackle physical problems of increasing difficulty and the numerical models increase in complexity, the coupling between numerical approximation and finite precision plays a crucial role in assessing the domain of computability of the solutions. The correlated question of the robustness of the convergence conditions to perturbations in the data has emerged in recent years, through the notion of pseudo-solution which is now being used for numerical approximation as well as for finite precision computations. It is likely that many examples of applications will continue to emerge.

Finally some numerical methods seem to be particularly robust to perturbations, provided these perturbations are applied in a proper way. This seems to be the case of the Krylov methods, which, unlike Newton methods, can bear perturbations of increasing sizes through the course of the convergence. As a consequence, embedded iterative solvers can be made less expensive when a Krylov outer scheme is used. Understanding the robustness of Krylov methods with respect to inexact computation is still a largely open question, and to our taste, one of the most exciting in linear algebra. There is no doubt that backward error analysis will prove again to be a key factor in answering this question.

To end on a more personal touch, I would like to emphasise that CERFACS, its Parallel Algorithms Project, and the Qualitative Computing Group within it, play an active role in promoting backward error analysis within the academic and the industrial worlds. It is not only the core of our research, but also a major topic in our educational activities and a key component of our industrial collaborations, and it should remain so in the future.

Bibliography

[1] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK User's Guide*. SIAM, Philadelphia, second edition, 1995.

[2] M. Arioli, I. S. Duff, and D. Ruiz. Stopping criteria for iterative solvers. *SIAM J. Matrix Anal. Appl.*, 13:138–144, January 1992.

[3] J. P. Aubin. *Approximation of elliptic boundary value problems*. Wiley, New York, 1972.

[4] F. L. Bauer. Genauigkeitsfragen bei der Lösung linearer Gleichungssysteme. *ZAMM*, 46:409–421, 1966.

[5] 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.

[6] Å. Björck and C. C. Paige. Solution of augmented linear systems using orthogonal factorizations. *BIT*, 34:1–24, 1994.

[7] Å. Björck and V. Pereyra. Solution of Vandermonde systems of equations. *Math. Comp.*, 24:893–903, 1970.

[8] A. Böttcher. Pseudospectra and singular values of large convolution operators. *J. Integral Equations Applics.*, 6:267–301, 1994.

[9] A. Bouras. *Contrôle de convergence de solveurs emboîtés pour le calcul de valeurs propres avec inversion*. Ph.D. dissertation, Université Toulouse I, 2000. in preparation.

[10] A. Bouras and V. Frayssé. A relaxation strategy for inexact matrix-vector products for Krylov methods. Technical Report TR/PA/00/15, CERFACS, Toulouse, France, 2000.

[11] A. Bouras and V. Frayssé. A relaxation strategy for the Arnoldi method in eigenproblems. Technical Report TR/PA/00/16, CERFACS, Toulouse, France, 2000.

[12] A. Bouras, V. Frayssé, and L. Giraud. A relaxation strategy for inner-outer linear solvers in domain decomposition methods. Technical Report TR/PA/00/17, CERFACS, Toulouse, France, 2000.

[13] T. Braconnier. *Sur le calcul de valeurs propres en précision finie*. Ph.D. dissertation, Université H. Poincaré, Nancy, May 1994. TH/PA/94/24, CERFACS.

[14] 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.

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

[16] F. Chaitin-Chatelin. Comprendre les méthodes de Krylov en précision finie : le programme du Groupe Qualitative Computing au CERFACS. Technical Report TR/PA/00/11, CERFACS, Toulouse, France, 2000.

[17] F. Chaitin-Chatelin, S. Dallakyan, and V. Frayssé. On the GPS carrier phase ambiguity resolution. the LAMBDA method: an analysis of speed, efficiency and numerical robustness. Contract Report IR/PA/99/23, CERFACS, Toulouse, France, 1999.

[18] F. Chaitin-Chatelin, S. Dallakyan, and V. Frayssé. An overview of carrier phase differential GPS. Contract Report IR/PA/99/50, CERFACS, Toulouse, France, 1999.

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

[20] 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.

[21] F. Chaitin-Chatelin, V. Frayssé, and S. Gratton. Fiabilité des calculs en précision finie : principes et mise en œuvre. Contract Report FR/PA/97/55, CERFACS, Toulouse, France, 1997.

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

[23] F. Chaitin-Chatelin and S. Gratton. Etude de la non convergence de gauss-newton. Contract Report IR/PA/97/14, CERFACS, Toulouse, France, 1997.

[24] F. Chaitin-Chatelin and S. Gratton. On the condition numbers associated with the polar factorization of a matrix. *Journal of Numerical Linear Algebra with Applications*, 2000. to appear.

[25] F. Chaitin-Chatelin and A. Harrabi. About definitions of pseudospectra of closed operators in Banach spaces. Tech. Rep. TR/PA/98/08, CERFACS, Toulouse, France, 1998.

[26] F. Chaitin-Chatelin, A. Harrabi, and A. Ilahi. About Hölder condition numbers and the stratification diagram for defective eigenvalues. Tech. Rep. TR/PA/99/19, CERFACS, Toulouse, France, 1999. to appear in IMACS Journal on Mathematics of Computation.

[27] F. Chaitin-Chatelin, V. Toumazou, and E. Traviesas. Accuracy assessment for eigencomputations : variety of backward errors and pseudospectra. *Linear Algebra and Appl.*, 309:73–83, 2000.

[28] F. Chaitin-Chatelin, E. Traviesas, and L. Plantié. Understanding krylov methods in finite precision, 2000. To appear in Proceedings of the Second Conference on Numerical Analysis and Applications, Rousse, Bulgaria.

[29] X. W. Chang, C. C. Paige, and G.W. Stewart. New perturbation analyses for the Cholesky factorization. *IMA J. Numer. Anal.*, 14:1–28, 1995.

[30] F. Chatelin. *Spectral approximation of linear operators.* Academic Press, New York, 1983.

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

[32] F. Chatelin. *Eigenvalues of matrices.* Wiley, Chichester, 1993. Enlarged Translation of the French Edition with Masson.

[33] F. Chatelin. The influence of nonnormality on matrix computations. In R. J. Plemmons and C. D. Meyer, editors, *Linear Algebra, Markov Chains and Queueing Models*, pages 13–19. Springer, New York, 1993.

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

[35] B. A. Cipra. Are eigenvalues overvalued ? *SIAM News*, 28:8, January 1995.

[36] G. Dahlquist and Å. Björck. *Numerical methods.* Prentice-Hall, Englewood Cliffs, NJ, 1980.

[37] E. B. Davies. Pseudospectra of differential operators. Technical Report Preprint KCL-MATH-97-67, King's College, London, 1997.

[38] R. S. Dembo, S.C. Eisenstat, and T. Steihaug. Inexact Newton methods. *SIAM J. Numer. Anal.*, 19(2):400–408, 1982.

[39] J. W. Demmel. The componentwise distance to the nearest singular matrix. *SIAM J. Matrix Anal. Appl.*, 13:10–19, January 1992.

[40] J. W. Demmel and B. Kågström. The generalized Schur decomposition of an arbitrary pencil $A - \lambda B$: robust software with error bounds and applications. Part I: Theory and Algorithms. *ACM Trans. Math. Software*, 19:160–174, 1993.

[41] J. W. Demmel and B. Kågström. The generalized Schur decomposition of an arbitrary pencil $A - \lambda B$: robust software with error bounds and applications. Part II: Software and Applications. *ACM Trans. Math. Software*, 19:175–201, 1993.

[42] P. G. Drazin and W. H. Reid. *Hydrodynamic Stability.* Cambridge University Press, Cambridge, 1981.

[43] S. C. Eisenstat and H. F. Walker. Globally convergent inexact Newton methods. *SIAM J. Optimization*, 4(2):393–422, 1994.

[44] L. Fox, H. D. Huskey, and J. H. Wilkinson. Notes on the solution of algebraic linear simultaneous equations. *Quart. J. Mech. and Applied Math.*, 1:149–173, 1948.

[45] 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, Toulouse, France, 1996.

[46] V. Frayssé, L. Giraud, and S. Gratton. A set of GMRES routines for real and complex arithmetics. Technical Report TR/PA/97/49, CERFACS, Toulouse, France, 1997.

[47] V. Frayssé, L. Giraud, and S. Gratton. A set of Flexible-GMRES routines for real and complex arithmetics. Technical Report TR/PA/98/20, CERFACS, Toulouse, France, 1998.

[48] 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.

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

[50] V. Frayssé, S. Gratton, and V. Toumazou. Structured backward error and condition number for linear systems of the type $A^*Ax = b$. *BIT*, 40:74–83, 2000.

[51] V. Frayssé, M. Gueury, F. Nicoud, and V. Toumazou. Spectral portraits for matrix pencils: a physical application. Technical Report TR/PA/96/19, CERFACS, Toulouse, France, 1996.

[52] V. Frayssé and V. Toumazou. A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$. *Journal of Numerical Linear Algebra with Applications*, 5(1):1–10, 1998.

[53] E. Gallestey, D. Hinrichsen, and A. Pritchard. On the spectral value sets in infinite dimension. In *Proc. 4th International Symposium on Methods and Models in Automation and Robotics*, pages 109–11436, 1997.

[54] E. Giladi, G. H. Golub, and J. B. Keller. Inner and outer iterations for the Chebyshev algorithm. *SIAM J. Numer. Anal.*, 35(1):300–319, 1998.

[55] W. Givens. Numerical computation of the characteristic values of a real symmetric matrix. Technical Report ORNL-1574, Oak Ridge National Laboratory, 1954.

[56] S. Godet-Thobie. *Eigenvalues of large highly nonnormal matrices (Valeurs propres de matrices hautement non normales en grande dimension).* Ph.D. dissertation, Université Paris IX Dauphine, December 1992.

[57] S. K. Godunov. Spectral portraits of matrices and criteria of spectrum dichotomy. In J. Herzberger and L. Atanassova, editors, *Computer arithmetic and enclosure methods.* North-Holland and IMACS, October 1992.

[58] S. K. Godunov and V. S. Ryabenki. *Theory of Difference Schemes: an Introduction.* North-Holland, Amsterdam, 1964. Translation by E. Godfredsen.

[59] G. H. Golub and M. L. Overton. The convergence of inexact Chebyshev and Richardson iterative methods for solving linear systems. *Numer. Math.*, 53:571–593, 1988.

[60] G. H. Golub and Q. Ye. Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM J. Sci. Comput.*, 21(4):1305–1320, 1999.

[61] G. H. Golub, Z. Zhang, and H. Zha. Large sparse symmetric eigenvalue problems with homogeneous linear constraints: the Lanczos process with inner-outer iterations. *Linear Algebra and Appl.*, 309:289–306, 2000.

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

[63] S. Gratton. *Outils théoriques d'analyse du calcul à précision finie.* Ph.D. dissertation, Institut National Polytechnique de Toulouse, June 1998. TH/PA/98/30, CERFACS.

[64] L. H. Gustavsson. Energy growth of three-dimensional disturbances in plane Poiseuille flow. *J. Fluid Mech.*, 224:241–260, 1991.

[65] A. Harrabi. Pseudospectre d'une suite d'operateurs bornés. Tech. Rep. TR/PA/97/48, CERFACS, Toulouse, France, 1997.

[66] A. Harrabi. Défaut de normalité. Tech. Rep. TR/PA/98/36, CERFACS, Toulouse, France, 1998.

[67] A. Harrabi. On the approximation of pseudospectra of nonnormal operators by discretization, part I: the first derivative operator. Tech. Rep. TR/PA/98/37, CERFACS, Toulouse, France, 1998.

[68] A. Harrabi. Pseudospectre d'une suite d'operateurs bornés. *Mathematical Modelling and Numerical Analysis*, 32(6):671–680, 1998.

[69] A. Harrabi. *Pseudospectres d'opérateurs intégraux et différentiels. Application à la Physique mathématique.* Ph.D. dissertation, Université Toulouse I, May 1998. TH/PA/98/29, CERFACS.

[70] D. J. Higham and N. J. Higham. Structured backward error and condition of generalized eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 20(2):493–512, 1998.

[71] N. J. Higham. Error analysis of the Björck-Pereyra algorithm for solving Vandermonde systems. *Numer. Math.*, 50:613–632, 1987.

[72] N. J. Higham. *Accuracy and Stability of Numerical Algorithms.* SIAM, Philadelphia, 1996.

[73] Nicholas J. Higham. Optimization by direct search in matrix computations. *SIAM J. Matrix Anal. Appl.*, 14(2):317–333, April 1993.

[74] D. Hinrichsen and B. Kelb. Spectral value sets : a graphical tool for robustness analysis. *Systems Control Lett.*, 21:127–136, 1993.

[75] Harold Hotelling. Some new methods in matrix calculation. *Ann. Math. Statist.*, 14(1):1–34, 1943.

[76] A. Ilahi. *Validation du calcul sur ordinateur: application de la théorie des singularités algébriques.* Ph.D. dissertation, Université Toulouse I, June 1998. TH/PA/98/31, CERFACS.

[77] A. Ilahi, A. Bouras, and F. Chaitin-Chatelin. Nouvelles analyses de la méthode de la puissance. Partie 1 : Etude asymptotique générale. Technical Report TR/PA/99/35, CERFACS, Toulouse, France, 1999.

[78] W. Kerner. Computing the complex eigenvalue spectrum for resistive magnetohydrodynamics. In J. Cullum and R. A. Willoughby, editors, *Large scale eigenvalue problems*, pages 240–264. North-Holland, Amsterdam, 1986.

[79] W. Kerner. Large scale complex eigenvalues problems. *J. Comp. Phys.*, 85:1–85, 1989.

[80] V. B. Lidskii. Perturbation theory of non-conjugate operators. *U.S.S.R. Comput. Maths. Math. Phys.*, 1:73–85, 1965.

[81] C. Van Loan. The sensitivity of the matrix exponential. *SIAM J. Numer. Anal.*, 14:971–981, 1977.

[82] R. A. McCoy and V. Toumazou. PRECISE User's Guide - Version 1.0. Tech. Rep. TR/PA/97/38, CERFACS, Toulouse, France, 1997.

[83] J. Moro, J. V. Burke, and M. L. Overton. On the Lidskii-Vishik-Lyusternik perturbation theory for eigenvalues of matrices with arbitrary Jordan structure. *SIAM J. Matrix Anal. Appl.*, 1996. To appear.

[84] R. G. Mosier. Root neighborhoods of a polynomial. *Math. Comp.*, 47:265–273, 1986.

[85] H. B. Nielsen. AAFAC: a package of Fortran 77 subprograms for solving $A^T Ax = c$. Report NI 90-01, Institute for Numerical Analysis, Technical University of Denmark, 1990.

[86] H. B. Nielsen. Computing a minimizer of a piecewise quadratic. Technical Report IMM-REP-1998-14, Department of Mathematical Modelling, Technical University of Denmark, 1998.

[87] W. Oettli and W. Prager. Compatibility of approximate solution of linear equations with given error bounds for coefficients and right-hand sides. *Numer. Math.*, 6:405–409, 1964.

[88] S. Poljak and J. Rohn. Checking robust nonsingularity is np-hard. *Math. Control Signals Systems*, 6:1–9, 1993.

[89] S. C. Reddy. *Pseudospectra of operators and discretization matrices and an application to stability of the method of lines.* Ph.D. dissertation, Dept. of Mathematics, Massachusetts Institute of Technology, August 1991.

[90] S. C. Reddy and L. N. Trefethen. Lax-stability of fully discrete spectral methods via stability regions and pseudo-eigenvalues. *Comp. Meth. Appl. Mech. Eng.*, 80:147–164, 1990.

[91] S. C. Reddy and L. N. Trefethen. Pseudospectra of the convection-diffusion operator. *SIAM J. Appl. Math.*, 54:1634–1649, 1994.

[92] J. R. Rice. A theory of condition. *SIAM J. Numer. Anal.*, 3:287–310, 1966.

[93] J.-L. Rigal and J. Gaches. On the compatibility of a given solution with the data of a linear system. *J. Assoc. Comput. Mach.*, 14:543–548, July 1967.

[94] J. Rohn. Systems of linear interval equations. *Linear Algebra and Appl.*, 126:39–78, 1989.

[95] S. M. Rump. Almost sharp bounds for the componentwise distance to the nearest singular matrix. *Linear and Multilinear Algebra*, 42:93–107, 1997.

[96] S. M. Rump. Bounds for the componentwise distance to the nearest singular matrix. *SIAM J. Matrix Anal. Appl.*, 18(1):83–103, 1997.

[97] S. M. Rump. Structured perturbations and symmetric matrices. *Linear Algebra and Appl.*, 278:121–132, 1998.

[98] P. J. Schmid and D. S. Henningson. A new mechanism for rapid transition involving a pair of oblique waves. *Phys. Fluids A*, 4:1986–1989, 1992.

[99] R. D. Skeel. Scaling for numerical stability in Gaussian Elimination. *J. Assoc. Comput. Mach.*, 26:494–526, July 1979.

[100] P. Smit and M.-H.-C. Paardekooper. The effects of inexact solvers in algorithms for symmetric eigenvalue problems. *Linear Algebra and Appl.*, 287:337–357, 1999.

[101] G. W. Stewart and J. Sun. *Matrix Perturbation Theory.* Academic Press, New York, 1991.

[102] F. Tisseur. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra and Appl.*, 309(1–3):339–361, 2000.

[103] F. Tisseur and N. J. Higham. Structured pseudospectra for polynomial eigenvalue problems, with applications. Technical Report Numerical Analysis Report No. 359, University of Manchester, April 200.

[104] K.-C. Toh and L. N. Trefethen. Pseudozeros of polynomials and pseudospectra of companion matrices. *Numer. Math.*, 68:403–425, 1994.

[105] V. Toumazou. *Portraits spectraux de matrices: un outil d'analyse de la stabilité.* Ph.D. dissertation, Université H. Poincaré, Nancy, 1996. TH/PA/96/46, CERFACS.

[106] E. Traviesas. *Sur le déploiement du champ spectral d'une matrice.* Ph.D. dissertation, Université Toulouse I, TH/PA/00/30, CERFACS, Mai 2000.

[107] L. N. Trefethen. Approximation theory and numerical linear algebra. In J. C. Mason and editors M. G. Cox, editors, *Algorithms for Approximation II*, pages 336–359, London, 1990. Chapman and Hall.

[108] L. N. Trefethen. Pseudospectra of matrices. In D. F. Griffiths and G. A. Watson, editors, *Numerical Analysis 1991.* Longman, Harlow, 1992.

[109] L. N. Trefethen. Pseudospectra of linear operators. *SIAM Review*, 39(3):383–406, 1997.

[110] L. N. Trefethen. Computation of pseudospectra. *Acta Numerica*, pages 247–295, 1999.

[111] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamics stability without eigenvalues. *Science*, 261:578–584, July 1993.

[112] L. N. Trefethen and M. Trummer. An instability phenomenon in spectral methods. *SIAM J. Numer. Anal.*, 24:1008–1023, 1987.

[113] A. van der Sluis. Domain of uncertainty for perturbed operator equations. *Computing*, 5:312–323, 1970.

[114] J. M. Varah. On the separation of two matrices. *SIAM J. Numer. Anal.*, 16:216–222, 1979.

[115] J. H. Wilkinson. Error analysis of direct methods of matrix inversion. *J. Assoc. Comput. Mach.*, 8:281–330, July 1961.

[116] J. H. Wilkinson. *Rounding errors in algebraic processes*, volume 32. Her Majesty's stationery office, London, 1963.

[117] J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Oxford University Press, Oxford, 1965.

List of publications

Book

F. Chaitin-Chatelin and V. Frayssé. *Lectures on Finite Precision Computations.* SIAM, Philadelphia, 1996 ($2^{\text{ème}}$ tirage en cours).

Articles from international journals

V. Frayssé, S. Gratton, and V. Toumazou. Structured backward error and condition number for linear systems of the type $A^*Ax = b$. *BIT*, 40:74–83, 2000.

V. Frayssé and V. Toumazou. A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$. *J. Numer. Linear Algebra Appl.*, 5(1):1–10, 1998.

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

Proceedings of international conferences

V. Frayssé, L. Giraud, and H. Kharraz-Aroussi. On the influence of the orthogonalization scheme on the parallel performance of GMRES. In D. Pritchard and J. Reeve, editors, *EUROPAR'98 Parallel Processing*, volume 1470, pages 751–762. Springer, 1998.
Presented at EUROPAR'98, Southampton, U.K., September 2-4 1998.

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.

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.
Presented at IMACS'95 Conference on Iterative Methods, Blaegovgrad, Bulgaria, June 17-20 1997.

F. Chatelin and V. Frayssé. Elements of a condition theory for the computational analysis of algorithms. In R. Beauwens and P. De Groen, editors, *Iterative methods in linear algebra*, pages 15–25, Elsevier Science Publishers B.V., North-Holland, 1992.

F. Chatelin and V. Frayssé. Arithmetic reliability of algorithms. In M. Durand and F. El Dabaghi, editors, *High Performance Computing II*, pages 441–450, North-Holland, Amsterdam, 1991.
Presented at the 2nd Symposium on High Performance Computing. Montpellier, France, October 7-9 1991.

F. Chatelin and V. Frayssé. Analysis of arithmetic algorithms : a statistical study. In P. Kornerup and D. Matula, editors, *10th IEEE Symposium on Computer Arithmetic*, pages 10–16, June 1991.

Technical reports

V. Frayssé and L. Giraud A Set of Conjugate Gradient Routines for Real and Complex Arithmetics. Tech. Rep. TR/PA/00/47, CERFACS, 2000.

A. Bouras, V. Frayssé and L. Giraud. A relaxation strategy for inner-outer linear solvers in domain decomposition methods. Tech. Rep. TR/PA/00/15, CERFACS, 2000.

A. Bouras and V. Frayssé. A relaxation strategy for inexact matrix-vector products for Krylov methods. Tech. Rep. TR/PA/00/15, CERFACS, 2000.

A. Bouras and V. Frayssé. A relaxation strategy for the Arnoldi method in eigenproblems. Tech. Rep. TR/PA/00/16, CERFACS, 2000.

S. Gratton, V. Frayssé, and V. Toumazou. Structured backward error and condition number for linear systems of the type $A^*Ax = b$. Tech. Rep. TR/PA/99/05, CERFACS, 1999.

S. Gratton, V. Frayssé, and V. Toumazou. On a structured backward error analysis for linear systems of the type $A^*Ax = b$. Technical Report TR/PA/98/01, CERFACS, 1998.

V. Frayssé, L. Giraud, and S. Gratton. A set of Flexible-GMRES routines for real and complex arithmetics. Technical Report TR/PA/98/20, CERFACS, 1998.

V. Frayssé and L. Giraud. Comparative study of QMR versus block-QMR for J-symmetric matrices in electromagnetism applications, Tech. Rep. TR/PA/98/11, CERFACS, 1998.

V. Frayssé, L. Giraud, and S. Gratton. A set of GMRES routines for real and complex arithmetics. Technical Report TR/PA/97/49, CERFACS, 1997.

T. Braconnier, V. Frayssé, and J.-C. Rioual. ARNCHEB users' guide : Solution of large non symmetric or non hermitian eigenvalue problems by the Arnoldi-Tchebycheff method. Tech. Rep. TR/PA/97/50, CERFACS, 1997.

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.

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.

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.

F. Chatelin, V. Frayssé, and T. Braconnier. The influence of large nonnormality on the quality of convergence of iterative methods in linear algebra. TR/PA/94/07, CERFACS, 1994.

F. Chatelin, V. Frayssé, and T. Braconnier. Qualitative Computing : elements of a theory for finite precision computation. Tech. Rep. TR/PA/93/12, CERFACS, 1993. Lecture Notes for the Workshop on Reliability of Computations, March 30-April 1, Toulouse, France.

F. Chatelin and V. Frayssé. A statistical study of the stability of linear systems. Tech. Rep. TR/PA/91/43, CERFACS, 1991.

Contract reports

F. Chaitin-Chatelin, S. Dallakyan and V. Frayssé. GPS Carrier Phase Ambiguity Resolution with the LAMBDA method: 1. A stability analysis 2. An exponential speed-up. Contract Report CR/PA/00/52, CERFACS, 2000.

A. Bouras and F. Chaitin-Chatelin and V. Frayssé. Solveurs itératifs imbriqués. Contract Report FR/PA/00/19, CERFACS, 2000.

F. Chaitin-Chatelin, S. Dallakyan and V. Frayssé An overview of Carrier Phase Differential GPS. Contract Report IR/PA/99/50, CERFACS, 1999.

F. Chaitin-Chatelin, S. Dallakyan and V. Frayssé. On the GPS Carrier Phase Ambiguity Resolution. The LAMBDA method: an analysis of speed, efficiency and numerical robustness. Contract Report IR/PA/99/23, CERFACS, 1999.

F. Chaitin-Chatelin, V. Frayssé, and S. Gratton. Traitement d'ambiguïtés entières. Contract Report IR/PA/98/54, CERFACS, 1998.

V. Frayssé and L. Giraud. An implementation of block QMR for J-symmetric matrices. Tech. Rep. FR/PA/97/57, CERFACS, 1997.

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.

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

Thesis manuscript

V. Frayssé. *Reliability of computer solutions (Sur la fiabilité des calculs sur ordinateurs).* Thèse de Doctorat, Institut National Polytechnique de Toulouse, No d'ordre 589, Juillet 1992. (TH/PA/92/11, CERFACS).

Talks given at international conferences

A. Bouras and V. Frayssé and L. Giraud. A relaxation strategy for inner-outer linear solvers in domain decomposition methods. 13th International Conference on Domain Decomposition Methods, Lyon, France, October 9-12, 2000.

V. Frayssé. Pseudospectra at CERFACS. ICIAM'99, Edimbourg, U.K., July 5-9, 1999.

F. Chaitin-Chatelin, V. Frayssé and S. Gratton. Use of Kronecker products in Backward Error Analysis. SIAM Annual Meeting, Atlanta, USA, May 8–15 1999.

S. Gratton, J. C. Bergès and V. Frayssé. Least Squares Problems in Aerospace Industry. Conference on Least squares methods: theory, algorithms and applications. Linköping University, January 9–10, 1995.

F. Chatelin and V. Frayssé. About the distance to singularity for nonlinear problems. SCAN'93, Technical University of Vienna, Austria, September 26-29 1993.

F. Chatelin and V. Frayssé. Distances to singularity viewed by computers. XII Householder Symposium on Numerical Algebra, UCLA Conference Center, Lake Arrowhead, CA, U.S.A., June 13-18 1993.

Other talks

V. Frayssé et L. Giraud. Introduction aux techniques performantes en algèbre linéaire. Invited talk tin the framework of the "Centre de Compétences Techniques (CCT) Calcul Scientifique et Modélisation", CNES, December 7 1999.

V. Frayssé. Recent advances in backward error analysis. Invited talk at the Faculty of Sciences of Rabat, Morocco, May 1999.

V. Frayssé et V. Toumazou. Calcul de Pseudospectres. Invited talk at the Faculty of Sciences of Rabat, Morocco, March 1997.

V. Frayssé. Calcul d'éléments propres de matrices symétriques. Université Paul Sabatier, April 1995.

T. Braconnier, V. Frayssé et O. Marques. Calculs de Valeurs Propres au CERFACS. Applications industrielles. Journées du PSNM. Ecole Normale Supérieure de Lyon, March 30-31 1995.

V. Frayssé. Conditionnement de singularités algébriques. Congrès National d'Analyse Numérique, Giens, May 24-28 1993.

Valérie Frayssé

Adresse professionnelle
CERFACS
42 av. Gaspard Coriolis
31057 Toulouse cedex 01
Tel : 05 61 19 30 18
Fax : 05 61 19 30 00
Email : `fraysse@cerfacs.fr`
Internet : `http://www.cerfacs.fr/~fraysse`

Adresse personnelle
30 rue Pargaminières
31000 Toulouse
Tel : 05 61 21 16 82

Née le 26 Février 1967.

FORMATION

1994 Inscription sur les listes de qualification aux fonctions de maître de conférences pour les sections 26 (Mathématiques appliquées et applications des mathématiques) et 27 (Informatique).

1992 **Doctorat** en Informatique et Mathématiques Appliquées de l'Institut National Polytechnique de Toulouse.
Titre : *Sur la fiabilité des calculs sur ordinateurs.* Direction: F. Chatelin et P. Spiteri. Soutenue le 3 Juillet 1992 devant le jury composé de : M. Arioli (Rapporteur, Directeur de Recherche CNR, Pavie, Italie), F. Chatelin (Professeur Université Paris IX Dauphine, Chef du groupe M&I, Thomson-CSF), I.S. Duff (Directeur de Recherche Rutherford Appleton Laboratory, Oxford et Chef de Projet au CERFACS), R. Glowinski (Professeur Université Paris IX et Université de Houston, Directeur du CERFACS), W. Kahan (Professeur Université de Californie, Berkeley), B. Lemaire (Rapporteur, Professeur Université du Languedoc, Montpellier), G. Ruget (Professeur Université Paris XI Orsay, Directeur Recherche et Développement, Thomson-CSF), P. Spiteri (Professeur INP Toulouse).

1989 **D**iplôme d'**E**tudes **A**pprofondies en Informatique, ENSEEIHT, mention Bien.

1989 **Diplôme d'Ingénieur** en Informatique et Mathématiques Appliquées de l'**ENSEEIHT**.

1984-1986 Classes préparatoires, lycée Fermat, Toulouse.

1984 Baccalauréat Série C, mention Très Bien.

EXPÉRIENCE PROFESSIONNELLE

Depuis Novembre 1994, **Chercheur Senior** dans l'équipe Algorithmes Parallèles du CERFACS (Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique), Toulouse.

De Novembre 1992 à Novembre 1994, **Chercheur post-doctoral** au CERFACS : rédaction du livre *Lectures on Finite Precision Computations* avec F. Chaitin-Chatelin publié en 1996 par SIAM, Philadelphie.

De Septembre 1989 à Juillet 1992, **Thèse de Doctorat** de l'Institut National Polytechnique de Toulouse en Informatique et Mathématiques Appliquées sur le thème de la fiabilité des calculs sur ordinateurs, au CERFACS.
Moniteur en informatique à l'ENSEEIHT (École Nationale Supérieure d'Electronique, d'Electrotechnique, d'Informatique et d'Hydraulique de Toulouse).

LANGUES

Anglais lu, écrit et parlé couramment. Langue de travail.
Espagnol lu, écrit et parlé couramment.

ACTIVITÉS DE RECHERCHE

Problématique

La compréhension du comportement des algorithmes en précision finie, l'appréciation de l'erreur commise sur un résultat de calcul sont de réels enjeux industriels, d'autant plus que la simulation numérique intervient de plus en plus souvent comme substitut de l'expérimentation physique. Parmi les approches qui se sont développées pour mieux évaluer l'impact de la précision finie des ordinateurs, l'analyse inverse des erreurs est celle autour de laquelle s'articulent mes activités de recherche. Introduit par J. Wilkinson dans les années soixante, ce concept s'est avéré très puissant pour analyser le comportement des méthodes numériques exécutées sur ordinateur en précision finie, car il permet de distinguer les instabilités mathématiques (notion de conditionnement) des instabilités numériques (notion d'erreur inverse). Ses applications sont nombreuses : qualification de méthodes numériques, estimateurs de conditionnement et d'erreur (LAPACK), ... Le principe de l'analyse inverse des erreurs est de considérer qu'une solution approchée est la solution exacte d'un problème perturbé. Il repose donc sur un modèle sous-jacent des perturbations des données induites implicitement par l'algorithme exécuté en précision finie. C'est en Algèbre Linéaire que ce concept s'est essentiellement développé jusqu'à présent. Les premiers modèles de perturbation utilisés sont souvent trop pessimistes, et il faut les

raffiner en utilisant des perturbation structurées sur des donnés pertinentes, afin de mieux comprendre les algorithmes complexes.

Mes contributions

Depuis ma thèse, mes contributions scientifiques s'inscrivent dans le contexte de l'analyse inverse des erreurs, et concernent l'analyse de perturbations structurées pour les systèmes linéaires ou de moindres carrés (structures du type A^*A intervenant dans les calculs de valeurs singulières de grandes matrices creuses) ou pour le problème généralisé de valeurs propres par exemple (critère d'arrêt de méthodes de type Krylov). Je me suis également intéressée aux instabilités spectrales de matrices fortement non-normales : les outils développés dans ce domaine (spectres perturbés, pseudo-spectres) sont aussi liés à l'analyse inverse des erreurs. Actuellement, je travaille sur la convergence de solveurs itératifs emboîtés, pour déterminer des stratégies d'arrêt des itérations internes afin de garantir la convergence interne mais au moindre coût. Les applications visées sont, en particulier, les calcul de valeurs propres par technique de *shift-and-invert* ou les méthodes de décomposition de domaine de type Schur. Ici encore, l'analyse inverse des erreurs fournit un cadre conceptuel particulièrement riche.

Le choix de ces thèmes a toujours procédé du souci constant d'apporter une expertise aux utilisateurs dans le choix de méthodes numériques ou de développer des codes offrant une robustesse éprouvée. Les collaborations contractuelles décrites plus loin soulignent également la pertinence de ces thèmes dans un contexte industriel. Enfin, le livre *Lectures on Finite Precision Computations*, co-écrit avec F. Chaitin-Chatelin et publié par SIAM en 1996, fait le point sur les notions mathématiques fondamentales intervenant dans l'analyse inverse des erreurs et montre la puissance de cette approche sur des applications numériques très variées. Ce livre, destiné à la fois aux chercheurs et aux ingénieurs, connaît actuellement son deuxième tirage.

Collaborations suivies

- Internationales
  - Depuis 1995, collaboration avec Dr. L. Valdettaro (Chercheur au Politecnico di Milano, Italie) et Dr. M. Rieutord (Observatoire Midi-Pyrénées, Toulouse) pour le calcul de valeurs propres de très grandes matrices issues de la modélisation d'intérieurs stellaires.
  - D'Octobre 1996 à Mai 1999, collaboration avec Dr. M. Bennani et H. Kharraz-Aroussi (ENSIAS, Rabat, Maroc) pour la parallélisation dans un environnement à mémoire distribuée d'un solveur linéaire et d'un solveur de valeurs propres.
  - En 1996, collaboration avec Dr. Roldan Pozo (Chef du groupe Mathematical Software, National Institute of Standards and Technology, U.S.A.) pour le

calcul des portraits spectraux de matrices dans la collection Matrix Market accessible par Internet
(`http://math.nist.gov/MatrixMarket/`). Cette collection permet à la communauté scientifique de sélectionner des matrices-tests vérifiant des critères prescrits et d'en connaître certaines caractéristiques telles que la structure creuse, le conditionnement ou la sensibilité spectrale.

- Nationales
  - Depuis Octobre 1997, collaboration avec B. Philippe (Directeur de Recherches à l'INRIA-IRISA, Rennes) pour la mise en œuvre de techniques de vérification de l'obtention d'un nombre donné de valeurs propres dans une région du plan pré-sélectionnée. Ces techniques sont basées sur le couplage de méthodes de type Arnoldi et d'itérations de sous-espaces.
  - Et bien sûr, collaborations récurrentes avec les chercheurs du Projet Algorithmes Parallèles et ceux des Projets applicatifs du CERFACS.

Sélection de publications

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

V. Frayssé, S. Gratton, and V. Toumazou. Structured backward error and condition number for linear systems of the type $A^*Ax = b$. *BIT*, 40:74–83, 2000.

V. Frayssé and V. Toumazou. A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$. *J. Numer. Linear Algebra Appl.*, 5(1):1–10, 1998.

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

Réalisations logicielles

Mon travail est indissociable d'une pratique logicielle qui comprend le développement

- d'outils pour l'analyse de la stabilité de méthodes numériques (boîte à outils PRECISE) : cette boite à outils est décrite dans F. Chaitin-Chatelin and V. Frayssé, *Lectures on Finite Precision Computations*, SIAM, Philadelphia, 1996. Les codes FORTRAN ont été réalisés en collaboration avec V. Toumazou et A. McCoy dans le cadre du projet européen PINEAPL.
- de routines implantant des solveurs linéaires efficaces et robustes (GMRES, FGMRES, Gradient Conjugué). Les codes FORTRAN ont été réalisés en collaboration avec L. Giraud et S. Gratton et sont décrits dans les rapports suivants :

V. Frayssé, L. Giraud, and S. Gratton. A set of GMRES routines for real and complex arithmetics. Technical Report TR/PA/97/49, CERFACS, 1997.
V. Frayssé, L. Giraud, and S. Gratton. A set of Flexible-GMRES routines for real and complex arithmetics. Technical Report TR/PA/98/20, CERFACS, 1998.
V. Frayssé and L. Giraud. A set of Conjugate Gradient routines for real and complex arithmetics. Technical Report TR/PA/00/47, CERFACS, 2000.

J'ai également contribué à la parallélisation du code ARNCHEB qui implante la méthode d'Arnoldi avec accélération de Tchébycheff pour le calcul des éléments propres de grandes matrices.

Ces codes sont adaptés à un environnement parallèle à mémoire distribuée et, autant que possible, sont mis à la disposition de la communauté scientifique pour un usage non-commercial à travers le serveur Web (`http://www.cerfacs.fr/algor/Softs/index.html`). Une licence d'exploitation commerciale pour le code GMRES a été vendue à une société de services allemande qui souhaite l'intégrer dans son code de calcul de structures.

Encadrement scientifique

Depuis Novembre 1994, j'assure l'animation scientifique du groupe Calcul Qualitatif sous la direction scientifique de F. Chaitin-Chatelin. Outre l'encadrement technique des post-docteurs travaillant sur les contrats récapitulés ci-dessus, je participe à l'encadrement scientifique des chercheurs du groupe Calcul Qualitatif. J'ai personnellement assuré les encadrements suivants (ou du moins j'y ai pris une part essentielle) :

- Post-docteurs
  - Sargis Dallakyan, Novembre 1998 à Novembre 2000.
    Sujet : traitement d'ambiguïtés entières pour des données GPS dans le cadre d'une collaboration contractuelle avec le CNES.
  - Vincent Toumazou, Novembre 1996 à Novembre 1998.
    Sujet : développement de la boîte à outils PRECISE dans le cadre du projet européen PINEAPL.
  - Alan McCoy, Mai 1996 à Mai 1998.
    Sujet : développement de la boîte à outils PRECISE dans le cadre du projet européen PINEAPL.
- Doctorants
  - Amina Bouras (Université Toulouse I) : thèse soutenue le 22 Septembre 2000.
    Sujet : étude de la convergence de solveurs itératifs emboîtés. Application aux calculs des valeurs propres de grandes matrices creuses par "shift-and-invert", et à la résolution des systèmes linéaires issus des techniques de décomposition de domaines de type Schur.

J'ai également participé de façon active à l'encadrement de six autre thésards du Groupe Calcul Qualitatif (thèses soutenues : V. Toumazou, S. Gratton, A. Harrabi, A. Ilahi et E. Traviesas; thèses en cours : A. Zaoui) au travers de leur formation, de l'aide scientifique quotidienne et de l'aide à la rédaction de leur manuscrit.
J'ai été membre du jury des thèses de L. Grammont (Université de St Etienne, 1994), V. Toumazou (ERIN, Nancy, 1996), S. Gratton (INP, Toulouse, 1998), A. Harrabi (Toulouse I, 1998) et A. Ilahi (Toulouse I, 1998).

- Stagiaires de $3^{\text{ème}}$ cycle

  - H. Kharraz-Aroussi (ENSIAS, Rabat, Maroc), Octobre 1996 à Mai 1999.
    Cadre : Programme Formation des Formateurs au Maroc (équivalent à une thèse de $3^{\text{ème}}$ cycle).
    Sujet : parallélisation d'un solveur linéaire et d'un solveur de valeurs propres itératifs dans un environnement parallèle distribué. Application à la résolution d'un problème issu de l'astrophysique.
  - J.-C. Rioual (ENSEEIHT $3^{\text{ème}}$ année Informatique et DEA), Octobre 1996 à Août 1997.
    Sujet : implantation de la méthode d'Arnoldi-Tchébycheff pour le calcul de valeurs propres de grandes matrices creuses.
  - A. Guermeur (ENSEEIHT $3^{\text{ème}}$ année Informatique), Octobre 1995 à Juin 1996.
    Sujet : conditionnements de factorisations en algèbre linéaire
  - S. Gratton (ENSEEIHT $3^{\text{ème}}$ année Informatique et DEA), Octobre 1993 à Août 1994.
    Sujet : audit (stabilité, précision) des méthodes numériques employées dans un logiciel de restitution d'orbite du CNES.

- Stagiaires de 2^{nd} cycle

  - S. Goldstein (CUST, Clermont-Ferrand), Mai à Juillet 1999.
    Sujet : traitement d'ambiguïtés entières pour GPS - Etude de Sensibilité.
  - E. Traviesas (IUP-MICSS, Université Paul Sabatier), Mai à Août 1996.
    Sujet : conditionnement de la factorisation de Cholesky - erreur inverse optimale pour le calcul de valeurs propres.

Organisation de conférences

- Industrial days at CERFACS on Inner-Outer iterations – Numerical quality of software coupling. 11–12 Septembre 2000. CERFACS, Toulouse, France. Membre du Comité Scientifique.

- Second Conference on Numerical Analysis and Applications, Rousse, Bulgarie, 11-15 Juin 2000, (co-sponsorisée par SIAM et ILAS). Membre du Comité International de Programme.

- Euro-Par'99, Toulouse, 31 Août - 3 Septembre 1999. Membre du comité local d'organisation et "local chair" d'une session sur la fiabilité des calculs parallèles. Co-éditrice des proceedings publiés par Springer dans la série Lecture Notes in Computer Sciences.
- Congrès National d'Analyse Numérique, Ardèche, 1997. Organisatrice du minisymposium "Difficultés numériques pour le calcul des valeurs propres de matrices de grandes tailles".
- Workshop "Eigenvalues and Beyond", Toulouse, 17-20 Octobre 1995. Co-organisatrice avec F. Chaitin-Chatelin. Participation active dans le processus de sélection des articles pour publication dans BIT.
- International Linear Algebra Year (ILAY), CERFACS, 1995-1996. Membre du comité local d'organisation.
- IMACS Conference on Iterative Methods, Blagoevgrad, Bulgarie, 17-20 Juin 1995. Co-organisatrice avec F. Chaitin-Chatelin du minisymposium "Influence of High Nonnormality on the reliability of Iterative Methods in Computational Linear Algebra".
- Workshop *Reliability of Computations*, CERFACS, 30 Mars - 1er Avril 1993. Membre du comité local d'organisation et du Comité Scientifique.

Referee

Je suis régulièrement sollicitée comme rapporteur sur des articles soumis à des revues internationales : BIT (2 articles), SIAM Journal on Matrix Analysis and Applications (3 articles), SIAM Journal on Applied Mathematics (1 article), SIAM Journal on Scientific Computing (1 article), Linear Algebra and Its Applications (2 articles), Numerical Linear Algebra with Applications (1 article), IEEE Journal on Signal Processing (1 article), IMA Journal on Numerical Analysis (1 article).

Je réalise également des rapports d'évaluation pour des articles soumis à des conférences internationales telles que SCAN, HPCN, Euro-Par, VecPar ou pour le prix Leslie Fox.

ACTIVITÉS CONTRACTUELLES

La recherche de financements, et le suivi technique, administratif et budgétaire de collaborations contractuelles avec l'industrie sont des activités fortement encouragées au CERFACS, dans le but d'obtenir des ressources extérieures permettant d'accroître le potentiel humain et matériel consacré aux recherches plus en amont. Dans la mesure du possible, nous nous efforçons de concevoir des collaborations industrielles dont le contenu scientifique exploite et peut faire progresser l'état de nos recherches. Les contrats dont j'ai

eu ou partagé la responsabilité, tant sur le plan technique que sur le plan administratif, sont les suivants :

- *méthodes numériques en orbitographie* : un contrat avec le CNES (Novembre 1998 à Mars 2000). Je suis en charge d'une étude concernant la convergence de solveurs linéaires itératifs emboîtés. Un état de l'art dans ce domaine émergent ainsi que des applications dans le domaine des systèmes linéaires et des valeurs propres doivent être fournis.

- *orbitographie précise* : un contrat avec le CNES (Novembre 1999 à Septembre 2000). Ce projet concerne l'évaluation de la qualité numérique des méthodes utilisées pour le traitement des ambiguïtés entières apparaissant dans les données de type GPS.

- *fiabilité des calculs* : responsable CERFACS du projet européen Esprit intitulé PINEAPL (Janvier 1996 à Décembre 1998). Ce projet, coordonné par NAG (U.K.), incluait dans son consortium British Aerospace, CERFACS, Thomson-CSF, CPS (Naples, Italie), Danish Hydraulic Institute, IBM SEMEA, Université de Manchester, Math-Tech, et PIAGGIO. Le but était de produire une librairie parallèle robuste. Le CERFACS était particulièrement en charge du développement de la boîte à outils PRECISE à partir de laquelle la robustesse des routines de la librairie parallèle a été testée et qui a permis d'apporter une expertise sur la qualité de la résolution des problèmes soumis par les industriels partenaires.

- *Solveurs linéaires itératifs* : un contrat avec Aerospatiale-Centre Commun de Recherches (1997). J'ai participé à la réalisation technique de ce contrat portant sur l'étude et l'implantation d'une variante de la méthode Block-QMR pour matrices J-symétriques.

- *problèmes de moindres carrés non-linéaires* : 3 contrats avec le CNES (Décembre 1994 à Août 1995; Décembre 1995 à Août 1996; Janvier 1997 à Octobre 1997). J'ai assuré la gestion et participé à la réalisation de ces trois contrats qui s'articulaient autour de la fiabilité de la résolution de problèmes de moindres carrés non-linéaires. A partir d'un audit effectué sur un code d'orbitographie du CNES, le CERFACS devait étudier et proposer des améliorations aux techniques utilisées pour les moindres carrés linéaires et non-linéaires. Ces contrats ont servi au financement d'une thèse (S. Gratton) qui a a permis des avancées théoriques importantes dans le domaine de l'analyse inverse des erreurs.

- *Solveurs linéaires itératifs* : un contrat avec le CNES (Décembre 1995 à Août 1996). Ce projet visait à développer un solveur de type GMRES en arithmétique complexe pour la résolution de systèmes complexes non hermitiens creux de grande taille.

ACTIVITÉS D'ENSEIGNEMENT ET DE FORMATION

Parce que je considère que les activités d'enseignement sont un complément enrichissant aux activités de recherche, j'ai choisi de participer régulièrement à des formations (initiales

ou continues). J'enseigne principalement l'algèbre linéaire appliquée (ou comment aboutir à des bilbiothèques numériques efficaces et robustes sur des architectures scalaires et parallèles) et la fiabilité des calculs sur ordinateur (en particulier l'analyse inverse des erreurs).

- Co-organisation avec L. Giraud d'une formation de trois jours intitulée "Outils de programmation efficace et robuste pour le logiciel scientifique" dispensée à un groupe d'ingénieurs et chercheurs du CNES en Mars 2000.
- Conférencière invitée à l'école d'été Calcul Numérique et Symbolique : cours de 2 heures sur la résolution des problèmes de valeurs propres de grande taille (Rabat, Maroc, 13-17 Septembre 1999).
- Intervenante dans l'option Mécanique des Fluides Numérique ($3^{\text{ème}}$ année) de l'ENSEEIHT pour un cours sur la fiabilité des calculs en précision finie (2 heures en Février 1999, 6 heures prévues en 2000).
- Depuis Octobre 1996, j'organise tous les ans une formation de base pour les nouveaux arrivants au CERFACS. Je donne un cours de 3 heures sur l'Algèbre Linéaire Appliquée.
- Depuis 1997, j'interviens chaque année dans le Mastère de Météorologie de Météo-France. Avec L. Giraud et B. Cuenot, nous avons conçu un enseignement de calcul scientifique (cours et travaux pratiques) qui part de l'équation différentielle discrétisée par éléments finis pour arriver à sa résolution parallèle. J'interviens dans cette formation pour 3 heures de cours et 3 heures de travaux pratiques.
- Co-organisation de deux journées de formation au CERFACS en Avril 1994 intitulées "Calcul Distribué sur Réseaux de Station de Travail". Les cours dispensés durant de ces deux journées étaient destinés aux ingénieurs des organismes partenaires du CERFACS ou ayant des collaborations avec le CERFACS.
- Organisation et réalisation, avec F. Chaitin-Chatelin, d'un cours sur le "Calcul Qualitatif", du 8 au 10 Juin 1993, dans le cadre du programme européen COMETT-MATARI (volume horaire assuré : 7 heures).
- De Septembre 1989 à Juin 1992, j'ai bénéficié d'un poste de moniteur de l'enseignement à l'UPS. Détachée à l'ENSEEIHT, j'ai assuré des travaux pratiques d'algorithmique et de programmation en $1^{\text{ère}}$ année Informatique.

RESPONSABILITÉS ADMINISTRATIVES

Une des caractéristiques originales du CERFACS est d'avoir choisi des Chefs de Projet dont l'activité principale n'est pas au CERFACS, et qui ne sont donc présents qu'à temps partiel. Les chercheurs seniors ont donc un rôle très important à jouer dans l'organisation

et la gestion des équipes de recherche.
Depuis Novembre 1994, je suis l'un des deux chercheurs seniors (permanents) dans le projet Algorithmique Parallèle dirigée par I. Duff (présent à quart-temps environ) et qui réunit une quinzaine de chercheurs non-permanents (thésitifs et post-docteurs). Mes responsabilités au sein de cette équipe comprennent :

- la gestion administrative et budgétaire (gestion de contrats industriels, achat de matériel et de documentation, recrutement, rédaction de réponses aux appels d'offre européens, représentation de l'équipe ...)
- l'organisation des séminaires (internes et externes), de formations (internes, pour les partenaires du CERFACS, ou externes), et de conférences internationales,
- l'animation du groupe de recherches Calcul Qualitatif, dirigé par F. Chaitin-Chatelin (également présente à temps partiel). Ce groupe a un effectif moyen de 6 chercheurs (thésitifs et post-docteurs) et accueille régulièrement des stagiaires.

Je suis également membre du bureau du Centre de Compétences Techniques "Calcul Scientifique et Modélisation" du CNES, et à ce titre, représentante du CERFACS dans ce CCT.
Enfin, j'ai occupé la fonction de déléguée du personnel (suppléante) de Juin 1997 à Décembre 1998 (mise en place du règlement intérieur, passage au 35 heures, ...).

Computations in the neighbourhood of algebraic singularities

F. Chatelin, V. Frayssé and T. Braconnier

Num. Funct. Anal. Opt., 16:287–302, 1995

A note on the normwise perturbation theory for the regular generalized eigenproblem $Ax = \lambda Bx$

V. Frayssé and V. Toumazou

J. Numer. Linear Algebra Appl., 5:1–10, 1998

Structured backward error and condition number for linear systems of the type $A^*Ax = b$

V. Frayssé, S. Gratton and V. Toumazou

BIT, 40:74–83, 2000

A relaxation strategy for inexact matrix-vector products for Krylov methods

A. Bouras and V. Frayssé

CERFACS Technical Report TR/PA/00/15, 2000.

This report has been submitted to the Journal of Numerical Linear Algebra with Applications. It can be downloaded from the Parallel Algorithms Project Report Web page under the reference TR/PA/00/15 at the URL
`http://www.cerfacs.fr/algor/algo_reports_2000.html`