

The Parallel Algorithms Team at CERFACS

By Iain S. Duff and Serge Gratton

CERFACS—the European Centre for Research and Advanced Training in Scientific Computation—is a research organisation based in Toulouse, France; it is supported by five industrial partners: CNES, the French Space Agency; EADS France, the European Aeronautic and Defence Space Company; EDF, Electricité de France; Météo-France, the French meteorological service; and SNECMA, Société Nationale d'Etude et de Construction de Moteurs d'Aviation.

Within CERFACS are research teams working in specific application areas: aerodynamics, combustion, climate and environment, data assimilation, and electromagnetism. The activities of another group, the more central Parallel Algorithms Project, are the focus of this article.

The research programme of the Parallel Algorithms Project combines the excitement of basic research discoveries with the satisfaction of their application in the solution of large-scale science and engineering problems that arise in academic research, commerce, and industry. We are concerned not only with underlying mathematical and computational science research, but also with the development of new techniques and algorithms, and their implementation on a range of high-performance computing platforms. The research conducted by the project team is in some cases academic (we currently have five PhD students); in other cases the team's results are applied to industrial problems of concern to CERFACS partners or other CERFACS teams. In this article we briefly describe a selection of applied research projects now in progress.

Preconditioners. In an ongoing project, we have been working since 1995 with EADS, our colleagues from ENSEEIHT, and the CERFACS electromagnetics team on preconditioning techniques for the Maxwell equation. The core of this project is the design of solvers for scattering problems discretized by boundary elements. The discretization leads to large dense systems with multiple right-hand sides, for which the action of the coefficient matrix on a vector can be efficiently performed by the fast multipole method.

Outcomes have included the successful use of efficient variants of well-known Krylov methods, where the key ingredients for an efficient implementation are the use of a combination of spectral and approximate inverse preconditioners, and a careful tuning of stopping criteria in the FGMRES variant of the GMRES algorithm. These techniques make it possible to solve systems with up to ten million degrees of freedom. On a more academic side, this work has stimulated basic research on inexact Krylov methods, and on the development of preconditioners based on low-rank corrections of existing preconditioners.

Gradiometry. An interesting area in which we have collaborated with our shareholder CNES since 2001 is determination of the Earth's gravity field by gradiometry, i.e., from gravity-gradient satellite observations. Highly accurate measurements of the Earth's gravity field with appropriate spatial sampling are essential to a better understanding of the processes that take place within the Earth, as well as on and above its surface. Better gravity-field models could be used in a wide range of research and application areas, including global ocean circulation and physical studies of the Earth's interior. In oceanography, for instance, new models would contribute to the recomputation and correction of a major deficiency—the lack of an accurate and precise reference surface for the interpretation of the mean sea level as mapped by satellite altimetry—allowing longer-term studies of the variations of the sea surface.

In the case of the impending GOCE mission (the GOCE satellite will be launched in 2007), coefficients of an expansion of the gravitational potential V are estimated. The considered expansion is

$$V(r, \theta, \lambda) = \frac{GM}{R} \sum_{l=0}^L \left(\frac{r}{R}\right)^{l+1} \sum_{m=0}^l P_{lm}(\cos\theta) [C_{lm} \cos m\lambda + S_{lm} \sin m\lambda],$$

where G is the gravitational constant, M is the Earth's mass, R is the Earth's reference radius, P_{lm} represents the fully normalized l -degree Legendre polynomials of order m , and C_{lm}, S_{lm} are the corresponding normalized harmonic coefficients. For the mission, the chosen value for L is 300, and we consider the following parameter estimation problem: Find the harmonic coefficients C_{lm} and S_{lm} as accurately as possible, using the satellite observations. This results in a linear least-squares problem involving millions of equations and 90,000 unknowns that engineers will need to solve on a daily basis on an eight-processor Power 5 IBM machine.

Envisioning the variety of situations that might arise during the mission, we proposed a fast solver based on normal equations and a more robust QR algorithm. The need to keep the memory requirements as low as possible for the target computer stimulated the development of a parallel distributed packed format, and of a full linear algebra library based on this format; this was done in collaboration with the ScaLAPACK development team.

Another request, from GOCE physicists, was that we study the sensitivity of a linear combination of the solution components to perturbations of the matrix and the right-hand side of the linear least-squares problem. Using a generalization of the condition number for least squares, our team answered the physicists' question and proposed several approximations with a range of computational costs and accuracies. These results could be used for any other application area to which such sensitivity issues are relevant.

Data assimilation. Data assimilation is the combination of physical observations and mathematical models of a system to produce forecasts. It is used intensively in numerical weather/ocean forecasting by Météo France, as well as in the work of the CERFACS climate modelling team. After time and space discretization and under the Gaussian assumption on the observation errors, the problem to be solved is a nonlinear least-squares problem with millions of unknowns.

A crucial requirement is the ability to solve this problem in as close to real time as possible. As shown ten years ago, the Gauss–Newton process is the most effective solver for large-scale problems of this type. Gauss–Newton, in other words, is the solver that provides the best reduction of the function being minimized in the time allowed, and is therefore the algorithm used by most operational centres (e.g., ECMWF, Météo France, . . .) to produce their forecasts.

Using the model developed by our CERFACS colleagues from the climate modelling team as a real-case example, we have developed new (limited-memory) preconditioners for the iterative solution of the linear least-squares problem that arises in the Gauss–Newton process. In theory, these preconditioners possess very interesting spectral properties, and our hope is that they will improve on existing preconditioners in accelerating the convergence of the unpreconditioned Gauss–Newton process.

This work, done in collaboration with researchers at the University of Namur, has a variety of possible extensions. A promising track that we are now exploring is the introduction of a combination of multigrid and trust region techniques for the optimization.

The Parallel Algorithms Project runs an extensive visitor programme, and each year in June we hold a “Sparse Days at CERFACS” meeting. A photograph taken at the 2006 meeting appears in “The Diverse and Active European Scene.” Further information about the group’s activities can be found at www.cerfacs.fr/algorithm/.

Iain Duff, the current chair of the SIAM Board of Trustees, is group leader for numerical analysis in the Department of Computational Science and Engineering at Rutherford Appleton Laboratory and project leader of the Numerical Algorithms Group at CERFACS. Serge Gratton, a senior researcher at CERFACS, oversees the day-to-day activities of the group. Previously, from 1999 to 2002, he worked at CNES on determining satellite orbits using GPS.