Science & Technology Facilities Council

Technical Report

# Numerical Analysis Group Progress Report: January 2008 – December 2009

**Jennifer A. Scott (Editor)**

January 2010

**RAL-TR-2010-001**

Library and Information Services
SFTC Rutherford Appleton Laboratory
Harwell Science and Innovation Campus
Didcot
OX11 0QX
UK
Tel: +44 (0)1235 445384
Fax: +44(0)1235 446403
Email: library@rl.ac.uk



**ISSN 1358-6254**

**RAL-TR-2010-001**

# Numerical Analysis Group Progress Report January 2008 – December 2009

Jennifer A. Scott (Editor)

## ABSTRACT

We discuss the research activities of the Numerical Analysis Group in the Computational Science and Engineering Department at the Rutherford Appleton Laboratory of the Science and Technology Facilities Council for the period January 2008 to December 2009. This work was principally supported by EPSRC grants EP/E053351/1, EP/F006535/1 and EP/F005369/1.

**Keywords:** large-scale optimization, sparse matrices, direct methods, iterative methods, ordering techniques, stopping criteria, parallel, multicore, numerical linear algebra, HSL, GALAHAD, CUTEr.

**AMS(MOS) subject classifications:** 65F05, 65F50.

---

Current reports are available from `http://www.numerical.stfc.ac.uk/reports/reports.html`.

Computational Science and Engineering Department
Atlas Centre
Rutherford Appleton Laboratory
Oxon OX11 0QX

January 20, 2010

## Contents

# Personnel in Numerical Analysis Group

## Staff

Jennifer Scott.
Group Leader.
Sparse linear systems and sparse eigenvalue problems, high-performance computing.

Mario Arioli
Numerical linear algebra, numerical solution of PDEs, error analysis.

Iain Duff
Sparse matrices and high-performance computing.

Nick Gould
Large-scale optimization, nonlinear equations and inequalities.

Jonathan Hogg
HSL, parallel linear algebra, optimization.

Sue Thorne (née Dollar)
Preconditioners, saddle-point problems and sparse linear systems.

Karen McIntyre
Part-time administrative and secretarial support.

Jill Snowdon (from November 2008 until September 2009)
Part-time administrative and secretarial support.

**Consultant:** John Reid
HSL, sparse matrices, automatic differentiation, and Fortran.

**Post-doctoral student:** Daniel Robinson (Oxford)
Nonlinear optimization.

**CASE student:** Philip Brown (Bath)
Structural optimization.

**Visiting Scientist:** Coralia Cartis (Edinburgh)
Interior-point methods, complexity analysis.

## Long term visitors:

Emilie Auzemery (ENSEEIHT)
Software.

Guillaume Martine (ENSEEIHT)
Software.

Philippe Toint (Namur)
Optimization.

Fred Wubs (Groningen)
Linear algebra.

# 1 Overview

This report covers the period January 2008 to December 2009 and summarises the activities of the Numerical Analysis Group within the Computational Science and Engineering Department at the STFC Rutherford Appleton Laboratory. This work was principally supported by EPSRC grants EP/E053351/1, EP/F006535/1, and EP/F005369/1.

On the staff front, in March 2008 we were pleased to be able to welcome Nick Gould back as a full-time member of the Group. Of course, we sympathise with him over the stress he experienced while working at the University of Oxford and trust that he will continue on the road to recovery. While at Oxford Nick was awarded an EPSRC grant EP/F005369/1 that enabled him to take on a post-doctoral student Daniel Robinson. Although formerly based in Oxford, Daniel is very much part of the RAL Group and spends most of his time at RAL collaborating with Nick (see Section 2.3).

Jonathan Hogg spent the calendar year 2008 with us, working on the EPSRC HPC grant EP/F006535/1 that had been awarded to Jennifer the previous spring. Jonathan was a PhD student at Edinburgh and he took leave from Edinburgh to work on this grant at RAL. At the end of the year, he returned to Edinburgh to continue his PhD studies and was pleased to be able to use much of the work he had done at RAL as part of his PhD. In the summer of 2009 the opportunity arose for the Group to appoint a new member whose role would be to combine research with management of the HSL software library. We were delighted to be able to appoint Jonathan to this post and he rejoined the Group in October 2009. He has quickly slipped back into the life of the Group and is gradually taking over from John Reid the responsibilities for the everyday running and maintenance of HSL, as well as becoming one of the key contacts for technical queries from users. Once STFC assumes full responsibility for HSL (see Section 2.5 for more details on this), Jonathan's HSL role will increase but he will, nevertheless, continue to make a significant contribution to the research work of the Group.

We now have a CASE student associated with the Group who is partially funded by sales of HSL software. Phil Brown is a second year graduate student at the University of Bath. His supervisors at Bath are Professor Chris Budd from the Department of Mathematics and Dr Alicia Kim from the Department of Mechanical Engineering. Phil's research work is on structural optimization. He is looking at a formulation of compliance minimisation subject to buckling constraints in order to reach a solution that is not structurally unstable. Phil spent a profitable month during the summer of 2009 at RAL, working with Jennifer, Nick and Daniel, and we look forward to welcoming him again in 2010.

During the summers of 2008 and 2009 we were joined by two French students who were getting some work experience as "stagiaires" as part of their undergraduate degree at ENSEEIHT (Toulouse). During their two-month internships, Guillaume Martine (summer 2008) and Emilie Auzemery (summer 2009), in collaboration with Mario and Iain, implemented MATLAB interfaces for some of the well-known HSL packages.

The other main staff-related issues involved happy domestic events. Our administrative assistant Karen McIntyre gave birth to twin boys in late November 2008 and we welcomed Jill Snowdon, who provided maternity cover until Karen's return in September 2009. Then in October 2008, Sue Dollar married James Thorne, a member of the e-Science department within STFC. Sue has chosen to take her married name and so for a while there may be some confusion around the names Dollar and Thorne.

Most of the Group members are currently financed by EPSRC grant EP/E053351/1. This is a 4 year grant that commenced on 1 October 2007. A condition of the Grant was that we should establish and run a Scientific Advisory Panel (SAP) with membership outside our immediate community to include industrial representation. The remit of the SAP includes:

1. advising us on the scientific direction of our research;
2. making us aware of the strategic challenges in the field of computational engineering;

3. advising us how to maximise our impact outside academia;
4. advising on how to best disseminate our results to increase the uptake of the software outside academia.

We are delighted that the following individuals kindly agreed to serve on the Panel:

- Andrew Cliffe (Nottingham)
- Jack Dongarra (Tennessee, Oak Ridge and Manchester)
- Chris Farmer (Oxford and Schlumberger)
- Leigh Lapworth (Rolls Royce)
- Wil Schilders (Eindhoven)
- Anne Trefethen (Oxford)

The SAP met with us in July 2008 and July 2009. The Group greatly appreciated these visits and are benefiting from the advice and suggestions made by the Panel.

During the last two years, we have welcomed a number of visitors to the Group. In particular, Fred Wubs (Groningen) visited us in late 2008/early 2009. He was partially supported by a small grant that the Group was awarded by the LMS. Philippe Toint (Namur) spent 3 months at Rutherford in early 2008, funded by EPSRC grant EP/G038643/1. Shorter visits were made by Marc Baboulin (Coimbra), Timo Betcke (Reading), Chris Budd and Alicia Kim (Bath), Coralia Cartis (Edinburgh), Ke Chen (Liverpool), Didier Henrion (Toulouse), Michiel Hochstenbach (Eindhoven), Yifan Hu (AT & T), Leigh Lapworth (Rolls Royce), Amos Lawless (Reading), Michal Kočvara (Birmingham), Volker Mehrmann (Berlin), Jorge Nocedal (Northwestern), Wil Schilders (Eindhoven), Mirek Tůma (Academy of Sciences of the Czech Republic), and Paul Van Dooren (Louvain).

The remainder of this report is organised as follows. We end this overview with a report of the conference organised in July 2009 in honour of John Reid's 70th birthday. In Section 2, Group members present brief personal statements. We follow this in Section 3 by a list of technical reports written by Group members, each accompanied by its abstract; this serves to summarize our research. Next we list our journal and conference publications in Section 4. Since producing software is our other main preoccupation, we then list new packages from HSL (formerly the Harwell Subroutine Library) and GALAHAD along with a brief synopsis of their purposes. Briefer lists of conferences attended, seminars presented, and teaching and tutorial activities are given in Sections 6 to 8. We finish with a list of seminars held at RAL (in the joint series with Oxford), in Section 9.

Current information on the research and software-related activities of the Group and on individual Group members can be found at `https://www.cse.scitech.ac.uk/nag`

Jennifer Scott (jennifer.scott@stfc.ac.uk)

## 1.1 Conference on Sparse Matrices for Scientific Computation

In August 2008, John Reid turned 70. The Group felt that this was an occasion we had to celebrate and so we decided to organise a two-day event the following summer with the aim of bringing together leading numerical analysts from the UK and abroad who had worked with John or been influenced by his work. John is an eminent UK numerical analyst who, over a career spanning five decades, has made many important contributions, particularly in the development of sparse matrix technology. His main contributions include: conjugate gradients as an iterative method, Markowitz pivoting for sparse matrices, estimating sparse Jacobians, factorizing and updating linear programming bases, steepest-edge simplex algorithm, the multifrontal method for sparse linear systems, using $2 \times 2$ pivots for sparse symmetric indefinite matrices, and ordering sparse matrices for small wavefront and profile.

Since 1969, John has been a member of the Numerical Analysis Group that began at the Harwell Laboratory and moved to the RAL in 1990. John formally retired when he reached the then compulsory retirement age of 60. Since then, income from the sales of HSL software has enabled the Group to employ John as a part-time consultant. One of the key activities of the Group and one for which it is internationally renowned is the HSL mathematical software library. Amongst its best known packages are those for solving sparse linear systems, optimization, and sparse eigenvalues. Over the years, John has contributed to all these areas and many of his algorithms and packages are widely used today.

The meeting "Sparse Matrices for Scientific Computation" was held over two days (15-16 July 2009) with all talks by invitation. On Day 1, the talks were related to the areas of numerical analysis that John has been involved in, including optimization, sparse direct methods, numerical linear algebra, iterative methods, and automatic differentiation. Day 2 concentrated on algorithms and software for large-scale systems. It included talks by leading experts on the history and importance of mathematical software libraries, and looked to the future and to new challenges for such libraries. Particular emphasis was placed on software for sparse problems that has been developed by the Group and is used to solve practical problems from a range of application areas. The speakers and their titles were:

- Andrew Cliffe (Nottingham). *Flow in a pipe with a sudden expansion.*
- Jack Dongarra (Tennessee, Oak Ridge and Manchester). *PLASMA and Scheduling Dense Linear Algebra on Multicore Chips.*
- Iain Duff (RAL). *Multifrontal Methods.*
- Al Erisman (Seattle Pacific University). *Tacit Assumptions in the Field of Mathematical Software.*
- Roger Fletcher (Dundee). *L-Implicit-U factorization and the Simplex Update.*
- Shaun Forth (Shrivenham). *Automatic Differentiation and Sparse Matrices.*
- Ian Gladwell (Southern Methodist University). *ACM Transactions on Mathematical Software Past, Present and Future.*
- Sven Hammarling (NAG and Manchester). *A Brief History of Numerical Libraries.*
- Kaj Madsen (Technical University of Denmark). *John Reid's influence in informatics and mathematical modelling.*
- Nancy Nichols (Reading). *Conditioning and Preconditioning of Problems in Data Assimilation.*
- Beresford Parlett (Berkeley). *A Simple Explanation of the MRRR Algorithm.*
- Michael Powell (Cambridge). *How many function values are sufficient to estimate the least value of a quadratic function.*
- John Reid (RAL). *Direct solution of large sparse sets of equations on a multicore computer.*
- Michael Saunders (Stanford). *Some sparse LU considerations.*
- Jennifer Scott (RAL). *Too large to handle directly?*
- Nick Trefethen (Oxford). *Spectral accuracy and conformal maps.*

The meeting was held at Cosener's House in Abingdon, just a few miles south of Oxford. Cosener's House occupies a beautiful location on the River Thames and this, combined with wonderful food and lovely weather, contributed to the great atmosphere of the event. There were approximately 60 attendees, including former colleagues and friends as well as a number of students from local universities.

We are very grateful to the LMS and the IMA for providing financial support for this conference.

# 2 Personal statements

## 2.1 Mario Arioli

Mario's work revolves around linear algebra and the numerical solution of PDEs. Mario has collaborated with Iain Duff on the error analysis of Krylov methods when mixed precision arithmetic is used. In particular, FGMRES has been shown both theoretically and numerically to be the best method to recover full backward stability of the solution when a linear system is solved by Gaussian elimination in single precision. The results have been used by Jonathan Hogg and Jennifer Scott to develop `HSL_MA79` (see Section 2.4). Mario's interest in the GMRES class of methods has also produced a novel characterization of the class of real $n \times n$ matrices for which there exists a starting point such that convergence is achieved in exactly $n - 1$ steps.

Mario has continued his collaboration with Serge Gratton (INPT, CERFACS) around the study of the relations between linear algebra and statistics. In particular, they have analysed the case when the least squares method is used to compute reliable realizations of the minimum-variance unbiased estimates for linear regression models and the conjugate gradient method is applied for solving the related normal equations system. Mario and Serge have obtained reliable stopping criteria for the conjugate gradient algorithm. Moreover, they have explained the link between the energy norm of the error and the $\chi^2$ and F distribution and they have used this to define reliable probabilistic thresholds for stopping criteria.

The collaboration with Daniel Loghin (Birmingham), investigating approximating the square roots of special elliptic operators using the generalised Lanczos method and its use in computing preconditioners for Krylov methods applied to Steklov-Poincaré operators, has continued. The theoretical results have been published and their application to domain decomposition has been studied in collaboration with Drosos Kourounis (Ioannina and Stanford) and Daniel.

More recently, in collaboration with Daniel and Emanouil Georgoulis (Leicester), Mario has studied the problem of proving global convergence of adaptive finite element methods where the solutions on each grid are only approximated by the conjugate gradient method. The use of energy norm stopping criteria has been proved theoretically as the key ingredient to global convergence. The numerical results show that the resulting inexact adaptive method is faster than the exact method without loss of global accuracy.

During the summers of 2008 and 2009, Mario helped supervise the ENSEEIHT students Guillaume Martine and Emilie Auzemery. This led to the development of MATLAB interfaces for the HSL packages `MA57, ME57, HSL_MC64, HSL_MC73,` and `HSL_EA19`. They also extensively tested several facilities of a Fortran 95 interface designed by Nick Gould and Sue Thorne to facilitate and standardize the use of HSL packages within MATLAB. An internal report describing guidelines for this standardization process will shortly be made available.

Finally, this autumn, Mario was responsible for organizing the $8^{th}$ Bath/RAL Numerical Analysis Day, which was held at Rutherford on 29th September 2009.

## 2.2 Iain Duff

Iain continues in his joint roles of research associated with the Group's EPSRC grant and as an ambassador for STFC in his capacity as an STFC Senior Fellow. Iain still leads a project at the European Centre for Research and Advanced Training in Scientific Computation (CERFACS) at Toulouse in France, where he supervises PhD students and organizes a regular conference "Sparse Days at CERFACS" each year that is usually attended by other members of the Group. With his French hat, Iain sat on four French PhD defence juries, twice as chairman and once as a "rapporteur". Iain also continues to visit Glasgow regularly in his capacity as a Visiting Professor at Strathclyde. His research interests continue to be in all aspects of sparse matrices, including iterative and hybrid methods as well as direct methods, and in the exploitation of parallel computers. One of his current interests is in developing combinatorial techniques both in the solution of sparse equations by direct methods and in the preprocessing of matrices prior to the solution of the equations. In the latter case, matching and scaling algorithms have particular prominence.

In joint work with Dubravka Mijuca from Serbia, he showed the power of good scaling prior to numerical factorization and solution where really tough problems in multiscale finite-element analysis were solved using standard HSL codes. One of his main projects at CERFACS over the last two years has been the development of algorithms for an efficient solve phase in a parallel distributed memory environment and the further extension of this work to enable the efficient generation of explicit entries of the inverse and the computation of a basis for the null space of sparse matrices. He is also involved in ongoing research concerning numerical pivoting for saddle-point problems and continues to work on this with Fred Wubs (Groningen), who visited RAL in 2008/2009.

Iain is on the Institute of Mathematics and its Applications Council, is a member of the Research Group, is chairman of the Journals Board of Management, is an IMA representative on the International Committee that oversees the ICIAM international conferences on applied mathematics, and is an editor of the IMA Journal of Numerical Analysis. He was reelected for a fifth year as Chairman of the Board of Trustees of SIAM in December 2009 and is the first non-US resident to hold that position. Iain is on several Advisory Boards and Prize Committees, has been on the Programme and Organizing Committee for several international meetings and has given invited talks at meetings in Beijing, Berlin, Brunei, Dagstuhl, Hanoi, Monterey, Nanjing, Rostov-on-Don, and Wroclaw, and has presented seminars in Berkeley, Strathclyde, and RAL.

A major preoccupation that caused a severe dent in research output was being a member of the Applied Mathematics Panel for the Research Assessment Exercise (RAE 2008), the only non-University person on either Maths Panel. All duties with respect to this terminated early in 2009 and the panel happily did not receive too much adverse criticism. He was one of the principal people involved in a roadmapping exercise for EPSRC on numerical analysis and high performance computing and has recently been asked to head a panel on numerical methods and exascale computing for a pan-European initiative, EESI.

## 2.3 Nick Gould

Nick's adventures as Professor of Numerical Optimisation at the University of Oxford and as a Tutorial Fellow in Mathematics at Exeter College ended in March 2008 because of stress-related illness. He was delighted to return full-time to RAL, and is gradually recovering now that the pressures of teaching have evaporated. Nevertheless, he still finds it nearly impossible to lecture and relies on his colleagues to talk about joint work.

Nick's work still revolves around optimization and related linear algebraic issues. GALAHAD continues to evolve and the long-awaited replacements for LANCELOT are in the early testing and evaluation stages. Much of the activity here is in conjunction with Daniel Robinson, a post-doctoral student formerly from the University of San Diego, and now supported by the EPSRC grant Nick obtained while at Oxford. Daniel and Nick are revisiting SQP methods in their natural second-derivative setting. In particular, care has to be taken when far from a solution since second derivatives may mislead iterates and control using simpler first-derivative methods may be necessary. A number of variants are under consideration, and the most promising will ultimately be available in the GALAHAD package `s2qp`. Nick's long-time collaborator Philippe Toint (Namur) is involved with Nick and Daniel in a rival, so-called "funnel" proposal which has SQP aspects but is appealing since it does not require access to matrix (Jacobian or Hessian) elements, merely matrix-vector products. Such matrix-free methods are particularly interesting in applications (such as those arising from continuous problems) involving discretized operators. While a coherent framework has been devised for equality-constrained problems, dealing with inequalities is still in its infancy.

At a different level, Coralia Cartis (Edinburgh), Philippe and Nick have continued their exploration of how bad popular optimization methods might be. The main theoretical breakthrough has been the realisation that popular unconstrained optimization methods based on standard regularisations of Newton's method may actually be as bad as simpler first-order methods such as steepest descent. Fortunately, the newer class of cubic regularisations are provably better in the worst case. Investigations of various levels of constrained optimization are currently underway.

Nick has also focused on basic optimization subproblems. He, Daniel and Sue Thorne have developed improved methods for minimizing general quadratic functions regularised by both trust-regions and cubic terms; GALAHAD packages TRS and RQS are the result. Daniel, Nick and Jorge Nocedal (Northwestern) are also investigating matrix-free quadratic programming methods with a view to the rival SQP approaches mentioned above. Furthermore, Philippe, Nick, Daniel and Michal Kočvara (Birmingham) have been sponsored by a small EPSRC grant (EP/G038643/1) to investigate algebraic multi-level optimization methods. Finally, Nick has ported several popular GALAHAD packages for MATLAB.

On leaving Oxford, Nick was appointed as a Visiting Professor in Numerical Optimisation, and this enables him to spend time with his former colleagues and to continue supervisory and examination roles. In 2009 he was elected as one of 191 inaugural Fellows of the Society for Industrial and Applied Mathematics (SIAM). He is about to start his final year as Editor-in-Chief of the *SIAM Journal on Optimization*, he was appointed as a startup member of the Advisory Board for *Mathematical Programming Computation* last year, and continues on the boards of several other high-impact journals.

## 2.4 Jonathan Hogg

Jonathan worked with the group in 2008 under a 12-month grant for numerical analysis and high performance computing, as time out from his doctoral studies under Julian Hall (Edinburgh). His initial project involved the development of a mixed-precision sparse solver, building on earlier work of Mario Arioli and Iain Duff and in collaboration with Jennifer Scott. This resulted in HSL_MA79, building on the core library codes MA57 and HSL_MA77, combined with multiple scaling and iterative error recovery options.

Jonathan then went on to look at recent advances in dense linear algebra and how they could be applied to sparse work. While becoming familiar with the technology, he wrote the parallel dense Cholesky code HSL_MP54 that used Directed Acyclic Graphs (DAGs) to determine parallelism. He experimented with various scheduling options based on the critical path, but met with only a modest performance improvement. The work was extended to the sparse case, with assistance from Jennifer and John Reid. It resulted in the parallel sparse Cholesky factorization HSL_MA87. This solver has been shown to be competitive with, and for large problems to significantly outperform, other state-of-the-art solvers.

After returning to the University of Edinburgh to finish his PhD, he rejoined the group in October 2009 to undertake support of the HSL software library in addition to research. Since then he has performed further improvements to HSL_MA87, and is currently working with Jennifer on a version for the symmetric indefinite case.

## 2.5 Jennifer Scott

Jennifer has continued her role as Group Leader throughout the last two years and has enjoyed the challenges of combining management responsibilities with research activities. The most challenging and time consuming management issue has involved long negotiations with Aspentech over the future of the Group's software library HSL. This has been necessary because Aspentech holds the marketing rights for HSL but is no longer in a position to fulfil its responsibilities relating to this role. Getting decisions from Aspentech has proved very difficult, which has made negotiating a mutually acceptable new agreement exceptionally frustrating. These negotiations have dragged on for two years, but STFC is hopeful that the situation will shortly be resolved and the Group will be able to take full control of all future marketing and development of HSL.

Much of Jennifer's recent research and software development has been in collaboration with Jonathan Hogg and John Reid. This has turned out to be an excellent working partnership, combining the experience of John and Jennifer with the enthusiasm and new ideas that Jonathan has brought to the Group. Their work has led to the development of a number of important new solvers for inclusion within HSL, notably a mixed-precision sparse direct solver and a DAG-based sparse Cholesky code for use on multicore machines. For large problems, the latter is achieving impressive speedups in excess of 7 on the Group's 8-core machine.

The next step is to extend the code to solve more challenging symmetric indefinite systems. Preliminary results are encouraging.

A very different project has involved working with Mirek Tůma of the Academy of Sciences of the Czech Republic on the development of incomplete Cholesky factorization preconditioners. The emphasis is on the importance of structure within such preconditioners. Jennifer has enjoyed a couple of short but useful visits to Prague (in January and October 2008) and we were pleased to welcome Mirek to Rutherford for a visit in July 2009,

Jennifer has also worked with Sue Thorne on approximate minimum degree orderings, in particular, addressing the case of the matrix having one or more dense rows. Most recently, Jennifer and Sue have started a collaboration with Iain Duff and Jonathan on the development of a nested dissection ordering routine that will be included as part of the suite of sparse matrix ordering routines within HSL. The intention is that this will remove the dependency of some of our sparse direct solvers on external packages and, by focusing only on computing orderings for our solvers, we hope to design and develop a routine that is both efficient and can easily be integrated within other HSL packages.

Jennifer is the editor of the IMA Numerical Analysis Newsletter, producing three newsletters each year. In addition, Jennifer has continued to take an active role in promoting women in mathematics in the UK. She is the UK coordinator for European Women in Mathematics, and is a member of both the Women's Committee of the LMS and the national coordinating committee for WISE (Women in Science, Engineering and Construction). In April 2008, she organised the Women in Mathematics Day at the LMS. This turned out to be useful practice for organising the conference for John Reid in July 2009.

## 2.6 Sue Thorne (née Dollar)

Sue's work continues to revolve around linear algebra and optimisation. Much of Sue's recent work has been on the linear algebraic issues within PDE-constrained optimisation. The linear systems have submatrices that contain discretized PDEs and, as a result, quickly become very large and increasingly ill-conditioned. In collaboration with Tyrone Rees and Andy Wathen (Oxford), Sue has devised effective preconditioners for a class of distributed control problems. These preconditioners give mesh-independent convergence and if a reasonable level of regularisation is used within the optimisation problem, then the relevant iterative methods converge in a very low number of iterations. Recently, Sue has been considering boundary control problems and further classes of distributed control problems. In particular, she has been investigating the spectral properties of the linear systems with respect to the level of regularisation used within the optimisation problem and the mesh sizes used within the discretization. This work has produced interesting results that will be utilised when forming further preconditioners.

As mentioned in Section 2.5, Sue has worked with Jennifer Scott on approximate minimum degree methods that identify and treat any dense rows within a matrix. For some problems, the CPU time to find a pivot ordering has been reduced by several orders of magnitudes without effecting the quality of the ordering. Most recently, Jennifer and Sue have started a collaboration with Iain Duff and Jonathan Hogg to develop a nested dissection routine that will become part of HSL. Currently, some of the HSL packages are dependent on MeTiS, an external package. By developing this new package, we will remove these dependencies. By concentrating on the generation of quality pivot orderings and not partitions, we hope to develop a package that is efficient and effective.

Sue's work has also focused on basic optimization subproblems. Nick Gould, Daniel Robinson (Oxford) and Sue have developed improved methods for minimizing general quadratic functions that have been regularised by both trust-regions and cubic terms. As a result, the packages TRS and RQS have been developed and added to GALAHAD.

Sue has developed a MATLAB interface for HSL_MI20, a package developed by Jennifer and Jonathan Boyle (Manchester) to generate and apply an algebraic multigrid preconditioner. This has been well received by the community and feedback has been very good. Collectively, the Group has developed MATLAB interfaces for a number of packages. Sue has been part of the team that has been unifying the development and structure of these interfaces.

At the beginning of 2008, Sue was elected to the position of Secretary and Treasurer for the UK and Republic of Ireland Section of SIAM (SIAM UKIE); her term started in March 2008. As part of this position, she organised the Annual Meeting in 2009 and is in the process of organising the 2010 Annual Meeting. The 2009 meeting took place in January and was held at the University of Limerick; the 2010 meeting is due to take place on 8 January at the Nation e-Science Centre, Edinburgh. Sue steps down from the position of Secretary and Treasurer for SIAM UKIE at the end of March 2010. However, her association with SIAM will not diminish. At the end of 2009, she was elected to the position of Secretary for the SIAM Activity Group on Linear Algebra.

# 3 Technical Reports

RAL-TR-2008-004 *Roundoff error analysis of orthogonal factorizations of upper Hessenberg rectangular matrices.*
M. Arioli.

Krylov space methods minimizing the 2-norm of the residual (GMRES and MINRES are classical examples) requires the solution of relative small linear least squares problems. The matrix modelling this least square problem is of upper Hessenberg type and the right-hand side is a multiple of the first column of the identity. We specialize some classical roundoff results for Givens (Householder) method to this case pointing out some peculiarities that are useful in the error analysis of Krylov methods such as GMRES, MINRES, and Flexible GMRES.

RAL-TR-2008-005 *Trust-region and other regularisations of linear least-squares problems.*
C. Cartis, N. I. M. Gould and Ph. L. Toint.

We consider methods for regularising the least-squares solution of the linear system $Ax = b$. In particular, we propose iterative methods for solving large problems in which a trust-region bound $\|x\| \leq \Delta$ is imposed on the size of the solution, and in which the least value of linear combinations of $\|Ax - b\|_2^q$ and a regularisation term $\|x\|_2^p$ for various $p$ and $q = 1, 2$ is sought. In each case, one of more "secular" equations are derived, and fast Newton-like solution procedures are suggested. The resulting algorithms are available as part of the GALAHAD optimization library.

RAL-TR-2008-006 *Using FGMRES to obtain backward stability in mixed precision.*
M. Arioli and I. S. Duff.

We consider the triangular factorization of matrices in single precision arithmetic and show how these factors can be used to obtain a backward stable solution. Our aim is to obtain double precision accuracy even when the system is ill-conditioned. We examine the use of iterative refinement and show by example that it may not converge. We then show both theoretically and practically that the use of FGMRES will give us the result that we desire with fairly mild conditions on the matrix and the direct factorization. We perform extensive experiments on dense matrices using MATLAB and indicate how our work extends to sparse matrix factorization and solution.

RAL-TR-2008-007 *The effects of scalings on the performance of a sparse symmetric indefinite solver.*
J. D. Hogg and J. A. Scott.

Scaling is an important part of solving large sparse symmetric linear systems. In a direct method where the analysis is based only on structure it can help by reducing the number of delayed pivots and hence the memory required, the size of the computed factors, and total solution time. In this paper, we examine the effects of scaling on the performance of a sparse symmetric indefinite solver than implements a multifrontal algorithm. We compare several scalings from the mathematical software library HSL, using a large test set of 367 problems from a wide range of practical applications.

RAL-TR-2008-008 *Least-squares problems, normal equations, and stopping criteria for the conjugate gradient method.*
M. Arioli and S. Gratton.

The Conjugate Gradient method can be successfully used in solving the symmetric and positive definite normal equations obtained from least-squares problems. Taking into account the results in the literature which make it possible to approximate the energy norm of the error during the conjugate gradient iterative process, we adapt the energy-norm stopping criterion to the normal equations case. Moreover, we show how the energy norm of the error is linked to the statistical properties of the least-squares problem and to the $\chi^2$-distribution and to the Fisher-Snedecor distribution. Finally, we present the results of several numerical tests that experimentally validate the effectiveness of our stopping criteria.

RAL-TR-2008-009 *On the efficient solution of mixed finite element equations in geometrically multiscale thermal stress analysis.*
I. S. Duff and D. Mijuca.

In order to identify the best technique to solve poorly scaled systems of linear equations arising in a primal-mixed finite-element approach (*FEMIX*) in geometrically multiscale thermoelasticity, we examine the combination of HSL direct sparse solvers and matrix scaling routines. The criteria for optimality were robustness, accuracy and execution time. It will be shown that the present approach *FEMIX-HSL* enables the reliable solution of finite-element model problems where finite elements can differ in size by several orders of magnitude. In addition, it will be shown that by the use of the HSL `MA57` sparse solver and one of the scaling routines `MC64` or `MC30` during the factorization process, the execution time is at least two orders of magnitude faster than using a previous solver based on simple Gaussian elimination and that the accuracy of the solution is maintained even if the system matrix is poorly scaled. A number of pathological tests in elasticity and thermoelasticity are examined to test the robustness and execution time of the solvers. Model problems are examined in nanoindentation and microsized coating, where accurate and straightforward calculation of stress in the solid body and on surfaces of material discontinuity is of the utmost importance. The numerical experiments are performed on a standard PC computing platform.

RAL-TR-2008-010 *Numerical study on incomplete orthogonal factorization preconditioners.*
Z.-Z. Bai, I. S. Duff, and J.-F. Yin.

We design, analyse and test a class of incomplete orthogonal factorization preconditioners constructed from Givens rotations, incorporating some dropping strategies and updating tricks, for the solution of large sparse systems of linear equations. Comprehensive accounts about how the preconditioners are coded, what storage is required and how the computation is executed for a given accuracy are presented. A number of numerical experiments show that these preconditioners are competitive with standard incomplete triangular factorization preconditioners when they are applied to accelerate Krylov subspace iteration methods such as GMRES and BiCGSTAB.

RAL-TR-2008-012 *Discrete interpolation norms with applications.*
M. Arioli and D. Loghin.

We describe norm representations for interpolation spaces generated by finite-dimensional subspaces of Hilbert spaces. These norms are products of integer and noninteger powers of the Grammian matrices associated with the generating pair of spaces for the interpolation space. We include a brief description of some of the algorithms which allow the efficient computation of matrix powers. We consider in some detail the case of fractional Sobolev spaces both for positive and negative indices together with applications arising in preconditioning techniques. Numerical experiments are included.

RAL-TR-2008-013 *A parallel matrix scaling algorithm.*
P. R. Amestoy, I. S. Duff, D. Ruiz, and B. Uçar.

We recently proposed an iterative procedure which asymptotically scales the rows and columns of a given matrix to one in a given norm. In this work, we briefly mention some of the properties of that algorithm and discuss its efficient parallelization. We report on a parallel performance study of our implementation on a few computing environments.

RAL-TR-2008-014 *On the block triangular form of symmetric matrices.*
I. S. Duff and B. Uçar.

We present some observations on the block triangular form (btf) of structurally symmetric, square, sparse matrices. If the matrix is structurally rank deficient, its canonical btf has at least one underdetermined and one overdetermined block. We prove that these blocks are transposes of each other. We further prove that the square block of the canonical btf, if present, has a special fine structure. These findings help us recover symmetry around the anti-diagonal in the block triangular matrix. The uncovered symmetry helps us to permute the matrix in a special form which is symmetric along the main diagonal while exhibiting the blocks of the original btf. As the square block of the canonical btf has full structural rank, the observation relating to the square block applies to structurally nonsingular, square symmetric matrices as well.

RAL-TR-2008-016 *Scaling and pivoting in an out-of-core sparse direct solver.*
J. A. Scott.

Out-of-core sparse direct solvers reduce the amount of main memory needed to factorize and solve large sparse linear systems of equations by holding the matrix data, the computed factors and some of the work arrays in files on disk. The efficiency of the factorization and solution phases is dependent upon the number of entries in the factors. For a given pivot sequence, the level of fill in the factors beyond that predicted on the basis of the sparsity pattern alone depends on the number of pivots that are delayed (that is, the number of pivots that are used later than expected because of numerical stability considerations). Our aim is to limit the number of delayed pivots, while maintaining robustness and accuracy. In this paper, we consider a new out-of-core multifrontal solver that is designed to solve efficiently the systems of linear equations that arise from finite element applications. We consider how equilibration can be built into the solver without requiring the system matrix to be held in main memory. We also examine the effects of different pivoting strategies, including threshold partial pivoting, threshold rook pivoting and static pivoting. Numerical experiments are reported for problems arising from a range of practical applications.

RAL-TR-2008-017 *A Bramble-Pasciak-like method with applications in optimization.*
H. S. Dollar, N. I. M. Gould, M. Stoll and A. J. Wathen.

Saddle-point systems arise in many applications areas, in fact in any situation where an extremum principle arises with constraints. The Stokes problem describing slow viscous flow of an incompressible fluid is a classic example coming from partial differential equations and in the area of optimization such problems are ubiquitous.

In this manuscript we present a framework into which many well-known methods for solving saddle point systems fit. Based on this description we show how new approaches for the solution of saddle-point systems arising in optimization can be derived from the Bramble-Pasciak Conjugate Gradient approach widely used in PDEs and more recent generalizations thereof. In particular we derive a class of new solution methods based on the use of Preconditioned Conjugate Gradients in non-standard inner products and demonstrate how these can be understood through more standard machinery. We show connections to Constraint Preconditioning and give the results of numerical computations on a number of standard optimization test examples.

RAL-TR-2008-018 *Optimal solvers for PDE-Constrained Optimization.*
T. Rees, H. S. Dollar and A. J. Wathen.

Optimization problems with constraints which require the solution of a partial differential equation arise widely in many areas of the sciences and engineering, in particular in problems of design. The solution of such PDE-constrained optimization problems is usually a major computational task. Here we consider simple problems of this type: distributed control problems in which the 2- and 3-dimensional Poisson problem is the PDE. The large dimensional linear systems which result from discretization and which need to be solved are of saddle-point type. We introduce two optimal preconditioners for these systems which lead to convergence of symmetric Krylov subspace iterative methods in a number of iterations which does not increase with the dimension of the discrete problem. These preconditioners are block structured and involve standard multigrid cycles. The optimality of the preconditioned iterative solver is proved theoretically and verified computationally in several test cases. The theoretical proof indicates that these approaches may have much broader applicability for other partial differential equations.

RAL-TR-2008-019 *A second derivative SQP method with imposed descent.*
N. I. M. Gould and D. P. Robinson.

Sequential quadratic programming (SQP) methods form a class of highly efficient algorithms for solving nonlinearly constrained optimization problems. Although second derivative information may often be calculated, there is little practical theory that justifies exact-Hessian SQP methods. In particular, the resulting quadratic programming (QP) subproblems are often nonconvex, and thus finding their global solutions may be computationally nonviable. This paper presents a second-derivative S$\ell_1$QP method based on quadratic subproblems that are either convex, and thus may be solved efficiently, or need not be solved

globally. Additionally, an explicit descent constraint is imposed on certain QP subproblems, which "guides" the iterates through areas in which nonconvexity is a concern. Global convergence of the resulting algorithm is established.

RAL-TR-2008-023 *A fast and robust mixed precision solver for the solution of sparse symmetric linear systems.*
J. D. Hogg and J. A. Scott.

The main bottleneck for emerging computing architectures is memory bandwidth. The amount of data moved around within a sparse direct solver can be approximately halved by using single precision arithmetic. However, the cost of this is a potential loss of accuracy in the solution of the linear systems. Double precision iterative methods preconditioned by a single precision factorization can enable the recovery of high precision solutions more quickly than a sparse direct solver run using double precision arithmetic. The gains from the reduced memory bandwidth are expected to be particularly prominent on multicore machines where the ratio of computational power to memory bandwidth is higher.

In this paper, we develop a practical algorithm to apply such a mixed precision approach and suggest parameters and techniques to minimize the number of solves required by the iterative recovery process. These experiments provide the basis for our new code `HSL_MA79` - a fast, robust, mixed precision sparse symmetric solver that will be included in the mathematical software library HSL.

Numerical results for a wide range of problems from practical applications are presented.

RAL-TR-2008-024 *An efficient out-of-core sparse symmetric indefinite direct solver.*
J. K. Reid and J A. Scott

The fast and accurate solution of large sparse linear systems of equations is important in many problems from computational science and engineering. `HSL_MA77` is a state-of-the-art high-performance, robust, Fortran 95 software package that implements a multifrontal algorithm. It can be used to solve both positive-definite and indefinite linear systems and, by holding the system matrix and its factors out-of-core, it is able to solve very large problems. In this paper, we focus on features of the code that are specifically designed for the more challenging indefinite case and present numerical results for indefinite problems that arise from a range of practical applications.

RAL-TR-2008-026 *Convergence of a regularized Euclidean residual algorithm for nonlinear least-squares.*
S. Bellavia, C. Cartis, N. I. M. Gould, B. Morini and Ph. L. Toint.

The convergence properties of the new Regularized Euclidean Residual method for solving general nonlinear least-squares and nonlinear equations problems are investigated. This method, derived from a proposal by Nesterov (2007). uses a model of the objective function consisting of the unsquared Euclidean linearized residual regularized by a quadratic term. At variance with previous analysis, its convergence properties are here considered without assuming uniformly nonsingular globally Lipschitz continuous Jacobians, nor exact subproblem solution. It is proved that the method is globally convergent to first-order critical points, and, under stronger assumptions, to roots of the underlying system of nonlinear equations. The rate of convergence is also shown to be quadratic under stronger assumptions.

RAL-TR-2008-027 *Guidelines for the development of HSL software, 2008 Version.*
J. K. Reid and J. A. Scott

HSL is a collection of portable, fully documented, and tested Fortran packages for large-scale scientific computation. It has been developed by the Numerical Analysis Group at the Rutherford Appleton Laboratory, with additional input from other experts and collaborators.

The aim of this report is to provide clear and comprehensive guidelines for those involved in the design, development and maintenance of software for HSL. It explains the organisation of HSL, including the use of version numbers and naming conventions, the aims and format of the user documentation, the programming language standards and style, and the verification and testing procedures.

This version supersedes RAL-TR-2006-031.

RAL-TR-2008-028 *Spectral analysis of saddle point matrices with indefinite leading blocks.*
N. I. M. Gould and V. Simoncini.

We provide eigenvalue intervals for symmetric saddle-point and regularised saddle-point matrices in the case where the (1,1) block may be indefinite. These generalise known results for the definite (1,1) case. We also study the spectral properties of the equivalent augmented formulation, which is an alternative to explicitly dealing with the indefinite (1,1) block. Such an analysis may be used to assess the convergence of suitable Krylov subspace methods. We conclude with spectral analyses of the effects of common block-diagonal preconditioners.

RAL-TR-2008-029 *A DAG-based parallel Cholesky factorization for multicore systems*
J. D. Hogg.

Modern processors have multiple cores, making multiprocessing essential for competitive desktop linear algebra. Asynchronous processing with much inherent parallelism can be derived by using a directed acyclic graph (DAG) to represent the data dependencies between tasks.

In this paper, we present our implementation of a DAG-based Cholesky factorization, comparing several different scheduling approaches for the prioritisation of tasks. We demonstrate that complex scheduling approaches offer only a small performance improvement over very simple heuristics.

Our factorization is implemented in Fortran 95 using OpenMP.

RAL-TR-2008-31 *Discrete fractional Sobolev norms for domain decomposition preconditioning*
M. Arioli, D. Kourounis, and D. Loghin.

We present a new approach for preconditioning the interface Schur complement arising in domain decomposition of second-order scalar elliptic problems. The preconditioners are discrete interpolation norms recently introduced by the authors [RAL-TR-2008-012,SINUM 2009]. In particular, we employ discrete representations of norms for the Sobolev space of index 1/2 to approximate the Steklov-Poincaré operators arising from non-overlapping one-level domain decomposition methods. We use the coercivity and continuity of the Schur complement with respect to the preconditioning norm to derive mesh-independent bounds on the convergence of iterative solvers. We also address the case of non-constant coefficients by considering the interpolation of weighted spaces and the corresponding discrete norms.

RAL-TR-2009-001 *A second derivative SQP method: global convergence.*
N. I. M. Gould and D. P. Robinson.

Sequential quadratic programming (SQP) methods form a class of highly efficient algorithms for solving nonlinearly constrained optimization problems. Although second derivative information may often be calculated, there is little practical theory that justifies exact-Hessian SQP methods. In particular, the resulting quadratic programming (QP) subproblems are often nonconvex, and thus finding their global solutions may be computationally nonviable. This paper presents a second-derivative SQP method based on quadratic subproblems that are either convex, and thus may be solved efficiently, or need not be solved globally. Additionally, an explicit descent-constraint is imposed on certain QP subproblems, which "guides" the iterates through areas in which nonconvexity is a concern. Global convergence of the resulting algorithm is established.

RAL-TR-2009-002 *A second derivative SQP method: local convergence.*
N. I. M. Gould and D. P. Robinson.

Gould and Robinson (RAL-TR-2009-001) gave global convergence results for a second-derivative SQP method for minimizing the exact $\ell_1$-merit function for a fixed value of the penalty parameter. To establish this result, we used the properties of the so-called Cauchy step, which was itself computed from the so-called predictor step. In addition, we allowed for the computation of a variety of (optional) SQP steps that were intended to improve the efficiency of the algorithm.

Although we established global convergence of the algorithm, we did not discuss certain aspects that are critical when developing software capable of solving general optimization problems. In particular, we must have strategies for updating the penalty parameter and better techniques for defining the positive-definite matrix $B_k$ used in computing the predictor step. In this paper we address both of these issues. We consider two techniques for defining the positive-definite matrix $B_k$—a simple diagonal approximation and a more sophisticated limited-memory BFGS update. We also analyse a strategy for updating the penalty parameter based on approximately minimizing the $\ell_1$-penalty function over a sequence of increasing values of the penalty parameter.

Algorithms based on exact penalty functions have certain desirable properties. To be practical, however, these algorithms must be guaranteed to avoid the so-called Maratos effect. We show that a nonmonotone variant of our algorithm avoids this phenomenon and, therefore, results in asymptotically superlinear local convergence; this is verified by preliminary numerical results on the Hock and Shittkowski test set.

RAL-TR-2009-003 *On solving trust-region and other regularised subproblems in optimization.*
H. S. Dollar, N. I. M. Gould and D. P. Robinson.

The solution of trust-region and regularisation subproblems which arise in unconstrained optimization is considered. Building on the pioneering work of Gay, Moré and Sorensen, methods which obtain the solution of a sequence of parametrized linear systems by factorization are used. Enhancements using high-order polynomial approximation and inverse iteration ensure that the resulting method is both globally and asymptotically at least superlinearly convergent in all cases, including in the notorious hard case. Numerical experiments validate the effectiveness of our approach. The resulting software is available as packages TRS and RQS as part of the GALAHAD optimization library, and is especially designed for large-scale problems.

RAL-TR-2009-004 *A DAG-based sparse Cholesky solver for multicore architectures.*
J. D. Hogg, J. K. Reid and J A. Scott

In this paper, we describe the design and development of a new code for the solution of sparse symmetric positive-definite linear systems aimed primarily at multicore architectures. Our new Fortran 95/OpenMP code, HSL_MA87, is available as part of the software library HSL and extends to the sparse case the task DAG-based approach that is becoming popular in the design and development of dense linear algebra kernels.

Comparisons are made with existing parallel solvers, using problems arising from a range of practical applications. We demonstrate that HSL_MA87 obtains good serial and parallel times on an 8-core machine.

RAL-TR-2009-006 *An adaptive cubic regularization algorithm for nonconvex optimization with convex constraints and its function-evaluation complexity.*
C. Cartis, N. I. M. Gould and Ph. L. Toint.

The adaptive cubic overestimation algorithm described by Cartis, Gould and Toint (RAL-TR-2007-007) is adapted to the problem of minimizing a nonlinear, possibly nonconvex, smooth objective function over a convex domain. Convergence to first-order critical points is shown under standard assumptions, but without any Lipschitz continuity requirement on the objective's Hessian. A worst-case complexity analysis in terms of evaluations of the problem's function and derivatives is also presented for the Lipschitz continuous case and for a variant of the resulting algorithm. This analysis extends the best known bound for general unconstrained problems to nonlinear problems with convex constraints.

RAL-TR-2009-007 *A note on a simple constrained ordering for saddle-point systems.*
J. A. Scott

A well-known problem with sparse direct solvers is that, if numerical pivoting is required, the number of entries in the computed factors can be significantly greater than the number predicted on the basis of the sparsity pattern alone. In this note, we review a simple constrained ordering recently proposed by Bridson for saddle-point systems. Bridson's approach allows the factorization to be computed without numerical pivoting but numerical experiments show that the computed factors are generally significantly denser than those obtained by prescaling the matrix and then using an unconstrained ordering combined with threshold partial pivoting.

RAL-TR-2009-013 *On accurate and time efficient solution of primal-mixed finite-element equations in multiscale solid mechanics.*
I. S. Duff and D. Mijuca.

In order to identify the best technique to solve a class of geometrically multiscale model problems in thermoelasticity, we examine a combination of a primal–mixed finite element approach and direct sparse solvers and matrix scaling routines. The criteria for optimality are robustness, accuracy and execution time. It will be shown that the present finite element approach, where displacement and stress variables are simultaneously solved from large scale indefinite poorly scaled systems of equations using the sparse HSL solver MA57 with the aid of the matrix scaling routines MC64 or MC30 during the factorization process, enables a reliable solution even if hexahedral finite elements in a mesh differ in size up to six orders of magnitude. A number of tests in multiscale elasticity and thermoelasticity are examined to test the accuracy and execution time efficiency of the proposed solution approach on a standard PC computing platform.

RAL-TR-2009-014 *Flexible deflation in Krylov methods with Chebyshev-based polynomial filters*
M. Arioli and D. Ruiz.

We consider the solution of ill-conditioned symmetric and positive definite large sparse linear systems of equations. These arise, for instance, when using some symmetrizing preconditioning technique for solving a general (possibly unsymmetric) ill-conditioned linear system, or in domain decomposition of a numerically difficult elliptic problem. Combining Chebyshev iterations with the Lanczos algorithm, we propose a way to identify and extract precise information related to the ill-conditioned part of the given linear system. This approach is equivalent to a flexible deflation based on Chebyshev filters. The potential of this combination, which can be related to the factorization and direct solution of linear systems, is illustrated numerically and theoretically. In particular, we also present a general theory that relates the level of filtering to the accuracy of the computed solution.

RAL-TR-2009-015 *Partial factorization of a dense symmetric indefinite matrix.*
J. K. Reid and J. A. Scott.

At the heart of a frontal or multifrontal solver for the solution of sparse symmetric sets of linear equations, there is the need to partially factorize dense matrices (the frontal matrices) and to be able to use their factorizations in subsequent forward and backward substitutions. For a large problem, packing (holding only the lower or upper triangular part) is important to save memory. It has long been recognized that blocking is the key to efficiency and this has become particularly relevant on modern hardware. For stability in the indefinite case, the use of interchanges and $2 \times 2$ pivots as well as $1 \times 1$ pivots is equally well established. It is shown here that it is possible to use these three ideas together to achieve stable factorizations of large real-world problems with good execution speed. The ideas are not restricted to frontal and multifrontal solvers and are applicable whenever partial or complete factorizations of dense symmetric indefinite matrices are needed.

RAL-TR-2008-016 *Combinatorial problems in solving linear systems.*
I. S. Duff and B. Uçar.

Numerical linear algebra and combinatorial optimization are vast subjects; as is their interaction. In virtually all cases there should be a notion of sparsity for a combinatorial problem to arise. Sparse matrices therefore form the basis of the interaction of these two seemingly disparate subjects. As the core of many of today's numerical linear algebra computations consists of the solution of sparse linear system by direct or iterative methods, we survey some combinatorial problems, ideas, and algorithms relating to these computations. On the direct methods side, we discuss issues such as matrix ordering; bipartite matching and matrix scaling for better pivoting; task assignment and scheduling for parallel multifrontal solvers. On the iterative method side, we discuss preconditioning techniques including incomplete factorization preconditioners, support graph preconditioners, and algebraic multigrid. In a separate part, we discuss the block triangular form of sparse matrices.

RAL-TR-2009-018 *Properties of linear systems in PDE-constrained optimization. Part I: Distributed control.*
H. S. Dollar.

Optimization problems with constraints that contain a partial differential equation arise widely in many areas of science. In this paper, we consider distributed control problems in which the 2- and 3-dimensional Poisson problem is the PDE. If a discretize-the-optimization approach is used to solve the optimization problem, then a large dimensional, symmetric and indefinite linear system must be solved. In general, distributed control problems include a regularization term, the size of which is determined by a real value known as the regularization parameter. The spectral properties and, hence, the condition number of the linear system are highly dependent on the size of this regularization parameter. We derive intervals that contain the eigenvalues of the linear systems and, using these, we are able to show that if the regularization parameter is larger than a certain value, then backward-stable direct methods will compute solutions to the discretized optimization problem that have relative errors of the order of machine precision: changing the value of the regularization parameter within this interval will have negligible effect on the accuracy but the condition number of the system may have dramatically changed. We also analyse the spectral properties of the Schur complement and reduced systems derived via the nullspace method. Throughout the paper, we complement the theoretical results with numerical results.

RAL-TR-2009-020 *An analysis of GMRES worst case convergence*
M. Arioli.

Krylov space methods minimizing the 2-norm of the residual, such as GMRES, used in solving a linear system with an unsymmetric matrix of order $n \times n$ can present pathological cases where the convergence will be achieved only after $n-1$ steps. Here we will characterize the class of real matrices for which a starting point inducing this worst case convergence exists always.

RAL-TR-2009-21 *Convergence of inexact adaptive nite element solvers for elliptic problems*
M. Arioli, E. H. Georgoulis, and D. Loghin.

We address the question of convergence of practical adaptive finite element solvers for symmetric elliptic problems, in the case where the resulting linear systems on each level of the algorithm are solved approximately. In particular, we show that if a particular smallness criterion (involving residuals of the linear solver and the a-posteriori bounds used in the adaptive finite element algorithm) is satisfied, then the adaptive algorithm satisfies a contraction property between two consecutive levels of refinement. This generalises the current convergence analyses of adaptive algorithms for elliptic problems whereby the resulting linear systems are assumed to be solved exactly. Moreover, based on known and new results for the estimation of the residual of the conjugate gradient method, we show that the smallness criterion gives rise to a practical stopping criterion for the iterations of the linear solver, which guarantees that the (inexact) adaptive algorithm converges. A series of numerical experiments highlights the practicality of the theoretical developments.

RAL-TR-2009-023 *On the complexity of steepest descent, Newton's method and regularized Newton methods for nonconvex unconstrained optimization.*
C. Cartis, N. I. M. Gould and Ph. L. Toint.

It is shown that the steepest descent and Newton's method for unconstrained nonconvex optimization under standard assumptions may require numbers of iterations and function evaluations arbitrarily close to $O(\epsilon^{-2})$ to drive the norm of the gradient below $\epsilon$. This shows that the upper bound of $O(\epsilon^{-2})$ evaluations known for the steepest descent method is tight, and that Newton's method may be as slow as steepest descent in the worst case. The improved evaluation complexity bound of $O(\epsilon^{-3/2})$ evaluations known for cubically-regularised Newton methods is also shown to be tight.

RAL-TR-2009-024 *A second-derivative trust-region SQP method with a "trust-region-free" predictor step.*
N. I. M. Gould and D. P. Robinson.

In RAL-TR-2009-001/002 we introduced a second-derivative SQP method (`S2QP`) for solving nonlinear nonconvex optimization problems. We proved that the method is globally convergent and locally superlinearly convergent under standard assumptions. A critical component of the algorithm is the so-called predictor step, which is computed from a strictly convex quadratic program with a trust-region constraint. This step is essential for proving global convergence, but its propensity to identify the optimal active set is paramount for recovering fast local convergence. Thus the global *and* local efficiency of the method is intimately coupled with the quality of the predictor step.

In this paper we study the effects of removing the trust-region constraint from the computation of the predictor step; this is reasonable since the resulting problem is still strictly convex and thus well-defined. Although this is an interesting theoretical question, our motivation is based on practicality. Our preliminary numerical experience with `S2QP` indicates that the trust-region constraint occasionally degrades the quality of the predictor step and diminishes its ability to correctly identify the optimal active set. Moreover, removal of the trust-region constraint allows for re-use of the predictor step over a sequence of failed iterations thus reducing computation. We show that the modified algorithm remains globally convergent and preserves local superlinear convergence provided a nonmonotone strategy is incorporated.

RAL-TR-2009-027 *Design of a multicore sparse Cholesky factorization using DAGs.*
J. D. Hogg, J. K. Reid and J. A. Scott.

The rapid emergence of multicore machines has led to the need to design new algorithms that are efficient on these architectures. Here, we consider the solution of sparse symmetric positive-definite linear systems by Cholesky factorization. We were motivated by the successful division of the computation in the dense case into tasks on blocks and use of a task manager to exploit all the parallelism that is available between these tasks, whose dependencies may be represented by a directed acyclic graph (DAG). Our algorithm is built on the assembly tree and subdivides the work at each node into tasks on blocks, whose dependencies may again be represented by a DAG. To limit memory requirements, updates of blocks are performed directly.

Our algorithm is implemented within a new solver `HSL_MA87`. It is written in Fortran 95 plus OpenMP and is available as part of the software library HSL. Using problems arising from a range of practical applications, we present experimental results that support our design choices and demonstrate `HSL_MA87` obtains good serial and parallel times on our 8-core test machines. Comparisons are made with existing modern solvers and show that `HSL_MA87` generally outperforms these solvers, particularly in the case of very large problems.

This version supersedes RAL-TR-2009-004.

# 4 Publications

P. Amestoy, I.S. Duff, A. Guermouche, and Tz. Slavova, "Analysis of the solution phase of a parallel multifrontal approach". *Parallel Computing*, 2009. DOI: 10.1016/j.parco.2009.06.001.

P. R. Amestoy, I. S. Duff, D. Ruiz, and B. Uçar. "A parallel matrix scaling algorithm". In *Proceedings of VECPAR'08 Conference. Lecture Notes in Computer Science 5336*, pages 301–313. Springer-Verlag, 2008.

M. Arioli and I. S. Duff. "Using FGMRES to obtain backward stability in mixed precision". *Electronic Transactions on Numerical Analysis*, 33:31–44, 2008.

M. Arioli and D. Loghin. "Stopping criteria for mixed finite element problems". *Electronic Transactions on Numerical Analysis*, 29:178–192, 2008.

M. Arioli and D. Loghin. "Discrete Interpolation Norms with Applications". *SIAM Journal on Numerical Analysis*, 47(4):2924–2951, 2009.

Z.-Z. Bai, I. S. Duff, and J.-F. Yin. "Numerical study on incomplete orthogonal factorization preconditioners". *J. Computational and Applied Mathematics*, 226:22–41, 2009.

S. Bellavia, C. Cartis, N. I. M. Gould, B. Morini and Ph. L. Toint, "Convergence of a regularized Euclidean residual algorithm for nonlinear least-squares." To appear *SIAM Journal on Numerical Analysis.*

C. Cartis, N. I. M. Gould and Ph. L. Toint, "Trust-region and other regularisations of linear least-squares problems." *BIT* 49(1):21–53, 2009.

C. Cartis, N. I. M. Gould and Ph. L. Toint, "Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results." *Mathematical Programming* DOI: 10.1007/s10107-009-0286-5, 2009.

C. Cartis, N. I. M. Gould and Ph. L. Toint "Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case iteration complexity." To appear *Mathematical Programming*.

H. S. Dollar, N. I. M. Gould, M. Stoll and A. J. Wathen, "A Bramble-Pasciak-like method with applications in optimization." To appear *SIAM Journal on Scientific Computing*.

H. S. Dollar and J. A. Scott, "A note on fast approximate minimum degree orderings for matrices with some dense rows." *Numerical Linear Algebra with Applications*, 17(1):43–55, 2010.

I. S. Duff. "The design and use of a sparse direct solver for skew symmetric matrices" *J. Computational and Applied Mathematics*, 226:50–54, 2009.

I. S. Duff and D. Mijuca. "On accurate and time efficient solution of primal-mixed finite-element equations in multiscale solid mechanics". To appear *Communications in Numerical Methods in Engineering*.

I. S. Duff and B. Uçar. "On the block triangular form of symmetric matrices". To appear *SIAM Review*.

N. I. M. Gould, "How good are projection methods for convex feasibility problems?" *Computational Optimization and Applications*, 40:1–12, 2008.

N. I. M. Gould and D. P. Robinson, "A second derivative SQP method: local convergence." To appear *SIAM Journal on Optimization*.

N. I. M. Gould and D. P. Robinson, "A second derivative SQP method: global convergence." To appear *SIAM Journal on Optimization*.

N. I. M. Gould and V. Simoncini, "Spectral analysis of saddle point matrices with indefinite leading blocks." *SIAM Journal on Matrix Analysis and Applications*, 31(3):1152–1171, 2009.

N. I. M. Gould and Ph. L. Toint, "Nonlinear programming without a penalty function or a filter." *Mathematical Programming* DOI: 10.1007/s10107-008-0244-7, 2009.

X.-P. Guo and I. S. Duff. "Semilocal and global convergence of the newton-hss method for systems of nonlinear equations". To appear *Numerical Linear Algebra with Applications*.

J. D. Hogg and J. A. Scott, "A fast and robust mixed precision solver for the solution of sparse symmetric linear systems". To appear *ACM Transactions on Mathematical Software*.

T. Rees, H. S. Dollar and A. J. Wathen, "Optimal solvers for PDE-Constrained Optimization". To appear *SIAM Journal on Scientific Computing*.

J. K. Reid and J. A. Scott, "Algorithm 891: a Fortran Virtual Memory System." *ACM Transactions on Mathematical Software*, 36(1): Article 5, 12 pages, 2009.

J. K. Reid and J. A. Scott, "An out-of-core sparse Cholesky solver." *ACM Transactions on Mathematical Software*, 36(2): Article 9, 33 pages, 2009.

J. K. Reid and J. A. Scott, "An efficient out-of-core multifrontal solver for large-scale unsymmetric element problem." *Inter. Journal Numerical Methods in Engineering*, 77(7):901–921, 2009.

J. A. Scott, "Scaling and pivoting in an out-of-core sparse direct solver". To appear *ACM Transactions on Mathematical Software*.

A. Trefethen, N. Higham, I. Duff, and P. V. Coveney. "Developing a high- performance computing/numerical analysis roadmap". *Int. J. of High Performance Computing Applications*, 23(4):423–426, 2009.

# 5 HSL and GALAHAD

## 5.1 HSL

There has been no formal release of HSL within the last two years. In this section, we give details of the new packages made available since January 2008 and the major changes to older packages since the same date.

### 5.1.1 New packages

`HSL_MA64` For a full symmetric matrix that need not be positive definite, this module performs **partial factorizations and solutions of corresponding sets of equations, paying special attention to the efficient use of cache memory**. It performs symmetric interchanges for stability and uses both 1×1 and 2×2 pivots. It is suitable for use in a frontal or multifrontal solver, but may also be used for the direct solution of a full set of equations. Optionally, it may be compiled to use OpenMP.

Eliminations are limited to the leading $p$ rows and columns. Stability considerations may lead to $q \leq p$ eliminations being performed, but there is an option to force all $p$ eliminations to be performed. The factorization takes the form

$$PAP^T = \begin{pmatrix} L_{11} & \\ L_{21} & I \end{pmatrix} \begin{pmatrix} D & \\ & S_{22} \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix}$$

where $A$ has order $n$, $P$ is a permutation matrix, $L_{11}$ is a unit lower triangular matrix of order $q$, $D$ is block diagonal of order $q$, and $S_{22}$ is a symmetric matrix of order $n-q$. The permutation matrix $P$ has the form

$$P = \begin{pmatrix} P_1 & \\ & I \end{pmatrix}$$

where $P_1$ is of order $p$. Each diagonal block of $D$ has size one or two.

Subroutines are provided for partial solutions, that is, solving equations of the form $\begin{pmatrix} L_{11} & \\ L_{21} & I \end{pmatrix} \begin{pmatrix} D & \\ & I \end{pmatrix} X = B$, $\begin{pmatrix} D & \\ & I \end{pmatrix} X = B$, $\begin{pmatrix} D & \\ & I \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix} X = B$, and $\begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix} X = B$, and the corresponding equations for a single right-hand side $b$ and solution $x$.

`HSL_MA79` is a **mixed-precision** sparse symmetric solver for solving one or more linear systems $AX = B$. A factorization of $A$ using single precision (that is, 32-bit real arithmetic) is performed using a direct solver (`MA57` or `HSL_MA77`) and then refinement (iterative refinement and, in some cases, FGMRES) in double precision (that is, 64-bit real arithmetic) is used to recover higher accuracy. This technique is termed a mixed-precision approach. If refinement fails to achieve the requested accuracy, a double-precision factorization is performed.

Use of single-precision arithmetic substantially reduces the amount of data that is moved around within a sparse direct solver, and on a number of modern architectures, it is currently significantly faster than double-precision computation. Thus `HSL_MA79` offers the potential of obtaining a solution to $AX = B$ to double-precision accuracy more rapidly than using a direct solver from HSL in double precision. `HSL_MA79` is primarily designed for solving very large systems.

`HSL_MA84` For a matrix that is full, symmetric and positive definite, this package performs **parallel partial and complete factorisations and solutions of corresponding sets of equations, using OpenMP**. It performs the factorization

$$A = \begin{pmatrix} A_{11} & A_{21}^T \\ A_{21} & A_{22} \end{pmatrix} = \begin{pmatrix} L_{11} & \\ L_{21} & I \end{pmatrix} \begin{pmatrix} I & \\ & S_{22} \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix} = LSL^T$$

where $A$ is order $n$, $L_{11}$ is lower triangular and both $A_{11}$ and $L_{11}$ have order $p \leq n$.

Subroutines are also provided for the complementary partial forward and backward substitutions, that is, solving

$$LX = B \quad \text{and} \quad L^T X = B.$$

HSL_MA87 Given a sparse matrix $A$ that is real, symmetric and positive definite or complex Hermitian and positive definite, HSL_MA87 computes the **sparse Cholesky factorization**

$$A = PL(PL)^T \quad \text{(real symmetric)} \quad \text{or} \quad A = PL(PL)^H \quad \text{(complex Hermitian)}$$

where $P$ is a permutation matrix and $L$ is lower triangular. This is used to solve one or more linear systems $AX = B$. Subroutines are also provided for the complementary forward and backward substitutions, that is, for solving

$$PLX = B \quad \text{and} \quad (PL)^T X = B \quad \text{(real symmetric)} \quad \text{or} \quad (PL)^H X = B \quad \text{(complex Hermitian)}.$$

**This package uses OpenMP and is designed for multicore architectures.**

The efficiency of HSL_MA87 is dependent on the elimination order that the user supplies. The HSL routine HSL_MC68 may be used to obtain a suitable ordering.

HSL_MC34 This subroutine **generates the expanded structure for a matrix $A$ with a symmetric sparsity pattern** given the structure for the lower triangular part. Diagonal entries need not be present.

HSL_MP01 is a simple package that encapsulates all necessary include files for the use of MPI. It is intended to be used in place of an `include 'mpif.h'` directive, thus allowing other library MPI packages to be written in free source form.

HSL_MP54 For a matrix that is full, symmetric and positive definite, this package performs **parallel partial and complete factorisations and solutions of corresponding sets of equations, using OpenMP**. It performs the factorization

$$A = \begin{pmatrix} A_{11} & A_{21}^T \\ A_{21} & A_{22} \end{pmatrix} = \begin{pmatrix} L_{11} & \\ L_{21} & I \end{pmatrix} \begin{pmatrix} I & \\ & S_{22} \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix} = LSL^T$$

where $A$ is order $n$, $L_{11}$ is lower triangular and both $A_{11}$ and $L_{11}$ have order $p \leq n$.

Subroutines are also provided for the complementary partial forward and backward substitutions, that is, solving

$$LX = B \quad \text{and} \quad L^T X = B.$$

### 5.1.2 Changed packages

HSL_MA77 and HSL_MA78 solve one or more sets of sparse symmetric/unsymmetric equations using an out-of-core multifrontal method. The codes now include an option to scale the system matrix out-of-core. After an experience of an unrecoverable stack overflow, the codes have been revised so that recursive subroutines are no longer used. Revisions have also been made to allow the front size to be so large that it can be handled only on a 64-bit machine.

HSL_MC68 Given a symmetric sparse matrix $A = \{a_{ij}\}_{n \times n}$, HSL_MC68 **computes elimination orderings** that are suitable for use with a sparse direct solver. The approximate minimum degree method that the package uses has been changed. In Version 1.0.0, the package is called MC47. In Version 2.0.0, the package uses the AMDD method described in RAL-TR-2007-022. Accordingly, the `restarts` component has been removed from the data type for holding information and the `n_dense_rows` component has been added.

HSL_MP42, HSL_MP43, HSL_MP48, and HSL_MP62 These packages have all been changed to allow the number of processors to exceed the number of subdomains. The addition of HSL_MP01 to HSL has allowed the source form to be changed to free format and means that the user of each of these packages no longer needs an INCLUDE line for the MPI constants. All the pointer array components have been changed to allocatable components, which should be more efficient and avoids any danger of memory leakage.

## 5.2 GALAHAD

Version 2.2 of GALAHAD was released in June 2008, and 2.3 in January 2009. New GALAHAD packages include:

GLRT Given real $n$ by $n$ symmetric matrices $\mathbf{H}$ and $\mathbf{M}$ (with $\mathbf{M}$ positive definite), real $n$ vectors $\mathbf{c}$ and $\mathbf{o}$, and scalars $\sigma \geq 0$, $\epsilon \geq 0$ and $f_0$, this package finds an **approximate minimizer of the regularised objective function**

$$1/\mathrm{p}\,\sigma[\|\mathbf{x}+\mathbf{o}\|_{\mathbf{M}}^2 + \epsilon]^{p/2} + \tfrac{1}{2}\mathbf{x}^T\mathbf{H}\mathbf{x} + \mathbf{c}^T\mathbf{x} + f_0,$$

where $\|\mathbf{v}\|_{\mathbf{M}} = \sqrt{\mathbf{v}^T\mathbf{M}\mathbf{v}}$ is the $\mathbf{M}$-norm of $\mathbf{v}$. This problem commonly occurs as a subproblem in nonlinear optimization calculations involving cubic regularisation. The method may be suitable for large $n$ as no factorization of $\mathbf{H}$ is required. Reverse communication is used to obtain matrix-vector products of the form $\mathbf{H}\mathbf{z}$ and $\mathbf{M}^{-1}\mathbf{z}$.

LSRT Given a real $m$ by $n$ matrix $\mathbf{A}$, a real $m$ vector $\mathbf{b}$ and scalars $\sigma > 0$ and $p \geq 2$, this package finds an **approximate minimizer of the regularised linear-least-squares objective function**

$\frac{1}{2}\|\mathbf{A}\mathbf{x}-\mathbf{b}\|_2^2 + 1/\mathbf{p}\,\sigma\|\mathbf{x}\|_2^p$**.** This problem commonly occurs as a subproblem in nonlinear optimization calculations involving cubic regularisation, and may be used to regularise the solution of under-determined or ill-conditioned linear least-squares problems. The method may be suitable for large $m$ and/or $n$ as no factorization involving $\mathbf{A}$ is required. Reverse communication is used to obtain matrix-vector products of the form $\mathbf{u}+\mathbf{A}\mathbf{v}$ and $\mathbf{v}+\mathbf{A}^T\mathbf{u}$.

LSTR Given a real $m$ by $n$ matrix $\mathbf{A}$, a real $m$ vector $\mathbf{b}$ and a scalar $\Delta > 0$, this package finds an **approximate minimizer of $\|\mathbf{A}\mathbf{x}-\mathbf{b}\|_2$, where the vector $\mathbf{x}$ is required to satisfy the "trust-region" constraint $\|\mathbf{x}\|_2 \leq \Delta$.** This problem commonly occurs as a trust-region subproblem in nonlinear optimization calculations, and may be used to regularize the solution of under-determined or ill-conditioned linear least-squares problems. The method may be suitable for large $m$ and/or $n$ as no factorization involving $\mathbf{A}$ is required. Reverse communication is used to obtain matrix-vector products of the form $\mathbf{u}+\mathbf{A}\mathbf{v}$ and $\mathbf{v}+\mathbf{A}^T\mathbf{u}$.

L2RT Given a real $m$ by $n$ matrix $\mathbf{A}$, a real $m$ vector $\mathbf{b}$ and scalars $\sigma > 0$, $\mu \geq 0$ and $p \geq 2$, this package finds an **approximate minimizer of the regularised linear-least-$\ell_2$-norm objective function** $\sqrt{\|\mathbf{A}\mathbf{x}-\mathbf{b}\|_2^2 + \mu\|\mathbf{x}\|_2^2} + 1/\mathbf{p}\,\sigma\|\mathbf{x}\|_2^p$**.** This problem commonly occurs as a subproblem in nonlinear optimization calculations involving quadratic or cubic regularisation, and may be used to regularise the solution of under-determined or ill-conditioned linear least-squares problems. The method may be suitable for large $m$ and/or $n$ as no factorization involving $\mathbf{A}$ is required. Reverse communication is used to obtain matrix-vector products of the form $\mathbf{u}+\mathbf{A}\mathbf{v}$ and $\mathbf{v}+\mathbf{A}^T\mathbf{u}$.

RQS Given real $n$ by $n$ symmetric matrices $\mathbf{H}$ and $\mathbf{M}$ (with $\mathbf{M}$ diagonally dominant), , another real $m$ by $n$ matrix $\mathbf{A}$, a real $n$ vector $\mathbf{c}$ and scalars $\sigma > 0$, $p > 2$ and $f$, this package finds an **approximate minimizer of the regularised quadratic objective function $\frac{1}{2}\mathbf{x}^T\mathbf{H}\mathbf{x} + \mathbf{c}^T\mathbf{x} + f + 1/\mathbf{p}\,\sigma\|\mathbf{x}\|_{\mathbf{M}}^p$, where the vector $\mathbf{x}$ may additionally be required to satisfy $\mathbf{A}\mathbf{x} = \mathbf{0}$,** and where the $\mathbf{M}$-norm of $\mathbf{x}$ is $\|\mathbf{x}\|_{\mathbf{M}} = \sqrt{\mathbf{x}^T\mathbf{M}\mathbf{x}}$. This problem commonly occurs as a subproblem in nonlinear optimization calculations. The matrix $\mathbf{M}$ need not be provided in the commonly-occurring $\ell_2$-regularisation case for which $\mathbf{M} = \mathbf{I}$, the $n$ by $n$ identity matrix.

Factorization of matrices of the form $\mathbf{H} + \lambda\mathbf{M}$—or

$$\begin{pmatrix} \mathbf{H} + \lambda\mathbf{M} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{pmatrix} \tag{1}$$

in cases where $\mathbf{Ax} = \mathbf{0}$ is imposed—for a succession of scalars $\lambda$ will be required, so this package is most suited for the case where such a factorization may be found efficiently. If this is not the case, the package `GLRT` may be preferred.

`SLS` This package **solves dense or sparse symmetric systems of linear equations** using variants of Gaussian elimination. Given a sparse symmetric matrix $\mathbf{A} = \{a_{ij}\}_{n\times n}$, and an $n$-vector $\mathbf{b}$ or a matrix $\mathbf{B} = \{b_{ij}\}_{n\times r}$, this subroutine solves the system $\mathbf{Ax} = \mathbf{b}$ or the system $\mathbf{AX} = \mathbf{B}$ . The matrix $\mathbf{A}$ need not be definite.

The method provides a common interface to a variety of well-known solvers from HSL and elsewhere. Currently supported solvers include `MA27/SILS`, `HSL_MA57`, `HSL_MA77` and `HSL_MA87` from HSL and `PARDISO` from the PARDISO Project `http://www.pardiso-project.org/`. Note that **the solvers themselves do not form part of this package and must be obtained separately.** Dummy instances are provided for solvers that are unavailable. Also note that additional flexibility may be obtained by calling the solvers directly rather that via this package.

`TRS` Given real $n$ by $n$ symmetric matrices $\mathbf{H}$ and $\mathbf{M}$ (with $\mathbf{M}$ diagonally dominant), another real $m$ by $n$ matrix $\mathbf{A}$, a real $n$ vector $\mathbf{c}$ and scalars $\Delta > 0$ and $f$, this package finds an **approximate minimizer of the quadratic objective function $\frac{1}{2}\mathbf{x}^T\mathbf{Hx} + \mathbf{c}^T\mathbf{x} + f$, where the vector $\mathbf{x}$ is required to satisfy the constraint $\|\mathbf{x}\|_\mathbf{M} \leq \Delta$ and possibly $\mathbf{Ax} = \mathbf{0}$**, and where the $\mathbf{M}$-norm of $\mathbf{x}$ is $\|\mathbf{x}\|_\mathbf{M} = \sqrt{\mathbf{x}^T\mathbf{Mx}}$. This problem commonly occurs as a trust-region subproblem in nonlinear optimization calculations. The package may also be used to solve the related problem in which $\mathbf{x}$ is instead required to satisfy the **equality constraint** $\|\mathbf{x}\|_\mathbf{M} = \Delta$. The matrix $\mathbf{M}$ need not be provided in the commonly-occurring $\ell_2$-trust-region case for which $\mathbf{M} = \mathbf{I}$, the $n$ by $n$ identity matrix.

Factorization of matrices of the form $\mathbf{H} + \lambda\mathbf{M}$—or

$$\begin{pmatrix} \mathbf{H} + \lambda\mathbf{M} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{pmatrix} \tag{2}$$

in cases where $\mathbf{Ax} = \mathbf{0}$ is imposed—for a succession of scalars $\lambda$ will be required, so this package is most suited for the case where such a factorization may be found efficiently. If this is not the case, the package `GLTR` may be preferred.

### 5.3 MATLAB Interfaces

Mex interfaces have been written for the HSL packages: `HSL_EA19`, `MA57`, `ME57`, `HSL_MC64`, `HSL_MC73`, and `HSL_MI20`.

Mex interfaces have also been written for ten GALAHAD packages: `GLRT`, `GLTR`, `LSTR`, `L2RT`, `QPC`, `RQS`, `SILS`, `TRS`, and `WCP`. Installation (under g95) is performed automatically as part of the more general GALAHAD installation process.

## 6 Conference and workshop presentations

31 March - 4 April, 2008, GAMM 2008, Bremen, Germany.
J. A. Scott, *On the development of out-of-core direct solvers.*

6 - 11 April, 2008, Tenth Copper Mountain Conference on Iterative Methods, Copper Mountain, US.
H. S. Dollar, *The challenges of PDE-constrained optimization.*
I. S. Duff, *Using FGMRES to obtain backward stability in mixed precision.*

10–13 May, 2008, SIAM Conference on Optimization, Boston, USA.
H. S. Dollar, *The challenges of PDE-constrained optimization.*
D. P. Robinson, *A second derivative SQP method with imposed descent.*

1–6 June, 2008, Householder Symposium XVII, Berlin, Germany.
M. Arioli *Matrix square root and interpolation spaces.*
H. S. Dollar, *Fast AMD orderings for matrices with some dense rows.*
I. S. Duff, *Practical rank determination for square and rectangular sparse matrices.*
J. A. Scott, *The world of (HSL) sparse direct solvers.*

20–22 June, 2008, Parallel Matrix Algorithms and Applications, Neuchâtel, Switzerland
M. Arioli, *Backward stability of FGMRES.*
J. D. Hogg, *A fast robust sparse mixed precision solver.*

23–24 June, 2008, Sparse Days at CERFACS, Toulouse, France.
M. Arioli, *Backward stability of FGMRES.*
M. Arioli, *Matrix square root and interpolation spaces*
J. D. Hogg, *A fast robust sparse mixed precision solver.*
J. A. Scott, *Gene meets Gronwall.*

14–24 July, 2008, LMS Durham Symposium on Computational Linear Algebra for Partial Differential Equations.
H.Ṡ. Dollar, *Symmetric iterative solvers for symmetric saddle-point problems.*
N.İ. M. Gould, *Challenges for CLAPDE from Optimization: a personal view.*

23 September, 2008, Seventh Bath-RAL Numerical Analysis Day, Bath.
M. Arioli, *Least-squares problems, normal equations, and stopping criteria for the conjugate gradient method*
J. D. Hogg, *A fast robust sparse mixed precision solver.*

13-17 October 2008, International Workshop on Numerical Analysis and Scientific Computing (NASCom08), Rostov-on-Don, Russia.
I. S. Duff, *Exploiting the backward stability of FGMRES.*

1-5 November 2008, Second International Conference on Numerical Algebra and Scientific Computing (NASC08), Beijing, China.
I. S. Duff, *Using FGMRES to obtain backward stability in mixed precision.*

27-30 January 2009, Complex Networks across the Natural and Technological Sciences, University of Strathclyde, Glasgow.
I. S. Duff, *Combinatorial problems in numerical linear algebra.*

1–6 February, 2009 Dagstuhl Seminar on Combinatorial Scientific Computing, Dagstuhl, Germany.
I. S. Duff, *Combinatorial problems in numerical linear algebra.*
J. K. Reid, *Reducing the total bandwidth of a sparse unsymmetric matrix.*
J. A. Scott, *Fast AMD orderings for matrices with some dense rows.*

2-6 March 2009, Fourth International Conference on High Performance Scientific Computing. Modeling, Simulation and Optimization of Complex Processes, Hanoi, Vietnam.
I. S. Duff, *The solution of really large sparse linear equations from three-dimensional modelling.*

13–17 April, 2009, Advances and Perspectives on Numerical Methods for Saddle Point Problems, Banff International Research Station, Canada.
H. S. Dollar, *Projected Krylov methods and PDE-constrained optimization.*

8-13 June 2009, The 2nd International Conference on Mathematical Modelling and Computation and The 5th East Asia SIAM Conference, Universiti Brunei Darussalam, Brunei.

I. S. Duff, *Solving very large sparse linear equations from three-dimensional modelling.*

18–19 June, 2009, Sparse Days at CERFACS, Toulouse, France.
M. Arioli, *Statistical stopping criteria for regression problems.*
J. A. Scott, *A DAG-based sparse Cholesky solver for multicore architectures.*

23–26 June, 2009, Biennial Conference on Numerical Analysis, University of Strathclyde.
H. S. Dollar, *Solving trust-region and other regularised subproblems in optimization.*
I. S. Duff, *Exploiting sparsity in the solution phase for large sparse systems.*
J. D. Hogg, *DAG-driven sparse Cholesky factorization for multicore systems*
J. K. Reid, *On Aasen's tridiagonalization method.*
D. P. Robinson, *A trust-funnel algorithm for general nonlinear optimization.*
J. A. Scott, *Exploiting mixed precision in the solution of sparse symmetric linear systems.*

15–16 July, 2009, Sparse Matrices for Scientific Computation, Coseners House, Abingdon, Oxfordshire.
I. S. Duff, *Multifrontal methods.*
J. K. Reid, *Direct solution of large sparse sets of equations on a multicore computer.*
J. A. Scott, *Too large to handle directly?*

23–28 August, 2009, International Symposium of Mathematical Programming, Chicago, USA.
D. P. Robinson, *S2QP - a second derivative SQP method for nonlinear optimization.*

13-16 September 2009, The 8th International Conference on Parallel Processing and Applied Mathematics (PPAM 2009), Wroclaw, Poland.
I. S. Duff, *Solving large sparse linear equations from discretizations of three-dimensional PDEs.*

29 September, 2009. Eighth Bath-RAL Numerical Analysis Day, Rutherford Appleton Laboratory.
I. S. Duff, *Exploiting sparsity in the solution phase for large sparse equations.*
D. P. Robinson, *Second derivative SQP methods for solving large nonlinear optimization problems.*

28 September-2 October , 2009, Warwick, CY902N: CSC / NAG Autumn School in Core Algorithms for High Performance Scientific Computing
M. Arioli, *An introduction to HSL: Design of reliable software.*

11–14 October, 2009, INFORMS annual meeting, San Diego, USA.
D. P. Robinson, *An interior-point trust-funnel algorithm for large-scale nonconvex optimization.*

26–29 October, 2009, SIAM Conference on Applied Linear Algebra 2009, Seaside, Monterey, USA.
I. S. Duff, *Development and history of sparse direct methods.*
I. S. Duff, *Hybrid techniques in the solution of large scale problems.*
J. A. Scott, *The importance of structure in algebraic preconditioners.*

29–31 October, SIAM Workshop on Combinatorial Scientific Computing, Seaside, Monterey, USA.
J. A. Scott, *A DAG-based sparse Cholesky solver for multicore architectures.*

2 November 2009, Sparse Day at Lawrence Berkeley National Lab, Berkeley, California, USA.
I. S. Duff, *Exploiting sparsity in the solution phase for large sparse equations.*

# 7 Seminars and other talks

16 January 2008, CASA Colloquium, Technical University of Eindhoven, The Netherlands.
H. S. Dollar, *The problem with being dense and how to get around it.*

9 February 2008, Gene Around the World, Oxford.
I.S. Duff, *After dinner talk on Gene as a host, guest, colleague and friend.*

17 April 2008, Department of Mathematics, University of Nottingham.
J. A. Scott, *An introduction to the world of sparse direct solvers and the HSL library.*

18 September 2008, Computational Science and Engineering Department, Daresbury Laboratory.
H. S. Dollar, *The problem with being dense and how to get around it.*

12 November 2008, ERGO Seminar, University of Edinburgh.
J. D. Hogg, *Mixed precision solution of linear systems.*

4 December 2008, CMA Seminar, University of Oxford.
J. D. Hogg, *Cholesky factorizations for multicore systems.*

# 8 Lecture Courses, Teaching, Supervision

University of Oxford. 12 lecture MSc course on Sparse Direct Methods, January - March 2009.
I. S. Duff and J. A. Scott

University of Reading. 10 lecture MSc course on Advanced Numerical Solution of Differential Equations, February - March 2008.
H. S. Dollar

University of Reading. 10 lecture MSc course on Advanced Numerical Solution of Differential Equations, February - March 2009.
H. S. Dollar

CEA-EDF-INRIA school: Methods and algorithms for solving large algebraic systems on modern high performance computing systems, Centre de recherche INRIA Sophia-Antipolis-Méditerranée, 30 March - 3 April 2009. M. Arioli and I. S. Duff

Post-doctoral student (University of Oxford): Daniel Robinson (since October 2007).
Supervised by N. I. M. Gould.

Doctoral student (University of Edinburgh): Jonathan Hogg (January - December 2008).
Supervised by J. A. Scott.

CASE student (University of Oxford): Jari Fawkes (since October 2007).
Jointly supervised by N. I. M. Gould and Chris Farmer (Schlumberger).

CASE student (University of Bath): Philip Brown (since October 2008).
Jointly supervised by N. I. M. Gould and J. A. Scott with Professor Chris Budd and Dr Alicia Kim (Bath).

Undergraduates from ENSEEIHT: Guillaume Martine (summer 2008) and Emilie Auzemery (summer 2009).
Jointly supervised by M. Arioli and I. S. Duff.

# 9 Seminars at RAL

| | |
|---|---|
| 7 February 2008 | Prof Paul Van Dooren (Universite Catholique de Louvain, Belgium)<br>Some graph optimization problems in data mining |
| 6 March 2008 | Prof Volker Mehrmann (Technical University of Berlin, Germany)<br>Nonlinear eigenvalue problems with structure. A challenge for current computational methods |

| | |
|---|---|
| 22 May 2008 | Dr Michiel Hochstenbach (Technical University Eindhoven, Netherlands)<br>An overview of the Jacobi-Davidson method |
| 23 October 2008 | Dr Marc Baboulin (University of Coimbra, Portugal)<br>Some issues in dense linear algebra algorithms for multicore and new architectures |
| 27 November 2008 | Dr Alicia Kim (University of Bath)<br>Topology Optimisation: Achievements and Challenges |
| 22 January 2009 | Dr Fred Wubs (University of Groningen, Netherlands)<br>Preconditioning of linear systems in an ocean flow model |
| 12 March 2009 | Prof Ke Chen (University of Liverpool)<br>On fast multilevel algorithms for nonlinear variational imaging models |
| 23 April 2009 | Dr Coralia Cartis (University of Edinburgh)<br>How sharp is the restricted isometry property? An investigation into sparse approximation techniques |
| 4 June 2009 | Dr Amos Lawless (University of Reading)<br>Approximate Gauss-Newton methods using reduced order models |
| 12 November 2009 | Dr Leigh Lapworth (Rolls Royce, Derby)<br>CFD in the Gas Turbine Industry |
| 26 November 2009 | Dr Timo Betcke (University of Reading)<br>Invariant pairs of matrix polynomials |