

Some numerical experiments on scaling and updating L-BFGS diagonal preconditioners

Fabrice Veersé and Didier Auroux

N° 3858

January 13, 2000

———— THÈME 4 ————



*Rapport
de recherche*

Some numerical experiments on scaling and updating L-BFGS diagonal preconditioners

Fabrice Veersé* and Didier Auroux

Thème 4 — Simulation et optimisation
de systèmes complexes
Projet IDOPT

Rapport de recherche n° 3858 — January 13, 2000 — 25 pages

Abstract: A numerical study is performed to assess the impact of different limited-memory BFGS (L-BFGS) diagonal preconditioner update formulae and scaling strategies on the minimization performance. The formulae studied are those of Gilbert and Lemaréchal (1989) and a generalized version of the recently proposed quasi-Cauchy formula of Zhu *et al.* (1999). The scaling strategies are those of Gilbert and Lemaréchal (1989) and a new approach that is proposed.

This numerical study uses a large number of test problems from the MODULOPT, MINPACK-2 and CUTE collections. Some rather stringent criteria are used for the line-search and the convergence, and the minimization is often performed up to the point where no more progress is achievable.

It is found that the quasi-Cauchy formula overall performs poorly and suffers from a tendency to generate search directions numerically orthogonal to the gradient one. The good performance and robustness of the scaled direct BFGS diagonal-preconditioner update formula proposed by Gilbert and Lemaréchal (1989) is confirmed by the results of our experiments. The direct BFGS formula with the new scaling approach proposed is found to be much less robust. However, it significantly outperforms all the other update formulae in some cases.

Key-words: large-scale minimization, quasi-Newton, limited-memory BFGS, diagonal preconditioning, quasi-Cauchy

(Résumé : *tsvp*)

* Fabrice.Veerse@imag.fr

Expériences numériques sur la mise à l'échelle et la mise à jour des préconditionneurs diagonaux dans L-BFGS

Résumé : Une étude numérique est effectuée pour évaluer l'impact de différentes formules de mise à jour et de différentes stratégies de mise à l'échelle du préconditionneur diagonal dans l'algorithme BFGS à mémoire limitée (L-BFGS), sur la performance de la minimisation. Les formules étudiées sont celles de Gilbert et Lemaréchal (1989) et une version généralisée de la formule quasi-Cauchy proposée récemment par Zhu *et al.* (1999). Les stratégies de mise à l'échelle sont celles de Gilbert et Lemaréchal (1989) et une nouvelle approche que nous proposons.

Cette étude numérique utilise un grand nombre de problèmes issus des collections de cas tests MODULOPT, MINPACK-2 et CUTE. Des critères plutôt sévères sont utilisés pour la recherche linéaire et la convergence, et la minimisation est souvent effectuée jusqu'au point où plus aucun progrès n'est réalisable.

On trouve que la formule quasi-Cauchy a une mauvaise performance dans l'ensemble et souffre d'une tendance à générer des directions de descente numériquement orthogonales au gradient. Les résultats de nos expériences numériques confirment la bonne performance et la robustesse de la formule BFGS directe avec mise à l'échelle proposée par Gilbert et Lemaréchal (1989). La formule BFGS directe associée à la nouvelle approche pour la mise à l'échelle proposée se révèle bien moins robuste. Cependant, sa performance surpasse largement celles de toutes les autres formules dans quelques cas.

Mots-clé : minimisation de grande échelle, quasi-Newton, BFGS à mémoire limitée, préconditionneur diagonal, quasi-Cauchy

Contents

1	Introduction	4
2	Update formulae and scaling strategies	5
2.1	Oren-Spedicato formula	5
2.2	Inverse BFGS formula	5
2.3	Direct BFGS formula	5
2.4	Inverse DFP formula	5
2.5	Quasi-Cauchy formula	6
2.6	Scaling strategies	6
3	MODULOPT experiments	7
3.1	U1BG1	7
3.2	U1TS0	9
3.3	U1MT1	11
4	MINPACK-2 experiments	13
5	CUTE experiments	19
6	Discussion	20
A	Results of the CUTE experiments	22

1 Introduction

Recently Veersé *et al.* (1999) assessed the impact of different limited-memory BFGS (L-BFGS) diagonal preconditioner update formulae on the quality of the L-BFGS inverse Hessian approximation and on the performance of the minimization, for two variants of a simple inverse problem in meteorology. They tested three diagonal update formulae from Gilbert and Lemaréchal (1989; hereafter denoted by GL89) and the quasi-Cauchy update formula of Zhu *et al.* (1999) generalized to accommodate non-Euclidean scalar products. For this problem the quasi-Cauchy update formula performed poorly for both points of view. As in GL89, scaled versions of the inverse BFGS, inverse DFP and direct BFGS diagonal-preconditioner update formulae were found to perform well in terms of number of iterations and simulations (evaluations of the cost function and its gradient) required to achieve convergence. The latter formula was found to have the best performance. However this scaling strategy, leading to a substantial improvement of the minimization performance, was found to alter significantly the quality of the L-BFGS inverse Hessian approximation. This led Veersé *et al.* (1999) to propose an alternate approach, where a scaled diagonal preconditioner is used for the computation of the descent direction and its original non-scaled version is updated by one of the formulae aforementioned. This strategy not only allowed to recover a good L-BFGS inverse Hessian approximation but also led to a substantial improvement of the minimization performance (about 10 % less simulations required to achieve convergence).

The present paper is an extension of this study to a large class of unconstrained minimization problems from the MODULOPT, MINPACK-2 (Averick *et al.*, 1992) and CUTE (Bongartz *at al.*, 1995) test-problem collections. Moreover the focus is directed toward medium-scale and large-scale problems, for which conjugate-gradient and L-BFGS methods are sometimes the only affordable minimization alternatives. Consequently the attention is essentially restricted to minimization performance, since the cost of assessing the quality of the L-BFGS inverse Hessian approximations using eigen-decompositions is largely prohibitive when the dimension of the control vector is large.

The following section recalls the various diagonal-preconditioner update formulae and scaling strategies studied. Numerical experiments with the MODULOPT test problems are reported in section 3. The next section provides the results of experiments using MINPACK-2 unconstrained test cases. Then section 5 reports some experiments with a number of medium-scale and large-scale problems from the CUTE collection. Based on these results the final discussion provides some recommendations on which diagonal-preconditioner formula to use, and suggests some reasons to explain part of the results obtained.

2 Update formulae and scaling strategies

2.1 Oren-Spedicato formula

With this formula, the diagonal preconditioner is obtained by the usual Oren-Spedicato scaling of the identity matrix (eg GL89, Liu and Nocedal, 1989) :

$$\mathbf{D}_+ = \langle \mathbf{y}, \mathbf{s} \rangle / \langle \mathbf{y}, \mathbf{y} \rangle \mathbf{I}$$

where $\mathbf{s} = \mathbf{x}_+ - \mathbf{x}$ is the difference between the new iterate and the previous one, $\mathbf{y} = \mathbf{g}_+ - \mathbf{g}$ is the corresponding gradient difference, $\langle \cdot, \cdot \rangle$ is the scalar product with respect to which the gradient is defined and the minimization is to be performed, and \mathbf{I} is the identity matrix.

2.2 Inverse BFGS formula

The inverse BFGS diagonal-preconditioner update formula (Eq. (4.6) in GL89) is obtained by taking the diagonal of the matrix \mathbf{D} updated with the inverse BFGS formula. The i -th updated diagonal component is:

$$D_+^{(i)} = D^{(i)} + \left(\frac{1}{\langle \mathbf{y}, \mathbf{s} \rangle} + \frac{\langle \mathbf{D}\mathbf{y}, \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{s} \rangle^2} \right) \langle \mathbf{s}, \mathbf{e}_i \rangle^2 - \frac{2D^{(i)} \langle \mathbf{y}, \mathbf{e}_i \rangle \langle \mathbf{s}, \mathbf{e}_i \rangle}{\langle \mathbf{y}, \mathbf{s} \rangle}. \quad (1)$$

Here $(\mathbf{e}_i)_{1 \leq i \leq n}$ is an orthonormal basis of \mathbb{R}^n for the scalar product $\langle \cdot, \cdot \rangle$.

2.3 Direct BFGS formula

The direct BFGS diagonal-preconditioner update formula (Eq. (4.7) in GL89) results from taking the inverse of the diagonal of the matrix obtained by updating \mathbf{D}^{-1} with the direct BFGS formula. The i -th updated diagonal component is given by :

$$D_+^{(i)} = \left(\frac{1}{D^{(i)}} + \frac{\langle \mathbf{y}, \mathbf{e}_i \rangle^2}{\langle \mathbf{y}, \mathbf{s} \rangle} - \frac{(\langle \mathbf{s}, \mathbf{e}_i \rangle / D^{(i)})^2}{\langle \mathbf{D}^{-1} \mathbf{s}, \mathbf{s} \rangle} \right)^{-1} \quad (2)$$

2.4 Inverse DFP formula

The inverse DFP diagonal-preconditioner update formula (Eq. (4.8) in GL89) corresponds to the diagonal of the matrix obtained by updating \mathbf{D} with the inverse DFP formula. The i -th updated diagonal component is:

$$D_+^{(i)} = D^{(i)} + \frac{\langle \mathbf{s}, \mathbf{e}_i \rangle^2}{\langle \mathbf{y}, \mathbf{s} \rangle} - \frac{(D^{(i)} \langle \mathbf{y}, \mathbf{e}_i \rangle)^2}{\langle \mathbf{D}\mathbf{y}, \mathbf{y} \rangle} \quad (3)$$

2.5 Quasi-Cauchy formula

The quasi-Cauchy diagonal-preconditioner update formula is an extension of Eq. (9) in Zhu *et al.* (1999) to the case of a general scalar product $\langle \cdot, \cdot \rangle$:

$$\mathbf{D}_+ = \begin{cases} \mathbf{D} & \text{if } \langle \mathbf{D}\mathbf{y}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{s} \rangle \\ (\mathbf{I} + \nu\mathbf{G})^{-2} \mathbf{D} & \text{if } \langle \mathbf{D}\mathbf{y}, \mathbf{y} \rangle \neq \langle \mathbf{y}, \mathbf{s} \rangle \end{cases} \quad (4)$$

where \mathbf{G} is the diagonal matrix with i -th diagonal component equal to $\langle \mathbf{y}, \mathbf{e}_i \rangle^2$, and ν is the largest solution of $F(\nu) = \langle \mathbf{y}, \mathbf{s} \rangle$ with

$$F(\nu) = \langle (\mathbf{I} + \nu\mathbf{G})^{-2} \mathbf{D}\mathbf{y}, \mathbf{y} \rangle.$$

This diagonal-preconditioner update formula is obtained by solving the minimization problem

$$\min \langle \mathbf{w}, \mathbf{w} \rangle \text{ such that } \langle (\mathbf{D}^{1/2} + \Omega)^2 \mathbf{y}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{s} \rangle$$

where Ω is a diagonal matrix with nonzero components given by the corresponding components of the vector \mathbf{w} .

Two variants of the quasi-Cauchy formula will be considered: one uses the *oldest* (\mathbf{s}, \mathbf{y}) pair – that is about to be removed and replaced by the newly computed pair – and the other one uses this *newest* (\mathbf{s}, \mathbf{y}) pair. In the former case, the quasi-Cauchy formula is used for iteration numbers greater than the number m of pairs stored for the L-BFGS method; whereas it is used from the second iteration onwards in the latter case. The *oldest*-pair option is considered as one variant of the quasi-Cauchy formula, as it revealed more efficient than the *newest*-pair one in the case study of Veersé *et al.* (1999); only the *newest*-pair option is retained for the other update formulae.

2.6 Scaling strategies

Gilbert and Lemaréchal (1989) showed that a large number of simulations may be saved by scaling the diagonal matrix \mathbf{D} before updating it, that is multiplying it by $\langle \mathbf{y}, \mathbf{s} \rangle / \langle \mathbf{D}\mathbf{y}, \mathbf{y} \rangle$. Such a scaling is irrelevant for the Oren-Spedicato and quasi-Cauchy formulae since the corresponding diagonal preconditioners are always scaled.

Three variants of this scaling will be considered for the direct BFGS formula:

- the diagonal preconditioner is scaled *before* it is updated. This corresponds to the current implementation of L-BFGS in the minimization code M1QN3 of the INRIA MODULOPT library (GL89), using the newest (\mathbf{s}, \mathbf{y}) pair;
- the diagonal preconditioner is scaled *after* it is updated;
- a *new approach* where the diagonal preconditioner that is updated is not scaled, but a scaled version of it is used for the computation of the descent direction.

Based on the results of Veersé *et al.* (1999), the inverse BFGS and inverse DFP formulae are implemented with the new approach only, using the newest (\mathbf{s}, \mathbf{y}) pair. All the experiments reported in this paper were performed with modified versions of the M1QN3 minimization code from the INRIA MODULOPT library.

3 MODULOPT experiments

At present time, 3 unconstrained test cases are available from the MODULOPT test-problem collection:

- U1BG1, an inverse problem in meteorology using a 1D periodic viscous Burgers' equation on a latitude circle;
- U1TS0, a problem in transonic fluid mechanics; and
- U1MT1, an inverse problem in meteorology using a shallow-water equation model on a plane limited area tangent to the Earth surface.

The first problem is described in Veersé *et al.* (1999) where it was used as a case study; the remaining two problems are described in GL89.

3.1 U1BG1

The main results of Veersé *et al.* (1999) with a double-precision version of U1BG1 are summarized in Table 1. The number m of (\mathbf{s}, \mathbf{y}) pairs used is 5. The Wolfe's parameters for the line-search are set to 10^{-4} and 0.9. The minimal distance in sup-norm between two successive iterates was set to the epsilon machine, $\varepsilon_m \approx 2.20 \times 10^{-16}$. When the minimization stopped on this line-search safeguard, the corresponding numbers of iterations and simulations are given between parentheses in the table, this is the case for both options of the quasi-Cauchy formula. The convergence criterion on the ratio of the final to the initial gradient norms was set to the square root of this value, $\varepsilon_g \approx 1.5 \times 10^{-8}$. The dimension n of the control variable is 258.

Surprisingly the quasi-Cauchy formula (Exp. 6 and 7) required more simulations than does the Oren-Spedicato formula, and the corresponding minimizations were stopped during the line search. All the other approaches converged with less iterations and simulations than the Oren-Spedicato formula. The new scaling approach yields better performance than that currently implemented in M1QN3 (Exp. 1), the better gain being obtained in conjunction with the direct BFGS diagonal-preconditioner update formula.

As the quasi-Cauchy experiments attained the minimum value of the cost function, one may argue that their poor performance result from the convergence criterion on the gradient decrease being set too small. However, this minimal value of the cost function was attained at simulation 43 for Exp. 0, simulation 60 for Exp. 6 and simulation 164 for Exp. 7. All the other experiments reached this value in less than 33 simulations. The results are thus not an artifact of a too stringent convergence criterion.

Table 1: UIBG1 experiments, n=258

Experiment	update formula	scaling strategy	(s,y) pair	iterations/simulations final cost function final gradient norm
0	Oren-Spedicato		newest	$\frac{74}{77}$ 0.92805792E+02 0.10655716E-05
1	direct BFGS	before	newest	$\frac{52}{56}$ 0.92805792E+02 0.78712288E-06
2	direct BFGS	after	newest	$\frac{47}{49}$ 0.92805792E+02 0.80882239E-06
3	direct BFGS	new approach	newest	$\frac{43}{45}$ 0.92805792E+02 0.12906082E-05
4	inverse BFGS	new approach	newest	$\frac{48}{52}$ 0.92805792E+02 0.11491123E-05
5	inverse DFP	new approach	newest	$\frac{48}{51}$ 0.92805792E+02 0.12427932E-05
6	quasi-Cauchy		oldest	$\frac{(98)}{139}$ 0.92805792E+02 0.31040104E-05
7	quasi-Cauchy		newest	$\frac{(188)}{248}$ 0.92805792E+02 0.16695687E-05

3.2 U1TS0

Since this experiment is performed in single precision, the minimal distance between successive iterates is set to the corresponding epsilon machine $\varepsilon_m \approx 1.19 \times 10^{-7}$, and the convergence criterion is set to the square root of this value $\varepsilon_g = \sqrt{\varepsilon_m} \approx 3.45 \times 10^{-4}$. The number of update pairs used is $m = 5$. The dimension of the problem n being equal to 403, a maximum of 500 minimization iterations has been allowed. The minimization with the quasi-Cauchy formula using the oldest pair failed to converge within this limit; this is indicated by the brackets for Exp. 6 in Table 2.

The inverse BFGS formula (Exp. 4) provides the smallest final value of the cost function but requires twice more iterations than the Oren-Spedicato formula. Moreover the final value of the cost function with the latter formula is attained only at iteration 427 with the inverse BFGS one. The quasi-Cauchy formula with the newest pair (Exp. 7) requires also more iterations than the Oren-Spedicato one (Exp. 0). The direct BFGS formula in combination with the new scaling approach (Exp. 3) gives the smallest final cost function value among those experiments that require less iterations than the Oren-Spedicato one.

Table 2: U1TS0 experiments, n=403

Experiment	update formula	scaling strategy	(s,y) pair	iterations/simulations final cost function final gradient norm
0	Oren-Spedicato		newest	193/202 0.12423874E-09 0.15033718E-05
1	direct BFGS	before	newest	153/158 0.13943750E-09 0.15387420E-05
2	direct BFGS	after	newest	158/165 0.95612407E-10 0.12559432E-05
3	direct BFGS	new approach	newest	161/169 0.48506216E-10 0.14082593E-05
4	inverse BFGS	new approach	newest	454/464 0.33084030E-11 0.15133713E-05
5	inverse DFP	new approach	newest	184/188 0.50225941E-10 0.15388894E-05
6	quasi-Cauchy		oldest	[500/569] 0.27667835E-04 0.10715576E-02
7	quasi-Cauchy		newest	340/387 0.54465082E-10 0.15578283E-05

3.3 U1MT1

This test problem is available in single floating-point precision and is performed using the same parameters as U1TS0, except that the maximum number of iterations allowed has been increased because of the higher problem dimension ($n=1875$).

The minimization with the inverse BFGS formula was stopped on the minimum authorized distance between two successive iterates, during the computation of the step size. This failure is indicated by the parentheses for Exp. 4 in Table 3. This was already noticed in GL89 (see their Table 4).

The quasi-Cauchy formula using the oldest pair (Exp. 6) performed poorly. It needed more iterations to converge than the Oren-Spedicato formula, and the corresponding final value of the cost function is higher. The experiment with the quasi-Cauchy formula using the newest pair failed; it tends to generate a search direction that is orthogonal to the gradient one.

On the contrary, the performance of the direct BFGS formula combined with the new scaling approach (Exp. 3) is impressive on this test problem : it gives the smallest final cost-function value in 3 times less simulations than the Oren-Spedicato formula (Exp. 0).

Table 3: U1MT1 experiments, n=1875

Experiment	update formula	scaling strategy	(s,y) pair	iterations/simulations final cost function final gradient norm
0	Oren-Spedicato		newest	632/647 0.57942863E+05 0.88985367E+02
1	direct BFGS	before	newest	577/596 0.57561266E+05 0.88255081E+02
2	direct BFGS	after	newest	673/689 0.57194254E+05 0.71056465E+02
3	direct BFGS	new approach	newest	171/173 0.57092719E+05 0.77145439E+02
4	inverse BFGS	new approach	newest	(165/209) 0.65327090E+05 0.76998322E+03
5	inverse DFP	new approach	newest	502/511 0.58321836E+05 0.89816109E+02
6	quasi-Cauchy		oldest	755/836 0.58621035E+05 0.89860970E+02
7	quasi-Cauchy		newest	

4 MINPACK-2 experiments

From the results obtained with the three MODULOPT test problems and in order to reduce the computational burden, it was chosen to retain only five configurations for the experiments using the MINPACK-2 and CUTE test problems:

- the **scalar** configuration, using the Oren-Spedicato update formula with the newest (\mathbf{s}, \mathbf{y}) pair, as it gives a reference for assessing the performance of more complicated update formulae;
- the **original** configuration using the direct BFGS formula with the newest pair, scaling the diagonal preconditioner before updating it. This is currently implemented in the M1QN3 L-BFGS minimization code from INRIA and is the configuration we would like to improve upon.
- the **modified** configuration using the direct BFGS formula with the newest pair and the new scaling strategy;
- the **QC oldest** configuration implementing the quasi-Cauchy formula with the oldest pair;
- the **QC newest** configuration using the quasi-Cauchy formula with the newest pair.

The quasi-Cauchy update formula has been retained because it derives from different principles than the three formulae of GL89, and the promising results of Zhu *et al.* (1999) inclined us to perform such a further comparison. Both options of the quasi-Cauchy formula were retained since it is not clear from the MODULOPT experiments of the previous section which should be preferred.

The unconstrained MINPACK-2 test problems used in our experiments are given in Table 4. The first column indicates the number associated to the test problem, that will be used in the forthcoming tables to describe the results of our experiments. The central column indicates the name of the test problem; and the third column gives the value of the problem parameters used in our numerical experiments. The reader is referred to (Averick *et al.*, 1992) for a description of the test problems.

The results of the numerical experiments are reported in Tables 5 to 8 for problem dimensions equal to 400, 2500, 10000 and 40000 respectively. Each row corresponds to the test-problem number given in the first column. Each column corresponds to one of the configurations for updating the diagonal preconditioner described above. For each configuration and test problem, the top couple of numbers indicates the number of iterations and simulations performed. Immediately below is the final value of the cost function, and the bottom value is the corresponding final gradient norm.

The experiments are performed in double precision. The minimal distance allowed between successive iterates is the epsilon machine, $\varepsilon_m \approx 2.20 \times 10^{-16}$. When the minimization was stopped on this criterion the corresponding number of iterations/simulations are given between parentheses in the tables. The minimization convergence criterion on the relative

Table 4: Unconstrained MINPACK-2 test problems

Problem number	Problem name	Parameters
1	Elastic-plastic torsion	$c=0.5$
2	Pressure distribution in a journal bearing	$ecc=0.1$ $b=10.$
3	Minimal surfaces	
4	Optimal design with composite materials	$\lambda=0.008$
5	Steady-state combustion	$\lambda=1.$
6	Homogeneous superconductors: 2-D Ginzburg-Landau	$numvor=2$

decrease of the gradient norm is set to $\varepsilon_g = \sqrt{\varepsilon_m} \approx 1.5 \times 10^{-8}$. A maximum of 40000 minimization iterations has been allowed. When the minimization was stopped on this maximum value, the corresponding number of iterations/simulations in the tables are written between brackets.

It is noticeable that the quasi-Cauchy formula required systematically more iterations and simulations than the Oren-Spedicato one, except for QC newest on test problem 6. Since the same problem parameters as those reported in Zhu *et al.* (1999) were used and our generalized quasi-Cauchy formula reduces to theirs with the Euclidean scalar product used here, this contradicts somewhat their results. Our first reaction was to suspect an error and to perform a careful check of our implementation of the quasi-Cauchy formula. This implementation follows the lines of Zhu *et al.* (1999) using a bisection root-finding algorithm. The updated diagonal was checked to satisfy the quasi-Cauchy relation to machine precision, making us confident in the correctness of our implementation. The reasons for this contradiction must therefore originate from somewhere else. The main differences between our experiments and those of Zhu *et al.* (1999) are a different value of the Wolfe's parameter on the cost function decrease (10^{-4} for us, 10^{-3} for them) and the different convergence criterion. Whether the quasi-Cauchy is sensitive to such implementation details is a question we leave open.

Clearly the QC newest configuration performs better than the QC oldest one on these MINPACK-2 test problems. Only for problem 3 with dimension 2500 and problem 1 with dimension 40000 does the Oren-Spedicato formula (scalar configuration) need less simulations than all the other configurations. Rarely does the modified configuration require more simulations than the scalar one; on the contrary, its performance can be impressive such as for problem 6 or problem 4 with dimension 40000. Except for the latter cases it is difficult to choose between the scalar, the original and the modified configurations because the smallest number of simulations often does not correspond to the smallest final gradient norm.

Table 5: MINPACK-2 experiments, n=400

Number	scalar	original	modified	QC oldest	QC newest
1	(74/113) -4.3609E-01 3.5827E-08	(67/108) -4.3609E-01 7.2236E-08	(69/101) -4.3609E-01 4.6096E-08	(144/210) -4.3609E-01 9.1750E-08	(130/188) -4.3609E-01 5.2960E-08
2	131/136 -2.8144E-01 5.4998E-08	126/129 -2.8144E-01 3.4302E-08	126/132 -2.8144E-01 3.9758E-08	(198/273) -2.8144E-01 6.0449E-08	(268/309) -2.8144E-01 7.0639E-08
3	(57/92) 1.4206E+00 4.9509E-08	(49/81) 1.4206E+00 9.0339E-08	(54/88) 1.4206E+00 2.7292E-08	(8271/9209) 1.4206E+00 9.1238E-08	(110/151) 1.4206E+00 3.3995E-08
4	(173/203) -1.1248E-02 8.5895E-09	(142/173) -1.1248E-02 1.4067E-08	(144/195) -1.1248E-02 1.0963E-08	(242/302) -1.1248E-02 7.6148E-09	(278/326) -1.1248E-02 1.1909E-08
5	(79/131) -1.0184E+00 1.1135E-07	(87/128) -1.0184E+00 6.4969E-08	(96/150) -1.0184E+00 3.7502E-08	(184/245) -1.0184E+00 1.1236E-07	(150/186) -1.0184E+00 7.6114E-08
6	(1374/1469) 1.6227E+01 5.2105E-07	(541/606) 1.6227E+01 7.3390E-08	(305/369) 1.6227E+01 5.2948E-08	(4247/4753) 1.6227E+01 3.5766E-07	(432/502) 1.6227E+01 4.8699E-07

Table 6: MINPACK-2 experiments, n=2500

Number	scalar	original	modified	QC oldest	QC newest
1	(173/208) -4.3875E-01 7.8437E-08	(182/221) -4.3875E-01 6.8515E-08	(166/204) -4.3875E-01 1.6923E-07	(848/1040) -4.3875E-01 3.3925E-07	(331/360) -4.3875E-01 9.2032E-08
2	(304/359) -2.8264E-01 1.0160E-07	(267/296) -2.8264E-01 1.2181E-07	(295/344) -2.8264E-01 6.3027E-08	(12143/13537) -2.8264E-01 1.0677E-07	(682/743) -2.8264E-01 2.3315E-07
3	(150/180) 1.4212E+00 5.9039E-08	(169/212) 1.4212E+00 5.6109E-08	(172/214) 1.4212E+00 1.2127E-07	[40000/45154] 1.4301E+00 1.5124E-01	(447/509) 1.4212E+00 2.1652E-07
4	(376/414) -1.1359E-02 2.6636E-08	(349/392) -1.1359E-02 1.5903E-08	(330/377) -1.1359E-02 5.6837E-08	(2179/2498) -1.1359E-02 7.8607E-08	(758/817) -1.1359E-02 4.5963E-08
5	(242/277) -1.0185E+00 7.7428E-08	(247/288) -1.0185E+00 5.1127E-08	(223/267) -1.0185E+00 4.2981E-08	(5580/6379) -1.0185E+00 1.2225E-07	(813/896) -1.0185E+00 1.2598E-07
6	(3048/3181) 1.6230E+01 7.1607E-07	(1676/1780) 1.6230E+01 2.3294E-07	(1178/1257) 1.6230E+01 2.3263E-07	(36800/41628) 1.6230E+01 6.8296E-07	(1074/1156) 1.6230E+01 9.6192E-07

Table 7: MINPACK-2 experiments, n=10000

Number	scalar	original	modified	QC oldest	QC newest
1	(365/411) -4.3916E-01 1.1418E-07	(366/398) -4.3916E-01 1.1259E-07	(353/396) -4.3916E-01 1.1605E-07	(5754/6645) -4.3916E-01 6.7105E-07	(836/900) -4.3916E-01 2.9099E-07
2	(586/634) -2.8284E-01 1.8070E-07	(604/658) -2.8284E-01 9.2910E-08	(586/631) -2.8284E-01 1.1956E-07	[40000/44853] 2.3093E-01 2.3295E+00	(1370/1460) -2.8284E-01 2.9978E-07
3	(323/379) 1.4213E+00 1.2740E-07	(309/355) 1.4213E+00 2.3909E-07	(319/369) 1.4213E+00 2.6799E-07	[40000/45075] 1.4610E+00 1.3721E-01	(1059/1141) 1.4213E+00 3.3456E-07
4	(654/678) -1.1377E-02 4.5625E-08	(683/723) -1.1377E-02 2.8322E-08	(609/640) -1.1377E-02 3.6203E-08	[40000/45387] -1.1295E-02 2.3397E-02	(3877/4043) -1.1377E-02 1.2167E-07
5	(427/481) -1.0185E+00 3.4100E-07	(406/442) -1.0185E+00 1.3086E-07	(423/486) -1.0185E+00 7.7086E-08	[40000/45644] -1.0184E+00 1.6276E-02	(3670/3836) -1.0185E+00 4.1118E-07
6	(11306/11668) 1.6230E+01 5.3772E-07	(4519/4685) 1.6230E+01 6.9923E-07	(2886/3002) 1.6230E+01 3.6633E-07	[40000/45429] 1.6280E+01 1.0746E+00	(3011/3205) 1.6230E+01 1.8443E-06

Table 8: MINPACK-2 experiments, n=40000

Number	scalar	original	modified	QC oldest	QC newest
1	(641/691) -4.3927E-01 1.7503E-07	(691/736) -4.3927E-01 1.2688E-07	(695/752) -4.3927E-01 1.8309E-07	[40000/45728] -4.2928E-01 1.4330E-01	(2967/3119) -4.3927E-01 1.4245E-06
2	(1141/1194) -2.8289E-01 1.6650E-07	(1170/1223) -2.8289E-01 2.1884E-07	(1080/1128) -2.8289E-01 2.1218E-07	[40000/44622] 4.1964E+00 3.9702E+00	(5729/5932) -2.8289E-01 5.9404E-07
3	(650/708) 1.4214E+00 2.0670E-07	(607/658) 1.4214E+00 2.1234E-07	(680/726) 1.4214E+00 1.5159E-07	[40000/45317] 1.5004E+00 1.6258E-01	(9669/10047) 1.4214E+00 3.8130E-06
4	(1525/1597) -1.1381E-02 6.4133E-08	(1521/1570) -1.1381E-02 4.9154E-08	(609/640) -1.1377E-02 3.6203E-08	[40000/44414] 1.5798E-02 1.7597E-01	[40000/40581] -1.1266E-02 2.0674E-02
5	(852/903) -1.0185E+00 3.8348E-07	(827/874) -1.0185E+00 2.0429E-07	(864/923) -1.0185E+00 1.5692E-07	[40000/45888] -9.3799E-01 6.9588E-01	(23117/24014) -1.0185E+00 3.6977E-07
6	(30285/31154) 1.6230E+01 1.0307E-06	(15734/16166) 1.6230E+01 9.4727E-07	(9455/9728) 1.6230E+01 6.6083E-07	[40000/45849] 1.6537E+01 2.6935E+00	(7850/8278) 1.6230E+01 4.3532E-06

5 CUTE experiments

A large number of numerical experiments using CUTE test problems have been performed, with the five configurations aforementioned. The detailed results are given in the tables of the appendix, where the first column indicates the SIF name of the test problem (see Bongartz *et al.*, 1995) and the second column gives the dimension of the problem. Double-precision arithmetic has been used. The same minimization parameters as for the MINPACK-2 experiments have been used, except that the maximum number of iterations and simulations have been set to 10000 and 20000 respectively according to the maximal problem dimension used. The format of the results is identical to that used for the MINPACK-2 experiments; when the minimization was stopped because of a too small step during the line-search, this is indicated by parentheses; square and curly brackets are used for minimizations that were stopped on the maximum number of iterations and simulations respectively allowed; a blank box indicates a failure during the minimization, associated with search directions becoming numerically orthogonal to the gradient direction.

The exploitation of such an amount of information is difficult and depends on the quantity of major relevance for the user: number of simulations, final value of the cost function or final value of the gradient norm. Consequently only a few constatations are now given, using the number of simulations as main guidance but taking also (subjectively) into account the other quantities:

- It was decided to consider only those experiments that require less simulations than the scalar configuration. It appeared that for some test problems (MSQRTBLS, VAREIGVL), the scalar configuration is the best alternative;
- the quasi-Cauchy configurations usually require much more simulations than the scalar one. In rare cases one of them behaves very nicely, such as for test problems FREUROTH, SINQUAD and WOODS. But they are also the only configurations that fail because of search directions numerically orthogonal to the gradient one. Besides there is not a quasi-Cauchy configuration clearly superior to the other one. It is worth noting that, although it may require more simulations than the scalar configuration, on some problems (CHAINWOO, DIXON3DQ, LIARWHD) QC oldest is able to compute a solution corresponding to a smaller cost function and possibly a higher gradient norm. On the contrary, it may be trapped on a local minimum with a higher cost function value such as for test problems NONMSQRT, SPMSRTLS and TRIDIA;
- the original configuration usually requires less simulations than the scalar one;
- the modified configuration behaves badly for some test problems (BRYBND, CRAGGLVY, DQRTIC, EDENSCH, ENGVAL1, FREUROTH, HILBERTA, HILBERTB, SCHMVETT, SROSENBR, VAREIGVL). In other cases where it requires more simulations than the scalar configuration, it provides either a smaller final cost function or a smaller final gradient norm. And there is a spectacular benefit in using it for other problems (SBRYBND, TESTQUAD, TRIDIA).

6 Discussion

A number of L-BFGS numerical experiments have been performed with MODULOPT, MINPACK-2 and CUTE test problems, using customized versions of the M1QN3 minimization code from INRIA. Several methods for specifying the diagonal preconditioner, presented in section 2, have been assessed. A new scaling strategy proposed by Veersé *et al.* (1999) has been studied, where the original (unscaled) diagonal preconditioner is updated but its scaled version is used for the computation of the descent direction.

The MODULOPT experiments confirmed the findings of GL89: that the direct BFGS formula is more robust than the inverse BFGS and inverse DFP ones, with the new scaling approach. The quasi-Cauchy formula performed overall badly on this large set of experiments.

For the MINPACK-2 and CUTE experiments five different diagonal-preconditioner update approaches were retained: a scalar configuration using the classical Oren-Spedicato multiple of the identity; the original version of M1QN3 that implements a scaled version of the direct BFGS formula from GL89; a modified version using the direct BFGS formula in conjunction with the new scaling approach; and two versions of the quasi-Cauchy formula of Zhu *et al.* (1999).

Overall, the quasi-Cauchy configurations performed poorly and tend on some problems to generate search directions numerically orthogonal to the gradient direction. The original configuration is the most robust one and is almost always better than the Oren-Spedicato formula. The modified version of M1QN3 does not share this robustness and was found worse or better than the original version in roughly the same number of cases, with spectacular improvements for some problems.

Consequently, we would recommend to implement the scalar, the original and the modified configurations as options in the same minimization code. Compared to the original version, the modified configuration requires either an additional real work array of size n or n additional scalar multiplications. Because of the benefit that can be obtained for some problems, it may be worth providing it as an option to the user.

Our results with the quasi-Cauchy formula are clearly different from those of Zhu *et al.* (1999). Consequently, the correctness of our implementation has been checked carefully and the differences between our experiments and theirs have been identified. But it is not clear whether they could explain such discrepancies.

In a sense, it is not surprising that the direct BFGS formula performs better than the quasi-Cauchy ones. Indeed, for the problem to be well-posed the (inverse) Hessian at the minimum must be definite positive. In practice, this usually means it has a dominant diagonal. Therefore it is sensible to create a diagonal-preconditioner updating formula as in GL89, by restricting the quasi-Newton formulae to the diagonal of the approximate inverse Hessian matrix. The more diagonal-dominant the inverse Hessian matrix is, the better should be the diagonal preconditioner. On the contrary, the quasi-Cauchy formula updates the diagonal preconditioner by adding to it the smallest correction (in the sense of some arbitrary norm) such that the corrected preconditioner satisfies the quasi-Cauchy relation, but does not try to accumulate any information available during the minimization process.

Among the diagonal-preconditioner update formulae used in our experiments, none outperforms the others for all the test problems. It would be useful to identify which type of problem a particular formula is optimal for. This could be done according to the nature of the problem (quadratic, sum of squares or non-convex) and/or the structure of the Hessian eigen-spectrum.

References

- [1] Averick, B.M., Carter, R.G., Moré, J.J and Xue, G.-L. (1992), *The MINPACK-2 Test Problem Collection*, Preprint MCS-P153-0692, Argonne National Laboratory, Argonne, IL.
- [2] Bongartz, I., Conn, A.R., Gould, N. and Toint, Ph.L. (1995), CUTE: Constrained and Unconstrained Testing Environment, *ACM Trans. Math. Software*, 21, pp. 123-160.
- [3] Gilbert, J.-Ch. and Lemaréchal, C. (1989), Some numerical experiments with variable storage quasi-Newton algorithms, *Math. Prog.*, 45, pp. 407-435.
- [4] Liu, D.C. and Nocedal, J. (1989), On the limited memory BFGS method for large scale optimization, *Math. Prog.*, 45, pp. 503-528.
- [5] Veersé, F., Auroux, D. and Fisher, M. (1999), Limited-memory BFGS diagonal preconditioners for a data assimilation problem in meteorology, submitted to *Optimization and Engineering*.
- [6] Zhu, M., Nazareth, J.L. and Wolkowicz, H. (1999), The Quasi-Cauchy Relation and Diagonal Updating, *SIAM J. Optim.*, 9, pp. 1192-1204.

A Results of the CUTE experiments

SIF name	n	scalar	original	modified	QC oldest	QC newest
ARGLINB	100	2/3 4.9626E+01 6.7310E-06	2/3 4.9626E+01 6.7310E-06	2/3 4.9626E+01 6.2967E-06	2/3 4.9626E+01 5.7369E-06	2/3 4.9626E+01 8.1770E-06
ARGLINC	100	2/3 5.1126E+01 1.3367E-05	2/3 5.1126E+01 1.3367E-05	2/3 5.1126E+01 1.4812E-05	2/3 5.1126E+01 1.3496E-05	2/3 5.1126E+01 1.0900E-05
ARWHEAD	5000	12/15 7.7700E-12 2.0271E-08	12/15 1.2210E-11 3.1197E-08	8/10 2.2200E-12 6.4414E-08	14/17 1.1100E-12 2.4198E-08	9/10 2.2200E-12 8.6294E-09
BDQRTC	1000	117/136 3.9838E+03 1.1255E-05	26/27 3.9838E+03 2.7274E-07	66/76 3.9838E+03 1.7593E-05	529/596 3.9838E+03 1.8952E-05	81/91 3.9838E+03 1.5181E-05
BROYDN7D	1000	368/373 3.9356E+02 2.3445E-11	371/378 3.9361E+02 3.8920E-11	246/252 4.0061E+02 2.1822E-11	1013/1251 3.9395E+02 1.5304E-10	424/440 3.8455E+02 4.1980E-11
BRYBND	5000	45/55 2.4527E-10 7.6052E-09	48/58 2.0052E-10 6.5469E-09	134/140 4.6415E-10 1.0826E-08	84/89 3.2176E-10 9.6236E-09	180/210 1.1004E-10 1.0153E-08
CHAINWOO	10000	[10000/10483] 2.9646E+04 1.1025E+03	[10000/10467] 2.9698E+04 2.1794E+03	[10000/10311] 1.4226E+03 1.4087E-03	[10000/10954] 4.5875E+00 2.3871E-01	[10000/10879] 3.8458E+04 4.0863E+03
CHNROSNB	50	247/262 1.5650E-10 1.4530E-09	232/253 2.0407E-11 2.3746E-09	142/160 9.2634E-11 2.1168E-09	311/343 2.6020E-11 2.7623E-09	365/388 4.1096E-12 1.7386E-09
COSINE	10000	(11/68) -9.9990E+03 1.8248E-12	(10/57) -9.9990E+03 5.2221E-11	(26/85) -9.9990E+03 6.0094E-11	(17/95) -9.9990E+03 1.3139E-12	(12/75) -9.9990E+03 1.0820E-11
CRAGGLVY	5000	58/72 1.6882E+03 1.2088E-05	48/59 1.6882E+03 1.7804E-05	234/273 1.6882E+03 1.6079E-05	1067/1133 1.6882E+03 1.6213E-05	(6402/7468) 1.7944E+03 1.2574E+00
CURLY10	10000	[10000/10292] -1.0032E+06 8.9215E-02	[10000/10248] -1.0032E+06 1.6493E-01	[10000/10212] -1.0032E+06 2.6743E-01	[10000/10939] -1.0032E+06 1.0223E+00	[10000/10193] -1.0031E+06 2.6933E+03
CURLY20	10000	[10000/10261] -1.0032E+06 1.5313E-01	[10000/10232] -1.0032E+06 9.1182E+00	[10000/10222] -1.0032E+06 3.2401E+00	[10000/10984] -1.0032E+06 2.2065E+01	[10000/10206] -1.0026E+06 4.2706E+06
CURLY30	10000	[10000/10300] -1.0032E+06 1.9678E+00	[10000/10268] -1.0032E+06 6.7753E-01	[10000/10222] -1.0032E+06 2.1435E+01	[10000/10899] -1.0032E+06 3.4557E+01	[10000/10145] -1.0025E+06 8.8833E+05
DIXMAANA	3000	11/12 1.0000E+00 8.5386E-16	11/12 1.0000E+00 8.3519E-16	11/12 1.0000E+00 6.3154E-16	16/17 1.0000E+00 2.1321E-10	11/12 1.0000E+00 2.8120E-11
DIXMAANC	3000	11/12 1.0000E+00 1.8660E-10	11/12 1.0000E+00 1.8389E-10	11/12 1.0000E+00 1.3557E-10	26/28 1.0000E+00 2.5922E-10	63/68 1.0000E+00 2.4863E-09
DIXMAAND	3000	13/14 1.0000E+00 8.0456E-10	13/14 1.0000E+00 8.4189E-10	13/14 1.0000E+00 8.2228E-10	30/32 1.0000E+00 1.1670E-09	95/101 1.0000E+00 6.7614E-09
DIXMAANE	3000	251/257 1.0000E+00 2.3543E-10	105/108 1.0000E+00 9.2708E-11	104/106 1.0000E+00 1.2430E-10	2009/2148 1.0000E+00 2.4527E-10	662/692 1.0000E+00 2.2868E-10
DIXMAANF	3000	200/206 1.0000E+00 7.3165E-10	87/91 1.0000E+00 6.5717E-10	89/93 1.0000E+00 7.1144E-10	1401/1616 1.0000E+00 7.0336E-10	1321/1388 1.0000E+00 7.8050E-10

SIF name	n	scalar	original	modified	QC oldest	QC newest
DIXMAANG	3000	158/165 1.0000E+00 2.7413E-09	81/85 1.0000E+00 2.6758E-09	79/83 1.0000E+00 1.8293E-09	1112/1247 1.0000E+00 2.6706E-09	[10000/10420] 1.0000E+00 4.2022E-08
DIXMAANH	3000	148/153 1.0000E+00 1.1806E-08	77/79 1.0000E+00 6.2129E-09	74/76 1.0000E+00 9.7004E-09	1007/1123 1.0000E+00 1.1927E-08	493/521 1.0000E+00 1.0367E-08
DIXMAANI	3000	1846/1900 1.0000E+00 1.9095E-10	495/507 1.0000E+00 9.4287E-11	486/497 1.0000E+00 1.4882E-10	6267/6888 1.0000E+00 2.3202E-10	3611/3801 1.0000E+00 1.7551E-10
DIXMAANJ	3000	242/245 1.0000E+00 5.3689E-10	196/200 1.0000E+00 5.3509E-10	194/197 1.0000E+00 5.7723E-10	600/671 1.0000E+00 7.0704E-10	967/1003 1.0000E+00 6.5716E-10
DIXMAANK	3000	171/174 1.0000E+00 2.4910E-09	148/151 1.0000E+00 2.8306E-09	140/142 1.0000E+00 2.6429E-09	415/476 1.0000E+00 2.5665E-09	1025/1075 1.0000E+00 2.8105E-09
DIXMAANL	3000	117/121 1.0000E+00 1.1274E-08	110/112 1.0000E+00 1.1545E-08	107/110 1.0000E+00 8.5157E-09	292/314 1.0000E+00 9.9733E-09	725/759 1.0000E+00 1.1500E-08
DIXON3DQ	10000	[10000/10294] 6.7635E-04 3.2121E-08	[10000/10155] 7.5156E-04 1.4864E-07	[10000/10002] 4.6088E-04 5.2271E-07	[10000/11063] 1.3249E-03 3.1578E-07	[10000/10001] 4.1303E-04 9.3343E-07
DQDRTIC	5000	12/16 2.4705E-10 3.3465E-08	11/14 3.3184E-10 1.3220E-07	9/12 2.7321E-11 1.3073E-08	16/19 2.5251E-11 1.2182E-08	8/9 4.1047E-10 2.6999E-07
DQRTIC	5000	22/23 1.4042E+07 1.3099E+10	22/23 1.4041E+07 1.3100E+10	86/273 3.0200E+06 2.1901E+10	542/617 1.1090E+07 3.9297E+10	[10000/12409] 1.5201E+17 2.3810E+25
EDENSCH	2000	26/29 1.2003E+04 1.3713E-07	24/26 1.2003E+04 4.2129E-07	32/33 1.2003E+04 1.4996E-06	52/56 1.2003E+04 1.1716E-06	59/72 1.2003E+04 1.6571E-06
ENGVAL1	5000	17/21 5.5487E+03 1.5861E-08	17/18 5.5487E+03 4.1016E-10	(25/59) 5.5487E+03 1.7547E-08	24/25 5.5487E+03 1.0134E-08	42/49 5.5487E+03 1.1962E-08
EXTROSNB	10	788/993 4.8840E-12 1.3077E-09	1279/1476 7.6592E-13 7.5167E-10	957/1133 2.2676E-12 2.0007E-09	867/1070 9.9149E-13 1.2564E-09	856/1023 5.4103E-11 1.7868E-09
FLETCHV2	10000	[10000/10266] -5.0014E-01 6.7104E-11	[10000/10203] -5.0014E-01 4.4553E-10	[10000/10191] -5.0014E-01 1.0211E-09	(24/426) -5.0013E-01 2.1851E-08	[10000/10298] -5.0013E-01 3.4519E-07
FLETCHV3	10000	[10000/10058] -1.3062E+19 5.1364E+06	[10000/10063] -1.3997E+19 3.3275E+06	[10000/10436] -4.6882E+18 2.2562E+05		(77/175) -2.5509E+16 1.2714E+08
FLETCHBV	100	(387/575) -1.7810E+13 4.9975E-01	(260/361) -1.7805E+13 8.0860E-01	(429/632) -1.7807E+13 6.9032E-02	{8876/20001} -3.5104E+12 5.8472E+11	{8963/20001} -4.8759E+12 1.2966E+11
FLETCHCR	100	545/606 1.1125E-16 8.5539E-14	856/924 1.0906E-16 8.5868E-14	453/542 2.1599E-16 8.5864E-14	532/637 2.6810E-16 5.7822E-14	481/558 6.8109E-17 5.3872E-14
FMINSRF2	10000	(1833/1932) 1.0000E+00 3.7445E-15	(2568/2666) 1.0000E+00 3.9610E-15	(2595/2700) 1.0000E+00 3.9498E-15	[10000/11278] 2.6069E+00 6.9096E-02	(2882/3045) 1.0000E+00 4.7262E-15
FMINSURF	10000	(1134/1214) 1.0000E+00 1.4000E-15	(1160/1227) 1.0000E+00 2.3810E-14	(982/1050) 1.0000E+00 3.6659E-15	[10000/11131] 2.6568E+00 9.4708E-02	(1430/1530) 1.0000E+00 3.1296E-15
FREUROTH	5000	(20/52) 6.0816E+05 3.0718E-05	(20/53) 6.0816E+05 1.4952E-05	(115/195) 6.0816E+05 4.6934E-05	23/29 6.0816E+05 4.6634E-07	22/31 6.0816E+05 2.2403E-09
HILBERTA	10	36/42 9.9396E-10 3.4903E-14	36/40 6.1670E-10 3.3023E-14	45/51 1.9200E-09 3.8437E-14	164/175 2.5088E-09 3.9200E-14	199/227 7.5108E-10 1.4332E-14
HILBERTB	50	6/7 5.6013E-14 1.1392E-12	6/7 1.3191E-14 2.6512E-13	12/15 7.4476E-14 1.4976E-12	10/11 4.7355E-14 9.4841E-13	12/13 6.5274E-14 1.3129E-12

SIF name	n	scalar	original	modified	QC oldest	QC newest
LIARWHD	10000	21/25 3.5235E-10 3.4026E-05	20/23 8.4450E-09 9.1092E-05	20/24 1.2743E-10 6.7207E-06	24/28 1.0350E-13 1.6129E-08	19/21 2.6590E-08 1.9652E-04
MOREBV	5000	[10000/10296] 9.6926E-12 2.3442E-16	[10000/10111] 9.7241E-12 2.2806E-15	[10000/10123] 9.7267E-12 2.1684E-16	[10000/11055] 1.5786E-10 3.0476E-09	[10000/10273] 1.1068E-08 2.2361E-08
MSQRTALS	1024	4042/4148 2.2611E-10 2.2615E-11	3990/4086 1.7359E-10 1.9759E-11	4294/4376 3.4737E-10 1.9896E-11	[10000/11694] 5.3718E-07 2.6497E-08	9492/9926 1.2961E-10 2.1932E-11
MSQRTBLS	1024	2665/2730 6.6489E-11 2.3597E-11	2878/2942 8.8078E-11 2.4363E-11	2868/2928 1.2374E-10 1.5960E-11	[10000/11688] 5.7921E-10 1.2277E-10	7711/8070 4.0239E-11 2.2122E-11
NCB20B	2000	(1777/1965) 3.1395E+03 1.2267E-09	(1684/1816) 3.1395E+03 2.0638E-09	(1952/2023) 3.1395E+03 2.4129E-09	(4924/5460) 3.1402E+03 3.8728E+01	
NONCVXU2	10000	1699/1734 2.3174E+04 1.9234E-02	1883/1916 2.3170E+04 1.7476E-02	1767/1796 2.3173E+04 1.8308E-02	[10000/11516] 2.3465E+04 4.8358E+00	3862/4080 2.3172E+04 1.6505E-02
NONCVXUN	10000	2335/2380 2.3218E+04 1.5425E-02	2229/2264 2.3211E+04 2.2222E-02	2505/2530 2.3207E+04 2.0238E-02	4468/4961 2.3224E+04 1.7885E-02	8052/8360 2.3200E+04 2.0562E-02
NONDIA	10000	4/5 6.4454E-05 1.0371E-05	4/5 6.4454E-05 1.0371E-05	4/5 1.0054E-04 1.6168E-05	4/5 6.4447E-05 1.0370E-05	4/5 1.0059E-04 1.6176E-05
NONDQUAR	10000	259/285 1.2182E-04 2.9481E-07	112/114 9.1847E-05 1.4884E-07	95/96 1.1918E-04 3.3432E-07	3931/4439 1.6550E-04 3.0087E-07	671/748 2.1712E-04 2.7467E-07
NONMSQRT	1024	[10000/10365] 9.1507E+01 1.0126E+00	[10000/10328] 9.2219E+01 2.1253E-01	[10000/10348] 9.2002E+01 1.3019E+00	[10000/10731] 1.2062E+03 1.1982E+06	[10000/10164] 1.5214E+03 1.0496E+05
PENALTY2	100	27/28 9.7096E+04 3.2526E-04	25/26 9.7096E+04 3.4481E-04	34/41 9.7096E+04 2.7720E-04	(1035/1151) 9.7096E+04 1.3215E-03	
PENALTY3	100	(63/131) 9.9998E-04 8.1851E-04	(48/109) 9.9996E-04 8.6137E-05	(49/119) 9.9999E-04 4.3544E-04	(94/147) 1.0000E-03 5.9679E-04	(86/136) 9.9998E-04 7.0917E-05
POWELLSG	10000	39/45 6.0871E-07 2.7410E-09	38/44 1.6361E-06 1.0001E-08	35/43 3.0770E-07 1.0897E-08	49/55 3.2559E-08 1.9929E-08	34/42 6.9455E-06 9.9397E-08
POWER	1000	29/31 8.5190E+01 2.4148E+05	30/31 3.2787E+01 7.7755E+04	23/24 1.6203E+02 2.3140E+05	4134/4284 1.4950E+02 2.8599E+05	8907/9448 8.8269E+00 2.8595E+05
QUARTC	10000	22/23 4.4967E+08 1.6785E+12	22/23 4.4966E+08 1.6785E+12	89/295 1.0380E+08 3.6727E+12	599/694 3.8978E+08 4.9654E+12	[10000/12105] 5.7894E+18 3.5030E+27
SBRVBND	5000	[10000/10121] 2.4243E+04 2.0310E+11	7317/7363 4.5252E-07 1.3049E+01	226/253 1.4050E-09 1.3036E+01		
SCHMVE'ET	10000	(45/100) -2.9994E+04 3.1617E-10	(34/100) -2.9994E+04 8.9321E-11	(81/146) -2.9994E+04 7.5106E-10	(4140/7778) -2.9994E+04 9.4161E-09	(2910/3370) -2.9994E+04 7.0637E-10
SCOSINE	10000	[10000/10274] 2.4331E+03 2.3424E+14	[10000/10150] -4.3122E+03 2.3871E+12	(108/237) -9.9713E+03 1.1581E+04		
SINQUAD	10000	368/483 3.7518E-16 3.8686E-16	476/613 9.6199E-18 8.4638E-16	142/176 6.5111E-20 1.2555E-16	483/655 6.5060E-17 5.6544E-16	103/142 2.8468E-18 1.0540E-15
SPARSINE	10000	1526/1580 2.3449E-04 1.4560E-02	1318/1345 7.2807E-05 1.5263E-02	1750/1789 6.5213E-04 1.3308E-02	[10000/11956] 9.3458E-06 9.1590E-02	
SPARSQR	10000	22/23 3.2842E-04 1.1303E-04	22/23 3.2840E-04 1.1304E-04	22/23 3.0842E-04 1.9003E-04	923/1122 4.0113E-04 3.1461E-04	6550/7363 2.1253E-04 3.4177E-04
SPMSR'TLS	10000	271/278 2.4655E-12 2.1664E-12	181/188 1.1484E-12 1.2519E-12	169/174 2.8334E-12 2.2639E-12	[10000/11160] 3.1165E+01 7.8428E+02	[10000/10299] 1.3423E+01 5.8091E+02

SIF name	n	scalar	original	modified	QC oldest	QC newest
SROSENBR	10000	17/18	17/18	18/19	20/22	13/15
		9.7462E-13	9.6945E-13	2.8877E-11	7.7844E-15	1.9032E-09
		7.3433E-10	7.3099E-10	5.4363E-09	9.5442E-12	4.2689E-09
TESTQUAD	1000	1185/1215	1014/1040	53/60	1138/1251	(9135/10505)
		2.3739E-03	1.1351E-05	1.8453E-07	6.8783E-03	2.6694E+06
		8.5825E-02	7.5109E-02	8.5722E-02	3.9191E-02	3.3720E+12
TOINTGSS	10000	19/26	21/33	22/30	15/18	17/19
		1.0001E+01	1.0001E+01	1.0001E+01	1.0001E+01	1.0001E+01
		3.5634E-11	7.2437E-12	8.3348E-12	5.6912E-15	6.5929E-12
TQUARTIC	10000	26/34	25/30	15/26	17/24	15/22
		1.2326E-30	6.4193E-23	0.0000E+00	1.0821E-23	9.3894E-28
		4.9304E-30	8.9399E-18	0.0000E+00	8.2180E-20	3.7558E-27
TRIDIA	10000	1578/1615	369/379	156/159	[10000/11020]	[10000/10265]
		4.5482E-07	4.9007E-08	5.4159E-09	9.3229E-01	8.0275E+04
		2.9146E-04	2.5147E-04	2.1295E-04	2.5196E+05	7.8197E+09
VARDIM	100	22/23	22/23	22/23	22/23	22/23
		2.3262E+03	2.3262E+03	2.3262E+03	2.3262E+03	2.3262E+03
		6.0116E+11	6.0116E+11	6.0116E+11	6.0116E+11	6.0116E+11
VAREIGVL	5000	23/25	119/123	83/85	543/594	218/235
		9.8324E-11	4.0310E-08	4.4423E-08	5.2682E-08	6.3107E-08
		5.8094E-09	1.7887E-08	2.1695E-08	2.3273E-08	2.2266E-08
WATSON	31	619/683	788/911	690/786		
		1.4633E-08	9.4057E-09	4.4895E-09		
		1.6175E-12	7.7828E-13	1.2245E-12		
WOODS	10000	91/116	87/112	82/112	38/41	82/104
		1.1613E-07	1.8284E-08	1.4632E-08	4.8820E-09	1.9805E-07
		4.7263E-06	3.3315E-05	1.0006E-05	3.7856E-06	8.1126E-05



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399