NUMERICAL LINEAR ALGEBRA WITH APPLICATIONS
Numer. Linear Algebra Appl. 2004; **11**:525–533 (DOI: 10.1002/nla.359)

Truncated Newton type solver with application to grid untangling problem

V. Garanzha, I. Kaporin*,† and I. Konshin

Center for Supercomputer and Massively Parallel Applications, Computing Center of Russian Academy of Sciences, Vavilova 40, Moscow 119991, Russia

SUMMARY

A practical stopping criterion for inner conjugate gradient iterations in the truncated Newton type unconstrained optimization method is proposed. Numerical results are given for a family of large-scale geometrical optimization problems. The comparison with the standard residual norm stopping rule demonstrates an essential gain in efficiency. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: unconstrained optimization; truncated Newton method; conjugate gradients; parallel preconditioning; grid untangling

1. INTRODUCTION

In this paper, we will construct a quasi-optimum termination rule for inner linear preconditioned conjugate gradient (PCG) iterations [1] used within a truncated Newton non-linear solver [2].

The problem of non-linear minimization is formulated as follows:

$$u_* = \arg\min_{u \in R^n} \varphi(u)$$

where the functional φ is assumed to be sufficiently smooth.

Such settings may arise, for instance, in variational formulations of mesh shape-quality improvement problems. In Reference [3], several families of optimization methods are compared on a rather simple mesh improvement problem, and the conclusion is drawn that the truncated Newton type method [2] far outperforms its competitors. In an earlier paper [4], the

*Correspondence to: I. E. Kaporin, Computing Center of Russian Academy of Sciences, Vavilova 40, Moscow 119991, Russia.
†E-mail: kaporin@ccas.ru

Contract/grant sponsor: Netherlands Technology Foundation (STW); contract/grant number: NNS.4683
Contract/grant sponsor: National Science Foundation (NWO); contract/grant number: 047.008.007
Contract/grant sponsor: Program for Fundamental Research of Russian Academy of Sciences (RAS); contract/grant numbers: 03, 16, and 17

Received 9 January 2003
Revised 12 May 2003
Accepted 10 June 2003

Published online 5 May 2004

truncated Newton type method combined with continuation techniques was successively used for a much harder grid untangling problem.

In such methods, an approximation to the solution u_* is found using the Newton type iterations

$$u_{k+1} = u_k + d_k, \quad k = 0, 1, \ldots$$

where d_k is a properly scaled approximation to the Newton direction $-(\varphi'')^{-1}\varphi'$.

For a certain solution error measure ε_k (the most usual choices are $\varepsilon_k = \varphi(u_k) - \varphi(u_*)$ or $\varepsilon_k = \|\varphi'(u_k)\|$) one typically has, cf. Reference [5],

$$\varepsilon_{k+1} \leqslant (1 - \tau\vartheta_k(s_k))\varepsilon_k \tag{1}$$

where

- $\tau \in (0, 1]$ characterizes the problem non-linearity ($\tau = 1$ if the problem is linear), and
- $\vartheta_k(s_k) \in (0, 1]$ is a special error measure for the solution x of an appropriate linear equation (e.g. $\varphi''(u_k)x = -\varphi'(u_k)$) using s_k iterations of the PCG method ($\vartheta_k = 1$ if the PCG method is converged).

We will present a simple practical rule for choosing the inner iteration numbers s_k, which does not assume the knowledge of τ and makes it possible to avoid making redundant PCG iterations within each non-linear step.

2. GENERAL IDEA UNDERLYING THE STOPPING RULE

Let us briefly describe the algorithmic construction presented in the paper.

Assume that the criterion for the outer (non-linear) iterations to be converged is

$$\varepsilon_m \leqslant \varepsilon\ \varepsilon_0, \quad 0 < \varepsilon \ll 1$$

Then the non-linear error estimate (1) yields the following sufficient condition for the convergence:

$$\sum_{k=0}^{m-1} \vartheta_k(s_k) \approx \tau^{-1} \log(\varepsilon^{-1})$$

Let then $\mathscr{P}$ and $\mathscr{I}$ be the computational costs of one outer iteration (including the PCG initialization cost) and the computational cost per one typical PCG iteration, respectively. Therefore, the total computational cost can be estimated as

$$\begin{aligned}
&\sum_{k=0}^{m-1} (\mathscr{P} + \mathscr{I} s_k) \\
&\quad \approx \left(\sum_{k=0}^{m-1} (\mathscr{P} + \mathscr{I} s_k) \Big/ \sum_{k=0}^{m-1} \vartheta_k(s_k) \right) \tau^{-1} \log(\varepsilon^{-1}) \\
&\quad \leqslant \left(\max_{0 \leqslant k \leqslant m-1} (\mathscr{P} + \mathscr{I} s_k)/\vartheta_k(s_k) \right) \tau^{-1} \log(\varepsilon^{-1})
\end{aligned}$$

Hence, the stopping criterion can now be formulated in terms of choosing the inner iteration numbers $s = s_k$ providing for a reasonably small value of each ratio

$$\psi_k(s) = (\mathscr{P} + \mathscr{I}s)/\vartheta_k(s)$$

Since the quantity $\vartheta_k(s)$ can be readily evaluated at any sth inner iteration (see Reference [1] and Section 4 below) one can stop at $s = s_k$ where the first local minimum of $\psi_k(s)$ is attained. In other words, as soon as the condition

$$\psi_k(s) \geqslant \psi_k(s-1)$$

has satisfied, one sets $s_k = s$ and quits the inner iterations.

3. THE UNCONSTRAINED MINIMIZATION PROBLEM

Let a differentiable functional

$$\varphi(u) : R^n \to R^1$$

be bounded from below, have gradient $g(u) \in R^n$, and

$$|\varphi(u+h) - \varphi(u) - h^{\mathrm{T}} g(u)| \leqslant \frac{\gamma}{2} h^{\mathrm{T}} A h, \quad \gamma \geqslant 1 \tag{2}$$

for certain symmetric positive definite $n \times n$ matrix $A = A(u)$ and all sufficiently small $\|h\|$. (The value of γ is not used in the actual calculations.)

Let $\|v\|_*$ be any vector norm appropriate for measuring $g(u)$. At the level of outer iterations, the method can be described by the following pseudo-code.

Step 1: Compute $g = g(u)$ and check the convergence, i.e. if

$$\|g\|_* \leqslant \varepsilon_0$$

then quit,

Step 2: Compute the matrix $A = A(u)$ and find $x \approx A^{-1}g$ such that the scaling condition

$$-x^{\mathrm{T}} g = x^{\mathrm{T}} A x \tag{3}$$

holds and $x^{\mathrm{T}} A x$ is sufficiently large,

Step 3: Set $\alpha = 1$ and

Check: **IF** $\varphi(u + \alpha x) \leqslant \varphi(u) - (\alpha/2) x^{\mathrm{T}} A x$
THEN $u := u + \alpha x$ and go to Step 1,
ELSE $\alpha = \alpha/2$ and repeat Check.

Note that (3) yields

$$x^{\mathrm{T}} A x = -(A^{1/2}x)^{\mathrm{T}}(A^{-1/2}g) \leqslant \sqrt{x^{\mathrm{T}} A x}\sqrt{g^{\mathrm{T}} A^{-1} g}$$

and therefore the quantity

$$\vartheta = x^{\mathrm{T}} A x / g^{\mathrm{T}} A^{-1} \ g \leqslant 1$$

is chosen to measure how close is x to $(-A^{-1}g)$.

Let x at Step 2 be obtained after s iterations of the (preconditioned) conjugate gradients with zero initial guess applied to the linear system $Ax = -g$. Then the required scaling condition (3) holds by $b = -g$ in the PCG method (see Section 4 below).

Theorem 1
The following estimate for the reduction in the functional value attained at

$$u_+ = u + \alpha x$$

by the descent along the direction x using stepsize α holds:

$$\varphi(u_+) \leqslant \varphi(u) - \frac{\vartheta}{4\gamma} g^{\mathrm{T}} A^{-1} g$$

with the above defined ϑ and γ.

Proof
Indeed, for any $0 < \tau < 2/\gamma$ one has by (2) and (3) the following estimate for the decrease in the functional φ along the direction x:

$$\begin{aligned} \varphi(u + \tau x) &= \varphi(u) + \tau x^{\mathrm{T}} g + (\varphi(u + \tau x) - \varphi(u) - \tau x^{\mathrm{T}} g) \\ &\leqslant \varphi(u) - \left(\tau - \frac{\gamma}{2}\tau^2\right) x^{\mathrm{T}} A x \end{aligned} \tag{4}$$

Next, we consider the following two cases.
Case 1: If one quits with $\alpha = 1$ at Step 3, it follows that

$$\varphi(u_+) \leqslant \varphi(u) - \frac{1}{2} x^{\mathrm{T}} A x \leqslant \varphi(u) - \frac{1}{4\gamma} x^{\mathrm{T}} A x$$

where the latter inequality holds by $\gamma \geqslant 1$.
Case 2: When an actual bisection of the stepsize takes place, one has at Step 3 using (4) with $\tau = 2\alpha$

$$\varphi(u) - \alpha x^{\mathrm{T}} A x < \varphi(u + 2\alpha x) \leqslant \varphi(u) - (2\alpha - 2\gamma\alpha^2) x^{\mathrm{T}} A x$$

which easily yields

$$\alpha > \frac{1}{2\gamma}$$

Since we also have

$$\varphi(u + \alpha x) \leqslant \varphi(u) - \frac{\alpha}{2} x^{\mathrm{T}} A x$$

the required estimate

$$\varphi(u + \alpha x) \leqslant \varphi(u) - \frac{1}{4\gamma} x^{\mathrm{T}} A x = \varphi(u) - \frac{\vartheta}{4\gamma} g^{\mathrm{T}} A^{-1} g$$

follows. □

Corollary 1
If ϑ is bounded below and γ is bounded above, then

$$\lim_{i\to\infty} \|g(u_i)\|_{A(u_i)^{-1}} = 0$$

where i is the outer iteration number and $\|v\|_A = \sqrt{v^{\mathrm{T}} A v}$.

Corollary 2
Let u_* be the minimizer of $\varphi(u)$ and assume that

$$g^{\mathrm{T}} A^{-1} g/(4\gamma(\varphi(u) - \varphi(u_*))) \geqslant \tau > 0$$

for all outer iterations. Then one has

$$\varphi(u_+) - \varphi(u_*) \leqslant (1 - \tau\vartheta)(\varphi(u) - \varphi(u_*))$$

which is exactly the convergence estimate (1).

Note that the latter readily follows from the result of Theorem 1 rewritten as

$$\varphi(u_+) - \varphi(u_*) \leqslant \left(1 - \frac{g^{\mathrm{T}} A^{-1} g/(4\gamma)}{\varphi(u) - \varphi(u_*)}\,\vartheta\right)(\varphi(u) - \varphi(u_*))$$

Therefore, in order to provide the quasi-optimum decrease in $\varphi(u) - \varphi(u_*)$ at each outer iteration, one can use the stopping criterion for inner PCG iterations as described in Section 2. In view of the explicit expression for ϑ obtained in the next Section, this stopping criterion can be written as follows:

$$\text{Iterate while} \quad s \geqslant -(\mathscr{P}/\mathscr{I}) + (\omega_0 + \cdots + \omega_s)/\omega_s \quad \text{holds true}$$

Here $s = s_k$ is the PCG iteration number, ω_s are the scalars calculated from the scalar products involved in the PCG recurrences, and (as defined above) $\mathscr{P}/\mathscr{I}$ is the ratio of the computational costs implied by evaluation of $\varphi(u), g(u), A(u)$, and the preconditioner C to the computational cost of one inner PCG iteration. Note that ω_s typically decreases as s grows while their sum remains bounded, and therefore the PCG iterations tend to terminate rather early.

4. THE CONJUGATE GRADIENT ITERATIONS

In our case, we use the PCG iterations for the solution of the problem $Ax = b$ with zero initial guess:

$$x_0 = 0$$

$$r_0 = b$$

$$p_0 = Cr_0$$

for $i = 0, 1, \ldots$

$$\alpha_i = r_i^{\mathrm{T}} C r_i / p_i^{\mathrm{T}} A p_i$$

$$\omega_i = r_i^{\mathrm{T}} C r_i \alpha_i$$

$$x_{i+1} = x_i + p_i \alpha_i$$

$$r_{i+1} = r_i - A p_i \alpha_i$$

$$\beta_i = r_{i+1}^{\mathrm{T}} C r_{i+1} / r_i^{\mathrm{T}} C r_i$$

$$p_{i+1} = C r_{i+1} + p_i \beta_i$$

Here C is a properly chosen SPD preconditioning matrix, which should approximate, in some sense, the matrix A^{-1}. The choice of the matrix C is subject to the requirement that the vector $w = Cr$ be easily calculated for any r. For instance, one of the best choices is the approximate Cholesky preconditioning, where $C = (U^{\mathrm{T}}U)^{-1}$ and $U^{\mathrm{T}}U \approx A$ with the upper triangular matrix U being much sparser than the exact Cholesky factor of A, cf. Reference [6] and references therein.

If x_s is obtained after s iterations of the above preconditioned conjugate gradient method with zero initial guess applied to the linear system $Ax = -g$, then

$$x_s \in K_s = \mathrm{span}\{Cg, CACg, \ldots, (CA)^{s-1} Cg\}$$

and

$$x_s = \arg\min_{x \in K_s} (g + Ax)^{\mathrm{T}} A^{-1} (g + Ax)$$

Since $\alpha x_s \in K_s$ for any scalar α, one gets

$$(g + \alpha A x_s)^{\mathrm{T}} A^{-1} (g + \alpha A x_s) \geqslant (g + A x_s)^{\mathrm{T}} A^{-1} (g + A x_s)$$

which readily gives, with $\alpha = -x_s^t g / x_s^{\mathrm{T}} A x_s$, the inequality

$$0 \geqslant (x_s^{\mathrm{T}} g + x_s^{\mathrm{T}} A x_s)^2 / x_s^{\mathrm{T}} A x_s$$

Hence one can obtain the required scaling condition (3) for $x = x_s$.

Using $x_0 = 0$ in the above method, one can also show (see Reference [1] for the proof) that the following explicit expression for A-norm of each PCG iterate x_i holds:

$$x_s^{\mathrm{T}} A x_s = \sum_{i=0}^{s-1} \omega_i$$

Therefore, the above-mentioned quantity $\vartheta_s = x_s^{\mathrm{T}} A x_s / g^{\mathrm{T}} A^{-1} g$ can be presented as

$$\vartheta_s = (g^{\mathrm{T}} A^{-1} g)^{-1} \left(\sum_{i=0}^{s-1} \omega_i \right), \quad s = 0, 1, \ldots, n-1$$

Using the latter identity and following the approach outlined above in Section 2 one can readily implement the new PCG stopping criterion.

5. PARALLEL PRECONDITIONING

For parallel computations, the preconditioner C was constructed using the block incomplete inverse Cholesky (BIIC) [7] and the second order incomplete Cholesky (IC2) [6] techniques.

Let A be reordered and split into $s \times s$ block form. For the tth diagonal block, each having the dimension n_t, the 'basic' index set $\{k_{t-1}+1,\ldots,k_t\}$ is defined, $(k_{t-1}=n_1+\cdots+n_{t-1}, k_0=0, k_s=n)$, and 'overlapping' index sets are constructed as $\{j_t(1),\ldots,j_t(m_t-n_t)\}$, $j_t(p)\leqslant k_{t-1}$.

The BIIC-IC2 preconditioner C is

$$C=C(\zeta)=\sum_{t=1}^{s} V_t U_t^{-1}(\zeta)\begin{bmatrix}0 & 0\\ 0 & I_{n_t}\end{bmatrix}U_t^{-T}(\zeta)V_t^{\mathrm{T}}$$

where V_t are rectangular matrices composed of unit n-vectors e_j as follows:

$$V_t=[e_{j_t(1)}|\cdots|e_{j_t(m_t-n_t)}|e_{k_{t-1}+1}|\cdots|e_{k_t}],\quad t=1,\ldots,s$$

and each upper triangular matrix $U_t(\zeta)$ is an incomplete Cholesky factor for the tth $m_t \times m_t$ submatrix $V_t^{\mathrm{T}}AV_t$:

$$V_t^{\mathrm{T}}AV_t=U_t^{\mathrm{T}}(\zeta)U_t(\zeta)+U_t^{\mathrm{T}}(\zeta)R_t(\zeta)+R_t^{\mathrm{T}}(\zeta)U_t(\zeta)$$

Here $R_t(\zeta)$ is a strictly upper triangular error matrix.

For each t, the 'overlapping' index set typically includes indices not greater than k_t and the most 'essentially' connected to the basic index set, e.g. in the sense of the sparse matrix graph adjacency relations. Here $m_t \geqslant n_t$ and, obviously, $m_1=n_1$, i.e. at least the first overlapping set is empty.

Here $0<\zeta\ll 1$ is the drop tolerance parameter which determines the quality of the incomplete factorization. The existence and correctness of such IC2 decompositions is guaranteed for any SPD matrix.

The recurrences for the calculation of IC2 factorization can easily be obtained from the above relation, especially in the case in which the sparsity patterns of U and R do not have coinciding non-zero positions and their non-zero elements are subject to the conditions $|U_{i,j}|\geqslant\zeta$ and $|R_{i,j}|<\zeta$, respectively, $i<j$.

The above described mathematical technique was implemented in a portable software with the use of the message passing interface (MPI) library for communications between processes. The special structure of the above described BIIC-IC2 preconditioning made it possible to run the PCG iterations very efficiently even on workstation clusters, see Reference [8].

6. TEST PROBLEM: GRID UNTANGLING

Rather hard-to-solve minimization problems arising in global untangling of computational grids using the continuation technique, were successfully solved in Reference [4] by an application of the inexact Newton-like minimization procedure.

Table I. Standard vs new inner PCG stopping rules.

Criterion	32×32	64×64	128×128	256×256
$\varepsilon = 1.\mathrm{d} - 08$	13.74	89.31	670.3	∞
$\varepsilon = 1.\mathrm{d} - 04$	16.86	60.30	400.1	4411
$\varepsilon = 1.\mathrm{d} - 02$	12.19	46.25	314.4	4415
$\varepsilon = 1.\mathrm{d} - 01$	12.57	42.88	290.4	3508
$\varepsilon = 3.\mathrm{d} - 01$	12.36	56.84	367.4	4702
$P/I = 5$	12.57	56.19	309.6	1861
$P/I = 10$	16.09	51.46	287.7	1937
$P/I = 20$	11.37	46.47	289.5	1902
$P/I = 40$	15.31	48.66	265.9	1804
$P/I = 70$	9.83	41.30	291.7	2008

Let us consider the following test problem: given a 2D region Ω, find a pair of functions $\{x(\xi,\eta), y(\xi,\eta)\}$ which performs the mapping

$$\mathscr{D} = (\xi_1, \xi_2) \times (\eta_1, \eta_2) \to \Omega$$

with prescribed mapping of the boundaries $\partial\mathscr{D} \to \partial\Omega$.

According to Reference [4], the solution is sought as

$$\{x(\xi,\eta), y(\xi,\eta)\} = \arg\min_{\{x,y\}} \int_{\mathscr{D}} (x_\xi^2 + x_\eta^2 + y_\xi^2 + y_\eta^2) \frac{2\,\mathrm{d}\xi\,\mathrm{d}\eta}{J + \sqrt{J^2 + \mu^2}}$$

where

$$J = x_\xi y_\eta - x_\eta y_\xi$$

is the Jacobian of the mapping and $0 < \mu \ll 1$ the small parameter used in the path-following procedure.

The above functional is then discretized using quadrilateral grid over $\mathscr{D}$ and bilinear finite elements for approximating the functions $x(\xi,\eta)$ and $y(\xi,\eta)$ and using the 2D trapezoidal quadrature rule to approximate the integrals over each cell, exactly as in Reference [4]. This yields a discrete minimization problem for some smooth functional $\varphi(u)$ with u representing the coordinates of the unknown grid.

The above test case was considered for an S-shaped region (a non-convex octagon) with the vertices

$$(0.0, 0.0),\quad (0.0, 1.0),\quad (0.6, 0.4),\quad (0.6, 1.4),\quad (1.6, 1.4),$$
$$(1.6, 0.4),\quad (1.0, 1.0),\quad (1.0, 0.0)$$

(as shown in Figure 7 in Reference [4]) and the grid sizes $32 \times 32, 64 \times 64, 128 \times 128, 256 \times 256$ and run on a Pentium PC (433 MHz, 128 Mb RAM). Table I presents the timing results (the total wall-clock time in seconds) obtained for the (standard) inner PCG stopping criterion Reference [2] with $\varepsilon = 10^{-8}, 10^{-4}, 10^{-2}, 0.1, 0.3$ and for the new criterion (see also Reference [1]) with $\mathscr{P}/\mathscr{I} = 5, 10, 20, 40, 70$. (The parameters of outer iterations were the same in all cases, and the same rule [4] was used for the reduction of the path-following parameter μ.)

The results presented in the Table I clearly indicate the advantage of the new rule for stopping CG iterations, especially for the largest test problem. In the latter case, the solution of linear problems with too high precision $\varepsilon = 10^{-8}$ have resulted in a stagnation of the non-linear iterations.

7. CONCLUSION

For the truncated Newton type unconstrained optimization method, the criterion for stopping the inner PCG iterations is formulated. It is independent of any characterization of non-linearity and is expressed in terms of scalar coefficients involved in PCG recursions. The stopping criterion also (slightly) depends on the ratio of the PCG startup cost to the regular PCG iteration cost. Our strategy is aimed towards a sufficient increase of $x^{\mathrm{T}}Ax$ rather than (standard) residual norm $\|g + Ax\|$ reduction. (Recall that g is the gradient and x is the approximate Newton direction.) It is worth noting here that the relative residual $\|g+Ax\|/\|g\|$ used for the construction of a similar truncated Newton method in Reference [2] is rather loosely related to the above quantity $\vartheta = x^{\mathrm{T}}Ax/g^{\mathrm{T}}A^{-1}g$ if the matrix A is not well-conditioned.

Thus, the proposed non-linear optimization solver:

- is well adapted to large-scale problems with sparse Jacobian,
- has theoretical justification for nontrivial classes of problems,
- demonstrates good performance for some hard-to-solve test problems,
- can be efficiently parallelized with the proposed techniques.

REFERENCES

1. Axelsson O, Kaporin I. Error norm estimation and stopping criteria in preconditioned conjugate gradient iterations. *Numerical Linear Algebra with Applications* 2001; **8**:265–286.
2. Dembo R, Steihaug T. Truncated Newton algorithms for large-scale unconstrained optimization. *Mathematical Programming* 1983; **26**:190–212.
3. Freitag L, Knupp P, Munson T, Shontz S. A comparison of optimization software for mesh shape-quality improvement problems. *Eleventh International Meshing Roundtable on Mesh Quality/Smoothing/Optimization*, Ithaca, New York, 15–18 September 2002.
4. Garanzha VA, Kaporin IE. Regularization of the barrier variational method of grid generation. *Computational Mathematics and Mathematical Physics* 1999; **39**:1426–1440.
5. Kaporin I, Axelsson O. On a class of non-linear equation solvers based on the residual norm reduction over a sequence of affine subspaces. *SIAM Journal on Scientific Computing* 1994; **16**:228–249.
6. Kaporin IE. High quality preconditioning of a general symmetric positive definite matrix based on its $U^{\mathrm{T}}U + U^{\mathrm{T}}R + R^{\mathrm{T}}U$-decomposition. *Numerical Linear Algebra with Applications* 1998; **5**:483–509.
7. Kaporin IE. New convergence results and preconditioning strategies for the conjugate gradient method. *Numerical Linear Algebra with Applications* 1994; **1**:179–210.
8. Kaporin IE, Konshin IN. A parallel block overlap preconditioning with inexact submatrix inversion for linear elasticity problems. *Numerical Linear Algebra with Applications* 2002; **9**:141–162.

