

A New Three-Term Conjugate Gradient Method

L. C. W. DIXON,¹ P. G. DUCKSBURY,² AND P. SINGH³

Abstract. In this paper, we describe an implementation and give performance results for a conjugate gradient algorithm for unconstrained optimization. The algorithm is based upon the Nazareth three-term formula and incorporates Allwright preconditioning matrices and restart tests. The performance results for this combination compare favorably with existing codes.

Key Words. Unconstrained optimization, conjugate gradients.

1. Introduction

The conjugate gradient approach was introduced as an algorithm for solving large systems of nonlinear equations by Hestenes and Stiefel (Ref. 1) and adapted for nonlinear optimization by Fletcher and Reeves (Ref. 2). For large-dimensional problems, it has the advantage over both modified Newton and variable metric methods of not requiring a matrix store and is therefore usually used for large-dimensional problems.

In this paper, we will first review the properties of conjugate gradient algorithms on large-dimensional problems and discuss the recent developments that have been suggested for improving the performance of such algorithms. We will then propose a new algorithm that includes these modifications and present numerical results obtained with the new algorithm, which indicate that it performs better than a selection of other algorithms on a number of test functions.

¹ Reader in Numerical Computation, Numerical Optimisation Centre, Hatfield Polytechnic, Hatfield, England.

² Research Student, School of Information Sciences, Hatfield Polytechnic, Hatfield, England. The support of the Science and Engineering Research Council is gratefully acknowledged.

³ Research Associate, Numerical Optimisation Centre, Hatfield Polytechnic, Hatfield, England.

2. Preliminaries

2.1. On Quadratic Functions

Definition 2.1. A set of directions $t^{(i)}$, $i = 1, \dots, k$, $k \leq n$, is said to be conjugate with respect to a matrix G and the associated quadratic function

$$f = a + \frac{1}{2}x^T Gx \tag{1}$$

if they possess the property

$$t^{(i)T} G t^{(j)} = 0, \quad i \neq j. \tag{2}$$

If we have a set of n independent conjugate directions, these form a base, and we may express any point x as a linear combination,

$$x = \sum_i \beta_i t^{(i)}. \tag{3}$$

So, if x^* represents the minimum point,

$$f(x) - f(x^*) = \frac{1}{2} \sum_i \sum_j \beta_i \beta_j t^{(i)T} G t^{(j)} - \beta_i^* \beta_j^* t^{(i)T} G t^{(j)}; \tag{4}$$

and, using Eq. (2),

$$f(x) - f(x^*) = \frac{1}{2} \sum_i (\beta_i^2 - \beta_i^{*2}) t^{(i)T} G t^{(i)}. \tag{5}$$

Theorem 2.1. Quadratic Termination. If we minimize $f(x)$ by searching along each conjugate direction in turn, then each search is independent of the other, and we must set $\beta_i = \beta_i^*$. The minimum of a quadratic function is therefore obtained after searching along n mutually conjugate directions.

This is, of course, the basis of the conjugate gradient approach; but, if this were its only property, the method would be hopelessly inefficient on most large problems.

In conjugate gradient algorithms, the sequence of directions $t^{(k)}$ and a sequence of points $x^{(k)}$ are generated iteratively,

$$x^{(k+1)} = x^{(k)} + \alpha t^{(k)}, \tag{6}$$

where the step α is chosen so that, if the function is the quadratic (1), then $x^{(k+1)}$ would be the minimum along the line (6). This implies that $x^{(k+1)}$ is then the minimum of the quadratic function in the space spanned by $[t^{(1)}, \dots, t^{(k)}]$; and, if the initial direction $t^{(1)} = g^{(1)}$ and the classic iterative method for generating $t^{(k)}$ is used, this space is the same as that spanned by $[g^{(1)}, Gg^{(1)}, \dots, G^{k-1}g^{(1)}]$. From this, two well-known theorems have been derived which account for the practical usefulness of the method on large systems.

Theorem 2.2. The number of iterations to obtain the minimum of (1) is bounded above by the number of independent vectors in the set $(g^{(1)}, Gg^{(1)}, \dots, G^{n-1}g^{(1)})$.

Theorem 2.3. The number of iterations to obtain the minimum of (1) is bounded above by the number of distinct eigenvalues of G .

2.2. On Nonquadratic Functions. When the function to be optimized is nonquadratic, the above framework can still be used to generate an efficient algorithm. Many codes have been written for this purpose, beginning with that described in Fletcher and Reeves (Ref. 2). In proposing a new conjugate gradient algorithm, we wished to ensure that its convergence pattern is based on Theorems 2.2 and 2.3 and not Theorem 2.1 above.

The structure of an early algorithm based on Theorem 2.1 consisted of seven basic steps.

Algorithm 2.1

- Step 1. Select $x^{(1)}$, ϵ_0 , j_{\max} . Set $k=1$, $j=1$. Calculate $f^{(1)}$ and $g^{(1)}$.
- Step 2. If $\|g^{(k)}\| < \epsilon_0$, stop. If $j \geq j_{\max}$, stop. If $\|x^{(k)} - x^{(k-n)}\| \leq \epsilon_0$, stop.
- Step 3. If $k=1$, set $p^{(k)} = -g^{(k)}$; otherwise, set $p^{(k)} = -g^{(k)} + bp^{(k-1)}$.
- Step 4. Perform an almost perfect line search along $p^{(k)}$.
- Step 5. Calculate $g^{(k+1)}$, b and set $j=j+1$, $k=k+1$.
- Step 6. If $k=n$, set $k=1$, to start a new cycle.
- Step 7. Go to Step 2.

In the above algorithm, Step 3 generates the set of conjugate directions on (1), provided the line search undertaken in Step 4 is perfect. Insofar as conjugacy is only defined for quadratic functions, the line search in Step 4 ought to contain a parabolic interpolation to guarantee that it is exact on quadratic functions.

To obtain a downhill direction, we need $p^{(k)T}g^{(k)} < 0$, where

$$p^{(k)T}g^{(k)} = -g^{(k)T}g^{(k)} + bp^{(k-1)T}g^{(k)}.$$

The last term is, of course, zero for a perfect line search, but otherwise this condition implies that the search must be sufficiently accurate for

$$p^{(k-1)T}g^{(k)} < g^{(k)T}g^{(k)}/b, \quad (7)$$

which imposes one limit on the accuracy needed in the line search.

The form of b , obtained immediately in Step 3 by insisting that $p^{(k)}$ is conjugate to $p^{(k-1)}$, namely,

$$b = \frac{g^{(k)T} y^{(k-1)}}{p^{(k-1)T} y^{(k-1)}}, \quad \text{where } y^{(k-1)} = g^{(k)} - g^{(k-1)}, \quad (8)$$

can be simplified using the conjugacy and orthogonality properties, and alternative formulas are often used.

By the definition of $y^{(k-1)}$,

$$b = \frac{g^{(k)T} (g^{(k)} - g^{(k-1)})}{p^{(k-1)T} (g^{(k)} - g^{(k-1)})}. \quad (9)$$

Using the perfect line search condition twice then gives the Polak-Ribière formula (Ref. 3)

$$b = \frac{g^{(k)T} (g^{(k)} - g^{(k-1)})}{g^{(k-1)T} g^{(k-1)}}; \quad (10)$$

and, applying the orthogonality property, then gives the Fletcher-Reeves formula

$$b = \frac{g^{(k)T} g^{(k)}}{g^{(k-1)T} g^{(k-1)}}. \quad (11)$$

Although both formulas are based on the assumption of a perfect line search, they are frequently used in conjunction with imperfect searches. There is little conclusive numerical evidence on which of the alternative formulas for b is preferable; but, in Ref. 4, Powell argues in favor of the Polak-Ribière form (10). He shows that, if the conjugacy property $g^{(k)T} g^{(k-1)} = 0$ has already been lost, then the Polak-Ribière form bounds the angle between $p^{(k)}$ and $g^{(k)}$ and can produce a much more downhill direction.

Finite termination of the above type of algorithm follows from Wolfe's theorem (Refs. 5 and 6), provided the line search in Step 4 is such that Conditions II and III are satisfied. Condition I is automatically satisfied on the regular subset of directions with $k=1$, as the search is then along the steepest descent direction.

Cohen (Ref. 7) first proved that the above class of algorithms is superlinearly convergent and Dixon (Ref. 8) gave one proof that there is a region around the solution in which it has n step second-order convergence. Most recent research work on the conjugate gradient algorithm has concentrated in four areas: (i) the avoidance of the implied perfect line search at Step 4; (ii) the avoidance of the steepest descent step at the beginning of each cycle at Step 3; (iii) the introduction of an improved resetting test at

Step 6, so that the cycle need not contain the full n iterations; and (iv) preconditioning the objective function to obtain a better eigenstructure and hence improved performance.

3. Inexact Line Searches

The proof of finite termination of the conjugate gradient algorithm only requires that the line search at Step 4 satisfies Conditions II and III of Wolfe's theorem when $k = 1$. If, however, the direction of search is not reset regularly to the steepest-descent direction, then, to retain finite termination, it is necessary to test the direction of search against Wolfe's Condition I regularly; in the implementations discussed, this is done at every iteration, and also Conditions II and III are satisfied at each iteration.

The n -step second-order convergence property can be retained if a parabolic or cubic interpolation is included in the line search and the accuracy steadily improved so that v , the angle between $g^{(k+1)}$ and $p^{(k)}$, satisfies

$$|\cos v| < R \|g^{(1)}\|, \tag{12}$$

where $g^{(1)}$ is the gradient at the start of the cycle.

As this is still a stringent condition to achieve as $\|g^{(1)}\| \rightarrow 0$, methods were sought for retaining n -step second-order convergence for more relaxed line searches. In the earliest of these, the *gradient prediction method* (Dixon, Ref. 9), this rate was achieved by the introduction of two additional vectors, thus not greatly increasing the storage requirement. The basis of the method was to generate the same set of conjugate directions that would have resulted from the use of perfect line searches and to store a correction vector.

Suppose that a step $d^{(k)}$ is taken along the direction $p^{(k)}$ to a new point $x^{(k+1)}$, where the new gradient is $g^{(k+1)}$, but that, if a perfect step had been taken, the step would have been $\tilde{d}^{(k)}$. Then, from the properties of a quadratic function, it can be shown that the following ratios are equal:

$$\theta_k = \frac{\|\tilde{d}^{(k)}\|}{\|d^{(k)}\|} = \frac{p^{(k)T} g^{(k)}}{p^{(k)T} (g^{(k)} - g^{(k+1)})} = \frac{p^{(k)T} g^{(k)}}{p^{(k)T} y^{(k)}}. \tag{13}$$

This enables us to predict both the point that would have been reached with a perfect line search and the gradient at that point. Also, due to the conjugacy of the directions generated, these corrections are independent, and we can define $x^{(k+1)*}$ and $g^{(k+1)*}$, the point and gradient that would have been reached after k perfect line searches, as

$$g^{(k+1)*} = g^{(k+1)} - w^{(k+1)}, \quad x^{(k+1)*} = x^{(k+1)} - z^{(k+1)}, \tag{14a}$$

where

$$w^{(k-1)} = w^{(k)} - (\theta - 1)(y^{(k)}), \quad w^{(1)} = 0, \quad (14b)$$

$$z^{(k-1)} = z^{(k)} - (\theta - 1)(d^{(k)}), \quad z^{(1)} = 0, \quad (14c)$$

$$\theta - 1 = \frac{p^{(k)T} g^{(k-1)}}{p^{(k)T} y^{(k)}}.$$

The prediction g^* is then used in the conjugate gradient algorithm to calculate p , with $y^{(k-1)}$ in the formula for b being replaced by $g^{(k)*} - g^{(k-1)*}$. On a quadratic function, if $g^{(k)*} = 0$, this indicates that, with perfect line searches, $x^{(k)*}$ would be the solution and that a step z is necessary to obtain the solution. On a nonquadratic function, it seems logical to include a line search along z if such a situation arises.

An alternative method for generating conjugate directions without performing accurate line searches was given by Nazareth (Ref. 10), namely,

$$t^{(k-1)} = -y^{(k)} + \frac{y^{(k)T} y^{(k)}}{t^{(k)T} y^{(k)}} t^{(k)} + \frac{y^{(k)T} y^{(k-1)}}{t^{(k-1)T} y^{(k-1)}} t^{(k-1)}; \quad (15)$$

in this form, the new direction $t^{(k+1)}$ does not depend (apart from a scalar multiplier) on the length of the step along $t^{(k)}$. The same set of conjugate directions is therefore generated for arbitrary step sizes, and the correction step z can be calculated as before. Again, two additional vector stores are needed, $t^{(k-1)}$ and z , and again the value of $g^{(k)*}$ can be introduced to indicate when the step z would predict the solution.

4. Restarting Procedures

In the early algorithms, the restarting strategy was usually to restart whenever $k = n$ or $n + 1$. When n is very large and the number of clusters of similar eigenvalues is very small, this can be very inefficient. It is therefore often felt desirable to reset more regularly. If the gradient prediction algorithm is being used, then it becomes natural to reset when

$$\|g^{(k)*}\| \leq \epsilon_0 \|g^{(1)}\|,$$

as this implies that enough iterations have been undertaken to approximately minimize a quadratic function closely related to a nonlinear objective function.

Powell (Ref. 4) has suggested restarting whenever

$$|g^{(k)T} g^{(k-1)}| \geq 0.2 \|g^{(k)}\|^2. \quad (16)$$

The left-hand side would be zero if the conjugate gradient algorithm were working with perfect line searches on a quadratic function. Its size is therefore an indicator of the nonconjugacy of the search directions, and so is an indicator that the cycle should be terminated.

Other measures could be similarly used; for instance, $t^{(k-1)T}y^{(k)}$ should be zero for conjugate directions and is available to the user of Nazareth's algorithm without additional store.

It also seems desirable to restart if the direction is not effectively downhill. Powell suggests restarting if

$$-1.2\|g^{(k)}\|^2 \leq d^{(k)T}g^{(k)} \leq -0.8\|g^{(k)}\|^2 \tag{17}$$

is not satisfied, while in the codes reported in the tables the cycle was also restarted whenever

$$|g^{(k)T}d^{(k)}| \leq \epsilon_1 \|g^{(k)}\| \|d^{(k)}\|, \quad \epsilon_1 = 0.001, \tag{18}$$

i.e., whenever Wolfe's Condition I was not satisfied.

5. Arbitrary Starting Directions

When it is decided to restart an algorithm, it is questionable whether it is always advisable to restart by setting $p^{(1)} = -g^{(1)}$.

Beale (Ref. 11) analyzed the modifications that arose in the generation of conjugate directions if this assumption was not made and showed that conjugate directions were still obtained if

$$t^{(k+1)} = -g^{(k+1)} + b_1 t^{(1)} + b_k t^{(k)}, \tag{19}$$

where

$$b_1 = \frac{g^{(k+1)T}y^{(1)}}{t^{(1)T}y^{(1)}}. \tag{20}$$

This approach allows a set of conjugate directions to be generated starting from any initial direction $t^{(1)}$. However, the fact that $t^{(1)}$ remains part of the formula for $t^{(k+1)}$ throughout the cycle may be undesirable, although the storing of $t^{(1)}$ would enable the conjugacy check

$$t^{(1)T}g^{(k+1)} = 0 \tag{21}$$

to be made and used as an effective restarting strategy.

An alternative way of allowing an arbitrary starting direction was suggested by Allwright (Ref. 12), who introduced a change of variables,

$$x = Lz, \tag{22}$$

where

$$H = LL^T.$$

Then, a quadratic function

$$F(x - x^*) = \frac{1}{2}(x - x^*)^T G(x - x^*)$$

is equivalent to

$$F(z - z^*) = \frac{1}{2}(z - z^*)^T L^T G L(z - z^*),$$

and

$$\nabla F_z = L^T G L(z - z^*) = L^T G(x - x^*) = L^T g, \quad (23)$$

which, when transformed back into x -space, becomes simply Hg .

If the conjugate direction method is applied to make the directions conjugate in the z -plane, then the directions in the x -plane are simply

$$p^{(1)} = t^{(1)} = -Hg^{(1)}, \quad (24)$$

$$p^{(k+1)} = t^{(k+1)} = -Hg^{(k+1)} + bt^{(k)}, \quad (25)$$

where

$$b = \frac{g^{(k)T} H y^{(k-1)}}{p^{(k-1)T} H y^{(k-1)}}. \quad (26)$$

There are many alternative ways of selecting H so that $p^{(1)} = -Hg^{(1)}$ has any preselected value p_h . One simple formula is to set

$$H = I - \frac{g^{(1)} g^{(1)T}}{g^{(1)T} g^{(1)}} - \frac{p_h p_h^T}{p_h^T g^{(1)}}, \quad (27)$$

but almost certainly this is not the best choice.

Indeed, if the Nazareth three-term formula (15) is used to generate the set of directions, then the directions generated are conjugate for an arbitrary initial direction $t^{(1)}$; see Nocedal, Ref. 13. However, if $t^{(1)}$ is not the gradient direction, $t^{(k)}$ can become zero away from the minimum. This implies that, provided the algorithm is restarted in the unlikely event that

$$\|t^{(k+1)}\| < \epsilon \|g^{(1)}\|, \quad (28)$$

we can use the Allwright matrix to precondition the problem to improve the eigenstructure of $L^T G L$ and thus improve convergence. However, this was not included in the algorithm when the numerical results were obtained, the matrix given by (27) being used.

6. New Conjugate Gradient Algorithm

The new conjugate gradient algorithm proposed combines the three-term conjugate direction formula proposed by Nazareth with a restarting procedure and includes a preconditioning matrix. The line search chosen is inexact, but incorporates a parabolic interpolation and terminates at a point that satisfies Wolfe's Conditions II and III. If any direction is generated that does not satisfy Wolfe's Condition I, the method is reset to ensure finite termination to a region around a solution defined by $g^T g \leq \epsilon$.

To be precise, the proposed method generates the new search direction by the following formula:

$$t^k = -Hg^k + b_k t^{k-1} + b_{k-1} t^{k-2}, \quad (29)$$

where

$$b_{k-1} = \frac{y^{kT} H y^{k-1}}{t^{k-1T} H y^{k-1}}, \quad k > 1, \quad (30)$$

$$b_k = \frac{y^{kT} H y^k}{t^{kT} H y^k}, \quad (31)$$

and H is a positive-definite matrix. Tests are reported for two alternative values of H :

$$(i) \quad H = I, \quad (32a)$$

$$(ii) \quad H = I - g_h g_h^T / g_h^T g_h - p_h p_h^T / p_h^T g_h, \quad (32b)$$

where $g_h = g^{(1)}$ and p_h is the step direction before resetting k .

The algorithm is restarted whenever one of the conditions below holds:

$$(i) \quad k \geq N + 1,$$

$$(ii) \quad (p_h^T H y^k)^2 > 0.2 (p_h^T H p_h) (y^k H y^k), \quad (33)$$

(iii) Wolfe's Condition I does not hold,

$$(iv) \quad \|g^*\| \leq \epsilon_0 \|g^k|. \quad (34)$$

$$(v) \quad \|t^{(k)}\| \leq \epsilon_0 \|g^k|. \quad (35)$$

The second test (33) checks whether the new direction is reasonably conjugate to the initial steepest descent direction in the frame being used.

The matrix H is updated at the restart of a cycle if

$$g_h^T p_h < -C_1 \sqrt{(p_h^T p_h)(g_h^T g_h)}, \quad C_1 = 0.001, \quad (36)$$

i.e., if the resulting matrix would be positive definite.

The structure of the algorithm based on the above philosophy consists of the following steps; at any stage, it requires $13N$ -vector stores for z , g^k , g^{k-1} , t^k , t^{k-1} , t^{k-2} , y^k , y^{k-1} , C^k , g_h , p_h , w^k and a work space. Let us assume the following notations:

$$v_1 = t^{k+1}, \quad v_0 = v^k, \quad v_{-1} = v^{k-1},$$

etc., where v is any given vector.

Algorithm 6.1 (OPCG)

Step 1. Select $x^{(0)}$, $k=0$, $i=0$.

Step 2. Evaluate $g_0 = g(x)^{(0)}$, $g_h = +g_0$, $p_h = -g_0$. Set $t^0 = -g_0$, $H = I$, and $a = 1/\sqrt{(g_0^T g_0)}$.

Step 3. Perform an approximate line search, $d_0 = at_0$, $x_1 = x_0 + d_0$, $i = i + 1$.

Put $k = k + 1$ and downrate all vectors $g_{-1} = g_0$, $x_{-1} = x_0$, $x_0 = x_1$, etc. Evaluate $g_0 = g(x_0)$. For $k > 2$, if the line search is along the correction vector z_0 , then do one of the following:

- (i) if $g_h^T p_h \geq -C\sqrt{[(p_h^T p_h)(g_h^T g_h)]}$, then go to Step 13;
- (ii) Otherwise, update the matrix $H = I - g_h g_h^T / g_h^T g_h - p_h p_h^T / p_h^T g_h$ or use matrix $H = I$, and continue.

Step 4. Calculate g_0 . If $\|g_0\| \leq \epsilon_0$, stop. If $i \geq i_{\max}$, stop.

Step 5. Calculate

$$y_0 = g_0 - g_{-1},$$

$$t_0 = -Hy_0 + \beta t_{-2} + \gamma t_{-1},$$

$$z_0 = z_{-1} + (\theta - 1)t_{-1}, \quad \text{with } z_{-1} = 0, \text{ if } k = 0,$$

$$w_0 = w_{-1} + (\theta - 1)y_{-1}, \quad \text{with } w_{-1} = 0, \text{ if } k = 0,$$

$$g_0^* = g_0 - w_0,$$

where

$$\theta - 1 = -g_0^T d_{-1} / y_0^T t_{-1}, \quad k \geq 1,$$

$$\beta = \begin{cases} y_0^T H y_{-1} / y_{-1}^T H t_{-2}, & k \geq 2, \\ 0, & k = 0, 1, \end{cases}$$

$$\gamma = \begin{cases} y_0^T H y_0 / y_0^T H t_{-1}, & k \geq 1, \\ 0, & k = 0. \end{cases}$$

Step 6. If $C_1 \|t_0\| \cdot \|g_0\| \geq -t_0^T g_0$, then go to Step 8. For $k > 1$, if $(p_h^T H y_0)^2 > 0.2(p_h^T H p_h)(y_0^T H y_0)$, then go to Step 8. If $\|g_0^*\| \leq C_1 \|g_0\|$ or $\|t_0\| \leq C_1 \|g_0\|$, then go to Step 8.

The first of these two conditions is Wolfe's Condition I, while the others are tests (33), (34), and (35), mentioned above.

Step 7. If $k \geq n + 1$, go to Step 8. If $k \leq n$, go to Step 3.

Step 8. Test z_0 for Wolfe's Condition I. If satisfied, go to Step 11.

Step 9. Set $g_h = g_0$, $p_h = z$. Update the matrix H .

Step 10. If $g_h^T p_h \geq -C_1 \sqrt{[(p_h^T p_h)(g_h^T g_h)]}$, i.e., H is unsatisfactory, then go to Step 13; otherwise, go to Step 4, with $k = 0$.

Step 11. Set $t_0 = z_0$, $g_h = g_0$, $p_h = z_0$, $a = 1.0$, $k = 0$.

Step 12. Go to Step 3.

Step 13. Set $g_h = g_0$, $p_h = t_{-1}$, $k = 0$. Update the matrix H .

Step 14. If $g_h^T p_h \geq -C_1 \sqrt{[(p_h^T p_h)(g_h^T g_h)]}$, i.e., H is unacceptable, then go to Step 15; otherwise, go to Step 4, with $k = 1$.

Step 15. Set $t_0 = -g_0$, $k = 0$, $g_h = -g_0$, $p_h = -g_0$; i.e., set $H = I$. Go to Step 3.

The line search at Step 3 consists of first performing a parabolic interpolation using the values of the function and its derivative at the starting point and the function value at an offset point. If the predicted function value does not satisfy Wolfe's Conditions II and III, then the Armijo (Ref. 14) procedure is adopted based on that step, i.e., the predicted step is halved or doubled until

$$f(x + at) < f(x) + 0.1af'(x)^T t,$$

$$f(x + 2at) > f(x) + 0.2af'(x)^T t.$$

7. Numerical Results

7.1. Comparison Algorithms. In order to justify proposing a new conjugate gradient code at this time, it was felt necessary to demonstrate that it was indeed an improvement over existing available codes (Table 1). We therefore chose for comparison purposes first the two standard Harwell routines VA08A and VA14A (Ref. 4), second the quoted results in Buckley and Le Nir (Ref. 15), and also the two earlier Hatfield codes [CONGRA, Dixon, Ref. 6; and CONLS, Dixon, Ref. 9].

Table 1. Algorithm codes.

OPCG1	= with preconditioner H and including test 33.
OPCGS	= without preconditioner H but including test 33
OPCG3	= with preconditioner H but excluding test 33.
OPCG4	= without preconditioner H and excluding test 33.
BLALN	= Buckley and Le Nir method (Ref. 15).
VA14A	= Harwell subroutines.
VA08A	= Harwell subroutines.
CONGRA	= Fletcher-Reeves (Ref. 2) as described in Dixon (Ref. 6).
CONLS	= Dixon (Ref. 9).
FRCG	= Fletcher-Reeves with the new line search.

Results for four versions of the new code are included for completeness, with and without the preconditioning matrix H and with and without the restart test (33), as it was felt of interest to indicate the effect of these innovations.

For each algorithm, the amount of computational effort is measured in terms of the effective number of function evaluations, given by $EFE = N_F + nN_G$, where n = number of unknowns, N_F = number of function calls, and N_G = number of gradient calls. In all cases, the termination criterion was set at

$$g^T g = 1E - 4,$$

and the constant C_1 in Wolfe's Condition 1 was set at $1E - 3$.

7.2. Test Problems. The following are the details of the test functions used.

$$(P1) \quad I(x) = \sum_{j=1}^{n/2} \sum_{i=1}^{10} [\exp(-x_{2j-1} - t_i) - 5 \exp(-x_{2j} t_i) - \exp(-t_i) + 5 \exp(-10t_i)]^2,$$

$$t_i = i/10, \quad i = 1, 2, \dots, 10,$$

$$x^{(0)} = (5, 0, 5, 0, \dots).$$

$$(P2) \quad I(x) = \sum_{i=1}^{10} [\exp(-x_1 z_i) - 5 \exp(-x_2 z_i) - \exp(-z_i) + 5 \exp(-10z_i)]^2,$$

$$z_i = i/10,$$

$$x^{(0)} = (1, 2)^T.$$

$$(P3) \quad I(x) = \sum_{i=1}^{10} [\exp(-x_1 z_i) - x_3 \exp(-x_2 z_i) - \exp(-z_i) + 5 \exp(-10z_i)]^2,$$

$$z_i = i/10,$$

$$x^{(0)} = (1, 2, 1)^T.$$

$$(P4) \quad I(x) = \sum_{i=1}^{10} [x_3 \exp(-x_1 z_i) - x_4 \exp(-x_2 z_i) - \exp(-z_i) + 5 \exp(-10z_i)]^2,$$

$$z_i = i/10,$$

$$x^{(0)} = (1, 2, 1, 1)^T.$$

$$(P5) \quad I(x) = 100(x_1^2 - x_2)^2 + (1 - x_1)^2,$$

$$x^{(0)} = (-1.2, 1)^T.$$

$$(P6) \quad I(x) = 100(x_1^2 - x_2)^8 + (1 - x_1)^8,$$

$$x^{(0)} = (-1.2, 1)^T.$$

$$(P7) \quad I(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2 + 90(x_4 - x_3^2)^2 + (1 - x_3)^2 + 10.1\{(x_2 - 1)^2 + (x_4 - 1)^2\} + 19.8(x_2 - 1)(x_4 - 1),$$

$$x^{(0)} = (-3, -1, -3, -1)^T.$$

$$(P8) \quad I(x) = \sum_{j=1}^{n/4} [(x_{4j-3} + 10x_{4j-2})^2 + 5(x_{4j-1} - x_{4j})^2 + (x_{4j-2} - 2x_{4j-1})^4 + 10(x_{4j-3} - x_{4j})^4],$$

$$x^{(0)} = (3, -1, 0, 3, \dots)^T.$$

$$(P9) \quad I(x) = \sum_{i=1}^{n/2} [100(x_{2i} - x_{2i-1}^2)^2 + (1 - x_{2i-1})^2],$$

$$x^{(0)} = (-1.2, 1, \dots)^T.$$

$$(P10) \quad I(x) = \sum_{i=2}^n [i(2x_i - x_{i-1})^2],$$

$$x^{(0)} = (1, 1, 1, 1, \dots)^T.$$

$$(P11) \quad I(x) = \sum_{i=2}^n [100(x_1 - x_i^2)^2 + (1 - x_2)^2],$$

$$x^{(0)} = (-1.2, 1, \dots)^T.$$

$$(P12) \quad I(x) = \sum_{i=1}^n \left[n + i - \sum_{j=1}^n (A_{ij} \sin(x_j) + B_{ij} \cos(x_j)) \right]^2,$$

$$A_{ij} = \delta_{ij}, \quad B_{ij} = i\delta_{ij} + 1,$$

$$\delta_{ij} = \begin{cases} 0, & \text{if } i \neq j, \\ 1, & \text{if } i = j, \end{cases}$$

$$x^{(0)} = (1/n, 1/n, \dots).$$

7.3. Numerical Results. The numerical results have been presented in two tables. In Table 2, the results are for the classic test problems with $n \leq 4$. In Table 3, results are presented for some problems with $10 \leq n \leq 60$. This latter set of results is the more interesting. Surprisingly, the results without the preconditioning matrix H (i.e., OPG2 and OPG4) are better than those with it, and this led to the rediscovery that the three-term conjugate gradient formula generates conjugate directions for any given starting direction t^0 [reported in Dixon, Ref. 16, and previously in Nocedal, Ref. 13]. As the choice of H used only had the effect of making the first direction the steepest descent direction in the transformed space, it is not necessary for this purpose and could be chosen to have other desirable properties. The results with the reset test were usually better, but not universally better, than those without the reset test, and it has been retained, the code OPG2 being adopted as the definitive version. On comparing the results of this code with those of other codes on the large-dimensional problems in Table 3, it obviously well outperforms VA08A and VA014A,

Table 2. Comparison of performance on problems with $n \leq 4$.

Problem	1	1	2	3	4	5	6	7	8	12
n	2	4	3	4	4	2	2	4	4	4
OPCG1	—	76	—	—	—	386	90	—	127	68
OPCG2	81	70	59	84	143	170	37	F4	119	56
OPCG3	—	136	—	—	—	386	90	—	127	53
OPCG4	81	70	59	67	140	170	37	F4	126	53
VA14A	36	80	45	92	120	195	F1	F2	345	44
VA08A	63	280	156	147	690	276	300	640	365	170
CONGRA	272	620	84	600	1251	366	99	390	446	1035
CONLS	80	690	62	134	362	170	247	1680	472	1060
FRCG	178	502	104	87	111	495	F3	1111	104	81

F1. The run terminated with $f = 0.018 (f^* = 0)$.

F2. The run terminated with $f = 1.6 (f^* = 0)$.

F3. The run terminated with $f = 1.93E - 3 (f^* = 0)$.

F4. The run terminated with $f = 7.8$, a stationary point which satisfies the convergency criteria.

Table 3. Comparison of performance on problems with $10 \leq n \leq 80$.

Problem	8	8	9	9	10	10	10	11	12
n	60	80	10	20	10	20	30	10	40
OPCG1	2111	6646	974	4111	119	439	831	—	225
OPCG2	1688	2055	393	801	119	439	831	702	225
OPCG3	2183	1566	782	2326	119	439	831	—	225
OPCG4	2122	2964	454	888	119	439	831	708	225
BLALN	7991	5913	671	1281	209	819	1891	858	—
VA14A	2623	3483	671	1281	297	819	1643	935	410
VA08A	9577	16443	957	2100	440	1029	2015	1100	4100
FRCG	3906	7463	1326	2511	119	439	863	F5	224

F5. The run terminated with $f = 1.8(f^* = 0)$.

and the results are also an improvement on those reported in Buckley and Le Nir (Ref. 15). On the less interesting small-dimensional problems, the situation is less clear. Out of interest, an algorithm was written combining the Fletcher-Reeves conjugate gradient formula with the line search strategy incorporated in the new code. This is reported in the tables as FRCG and, as expected, the new code (OPCG2) also regularly performed better than this modification, though this simple code did surprisingly well.

8. Conclusions

In this paper, a modified form of the conjugate gradient algorithm is motivated, having desirable restart properties, that outperforms a wide subset of current conjugate gradient codes. The new code contains provision for a preconditioning matrix that ought to lead to a further improvement in performance.

References

1. HESTENES, M., and STIEFEL, E., *Methods of Conjugate Gradients for Solving Linear Systems*, Journal of Research of National Bureau of Standards, Vol. 49, pp. 409-436, 1952.
2. FLETCHER, R., and REEVES, C. M., *Function Minimization by Conjugate Gradients*, Computer Journal, Vol. 17, pp. 149-154, 1964.
3. POLAK, E., and RIBIÈRE, G., *Note sur la Convergence de Methodes de Directions Conjugées*, Revue Française d'Informatique et Recherche Operationnelle, Vol. 3, pp. 35-43, 1969.

4. POWELL, M. J. D., *Restart Procedures for the Conjugate Gradient Method*, Mathematical Programming, Vol. 12, pp. 241-254, 1977.
5. WOLFE, P., *Convergence Conditions for Ascent Methods*, SIAM Review, Vol. 11, pp. 226-235, 1969.
6. DIXON, L. C. W., *Nonlinear Optimization: A Survey of the State of the Art*, Software for Numerical Mathematics, Edited by D. J. Evans, Academic Press, New York, New York, 1974.
7. COHEN, A., *Rate of Convergence of Several Conjugate Gradient Algorithms*, SIAM Journal, Vol. 9, pp. 248-259, 1972.
8. DIXON, L. C. W., *On Quadratic Termination and Second-Order Convergence, Two Properties of Unconstrained Optimization Algorithms*, Toward Global Optimization, Edited by L. C. W. Dixon and G. P. Szegö, North-Holland Press, Amsterdam, Holland, 1975.
9. DIXON, L. C. W., *Conjugate Directions without Line Searches*, Journal of the Institute of Mathematics and Its Applications, Vol. 11, pp. 317-328, 1973.
10. NAZARETH, L., *A Conjugate Gradient Algorithm without Line Searches*, Journal of Optimization Theory and Applications, Vol. 25, pp. 373-387, 1977.
11. BEALE, E. M. L., *A Derivation of Conjugate Gradients*, Numerical Methods for Nonlinear Optimization, Edited by F. A. Lootsma, Academic Press, London, England, 1972.
12. ALLWRIGHT, J. C., *Improving the Conditioning of Optimal Control Problems Using Simple Methods*, Recent Mathematical Developments in Control, Edited by O. J. Bell, Academic Press, London, England, 1972.
13. NOCEDAL, J., *On the Method of Conjugate Gradients for Function Minimization*, Rice University, PhD Thesis, 1978.
14. ARMIJO, L., *Minimization of Functions Having Lipschitz Continuous First Partial Derivative*, Pacific Journal of Mathematics, Vol. 16, pp. 1-16, 1966.
15. BUCKLEY, A., and LE NIR, X. X., *Q-N-Like Variable Storage Conjugate Gradients*, Mathematical Programming, Vol. 27, pp. 155-175, 1983.