

GLOBAL CONVERGENCE PROPERTIES OF CONJUGATE GRADIENT METHODS FOR OPTIMIZATION*

JEAN CHARLES GILBERT[†] AND JORGE NOCEDAL[‡]

Abstract. This paper explores the convergence of nonlinear conjugate gradient methods without restarts, and with practical line searches. The analysis covers two classes of methods that are globally convergent on smooth, nonconvex functions. Some properties of the Fletcher–Reeves method play an important role in the first family, whereas the second family shares an important property with the Polak–Ribière method. Numerical experiments are presented.

Key words. conjugate gradient method, global convergence, unconstrained optimization, large-scale optimization

AMS(MOS) subject classifications. 65, 49

1. Introduction. The object of this paper is to study the convergence properties of several conjugate gradient methods for nonlinear optimization. We consider only the case where the methods are implemented without regular restarts, and ask under what conditions they are globally convergent for general smooth nonlinear functions. The analysis will allow us to highlight differences among various conjugate gradient methods, and will suggest new implementations.

Our problem is to minimize a function of n variables,

$$(1.1) \quad \min f(x),$$

where f is smooth and its gradient g is available. We consider iterations of the form

$$(1.2) \quad d_k = \begin{cases} -g_k & \text{for } k = 1, \\ -g_k + \beta_k d_{k-1} & \text{for } k \geq 2, \end{cases}$$

$$(1.3) \quad x_{k+1} = x_k + \alpha_k d_k,$$

where β_k is a scalar, and α_k is a steplength obtained by means of a one-dimensional search. We call this iteration a *conjugate gradient method* if β_k is such that (1.2)–(1.3) reduces to the linear conjugate gradient method in the case when f is a strictly convex quadratic and α_k is the exact one-dimensional minimizer. Some of the results of this paper, however, also apply to methods of the form (1.2)–(1.3) that do not reduce to the linear conjugate gradient method.

The best-known formulas for β_k are called the Fletcher–Reeves (FR), Polak–Ribière (PR), and Hestenes–Stiefel (HS) formulas, and are given by

$$(1.4) \quad \beta_k^{\text{FR}} = \|g_k\|^2 / \|g_{k-1}\|^2,$$

$$(1.5) \quad \beta_k^{\text{PR}} = \langle g_k, g_k - g_{k-1} \rangle / \|g_{k-1}\|^2,$$

* Received by the editors July 9, 1990; accepted for publication (in revised form) February 15, 1991.

[†] Institut National de Recherche d'Informatique et d'Automatique, Domaine de Voluceau, Rocquencourt, B.P. 105, 78153 Le Chesnay, France.

[‡] Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois 60208. The work of this author was supported by United States Department of Energy grant DE-FG02-87ER25047-A001 and National Science Foundation grant CCR-8902096.

$$(1.6) \quad \beta_k^{\text{HS}} = \langle g_k, g_k - g_{k-1} \rangle / \langle d_{k-1}, g_k - g_{k-1} \rangle.$$

Here, $\langle \cdot, \cdot \rangle$ is the scalar product used to compute the gradient and $\| \cdot \|$ denotes its associated norm. The numerical performance of the Fletcher–Reeves [6] method is somewhat erratic: it is sometimes as efficient as the Polak–Ribi  re and Hestenes–Stiefel methods, but it is often much slower. Powell [18] gives an argument showing that, under some circumstances, the Fletcher–Reeves method with exact line searches will produce very small displacements, and will normally not recover unless a restart along the gradient direction is performed. In spite of these drawbacks, Zoutendijk [27] has shown that the method cannot fail. He proved that the Fletcher–Reeves method with exact line searches is globally convergent on general functions. Al-Baali [1] extended this result to inexact line searches.

The Hestenes–Stiefel and Polak–Ribi  re methods appear to perform very similarly in practice, and are to be preferred over the Fletcher–Reeves method. Nevertheless, in a remarkably laborious paper, Powell [19] was able to show that the Polak–Ribi  re method with exact line searches can cycle infinitely without approaching a solution point. The same result applies to the Hestenes–Stiefel method, since the two methods are identical when $\langle g_k, d_{k-1} \rangle = 0$, which holds when line searches are exact. Since the steplength of Powell’s example would probably be accepted by any practical line search, it appears unlikely that a satisfactory global convergence result can be found for the Polak–Ribi  re and Hestenes–Stiefel methods. In contrast, Al-Baali’s convergence result for the less efficient Fletcher–Reeves method is very satisfactory. This disconcerting state of affairs motivated the present study.

In this paper we will consider various choices of β_k and various line search strategies that result in globally convergent methods. In §2 we describe the approach used in our analysis, and summarize some of the previous work in the area. Section 3 establishes global convergence for the class of methods with $|\beta_k| \leq \beta_k^{\text{FR}}$, and describes a modification of the Polak–Ribi  re formula. In §4 we consider methods that use only nonnegative values for β_k , and which are, in some sense, related to the Polak–Ribi  re method. In particular, we show that a suggestion of Powell [20] to set $\beta_k = \max\{\beta_k^{\text{PR}}, 0\}$ results in global convergence, even for inexact line searches. Further remarks on the convergence results are made in §5, and the results of some numerical experiments are presented in §6.

We note that this paper does not study the rate of convergence of conjugate gradient methods. For some results on this subject, see Crowder and Wolfe [5], Cohen [4], Powell [17], Baptist and Stoer [2], and Stoer [22].

2. Preliminaries. Some important global convergence results for conjugate gradient methods have been given by Polak and Ribi  re [16], Zoutendijk [27], Powell [19], and Al-Baali [1]. In this section we will see that the underlying approach used for these analyses is essentially the same, and we will describe it in detail, since it is also the basis for the results presented in this paper. Before doing so, we describe our notation, state the assumptions we make about the objective function, and consider the line search strategy.

Notation and definitions. We denote the starting point by x_1 , and define $s_k := x_{k+1} - x_k$ and $y_k := g_{k+1} - g_k$. We say that d_k is a *descent direction* if $\langle g_k, d_k \rangle < 0$. We will also make use of the angle θ_k between $-g_k$ and d_k :

$$(2.1) \quad \cos \theta_k := -\langle g_k, d_k \rangle / \|g_k\| \|d_k\|.$$

The Fletcher–Reeves, Polak–Ribi  re, and Hestenes–Stiefel methods will be abbreviated as FR, PR, and HS, respectively. For a derivation of these methods and a

discussion of some of their properties, see Gill, Murray, and Wright [11] and Fletcher [7].

ASSUMPTIONS 2.1. (i) *The level set $\mathcal{L} := \{x : f(x) \leq f(x_1)\}$ is bounded.*
(ii) *In some neighborhood \mathcal{N} of \mathcal{L} , the objective function f is continuously differentiable, and its gradient is Lipschitz continuous, i.e., there exists a constant $L > 0$ such that*

$$(2.2) \quad \|g(x) - g(\tilde{x})\| \leq L\|x - \tilde{x}\|,$$

for all $x, \tilde{x} \in \mathcal{N}$.

Note that these assumptions imply that there is a constant $\bar{\gamma}$ such that

$$(2.3) \quad \|g(x)\| \leq \bar{\gamma}, \quad \text{for all } x \in \mathcal{L}.$$

Let us now turn our attention to the line search. An efficient strategy, studied by Wolfe [25], consists in accepting a positive steplength α_k if it satisfies the two conditions:

$$(2.4) \quad f(x_k + \alpha_k d_k) \leq f(x_k) + \sigma_1 \alpha_k \langle g_k, d_k \rangle$$

$$(2.5) \quad \langle g(x_k + \alpha_k d_k), d_k \rangle \geq \sigma_2 \langle g_k, d_k \rangle,$$

where $0 < \sigma_1 < \sigma_2 < 1$. We will sometimes also refer to more ideal line search conditions. To this end let us define the following strategy: a positive steplength α_k is accepted if

$$(2.6) \quad f(x_k + \alpha_k d_k) \leq f(x_k + \hat{\alpha}_k d_k),$$

where $\hat{\alpha}_k$ is the *smallest* positive stationary point of the function $\xi_k(\alpha) := f(x_k + \alpha d_k)$. Assumptions 2.1 ensure that $\hat{\alpha}_k$ exists. Note that both the first local minimizer and the global minimizer of f along the search direction satisfy (2.6).

Any of these line search strategies is sufficient to establish the following very useful result.

THEOREM 2.1. *Suppose that Assumptions 2.1 hold, and consider any iteration of the form (1.3), where d_k is a descent direction and α_k satisfies one of the following line search conditions:*

- (i) *the Wolfe conditions (2.4)–(2.5), or*
- (ii) *the ideal line search condition (2.6).*

Then

$$(2.7) \quad \sum_{k \geq 1} \cos^2 \theta_k \|g_k\|^2 < \infty.$$

This result was essentially proved by Zoutendijk [27] and Wolfe [25], [26]. We shall call (2.7) the *Zoutendijk condition*.

We can now describe the basic ideas used for the convergence analysis. The first results, by Polak and Ribi  re [16] and Zoutendijk [27], assume exact line searches. The term *exact line search* can be ambiguous. Sometimes, it implies that a one-dimensional minimizer is found, but often it simply means that the orthogonality condition

$$(2.8) \quad \langle g_k, d_{k-1} \rangle = 0$$

is satisfied. Throughout the paper we will indicate in detail the conditions required of the line search. Let us suppose that d_{k-1} is a descent direction and that the line search satisfies Zoutendijk's condition and condition (2.8). From (1.2) and (2.8) we have that

$$(2.9) \quad \cos \theta_k = \frac{\|g_k\|}{\|d_k\|},$$

which shows that d_k is a descent direction. Substituting this relation in Zoutendijk's condition (2.7) we obtain

$$(2.10) \quad \sum_{k \geq 1} \frac{\|g_k\|^4}{\|d_k\|^2} < \infty.$$

If one can show that $\{\|d_k\|/\|g_k\|\}$ is bounded, which means that $\{\cos \theta_k\}$ is bounded away from zero, then (2.10) immediately gives

$$(2.11) \quad \lim_{k \rightarrow \infty} g_k = 0.$$

This is done by Polak and Ribière [16] for their method, assuming that f is *strongly convex*, i.e., $\langle g(x) - g(\tilde{x}), x - \tilde{x} \rangle \geq c \|x - \tilde{x}\|^2$, for some positive constant c and for all x and \tilde{x} in \mathcal{L} .

For general functions, however, it is usually impossible to bound $\{\|d_k\|/\|g_k\|\}$ a priori, and only a weaker result than (2.11) can be obtained, namely,

$$(2.12) \quad \liminf_{k \rightarrow \infty} \|g_k\| = 0.$$

To obtain this result one proceeds by contradiction. Suppose that (2.12) does not hold, which means that the gradients remain bounded away from zero: there exists $\gamma > 0$ such that

$$(2.13) \quad \|g_k\| \geq \gamma$$

for all $k \geq 1$. Then (2.10) implies that

$$(2.14) \quad \sum_{k \geq 1} \frac{1}{\|d_k\|^2} < \infty.$$

We conclude that the iteration can fail only if $\|d_k\| \rightarrow \infty$ sufficiently rapidly. The method of proof used by Zoutendijk for the FR method consists in showing that, if (2.13) holds, then $\|d_k\|^2$ can grow at most linearly, i.e.,

$$\|d_k\|^2 \leq c k,$$

for some constant c . This contradicts (2.14), proving (2.12).

The analysis for *inexact* line searches that satisfy Zoutendijk's condition can proceed along the same lines if one can show that the iteration satisfies

$$(2.15) \quad \cos \theta_k \geq c \|g_k\|/\|d_k\|,$$

for some positive constant c . Then, this relation can be used instead of (2.9) to give (2.10), and the rest of the analysis is as in the case of exact line searches.

Al-Baali [1] shows that the FR method gives (2.15) if the steplength satisfies the *strong* Wolfe conditions:

$$(2.16) \quad f(x_k + \alpha_k d_k) \leq f(x_k) + \sigma_1 \alpha_k \langle g_k, d_k \rangle$$

$$(2.17) \quad |\langle g(x_k + \alpha_k d_k), d_k \rangle| \leq -\sigma_2 \langle g_k, d_k \rangle,$$

where $0 < \sigma_1 < \sigma_2 < 1$. In fact, it is necessary to require that $\sigma_2 < \frac{1}{2}$ for the result to hold. He thus shows that (2.12) holds for the FR method.

Al-Baali's result is also remarkable in another respect. By establishing (2.15), which by (2.1) is equivalent to

$$(2.18) \quad \langle g_k, d_k \rangle \leq -c \|g_k\|^2,$$

he proved that the FR method using the strong Wolfe conditions (with $\sigma_2 < \frac{1}{2}$) always generates descent directions. Prior to this result it was believed that it was necessary to enforce the descent condition while doing the line search.

In this paper we use the approach described above to establish the global convergence of various algorithms with inexact line searches. As we do so, we will repeatedly encounter (2.18), which appears to be a natural way of guaranteeing descent for conjugate gradient methods. We call (2.18) the *sufficient descent condition*. The first class of methods we consider, in §3, is related to the FR method. We show that any method of the form (1.2)–(1.3) is globally convergent if β_k satisfies $|\beta_k| \leq \beta_k^{\text{FR}}$. The result readily suggests a new implementation of the PR method that preserves its efficiency and assures its convergence.

In §4, we study methods with $\beta_k \geq 0$ that are, in some sense, related to the PR method. A particular case is the following adaptation of the PR method, which consists in restricting β_k to positive values: we let

$$(2.19) \quad \beta_k = \max\{\beta_k^{\text{PR}}, 0\}.$$

The motivation for this strategy arises from Powell's analysis of the PR method. Powell [19] assumes that the line search always finds the first stationary point, and shows that there is a twice continuously differentiable function and a starting point such that the sequence of gradients generated by the PR method stays bounded away from zero. Since Powell's example requires that some consecutive search directions become almost contrary, and since this can only be achieved (in the case of exact line searches) when $\beta_k < 0$, Powell [20] suggests modifying the PR method as in (2.19). In §4 we show that this choice of β_k does indeed result in global convergence, both for exact and inexact line searches. Moreover, we show that the analysis also applies to a family of methods with $\beta_k \geq 0$ that share a common property with the PR method—we call this Property (*).

3. Iterations constrained by the FR method. In this section we will see that it is possible to obtain global convergence if the parameter β_k is appropriately bounded in magnitude. We consider a method of the form (1.2)–(1.3), where β_k is any scalar such that

$$(3.1) \quad |\beta_k| \leq \beta_k^{\text{FR}},$$

for all $k \geq 2$, and where the steplength satisfies the strong Wolfe conditions (2.16)–(2.17) with $\sigma_2 < \frac{1}{2}$. Note that Zoutendijk's result, Theorem 2.1, holds in this case, since the strong Wolfe conditions imply the Wolfe conditions (2.4)–(2.5). The next

two results are based upon the work of Al-Baali [1] for the FR method, and are slightly stronger than those given by Touati-Ahmed and Storey [24].

LEMMA 3.1. *Suppose that Assumptions 2.1 hold. Consider any method of the form (1.2)–(1.3), where β_k satisfies (3.1), and where the steplength satisfies the Wolfe condition (2.17) with $0 < \sigma_2 < \frac{1}{2}$. Then, the method generates descent directions d_k satisfying*

$$(3.2) \quad -\frac{1}{1-\sigma_2} \leq \frac{\langle g_k, d_k \rangle}{\|g_k\|^2} \leq \frac{2\sigma_2 - 1}{1-\sigma_2}, \quad k = 1, \dots$$

Proof. The proof is by induction. The result clearly holds for $k = 1$ since the middle term equals -1 and $0 \leq \sigma_2 < 1$. Assume that (3.2) holds for some $k \geq 1$. This implies that $\langle g_k, d_k \rangle < 0$, since

$$(3.3) \quad \frac{2\sigma_2 - 1}{1-\sigma_2} < 0,$$

by the condition $0 < \sigma_2 < \frac{1}{2}$. From (1.2) and (1.4) we have

$$(3.4) \quad \frac{\langle g_{k+1}, d_{k+1} \rangle}{\|g_{k+1}\|^2} = -1 + \beta_{k+1} \frac{\langle g_{k+1}, d_k \rangle}{\|g_{k+1}\|^2} = -1 + \frac{\beta_{k+1}}{\beta_{k+1}^{\text{FR}}} \frac{\langle g_{k+1}, d_k \rangle}{\|g_k\|^2}.$$

Using the line search condition (2.17) we have

$$|\beta_{k+1} \langle g_{k+1}, d_k \rangle| \leq -\sigma_2 |\beta_{k+1}| \langle g_k, d_k \rangle,$$

which, together with (3.4), gives

$$-1 + \sigma_2 \frac{|\beta_{k+1}| \langle g_k, d_k \rangle}{\|g_k\|^2} \leq \frac{\langle g_{k+1}, d_{k+1} \rangle}{\|g_{k+1}\|^2} \leq -1 - \sigma_2 \frac{|\beta_{k+1}| \langle g_k, d_k \rangle}{\|g_k\|^2}.$$

From the left-hand side of the induction hypothesis (3.2), we obtain

$$-1 - \frac{|\beta_{k+1}|}{\beta_{k+1}^{\text{FR}}} \frac{\sigma_2}{1-\sigma_2} \leq \frac{\langle g_{k+1}, d_{k+1} \rangle}{\|g_{k+1}\|^2} \leq -1 + \frac{|\beta_{k+1}|}{\beta_{k+1}^{\text{FR}}} \frac{\sigma_2}{1-\sigma_2}.$$

Using the bound (3.1), we conclude that (3.2) holds for $k + 1$. \square

Lemma 3.1 achieves three objectives: (i) it shows that all search directions are descent directions, and the upper bound in (3.2) shows that the sufficient descent condition (2.18) holds; (ii) the bounds on $\langle g_k, d_k \rangle$ impose a limit on how fast $\|d_k\|$ can grow when the gradients are not small, as we will see in the next theorem; (iii) from (2.1) and (3.2) we see that there are positive constants c_1 and c_2 such that

$$(3.5) \quad c_1 \frac{\|g_k\|}{\|d_k\|} \leq \cos \theta_k \leq c_2 \frac{\|g_k\|}{\|d_k\|}.$$

Therefore, for the FR method or any method with $|\beta_k| \leq \beta_k^{\text{FR}}$, we have that $\cos \theta_k$ is proportional to $\|g_k\|/\|d_k\|$. We will make good use of this fact later on.

THEOREM 3.2. *Suppose that Assumptions 2.1 hold. Consider any method of the form (1.2)–(1.3), where β_k satisfies (3.1), and where the steplength satisfies the strong Wolfe conditions (2.16)–(2.17), with $0 < \sigma_1 < \sigma_2 < \frac{1}{2}$. Then*

$$\liminf_{k \rightarrow \infty} \|g_k\| = 0.$$

Proof. From (2.17) and Lemma 3.1 we have

$$(3.6) \quad |\langle g_k, d_{k-1} \rangle| \leq -\sigma_2 \langle g_{k-1}, d_{k-1} \rangle \leq \frac{\sigma_2}{1 - \sigma_2} \|g_{k-1}\|^2.$$

Thus from (1.2) and (3.1),

$$\begin{aligned} \|d_k\|^2 &\leq \|g_k\|^2 + 2|\beta_k| |\langle g_k, d_{k-1} \rangle| + \beta_k^2 \|d_{k-1}\|^2 \\ &\leq \|g_k\|^2 + \frac{2\sigma_2}{1 - \sigma_2} |\beta_k| \|g_{k-1}\|^2 + \beta_k^2 \|d_{k-1}\|^2 \\ &\leq \left(\frac{1 + \sigma_2}{1 - \sigma_2} \right) \|g_k\|^2 + \beta_k^2 \|d_{k-1}\|^2. \end{aligned}$$

Applying this relation repeatedly, defining $\hat{\sigma} := (1 + \sigma_2)/(1 - \sigma_2) \geq 1$, and using the condition $|\beta_k| \leq \beta_k^{\text{FR}}$, we have

$$\begin{aligned} \|d_k\|^2 &\leq \hat{\sigma} \|g_k\|^2 + \beta_k^2 (\hat{\sigma} \|g_{k-1}\|^2 + \beta_{k-1}^2 \|d_{k-2}\|^2) \\ &\leq \hat{\sigma} \|g_k\|^4 \sum_{j=1}^k \|g_j\|^{-2}. \end{aligned}$$

Let us now assume that $\|g_k\| \geq \gamma > 0$ for all k . This implies, by (2.3), that

$$(3.7) \quad \|d_k\|^2 \leq \frac{\hat{\sigma} \gamma^4}{\gamma^2} k.$$

We now follow the reasoning described in §2. From the left inequality in (3.5) and Zoutendijk's result (2.7), we obtain (2.10). If the gradients are bounded away from zero, (2.10) implies (2.14). We conclude the proof by noting that (3.7) and (2.14) are incompatible. \square

This theorem suggests the following globally convergent modification of the PR method. It differs from that considered by Touati-Ahmed and Storey [24] in that it allows for negative values of β_k . For all $k \geq 2$ let

$$(3.8) \quad \beta_k = \begin{cases} -\beta_k^{\text{FR}} & \text{if } \beta_k^{\text{PR}} < -\beta_k^{\text{FR}}, \\ \beta_k^{\text{PR}} & \text{if } |\beta_k^{\text{PR}}| \leq \beta_k^{\text{FR}}, \\ \beta_k^{\text{FR}} & \text{if } \beta_k^{\text{PR}} > \beta_k^{\text{FR}}. \end{cases}$$

This strategy avoids one of the main disadvantages of the FR method, as we will now discuss.

We have observed in numerical tests that the FR method with inexact line searches sometimes slows down away from the solution: the steps become very small and this behavior can continue for a very large number of iterations, unless the method is restarted. This behavior was observed earlier by Powell [18], who provides an explanation, under the assumption of exact line searches. It turns out that his argument can be extended to the case of inexact line searches, due to (3.5). The argument is as follows. Suppose that at iteration k an unfortunate search direction is generated, such that $\cos \theta_k \approx 0$, and that $x_{k+1} \approx x_k$. Thus $\|g_{k+1}\| \approx \|g_k\|$, and

$$(3.9) \quad \beta_{k+1}^{\text{FR}} \approx 1.$$

Moreover, by (3.5),

$$\|g_{k+1}\| \approx \|g_k\| \ll \|d_k\|.$$

From this relation, (3.9), and (1.2), we see that $\|d_{k+1}\| \approx \|d_k\| \gg \|g_{k+1}\|$, which by (3.5) implies that $\cos \theta_{k+1} \approx 0$. The argument can therefore start all over again. In §6 we give a numerical example demonstrating this behavior.

The PR method would behave quite differently from the FR method in this situation. If $g_{k+1} \approx g_k$, then $\beta_{k+1}^{\text{PR}} \approx 0$, so that by (1.2) and (3.5), $\cos \theta_{k+1} \gg \cos \theta_k$. Thus the PR method would recover from that situation. Let us now consider the behavior of method (3.8) in these circumstances. We have seen that $\beta_{k+1}^{\text{FR}} \approx 1$, and $\beta_{k+1}^{\text{PR}} \approx 0$, in this case. The method (3.8) will thus set $\beta_{k+1} = \beta_{k+1}^{\text{PR}}$, as desired. It is reassuring that the modification (3.8), which falls back on the FR method to ensure global convergence, avoids the inefficiencies of this method.

The previous discussion highlights a property of the PR method that is not shared by the FR method: when the step is small, β_k^{PR} will be small. This property is essential for the analysis given in the next section, where a method that possesses it will be said to have Property (*).

It is natural to ask if the bound $|\beta_k| \leq \beta_k^{\text{FR}}$ can be replaced by

$$(3.10) \quad |\beta_k| \leq c \beta_k^{\text{FR}},$$

where $c > 1$ is some suitable constant. We have not been able to establish global convergence in this case (although, by modifying Lemma 3.1, one can show that the descent property of the search directions can still be obtained provided $\sigma_2 < 1/(2c)$). In fact, one can prove the following negative result.

PROPOSITION 3.3. *Consider the method (1.2)–(1.3), with a line search that always chooses the first positive stationary point of $\xi_k(\alpha) := f(x_k + \alpha d_k)$. There exists a twice continuously differentiable objective function of three variables, a starting point, and a choice of β_k satisfying (3.10) for some constant $c > 1$, such that the sequence of gradients $\{\|g_k\|\}$ is bounded away from zero.*

Proof. The objective function is taken from the fourth example of Powell [19]. It is twice continuously differentiable. For this function, there is a starting point from which the PR method with a line search providing the first stationary point fails to converge, in the sense that $\|g_k\| \geq \gamma > 0$ for all k . Therefore, using (1.5) and (2.3), we have for all $k \geq 2$,

$$|\beta_k^{\text{PR}}| \leq \frac{2\bar{\gamma}^2}{\gamma^2}.$$

Now, suppose that we computed (but did not use) β_k^{FR} . We would see that for all $k \geq 2$,

$$\beta_k^{\text{FR}} \geq \frac{\gamma^2}{\bar{\gamma}^2}.$$

Combining the two inequalities we obtain

$$|\beta_k^{\text{PR}}| \leq \frac{2\bar{\gamma}^4}{\gamma^4} \beta_k^{\text{FR}}.$$

Therefore, if the constant c in (3.10) is chosen larger than $2\bar{\gamma}^4/\gamma^4$, the PR parameter β_k^{PR} in Powell's example would always satisfy (3.10). \square

We end this section by making an observation about the restart criterion of Powell [18]. Even though this criterion was designed to ensure the convergence of Beale's method, we will apply it to the PR method, and see that it has some of the flavor

of the modifications described in this section. Powell [18] suggests restarting if the following inequality is violated

$$|\langle g_k, g_{k-1} \rangle| \leq \nu \|g_{k-1}\|^2,$$

where ν is a small positive constant. (Powell actually uses g_k instead of g_{k-1} in the right-hand side, but one can argue for either choice.) From (1.4) and (1.5),

$$\beta_k^{\text{PR}} = \beta_k^{\text{FR}} - \frac{\langle g_k, g_{k-1} \rangle}{\|g_{k-1}\|^2}.$$

Applying the restart criterion to the PR method we see that a restart is *not* necessary as long as

$$\beta_k^{\text{FR}} - \nu \leq \beta_k^{\text{PR}} \leq \beta_k^{\text{FR}} + \nu.$$

Once more, β_k^{FR} appears as a measure of the adequacy of β_k^{PR} , but this measure is quite different from (3.1). In the next section we will view Powell's restart criterion from a somewhat different angle.

4. Methods related to the PR method with nonnegative β_k . We now turn our attention to methods with $\beta_k \geq 0$ for all k . In §2 we mentioned that a motivation for placing this restriction comes from the example of Powell, in which the PR method cycles without obtaining the solution. Another reason for keeping $\beta_k \geq 0$ is that it allows us to easily enforce the descent property of the algorithm, as we will now discuss.

Let us consider the iteration (1.2)–(1.3) with any $\beta_k \geq 0$. We will require the *sufficient descent condition*

$$(4.1) \quad \langle g_k, d_k \rangle \leq -\sigma_3 \|g_k\|^2,$$

for some $0 < \sigma_3 \leq 1$ and for all $k \geq 1$. In contrast to the FR method, the strong Wolfe conditions (2.16)–(2.17) no longer ensure (4.1). Note, from (1.2), that

$$(4.2) \quad \langle g_k, d_k \rangle = -\|g_k\|^2 + \beta_k \langle g_k, d_{k-1} \rangle.$$

Therefore, to obtain descent for an inexact line search algorithm, one needs to ensure that the last term is not too large. Suppose that we perform a line search along d_{k-1} , enforcing the Wolfe (or strong Wolfe) conditions, to obtain x_k . If $\langle g_k, d_{k-1} \rangle \leq 0$, the nonnegativity of β_k implies that the sufficient descent condition (4.1) holds. Moreover, if (4.1) is not satisfied, then $\langle g_k, d_{k-1} \rangle > 0$, which means that a one-dimensional minimizer has been bracketed. In this case it is easy to apply a line search algorithm, such as that given by Lemaréchal [12], Fletcher [7], or Moré and Thuente [15], to reduce $|\langle g_k, d_{k-1} \rangle|$ sufficiently and obtain (4.1). This will be discussed in detail in §6.

We now prove a global convergence result for methods that are related to the PR method, but that allow only nonnegative values of β_k . The idea of our analysis is simple, but is somewhat concealed in the proofs. We establish the results by contradiction, assuming that the gradients are bounded away from zero:

$$(4.3) \quad \text{for some } \gamma > 0, \quad \|g_k\| \geq \gamma \quad \text{for all } k \geq 1.$$

Lemma 4.1 shows that in this case the direction of search changes slowly, asymptotically, and Lemma 4.2 proves that a certain fraction of the steps are not too small.

In Theorem 4.3 we show that these two results contradict the assumption that the iterates stay in the bounded level set \mathcal{L} . We conclude that a subsequence of the iterates converges to a stationary point.

For the results that follow, we do not specify a particular line search strategy. We only assume that the line search satisfies the following three properties:

- (i) all iterates remain in the level set \mathcal{L} defined in Assumptions 2.1;

$$(4.4) \quad \{x_k\} \subset \mathcal{L};$$

- (ii) the Zoutendijk condition (2.7) holds; and
- (iii) the sufficient descent condition (4.1) holds.

We mentioned in §2 that the Wolfe line search, as well as the ideal line search (2.6), ensure Zoutendijk's condition and reduce f at each step, which implies (4.4). An exact line search satisfies the sufficient descent condition (4.1), because in this case $\langle g_k, d_k \rangle = -\|g_k\|^2$, and in §6 we describe an inexact line search procedure that satisfies the Wolfe conditions and (4.1) when $\beta_k \geq 0$. Therefore the results of this section apply to both ideal and practical line searches.

For the rest of the section, we assume that convergence does not occur in a finite number of steps, i.e., $g_k \neq 0$ for all k .

LEMMA 4.1. *Suppose that Assumptions 2.1 hold. Consider the method (1.2)–(1.3), with $\beta_k \geq 0$, and with any line search satisfying both the Zoutendijk condition (2.7) and the sufficient descent condition (4.1). If (4.3) holds, then $d_k \neq 0$ and*

$$(4.5) \quad \sum_{k \geq 2} \|u_k - u_{k-1}\|^2 < \infty,$$

where $u_k := d_k / \|d_k\|$.

Proof. First, note that $d_k \neq 0$, for otherwise (4.1) would imply $g_k = 0$. Therefore, u_k is well defined. Now, let us define

$$(4.6) \quad r_k := \frac{-g_k}{\|d_k\|} \quad \text{and} \quad \delta_k := \frac{\beta_k \|d_{k-1}\|}{\|d_k\|}.$$

From (1.2), we have for $k \geq 2$:

$$(4.7) \quad u_k = r_k + \delta_k u_{k-1}.$$

Using the identity $\|u_k\| = \|u_{k-1}\|$ and (4.7), we have

$$(4.8) \quad \|r_k\| = \|u_k - \delta_k u_{k-1}\| = \|\delta_k u_k - u_{k-1}\|$$

(the last equality can be verified by squaring both sides). Using the condition $\delta_k \geq 0$, the triangle inequality, and (4.8), we obtain

$$\begin{aligned} \|u_k - u_{k-1}\| &\leq \|(1 + \delta_k)u_k - (1 + \delta_k)u_{k-1}\| \\ &\leq \|u_k - \delta_k u_{k-1}\| + \|\delta_k u_k - u_{k-1}\| \\ (4.9) \quad &= 2\|r_k\|. \end{aligned}$$

Now, by (2.1) and (4.1), we have

$$\cos \theta_k \geq \sigma_3 \frac{\|g_k\|}{\|d_k\|}.$$

This relation, Zoutendijk's condition (2.7), and (4.6) imply

$$\sum_{k \geq 2} \frac{\|g_k\|^4}{\|d_k\|^2} = \sum_{k \geq 2} \|r_k\|^2 \|g_k\|^2 < \infty.$$

Using (4.3), we obtain

$$\sum_{k \geq 2} \|r_k\|^2 < \infty,$$

which together with (4.9) completes the proof. \square

Of course, condition (4.5) does not imply the convergence of the sequence $\{u_k\}$, but shows that the search directions u_k change slowly, asymptotically.

Lemma 4.1 applies to any choice of $\beta_k \geq 0$. To proceed, we need to require, in addition, that β_k be small when the step $s_{k-1} = x_k - x_{k-1}$ is small. We saw in §3 that the PR method possesses this property and that it prevents the inefficient behavior of the FR method from occurring. We now state this property formally.

PROPERTY (*). *Consider a method of the form (1.2)–(1.3), and suppose that*

$$(4.10) \quad 0 < \gamma \leq \|g_k\| \leq \bar{\gamma},$$

for all $k \geq 1$. Under this assumption we say that the method has Property () if there exist constants $b > 1$ and $\lambda > 0$ such that for all k :*

$$(4.11) \quad |\beta_k| \leq b,$$

and

$$(4.12) \quad \|s_{k-1}\| \leq \lambda \implies |\beta_k| \leq \frac{1}{2b}.$$

It is easy to see that under Assumptions 2.1 the PR and HS methods have Property (*). For the PR method, using the constants γ and $\bar{\gamma}$ in (4.10), we can choose $b := 2\bar{\gamma}^2/\gamma^2$ and $\lambda := \gamma^2/(2L\bar{\gamma}b)$. Then we have, from (1.5) and (4.10),

$$|\beta_k^{\text{PR}}| \leq \frac{(\|g_k\| + \|g_{k-1}\|)\|g_k\|}{\|g_{k-1}\|^2} \leq \frac{2\bar{\gamma}^2}{\gamma^2} = b,$$

and when $\|s_{k-1}\| \leq \lambda$, we have from (2.2),

$$|\beta_k^{\text{PR}}| \leq \frac{\|y_{k-1}\| \|g_k\|}{\|g_{k-1}\|^2} \leq \frac{L\lambda\bar{\gamma}}{\gamma^2} = \frac{1}{2b}.$$

For the HS method, we assume that the descent condition (4.1) and the second Wolfe condition (2.5) are satisfied. Then

$$\begin{aligned} \langle d_{k-1}, y_{k-1} \rangle &= \langle d_{k-1}, g_k \rangle - \langle d_{k-1}, g_{k-1} \rangle \\ &\geq -(1 - \sigma_2) \langle g_{k-1}, d_{k-1} \rangle \\ &\geq (1 - \sigma_2) \sigma_3 \|g_{k-1}\|^2 \\ &\geq (1 - \sigma_2) \sigma_3 \gamma^2. \end{aligned}$$

Using this in (1.6) we obtain

$$|\beta_k^{\text{HS}}| \leq \frac{2\bar{\gamma}^2}{(1 - \sigma_2)\sigma_3\gamma^2} =: b.$$

Now define $\lambda := (1 - \sigma_2)\sigma_3\gamma^2/(2L\bar{\gamma}b)$. Using (2.2) we see that if $\|s_{k-1}\| \leq \lambda$, then

$$|\beta_k^{\text{HS}}| \leq \frac{L\lambda\bar{\gamma}}{(1 - \sigma_2)\sigma_3\gamma^2} = \frac{1}{2b}.$$

It is clear that many other choices of β_k give rise to algorithms with Property (*). For example, if β_k has Property (*), so do $|\beta_k|$ and $\beta_k^+ := \max\{\beta_k, 0\}$.

The next lemma shows that if the gradients are bounded away from zero, and if the method has Property (*), then a fraction of the steps cannot be too small. We let \mathbf{N}^* denote the set of positive integers, and for $\lambda > 0$ we define

$$\mathcal{K}^\lambda := \{i \in \mathbf{N}^* : i \geq 2, \|s_{i-1}\| > \lambda\},$$

i.e., the set of integers corresponding to steps that are larger than λ . We will need to consider groups of Δ consecutive iterates, and for this purpose we define

$$\mathcal{K}_{k,\Delta}^\lambda := \{i \in \mathbf{N}^* : k \leq i \leq k + \Delta - 1, \|s_{i-1}\| > \lambda\}.$$

Let $|\mathcal{K}_{k,\Delta}^\lambda|$ denote the number of elements of $\mathcal{K}_{k,\Delta}^\lambda$ and let $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote, respectively, the floor and ceiling operators.

LEMMA 4.2. *Suppose that Assumptions 2.1 hold. Consider the method (1.2)–(1.3), with any line search satisfying (4.4), the Zoutendijk condition (2.7), and the sufficient descent condition (4.1), and assume that the method has Property (*). Suppose also that (4.3) holds. Then there exists $\lambda > 0$ such that, for any $\Delta \in \mathbf{N}^*$ and any index k_0 , there is a greater index $k \geq k_0$ such that*

$$|\mathcal{K}_{k,\Delta}^\lambda| > \frac{\Delta}{2}.$$

Proof. We proceed by contradiction. Suppose that

$$(4.13) \quad \begin{cases} \text{for any } \lambda > 0, \text{ there exists } \Delta \in \mathbf{N}^* \text{ and } k_0 \text{ such that} \\ \text{for any } k \geq k_0, \text{ we have } |\mathcal{K}_{k,\Delta}^\lambda| \leq \frac{\Delta}{2}. \end{cases}$$

Assumptions 2.1 and equations (4.4) and (4.3) imply that (4.10) holds. Since the method has Property (*), there exists $\lambda > 0$ and $b > 1$ such that (4.11) and (4.12) hold for all k . For this λ , let Δ and k_0 be given by (4.13).

For any given index $l \geq k_0 + 1$, we have

$$\begin{aligned} \|d_l\|^2 &\leq (\|g_l\| + |\beta_l| \|d_{l-1}\|)^2 \\ &\leq 2\|g_l\|^2 + 2\beta_l^2 \|d_{l-1}\|^2 \\ &\leq 2\bar{\gamma}^2 + 2\beta_l^2 \|d_{l-1}\|^2, \end{aligned}$$

where the second inequality follows from the fact that, for any scalars a and b , we have $2ab \leq a^2 + b^2$, and hence $(a + b)^2 \leq 2a^2 + 2b^2$. By induction, we obtain

$$(4.14) \quad \|d_l\|^2 \leq c(1 + 2\beta_l^2 + 2\beta_l^2 2\beta_{l-1}^2 + \cdots + 2\beta_l^2 2\beta_{l-1}^2 \cdots 2\beta_{k_0}^2),$$

where c depends on $\|d_{k_0-1}\|$, but not on the index l . Let us consider a typical term in (4.14):

$$(4.15) \quad 2\beta_l^2 2\beta_{l-1}^2 \cdots 2\beta_k^2,$$

where

$$(4.16) \quad k_0 \leq k \leq l.$$

We now divide the $2(l - k + 1)$ factors of (4.15) into groups of 2Δ elements, i.e., if $N := \lfloor (l - k + 1)/\Delta \rfloor$, then (4.15) can be divided into N or $N + 1$ groups, as follows:

$$(4.17) \quad (2\beta_{l_1}^2 \cdots 2\beta_{k_1}^2), \dots, (2\beta_{l_N}^2 \cdots 2\beta_{k_N}^2),$$

and possibly

$$(4.18) \quad (2\beta_{l_{N+1}}^2 \cdots 2\beta_k^2),$$

where $l_i = l - (i - 1)\Delta$, for $i = 1, \dots, N + 1$, and $k_i = l_{i+1} + 1$, for $i = 1, \dots, N$. Note from (4.16) that $k_i \geq k_0$, for $i = 1, \dots, N$, so that we can apply (4.13) for $k = k_i$. We thus have

$$(4.19) \quad p_i := |\mathcal{K}_{k_i, \Delta}^\lambda| \leq \frac{\Delta}{2}, \quad i = 1, \dots, N.$$

This means that in the range $[k_i, k_i + \Delta - 1]$ there are exactly p_i indices j such that $\|s_{j-1}\| > \lambda$, and thus there are $(\Delta - p_i)$ indices with $\|s_{j-1}\| \leq \lambda$. Using this fact, (4.11), and (4.12), we examine a typical factor in (4.17),

$$\begin{aligned} 2\beta_{l_i}^2 \cdots 2\beta_{k_i}^2 &\leq 2^\Delta b^{2p_i} \left(\frac{1}{2b}\right)^{2(\Delta-p_i)} \\ &= 2^{\Delta-2\Delta+2p_i} b^{2p_i-2\Delta+2p_i} \\ &= (2b^2)^{2p_i-\Delta} \\ &\leq 1, \end{aligned}$$

since by (4.19), $2p_i - \Delta \leq 0$ and $2b^2 > 1$. Therefore each of the factors in (4.17) is less than or equal to 1, and so is their product. For the last group of factors, given in (4.18), we simply use (4.11):

$$2\beta_{l_{N+1}}^2 \cdots 2\beta_k^2 < (2b^2)^\Delta.$$

We conclude that each term on the right-hand side of (4.14) is bounded by $(2b^2)^\Delta$, and as a result we have

$$(4.20) \quad \|d_l\|^2 \leq c(l - k_0 + 2),$$

for a certain positive constant c independent of l . In other words, we have shown that $\|d_l\|^2$ grows at most linearly, and we now obtain a contradiction as described in §2. Recalling that (4.1) implies condition (2.15) and using the Zoutendijk condition (2.7), we obtain that

$$\gamma^4 \sum_{k \geq 1} \frac{1}{\|d_k\|^2} \leq \sum_{k \geq 1} \frac{\|g_k\|^4}{\|d_k\|^2} < \infty.$$

This contradicts (4.20), concluding the proof. \square

THEOREM 4.3. *Suppose that Assumptions 2.1 hold. Consider the method (1.2)–(1.3) with the following three properties:*

- (i) $\beta_k \geq 0$ for all k ;
- (ii) the line search satisfies (4.4), the Zoutendijk condition (2.7), and the sufficient descent condition (4.1);
- (iii) Property (*) holds.

Then $\liminf \|g_k\| = 0$.

Proof. We proceed by contradiction, assuming (4.3). Therefore, the conditions of Lemmas 4.1 and 4.2 hold. Defining $u_i := d_i/\|d_i\|$, as before, we have for any two indices l, k , with $l \geq k$:

$$\begin{aligned} x_l - x_{k-1} &= \sum_{i=k}^l \|s_{i-1}\| u_{i-1} \\ &= \sum_{i=k}^l \|s_{i-1}\| u_{k-1} + \sum_{i=k}^l \|s_{i-1}\| (u_{i-1} - u_{k-1}). \end{aligned}$$

Taking norms,

$$\sum_{i=k}^l \|s_{i-1}\| \leq \|x_l - x_{k-1}\| + \sum_{i=k}^l \|s_{i-1}\| \|u_{i-1} - u_{k-1}\|.$$

By (4.4) and Assumptions 2.1 we have that the sequence $\{x_k\}$ is bounded, and thus there exists a positive constant B such that $\|x_k\| \leq B$, for all $k \geq 1$. Thus

$$(4.21) \quad \sum_{i=k}^l \|s_{i-1}\| \leq 2B + \sum_{i=k}^l \|s_{i-1}\| \|u_{i-1} - u_{k-1}\|.$$

Let $\lambda > 0$ be given by Lemma 4.2. Following the notation of this lemma, we define $\Delta := \lceil 8B/\lambda \rceil$. By Lemma 4.1, we can find an index k_0 such that

$$(4.22) \quad \sum_{i \geq k_0} \|u_i - u_{i-1}\|^2 \leq \frac{1}{4\Delta}.$$

With this Δ and k_0 , Lemma 4.2 gives an index $k \geq k_0$ such that

$$(4.23) \quad |\mathcal{K}_{k,\Delta}^\lambda| > \frac{\Delta}{2}.$$

Next, for any index $i \in [k, k + \Delta - 1]$, we have, by the Cauchy–Schwarz inequality and (4.22),

$$\begin{aligned} \|u_{i-1} - u_{k-1}\| &\leq \sum_{j=k}^{i-1} \|u_j - u_{j-1}\| \\ &\leq (i - k)^{1/2} \left(\sum_{j=k}^{i-1} \|u_j - u_{j-1}\|^2 \right)^{1/2} \\ &\leq \Delta^{1/2} \left(\frac{1}{4\Delta} \right)^{1/2} = \frac{1}{2}. \end{aligned}$$

Using this relation and (4.23) in (4.21), with $l = k + \Delta - 1$, we have

$$2B \geq \frac{1}{2} \sum_{i=k}^{k+\Delta-1} \|s_{i-1}\| > \frac{\lambda}{2} |\mathcal{K}_{k,\Delta}^\lambda| > \frac{\lambda\Delta}{4}.$$

Thus $\Delta < 8B/\lambda$, which contradicts the definition of Δ . \square

Since the PR and HS methods have Property (*), the previous theorem applies to them provided we restrict β_k to be nonnegative. This suggests, among other things, the following formulae:

$$(4.24) \quad \beta_k = \max\{\beta_k^{\text{PR}}, 0\},$$

$$(4.25) \quad \beta_k = |\beta_k^{\text{PR}}|,$$

and the corresponding formulae for the HS method. Of particular interest are inexact line searches, such as the Wolfe search. We formally state the convergence result for (4.24)—a choice of β_k suggested by Powell [20].

COROLLARY 4.4. *Suppose that Assumptions 2.1 hold. Consider the method (1.2)–(1.3) with $\beta_k = \max\{\beta_k^{\text{PR}}, 0\}$, and with a line search satisfying the Wolfe conditions (2.4)–(2.5) and the sufficient descent condition (4.1). Then $\liminf \|g_k\| = 0$.*

We conclude this section by noting the relationship between (4.24), which can be viewed as an automatic restarting procedure, and Powell's restarting criterion. The latter states that a restart is not needed as long as

$$(4.26) \quad |\langle g_k, g_{k-1} \rangle| \leq \nu \|g_k\|^2,$$

where we now use g_k and not g_{k-1} in the right-hand side, and where ν is a small positive constant. By (1.5) the condition $\beta_k^{\text{PR}} \geq 0$ is equivalent to

$$\langle g_k, g_{k-1} \rangle \leq \|g_k\|^2.$$

Thus (4.24) can be viewed as a less restrictive restarting test than (4.26). It follows that the global convergence result of Corollary 4.4 also applies to the PR method with Powell's restart (4.26), provided $\nu \leq 1$.

5. Discussion. In §3 we saw that global convergence is obtained for any β_k in the interval $\mathcal{I}_1 = [-\beta_k^{\text{FR}}, \beta_k^{\text{FR}}]$, and in §4 we proved global convergence for any β_k with Property (*) contained in the interval $\mathcal{I}_2 = [0, \infty)$. We now ask whether these results can be combined to obtain larger intervals of admissible β_k . In particular, since the PR method has Property (*), we ask whether global convergence is obtained by restricting β_k^{PR} to the larger interval $\mathcal{I}_1 \cup \mathcal{I}_2$, i.e., by letting

$$\beta_k = \begin{cases} \beta_k^{\text{PR}} & \text{if } \beta_k^{\text{PR}} \geq -\beta_k^{\text{FR}}, \\ -\beta_k^{\text{FR}} & \text{otherwise.} \end{cases}$$

Interestingly enough, global convergence cannot be guaranteed, and this is shown by the fourth example of Powell [19]. In this example, the sequence $\{\beta_k^{\text{PR}}/\beta_k^{\text{FR}}\}$ has exactly three accumulation points:

$$-\frac{1}{3}, 1, \text{ and } 10.$$

Therefore, there exists an index k_0 such that $\beta_k = \beta_k^{\text{PR}} \geq -\beta_k^{\text{FR}}$, for all $k \geq k_0$. Now the function can be modified and the starting point can be changed so that the PR method generates, from the new initial point \tilde{x}_1 , a sequence $\{\tilde{x}_k\}$ with $\tilde{x}_k = x_{k+k_0-2}$, for $k \geq 2$. In this modified example, we have $\tilde{\beta}_k^{\text{PR}} \geq -\tilde{\beta}_k^{\text{FR}}$, for all $k \geq 2$, but the sequence of gradients is bounded away from zero.

There is another example in which intervals of admissible β_k cannot be combined. Any method of the form (1.2)–(1.3) with a line search giving $\langle g_k, d_{k-1} \rangle = 0$ for all k , and with $\beta_k \in \mathcal{I}_3 = [-1, 1]$, is globally convergent. This is easy to see, since in this case

$$\|d_k\|^2 \leq \|g_k\|^2 + \|d_{k-1}\|^2 \leq \dots \leq \bar{\gamma}^2 k,$$

where $\bar{\gamma}$ is an upper bound on $\|g(x)\|$. Therefore $\|d_k\|^2$ grows at most linearly, and global convergence follows by the arguments given in §2. On the other hand, Corollary 4.4 shows that the PR method is convergent if restricted to $\mathcal{I}_2 = [0, \infty)$. However, the PR method may not converge if β_k^{PR} is restricted to $\mathcal{I}_3 \cup \mathcal{I}_2 = [-1, \infty)$. The argument is again based on the counterexample of Powell and on the fact that $\beta_k^{\text{PR}} \geq -\frac{1}{4}$ for all k (this is proved by means of the Cauchy–Schwarz inequality; see Powell [19]). Therefore, in this example $\beta_k^{\text{PR}} \in [-1, \infty)$, but convergence is not obtained.

Therefore we are not able to generalize the results of §§3 and 4, and instead look more closely at the conditions used in these sections. We ask under what conditions is $\beta_k^{\text{PR}} \geq 0$, or $\beta_k^{\text{PR}} \geq -\beta_k^{\text{FR}}$. For strictly convex quadratic functions and exact line searches, the PR method coincides with the FR method. Since β_k^{FR} is always positive, so is β_k^{PR} . Let us now consider strongly convex functions. It turns out that in this case β_k^{PR} can be negative, and in fact can be less than $-\beta_k^{\text{FR}}$.

PROPOSITION 5.1. *There exists a C^∞ strongly convex function of two variables and a starting point x_1 for which the PR method with exact line searches gives $\beta_3^{\text{PR}} < -\beta_3^{\text{FR}} < 0$.*

Proof. Let us introduce the following strictly convex quadratic function \tilde{f} of two variables $x = (x_{(1)}, x_{(2)})$:

$$\tilde{f}(x) := x_{(1)}^2 + \frac{1}{2}x_{(2)}^2,$$

with gradient and Hessian (the Euclidean scalar product is assumed)

$$\nabla \tilde{f}(x) = \begin{pmatrix} 2x_{(1)} \\ x_{(2)} \end{pmatrix}, \quad \nabla^2 \tilde{f}(x) = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}.$$

Starting from the point $x_1 = (-3, 3)$, the PR method with exact line searches gives

$$\nabla \tilde{f}(x_1) = \begin{pmatrix} -6 \\ 3 \end{pmatrix}, \quad \tilde{\alpha}_1 = \frac{5}{9}, \quad \text{and} \quad \tilde{x}_2 = \frac{1}{3} \begin{pmatrix} 1 \\ 4 \end{pmatrix}.$$

Next, it finds

$$\nabla \tilde{f}(\tilde{x}_2) = \frac{2}{3} \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \quad \tilde{\beta}_2^{\text{PR}} = \frac{4}{81}, \quad \tilde{d}_2 = -\frac{10}{27} \begin{pmatrix} 1 \\ 4 \end{pmatrix}, \quad \text{and} \quad \tilde{\alpha}_2 = \frac{9}{10}.$$

The third point is the solution point $x_* = (0, 0)$.

We now perturb the function \tilde{f} inside the ball $B(0, 1) := \{x : x_{(1)}^2 + x_{(2)}^2 < 1\}$, defining

$$f(x) := \tilde{f}(x) + \epsilon \psi(x),$$

where the function ψ will be such that

$$(5.1) \quad \psi(x) = 0 \quad \forall x \notin B(0, 1),$$

and ϵ will be a small positive number. As the line joining x_1 and \tilde{x}_2 does not intersect the closure of $B(0, 1)$, we see that the PR method on this new function f , starting from the same point x_1 , will give $x_2 = \tilde{x}_2$ and $d_2 = \tilde{d}_2$. We now show how to choose the function ψ and the number $\epsilon > 0$ so that f is strongly convex and β_3^{PR} is negative.

We take for ψ a function of the form

$$\psi(x) := \eta(x)\ell(x),$$

where ℓ is the linear function

$$\ell(x) := 4x_{(1)} - x_{(2)},$$

and η is a C^∞ function satisfying

$$\eta(x) = \begin{cases} 1 & \text{if } x \in B(0, \frac{1}{2}) \\ 0 & \text{if } x \notin B(0, 1). \end{cases}$$

Clearly, ψ satisfies (5.1), and has bounded second-order derivatives. Therefore, by choosing ϵ sufficiently small, say $0 < \epsilon < \epsilon_1$, the Hessian of f will be uniformly positive definite and f will be a C^∞ strongly convex function.

Now, when the function f is determined in this manner, there is a unique minimum of f from x_2 in the direction d_2 . As

$$\nabla f(0) = \nabla \tilde{f}(0) + \epsilon \nabla \psi(0) = \epsilon \begin{pmatrix} 4 \\ -1 \end{pmatrix}$$

is orthogonal to $d_2 = \tilde{d}_2$, the one-dimensional minimum is still obtained at $x_3 = (0, 0)$ (but this is no longer the solution point). Therefore,

$$\beta_3^{\text{PR}} + \beta_3^{\text{FR}} = \frac{2|\nabla f(0)|^2 - \langle \nabla f(0), \nabla \tilde{f}(x_2) \rangle}{|\nabla \tilde{f}(x_2)|^2} = \frac{34\epsilon^2 - 4\epsilon/3}{20/9}.$$

We see that $\beta_3^{\text{PR}} < -\beta_3^{\text{FR}} < 0$, if $0 < \epsilon < \epsilon_2 := 2/51$. By taking $\epsilon \in (0, \min(\epsilon_1, \epsilon_2))$, we obtain the desired result. \square

This proposition shows that the convergence result given by Polak and Ribi  re [16], which was obtained for strongly convex functions and exact line searches, is not a consequence of Theorem 4.3, since the latter requires $\beta_k \geq 0$. Nor is it a consequence of Theorem 3.2, because Proposition 5.1 shows that β_k^{PR} can lie outside the interval $[-\beta_k^{\text{FR}}, \beta_k^{\text{FR}}]$.

6. Numerical experiments. We have tested several of the algorithms suggested by the convergence analysis of this paper, on the collection of large test problems given in Table 1.

The starting points used are those given in the references. For the problems of Mor  , Garbow, and Hillstrom [14], we set the parameter factor equal to 1; for test problems 8, 9 and 10, starting point 3 from the reference was used. We verified that, in each run, all the methods converged to the same solution point; otherwise the problem was not included in the test set. The problems are not numbered consecutively because they belong to a larger test set. Since conjugate gradient methods are mainly useful for large problems, our test problems have at least 100 variables.

The following are the methods tested; they differ only in the choice of β_k and, possibly, in the line search.

TABLE 1
List of test functions.

Problem	Name	Reference	<i>n</i>
2	Calculus of variations 2	Gill and Murray [9]	100, 200
3	Calculus of variations 3	Gill and Murray [9]	100, 200
6	Generalized Rosenbrock	Moré et al. [14]	100, 500
8	Penalty 1	Gill and Murray [9]	100, 1000
9	Penalty 2	Gill and Murray [9]	100
10	Penalty 3	Gill and Murray [9]	100, 1000
28	Extended Powell singular	Moré et al. [14]	100, 1000
31	Brown almost linear	Moré et al. [14]	100, 200
38	Tridiagonal 1	Buckley and LeNir [3]	100, 1000
39	Linear minimal surface	Toint [23]	121, 961
40	Boundary-value problem	Toint [23]	100
41	Broyden tridiagonal nonlinear	Toint [23]	100
42	Extended ENGLV1	Toint [23]	1000, 10000
43	Extended Freudenstein and Roth	Toint [23]	100, 1000
45	Wrong extended Wood	Toint [23]	100
46(1)	Matrix square root (ns=1)	Liu and Nocedal [13]	100
46(2)	Matrix square root (ns=2)	Liu and Nocedal [13]	100
47	Sparse matrix square root	Liu and Nocedal [13]	100, 1000
48	Extended Rosenbrock	Moré et al. [14]	1000, 10000
49	Extended Powell	Moré et al. [14]	100, 1000
50	Tridiagonal 2	Toint [23]	100, 1000
51	Trigonometric	Moré et al. [14]	100, 1000
52	Penalty 1 (2nd version)	Moré et al. [14]	1000, 10000
53	INRIA u1ts0.4 (u0=0.95)	Gilbert and Lemaréchal [8]	403
54	INRIA u1cr1.2	Gilbert and Lemaréchal [8]	455
55	INRIA u1cr1.3	Gilbert and Lemaréchal [8]	1559

1. FR: The Fletcher–Reeves method.
2. PR-FR: The Polak–Ribi  re method constrained by the FR method, as in (3.8).
3. PR: The Polak–Ribi  re method.
4. PR⁺: The Polak–Ribi  re method allowing only positive values of β_k^{PR} , as in (4.24).

For the line search we used the algorithm of Mor   and Thuente [15]. This algorithm finds a point satisfying the strong Wolfe conditions (2.16)–(2.17). We used the values $\sigma_1 = 10^{-4}$ and $\sigma_2 = 0.1$, which, by Theorem 3.2, ensure that methods FR and PR-FR are globally convergent. The line search for the PR and PR⁺ methods was performed as follows. We first found a point satisfying the strong Wolfe conditions, using the values of σ_1 and σ_2 mentioned above. If at this point the directional derivative of f is negative, we know that the sufficient descent condition (4.1) holds for the PR⁺ method, and we terminate the line search (this was discussed at the beginning of §4). On the other hand, if the directional derivative is positive, the algorithm of Mor   and Thuente has bracketed a one-dimensional minimizer, and if the line search iteration is continued it will give, in the limit, a point x_k with $\langle g_k, d_{k-1} \rangle = 0$. By continuity and (4.2) it is clear that the line search will find a point satisfying the sufficient descent condition (4.1) in a finite number of iterations. In the numerical tests we set $\sigma_3 = 10^{-2}$ in (4.1). This line search can fail to produce a descent direction for

TABLE 2
Smaller problems.

P	N	FR	PR-FR		PR	PR ⁺	
		it/f-g	it/f-g	mod	it/f-g	it/f-g	mod
2	100	405/827	405/820	351	400/812	400/812	0
3	100	1313/2627	1313/2627	1313	1299/2599	1299/2599	0
6	100	*	261/547	95	256/529	254/525	1
8	100	10/36	15/49	12	9/39	12/47	2
9	100	7/20	8/22	6	8/25	7/20	2
10	100	116/236	93/191	91	118/244	119/244	1
28	100	1426/2855	1291/2584	1289	120/280	168/382	3
31	100	2/3	2/3	1	1/4	1/4	0
38	100	70/142	70/142	47	71/144	71/144	0
39	121	*	59/122	4	59/122	59/122	0
40	100	175/351	175/351	175	132/266	132/266	0
41	100	29/60	24/50	1	24/50	24/50	0
42	1000	10/27	9/25	8	10/34	9/30	2
43	100	16/41	14/39	13	16/44	13/37	1
45	100	*	74/166	66	37/90	45/109	3
46(1)	100	617/1238	253/510	248	257/518	257/518	0
46(2)	100	886/1776	251/506	243	251/506	251/506	0
47	100	151/306	59/122	50	60/124	60/124	0
48	1000	79/185	71/172	66	26/73	23/70	3
49	100	1426/2855	1291/2584	1289	117/281	168/382	3
50	100	72/146	72/146	52	72/146	72/146	0
51	100	202/409	42/94	12	45/103	45/103	0
52	1000	3/10	3/10	2	4/12	4/12	2

the PR method if it terminates at a point with negative directional derivative, and if $\beta_k < 0$ (see the discussion in §4). We used it, nevertheless, because we know of no line search algorithm that is guaranteed to satisfy the strong Wolfe conditions and also provide the descent property for the PR method. Fortunately, in our tests the line search strategy described above always succeeded for the PR method.

Our numerical experience with conjugate gradient methods indicates that it is advantageous to perform a reasonably accurate line search. Therefore, in addition to setting σ_2 to the small number 0.1, we ensured that the line search evaluated the function at least twice. The choice of the initial trial value for the line search is also important. For the first iteration we set it to $1/\|g_1\|$, and for subsequent iterations we used the formula recommended by Shanno and Phua [21], which is based on quadratic interpolation.

The tests were performed on a SPARCstation 1, using FORTRAN in double precision. All runs were stopped when

$$\|g(x_k)\|_\infty < 10^{-5}(1 + |f(x_k)|),$$

except for the INRIA problems for which the runs were stopped when the value of the function had reached a given threshold ($f_{\text{stop}} = 10^{-12}$ for u1ts0.4, $f_{\text{stop}} = -0.8876 \cdot 10^{-2}$ for u1cr1.2 and $f_{\text{stop}} = -0.10625 \cdot 10^{-1}$ for u1cr1.3). The results in Tables 2 and 3 are given in the form: (number of iterations)/(number of function evaluations). The number given under the column “mod” for method PR-FR denotes

TABLE 3
Larger problems.

P	N	FR	FR-PR		PR	PR ⁺	
		it/f-g	it/f-g	mod	it/f-g	it/f-g	mod
2	200	703/1424	701/1420	596	701/1420	701/1420	0
3	200	2808/5617	2808/5617	2808	2631/5263	2631/5263	0
6	500	*	1107/2231	433	1068/2151	1067/2149	1
8	1000	12/39	9/34	7	6/28	10/42	2
10	1000	138/281	145/299	142	165/338	165/338	0
28	1000	533/1102	1369/2741	1366	212/473	97/229	3
31	200	2/4	2/4	1	1/5	1/5	0
38	1000	264/531	263/529	217	262/527	262/527	0
39	961	*	143/290	5	142/287	142/287	0
42	10000	6/26	6/26	5	7/28	6/26	1
43	1000	10/27	15/38	15	10/33	9/29	2
47	1000	422/849	114/233	92	113/231	113/231	0
48	10000	61/143	130/283	123	24/73	19/62	4
49	1000	568/1175	1369/2741	1366	212/473	97/229	3
50	1000	274/551	273/549	245	274/551	274/551	0
51	1000	231/467	40/91	5	40/92	40/92	0
52	10000	4/15	4/15	4	3/13	3/13	1
53	403	**	233/494	130	237/508	237/508	0
54	455	**	44/91	7	44/87	44/87	0
55	1559	23/47	23/49	15	23/47	23/47	0

TABLE 4
Relative performance, in terms of function evaluations.

FR	FR-PR	PR	PR ⁺
> 4.07	1.55	1.02	1.00

the number of iterations for which $|\beta_k^{\text{PR}}| > \beta_k^{\text{FR}}$. For method PR⁺, “mod” denotes the number of iterations for which $\beta_k^{\text{PR}} < 0$. If the limit of 9999 function evaluations was exceeded the run was stopped; this is indicated by “*.” The sign “**” means that the run stopped because the line search procedure described above failed to find a steplength. This occurred when the stopping criterion was very demanding.

It is interesting to note that β_k^{PR} was constrained in most of the iterations of the method PR-FR, but was only rarely modified in the PR⁺ method. Many of the problems were run again for a larger number of variables. The results are given in Table 3.

In these runs the methods were implemented without restarting. We also performed tests in which the methods were restarted along the steepest descent direction every n iterations. (Since n is large, very few restarts were performed.) The FR method improved substantially, but this method was still the least efficient of the four. The other three methods performed similarly with and without restarts, and we will not present the results here.

In Table 4 we summarize the results of Tables 2 and 3 by giving the relative number of function evaluations required by the four methods. We have normalized the numbers so that PR⁺ corresponds to 1. The symbol $>$ means that FR requires more function evaluations than the number given, since for some runs the method was

stopped prematurely; also, problems 53 and 54, in which FR failed, were not taken into account.

The FR method is clearly the least efficient, requiring a very large number of function evaluations in some problems. The performance of methods PR-FR, PR, and PR⁺ appears to be comparable, but we would not like to draw any firm conclusions from our experiments. PR-FR appears to be preferable to FR, but we have no explanation for its poor performance on some problems. A close examination of the runs provided no new insights about the behavior of the methods. The global convergence analysis of this paper has not suggested a method that is clearly superior to PR. For that it may be necessary to study the convergence rate or other measures of efficiency of the methods. We leave this for a future study.

We conclude by giving an example that illustrates the inefficient behavior of the FR method, as predicted in §3. For problem 45 with $n = 100$, we observed that for hundreds of iterations $\cos \theta_k$ stays fairly constant, and is of order 10^{-2} , while the steps $\|x_k - x_{k-1}\|$ are of order 10^{-2} to 10^{-3} . This causes the algorithm to require a very large number of iterations to approach the solution. A restart along the steepest descent direction terminates this cycle of bad search directions and tiny steps. A similar behavior was observed in several other problems.

Acknowledgment. We are grateful to Michael Powell for showing us how to shorten the proof of Lemma 4.1, and for several other helpful comments.

REFERENCES

- [1] M. AL-BAALI, *Descent property and global convergence of the Fletcher-Reeves method with inexact line search*, IMA J. Numer. Anal., 5 (1985), pp. 121–124.
- [2] P. BAPTIST AND J. STOER, *On the relation between quadratic termination and convergence properties of minimization algorithms, Part II, Applications*, Numer. Math., 28 (1977), pp. 367–392.
- [3] A. BUCKLEY AND A. LENIR, *QN-like variable storage conjugate gradients*, Math. Programming, 27 (1983), pp. 155–175.
- [4] A. COHEN, *Rate of convergence of several conjugate gradient algorithms*, SIAM J. Numer. Anal., 9 (1972), pp. 248–259.
- [5] H. P. CROWDER AND P. WOLFE, *Linear convergence of the conjugate gradient method*, IBM J. Res. Develop., 16 (1972), pp. 431–433.
- [6] R. FLETCHER AND C. M. REEVES, *Function minimization by conjugate gradients*, Comput. J., 7 (1964), pp. 149–154.
- [7] R. FLETCHER, *Practical Methods of Optimization*, John Wiley, New York.
- [8] J. CH. GILBERT AND C. LEMARÉCHAL, *Some numerical experiments with variable-storage quasi-Newton algorithms*, Math. Programming B, 45 (1989), pp. 407–436.
- [9] P. E. GILL AND W. MURRAY, *The numerical solution of a problem in the calculus of variations*, in Recent Mathematical Developments in Control, D.J. Bell, ed., Academic Press, New York, 1973, pp. 97–122.
- [10] ———, *Conjugate-gradient methods for large-scale nonlinear optimization*, Technical report SOL 79-15, Department of Operations Research, Stanford University, Stanford, CA, 1979.
- [11] P. E. GILL, W. MURRAY AND M. H. WRIGHT, *Practical Optimization*, Academic Press, New York, 1981.
- [12] C. LEMARÉCHAL, *A view of line searches*, in Optimization and Optimal Control, A. Auslander, W. Oettli, and J. Stoer, eds., Lecture Notes in Control and Information Science 30, Springer-Verlag, Berlin, 1981, pp. 59–78.
- [13] D. C. LIU AND J. NOCEDAL, *Test results of two limited memory methods for large scale optimization*, Report NAM 04, Dept. of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL, 1988.
- [14] J. J. MORÉ, B. S. GARROW, AND K. E. HILLSTROM, *Testing unconstrained optimization software*, ACM Trans. Math. Software, 7 (1981), pp. 17–41.

- [15] J. J. MORÉ AND D. J. THUENTE, *On line search algorithms with guaranteed sufficient decrease*, Mathematics and Computer Science Division Preprint MCS-P153-0590, Argonne National Laboratory, Argonne, IL, 1990.
- [16] E. POLAK AND G. RIBIÈRE, *Note sur la convergence de méthodes de directions conjuguées*, Revue Française d'Informatique et de Recherche Opérationnelle, 16 (1969), pp. 35–43.
- [17] M. J. D. POWELL, *Some convergence properties of the conjugate gradient method*, Math. Programming, 11 (1976), pp. 42–49.
- [18] ———, *Restart procedures for the conjugate gradient method*, Math. Programming, 12 (1977), pp. 241–254.
- [19] ———, *Nonconvex minimization calculations and the conjugate gradient method*, in Lecture Notes in Mathematics 1066, Springer-Verlag, Berlin, 1984, pp. 122–141.
- [20] ———, *Convergence properties of algorithms for nonlinear optimization*, SIAM Rev., 28 (1986), pp. 487–500.
- [21] D. F. SHANNO AND K. H. PHUA, *Remark on algorithm 500 : Minimization of unconstrained multivariate functions*, ACM Trans. Math. Software, 6 (1980), pp. 618–622.
- [22] J. STOER, *On the relation between quadratic termination and convergence properties of minimization algorithms, Part I, Theory*, Numer. Math., 28 (1977), pp. 343–366.
- [23] PH. L. TOINT, *Test problems for partially separable optimization and results for the routine PSPMIN*, Report Nr 83/4, Department of Mathematics, Facultés Universitaires de Namur, Namur, Belgium, 1983.
- [24] D. TOUATI-AHMED AND C. STOREY, *Efficient hybrid conjugate gradient techniques*, J. Optim. Theory Appl., 64 (1990), pp. 379–397.
- [25] P. WOLFE, *Convergence conditions for ascent methods*, SIAM Rev., 11 (1969), pp. 226–235.
- [26] ———, *Convergence conditions for ascent methods II: some corrections*, SIAM Rev., 13 (1971), pp. 185–188.
- [27] G. ZOUTENDIJK, *Nonlinear programming, computational methods*, in Integer and Nonlinear Programming, J. Abadie, ed., North-Holland, Amsterdam, 1970, pp. 37–86.