On the Q-Superlinear Convergance of Self-Scaling Quasi-Newton Methods

Robert B. Schnabel

CU-CS-144-78 August 1978

University of Colorado at Boulder
DEPARTMENT OF COMPUTER SCIENCE

Abstract

Self-scaling updates have been proposed by Luenberger, Oren, and Spedicato for use in quasi-Newton minimization algorithms. Their departure from other updates is that the intermediate update $\hat{H}_i = \gamma_i H_i$ to the inverse Hessian approximation is performed before each regular update. In recent computational tests by Brodlie and Shanno and Phua, they performed less well than the BFGS update except on problems with a singular Hessian at the solution. In this paper we examine the self-scaling updates in an attempt to explain this behavior. We find that for the self-scaling BFGS update to retain the Q-superlinear convergence of the normal BFGS on problems with a non-singular Hessian at the solution, it is necessary that γ_i converge to 1; a somewhat stronger condition is sufficient. This indicates that asymptotically, use of the scaling parameter is unlikely to be advantageous on non-singular problems. On the other hand, on problems with a singular Hessian at the solution, where only linear convergence is expected in general, γ_i does not necessarily converge to 1, so that the self-scaling update may differ from the BFGS even asymptotically.

1. Introduction

This paper analyzes a class of quasi-Newton methods for solving the unconstrained minimization problem

$$\min_{x \in R^n} f: \quad R^n \to R \tag{1.1}$$

where f is assumed twice continuously differentiable. We give a brief overview of quasi-Newton methods, but assume the reader is basically familiar with them. Recent references include Brodlie [2] and Dennis and Moré [7].

Quasi-Newton methods generate a sequence of points $x_i \in R^n$ which hopefully converge to the solution x_* of (1.1). They are related to Newton's method, in which this sequence is produced by the iteration

$$x_{i+1} = x_i - \nabla^2 f(x_i)^{-1} \nabla f(x_i). \tag{1.2}$$

In the type of quasi-Newton methods with which we are concerned, due to the Hessian matrix $\nabla^2 f(x)$ being expensive or impossible to compute, (1.2) is modified by replacing $\nabla^2 f(x_i)^{-1}$ with an approximation H_i which is updated following each iteration. In addition, a line search parameter λ_1 may be included, so that the iteration becomes

$$x_{i+1} = x_i - \lambda_i H_i \nabla f(x_i), \lambda_i > 0.$$

If H_i is positive definite, λ_i can be chosen so that $f(x_{i+1}) < f(x_i)$. The value $\lambda_i = 1$ is usually attempted first, and is called the direct prediction value of λ_i.

The most successful quasi-Newton methods have chosen H_{i+1} by the Broyden-Fletcher-Goldfarb-Shanno (BFGS) update, [4,8,9,17],

$$H_{i+1} = H_i + \frac{(s_i - H_i y_i) s_i^T + s_i (s_i - H_i y_i)^T}{s_i^T y_i} - \frac{< s_i - H_i y_i, y_i > s_i s_i^T}{(s_i^T y_i)^2} \tag{1.3}$$

$$s_i \triangleq x_{i+1} - x_i$$
$$y_i \triangleq \nabla f(x_i) - \nabla f(x_i).$$

This is one of the many updates which obey the secant equation,

$$H_{i+1} y_i = s_i \tag{1.4}$$

which is the way H_{i+1} is made to resemble $\nabla^2 f(x_{i+1})^{-1}$. Equation (1.4) accurately models $\nabla^2 f(x_*)^{-1}$ if f is quadratic, because then $\nabla^2 f(x_{i+1}) s_i = y_i$, and otherwise models it approximately because

$$\left[\int_{\tau=0}^{1} \nabla^2 f(x_i + \tau s_i) d\tau \right] s_i = y_i.$$

Update (1.3) is the one which, given $H_i \varepsilon R^{n \times n}$ symmetric, solves the problem

$$\min_{H_{i+1} \varepsilon R^{n\times n}} \quad \| H_{i+1} - H_i \|_{\overline{H}^{-1/2}}$$

subject to $H_{i+1} y_i = s_i$, H_{i+1} symmetric

where $\overline{H}$ is any fixed matrix obeying $\overline{H} y_i = s_i$. Here $\|A\|_M$, $A, M \varepsilon R^{n\times n}$ denotes the Frobenius norm of A weighted by M, $\|MAM\|_F$, where the Frobenius norm $\|B\|_F$ of any $B \varepsilon R^{n\times n}$ is the square root of the sum of the squares of the elements of B.

In this paper we are concerned with a variant of update (1.3), the self-scaling update introduced by Oren and Luenberger [10,12], and studied by Oren and Spedicato in many papers including [11,13,19,20]. In these methods, the matrix H_i is multiplied by a positive scalar γ_i before each update, so that

$$\hat{H}_i = \gamma_i H_i. \tag{1.5}$$

$\hat{H}_i$ is then updated by a normal quasi-Newton update which obeys secant equation (1.4); while Oren and Spedicato have proposed many updates, the most successful still seems to be the BFGS, (1.3) with $\hat{H}_i$ in place of H_i. The entire updating process can be combined as a single update to H_i,

$$H_{i+1} = \gamma_i H_i + \frac{(s_i - \gamma_i H_i y_i)s_i^T + s_i(s_i - \gamma_i H_i y_i)^T}{s_i^T y_i} - \frac{\langle s_i - \gamma_i H_i y_i, y_i \rangle s_i s_i^T}{(s_i^T y_i)^2} \tag{1.6}$$

which we will refer to as the self-scaling BFGS update.

The original motivation for the scaling (1.5) (Oren and Luenberger [12]) concerned the performance of a perfect line search algorithm on quadratic problems; its pros and cons are well discussed by Brodlie [3]. In any case this motivation is not directly relevant to current algorithms or most problems. A more general motivation is that (1.5) enables a rough rank n change to H_i, which may enable the general size of H_{i+1} to more accurately resemble $\nabla^2 f(x_{i+1})^{-1}$ than the typical quasi-Newton update, which allows only a rank two modification of H_i. Specific choices of γ_i, as suggested by Luenberger, Oren and Spedicato, are explained in Section 4.

Quasi-Newton algorithms using self-scaling updates have recently been tested extensively on a variety of problems, and compared with the same algorithms using the BFGS and other updates (Brodlie [3], Shanno and Phua [18]). The general conclusion is that the self-scaling update seems to be less effective than the BFGS, except perhaps on problems where the Hessian at the solution is singular. (Shanno and Phua also find that scaling at the first iteration only is helpful). Our objective in this paper is to provide a theoretical analysis of self-scaling methods which will hopefully cast additional light on these results.

In particular, Broyden, Dennis and Moré [5] and Powell [14] have proven that the points generated by (two different) quasi-Newton algorithms using the BFGS update converge Q-superlinearly to the minimum x_* for a large class of problems, as long as $\nabla^2 f(x_*)$ is nonsingular. (Q-superlinear convergence of $\{x_i\}$ to x_* means that

$$\lim_{i \to \infty} \frac{\| x_{i+1} - x_* \|}{\| x_i - x_* \|} = 0$$

where $\| \cdot \|$ is any vector norm.) In Section 2 we show that the same results apply using self-scaling BFGS update (1.6) if $\sum_{i=0}^{\infty} | \gamma_i - 1 |$ is bounded above by a suitable constant. In Section 3 we show that a necessary condition for Q-superlinear convergence on the self-scaling BFGS update is $\liminf_{i \to \infty} \gamma_i = 1$ (for a subsequence containing at least every second iteration). These results cast some doubt upon the asymptotic value of self-scaling updates for problems with $\nabla^2 f(x_*)$ nonsingular, because they show that if the self-scaling BFGS update is to work as well as the BFGS on such problems, then asymptotically it must become the BFGS.

We then examine the specific choices of γ_i (in (1.5-6)) proposed by Luenberger, Oren and Spedicato [10-13,19,20]. In Section 4 we give a new derivation of the most commonly used choice. It is strongly related to the derivation by Biggs [1] of a different variant of the BFGS, which Brodlie [3] finds to be just as effective as the BFGS. The connection gives some indication of why the update of Biggs may be superior to the self-scaling update. In Section 5 we discuss the relation of our analysis to the updates of Luenberger, Oren and

Spedicato. It is an open question whether these choices obey our necessary or sufficient condition for Q-superlinear convergence on problems with $\nabla^2 f(x_*)$ non-singular. However, Q-superlinear convergence with any of the commonly used choices of γ_i would require $\lim_{i \to \infty} \gamma_1 = 1$, so that the BFGS and self-scaling BFGS updates would become the same. On problems with $\nabla^2 f(x_*)$ singular the situation is quite different. On such problems, where no better than linear convergence is expected and the self-scaling updates have performed well in practice, we show by example that γ_i does not in general approach 1 as $i \to \infty$. Thus the self-scaling BFGS algorithm may differ substantially from the normal BFGS in this case. This supports the continued consideration of self-scaling updates for problems with $\nabla^2 f(x_*)$ singular.

We briefly summarize our results in Section 6.

2. Sufficient conditions for Q-superlinear convergence of self-scaling methods

In this section we show that the Q-superlinear convergence results of both Broyden, Dennis and Moré [5] and Powell [14] for quasi-Newton algorithms using the BFGS update (1.3) extend to the self-scaling BFGS update (1.6) if $\sum_{i=0}^{\infty} |\gamma_i - 1|$ is bounded above by a suitable constant. The theorems of Broyden-Dennis-Moré and Powell involve two different quasi-Newton algorithms and different assumptions on the function f, in a way that neither result implies the other. However, they both demonstrate that an algorithm using the BFGS update and eventually only the direct prediction step

$$x_{i+1} = x_i - H_i \nabla f(x_i) \qquad (2.1)$$

will converge locally Q-superlinearly to the minimum of many functions whose Hessian at the solution is nonsingular.

Broyden, Dennis and Moré [5] show that if (2.1) together with BFGS update (1.3) is applied to minimize $f(x)$ with x_0 sufficiently close to the solution x_*, H_0 sufficiently close to $\nabla^2 f(x_*)^{-1}$, $\nabla^2 f(x_*)$ positive definite and some lesser conditions on f, then $\{x_i\}$ converges Q-superlinearly to x_*. In Theorem 2.2 we show that the same result applies using the self-scaling BFGS update (1.6) if $\sum_{i=0}^{\infty} |\gamma_i - 1| \le \sigma$ for σ sufficiently small. The proof is just an extension of the techniques of Broyden, Dennis and Moré for the BFGS. It is based on Theorem 2.1, which gives a general condition on a quasi-Newton update for (2.1) to be locally Q-linearly convergent. This theorem differs from the original Theorem 3.1 in Broyden, Dennis and Moré [5] only by the addition of the $w(x,H)$ term in (2.3). For its proof, which is a straightforward extension of the techniques of Broyden, Dennis and Moré, see Schnabel [15], Theorems 9.2.2 and 9.2.3.

For the remainder of the paper, $\|\cdot\|$ (without subscript) will denote the ℓ_2 vector norm or its induced matrix norm.

<u>Theorem 2.1</u> Let $f\colon R^n \to R$ be twice continuously differentiable in the open convex set D, and assume for some $x_* \in D$ and $\ell \ge 0$,

$$\|\nabla^2 f(x) - \nabla^2 f(x_*)\| \le \ell \|x - x_*\| \tag{2.2}$$

for all $x \in D$, where $\nabla f(x_*) = 0$ and $\nabla^2 f(x_*)$ is positive definite. Let $U\colon R^n \times R^{n\times n} \to 2^{R^{n\times n}} \sim \{\}$ be defined in a neighborhood $N = N_1 \times N_2$ of $(x_*, \nabla^2 f(x_*)^{-1})$ where $N_1 \in D$. Suppose there exist nonnegative constants α_1, α_2 and a nonsingular symmetric $M \in R^{n\times n}$ such that for any $(x,H) \in N$ and $x_+ = x - H\nabla f(x)$, the function U satisfies

$$\|H_+ - \nabla^2 f(x_*)^{-1}\|_M \le (1+\alpha_1 m)\|H - \nabla^2 f(x_*)^{-1}\|_M + \alpha_2 m + w(x,H) \tag{2.3}$$

for each $H_+ \in U(x,H)$, where $m = \max\{\|x_+ - x_*\|, \|x - x_*\|\}$ and w: $R^n \times R^{n\times n} \to R$. Consider the sequences $x_i \in R^n$, $H_i \in R^{n\times n}$ generated from $(x_0,H_0) \in N$ by (2.1) with $H_{i+1} \in U(x_i,H_i)$. Then for each $r \in (0,1)$ and any nonnegative constant ω, there exist positive constants $\varepsilon(r)$, $\delta(r)$ such that if $\|x_0 - x_*\| \le \varepsilon(r)$, $\|H_0 - \nabla^2 f(x_*)^{-1}\|_M \le \delta(r)$ and $\sum_{i=0}^{j} w(x_i,H_i) \le \omega\,\delta(r)$ for all $j \le 0$, then the sequence $\{H_i\}$ is well defined and $\{x_i\}$ converges to x_*. Furthermore,

$$\|x_{i+1} - x_*\| \le r\,\|x_i - x_*\|$$

for each $i \ge 0$, and $\{\|H_i\|\}$, $\{\|H_i^{-1}\|\}$ are uniformly bounded.

The Q-superlinear convergence proof of Broyden, Dennis and Moré [5] for the BFGS is extended to the self-scaling BFGS below. The main step is to show that under its conditions, the norm of the difference of these two updates is bounded above by a constant multiple of $|\gamma_i - 1|$. This enables the application of Theorem 2.1 to prove linear convergence, and then Q-superlinear convergence follows directly from the work of Broyden, Dennis and Moré.

<u>Theorem 2.2</u> Let $f: R^n \to R$, $x_* \in R^n$ satisfy the conditions of Theorem 2.1, and define $M = \nabla^2 f(x_*)^{1/2}$. Consider the sequences $x_i \in R^n$, $H_i \in R^{n\times n}$, $i = 1,2,\ldots$ generated from $x_0 \in R^n$ and a symmetric $H_0 \in R^{n\times n}$ by (2.1) with $H_{i+1} = H_{i+1}(\gamma_i)$ given by (1.6), s_i, y_i defined as in (1.3) and γ_i a nonzero real number. Then there exist positive constants ε, δ, σ such that if

$\|x_0 - x_*\| \le \varepsilon$, $\|H_0 - \nabla^2 f(x_*)^{-1}\|_M \le \delta$ and $\sum_{i=0}^{j} |\gamma_i - 1| \le \sigma$ for all $j \ge 0$, then the sequence $\{H_i\}$ is well defined and $\{x_i\}$ converges Q-superlinearly to x_*.

Proof: Broyden, Dennis and Moré [5] have proven this theorem in the case $\rho_i \equiv 1$. With a slight expansion of their techniques we show that update (1.6) satisfies (2.3) with $w(x_i,H_i) = \frac{9}{2} \mid \gamma_i - 1 \mid$. This proves the linear convergence of $\{x_i\}$ by Theorem 3.1, and then the Q-superlinear convergence follows directly from the work of Broyden, Dennis and Moré.

Broyden, Dennis and Moré show that under the assumptions of this theorem, there exist open neighborhoods N_1 around x_* and N_2 around $\nabla^2 f(x_*)^{-1}$ such that if $x_i \in N_1$ and $H_i \in N_2$ is symmetric, then there exist nonnegative constants α_1, α_2, for which

$$\| H_{i+1}(1) - H \|_M \leq (1+\alpha_1 m_i) \| H_i - H \|_M + \alpha_2 m_i \tag{2.4}$$

where $m_i \overset{\Delta}{=} \max \{ \| x_{i+1} - x_* \|, \| x_i - x_* \| \}$, $H \overset{\Delta}{=} \nabla^2 f(x_*)^{-1}$. ($H_{i+1}(1)$, (1.6) with $\gamma_i = 1$, is simply the BFGS). They also show that under these assumptions,

$$\left\| \left(I - \frac{y_i s_i^T}{s_i^T y_i}\right) A \left(I - \frac{s_i y_i^T}{s_i^T y_i}\right) \right\|_M \leq \frac{9}{4} \| A \|_M$$

for any $A \in R^{n \times n}$. Now, if necessary further restrict N_2 so that

$$\max \{ \| A \|_M , A \in N_2 \} \leq 2.$$

(The constant 2 is arbitrary. Note that as N_2 becomes small, $\| A \|_M \to 1$ for all $A \in N_2$.) Then since update (1.6) can be arranged as

$$H_{i+1}(\gamma_i) = H_{i+1}(1) + (\gamma_i - 1) \left(I - \frac{y_i s_i^T}{s_i^T y_i}\right) H_i \left(I - \frac{s_i y_i^T}{s_i^T y_i}\right), \tag{2.5}$$

the triangle inequality, (2.4), (2.5) and $\| H_i \|_M \leq 2$ yield

$$\| H_{i+1}(\gamma_i) - H \|_M \leq (1+\alpha_1 m_i) \| H_i - H \|_M + \alpha_2 m_i + (9/2)(\gamma_i - 1).$$

Therefore by Theorem 2.1 there exist positive constants ε, δ, and σ such that if $\| x_0 - x_* \| \leq \varepsilon$, $\| H_0 - H \|_M \leq \delta$ and

$\sum_{i=0}^{j} |\gamma_i - 1| \leq \sigma$ for all $j \geq 0$, the sequence $\{H_i\}$ is well defined and $\{x_i\}$ converges Q-linearly to x_*. The proof of Q-superlinear convergence is then identical to that of Broyden, Dennis and Moré [5] for the case $\gamma_i \equiv 1$.

It should be noted that Theorem 2.2 is not peculiar to the BFGS update; it can readily be proven for other updates,such as the DFP or Broyden's method, which have been proven Q-superlinearly convergent by the techniques of Broyden, Dennis and Moré [5].

Powell [14] examines a different algorithm, in which each x_{i+1} is chosen by the line search iteration

$$x_{i+1} = x_i - \lambda_i H_i \Delta f(x_i). \tag{2.6}$$

Parameter λ_i is chosen to obey conditions (2.7), which basically assure that each step results in an adequate reduction in f(x) and a step which is not too small. Wolfe [21] has shown that these conditions are sufficient to guarantee global convergence in many cases, and they are used in a number of computational routines. Powell [14] proves that if f is convex and some lesser conditions are met, then the points produced by (2.6) using the BFGS update and λ_i obeying (2.7) converge to a minimum x_* of f from any x_0 and H_0. If in addition $\nabla^2 f(x_*)$ is positive definite, (2.2) holds and $\lambda_i = 1$ is used whenever it obeys (2.7), he shows that the convergence is Q-superlinear. In Theorem 2.3 we extend these results to the self-scaling BFGS update. The first part, establishing convergence, goes through without much change as long as $\left\{ \prod_{i=0}^{j} \gamma_i \right\}$ is uniformly bounded above and below by positive constants.

The Q-superlinear portion requires the slightly stronger condition

$\sum_{i=0}^{\infty} |\gamma_i - 1| < \infty$; the extension of this portion of Powell's proof depends mainly on the relation between the BFGS and self-scaling BFGS updates established in the proof of Theorem 2.2.

<u>Theorem 2.3</u> Let $f: R^n \rightarrow R$ be convex, $x_0 \in R^n$, and assume that $D = \{x | f(x) \leq f(x_0)\}$ is bounded and f is twice continuously differentiable in D. Let the sequence $x_i \in R^n$, $i = 1,2,\ldots$ be generated from x_0 and a symmetric $H_0 \in R^{n\times n}$ by (2.6), where $H_{i+1} = H_{i+1}(\gamma_i)$ given by (1.6), γ_i is a nonzero real number, and $\lambda_i > 0$ is chosen so that

$$f(x_{i+1}) \leq f(x_i) + \beta_1 \nabla f(x_i)^T (x_{i+1} - x_i) \tag{2.7}$$

$$\nabla f(x_{i+1})^T (x_{i+1} - x_i) \geq \beta_2 \nabla f(x_i)^T (x_{i+1} - x_i)$$

β_1, β_2 constants obeying $0 < \beta_1 < \beta_2 < 1$, $\beta_1 < 1/2$. If there exist constants $0 < \gamma\min \leq 1 \leq \gamma\max$ such that

$$\prod_{i=0}^{j} \gamma_i \in [\gamma\min, \gamma\max] \tag{2.8}$$

for all $j > o$, then the sequence $\{x_i\}$ is well defined and converges to a minimum x_* of f. If in addition $\nabla^2 f(x_*)$ is positive definite, (2.2) holds for all x in some open neighborhood around x_*, and there exists some $\nu > 0$ such that $\sum_{i=0}^{j} |\gamma_i - 1| < \nu$ for all $j \geq 0$, then the rate of convergence is Q-superlinear.

<u>Proof</u>: The proof of convergence is accomplished by suitable modification of the proof of Powell [14]. Perhaps the easiest way is to use the fact that the sequences $\{x_i\}$ which can be produced by the algorithm under consideration and γ_i obeying (2.8) are exactly the same as can be produced if H_i in (2.6) is replaced by $\overline{H}_i$ and the update is replaced by

$$\bar{H}_{i+1} = \bar{H}_{i+1}(\rho_i) = \bar{H}_i + \frac{(\rho_i s_i - \bar{H}_i y_i) s_i^T + s_i(\rho_i s_i - \bar{H}_i y_i)^T}{s_i^T y_i} - \frac{\langle \rho_i s_i - \bar{H}_i y_i, y_i \rangle s_i s_i^T}{(s_i^T y_i)^2} \tag{2.9}$$

with

$$\rho_i \in [1/\gamma \max, 1/\gamma \min] \tag{2.10}$$

(and $\bar{H}_0 = H_0$). This is proven below. The proof of Q-linear convergence is then complete because convergence of the above algorithm using update (2.9) and ρ_i obeying (2.10) has been proven in Schnabel [16], under the same conditions on f. It should be noted that this proof required only one trivial modification of the proof in Powell [14].

The proof that the sequences which can be generated by algorithm (2.6-7) using updates (1.6) or (2.8) are equivalent is an extension of a result in Spedicato [19]. It rests mainly on the observation that if s_i and y_i are the same in (1.6) and (2.8) and $\bar{H}_i = \xi_i H_i$ for any constant ξ, then $\bar{H}_{i+1}(\xi/\gamma_i) = (\xi/\gamma_i) H_{i+1}(\gamma_i)$. Using this fact we first show by induction on k that if sequences $\{x_0,\ldots,x_k\}$ and $\{H_0,\ldots,H_k\}$ can be generated using (1.6) with γ_i obeying (2.8), then the sequences $\{x_0,\ldots,x_k\}$ and $\{H_0,\rho_0 H_1,\ldots \rho_{k-1} H_k\}$ can be generated by (2.9-10) using $\rho_i = 1/\prod_{j=0}^{i} \gamma_j$. This is clearly true for k = 0. Assume it is true for k = i. Then since $\bar{H}_i = \rho_{i-1} H_i$, the search directions are equivalent and the possibilities for x_{i+1} are clearly equivalent. If the same x_{i+1} is chosen by both algorithms, then their s_i's and y_i's are the same. Furthermore, if γ_i obeys (2.8) then ρ_i obeys (2.10). Since $\rho_i = \rho_{i-1}/\gamma_i$ we have $\bar{H}_{i+1}(\rho_i) = \bar{H}_{i+1}(\rho_{i-1}/\gamma_i) = (\rho_{i-1}/\gamma_i) H_{i+1}(\gamma_i)$ $= \rho_i H_{i+1}(\gamma_i)$ and the induction is complete. To complete the proof, one shows in a virtually identical manner that if $\{x_0,\ldots,x_k\}$, $\{\overline{H}_0,\ldots,\overline{H}_k\}$ can be generated by (2.9-10), then $\{x_0,\ldots,x_k\}$, $\{\overline{H}_0,\overline{H}_1/\rho_0,\ldots,\overline{H}_k/\rho_{k-1}\}$ can be generated by (1.6-2.8) using $\gamma_i = \rho_{i-1}/\rho_i$, $\rho_{-1} \triangleq 1$.

The extension of Powell's [14] proof of Q-superlinear convergence to update (1.6) is based on our analysis of this update in proving Theorem 2.2. The first portion of his proof, that $\sum_{i=0}^{\infty} \| x_i - x_* \|$ is bounded, goes through without change. It is then only necessary to extend Theorem 8.7 of Dennis and Moré [7] to update (1.6). To do this, we use the fact from the proof of Theorem 2.2 that given $\{x_i\} \to x_*$, there exists an $i_0 \geq$ such that for all $i \geq i_0$,

$$\| H_{i+1} - H\|_M \leq (1+\alpha_1 m_i) \|H_i - H\|_M + \alpha_2 m_i + (9/4) \|H_i\|_M \,|\gamma_i - 1| \qquad (2.11)$$

where $H = \nabla^2 f(x_*)^{-1}$, $M = H^{-1/2}$, $m_i = \max\{\| x_{i+1} - x_*\|, \|x_i - x_*\|\}$ and α_1, α_2 are nonnegative constants. (Note that consulting the original techniques of Broyden, Dennis and Moré [5], (2.11) requires $\| x_i - x_*\|$ sufficiently small but does not require anything about $\|H_i - H\|$.) To prove Theorem 8.7, it is now necessary to show that (2.11) implies $\{\|H_i\|_M\}$ uniformly bounded. To do this, select some $j \geq i_0$ such that

$$2\,\alpha_1 m + (9/2)\gamma < 1$$

where $m = \sum_{i=j}^{\infty} m_i$, $\gamma = \sum_{i=j}^{\infty} |\gamma_i - 1|$; and then select τ so that

$$\tau \geq \max\{\|H_j - H\|, \frac{\alpha_2 m + (9/4)\|H\|_M \gamma}{1 - 2\alpha_1 m - (9/2)\gamma}\} \qquad (2.12)$$

We show by induction that

$$\|H_i - H\|_M \leq 2\tau, \quad \|H_i\|_M \leq 2\tau + \|H\|_M \qquad (2.13)$$

for all $i \geq j$. For $i = j$ this is clearly true, the second inequality resulting from the triangle inequality. Assume (2.13) true for $i = j, \ldots, k$. Then summing (2.11) for $i = j$ to k gives

$$\begin{aligned}\|H_{k+1} - H\|_M &\le \|H_j - H\|_M + \\ &\qquad \sum_{i=j}^{k} (\alpha_1 m_i \|H_i - H\|_M + \alpha_2 m_i + (9/4)\,\|H_i\|_M |\gamma_i - 1|) \\ &\le \|H_j - H\|_M + \alpha_1 m\, 2.\tau + \alpha_2 m + (9/4)\,(2\tau + \|H\|_M)\,\gamma \\ &\le \|H_j - H\|_M + \tau \le 2\tau\end{aligned}$$

the last two inequalities coming from the application of (2.12). This completes the induction as $\|H_{k+1}\|_M \le 2\tau + \|H\|_M$ again results from the triangle inequality. Now consulting the original proof of Theorem 8.7 of Dennis and Moré [7] in Theorems 3.1-3.4 of Dennis and Moré [6], it is seen that given (2.11) and (2.13), the remainder of the proof of Theorem 8.7 goes through without change. As shown in Powell [14], this completes the proof of Theorem 2.3.

We note that Brodlie states in [3] that Theorem 2.3 is true if restriction (2.8) is relaxed to $\gamma_i \in [\gamma\min, \gamma\max]$, $\gamma\min$, $\gamma\max$ defined as above. We have not been able to show this. Indeed, from our techniques of proof, if Brodlie's statement is correct, then Theorem 2.3 would also apply to update (2.9) as long as $\rho_i \in [(\rho\min)^i | (\rho\max)^i]$ for some constants $o < \rho\min \le 1 \le \rho\min$. We doubt whether this is correct.

3. Necessary conditions for Q-superlinear convergence of self-scaling methods

In this section, we show that for the direct-prediction quasi-Newton method (2.1) using the self-scaling BFGS update to be Q-superlinearly convergent on the class of "well-behaved" functions considered by Broyden, Dennis and Moré [5] or by Powell [14], one must have $\lim_{i\to\infty} \gamma_i = 1$ for all i in some subsequence of R which includes at least every second integer. We accomplish this in Theorem 3.1 by showing that this condition is necessary for Q-superlinear convergence on the nicest function of all, $f(x) = (1/2)x^2$ (starting from any $x_0 \neq 0$ and $H_0 \neq I$). Unfortuantely, we must consider the two dimensional case, since in one dimension H_{i+1} is completely determined by the secant equation and the parameter γ_i has no effect. This makes the analysis of our example quite complex; the reader may wish to skip it. We believe the techniques used are of some further interest, and comment on this following the proof. In Section 5 we show that for any of the commonly used choices of γ_i, the stronger condition $\lim_{i \to \infty} \gamma_i = 1$ can trivially be shown necessary for Q-superlinear convergence on any function obeying the assumptions of Theorem 2.1

To avoid an even messier proof of Theorem 3.1, we modify the update being considered very slightly: if the angle between s_i and s_{i-1} is less than some arbitrarily small constant, we let $H_{i+1} = H_i$. (The actual condition is slightly more complex - see below.) In practice this would be a reasonable implementation condition for the update. However, is should be realized that Theorem 3.1 is still basically true without this restriction (there may be an exception in one weird situation). Besides, from the proof of Theorem 3.1 we see that if Q-superlinear convergence is attained, then even with this condition eventually there will be an update at least every second iteration. We also assume that $\{\|H_i\|\}$ is uniformly bounded. This condition can similarly be omitted at the cost of the same possible exception to Theorem 3.1, but since it is satisfied by all successful quasi-Newton updates we know of on functions for which they Q-superlinearly convergent, we prefer to include it.

We use $\measuredangle(v,w)$ to denote the smallest positive angle from v to w, $v, w \in R^n$, where in measuring this angle we extend v and w to be full lines (see Fig. 3.1). Therefore $\measuredangle(v,w) \in [0, \pi/2]$ for all $v, w \in R^n$.

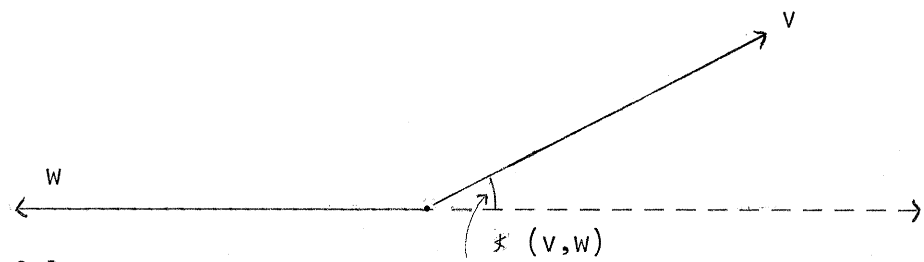

Fig. 3.1

<u>Theorem 3.1</u> Let $f: R^n \to R$, $x_* \in D$ satisfy the conditions of Theorem 2.1. Let the sequences $x_i \in R^n$, $H_i \in R^{n\times n}$, $i = 1, 2, \ldots$ be generated from $x_0 \in R^n$ and a symmetric $H_0 \in R^{n\times n}$ as in Theorem 2.2, with this modification: given some ε in $(0, \frac{\pi}{6})$ and $j(0) = 0$, define at each iteration

$$j(i) = \begin{cases} i & \measuredangle(s_{j(i-1)}, s_i) \geq \varepsilon \\ j(i-1) & \text{otherwise} \end{cases}$$

and if $j(i) = j(i-1)$, then $H_{i+1} = H_i$. Then there exists some such f such that if $x_0 \neq x_*$ and $H_0 \neq \nabla^2 f(x_*)$, then $\{x_i\}$ converges Q-superlinearly to x_* only if $\lim_{i\to\infty} \inf |\gamma_i - 1| = 0$. Furthermore, the subsequence of γ_i which converges to 1 must contain at least every second γ_i.

<u>Proof:</u> The proof is via an example: Q-superlinear convergence is shown to require $\lim_{i\to\infty} \inf |\gamma_i - 1| = 0$ in the case $f(x) = x^T x$, $n = 2$. Note that by a linear transformation, this example can be extended to any positive difinite two dimensional quadratic.

The proof for this example is rather complicated and probably not too instructive. The possiblility that $H_{i+1} = H_i$ makes the notation and proof harder to follow, and so we first show how the proof proceeds if $\measuredangle(s_{i-1},s_i) \geq \varepsilon$ at each iteration. Then we readily extend our techniques to the general case.

Since $\nabla f(x) = x$, $y_i = s_i$ for all i. Thus we only use s_i in what follows. For any $v \in R^2$, $v^\perp$ will denote any vector in R^2 such that $\langle v^\perp, v\rangle = 0$. (By not specifying the orientation of $v^\perp$ we will get a few $\pm$ signs below, but they won't matter.) Finally for the description of convergence to be meaningful we assume that no $x_i = 0$. (In general, this is the case for our example even if $\gamma_i \equiv 1$.)

The main technique of the proof is to express any x_{i+1} as a linear combination of s_i and s_{i-1}, calculate x_{i+2} as a function of s_i and s_{i-1}, and then show that Q-superlinear convergence is possible only if $\liminf_{i \to \infty} |\gamma_i - 1| = 0$.

Define

$$w_i = \measuredangle(s_i, x_{i+1})$$

$$\emptyset_i = \measuredangle(s_{i-1}, s_i)$$

for $i \geq 0$, where s_{-1} is an eigenvector of H_0 which makes $\emptyset_0 \geq \varepsilon$. Note that we are assuming for now that each $\emptyset_i \geq \varepsilon$. Then we can express

$$x_{i+1} = \alpha_i s_{i-1} + \beta_i s_i, \; |\alpha_i| = \frac{\|x_{i+1}\| \sin w_i}{\|s_{i-1}\| \sin \emptyset_i} . \tag{3.1}$$

From (3.1), $x_{i+2} = x_{i+1} - H_{i+1}x_{i+1} = \alpha_i(s_{i-1} - H_{i+1}s_{i-1})$ since $H_{i+1}s_i = s_i$. Now substituting (1.6) for H_{i+1}, using $H_i s_{i-1} = s_{i-1}$) and doing a fair bit of rearranging,

$$x_{i+2} = \pm\alpha_i \frac{s_i^{\perp}}{\|s_i^{\perp}\|} \|s_{i-1}\| \left[(1-\gamma_i) \sin \phi_i \pm \gamma_i \frac{\| H_i s_i \|}{\|s_i\|} \cos \phi_i \sin \measuredangle (s_i, H_i s_i) \right] \tag{3.2}$$

so that using also formula (3.1) for α_i,

$$\frac{\|x_{i+2}\|}{\|x_{i+1}\|} = \sin w_i \left[(1-\gamma_i) \pm \gamma_i \frac{\|H_i s_i\| \cos \phi_i}{\|s_i\| \sin \phi_i} \sin \measuredangle (s_i, H_i s_i) \right] \triangleq r_i . \tag{3.3}$$

We now show that $\lim_{i \to \infty} r_i = 0$ implies $\liminf_{i \to \infty} |\gamma_i - 1| = 0$.

Assume Q-superlinear convergence, i.e., $\lim_{i \to \infty} r_i = 0$. The main step is to show that then $\limsup_{i \to \infty} w_i \neq 0$. From (3.2), the direction of x_{i+1} and $s_{i-1}^{\perp}$ are the same, so that

$$w_i + \phi_i = \pi/2. \tag{3.4}$$

Also define $\delta_i = \measuredangle (x_i s_i)$, so that

$$\phi_i = w_{i-1} \pm \delta_i ,$$

and from (3.4),

$$w_i + w_{i-1} \pm \delta_i = \pi/2, \; i = 1,2,\ldots \quad . \tag{3.5}$$

Now Dennis and Moré [6] show that Q-superlinear convergence implies $\lim_{i \to \infty} \delta_i = 0$, so that from (3.5), $\limsup_{i \to \infty} w_i \neq 0$.

Thus Q-superlinear convergence implies

$$\liminf_{i \to \infty} |r_i / \sin w_i| = 0 . \tag{3.6}$$

Dennis and Moré [6] also show that Q-superlinear convergence and $\{\|H_i\|\}$ uniformly bounded imply $\lim_{i \to \infty} (\|H_i s_i - s_i\| / \|s_i\|) = 0$, so that

$$\lim_{i \to \infty} \frac{\|H_i s_i\|}{\|s_i\|} \sin \measuredangle (s_i, H_i s_i) = 0. \tag{3.7}$$

Since also $\{\cos\emptyset_i / \sin\emptyset_i\}$ is bounded above due to each $\emptyset_i \geq \varepsilon$, (3.3), (3.7) and Q-superlinear convergence imply that

$$\lim_{i\to\infty} |r_i/\sin w_i| = \lim_{i\to\infty} |\gamma_i - 1|. \tag{3.8}$$

Therefore from (3.6) and (3.8), $\liminf_{i\to\infty} |\gamma_i - 1| = \liminf_{i\to\infty} |r_i/\sin w_i| = 0$. From (3.8), Q-superlinear convergence also implies that $\lim_{i\to\infty} (\sin w_i)(1-\gamma_i) = 0$, which with (3.5) and $\lim \delta_i = 0$ shows that the subsequence of $\{\gamma_i\}$ which converges to 1 must contain at least every second γ_i.

This completes the proof under the assumption that each $\emptyset_i \geq \varepsilon$. In the general case, the proof is very similar if we express x_{i+1} as a linear combination of $s_{j(i)}$ and $s_{j(j(i)-1)}$, the last two values of s_i which caused updates to occur. Define

$$\mu_i = \measuredangle(s_{j(i)}, x_{i+1})$$
$$\sigma_i = \measuredangle(s_{j(j(i)-1)}, s_{j(i)}).$$

By the construction of our algorithm, each $\sigma_i \geq \varepsilon$, so we can express

$$x_{i+1} = \hat{\alpha}_i s_{j(j(i)-1)} + \hat{\beta}_i\, s_{j(i)},\ |\hat{\alpha}_i| = \frac{\|x_{i+1}\| \quad \sin\mu_i}{\|s_{j(j(i)-1)}\| \sin\sigma_i}.$$

(We define s_{-1} as before and $j(-1) = -1$.) The same algebra as above then gives the analogous equation to (3.2),

$$x_{i+2} = \pm\hat{\alpha}_i \frac{s_{j(i)}^{\perp}}{\|s_{j(i)}^{\perp}\|} \|s_{j(j(i)-1)}\| \left[(1-\gamma_{j(i)}) \sin\sigma_i \pm \gamma_{j(i)} \frac{\|H_{j(i)}s_{j(i)}\|}{\|s_{j(i)}\|} \cos\sigma_i \sin\measuredangle(s_{j(i)}, H_{j(i)}s_{j(i)}) \right] \tag{3.9}$$

so that

$$\frac{\|x_{i+2}\|}{\|x_{i+1}\|} = \sin \mu_i \left[(1-\gamma_{j(i)}) \pm \gamma_{j(i)} \frac{\|H_{j(i)} s_{j(i)}\| \cos \sigma_i}{\|s_{j(i)}\| \sin \sigma_i} \sin \measuredangle (s_{j(i)}, H_{j(i)} s_{j(i)}) \right]. \tag{3.10}$$

From (3.10), the proof that Q-superlinear convergence requires $\liminf_{i \to \infty} \gamma_i = 0$ is similar to the above. Define

$$\theta_i = \measuredangle (s_{j(i-1)}, s_i),$$

$$\tau_i = \measuredangle (s_{j(i)}, s_i).$$

Using still $w_i = \measuredangle (s_i, x_{i+1})$, (3.9) shows that

$$w_i + \theta_i = \pi/2.$$

Also

$$w_i = \mu_i \pm \tau_i$$

and from the construction of the algorithm, either $j(i) = i$ so that $\tau_i = 0$, or $j(i) = j(i-1)$ and $\tau_i < \varepsilon$. Using again $\delta_i = \measuredangle (x_i, s_i)$,

$$\theta_i = \mu_{i-1} \pm \delta_i$$

and again Q-superlinear convergence implies $\lim_{i \to \infty} \delta_i = 0$. Thus from the above equations

$$\mu_i + \mu_{i-1} \pm (\tau_i \pm \delta_i) = \pi/2$$

and $\limsup_{i \to \infty} |\tau_i \pm \delta_i| \le \varepsilon$. This shows that $\limsup_{i \to \infty} \mu_i \neq 0$ which is again the main result. These equations also show that

$$\theta_{i+1} + \theta_i \pm (\tau_i \pm \delta_i) = \pi/2$$

which implies that since $\varepsilon < \frac{\pi}{6}$, under Q-superlinear convergence eventually at least every second $\theta_i > \frac{\pi}{6} > \varepsilon$, so that an update is made at that iteration. The remainder of the proof proceeds exactly as above. Since $\limsup_{i \to \infty} \mu_i \neq 0$ and each $\mu_i \ge \varepsilon$, the results of Dennis and Moré [6] show that (3.10) can only converge to zero if

$$\liminf_{i \to \infty} |\gamma_{j(i)} - 1| = 0.$$

The techniques used in the proof of Theorem 3.1, and much of the analysis, are independent of the update used. For instance, Schnabel [15] uses basically the same techniques to prove a version of Theorem 3.1 for a variant of Broyden's method. We can also use the above analysis with $\gamma_i \equiv 1$ to examine the direct prediction BFGS algorithm. For example, the interesting fact that

$$< x_{i+2} - x_*, y_i > = 0 \qquad (3.11)$$

holds for general quadratics of any dimension. ((3.11) is the general form of the condition $< x_{i+2}, s_i > = 0$ which was shown above in the case $f = x^T x$; it is obtained by expressing $x_{i+1} = x_* + \alpha_i s_i^{\perp} + \beta_i s_i$ for some $s_i^{\perp}$.) Therefore if Q-superlinear convergence is attained, $\lim_{i \to \infty} < s_{i+2}, \nabla^2 f(x_*) s_i > = 0$, which contradicts a conjecture we have sometimes heard that the final steps of a Q-superlinearly convergent quasi-Newton algorithm will be in almost the same direction. Asymptotic results like this are also readily extended to non-quadratic functions $f(x)$ with $\nabla^2 f(x_*)$ positive definite.

4. Relation of Luenberger, Oren and Spedicato's self-scaling update to Biggs' update

We now introduce the choices of self-sealing parameter γ_i suggested by Luenberger, Oren and Spedicato [10-13,19,20], and give a new derivation of a particular choice used by Brodlie [3] and Shanno and Phua [18]. Our derivation is based on a cubic model of f obtainable from the last iteration, and is closely related to the derivation of a successful update by Biggs [1]. It suggests a reason for the observed superior performance of Biggs' update on most problems.

Luenberger, Oren and Spedicato suggest choosing γ_i in (1.6) in the range

$$\gamma_i \ \varepsilon \left[\frac{<s_i,y_i>}{<y_i,H_iy_i>} \ , \ \frac{<s_i,H_i^{-1}s_i>}{<s_i,y_i>} \right] . \tag{4.1}$$

The original justification [12] stems from the performance of a perfect line search algorithm on quadratic functions. A rougher reason is that since $\nabla^2 f(x_i)\, s_i = y_i$ for quadratic f, it would be nice if

$$\gamma_i H_i y_i = s_i . \tag{4.2}$$

If (4.2) is true then $<y_i, \gamma_i H_i y_i> = <y_i, s_i>$ and $<s_i, \gamma_i y_i> = <s_i, H_i^{-1} y_i>$, which motivates the two limits of range (4.1). Brodlie [3] chooses

$$\gamma_i = <s_i,H_i^{-1}s_i> / <s_i,y_i> \tag{4.3}$$

because his algorithm stores and updates H_i^{-1} instead of H_i, making (4.3) the only readily available member of (4.1). Shanno and Phua [18] first consider $\gamma_i = (<s_i,H_i^{-1}s_i> / <y_i,H_iy_i>)^{1/2}$, the geometric mean of (4.1), but later consider using (4.3) at the first iteration and $\gamma_i = 1$ thereafter.

A different motivation of (4.3), following the derivation of a different update by Biggs [1], is as follows. From the step just taken we have four pieces of information about f: $f(x_i)$, $f(x_{i+1})$, $\nabla f(x_i)$, $\nabla f(x_{i+1})$. These can be used to approximate the restriction of f to the line connecting x_i and x_{i+1}, $\hat{f}(\tau) = f(x_i + \tau_i s_i)$, by the unique cubic $q(\tau)$ which interpolates $\hat{f}$ and $\hat{f}'$ at x_i and x_{i+1} $(\tau = 0,1)$. Note that $\hat{f}'(\tau) = \nabla f(x_i + \tau s_i)^T s_i$, $\hat{f}''(\tau) = s_i^T \nabla^2 f(x_i + \tau s_i)\, s_i$. Function $q(\tau)$ is found by standard interpolation techniques; in particular

$$q''(1/2) = s_i^T y_i . \tag{4.4}$$

Now since $\hat{H}_i = \gamma_i H_i$ is an intermediate matrix between approximating $\nabla^2 f(x_i)^{-1}$ and $\nabla^2 f(x_{i+1})^{-1}$, it seems reasonable that it approximate $\nabla^2 f(x_i + (1/2)\, s_i)^{-1}$. Since also

$$q''(1/2) \simeq s_i^T \nabla^2 f(x_i + (1/2) s_i)\, s_i,$$

this may be partially accomplished by setting

$$q''(1/2) = s_i^T \hat{H}_i^{-1} s_i = s_i^T H_i^{-1} s_i / \gamma_i$$

which from (4.4) results in the choice (4.3).

Biggs uses the same model to a different end. He reasons that $H_{i+1} y_i$ should be in the direction of s_i, but allows

$$H_{i+1} y_i = \rho_i s_i$$

for some parameter ρ_i. He then reasons that since H_{i+1} approximates $\nabla^2 f(x_{i+1})^{-1}$ and $q''(1) \simeq s_i \nabla^2 f(x_{i+1})\, s_i$, it would be good if $s_i^T H_{i+1}^{-1} s_i = q''(1)$. Since $H_{i+1}^{-1} s_i = y_i / \rho_i$, this is accomplished by setting $\rho_i = s_i^T y_i / q''(1)$ ($q''(1)$ is easily calculated). The update then used is the BFGS (1.3) with s_i replaced by $\rho_i s_i$ throughout.

To us, Biggs' use of the cubic model is more appealing than the self-scaling idea. This is because the step from x_i to x_{i+1}, and the cubic model, give information about f only in the direction s_i. Biggs uses this information to affect the model Hessian H_{i+1}^{-1} specifically as it acts on s_i. However, the self-scaling update, no matter how it is motivated, incorporates the information in direction s_i by making an equal change to the product which H_{i+1}^{-1} makes with any vector. This seems like a fundamental weakness of the self-scaling idea. Perhaps it accounts for its inferior performance to Biggs' method and the BFGS on most problems (Brodlie [3]).

We finally comment on the observation of Shanno and Phua [18] that when $H_0 = I$ (a common choice), the scaling $\hat{H}_0 = \gamma_0 H_0$, γ_0 given by (4.3), is helpful while scaling at subsequent iterations usually is not. While this initial scaling would seem to have the same weakness mentioned above, namely that information from one direction is used to influence $\hat{H}_0$ in all directions, this may help explain its success at the first iteration. This is because, if the range of eigenvalues of $\nabla^2 f(x_0)$ is not too broad and doesn't include 1, then multiplying $H_0^{-1} = I$ by a constant which scales it well in direction s_0 will probably also improve its accuracy in the other directions.

5. Relation of this analysis to the self-scaling methods of Luenberger, Oren and Spedicato

We now discuss the relation of our analysis in Sections 2 and 3 to the particular self-scaling BFGS updates proposed by Luenberger, Oren and Spedicato [10-13,19,20]. We consider separately the cases when the Hessian at the solution x_* is nonsingular and singular.

When $\nabla^2 f(x_*)$ is non-singular, successful quasi-Newton methods converge Q-superlinearly on a wide class of problems (Broyden, Dennis and Moré [5], Powell [14]). In Sections 2 and 3 we exhibited necessary and sufficient conditions for methods using the self-scaling BFGS update (1.6) to also attain this rate of convergence. They indicate that a sequence containing at least every second γ_i will need to approach 1, perhaps at a fast rate. We have not been able to show whether or not the choices (4.1) suggested by Luenberger, Oren and Spedicato obey either the necessary or the sufficient condition. (The problem seems to be obtaining a measure of how close $H_i^{-1} s_i$ is to y_i.) However, from Theorem 2.2 of Dennis and Moré [6], it is trivial to show that if the points generated by direct prediction iteration (2.1) converge Q-superlinearly to the minimum x_* of a function f obeying the assumptions of our Theorem 2.1 and all the H_i's are nonsingular, then for the most common choice $\gamma_i = < s_i, H_i^{-1} s_i > / < s_i, y_i >$, $\lim_{i \to \infty} \gamma_i$ must equal 1 regardless of the update used. This stems from Dennis and Moré's [6] result that Q-superlinear convergence requires

$$\lim \frac{\| (H_i^{-1} - \nabla^2 f(x_*)) s_i \|}{\| s_i \|} = 0 \tag{5.1}$$

and the commonly used lemma that under the assumptions of Theorem 2.1

$$\frac{\| \nabla^2 f(x_*) s_i - y_i \|}{\| s_i \|} \leq \ell \max \{ \| x_i - x_* \| , \| x_{i+1} - x_* \| \} \tag{5.2}$$

(see e.g., Broyden, Dennis and Moré [5]). If in addition $\{ \| H_i \| \}$ is uniformly bounded, then (5.1) and (5.2) show that given Q-superlinear convergence, any other γ_i in the range (4.1) must also converge to 1. Therefore, if the self-scaling BFGS update with any γ_i given by (4.1) is to be successful on problems with $\nabla^2 f(x_*)$ nonsingular, it must asymptotically become the same as the BFGS.

The situation on problems with $\nabla^2 f(x_*)$ singular is markedly different. It is on these problems that self-scaling methods have been successful in practice (Brodlie [3], Shanno and Phua [18]). On such problems, neither quasi-Newton nor Newton's method achieve better than linear convergence in general. Furthermore, we show in Example 5.1 below that any γ_i in the range (4.1) does not in general converge to 1 on such problems (using either direct prediction or the line search of Theorem 2.3). Therefore the self-scaling BFGS update is expected to differ from the BFGS on singular problems. This does not indicate whether it will be better or worse, but at least that we are considering fundamentally different updates. Shanno and Phua [18] give some justification why the self-scaling updates should be helpful on a special class of singular problems, the homogeneous functions. In our example, f is homogeneous as well as being singular.

Example 5.1 may be somewhat misleading. For ease of analysis, we use a one dimensional function f. So while we show that γ_i given by (4.3) (or any $\gamma_i \in (4.1)$, as they are equivalent in this case) does not converge to 1, this has no impact on the points generated by the algorithm in this case, because H_{i+1} is totally determined by $H_{i+1}y_i = s_i$. However, it should be clear from this example that $\gamma_i \in (4.1)$ will in general also not converge to 1 when $\nabla^2 f(x_*)$ is singular and $n > 1$, and then the choices of γ_i will affect the matrices H_i and points x_i which are produced.

Example 5.1 Suppose the direct prediction quasi-Newton iteration (2.1), using update (1.6) with γ_i given by (4.3) is applied to $f : R \to R$, $f = x^{2k}$, $k > 1$ a positive integer. Let r be the root of $r^{2k-1} + r^{2k-2} = 1$ in the interval (0,1). If $x_0 \neq 0$ and $H_0 = (1-r) / 2kx_0^{2k-2}$, then at each iteration

$$x_{i+1} = r \cdot x_i,\ \gamma_i = 1/r^{2k-2} \triangleq \gamma,\ H_{i+1} = (1-r)/2\,k\,x_{i+1}^{\ 2k-2}.$$

For example, if $k = 2$, $r \simeq .755$ and $\gamma \simeq 1.75$. As $k \to \infty$, $r \to 1$ and $\gamma \to 2$, both from below.

Furthermore, if line search iteration (2.6-7) is used instead of (2.1), then if

$$\beta_1 \leq \frac{1-r^{2k}}{2k(1-r)} \triangleq \hat{\beta}_1\ (k),$$

$$\beta_2 \geq r^{2k-1} \qquad \triangleq \hat{\beta}_2\ (k),$$

the choice $\lambda_i = 1$ always satisfies the line search condition. Since $\hat{\beta}_1(k) > 1/2 > \hat{\beta}_2(k)$ for all $k > 1$, this means that there exist permissible values of β_1 and β_2 in (2.7) such that if the line search is implemented to use $\lambda_i = 1$ whenever it satisfies (2.7), then the above also applies to line search iteration (2.6-7).

Proof: The first part of the example, concerning the direct prediction method, follows from straightforward algebra. Asymptotically, $(1-r) \varepsilon [1/(4k-2), 1/(4k-4)]$ and $r^{2k-2} \simeq r^{2k-1} \simeq 1/2$.

For the line search result, we just need to show that $\hat{\beta}_1(k) > 1/2 > \hat{\beta}_2(k)$. We use the fact that since $r^{2k-1} + r^{2k-2} = 1$ and $r < 1$,

$$r^i > 1/2 \text{ for all } i \le 2k-2 \tag{5.3}$$

and $r^{2k-1} < 1/2$. Thus $\hat{\beta}_2(k) < 1/2$ is immediate. Now express

$$\hat{\beta}_1(k) = (1/2k) \sum_{j=0}^{k-1} (r^j + r^{(2k-1)-j}),$$

and from (5.3), $r^j + r^{(2k-1)-j} > 1$ for all $j \varepsilon [0,k-1]$ so that $\hat{\beta}_1(k) > 1/2$. For $k = 2$, $\hat{\beta}_1(k) \simeq .68$, $\hat{\beta}_2(k) \simeq .43$; asymptotically, $\hat{\beta}_1(k) \to 1$ and $\hat{\beta}_2(k) \to 1/2$.

Example 5.1 can be shown to indicate the asymptotic value of γ_i when $H_0 \neq (1-r)/2kr^{2k-2}$. We also expect that it is indicative of the asymptotic behavior of γ_i on other functions f with $\nabla^2 f(x_*)$ singular. Thus it gives good indication that the self-scaling BFGS and normal BFGS updates differ, even asymptotically, on singular problems.

6. Summary and concluding remarks

We have shown that for the self-scaling BFGS method to be Q-superlinearly convergent under normal assumptions (including $\nabla^2 f$ nonsingular at the solution x_*) it is necessary that $\liminf_{i \to \infty} \gamma_i = 1$ and sufficient that $\sum_{i=0}^{\infty} |\gamma_i - 1|$ is bounded by a suitable constant. It is an open question whether the choices of γ_i suggested by Luenberger, Oren and Spedicato [10-13,19,20] and used by Brodlie [3] and Shanno and Phua [18] satisfy either condition. However, from our theory it is seen that if any self-scaling BFGS update in use is to be as successful as the BFGS on nonsingular problems, then asymptotically it must be the same as the BFGS. Along with the computational results of Broyden and Shanno and Phua, which show the self-scaling update faring less well than the BFGS on problems with $\nabla^2 f(x_*)$ nonsingular, this discourages the use of self-scaling updates on nonsingular problems. Another possible disadvantage of self-scaling methods is indicated in Section 4: they use information about the effect of the Hessian in one direction to alter the Hessian approximation in all directions.

We have also considered the behavior of the self-scaling update on problems with $\nabla^2 f(x_*)$ singular, where it has proven attractive in practice, and where only linear convergence can be expected in general. We give an example of a class of such problems in which the commonly used self-scaling parameter γ_i does not converge to 1. Our example indicates that the normal and self-scaling BFGS updates are likely to be fundamentally different on such problems.

References

[1] M. C. Biggs, "A note on minimization algorithms which make use of non-quadratic properties of the objective function", Journal of the Institute of Mathematics and its Applications 12 (1973) 337-8.

[2] K. W. Brodlie, "Unconstrained minimization", in: D. Jacobs, ed., The State of the Art in Numerical Analysis (Academic Press, London, 1977).

[3] K. W. Brodlie, "An assessment of two approaches to variable metric methods", Mathematical Programming 12 (1977) 344-355.

[4] C. G. Broyden, "The convergence of a class of double rank minimization algorithms", Journal of the Institute of Mathematics and its Applications 6 (1970) 76-90.

[5] C. G. Broyden, J. E. Dennis and J. J. Moré, "On the local and superlinear convergence of quasi-Newton methods", Journal of the Institute of Mathematics and its Applications 12 (1973) 223-246.

[6] J. E. Dennis and J. J. Moré, "A characterization of superlinear convergence and its application to quasi-Newton methods", Mathematics of Computation 28 (1974) 549-560.

[7] J. E. Dennis and J. J. Moré, "Quasi-Newton methods, motivation and theory", SIAM Review 19 (1977) 46-89.

[8] R. Fletcher, "A new approach to variable metric methods ", Computer Journal 13 (1970) 317-322.

[9] D. Goldfarb, "A family of variable-metric methods derived by variational means", Mathematics of Computation 24 (1970) 23-26.

[10] S. S. Oren, "Self-scaling variable metric (SSVM) algorithms II: implementation and experiments", Management Science 20 (1974) 863-874.

[11] S. S. Oren, "On the selection of parameters in self-scaling variable metric algorithms", Mathematical Programming 7 (1974) 351-367.

[12] S. S. Oren and D. G. Luenberger, "Self-scaling variable metric (SSVM) algorithms I: criteria and sufficient conditions for scaling a class of algorithms", Management Science 20 (1974) 845-862.

[13] S. S. Oren and E. Spedicato, "Optimal conditioning of self-scaling variable metric algorithms", Mathematical Programming 10 (1976) 70-90.

[14] M. J. D. Powell, "Some global convergence properties of a variable metric algorithm for minimization without exact line searches", in: R. W. Cottle and C. E. Lemke, eds., Nonlinear Programming (SIAM-AMS proceeding, vol. IX, 1976).

[15] R. B. Schnabel, "Analyzing and improving quasi-Newton methods for unconstrained optimization", Ph.D. thesis, Department of Computer Science, Cornell University, Ithaca, New York (1977).

[16] R. B. Schnabel, "Q-superlinear convergence of Biggs' method and related methods for unconstrained optimization", report no. CU-CS-133-78, Department of Computer Science, University of Colorado, Boulder, Colorado (1978).

[17] D. F. Shanno, "Conditioning of quasi-Newton methods for function minimization", Mathematics of Computation 24 (1970) 647-656.

[18] D. F. Shanno and K. H. Phua, "Matrix conditioning and nonlinear optimization", Mathematical Programming 14 (1978) 149-160.

[19] E. Spedicato, "Computational experience with quasi-Newton algorithms for minimization problems of moderately large size", report no. CISE-N-175, Segrate, Italy (1975).

[20] E. Spedicato, "A variable metric method derived from invariancy to nonlinear scaling", Journal of Optimization Theory and its Applications 20 (1976) 315-329.

[21] P. Wolfe, "Convergence conditions for ascent methods", SIAM Review 11 (1969) 226-235.