

COMPUTING MODIFIED NEWTON DIRECTIONS USING A PARTIAL CHOLSKY FACTORIZATION*

A. FORSGREN[†], P. E. GILL[‡], AND W. MURRAY[§]

Abstract. The effectiveness of Newton’s method for finding an unconstrained minimizer of a strictly convex twice continuously differentiable function has prompted the proposal of various *modified* Newton methods for the nonconvex case.

Line search modified Newton methods utilize a linear combination of a descent direction and a direction of negative curvature. If these directions are *sufficient* in a certain sense, and a suitable linesearch is used, the resulting method will generate limit points that satisfy the second-order necessary conditions for optimality.

We propose an efficient method for computing a descent direction and a direction of negative curvature that is based on a partial Cholesky factorization of the Hessian. This factorization not only gives theoretically satisfactory directions, but also requires only a *partial* pivoting strategy, i.e., the equivalent of only two rows of the Schur complement need be examined at each step.

Key words. Unconstrained minimization, modified Newton method, descent direction, negative curvature, Cholesky factorization.

AMS subject classifications. 49D37, 65F05, 65K05, 90C30

1. Introduction. We consider the unconstrained minimization of a twice continuously differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. If f is strictly convex, the excellent local convergence properties of Newton’s method make it one of the most effective methods for minimization (see, e.g., Ortega and Rheinboldt [19]).

In the non-convex case, various *modified* Newton methods have been proposed that ensure convergence from an arbitrary starting point. Here we focus on the class of *linesearch* modified Newton methods (for a complete discussion of modified Newton methods and their relative merits, see, e.g., Shultz *et al.* [21], Dennis and Schnabel [4]). Linesearch modified Newton methods generate a sequence $\{x_k\}_{k=0}^{\infty}$ of improving estimates of a local minimizer. At iteration k , a linesearch is performed along a path formed from a linear combination of two directions s_k and d_k , where either s_k or d_k can be zero. The directions s_k and d_k are chosen such that $g_k^T s_k \leq 0$ and $d_k^T H_k d_k \leq 0$, where g_k and H_k denote the gradient $\nabla f(x)$ and Hessian $\nabla^2 f(x)$ evaluated at x_k . (Implicitly, we also assume the condition $g_k^T d_k \leq 0$, which can be imposed with a trivial sign change of d_k .) Each nonzero s_k satisfies $g_k^T s_k < 0$ and is known as a *descent direction*. Each nonzero d_k satisfies $d_k^T H_k d_k < 0$ and is known as a *direction of negative curvature*. If d_k is nonzero, H_k must have at least one negative eigenvalue. (Henceforth we will sacrifice precision for the sake of brevity and refer to the sequences $\{s_k\}$ and $\{d_k\}$ as sequences of “descent directions” and “directions of negative curvature”.) Linesearch methods of this type have been proposed by Gill

*Received by the editors March 9, 1993; accepted for publication (in revised form) November 19, 1993. Research supported by the Department of Energy Contract DE-FG03-92ER25117, the National Science Foundation Grants DDM-9204208, DDM-9204547, the Office of Naval Research Grant N00014-90-J-1242, the Göran Gustafsson Foundation, the Royal Swedish Academy of Sciences, the Swedish National Board for Technical Development, and the Swedish Natural Science Research Council.

[†]Optimization and Systems Theory, Department of Mathematics, Royal Institute of Technology, S-100 44 Stockholm, Sweden (andersf@math.kth.se).

[‡]Department of Mathematics, University of California, San Diego, La Jolla, California, 92093-0012 (peg@optimal.ucsd.edu).

[§]Department of Operations Research, Stanford University, Stanford, California 94305-4022 (walter@sol-walter.stanford.edu).

and Murray [10], Fletcher and Freeman [7], McCormick [16], Mukai and Polak [18], Kaniel and Dax [14], and Goldfarb [11].

Moré and Sorensen [17] have shown that if: (i) a modified Newton method is used in conjunction with a suitable linesearch; and (ii) the directions s_k and d_k are *sufficient* in the sense that the sequences $\{s_k\}$ and $\{d_k\}$ are bounded and satisfy

$$g_k^T s_k \rightarrow 0 \implies g_k \rightarrow 0 \quad \text{and} \quad s_k \rightarrow 0, \quad (1.1a)$$

and

$$d_k^T H_k d_k \rightarrow 0 \implies \min\{\lambda_{\min}(H_k), 0\} \rightarrow 0 \quad \text{and} \quad d_k \rightarrow 0, \quad (1.1b)$$

then every limit point of the resulting sequence $\{x_k\}_{k=0}^\infty$ will satisfy the second-order necessary conditions for optimality.

It has been observed in practice that the number of iterates at which the Hessian is positive definite is large compared to the total number of iterations. Since linesearch methods revert to Newton's method when the Hessian is sufficiently positive definite, it would seem sensible to use a modified Newton method based on the most efficient method for solving a symmetric positive-definite system. This is the motivation for the modified Cholesky factorization proposed by Gill and Murray [10]. However, it has been shown by Moré and Sorensen [17] that this factorization may not give directions of negative curvature that are sufficient in the sense of (1.1b). This paper is motivated by the need for an algorithm with the efficiency and simplicity of the Cholesky factorization, but with the guarantee of convergence when used with a suitable linesearch. It is shown in Section 3 that a *partial* Cholesky factorization can give search directions that are sufficient in the sense of (1.1).

To simplify the notation, we will drop the subscript k when referring to the quantities g_k , H_k , s_k and d_k at a specific iteration. Unless otherwise stated, $\|\cdot\|$ refers to the vector two-norm or its induced matrix norm. The vector e_j denotes the j th unit vector whose dimension is determined by the context.

2. The partial Cholesky factorization. The partial Cholesky factorization of H is a variant of the standard Cholesky factorization with diagonal pivoting. The algorithm is stated in *outer-product form*, where the Schur complement associated with the unfactorized part of H is updated explicitly at each step (see, e.g., Golub and Van Loan [12, page 143] and Higham [13]).

At each step, the largest diagonal¹ is selected as pivot and is used to eliminate a row and column from the Schur complement. The algorithm continues until either all the matrix has been factorized or the pivot is considered unacceptable. The final factors are therefore uniquely determined by the rule used to accept the pivot (i.e., the rule used to terminate the elimination). Termination is controlled by a preassigned scalar parameter ν ($0 < \nu < 1$). A pivot is acceptable if it is both positive and larger in absolute value than ν times the off-diagonal of largest magnitude in the pivot row and column. At each step, the determination of an acceptable pivot requires the examination of the diagonals and a single row of the Schur complement. (For a similar scheme in the context of quadratic programming, see Casas and Pola [3].)

It will be shown below that once a pivot is deemed unacceptable (and hence the factorization is terminated), a suitable direction of negative curvature can be determined from the elements of the remaining Schur complement.

¹In the event of a tie, the largest diagonal with least index is selected.

Let P denote the permutation matrix representing the symmetric interchanges performed during the factorization. If n_1 denotes the number of steps needed before termination, the factorization implicitly identifies a leading $n_1 \times n_1$ positive-definite submatrix of the permuted matrix $P^T H P$. In terms of a partition H_{11} , H_{12} , H_{21} and H_{22} of $P^T H P$, we have

$$\begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix} = \begin{pmatrix} L_{11} & \\ & I \end{pmatrix} \begin{pmatrix} B_1 & \\ & B_2 \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ & I \end{pmatrix}, \quad (2.1)$$

where L_{11} is unit lower triangular and B_1 is a positive-definite diagonal matrix. The submatrix H_{11} is positive definite, and $H_{11} = L_{11} B_1 L_{11}^T$ is its usual Cholesky factorization obtained using diagonal pivoting. The factorization may be written briefly as $H = L B L^T$, where L is a row-permuted lower-triangular matrix with

$$L = P \begin{pmatrix} L_{11} & \\ & I \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} B_1 & \\ & B_2 \end{pmatrix}. \quad (2.2)$$

We will use n_2 to denote the size of H_{22} , so that $n_1 + n_2 = n$. A “pseudo-matlab” version of the partial Cholesky algorithm is given in Algorithm 2.1.

```
% PARTCHOL Partial Cholesky factorization routine for a real symmetric matrix H.
%           [L,B,perm,n1] = partchol(H)
%           forms a permutation perm, a unit lower-triangular matrix
%           L(perm,:) and a block diagonal matrix B such that
%           LBL'=H
%           using the partial Cholesky factorization with diagonal pivoting.
%           The size of the positive-definite principal submatrix obtained
%           in the factorization is n1.
function [L,B,perm,n1] = partchol(H)
n = length(H);
perm = 1:n;
B = H;
L = zeros(n);
nu = (0,1);
k = 1;
n1 = 0;
while k <= n
    [mu_r,r] = max([zeros(1,k-1) diag(B(k:n,k:n))]);
    if k < n
        mu_pr = max(abs(B(r,[1:r-1 r+1:n])));
    else
        mu_pr = 0;
    end
    if mu_r > 0 and mu_r >= nu * mu_pr
        n1 = k;
        perm([k r]) = perm([r k]);
        B([k r],:) = B([r k],:);
        B(:,[k r]) = B(:,[r k]);
        L(perm(k:n),k) = B(k:n,k)/B(k,k);
    end
    if k < n
```

```

    B(k+1:n, k+1:n) = B(k+1:n, k+1:n) - L(perm(k+1:n), k) × B(k, k+1:n);
    B(k+1:n, k) = zeros(n-k, 1);
    B(k, k+1:n) = zeros(1, n-k);
end
k = k+1;
else
    L(perm(k:n), k:n) = eye(n-k+1);
    k = n+1;
end
end
end

```

ALGORITHM 2.1. *An algorithm for the partial Cholesky factorization*

The curvature along any direction d computed from the partial Cholesky factorization is related to the magnitude of the smallest eigenvalue of the Schur complement B_2 . The following lemma relates the smallest eigenvalue of B_2 to the smallest eigenvalue of H .

LEMMA 2.1. *Let H be a symmetric $n \times n$ matrix with at least one negative eigenvalue. Let the partial Cholesky factorization of H be denoted by $H = LBL^T$, where $P^T H P$ is partitioned as in (2.1). Then*

$$\lambda_{\min}(B_2) \leq \lambda_{\min}(H) \quad \text{and} \quad B_2 = Y^T H Y,$$

where

$$Y = P \begin{pmatrix} -L_{11}^{-T} L_{21}^T \\ I \end{pmatrix} = P \begin{pmatrix} -H_{11}^{-1} H_{12} \\ I \end{pmatrix}. \quad (2.3)$$

Proof. The inequality $\lambda_{\min}(B_2) \leq \lambda_{\min}(H)$ can be established using the identity

$$\begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix} = \begin{pmatrix} 0 & \\ & B_2 \end{pmatrix} + \begin{pmatrix} L_{11} \\ L_{21} \end{pmatrix} B_1 \begin{pmatrix} L_{11}^T & L_{21}^T \end{pmatrix}, \quad (2.4)$$

which is a rearrangement of the factorization (2.1). The eigenvalues of H and $P^T H P$ are identical. Moreover, the positive-definiteness of B_1 implies that the second term on the right-hand-side of (2.4) is positive semidefinite. Since the eigenvalues of $P^T H P$ cannot increase on subtraction of a positive semidefinite matrix, it must follow that $\min\{0, \lambda_{\min}(B_2)\} \leq \lambda_{\min}(H)$ (see e.g., Golub and Van Loan [12, page 411]). From the assumption $\lambda_{\min}(H) < 0$, we conclude that $\lambda_{\min}(B_2) \leq \lambda_{\min}(H)$, as required.

To show that the matrix Y (2.3) is well defined, it is sufficient to verify that $H_{11}^{-1} H_{12} = L_{11}^{-T} L_{21}^T$. This is an immediate consequence of multiplying the partitioned right-hand-side matrix from (2.1) to obtain $H_{11} = L_{11} B_1 L_{11}^T$ and $H_{12} = L_{11} B_1 L_{21}^T$.

Finally, the identity $Y^T H Y = B_2$ may be verified by expressing $L^{-1} H L^{-T} = B$ in the partitioned form

$$\begin{pmatrix} L_{11}^{-1} & \\ -L_{21} L_{11}^{-1} & I \end{pmatrix} \begin{pmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{pmatrix} \begin{pmatrix} L_{11}^{-T} & -L_{11}^{-T} L_{21}^T \\ & I \end{pmatrix} = \begin{pmatrix} B_1 & \\ & B_2 \end{pmatrix},$$

from which the result follows. \square

Note that the matrix Y (2.3) consists of the last n_2 columns of L^{-T} . Our analysis requires bounds on the norms of Y , L and L^{-1} , which are provided by the following lemma given by Higham [13].

LEMMA 2.2. *Let H be factorized using the partial Cholesky factorization described in Algorithm 2.1. If $P^T H P$ is partitioned as in (2.1), then*

- (a) $\|L_{11}^{-T} L_{21}^T\| \leq \frac{1}{\nu} \sqrt{(n - n_1)(4^{n_1} - 1)}/3$;
- (b) $\|L_{11}^{-T} L_{21}^T e_i\| \leq \frac{1}{\nu} \sqrt{(4^{n_1} - 1)}/3, \quad i = 1, 2, \dots, n_2$;
- (c) $\|L\| \leq n/\nu$;
- (d) $\|L^{-1}\| \leq n2^{n_1-1}/\nu$.

Proof. Part (a) follows immediately from Lemma 9.4 of Higham [13] and the fact that the elements of L_{21} are bounded in absolute value by $1/\nu$. Part (b) is a consequence of part (a), since $L_{21}^T e_i$ is an n_2 -vector whose elements are bounded in absolute value by $1/\nu$. Part (c) follows from the fact that all elements of L are bounded by $1/\nu$ in absolute value. Similarly, part (d) is a consequence of the fact that all elements of L^{-1} are bounded by $2^{n_1-1}/\nu$ (see Higham [13] for details). \square

2.1. Computation of the descent direction. We now discuss the application of the partial Cholesky factorization to the calculation of a descent direction s_k satisfying (1.1a). Let \bar{B} be any positive-definite modification of B , i.e., \bar{B} is a positive-definite matrix with $\|B - \bar{B}\|$ “small” and $\bar{B} = B$ when B is sufficiently positive definite. There are many choices for \bar{B} —for example, consider the block-diagonal matrix $\bar{B} = \text{diag}(B_1, I)$, where I is the identity matrix of order n_2 . With this definition, when $n_1 = n$ and H is sufficiently positive definite, $\bar{B}_1 = B_1$ and s satisfies the usual Newton equations $Hs = -g$.

LEMMA 2.3. *Let H be factorized using the partial Cholesky factorization described in Algorithm 2.1 and assume that $P^T H P$ is partitioned as in (2.1). Let \bar{B} be a positive-definite modification of B , and let s satisfy*

$$L\bar{B}L^T s = -g. \quad (2.5)$$

Then,

$$-g^T s \geq \frac{\nu^2}{n^2 \lambda_{\max}(\bar{B})} \|g\|^2 \quad \text{and} \quad \|s\| \leq \frac{n^2 4^{n_1-1}}{\nu^2 \lambda_{\min}(\bar{B})} \|g\|.$$

Proof. From the definition of s in (2.5) we have

$$s = -L^{-T} \bar{B}^{-1} L^{-1} g. \quad (2.6)$$

Premultiplying (2.6) by g^T gives

$$-g^T s = g^T L^{-T} \bar{B}^{-1} L^{-1} g \geq \frac{1}{\|L\|^2 \lambda_{\max}(\bar{B})} g^T g,$$

and the required lower bound on $-g^T s$ follows from part (c) of Lemma 2.2. To obtain the bound on $\|s\|$ we derive the inequality $\|s\| \leq \lambda_{\max}(\bar{B}^{-1}) \|L^{-1}\|^2 \|g\|$, by taking norms of both sides of (2.6) and using norm inequalities. The required upper bound follows from part (d) of Lemma 2.2. \square

2.2. Computation of the direction of negative curvature. The formula for d is derived from a method for computing directions of negative curvature in quadratic programming (see Forsgren *et al.* [9]). The approach is based on the observation that, in the positive-definite case, the Newton direction is a minimizer of a quadratic model with gradient g and Hessian H . In particular, the Newton direction can be found by a quadratic programming algorithm that minimizes the model function while successively releasing variables from temporarily fixed values. This analogy can be extended to the indefinite case, where the variables corresponding to H_{22} are temporarily fixed at their current values, and a direction of negative curvature is defined by releasing either one or two of the fixed variables. This scheme corresponds to using a direction of negative curvature that is a multiple of either y_i or $y_i \pm y_j$, where y_i and y_j denote columns i and j of the matrix Y (2.3). The following lemma shows how the indices i and j are determined from the elements of $B_2 = Y^T H Y$.

LEMMA 2.4. *On termination of the partial Cholesky factorization with diagonal pivoting, let $P^T H P$ be partitioned as in (2.1). If $n_1 = n$, define $d = 0$. Otherwise, if $n_1 < n$, define d as follows. Given $\rho = \max_{i > n_1, j > n_1} |b_{ij}|$ and any pair of indices q ($q > n_1$) and r ($r > n_1$) such that $|b_{qr}| = \rho$, let d be the solution of*

$$L^T d = \sqrt{\rho} v, \quad \text{where} \quad v = \begin{cases} e_q & \text{if } q = r, \\ \frac{1}{\sqrt{2}}(e_q - \text{sgn}(b_{qr})e_r) & \text{otherwise.} \end{cases}$$

Then, if $\lambda_{\min}(H) \geq 0$, then $d = 0$. Otherwise, if $\lambda_{\min}(H) < 0$, then

$$-\frac{1}{n_2} \lambda_{\min}(H) \leq d^T d \leq -\frac{1}{1-\nu} \left(1 + \frac{2(4^{n_1} - 1)}{3\nu^2}\right)^2 \lambda_{\min}(H)$$

and

$$\frac{d^T H d}{d^T d} \leq \frac{3\nu^2(1-\nu)}{n_2(3\nu^2 + 2(4^{n_1} - 1))} \lambda_{\min}(H).$$

Proof. If $n_1 = n$, then $\lambda_{\min}(H) > 0$, and the lemma holds from the definition $d = 0$. For the remainder of the proof, assume that $n_1 < n$.

First, it is necessary to show that $\gamma \leq \nu\rho$, where $\gamma = \max\{\{\max_{i > n_1} b_{ii}\}, 0\}$. If the factorization terminates with $\gamma = 0$, the inequality $\gamma \leq \nu\rho$ is trivially satisfied. If the factorization terminates with $\gamma > 0$, there exists an index t ($t > n_1$) such that $b_{tt} = \gamma$. Since γ must be an unacceptable pivot, we can infer that $\gamma < \nu \max_{i \neq t, i > n_1} |b_{it}|$. Consequently, if $n_1 < n$, it must hold that $\gamma \leq \nu\rho$.

Define $\tilde{d} = P^T d$. Let \tilde{d}_1 and v_1 denote the first n_1 components of \tilde{d} and v respectively. Similarly, let \tilde{d}_2 and v_2 denote the last n_2 components of \tilde{d} and v . The definitions of \tilde{d} and v imply that $\|v_1\| = 0$, $\|v_2\| = 1$, and $\tilde{d}_2 = \sqrt{\rho} v_2$. Therefore,

$$d^T d = \tilde{d}^T \tilde{d} = \tilde{d}_1^T \tilde{d}_1 + \tilde{d}_2^T \tilde{d}_2 \geq \rho v_2^T v_2 = \rho. \quad (2.7)$$

Similarly, the definition of d and (2.2) imply that

$$d^T d \leq (1 + \|L_{11}^{-T} L_{21}^T v_2\|^2) \rho \leq \left(1 + \frac{2(4^{n_1} - 1)}{3\nu^2}\right) \rho, \quad (2.8)$$

where the last inequality follows from Lemma 2.2. Combining (2.7) and (2.8) yields

$$\rho \leq d^T d \leq \left(1 + \frac{2(4^{n_1} - 1)}{3\nu^2}\right) \rho. \quad (2.9)$$

Consider the case $\rho = 0$, which is equivalent to H being positive semidefinite and singular with $\lambda_{\min}(H) = 0$. In this case, (2.9) implies $d = 0$, as required.

Now assume that $\rho > 0$. First, if $q = r$, then $|b_{qq}| = \rho$. Since $b_{qq} \leq \gamma \leq \nu\rho < \rho$, it must hold that $b_{qq} = -\rho$, and from the definition of d we obtain the bound

$$d^T H d = \rho b_{qq} = -\rho^2 \leq -(1 - \nu)\rho^2. \quad (2.10)$$

Alternatively, if $q \neq r$, then the definition of d yields

$$d^T H d = \frac{\rho}{2}(b_{qq} + b_{rr} - 2|b_{qr}|) \leq \rho(\gamma - \rho) \leq -(1 - \nu)\rho^2, \quad (2.11)$$

where the inequalities follow from the conditions $b_{qq} \leq \gamma$, $b_{rr} \leq \gamma$ and $\rho \geq \gamma/\nu$.

Since the magnitude of every element in B_2 is bounded by ρ , the Gershgorin circle theorem and Lemma 2.1 imply

$$\rho \geq -\frac{1}{n_2} \lambda_{\min}(B_2) \geq -\frac{1}{n_2} \lambda_{\min}(H). \quad (2.12)$$

Combining (2.9), (2.10), (2.11) and (2.12) we obtain

$$\frac{d^T H d}{d^T d} \leq -\frac{3\nu^2(1 - \nu)}{3\nu^2 + 2(4^{n_1} - 1)} \rho \leq \frac{3\nu^2(1 - \nu)}{n_2(3\nu^2 + 2(4^{n_1} - 1))} \lambda_{\min}(H), \quad (2.13)$$

as required.

Since, by definition, $\lambda_{\min}(H) \leq d^T H d / d^T d$, the left-most inequality of (2.13) gives an upper bound on ρ , which in conjunction with (2.9) and (2.12) give the bounds on $d^T d$ as

$$-\frac{1}{n_2} \lambda_{\min}(H) \leq d^T d \leq -\frac{1}{1 - \nu} \left(1 + \frac{2(4^{n_1} - 1)}{3\nu^2}\right)^2 \lambda_{\min}(H).$$

□

This lemma gives a relation between the curvature along d and the smallest eigenvalue of H , which is the “best possible” curvature. The bound is exponential in n_1 , but the computational experiments discussed below imply that the bound is unlikely to be tight in practice. However, as in Higham [13], we observe that there do exist matrices whose bound is “almost” tight. For given n ($n \geq 3$) and θ , define $L(\theta)$ and $B(\theta)$ as

$$L(\theta) = \begin{pmatrix} 1 & & & & & \\ -c & 1 & & & & \\ \vdots & \vdots & \ddots & & & \\ -c & -c & \cdots & 1 & & \\ -c & -c & \cdots & -c & 1 & \\ -c & -c & \cdots & -c & 0 & 1 \end{pmatrix}, \quad B(\theta) = \begin{pmatrix} 1 & & & & & \\ & s^2 & & & & \\ & & \ddots & & & \\ & & & s^{2(n-3)} & & \\ & & & & 0 & -1 \\ & & & & -1 & 0 \end{pmatrix},$$

where $c = \cos \theta$ and $s = \sin \theta$. Define $H(\theta) = L(\theta)B(\theta)L(\theta)^T$. If $\theta = 0$, it is shown in Lemma A.1 of Appendix A that $\lambda_{\min}(H(0)) = -\frac{1}{2}(\sqrt{n^2 + 2n - 7} - n + 1)$, where

$$-1 \leq \lambda_{\min}(H(0)) \leq -1 + \frac{4}{n+1}.$$

If $\theta = 0$, the partial Cholesky factorization with diagonal pivoting gives $n_1 = 1$. If $d(\theta)$ denotes the direction of negative curvature associated with $H(\theta)$, we obtain

$$\frac{d(0)^T H(0) d(0)}{d(0)^T d(0)} = -\frac{1}{3}, \quad (2.14)$$

and $d(0)$ is a satisfactory direction of negative curvature. However, if θ is nonzero, it follows from the analysis of Higham [13] that the partial Cholesky factorization with diagonal pivoting will define $L(\theta)$ and $B(\theta)$ as factors with $n_1 = n - 2$ for all $\theta \neq 0$. Moreover,

$$\lim_{\theta \rightarrow 0} \frac{d(\theta)^T H(\theta) d(\theta)}{d(\theta)^T d(\theta)} = -\frac{3}{1 + 2 \times 4^{n-2}},$$

and for θ near zero, the curvature along $d(\theta)$ is close to the worst possible value predicted by Lemma 2.4 (see Higham [13] for the details). This “pathological” example arises because the principal submatrix of order $n - 2$ of $H(\theta)$ is positive definite but arbitrarily close to being singular so that $\|H_{11}^{-1}H_{12}\|$ (or equivalently $\|L_{11}^{-T}L_{21}^T\|$) is very large. This is reflected in arbitrarily small pivot elements.

A numerical experiment was devised to investigate if the bound of Lemma 2.4 is likely to be sharp in practice. Using MATLAB 4.1 [15] on a SUN SPARCstation, the direction d of Lemma 2.4 was computed for a set of random 50×50 indefinite symmetric matrices with specified spectra. Each matrix was of the form $Q\Lambda Q^T$, with Q a random orthogonal matrix and $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$ a diagonal matrix with t negative elements ($1 \leq t \leq 20$, see below). The matrix Q was obtained from the QR -factorization of a 50×50 matrix whose elements were taken from an independent normal distribution with zero mean and unit variance. The method for generating the spectrum Λ was similar to that proposed by Higham [13]. The eigenvalues were chosen to allow variation of the condition number, the magnitude of the most negative eigenvalue and the number of negative eigenvalues. In this way, a wide variety of matrices were included in the test set.

We used two distributions of the eigenvalues:

$$\lambda_i = \begin{cases} 1 & i = 1, \dots, n - t; \\ -\alpha^{1/(n+1-i)} & i = n - t + 1, \dots, n, \end{cases} \quad (\text{the } \alpha \text{ distribution}),$$

and

$$\lambda_i = \begin{cases} \beta^{i-1} & i = 1, \dots, n - t; \\ -\beta^{i-1} & i = n - t + 1, \dots, n, \end{cases} \quad (\text{the } \beta \text{ distribution}),$$

α ($0 < \alpha \leq 1$) and β ($0 < \beta \leq 1$) being used to vary $\kappa_2(H) = |\lambda_1/\lambda_n|$. For a particular value of ν , 100 different matrices were generated for each distribution by taking all combinations of $t = \{1, 2, \dots, 20\}$, and $\kappa_2(H) = \{1, 10^3, 10^6, 10^9, 10^{12}\}$. These runs were repeated for each of the ν -values $\{\sqrt{\epsilon}, 0.05, 0.10, \dots, 0.95, 1 - \sqrt{\epsilon}\}$, where ϵ

denotes the machine precision ($\epsilon \approx 2.2 \times 10^{-16}$). A new Q was generated for each H . Figure 1 gives the outcome of the computational experiment for the two eigenvalue distributions. The three lines depict the maximum, mean, and minimum values of the ratio r of $d^T H d / d^T d$ to $\lambda_{\min}(H)$. Each “+” represents the value of r for a particular value of the parameter ν .

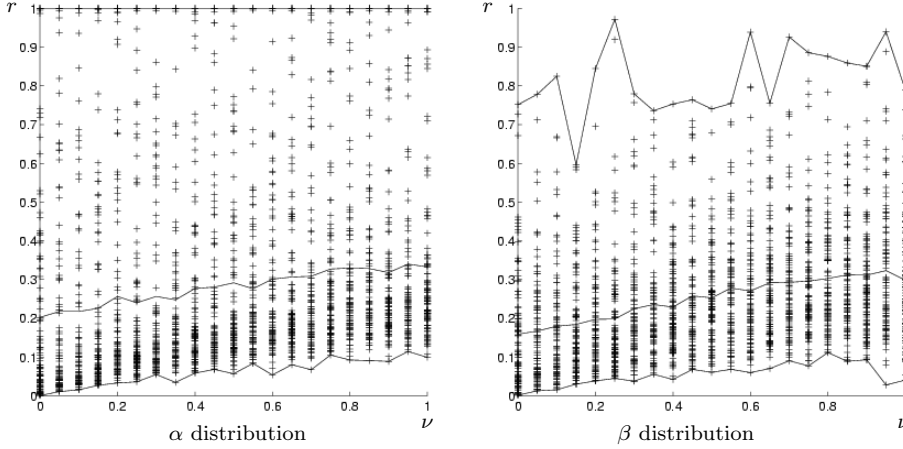


FIG. 1. Ratio r as a function of ν for the two distributions.

The bound on r given by Lemma 2.4 is approximately maximized for $\nu = 2/3$. If $n = 50$ and $n_1 = 49$, the theoretical bound is approximately 7×10^{-31} for this value of ν . This should be compared with the computed values of r , for which the minimum value of r attained a maximum of 0.092 for $\nu = 0.8$. Based on these results, we would recommend a value of ν in the range $(0.5, 0.9)$. In this interval, r never fell below 0.05. Note that the larger the value of ν , the smaller the value of n_1 and consequently, the smaller the amount of computation.

3. Theoretical results. The partial Cholesky factorization can be used as the basis for a descent method for minimizing a twice-continuously differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. This method defines a sequence $\{x_k\}_{k=0}^{\infty}$ of improving estimates of a local minimizer.

Let x_0 be any starting point such that the level set $\{x \mid f(x) \leq f(x_0)\}$ is compact. Let $\{s_k\}$ and $\{d_k\}$ be bounded sequences such that each s_k is a descent direction that satisfies (1.1a) and each d_k is a direction of negative curvature that satisfies (1.1b). Moré and Sorensen [17] show that with an appropriate linesearch, certain linear combinations of s_k and d_k define x_{k+1} so that every limit point of $\{x_k\}_{k=0}^{\infty}$ will satisfy the second-order necessary conditions for optimality—i.e., at every limit point \bar{x} , $\nabla f(\bar{x})$ is zero and $\nabla^2 f(\bar{x})$ is positive semidefinite. The main result of this paper—that the search directions obtained using the partial Cholesky factorization are sufficient in the sense of Moré and Sorensen [17]—is stated in the following theorem.

THEOREM 3.1. *Let $\{x_k\}_{k=0}^{\infty}$ be a sequence of iterates contained in a compact region of \mathbb{R}^n , and assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a twice-continuously differentiable function. For each k , define $g_k = \nabla f(x_k)$ and $H_k = \nabla^2 f(x_k)$, and let $H_k = L_k B_k L_k^T$ be the partial Cholesky factorization of H_k as described in Algorithm 2.1. Given positive constants c_1 and c_2 ($c_1 < c_2$), let s_k be defined from Lemma 2.3 with the additional requirement that $c_1 \leq \lambda_{\min}(\bar{B}_k) \leq \lambda_{\max}(\bar{B}_k) \leq c_2$. Finally, let d_k be*

defined from Lemma 2.4. Then, $\{s_k\}$ and $\{d_k\}$ are bounded sequences such that

$$g_k^T s_k \rightarrow 0 \implies g_k \rightarrow 0 \quad \text{and} \quad s_k \rightarrow 0$$

and

$$d_k^T H_k d_k \rightarrow 0 \implies \min\{\lambda_{\min}(H_k), 0\} \rightarrow 0 \quad \text{and} \quad d_k \rightarrow 0.$$

Proof. Since $\{x_k\}$ lies in a compact region, the smoothness of f implies that $\{\|g_k\|\}$ and $\{\|H_k\|\}$ are bounded.

With the existence of c_1 and c_2 , and the boundedness of $\|g_k\|$, Lemma 2.3 implies that $\{s_k\}$ is a bounded sequence, and $g_k^T s_k \rightarrow 0$ implies $g_k \rightarrow 0$ and $s_k \rightarrow 0$, as required.

Lemma 2.4 and the boundedness of $\|H_k\|$ imply that $\{d_k\}$ is a bounded sequence, and $d_k^T H_k d_k \rightarrow 0$ implies $d_k \rightarrow 0$ and $\min\{\lambda_{\min}(H_k), 0\} \rightarrow 0$, as required. \square

If $\nabla^2 f(x_k)$ is sufficiently positive definite, all pivots will be acceptable and the partial Cholesky factorization will terminate with $n_1 = n$. This implies that if $\{x_k\}_{k=0}^\infty$ has a limit point \bar{x} at which $\nabla^2 f(\bar{x})$ is sufficiently positive definite, then the iterates will be identical to those of Newton's method for k sufficiently large.

4. Discussion. Given an indefinite matrix H , almost all factorizations² of the form $H = LDL^T$ can be the basis of a method for finding s and d . However, each factorization must incorporate a pivoting strategy that ensures s and d satisfy (1.1). When D is block diagonal and L is a row-permuted unit lower-triangle, the pivoting rules of Bunch and Parlett [2], Fletcher [6], and Duff *et al.* [5], all provide suitable directions. However, all of these methods incorporate a *complete* pivoting strategy in the sense that it may be necessary to examine all rows of the Schur complement at any step (see Forsgren [8]). By contrast, the partial Cholesky factorization needs only a *partial* pivoting strategy, i.e., the equivalent of only two rows of the Schur complement need be examined at each step.

So-called *modified* factorizations of the form $LDL^T = H + E$ may also be used to define s and d (see, e.g., Gill and Murray [10]). As yet, none of these methods has been shown to generate d satisfying (1.1b).

The partial Cholesky factorization may be implemented in other ways. For example, the calculation of the matrix H_{11} can be made independent of the calculation of the descent direction s_k . Once a direction of negative curvature has been defined, a descent direction can be calculated by forming a modified Cholesky factorization of B_2 (see, e.g., Gill and Murray [10], Schnabel and Eskow [20]).

The algorithm of Section 2.2 requires the examination of the diagonals and a single row of the Schur complement at each step. Alternative strategies can be devised in which the complete Schur complement is examined under certain exceptional circumstances. For example, if a pivot is small, the pivot acceptance criterion could be strengthened so that a pivot is acceptable if, in addition to the requirements of Algorithm 2.1, it is larger in absolute value than νb_{\max} , where b_{\max} is either the diagonal of largest magnitude in the Schur complement or the element of largest magnitude in the full Schur complement. Each of these modifications gives an algorithm with identical theoretical properties, but a potentially smaller value of n_1 . However, this potential improvement is at the expense of an increase in the number of comparisons during the factorization. The pivot criterion that requires the examination of the full Schur

²One exception to this rule is Aasen's method [1], for which no $O(n^3)$ method is known for computing suitable s and d .

complement would cope successfully with the “pathological” $H(\theta)$ of Section 2.2 since the factorization would terminate after one step for θ sufficiently small.

5. Summary. We have shown how a partial Cholesky factorization can be used to define search directions suitable for a linesearch-based modified Newton method. The resulting directions are sufficient in the sense that it is possible to generate a sequence $\{x_k\}_{k=0}^{\infty}$ with limit points having a zero gradient and a positive-semidefinite Hessian.

To our knowledge, this is the first triangular factorization that not only gives theoretically satisfactory directions, but also requires only a *partial* pivoting strategy, i.e., the equivalent of only two rows of the Schur complement need be examined at each step.

Acknowledgements. We thank the Numerical Algorithms Group, Oxford, for providing the computing facilities that enabled work on this paper to be completed. We also thank the referees for their careful reading of the manuscript and their helpful suggestions.

Appendix A. Eigenvalues of $H(0)$.

LEMMA A.1. *Let the $n \times n$ -matrices $L(0)$ and $B(0)$ be defined as in Section 2 for $\theta = 0$ and $n \geq 3$. Define $H(0) = L(0)B(0)L(0)^T$. Then $\lambda = -\frac{1}{2}(\sqrt{n^2 + 2n - 7} - n + 1)$ is the smallest eigenvalue of $H(0)$, and $-1 \leq \lambda \leq -1 + 4/(n + 1)$.*

Proof. It is straightforward to verify that

$$H(0) = \begin{pmatrix} 1 & -1 & \cdots & -1 & -1 & -1 \\ -1 & 1 & \cdots & 1 & 1 & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ -1 & 1 & \cdots & 1 & 1 & 1 \\ -1 & 1 & \cdots & 1 & 1 & 0 \\ -1 & 1 & \cdots & 1 & 0 & 1 \end{pmatrix}.$$

Since $B(0)$ has one negative eigenvalue and $L(0)$ is nonsingular, Sylvester’s law of inertia implies that $H(0)$ has one negative eigenvalue (see e.g., Golub and Van Loan [12, page 416]). Consequently, since λ is negative for $n \geq 3$, it is enough to show that it is an eigenvalue.

Assume that $v = (1 \ -1 \ -1 \ \cdots \ -1 \ a \ a)^T$ is an eigenvector of $H(0)$ for some scalar a . Then, if v is an eigenvector, there must exist a λ such that

$$n - 2 - 2a = \lambda \quad \text{and} \quad -n + 2 + a = \lambda a.$$

It is straightforward to show that for $n \geq 3$, these equations have a solution given by

$$\lambda = -\frac{\sqrt{n^2 + 2n - 7} - n + 1}{2} \quad \text{and} \quad a = \frac{\sqrt{n^2 + 2n - 7} + n - 3}{4}.$$

The upper and lower bounds on λ follow from the sequence of inequalities

$$n + 1 \geq \sqrt{(n + 1)^2 - 8} = (n + 1) \sqrt{1 - \frac{8}{(n + 1)^2}} \geq n + 1 - \frac{8}{n + 1}.$$

(Note that the lower bound can also be obtained directly from Lemma 2.1.) \square

REFERENCES

- [1] J. O. AASEN, *On the reduction of a symmetric matrix to tridiagonal form*, Nordisk Tidskr. Informationsbehandling (BIT), 11 (1971), pp. 233–242.
- [2] J. R. BUNCH AND B. N. PARLETT, *Direct methods for solving symmetric indefinite systems of linear equations*, SIAM J. Numer. Anal., 8 (1971), pp. 639–655.
- [3] E. CASAS AND C. POLA, *An algorithm for indefinite quadratic programming based on a partial Cholesky factorization*. Preprint. Universidad de Cantabria, Santander, Spain, 1990.
- [4] J. E. DENNIS, JR. AND R. B. SCHNABEL, *A view of unconstrained optimization*, in Optimization, North-Holland, Amsterdam, 1989, pp. 1–72.
- [5] I. S. DUFF, J. K. REID, N. MUNKSGAARD, AND H. B. NIELSEN, *Direct solution of sets of linear equations whose matrix is sparse, symmetric and indefinite*, J. Inst. Math. Applics., 23 (1979), pp. 235–250.
- [6] R. FLETCHER, *Factorizing symmetric indefinite matrices*, Linear Algebra Appl., 14 (1976), pp. 257–272.
- [7] R. FLETCHER AND T. L. FREEMAN, *A modified Newton method for minimization*, J. Optim. Theory Appl., 23 (1977), pp. 357–372.
- [8] A. FORSGREN, *On the computation of search directions in modified Newton methods for optimization*, February 15–17 1993. Presented at Conference on Large-Scale Optimization, University of Florida, Gainesville.
- [9] A. FORSGREN, P. E. GILL, AND W. MURRAY, *On the identification of local minimizers in inertia-controlling methods for quadratic programming*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 730–746.
- [10] P. E. GILL AND W. MURRAY, *Newton-type methods for unconstrained and linearly constrained optimization*, Math. Program., 7 (1974), pp. 311–350.
- [11] D. GOLDFARB, *Curvilinear path steplength algorithms for minimization which use directions of negative curvature*, Math. Program., 18 (1980), pp. 31–40.
- [12] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, Maryland, second ed., 1989.
- [13] N. J. HIGHAM, *Analysis of the Cholesky decomposition of a semi-definite matrix*, in Reliable Numerical Computation, M. G. Cox and S. Hammarling, eds., Oxford University Press, 1990, pp. 161–185.
- [14] S. KANIEL AND A. DAX, *A modified Newton's method for unconstrained minimization*, SIAM J. Numer. Anal., 16 (1979), pp. 324–331.
- [15] MATHWORKS INC., *Matlab User's Guide*, Natick, Massachusetts, 1992.
- [16] G. P. MCCORMICK, *A modification of Armijo's step-size rule for negative curvature*, Math. Program., 13 (1977), pp. 111–115.
- [17] J. J. MORÉ AND D. C. SORESENSEN, *On the use of directions of negative curvature in a modified Newton method*, Math. Program., 16 (1979), pp. 1–20.
- [18] H. MUKAI AND E. POLAK, *A second-order method for unconstrained optimization*, J. Optim. Theory Appl., 26 (1978), pp. 501–513.
- [19] J. M. ORTEGA AND W. C. RHEINBOLDT, *Iterative solution of nonlinear equations in several variables*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. Reprint of the 1970 original.
- [20] R. B. SCHNABEL AND E. ESKOW, *A new modified Cholesky factorization*, SIAM J. Sci. and Statist. Comput., 11 (1990), pp. 1136–1158.
- [21] G. A. SHULTZ, R. B. SCHNABEL, AND R. H. BYRD, *A family of trust-region based algorithms for unconstrained minimization with strong global convergence properties*, SIAM J. Numer. Anal., 22 (1985), pp. 47–67.