STENMIN: A Software Package for Large, Sparse, Unconstrained Optimization Using Tensor Methods

Ali Bouaricha

CRPC-TR94584
October 1994

Center for Research on Parallel Computation
Rice University
6100 South Main Street
CRPC - MS 41
Houston, TX 77005

Also available as *MCS-P451-0794* from Argonne National Laboratory, Mathematics and Computer Science Division.

STENMIN: A Software Package for Large, Sparse Unconstrained Optimization Using Tensor Methods *

Ali Bouaricha†

Argonne National Laboratory

We describe a new package for minimizing an unconstrained nonlinear function where the Hessian is large and sparse. The software allows the user to select between a tensor method and a standard method based upon a quadratic model. The tensor method models the objective function by a fourth–order model, where the third– and fourth–order terms are chosen such that the extra cost of forming and solving the model is small. The new contribution of this package consists of the incorporation of an entirely new way of minimizing the tensor model that makes it suitable for solving large, sparse optimization problems efficiently. The test results indicate that, in general, the tensor method is significantly more efficient and more reliable than the standard Newton method for solving large, sparse unconstrained optimization problems.

Categories and Subject Descriptors: G.1.3 [**Numerical Analysis**]: Numerical Linear Algebra–*sparse and very large systems*; G.1.6 [**Numerical Analysis**]: Optimization–*unconstrained optimization*; G.4 [**Mathematics of Computing**]: Mathematical Software

General Terms: Algorithms

Additional Key Words and Phrases: tensor methods, sparse problems, large-scale optimization, rank-deficient matrices

*Part of this work was performed while the author was research associate at CERFACS (Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, Toulouse, France).

†Author's address: Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Illinois, 60439. bouarich@mcs.anl.gov. This work was supported in part by the Office of Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

1. Introduction

This paper describes a software package for solving the unconstrained optimization problem

$$\text{given } f : \Re^n \to \Re, \text{ find } x_* \in \Re^n \text{ such that } f(x_*) \leq f(x) \text{ for all } x \in D, \qquad (1.1)$$

using tensor methods, where D is some open set containing x_*. We assume that f is at least twice continuously differentiable and $\nabla^2 f(x_c)$ is large and sparse.

Tensor methods for unconstrained optimization are general–purpose methods primarily intended to improve upon the performance of standard methods especially on problems where $\nabla^2 f(x_*)$ has a small rank deficiency, and to be at least as efficient as standard methods on problems where $\nabla^2 f(x_*)$ is nonsingular. Tensor methods for unconstrained optimization base each iteration upon the fourth–order model of the objective function $f(x)$

$$M_T(x_c + d) = f(x_c) + \nabla f(x_c) \cdot d + \frac{1}{2}\nabla^2 f(x_c) \cdot d^2 + \frac{1}{6} T_c \cdot d^3 + \frac{1}{24} V_c \cdot d^4, \qquad (1.2)$$

where $d \in \Re^n$, x_c is the current iterate, $\nabla f(x_c)$ and $\nabla^2 f(x_c)$ are the first and second analytic derivatives of f at x_c, or finite difference approximations to them, and the tensor terms at x_c, $T_c \in \Re^{n \times n \times n}$ and $V_c \in \Re^{n \times n \times n \times n}$, are symmetric. (We use the notation $\nabla f(x_c) \cdot d$ for $\nabla f(x_c)^T d$, and $\nabla^2 f(x_c) \cdot d^2$ for $d^T \nabla^2 f(x_c) d$ to be consistent with the tensor notation $T_c \cdot d^3$ and $V_c \cdot d^4$. We abbreviate terms of the form dd, ddd, and $dddd$ by d^2, d^3, and d^4, respectively.)

Schnabel and Chow [11] select T_c and V_c such that the model interpolates function and gradient values from p past iterates, where p is a small number. This strategy results in T_c and V_c being low-rank tensors, which is crucial for the efficiency of the tensor method. Here, we consider only the case where the tensor model interpolates $f(x)$ and $\nabla f(x)$ at the previous iterate (i.e., $p = 1$). The reasons for this choice are that the performance of the tensor version that allows $p \geq 1$ is similar overall to that constraining p to be 1, and that the method is simpler and less expensive to implement in this case.

The above choice of T_c and V_c yields the tensor model

$$M_T(x_c + d) = f(x_c) + \nabla f(x_c) \cdot d + \frac{1}{2}\nabla^2 f(x_c) \cdot d^2 + \frac{1}{2}(b^T d)(s^T d)^2 + \frac{\gamma}{24}(s^T d)^4, \quad (1.3)$$

where $s \in \Re^n$ is the step from x_c to the previous iterate x_{-1} (i.e., $s = x_{-1} - x_c$) and $b \in \Re^n$ and $\gamma \in \Re$ are uniquely determined by the requirements $M_T(x_{-1}) = f(x_{-1})$ and $\nabla M_T(x_{-1}) = \nabla f(x_{-1})$. The whole process of forming the tensor model requires only $O(n^2)$ arithmetic operations. The storage needed for forming and storing the tensor model is only a total of $6n$.

The tensor algorithms described in [11] are QR-based algorithms involving orthogonal transformations of the variable space. These algorithms are very effective for minimizing the tensor model when the Hessian is dense because they are very stable numerically, especially when the Hessian is singular. They are not efficient for sparse problems, however, because they destroy the sparsity of the Hessian due to the orthogonal transformation of the variable space. To preserve the sparsity of the Hessian, we developed in [4] an entirely new way of minimizing the tensor model that employs a sparse variant of the Cholesky decomposition. This makes the new algorithms very well suited for sparse problems. In this new approach, we show that the minimization of (1.3) can be reduced to the solution of a third–order polynomial in one unknown,

plus the solution of three systems of linear equations that all involve the same coefficient matrix $\nabla^2 f(x_c)$. The STENMIN package is essentially based on this new approach.

The remainder of this paper is organized as follows. In §2 an iteration of tensor methods for large, sparse unconstrained optimization is outlined. In §3 we give an overview of the input, output, and important options provided by the software package. We describe the user interface to the package in §4, which includes both a simplified (default) and a longer calling sequence. In §5 we describe the meaning of the input, input-output, and output parameters for the package. In §6 we present the default values provided by the package. A few implementation dependencies are described in §7. In §8 we give an example of the use of the package. Finally, in §9 we describe comparative testing for an implementation based on the tensor method versus an implementation based on the Newton's method, and we present summary statistics of the test results.

2. An Iteration of Tensor Methods

In this section, we present the overall algorithm for tensor methods for large, sparse unconstrained optimization. Algorithm 2.1 is a slightly modified version of the algorithm described in [4] in the way the tensor step is selected when the β equation (see algorithm below) has more than one root. In general, this new way of computing the tensor step appears to perform better than the strategy described in [4], in both function evaluations and execution times. A summary of the experimental results for this implementation is presented in §9.

Algorithm 2.1. An Iteration of Tensor Methods for Large, Sparse Unconstrained Optimization

Let x_c be the current iterate, x_+ the next iterate, d_t the tensor step, and d_n the Newton step.

1. Calculate $\nabla f(x_c)$, and decide whether to stop. If not:

2. Calculate $\nabla^2 f(x_c)$

3. Calculate b and γ in the tensor model (1.3), so that the tensor model interpolates $f(x)$ and $\nabla f(x)$ at x_{-1}

4. Find a potential minimizer d_t of the tensor model

4.1. Factorize $\nabla^2 f(x_c)$ using the MA27 package [8]

4.2. **if** $\nabla^2 f(x_c)$ has full rank **then**

4.2.1. Form the β equation ($\beta \in \Re$) :

$$-u + (yw - uv - 1)\beta - \tfrac{3}{2}v\beta^2 + (\tfrac{1}{2}wz - \tfrac{\gamma}{6}w - \tfrac{1}{2}v^2)\beta^3,$$

where $u = s^T\nabla^2 f(x_c)^{-1}\nabla f(x_c)$, $v = s^T\nabla^2 f(x_c)^{-1}b$, $w = s^T\nabla^2 f(x_c)^{-1}s$, $y = b^T\nabla^2 f(x_c)^{-1}\nabla f(x_c)$, and $z = b^T\nabla^2 f(x_c)^{-1}b$

4.2.2. Compute the roots of the β equation

4.2.3. Select $\beta_* = \min(|\beta_i|)$ where β_i are the roots of the β equation

4.2.4. Substitute β_* into

$$\theta_* = -\frac{(u + \beta_* + \frac{1}{2}v\beta_*^2 + \frac{\gamma}{6}w\beta_*^3)}{w\beta_*}$$

4.2.5. Calculate the tensor step:

$$d_t = -\nabla^2 f(x_c)^{-1}(\nabla f(x_c) + \theta_*\beta_* s + \tfrac{1}{2}\beta_*^2 b + \tfrac{\gamma}{6}\beta_*^3 s)$$

4.3. **elseif** $\nabla^2 f(x_c)$ is singular with $\text{rank}(\nabla^2 f(x_c)) = n - 1$ **then**

4.3.1. Form the β equation ($\beta \in \Re$):

$u + (1 + \hat{\beta}v)\beta + (\frac{1}{2}v + \frac{\gamma}{2}w\hat{\beta})\beta^2 + \frac{\gamma}{6}w\beta^3$,

where $u = s^T\hat{\nabla}^2 f(x_c)^{-1}\hat{\nabla} f(x_c)$, $\hat{\nabla}^2 f(x_c) = \nabla^2 f(x_c) + ss^T$,
$\hat{\nabla} f(x_c) = \nabla f(x_c) + \nabla^2 f(x_c)\hat{d} + \hat{\theta}\hat{\beta}s + \frac{1}{2}\hat{\beta}^2 b + \frac{\gamma}{6}\hat{\beta}^3 s$, $\hat{\beta} = s^T\hat{d}$, $\hat{\theta} = b^T\hat{d}$, $\hat{d}$ is the global step computed in the previous iteration, $v = s^T\hat{\nabla}^2 f(x_c)^{-1}b$, and $w = s^T\hat{\nabla}^2 f(x_c)^{-1}s$.
The matrix $\hat{\nabla}^2 f(x_c)$ is factorized using the augmented system approach described in [4]

4.3.2. Compute the roots of the β equation

4.3.3. Select $\beta_* = \min(|\beta_i|)$ where β_i are the roots of the β equation

4.3.4. Substitute β_* into

$\theta_* = \frac{1}{w(\hat{\beta}+\beta_*)}(yw\hat{\beta} - u - uv\hat{\beta} + (yw + zw\hat{\beta}^2 - 2v\hat{\beta} - v^2\hat{\beta}^2 - uv - 1)\beta_*$
$+(\frac{3}{2}zw\hat{\beta} - \frac{\gamma}{2}w\hat{\beta} - \frac{3}{2}v - \frac{3}{2}v^2\hat{\beta}) + \frac{1}{2}zw - \frac{\gamma}{6}w - \frac{v^2}{2})\beta_*^3)$,

where $y = b^T\hat{\nabla}^2 f(x_c)^{-1}\hat{\nabla} f(x_c)$, and $z = b^T\hat{\nabla}^2 f(x_c)^{-1}b$

4.3.5. Calculate the tensor step of the transformed tensor model (2.1) below

$\delta = -\hat{\nabla}^2 f(x_c)^{-1}(\hat{\nabla} f(x_c) + \hat{\beta}\beta_* b + \hat{\beta}\theta_* s + \beta_*\theta_* s + (\frac{1}{2}b + \frac{\gamma}{2}\hat{\beta}s)\beta_*^2 + \frac{\gamma}{6}\beta_*^3 s)$

4.3.6. Calculate the tensor step of the original model (1.3):

$d_t = \delta + \hat{d}$

4.4. else $\{rank(\nabla^2 f(x_c)) < n - 1\}$

4.4.1. Modify the eigencomponents of $\nabla^2 f(x_c)$

4.4.2. Perform steps 4.2.1–4.2.5

endif

5. Test whether the tensor step is descent. If it is not compute the Newton step

5.1. if $\nabla^T f(x_c)d_t > 0$ **then**

5.1.1. Compute the Newton step

if $rank(\nabla^2 f(x_c)) < n - 1$ **then**

$dn = \nabla_m^2 f(x_c)^{-1}\nabla f(x_c)$, where $\nabla_m^2 f(x_c)$ is $\nabla^2 f(x_c)$ with the eigencomponents modified, (d_n is obtained for free)

else

Modify the eigencomponents of $\nabla^2 f(x_c)$

if all the eigencomponents of $\nabla^2 f(x_c)$ remain unchanged
$\{\nabla^2 f(x_c)$ is already positive definite$\}$ **then**

$dn = \nabla^2 f(x_c)^{-1}\nabla f(x_c)$, ($d_n$ is obtained for free)

else

$dn = \nabla_m^2 f(x_c)^{-1}\nabla f(x_c)$

endif

endif

endif

6. Compute a next iterate x_+

6.1. if d_t is descent **then**

$x_+^t = x_c + d_t$

if $f(x_+^t) < f(x_c) + 10^{-4}\cdot\nabla f(x_c)^T d_t$ **then**

$x_+ = x_+^t$

else

Find an acceptable x_+^n in the Newton direction d_n
using the line search Algorithm A6.3.1, page 325 [7]

Find an acceptable x_+^t in the tensor direction d_t
using the line search Algorithm A6.3.1, page 325 [7]
if $f(x_+^n) < f(x_+^t)$ **then**
$x_+ = x_+^n$
else
$x_+ = x_+^t$
endif
endif

6.2. else

Find an acceptable x_+^n in the Newton direction d_n
using Algorithm A6.3.1, page 325 [7]
$x_+ = x_+^n$
endif

7. $x_c = x_+$, $f(x_c) = f(x_+)$, **go to** step 1

In step 1, the gradient is either computed analytically or approximated by the algorithm A5.6.3 given in Dennis and Schnabel [7]. In step 2, the Hessian matrix is either calculated analytically or approximated by a graph coloring algorithm described in [6]. In step 4.3, we first compute the tensor step δ of the transformed model (obtained by substituting $\hat{d} + \delta$ for d in (1.3), where $\hat{d}$ is the global step computed in the previous iteration)

$$\begin{aligned} M_T(x_c + d) \quad = \quad & f(x_c) \; + \; \nabla f(x_c) \cdot \hat{d} \; + \; \frac{1}{2}\nabla^2 f(x_c) \cdot \hat{d}^2 \; + \; \frac{1}{2}(b^T \hat{d})(s^T \hat{d})^2 \\ & + \; \frac{\gamma}{24}(s^T \hat{d})^4 \; + \; (\nabla f(x_c) \; + \; \nabla^2 f(x_c)\hat{d} \; + \; (b^T \hat{d})(s^T \hat{d})s \\ & + \; \frac{1}{2}(s^T \hat{d})^2 b \; + \; \frac{\gamma}{24}(s^T \hat{d})^3 s) \cdot \delta \; + \; \frac{1}{2}(\nabla^2 f(x_c) \\ & + \; (b^T \hat{d} \; + \; \frac{\gamma}{2} s s^T) \cdot \delta^2 \; + \; (s^T \hat{d})(b^T \delta)(s^T \delta) \; + \; \frac{1}{2}(b^T \delta)(s^T \delta)^2 \\ & + \; \frac{\gamma}{6}(s^T \hat{d})(s^T \delta)^3 \; + \; \frac{\gamma}{24}(s^T \delta)^4. \end{aligned} \tag{2.1}$$

Then we set the tensor step d of the original tensor model (1.3) to $\hat{d} + \delta$. In step 4.4, we obtain a perturbation μ such as $\nabla^2 f(x_c) + \mu I$ is safely positive definite by using the Gill, Murray, Ponceleon, and Saunders method [9]. After we compute the $\mathrm{LDL}^{\mathrm{T}}$ of the Hessian matrix using the MA27 package [8], we change the block diagonal matrix D to D + E. The modified matrix is block diagonal positive definite. This guarantees that the decomposition $\mathrm{L(D + E)L}^{\mathrm{T}}$ is positive definite as well. Note that the Hessian matrix is not modified if it is already positive definite. In step 5, we test whether or not the tensor step is descent. If it is not, then we compute the Newton step d_n as a by–product of the minimization of the tensor model. That is, if $rank(\nabla^2 f(x_c)) < n - 1$ or all the eigencomponents of D turn out to be positive, i.e., $\nabla^2 f(x_c)$ is positive definite, then the Newton step is obtained for free; otherwise we perform another solve after we have modified the eigencomponents of D. Thus, d_n is the modified Newton step $(\nabla^2 f(x_c) + \mu I)^{-1} \nabla f(x_c)$, where $\mu = 0$ if $\nabla^2 f(x_c)$ is safely positive definite, and $\mu > 0$ otherwise. In step 6, we perform a standard backtracking line search global strategy to compute a next iterate x_+. The line search tensor method is much simpler to implement and to understand

than the two-dimensional trust region tensor method introduced in [4], and is appreciably faster. For these reasons, this software uses a line search method. The global framework for the line search method we used in conjunction with our tensor method for large, sparse unconstrained optimization is similar to the one used for systems of nonlinear equations [3, 5]. This strategy has proved very successful for large, sparse systems of nonlinear equations. This approach always tries the full tensor step first. If this provides enough decrease in the objective function, then we terminate; otherwise we find acceptable next iterates in both the Newton and tensor directions and select the one with the lower function value as the next iterate. The stopping criteria of Algorithm 2.1 are described by the parameter `TERMCD` in §5.

3. Overview of the Software Package

The required input to the software is the number of variables `N`, the function `FCN` that computes $f(x)$, an initial guess x_0, the number of nonzeros `NZ` stored in the lower or upper triangular part of the Hessian matrix, and the row and column indices of these nonzeros given in any order.

Two methods of calling the package are provided. In the short version, the user supplies only the above information, and default values of all other options are used. These include the calculation of the gradient and Hessian matrix by finite differences, and the use of the tensor rather than the standard Newton method. In the other method for calling the package, the user may override any default values of the package options.

The user has the option to choose between the tensor method and the standard Newton method. If the flag `METHOD` is set to 0, the package will use the standard method. The tensor method is used otherwise.

Upon completion, the program returns with an approximation `XPLS` to the minimizer x_*, the value of the objective function `FPLS` at `XPLS`, the value of the gradient `GPLS(XPLS)`, the Hessian `HESS(XPLS)`, and a flag specifying under which stopping condition the algorithm has terminated.

The software package is coded so that if the user inputs the typical magnitude $\texttt{TYPX}_i$ of each component of x, the performance of the package is the equivalent to what would result from redefining the independent variable x with

$$x_{\text{scaled}} = \begin{bmatrix} 1/\texttt{TYPX}_1 & & \\ & \ddots & \\ & & 1/\texttt{TYPX}_n \end{bmatrix} \cdot x \tag{3.1}$$

and then running the package without scaling. The default value of each $\texttt{TYPX}_i$ is 1. Scaling is often important to use for problems in which the variable components are widely different in magnitudes.

The user may supply analytic routines for the gradient and/or the Hessian. If they are not supplied the package computes them by finite differences. The parameters `GRDFLG` and `HSNFLG`

specify whether analytic gradient and Hessian have been provided, respectively. When the analytic gradient and/or Hessian are supplied, the user has the option of checking the supplied analytic routines against the package's finite difference routines.

The standard (default) output from this package consists of printing the input parameters and the final results. The printed input parameters are those used by the algorithm and hence include any corrections made by the program module `OPTCHK`, which examines the input specifications for illegal entries and consistency. The program will provide an error message if it terminates as a result of input errors. The printed results include a message indicating the reason for termination, an approximation `XPLS` to the solution x_*, the function value at `XPLS`, and the gradient vector `GPLS`. The package provides an additional means for the control of output via the variable `MSG` described in §5. The standard output is the input state, the final results, and the stopping conditions. The user may suppress all output or may print the intermediate iteration results in addition to the standard output.

If the user sets the variable `INFORM` to 1, then the package uses reverse communication to obtain the multiplication of the Hessian matrix at the current iterate by a given vector. If `INFORM` is set to 0, then this quantity is computed by the subroutine `MATMV` provided by the package.

4. Interfaces and Usage

Two interfaces have been provided with the package. If the user wishes to use all the defaults options provided by the package, then he should call `TENSDO` (`TENSSO` if single–precision is used). Only the required input described in §3 needs to be supplied. The other interface, `TENSD` (`TENSS` if single–precision is used), requires the user to supply all parameters. The user may specify selected parameters only by first invoking the subroutine `DFAULT`, which sets all parameters to their default values, and then overriding only the desired values.

The two calling sequences are as follows.

```
1.    CALL TENSDO(N, X0, NZ, IRN, LIRN, ICN, LICN, FCN, TYPX, MSG,
     * XPLS, FPLS, GPLS, HESS, WRK, LWRK, IWRK, LIWRK, HTV, TERMCD)


2.    CALL DFAULT(N, TYPX, FSCALE, GRADTL, STEPTL, ILIM, STEPMX,
     *            IPR, METHOD, GRDFLG, HSNFLG, NDIGIT, INFORM, MSG)

C  USER OVERRIDES SPECIFIC DEFAULT VALUES PARAMETERS, E.G.

               GRADTL   = 1.0D-6
               ILIM     = 500
               GRDFLG   = 1
               HSNFLG   = 1

      CALL TENSD(N, X0, NZ, IRN, LIRN, ICN, LICN, FCN, UGR,
```

```
*  USH, TYPX, FSCALE, GRADTL, STEPTL, ILIM, STEPMX, IPR,
*  METHOD, GRDFLG, HSNFLG, NDIGIT, MSG, XPLS, FPLS, GPLS,
*  HESS, WRK, LWRK, IWRK, LIWRK, TERMCD, HTV, INFORM)
```

5. Parameters and Default Values

The parameters used in the calling sequences of §4 are fully described here. TENSDO uses only those parameters that are preceded by an asterisk. When it is noted that module DFAULT returns a given value, this is the default employed by interface TENSDO. The user may override the default value by utilizing TENSD.

Following each variable name in the list below appears a one– or a two–headed arrow symbol of forms $\rightarrow$, $\leftarrow$, and $\longleftrightarrow$. These symbols signify that the variable is for input, output, and input-output, respectively.

*N$\rightarrow$: A positive integer variable specifying the number of variables in the problem. **Restriction:** N ≥ 1.

*X0$\rightarrow$: An array of length N that contains an initial estimate of the minimizer x_*.

*NZ$\rightarrow$: An integer variable that must be set by the user to the number of nonzeros stored in the lower or upper triangular part of the Hessian matrix. It is not altered by the program. **Restriction:** NZ ≥ 1.

*IRN$\rightarrow$: An integer array of length LIRN. On entry, it must hold the row index of each nonzero stored in the lower or upper triangular part of the Hessian matrix.

*LIRN$\rightarrow$: An integer variable that must be set by the user to the length of array IRN. LIRN need not be as large as LICN; normally it need not be very much greater than NZ. It is not altered by the program. **Restriction:** LIRN $\geq$ NZ.

*ICN$\rightarrow$: An integer array of length LICN. On entry, it must hold the column index of the nonzeros stored in lower or upper triangular part of the Hessian matrix.

*LICN$\rightarrow$: An integer variable that must be set by the user to the length of the Hessian array HESS and ICN. LICN should ordinarily be 2 to 4 times as large as NZ. It is not altered by the program. **Restriction:** LICN $\geq$ NZ.

*FCN$\rightarrow$: The name of a user supplied subroutine that evaluates the function f at an arbitrary vector x. The subroutine must be declared EXTERNAL in the user's calling program and must conform to

```
CALL FCN(N, X, F),
```

where X is a vector of length N. The subroutine must not alter the values of X.

UGR→: The name of a user supplied subroutine that returns in G the value of the gradient. UGR must be declared EXTERNAL in the user's calling program and must conform to the usage

```
CALL UGR(N, X, G),
```

where N is the dimension of the problem, X is a vector of length N, and G is the gradient at X. UGR must not alter the values of N and X. When using the interface TENSD, if no analytic gradient is supplied (GRDFLG = 0), the user must use the dummy name DUGR.

USH→: The name of a user supplied subroutine that returns in HESS the value of the Hessian $\nabla^2 f(x)$ at the current point X. USH must be declared EXTERNAL in the user's calling program and must conform to the usage

```
CALL USH(N, X, NZ, LICN, HESS, IRN, ICN)
```

where N is the dimension of the problem, X is a vector of length N, HESS is the Hessian at X, LICN is the length of HESS, NZ is the number of nonzeros in HESS, and IRN and ICN are the row and column indices of the nonzeros in HESS. USH must not alter the values of N, X, and LICN. Only the lower or upper triangular part of HESS should be given. When using the interface TENSD, if no analytic gradient is supplied (HSNFLG = 0), the user must use the dummy name DUSH.

*TYPX→: An array of length N in which the typical size of the components of X are specified. The typical component sizes should be positive real scalars. If a negative value is specified, its absolute value will be used. When 0. is specified, 1. will be used. The program will not abort. This vector is used by the the package to determine the scaling matrix D_x. Although the package may work reasonably well in a large number of instances without scaling, it may fail when the components of x_* are of radically different magnitude and scaling is not invoked. If the sizes of the parameters are known to differ by many orders of magnitude, then the scale vector TYPX should definitely be used. Module DFAULT returns TYPX = (1.0, ..., 1.0). For example, if it is anticipated that the range of values for the iterates x_k is

$$\begin{aligned} x_1 &\in [-10^{10}\ ,\ 10^{10}] \\ x_2 &\in [-10^{2}\ ,\ 10^{4}] \\ x_3 &\in [-6\times 10^{-6},\ 9\times 10^{-6}] \end{aligned}$$

then an appropriate choice will be TYPX = (1.0E+10, 1.0E+3, 7.0E-6).

FSCALE→: A positive real number estimating the magnitude of $f(x)$ near the minimizer x_*. It is used in the gradient stopping condition given below. If $f(x_0)$ is much greater than $f(x_*)$, FSCALE should be approximately $f(x_*)$. If a negative value is specified for FSCALE, its absolute value is used. When 0. is specified, 1. will be used. The program will not abort.

GRADTL→: Positive scalar giving the tolerance at which the scaled gradient of $f(x)$ is considered close enough to zero to terminate the algorithm. The scaled gradient is a measure of the relative change in f in each direction x_i divided by the relative change in x_i. More precisely,

the test used by the program is

$$\max_i \left\{ \frac{|\ \nabla f(x)\ |_i\ \max\{|\ x_i\ |, \texttt{TYPX}_i\}}{\max\{|\ f\ |, \texttt{FSCALE}\}} \right\} \leq \texttt{GRADTL}.$$

The module DFAULT returns the value $\epsilon^{1/3}$. If the user specifies a negative value, the default value is used instead.

STEPTL→: A positive scalar providing the minimum allowable relative step length. STEPTL should be at least as small as 10^{-d}, where d is the number of accurate digits the user desires in the solution x_*. The actual test used is

$$\max_i \left\{ \frac{|x_i{}^k - x_i{}^{k-1}|}{\max\{|x_i{}^k, \texttt{TYPX}_i|\}} \right\} \leq \texttt{STEPTL},$$

where x^k and x^{k-1} are the new and old iterates, respectively. The program may terminate prematurely if STEPTL is too large. Module DFAULT returns the value $\epsilon^{2/3}$. If the user specifies a negative value, then the default value is used instead.

ILIM→: Positive integer specifying the maximum iterations to be performed before the program is terminated. Module DFAULT returns ILIM = 150. If the user specifies ILIM ≤ 0, the default value is used instead.

STEPMX→: A positive scalar providing the maximum allowable scaled step length $||D_x(x_+ - x_c)||_2$, where $D_x = diag(\texttt{1/TYPX}_1, \ \ldots, \ \texttt{1/TYPX}_n)$. STEPMX is used to prevent steps that would cause the optimization problem to overflow, to prevent the algorithm from leaving the area of interest in parameter space, or to detect divergence in the algorithm. STEPMX should be chosen small enough to prevent these occurrences but should be larger than any anticipated "reasonable" step. The algorithm will halt and provide a diagnostic if it attempts to exceed STEPMX on five successive iterations. If a nonpositive value is specified for STEPMX, the default is used. Module DFAULT returns the value STEPMX $= \max\{||x_0||_2 \cdot 10^3, 10^3\}$, where x_0 is the initial approximation provided by the user.

IPR→: The unit on which the routine outputs information. DFAULT returns the value 6, which is the standard FORTRAN unit for the printer.

METHOD→: An integer flag designating which method to use.

- METHOD = 0 : Use Newton's method.
- METHOD = 1 : Use the tensor method.

Module DFAULT returns value 1. If the user specifies an illegal value, module OPTCHK will set METHOD to 1; the program will not abort.

GRDFLG→: Integer flag designating whether or not analytic Hessian has been supplied by the user.

- GRDFLG = 0 : No analytic gradient supplied.
- GRDFLG = 1 : Analytic gradient supplied (will be checked against finite difference gradient.)

- GRDFLG = 2 : Analytic gradient supplied (will not be checked against finite difference gradient.)

When GRDFLG = 0, the gradient is obtained by finite differences. When GRDFLG = 1 or 2, the name of the user supplied routine that evaluates $\nabla f(x)$ must be supplied in UGR. When GRDFLG = 1, the program compares the value of the user's analytic gradient routine at x_0 with a finite difference estimate and aborts if the relative difference between any two components is greater than 0.01. The module DFAULT returns GRDFLG = 0. If the user specifies an illegal value, the module OPTCHK supplies the value 0.

HSNFLG→: Integer flag designating whether or not analytic Hessian has been supplied by the user.

- HSNFLG = 0 : No analytic Hessian supplied.
- HSNFLG = 1 : Analytic Hessian supplied (will be checked against finite difference Hessian.)
- HSNFLG = 2 : Analytic Hessian supplied (will not be checked against finite difference Hessian.)

When HSNFLG = 0, the Hessian values are computed by finite differences. When HSNFLG = 1 or 2, the name of the user–supplied routine that evaluates $\nabla^2 f(x)$ must be supplied in USH. When HSNFLG = 1, the program compares the value of the user's analytic Hessian routine at x_0 with a finite difference estimate and aborts if the relative difference between any two components is greater than 0.01. The module DFAULT returns HSNFLG = 0. If the user specifies an illegal value, the module OPTCHK supplies the value 0.

NDIGIT→: Integer estimating the number of accurate digits on the objective function $f(x)$. DFAULT returns the value -LOG$_{10}(\epsilon)$, where ϵ is machine precision. If NDIGIT ≤ 0 then the default value is used instead.

*MSG⟷: An integer variable that the user may set on input to inhibit certain automatic checks or override certain default characteristics of the package. Currently, three "message" features can be used individually or in combination.

- MSG = 0 : No output will be produced.
- MSG = 1 : Print the input state, the final results, and the stopping conditions.
- MSG = 2 : Print the intermediate results, that is, the input state, each iteration

including the current value of the objective function and the scaled gradient, and the final results including the stopping conditions and the number of function, gradient, and Hessian evaluations. The module DFAULT returns a value of 1. On output, if the program has terminated because of erroneous input, MSG contains an error code indicating the reason:

- MSG = -1 : Illegal dimension N; N ≤ 0. The program aborts.
- MSG = -2 : Illegal length of LIRN or LICN; LIRN ≤ 0 or LICN ≤ 0. The program aborts.
- MSG = -3 : Illegal length of LIWRK or LWRK; LIWRK $<$ 2*LIRN+12*N+2 or LWRK $<$ 7*N. The program aborts.
- MSG = -4 : Illegal number of nonzeros NZ; NZ ≤ 0. The program aborts.
- MSG = -5 : The K-th element of IRN or the K-th element of ICN is not an integer between 1 and N; (IRN(K) $<$ 1 or IRN(K) $>$ N) or (ICN(K) $<$ 1 or ICN(K) $>$ N). The program aborts.

- MSG = -6 : The K-th diagonal element is not in the sparsity pattern. This is checked only if HSNFLG = 0 because the finite difference Hessian approximation require that diagonal elements be in the sparsity pattern. The program aborts.
- MSG = -7 : Redundant entries in sparsity pattern was encountered. When HSNFLG = 1 or HSNFLG = 2, the program aborts. When HSNFLG = 0, the program eliminates the redundant entries and continue the execution (no error message is reported in this case).
- MSG = -8 : Probable coding error in the user's analytic gradient routine . Analytic and finite difference gradient do not agree within a tolerance of 0.01. The program aborts. (This check can be overridden by setting GRDFLG = 2.)
- MSG = -9 : Probable coding error in the user's analytic Hessian routine USH. Analytic and finite difference Hessian do not agree within a tolerance of 0.01. The program aborts. (This check can be overridden by setting HSNFLG = 2.)

*XPLS←: An array of length N containing the best approximation to the minimizer x_* upon return. (If the algorithm has not converged, the last iterate is returned.)

*FPLS←: A scalar variable that contains the function value at the final iterate XPLS.

*GPLS←: An array of length N containing the gradient value at XPLS.

HESS→: An array that is used to store the Hessian matrix at each iteration. It needs to be at least of dimension LICN. On exit, HESS contains the Hessian value at the minimizer x_*.

*WRK→: An array of length LWRK. This is used as workspace by the package. Its length must be at least 7*N.

*LWRK→: An integer variable. It must be set by the user to the length of array WRK and is not altered by the package.

*IWRK→: An integer array of length LIWRK. This is used as workspace by the package. Its length must be at least 2*LIRN+12*N+2.

*LIWRK→: An integer variable. It must be set by the user to the length of array IWRK and is not altered by the package.

*TERMCD←: An integer that specifies the reason why the algorithm has terminated.
- TERMCD = 1 : The scaled gradient at the final iterate was less than GRADTL.
- TERMCD = 2 : The length of the last step was less than STEPTL.
- TERMCD = 3 : Last global step failed to locate a point lower than XPLS. It is likely that either XPLS is an approximate solution of the function or STEPTL is too large.
- TERMCD = 4 : The iteration limit has been exceeded.
- TERMCD = 5 : Five consecutive steps of length STEPMX have been taken.

*HTV↔: An array of length N. It need not be set by the user on entry. If INFORM is set to 1, a re-entry must be made with HTV set to HESS times HTV (see INFORM.)

INFORM⟷: An integer variable. If it is set to 1, the user must obtain HESS times HTV and re-enter TENSD (TENSS if single–precision is used) with INFORM unchanged. The result of HESS times HTV must be stored in HTV. The default value of INFORM is 0, meaning that HESS times HTV is computed by the package.

6. Summary of Default Values

The following parameters are returned by the module DFAULT:

ILIM = 150
GRDFLG = 0
HSNFLG = 0
IPR = 6
GRADTL = $\epsilon^{1/3}$ (ϵ is machine precision)
STEPTL = $\epsilon^{2/3}$
METHOD = 1
NDIGIT = -LOG$_{10}(\epsilon)$
STEPMX = 0.0
TYPX = (1.0, ..., 1.0)
FSCALE = 1.0
MSG = 1
INFORM = 0

7. Implementation Details

This software package has been coded in Fortran 77. The user has the choice between single– and double–precision versions. The user must then preprocess the package at compile time using either the **tosngl** or **todble** tools from CUTE [2], for the single– and double–precision versions, respectively. The **tosngl** program picks up the appropriate version by selecting any statement that begins with CS in the first column, where the S character means that this is a single–precision version. On the other hand, the **todble** program picks up the appropriate version by selecting any statement that begins with CD in the first column, with D meaning that this is a double–precision version. Note that a statement that begins by neither CS nor CD will be picked by both tools.

The following software are included in the package:

1. Harwell MA27 package [8], which is used to compute the $L^T DL$ factorization of the sparse Hessian matrix.

2. The Coleman and Moré graph coloring algorithm [6], which is used for estimating a finite-difference approximation of a sparse Hessian matrix.

The program was developed and tested on a Sun SPARC 10 Model 40 computer.

The machine precision is calculated by the package and used in several places including finite differences stepsizes and stopping criteria. On some computers, the returned value may be incorrect because of compiler optimizations. The user may wish to check the computer value of the machine epsilon and, if it is incorrect, replace the code in the subroutine `MCHEPS` with the following statement

```
EPS = correct value of machine epsilon
```

8. Example of Use

In the example code shown in Figure 1, we first call the routine `DFAULT`, which returns the default values. We then override the values of `ILIM`, `GRADTL`, `GRDFLG` and `HSNFLG`. Next we call either the interface `TENSS` or `TENSD` for the single– and double–precision version, respectively, to solve the sparse unconstrained optimization problem coded in `FCN` and whose gradient and Hessian are given by `UGRAD` and `UHESS`, respectively.

```
C
C EXAMPLE OF USE FOR STENMIN.  THE TEST PROBLEM IS THE
C THE BROYDEN TRIDIAGONAL [10].
C
C  ALI BOUARICHA, OCTOBER 1994.
C  MCS DIVISION, ARGONNE NATIONAL LAB.
C
      INTEGER          NMAX, N, NZ, LIRN, LICN, ILIM, IPR, METHOD
      INTEGER          GRDFLG, HSNFLG, NDIGIT, MSG, LWRK, LIWRK
      INTEGER          TERMCD, INFORM, I
CD    DOUBLE PRECISION FSCALE, GRADTL, STEPTL, FPLS, STEPMX, ONE
CS    REAL             FSCALE, GRADTL, STEPTL, FPLS, STEPMX, ONE
      PARAMETER      ( NMAX = 10000, LIRN = 50000, LICN = 500000 )
      PARAMETER      ( LIWRK = 2 * LIRN + 12 * NMAX + 2 )
      PARAMETER      ( LWRK = 7 * NMAX )
      INTEGER          IRN ( LIRN  ), ICN ( LICN )
      INTEGER          IWRK( LIWRK )
CD    DOUBLE PRECISION X   ( NMAX ),  TYPX( NMAX ), XPLS( NMAX )
CD    DOUBLE PRECISION GPLS( NMAX ),  HESS( LICN ), WRK ( LWRK )
CD    DOUBLE PRECISION HTV ( NMAX )
CS    REAL             X   ( NMAX ),  TYPX( NMAX ), XPLS( NMAX )
CS    REAL             GPLS( NMAX ),  HESS( LICN ), WRK ( LWRK )
CS    REAL             HTV ( NMAX )
      EXTERNAL         FCN, UGRAD, UHESS
CD    DATA ONE / 1.0D0 /
CS    DATA ONE / 1.0E0 /
C
C READ DATA
```

```
C
      READ(5,*) N
C
C  COMPUTE THE STANDARD STARTING POINT.
C
      DO 10 I = 1, N
         X(I) = -ONE
 10   CONTINUE
C
      CALL DFAULT(N, TYPX, FSCALE, GRADTL, STEPTL, ILIM, STEPMX,
     *            IPR, METHOD, GRDFLG, HSNFLG, NDIGIT, INFORM, MSG)
C
      ILIM = 500
CD    GRADTL = 1.0D-5
CS    GRADTL = 1.0E-3
      GRDFLG = 1
      HSNFLG = 1
C
C  CALL THE SPARSE OPTIMIZER
C
CD    CALL TENSD(N,X,NZ,IRN,LIRN,ICN,LICN,FCN,UGRAD,
CS    CALL TENSS(N,X,NZ,IRN,LIRN,ICN,LICN,FCN,UGRAD,
     *  UHESS,TYPX,FSCALE,GRADTL,STEPTL,ILIM,STEPMX,IPR,
     *  METHOD,GRDFLG,HSNFLG,NDIGIT,MSG,XPLS,FPLS,GPLS,
     *  HESS,WRK,LWRK,IWRK,LIWRK,TERMCD,HTV,INFORM)
C
      STOP
      END
C
C THE FOLLOWING IS A SUBROUTINE FOR THE BROYDEN TRIDIAGONAL
C PROBLEM
C
        SUBROUTINE FCN(N, X, F)
        INTEGER N
CD      DOUBLE PRECISION X(N), F
CS      REAL X(N), F
C
C  LOCAL VARIABLES
C
        INTEGER I
CD      DOUBLE PRECISION ONE, TWO, THREE
CS      REAL ONE, TWO, THREE
CD      DATA ONE, TWO, THREE / 1.0D0, 2.0D0, 3.0D0 /
CS      DATA ONE, TWO, THREE / 1.0E0, 2.0E0, 3.0E0 /
```

```
C
        F = ((THREE - TWO * X(1)) * X(1) - TWO * X(2) + ONE) ** 2
        DO 10 I = 2, N-1
           F = F + ((THREE - TWO * X(I)) * X(I) - X(I-1) -
     *         TWO * X(I+1) + ONE) ** 2
 10     CONTINUE
        F = F + ((THREE - TWO * X(N)) * X(N) - X(N-1) + ONE) ** 2
        RETURN
        END
C
C THE FOLLOWING IS A SUBROUTINE FOR THE GRADIENT OF THE BROYDEN
C TRIDIAGONAL PROBLEM
C
        SUBROUTINE UGRAD(N, X, G)
        INTEGER N
CD      DOUBLE PRECISION X(N), G(N)
CS      REAL X(N), G(N)
C
C LOCAL VARIABLES
C
        INTEGER I
CD      DOUBLE PRECISION RL, RM, RR, ONE, TWO, THREE, FOUR
CS      REAL RL, RM, RR, ONE, TWO, THREE, FOUR
CD      DATA ONE, TWO, THREE, FOUR/1.0D0, 2.0D0, 3.0D0, 4.0D0/
CS      DATA ONE, TWO, THREE, FOUR/ 1.0E0, 2.0E0, 3.0E0, 4.0E0/
C
        RL = (THREE - TWO * X(1)) * X(1) - TWO * X(2) + ONE
        RR = (THREE - TWO * X(2)) * X(2) - X(1) - TWO * X(3) + ONE
        G(1) = TWO * (RL * (THREE - FOUR * X(1)) - RR)
        DO 10 I = 2, N-1
           IF(I .NE. 2) THEN
              RL = (THREE - TWO * X(I-1)) * X(I-1) - X(I-2) -
     *             TWO * X(I) + ONE
           ENDIF
           RM = (THREE - TWO * X(I)) * X(I) - X(I-1) -
     *           TWO * X(I+1) + ONE
           IF(I .EQ. N-1) THEN
              RR = (THREE - TWO * X(N)) * X(N) - X(N-1) + ONE
           ELSE
              RR = (THREE - TWO * X(I+1)) * X(I+1) - X(I) -
     *             TWO * X(I+2) + ONE
           ENDIF
        G(I) = -TWO * (TWO * RL - RM * (THREE - FOUR * X(I)) + RR)
 10     CONTINUE
```

```
        G(N) = -TWO * (TWO * RM - RR * (THREE - FOUR * X(N)))
        RETURN
        END
C
C THE FOLLOWING IS A SUBROUTINE FOR THE HESSIAN OF THE BROYDEN
C TRIDIAGONAL PROBLEM
C
        SUBROUTINE UHESS(N,X,NZ,LICN,HESS,IRN,ICN)
        INTEGER N, NZ, LICN
        INTEGER IRN(NZ), ICN(LICN)
CD      DOUBLE PRECISION X(N), HESS(LICN)
CS      REAL X(N), HESS(LICN)
C
C LOCAL VARIABLES
C
        INTEGER I
CD      DOUBLE PRECISION RL,RM,RR,DRLIM1,DRMI
CD      DOUBLE PRECISION ONE,TWO,THREE,FOUR
CS      REAL RL,RM,RR,DRLIM1,DRMI
CS      REAL ONE,TWO,THREE,FOUR
CD      DATA ONE, TWO, THREE, FOUR/1.0D0, 2.0D0, 3.0D0, 4.0D0/
CS      DATA ONE, TWO, THREE, FOUR/1.0E0, 2.0E0, 3.0E0, 4.0E0/
C
        NZ = 1
        RL = (THREE - TWO * X(1)) * X(1) - TWO * X(2) + ONE
        HESS(NZ) = TWO * ((THREE - FOUR * X(1))**2 -
     *            FOUR * RL + ONE)
        IRN(NZ) = 1
        ICN(NZ) = 1
        DO 10 I = 2, N-1
           DRLIM1 = THREE - FOUR * X(I-1)
           DRMI = THREE - FOUR * X(I)
           IF(I .NE. 2) THEN
              NZ = NZ + 1
              HESS(NZ) = FOUR
              IRN(NZ) = I
              ICN(NZ) = I-2
           ENDIF
           NZ = NZ + 1
           HESS(NZ) = -TWO * (TWO * (THREE - FOUR * X(I-1)) +
     *                ONE * (THREE - FOUR * X(I)))
           IRN(NZ) = I
           ICN(NZ) = I-1
           RM = (THREE - TWO * X(I)) * X(I) - X(I-1) -
```

```
*          TWO * X(I+1) + ONE
           NZ = NZ + 1
           HESS(NZ) = -TWO * (-FOUR - (THREE - FOUR * X(I))**2 +
*                 FOUR * RM  - ONE)
           IRN(NZ) = I
           ICN(NZ) = I
10      CONTINUE
        RR = (THREE - TWO * X(N)) * X(N) - X(N-1) + ONE
        NZ = NZ + 1
        HESS(NZ) = FOUR
        IRN(NZ) = N
        ICN(NZ) = N-2
        NZ = NZ + 1
        HESS(NZ) = -TWO * (TWO * (THREE - FOUR * X(N-1)) +
*              THREE - FOUR * X(N))
        IRN(NZ) = N
        ICN(NZ) = N-1
        NZ = NZ + 1
        HESS(NZ) = TWO * (FOUR + (THREE - FOUR * X(N))**2 - FOUR * RR)
        IRN(NZ) = I
        ICN(NZ) = I
        RETURN
        END
```

Figure 1: Code to solve the Broyden tridiagonal problem

If we use the double–precision version of the package to solve the Broyden tridiagonal problem given by `FCN`, for `N = 10000`, we obtain the following output:

```
 OPTIM      GRADIENT FLAG      = 1
 OPTIM      HESSIAN FLAG       = 1
 OPTIM      METHOD             = 1
 OPTIM      ITERATION LIMIT    =  500
 OPTIM      MACHINE EPSILON    =  0.2220446049250E-15
 OPTIM      STEP TOLERANCE     =  0.3666852862501E-10
 OPTIM      GRADIENT TOLERANCE =  0.1000000000000E-04
 OPTIM      MAXIMUM STEP SIZE  =  0.1000000000000E+06

 ----------------------------------------
 RESULT    ITERATION K   =    0
 RESULT    FUNCTION AT X(K)
 RESULT        0.1001100000000E+05
 RESULT    SCALED GRADIENT AT X(K)
 RESULT        0.3800000000000E+02
 ----------------------------------------
```

```
OPTSTP    RELATIVE GRADIENT CLOSE TO ZERO
OPTSTP    CURRENT ITERATE IS PROBABLY SOLUTION

-----------------------------------------
RESULT    ITERATION K   =    4
RESULT    FUNCTION AT X(K)
RESULT        0.1884575867777E-13
RESULT    SCALED GRADIENT AT X(K)
RESULT        0.1113397081739E-05
-----------------------------------------

RESULT    NUMBER OF FUNCTION EVALUATIONS    5
RESULT    NUMBER OF GRADIENT EVALUATIONS    5
RESULT    NUMBER OF HESSIAN EVALUATIONS    4
```

In the Appendix, we give another example of use–the optimal design with composite materials problem–from the MINPACK-2 collection [1].

9. Test Results

We tested our tensor and standard methods on the set of unconstrained optimization problems from the CUTE [2] and the MINPACK-2 [1] collections. Most of these problems have nonsingular Hessians at the solution. We also created singular test problems as proposed in [3, 12] by modifying the nonsingular test problems from the CUTE collection. The dimensions of these problems range from 100 to 10000. All our computations were performed on a Sun SPARC 10 Model 40 machine using double–precision arithmetic.

A summary for the test problems whose Hessians at the solution have ranks n, $n-1$, and $n-2$ is presented in Table 1. The descriptions of the test problems and the detailed results are given in [4]. In Table 1 the columns "better" and "worse" represent the number of times the tensor method was better and worse, respectively, than Newton's method by more than one gradient evaluation. The "tie" column represents the number of times the tensor and standard methods required within one gradient evaluation of each other. For each set of problems, we summarize the comparative costs of the tensor and standard methods using average ratios of three measures: gradient evaluations, function evaluations, and execution times. The average gradient evaluation ratio (geval) is the total number of gradients evaluations required by the tensor method, divided by the total number of gradients evaluations required by the standard method on these problems. The same measure is used for the average function evaluation (feval) and execution time (time) ratios. These average ratios include only problems that were successfully solved by both methods. We have excluded all cases where the tensor and standard methods converged to a different minimizer. However, the statistics for the "better," "worse," and "tie" columns include the cases where only one of the two methods converges, and exclude the cases where both methods do not converge. We also excluded problems requiring a number

of gradient evaluations less or equal than 3 by both methods. Finally, columns "t/s" and "s/t" show the number of problems solved by the tensor method but not by the standard method and the number of problems solved by the standard method but not by the tensor method, respectively.

The improvement by the tensor method over the standard method on problems with rank $n-1$ is dramatic, averaging 49% in function evaluations, 52% in gradient evaluations, and 60% in execution times. This is due in part to the rate of convergence of the tensor method being faster than that of Newton's method, which is known to be only linearly convergent with constant $\frac{2}{3}$. A typical convergence rate of the tensor method on rank $n-1$ problems is around 0.01. Whether this is a superlinear convergence remains to be proved. On problems with rank $n-2$, the improvement by the tensor method over the standard method is also substantial, averaging 34% in function evaluations, 37% in gradient evaluations, and 38% in execution times. In the test results obtained for the nonsingular problems, the tensor method is only 2% better than the standard method in function evaluations, but 32% and 37% better in gradient evaluations and in execution times, respectively. The tensor method requires on the average more function evaluations than the standard method on some nonsingular problems. This is because the full tensor step does not provide sufficient decrease in the objective function, and therefore the tensor method has to perform a line search method in both the Newton and tensor directions, which causes the number of function evaluations required by the tensor method to be inflated.

The tensor method solved a total of four nonsingular problems, five rank $n-1$ problems, and 7 rank $n-2$ problems, stil Newton's method failed to solve. The reverse never occurred. This clearly indicates that the tensor method is most likely to be more robust than Newton's method.

The overall results show that the tensor method is more efficient than the standard method in solving large, sparse unconstrained optimization problems. Furthermore, the tensor method is likely to solve a wider range of problems.

Table 1: Summary of the CUTE and MINPACK-2 test problems using line search

Rank	Tensor/Standard			Pbs Solved		Average Ratio–Tensor/Standard		
$\nabla^2 f(x_*)$	better	tie	worse	t/s	s/t	feval	geval	time
n	54	38	4	4	0	0.98	0.68	0.63
$n-1$	18	2	0	5	0	0.51	0.48	0.40
$n-2$	18	1	1	7	0	0.66	0.63	0.62

Acknowledgments. I am grateful to Nick Gould for his assistance and encouragements. I also thank my CERFACS colleague Jacko Koster for reviewing this paper and Gail Pieper from the MCS division at Argonne National Laboratory for her suggestions for improvement.

References

[1] B. M. Averick, R. G. Carter, J. J. Moré, and G. L. Xue. The MINPACK-2 test problem collection. Technical Report ANL/MCS-P153-0692, Argonne National Laboratory, 1992.

[2] I. Bongartz, A. R. Conn, N. I. M. Gould, and Ph. L. Toint. CUTE: Constrained and Unconstrained Testing Environment. Technical Report TR/PA/93/10, Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), Toulouse, France, 1993.

[3] A. Bouaricha. *Solving large sparse systems of nonlinear equations and nonlinear least squares problems using tensor methods on sequential and parallel computers.* Ph.D. thesis, Computer Science Department, University of Colorado at Boulder, 1992.

[4] A. Bouaricha. Tensor methods for large, sparse unconstrained optimization. Technical Report ANL/MCS-P452-0794, Argonne National Laboratory, 1994.

[5] A. Bouaricha and R. B. Schnabel. TENSOLVE: a software package for solving systems of nonlinear equations and nonlinear least squares problems using tensor methods. Technical Report CU-CS-735-94, Department of Computer Science, University of Colorado at Boulder, 1994.

[6] T. F. Coleman, B. S. Garbow, and J. J. Moré. Estimating sparse Hessian matrices. *ACM Trans. Math. Software*, 11:363–377, 1985.

[7] J. E. Dennis and R. B. Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations.* Prentice-Hall, Englewood Cliffs, N.J., 1983.

[8] I. S. Duff and J. K. Reid. MA27: A set of Fortran subroutines for solving sparse symmetric sets of linear equations. Technical Report R-10533, AERE Harwell Laboratory, Harwell, UK, 1983.

[9] P. E. Gill, W. Murray, D. B. Ponceleon, and M. A. Saunders. Preconditioners for indefinite systems arising in optimization and nonlinear least squares problems. Technical Report SOL 90-8, Department of Operations Research, Stanford University, California, 1990.

[10] J. J. Moré, B. S. Garbow, and K. E. Hillstrom. Testing unconstrained optimization software. *ACM Trans. Math. Software*, 7(1):17–41, 1981.

[11] R. B. Schnabel and T. Chow. Tensor methods for unconstrained optimization using second derivatives. *SIAM J. Optimization*, 1:293–315, 1991.

[12] R. B. Schnabel and P. D. Frank. Tensor methods for nonlinear equations. *SIAM J. Numer. Anal.*, 21:815–843, 1984.

Appendix: Another Example of Use: The Optimal Design Problem

In this example, we first call the routine DFAULT, which returns the default values. We then override the values of ILIM, GRADTL and GRDFLG. Next we call either the interface TENSS or TENSD for the single– and double–precision version, respectively, to solve the optimal design problem coded in DODCF and whose gradient is given by DODCG.

```
C
C EXAMPLE OF USE FOR STENMIN.  THE TEST PROBLEM IS THE
C OPTIMAL DESIGN WITH COMPOSITE MATERIALS PROBLEM FROM
C THE MINPACK-2 TEST PROBLEM COLLECTION.
C
C  ALI BOUARICHA, OCTOBER 1994.
C  MCS DIVISION, ARGONNE NATIONAL LAB.
C
      INTEGER          NMAX, N, NZ, LIRN, LICN, ILIM, IPR, METHOD
      INTEGER          GRDFLG, HSNFLG, NDIGIT, MSG, LWRK, LIWRK
      INTEGER          TERMCD, INFORM, I, J, K, NX, NY
CD    DOUBLE PRECISION FSCALE, GRADTL, STEPTL, FPLS, STEPMX
CD    DOUBLE PRECISION LAMBDA, HX, HY, TEMP, ONE
CS    REAL             FSCALE, GRADTL, STEPTL, FPLS, STEPMX
CS    REAL             LAMBDA, HX, HY, TEMP, ONE
      PARAMETER      ( NMAX = 10000, LIRN = 50000, LICN = 500000 )
      PARAMETER      ( LIWRK = 2 * LIRN + 12 * NMAX + 2 )
      PARAMETER      ( LWRK = 7 * NMAX )
      INTEGER          IRN ( LIRN  ), ICN ( LICN )
      INTEGER          IWRK( LIWRK )
CD    DOUBLE PRECISION X   ( NMAX ),  TYPX( NMAX ), XPLS( NMAX )
CD    DOUBLE PRECISION GPLS( NMAX ),  HESS( LICN ), WRK ( LWRK )
CD    DOUBLE PRECISION HTV ( NMAX )
CS    REAL             X   ( NMAX ),  TYPX( NMAX ), XPLS( NMAX )
CS    REAL             GPLS( NMAX ),  HESS( LICN ), WRK ( LWRK )
CS    REAL             HTV ( NMAX )
      COMMON / PARAM / NX, NY
      COMMON / OTHER / LAMBDA
      EXTERNAL         DODCF, DODCG, DUSH
CD    INTRINSIC        DBLE, MIN
CS    INTRINSIC        FLOAT, MIN
CD    DATA ONE / 1.0D0 /
CS    DATA ONE / 1.0E0 /
C
C READ DATA
C
      READ(5,*) NX, NY, LAMBDA
      N = NX * NY
```

```
C
C  COMPUTE THE STANDARD STARTING POINT.
C
CD    HX = ONE/DBLE(NX+1)
CD    HY = ONE/DBLE(NY+1)
CS    HX = ONE/FLOAT(NX+1)
CS    HY = ONE/FLOAT(NY+1)
      DO 20 J = 1, NY
CD       TEMP = DBLE(MIN(J,NY-J+1))*HY
CS       TEMP = FLOAT(MIN(J,NY-J+1))*HY
         DO 10 I = 1, NX
            K = NX*(J-1) + I
CD          X(K) = -(MIN(DBLE(MIN(I,NX-I+1))*HX,TEMP))**2
CS          X(K) = -(MIN(FLOAT(MIN(I,NX-I+1))*HX,TEMP))**2
   10    CONTINUE
   20 CONTINUE
C
C  DEFINE THE SPARSITY STRUCTURE OF THE HESSIAN.
C
      CALL DODCSP(NX,NY,NZ,IRN,ICN)
C
C SET THE DEFAULT VALUES OF THE PACKAGE.
C
      CALL DFAULT(N, TYPX, FSCALE, GRADTL, STEPTL, ILIM, STEPMX,
     *            IPR, METHOD, GRDFLG, HSNFLG, NDIGIT, INFORM, MSG)
C
      ILIM = 500
CD    GRADTL = 1.0D-5
CS    GRADTL = 1.0E-3
      GRDFLG = 2
C
C  CALL THE SPARSE OPTIMIZER.
C
CD    CALL TENSD(N,X,NZ,IRN,LIRN,ICN,LICN,DODCF,DODCG,
CS    CALL TENSS(N,X,NZ,IRN,LIRN,ICN,LICN,DODCF,DODCG,
     *  DUSH,TYPX,FSCALE,GRADTL,STEPTL,ILIM,STEPMX,IPR,
     *  METHOD,GRDFLG,HSNFLG,NDIGIT,MSG,XPLS,FPLS,GPLS,
     *  HESS,WRK,LWRK,IWRK,LIWRK,TERMCD,HTV,INFORM)
C
      STOP
      END
```

```
      SUBROUTINE DODCF(N,X,F)

C     **********
C
C  PURPOSE
C  -------
C
C     THIS SUBROUTINE COMPUTES THE FUNCTION OF THE
C     OPTIMAL DESIGN WITH COMPOSITE MATERIALS PROBLEM.
C
C  PARAMETERS
C  ----------
C
C       N IS THE DIMENSION OF THE PROBLEM
C
C       NX IS AN INTEGER VARIABLE.
C         ON ENTRY NX IS THE NUMBER OF GRID POINTS IN THE FIRST
C            COORDINATE DIRECTION.
C         ON EXIT NX IS UNCHANGED.
C
C       NY IS AN INTEGER VARIABLE.
C         ON ENTRY NY IS THE NUMBER OF GRID POINTS IN THE SECOND
C            COORDINATE DIRECTION.
C         ON EXIT NY IS UNCHANGED.
C
C       X IS A DOUBLE PRECISION (REAL) ARRAY OF DIMENSION N = NX*NY.
C
C       F IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON EXIT F IS SET TO THE FUNCTION EVALUATED AT X.
C
C       LAMBDA IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY LAMBDA IS THE LAGRANGE MULTIPLIER.
C         ON EXIT LAMBDA IS UNCHANGED.
C
C     SUBPROGRAMS CALLED
C
C       MINPACK-SUPPLIED   ...   DODCPS
C
C     MINPACK-2 PROJECT. NOVEMBER 1993.
C     ARGONNE NATIONAL LABORATORY AND UNIVERSITY OF MINNESOTA.
C     BRETT M. AVERICK. MODIFIED BY ALI BOUARICHA ON OCTOBER 1994.
C
C     **********
C
```

```
      INTEGER N, NX, NY
CD    DOUBLE PRECISION F, LAMBDA
CD    DOUBLE PRECISION X(N)
CS    REAL F, LAMBDA
CS    REAL X(N)
      COMMON/PARAM/NX,NY
      COMMON/OTHER/LAMBDA
C
C  LOCAL VARIABLES
C
      INTEGER I, J, K
CD    DOUBLE PRECISION MU1, MU2, ONE, P5, TWO, ZERO
CD    DOUBLE PRECISION AREA, DPSI, DVDX, DVDY, GRADV, HX, HXHY
CD    DOUBLE PRECISION HY, TEMP, T1, T2, V, VB, VL, VR, VT
CS    REAL MU1, MU2, ONE, P5, TWO, ZERO
CS    REAL AREA, DPSI, DVDX, DVDY, GRADV, HX, HXHY
CS    REAL HY, TEMP, T1, T2, V, VB, VL, VR, VT
CD    PARAMETER (ZERO=0.0D0,P5=0.5D0,ONE=1.0D0,TWO=2.0D0)
CS    PARAMETER (ZERO=0.0E0,P5=0.5E0,ONE=1.0E0,TWO=2.0E0)
      PARAMETER (MU1=ONE,MU2=TWO)
CD    INTRINSIC DBLE, SQRT
CS    INTRINSIC FLOAT, SQRT
      EXTERNAL DODCPS
C
C  INITIALIZATION.
C
CD    HX = ONE/DBLE(NX+1)
CD    HY = ONE/DBLE(NY+1)
CS    HX = ONE/FLOAT(NX+1)
CS    HY = ONE/FLOAT(NY+1)
      HXHY = HX*HY
      AREA = P5*HXHY
C
C  COMPUTE THE BREAK POINTS.
C
      T1 = SQRT(TWO*LAMBDA*MU1/MU2)
      T2 = SQRT(TWO*LAMBDA*MU2/MU1)
      F = ZERO
C
C  COMPUTATION OF THE FUNCTION OVER THE LOWER
C  TRIANGULAR ELEMENTS.
C
      DO 50 J = 0, NY
         DO 40 I = 0, NX
```

```
            K = NX*(J-1) + I
            V = ZERO
            VR = ZERO
            VT = ZERO
            IF (J .GE. 1 .AND. I .GE. 1) V = X(K)
            IF (I .LT. NX .AND. J .GT. 0) VR = X(K+1)
            IF (I .GT. 0 .AND. J .LT. NY) VT = X(K+NX)
            DVDX = (VR-V)/HX
            DVDY = (VT-V)/HY
            GRADV = DVDX**2 + DVDY**2
            CALL DODCPS(GRADV,MU1,MU2,T1,T2,DPSI,0,LAMBDA)
            F = F + DPSI
   40    CONTINUE
   50 CONTINUE
C
C  COMPUTATION OF THE FUNCTION OVER THE UPPER
C  TRIANGULAR ELEMENTS.
C
      DO 70 J = 1, NY + 1
         DO 60 I = 1, NX + 1
            K = NX*(J-1) + I
            VB = ZERO
            VL = ZERO
            V = ZERO
            IF (I .LE. NX .AND. J .GT. 1) VB = X(K-NX)
            IF (I .GT. 1 .AND. J .LE. NY) VL = X(K-1)
            IF (I .LE. NX .AND. J .LE. NY) V = X(K)
            DVDX = (V-VL)/HX
            DVDY = (V-VB)/HY
            GRADV = DVDX**2 + DVDY**2
            CALL DODCPS(GRADV,MU1,MU2,T1,T2,DPSI,0,LAMBDA)
            F = F + DPSI
   60    CONTINUE
   70 CONTINUE
C
C  SCALE THE FUNCTION.
C
      F = AREA*F
C
C  INTEGRATE V OVER THE DOMAIN.
C
      TEMP = ZERO
      DO 80 K = 1, NX*NY
         TEMP = TEMP + X(K)
```

```
   80 CONTINUE
      F = F + HXHY*TEMP
      END
C
      SUBROUTINE DODCG(N,X,FGRAD)

C     **********
C
C  PURPOSE
C  -------
C
C     THIS SUBROUTINE COMPUTES THE GRADIENT OF THE
C     OPTIMAL DESIGN WITH COMPOSITE MATERIALS PROBLEM.
C
C  PARAMETERS
C  ----------
C
C       N IS THE DIMENSION OF THE PROBLEM
C
C       NX IS AN INTEGER VARIABLE.
C         ON ENTRY NX IS THE NUMBER OF GRID POINTS IN THE FIRST
C            COORDINATE DIRECTION.
C         ON EXIT NX IS UNCHANGED.
C
C       NY IS AN INTEGER VARIABLE.
C         ON ENTRY NY IS THE NUMBER OF GRID POINTS IN THE SECOND
C            COORDINATE DIRECTION.
C         ON EXIT NY IS UNCHANGED.
C
C       X IS A DOUBLE PRECISION (REAL) ARRAY OF DIMENSION N = NX*NY.
C
C       FGRAD IS A DOUBLE PRECISION (REAL) ARRAY OF DIMENSION N = NX*NY.
C         ON ENTRY FGRAD NEED NOT BE SPECIFIED.
C         ON EXIT FGRAD CONTAINS THE GRADIENT EVALUATED AT X.
C
C       LAMBDA IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY LAMBDA IS THE LAGRANGE MULTIPLIER.
C         ON EXIT LAMBDA IS UNCHANGED.
C
C     SUBPROGRAMS CALLED
C
C       MINPACK-SUPPLIED   ...   DODCPS
C
C     MINPACK-2 PROJECT. NOVEMBER 1993.
```

```
C     ARGONNE NATIONAL LABORATORY AND UNIVERSITY OF MINNESOTA.
C     BRETT M. AVERICK. MODIFIED BY ALI BOUARICHA ON OCTOBER 1994.
C
C     **********
C
      INTEGER N, NX, NY
CD    DOUBLE PRECISION X(N), FGRAD(N), LAMBDA
CS    REAL X(N), FGRAD(N), LAMBDA
      COMMON/PARAM/NX,NY
      COMMON/OTHER/LAMBDA
C
C  LOCAL VARIABLES
C
      INTEGER I, J, K
CD    DOUBLE PRECISION MU1, MU2, ONE, P5, TWO, ZERO
CD    DOUBLE PRECISION AREA, DPSIP, DVDX, DVDY, GRADV, HX, HXHY
CD    DOUBLE PRECISION HY, T1, T2, V, VB, VL, VR, VT
CS    REAL MU1, MU2, ONE, P5, TWO, ZERO
CS    REAL AREA, DPSIP, DVDX, DVDY, GRADV, HX, HXHY
CS    REAL HY, T1, T2, V, VB, VL, VR, VT
CD    PARAMETER (ZERO=0.0D0,P5=0.5D0,ONE=1.0D0,TWO=2.0D0)
CS    PARAMETER (ZERO=0.0E0,P5=0.5E0,ONE=1.0E0,TWO=2.0E0)
      PARAMETER (MU1=ONE,MU2=TWO)
CD    INTRINSIC DBLE, SQRT
CS    INTRINSIC FLOAT, SQRT
      EXTERNAL DODCPS

C  INITIALIZATION.

CD    HX = ONE/DBLE(NX+1)
CD    HY = ONE/DBLE(NY+1)
CS    HX = ONE/FLOAT(NX+1)
CS    HY = ONE/FLOAT(NY+1)
      HXHY = HX*HY
      AREA = P5*HXHY
C
C  COMPUTE THE BREAK POINTS.
C
      T1 = SQRT(TWO*LAMBDA*MU1/MU2)
      T2 = SQRT(TWO*LAMBDA*MU2/MU1)

      DO 30 K = 1, NX*NY
         FGRAD(K) = ZERO
   30 CONTINUE
```

```
C
C  COMPUTATION OF THE THE GRADIENT OVER THE LOWER
C  TRIANGULAR ELEMENTS.
C
      DO 50 J = 0, NY
         DO 40 I = 0, NX
            K = NX*(J-1) + I
            V = ZERO
            VR = ZERO
            VT = ZERO
            IF (J .GE. 1 .AND. I .GE. 1) V = X(K)
            IF (I .LT. NX .AND. J .GT. 0) VR = X(K+1)
            IF (I .GT. 0 .AND. J .LT. NY) VT = X(K+NX)
            DVDX = (VR-V)/HX
            DVDY = (VT-V)/HY
            GRADV = DVDX**2 + DVDY**2
            CALL DODCPS(GRADV,MU1,MU2,T1,T2,DPSIP,1,LAMBDA)
            IF (I .GE. 1 .AND. J .GE. 1)
     +          FGRAD(K) = FGRAD(K) - TWO*(DVDX/HX+DVDY/HY)*DPSIP
            IF (I .LT. NX .AND. J .GT. 0)
     +          FGRAD(K+1) = FGRAD(K+1) + TWO*(DVDX/HX)*DPSIP
            IF (I .GT. 0 .AND. J .LT. NY)
     +          FGRAD(K+NX) = FGRAD(K+NX) + TWO*(DVDY/HY)*DPSIP
   40    CONTINUE
   50 CONTINUE
C
C  COMPUTATION OF THE GRADIENT OVER THE UPPER
C  TRIANGULAR ELEMENTS.
C
      DO 70 J = 1, NY + 1
         DO 60 I = 1, NX + 1
            K = NX*(J-1) + I
            VB = ZERO
            VL = ZERO
            V = ZERO
            IF (I .LE. NX .AND. J .GT. 1) VB = X(K-NX)
            IF (I .GT. 1 .AND. J .LE. NY) VL = X(K-1)
            IF (I .LE. NX .AND. J .LE. NY) V = X(K)
            DVDX = (V-VL)/HX
            DVDY = (V-VB)/HY
            GRADV = DVDX**2 + DVDY**2
            CALL DODCPS(GRADV,MU1,MU2,T1,T2,DPSIP,1,LAMBDA)
            IF (I .LE. NX .AND. J .GT. 1)
     +          FGRAD(K-NX) = FGRAD(K-NX) - TWO*(DVDY/HY)*DPSIP
```

```
            IF (I .GT. 1 .AND. J .LE. NY)
     +          FGRAD(K-1) = FGRAD(K-1) - TWO*(DVDX/HX)*DPSIP
            IF (I .LE. NX .AND. J .LE. NY)
     +          FGRAD(K) = FGRAD(K) + TWO*(DVDX/HX+DVDY/HY)*DPSIP
   60    CONTINUE
   70 CONTINUE
C
C  INTEGRATE V OVER THE DOMAIN.
C
      DO 90 K = 1, NX*NY
         FGRAD(K) = AREA*FGRAD(K) + HXHY
   90 CONTINUE
      END
C
      SUBROUTINE DODCPS(T,MU1,MU2,T1,T2,RESULT,OPTION,LAMBDA)
C
C     **********
C
C     THIS SUBROUTINE COMPUTES THE FUNCTION PSI(T) AND THE SCALED
C     FUNCTIONS PSI'(T)/T AND PSI''(T)/T FOR THE OPTIMAL DESIGN
C     WITH COMPOSITE MATERIALS PROBLEM.
C
C     THE SUBROUTINE STATEMENT IS
C
C       SUBROUTINE DODCPS(T,MU1,MU2,T1,T2,RESULT,OPTION,LAMBDA)
C
C     WHERE
C
C       T IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY T IS THE VARIABLE T
C         ON EXIT T IS UNCHANGED
C
C       MU1 IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY MU1 IS THE RECIPROCAL SHEAR MODULUS OF MATERIAL 1.
C         ON EXIT MU1 IS UNCHANGED.
C
C       MU2 IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY MU2 IS THE RECIPROCAL SHEAR MODULUS OF MATERIAL 2.
C         ON EXIT MU2 IS UNCHANGED.
C
C       T1 IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY T1 IS THE FIRST BREAKPOINT.
C         ON EXIT T1 IS UNCHANGED.
C
```

```
C       T2 IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY T2 IS THE SECOND BREAKPOINT.
C         ON EXIT T2 IS UNCHANGED.
C
C       RESULT IS A DOUBLE PRECISION (REAL) VARIABLE.
C         ON ENTRY RESULT NEED NOT BE SPECIFIED.
C         ON EXIT RESULT IS SET ACCORDING TO TASK.
C
C       OPTION IS AN INTEGER VARIABLE.
C         ON ENTRY OPTION SPECIFIES THE ACTION OF THE SUBROUTINE:
C
C            IF OPTION = 0 THEN EVALUATE THE FUNCTION PSI(T).
C            IF OPTION = 1 THEN EVALUATE THE SCALED FUNCTION PSI'(T)/T.
C            IF OPTION = 2 THEN EVALUATE THE SCALED FUNCTION PSI''(T)/T.
C
C        ON OPTION TASK IS UNCHANGED.
C
C       LAMBDA IS A DOUBLE PRECISION (REAL) VARIABLE
C         ON ENTRY LAMBDA IS THE LAGRANGE MULTIPLIER.
C         ON EXIT LAMBDA IS UNCHANGED.
C
C     MINPACK-2 PROJECT. NOVEMBER 1993.
C     ARGONNE NATIONAL LABORATORY AND UNIVERSITY OF MINNESOTA.
C     BRETT M. AVERICK. MODIFIED BY ALI BOUARICHA ON OCTOBER 1994.
C
C     **********
C
      INTEGER OPTION
CD    DOUBLE PRECISION T, MU1, MU2, T1, T2, RESULT, LAMBDA
CS    REAL T, MU1, MU2, T1, T2, RESULT, LAMBDA
C
C  LOCAL VARIABLES
C
CD    DOUBLE PRECISION P25, P5, ZERO
CD    PARAMETER (ZERO=0.0D0,P25=0.25D0,P5=0.5D0)
CS    REAL P25, P5, ZERO
CS    PARAMETER (ZERO=0.0E0,P25=0.25E0,P5=0.5E0)
CD    DOUBLE PRECISION SQRTT
CS    REAL SQRTT
      INTRINSIC SQRT
C
      SQRTT = SQRT(T)
      IF (OPTION .EQ. 0) THEN
         IF (SQRTT .LE. T1) THEN
```

```
            RESULT = P5*MU2*T
         ELSE IF (SQRTT .GT. T1 .AND. SQRTT .LT. T2) THEN
            RESULT = MU2*T1*SQRTT - LAMBDA*MU1
         ELSE IF (SQRTT .GE. T2) THEN
            RESULT = P5*MU1*T + LAMBDA*(MU2-MU1)
         END IF
      ELSE IF (OPTION .EQ. 1) THEN
         IF (SQRTT .LE. T1) THEN
            RESULT = P5*MU2
         ELSE IF (SQRTT .GT. T1 .AND. SQRTT .LT. T2) THEN
            RESULT = P5*MU2*T1/SQRTT
         ELSE IF (SQRTT .GE. T2) THEN
            RESULT = P5*MU1
         END IF
      ELSE IF (OPTION .EQ. 2) THEN
         IF (SQRTT .LE. T1) THEN
            RESULT = ZERO
         ELSE IF (SQRTT .GT. T1 .AND. SQRTT .LT. T2) THEN
            RESULT = -P25*MU2*T1/(SQRTT*T)
         ELSE IF (SQRTT .GE. T2) THEN
            RESULT = ZERO
         END IF
      END IF
      END
C
      SUBROUTINE DODCSP(NX,NY,NNZ,INDROW,INDCOL)
C
C     **********
C
C     SUBROUTINE DODCSP
C
C     THIS SUBROUTINE DEFINES THE SPARSITY STRUCTURE OF THE HESSIAN
C     MATRIX FOR THE OPTIMAL DESIGN WITH COMPOSITES PROBLEM.
C
C     THE SUBROUTINE STATEMENT IS
C
C       SUBROUTINE DODCSP(NX,NY,NNZ,INDROW,INDCOL)
C
C     WHERE
C
C       NX IS AN INTEGER VARIABLE.
C         ON ENTRY NX IS THE NUMBER OF GRID POINTS IN THE FIRST
C            COORDINATE DIRECTION.
C         ON EXIT NX IS UNCHANGED.
```

```
C
C       NY IS AN INTEGER VARIABLE.
C         ON ENTRY NY IS THE NUMBER OF GRID POINTS IN THE SECOND
C            COORDINATE DIRECTION.
C         ON EXIT NY IS UNCHANGED.
C
C       NNZ IS AN INTEGER VARIABLE.
C         ON ENTRY NNZ NEED NOT BE SPECIFIED.
C         ON EXIT NNZ IS SET TO THE NUMBER OF NONZEROS IN THE
C            LOWER TRIANGLE OF THE HESSIAN MATRIX.
C
C       INDROW IS AN INTEGER ARRAY OF DIMENSION AT LEAST NNZ.
C         ON ENTRY INDROW NEED NOT BE SPECIFIED.
C         ON EXIT INDROW CONTAINS THE ROW INDICES OF THE NONZEROS
C            IN THE LOWER TRIANGLE OF THE HESSIAN MATRIX.
C
C       INDCOL IS AN INTEGER ARRAY OF DIMENSION AT LEAST NNZ.
C         ON ENTRY INDCOL NEED NOT BE SPECIFIED.
C         ON EXIT INDCOL CONTAINS THE COLUMN INDICES OF THE NONZEROS
C            IN THE LOWER TRIANGLE OF THE HESSIAN MATRIX.
C
C     MINPACK-2 PROJECT. NOVEMBER 1993.
C     ARGONNE NATIONAL LABORATORY AND UNIVERSITY OF MINNESOTA.
C     BRETT M. AVERICK. MODIFIED BY ALI BOUARICHA ON OCTOBER 1994.
C
C     **********
C
      INTEGER NX, NY, NNZ
      INTEGER INDROW(*), INDCOL(*)
C
C  LOCAL VARIABLES
C
      INTEGER I, J
C
C     COMPUTE THE SPARSITY STRUCTURE.
C
      NNZ = 0
      DO 20 J = 1, NY
         DO 10 I = 1, NX
            NNZ = NNZ + 1
            INDROW(NNZ) = (J-1)*NX + I
            INDCOL(NNZ) = (J-1)*NX + I
            IF (I .NE. NX) THEN
               NNZ = NNZ + 1
```

```
            INDROW(NNZ) = (J-1)*NX + I + 1
            INDCOL(NNZ) = (J-1)*NX + I
         END IF
         IF (J .NE. NY) THEN
            NNZ = NNZ + 1
            INDROW(NNZ) = (J-1)*NX + I + NX
            INDCOL(NNZ) = (J-1)*NX + I
            IF (I .NE. 1) THEN
               NNZ = NNZ + 1
               INDROW(NNZ) = (J-1)*NX + I + NX - 1
               INDCOL(NNZ) = (J-1)*NX + I
            END IF
         END IF
10    CONTINUE
20 CONTINUE
   END
```

Figure 2: Code to solve the optimal design with composite materials problem

If we use the double–precision version of the package to solve the optimal design problem given by DODCF, for the following input:

```
NX, NY, LAMBDA :   100 100 0.008,
```

we obtain the following output:

```
 OPTIM      GRADIENT FLAG      = 2
 OPTIM      HESSIAN FLAG       = 0
 OPTIM      METHOD             = 1
 OPTIM      ITERATION LIMIT    =  500
 OPTIM      MACHINE EPSILON    =  0.2220446049250E-15
 OPTIM      STEP TOLERANCE     =  0.3666852862501E-10
 OPTIM      GRADIENT TOLERANCE =  0.1000000000000E-04
 OPTIM      MAXIMUM STEP SIZE  =  0.6521118878154E+04

 ----------------------------------------
 RESULT    ITERATION K   =    0
 RESULT    FUNCTION AT X(K)
 RESULT        0.4823420295546E-01
 RESULT    MAXIMUM RELATIVE GRADIENT AT X(K)
 RESULT        0.1931183217332E-01
 ----------------------------------------
```

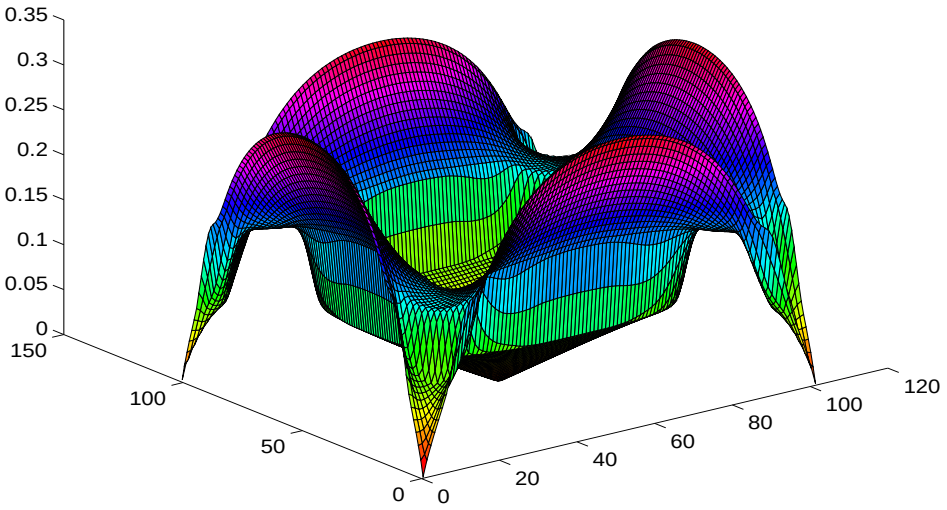

Figure 3: Norm of $||\nabla v||$ for the stress field v in a design with composite materials

```
OPTSTP    RELATIVE GRADIENT CLOSE TO ZERO
OPTSTP    CURRENT ITERATE IS PROBABLY SOLUTION

-----------------------------------------
RESULT    ITERATION K   =   20
RESULT    FUNCTION AT X(K)
RESULT       -0.1137724408643E-01
RESULT    SCALED GRADIENT AT X(K)
RESULT        0.3938142592477E-05
-----------------------------------------

RESULT    NUMBER OF FUNCTION EVALUATIONS   67
RESULT    NUMBER OF GRADIENT EVALUATIONS   21
RESULT    NUMBER OF HESSIAN EVALUATIONS   20
```

A plot of the norm $||\nabla v||$ of the gradient of the stress field v in the bounded domain $D = (0,1) \times (0,1)$ where LAMBDA $= 0.008$ is given in Figure 3. Figure 4 shows the contour plot for this surface.

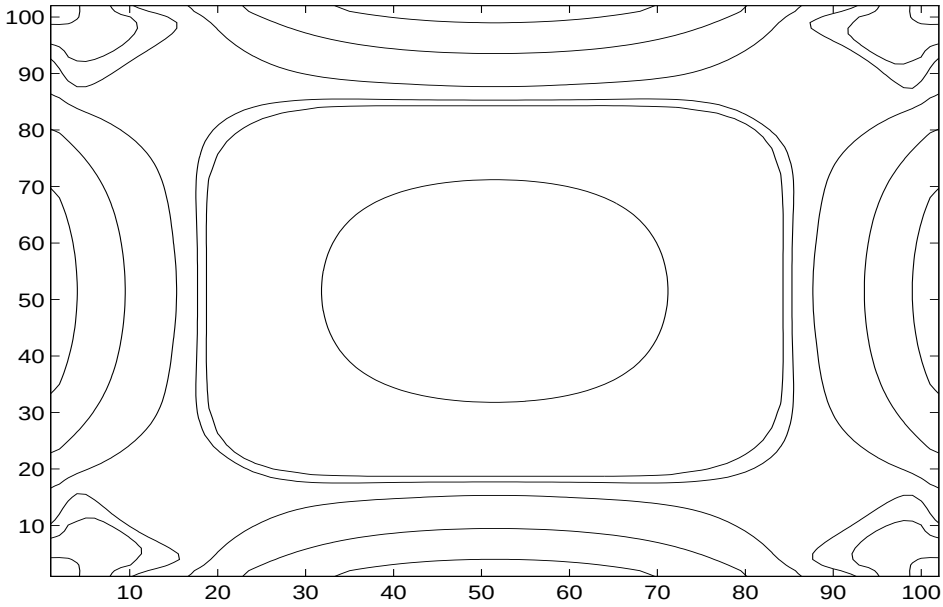

Figure 4: Contours of $||\nabla v||$ for the stress field v in a design with composite materials