Ladislav Lukšan; Ctirad Matonoha; Jan Vlček
On Lagrange multipliers of trust-region subproblems

In: Jan Chleboun and Petr Přikryl and Karel Segeth and Tomáš Vejchodský (eds.): Programs and Algorithms of Numerical Mathematics, Proceedings of Seminar. Dolní Maxov, June 1-6, 2008. Institute of Mathematics AS CR, Prague, 2008. pp. 130–137.

Persistent URL: http://dml.cz/dmlcz/702866

ON LAGRANGE MULTIPLIERS OF TRUST-REGION SUBPROBLEMS*

Ladislav Lukšan, Ctirad Matonoha, Jan Vlček

1. Introduction

Consider the problem

$$\min F(x), \quad x \in \mathcal{R}^n,$$

where $F : \mathcal{R}^n \to \mathcal{R}$ is a twice continuously differentiable objective function. Basic optimization methods (trust-region and line-search methods) [6] generate points $x_i \in \mathcal{R}^n, i \in \mathcal{N}$, in such a way that x_1 is arbitrary and

$$x_{i+1} = x_i + \alpha_i d_i, \quad i \in \mathcal{N},$$

where $d_i \in \mathcal{R}^n$ are direction vectors and $\alpha_i > 0$ are step sizes.

Trust-region methods [1] are globally convergent techniques widely used, for example, in connection with the Newton's method for unconstrained optimization. They can be advantageously used when the Hessian matrix (or its approximation) of the objective function is indefinite, ill-conditioned or singular. This situation often arises in connection with the Newton's method for general objective function (indefiniteness) or with the Gauss-Newton's method for nonlinear least-squares problems (near singularity).

For a description of trust-region methods we define the quadratic function

$$Q_i(d) = \frac{1}{2} d^T B_i d + g_i^T d$$

which locally approximates the difference $F(x_i + d) - F(x_i)$, the vector

$$\omega_i(d) = \frac{(B_i d + g_i)}{\|g_i\|}$$

for the accuracy of computed direction, and the number

$$\rho_i(d) = \frac{F(x_i + d) - F(x_i)}{Q_i(d)}$$

*This work was supported by the Grant Agency of the Czech Academy of Sciences, project No. IAA1030405, the Grant Agency of the Czech Republic, project No. 201/06/P397, and the institutional research plan No. AV0Z10300504.

for the ratio of actual and predicted decrease of the objective function. Here $g_i = g(x_i) = \nabla F(x_i)$ and $B_i \approx \nabla^2 F(x_i)$ is an approximation of the Hessian matrix of the function F at the point $x_i \in \mathcal{R}^n$.

Trust-region methods are based on approximate minimizations of $Q_i(d)$ on the balls $\|d\| \leq \Delta_i$ followed by updates of radii $\Delta_i > 0$. Thus direction vectors $d_i \in \mathcal{R}^n$ are chosen to satisfy the conditions

$$\begin{aligned} \|d_i\| &\leq \Delta_i, && (1)\\ \|d_i\| &< \Delta_i \quad \Rightarrow \quad \|\omega_i(d_i)\| \leq \overline{\omega}, && (2)\\ -Q_i(d_i) &\geq \underline{\sigma}\|g_i\| \min(\|d_i\|, \|g_i\|/\|B_i\|), && (3) \end{aligned}$$

where $0 \leq \overline{\omega} < 1$ and $0 < \underline{\sigma} < 1$. Step sizes $\alpha_i \geq 0$ are selected so that

$$\begin{aligned} \rho_i(d_i) \leq 0 &\Rightarrow \alpha_i = 0, && (4)\\ \rho_i(d_i) > 0 &\Rightarrow \alpha_i = 1. && (5) \end{aligned}$$

Trust-region radii $0 < \Delta_i \leq \overline{\Delta}$ are chosen in such a way that $0 < \Delta_1 \leq \overline{\Delta}$ is arbitrary and

$$\begin{aligned} \rho_i(d_i) < \underline{\rho} &\Rightarrow \underline{\beta}\|d_i\| \leq \Delta_{i+1} \leq \overline{\beta}\|d_i\|, && (6)\\ \rho_i(d_i) \geq \underline{\rho} &\Rightarrow \Delta_i \leq \Delta_{i+1} \leq \overline{\Delta}, && (7) \end{aligned}$$

where $0 < \underline{\beta} \leq \overline{\beta} < 1$ and $0 < \underline{\rho} < 1$.

2. Direction determination

A crucial part of each trust-region method is the direction determination. There are various commonly known methods for computing direction vectors satisfying conditions (1)-(3). To simplify the notation, we omit the major index i and use the inner index j.

The most sophisticated method is based on a computation of the optimal locally constrained step. In this case, the vector $d \in \mathcal{R}^n$ is obtained by solving the subproblem

$$\min \; Q(d) = \frac{1}{2} d^T B d + g^T d \quad \text{subject to} \quad \|d\| \leq \Delta. \tag{8}$$

Necessary and sufficient conditions for this solution are

$$\|d\| \leq \Delta, \quad (B + \lambda I)d = -g, \quad B + \lambda I \succeq 0, \quad \lambda \geq 0, \quad \lambda(\Delta - \|d\|) = 0,$$

where λ is the optimal Lagrange multiplier. The Moré-Sorensen method [5] is based on solving the nonlinear equation $1/\|d(\lambda)\| = 1/\Delta$ with $(B + \lambda I)d(\lambda) + g = 0$ by the Newton's method, possibly the modified Newton's method [8] using the Choleski decomposition of $B + \lambda I$.

Steihaug [7] and Toint [9] proposed a technique for finding an approximate solution of (8). This implementation is based on the conjugate gradient algorithm [6] for

solving the linear system $Bd = -g$. We either obtain an unconstrained solution with a sufficient precision or stop on the trust-region boundary. The latter possibility occurs if either a negative curvature is encountered or the constraint is violated. This method is based on the fact that $Q(d_{j+1}) < Q(d_j)$ and $\|d_{j+1}\| > \|d_j\|$ hold in the subsequent CG iterations if the CG coefficients are positive and no preconditioning used. For a general preconditioner C (symmetric and positive definite), we have $\|d_{j+1}\|_C > \|d_j\|_C$, where $\|d_j\|_C^2 = d_j^T C d_j$.

There are two possibilities how the Steihaug-Toint method can be preconditioned. The first way uses the norms $\|d_i\|_{C_i}$ (instead of $\|d_i\|$) in (1)–(7), where C_i are preconditioners chosen. This possibility is not always efficient because the norms $\|d_i\|_{C_i}$, $i \in \mathcal{N}$, vary considerably in the major iterations and preconditioners C_i, $i \in \mathcal{N}$, can be ill-conditioned. The second way uses Euclidean norms in (1)–(7), even if arbitrary preconditioners C_i, $i \in \mathcal{N}$, are used. In this case, the trust-region can be leaved prematurely and the direction vector obtained can be farther from the optimal locally-constrained step than that obtained without preconditioning. This shortcoming is usually compensated by the rapid convergence of the preconditioned CG method. Our computational experiments indicate that the second way is more efficient in general.

Another approach, the GLRT method [2], approximately solves (8) iteratively by using the symmetric Lanczos process. A vector d_j which is the j-th approximation of optimal d is contained in the Krylov subspace $\mathcal{K}_j = \mathrm{span}\{g, Bg, \ldots, B^{j-1}g\}$ of dimension j defined by the matrix B and the vector g. In this case, $d_j = Z\tilde{d}$, where $\tilde{d}$ is obtained by minimizing the quadratic function

$$\frac{1}{2}\,\tilde{d}^T T \tilde{d} + \|g\| e_1^T \tilde{d}$$

subject to $\|\tilde{d}\| \leq \Delta$. Here, $T = Z^T B Z$ (with $Z^T Z = I$) is the Lanczos tridiagonal matrix and e_1 is the first column of the unit matrix. Using preconditioner C, the preconditioned Lanczos method generates basis such that $Z^T C Z = I$. Thus, we have to use the norms $\|d_i\|_{C_i}$ in (1)–(7), i.e. the first way of preconditioning, which can be inefficient when C_i vary considerably in the trust-region iterations or are ill-conditioned.

3. A shifted Steihaug-Toint method

In this contribution, we consider a sequence of subproblems

$$d_j = \arg\min_{d \in \mathcal{K}_j} Q(d) \quad \text{subject to} \quad \|d\| \leq \Delta, \quad \text{where} \quad Q(d) = \frac{1}{2}\,d^T B d + g^T d,$$

with corresponding Lagrange multipliers λ_j, $j \in \{1, \ldots, n\}$. The method [3] uses the conjugate gradient method applied to the linear system $(B + \tilde{\lambda} I)d + g = 0$, where $\tilde{\lambda}$ is an approximation to the optimal Lagrange multiplier λ. For this reason, we need to investigate the properties of the Lagrange multipliers corresponding to the trust-region subproblems used.

Before showing the main result, we first present several lemmas, which lead to the main theorem. The first lemma, coming from [7], shows a simple property of the conjugate gradient method, the second one compares Krylov subspaces of the matrices B and $B + \lambda I$, and the last one states a relation between the values of the Lagrange multipliers and the norms of the direction vectors.

Lemma 1. *Let B be a symmetric and positive definite matrix, let*

$$\mathcal{K}_j = \operatorname{span}\{g, Bg, \ldots, B^{j-1}g\}, \quad j \in \{1, \ldots, n\},$$

be the j-th Krylov subspace given by the matrix B and the vector g. Let

$$d_j = \arg\min_{d \in \mathcal{K}_j} Q(d), \quad \text{where} \quad Q(d) = \frac{1}{2} d^T B d + g^T d.$$

If $1 \leq k \leq l \leq n$, then

$$\|d_k\| \leq \|d_l\|.$$

Especially,

$$\|d_k\| \leq \|d_n\|, \quad \text{where} \quad d_n = \arg\min_{d \in \mathcal{R}^n} Q(d).$$

Lemma 2. *Let $\lambda \in \mathcal{R}$ and*

$$\mathcal{K}_j(\lambda) = \operatorname{span}\{g, (B + \lambda I)g, \ldots, (B + \lambda I)^{j-1}g\}, \quad j \in \{1, \ldots, n\},$$

be the j-dimensional Krylov subspace generated by the matrix $B+\lambda I$ and the vector g. Then

$$\mathcal{K}_j(\lambda) = \mathcal{K}_j(0).$$

Lemma 3. *Let $Z_j^T B Z_j + \lambda_k I$, $\lambda_k \in \mathcal{R}$, $k \in \{1, 2\}$, where $Z_j \in \mathcal{R}^{n \times j}$ is a matrix whose columns form an orthonormal basis for $\mathcal{K}_j$, be symmetric and positive definite. Let*

$$d_j(\lambda_k) = \arg\min_{d \in \mathcal{K}_j} Q_{\lambda_k}(d), \quad \text{where} \quad Q_\lambda(d) = \frac{1}{2} d^T (B + \lambda I) d + g^T d.$$

Then

$$\lambda_2 \leq \lambda_1 \quad \Leftrightarrow \quad \|d_j(\lambda_2)\| \geq \|d_j(\lambda_1)\|.$$

Now we are in a position to present the main theorem.

Theorem 1. *Let d_j, $j \in \{1, \ldots, n\}$, be solutions of the minimization problems*

$$d_j = \arg\min_{d \in \mathcal{K}_j} Q(d) \quad \text{subject to} \quad \|d\| \leq \Delta, \quad \text{where} \quad Q(d) = \frac{1}{2} d^T B d + g^T d,$$

with corresponding Lagrange multipliers λ_j, $j \in \{1, \ldots, n\}$. If $1 \leq k \leq l \leq n$, then

$$\lambda_k \leq \lambda_l.$$

4. Applications

The result of Theorem 1 can be applied to the following idea. We apply the Steihaug-Toint method to a shifted subproblem

$$\min\ \tilde{Q}(d) = Q_{\tilde{\lambda}}(d) = \frac{1}{2}\, d^T(B + \tilde{\lambda}I)d + g^T d \quad \text{subject to} \quad \|d\| \leq \Delta, \tag{9}$$

where $\tilde{\lambda}$ is an approximation to the optimal λ. If we set $\tilde{\lambda} = \lambda_j$ for some $j \leq n$, then Theorem 1 implies that $0 \leq \tilde{\lambda} = \lambda_j \leq \lambda_n = \lambda$. As a consequence of this inequality, one has that $\lambda = 0$ implies $\tilde{\lambda} = 0$ so that $\|d\| < \Delta$ implies $\tilde{\lambda} = 0$. Thus, the shifted Steihaug-Toint method proposed in [3] reduces to the standard Steihaug-Toint method in this case. At the same time, if B is positive definite and $0 < \tilde{\lambda} \leq \lambda$, then one has $\Delta = \|(B + \lambda I)^{-1}g\| \leq \|(B + \tilde{\lambda}I)^{-1}g\| < \|B^{-1}g\|$ by Lemma 3. Thus, the unconstrained minimizer of (9) is closer to the trust-region boundary than the unconstrained minimizer of (8) and we can expect that $d(\tilde{\lambda})$ is closer to the optimal locally constrained step than $d(0)$. Finally, if B is positive definite and $\tilde{\lambda} > 0$, then the matrix $B + \tilde{\lambda}I$ is better conditioned than B and we can expect that the shifted Steihaug-Toint method will converge more rapidly than the standard Steihaug-Toint method.

The shifted Steihaug-Toint method for solving subproblem (8) consists of the three major steps.

1. Carry out $j \ll n$ steps (usually $j = 5$) of the unpreconditioned Lanczos method to obtain the tridiagonal matrix $T \equiv T_j = Z_j^T B Z_j$.

2. Solve the subproblem

$$\min\ \frac{1}{2}\, \tilde{d}^T T \tilde{d} + \|g\| e_1^T \tilde{d} \quad \text{subject to} \quad \|\tilde{d}\| \leq \Delta,$$

   using the method of Moré and Sorensen, to obtain the Lagrange multiplier $\tilde{\lambda}$.

3. Apply the (preconditioned) Steihaug-Toint method to the shifted subproblem

$$\min\ \tilde{Q}(d) \quad \text{subject to} \quad \|d\| \leq \Delta$$

   to obtain the direction vector $d = d(\tilde{\lambda})$, a suitable approximation to the solution of problem (8).

5. Numerical experiments

We present a numerical comparison of methods for computing direction vectors satisfying conditions (1)-(3):

- MS – the method of Moré and Sorensen [5] for computing the optimal locally constrained step.

- ST – the basic (unpreconditioned) Steihaug [7] and Toint [9] method.
- SST – the basic (unpreconditioned) shifted Steihaug-Toint method [3].
- GLRT – the method of Gould, Lucidi, Roma, and Toint [2] which combines CG method with the Lanczos process to give a good approximation to the optimal locally constrained step.
- PST – the preconditioned Steihaug-Toint method.
- PSST – the preconditioned shifted Steihaug-Toint method.

Note that the incomplete Choleski preconditioner is used for methods PST and PSST, the number of extra CG or Lanczos steps in SST and PSST methods is equal to 5, and the number of Lanczos vectors in the GLRT method is bounded from above by 100.

The methods were tested by using two collections of 22 sparse test problems with 1000 and 5000 variables (subroutines `TEST14` and `TEST15` described in [4], which can be downloaded from `www.cs.cas.cz/luksan/test.html`). The results are given in Tables 1 and 2, where `NIT` is the total number of iterations, `NFV` is the total number of function evaluations, `NFG` is the total number of gradient evaluations, `NDC` is the total number of Choleski-type decompositions (complete for method MS and incomplete for methods PST, PSST), `NMV` is the total number of matrix-vector multiplications, and `Time` is the total computational time in seconds. Note that `NFG` is much greater than `NFV` in Table 1, since the Hessian matrices are computed by using gradient differences. At the same time, the problems referred in Table 2 are the sums of squares having the form $F(x) = (1/2)f^T(x)f(x)$ and `NFV` denotes the total number of vector $f(x)$ evaluations. Since $f(x)$ is used in the expression $g(x) = J^T(x)f(x)$, where $J(x)$ is the Jacobian matrix of $f(x)$, `NFG` is comparable with `NFV` in this case.

Results in Tables 1 and 2 require several comments. All problems are sparse with a simple sparsity pattern. For this reason, the MS method based on complete Choleski-type decompositions (CD) is very efficient, much better than unpreconditioned methods based on matrix-vector multiplications (MV). Since `TEST14` contains reasonably conditioned problems, the preconditioned MV methods are competitive with the CD method. On the contrary, `TEST15` contains several very ill-conditioned problems (one of them had to be removed) and thus, the CD method works better than the MV methods.

6. Conclusion

Our conclusions concern large-scale problems where the sparsity pattern plays a considerable role. The Moré-Sorensen method is very efficient for ill-conditioned but reasonably sparse problems. If the problems do not have sufficiently sparse Hessian matrices, then this method can be much worse than the Steihaug-Toint method whose efficiency also strongly depends on suitable preconditioning. There are two possibilities of preconditioning mentioned in Section 2. The first one changes the trust-region problem whereas the second one deforms the trust-region path in the

N	Method	NIT	NFV	NFG	NDC	NMV	Time
1000	MS	1911	1952	8724	3331	1952	3.13
	ST	3475	4021	17242	0	63016	5.44
	SST	3149	3430	15607	0	75044	5.97
	GLRT	3283	3688	16250	0	64166	5.40
	PST	2608	2806	12802	2609	5608	3.30
	PSST	2007	2077	9239	2055	14440	2.97
5000	MS	8177	8273	34781	13861	8272	49.02
	ST	16933	19138	84434	0	376576	134.52
	SST	14470	15875	70444	0	444142	146.34
	GLRT	14917	16664	72972	0	377588	132.00
	PST	11056	11786	53057	11057	23574	65.82
	PSST	8320	8454	35629	8432	59100	45.57

Tab. 1: *Comparison of methods using TEST14.*

N	Method	NIT	NFV	NFG	NDC	NMV	Time
1000	MS	1946	9094	9038	3669	2023	5.86
	ST	2738	13374	13030	0	53717	11.11
	SST	2676	13024	12755	0	69501	11.39
	GLRT	2645	12831	12547	0	61232	11.30
	PST	3277	16484	16118	3278	31234	11.69
	PSST	2269	10791	10613	2446	37528	8.41
5000	MS	7915	33607	33495	14099	8047	89.69
	ST	11827	54699	53400	0	307328	232.70
	SST	11228	51497	50333	0	366599	231.94
	GLRT	10897	49463	48508	0	300580	214.74
	PST	9360	41524	41130	9361	179166	144.40
	PSST	8634	37163	36881	8915	219801	140.44

Tab. 2: *Comparison of methods using TEST15.*

original trust-region problem. Note that the GLRT method cannot be preconditioned in the second way since the preconditioned Lanczos process does not generate the orthonormal basis related to the original trust-region subproblem. Our preliminary tests have shown that the first preconditioning technique is less efficient because it failed in many cases.

To sum up, the shifted Steihaug-Toint method combines good properties of the Moré-Sorensen and the Steihaug-Toint methods. Our computational experiments indicate that this method works well in connection with the second way of preconditioning. The point on the trust-region boundary obtained by this method is usually closer to the optimal solution in comparison with the point obtained by the original Steihaug-Toint method.

References

[1] A.R. Conn, N.I.M. Gould, P.L. Toint: *Trust-region methods*. SIAM, Philadelphia, 2000.

[2] N.I.M. Gould, S. Lucidi, M. Roma, P.L. Toint: *Solving the trust-region subproblem using the Lanczos method*. Report No. RAL-TR-97-028, Rutherford Appleton Laboratory, 1997.

[3] L. Lukšan, C. Matonoha, J. Vlček: *A shifted Steihaug-Toint method for computing a trust-region step*. Report No. V914-04, Institute of Computer Science, Academy of Sciences, Czech Republic, 2004.

[4] L. Lukšan, J. Vlček: *Sparse and partially separable test problems for unconstrained and equality constrained optimization*. Report No. V767-98, Institute of Computer Science, Academy of Sciences, Czech Republic, 1998.

[5] J.J. Moré, D.C. Sorensen: *Computing a trust region step*. SIAM J. Sci. Statist. Comput. **4** (1983), 553–572.

[6] J. Nocedal, S.J. Wright: *Numerical optimization*. Springer, New York, 1999.

[7] T. Steihaug: *The conjugate gradient method and trust regions in large-scale optimization*. SIAM J. Numer. Anal. **20** (1983), 626–637.

[8] D.C. Sorensen: *Newton's method with a model trust region modification*. SIAM J. Numer. Anal. **19(2)** (1982), 409–426.

[9] P.L. Toint: *Towards an efficient sparsity exploiting Newton method for minimization*. In: I.S. Duff (Ed.), Sparse Matrices and Their Uses. Academic Press, London, 1981, pp. 57–88.