SANDIA REPORT

SAND2004-1944
Unlimited Release
Printed September 2004

On the Performance of Tensor Methods for Solving Ill-conditioned Problems

Brett W. Bader and Robert B. Schnabel

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

SAND2004-1944
Unlimited Release
Printed September 2004

On the Performance of Tensor Methods for Solving Ill-conditioned Problems

Brett W. Bader
Computational Sciences Department
Sandia National Laboratories
P.O. Box 5800, MS-0316
Albuquerque, NM 87185-0316

Robert B. Schnabel
Department of Computer Science
University of Colorado at Boulder
Boulder, CO 80309-0040

Abstract

This paper investigates the performance of tensor methods for solving small- and large-scale systems of nonlinear equations where the Jacobian matrix at the root is ill-conditioned or singular. This condition occurs on many classes of problems, such as identifying or approaching turning points in path following problems. The singular case has been studied more than the highly ill-conditioned case, for both Newton and tensor methods. It is known that Newton-based methods do not work well with singular problems because they converge linearly to the solution and, in some cases, with poor accuracy. On the other hand, direct tensor methods have performed well on singular problems and have superlinear convergence on such problems under certain conditions. This behavior originates from the use of a special, restricted form of the second-order term included in the local tensor model that provides information lacking in a (nearly) singular Jacobian. With several implementations available for large-scale problems, tensor methods now are capable of solving larger problems. We compare the performance of tensor methods and Newton-based methods for both small- and large-scale problems over a range of conditionings, from well-conditioned to ill-conditioned to singular. Previous studies with tensor methods only concerned the ends of this spectrum. Our results show that tensor methods are increasingly superior to Newton-based methods as the problem grows more ill-conditioned.

Acknowledgment

The work of both authors was supported in part by Air Force Office of Scientific Research grant F49620-00-1-0162 and Army Research Office grant DAAG55-98-1-0176 while in the Department of Computer Science at the University of Colorado, Boulder.

While at Sandia, the first author also received support from DOE's Office of Science MICS program and Sandia's Laboratory Directed Research and Development program under LDRD 04-1024.

Contents

1 Introduction **7**

2 Algorithms **9**
2.1 Standard Methods 9
2.2 Tensor methods 9
2.3 Newton-Krylov methods 10
2.4 Tensor-Krylov methods 11
2.5 Tensor-GMRES method 12

3 Numerical experiments on ill-conditioned problems **13**
3.1 Small problems solved with direct methods 13
3.2 Moderate-size problems solved with inexact methods 18

4 Summary and conclusions **28**

References **31**

Figures

1 Example iteration profiles on ill-conditioned problems 16
2 Comparison of tensor vs. Newton on ill-conditioned problems approaching rank $n-1$ 17
3 Comparison of tensor vs. Newton on ill-conditioned problems approaching rank $n-2$ 19
4 Comparison of tensor vs. Newton on ill-conditioned problems approaching rank $n-3$ 20
5 Ill-conditioning tests on the Bratu problem 22
6 Example iteration history on the Bratu problem 23
7 Ill-conditioning tests on the discrete boundary value problem 24
8 Example iteration history on the discrete boundary value problem 26
9 Ill-conditioning tests on the modified Broyden tridiagonal problem 27
10 Example iteration history on the modified Broyden tridiagonal problem 28

Tables

1 Ill-conditioned test problems 15

On the Performance of Tensor Methods for Solving Ill-conditioned Problems

1 Introduction

This paper examines two classes of methods for solving the nonlinear equations problem

$$\text{given } F : \mathbb{R}^n \to \mathbb{R}^n, \quad \text{find } x_* \in \mathbb{R}^n \text{ such that } F(x_*) = 0, \tag{1}$$

where it is assumed that $F(x)$ is at least once continuously differentiable. General systems of nonlinear equations defined by (1) arise in many practical situations, including systems produced by finite-difference or finite-element discretizations of boundary value problems for ordinary and partial differential equations.

As a subset of the general nonlinear equations problem (1), there is a class of important problems where $F'(x_*)$ is singular or, at least, very ill-conditioned. Such examples arise in bifurcation tracking and path following problems where the goal is to locate turning points, such as the ignition and extinction points in chemical combustion. Resolving these features is important to engineers, who, for instance, may be designing control systems for such applications and may need to know important operating boundaries.

Standard methods for solving (1), such as Newton's method, base each iteration upon a local, linear model $M(x_k + d)$ of the function $F(x)$ around the current iterate $x_k \in \mathbb{R}^n$. Standard methods work well for problems where the Jacobian at the solution, $F'(x_*)$, is well-conditioned; but they face difficulties when the Jacobian is singular, or even nearly singular, at the solution. Many authors have analyzed the behavior of Newton's method on singular problems and have proposed acceleration techniques as remedies (see, e.g., Decker, Keller, and Kelley [8]; Decker and Kelley [9, 10, 11]; Griewank [17]; Griewank and Osborne [18]; Kelley and Suresh [20]; and Reddien [24]). Their collective analysis shows that, from many starting points, Newton's method is locally q-linearly convergent with constant converging to $\frac{1}{2}$ on singular problems where the second-order term $F''(x_k)$ contains appropriate null space information. Acceleration techniques can improve this behavior; however, they require *a priori* knowledge that the problem is singular, which is not practical for general problem solving.

Tensor methods, however, do not require a priori knowledge of whether the problem is singular or not. These methods were introduced by Schnabel and Frank in [26] and base each iteration on a simplified quadratic model of $F(x)$ such that the quadratic term is a low-rank secant approximation that augments the standard linear

model. Tensor methods also have been extended by other authors to utilize iterative solvers, making the methods appropriate for solving large-scale problems (see [1, 2], [3], and [15]).

The analysis in [14] proves that direct tensor methods have quadratic convergence on nonsingular problems and a faster convergence rate on problems where the Jacobian matrix at the solution is singular. Specifically, when the rank of the Jacobian at the root is $n-1$, "practical" tensor methods (i.e., those using secant approximations for the tensor term) have three-step superlinear convergence behavior with q-order $\frac{3}{2}$. In practice, one-step superlinear convergence frequently is observed on these problems, which makes the method even more attractive. The second-order term provides higher order information in recent step directions, which aids in cases where the Jacobian is (nearly) singular at the solution. As the iterates approach the solution, the Jacobian lacks information in the null space direction, and the second-order term supplies useful information for a better quality step. Computational evidence in [26] on small problems shows that tensor methods have about 20% average improvement over standard methods on nonsingular problems and about 40% improvement on singular problems with $\text{rank}(F'(x_*)) = n-1$.

While tensor methods have encouraging theoretical and computational results on singular problems, less is known about their performance relative to Newton's method on ill-conditioned problems. Do tensor methods outperform Newton's method due to the close relationship of ill-conditioned matrices with singular matrices? Or do tensor methods only exhibit superior behavior when the problem is truly singular? Does the computational performance of Newton's method degrade gradually as the problem becomes more singular, or sharply at the singularity? The performance comparison over a spectrum of ill-conditioned problems was previously unknown. Thus, this paper examines the performance of tensor methods versus standard methods as the problem becomes more ill-conditioned. We consider tensor methods using direct factorizations of the Jacobian matrix for small-scale problems in addition to Krylov-based iterative tensor methods for large-scale problems.

The organization of this paper is as follows. Because this research involves methods for solving small- to large-scale problems, this paper includes background for both types in section 2. Specifically, we review direct methods for solving small-scale problems, and we review Krylov-based iterative methods for solving large-scale problems, including the relevant algorithms from [1, 2] and [15]. Section 3 presents numerical results on several small- and large-scale ill-conditioned test problems to examine the performance of tensor methods on problems over a range of conditionings, from well-conditioned to singular. Finally, section 4 summarizes the numerical results and provides some concluding remarks.

Throughout this paper, a subscript k refers to the current iterate of a nonlinear solver. We denote the Jacobian $F'(x)$ by $J(x)$ and abbreviate $J(x_k)$ as J_k. Similarly, $F(x_k)$ is abbreviated often as F_k. When the context is clear, we may drop the subscript k while still referring to the "current" values at an iteration.

2 Algorithms

In this section, we introduce the relevant methods for solving systems of nonlinear equations. We start with methods that use a direct factorization of the Jacobian matrix, and then we discuss inexact methods that use Krylov subspace projection techniques. General references for these topics in nonlinear solvers include [12], [19], [23], and [25].

2.1 Standard Methods

In this paper, we denote by standard methods the class of methods for solving (1) that uses a linear approximation to $F(x)$ at each iterate around the current iterate $x_k \in \mathbb{R}^n$. Most notable among these methods is Newton's method, which uses the linear local model

$$M_N(x_k + d) = F_k + J_k d, \tag{2}$$

where $d \in \mathbb{R}^n$ is the step and $J_k \in \mathbb{R}^{n \times n}$ is either the current Jacobian matrix or an approximation to it. A root of this local model provides the Newton step

$$d_N = -J_k^{-1} F_k,$$

which is used to reach the next trial point. Thus, Newton's method is defined when J_k is nonsingular and consists of updating the current point with the Newton step,

$$x_+ = x_k + d_N. \tag{3}$$

If the Jacobian $J(x_k)$ is Lipschitz continuous in a neighborhood containing the root x_* and $J(x_*)$ is nonsingular, then the sequence of iterates produced by (3) converges locally and q-quadratically to x_*. That is, there exists constants $\delta > 0$ and $c \geq 0$ such that the sequence of iterates x_k produced by Newton's method obeys

$$\|x_{k+1} - x_*\| \leq c \|x_k - x_*\|^2$$

if $\|x_0 - x_*\| \leq \delta$.

When these standard approaches use direct factorizations of the Jacobian matrix, we will refer to these methods as direct methods. Due to the storage and linear algebra costs, direct methods are only practical for solving small, dense problems.

2.2 Tensor methods

Tensor methods [26] solve (1) by including a second-order term in the local model at each iteration. The local tensor model has the general form

$$M_T(x_k + d) = F_k + J_k d + \tfrac{1}{2} T_k dd, \tag{4}$$

where $T_k \in \mathbb{R}^{n\times n\times n}$ is the tensor term at x_k and is selected so that the model interpolates a small number p of function values in the recent history of iterates. By choosing the smallest T_k in the Frobenius norm, T_k has rank p and $T_k dd$ is both simple in form and inexpensive to find. Because (4) may not have a root, one solves the minimization subproblem

$$\min_{d\in\mathbb{R}^n} \|M_T(x_k + d)\|_2 \,, \tag{5}$$

and a root or minimizer of the model is the tensor step.

The additional cost of forming, storing, or solving the model is minor compared to Newton's method. Specifically, the additional cost is about n^2p multiplications (QR implementation) and $2p$ vectors of length n in storage. For our numerical experiments, we will only consider the simplest case of $p = 1$. Computational evidence in [26] suggests that additional past iterates add little benefit to the computational performance of the tensor method.

Tensor methods are considerably more efficient and robust than standard methods on singular problems and, to a lesser extent, on nonsingular problems. The second-order term provides higher order information in recent step directions, which aids in cases where the Jacobian is (nearly) singular at the solution. As the iterates approach the solution, the Jacobian lacks information in these directions, and the second-order term supplies useful information for a better step.

The analysis in [14] confirms that tensor methods have at least quadratic convergence on nonsingular problems. In addition, [14] also shows that tensor methods have local superlinear convergence for a large class of singular problems with $\text{rank}(F'(x_*)) = n - 1$ under mild conditions. In contrast, Newton's method without any acceleration techniques on such problems exhibits only q-linear convergence with constant converging to $\frac{1}{2}$.

Computational evidence in [26] on small problems shows that tensor methods hold 21–23% average improvement over standard methods on nonsingular problems and 40–43% improvement on problems with $\text{rank}(F'(x_*)) = n - 1$. Thus, tensor methods outperform standard methods on many problems, especially on singular problems.

2.3 Newton-Krylov methods

Up to this point, we have discussed direct methods for the solution of small, dense problems, such that the local model is solved with direct factorizations of the Jacobian matrix. Standard direct methods, such as Newton's method, are impractical on large-scale problems because of their high linear algebra costs and large memory requirements. Thus, most large systems often are solved successfully using a class of "inexact" Newton methods:

$$x_{k+1} = x_k + d_k, \quad \text{where} \quad F'(x_k)d_k = -F(x_k) + r_k, \quad \|r_k\| \le \eta_k \|F(x_k)\| \,, \tag{6}$$

such that the local model typically is solved only approximately at each step using a less expensive approach. These "inexact" steps then locate the next trial point. Successively better approximations to the linear model at each iteration preserve the rapid convergence behavior of Newton's method when nearing the solution. The computational savings reflected in this less expensive inner iteration is usually partially offset with more outer iterations, but the overall savings still is quite significant on large-scale problems by avoiding the direct methods that solve the local model exactly.

The most common methods for approximately solving the local Newton model are Krylov-based methods, which iteratively solve the linear system projected onto the Krylov subspace $\mathcal{K}$. A linear Krylov subspace method is a projection method that seeks an approximate solution x_m to the linear system $Ax = b$ from an m-dimensional affine subspace $x_0 + \mathcal{K}_m$. Here $\mathcal{K}_m$ is the subspace

$$\mathcal{K}_m(A, r_0) = span\{r_0, Ar_0, A^2 r_0, \ldots, A^{m-1} r_0\},$$

where $r_0 = b - Ax_0$ is the residual at an initial guess x_0. A popular Krylov subspace method is the Generalized Minimum Residual method (GMRES), which computes a solution $x_m \in x_0 + \mathcal{K}_m$ such that the residual norm over all vectors in $x_0 + \mathcal{K}_m$ is minimized. That is, at the mth step, GMRES finds x_m such that $\|b - Ax_m\|_2$ is minimized for all $x_m \in x_0 + \mathcal{K}_m$.

Newton-GMRES is one specific method in the class of Newton-Krylov methods. Here, the linear system is the Newton equation $J_k d = -F_k$, and the system is solved via GMRES according to the tolerance η in (6). Krylov subspace methods have the appeal of requiring almost no matrix storage due to their exclusive use of Jacobian-vector products, which may be calculated by a finite-difference directional derivative. For this reason and others, Newton-GMRES is a popular algorithm for solving large-scale problems, and it will be the standard large-scale Newton-based algorithm for comparisons in our numerical experiments.

Newton-Krylov methods have been considered by many authors, including Brown and Saad [5, 6], Chan and Jackson [7], and Brown and Hindmarsh [4]. Their computational results show that Newton-Krylov methods can be quite effective for many classes of problems in the context of systems of partial differential equations and ordinary differential equations.

2.4 Tensor-Krylov methods

Direct methods cannot efficiently solve large-scale problems due to large storage considerations and the expensive direct solution of the local model. To this end, the three tensor-Krylov methods described in [1, 2] combine the concepts from direct tensor methods with concepts from inexact Newton methods using Krylov-based linear solvers. The tensor-Krylov methods calculate an inexact tensor step from a specially

chosen Krylov subspace that facilitates the solution of a minimization subproblem at each step. Here we just give a very brief overview of these methods.

All large-scale tensor methods in this paper only consider a rank-one tensor model, which only interpolates the function value at the previous iterate. Thus the rank-p model of (4) reduces to

$$M_T(x_k + d) = F_k + J_k d + \tfrac{1}{2} a_k (s_k^T d)^2, \tag{7}$$

such that

$$a_k \in \mathbb{R}^n = \frac{2(F_{k-1} - F_k - J_k s_k)}{(s_k^T s_k)^2} \tag{8}$$

$$s_k \in \mathbb{R}^n = x_{k-1} - x_k. \tag{9}$$

In each of the tensor-Krylov methods, the tensor step is found by approximately solving the minimization subproblem

$$\min_{d \in \mathcal{K}_m} \left\| F_k + J_k d + \tfrac{1}{2} a_k (s_k^T d)^2 \right\|_2 \tag{10}$$

where $\mathcal{K}_m$ is a specially chosen Krylov subspace that facilitates the solution of the quadratic model. The three methods described in [1, 2] differ in their choice of $\mathcal{K}_m$, and they are identified by this characteristic difference. TK2 and TK2+ use a Krylov-based local solver that starts with an initial block of two vectors (TK2+ also augments the Krylov subspace in a special way). Similarly, TK3 uses an initial block of size three. The three methods share the ability to calculate an approximate tensor step that satisfies the tensor model to within a specified tolerance. Their cost per nonlinear iteration exceeds that of Newton-GMRES by at most $10n + 4mn + 6m^2$ multiplications (cf. GMRES costs $\mathcal{O}(nm^2)$ multiplications). The methods can be readily combined with either left or right preconditioning. More details of these Krylov-subspace methods for solving the tensor model may be found in [1] and [2].

2.5 Tensor-GMRES method

Another large-scale tensor method is that of Feng and Pulliam [15], which uses Krylov subspace projection techniques for solving the local tensor model. In particular, it uses GMRES to first find the approximate Newton step $d_N = d_0 + V_m y_m$. The columns of V_m form an orthonormal basis for the Krylov subspace $\mathcal{K}_m$ generated by the corresponding Arnoldi process, and the Hessenberg matrix H_m is also generated from the Arnoldi process. Given these key matrices, their tensor-GMRES algorithm proceeds to solve a projected version of the tensor model (7) along a subspace that spans the Newton step direction and the Krylov subspace from the Newton step solution. (That is, the approximate tensor step is in the span of the Krylov subspace $\mathcal{K}_m^N$ and d_0, or equivalently the span of the matrix $[V_m, d_0]$). Thus, their algorithm solves the least-squares problem

$$\min_{d \in \{d_0\} \cup \mathcal{K}_m^N} \left\| F_k + J_k d + \tfrac{1}{2} P a (s^T d)^2 \right\|, \tag{11}$$

where P is the projection matrix

$$P = Y(Y^T Y)^{-1} Y^T, \quad \text{where } Y = J_k[V_m, d_0]. \tag{12}$$

The analysis in [15] shows that the same superlinear convergence properties for the unprojected tensor model considered in [14] also hold for the projected tensor model (11). The complete tensor-GMRES algorithm for solving (11) at the kth nonlinear iteration is listed in [15].

Despite the algorithm's difficult algebra, the design actually is rather straightforward. The algorithm may be viewed as an extension of Newton-GMRES, where the inexact Newton step is calculated via GMRES in the standard way. The tensor step is calculated subsequently using the Krylov subspace information generated for the Newton step. In this way, the method also is consistent with preconditioning techniques and a matrix-free implementation, which makes it appealing for general use.

The extra work and storage beyond GMRES for computing the tensor step is quite small. The extra work is at most $4mn + 5n + 2m^2 + \mathcal{O}(m)$ multiplications plus a single Jacobian-vector product for evaluating the tensor term a_k. The extra storage amounts to two extra n-vectors for a and s plus a few smaller working vectors of length m.

The results in [15] show the superlinear convergence behavior of tensor-GMRES on three singular and nearly singular problems, where the Newton-GMRES method exhibits linear convergence due to a lack of sufficient first-order information. The margin of improvement (in terms of reduction of nonlinear iterations over Newton's method) varied from 20% to 55% on the simpler problems and 32% to 60% improvement on the more difficult Euler problem. Running times and the total number of Jacobian-vector products for each method were not reported in [15], but from our own experience with the algorithm, we assume that these performance metrics are correlated with the number of nonlinear iterations.

3 Numerical experiments on ill-conditioned problems

3.1 Small problems solved with direct methods

This section investigates the performance of direct tensor methods as well as Newton's method on a set of small problems that include a parameter for adjusting the ill-conditioning of the Jacobian matrix at the root. The results show that Newton's method requires increasingly more iterations as the problem ill-conditioning grows, whereas the direct tensor methods are only mildly affected.

Following the approach in [26], we created ill-conditioned problems by modifying nonsingular test problems to be of the form

$$\hat{F}(x,\lambda) = F(x) - \lambda F'(x_*)A(A^TA)^{-1}A^T(x - x_*), \tag{13}$$

where $F(x)$ is the standard nonsingular test function, x_* is its root, $A \in \mathbb{R}^{n\times j}$ is an arbitrary matrix that has full column rank with $1 \le j \le n$, and $\lambda \in [0,1]$ is a parameter for ill-conditioning. We denote by $\hat{J}(x,\lambda)$ as the Jacobian matrix of $\hat{F}(x,\lambda)$ with respect to x:

$$\hat{J}(x,\lambda) \equiv \frac{\partial}{\partial x}\hat{F}(x,\lambda).$$

These new test problems are similar to problems in continuation or homotopy methods, except that $\hat{F}(x,\lambda)$ has the same root as $F(x)$ (i.e., x_*) for all values of λ. While the problem becomes harder to solve as λ approaches 1, the idea is not to follow the path over a sequence of values of λ. Rather, we solve the modified problem for each value of λ from the same starting point across all tests and record the number of iterations required to reach the solution.

One special quality of this modified problem is that if $\lambda = 1$ and $F'(x_*)$ has full rank, then the rank of $\hat{J}(x_*,\lambda)$ equals $\text{rank}(F'(x_*)) - \text{rank}(A) = n - \text{rank}(A)$. Thus, as λ approaches 1, the rank deficiency of $\hat{J}(x_*,\lambda)$ approaches the rank of A. Stated another way, a set of the smallest singular values of $\hat{J}(x_*,\lambda)$ equal to the rank of A will approach 0 as λ approaches 1.

Different sets of singular and ill-conditioned problems may be created using the matrix A in (13); adding more independent columns to A serves to decrease the rank of $\hat{J}(x_*,1)$. We routinely used the matrices

$$A \in \mathbb{R}^{n\times 1}, \quad A^T = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \end{pmatrix},$$

$$A \in \mathbb{R}^{n\times 2}, \quad A^T = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & -1 & 1 & \dots & \pm 1 \end{pmatrix},$$

and

$$A \in \mathbb{R}^{n\times 3}, \quad A^T = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & -1 & 1 & \dots & \pm 1 \\ 1 & 1 & -1 & \dots & \pm 1, \end{pmatrix}$$

because they provide "balanced" problems by acting equally on the whole Jacobian $F'(x_*)$. Letting A equal the unit vector e_1, for example, would only operate on the $(1,1)$ element of $F'(x)$.

We chose six problems for testing from the standard small dimensional test set of Moré, Garbow, and Hillstrom [22] and modified them according to (13). Table 1 lists the problems used in our numerical tests, along with the corresponding dimensions. For all problems, the starting vector x_0 was the standard starting point for each problem published in [22]. Except for the cases mentioned below, these starting

Table 1. Ill-conditioned test problems.

Problem	Size
Broyden Tridiagonal	50
Broyden Banded	50
Discrete Integral Equation	50
Discrete Boundary Value Function	50
Rosenbrock's Function	2
Brown Almost Linear	10

points did not require a global strategy to reach the solution (i.e., the full step was accepted at all iterations).

We chose the parameter λ in (13) to asymptotically approach $\lambda = 1$ by using the values $\lambda = 1 - 10^{-j}, j = 0, 1, 2, \ldots$. Thus, for $j = 0$ the problem is the original, unmodified problem, and each subsequent value of j makes the problem more ill-conditioned. At some point, round-off errors in the evaluation of $\hat{F}(x, \lambda)$ as well as the numerical precision of the root x_* make the numerical solution indistinguishable from the case of $\lambda = 1$. All subsequent values of λ would produce the same result, so we collected only the results up to these points and indicated the case of $\lambda = 1$ at the rightmost extent of our plots.

We used the following two stopping conditions in all these tests: $\|F(x_k)\|_\infty \leq 10^{-12}$ or $\|x_k - x_{k-1}\|_\infty \leq 10^{-12}$, which are the same conditions Eisenstat and Walker used when analyzing inexact Newton methods in [13]. In many practical applications, less stringent convergence tolerances are commonly used, but these tight tolerances were used in this experiment (and later experiments) to differentiate results at higher condition numbers and to allow asymptotic convergence behavior to become evident. The numerical differences are still present but less striking at looser stopping tolerances.

As a prelude to these results and to help explain what is happening in these experiments, we first provide a graphical description. Figure 1 shows the typical iteration profiles for different values of λ that we observed. This figure graphs the function value at each iteration for the tensor and Newton methods on a typical problem approaching rank $n - 1$ with various values of $\lambda = 0, 0.9, 0.99, 0.999, \ldots, 1$. All of the tensor method profiles are bunched together on the left, requiring few iterations even as λ nears 1, whereas the profiles of Newton's method are spread out and require increasingly more iterations for convergence as the problem becomes more ill-conditioned. This plot of iteration profiles is typical for all problems. Thus, while all of the tensor runs display superlinear convergence throughout the iterations, it is evident that Newton's method converges linearly for a number of iterations before accelerating to quadratic convergence. With increasing ill-conditioning, the region of quadratic convergence for Newton's method shrinks in size, acting as if the problem were singular outside of this region.

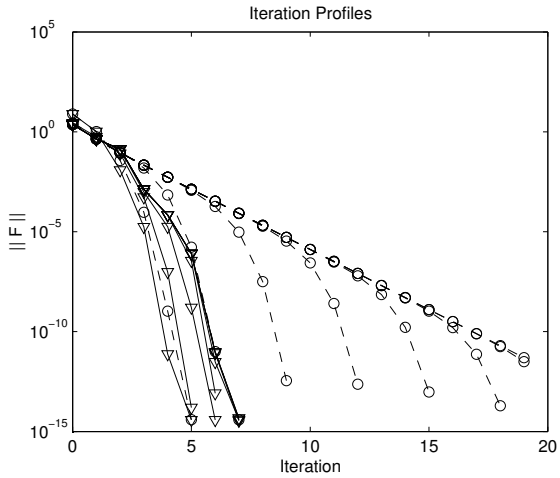

Figure 1. Superimposed iteration profiles of the tensor method ($\triangledown$, solid lines) and Newton's method ($\circ$, dashed lines) when solving the Modified Broyden Tridiagonal function as it approaches rank $n-1$. For both methods and their corresponding set of iteration profiles, problem difficulty (λ) increases from left to right, making the problem more ill-conditioned and requiring more iterations to solve.

Figure 2 shows the relative performance of Newton's method versus a rank-one tensor method on the six problems, in the case where each Jacobian at the solution approaches rank $n-1$. The condition number $\hat{J}(x_*, \lambda)$ is plotted on the abscissa, while the number of iterations for both methods is plotted on the ordinate axis. One tensor method iteration is just slightly more expensive than a Newton iteration. There were no linesearches in all cases except on Rosenbrock's function at $\lambda = 0$, which explains the unusual spike on the left in that plot. Also, both methods arrived at different solutions in the Brown Almost Linear problem at $\lambda = 0$, so these points were not included in that plot. For all other points, both methods arrived at the same x_*.

The six plots of Figure 2 show that as the ill-conditioning of a particular problem grows, Newton's method requires more iterations whereas the tensor method is very mildly affected. This is a key result. Previously, tensor methods were suggested to behave reasonably well on ill-conditioned problems, but there was no numerical testing of this conjecture.

However, because the condition number is just a single value representing the ratio of the largest to the smallest magnitude singular values, it does not capture the singular value spectrum of the Jacobian. Say, for instance, that there are multiple singular values approaching zero, creating a very ill-conditioned matrix. One might speculate that a tensor method using a rank-one tensor term would be less advanta-

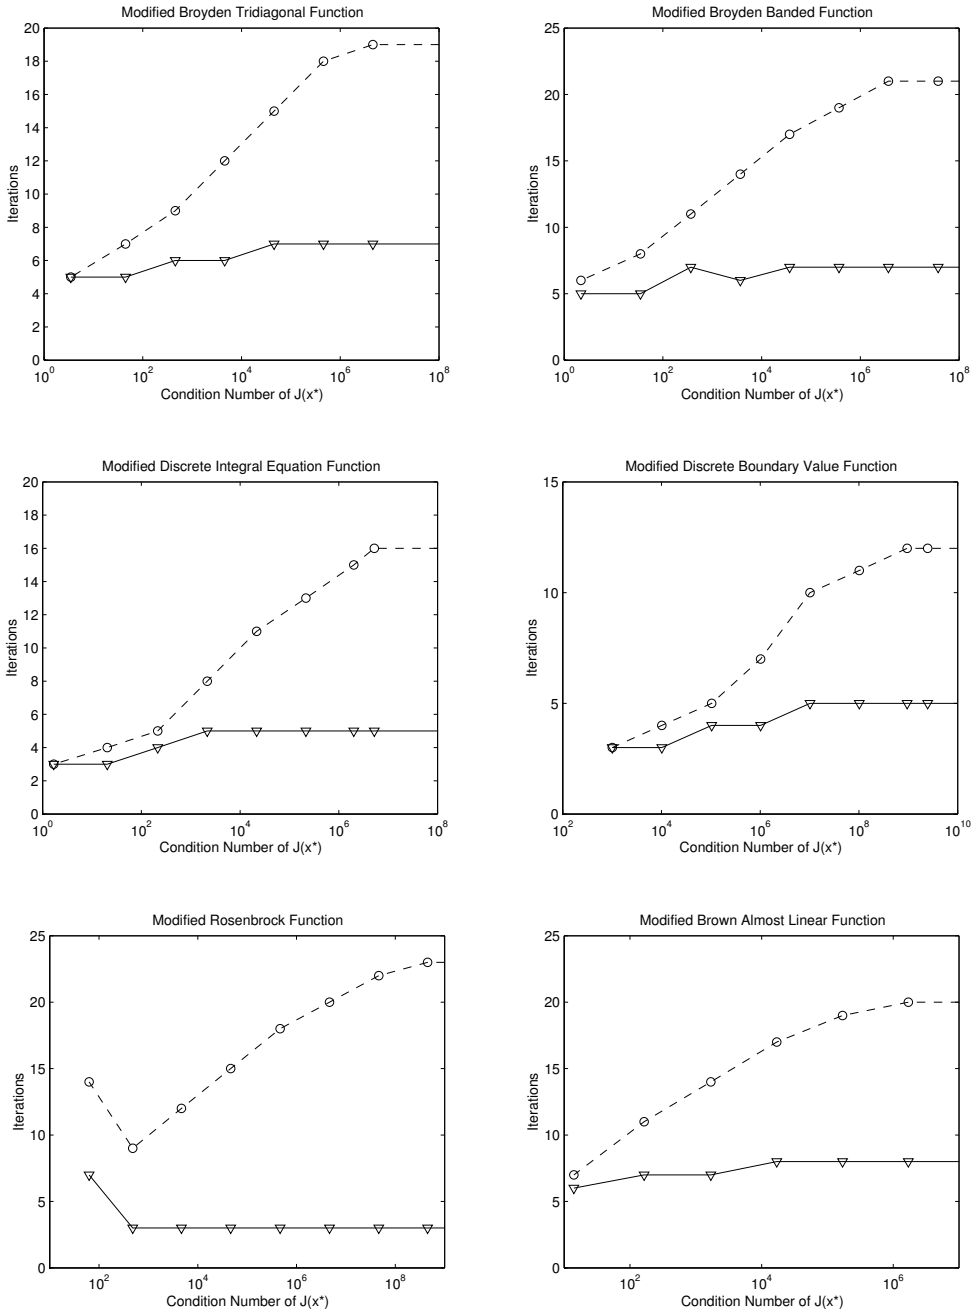

Figure 2. Comparison of a rank-one tensor method ($\triangledown$, solid line) with Newton's method ($\circ$, dashed line) on variably ill-conditioned test problems that approach rank $n-1$.

geous in this case because it could "handle" only one direction, perhaps the direction associated with the smallest singular value, leaving other (near) singular directions to impede convergence. Figures 3 and 4 consider exactly these cases where the Jacobians

approach rank $n-2$ and $n-3$, respectively, while still using a rank-one tensor model. Two problems were not included in these figures for different reasons. The Modified Rosenbrock Function had excessive linesearches that obscured the results in the rank $n-2$ problems, and the rank $n-3$ problem is not possible with the Rosenbrock Function due to its small dimension of $n = 2$. For the Brown Almost Linear problem, each method arrived at a solution that was different from the root x_* used in (13). For these reasons, these two problems were not included in Figures 3 and 4.

It is evident that the results for higher rank deficiencies are not as striking as in Figure 2, but the general trend still remains—tensor methods perform better than Newton's method on ill-conditioned problems, even when the Jacobian approaches a rank deficiency greater than the rank of the tensor term. A few peculiarities exist in the results that warrant explanation. We attribute the "humps" in both curves to the random paths each method takes to arrive at the solution, and we believe this shape is coincidental. The "spikes" in both methods for the two Discrete problems in Figure 4 are due to linesearches occurring at a specific value of λ. The linesearches force a few more outer iterations to eventually solve the problems.

3.2 Moderate-size problems solved with inexact methods

This subsection investigates the performance of Newton-GMRES, tensor-GMRES, and the tensor-Krylov methods on three ill-conditioned problems of moderate size and complexity: the Bratu problem, a modified discrete boundary value function, and a modified Broyden tridiagonal problem. The tests on the ill-conditioned problems have a similar design as the study performed in section 3.1, and we investigate to what extent the benefits of using direct tensor methods for solving small ill-conditioned/singular problems that were shown section 3.1 extend to large-scale, inexact versions of tensor methods.

We implemented all of the methods in MATLAB using double precision arithmetic. For objective comparisons, we used the same level of basic linear algebra routines in MATLAB in all of our tests. Thus, the results in this section do not reflect the most efficient implementations that are available, but the statistics that we collected (nonlinear iterations, function evaluations) are invariant to optimal implementations.

For the numerical tests, we used the same stopping conditions and parameters as in section 3.1. For the inner method, we solved the local model to a constant relative tolerance of 10^{-4}, which requires uniformly close approximations to Newton and/or tensor steps at each iterate and results in fast local q-linear convergence. Using a constant relative tolerance may not be ideal in practice (see [13] for alternatives), but it provides for more unbiased comparisons among the methods.

In practice, the Krylov subspace dimension should be kept small to reduce arithmetic and storage costs, which may be accomplished via restarted Krylov methods

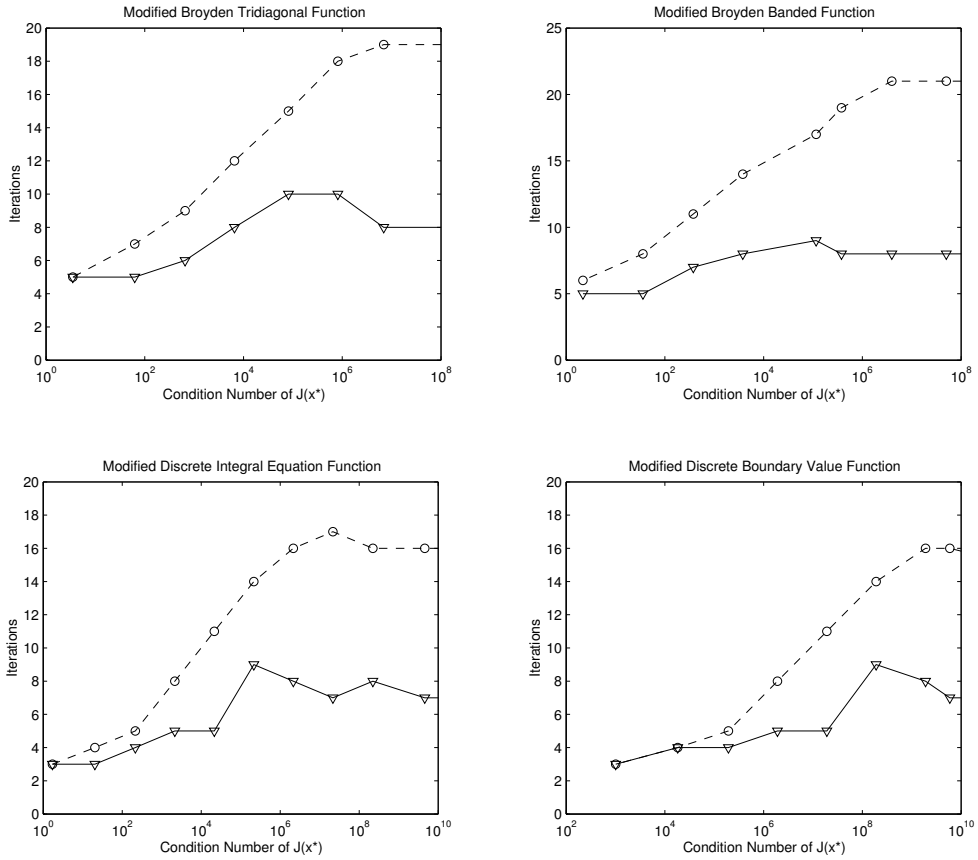

Figure 3. Comparison of a rank-one tensor method ($\triangledown$, solid line) with Newton's method ($\circ$, dashed line) on variably ill-conditioned test problems that approach rank $n-2$.

and preconditioning. However, to eliminate any ill effects of small Krylov subspaces preventing convergence in the local model and affecting the outer iterations, the maximum Krylov subspace was set to the problem dimension, $m_{max} = n$, and the solver was not restarted. We used preconditioners that are appropriate for the problem, and they will be discussed with each problem.

In the next three subsections, we present numerical results on the ill-conditioned problems. The results of the tensor-Krylov method TK2 are virtually identical to TK2+, so we do not include them here.

3.2.1 Bratu problem

The Bratu problem is a simplified model for nonlinear diffusion phenomena occurring, for example, in semiconductors and combustion, where the source term is related to

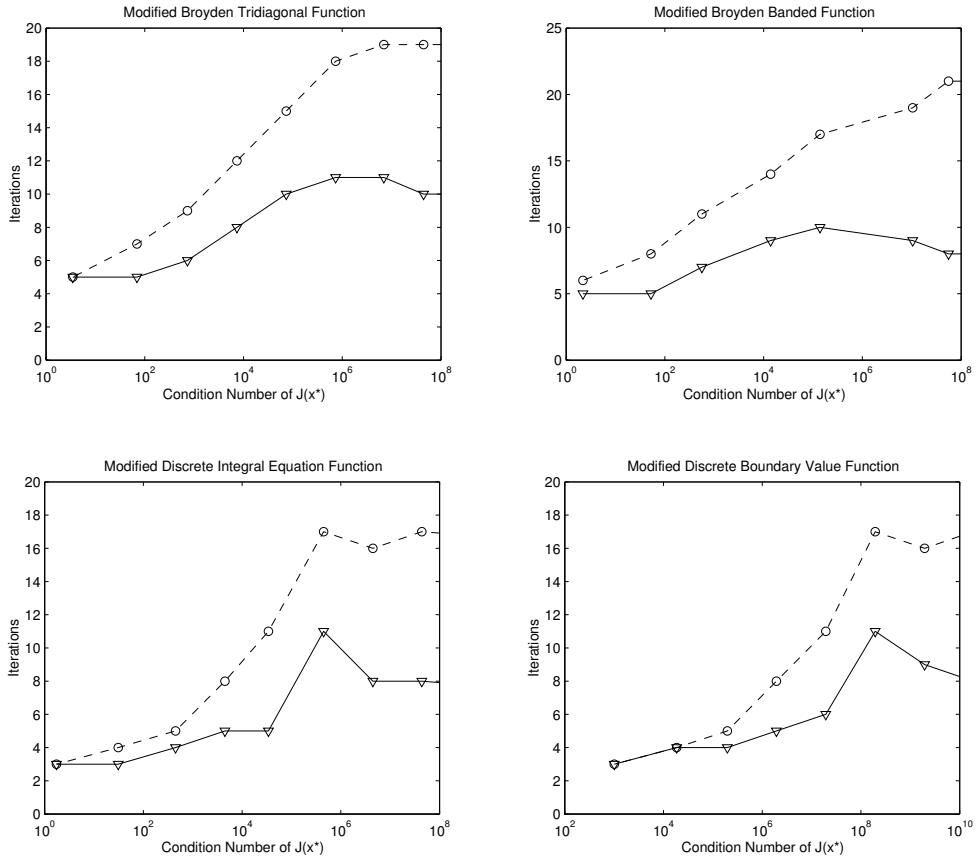

Figure 4. Comparison of a rank-one tensor method ($\triangledown$, solid line) with Newton's method ($\circ$, dashed line) on variably ill-conditioned test problems that approach rank $n-3$.

the Arrhenius law for modeling exothermic reactions. The following version is taken from the set of nonlinear model problems collected by Moré [21]. The problem is the nonlinear partial differential equation in u

$$-\nabla^2 u = \lambda e^u \text{ in } \Omega, \quad u = 0 \text{ on } \partial\Omega, \tag{14}$$

where $\nabla^2 = \sum_{i=1}^n \partial^2/\partial x_i^2$ is the Laplace operator, $\lambda \in \mathbb{R}$ is a parameter, Ω is the bounded domain $(0,1) \times (0,1)$, and $\partial\Omega$ the boundary of Ω.

Problem (14) has a unique solution for $\lambda \leq 0$, but for $\lambda > 0$, there may be zero, one, or two solutions (cf., [16]). The critical value $\lambda^* = 6.80812$ is a limit point such that for $0 < \lambda < \lambda^*$, problem (14) has two solutions; and for $\lambda > \lambda^*$, it has no solutions. Also, the problem at the limit point is singular with a rank of $n-1$, and as λ approaches the limit point, the discretized problem becomes harder to solve. To investigate the effects of ill-conditioning on the inexact algorithms, we increased λ over the range $\lambda \in [5, 6.806652]$, which increased the condition number of $J(x_*)$ from

about 10^3 to 10^6.

When testing the Bratu problem, the initial approximate solution was zero on a uniform grid of size 31×31. The Laplace operator was discretized using centered differences (5-point stencil), and our preconditioner was also the centered differences discretization of the Laplace operator. We computed Jacobian-vector products using first-order forward differences. Hence, the number of function evaluations is the sum of the total number of Arnoldi iterations, the number of linesearch backtracks (if any), and the total number of nonlinear iterations. Thus, the number of function evaluations provides a relative measure of overall work for each algorithm.

Figure 5 presents the results of these tests, comparing the number of function evaluations computed for each method using the three choices of preconditioning (none, left, and right preconditioning). If one considers the condition number of the preconditioned Jacobian at the root (i.e., $M^{-1}J(x_*)$ or $J(x_*)M^{-1}$), then the condition numbers in the bottom two plots span the range of 3×10^0 to about 10^4 instead of 10^3 to 10^6 for the Jacobian itself. The results in the figure are similar to those in figures in section 3, where Newton-GMRES requires increasingly more function evaluations as the problem becomes more ill-conditioned and difficult. The tensor-Krylov methods required from 4 to 7 outer iterations, whereas Newton-GMRES required from 4 to 13 outer iterations. That is, Newton-GMRES required almost double the number of tensor-Krylov iterations on the most ill-conditioned problem. The three plots in Figure 5 uphold the prior conclusion that increased ill-conditioning only mildly increases the number of outer iterations for tensor methods.

In all cases, when the problem is not overly difficult (e.g., the experiment at condition number of 10^3 on the left-hand side of the figure), Newton-GMRES is more efficient in terms of nonlinear iterations than the tensor methods. Here, the number of iterations for all of the methods is equal, but GMRES is more efficient at solving the local Newton model, which results in fewer total function evaluations. When the problem is more difficult, however, the tensor methods are superior to Newton-GMRES by a factor of up to 1.9. Approaching the limit point even closer would have given even more of an advantage to the tensor-Krylov methods.

When comparing only the tensor methods, the tensor-GMRES method of Feng and Pulliam is more efficient than the TK3 and TK2+ methods. It appears that the step produced from the projected tensor model of the Feng-Pulliam method is nearly the same as the steps calculated from the more precise local models of the tensor-Krylov methods, and it is less expensive to compute. This behavior may be understood more clearly by investigating the iteration history of all of the methods, which we describe next.

To investigate the step quality of each tensor-Krylov method, we tested the Bratu problem with $\lambda = 6.806652$, which has a condition number of roughly 1.5×10^6. Using right preconditioning, the results in the first plot of Figure 6 show the faster outer (nonlinear) iterations of the three tensor methods, which display very similar

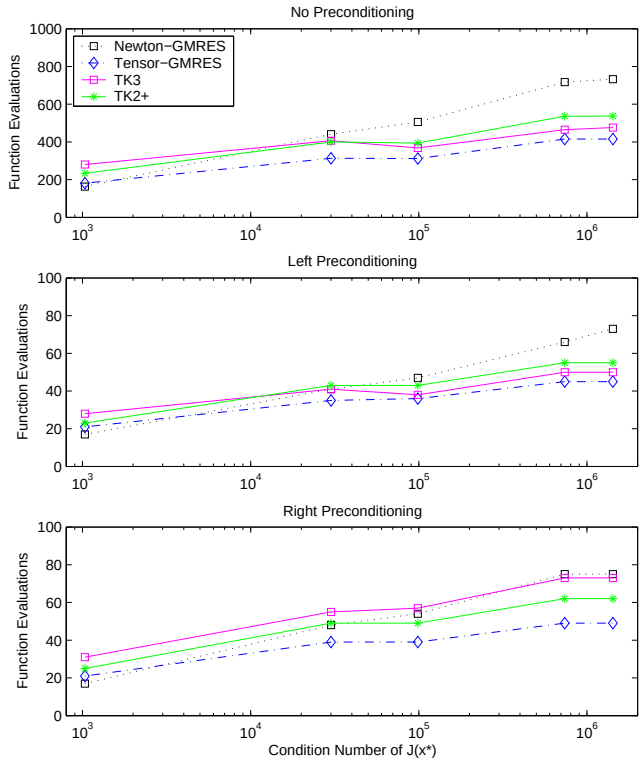

Figure 5. Effects of ill-conditioning on the inexact algorithms as seen in the Bratu problem.

performance. All methods start by exhibiting linear convergence until the respective method can overcome the near singularity and accelerate convergence. The tensor methods accelerate convergence sooner than Newton-GMRES (iteration 2 versus iteration 5), and this is typical behavior for ill-conditioned problems—Newton-type methods branch into superlinear convergence later and later as the ill-conditioning grows. Because the forcing term for the inner iterative method is constant (instead of decreasing each iteration), all methods exhibit asymptotic linear convergence, as evidenced by their straight trajectories near the solution.

When we consider function evaluations in the bottom plot, the tensor-Krylov methods separate from the Feng-Pulliam method. The block size of the method is a clear indicator of the relative efficiency of the method. Specifically, tensor-GMRES, which uses the scalar (block size one) implementation of GMRES, is more efficient than TK2+ (block size two) and TK3 (block size three). The Bratu problem is unique among our tests in that the steps computed from a projected local tensor model are of roughly the same quality as the steps from TK2+ and TK3. Therefore, the number of outer iterations are the same, but because tensor-GMRES is more efficient in solving the local model, tensor-GMRES has the advantage.

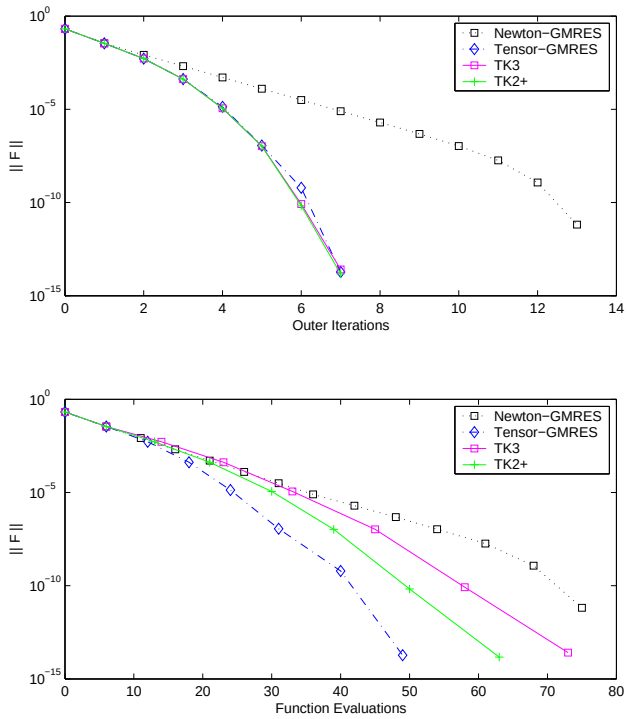

Figure 6. Example iteration history on the Bratu problem at $\lambda = 6.806652$.

3.2.2 Modified discrete boundary value problem

The discrete boundary value problem is a simple test problem from [22]. The standard discrete boundary value problem is

$$f_i(x) = 2x_i - x_{i-1} - x_{i+1} + \tfrac{1}{2}h^2(x_i + t_i + 1)^3, \quad 1 \le i \le n, \tag{15}$$

where $h = \frac{1}{n+1}$, $t_i = ih$, and $x_0 = x_{n+1} = 0$. The initial approximate solution is zero on a problem size of $n = 100$ equations. For our tests, we have modified (15) in accordance with equation (13) of section 3.1.

When testing this problem, we used a preconditioner that corresponds to the Jacobian of (15) but without the term t_i. We computed Jacobian-vector products using an analytic evaluation of the Jacobian, and we tallied the number of "function evaluation equivalents," which is the sum of function evaluations and Jacobian-vector products. If these products were approximated by first-order forward differences, as is generally the case with complex problems, then "function evaluation equivalents" would be equal to the number of function evaluations. This number also provides a relative measure of overall work for each algorithm.

Figure 7 presents the results of these tests, comparing the number of function evaluation equivalents computed for each method using the three choices of preconditioning (none, left, and right preconditioning) for the various values of $\lambda = 1 - 10^{-j}, j = 0, 1, 2, \ldots$ in equation (13). If one considers the condition number of the preconditioned Jacobian at the root (i.e., $M^{-1}J(x_*)$ or $J(x_*)M^{-1}$), then the condition numbers in the bottom two plots span the range of 1×10^0 to about 2×10^6 (left preconditioning) or 2×10^8 (right preconditioning) instead of 4×10^3 to 2×10^9 for the Jacobian, not counting the singular case.

The results for the modified discrete boundary value problem are similar to those of the Bratu problem in Figure 5: Newton-GMRES requires increasingly more function evaluations as the problem becomes more ill-conditioned, whereas the number of function evaluations required by the tensor methods increases to a lesser extent. It is interesting to note that the tensor methods are virtually identical in performance without preconditioning and resemble the results in section 3.1. However, once preconditioning is used, the tensor methods are affected by the ill-conditioning of the problem, yet still to a lesser extent than Newton-GMRES.

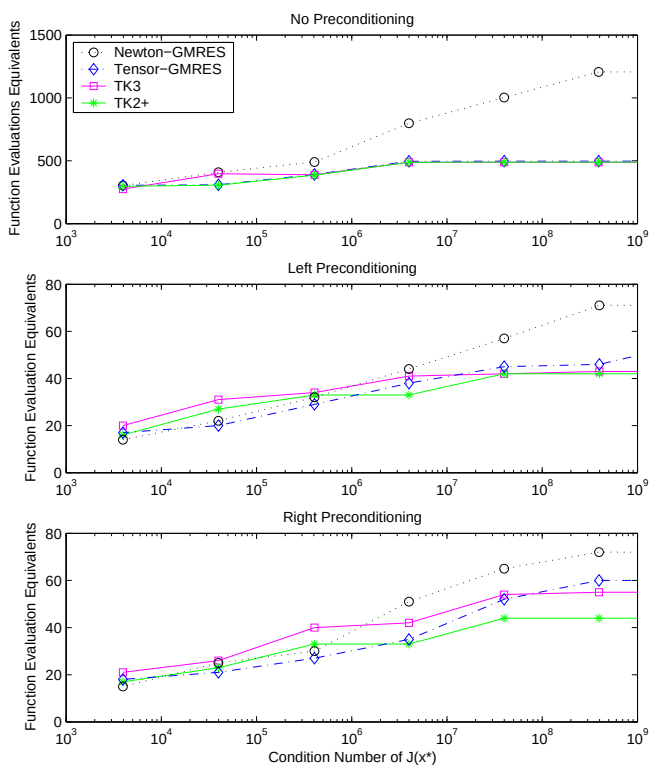

Figure 7. Effects of ill-conditioning on the inexact algorithms as seen in the discrete boundary value problem.

When comparing only the tensor methods while using left and right precondition-

ing, the tensor-GMRES method of Feng and Pulliam is more efficient than the TK3 and TK2+ methods on the easier problems but less efficient on the hardest problems, which is different from the experience with the Bratu problem above. We explain this relative difference by noting that Tensor-GMRES required more nonlinear iterations as the problem grew more ill-conditioned. Because the tensor-Krylov methods generally require more Arnoldi (inner) iterations to solve the local tensor model due to the less efficient block-Arnoldi process, any computational savings must come from fewer nonlinear (outer) iterations.

To better illustrate this behavior, Figure 8 presents the iteration history of all methods for one test case in Figure 7. We present the results of the modified problem using right preconditioning with $\lambda = 1$, which corresponds to the right-most set of points in the bottom plot of Figure 7. Here, the top plot of Figure 8 shows that Newton-GMRES has linear convergence, and after one good step on the second iteration, Tensor-GMRES appears to have linear convergence that is slightly faster than Newton-GMRES. The TK2+ and TK3 methods appear to have superlinear convergence with nearly the same steps. This profile indicates that the tensor step found from a projected tensor model is not as good for this problem as the tensor-Krylov steps which use the full tensor model. That is, the projection of the tensor model loses some information that is important for achieving superlinear convergence.

When we consider function evaluations in the bottom plot of Figure 8, two features are evident. First,the slope of the Newton-GMRES and tensor-GMRES lines are nearly identical, which indicates that the cost of solving the projected tensor model in terms of function evaluations is roughly the same as solving the Newton model in Newton-GMRES. Second, the TK2+ and TK3 trajectories separate, which indicates that the block-2+ local solver in TK2+ is more efficient than the block-3 solver in TK3.

3.2.3 Modified Broyden tridiagonal problem

The Broyden tridiagonal problem is another test problem from [22]. The function is defined as

$$f_i(x) = (3 - 2x_i)x_i - x_{i-1} - 2x_{i+1} + 1, \quad 1 \leq i \leq n, \tag{16}$$

where $x_0 = x_{n+1} = 0$. Here again, we have modified the problem to be ill-conditioned, but in this case, we modify the last function to be

$$f_n(x) = (1 - \lambda)\left[(3 - 2x_n)x_n - x_{n-1} + 1\right] + \lambda\left[(3 - 2x_n)x_n - x_{n-1} + 1\right]^2.$$

This modification is similar to the problem studied by Feng and Pulliam in [15] and makes the problem more ill-conditioned as λ asymptotically approaches one and singular at $\lambda = 1$ with a rank $n - 1$. We used the values $\lambda = 1 - 10^{-j}, j = 0, 1, 2, \ldots,$ until the results were indistinguishable from $\lambda = 1$ due to round-off error, at which point we just use $\lambda = 1$. For our tests, we set $n = 100$ and used $(-1, \ldots, -1)^T$ as the

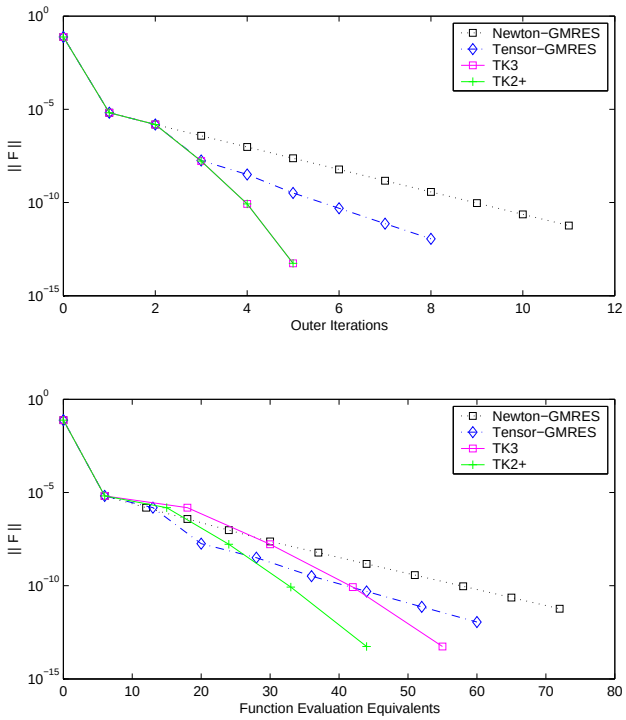

Figure 8. Example iteration history on the discrete boundary value problem at $\lambda = 1$.

starting vector. We used the Jacobian of (16) as our preconditioner for all values of λ.

Figure 9 presents the results of the ill-conditioning tests over the span of problems $\lambda \in [0, 1]$, which affects the condition number of $J(x_*)$. The three plots represent the three choices of preconditioning (none, left, and right preconditioning), and each compares the number of function evaluation equivalents computed for the different methods. If one considers the condition number of the preconditioned Jacobian at the root (i.e., $M^{-1}J(x_*)$ or $J(x_*)M^{-1}$), then the condition numbers in the bottom two plots (not counting the singular case) span the range of 1×10^0 to 2×10^6 instead of 3×10^0 to 3×10^6 for the Jacobian.

Here again the results show that Newton-GMRES requires increasingly more function evaluations as the problem becomes more ill-conditioned. The tensor methods are affected somewhat at the low condition numbers, but there is a plateau where the number of function evaluations required by the tensor methods are no longer affected by the ill-conditioning of the problem. Near $\lambda = 1$, Newton-GMRES requires almost twice the number of function evaluations as the tensor methods.

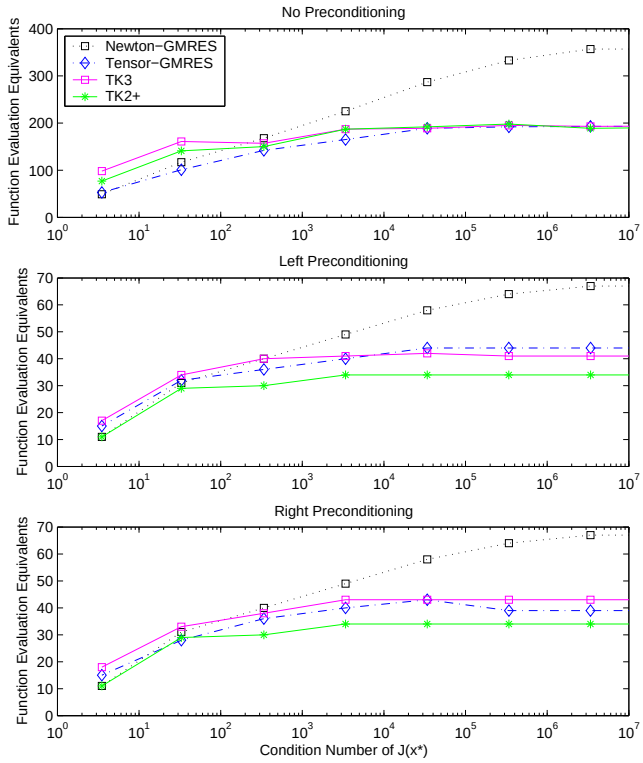

Figure 9. Effects of ill-conditioning on the inexact algorithms as seen in the modified Broyden Tridiagonal problem.

When comparing only the tensor methods while using left and right preconditioning, the tensor-GMRES method of Feng and Pulliam is roughly comparable to TK3. The TK2+ method is more efficient on all problems except one, where tensor-GMRES is the best. In comparison with Newton-GMRES on the easier problems with left or right preconditioning, TK2+ requires the same number or fewer function evaluation equivalents.

Figure 10 presents the iteration history of all methods for the right-most test case ($\lambda = 1$) in the right preconditioning plot in Figure 9. The top plot of Figure 10 shows that Newton-GMRES has linear convergence on this singular problem, and the tensor methods have superlinear convergence. Once again, the two tensor-Krylov methods have the same quality of steps, but the tensor-GMRES method no longer shares the same trajectory as TK2+ and TK3, which indicates that the projected tensor model loses some critical directional information.

The bottom plot of Figure 10 considers the function evaluation equivalents on the x-axis. It shows that tensor-GMRES is generally more efficient than TK3 at solving the local model. Thus, while TK3 is more efficient in total nonlinear iterations, tensor-

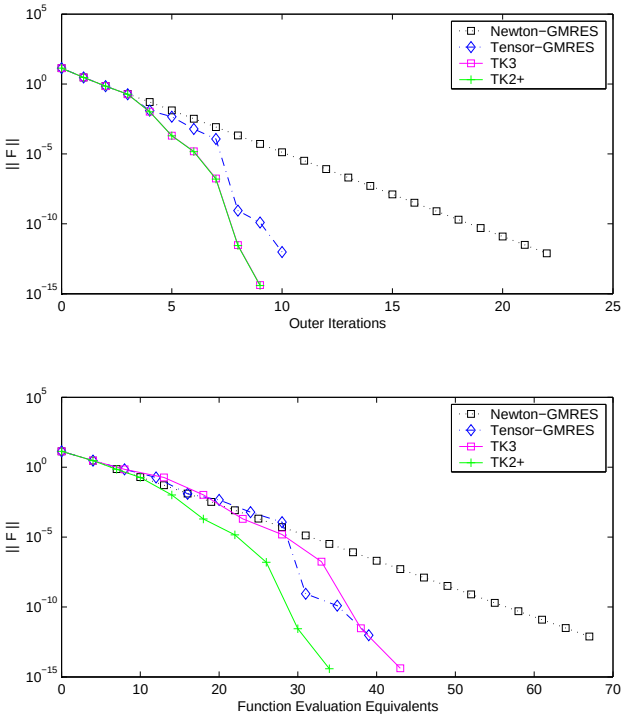

Figure 10. Example iteration history on the modified Broyden tridiagonal problem at $\lambda = 1$.

GMRES is more efficient at solving the local model, which accounts for the difference. Another feature evident in the bottom plot is that the TK2+ and TK3 trajectories separate. As with the discrete boundary value problem above, this indicates that the block-2+ local solver in TK2+ is more efficient than the block-3 solver in TK3.

4 Summary and conclusions

This paper has investigated the performance of small- and large-scale tensor methods on problems over a range of conditionings, from well-conditioned to ill-conditioned to singular. Our results showed that tensor methods outperform Newton's method as the problems become more ill-conditioned. Prior to this investigation, studies on direct tensor methods only focused on singular problems or on general problems that are well-conditioned.

Specifically, our results show that, eventual quadratic convergence notwithstanding, the performance of Newton's method will degrade as the ill-conditioning grows,

whereas tensor methods appear to be relatively unaffected or only mildly affected (in the case of larger rank deficiencies). Newton-based methods do not handle these singular problems well because they converge linearly to the solution and, in some cases, with poor accuracy.

For the large-scale methods, despite the use of an iterative inner method with an approximate solve, the tensor-Krylov methods appear to retain superlinear convergence properties on ill-conditioned problems. On the other hand, Newton-GMRES is affected by the ill-conditioning and branches into superlinear convergence later and later as the problems become more ill-conditioned. Thus, tensor methods are especially useful for large-scale problems that are highly ill-conditioned or singular, where Newton-based algorithms exhibit very slow convergence.

There are many important and practical problems that have ill-conditioned or singular Jacobian matrices at the solution, such as large PDE problems that exhibit "turning points" and/or shocks. The conclusions of this research indicate that concepts from tensor methods may benefit algorithms for bifurcation tracking and stability analysis. We intend to investigate these applications in future research.

References

[1] Brett W. Bader. Tensor-Krylov methods for solving large-scale systems of nonlinear equations. *SIAM J. Numer. Anal.* submitted.

[2] Brett W. Bader. *Tensor-Krylov methods for solving large-scale systems of nonlinear equations*. PhD thesis, University of Colorado, Boulder, Department of Computer Science, 2003.

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

[4] Peter N. Brown and Alan C. Hindmarsh. Reduced storage methods in stiff ODE systems. *J. Appl. Math. Comput.*, 31:40–91, 1989.

[5] Peter N. Brown and Yousef Saad. Hybrid Krylov methods for nonlinear systems of equations. *SIAM J. Sci. Statist. Comput.*, 11:450–481, 1990.

[6] Peter N. Brown and Yousef Saad. Convergence theory of nonlinear Newton-Krylov algorithms. *SIAM J. Optim.*, 4:297–330, 1994.

[7] T. F. Chan and K. R. Jackson. The use of iterative linear equation solvers in codes for large systems of stiff IVPs for ODEs. *SIAM J. Sci. Statist. Comput.*, 7:378–417, 1986.

[8] D. W. Decker, H. B. Keller, and C. T. Kelley. Convergence rate for Newton's method at singular points. *SIAM J. Numer. Anal.*, 20:296–314, 1983.

[9] D. W. Decker and C. T. Kelley. Newton's method at singular points I. *SIAM J. Numer. Anal.*, 17:66–70, 1980.

[10] D. W. Decker and C. T. Kelley. Newton's method at singular points II. *SIAM J. Numer. Anal.*, 17:465–471, 1980.

[11] D. W. Decker and C. T. Kelley. Convergence acceleration for Newton's method at singular points. *SIAM J. Numer. Anal.*, 19:219–229, 1982.

[12] J. E. Dennis, Jr. and Robert B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, Englewood Cliffs, NJ, 1983.

[13] Stanley C. Eisenstat and Homer F. Walker. Choosing the forcing terms in an inexact Newton method. *SIAM J. Sci. Comput.*, 17:16–32, 1996.

[14] Dan Feng, Paul D. Frank, and Robert B. Schnabel. Local convergence analysis of tensor methods for nonlinear equations. *Math. Programming*, 62:427–459, 1993.

[15] Dan Feng and Thomas H. Pulliam. Tensor-GMRES method for large systems of nonlinear equations. *SIAM J. Optim.*, 7:757–779, 1997.

[16] R. Glowinski, H. B. Keller, and L. Reinhart. Continuation-conjugate gradient methods for the least squares solution of nonlinear boundary value problems. *SIAM J. Sci. Statist. Comput.*, 6:793–832, 1985.

[17] A. Griewank. On solving nonlinear equations with simple singularities or nearly singular solutions. *SIAM Review*, 27:537–563, 1985.

[18] A. Griewank and M. R. Osborne. Analysis of Newton's method at irregular singularities. *SIAM J. Numer. Anal.*, 20:747–773, 1983.

[19] C. T. Kelley. *Iterative methods for linear and nonlinear equations*. Society for Industrial and Applied Mathematics, Philadelphia, 1995.

[20] C. T. Kelley and R. Suresh. A new acceleration method for Newton's method at singular points. *SIAM J. Numer. Anal.*, 20:1001–1009, 1983.

[21] Jorge J. Moré. A collection of nonlinear model problems. *Lectures Appl. Math.*, 26:723–762, 1990.

[22] Jorge J. Moré, Burton S. Garbow, and Kenneth E. Hillstrom. Testing unconstrained optimization software. *ACM Trans. Math. Software*, 7:17–41, 1981.

[23] Jorge Nocedal and S. J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 1999.

[24] G. W. Reddien. On Newton's method for singular problems. *SIAM J. Numer. Anal.*, 15:993–996, 1978.

[25] Yousef Saad. Analysis of augmented Krylov subspace methods. *SIAM J. Matrix Anal. Appl.*, 18:435–449, 1997.

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

DISTRIBUTION:

5	MS 0316	Brett W. Bader, 9233
2	MS 9018	Central Technical Files, 8945-1
1	MS 0899	Technical Library, 9616