ON BACKTRACKING FAILURE IN NEWTON-GMRES METHODS WITH A DEMONSTRATION FOR THE NAVIER-STOKES EQUATIONS

RAYMOND S. TUMINARO*, HOMER F. WALKER†, AND JOHN N. SHADID‡

Accepted for publication in JCP, 180, pp. 549-558 (2002).

Abstract. In an earlier study of inexact Newton methods (JCP, 1997), we pointed out that certain counter-intuitive behavior may occur when applying residual backtracking to the Navier–Stokes equations with heat and mass transport. Specifically, it was observed that a Newton–GMRES method globalized by backtracking (linesearch, damping) may be less robust when high accuracy is required of each linear solve in the Newton sequence than when less accuracy is required. In this brief discussion, we offer a possible explanation for this phenomenon, together with an illustrative numerical experiment involving the Navier–Stokes equations.

Key words. Navier–Stokes equations, Newton's method, backtracking, linesearch, damping, Newton–GMRES methods, Newton–Krylov methods, Newton iterative methods, truncated Newton methods, inexact Newton methods.

AMS subject classifications. 65H10, 65F10

1. Introduction. We consider the simulation of fluid flow governed by the steady transport equations for momentum, heat, and mass transfer. Discretization of these equations gives rise to a system of nonlinear algebraic equations, the numerical solution of which can be very challenging. In most nontrivial calculations, the solution process is computationally intensive and requires sophisticated algorithms to cope with high nonlinearity, strong PDE coupling, and a large degree of nonsymmetry.

Newton's method is a potentially attractive nonlinear solution method because of its ability to address fully the coupling of the variables. In addition, it enjoys rapid (typically q-quadratic) convergence near a solution that is not hindered by bad scaling of the variables. However, the implementation of Newton's method involves special considerations. Determining Newton steps requires the solution of very large linear systems, and iterative linear algebra methods are typically preferred for this. Consequently, obtaining exact solutions of these systems is infeasible, and the appropriate method is an *inexact Newton method* [4].

In this paper, we explore the counter-intuitive relationship between robustness and solution accuracy of linear subproblems within an inexact Newton method. Specif-

*Sandia National Laboratories, MS 9214, P.O. Box 969, Livermore, CA 94551. Author's e-mail address: `tuminaro@cs.sandia.gov`. The work of this author was funded by the United States Department of Energy, Mathematical Information, and Computational Sciences Division, and was carried out at Sandia National Laboratories operated for the U.S. Department of Energy under contract no. DE-ACO4-94AL85000.

†Department of Mathematical Sciences, Worcester Polytechnic Institute, Worcester, MA 01609-2280. Author's e-mail address: `walker@wpi.edu`. The work of this author was supported in part through a summer visit to Sandia National Laboratories and also in part by National Science Foundation Grant DMS-9727128 and by Subcontract No. B341493, University of California, Lawrence Livermore National Laboratory.

‡Sandia National Laboratories, MS 1111, P.O. Box 5800, Albuquerque, NM 87185-1111. Author's e-mail address: `jnshadi@cs.sandia.gov`. The work of this author was funded by the United States Department of Energy, Mathematical Information, and Computational Sciences Division, and was carried out at Sandia National Laboratories operated for the U.S. Department of Energy under contract no. DE-ACO4-94AL85000.

ically, we focus on *Newton–GMRES methods*, in which the iterative method GMRES[1] [12] is used to solve approximately for steps of Newton's method. Newton–GMRES methods are representative of the broader class of *Newton–Krylov methods* and have enjoyed considerable success in fluid flow applications. Following [15], we assume the primary mechanism for enhancing robustness is a *backtracking* (*linesearch*, *damping*) technique that shortens steps as necessary to ensure adequate decrease in the residual of the nonlinear system.[2] See [7],[11], and [15] for discussions of the implementation of backtracking in Newton–GMRES methods and experiments that show its effectiveness on fluid flow applications and other large-scale problems. See also [5] for a general treatment of backtracking for exact Newton's method and [6] for a discussion of backtracking for general inexact Newton methods.

The accuracy with which Newton linear subproblems are solved effect both the solution efficiency and robustness of the inexact Newton method. This paper discusses robustness issues as solution efficiency is fairly well understood. In particular, efficiency suffers when requiring too much subproblem accuracy away from the solution as this leads to oversolving (cf. [7]), i.e., reducing the linear residual norm without achieving a commensurate reduction in the nonlinear residual norm. Since the purpose of a step is to reduce the nonlinear residual norm, oversolving is clearly associated with unproductive iterations of the linear solver. It should, however, be understood that considerable accuracy may be appropriate near the solution in order to realize the rapid local convergence of Newton's method. An effective strategy for minimizing oversolving is to use nonlinear residual information to determine adaptively the accuracy with which the linear subproblems are solved. That is, the accuracy required in solving the linear subproblems varies as the nonlinear algorithm proceeds, and the accuracy requirement at each step is based on how well the residual of the linear system reflects the behavior of the nonlinear residual. (See [7] for an extensive discussion.) This strategy often drastically improves computational efficiency by reducing oversolving in the sequence of linear subproblems.

It was demonstrated in [15], as well as in previous experiments in [7] and [11], that requiring too much accuracy in solving the linear subproblems can cause Newton–GMRES methods to become less robust, even with a backtracking globalization. In particular, it was shown that requiring too much accuracy may result in more frequent failure of the backtracking routine to determine an acceptable step, as well as inefficiency. This might seem counter-intuitive: one might expect overly-stringent accuracy requirements to cause the method to work harder but not to fail altogether.[3] This counter-intuitive behavior is the subject of this brief note. In §2, we develop heuristics that suggest why this phenomenon occurs. In essence, these indicate that, as the GMRES iterations proceed, the resulting inexact Newton step is likely to grow increasingly long, especially if the Jacobian (matrix) of the nonlinear residual has small singular values. Furthermore, the step may become increasingly nearly orthogonal to the gradient of the nonlinear residual norm, especially if the Jacobian is ill-conditioned. If the problem is more than mildly nonlinear, then a long step that

[1]For convenience, we usually do not distinguish between GMRES and the restarted version GMRES(m).

[2]Trust-region methods, which are also popular techniques for enhancing robustness, are not considered here. While these methods have many attractive features, their implementation tends to be more problematical when solving large-scale nonlinear equations because of the need to evaluate products involving the transpose of the Jacobian.

[3]Indeed, if f is continuously differentiable, then reducing the length of an inexact newton step sufficiently always yields an acceptable step [6].

is nearly orthogonal to the nonlinear residual norm gradient may have to be reduced in length many times by backtracking in order to achieve acceptable reduction of the nonlinear residual norm. Indeed, so many backtracking reductions may be necessary that a practical backtracking routine may declare failure before an acceptable step is found. Following these developments, we offer in §3 an illustrative experiment in which overly stringent accuracy requirements lead to backtracking failure, while more relaxed tolerances would have resulted in a successful step. A concluding discussion is given in §4.

2. A heuristic analysis. We write the nonlinear problem to be solved as $F(x) = 0$, where $F : \boldsymbol{R}^n \to \boldsymbol{R}^n$. At a step of a Newton–GMRES method, GMRES is applied to determine an approximate solution of the *Newton equation*

$$J(x)s = -F(x), \tag{2.1}$$

where x is the current approximate solution and $J(x)$ denotes the Jacobian (matrix) of F at x. In the inexact Newton framework, one first chooses a *forcing term* $\eta \in [0, 1)$ (cf. [7]) and then applies GMRES until an iterate s_k satisfies the *inexact Newton condition*

$$\|F(x) + J(x)s_k\| \leq \eta \|F(x)\|. \tag{2.2}$$

It is shown in [4] that local convergence to a solution is controlled by the forcing terms; in particular, by choosing the forcing terms to be sufficiently small, one can obtain local convergence that is as fast as desired, up to the (typically q-quadratic) rate of convergence of Newton's method. However, as noted in §1 above, choosing forcing terms that are too small may reduce the robustness of the method as well as its efficiency. We now develop a heuristic analysis to offer an explanation of this.

In the following, we consider only the Euclidean norm $\|\cdot\|_2$, although there are extensions to other inner-product norms. If M is any matrix or operator, then we denote its largest and smallest singular values by $\sigma_{\max}(M)$ and $\sigma_{\min}(M)$, respectively. (See [9] for matters pertaining to singular values and the singular value decomposition.) Then $\kappa_2(M) = \sigma_{\max}(M)/\sigma_{\min}(M)$ is the condition number of M with respect to $\|\cdot\|_2$, provided $\sigma_{\min}(M) \neq 0$.

We begin with two observations about a step of (exact) Newton's method. At a current approximate solution x, the Newton step is given by $s^N \equiv -J(x)^{-1}F(x)$. This immediately yields our first observation: Away from a solution, where $F(x)$ is not small, s^N is likely to be long if $J(x)$ has small singular values.

Now let θ^N denote the angle between s^N and the negative gradient direction for $\|F\|_2$ at x. We estimate how small $\cos\theta^N$ can be, treating for convenience $f \equiv \frac{1}{2}\|F\|_2^2$, for which $\nabla f = J^T F$ has the same direction everywhere as $\nabla \|F\|_2$. Letting $J(x) = U\Sigma V^T$ be the singular value decomposition of $J(x)$, we have

$$\begin{aligned} \cos\theta^N &= \frac{-\nabla f(x)^T s^N}{\|\nabla f(x)\|_2 \, \|s^N\|_2} = \frac{(-J(x)^T F(x))^T(-J(x)^{-1}F(x))}{\|J(x)^T F(x)\|_2 \, \|J(x)^{-1}F(x)\|_2} \\ &= \frac{\|F(x)\|_2^2}{\|J(x)^T F(x)\|_2 \, \|J(x)^{-1}F(x)\|_2} = \frac{\|U^T F(x)\|_2^2}{\|V\Sigma U^T F(x)\|_2 \, \|V\Sigma^{-1}U^T F(x)\|_2} \\ &= \frac{u^T u}{\|\Sigma u\|_2 \, \|\Sigma^{-1} u\|_2}, \end{aligned} \tag{2.3}$$

where $u = U^T F(x)$.

We claim that

$$\frac{1}{\kappa_2(J(x))} \le \min_{u\neq 0} \frac{u^T u}{\|\Sigma u\|_2\,\|\Sigma^{-1}u\|_2} \le \frac{2}{\kappa_2(J(x))}\,. \tag{2.4}$$

Indeed, one sees immediately that, for $u \neq 0$,

$$\frac{u^T u}{\|\Sigma u\|_2\,\|\Sigma^{-1}u\|_2} \ge \frac{1}{\kappa_2(\Sigma)} = \frac{1}{\kappa_2(J(x))}\,,$$

and the left-hand inequality in (2.4) holds. Taking $\hat{u} = (1, 0, \ldots, 0, 1)^T \in \boldsymbol{R}^k$, we have

$$\begin{aligned}
\min_{u\neq 0} \frac{u^T u}{\|\Sigma u\|_2\,\|\Sigma^{-1}u\|_2} &\le \frac{\hat{u}^T \hat{u}}{\|\Sigma \hat{u}\|_2\,\|\Sigma^{-1}\hat{u}\|_2} \\
&= \frac{2}{\sqrt{\sigma_{\max}(J(x))^2 + \sigma_{\min}(J(x))^2}\sqrt{\sigma_{\max}(J(x))^{-2} + \sigma_{\min}(J(x))^{-2}}} \\
&\le \frac{2}{\sigma_{\max}(J(x))\sigma_{\min}(J(x))^{-1}} = \frac{2}{\kappa_2(J(x))}\,,
\end{aligned}$$

and the right-hand inequality also holds.

It follows from (2.3) and (2.4) that an unfortunate combination of $F(x)$ and $J(x)$ can result in

$$\frac{1}{\kappa_2(J(x))} \le \cos\theta^N \le \frac{2}{\kappa_2(J(x))}\,.$$

This yields our second observation: The Newton step can be nearly orthogonal to the gradient of $\|F\|_2$ when $\kappa_2(J(x))$ is large, i.e., when J is ill-conditioned.

Since GMRES steps approach the Newton step in the limit, it is plausible that similar observations about GMRES steps hold to increasing degrees as the GMRES iterations proceed. In the following, we establish this analytically. To simplify the discussion, we assume that GMRES is not restarted and that no preconditioning is used. We also assume that GMRES is applied to (2.1) with zero as the initial iterate, which is typically the case with Newton–GMRES methods.

For convenience, we usually denote $F = F(x)$, $J = J(x)$, etc. We denote the kth *Krylov subspace* by $\mathcal{K}_k = \mathcal{K}_k(F, J) \equiv \text{span}\,\{F, JF, \ldots, J^{k-1}F\}$ and the operator restriction of J to $\mathcal{K}_k$ by $J|_{\mathcal{K}_k}$. See [12] for these and other GMRES-related matters raised here.

If s_k is the kth GMRES iterate, then s_k minimizes $\|F + Js\|_2$ over all $s \in \mathcal{K}_k$. We have the fundamental relation

$$JV_k = V_{k+1}H_k, \tag{2.5}$$

where $H_k \in \boldsymbol{R}^{(k+1)\times k}$ is upper Hessenberg and $V_k \in \boldsymbol{R}^{n\times k}$ and $V_{k+1} \in \boldsymbol{R}^{n\times(k+1)}$ are matrices the columns of which are the orthonormal Arnoldi basis vectors. With (2.5), one can show that $s_k = V_k y_k$, where y_k minimizes $\big\|\|F\|_2 e_1 + H_k y\big\|_2$ over all $y \in \boldsymbol{R}^k$ and $e_1 = (1, 0 \ldots, 0)^T \in \boldsymbol{R}^{k+1}$. Then $y_k = -\|F\|_2 H_k^+ e_1$, where "+" denotes pseudo-inverse, and

$$\|s_k\|_2 = \|y_k\|_2 = \|F\|_2\,\|H_k^+ e_1\|_2. \tag{2.6}$$

We assume that J is nonsingular and $Js_k \neq -F$ for $k < n$. Then, for $k \leq n$, H_k has rank k, and $\|H_k^+\|_2 = \sigma_{\min}(H_k)^{-1}$. Since

$$\sigma_{\min}(H_k) = \min_{y \in R^k, \|y\|_2 = 1} \|H_k y\|_2 = \min_{v \in \mathcal{K}_k, \|v\|_2 = 1} \|Jv\|_2 = \sigma_{\min}(J|_{\mathcal{K}_k})$$

and since $\mathcal{K}_k \subseteq \mathcal{K}_{k+1}$ for each k, it follows that $\|H_k^+\|_2$ grows monotonically in k to $\|J^{-1}\|_2$. Thus, in view of (2.6), we have the following:

Observation: As k grows, it is likely that the GMRES step s_k will become increasingly long, especially if $J(x)$ has small singular values.

We now consider the angle between s_k and the negative gradient direction for $\|F\|_2$, which we denote by θ_k. In the Appendix, we offer a heuristic development that suggests $\cos\theta_k$ may be comparable in size to $\kappa_2(J|_{\mathcal{K}_k})^{-1}$. Since $\kappa_2(J|_{\mathcal{K}_k})$ grows monotonically in k to $\kappa_2(J)$, we have the following:

Observation: As k grows, the GMRES step s_k may become increasingly nearly orthogonal to the gradient of $\|F\|_2$, especially if $J(x)$ is ill-conditioned.

Remark. These developments apply to any more general Newton–Krylov method that uses a “residual-minimizing” Krylov subspace method equivalent to GMRES. See [8] for a discussion of equivalent Krylov subspace methods.

3. An illustrative experiment. In our earlier study [15], we implemented a Newton–GMRES method with a backtracking globalization in the parallel finite-element reacting flow code MPSalsa [13]. Briefly, MPSalsa uses a Galerkin Least Squares (GLS) formulation for the spatial discretization, an inexact Newton scheme, and various parallel Krylov iterative methods as implemented in the Aztec library [10] for the linear solve. Details of this implementation can be found in [15, §4]; see also [14]. To illustrate the developments in §2 above, we applied this Newton–GMRES implementation to the well-known 2D driven cavity problem. This was posed in primitive-variable form and discretized using a 100×100 equally spaced mesh, resulting in 30,486 unknowns for the discretized problem. We took the GMRES restart value to be 200 and allowed a maximum of three restarts at each inexact Newton step, whether or not the inexact Newton condition (2.2) was satisfied. We used right preconditioning with a domain-based (overlapping Schwarz) ILU preconditioner. See [15, §4] for more details of the test problem, the discretization, the preconditioner, and the computing environment.

In this experiment, we used Reynolds number 2000 and a forcing term $\eta = 10^{-4}$ at each inexact Newton step.[4] In this case, the code declared failure at the thirteenth inexact Newton step after the backtracking routine reduced the steplength five times (for an overall reduction factor of 1.05×10^{-5}) without obtaining an acceptable step. In Table 3.1 below, we consider data associated with an illustrative subset of the GMRES iterations at this thirteenth inexact Newton step with an eye toward assessing not only the specific nature of the failure but also what would have happened if GMRES had terminated at an earlier stage, with subsequent backtracking. In Table 3.1, k is the GMRES iteration number, $\|F + Js_k\|_2 / \|F\|_2$ is the GMRES relative residual at the iterate s_k, $\cos\theta_k$ is the cosine of the angle between s_k and the negative gradient of $\|F\|_2$, λ_k is the step reduction factor resulting from backtracking from s_k, and $\|F(x + \lambda_k s_k)\|_2$ is the resulting nonlinear residual norm.

[4]It is important to note that this specific example is typical of many such runs and was not contrived in any way.

k	$\frac{\|F+Js_k\|_2}{\|F\|_2}$	$\cos\theta_k$	$\|s_k\|_2$	λ_k	$\|F(x+\lambda_k s_k)\|_2$
0	1.0		0	1.0	2.446
5	6.72e-01	6.24e-04	2.11e+06	1.0	1.673
10	5.06e-01	5.69e-04	3.14e+06	1.0	1.380
15	4.40e-01	5.51e-04	3.51e+06	1.0	1.385
20	3.72e-01	5.39e-04	3.83e+06	1.0	1.499
25	3.33e-01	5.26e-04	4.05e+06	1.0	1.645
50	2.78e-01	4.48e-04	4.94e+06	1.0	2.250
75	9.19e-02	4.47e-04	5.32e+06	1.0	1.603
100	7.42e-02	4.30e-04	5.55e+06	1.0	1.907
150	6.46e-02	3.17e-04	7.54e+06	4.98e-01	1.789
200	5.55e-02	4.78e-05	5.00e+07	2.15e-02	2.435
300	5.55e-02	4.79e-05	4.99e+07	2.15e-02	2.435
400	1.20e-02	2.16e-06	1.11e+09	1.05e-05	2.446
500	1.09e-02	2.16e-06	1.11e+09	1.05e-05	2.446
600	2.96e-03	2.23e-06	1.08e+09	1.05e-05	2.446

TABLE 3.1

Data associated with the GMRES iterations at inexact Newton step 13, at which backtracking failure occurred.

¿From Table 3.1, one sees that GMRES reduced the residual norm relatively rapidly in the early going but was ultimately unable to produce the requested reduction by a factor of 10^{-4} within the allowed 600 iterations. During the iterations, $\cos\theta_k$ was reduced by a factor of about $\frac{1}{280}$, while $\|s_k\|_2$ grew by a factor of more than 500. By the 400th iteration, the length of the GMRES step and its near-orthogonality to the gradient of $\|F\|_2$ combined to preclude the backtracking from determining an acceptable step within the allowable five steplength reductions. Even at the outset, $\cos\theta_k$ is small and $\|s_k\|_2$ is large in an absolute sense; however, the last column of Table 3.1 shows that backtracking would have produced an acceptable step through the 300th GMRES iteration. In fact, the GMRES step would have been acceptable without backtracking through the 100th iteration.

It is perhaps most notable in Table 3.1 that terminating GMRES around the tenth iteration, corresponding to $\eta \approx .5$, would have resulted in optimal reduction of $\|F\|_2$ with only a very modest number of GMRES iterations. Similar observations are valid for all previous inexact Newton steps. This is seen in Table 3.2 below, which contrasts the optimal and final GMRES iterates at each inexact Newton step. We note that the optimal GMRES iterate was acceptable without backtracking at every inexact Newton step, although this is not indicated in the table. We also note that, as seen in the third column of the table, a value of $\eta \approx .5$ would have resulted in near-optimal reduction of $\|F\|_2$ with very few GMRES iterations at every inexact Newton step. Of course, this value of η would have resulted in a different inexact Newton sequence, for which other values of η might have been more appropriate. Also, no fixed value of η is likely to be suitable for a broad range of problems. Adaptive choices that incorporate information about F have been shown to be most widely effective in the numerical experiments in [15]. Indeed, in the experiments in [15], the problem discussed here was successfully solved using such choices.

Inexact Newton Step	Optimal GMRES Iterate			Final GMRES Iterate		
	k	$\frac{\|F+Js_k\|_2}{\|F\|_2}$	$\frac{\|F(x+\lambda_k s_k)\|_2}{\|F\|_2}$	k	$\frac{\|F+Js_k\|_2}{\|F\|_2}$	$\frac{\|F(x+\lambda_k s_k)\|_2}{\|F\|_2}$
1	4	3.66e-01	.6483	149	9.26e-05	.8915
2	2	6.65e-01	.6578	347	9.94e-05	.9400
3	12	3.52e-01	.4872	377	9.88e-05	.8946
4	11	4.57e-01	.5808	589	9.93e-05	.9377
5	11	4.67e-01	.5796	600	1.04e-03	.9988
6	11	4.67e-01	.5877	600	2.49e-04	.9988
7	11	4.67e-01	.5762	600	1.15e-03	.9992
8	11	4.67e-01	.5821	600	4.08e-04	.9996
9	11	4.67e-01	.5753	600	1.25e-03	.9996
10	11	4.67e-01	.5834	600	1.81e-04	.9981
11	11	4.67e-01	.5629	600	1.08e-03	.9925
12	10	4.94e-01	.5927	600	2.09e-04	.9867
13	12	4.65e-01	.5519	600	2.96e-03	1.000

TABLE 3.2

Relative linear and nonlinear residual norm reduction at the optimal and final GMRES iterates.

4. Concluding discussion. Previous studies have shown that Newton–GMRES methods globalized with backtracking suffer loss of robustness in some cases when overly small forcing terms are used to terminate the GMRES iterations. In the foregoing, we have offered a heuristic analysis and an illustrative numerical experiment that suggest a possible explanation for this. The key observations are the following:

1. The step produced by GMRES is likely to become increasingly long as the iterations proceed, especially if the Jacobian has small singular values.
2. In addition, the step may become increasingly nearly orthogonal to the gradient of $\|F\|$, especially if the Jacobian is ill-conditioned.

These observations also apply to any other Newton–Krylov method that uses a residual minimizing Krylov subspace method.

¿From these observations, one sees that, while a small forcing term may result in a step s_k that considerably reduces $\|F(x)+J(x)s_k\|$, such a step may be so long that $F(x+s_k)$ differs significantly from the linear model $F(x)+J(x)s_k$. Then $\|F(x+s_k)\|$ may not be sufficiently reduced for s_k to be acceptable, in which case backtracking is necessary. If s_k is also nearly orthogonal to the gradient of $\|F\|$, then it is only a weak descent direction, and so many steplength reductions may be required to produce an acceptable step that the backtracking routine may declare failure before one is found. Conversely, a larger forcing term may yield a step s_k that reduces $\|F(x)+J(x)s_k\|$ less but is sufficiently short that $F(x)+J(x)s_k$ approximates $F(x+s_k)$ well. Then $\|F(x+s_k)\|$ is more likely to be reduced enough to accept s_k. Even if backtracking is required, s_k is likely to be a more robust descent direction, and backtracking is more likely to succeed. Of course, these considerations are most important away from a solution of the nonlinear problem. Near a solution of the nonlinear problem, $\|F\|$ is small and an accurate solution of (2.1) is relatively short and, therefore, appropriate.

These developments suggest that there might be benefits to an approach in which an estimate of $\kappa_2(J|_{\mathcal{K}_k})$ is monitored to help guide decisions on terminating the iterative solution of the linear subproblems. In a Newton–GMRES method, such an estimate can be economically provided by using incremental condition estimation (ICE) [1, 2] to estimate the condition numbers of the matrices H_k (see [3] for a discussion of this application of ICE). We leave this issue for future work.

Appendix. We develop a heuristic anaysis suggesting that $\cos\theta_k$ may be comparable in size to $\kappa_2(J|_{\mathcal{K}_k})^{-1}$. Again treating $f(x) \equiv \frac{1}{2}\|F(x)\|_2^2$, we denote the projection of ∇f onto $\mathcal{K}_k$ by $\Pi_k \nabla f$ and the angle between s_k and $-\Pi_k\nabla f$ by θ_{Π_k}. Since $s_k \in \mathcal{K}_k$, we have

$$\begin{aligned}\cos\theta_k &= \frac{-\nabla f^T s_k}{\|\nabla f\|_2\,\|s_k\|_2} = \frac{(-\Pi_k\nabla f)^T s_k}{\|\Pi_k\nabla f\|_2\,\|s_k\|_2}\cdot\frac{\|\Pi_k\nabla f\|_2}{\|\nabla f\|_2} \\ &= \cos\theta_{\Pi_k}\cdot\frac{\|\Pi_k\nabla f\|_2}{\|\nabla f\|_2} \le \cos\theta_{\Pi_k}\,. \end{aligned} \tag{A.1}$$

To complete our analysis, we show that $\cos\theta_{\Pi_k}$ may be comparable in size to $\kappa_2(J|_{\mathcal{K}_k})^{-1}$. Note that $\Pi_k\nabla f = V_k V_k^T J^T F$, that $V_k^T V_k = I \in \boldsymbol{R}^{k\times k}$, and that $\|V_k y\|_2 = \|y\|_2$ for all $y \in \boldsymbol{R}^k$. Also, recall that $s_k = V_k y_k$, where $y_k = -\|F\|_2 H_k^+ e_1$. Then with (2.5), we have

$$\begin{aligned}\cos\theta_{\Pi_k} &= \frac{(-\Pi_k\nabla f)^T s_k}{\|\Pi_k\nabla f\|_2\,\|s_k\|_2} = -\frac{(V_k V_k^T J^T F)^T s_k}{\|V_k V_k^T J^T F\|_2\,\|s_k\|_2} \\ &= -\frac{(H_k^T V_{k+1}^T F)^T V_k^T V_k y_k}{\|H_k^T V_{k+1}^T F\|_2\,\|V_k y_k\|_2} = -\frac{(H_k^T V_{k+1}^T F)^T y_k}{\|H_k^T V_{k+1}^T F\|_2\,\|y_k\|_2}\,. \end{aligned} \tag{A.2}$$

Set $h_1 \equiv H_k^T e_1$, where $e_1 = (1, 0\ldots, 0)^T \in \boldsymbol{R}^{k+1}$ as before. Since the first column of V_{k+1} is $F/\|F\|_2$, we have $H_k^T V_{k+1}^T F = \|F\|_2 h_1$. Furthermore,

$$y_k = -\|F\|_2 H_k^+ e_1 = -\|F\|_2 (H_k^T H_k)^{-1} H_k^T e_1 = -\|F\|_2 (H_k^T H_k)^{-1} h_1.$$

Then (A.2) yields

$$\cos\theta_{\Pi_k} = \frac{\|F\|_2^2\, h_1^T (H_k^T H_k)^{-1} h_1}{\|F\|_2^2\,\|h_1\|_2\,\|(H_k^T H_k)^{-1} h_1\|_2} = \frac{h_1^T (H_k^T H_k)^{-1} h_1}{\|h_1\|_2\,\|(H_k^T H_k)^{-1} h_1\|_2}\,.$$

Letting $H_k = U_k \begin{pmatrix}\Sigma_k \\ 0\end{pmatrix} W_k^T$ be the singular value decomposition of H_k, where $\Sigma_k \in \boldsymbol{R}^{k\times k}$, we obtain

$$\begin{aligned}\cos\theta_{\Pi_k} &= \frac{h_1^T W_k {\Sigma_k}^{-2} W_k^T h_1}{\|h_1\|_2\,\|W_k {\Sigma_k}^{-2} W_k^T h_1\|_2} \\ &= \frac{({\Sigma_k}^{-1} W_k^T h_1)^T ({\Sigma_k}^{-1} W_k^T h_1)}{\|\Sigma_k ({\Sigma_k}^{-1} W_k^T h_1)\|_2\,\|{\Sigma_k}^{-1}({\Sigma_k}^{-1} W_k^T h_1)\|_2} \\ &= \frac{u^T u}{\|\Sigma_k u\|_2\,\|{\Sigma_k}^{-1} u\|_2}, \end{aligned} \tag{A.3}$$

where $u = {\Sigma_k}^{-1} W_k^T h_1$.

Adapting the argument leading to (2.4), we have

$$\text{(A.4)} \qquad \frac{1}{\kappa_2(J|_{\mathcal{K}_k})} \le \min_{u \ne 0} \frac{u^T u}{\|\Sigma_k u\|_2 \, \|{\Sigma_k}^{-1} u\|_2} \le \frac{2}{\kappa_2(J|_{\mathcal{K}_k})} \ .$$

Then (A.3) and (A.4) suggest that $\cos\theta_{\Pi_k}$ may be comparable in size to $\kappa_2(J|_{\mathcal{K}_k})^{-1}$. We note, however, that since $u = {\Sigma_k}^{-1} W_k^T h_1$ in (A.3) cannot be arbitrarily specified, (A.3) and (2.4) do not *guarantee* that $\cos\theta_{\Pi_k}$ can be comparable in size to $\kappa_2(J|_{\mathcal{K}_k})^{-1}$.

REFERENCES

[1] C. H. Bischof, *Incremental condition estimation*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 312–322.

[2] C. H. Bischof and P. T. P. Tang, *Robust incremental condition estimation*, Tech. Report CS-91-133, LAPACK Working Note 33, Computer Science Department, University of Tennessee, May 1991.

[3] P. N. Brown and H. F. Walker, GMRES *on (nearly) singular systems*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 37–51.

[4] R. S. Dembo, S. C. Eisenstat, and T. Steihaug, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.

[5] J. E. Dennis, Jr. and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Series in Automatic Computation, Prentice-Hall, Englewood Cliffs, NJ, 1983.

[6] S. C. Eisenstat and H. F. Walker, *Globally convergent inexact Newton methods*, SIAM J. Optimization, 4 (1994), pp. 393–422.

[7] ———, *Choosing the forcing terms in an inexact Newton method*, SIAM J. Sci. Comput., 17 (1996), pp. 16–32.

[8] R. W. Freund, G. H. Golub, and N. M. Nachtigal, *Iterative solution of linear systems*, Acta Numerica 1992, 1 (1992), pp. 57–100.

[9] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, third ed., 1989.

[10] S. A. Hutchinson, L. Prevost, J. N. Shadid, C. Tong, and R. S. Tuminaro, *Aztec user's guide, version 2.0*, Tech. Report Sand99-8801J, Sandia National Laboratories, Albuquerque NM, 87185, Oct. 1999.

[11] M. Pernice and H. F. Walker, NITSOL*: a Newton iterative solver for nonlinear systems*, SIAM J. Sci. Comput., 19 (1998), pp. 302–318.

[12] Y. Saad and M. H. Schultz, GMRES*: a generalized minimal residual method for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.

[13] A. G. Salinger, K. D. Devine, G. L. Hennigan, H. K. Moffat, S. A. Hutchinson, and J. N. Shadid, **MPSalsa**, *A finite element computer program for reacting flow problems. Part 2 – User's Guide*, Tech. Report SAND96-2331, Sandia National Laboratories, Albuquerque, NM, 1996.

[14] J. N. Shadid, *A fully coupled Newton–Krylov solution method for parallel unstructured finite element fluid flow, heat and mass transport simulations*, Int. J. Comput. Fluid Dynamics, 12 (1999), pp. 199–211.

[15] J. N. Shadid, R. S. Tuminaro, and H. F. Walker, *An inexact Newton method for fully-coupled solution of the Navier–Stokes equations with heat and mass transport*, J. Comput. Phys., 137 (1997), pp. 155–185.