Slide 1

Numerical Methods for Large-Scale Nonlinear Equations

Homer Walker

MA 512

April 28, 2005

Slide 2

Inexact Newton and Newton–Krylov Methods

a. Newton-iterative and inexact Newton methods.
   i. Formulation and local convergence.
   ii. Globally convergent methods.
   iii. Choosing the forcing terms.

b. Newton–Krylov methods.
   i. General considerations.
   ii. Matrix-free implementations.

Slide 3

a. Newton-iterative and inexact Newton methods.

The model method will be . . .

Newton's Method:

Given an initial x.

Iterate:

Decide whether to stop or continue.

Solve $J(x)s = -F(x)$.

Update $x \leftarrow x + s$.

Here, $J(x) = F'(x) = \left(\frac{\partial F_i(x)}{\partial x_j}\right) \in \mathbb{R}^{n\times n}$.

Slide 4

About Newton's method, recall . . .

- Major strength: *quadratic local convergence*, which is often *mesh-independent* on discretized PDE problems [6], [1]. .
- For general discussions of stopping, scaling, "globalization" procedures, local and global convergence, etc., see [6].

Assume throughout: F is continuously differentiable.

Slide 5

Suppose that *iterative linear algebra methods* are preferred for solving

$$J(x)s = -F(x).$$

The resulting method is a **Newton iterative (truncated Newton) method**.

Key aspect: $J(x)s = -F(x)$ is *solved only approximately*.

Key issues:

- When should we stop the linear iterations?
- How should we globalize the method?
- Which linear solver should we use?

The first two can be well treated in the *strictly more general context* of *inexact Newton methods*.

Slide 6

An **inexact Newton method** [4] is *any* method each step of which reduces the norm of the local linear model of F.

Inexact Newton Method [4]:

Given an initial x.

Iterate:

Decide whether to stop or continue.

Find **some** $\eta \in [0, 1)$ and s that satisfy

$$\|F(x) + J(x)\,s\| \le \eta \|F(x)\|.$$

Update $x \leftarrow x + s$.

Slide 7

- A Newton iterative method fits naturally into this framework:
  - — Choose $\eta \in [0, 1)$.
  - — Apply the iterative linear solver until $\|F(x) + J(x)\, s\| \le \eta \|F(x)\|$.

▷ Used in this way, η is called a *forcing term*.

▷ The issue of stopping the linear iterations becomes the issue of *choosing the forcing terms*.

Slide 8

Local convergence is controlled by choices of η [4].

Theorem [4]: *Suppose $F(x_*) = 0$ and $J(x_*)$ is invertible. If $\{x_k\}$ is an inexact Newton sequence with x_0 sufficiently near x_*, then*

- $\eta_k \le \eta_{\max} < 1 \quad \Longrightarrow \quad x_k \to x_*$ *q-linearly**,
- $\eta_k \to 0 \quad \Longrightarrow \quad x_k \to x_*$ *q-superlinearly***,

If also J is Lipschitz continuous*** at x_*, then

- $\eta_k = O(\|F(x_k)\|) \quad \Longrightarrow \quad x_k \to x_*$ *q-quadratically*****.

* For some $\beta < 1$, $\|x_{k+1} - x_*\|_{J(x_*)} \le \beta \|x_k - x_*\|_{J(x_*)}$ for sufficiently large k, where $\|w\|_{J(x_*)} \equiv \|J(x_*)\, w\|$.

** $\|x_{k+1} - x_*\| \le \beta_k \|x_k - x_*\|$, where $\beta_k \to 0$.

*** For some λ, $\|J(x) - J(x_*)\| \le \lambda \|x - x_*\|$ for x near x_*.

**** For some C, $\|x_{k+1} - x_*\| \le C \|x_k - x_*\|^2$ for all k.

Slide 9

Proof idea:

Suppose $\|F(x) + J(x)\, s\| \le \eta \|F(x)\|$. Set $x_+ = x + s$.

We have $F(x_+) \approx F(x) + J(x)\, s \quad \Longrightarrow \quad \|F(x_+)\| \lessapprox \eta \|F(x)\|$.

Near $x_* \ldots$

$$\begin{aligned} F(x) &= F(x) - F(x_*) \approx J(x_*)\,(x - x_*), \\ F(x_+) &= F(x_+) - F(x_*) \approx J(x_*)\,(x_+ - x_*). \end{aligned}$$

So $\|J(x_*)\,(x_+ - x_*)\| \lessapprox \eta \|J(x_*)\,(x - x_*)\|$, i.e.,

$$\|x_+ - x_*\|_{J(x_*)} \lessapprox \eta \|x - x_*\|_{J(x_*)}$$

Slide 10

A *globally convergent* algorithm is ...

Inexact Newton Backtracking (INB) Method [7]:

Given an initial x and $t \in (0,1)$, $\eta_{\max} \in [0,1)$, $t \in (0,1)$, and $0 < \theta_{\min} < \theta_{\max} < 1$.

Iterate:

Decide whether to stop or continue.

Choose *initial* $\eta \in [0, \eta_{\max}]$ and s such that
$\|F(x) + J(x)\, s\| \le \eta \|F(x)\|$.

Evaluate $F(x+s)$.

While $\|F(x+s)\| > [1 - t(1-\eta)]\|F(x)\|$, do:
Choose $\theta \in [\theta_{\min}, \theta_{\max}]$.
Update $s \leftarrow \theta s$ and $\eta \leftarrow 1 - \theta(1-\eta)$.
Revaluate $F(x+s)$.

Update $x \leftarrow x + s$ and $F(x) \leftarrow F(x+s)$.

Slide 11

The global convergence result is . . .

Theorem [7, Th.6.1]: *Suppose $\{x_k\}$ is produced by the INB method. If $\{x_k\}$ has a limit point x_* such that $J(x_*)$ is nonsingular, then $F(x_*) = 0$ and $x_k \to x_*$. Furthermore, the initial s_k and η_k are accepted for all sufficiently large k.*

Possibilities:

- $\|x_k\| \to \infty$.
- $\{x_k\}$ has limit points, and J is singular at each one.
- $\{x_k\}$ converges to x_* such that $F(x_*) = 0$, $J(x_*)$ is nonsingular, and *asymptotic convergence is determined by the initial η_k's.*

Slide 12

Practical implementation of the INB method.

Minor details (most as before):

- Choose $\eta_{\max}$ near 1, e.g., $\eta_{\max} = .9$.
- Choose t small, e.g., $t = 10^{-4}$.
- Choose $\theta_{\min} = .1$, $\theta_{\max} = .5$.
- Take $\|\cdot\|$ to be an inner-product norm, e.g., $\|\cdot\| = \|\cdot\|_2$.
- Choose $\theta \in [\theta_{\min}, \theta_{\max}]$ to minimize a quadratic or cubic that interpolates $\|F(x_k + \theta s_k)\|$.

Slide 13

Choosing the forcing terms.

From [4], we know ...

- $\eta_k \leq$ constant $< 1 \implies$ local *linear* convergence.
- $\eta_k \to 0 \implies$ local *superlinear* convergence.
- $\eta_k = O(\|F(x_k)\|) \implies$ local *quadratic* convergence.

These allow practically implementable choices of the η_k's that lead to desirable asymptotic convergence rates.

But there remains the danger of **oversolving**, i.e., *imposing an accuracy on an approximate solution s of the Newton equation that leads to significant disagreement between $F(x+s)$ and $F(x) + J(x)\, s$*.

Slide 14

Example: **The driven cavity problem.**

$$(1/Re)\Delta^2\psi + \frac{\partial\psi}{\partial x_1}\frac{\partial}{\partial x_2}\Delta\psi - \frac{\partial\psi}{\partial x_2}\frac{\partial}{\partial x_1}\Delta\psi = 0 \quad \text{in } \mathcal{D} = [0,1]\times[0,1],$$

$$\text{On } \partial\mathcal{D},\ \psi = 0 \text{ and } \frac{\partial\psi}{\partial n} = \begin{cases} 1 & \text{on top.} \\ 0 & \text{on the sides and bottom.} \end{cases}$$

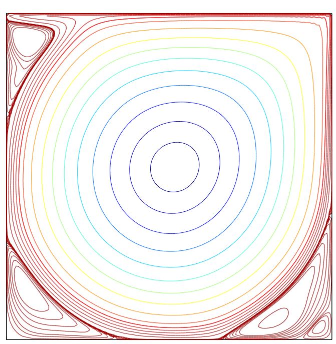

Streamlines for $Re = 10,000$.

Slide 15

For $\eta_k = \min\left\{\|F(x_k)\|_2, \dfrac{1}{k+2}\right\}$ (from [5]), ...

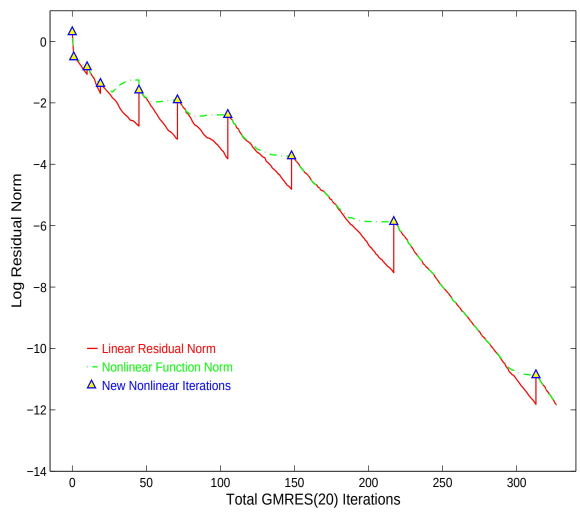

Performance on the driven cavity problem, $Re = 500$. *"Gaps" indicate oversolving.*

Slide 16

Forcing term choices have been proposed in [8] that are aimed at reducing oversolving. Here's one ...

Choice 1: Set $\eta_k = \min\{\eta_{\max}, \tilde{\eta}_k\}$, where

$$\tilde{\eta}_k = \frac{\Big| \|F(x_k)\| - \|F(x_{k-1}) + J(x_{k-1})\, s_{k-1}\| \Big|}{\|F(x_{k-1})\|}$$

- This directly reflects the (dis)agreement between F and its local linear model at the previous step.
- This is invariant under multiplication of F by a scalar.

Slide 17

If we use η_k given by Choice 1 in the backtracking method, we can combine a local convergence result in [8] with the previous global result to obtain ...

Theorem: *Suppose $\{x_k\}$ is produced by the INB method with each η_k given by Choice 1. If $\{x_k\}$ has a limit point x_* such that $J(x_*)$ is nonsingular and J is Lipschitz continuous at x_*, then $F(x_*) = 0$ and $x_k \to x_*$ with*

$$\|x_{k+1} - x_*\| \leq \beta \|x_k - x_*\| \|x_{k-1} - x_*\|.$$

for some β independent of k.

- It follows that the convergence is ...
  - ▷ r-order $(1+\sqrt{5})/2$,
  - ▷ q-superlinear,
  - ▷ two-step q-quadratic.

Slide 18

This and other choices in [8] may become too small too quickly away from a solution.

We recommend safeguards that work against this.

Rationale: If large forcing terms are appropriate at some point, then dramatically smaller forcing terms should be justified over several iterations before usage.

Choice 1 safeguard [8]: Modify η_k by

$$\eta_k \leftarrow \max\{\eta_k, \eta_{k-1}^{(1+\sqrt{5})/2}\}$$

whenever $\eta_{k-1}^{(1+\sqrt{5})/2} > .1$.

Slide 19

For *safeguarded Choice 1 η_k's*, ...

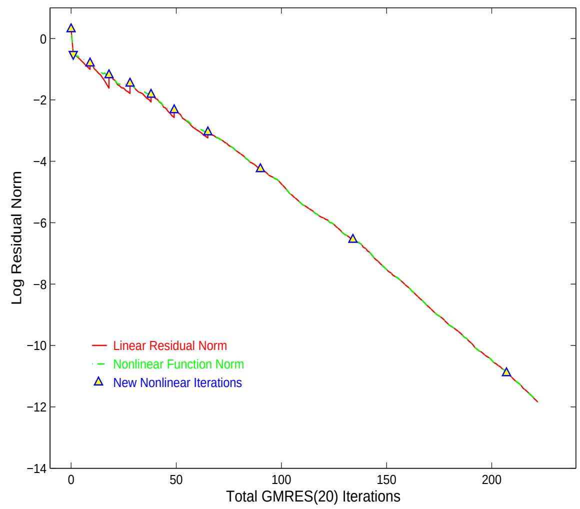

Performance on the driven cavity problem, $Re = 500$.
The inverted triangle indicates the safeguard value was used.

Slide 20

b. Newton–Krylov methods.

Idea: Implement a Newton iterative method using a *Krylov subspace method* as the linear solver.

- The term appears to have originated with [3].
- Naming conventions: Newton-GMRES, Newton–Krylov–Schwarz (NKS), Newton-Krylov-Multigrid (NKMG), ...
- "Truncated Newton" originated with [5], which outlined an implementation of Newton with CG.

Slide 21

General considerations.

- The linear system is $J(x)\,s = -F(x)$. The usual initial approximate solution is $s_0 = 0$.
- *The linear residual norm $\|F(x) + J(x)\,s\|$ is just the local linear model norm.*
- About **preconditioning** . . .
  - ▷ Preconditioning on the right retains compatibility between the norms used in the linear and nonlinear inexact Newton strategies.
  - ▷ *Preconditioning on the left may introduce incompatibilities.*
  - ▷ It *is* safe to "precondition the problem" on the left, i.e., to solve $M^{-1}F(x) = 0$ for an M that is *used without change throughout the solution process.*

Slide 22

Matrix-free implementations.

Krylov subspace methods require only products of $J(x)$ — and sometimes $J(x)^T$ — with vectors.

There are possibilities for producing these *without creating and storing $J(x)$.*

One possibility for products involving either $J(x)$ or $J(x)^T$ is **automatic differentiation**.

This is actively being explored in the Mathematics and Computer Science Division, Argonne National Lab. See the ANL Computational Differentiation Project web page , `www-unix.mcs.anl.gov/autodiff/index.html`.

Slide 23

A very widely used technique, applicable when only products involving $J(x)$ are needed, is **finite-difference approximation**.

For a local convergence analysis, see [2].

Given $v \in I\!R^n$, formulas for approximating $J(x)v$ to 1st, 2nd, 4th, and 6th order are ...

$$\frac{1}{\delta}[F(x+\delta v) - F(x)],$$

$$\frac{1}{2\delta}[F(x+\delta v) - F(x-\delta v)],$$

$$\frac{1}{6\delta}\left[8F(x+\frac{\delta}{2}v) - 8F(x-\frac{\delta}{2}v) - F(x+\delta v) + F(x-\delta v)\right],$$

$$\frac{1}{90\delta}\left[256F(x+\frac{\delta}{4}v) - 256F(x-\frac{\delta}{4}v) - 40F(x+\frac{\delta}{2}v) + 40F(x-\frac{\delta}{2}v) + F(x+\delta v) - F(x-\delta v)\right].$$

- The 1st-order formula is very commonly used; the others very rarely used (although sometimes they're needed, see [13], [12]).

Slide 24

Choosing δ.

- As before ...
  - ▷ We try to choose δ to roughly balance truncation and floating point error.
  - ▷ Fairly well-justified choices can be made for scalar functions. The justifications weaken with vector functions. Nothing is foolproof.
- A choice used in [12] that approximately minimizes a bound on the relative error in the difference approximation is based on ...

$$\delta = \frac{\left[(1+\|x\|)\epsilon_F\right]^{1/(p+1)}}{\|v\|},$$

where p is the difference order and ϵ_F is the relative error in F-evaluations ("function precision"). The main underlying assumption is that F and its derivatives up to order $p+1$ have about the same scale.

- A crude heuristic is $\delta = \epsilon^{1/(p+1)}$, where ϵ is machine epsilon.

Slide 25

Additional references and software.

In addition to the classic book of Dennis and Schnabel [6], I recommend the books by Kelley, especially [9] and also [10] and [11]. These come with MATLAB software. You can download PDF files of [9] and [10] as well as the associated software from the SIAM sites `http://www.siam.org/books/kelley/kelley.html` and `http://www.siam.org/books/fr18/`, resp.

See also Kelley's website `http://www4.ncsu.edu/eos/users/c/ctkelley/www/tim.html`.

For challenging large-scale applications, I recommend the following publicly available, downloadable software:

- NITSOL This is a flexible Newton–Krylov code written in Fortran. It is described in [12] and downloadable at `http://users.wpi.edu/~walker/NITSOL/`.
- NOX and KINSOL These are powerful codes for parallel solution of large-scale nonlinear systems, downloadable at `http://software.sandia.gov/trilinos/packages/nox/` and `http://www.llnl.gov/CASC/sundials/`, resp.
- PETSc This is a broad suite of codes for parallel solution of large-scale linear and nonlinear systems plus ancillary tasks such as preconditioning. It is downloadable at `http://www-unix.mcs.anl.gov/petsc/petsc-as/`.

References

[1] E. L. Allgower, K. Böhmer, F. A. Potra, and W. C. Rheinboldt, *A mesh-independence principle for operator equations and their discretizations*, SIAM J. Numer. Anal., 23 (1986), pp. 160–169.

[2] P. N. Brown, *A local convergence theory for combined inexact-newton/finite difference projection methods*, SIAM J. Numer. Anal., 24 (1987), pp. 407–434.

[3] P. N. Brown and Y. Saad, *Hybrid Krylov methods for nonlinear systems of equations*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 450–481.

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

[5] R. S. Dembo and T. Steihaug, *Truncated Newton algorithms for large-scale optimization*, Math. Prog., 26 (1983), pp. 190–212.

[6] J. E. Dennis, Jr. and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, vol. 16 of Classics in Applied Mathematics, SIAM, Philadelphia, PA, 1996.

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

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

[9] C. T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, vol. 16 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1995.

[10] ———, *Iterative Methods for Optimization*, vol. 18 of Frontiers in Applied Mathematics, SIAM, Philadelphia, PA, 1999.

[11] ———, *Solving Nonlinear Equations with Newton's Method*, vol. 1 of Fundamentals of Algorithms, SIAM, Philadelphia, PA, 2003.

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

[13] K. Turner and H. F. Walker, *Efficient high accuracy solutions with* GMRES(m), SIAM J. Sci. Stat. Comput., 13 (1992), pp. 815–825.