Classical iterative methods for solving linear systems

Ed Bueler

MATH 614 Numerical Linear Algebra

9 October 2015

examples: two concrete linear systems

- Suppose we want to solve the usual kind of linear system, namely

$$A\mathbf{x} = \mathbf{b} \tag{1}$$

  where $A \in \mathbb{C}^{m\times m}$ and $\mathbf{b} \in \mathbb{C}^{m\times 1}$, to find $\mathbf{x} \in \mathbb{C}^{m\times 1}$.
- Throughout these notes we will use just two examples,

  LS1

$$\begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 1 \\ 1 & 0 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 4 \end{bmatrix}$$

  LS2

$$\begin{bmatrix} 1 & 2 & 3 & 0 \\ 2 & 1 & -2 & -3 \\ -1 & 1 & 1 & 0 \\ 0 & 1 & 1 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 7 \\ 1 \\ 1 \\ 3 \end{bmatrix}$$

- On the first problem in Assignment #5 you will check that these are well-conditioned linear systems, and find their (unique) solutions.

residual

- By definition, the residual of a vector $\mathbf{x}_0$ in linear system (1) is the vector

$$\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0. \tag{2}$$

  - Making the residual zero is the same as solving the system.
  - We call $\mathbf{x}_0$ a "test vector" or "approximate solution".

- Evaluating the residual (2) requires a matrix-vector product and a vector subtraction. This requires in total $O(m^2)$ operations at worst.
  - But many practical linear problems involve matrices A for which the majority, often 99% or more, of the entries are zero.[1] Such a matrix is called *sparse*. The residual often requires $O(m)$ operations.
  - Thus, in many applications, evaluating the residual is much cheaper than applying Gauss elimination to solve linear system (1), which is an $O(m^3)$ operation in general.

[1] In such cases, A^{-1} is generally *dense*, i.e. most entries are nonzero.

Richardson iteration

- *Iterative methods* for linear system (1) attempt to solve it by operations like computing the residual, or applying A to a vector, plus other operations, to update the current vector approximation to $\mathbf{x}$.
  - One wants the sequence of approximations, the iterates, to *converge* to the solution $\mathbf{x} = A^{-1}\mathbf{b}$.
  - Iterative methods always require an initial iterate, the starting "guess."
- *Richardson iteration*, sometimes called *simple iteration*, adds a multiple ω of the last residual at each step,

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \omega(\mathbf{b} - A\mathbf{x}_k). \tag{3}$$

- For system LS1, using initial iterate $\mathbf{x}_0 = 0$ and $\omega = 1/5$, (3) gives:

$$\mathbf{x}_0 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{x}_1 = \begin{bmatrix} 0.4 \\ 0.2 \\ 0.8 \end{bmatrix}, \quad \mathbf{x}_2 = \begin{bmatrix} 0.6 \\ 0.16 \\ 1.04 \end{bmatrix}, \quad \mathbf{x}_3 = \begin{bmatrix} 0.728 \\ 0.088 \\ 1.096 \end{bmatrix},$$

$$\ldots, \quad \mathbf{x}_{10} = \begin{bmatrix} 0.99817 \\ -0.01167 \\ 1.00667 \end{bmatrix}, \quad \ldots$$

These iterates seem to be converging to $\mathbf{x} = [1\ 0\ 1]^*$.

eigenvalues and vectors

Recall:

- A complex number $\lambda \in \mathbb{C}$ is an *eigenvalue* of a square matrix $B \in \mathbb{C}^{m \times m}$ if there is a nonzero vector $\mathbf{v} \in \mathbb{C}^m$ so that $B\mathbf{v} = \lambda\mathbf{v}$.
- The set of all eigenvalues of B is the *spectrum* $\sigma(B)$ of B.
- The *spectral radius* $\rho(B)$ is the maximum absolute value of an eigenvalue:

$$\rho(B) = \max_{\lambda \in \sigma(B)} |\lambda|$$

- Even if B is real, λ may be complex.[2]
- Recall that *singular values* are also eigenvalues:

$$s \text{ is a singular value of } B \iff s \in \sigma(B^*B) = \sigma(BB^*)$$

[2] The roots of a polynomial with real coefficients may be complex.

spectral properties and convergence of iterations

- Properties of B described in terms of its eigenvalues or singular values are generically called *spectral properties*. Some examples:
  - $\rho(B)$
  - $\sigma_1(B) = \|B\|_2$
  - the 2-norm condition number

$$\kappa(B) = \|B\|_2 / \|B^{-1}\|_2 = \sigma_1(B)/\sigma_m(B)$$

- *the general idea*: The conditions determining whether an iterative method for solving $A\mathbf{x} = \mathbf{b}$ converges, or not, generally depend on spectral properties of A or associated matrices.
- The right-hand side $\mathbf{b}$ and the initial iterate $\mathbf{x}_0$ generally do not determine whether the iteration converges, though a good choice of $\mathbf{x}_0$ can speed up convergence.

convergence of the Richardson iteration

- Rewrite the Richardson iteration (3) as

$$\mathbf{x}_{k+1} = (I - \omega A)\mathbf{x}_k + \omega \mathbf{b} \tag{4}$$

- The lemma on the next slide shows that Richardson iteration converges if and only if all the eigenvalues of the matrix on the right side of (4) are inside the unit circle:

$$\text{(3) converges if and only if } \rho(I - \omega A) < 1. \tag{5}$$

- Recall $\rho(B) \leq \|B\|$ in any induced matrix norm—Exercise 3.2 in Trefethen & Bau. Thus (3) converges if $\|I - \omega A\| < 1$.

convergence lemma

Lemma

The iteration

$$\mathbf{y}_{k+1} = M\mathbf{y}_k + \mathbf{c}$$

converges to the solution of $\mathbf{y} = M\mathbf{y} + \mathbf{c}$ *for all initial* $\mathbf{y}_0$ *if and only if*

$$\rho(M) < 1.$$

Proof.

It is easy to solve this iteration by writing out a few cases:

$$\mathbf{y}_2 = M(M\mathbf{y}_0 + \mathbf{c}) + \mathbf{c} = M^2\mathbf{y}_0 + (I + M)\mathbf{c},$$
$$\mathbf{y}_3 = M(M^2\mathbf{y}_0 + (I + M)\mathbf{c}) + \mathbf{c} = M^3\mathbf{y}_0 + (I + M + M^2)\mathbf{c},$$

and so on. By induction we get

$$\mathbf{y}_k = M^k\mathbf{y}_0 + p_k(M)\mathbf{c}$$

where $p_k(x) = 1 + x + x^2 + \cdots + x^{k-1}$. But $p_k(x) \to 1/(1-x)$ as $k \to \infty$ iff $x \in (-1, 1)$. Also, $\rho(M) < 1$ iff $M^k \to 0$. Thus $\mathbf{y}_k \to (I - M)^{-1}\mathbf{c}$ iff $\rho(M) < 1$. □

the Richard iteration spectral condition

- Since Richardson iteration converges if $\rho(I - \omega A) < 1$ or $\|I - \omega A\| < 1$, we choose ω based on the principle that

  $$\omega A \text{ should be close to the identity } I$$

  - Not always possible! Some matrices *cannot* be rescaled to be close to I
- We can graph $f(\omega) = \rho(I - \omega A)$:

```
omega = -1:.01:1;
rho = zeros(size(omega));
for j = 1:length(omega)
    M = eye(n) - omega(j) * A;
    rho(j) = max(abs(eig(M)));
end
plot(omega,rho)
```

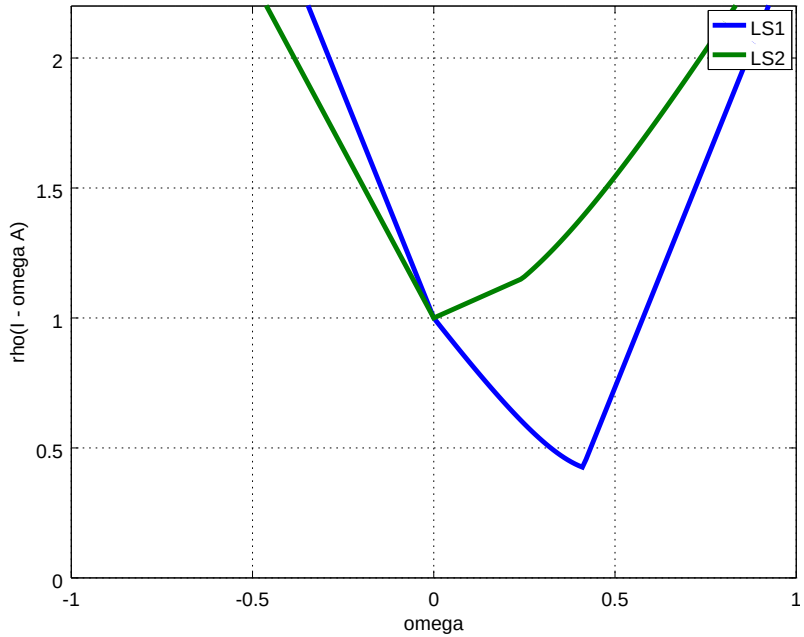

LS1: $\rho(I - \omega A)$ dips below 1 for $0 < \omega \lesssim 0.6$

LS2: $\rho(I - \omega A) \geq 1$ always

- Note $\rho(I - 0A) = 1$ when $\omega = 0$
- For LS1 we predict $\omega \approx 0.4$ gives fastest convergence.

Jacobi iteration

- Other classical iteration methods "split" the matrix A before iterating. The simplest is *Jacobi iteration*. It extracts the diagonal of A and inverts it.
- For this and the Gauss-Seidel iteration (next slide) we must assume
  all diagonal entries of A are nonzero: $a_{ii} \neq 0$
  This *is* a restriction; many invertible matrices have zeros on the diagonal.
- Split $A = D + L + U$ where D is diagonal, L is strictly lower triangular (so $\ell_{ij} = 0$ if $i \leq j$), and U is strictly upper triangular ($u_{ij} = 0$ if $i \geq j$).
- The Jacobi iteration is

$$D\mathbf{x}_{k+1} = b - (L + U)\mathbf{x}_k \tag{6}$$

- We could also write it $\mathbf{x}_{k+1} = D^{-1}\left(b - (L + U)\mathbf{x}_k\right)$ or as

$$x_i^{(k+1)} = \frac{1}{a_{ii}}\left(b_i - \sum_{j \neq i} a_{ij} x_j^{(k)}\right)$$

where $x_j^{(k)}$ denotes the jth entry of the kth iterate $\mathbf{x}_k$.

Gauss-Seidel iteration

- *Gauss-Seidel iteration* extracts the non-strict lower-triangular part of A and inverts it.
- Again if $A = D + L + U$, with the same meanings as before, the Gauss-Seidel iteration is

$$(D + L)\mathbf{x}_{k+1} = b - U\mathbf{x}_k \tag{7}$$

- We could also write it $\mathbf{x}_{k+1} = (D + L)^{-1}(b - U\mathbf{x}_k)$ ...but that would miss the point!
- Instead write it as $D\mathbf{x}_{k+1} = b - U\mathbf{x}_k - L\mathbf{x}_{k+1}$ or equivalently:

$$x_i^{(k+1)} = \frac{1}{a_{ii}}\left(b_i - \sum_{j<i} a_{ij}x_j^{(k)} - \sum_{j>i} a_{ij}x_j^{(k+1)}\right) \tag{8}$$

- Note the lower triangular entries apply to entries of $\mathbf{x}_{k+1}$ which have already been computed.
- Form (8) is actually *easier* to implement than Jacobi.

convergence conditions for Jacobi and Gauss-Seidel

- We know enough already to say that
  - Jacobi iteration converges if and only if $\rho(D^{-1}(L+U)) < 1$
  - Gauss-Seidel iteration converges if and only if $\rho((D+L)^{-1}U) < 1$
- However, these conditions are hard to use in practice, because computing these spectral radii (i.e. maximum eigenvalues) is harder than solving the original system.
- *Definition*. A is *strictly diagonally-dominant* if $|a_{ii}| > \sum_{j \neq i} |a_{ij}|$.
- Two of the relatively-famous[3] theorems are these:
  - *Theorem*. Both the Jacobi and Gauss-Seidel iterations converge if A is strictly diagonally-dominant.
  - *Theorem*. Gauss-Seidel iteration converges if A is symmetric positive definite.
- Unlike the "$\rho(\dots) < 1$" conditions above, these are only *sufficient* conditions.
  - E.g. there are nonsymmetric A, which are *not* diagonally-dominant, for which the iterations converge.

[3] Though I do not know who first proved these.

history, past and future

- The Jacobi and Gauss-Seidel iterations are from the 19th century.
  - Probably Richardson was too, though it appears in a 1910 publication.
- The early history of numerical partial differential equations, e.g. in the 1920 to 1970 period, heavily used these classical iterations.
  - A generalization of Gauss-Seidel iteration called *successive over-relaxation*, was a particular favorite around 1970.
- But none of these iterations work on system LS2.
- There are other iterative ideas, and they flourished in the 1990s. One of the best known, “GMRES” from a 1992 paper, is in Lecture FIXME of Trefethen & Bau. It is in widespread use. It works on LS2.
- Iteration is the future:
  - It’s obligatory on sufficiently-big systems.
  - It works better in parallel than direct methods.

biographies

- Gauss (1777–1855) did big stuff, while the Gauss-Seidel method is a small thing:
  `en.wikipedia.org/wiki/Carl_Friedrich_Gauss`
- Jacobi (1804–1851) was a major mathematician. He also has his name on the "Jacobian", the matrix of derivatives appearing in Newton's method for systems of equations:
  `en.wikipedia.org/wiki/Carl_Gustav_Jacob_Jacobi`
- Richardson (1881–1953) is the most interesting.[4] He invented numerical weather forecasting, doing it by-hand for fun during WWI. Later, as a pacifist and quaker, he quit the subject entirely when he found his meteorological work was of most value to chemical weapons designers and the British Air Force:
  `en.wikipedia.org/wiki/Lewis_Fry_Richardson`
- Seidel (1821–1896) is relatively little known:
  `en.wikipedia.org/wiki/Philipp_Ludwig_von_Seidel`

[4] To me, he is the best answer to "Who was the last person to *not* specialize and still be important in math and physics?" Same for Gauss, and perhaps Jacobi, but they were earlier.