THE CHEBYSHEV ITERATION REVISITED

Martin H. Gutknecht
Seminar for Applied Mathematics, ETH-Zentrum HG, CH-8092 Zürich, Switzerland,
Email: mhg@sam.math.ethz.ch

Stefan Röllin
Integrated Systems Laboratory, ETH-Zentrum ETZ, CH-8092 Zürich, Switzerland,
Email: roellin@iis.ee.ethz.ch

Abstract

C ompared to Krylov space methods based on orthogonal or oblique projection, the Chebyshev iteration does not require inner products and is therefore particularly suited for massively parallel computers with high communication cost. We compare six different algorithms that implement this methods and compare them with respect to roundoff effects, in particular, the ultimately achievable accuracy. Two of these algorithms replace the three-term recurrences by more accurate coupled two-term recurrences and seem to be new. But we also show that, for real data, the classical three-term Chebyshev iteration is never seriously affected by roundoff, in contrast to the corresponding version of the conjugate gradient method. Even for complex data, strong roundoff effects are seen to be limited to very special situations where convergence is anyway slow.

The Chebyshev iteration is applicable to symmetric definite linear systems and to nonsymmetric matrices whose eigenvalues are known to be confined to an elliptic domain that does not include the origin. We also consider a corresponding stationary 2-step method, which has the same asymptotic convergence behavior and is additionally suitable for mildly nonlinear problems.

THE CHEBYSHEV ITERATION REVISITED

Martin H. Gutknecht
Seminar for Applied Mathematics, ETH-Zentrum HG, CH-8092 Zürich, Switzerland,
Email: mhg@sam.math.ethz.ch

Stefan Röllin
Integrated Systems Laboratory, ETH-Zentrum ETZ, CH-8092 Zürich, Switzerland,
Email: roellin@iis.ee.ethz.ch

Abstract

T he Chebyshev iteration has been one of the favorite Krylov space methods for solving a large sparse linear system of equations in a parallel environment, since, unlike methods based on orthogonalization (such as the conjugate and biconjugate gradient methods and GMRES—to name a few), it does not require to compute communication-intensive inner products for the determination of the recurrence coefficients. Only the monitoring of the convergence, that is, the determination of the norm of the residuals requires inner products, and even this norm needs to be evaluated only occasionally since its time-dependence can be forecast quite reliably.

The Chebyshev iteration requires some preliminary knowledge about the spectrum $\sigma(\mathsf{A})$ of the coefficient matrix A: an elliptic domain $\mathcal{E} \supset \sigma(\mathsf{A})$ with $0 \notin \mathcal{E}$ is normally assumed to be known in advance. Denote the center of the ellipse by d, its foci by $d \pm c$, and the lengths of the large and the small semi-axes by a and b. When $b = 0$, the elliptic domain turns into an interval. Recursions for the residuals r_n and the iterates x_n are found from the standard three-term recursions for the classical Chebyshev polynomials T_n by translating the latter from the interval $[-1, 1]$ to the interval $\{d + c\tau : -1 \leq \tau \leq +1\}$ and scaling them so that the value at 0 is 1: we obtain

$$\begin{aligned} \mathsf{r}_{n+1} &= (\mathsf{A}\mathsf{r}_n - d\mathsf{r}_n - \beta_{n-1}\mathsf{r}_{n-1})/\gamma_n\,, && (1)\\ \mathsf{x}_{n+1} &= (\mathsf{r}_n + d\mathsf{x}_n + \beta_{n-1}\mathsf{x}_{n-1})/\gamma_n\,, && (2) \end{aligned}$$

where, for $n \geq 2$, if we let $\eta := -d/c$,

$$\gamma_n := \frac{c}{2}\frac{T_{n+1}(\eta)}{T_n(\eta)} = -\left(d + \left(\frac{c}{2}\right)^2 \frac{1}{\gamma_{n-1}}\right), \qquad \beta_{n-1} := \frac{c}{2}\frac{T_{n-1}(\eta)}{T_n(\eta)} = \left(\frac{c}{2}\right)^2 \frac{1}{\gamma_{n-1}},$$

while $\beta_{-1} := 0,\ \beta_0 := -\frac{c^2}{2d},\ \gamma_0 := -d,\ \gamma_1 := -\left(d + \frac{c^2}{2d}\right)$.

Chebyshev iteration is known to be optimal in the sense that it yields the smallest nth maximum residuals, if the maximum is taken over all normal matrices whose eigenvalues lie on a real interval not containing 0. It is asymptotically optimal if they lie inside the mentioned ellipse.

However, as has recently been shown by Gutknecht and Strakoš, implementations of methods based on recurrences of the form (1) and (2) may suffer from a large gap between the updated residuals r_n and true residuals $\mathsf{b} - \mathsf{A}\mathsf{x}_n$; this means that the algorithm may stagnate early with relatively large true residuals, that is, low ultimate accuracy. This holds, in particular, for corresponding implementations of the conjugate gradient (CG) and biconjugate gradient (BiCG) methods. However, other implementations produce more accurate solutions, that is, stagnate ultimately with smaller true residuals.

Here we compare six different algorithms for the Chebyshev iteration and compare them with respect to the ultimately achievable accuracy. Two of these algorithms replace the three-term

recurrences by more accurate coupled two-term recurrences and are perhaps new. But we also show that, for real data, the classical three-term Chebyshev iteration using (1) and (2) is never seriously affected by roundoff, in contrast to the corresponding version of the conjugate gradient method. Even for complex data, strong roundoff effects are seen to be limited to very special situations where convergence is anyway slow.

The Chebyshev iteration is applicable to symmetric definite linear systems and to nonsymmetric matrices whose eigenvalues are known to be confined to an elliptic domain that does not include the origin. We also consider a corresponding stationary 2-step method, which has the same asymptotic convergence behavior and is additionally suitable for mildly nonlinear problems.