IMA Journal of Numerical Analysis (2014) **34**, 462–479
doi:10.1093/imanum/drt025
Advance Access publication on July 25, 2013

GMRES convergence bounds that depend on the right-hand-side vector

DAVID TITLEY-PELOQUIN, JENNIFER PESTANA* AND ANDREW J. WATHEN
Mathematical Institute, University of Oxford, Oxford OX1 3LB, UK
*Corresponding author: Jennifer.Pestana@maths.ox.ac.uk

[Received on 29 January 2012; revised on 12 April 2013]

We consider the convergence of the algorithm GMRES of Saad and Schultz for solving linear equations $Bx = b$, where $B \in \mathbb{C}^{n\times n}$ is nonsingular and diagonalizable, and $b \in \mathbb{C}^n$. Our analysis explicitly includes the initial residual vector r_0. We show that the GMRES residual norm satisfies a weighted polynomial least-squares problem on the spectrum of B, and that GMRES convergence reduces to an ideal GMRES problem on a rank-1 modification of the diagonal matrix of eigenvalues of B. Numerical experiments show that the new bounds can accurately describe GMRES convergence.

Keywords: GMRES; convergence analysis; iterative methods; linear systems; Krylov subspace methods.

1. Introduction

Let x_k denote the kth iterate of the algorithm GMRES (Saad & Schultz, 1986) applied to $Bx = b$, $B \in \mathbb{C}^{n\times n}$, $b \in \mathbb{C}^n$ with the corresponding residual vector $r_k \equiv b - Bx_k$. By definition,

$$\|r_k\|_2 = \min_{\substack{q\in\Pi_k\\ q(0)=1}} \|q(B)r_0\|_2, \tag{1.1}$$

where Π_k denotes the set of polynomials of degree at most k. For a more thorough description of the algorithm and its implementation details, see Saad & Schultz (1986) or, for example, the textbook of Saad (1996) or Greenbaum (1997).

Obtaining generally descriptive convergence bounds for GMRES has been an active research topic since the algorithm was introduced in Saad & Schultz (1986) and is still to date considered a difficult open problem. A typical first step in deriving convergence bounds is the following, which is an immediate consequence of (1.1):

$$\frac{\|r_k\|_2}{\|r_0\|_2} \leqslant \min_{\substack{q\in\Pi_k\\ q(0)=1}} \|q(B)\|_2. \tag{1.2}$$

The polynomial approximation problem in (1.2) is widely referred to as the *ideal* GMRES problem (Greenbaum & Trefethen, 1994). Bounds on (1.2) can be obtained using the spectral decomposition of B (if B is diagonalizable) (Saad & Schultz, 1986), its field of values (Eisenstat *et al.*, 1983; Eiermann, 1993; Eiermann & Ernst, 2001; Beckermann *et al.*, 2006; Liesen & Tichý, 2012) or its pseudospectra (Trefethen, 1990). An overview of these approaches is given in Simoncini & Szyld (2007, Section 6).

The bound (1.2) is a worst-case bound that holds over all $r_0 \in \mathbb{C}^n$. In practice, however, one has a specific B and a specific r_0, most likely *not* the worst case. Consequently, (1.2) may not be descriptive of actual GMRES convergence. In short, practical convergence bounds should take into account the effect of the initial residual vector r_0. Indeed, r_0 determines the Krylov subspace from which the approximation to x is sought. Arguments along these general lines have also been made, for instance, in Toh (1997), Driscoll *et al.* (1998, Section 8), Liesen & Strakoš (2004, Section 3.1; 2013, Section 3) and Duintjer Tebbens & Meurant (2012). Some analyses that explicitly include r_0 can be found in Ipsen (2000) (when B is a Jordan block), Liesen & Strakoš (2004) and Li & Zhang (2009) (when B is Toeplitz tridiagonal), and for more general matrices B in Liesen (2000). In this note, by considering the influence of the right-hand-side vector, we attempt to further identify situations in which convergence is much faster than would be predicted by (1.2).