# ‘Superlinear’ Convergence of GMRES for Non-Normal Matrices

Mark Embree
embree@rice.edu
Department of Computational and Applied Mathematics
Rice Univeristy
Houston, Texas 77005-1892

Bounds for the GMRES algorithm for solving nonsymmetric linear systems typically describe convergence in terms of a polynomial approximation problem posed on a continuum in the complex plane (containing the spectrum, or the numerical range, or a pseudospectrum) leading to predictions of convergence at a linear rate.

In practice, this linear rate often improves as the iteration progresses (called, in a mild abuse of terminology, 'superlinear' convergence). Van der Sluis and van der Vorst explained this rigorously for the conjugate gradient algorithm: At some point in the iteration, the optimal iteration polynomial exploits the discrete nature of some part of the spectrum (most notably those eigenvalues closest to the origin). This work was adapted to GMRES convergence by van der Vorst and Vuik, but the resulting bound is complicated by the fact that eigenvectors of nonsymmetric matrices typically do not form an orthonormal set. This bound involves the condition number of an eigenvector matrix, making it unclear if superlinear behavior is possible for matrices that are far from normal.

In this talk, we address superlinear convergence for systems involving nonnormal matrices. While, generally speaking, increased distance from normality leads to slower convergence overall, we show that nonnormality can actually enhance superlinear effects. Our rigorous explanation involves resolvent norms and pseudospectra, but the intuition is rather natural: Suppose the matrix exhibits ‘low dimensional non-normality’ — some small subset of the spectrum has nearly parallel eigenvectors, or more generally, the pseudospectra associated with some eigenvalues are much larger than for a normal matrix with the same spectrum. Such eigenvalues act as a thorn in the side of GMRES: They slow convergence more acutely than if they were normal eigenvalues, and thus GMRES benefits all the more if they are eliminated. While one expects superlinear convergence to occur because of well-separated eigenvalues near the origin, we show that it can result from eigenvalues anywhere in the spectrum, provided they are sufficiently ill-conditioned.