

The effect of aggressive early deflation on the convergence of the QR algorithm

Daniel Kressner

TU Berlin

kressner@math.ethz.ch

Harrachov 2007

Computational Methods with Applications

08/25/07

QR and QZ algorithms

QR algorithm	standard method for solving dense nonsymmetric eigenvalue problems ($\text{eig}(A)$)
QZ algorithm	standard method for solving dense generalized eigenvalue problems ($\text{eig}(A, B)$)

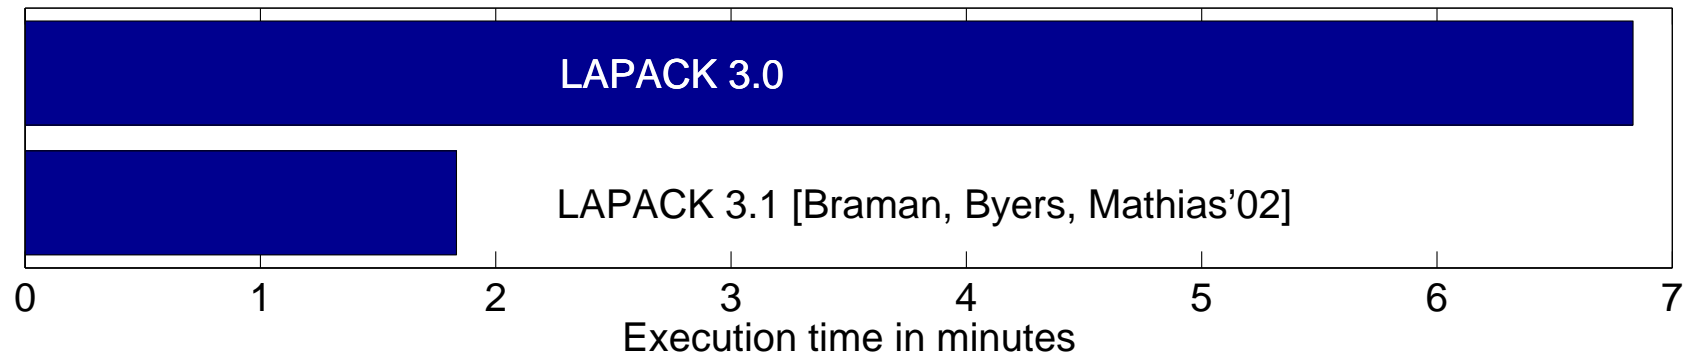
Both have recently undergone [significant improvements](#).

QR and QZ algorithms

QR algorithm	standard method for solving dense nonsymmetric eigenvalue problems ($\text{eig}(A)$)
QZ algorithm	standard method for solving dense generalized eigenvalue problems ($\text{eig}(A, B)$)

Both have recently undergone **significant improvements**.

QR for 2961×2961 matrix (Matrix Market's PDE2961):

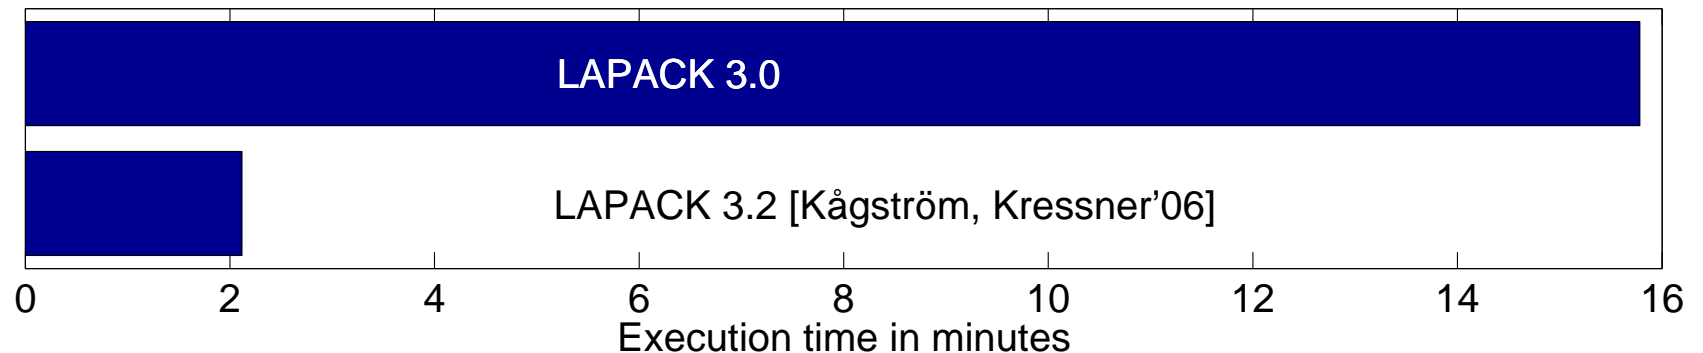


QR and QZ algorithms

QR algorithm	standard method for solving dense nonsymmetric eigenvalue problems ($\text{eig}(A)$)
QZ algorithm	standard method for solving dense generalized eigenvalue problems ($\text{eig}(A, B)$)

Both have recently undergone **significant improvements**.

QZ for 3600×3600 matrix pencil (Matrix Market's BCSST21):



New ingredients: block multishift + **advanced deflation techniques**.

Goal of QR algorithm

For real $n \times n$ matrix A , compute orthogonal Q s.t.

$$Q^T A Q = T = \begin{bmatrix} \blacksquare & & & \\ & \blacksquare & & \\ & & \blacksquare & \\ & & & \blacksquare \end{bmatrix},$$

where T is in real Schur form.

- Diagonal of T yields **eigenvalues** of A .
- First k columns of Q span **invariant subspace** \mathcal{X} : $A\mathcal{X} \subseteq \mathcal{X}$.

Goal of QR algorithm

For real $n \times n$ matrix A , compute orthogonal Q s.t.

$$Q^T A Q = T = \begin{bmatrix} \blacksquare & & & \\ & \blacksquare & & \\ & & \blacksquare & \\ & & & \blacksquare \end{bmatrix},$$

where T is in real Schur form.

- Diagonal of T yields **eigenvalues** of A .
- First k columns of Q span **invariant subspace** \mathcal{X} : $A\mathcal{X} \subseteq \mathcal{X}$.

Goal of this talk

- Show intimate relation between QR algorithm and Krylov subspace techniques.
- Provide intuitive explanation and convergence bounds capturing observed improvements.

The basic multishift QR algorithm

QR generates sequence of orthogonally similar matrices

$$A_0, A_1, A_2, \dots$$

The basic multishift QR algorithm

QR generates sequence of orthogonally similar matrices

$$A_0, A_1, A_2, \dots$$

(i) Initial reduction to **Hessenberg form**

$$A_0 \leftarrow Q_0^T A Q_0 = \begin{bmatrix} \square & & \\ & \square & \\ & & \square \end{bmatrix}$$

The basic multishift QR algorithm

QR generates sequence of orthogonally similar matrices

$$A_0, A_1, A_2, \dots$$

(i) Initial reduction to **Hessenberg form**

$$A_0 \leftarrow Q_0^T A Q_0 = \begin{bmatrix} \square & & \\ & \square & \\ & & \square \end{bmatrix}$$

(ii) **QR iterations** (preserve Hessenberg form)

for $i \leftarrow 1, 2, \dots$

Select $m \ll n$ shifts $\sigma_1, \dots, \sigma_m$.


 QR factorization $p_i(A_{i-1}) = Q_i R_i$ with $p_i(z) = \prod_{j=1}^m (z - \sigma_j)$.

 Update $A_i \leftarrow Q_i^T A_{i-1} Q_i$.

end for

Careful implementation: implicit, bulge chasing.

Idea behind blocked multishift QR



Do computations only locally (**red area**), delay and accumulate updates of rest (**blue area**).

- Bulk of computation becomes level BLAS 3 (matrix-matrix multiplications).
- Typically reduces execution time by factor 2–3.

[Lang'99], [Braman, Byers, Mathias'02].

Progress made by multishift QR iterations

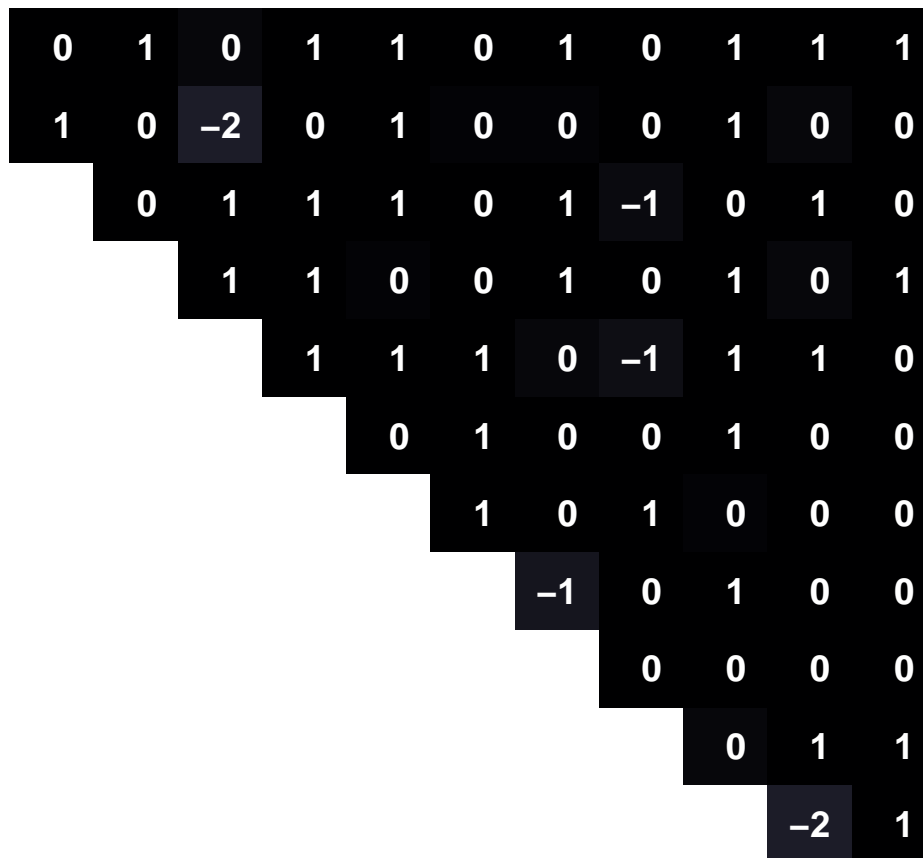
`A = hess(randn(11));` After 1 iteration ...

0	1	1	1	0	1	1	1	1	1	0
1	0	1	0	0	1	0	0	0	1	1
	1	0	0	-1	1	1	1	1	1	1
		1	0	1	1	0	0	1	0	0
			1	1	0	0	1	1	1	1
				1	1	1	0	1	0	0
					1	-1	1	0	1	0
						0	0	0	1	0
							1	0	0	1
								0	0	0
									-1	1

shifts = eigenvalues of bottom right 4×4 submatrix.

Progress made by multishift QR iterations


`A = hess(randn(11));` After 2 iterations ...



shifts = eigenvalues of bottom right 4×4 submatrix.

Progress made by multishift QR iterations


`A = hess(randn(11));` After 3 iterations ...



shifts = eigenvalues of bottom right 4×4 submatrix.

Progress made by multishift QR iterations


`A = hess(randn(11))`; After 4 iterations ...



shifts = eigenvalues of bottom right 4×4 submatrix.

Progress made by multishift QR iterations

`A = hess(randn(11))`; After 5 iterations ...



shifts = eigenvalues of bottom right 4×4 submatrix.

Idea of aggressive early deflation

`A = hess(randn(11));` After 4 iterations ...

0	1	0	0	1	0	1	0	0	1	1
1	0	0	-2	1	0	0	0	1	1	0
	0	1	1	1	1	1	0	0	1	0
		1	1	0	1	0	1	1	0	1
			-1	1	1	1	1	0	1	1
				1	0	0	1	0	-1	0
					0	1	-1	1	0	0
						-10	0	0	0	1
							0	0	-1	0
								-7	1	1
									-3	1

It turns out that some eigenvalues have already converged.

Idea of aggressive early deflation

Compute Schur form of bottom 4×4 submatrix ...

0	1	0	0	1	0	1	0	0	1	1
1	0	0	-2	1	0	0	1	-1	1	0
	0	1	1	1	1	1	0	-1	1	0
		1	1	0	1	0	1	1	0	1
			-1	1	1	1	1	1	1	1
				1	0	0	0	1	-1	0
					0	1	1	0	0	0
						-11	0	0	0	0
						-11	1	0	0	1
							-19		1	1
									-4	1

Two eigenvalues can be deflated!

Idea of aggressive early deflation

Or: Compute Schur form of bottom 9×9 submatrix ...

0	1	0	-1	1	1	0	0	0	1	1
1	0	0	0	1	0	0	1	-1	1	0
	0	1	1	1	0	1	1	0	1	-2
	0	1	1	1	1	1	1	0	1	1
	-2			1	1	1	1	1	1	1
	-3				1	1	0	0	0	0
	-4				0	1	0	0	0	0
	-14						0	0	0	0
	-14						1	0	0	1
	-22								1	1
	-24								-4	1

Four eigenvalues can be deflated!

A closer look at the Hessenberg matrix

After i iterations ...

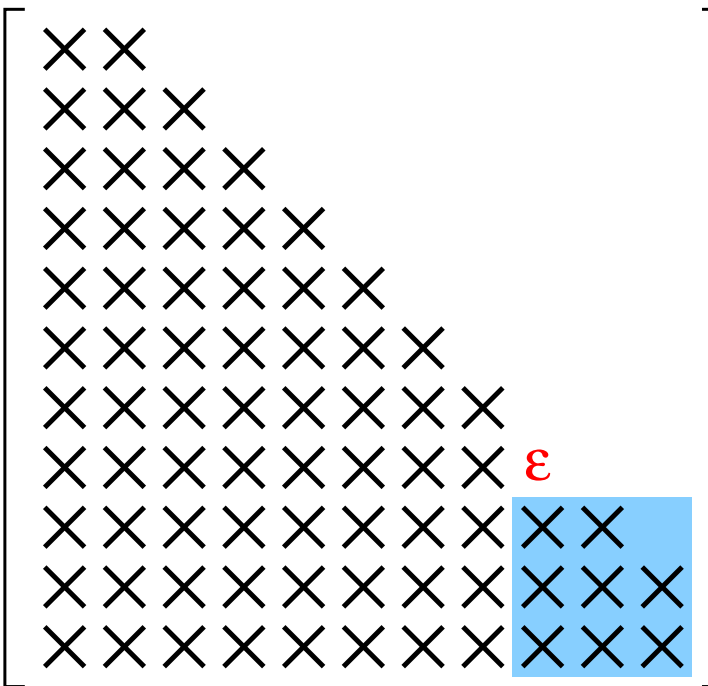
$$A\hat{Q}_i = \hat{Q}_i \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times & \times & \times \\ & \times & \times & \times & \times & \times & \times & \times & \times & \times & \times \\ & & \times & \times & \times & \times & \times & \times & \times & \times & \times \\ & & & \times & \times & \times & \times & \times & \times & \times & \times \\ & & & & \times & \times & \times & \times & \times & \times & \times \\ & & & & & \times & \times & \times & \times & \times & \times \\ & & & & & & \times & \times & \times & \times & \times \\ & & & & & & & \times & \times & \times & \times \\ & & & & & & & & \times & \times & \times \\ & & & & & & & & & \times & \times \\ & & & & & & & & & & \times \\ & & & & & & & & & & & \times \\ & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & & & & & \times \\ & & & & & & & & & & & & & & & & & & \times \end{bmatrix}$$

ε

Observed: If shifts are eigenvalues of **bottom right** $m \times m$ **submatrix**:
 $\varepsilon \rightarrow 0$ locally quadratically.

A closer look at the Hessenberg matrix

After i iterations ...


$$A^T \hat{Q}_i = \hat{Q}_i$$


For notational convenience: transpose

A closer look at the Hessenberg matrix

After i iterations ...

$$A^T \hat{Q}_i F = \hat{Q}_i F$$



For notational convenience: transpose and flip rows+columns .

A closer look at the Hessenberg matrix

After i iterations ...

$$A^T \hat{Q}_i F = \hat{Q}_i F$$

The diagram shows a matrix structure with 'x' marks representing non-zero entries. The top-left 3x3 submatrix is highlighted in blue. A red ϵ symbol is placed below the first column of the matrix, indicating a small value in the sub-diagonal element.

For notational convenience: transpose and flip rows+columns .

Set

$$\hat{Q}_i F = \left[u_1^{(i)}, u_2^{(i)}, \dots, u_n^{(i)} \right].$$

Connection to Krylov subspaces

Then

$$A^T \begin{bmatrix} u_1^{(i)} \\ \dots \\ u_m^{(i)} \end{bmatrix} = \begin{bmatrix} u_1^{(i)} \\ \dots \\ u_m^{(i)} \end{bmatrix} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix} + u_{m+1}^{(i)} \begin{bmatrix} 0 & 0 & \epsilon \end{bmatrix}$$

Connection to Krylov subspaces

Then

$$A^T \begin{bmatrix} u_1^{(i)} \\ \vdots \\ u_m^{(i)} \end{bmatrix} = \begin{bmatrix} u_1^{(i)} \\ \vdots \\ u_m^{(i)} \end{bmatrix} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ & \times & \times \end{bmatrix} + u_{m+1}^{(i)} \begin{bmatrix} 0 & 0 & \varepsilon \end{bmatrix}$$

This is an Arnoldi decomposition and reveals

$$\text{span} \left\{ u_1^{(i)}, \dots, u_m^{(i)} \right\} = \text{span} \left\{ u_1^{(i)}, A^T u_1^{(i)}, \dots, (A^T)^{m-1} u_1^{(i)} \right\} := \mathcal{X}_i.$$

Moreover, $\varepsilon \rightarrow 0$ iff \mathcal{X}_i converges to an invariant subspace \mathcal{X} of A^T .

Formal convergence result [Watkins/Elsner'91]

Let cols of X form orthonormal basis for \mathcal{X} ,
 X_{\perp} form orthonormal basis for \mathcal{X}^{\perp} .

Corresponding block Schur decomposition:

$$A [X_{\perp}, X] = [X_{\perp}, X] \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}.$$

If QR algorithm converges then

$$d(\mathcal{X}, \mathcal{X}_i) \leq C \|p_i(A_{11})^{-1}\| \|p_i(A_{22})\| d(\mathcal{X}, \mathcal{X}_{i-1}),$$

where p_i is the shift polynomial used in iteration i .

Formal convergence result [Watkins/Elsner'91]

Let cols of X form orthonormal basis for \mathcal{X} ,
 X_{\perp} form orthonormal basis for \mathcal{X}^{\perp} .

Corresponding block Schur decomposition:

$$A [X_{\perp}, X] = [X_{\perp}, X] \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}.$$

If QR algorithm converges then

$$d(\mathcal{X}, \mathcal{X}_i) \leq C \|p_i(A_{11})^{-1}\| \|p_i(A_{22})\| d(\mathcal{X}, \mathcal{X}_{i-1}),$$

where p_i is the shift polynomial used in iteration i .

Chosen shifts converge to eigenvalues of A_{22}

$$\rightsquigarrow \|p_i(A_{22})\| \leq \tilde{C} d(\mathcal{X}, \mathcal{X}_{i-1}) \rightsquigarrow d(\mathcal{X}, \mathcal{X}_i) \leq \hat{C} d(\mathcal{X}, \mathcal{X}_{i-1})^2.$$

Intermediate summary

- QR generates sequence of m -dimensional Krylov subspaces \mathcal{X}_i converging locally quadratically to invariant subspace \mathcal{X} of A^T .
- Convergence bound determined by maximal distance between *all* shifts and the eigenvalues belonging to \mathcal{X} .

Intermediate summary


- QR generates sequence of m -dimensional Krylov subspaces \mathcal{X}_i converging locally quadratically to invariant subspace \mathcal{X} of A^T .
- Convergence bound determined by maximal distance between *all* shifts and the eigenvalues belonging to \mathcal{X} .

Questions:

- Does \mathcal{X}_i contain much better approximations to individual eigenvectors?
- Does it make sense to consider Krylov subspaces larger than m ?


Example ($n = 250, m = 4$)

Convergence of individual Ritz vectors x_1, \dots, x_4 from \mathcal{X}_i to eigenvectors vs. ε :



Example ($n = 250, m = 4$)


Convergence of individual Ritz vectors x_1, \dots, x_4 from \mathcal{X}_i to eigenvectors vs. ε :



- Some Ritz vectors converge much faster than others.
- ε determined by the poorest Ritz vector approx.

Example ($n = 250, m = 4$)

Convergence of **individual Ritz vectors** x_1, \dots, x_4 from \mathcal{X}_i to eigenvectors vs. ε :



- **Classical deflation** (LAPACK 3.0) based on ε ; cannot benefit from faster convergence of individual Ritz vectors.

Convergence result for individual eigenvectors

Let v_1, \dots, v_m be left eigenvectors of A belonging to eigenvalues $\lambda_1, \dots, \lambda_m$ contained in A_{22} .

[K.'06]: For each v_j ,

$$d(v_j, \mathcal{X}_i) \leq C \|p_i(A_{11})^{-1}\| |p_i(\lambda_j)| d(v_j, \mathcal{X}_{i-1}).$$

- $|p_i(\lambda_j)|$ can be much smaller than $\|p_i(A_{22})\|$;
- only one shift needs to converge in order to obtain rapid convergence.

The Krylov-Schur algorithm [Stewart'01]


$$A^T \begin{bmatrix} u_1^{(i)} \\ \vdots \\ u_m^{(i)} \end{bmatrix} = \begin{bmatrix} u_1^{(i)} \\ \vdots \\ u_m^{(i)} \end{bmatrix} \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & & \times \end{bmatrix} + u_{m+1}^{(i)} \begin{bmatrix} 0 & 0 & \epsilon \end{bmatrix}$$

Compute Schur decomposition of Hessenberg factor:


$$A^T \begin{bmatrix} \tilde{u}_1^{(i)} \\ \vdots \\ \tilde{u}_m^{(i)} \end{bmatrix} = \begin{bmatrix} \tilde{u}_1^{(i)} \\ \vdots \\ \tilde{u}_m^{(i)} \end{bmatrix} \begin{bmatrix} \times & \times & \times \\ & \times & \times \\ & & \times \end{bmatrix} + u_{m+1}^{(i)} \begin{bmatrix} \epsilon_1 & \epsilon_2 & \epsilon_3 \end{bmatrix}$$

- If $|\epsilon_1| \leq \mathbf{u} \|A\|_F$, converged Ritz vector $\tilde{u}_1^{(i)}$ is deflated by setting $\epsilon_1 = 0$.
- If $|\epsilon_1| > \mathbf{u} \|A\|_F$, test other Ritz vectors by reordering Schur form.

The Krylov-Schur algorithm on A




The Krylov-Schur algorithm on A



Re-interpreting Krylov-Schur algorithm for \mathcal{X}_i in terms of orthogonal operations on A reveals theoretical and numerical equivalence to aggressive early deflation [Braman/Byers/Mathias'02] as implemented in LAPACK 3.1.


Example ($n = 250, m = 4$)

Convergence of **individual Ritz vectors** x_1, \dots, x_4 from $\text{span}\{u_1^{(i)}, \dots, u_4^{(i)}\}$ to eigenvectors vs. ε :



Example ($n = 250, m = 4$)

Convergence of **individual Ritz vectors** x_1, \dots, x_8 from Krylov subspace $\text{span}\{u_1^{(i)}, \dots, u_8^{(i)}\}$ to eigenvectors vs. ε :



Larger \mathcal{X}_i add Krylov subspace acceleration.

Convergence with Krylov subspace acceleration

Let v_1, \dots, v_m be left eigenvectors of A belonging to eigenvalues $\lambda_1, \dots, \lambda_m$ contained in A_{22} . Let $w \geq m$ be dim of Krylov subspace

$$\mathcal{X}_i = \text{span}\{u_1^{(i)}, \dots, u_w^{(i)}\}.$$

Then [K'06]:

$$d(v_j, \mathcal{X}_i) \leq C \|p_i(A_{11})^{-1}\| |p_i(\lambda_j)| \inf_{\phi \in \mathcal{P}_{w-m}} \frac{\|\phi(A_{11})\|}{|\phi(\lambda_j)|} d(v_j, \mathcal{X}_{i-1}).$$

Convergence with Krylov subspace acceleration

Let v_1, \dots, v_m be left eigenvectors of A belonging to eigenvalues $\lambda_1, \dots, \lambda_m$ contained in A_{22} . Let $w \geq m$ be dim of Krylov subspace

$$\mathcal{X}_i = \text{span}\{u_1^{(i)}, \dots, u_w^{(i)}\}.$$

Then [K'06]:

$$d(v_j, \mathcal{X}_i) \leq C \|p_i(A_{11})^{-1}\| |p_i(\lambda_j)| \inf_{\phi \in \mathcal{P}_{w-m}} \frac{\|\phi(A_{11})\|}{|\phi(\lambda_j)|} d(v_j, \mathcal{X}_{i-1}).$$

- Extra term $\inf_{\phi \in \mathcal{P}_{w-m}} \|\phi(A_{11})\|/|\phi(\lambda_j)|$ can be expected to become very small as $w \gg m$.
- Need extra assumptions on eigenvalue distribution of A_{11} to obtain quantitative results using, e.g., elementary potential theory [Beattie/Embree/Rossi'04].

Aggressive early deflation in practice

Choice of m and w critical for obtaining good performance.

Current setting in LAPACK's QR and QZ based on experiments with random matrices, e.g.,

$$500 \leq n \leq 3000 : \quad m = 64, \quad w = 96.$$

Usually not optimal for practical examples.

Aggressive early deflation in practice

Choice of m and w critical for obtaining good performance.

Current setting in LAPACK's QR and QZ based on experiments with random matrices, e.g.,

$$500 \leq n \leq 3000 : \quad m = 64, \quad w = 96.$$

Usually not optimal for practical examples.

Ongoing work

- develop heuristics for dynamical choice of w based on convergence bounds;
- replace ScaLAPACK's QR by two-level recursive aggressive early deflation approach.

Conclusions

- Aggressive early deflation = extraction of Ritz vectors from Krylov subspaces associated with the QR algorithm.
- Can be turned into improved convergence bounds and intuitive explanations why aggressive early deflation works so well.

References:

C. A. Beattie, M. Embree, and J. Rossi. Convergence of restarted Krylov subspaces to invariant subspaces. *SIAM J. Matrix Anal. Appl.*, 25(4):1074–1109, 2004.

K. Braman, R. Byers, and R. Mathias. The multishift QR algorithm. II. Aggressive early deflation. *SIAM J. Matrix Anal. Appl.*, 23(4):948–973, 2002.

B. Kågström and D. Kressner. Multishift variants of the QZ algorithm with aggressive early deflation. *SIAM J. Matrix Anal. Appl.*, 29(1):199-227, 2006.

D. Kressner. The effect of aggressive early deflation on the convergence of the QR algorithm. Uminf report, Department of Computing Science, Umeå University, 2006.