MANCHESTER 1824

The University of Manchester

NLA Group Manchester

Computing accurate eigenvalues and singular values using a mixed-precision (one-sided) Jacobi algorithm.

Zhengbo Zhou
Department of Mathematics
The University of Manchester

zhengbo.zhou@student.manchester.ac.uk

Scientific Computation Seminar, University of Leeds
March 17, 2026

Joint work with Nick Higham, Françoise Tisseur and Marcus Webb.

Symmetric Eigenvalue Problem

$$Ax = \lambda x, \qquad A \in \mathbb{R}^{n\times n}, x \in \mathbb{R}^n, \lambda \in \mathbb{R}.$$

Eigenvalue Problem

Find a factorization $A = Q\Lambda Q^T$:

- $Q \in \mathbb{R}^{n\times n}$ orthogonal, and
- $\Lambda \in \mathbb{R}^{n\times n}$ is diagonal..

Eigensolvers:

- Tridiagonalization based methods. E.g. QR algorithm, D&C algorithm, and RRR algorithm.
- Jacobi algorithm. (**our focus**)

Jacobi Algorithm

Jacobi algorithm (Jacobi, 1846):

- Oldest way for computing all eigenvalues of a **symmetric** matrix.

Jacobi Algorithm

Jacobi algorithm (Jacobi, 1846):

- Oldest way for computing all eigenvalues of a **symmetric** matrix.
- 1. Starting with a symmetric matrix $A^{(0)} = A \in \mathbb{R}^{n \times n}$,
  2. Generating a sequence of symmetric matrices $A^{(k+1)} = (V^{(k)})^T A^{(k)} V^{(k)}$, where $V^{(k)}$ is a plane rotation that differs from the identity only at four chosen positions (p_k, p_k), (p_k, q_k), (q_k, p_k), and (q_k, q_k).
  3. $V^{(k)}$ is specifically chosen such that positions (p_k, q_k) and (q_k, p_k) of $A^{(k+1)}$ are zero.

Jacobi Algorithm

Jacobi algorithm (Jacobi, 1846):

- Oldest way for computing all eigenvalues of a **symmetric** matrix.
- 1. Starting with a symmetric matrix $A^{(0)} = A \in \mathbb{R}^{n\times n}$,
  2. Generating a sequence of symmetric matrices $A^{(k+1)} = (V^{(k)})^T A^{(k)} V^{(k)}$, where $V^{(k)}$ is a plane rotation that differs from the identity only at four chosen positions (p_k, p_k), (p_k, q_k), (q_k, p_k), and (q_k, q_k).
  3. $V^{(k)}$ is specifically chosen such that positions (p_k, q_k) and (q_k, p_k) of $A^{(k+1)}$ are zero.
- Define $\mathrm{off}(A) = \|A - \mathrm{diag}(A)\|_F$,

$$\mathrm{off}(A^{(k+1)})^2 = \mathrm{off}(A^{(k)})^2 - 2(a^{(k)}_{p_k,q_k})^2.$$

Convergence of Jacobi algorithm

Cyclic Jacobi Algorithm

If the pivots $(p_k, q_k)_{i=1}^{\infty}$ are chosen to be

$$\underbrace{(1,2),(1,3),\ldots,(1,n),(2,3),\ldots,(2,n),\ldots,(n-1,n)}_{\text{one sweep}},(1,2),\ldots$$

then $A^{(\infty)} = \Lambda$.

Convergence of Jacobi algorithm

Cyclic Jacobi Algorithm

If the pivots $(p_k, q_k)_{i=1}^{\infty}$ are chosen to be

$$\underbrace{(1,2),(1,3),\ldots,(1,n),(2,3),\ldots,(2,n),\ldots,(n-1,n)}_{\text{one sweep}},(1,2),\ldots$$

then $A^{(\infty)} = \Lambda$.

$$\begin{bmatrix} \times & \times & \times & \times \\ & \times & \times & \times \\ & & \times & \times \\ & & & \times \end{bmatrix}$$

Convergence of Jacobi algorithm

Cyclic Jacobi Algorithm

If the pivots $(p_k, q_k)_{i=1}^{\infty}$ are chosen to be

$$\underbrace{(1,2),(1,3),\ldots,(1,n),(2,3),\ldots,(2,n),\ldots,(n-1,n)}_{\text{one sweep}},(1,2),\ldots$$

then $A^{(\infty)} = \Lambda$.

Theorem ([Hari, 1991])

Let $N = n(n-1)/2$, which is the number of rotation in one sweep. If $\mathrm{off}(A^{(k)}) \leq \delta/3$, where $\delta = \min_{\lambda_i \neq \lambda_j} |\lambda_i - \lambda_j|$, then

$$\mathrm{off}(A^{(k+N)}) \leq c\,\mathrm{off}(A^{(k)})^2. \quad \Leftarrow \textit{Quadratic convergence!}$$

Small Example

```
A = gallery("randsvd", 500, -1e5, 3);
```

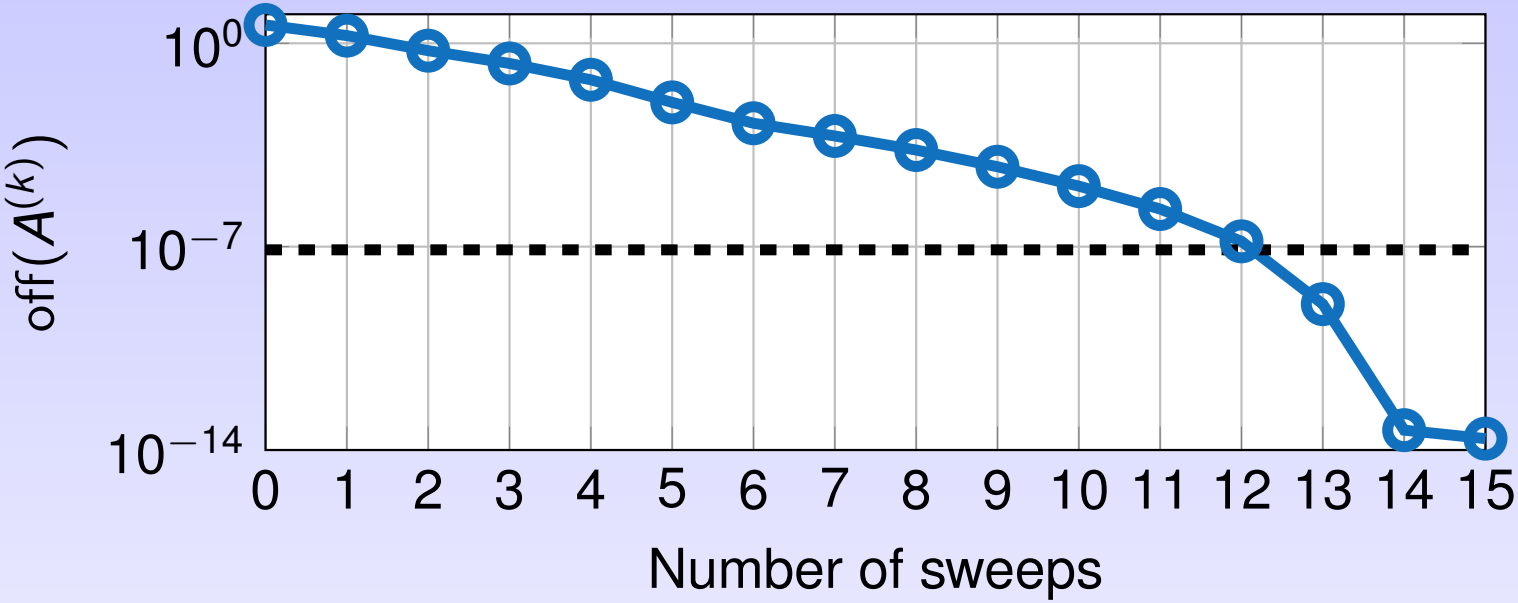

Small Example

```
A = gallery("randsvd", 500, -1e5, 3);
```

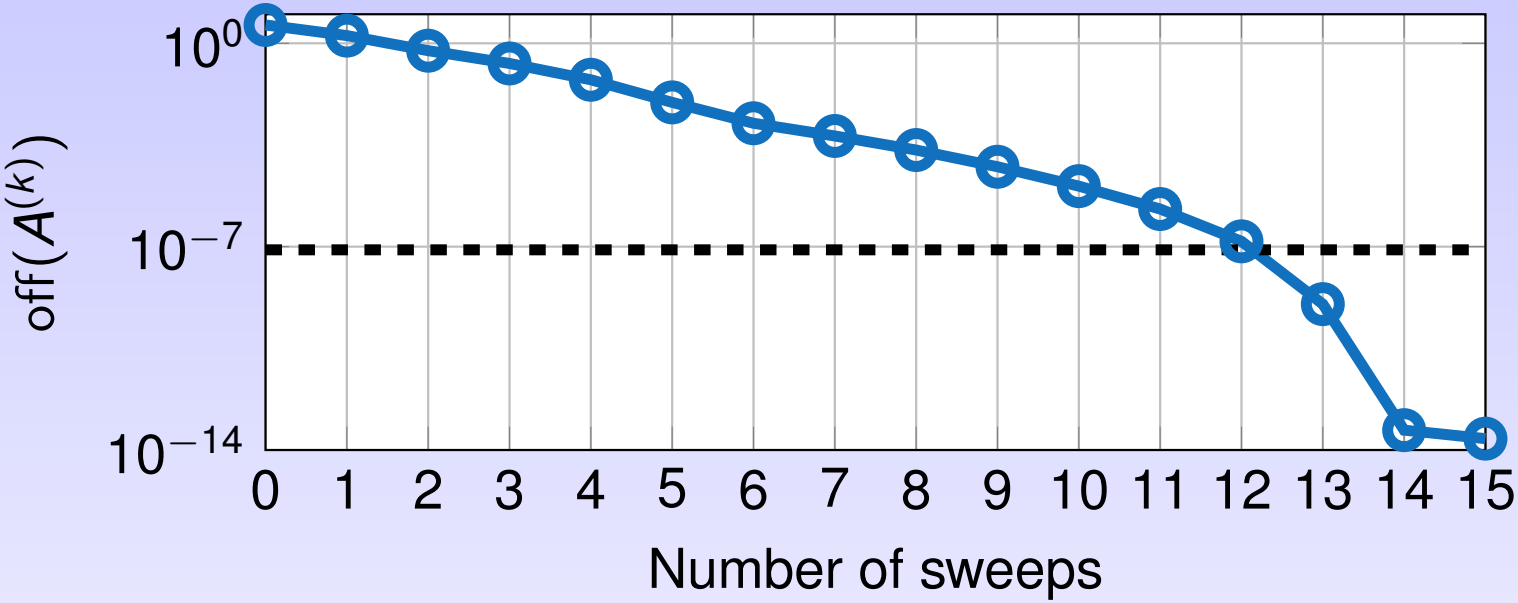

The cyclic methods waste time annihilating small elements in the early sweeps. — Parlett (1998)

Small Example

```
A = gallery("randsvd", 500, -1e5, 3);
```

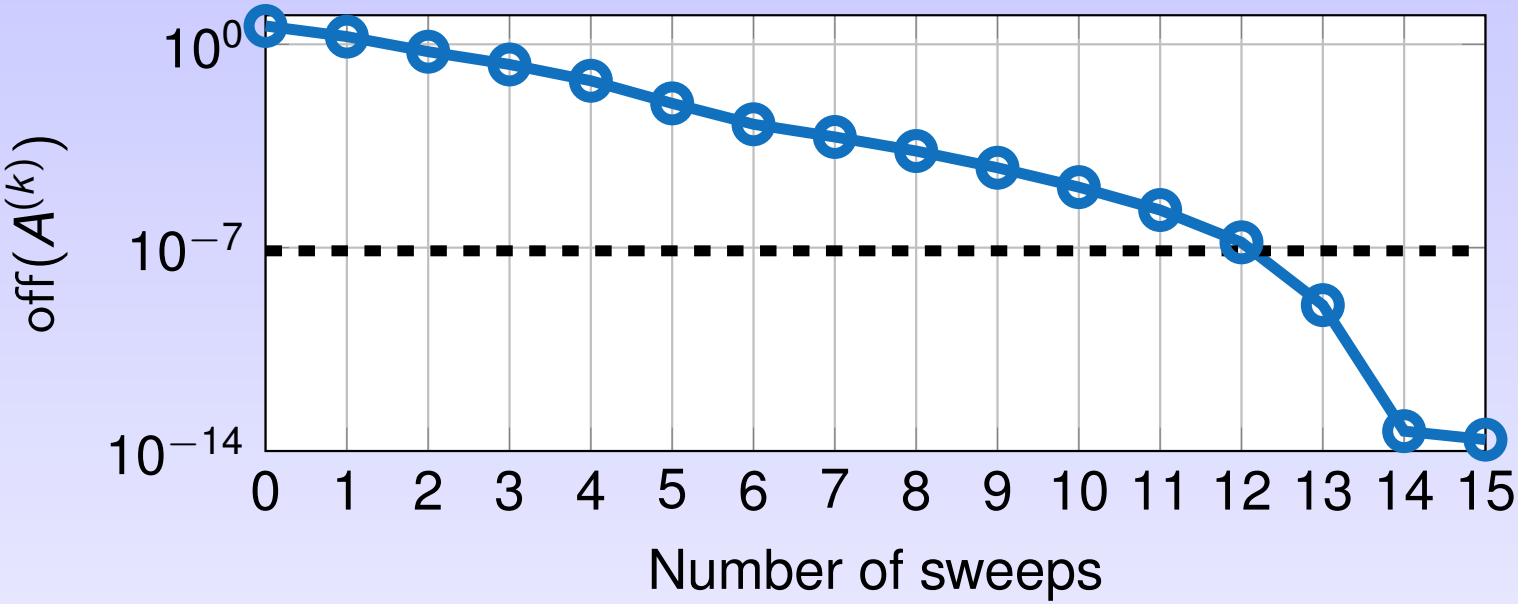

The cyclic methods waste time annihilating small elements in the early sweeps. — Parlett (1998)

Question: Can we skip these early sweeps?

Preconditioning the Jacobi algorithm

Requirements

- Preconditioning $A \mapsto \widetilde{A}$ such that $\mathrm{off}(\widetilde{A})$ is small.
- Preconditioning should not change the eigenvalues more than $O(u)$.
- Make use of low precision u_ℓ $(< u)$ to speed up the construction of preconditioner.

Preconditioning the Jacobi algorithm

Requirements

- Preconditioning $A \mapsto \widetilde{A}$ such that $\mathrm{off}(\widetilde{A})$ is small.
- Preconditioning should not change the eigenvalues more than $O(u)$.
- Make use of low precision u_ℓ $(< u)$ to speed up the construction of preconditioner.

Idea

- Let QDQ^T be an spectral decomposition of A computed at precision u_ℓ, then $\|Q^TAQ - D\|_F = O(u_\ell)$.
- $$\begin{aligned}\mathrm{off}(Q^TAQ) &\overset{\text{def}}{=} \|(Q^TAQ) - \mathrm{diag}(Q^TAQ)\|_F \\ &\leq \|Q^TAQ - D\|_F = O(u_\ell)\end{aligned}$$
- Q need to be numerically **orthogonal** at u!

Algorithms

Approach 1: Orthogonalization method

[Zhou, 2022], **[Zhang and Bai, 2022]**

1. Compute an eigenvector matrix Q_ℓ. (u_ℓ)
2. Orthogonalize Q_ℓ to $\widetilde{Q}$. (u)

Algorithms

Approach 1: Orthogonalization method

[Zhou, 2022], [Zhang and Bai, 2022]

1. Compute an eigenvector matrix Q_ℓ. (u_ℓ)
2. Orthogonalize Q_ℓ to $\widetilde{Q}$. (u)

Approach 2: Modified tridiagonalization method

[Higham et al., 2025]

1. Perform Tridiag at u_ℓ. Store Householder vectors and construct transformation matrix at u. T_ℓ, Q_T
2. Apply any eigensolver to T_ℓ at u. Q_S
3. Obtain $\widetilde{Q} = Q_T Q_S$ at u.

How small is $\mathrm{off}(\widetilde{Q}^T A \widetilde{Q})$?

Theorem

Let $\widetilde{Q}$ be generated using

- *Approach 1 with MGS,*
- *Approach 1 with HHQR,*
- *Approach 1 with Newton–Schulz iteration, or*
- *Approach 2,*

then

$$\mathrm{off}(\widetilde{Q}^T A \widetilde{Q})/\|A\|_F \leq p(n) u_\ell,$$

$p(n)$ is a low degree polynomial.

[Zhang and Bai, 2022] proved the first one,
[Higham et al., 2025] proved the rest.

Example

$A \in \mathbb{R}^{n\times n}$, $\kappa(A) = 10^6$, geometrically distributed evals.

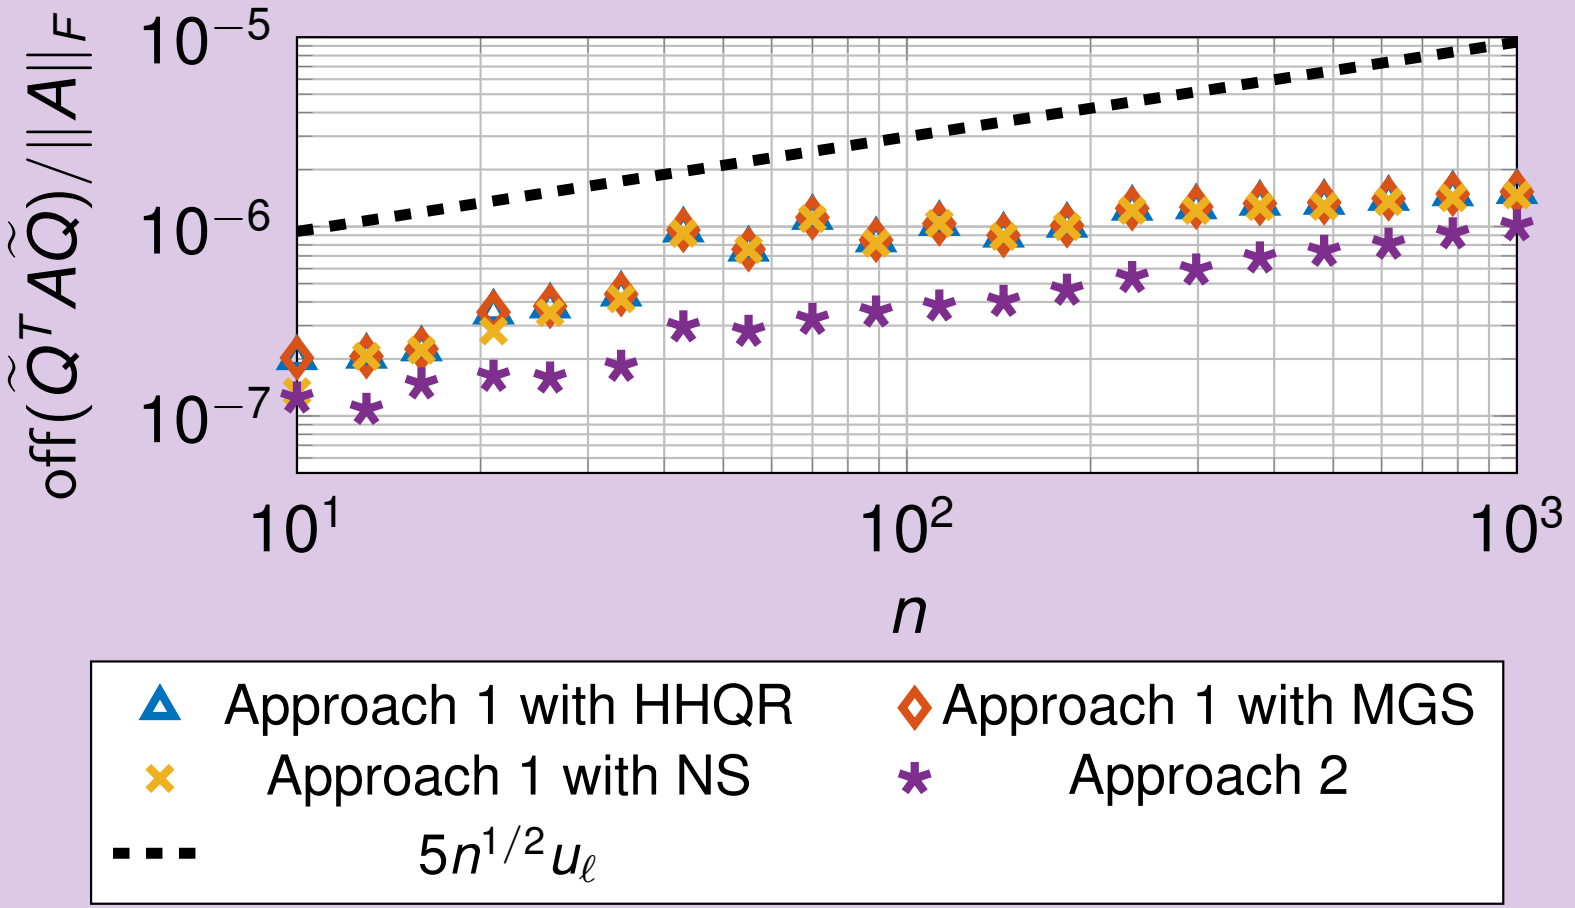

Comparison example

- `A = gallery("randsvd", 500, -1e5, 3);`
- Approach 1 + HHQR.

```
[Ql, ] = eig(single(A));
[Qt, ] = qr(double(Ql));
At = Qt'*A*Qt;
```

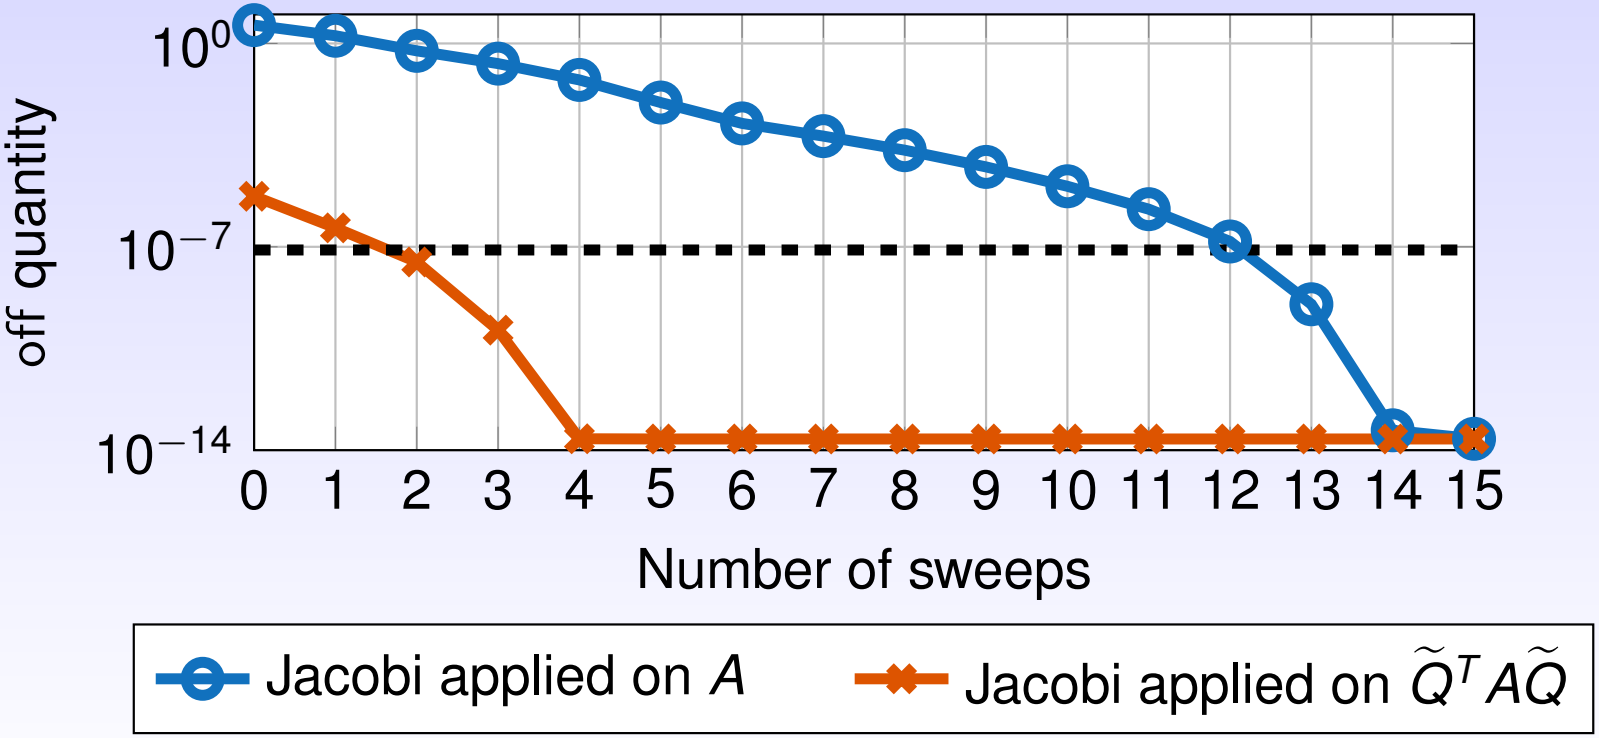

High Accuracy of the Jacobi Algorithm

Interested in **symmetric positive definite problem**

High Accuracy of the Jacobi Algorithm

Interested in **symmetric positive definite problem**

Tridiagonalization based methods

$$\frac{|\lambda_i(A) - \widehat{\lambda}_i(A)|}{\lambda_i(A)} \leq p(n)\, u\, \kappa(A).$$

- $p(n)$ = low deg. poly., u = working precision.
- $\kappa(A)$ = 2 norm condition number of A.
- $\lambda_i(A), \widehat{\lambda}_i(A)$ = ith largest exact, computed eigenvalue.

High Accuracy of the Jacobi Algorithm

Interested in **symmetric positive definite problem**

Tridiagonalization based methods

$$\frac{|\lambda_i(A) - \widehat{\lambda}_i(A)|}{\lambda_i(A)} \leq p(n)\, u\, \kappa(A).$$

Jacobi algorithm is "more accurate" for SPD matrix

[Demmel and Veselić, 1992]

By implementing a specific stopping criterion,

$$\frac{|\lambda_i(A) - \widehat{\lambda}_i(A)|}{\lambda_i(A)} \leq p(n)\, u\, \kappa_S(A)$$

$$\kappa_S(A) = \kappa(DAD),\ D = \mathrm{diag}(a_{ii}^{-1/2})$$

Example by Demmel & Veselić

$$A = \begin{bmatrix} 10^{40} & 10^{29} & 10^{19} \\ 10^{29} & 10^{20} & 10^{9} \\ 10^{19} & 10^{9} & 1 \end{bmatrix}, \qquad \kappa(A) \approx 10^{40}$$

Use $D = \mathrm{diag}(10^{-20}, 10^{-10}, 1)$,

$$DAD = \begin{bmatrix} 1 & 0.1 & 0.1 \\ 0.1 & 1 & 0.1 \\ 0.1 & 0.1 & 1 \end{bmatrix}, \qquad \kappa_S(A) = \kappa(DAD) \approx 1.33.$$

	Jacobi	`eig`
$\lambda_1 = 10^{40}$	10^{40}	10^{40}
$\lambda_2 = 9.9 \times 10^{19}$	9.9×10^{19}	9.9×10^{-1}
$\lambda_3 = 9.82 \times 10^{-1}$	9.82×10^{-1}	-1.93×10^{23}

High Precision Preconditioning

Even more: *Compute* $\widetilde{Q}^T A \widetilde{Q}$ *at* u_h *instead of* u.

High Precision Preconditioning

Even more: *Compute* $\widetilde{Q}^T A \widetilde{Q}$ *at* u_h *instead of* u.

Proposed algorithm: MP3Jacobi

Given a SPD $A \in \mathbb{R}^{n \times n}$.

- Construct preconditioner $\widetilde{Q}$. (u_ℓ, u)
- Obtain preconditioned matrix $\widetilde{A} = \widetilde{Q}^T A \widetilde{Q} \rightsquigarrow \widetilde{A}_{\text{comp}}$. (u_h)
- Apply Jacobi to $\widetilde{A}_{\text{comp}}$. (u)

High Precision Preconditioning

Even more: *Compute $\widetilde{Q}^T A \widetilde{Q}$ at u_h instead of u.*

Proposed algorithm: MP3Jacobi

Given a SPD $A \in \mathbb{R}^{n \times n}$.

- Construct preconditioner $\widetilde{Q}$. (u_ℓ, u)
- Obtain preconditioned matrix $\widetilde{A} = \widetilde{Q}^T A \widetilde{Q} \rightsquigarrow \widetilde{A}_{\text{comp}}$. (u_h)
- Apply Jacobi to $\widetilde{A}_{\text{comp}}$. (u)

Theorem

[Higham et al., 2025]

$$\frac{|\widehat{\lambda}_i(\widetilde{A}_{\text{comp}}) - \lambda_i(A)|}{\lambda_i(A)} \leq p(n)\, u\, \kappa_S(\widetilde{A}).$$

Insights on the size of $\kappa_S(\widetilde{A})$

Why $\kappa_S(\widetilde{A})$?

$\kappa_S(\widetilde{A})$ can be **significantly smaller** than $\kappa_S(A)$ and $\kappa(A)$.

Insights on the size of $\kappa_S(\widetilde{A})$

Why $\kappa_S(\widetilde{A})$?

$\kappa_S(\widetilde{A})$ can be **significantly smaller** than $\kappa_S(A)$ and $\kappa(A)$.

Define the **reverse Hadamard measure**:

$$\mathcal{H}(A) = \prod_i \{a_{ii}/\lambda_i(A)\},$$

then $\kappa_S(A) \leq n \cdot \mathrm{e} \cdot \mathcal{H}(A)$. **[Demmel and Veselić, 1992]**

Insights on the size of $\kappa_S(\widetilde{A})$

Why $\kappa_S(\widetilde{A})$?

$\kappa_S(\widetilde{A})$ can be **significantly smaller** than $\kappa_S(A)$ and $\kappa(A)$.

Define the **reverse Hadamard measure**:

$$\mathcal{H}(A) = \prod_i \{a_{ii}/\lambda_i(A)\},$$

then $\kappa_S(A) \leq n \cdot \mathrm{e} \cdot \mathcal{H}(A)$. **[Demmel and Veselić, 1992]**

- **Observation:** Diagonals of $\widetilde{A}$ can act as **approximate eigenvalues at precision** u_ℓ.
- $\mathcal{H}(\widetilde{A})$ close to 1; $\kappa_S(\widetilde{A}) \approx n \cdot \mathrm{e}$.
- $\mathcal{H}(A)$ is arbitrary; $\kappa_S(A)$ can be as large as possible.

How small $\kappa_S(\widetilde{A})$ can be?

If $\mathrm{off}(\widetilde{A})$ is small such that

$$\mathrm{off}(\widetilde{A}) < \frac{1}{2} \min_i \widetilde{a}_{ii},$$

then $\kappa_S(\widetilde{A}) \approx 1$.

How small $\kappa_S(\widetilde{A})$ can be?

If $\mathrm{off}(\widetilde{A})$ is small such that

$$\mathrm{off}(\widetilde{A}) < \frac{1}{2} \min_i \widetilde{a}_{ii},$$

then $\kappa_S(\widetilde{A}) \approx 1$.

Special matrix examples

Matrix type	$\kappa(A)$	$\kappa_S(A)$	$\kappa_S(\widetilde{A})$
`hilb(7)`	5e8	2e8	3
`pascal(15)`	3e15	6e12	1e4

Why u_h?

Decompose the relative forward error

$$\frac{|\widehat{\lambda}_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A}_{\text{comp}})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}) - \lambda_i(A)|}{\lambda_i(A)}$$

Why u_h?

Decompose the relative forward error

$$\frac{|\widehat{\lambda}_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A}_{\text{comp}})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}) - \lambda_i(A)|}{\lambda_i(A)}$$

First term:

- Relative forward error analysis of the Jacobi algorithm.
  [Demmel and Veselić, 1992] [Mathias, 1995]
- Bounded by $p(n)u\kappa_S(\widetilde{A}_{\text{comp}})$.

Why u_h?

Decompose the relative forward error

$$\frac{|\widehat{\lambda}_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A}_{\text{comp}})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A})|}{\lambda_i(A)} + \frac{|\lambda_i(\widetilde{A}) - \lambda_i(A)|}{\lambda_i(A)}$$

First term:

- Relative forward error analysis of the Jacobi algorithm.
  [Demmel and Veselić, 1992] [Mathias, 1995]
- Bounded by $p(n)u\kappa_S(\widetilde{A}_{\text{comp}})$.

Third term:

- **Exact** eigenvalues & matrix–matrix multiplications.
- Using Ostrowski's theorem: $\lambda_k(\widetilde{A}) \leq \lambda_1(\widetilde{Q}^T\widetilde{Q})\lambda_k(A)$.
- Bounded by $p(n)u \Leftarrow$ comes from orthogonality of $\widetilde{Q}$.

Analyzing the Second Term I

$$\frac{|\lambda_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A})|}{\lambda_i(A)}$$

- Need relationship between $\widetilde{A}_{\text{comp}}$ (computed $\widetilde{A}$) and $\widetilde{A}$:

$$\widetilde{A}_{\text{comp}} = \widetilde{A} + \Delta\widetilde{A}$$

- $\Delta\widetilde{A}$ comes from
  1. Matrix–Matrix products at precision u_h.
  2. Rounding to u.

Lemma (Componentiwse error bound)

$$|\Delta\widetilde{a}_{ij}| \leq u|\widetilde{a}_{ij}| + p(n)u_h\left(|\widetilde{Q}^T||A||\widetilde{Q}|\right)_{ij}.$$

Analyzing the Second Term II

Theorem (Demmel & Veselić)

If $|\Delta\widetilde{a}_{ij}|/\sqrt{\widetilde{a}_{ii}\widetilde{a}_{jj}} \leq \varepsilon$, *then*

$$\frac{|\lambda_i(\widetilde{A}_{\text{comp}}) - \lambda_i(\widetilde{A})|}{\lambda_i(A)} \leq \varepsilon \cdot n \cdot \kappa_S(\widetilde{A}).$$

From previous page's bound on $|\Delta\widetilde{a}_{ij}|$,

$$\frac{|\Delta\widetilde{a}_{ij}|}{\sqrt{\widetilde{a}_{ii}\widetilde{a}_{jj}}} \leq u + p(n)u_h\alpha_{ij}, \qquad \alpha_{ij} := \frac{\left(|\widetilde{Q}^T||A||\widetilde{Q}|\right)_{ij}}{\sqrt{\widetilde{a}_{ii}\widetilde{a}_{jj}}}.$$

Reason for using high precision: α_{ij} can be as large as $\kappa(A)$! A small u_h is required to balance α_{ij}.

One-sided Jacobi algorithm

Idea. Applying one-sided Jacobi to G is equivalent as applying two-sided Jacobi to $G^T G$.

Two-sided Jacobi algorithm: Given plane rotation $V^{(k)}$ constructed using $(G^{(k)})^T G^{(k)}$ at each step

$$\ldots (V^{(3)})^T (V^{(2)})^T (V^{(1)})^T G^T \; GV^{(1)} V^{(2)} V^{(3)} \ldots$$

One-sided Jacobi algorithm

Idea. Applying one-sided Jacobi to G is equivalent as applying two-sided Jacobi to $G^T G$.
Two-sided Jacobi algorithm: Given plane rotation $V^{(k)}$ constructed using $(G^{(k)})^T G^{(k)}$ at each step

$$\ldots (V^{(3)})^T (V^{(2)})^T (V^{(1)})^T G^T \; GV^{(1)} V^{(2)} V^{(3)} \ldots$$

- Zero out (p_k, q_k) and (q_k, p_k) entries of $(G^{(k)})^T G^{(k)}$.
  One-sided Jacobi: make p_kth and q_kth columns of $G^{(k)}$ orthogonal.

One-sided Jacobi algorithm

Idea. Applying one-sided Jacobi to G is equivalent as applying two-sided Jacobi to G^TG.
Two-sided Jacobi algorithm: Given plane rotation $V^{(k)}$ constructed using $(G^{(k)})^TG^{(k)}$ at each step

$$\ldots (V^{(3)})^T (V^{(2)})^T (V^{(1)})^T G^T \; GV^{(1)}V^{(2)}V^{(3)}\ldots$$

- Zero out (p_k, q_k) and (q_k, p_k) entries of $(G^{(k)})^TG^{(k)}$.
  One-sided Jacobi: make p_kth and q_kth columns of $G^{(k)}$ orthogonal.
- $\mathrm{off}(G^TG)$ translates to orthonormality of G.
- **Preconditioner:** reduce $\mathrm{off}(G^TG)$ translates to make G more orthonormal.

Preconditioning the one-sided Jacobi

Idea.
Copying the algorithms for two-sided Jacobi!

Orthogonalization method (modified)

[Zhou et al., 2026]
Given a general matrix $G \in \mathbb{R}^{m\times n}$,

1. Compute the right singular vectors $V_\ell \in \mathbb{R}^{n\times n}$. (u_ℓ)
2. Orthogonalize V_ℓ to $\widetilde{V}$. (u)

Preconditioning. $\widetilde{G} = G\widetilde{V}$

Accurate Singular Values!

Proposed algorithm: MP3JacobiSVD

Given a general matrix $G \in \mathbb{R}^{m\times n}$.

- Construct preconditioner $\widetilde{V} \in \mathbb{R}^{n\times n}$. (u_ℓ, u)
- Obtain preconditioned matrix $\widetilde{G} = G\widetilde{V} \rightsquigarrow \widetilde{G}_{\text{comp}}$. (u_h)
- Apply one-sided Jacobi to $\widetilde{G}$. (u)

Theorem

[Zhou et al., 2026]

$$\frac{|\widehat{\sigma}_i(\widetilde{G}_{\text{comp}}) - \sigma_i(G)|}{\sigma_i(G)} \leq p(m,n)\, u\, \kappa^D(\widetilde{G}),$$

$$\kappa^D(\widetilde{G}) := \kappa(\widetilde{G}D), \quad D_{ii} = \text{diag}(\|\widetilde{G}(:,i)\|_2^{-1}).$$

How small is $\kappa^D(\widetilde{G})$?

Introduce another metric:

$$\text{obliq}(\widetilde{G}) := \|\widetilde{G}D - U\|_2,$$

- $D = \text{diag}(\|\widetilde{G}(:, i)\|_2^{-1})$.
- U is an orthogonal polar factor of $\widetilde{G}D$.
- Measure the distance between $\widetilde{G}$ and the nearest orthonormal matrix.

Remarks

- If $\text{obliq}(\widetilde{G}) < \theta$ for some $\theta < 1$, then $\kappa^D(\widetilde{G}) \leq (1 + \theta)/(1 - \theta)$.
- If u_ℓ is small enough, or the input matrix is well–conditioned such that $u_\ell \kappa(A) \ll 1$, then $\kappa^D(\widetilde{G}) \approx 1$.

Experiment I: Two-sided Jacobi

Setup:

- Random matrix $A \in \mathbb{R}^{100\times 100}$ SPD.
- $\kappa(A) = 10^8$.
- Geometrically distributed eigenvalues.
- (u_ℓ, u, u_h) = (single, double, quadruple).
- **MP2Jacobi**: $u_h = u$.

Relative forward error

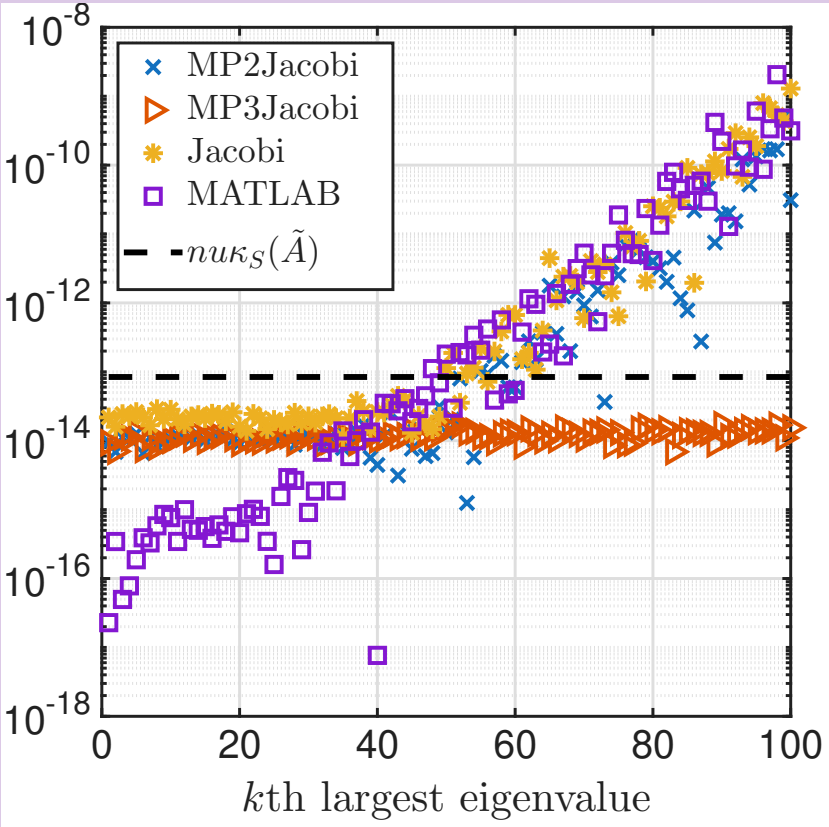

Experiment II: Two-sided Jacobi

Setup:

- Random matrix $A \in \mathbb{R}^{100\times 100}$ SPD.
- $\kappa(A) = 10^{14}$.
- Geometrically distributed eigenvalues.
- (u_ℓ, u, u_h) = (single, double, quadruple).
- **MP2Jacobi**: $u_h = u$.

Relative forward error

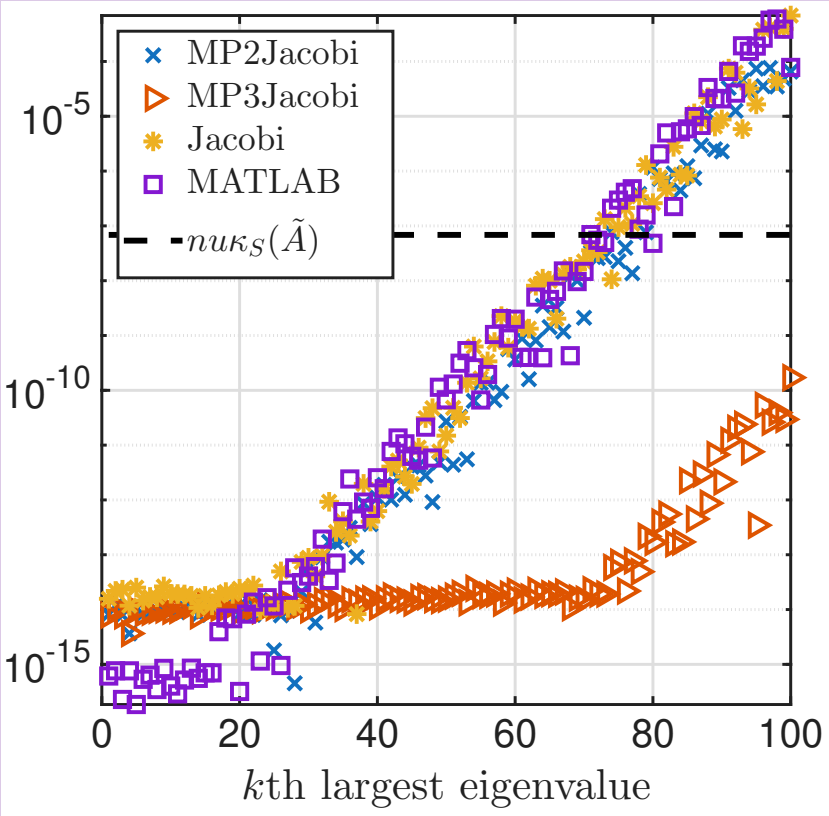

Experiment III: Our Algorithm VS LAPACK

- General random matrix $G \in \mathbb{R}^{3000\times n}$, n is varying.
- $\kappa(G) = 10^6$.
- Geometrically distributed singular values.
- (u_ℓ, u, u_h) = (single, single, double).
- **SGESVJ**: one-sided Jacobi (LAPACK);
  SGEJSV: preconditioned version (LAPACK).

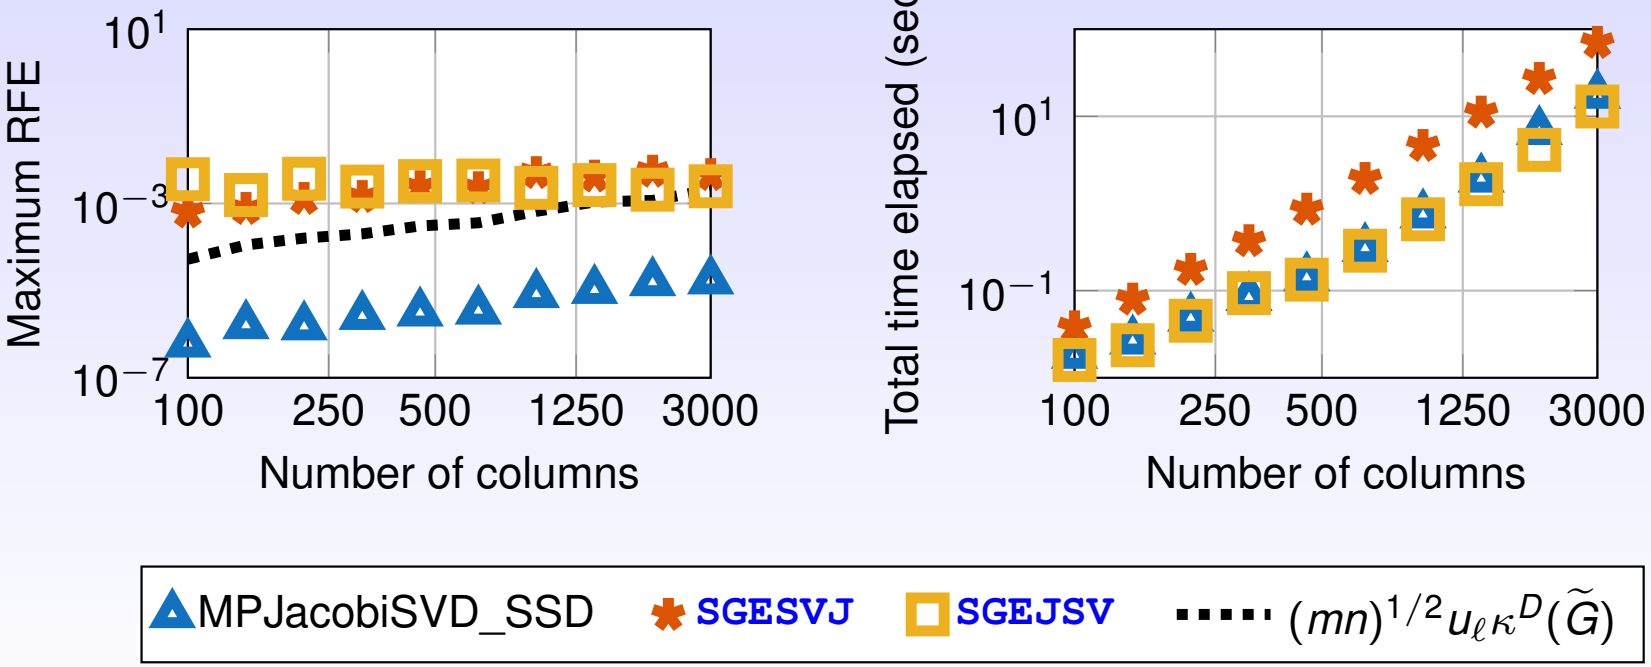

Summary I

We proposed and analyzed:

- Ways to construct preconditioner for the (one-sided) Jacobi algorithm in order to skip the slow convergence phase.
- A mixed-precision preconditioned (one-sided) Jacobi algorithm with much more accurate computed (singular values) eigenvalues. The cost will be (one) two matrix multiplications at u_h.

Summary II

Two papers online:

1. Nicholas J. Higham, Françoise Tisseur, Marcus Webb, and Zhengbo Zhou. Computing accurate eigenvalues using a mixed–precision Jacobi algorithm. *SIAM J. Matrix Anal. Appl.*, 46(4):2423–2448, 2025.
2. Zhengbo Zhou, Françoise Tisseur, and Marcus Webb. Computing accurate singular values using a mixed-precision one-sided Jacobi algorithm. arXiv:2602.18134, February 2026.

Summary II

Two papers online:

1. Nicholas J. Higham, Françoise Tisseur, Marcus Webb, and Zhengbo Zhou. Computing accurate eigenvalues using a mixed–precision Jacobi algorithm. *SIAM J. Matrix Anal. Appl.*, 46(4):2423–2448, 2025.
2. Zhengbo Zhou, Françoise Tisseur, and Marcus Webb. Computing accurate singular values using a mixed-precision one-sided Jacobi algorithm. arXiv:2602.18134, February 2026.

Thanks for your listening!

References I

James Demmel and Krešimir Veselić. Jacobi's method is more accurate than QR. *SIAM J. Matrix Anal. Appl.*, 13 (4):1204–1245, October 1992. doi: 10.1137/0613074.

Vjeran Hari. On sharp quadratic convergence bounds for the serial Jacobi methods. *Numer. Math.*, 60(1):375–406, December 1991. doi: 10.1007/BF01385728.

Nicholas J. Higham, Françoise Tisseur, Marcus Webb, and Zhengbo Zhou. Computing accurate eigenvalues using a mixed–precision Jacobi algorithm. *SIAM J. Matrix Anal. Appl.*, 46(4):2423–2448, October 2025. doi: 10.1137/25m1723748.

Roy Mathias. Accurate eigensystem computations by Jacobi methods. *SIAM J. Matrix Anal. Appl.*, 16(3): 977–1003, July 1995. doi: 10.1137/s089547989324820x.

References II

Beresford N. Parlett. *The Symmetric Eigenvalue Problem*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, January 1998. ISBN 978-0-89871-402-9. doi: 10.1137/1.9781611971163.

Zhiyuan Zhang and Zheng-Jian Bai. A mixed precision Jacobi method for the symmetric eigenvalue problem. arXiv:2211.03339v1, November 2022. URL https://arxiv.org/abs/2211.03339v1.

Zhengbo Zhou. A mixed precision eigensolver based on the Jacobi algorithm. M.Sc. thesis, The University of Manchester, Manchester, UK, September 2022. URL https://arxiv.org/abs/2509.00495.

References III

Zhengbo Zhou, Françoise Tisseur, and Marcus Webb. Computing accurate singular values using a mixed-precision one-sided jacobi algorithm. arXiv:2602.18134, February 2026. URL https://arxiv.org/abs/2602.18134.