Sparse Days 2025

2-3 June 2025

Mixed precision preconditioned GMRES

Theo Mary

Sorbonne Université, CNRS, LIP6

Alfredo Buttari, Nicholas J. Higham, T. M., and Bastien Vieublé. A modular framework for the backward error analysis of GMRES. HAL EPrint hal-04525918, 2024, to appear in IMAJNA.
https://hal.science/hal-04525918

Alfredo Buttari, Xin Liu, T. M., and Bastien Vieublé. Mixed precision strategies for preconditioned GMRES: a comprehensive analysis. HAL EPrint hal-05071696, 2025.
https://hal.science/hal-05071696

GMRES

- Krylov-based iterative solver for the solution of general square linear systems $Ax = b$
- Compute an orthonormal Krylov basis V_k through the Arnoldi iteration
- Find the vector x_k in span$\{V_k\}$ that minimizes $\|Ax_k - b\|$
- Reiterate until x_k is a satisfactory approximation of x

GMRES

Input: $A \in \mathbb{R}^{n \times n}$, $b \in \mathbb{R}^n$, $x_0 \in \mathbb{R}^n$, $\tau > 0$

$r_0 = b - Ax_0$
$\beta = \|r_0\|, \; v_1 = r_0/\beta, \; k = 1$
repeat
    $w_k = Av_k$
    for $i = 1, \dots, k$ **do**
        $h_{i,k} = v_i^T w_k$
        $w_k = w_k - h_{i,k} v_i$
    end for
    $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k / h_{k+1,k}$
    $V_k = [v_1, \dots, v_k]$
    $H_k = \{h_{i,j}\}_{1 \leq i \leq j+1; 1 \leq j \leq k}$
    $y_k = \mathrm{argmin}_y \|\beta e_1 - H_k y\|$
    $k = k + 1$
until $\|\beta e_1 - H_k y_k\| \leq \tau$
$x_k = x_0 + V_k y_k$

Output: $x_k \in \mathbb{R}^n$ such that $\|Ax_k - b\| \leq \tau$

Preconditioned GMRES

$$M^{-1}Ax = M^{-1}b$$

Left GMRES

$r_0 = b - Ax_0$
$s_0 = M^{-1}r_0$
$\beta = \|s_0\|,\ v_1 = s_0/\beta,\ k = 1$
repeat
  $z_k = Av_k$
  $w_k = M^{-1}z_k$
  for $i = 1, \ldots, k$ **do**
    $h_{i,k} = v_i^T w_k$
    $w_k = w_k - h_{i,k}v_i$
  end for
  $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k/h_{k+1,k}$
  $V_k = [v_1, \ldots, v_k]$
  $H_k = \{h_{i,j}\}_{1\le i\le j+1; 1\le j\le k}$
  $y_k = \operatorname{argmin}_y \|\beta e_1 - H_k y\|$
  $k = k + 1$
until $\|\beta e_1 - H_k y_k\| \le \tau$
$x_k = x_0 + V_k y_k$

$$AM^{-1}y = b,\ x = M^{-1}y$$

Right GMRES

$r_0 = b - Ax_0$
$\beta = \|r_0\|,\ v_1 = r_0/\beta,\ k = 1$
repeat
  $z_k = M^{-1}v_k$
  $w_k = Az_k$
  for $i = 1, \ldots, k$ **do**
    $h_{i,k} = v_i^T w_k$
    $w_k = w_k - h_{i,k}v_i$
  end for
  $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k/h_{k+1,k}$
  $V_k = [v_1, \ldots, v_k]$
  $H_k = \{h_{i,j}\}_{1\le i\le j+1; 1\le j\le k}$
  $y_k = \operatorname{argmin}_y \|\beta e_1 - H_k y\|$
  $k = k + 1$
until $\|\beta e_1 - H_k y_k\| \le \tau$
$d_k = V_k y_k$
$x_k = x_0 + M^{-1}d_k$

$$M^{-1}Ax = M^{-1}b$$

Left GMRES

$r_0 = b - Ax_0$
$s_0 = M^{-1}r_0$
$\beta = \|s_0\|,\ v_1 = s_0/\beta,\ k = 1$
repeat
    $z_k = Av_k$
    $w_k = M^{-1}z_k$
    for $i = 1, \dots, k$ **do**
        $h_{i,k} = v_i^T w_k$
        $w_k = w_k - h_{i,k}v_i$
    end for
    $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k/h_{k+1,k}$
    $V_k = [v_1, \dots, v_k]$
    $H_k = \{h_{i,j}\}_{1 \le i \le j+1; 1 \le j \le k}$
    $y_k = \text{argmin}_y \|\beta e_1 - H_k y\|$
    $k = k + 1$
until $\|\beta e_1 - H_k y_k\| \le \tau$
$x_k = x_0 + V_k y_k$

$$AM^{-1}y = b,\ x = M^{-1}y$$

Flexible GMRES

$r_0 = b - Ax_0$
$\beta = \|r_0\|,\ v_1 = r_0/\beta,\ k = 1$
repeat
    $z_k = M^{-1}v_k$
    $w_k = Az_k$
    for $i = 1, \dots, k$ **do**
        $h_{i,k} = v_i^T w_k$
        $w_k = w_k - h_{i,k}v_i$
    end for
    $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k/h_{k+1,k}$
    $V_k = [v_1, \dots, v_k],\ Z_k = [z_1, \dots, z_k]$
    $H_k = \{h_{i,j}\}_{1 \le i \le j+1; 1 \le j \le k}$
    $y_k = \text{argmin}_y \|\beta e_1 - H_k y\|$
    $k = k + 1$
until $\|\beta e_1 - H_k y_k\| \le \tau$
$x_k = x_0 + Z_k d_k$

Mixed precision

	Signif. bits	Exp. bits	Range ($f_{\max}/f_{\min}$)	Unit roundoff u
fp128	113	15	$2^{32766} \approx 10^{9863}$	$2^{-114} \approx 1 \times 10^{-34}$
fp64	52	11	$2^{2046} \approx 10^{308}$	$2^{-53} \approx 1 \times 10^{-16}$
fp32	23	8	$2^{254} \approx 10^{76}$	$2^{-24} \approx 6 \times 10^{-8}$
tfloat32	10	8	$2^{254} \approx 10^{76}$	$2^{-11} \approx 5 \times 10^{-4}$
fp16	10	5	$2^{30} \approx 10^{9}$	$2^{-11} \approx 5 \times 10^{-4}$
bfloat16	7	8	$2^{254} \approx 10^{76}$	$2^{-8} \approx 4 \times 10^{-3}$
fp8 (E4M3)	3	4	$2^{15} \approx 3 \times 10^{4}$	$2^{-4} \approx 6 \times 10^{-2}$
fp8 (E5M2)	2	5	$2^{30} \approx 10^{9}$	$2^{-3} \approx 1 \times 10^{-1}$
fp6 (E2M3)	3	2	$2^{3} \approx 8$	$2^{-4} \approx 6 \times 10^{-2}$
fp6 (E3M2)	2	3	$2^{7} \approx 128$	$2^{-3} \approx 0.125$
fp4 (E2M1)	1	2	$2^{3} \approx 8$	$2^{-2} \approx 0.25$

Lower precisions:

- ☺ Faster, consume less memory and energy
- ☹ Lower accuracy and narrower range
- $\Rightarrow$ Mixed precision algorithms

N. J. Higham and T. M. Mixed precision algorithms in numerical linear algebra. *Acta Numerica* 2022. https://www.doi.org/10.1017/S0962492922000022

Mixed precision preconditioned GMRES

Left GMRES

$r_0 = b - Ax_0$	u_a
$s_0 = M^{-1} r_0$	u_m
$\beta = \lVert s_0 \rVert, \ v_1 = s_0/\beta, \ k = 1$	u_g
repeat	
$z_k = A v_k$	u_a
$w_k = M^{-1} z_k$	u_m
for $i = 1, \ldots, k$ **do**	
$h_{i,k} = v_i^T w_k$	u_g
$w_k = w_k - h_{i,k} v_i$	u_g
end for	
$h_{k+1,k} = \lVert w_k \rVert, v_{k+1} = w_k / h_{k+1,k}$	u_g
$V_k = [v_1, \ldots, v_k]$	
$H_k = \{h_{i,j}\}_{1 < i < j+1; 1 < j < k}$	
$y_k = \mathrm{argmin}_y \lVert \beta e_1 - H_k y \rVert$	u_g
$k = k + 1$	
until $\lVert \beta e_1 - H_k y_k \rVert \leq \tau$	
$x_k = x_0 + V_k y_k$	u_g

Right GMRES

$r_0 = b - Ax_0$	u_a
$\beta = \lVert r_0 \rVert, \ v_1 = r_0/\beta, \ k = 1$	u_g
repeat	
$z_k = M^{-1} v_k$	u_m
$w_k = A z_k$	u_a
for $i = 1, \ldots, k$ **do**	
$h_{i,k} = v_i^T w_k$	u_g
$w_k = w_k - h_{i,k} v_i$	u_g
end for	
$h_{k+1,k} = \lVert w_k \rVert, v_{k+1} = w_k / h_{k+1,k}$	u_g
$V_k = [v_1, \ldots, v_k]$	
$H_k = \{h_{i,j}\}_{1 < i < j+1; 1 < j < k}$	
$y_k = \mathrm{argmin}_y \lVert \beta e_1 - H_k y \rVert$	u_g
$k = k + 1$	
until $\lVert \beta e_1 - H_k y_k \rVert \leq \tau$	
$d_k = V_k y_k$	u_g
$x_k = x_0 + M^{-1} d_k$	u_m

Mixed precision preconditioned GMRES

Left GMRES

$r_0 = b - Ax_0$	u_a
$s_0 = M^{-1} r_0$	u_m
$\beta = \|s_0\|, \ v_1 = s_0/\beta, \ k = 1$	u_g
repeat	
$\quad z_k = Av_k$	u_a
$\quad w_k = M^{-1} z_k$	u_m
$\quad$ **for** $i = 1, \dots, k$ **do**	
$\qquad h_{i,k} = v_i^T w_k$	u_g
$\qquad w_k = w_k - h_{i,k} v_i$	u_g
$\quad$ **end for**	
$\quad h_{k+1,k} = \|w_k\|, v_{k+1} = w_k / h_{k+1,k}$	u_g
$\quad V_k = [v_1, \dots, v_k]$	
$\quad H_k = \{h_{i,j}\}_{1<i<j+1; 1<j<k}$	
$\quad y_k = \mathrm{argmin}_y \|\beta e_1 - H_k y\|$	u_g
$\quad k = k + 1$	
until $\|\beta e_1 - H_k y_k\| \leq \tau$	
$x_k = x_0 + V_k y_k$	u_g

Flexible GMRES

$r_0 = b - Ax_0$	u_a
$\beta = \|r_0\|, \ v_1 = r_0/\beta, \ k = 1$	u_g
repeat	
$\quad z_k = M^{-1} v_k$	u_m
$\quad w_k = Az_k$	u_a
$\quad$ **for** $i = 1, \dots, k$ **do**	
$\qquad h_{i,k} = v_i^T w_k$	u_g
$\qquad w_k = w_k - h_{i,k} v_i$	u_g
$\quad$ **end for**	
$\quad h_{k+1,k} = \|w_k\|, v_{k+1} = w_k / h_{k+1,k}$	u_g
$\quad V_k = [v_1, \dots, v_k], \ Z_k = [z_1, \dots, z_k]$	
$\quad H_k = \{h_{i,j}\}_{1<i<j+1; 1<j<k}$	
$\quad y_k = \mathrm{argmin}_y \|\beta e_1 - H_k y\|$	u_g
$\quad k = k + 1$	
until $\|\beta e_1 - H_k y_k\| \leq \tau$	
$x_k = x_0 + Z_k d_k$	u_m

Mixed precision GMRES: state-of-the-art

- $u_a = u_m \ll u_g$ Applying A and M^{-1} in high precision to improve accuracy.
  Existing studies are dedicated to left-preconditioned GMRES

E. Carson and N. J. Higham. A New Analysis of Iterative Refinement and Its Application to Accurate Solution of Ill-Conditioned Sparse Linear Systems. SISC 2017.

E. Carson and N. J. Higham. Accelerating the solution of linear systems by iterative refinement in three precisions. SISC 2018.

E. Carson, N. J. Higham, and S. Pranesh. Three-Precision GMRES-Based Iterative Refinement for Least Squares Problems. SISC 2020.

P. R. Amestoy, A. Buttari, N. J. Higham, J.-Y. L'Excellent, T. M., and B. Vieublé. Five-Precision GMRES-Based Iterative Refinement. SIMAX 2024.

P. R. Amestoy, A. Buttari, N. J. Higham, J.-Y. L'Excellent, T. M., and B. Vieublé. Combining sparse approximate factorizations with mixed-precision iterative refinement. TOMS 2023.

E. Carson and N. Khan. Mixed Precision Iterative Refinement with Sparse Approximate Inverse Preconditioning. SISC 2023.

Mixed precision GMRES: state-of-the-art

- $u_a = u_m \ll u_g$ Applying A and M^{-1} in high precision to improve accuracy. Existing studies are dedicated to left-preconditioned GMRES
- $u_a = u_g \ll u_m$ Applying M^{-1} in a lower precision to improve performance. Existing studies are dedicated to flexible GMRES

M. Arioli and I. S. Duff. Using FGMRES to obtain backward stability in mixed precision. ETNA 2008.

J. D. Hogg and J. A. Scott. A fast and robust mixed-precision solver for the solution of sparse symmetric linear systems. TOMS 2010.

E. Carson and I. Daužickaitė. The stability of split-preconditioned FGMRES in four precisions. ETNA 2024.

E. Carson and I. Daužickaitė. Mixed precision sketching for least-squares problems and its application in GMRES-based iterative refinement. arXiv preprint arXiv:2410.06319, 2024.

Goal

Still no comprehensive understanding of GMRES in mixed precision:

- How should we set u_g, u_m, and u_a?
- Which preconditioning style (left, right, flexible) should we use?
- In short,

What are (*all*) the meaningful
mixed precision GMRES strategies?

- *Meaningful* refers to a strategy for which lowering any of its precisions also lowers the bound on its attainable accuracy

Attainable accuracy of GMRES

- The numerical behavior of GMRES is characterized by two main properties: **attainable accuracy** and **convergence rate**
- Unfortunately, no strong theoretical results for the convergence rate (even in exact arithmetic)

  A. Greenbaum, V. Pták and Z. Strakoš. Any nonincreasing convergence curve is possible for GMRES. SIMAX 1996.

$\Rightarrow$ We focus on the attainable accuracy

Restarted GMRES as iterative refinement

Restarted GMRES

$x = x_0$
$r = b - Ax$
repeat
    $s_0 = M^{-1} r_0$
    $\beta = \|s_0\|,\ v_1 = s_0/\beta,\ k = 1$
    for $k = 1, \dots, m$ **do**
        $z_k = A v_k$
        $w_k = M^{-1} z_k$
        for $i = 1, \dots, k$ **do**
            $h_{i,k} = v_i^T w_k$
            $w_k = w_k - h_{i,k} v_i$
        end for
        $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k / h_{k+1,k}$
        $V_k = [v_1, \dots, v_k]$
        $H_k = \{h_{i,j}\}_{1<i<j+1;1<j<k}$
        $y_k = \operatorname{argmin}_y \|\beta e_1 - H_k y\|$
    end for
    $x = x + V_k y_k$
    $r = b - Ax$
until $\|r\| \le \tau$

Restarted GMRES as iterative refinement

Restarted GMRES

$x = x_0$
$r = b - Ax$
repeat
  $s_0 = M^{-1} r_0$
  $\beta = \|s_0\|, \; v_1 = s_0/\beta, \; k = 1$
  for $k = 1, \dots, m$ **do**
    $z_k = A v_k$
    $w_k = M^{-1} z_k$
    for $i = 1, \dots, k$ **do**
      $h_{i,k} = v_i^T w_k$
      $w_k = w_k - h_{i,k} v_i$
    end for
    $h_{k+1,k} = \|w_k\|, v_{k+1} = w_k / h_{k+1,k}$
    $V_k = [v_1, \dots, v_k]$
    $H_k = \{h_{i,j}\}_{1<i<j+1; 1<j<k}$
    $y_k = \text{argmin}_y \|\beta e_1 - H_k y\|$
  end for
  $x = x + V_k y_k$
  $r = b - Ax$
until $\|r\| \le \tau$

Iterative refinement

$x = x_0$
$r = b - Ax$
repeat
  Solve $Ad = r$
  $x = x + d$
  $r = b - Ax$
until $\|r\| \le \tau$

- Restarted GMRES $\Leftrightarrow$ iterative refinement with m iterations of GMRES as correction solver
- Attainable accuracy of IR: $\mathrm{cond}(A, x)\boldsymbol{u_r} + \boldsymbol{u}$
- Unlike GMRES, we also have a good understanding of IR's convergence rate: $\frac{\|\widehat{d} - d\|}{\|d\|}$

$\Rightarrow$ GMRES's attainable accuracy is IR's (and thus restarted GMRES's) convergence rate!

Modular framework for the error analysis of GMRES

Error model and assumptions

We model GMRES as the following process:

1. Compute $C_k = M_L^{-1} A Z_k$ such that
$\widehat{C}_k = M_L^{-1} A Z_k + \Delta C_k, \quad \|\Delta C_k\| \le \varepsilon_c \|M_L^{-1} A Z_k\|.$
2. Compute $\widetilde{b} = M_L^{-1} b$ such that
$\widehat{b} = \widetilde{b} + \Delta b, \quad \|\Delta b\| \le \varepsilon_b \|\widetilde{b}\|.$
3. Solve $y_k = \arg\min_y \|\widetilde{b} - \widehat{C}_k y\|$ such that
$\widehat{y}_k = \arg\min_y \|\widehat{b} + \Delta_{\mathrm{ls}}^b - (\widehat{C}_k + \Delta_{\mathrm{ls}}^c) y\|$
$\|[\Delta_{\mathrm{ls}}^b, \Delta_{\mathrm{ls}}^c] e_j\| \le \varepsilon_{\mathrm{ls}} \|[\widehat{b}, \widehat{C}_k] e_j\|, \quad j = 1\colon k+1.$
4. Compute $x_k = Z_k \widehat{y}_k$ such that
$\widehat{x}_k = Z_k \widehat{y}_k + \Delta x_k, \quad \|\Delta x_k\| \le \varepsilon_x \|Z_k \widehat{y}_k\|.$

We also make the following assumptions:

5. All accuracy parameters are sufficiently less than 1: $0 \le \varepsilon_c, \varepsilon_b, \varepsilon_{\mathrm{ls}}, \varepsilon_x \ll 1$.
6. Z_k is not numerically singular to accuracy ε_x: $\kappa(Z_k)\varepsilon_x \ll 1$.
7. There exists a "key" iteration $k \le n$ at which $\widetilde{b}$ lies in the range of $M_L^{-1} A Z_k$ and $M_L^{-1} A Z_k$ is not numerically singular w.r.t. $\varepsilon = \varepsilon_c + \varepsilon_b + \varepsilon_{\mathrm{ls}}$: for all $\phi > 0$,
$\sigma_{\min}\left([\widetilde{b}\phi, M_L^{-1} A Z_k]\right) \ll \varepsilon \|[\widetilde{b}\phi, M_L^{-1} A Z_k]\|$
$\sigma_{\min}\left(M_L^{-1} A Z_k\right) \gg \varepsilon \|M_L^{-1} A Z_k\|.$

Bound on the attainable accuracies

At the key iteration k, $\widehat{x}_k$ satisfies $\frac{\|M_L^{-1} b - M_L^{-1} A \widehat{x}_k\|}{\|M_L^{-1} A\| \|\widehat{x}_k\| + \|M_L^{-1} b\|} \lesssim c(n,k)\xi$ and $\frac{\|\widehat{x}_k - x\|}{\|x\|} \lesssim c(n,k)\kappa(M_L^{-1} A)\xi$,

where $\xi = \alpha\varepsilon_c + \beta\varepsilon_b + \beta\varepsilon_{\mathrm{ls}} + \lambda\varepsilon_x, \quad \alpha = \frac{\|M_L^{-1} A Z_k\|}{\|M_L^{-1} A\| \sigma_{\min}(Z_k)}, \quad \beta = \max(1, \alpha), \quad \lambda = \kappa(Z_k)$

Recovering known results with the framework

- Unpreconditioned Householder-GMRES is backward stable, i.e., $\xi = O(u)$

  J. Drkošová, A. Greenbaum, M. Rozložník, Z. Strakoš. Numerical stability of GMRES. BIT 1995.

- Unpreconditioned MGS-GMRES is backward stable, i.e., $\xi = O(u)$

  C. C. Paige, M. Rozložník, and Z. Strakoš. Modified Gram-Schmidt (MGS), least squares, and backward stability of MGS-GMRES. SIMAX 2006.

- Flexible GMRES achieves $\xi = O(u\kappa(Z_k))$, hence:
  - FGMRES is *not* backward stable for *general* preconditioners
  - For a low precision LU preconditioner, under suitable assumptions, FGMRES is backward stable

  M. Arioli and I. S. Duff. Using FGMRES to obtain backward stability in mixed precision. ETNA 2008.

- Left and right preconditioned GMRES are also *not* backward stable in general:
  $\xi = O(\frac{\|M^{-1}\|\|A\|}{\|M^{-1}A\|}u) \leq O(\kappa(M)u)$

  P. R. Amestoy, A. Buttari, N. J. Higham, J.-Y. L'Excellent, T. M., and B. Vieublé. Five-Precision GMRES-Based Iterative Refinement. SIMAX 2024.

Deriving new results with the framework

- Unpreconditioned CGS2-GMRES is backward stable, i.e., $\xi = O(u)$

  A. Buttari, N. J. Higham, T. M., and B. Vieublé. A modular framework for the backward error analysis of GMRES. IMAJNA 2025.

- s-step (communication-avoiding) GMRES achieves $\xi = O(u\kappa(Z_k))$

  E. Carson and Y. Ma. On the backward stability of s-step GMRES. arXiv preprint arXiv:2409.03079, 2024.

- Sketched GMRES also achieves $\xi = O(u\kappa(Z_k))$

  L. Burke, E. Carson and Y. Ma. On the numerical stability of sketched GMRES. arXiv preprint arXiv:2503.19086, 2025.

- Mixed precision preconditioned GMRES

  A. Buttari, X. Liu, T. M., and B. Vieublé. Mixed precision strategies for preconditioned GMRES: a comprehensive analysis. HAL EPrint hal-05071696, 2025.

Applying the framework to mixed precision GMRES

We obtain the following bounds on the attainable forward error $\frac{\|\widehat{x}_k - x\|}{\|x\|}$:

- Left:
$$\kappa(M^{-1}A)\boldsymbol{u_g} + \max(\rho, \kappa(M^{-1}A))\boldsymbol{u_m} + \kappa(A)\boldsymbol{u_a}$$
$$(\rho \le \kappa(M^{-1}A)\kappa(M)\|Av_j\|/\|A\|)$$

- Right:
$$\kappa(AM^{-1})\kappa(M)\boldsymbol{u_g} + \kappa(M)\boldsymbol{u_m} + \kappa(A)\boldsymbol{u_a}$$

- Flexible:
$$\kappa(AM^{-1})\kappa(M)\boldsymbol{u_g} + \kappa(A)\boldsymbol{u_a}$$

Key observations

- For a given preconditioning style, the term in front of each precision is different
- For a given precision, the term in front of it depends on the preconditioning style

Mixed precision strategies

	Left	Right	Flexible
$u_a = u_m \ll u_g$	exists	new	new
$u_a = u_g \ll u_m$	new	new	exists
$u_a \ll u_g = u_m$	new	new	new
$u_a \ll u_g \ll u_m$	new	new	new
$u_a \ll u_m \ll u_g$	new	new	new
$u_g \ll u_a = u_m$	new	new	new
$u_g \ll u_a \ll u_m$	new	new	new
$u_m \ll u_a = u_g$	new	new	new
$u_m \ll u_a \ll u_g$	new	new	new
$u_g = u_m \ll u_a$	new	new	new
$u_g \ll u_m \ll u_a$	new	new	new
$u_m \ll u_g \ll u_a$	new	new	new

Theoretically meaningful
Not theoretically meaningful

Mixed precision strategies

	Left	Right	Flexible
$u_a = u_m \ll u_g$	exists	new	new
$u_a = u_g \ll u_m$	new	new	exists
$u_a \ll u_g = u_m$	new	new	new
$u_a \ll u_g \ll u_m$	new	new	new
$u_a \ll u_m \ll u_g$	new	new	new
$u_g \ll u_a = u_m$	new	new	new
$u_g \ll u_a \ll u_m$	new	new	new
$u_m \ll u_a = u_g$	new	new	new
$u_m \ll u_a \ll u_g$	new	new	new
$u_g = u_m \ll u_a$	new	new	new
$u_g \ll u_m \ll u_a$	new	new	new
$u_m \ll u_g \ll u_a$	new	new	new

Theoretically meaningful
Not theoretically meaningful

Experimental setting

- Three floating-point arithmetics : fp64 (D), fp32 (S), and bfloat16 (H)
- Restarted GMRES with $u_r = u =$ D, with reasonable tuning of the restart criterion. We report the total iteration count required to achieve

$$\frac{\|\widehat{x} - x\|_2}{\|x\|_2} \leq 10^{-10}$$

- Various matrices and preconditioners. A few representatives will be shown here; see paper for the full range of experiments:

  A. Buttari, X. Liu, T. M., and B. Vieublé. Mixed precision strategies for preconditioned GMRES: a comprehensive analysis. HAL EPrint hal-05071696, 2025.

Experimental illustration

1138_bus matrix, ILU preconditioner

			Iteration count		
u_a	u_g	u_m	Left	Right	Flexible
D	D	D	17	18	18
S	S	S	34	37	38

Theoretically meaningful
Not theoretically meaningful

Experimental illustration

1138_bus matrix, ILU preconditioner

			Iteration count		
u_a	u_g	u_m	Left	Right	Flexible
D	D	D	17	18	18
S	S	S	34	37	38
D	S	D	**26**	**25**	**31**

Theoretically meaningful
Not theoretically meaningful

Experimental illustration

1138_bus matrix, ILU preconditioner

			Iteration count		
u_a	u_g	u_m	Left	Right	Flexible
D	D	D	17	18	18
S	S	S	34	37	38
D	S	D	**26**	**25**	**31**
D	D	S	**25**	**28**	**18**

Theoretically meaningful

Not theoretically meaningful

Experimental illustration

1138_bus matrix, ILU preconditioner

			Iteration count		
u_a	u_g	u_m	Left	Right	Flexible
D	D	D	17	18	18
S	S	S	34	37	38
D	S	D	**26**	**25**	**31**
D	D	S	**25**	**28**	**18**
D	S	S	**27**	**31**	**31**

Theoretically meaningful

Not theoretically meaningful

Three-precision strategies: $u_a \ll u_m \ll u_g$

bcsstk19 matrix, LU preconditioner

			Iteration count		
u_a	u_g	u_m	Left	Right	Flexible
D	D	D	9	10	10
S	S	S	19	25	26
H	H	H	—	—	—
D	H	H	—	—	—
S	H	S	21	—	—
D	H	S	14	23	17
D	S	S	14	23	17
D	H	D	13	—	—

Three-precision strategies: $u_a \ll u_g \ll u_m$

bwm200 matrix, LU preconditioner

u_a	u_g	u_m	Iteration count Left	Right	Flexible
D	D	D	24	24	24
S	S	S	39	49	45
H	H	H	274	729	—
S	S	H	216	260	57
D	H	H	277	389	279
D	S	H	216	250	44
D	D	H	183	250	29
D	S	S	38	39	39

Caveats

- The comparison between strategies is based on the *bounds* on their attainable accuracies, which are not necessarily sharp in practice
- The quantities appearing in these bounds are not always known or easy to compute
- The attainable accuracy of GMRES only becomes descriptive of the convergence rate of restarted GMRES for large restart sizes

  Y. Zhao, T. Fukaya, L. Zhang, and T. Iwashita. Numerical Investigation into the Mixed Precision GMRES(m) Method Using FP64 and FP32. Journal of Information Processing, 2022.
- Precisions are not allowed to vary across iterations

  L. Giraud, S. Gratton, and J. Langou. Convergence in backward error of relaxed GMRES. SISC 2007.

  S. Gratton, E. Simon, D. Titley-Peloquin, P. L. Toint. A note on inexact inner products in GMRES. SIMAX 2022.

Guidelines

- If your preconditioner is costly to apply:
  use $u_a = u_g \ll u_m$ with flexible GMRES...
  ...or use a memory accessor with $u_a = u_g = u_m$
  P. R. Amestoy, A. Jego, J.-Y. L'Excellent, T.M., and G. Pichon. BLAS-based Block Memory Accessor with Applications to Mixed Precision Sparse Direct Solvers. HAL EPrint hal-05019106, 2025.

- If the orthonormalization is expensive / if a large restart size is required: use $u_a = u_m \ll u_g$ with left-preconditioned GMRES.

- If both of the above:
  use $u_a \ll u_m = u_g$ with left-preconditioned GMRES

Conclusions

- Preconditioned GMRES is not backward stable (due to cancellation in the preconditioned matrix–vector product)
- Restarted GMRES is more accurate than GMRES (due to its iterative refinement properties)
- The precisions of the matrix–vector product, preconditioner application, and orthonormalization play very different roles
- Left, right, and flexible preconditioning are all very different, with no clear winner. Each tends to favor using low precision for different operations

$\Rightarrow$ Many (new) mixed precision strategies, with various levels of performance–accuracy tradeoff

Thanks! Questions?