The effect of the ILU (0) Low-Precision Precondition to a variety of problems

KOGAKUIN UNIVERSITY | KUTE-TOKYO Kogakuin University of Technology & Engineering

Jin Kukita, Akihiro Fujii, Teruo Tanaka Kogakuin University

1. Introduction

- Mixed-precision arithmetic has been actively studied to accelerate iterative solvers for large sparse linear systems.
- In Krylov subspace methods, preconditioning plays a crucial role in convergence behavior, but its computational cost often dominates the overall solver performance.
- In our previous study [1], we obtained the results shown in Fig.1, and observed that even when low precision is applied to the ILU(0) preconditioner, convergence does not significantly degrade even when precision is aggressively reduced for the two-dimensional steady advection–diffusion equation. These results suggest that the ILU(0) preconditioner may be tolerant to extreme precision degradation.
- Although such extremely low precision is not supported by current hardware, software-based precision reduction[2] can still be useful for reducing memory traffic and for exploring the numerical tolerance of ILU(0).
- However, it remains unclear whether such robustness holds for a wide variety of sparse matrices, and how far the precision can be reduced without losing convergence.

2. Research Purpose

- The purpose of this study is to investigate the effect of extremely low-precision ILU(0) preconditioning on solver convergence.
- Specifically, we focus on ILU(0)-preconditioned BiCGSTAB and aim to:
  - evaluate whether convergence is preserved when aggressive precision reduction is applied to ILU(0),
  - analyze changes in the number of iterations caused by precision reduction,
  - clarify matrix characteristics suitable for low-precision ILU(0) preconditioning.

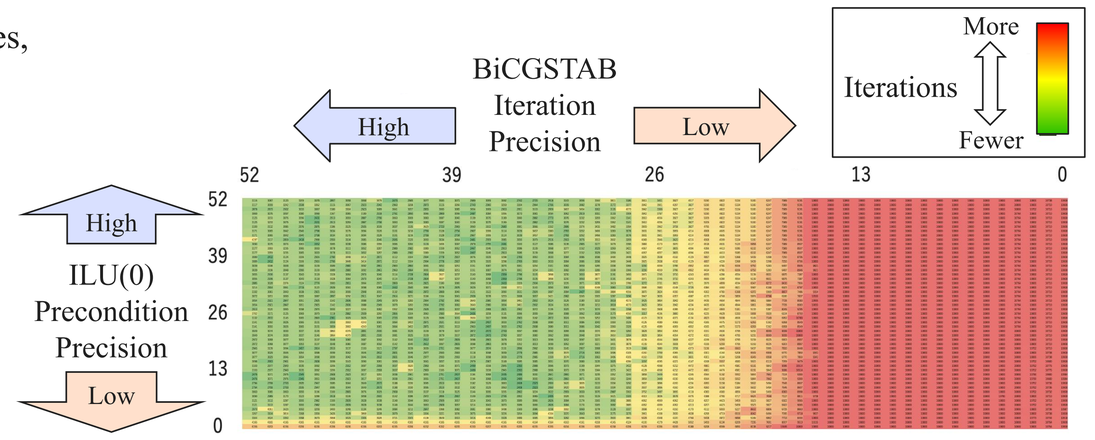

Fig.1: Low-precision in both preconditioning and iteration

3. Target Problem

- Sparse linear system: $A\boldsymbol{x} = \boldsymbol{b}$
  - Initial vector $\boldsymbol{x}_0$: 1.0
  - Reference solution $\boldsymbol{x}_{ref}$: random numbers in [-1, 1]
  - Right-hand side vector $\boldsymbol{b}$: generated as $\boldsymbol{b} = A\boldsymbol{x}_{ref}$
- Test matrices
  - 35 matrices from the SuiteSparse Matrix Collection
    - All selected matrices converge in FP64

4. Experimental Setup

- Numerical conditions
  - Stopping criterion: $\frac{||\boldsymbol{b}-A\boldsymbol{x}||}{||\boldsymbol{b}||} < 10^{-7}$
  - Iteration limit: $3n$ (n: the number of rows)
- Metrics
  - Convergence or divergence
  - Number of iterations to convergence

5. Experimental Result

- Fig.2 shows the convergence rate of the tested matrices for each ILU(0) precision. The value labeled “Median Rate” indicates the median ratio of the number of iterations required for convergence, normalized by the iteration count in double precision, among the matrices that successfully converged at each precision level.
- All tested matrices converged when 23 bits equivalent to single precision was used for the ILU(0) preconditioner.
- Even with 6-bit precision, more than half of the matrices successfully converged.
- As the precision was further reduced, the number of convergent matrices gradually decreased.
- For matrices that converged under low precision, the number of iterations often increased, but convergence was still achieved.

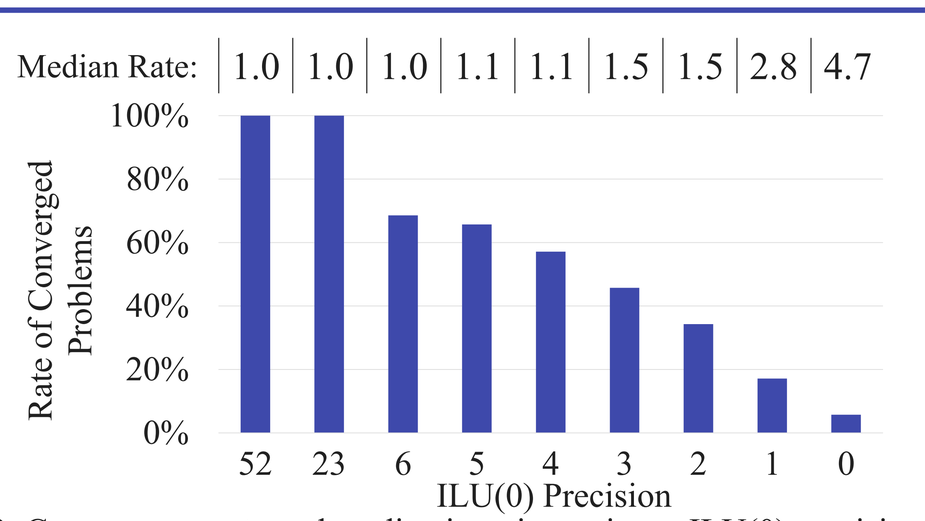

Fig.2: Convergence rate and median iteration ratio vs. ILU(0) precision

- Fig.3 shows the relationship between the matrix condition number and convergence robustness under reduced-precision ILU(0).
- The vertical axis represents the condition number of each matrix, and the horizontal axis indicates the number of precision levels at which convergence was achieved.
- Overall, matrices with relatively smaller condition numbers tended to converge at a larger number of reduced-precision levels.
- This result suggests that matrix numerical properties influence the tolerance of ILU(0) preconditioning to aggressive precision reduction.

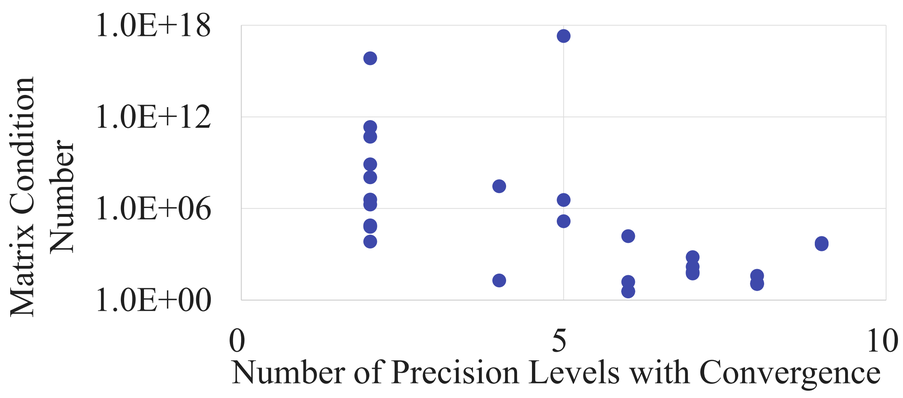

Fig.3: Convergence Behavior under Reduced-Precision ILU(0) and Matrix Condition Number

6. Conclusion

- This study demonstrated that ILU(0) preconditioning can tolerate aggressive precision reduction, with all matrices converging in single precision and more than half converging even at 6-bit precision.
- Matrices with smaller condition numbers tended to remain convergent across a wider range of reduced-precision levels.
- These results suggest that software-based extreme low-precision ILU(0) can be partially useful for certain classes of sparse matrices.
- Performance benefits (execution time, memory traffic) were not measured in this study and are left for future work.

7. Future Work

- Based on the observed robustness of ILU(0) to extreme precision reduction, future work will investigate which matrix properties govern this tolerance beyond condition number.
- The approach will be extended to other preconditioners to examine whether similar behavior can be observed.
- Investigating whether the precision of the ILU(0) preconditioner can be set automatically according to matrix characteristics, as clear selection criteria are not yet available.

Acknowledgements: This study was partially supported by JSPS KAKENHI Grand Number JP 23K11126.

References

[1] Jin Kukita, Akihiro Fujii, Teruo Tanaka: Analysis of the effect of low precision realization methods on the Mixed Precision ILU(0) Preconditioned BiCGSTAB Method (in Japanese), Annual Symposium of JSIAM 2025.

[2] Giuseppe Tagliavini, Stefan Mach, Davide Rossi, Andrea Marongiu, and Luca Benini: A Transprecision Floating-Point Platform for Ultra-Low Power Computing, DATE2018.