

Global and local growth behavior of GEPP and GECP

John Peca-Medlin

Abstract

Gaussian elimination (GE) remains the most used approach to solve dense linear systems in modern applications. For instance, GE with partial pivoting (GEPP) is the default solver in MATLAB when using the backslash operator with a general input matrix. Additionally, GE is a staple of introductory linear algebra courses (although, based on anecdotal evidence, it may not be a favorite among all students).

GE iteratively uses elimination updates below the diagonal on an initial input matrix $A \in \mathbb{R}^{n \times n}$ to transform A into an upper triangular linear system, which eventually builds the matrix factorization $A = LU$, where L is a unit lower triangular matrix and U is upper triangular. General nonsingular A will not have a LU factorization if any leading minors are singular, i.e., if $\det(A_{ij})_{i,j=1}^k = 0$ for some $k \leq n - 1$. Moreover, numerical stability of GE when using floating-point arithmetic is sensitive to any elimination steps that involve division by numbers close to 0. Hence, combining GE with certain pivoting schemes is preferable even if not necessary. A pivoting strategy results instead in the factorization $PAQ = LU$ where P and Q are permutation matrices (see [10, 13] for studies of random permutations generated using GEPP). I will focus on the two particular pivoting strategies GEPP and GE with complete pivoting (GECP), as well as the standard GE with no alternative pivoting strategy (GENP). I am interested in studying how GEPP and GECP behave on the same linear systems as well as studying large growth (see below) on particular subclasses of matrices, including orthogonal matrices. Moreover, as a means to better address the question of why large growth is rarely encountered, I further study matrices with a large difference in growth between using GEPP and GECP, and I explore how the smaller growth strategy dominates behavior in a small neighborhood of the initial matrix. This is a summary overview for my recently published results in [11].

Understanding the behavior of GE under floating-point arithmetic has been an ongoing focus of numerical analysis for over 60 years. Early work by Wilkinson in [17], which led to the start of modern error analysis, considered studying the relative errors of computed solutions $\hat{\mathbf{x}}$ to $A\mathbf{x} = \mathbf{b}$ through the bound

$$\frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_\infty}{\|\mathbf{x}\|_\infty} \leq 4n^2 \epsilon_{\text{machine}} \kappa_\infty(A) \rho(A), \quad \text{where } \rho(A) = \frac{\max_{i,j,k} |A_{ij}^{(k)}|}{\max_{i,j} |A_{ij}|}$$

is the *growth factor*, $\kappa_\infty(A) = \|A\|_\infty \|A^{-1}\|_\infty$ is the ∞ -condition number, and $A^{(k)}$ is the intermediate form of A after $k - 1$ GE steps, with $A^{(1)} = A$ and $A^{(n)} = U$. The growth factor returns the relative largest entry ever encountered running through the entire GE process. Using GECP, $\rho(A)$ takes the simpler form $\rho(A) = \max_k |U_{kk}|/|U_{11}|$ while in general the largest entry is not necessarily captured by U . In well-conditioned systems, error analysis of GE can then be reduced to analysis of the growth factor.

Understanding worst-case bounds on ρ (i.e., the *growth problem*) using different pivoting strategies has been a continued area of research since Wilkinson's initial analysis in the early 1960s. Wilkinson resolved the worst-case behavior of GEPP, showing the bound $\rho^{\text{GEPP}}(A) \leq 2^{n-1}$ for $A \in \mathbb{R}^{n \times n}$ is sharp [17, 18]. For GECP, however, Wilkinson in the same work only provided an upper bound of

$$\rho^{\text{GECP}}(A) \leq (n \cdot 2 \cdot 3^{1/2} \dots n^{1/(n-1)})^{1/2} \leq 2 \cdot n^{0.25 \ln n + 0.5},$$

which is believed to be very pessimistic. For a long time a conjecture (of apocryphal origins; cf. [5]) ventured the bound $\rho^{\text{GEC}} \leq n$ (for real matrices). This conjecture survived for almost three decades, until Gould [6] found a “counterexample” for $n = 13$ by producing a matrix using floating-point arithmetic with GEC growth of 13.0205 in 1991. Edelman [4] confirmed a true counterexample in exact arithmetic one year later. No substantial progress on the GEC growth problem was made for another 30 years later until 2023. Edelman and Urschel [5] establish a linear lower bound on ρ^{GEC} at least $1.0045n$ for $n > 10$, by upgrading Edelman’s computational technique to upgrade Gould’s floating-point “counterexample” into an exact arithmetic counterexample (by updating one entry by 10^{-7}) into a theorem establishing floating-point and exact arithmetic growth factors cannot be too far from one another (cf. Theorem 4.2). Moreover, later that same year, Edelman and Urschel along with now Bisain [3] provide the first substantial improvement to Wilkinson’s upper bound for ρ^{GEC} , where (using a modified Hadamard inequality) they provide a bound of $n^{c \log n + 0.91}$ using $c \approx 0.20781$, which beats Wilkinson’s original exponential $\log n$ coefficient of 0.25. The huge gulf between these lower and upper bounds leaves much on the table for future improvements.

More recent analysis of matrix algorithm efficiency and accuracy has shifted away from worst-case analysis to more modern approaches, such as smoothed analysis (cf. [15] and references therein), which studies behavior under random perturbations. A full smoothed analysis using Gaussian (additive) perturbations has been successfully implemented in the case of GENP growth factors by Sankar, Spielman, and Teng [14], but has remained out of reach for both the GEPP and GEC growth problem. The closest such result for GEPP was the recent average-case analysis work of Huang and Tikhomirov [8], which established high probability polynomial growth bounds using input matrices with independent and identically distributed (iid) standard normal entries, but their proof strategy cannot be upgraded to a smoothed analysis approach (i.e., they only establish bounds on Gaussian perturbations of the zero matrix). No such (non-empirical) result for GEC has come close to smoothed analysis nor a full average-case analysis (other than with very particular structured random matrices in [9]). So worst-case analysis remains relevant for these pivoting strategies.

Our focus will be on studying the growth behavior of using both GEPP and GEC on the same linear system, and on how each strategy can inform growth behavior about the other. In particular, I am interested in local growth behavior around a system with large differences in growth behavior between both strategies. For instance, I am interested in the question of whether small perturbations on an initial system with a large discrepancy in growth behavior between both strategies concentrates toward the smaller growth for both strategies? For example, if A has large GEPP growth and small GEC growth, and if G is an iid Gaussian matrix, I am interested in how often $\rho^{\text{GEPP}}(A + \varepsilon G) \approx \rho^{\text{GEC}}(A)$ for sufficiently small ε . To move toward addressing this question, I will establish bounds on how far apart differences in growth between both strategies can be when used on the same input matrix.

I will further focus on large growth for particularly structured matrices, including orthogonal matrices. This will include a refinement to the largest possible GEPP growth on orthogonal matrices, while also establishing a rich set of matrices for further study with large growth difference behavior. Growth for structured systems has proven fruitful in recent studies. For instance, growth using GENP, GEPP and GE with rook pivoting (GERP) for matrices formed using the Kronecker products of rotation matrices (i.e., $\otimes^n \text{SO}(2)$) is now completely understood [12], with even a full picture understood using GEC on a further subclass of these Kronecker product matrices [9]. Although GE should not be a first choice for solving orthogonal linear systems (viz., $Q\mathbf{x} = \mathbf{b}$ has

the solution $\mathbf{x} = Q^T \mathbf{b}$ when Q is orthogonal), there are situations when applying GE to orthogonal matrices makes sense. For example, Barlow needed to understand the effect of GEPP on orthogonal matrices to carry out error analysis of bidiagonal reduction [1]; this led to Barlow and Zha's analysis in [2] using GEPP on orthogonal matrices, which they showed maximized a different L^2 -growth factor. Additionally, while original studies of random growth factors tended to focus on ensembles with iid entries (cf. [16]), many authors noted and explored the potential that orthogonal matrices can produce large growth [7, 16]. Hence, orthogonal matrices remain a rich source of study for potential large growth factors.

References

- [1] J. L. Barlow. More accurate bidiagonal reduction for computing the singular value decomposition. *SIAM Journal of Matrix Analysis and Applications*, 23:761–798, 2002.
- [2] J. L. Barlow and H. Zha. Growth in Gaussian elimination, orthogonal matrices, and the 2-norm. *SIAM Journal of Matrix Analysis and Applications*, 19:807–815, 1998.
- [3] A. Bisain, A. Edelman, and J. Urschel. A new upper bound for the growth factor in Gaussian elimination with complete pivoting. *arXiv preprint arXiv:2312.00994*, 2023.
- [4] A. Edelman. The complete pivoting conjecture for Gaussian elimination is false. *The Mathematica Journal*, 2:58–61, 1992.
- [5] A. Edelman and J. Urschel. Some new results on the maximum growth factor in Gaussian elimination. *SIAM Journal of Matrix Analysis and Applications*, 45(2):967–991, 2024.
- [6] N. Gould. On growth in Gaussian elimination with complete pivoting. *SIAM Journal of Matrix Analysis and Applications*, 12:354–361, 1991.
- [7] N. J. Higham and D. Higham. Large growth factors in Gaussian elimination with pivoting. *SIAM Journal of Matrix Analysis and Applications*, 10:155–164, 1989.
- [8] H. Huang and K. Tikhomirov. Average-case analysis of the Gaussian elimination with partial pivoting. *Probability Theory and Related Fields*, 189:501–567, 2024.
- [9] J. Peca-Medlin. Complete pivoting growth of butterfly matrices and butterfly Hadamard matrices. *arXiv preprint arXiv:2410.06477*, 2024.
- [10] J. Peca-Medlin. Distribution of the number of pivots needed using Gaussian elimination with partial pivoting on random matrices. *Annals of Applied Probability*, 24(2):2294–2325, 2024.
- [11] J. Peca-Medlin. Growth factors of orthogonal matrices and local behavior of Gaussian elimination with partial and complete pivoting. *SIAM Journal on Matrix Analysis and Applications*, 45(3):1599–1620, 2024.
- [12] J. Peca-Medlin and T. Trogdon. Growth factors of random butterfly matrices and the stability of avoiding pivoting. *SIAM Journal of Matrix Analysis and Applications*, 44(3):945–970, 2023.
- [13] J. Peca-Medlin and C. Zhong. On the longest increasing subsequence and number of cycles of butterfly permutations. *arXiv preprint arXiv:2410.20952*, 2024.

- [14] A. Sankar, D. Spielman, and S.-H. Teng. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM Journal of Matrix Analysis and Applications*, 11:335–360, 1990.
- [15] D. Spielman and S.-H. Teng. *Smoothed analysis: an attempt to explain the behavior of algorithms in practice*, Volume 52 (10). ACM, New York, 2009.
- [16] L. Trefethen and R. Schreiber. Average case stability of Gaussian elimination. *SIAM Journal of Matrix Analysis and Applications*, 11:335–360, 1990.
- [17] J. Wilkinson. Error analysis of direct methods of matrix inversion. *J. Assoc. Comput. Mach.*, 8:281–330, 1961.
- [18] J. Wilkinson. *The Algebraic Eigenvalue Problem*. Oxford University Press, London, UK, 1965.