Symmetric Pseudospectral Shattering and Fast Divide-and-Conquer for the Definite Generalized Eigenvalue Problem

James Demmel, Ioana Dumitriu, and Ryan Schneider

Abstract

Overview: We adapt the asymptotically fastest-known algorithm for diagonalizing arbitrary matrix pencils – as well as the related phenomenon of pseudospectral shattering – to the definite generalized eigenvalue problem. Put simply, we obtain significant efficiency gains by preserving and exploiting structure, in this case symmetry. In doing so, our work provides a general road map for tailoring fast diagonalization to structured problems.

Recent work in randomized numerical linear algebra produced the first sub-$O(n^3)$ algorithms for diagonalizing any matrix A or matrix pencil (A, B) [3, 5]. The key insight of this work is the phenomenon of *pseudospectral shattering*, where a random perturbation to a matrix, or pencil, has a regularizing effect on its (pseudo)spectrum. A result of smoothed analysis [11], shattering is characterized by a minimum eigenvalue gap and minimally well-conditioned eigenvectors. Moreover, it implies success for fast divide-and-conquer eigensolvers, which can diagonalize a perturbed matrix/pencil with essentially optimal complexity (that is, complexity equal to matrix multiplication up to log factors). The name "pseudospectral shattering" is derived from the fact that a random grid covering the ϵ-pseudospectra of the perturbed problem separates its disjoint components, and the eigenvalues they contain, into separate grid boxes for ϵ sufficiently small – i.e., inverse polynomial in n.

Pseudospectral shattering suggests a simple, high-level approach to eigenvalue problems: apply a random perturbation and run a fast version of divide-and-conquer, where the shattering grid can be used to reliably divide the spectrum at each step. The result is an accurate diagonalization, in the backward-error sense, provided the initial perturbation is small. Prior to [3], which introduced pseudospectral shattering in the context of the standard eigenvalue problem, no way of leveraging divide-and-conquer's natural parallelization to obtain fast diagonalizations of arbitrary matrices (or pencils) was known. Importantly, [5] established that this approach can be implemented without relying on matrix inversion, thereby promoting stability while also minimizing associated communication costs (following Ballard et al. [2]).

These randomized eigensolvers, which we refer to collectively as pseudospectral divide-and-conquer, are fully general. In particular, both [3] and [5] allow matrices to be arbitrary and apply Ginibre perturbations to obtain a guarantee of pseudospectral shattering. This begs the question: how can we adapt these algorithms to better handle symmetric or sparse inputs, for which dense Gaussian perturbations are structure-destroying? Going further: if we can achieve pseudospectral shattering while maintaining structure – i.e., via structured perturbations – how can we translate that into efficiency gains in divide-and-conquer?

We answer these question for the definite generalized eigenvalue problem, which corresponds to pencils (A, B) in which A and B are Hermitian and the *Crawford number* $\gamma(A, B)$ satisfies

$$\gamma(A, B) = \min_{||x||_2=1} |x^H(A + iB)x| > 0. \tag{1}$$

Pencils arising in scientific computing and machine learning are often definite [7, 6]. We note two important sub-problems in particular: (1) the Hermitian eigenvalue problem, corresponding to $B = I$, and (2) the generalized symmetric definite eigenvalue problem, in which B is positive definite.

As inputs to divide-and-conquer, definite pencils exhibit a number of properties that can be leveraged for improved efficiency. Most notably, the eigenvalues of a definite pencil (A, B) are real (and in fact any ϵ-pseudospectrum that considers only Hermitian perturbations to A and B will be constrained to the real axis for ϵ sufficiently small). Additionally, definite pencils are regular and satisfy stronger eigenvalue/eigenvector perturbation bounds than the generic case (see e.g., [12]). Finally, where an arbitrary pencil (A, B) is diagonalized by a pair of eigenvector matrices – if it is diagonalizable at all – a definite pencil can always be diagonalized by a single matrix. That is, for any definite pencil (A, B) there exists invertible X such that

$$(X^H AX, X^H BX) = (\Lambda_A, \Lambda_B) \tag{2}$$

for Λ_A and Λ_B diagonal.

Motivated by these observations, we devise a version of pseudospectral divide-and-conquer that pursues efficiency by maintaining definiteness through both the initial random perturbation and the subsequent recursive procedure. The main ingredients are the following:

1. We prove shattering for a symmetric pseudospectrum

$$\Lambda_\epsilon^{\mathrm{sym}}(A, B) = \left\{ z : \begin{array}{l} (A+E)u = z(B+F)u \;\text{ for }\; u \neq 0 \;\text{ and} \\ E, F \text{ Hermitian with } \sqrt{||E||_2^2 + ||F||_2^2} \leq \epsilon \end{array} \right\} \tag{3}$$

   under random perturbations that are either diagonal or sampled from the Gaussian Unitary Ensemble (GUE). The diagonal case builds on work of Minami [8] and implies a remarkably simple path to structured shattering for (Hermitian) banded matrices. The GUE case, meanwhile, leverages recent results of Aizenman et al. [1]. In both settings, the key insight is a bound on the probability that a perturbed Hermitian matrix has a certain number of eigenvalues in a given interval of the real axis.

2. Next, we demonstrate that (inverse-free) iterative methods for computing spectral projectors of (A, B) – i.e., projectors onto deflating subspaces corresponding to sets of eigenvalues, which are the key to the recursive splits of divide-and-conquer – can be optimized for fast convergence on problems with real eigenvalues [4]. This is the primary advantage we gain access to by preserving definiteness (and itself generalizes work of Nakatsukasa et al. [10]).

Combining points one and two yields a specialized version of pseudospectral divide-and-conquer that is significantly more efficient on definite inputs. Ongoing work seeks high performance implementations of both standard pseudospectral divide-and-conquer and this specialization. Accordingly, and in parallel with broader efforts to deploy randomized algorithms in numerical linear algebra [9], our work represents an important step towards bringing fast, randomized diagonalization to practice.

References

[1] M. Aizenman, R. Peled, J. Schenker, M. Shamis, and S. Sodin. Matrix regularizing effects of Gaussian perturbations. *Communications in Contemporary Mathematics*, 19(03):1750028, 2017.

[2] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. Minimizing Communication in Numerical Linear Algebra. *SIAM Journal on Matrix Analysis and Applications*, 32(3):866–901, 2011.

[3] J. Banks, J. Garza-Vargas, A. Kulkarni, and N. Srivastava. Pseudospectral Shattering, the Sign Function, and Diagonalization in Nearly Matrix Multiplication Time. *Foundations of Computational Mathematics*, 23:1959–2047, 2023.

[4] J. Demmel, I. Dumitriu, and R. Schneider. Fast and Inverse-Free Algorithms for Deflating Subspaces. arXiv:2310.00193, 2024.

[5] J. Demmel, I. Dumitriu, and R. Schneider. Generalized Pseudospectral Shattering and Inverse-Free Matrix Pencil Diagonalization. *Foundations of Computational Mathematics*, 2024.

[6] B. Ford and G. Hall. The generalized eigenvalue problem in quantum chemistry. *Computer Physics Communications*, 8(5):337–348, 1974.

[7] O. Mangasarian and E. Wild. Multisurface proximal support vector machine classification via generalized eigenvalues. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(1):69–74, 2006.

[8] N. Minami. Local fluctuation of the spectrum of a multidimensional Anderson tight binding model. *Communications in Mathematical Physics*, 177(3):709–725, 1996.

[9] R. Murray, J. Demmel, M. W. Mahoney, N. B. Erichson, M. Melnichenko, O. A. Malik, L. Grigori, P. Luszczek, M. Derezinski, M. E. Lopes, T. Liang, H. Luo, and J. Dongarra. Randomized Numerical Linear Algebra: A Perspective on the Field with an Eye to Software. Technical Report UCB/EECS-2023-19, EECS Department, University of California, Berkeley, Feb 2023.

[10] Y. Nakatsukasa, Z. Bai, and F. Gygi. Optimizing Halley's Iteration for Computing the Matrix Polar Decomposition. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2700–2720, 2010.

[11] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.

[12] G. W. Stewart. Perturbation bounds for the definite generalized eigenvalue problem. *Linear Algebra and its Applications*, 23:69–85, 1979.