Selected Eigenpairs of Fixed-Bandwidth Positive Definite Matrices in Bit Complexity

August 4, 2026

Abstract

Let A be an $n \times n$ real symmetric positive definite matrix of fixed half-bandwidth w, with nonzero integer entries of $O(\log n)$ bits and spectrum contained in $[n^{-10}, n^{10}]$. Given $1 \le k \le n$ and dyadic $0 < \epsilon, \delta \le 1/4$, we give a randomized finite-bit algorithm that terminates on every random branch and, with probability at least $1 - 9\delta/64 > 1 - \delta$, returns approximations to the leading k eigenpairs, equivalently the leading symmetric singular triples. Each returned unit vector q_i is represented by a nonzero finite dyadic vector y_i, interpreted as $y_i/\sqrt{y_i^T y_i}$, and the nonnegative dyadic eigenvalue estimate $\widehat{\sigma}_i$ satisfies

$$\|Aq_i - \widehat{\sigma}_i q_i\|_2 \le C_w \epsilon \|A\|_2, \qquad |q_i^T q_j| \le C_w \epsilon \quad (i \ne j), \qquad |\widehat{\sigma}_i - \lambda_i(A)| \le C_w \epsilon \|A\|_2.$$

For $Q = [q_1 \ \cdots \ q_k]$ the algorithm also guarantees $\|Q^T Q - I\|_2 \le C_w k\epsilon$. If $L = 1 + b_\epsilon + b_\delta + \lceil \log_2(n/(\epsilon\delta)) \rceil$, where b_ϵ and b_δ are the canonical binary encoding lengths, its worst-case cost is $O_w(nkL^3)$ bit operations, it uses $O(nL + kL^2)$ random bits, and its output uses $O(nkL)$ bits. The special case in which ϵ^{-1} and δ^{-1} are polynomial in n and their encodings have $O(\log n)$ bits therefore costs $O_w(nk \log^3 n)$ bit operations. No input eigengap is assumed. The construction uses a finite-bit diagonal perturbation, guarded banded Sturm bisection, and a multishift bank of shifted-inverse polynomial filters. Constants denoted by C_w or hidden by $O_w(\cdot)$ depend only on w.

Note. This paper was fully AI generated by GPT 5.6 Sol at ultra reasoning effort using an internal harness, from a fairly detailed prompt by Richard Peng that outlined the projection and iteration approach and explicitly requested a bit-complexity analysis rather than a word-complexity analysis.

Contents

1 Introduction

Computing a selected part of the spectrum of a structured symmetric matrix is an output-sensitive numerical problem. An $n \times n$ matrix of fixed half-bandwidth has only $O_w(n)$ nonzero entries, whereas k explicit eigenvectors already contain nk coordinates. At the same time, locating the corresponding eigenvalues is not enough: the returned vectors should have small residuals and be mutually nearly orthogonal, even when the input spectrum is clustered. In a bit model, the precision used by the numerical primitives, the randomness, and the finite representation of the output must all be included in the complexity bound.

For half-bandwidth one, this is the classical selected symmetric tridiagonal eigenproblem. Sturm bisection and inverse iteration give the basic localization-and-recovery paradigm [BMW67, Wil62]. MRRR methods target k selected eigenpairs in $O(nk)$ floating-point arithmetic when the required relatively robust representation properties hold [PD00, DP04, DPV06, WL13]. Subset calculations and tightly clustered spectra introduce additional representation and orthogonality issues [MPV06,DPV05,DMPV08]. These results treat the half-bandwidth-one special case in floating-point arithmetic, rather than giving an end-to-end finite-bit bound for the wider fixed-bandwidth input considered here. For symmetric band and fixed-block-width matrices more generally, direct divide-and-conquer and block-tridiagonal methods compute full or approximate eigensystems in arithmetic models [Arb92, GSU01, GWMG03]; structured spectral divide-and-conquer can instead keep the eigenvector matrix implicit or compressed [VXCB16, ŠK21]. Those output models do not by themselves give a bit bound for k explicit vectors.

This question is also a structured analogue of the bidiagonal singular-value problem posed by Amsel et al. [ABB+26, Problem 3.8], which asks for k singular triplets of an arbitrary upper-bidiagonal matrix in $O(nk)$ operations, together with finite-precision residual and orthogonality guarantees. MRRR-based and associated-tridiagonal approaches have been developed for selected bidiagonal singular triplets [WLV06, WL12, MDV20]. This is a nearby but different problem: a general bidiagonal input requires two coupled families of left and right singular vectors. For a positive definite symmetric matrix, by contrast, an eigenpair (λ_i, q_i) is the symmetric singular triple (λ_i, q_i, q_i). We study this symmetric problem under fixed bandwidth, polynomial conditioning, and finite integer input; the result does not solve the arbitrary-bidiagonal problem.

Our contribution. We give and analyze BandedPSD-Eigenpairs, a randomized finite-bit algorithm for the leading k eigenpairs in this setting. All complexity bounds below count bit operations. Constants denoted by C_w or hidden by $O_w(\cdot)$ may depend only on the fixed half-bandwidth w.

Theorem 1.1 (Top eigenpairs in bit complexity)**.** *Let A be an $n \times n$ real symmetric positive definite matrix with $A_{ij} = 0$ whenever $|i - j| > w$, where w is fixed. Suppose every nonzero entry of A is an integer with $O(\log n)$ bits and*

$$n^{-10} \le \lambda_n(A) \le \lambda_1(A) \le n^{10}.$$

Given $1 \le k \le n$ and dyadic $0 < \epsilon, \delta \le 1/4$, the randomized algorithm BandedPSD-Eigenpairs terminates on every random branch and outputs k triples $(\widehat{\sigma}_i, q_i, q_i)$, where $\widehat{\sigma}_i$ is a nonnegative dyadic and q_i is a normalized dyadic-vector encoding of a real unit vector. With probability at least $1 - 9\delta/64$, and therefore at least $1 - \delta$, simultaneously for $1 \le i \le k$ and distinct $i, j \in \{1, \dots, k\}$,

$$\begin{aligned} \|Aq_i - \widehat{\sigma}_i q_i\|_2 &\le C_w \epsilon \|A\|_2, \\ |q_i^T q_j| &\le C_w \epsilon \qquad (i \ne j), \end{aligned}$$

$$|\widehat{\sigma}_i - \lambda_i(A)| \le C_w \epsilon \|A\|_2.$$

Moreover, for $Q = [q_1 \ \cdots \ q_k]$, $\|Q^T Q - I\|_2 \le C_w k\epsilon$. *Writing*

$$L = 1 + b_\epsilon + b_\delta + \left\lceil \log_2 \frac{n}{\epsilon\delta} \right\rceil,$$

where b_ϵ *and* b_δ *are the canonical binary encoding lengths of the two dyadic inputs, the algorithm uses* $O_w(nkL^3)$ *bit operations, uses* $O(nL + kL^2)$ *random bits, and produces an* $O(nkL)$*-bit output. In particular, if* ϵ^{-1} *and* δ^{-1} *are polynomial in* n *and their encodings have* $O(\log n)$ *bits, the cost is* $O_w(nk \log^3 n)$ *bit operations.*

The theorem assumes no eigengap for A. In particular, when eigenvalues are repeated, it does not single out a prescribed basis of the corresponding eigenspace; the top-k guarantee is expressed through indexed eigenvalue estimates, residuals, and near orthogonality. When ϵ^{-1} and δ^{-1} are polynomial in n and their canonical encodings have $O(\log n)$ bits, the running time is $O_w(nk \log^3 n)$ and the explicit output occupies $O(nk \log n)$ bits.

On the joint gap, safe-localization, and random-overlap event used in the proof, which has probability at least $1 - 9\delta/64$, a small finite-bit random diagonal perturbation creates quantitative separation. Guarded banded Sturm bisection then localizes each leading eigenvalue of the perturbed matrix. Around every localization, an even power of a shifted inverse acts as an almost rank-one spectral filter; the sum of these filters approximates the leading spectral projector. Applying the individual filters to a finite random probe recovers the k vectors without forming and orthogonalizing a dense range basis. Finally, certified Rayleigh quotients and perturbation bounds transfer the residuals and indexed eigenvalue estimates back to the original matrix. Fixed bandwidth makes every inertia computation and shifted solve linear in n up to the explicitly counted precision factors; explicit guards ensure termination on every random branch.

Nearby techniques and complexity results. Polynomial and rational filtering, FEAST, and spectrum-slicing methods also construct approximate projectors from matrix functions or shifted solves [TP14, LXV+16, XS16, LXES19]. They typically recover a block invariant subspace and use a Ritz or orthogonalization stage, which is a different output mechanism from the individually centered filters used here. The level-spacing estimate behind our diagonal perturbation comes from random-operator correlation bounds [BHS07]; these are probabilistic spectral estimates, not eigensolvers. Work on stable or bit-complexity diagonalization for general matrices and dense Hermitian inputs provides a broader complexity context [BGVKS23, Sha25, Sob25], but concerns full, dense, or values/projector output models rather than the fixed-bandwidth selected output considered here.

Organization. Section 2 collects spectral and perturbation notation, block-banded inertia facts, and the finite-precision model. Section 3 explains the separation, localization, filtering, recovery, and certification steps. Section 4 constructs the numerical primitives and the main algorithm and proves its correctness, probability, and bit-complexity bounds.

2 Preliminaries

Spectral notation. All vector spaces are real. The notation $\|\cdot\|_2$ denotes the Euclidean norm for vectors and the induced spectral norm for matrices. For a non-singular matrix B, write $\kappa_2(B) = \|B\|_2 \|B^{-1}\|_2$. For a real symmetric $n \times n$ matrix T, order the eigenvalues as $\lambda_1(T) \ge \cdots \ge \lambda_n(T)$ and choose a corresponding orthonormal eigenbasis $u_1(T), \ldots, u_n(T)$. We write

$$\mathrm{spec}(T) = \{\lambda_1(T), \ldots, \lambda_n(T)\}, \qquad \mathrm{dist}(z, \mathrm{spec}(T)) = \min_i |z - \lambda_i(T)|,$$

and let e_i be the ith coordinate vector. For symmetric matrices, $S \preceq T$ means that $T - S$ is positive semidefinite. Under the hypotheses of Theorem 1.1, A is positive definite.

For a scalar function f defined on $\operatorname{spec}(T)$, the spectral calculus gives

$$f(T) = \sum_{i=1}^{n} f(\lambda_i(T)) u_i(T) u_i(T)^T, \qquad \|f(T)\|_2 = \max_i |f(\lambda_i(T))|. \tag{1}$$

In particular, for an interval $J \subseteq \mathbb{R}$ we use

$$\mathbf{1}_J(T) = \sum_{\lambda_i(T) \in J} u_i(T) u_i(T)^T, \qquad \|(T - zI)^{-1}\|_2 = \operatorname{dist}(z, \operatorname{spec}(T))^{-1}$$

when $z \notin \operatorname{spec}(T)$. We use Weyl's standard perturbation bound

$$|\lambda_i(T + E) - \lambda_i(T)| \le \|E\|_2 \qquad (1 \le i \le n) \tag{2}$$

for symmetric E. We also use the elementary normalization estimate: if u is unit, $\alpha \neq 0$, and $\|y - \alpha u\|_2 \le \eta < |\alpha|$, then

$$\left\| \frac{y}{\|y\|_2} - \operatorname{sgn}(\alpha) u \right\|_2 \le \frac{2\eta}{|\alpha|}. \tag{3}$$

Banded matrices and inertia. A matrix has half-bandwidth w if $T_{ij} = 0$ whenever $|i - j| > w$. Set $b = \max\{1, w\}$ and partition the coordinates into consecutive blocks of size b, with a possibly shorter final block. Then every matrix of half-bandwidth w is block tridiagonal. We write $T^{[j]}$ for the leading principal matrix through block j, and use the same superscript for a conformally blocked matrix. For $\eta > 0$, a real query x is *η-safe for T* if

$$\min_j \operatorname{dist}(x, \operatorname{spec}(T^{[j]})) \ge \eta. \tag{4}$$

For a symmetric matrix P, let $\iota_+(P)$ be its number of positive eigenvalues, counted with multiplicity. Thus

$$\iota_+(T - xI) = |\{i : \lambda_i(T) > x\}|. \tag{5}$$

We use Sylvester's law of inertia in its block Schur-complement form. Write B_j and C_j for the diagonal and superdiagonal blocks of T, respectively. If the block LDL^T elimination of $T - xI$ has nonsingular pivots

$$P_1 = B_1 - xI, \qquad P_j = B_j - xI - C_{j-1}^T P_{j-1}^{-1} C_{j-1}, \tag{6}$$

then

$$\iota_+(T - xI) = \sum_j \iota_+(P_j). \tag{7}$$

Safety in Equation (4) guarantees that these exact pivots are nonsingular.

For the input matrix A, define the scale

$$W = \max_{1 \le i \le n} \sum_{j=1}^{n} |A_{ij}|.$$

Symmetry and fixed bandwidth give

$$\|A\|_2 \le W \le (2w + 1)\|A\|_2. \tag{8}$$

The first inequality is the symmetric maximum-row-sum bound, while bandwidth gives the second. Moreover, $W \ge 1$ and W has $O(\log n)$ bits.

Finite representations and guarded arithmetic. A finite dyadic is a number $m2^e$ with integers m, e. A *normalized dyadic-vector encoding* is a nonzero finite dyadic vector y, interpreted as the exact real unit vector $y/\sqrt{y^T y}$. Only y and this interpretation are stored. Homogeneous expressions are evaluated without expanding the generally irrational normalization; for example,

$$\left(\frac{y}{\sqrt{y^T y}}\right)^T A \left(\frac{y}{\sqrt{y^T y}}\right) = \frac{y^T A y}{y^T y}.$$

All running times count bit operations. We use schoolbook arithmetic, so an arithmetic operation on $O(p)$-bit integer data costs $O(p^2)$ bit operations. The canonical form of a nonzero dyadic is its reduced mantissa–exponent pair $m2^e$, with m odd. For the dyadic inputs ϵ and δ, let b_ϵ and b_δ denote the total binary lengths of the signed integers in these canonical forms, and put

$$L = 1 + b_\epsilon + b_\delta + \left\lceil \log_2 \frac{n}{\epsilon\delta} \right\rceil.$$

Fix the constant

$$K_w = 2^{20(w+1)^3}.$$

The following rounding model is shared by all numerical primitives.

Definition 2.1 (Guarded dyadic arithmetic)**.** *For integers $p \geq 8$ and $E \geq p$, a guarded p-bit dyadic number is zero or a number $\pm m2^e$, where $m, e \in \mathbb{Z}$, $2^{p-1} \leq m < 2^p$, and $|e| \leq E$. Every point operation $+, -, \times, \div, \sqrt{\ }$ returns the nearest such dyadic to the exact result, with ties resolved toward an even significand. Put*

$$u_p = 2^{1-p}, \qquad \tau_{p,E} = 2^{p-1-E}.$$

Absent exponent overflow, the point-rounding map satisfies

$$|\operatorname{rnd}_{p,E}(s) - s| \leq u_p|s| + \tau_{p,E}.$$

A result whose magnitude is below the least positive representable magnitude may underflow to zero; this is not a failure. A zero divisor, a negative square-root argument, or exponent overflow returns FAIL.

For interval certificates, lower and upper endpoints are rounded outward. They enclose the exact result, and each endpoint differs from it by at most $2u_p$ times its magnitude plus $\tau_{p,E}$ while the exponent guard is respected; interval division through zero returns FAIL. *For an interval $I = [\ell, u]$, write* $\inf I = \ell$, $\sup I = u$, *and* $\operatorname{width}(I) = u - \ell$. *From a certified interval, we select* $\operatorname{rnd}_{p,E}((\ell+u)/2)$ *and use the larger endpoint distance as a certified radius; for a box this rule is applied coordinatewise, with an upward-rounded Euclidean norm of the radii.*

An exact finite dyadic input longer than p bits may be used in its first rounded operation when the input-size and exponent-margin conditions of the routine hold; the result is guarded and obeys the same point-rounding bound. Block-Sturm computations use $E_{\mathrm{blk}} = K_w p$, while Givens and triangular-solve computations use $E_{\mathrm{sol}} = K_w np$. On any failure, a routine returns its specified finite zero output.

3 Overview

The proof of Theorem 1.1 analyzes BandedPSD-Eigenpairs. On input (A, k, ϵ, δ), this algorithm produces normalized dyadic encodings of k approximate top eigenvectors and nonnegative dyadic

Rayleigh estimates obtained from certified intervals. Two issues drive the construction. First, A need not have any eigengap, whereas isolating individual eigenvectors with shifted-inverse filters requires quantitative separation. Second, finite-precision Sturm localization must avoid every leading-principal spectrum, and all subsequent spectral operations must be certified at a common precision while retaining a cost essentially linear in both n and k. The proof resolves these issues by creating a harmless gap, locating the perturbed eigenvalues with safe inertia queries, isolating them with a bank of shifted-inverse filters, and certifying the resulting vectors with banded finite-precision computations.

Creating a usable eigengap. We measure the matrix at the row-sum scale $W = \max_i \sum_j |A_{ij}|$, which satisfies $\|A\|_2 \le W \le (2w+1)\|A\|_2$. We then sample a finite dyadic diagonal matrix D and put

$$H = A + D, \qquad 0 \preceq D \preceq \rho I, \qquad \rho \le \frac{\epsilon W}{128}.$$

The finite-dimensional Minami estimate first shows that a continuous diagonal perturbation has pairwise eigenvalue gaps at least 4Δ with high probability. Sampling instead on a dyadic grid of spacing $h \le \Delta/8$ changes every eigenvalue by less than h, so Weyl's bound preserves a gap of at least 3Δ while using only $O(n \log(n/(\epsilon\delta)))$ random bits. This perturbation is introduced only to split spectral clusters: because $\|D\|_2 \le \rho$, its effect fits within the final error budget. In particular, the proof later transfers residuals and indexed eigenvalue estimates from H to A rather than claiming that a chosen eigenbasis of a clustered A is preserved.

Stable spectral localization. Fixed bandwidth makes $H - xI$ block tridiagonal with constant-size blocks. When its leading Schur pivots are nonsingular, the block Sturm recurrence reduces eigenvalue counting to

$$P_1 = B_1 - xI, \qquad P_j = B_j - xI - C_{j-1}^T P_{j-1}^{-1} C_{j-1}, \qquad \iota_+(H - xI) = \sum_j \iota_+(P_j) = |\{i : \lambda_i(H) > x\}|.$$

The numerical difficulty is that a query can be separated from $\mathrm{spec}(H)$ but extremely close to the spectrum of a leading principal block, making one of these Schur pivots nearly singular. RandomizedBandedBisect therefore chooses each query from a finite dyadic lattice in the middle half of its current interval. With high probability all adaptive queries avoid every leading-principal spectrum. At such a safe query, ExactSmallBlock computes the inertia of each rounded constant-size pivot exactly and its inverse to guarded precision; the collection of rounded pivots is the exact recurrence for a perturbation too small to change the full inertia. Each query consequently gives the correct count and shrinks the interval by a constant factor. The algorithm first localizes $\lambda_k(H)$ to set a rank-k threshold, and localizes the other top eigenvalues because each will center one spectral slice below.

A multishift spectral step. On the gap and safe-localization event, write $|\widetilde{\lambda}_i - \lambda_i(H)| \le a$ and set

$$z_i = \widetilde{\lambda}_i + \frac{\Delta}{2}, \qquad F_i = \left[-\frac{\Delta}{2}(H - z_i I)^{-1}\right]^d, \qquad \gamma = \widetilde{\lambda}_k - \Delta, \qquad S_\gamma = \sum_{i=1}^{k} F_i.$$

Here d is the least positive even integer with $4^{-d} \le \zeta/(128n)$, where $\zeta \le n^{-10}$, and the localization accuracy is chosen so that $a \le \zeta\Delta/(2^{10}d)$. The threshold γ has exactly k eigenvalues of H above it. In the eigenbasis of H, the multiplier of F_i at $\lambda_j(H)$ is $(-\Delta/(2(\lambda_j(H) - z_i)))^d$. It differs from 1 by at most $\zeta/128$ when $j = i$, while the 3Δ gap makes its magnitude at most $4^{-d} \le \zeta/(128n)$

when $j \neq i$. Thus, writing $u_i = u_i(H)$,

$$\|F_i - u_i u_i^T\|_2 \leq \frac{\zeta}{128}, \qquad \left\|S_\gamma - \mathbf{1}_{[\gamma,\infty)}(H)\right\|_2 \leq \zeta \leq n^{-10}.$$

The sum S_γ is therefore the desired approximate step function, while its individual summands are the nearly rank-one objects used to recover the eigenvectors.

One-vector recovery and certified inverse powers. Independently of H, the algorithm samples a finite random integer vector r and uses its normalized encoding $x = r/\|r\|_2$ when $r \neq 0$, with $x = e_1$ when $r = 0$. A discrete anti-concentration argument gives $|u_i(H)^T x| \geq \beta$ simultaneously for the entire eigenbasis with probability at least $1 - \delta/16$. For each i, ShiftedInversePower takes $(H, W, z_i, -\Delta/2, \Delta, \zeta, d, x, p)$, factors the banded matrix $H - z_i I$ once by Givens QR, and performs d banded triangular solves. It returns a dyadic vector y_i and a certified radius R_i satisfying $\|y_i - F_i x\|_2 \leq R_i \leq \zeta/128$, using $O_w(ndp^2)$ bit operations. Its central error certificate propagates the directed residual radii $\widehat{r}_\ell, \widehat{s}_\ell$ through

$$\eta_\ell = \widehat{s}_\ell + g_+ \widehat{r}_\ell, \qquad R_\ell = g_+ R_{\ell-1} + \eta_\ell.$$

Here g_+ is a directed upper rounding of $1 + \zeta/(256d)$. This converts local solve residuals into an end-to-end bound for the power $F_i x$. Writing $\alpha_i = u_i(H)^T x$, the overlap event gives $|\alpha_i| \geq \beta$, while the exact-filter and numerical-application bounds give $\|y_i - \alpha_i u_i(H)\|_2 \leq \zeta/64$. The choice $\zeta \leq \epsilon\beta/64$ therefore makes the normalized encoding $q_i = y_i/\sqrt{y_i^T y_i}$ satisfy $\|q_i - \operatorname{sgn}(\alpha_i) u_i(H)\|_2 \leq \epsilon/8$. Recovering the vectors from the individual slices avoids constructing and orthogonalizing a dense $n \times k$ range basis.

Transfer, certification, and cost. Closeness to $u_i(H)$ first gives $\|Hq_i - \lambda_i(H) q_i\|_2 \leq \epsilon W/2$; since $A = H - D$,

$$\|Aq_i - \lambda_i(H) q_i\|_2 \leq \|Hq_i - \lambda_i(H) q_i\|_2 + \|D\|_2 < \epsilon W.$$

Directed intervals for the homogeneous Rayleigh quotient $y_i^T A y_i/(y_i^T y_i)$ then produce the stored nonnegative dyadic $\widehat{\sigma}_i$. The preceding residual, Weyl's bound, and the norm comparison in Equation (8) give the claimed residual and indexed eigenvalue errors for A. Because the $u_i(H)$ are orthonormal, the same vector approximation gives both the pairwise inner-product bound and $\|Q^T Q - I\|_2 \leq C_w k\epsilon$.

Finally, all scalar parameters have $O(L)$-bit encodings, the common precision is $p = O_w(L)$, and both the bisection depth and filter degree are $O(L)$. A fixed-bandwidth inertia query or shifted solve uses $O_w(np^2)$ bit operations, so the k localizations and k filter applications are bounded by $O_w(nkL^3)$ bit operations. Perturbation, overlap, and query sampling use $O(nL + kL^2)$ random bits, and the normalized vector records use $O(nkL)$ output bits. The gap, safe-query, and overlap events hold jointly with probability at least $1 - 9\delta/64$; on this event every claimed guarantee holds, while explicit guards and finite fallback records make the algorithm terminate on every random branch.

4 Proof of the Main Theorem

The case $n = 1$ is immediate, so until the final algorithm we assume $n \geq 2$. We now fix the dyadic scales used throughout the proof: ρ and Δ control the perturbation and its induced eigengap, β and ζ control random overlap and filter accuracy, and a and ν control localization accuracy and query safety. Let ρ be the largest power of two not exceeding $\epsilon W/128$, and let Δ be the largest power of

two not exceeding $\delta\rho^2/(2^{16}Wn^2)$. Let β be the largest power of two not exceeding $\delta/(32n^2)$, and let ζ be the largest power of two not exceeding $\min\{n^{-10}, \epsilon\beta/64\}$. Thus

$$\frac{\epsilon W}{256} < \rho \le \frac{\epsilon W}{128}, \tag{9}$$

$$\frac{\delta\rho^2}{2^{17}Wn^2} < \Delta \le \frac{\delta\rho^2}{2^{16}Wn^2}, \qquad \frac{\delta}{64n^2} < \beta \le \frac{\delta}{32n^2}, \tag{10}$$

$$\frac{1}{2}\min\left\{n^{-10}, \frac{\epsilon\beta}{64}\right\} < \zeta \le \min\left\{n^{-10}, \frac{\epsilon\beta}{64}\right\}. \tag{11}$$

Let d be the least positive even integer such that

$$4^{-d} \le \frac{\zeta}{128n}, \tag{12}$$

and let a be the largest power of two not exceeding $\zeta\Delta/(2^{10}d)$. Hence

$$\frac{\zeta\Delta}{2^{11}d} < a \le \frac{\zeta\Delta}{2^{10}d}. \tag{13}$$

For the randomized inertia bisections define

$$m_{\mathrm{bis}} = \left\lceil \log_{4/3}\frac{2W}{a} \right\rceil + 2, \qquad N_{\mathrm{qry}} = (k+1)m_{\mathrm{bis}}. \tag{14}$$

Finally, let ν be the largest power of two not exceeding $\min\{1, \delta a/(2^{20}n^2N_{\mathrm{qry}})\}$. Then

$$\frac{1}{2}\min\left\{1, \frac{\delta a}{2^{20}n^2N_{\mathrm{qry}}}\right\} < \nu \le \min\left\{1, \frac{\delta a}{2^{20}n^2N_{\mathrm{qry}}}\right\}. \tag{15}$$

These definitions give

$$\log_2(W/\Delta) + \log_2(1/\zeta) + \log_2(W/a) + \log_2(W/\nu) = O\Big(\log\frac{n}{\epsilon\delta}\Big), \tag{16}$$

$$d + m_{\mathrm{bis}} = O\Big(\log\frac{n}{\epsilon\delta}\Big). \tag{17}$$

We first create the gap used by the multishift construction. The only probabilistic input is the following finite-dimensional two-level estimate.

Lemma 4.1 (Finite-dimensional Minami estimate)**.** *Let T be a fixed real symmetric $n \times n$ matrix, let $\xi_1, \dots, \xi_n$ be independent real random variables with a common density bounded by K, and put $T_\xi = T + \mathrm{diag}(\xi_1, \dots, \xi_n)$. For every interval $J \subseteq \mathbb{R}$,*

$$\Pr\{\operatorname{tr}\mathbf{1}_J(T_\xi) \ge 2\} \le \frac{\pi^2}{2}(Kn|J|)^2.$$

Proof. This is the $r = 2$ finite-volume specialization of the correlation estimate of Bellissard, Hislop, and Stolz [BHS07, Theorem 4]. To obtain the displayed finite-dimensional form, embed T as the finite restriction of the bounded background operator and apply their estimate on the coordinate interval $\{1, \dots, n\}$. Their bound is $\Pr\{\operatorname{tr}\mathbf{1}_J(T_\xi) \ge r\} \le \pi^r(Kn|J|)^r/r!$; substituting $r = 2$ gives the displayed inequality. □

The continuous perturbation supplied by the preceding estimate must be rounded without losing its gap. The next lemma carries out that finite-bit step and bounds both its randomness and its norm.

Lemma 4.2 (Finite-bit diagonal gap perturbation)**.** *There is a distribution on diagonal dyadic matrices D, samplable with $O(n\log(n/(\epsilon\delta)))$ random bits, such that $0 \preceq D \preceq \rho I$ always and, with probability at least $1-\delta/16$, the eigenvalues of $H = A + D$ satisfy*

$$|\lambda_i(H) - \lambda_j(H)| \geq 3\Delta \qquad (i \neq j).$$

Proof. Let $\xi_1, \ldots, \xi_n$ be independent uniform random variables on $[0, \rho]$, and set $H_c = A + \mathrm{diag}(\xi_1, \ldots, \xi_n)$. Its spectrum lies in $[0, W + \rho] \subseteq [0, 2W]$. If two eigenvalues are at distance less than 4Δ, they lie together in one interval of the overlapping cover

$$[4j\Delta, 4(j+2)\Delta], \qquad 0 \leq j \leq \left\lceil \frac{2W}{4\Delta} \right\rceil.$$

There are at most $W/\Delta + 2$ such intervals, each of length 8Δ. Applying Lemma 4.1 with $K = 1/\rho$ and taking a union bound gives

$$\begin{aligned}\Pr\left\{\min_{i\neq j} |\lambda_i(H_c) - \lambda_j(H_c)| < 4\Delta\right\} &\leq \left(\frac{W}{\Delta} + 2\right)\frac{\pi^2}{2}\left(\frac{8n\Delta}{\rho}\right)^2 \\ &\leq 64\pi^2 \frac{Wn^2\Delta}{\rho^2} \leq \frac{\pi^2\delta}{2^{10}} < \frac{\delta}{16},\end{aligned}$$

where the parameter bounds give $W/\Delta > 2$, so $W/\Delta + 2 \leq 2W/\Delta$, and Equation (10) is used in the last line.

Choose the least s such that $h = \rho 2^{-s} \leq \Delta/8$. Sample independent uniform $J_i \in \{0, \ldots, 2^s - 1\}$ and set $D_{ii} = hJ_i$. Couple this sample to the continuous one by writing $\xi_i = h(J_i + U_i)$ with independent U_i uniform on $[0, 1)$. Then $\|D - \mathrm{diag}(\xi_i)\|_2 < h$. By Weyl's inequality in Equation (2), each eigenvalue moves by less than h, so a 4Δ gap of H_c becomes a gap larger than $4\Delta - 2h > 3\Delta$ for $H = A + D$. Finally, $s = O(\log(n/(\epsilon\delta)))$ by Equations (9) and (10). □

We now specify the finite-precision primitives using the block partition and guarded arithmetic fixed in the preliminaries.

For later reference, the common sufficient precision condition is

$$u_p \leq \min\left\{\frac{\nu^2}{2^{20}K_w W^2}, \frac{\zeta\Delta}{2^{20}K_w n^2 dW}, \frac{1}{2^{20}K_w n^2}\right\}. \tag{18}$$

The first primitive resolves constant-size singular-pivot ambiguities in the block recurrence by exact rational arithmetic.

EXACT**SmallBlock**: Exact inversion and inertia of a constant-size symmetric block.
Input: A symmetric dyadic matrix P of dimension at most b whose entries are guarded p-bit numbers, followed by an integer precision $p \geq 8$; the exponent guard is $E_{\mathrm{blk}} = K_w p$.
Output: Either $(\mathrm{OK}, X, \iota_+(P))$, where $\iota_+(P)$ is the exact number of positive eigenvalues of P and X approximates P^{-1}, or $(\mathrm{FAIL}, 0, 0)$.
Guarantee: An OK result has the exact inertia and satisfies

$$\|X - P^{-1}\|_2 \leq K_w\big(u_p\|P^{-1}\|_2 + \tau_{p,E_{\mathrm{blk}}}\big).$$

1. Interpret the entries of P as exact rationals. Compute $\det(P)$ and $\mathrm{adj}(P)$ from their finite permutation formulas. If $\det(P) = 0$, or an exact numerator or denominator exceeds $K_w^2 p$ bits, return $(\mathrm{FAIL}, 0, 0)$. Otherwise form $P^{-1} = \mathrm{adj}(P)/\det(P)$ exactly and round its entries to obtain X; an exponent-guard violation returns failure.

2. Compute inertia by exact symmetric elimination. At each stage, symmetrically permute the least-indexed nonzero diagonal entry into the first position and use it as a 1-by-1 pivot. If all diagonal entries vanish, use the lexicographically first nonzero off-diagonal entry c in the pivot $\left(\begin{smallmatrix} 0 & c \\ c & 0 \end{smallmatrix}\right)$. Add the positive inertia of the pivot and recurse on its exact rational Schur complement. Return failure if the bit guard is exceeded; otherwise return the accumulated count together with X.

Lemma 4.3 (Exact constant-size block kernel)**.** EXACTSMALLBLOCK *is total, never returns status* OK *with an incorrect positive inertia, satisfies its stated inverse-error bound whenever it returns* OK*, and uses $O_w(p^2)$ bit operations. For the safe block-Sturm pivots below, its status is* OK*.*

Proof. A nonsingular real symmetric matrix always admits the selected pivot. If no diagonal entry is nonzero, some off-diagonal entry c is nonzero and the displayed 2-by-2 pivot has determinant $-c^2$. A nonsingular principal pivot in a nonsingular matrix has a nonsingular Schur complement. Repeated congruence and inertia additivity therefore prove that the returned inertia is exact. The determinant–adjugate identity proves the inverse formula, and the rounding bound in Definition 2.1, in the fixed dimension b, gives the stated inverse error, including underflow. The $K_w^2 p$-bit cutoff makes every branch finite. Only $O_w(1)$ rational operations are performed, so schoolbook arithmetic on the permitted operands costs $O_w(p^2)$ bit operations. Success on safe pivots follows from the magnitude estimates in Lemma 4.4. □

The next primitive uses these exact block inertias. Its query is chosen from a finite random lattice in the middle half of the current interval, so with high probability it remains separated from every leading-principal spectrum. For an interval $I = [\ell, u]$ of length $h_I > 2a$, define

$$j^-(I) = \left\lceil \frac{8(\ell + h_I/4)}{\nu} \right\rceil, \qquad j^+(I) = \left\lfloor \frac{8(\ell + 3h_I/4)}{\nu} \right\rfloor, \tag{19}$$

$$R(I) = j^+(I) - j^-(I) + 1, \qquad \mathcal{X}(I) = \left\{ \frac{\nu}{8}(j^-(I) + J) : 0 \leq J < 2^{\lfloor \log_2 R(I) \rfloor} \right\}. \tag{20}$$

A query is safe for the bisection when it is 2ν-safe for T in the sense of Equation (4). If B_j and C_j are the diagonal and adjacent off-diagonal blocks of T, respectively, the point Schur-pivot recurrence used at query x, the rounded counterpart of Equation (6), is

$$\widehat{P}_1 = \mathrm{rnd}_{p,E_{\mathrm{blk}}}(B_1 - xI), \qquad \widehat{P}_j = \mathrm{rnd}_{p,E_{\mathrm{blk}}}(B_j - xI - C_{j-1}^T X_{j-1} C_{j-1}), \tag{21}$$

where X_{j-1} is the certified rounded inverse returned for the preceding pivot. Call a tuple $(T, i, W, a, \nu, m_{\mathrm{bis}}, p)$ *admissible* if p and m_{bis} are integers with $p \geq 8$, W, a, ν are finite dyadics, $W \geq 1$, T is real symmetric of half-bandwidth w with spectrum in $[0, 2W]$, $i \in \{1, \ldots, n\}$, $0 < \nu \leq \min\{a, 1\} \leq W$, and

$$m_{\mathrm{bis}} \geq \left\lceil \log_{4/3} \frac{W}{a} \right\rceil + 1,$$

with every dyadic scalar and nonzero entry of T using at most $K_w p$ bits and exponent magnitude at most $E_{\mathrm{blk}}/16$, and every integer input using at most $K_w p$ bits.

RandomizedBandedBisect: Guarded Sturm bisection for one ordered eigenvalue.
Input: An admissible tuple $(T, i, W, a, \nu, m_{\mathrm{bis}}, p)$ as defined above.
Output: A pair $(\mathrm{status}, \widetilde{\lambda}_i)$ with status $\in \{\mathrm{OK}, \mathrm{FAIL}\}$; the failure output is $(\mathrm{FAIL}, 0)$.
Guarantee: Under Equation (18), if every query is 2ν-safe, then the status is OK and $|\widetilde{\lambda}_i - \lambda_i(T)| \leq a$.

1. Initialize $[\ell_0, u_0] = [0, 2W]$.
2. For $r = 0, \ldots, m_{\mathrm{bis}} - 1$, return the midpoint with status OK if $u_r - \ell_r \leq 2a$; otherwise sample x_r uniformly from $\mathcal{X}([\ell_r, u_r])$. In block order form Equation (21) with $x = x_r$ and set
$$(s_{r,j}, X_j, \iota_{r,j}) = \textsc{ExactSmallBlock}(\widehat{P}_j, p).$$
If some $s_{r,j} = \mathrm{FAIL}$, return $(\mathrm{FAIL}, 0)$. Set $N_r = \sum_j \iota_{r,j}$ and update to $[x_r, u_r]$ when $N_r \geq i$, or to $[\ell_r, x_r]$ otherwise; name the selected interval $[\ell_{r+1}, u_{r+1}]$.
3. After the last update, return the midpoint with status OK if the width is at most $2a$, and return $(\mathrm{FAIL}, 0)$ otherwise.

Lemma 4.4 (Stable randomized banded Sturm bisection)**.** *Under Equation (18),* RandomizedBandedBisect *has its stated output guarantee and uses $O_w(n m_{\mathrm{bis}} p^2)$ bit operations and $O(m_{\mathrm{bis}} \log(W/\nu))$ random bits on every branch. With the global parameters in Equations (10) and (15), all at most N_{qry} adaptive queries are simultaneously safe with probability at least $1 - \delta/64$.*

Proof. Fix one invocation and condition on its adaptive history. If the current interval has length $h_r > 2a$, the full middle-half lattice contains at least $4h_r/\nu - 1$ points. Its retained power-of-two prefix has at least $2h_r/\nu - 1/2 \geq 7h_r/(4\nu)$ points. At most 33 lattice points can be within 2ν of one fixed real number. Thus at most a $64\nu/a$ fraction of the retained points is unsafe for it. The sum of the dimensions of all leading block matrices is at most n^2, so the conditional probability that one query lies within 2ν of any eigenvalue of any $T^{[j]}$ is at most $64n^2\nu/a$. The adaptive union bound and Equation (15) give

$$\Pr\{\text{some one of the at most } N_{\mathrm{qry}} \text{ queries is unsafe}\} \leq \frac{64n^2 N_{\mathrm{qry}} \nu}{a} < \frac{\delta}{64}.$$

Now fix a safe query x. For the point blocks generated by the routine, define exact symmetric residual blocks

$$E_1 = \widehat{P}_1 - (B_1 - xI), \qquad E_j = \widehat{P}_j - (B_j - xI - C_{j-1}^T \widehat{P}_{j-1}^{-1} C_{j-1}).$$

They include inverse and product rounding. Inductively, the $\widehat{P}_j$ are exactly the block LDL^T pivots of $T + E - xI$, with E block diagonal. The first residual satisfies

$$\|E_1\|_2 \le K_w(u_p W + \tau_{p,E_{\mathrm{blk}}}) < \nu/2^{18}.$$

Safety for $T^{[1]}$ makes $\widehat{P}_1$ nonsingular with inverse norm below ν^{-1} and ensures that its constant-size kernel succeeds. Suppose the claim holds through block $j-1$. Safety and Weyl's inequality put $T^{[j-1]} + E^{[j-1]} - xI$ at distance at least $31\nu/16$ from singularity. Its lower-right inverse block is $\widehat{P}_{j-1}^{-1}$, and therefore $\|\widehat{P}_{j-1}^{-1}\|_2 < \nu^{-1}$. The inverse certificate and the rounded block product give

$$\|E_j\|_2 \le K_w\left(u_p\frac{W^2}{\nu} + \tau_{p,E_{\mathrm{blk}}}W^2\right) + K_w\tau_{p,E_{\mathrm{blk}}} < \frac{\nu}{2^{18}}.$$

Safety at the next leading principal matrix now gives

$$\|\widehat{P}_j^{-1}\|_2 < \nu^{-1}, \qquad \|\widehat{P}_j\|_2 \le K_w(W + W^2/\nu).$$

These bounds, Equation (18), and fixed block dimension show that all exact rational intermediates have $O_w(p)$ bits, so the block kernel succeeds and the induction closes. Consequently $\|E\|_2 < \nu/16$.

Equations (5) and (7) say that the computed count is the number of eigenvalues of $T+E$ above x. Since $T - xI$ is at distance at least 2ν from singularity, Weyl's inequality Equation (2) shows that this is also the exact count for T. Every update therefore retains $\lambda_i(T)$. A query in the middle half reduces the interval length by at least a factor $3/4$, and the assumed iteration bound gives final width at most $2a$. This proves the output guarantee.

The block recurrence performs $O_w(n)$ guarded scalar operations per query, each costing $O(p^2)$ bit operations under schoolbook integer arithmetic. The candidate lattice uses $O(\log(W/\nu))$ random bits per query. All loops and failure branches have the stated worst-case bounds. □

We next certify repeated application of a shifted inverse. For an input matrix T, shift z, and scalar c, define

$$B = T - zI, \qquad G = cB^{-1}. \tag{22}$$

Given a point solution v_ℓ to $Bv = y_{\ell-1}$, first let y_ℓ be the point rounding of cv_ℓ. Then let directed banded matrix-vector evaluation produce coordinate intervals $J_{\ell,j} \ni (Bv_\ell - y_{\ell-1})_j$ and $K_{\ell,j} \ni (y_\ell - cv_\ell)_j$. With ru_p denoting upward rounding under E_{sol}, define

$$\widehat{r}_\ell = \mathrm{ru}_p\left(\sqrt{\sum_{j=1}^{n}\max\{|\inf J_{\ell,j}|, |\sup J_{\ell,j}|\}^2}\right), \tag{23}$$

$$\widehat{s}_\ell = \mathrm{ru}_p\left(\sqrt{\sum_{j=1}^{n}\max\{|\inf K_{\ell,j}|, |\sup K_{\ell,j}|\}^2}\right). \tag{24}$$

These quantities certify

$$\widehat{r}_\ell \ge \|Bv_\ell - y_{\ell-1}\|_2, \qquad \widehat{s}_\ell \ge \|y_\ell - cv_\ell\|_2. \tag{25}$$

Let g_+ be the directed upper rounding of $1 + \zeta/(256d)$ and define

$$\eta_\ell = \widehat{s}_\ell + g_+\widehat{r}_\ell, \qquad R_\ell = g_+R_{\ell-1} + \eta_\ell. \tag{26}$$

These objects turn local residuals into an end-to-end error certificate for G^d. Call $(T, W, z, c, \Delta, \zeta, d, x, p)$ an *admissible shifted-inverse input* if p and d are integers with $p \ge 8$ and $d \ge 1$, W, z, c, Δ, ζ

are finite dyadics with $W \geq 1$, T is real symmetric of half-bandwidth w with $\|T\|_2 \leq 2W$, x is a normalized dyadic-vector encoding, every dyadic scalar, nonzero entry of T, and coordinate of the representative for x has at most $K_w p$ bits and exponent in $[-E_{\mathrm{sol}}/16, E_{\mathrm{sol}}/16]$, the integer d has at most $K_w p$ bits, and

$$0 < \Delta \leq W, \quad 0 < \zeta \leq 1, \quad |c| \leq \Delta, \quad |z| \leq 3W, \quad \operatorname{dist}(z, \operatorname{spec}(T)) \geq \Delta/3, \quad \|G\|_2 \leq 1 + \frac{\zeta}{256d}.$$

SHIFTED**InversePower**: Certified application of an integer power of a shifted inverse.
Input: An admissible tuple $(T, W, z, c, \Delta, \zeta, d, x, p)$ as defined above.
Output: A triple (status, y, R), where y is a dyadic vector and R is a nonnegative dyadic error radius; the failure output is $(\mathrm{FAIL}, 0, 0)$.
Guarantee: Under Equation (18), the status is OK and $\|y - G^d x\|_2 \leq R \leq \zeta/128$.

1. Let r^{enc} be the nonzero dyadic representative of x. Rescale it by a common exact power of two so that its largest coordinate has magnitude in $[1/2, 1)$. Using directed arithmetic, enclose $r^{\mathrm{enc}}/\sqrt{(r^{\mathrm{enc}})^T r^{\mathrm{enc}}}$ in a coordinate box, select its midpoint y_0, and let R_0 be the certified Euclidean radius. If $R_0 > \zeta/2048$, return failure.

2. Form B from Equation (22). Compute a point-valued banded Givens QR factorization of B: after a common exact power-of-two rescaling, round each hypotenuse, cosine, and sine to nearest, apply the point rotation, using the identity rotation when both entries are zero, and set the targeted subdiagonal entry to zero. Store the $O_w(n)$ rotations and the resulting upper-banded matrix $\widehat{R}$.

3. For $\ell = 1, \ldots, d$, apply the stored rotations to $y_{\ell-1}$ and solve the upper-banded system with $\widehat{R}$ by point back substitution, obtaining v_ℓ. Compute $\widehat{r}_\ell, \widehat{s}_\ell, y_\ell$ from Equations (24) and (25), and update η_ℓ, R_ℓ by Equation (26). Return failure if $\eta_\ell > \zeta/(4096d)$, $R_\ell > \zeta/128$, a divisor is zero, or an exponent guard is violated.

4. Return (OK, y_d, R_d).

Lemma 4.5 (Certified shifted-inverse powering)**.** *Under Equation (18),* SHIFTEDINVERSEPOWER *has its stated output guarantee and uses $O_w(ndp^2)$ bit operations on every branch. Successful factorization and solve quantities have exponent magnitude $O_w(np)$.*

Proof. The promises on B imply

$$\|B\|_2 \leq 5W, \qquad \|B^{-1}\|_2 \leq 3/\Delta, \qquad \kappa_2(B) \leq 15W/\Delta.$$

We first bound the point-valued QR solve. One rounded hypotenuse and its two quotients perturb a 2-by-2 rotation by at most $K_w u_p$ in norm. Applying it to the $O_w(1)$ live entries in two banded rows introduces a residual at most $K_w u_p$ times their norm plus $K_w \tau_{p,E_{\mathrm{sol}}}$. Telescoping over $O_w(n)$ rotations and normalizing each computed cosine–sine pair to an exact rotation gives

$$\widehat{R} = \widehat{Q}^T(B + \Delta B), \qquad \widehat{g} = \widehat{Q}^T(g + \delta g),$$

where $\widehat{Q}$ is exactly orthogonal and

$$\|\Delta B\|_2 \leq K_w n^2 (u_p W + \tau_{p,E_{\mathrm{sol}}}),$$

$$\|\delta g\|_2 \le K_w n^2 (u_p \|g\|_2 + \tau_{p,E_{\mathrm{sol}}}).$$

Each back-substitution row has $O_w(1)$ live entries. Assigning relative rounding to the corresponding coefficient or right-hand-side entry, and charging an underflowed quotient after multiplication by its diagonal, gives

$$(\widehat{R} + \Delta R)\widehat{v} = \widehat{g} + \delta h,$$

with

$$\begin{aligned} \|\Delta R\|_2 &\le K_w n^2 u_p W, \\ \|\delta h\|_2 &\le K_w n^2 \big(u_p \|g\|_2 + W \tau_{p,E_{\mathrm{sol}}}\big). \end{aligned}$$

The input exponent margin implies

$$W \tau_{p,E_{\mathrm{sol}}} \le u_p.$$

Indeed, $\log_2 W \le K_w p + E_{\mathrm{sol}}/16 \le 9E_{\mathrm{sol}}/16$ for $n \ge 2$, and hence

$$\log_2 \frac{W \tau_{p,E_{\mathrm{sol}}}}{u_p} \le 2p - 2 - \frac{7E_{\mathrm{sol}}}{16} < 0.$$

Multiplication by $\widehat{Q}$ yields the local backward-error identity

$$(B + F)\widehat{v} = g + f, \qquad \|F\|_2 \le 4K_w n^2 u_p W, \qquad \|f\|_2 \le 2K_w n^2 \big(u_p \|g\|_2 + W \tau_{p,E_{\mathrm{sol}}}\big). \tag{27}$$

By Equation (18), $\|B^{-1}F\|_2 < 2^{-12}$. Whenever $\|g\|_2 \le 2$, Equation (27) implies

$$\|\widehat{v}\|_2 \le \frac{12}{\Delta}, \qquad \|B\widehat{v} - g\|_2 \le \frac{54\zeta}{2^{20}d} < \frac{\zeta}{2^{14}d}.$$

A directed evaluation of the fixed-bandwidth residual, and point rounding of $c\widehat{v}$, add no more than the same amount. Thus the certificates satisfy

$$\widehat{r}_\ell \le \frac{\zeta}{2^{13}d}, \qquad \widehat{s}_\ell \le \frac{\zeta}{2^{14}d}, \qquad \eta_\ell \le \frac{\zeta}{4096d}. \tag{28}$$

The same directed error summation in the normalization step gives $R_0 \le K_w n(u_p + \tau_{p,E_{\mathrm{sol}}}) \le \zeta/2048$.

We prove the radius certificate by induction. At stage ℓ, the exact identity

$$y_\ell - G y_{\ell-1} = (y_\ell - c v_\ell) + c B^{-1}(B v_\ell - y_{\ell-1})$$

and $\|cB^{-1}\|_2 \le 1 + \zeta/(256d) \le g_+$ give $\|y_\ell - G y_{\ell-1}\|_2 \le \eta_\ell$. Therefore $\|y_\ell - G^\ell x\|_2 \le R_\ell$ whenever the preceding certificate is valid. Moreover,

$$\|y_{\ell-1}\|_2 \le \|G^{\ell-1}x\|_2 + R_{\ell-1} \le e^{1/256} + \zeta/128 < 2,$$

so the solve estimate used above remains valid at every stage. Directed rounding gives $g_+ \le 1 + \zeta/(256d) + 4u_p$, while Equation (18) gives $du_p < 2^{-20}\zeta$. Hence

$$g_+^d < 2, \qquad R_d \le 2(1 + 8u_p)^d \left(\frac{\zeta}{2048} + d\frac{\zeta}{4096d} + d\tau_{p,E_{\mathrm{sol}}} \right) < \frac{\zeta}{128}.$$

Thus no radius guard fires on a promised input.

For completeness, the rounded factorization represents a matrix within $\Delta/12$ of B, and its accumulated rotations are within 2^{-12} of an isometry. Its triangular factor therefore has smallest singular value at least $\Delta/5$, so no back-substitution diagonal vanishes. A crude reverse row recurrence places every intermediate below $(K_w W/\Delta)^{O(n)}$; its exponent has magnitude $O_w(np)$ and fits E_{sol}. There are $O_w(n)$ point operations in the factorization and each solve, and each costs $O(p^2)$ bit operations. With d solves, this is $O_w(ndp^2)$ on every branch. □

The parameter bit bounds now show that one common polylogarithmic precision supports both numerical primitives.

Corollary 4.6 (Polylogarithmic guard precision)**.** *For the global parameters above, a precision $p = O_w(L)$ satisfies the precision condition in Equation (18). Successful block quantities have exponent magnitude $O_w(p)$, successful shifted-solve quantities have exponent magnitude $O_w(np)$, and the signed exponents themselves use $O(\log n + \log p)$ bits.*

Proof. Each reciprocal on the right side of Equation (18) has $O(\log(n/(\epsilon\delta)))$ bits by Equation (16), while the exact scalar inputs have at most $O(L)$ bits. Choosing p larger than these quantities by the fixed factor implicit in K_w proves the precision assertion. The exponent bounds were established in Lemmas 4.4 and 4.5; encoding a signed integer exponent of magnitude $O_w(np)$ takes $O(\log n+\log p)$ bits. □

The localized eigenvalues now determine well-separated shifts, and the certified inverse-power routine implements their spectral slices. We define the exact filters before analyzing their numerical application.

Definition 4.7 (Shifted-inverse filter bank)**.** *Let H have eigenpairs (λ_i, u_i) in nonincreasing eigenvalue order, with every two eigenvalues separated by at least 3Δ. Suppose $|\widetilde{\lambda}_i - \lambda_i| \le a$ for $1 \le i \le k$. Define*

$$z_i = \widetilde{\lambda}_i + \frac{\Delta}{2}, \qquad F_i = \left[-\frac{\Delta}{2}(H - z_i I)^{-1}\right]^d, \qquad \gamma = \widetilde{\lambda}_k - \Delta, \qquad S_\gamma = \sum_{i=1}^{k} F_i. \tag{29}$$

Thus each F_i is a degree-d polynomial in its own shifted inverse; the sum is a multishift construction for H.

Lemma 4.8 (Spectral-step approximation and random eigenvector recovery)**.** *Under Definition 4.7, exactly k eigenvalues of H are at least γ, and*

$$\left\| S_\gamma - \mathbf{1}_{[\gamma,\infty)}(H) \right\|_2 \le \zeta \le n^{-10}.$$

Let M be the least power of two with $M \ge 16n/\delta$, sample independent r_j uniformly from $\{-M, -M+1, \dots, M-1\}$, and let $x = r/\|r\|_2$ when $r \ne 0$, with $x = e_1$ when $r = 0$. Independently of H, with probability at least $1 - \delta/16$,

$$|u_i^T x| \ge \beta \qquad (1 \le i \le n).$$

On this overlap event, if $\|\widehat{y}_i - F_i x\|_2 \le \zeta/128$ and $q_i = \widehat{y}_i/\|\widehat{y}_i\|_2$, then, after a possible sign change,

$$\|q_i - u_i\|_2 \le \epsilon/8 \qquad (1 \le i \le k). \tag{30}$$

Proof. The gap and $a < \Delta/100$ imply

$$\lambda_k - \gamma \geq \Delta - a > 0, \qquad \gamma - \lambda_{k+1} \geq 2\Delta - a > 0$$

when $k < n$; only the first inequality is needed for $k = n$. Thus the threshold has the claimed rank.

By Equation (1), the multiplier of F_i at λ_j is

$$f_i(\lambda_j) = \left(\frac{-\Delta/2}{\lambda_j - z_i}\right)^d. \tag{31}$$

For $j = i$, write $\varepsilon_i^{\mathrm{loc}} = \lambda_i - \widetilde{\lambda}_i$. Then

$$f_i(\lambda_i) = \left(1 - \frac{2\varepsilon_i^{\mathrm{loc}}}{\Delta}\right)^{-d}.$$

Since $|\varepsilon_i^{\mathrm{loc}}| \leq a \leq \zeta\Delta/(2^{10}d)$, the inequality $|(1-s)^{-d} - 1| \leq 2d|s|$ for $d|s| \leq 1/4$ gives

$$|f_i(\lambda_i) - 1| \leq \zeta/128. \tag{32}$$

If $j \neq i$, the 3Δ gap gives

$$|\lambda_j - z_i| \geq 3\Delta - \Delta/2 - a > 2\Delta, \qquad |f_i(\lambda_j)| \leq 4^{-d} \leq \frac{\zeta}{128n}.$$

The one-step operator

$$G_i = -\frac{\Delta}{2}(H - z_i I)^{-1}$$

has its target multiplier equal to $(1 - 2\varepsilon_i^{\mathrm{loc}}/\Delta)^{-1}$ and every off-target multiplier below $1/4$ in absolute value. Hence

$$\|G_i\|_2 \leq 1 + \frac{\zeta}{256d}. \tag{33}$$

Equations (12) and (32), summed over the slices, prove the spectral-step estimate. They also give the stronger individual bound

$$\|F_i - u_i u_i^T\|_2 \leq \zeta/128. \tag{34}$$

For the overlap assertion, fix a unit vector u and choose j with $|u_j| \geq n^{-1/2}$. Conditional on all r_ℓ with $\ell \neq j$, the number of values of r_j for which $|u^T r| < \beta M\sqrt{n}$ is at most $2\beta M\sqrt{n}/|u_j| + 1$. Consequently

$$\Pr\{|u^T r| < \beta M\sqrt{n}\} \leq \beta n + \frac{1}{2M} \leq \frac{\delta}{16n}.$$

Since $\|r\|_2 \leq M\sqrt{n}$, the complementary event gives $|u^T x| \geq \beta$. A union bound over the eigenbasis proves the simultaneous claim and excludes $r = 0$.

On that event put $\alpha_i = u_i^T x$. By Equation (34), $\|F_i x - \alpha_i u_i\|_2 \leq \zeta/128$. Including the numerical application error makes the right side at most $\zeta/64$. Since $|\alpha_i| \geq \beta$ and $\zeta \leq \epsilon\beta/64$, Equation (3) bounds the change in direction by $2(\zeta/64)/\beta \leq \epsilon/16$, which is stronger than Equation (30). □

The slices are already nearly rank one, so normalization replaces a dense range orthogonalization. It remains to assemble the perturbation, localization, filter application, and certified Rayleigh values into one total algorithm.

Let s be the least integer for which

$$h = \rho 2^{-s} \le \Delta/8. \tag{35}$$

For the input matrix let b_A be the largest binary encoding length of an entry, and choose the literal integer precision

$$p = \max \left\{ \begin{array}{l} 8,\ b_A + 8,\ b_\epsilon + 8,\ b_\delta + 8,\ s + 8, \\ \left\lceil 1 + \log_2 \dfrac{2^{20} K_w W^2}{\nu^2} \right\rceil, \quad \left\lceil 1 + \log_2 \dfrac{2^{20} K_w n^2 dW}{\zeta \Delta} \right\rceil, \\ \left\lceil 1 + \log_2 \dfrac{2^{20} K_w n}{\epsilon} \right\rceil, \quad \left\lceil 1 + \log_2 (2^{20} K_w n^2) \right\rceil \end{array} \right\}. \tag{36}$$

The last displayed term makes the third inequality in Equation (18) explicit; it does not alter the $O_w(\log(n/(\epsilon\delta)))$ precision bound.

The perturbation law is the one from Lemma 4.2: sample independent uniform $J_j \in \{0, \dots, 2^s - 1\}$, set

$$D = \operatorname{diag}(hJ_1, \dots, hJ_n), \qquad H = A + D. \tag{37}$$

The random vector law is the one in Lemma 4.8. Define the finite fallback record

$$\mathcal{R}_k = ((0, e_1, e_1), \dots, (0, e_k, e_k)). \tag{38}$$

We also fix the Rayleigh output rule outside the algorithm. Given a nonzero dyadic vector y_i, rescale it by a common exact power of two until its largest coordinate has magnitude in $[1/2, 1)$. Under the exponent guard $E_{\text{ray}} = 2E_{\text{sol}} + 2p$, let directed evaluation produce intervals

$$\mathcal{N}_i \ni y_i^T A y_i, \qquad \mathcal{D}_i \ni y_i^T y_i, \qquad I_i = \mathcal{N}_i / \mathcal{D}_i. \tag{39}$$

When $\inf \mathcal{D}_i > 0$ and $\operatorname{width}(I_i) \le \epsilon W/128$, define

$$\widetilde{\sigma}_i = \frac{\inf I_i + \sup I_i}{2}, \qquad \widehat{\sigma}_i = \max\{0, \widetilde{\sigma}_i\}. \tag{40}$$

Call (A, k, ϵ, δ) an *admissible eigenpair input* if A is real symmetric positive definite of fixed half-bandwidth w, its nonzero entries are integers of $O(\log n)$ bits, its spectrum is contained in $[n^{-10}, n^{10}]$, $1 \le k \le n$, and $0 < \epsilon, \delta \le 1/4$ are finite dyadics.

BandedPSD-**Eigenpairs**: Top eigenpairs of a fixed-bandwidth positive definite matrix.
Input: An admissible tuple (A, k, ϵ, δ) as defined above.
Output: The k triples $(\widehat{\sigma}_i, q_i, q_i)$, with $\widehat{\sigma}_i$ nonnegative dyadic and q_i a normalized dyadic-vector encoding.
Guarantee: With probability at least $1 - 9\delta/64$, every bound in Theorem 1.1 holds.
Totality: Every branch terminates, and any detected failure returns $\mathcal{R}_k$ from Equation (38).

1. If $n = 1$, return $((A_{11}, e_1, e_1))$. Otherwise compute $W, b_A, \rho, \Delta, \beta, \zeta, d, a, m_{\mathrm{bis}}, N_{\mathrm{qry}}, \nu, s, h$, and p from Equations (9) and (36).
2. Sample D and form the named perturbed matrix H according to Equation (37).
3. First for $i = k$, then for $i = 1, \ldots, k-1$, invoke RandomizedBandedBisect$(T = H, i = i, W = W, a = a, \nu = \nu, m_{\mathrm{bis}} = m_{\mathrm{bis}}, p = p)$ with fresh randomness, obtaining $(s_i, \widetilde{\lambda}_i)$. If $s_i = \mathrm{FAIL}$, return $\mathcal{R}_k$; after the first call set $\gamma = \widetilde{\lambda}_k - \Delta$.
4. Compute M, sample r, and form its normalized encoding x by the law in Lemma 4.8. For each $i = 1, \ldots, k$, set $z_i = \widetilde{\lambda}_i + \Delta/2$ and invoke ShiftedInversePower$(T = H, W = W, z = z_i, c = -\Delta/2, \Delta = \Delta, \zeta = \zeta, d = d, x = x, p = p)$, obtaining $(s_i^{\mathrm{inv}}, y_i, R_i^{\mathrm{filt}})$. If $s_i^{\mathrm{inv}} = \mathrm{FAIL}$, $R_i^{\mathrm{filt}} > \zeta/128$, or $y_i = 0$, return $\mathcal{R}_k$; otherwise store y_i as the normalized encoding of q_i.
5. For every i, evaluate the intervals in Equation (39). If evaluation fails, $\inf \mathcal{D}_i \le 0$, or $\mathrm{width}(I_i) > \epsilon W/128$, return $\mathcal{R}_k$. Otherwise compute $\widehat{\sigma}_i$ by Equation (40) and return the k triples.

Lemma 4.9 (Correctness and bit complexity of the banded eigensolver)**.** BandedPSD-Eigenpairs *has the output guarantee, totality, worst-case* $O_w(nkL^3)$ *bit complexity,* $O(nL+kL^2)$ *random-bit usage, and* $O(nkL)$ *output size stated in Theorem 1.1.*

Proof. Let $\mathcal{E}_{\mathrm{gap}}$ be the gap event in Lemma 4.2, let $\mathcal{E}_{\mathrm{bis}}$ be the event that all adaptive bisection queries are 2ν-safe, and let $\mathcal{E}_{\mathrm{ov}}$ be the overlap event in Lemma 4.8. By Lemmas 4.2 and 4.4,

$$\Pr(\mathcal{E}_{\mathrm{gap}} \cap \mathcal{E}_{\mathrm{bis}}) \ge 1 - \frac{\delta}{16} - \frac{\delta}{64} = 1 - \frac{5\delta}{64}. \tag{41}$$

The overlap vector is independent, and a union bound gives

$$\Pr(\mathcal{E}) \ge 1 - \frac{9\delta}{64} > 1 - \delta, \qquad \mathcal{E} = \mathcal{E}_{\mathrm{gap}} \cap \mathcal{E}_{\mathrm{bis}} \cap \mathcal{E}_{\mathrm{ov}}. \tag{42}$$

The choice in Equation (36) implies the sufficient condition in Equation (18). All exact encodings of D, H, the scalar parameters, bisection lattice endpoints, and shifts have at most

$$b_A + s + 8 + \log_2(W/\Delta) + \log_2(1/\zeta) + \log_2(W/a) + \log_2(W/\nu)$$

bits. By Equation (16) and the explicit terms in Equation (36), this is at most $8p$, hence at most $K_w p/32$. Their exponent magnitudes have the same margin below $E_{\mathrm{blk}}/16$ and $E_{\mathrm{sol}}/16$. The representative r and the integer d satisfy these bounds as well. Thus every primitive call meets its literal input premises.

On $\mathcal{E}$, every bisection returns an a-accurate $\widetilde{\lambda}_i$. By Lemma 4.8, the interval $[\gamma, 2W]$ contains exactly the top k eigenvalues of H and the exact filter sum satisfies

$$\left\| S_\gamma - \mathbf{1}_{[\gamma,\infty)}(H) \right\|_2 \le \zeta \le n^{-10}. \tag{43}$$

Together with Equation (41), this also proves the conditional projector guarantee independently of the overlap event. For every slice,

$$\operatorname{dist}(z_i, \operatorname{spec}(H)) \ge \Delta/2 - a > \Delta/3,$$

while $|z_i| \le 3W$, $|c| = \Delta/2$, and the power promise is precisely Equation (33). Therefore Lemma 4.5 gives a numerical application error at most $\zeta/128$. The recovery part of Lemma 4.8 then gives, after signs are chosen,

$$\|q_i - u_i(H)\|_2 \le \epsilon/8. \tag{44}$$

In particular, no y_i is zero: before numerical error, Equation (34) and $|u_i^T x| \ge \beta$ give $\|F_i x\|_2 \ge \beta - \zeta/128$, and the application error is less than $\beta/4$.

Put $\theta_i = \lambda_i(H)$. Since $\|H\|_2 \le W + \rho < 2W$, the spectral theorem and Equation (44) give

$$\|Hq_i - \theta_i q_i\|_2 \le 2\|H\|_2 \|q_i - u_i(H)\|_2 \le \epsilon W/2. \tag{45}$$

Because $A = H - D$ and $\|D\|_2 \le \rho$,

$$\|Aq_i - \theta_i q_i\|_2 \le \epsilon W/2 + \rho < \epsilon W. \tag{46}$$

We next show that every Rayleigh certificate succeeds on $\mathcal{E}$. After its exact common rescaling, $D_i = y_i^T y_i$ satisfies $1/4 \le D_i \le n$. The row-sum definition of W and $2|y_{i,r} y_{i,s}| \le y_{i,r}^2 + y_{i,s}^2$ imply

$$\sum_{r,s} |A_{rs} y_{i,r} y_{i,s}| \le W D_i.$$

Sequential directed dot products under E_{ray} enclose the numerator and denominator with endpoint errors at most

$$8K_w n(u_p + \tau_{p,E_{\text{ray}}}) W D_i \quad \text{and} \quad 8K_w n(u_p + \tau_{p,E_{\text{ray}}}) D_i,$$

respectively. Since $\tau_{p,E_{\text{ray}}} \le u_p$, the Rayleigh term in Equation (36) makes each relative error at most $\epsilon/2^{15}$. Thus $\inf \mathcal{D}_i > D_i/2 \ge 1/8$, and directed division gives $\operatorname{width}(I_i) \le \epsilon W/128$. Hence no Rayleigh fallback occurs and

$$|\widetilde{\sigma}_i - q_i^T A q_i| \le \epsilon W/256. \tag{47}$$

For the exact Rayleigh quotient $\sigma_i = q_i^T A q_i$, write $r_i = Aq_i - \theta_i q_i$. Then

$$Aq_i - \sigma_i q_i = r_i - (q_i^T r_i) q_i,$$

so Equation (46) gives $\|Aq_i - \sigma_i q_i\|_2 < 2\epsilon W$. Since $A \succeq 0$, $\sigma_i \ge 0$; clamping the approximate Rayleigh value to zero cannot increase its error. By Equation (47) and Equation (8),

$$\|Aq_i - \widehat{\sigma}_i q_i\|_2 \le C_w \epsilon \|A\|_2. \tag{48}$$

The eigenvalue estimate follows explicitly from

$$|\widehat{\sigma}_i - \lambda_i(A)| \le |\widehat{\sigma}_i - q_i^T A q_i| + |q_i^T A q_i - q_i^T H q_i|$$

$$+ |q_i^T H q_i - \lambda_i(H)| + |\lambda_i(H) - \lambda_i(A)|.$$

The four terms are at most $\epsilon W/256$, ρ, $2\|H\|_2\|q_i - u_i(H)\|_2$, and ρ, respectively; the last bound follows from Weyl's inequality in Equation (2). Equations (8), (9) and (44) therefore yield

$$|\widehat{\sigma}_i - \lambda_i(A)| \le C_w \epsilon \|A\|_2. \tag{49}$$

For $i \neq j$, the simple-spectrum eigenvectors of H are orthogonal, and Equation (44) gives

$$|q_i^T q_j| \le \|q_i - u_i(H)\|_2 + \|q_j - u_j(H)\|_2 + \|q_i - u_i(H)\|_2\|q_j - u_j(H)\|_2 \le C\epsilon.$$

The encoded columns are exactly unit vectors. Hence the maximum absolute row sum of $Q^TQ - I$ is at most $C(k-1)\epsilon$, and symmetry gives

$$\|Q^TQ - I\|_2 \le Ck\epsilon \le Cn\epsilon. \tag{50}$$

It remains to count bits. Computing the scalar parameters costs $O(nL + L^2)$ bit operations. The input assumption gives $b_A = O(\log n)$, while Equation (16) and Equation (36) give $s = O(L)$ and $p = O_w(L)$. Each of the k bisections has $O(L)$ iterations, and each of the k filters has $d = O(L)$ shifted solves. By Lemmas 4.4 and 4.5, these operations cost

$$O_w(nkLp^2) = O_w(nkL^3)$$

bit operations. The primitive bounds count every directed endpoint, division, square root, exponent-guard update, and residual-certificate operation among these $O_w(nkL)$ guarded scalar operations; each such operation costs $O(p^2)$ bits. Scalar bookkeeping and the k certified Rayleigh quotients are smaller. Perturbation and overlap sampling use $O(nL)$ random bits, while Lemma 4.4 gives $O(L^2)$ random bits for each of the k localizations. Thus the total is $O(nL + kL^2)$ random bits. All iteration counts are explicit and every exceptional arithmetic operation returns the finite fallback, so the bit-operation bound holds on every random branch.

Every returned coordinate has a p-bit significand and an exponent of magnitude $O_w(np)$. Encoding the signed exponent uses $O(\log n + \log p) = O(L)$ bits. Exact common rescaling, directed endpoints, and the dyadic midpoint preserve this bound, so the k normalized-vector records and eigenvalue records use $O(nkL)$ bits. Combining this accounting with Equation (42) proves every claim. □

Lemma 4.9 proves Theorem 1.1.

References

[ABB+26] Noah Amsel, Yves Baumann, Paul Beckman, Peter Bürgisser, Chris Camaño, Tyler Chen, Edmond Chow, Anil Damle, Michal Derezinski, Mark Embree, Ethan N. Epperly, Robert Falgout, Mark Fornace, Anne Greenbaum, Chen Greif, Diana Halikias, Zhen Huang, Elias Jarlebring, Yiannis Koutis, Daniel Kressner, Rasmus Kyng, Jörg Liesen, Jackie Lok, Raphael A. Meyer, Yuji Nakatsukasa, Kate Pearce, Richard Peng, David Persson, Eliza Rebrova, Ryan Schneider, Rikhav Shah, Edgar Solomonik, Nikhil Srivastava, Alex Townsend, Robert J. Webber, and Jess Williams. Linear systems and eigenvalue problems: Open questions from a Simons workshop, 2026. arXiv:2602.05394. 1

[Arb92] Peter Arbenz. Divide and conquer algorithms for the bandsymmetric eigenvalue problem. *Parallel Computing*, 18(10):1105–1128, 1992. 1

[BGVKS23] Jess Banks, Jorge Garza-Vargas, Archit Kulkarni, and Nikhil Srivastava. Pseudospectral shattering, the sign function, and diagonalization in nearly matrix multiplication time. *Foundations of Computational Mathematics*, 23:1959–2047, 2023. 2

[BHS07] Jean V. Bellissard, Peter D. Hislop, and Günter Stolz. Correlation estimates in the Anderson model. *Journal of Statistical Physics*, 129(4):649–662, 2007. 2, 7

[BMW67] W. Barth, R. S. Martin, and J. H. Wilkinson. Calculation of the eigenvalues of a symmetric tridiagonal matrix by the method of bisection. *Numerische Mathematik*, 9:386–393, 1967. 1

[DMPV08] James W. Demmel, Osni A. Marques, Beresford N. Parlett, and Christof Vömel. Performance and accuracy of LAPACK's symmetric tridiagonal eigensolvers. *SIAM Journal on Scientific Computing*, 30(3):1508–1526, 2008. 1

[DP04] Inderjit S. Dhillon and Beresford N. Parlett. Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices. *Linear Algebra and its Applications*, 387:1–28, 2004. 1

[DPV05] Inderjit S. Dhillon, Beresford N. Parlett, and Christof Vömel. Glued matrices and the MRRR algorithm. *SIAM Journal on Scientific Computing*, 27(2):496–510, 2005. 1

[DPV06] Inderjit S. Dhillon, Beresford N. Parlett, and Christof Vömel. The design and implementation of the MRRR algorithm. *ACM Transactions on Mathematical Software*, 32(4):533–560, 2006. 1

[GSU01] Wilfried Gansterer, Josef Schneid, and Christoph Ueberhuber. A low-complexity divide-and-conquer method for computing eigenvalues and eigenvectors of symmetric band matrices. *BIT Numerical Mathematics*, 41(5):967–976, 2001. 1

[GWMG03] Wilfried N. Gansterer, Robert C. Ward, Richard P. Muller, and William A. Goddard, III. Computing approximate eigenpairs of symmetric block tridiagonal matrices. *SIAM Journal on Scientific Computing*, 25(1):65–85, 2003. 1

[LXES19] Ruipeng Li, Yuanzhe Xi, Lucas Erlandson, and Yousef Saad. The eigenvalues slicing library (EVSL): Algorithms, implementation, and software. *SIAM Journal on Scientific Computing*, 41(4):C393–C415, 2019. 2

[LXV+16] Ruipeng Li, Yuanzhe Xi, Eugene Vecharynski, Chao Yang, and Yousef Saad. A thick-restart Lanczos algorithm with polynomial filtering for Hermitian eigenvalue problems. *SIAM Journal on Scientific Computing*, 38(4):A2512–A2534, 2016. 2

[MDV20] Osni A. Marques, James W. Demmel, and Paulo B. Vasconcelos. Bidiagonal SVD computation via an associated tridiagonal eigenproblem. *ACM Transactions on Mathematical Software*, 46(2):14:1–14:25, 2020. 1

[MPV06] Osni A. Marques, Beresford N. Parlett, and Christof Vömel. Computations of eigenpair subsets with the MRRR algorithm. *Numerical Linear Algebra with Applications*, 13(8):643–653, 2006. 1

[PD00] Beresford N. Parlett and Inderjit S. Dhillon. Relatively robust representations of symmetric tridiagonals. *Linear Algebra and its Applications*, 309(1–3):121–151, 2000. 1

[Sha25] Rikhav Shah. Hermitian diagonalization in linear precision. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5599–5615. Society for Industrial and Applied Mathematics, 2025. 2

[ŠK21] Ana Šušnjara and Daniel Kressner. A fast spectral divide-and-conquer method for banded matrices. *Numerical Linear Algebra with Applications*, 28(4):e2365, 2021. 1

[Sob25] Aleksandros Sobczyk. Deterministic complexity analysis of hermitian eigenproblems. In *52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025)*, volume 334 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 131:1–131:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. 2

[TP14] Ping Tak Peter Tang and Eric Polizzi. FEAST as a subspace iteration eigensolver accelerated by approximate spectral projection. *SIAM Journal on Matrix Analysis and Applications*, 35(2):354–390, 2014. 2

[VXCB16] James Vogel, Jianlin Xia, Stephen Cauley, and Venkataramanan Balakrishnan. Superfast divide-and-conquer method and perturbation analysis for structured eigenvalue solutions. *SIAM Journal on Scientific Computing*, 38(3):A1358–A1382, 2016. 1

[Wil62] J. H. Wilkinson. Calculation of the eigenvectors of a symmetric tridiagonal matrix by inverse iteration. *Numerische Mathematik*, 4:368–376, 1962. 1

[WL12] Paul R. Willems and Bruno Lang. The MR^3-GK algorithm for the bidiagonal SVD. *Electronic Transactions on Numerical Analysis*, 39:1–21, 2012. 1

[WL13] Paul R. Willems and Bruno Lang. A framework for the MR^3 algorithm: Theory and implementation. *SIAM Journal on Scientific Computing*, 35(2):A740–A766, 2013. 1

[WLV06] Paul R. Willems, Bruno Lang, and Christof Vömel. Computing the bidiagonal SVD using multiple relatively robust representations. *SIAM Journal on Matrix Analysis and Applications*, 28(4):907–926, 2006. 1

[XS16] Yuanzhe Xi and Yousef Saad. Computing partial spectra with least-squares rational filters. *SIAM Journal on Scientific Computing*, 38(5):A3020–A3045, 2016. 2