Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Exact Spectral Certificates and Ordering Geometry for Overlapping Block Matrix Inversion

**Binh T. Nguyen**

ngtbinh@hcmus.edu.vn

Ho Chi Minh City University of Science

**Research Article**

**Keywords:** iterative block matrix inversion, symmetric positive definite matrix, overlapping block methods, return operator, subspace correction, matrix inverse approximation

**Posted Date:** November 3rd, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11357445/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Exact Spectral Certificates and Ordering Geometry for Overlapping Block Matrix Inversion

Binh T. Nguyen[1*]

[1*]Faculty of Mathematics and Computer Sciences, University of Science, Vietnam National University Ho Chi Minh City, Ho Chi Minh City, Vietnam.

Corresponding author(s). E-mail(s): ngtbinh@hcmus.edu.vn;

**Abstract**

Iterative block matrix inversion (IBMI) approximates the inverse of a symmetric positive definite matrix through ordered block-inverse updates. We establish an exact pre-iteration certificate for the asymptotic convergence factor of the full inverse-error iteration. Each block update induces an energy-orthogonal projection, and a complete sweep admits a reduced cyclic return representation whose spectral radius determines the inverse-error factor. We show that the return matrix can be constructed entirely from the original matrix $\boldsymbol{A}$ through local principal-block solves, yielding the exact factor $\boldsymbol{\rho(M_A)^2}$ without forming $\boldsymbol{A^{-1}}$ or executing an IBMI sweep.

The return framework provides finite-termination criteria, perturbation bounds, and a first-order sensitivity formula under a simple dominant-eigenvalue condition. For two update blocks, nested enlargement cannot increase the exact convergence factor; an explicit rational SPD counterexample shows that this monotonicity fails for three blocks. Nevertheless, a monotone canonical-correlation envelope holds for arbitrary block counts, with a quantitative condition-number contraction bound for separated adjacent complements.

For singleton complements, the exact factor reduces to a weighted Hamiltonian-cycle product of inverse couplings. This characterization yields exact local ordering comparisons, an NP-hardness result for unrestricted SPD instances, and an explicit globally optimal ordering for exponential-line coupling structures. Numerical experiments on covariance models, public sparse SPD matrices, and controlled ordering instances confirm the spectral identities to numerical accuracy and demonstrate the matrix-dependent effects of overlap and ordering.

**Keywords:** iterative block matrix inversion, symmetric positive definite matrix, overlapping block methods, return operator, subspace correction, matrix inverse approximation

# 1 Introduction

Iterative block matrix inversion (IBMI) [1] constructs an approximation to the inverse of a symmetric positive definite (SPD) matrix by repeatedly applying exact block-inverse identities. Its initial analysis addresses two nonoverlapping blocks, while reported multi-block experiments indicate that both overlap and sweep ordering affect convergence. These observations motivate a more precise question: given the matrix and an ordered collection of update blocks, can the asymptotic inverse-error factor be determined *before* the IBMI iteration is run?

Let $A = A^\top \succ 0$, $H = A^{-1}$, and $E = X - H$ denote the inverse error. Write $[n] = \{1, \dots, n\}$, and let $E_S$ be the coordinate embedding for an index set $S \subseteq [n]$. For a nonempty proper update block $I \subsetneq [n]$, with complement $J = I^c$, one IBMI update has the exact error representation

$$E^+ = Q_I E Q_I^\top, \qquad Q_I = H E_J H_{JJ}^{-1} E_J^\top,$$

The matrix $Q_I$ is an $A$-orthogonal projector. Consequently, an ordered sweep with $K \geq 2$ blocks has the form $E^+ = QEQ^\top$, where $Q = Q_{I_K} \cdots Q_{I_1}$. The connection to products of energy projections and successive subspace correction is classical [2–4].

For an ordered cover $I_1, \dots, I_K$ of $[n]$, write $J_k = I_k^c$. The complement-space structure of these projections yields a reduced cyclic return matrix $M$ whose nonzero eigenvalues, including algebraic multiplicities, coincide with those of $Q$. For the induced linear operator $\mathcal{L}(E) = QEQ^\top$ on symmetric inverse errors, this gives the exact worst-case asymptotic spectral factor

$$\rho(\mathcal{L}) = \rho(Q)^2 = \rho(M)^2.$$

A representation of $M$ in terms of blocks of $H$, however, would require access to the inverse that IBMI is intended to approximate.

Our central result makes the return matrix computable from $A$ alone. For an update block $I$, its complement $J$, and any target index set $L$, the transfer between complementary coordinates satisfies

$$H_{LJ} H_{JJ}^{-1} = E_L^\top E_I (-A_{II}^{-1} A_{IJ}) + E_L^\top E_J. \tag{1}$$

Thus, each transfer can be obtained by a principal-block solve with $A$; composing the transfers around the ordered sweep produces a matrix $M_A$ that agrees exactly with $M$. In particular, $\rho(\mathcal{L}) = \rho(M_A)^2$ without forming $A^{-1}$ or generating an IBMI trajectory. The local block identity itself follows from standard block elimination. The IBMI-specific consequence is that its cyclic composition recovers the exact spectral factor of the *full* inverse-error operator. This is an asymptotic certificate, not a bound on every finite-sweep error or a claim that certificate construction is cheaper than an individual IBMI sweep.

The return representation has further consequences for overlap and ordering. For two update blocks, the exact factor is governed by a largest canonical correlation and cannot increase when either block is enlarged. This monotonicity fails for a three-block

SPD example; for arbitrary block counts, a monotone product of adjacent canonical correlations remains a valid, generally nonsharp upper bound. When at least one cyclically adjacent pair of complements is disjoint, a Wielandt bound converts this envelope into a strict $A$-computable contraction estimate involving Schur-complement condition numbers. A dimension-independent rate follows for families with a uniformly bounded condition number and at least one such separated pair. Nilpotent return matrices also identify finite-sweep termination mechanisms.

For singleton complements, the exact factor becomes the product of normalized inverse couplings along a Hamiltonian cycle. Its negative logarithm converts rate-optimal ordering into a symmetric maximum-weight traveling-salesperson (Max-TSP) objective. We establish an SPD realization of arbitrary finite nonnegative edge weights up to an order-independent constant, a rational NP-hardness reduction, and an exact four-coupling certificate for each admissible 2-opt exchange. Under a log-metric assumption, an existing metric Max-TSP approximation guarantee [5] transfers to the asymptotic contraction exponent. For exponential couplings on an ordered line, where $A$ is tridiagonal, classical longest-tour geometry [6] yields a globally optimal ordering and its factor in closed form. For fixed general block collections, local graph maps can instead be cached to evaluate different sweep orders without repeating the principal-block factorizations.

***Contributions and scope.***

The principal contribution is the exact $A$-only return certificate for the full inverse-error operator. Its accompanying results are (i) finite and first-order perturbation bounds, including a Hölder spectral-radius estimate that does not assume diagonalizability; (ii) two-block overlap monotonicity, its three-block limitation, and a quantitative multi-block contraction bound; (iii) singleton ordering geometry, its realization and complexity consequences, exact local exchanges, and a solvable structured class; and (iv) reuse of local solves across permutations of a fixed block collection. Numerical experiments with structured covariance matrices, public SPD stiffness matrices, and controlled ordering instances assess the certificate and ordering predictions, while also documenting finite-precision limitations.

***Relation to prior work.***

The use of energy projections, canonical correlations, and subspace correction is established [2–4, 7]. Ordering effects in related sequential correction methods have also been studied through greedy, randomized, and expected-rate analyses [8–11]. Likewise, Hamiltonian-cycle optimization, 2-opt local search, and metric Max-TSP approximation are classical; the hardness argument uses the Hamiltonian-cycle problem [12]. Our contribution is not a new projection principle or combinatorial algorithm, but the exact identification of the IBMI return factor with these geometric and ordering objects, together with its $A$-only construction and consequences. Related work on Schwarz methods, selected and randomized inversion, and structured covariance approximation [13–20] addresses different computational objectives or update mechanisms. The present results concern the exact asymptotic factor and ordering of the

IBMI inverse-error iteration, rather than a general acceleration guarantee for matrix inversion.

# 2 Energy Projections and Inverse-Error Dynamics

Throughout, let $A \in \mathbb{R}^{n\times n}$ be symmetric positive definite (SPD), and set $H = A^{-1}$. For an index set $S \subseteq [n] := \{1, \dots, n\}$, let $E_S$ denote the coordinate embedding whose columns are the standard basis vectors indexed by $S$. Thus, $H_{ST} = E_S^\top H E_T$. We use the energy norm $\|x\|_A = (x^\top Ax)^{1/2}$ and write

$$V_I := \mathrm{range}(E_I)$$

for the coordinate subspace supported on $I$.

## 2.1 One exact block update

Let $I \subsetneq [n]$ be nonempty and set $J = I^c$. After a simultaneous permutation of rows and columns, partition

$$A = \begin{bmatrix} A_{II} & A_{IJ} \\ A_{JI} & A_{JJ} \end{bmatrix}. \tag{2}$$

An IBMI update retains the current approximation $X_{JJ}$ on the complementary principal block and reconstructs the remaining blocks using the block-inverse identities [1]. With

$$R_I := A_{II}^{-1} A_{IJ},$$

the update is

$$\Phi_I(X_{JJ}) = \begin{bmatrix} A_{II}^{-1} + R_I X_{JJ} R_I^\top & -R_I X_{JJ} \\ -X_{JJ} R_I^\top & X_{JJ} \end{bmatrix}. \tag{3}$$

Setting $X_{JJ} = H_{JJ}$ recovers the exact inverse $H$.

**Theorem 2.1** (IBMI inverse-error congruence)**.** *Let $X$ be symmetric and let $E = X - H$. One exact IBMI update on $I$ satisfies*

$$E^+ = Q_I E Q_I^\top, \qquad Q_I := H E_J H_{JJ}^{-1} E_J^\top. \tag{4}$$

*Equivalently, defining*

$$Z_I = \begin{bmatrix} -R_I \\ I_{|J|} \end{bmatrix}$$

*in the $(I, J)$ ordering, we have*

$$E^+ = Z_I E_{JJ} Z_I^\top, \qquad Q_I = Z_I E_J^\top.$$

*Proof* Subtracting the exact block-inverse identity from (3) yields

$$E^+ = \begin{bmatrix} R_I E_{JJ} R_I^\top & -R_I E_{JJ} \\ -E_{JJ} R_I^\top & E_{JJ} \end{bmatrix} = Z_I E_{JJ} Z_I^\top.$$

Since $E_{JJ} = E_J^\top E E_J$, it remains to identify $Z_I E_J^\top$. The $(I, J)$ block of $AH = I_n$ gives

$$A_{II} H_{IJ} + A_{IJ} H_{JJ} = 0.$$

As $A_{II}$ and $H_{JJ}$ are SPD,

$$H_{IJ} H_{JJ}^{-1} = -A_{II}^{-1} A_{IJ} = -R_I.$$

Consequently,

$$H E_J H_{JJ}^{-1} = \begin{bmatrix} H_{IJ} \\ H_{JJ} \end{bmatrix} H_{JJ}^{-1} = \begin{bmatrix} -R_I \\ I_{|J|} \end{bmatrix} = Z_I,$$

which establishes (4). □

Theorem 2.1 expresses the full inverse-error update as a two-sided congruence transformation. This identity allows the sweep dynamics to be analyzed through products of energy projections.

## 2.2 Energy-projection structure

**Proposition 2.2** (Energy projection)**.** *For every nonempty proper block $I$,*

$$Q_I^2 = Q_I, \qquad Q_I^\top A = A Q_I. \tag{5}$$

*Consequently, $P_I := I_n - Q_I$ is the $A$-orthogonal projector onto $V_I$, whereas $Q_I$ is the $A$-orthogonal projector onto $V_I^{\perp_A}$.*

*Proof* Using $E_J^\top H E_J = H_{JJ}$ gives

$$Q_I^2 = H E_J H_{JJ}^{-1} (E_J^\top H E_J) H_{JJ}^{-1} E_J^\top = Q_I.$$

Moreover,

$$A Q_I = E_J H_{JJ}^{-1} E_J^\top,$$

which is symmetric. Hence $Q_I^\top A = A Q_I$, so $Q_I$ is an $A$-orthogonal projector.

Since $E_J^\top E_I = 0$, we have $V_I \subseteq \ker(Q_I)$. Also,

$$\operatorname{rank}(Q_I) = |J|,$$

because $H E_J H_{JJ}^{-1}$ has full column rank and $E_J^\top$ has full row rank. Thus $\dim \ker(Q_I) = |I|$, and

$$\ker(Q_I) = V_I.$$

The range of an $A$-orthogonal projector is the $A$-orthogonal complement of its kernel. Therefore, $\operatorname{range}(Q_I) = V_I^{\perp_A}$, and $P_I = I_n - Q_I$ projects $A$-orthogonally onto $V_I$. □

This projection interpretation connects IBMI with classical successive subspace-correction theory [3, 4]. The additional structure used below comes from the coordinate representation of $Q_I$.

## 2.3 Ordered sweeps and fixed-cover convergence

Let $\mathcal{I} = (I_1, \ldots, I_K)$ be an ordered family of nonempty proper blocks, and define $Q := Q_{I_K} \cdots Q_{I_1}$. One complete sweep propagates the inverse error according to

$$E^{(r+1)} = QE^{(r)}Q^\top. \tag{6}$$

Repeated sweeps with the same ordering satisfy

$$E^{(r)} = Q^r E^{(0)} (Q^\top)^r. \tag{7}$$

**Corollary 2.3** (General fixed-cover convergence)**.** *Suppose that*

$$\bigcup_{k=1}^{K} I_k = [n]. \tag{8}$$

*Then,*

$$\|Q\|_A < 1, \qquad \rho(Q) < 1,$$

*and the fixed-order IBMI iteration converges to $H = A^{-1}$ for every symmetric initial approximation $X^{(0)}$.*

*More precisely, with*

$$\|E\|_{A,\mathrm{op}} := \|A^{1/2} E A^{1/2}\|_2,$$

*the inverse error satisfies*

$$\|E^{(r)}\|_{A,\mathrm{op}} \le \|Q\|_A^{2r} \|E^{(0)}\|_{A,\mathrm{op}}.$$

*No disjointness assumption is required.*

*Proof* Each $Q_{I_k}$ is an $A$-orthogonal projector and hence nonexpansive in the energy norm. Suppose that $\|Qx\|_A = \|x\|_A$ for some $x \neq 0$. Then, every inequality in

$$\|Q_{I_K} \cdots Q_{I_1} x\|_A \le \cdots \le \|Q_{I_1} x\|_A \le \|x\|_A$$

must be an equality.

Equality for an orthogonal projection holds precisely on its range. Applying this observation successively shows that

$$x \in \bigcap_{k=1}^{K} V_{I_k}^{\perp_A} = \left( \sum_{k=1}^{K} V_{I_k} \right)^{\perp_A}.$$

The cover condition implies $\sum_{k=1}^{K} V_{I_k} = \mathbb{R}^n$, so the intersection is $\{0\}$, a contradiction. Compactness of the $A$-unit sphere therefore gives $\|Q\|_A < 1$, and consequently, $\rho(Q) < 1$. Finally,

$$\begin{aligned} \|E^{(r)}\|_{A,\mathrm{op}} &= \left\| (A^{1/2} Q A^{-1/2})^r (A^{1/2} E^{(0)} A^{1/2}) \cdot \left( (A^{1/2} Q A^{-1/2})^\top \right)^r \right\|_2 \\ &\le \|Q\|_A^{2r} \|E^{(0)}\|_{A,\mathrm{op}}, \end{aligned}$$

which proves convergence. □

The cover condition gives a strict contraction for every fixed ordering, but the bound $\|Q\|_A^2$ need not equal the asymptotic inverse-error factor. The latter is determined by the sweep spectrum and can depend substantially on the ordered block geometry.

## 2.4 Transfer coordinates for ordered energy projections

We record a standard finite-dimensional factorization that will be used to reduce the IBMI sweep dynamics. Let $A = A^\top \succ 0$, and consider $K \geq 2$ $A$-orthogonal projectors

$$P_k = Z_k G_k^{-1} Z_k^\top A, \qquad G_k := Z_k^\top A Z_k,$$

onto $\mathcal{W}_k = \mathrm{range}(Z_k)$, where each $Z_k$ has full column rank. Define the successive transfer matrices

$$T_{k+1,k} := G_{k+1}^{-1} Z_{k+1}^\top A Z_k, \qquad k = 1, \ldots, K-1,$$

and the closing transfer

$$T_{1,K} := G_1^{-1} Z_1^\top A Z_K.$$

Their cyclic product is the return matrix

$$M_{\mathcal{W}} := T_{1,K} T_{K,K-1} \cdots T_{2,1}.$$

Set

$$Q := P_K \cdots P_1, \qquad B := T_{K,K-1} \cdots T_{2,1}.$$

Direct multiplication yields

$$Q = Z_K B G_1^{-1} Z_1^\top A. \tag{9}$$

Writing $Q = UV^\top$, where

$$U := Z_K B, \qquad V^\top := G_1^{-1} Z_1^\top A,$$

gives

$$V^\top U = M_{\mathcal{W}}.$$

The standard $UV^\top$–$V^\top U$ spectral identity therefore implies

$$\sigma(Q) \setminus \{0\} = \sigma(M_{\mathcal{W}}) \setminus \{0\}, \qquad \rho(Q) = \rho(M_{\mathcal{W}}). \tag{10}$$

The nonzero eigenvalues agree with their algebraic multiplicities. Moreover, for every integer $s \geq 1$,

$$Q^s = Z_K B M_{\mathcal{W}}^{s-1} G_1^{-1} Z_1^\top A. \tag{11}$$

Under changes of basis $Z_k \mapsto Z_k S_k$, with each $S_k$ invertible, the transfer matrices transform as

$$T_{k+1,k} \mapsto S_{k+1}^{-1} T_{k+1,k} S_k,$$

and consequently,

$$M_{\mathcal{W}} \mapsto S_1^{-1} M_{\mathcal{W}} S_1.$$

Thus, the return spectrum depends only on the ordered subspaces, not on their chosen bases. In $A$-orthonormal coordinates, the transfers are cross-Gram matrices; for two subspaces, their product recovers the classical squared-cosine principal-angle spectrum.

This factorization is a standard consequence of projection coordinates and the $UV^\top$–$V^\top U$ identity, rather than a new result about products of projections [2–4, 7]. Its role here is to identify the return-coordinate mechanism. Section 4 exploits the specific structure of IBMI to compute these transfers from local solves with $A$, without forming $A^{-1}$.

# 3 Return-operator spectral theory

The projection representation in Section 2 establishes convergence for a fixed ordered cover. To determine the asymptotic inverse-error factor, we exploit the coordinate structure of the IBMI projectors and reduce the nonzero sweep spectrum to a return matrix acting on the complement of the first update block.

Let $\mathcal{I} = (I_1, \dots, I_K)$ be an ordered family of nonempty proper blocks, with $K \geq 2$, and set $J_k = I_k^c$. Write

$$Q_k := Q_{I_k} = HE_{J_k} H_{J_k J_k}^{-1} E_{J_k}^\top, \qquad Q := Q_K \cdots Q_1.$$

Define the transfer product

$$\begin{aligned} B := H_{J_K J_K}^{-1} H_{J_K J_{K-1}} H_{J_{K-1} J_{K-1}}^{-1} \cdots \\ \cdots H_{J_2 J_2}^{-1} H_{J_2 J_1} H_{J_1 J_1}^{-1}. \end{aligned} \tag{12}$$

Thus, $B \in \mathbb{R}^{|J_K| \times |J_1|}$. The associated return matrix is

$$M := H_{J_1 J_K} B \in \mathbb{R}^{|J_1| \times |J_1|}. \tag{13}$$

The return matrix closes the sequence of transfers through $J_1, \dots, J_K$ by mapping the final complement coordinates back to those of $J_1$.

**Theorem 3.1** (Return factorization)**.** *For every ordered IBMI sweep with $K \geq 2$,*

$$Q = HE_{J_K} BE_{J_1}^\top. \tag{14}$$

*Set*

$$U := HE_{J_K} B, \qquad V := E_{J_1}.$$

*Then, $Q = UV^\top$, $M = V^\top U$, and, for every integer $s \geq 1$,*

$$Q^s = UM^{s-1}V^\top = HE_{J_K} BM^{s-1} E_{J_1}^\top. \tag{15}$$

*Moreover,*

$$\det(\lambda I_n - Q) = \lambda^{n-|J_1|} \det(\lambda I_{|J_1|} - M). \tag{16}$$

*Consequently,*

$$\sigma(Q) \setminus \{0\} = \sigma(M) \setminus \{0\}, \tag{17}$$

*where nonzero eigenvalues agree with their algebraic multiplicities. In particular,*

$$\rho(Q) = \rho(M). \tag{18}$$

*Proof* For consecutive projectors,

$$Q_{k+1}Q_k = HE_{J_{k+1}}H^{-1}_{J_{k+1}J_{k+1}}H_{J_{k+1}J_k}H^{-1}_{J_kJ_k}E^\top_{J_k}.$$

Successive multiplication yields (14) with $B$ defined in (12). Hence, $Q = UV^\top$, and

$$V^\top U = E^\top_{J_1}HE_{J_K}B = H_{J_1J_K}B = M.$$

Sylvester's determinant identity gives, initially for $\lambda \neq 0$,

$$\det(\lambda I_n - UV^\top) = \lambda^n \det(I_{|J_1|} - \lambda^{-1}V^\top U) = \lambda^{n-|J_1|}\det(\lambda I_{|J_1|} - M).$$

Both sides are polynomials in $\lambda$, so the identity holds for all $\lambda$. This proves (16), including the assertion about algebraic multiplicities, and yields (17)–(18).

Finally, $Q = UV^\top$ and $V^\top U = M$ imply

$$Q^{s+1} = UM^{s-1}(V^\top U)V^\top = UM^sV^\top,$$

which proves (15) by induction. □

Theorem 3.1 identifies an explicit complement-space return matrix for the ordered IBMI sweep. The $UV^\top$–$V^\top U$ spectral identity is classical; the IBMI-specific content is the factorization (14) and the resulting return matrix (13). When $|J_1| \ll n$, this representation substantially reduces the dimension of the spectral problem. The matrix $M$ is generally nonsymmetric, even though $A$ and $H$ are SPD.

## 3.1 Exact inverse-error dynamics and the square law

Let $\mathbb{S}^n$ denote the real symmetric matrices and define the linear inverse-error operator $\mathcal{L} : \mathbb{S}^n \to \mathbb{S}^n$ by

$$\mathcal{L}(E) := QEQ^\top. \tag{19}$$

The return factorization gives an exact representation of every repeated sweep.

**Corollary 3.2** (Exact return dynamics)**.** *For every integer $s \geq 1$ and every symmetric initial error $E^{(0)}$,*

$$E^{(s)} = UM^{s-1}V^\top E^{(0)}V(M^\top)^{s-1}U^\top. \tag{20}$$

*Thus, the inverse-error dynamics after the first sweep are expressed entirely through powers of $M$ and $M^\top$, together with the fixed endpoint factors.*

*Proof* Substitute (15) and its transpose into (7). □

**Theorem 3.3** (Full inverse-error square law)**.** *The spectral radius of $\mathcal{L}$, restricted to the real symmetric matrix space, satisfies*

$$\rho(\mathcal{L}) = \rho(Q)^2 = \rho(M)^2. \tag{21}$$

*Thus, $\rho(M)^2$ is the exact asymptotic spectral factor of the full inverse-error operator.*

*Proof* On the full matrix space, the vectorization identity gives

$$\operatorname{vec}(QEQ^\top) = (Q \otimes Q)\operatorname{vec}(E).$$

The unrestricted operator therefore has spectral radius $\rho(Q \otimes Q) = \rho(Q)^2$. Since $\mathbb{S}^n$ is invariant under $\mathcal{L}$, the restriction satisfies

$$\rho(\mathcal{L}) \le \rho(Q)^2.$$

For the reverse inequality, complexify the symmetric matrix space. Choose an eigenvalue $\lambda$ of $Q$ such that $|\lambda| = \rho(Q)$, and let $v \neq 0$ satisfy $Qv = \lambda v$ over $\mathbb{C}$. The matrix $vv^\top$ is nonzero and complex symmetric. Moreover,

$$\mathcal{L}_{\mathbb{C}}(vv^\top) = (Qv)(Qv)^\top = \lambda^2 vv^\top.$$

Hence, $\lambda^2$ belongs to the spectrum of the complexified restriction, giving

$$\rho(\mathcal{L}) \ge |\lambda|^2 = \rho(Q)^2.$$

The spectral radius is unchanged by complexification. Together with (18), this proves (21). □

The square law is an operator-level statement. An individual initial error may have no component in a dominant spectral mode and may therefore converge faster. Furthermore, the spectral radius determines the asymptotic operator factor but does not, by itself, bound finite-sweep amplification for a nonnormal return matrix.

## 3.2 Finite-time bounds and nonnormal return dynamics

The exact return dynamics also provide finite-sweep bounds that do not require diagonalizability.

**Proposition 3.4** (Finite-time return bound)**.** *For every integer $s \ge 1$,*

$$\|E^{(s)}\|_2 \le \|U\|_2^2 \|M^{s-1}\|_2^2 \|E^{(0)}\|_2. \tag{22}$$

*If $M = X\Lambda X^{-1}$ is diagonalizable over $\mathbb{C}$, then*

$$\|E^{(s)}\|_2 \le \|U\|_2^2 \kappa_2(X)^2 \rho(M)^{2(s-1)} \|E^{(0)}\|_2. \tag{23}$$

*More generally, suppose that $\rho(M) > 0$, and let $p$ be the largest size of a Jordan block associated with an eigenvalue of modulus $\rho(M)$. Then,*

$$\|E^{(s)}\|_2 = O\Big(s^{2(p-1)} \rho(M)^{2s}\Big) \qquad (s \to \infty), \tag{24}$$

*where the implicit constant may depend on $M$, $U$, and $E^{(0)}$, but not on $s$.*

*Proof* Since $V = E_{J_1}$ has orthonormal columns, $\|V\|_2 = 1$. Taking spectral norms in (20) yields (22). If $M = X\Lambda X^{-1}$, then

$$\|M^k\|_2 \le \|X\|_2\|X^{-1}\|_2\|\Lambda^k\|_2 = \kappa_2(X)\rho(M)^k.$$

Substitution with $k = s-1$ proves (23).

For the general case, write $M = XJX^{-1}$ in Jordan canonical form over $\mathbb{C}$. A Jordan block $J_\lambda = \lambda I + N$ of size $m$ satisfies

$$J_\lambda^k = \sum_{j=0}^{m-1} \binom{k}{j} \lambda^{k-j} N^j,$$

where terms with $j > k$ are understood to vanish. For $|\lambda| = \rho(M) > 0$, the contribution of such a block is $O(k^{m-1}\rho(M)^k)$. Blocks with $|\lambda| < \rho(M)$ contribute terms that decay exponentially relative to $\rho(M)^k$, even when their Jordan sizes are larger. Consequently,

$$\|M^k\|_2 = O(k^{p-1}\rho(M)^k).$$

Applying (22) with $k = s-1$ gives (24). □

The return matrix therefore serves two distinct purposes: its spectral radius gives the exact asymptotic inverse-error factor, while the norms of its powers control finite-sweep behavior. The latter distinction is examined numerically in Section 9.

# 4 An exact *A*-only return certificate

The return matrix in Section 3 is expressed through blocks of $H = A^{-1}$. We next construct the same matrix using local solves with $A$, without forming $H$ or performing an IBMI sweep.

For a nonempty proper block $I$, set $J = I^c$. For any index set $L \subseteq [n]$, define the transfer

$$\mathcal{R}_{L\leftarrow J} := H_{LJ}H_{JJ}^{-1}. \tag{25}$$

**Theorem 4.1** ($A$-only local-transfer identity)**.** *For every nonempty proper $I$, with $J = I^c$, and every $L \subseteq [n]$,*

$$\mathcal{R}_{L\leftarrow J} = E_L^\top E_I(-A_{II}^{-1}A_{IJ}) + E_L^\top E_J. \tag{26}$$

*Thus, the transfer is obtained by solving a system with the SPD principal block $A_{II}$ and restricting the resulting graph map to $L$.*

*Proof* The $(I, J)$ block equation in $AH = I_n$ reads

$$A_{II}H_{IJ} + A_{IJ}H_{JJ} = 0.$$

Since $A_{II}$ and $H_{JJ}$ are invertible,

$$H_{IJ}H_{JJ}^{-1} = -A_{II}^{-1}A_{IJ}.$$

Together with $H_{JJ}H_{JJ}^{-1} = I_{|J|}$, this gives

$$HE_JH_{JJ}^{-1} = E_I(-A_{II}^{-1}A_{IJ}) + E_J.$$

Left multiplication by $E_L^\top$ proves the claim. □

The local identity is a standard consequence of block elimination. Its role here is to realize every transfer in the IBMI return cycle using the original matrix $A$.

**Theorem 4.2** (Exact pre-iteration return certificate)**.** *Let $\mathcal{I} = (I_1, \dots, I_K)$ be an ordered IBMI cover with $K \geq 2$, and set $J_k = I_k^c$. Define*

$$\mathcal{R}_{k+1,k} := \mathcal{R}_{J_{k+1} \leftarrow J_k} \quad (k = 1, \dots, K-1), \qquad \mathcal{R}_{1,K} := \mathcal{R}_{J_1 \leftarrow J_K},$$

*with each transfer evaluated by* (26)*. Then*

$$M_A := \mathcal{R}_{1,K} \mathcal{R}_{K,K-1} \cdots \mathcal{R}_{2,1} = M, \tag{27}$$

*where $M$ is the return matrix in* (13)*. Consequently, for the full symmetric inverse-error sweep operator,*

$$\rho(\mathcal{L}) = \rho(M_A)^2. \tag{28}$$

*The right-hand side is determined by $A$ and the ordered cover, without forming $A^{-1}$ or running IBMI.*

*Proof* By Theorem 4.1,

$$\mathcal{R}_{J_{k+1} \leftarrow J_k} = H_{J_{k+1} J_k} H_{J_k J_k}^{-1},$$

and the closing transfer has the analogous form. Their ordered product is exactly the matrix $M$ defined in (12)–(13). The square law (21) gives (28). □

*Remark* 4.3 (Scope of the certificate). The certificate is exact for the asymptotic spectral factor of the full error operator, not necessarily for the rate of a particular initial error. It does not bound finite-sweep transients from the spectral radius alone: the bounds in Proposition 3.4 involve powers of $M_A$. The term "pre-iteration" refers to information availability, not to negligible construction cost.

## 4.1 Cost of the certificate

Let $b_k = |I_k|$ and $r_k = |J_k|$. The local graph map

$$T_k := H E_{J_k} H_{J_k J_k}^{-1} = E_{I_k}(-A_{I_k I_k}^{-1} A_{I_k J_k}) + E_{J_k}$$

can be computed by factoring $A_{I_k I_k}$ and solving for $r_k$ right-hand sides. For dense Cholesky factorization and dense triangular solves, the leading arithmetic cost is

$$\frac{1}{3} b_k^3 + 2 b_k^2 r_k \tag{29}$$

flops, excluding lower-order terms and data movement. The transfers are obtained by restricting $T_k$ to the required target coordinates; the return product and its spectral radius entail additional reduced matrix multiplications and an eigenvalue computation. These costs, and any sparse fill-in, are not included in (29).

The construction avoids forming the full inverse, but it need not be cheaper than one IBMI sweep. Small complements reduce the dimensions of the transfer matrices while increasing the sizes of the factored principal blocks. The main computational opportunity is therefore reuse: for a fixed collection of update blocks, the local graph maps can be cached and the return matrices for different orderings assembled from their restrictions.

# 5 Perturbation and sensitivity of the $A$-only certificate

The certificate in Theorem 4.2 is exact for a fixed SPD matrix and ordered cover. We now quantify its sensitivity to perturbations of $A$, keeping the index sets and their ordering fixed. A finite perturbation bound requires no diagonalizability of the return matrix; a first-order formula is available when the dominant eigenvalue is simple and spectrally separated. The spectral-variation ingredients are classical [21, 22]. Their application here follows the dependence of the IBMI return product on local principal-block solves.

Fix a nonempty proper block $I$, set $J = I^c$, and write

$$B = A_{II}, \qquad C = A_{IJ}, \qquad F = -B^{-1}C.$$

For a symmetric perturbation $\Delta A$, define $\Delta B = (\Delta A)_{II}$ and $\Delta C = (\Delta A)_{IJ}$. Quantities associated with $\widetilde{A} = A + \Delta A$ are denoted by tildes.

**Lemma 5.1** (Local transfer perturbation)**.** *Suppose that*

$$\eta := \|B^{-1}\|_2 \|\Delta B\|_2 < 1.$$

*Then, $\widetilde{B} = B + \Delta B$ is nonsingular and*

$$\|\widetilde{F} - F\|_2 \le \delta_I := \frac{\|B^{-1}\|_2}{1-\eta} \left(\|\Delta B\|_2 \|F\|_2 + \|\Delta C\|_2\right). \tag{30}$$

*For every restriction set $L$,*

$$\|\widetilde{\mathcal{R}}_{L\leftarrow J} - \mathcal{R}_{L\leftarrow J}\|_2 \le \delta_I. \tag{31}$$

*Moreover, the local graph map is Fréchet differentiable, with*

$$DF[A](\Delta A) = -B^{-1}(\Delta B\, F + \Delta C). \tag{32}$$

*Proof* Since

$$(B + \Delta B)\widetilde{F} = -(C + \Delta C), \qquad BF = -C,$$

subtraction gives

$$(B + \Delta B)(\widetilde{F} - F) = -(\Delta B\, F + \Delta C).$$

Therefore,

$$\widetilde{F} - F = -(B + \Delta B)^{-1}(\Delta B\, F + \Delta C).$$

The assumption $\eta < 1$ and the Banach perturbation lemma imply

$$\|(B + \Delta B)^{-1}\|_2 \le \frac{\|B^{-1}\|_2}{1 - \eta},$$

which proves (30). By Theorem 4.1,

$$\mathcal{R}_{L \leftarrow J} = E_L^\top E_I F + E_L^\top E_J.$$

Since coordinate restriction matrices have operator norm at most one, (31) follows. Finally, applying the preceding identity to $A + t\Delta A$, dividing by $t$, and taking $t \to 0$ gives (32). □

We next propagate these local estimates through the cyclic return product. Enumerate the $K$ transfers in multiplication order as $R_1, \dots, R_K$, so that

$$M_A = R_K \cdots R_1.$$

Let $a_j = \|R_j\|_2$, and let $\delta_j$ denote the bound (31) for the update block generating $R_j$.

**Theorem 5.2** (Finite stability of the return certificate)**.** *Let $A \succ 0$ and $\widetilde{A} = A + \Delta A \succ 0$ share the same ordered IBMI cover. Suppose that the local smallness conditions in Lemma 5.1 hold for every update block. Define*

$$\varepsilon_M := \sum_{j=1}^{K} \delta_j \left( \prod_{\ell=j+1}^{K} (a_\ell + \delta_\ell) \right) \left( \prod_{\ell=1}^{j-1} a_\ell \right). \tag{33}$$

*Then,*

$$\|M_{\widetilde{A}} - M_A\|_2 \le \varepsilon_M. \tag{34}$$

*If $r = |J_1|$ is the return dimension, then*

$$|\rho(M_{\widetilde{A}}) - \rho(M_A)| \le \left(\|M_{\widetilde{A}}\|_2 + \|M_A\|_2\right)^{1-1/r} \varepsilon_M^{1/r}. \tag{35}$$

*Equivalently, for the full inverse-error sweep operators,*

$$\left| \sqrt{\rho(\mathcal{L}_{\widetilde{A}})} - \sqrt{\rho(\mathcal{L}_A)} \right| \le \left(\|M_{\widetilde{A}}\|_2 + \|M_A\|_2\right)^{1-1/r} \varepsilon_M^{1/r}. \tag{36}$$

*These bounds require neither normality nor diagonalizability of the return matrices.*

*Proof* Write

$$M_{\widetilde{A}} = \widetilde{R}_K \cdots \widetilde{R}_1.$$

A telescoping expansion gives

$$M_{\widetilde{A}} - M_A = \sum_{j=1}^{K} \widetilde{R}_K \cdots \widetilde{R}_{j+1} (\widetilde{R}_j - R_j) R_{j-1} \cdots R_1.$$

Using $\|\widetilde{R}_j - R_j\|_2 \le \delta_j$ and $\|\widetilde{R}_j\|_2 \le a_j + \delta_j$ proves (34). Elsner's spectral-variation inequality [21] gives

$$d_H\big(\sigma(M_{\widetilde{A}}), \sigma(M_A)\big) \le \big(\|M_{\widetilde{A}}\|_2 + \|M_A\|_2\big)^{1-1/r} \, \|M_{\widetilde{A}} - M_A\|_2^{1/r},$$

where $d_H$ denotes the Hausdorff distance between the spectra. Since

$$|\rho(M_{\widetilde{A}}) - \rho(M_A)| \le d_H\big(\sigma(M_{\widetilde{A}}), \sigma(M_A)\big),$$

we obtain (35). The square law (28) applied to both matrices yields (36). □

The exponent $1/r$ is a worst-case spectral-variation exponent. Defective nonnormal matrices can exhibit only Hölder-continuous eigenvalue variation under unrestricted perturbations. A sharper local formula is available when one simple eigenvalue uniquely determines the spectral radius.

**Theorem 5.3** (First-order sensitivity of the convergence factor)**.** *Let $A \succ 0$, and suppose that $M_A$ has a simple eigenvalue $\lambda_\star \neq 0$ that uniquely attains $\rho(M_A)$. Let $x$ and $y$ be corresponding right and left eigenvectors, normalized by $y^\top x = 1$.*

*For a symmetric direction $\Delta A$, form the local transfer derivatives using* (32)*, and define*

$$DM_A[\Delta A] = \sum_{j=1}^{K} R_K \cdots R_{j+1} \, DR_j[\Delta A] \, R_{j-1} \cdots R_1. \tag{37}$$

*Then, $\rho(M_A)$ is differentiable at $A$, and*

$$D\rho(M_A)[\Delta A] = \operatorname{sign}(\lambda_\star) \, y^\top DM_A[\Delta A] x. \tag{38}$$

*The full inverse-error factor $\gamma(A) := \rho(\mathcal{L}_A) = \rho(M_A)^2$ therefore satisfies*

$$D\gamma(A)[\Delta A] = 2|\lambda_\star| \operatorname{sign}(\lambda_\star) \, y^\top DM_A[\Delta A] x. \tag{39}$$

*In particular,*

$$|D\gamma(A)[\Delta A]| \le 2\rho(M_A)\|x\|_2\|y\|_2\|DM_A[\Delta A]\|_2. \tag{40}$$

*Proof* Because $M_A$ is real, its nonreal eigenvalues occur in conjugate pairs. The assumption that $\lambda_\star$ uniquely attains the spectral radius therefore implies $\lambda_\star \in \mathbb{R}$.

The simplicity and strict spectral dominance of $\lambda_\star$ ensure that its eigenvalue branch remains simple and dominant under sufficiently small real perturbations. Classical eigenvalue perturbation theory [22] gives

$$D\lambda_\star = y^\top DM_A[\Delta A] x.$$

Since $\lambda_\star \neq 0$, differentiating $|\lambda_\star|$ proves (38). The square law gives (39), and Cauchy–Schwarz yields (40). □

Theorems 5.2 and 5.3 describe complementary perturbation regimes. The finite bound remains valid for defective return matrices, whereas the derivative identifies the local conditioning of a simple dominant mode. Together with Remark 4.3, these results distinguish sensitivity of the asymptotic spectral factor from finite-sweep effects governed by powers of the return matrix.

# 6 Finite and near-finite convergence

The return representation also characterizes exact finite termination of a fixed IBMI sweep. We first relate nilpotency of the sweep and return matrices, then quantify the behavior of return dynamics close to nilpotency.

## 6.1 Nilpotent return dynamics

For a nilpotent square matrix $T$, define its nilpotency index by

$$\nu(T) := \min\{m \geq 1 : T^m = 0\}.$$

For an initial inverse error $E^{(0)}$, define the trajectory termination index, when it exists, by

$$\nu(Q, E^{(0)}) := \min\{m \geq 1 : Q^m E^{(0)} (Q^\top)^m = 0\}.$$

**Theorem 6.1** (Finite termination through the return operator)**.** *For every integer $p \geq 1$,*

$$Q^{p+1} = 0 \quad \Longleftrightarrow \quad BM^p = 0. \tag{41}$$

*Consequently, if $M^p = 0$, then $Q^{p+1} = 0$, and IBMI is exact after at most $p+1$ complete sweeps for every symmetric initial approximation:*

$$E^{(p+1)} = 0. \tag{42}$$

*If $B$ has full column rank, then*

$$Q^{p+1} = 0 \quad \Longleftrightarrow \quad M^p = 0. \tag{43}$$

*Moreover, for every integer $s \geq 1$,*

$$\mathcal{L}^s = 0 \text{ on } \mathbb{S}^n \quad \Longleftrightarrow \quad Q^s = 0. \tag{44}$$

*Proof* By (15),

$$Q^{p+1} = HE_{J_K} BM^p E_{J_1}^\top.$$

The matrix $HE_{J_K}$ has full column rank, whereas $E_{J_1}^\top$ has full row rank. Applying suitable left and right inverses therefore gives

$$Q^{p+1} = 0 \quad \Longleftrightarrow \quad BM^p = 0.$$

The stated consequences follow immediately, including the equivalence under the additional full-column-rank assumption on $B$. If $Q^s = 0$, then $\mathcal{L}^s = 0$. Conversely, if $\mathcal{L}^s$ vanishes on all symmetric matrices, applying it to $I_n$ yields

$$Q^s (Q^s)^\top = 0.$$

Hence, $Q^s = 0$, proving (44). □

The condition $BM^p = 0$ is exact without a rank assumption on $B$. Nilpotency of $M$ is sufficient for finite termination, but the nilpotency indices of $M$ and $Q$ need not coincide. Likewise, termination of the operator does not imply that every individual trajectory first terminates at the same sweep. When $Q$ is nilpotent,

$$\nu(Q, E^{(0)}) \leq \nu(Q),$$

and strict inequality may occur.

The simplest case is $M = 0$, which gives $Q^2 = 0$. The factorization $M = H_{J_1 J_K} B$ yields a structural sufficient condition.

**Corollary 6.2** (Endpoint inverse decoupling)**.** *If*

$$H_{J_1 J_K} = 0, \tag{45}$$

*then* $M = 0$, $Q^2 = 0$, *and*

$$E^{(2)} = 0$$

*for every symmetric initial error.*

*If* $B \neq 0$, *then* $Q \neq 0$ *and* $Q^2 = 0$*; thus the sweep operator has nilpotency index exactly two.*

*Proof* Equation (13) gives $M = 0$, and Theorem 6.1 then yields $Q^2 = 0$.

If $B \neq 0$, the full-column-rank property of $HE_{J_K}$ implies $HE_{J_K} B \neq 0$. Since $E_{J_1}^\top$ has full row rank,

$$Q = HE_{J_K} BE_{J_1}^\top \neq 0.$$

Therefore, $\nu(Q) = 2$. □

Since $H \succ 0$, condition (45) necessarily implies $J_1 \cap J_K = \varnothing$: a shared index would contribute a positive diagonal entry to $H_{J_1 J_K}$. The condition thus represents exact decoupling of two disjoint complement sets in the inverse matrix $H$.

The corresponding two-sweep termination mechanism is specific to the ordered IBMI return product. Section A examines its consequences for singleton complements.

## 6.2 Near-nilpotency and quadratic inverse-error collapse

We next quantify finite-sweep behavior when the return matrix is close to a nilpotent matrix. The bound is nonasymptotic and does not require diagonalizability.

**Theorem 6.3** (Near-nilpotent return bound)**.** *Suppose*

$$M = N + \Delta, \qquad N^p = 0, \qquad \|\Delta\|_2 \leq \varepsilon, \tag{46}$$

*and*

$$\max\{\|M\|_2, \|N\|_2\} \leq L. \tag{47}$$

*Then,*

$$\|M^p\|_2 \leq pL^{p-1}\varepsilon. \tag{48}$$

*For every symmetric initial error,*

$$\|E^{(p+1)}\|_2 \le \|U\|_2^2 p^2 L^{2p-2} \varepsilon^2 \|E^{(0)}\|_2. \tag{49}$$

*Proof* The telescoping identity gives

$$M^p - N^p = \sum_{j=0}^{p-1} M^{p-1-j}(M-N)N^j.$$

Since $N^p = 0$,

$$\begin{aligned}\|M^p\|_2 &\le \sum_{j=0}^{p-1} L^{p-1-j}\varepsilon L^j \\ &= pL^{p-1}\varepsilon.\end{aligned}$$

By Corollary 3.2,

$$E^{(p+1)} = UM^pV^\top E^{(0)} V (M^\top)^p U^\top.$$

Using $\|V\|_2 = 1$ and (48) proves (49). □

For a parameterized family, the $O(\varepsilon^2)$ interpretation requires $\|U\|_2$ and $L$ to remain uniformly bounded as $\varepsilon \to 0$. Without this uniformity, the bound does not assert a uniform quadratic collapse.

**Corollary 6.4** (Approximate endpoint decoupling)**.** *Let*

$$C := H_{J_1 J_K}, \qquad M = CB.$$

*If* $\|C\|_2 \le \varepsilon$*, then*

$$\rho(M) \le \|M\|_2 \le \varepsilon \|B\|_2, \tag{50}$$

*and hence,*

$$\rho(\mathcal{L}) = \rho(M)^2 \le \varepsilon^2 \|B\|_2^2. \tag{51}$$

*Moreover,*

$$\|E^{(2)}\|_2 \le \|U\|_2^2 \|B\|_2^2 \varepsilon^2 \|E^{(0)}\|_2. \tag{52}$$

*Proof* Since $M = CB$, submultiplicativity gives

$$\|M\|_2 \le \|C\|_2 \|B\|_2 \le \varepsilon \|B\|_2.$$

The square law proves (51). Finally, the exact return dynamics at $s = 2$ give

$$E^{(2)} = UMV^\top E^{(0)} V M^\top U^\top.$$

Taking spectral norms proves (52). □

The corollary identifies a quantitative near-finite regime. If the endpoint inverse coupling is $O(\varepsilon)$ and $\|B\|_2$ remains bounded, the return spectral radius is $O(\varepsilon)$, while the full inverse-error spectral factor is $O(\varepsilon^2)$. The two-sweep error bound is likewise quadratic provided $\|U\|_2$ is uniformly bounded.

For general multi-block covers, overlap changes both the endpoint coupling and the intermediate transfer product. The exact ordering-dependent spectral factor remains

$$\rho\big(H_{J_1J_K}H_{J_KJ_K}^{-1}H_{J_KJ_{K-1}}\cdots H_{J_2J_2}^{-1}H_{J_2J_1}H_{J_1J_1}^{-1}\big)^2. \tag{53}$$

Thus, increasing overlap need not improve the convergence factor for an arbitrary ordered multi-block cover. This dependence motivates the matched-overlap experiments in Section 9.

# 7 Overlap laws and their boundary

The $A$-only return certificate makes the dependence on an ordered block cover explicit. We first establish monotonicity under nested enlargement of two update blocks, then show by an exact three-block example that this law does not extend to arbitrary ordered covers. A canonical-correlation envelope provides a monotone upper bound for any number of blocks. The underlying principal-angle theory is classical [7]; its role here is to characterize the IBMI inverse-error factor.

**Theorem 7.1** (Two-block nested-overlap monotonicity)**.** *For a two-block sweep with nonempty complements $J_1, J_2$, let*

$$C := H_{J_1J_1}^{-1/2}H_{J_1J_2}H_{J_2J_2}^{-1/2}. \tag{54}$$

*Then,*

$$\rho(M) = \|C\|_2^2, \qquad \rho(\mathcal{L}) = \|C\|_2^4. \tag{55}$$

*If the update blocks are enlarged to nonempty proper blocks $I_k \subseteq I_k'$ for $k = 1, 2$, equivalently $J_k' \subseteq J_k$, then*

$$\rho(M') \le \rho(M), \qquad \rho(\mathcal{L}') \le \rho(\mathcal{L}). \tag{56}$$

*Proof* For $K = 2$, $M = H_{J_1J_2}H_{J_2J_2}^{-1}H_{J_2J_1}H_{J_1J_1}^{-1}$, which is similar to $CC^\top$ and proves (55). For the monotonicity statement, define

$$U_k := H^{1/2}E_{J_k}H_{J_kJ_k}^{-1/2}.$$

The columns of $U_k$ are orthonormal and $U_1^\top U_2 = C$. Hence, $\|C\|_2$ is the largest canonical correlation, equivalently the maximum absolute inner product between unit vectors in

$$\mathcal{S}_k := \operatorname{range}(H^{1/2}E_{J_k}).$$

If $J_k' \subseteq J_k$, then $\mathcal{S}_k' \subseteq \mathcal{S}_k$. The variational domain defining the largest canonical correlation therefore only shrinks:

$$\max_{\substack{x\in\mathcal{S}_1',\, y\in\mathcal{S}_2' \\ \|x\|_2=\|y\|_2=1}} |x^\top y| \le \max_{\substack{x\in\mathcal{S}_1,\, y\in\mathcal{S}_2 \\ \|x\|_2=\|y\|_2=1}} |x^\top y|.$$

Thus, $\|C'\|_2 \le \|C\|_2$, and (56) follows. □

This monotonicity concerns nested block enlargement, not arbitrary changes in overlap percentage. It can fail already for a three-block ordered cover.

**Theorem 7.2** (Three-block nonmonotonicity)**.** *There exists an SPD matrix and an ordered three-block cover for which enlarging one update block strictly increases the return spectral radius.*

*Proof* Let

$$H = \frac{1}{100}\begin{bmatrix} 784 & -163 & -265 & -24 & -357 & -115 \\ -163 & 598 & 192 & 101 & -427 & 235 \\ -265 & 192 & 260 & 34 & 5 & 159 \\ -24 & 101 & 34 & 855 & 113 & 197 \\ -357 & -427 & 5 & 113 & 842 & -17 \\ -115 & 235 & 159 & 197 & -17 & 217 \end{bmatrix}.$$

Its leading principal minors are

$$\frac{196}{25}, \quad \frac{442263}{10000}, \quad \frac{30339667}{500000}, \quad \frac{50841748123}{10^8}, \quad \frac{4880230217611}{5 \cdot 10^9}, \quad \frac{232688511158651}{10^{12}},$$

all positive, so $H \succ 0$. Take

$$J_1 = \{1, 3\}, \qquad J_2 = \{2, 3, 6\}, \qquad J_3 = \{4, 5\},$$

and enlarge the middle update block by replacing $J_2$ with $J_2' = \{2, 6\} \subset J_2$. The associated return matrices have the following characteristic polynomials, obtained by exact rational arithmetic:

$$p_{\text{before}}(x) = \frac{167100723926068535x^2 - 1338568090325562x + 1688270156316}{167100723926068535},$$

whereas

$$p_{\text{after}}(x) = \frac{2347659967233605x^2 - 29322446429644x - 1266602954844}{2347659967233605}.$$

The strict comparison can be certified without floating-point eigenvalues. Let the numerators of the two displayed polynomials be $q_{\text{before}}$ and $q_{\text{after}}$. Direct integer arithmetic gives

$$q_{\text{before}}(0) > 0, \qquad q_{\text{before}}\left(\frac{1}{500}\right) < 0, \qquad q_{\text{before}}\left(\frac{1}{100}\right) > 0.$$

Because $q_{\text{before}}$ is quadratic with positive leading coefficient, these two sign changes isolate its two real roots in $\left(0, \frac{1}{500}\right)$ and $\left(\frac{1}{500}, \frac{1}{100}\right)$, respectively. Hence, $\rho(M_{\text{before}}) < \frac{1}{100}$. For the enlarged middle block,

$$q_{\text{after}}\left(-\frac{1}{50}\right) > 0, \qquad q_{\text{after}}(0) < 0, \qquad q_{\text{after}}\left(\frac{1}{50}\right) < 0.$$

Since $q_{\text{after}}$ has positive leading coefficient and negative constant term, it has one negative and one positive root. The first two signs place the negative root in $(-1/50, 0)$, while $q_{\text{after}}(1/50) < 0$ forces the positive root to lie in $(1/50, \infty)$. Therefore,

$$\rho(M_{\text{after}}) > \frac{1}{50} > \frac{1}{100} > \rho(M_{\text{before}}).$$

Thus, enlarging the middle update block strictly worsens the exact return spectral radius. Numerically,

$$\rho(M_{\text{before}}) \approx 0.0064422599637, \qquad \rho(M_{\text{after}}) \approx 0.0302974317480.$$

□

*Remark* 7.3. Theorem 7.2 establishes that the two-block monotonicity law can fail already for $K = 3$. This example establishes failure for $K = 3$; it does not, by itself, provide a counterexample for each larger value of $K$.

**Theorem 7.4** (Multi-block canonical-correlation envelope)**.** *For an ordered $K$-block IBMI sweep with $K \geq 2$, define*

$$U_k := H^{1/2} E_{J_k} H_{J_k J_k}^{-1/2}, \qquad C_{k+1,k} := U_{k+1}^{\top} U_k,$$

*with cyclic indexing $U_{K+1} = U_1$, and set $c_{k+1,k} := \|C_{k+1,k}\|_2$. Then,*

$$\rho(M) \leq \prod_{k=1}^{K} c_{k+1,k}, \qquad \rho(\mathcal{L}) \leq \left( \prod_{k=1}^{K} c_{k+1,k} \right)^2. \tag{57}$$

*Each $c_{k+1,k}$ is the largest canonical correlation between $\mathcal{S}_k = \mathrm{range}(H^{1/2} E_{J_k})$ and $\mathcal{S}_{k+1} = \mathrm{range}(H^{1/2} E_{J_{k+1}})$. If one or more update blocks are enlarged, so that $J'_k \subseteq J_k$ for the affected indices, the product of canonical correlations on the right-hand side of* (57) *cannot increase. This does not imply monotonicity of the exact return spectral radius for $K \geq 3$. For $K = 2$ the first inequality is an equality and recovers Theorem 7.1.*

*Proof* The columns of every $U_k$ are orthonormal. The return matrix $M$ is similar to the normalized cyclic product

$$\widehat{M} = C_{1,K} C_{K,K-1} \cdots C_{2,1}, \qquad \widehat{M} = H_{J_1 J_1}^{-1/2} M H_{J_1 J_1}^{1/2}.$$

Consequently,

$$\rho(M) = \rho(\widehat{M}) \leq \|\widehat{M}\|_2 \leq \prod_{k=1}^{K} \|C_{k+1,k}\|_2,$$

and Theorem 3.3 yields the second bound. If $J'_k \subseteq J_k$, then $\mathcal{S}'_k \subseteq \mathcal{S}_k$. The variational characterization of largest canonical correlation implies that both adjacent factors $c_{k,k-1}$ and $c_{k+1,k}$ cannot increase; all nonadjacent factors are unchanged. Repeating this argument proves the monotonicity of the envelope. For $K = 2$, the two factors are transposes of one another and their product is $\|C_{12}\|_2^2 = \rho(M)$ by Theorem 7.1. □

**Theorem 7.5** ($A$-only quantitative contraction from separated complements)**.** *For an ordered cover with $K \geq 2$, use cyclic indexing $J_{K+1} = J_1$. For each adjacent pair $(J_k, J_{k+1})$, let $S_k := J_k \cup J_{k+1}$ and $T_k := S_k^c$. Whenever $J_k \cap J_{k+1} = \varnothing$, define the Schur complement*

$$\Sigma_k := \begin{cases} A_{S_k S_k} - A_{S_k T_k} A_{T_k T_k}^{-1} A_{T_k S_k}, & T_k \neq \varnothing, \\ A, & T_k = \varnothing, \end{cases} \qquad \vartheta_k := \frac{\kappa_2(\Sigma_k) - 1}{\kappa_2(\Sigma_k) + 1}.$$

*Let* $\mathcal{D} := \{k : J_k \cap J_{k+1} = \varnothing\}$*, with cyclic indexing. Then*

$$\rho(\mathcal{L}) \le \prod_{k\in\mathcal{D}} \vartheta_k^2 \le \left(\frac{\kappa_2(A)-1}{\kappa_2(A)+1}\right)^{2|\mathcal{D}|}. \tag{58}$$

*If* $\mathcal{D} \neq \varnothing$*, at least one separated adjacent complement pair gives the strict* $A$*-computable bound* $\rho(\mathcal{L}) < 1$*. When* $\mathcal{D} = \varnothing$*, the empty product is one and this envelope gives no strict contraction. If a family of problems satisfies* $|\mathcal{D}| \ge d_0 \ge 1$ *and* $\kappa_2(A) \le \kappa_0 < \infty$*, then its IBMI factors obey the dimension-independent bound*

$$\rho(\mathcal{L}) \le \left(\frac{\kappa_0-1}{\kappa_0+1}\right)^{2d_0} < 1. \tag{59}$$

*Proof* Fix $k \in \mathcal{D}$. Since $J_k$ and $J_{k+1}$ are disjoint coordinate sets, every $x \in \operatorname{range}(E_{J_k})$ and $y \in \operatorname{range}(E_{J_{k+1}})$ satisfy $x^\top y = 0$. By the variational characterization of canonical correlation,

$$c_{k+1,k} = \sup_{0\neq x\in\operatorname{range}(E_{J_k}),\ 0\neq y\in\operatorname{range}(E_{J_{k+1}})} \frac{|x^\top H y|}{(x^\top H x)^{1/2}(y^\top H y)^{1/2}}.$$

Only the principal block $H_{S_kS_k}$ enters this quotient. The classical Wielandt inequality for an SPD matrix $G$ and Euclidean-orthogonal vectors gives

$$\frac{|x^\top G y|}{(x^\top G x)^{1/2}(y^\top G y)^{1/2}} \le \frac{\kappa_2(G)-1}{\kappa_2(G)+1};$$

see, for example, the matrix extensions discussed in [23]. Taking $G = H_{S_kS_k}$ therefore yields $c_{k+1,k} \le \vartheta_k$, because the block inverse identity gives

$$H_{S_kS_k} = \Sigma_k^{-1}, \qquad \kappa_2(H_{S_kS_k}) = \kappa_2(\Sigma_k).$$

For $k \notin \mathcal{D}$ we use only $c_{k+1,k} \le 1$. Theorem 7.4 now gives the first inequality in (58). Since $H_{S_kS_k}$ is a principal submatrix of $H$, eigenvalue interlacing implies $\kappa_2(\Sigma_k) = \kappa_2(H_{S_kS_k}) \le \kappa_2(H) = \kappa_2(A)$, which proves the second inequality. Equation (59) follows immediately. □

*Remark* 7.6. If $J_k \cap J_{k+1} \neq \varnothing$, then

$$H^{1/2}\operatorname{range}(E_{J_k\cap J_{k+1}}) \subseteq \mathcal{S}_k \cap \mathcal{S}_{k+1}.$$

Consequently, the largest canonical correlation satisfies $c_{k+1,k} = 1$. Thus, shared complement directions necessarily contribute a unit factor to the largest-canonical-correlation envelope.

Theorem 7.5 identifies separated adjacent complements as the pairs that yield explicit strict contraction factors through $A$-only Schur-complement bounds. For intersecting complements, strict contraction is governed by the geometry beyond their common directions, which can be captured through quotient-space or Friedrichs-angle analysis rather than the largest canonical correlation alone.

**Corollary 7.7** (Two-sweep criterion for multiple blocks)**.** *If $c_{k+1,k} = 0$ for at least one cyclic adjacent pair, then $M = 0$, $Q^2 = 0$, and every inverse error is annihilated after at most two complete sweeps.*

*Proof* A zero factor means $C_{k+1,k} = 0$, so the cyclic product $\widehat{M} = 0$. Since $M$ is similar to $\widehat{M}$, $M = 0$. Equation (15) gives $Q^2 = 0$, and hence $E^{(2)} = Q^2 E^{(0)} (Q^\top)^2 = 0$. □

Theorem 7.2 and Theorem 7.4 identify a useful boundary of the overlap law. Beyond two blocks, the exact spectral radius need not be monotone, because cyclic alignment of the transfer factors can change. What survives for arbitrary $K$ is a monotone geometric upper envelope built from adjacent canonical correlations. Thus nested block enlargement can worsen the exact rate while making the pairwise envelope no larger. The distinction reflects the cyclic alignment of the transfer factors, which their individual norms do not capture.

# 8 Ordering structure and certificate-guided design

For singleton complements, the return certificate reduces to a weighted cycle product. Assume $n \geq 3$ for the Hamiltonian-cycle statements, and let $J_i = \{i\}$, so $I_i = [n] \setminus \{i\}$.

**Theorem 8.1** (Singleton cycle formula)**.** *For an ordered sequence $\pi = (i_1, \ldots, i_s)$ of distinct indices,*

$$Q_\pi = \frac{\prod_{r=1}^{s-1} h_{i_{r+1} i_r}}{\prod_{r=1}^{s} h_{i_r i_r}} H e_{i_s} e_{i_1}^\top.$$

*For a permutation of $[n]$, the only possible nonzero eigenvalue is*

$$\lambda_\pi = \frac{h_{i_1 i_n} \prod_{r=1}^{n-1} h_{i_{r+1} i_r}}{\prod_{r=1}^{n} h_{i_r i_r}}, \qquad \rho(\mathcal{L}_\pi) = |\lambda_\pi|^2. \tag{60}$$

*Equivalently, with $c_{ij} = h_{ij}/\sqrt{h_{ii} h_{jj}}$,*

$$|\lambda_\pi| = |c_{i_1 i_n}| \prod_{r=1}^{n-1} |c_{i_{r+1} i_r}|.$$

*Proof* For a singleton complement, $Q_i = H e_i e_i^\top / h_{ii}$. Successive multiplication produces one factor $e_j^\top H e_i = h_{ji}$ at each transition. Closing the rank-one product produces the endpoint factor $h_{i_1 i_n}$. □

**Corollary 8.2** (Path–cycle dichotomy)**.** *If all consecutive couplings along $\pi$ are nonzero, then $Q_\pi \neq 0$ and*

$$Q_\pi^2 = 0 \quad \Longleftrightarrow \quad h_{i_1 i_n} = 0.$$

*Thus, an ordered Hamiltonian path of nonzero couplings in the graph of $H$ whose endpoints are not adjacent yields genuine two-sweep exactness.*

The singleton return space is one-dimensional. Hence, any finite-terminating singleton-complement sweep has nilpotency depth at most two; higher finite-sweep indices, if realizable, require a multidimensional return space.

## 8.1 Logarithmic inverse-coupling geometry

The cycle product in Theorem 8.1 exposes a useful additive coordinate system for ordering. The logarithmic transformation itself is classical; for example, $-\log|\rho_{ij}|$ is an information distance in Gaussian latent-tree models and is additive on tree-structured models [24]. Here, its role is different: it is induced by the exact IBMI convergence factor.

For $i \neq j$, set $w_{ij} = |c_{ij}|$. When $w_{ij} > 0$, define

$$d_{ij} := -\log w_{ij}.$$

If $w_{ij} = 0$, use the extended value $d_{ij} = +\infty$.

**Theorem 8.3** (Exact log-cycle ordering geometry)**.** *For every singleton-complement permutation $\pi$ whose cycle uses only positive couplings,*

$$\rho(\mathcal{L}_\pi) = \exp\{-2D(\pi)\}, \qquad D(\pi) := \sum_{\{i,j\}\in C_\pi} d_{ij}. \tag{61}$$

*Consequently, minimizing the exact IBMI asymptotic factor is equivalent to maximizing the symmetric Hamiltonian-cycle weight $D(\pi)$. For two such orders $\pi$ and $\pi'$, the exact comparison is*

$$\frac{\rho(\mathcal{L}_{\pi'})}{\rho(\mathcal{L}_\pi)} = \exp\{-2[D(\pi') - D(\pi)]\}. \tag{62}$$

*If a cycle contains a zero coupling, its factor is zero and the corresponding singleton sweep is finite terminating.*

*Proof* Theorem 8.1 gives $\rho(\mathcal{L}_\pi) = \prod_{\{i,j\}\in C_\pi} w_{ij}^2$. Taking logarithms proves (61) and (62). A zero factor gives $\rho(\mathcal{L}_\pi) = 0$ directly from the same product formula. □

Theorem 8.3 identifies the exact deterministic ordering objective with a maximum-weight Hamiltonian cycle in the induced logarithmic couplings. This is an IBMI consequence of the return formula, not a claim of a new Max-TSP formulation or approximation algorithm.

**Theorem 8.4** (SPD realization of arbitrary ordering geometry)**.** *Let $n \geq 3$ and let $q_{ij} = q_{ji} \geq 0$ be arbitrary finite edge weights on the complete graph on $n$ vertices.*

*There exists an SPD matrix $A$ and a constant $\kappa$ independent of the permutation such that its singleton-complement IBMI factors satisfy*

$$-\frac{1}{2}\log\rho(\mathcal{L}_\pi) = \kappa + \sum_{\{i,j\}\in C_\pi} q_{ij} \qquad \textit{for every } \pi. \tag{63}$$

*Hence, arbitrary symmetric maximum-weight Hamiltonian-cycle ordering geometry can be realized exactly by an SPD singleton-IBMI instance.*

*Proof* Choose $L > 0$ sufficiently large that $(n-1)e^{-L} < 1$, and define a symmetric matrix $H$ by

$$h_{ii} = 1, \qquad h_{ij} = \exp[-(L + q_{ij})], \quad i \neq j.$$

For every row, $\sum_{j\neq i} |h_{ij}| \leq (n-1)e^{-L} < 1$; hence $H$ is SPD by strict diagonal dominance. Set $A = H^{-1} \succ 0$. Since $h_{ii} = 1$, the normalized inverse couplings satisfy $c_{ij} = h_{ij} = e^{-(L+q_{ij})}$, so $d_{ij} = L + q_{ij}$. Every Hamiltonian cycle has exactly $n$ edges, and Theorem 8.3 yields

$$-\frac{1}{2}\log\rho(\mathcal{L}_\pi) = nL + \sum_{\{i,j\}\in C_\pi} q_{ij}.$$

Thus, (63) holds with $\kappa = nL$. □

The realization theorem establishes the full expressive power of SPD singleton-IBMI ordering. In particular, every symmetric Max-TSP instance can be embedded into an SPD inverse-coupling matrix, with the resulting logarithmic cycle objective agreeing with the prescribed edge-weight objective up to an order-independent constant. Thus, positive definiteness alone imposes no additional combinatorial restrictions on the ordering problem. Metric approximation guarantees become applicable only when the induced logarithmic edge weights satisfy the corresponding metric assumptions.

**Theorem 8.5** (Exact 2-opt convergence certificate)**.** *Suppose a positive-coupling cycle $\pi$ contains nonadjacent boundary edges $\{a, b\}$ and $\{c, d\}$. Let $\pi'$ be the 2-opt reversal replacing them by $\{a, c\}$ and $\{b, d\}$. If both new boundary couplings are positive, then*

$$\frac{\rho(\mathcal{L}_{\pi'})}{\rho(\mathcal{L}_\pi)} = \left(\frac{w_{ac}w_{bd}}{w_{ab}w_{cd}}\right)^2 = e^{-2\Delta}, \tag{64}$$

*where*

$$\Delta = d_{ac} + d_{bd} - d_{ab} - d_{cd}. \tag{65}$$

*Thus, when the new couplings are positive, the move strictly improves the exact asymptotic factor if and only if $\Delta > 0$. If either new boundary coupling vanishes, the new order has factor zero and the logarithmic gain is not used.*

*Proof* Because $w_{ij} = w_{ji}$, reversing the interior path preserves every interior edge factor. Only the two boundary factors change. Cancelling the common product in Theorem 8.1 gives

(64); taking logarithms gives (65). The zero-coupling case follows directly from the cycle product and is stated separately to avoid a zero denominator. □

**Corollary 8.6** (Finite monotone certificate-guided 2-opt)**.** *Starting from any singleton order with positive factor, repeatedly accepting strict positive-coupling 2-opt improvements terminates after finitely many moves. If the accepted gains are* $\Delta_0, \dots, \Delta_{m-1} > 0$*, then*

$$\rho(\mathcal{L}_{\pi_m}) = \rho(\mathcal{L}_{\pi_0}) \exp\left( -2 \sum_{t=0}^{m-1} \Delta_t \right). \tag{66}$$

*At a terminal order for this positive-coupling search, no admissible positive-coupling 2-opt move has* $\Delta > 0$*. A move introducing a zero coupling can instead be accepted as an immediate zero-factor termination. Once the* $d_{ij}$ *are cached, each positive-coupling candidate gain requires only four edge weights.*

*Proof* Theorem 8.5 gives strict decrease at every accepted move and (66) by multiplication. There are finitely many cyclic orders, so strict descent cannot continue indefinitely. The terminal condition and constant-size gain evaluation are immediate. □

**Corollary 8.7** (Structured metric guarantee)**.** *Assume all couplings are positive and the induced logarithmic weights satisfy*

$$d_{ij} \le d_{ik} + d_{kj} \qquad \text{for all distinct } i, j, k, \tag{67}$$

*equivalently* $|c_{ij}| \ge |c_{ik}||c_{kj}|$*. Let* $\pi_\star$ *minimize the IBMI factor. A deterministic* 7/8*-approximation algorithm for metric Max-TSP [5] returns an order* $\widehat{\pi}$ *satisfying*

$$-\log \rho(\mathcal{L}_{\widehat{\pi}}) \ge \frac{7}{8}[-\log \rho(\mathcal{L}_{\pi_\star})], \qquad \rho(\mathcal{L}_{\widehat{\pi}}) \le \rho(\mathcal{L}_{\pi_\star})^{7/8}. \tag{68}$$

*The guarantee concerns the asymptotic contraction exponent; it is not a uniform finite-sweep bound in the presence of nonnormal transients.*

*Proof* Condition (67), together with positivity of $d_{ij}$ for distinct indices (which follows from $H \succ 0$), makes $d$ a metric. The cited metric Max-TSP algorithm gives $D(\widehat{\pi}) \ge (7/8)D(\pi_\star)$. Apply Theorem 8.3 and exponentiate. The final qualification follows from Proposition 3.4. □

**Theorem 8.8** (Exact global ordering for exponential line couplings)**.** *Let* $x_1 < \cdots < x_n$*,* $\beta > 0$*, and let*

$$H = DCD, \qquad C_{ij} = \exp(-\beta|x_i - x_j|),$$

*where* $D$ *is any positive diagonal matrix. Then* $H \succ 0$*, and the singleton-complement IBMI ordering problem has the exact objective* $d_{ij} = \beta|x_i - x_j|$*. Writing* $g_k = x_{k+1} - x_k$*,*

*its globally optimal log-cycle value is*

$$D_\star = 2\beta \sum_{k=1}^{n-1} \min\{k, n-k\}\, g_k, \tag{69}$$

*and the globally minimal inverse-error factor is*

$$\rho_\star = \exp\left( -4\beta \sum_{k=1}^{n-1} \min\{k, n-k\}\, g_k \right). \tag{70}$$

*An optimal order is obtained by alternating the lower and upper halves. For $n = 2m$ one may take*

$$\pi_\star = (1, m+1, 2, m+2, \ldots, m, 2m),$$

*and for $n = 2m+1$ one may take*

$$\pi_\star = (1, m+2, 2, m+3, \ldots, m, 2m+1, m+1).$$

*Moreover, $A = H^{-1}$ is tridiagonal. Thus this is an exactly solvable IBMI ordering class with a sparse tridiagonal input matrix $A$.*

*Proof* Put $r_k = \exp(-\beta g_k) \in (0,1)$. The Gaussian Markov recursion

$$z_1 = \varepsilon_1, \qquad z_{k+1} = r_k z_k + \sqrt{1 - r_k^2}\, \varepsilon_{k+1},$$

with independent standard normal $\varepsilon_k$ has covariance $\mathrm{Cov}(z_i, z_j) = \exp(-\beta|x_i - x_j|) = C_{ij}$. Hence, $C \succ 0$. The quadratic form of the joint density is a sum of terms $(z_{k+1} - r_k z_k)^2/(1 - r_k^2)$, so $C^{-1}$ is tridiagonal. Positive diagonal congruence gives $H \succ 0$ and $A = D^{-1}C^{-1}D^{-1}$, which is again tridiagonal. The normalized couplings are $c_{ij} = C_{ij}$, and therefore $d_{ij} = \beta|x_i - x_j|$.

For a Hamiltonian cycle $\pi$, let $N_k(\pi)$ be the number of tour edges crossing the cut $\{1, \ldots, k\} \mid \{k+1, \ldots, n\}$. Decomposing every line distance over the gaps gives

$$D(\pi) = \beta \sum_{k=1}^{n-1} N_k(\pi) g_k.$$

A cycle has degree two at every vertex, so $N_k(\pi) \le 2\min\{k, n-k\}$. The displayed alternating-half orders attain equality simultaneously for every cut $k$. They therefore attain the upper bound for every choice of positive gaps, proving (69). Equation (70) follows from Theorem 8.3. □

*Remark* 8.9. The underlying longest-tour structure for this one-dimensional distance geometry is classical and belongs to the Supnick family [6]. The new point here is its exact transfer through the IBMI return certificate: Theorem 8.8 identifies a natural sparse precision class for which both the globally rate-optimal sweep order and its optimal contraction factor are available in closed form. This provides a positive structured counterpart to the unrestricted SPD hardness result below.

The preceding realization theorem uses real exponential weights. The next result has a different purpose: it gives a polynomial-size rational encoding suitable for a complexity-theoretic reduction. Hamiltonian Circuit remains NP-complete even under strong graph restrictions [12]; we use it only as the source problem.

**Theorem 8.10** (NP-hard exact singleton ordering)**.** *Consider the decision problem whose input is a rational SPD matrix $A$ and a rational threshold $\tau > 0$, and which asks whether some singleton-complement sweep order $\pi$ satisfies $\rho(\mathcal{L}_\pi) \le \tau$. This problem is NP-hard, even on instances for which the inverse matrix $H = A^{-1}$ has unit diagonal.*

*Proof* Reduce from undirected Hamiltonian Cycle. Given $G = (V, E)$ with $n = |V|$, set $\varepsilon = 1/(4n)$ and define the rational unit-diagonal matrix

$$h_{ij} = \begin{cases} \varepsilon, & \{i,j\} \in E, \\ 2\varepsilon, & \{i,j\} \notin E, \end{cases} \qquad i \neq j.$$

The off-diagonal absolute row sum is at most $2\varepsilon(n-1) = (n-1)/(2n) < 1$, so $H$ is SPD by strict diagonal dominance. Its entries have $O(\log n)$ bits. Set $A = H^{-1}$. By the adjugate formula (equivalently, exact rational Gaussian elimination), every entry of $A$ is a ratio of determinants of rational matrices of order at most $n$; standard determinant bit bounds therefore give a representation of $A$ with polynomial encoding length, computable in polynomial time.

Because $h_{ii} = 1$, the full inverse-error factor is

$$\rho(\mathcal{L}_\pi) = |\lambda_\pi|^2 = \prod_{\{i,j\} \text{ on the cycle induced by } \pi} h_{ij}^2.$$

If $G$ has a Hamiltonian cycle, its corresponding order has factor exactly $\varepsilon^{2n}$. Conversely, if the induced cycle contains a nonedge, its factor is at least

$$(2\varepsilon)^2 \varepsilon^{2(n-1)} = 4\varepsilon^{2n}.$$

Set the rational threshold $\tau = \varepsilon^{2n}$. Its encoding length is $O(n \log n)$. The constructed IBMI instance is therefore a YES instance if and only if $G$ has a Hamiltonian cycle, giving a polynomial-time many-one reduction. □

Theorem 8.10 concerns the ordering problem induced specifically by the IBMI singleton return cycle. Ordering effects, adaptive strategies, and randomized reorderings are well established in the literature on multiplicative Schwarz, subspace-correction, SOR, and related methods [8–11]. The contribution here is the exact inverse-coupling cycle objective arising from the IBMI return representation and the complexity characterization of its rate-optimal ordering problem.

## 8.2 Factor once, score many

For a fixed block collection, Theorem 4.1 permits reuse of local factorizations across candidate orderings.

**Proposition 8.11** (Cached order scoring)**.** *Fix update blocks $\{I_1, \dots, I_K\}$ and form once*

$$T_k = E_{I_k}(-A_{I_k I_k}^{-1} A_{I_k J_k}) + E_{J_k}.$$

*For any permutation $\pi$ of the fixed blocks,*

$$\mathcal{R}_{J_{\pi_{r+1}} \leftarrow J_{\pi_r}} = E_{J_{\pi_{r+1}}}^{\top} T_{\pi_r}.$$

*Hence, after the $K$ local graph maps are available, an arbitrary permutation can be scored through reduced restrictions, products, and a numerical evaluation of $\rho(M_A)$, with no additional solve involving $A_{I_k I_k}$. The resulting score equals the exact asymptotic factor in exact arithmetic; numerical eigenvalue errors remain possible.*

*Proof* Apply Theorem 4.1 with source complement $J_{\pi_r}$ and target $L = J_{\pi_{r+1}}$. □

We refer to budgeted search over these cached exact scores as *certificate-guided sweep ordering* (CGSO). The NP-hardness and realization theorems preclude interpreting a generic local search as an exact global optimizer. For singleton complements, Theorem 8.5 provides the sharper alternative: local proposals can be evaluated from four cached logarithmic inverse couplings with an exact algebraic change in the IBMI asymptotic factor.

# 9 Numerical experiments

The numerical study examines three consequences of the return-operator theory. First, the reduced $A$-only certificate must reproduce the spectrum of an independently assembled complete-sweep operator. Second, its asymptotic factor should explain the observed inverse-error trajectories, while remaining distinct from finite-sweep behavior. Third, overlap and ordering should affect convergence through the transfer geometry identified by the theory, rather than through block size or permutation alone. We distinguish these questions using covariance models, public SPD matrices, and deliberately structured inverse-coupling ensembles. All reported numerical values are drawn from the consolidated experimental records. Failed evaluations and duplicate matrix inputs are explicitly identified rather than silently excluded.

## 9.1 Covariance models: independent validation of the spectral certificate

We consider EXP, RBF, inverse-quadratic, Matérn-3/2, and Matérn-5/2 covariance families, following [1]. The experiment comprises 30 configurations with $n = 64, 128, 256$ in one dimension and $n = 64, 144, 256$ on two-dimensional grids, using four update blocks and nominal 5% overlap.

The condition-aware implementation successfully evaluates the certificate in 25 configurations. The remaining five encounter numerical nonpositive failures: all three two-dimensional RBF configurations and the two-dimensional inverse-quadratic configurations at $n = 144, 256$. No nugget regularization is applied. These outcomes identify the numerical SPD admissibility of covariance discretizations as an important consideration when applying the certificate in floating-point arithmetic.

We independently assemble the complete-sweep operator for 20 designated smaller configurations, obtaining 17 successful comparisons. Across these comparisons, the

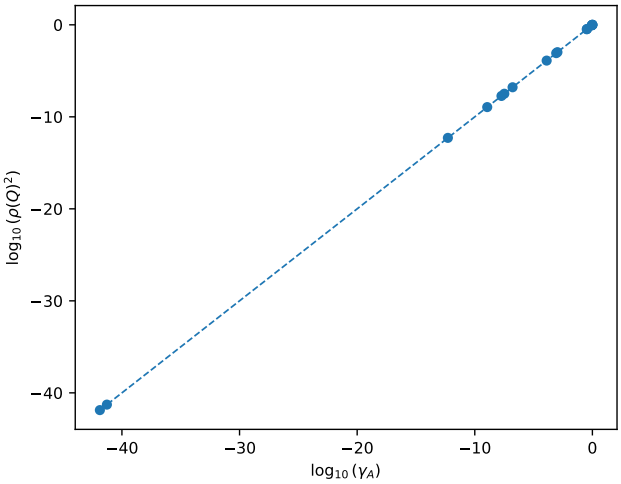



**Fig. 1** Independent spectral-identity check on the 17 successful covariance configurations. The $A$-only certificate is compared with the spectral factor of the separately assembled sweep operator (both axes logarithmic); the dashed line denotes equality.

maximum absolute discrepancy between the logarithmic spectral factors is $4.9434 \times 10^{-10}$, as illustrated in Fig. 1. The close agreement between the independently assembled sweep operator and the $A$-only local-transfer construction provides direct numerical validation of the spectral identity in Theorem 4.2. In particular, the reduced certificate reproduces the persistent sweep spectrum to the reported numerical accuracy without explicitly forming the full inverse.

## 9.2 Public SPD matrices: predictive accuracy and overlap effects

We evaluate the spectral certificate on eight named entries `HB/bcsstk01`–`HB/bcsstk08` from the SuiteSparse Matrix Collection [25]. The experiments use diagonal congruence scaling, four update blocks, overlap levels $0\%, 5\%, 10\%, 20\%$, and three sweep orderings: native, fixed random, and certificate-minimizing. Since `bcsstk07` duplicates `bcsstk06`, the 96 archived trajectories correspond to seven distinct matrix inputs. Both named entries are retained for reproducibility.

Across the 96 trajectories, the median absolute discrepancy between the certificate and the observed pre-floor propagation factor is $5.884 \times 10^{-7}$, with a 90th-percentile discrepancy of $6.317 \times 10^{-5}$. These results demonstrate close agreement between the pre-iteration spectral certificate and the contraction factors estimated from the actual inverse-error trajectories. The agreement supports the predictive relevance of the return spectrum for persistent convergence behavior. Finite-sweep transients remain governed by the powers of the sweep operator, as characterized in Proposition 3.4.

Table 1 reveals a pronounced dependence of overlap effectiveness on the underlying matrix. The native-order certificate is nonincreasing across the four tested overlap levels for every named entry, but the magnitude of improvement varies substantially.

**Table 1** Native-order spectral certificate $\gamma_A$ for scaled SuiteSparse entries. The last column is $\gamma_A(0\%)/\gamma_A(20\%)$, a ratio of asymptotic factors rather than a runtime speedup. The entries `bcsstk06` and `bcsstk07` are duplicate inputs.

| Matrix | 0% | 5% | 10% | 20% | Ratio |
|---|---|---|---|---|---|
| bcsstk01 | .993623 | .993467 | .993467 | .993467 | 1.000 |
| bcsstk02 | .986422 | .978287 | .971695 | .940730 | 1.049 |
| bcsstk03 | .983275 | .891011 | .678239 | .048749 | 20.170 |
| bcsstk04 | .985801 | .980258 | .974219 | .937667 | 1.051 |
| bcsstk05 | .987576 | .977014 | .959700 | .914594 | 1.080 |
| bcsstk06 | .996790 | .994596 | .991774 | .986704 | 1.010 |
| bcsstk07 | .996790 | .994596 | .991774 | .986704 | 1.010 |
| bcsstk08 | .996312 | .995799 | .989490 | .970852 | 1.026 |

For `bcsstk01`, increasing overlap from 0% to 20% changes the certificate only from 0.993623 to 0.993467. In contrast, the same overlap increase reduces the certificate for `bcsstk03` from 0.983275 to 0.048749, a ratio of 20.170.

This contrast illustrates a central implication of the transfer-space analysis: the convergence effect of overlap depends on the geometry of the complementary subspaces and their ordered interactions, rather than on the nominal overlap percentage alone. The observed improvements are consistent with the canonical-correlation framework, while the three-block counterexample in Theorem 7.2 establishes that monotonic improvement of the exact spectral factor does not hold for arbitrary overlapping sweeps.

The corresponding `bcsstk03` trajectories provide an iteration-level interpretation of this spectral contrast. Under native ordering, the archived 0%-overlap run reaches the 300-sweep cap, while the $5\%, 10\%, 20\%$ runs stop after 150, 45, and 7 sweeps, respectively (Fig. 2). The direction of the observed change agrees with the large reduction in the certificate. The factor ratio of 20.170 should not, however, be interpreted as a 20.170-fold reduction in iteration count: spectral factors, stopping tolerances, initial-error components, and finite-sweep transients enter differently.

This observed monotonicity across the tested public configurations is not a general overlap theorem. Theorem 7.1 guarantees monotonicity in the two-block setting, whereas Theorem 7.2 constructs an SPD three-block enlargement that worsens the exact factor. The experiments and the counterexample therefore have complementary roles: the former identifies cases where overlap is beneficial, and the latter establishes why that benefit cannot be inferred from enlargement alone.

## 9.3 Ordering sensitivity and inverse-coupling geometry

The public SPD experiments exhibit limited ordering sensitivity. Across the 32 matrix–overlap pairs, the median native/best certificate ratio is 1.000000, with a maximum of 1.000517. These records include the documented duplicate matrix input. Thus, for the tested public matrices and block configurations, optimizing the sweep order produces only marginal improvements in the asymptotic spectral factor.

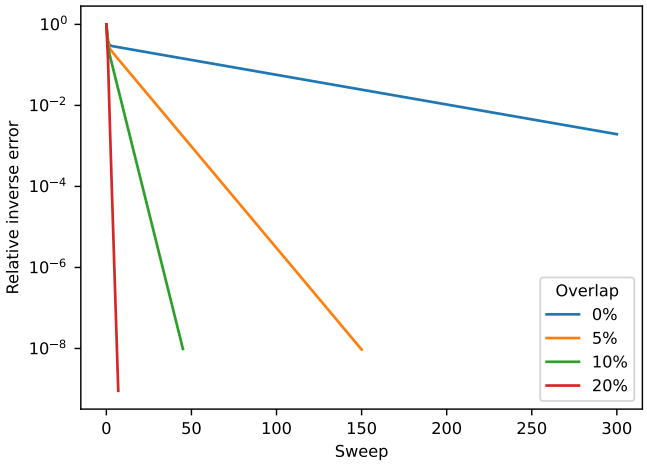



**Fig. 2** Archived relative inverse-error trajectories for `bcsstk03` under native ordering. The marked change in observed convergence across the four overlap levels is consistent with the corresponding spectral certificates in Table 1.

To examine the ordering mechanism under more heterogeneous coupling structures, we construct a controlled ensemble by prescribing SPD inverse-coupling matrices $H$ and setting $A = H^{-1}$. Each instance has dimension $n = 48$ and a fixed collection of $K = 6$ update blocks. We exhaustively evaluate all $6! = 720$ sweep orders for 20 instances in each of three inverse-coupling regimes: dense, heterogeneous, and sparse.

The median native/best certificate ratios are 6.7645, 19.6567, and 31.0891 in the three regimes, respectively, with a maximum of 115.7565 across all 60 instances (Fig. 3). Since the update blocks remain fixed within each instance, these differences arise solely from their ordering. In particular, changing the permutation can substantially alter the persistent convergence factor without changing the matrix or the collection of local update spaces.

The contrast between the two experiment families highlights the geometric nature of ordering sensitivity. The spectral effect of a permutation depends on the ordered composition of local transfers, which is determined by the inverse-coupling structure and the selected block complements. This is precisely the dependence captured by the return operator in Theorem 4.2. The exact $A$-only certificate therefore provides a principled criterion for comparing sweep orders before iteration. The controlled instances establish that substantial ordering effects are attainable within SPD IBMI, while the public-matrix results show that such effects need not be pronounced in every application setting.

The singleton-complement experiments further connect this numerical observation to the cycle formula in Section 8. At $n = 8$, the archived exhaustive search evaluates all $8! = 40320$ permutations, corresponding to 2520 distinct unoriented Hamiltonian cycles. The dense and heterogeneous regimes have median native/optimal certificate ratios of $1.579 \times 10^4$ and $2.712 \times 10^8$, respectively. These ratios can be large because the certificates themselves may be extremely small; they are not iteration-count or

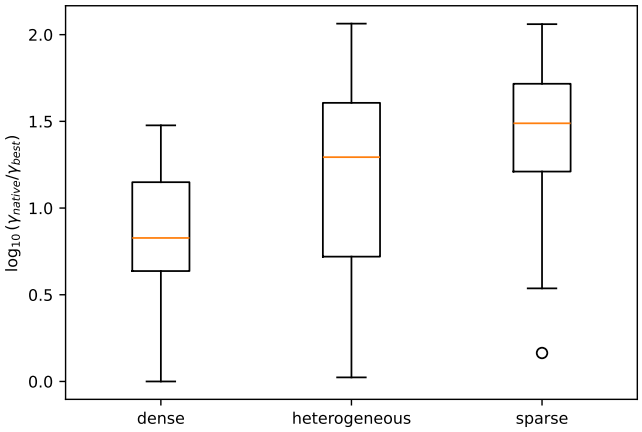



**Fig. 3** Native/best spectral-certificate ratios for the 60 general fixed-block ordering instances, grouped by inverse-coupling regime. The block collection is unchanged across permutations within each instance.

wall-clock speedup factors. Zero products in the sparse regime correspond to the exact finite-termination mechanism and are excluded from positive-ratio summaries. This behavior is precisely what the weighted-cycle characterization predicts: order selection changes the return product, and missing cycle edges annihilate it.

Exhaustive search also supplies a useful check on optimization claims. The targeted $n = 8$ instance in Appendix B.1 has a 2-opt factor 2.32277 times the global optimum. Hence exact local 2-opt gain formulas certify the effect of a proposed exchange, but they do not convert local search into a globally optimal ordering method. This is consistent with the unrestricted Max-TSP realization and NP-hardness results.

## 9.4 Theorem diagnostics and computational implications

We complement the convergence experiments with numerical checks of the contraction bounds and first-order sensitivity formula. In four tested instances, the computed quantities satisfy

$$\gamma_A \leq b_{\text{canonical}} \leq b_{\text{local}} \leq b_{\text{global}}.$$

These results confirm the predicted hierarchy of the exact spectral factor and its successive upper bounds in the tested configurations. The hierarchy also illustrates the distinction between the exact return certificate and more readily computable contraction estimates, whose sharpness depends on the underlying subspace geometry.

The sensitivity experiment examines the first-order approximation of the spectral factor under matrix perturbations. For

$$t = 10^{-2}, \quad 3 \times 10^{-3}, \quad 10^{-3}, \quad 3 \times 10^{-4},$$

the absolute Taylor residual normalized by $t^2$ remains between 0.03391 and 0.03393. The stable normalized residual is consistent with a quadratic remainder in the tested simple-dominant-eigenvalue regime, supporting the local sensitivity formula

in Theorem 5.3. This behavior is a local numerical observation; the differentiability assumptions remain essential near dominant-eigenvalue crossings or multiplicities.

We also evaluate the computational benefit of reusing local factorizations during certificate-based ordering comparisons. Across three repeated timing measurements of 24 order evaluations at $n = 40$, the median cached-scoring speedup is $2.960\times$, while the median amortized speedup is $2.738\times$. These measurements demonstrate a reduction in the cost of evaluating multiple sweep orders within the tested implementation.

The computational implication is specific: once the local block data have been constructed, alternative orderings can be compared through their reduced return operators without independently rebuilding the full sweep operator for every permutation. The reported timing gains concern certificate evaluation and do not represent end-to-end acceleration of the IBMI iteration, which also depends on ordering-search overhead and the number of sweeps required. Additional ordering and finite-termination diagnostics are reported in Appendix B.

***Synthesis.***

The numerical results establish a consistent connection between the return-operator theory and the observed behavior of overlapping IBMI. Independent operator comparisons validate the reduced $A$-only spectral identity to the reported numerical accuracy, while inverse-error trajectories demonstrate the predictive value of the spectral certificate for persistent convergence. The overlap experiments reveal substantial matrix-dependent differences in convergence improvement, and the controlled ordering experiments show that permutations of a fixed block collection can produce markedly different asymptotic factors.

Together, these findings demonstrate the practical relevance of the exact return certificate for analyzing convergence and comparing sweep designs. They also identify the geometric dependence underlying these effects: neither the magnitude of overlap nor the choice of ordering alone determines the convergence factor independently of the matrix and its associated transfer operators.

***Reproducibility.***

The supplementary experiment archive provides the merged raw CSV files and 96 inverse-error trajectory files underlying the reported numerical results. The implementation is publicly available at the project repository. Together, these resources support independent inspection of the numerical evidence and verification of the reported certificate comparisons, convergence trajectories, and ordering statistics.

## 9.5 Discussion

The return-operator representation connects exact convergence analysis with the design of overlapping IBMI sweeps. Theorem 4.2 expresses the asymptotic spectral factor entirely in terms of the original SPD matrix $A$ and the ordered block cover. The resulting certificate can be evaluated before performing an IBMI sweep, using local principal-block solves and a reduced spectral computation, without explicitly forming $A^{-1}$. This separates the assessment of asymptotic convergence from the execution of the inverse iteration itself. For a fixed block collection, the local transfer data can also

be reused across candidate orderings, making the return representation suitable for repeated sweep comparisons.

A central distinction is that the exact spectral factor and the available contraction bounds describe different levels of geometric information. The return certificate captures the complete ordered interaction of the complementary subspaces, whereas the canonical-correlation envelope and its $A$-only Schur-complement bounds control this interaction through pairwise quantities. Separated adjacent complements yield explicit strict contraction factors, and uniformly conditioned matrix families admit dimension-independent estimates under the stated separation assumptions. When adjacent complement spaces intersect, the largest canonical correlation equals one, and the pairwise envelope may lose information about contraction in directions outside their intersection. This explains why a valid a priori bound can be substantially less informative than the exact return spectrum.

The overlap results identify a corresponding limitation of block-size-based reasoning. For two-block sweeps, nested enlargement improves or preserves the exact spectral factor. For three blocks, however, an explicit SPD construction shows that enlarging one update block can increase the asymptotic factor. The monotonicity of the canonical-correlation envelope therefore cannot be transferred to the exact factor for general multi-block sweeps. The decisive quantity is the ordered composition of the local transfer operators, rather than overlap cardinality alone. This distinction provides a mathematical explanation for the strongly matrix-dependent overlap effects observed in the public SPD experiments.

Ordering introduces a separate but related geometric design problem. The public SPD instances exhibit limited sensitivity to block permutations, while the controlled inverse-coupling ensembles produce substantial differences in the exact spectral factor with the matrix and block collection held fixed. These findings show that ordering effects can be significant without being uniformly pronounced across problem families. The return certificate makes this dependence explicit and provides an exact criterion for comparing candidate orders.

The singleton-complement theory further characterizes the combinatorial structure behind ordering. The spectral factor reduces to a weighted Hamiltonian-cycle product, whose logarithmic form yields a symmetric Max-TSP objective. The SPD realization theorem shows that positive definiteness alone does not restrict the combinatorial expressivity of this objective, and the associated optimization problem is NP-hard. Consequently, efficient ordering guarantees require additional structure. Metric logarithmic couplings permit the transfer of appropriate Max-TSP approximation guarantees to the contraction exponent, while exponential-line couplings admit an explicit globally optimal ordering. The exact 2-opt exchange formula supports inexpensive local improvements, without guaranteeing global optimality.

The perturbation theory complements these structural results by quantifying the response of the spectral certificate to changes in $A$. The return spectrum admits a worst-case Hölder perturbation bound, while a simple, uniquely dominant nonzero return eigenvalue yields an explicit first-order sensitivity formula. These statements concern the asymptotic spectral factor. Finite-sweep behavior also depends on powers

of the generally nonnormal return operator, so asymptotic ordering comparisons need not predict every transient trajectory.

From a computational perspective, the principal opportunity is certificate-guided comparison of alternative sweep designs. The numerical experiments demonstrate agreement between independently constructed spectral factors, correspondence with observed persistent inverse-error behavior, and measurable savings from caching local factorizations during repeated ordering evaluations. The certificate is nevertheless not a cost-free convergence predictor: its construction requires local solves and a reduced eigenvalue computation. The reported timing improvements concern certificate scoring rather than the total cost of computing an approximate inverse. End-to-end benefits depend on the balance between preprocessing, ordering search, and the reduction in subsequent iteration cost.

Taken together, the results establish the return certificate as a framework for distinguishing exact asymptotic convergence, computable geometric bounds, and the effects of overlap and ordering. Its main contribution is to make these properties mathematically explicit and accessible from the original SPD matrix, while identifying the structural assumptions under which stronger convergence or ordering guarantees are available.

# 10 Conclusion

We developed an exact spectral framework for analyzing the convergence of overlapping iterative block matrix inversion. The inverse-error dynamics admit a structured congruence representation, and their asymptotic convergence factor is determined by a reduced return operator. An explicit $A$-only construction evaluates this factor as $\rho(M_A)^2$ using local principal-block solves, without forming the full inverse or executing an IBMI sweep.

The return-operator framework also yields finite-termination criteria, perturbation bounds, and a first-order sensitivity formula under a simple dominant-eigenvalue condition. For overlapping sweeps, it establishes exact two-block monotonicity, identifies the failure of general monotonicity in three-block systems, and provides canonical-correlation contraction bounds with an $A$-only specialization for separated complements. These results distinguish exact convergence behavior from the geometric bounds available before iteration.

The framework further exposes the combinatorial structure of sweep ordering. For singleton complements, the exact spectral factor reduces to a weighted Hamiltonian-cycle objective, leading to explicit local exchange formulas, NP-hardness for unrestricted SPD instances, and a globally optimal ordering for exponential-line coupling structures. For general fixed block collections, reusable local transfer data enable repeated exact evaluation of candidate orders.

The numerical experiments support the spectral identities and demonstrate their relevance to observed inverse-error convergence. They also show that the effects of overlap and ordering depend strongly on the underlying matrix geometry. Together, the theoretical and numerical results establish the return operator as a basis for exact convergence certification and structure-aware sweep design in overlapping IBMI.

## Statements and Declarations

**Funding.** This research received no external funding.

**Competing interests.** The author declares that there are no competing interests.

**Authors' contributions.** The sole author is responsible for the research, mathematical analysis, numerical investigation, and preparation of the manuscript.

**Data and code availability.** The input matrix identities, archived numerical results, and reproducibility materials are described in Section 9 and in the archived experiment package. Implementation materials are available at https://github.com/aisia-research-lab/ibmi-exact-convergence. The archived CSV records document the reported experiments; a complete independent source-to-output rerun of every historical diagnostic is not claimed.

**Use of generative AI-assisted tools.** Generative AI-assisted tools were used to assist with language editing, code refactoring, and code review. The author independently reviewed and verified the resulting materials and takes full responsibility for the mathematical statements, proofs, numerical results, and final manuscript.

## References

[1] Paterson, A., Pestana, J., Dolean Maini, V.: An iterative block matrix inversion (IBMI) algorithm for symmetric positive definite matrices with applications to covariance matrices. SIAM Journal on Matrix Analysis and Applications **47**(3), 1605–1627 (2026) https://doi.org/10.1137/25M1732076

[2] Benzi, M., Frommer, A., Nabben, R., Szyld, D.B.: Algebraic theory of multiplicative schwarz methods. Numerische Mathematik **89**(4), 605–639 (2001) https://doi.org/10.1007/s002110100275

[3] Xu, J.: Iterative methods by space decomposition and subspace correction. SIAM Review **34**(4), 581–613 (1992) https://doi.org/10.1137/1034116

[4] Xu, J., Zikatanov, L.: The method of alternating projections and the method of subspace corrections in Hilbert space. Journal of the American Mathematical Society **15**(3), 573–597 (2002) https://doi.org/10.1090/S0894-0347-02-00398-3

[5] Kowalik, Ł., Mucha, M.: Deterministic 7/8-approximation for the metric maximum tsp. Theoretical Computer Science **410**(47–49), 5000–5009 (2009) https://doi.org/10.1016/j.tcs.2009.07.051

[6] Supnick, F.: Extreme hamiltonian lines. Annals of Mathematics **66**(1), 179–201 (1957) https://doi.org/10.2307/1970124

[7] Björck, Å., Golub, G.H.: Numerical methods for computing angles between linear subspaces. Mathematics of Computation **27**(123), 579–594 (1973) https://doi.org/10.1090/S0025-5718-1973-0348991-3

[8] Griebel, M., Oswald, P.: Greedy and randomized versions of the multiplicative schwarz method. Linear Algebra and its Applications **437**(7), 1596–1610 (2012) https://doi.org/10.1016/j.laa.2012.04.052

[9] Hu, X., Xu, J., Zikatanov, L.T.: Randomized and fault-tolerant method of subspace corrections. Research in the Mathematical Sciences **6**(3), 29 (2019) https://doi.org/10.1007/s40687-019-0187-z

[10] Oswald, P., Zhou, W.: Random reordering in SOR-type methods. Numerische Mathematik **135**(4), 1207–1220 (2017) https://doi.org/10.1007/s00211-016-0829-7

[11] Oswald, P., Zhou, W.: Correction: Random reordering in SOR-type methods. Numerische Mathematik **154**(3–4), 521–525 (2023) https://doi.org/10.1007/s00211-023-01365-9

[12] Garey, M.R., Johnson, D.S., Tarjan, R.E.: The planar hamiltonian circuit problem is NP-complete. SIAM Journal on Computing **5**(4), 704–714 (1976) https://doi.org/10.1137/0205049

[13] Gander, M.J.: Optimized schwarz methods. SIAM Journal on Numerical Analysis **44**(2), 699–731 (2006) https://doi.org/10.1137/S0036142903425409

[14] Gander, M.J., Loisel, S., Szyld, D.B.: An optimal block iterative method and preconditioner for banded matrices with applications to PDEs on irregular domains. SIAM Journal on Matrix Analysis and Applications **33**(2), 653–680 (2012) https://doi.org/10.1137/100796194

[15] Lin, L., Yang, C., Lu, J., Ying, L., E, W.: A fast parallel algorithm for selected inversion of structured sparse matrices with application to 2d electronic structure calculations. SIAM Journal on Scientific Computing **33**(3), 1329–1351 (2011) https://doi.org/10.1137/09077432X

[16] Gower, R.M., Richtárik, P.: Randomized quasi-newton updates are linearly convergent matrix inversion algorithms. SIAM Journal on Matrix Analysis and Applications **38**(4), 1380–1409 (2017) https://doi.org/10.1137/16M1062053

[17] Schäfer, F., Sullivan, T.J., Owhadi, H.: Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity. Multiscale Modeling & Simulation **19**(2), 688–730 (2021) https://doi.org/10.1137/19M129526X

[18] Schäfer, F., Katzfuss, M., Owhadi, H.: Sparse cholesky factorization by kullback–leibler minimization. SIAM Journal on Scientific Computing **43**(3), 2019–2046 (2021) https://doi.org/10.1137/20M1336254

[19] Frangella, Z., Tropp, J.A., Udell, M.: Randomized nyström preconditioning. SIAM

Journal on Matrix Analysis and Applications **44**(2), 718–752 (2023) https://doi.org/10.1137/21M1466244

[20] Huan, S., Guinness, J., Katzfuss, M., Owhadi, H., Schäfer, F.: Sparse inverse cholesky factorization of dense kernel matrices by greedy conditional selection. SIAM/ASA Journal on Uncertainty Quantification **13**(3), 1649–1679 (2025) https://doi.org/10.1137/23M1606253

[21] Elsner, L.: An optimal bound for the spectral variation of two matrices. Linear Algebra and its Applications **71**, 77–80 (1985) https://doi.org/10.1016/0024-3795(85)90236-8

[22] Stewart, G.W., Sun, J.-g.: Matrix Perturbation Theory. Academic Press, Boston (1990)

[23] Wu, J., Yi, W.: Matrix spectral norm wielandt inequalities with statistical applications. Journal of Inequalities and Applications **2014**, 110 (2014) https://doi.org/10.1186/1029-242X-2014-110

[24] Choi, M.J., Tan, V.Y.F., Anandkumar, A., Willsky, A.S.: Learning latent tree graphical models. Journal of Machine Learning Research **12**, 1771–1812 (2011)

[25] Davis, T.A., Hu, Y.: The university of florida sparse matrix collection. ACM Transactions on Mathematical Software **38**(1), 1–25 (2011) https://doi.org/10.1145/2049662.2049663

[26] Dempster, A.P.: Covariance selection. Biometrics **28**(1), 157–175 (1972) https://doi.org/10.2307/2528966

# Appendix A Inverse-coupling graph consequences for finite termination

The singleton cycle formula and the path–cycle dichotomy in Section 8 connect finite termination to the sparsity pattern of $H = A^{-1}$. Classical relationships between matrix sparsity and inverse structure are discussed, for example, in [26]. For singleton-complement IBMI, the decisive feature is whether the ordered sequence of nonzero inverse couplings forms a closed cycle.

## A.1 An arbitrary-dimensional exact family

The path–cycle criterion yields an explicit family of genuine two-sweep termination examples in every dimension $n \geq 3$.

**Corollary A.1** (An arbitrary-dimensional two-sweep family)**.** *Let $n \geq 3$, and let $H \in \mathbb{R}^{n\times n}$ be an SPD tridiagonal matrix with nonzero first sub- and super-diagonal*

*entries:*

$$H = \begin{bmatrix} d_1 & a_1 & & & \\ a_1 & d_2 & a_2 & & \\ & a_2 & \ddots & \ddots & \\ & & \ddots & d_{n-1} & a_{n-1} \\ & & & a_{n-1} & d_n \end{bmatrix} \succ 0, \qquad a_r \neq 0. \tag{A1}$$

*Set $A = H^{-1}$ and consider the singleton-complement update blocks*

$$I_r = [n] \setminus \{r\}, \qquad r = 1, \ldots, n,$$

*in the order $1, 2, \ldots, n$. Then the complete-sweep propagator satisfies*

$$Q \neq 0, \qquad Q^2 = 0. \tag{A2}$$

*Consequently, the inverse-error operator has nilpotency index exactly two: every symmetric initial inverse error is annihilated after at most two complete sweeps, but not every initial error is annihilated after one sweep.*

*Proof* For singleton complements, the complete-sweep propagator admits the rank-one return representation

$$Q = uv^\top, \qquad Q^2 = (v^\top u)Q,$$

where the scalar return coefficient is

$$v^\top u = \prod_{r=1}^{n} \frac{h_{r+1,r}}{h_{rr}}, \qquad r+1 \text{ interpreted modulo } n.$$

The tridiagonal structure gives $h_{r+1,r} = a_r \neq 0$ for $r = 1, \ldots, n-1$, while $h_{1n} = 0$ because $n \geq 3$. Therefore, $v^\top u = 0$ and $Q^2 = 0$. The rank-one factorization of $Q$ contains the product of the $n-1$ nonzero consecutive couplings $a_1, \ldots, a_{n-1}$, together with invertible diagonal scaling factors. Hence, $Q \neq 0$.

Finally, for the inverse-error operator $\mathcal{L}(E) = QEQ^\top$,

$$\mathcal{L}(I_n) = QQ^\top \neq 0, \qquad \mathcal{L}^2(E) = Q^2 E (Q^\top)^2 = 0.$$

Thus, $\mathcal{L}$ has nilpotency index exactly two. □

The mechanism in Corollary A.1 is determined by the ordered inverse-coupling pattern, not by chordality alone. The path graph of a tridiagonal matrix is chordal, but a complete graph is also chordal and has no missing endpoint coupling. What matters here is a sequence of nonzero consecutive couplings whose closing edge vanishes.

More generally, the cycle formula (60) separates the roles of sparsity and coupling strength. A missing edge in the ordered return cycle produces a zero spectral factor and finite termination. When every cycle edge is nonzero, the magnitudes of the normalized couplings determine the asymptotic factor.

For singleton complements, the reduced return matrix is scalar. If finite termination occurs, its return coefficient must vanish, and the rank-one sweep representation then gives $Q^2 = 0$. Thus, every finitely terminating singleton-complement sweep terminates in at most two complete sweeps. Higher nilpotency indices, if realizable within the IBMI framework, require a return space of dimension greater than one.

# Appendix B Additional numerical diagnostics

This appendix provides additional ordering and finite-termination diagnostics supporting Section 9. The reported results are drawn from archived experiment outputs. The targeted ordering examples and finite-termination diagnostics use different experimental configurations and are interpreted separately.

## B.1 Targeted ordering

We consider a synthetic SPD inverse-coupling instance with $n = 8$ (seed 111). Exhaustive enumeration gives a native-to-optimal spectral-factor ratio of $1.41905 \times 10^5$. The archived 2-opt procedure produces a factor 2.32277 times the global optimum.

Table B1 reports the corresponding spectral certificates. The example establishes that exact evaluation of local 2-opt exchanges does not guarantee a globally optimal ordering. The ratios compare asymptotic spectral factors rather than iteration counts or execution times.

**Table B1** Targeted ordering diagnostic for an SPD inverse-coupling instance ($n = 8$, seed 111). All reported factors are spectral certificates.

| Quantity | Value |
|---|---|
| $\lambda_{\min}(H)$ | 0.868124 |
| Native factor | $1.00924 \times 10^{-21}$ |
| Exact optimal factor | $7.11209 \times 10^{-27}$ |
| 2-opt factor | $1.65198 \times 10^{-26}$ |
| Native/optimal | $1.41905 \times 10^5$ |
| 2-opt/optimal | 2.32277 |

## B.2 Finite-precision scope

In experiments, we use the dimensions $n = 8, 16, 32$. For the larger instances, the recorded inverse errors after the first sweep are already close to floating-point roundoff. Consequently, these trajectories do not reliably distinguish the predicted two-sweep termination from the effects of finite precision.

Exact finite termination follows from the algebraic return-operator identities and does not depend on numerical evidence. The experiments illustrate the corresponding behavior in floating-point arithmetic. Additional historical examples, including the tree log-metric instance, are documented separately because their independent reproduction was not established in the consolidated experimental study.