# Gaussian Elimination in Komorebi

## Collected source pages

Block reduction, exact reconstruction, nonlinear elimination, Lean formalization, and exact-arithmetic verification.

### What this collection contains

Original PDF pages from four Komorebi sources, followed by two arithmetic notes reflowed from their Markdown originals. Equations, figures, theorem numbering, and printed page numbers on the source pages are retained.

The source map on the next page identifies both the page positions in this collection and the page positions in each original PDF. Use the PDF bookmarks to move between subjects.

### Reading the inverse formulas

Eliminating a block requires that block to be invertible. A projected-inverse formula additionally requires the full matrix, equivalently its Schur complement under that hypothesis, to be invertible. The Paper VI excerpt states these assumptions explicitly.

Verification reports and numerical results are reproduced from the source documents. This compilation did not rerun the source experiments or Lean builds.

Komorebi P. | Compiled 7 October 2026 | 45 pages

# Source map and reading order

Collection pages below are the sequential PDF page positions. The original printed numbers remain visible on each extracted page.

| Pages | Source and coverage | Original PDF |
|---|---|---|
| 3-4 | **Projector-Resolved Quantum Spacetime**<br>Schur-Feshbach reduction, the projected inverse, and nested elimination. | 9-10<br>Printed 9-10 |
| 5-32 | **From Hidden Packets to Exact Memory**<br>Scalar elimination, continued fractions, reconstruction, boundary rank, memory, polyphase refinement, numerical examples, and the formula sheet. | 4-31<br>Printed 4-31 |
| 33-36 | **Paper VI - Lean Guide, final edition**<br>Section 7: SchurFeshbach, nonnormal decay, Laplace inversion, and rectangular memory. | 25-28<br>Printed 24-27 |
| 37-42 | **Komorebi Math User Guide - Technical Manual, Block 36**<br>Sections 41.4-41.7: derivatives, symbolic and numerical examples, API contracts, result records, and failure conditions. | 731-736<br>Printed 661-666 |
| 43-45 | **Exact-arithmetic source notes**<br>Rational Gaussian-elimination audit; independent fraction-free Bareiss and modular rank checks. | Markdown sources<br>Reflowed here |

## Source boundaries

The worked example includes adjacent discussion and its existing verification ledger so that the elimination calculation retains its mathematical context. Original page boundaries can include the end of a preceding section or the start of a following one.

The nonlinear material is from version 1.0.0b4 of the technical manual. The arithmetic appendix preserves the source-specific finite scope and recorded failures. Historical counts are not presented as current project totals.

# 4 Exact elimination and effective dynamics

A coarse model that simply drops fine packets is generally wrong. The fine packets return through an energy-dependent self-energy. This is the finite-dimensional mechanism by which exact coarse graining differs from truncation.

## 4.1 Schur–Feshbach reduction

Let $P$ be an orthogonal projector and $Q = I - P$. Relative to $\mathcal{H} = P\mathcal{H} \oplus Q\mathcal{H}$, write

$$H = \begin{pmatrix} H_{PP} & H_{PQ} \\ H_{QP} & H_{QQ} \end{pmatrix}.$$

**Theorem 4.1** (Exact projected resolvent)**.** *Let $z - H_{QQ}$ be invertible. Define*

$$H_{\text{eff}}(z) = H_{PP} + H_{PQ}(z - H_{QQ})^{-1} H_{QP}. \tag{18}$$

*Whenever the displayed inverses exist,*

$$\boxed{P(z - H)^{-1} P = (z - H_{\text{eff}}(z))^{-1}.} \tag{19}$$

*Proof.* Solve $(z - H)(u, v) = (f, 0)$ in block form:

$$\begin{aligned} (z - H_{PP})u - H_{PQ}v &= f, \\ -H_{QP}u + (z - H_{QQ})v &= 0. \end{aligned}$$

The second equation gives

$$v = (z - H_{QQ})^{-1} H_{QP} u.$$

Substitution into the first gives

$$\Big(z - H_{PP} - H_{PQ}(z - H_{QQ})^{-1} H_{QP}\Big)\, u = f.$$

This is exactly (19). □

The self-energy

$$\Sigma_P(z) = H_{PQ}(z - H_{QQ})^{-1} H_{QP} \tag{20}$$

contains the complete influence of the eliminated packet. Bare truncation replaces $H_{\text{eff}}$ by $H_{PP}$ and therefore discards $\Sigma_P$.

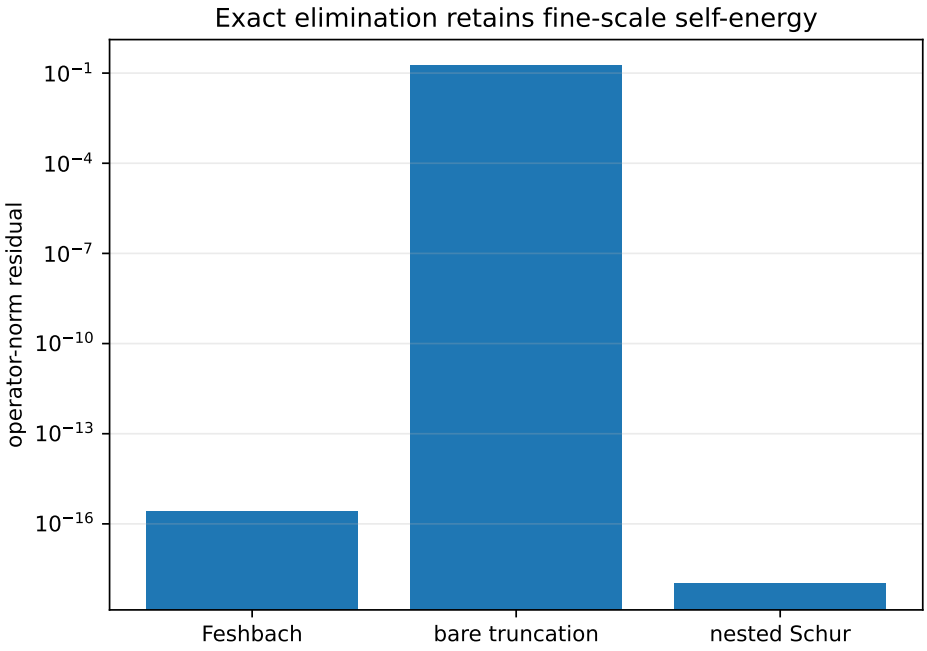



Figure 5: The projected-resolvent identity and nested Schur elimination agree to machine precision. Bare truncation, which omits the self-energy, has a residual of approximately $1.79 \times 10^{-1}$ in the same model.

### 4.2 Elimination order

**Theorem 4.2** (Associativity of nested Schur elimination)**.** *Suppose a matrix is partitioned into retained, middle, and outer coordinates, and every required inverse exists. Eliminating the outer coordinates and then the middle coordinates gives the same Schur complement as eliminating the union of middle and outer coordinates in one step.*

*Proof.* Both procedures solve the same block linear system for the retained coordinates. Equivalently, block Gaussian elimination factors the matrix into triangular factors and the retained Schur complement. Grouping two successive elimination steps changes the factorization order but not the unique reduced equation on the retained variables. □

This property is important for multiscale refinement. An exact hierarchy should not depend on whether unresolved scales are removed individually or in batches.

### 4.3 A fixed-point target

For nested resolutions $P_0 \le P_1 \le \cdots$, define the exact effective operators

$$H_\ell^{\mathrm{eff}}(z) = P_\ell H P_\ell + P_\ell H Q_\ell (z - Q_\ell H Q_\ell)^{-1} Q_\ell H P_\ell. \tag{21}$$

The finite identity is exact at each $\ell$. The continuum question is whether these operators approach a scale-independent limit in a topology strong enough to determine dynamics.

**Conjecture 4.3** (Uniform Feshbach continuum control)**.** There exists a physically meaningful nested projector family for which, on a common resolvent set,

$$H_\ell^{\mathrm{eff}}(z) \longrightarrow H_*(z) \tag{22}$$

locally uniformly in $z$, with bounds independent of volume and ultraviolet cutoff.

The conjecture is not implied by finite exactness. It requires uniform control of the eliminated resolvents and couplings.

17 Exact Schwarzschild exterior self-energy. Only the two retained boundary coordinates receive nonzero exterior correction. The third and later singular values vanish at numerical precision. . . . . 26

18 Response and reconstruction certificate. The finite-domain response peaks near the middle of the sampled band. Exact Feshbach reconstruction remains at floating-point accuracy through the peak. 27

19 The final reconstruction cycle. The latent packet is compressed into boundary transfer and memory, not discarded. Exact synthesis recovers the full coordinate state. . . . . 28

# 1 Why this area benefits from a complete example

The hypergeometric example showed how one factorial expression can be repeatedly rewritten until its arithmetic cancellation becomes visible. The present example treats a different kind of hidden structure. Here the difficulty is not a denominator. It is the influence of variables that have been removed from the displayed model.

A naive coarse model keeps a subset of coordinates and drops the rest. That operation is usually wrong. The latent coordinates continue to affect the retained variables. The Schur–Feshbach construction records that effect exactly. In the frequency domain it appears as a self-energy. In the time domain it appears as memory.

This is one of the most important parts of the calculus for four reasons.

1. It distinguishes exact coarse graining from bare truncation.
2. It gives a precise object for the effect of unresolved scales.
3. It connects operator algebra, signal theory, control, quantum mechanics, numerical analysis, and wave propagation.
4. It supplies a natural place for polyphase refinement because boundary channels can be assigned to exact phase packets.

# 2 The one chain used throughout

## 2.1 The operator

Consider the real symmetric Jacobi matrix

$$H = \begin{pmatrix} 8 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 7 & 2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 6 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 4 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 3 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 5 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 7 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 9 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 10 \end{pmatrix}.$$

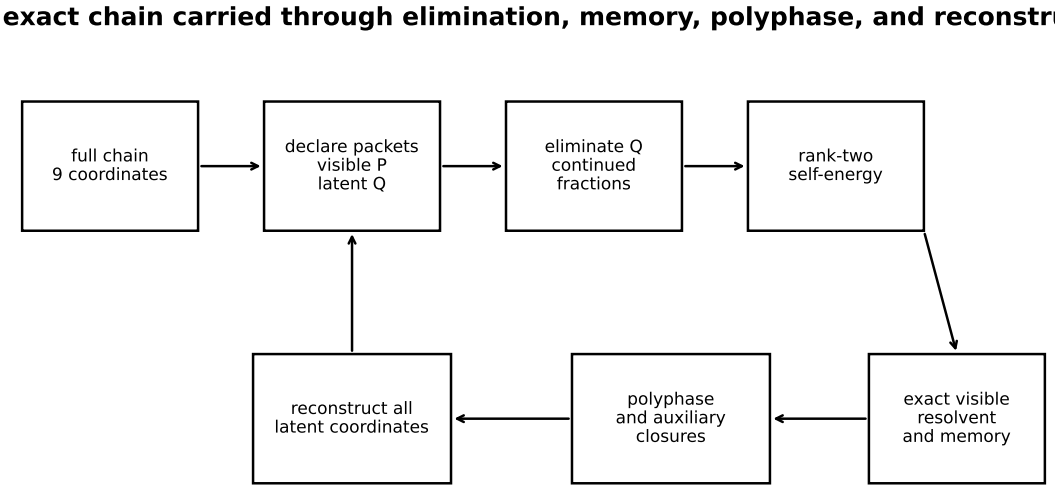



**Figure 1:** The complete calculation. The same finite object is carried through packet declaration, nested elimination, exact resolvent reconstruction, time-domain memory, auxiliary-state closure, polyphase refinement, and physical scaling.

The diagonal entries act as local on-site terms. The off-diagonal entries are nearest-neighbor transfers. Every eigenvalue is positive. Numerically,

$$\operatorname{spec}(H) \approx \{1.17559, 3.94687, 4.57492, 5.70387, 6.42207, 7.61638, 9.18889, 9.35170, 11.01970\}.$$

The positivity lets the same matrix generate stable decay through

$$\dot{x}(t) = -Hx(t).$$

## 2.2 The packet declaration

Retain the middle coordinates

$$P = \{3, 4, 5\}$$

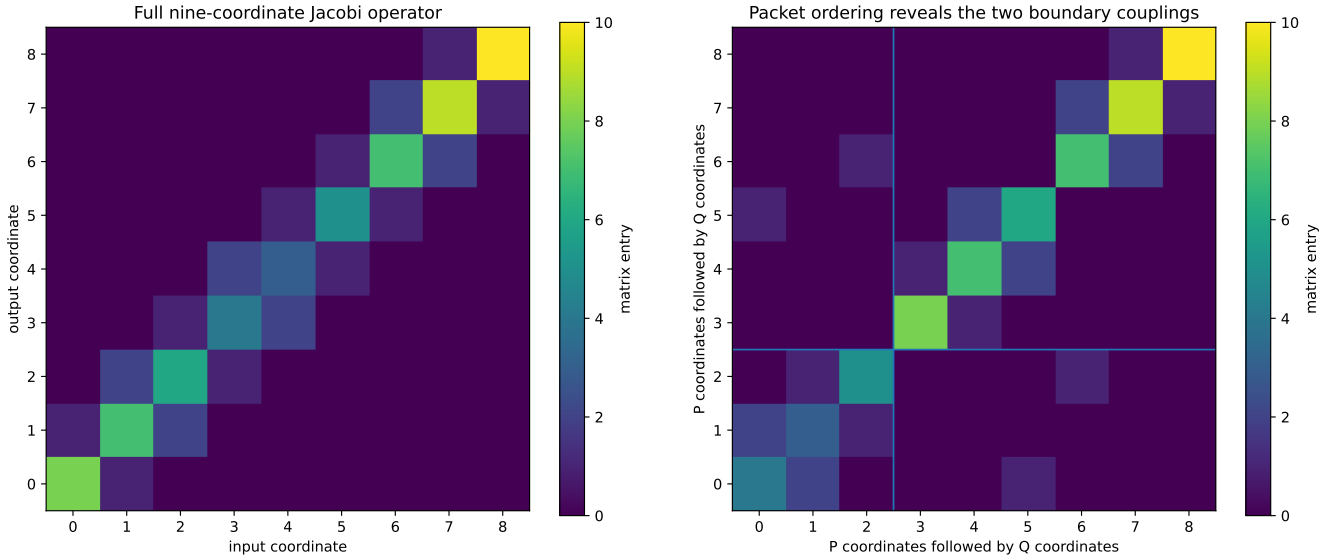
and eliminate

$$Q = \{0, 1, 2, 6, 7, 8\}.$$

The retained packet can be read as a small interaction region. The left latent packet is an incoming or near-horizon exterior. The right latent packet is an outgoing or far-field exterior. The labels are suggestive only. Every proof below is finite linear algebra.

In the order

$$(3, 4, 5 \mid 0, 1, 2, 6, 7, 8),$$




**(a)** Full coordinate matrix

**(b)** Visible and latent packet order

**Figure 2:** The operator before and after packet ordering. The nearest-neighbor matrix looks like one chain in ordinary coordinates. Moving the retained coordinates to the front exposes the visible block, the latent block, and only two cross-packet couplings.

the resolvent matrix

$$A(z) = zI - H$$

has block form

$$A(z) = \begin{pmatrix} A_{PP}(z) & A_{PQ} \\ A_{QP} & A_{QQ}(z) \end{pmatrix}.$$

The retained block is

$$A_{PP}(z) = \begin{pmatrix} z-4 & -2 & 0 \\ -2 & z-3 & -1 \\ 0 & -1 & z-5 \end{pmatrix}.$$

The latent block is a direct sum of two three-coordinate chains,

$$A_{QQ}(z) = (zI - H_L) \oplus (zI - H_R),$$

where

$$H_L = \begin{pmatrix} 8 & 1 & 0 \\ 1 & 7 & 2 \\ 0 & 2 & 6 \end{pmatrix}, \qquad H_R = \begin{pmatrix} 7 & 2 & 0 \\ 2 & 9 & 1 \\ 0 & 1 & 10 \end{pmatrix}.$$

Only the first retained coordinate touches the left exterior. Only the last retained coordinate touches the right exterior.

# 3 Eliminate the left exterior one coordinate at a time

The first purpose of the long calculation is to show exactly where the boundary transfer function comes from.

## 3.1 First elimination

Start at the far-left coordinate. Its scalar resolvent factor is

$$d_0(z) = z - 8.$$

Eliminating it changes the next diagonal from

$$z - 7$$

to

$$\boxed{d_1(z) = z - 7 - \frac{1}{z-8}.}$$

Nothing has been approximated. The removed coordinate returns through the rational correction

$$\frac{1}{z-8}.$$

## 3.2 Second elimination

The middle left coordinate couples to the boundary coordinate with strength two. Eliminating it changes the boundary diagonal from

$$z - 6$$

to

$$\boxed{d_2(z) = z - 6 - \frac{4}{z - 7 - \frac{1}{z-8}}.}$$

The surface Green function seen from the retained packet is the reciprocal

$$\boxed{g_L(z) = \frac{1}{z - 6 - \dfrac{4}{z - 7 - \dfrac{1}{z-8}}}.}$$

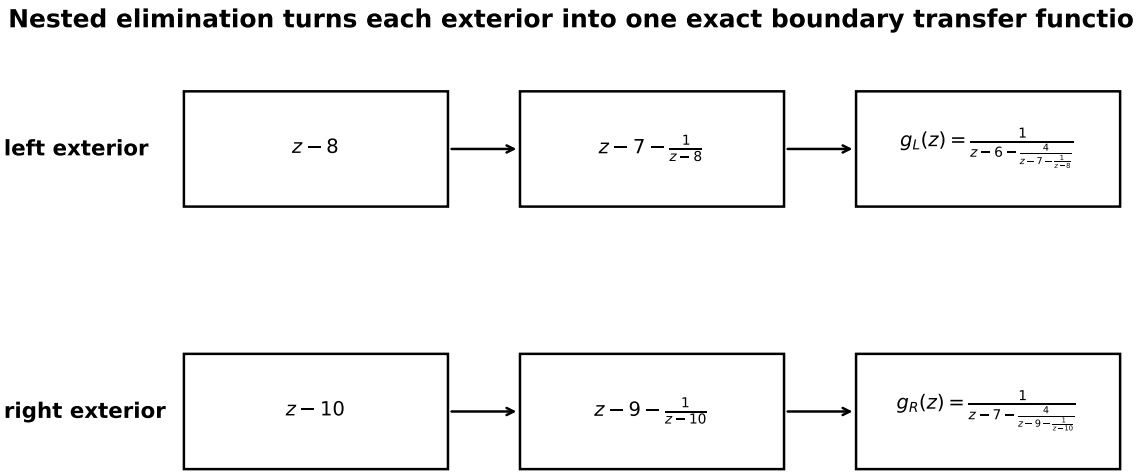



**Figure 3:** Nested elimination of both exteriors. Each removed coordinate modifies the next boundary equation. The final nested fraction is the exact transfer function of the complete exterior.

### 3.3 Determinant form

Define the determinant recursion

$$D_0 = 1, \qquad D_1 = z - 8,$$

$$D_2 = (z-7)D_1 - D_0,$$

$$D_3 = (z-6)D_2 - 4D_1.$$

Direct expansion gives

$$D_2 = z^2 - 15z + 55$$

and

$$D_3 = z^3 - 21z^2 + 141z - 298.$$

Therefore

$$\boxed{g_L(z) = \frac{D_2(z)}{D_3(z)} = \frac{z^2 - 15z + 55}{z^3 - 21z^2 + 141z - 298}.}$$

The continued fraction and determinant quotient are the same object in two coordinate systems.

## 4 Eliminate the right exterior one coordinate at a time

The right calculation is similar but not identical. Starting at the far-right coordinate gives

$$e_0(z) = z - 10.$$

After removing that coordinate,

$$e_1(z) = z - 9 - \frac{1}{z-10}.$$

The boundary coordinate couples with strength two, so

$$e_2(z) = z - 7 - \frac{4}{z - 9 - \frac{1}{z-10}}.$$

Hence

$$\boxed{g_R(z) = \frac{1}{z - 7 - \dfrac{4}{z - 9 - \dfrac{1}{z-10}}}.}$$

The determinant polynomials are

$$E_2 = z^2 - 19z + 89$$

and

$$E_3 = (z-11)(z^2 - 15z + 53).$$

Thus

$$\boxed{g_R(z) = \frac{z^2 - 19z + 89}{(z-11)(z^2 - 15z + 53)}.}$$

## 5 Assemble the exact self-energy

**Definition 5.1** (Boundary embedding)**.** Let

$$U = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

The columns of $U$ select the left and right retained boundaries.

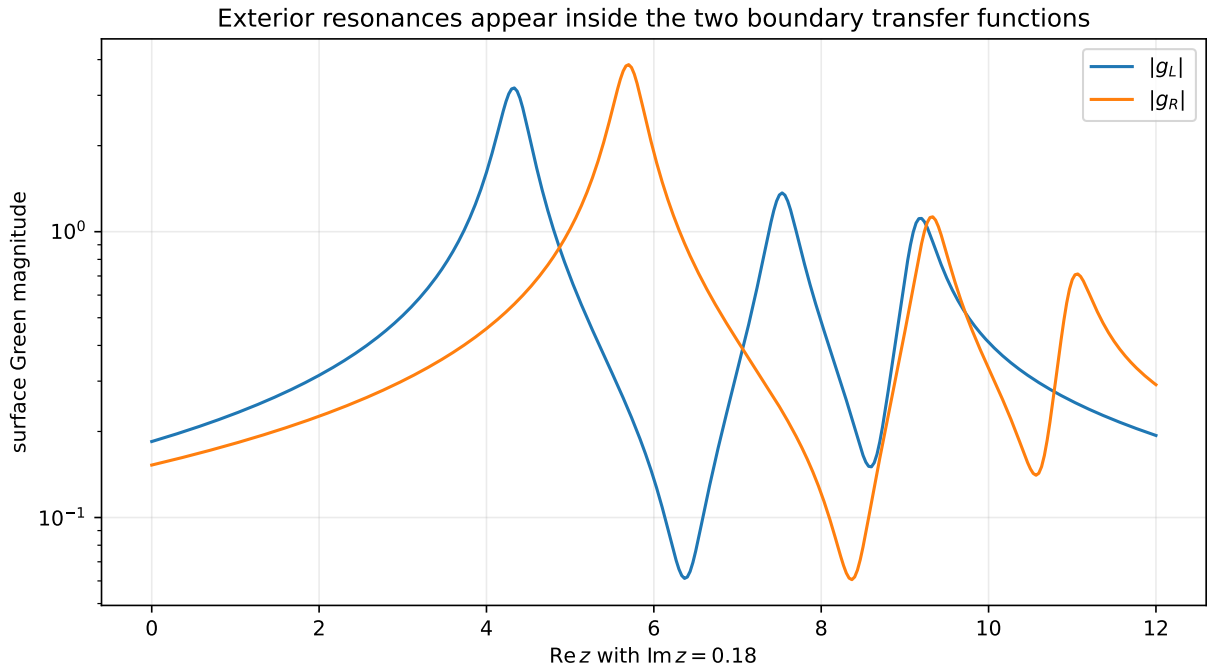



**Figure 4:** The two surface Green functions. The peaks are latent exterior resonances. The retained packet does not lose those spectral features when the exterior is eliminated. They return through the boundary transfer functions.

The exterior Green matrix at the two interfaces is

$$G_\partial(z) = \begin{pmatrix} g_L(z) & 0 \\ 0 & g_R(z) \end{pmatrix}.$$

The exact self-energy is

$$\boxed{\Sigma(z) = UG_\partial(z)U^{\mathsf{T}} = \begin{pmatrix} g_L(z) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & g_R(z) \end{pmatrix}.}$$

The exact Schur operator is

$$\boxed{S(z) = A_{PP}(z) - \Sigma(z).}$$

**Theorem 5.2** (Exact boundary-rank theorem)**.** *For the nine-coordinate chain,*

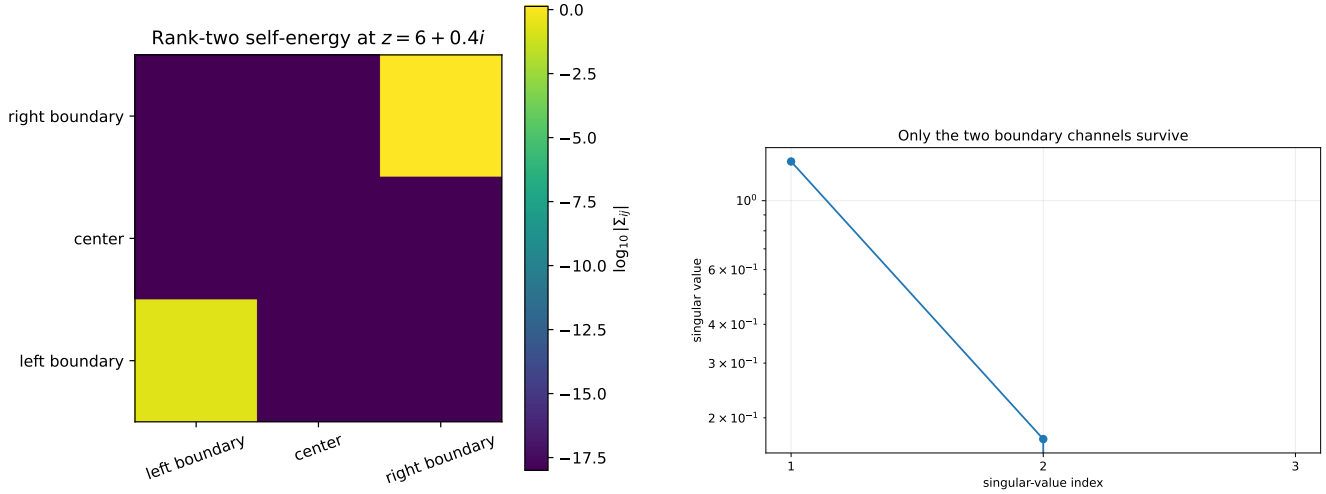
$$\operatorname{rank} \Sigma(z) \leq 2$$

*whenever the latent resolvent exists. The self-energy is supported only at the first and last retained coordinates.*

*Proof.* The cross-block matrix $A_{QP}$ has nonzero columns only at the two retained boundary coordinates. Therefore rank $A_{QP} \leq 2$. Since

$$\Sigma = A_{PQ} A_{QQ}^{-1} A_{QP},$$

its rank cannot exceed the rank of $A_{QP}$. The support statement follows from the same two columns and the corresponding two rows of $A_{PQ}$. □




**(a)** Self-energy support at $z = 6 + 0.4i$ **(b)** Singular values

**Figure 5:** The exterior returns through two channels. The center retained coordinate receives no direct self-energy term. Two singular values survive and the third is zero to numerical precision.

At

$$z = 6 + 0.4i,$$

the two nonzero entries are approximately

$$g_L(z) = 0.11753807 - 0.12398961i$$

and

$$g_R(z) = 0.75014451 - 1.10713376i.$$

The reduced operator is

$$S(z) \approx \begin{pmatrix} 1.88246193 + 0.52398961i & -2 & 0 \\ -2 & 3 + 0.4i & -1 \\ 0 & -1 & 0.24985549 + 1.50713376i \end{pmatrix}.$$

## 6 Exact resolvent reconstruction

Let

$$A(z)u = b.$$

Split

$$u = \begin{pmatrix} u_P \\ u_Q \end{pmatrix}, \qquad b = \begin{pmatrix} b_P \\ b_Q \end{pmatrix}.$$

The latent equation gives

$$u_Q = A_{QQ}^{-1}(b_Q - A_{QP}u_P).$$

Substitution into the visible equation yields

$$\boxed{S(z)u_P = b_P - A_{PQ}A_{QQ}^{-1}b_Q.}$$

After solving the three-coordinate visible equation, the six latent coordinates are reconstructed by back-substitution.

**Theorem 6.1** (Projected resolvent identity)**.** *Whenever the displayed inverses exist,*

$$\boxed{P(zI - H)^{-1}P = S(z)^{-1}.}$$

*Proof.* The block equations above show that the visible component of the full resolvent is obtained by applying $S(z)^{-1}$ to the effective visible source. When the source lies entirely in the retained packet, this is exactly the projected resolvent. □

### 6.1 An exact rational calculation at $z = 0$

Use the source

$$b = e_4,$$

which acts at the center retained coordinate. At $z = 0$,

$$g_L(0) = -\frac{55}{298}, \qquad g_R(0) = -\frac{89}{583}.$$

Therefore

$$S(0) = \begin{pmatrix} -\frac{1137}{298} & -2 & 0 \\ -2 & -3 & -1 \\ 0 & -1 & -\frac{2826}{583} \end{pmatrix}.$$

Solving the three-coordinate reduced system gives

$$\boxed{u_P = \frac{1}{1869341}\begin{pmatrix} 561432 \\ -1071054 \\ 220957 \end{pmatrix}.}$$

Back-substitution gives

$$\boxed{u_Q = \frac{1}{1869341}\begin{pmatrix} -3768 \\ 30144 \\ -103620 \\ -33731 \\ 7580 \\ -758 \end{pmatrix}.}$$

In original coordinate order,

$$\boxed{u = \frac{1}{1869341}\begin{pmatrix} -3768 \\ 30144 \\ -103620 \\ 561432 \\ -1071054 \\ 220957 \\ -33731 \\ 7580 \\ -758 \end{pmatrix}.}$$

Every entry agrees exactly with the direct inverse of $-H$.

### 6.2 A complex spectral calculation

At $z = 6 + 0.4i$, the full nine-coordinate solve and the reduced-plus-reconstructed solve agree with relative error

$$\boxed{1.65 \times 10^{-16}}.$$

## 7 The determinant identity

Block determinant factorization gives

$$\boxed{\det(zI - H) = \det(zI - H_Q)\det S(z).}$$

Since the latent block is the direct sum of the left and right exteriors,

$$\det(zI - H_Q) = D_3(z)E_3(z).$$

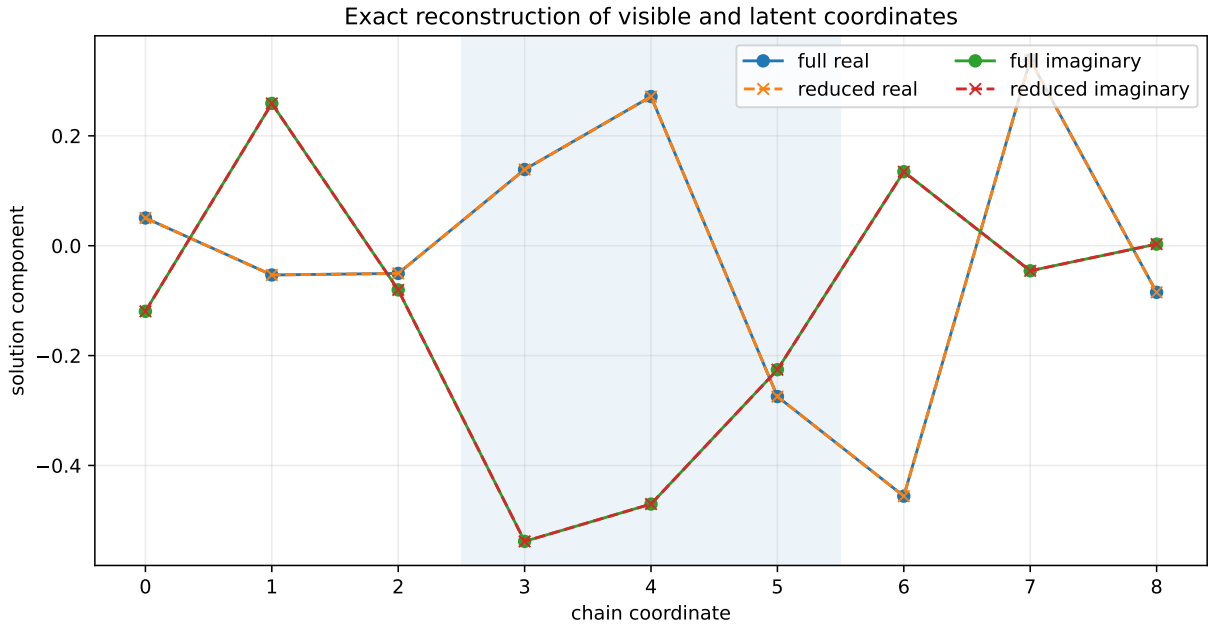



**Figure 6:** Full and reconstructed solution. The retained coordinates are solved in the three-dimensional Schur system. Every latent coordinate is then reconstructed. The two curves lie on top of one another.

The full characteristic polynomial is

$$\begin{aligned}\det(zI - H) = z^9 - 59z^8 + 1509z^7 - 21899z^6 + 198010z^5 \\ - 1151052z^4 + 4269941z^3 - 9631059z^2 \\ + 11728723z - 5608023.\end{aligned}$$

The rational poles of $S(z)$ at the latent eigenvalues are not spurious. They cancel against $\det(zI - H_Q)$ in the full determinant identity.

## 8 A second compression to two boundary coordinates

The self-energy already has the factorization

$$\Sigma = UG_\partial U^\mathsf{T}.$$

Assume $A_{PP}$ is invertible. The Woodbury identity gives

$$\boxed{S^{-1} = A_{PP}^{-1} + A_{PP}^{-1}U\left(G_\partial^{-1} - U^\mathsf{T}A_{PP}^{-1}U\right)^{-1}U^\mathsf{T}A_{PP}^{-1}.}$$

The full latent exterior has been compressed into a two-by-two boundary correction. This is a stronger statement than saying the exterior is low rank. It identifies the exact coordinates through which it acts.

*Remark* 8.1. For this example the left and right exteriors are disconnected inside $Q$, so $G_\partial$ is diagonal. If the eliminated region connects the two retained boundaries, the same formula holds with a full two-by-two boundary Green matrix.

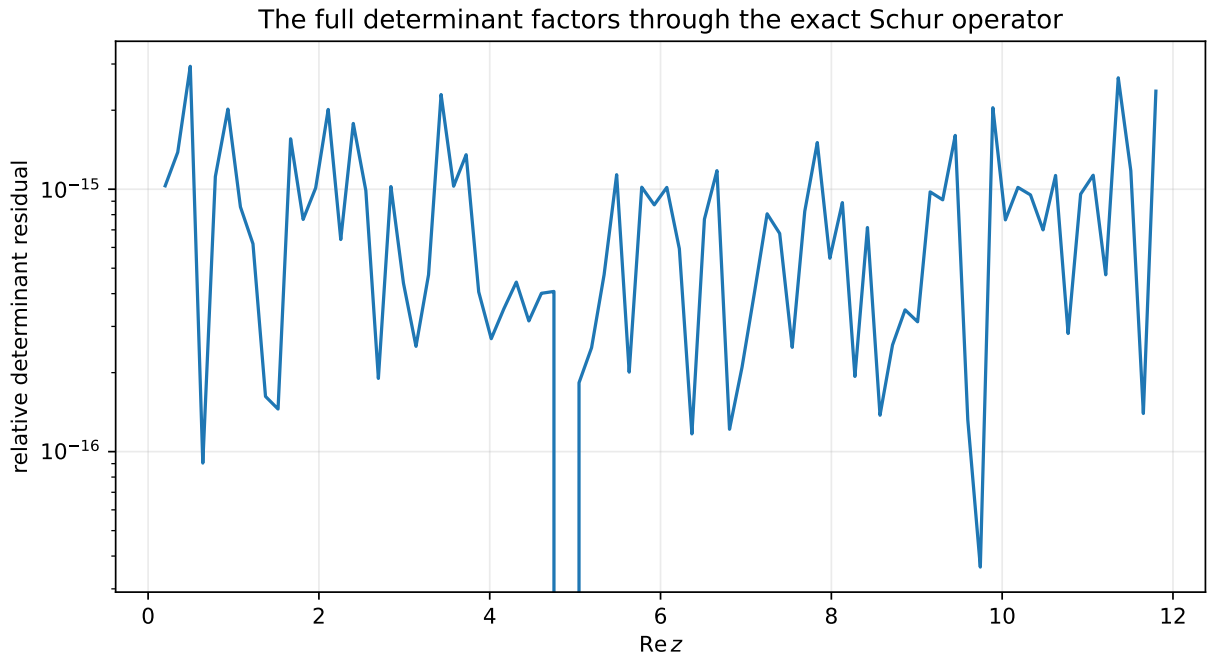



**Figure 7:** Numerical determinant certificate. Across the sampled complex line, the relative residual of the exact determinant factorization stays near floating-point roundoff.

## 9 Nested elimination is associative

The left exterior was removed in three scalar steps. It could instead be removed as one three-coordinate block. Both procedures produce the same $g_L(z)$. The same is true on the right.

More generally, suppose the latent packet is divided into an outer packet $Q_2$ and an intermediate packet $Q_1$. Eliminating $Q_2$ and then $Q_1$ gives the same retained Schur complement as eliminating $Q_1 \oplus Q_2$ in one step, provided every required inverse exists.

**Theorem 9.1** (Associativity of exact elimination)**.** *Exact Schur elimination is independent of how the eliminated coordinates are grouped.*

*Proof.* Every elimination order solves the same block linear system for the retained coordinates. Algebraically, the different routes are regroupings of block Gaussian elimination. The unique retained equation is therefore the same. □

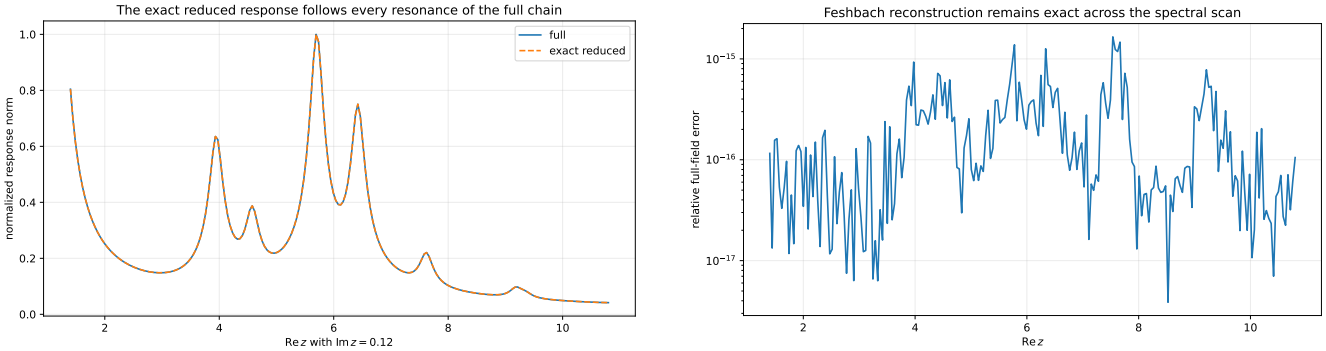
This property is essential for multiscale refinement. A coarse model should not depend on whether fine packets are removed one level at a time or in one batch.

## 10 Frequency sweep and resonant sensitivity

The source is again concentrated at the middle retained coordinate. Along

$$z = x + 0.12i,$$

the exact reduced response follows every peak of the full response.




**(a)** Full and reduced response **(b)** Reconstruction error

**Figure 8:** Exact response across the spectral band. The self-energy retains the latent poles, so the reduced model reproduces the resonant structure rather than smoothing it away.

Across 220 spectral samples, the largest full-field reconstruction error was

$$\boxed{1.66 \times 10^{-15}}.$$

Near a resonance, an approximate self-energy becomes more dangerous. Let

$$\widetilde{S} = S - (\widetilde{\Sigma} - \Sigma).$$

If

$$\left\| S^{-1}(\widetilde{\Sigma} - \Sigma) \right\| < 1,$$

then the resolvent identity gives

$$\boxed{\left\| \widetilde{S}^{-1} - S^{-1} \right\| \leq \frac{\|S^{-1}\|^2 \left\| \widetilde{\Sigma} - \Sigma \right\|}{1 - \|S^{-1}\| \left\| \widetilde{\Sigma} - \Sigma \right\|}}.$$

The same absolute self-energy error can therefore cause a much larger solution error when $\|S^{-1}\|$ is large.

## 11 A hierarchy of boundary closures

The exact left transfer function is

$$g_L(z) = \cfrac{1}{z - 6 - \cfrac{4}{z - 7 - \cfrac{1}{z - 8}}}.$$

Truncating the continued fraction after zero, one, two, or three latent coordinates gives

$$
\begin{aligned}
g_L^{(0)}(z) &= 0, \\
g_L^{(1)}(z) &= \frac{1}{z-6}, \\
g_L^{(2)}(z) &= \frac{1}{z-6-\dfrac{4}{z-7}}, \\
g_L^{(3)}(z) &= g_L(z).
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
g_R^{(0)}(z) &= 0, \\
g_R^{(1)}(z) &= \frac{1}{z-7}, \\
g_R^{(2)}(z) &= \frac{1}{z-7-\dfrac{4}{z-9}}, \\
g_R^{(3)}(z) &= g_R(z).
\end{aligned}
$$

These are not arbitrary curve fits. They are the successive boundary Krylov or Jacobi truncations generated from the interface inward.

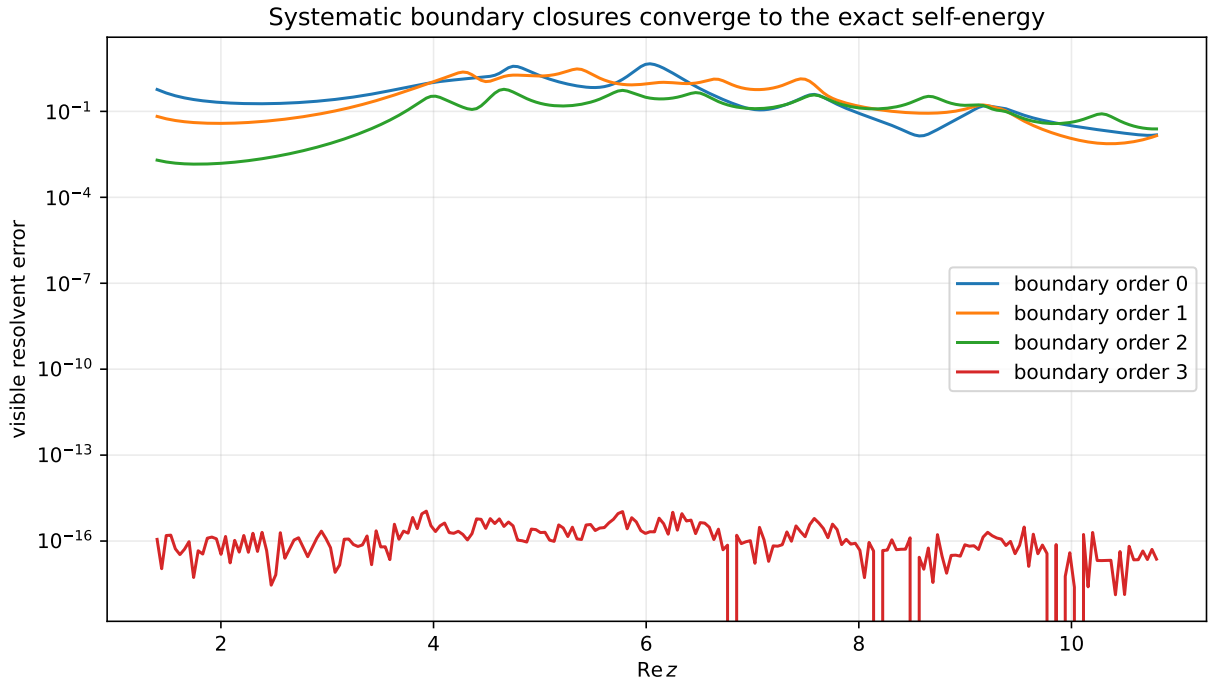



**Figure 9:** Visible resolvent error for the closure hierarchy. Order zero is bare truncation. Orders one and two retain progressively more boundary moments. Order three is exact because each exterior has three active states.

*Boundary* 11.1 (No universal low-order closure). A small self-energy rank does not imply a small dynamical realization order. Rank counts boundary channels. The number of poles counts latent memory modes. This example has two boundary channels but six generic latent decay modes.

## 12 The same object in time

Now use the stable evolution

$$\dot{x} = -Hx.$$

The packet equations are

$$\dot{x}_P = -H_{PP}x_P - H_{PQ}x_Q$$

and

$$\dot{x}_Q = -H_{QP}x_P - H_{QQ}x_Q.$$

Variation of constants gives

$$x_Q(t) = \mathrm{e}^{-tH_{QQ}}x_Q(0) - \int_0^t \mathrm{e}^{-(t-s)H_{QQ}}H_{QP}x_P(s)\,ds.$$

Substituting into the visible equation yields

$$\boxed{\begin{aligned}\dot{x}_P(t) = &-H_{PP}x_P(t) - H_{PQ}\mathrm{e}^{-tH_{QQ}}x_Q(0)\\ &+\int_0^t K(t-s)x_P(s)\,ds,\end{aligned}}$$

where

$$\boxed{K(t) = H_{PQ}\mathrm{e}^{-tH_{QQ}}H_{QP}.}$$

When the latent initial condition vanishes, only the convolution remains.

### 12.1 The memory kernel is boundary supported

Because the exteriors are disconnected and each touches only one retained boundary,

$$\boxed{K(t) = \begin{pmatrix} k_L(t) & 0 & 0\\ 0 & 0 & 0\\ 0 & 0 & k_R(t)\end{pmatrix}.}$$

The two scalar kernels are

$$k_L(t) = e_1^{\mathsf{T}}\mathrm{e}^{-t\widetilde{H}_L}e_1$$

and

$$k_R(t) = e_1^{\mathsf{T}}\mathrm{e}^{-t\widetilde{H}_R}e_1,$$

where the side matrices are ordered from the retained boundary outward,

$$\widetilde{H}_L = \begin{pmatrix} 6 & 2 & 0 \\ 2 & 7 & 1 \\ 0 & 1 & 8 \end{pmatrix}, \qquad \widetilde{H}_R = \begin{pmatrix} 7 & 2 & 0 \\ 2 & 9 & 1 \\ 0 & 1 & 10 \end{pmatrix}.$$

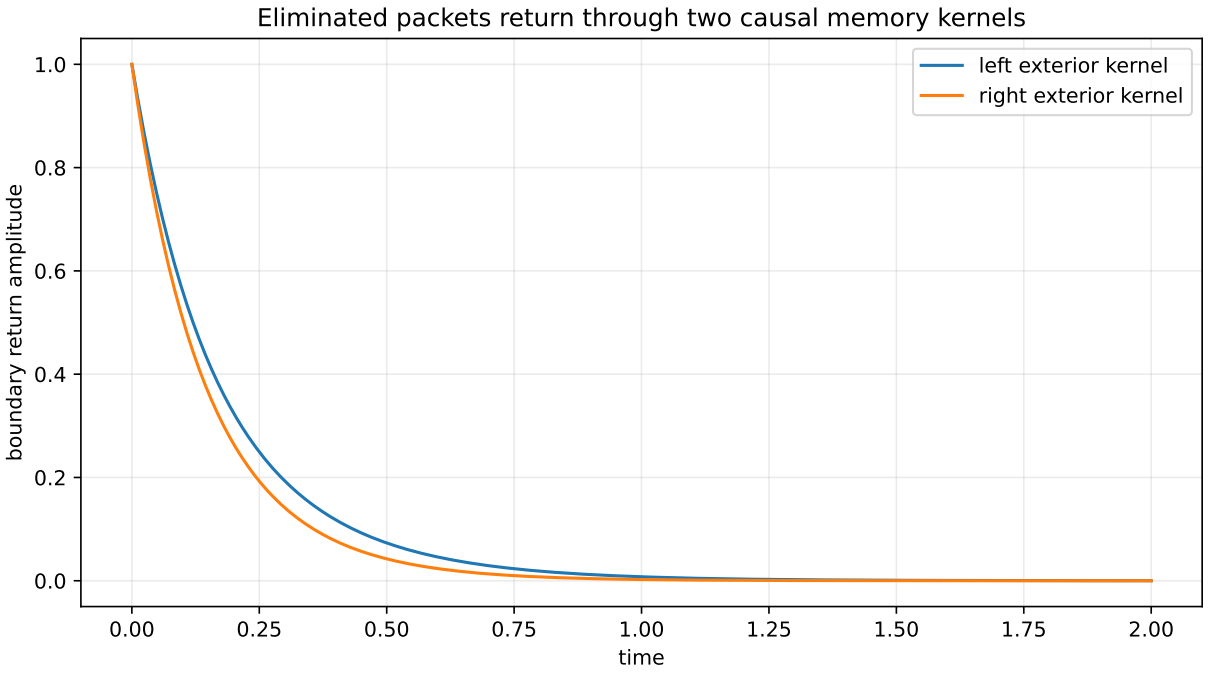



**Figure 10:** The exact left and right memory kernels. Both kernels are causal and decaying. Their different shapes record the different latent spectra on the two sides.

## 13 Modal decomposition of the exact memory

Diagonalize

$$\widetilde{H}_L = V_L \Lambda_L V_L^{\mathsf{T}}.$$

Then

$$\boxed{k_L(t) = \sum_{r=1}^{3} w_{L,r} \mathrm{e}^{-\lambda_{L,r} t}, \qquad w_{L,r} = |(V_L)_{1r}|^2.}$$

The right side has the same form.

The exact numerical modes are

| side and mode | $\lambda$ | $w$ |
|---|---|---|
| $L1$ | 4.330920912 | 0.572019980 |
| $L2$ | 7.523976397 | 0.241372019 |
| $L3$ | 9.145102691 | 0.186608001 |
| $R1$ | 5.697224362 | 0.690977865 |
| $R2$ | 9.302775638 | 0.197911024 |
| $R3$ | 11.000000000 | 0.111111111 |

Each side's weights sum to one, which gives

$$k_L(0) = k_R(0) = 1.$$

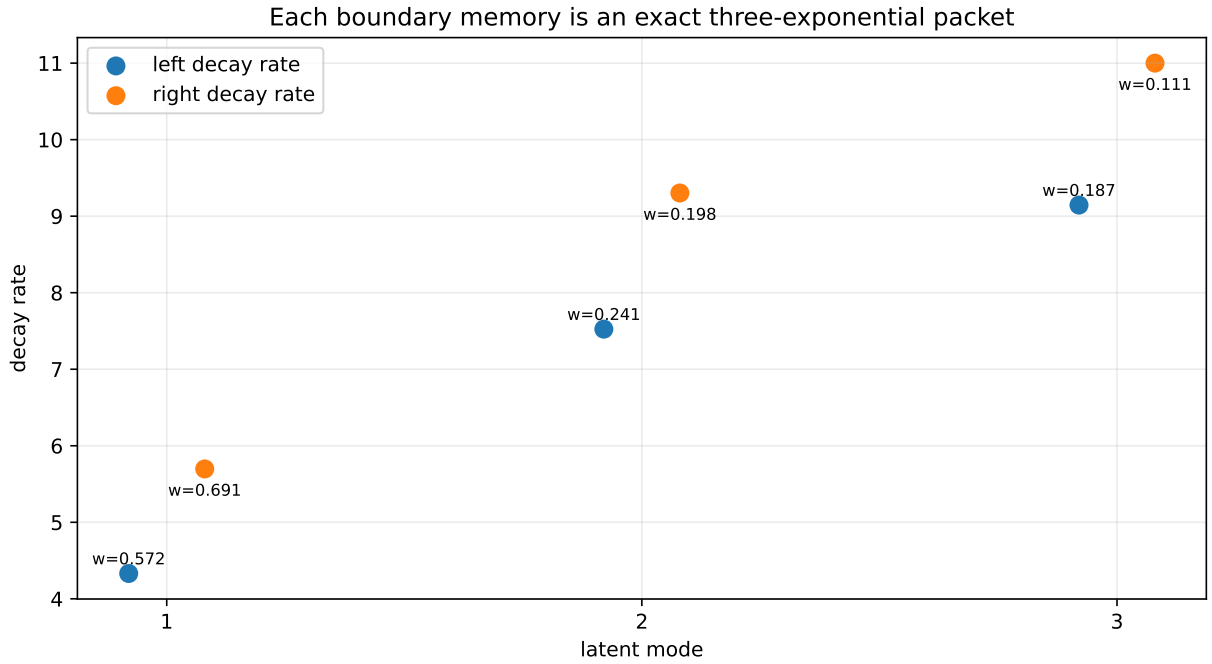



**Figure 11:** Decay rates and boundary weights. The latent packet is not represented by one relaxation time. Each exterior contributes three exact memory modes with different weights.

## 14 The self-energy is the Laplace transform of memory

For $\operatorname{Re} s > 0$,

$$\begin{aligned}\widehat{K}(s) &= \int_0^\infty \mathrm{e}^{-st} K(t)\, dt \\ &= H_{PQ}(sI + H_{QQ})^{-1} H_{QP}.\end{aligned}$$

Hence

$$\boxed{\widehat{k}_L(s) = \sum_{r=1}^{3} \frac{w_{L,r}}{s+\lambda_{L,r}} = \cfrac{1}{s+6-\cfrac{4}{s+7-\cfrac{1}{s+8}}}.}$$

Similarly,

$$\boxed{\widehat{k}_R(s) = \sum_{r=1}^{3} \frac{w_{R,r}}{s+\lambda_{R,r}} = \cfrac{1}{s+7-\cfrac{4}{s+9-\cfrac{1}{s+10}}}.}$$

The frequency-domain and time-domain objects are therefore the same transfer law under a change of spectral convention. The resolvent self-energy records the effect algebraically. The memory kernel records it causally.

The implementation checked the modal and matrix formulas at forty values of $s$. The largest discrepancy was

$$\boxed{1.12 \times 10^{-16}}.$$

## 15 Finite auxiliary-state realizations

A sum of exponentials can be represented by auxiliary states. For the left kernel, define

$$\dot{y}_{L,r} = -\lambda_{L,r} y_{L,r} - b_{L,r}^{\mathsf{T}} x_P.$$

The visible equation reads

$$\dot{x}_P = -H_{PP} x_P - C_L y_L - C_R y_R.$$

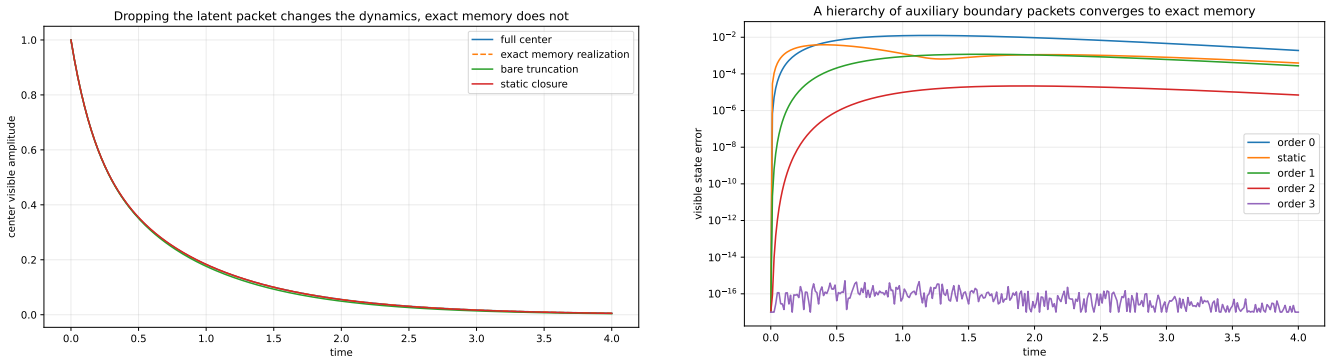
Eliminating $y_L$ and $y_R$ recovers the convolution kernel exactly.

The chosen chain admits a particularly transparent hierarchy. Retain the first $m$ Jacobi coordinates from each boundary. The auxiliary dimension is $2m$.

| $m$ | auxiliary dimension | largest visible-state error on $0 \le t \le 4$ |
|---|---|---|
| 0 | 0 | $1.2533 \times 10^{-2}$ |
| 1 | 2 | $1.1783 \times 10^{-3}$ |
| 2 | 4 | $2.1988 \times 10^{-5}$ |
| 3 | 6 | $5.1479 \times 10^{-16}$ |

A separate static approximation replaces the full memory by its zero-frequency value. Its largest error is

$$3.8933 \times 10^{-3}.$$




**(a)** Visible center dynamics **(b)** Error curves

**Figure 12:** Exact memory versus approximate closure. Bare truncation changes the visible trajectory. Static closure improves it. One- and two-layer dynamic closures improve it further. The complete three-layer realization agrees with the full system to roundoff.

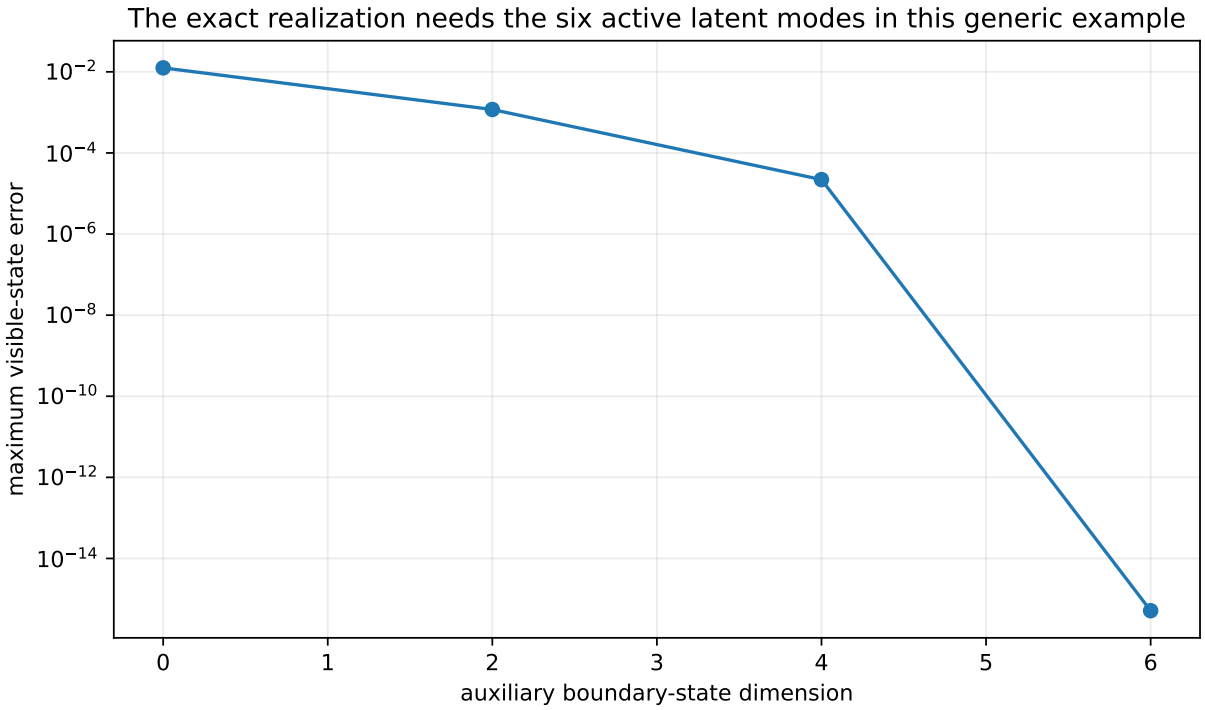



**Figure 13:** Accuracy versus auxiliary dimension. The boundary channel count is two, while the generic exact memory realization requires six latent modes. Rank and dynamical order measure different kinds of complexity.

## 16 Polyphase refinement of the same chain

Let $M = 4$. Group coordinates by their residue modulo four:

$$\Gamma_0 = \{0, 4, 8\},$$
$$\Gamma_1 = \{1, 5\},$$
$$\Gamma_2 = \{2, 6\},$$
$$\Gamma_3 = \{3, 7\}.$$

The phase-major order is

$$(0, 4, 8, 1, 5, 2, 6, 3, 7).$$

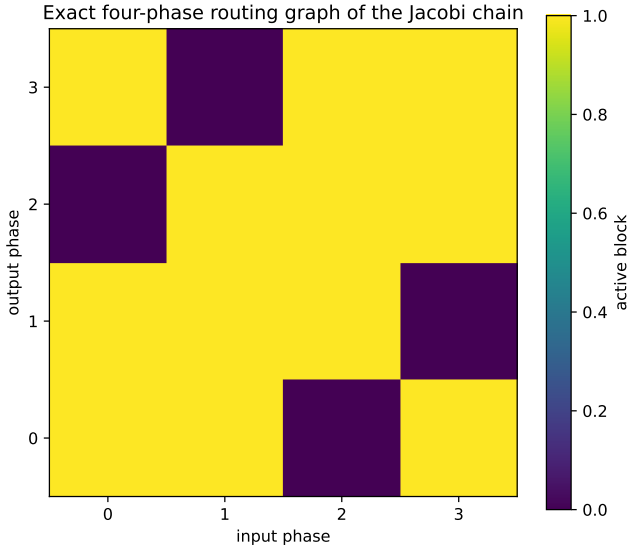
Let $\Pi_4$ be the corresponding permutation matrix. The refined operator is

$$H^{(4)} = \Pi_4 H \Pi_4^{\mathsf{T}}.$$

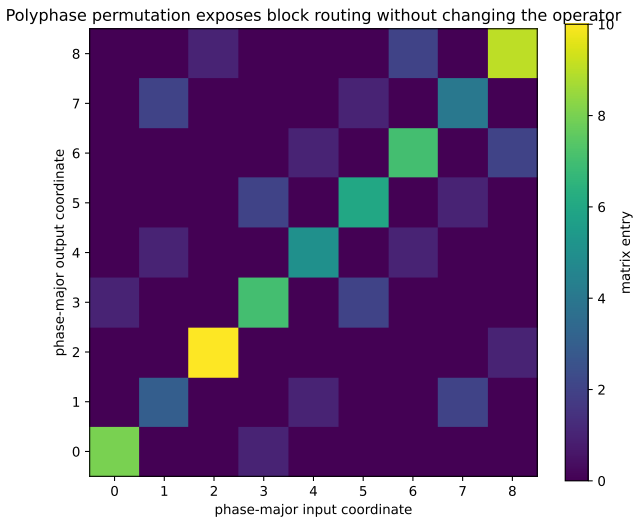
**Theorem 16.1** (No-free polyphase theorem)**.** *The polyphase representation is permutation similar to the original operator. Therefore it preserves spectrum, rank, determinant, and scalar nonzero count.*

*Proof.* Permutation similarity is an exact coordinate change. All similarity invariants are unchanged, and permutation only moves nonzero entries. □

The refinement still matters. It exposes which phase packets exchange information. For a nearest-neighbor stencil, a phase can couple only to itself and its neighboring residues.




**(a)** Four-phase transfer graph




**(b)** Phase-major operator

**Figure 14:** Polyphase structure of the exact chain. The phase graph exposes routing. It does not reduce the scalar stencil cost when every active block is evaluated.

The retained boundaries lie at coordinates three and five. Their phases are

$$3 \pmod 4$$

and

$$1 \pmod 4.$$

After retained-phase ordering, the self-energy occupies only those two phase packets. This is the exact compatibility between boundary elimination and polyphase refinement.

## 17 The boundary-rank theorem beyond tridiagonal chains

Suppose a banded matrix has half-bandwidth $b$, and one contiguous interval is retained. Each boundary crosses at most $b$ coordinate channels. Therefore the cross-block matrix has rank at most $2b$.

**Corollary 17.1** (Boundary-width rank)**.** *For a one-dimensional operator of half-bandwidth $b$,*

$$\boxed{\operatorname{rank} \Sigma(z) \le 2b.}$$

The code generated random banded matrices and observed exact ranks

$$2, 4, 6, 8$$

for half-bandwidths

$$1, 2, 3, 4.$$

In higher dimensions the analogous object is a boundary operator. Its dimension scales with resolved boundary area rather than eliminated volume.

## 18 A Schwarzschild scalar-wave realization

The exact nine-coordinate example is algebraically complete. To show that the same packet law survives in a more physical operator, consider the massless scalar Schwarzschild potential

$$V_\ell(r) = \left(1 - \frac{2M}{r}\right)\left[\frac{\ell(\ell+1)}{r^2} + \frac{2M}{r^3}\right].$$

Use the tortoise coordinate

$$r_* = r + 2M \log\left(\frac{r}{2M} - 1\right)$$

and discretize

$$H_\ell = -\frac{d^2}{dr_*^2} + V_\ell(r) - i\sigma(r)$$

on 240 points. The absorber $\sigma$ keeps the sampled finite-frequency matrices invertible. The retained packet is

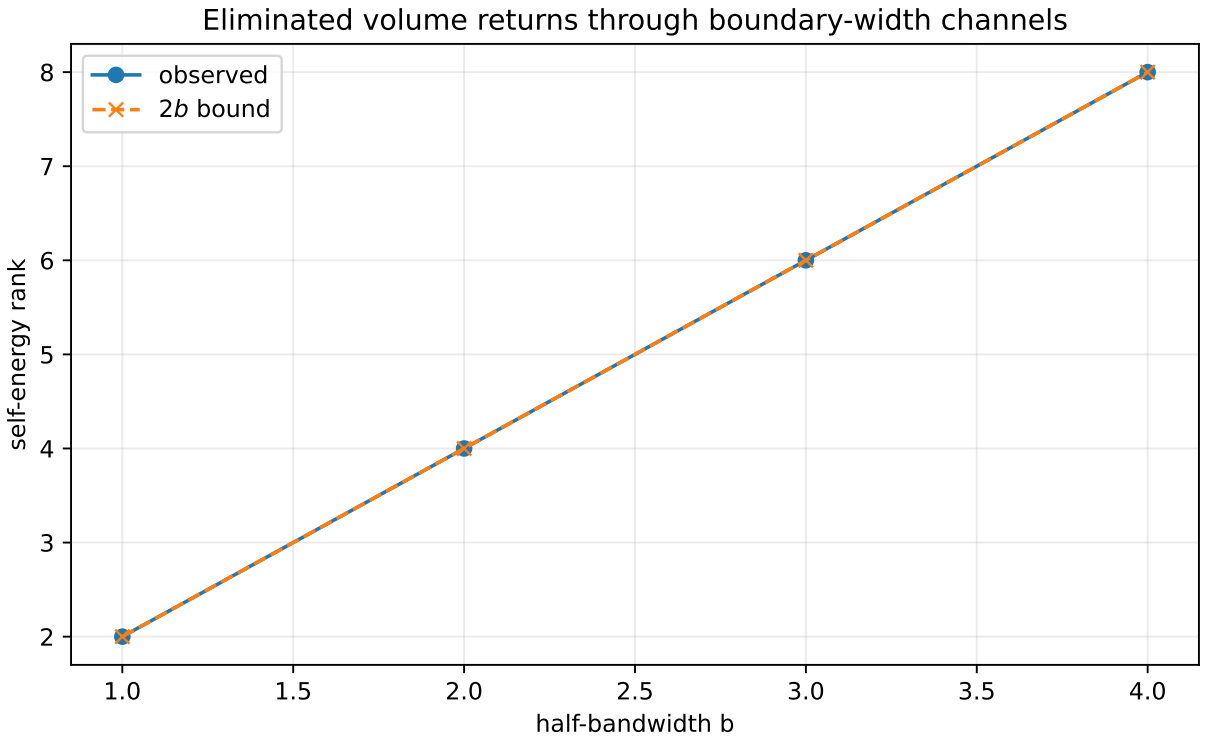



**Figure 15:** Self-energy rank follows boundary width. The eliminated volume may grow without changing this bound. What matters is how many channels cross the retained boundary.

$$-12 \le r_*/M \le 18.$$

Everything to the left and right is eliminated.

The exterior contains many coordinates, but a tridiagonal stencil crosses the retained interval at only two grid edges. Therefore the exact self-energy still has rank at most two.

Across nineteen frequencies from

$$M\omega = 0.28$$

to

$$M\omega = 0.72,$$

the self-energy rank was two at every sample. The largest full-field reconstruction error was

$$\boxed{2.53 \times 10^{-15}}.$$

*Boundary* 18.1 (Physical claim boundary). The frequency peak is a finite-domain response feature controlled by the grid, absorber, source, and observable. It is not presented as a rigorous quasinormal pole. The worked theorem concerns exact finite-dimensional elimination and boundary rank.

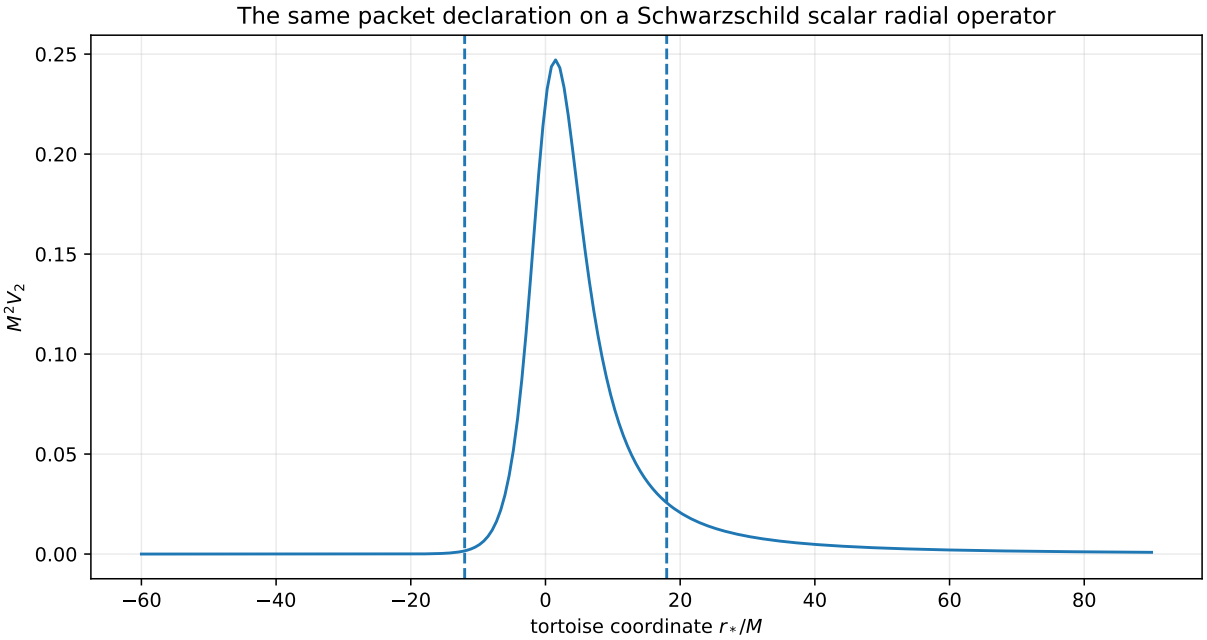
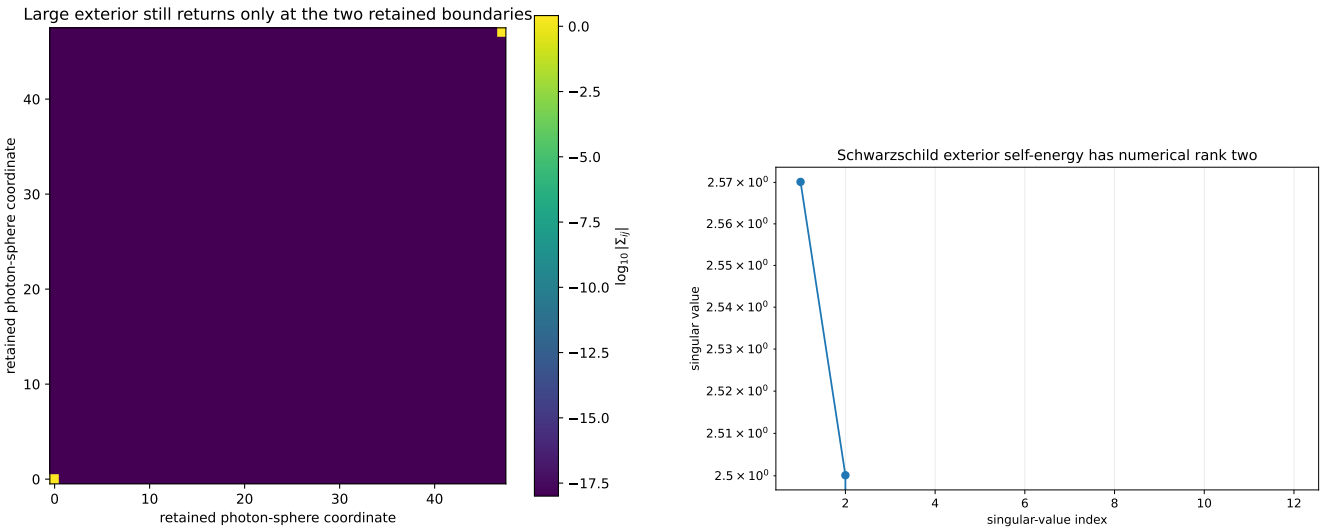



**Figure 16:** Schwarzschild potential and retained packet. The dashed lines mark the retained photon-sphere interval. The near-horizon and far-field regions form the latent exterior.




**(a)** Boundary support **(b)** Singular values

**Figure 17:** Exact Schwarzschild exterior self-energy. Only the two retained boundary coordinates receive nonzero exterior correction. The third and later singular values vanish at numerical precision.

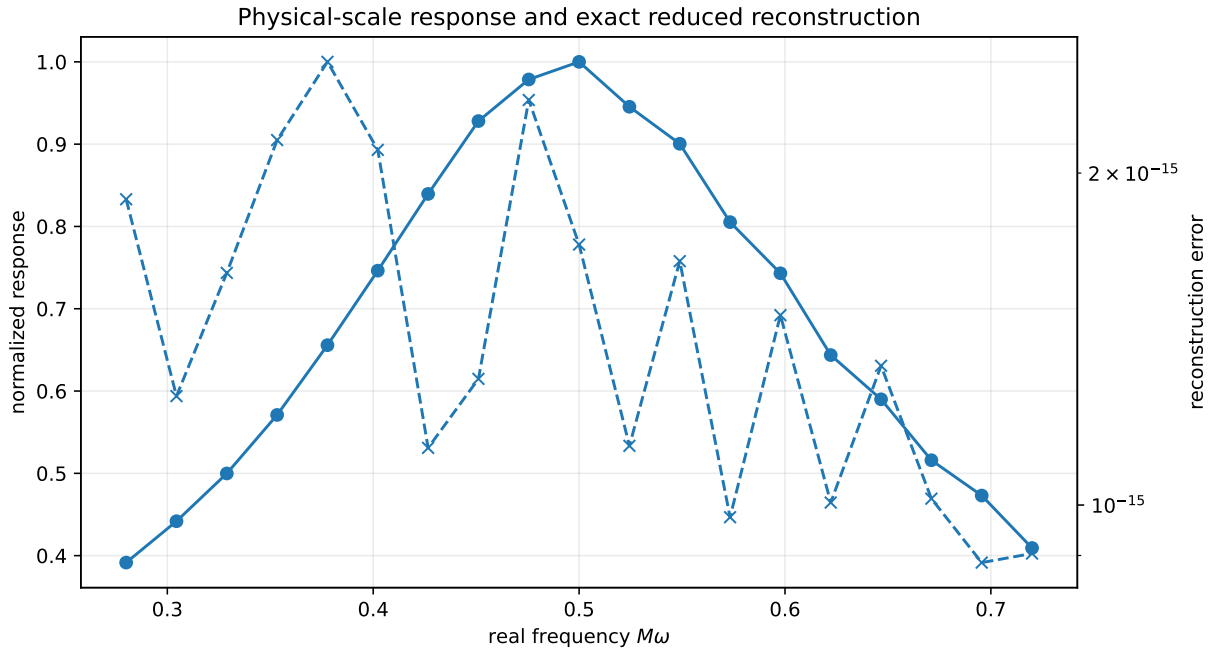



**Figure 18:** Response and reconstruction certificate. The finite-domain response peaks near the middle of the sampled band. Exact Feshbach reconstruction remains at floating-point accuracy through the peak.

# 19 What the single example teaches

## 19.1 Exact coarse graining is frequency dependent

A fixed visible matrix cannot generally reproduce the full resolvent at every spectral parameter. The exact reduced operator is

$$S(z) = A_{PP}(z) - \Sigma(z).$$

The eliminated packet returns through a rational function of $z$.

## 19.2 Exact coarse graining is nonlocal in time

The same elimination produces

$$\int_0^t K(t-s)x_P(s)\,ds.$$

A Markovian reduced model is exact only if the memory transfer has a finite realization that has been retained as auxiliary state.

## 19.3 Boundary rank and memory order are different

This example has two boundary channels and six generic memory modes. Boundary rank controls how the latent packet touches the visible packet. Memory order controls how much internal latent dynamics must be retained to reproduce that transfer exactly.

## 19.4 Polyphase refinement exposes routing, not free compression

The phase graph is useful for block scheduling, multirate organization, and nonlinear routing. It does not erase the self-energy or reduce scalar work by itself.

## 19.5 The same algebra appears across fields

The object

$$\Sigma(z) = A_{PQ}(zI - A_{QQ})^{-1}A_{QP}$$

can be read as

- a self-energy in quantum mechanics
- a Dirichlet-to-Neumann or boundary map in partial differential equations
- a transfer function in control
- an exact closure term in model reduction
- a memory kernel after inverse Laplace transform
- a packet-return operator in the projector calculus

The terminology changes. The algebra does not.

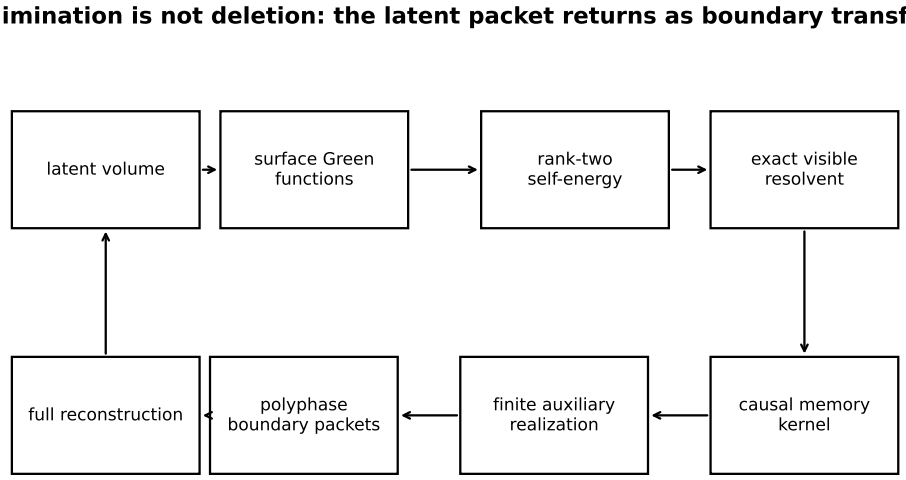



**Figure 19:** The final reconstruction cycle. The latent packet is compressed into boundary transfer and memory, not discarded. Exact synthesis recovers the full coordinate state.

## 20 General finite-dimensional statement

Let $A$ act on

$$\mathcal{H} = P\mathcal{H} \oplus Q\mathcal{H}.$$

Write

$$A = \begin{pmatrix} A_{PP} & A_{PQ} \\ A_{QP} & A_{QQ} \end{pmatrix}.$$

**Theorem 20.1** (Resolvent and memory pair)**.** *Whenever $zI - A_{QQ}$ is invertible, define*

$$A_{\mathrm{eff}}(z) = A_{PP} + A_{PQ}(zI - A_{QQ})^{-1}A_{QP}.$$

*Then*

$$P(zI - A)^{-1}P = (zI - A_{\mathrm{eff}}(z))^{-1}.$$

*For $\dot{x} = Ax$, the visible component satisfies*

$$\dot{x}_P(t) = A_{PP}x_P(t) + A_{PQ}\mathrm{e}^{tA_{QQ}}x_Q(0) + \int_0^t A_{PQ}\mathrm{e}^{(t-s)A_{QQ}}A_{QP}x_P(s)\,ds.$$

*Proof.* The resolvent identity follows by solving the latent block equation and substituting it into the visible block equation. The memory equation follows from variation of constants in the latent time equation and the same substitution. □

This theorem is the reusable core. The worked chain shows every matrix, rational function, pole, kernel mode, approximation, and reconstruction step explicitly.

## 21 Verification ledger

The single implementation generated every figure and table and performed twenty-three named checks.

**Table 1:** Complete verification ledger.

| Check | Result | Detail |
|---|---|---|
| Symmetric chain | pass | exact matrix symmetry |
| Positive chain | pass | smallest eigenvalue 1.175591661 |
| Left continued fraction | pass | symbolic equality |
| Right continued fraction | pass | symbolic equality |
| Determinant factorization | pass | symbolic residual zero |
| Exact rational reconstruction | pass | all nine fractions agree |

| Check | Result | Detail |
|---|---|---|
| Self-energy support | pass | only two boundary entries |
| Numerical self-energy rank | pass | rank two |
| Complex resolvent reconstruction | pass | $1.65 \times 10^{-16}$ |
| Sampled determinant identity | pass | maximum $2.95 \times 10^{-15}$ |
| Frequency reconstruction scan | pass | maximum $1.66 \times 10^{-15}$ |
| Left memory weights | pass | sum equals one |
| Right memory weights | pass | sum equals one |
| Memory Laplace transform | pass | maximum $1.12 \times 10^{-16}$ |
| Order-three time closure | pass | maximum $5.15 \times 10^{-16}$ |
| Polyphase spectrum | pass | permutation similarity |
| Polyphase reconstruction | pass | exact inverse permutation |
| Bandwidth rank, $b = 1$ | pass | rank $2 \leq 2b$ |
| Bandwidth rank, $b = 2$ | pass | rank $4 \leq 2b$ |
| Bandwidth rank, $b = 3$ | pass | rank $6 \leq 2b$ |
| Bandwidth rank, $b = 4$ | pass | rank $8 \leq 2b$ |
| Schwarzschild reconstruction | pass | maximum $2.53 \times 10^{-15}$ |
| Schwarzschild rank | pass | rank two at every frequency |

**Certificate 21.1** (Finite release certificate)**.** The source script produced twenty-five vector figures, twenty-five matching PNG figures, fifteen CSV ledgers, one machine-readable result file, and twenty-three passing checks. Exact symbolic calculations use SymPy. Numerical calculations use double-precision NumPy and SciPy. The black-hole realization is a finite discretization with an absorbing layer.

## 22 Conclusion

The example began with nine coordinates and ended with two boundary transfer functions. No latent effect was lost.

The left exterior became

$$g_L(z) = \frac{z^2 - 15z + 55}{z^3 - 21z^2 + 141z - 298}.$$

The right exterior became

$$g_R(z) = \frac{z^2 - 19z + 89}{(z - 11)(z^2 - 15z + 53)}.$$

Together they formed a rank-two self-energy. The exact reduced resolvent reconstructed the full nine-coordinate solution. In time, the same exteriors became two memory kernels with six total decay modes. A hierarchy of auxiliary boundary packets converged from visible errors near $10^{-2}$ to roundoff. Polyphase ordering exposed the routing but did not change the operator. A much larger Schwarzschild radial discretization showed the same rank-two boundary law.

The durable conclusion is

$$\boxed{\text{latent packet} \longrightarrow \text{boundary self-energy} \longleftrightarrow \text{causal memory} \longrightarrow \text{exact reconstruction.}}$$

That chain is one of the central closure mechanisms of the calculus.

# A Exact formula sheet

## A.1 Left exterior

$$\boxed{D_L(z) = z^3 - 21z^2 + 141z - 298}$$

$$\boxed{N_L(z) = z^2 - 15z + 55}$$

$$\boxed{g_L(z) = \frac{N_L(z)}{D_L(z)}.}$$

## A.2 Right exterior

$$\boxed{D_R(z) = (z - 11)(z^2 - 15z + 53)}$$

$$\boxed{N_R(z) = z^2 - 19z + 89}$$

$$\boxed{g_R(z) = \frac{N_R(z)}{D_R(z)}.}$$

## A.3 Exact zero-spectral solution

$$\boxed{u(0) = \frac{1}{1869341}(-3768, 30144, -103620, 561432, -1071054, 220957, -33731, 7580, -758)^{\mathsf{T}}.}$$

## A.4 Time kernels

$$\boxed{k_L(t) = 0.572019980\mathrm{e}^{-4.330920912t} + 0.241372019\mathrm{e}^{-7.523976397t} + 0.186608001\mathrm{e}^{-9.145102691t}.}$$

$$\boxed{k_R(t) = 0.690977865\mathrm{e}^{-5.697224362t} + 0.197911024\mathrm{e}^{-9.302775638t} + 0.111111111\mathrm{e}^{-11t}.}$$

# B Selected machine-readable ledgers

The release includes the following CSV files.

# 7 Schur elimination and nonnormal memory

## 7.1 A projected resolvent needs two invertibility facts

For retained and hidden coordinates $P \oplus Q$, the Schur complement is

$$S(z) = zI - A_{PP} - A_{PQ}(zI - A_{QQ})^{-1}A_{QP}.$$

`SchurFeshbach` uses literal complex block matrices. The projected inverse formula assumes invertibility of both the full resolvent matrix and the hidden block:

Source: `KomorebiPaperVI/SchurFeshbach.lean`.

```
theorem projected_resolvent (App : Matrix P P ℂ) (Apq : Matrix P Q ℂ)
    (Aqp : Matrix Q P ℂ) (Aqq : Matrix Q Q ℂ) (z : ℂ)
    [Invertible (Matrix.scalar Q z-Aqq)]
    [Invertible (Matrix.scalar (P ⊕ Q) z-Matrix.fromBlocks App Apq Aqp Aqq)] :
    ((Matrix.scalar (P ⊕ Q) z-Matrix.fromBlocks App Apq Aqp Aqq)⁻¹).submatrix
      Sum.inl Sum.inl = (schur App Apq Aqp Aqq z)⁻¹ := by
  let D := Matrix.scalar Q z-Aqq
  let A := Matrix.scalar P z-App
  have he := resolvent_blocks App Apq Aqp Aqq z
  let : Invertible (Matrix.fromBlocks A (-Apq) (-Aqp) D) := Invertible.copy ‹_› _ he.symm
  let : Invertible (A-(-Apq)*⅟D*(-Aqp)) := Matrix.invertibleOfFromBlocks₂₂Invertible A (-Apq)
   ↪  (-Aqp) D
  have hinv := Matrix.invOf_fromBlocks₂₂_eq A (-Apq) (-Aqp) D
  have hs : A-(-Apq)*⅟D*(-Aqp)=schur App Apq Aqp Aqq z := by
    simp only [Matrix.neg_mul,Matrix.mul_neg,neg_neg,Matrix.invOf_eq_nonsing_inv,schur,D,A]
  rw [he,← Matrix.invOf_eq_nonsing_inv,hinv]
  change ⅟(A-(-Apq)*⅟D*(-Aqp))=(schur App Apq Aqp Aqq z)⁻¹
  rw [Matrix.invOf_eq_nonsing_inv]
  congr 1
```

The module also proves the invertibility equivalence between the full block matrix and its Schur complement when the hidden block is invertible. The retained block is an inverse Schur complement, not simply the resolvent of $A_{PP}$. The difference is the self-energy $A_{PQ}(zI - A_{QQ})^{-1}A_{QP}$.

## 7.2 The fifth checked block removes a normality shortcut

The earlier memory layer establishes Laplace inversion from a supplied exponential growth estimate and also treats normal elements. The latest core 13.2 goes further for arbitrary finite complex matrices: a strict spectral half-plane alone yields integrability and decay. Nonnormality can create polynomial factors in time, so the proof must account for nilpotent generalized eigenspace contributions.

The scalar input is integrability of a complex exponential times a polynomial whenever the real part of the exponent is negative:

Source: `KomorebiPaperVI/ScalarDecay.lean`.

```
theorem exp_pow_integrable (a : ℂ) (ha : a.re < 0) (k : ℕ) :
    IntegrableOn (fun t : ℝ => Complex.exp (a*t)*(t:ℂ)^k) (Ioi 0) := by
  have hi := integrableOn_rpow_mul_exp_neg_mul_rpow
    (s := (k:ℝ)) (p := (1:ℝ)) (b := -a.re) (by linarith [Nat.cast_nonneg (α := ℝ) k])
     ↪  zero_lt_one (by linarith)
  simp only [Real.rpow_natCast,Real.rpow_one,neg_neg] at hi
```

```
  apply hi.mono' (by fun_prop)
  filter_upwards [ae_restrict_mem measurableSet_Ioi] with t ht
  rw [norm_exp_pow a k t ht.le]
```

There is also a separate zero-limit theorem. Neither a numerical quadrature nor an assertion about eigenvalue locations substitutes for these analytic steps.

## 7.3 From generalized eigenspaces to a finite sum

`NonnormalMemory` studies left multiplication by $A$ on a finite-dimensional complex $C^*$-algebra. A complex matrix algebra is a specialization. The source decomposes this vector space into generalized eigenspaces. On one component the exponential is a finite polynomial times $e^{\mu t}$:

Source: `KomorebiPaperVI/NonnormalMemory.lean`.

```
theorem exp_mul_component (A : R) (μ : ℂ) (x : R)
    (hx : x ∈ (Module.End.maxGenEigenspace (LinearMap.mulLeft ℂ A)) μ) (t : ℝ) :
    transfer A t*x=Complex.exp (μ*t) •
      ∑ k ∈ Finset.range (finrank ℂ R),((t:ℂ)^k/(k.factorial:ℂ)) • ((A-μ • 1)^k*x) := by
  classical
  let : NormedAlgebra ℚ R := NormedAlgebra.restrictScalars ℚ ℝ R
  let D := A-μ • 1
  have hkill (k : ℕ) (hk : finrank ℂ R ≤ k) : D^k*x=0 := by
    rw [← left_shift_power]
    exact CanonicalJets.power_kills _ μ x hx k hk
  have hs := (exp_series_hasSum_exp' (𝕂 := ℂ) ((t:ℂ) • D)).mul_right x
  have he : exp ((t:ℂ) • D)*x=
      ∑ k ∈ Finset.range (finrank ℂ R),((t:ℂ)^k/(k.factorial:ℂ)) • (D^k*x) := by
    have hs' : HasSum (fun k : ℕ => ((t:ℂ)^k/(k.factorial:ℂ)) • (D^k*x))
        (exp ((t:ℂ) • D)*x) := by
      simpa only [smul_pow,Algebra.smul_mul_assoc,smul_smul,div_eq_mul_inv,mul_comm] using hs
    exact hs'.unique (hasSum_sum_of_ne_finset_zero (fun k hk => by
      rw [hkill k (by simpa using hk),smul_zero]))
  have ht : t • A=algebraMap ℂ R (μ*t)+(t:ℂ) • D := by
    simp only [D,smul_sub,Algebra.algebraMap_eq_smul_one,smul_smul]
    rw [RCLike.real_smul_eq_coe_smul (K := ℂ),mul_comm μ (t:ℂ)]
    abel
  rw [transfer,ht,exp_add_of_commute (Algebra.commutes _ _),
    ← algebraMap_exp_comm,← Complex.exp_eq_exp_ℂ,Algebra.algebraMap_eq_smul_one,
    Algebra.smul_mul_assoc,one_mul,Algebra.smul_mul_assoc,he]
```

The cutoff is the complex dimension of the ambient algebra. For an $n \times n$ matrix algebra that can be $n^2$; it need not be a minimal Jordan-block length. The proof only needs a valid nilpotence cutoff. Terms beyond that cutoff vanish on the component.

After damping by $e^{-zt}$, each coefficient has exponent $\mu - z$ with negative real part. The source combines the scalar integrability and limit results over finite sums and then over the generalized-eigenspace decomposition. Applying the resulting identity to the algebra unit gives the operator-valued decay:

Source: `KomorebiPaperVI/NonnormalMemory.lean`.

```
theorem spectral_decay (A : R) (z : ℂ) (hz : ∀ μ∈spectrum ℂ A,μ.re < z.re) :
    IntegrableOn (damped A z) (Ioi 0) ∧ Tendsto (damped A z) atTop (𝓝 0) := by
  have h (x : R) : IntegrableOn (fun t => damped A z t*x) (Ioi 0) ∧
      Tendsto (fun t => damped A z t*x) atTop (𝓝 0) := by
```

```
    apply Submodule.iSup_induction (Module.End.maxGenEigenspace (LinearMap.mulLeft ℂ A))
      (x := x) (motive := fun x => IntegrableOn (fun t : ℝ => damped A z t*x) (Ioi 0) volume
      ↪ ∧
        Tendsto (fun t => damped A z t*x) atTop (𝓝 0))
      (by rw [Module.End.iSup_maxGenEigenspace_eq_top]; trivial)
    · exact component_decay A z hz
    · simp
    · intro x y hx hy
      simpa only [Pi.add_def,mul_add,add_zero] using And.intro (hx.1.add hy.1) (hx.2.add
      ↪ hy.2)
  simpa using h 1
```

The condition is strict: $\operatorname{Re}\mu < \operatorname{Re}z$ for every spectral value. The conclusion includes both Bochner integrability and the limit zero. It does not rely on a claimed universal bound $\|e^{tA}\| \le e^{t\max \operatorname{Re}\sigma(A)}$, which would fail for general nonnormal matrices.

## 7.4 Why both derivative identities appear

The improper-integral fundamental theorem is applied once to left multiplication and once to right multiplication. These establish both inverse equations. This is what connects the integral to the actual inverse, rather than just presenting a formal solution to one equation.

Source: `KomorebiPaperVI/NonnormalMemory.lean`.

```
theorem spectral_laplace_inverse (A : R) (z : ℂ)
    (hz : ∀ μ∈spectrum ℂ A,μ.re < z.re) :
    IntegrableOn (damped A z) (Ioi 0) ∧ IsUnit (algebraMap ℂ R z-A) ∧
      (∫ t in Ioi 0,damped A z t)=Ring.inverse (algebraMap ℂ R z-A) := by
  have hs := spectral_decay A z hz
  let B := A-algebraMap ℂ R z
  have he : damped A z=transfer B := funext (damped_eq_shifted A z)
  rw [he] at hs ⊢
  have hl := integral_Ioi_of_hasDerivAt_of_tendsto'
    (fun t (_ : t∈Ici (0:ℝ)) => transfer_derivative_left B t)
    ((ContinuousLinearMap.mul ℝ R B).integrable_comp hs.1) hs.2
  have hr := integral_Ioi_of_hasDerivAt_of_tendsto'
    (fun t (_ : t∈Ici (0:ℝ)) => transfer_derivative_right B t)
    (((ContinuousLinearMap.mul ℝ R).flip B).integrable_comp hs.1) hs.2
  rw [integral_const_mul_of_integrable hs.1,transfer_zero,zero_sub] at hl
  rw [integral_mul_const_of_integrable hs.1,transfer_zero,zero_sub] at hr
  have hn : algebraMap ℂ R z-A= -B := by dsimp [B]; abel
  have hl' : (algebraMap ℂ R z-A)*(∫ t in Ioi 0,transfer B t)=1 := by
    rw [hn,neg_mul,hl,neg_neg]
  have hr' : (∫ t in Ioi 0,transfer B t)*(algebraMap ℂ R z-A)=1 := by
    rw [hn,mul_neg,hr,neg_neg]
  have hu : IsUnit (algebraMap ℂ R z-A) := ⟨⟨_,_,hl',hr'⟩,rfl⟩
  refine ⟨hs.1,hu,?_⟩
  calc
    _ = Ring.inverse (algebraMap ℂ R z-A)*((algebraMap ℂ R z-A)*(∫ t in Ioi 0,transfer B t))
    ↪ :=
      (Ring.inverse_mul_cancel_left _ _ hu).symm
    _ = _ := by rw [hl',mul_one]
```

The result supplies invertibility as part of the conclusion. The total operation `Ring.inverse` is therefore being used at an established unit. This is the formal endpoint connecting the strict spectral half-plane with a genuine resolvent formula.

## 7.5 Retained and hidden dimensions can differ

The final matrix theorem uses $A$ on hidden coordinates $J$, $U : J \to I$, and $V : I \to J$. It wraps $B \mapsto UBV$ as a continuous linear map and commutes that map with the Bochner integral. The result is the literal rectangular memory/self-energy identity:

Source: `KomorebiPaperVI/MatrixMemory.lean`.

```
theorem rectangular_memory_self_energy (A : Matrix J J ℂ)
    (U : Matrix I J ℂ) (V : Matrix J I ℂ) (z : ℂ)
    (hz : ∀ μ∈spectrum ℂ A,μ.re < z.re) :
    IntegrableOn (fun t : ℝ => Complex.exp (-z*t) • (U*transfer A t*V)) (Ioi 0) ∧
    (∫ t : ℝ in Ioi 0,Complex.exp (-z*t) • (U*transfer A t*V))=
      U*(Matrix.scalar J z-A)⁻¹*V := by
  have h := NonnormalMemory.spectral_laplace_inverse A z hz
  have he : (fun t : ℝ => Complex.exp (-z*t) • (U*transfer A t*V))=
      (fun t => sandwich U V (damped A z t)) := by
    ext t i j
    simp [sandwich,damped]
  rw [he]
  refine ⟨(sandwich U V).integrable_comp h.1,?_⟩
  rw [(sandwich U V).integral_comp_comm h.1,h.2.2]
  change U*Ring.inverse (algebraMap ℂ (Matrix J J ℂ) z-A)*V=_
  rw [Matrix.nonsing_inv_eq_ringInverse]
  rfl
```

The source imposes no equality of retained and hidden dimensions and no nonemptiness assumption on the finite index types. The theorem includes nilpotent Jordan parts and zero-dimensional cases. It proves a finite-dimensional Laplace identity; general infinite hidden systems, globally optimal model reduction, and arbitrary memory closure remain separate questions.

Observed output:

```
1 1 1
[0.0]
[0.09] True
[2.0] True
```

The valid point at $\kappa = 0.3$ is not a root. It lies on the eliminated range branch and retains the nonzero scalar obstruction $0.09$.

# 41.4 Nonlinear Schur elimination

## 41.4.1 Retained and hidden equations

Split a variable as $x = (u, v)$, where $u \in \mathbb{K}^{d_P}$ is retained and $v \in \mathbb{K}^{d_Q}$ is hidden. Split the equation as

$$F(u, v) = \big(F_P(u, v), F_Q(u, v)\big).$$

At a base point $(u_0, v_0)$, require

$$F_Q(u_0, v_0) = 0, \qquad D_v F_Q(u_0, v_0) \text{ invertible.}$$

The local implicit branch $v = \phi(u)$ then satisfies

$$Dv = -(D_v F_Q)^{-1} D_u F_Q, \qquad F_{\mathrm{red}}(u) = F_P(u, v(u)).$$

Writing that branch as $v = \phi(u)$, the first identity becomes

$$D\phi = -(D_v F_Q)^{-1} D_u F_Q.$$

The reduced retained map

$$F_{\mathrm{red}}(u) = F_P(u, \phi(u))$$

has derivative

$$DF_{\mathrm{red}} = D_u F_P + D_v F_P Dv.$$

Equivalently, in the $\phi$ notation,

$$DF_{\mathrm{red}} = D_u F_P + D_v F_P D\phi = D_u F_P - D_v F_P (D_v F_Q)^{-1} D_u F_Q.$$

This is the nonlinear analogue of a Schur complement. The finite derivative formula is **E** when both providers are exact SymPy providers, and **N** for numerical providers. The assertion that these derivatives arise from a local branch remains **C**.

### 41.4.2 Second-order formula

For retained directions $h, k$, lift them to tangent directions on the branch:

$$\widetilde{h} = (h, D\phi[h]), \qquad \widetilde{k} = (k, D\phi[k]).$$

Differentiating the hidden equation twice gives

$$D^2 v[h,k] = -(D_v F_Q)^{-1} D^2 F_Q[(h, Dv[h]), (k, Dv[k])],$$

In the $\phi$ notation this is

$$D^2\phi[h,k] = -(D_v F_Q)^{-1} D^2 F_Q[\widetilde{h}, \widetilde{k}].$$

Therefore

$$D^2 F_{\mathrm{red}}[h,k] = D^2 F_P[(h, Dv[h]), (k, Dv[k])] + D_v F_P D^2 v[h,k].$$

The equivalent lifted-direction form used by the implementation is

$$D^2 F_{\mathrm{red}}[h,k] = D^2 F_P[\widetilde{h}, \widetilde{k}] + D_v F_P D^2\phi[h,k].$$

The implementation also assembles the equivalent displayed subtraction form and certifies that both finite expressions agree. This check cannot detect a shared error in the supplied second-derivative providers.

### 41.4.3 Exact and numerical modes

The constructor decides its mode from the two providers:

- if both providers declare `Exactness.EXACT`, inputs and derivatives are represented as SymPy matrices and the certificate exactness is `exact`
- otherwise, both blocks are evaluated through complex NumPy arrays and the certificate exactness is `numerical`

Both providers must use the concatenated source dimension $d_P + d_Q$. The hidden provider's target dimension must equal $d_Q$. The retained provider may have an arbitrary positive target dimension.

`first_derivatives()` returns the hidden Jacobian, its inverse, the implicit derivative, the reduced derivative, and a certificate with theorem identifier `PSC-III-T41`. `second_derivative(left, right)` returns the two lifted directions, the implicit second derivative, the reduced second derivative, and a certificate with theorem identifier `PSC-III-T42`.

`solve_hidden` is intentionally numerical only. For a fixed retained value it uses damped Newton iteration on $F_Q(u, v) = 0$, with at most twenty-four backtracking trials per outer iteration. Exact-mode callers must supply or manipulate a symbolic branch externally. `reduced_value` first requires a valid hidden solve and then evaluates $F_P$ on that point. The numerical hidden-solve and reduced-value certificates use the historical theorem identifier `PSC-III-T40`.

### 41.4.4 Exact second-order example

This example requires the optional SymPy dependency.

```
import sympy as sp

from komorebi_math.nonlinear import (
    NonlinearSchurElimination,
    SympyDerivativeProvider,
)

u, v = sp.symbols("u v")

retained = SympyDerivativeProvider(
    (3 * u + 2 * v + u * v + v**2,),
    (u, v),
)
hidden = SympyDerivativeProvider(
    (v - 4 * u + u**2 + u * v,),
    (u, v),
)

problem = NonlinearSchurElimination(
    retained,
    hidden,
    1,
    1,
    sp.Matrix([0]),
    sp.Matrix([0]),
)

first = problem.first_derivatives()
second = problem.second_derivative(sp.Matrix([1]), sp.Matrix([1]))

print(list(first.implicit_derivative))
print(list(first.reduced_derivative))
print(list(second.implicit_second_derivative))
print(list(second.reduced_second_derivative))
print(first.certificate.exactness.value)
```

Observed with SymPy `1.14.0`:

```
[4]
[11]
[-10]
[20]
exact
```

### 41.4.5 Numerical hidden branch

For the same hidden equation, the branch is available explicitly as

$$v(u) = \frac{4u - u^2}{1 + u}$$

near zero. The focused package test evaluates `solve_hidden([0.05])` and matches this expression to absolute tolerance $10^{-12}$. A valid numerical certificate checks only the returned hidden residual. It leaves uniqueness, the size of the implicit neighborhood, and absence of remote solutions unchecked.

## 41.5 Public signature ledger

These signatures are the reflected version `1.0.0b4` contracts. Type annotations are descriptive runtime annotations, not automatic validators.

### 41.5.1 Correction surface

```
HilbertObstructionProjector(
    correction_map: Any,
    *,
    tolerance: float = 1e-10,
)
    .solve(defect: Any) -> HilbertObstructionResult
```

```
RightInverseCorrectionProblem(
    linear_map: Any,
    right_inverse: Any,
    nonlinear: Callable[[np.ndarray], Any],
    target: Any,
    *,
    radius: float,
    lipschitz_bound: float,
    tolerance: float = 1e-10,
)
    .hypothesis_certificate() -> Certificate
    .solve(
        *,
        max_iterations: int = 10000,
    ) -> RightInverseCorrectionResult
    .target_stability_bound(
        target_distance: float,
    ) -> float
```

```
LyapunovSchmidtReduction(
    provider: DerivativeProvider,
    base_point: Any,
    *,
    tolerance: float = 1e-10,
)
    .solve_range(
        kernel_coordinates: Any,
        *,
        initial_complement: Any | None = None,
        max_iterations: int = 50,
    ) -> LyapunovSchmidtPoint
    .base_derivative_certificate() -> Certificate
    .reduced_second_derivative(
        left: Any,
        right: Any,
    ) -> ComputationResult[np.ndarray]
```

### 41.5.2 Elimination surface

```
NonlinearSchurElimination(
    retained_provider: DerivativeProvider,
    hidden_provider: DerivativeProvider,
    retained_dimension: int,
    hidden_dimension: int,
    base_retained: Any,
    base_hidden: Any,
    *,
    tolerance: float = 1e-10,
)
    .first_derivatives() -> SchurDerivativeResult
    .second_derivative(
        left: Any,
        right: Any,
    ) -> SchurSecondDerivativeResult
    .solve_hidden(
        retained_value: Any,
        *,
        initial_hidden: Any | None = None,
        max_iterations: int = 50,
    ) -> HiddenSolveResult
    .reduced_value(
        retained_value: Any,
        *,
        initial_hidden: Any | None = None,
        max_iterations: int = 50,
    ) -> ComputationResult[np.ndarray]
```

## 41.6 Result-record ledger

All named records below are frozen, slotted dataclasses. Their NumPy arrays remain mutable objects, so `frozen=True` prevents field reassignment but does not make array storage immutable.

- `HilbertObstructionResult`
  - `obstruction`
  - `obstruction_norm`
  - `vanishes`
  - `least_norm_correction`
  - `correction_residual`
  - `certificate`
- `RightInverseCorrectionResult`
  - `solution`
  - `iterations`
  - `contraction_factor`
  - `equation_residual`
  - `fixed_point_residual`
  - `certificate`
- `LyapunovSchmidtDecomposition`
  - `linearization`
  - `kernel_basis`
  - `complement_basis`
  - `range_basis`
  - `obstruction_basis`
  - `rank`
  - `certificate`
- `LyapunovSchmidtPoint`
  - `kernel_coordinates`

    - complement_coordinates
    - point
    - reduced_obstruction
    - range_residual
    - iterations
    - certificate
- SchurDerivativeResult
    - hidden_jacobian
    - hidden_inverse
    - implicit_derivative
    - reduced_derivative
    - certificate
- SchurSecondDerivativeResult
    - lifted_left
    - lifted_right
    - implicit_second_derivative
    - reduced_second_derivative
    - certificate
- HiddenSolveResult
    - hidden
    - point
    - residual
    - iterations
    - certificate

There are no registered package portable codecs for these Chapter 41 records in this release. If they must cross a process boundary, serialize a deliberate schema containing array dtype, shape, real and imaginary components, theorem identifier, and certificate fields. A raw dataclass conversion is not by itself a stable persistence contract.

## 41.7 Error and certificate interpretation

### 41.7.1 Exception families

- `DimensionError` reports matrix rank, vector length, or concatenated-domain shape disagreement.
- `ValidationError` reports nonfinite values, invalid tolerances, missing derivative orders, an uncertified contraction, a bad base root, a singular numerical Jacobian, a failed damped step, or an exact-mode request for the numerical hidden solver.
- Calling `.require()` on an invalid `Certificate` also raises `ValidationError` with the certificate message.
- SymPy is optional. Exact provider construction without it follows that provider's `BackendUnavailableError` path.

An iteration limit is not uniformly an exception boundary. Both the right-inverse solver and the two Newton solvers may return an invalid result certificate after reaching the cap. Singular linear systems and failed backtracking steps raise immediately.

# Exact-arithmetic evidence

Source 1: QUANTITATIVE_OVERLAP_REPAIR.md, section 8. Komorebi Space Theory v0.17.0. Original wording, reflowed for this collection.

## 8. Computation, formal scope, and open boundaries

Five new proof modules add 78 named theorems and 42 indexed constructions. The cumulative library contains 107 proof modules, 1,071 named theorems, and 601 indexed constructions. All 102 v0.16 proof modules are unchanged. Counts include elementary and auxiliary declarations, not claimed discoveries.

The separate rational audit uses signed permutation frames and checks 1,080 deterministic graph/frame/candidate configurations, 2,160 independent edge residual assignments, and 40 chain lengths. It records 41,301 assertions and five rejected false-claim controls. An independent Gaussian-elimination test is compared with the route integrability criterion. These tests do not enumerate all graphs or frames and are not the proof of the universal conclusions.

The executable audit initially failed because integer entries divided by integer pivots introduced floating-point values into the Gaussian checker. Every matrix entry and right-hand side is now explicitly converted to `Fraction` before elimination. The failed attempt is retained and not counted as evidence of correctness. Lean development errors and the interrupted first baseline launcher are separately recorded.

Closed within this model: actual continuous linear residual and correction operators; a graph-derived sharp-in-chains error bound; preservation of compatible prescribed values; explicit ambient reconstruction; connection to the old correction scheme; an image/integrability diagnostic; and isometric coordinate covariance.

Not proved: existence of globally trivializing isometric frames for arbitrary local systems; repair of nonflat or nonlinear cocycles; a numerical structure on arbitrary stacks; uniform conditioning on arbitrary growing graphs; spectral/optimal estimates on general graphs; orthogonal or unanchored best approximation; nonisometric presentation-independent norms; infinite-graph estimates; executable extraction or complexity of the Lean noncomputable definitions; or unrestricted arithmetic correspondence.

The Lean compiler/kernel, pinned imported caches, and `propext`, `Classical.choice`, `Quot.sound` remain trusted. The current and fresh build reports control verification status. The archive does not include the multi-gigabyte runtime.

# Fraction-free and modular elimination

Source 2: FINITE_CERTIFICATE.md. Complete source text, including the declared finite scope and independent-check description. Original wording, reflowed.

## Exact finite certificate and its reduction

The universal theorem B55.3 is not an enumeration over Diophantine solutions. For a declared set of possible high-congruence prime divisors it reduces actual eta triviality to exact residue symbols and restricted matrix ranks. The tested finite universe U has ten rational primes and twelve places. It includes the inert h=3 prime1373 and is NOT asserted to cover arbitrary C.

## Primary implementation

maximal_layer.py imports a byte-identical copy of the previous phi-basis residue arithmetic. For every place it computes complete seven-primary logs, then reduces them modulo7. It constructs the five global rows and the three natural Kummer rows in matching coordinates. For each of42 normalized F7 candidates it selects the Kummer zero columns and computes exact F7 rank. It checks all1024 rational subsets of U.

It also enumerates all1585 prime-ideal subsets of cardinality1 through5. A column subset is a circuit precisely when its kernel has dimension1 and a kernel generator has no zero coordinate. A compatible circuit additionally admits a normalized Kummer coefficient vector with nonzero ratio. The universal circuit bound ensures these subsets suffice for the criterion. The circuit-containment classification agrees with all1024 direct rank tests.

All2352 normalized F49 directions are separately enumerated. Their zero sets are checked to be contained in the corresponding valid F7 zero set from the proof of B55.4. Of these,2310 normalized directions are outside F7. None is silently discarded or identified with the actual coefficients.

## Independent implementation

independent_certificate.py imports no primary arithmetic module. It uses sqrt5-basis residue fields and direct seventh-power symbols instead of full seven-primary logs. Each entire column is compared to the primary column up to a single nonzero scalar. Prime tests, split roots, symbol orders, every pairwise product and inverse of the five named units are checked directly.

Ranks are obtained from exact integer minors by fraction-free Bareiss elimination, not the primary modular row-reduction procedure. All1585 subsets are examined. Circuit minimality is tested by deleting each column. The compatible circuits are tested against all2352 F49 directions, without using the42-direction reduction. Their containment independently reproduces the classification of all1024 rational supports. The coefficient bases differ, but the complete sets of zero patterns and ranks agree.

## Observed finite results

124 rational supports pass the criterion;900 leave a first-layer survivor. There are8 maximal passing rational sets and15 minimal surviving rational sets. There are23 compatible prime-ideal circuits:6 of size2,8 of size3,9 of size5. Natural Kummer rank is1 for6 circuits and2 for17 circuits.

The main H_* matrix passes all42 tests. Its seven columns have ranks checked on every relevant Kummer zero set; the complete matrix is NOT required to have independent columns. That distinction is essential: seven arbitrary columns in a five-row matrix cannot all be independent.

The independent program checks360 named symbol products/inverses and625 exact determinant benchmarks, besides the complete direction and subset classifications. Independent code is not an independent mathematical reviewer.

## Surviving certificates

The minimal two-place example97,293 has alpha=(0,1,1) and character vector (6,1). The full JSON also records a five-place circuit over rational primes 883,1373,1471,1667. Its natural Kummer rank2 forces alpha=(0,1,1) if that circuit occurs inside the actual maximal support. Neither fixture supplies full actual Frobenius data or a Frey solution.

## Local arithmetic audit

local_maximal_checks.py checks8000 nonrational matrix instances, including 6080 with maximal eta ramification but unramified beta. The natural upper-right valuation bound is tight in4000 cases. All1440 declared principal-unit power tests verify that the highest character layer is supported exactly at maximal inertia orders.

An independent direct integer calculation rechecks the Q491 negative control: nonmaximal eta inertia can coexist with nonzero local alpha. No universal statement is inferred merely from the finite sample counts.

## Eight small minors for the main H_* result

Use one-based column numbering in the seven-column H_* matrix in ARITHMETIC_CONSEQUENCES.md, and row numbering(phi,2,3,sqrt5,7). Every admissible Kummer zero set is contained in one of the eight sets below. Thus these eight nonzero minors, together with the42-test zero-set coverage, are a compact certificate for full diagonal triviality on H_*.

| Maximal zero columns | Selected rows | Integer determinant | Modulo7 |
|---|---|---|---|
| 1,2,4,6 | 1,2,3,5 | -66 | 4 |
| 1,3,4,5 | 1,2,3,5 | -18 | 3 |
| 2,3,4 | 1,2,3 | -5 | 2 |
| 2,4,5 | 1,2,3 | -5 | 2 |
| 2,4,7 | 1,2,3 | 12 | 5 |
| 3,4,6 | 1,2,3 | 5 | 5 |
| 3,4,7 | 1,2,3 | 9 | 2 |
| 4,5,6,7 | 1,2,3,5 | 6 | 6 |

The primary code computes these determinants by signed permutations. The independent implementation recomputes them using fraction-free elimination.