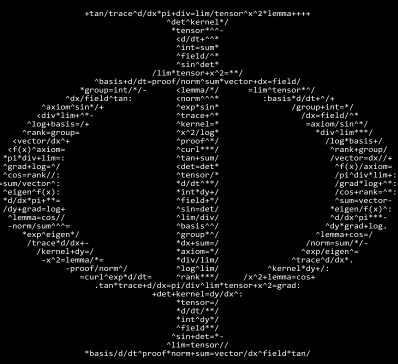
PHI TECHNICAL MONOGRAPHS

Analytic Models and Verification

Complex analysis, mathematical physics and proof specifications



PHI research collective

Abstract

Seventy-two independently formulated studies develop complex-analytic methods across approximation, algorithms, stochastic systems, propagation, operator theory and mathematical physics. Each study specifies a mathematical task, a conventional calculation, an analytic representation, a quantitative consequence and a boundary of applicability. The studies are organized by mathematical dependence and output type.

The analysis includes explicit remainder estimates, zero-free logarithms, separated-interaction compression, phase-based translation, implicit-function sensitivity, coupled propagation branches, finite-rank Fredholm models, canonical turning-point equations and asymptotic truncation. Scope-limited negative results are retained: missing regularity, inaccessible information, unsuitable constitutive assumptions and changed output requirements can prevent a proposed reduction.

Reported timings and numerical examples reproduce the recorded study data. They are distinguished from theorems concerning error, rank, memory and admissible competitors. New graphical rendering in this edition changes neither the underlying measurements nor their experimental scope. Conventional methods remain the reference where no analytical improvement has been demonstrated.

A register of 181 prospective formalization developments connects the mathematical cases to precise proof and implementation obligations. A final chapter develops finite certificate semantics, canonical ownership, obligation coverage, assumption propagation, exact normalization, interval and Bernstein enclosures, repeated-evaluation costs and graphical envelopes. Those new written arguments specify future verification interfaces; no new kernel acceptance is reported.

Reproduction material. The public source package contains the LaTeX manuscript and authored vector assets required to rebuild this volume. Its executable supplement covers selected numerical studies and certificate calculations. External textbook page reproductions and original externally acquired datasets are not distributed. References identify the scholarly works needed for independent source comparison.

Evidence convention. Established mathematical results retain their scholarly attribution. The exposition states the assumptions of each calculation and comparison. Recorded numerical results retain their original experimental scope. Formal specifications describe prospective proof developments unless an accepted proof object is explicitly identified. This public edition does not establish new proof-assistant kernel acceptance.

Contents

Abstract **i**

1 Approximation, singularities and boundary estimates **1**

1.1 MA01. Parameter-dependent rational integration 1
1.2 MA02. Exponential quadrature under a quantified analytic strip 2
1.3 MA03. Derivative outputs, directional information and the cost of a full gradient 2
1.4 MA04. An exact Chebyshev expansion and its geometric convergence parameter . 3
1.5 MA05. Padé time stepping and admissible poles 4
1.6 MA06. Analytic continuation and a minimax information obstruction 5
1.7 MA07. Parameter-uniform spectral exclusion via the resolvent 5
1.8 MA08. Contour projectors with a complete output specification 6
1.9 MA09. Binomial coefficients, saddle points and a rigorous relative envelope . . . 7
1.10 MA10. Zero-free logarithms and approximate counting 8
1.11 MA11. Several-variable coefficient extraction and anisotropic resolution 9
1.12 MA12. Characteristic functions when moment-generating functions fail 10
1.13 MA13. Contour displacement for oscillatory Gaussian integration 10
1.14 MA14. Elliptic integrals and an arithmetic–geometric mean solver 11
1.15 MA15. An exact Poisson model for elliptical torsion 12
1.16 MA16. From a potential certificate to an interior derivative certificate 13

2 Analytic structure in algorithms and computation **14**

2.1 CP01. Ordered ternary trees and Lagrange inversion 14
2.2 CP02. Marked automata and derivatives of a transfer matrix 15
2.3 CP03. Polynomial multiplication and a contract for FFT arithmetic 15
2.4 CP04. Closed walks, graph zeta data and the meaning of a cycle 16
2.5 CP05. Moment compression of separated Cauchy interactions 17
2.6 CP06. Graph heat propagation and a strong structured baseline 18
2.7 CP07. Image translation from Fourier phase correlation 19
2.8 CP08. Beltrami distortion and local geometric guarantees 20
2.9 CP09. Rational approximation with a pole-exclusion certificate 21
2.10 CP10. Gradient descent as a polynomial of a linear operator 22
2.11 CP11. Wirtinger differentiation of a complex least-squares objective 22
2.12 CP12. Holomorphic trial spaces and model error in PDE learning 23
2.13 CP13. Implicit neural layers and branch-dependent derivatives 23

2.14 CP14. Analytic kernels and finite-feature error 24
2.15 CP15. Quantum propagation and polynomial approximation on a real spectrum . 25
2.16 CP16. Partition functions on a path and the cost of an exact baseline 25

3 Stochastic models, estimation and service systems 27
3.1 ST01. Service-time variability in a single-server queue 27
3.2 ST02. Independent latency stages and repeated transform poles 28
3.3 ST03. Absorption, reachability probabilities and resolvent remainders 29
3.4 ST04. Distribution convolution and exact support accounting 29
3.5 ST05. Characteristic functions and bounded distributional features 30
3.6 ST06. Tail-latency mixtures and model identifiability 31
3.7 ST07. Optimal timeout allocation under a declared resource budget 32
3.8 ST08. Circular Wiener estimation and the distinction between batch and causal outputs 32
3.9 ST09. PageRank residuals, parameter sensitivity and sparse computation 34
3.10 ST10. Streaming transforms, cumulants and a zero-sensitive logarithm 34
3.11 ST11. Analytic kernels inside verified software workflows 35
3.12 ST12. Amortized verification and a measurable deployment threshold 36

4 Analytic response, propagation and stable realization 37
4.1 EL01. Coalescing circuit poles and a regular transient representation 37
4.2 EL02. A continuous-contour Nyquist certificate 38
4.3 EL03. A delay-independent exclusion of infinitely many characteristic roots . . . 38
4.4 EL04. Small gain as an analytic norm certificate 39
4.5 EL05. Schur recursion and stable bounded response parameterization 40
4.6 EL06. Interpolation feasibility from two measurements 40
4.7 EL07. Spectral factorization and magnitude ambiguity 41
4.8 EL08. Transmission-line square roots with a coupled branch contract 42
4.9 EL09. Array factors as polynomials on the unit circle 43
4.10 EL10. Dispersive Gaussian pulses and analytic propagation 44
4.11 EL11. OFDM orthogonality and carrier-frequency leakage 45
4.12 EL12. Bessel modes, entire series and cutoff calculations 45
4.13 EL13. Fractional impedance as a positive continuum of relaxations 46
4.14 EL14. Causality, Hilbert transforms and missing frequency bands 47
4.15 EL15. Conformal field reduction and the transformed wave equation 48
4.16 EL16. Hankel singular values in a separable model-reduction problem 48
4.17 EL17. A winding invariant and an edge-state recurrence 49
4.18 EL18. Switching between stable subsystems and Floquet growth 49
4.19 EL19. Finite-state verification of a rounded recurrence 50
4.20 EL20. Harmonic measure and optimal allocation of boundary observations 51

5 Operator factorization, inverse problems and singular limits 52
5.1 SP01. Wiener–Hopf factors and the boundary term of a unilateral shift 52
5.2 SP02. Contour moments and the conditioning of pole reconstruction 53
5.3 SP03. Discrete Cauchy–Riemann consistency on a square 53
5.4 SP04. Turning points and a uniform Airy representation 54
5.5 SP05. Analytic Fredholm behaviour in an exactly finite-rank integral operator . . 55
5.6 SP06. Entire-function coordinates for the harmonic oscillator 56
5.7 SP07. Optimal truncation of an asymptotic expansion 57
5.8 SP08. Jensen bounds as global zero-count information 58

6 Formal proof specifications 59
6.1 Source integrity and proof governance 59
6.2 Cantor coding, geometry and topology 62
6.3 Functional analysis and finite-vector foundations 65
6.4 Proper holomorphic maps and branched coverings 69
6.5 Normal families, mapping and distortion 72
6.6 Certified complex and numerical analysis 76
6.7 Discrete mathematics, algorithms and exact computation 81
6.8 Learning, kernels and statistical computation 84
6.9 Probability, information and wireless optimization 88
6.10 Signals, circuits and implementation-aware physics 91
6.11 Specialist analytic and geometric extensions 95
6.12 Software, hardware and distributed-state specifications 98
6.13 Efficiency, data, figures and publication verification 100

7 Finite certificate semantics and publication evidence 103
7.1 Versioned identity and canonical ownership 103
7.2 Questions, subparts and positive-weight coverage 104
7.3 Dependency closure and assumption propagation 105
7.4 Nonvacuity and an attaining optimum 105
7.5 Sound checkers and composition 106
7.6 Exact polynomial normalization 106
7.7 Interval enclosure and reciprocal domains 107
7.8 Bernstein coefficients and continuous-band certificates 108
7.9 Cost semantics and a break-even theorem 108
7.10 Certified graphical envelopes 109
7.11 Evidence paths for independent review 109

References 111

Chapter 1

Approximation, singularities and boundary estimates

1.1 MA01. Parameter-dependent rational integration

Analytical references: [1, 2].

Problem. For positive parameters a, b, evaluate

$$I(a,b) = \int_{-\infty}^{\infty} \frac{dx}{(x^2+a^2)(x^2+b^2)}.$$

The required output is a formula valid through the coalescing case $a = b$. A numerical routine must also evaluate that formula without cancellation when the parameters are close.

Conventional derivation. For $a \neq b$, partial fractions give

$$\frac{1}{(x^2+a^2)(x^2+b^2)} = \frac{1}{b^2-a^2}\left(\frac{1}{x^2+a^2} - \frac{1}{x^2+b^2}\right).$$

Integrating each term gives $\pi(1/a - 1/b)/(b^2-a^2)$. Algebraic cancellation produces $\pi/[ab(a+b)]$. Substituting nearly equal machine numbers into the uncancelled expression can magnify rounding error. The simplified expression has no such subtractive cancellation for positive a, b.

Complex derivation. Close the real line with an upper semicircle of radius $R > \max(a,b)$. The arc integral is $O(R^{-3})$. The enclosed simple poles are ia and ib, with residues $[2ia(b^2-a^2)]^{-1}$ and $[2ib(a^2-b^2)]^{-1}$. Their sum yields

$$\boxed{I(a,b) = \frac{\pi}{ab(a+b)}}.$$

At $a = b$, either compute the double-pole residue or take the limit using dominated convergence on a compact positive parameter interval. The result is $\pi/(2a^3)$.

Complexity and conditioning. Both symbolic constructions have constant arithmetic cost for this two-factor family. Residues offer systematic bookkeeping when the number of factors grows or when inverse-transform terms are required. They do not improve the constant-cost answer already obtained by partial fractions. The relative logarithmic derivatives of the simplified answer are $-1/a - 1/(a+b)$ and $-1/b - 1/(a+b)$; large sensitivity near zero is a property of the integral, independent of the algebra used to evaluate it.

Extension and solution. For the half-line integral divide the formula by two, since the integrand is even. For $a = 1, b = 2$, the full-line answer is $\pi/6$. If one parameter is zero, the ordinary integral diverges. An analytic continuation of the formula in parameters does not supply the missing improper-integral convergence.

1.2 MA02. Exponential quadrature under a quantified analytic strip

Analytical references: [3, 4].

Problem. Let g be 2π-periodic, holomorphic on a neighbourhood of $|\operatorname{Im} z| \le a$, and bounded there by M. Approximate $I = \int_0^{2\pi} g(t)\,dt$ using N equally spaced evaluations. The specified input includes the strip width and the bound; smooth real samples alone do not establish them.

Conventional baseline. Composite quadrature estimates based on a finite number of real derivatives produce algebraic error rates. Adaptive Gauss-type rules are strong practical controls. For a periodic analytic integrand, the periodic trapezoidal rule exploits additional structure while retaining a particularly simple sampling pattern.

Derivation. Write $g(t) = \sum_{k\in\mathbb{Z}} c_k e^{ikt}$. Shift the coefficient contour upward or downward according to the sign of k. Periodicity cancels the vertical sides and gives $|c_k| \le Me^{-a|k|}$. Root-of-unity orthogonality implies

$$Q_N = \frac{2\pi}{N}\sum_{j=0}^{N-1} g(2\pi j/N) = 2\pi \sum_{m\in\mathbb{Z}} c_{mN}.$$

Therefore

$$|Q_N - I| \le 4\pi M \frac{e^{-aN}}{1-e^{-aN}}.$$

A sufficient sample count for exact-arithmetic error ϵ is

$$N \ge a^{-1}\log(1 + 4\pi M/\epsilon).$$

If each evaluation has absolute error at most η, the quadrature contributes at most $2\pi\eta$ of evaluation error. Summation error requires an additional arithmetic bound.

Efficiency comparison. The guarantee uses $O(\log(1/\epsilon))$ evaluations for a fixed strip, whereas a fixed-order algebraic bound gives a power of $1/\epsilon$. This compares proved worst-case error models, not measured universal superiority over adaptive software. As singularities approach the integration line, a decreases and the evaluation count increases.

Extension and solution. For $g(t) = (1 - qe^{it})^{-1}$ with $0 < q < 1$, the exact average is one and the discrete average is $1/(1-q^N)$. Thus the error is exactly $q^N/(1-q^N)$ for the average. This example independently checks the aliasing mechanism and shows why periodicity and complex singularity distance determine the rate. An endpoint jump violates the analytic-periodic contract.

1.3 MA03. Derivative outputs, directional information and the cost of a full gradient

Analytical references: [5, 6, 7].

Problem. For $x \in \mathbb{R}^d$, compute all components of the gradient of

$$F(x) = \sum_{j=1}^{d} \sin x_j + \exp\left(d^{-1}\sum_{j=1}^{d} x_j\right).$$

An independent directional derivative check is a different output contract from a complete gradient.

Native solution. Set $m = d^{-1}\sum_j x_j$. One reduction and one vector evaluation give $\partial_j F = \cos x_j + e^m/d$. The cost is $\Theta(d)$ arithmetic operations and $\Theta(d)$ output storage. Reverse-mode differentiation of a compatible expression graph can attain the same order, with graph-dependent storage.

Analytic verification. For a real direction v, holomorphy and reality on real inputs give

$$\frac{\operatorname{Im} F(x+ihv)}{h} = DF(x)v - \frac{h^2}{6}D^3F(x)[v,v,v] + O(h^4).$$

One complex evaluation checks a directional derivative in $O(d)$ work. Repeating with every coordinate direction uses d function evaluations, each costing $O(d)$, hence $O(d^2)$ work. Centered differences similarly use $2d$ evaluations but also incur subtractive cancellation. Complex step requires a compatible complex implementation; magnitude, conjugation, clipping and data-dependent discrete branches can change the derivative being computed.

Measured comparison. The recorded experiment evaluates the analytic vector formula, coordinate complex steps and centered differences at dimensions 8, 32, 128 and 512. It uses the same complete gradient target for all methods. The vector formula is the strong baseline. Coordinate complex step is retained for independent checking, not asserted to be the fastest full-gradient construction. Its finite error is checked against the explicit derivative.

Extension and solution. For $v = (1,\dots,1)$, the directional derivative is $\sum_j \cos x_j + e^m$. A single directional check can pass despite a nonzero gradient error orthogonal to v. At least d linearly independent exact directions are needed to identify an arbitrary gradient in $\mathbb{R}^d$. Randomized checks have a probabilistic interpretation and require a declared error/noise model.

1.4 MA04. An exact Chebyshev expansion and its geometric convergence parameter

Analytical references: [5, 8].

Problem. Approximate $f_c(x) = 1/(c-x)$ uniformly on $[-1,1]$ for $c > 1$. Compare monomial Taylor expansion about zero, Chebyshev truncation and the known rational formula. The rational formula is the strongest baseline when c is explicitly given.

Derivation. Set $q = c - \sqrt{c^2-1} \in (0,1)$, so $c = (q+q^{-1})/2$. The geometric Poisson identity gives

$$\frac{1}{c-x} = \frac{1}{\sqrt{c^2-1}}\left(1 + 2\sum_{n\ge 1} q^n T_n(x)\right).$$

To verify it, substitute $x = \cos\theta$ and sum the two geometric series in $qe^{\pm i\theta}$. The constant factor follows from $\sqrt{c^2-1} = (q^{-1}-q)/2$. Since $|T_n(x)| \le 1$,

$$\|f_c - p_m\|_\infty \le \frac{2q^{m+1}}{\sqrt{c^2-1}(1-q)}.$$

The nearest Bernstein ellipse meeting the pole has parameter $\rho = q^{-1}$. The monomial expansion has tail at most $c^{-m-1}/(c-1)$ after degree m, obtained from the geometric series in x/c.

Complexity and representation. A Chebyshev polynomial can be evaluated in $O(m)$ operations by Clenshaw recurrence. The rational formula uses one subtraction and one division. The polynomial becomes relevant when a downstream hardware, approximation-space or certification contract requires a polynomial. The coefficient-decay theorem gives an explicit degree, independent of a training grid. As $c \downarrow 1$, both the function magnitude and required polynomial degree increase.

Extension and solution. For $c = 5/4$, $q = 1/2$ and $\sqrt{c^2 - 1} = 3/4$, so the displayed error bound is $(16/3)2^{-(m+1)}$. Degree $m = 13$ suffices for error below 10^{-3}. The derivative error has a different tail because $\|T_n'\|_\infty = n^2$; a function-value certificate cannot be reused unchanged for a gradient calculation.

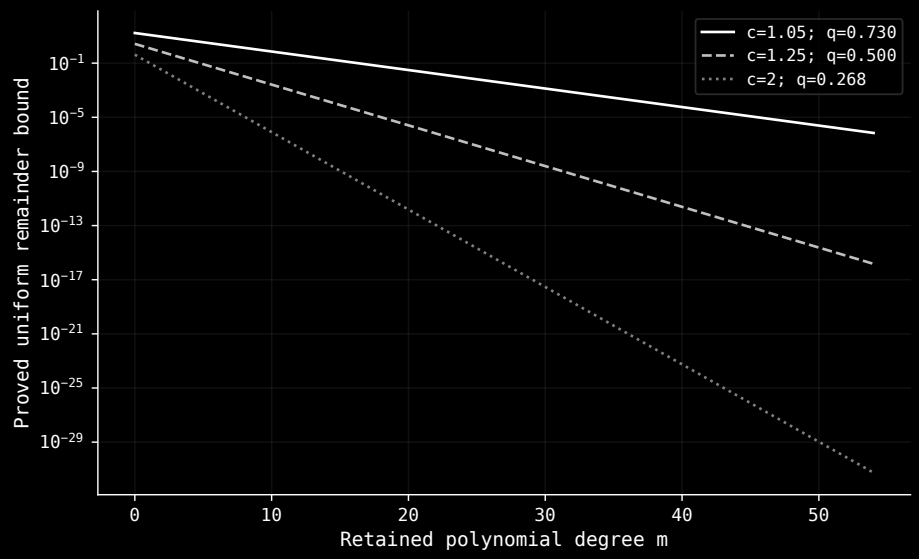


Figure 1.1: Complex-pole distance and the exact Chebyshev coefficient decay for the family $1/(c - x)$. The plotted upper envelopes follow the proved geometric tail; they are not fitted learning curves.

1.5 MA05. Padé time stepping and admissible poles

Analytical references: [5, 2].

Problem. Advance $y' = \lambda y$ by one step h with $z = h\lambda$. Forward Euler uses $1 + z$; the trapezoidal method uses

$$R(z) = \frac{1 + z/2}{1 - z/2}.$$

The exact factor is e^z. The rational formula is the [1/1] Padé approximant.

Local derivation. Expand the denominator geometrically near zero. One obtains $R(z) = 1 + z + z^2/2 + z^3/4 + \cdots$, so $R(z) - e^z = z^3/12 + O(z^4)$. For $\operatorname{Re} z \le 0$,

$$|1 - z/2|^2 - |1 + z/2|^2 = -2 \operatorname{Re} z \ge 0,$$

which proves $|R(z)| \le 1$. The pole $z = 2$ is outside the closed left half-plane. Forward Euler requires $|1 + z| \le 1$, a bounded disk.

Interpretation. The rational step is stable throughout the scalar left half-plane. However $R(z) \to -1$ along the negative real axis, so unresolved stiff modes can alternate without strong damping. Backward Euler, $1/(1 - z)$, tends to zero and offers a different accuracy/damping trade-off. Both rational methods require a solve for matrix problems; a matrix's spectral inclusion alone does not bound transient norm growth when the matrix is nonnormal.

Cost. For a fixed sparse matrix, factorization setup can be reused over many steps. The comparison cost is factorization plus triangular solves, not a scalar division. For a dense $d \times d$ matrix a direct factorization costs $O(d^3)$ and each subsequent right-hand side $O(d^2)$. A polynomial step avoids that setup but may require a much smaller step size.

Extension and solution. For $z = -10$, Euler gives -9, trapezoidal gives $-2/3$ and backward Euler gives $1/11$, while the exact factor is e^{-10}. Stability, transient fidelity and damping are three separate outputs. The exact reciprocal identity $R(z)R(-z) = 1$ explains the trapezoidal scheme's time-reversal property where neither factor has a pole.

1.6 MA06. Analytic continuation and a minimax information obstruction

Analytical references: [9, 8].

Problem. A holomorphic function is observed with uniform error ϵ on $|z| \le r$. Predict its value at $R > r$. Assume first that no bound on a larger domain is available.

Obstruction. The functions $f_n(z) = \epsilon(z/r)^n$ satisfy $\|f_n\|_{|z|\le r} \le \epsilon$, so they are observationally compatible with zero. Their values at R are $\epsilon(R/r)^n$, which are unbounded as $n \to \infty$. No algorithm can provide a finite uniform continuation bound over this unrestricted input class. Increasing the number of samples inside the observation disk does not remove this worst-case ambiguity.

Regularized class. Now impose $\|f\|_{|z|\le S} \le M$ for $r < R < S$. Hadamard's three-circles inequality yields

$$\|f\|_R \le \|f\|_r^{\theta} \|f\|_S^{1-\theta}, \qquad \theta = \frac{\log(S/R)}{\log(S/r)}.$$

It follows by applying the maximum principle to $\log|f(z)|$ minus the harmonic radial affine function of $\log|z|$ on the annulus. Zeros are handled by subharmonicity or a regularization of the logarithm. This proves a Hölder dependence $M^{1-\theta}\epsilon^{\theta}$ under a quantitative outer-domain hypothesis.

Efficiency criterion. A fast extrapolator that ignores the outer bound does not solve the regularized problem. The difficult input is the analytic-domain and growth certificate. Once it is supplied, degree selection, truncation and observation precision can be balanced. The exponent tends to zero as the target approaches the outer boundary, explaining deterioration even when uniqueness remains valid.

Extension and solution. Take $r = 1/4$, $R = 1/2$, $S = 1$. Then $\theta = 1/2$ and the bound is $\sqrt{M\epsilon}$. Two feasible functions consistent with the same data differ by at most 2ϵ on the inner disk and $2M$ on the outer disk; the corresponding pairwise continuation diameter is at most $2\sqrt{M\epsilon}$. The factors must be retained when translating a one-function theorem into an inverse-problem uncertainty set.

1.7 MA07. Parameter-uniform spectral exclusion via the resolvent

Analytical references: [5, 7].

Problem. Let A be a matrix with no eigenvalue on a positively oriented contour Γ. Certify that every $A + E$ in a norm-bounded perturbation set has the same number of eigenvalues inside Γ, counted with algebraic multiplicity.

Derivation. For $z \in \Gamma$,

$$zI - A - E = (zI - A)\left[I - (zI - A)^{-1}E\right].$$

If $\sup_{\Gamma} \|(zI - A)^{-1}\| \, \|E\| < 1$, the Neumann series makes the second factor invertible. The same strict inequality holds for tE, $0 \le t \le 1$. The determinant therefore never vanishes on the contour

during this homotopy, and its winding number is constant. The argument principle gives the desired invariant eigenvalue count.

Native baseline. For a small explicitly known characteristic polynomial, Routh, Schur or a direct exact root calculation can be simpler. Computing all eigenvalues at a finite grid of parameter values does not establish a continuous-family exclusion. The resolvent certificate is appropriate when a uniform operator bound is available and the parameter set is large.

Cost and conditioning. If A is normal, the resolvent norm is $1/\operatorname{dist}(z, \sigma(A))$ in the spectral norm. For a general matrix it can be much larger. A finite contour mesh requires a bound between nodes. One may use a Lipschitz bound from $R'(z) = -R(z)^2$ on certified resolvent disks; sampled singular values alone provide a diagnostic.

Extension and solution. Let $A = \operatorname{diag}(-2, 1)$ and let Γ be the circle centred at -2 of radius $1/2$. The nearest eigenvalue has distance $1/2$, so every perturbation with $\|E\|_2 < 1/2$ preserves exactly one enclosed eigenvalue. The strict inequality excludes contact with the contour. Replacing a normal matrix by a similar matrix with a badly conditioned eigenvector basis can destroy this margin without moving the spectrum.

1.8 MA08. Contour projectors with a complete output specification

Analytical references: [5, 10].

Problem. Compute the spectral subspace associated with eigenvalues in a prescribed contour. The target includes the contour, orientation and enclosed algebraic multiplicity. Merely returning an idempotent matrix does not identify the intended subspace.

Analytic representation. For a contour separated from the spectrum,

$$\Pi = \frac{1}{2\pi i} \int_\Gamma (zI - A)^{-1}\, dz.$$

Jordan-block expansion of the resolvent shows that the integral is the identity on blocks inside the contour and zero on the remaining blocks. Thus $\Pi^2 = \Pi$, $A\Pi = \Pi A$, and $\operatorname{tr} \Pi$ is the enclosed multiplicity. The formula avoids selecting a basis in a nearly degenerate eigenspace.

Numerical construction. Quadrature yields $\widehat{\Pi} = \sum_j w_j (z_j I - A)^{-1}$. A certified contour quadrature bound and certified linear solves yield an operator error η. When A is Hermitian, the true projector is orthogonal; a subspace distance can be deduced from a sufficiently small operator error. For nonnormal matrices the oblique projection may have norm much larger than one.

Comparison. Dense eigendecomposition and Schur methods are strong baselines for small matrices. Shifted solves are attractive for large sparse problems, repeated right-hand sides and parallel hardware. Their cost includes factorization at every distinct shift, communication and solve accuracy. Idempotence and commutator residuals alone cannot certify the selected contour: both zero and the identity satisfy those algebraic identities.

Extension and solution. For $A = \operatorname{diag}(-2, -1, 3)$ and the circle $|z + 3/2| = 1$, the integral equals $\operatorname{diag}(1, 1, 0)$. A trace close to two, a small idempotence residual and a small commutator become meaningful only after the contour-to-spectrum separation has been established. This gives a finite test design with an independently known correct subspace.

1.9 MA09. Binomial coefficients, saddle points and a rigorous relative envelope

Analytical references: [11, 12, 13, 14].

Problem. Evaluate $\mathbb{P}(X = k)$ for $X \sim \mathrm{Bin}(n, p)$ when n is large. Exact integer combinations, log-gamma evaluation and asymptotic formulas are distinct admissible outputs. This comparison seeks a logarithmic probability with an explicit approximation envelope.

Saddle derivation. For $x = k/n \in (0, 1)$, coefficient extraction from $(1 - p + pz)^n$ has a positive saddle $z_* = x(1 - p)/[p(1 - x)]$. Expanding the logarithm of the integrand near the saddle gives

$$S = \frac{e^{-nD(x\|p)}}{\sqrt{2\pi n x(1-x)}}, \quad D(x\|p) = x \log \frac{x}{p} + (1-x) \log \frac{1-x}{1-p}.$$

For the binomial family, an error envelope follows from the classical Robbins factorial inequality, explicitly recorded in R54, p. 1. This inequality is an imported theorem in the present monograph. Write

$$n! = \sqrt{2\pi n}(n/e)^n e^{\delta_n}, \qquad \frac{1}{12n+1} < \delta_n < \frac{1}{12n}.$$

Then $\log(\mathbb{P}(X = k)/S) = \delta_n - \delta_k - \delta_{n-k}$. Substituting the lower and upper bounds separately gives the explicit two-sided interval used in the experiments.

Baseline and cost. A library log-PMF is normally the most convenient floating baseline. Exact $\binom{n}{k}$ has a growing output bit length. The saddle expression has a constant number of elementary function evaluations at a fixed precision, but its approximation error grows near $k = 0$ or n unless a different boundary regime is used. No numerical underflow is necessary when the result is kept in logarithmic form.

Extension and solution. For $k = n/2$, the correction interval shrinks like $O(1/n)$. For fixed k with $n \to \infty$, the central saddle expression is not uniformly accurate and a Poisson or boundary expansion may be preferable. The recorded diagnostics check the logarithmic interval at twelve finite (n, k) pairs; the factorial inequalities supply the general proof.

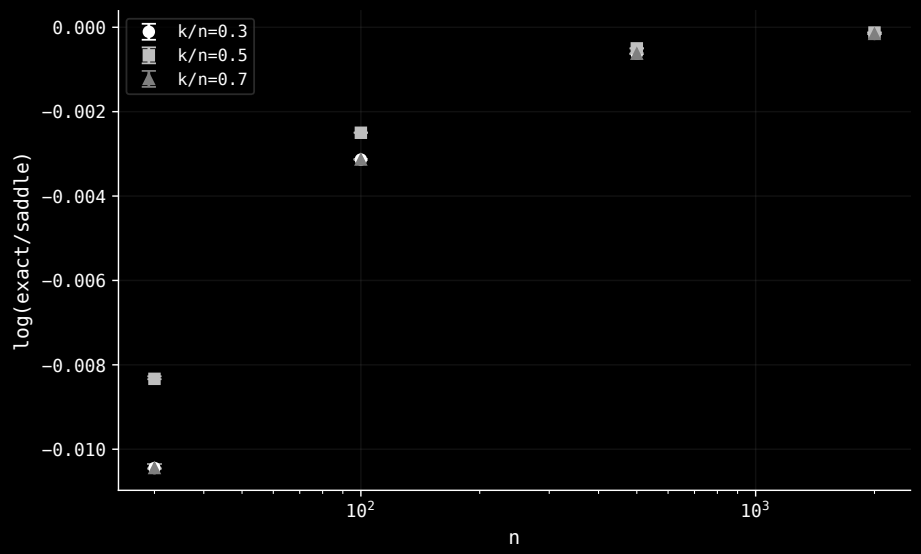


Figure 1.2: Logarithmic relative correction to the binomial saddle approximation and the factorial-remainder envelope. The values correspond to specified finite probabilities; the error bars are mathematical bounds.

1.10 MA10. Zero-free logarithms and approximate counting

Analytical references: [11].

Problem. Let $Z(0) = 1$ be a polynomial of degree d whose roots satisfy $|\zeta_j| \geq R$. Approximate $Z(\lambda)$ for $|\lambda| < R$ from finitely many coefficients of its analytic logarithm.

Derivation. Factor $Z(\lambda) = \prod_j (1 - \lambda/\zeta_j)$ and choose the branch of $\log Z$ with value zero at the origin. The power series gives

$$\log Z(\lambda) = -\sum_{k \geq 1} \frac{\lambda^k}{k} \sum_j \zeta_j^{-k}.$$

Writing $q = |\lambda|/R < 1$, the degree-m remainder is bounded by

$$\epsilon_m \leq \frac{dq^{m+1}}{(m+1)(1-q)}.$$

Exponentiation converts an additive logarithmic bound into relative error at most $e^{\epsilon_m} - 1$ for the complex value. A real positive partition function fixes the physically meaningful branch at the target.

Worked model. For independent sets of a three-vertex path, $Z(\lambda) = 1 + 3\lambda + \lambda^2$. The nearest root has modulus $(3 - \sqrt{5})/2 > 1/3$, so $R = 1/3$ is a simple certified radius. Reciprocal-root power sums satisfy $s_0 = 2$, $s_1 = -3$, $s_k = -3s_{k-1} - s_{k-2}$. Therefore the logarithmic coefficients can be computed by exact rational recurrence without numerically finding roots.

Complexity. A small path has a linear-time exact dynamic program, which is the stronger baseline. The zero-free formulation becomes algorithmically useful for a class where a short coefficient prefix is substantially cheaper than exact enumeration. The cost of computing that prefix remains part of the algorithm. A claimed short truncation degree does not make coefficient computation free.

Extension and solution. For $\lambda = 1/10$, $q = 3/10$; insert $m = 8$ to obtain an explicit rational logarithmic error. For $\lambda \geq 1/3$, this particular radius certificate no longer proves convergence. A different zero-free domain or exact calculation is then required. The experiment retains cases up to $\lambda = 1/4$ and records actual errors alongside the proved bounds.

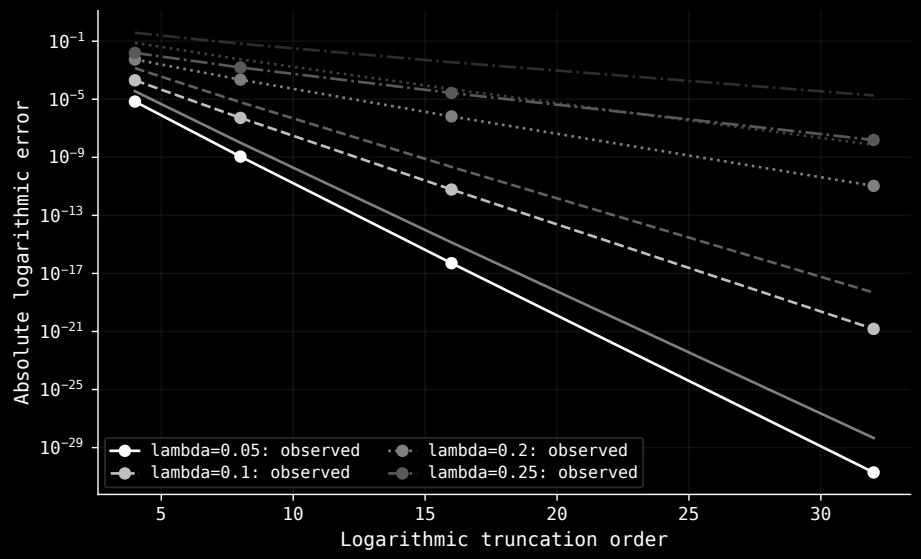


Figure 1.3: Exact-recurrence logarithmic approximations to $1 + 3\lambda + \lambda^2$ and the certified $R = 1/3$ tail. The deterioration as λ approaches the certified boundary is explicit.

1.11 MA11. Several-variable coefficient extraction and anisotropic resolution

Analytical references: [9].

Problem. Let f be holomorphic on a neighbourhood of the closed polydisk $|z_j| \le R_j$ and bounded there by M. Recover a finite multivariate Taylor table for evaluations on $|z_j| \le r_j < R_j$.

Derivation. Repeated Cauchy integration gives $|a_\alpha| \le M \prod_j R_j^{-\alpha_j}$. A tensor Fourier rule extracts coefficient congruence classes in each coordinate. For truncation degrees m_j, let $q_j = r_j/R_j$. Summing the coefficient tail over the union of events $\alpha_j > m_j$ gives

$$\|f - p\| \le M \prod_{j=1}^{d} (1 - q_j)^{-1} \sum_{j=1}^{d} q_j^{m_j+1}.$$

The union bound may overcount multi-index corners but remains valid. It directly supports unequal degrees when different complex singularity distances give different q_j.

Comparison. Repeated real finite differentiation becomes unstable for high orders and requires mixed-derivative bookkeeping. Tensor Cauchy sampling computes many coefficients from one structured evaluation array. Its sample count is still $\prod_j N_j$; analyticity alone does not eliminate exponential dependence on the number of parameters. Low-rank tensor or sparse-index methods need additional summability or structural assumptions.

Input model. A multivariate real loss containing conjugation, absolute values, sorting or thresholds does not automatically extend holomorphically. The theorem applies to an explicitly analytic parameter map. A complexification of physical variables must retain its relation to the real model and any conjugate constraints.

Extension and solution. For two variables with $q_1 = 1/2$, $q_2 = 1/8$ and equal M, the second direction needs about one third as many Taylor orders to attain the same geometric tail. The total product cost must nevertheless be counted. A rank-one separable function can be much cheaper, but that gain derives from separability; the general polydisk bound does not prove it.

1.12 MA12. Characteristic functions when moment-generating functions fail

Analytical references: [12, 15].

Problem. Represent a Cauchy-distributed random variable with density $p(x) = 1/[\pi(1+x^2)]$. A moment-generating function does not exist on any open real interval about zero, so an MGF-based cumulant calculation is inadmissible.

Contour calculation. The characteristic function is

$$\varphi(t) = \int_{\mathbb{R}} e^{itx} p(x)\, dx = e^{-|t|}.$$

For $t > 0$, close the contour in the upper half-plane and use the pole at i; for $t < 0$, close below with the corresponding orientation. At $t = 0$ the integral is one. The exponential factor determines the decaying half-plane, and the rational denominator controls the closing arcs.

Consequence. For independent Cauchy variables, the characteristic function of the sample mean is $[\varphi(t/n)]^n = e^{-|t|}$. The mean has the same Cauchy law for every n. A variance-based concentration claim is therefore false. The exact transform reveals this immediately, without a Monte Carlo convergence illusion.

Computational comparison. Empirical characteristic functions are defined for every sample because $|e^{itx}| = 1$. At a fixed frequency, each real and imaginary sample component is bounded, so finite-sample concentration can be established without moments of X. Uniform frequency-band claims require an additional covering or regularity argument. Fourier inversion also requires truncation and quadrature control; finitely sampled transforms do not uniquely identify an arbitrary distribution.

Extension and solution. For an independent sum of a Cauchy variable and a centred Gaussian with variance σ^2, the transform is $e^{-|t|-\sigma^2 t^2/2}$. Its cusp at zero remains, so the Gaussian convolution does not create a finite first absolute moment. This separates smoothness of the density from existence of moments and analytic continuation of the transform.

1.13 MA13. Contour displacement for oscillatory Gaussian integration

Analytical references: [1, 13].

Problem. Compute $J(\omega) = \int_{\mathbb{R}} e^{-x^2} e^{i\omega x}\, dx$. For large $|\omega|$, the result is much smaller than the absolute integral of the integrand. Absolute and relative error targets have very different difficulty.

Analytic deformation. Complete the square:

$$-z^2 + i\omega z = -(z - i\omega/2)^2 - \omega^2/4.$$

The integrand is entire. On a rectangular contour connecting the real axis to $\operatorname{Im} z = \omega/2$, the vertical-side contributions vanish as the horizontal extent tends to infinity. Shifting to that line gives

$$\boxed{J(\omega) = \sqrt{\pi}\, e^{-\omega^2/4}}.$$

The complex deformation converts a highly cancelling real-line integral into a nonoscillatory Gaussian times a known scale factor.

Baseline. A quadrature routine with an oscillatory cosine weight is a stronger numerical control than an unadapted real-grid sum. The closed formula is best for this exact family. For nearby

analytic amplitudes, a saddle-adapted contour can retain the advantage, provided that crossed poles, endpoint contributions and decay sectors are checked.

Error model. If the shifted amplitude is integrated to absolute error δ, the final error is $e^{-\omega^2/4}\delta$. Direct real quadrature with absolute tolerance 10^{-12} need not give any relative information once $J(\omega)$ is much smaller. Underflow in the final exponential can be avoided by returning $\log J = \frac{1}{2}\log\pi - \omega^2/4$ when the output contract permits it.

Extension and solution. Differentiating under the integral gives $\int xe^{-x^2}e^{i\omega x}dx = (i\omega/2)J(\omega)$. Domination by $|x|e^{-x^2}$ justifies the differentiation. This provides a second exact diagnostic for a numerical contour implementation and checks the sign of the Fourier convention.

1.14 MA14. Elliptic integrals and an arithmetic–geometric mean solver

Analytical references: [16, 17].

Problem. For an undamped pendulum with amplitude $\theta_0 \in (0, \pi)$, compute

$$T = 4\sqrt{L/g}\,K(k), \qquad k = \sin(\theta_0/2), \quad K(k) = \int_0^{\pi/2} (1 - k^2\sin^2\theta)^{-1/2}\,d\theta.$$

The dimensionless task is $K(k)$. A small-amplitude period is an asymptotic approximation, not the exact output.

Algebraic-geometric representation. The energy integral lies on the curve $v^2 = (1-u^2)(1-k^2u^2)$. Its branch-point collision as $k \uparrow 1$ explains the logarithmic period divergence. A quadratic transformation of the elliptic integral yields

$$K(k) = \frac{\pi}{2\,\mathrm{AGM}(1, k')}, \qquad k' = \sqrt{1-k^2}.$$

The AGM iterates $a_{j+1} = (a_j + b_j)/2$, $b_{j+1} = \sqrt{a_j b_j}$. Positivity gives $b_j \le \mathrm{AGM} \le a_j$, so $\pi/(2a_j) \le K \le \pi/(2b_j)$ in exact arithmetic. The quadratic transformation identity is recorded in R49, equations 19.8.1–19.8.5.

Convergence. The difference satisfies

$$a_{j+1} - b_{j+1} = \frac{(a_j - b_j)^2}{2(\sqrt{a_j} + \sqrt{b_j})^2}.$$

After the iterates become comparable, the number of iterations needed for p correct digits grows logarithmically in p. Finite-precision square roots require outward rounding to turn the exact interval into a machine enclosure.

Comparison. The experiment compares quadrature, a transparent AGM implementation, and `ellipkm1`, which accepts the complementary parameter and avoids cancellation near $k = 1$. All evaluate the same K. A robust library routine is retained as a conventional control; the AGM is not claimed to outperform it. Initial data use k' directly, so forming $1 - k^2$ cannot silently discard the small parameter.

Extension and solution. As $k' \downarrow 0$, $K(k) \sim \log(4/k')$. For $k' = 10^{-5}$ this asymptotic is an informative control, while $\pi/2$ is not. Forcing or damping invalidates the conserved-energy derivation of a single amplitude-dependent period.

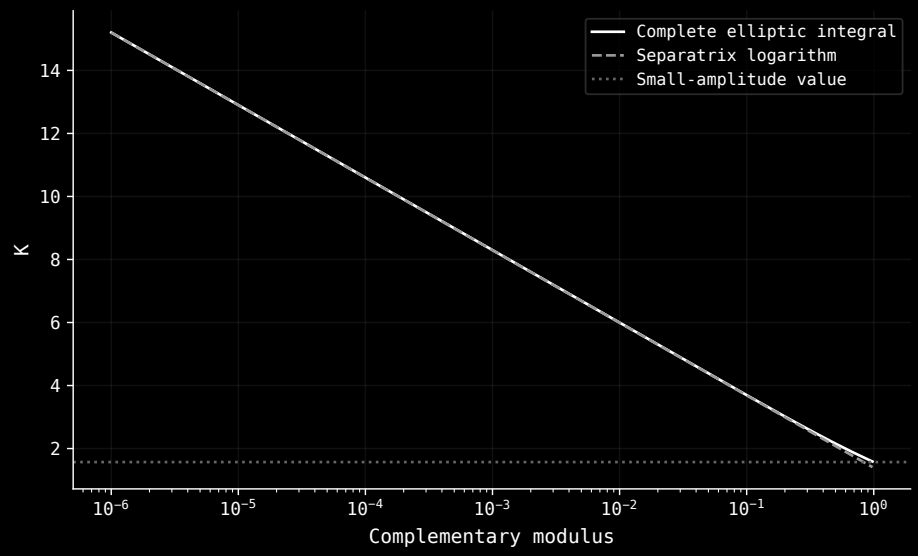


Figure 1.4: Pendulum elliptic integral as the complementary modulus decreases. Exact special-function values, the separatrix approximation and the small-amplitude constant answer different parameter regimes.

1.15 MA15. An exact Poisson model for elliptical torsion

Analytical references: [2].

Problem. On the ellipse $x^2/a^2 + y^2/b^2 < 1$, solve $-\Delta u = 2$ with $u = 0$ on the boundary. This is the dimensionless scalar stress-function model; elastic constants and engineering units must be supplied before interpreting u physically.

Native construction. A quadratic ansatz $u = C(1 - x^2/a^2 - y^2/b^2)$ satisfies the boundary condition. Substitution gives $2C(a^{-2} + b^{-2}) = 2$, hence

$$u(x, y) = \frac{a^2 b^2}{a^2 + b^2} \left(1 - \frac{x^2}{a^2} - \frac{y^2}{b^2}\right).$$

Uniqueness follows from the maximum principle applied to the harmonic difference of two solutions.

Complex representation. Write $u = -(x^2 + y^2)/2 + \operatorname{Re} F(z) + c$, where the first term has Laplacian -2 and F is holomorphic. A quadratic F captures the anisotropic harmonic correction. This decomposes a Poisson problem into a known particular solution and a harmonic boundary problem, a useful pattern for boundary approximation.

Comparison. For this ellipse, the polynomial ansatz already gives a constant-size exact solution. A conformal map or finite-element mesh would add setup cost. For a perturbed cross-section, a particular-solution-plus-harmonic representation can reduce the unknown to boundary coefficients, but the altered boundary must be controlled over its entire extent. The method applies to the scalar model; arbitrary vector elasticity requires its own constitutive and boundary equations.

Extension and solution. The integral is $\int_\Omega u \, dA = \pi a^3 b^3/[2(a^2 + b^2)]$, obtained by the change $x = ar\cos\theta$, $y = br\sin\theta$. This exact integral supplies a verification quantity for a mesh-based solution. It also demonstrates that potential and integrated torque-related quantities have different units and normalization factors.

1.16 MA16. From a potential certificate to an interior derivative certificate

Analytical references: [1, 18].

Problem. Suppose u and v are harmonic on a bounded planar domain and continuous on its closure, with $|u - v| \le \epsilon$ on the complete boundary. Bound both the potential error and its gradient at an interior point.

Potential bound. The maximum principle gives $|e| \le \epsilon$ throughout the domain, where $e = u - v$. Let the closed disk $\overline{D(z, \delta)}$ lie inside the domain. Differentiate the Poisson representation at its centre. For any unit direction ν,

$$\partial_\nu e(z) = \frac{1}{\pi\delta} \int_0^{2\pi} e(z + \delta e^{it}) \cos(t - \arg \nu)\, dt.$$

It follows that $|\nabla e(z)| \le 4\epsilon/(\pi\delta)$. The direction attaining the gradient converts the directional inequality into the vector bound. A simpler upper bound is $2\epsilon/\delta$.

Why the distance appears. The family $\epsilon r^n \cos(n\theta)$ has potential bounded by ϵ in the unit disk but boundary gradients of magnitude $n\epsilon$. Uniform potential control cannot give a boundary gradient bound independent of frequency. On every fixed smaller disk, the factor nr^{n-1} is uniformly bounded, consistently with the interior estimate.

Cost and implementation. An analytic or rational harmonic ansatz can certify the PDE exactly when all poles lie outside the domain. Bounding the boundary residual is still a continuum task. A mesh of residual samples becomes a certificate only with a between-sample derivative or interval enclosure. If v has a nonzero Poisson residual, a particular-solution estimate and extra regularity enter the gradient analysis.

Extension and solution. At an interior distance $\delta = 1/10$ and boundary error 10^{-4}, the proved gradient error is at most $4 \times 10^{-3}/\pi$. Reporting the original 10^{-4} as a field-gradient tolerance would be dimensionally and mathematically incorrect. This is a direct bridge from a mathematical verification quantity to an engineering observable.

Chapter 2

Analytic structure in algorithms and computation

2.1 CP01. Ordered ternary trees and Lagrange inversion

Analytical references: [11].

Problem. Count rooted ordered trees whose internal vertices have exactly three children, with n internal vertices. Let $T(z)$ include the single-leaf tree. The decomposition gives $T = 1 + zT^3$. Direct generation produces exponentially many objects; the required output is only their number.

Coefficient derivation. Put $Y = T - 1$, so $Y = z(1 + Y)^3$. Lagrange inversion gives, for $n \geq 1$,

$$[z^n]Y = \frac{1}{n}[u^{n-1}](1 + u)^{3n} = \frac{1}{2n + 1}\binom{3n}{n}.$$

The first coefficients of T are $1, 1, 3, 12, 55, 273$. The formula can also be proved by a cyclic-word argument, providing an independent combinatorial route. The inverse-function relation has a critical point when $1 = 3zT^2$; combining it with $T = 1 + zT^3$ gives $T_* = 3/2$, $z_* = 4/27$. A local quadratic expansion explains the square-root singularity and the $n^{-3/2}(27/4)^n$ coefficient scale.

Baseline and complexity. A coefficient recurrence using three convolutions is polynomial-time and produces the whole prefix. The binomial formula is a stronger baseline for a single coefficient. Exact integers have $\Theta(n)$ bits, so a constant number of symbolic factorial symbols is not constant bit cost. Fast factorial or product algorithms and their multiplication costs must be counted.

Application boundary. A grammar producing the same word by multiple trees has ambiguity. Its tree generating function counts derivations, not distinct accepted strings. This distinction matters in parser analysis and probabilistic syntax models. Lagrange inversion does not decide language equivalence or remove ambiguity.

Extension and solution. For ordered full m-ary trees the same argument gives $[z^n]T = ((m - 1)n + 1)^{-1}\binom{mn}{n}$. For $m = 2$ it reduces to the Catalan sequence. The formula assumes ordered child positions; unordered trees require symmetry analysis and generally a different generating equation.

2.2 CP02. Marked automata and derivatives of a transfer matrix

Analytical references: [11, 19].

Problem. Among binary words of length n containing no adjacent ones, count words by the number of ones. Let $F_n(u)$ be the generating polynomial. A word ends in zero after any admissible prefix, or in one after a prefix ending in zero. Thus $F_n(u) = F_{n-1}(u) + uF_{n-2}(u)$ with $F_0 = 1$ and $F_1 = 1 + u$.

Matrix formulation. With a state for the last symbol,

$$A(u) = \begin{pmatrix} 1 & u \\ 1 & 0 \end{pmatrix}, \qquad \sum_{n \geq 0} F_n(u) z^n = \frac{1 + uz}{1 - z - uz^2}.$$

At $u = 1$, $F_n(1)$ is the total count. The expected number of ones under the uniform measure on admissible words is $F_n'(1)/F_n(1)$. A second derivative gives the factorial second moment and hence the variance.

Efficient differentiation. Differentiate the recurrence jointly with its value. The derivative satisfies $G_n = G_{n-1} + F_{n-2} + uG_{n-2}$, allowing exact integer evaluation at $u = 1$ in linear recurrence steps. For a single large n, powering the block matrix

$$\begin{pmatrix} A & A' \\ 0 & A \end{pmatrix}$$

retains both powers and their derivative in $O(\log n)$ matrix multiplications. This is finite automatic differentiation of an analytic parameter representation.

Worked result. At $n = 3$, $F_3(u) = 1 + 3u + u^2$. There are five words and the mean number of ones is one. The variance is $2/5$. Direct enumeration gives the same values and serves as an exact small-instance check.

Limits. The weights must encode the intended probability law. A uniform measure on states is not a uniform measure on accepted words. A spectral asymptotic requires separation of the dominant eigenvalue and control of periodicity. The finite recurrence is valid without those asymptotic assumptions.

2.3 CP03. Polynomial multiplication and a contract for FFT arithmetic

Analytical references: [20, 21, 22].

Problem. Multiply two coefficient arrays of length n, returning all $2n - 1$ coefficients. Linear convolution and length-n circular convolution are different outputs. The comparison uses equal full linear outputs with real floating-point coefficients.

Native and Fourier constructions. Direct convolution uses n^2 scalar products. A discrete Fourier transform diagonalizes circular convolution. Padding both arrays to length $L \geq 2n - 1$ prevents wraparound, yielding

$$c = \mathrm{IDFT}_L(\mathrm{DFT}_L(a) \odot \mathrm{DFT}_L(b)).$$

Fast transforms give $O(L \log L)$ arithmetic operations and $O(L)$ storage, with padding, transforms and output extraction included. The recorded experiment compares direct convolution, FFT convolution and overlap-add at four sizes.

Exact integer contract. For integer inputs, a floating FFT followed by rounding is exact only if the total coefficient error is proved below $1/2$. A finite-field number-theoretic transform with enough compatible primes and Chinese remaindering can provide exact arithmetic. If $|a_j|, |b_j| \le B$, then $|c_k| \le nB^2$; a modulus product greater than $2nB^2$ permits unique signed reconstruction. The existence of the required roots of unity and the reconstruction cost are part of that method.

Measured result. For the recorded $n = 4096$ real arrays, the median direct-convolution time is about 2.254 ms; FFT convolution takes about 0.150 ms. The maximum output difference is approximately 1.21×10^{-13}. This is a floating diagnostic on a synthetic input, not a uniform roundoff theorem. At small sizes setup can dominate.

Extension and solution. For $a = (1, 2)$ and $b = (3, 4)$, linear convolution is $(3, 10, 8)$. A transform of length two returns $(11, 10)$ after wraparound. This exact counterexample detects missing padding independently of numerical precision. Fourier acceleration is established mathematics; the monograph's contribution is the explicit output and verification comparison.

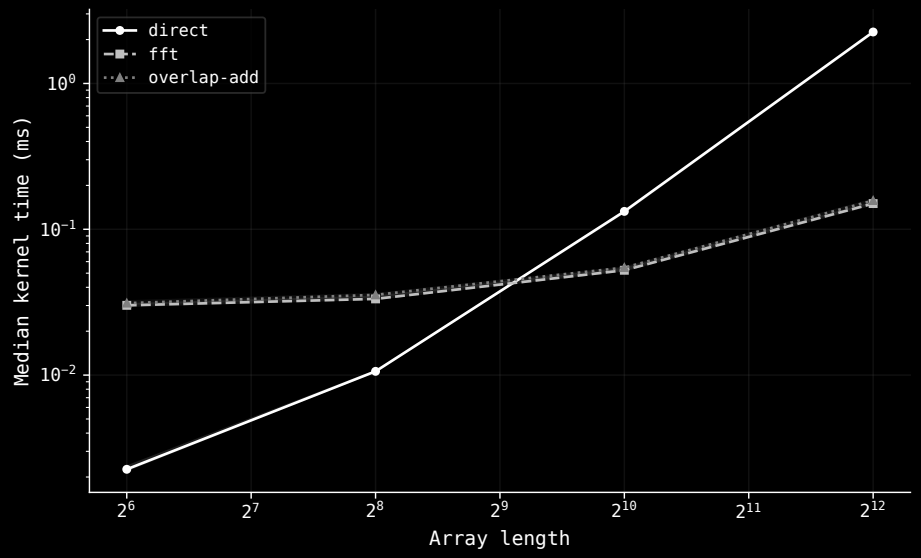


Figure 2.1: Recorded full linear-convolution timings. Input generation is excluded; padding and transform work are included. The overlap-add implementation is retained as an additional conventional control.

2.4 CP04. Closed walks, graph zeta data and the meaning of a cycle

Analytical references: [11, 20].

Problem. For a directed graph with adjacency matrix A, relate the counts of closed walks to an analytic invariant. Direct graph algorithms remain the baseline for connectivity or the existence of a directed cycle.

Analytic identity. For $|z| < 1/\rho(A)$,

$$-\log\det(I - zA) = \sum_{k \ge 1} \frac{z^k}{k} \operatorname{tr}(A^k).$$

One proof expands $-\log(1 - z\lambda_j)$ for each eigenvalue with algebraic multiplicity. Another differentiates the determinant logarithm and uses the resolvent geometric series, then fixes the value at zero. The trace counts closed walks with a marked starting vertex.

Nonbacktracking refinement. For an undirected graph, construct a matrix on oriented edges that permits $(u, v) \to (v, w)$ only when $w \neq u$. Traces of its powers exclude immediate reversal. This representation distinguishes actual circulation from the two-step backtracking present on every undirected edge. It enlarges the state space from vertices to oriented edges.

Cost and comparison. A depth-first search detects a cycle in $O(V + E)$ time. A determinant or many trace powers would be excessive for that Boolean task. The analytic representation becomes relevant when one requests weighted closed-walk generating data, spectral surrogates or a collection of length-dependent statistics. For a sparse graph, materializing dense powers can destroy the intended efficiency.

Extension and solution. The graph consisting of one undirected edge has adjacency eigenvalues ± 1, so $\operatorname{tr}(A^2) = 2$. Its nonbacktracking matrix is zero, so all positive-length nonbacktracking closed-walk counts vanish. The example prevents an incorrect inference of a geometric cycle from $\operatorname{tr}(A^2) > 0$.

2.5 CP05. Moment compression of separated Cauchy interactions

Analytical references: [5].

Problem. Evaluate $F(z) = \sum_{j=1}^{m} w_j/(z - a_j)$ at n targets, where $|a_j| \leq r < R \leq |z|$. The input includes a proved geometric separation and $Q = \sum_j |w_j|$.

Derivation. The finite geometric identity is

$$\frac{1}{z-a} = \sum_{k=0}^{p-1} \frac{a^k}{z^{k+1}} + \frac{a^p}{z^p(z-a)}.$$

Define moments $M_k = \sum_j w_j a_j^k$. Their truncated evaluation gives

$$\left| F(z) - \sum_{k=0}^{p-1} \frac{M_k}{z^{k+1}} \right| \leq \frac{Q}{R-r} \left(\frac{r}{R}\right)^p.$$

The moments cost $O(mp)$ arithmetic operations to construct and the target evaluations cost $O(np)$. Direct all-pairs evaluation costs $O(mn)$. For fixed separation and tolerance, $p = O(\log(1/\epsilon))$.

Matched experiment. The source cluster is inside radius 0.2, all targets have radius at least one, and weights sum in absolute value to one. At $p = 14$, the analytic truncation bound is 2.048×10^{-10}. With 4096 sources and targets, blocked direct evaluation takes approximately 130.934 ms and moment construction plus all evaluations approximately 0.287 ms. The measured maximum error is about 4.12×10^{-13}. Floating roundoff is a separate contribution; the displayed bound is for exact arithmetic.

Failure boundary. As a target approaches the source cluster, the geometric factor approaches one and the degree becomes large. A practical fast multipole method introduces a hierarchical near/far partition and local translations; none of those implementation costs is hidden inside the one-cluster claim. Cancellation in signed weights can make relative error large even when the absolute bound is small.

Extension and solution. For $r/R = 1/2$, $Q = 1$, $R = 1$ and tolerance 10^{-6}, $2^{1-p} \leq 10^{-6}$ requires $p \geq 21$. The same target error at separation $r/R = 1/5$ requires far fewer terms. Geometry is therefore part of the algorithm's complexity input.

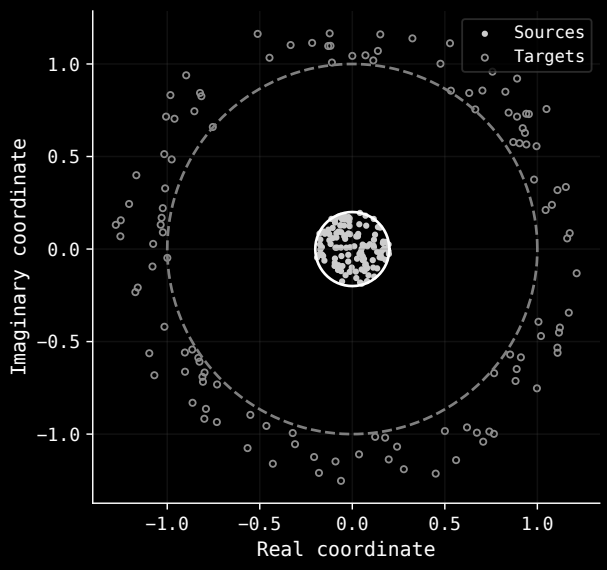


Figure 2.2: Separated Cauchy interaction: source and target geometry. The diagram represents the benchmark separation contract; the truncation estimate is derived in the text.

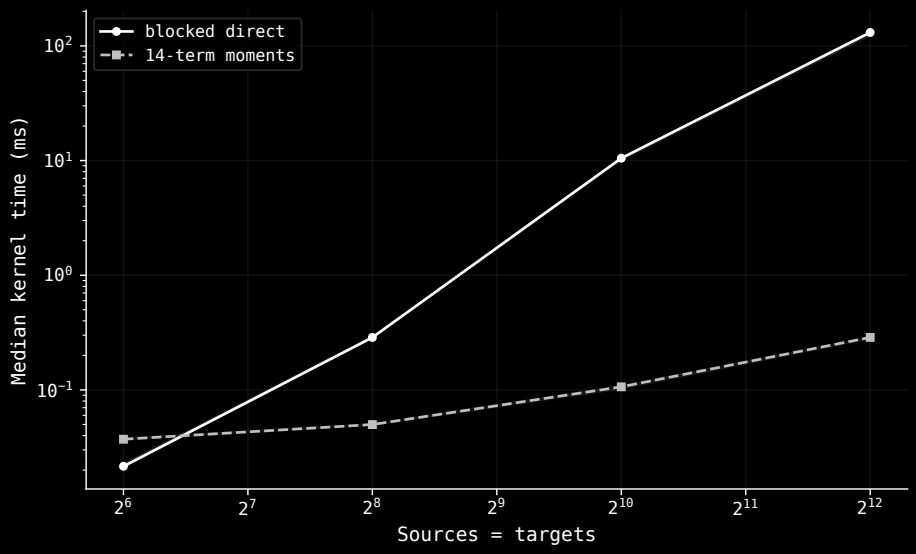


Figure 2.3: Recorded one-cluster Cauchy-kernel timings including moment construction. This experiment does not measure a complete hierarchical fast multipole algorithm.

2.6 CP06. Graph heat propagation and a strong structured baseline

Analytical references: [5, 10].

Problem. Compute $u(t) = e^{-tL}u_0$ for the Laplacian of an n-vertex cycle. The eigenvalues are $\lambda_k = 2 - 2\cos(2\pi k/n)$ and the eigenvectors are Fourier modes.

Structured solution. The exact finite-dimensional solution is

$$u(t) = \mathrm{IDFT}\left(e^{-t\lambda_k}\,\mathrm{DFT}(u_0)_k\right).$$

It costs $O(n \log n)$ arithmetic operations. A sparse matrix-exponential action and a dense scaling-and-squaring exponential are independent conventional controls. For a general sparse graph, the cycle Fourier formula is unavailable, while Krylov or polynomial methods remain applicable.

Polynomial certificate. For a symmetric Laplacian with spectrum in $[0, \Lambda]$, approximate $e^{-t\lambda}$ by p_m on that interval. The spectral theorem gives

$$\|(e^{-tL} - p_m(L))u_0\|_2 \le \max_{[0,\Lambda]} |e^{-t\lambda} - p_m(\lambda)| \, \|u_0\|_2.$$

Chebyshev recurrence evaluates the action in $O(m \operatorname{nnz}(L))$ work and $O(n)$ vector storage. Complex ellipses bound the scalar coefficients; this is a genuine analytic certificate for a sparse polynomial algorithm.

Measurement. At $n = 2048$, $t = 0.7$, the FFT action takes about 0.070 ms and sparse `expm_multiply` about 0.756 ms in the recorded environment. Both use the same preprocessed graph and input vector; their maximum difference is below 10^{-15}. Preprocessing is stated separately. The periodic FFT is retained as the strongest control for this special graph.

Extension and solution. For the constant vector, $Lu_0 = 0$ and the solution is unchanged. A numerical solver that decays the constant mode violates mass preservation. For a directed graph operator, the scalar spectral error bound above generally fails; a resolvent or an appropriate contraction structure is required.

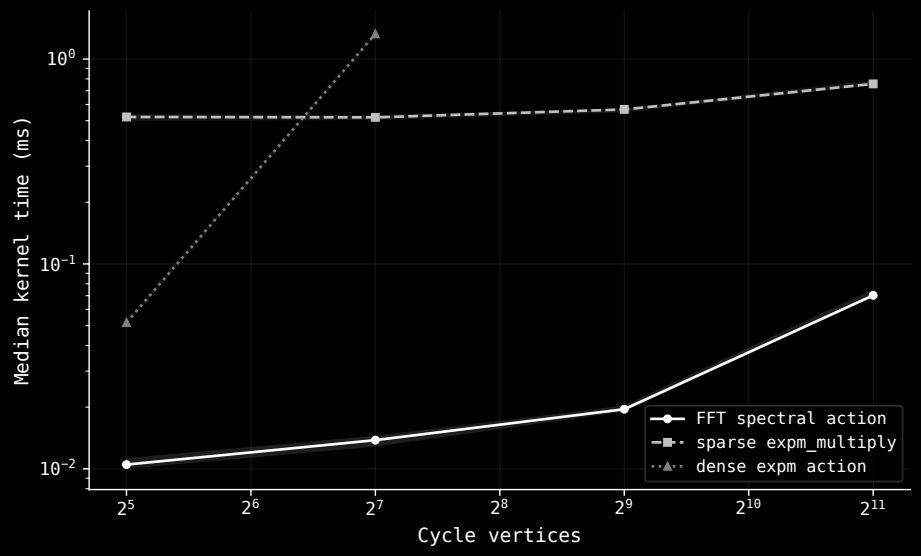


Figure 2.4: Recorded periodic-graph heat-action timings against sparse and dense general matrix-function methods. The periodic geometry is an explicit algorithmic assumption.

2.7 CP07. Image translation from Fourier phase correlation

Analytical references: [23, 21].

Problem. Let $v_j = u_{j-s}$ on a finite cyclic grid. Recover the integer shift s. Direct correlation evaluates every shift and can cost $O(n^2)$ in one dimension; FFT correlation reduces that to $O(n \log n)$.

Phase derivation. With $\widehat{u}_k = \sum_j u_j e^{-2\pi i j k/n}$,

$$\widehat{v}_k = e^{-2\pi i k s/n} \widehat{u}_k.$$

Where $\widehat{u}_k \neq 0$, the normalized cross-spectrum $\widehat{v}_k \overline{\widehat{u}_k} / |\widehat{v}_k \overline{\widehat{u}_k}|$ equals the phase ramp. Its inverse transform is a translated point mass if all required frequencies are available. Two-dimensional translations use the product Fourier group in the same way.

Identifiability. If u is periodic with a smaller period, multiple shifts are identical. If only a subset of frequencies is nonzero, the recovered shift may be determined only modulo a subgroup. Dividing by small spectral magnitudes also amplifies noise. A thresholded or weighted cross-spectrum changes the exact impulse theorem and requires a new error analysis.

Complexity and metric. Compare identical boundary conventions and outputs: periodic integer translation, subpixel interpolation, and a finite image with occlusion are distinct problems. The certificate consists of reconstructing the shifted input and bounding its residual under the declared noise model. An apparent sharp correlation peak alone does not prove the physical registration.

Extension and solution. For a one-dimensional signal supported only at even Fourier indices, shifts s and $s + n/2$ produce the same observed phase when n is even. This exact ambiguity is a Fourier-group property. Adding independent odd-frequency information removes that particular ambiguity, provided the corresponding coefficients are nonzero.

2.8 CP08. Beltrami distortion and local geometric guarantees

Analytical references: [23].

Problem. Study the real-linear planar map $f(z) = az + b\overline{z}$, $a, b \in \mathbb{C}$, with $|a| > |b|$. It is generally nonholomorphic but has a precise quasiconformal distortion.

Derivation. Rotations of source and target reduce the differential to principal stretches $|a| + |b|$ and $|a| - |b|$. The Jacobian is $J = |a|^2 - |b|^2 > 0$. The Beltrami coefficient is $\mu = f_{\overline{z}}/f_z = b/a$, with

$$K = \frac{1 + |\mu|}{1 - |\mu|}.$$

Thus the maximum aspect-ratio distortion of an infinitesimal circle is explicit. The conformal case is $b = 0$.

Comparison. A generic coordinate warp can be assessed by singular values of its Jacobian. The complex representation packages the two-dimensional anisotropy into one coefficient and supports composition and global PDE formulations. It does not lower the cost of checking an arbitrary image pipeline automatically. On a mesh, local Jacobian positivity does not establish global injectivity without boundary and overlap conditions.

Worked model. For $f(z) = z + \overline{z}/3$, real and imaginary coordinates scale by $4/3$ and $2/3$, respectively; $J = 8/9$ and $K = 2$. The inverse is real-linear and is explicitly $(\overline{a}w - b\overline{w})/(|a|^2 - |b|^2)$. As $|b| \uparrow |a|$, inverse conditioning diverges.

Extension and solution. Given an affine-triangle map with singular values $\sigma_1 \geq \sigma_2 > 0$, its local Beltrami magnitude is $(\sigma_1 - \sigma_2)/(\sigma_1 + \sigma_2)$. A uniform bound below one certifies local distortion. Adjacent triangles still need consistent shared edges, a globally admissible boundary map and a nonoverlap argument.

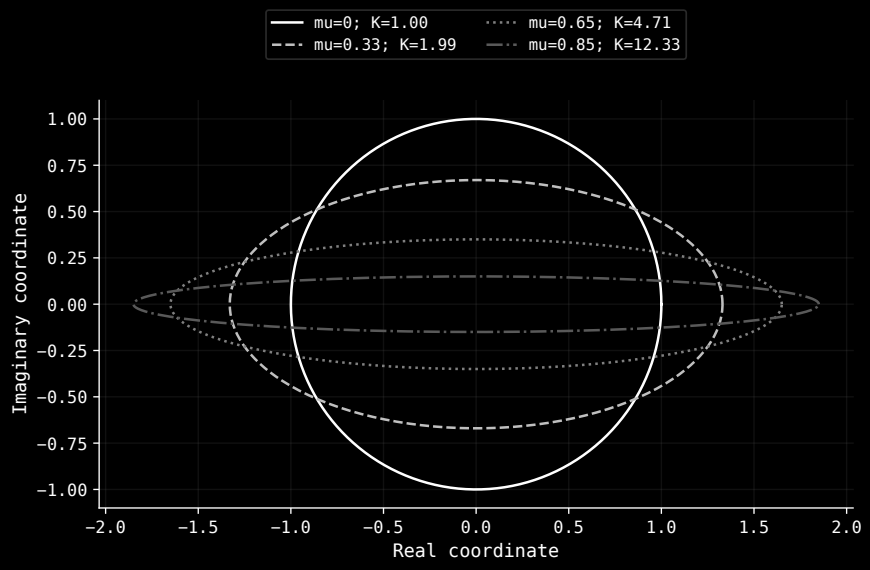


Figure 2.5: Images of concentric circles under $z + \mu\bar{z}$ for real μ. The principal axes and exact distortion increase as $|\mu|$ approaches one.

2.9 CP09. Rational approximation with a pole-exclusion certificate

Analytical references: [24, 8].

Problem. Fit a rational function to data and certify its error on a continuum. A finite sample residual is insufficient when a fitted denominator can vanish between samples.

Representation. A barycentric rational candidate has the form

$$r(z) = \frac{\sum_{j=1}^{m} w_j f_j/(z - z_j)}{\sum_{j=1}^{m} w_j/(z - z_j)}.$$

At nonzero-weight support points, removable limits interpolate f_j. The AAA method uses this representation with adaptive support selection; its methodological provenance is R46. The present task concerns subsequent certification, not a reimplementation or novelty claim for AAA.

Algebraic certificate. After clearing support factors, write $r = P/Q$. On a specified compact set, establish $|Q| \geq q_0 > 0$ and $|fQ - P| \leq \eta$. Then $|f - r| \leq \eta/q_0$. For polynomial residuals on segments, Bernstein coefficients or interval arithmetic can provide complete enclosures. For a holomorphic f, a boundary bound and an interior pole exclusion allow a maximum-principle argument.

Native baseline. For $f(x) = 1/(1 + 25x^2)$ with its formula known, the exact rational representation is already available. A numerical fit adds work. Rational fitting is meaningful for black-box samples or complicated analytic observables; in that setting, denominator conditioning and continuation assumptions become explicit costs.

Extension and solution. Suppose $Q(x) = 1 + x^2 + \delta q(x)$ on $[-1, 1]$ and $\|\delta q\|_\infty \leq 1/4$. Then $|Q| \geq 3/4$. If $\|fQ - P\|_\infty \leq 10^{-6}$, the error is at most $4 \times 10^{-6}/3$. Checking Q only at endpoints cannot establish the same bound for an unrestricted high-degree perturbation.

2.10 CP10. Gradient descent as a polynomial of a linear operator

Analytical references: [6, 7].

Problem. Minimize $F(x) = \frac{1}{2}x^T Hx - b^T x$ with H real symmetric positive definite and spectrum in $[m, L]$. The error of fixed-step gradient descent satisfies $e_{k+1} = (I - hH)e_k$.

Optimal fixed step. The spectral theorem gives

$$\|e_k\|_2 \leq \max_{\lambda \in [m,L]} |1 - h\lambda|^k \|e_0\|_2.$$

The minimax fixed step equalizes the two endpoint magnitudes, giving $h_* = 2/(L + m)$ and contraction $q = (L - m)/(L + m)$. This proof is elementary real linear algebra. The functional-calculus perspective identifies iteration as polynomial approximation of the inverse on a spectral set.

Analytic extension. For a nonnormal linear fixed-point map, spectral radius below one ensures eventual decay in finite dimension but does not give the same finite-time Euclidean bound. A contour representation or a suitable energy norm controls amplification. For a nonlinear objective, the Hessian varies and one must establish the assumed convexity and Lipschitz constants along the iterates.

Comparison. Conjugate gradients exploit symmetry and often outperform stationary iteration; direct factorizations may be best for many repeated right-hand sides. A complex-analytic interpretation is not itself a replacement algorithm. It supplies a common language for polynomial filters, rational acceleration and certified iteration factors.

Extension and solution. For $m = 1, L = 9$, the optimal fixed step is $1/5$ and the factor is $4/5$. Ten steps reduce the proved norm bound by $(4/5)^{10}$. A step $h = 1/9$ has worst-case factor $8/9$ and is slower under this bound. If the target is a function-value error, square the norm factor and multiply by the appropriate quadratic-form constants.

2.11 CP11. Wirtinger differentiation of a complex least-squares objective

Analytical references: [6, 7, 15].

Problem. Minimize $L(z) = \|Az - b\|_2^2$ over $z \in \mathbb{C}^d$. The real-valued loss depends on z and $\overline{z}$ and is not holomorphic except in degenerate cases.

Differential. For a perturbation h,

$$L(z + h) - L(z) = 2\,\mathrm{Re}\langle Ah, Az - b\rangle + \|Ah\|^2.$$

Hence the real gradient represented in complex coordinates is $2A^*(Az - b)$, and the Wirtinger derivative $\partial L/\partial \overline{z}$ is $A^*(Az - b)$. Gradient conventions differ by this factor of two; the convention must be fixed before comparing software.

Baseline. Realifying the system into $2d$ real variables gives the same normal equations. A complex QR or SVD method avoids the conditioning square introduced by explicitly forming A^*A. For iterative optimization, adjoint products compute the gradient at the cost of one forward and one adjoint application.

Invalid analytic check. A direct complex-step formula applied to $|z|^2$ can return zero because the implementation produces a real value. The correct directional derivative is $2\,\mathrm{Re}(\overline{z}h)$. The failure is in the assumed holomorphic continuation of the loss, not in arithmetic precision.

Extension and solution. For a scalar $A = 1 + i$, $b = 2$, the unique minimizer is $z = 1 - i$. At $z = 0$, the complex real-gradient representation is $-4 + 4i$. Substitution checks the minimizer exactly. A penalty $\lambda|z|^2$ adds $2\lambda z$ to the real gradient and changes the normal equation to $(A^*A + \lambda I)z = A^*b$.

2.12 CP12. Holomorphic trial spaces and model error in PDE learning

Analytical references: [18].

Problem. Approximate a harmonic potential on a planar domain by $u_\theta = \operatorname{Re} f_\theta(z)$, with f_θ holomorphic there. Such a trial space satisfies $\Delta u_\theta = 0$ identically, independently of training residual samples.

Construction and certificate. A polynomial or rational function with poles outside the domain provides an explicit analytic class. Fit boundary data by least squares. A complete boundary-supremum bound $\|u_\theta - g\|_{\partial\Omega} \le \epsilon$ then gives an interior error at most ϵ by uniqueness and the maximum principle. Sampled boundary fitting alone does not establish that supremum bound.

Model counterexample. On the unit disk, the solution of $-\Delta u = 4$ with zero boundary data is $u = 1 - |z|^2$. The holomorphic-real-part trial $u_\theta = 0$ matches all boundary data exactly and has zero Laplace residual, but its error at the centre is one. The model equation used in the certificate is wrong for the stated Poisson problem.

Repair. Subtract a known particular solution for the source term and fit the harmonic remainder. For more general source terms, a Green-function or volume-potential approximation introduces additional errors that must be bounded. The representation gain consists of satisfying a correct structural constraint exactly; it cannot compensate for an incorrect physical model.

Comparison. Finite elements and finite differences naturally accommodate inhomogeneous sources and complicated material coefficients. Rational boundary methods can be compact for planar homogeneous problems with controlled geometry. Performance comparisons must include map or pole placement, least-squares conditioning, boundary certification and evaluation count.

Extension and solution. For $-\Delta u = 4$ with boundary data $\cos 2\theta$, the exact solution is $1 - |z|^2 + \operatorname{Re} z^2$. Its source, boundary and uniqueness checks are independent. The example gives an exact regression target for a hybrid particular-plus-holomorphic implementation.

2.13 CP13. Implicit neural layers and branch-dependent derivatives

Analytical references: [6].

Problem. Define an equilibrium by $u - u^2 - \lambda = 0$, selecting the branch $u(0) = 0$. The implicit layer is

$$u(\lambda) = \frac{1 - \sqrt{1 - 4\lambda}}{2}.$$

Its derivative is $(1 - 4\lambda)^{-1/2}$ on the chosen branch. The nearest complex singularity to zero is $\lambda = 1/4$.

Local theorem. The implicit-function theorem applies whenever $1 - 2u \neq 0$. Differentiating the residual gives $(1 - 2u)u' = 1$. Thus derivative blow-up near the branch point is intrinsic to the selected solution map. Replacing the optimizer or increasing arithmetic precision cannot remove this condition number.

Branch verification. A small residual permits both roots. For a real target $0 \le \lambda < 1/4$, the required branch has $0 \le u < 1/2$. A bracket in this interval together with monotonicity of $u - u^2$ gives a native real certificate. A complex contour or continued square root supplies a broader parameter-domain certificate when complex continuation is genuinely required.

Complexity. Unrolling K fixed-point iterations incurs K forward and reverse operations and introduces truncation bias. Implicit differentiation uses a linear solve with Jacobian $1 - 2u$ in the scalar case. The solve can be cheap but ill-conditioned near the fold. For a high-dimensional layer, factorization or Krylov costs and approximate-adjoint residuals must be included.

Extension and solution. For $\lambda = 3/16$, the two roots are $1/4$ and $3/4$; the selected branch is $1/4$ and its derivative is two. Both roots have zero residual. The branch label is therefore part of the model contract and of any benchmark claiming correct outputs.

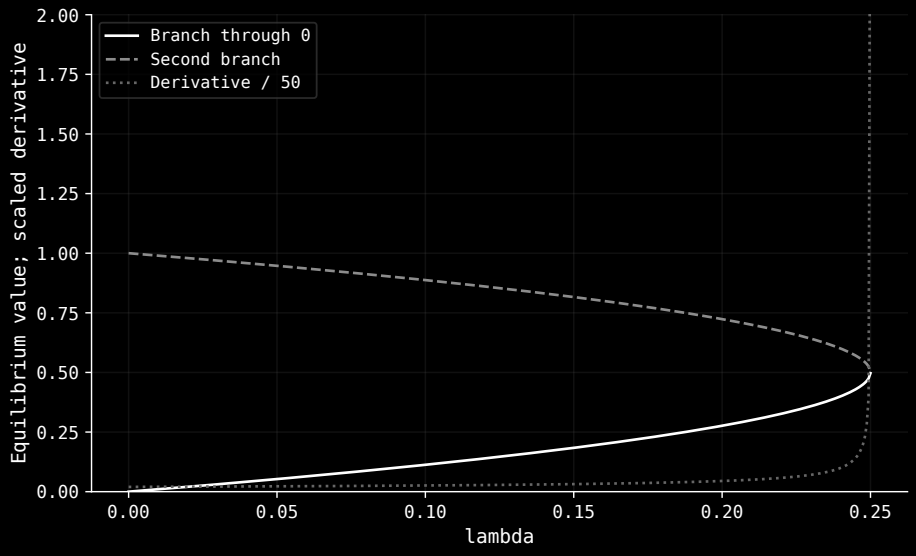


Figure 2.6: Two branches of the implicit equation $u - u^2 = \lambda$ and the derivative growth on the branch through the origin. Equal residuals do not identify the intended branch.

2.14 CP14. Analytic kernels and finite-feature error

Analytical references: [9, 7].

Problem. On the unit disk consider the Hardy kernel $K(z, w) = (1 - z\overline{w})^{-1}$. For nodes with $|z|, |w| \le r < 1$, approximate it by the finite feature sum $K_m(z, w) = \sum_{k=0}^{m-1} z^k \overline{w}^k$.

Derivation. The geometric series gives the uniform kernel error

$$|K - K_m| \le \frac{r^{2m}}{1 - r^2}.$$

For nodes $z_1, \ldots, z_n$, the Gram matrix of K_m is VV^* with $V_{jk} = z_j^{k-1}$, so it is positive semidefinite and has rank at most m. The entrywise bound yields the conservative spectral bound $\|G - G_m\|_2 \le nr^{2m}/(1 - r^2)$.

Inference cost. With regularization λI, a low-rank factorization can reduce repeated solve costs through the Woodbury identity. The resolvent identity bounds the inverse perturbation by $\|(G + \lambda I)^{-1}\| \, \|G - G_m\| \, \|(G_m + \lambda I)^{-1}\| \le \|G - G_m\|/\lambda^2$. This translates approximation accuracy into a prediction-solve guarantee.

Limits. The kernel encodes a specific analytic hypothesis class. It is not automatically suitable for arbitrary image, text or tabular data. Near the unit-circle boundary, the feature degree needed for a fixed tolerance increases. Changing the coordinate map changes the hypothesis class and the interpretation of regularization.

Extension and solution. For $r = 1/2$ and $m = 10$, the entrywise error is at most $4/(3 \cdot 4^{10})$. For $n = 100$ the displayed matrix-norm bound multiplies that value by 100. An empirical low-rank spectrum can be better, but the guaranteed rank selection must use the declared uniform bound or an independently certified sharper estimate.

2.15 CP15. Quantum propagation and polynomial approximation on a real spectrum

Analytical references: [25, 10].

Problem. For Hermitian H, approximate $e^{-itH}\psi$ while controlling the vector error. A physical state requires normalization, but an operator-error certificate is stronger because it applies to every input state.

Spectral reduction. If the spectrum lies in $[a, b]$, a polynomial p_m satisfying $\max_{[a,b]} |e^{-itx} - p_m(x)| \le \epsilon$ yields

$$\|(e^{-itH} - p_m(H))\psi\|_2 \le \epsilon \|\psi\|_2.$$

Chebyshev recurrence computes the action using matrix-vector products. Dense diagonalization is a strong baseline when many times or input vectors share the same small matrix; Krylov exponential action is a strong sparse baseline.

Unitarity and interpretation. The exact exponential is unitary because the spectral phases have modulus one. A polynomial approximant need not be exactly unitary. If its operator error is ϵ, its norm defect on a unit state is at most $2\epsilon + \epsilon^2$. Renormalizing the result changes its error and must be included in a comparison.

Complex-analytic role. Contour calculus covers more general analytic matrix functions and separated spectral projectors, while the Hermitian spectral theorem gives the sharper interval bound here. Branch choices matter for logarithms of unitary matrices and cannot be inferred uniquely from the propagator. None of these classical calculations establishes a quantum computational advantage.

Extension and solution. For $H = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $H^2 = I$ and $e^{-itH} = \cos t\, I - i \sin t\, H$. Two scalar functions suffice exactly. This quotient-algebra formula is the preferred baseline for that model, while it remains unavailable for a general large Hamiltonian without additional structure.

2.16 CP16. Partition functions on a path and the cost of an exact baseline

Analytical references: [11, 26].

Problem. For the n-vertex path, let $Z_n(\lambda)$ sum $\lambda^{|I|}$ over independent vertex sets. The endpoint is either absent or present. Therefore

$$Z_n = Z_{n-1} + \lambda Z_{n-2}, \quad Z_0 = 1, \quad Z_1 = 1 + \lambda.$$

The exact recurrence computes one parameter value in $O(n)$ arithmetic operations and constant working memory, excluding output bit size.

Analytic data. The generating function in the path length is $(1 + \lambda z)/(1 - z - \lambda z^2)$. Its poles determine large-n growth. Logarithmic derivatives in λ yield expected occupancy and variance under the normalized Gibbs weights. A second variable can record spatial observables, with a transfer matrix preserving the exact finite model.

Comparison. Enumerating 2^n subsets is an unnecessarily weak baseline for a path. A zero-free approximation must be compared with the recurrence and, for very large single-index queries, matrix powering. On unrestricted graph classes the exact partition problem can be much harder; that change of graph class must be stated before discussing a benefit from zero-free truncation.

Failure analysis. A polynomial positive for all positive real λ may still have complex zeros near the expansion point. Positivity at the target does not prove convergence of a Taylor expansion of $\log Z$. Conversely, a visible nearby zero does not by itself quantify a training optimizer's behavior. The coefficient cost, zero-free domain and output tolerance are independent obligations.

Extension and solution. For $n = 3$, $Z_3 = 1+3\lambda+\lambda^2$, matching M10. For $\lambda = 1$, the recurrence counts Fibonacci-shifted path sets. The derivative recurrence produces expected occupancy in the same order of arithmetic work. This finite calculation supplies a transparent control for a more elaborate approximate-counting pipeline.

Chapter 3

Stochastic models, estimation and service systems

3.1 ST01. Service-time variability in a single-server queue

Analytical references: [27, 15, 12].

Problem. Arrivals form a Poisson process of rate λ. Service requirements are independent and identically distributed copies of $B \geq 0$, independent of arrivals, with mean b and finite second moment. A first-come, first-served, work-conserving single server processes one unit of work per unit time, with $\rho = \lambda b < 1$. Compute the mean stationary waiting workload. An exponential-service formula applied to an arbitrary service distribution is an invalid baseline.

Transform derivation. Let V be the stationary workload, $W(s) = \mathbb{E}e^{-sV}$ and $\beta(s) = \mathbb{E}e^{-sB}$, $s > 0$. Between arrivals the positive workload decreases at unit speed; an arrival adds an independent copy of B. The stationary infinitesimal balance for e^{-sV} is

$$0 = s\{W(s) - \Pr(V = 0)\} + \lambda\{\beta(s) - 1\}W(s).$$

Work conservation gives $\Pr(V = 0) = 1 - \rho$. Consequently

$$W(s) = \frac{(1-\rho)s}{s - \lambda + \lambda\beta(s)}.$$

Expansion $\beta(s) = 1 - bs + \mathbb{E}(B^2)s^2/2 + o(s^2)$ gives

$$\mathbb{E}V = \frac{\lambda\mathbb{E}(B^2)}{2(1-\rho)}.$$

Poisson arrivals see the stationary workload, so this is the mean waiting time before service. The arrival-seen statement requires the Poisson sampling property; it is not automatic for feedback-controlled arrivals.

Comparison. With $b = 1$ and $\lambda = 4/5$, deterministic service gives $\mathbb{E}V = 2$, whereas exponential service gives $\mathbb{E}V = 4$. Both servers have utilization $4/5$. Computing the correct second moment avoids a factor-of-two error without simulating long trajectories. A discrete-event simulator remains a useful independent model check, especially when the assumptions are changed.

Costs and limits. The mean calculation requires two service moments and constant scalar work. Estimating those moments from data has statistical cost. If $\mathbb{E}B^2 = \infty$, the displayed finite mean is unavailable despite $\rho < 1$. The full waiting distribution requires transform inversion or a distribution-specific argument.

Extension and solution. If service variance is $c_b^2 b^2$, then $\mathbb{E}V = \rho b(1 + c_b^2)/(2(1 - \rho))$. At fixed load and mean, reducing variance reduces this model's mean delay linearly. The result does not establish a universal scheduling rule for heavy-tailed or priority queues.



Figure 3.1: Mean waiting workload for deterministic and exponential service at equal mean. The divergence near saturation belongs to the stated stationary queue models.

3.2 ST02. Independent latency stages and repeated transform poles

Analytical references: [27, 15].

Problem. Let $X_1, \ldots, X_m$ be independent exponential stage times of common rate ν. Determine the distribution of $S = \sum X_j$. The model assumes independence of stage times; an arbitrary tandem queue does not satisfy this specification merely because each server has an exponential service law.

Solution. Independence gives $\mathbb{E}e^{-sS} = (\nu/(s + \nu))^m$. Inverting its pole of order m yields

$$p_S(t) = \frac{\nu^m t^{m-1} e^{-\nu t}}{(m-1)!}, \qquad \Pr(S > t) = e^{-\nu t} \sum_{j=0}^{m-1} \frac{(\nu t)^j}{j!}.$$

The density follows directly by m-fold convolution as well. Repeated integration by parts proves the survival formula and supplies a real-variable verification independent of residues.

Complexity. For one tail probability the finite sum costs $O(m)$ scalar operations with a recurrence for consecutive terms. Monte Carlo with absolute standard error ϵ typically requires order ϵ^{-2} independent replications; rare-event relative error can be substantially more expensive. A library incomplete-gamma routine is the appropriate strong numerical baseline and must be retained in performance claims.

Adversarial case. Replacing X_2 by X_1 preserves both marginal distributions but gives $S = 2X_1$ when $m = 2$. Its variance is twice the independent sum's variance. A product of marginal transforms is therefore a testable independence assumption.

Extension and solution. For $m = 3$, $\nu = 2$, the mean is $3/2$, variance $3/4$, and survival at $t = 1$ is $5e^{-2}$. Distinct stage rates give a sum of simple-pole contributions. When rates nearly

coincide, the repeated-pole representation or divided differences avoid cancellation in nearly singular partial-fraction coefficients.

3.3 ST03. Absorption, reachability probabilities and resolvent remainders

Analytical references: [15, 12].

Problem. A finite transient state set has transition matrix $Q \geq 0$; a vector $r \geq 0$ records one-step absorption in a selected terminal set. Compute eventual absorption probabilities u. First-step conditioning gives $u = Qu + r$. Under $\rho(Q) < 1$,

$$u = (I - Q)^{-1} r = \sum_{j \geq 0} Q^j r.$$

The same formula is a generating resolvent evaluated at one. A Boolean reachability question uses graph search and has a different output contract.

Certified truncation. If $\|Q\|_\infty \leq q < 1$, then the m-step approximation satisfies

$$\left\| u - \sum_{j=0}^{m} Q^j r \right\|_\infty \leq \frac{q^{m+1}}{1-q} \|r\|_\infty.$$

For any candidate $\widetilde{u}$, the residual certificate is $\|u - \widetilde{u}\|_\infty \leq \|r - (I - Q)\widetilde{u}\|_\infty / (1 - q)$. The proof sums a geometric operator series. Some absorbing chains have $\|Q\|_\infty = 1$; one may instead establish $\|Q^k\|_\infty < 1$ for a block length k.

Worked instance. For

$$Q = \begin{pmatrix} 1/2 & 1/4 \\ 0 & 1/2 \end{pmatrix}, \qquad r = \begin{pmatrix} 1/4 \\ 1/2 \end{pmatrix},$$

the exact solution is $(1, 1)^T$. Substitution is an exact verification. The infinity-norm bound $q = 3/4$ is sufficient, although spectral radius $1/2$ permits sharper structure-sensitive estimates.

Cost comparison. A sparse recurrence costs $O(m \operatorname{nnz} Q)$ and needs linear vector storage. A sparse linear solve can be substantially faster for many right-hand sides after factorization. The residual theorem applies to both constructions. Strong conditioning is lost as absorption becomes rare; a small residual alone then has less forward significance.

Extension and solution. Replacing the first row's absorption entry by zero while retaining all other values gives $u_2 = 1$, $u_1 = 1/2$. The leftover probability represents absorption outside the selected set. Renormalizing every row to sum to one would change the problem.

3.4 ST04. Distribution convolution and exact support accounting

Analytical references: [21, 12].

Problem. Independent integer-valued variables X, Y have supports $0, \ldots, a$ and $0, \ldots, b$. Their probability generating polynomials are $P(z)$ and $Q(z)$. The coefficient of z^k in $P(z)Q(z)$ is $\Pr(X + Y = k)$.

Fourier formulation. On Lth roots of unity, multiplication corresponds to circular convolution. Choosing $L \geq a + b + 1$ recovers the full linear convolution without wrapping. This is the distributional instance of CP03. The normalization identity $P(1)Q(1) = 1$, nonnegativity and support are additional probabilistic checks.

Repeated sums. For m independent copies, $P(z)^m$ determines the sum. Binary exponentiation of polynomial products and FFT methods can reduce arithmetic cost. A sequential probability recurrence is competitive for small support or when only a narrow tail is needed. If rational probabilities are required exactly, rational convolution or integer scaling plus a certified transform is necessary.

Worked instance. A variable taking $0, 1, 2$ with probabilities $1/4, 1/2, 1/4$ has $P(z) = (1+z)^2/4$. Three copies sum to a binomial law with six trials: $P(z)^3 = (1+z)^6/64$. The probability of total three is $20/64 = 5/16$. This result can be checked by counting six independent fair bits.

Failure and efficiency. Insufficient Fourier length computes the sum modulo L, not the ordinary sum. Floating coefficients can contain small negative roundoff values; clipping and renormalization change the computed object and require a new error account. The transform method accelerates a mathematical aggregate under independence; it does not establish SQL bag semantics or distributed transaction correctness.

Extension and solution. If a Poisson number of independent marks with generating function P is accumulated, the resulting generating function is $\exp(\lambda(P(z) - 1))$. Its support is generally unbounded. A finite FFT then requires an explicit tail and aliasing bound.

3.5 ST05. Characteristic functions and bounded distributional features

Analytical references: [6, 15, 12].

Problem. Compare two real distributions P, Q when their means or variances may not exist. Characteristic functions $\phi_P(t) = \mathbb{E}e^{itX}$ always exist because the integrand has unit magnitude. A mean-based comparison is undefined for important heavy-tailed models.

Analytic metric. For $\ell > 0$ let $w(t) = \ell e^{-\ell^2 t^2/2}/\sqrt{2\pi}$ and define

$$D^2(P, Q) = \int_{\mathbb{R}} |\phi_P(t) - \phi_Q(t)|^2 w(t)\, dt.$$

Fubini is justified by boundedness. The Gaussian Fourier integral gives

$$D^2 = \mathbb{E}k(X, X') + \mathbb{E}k(Y, Y') - 2\mathbb{E}k(X, Y), \quad k(x, y) = e^{-(x-y)^2/(2\ell^2)}.$$

Independent primed copies are required. The identity gives an exact connection between Fourier features and a Gaussian-kernel discrepancy.

Computation. An empirical all-pairs kernel statistic costs quadratic work in sample size. Approximating the frequency integral by m sampled frequencies gives $O(mn)$ feature construction, followed by vector operations. This introduces quadrature or Monte Carlo error in addition to sampling error. The two errors need separate uncertainty statements.

Limit of finite observations. Agreement at finitely many frequencies does not force equality of distributions. Agreement at every real frequency does. A frequency design should therefore state the alternative class it aims to distinguish; a small feature discrepancy is not a universal distribution-identity certificate.

Extension and solution. For point masses at a and b, the exact value is $2(1 - e^{-(a-b)^2/(2\ell^2)})$. It vanishes only when $a = b$. The same expression shows saturation for widely separated locations, so the bandwidth changes the scale of sensitivity without changing the underlying physical data.

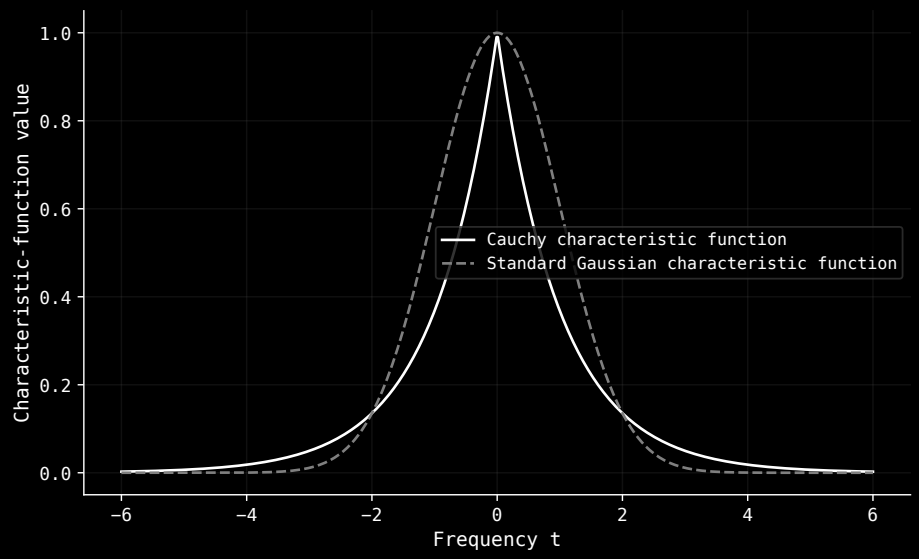


Figure 3.2: Characteristic-function behaviour for light- and heavy-tailed distributions. Existence of the transform does not require existence of its derivatives at zero.

3.6 ST06. Tail-latency mixtures and model identifiability

Analytical references: [27, 28].

Problem. A service time is exponential with rate one with probability 99/100, and exponential with rate 1/100 with probability 1/100. Compare its tail with an exponential model having the same mean.

Exact transforms and tails. The Laplace transform and survival function are

$$L(s) = \frac{99}{100(s+1)} + \frac{1}{100(100s+1)}, \qquad \Pr(X > t) = \frac{99}{100}e^{-t} + \frac{1}{100}e^{-t/100}.$$

The mean is 199/100. At $t = 100$, the mixture survival is greater than $e^{-1}/100$, whereas the matched exponential survival is $e^{-10000/199}$. The closest transform pole to the origin identifies the slow tail scale.

Comparative efficiency. Partial fractions supply a family of tail values with constant scalar work per query. A sample mean cannot identify the slow component; fitting one exponential from that mean conceals the tail. A two-component mixture fitted to finite data is also uncertain, particularly when rare observations are scarce. The exact calculation validates the model's implications, not its empirical selection.

Systems interpretation. Caching, retries, background compaction and network stalls can motivate mixture models, but these mechanisms must not be inferred from the fitted poles alone. A confidence interval for a quantile requires a sampling design and a distributional or order-statistic argument. The observed maximum is not a hard upper bound.

Extension and solution. For any fixed mean m, distributions with arbitrarily large second moment can be formed by placing a sufficiently small mass at a sufficiently large value. Consequently mean service time alone cannot determine the queueing mean in ST01. A cost report should retain at least the variability assumptions and the tail output requested.

3.7 ST07. Optimal timeout allocation under a declared resource budget

Analytical references: [27, 7].

Problem. Independent request streams have Poisson rates $\lambda_j > 0$. A timer of length τ_j succeeds when the next request arrives before expiry, with probability $1 - e^{-\lambda_j \tau_j}$. Maximize $\sum_j w_j(1 - e^{-\lambda_j \tau_j})$ under $\tau_j \geq 0$ and $\sum_j \tau_j \leq T$, with $w_j > 0$.

Solution. The objective is strictly concave. The Karush–Kuhn–Tucker conditions give

$$\tau_j = \frac{1}{\lambda_j}\left[\log \frac{w_j \lambda_j}{\nu}\right]_+, \qquad \sum_j \tau_j = T,$$

for a unique multiplier $\nu > 0$ when $T > 0$. The sum is continuous and decreasing in ν, so bisection determines the multiplier with an explicit bracket. The exponential survival law is the analytic input; the optimization is a standard convex-resource calculation.

Cost comparison. A general constrained optimizer is a strong baseline. The specialized representation reduces a vector optimization to a monotone scalar solve, taking $O(d \log(1/\epsilon))$ arithmetic work for d timers at multiplier tolerance ϵ. Converting multiplier tolerance to objective tolerance requires the derivative bound on the chosen bracket.

Worked instance. If two streams have equal rates and weights, strict concavity gives $\tau_1 = \tau_2 = T/2$. If one weight becomes small enough, its optimal timer is zero. This active-set transition is captured by the positive-part operation and cannot be described by a globally holomorphic map of all parameters.

Model boundary. A real cache typically constrains expected occupancy or a hard number of objects. The timer sum used here is a different resource. Substituting it for capacity without a physical accounting model invalidates any policy recommendation. The formula is suitable for a controlled coursework experiment whose budget is explicitly the sum of timers.

3.8 ST08. Circular Wiener estimation and the distinction between batch and causal outputs

Analytical references: [21, 15].

Problem. Observe $y = x + e$ on n periodic samples, where x has circulant covariance $C \geq 0$, e is independent with covariance $\sigma^2 I$, and both means vanish. The optimal linear mean-square estimator is

$$\widehat{x} = C(C + \sigma^2 I)^{-1} y.$$

The normal equations follow by differentiating expected squared error or by orthogonal projection.

Spectral representation. With unitary Fourier matrix F, write $C = F^* \operatorname{diag}(\lambda_k) F$. Then

$$\widehat{x} = F^* \operatorname{diag}\left(\frac{\lambda_k}{\lambda_k + \sigma^2}\right) F y.$$

Two FFTs and a pointwise product cost $O(n \log n)$ work. A dense positive-definite solve costs $O(n^3)$ factorization work; reuse reduces subsequent solves to quadratic work. The benchmark includes a dense solve for each call and assumes covariance construction has been completed for both methods.

Recorded result. At $n = 512$, the recorded dense solve takes approximately 4.187 ms and the FFT application 0.0204 ms. The maximum output difference is about 2.67×10^{-15}. This

comparison concerns a known circulant matrix; a general covariance matrix need not admit Fourier diagonalization.

Adversarial distinction. The output uses the entire observation vector, including future samples under a temporal interpretation. It is a circular batch estimator. A causal streaming Wiener filter has a different admissible information set and requires factorization and projection. Spectral multiplication by a magnitude ratio does not establish causal realizability.

Extension and solution. If $\lambda_k = 0$, the estimated mode is exactly zero. If $\sigma^2 \to 0$ and $\lambda_k > 0$, its gain tends to one. At $\lambda_k = \sigma^2 = 0$, the displayed fraction is undefined; the mode is deterministically zero in the model and can be assigned estimator value zero by the corresponding pseudoinverse convention.

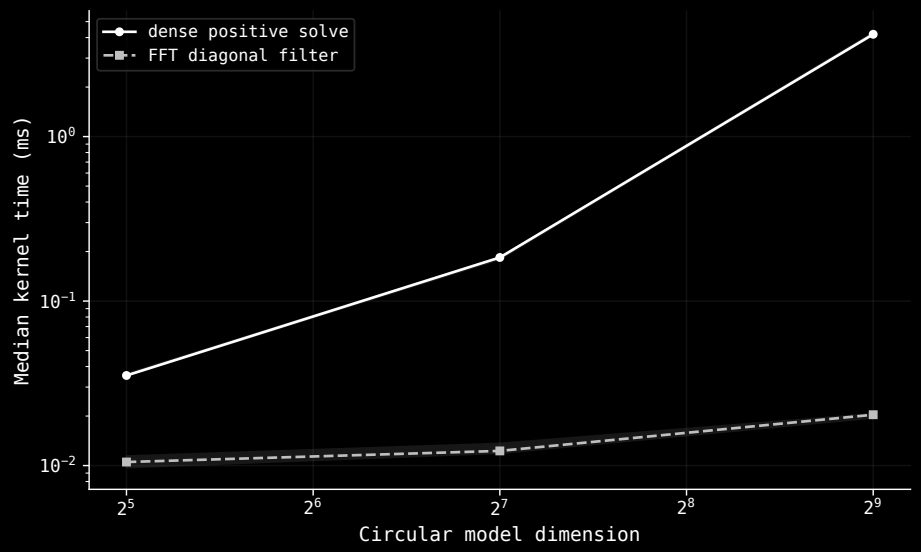


Figure 3.3: Recorded circular Wiener-estimation timings under the stated preprocessed-covariance contract. Reusing a dense factorization changes the baseline and is not included in this curve.

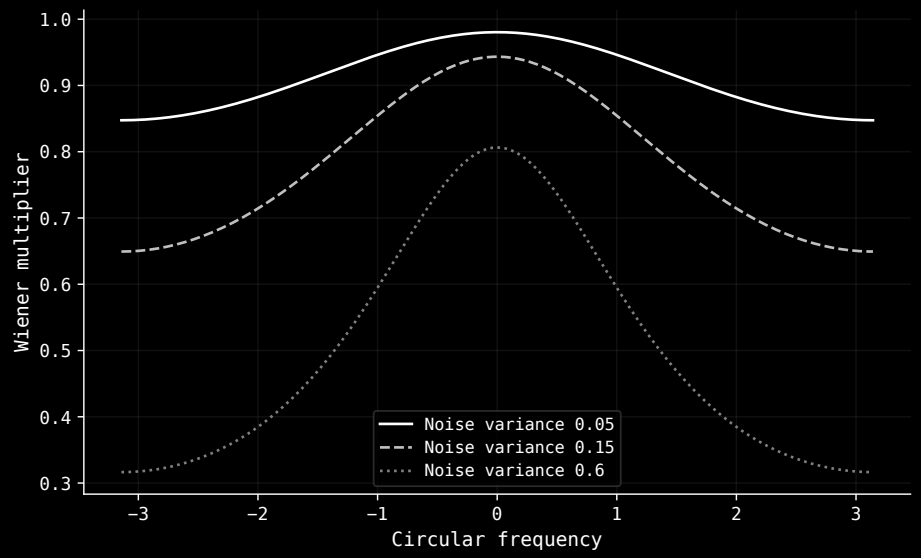


Figure 3.4: Spectral attenuation in a circular Wiener estimator. Mode-dependent gains follow from the assumed signal and noise covariances.

3.9 ST09. PageRank residuals, parameter sensitivity and sparse computation

Analytical references: [29, 10].

Problem. Let $P \geq 0$ be column stochastic, $v \geq 0$ with $\mathbf{1}^T v = 1$, and $0 < \alpha < 1$. The score vector satisfies $x = \alpha P x + (1-\alpha)v$. Its resolvent representation is

$$x = (1-\alpha)(I - \alpha P)^{-1} v.$$

The geometric series proves positivity and normalization. Sparse power iteration is the conventional baseline; forming a dense inverse is unnecessary.

Residual certificate. For a numerical candidate $\widetilde{x}$ and residual $r = (1-\alpha)v - (I - \alpha P)\widetilde{x}$, the induced one-norm gives

$$\|x - \widetilde{x}\|_1 \leq \frac{\|r\|_1}{1-\alpha}.$$

This bound applies to power, Krylov and direct methods alike. A top-two ranking is certified when the computed score gap exceeds twice an appropriate component error bound; small global residual alone does not resolve a smaller ranking gap.

Sensitivity. Differentiating the defining equation gives

$$(I - \alpha P)x' = Px - v.$$

One additional solve supplies the derivative after the score is known. Analytic dependence holds for $|\alpha| < 1$ in this norm bound. The growth near one reflects a loss of contraction; graph structure can permit a sharper estimate on a mean-zero subspace.

Cost and counterexample. Each power step costs $O(\operatorname{nnz} P)$ and residual evaluation has the same order. A contour integral can express the solution but supplies no automatic saving for this one resolvent. A nonstochastic matrix or unhandled dangling column invalidates the normalization proof.

Extension and solution. For $P = I$, the exact score is v for every $\alpha < 1$, so $x' = 0$. The bound $1/(1-\alpha)$ can therefore be conservative. The example illustrates the distinction between a uniform worst-case certificate and parameter-specific conditioning.

3.10 ST10. Streaming transforms, cumulants and a zero-sensitive logarithm

Analytical references: [26, 6, 12].

Problem. For observations x_j, maintain the empirical characteristic function at fixed frequencies $t_1, \ldots, t_m$:

$$\widehat{\phi}_n(t_k) = \frac{1}{n} \sum_{j=1}^{n} e^{i t_k x_j}.$$

The sufficient aggregate is the pair consisting of the count and the complex sums. Combining partitions by adding these quantities preserves the exact empirical feature, including unequal partition sizes.

Complexity. One observation costs $O(m)$ elementary-function evaluations and storage is $O(m)$. The choice of frequencies fixes the information retained. The features are bounded, which supports distribution-free concentration for each real and imaginary component. A finite set of frequencies does not recover arbitrary data or all distributional properties.

Cumulant limitation. When a genuine analytic moment-generating function exists near zero, derivatives of its logarithm are cumulants. A characteristic function may fail differentiability at zero or vanish away from zero. Computing a complex logarithm of a near-zero empirical characteristic value is ill-conditioned and branch-dependent. The central limit theorem alone does not supply an analytic neighbourhood for that logarithm.

Worked counterexample. The symmetric two-point law on $\{-1, 1\}$ has $\phi(t) = \cos t$, which vanishes at $\pi/2$. Its moment-generating function $\cosh s$ is entire, but its logarithm still has complex singularities. A polynomial surrogate for the log-transform has a domain determined by those zeros.

Extension and solution. Partitions with counts two and one and real feature sums two and zero have global mean $2/3$. The unweighted mean of their local means is $1/2$. This exact discrepancy is a distributed-aggregation defect that precedes any analytic approximation.

3.11 ST11. Analytic kernels inside verified software workflows

Analytical references: [30, 28, 31].

Problem. A pipeline parses complex observations, applies an analytic transform, aggregates outputs and commits them to storage. State which parts of the final result follow from the analytic theorem and which require software semantics.

Decomposition. Represent the pipeline as $C \circ A \circ P$, with parser P, numerical algorithm A and commit operation C. A proof that $\|A(x) - F(x)\| \leq \epsilon$ applies only to inputs x satisfying its domain contract. Parsing must establish that the stored bytes represent such an x. The commit protocol must establish the intended atomicity, order and durability properties. No Cauchy estimate supplies those protocol assertions.

Finite counterexample. Binary floating arithmetic can distinguish $(10^{16} + (-10^{16})) + 1$ from $10^{16} + ((-10^{16}) + 1)$. The exact real sum is one, but a reassociated reduction can return zero. An algebraically associative specification does not make the implementation associative.

Quantitative repair. Under the standard rounding model, no overflow and suitable small nu, sequential summation has absolute error bounded by $\gamma_{n-1} \sum |x_j|$, where $\gamma_k = ku/(1 - ku)$. A balanced pairwise tree replaces the depth by approximately $\lceil \log_2 n \rceil$. Cancellation is retained in the absolute sum. An exact superaccumulator changes the implementation and has its own storage and cost requirements.

Efficiency comparison. Parallel reduction can shorten dependency depth, but communication, serialization and scheduling must enter a wall-clock comparison. Maintaining a numerical error certificate need not certify the entire pipeline. Functional correctness of parsing, schema constraints, transactions and authentication remains attached to the corresponding native methods.

Extension and solution. A database average can be aggregated by a pair (sum, count) with exact count arithmetic and a declared sum representation. Applying a nonlinear function independently to each partition mean generally fails to reproduce its value at the global mean. An explicit counterexample is squaring the means of (-1) and (1).

3.12 ST12. Amortized verification and a measurable deployment threshold

Problem. A conventional method costs b per query. A prepared analytic model has setup cost s, evaluation cost a, and one certification cost c. Determine when it reduces total elapsed work while meeting the same error contract.

Exact threshold. For m queries, the costs are mb and $s+c+ma$. If $b > a$, the prepared method is cheaper exactly when

$$m > \frac{s+c}{b-a}.$$

If $b \leq a$, setup cannot be amortized into a benefit under this model. Costs dependent on tolerance must be written $s(\epsilon), a(\epsilon), c(\epsilon)$ before comparing query counts.

Worked instance. For $b = 5$, $a = 1$, $s = 30$, $c = 10$ in a common time unit, equality occurs at ten queries and a strict saving starts at eleven. Changing the input family can invalidate the certificate and force setup to be repeated. The relevant count is reuse within a certified family, not the total number of unrelated calls made by a service.

Experimental design. A paired benchmark must retain the same input set, output dimension, error tolerance, failure policy, thread count and warm/cold convention. Report setup and per-query timings separately, the number of failures, complete input generation, and accuracy against an independent reference. Quantiles of repeated timing batches describe that execution experiment; they do not establish population performance on unknown workloads.

Extension and solution. For a rank-limited surrogate, representation length may decrease while certification cost increases. A Pareto comparison preserves both coordinates. The tuple of time, memory, approximation error and proof status has no canonical scalar weighting. A project can prove an optimum only after defining its admissible methods and its objective, as in the Hardy restriction theorem in the retained core.

Chapter 4

Analytic response, propagation and stable realization

4.1 EL01. Coalescing circuit poles and a regular transient representation

Analytical references: [3, 15, 32].

Problem. Invert $H(s) = ((s+a)^2 - d^2)^{-1}$ for $a > |d| \geq 0$ and $t \geq 0$. Partial fractions give two exponentials with coefficients proportional to $1/d$; their individual magnitudes diverge as $d \to 0$.

Regular solution. Combining the residues before evaluation gives

$$h(t) = e^{-at}\frac{\sinh(dt)}{d} = te^{-at}\sum_{k=0}^{\infty}\frac{(dt)^{2k}}{(2k+1)!}.$$

The limit at $d = 0$ is te^{-at}. The series is entire in d^2, so pole coalescence does not make the impulse response singular. Direct transformation of the series or differentiation of sinh verifies the differential equation and initial data.

Comparison. For moderate $|dt|$, a specialized hyperbolic-sinc evaluation avoids subtracting nearly equal exponentials. For large negative exponents, scaled exponential forms can avoid overflow in sinh followed by underflow in e^{-at}. A robust state-space matrix exponential is an independent baseline. Conditioning of the response must be separated from conditioning of the partial-fraction coefficients.

Quantitative remainder. If $|dt| \leq R$, truncating after $k = m$ gives an absolute error bounded by $te^{-at}\sum_{k>m} R^{2k}/(2k+1)!$. Once $R^2/((2m+4)(2m+5)) < 1$, the first omitted term and a geometric ratio bound enclose the remaining tail.

Extension and solution. For an underdamped model put $d = i\omega$; the same entire formula becomes $e^{-at}\sin(\omega t)/\omega$. This is a representation invariant across critical damping. Physical coefficients and input normalization must be restored when the model is derived from an RLC circuit.

4.2 EL02. A continuous-contour Nyquist certificate

Analytical references: [3, 15, 32].

Problem. For a meromorphic loop transfer $L(s)$ and a positively oriented contour Γ bounding a region D, determine the number of zeros of $1 + L$ in D. Assume no poles or zeros occur on Γ. The argument principle gives

$$Z_D(1 + L) - P_D(1 + L) = \operatorname{wind}((1 + L) \circ \Gamma, 0).$$

The contour orientation fixes the sign. A conventional engineering Nyquist contour may have the opposite orientation; its formula must be adjusted consistently.

Finite certificate. Subdivide Γ into arcs. For a sample s_j and a bound $M_j \geq \sup |L'|$ on an arc whose arclength from s_j to every point is at most h_j, the inequality $|1 + L(s_j)| > M_j h_j$ excludes zero throughout that arc. A compatible homotopy to a polygon then permits an exact winding count. Interval or rational bounds must include the sample error.

Comparison. For a low-degree real polynomial, Routh–Hurwitz or exact root isolation can be shorter and stronger. Nyquist becomes useful for frequency-domain models, delays, meromorphic transfer functions and parameter-dependent families. A plotted curve with arbitrarily many unvalidated samples has no universal exclusion guarantee.

Open-region completion. A finite contour certifies a bounded region. A right-half-plane stability conclusion additionally needs a high-frequency exclusion bound and treatment of poles on the imaginary axis. For delay systems there are infinitely many characteristic roots, so a finite plot without such control cannot establish global stability.

Extension and solution. If $L(s) = k/(s + a)$ with $a > 0$, then $1 + L = (s + a + k)/(s + a)$. Its only closed-loop zero is $-a - k$. The exact condition is $k > -a$. This elementary calculation checks the contour convention before a more complicated transfer family is analysed.

4.3 EL03. A delay-independent exclusion of infinitely many characteristic roots

Analytical references: [15, 32].

Problem. Study $\dot{x}(t) = -ax(t) - bx(t - \tau)$ with $a > 0$, real b, and $\tau \geq 0$. The characteristic equation is $s + a + be^{-s\tau} = 0$.

Uniform certificate. For $\Re s \geq 0$, $|s + a| \geq a$ and $|be^{-s\tau}| \leq |b|$. Hence $a > |b|$ excludes all characteristic roots from the closed right half-plane for every nonnegative delay. This is a strict uniform zero-exclusion statement. The standard spectral criterion for retarded delay equations then yields asymptotic exponential stability for each fixed delay; a decay rate uniform in unbounded τ is not asserted.

First imaginary crossing. For $b > a > 0$, set $s = i\omega$. Separating real and imaginary parts gives

$$a + b\cos(\omega\tau) = 0, \qquad \omega - b\sin(\omega\tau) = 0.$$

Thus $\omega = \sqrt{b^2 - a^2}$ and the smallest positive candidate delay is

$$\tau_c = \frac{\arccos(-a/b)}{\sqrt{b^2 - a^2}}.$$

At $\tau = 0$ the ordinary differential equation is stable. A full stability interval argument also checks crossing direction and absence of earlier crossings; differentiating the characteristic equation

provides that check. The plot displays the candidate boundary derived from the imaginary-axis calculation.

Efficiency. The strict inequality $a > |b|$ replaces an unbounded search over roots by one scalar margin. It is sufficient and not necessary at a specified finite delay. When it fails, small-delay stability can still hold. The boundary formula uses the exact scalar structure and does not apply unchanged to matrix or nonlinear delay equations.

Extension and solution. For $a = 1$, $b = 2$, the crossing frequency is $\sqrt{3}$ and the first delay is $2\pi/(3\sqrt{3})$. For $b = 1/2$, the delay-independent exclusion margin is $1/2$.

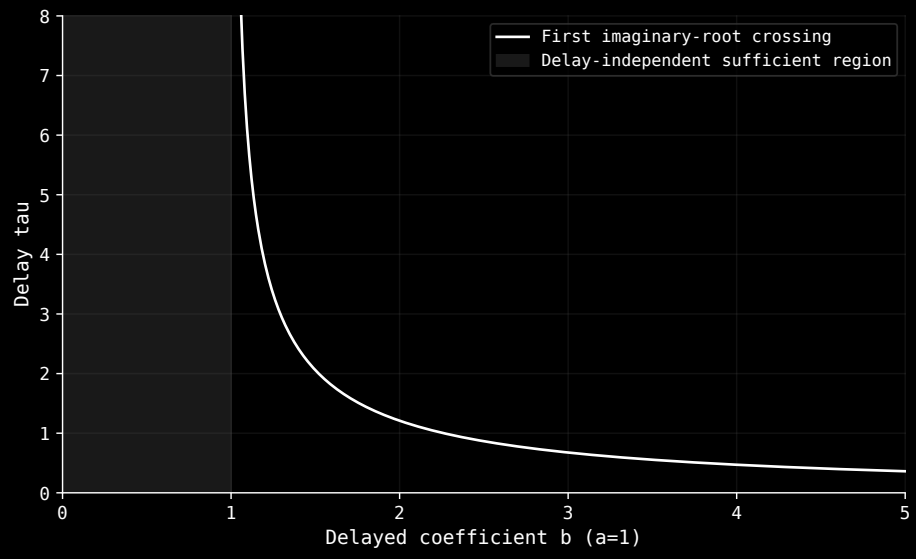


Figure 4.1: The first imaginary-axis crossing curve for the scalar delay family with $a = 1$. The region $|b| < 1$ has a separate delay-independent exclusion proof.

4.4 EL04. Small gain as an analytic norm certificate

Analytical references: [15, 32].

Problem. Let $G(s) = g/(s + a)$ with $a > 0$ and let a stable causal perturbation Δ have induced L^2 gain at most δ. Under zero initial conditions, determine a sufficient condition for invertibility of $I - G\Delta$ on the signal space.

Proof. The Fourier multiplier norm of G is $\|G\|_\infty = |g|/a$. If $q = |g|\delta/a < 1$, the Neumann series converges in operator norm and

$$\|(I - G\Delta)^{-1}\| \le \frac{1}{1 - q}.$$

The proof uses bounded causal operators and the submultiplicative norm. For a frequency-domain statement the stability and realizability assumptions on both factors must be supplied; magnitude data at finitely many frequencies do not establish them.

Cost and interpretation. The certificate consists of a norm bound and a strict margin. It applies to an entire perturbation class, while simulation examines individual members. It can be conservative because it discards phase and interconnection structure. A nominal root calculation may give a larger stability region for a specified parametric perturbation.

Worked comparison. If $a = 2$, $g = 1$, $\delta = 1$, then $q = 1/2$ and the inverse gain is at most two. Increasing δ to two removes the strict certificate. This does not prove every member is unstable; it identifies the boundary of this sufficient theorem.

Extension and solution. For $G = 0$, the exact inverse is the identity and the bound equals one. Saturation, scheduling jitter, packet loss and fixed-point arithmetic need models compatible with the chosen operator class. An analytic transfer certificate cannot certify an unspecified implementation.

4.5 EL05. Schur recursion and stable bounded response parameterization

Analytical references: [21, 15].

Problem. Construct a holomorphic function on $\mathbb{D}$ bounded in magnitude by one from a free Schur-class function g and a scalar a with $|a| < 1$.

Identity and solution. Define $f(z) = (a + zg(z))/(1 + \bar{a}zg(z))$. Direct algebra gives

$$1 - |f(z)|^2 = \frac{(1 - |a|^2)(1 - |zg(z)|^2)}{|1 + \bar{a}zg(z)|^2}.$$

The denominator is nonzero in the disk because $|\bar{a}zg(z)| < 1$. Hence f is holomorphic and disk-valued, and $f(0) = a$. Iterating the construction gives a constraint-preserving rational parameterization when the final free function is rational.

Comparison. Optimizing unconstrained coefficients and testing stability afterward can repeatedly generate inadmissible models. The recursion encodes a norm constraint structurally. It changes the optimization geometry and can introduce parameter sensitivity near $|a| = 1$; it does not guarantee an easy global optimization problem.

Digital convention. For a causal discrete-time transfer, one commonly sets z here equal to the delay variable $q = z_{\text{physical}}^{-1}$. Analyticity in the unit q-disk corresponds to absence of poles in that disk; translating this statement to pole locations in the physical z-plane requires the reciprocal convention. The Schur bound gives a bounded ℓ^2 input–output operator. Absolute summability of the impulse response, required for general BIBO stability, is an additional condition; for reduced rational responses it follows from poles strictly outside the closed delay disk. Quantized-state behaviour requires a separate model.

Extension and solution. For $g(z) = b$ constant with $|b| \le 1$, the recursion gives $(a+bz)/(1+\bar{a}bz)$. If $b = 0$, the response is constant. If $a = 0$, one obtains $zg(z)$, showing that the first parameter specifies the zero-delay value. The proof fails for an arbitrary nonholomorphic learned function g even if a training sample happens to lie inside the disk.

4.6 EL06. Interpolation feasibility from two measurements

Analytical references: [15].

Problem. Does a holomorphic $f : \mathbb{D} \to \overline{\mathbb{D}}$ exist with $f(0) = v$ and $f(r) = w$, where $0 < r < 1$ and $|v| < 1$? Schwarz–Pick gives the necessary inequality

$$\left| \frac{w - v}{1 - \bar{v}w} \right| \le r.$$

It is sufficient: choose $c = (w - v)/(r(1 - \bar{v}w))$, $|c| \le 1$, and set $f(z) = (v + cz)/(1 + \bar{v}cz)$.

Feasible disk. Rearrangement yields a Euclidean disk for w with centre and radius

$$c_v = \frac{(1-r^2)v}{1-r^2|v|^2}, \qquad R_v = \frac{r(1-|v|^2)}{1-r^2|v|^2}.$$

This describes all second observations compatible with the analytic norm constraint. The equivalent two-by-two Pick matrix has entries $(1 - w_i\bar{w}_j)/(1 - z_i\bar{z}_j)$ and must be positive semidefinite.

Comparison. A polynomial fit always interpolates two values but can violate the global unit bound. The feasibility inequality can reject impossible data or quantify uncertainty before any high-order fitting. For many points, positive-semidefinite feasibility is a finite test under the exact interpolation theorem, and matrix conditioning must be monitored.

Worked instance. For $v = 1/2$, $r = 1/2$, the feasible disk has centre $2/5$ and radius $2/5$. Thus the real value $w = 9/10$ is impossible, while $w = 4/5$ is extremal. The rejection is a mathematical statement about the model class, not evidence that the instrument malfunctioned; the physical response might not belong to that class.

Extension and solution. Measurement intervals lead to an existential feasibility problem over the intervals. Checking only their midpoints does not decide feasibility for all possible measurements. A robust version needs interval or semidefinite uncertainty analysis.

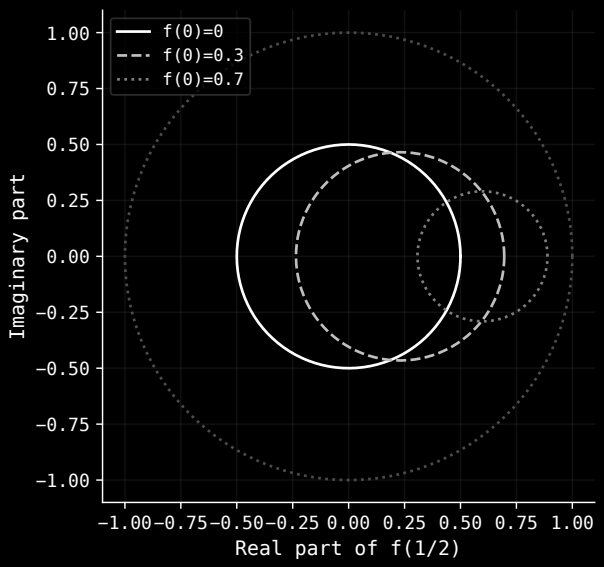


Figure 4.2: Two-point feasible disks as the first value and node separation change. Points outside a disk violate the stated analytic interpolation contract.

4.7 EL07. Spectral factorization and magnitude ambiguity

Analytical references: [21, 15].

Problem. Find a causal finite filter with power spectrum $S(\theta) = 5/4 - \cos\theta$. The identity

$$S(\theta) = |1 - \tfrac{1}{2}e^{-i\theta}|^2$$

gives impulse response $(1, -1/2)$. Direct expansion verifies the entire frequency continuum, including positivity $S \geq 1/4$.

Nonuniqueness. Multiplication by any causal stable all-pass factor preserves magnitude on the unit circle and changes phase. A delay also preserves magnitude. Uniqueness therefore needs a normalization and a minimum-phase or outer-factor convention. A finite list of power measurements provides even less information than the full spectrum.

Causal estimation. The batch estimator in ST08 uses a pointwise spectral ratio. Causal Wiener filtering additionally factors the observed spectrum and projects an appropriate cross-spectrum onto the causal Hardy subspace. These operations encode the time-direction constraint. The finite identity above is the simplest example of the factorization stage.

Comparison and costs. For low-degree rational spectra, factoring roots inside/outside the unit circle is efficient. Toeplitz or state-space algorithms provide stronger general baselines than a dense unconstrained frequency fit. Near unit-circle zeros, the factorization becomes sensitive and a strict spectral lower bound is useful.

Extension and solution. The factor $(1/2 - e^{-i\theta})$ has the same magnitude as $(1 - (1/2)e^{-i\theta})$. Their zeros occupy reciprocal locations in the physical z variable. The equality gives an explicit ambiguity that a magnitude-only identification routine must report.

4.8 EL08. Transmission-line square roots with a coupled branch contract

Analytical references: [33, 34].

Problem. For $R, L, G, C > 0$, define propagation and characteristic impedance on $\Re s > 0$ by

$$\gamma(s)^2 = (R + sL)(G + sC), \qquad Z_c(s)^2 = \frac{R + sL}{G + sC}.$$

The physical branch must be analytic, positive on the positive real axis, and compatible across both quantities.

Construction. Each factor has positive real part on the right half-plane and admits an analytic logarithm there. Set

$$\gamma = \exp\left(\tfrac{1}{2}\log(R + sL) + \tfrac{1}{2}\log(G + sC)\right), \quad Z_c = \exp\left(\tfrac{1}{2}\log(R + sL) - \tfrac{1}{2}\log(G + sC)\right).$$

The logarithms have imaginary parts in $(-\pi/2, \pi/2)$. Thus $\Re\gamma > 0$, and the coupled identities $\gamma Z_c = R + sL$ and $\gamma/Z_c = G + sC$ follow exactly.

Comparison. Independent principal-square-root calls at unrelated frequencies can obscure the required continuation and sign consistency. The logarithmic construction supplies a global branch theorem on the model domain. A stable library routine with an explicit branch convention remains an appropriate computational implementation.

Lossless limit. As $R, G \downarrow 0$, the limiting expressions are $\gamma = s\sqrt{LC}$ and $Z_c = \sqrt{L/C}$ in the right half-plane. On the imaginary boundary, propagation is oscillatory. Selecting $-\gamma$ produces exponential growth into a passive line and violates the chosen directional boundary condition.

Extension and solution. A sign reversal applied simultaneously to both γ and Z_c preserves one product identity but changes the decay and impedance convention. Therefore algebraic squaring checks alone are insufficient. Positive-real-axis normalization and analytic continuation supply the missing information.

4.9 EL09. Array factors as polynomials on the unit circle

Analytical references: [35, 33].

Problem. Compare uniform weights $(1, 1, 1, 1)/4$ with binomial weights $(1, 3, 3, 1)/8$. With inter-element phase ψ, the array factor is a polynomial in $z = e^{i\psi}$.

Exact geometry. The binomial factor is $(1 + z)^3/8$, so

$$|A_b(e^{i\psi})| = |\cos(\psi/2)|^3.$$

The uniform factor is $(1 + z + z^2 + z^3)/4$ and has roots at $-1, i, -i$. The binomial weights place a triple zero at $z = -1$, changing null multiplicity and beam shape. These are exact polynomial statements for the ideal uncoupled array factor.

Cost and interpretation. Direct evaluation at m angles costs $O(mn)$ for n elements; FFT evaluation costs $O(n \log n)$ on a compatible uniform phase grid. Arbitrary physical angle grids need resampling or a nonuniform transform. Polynomial factorization supplies explanatory design information but does not prove a universal optimal radiation pattern.

Physical qualifications. The total field includes the element pattern, spacing, polarization and mutual coupling. The map from angle to phase can create grating lobes. Normalizing the on-axis amplitude does not impose equal total radiated power, so efficiency comparisons must state their normalization.

Extension and solution. For $n + 1$ binomially weighted elements, $A(z) = (1 + z)^n/2^n$ and the null at $\psi = \pi$ has order n. A small phase displacement δ near that null gives amplitude asymptotic to $|\delta/2|^n$. Weight errors can reduce this exact multiplicity, producing a sensitivity question distinct from the nominal design.

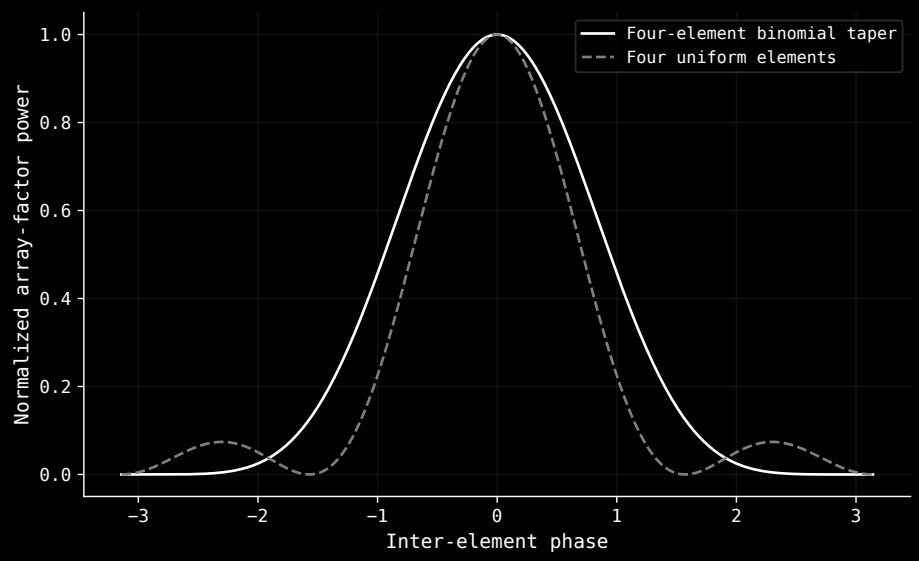


Figure 4.3: Ideal uniform and binomial four-element array factors under equal on-axis amplitude normalization. Element-pattern and coupling effects are outside the model.

4.10 EL10. Dispersive Gaussian pulses and analytic propagation

Analytical references: [35, 33].

Problem. Let $A(0,t) = e^{-t^2/(2T_0^2)}$ and adopt Fourier transform $\widehat{A}(\omega) = \int A(t)e^{-i\omega t}\,dt$. Propagation multiplies the spectrum by $e^{-i\beta_2 z\omega^2/2}$.

Gaussian integration. With $D = \beta_2 z/T_0^2$,

$$A(z,t) = \frac{1}{\sqrt{1+iD}} \exp\left[-\frac{(t/T_0)^2}{2(1+iD)}\right],$$

where the square root is continued from $D = 0$. Thus

$$|A(z,t)|^2 = \frac{1}{\sqrt{1+D^2}} \exp\left[-\frac{(t/T_0)^2}{1+D^2}\right].$$

The intensity width grows by $\sqrt{1+D^2}$ and the integrated energy is constant. Direct Gaussian integration proves both statements.

Comparison. The closed form gives constant work per requested time sample and provides an exact reference for an FFT propagator. A time-domain finite-difference calculation introduces temporal and propagation discretization errors. For arbitrary initial pulses, Fourier propagation remains the strong baseline; the Gaussian formula is a specialized solution, not a general simulation replacement.

Boundary conditions. Attenuation, nonlinearity, higher-order dispersion and polarization coupling alter the model. A nonlinear Schrödinger equation needs additional terms. The chosen Fourier sign determines the chirp sign; intensity broadening alone cannot detect an inconsistent phase convention.

Extension and solution. At $D = 3$, the width factor is $\sqrt{10}$ and peak intensity is $1/\sqrt{10}$. Their product preserves the integral of the Gaussian profile. A reported simultaneous increase of both quantities under this lossless model would violate energy conservation.

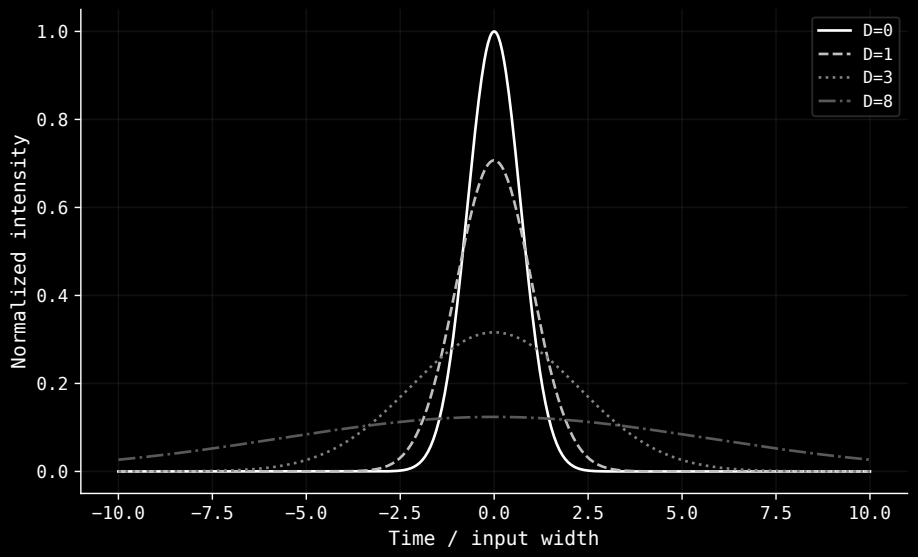


Figure 4.4: Gaussian intensity broadening under normalized linear dispersion. The area under each profile is invariant; the phase convention is stated in the text.

4.11 EL11. OFDM orthogonality and carrier-frequency leakage

Analytical references: [21, 35].

Problem. For N complex samples, prove orthogonality of the tones $e^{2\pi ikn/N}$ and quantify leakage under a normalized carrier-frequency offset δ.

Root-of-unity identity. The inner product of tones k and ℓ is N when $k = \ell \pmod N$ and zero otherwise. With offset δ, the normalized response in bin ℓ is

$$D_N(k - \ell + \delta) = \frac{1}{N} \sum_{n=0}^{N-1} e^{2\pi i(k-\ell+\delta)n/N}.$$

A finite geometric sum gives magnitude $|\sin(\pi x)/(N \sin(\pi x/N))|$, $x = k - \ell + \delta$, with removable values at denominator zeros. The leakage is a precise failure of the nominal orthogonality relation.

Channel condition. A cyclic prefix at least as long as the channel memory converts the retained block's linear convolution into a circular convolution under the standard block model. Fourier diagonalization then permits per-bin equalization. Insufficient prefix, time variation or timing mismatch destroys that diagonal model.

Efficiency. The DFT organizes N inner products into an $O(N \log N)$ FFT. A direct filter-bank calculation costs quadratic work. The algebraic identity is exact, while finite arithmetic, channel estimation and synchronization contribute separate errors. The signal-processing gain is already part of conventional OFDM practice.

Extension and solution. For $\delta = 1$, a tone shifts by one full bin with no fractional leakage. For $\delta = 1/2$, energy spreads across bins. Summed bin energy remains equal to input energy under a unitary DFT, providing an independent implementation check.

4.12 EL12. Bessel modes, entire series and cutoff calculations

Analytical references: [33, 36].

Problem. The radial part of a rotationally symmetric Helmholtz mode satisfies $r^2u'' + ru' + k^2r^2u = 0$. The regular solution is $J_0(kr)$, where

$$J_0(z) = \sum_{j=0}^{\infty} \frac{(-1)^j (z^2/4)^j}{(j!)^2}.$$

The series follows by a power-series substitution and is entire by the ratio test. Boundary conditions select zeros or derivative zeros depending on the physical mode; these conditions cannot be interchanged.

Finite error bound. On $|z| \le R$, the ratio of consecutive absolute terms is at most $R^2/(4(j+1)^2)$. Once this is bounded by $q < 1$ beyond the first omitted index, the remaining absolute tail is bounded by that omitted term divided by $1 - q$. The bound gives a certified evaluation on a fixed disk in exact arithmetic, with roundoff added separately.

Comparison. A standard Bessel-function implementation and bracketed root finder are strong baselines for numerical cutoffs. A large-argument asymptotic can be inaccurate for the first mode. Complex continuation and zero-exclusion estimates are useful for parameter families, root multiplicity and uniform validity, not merely for reproducing a tabulated first zero.

Worked scaling. For a Dirichlet disk of radius a, any zero j of J_0 gives wavenumber $k = j/a$. Therefore doubling the radius halves every wavenumber. This exact dimensional statement can be checked independently of the zero algorithm.

Extension and solution. The regularity condition at $r = 0$ rejects the singular Y_0 solution. On an annulus the centre is absent and both independent radial solutions can enter. Changing domain topology therefore changes the admissible solution basis.

4.13 EL13. Fractional impedance as a positive continuum of relaxations

Analytical references: [15, 36].

Problem. Represent $s^{-\alpha}$ analytically on the right half-plane and approximate it by a passive rational model.

Integral derivation. The beta integral on the positive real axis gives

$$s^{-\alpha} = \frac{\sin(\pi\alpha)}{\pi} \int_0^\infty \frac{t^{-\alpha}}{s+t}\, dt, \qquad \Re s > 0.$$

Both sides are holomorphic on this domain and agree for positive s, so the identity theorem extends the formula. The branch is fixed by positivity on that axis. Since $\Re(1/(s+t)) > 0$, the representation is positive real.

Rational approximation. Positive quadrature weights produce $R(s) = \sum_j c_j/(s+t_j)$ with $c_j, t_j > 0$. Each term has a passive first-order realization, and their sum retains positivity. A generic least-squares rational fit can generate right-half-plane poles or negative residues unless additional constraints are enforced.

Tail control. For $\Re s \geq \sigma > 0$, omitting $0 < t < \epsilon$ contributes at most $\epsilon^{1-\alpha}/((1-\alpha)\sigma)$ before the sine prefactor. Omitting $t > T$ contributes at most $T^{-\alpha}/\alpha$. Quadrature error on the finite interval is a third term. These estimates give a full-domain error decomposition; a frequency plot alone does not.

Limits and units. The ideal power law has unbounded bandwidth and a branch point. A physical fractional model requires scale and units, and finite-band empirical support. The exponent does not uniquely identify a microscopic mechanism.

Extension and solution. For $\alpha = 1/2$, the phase on $s = i\omega$, $\omega > 0$, is $-\pi/4$ and magnitude is $\omega^{-1/2}$. A finite positive relaxation sum can approximate this behaviour on a compact band, but its asymptotic behaviour at zero and infinity follows from its finite poles.

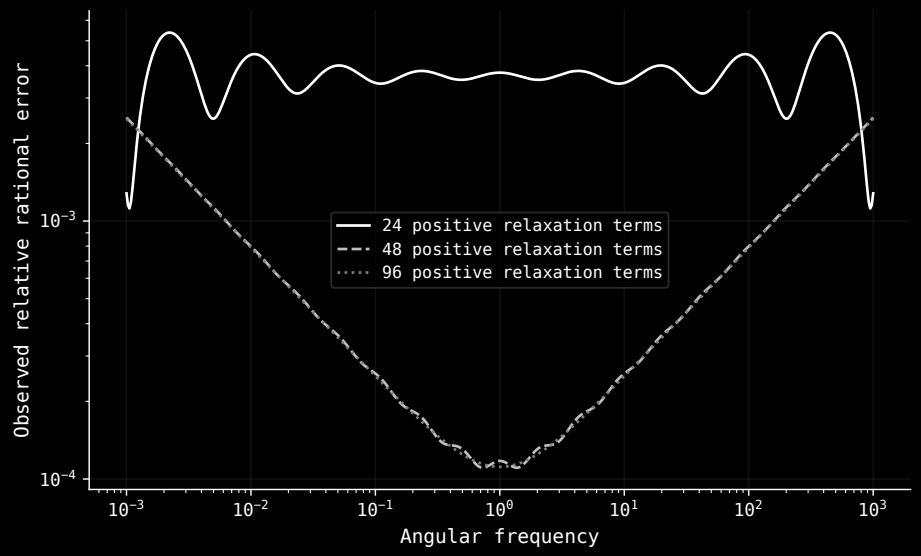


Figure 4.5: A fractional power response and positive-relaxation approximations. The rational family preserves positive-real structure; its numerical fit and finite-band error remain separate questions.

4.14 EL14. Causality, Hilbert transforms and missing frequency bands

Analytical references: [21, 15].

Problem. For the causal impulse $h(t) = \tau^{-1} e^{-t/\tau} \mathbf{1}_{t \geq 0}$, use $H(\omega) = \int h(t) e^{-i\omega t}\, dt$. Then $H(\omega) = 1/(1 + i\omega\tau)$, analytic in a half-plane below the real frequency axis.

Dispersion pair. The real and imaginary parts are

$$H_R(\omega) = \frac{1}{1 + \omega^2\tau^2}, \qquad H_I(\omega) = -\frac{\omega\tau}{1 + \omega^2\tau^2}.$$

With $\mathcal{H}f(\omega) = \pi^{-1}\,\mathrm{PV} \int f(\xi)/(\omega - \xi)\, d\xi$, the pair satisfies $H_I = -\mathcal{H}H_R$. Contour integration or direct partial fractions proves the identity. Reversing the Fourier sign changes the corresponding half-plane and sign relation.

Finite-band issue. A Hilbert transform integrates all real frequencies. Truncating to $[-W, W]$ introduces a missing-tail term. For the Debye real part, the tail decays quadratically, allowing an explicit bound away from the omitted singular denominator region. For an arbitrary measured response, unobserved high-frequency behaviour must be modelled or bounded.

Comparison. A complex analytic response model can enforce a causal dispersion relation across a continuum. A numerical Hilbert transform of finite noisy samples is an estimator with windowing and discretization error. Neither technique alone proves that an experimental device is linear and time invariant.

Extension and solution. Two responses differing by a pure delay have equal magnitudes and different phases. A minimum-phase reconstruction from magnitude requires exclusion or separate treatment of such all-pass factors. This connects the causal reconstruction problem to EL07.

4.15 EL15. Conformal field reduction and the transformed wave equation

Analytical references: [33, 34].

Problem. Let $z = \phi(w)$ be a conformal map with $\phi' \neq 0$, and let $v = u \circ \phi$. The chain rule and Cauchy–Riemann equations give

$$\Delta_w v = |\phi'(w)|^2 (\Delta_z u) \circ \phi.$$

Harmonic functions therefore pull back to harmonic functions. Boundary data must be composed with the same map, and normal derivatives acquire a metric factor.

Wave qualification. If $(\Delta_z + k^2)u = 0$, then

$$\Delta_w v + k^2 |\phi'(w)|^2 v = 0.$$

The transformed coefficient is generally spatially variable. Applying a Laplace solution unchanged to a finite-frequency Helmholtz problem removes this term and changes the governing equation.

Efficiency. For a domain with a known conformal chart, harmonic boundary data can yield a low-dimensional analytic solution. Constructing an unknown map can be expensive and ill-conditioned near crowded geometry. Finite elements, boundary elements and finite differences are the relevant conventional controls. Their cost should be compared at equal boundary and interior error, including map construction.

Physical interpretation. A scalar quasistatic approximation requires an operating regime justifying the neglected dynamics. Three-dimensional vector Maxwell equations do not reduce globally to a single planar holomorphic function. Anisotropy and inhomogeneous material tensors also modify the constitutive model.

Extension and solution. For $\phi(w) = cw$, the wave coefficient becomes $k^2|c|^2$, exactly matching the inverse relation between spatial scale and wavelength. For a non-affine map, retaining constant k in the w coordinate is generally incorrect.

4.16 EL16. Hankel singular values in a separable model-reduction problem

Analytical references: [15, 32].

Problem. Consider independent channels $G_j(s) = g_j/(s + a_j)$, with $g_j, a_j > 0$. For each channel, the Hankel operator maps a past-input profile to a future output and has kernel $g_j e^{-a_j(t+u)}$, $t, u \geq 0$.

Exact singular value. The operator is rank one: $H_j f = g_j e^{-a_j t} \int_0^\infty e^{-a_j u} f(u)\, du$. Since $\|e^{-a_j t}\|_2^2 = 1/(2a_j)$, its unique singular value is

$$\sigma_j = \frac{g_j}{2a_j}.$$

Orthogonal channel separation makes the full Hankel singular values the collection of these numbers. This calculation derives the memory scale directly from an integral operator.

Truncation comparison. Deleting selected independent channels gives exact induced gain error $\max_{j \in J} g_j/a_j$ for Euclidean vector signals, while the sum $2\sum_{j \in J} \sigma_j$ is a valid, generally looser upper bound in this diagonal example. The general balanced-truncation theorem requires controllability, observability and stability hypotheses; it is not proved by this special calculation.

Efficiency. After the parameters are known, channel selection is a sorting task. For coupled systems, Lyapunov equations and balancing can dominate setup. A model with a small state dimension may reduce repeated simulation cost, but the approximation norm and frequency/time interpretation must be stated.

Extension and solution. For channels $(g, a) = (2, 1)$ and $(1, 4)$, the singular values are one and $1/8$. Dropping the second channel gives exact transfer norm error $1/4$. It can still be unacceptable for an output that isolates that channel, illustrating why a global norm must be matched to the actual observable.

4.17 EL17. A winding invariant and an edge-state recurrence

Analytical references: [33, 25].

Problem. For positive couplings t_1, t_2, consider the off-diagonal Bloch symbol $q(k) = t_1 + t_2 e^{ik}$. Its image is a circle of radius t_2 centred at t_1.

Winding calculation. If $t_2 > t_1$, the circle contains zero and has winding one. If $t_2 < t_1$, the winding is zero. At equality the symbol vanishes and the invariant is undefined. The spectral gap of the two-band Hamiltonian is controlled by $|q(k)|$.

Boundary solution. For a semi-infinite chain with the specified sublattice termination, a zero-energy mode has amplitudes $a_n = (-t_1/t_2)^n a_0$ on one sublattice and zero on the other. It is square summable exactly when $t_1 < t_2$. Normalization gives $|a_0| = \sqrt{1 - (t_1/t_2)^2}$. This recurrence supplies an explicit state compatible with the winding result.

Comparison. Finite matrix diagonalization produces approximate modes for a chosen chain length. The analytic recurrence explains localization length and the parameter threshold for the semi-infinite model. A finite chain can exhibit exponentially small splitting between two boundary modes; the infinite zero eigenvalue should not be asserted literally for every finite realization.

Hypotheses. Chiral structure, nearest-neighbour couplings, a gap and a chosen boundary termination are essential. Generic diagonal perturbations can move the zero mode. A winding number without the operator and boundary assumptions is not a device guarantee.

Extension and solution. If $t_1/t_2 = 1/2$, the probability beyond site $m - 1$ is 4^{-m}. This gives an exact truncation-error certificate for the semi-infinite state represented on a finite array.

4.18 EL18. Switching between stable subsystems and Floquet growth

Analytical references: [15, 32].

Problem. Let

$$A_1 = \begin{pmatrix} -1 & K \\ 0 & -1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} -1 & 0 \\ K & -1 \end{pmatrix}, \quad K > 0.$$

Each subsystem has eigenvalue -1. Alternate the two flows for duration τ each.

One-cycle propagator. Writing $q = K\tau$ and choosing the cycle order $e^{\tau A_1} e^{\tau A_2}$ gives

$$M = e^{-2\tau} \begin{pmatrix} 1 + q^2 & q \\ q & 1 \end{pmatrix}.$$

Its larger multiplier is

$$\lambda_+ = e^{-2\tau} \frac{2 + q^2 + q\sqrt{q^2 + 4}}{2}.$$

The system grows from cycle to cycle when $\lambda_+ > 1$. For $K = 4$, $\tau = 1$, the inner factor exceeds 17 and $17e^{-2} > 1$, proving growth even though both continuous subsystems are individually asymptotically stable.

Comparison. Checking subsystem eigenvalues misses the product geometry. Multiplying exact or certified phase propagators answers the periodic linear problem. A nonlinear converter additionally requires a periodic orbit and its variational equations, switching-time sensitivity and physical constraints.

Efficiency. For low-dimensional phase matrices, exact nilpotent formulas are inexpensive. For many phases, error propagation through matrix products must be bounded. Sparse exponentials and time-domain solvers remain conventional controls. The benefit is a correct stability object, not a blanket reduction in arithmetic.

Extension and solution. Reversing the two phases changes the matrix entries but preserves the nonzero eigenvalues of the two-factor product. State trajectories and peak norms can still differ. Floquet stability alone does not bound within-cycle transient amplification.

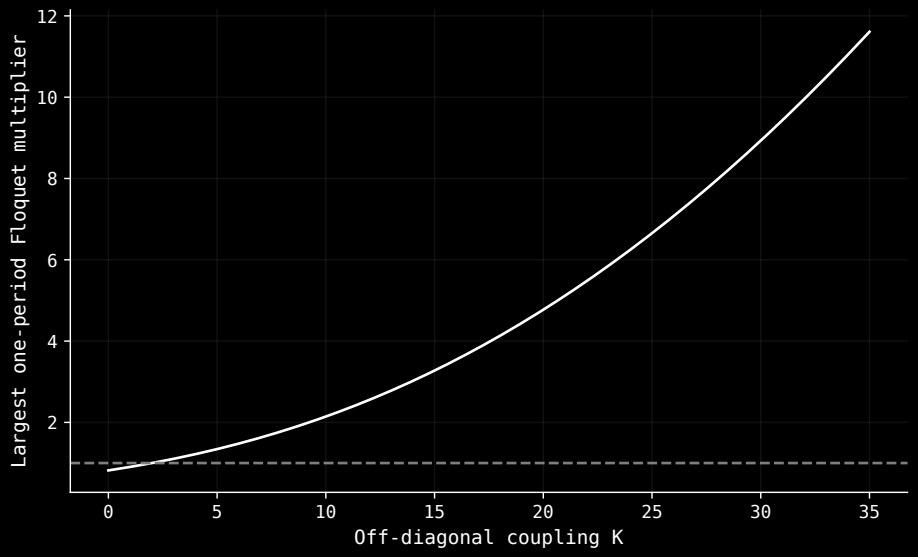


Figure 4.6: Largest one-cycle multiplier for the specified switched family. The unit level marks the boundary of discrete cycle-to-cycle growth.

4.19 EL19. Finite-state verification of a rounded recurrence

Analytical references: [21, 37].

Problem. On integer states $k \in \{-8, \ldots, 8\}$, let the next state be nearest-integer rounding of $3k/4$, with ties away from zero. This is a deterministic finite-state model of a scaled rounded contraction.

Exact analysis. The states $0, 1, 2$ are fixed: $3/4$ rounds to one and $3/2$ rounds to two. Their negatives are also fixed. For $|k| \geq 3$, direct inequalities or enumeration show that the absolute state decreases. Thus every orbit reaches one of five fixed points, and no longer cycle exists. The real-valued recurrence $x_{n+1} = 3x_n/4$ instead converges uniquely to zero.

Verification cost. There are 17 states and 17 transitions. A complete graph traversal proves the finite result with linear work in this explicit state count. For a d-dimensional b-bit model the state space can contain 2^{bd} elements, motivating symbolic invariants and reachability methods. A coefficient-root certificate does not remove this combinatorial growth.

Comparison. The analytic contraction identifies the intended unrounded dynamics. The finite transition relation identifies implemented behaviour under the declared rounding rule. Testing a few initial states can miss a basin. An exact state-space proof is possible here because the entire bounded universe is small.

Extension and solution. Rounding toward zero makes every nonzero state strictly decrease in magnitude, so zero becomes the unique fixed point. Changing one rounding convention changes the theorem. Saturation and overflow modes must be included explicitly when translating the model into a processor or HDL implementation.

4.20 EL20. Harmonic measure and optimal allocation of boundary observations

Analytical references: [15, 34].

Problem. An interior harmonic observable is approximated by $\widehat{u} = \sum_j w_j \widehat{g}_j$, with known quadrature or harmonic-measure weights. Estimate each boundary value using n_j independent measurements of variance σ_j^2, under $\sum_j n_j = N$.

Continuous optimum. For positive n_j, independence gives $\operatorname{Var} \widehat{u} = \sum_j w_j^2 \sigma_j^2 / n_j$. Cauchy–Schwarz yields

$$\left(\sum_j |w_j| \sigma_j \right)^2 \leq \left(\sum_j \frac{w_j^2 \sigma_j^2}{n_j} \right) \left(\sum_j n_j \right).$$

Equality is attained by $n_j = N|w_j|\sigma_j / \sum_k |w_k|\sigma_k$. The minimum variance is therefore $(\sum |w_j|\sigma_j)^2 / N$ over continuous allocations.

Geometry and efficiency. The analytic field representation supplies the weights, which can concentrate near the nearest boundary arc. Uniform sampling spends observations inefficiently when the weights or noise levels differ. The allocation theorem is a standard convex-resource result and assumes the weights, variances and independence are known.

Discrete and deterministic errors. Integer sample counts require rounding and a re-evaluation of the achieved variance. Boundary quadrature bias, model mismatch and correlated noise are separate terms. A variance optimum does not bound deterministic boundary error or the total mean-square error without those contributions.

Extension and solution. For two weights $(3/4, 1/4)$ and equal noise variances, the continuous optimal ratio is $3 : 1$. At $N = 100$, the variance is $\sigma^2/100$. Uniform allocation gives $5\sigma^2/400$, 25% larger. This is an exact variance comparison for the stated model, not a measured improvement in a physical sensor system.

Chapter 5

Operator factorization, inverse problems and singular limits

5.1 SP01. Wiener–Hopf factors and the boundary term of a unilateral shift

Analytical references: [21, 15].

Problem. Let S be the unilateral shift on $\ell^2(\mathbb{N}_0)$, $S(x_0, x_1, \ldots) = (0, x_0, x_1, \ldots)$, and let P_0 project onto the first coordinate. Then $S^*S = I$ and $SS^* = I - P_0$. Factor a half-line tridiagonal operator.

Exact factorization. For $|a|, |b| < 1$,

$$(I - bS^*)(I - aS) = (1 + ab)I - aS - bS^*.$$

Both factors are invertible by geometric operator series. The reverse product is

$$(I - aS)(I - bS^*) = (1 + ab)I - aS - bS^* - abP_0.$$

The rank-one discrepancy is an exact boundary effect. Replacing S by the bilateral shift removes it because both shift identities become the identity.

Complex interpretation. On the circle, the associated Laurent symbol factors as $(1 - bz^{-1})(1 - az)$. The factors correspond to opposite one-sided directions. Half-line compression retains a boundary term that pointwise commutative multiplication alone cannot detect. This is an elementary entry to Wiener–Hopf and Toeplitz theory.

Complexity comparison. Finite tridiagonal systems already admit linear-time direct solves. Factorization contributes analytic invertibility, norm bounds and an infinite-domain interpretation. It does not automatically beat the Thomas algorithm on a finite input. The inverse norm bound is at most $1/((1 - |a|)(1 - |b|))$.

Extension and solution. For $a = b = 1/2$, reversing the factors subtracts $P_0/4$. Applying both products to the first basis vector detects the discrepancy immediately. A periodic FFT model would impose a different boundary condition and cannot serve as an exact solution of the half-line problem.

5.2 SP02. Contour moments and the conditioning of pole reconstruction

Analytical references: [1, 5, 36].

Problem. Let $F(z) = r_1/(z - a_1) + r_2/(z - a_2)$ with distinct poles inside a positively oriented contour Γ and no poles on the contour. Residues give

$$m_k = \frac{1}{2\pi i} \int_\Gamma z^k F(z)\, dz = r_1 a_1^k + r_2 a_2^k.$$

Recover poles from m_0, m_1, m_2, m_3.

Recurrence. If $q(z) = z^2 - s_1 z + s_2$ with $s_1 = a_1 + a_2$, $s_2 = a_1 a_2$, then

$$m_{k+2} = s_1 m_{k+1} - s_2 m_k.$$

The first two instances form a two-by-two linear system. Its determinant has magnitude proportional to $|r_1 r_2 (a_1 - a_2)^2|$. Nonzero residues and separated poles yield identifiability; coalescence or a vanishing residue makes the inverse unstable or degenerate.

Worked instance. For poles one and two with residues one, $(m_0, m_1, m_2, m_3) = (2, 3, 5, 9)$. The recurrence gives $s_1 = 3$, $s_2 = 2$, hence $q(z) = (z - 1)(z - 2)$. A Vandermonde solve recovers the residues. Each algebraic step admits exact rational checking.

Comparison. When the poles are known, direct residue fitting is simpler. When the rational function itself is given symbolically, denominator factorization is a strong baseline. Moment reconstruction becomes useful when contour evaluations or transform samples are the available observations. Numerical integration, moment perturbation and nonlinear root conditioning all enter the error budget.

Extension and solution. For d simple poles, moments satisfy a recurrence of order d and a Hankel system determines its coefficients under a full-rank condition. Increasing d without a separation and noise model can fit spurious poles. A low algebraic residual is not a uniqueness certificate for noisy inverse data.

5.3 SP03. Discrete Cauchy–Riemann consistency on a square

Analytical references: [1, 5].

Problem. For square vertices $0, h, h + ih, ih$, compare the two diagonal quotients

$$D_1 f = \frac{f(h + ih) - f(0)}{h + ih}, \qquad D_2 f = \frac{f(ih) - f(h)}{ih - h}.$$

Equality is one discrete Cauchy–Riemann relation. Determine what it preserves.

Exact polynomial tests. For $f(z) = 1$, both quotients vanish. For $f(z) = z$, both equal one. For $f(z) = z^2$, both equal $h(1 + i)$. For $f(z) = z^3$, the first quotient is $2ih^2$ and the second is ih^2, so their difference is ih^2. Thus the restriction of every holomorphic function is not exactly discrete holomorphic under this stencil.

Analytical meaning. Taylor expansion about the cell centre shows a consistency error controlled by higher derivatives and mesh size. A convergence theorem additionally requires stability, admissible mesh geometry, boundary conditions and global compatibility. Testing a few monomials identifies local exactness but does not prove those global conditions.

Topological obstruction. A discrete potential integrated around a hole may accumulate a nonzero period. A spanning-tree integration defines values along chosen paths but cannot remove inconsistent loop sums. This is the discrete counterpart of the logarithmic period in annular harmonic problems.

Comparison. Finite-element and finite-difference solvers have their own consistency and stability theories. A discrete analytic formulation is useful when it preserves a relevant conjugacy or energy identity and improves the specified error or representation. It must be compared on the same mesh and boundary problem.

Extension and solution. For $f(z) = \bar{z}$, the first quotient is $-i$ and the second is i, giving a nonvanishing discrepancy $-2i$ independent of h. The stencil distinguishes this antiholomorphic linear map even as the mesh is refined.

5.4 SP04. Turning points and a uniform Airy representation

Analytical references: [25, 36, 38, 13].

Problem. Solve $y''(x) = \lambda^2 x y(x)$ for $\lambda > 0$ through $x = 0$. Put $\xi = \lambda^{2/3} x$. The chain rule gives $d^2y/d\xi^2 = \xi y$, so the solutions are linear combinations of $\mathrm{Ai}(\xi)$ and $\mathrm{Bi}(\xi)$.

Comparison with outer asymptotics. For positive large ξ, the decaying Airy solution has leading behaviour $\mathrm{Ai}(\xi) \sim (2\sqrt{\pi})^{-1} \xi^{-1/4} e^{-2\xi^{3/2}/3}$. The prefactor diverges as $\xi \downarrow 0$, while the exact Airy function is finite there. The outer expansion therefore cannot be extrapolated across the turning point.

Uniform core. The scaled special function treats both oscillatory and exponential regimes in one analytic solution. For this exact equation it is a closed-form representation. For a general equation $y'' = \lambda^2 q(x) y$ with a simple zero of q, a transformed Airy approximation has additional remainder terms requiring a separate uniform error analysis.

Costs. A robust Airy library is the conventional numerical baseline. Direct integration can be unstable when a growing solution contaminates the desired decaying branch. Scaled functions and boundary-value formulations address that issue. The complex asymptotic method supplies regime selection and conditioning information; no universal runtime advantage is asserted.

Extension and solution. The turning region has width $x = O(\lambda^{-2/3})$. Increasing λ compresses that region, so a fixed coarse grid can miss it. A mesh or basis adapted to this scaling follows from the differential equation before a numerical experiment is run.

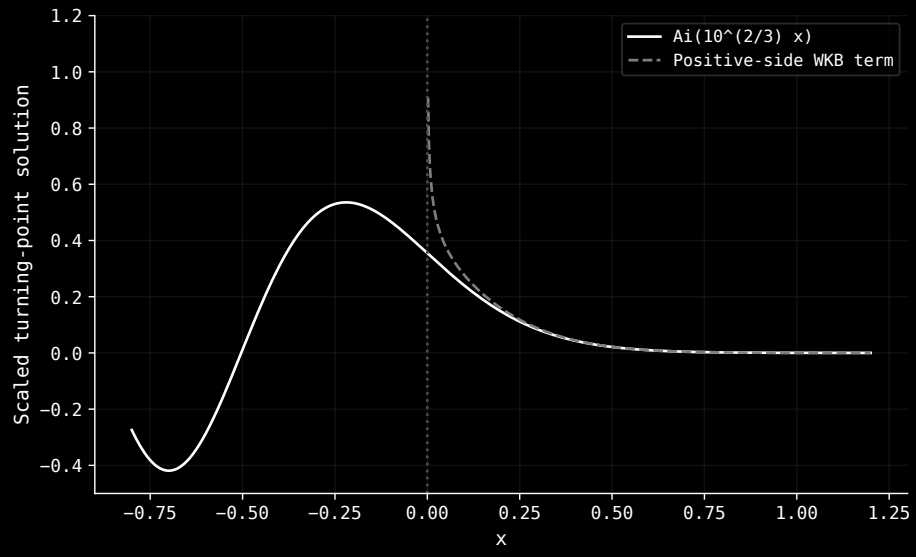


Figure 5.1: The Airy solution and its positive-axis outer asymptotic. The asymptotic approximation is not used at the turning point where its prefactor is singular.

5.5 SP05. Analytic Fredholm behaviour in an exactly finite-rank integral operator

Analytical references: [5, 15].

Problem. Define

$$(Kf)(x) = \alpha \int_{-1}^{1} f(t)\,dt + \beta x \int_{-1}^{1} t f(t)\,dt,$$

with nonzero real α, β. Solve $(I - zK)f = g$ as a function of complex z.

Invariant decomposition. The constant function and x are orthogonal. Their eigenvalues are 2α and $2\beta/3$, respectively; K vanishes on their orthogonal complement. Let $P_0 g = \frac{1}{2}\int g$ and $P_1 g = \frac{3}{2} x \int tg$. Then

$$(I - zK)^{-1} g = (I - P_0 - P_1)g + \frac{P_0 g}{1 - 2\alpha z} + \frac{P_1 g}{1 - 2\beta z/3}.$$

This is a meromorphic resolvent with at most two poles. At a pole, solvability depends on whether the corresponding component of g vanishes.

Complexity. A quadrature representation with n samples needs only two moments, hence $O(n)$ work and constant auxiliary scalar storage. A generic dense discretization and factorization costs far more but ignores the explicitly rank-two kernel. A strong low-rank baseline has the same advantage; the analytic expression adds pole locations and parameter-uniform error information.

Verification. Substitute the decomposition and the two eigenvalues. Orthogonal components remain unchanged, and each range component is divided by its scalar factor. This verifies the infinite-dimensional formula without inferring a theorem from a large sampled matrix.

Extension and solution. If $2\alpha = 2\beta/3$, the two pole locations coincide, but the resolvent still has a simple pole because the operator is diagonal on the range. A repeated characteristic value need not generate a higher-order pole; Jordan structure is the decisive additional datum.

5.6 SP06. Entire-function coordinates for the harmonic oscillator

Analytical references: [25, 39].

Problem. On entire functions define

$$\langle f, g\rangle = \frac{1}{\pi}\int_{\mathbb{C}} f(z)\overline{g(z)}e^{-|z|^2}\,dA(z).$$

Polar integration gives $\langle z^n, z^m\rangle = \delta_{nm}n!$. Hence $e_n = z^n/\sqrt{n!}$ is the normalized monomial family.

Operators. Multiplication by z gives $a^\dagger e_n = \sqrt{n+1}e_{n+1}$, while differentiation gives $ae_n = \sqrt{n}e_{n-1}$. On polynomials, $aa^\dagger - a^\dagger a = I$. The domains of these unbounded operators are part of their infinite-dimensional definition; polynomial identities do not establish boundedness on the full Hilbert space.

Finite truncation. Retain $e_0, \ldots, e_N$ and truncate the raising operator at the top. The finite matrices satisfy

$$[a_N, a_N^\dagger] = I - (N+1)|e_N\rangle\langle e_N|.$$

They cannot satisfy the exact infinite commutator because the trace of every finite commutator is zero. This supplies a precise boundary correction for numerical oscillator simulations.

Comparison. The Fock representation converts state operations into polynomial calculus and sparse ladder matrices. A direct position-grid discretization uses a different truncation and boundary model. Costs should be compared for an observable and state class, including tail probability and energy control. A beautiful entire-function plot is not a truncation certificate.

Extension and solution. For a state supported below e_N, the expectation of the finite commutator is one. A state with probability p_N in the top basis vector gives expectation $1-(N+1)p_N$. Thus a small top-mode population multiplied by $N+1$ determines this particular observable's error.

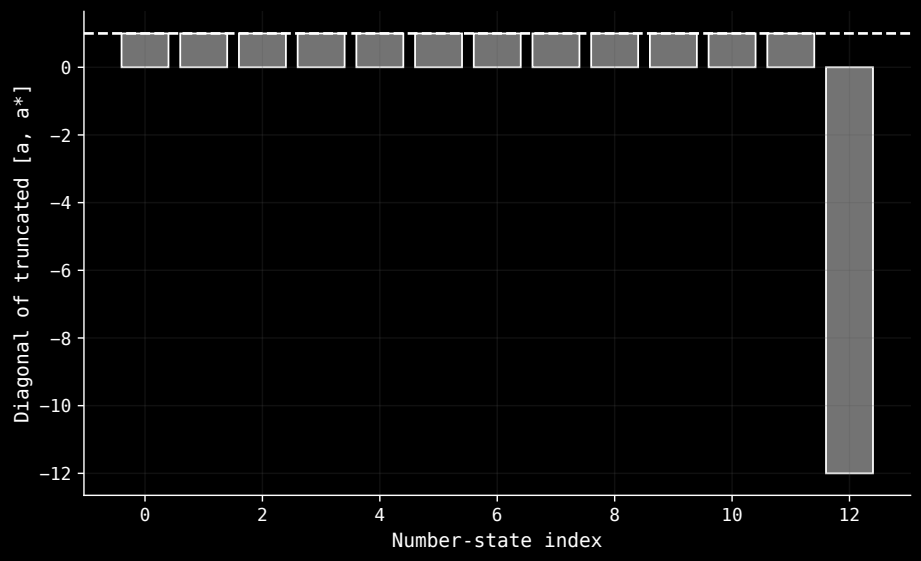


Figure 5.2: Selected normalized entire-function oscillator coordinates. The Gaussian measure and the finite-truncation correction are specified analytically in the text.

5.7 SP07. Optimal truncation of an asymptotic expansion

Analytical references: [36, 13].

Problem. For $\epsilon > 0$, consider

$$F(\epsilon) = \int_0^\infty \frac{e^{-t}}{1+\epsilon t}\, dt.$$

A formal geometric expansion suggests $\sum_{k\geq 0}(-\epsilon)^k k!$, which diverges for every fixed positive ϵ.

Exact remainder. The finite geometric identity gives

$$F(\epsilon) = \sum_{k=0}^{N-1}(-\epsilon)^k k! + (-\epsilon)^N \int_0^\infty \frac{t^N e^{-t}}{1+\epsilon t}\, dt.$$

Therefore the remainder has sign $(-1)^N$ and absolute value at most $\epsilon^N N!$. This is a rigorous one-sided alternating enclosure for every finite N. Consecutive term magnitudes have ratio $(N+1)\epsilon$, so the least-term region is near $N \approx 1/\epsilon$.

Complex interpretation. The formal Borel transform is $1/(1+t)$, with its pole on the negative axis. The positive-axis Laplace integral is unambiguous. Moving a Borel singularity onto the integration ray introduces a different lateral-continuation problem; that case is not silently included here.

Efficiency comparison. For small ϵ, a short optimally truncated expansion provides a useful approximation and an explicit error. For moderate ϵ, quadrature or a scaled exponential-integral routine is a stronger baseline. Summing all available terms eventually worsens the answer because the series is asymptotic, not convergent.

Extension and solution. At $\epsilon = 1/10$, truncation after ten terms has certified error at most $10!/10^{10} = 0.00036288$. The actual error can be smaller. The bound is an exact arithmetic statement and does not include cancellation or roundoff in evaluating the alternating partial sum.

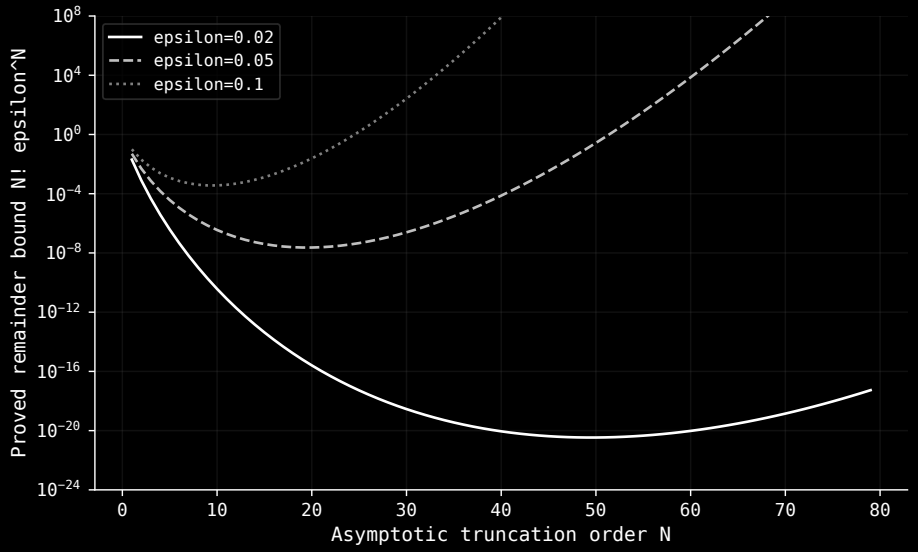


Figure 5.3: Term magnitudes and finite partial-sum behaviour for the stated asymptotic model. The least-term scale follows from the exact ratio test.

5.8 SP08. Jensen bounds as global zero-count information

Analytical references: [1, 9].

Problem. Let f be holomorphic on a neighbourhood of $\overline{D_R}$, with $f(0) \neq 0$ and $|f(z)| \leq M$ on $|z| = R$. Bound the number of zeros in $|z| \leq r < R$, including multiplicity.

Derivation. Jensen's formula gives

$$\sum_{|a_j|<R} \log \frac{R}{|a_j|} = \frac{1}{2\pi} \int_0^{2\pi} \log |f(Re^{i\theta})| \, d\theta - \log |f(0)|.$$

Boundary zeros can be handled by a limiting radius argument; a chosen zero-free boundary simplifies the formula. Each zero with $|a_j| \leq r$ contributes at least $\log(R/r)$. Consequently

$$N(r) \leq \frac{\log M - \log |f(0)|}{\log(R/r)}.$$

The integer floor yields a count bound. A very small central value weakens the estimate; it does not directly prove a nearby root.

Comparison. Rouché or argument-principle certificates can determine an exact count, while Jensen supplies an upper bound from coarse magnitude data. The latter can be valuable when boundary growth is easy to control but phase is difficult. A sampled maximum without an enclosure does not establish M.

Worked result. If $R = 2r$ and $M/|f(0)| < 4$, then fewer than two zeros can lie in $|z| \leq r$, so there is at most one. If f has a known zero of order m at the origin, factor z^m first and apply the bound to the remaining nonzero central value.

Extension and solution. The constant function achieves zero count for every radius although the displayed bound may be positive if M is loose. This demonstrates why a count upper bound is not an existence theorem. Combining it with a local existence certificate can establish uniqueness on a larger region.

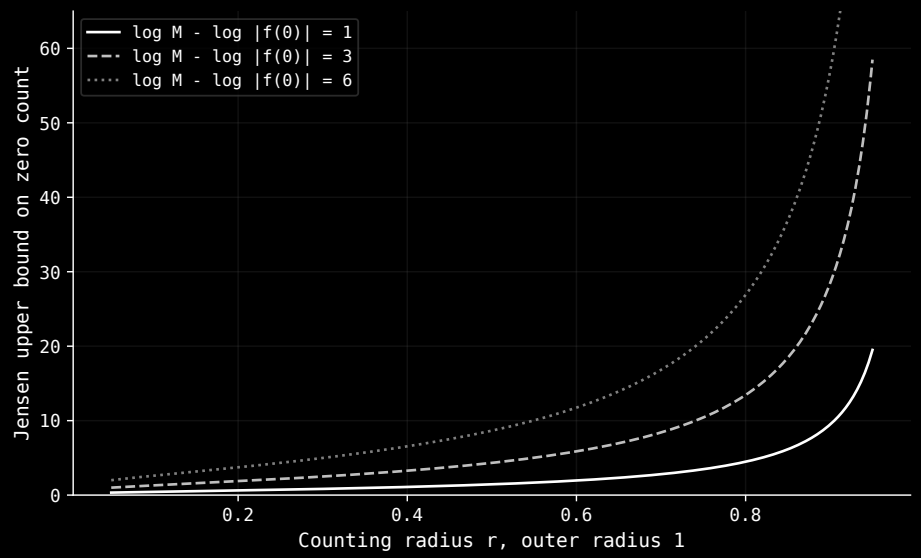


Figure 5.4: The Jensen upper bound as central-to-boundary growth and radius ratio vary. The plotted real bound must be converted to an integer statement with its strictness retained.

Chapter 6

Formal proof specifications

This chapter records 181 prospective developments of definitions, propositions, algorithms and certificate checkers. A development can contain several theorems and implementation obligations. The entries specify future acceptance conditions; their presence in this volume supplies no kernel-accepted proof. Mathematical hypotheses, checker semantics and implementation refinements are listed separately.

The identifiers are local to this register. Dependencies identify mathematical prerequisites, not a completed library or a mechanically checked dependency graph. Source interpretation must be reviewed before any formal statement is frozen. A counterexample is the appropriate completed result for a false assertion; an underdetermined task requires a parameterized formulation. Acceptance of a mathematical abstraction does not establish that an unexecuted physical or data-dependent experiment took place.

Each entry contributes a distinct specification. Common proof-integrity requirements apply throughout: fixed statements, explicit foundational assumptions, an actual successful checker run, and a reproducible dependency audit. The next chapter supplies written proofs for several finite governance and certificate principles used by these specifications.

6.1 Source integrity and proof governance

6.1.1 G01. Frozen identifiers and versioned ownership

Specification. Define a finite source-occurrence type; prove uniqueness, referential integrity and unambiguous ownership after namespace qualification.

Hypotheses. File version and edition are part of identity; repeated M01 labels from different volumes are distinct.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Every source occurrence resolves to one versioned owner; duplicate-ID and missing-owner counterexamples are rejected.

6.1.2 G02. Question and subpart obligation ledger

Specification. Define required subparts, baseline correctness, structural correctness, cost and implementation obligations for each selected question.

Hypotheses. The required set is fixed before proof counts are calculated; original statements must be human-adjudicated.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Prove that a question is complete exactly when every required obligation is closed; unresolved execution rows stay explicit.

6.1.3 G03. Source-to-formal-statement relation

Specification. Specify the correspondence among original wording, adopted corrections, formal domains and target conclusions.

Hypotheses. Source transcription and interpretive fidelity are reviewed premises, not consequences of type checking.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Store original and corrected assertions separately and require an identified reviewer or explicit unresolved status.

6.1.4 G04. Proof dependency closure

Specification. Prove graph reachability and acyclicity checks; calculate the transitive dependency and axiom set of each exported target.

Hypotheses. External libraries are pinned; an axiom audit complements the graph checker.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. No goal is discharged by a cycle, by assuming its conclusion, or by an undeclared imported claim.

6.1.5 G05. Nonvacuous hypotheses

Specification. Record an admissible instance or prove consistency for proposed theorem hypotheses when feasible; reject incompatible numeric premises.

Hypotheses. An implication with impossible hypotheses is logically true but does not solve the source problem.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. For every repaired task distinguish a constructive theorem, impossibility proof and conditional specification.

6.1.6 G06. Fixed-universe coverage

Specification. Prove 0 <= covered/required <= 1 and exact attainment when the finite covered set equals the declared obligation universe.

Hypotheses. Required set is nonempty; coverage is a count of specified obligations, not mathematical comprehensiveness.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Counts use unique identifiers; no inference to absolute maximality, research priority or learning outcomes.

6.1.7 G07. Duplicate-invariant metrics

Specification. Define metrics on canonical owners and prove invariance under repeated presentation, aliases and duplicate files.

Hypotheses. Semantic equivalence across different theorem statements requires an additional proof or human decision.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Repeating one result in nine branches changes presentation counts but leaves the canonical verified count unchanged.

6.1.8 G08. Premise-aware release gate

Specification. Specify separate states for written proof, executable test, checked certificate, accepted theorem and original-target experiment.

Hypotheses. A successful build does not certify source correctness; exact computation does not automatically yield a universal theorem.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Prove that the release gate cannot promote one state to another without the required witness.

6.1.9 G09. Extraction and numbering invariants

Specification. Prove selected parser functions preserve order, labels and subpart nesting under a specified document grammar.

Hypotheses. PDF extraction grammar is an abstraction; OCR, unknown fonts and reading-order errors remain input concerns.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Test and prove termination and conservation for the supported grammar; retain exception records outside the success claim.

6.1.10 G10. Archive and dataset accounting

Specification. Verify the master identity 1901-347+4351=5905, statement accounting, selected-ID membership and deterministic data transformations.

Hypotheses. Matching hashes establishes correspondence to a byte manifest, not source authenticity or collision-free mathematics.

Dependencies. None; finite records and lists

Verification route. Lean; Rocq; independent exact checker

Required evidence. Kernel-replay finite integer facts; keep byte comparison, hash computation and external authenticity separate.

6.2 Cantor coding, geometry and topology

6.2.1 C01. Finite words and affine coordinates

Specification. Prove A(wb)=3A(w)+2b, the concatenation identity, endpoint formulas and the corresponding affine-matrix representation.

Hypotheses. Bits belong to {0,1}; composition order and the empty word are fixed.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Theorems for all word lengths plus exact examples; no machine-integer overflow is hidden.

6.2.2 C02. Interval-tree separation

Specification. Prove interval nesting, sibling disjointness, coordinate bounds and depth-tagged injectivity.

Hypotheses. Intervals use the corrected endpoints; distinct depths are not silently identified.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Derive every finite invariant from the coordinate recurrence and prove exact rational certificate soundness.

6.2.3 C03. Infinite address map

Specification. Define the convergent ternary series, prove its tail estimate and show correspondence with the nested construction.

Hypotheses. Product topology, convergence radius and endpoint conventions are explicit.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove existence, uniqueness and continuity of the coding map, including endpoint cases.

6.2.4 C04. Classical Cantor properties

Specification. Prove nonemptiness, compactness, perfectness, total disconnectedness and zero Lebesgue measure.

Hypotheses. Infinite intersections and measure limits require their hypotheses, not finite pictures.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Each property has a theorem with all infinite-limit premises discharged.

6.2.5 C05. Quantitative approximation errors

Specification. Establish the distinct errors for interval unions, endpoint sets and finite sample representations.

Hypotheses. Hausdorff distance and the chosen approximating sets are specified separately.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove upper bounds and matching witnesses for each claimed exact error.

6.2.6 C06. Hausdorff dimension

Specification. Prove the upper-cover and lower-mass inequalities yielding dimension $\log(2)/\log(3)$.

Hypotheses. Use a probability measure with proved cylinder masses and a valid interval covering estimate.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Both dimension inequalities are checked; cylinder counting alone closes only the upper bound.

6.2.7 C07. Generalized interval trees

Specification. Formalize the mesh and branching hypotheses for compactness and the perfectness criterion.

Hypotheses. Finite branching, nondegenerate child intervals and uniform mesh decay are separately stated.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Include a counterexample when an essential hypothesis is omitted.

6.2.8 C08. Bounded geometry and thin/fat sets

Specification. Formalize the source comparison of bounded geometry and combinatorics and the explicit zero/positive-measure constructions.

Hypotheses. Summability or product conditions must match the construction; no converse is assumed without proof.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove the product/measure conclusions and the stated counterexamples.

6.2.9 C09. Cantor-space classification

Specification. Construct a homeomorphism using compatible clopen refinements.

Hypotheses. Nonempty compact metrizable perfect zero-dimensional spaces; distinguish this from an arbitrary totally disconnected space.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Construct inverse maps and prove continuity and mutual inverse identities.

6.2.10 C10. Ambient homeomorphism

Specification. Extend an order-compatible Cantor correspondence across complementary gaps to a homeomorphism of the real line.

Hypotheses. Order, endpoint matching and unbounded complementary components are handled.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove strict monotonicity, bijectivity and continuity of the extension and inverse.

6.2.11 C11. Inverse limits and observation algebras

Specification. Prove that compatible finite addresses recover the infinite space; identify locally constant functions and clopen Boolean operations.

Hypotheses. Finite-stage transition maps and topology are part of the specification.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove the universal compatibility maps and their inverse relation to the coding model.

6.2.12 C12. Homotopy and homology distinctions

Specification. Formalize constant paths, higher homotopy triviality and the source's separate singular and Cech degree-zero constructions.

Hypotheses. Fix homology/cohomology theories and coefficient rings; singular H0 is not identified with finite cylinder approximations.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Verify each claimed invariant in its own theory; prove any comparison maps separately.

6.2.13 C13. Rational membership certificates

Specification. Specify the rational-input membership algorithm, eventually periodic digit certificates and rejection witnesses.

Hypotheses. Input representation and both endpoint encodings are explicit; reduction to finite remainders must be proved.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove total correctness and replayable certificate soundness for all rational inputs in scope.

6.2.14 C14. Computable-real membership boundary

Specification. Formalize the precise representation-dependent impossibility result for total exact membership decisions.

Hypotheses. State Cauchy-name/oracle conventions and the permitted algorithms; the rational decision theorem remains distinct.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Give the reduction or continuity obstruction in the declared computation model.

6.2.15 C15. Persistent finite components

Specification. Prove the gap thresholds, component counts and inter-scale maps for the particular finite observations.

Hypotheses. Specify open/closed neighborhoods and equality cases at merging thresholds.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Prove all breakpoints and justify correspondence to the plotted barcode.

6.2.16 C16. Quadratic repeller coding

Specification. Construct the inverse branches for x^2-3, certified interval images and contraction coding.

Hypotheses. Rational square-root enclosures, orientation reversals and invariant interval are explicit.

Dependencies. G01–G05; finite words; real series; metric topology

Verification route. Lean/mathlib; Rocq; Isabelle/HOL

Required evidence. Derive an actual conjugacy or semiconjugacy with its exact domain and approximation bounds.

6.3 Functional analysis and finite-vector foundations

6.3.1 F01. Weighted finite-vector refinement

Specification. Define weighted inner products and prove that duplication under refined cell weights is an isometry.

Hypotheses. Weights are nonnegative and normalized; zero-weight coordinates use the stated quotient convention.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Exact finite proof for every refinement level and an implementation index correspondence.

6.3.2 F02. Completion from finite prefixes

Specification. Construct the completion of finite sequences under the specified norm and prove the geometric-tail example.

Hypotheses. A finite animation cannot establish completeness; distinguish c00, ell2 and ell-infinity.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Prove density, Cauchy convergence and uniqueness in the completed space.

6.3.3 F03. Cylinder-function approximation

Specification. Prove uniform approximation of continuous Cantor observables by finite cylinder functions.

Hypotheses. Compactness and a modulus of continuity or uniform continuity are supplied.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Obtain a constructive modulus-based error statement and the qualitative density theorem.

6.3.4 F04. Cantor probability and moments

Specification. Construct the product/pushforward measure and derive cylinder probabilities and moment recurrences.

Hypotheses. Measurability, countable additivity and normalization are explicit.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Finite moment calculations follow from the general integral identity.

6.3.5 F05. Conditional expectation as projection

Specification. Prove that cylinder averaging is the L2 orthogonal projection, with contraction and tower identities.

Hypotheses. A finite sigma-algebra and measure are fixed; expectations are defined up to almost-everywhere equality.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Orthogonality, idempotence, norm bounds and refinement compatibility all hold.

6.3.6 F06. Haar orthonormality and completeness

Specification. Prove normalized Haar orthonormality, scale-wise energy decomposition and completeness.

Hypotheses. Haar normalization follows the measure; finite orthogonality and infinite completeness are separate tasks.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Parseval and approximation theorems follow from proved density.

6.3.7 F07. Koopman and transfer adjoints

Specification. Formalize symbolic shift/composition and preimage averaging; prove adjoint and invariant-measure statements.

Hypotheses. Norm and invariant measure are specified; constants remain fixed.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Distinguish L2 contraction from the stronger seminorm estimates on variation.

6.3.8 F08. Convergence on a quotient by constants

Specification. Prove the qualified decay estimates for Lipschitz or variation seminorms and convergence to an invariant mean.

Hypotheses. State the function space, regularity and contraction coefficient; exclude an impossible strict contraction on constants.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Establish the advertised rate with the stationary subspace retained.

6.3.9 F09. Certified quadrature and optimization

Specification. Prove finite cylinder quadrature errors and the specific nonconvex optimization certificates.

Hypotheses. Regularity constants, feasible Cantor set and outward bounds are fixed.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Global lower/upper bounds and attaining or enclosing witnesses agree with the required tolerance.

6.3.10 F10. Hardy inner multiplication

Specification. Prove isometry of multiplication by a finite Blaschke product, model-space defect dimension and the declared index convention.

Hypotheses. H2, boundary values, multiplier boundedness and Fredholm definitions are developed.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Derive dimension/index from zeros with multiplicity and verify the sign convention.

6.3.11 F11. Finite group averaging

Specification. Prove deck-group averaging is a self-adjoint idempotent with image the invariant subspace.

Hypotheses. The representation is unitary; group finiteness and inner product are explicit.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Prove norm and invariance claims and instantiate the finite closure examples.

6.3.12 F12. Dirichlet energy pullback

Specification. Prove the degree multiplier for proper holomorphic pullbacks and conformal invariance in degree one.

Hypotheses. Integrability, multiplicity, change of variables and branch-set measure are justified.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Match domains, units and multiplicity factors in the mathematical and field applications.

6.3.13 F13. Hardy restriction optimality

Specification. Prove inf over rank-at-most-N bounded complex-linear A of norm(T_r-A) equals r^N.

Hypotheses. 0<r<1; the coefficient ell2 realization of H2 is established.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Prove both truncation upper bound and the finite-dimensional nullspace lower bound.

6.3.14 F14. Bergman and conformal radius

Specification. Prove the specified unitary pullback, evaluation norm and conformal-radius identities.

Hypotheses. Area normalization and the derivative multiplier are fixed; point evaluation is bounded.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Derive the kernel/evaluation formula without a missing Jacobian factor.

6.3.15 F15. Graph heat and dissipative semigroups

Specification. Prove exponential evolution, mass preservation and energy decay for the finite graph models.

Hypotheses. Laplacian sign, edge weights and connectivity are stated.

Dependencies. C01–C03; finite inner products; measure and integration

Verification route. Lean/mathlib; Rocq/Coquelicot; Isabelle/HOL

Required evidence. Connect spectral multiplication to the finite matrix exponential and certify truncation error if used.

6.4 Proper holomorphic maps and branched coverings

6.4.1 H01. Local holomorphic power model

Specification. Construct local coordinates in which a nonconstant holomorphic map is z raised to its zero multiplicity.

Hypotheses. Local nonvanishing factor, analytic root and inverse-function hypotheses are required.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove coordinate invertibility and preservation of local degree.

6.4.2 H02. Properness and global degree

Specification. Prove discrete fibres are finite and weighted fibre cardinality is constant on a connected target.

Hypotheses. Connected Riemann surfaces, nonconstant map and properness are retained.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove surjectivity and constancy with all compactness and multiplicity arguments.

6.4.3 H03. Finite-type ideal ends

Specification. Formalize extension over the specific finite-type compactification used in the source repair.

Hypotheses. Do not quantify over an undefined ideal boundary for arbitrary surfaces.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. A typed compactification, extension theorem and boundary case cover the adopted statement.

6.4.4 H04. Proper disk-map classification

Specification. Prove the finite Blaschke representation and its converse.

Hypotheses. Boundary behavior and zero multiplicities are justified from properness.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove zero removal, constant residual factor and properness of the constructed product.

6.4.5 H05. Half-plane rational models

Specification. Prove the positive-residue representation and its proper-map or passive-network interpretation under the source conditions.

Hypotheses. Pole locations, orientation and normalization are stated.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. The algebraic identity and positivity inequalities are proved for the complete domain.

6.4.6 H06. Cellular Riemann–Hurwitz

Specification. Prove Euler characteristic equals degree-scaled target Euler characteristic minus total ramification defect.

Hypotheses. Finite branched covering, compatible triangulation and every branch index are explicit.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Derive vertex/edge/face counts and sum defects with multiplicity.

6.4.7 H07. Divisorial Riemann–Hurwitz

Specification. Prove the pullback relation for canonical divisors and derive the genus equation.

Hypotheses. Compact smooth complex curves and canonical degree theorem are imported or proved explicitly.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Link the algebraic-geometric and cellular statements through compatible genus conventions.

6.4.8 H08. Bordered and punctured variants

Specification. Prove finite-type corrections and the branch-budget disk criterion.

Hypotheses. Boundary and puncture terms are included; classification of surfaces is a dependency.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Negative examples demonstrate why omitted topology invalidates the inference.

6.4.9 H09. Explicit algebraic covers

Specification. Verify normalized double covers, branch points including infinity, genus and discriminant calculations.

Hypotheses. Singular affine models are normalized before genus is inferred.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove finite algebra calculations and their connection to the actual smooth covering.

6.4.10 H10. Winding and local-to-global degree

Specification. Prove homotopy invariance, winding of power maps, value existence and the sum of local indices.

Hypotheses. Loops need not be simple; contour orientation and 2pi factors are fixed.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Topological degree and analytic integral agree under their shared hypotheses.

6.4.11 H11. Quantitative local surjectivity

Specification. Prove that a boundary-clearance bound forces a target ball into the image.

Hypotheses. Continuous boundary map and nonzero degree; no boundary zeros.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Convert the degree theorem into a uniform perturbation radius.

6.4.12 H12. Covering-space lifting criterion

Specification. Prove the basepoint-dependent subgroup criterion, uniqueness and the source's lift count.

Hypotheses. Domain connectedness/local path-connectedness and covering hypotheses are explicit.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Construct the lift, prove independence of path choices and verify its full domain.

6.4.13 H13. Branched lifting and reflection

Specification. Prove the qualified double-cover and real-symmetry cases; formalize counterexamples without the side condition.

Hypotheses. Branch locations, local degrees and chosen lift/basepoint are recorded.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. The reflected lift is identified by uniqueness only when the hypotheses permit it.

6.4.14 H14. Deck transformations and analytic quotients

Specification. Prove deck maps are holomorphic and quotient charts have the asserted local power structure.

Hypotheses. Proper discontinuity/finiteness and stabilizers are handled.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove the quotient topology, chart compatibility and covering properties.

6.4.15 H15. Normal core and Galois closure

Specification. Formalize subgroup cores, permutation monodromy, factorial bound and the explicit degree-six cubic closure.

Hypotheses. Covering classification is separate from finite group arithmetic.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove normality/index bounds and the actual commuting rational diagram.

6.4.16 H16. Nonregular orbifold counterexample

Specification. Verify the explicit covering is an orbifold covering while its deck action is not regular.

Hypotheses. Keep local orbifold data distinct from the global normality condition.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Prove both sides of the counterexample with a concrete monodromy computation.

6.4.17 H17. Typed telescoping and ramification composition

Specification. Prove nested pullback domains, admissibility at every stage and the ramification divisor composition law.

Hypotheses. Every intermediate target restriction is part of the map type.

Dependencies. N01,N02; topology of surfaces; finite groups; F10–F12

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with topology infrastructure

Required evidence. Derive the corrected nest and compositional error estimate with all domain guards.

6.5 Normal families, mapping and distortion

6.5.1 R01. Cauchy derivative and finite-jet bounds

Specification. Prove derivative, coefficient and truncation estimates including the scale-correct distance-to-boundary exponent.

Hypotheses. Holomorphy on the stated neighborhood and a justified uniform bound are required.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Include the source's incorrect-scale counterexample and the corrected estimate.

6.5.2 R02. Little Montel compactness

Specification. Prove locally bounded holomorphic families have locally uniformly convergent subsequences with holomorphic limit.

Hypotheses. Countable compact exhaustion, equicontinuity and limit-integral exchange are discharged.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Formalize the diagonal construction and invoke only established compactness theorems.

6.5.3 R03. Bounded conformal reduction

Specification. Construct the logarithm/root and fractional map reducing a proper simply connected domain to a bounded model.

Hypotheses. Domain is nonempty, proper and simply connected; branches are specified.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Prove injectivity, boundedness and the topology required for the branch.

6.5.4 R04. Univalent limits and Riemann mapping

Specification. Prove the limit alternative, extremal attainment and normalized conformal equivalence with the disk.

Hypotheses. Normalize at a point and derivative; compactness, zero exclusion and extremal improvement are justified.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Produce the bijection, holomorphic inverse and normalization uniqueness.

6.5.5 R05. Green rays and the degree-two map

Specification. Prove the explicit critical-point model and the transported Green/equipotential constructions.

Hypotheses. Branch, domain and Green normalization are fixed.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Verify the formula globally on the intended domain and identify exceptional points.

6.5.6 R06. Spherical compact-open metric

Specification. Construct the complete mapping-space metric and compare chordal and spherical conventions.

Hypotheses. Target is the Riemann sphere; compact exhaustion weights and metrics are explicit.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Prove metric axioms, completeness and equivalence with local uniform convergence.

6.5.7 R07. Escape and local normality

Specification. Prove holomorphic escape alternatives, local-to-global normality and the subsequence criterion.

Hypotheses. Distinguish convergence to infinity from a finite-valued holomorphic limit.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. All dichotomies are proved, including exclusions at poles and collisions.

6.5.8 R08. Marty and Big Montel

Specification. Prove spherical-derivative normality and the three-fixed-omitted-values theorem.

Hypotheses. Uniform local bounds and distinct omitted targets are essential.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Record the exact theorem statements and library dependencies; author-reported external developments require rebuild.

6.5.9 R09. Moving targets and collisions

Specification. Prove normality with uniformly separated moving values and the qualified isolated-collision result.

Hypotheses. Quantitative target separation is preserved; coincident values require the separate source formulation.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Verify Mobius normalization and a failure example when separation is lost.

6.5.10 R10. Normalized univalent compactness

Specification. Prove the omitted-circle point lemma, precompactness criteria and universal interior bounds.

Hypotheses. Normalizations are nondegenerate; limits that become constant are treated correctly.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Qualitative constants and later sharp quantitative constants remain distinct theorems.

6.5.11 R11. Area theorem and second coefficient

Specification. Prove the exterior area estimate and the derived coefficient inequality.

Hypotheses. Univalence, Laurent convergence, boundary approximation and area signs are justified.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. The coefficient bound follows from a complete area argument or an audited imported theorem.

6.5.12 R12. Koebe distortion and quarter theorem

Specification. Prove derivative/growth bounds and image-disk inclusion, with sharp extremal examples.

Hypotheses. Normalized univalence and distance/radius conventions are fixed.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Upper bounds, lower bounds and sharpness are separate checked targets.

6.5.13 R13. Capacity, collars and slit extremality

Specification. Prove condenser/capacity relations, collar-to-radius estimates and the radial-slit extremal argument.

Hypotheses. Full-continuum and annular-topology hypotheses; rearrangement approximation and equality conventions.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Discharge the capacity foundations before using the claimed global geometric conclusion.

6.5.14 R14. Injectivity under approximation

Specification. Prove coefficient-specific chart bounds and a global lower distance bound after controlled derivative perturbation.

Hypotheses. A convex source or another justified path bound and certified derivative lower margin are required.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Function error is converted to derivative error using a valid interior-radius estimate.

6.5.15 R15. Schwarzian projective invariant

Specification. Formalize the chain rule and the stated univalence criterion or bound.

Hypotheses. Nonzero derivative and exact criterion hypotheses; Mobius invariance is separate from global injectivity.

Dependencies. N01–N04; F03,F13,F14; H10; compact metric spaces

Verification route. Lean/mathlib; Isabelle/HOL; Rocq with analysis infrastructure

Required evidence. Prove algebraic identities and the analytic theorem connecting them to geometry.

6.6 Certified complex and numerical analysis

6.6.1 N01. Complex arithmetic and branches

Specification. Prove real/imaginary algebra, Cauchy–Riemann criteria, exponential identities and branch-restricted roots/logarithms.

Hypotheses. Complex-valued output alone does not imply holomorphy; domains exclude specified cuts and poles.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Every branch-sensitive operation has a precondition and an inverse/continuation property.

6.6.2 N02. Power series and Cauchy integrals

Specification. Formalize analytic expansions, coefficient extraction, residues, removable singularities and the needed contour estimates.

Hypotheses. Uniform convergence and valid contours are included; multiplicities and orientations are explicit.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Prove the exact integral identities and all exchanges of limits, derivatives and sums.

6.6.3 N03. Rouche and argument-principle certificates

Specification. Prove that accepted nonvanishing boundary certificates imply the correct root count with multiplicity.

Hypotheses. Domain, contour, pole subtraction and certificate-to-function connection are checked.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Finite checker acceptance is connected to a general complex-analysis soundness theorem.

6.6.4 N04. Rational interval arithmetic

Specification. Prove enclosure soundness for rational interval operations and the supported transcendental bounds.

Hypotheses. Rounding is outward; division excludes zero; invalid inputs produce failure, not a bound.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. All primitive operations and their composition preserve set containment.

6.6.5 N05. Bernstein whole-interval positivity

Specification. Prove conversion, convex-hull inclusion, subdivision and positivity of each accepted polynomial interval.

Hypotheses. Coefficients and parameter endpoints are exact; strict/non-strict inequalities match the claim.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. A proved checker replays every interval certificate and proves full union coverage.

6.6.6 N06. Whole-arc winding witnesses

Specification. Prove rational or Bernstein images remain in a chosen half-plane, including rotated witnesses and parameter families.

Hypotheses. Witness orientation and origin clearance are verified on the entire segment.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Derive homotopy-preserving polygon replacement and integer winding count.

6.6.7 N07. Checker termination and incompleteness

Specification. Prove termination for the fixed zero-free case and formalize the parameter-family counterexample to axis-only completeness.

Hypotheses. Pointwise zero-freeness differs from existence of a uniform witness in one restricted direction.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. A failed checker is not interpreted as a counterexample to the mathematical property.

6.6.8 N08. Validated root enclosures

Specification. Prove interval-Newton or separation-based enclosure conditions for the exact problem class.

Hypotheses. Derivative nonsingularity, containing box and contraction hypotheses are explicit.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. A small residual yields a forward error only through the proved inverse/separation bound.

6.6.9 N09. Newton iteration and exceptional behavior

Specification. Prove the repaired local error theorem, multiplicity-dependent rates and the explicit nonroot cycle.

Hypotheses. Differentiability order, root simplicity and denominator exclusion are specified.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Include valid local convergence domains and exact counterexamples to global claims.

6.6.10 N10. Cauchy differentiation and complex step

Specification. Prove derivative formula, truncation bounds and admissibility of analytic extensions used by code.

Hypotheses. Conjugation, absolute value, comparisons and branch switches need separate semantic handling.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Establish exact derivative and rounded evaluation error as separate claims.

6.6.11 N11. Analytic-strip quadrature and aliasing

Specification. Prove exponential trapezoidal error, periodic coefficient alias sums and finite sample conditions.

Hypotheses. Analytic strip width and boundary bound are quantified; an unseen singularity invalidates the estimate.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Bound both discretization and data error and connect to the specified returned integral.

6.6.12 N12. Chebyshev approximation

Specification. Prove ellipse-based coefficient decay and the explicit rational-function expansion.

Hypotheses. Ellipse lies strictly inside a pole-free region; endpoint/normalization conventions are fixed.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Prove a summable tail with a finite bound on that ellipse.

6.6.13 N13. Certified sampling-radius optimum

Specification. Prove the unique minimizer of the noise-plus-alias bound and replay the rational stationary-polynomial enclosure.

Hypotheses. Noise level, coefficient bound, sampling order and radius domain are declared.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Global uniqueness and endpoint sign certificates justify the optimum interval.

6.6.14 N14. Compositional error allocation

Specification. Prove the propagated error formula and minimum precision for the declared integer-allocation model.

Hypotheses. Costs and gains are fixed; integer bit variables and minimum precision are part of the feasible set.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Prove a global lower bound and an attaining allocation, not only an enumerated local candidate.

6.6.15 N15. Resolvent and matrix functions

Specification. Prove resolvent identities, spectrum exclusion and holomorphic functional calculus for the selected matrices/operators.

Hypotheses. Contour encloses the required spectrum and avoids poles and branch cuts; nonnormality is retained.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Derive the function or projector and a norm error certificate from the same domain assumptions.

6.6.16 N16. Near-defective analytic matrix models

Specification. Prove nilpotent and quotient-algebra exponential formulas and uniform series tails through eigenvalue coalescence.

Hypotheses. Finite-algebra relation such as A^2=lambda I is verified exactly.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Separate stable representation guarantees from runtime comparisons with a robust general method.

6.6.17 N17. Floating-point and fixed-point refinement

Specification. Prove implemented arithmetic refines the ideal calculation with the stated forward/backward error and range bounds.

Hypotheses. Format, rounding mode, subnormals, overflow, fused operations and compiler flags are fixed.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Combine machine-step bounds with the ideal theorem and reject executions outside range.

6.6.18 N18. Analytic continuation information bounds

Specification. Formalize indistinguishable data examples, unique continuation and minimax instability bounds.

Hypotheses. Observation model, admissible function class, noise ball and target norm are fixed.

Dependencies. G01–G08; rational arithmetic; real/complex analysis

Verification route. Lean/mathlib; Rocq/Interval/Gappa/Flocq; Isabelle/HOL; proof-replayed SMT

Required evidence. Exhibit witnesses proving a lower bound against every algorithm in the specified class.

6.7 Discrete mathematics, algorithms and exact computation

6.7.1 D01. Rational generating functions

Specification. Prove recurrence-to-series conversion, initial-condition numerators, shifts, differences, partial sums and repeated-pole coefficients.

Hypotheses. Formal power series avoid unsupported analytic convergence; analytic interpretations add their radius hypotheses.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Each selected numerical sequence follows by coefficient extraction and uniqueness of the recurrence.

6.7.2 D02. Fibonacci and exact rounding

Specification. Prove Binet identities, rounding margins, fast doubling and output bit-length bounds.

Hypotheses. Exact algebraic quantities and fixed-precision approximations are distinguished.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Show that the algorithm returns the required integer and state arithmetic versus bit complexity separately.

6.7.3 D03. Catalan and tree enumeration

Specification. Prove the bijection/recurrence, closed coefficient formula and stated asymptotic.

Hypotheses. Ordered trees, leaf/internal-node indexing and empty-tree conventions are fixed.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Prove exact counting before any asymptotic or computational shortcut.

6.7.4 D04. Join-order enumeration and optimization

Specification. Prove oriented labelled join counts and dynamic-programming bounds for the allowed plan class.

Hypotheses. Associativity/commutativity, relation labels and cost-lookup assumptions are explicit.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Count only distinct plans in the specified semantics and prove the DP attains the modeled minimum.

6.7.5 D05. Automata construction and minimization

Specification. Prove regular expression, NFA and DFA language equivalence, closure constructions and divisibility invariants.

Hypotheses. Alphabet, leading zeros, empty word and dead states are stated.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Verify all states or supply a general induction; minimality requires distinguishing suffixes.

6.7.6 D06. Transfer matrices and character filters

Specification. Prove accepted-word counts from matrix powers, rational resolvents and roots-of-unity filters.

Hypotheses. Count accepted words or paths according to the declared determinism/ambiguity conventions.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Connect the algebra to the exact combinatorial object and handle modular boundaries.

6.7.7 D07. Polynomial evaluation and interpolation

Specification. Prove Horner, synthetic division, uniqueness from values and the Vandermonde determinant.

Hypotheses. Field/ring assumptions and distinct interpolation nodes are required.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Prove total correctness, divisibility and evaluation identities together with the stated operation counts.

6.7.8 D08. Convolution and Cartesian sums

Specification. Prove linear/circular convolution, zero-padding and integer coefficient counting.

Hypotheses. Full support and multiplicities are preserved; circular aliasing is excluded where linear output is required.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Exact output equivalence plus a numerical reconstruction threshold when FFT arithmetic is used.

6.7.9 D09. DFT and radix recursions

Specification. Prove orthogonality, inversion, conjugate symmetry and radix-two/radix-three algorithms.

Hypotheses. Sign and normalization are source-specific; recursion length assumptions are checked.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. The full algorithm is derived from the finite sum and its exact operation recurrence.

6.7.10 D10. Bluestein and chirp transforms

Specification. Prove reduction to convolution, embedding length and the exceptional zero-parameter case.

Hypotheses. Roots and phase factors are chosen consistently; indexing and negative chirp positions are validated.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Return all requested transform outputs and prove the cost for every supported input length.

6.7.11 D11. Product and remainder trees

Specification. Prove polynomial reconstruction from roots, multipoint evaluation and all derivatives at one point.

Hypotheses. Multiplication/division costs are modeled; requested values and multiplicities are explicit.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Structural induction proves correctness and the recurrence proves the complexity bound.

6.7.12 D12. Finite-field and composite-ring transforms

Specification. Prove root-of-unity orthogonality and invertibility over each stated algebra.

Hypotheses. Primitive order alone can be insufficient in a ring; needed units/non-zero-divisors must be justified.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Kernel-check number-theoretic preconditions and reconstruction bounds.

6.7.13 D13. Matrix-chain optimal association

Specification. Prove the exact arithmetic counts, peak intermediate sizes and minimum among permitted parenthesizations.

Hypotheses. Matrix shapes and multiplication algorithm are fixed; views/copies are modeled separately.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. A finite DP or complete comparison proves optimality within the specified class.

6.7.14 D14. Data structures and amortization

Specification. Prove invariants of linked lists, queues, hashes, BST/AVL operations and exact semantics of undo/redo.

Hypotheses. Memory model, aliasing, resize policy and partial-operation preconditions are specified.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Functional correctness plus termination and a potential-function cost bound where claimed.

6.7.15 D15. Graph algorithm certificates

Specification. Prove BFS/DFS/SCC outputs, shortest-path potential certificates and MST cut/cycle optimality witnesses.

Hypotheses. Directedness, weights, disconnected components and negative-cycle assumptions match the task.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Verify certificates independently of the algorithm that produced them.

6.7.16 D16. Finite optimization and symmetry reduction

Specification. Prove DP recurrence optimality, bitonic tour result, orbit quotient soundness and counterexamples for nonintegral relaxations.

Hypotheses. The admissible path/plan space and objective are explicit; ILP integrality is not inferred from LP duality.

Dependencies. G01–G08; finite sums; polynomial algebra; explicit cost semantics

Verification route. Lean; Rocq/MathComp; Isabelle/HOL; Why3 or Dafny for programs

Required evidence. Prove a feasible attaining witness and a lower bound, or a valid obstruction.

6.8 Learning, kernels and statistical computation

6.8.1 L01. Gram kernels and feature maps

Specification. Prove positive semidefiniteness through explicit feature, sum, product and integral representations.

Hypotheses. Real/complex convention, kernel symmetry, integrability and positive weights are stated.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Cover every selected kernel with its own feature identity or a finite counterexample.

6.8.2 L02. Negative-definite kernels

Specification. Prove the conditional-zero-sum inequalities for distance kernels and distinguish them from positive definiteness.

Hypotheses. Metric exponent and Hilbert-space hypotheses are retained.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Verify two-point obstructions and the general valid inequality separately.

6.8.3 L03. Least-squares geometry

Specification. Prove normal equations, nullspace family, projection identity, minimum norm and unpenalized-intercept centring.

Hypotheses. Rank deficiency and uniqueness conditions are explicit.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Identify all minimizers and prove existence and any asserted uniqueness.

6.8.4 L04. Representer theorem and ridge duality

Specification. Prove projection to the data span and equivalence of feature-space and kernel-space regularized solutions.

Hypotheses. Loss depends only on declared observations; penalty monotonicity/coercivity and lambda positivity are retained.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Both formulations return identical predictions; coefficient differences are accounted for.

6.8.5 L05. Ridge perturbation and limits

Specification. Prove resolvent sensitivity, duplicate-predictor scaling, large-penalty directions and row-space reductions.

Hypotheses. Norms, data normalization and regularization conventions are fixed.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Prove the uniform bound or corrected limiting statement and expose unstable unregularized cases.

6.8.6 L06. PRESS and exact deleted fits

Specification. Prove the Schur-complement leave-one-out identity, denominator nonvanishing and applicable smoother variants.

Hypotheses. Training design, weights and regularization convention remain compatible after deletion.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Prove full refit equality for every deletion; the incorrect printed denominator has a checked counterexample.

6.8.7 L07. SVD and finite-rank optimality

Specification. Prove the selected eigenspaces, singular factorizations, PCA variance maxima and rank-constrained approximation errors.

Hypotheses. Spectral versus Frobenius norm and ties in singular values are explicit.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Upper bound from truncation and lower bound over all admissible competitors.

6.8.8 L08. Kernel PCA and centring

Specification. Prove empirical centring, normalized principal coordinates, out-of-sample formula and the population variational result.

Hypotheses. Finite matrices and infinite integral operators use separate existence/compactness assumptions.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Every coefficient and normalization agrees with the declared feature inner product.

6.8.9 L09. Nystrom as a projection

Specification. Prove CW-dagger-C* equals the projected Gram matrix, PSD residual and exactness under the landmark-span condition.

Hypotheses. Pseudoinverse and rank-deficient landmarks are handled.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Derive residual and exact storage formulas; random approximation guarantees need their own sampling theorem.

6.8.10 L10. Automatic differentiation semantics

Specification. Prove expression-graph evaluation, chain rule, forward/reverse accumulation and operation counts.

Hypotheses. Primitive derivatives and domain guards are supplied; nondifferentiable ReLU boundaries are explicit.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. The generated derivative equals the mathematical derivative on the stated admissible set.

6.8.11 L11. Complex gradients and implicit layers

Specification. Prove Wirtinger identities, implicit derivatives and branch-conditioned gradient bounds.

Hypotheses. Real-valued objectives can depend on conjugates; Jacobian invertibility and selected branch are required.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Do not use holomorphic differentiation across inadmissible operations or singular branches.

6.8.12 L12. Logistic and sparse estimators

Specification. Prove convexity/KKT conditions, selected closed forms, orthogonal ridge/lasso behavior and uniqueness conditions.

Hypotheses. Logistic separation, nonsmooth subgradients and chosen regularization are included.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. A computed iterate becomes certified only with a residual/error or optimization-gap bound.

6.8.13 L13. Finite empirical pipelines

Specification. Specify seeded samples, transformations, train/test separation, deleted-fit scoring and exact output summaries.

Hypotheses. Dataset bytes and random generator are fixed; observed accuracy and runtime are instance facts.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Prove implementation/scoring correctness and attach replay records without asserting population performance.

6.8.14 L14. Random projections and neighbor preservation

Specification. Formalize the distortion theorem and implications for distance ordering under explicit gaps; verify overlap-score implementation.

Hypotheses. Distribution, failure probability and dimension are specified; PCA and random projection have different hypotheses.

Dependencies. F05,F13; N15; finite-dimensional spectral theory; probability

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; exact linear-algebra checkers

Required evidence. Original-MNIST empirical overlaps remain unavailable until its stated images are supplied and executed.

6.9 Probability, information and wireless optimization

6.9.1 P01. Finite probability and concentration

Specification. Prove Bayes tables, variances of shared samples, binomial estimates, Chebyshev/Hoeffding bounds and their independence requirements.

Hypotheses. Conditional probabilities use positive conditioning mass; dependence is not erased.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Prove exact finite results and inequality hypotheses; statistical interpretations remain conditional on the model.

6.9.2 P02. Densities and transforms

Specification. Prove Gaussian/chi-square densities, characteristic-function identities and repeated-pole latency convolutions.

Hypotheses. Normalization, support, convergence strips and transform inversion hypotheses are fixed.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Verify density positivity/integral one and justify any limit interchange.

6.9.3 P03. Complex Gaussian structure

Specification. Prove covariance/pseudocovariance, properness, circular symmetry and unitary-invariance characterizations.

Hypotheses. Gaussianity is required for the relevant equivalences; dimension and variance normalization are explicit.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Establish the exact iff statement and handle improper laws separately.

6.9.4 P04. Matched filter and MMSE sufficiency

Specification. Prove projection-based sufficient statistics and the selected MMSE identities.

Hypotheses. Noise covariance, likelihood, signal model and independence are fixed.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Factor the likelihood and prove the conditional-distribution conclusion.

6.9.5 P05. Entropy and information inequalities

Specification. Prove discrete/differential maximum entropy, KL nonnegativity, conditional identities and equality cases.

Hypotheses. Density/moment existence, reference measures and finite entropies are included.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. All equality and singular-covariance cases are addressed; operational coding is an explicit deeper dependency.

6.9.6 P06. PageRank and stochastic resolvents

Specification. Prove existence, reachable support, ergodicity conditions, affine personalization and residual-to-error bounds.

Hypotheses. Fixed transition/dangling-node policy, restart probability and stochastic orientation are declared.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Prove mixture and contraction theorems; certify the nonmonotonicity counterexample.

6.9.7 P07. Queues and absorption

Specification. Prove the specified queue stationary/transform formulas and finite Markov absorption probabilities.

Hypotheses. Stability, independence, service moments and reachability hypotheses match the model.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Separate exact finite linear solves from infinite-state queueing theorems.

6.9.8 P08. Beamforming and covariance optima

Specification. Prove MISO total-power optimum, isotropic-input comparison and MIMO reciprocity.

Hypotheses. Gaussian channel, known channel state and allowed covariance constraints are fixed.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Produce feasible covariance and a global upper bound attained by it.

6.9.9 P09. Waterfilling

Specification. Prove finite-mode and joint time-frequency allocation, active-set formulas and uniqueness where valid.

Hypotheses. Gains, noise, total/peak constraints and measure-theoretic admissibility are fixed.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Verify KKT sufficiency by concavity; handle inactive/zero-gain modes and boundary budgets.

6.9.10 P10. Multiple-access regions

Specification. Prove corner points, sum/symmetric constraints and orthogonal/grouped-access comparisons for the selected channel models.

Hypotheses. Per-user versus aggregate power and CSI assumptions are explicit.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Theoretical rate inequalities are distinct from Monte Carlo estimates and their uncertainty.

6.9.11 P11. Toeplitz spectral limits

Specification. Prove the finite-memory determinant limit, spectral integral and exact two-tap continuant bound.

Hypotheses. Fixed memory, positive noise floor, trace normalization and limiting input spectrum are stated.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. General limiting theorem and finite special-case error each have complete dependencies.

6.9.12 P12. Power-control fixed points

Specification. Prove Perron–Frobenius feasibility, minimal feasible vector and synchronous/asynchronous convergence.

Hypotheses. Positivity, association rules, fairness and delay/update assumptions are explicit.

Dependencies. L03,L07; finite/infinite probability; integration; convexity

Verification route. Lean/mathlib; Rocq; Isabelle/HOL; PRISM for finite models

Required evidence. Convergence follows for every admissible schedule; arbitrary stale or unfair schedules are excluded.

6.10 Signals, circuits and implementation-aware physics

6.10.1 E01. Sampling and alias equivalence

Specification. Prove sampled-tone identities, period conditions, Nyquist endpoint ambiguity and two-period identification.

Hypotheses. Frequency units, phase and admissible band are fixed; a sample sequence does not determine an arbitrary analog signal.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Give an exact equivalence class and prove any claimed uniqueness within the stated band.

6.10.2 E02. Fourier invariants and Parseval

Specification. Prove time/frequency shifts, conjugation, symmetry, finite energy and real/imaginary spectral reconstruction.

Hypotheses. Absolute/L2 convergence and finite/infinite transform conventions are explicit.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Verify invariants for the full domain, including contradictory real-signal data.

6.10.3 E03. Z transforms and regions of convergence

Specification. Prove each rational transform with its support-dependent annulus and inverse residues/partial fractions.

Hypotheses. Pole locations, causality, sided support and annular boundary are included.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Two identical rational expressions with different ROCs remain distinct signals; impossible ROCs are refuted.

6.10.4 E04. LTI, causality and BIBO

Specification. Prove moving-average/difference properties, convolution operator norm and absolute-summability equivalence.

Hypotheses. Input/output norm, zero initial state and two-sided/unilateral indexing are fixed.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Prove both directions and any sharp norm using a valid sign/phase-aligned input.

6.10.5 E05. Transfer and state realizations

Specification. Prove direct, cascade and parallel realizations generate the same zero-state response and account for hidden modes.

Hypotheses. Initial states, controllability/observability and pole-zero cancellation are explicit.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Establish simulation/equivalence of state machines, not only equality of rational expressions.

6.10.6 E06. Bilinear and impulse-invariant mappings

Specification. Prove half-plane/disk stability mapping, prewarping, trapezoidal equivalence and the impulse-invariance obstruction.

Hypotheses. Sampling interval positive; frequencies and highpass distributional terms are handled.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Verify coefficients, transformed poles and the exact frequency mapping.

6.10.7 E07. Analog prototype design

Specification. Prove the adopted Butterworth/Chebyshev/elliptic order criteria and verify delivered coefficients/specifications.

Hypotheses. Prototype family and ripple/attenuation conventions are fixed; special-function formulas need proven prerequisites.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Feasibility, integer order rounding and family-specific minimality are separate targets.

6.10.8 E08. FIR continuous-band certification

Specification. Prove symmetric coefficient realization, exact polynomial amplitude formula and all passband/stopband inequalities.

Hypotheses. Delivered coefficient bytes are tied to exact rationals; frequency-to-polynomial substitution is sound.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Kernel-replay Bernstein certificates and establish coverage of every specified band.

6.10.9 E09. Symmetry and multiplier counts

Specification. Prove linear phase where valid and derive multiplication savings for symmetric FIR implementation.

Hypotheses. Odd/even length, symmetry sign, special taps and zero-phase sign changes are handled.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Match exact input/output behavior and count both additions and multiplications.

6.10.10 E10. Coefficient robustness and root location

Specification. Prove uniform unit-circle clearance and Rouche/Schur/Jury stability for the stated coefficient region.

Hypotheses. Strict margins, uncertainty correlations and leading-coefficient assumptions are retained.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Certify every point in the uncertainty set; nominal pole computations are insufficient.

6.10.11 E11. Fixed-point realizations and limit cycles

Specification. Define quantization/wrap/saturation and prove range safety, error bounds, invariant sets and cycle properties.

Hypotheses. Word length and rounding are explicit; coefficient stability is separate from state recurrence.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Exhaustive finite-state certificate or inductive invariant covers all admissible initial states.

6.10.12 E12. All-pass phase and spectral factorization

Specification. Prove magnitude preservation, delay/phase identities and minimum-phase spectral factor conditions.

Hypotheses. Fourier convention, causal stable factors and zeros on the boundary are handled.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Recover only information justified by magnitude data; prove the ambiguity class.

6.10.13 E13. Nyquist, delay and small gain

Specification. Prove orientation-correct stability criteria, delay-independent exclusions and norm-based feedback invertibility.

Hypotheses. Open-loop unstable poles, contour indentations, analytic domain and gain margin are explicit.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Connect the analytic theorem to whole-contour and norm certificates.

6.10.14 E14. Passivity, Schur and Pick feasibility

Specification. Prove positive-real/Schur transformations, finite Pick conditions, spectral factors and passive realization conditions.

Hypotheses. Model class, normalization, measurement uncertainty and minimality are fixed.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Finite matrix certificates are connected to the global interpolation/realization theorem.

6.10.15 E15. Circuit storage and hybrid switching

Specification. Prove power balance, storage inequalities and valid switching stability conditions; formalize the stable-subsystem counterexample.

Hypotheses. Constitutive laws, control switches, reset maps and initial conditions are explicit.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Check a common invariant/Lyapunov argument or a full hybrid proof, including switching restrictions.

6.10.16 E16. Two-dimensional fields and capacity

Specification. Prove harmonic potential, boundary data, energy/capacitance and gradient scaling under conformal coordinates.

Hypotheses. Homogeneity, dimension, voltage/length units and Jacobian factors are retained.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Model PDE correctness and bound-to-physical-unit translation are separate verified steps.

6.10.17 E17. Dispersive propagation and branches

Specification. Prove coupled transmission-line roots, Gaussian pulse dispersion, fractional impedance and causal transform identities.

Hypotheses. Branches must jointly satisfy propagation/impedance identities; constitutive and convergence hypotheses apply.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Prove positivity/causality in the stated complex domain and derived observable bounds.

6.10.18 E18. RF, optical and sampled-array identities

Specification. Prove antenna polynomial response, OFDM orthogonality/leakage and Bessel-mode truncation/cutoff enclosures.

Hypotheses. Array geometry, carrier offset, boundary conditions and units are fixed.

Dependencies. N01–N17; finite convolution; ODEs; units and state semantics

Verification route. Lean; Rocq/Interval/Flocq; Isabelle/HOL; KeYmaera X; CBMC/SymbiYosys

Required evidence. Give exact identities or certified intervals with explicit tail errors.

6.11 Specialist analytic and geometric extensions

6.11.1 X01. Scalar and matrix Riemann–Hilbert problems

Specification. Formalize Cauchy boundary operators, jump relations, scalar factors and small-norm inverse estimates.

Hypotheses. Contour regularity, orientation, function spaces and operator norm bounds are explicit.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. The selected worked problem is solved without assuming its solvability; general theory is a separate target.

6.11.2 X02. Wiener–Hopf and unilateral boundaries

Specification. Prove factorization and boundary correction for the chosen half-line convolution/shift model.

Hypotheses. Index, support and invertibility conditions are fixed.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Full-line multiplication is not substituted for the unilateral operator without the proved correction.

6.11.3 X03. Fredholm pole structure

Specification. Prove finite-rank operator determinant identities and the analytic Fredholm conclusion at the claimed level.

Hypotheses. Compactness/analyticity and existence of an invertible point for the general theorem are explicit.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Distinguish exactly finite-rank proof from an infinite-dimensional generalization.

6.11.4 X04. Graph zeta and partition zeros

Specification. Prove nonbacktracking transfer identities, zero-free logarithmic truncation and declared graph-count interpretations.

Hypotheses. Walk equivalence, backtracking, graph loops/multiple edges and zero-free region are specified.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Demonstrate that the invariant counts the intended objects and carries its claimed error bound.

6.11.5 X05. Separated Cauchy compression

Specification. Prove geometric moment expansion, translation identities, truncation error and construction/evaluation costs.

Hypotheses. Source-target separation and charge norms are explicit; one cluster is not a full hierarchical solver.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Generalize to a hierarchy only after proving partition, interaction-list and cumulative-error invariants.

6.11.6 X06. Contour moments and inverse reconstruction

Specification. Prove residue moments, Prony/Hankel recovery and separation-dependent conditioning.

Hypotheses. Distinct poles, multiplicities, number of samples and nonsingularity are stated.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Certify the forward moment identity and the inverse stability bound independently.

6.11.7 X07. Beltrami and discrete conformal models

Specification. Prove distortion identities, discrete CR consistency and compatible continuum limiting claims where developed.

Hypotheses. Mesh, boundary conditions, regularity and anisotropy are explicit.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Finite consistency does not imply continuum convergence without a stability/-convergence proof.

6.11.8 X08. Brownian harmonic representation

Specification. Prove the harmonic-measure or exit-distribution formula for the specified bounded-domain setting.

Hypotheses. Brownian construction, stopping/integrability and boundary regularity are necessary.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Probabilistic representation and any numerical sampling approximation receive separate theorems.

6.11.9 X09. Elliptic pendulum and AGM

Specification. Prove period integral, elliptic-curve relation, separatrix behavior and AGM enclosure algorithm.

Hypotheses. Energy/amplitude domain, branch and parameter convention are fixed.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Formalize invariant integral, convergence and explicit error enclosure.

6.11.10 X10. Stokes, Airy and asymptotic truncation

Specification. Prove factorial-term optimum, remainder bounds, exponentially small homogeneous ambiguity and the turning-point model.

Hypotheses. Sector, contour/lateral prescription, parameter domain and remainder class are stated.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Match the solution's global boundary data; formal coefficients alone do not identify it.

6.11.11 X11. Scattering and topological states

Specification. Prove the model's pole classifications, winding invariant and edge-state recurrence/localization.

Hypotheses. Hamiltonian, boundary conditions, self-adjointness or resonance continuation and gap hypotheses are fixed.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. A mathematical invariant is linked to the precise physical observable, not an unrestricted material claim.

6.11.12 X12. Bargmann–Fock and quantum evolution

Specification. Prove entire-function basis, ladder operators, domain conditions and unitary/polynomial propagation bounds.

Hypotheses. Unbounded operators require common dense domains and adjoint/closure conditions.

Dependencies. F and N modules; specialist Sobolev/spectral/probability foundations

Verification route. Lean; Isabelle/HOL; Rocq; independently replayed interval certificates

Required evidence. Separate formal commutator algebra from operator identities on actual function spaces.

6.12 Software, hardware and distributed-state specifications

6.12.1 S01. Programming-language state and objects

Specification. Specify copying, aliasing, mutation, exceptions and constructors/destructors for the supported code fragments.

Hypotheses. A source-language semantics or a documented executable subset is required.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Prove memory/state invariants and postconditions under normal and exceptional exits.

6.12.2 S02. Parsing, XML and compiler transformations

Specification. Prove lexer/parser language acceptance, AST construction and preservation through specified translations.

Hypotheses. Ambiguity, errors, tokens, encoding and grammar version are included.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Every accepted parse corresponds to a grammar derivation; rejection guarantees require completeness.

6.12.3 S03. Database semantics and dependencies

Specification. Prove relational rewrites, functional-dependency closure, lossless decomposition and aggregate preservation.

Hypotheses. Set versus bag semantics, nulls, keys, multiplicities and functional assumptions are fixed.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Verify equivalence for every admissible database, including counterexamples when assumptions fail.

6.12.4 S04. Persistent state and parallel reductions

Specification. Prove serialization round trips, audit update invariants and associative MapReduce aggregation.

Hypotheses. Encoding, key equality, overflow and failure/retry semantics are explicit.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Mathematical monoid properties are connected to the actual distributed implementation model.

6.12.5 S05. Concurrent protocols and schedules

Specification. Formalize mutual exclusion, queue/IO safety, deadlines and fair progress for selected finite-state models.

Hypotheses. Fairness, scheduler, communication failures and model size are fixed.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. TLC gives a finite-model result; unbounded parameterized claims need induction or a proof system.

6.12.6 S06. RTL and bit-vector circuits

Specification. Verify arithmetic units, FSM equivalence, pipeline/handshake invariants and reset/range properties.

Hypotheses. Clock/reset semantics, widths, sign extension and wrap/saturation are declared.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Record whether the result is bounded or inductive/unbounded and retain solver proof obligations.

6.12.7 S07. C numerical-kernel refinement

Specification. Verify loops, bounds, memory safety and equivalence of C implementations to finite mathematical functions.

Hypotheses. Compiler behavior, undefined behavior, alias rules and floating-point semantics are explicit.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Frama-C/Why3 obligations or bounded checks are linked to the exact implementation and scope.

6.12.8 S08. Legacy DSP practical abstraction

Specification. Prove host reference computations, scheduling recurrences and memory-placement constraints for the four practical records.

Hypotheses. Physical device traces, proprietary sessions and actual timing measurements remain external observations.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. A proved abstract model is not labelled as the completed original hardware laboratory.

6.12.9 S09. Security and native mathematical algebra

Specification. Formalize finite-field coding, named cryptographic arithmetic and declared access/-control invariants where specified.

Hypotheses. A protocol or security game must be explicit; no security claim follows from unrelated complex analysis.

Dependencies. G01–G08; finite state and operational semantics

Verification route. Lean; Rocq; Why3; Dafny; Frama-C/WP; CBMC; TLA+/TLC; SymbiYosys

Required evidence. Prove the specified property in its native semantics with threat/model assumptions documented.

6.13 Efficiency, data, figures and publication verification

6.13.1 V01. Equivalent-output comparison

Specification. Define identical input, output, accuracy and precomputation contracts for the baseline and proposed methods.

Hypotheses. Counts for a root list cannot be compared directly with counts for only the number of roots.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. Prove both methods satisfy the same postcondition before a preference label is allowed.

6.13.2 V02. Arithmetic, bit and memory cost

Specification. Prove exact counters, asymptotic bounds, peak live storage and preprocessing amortization.

Hypotheses. Unit-cost arithmetic, word-RAM and bit complexity are different models; output cost is included.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. A proved cost theorem cites its model and distinguishes setup, solve and certification work.

6.13.3 V03. Scoped optimality and break-even

Specification. Prove lower bounds and attaining constructions for finite rank, precision allocation, parenthesization and concave physical designs.

Hypotheses. The competitor class and objective are explicitly restricted and nonempty.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. A speedup against one implementation does not imply global algorithmic optimality.

6.13.4 V04. Benchmark summaries and uncertainty

Specification. Verify medians, ratios, iteration totals and confidence-interval calculations from frozen experimental records.

Hypotheses. Timing values are measured inputs; independence and sampling assumptions are explicit.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. Prove arithmetic of the report and implementation correctness without universalizing one machine's timing.

6.13.5 V05. Certified plot envelopes

Specification. Prove plotted analytic functions lie inside interval envelopes over each drawn segment or parameter cell.

Hypotheses. Pixel coordinates are obtained by a declared map; rounding and clipping are bounded.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. Certified mathematical envelopes are separated from unverified renderer behavior.

6.13.6 V06. Animation invariants

Specification. Prove each scene parameter satisfies algebraic/geometric invariants, branch continuity and any interpolation error bound.

Hypotheses. All-frame claims require a uniform parameter theorem or finite-frame enumeration with its exact scope.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. A valid geometric construction does not certify aesthetic quality or audience comprehension.

6.13.7 V07. Typography and layout predicates

Specification. Check finite text boxes, rule clearances, image bounds, links and glyph inventories against explicit thresholds.

Hypotheses. Font metrics and extracted bounding boxes are input measurements; actual rasterization has its own trusted components.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. Produce a finite certificate for declared layout constraints; no assertion that all viewers render identically.

6.13.8 V08. Figures-to-data provenance

Specification. Prove data selection, units, log scales, legends and rounding preserve their declared quantitative meaning.

Hypotheses. Underlying dataset and axes are fixed; misleading axis choices require semantic review.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. Every displayed number and curve points to its computation or proved enclosing formula.

6.13.9 V09. Document and branch traceability

Specification. Prove every lecture/result link resolves to the intended theorem or question and that prerequsite links obey the prerequisite graph.

Hypotheses. Pedagogical sufficiency and actual teaching remain empirical; graph membership is only structural.

Dependencies. G01–G10; cost semantics; interval enclosures

Verification route. Lean; Rocq; Isabelle/HOL; exact geometry/checker pipeline

Required evidence. No orphan theorem, omitted selected ID or duplicate canonical count survives the gate.

Chapter 7

Finite certificate semantics and publication evidence

The preceding register describes a prospective formalization programme. This chapter develops a finite mathematical basis for several of its acceptance conditions. The propositions below are written proofs. They are not accompanied by newly accepted proof-assistant objects. The distinction is material: the syntax accepted by a compiler, the proposition checked by its kernel, and the intended source assertion are separate objects of review.

7.1 Versioned identity and canonical ownership

Let S be a finite set of source occurrences. An occurrence can be a theorem, an exercise subpart, a coefficient array, or a particular numerical experiment. Fix a map

$$c : S \longrightarrow K,$$

where K is a set of keys containing the work, edition, local identifier and statement version. Equality in K is a declared identity relation. It must not be inferred solely from a title shared by two files. Conversely, different keys do not imply mathematical independence: two separately attributed exercises may ask for the same calculation.

A canonical owner map assigns each public mathematical statement to one location. Its applications can refer to that location while adding different input data, hypotheses or output requirements. The following elementary result specifies the behavior expected of a duplicate-insensitive coverage measure.

Proposition 7.1 (Duplicate invariance)**.** *Let $S' = S \cup T$, and let $c' : S' \to K$ restrict to c on S. If $c'(T) \subseteq c(S)$, then $c'(S') = c(S)$.*

Proof. For $s \in S$, the equality $c'(s) = c(s)$ gives $c(S) \subseteq c'(S')$. For the reverse inclusion, write $k = c'(x)$ with $x \in S'$. If $x \in S$, then $k = c(x) \in c(S)$. If $x \in T$, the stated inclusion gives the same conclusion. These cases exhaust S', proving equality. □

This proposition concerns the set of represented keys. It does not justify a particular canonicalization algorithm. For example, deleting whitespace from two formulas is not a proof of equivalence if the notation uses whitespace to delimit tokens. A correct normalization must have its own semantics-preservation argument.

Concrete instance. Suppose three files contain references to keys $\{a, b, c\}$, $\{b, c\}$, and $\{c\}$. The number of files is three and the sum of per-file reference counts is six. The represented-key set has cardinality three. A claim of six distinct proved statements would use the wrong counting object. The owner of c can be the first file; the other files then contribute applications or alternative derivations, not additional instances of statement identity.

7.2 Questions, subparts and positive-weight coverage

Let Q be a finite set of questions. For each $q \in Q$, let U_q be a finite nonempty set of required obligations. The disjoint union

$$\mathcal{U} = \{(q, u) : q \in Q,\ u \in U_q\}$$

prevents identically numbered subparts in different questions from being conflated. Fix an accepted-obligation set $A \subseteq \mathcal{U}$. Define

$$\mathrm{Complete}_A(q) \quad \Longleftrightarrow \quad \{q\} \times U_q \subseteq A.$$

The word "accepted" here is parameterized by an explicit acceptance predicate. It can refer to a written-review gate, a kernel-proof gate, or an execution gate. Those predicates should have separate accepted sets.

Proposition 7.2 (Monotonicity and full completion)**.** *If $A \subseteq B \subseteq \mathcal{U}$, then* $\mathrm{Complete}_A(q)$ *implies* $\mathrm{Complete}_B(q)$*. Moreover, every question is complete with respect to A if and only if $A = \mathcal{U}$.*

Proof. The first assertion follows by transitivity of set inclusion. If every question is complete, take any $(q, u) \in \mathcal{U}$. Then $u \in U_q$, so $(q, u) \in \{q\} \times U_q \subseteq A$. Thus $\mathcal{U} \subseteq A$; the reverse inclusion is an assumption. If $A = \mathcal{U}$, every $\{q\} \times U_q$ is a subset of A by construction, proving the converse. □

Assume now $\mathcal{U} \neq \varnothing$, and assign strictly positive weights $w : \mathcal{U} \to \mathbb{R}_{>0}$. Define

$$C(A) = \frac{\sum_{u \in A} w(u)}{\sum_{u \in \mathcal{U}} w(u)}.$$

Weights are fixed before evaluation. Altering them after observing which obligations were completed changes the metric.

Proposition 7.3 (Sharp bounds for a fixed coverage universe)**.** *For every $A \subseteq \mathcal{U}$, one has $0 \leq C(A) \leq 1$. Furthermore, $C(A) = 1$ if and only if $A = \mathcal{U}$. If $A \subsetneq B \subseteq \mathcal{U}$, then $C(A) < C(B)$.*

Proof. The denominator is positive because $\mathcal{U}$ is nonempty and all weights are positive. The numerator is nonnegative. Decomposing $\mathcal{U}$ into the disjoint sets A and $\mathcal{U} \setminus A$ gives

$$\sum_{\mathcal{U}} w = \sum_{A} w + \sum_{\mathcal{U} \setminus A} w \geq \sum_{A} w,$$

which proves the upper bound. If $A \neq \mathcal{U}$, at least one positive term occurs in the second sum, making the inequality strict. The equality characterization follows. Replacing $\mathcal{U} \setminus A$ with $B \setminus A$ proves strict monotonicity for $A \subsetneq B$. □

The theorem proves completeness relative to a declared obligation universe. It supplies no claim that the universe contains every useful theorem, every interpretation of a problem, or every possible experiment. In particular, a document's page count is absent from the definition.

7.3 Dependency closure and assumption propagation

Consider a finite directed graph $G = (V, E)$. An edge $u \to v$ means that the proof of v uses u. Assign each vertex a nonnegative integer rank $r(v)$.

Proposition 7.4 (A sufficient acyclicity witness)**.** *If $u \to v$ implies $r(u) < r(v)$, then G contains no directed cycle.*

Proof. A directed cycle $v_0 \to v_1 \to \cdots \to v_k = v_0$ would imply

$$r(v_0) < r(v_1) < \cdots < r(v_k) = r(v_0),$$

contradicting irreflexivity of the strict order on integers. □

Suppose a topological order $v_1, \ldots, v_n$ is supplied and checked, so every dependency of v_i has smaller index. Let $L(v_i)$ denote the locally introduced assumptions. Set

$$A(v_i) = L(v_i) \cup \bigcup_{u \to v_i} A(u).$$

The interpretation of an assumption is not determined by this set operation: foundational axioms, physical premises and unproved target statements must carry different types.

Proposition 7.5 (Transitive assumption characterization)**.** *For every vertex v, $A(v)$ is the union of $L(u)$ over all vertices u that equal v or have a directed path to v.*

Proof. Induct on the topological position of v. If v has no predecessor, the recursive expression is $L(v)$, which is precisely the asserted union. For the induction step, every strict ancestor of v has a path ending with an edge $u \to v$, and its earlier vertices belong to the ancestor set of u. By the induction hypothesis, their local assumptions occur in $A(u)$. Thus every assumption in the asserted ancestor union occurs in the recursive definition of $A(v)$. Conversely, an element of $A(v)$ belongs either to $L(v)$ or to some $A(u)$ with $u \to v$. In the latter case the induction hypothesis gives an ancestor of u, and appending $u \to v$ makes it an ancestor of v. This proves the reverse inclusion. □

Acyclicity and assumption propagation can be checked on a finite manifest. Their sound application requires the manifest to contain the actual proof dependencies. A manually incomplete graph can pass both graph checks while concealing a proof assumption.

7.4 Nonvacuity and an attaining optimum

For a hypothesis H and conclusion P, a proof of $H \Rightarrow P$ does not establish that H is satisfiable. An inconsistent set of physical or numerical requirements may entail any conclusion. A useful optimization claim therefore records feasibility separately.

Proposition 7.6 (An attained global bound)**.** *Let $F \neq \varnothing$, let $J : F \to \mathbb{R}$, and let $L \in \mathbb{R}$. If $J(x) \geq L$ for every $x \in F$, and $x_* \in F$ satisfies $J(x_*) = L$, then x_* is a global minimizer and $\inf_{x \in F} J(x) = L$.*

Proof. For each $x \in F$, the assumed bound yields $J(x) \geq L = J(x_*)$, which is the definition of global minimality. The value L is a lower bound for $J(F)$, and $L = J(x_*)$ belongs to $J(F)$. Every lower bound of a set is no larger than any member of that set, so no lower bound exceeds L. Consequently L is its greatest lower bound. □

The operator-rank and precision-allocation results in Volume I use this pattern with different feasible classes. The feasible class is part of the result: bounded linear approximants, integer precision allocations, passive functions and covariance matrices are not interchangeable competitor sets.

7.5 Sound checkers and composition

Let D be a set of problem instances, C a set of finite witnesses, and $\mathrm{check} : D \times C \to \{0, 1\}$ an executable function. For a semantic property $P \subseteq D$, soundness is

$$\forall d, c, \qquad \mathrm{check}(d, c) = 1 \Longrightarrow P(d).$$

Completeness would assert that every d with $P(d)$ admits some accepted witness. Soundness and completeness are logically independent requirements.

Proposition 7.7 (Sequential certificate composition)**.** *Suppose*

$$\mathrm{check}_1(d, c_1) = 1 \Longrightarrow P(d)$$

and

$$P(d) \ \wedge \ \mathrm{check}_2(d, c_2) = 1 \Longrightarrow Q(d).$$

Then accepting both checks implies $Q(d)$.

Proof. Acceptance of the first check supplies $P(d)$. Together with acceptance of the second check, this supplies both premises of its soundness implication, whose conclusion is $Q(d)$. □

For a contour root count, the first checker can establish that the polynomial has no zero on the contour. The second can establish the winding number, with the semantic argument principle connecting that integer to the number of interior roots. The contour orientation, multiplicities and domain remain premises of the second implication.

A completeness obstruction. A checker restricted to separation by horizontal or vertical half-planes may reject a nonvanishing contour segment whose image fits a rotated half-plane. Its rejection does not disprove nonvanishing. Adding rotated witnesses extends the accepted certificate class while leaving the intended semantic property unchanged. Soundness of the enlarged checker must be proved for the new witness type.

7.6 Exact polynomial normalization

Fix a number d of variables. A sparse rational polynomial is a finitely supported map $a : \mathbb{N}^d \to \mathbb{Q}$. For x in a commutative $\mathbb{Q}$-algebra, define

$$\mathrm{ev}_x(a) = \sum_{\alpha} a_\alpha x^\alpha, \qquad x^\alpha = \prod_{j=1}^{d} x_j^{\alpha_j}.$$

All sums are finite. Define addition coefficientwise, and multiplication by

$$(a * b)_\gamma = \sum_{\alpha+\beta=\gamma} a_\alpha b_\beta.$$

Proposition 7.8 (Evaluation preserves the arithmetic)**.** *One has*

$$\mathrm{ev}_x(a + b) = \mathrm{ev}_x(a) + \mathrm{ev}_x(b), \qquad \mathrm{ev}_x(a * b) = \mathrm{ev}_x(a)\,\mathrm{ev}_x(b).$$

Proof. For addition, substitute the coefficientwise definition and distribute the finite sum. For multiplication, commutativity gives $x^{\alpha+\beta} = x^\alpha x^\beta$. Substitute the convolution coefficients and regroup the finite sum by the pair (α, β):

$$\sum_\gamma \sum_{\alpha+\beta=\gamma} a_\alpha b_\beta x^\gamma = \sum_{\alpha,\beta} a_\alpha x^\alpha b_\beta x^\beta = \left(\sum_\alpha a_\alpha x^\alpha\right)\left(\sum_\beta b_\beta x^\beta\right).$$

Finite distributivity justifies every rearrangement. □

Structural induction on an expression built from variables, rational constants, addition and multiplication now shows that normalization preserves evaluation. If normalization of the difference of two expressions produces the empty coefficient map, their evaluations coincide in every such algebra.

Division requires an additional domain condition. Clearing a denominator $q(x)$ proves a rational identity only where $q(x) \neq 0$. A normalization checker that forgets this premise would accept statements at points where the original expressions are undefined.

7.7 Interval enclosure and reciprocal domains

A real interval $[a, b]$ is valid when $a \leq b$. For $x \in [a, b]$ and $y \in [c, d]$,

$$x + y \in [a + c, b + d].$$

The claim follows by adding the two lower bounds and the two upper bounds.

Proposition 7.9 (Corner enclosure for interval multiplication)**.** *Let*

$$m = \min\{ac, ad, bc, bd\}, \qquad M = \max\{ac, ad, bc, bd\}.$$

Then $xy \in [m, M]$ for every $x \in [a, b]$, $y \in [c, d]$.

Proof. If $a < b$, write $x = (1 - s)a + sb$ with $s = (x - a)/(b - a) \in [0, 1]$. If $a = b$, the same representation holds with $s = 0$. Similarly write $y = (1 - t)c + td$ with $t \in [0, 1]$. Expanding the product expresses xy as a convex combination of the four corner products, with weights

$$(1 - s)(1 - t), \quad (1 - s)t, \quad s(1 - t), \quad st.$$

They are nonnegative and sum to one. Each corner lies between m and M; multiplying these inequalities by the weights and summing proves the enclosure. □

On either interval $(0, \infty)$ or $(-\infty, 0)$, the reciprocal function is strictly decreasing. For $0 \notin [a, b]$, it follows that

$$\{1/x : x \in [a, b]\} \subseteq [1/b, 1/a].$$

To verify the order directly, take $a \leq x \leq b$ with a common nonzero sign. The products ax and bx are positive, so $1/a - 1/x = (x - a)/(ax) \geq 0$ and $1/x - 1/b = (b - x)/(bx) \geq 0$. An interval containing zero must be rejected or split with a separately justified exclusion.

For complex rectangles, use $(a + ib)(c + id) = (ac - bd) + i(ad + bc)$ and apply the real operations to its components. Dependency overestimation can widen the enclosure, but does not invalidate inclusion. A failed strict inequality from a wide enclosure supplies no negative semantic conclusion.

7.8 Bernstein coefficients and continuous-band certificates

For an integer $n \geq 0$, define

$$B_{k,n}(t) = \binom{n}{k} t^k (1-t)^{n-k}, \qquad 0 \leq k \leq n, \quad 0 \leq t \leq 1.$$

Each basis function is nonnegative. The binomial theorem yields $\sum_{k=0}^{n} B_{k,n}(t) = (t + 1 - t)^n = 1$.

Proposition 7.10 (Bernstein convex-hull bound)**.** *If $p(t) = \sum_{k=0}^{n} b_k B_{k,n}(t)$ and $\ell \leq b_k \leq u$ for every k, then $\ell \leq p(t) \leq u$ for every $t \in [0, 1]$.*

Proof. Multiply the lower coefficient inequality by $B_{k,n}(t) \geq 0$ and sum. The result is $p(t) \geq \ell \sum_k B_{k,n}(t) = \ell$. The upper inequality follows identically from $b_k \leq u$. □

An affine coordinate $x = a + (b - a)t$ transports this theorem to $[a, b]$. To certify a polynomial on several subintervals, the checker must establish both the coefficient inequalities and complete interval coverage.

Proposition 7.11 (Adjacent subinterval coverage)**.** *If $a = t_0 \leq t_1 \leq \cdots \leq t_m = b$, then $\bigcup_{j=1}^{m} [t_{j-1}, t_j] = [a, b]$.*

Proof. Induct on m. The assertion for $m = 1$ is immediate. For the step, the induction hypothesis gives $[a, t_{m-1}]$ for the first $m - 1$ intervals. Its union with $[t_{m-1}, b]$ is $[a, b]$: a point $x \in [a, b]$ lies in the first interval when $x \leq t_{m-1}$, and in the second otherwise. Every point of the union is between a and b, proving equality. □

For the FIR certificates in Volume III, frequency substitution introduces algebraic interval endpoints and an exact polynomial derived from the delivered coefficient array. Endpoint enclosures, coefficient interpretation, basis conversion and covering are all independent checker obligations. A sampled frequency response does not replace any of them.

7.9 Cost semantics and a break-even theorem

Attach a natural-number cost to every elementary operation in an abstract machine. A procedure returns an output and the sum of costs along its executed path. Sequential composition adds costs. A conditional counts only its selected branch. A terminating loop has a cost recurrence obtained by adding its iteration cost to the remaining execution.

Induction on the number of loop iterations proves equality between the counter and this recurrence, provided termination and the meaning of each counted operation are already specified. Counting multiplications in exact arithmetic leaves integer bit lengths and memory traffic outside the model. Those require additional components.

Proposition 7.12 (Repeated-evaluation break-even point)**.** *Let $s \geq 0$, $b > p \geq 0$, and let m be a positive integer. For*

$$T_{\text{direct}}(m) = bm, \qquad T_{\text{prepared}}(m) = s + pm,$$

strict preference for the prepared method holds exactly when

$$m > \frac{s}{b - p}.$$

The least positive integer satisfying this condition is $\lfloor s/(b-p) \rfloor + 1$.

Proof. Subtracting pm from both sides of $s + pm < bm$ gives $s < (b - p)m$. Division by the positive number $b - p$ gives the displayed condition. For any real $x \geq 0$, the defining inequalities $\lfloor x \rfloor \leq x < \lfloor x \rfloor + 1$ show that $\lfloor x \rfloor + 1$ satisfies $m > x$. Every smaller integer is at most $\lfloor x \rfloor \leq x$, so cannot satisfy it. The proposed least integer is positive because $x \geq 0$. □

If $b \leq p$, then $s + pm \geq bm$, so this model gives no strict gain. For example, $s = 1000$, $b = 20$ and $p = 4$ give the exact threshold $m \geq 63$. These numbers may describe operation counts or a declared fitted time model. A proof about a fitted model does not establish that every future clock measurement lies on its fitted line.

7.10 Certified graphical envelopes

Let $f : [a, b] \to \mathbb{R}$, and suppose a finite collection of intervals I_j covers $[a, b]$. If an enclosure procedure proves $f(I_j) \subseteq Y_j$ for every j, then the union of rectangles $I_j \times Y_j$ contains the complete graph of f. Indeed, for every $x \in [a, b]$, choose an interval containing x; its enclosure then contains $(x, f(x))$. This statement concerns mathematical coordinates, before conversion to pixels.

A smoothness condition can bound the gap between a curve and its rendered linear segments.

Proposition 7.13 (Linear interpolation envelope)**.** *Let $f \in C^2([a, b])$, $a < b$, and let ℓ be the affine function agreeing with f at a and b. If $|f''| \leq M$, then*

$$|f(x) - \ell(x)| \leq \frac{M}{2}(x - a)(b - x) \leq \frac{M(b - a)^2}{8} \quad (a \leq x \leq b).$$

Proof. The endpoint cases are zero. For $a < x < b$, choose k so that $g(t) = f(t) - \ell(t) - k(t - a)(t - b)$ vanishes at $t = x$. It also vanishes at a and b. Applying Rolle's theorem on the two intervening intervals, and then to the derivative between its two zeros, gives $g''(\xi) = 0$ at some $\xi \in (a, b)$. Thus $k = f''(\xi)/2$. Substitution at x yields $|f(x) - \ell(x)| \leq M(x - a)(b - x)/2$. Finally, $(x - a)(b - x) = (b - a)^2/4 - (x - (a + b)/2)^2 \leq (b - a)^2/4$. □

Rolle's theorem is the only calculus prerequisite in this proof. A bound on f'' must hold across the entire interval; samples of the second derivative do not establish it.

For finite text rectangles, page containment and rule clearance are elementary inequalities between coordinates. Disjointness of two axis-aligned rectangles can be established by one separating horizontal or vertical inequality. Such geometric checks are useful release conditions. They establish neither correctness of a displayed formula nor instructional effectiveness.

7.11 Evidence paths for independent review

A result intended for independent examination needs an explicit evidence path:

$$\begin{aligned} \text{stated problem} &\longrightarrow \text{mathematical model} \longrightarrow \text{proved assertion}, \\ \text{proved assertion} &\longrightarrow \text{algorithm} \longrightarrow \text{reported output}. \end{aligned}$$

A displayed graph adds a final transformation from output to coordinates and then to a rendered asset. Each arrow can introduce a distinct error. Consequently, a checksum, a successful build and a mathematical proof have different objects.

A reproducible byte comparison establishes equality of two files. A collision-resistant hash is a practical identifier, but its matching values alone are not a mathematical proof of equality of arbitrary byte strings. A kernel checks a formal assertion under its declared foundations; it does

not infer the intended source interpretation. A numerical residual measures a particular equation discrepancy; turning it into solution error requires a conditioning or stability theorem.

This collection supplies written derivations, previously recorded numerical results and an explicit formalization programme. It provides no new independent peer review, no new target-hardware execution and no new proof-assistant kernel acceptance. The certificate semantics above identify concrete developments that can strengthen subsequent releases without changing the meaning of these evidence classes.

References

[1] J. Orloff. *Complex Variables with Applications.*

[2] H. Miller and A. Mattuck. *Differential Equations.*

[3] D. Freeman. *Signals and Systems, Fall 2011.*

[4] L. N. Trefethen and J. A. C. Weideman. *The Exponentially Convergent Trapezoidal Rule.*

[5] T. A. Driscoll and R. J. Braun. *Fundamentals of Numerical Computation.*

[6] M. P. Deisenroth, A. A. Faisal and C. S. Ong. *Mathematics for Machine Learning.*

[7] S. Boyd and L. Vandenberghe. *Convex Optimization.*

[8] L. N. Trefethen. *Approximation Theory and Approximation Practice.*

[9] J. Lebl. *Guide to Cultivating Complex Analysis.*

[10] G. Strang. *Linear Algebra.*

[11] P. Flajolet and R. Sedgewick. *Analytic Combinatorics.*

[12] OpenCourseWare. *Probabilistic Systems Analysis and Applied ProbabilitySC.*

[13] NIST DLMF. *DLMF 2.4: Contour Integrals.*

[14] N. H. Bshouty and co-authors. *Elementary Proofs of Some Stirling Bounds.*

[15] A. V. Oppenheim and G. Verghese. *Introduction to Communication, Control, and Signal Processing.*

[16] J. Lebl. *Notes on Diffy Qs: Differential Equations for Engineers.*

[17] NIST DLMF. *DLMF 19.8: Quadratic Transformations.*

[18] A. Gopal and L. N. Trefethen. *Solving Laplace problems with corner singularities via rational functions.*

[19] OpenCourseWare. *Automata, Computability, and ComplexityJ.*

[20] J. Erickson. *Algorithms.*

[21] S. W. Smith. *The Scientist and Engineer's Guide to Digital Signal Processing.*

[22] E. Demaine, S. Devadas and N. Lynch. *Design and Analysis of AlgorithmsJ.*

[23] R. Szeliski. *Computer Vision: Algorithms and Applications, second edition.*

[24] Y. Nakatsukasa, O. Sete and L. N. Trefethen. *The AAA algorithm for rational approximation.*

[25] OpenCourseWare. *Quantum Physics I.*

[26] I. Goodfellow, Y. Bengio and A. Courville. *Deep Learning.*

[27] P. L. Dordal. *An Introduction to Computer Networks.*

[28] R. H. Arpaci-Dusseau and A. C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces.*

[29] C. D. Manning, P. Raghavan and H. Schuetze. *Introduction to Information Retrieval.*

[30] A. Silberschatz, H. F. Korth and S. Sudarshan. *Database System Concepts: seventh-edition teaching site.*

[31] PDOS. *Distributed Systems, Spring 2026.*

[32] OpenCourseWare. *Dynamics and Control II.*

[33] OpenCourseWare. *Electromagnetics and Applications.*

[34] OpenCourseWare. *Electromagnetic Energy: From Motors to Lasers.*

[35] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication.*

[36] NIST and contributing authors. *Digital Library of Mathematical Functions.*

[37] OpenCourseWare. *Computation Structures.*

[38] NIST DLMF. *DLMF 10.20: Uniform Asymptotic Expansions for Large Order.*

[39] R. P. Feynman, R. B. Leighton and M. Sands. *The Feynman Lectures on Physics.*