PHI TECHNICAL MONOGRAPHS

Volume VI

Extremal Length and Variational Geometry

Conformal invariants, capacity and certified optimization



PHI research collective

Contents

1 Scope, conventions and status **1**

1.1 The mathematical object 1
1.2 Mathematics, physics and computation 2
1.3 Evidence classes and source corrections 2
1.4 Notation and elementary boundary 3
1.5 Comparison contract 3

2 Supporting inequalities and variational foundations **4**

2.1 Quadratic inequalities and equality cases 4
2.2 Completion and orthogonal projection 5
2.3 Coercive quadratic energies 5
2.4 A concrete Poincaré inequality 6
2.5 Representation of a bounded linear functional 6
2.6 Finite-dimensional spectral decomposition 6
2.7 Weak harmonicity and regularity boundary 7
2.8 Coarea in the regular smooth case 7
2.9 Two different kinds of density 8

3 Source §6.1: extremal length and admissibility **9**

3.1 Functions, curves and extended values 9
3.2 Monotonicity and overflowing 10
3.3 A literal obstruction to Exercise 6.1 10
3.4 Conformal covariance of the density 10
3.5 Convexity and uniqueness of normalized minimizers 11
3.6 Finite constraint systems as an exact model 11

4 Source §6.2: series, parallel and variational duality **13**

4.1 Series inequalities without assuming minimizers 13
4.2 A weighted construction and its equality conditions 13
4.3 Parallel inequalities and disjoint supports 14
4.4 Electrical network interpretation 14
4.5 A complete two-route calculation 14
4.6 Countable families 15

5 Source §6.3: annuli, rectangles and tori **16**
5.1 §6.3.1: the cylinder calculation . . . 16
5.2 Equality and quantitative rigidity . . . 16
5.3 Logarithmic coordinates and exact annular energy . . . 17
5.4 §6.3.2: the Grötzsch inequality . . . 17
5.5 §6.3.3: Euclidean separation and normalization . . . 18
5.6 Problem 6.15: an area inequality with a stronger bound . . . 18
5.7 §6.3.4 and §6.3.6: divergent nests . . . 19
5.8 §6.3.5: real-symmetric interval pairs . . . 19
5.9 §6.3.7: rectangles and laminations . . . 20
5.10 §6.3.8: tori and homology classes . . . 20

6 Source §6.4: energy, capacity and dual witnesses **22**
6.1 §6.4.1: normalization audit . . . 22
6.2 §6.4.2: proper harmonic functions on a finite cylinder . . . 22
6.3 The regular condenser and its terminal conditions . . . 23
6.4 §6.4.3: the separating family . . . 23
6.5 A continuum complementary-energy certificate . . . 24
6.6 Energy error and flux error as orthogonal quantities . . . 24
6.7 Physical capacity, resistance and stored energy . . . 24

7 Source §§6.5–6.6: crossing and holomorphic transformation **25**
7.1 Crossing as a quantified condition . . . 25
7.2 Equality and what the product alone does not prove . . . 26
7.3 A quantitative overlap estimate . . . 26
7.4 §6.6.1: images and multiplicity . . . 26
7.5 Sharpness on power maps . . . 28
7.6 Lifting and overflowing . . . 28
7.7 §6.6.2: holes under a finite proper map . . . 28
7.8 An independent covering proof using energy . . . 29

8 Source §§6.7–6.8: canonical annuli and arcs **30**
8.1 Maximal modulus and inclusion maximality . . . 30
8.2 A torus equality family . . . 31
8.3 Quadratic differentials and the missing global step . . . 31
8.4 Buffers and the disjointness mechanism . . . 31
8.5 Covering annuli and strict inequalities . . . 32
8.6 Canonical weighted arcs . . . 32
8.7 The Euler-characteristic count . . . 33
8.8 What is completed by the source audit . . . 33

9 Finite networks: complete primal and dual constructions **35**
9.1 Incidence, orientation and the energy . . . 35
9.2 The discrete maximum principle . . . 35

9.3 Unit flows and summation by parts . . . 36
9.4 A completed-square proof of the same certificate . . . 36
9.5 The path-density formulation on a graph . . . 37
9.6 Rayleigh monotonicity and conductance sensitivity . . . 37
9.7 Schur complements and boundary response . . . 37
9.8 Exact rational witnesses and bit complexity . . . 38
9.9 Failure modes addressed by the checker . . . 38

10 Quantitative variational stability and attained constants 40
10.1 A common affine constraint . . . 40
10.2 An attained perturbation bound . . . 40
10.3 A weaker resolvent estimate and the source of improvement . . . 41
10.4 Energy values under form perturbation . . . 41
10.5 Boundary perturbation as a separate mechanism . . . 42
10.6 A-posteriori error from an approximate Euler condition . . . 42
10.7 Primal–dual gap and two error components . . . 43
10.8 Where the generalization remains open . . . 43

11 Weighted and nonlinear modulus 44
11.1 Finite p-modulus . . . 44
11.2 Series and parallel nonlinear resistances . . . 44
11.3 A dual finite certificate . . . 44
11.4 Radial p-capacity in dimension n . . . 45
11.5 Why conformal invariance selects an exponent . . . 45
11.6 Anisotropic rectangle: a complete model . . . 46
11.7 Metric-measure and singular categories . . . 46

12 Periods, degeneration and Fourier boundary operators 47
12.1 Marked torus deformations . . . 47
12.2 Periods as finite-dimensional data . . . 47
12.3 A circular plumbing neck . . . 48
12.4 Harmonic extension on a cylinder . . . 48
12.5 Dirichlet-to-Neumann matrices . . . 49
12.6 Exponential information loss and an inverse lower bound . . . 49
12.7 Optimal finite-rank approximation of a diagonal transfer . . . 49
12.8 Elliptic-integral computation and certification boundary . . . 50

13 Physical models with explicit conformal and energy reductions 51
13.1 Coaxial electrostatics and thermal conduction . . . 51
13.2 Diffusion time is not conformally invariant . . . 51
13.3 Anti-plane elasticity . . . 52
13.4 Threshold sheets in a scattering coordinate . . . 52
13.5 A complete single-channel unitarity model . . . 53
13.6 A finite-range bound-state model . . . 53

13.7 Classical worldsheet energy . . . 53
13.8 Gauge holonomy on an annulus . . . 53
13.9 Spin boundary conditions and mode shifts . . . 54
13.10 Period geometry in a conditional gauge model . . . 54

14 Spectral thermodynamics, horizons and identifiability 55
14.1 Finite Gaussian systems . . . 55
14.2 Finite-mode determinants and the infinite limit . . . 55
14.3 A bosonic oscillator and its heat capacity . . . 56
14.4 Partition zeros and real-axis response . . . 56
14.5 A local Euclidean horizon calculation . . . 56
14.6 Area, entropy and ensemble stability . . . 57
14.7 A statistical obstruction for multiverse-model claims . . . 58
14.8 Physical data and remaining execution tasks . . . 58

15 Algorithms, information and distributed computation 59
15.1 A chain admits a sufficient statistic . . . 59
15.2 One-edge updates and their exact restriction . . . 59
15.3 Iterative solution with an explicit error contract . . . 60
15.4 Parallel work, depth and communication . . . 60
15.5 Consensus as a spectral statement . . . 61
15.6 Cryptographic and quantum boundaries . . . 61
15.7 Logical complexity and proof reuse . . . 62

16 Exact arithmetic, enclosures and certificate completeness 63
16.1 Finite witnesses and the logical quantifiers . . . 63
16.2 Rational existence for connected networks . . . 63
16.3 An algebraic positive-definiteness certificate . . . 64
16.4 Directed interval semantics . . . 64
16.5 A logarithm enclosure with an explicit tail . . . 64
16.6 Exact square-root enclosures . . . 65
16.7 Input intervals for conductance . . . 65
16.8 Termination and malformed certificates . . . 66

17 Source-selected problems: complete finite-dimensional treatments 67
17.1 Selection boundary and mathematical role . . . 67
17.2 QK01. An absolute-difference kernel is not positive semidefinite . . . 67
17.3 QK02. An explicit feature map and its converse . . . 68
17.4 QS01. Gauss–Markov optimality as an orthogonal decomposition . . . 69
17.5 QK03. Nyström approximation with singular landmarks . . . 70

18 Solved investigations in extremal geometry and computation 73
18.1 Problem-family conventions . . . 73
18.2 VI01. Optimal allocation in a layered conductor . . . 73

18.3 VI02. A balanced bridge and an unidentifiable edge . . . 74
18.4 VI03. Exact one-edge sensitivity beyond a derivative . . . 74
18.5 VI04. Path and cut witnesses bound every network . . . 74
18.6 VI05. Star elimination and the boundary response . . . 75
18.7 VI06. Uncertainty propagation on a three-layer chain . . . 75
18.8 VI07. An attained stability bound for a cylinder density . . . 75
18.9 VI08. A homological intersection inequality on a flat torus . . . 76
18.10VI09. A minimax shape for two marked torus classes . . . 76
18.11VI10. A rotated tensor invalidates a guessed linear solution . . . 77
18.12VI11. The limiting behavior of a finite p-modulus chain . . . 77
18.13VI12. A certificate for canonical-weight positivity . . . 77
18.14VI13. An exactly attained form-perturbation example . . . 78
18.15VI14. Choosing precision from the sharp error factor . . . 78
18.16VI15. Complementary modes of the cylinder boundary operator . . . 79
18.17VI16. An information lower bound for analytic continuation . . . 79
18.18VI17. The best uniform rank for a diagonal restriction . . . 79
18.19VI18. Unitary scattering and perturbation of a reaction matrix . . . 80
18.20VI19. Conformal invariance of steady heat but not its clock . . . 80
18.21VI20. A Gaussian determinant with a missing zero mode . . . 81
18.22VI21. A phase transition between certificate regimes . . . 81
18.23VI22. Resource-optimal repeated-query selection . . . 81
18.24VI23. A finite nonidentifiability result for model selection . . . 82
18.25VI24. A robust current certificate under a changed orientation . . . 82

19 Quantitative network design and perturbation geometry 83
19.1 The mathematical design problem . . . 83
19.2 A complete budget-optimality theorem . . . 83
19.3 A second proof through a current allocation . . . 84
19.4 Worked rational design . . . 84
19.5 First and second variations of fixed-network conductance . . . 85
19.6 Logarithmic sensitivity is an energy distribution . . . 85
19.7 A certificate for a proposed resource optimum . . . 86
19.8 Relation to continuum design . . . 86

20 Boundary operators and finite spectral reductions 87
20.1 Elimination as a variational operation . . . 87
20.2 Associativity of exact elimination . . . 87
20.3 Energy-preserving assembly . . . 88
20.4 Harmonic extensions are contractions in the energy quotient . . . 88
20.5 A two-terminal response determines one scalar . . . 88
20.6 Generalized eigenvalues and a controlled mass matrix . . . 89
20.7 Residual inclusion for a symmetric eigenproblem . . . 89
20.8 A finite heat-semigroup comparison . . . 89

20.9 Infinite-dimensional extension requirements 90

21 Constructive continuum comparisons and certified truncation 91
21.1 An exactly solvable variable medium 91
21.2 Piecewise-linear restriction and its precise defect 91
21.3 A manufactured discontinuous example 92
21.4 A two-dimensional layered reduction 92
21.5 An interface resistance as a singular-layer limit 93
21.6 Energy distortion under a controlled coordinate map 93
21.7 Complex dilatation and a computable bound 93
21.8 Fourier truncation with an explicit energy tail 94
21.9 A certified geometric-computational error split 94

22 Deficit bounds for finite extremal metrics 95
22.1 A convex admissible population 95
22.2 A scalar remainder inequality 95
22.3 A quantitative certificate for the optimizer 96
22.4 The quadratic case and active constraints 96
22.5 A sharp quadratic example 96
22.6 A degenerating pair of linear constraints 97
22.7 Turning a computed deficit into an explicit stopping rule 97

23 Exceptional curve families and attainment of the variational problem 99
23.1 The distinction resolved by the extension 99
23.2 A zero-width criterion 99
23.3 Strong and weak constraint infima agree 100
23.4 Convergence on almost every curve 100
23.5 Closedness, existence and uniqueness 100
23.6 Quantitative attainment and an exact equality family 101
23.7 A valid smooth-density conclusion for regular condensers 101

24 Explicit geometric estimates for ring domains 103
24.1 Complementary continua and the model convention 103
24.2 A complete tube proof of the linear upper bound 103
24.3 A logarithmic upper estimate 104
24.4 An inscribed annulus with explicit thresholds 104
24.5 A quantitative two-sided pinching statement 105
24.6 Real-symmetric intervals 106
24.7 Divergence of nested annuli 106

25 Area distortion and a sharp annular deficit identity 108
25.1 A proof that avoids smooth boundary assumptions 108
25.2 The signed boundary-area identity 108
25.3 The source area bound and its equality condition 109

25.4 The exact spectral deficit . . . 109
25.5 Representation and certified evaluation . . . 110
25.6 Comparative problem: near-circular annular images . . . 110

26 A matrix proof of the collar estimate 112
26.1 The geometric setting and the remaining distinction . . . 112
26.2 Elementary isometries and their fixed axes . . . 112
26.3 Distance between the imaginary axis and a semicircle . . . 112
26.4 Trace separation of equal-length axes . . . 113
26.5 Fermi coordinates and an embedded collar . . . 113
26.6 A stronger covering-modulus estimate . . . 114
26.7 Two counterexamples to the unrestricted crossing equality claim . . . 114
26.8 Computational verification and mathematical scope . . . 115

27 Attainment of maximal modulus on a Riemann surface 117
27.1 Separating attainment from uniqueness . . . 117
27.2 A noncollapse criterion from four omitted points . . . 117
27.3 Existence in a nonperipheral planar class . . . 118
27.4 An elementary cusp-neighborhood construction . . . 118
27.5 Attainment in a hyperbolic surface . . . 119
27.6 Spherical and flat cases . . . 120
27.7 Elementary exceptions in the canonical-arc formulation . . . 120
27.8 Constructive information in the existence proof . . . 121

28 Regular condenser duality from variational and flow coordinates 122
28.1 Boundary charts and the precise regular class . . . 122
28.2 Coercivity on a finite system of charts . . . 122
28.3 Existence and weak stationarity . . . 123
28.4 Interior regularity and analytic reflection . . . 123
28.5 Flux conservation and flow coordinates . . . 124
28.6 The two curve families and exact duality . . . 125
28.7 Uniqueness, error identities and the physical normalization . . . 126
28.8 What this proof does not transfer automatically . . . 126

29 Marked quadrilaterals and exact arithmetic for elliptic moduli 127
29.1 The standard half-plane problem . . . 127
29.2 An explicit rectangle coordinate . . . 127
29.3 Uniqueness and stronger elementary estimates . . . 128
29.4 A complete arithmetic–geometric-mean reduction . . . 128
29.5 A rational certificate for the transcendental modulus . . . 129
29.6 An independent validated quadrature representation . . . 130

30 A two-dimensional continuum certificate with independent implementations 132
30.1 The fixed mathematical problem . . . 132

30.2 Sobolev spaces and the terminal traces used here . . . 132
30.3 The conjugate-potential pairing . . . 133
30.4 Exact reciprocity without interface smoothness . . . 134
30.5 A continuum reference value for a checkerboard . . . 135
30.6 Conforming triangular witnesses . . . 135
30.7 What the finite checker establishes . . . 136
30.8 Completeness and termination of the specified fragment . . . 136
30.9 A-posteriori function error and coefficient uncertainty . . . 137

31 Work, precision and stopping rules for the continuum witnesses 138
31.1 The matrix problem and its relation to the continuum . . . 138
31.2 Conjugate gradients in exact arithmetic . . . 138
31.3 A complete polynomial error bound . . . 139
31.4 An explicit mesh-dependent condition estimate . . . 139
31.5 Preconditioning and the measured ablation . . . 140
31.6 Precision limits and a useful stopping decision . . . 140
31.7 Parallel and distributed interpretation . . . 141

32 New continuum experiments and exact elliptic certificates 142
32.1 What was executed and what a passing check means . . . 142
32.2 Mesh convergence and high-contrast limitations . . . 142
32.3 Separating rounding from spatial error . . . 143
32.4 A matched solver comparison with an optimized control . . . 145
32.5 Exact checking costs and implementation independence . . . 146
32.6 Quadrilateral certificates at large and small moduli . . . 146
32.7 Finite audit counts and their limits . . . 147
32.8 Reproduction sequence and preserved historical evidence . . . 147

33 Executed calculations and reproducibility 148
33.1 The evidence obtained in this reconstruction . . . 148
33.2 Exact potential–flow certificates . . . 148
33.3 Exact resource-design witnesses . . . 149
33.4 Chain timing experiment . . . 149
33.5 Independent exact forward-error evaluation . . . 150
33.6 Symbolic identities and selected source problems . . . 151
33.7 Reproduction order . . . 151
33.8 Unperformed experimental programmes . . . 151

34 Worked synthesis: geometry, uncertainty and certificates 153
34.1 VI25. An anisotropic rectangle with insulating sides . . . 153
34.2 VI26. Rotating a two-phase laminate . . . 153
34.3 VI27. A sharp interval with uncertain geometry . . . 154
34.4 VI28. An off-diagonal tensor changes the boundary problem . . . 154
34.5 VI29. Removing diagonal anisotropy by changing the domain . . . 154

34.6 VI30. The exact residual behind a primal–dual gap . . . 155
34.7 VI31. A field-error bound from an interval . . . 155
34.8 VI32. Which central estimate minimizes relative error? . . . 155
34.9 VI33. Reciprocation preserves a relative certificate radius . . . 155
34.10 VI34. Combining independent certificates . . . 156
34.11 VI35. A coefficient uncertainty bound valid for every admissible field . . . 156
34.12 VI36. What nested approximation spaces actually guarantee . . . 156
34.13 VI37. Exact admissibility of a rational tensor . . . 156
34.14 VI38. A complete rational network certificate by hand . . . 157
34.15 VI39. Duality returns the original tensor . . . 157
34.16 VI40. Separating material uncertainty from solve error . . . 157
34.17 Executable checks and the boundary of the examples . . . 158

35 Source obligations and formalization groundwork 159
35.1 Source locations and mathematical completion are separate . . . 159
35.2 Definitions, circuit laws and ring geometry . . . 159
35.3 Dirichlet duality, crossing and transformation . . . 160
35.4 Maximal annuli and the unresolved canonical geometry . . . 160
35.5 What the executable evidence establishes . . . 160
35.6 Prospective formalization and remaining source screening . . . 161
35.7 Publication status . . . 161

Chapter 1

Scope, conventions and status

1.1 The mathematical object

Extremal length assigns a conformal invariant to a family of curves by optimizing a ratio between a squared minimum length and area. The optimization is over metrics, not over coordinates alone. Its reciprocal, extremal width, has a convex formulation as a minimum quadratic energy under infinitely many path constraints. These two descriptions connect the geometry of Riemann surfaces to Dirichlet problems, finite electrical networks, and certified numerical optimization. They also make the assumptions of a comparison unusually visible: the admissible family, the meaning of a path, its multiplicity, the measure, and the normalization all influence the answer.

This manuscript follows Section 6 of Mikhail Lyubich's *Conformal Geometry and Dynamics of Quadratic Polynomials, volumes I–II*, using the supplied 702-page version. The source section begins on physical page 119 and ends on page 133 before Section 7. The eight main subsections and their nested divisions are preserved. Source items 6.1–6.45 include exercises, remarks and a problem as well as theorems and propositions. Four unnumbered named principles are separately tracked. Neither the count of source identifiers nor the length of this manuscript is a count of established new theorems.

This scholarly edition develops explicitly delimited finite-dimensional and continuum statements, source corrections and supporting proofs. Its mathematical core includes weak modulus, rough ring estimates, annular area deficits, collar geometry, maximal-annulus existence, regular condenser duality, exact elliptic-modulus certificates and independently checked two-dimensional continuum witnesses. Forty worked investigations connect these arguments with geometry, physical models, uncertainty and computation. The companion internal edition adds a teaching sequence and assessment material.

The completion scope follows the public and internal standard of PHI Volume V: written derivations with declared dependencies, reproducible finite checks and explicit boundaries. The earlier broad commission has been superseded for this release. This change of scope does not supply missing mathematical proofs. In particular, the global nonperipheral Jenkins–Strebel uniqueness and canonical-buffer/arc development remain conditional where identified below; arbitrary rough-boundary extensions are not inferred from regular cases. The source audit records those boundaries rather than presenting every source assertion as established. No proof-assistant implementation or independent peer review is claimed.

1.2 Mathematics, physics and computation

The principal mathematical relation is

$$\mathsf{L}(\Gamma_v) = \frac{1}{\mathrm{Cap}}, \qquad \mathsf{W}(\Gamma_v) = \mathrm{Cap}, \qquad \mathrm{Cap} = \inf_{u|_{B_0}=0,\ u|_{B_1}=1} \int |\nabla u|^2 \, \mathrm{d}A, \tag{1.1}$$

for a regular condenser with the specified conventions. The geometric theorem connecting the path problem to the potential problem needs a proof; the equality cannot be introduced by naming both quantities “capacity.” The opposite, separating curve family has reciprocal extremal behavior. In an electrical interpretation, conductance is proportional to the potential energy under unit voltage, whereas resistance is its reciprocal. Physical stored energy has an additional one-half factor and material parameters. Each factor is retained below.

For computer science the useful structure is a specification with a witness. A numerical solver can propose a potential and a flux. A separate exact checker tests boundary conditions, conservation and energy inequalities. Once those finite constraints hold, a written duality theorem supplies an upper and a lower bound. The candidate-generating solver need not itself be trusted to have converged. The finite representation and the checker arithmetic remain part of the trust boundary.

The physical chapters develop selected electrostatic, thermal, wave, scattering and field models. Relations to worldsheet actions, spectral thermodynamics and Euclidean horizon geometry are stated within their models. This material contains no measured nuclear or particle-data analysis and no experimental hardware implementation. High-dimensional compactification, full gauge-theory dynamics, quantum gravity and multiverse proposals are not established by a two-dimensional conformal invariant. These broader programmes remain in the historical obligation register and are outside the present completion scope.

1.3 Evidence classes and source corrections

A written proof establishes the mathematical statement under its premises. An exact rational computation establishes a finite identity within its arithmetic model. A floating-point comparison is a diagnostic unless an enclosure theorem and directed error analysis turn it into a certificate. A compiled native executable is not a verified compiler or a proof-assistant kernel. These categories occupy separate fields in the delivered registers.

Source-derived definitions and assertions are identified by source number. A proposed correction receives its own statement and explanation. Original authoring of a proof does not establish priority. The quantitative perturbation and certificate results proved here are recorded as derivations in this manuscript; a complete novelty assessment against the literature remains open. General modulus duality and network duality are established research areas [4, 5, 6]. The external references are mathematical context, not substitutes for missing proofs in the current development.

Three source issues require particular attention. First, an arbitrary family of curves can detect density values on area-null sets; Exercise 6.1, read literally for all such families, admits the counterexample in Chapter 3. Second, the energy formulas on source page 126 mix incompatible factors under the stated Wirtinger convention. Third, puncture classes admit many annuli of infinite modulus, so maximal-modulus uniqueness requires qualifications. These assertions are tested against the rendered source, and corrected statements are kept separate from the original wording.

1.4 Notation and elementary boundary

A Riemann surface is connected, Hausdorff and second countable, with holomorphic transition maps. A domain is an open connected set. Unless a chapter says otherwise, curves are continuous rectifiable parametrized curves and line integrals count traversal multiplicity. A multicurve is a finite union with the stated multiplicity convention. “Almost everywhere” refers to the named area or transverse measure, never to an unspecified exceptional population.

Write $A(r, R) = \{z : r < |z| < R\}$ and

$$\operatorname{mod} A(r, R) = \frac{1}{2\pi} \log \frac{R}{r}.$$

The cylinder $\mathbb{R}/\ell\mathbb{Z} \times (0, h)$ has modulus h/ℓ in the vertical direction. We use $\partial_z = (\partial_x - i\partial_y)/2$ and $\partial_{\bar{z}} = (\partial_x + i\partial_y)/2$. For a real potential,

$$\mathcal{E}(u) = \int (u_x^2 + u_y^2)\, \mathrm{d}A.$$

For a complex map $f = u + iv$, the same full real energy is $\mathcal{E}(f) = \mathcal{E}(u) + \mathcal{E}(v)$. The half-energy $\mathcal{E}(f)/2$ equals image area for an injective holomorphic map. We will not change between these conventions without notation.

The directly proved elementary boundary includes finite-dimensional linear algebra, differentiable calculus, ordinary integration, and the elementary complex-analytic prerequisites identified in previous volumes. Supporting Hilbert-space arguments are supplied here. General Sobolev trace theory, geometric measure theory and quadratic-differential existence are identified as imported dependencies where still needed. These named dependencies are part of the scholarly scope, as in Volume V; a transitive derivation of every prerequisite from basic undergraduate knowledge is not claimed.

1.5 Comparison contract

A comparison fixes the input representation, output, norm, accuracy and termination requirement. For an exact graph calculation the input is a finite list of vertices, edges and positive rational conductances. For a surface calculation the geometry also needs a finite description and a theorem connecting it to the continuum model. The latter obligation cannot be settled by drawing the boundary.

Setup, solve, memory, arithmetic precision, communication and certification costs are separated. An $O(m)$ arithmetic checker need not have $O(m)$ bit complexity: intermediate integer lengths and multiplication costs must be included. The experiments use their recorded environment and finite instances. They do not assert GPU, ARM, multi-node or user-workstation measurements. A disclosed experiment suite is not a hidden holdout or an independent physical validation.

Every unfinished mathematical item retains its status in the source register. Publication of this scoped edition does not mark the earlier commission complete. The sixteen new investigations and the internal teaching layer are editorial and expository additions; the accepted historical witnesses retain their original bytes.

Chapter 2

Supporting inequalities and variational foundations

2.1 Quadratic inequalities and equality cases

Lemma 2.1 (Weighted Cauchy–Schwarz)**.** [F01] *Let $a_i > 0$ and $x_i, y_i \in \mathbb{R}$. Then*

$$\left(\sum_i x_i y_i\right)^2 \le \left(\sum_i a_i x_i^2\right)\left(\sum_i \frac{y_i^2}{a_i}\right).$$

If both vectors are nonzero, equality holds precisely when $\sqrt{a_i}x_i$ is one common scalar multiple of $y_i/\sqrt{a_i}$.

Proof. Set $X_i = \sqrt{a_i}x_i$ and $Y_i = y_i/\sqrt{a_i}$. For every real t the sum $\sum_i (X_i - tY_i)^2$ is nonnegative. If $Y \neq 0$, its minimum at $t = (X \cdot Y)/\|Y\|^2$ equals $\|X\|^2 - (X \cdot Y)^2/\|Y\|^2$. Rearrangement proves the inequality. The minimum vanishes precisely when $X = tY$. Zero vectors give the remaining cases directly. □

The same proof, using $\int (X - tY)^2 \ge 0$, works in L^2 whenever the integrals exist. The equality statement is then an almost-everywhere statement. A function with a changed value on a null set represents the same L^2 vector; it need not give the same line integrals on all curves. The distinction between a function and its equivalence class is essential for extremal length.

Lemma 2.2 (Young's inequality)**.** [F02] *For $p > 1$, $q = p/(p-1)$ and $a, b \ge 0$,*

$$ab \le \frac{a^p}{p} + \frac{b^q}{q},$$

with equality exactly when $b = a^{p-1}$.

Proof. For fixed b minimize $a^p/p - ab$ over $a \ge 0$. The derivative is $a^{p-1} - b$, which has the unique zero $a = b^{1/(p-1)}$ when $b > 0$. The second derivative is nonnegative, and the minimum equals $-b^q/q$. The case $b = 0$ is direct. The derivative and strict convexity also give the equality condition. □

After normalizing f and g by their L^p and L^q norms and integrating Young's inequality, one obtains Hölder's inequality. This supplies the nonlinear energy bounds in Chapter 11; no analogy with quadratic energy is needed.

2.2 Completion and orthogonal projection

A real Hilbert space is a complete inner-product space. The completion of a normed inner-product space can be constructed as Cauchy sequences modulo sequences whose difference tends to zero. Addition and scalar multiplication act termwise. The inner product is the limit of the inner products of representatives; Cauchy–Schwarz shows that changing representatives does not change the limit. Completeness follows by selecting a diagonal representative from a Cauchy sequence of equivalence classes.

Theorem 2.3 (Nearest point in a closed affine subspace)**.** [F03] *Let H be a Hilbert space, V a closed linear subspace, and $K = g + V$. There is a unique $u \in K$ minimizing $\|u\|^2$. It is characterized by*

$$\langle u, v\rangle = 0 \quad (v \in V), \qquad \|u + v\|^2 = \|u\|^2 + \|v\|^2.$$

Proof. Let $d = \inf_{x\in K}\|x\|$ and choose $x_n \in K$ with $\|x_n\| \to d$. Since $(x_n + x_m)/2 \in K$, the parallelogram identity gives

$$\|x_n - x_m\|^2 = 2\|x_n\|^2 + 2\|x_m\|^2 - 4\|(x_n + x_m)/2\|^2 \le 2\|x_n\|^2 + 2\|x_m\|^2 - 4d^2 \to 0.$$

Completeness and closedness give a limit $u \in K$, with norm d. For $v \in V$, the quadratic $\|u + tv\|^2$ has a minimum at $t = 0$, so its linear coefficient $2\langle u, v\rangle$ vanishes. Orthogonality gives the displayed Pythagorean identity. If a second minimizer were $u + v$, the identity would force $v = 0$. □

This proof is existence-based: choosing a minimizing sequence is not an algorithm with a stopping criterion. In a finite-dimensional rational model, Gaussian elimination provides an effective construction. In an infinite-dimensional problem, an approximation space, density theorem and error estimator are additional requirements.

2.3 Coercive quadratic energies

Let a be a symmetric bilinear form with

$$\alpha\|v\|_H^2 \le a(v, v) \le \beta\|v\|_H^2, \quad 0 < \alpha \le \beta < \infty.$$

The a-norm is equivalent to the original norm, hence is complete and has the same closed subspaces. Applying the preceding theorem in this new inner product proves existence and uniqueness of a minimizer on $g + V$. Its Euler condition is $a(u, v) = 0$ for all $v \in V$.

Proposition 2.4 (Energy error identity)**.** [F04] *If u is this minimizer and $w \in g + V$, then*

$$a(w, w) - a(u, u) = a(w - u, w - u).$$

Proof. Set $e = w - u \in V$. Expand $a(u + e, u + e)$. The cross term is $2a(u, e) = 0$ by the Euler condition. The remaining terms give the result. □

An approximate energy supplies an upper bound on the minimum; it does not by itself supply a lower bound or a norm error, because $a(u, u)$ remains unknown. A dual witness supplies the missing side. The finite graph construction in Chapter 9 makes this statement entirely algebraic.

2.4 A concrete Poincaré inequality

For a continuously differentiable function v on $[0,\ell]\times[0,h]$ with $v(x,0)=0$,

$$v(x,y)=\int_0^y v_y(x,s)\,\mathrm{d}s, \qquad |v(x,y)|^2 \le y\int_0^y |v_y(x,s)|^2\,\mathrm{d}s.$$

Integrating and reversing the nonnegative integrals yields

$$\int |v|^2 \le \frac{h^2}{2}\int |v_y|^2 \le \frac{h^2}{2}\int |\nabla v|^2. \tag{2.1}$$

Thus gradient energy controls the full L^2 norm on this boundary class. The constant is sufficient for coercivity; it is not asserted to be optimal. Completing smooth functions in the gradient norm produces a Hilbert space of zero-trace candidates. The boundary trace and equality of this completion with other definitions of H_0^1 require the relevant trace/density theorem on a general domain.

On a rectangle, a smooth prescribed trace can be extended explicitly by interpolation between opposite sides. The affine variational class is then nonempty. For rough domains, "set the boundary values" must be replaced by a specified trace or quasi-continuous boundary condition. No generic regularity theorem is inferred from a diagram of the domain.

2.5 Representation of a bounded linear functional

Proposition 2.5. [F05] *Every bounded linear functional ℓ on a real Hilbert space has a unique representing vector y such that $\ell(x)=\langle x,y\rangle$ for all x, and $\|\ell\|=\|y\|$.*

Proof. The zero functional has the zero representing vector. Otherwise $M=\ker\ell$ is a closed proper subspace. Choose x_0 with $\ell(x_0)\neq 0$ and let $p=x_0-P_Mx_0$, using the projection already proved above. Then $p\perp M$ and $\ell(p)=\ell(x_0)\neq 0$. For every x, the vector $x-\ell(x)p/\ell(p)$ belongs to M, so $\langle x,p\rangle=\ell(x)\|p\|^2/\ell(p)$. Set $y=\ell(p)p/\|p\|^2$ to obtain the representation. If two vectors represent the same functional, their difference has zero inner product with every vector, in particular with itself, proving uniqueness. Cauchy–Schwarz gives $\|\ell\|\le\|y\|$, while evaluating at $y/\|y\|$ gives the reverse inequality when $y\neq 0$. □

Consequently a bounded symmetric bilinear perturbation $d(x,z)$ with $|d(x,z)|\le\varepsilon\|x\|\|z\|$ has an operator representation $d(x,z)=\langle Ex,z\rangle$. Apply the proposition to the functional $z\mapsto d(x,z)$. Uniqueness proves linearity in x, the functional-norm identity proves $\|E\|\le\varepsilon$, and symmetry proves E self-adjoint. This supplies the representation used by the later affine perturbation theorem without leaving Riesz representation as an unproved dependency.

If $b=a+d$ and $\varepsilon<1$, then b is a positive inner product with equivalent norm. A b-Cauchy sequence is a-Cauchy and converges in a; the upper norm bound then gives convergence in b. Thus the new inner-product space is complete, and a closed affine subspace remains closed under either norm. The projection theorem proves existence and uniqueness of both minimizers used in the perturbation argument.

2.6 Finite-dimensional spectral decomposition

Let A be real symmetric. The continuous Rayleigh quotient x^TAx attains a maximum on the Euclidean unit sphere by compactness. Let v be a maximizing unit vector. For every unit $w\perp v$, differentiate the value at $v\cos t+w\sin t$ at zero. The derivative is $2w^TAv$, so Av is parallel to

v. Its eigenvalue is $v^T Av$. Symmetry makes $v^\perp$ invariant: $v^T Aw = (Av)^T w = 0$. Induction on dimension produces an orthonormal eigenbasis.

This proves the finite spectral decomposition used throughout the manuscript from compactness and differentiable calculus. A positive semidefinite matrix has nonnegative eigenvalues by evaluating the quadratic form on eigenvectors, and conversely nonnegative eigenvalues imply positivity by expansion. Square roots, pseudoinverses and orthogonal range projectors are defined by their eigenvalues, with zeros treated explicitly. For a complex Hermitian matrix, varying in both w and iw gives the real and imaginary orthogonality conditions, and the same induction applies.

These finite constructions do not establish the spectral theorem for arbitrary unbounded operators. Any infinite-dimensional spectral statement later in the manuscript either concerns an explicitly diagonal operator or remains a named supporting obligation.

2.7 Weak harmonicity and regularity boundary

For a minimizer of $\int |\nabla u|^2$, perturbing by $t\varphi$ for a smooth compactly supported φ gives

$$\int \nabla u \cdot \nabla \varphi = 0.$$

If u is twice differentiable, integration by parts identifies this with $\Delta u = 0$. Conversely, a twice differentiable harmonic function with the prescribed boundary values minimizes energy because its cross term against zero-boundary perturbations vanishes.

The implication from a weak solution to an interior smooth harmonic function is an elliptic regularity theorem. One classical route mollifies u on compactly contained disks: u_ε is smooth and harmonic because convolution commutes with distributions and the Laplacian. Mean-value estimates for the harmonic mollifications produce a smooth limit agreeing almost everywhere with u. Boundary regularity requires separate hypotheses. This paragraph gives the local route. Lemmas 28.1–28.3 supply the subsequent weak, reflection, regularity and flux developments for compact surfaces with analytic boundary; Chapter 30 develops the square trace and discontinuous-tensor setting independently. The full measure/distribution construction and arbitrary rough trace theory remain separately identified foundations.

2.8 Coarea in the regular smooth case

Let u be smooth and $\nabla u \neq 0$ on a strip of level sets. Locally take coordinates (t, s) with $t = u$. The normal coordinate increment satisfies $\mathrm{d}n = \mathrm{d}t/|\nabla u|$, so the Jacobian identity is

$$\int_{a<u<b} F|\nabla u| \,\mathrm{d}A = \int_a^b \left(\int_{u=t} F \,\mathrm{d}s \right) \mathrm{d}t.$$

A partition of unity extends this equality over a compact region of regular levels. For isolated critical points of a nonconstant harmonic function, remove small disks around them, apply the regular identity and pass to limits under the stated integrability conditions. General measurable coarea is a stronger theorem. The distinction is important when the density is singular or when infinitely many critical points approach a boundary.

For a harmonic condenser potential u, the divergence theorem shows that

$$I(t) = \int_{u=t} |\nabla u| \,\mathrm{d}s$$

is constant on regular levels that separate the same terminal boundaries. Applying coarea to $F = |\nabla u|$ shows that this constant is $\mathcal{E}(u)$ when the boundary values are zero and one. This identifies flux with energy under unit potential difference.

2.9 Two different kinds of density

An area-density approximation in L^2 controls integrals against area-measurable test functions. It does not uniformly control restrictions to every curve. If a family contains curves supported on a null set, two equal L^2 classes can assign different lengths. Therefore a Hilbert-space closure of admissible densities may relax pointwise path constraints unless exceptional curve families are explicitly handled.

The convex geometric machinery above is consequently applied either to finite constraints, to smooth regular condensers with an established modulus–capacity correspondence, or to a precisely stated weak-modulus theory. Passing between these settings is never counted as an automatic change of notation.

Chapter 3

Source §6.1: extremal length and admissibility

Source correspondence. Lyubich, §6.1, Exercises 6.1–6.3 and Remark 6.4..

3.1 Functions, curves and extended values

Let U carry a reference conformal area element $\mathrm{d}A$. A density is an everywhere-defined Borel function $\rho : U \to [0, \infty]$. Its area energy is $M(\rho) = \int_U \rho^2 \, \mathrm{d}A$. For a rectifiable parametrized curve γ, define $\ell_\rho(\gamma) = \int_\gamma \rho \, \mathrm{d}s$ and $\ell_\rho(\Gamma) = \inf_{\gamma \in \Gamma} \ell_\rho(\gamma)$. Only $0 < M(\rho) < \infty$ is used in the quotient

$$\mathsf{L}(\Gamma) = \sup_\rho \frac{\ell_\rho(\Gamma)^2}{M(\rho)}.$$

For the dual definition, densities satisfying every path constraint are allowed to have zero energy:

$$\mathrm{Adm}(\Gamma) = \{\rho \geq 0 : \ell_\rho(\gamma) \geq 1 \ \forall \gamma \in \Gamma\}, \qquad \mathsf{W}(\Gamma) = \inf_{\rho \in \mathrm{Adm}(\Gamma)} M(\rho).$$

The reciprocal relation is interpreted in $[0, \infty]$, with $1/0 = \infty$ and $1/\infty = 0$. An empty family has infinite infimum length and no constraints, giving $\mathsf{L} = \infty$ and $\mathsf{W} = 0$. A family containing a constant curve admits no density with that curve's positive line integral, giving $\mathsf{W} = \infty$ and $\mathsf{L} = 0$.

Proposition 3.1 (Reciprocity)**.** [D01] *Under these extended-value conventions,* $\mathsf{W}(\Gamma) = 1/\mathsf{L}(\Gamma)$.

Proof. Whenever $0 < \ell_\rho(\Gamma) < \infty$, divide ρ by $\ell_\rho(\Gamma)$. The new density is admissible and has energy $M(\rho)/\ell_\rho(\Gamma)^2$. Hence $\mathsf{W} \leq 1/\mathsf{L}$ by approximation to the supremum. Conversely an admissible finite positive-energy density has quotient at least $1/M(\rho)$, so $\mathsf{L} \geq 1/\mathsf{W}$. If a finite-energy density has infinite minimum length, rescaling it by any small positive constant preserves admissibility and drives energy to zero. If an admissible density has zero energy, adding a small positive finite-energy reference density gives admissible densities of arbitrarily small positive energy on finite-area exhaustions; equivalently the definition directly assigns zero width. If no admissible finite-energy density exists, the quotient must vanish for every finite-energy density, since a positive minimum length would yield one by normalization. The remaining endpoint cases follow from these alternatives. □

For a noncompact infinite-area domain the zero-energy perturbation can be made with a strictly positive integrable reference density, for example one constructed on a countable relatively compact exhaustion. That extra construction is necessary; a constant density need not have finite area there.

3.2 Monotonicity and overflowing

Proposition 3.2. [D02] *If $\Gamma_1 \subseteq \Gamma_2$, then $\mathsf{L}(\Gamma_2) \le \mathsf{L}(\Gamma_1)$ and $\mathsf{W}(\Gamma_2) \ge \mathsf{W}(\Gamma_1)$. If every curve of Γ contains a subcurve in Δ, then $\mathsf{L}(\Gamma) \ge \mathsf{L}(\Delta)$.*

Proof. Taking an infimum over a larger population can only decrease the minimum length. This gives the first length inequality for every density and hence for the suprema. The admissible densities for Γ_2 form a subset of those for Γ_1, which proves the width inequality independently. For overflowing, nonnegativity gives $\ell_\rho(\gamma) \ge \ell_\rho(\delta)$ for the contained subcurve. Infima and then suprema give the result. □

Restricting a density to a measurable support containing every curve, or extending it by zero to a larger ambient surface, preserves the relevant path integrals and does not increase energy. Applying both operations proves ambient independence under compatible measure and support conventions. Endpoints have zero arclength and do not alter the conclusion.

3.3 A literal obstruction to Exercise 6.1

The source asks that smooth densities give the same answer. For unrestricted families of curves this is false. The following example fixes both the family and the domain and does not use a limiting change of the question.

Proposition 3.3 (A null-support curve family)**.** [D03] *Let $U = (-1,1)^2$ and, for $n \ge 2$, let*

$$\gamma_n(t) = (1/n, t/n^2), \qquad 0 \le t \le 1, \quad \Gamma = \{\gamma_n : n \ge 2\}.$$

The supremum over smooth densities is zero, whereas the supremum over Borel densities is infinite.

Proof. All curves lie in the compact rectangle $K = [0,1/2] \times [0,1/4] \Subset U$. For any smooth density ρ, compactness gives a finite upper bound M_K. Thus $\ell_\rho(\gamma_n) \le M_K/n^2 \to 0$, so $\ell_\rho(\Gamma) = 0$. Every smooth quotient is zero.

The vertical segments are pairwise disjoint and their union E has planar measure zero. Define $\sigma = n^2$ on $\gamma_n([0,1])$ and $\sigma = 0$ elsewhere. This is a finite-valued Borel function and $\ell_\sigma(\gamma_n) = 1$ for every n, while $M(\sigma) = 0$. For $\varepsilon > 0$, set $\rho_\varepsilon = \sigma + \varepsilon$. Its energy is exactly $4\varepsilon^2$, because the difference from the constant density is supported on E. Its minimum curve length is 1. The quotient $1/(4\varepsilon^2)$ is unbounded. This proves the two claims. □

The example does not invalidate smooth-density formulas for regular condensers or the explicitly computed cylinder families. It refutes the unrestricted quantifier. A repair needs a specified curve class or an exceptional-family convention and its own approximation theorem. We do not cite ordinary L^2 density as that theorem. General measure-system modulus provides a framework for treating this distinction [4, 5].

3.4 Conformal covariance of the density

Theorem 3.4 (Conformal invariance)**.** [D04] *Let $f : U \to V$ be a biholomorphism and let Γ be a rectifiable curve family with the compatible transformed line-integral convention. Then $\mathsf{L}(f\Gamma) = \mathsf{L}(\Gamma)$ and $\mathsf{W}(f\Gamma) = \mathsf{W}(\Gamma)$.*

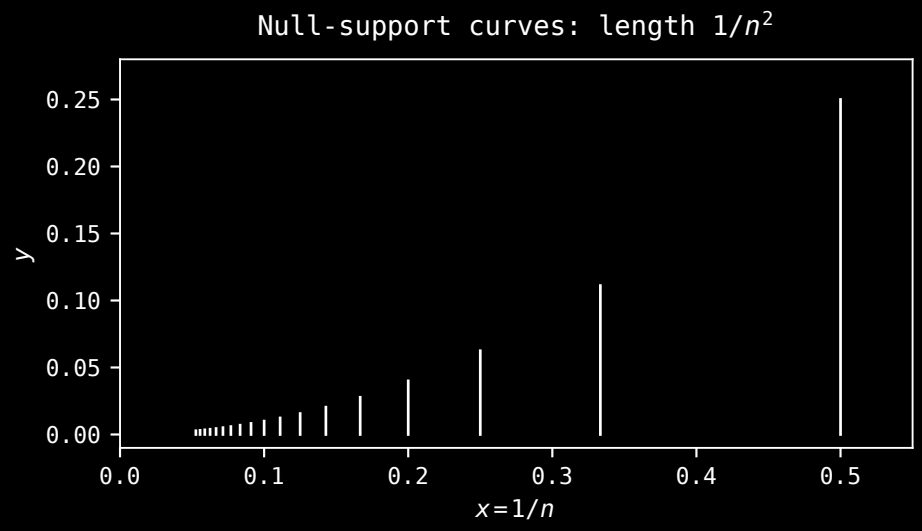


Figure 3.1: The counterexample uses shorter disjoint curves accumulating inside a compact set. A smooth density is bounded there. An unbounded Borel density supported on their area-null union can assign unit length to each curve.

Proof. For a target density τ, define $\rho(z) = \tau(f(z))|f'(z)|$. The chain rule for arclength gives $\ell_\rho(\gamma) = \ell_\tau(f \circ \gamma)$. The real Jacobian of a holomorphic local diffeomorphism is $|f'|^2$, so change of variables gives $M(\rho) = M(\tau)$. Pullback is bijective, with inverse induced by f^{-1}. It therefore preserves the complete sets of quotient values and admissible energies. The equality of their suprema and infima follows. □

This proof uses a change-of-variables theorem and curve regularity sufficient for the chain rule. A general family containing nonrectifiable curves needs conventions checked under the map. In every explicit example below the coordinate maps are smooth local diffeomorphisms on the relevant compact subsets, and improper limits are stated separately.

3.5 Convexity and uniqueness of normalized minimizers

If ρ and σ are admissible, so is $t\rho + (1-t)\sigma$ for $0 \le t \le 1$, because each path integral is linear and both are at least one. Moreover,

$$M\left(\frac{\rho+\sigma}{2}\right) = \frac{M(\rho)+M(\sigma)}{2} - \frac{1}{4}\|\rho-\sigma\|_{L^2}^2.$$

Consequently, if both attain the finite minimum width, they agree almost everywhere. This statement is uniqueness in the Hilbert norm, not pointwise uniqueness. It also does not prove existence: the admissible population need not be norm closed, as null-supported constraints demonstrate.

Scaling leaves the length quotient invariant. Thus a assertion that an extremal density is uniquely equal to one particular density must at least fix scale. Normalized width constraints supply that scale naturally, but pointwise representatives still need separate treatment.

3.6 Finite constraint systems as an exact model

For a finite set of m weighted edges and a finite path-incidence matrix N, the admissible set is

$$\{\rho \in \mathbb{R}^m : \rho \ge 0,\ N\rho \ge \mathbf{1}\}.$$

It is a closed polyhedron. For positive edge weights a_e, the energy $\sum_e a_e \rho_e^2$ is coercive and strictly convex. Existence follows by restricting a minimizing sequence to a compact sublevel set, and uniqueness follows from the displayed identity. Here no trace or exceptional-family issue occurs. A finite formulation with an exponential number of paths may still be computationally expensive; the graph-potential reduction later removes those constraints only after its exact equivalence is proved.

Chapter 4

Source §6.2: series, parallel and variational duality

Source correspondence. Lyubich, §6.2, Series Law, Parallel Law and Remark 6.5..

4.1 Series inequalities without assuming minimizers

A family Γ disjointly overflows $\Gamma_1, \dots, \Gamma_n$ when every $\gamma \in \Gamma$ contains pairwise disjoint subarcs $\gamma_i \in \Gamma_i$. Disjointness concerns arclength along the parametrized curve; a repeated traversal needs to be counted with its multiplicity. The ambient supports of the separate families need not be disjoint.

Theorem 4.1 (Series law)**.** [S01] *For finitely many disjointly overflowed families,*

$$\mathsf{L}(\Gamma) \geq \sum_{i=1}^{n} \mathsf{L}(\Gamma_i).$$

Proof. First choose densities σ_i with finite positive energy and finite positive minimum lengths a_i. Set $b_i = M(\sigma_i)$ and $\rho_i = (a_i/b_i)\sigma_i$. Then

$$\ell_{\rho_i}(\Gamma_i) = M(\rho_i) = a_i^2/b_i =: r_i.$$

Let $\rho = \max_i \rho_i$. On each of the disjoint subarcs, $\rho \geq \rho_i$, so $\ell_\rho(\gamma) \geq \sum_i r_i$. Pointwise,

$$\rho^2 = \max_i \rho_i^2 \leq \sum_i \rho_i^2,$$

whence $M(\rho) \leq \sum_i r_i$. The quotient is therefore at least $\sum_i r_i$. Taking independently chosen approximating densities gives the required supremum bound. A zero summand can be omitted. If a summand is infinite, choose its quotient arbitrarily large and repeat the finite argument; then the target extremal length is infinite. No extremal density is assumed to exist. □

The source proof's energy step must use $\max(\rho_1, \rho_2)^2 \leq \rho_1^2 + \rho_2^2$. The weaker inequality $\rho \leq \rho_1 + \rho_2$ does not itself imply the desired sum of energies, since squaring creates a cross term. The corrected algebra retains the source argument and its conclusion.

4.2 A weighted construction and its equality conditions

If the constituent families occupy disjoint regions and have attained extremal lengths $R_i > 0$, normalized densities can be chosen with minimum lengths and energies both equal to R_i. The

pasted density then has energy $\sum_i R_i$. Equality for the complete family additionally requires every relevant path to incur no unavoidable excess outside the chosen pieces and the lower-bound candidate to saturate the global extremal problem. Disjoint supports alone do not prove equality for an arbitrary constrained concatenation.

A concrete equality model is a stack of rectangles of the same width ℓ and heights h_i. Their union is a rectangle of height $H = \sum_i h_i$. Each vertical family has length h_i/ℓ and the global one has length H/ℓ. Both the geometry and the variational calculation give exact additivity. Altering the interface to a narrow aperture can only raise the resistance-like invariant; the interface constraint is not represented by a naive scalar sum.

4.3 Parallel inequalities and disjoint supports

Theorem 4.2 (Parallel law)**.** [S02] *For* $\Gamma = \bigcup_{i=1}^{n} \Gamma_i$,

$$\mathsf{W}(\Gamma) \le \sum_i \mathsf{W}(\Gamma_i).$$

If the families are supported in pairwise disjoint measurable sets, equality holds.

Proof. Choose admissible densities ρ_i of energy within a prescribed positive tolerance of each infimum. Their maximum is admissible for the union and has energy at most the sum of their energies. Passing to the infima proves the upper bound, with infinite endpoint cases immediate. For the reverse inequality in the disjoint-support case, restrict any global admissible ρ to each support X_i. The restrictions are admissible for Γ_i, and the sum of their energies is at most $M(\rho)$. Thus $M(\rho) \ge \sum_i \mathsf{W}(\Gamma_i)$. Taking the infimum proves equality. □

The reverse inequality generally fails if supports overlap. For example, taking $\Gamma_1 = \Gamma_2$ with finite positive width gives $\mathsf{W}(\Gamma_1 \cup \Gamma_2) = \mathsf{W}(\Gamma_1) < 2\mathsf{W}(\Gamma_1)$. This elementary control prevents an implementation from counting identical paths as independent parallel channels.

4.4 Electrical network interpretation

For two terminals at voltage difference V, a resistor of resistance R carries current $I = V/R$ and dissipates power $P = V^2/R$. If resistors are in series, conservation imposes the same current and their voltage drops add; hence $R = \sum_i R_i$. In parallel the voltage is common and currents add, so $1/R = \sum_i 1/R_i$. Extremal length and width reproduce these operations under the geometric hypotheses proved above.

The interpretation is not a proof that every geometric family is literally an electrical network. A conducting sheet requires material conductivity, thickness, terminal geometry and boundary conditions. For a two-dimensional homogeneous unit-conductivity model the normalized resistance equals the reciprocal of the unit-voltage Dirichlet minimum. Section 13 restores physical units.

4.5 A complete two-route calculation

Take two parallel branches, with resistances 1 and 2 in the first branch, and 2 and 4 in the second. The series–parallel calculation gives

$$R_1 = 3, \quad R_2 = 6, \quad C = \frac{1}{3} + \frac{1}{6} = \frac{1}{2}, \quad R = 2.$$

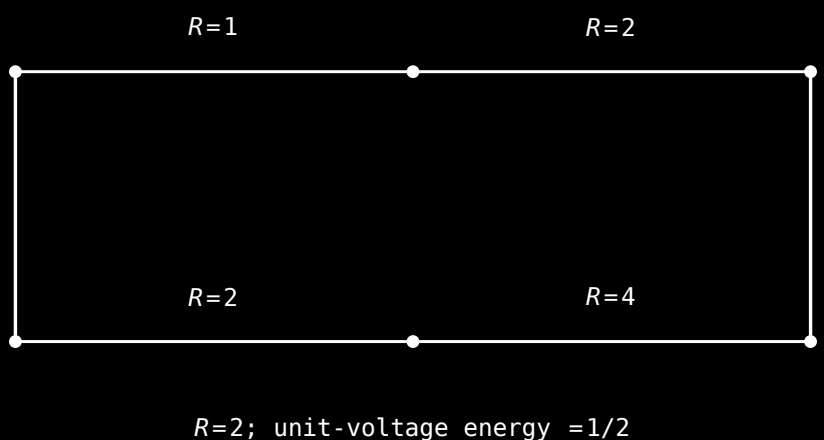


Figure 4.1: Two series branches followed by a parallel connection. The diagram labels an exact finite example; the proof of the graph reduction includes the actual terminal and conservation constraints.

An independent energy calculation fixes terminals at zero and one. Let x, y be the interior potentials. The energy is

$$E(x,y) = x^2 + \frac{(1-x)^2}{2} + \frac{y^2}{2} + \frac{(1-y)^2}{4}.$$

The stationary equations give $x = y = 1/3$. Substitution gives $E = 1/2$. Completing squares gives

$$E(x,y) = \frac{1}{2} + \frac{3}{2}(x - 1/3)^2 + \frac{3}{4}(y - 1/3)^2,$$

which proves global minimality and uniqueness without relying only on stationary points. A unit current splits as $2/3$ and $1/3$ between the branches; its resistance energy is $3(2/3)^2 + 6(1/3)^2 = 2$. The matching reciprocal values certify the optimum.

The scalar reduction uses a fixed number of rational operations. The direct quadratic route constructs two variables and solves a diagonal system. For this tiny instance neither route needs advanced complex analysis; its role is to calibrate the meaning of energy and reciprocity before using a surface theorem.

4.6 Countable families

For a countable union, finite partial unions and monotonicity give

$$\mathsf{W}\left(\bigcup_{i=1}^{\infty} \Gamma_i\right) \le \sum_i \mathsf{W}(\Gamma_i)$$

by choosing densities with summable error tolerances and taking their pointwise supremum. The energy bound follows from $\sup_i \rho_i^2 \le \sum_i \rho_i^2$ and monotone integration. For pairwise disjoint supports the reverse inequality follows from every finite partial family and then the supremum of the partial sums. Infinite series overflowing requires every finite collection of required subarcs; the finite series law then gives an infinite lower bound if the sum diverges.

This argument is about nonnegative integrals. It does not justify rearranging signed conditional sums or treating an infinite network as a finite one without specifying boundary behavior at infinity.

Chapter 5

Source §6.3: annuli, rectangles and tori

5.1 §6.3.1: the cylinder calculation

Source correspondence. Lyubich, Proposition 6.6 and Exercises 6.7–6.8..

Let $C_{\ell,h} = \mathbb{R}/\ell\mathbb{Z} \times (0,h)$. A vertical curve joins its ends; a horizontal curve represents one circuit of the circular direction. The metric $\rho = 1$ has area ℓh, vertical minimum length h, and horizontal minimum length ℓ.

Theorem 5.1 (Cylinder modulus)**.** [A01] *A vertical family containing almost every straight vertical fiber has*

$$\mathsf{L}(\Gamma_v) = \frac{h}{\ell}, \qquad \mathsf{W}(\Gamma_v) = \frac{\ell}{h}.$$

A horizontal family containing almost every circular fiber has $\mathsf{L}(\Gamma_h) = \ell/h$ *and* $\mathsf{W}(\Gamma_h) = h/\ell$.

Proof. The constant metric gives the lower vertical bound h/ℓ. For any finite-energy density, every included straight fiber has length at least $s = \ell_\rho(\Gamma_v)$. Integrating over the circle gives $\ell s \le \int_C \rho$. Integral Cauchy–Schwarz yields $(\ell s)^2 \le \ell h M(\rho)$, so the quotient is at most h/ℓ. The horizontal argument interchanges the roles of h and ℓ: integrate the circular-fiber lengths over height. The constant metric again attains equality. Reciprocity supplies the width formulas. □

A second proof minimizes the vertical energy among potentials u with values zero and one at the ends. For every x, $1 = \int_0^h u_y \,\mathrm{d}y$, so $\int_0^h u_y^2 \,\mathrm{d}y \ge 1/h$. Integrating over x gives $\mathcal{E}(u) \ge \ell/h$, with equality for $u = y/h$. The potential/curve equivalence is proved in the regular setting in Chapter 6; this route contributes the unique harmonic witness and a direct error identity.

5.2 Equality and quantitative rigidity

Proposition 5.2 (Normalized near-extremizers)**.** [A02] *Suppose* $M(\rho) = 1$ *on the cylinder and let* $r = \mathsf{L}_\rho(\Gamma_v)/(h/\ell)$. *Then* $0 \le r \le 1$, *and*

$$\left\| \rho - \frac{1}{\sqrt{\ell h}} \right\|_{L^2}^2 \le 2(1 - \sqrt{r}).$$

Proof. Write $A = \ell h$, $\mu = A^{-1} \int \rho$, and $s = \ell_\rho(\Gamma_v)$. The fiber inequality gives $A\mu \ge \ell s$, hence $\sqrt{A}\mu \ge \sqrt{\ell s^2/h} = \sqrt{r}$. Expanding the squared distance to the normalized constant gives $1 + 1 - 2\sqrt{A}\mu \le 2(1 - \sqrt{r})$. Cauchy–Schwarz gives $\sqrt{A}\mu \le 1$ and hence $r \le 1$. □

If $r = 1$, the normalized density equals the constant almost everywhere. Multiplying an extremal density by a positive scalar preserves its quotient, and null-set changes may preserve it too. Thus Exercise 6.7's uniqueness is correctly interpreted with scale and almost-everywhere qualifications. The displayed bound quantitatively strengthens that qualified equality statement; no research-priority claim is attached to it.

5.3 Logarithmic coordinates and exact annular energy

The covering coordinate $w = \log z = x + iy$ sends $A(r, R)$ to a cylinder of height $H = \log(R/r)$ and circumference 2π. Pulling back the constant cylinder density gives $\rho(z) = 1/|z|$. Therefore

$$\operatorname{mod} A(r, R) = \frac{H}{2\pi}, \quad u(z) = \frac{\log(|z|/r)}{H}, \quad \mathcal{E}(u) = \frac{2\pi}{H}. \tag{5.1}$$

Direct polar integration gives the last formula: $|\nabla u| = 1/(H|z|)$ and $\int_r^R \int_0^{2\pi} (H^2 s^2)^{-1} s \, \mathrm{d}\theta \, \mathrm{d}s = 2\pi/H$. It also gives the conjugate period $2\pi/H$. These independent representations prevent a hidden 2π from moving between modulus and capacity.

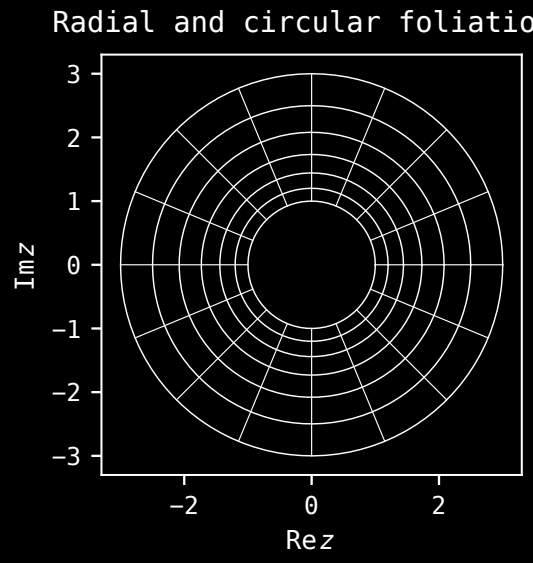


Figure 5.1: A round annulus and its radial/circular foliations. The logarithmic coordinate transforms their metric problem to a flat cylinder.

5.4 §6.3.2: the Grötzsch inequality

Source correspondence. Lyubich, Proposition 6.9..

Proposition 5.3. [A03] *If disjoint embedded conformal annuli A_i lie essentially in A, then* $\operatorname{mod} A \geq \sum_i \operatorname{mod} A_i$.

Proof. Every essential horizontal curve in A_i is horizontal in A. The union of these subfamilies has disjoint supports, hence its width is the sum of the widths $\operatorname{mod} A_i$. Monotonicity under inclusion in the full horizontal family of A gives $\sum_i \operatorname{mod} A_i \leq \operatorname{mod} A$. □

For concentric annuli this reduces to additivity of logarithms. If $r = r_0 < r_1 < \cdots < r_n = R$, their open interiors are disjoint and the sum $\sum_i (2\pi)^{-1} \log(r_i/r_{i-1})$ is exactly the parent modulus.

This is an equality case without any approximation. Missing shells make the sum smaller by their total logarithmic thickness.

5.5 §6.3.3: Euclidean separation and normalization

Source correspondence. Lyubich, Lemma 6.10, Corollaries 6.11–6.12, Exercise 6.13 and Proposition 6.14..

Let K be a compact connected inner complement of diameter $D > 0$, and let the outer complement Q be at distance $d > 0$. The density $\rho = d^{-1}$ on the d-neighborhood of K outside K, and zero elsewhere, is admissible for connecting curves. A path from K to Q must traverse distance d before its first contact with Q. If a disk of radius D about a point of K contains K, the support is contained in a disk of radius $D + d$. Consequently

$$\operatorname{mod} A \geq \frac{d^2}{\pi(D+d)^2}. \tag{5.2}$$

The value is not sharp, but it is explicit and dimensionless. It proves a positive lower bound when d/D is bounded below, without relying on an implicit unnamed function.

The source upper estimate is established with explicit constants in Proposition 24.1. Its tube witness treats arbitrary rectifiable winding curves without a Jordan-reduction assumption and gives $m \leq 2 + 2\pi + 4d/D$. Theorem 24.2 supplies the additional logarithmic upper bound used for diameter decay.

Exercise 6.13(i), interpreted in Euclidean coordinates without an outer-size normalization, cannot have absolute constants M, α in $D \leq Me^{-\alpha\mu}$. Scale a fixed annulus of modulus μ by $t > 0$. Its modulus is unchanged and its inner diameter tends to infinity with t. A correct estimate must include a geometric scale, such as an outer radius, or use a fixed spherical normalization. For round annuli, $D/R = 2e^{-2\pi\mu}$ is the exact scale-covariant identity.

For Corollary 6.12, a concentric annulus constructed from D and d has modulus $(2\pi)^{-1}\log(d/D)$ when $d > D$. Substituting only the preceding linear upper estimate gives a logarithm with the corresponding factor and additive constants. Those constants must be retained; the source's unqualified $\log\mu$ conclusion is not obtained by simply dropping them. No counterexample to a separately normalized stronger theorem is asserted here.

The pinching comparison in Proposition 6.14 concerns a spherical configuration. Spherical distance and diameter, the admissible continua, and the treatment of singleton complements must be fixed. Theorem 24.4 supplies explicit two-sided functions for arbitrary nondegenerate complementary continua using unit-sphere chordal distance. The statement excludes singleton components before forming the relative ratio.

5.6 Problem 6.15: an area inequality with a stronger bound

For a smooth bounded ring domain with inner filled component K of positive area, let u be its regular condenser potential, $C = \mathcal{E}(u)$, and $A(t) = \operatorname{area}(K \cup \{u < t\})$. Assume the usual coarea passage over critical levels. At a regular level, write $P(t)$ for total perimeter. Cauchy–Schwarz gives

$$P(t)^2 \leq \left(\int_{u=t} |\nabla u|\,\mathrm{d}s\right)\left(\int_{u=t} |\nabla u|^{-1}\,\mathrm{d}s\right) = CA'(t).$$

The planar isoperimetric inequality gives $P(t)^2 \geq 4\pi A(t)$, and thus

$$\frac{A'(t)}{A(t)} \geq \frac{4\pi}{C}.$$

Integration from zero to one yields

$$\frac{\operatorname{area}(A \cup K)}{\operatorname{area} K} \geq e^{4\pi \bmod A} \geq 1 + 4\pi \bmod A. \tag{5.3}$$

For concentric circular boundaries the first inequality is equality. The source asks only for the weaker linear lower bound.

This is a conditional coarea/isoperimetry route. Chapter 25 gives a second complete route through Laurent coefficients for arbitrary bounded ring complements of positive finite inner area and finite outer area. It proves the stronger exponential estimate and an attained deficit bound. A zero-area denominator remains outside the ratio formulation.

5.7 §6.3.4 and §6.3.6: divergent nests

Source correspondence. Lyubich, Proposition 6.16 and Corollary 6.20..

A decreasing nest of nonempty compact connected sets K_n in a fixed compact set has nonempty connected intersection K. The set in the source divergence statement must be the intersection; a union of decreasing sets is just the first set and need not shrink. Once the geometric theorem that a ring with two nondegenerate complementary continua has finite modulus is available, Grötzsch proves the divergence criterion: arbitrarily large finite sums of annular moduli force the outer-to-limit ring to have infinite modulus, so its inner continuum is a point.

For a complete constructive model, take $r_n = r_0 \exp(-2\pi \sum_{j<n} m_j)$ and $K_n = \{|z| \leq r_n\}$, where $m_j > 0$. The annulus between K_j and K_{j+1} has modulus m_j. If $\sum m_j = \infty$, then $r_n \to 0$ and $\bigcap K_n = \{0\}$. If the sum is finite, the intersection is a disk of positive radius. This gives both directions in the round model and a computable stopping criterion $\sum_{j<n} m_j \geq (2\pi)^{-1} \log(r_0/\varepsilon)$ for radius at most ε.

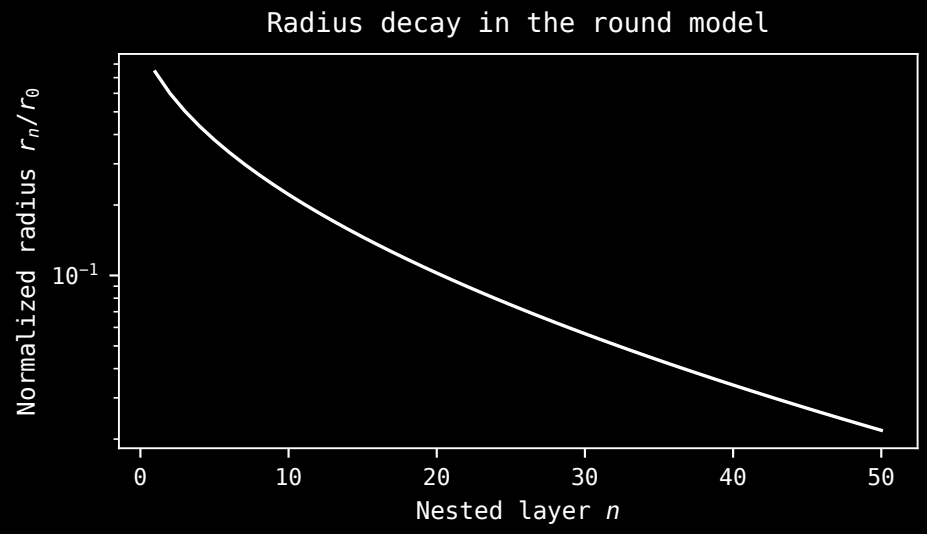


Figure 5.2: A logarithmic radius schedule gives exact annular moduli. Divergence of the sum is equivalent to point collapse for this round nested family.

5.8 §6.3.5: real-symmetric interval pairs

Source correspondence. Lyubich, Lemma 6.17, Corollary 6.18 and Exercise 6.19..

For $L = [a, b] \Subset I = [c, d]$, the dimensionless real buffer is

$$\delta = \frac{\min(a - c, d - b)}{b - a}.$$

Translation and positive rescaling preserve δ. The corresponding slit-surface modulus is a conformal invariant. A valid equivalence between a positive lower bound on δ and a positive lower bound on the slit modulus must quantify both directions uniformly over all normalized endpoint configurations. The positivity statement alone is weaker than an explicit sharp formula.

A purely real check of a proposed endpoint transformation is often preferable for a single affine map: if $\phi(x) = sx + t$ with $s > 0$, all three lengths are multiplied by s and δ is exactly invariant. A non-affine univalent map requires a distortion bound on a surrounding complex neighborhood. If $m \le |\phi'| \le M$ on the real interval and the map is increasing there, integration gives $\delta' \ge (m/M)\delta$. This is the explicit reduction requested by Exercise 6.19 once a suitable Koebe bound supplies M/m. The existence and size of that surrounding neighborhood are indispensable; monotonicity of an arbitrary real diffeomorphism does not supply a uniform ratio.

5.9 §6.3.7: rectangles and laminations

Source correspondence. Lyubich, Exercises 6.21 and 6.23; Proposition 6.22..

Let $E \subset [0, \ell]$ be measurable with measure $a > 0$, and let Γ_E consist of the vertical segments above E in a rectangle of height h. The metric $\rho = 1$ on $E \times (0, h)$ and zero elsewhere gives quotient h/a. Conversely,

$$a\,\ell_\rho(\Gamma_E) \le \int_{E\times(0,h)} \rho \le \sqrt{ahM(\rho)},$$

so $\mathsf{L}(\Gamma_E) = h/a$. The symbol denoting the projection in the source formula must mean its measure, not the set itself. For $a = 0$, a density supported on the union of the fibers can make the width zero under the literal Borel formulation.

The existence of a rectangular uniformization of a marked quadrilateral uses a Riemann map and an explicit four-point Schwarz–Christoffel integral. Uniqueness of its aspect ratio follows from conformal invariance and the rectangle formula. Thus existence and uniqueness have different prerequisites. This review uses the previous conformal-mapping developments as references and supplies the integral's numerical role in Chapter 12; it does not claim a new complete prime-end proof for arbitrary rough quadrilaterals.

The source bounds for $\operatorname{mod} \mathbb{H}(0, 1, R)$ must retain the side marking and the position of infinity. Different choices of opposite sides reciprocate the modulus. Chapter 29 constructs the rectangle coordinate and proves stronger inequalities on the whole interval $R > 1$. It also provides a transcendental-modulus enclosure from finite rational witnesses; no finite sampling is substituted for the universal inequality.

5.10 §6.3.8: tori and homology classes

Source correspondence. Lyubich, Exercise 6.24 and Proposition 6.25..

Let $T_\tau = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ with $\operatorname{Im}\tau = y > 0$. For an integral class $(m, n) \ne (0, 0)$, let $v = m + n\tau$. Closed representative curves lift to paths joining z to $z + v$, so their Euclidean length is at least $|v|$. Constant density therefore gives $\mathsf{L} \ge |v|^2/y$.

Proposition 5.4 (Torus class formula)**.** [A04] *For the family of representatives of (m, n), including all translated straight representatives with traversal multiplicity,*

$$\mathsf{L}_\tau(m, n) = \frac{|m + n\tau|^2}{\operatorname{Im}\tau}, \qquad \mathsf{W}_\tau(m, n) = \frac{\operatorname{Im}\tau}{|m + n\tau|^2}.$$

Proof. For a density ρ, integrate lengths of translated straight representatives over uniform starting points $z \in T_\tau$. Translation invariance gives their mean length $|v|y^{-1}\int_T \rho$. The minimum length over all representatives is no greater than this mean. Cauchy–Schwarz gives $\ell_\rho(\Gamma)^2 \le |v|^2 M(\rho)/y$. This proves the upper bound and matches the constant-density lower bound. Reciprocity gives width. The calculation includes nonprimitive classes because the straight representative is then multiply traversed, with length multiplied by the same integer. □

For disjoint embedded annuli representing a primitive class, their horizontal families lie in that torus class. Parallel additivity and monotonicity therefore prove $\sum_i \operatorname{mod} A_i \le y/|v|^2$. A nonprimitive class cannot be represented by the core of an embedded annulus; this topological requirement must be stated rather than hidden in the same notation.

Chapter 6

Source §6.4: energy, capacity and dual witnesses

6.1 §6.4.1: normalization audit

Source correspondence. Lyubich, Exercise 6.26 and the energy definitions on source page 126..

In a conformal metric $g = \lambda^2(dx^2+dy^2)$, the squared gradient norm of a scalar u is $\lambda^{-2}(u_x^2+u_y^2)$ and the area form is $\lambda^2\,\mathrm{d}x\,\mathrm{d}y$. Their product is independent of λ. Thus

$$\mathcal{E}_g(u) = \int |\nabla_g u|_g^2 \,\mathrm{d}A_g = \int (u_x^2 + u_y^2)\,\mathrm{d}x\,\mathrm{d}y.$$

The squared norm is essential. An integral of $|\nabla_g u|_g$ would acquire a factor λ under conformal scaling. The rendered source's first formula omits the square, while the subsequent coordinate expression uses a quadratic energy. This manuscript adopts the quadratic formula explicitly.

For a complex f, direct expansion gives

$$|f_x|^2 + |f_y|^2 = 2(|\partial_z f|^2 + |\partial_{\bar z} f|^2). \tag{6.1}$$

Indeed $f_x = f_z + f_{\bar z}$ and $f_y = i(f_z - f_{\bar z})$, so the cross terms cancel. If f is holomorphic, its full real energy is $2\int |f'|^2$, twice its area counted with multiplicity. The half-energy of a holomorphic embedding is image area. The source expressions cannot all be simultaneously correct under the standard half-derivative convention. Subsequent scalar capacity formulas use the full scalar energy, with no one-half factor.

6.2 §6.4.2: proper harmonic functions on a finite cylinder

Proposition 6.1. [E01] *On a cylinder of height h and circumference ℓ, the harmonic function with boundary values zero and one is $u = y/h$, uniquely. Its energy is ℓ/h.*

Proof. The displayed function is harmonic and has the prescribed traces. The difference of two continuous-up-to-boundary harmonic solutions has zero boundary values; the maximum principle applied to the difference and its negative makes it zero. Direct differentiation and integration give $\int h^{-2}\,\mathrm{d}A = \ell/h$. □

A proper harmonic map from a finite annulus to $(0,1)$ tends uniformly to one of the two endpoints along each end. To justify this, properness excludes values in any fixed compact subinterval sufficiently near an end; connectedness of a sufficiently small end collar forces a single choice of side of that subinterval. Nesting collars makes the choice consistent. The

two ends cannot both tend to the same endpoint, since the maximum principle on truncated cylinders would force a constant limit. The two assignments therefore give exactly u and $1-u$. A punctured disk has infinite modulus and requires separate treatment: a bounded harmonic function removable at the puncture cannot produce the same nonconstant proper finite-cylinder model.

6.3 The regular condenser and its terminal conditions

Let S be a compact smooth surface with boundary split into nonempty unions B_0 and B_1. Take real potentials with traces zero on B_0 and one on B_1. The capacity is the minimum of $\mathcal{E}$. Existence in an affine Sobolev class follows from the coercive projection argument once the trace and Poincaré properties are established. The resulting weak solution is harmonic. For smooth data on smooth boundary, the regular representative gives the classical condenser potential.

Every curve connecting the boundary sets has

$$1 = |u(\gamma(1)) - u(\gamma(0))| \le \int_\gamma |\nabla u| \, \mathrm{d}s.$$

Thus $|\nabla u|$ is a path-admissible density and $\mathsf{W}(\Gamma_v) \le \mathrm{Cap}$. The reverse inequality needs a construction converting an arbitrary admissible metric into a potential with no larger energy, or an equivalent modulus–capacity theorem. It is not obtained by only exhibiting the harmonic function.

For a continuous positive metric, define the truncated distance

$$v(x) = \min\left(1, \inf_{\gamma : B_0 \to x} \int_\gamma \rho \, \mathrm{d}s\right).$$

Path concatenation gives the local upper-gradient bound $|\nabla v| \le \rho$ almost everywhere, and admissibility gives trace one at B_1. Hence $\mathrm{Cap} \le \mathcal{E}(v) \le M(\rho)$. Extension to all Borel admissible metrics needs the appropriate weak-gradient and trace theorem. This manuscript records that theorem as an unresolved supporting dependency in full rough generality, rather than using the false unrestricted smooth-density exercise to bypass it.

6.4 §6.4.3: the separating family

For a regular harmonic condenser potential u, put $C = \mathcal{E}(u)$. The conserved flux through a regular separating level set is C. On any separating oriented multicurve, the divergence theorem gives signed total flux C, so $\int_\gamma |\nabla u| \ge C$. The density $|\nabla u|/C$ is therefore admissible for separating curves and has energy $1/C$.

Conversely, if ρ is admissible for every separating level multicurve, then $\int_{u=t} \rho \, \mathrm{d}s \ge 1$ for regular t. Integrating coarea gives

$$1 \le \int_S \rho |\nabla u| \, \mathrm{d}A \le \sqrt{M(\rho) C}.$$

Thus $M(\rho) \ge 1/C$, proving $\mathsf{W}(\Gamma_h) = 1/C$ in the regular setting. Together with the connecting-family identity,

$$\mathsf{L}(\Gamma_v) = \mathsf{W}(\Gamma_h) = 1/C, \quad \mathsf{W}(\Gamma_v) = \mathsf{L}(\Gamma_h) = C.$$

Chapter 28 proves the regularity, finite-critical-set, flow-coordinate and separating-multicurve statements for compact analytic-bordered condensers. Its precise class supplies a completed setting for the formula here; other boundary classes require their own trace/regularity argument. The source writes $D(h)$ after naming the potential χ; the normalized statement here uses one symbol consistently.

6.5 A continuum complementary-energy certificate

Here the certificate flux is oriented toward increasing numerical potential. Conventional physical current uses the opposite sign when the physical voltage is identified with that potential; energy and absolute conductance are unchanged by reversing this orientation.

Let $A(x)$ be a symmetric uniformly positive conductivity tensor. For a potential v with the specified traces and a vector field J with weak divergence zero and unit flux from the zero terminal to the one terminal, integration by parts gives

$$\int_S J \cdot \nabla v \, \mathrm{d}A = 1.$$

The pointwise weighted Cauchy–Schwarz inequality and its integrated form imply

$$1 \leq \left(\int \nabla v^T A \nabla v \right) \left(\int J^T A^{-1} J \right).$$

Consequently every feasible pair supplies

$$\frac{1}{\int J^T A^{-1} J} \leq C_A \leq \int \nabla v^T A \nabla v. \tag{6.2}$$

If u is the minimizer, the flux $J = A\nabla u / C_A$ attains equality and has the required unit total flux. This proves both variational directions for the conductivity model, assuming the integration-by-parts and trace framework.

The checker must establish both feasibility conditions. A small divergence residual is not exact feasibility. One may equilibrate a flux, or bound its residual pairing using a coercivity estimate, but the uncorrected flux cannot be substituted into the exact lower bound. Hypercircle methods are a standard context for this construction [9]; the finite graph version below supplies a fully explicit arithmetical checker.

6.6 Energy error and flux error as orthogonal quantities

Let u be the exact potential, v a feasible potential, and q a feasible physical flux with the same total boundary-flux convention as $A\nabla u$. If $e = v - u$ and $r = q - A\nabla u$, then $\int r \cdot \nabla e = 0$ by the zero-divergence and homogeneous-boundary conditions. Therefore

$$\|A^{1/2}\nabla v - A^{-1/2} q\|_{L^2}^2 = \|A^{1/2}\nabla e\|_{L^2}^2 + \|A^{-1/2} r\|_{L^2}^2.$$

This identity bounds the primal energy-norm error by a computable discrepancy. Its strength is not merely that two numerical methods agree; it is that feasible primal and dual spaces make the cross term vanish. If both fields are generated by the same unverified discretization and neither is equilibrated, agreement can conceal the same bias.

6.7 Physical capacity, resistance and stored energy

For homogeneous permittivity ϵ and a two-dimensional cross-section invariant along an axial length L_z, a dimensionless unit-potential capacity C yields capacitance $C_{\mathrm{phys}} = \epsilon L_z C$. With voltage V, stored energy is $\frac{1}{2} C_{\mathrm{phys}} V^2$ and total charge is $C_{\mathrm{phys}} V$. For conductivity σ in the corresponding stationary conduction model, conductance is $\sigma L_z C$.

The source's physical paragraph identifies an inverse energy with capacity and refers to uniform charge. In a general conductor geometry the equilibrium boundary charge need not be uniform. Under the conventions just specified, C is capacity/conductance-like and $1/C$ is resistance-like. These explicit units and definitions replace the ambiguous identification; no change is made silently.

Chapter 7

Source §§6.5–6.6: crossing and holomorphic transformation

7.1 Crossing as a quantified condition

Source correspondence. Lyubich, Non-Crossing Principle, Exercise 6.31 and Small Overlapping Principle..

Two families cross when every curve in one intersects every curve in the other in the specified geometric sense. This is stronger than the existence of two intersecting curves, and weaker than the assertion that every pair of arbitrary continuations has the same intersection number. The source uses horizontal and vertical foliations in a quadrilateral or annulus, where topological separation translates intersection into an overflowing statement.

Proposition 7.1 (Width obstruction)**.** [X01] *Let* Γ *be the complete horizontal foliation of a conformally marked rectangle or annulus. Suppose every locally rectifiable curve of* Δ *meets every open leaf of* Γ*. Then*

$$\mathsf{W}(\Gamma)\mathsf{W}(\Delta) \le 1.$$

Proof. Use flat coordinates with horizontal length ℓ and vertical height h on the supporting domain Q. The horizontal foliation has width h/ℓ. Extend its flat metric density by zero outside Q; its area is ℓh.

For a curve $\delta \in \Delta$, the inverse image of Q is a countable union of open parameter intervals. On each compact subinterval the flat height coordinate is absolutely continuous after arclength parametrization. The measure of its image is at most its total variation, which is bounded by the curve's length in the flat metric. Exhaust each open component by compact subintervals and use monotone limits of images and length. Since δ meets every open horizontal leaf, the union of these height images contains $(0, h)$. Subadditivity of one-dimensional measure therefore gives a total flat length at least h on the portions inside Q. Hence the density divided by h is admissible for Δ, and $\mathsf{W}(\Delta) \le \ell/h$. Multiplication proves the result. On an annulus the periodic horizontal coordinate causes no difficulty: its height is single valued and has unit gradient in the flat metric. □

The source's strict Non-Crossing Principle is the contrapositive. This proof does not require that one contiguous portion of a curve runs across the entire support: countably many interior portions are treated through their height images. The definition uses open leaves; a curve meeting only added endpoints is a different crossing relation.

7.2 Equality and what the product alone does not prove

In a rectangle $[0,\ell]\times[0,h]$, horizontal and vertical foliations have widths h/ℓ and ℓ/h, whose product is one. Thus the strict threshold cannot be decreased. The equality case in a general embedded domain requires equality in the overflowing and metric inequalities and a rigidity argument for the support. A product equal to one alone supplies none of those equality conditions individually.

The literal equality assertion of Exercise 6.31 is false: Proposition 26.5 gives both a planar mixed annulus–quadrilateral counterexample and a torus counterexample. The separately restricted equality-rigidity statement for two Jordan quadrilateral buffers, needed in the later canonical-rectangle construction, is not resolved by these counterexamples and remains open here. A correction must state its actual support class.

7.3 A quantitative overlap estimate

Let $E\subset[0,a]$ be Borel with measure k, and consider the vertical lamination $\Lambda=\{\{x\}\times(0,h):x\in E\}$ in a rectangle. Its width is k/h. Suppose $k=\kappa$, $h<1$, and every curve in another family Γ intersects leaves from a subset of E of measure at least one. Define $\rho=1$ on $E\times(0,h)$ and zero elsewhere. On the portions of every such curve inside the support, its attained horizontal values cover the measured set of intersected leaves. The one-dimensional area inequality bounds the measure of those values by the integral of the absolute horizontal derivative there. This is at most its ρ-length, which is therefore at least one. Its energy is $kh<\kappa$. Therefore $\mathsf{W}(\Gamma)<\kappa$.

The contrapositive is the Small Overlapping Principle under the given rectangle and transverse-measure model. A rectifiable curve has absolutely continuous coordinate functions after arclength parametrization, and the one-dimensional area formula bounds the measure of the attained horizontal values by their total variation. This is the analytical justification for the transverse-length estimate. The source's area line must read kh; replacing k by the width k/h changes the expression. We retain the correct units throughout.

For an annular supporting domain, replace the horizontal coordinate by its periodic circle coordinate. The one-dimensional length measure of its attained values is again at most the integral of its absolute metric derivative on the supported portions; lifting each parameter interval to a real coordinate proves this inequality. The same energy kh and contradiction therefore apply to the annular product-lamination case.

If an irregular lamination is not a full product $E\times(0,h)$ in the chosen uniformization, the reduction must establish the corresponding transverse measure and geometric identification. An arbitrary collection of partially broken vertical segments is a different family and need not satisfy the same formula.

7.4 §6.6.1: images and multiplicity

Source correspondence. Lyubich, Lemma 6.32 and Corollary 6.33..

Let $f:U\to V$ be a nonconstant holomorphic map. Images of paths are parametrized images, so a path that winds several times remains multiply traversed. For a target density τ, pull back $\rho=(\tau\circ f)|f'|$. Then $\ell_\rho(\gamma)=\ell_\tau(f\gamma)$ and the area formula gives

$$M(\rho)=\int_V N_f(w)\tau(w)^2\,\mathrm{d}A(w).$$

If $N_f\le d$ almost everywhere, then $M(\rho)\le dM(\tau)$.

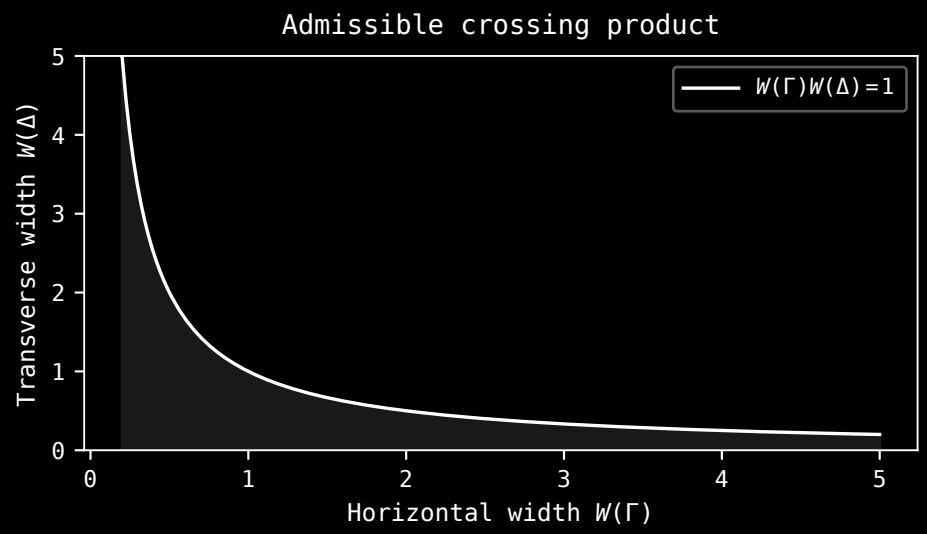


Figure 7.1: The threshold product for rectangular transverse families and its equality case. This is a geometric calculation; it is not a numerical test of all possible embedded supports.

Proposition 7.2 (Degree-dependent upper bound)**.** [XO2] *Under these regularity and multiplicity conventions,*

$$\mathsf{L}(f\Gamma) \le d\mathsf{L}(\Gamma).$$

Proof. For every target density with positive minimum length, the pullback has the same minimum length and energy at most d times as large. Hence $\mathsf{L}(\Gamma) \ge d^{-1}\ell_\tau(f\Gamma)^2/M(\tau)$. Taking the supremum over τ proves the assertion. Degenerate zero-length cases give the same inequality directly. □

Proposition 7.3 (Countable-sheet lower bound)**.** [T10] *For a nonconstant holomorphic map between second-countable Riemann surfaces and a family of locally rectifiable parametrized curves,*

$$\mathsf{L}(f\Gamma) \ge \mathsf{L}(\Gamma).$$

No finite degree bound is required for this lower inequality.

Proof. The critical points form a discrete countable set. Cover their complement by countably many injective coordinate neighborhoods and turn the cover into a disjoint Borel partition by removing the preceding pieces. Each piece has a continuous local inverse on its ambient image chart; its Borel restriction therefore defines a measurable summand. For a source Borel density define, off the countable set of critical values,

$$\tau(w)^2 = \sum_{z\in f^{-1}(w)} \frac{\rho(z)^2}{|f'(z)|^2}.$$

Set the density to zero at the exceptional target points. Tonelli's theorem and ordinary change of variables on the disjoint pieces give $M(\tau) = M(\rho)$. Every regular source point satisfies $\tau(f(z))|f'(z)| \ge \rho(z)$.

A locally rectifiable curve parametrized by arclength spends zero arclength at any prescribed point: an absolutely continuous coordinate map has derivative zero almost everywhere on the inverse image of that point, whereas arclength speed is one almost everywhere on its nonconstant parametrized part. The countable preimage of the countable critical-value set has the same zero-arclength property. Thus the pointwise inequality can be integrated on every curve without loss at these points. The chain rule and parametrized-image convention give $\ell_\tau(f\gamma) \ge \ell_\rho(\gamma)$.

Taking infima and then the extremal supremum proves the assertion. Exhaust noncompact parameter intervals to handle infinite line integrals. □

This proof uses locally rectifiable curves and their actual arclength integrals. It does not silently extend a different convention that assigns infinite length to every nonrectifiable curve regardless of its density. Critical-value exclusions change neither the area integral nor the line integrals in the stated class.

A constant map is excluded: it sends every curve to a constant curve, of extremal length zero, and cannot satisfy the source's lower inequality for a positive-length family. This supplies a simple audit of the implicit nonconstancy hypothesis.

7.5 Sharpness on power maps

Let $f(z) = z^d$ map $A(r, R)$ onto $A(r^d, R^d)$. For vertical families, target modulus is d times source modulus, attaining the upper multiplicity bound. For horizontal families, a source once-winding loop maps to a d-fold target loop. Its extremal length is d^2 times the once-winding target value, again equal to d times the source horizontal value. Forgetting traversal multiplicity would reverse this conclusion and produce an apparent contradiction.

The area change can be checked directly in polar coordinates, independently of covering terminology. A derivative has modulus $d|z|^{d-1}$ and its Jacobian is the square. The radial substitution $s = |z|^d$ and the d angular sheets yield the same multiplicity factor.

7.6 Lifting and overflowing

Suppose each target curve contains an arc that is the image of a source curve in Γ. Then the target family overflows $f\Gamma$, so its extremal length is at least $\mathsf{L}(f\Gamma)$. Combining with the lower transformation bound gives Corollary 6.33 in its qualified setting. This proof separates the topology of available lifts from the analytic area bound. A lift to a different component or one that fails the required terminal condition is not a qualifying source curve.

7.7 §6.6.2: holes under a finite proper map

Source correspondence. Lyubich, Proposition 6.34 and Exercise 6.35..

Let $f : U \to U'$ be a proper degree-D holomorphic map of disks. Let B' be a relatively compact Jordan disk and B one component of its inverse image, with local degree d. The target connecting family on $U' \setminus \overline{B'}$ has width C'. Extend its admissible density by zero over B' and pull it back to $U \setminus \overline{B}$. Every source connecting path has an initial segment whose image connects the outer boundary to B'; its pulled-back length is at least one. The pulled-back energy is at most DC'. Hence $C \le DC'$.

For the lower bound choose target vertical leaves avoiding critical values and cut the annulus into strips along the finitely many exceptional leaves, with an additional cut if necessary. Each strip lifts adjacent to B in d disjoint conformal copies. A lift remains outside every preimage of B' and, by properness, reaches the outer end. The widths of the target strips add to C'; the d lifted copies have total width dC'. They form a subfamily of the source connecting family, so monotonicity gives $C \ge dC'$.

This proof requires regular boundary/critical-leaf approximation to justify the strip decomposition. It explains why the local and total degrees enter different sides of the inequality. It does not replace d by D merely because the global map has degree D.

7.8 An independent covering proof using energy

If a finite covering of degree d maps an entire source condenser to a target condenser, respecting both terminal partitions, compose the target harmonic potential with the covering. The pullback is harmonic and satisfies the source boundary conditions; uniqueness identifies it with the source potential. Its energy is dC', so

$$C = dC', \qquad \operatorname{mod}(S') = d \operatorname{mod}(S).$$

This gives the equality part of Exercise 6.35 directly for multiple terminal components once the regular condenser theorem has been established. Ramification away from the boundary can also be handled through the harmonic pullback and area formula under an appropriate proper-map model, but a map that changes the terminal partition is not covered by this statement.

Chapter 8

Source §§6.7–6.8: canonical annuli and arcs

8.1 Maximal modulus and inclusion maximality

Source correspondence. Lyubich, Lemma 6.36, Exercise 6.37 and Theorem 6.38..

Fix an essential free homotopy class of simple closed curves in a Riemann surface S. An embedded annulus represents the class when its primitive core curve does. There are two different maximality notions: maximizing the modulus among all such embeddings, and being maximal under set inclusion. Neither should be substituted for the other. The source explicitly defines maximal annuli through their modulus.

An existence proof by taking a supremum of moduli must do more than invoke a normal family. It must normalize the maps, show that the limit remains nonconstant and injective, show that it remains in the intended surface and homotopy class, and show that its limiting domain attains the supremum. A sequence can converge to a constant or escape to a boundary in an unnormalized formulation. Chapter 27 discharges these steps: a four-point noncollapse argument, a direct cusp construction, and an injectivity check after cyclic-cover projection give attainment in the stated Riemann-surface setting, relative to the identified earlier uniformization and normal-family dependencies.

Proposition 8.1 (Puncture counterexample to unrestricted uniqueness)**.** [G01] *In* $S = \widehat{\mathbb{C}} \setminus \{0, 1, \infty\}$*, the free homotopy class surrounding* 0 *contains infinitely many distinct embedded annuli of maximal modulus, when infinite modulus is admitted.*

Proof. For every $0 < r < 1$, the punctured disk $A_r = \{0 < |z| < r\}$ is an embedded annulus in S surrounding 0. Its modulus is infinite: the compact subannuli $A(\varepsilon, r)$ have moduli $(2\pi)^{-1} \log(r/\varepsilon) \to \infty$, and monotonicity gives infinite modulus for A_r. Infinite modulus is the largest possible extended value, and the A_r are distinct. Thus there is no unique maximizing annulus in this class under the literal modulus definition. □

The example resolves the unrestricted assertion at punctures. It does not prove or disprove the usual nonperipheral finite-modulus uniqueness theorem. That theorem is related to Jenkins–Strebel quadratic differentials, whose specific existence and uniqueness development is a remaining supporting obligation. Primary sources on such differentials distinguish prescribed puncture behavior and moduli data [7, 8].

8.2 A torus equality family

On $T = \mathbb{C}/(\mathbb{Z} + ih\mathbb{Z})$, remove the horizontal closed geodesic at height s. The complement is an open cylinder of circumference one and height h, hence has modulus h. Its core represents the horizontal primitive class. The torus inequality in Chapter 5 bounds the modulus of every embedded annulus in that class by h. Thus every such cut cylinder is a maximizer. Varying s changes the embedded annulus.

This is the flat case of the general existence result in Corollary 27.5, together with its explicit failure of uniqueness. It also explains why a generalization from planar domains to all Riemann surfaces must treat tori separately. The variational value is unique even when the maximizing embedded region is not.

8.3 Quadratic differentials and the missing global step

A quadratic differential is locally $q(z)\,\mathrm{d}z^2$ and transforms by the square of the coordinate derivative. Away from zeros and poles choose a local square root and integrate it to obtain a coordinate $w = \int \sqrt{q}\,\mathrm{d}z$. In this coordinate the metric is $|\,\mathrm{d}w|$ and horizontal trajectories satisfy $\operatorname{Im} w = \text{constant}$. Closed trajectories can form flat cylinders. At a zero of order k, the integral behaves like $z^{(k+2)/2}$, producing a cone angle $(k+2)\pi$.

These local facts follow from the holomorphic power-series factorization and the chain rule. They do not produce a global differential with prescribed trajectory classes. A global existence theorem must establish compatibility across the surface, the behavior at punctures and boundaries, finite area where required, and the appropriate extremality condition. Identifying a plausible local differential is not enough.

The manuscript therefore retains a separate global obligation: derive the precise finite-type nonperipheral maximal-annulus theorem from the stated foundational boundary and then examine the extension to arbitrary planar domains. No subsequent completed finite-network theorem depends on this missing result. The conditional geometric propositions below are tagged accordingly.

8.4 Buffers and the disjointness mechanism

Source correspondence. Lyubich, Lemma 6.39..

On a maximal annulus of modulus $h > 2M$, choose cylinder coordinates of circumference one and height h. Removing bands of height M from both ends leaves a canonical middle cylinder of height $h - 2M$. Its modulus is therefore exactly $h - 2M$.

For two essential homotopy classes that intersect, horizontal curves in their supporting annuli cross, so the Non-Crossing Principle rules out simultaneous sufficiently wide canonical regions. For disjoint classes, the source's proof uses planar separation to show that an overlap of middle regions forces the outer buffer families to cross. Both buffer widths exceed one, contradicting the product bound.

The planar-separation claim depends on distinct nonhomotopic classes, proper complementary components and the choice of a point placed at infinity. It is not a Euclidean rectangle-overlap rule valid for arbitrary Riemann surfaces. The global existence component is supplied in Chapter 27. A complete topological proof of the overlap-to-buffer-crossing step and the appropriate choice/uniqueness of canonical maximizers are still required for the unrestricted source statement. The finite cylinder geometry and the width obstruction are established independently above.

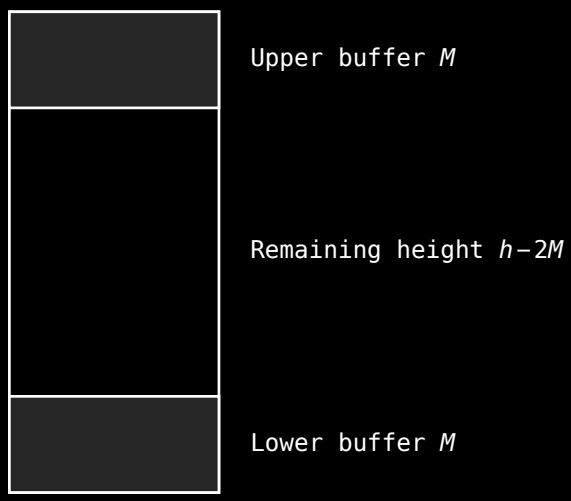


Figure 8.1: A cylinder of height h loses two buffers of height M. The remaining modulus is $h - 2M$. This local construction does not by itself prove that canonical regions from different classes are disjoint.

8.5 Covering annuli and strict inequalities

Source correspondence. Lyubich, Lemma 6.40 and its stated dependence on Lemma 7.15..

Let $\widetilde{A}_\gamma \to S$ denote the cyclic covering associated with the subgroup generated by an essential primitive class γ. Every embedded annulus in that class lifts conformally to this cyclic covering, so monotonicity gives

$$\operatorname{mod} A_\gamma \leq \operatorname{mod} \widetilde{A}_\gamma.$$

Strict inequality is not universal. If S itself is an annulus and γ is its core generator, the cyclic cover is the identity cover and the maximizing annulus is S. The two moduli are equal. Hence any strict version requires a non-annular hypothesis and a separate rigidity argument.

For a primitive simple closed geodesic of length ℓ in a torsion-free hyperbolic quotient, Chapter 26 proves the embedded collar and its metric directly. The resulting estimate is

$$\frac{\pi}{\ell} - 1 < M_\gamma \leq \frac{\pi}{\ell},$$

where M_γ is the supremum of embedded moduli. The existence argument in Chapter 27 identifies it with an attained modulus. The strict upper inequality printed in the source is not valid for the cyclic-annulus case. The additive lower constant one is sufficient in the stated geodesic setting; optimality of that additive constant is not asserted.

8.6 Canonical weighted arcs

Source correspondence. Lyubich, §6.8, Lemmas 6.41 and 6.44, Remark 6.42 and Corollaries 6.43–6.45..

For a non-elementary finite-type hyperbolic surface with ideal boundary circles, a proper arc lifts to the disk between two appropriate boundary gaps. The four endpoints define a marked quadrilateral of width a. The source defines the buffered weight

$$w(\alpha) = \max(a - 2, 0).$$

Changing a lift by a deck transformation applies a conformal automorphism of the disk and preserves the quadrilateral width. Thus the weight is independent of lift, once the gap construction and homotopy-class identification are established.

The elementary operation $a \mapsto \max(a-2,0)$ is 1-Lipschitz. If a is enclosed by $[a_-, a_+]$, its weight is enclosed by $[\max(a_- - 2, 0), \max(a_+ - 2, 0)]$. A certificate of positive weight requires $a_- > 2$; a floating-point value slightly above two without an error bound does not decide positivity. The distance of the interval from two is a conditioning margin for the combinatorial decision.

The geometric disjointness of buffered rectangles uses the equality form of the Non-Crossing Principle. Since that global equality proof has not been supplied here, the source's canonical-embedding and global disjointness conclusions remain conditional in the register. Their finite-dimensional threshold arithmetic is fully executable.

8.7 The Euler-characteristic count

Proposition 8.2 (Ideal arc bound)**.** [GO2] *Let S be an orientable finite-type surface of negative Euler characteristic whose boundary components are treated as ideal vertices. A finite family of pairwise disjoint, pairwise nonhomotopic essential arcs that extends to an ideal triangulation contains at most $-3\chi(S)$ arcs.*

Proof. In an ideal triangulation, every triangle has three side occurrences and every interior ideal edge is counted twice, so $3F = 2E$. Removing the ideal vertices from the Euler formula gives $\chi(S) = F - E$. Substitution yields $E = -3\chi(S)$. The original family is a subset of the triangulation, so it contains at most that many edges. □

For a disk with n holes, $\chi = 1 - n$, giving $3(n-1)$. The extendability and essential/nonparallel conditions matter. Repeating the same arc, allowing boundary-parallel arcs, or allowing intersections destroys the counting argument. A complete proof that the source's canonical family satisfies all conditions depends on its preceding geometric lemmas; the combinatorial count itself is independent of the modulus calculation.

8.8 What is completed by the source audit

The manuscript supplies explicit counterexamples to unqualified puncture uniqueness, torus uniqueness and universal strictness of the covering comparison. It proves the cylinder buffer calculation, lift invariance of width under its construction, threshold interval propagation and the ideal-triangulation count. The explicit rough-ring, collar, and annular-existence statements are proved in their supporting chapters. The global quadratic-differential uniqueness and the restricted equality-rigidity statements required for canonical rectangles remain open. These distinctions determine the status of every downstream geometric assertion.

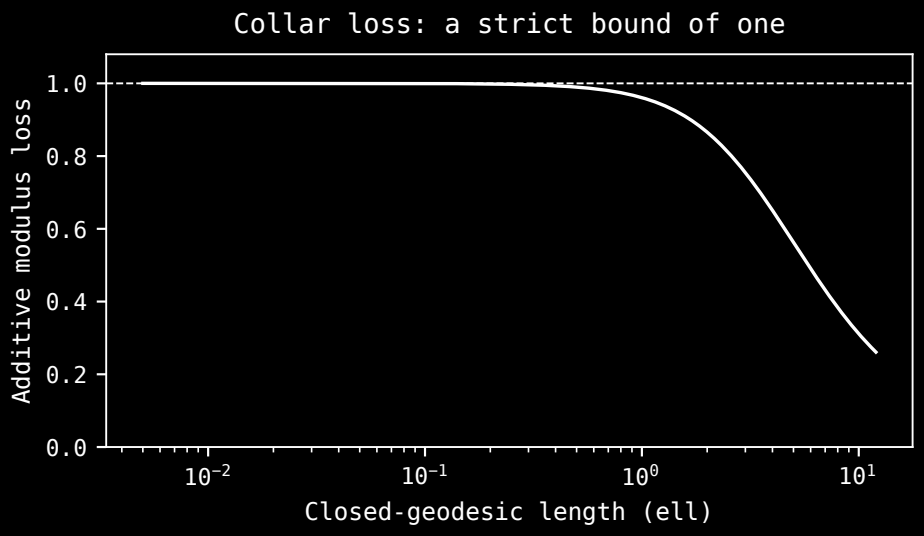


Figure 8.2: The exact collar-gap function from Corollary 26.4. The global embedding argument is supplied in its proof; the sampled graph illustrates the proved bound and is not an additional proof.

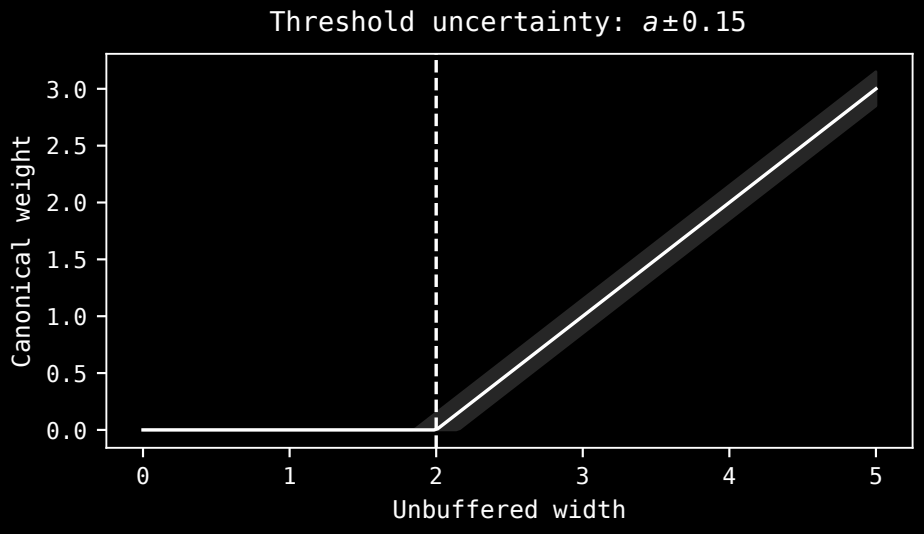


Figure 8.3: The positive-part rule propagates an input interval through the canonical-weight threshold. A straddling interval certifies a range but cannot decide strict positivity.

Chapter 9

Finite networks: complete primal and dual constructions

9.1 Incidence, orientation and the energy

Let $G = (V, E)$ be a finite connected undirected graph with n vertices and m edges. Choose an orientation of each edge once. For $e = (i, j)$ define $(Bu)_e = u_j - u_i$. Positive conductances form the diagonal matrix $C = \mathrm{diag}(c_e)$. The weighted Laplacian is $L = B^T C B$. Reversing an orientation negates the corresponding row of B and the corresponding flow entry; it leaves L and all energies unchanged.

Choose distinct terminals s, t, with $u_s = 0$ and $u_t = 1$. The primal energy is

$$E(u) = \sum_{e=(i,j)} c_e (u_j - u_i)^2 = u^T L u.$$

Its minimum C_{st} is effective conductance under unit terminal voltage. The symbol C for the diagonal matrix is distinguished by context from the scalar C_{st}. The reciprocal $R_{st} = 1/C_{st}$ is effective resistance.

Connectivity makes the reduced Dirichlet matrix positive definite: if a zero-boundary vector has zero energy, its values agree along every edge, hence on the connected graph, and its terminal zero makes it vanish. This gives existence and uniqueness of the finite minimizer by completing squares or solving the reduced linear system. A disconnected input needs a different model with infinite resistance or undetermined interior components; the checker in this release rejects it explicitly.

9.2 The discrete maximum principle

Proposition 9.1. [NO1] *At every nonterminal vertex, the minimizing potential satisfies*

$$u_i = \frac{\sum_{j \sim i} c_{ij} u_j}{\sum_{j \sim i} c_{ij}}, \qquad 0 \le u_i \le 1.$$

Proof. Differentiate the quadratic with respect to u_i. The stationarity equation is $\sum_{j \sim i} c_{ij}(u_i - u_j) = 0$, giving the weighted average. If a maximum exceeding one occurred at an interior vertex, every neighbor would have to attain the same maximum because all weights are positive. Connectivity propagates it to a terminal, a contradiction. The minimum argument is identical. Uniqueness follows from positive definiteness, so the stationary solution is the minimizer. □

An implementation returning a potential outside $[0, 1]$ on such a graph has failed at least one premise or calculation. Conversely, lying in $[0, 1]$ is not a sufficient certificate of harmonicity. The exact energy and divergence equations must still be checked.

9.3 Unit flows and summation by parts

A flow $j \in \mathbb{R}^m$ is given a sign according to the chosen edge orientation. We use outgoing divergence at each vertex: each j_e is added at its initial vertex and subtracted at its final vertex. A unit s–t flow has divergence $+1$ at s, -1 at t, and zero elsewhere. Since B is final-minus-initial, the vector relation is $B^T j = e_t - e_s$.

For every potential satisfying the terminal constraints,

$$\sum_e j_e (Bu)_e = u^T B^T j = u_t - u_s = 1.$$

This is a finite exact analogue of integration by parts. There are no hidden regularity or trace hypotheses: it is matrix multiplication. Define the flow energy $F(j) = \sum_e j_e^2 / c_e$.

Theorem 9.2 (Finite capacity bracket)**.** [N02] *For every feasible potential u and unit flow j,*

$$\boxed{\frac{1}{F(j)} \le C_{st} \le E(u)}.$$

Both bounds can be attained simultaneously, and the minimizing unit flow is unique.

Proof. The upper bound follows from the definition of the primal minimum. Weighted Cauchy–Schwarz applied to the summation-by-parts identity gives $1 \le E(v)F(j)$ for every feasible v. Taking v to be the minimizer gives the lower bound.

Let u_* be the minimizing potential and put $i = CBu_*$. Interior harmonicity gives zero divergence there. The terminal flux is C_{st}, since $u_*^T B^T i = E(u_*)$ and $u_s = 0, u_t = 1$. Hence $j_* = i / C_{st}$ is a unit flow. Its energy is $E(u_*)/C_{st}^2 = 1/C_{st}$. This proves attainment.

If $j = j_* + k$ is another unit flow, then $B^T k = 0$ and

$$\sum_e \frac{j_{*,e} k_e}{c_e} = \frac{1}{C_{st}} (Bu_*)^T k = 0.$$

Thus $F(j) = F(j_*) + F(k)$, which is strictly larger unless $k = 0$. This proves uniqueness and the flow error identity. □

The proof is self-contained finite linear algebra. It establishes a universal implication for exact input representations; it does not establish that floating-point equality tests are exact.

9.4 A completed-square proof of the same certificate

For every real scalar α,

$$\sum_e \left(\sqrt{c_e} (Bu)_e - \alpha \frac{j_e}{\sqrt{c_e}} \right)^2 = E(u) - 2\alpha + \alpha^2 F(j) \ge 0.$$

Minimizing the right-hand side over α gives $E(u) \ge 1/F(j)$. Equality holds exactly when $c_e (Bu)_e = j_e / F(j)$ on every edge. Together with conservation this identifies the harmonic potential and the optimum. This alternative route exposes an explicit nonnegative residual whose vanishing is an equality certificate. The two proofs share the same summation-by-parts identity, so they are not logically independent of its orientation convention.

9.5 The path-density formulation on a graph

Give every edge a nonnegative length ρ_e, with energy $\sum_e c_e \rho_e^2$. Impose $\sum_{e\in\gamma} \rho_e \geq 1$ on every s–t path, counting repeated traversals. For any feasible potential, $\rho_e = |u_j - u_i|$ is path admissible by telescoping and the triangle inequality.

Conversely, from a path-admissible ρ define

$$u(v) = \min\left(1, \inf_{\gamma: s\to v} \sum_{e\in\gamma} \rho_e\right).$$

Nonnegative lengths imply $|u(i) - u(j)| \leq \rho_e$, and path admissibility gives $u(t) = 1$. Thus the potential energy does not exceed the density energy. Minimizing both constructions proves exact equality of the graph path-density minimum and C_{st}.

Although the path family can be exponentially large, the equivalent potential problem has $n-2$ unknowns. Dijkstra's algorithm can test the constraints through a shortest-path calculation for nonnegative proposed lengths. The potential-flow certificate avoids enumerating paths entirely. A density constrained only on a sampled path subset generally underestimates the full width; the sampling error is a mathematical issue, not just a solver tolerance.

9.6 Rayleigh monotonicity and conductance sensitivity

Proposition 9.3. [N03] *Increasing any conductance while leaving all others fixed cannot decrease C_{st}. Deleting an edge cannot increase it. If $C_{st}(c)$ is evaluated in a strictly positive conductance region, then*

$$\frac{\partial C_{st}}{\partial c_e} = (u_{*,j} - u_{*,i})^2.$$

Proof. Every feasible potential energy increases when a conductance increases, so its infimum increases. Edge deletion sets one nonnegative contribution to zero, giving the opposite comparison. In the positive connected region the reduced inverse depends smoothly on conductances. Differentiate $E_c(u_*(c))$. The derivative in u_* vanishes against its zero-boundary variation by stationarity, leaving only the direct derivative with respect to c_e, namely the squared edge drop. □

The scalar capacity is the infimum of linear functions of c, so it is concave in the conductance vector. To verify this directly, for $0 \leq t \leq 1$,

$$C_{st}(tc + (1-t)d) = \inf_u [tE_c(u) + (1-t)E_d(u)] \geq tC_{st}(c) + (1-t)C_{st}(d).$$

Sensitivity is nonnegative, but second derivatives include redistribution of the minimizing potential. Ignoring that redistribution is permissible for the first derivative because of stationarity; it is not permissible for every higher derivative.

9.7 Schur complements and boundary response

Partition vertices into retained B and interior I. If every interior component is attached to B, then L_{II} is invertible. The interior minimizer for prescribed u_B is

$$u_I = -L_{II}^{-1} L_{IB} u_B.$$

Substitution gives the boundary-response matrix

$$\Lambda = L_{BB} - L_{BI} L_{II}^{-1} L_{IB}.$$

This is an exact elimination identity. Its quadratic form is nonnegative, it annihilates the constant boundary vector, and it gives the minimum energy for every retained boundary pattern. A reduced circuit is therefore equivalent at its retained terminals even though it may have different internal currents.

A single factorization of L_{II} can serve multiple boundary patterns. Setup and repeated queries must be separated in any performance claim. If the retained boundary changes, or a conductance changes enough to require refactorization, the previous setup cost cannot silently be reused. Approximate Schur complements require a spectral error contract; matching a few observed boundary responses does not prove uniform equivalence.

9.8 Exact rational witnesses and bit complexity

For rational conductances the reduced linear system has rational entries and its nonsingular solution is rational. Clearing denominators produces an integer system. Cramer's rule proves existence of rational coordinates, while Hadamard's determinant bound gives a finite bound on their integer lengths in terms of matrix size and entry lengths. Fraction-free elimination or modular reconstruction can avoid some intermediate growth, but those algorithms need their own cost proofs.

Once a potential and unit flow are supplied, checking boundary values, conservation and energies takes $O(n + m)$ rational additions and multiplications. If all intermediate numerators and denominators have at most B bits, the arithmetic operation count gives a bound of order $O((n + m)\,\mathsf{M}(B))$ plus normalization/gcd costs in the chosen rational representation. This is not an unconditional linear bit bound; B must be measured or bounded.

The companion implementation uses Python `Fraction` to generate witnesses and a separately written C++ multiprecision-rational program to check them. The programs do not share a numerical factorization routine. Their common mathematical specification and unverified compilers/runtimes remain explicit trust dependencies. No proof-assistant checker is claimed.

9.9 Failure modes addressed by the checker

The delivered schema requires $n \geq 2$, distinct valid terminals, connectedness, positive conductances, a fixed edge list, matching vector dimensions, exact terminal values, and exact divergence. Duplicate undirected edges are rejected by this schema; parallel conductances must be combined explicitly beforehand or represented by a revised schema. This is an implementation restriction, not a theorem that electrical parallel edges are invalid.

The lower bound is recomputed from flow energy and the upper bound from potential energy. Changing the printed bound without changing the witness is rejected. The checker also rejects a reversed flow without its corresponding orientation update, a missing edge, an invalid rational denominator, and a self-loop. A feasible nonoptimal witness is accepted with a positive gap. Acceptance proves the bracket, not necessarily exact optimality. Exact equality of the endpoints is a stronger property recorded separately.

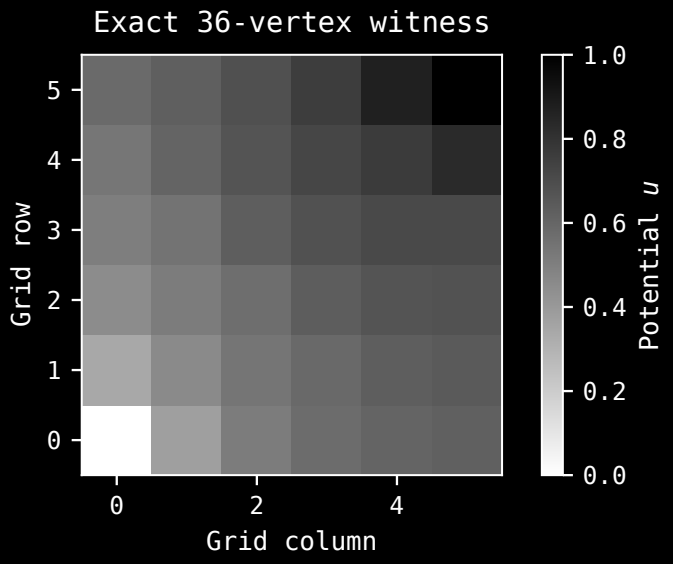


Figure 9.1: A rational potential on the 36-vertex test network. The grayscale display rounds the exact values for visualization; the certificate files retain exact numerators and denominators.

Chapter 10

Quantitative variational stability and attained constants

10.1 A common affine constraint

Let H be a real Hilbert space with inner product a, and let $K = g + V$ be a closed affine subspace. Assume its minimum-norm point u is nonzero. Let b be another symmetric form with

$$|b(x,y) - a(x,y)| \le \varepsilon \|x\|_a \|y\|_a, \qquad 0 < \varepsilon < 1. \tag{10.1}$$

Equivalently, $b(x,y) = a((I+E)x, y)$ for a bounded self-adjoint operator E of norm at most ε, by the representation proved in Proposition 2.5. In finite dimensions this is a symmetric matrix representation after an a-orthonormal basis change. Equivalent-norm completeness and projection, also proved in the supporting chapter, give the required existence statement.

The form b is coercive with lower bound $1-\varepsilon$, so it has a unique minimizer $v \in K$. Both minimizers satisfy Euler conditions in the same variation space. Changing the feasible affine set would add a separate boundary/constraint perturbation and invalidate the following estimate without correction.

10.2 An attained perturbation bound

Theorem 10.1 (Affine minimizer perturbation)**.** [Q01] *Under* (10.1)*,*

$$\boxed{\frac{\|v-u\|_a}{\|u\|_a} \le \frac{\varepsilon}{\sqrt{1-\varepsilon^2}}.}$$

The constant is attained by a two-dimensional real example for every $0 < \varepsilon < 1$*.*

Proof. Set $e = v - u \in V$. If $e = 0$ there is nothing to prove. Otherwise put $p = u/\|u\|_a$, $q = e/\|e\|_a$ and $t = \|e\|_a/\|u\|_a$. Since u is the a-minimizer, $a(p,q) = 0$. The b-stationarity equation against q gives

$$0 = b(p+tq, q) = a(Ep, q) + t(1 + a(Eq, q)).$$

Write $d = a(Eq,q)$ and $c = a(Eq,p) = a(Ep,q)$. Orthogonal components of Eq yield $c^2 + d^2 \le \|Eq\|_a^2 \le \varepsilon^2$. Therefore

$$t^2 = \frac{c^2}{(1+d)^2} \le \frac{\varepsilon^2 - d^2}{(1+d)^2} \le \frac{\varepsilon^2}{1-\varepsilon^2}.$$

The last inequality is equivalent, after positive denominators are cleared, to $(d+\varepsilon^2)^2 \geq 0$.

For attainment take $H = \mathbb{R}^2$ with the Euclidean form, $K = \{(1,t) : t \in \mathbb{R}\}$, and

$$E = \begin{pmatrix} \varepsilon^2 & -\varepsilon\sqrt{1-\varepsilon^2} \\ -\varepsilon\sqrt{1-\varepsilon^2} & -\varepsilon^2 \end{pmatrix}.$$

A multiplication gives $E^2 = \varepsilon^2 I$, so its norm is ε. The unperturbed minimizer is $(1,0)$, while the perturbed stationary equation is $-\varepsilon\sqrt{1-\varepsilon^2} + (1-\varepsilon^2)t = 0$. Thus $t = \varepsilon/\sqrt{1-\varepsilon^2}$, attaining the bound. □

This proof controls the difference between minimizers, not the condition number of an arbitrary matrix solve. It applies to any closed affine constraint in the stated Hilbert geometry. The sharpness claim is over the complete class of symmetric relative form perturbations with that norm bound. Priority relative to existing angle and perturbation inequalities has not been established.

10.3 A weaker resolvent estimate and the source of improvement

Let P be the a-orthogonal projection onto V. From $b(u+e, w) = 0$ for $w \in V$ one obtains

$$(I + PE|_V)e = -PEu.$$

For $T = PE|_V$ with $\|T\| \leq \varepsilon < 1$, the series $\sum_{k\geq 0}(-T)^k$ converges in operator norm by the geometric bound. Multiplying a finite partial sum by $I+T$ gives $I - (-T)^{N+1}$; passage to the limit proves that the series is its inverse, of norm at most $1/(1-\varepsilon)$. This yields

$$\|e\|_a \leq \frac{\varepsilon}{1-\varepsilon}\|u\|_a.$$

This route treats the forcing PEu and the perturbation on V independently. They are generated by the same self-adjoint operator and cannot simultaneously take their worst unconstrained values. The two-dimensional compression in Theorem 10.1 retains the constraint $c^2 + d^2 \leq \varepsilon^2$ and gives the smaller attained constant. The distinction is mathematical, not an empirical speedup.

At $\varepsilon = 3/5$, the sharp factor is $3/4$, whereas the resolvent factor is $3/2$. At $\varepsilon = 5/13$, the sharp factor is $5/12$, whereas the resolvent factor is $5/8$. Both can be represented exactly with rational extremizing matrices by parametrizing the unit circle. The certificate suite checks several such rational examples.

10.4 Energy values under form perturbation

Proposition 10.2. [Q02] *Let $C_a = \min_K a(x,x)$ and $C_b = \min_K b(x,x)$. Then*

$$(1-\varepsilon)C_a \leq C_b \leq (1+\varepsilon)C_a.$$

Proof. The pointwise bounds $(1-\varepsilon)a(x,x) \leq b(x,x) \leq (1+\varepsilon)a(x,x)$ follow from the norm bound. Infima give the lower inequality. Evaluating b at the a-minimizer gives the upper inequality. Both constants are attained by $b = (1 \pm \varepsilon)a$. □

The resistance-like reciprocal then lies between $C_a^{-1}/(1+\varepsilon)$ and $C_a^{-1}/(1-\varepsilon)$. Absolute error and relative error therefore have different behavior near a zero capacity. If a family degenerates and $C_a \to 0$, a small relative form perturbation can still leave the reciprocal very large. No uniform absolute resistance bound follows without a positive capacity floor.

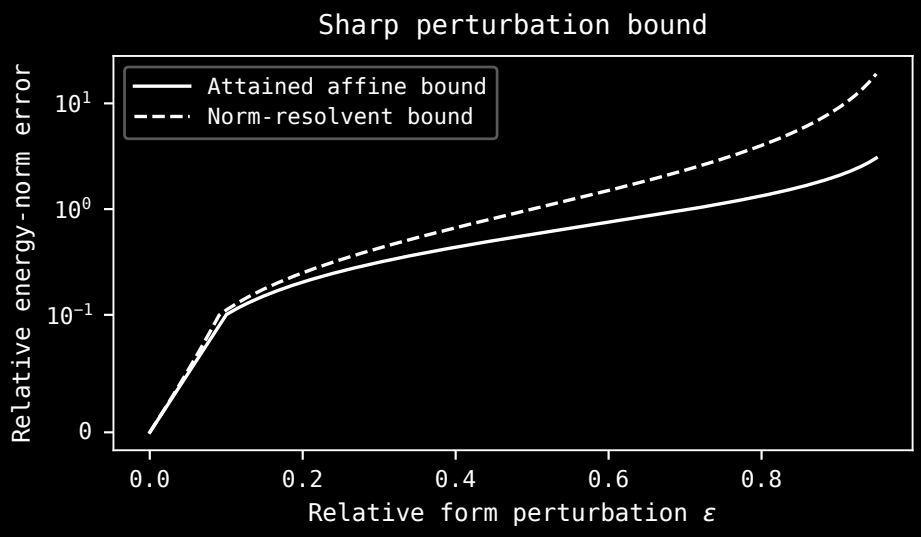


Figure 10.1: The attained affine-minimizer bound and the norm-resolvent bound. The vertical scale is symmetric-logarithmic to show the small-error and degenerating regimes together.

10.5 Boundary perturbation as a separate mechanism

Suppose the affine constraint changes from $g + V$ to $(g + \eta) + V$. With the form fixed, the minimizer is $u = P_{V^\perp} g$ and the new one is $u' = P_{V^\perp}(g + \eta)$. Hence

$$\|u' - u\|_a \leq \|\eta\|_a.$$

Equality is possible when $\eta \in V^\perp$. This is a projection contraction, unrelated to the relative form denominator in Theorem 10.1. For combined perturbations one can insert the intermediate minimizer with the new boundary and old form, then apply the triangle inequality and the two estimates with their respective normalizations.

In a conductivity problem, boundary measurement error and coefficient error therefore require separate error-budget entries. Treating both as a single residual can hide a physically incorrect boundary condition despite an accurate linear solve.

10.6 A-posteriori error from an approximate Euler condition

Let $w \in K$ be feasible and suppose its residual functional $r(v) = a(w, v)$ on V has dual norm at most δ. Since $e = w - u \in V$,

$$\|e\|_a^2 = a(w, e) = r(e) \leq \delta \|e\|_a.$$

Thus $\|e\|_a \leq \delta$. In a different reference norm, coercivity supplies the conversion factor. This is a precise residual-to-error implication. A raw coordinate residual norm without the dual norm or coercivity constant does not have the same interpretation.

For a finite SPD system $Ax = b$, an approximation $\widehat{x}$ with residual $r = b - A\widehat{x}$ satisfies $e = A^{-1}r$ and $\|e\|_A^2 = r^T A^{-1} r \leq \|r\|_2^2 / \lambda_{\min}(A)$. A certified lower eigenvalue bound turns a residual into a norm certificate. A floating-point estimate of the smallest eigenvalue is not automatically such a bound.

10.7 Primal–dual gap and two error components

For a graph feasible pair (u, j), let $C = C_{st}$, $e = E(u) - C \geq 0$, and $f = F(j) - 1/C \geq 0$. Then

$$E(u) - \frac{1}{F(j)} = e + \frac{C^2 f}{1 + Cf}.$$

The identity follows by substituting $F(j) = 1/C + f$. A vanishing gap forces both $e = f = 0$. A small gap controls the primal energy error directly, while the flux error additionally depends on C and the gap regime. This decomposes the role of the two witnesses quantitatively.

The condition number of the generator and the gap of the certificate are different data. A badly conditioned generator can still produce a rigorously checkable bracket; a small numerical residual without feasibility can fail to produce any valid bracket. The proposed separation of generation and verification is valuable precisely in this regime.

10.8 Where the generalization remains open

The arguments in this chapter use symmetric coercive forms and a common affine variation space. Nonsymmetric operators, indefinite energies, nonlinear constraints and nonconvex physical models require different hypotheses. A finite-dimensional counterexample is obtained from $b(x, y) = a((I + E)x, y)$ with a large skew part: the quadratic energy does not even determine the same Euler equation as the full bilinear form. Nonlinear p-energies retain convexity for $p > 1$ but lose the exact quadratic Pythagorean identity. Their estimates must be derived from uniform convexity or monotonicity, not by retaining the same constant without proof.

Chapter 11

Weighted and nonlinear modulus

11.1 Finite p-modulus

For a finite path-incidence system $N\rho \geq \mathbf{1}$ and positive weights a_e, define

$$M_p(\Gamma) = \min_{\rho \geq 0,\ N\rho \geq \mathbf{1}} \sum_e a_e \rho_e^p, \qquad 1 < p < \infty.$$

Coercivity and closedness give existence; strict convexity gives uniqueness. At $p = 2$ this is the width convention used previously. At other exponents the resistance-like normalization is $M_p^{-1/(p-1)}$ in series computations. Calling that quantity simply the reciprocal modulus would conceal the exponent.

11.2 Series and parallel nonlinear resistances

Consider a chain of edges and require $\sum_i \rho_i = 1$, the active path constraint. The Lagrangian derivative is $p a_i \rho_i^{p-1} = \lambda$, so

$$\rho_i = \frac{a_i^{-1/(p-1)}}{\sum_j a_j^{-1/(p-1)}}.$$

Substitution gives

$$M_p = \left(\sum_i a_i^{-1/(p-1)} \right)^{1-p}. \tag{11.1}$$

The same lower bound follows independently from Hölder:

$$1 = \sum_i (a_i^{1/p} \rho_i) a_i^{-1/p} \leq M_p^{1/p} \left(\sum_i a_i^{-1/(p-1)} \right)^{1/q}.$$

This second proof supplies global optimality and equality, rather than only a stationary point. For disjoint parallel branches, their independent path constraints make the minimum energies add. Thus the transformed resistance has additive series behavior and the capacity-like energy has additive parallel behavior.

11.3 A dual finite certificate

Let $\lambda \geq 0$ be multipliers for $N\rho \geq \mathbf{1}$. With $w = N^T \lambda$, minimizing the Lagrangian over $\rho_e \geq 0$ gives

$$a_e x^p - w_e x \text{ has minimum } -(p-1) p^{-q} a_e^{-1/(p-1)} w_e^q.$$

Consequently every nonnegative λ gives the lower bound

$$\sum_{\gamma} \lambda_\gamma - (p-1)p^{-q} \sum_{e} a_e^{-1/(p-1)} (N^T\lambda)_e^q \le M_p.$$

Every feasible density gives an upper bound. This weak-duality certificate follows from a one-variable minimization and finite summation. Equality requires complementary slackness and the coordinate stationarity relations. Strong duality for the nonempty finite system follows from convex separation under a strictly feasible scaling whenever every path has nonzero incidence; that separation theorem is an additional dependency, not needed for the weak certificate itself.

If p is rational, the powers in the certificate may be irrational. Exact rational arithmetic alone then cannot evaluate the dual bound. One must either choose an algebraically tractable fragment or enclose the required powers with proved algebraic inequalities. The quadratic graph checker is not silently described as a general p-modulus checker.

11.4 Radial p-capacity in dimension n

Consider a spherical annulus $r < |x| < R$ in $\mathbb{R}^n$, with radial potentials $u(r) = 0$, $u(R) = 1$. Its radial energy is

$$\omega_{n-1} \int_r^R |u'(s)|^p s^{n-1} \, \mathrm{d}s.$$

Set $I = \int_r^R s^{-(n-1)/(p-1)} \, \mathrm{d}s$. Weighted Hölder gives $1 \le (\mathcal{E}_p/\omega_{n-1})^{1/p} I^{1/q}$, so the minimum is $\omega_{n-1} I^{1-p}$, attained by

$$u'(s) = I^{-1} s^{-(n-1)/(p-1)}.$$

Rotational averaging and convexity show that radialization of any feasible smooth potential cannot increase its energy. Thus this is also the full smooth capacity minimum. For $p = n$, $I = \log(R/r)$; for $p \ne n$,

$$I = \frac{R^{(p-n)/(p-1)} - r^{(p-n)/(p-1)}}{(p-n)/(p-1)}.$$

In the plane at $p = 2$, the formula reduces to $2\pi/\log(R/r)$.

The limit $p \to n$ follows from $(x^\alpha - 1)/\alpha \to \log x$. The limiting logarithm is a geometrical consequence of the critical dimension, not a numerical singularity requiring a different physical model.

11.5 Why conformal invariance selects an exponent

Under a conformal coordinate change in dimension n with local scale λ, a gradient gains a factor λ and the volume change contributes λ^{-n} in pullback comparison. The p-energy therefore gains λ^{p-n}. Exact invariance is natural at $p = n$. In dimension two the Dirichlet exponent $p = 2$ is critical.

For a general linear change $y = Fx$, the transformed tensor for an energy $\int \nabla u^T A \nabla u \, \mathrm{d}x$ is

$$A_y(y) = \frac{F A_x(x) F^T}{|\det F|}, \qquad x = F^{-1}y.$$

This follows by $\nabla_x u = F^T \nabla_y v$ and the Jacobian. A scalar isotropic tensor is generally transformed into an anisotropic tensor. Only conformal linear maps in two dimensions preserve scalar isotropy with the same coefficient. Coordinate simplification can therefore make the material law more complicated.

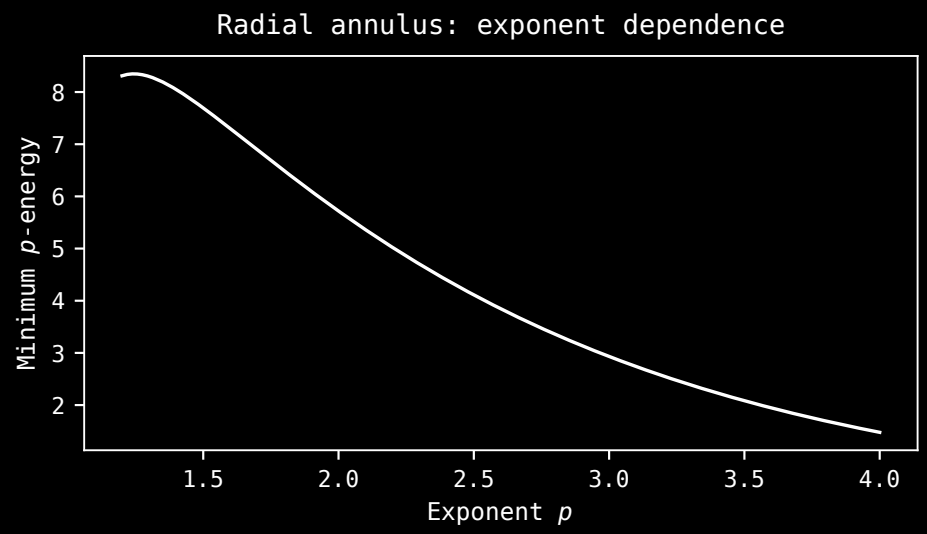


Figure 11.1: The radial p-capacity of the fixed planar annulus $A(1,3)$. The quadratic case is one point in an exponent-dependent family, not a universal reciprocal rule.

11.6 Anisotropic rectangle: a complete model

For $A = \mathrm{diag}(a_x, a_y)$ on $[0,\ell] \times [0,h]$ with fixed bottom/top potentials and insulating sides, $u = y/h$ solves the Euler equation. Its energy is $a_y \ell/h$. A direct lower bound integrates $\int_0^h a_y u_y^2 \geq a_y/h$ on each vertical fiber; the horizontal derivative term is nonnegative. Thus the formula is globally optimal within the feasible class. The coefficient a_x influences other boundary patterns but not this one. A single terminal experiment cannot identify both anisotropic coefficients.

This supplies an inverse-identifiability counterexample: all tensors with the same a_y and arbitrary $a_x > 0$ have the same observed two-terminal capacity for this boundary condition. Adding more precision to that one measurement cannot recover a_x.

11.7 Metric-measure and singular categories

In a metric-measure space, line length and area no longer arise from one smooth coordinate metric. A p-modulus requires a measure, rectifiable curves and upper-gradient conventions. Atoms, disconnected supports, and families of zero-modulus curves change admissibility. A dimension-specific conformal formula cannot be retained solely because the functional is still denoted M_p.

On a nodal surface one must decide whether functions have matching traces at the identified points and whether those points carry positive capacity in the selected exponent. In the planar quadratic theory a single point can have zero capacity, so a point identification need not transmit the same amount of energy as a narrow but finite neck. The circular neck model in the next chapter makes that limiting distinction explicit. A complete general theory on all approved singular and metric-measure categories remains outside the completed part of this review; the finite and radial statements above have their own proved domains.

Chapter 12

Periods, degeneration and Fourier boundary operators

12.1 Marked torus deformations

Write $\tau = x + iy$ with $y > 0$. The extremal length of a class (m, n) is

$$L_{m,n}(x, y) = \frac{(m + nx)^2}{y} + n^2 y.$$

Its derivatives are

$$\partial_x L = \frac{2n(m + nx)}{y}, \qquad \partial_y L = n^2 - \frac{(m + nx)^2}{y^2}.$$

The Hessian equals

$$\frac{2}{y^3} \begin{pmatrix} n^2 y^2 & -ny(m + nx) \\ -ny(m + nx) & (m + nx)^2 \end{pmatrix},$$

a positive semidefinite rank-one matrix unless the class is zero. This proves convexity in the displayed coordinates. It is not a proof of geodesic convexity for an arbitrary metric on moduli space. A coordinate claim and a geometric claim have different meanings.

Changing the marking changes the pair (m, n). Under $\tau' = \tau + 1$, take $(m', n') = (m - n, n)$; under $\tau' = -1/\tau$, take $(m', n') = (n, -m)$. Direct substitution proves $L_{m',n'}(\tau') = L_{m,n}(\tau)$. The first equality is algebraic cancellation. For the second, $m' + n'\tau' = (m + n\tau)/\tau$ and $\operatorname{Im} \tau' = y/|\tau|^2$. Thus the scalar invariant is preserved only after the homology coordinate is transformed with the geometry.

12.2 Periods as finite-dimensional data

On a flat torus, a constant real one-form $\alpha = a\,\mathrm{d}x + b\,\mathrm{d}y$ has periods determined by a two-dimensional vector. Requiring a fixed period along one primitive cycle imposes an affine linear constraint, and minimizing its L^2 norm is a quadratic projection problem. This provides a functional-analytic derivation of the torus formula and connects the curve representative to a dual cohomology class.

For a general compact Riemann surface, period matrices involve harmonic or holomorphic differentials, symplectic homology bases and positivity identities. Those statements require Hodge or Riemann–Roch theory. The torus calculation does not establish their general existence. The approved specialist algebraic-geometric programme remains in the open register beyond the explicit genus-one and local differential calculations given here.

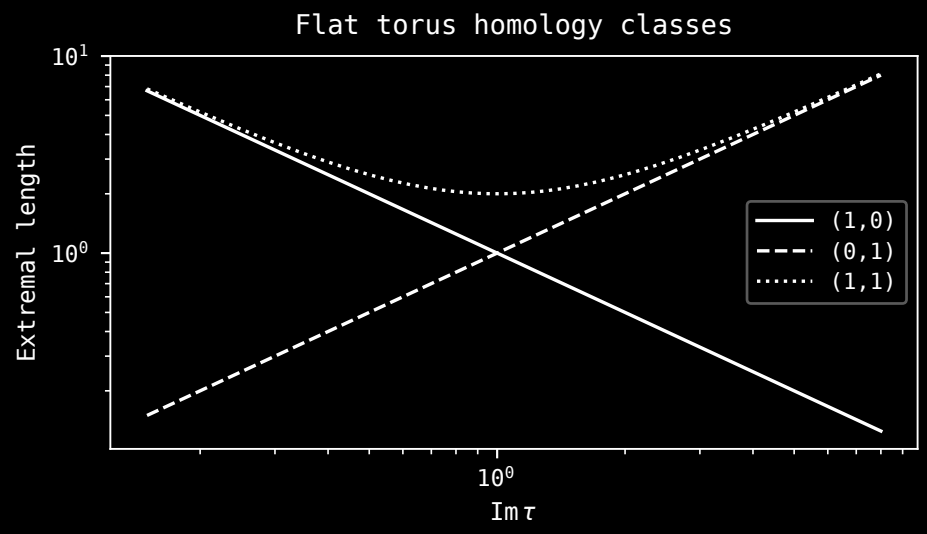


Figure 12.1: The same unmarked torus geometry supports different homology classes with different extremal lengths. The marking is part of the input.

12.3 A circular plumbing neck

A local nodal smoothing can be modelled by $zw = t$ with $0 < |t| < r^2$, restricting $|z| < r$ and $|w| < r$. In the z coordinate the overlap is

$$|t|/r < |z| < r,$$

so its modulus is exactly

$$m(t) = \frac{1}{2\pi} \log \frac{r^2}{|t|}.$$

The conductance of the neck is $1/m(t)$. Its logarithmic derivative is $\mathrm{d}m/\,\mathrm{d}\log|t| = -1/(2\pi)$, while its absolute derivative with respect to $|t|$ diverges as $|t| \to 0$. Conditioning therefore depends on the parameter representation.

At $t = 0$ the chart has two branches meeting at a node and the smooth annular neck no longer exists. A formula at positive t cannot be evaluated at zero by treating infinity as an ordinary real number. Whether the limiting energy couples the branches depends on the chosen trace and capacity model. This exact local calculation motivates, but does not complete, global degeneration theory.

12.4 Harmonic extension on a cylinder

Let u be periodic in x with period ℓ on $0 < y < h$. Use Fourier coefficients a_n at $y = 0$ and b_n at $y = h$, with wavenumber $k_n = 2\pi|n|/\ell$. For $n \neq 0$, separation of variables gives

$$u_n(y) = a_n \frac{\sinh(k_n(h-y))}{\sinh(k_n h)} + b_n \frac{\sinh(k_n y)}{\sinh(k_n h)}.$$

For $n = 0$, $u_0(y) = a_0 + (b_0 - a_0)y/h$. Differentiating twice verifies the Laplace equation; evaluating endpoints verifies the boundary values. For finite Fourier data this is a complete classical solution. Infinite sums require convergence assumptions, such as finite boundary energy in the norm below.

Multiplying $u_n'' = k_n^2 u_n$ by $\overline{u_n}$ and integrating gives

$$\int_0^h (|u_n'|^2 + k_n^2 |u_n|^2)\,\mathrm{d}y = [\overline{u_n} u_n']_0^h.$$

Substitution yields the exact mode energy

$$k_n \left[(|a_n|^2 + |b_n|^2) \coth(k_n h) - 2\,\mathrm{Re}(\overline{a}_n b_n) \operatorname{csch}(k_n h) \right].$$

Parseval multiplies the sum over integer modes by ℓ, and the zero mode contributes $\ell |b_0 - a_0|^2 / h$.

12.5 Dirichlet-to-Neumann matrices

For each nonzero mode the outward-normal derivatives at the two boundaries are obtained from

$$\Lambda_n = k_n \begin{pmatrix} \coth(k_n h) & -\operatorname{csch}(k_n h) \\ -\operatorname{csch}(k_n h) & \coth(k_n h) \end{pmatrix}.$$

Its eigenvectors $(1, 1)$ and $(1, -1)$ have eigenvalues $k_n \tanh(k_n h/2)$ and $k_n \coth(k_n h/2)$, both positive. These identities follow from the half-angle formulas for hyperbolic sine and cosine. The zero-mode matrix is $h^{-1}\left(\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right)$, whose constant-vector kernel represents an undetermined additive potential.

This is a complete boundary compression of a continuum PDE for the cylinder geometry. Computing the boundary response for finite Fourier input requires one independent two-by-two block per mode. FFT costs enter only when converting sampled physical-space boundary data to Fourier coefficients. Sampling and aliasing error are separate from the diagonal PDE solve.

12.6 Exponential information loss and an inverse lower bound

On a half-cylinder the decaying harmonic extension of a boundary Fourier mode has factor $e^{-k_n y}$. Suppose one observes the field at height h with deterministic error at most δ. Two boundary inputs $\pm A e^{2\pi i n x/\ell}$ produce interior data differing by $2Ae^{-k_n h}$. If $Ae^{-k_n h} \le \delta$, the same noisy observation zero is consistent with both inputs. Any reconstruction must then have error at least A for one of them.

This is an information-based obstruction independent of the reconstruction algorithm. It proves that inverse extension is ill-conditioned without a prior bound on high-frequency amplitude. It also supplies a resolution estimate $|n| \lesssim (\ell/(2\pi h)) \log(A/\delta)$ under a fixed amplitude model. Increasing arithmetic precision cannot recover information already below measurement uncertainty.

12.7 Optimal finite-rank approximation of a diagonal transfer

Consider a Hilbert space with orthonormal basis $(e_n)_{n \ge 0}$ and a compact diagonal operator $Te_n = s_n e_n$, where $s_0 \ge s_1 \ge \cdots \ge 0$. Truncating after the first N basis vectors gives a rank-N operator with norm error s_N.

For any rank-N linear operator A, the restriction of A to the $(N+1)$-dimensional span of $e_0, \ldots, e_N$ has nonzero kernel. Choose a unit vector v in that kernel. Then $\|(T - A)v\| = \|Tv\| \ge s_N$. Thus the best rank-$N$ norm error is exactly s_N. This proves both sides of the optimum over all bounded linear rank-N competitors, not only diagonal ones.

Applying this to one-sided cylinder Fourier modes with $s_n = e^{-2\pi n h/\ell}$ gives an explicit minimum rank for operator-norm tolerance ε. The result concerns linear approximation on the unit ball of the specified Hilbert space. Nonlinear approximation or an input class with different regularity needs a different lower bound.

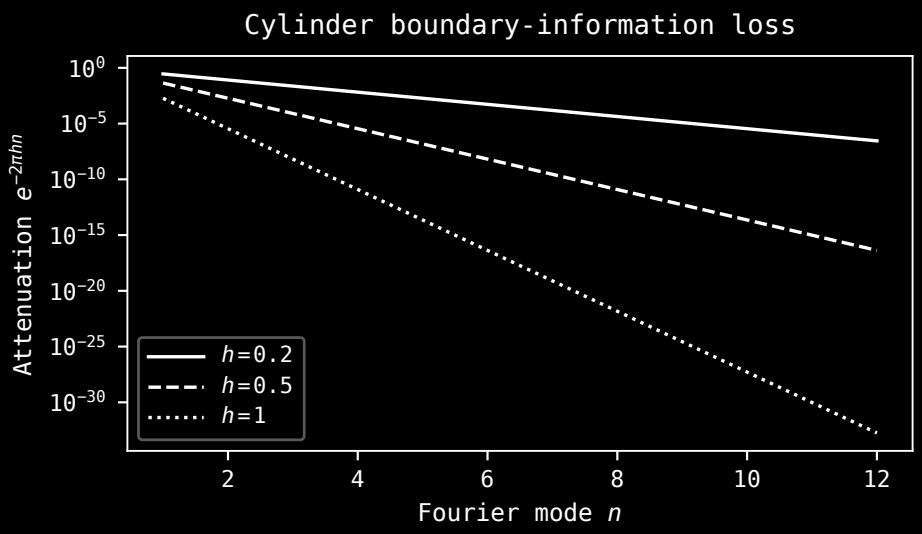


Figure 12.2: High Fourier modes are exponentially attenuated across a cylindrical distance. Inverse recovery requires a prior class and a measurement-error model.

12.8 Elliptic-integral computation and certification boundary

Marked quadrilateral uniformization often leads to complete elliptic integrals. One useful identity is $K(k) = \pi/[2\,\mathrm{AGM}(1, \sqrt{1-k^2})]$ for $0 \le k < 1$. Chapter 29 proves the integral transformation and supplies rational traces checked by independent implementations. The arithmetic–geometric iteration $a_{n+1} = (a_n + b_n)/2$, $b_{n+1} = \sqrt{a_n b_n}$ nevertheless has an elementary invariant bracket: $b_n \le b_{n+1} \le a_{n+1} \le a_n$ when $a_n \ge b_n > 0$.

The difference satisfies

$$a_{n+1} - b_{n+1} = \frac{(\sqrt{a_n} - \sqrt{b_n})^2}{2},$$

so the two sequences converge to a common positive limit. Rational lower and upper bounds on the square root preserve enclosing intervals. An implementation of the bracket is not a proof of the quoted identity with K; the analytical identity and numerical enclosure have distinct obligations.

Chapter 13

Physical models with explicit conformal and energy reductions

13.1 Coaxial electrostatics and thermal conduction

Consider a translationally invariant annular cross-section with radii $r < R$, homogeneous permittivity ϵ and prescribed conductor voltages $0, V$. The solution is $\phi(s) = V \log(s/r)/\log(R/r)$. Integrating the outward displacement flux on the inner conductor gives capacitance per axial length

$$C' = \frac{2\pi\epsilon}{\log(R/r)}.$$

Integrating $\epsilon|\nabla\phi|^2/2$ over the cross-section gives stored energy per length $C'V^2/2$. These two independent calculations agree on the normalizations. Fringing fields at finite axial ends are excluded by the translational-invariance assumption.

For steady conduction with thermal conductivity k and temperatures T_0, T_1, the same harmonic profile gives heat rate per length $2\pi k(T_1 - T_0)/\log(R/r)$. The physical dimensions differ from electrical capacitance even though the scalar boundary problem is identical. For temperature-dependent $k(T)$, the Kirchhoff transform $F(T) = \int^T k(s)\,\mathrm{d}s$ is harmonic, and one must invert F to recover T. Reusing the linear temperature profile is then generally wrong.

13.2 Diffusion time is not conformally invariant

Let a scalar concentration satisfy $\partial_t u = D\Delta_x u$. Under a planar conformal map $w = f(z)$, the Laplacian of the pullback gains the factor $|f'|^2$. Hence the transformed diffusion has a spatially varying coefficient unless $|f'|$ is constant. A stationary harmonic boundary solution is conformally transported, but the transient at the same physical time is not. This provides a concrete failure case for extending a stationary capacity reduction to a whole time-dependent experiment.

On a flat cylinder with zero Dirichlet conditions at both ends, the mode $\sin(\pi y/h)$ independent of the periodic direction decays as $e^{-D\pi^2 t/h^2}$. The stationary conductance scales as ℓ/h, while this decay time scales as h^2/D. Two geometries with the same ratio h/ℓ can therefore have the same conformal modulus and different physical diffusion times. Absolute scale matters to the time evolution.

13.3 Anti-plane elasticity

For a homogeneous isotropic solid invariant along one axis, assume an anti-plane displacement $w(x, y)$ with shear modulus μ. Its energy per axial length is $\mu \int |\nabla w|^2/2$. With prescribed displacements on two boundary sets and traction-free remaining boundaries, the Euler equation is Laplace's equation. The condenser potential therefore gives the exact scalar anti-plane model and its stiffness is proportional to capacity.

Plane strain and plane stress are vector systems with additional elastic constants. The scalar conformal reduction does not automatically solve them. Near a reentrant wedge of opening ω, a homogeneous Dirichlet harmonic mode behaves like $r^{\pi/\omega} \sin(\pi\theta/\omega)$. Its gradient behaves like $r^{\pi/\omega-1}$ and can diverge pointwise when $\omega > \pi$, while its local energy remains finite because $\int_0^\varepsilon r^{2\pi/\omega-1}\,\mathrm{d}r$ converges. A divergent mesh maximum is therefore compatible with convergent energy.

13.4 Threshold sheets in a scattering coordinate

Let s_0 be a threshold and choose $s_* < s_0$. On the cut plane $\mathbb{C} \setminus [s_0, \infty)$ choose the square root with positive real part and set $a = \sqrt{s_0 - s_*}$. Then

$$w(s) = \frac{a - \sqrt{s_0 - s}}{a + \sqrt{s_0 - s}}, \qquad s = s_0 - a^2 \left(\frac{1 - w}{1 + w}\right)^2 .$$

For $q = \sqrt{s_0 - s}$ with $\operatorname{Re} q > 0$, $|a - q| < |a + q|$, so $|w| < 1$. The inverse formula proves the coordinate relation. A power series in w may therefore organize analytic data relative to a cut more effectively than a Taylor series centered directly in s.

The map does not remove other thresholds, poles or crossed-channel singularities. A bound on a power-series tail requires a holomorphic class and a norm bound on a larger disk. Fitting coefficients from noisy data does not establish those bounds. Multichannel scattering needs one explicit sheet and threshold specification per channel. No nuclear-structure or measured scattering dataset is fitted in this review.

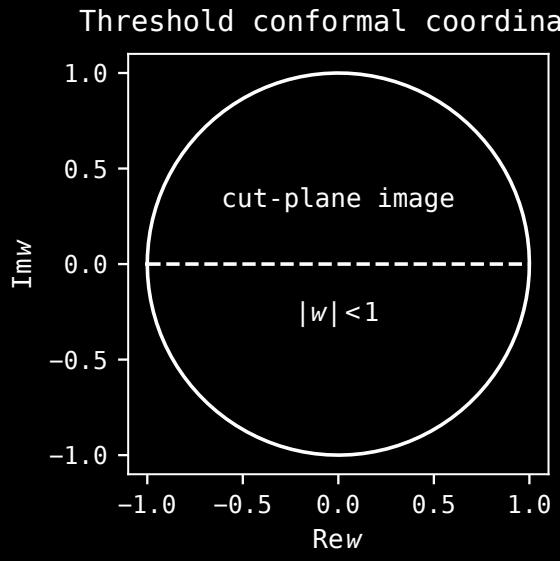


Figure 13.1: The chosen square-root sheet maps to the disk. The square-root branch and the location of additional singularities are part of the physical model.

13.5 A complete single-channel unitarity model

For real $K(E)$ on the real physical-energy interval define

$$S(E) = \frac{1 + iK(E)}{1 - iK(E)}.$$

Its numerator and denominator are complex conjugates, so $|S| = 1$ whenever the denominator is nonzero. Writing $K = \tan\delta$ gives $S = e^{2i\delta}$. If K is a rational function, poles of S satisfy $1 - iK = 0$ after cancellations are resolved. Thus imposing a real K parametrization can make single-channel elastic unitarity exact on the real axis while leaving analyticity and causality as separate obligations.

For a matrix $K = K^*$, the Cayley transform $(I + iK)(I - iK)^{-1}$ is unitary by direct multiplication, since the factors are commuting polynomials/rational functions of K. Positive definiteness of $I + K^2$ ensures invertibility of $I - iK$. A non-Hermitian fitted K has no such guarantee. This is a mathematical parametrization, not a complete physical scattering theory.

13.6 A finite-range bound-state model

Take a spherically symmetric square-well potential as an analytic test model for the radial s-wave Schrödinger equation. In units where $\hbar^2/(2\mu) = 1$, let $V = -V_0$ for $0 < r < a$ and $V = 0$ for $r > a$. For a bound state $E = -\kappa^2$ with $0 < \kappa^2 < V_0$, the regular interior reduced radial solution is $u = A\sin(kr)$, $k^2 = V_0 - \kappa^2$, and the decaying exterior solution is $u = Be^{-\kappa r}$.

Matching u and u' at a yields $k\cot(ka) = -\kappa$. Together with $k^2 + \kappa^2 = V_0$, this is a scalar root problem with explicit branch intervals. The equation is derived from a specified toy potential and represents no realistic evaluated nuclear interaction by itself. Increasing the energy range beyond the model's validity does not improve the physical approximation. Root enclosures can certify solutions of this equation while leaving empirical validity entirely separate.

13.7 Classical worldsheet energy

Let X map a two-dimensional parameter surface to Euclidean target space, and let

$$S[X, g] = \frac{T}{2} \int \sqrt{\det g}\, g^{ab} \partial_a X \cdot \partial_b X \, \mathrm{d}^2\sigma.$$

Under $g \mapsto e^{2\omega} g$, the square-root determinant gains $e^{2\omega}$ and the inverse metric gains $e^{-2\omega}$, so the classical action is invariant. This is an exact dimension-two cancellation. Varying X with fixed boundary values gives the harmonic-map equation componentwise for a flat target. The cylinder Fourier development supplies explicit classical solutions for finite boundary-mode data.

The quantum functional measure, anomaly cancellation, ghosts, non-flat targets and integration over moduli add separate mathematical and physical structures. Polyakov's formulation is a primary source for the quantum theory [11]; the classical cancellation proved here is not its complete derivation. The required full string and supersymmetric specialist programme remains an open commission item.

13.8 Gauge holonomy on an annulus

An elementary flat $U(1)$ connection on an annulus is $A = \alpha\, \mathrm{d}\theta$. Its curvature is zero away from the puncture, yet its loop holonomy is $e^{i2\pi\alpha}$. A single-valued gauge transformation $g = e^{ik\theta}$ with

$k \in \mathbb{Z}$ shifts α by an integer and leaves the holonomy unchanged. Thus local curvature does not determine the global gauge equivalence class.

The calculation follows directly from integrating A around a generator and requiring $g(\theta+2\pi) = g(\theta)$. It is a period obstruction on a multiply connected domain, closely aligned with the source's treatment of path classes. Non-Abelian holonomy requires path ordering and cannot be reduced to the same scalar integral without additional commutativity assumptions.

13.9 Spin boundary conditions and mode shifts

On a circle of circumference ℓ, periodic scalar modes have momenta $2\pi n/\ell$. A field with antiperiodic condition $\psi(x+\ell) = -\psi(x)$ instead has momenta $2\pi(n+1/2)/\ell$, obtained by solving $e^{ik\ell} = -1$. This elementary shift changes spectral sums and removes a zero mode. Naming a "fermionic" boundary condition does not define the full spinor bundle, statistics or functional measure; those structures are additional data.

A one-dimensional supersymmetric quantum-mechanics model can be specified algebraically by $A = \mathrm{d}/\mathrm{d}x + W(x)$ and its formal adjoint $A^* = -\mathrm{d}/\mathrm{d}x + W(x)$. Multiplication gives $A^*A = -\mathrm{d}^2/\mathrm{d}x^2 + W^2 - W'$ and $AA^* = -\mathrm{d}^2/\mathrm{d}x^2 + W^2 + W'$. Positivity holds on a common domain with vanishing boundary terms. Self-adjoint domain selection and spectral equivalence beyond nonzero paired eigenvalues require separate analysis. This explicit model does not stand for a derivation of supergeometry or M-theory.

13.10 Period geometry in a conditional gauge model

A family of algebraic curves with a chosen differential can define periods $a(u)$ and $a_D(u)$. On a region where $a'(u) \neq 0$, the ratio $\tau(u) = a'_D(u)/a'(u)$ is a well-defined local function. The condition $\operatorname{Im}\tau > 0$ gives positivity of the quadratic form $\operatorname{Im}\tau\,|\,\mathrm{d}a|^2$. A change of period basis acts by a fractional linear transformation on τ and preserves the upper half-plane.

Seiberg–Witten theory supplies a specific physical setting in which period geometry occurs [13]. Selecting an arbitrary curve and differential does not establish that physical setting or its spectrum. The corresponding M-theory and compactification interpretations in the commission require their own governing assumptions, derivations and literature reconciliation; they remain open rather than being inferred from the genus-one modulus alone.

Chapter 14

Spectral thermodynamics, horizons and identifiability

14.1 Finite Gaussian systems

Let K be a positive-definite real symmetric $N \times N$ matrix. Diagonalizing it orthogonally and applying the one-dimensional Gaussian integral gives

$$Z(\beta) = \int_{\mathbb{R}^N} e^{-\beta x^T K x/2}\,\mathrm{d}x = \left(\frac{2\pi}{\beta}\right)^{N/2} (\det K)^{-1/2}.$$

The mean quadratic energy is $-\partial_\beta \log Z = N/(2\beta)$. The covariance is $\beta^{-1}K^{-1}$, obtained either by differentiating a source-augmented Gaussian integral or by integrating each diagonal mode. These are exact finite-dimensional identities. They provide a mathematical relation between a stiffness matrix, fluctuations and spectral response.

A graph Laplacian has a constant zero mode, so the unrestricted Gaussian integral diverges. Pinning one vertex, imposing zero mean, or quotienting the constant mode makes a different well-defined model. The determinant used in the formula must correspond to that selected model. A numerical routine returning a tiny nonzero eigenvalue does not justify retaining the divergent mode.

14.2 Finite-mode determinants and the infinite limit

For positive eigenvalues $\lambda_1, \ldots, \lambda_N$, $\log\det K = \sum \log \lambda_j$. A spectral approximation can therefore compute a finite-mode free energy exactly within the finite model. As $N \to \infty$, this sum often diverges. Zeta or other regularizations prescribe a renormalized value through additional analytic continuation and subtraction rules. Distinct subtraction conventions can shift local or constant terms.

The present manuscript makes no unqualified assertion that a finite determinant converges to an infinite-dimensional determinant. The continuum limit needs a named operator, domain, boundary conditions, eigenvalue asymptotics and regularization scheme. The finite graph certificates prove positivity and energy bounds but do not themselves establish a quantum partition function on a continuum surface.

14.3 A bosonic oscillator and its heat capacity

For a single oscillator with energy spacing $\hbar\omega$, the thermal partition function after omitting the constant zero-point factor is $Z = (1 - e^{-\beta\hbar\omega})^{-1}$. The geometric series converges for $\beta\hbar\omega > 0$. Differentiating gives

$$U = \frac{\hbar\omega}{e^{\beta\hbar\omega} - 1}, \qquad \frac{C_V}{k_B} = \frac{x^2 e^x}{(e^x - 1)^2}, \quad x = \beta\hbar\omega.$$

The limit $x \to 0$ is one by the expansion $e^x - 1 = x + O(x^2)$; as $x \to \infty$, it tends to zero exponentially. A finite set of independent modes has the sum of these contributions. Coupled or interacting modes require diagonalization or a different statistical model.

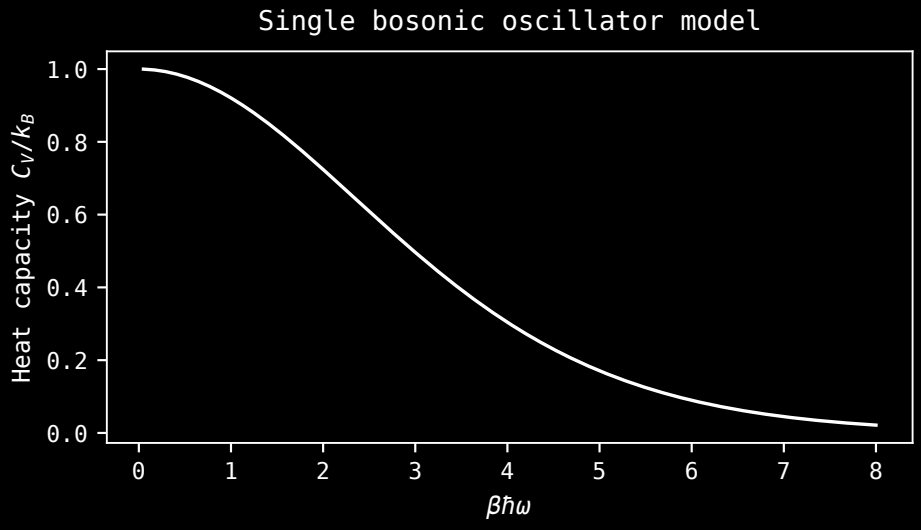


Figure 14.1: The heat capacity of one bosonic oscillator in its specified canonical ensemble. It is a model prediction, not a measured material response.

14.4 Partition zeros and real-axis response

For two energy levels $0, \Delta$ with degeneracies $1, g > 0$, $Z(\beta) = 1 + ge^{-\beta\Delta}$. Complex zeros satisfy

$$\beta\Delta = \log g + (2k + 1)i\pi,$$

up to the equivalent sign convention for the integer index. For real positive β, $Z > 0$ and the free energy is smooth. A thermodynamic phase transition cannot be inferred from this finite two-level example alone. A sequence of systems and a limit in which zeros approach the real axis are additional requirements.

The derivative $\partial_\beta^2 \log Z$ equals the energy variance by direct differentiation of normalized Boltzmann probabilities. Thus the convexity of $\log Z$ is an exact finite-model theorem. Complex zero locations organize analytic continuation, but a claim about a physical phase boundary still needs the actual model and limiting population.

14.5 A local Euclidean horizon calculation

Consider a static two-dimensional metric factor

$$\mathrm{d}s^2 = -f(r)\,\mathrm{d}t^2 + f(r)^{-1}\,\mathrm{d}r^2,$$

with a simple horizon $f(r_h) = 0$ and $f'(r_h) > 0$. Put $t = -i\tau$ and write $r - r_h = f'(r_h)\rho^2/4$. To leading order near the horizon,

$$\mathrm{d}s_E^2 = \mathrm{d}\rho^2 + \left(\frac{f'(r_h)}{2}\right)^2 \rho^2\,\mathrm{d}\tau^2 + O(\rho^2)\,\mathrm{d}\rho^2 + O(\rho^4)\,\mathrm{d}\tau^2.$$

Smooth polar geometry requires τ to have period $4\pi/f'(r_h)$. In a model identifying this Euclidean period with inverse temperature in units $\hbar = c = k_B = 1$, the temperature is $f'(r_h)/(4\pi)$.

This is a local regularity calculation. It assumes a nonextremal simple horizon and an appropriate equilibrium interpretation. An extremal zero of higher order gives a different limiting geometry. The calculation is not a proof of the full Hawking emission result [12], nor does it determine gray-body factors, backreaction, quantum-gravity corrections or empirical radiation from an observed black hole.

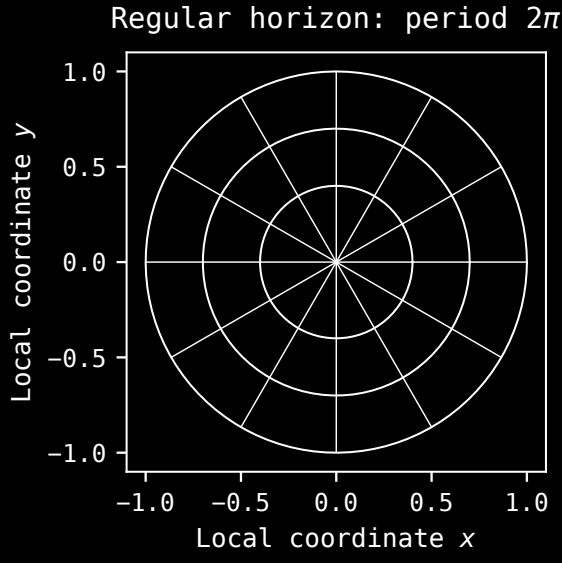


Figure 14.2: A regular polar origin requires a full angular period. The Euclidean horizon calculation uses this local condition together with separately stated thermodynamic assumptions.

14.6 Area, entropy and ensemble stability

For the nonrotating asymptotically flat Schwarzschild model, one may adopt $T \propto M^{-1}$ and $S \propto M^2$ under the specified semiclassical conventions. The derivative $\mathrm{d}M/\,\mathrm{d}T$ is negative. This is not automatically instability in every ensemble: a canonical heat bath, a finite cavity, an asymptotically anti-de Sitter boundary and a microcanonical constraint are distinct physical settings. The sign of a single derivative must be interpreted with the allowed fluctuations and conserved quantities.

A rigorous variational comparison therefore specifies the free-energy functional, boundary terms and admissible perturbations. The complete black-hole thermodynamic programme requested in the commission remains open beyond the explicitly derived local regularity model. No geometric surface capacity is identified with gravitational entropy without a separately derived physical relation.

14.7 A statistical obstruction for multiverse-model claims

Proposition 14.1 (Indistinguishable observation models)**.** [P01] *If two proposed models assign the same probability distribution P to all available observations, no decision rule based on those observations distinguishes the models with average success probability exceeding* $1/2$ *under equal priors.*

Proof. Let a possibly randomized decision rule output model one with probability $\varphi(x) \in [0,1]$. Its average success is

$$\frac{1}{2}\int \varphi \,\mathrm{d}P + \frac{1}{2}\int (1-\varphi)\,\mathrm{d}P = \frac{1}{2}.$$

The equality holds for every rule. Increasing computing resources does not alter the distributions and cannot change the conclusion. □

More generally the optimal average advantage is bounded by half the total-variation distance, obtained by bounding $\int \varphi \,\mathrm{d}(P-Q)$. A proposal involving unobserved sectors requires a defined measure, observation map and predictions before such a distance can even be evaluated. The proof identifies a general identifiability condition; it neither affirms nor denies any particular cosmological ontology.

14.8 Physical data and remaining execution tasks

No experimental nuclear, particle, plasma or astrophysical measurements are incorporated as fitted data in this manuscript. The mathematical test models are expressly synthetic or analytic. The commission's measured-data work, uncertainty/covariance handling, model validation and broader physical implementation remain uncompleted. An accurate calculation in a stated model is retained as such; it is not described as a national-laboratory validation or an experimental result.

Chapter 15

Algorithms, information and distributed computation

15.1 A chain admits a sufficient statistic

Let a chain have m positive conductances c_i, endpoints at potentials 0 and 1, and cumulative resistances $R_k = \sum_{i=1}^{k} c_i^{-1}$. Conservation at an interior vertex states that $c_i(u_i - u_{i-1})$ is independent of i. Summing the potential increments determines that current as $I = 1/R_m$, and hence

$$u_k = R_k/R_m, \qquad C = 1/R_m.$$

This derivation is simultaneously a direct solver and a certificate generator. The input is the entire conductance vector, not merely its minimum, average or number of entries. Any exact algorithm in an entry-query model must inspect every entry: if it omits c_j, two positive vectors agreeing on all inspected entries but differing at j have different R_m. The algorithm has the same transcript for both and must return a wrong value on at least one. Thus an $\Omega(m)$ input-access lower bound matches the $O(m)$ arithmetic construction for the scalar conductance. Producing all $m - 1$ interior potentials also has an $\Omega(m)$ output-size bound.

A dense Gaussian elimination representation performs unnecessary work on this particular topology. A tridiagonal elimination method already has linear complexity. The appropriate strong control is therefore tridiagonal elimination, not only the dense representation. The experiment in Chapter 33 includes both controls. The proof of the linear bound applies to arbitrary positive chain conductances and does not depend on the measured timings.

For exact rational inputs, the sums may acquire long denominators. Store fractions in reduced form and charge for multiplication, division and greatest-common-divisor calculations. The arithmetic lower bound does not settle the optimal bit complexity of all possible rational representations. For interval inputs, monotonicity supplies a different contract, developed below: an enclosure of C can often be obtained without recovering an exact fraction.

15.2 One-edge updates and their exact restriction

Proposition 15.1. [C01] *Let A be positive definite, $b \in \mathbb{R}^n$, and $\delta \geq 0$. Then*

$$(A + \delta bb^T)^{-1} = A^{-1} - \frac{\delta A^{-1} bb^T A^{-1}}{1 + \delta b^T A^{-1} b}.$$

Proof. Put $v = A^{-1}b$ and $s = b^T v \geq 0$. The denominator $1 + \delta s$ is positive. Multiplication of the right-hand side by $A + \delta bb^T$ gives I plus a multiple of bv^T whose coefficient is $\delta - \delta/(1 +$

$\delta s) - \delta^2 s/(1+\delta s) = 0$. Both matrices are square, so the right-hand side is the inverse. Positive definiteness also follows from $x^T(A+\delta bb^T)x = x^TAx + \delta(b^Tx)^2 > 0$ for $x \neq 0$. □

A grounded graph Laplacian changes by this form when one edge conductance changes and both endpoints are retained in the grounded coordinates; a boundary-edge update also changes the forcing term when the prescribed boundary potential is nonzero. The latter term must be updated as well. Applying the inverse identity to an unchanged right-hand side in that situation solves a different network.

After an existing factorization, computing v requires one solve. Applying the rank-one correction to another solution requires a dot product and a scaled vector update. For many right-hand sides this can avoid repeated factorizations. If the change is negative, invertibility requires $1+\delta s \neq 0$ and positive definiteness requires $1+\delta s > 0$; arbitrary negative conductances are not admitted by the physical model. Updating repeatedly without controlling rounding may lose numerical accuracy even though the exact identity remains valid.

15.3 Iterative solution with an explicit error contract

Let A be real symmetric positive definite with eigenvalues in $[a,b]$, $0 < a \leq b$. Consider the stationary iteration

$$x_{k+1} = x_k + \omega(f - Ax_k), \qquad \omega = \frac{2}{a+b}.$$

Proposition 15.2. [C02] *If $x_* = A^{-1}f$, then*

$$\|x_k - x_*\|_A \leq \left(\frac{b-a}{b+a}\right)^k \|x_0 - x_*\|_A.$$

Among real constant step sizes, this ω minimizes the worst scalar error multiplier on $[a,b]$.

Proof. An orthogonal eigenbasis reduces the error update to multiplication by $1-\omega\lambda$. At the chosen step size this affine function has endpoint values $(b-a)/(a+b)$ and its negative. Its absolute value nowhere exceeds their common magnitude. Squaring the coordinate errors with their positive eigenvalue weights proves the norm estimate. For any candidate step, if both endpoint magnitudes were smaller than $q = (b-a)/(b+a)$, then $1-\omega a < q$ would imply $\omega > 2/(a+b)$, whereas $1-\omega b > -q$ would imply $\omega < 2/(a+b)$, a contradiction. The case $a = b$ gives one exact-arithmetic step. □

This elementary result supplies a computable stopping estimate when certified spectral bounds and an initial-error bound are available. A residual alone gives

$$\|x - x_*\|_A^2 = r^TA^{-1}r \leq \|r\|_2^2/a, \qquad r = f - Ax.$$

The displayed bound follows from diagonalization and $1/\lambda_i \leq 1/a$. The value a must be established; replacing it by a numerical estimate without a lower-bound certificate changes the status of the conclusion. More advanced Krylov and multigrid methods require their own convergence and implementation analysis and are not benchmarked here.

15.4 Parallel work, depth and communication

The chain formula is a reduction followed by a prefix sum. A balanced binary tree on m inputs uses $m-1$ additions and depth $\lceil \log_2 m \rceil$ for a total sum. A prefix-scan construction uses $O(m)$ operations and $O(\log m)$ depth: an upward traversal stores subtree sums, and a downward

traversal sends to a right child the parent prefix plus its left sibling's sum. Induction on subtrees proves that every leaf receives the sum of entries preceding it. One final addition forms each inclusive prefix.

On p distributed workers, each worker can sum and scan its local consecutive block. A scan of the p block totals determines the offsets. The arithmetic work is $O(m)$, while the idealized depth is $O(m/p + \log p)$. The communication statement is conditional on the specified tree and network model: $O(p)$ transmitted scalar block values overall and $O(\log p)$ rounds along a critical path. It is not a wall-time prediction. Fraction bit lengths, message headers, network latency and synchronization costs can dominate the numerical arithmetic.

Ordinary floating-point addition is not associative. A fixed reduction tree yields repeatable results under a fixed arithmetic environment, but not generally the same bits as a serial sum. Exact rational sums are associative, while directed interval sums provide guaranteed enclosures whose widths depend on the operation order. A parallel implementation is accepted against its output contract, not by requiring accidental bit equality with an unrelated summation order.

15.5 Consensus as a spectral statement

Let a connected undirected graph have Laplacian L. Take $0 < \eta < 2/\lambda_{\max}(L)$ and $x_{k+1} = (I - \eta L)x_k$. Since $L\mathbf{1} = 0$, the mean $\bar{x} = n^{-1}\mathbf{1}^T x_0$ is invariant. On the orthogonal complement of constants, the eigenvalues are positive, and

$$\|x_k - \bar{x}\mathbf{1}\|_2 \le q^k \|x_0 - \bar{x}\mathbf{1}\|_2, \quad q = \max_{i \ge 2} |1 - \eta\lambda_i| < 1.$$

The proof is the same coordinate argument as in 15.2, with the null eigenspace removed. Disconnected graphs retain one independent constant on each component. The algorithm then converges to component means, not a global mean. This is a mathematical negative control for an implementation that silently assumes connectedness.

Message loss, asynchronous updates, adversarial participants and changing graphs are separate models. The synchronous recurrence does not prove their safety or liveness. No consensus protocol under malicious faults is implemented in this review. The relevance to Section 6 is the shared positive quadratic form and the energy of deviations, not an asserted equivalence between consensus and conformal modulus.

15.6 Cryptographic and quantum boundaries

An elliptic curve over $\mathbb{C}$ may be represented analytically by a lattice quotient, as developed in the preceding volume. A finite-field group used in cryptography is a different object. For example, on $y^2 = x^3 + 2x + 2$ over $\mathbb{F}_{17}$, $P = (5, 1)$ lies on the curve because $1 \equiv 125 + 10 + 2 \pmod{17}$. The tangent slope for doubling is $(3 \cdot 25 + 2)/(2 \cdot 1) \equiv 77 \cdot 9 \equiv 13$, giving $2P = (6, 3)$. This modular calculation contains no conformal modulus or ordered positive energy. An extremal-length estimate for complex periods is therefore not a security reduction for this finite group.

A useful public mathematical contribution could concern validated conversion of complex parameters, certified algebraic identities, or implementation invariants, but it would need an explicit theorem connecting its objects. Security additionally requires an adversary model, parameter regime and hardness assumption. None is supplied by the small example. Lattice, isogeny, code, zero-knowledge and multiparty constructions remain uncompleted specialist obligations.

Likewise, replacing a linear solve by a quantum-state preparation problem changes the output. The graph certificate here contains every potential and flow coefficient, so any claimed quantum

advantage must pay for producing or checking that classical witness. State preparation, input access, precision and readout cannot be omitted from a comparison. The manuscript makes no quantum speedup claim.

15.7 Logical complexity and proof reuse

For a fixed graph, one can verify a rational witness without reconstructing the path family, the optimizer's search history or a dense factorization. The proof dependence is compact: positivity and feasibility, a summation-by-parts identity, and Cauchy–Schwarz imply the bracket. The certificate size is proportional to the encoded graph, potential and flow, not to the number of possible source-to-target paths.

This observation is a representation statement. It is not a lower bound for every proof system, nor a proof that a reader understands the argument faster. The formalization plan should name the rational data type, incidence convention, arithmetic semantics and theorem statement. Kernel term sizes, compiler refinement and cross-system proofs remain unexecuted under the commission's present-phase decision.

Chapter 16

Exact arithmetic, enclosures and certificate completeness

16.1 Finite witnesses and the logical quantifiers

A certificate checker accepts a particular finite input. Its soundness theorem is universal over the admitted encoding: if the checker accepts the data and witness, then the mathematical predicate holds. Exhaustively trying many witnesses does not prove that implication. The written graph proof in Chapter 9 supplies the mathematical implication; the executable checker realizes its arithmetic steps with an explicitly named rational library.

There are three independent questions. Soundness asks whether every accepted witness is valid. Completeness asks whether every valid input admits a witness in the chosen format. Search termination asks whether a particular algorithm eventually finds one. For connected rational networks with positive conductances, all three can be addressed in finite exact arithmetic. For arbitrary continuum geometry, the same finite rational format does not even specify the domain, so no analogous completeness claim is made.

16.2 Rational existence for connected networks

Let the prescribed potential be zero at s and one at t. Remove both terminal coordinates from the quadratic energy. It has the form

$$E(x) = x^T A x - 2b^T x + c,$$

where A is positive definite on the free vertices and all entries are rational when the conductances are rational. Positive definiteness follows because a zero-energy extension with zero terminal values is constant along every edge and therefore zero by connectedness. Thus A is nonsingular. Gaussian elimination over $\mathbb{Q}$ terminates after finitely many exact field operations, using a nonzero pivot at each step after a permitted row interchange. The unique solution is rational.

The resulting terminal current C is positive and rational. Dividing the edge currents by C gives a rational unit flow. Its energy is $1/C$. Consequently an exact zero-gap witness exists for every finite connected positive rational network with distinct terminals. This is completeness of the witness format on that class. It does not bound the intermediate integers by a fixed machine word.

A fraction-free elimination or determinant representation can control denominators. After clearing denominators, Cramer's rule writes each coordinate as a ratio of integer determinants. If the resulting integer matrix has entries bounded by H, a determinant expansion gives the

elementary bound $n!H^n$. Its bit length is at most $\log_2(n!) + n\log_2 H + 1$. This is a crude finite bound, sufficient to establish that the rational witness has finite size. Sharper bit-complexity optimization is a separate algorithmic problem.

16.3 An algebraic positive-definiteness certificate

Proposition 16.1. [V01] *Let A be symmetric with rational entries. Suppose $A = LDL^T$, where L is lower triangular with diagonal entries one and D is diagonal with strictly positive rational entries. Then A is positive definite.*

Proof. The triangular matrix L is invertible, because forward substitution solves $Ly = z$ uniquely for every z. For $x \neq 0$, $L^T x \neq 0$, so $x^T Ax = (L^T x)^T D(L^T x) = \sum_i D_{ii}(L^T x)_i^2 > 0$. □

A checker can verify the equality and diagonal signs exactly. Conversely, a positive-definite matrix admits such a decomposition: eliminate its positive first diagonal entry and take the Schur complement. For any nonzero y, the vector $(-a^{-1}b^T y, y)$ proves positivity of that complement. Induction constructs the remaining factors. This proof supplies both existence and a terminating exact algorithm. It also shows why an indefinite matrix cannot be accepted merely because some numerical pivots happened to be positive after unrecorded rounding.

For a sparse matrix the full product check may introduce fill. A computational claim must charge for that fill or use an equivalent verified factorization trace. The sparse graph-energy representation already supplies positivity by connectedness and positive conductances; an LDL^T certificate is an alternative for more general forms, not a mandatory redundant step for every graph.

16.4 Directed interval semantics

An interval $[a, b]$ denotes all real values between its rational endpoints. Addition and subtraction have exact endpoint rules. Multiplication takes the minimum and maximum of the four endpoint products. Division is admitted only when the denominator interval excludes zero. These rules are sound because affine functions attain extrema at interval endpoints, applied successively to each variable. They are generally not exact when a variable occurs repeatedly: $x - x$ for $x \in [0, 1]$ evaluates to $[-1, 1]$ by independent interval subtraction, although the true range is $\{0\}$.

Algebraic simplification and variable dependencies therefore affect enclosure width. Soundness is retained under every rule, but a wide interval may be useless. A certificate must include the expression actually enclosed. Replacing x^2 with a fresh independent product interval may lose sharpness; replacing it with an unproved narrower interval loses soundness.

A rational interval implementation does not require directed binary rounding for its elementary operations. Conversion from measured decimal data or from binary floating-point output must still produce a containing interval. A nearest-rounded decimal string is not automatically a certified enclosure of the original physical value. Input uncertainty and arithmetic error occupy distinct fields.

16.5 A logarithm enclosure with an explicit tail

Proposition 16.2. [V02] *Let $x > 0$, $t = (x-1)/(x+1)$ and $N \geq 0$. Then*

$$\log x = 2\sum_{k=0}^{N} \frac{t^{2k+1}}{2k+1} + R_N, \qquad |R_N| \leq \frac{2|t|^{2N+3}}{(2N+3)(1-t^2)}.$$

Proof. For $|s| < 1$, the finite geometric identity gives $(1 - s^2)^{-1} = \sum_{k=0}^{N} s^{2k} + s^{2N+2}/(1 - s^2)$. Integrate from 0 to t and multiply by two. The integral of $2/(1 - s^2)$ is $\log((1 + t)/(1 - t)) = \log x$. On the integration segment, $1 - s^2 \geq 1 - t^2$, so the absolute remainder is bounded by $2 \int_0^{|t|} s^{2N+2}\, \mathrm{d}s/(1 - t^2)$, which is the displayed expression. □

For rational x, both the partial sum and bound are rational. The certificate is thus a finite list of exact arithmetic operations together with the proved tail. When x is very large or small, $|t|$ approaches one and convergence slows. Repeated square-root reduction can improve convergence but requires certified square-root intervals. A binary exponent decomposition can reduce the argument at the expense of a separate enclosure for $\log 2$.

For an annular modulus one also needs an enclosure of π. Alternating arctangent sums combined with Machin's identity provide a route: the identity itself must first be proved from the tangent addition formula and the relevant quadrant. This manuscript's executable numerical illustrations use the system transcendental functions and are labelled diagnostics; they are not described as implementations of this exact logarithm/ π certificate pipeline.

16.6 Exact square-root enclosures

For a nonnegative rational a, rational bounds $0 \leq l \leq u$ with $l^2 \leq a \leq u^2$ certify $l \leq \sqrt{a} \leq u$ by monotonicity of squaring on nonnegative reals. Starting with a finite containing interval, bisection reduces its width by a factor of two at each step; one exact rational square comparison selects the next half. After k steps, the interval width is 2^{-k} times the initial width. This proves termination for every positive prescribed rational absolute tolerance.

An equality test between an arbitrary real input and a threshold is different. If the true value equals the threshold, nested approximate intervals can continue straddling it forever. The canonical weight $w = \max(a - 2, 0)$ is a Lipschitz function of a real parameter, but deciding the strict predicate $w > 0$ from arbitrary approximation oracles requires a separation promise. A positive certificate $a_- > 2$ suffices; failure to obtain it is not a certificate of zero weight.

16.7 Input intervals for conductance

Suppose $c_e \in [c_e^-, c_e^+]$ with all $c_e^- > 0$. Rayleigh monotonicity gives

$$C(c^-) \leq C(c) \leq C(c^+).$$

If exact or certified solves give lower and upper bounds on the endpoint networks, their outer bounds enclose the entire uncertain network family. This includes correlated choices of the conductances because every permitted vector lies between the componentwise extremes. It may be conservative when correlations exclude those extreme vectors.

For a chain, the sharper formula follows directly:

$$\left(\sum_i \frac{1}{c_i^-}\right)^{-1} \leq C \leq \left(\sum_i \frac{1}{c_i^+}\right)^{-1}.$$

Every endpoint is attained when the uncertainty box allows its corner. Thus these bounds are optimal among bounds that use only the independent intervals and no additional correlation information. This is a precisely scoped information-optimality statement; it makes no claim about optimal hardware execution.

16.8 Termination and malformed certificates

The graph checker parses finite records with explicit dimension and edge-count limits. It rejects zero denominators, nonpositive conductances, identical terminals, loops, duplicate unordered edges, invalid vertex indices and disconnected graphs. Boundary values and every divergence equation are checked. Claimed bounds are compared with the actual rational energies, not trusted as supplied annotations.

The checker accepts feasible nonoptimal witnesses as well as optimizers. Requiring zero gap would turn a useful certified approximation procedure into an exact-solution recognizer. The positive-gap example in the supplement tests this distinction. Conversely, accepting a potential alone supplies only an upper bound on conductance; without a unit-flow witness there is no matching lower bound of the stated form.

The implementation's use of arbitrary-precision rational arithmetic prevents fixed-width integer overflow within that abstraction. Allocation failures, library defects, malicious resource exhaustion and compiler correctness remain external implementation concerns. The future formalization plan isolates these layers. The current files contain written soundness arguments and executable finite tests, with no claimed proof-assistant acceptance.

Chapter 17

Source-selected problems: complete finite-dimensional treatments

17.1 Selection boundary and mathematical role

Four source questions have received individual statement review for this manuscript. Their roles are deliberately distinct: a negative kernel example prevents misuse of an energy representation; an explicit feature construction connects a positive form to Euclidean geometry; a statistical minimum-variance result supplies a second projection setting; and a low-rank kernel approximation exposes the cost and exactness of a compressed representation. These are auxiliary finite-dimensional applications. They are not additional theorems of Section 6, and their inclusion is not a claim that every candidate in the complete question bank has been adjudicated.

The reference is the named book, edition and exercise. Statements below are concise mathematical specifications of the inspected tasks; original source-page reproductions are not distributed. The source's complete subpart structure is preserved. Proofs use real matrices; the Hermitian variant requires conjugate transposes and has the same geometric construction.

17.2 QK01. An absolute-difference kernel is not positive semidefinite

Reference. Mohri, Rostamizadeh and Talwalkar, *Foundations of Machine Learning*, second edition, Exercise 6.7 [14].

Task. For $K(x, y) = |x - y|$ on $\mathbb{R}$, determine whether every Gram matrix is positive semidefinite and prove the conclusion.

17.2.1 Direct quadratic-form solution

Take $x_1 = 0, x_2 = 1$. The Gram matrix is

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

For $a = (1, -1)^T$, $a^T G a = -2 < 0$. Therefore K is not a positive semidefinite kernel. The counterexample has only two sample points and integer arithmetic. No numerical eigenvalue computation is needed. Equivalently, its eigenvectors $(1, 1)$ and $(1, -1)$ have eigenvalues 1 and -1.

17.2.2 Structural solution from diagonal data

Suppose a positive semidefinite kernel admits the feature representation $K(x, y) = \langle \Phi(x), \Phi(y) \rangle$. Then $K(x, x) = \|\Phi(x)\|^2$. Here the diagonal vanishes everywhere, forcing every feature vector to be zero and every kernel value to be zero. This contradicts $K(0, 1) = 1$. The same implication can be obtained without assuming a feature theorem: positivity of the two-point quadratic form $a^2K(x, x) + 2abK(x, y) + b^2K(y, y)$ and its zero diagonal forces the off-diagonal term to vanish.

17.2.3 Comparison and limitation

For the task of refuting universal positive semidefiniteness, the two-point exact witness is sufficient and asymptotically constant size. Building an $m \times m$ Gram matrix and diagonalizing it for an arbitrary data set incurs unnecessary work and supplies no stronger universal disproof. The structural argument explains the entire failure mechanism and applies to any nonzero kernel with identically zero diagonal.

The conclusion concerns positive semidefiniteness, not conditional negative definiteness. Some distance functions support different constructions after centering, exponentiation or restriction to coefficients summing to zero. Such a transformation needs its own theorem. Replacing $|x - y|$ by another function would alter the source task.

17.3 QK02. An explicit feature map and its converse

Reference. Foundations of Machine Learning, second edition, Exercise 6.11(a)–(b) [14].

Task. Given a symmetric positive semidefinite $m \times m$ Gram matrix K, construct vectors whose pairwise inner products are its entries. Conversely, show that a Hilbert-space feature map yields a positive semidefinite Gram matrix.

17.3.1 Part (a): spectral construction

By the finite-dimensional spectral theorem, $K = Q\Lambda Q^T$ with Q orthogonal and $\Lambda = \operatorname{diag}(\lambda_i)$, $\lambda_i \geq 0$. Put $Z = \Lambda^{1/2}Q^T$ and define $\Phi(x_j)$ to be column j of Z. Then

$$\langle \Phi(x_i), \Phi(x_j) \rangle = (Z^T Z)_{ij} = K_{ij}.$$

Discarding zero eigenvalue coordinates realizes the map in $\mathbb{R}^r$, where $r = \operatorname{rank} K$. No realization in a smaller Euclidean dimension is possible because a Gram matrix of vectors in $\mathbb{R}^d$ has rank at most d. Thus the feature dimension r is optimal for exact realization of this finite Gram matrix.

If the input list repeats the same object, the corresponding kernel rows agree. The construction then assigns the same feature vector: the squared distance of the two vectors is $K_{ii}+K_{jj}-2K_{ij} = 0$. For an arbitrary positive matrix unaccompanied by a globally defined kernel, the construction specifies only the finite indexed data. It cannot determine unseen kernel values.

17.3.2 An independent quotient-space construction

Define a symmetric bilinear form on $\mathbb{R}^m$ by $[a, b] = a^T Kb$. Positivity gives a seminorm. Let $N = \ker K$. Quotient by N and set $\langle [a], [b] \rangle = a^T Kb$. This is well defined because adding a vector in N changes neither argument value. Positive definiteness on the quotient follows from $a^T Ka = 0 \Rightarrow Ka = 0$, seen directly in an eigenbasis. Map x_i to the class of the standard vector e_i. Its inner products are K_{ij}, so the quotient itself is the feature space. Its dimension is $m - \dim N = r$ by rank-nullity.

This construction identifies the intrinsic geometry without choosing coordinates. To store and compute feature vectors, a basis and an actual factorization are still required. It therefore does not by itself beat the arithmetic cost of the spectral construction.

17.3.3 Part (b): positivity from a feature map

For any finite real coefficient vector a,

$$a^T K a = \sum_{i,j} a_i a_j \langle \Phi(x_i), \Phi(x_j) \rangle = \left\| \sum_i a_i \Phi(x_i) \right\|^2 \geq 0.$$

Symmetry follows from the real inner product. This proves every requested finite Gram matrix is positive semidefinite. Completeness of the ambient Hilbert space is not needed for this finite sum, although it matters for further infinite-dimensional theory.

17.3.4 Computational comparison

A dense spectral factorization has the usual cubic arithmetic cost and quadratic storage. Given a verified rank-r factor directly, checking all m^2 inner products costs $O(m^2 r)$ arithmetic operations, and the feature representation stores mr numbers. A full explicit Gram output still requires $\Omega(m^2)$ scalar entries. Claims of subquadratic cost must change the output contract to factor access or matrix-vector products. This distinction mirrors the distinction between a path-family invariant and an explicit listing of every path.

17.4 QS01. Gauss–Markov optimality as an orthogonal decomposition

Reference. Hastie, Tibshirani and Friedman, *The Elements of Statistical Learning*, second edition, Exercise 3.3(a)–(b), referring to the linear-model setting of Section 3.2.2 [15].

Model. Let $Y = X\beta + \varepsilon$, with a fixed real $n \times p$ matrix X of full column rank, $\mathbb{E}\varepsilon = 0$ and $\mathrm{Cov}(\varepsilon) = \sigma^2 I_n$. Gaussian noise is not assumed. The comparisons are among estimators linear in Y and unbiased for every β in this model. Heteroscedasticity or correlated noise changes the inner product and requires a separate weighted formulation.

17.4.1 Part (a): scalar parameter

Let $a \in \mathbb{R}^p$. A linear estimator $c^T Y$ is unbiased for $a^T\beta$ exactly when $X^T c = a$. The least-squares estimator has coefficient

$$c_0 = X(X^T X)^{-1} a.$$

Any other feasible c is $c = c_0 + d$, with $X^T d = 0$. Since $c_0 \in \mathrm{range}\, X$, $c_0^T d = 0$. Therefore

$$\mathrm{Var}(c^T Y) = \sigma^2 \|c\|_2^2 = \sigma^2 \|c_0\|_2^2 + \sigma^2 \|d\|_2^2 \geq \mathrm{Var}(c_0^T Y).$$

For $\sigma^2 > 0$, equality requires $d = 0$. If $\sigma^2 = 0$, every admissible linear estimator has zero variance and uniqueness of the estimator is not inferred from the variance criterion.

A second route minimizes $\|c\|^2$ under $X^T c = a$ using a multiplier. Stationarity gives $c = X\lambda$ and $X^T X \lambda = a$, yielding c_0. The positive Hessian proves a global minimum on the affine constraint set. The orthogonal decomposition supplies the stronger exact excess-variance identity, which the multiplier equations alone do not display until the square is completed.

17.4.2 Part (b): covariance ordering

Write the least-squares matrix as $B_0 = (X^T X)^{-1} X^T$. Any linear unbiased estimator BY satisfies $BX = I_p$. Set $D = B - B_0$; then $DX = 0$ and $B_0 D^T = (X^T X)^{-1} X^T D^T = 0$. Consequently

$$\mathrm{Cov}(BY) = \sigma^2 BB^T = \sigma^2 (X^T X)^{-1} + \sigma^2 DD^T.$$

The last term is positive semidefinite because $v^T DD^T v = \|D^T v\|^2 \geq 0$. Thus $\widehat{V} \preceq \widetilde{V}$ in the source's matrix ordering. Conversely, applying this matrix inequality to an arbitrary a gives the scalar variance comparison. This completes both requested parts.

17.4.3 Relation to energy minimization

The feasible estimator coefficients form an affine subspace, and minimum variance is its minimum squared norm. In an electrical network the feasible unit flows also form an affine subspace, with a resistance-weighted squared norm. Both comparisons follow from orthogonal projection and yield an exact excess objective. Neither application needs a new theorem asserting that every optimizer in every statistical model is a least-squares estimator.

A conventional direct comparison expands all covariances and cross terms. The projection formulation proves the scalar and matrix statements together through the vanishing cross term. No execution-time or cognitive-time improvement is claimed for a written proof. The objective improvement is exact and model-wide: the covariance difference is the explicitly positive matrix $\sigma^2 DD^T$.

17.5 QK03. Nyström approximation with singular landmarks

Reference. Foundations of Machine Learning, second edition, Exercise 15.4(a)–(e) [14].

Task. With $K \succeq 0$, write

$$K = \begin{pmatrix} W & K_{21}^T \\ K_{21} & K_{22} \end{pmatrix}, \qquad C = \begin{pmatrix} W \\ K_{21} \end{pmatrix}, \qquad \widetilde{K} = CW^\dagger C^T.$$

The symbol $W^\dagger$ denotes the Moore–Penrose inverse. All parts below allow singular W.

17.5.1 Part (a): the exact error block

A principal block of a positive semidefinite matrix is positive semidefinite: for $z \in \mathbb{R}^l$, $z^T W z = (z, 0)^T K (z, 0) \geq 0$. If $z \in \ker W$, positivity applied to (z, tw) gives

$$2t\, w^T K_{21} z + t^2 w^T K_{22} w \geq 0 \quad \text{for all } t \in \mathbb{R}.$$

Its linear coefficient must vanish, or a sufficiently small t of the opposite sign makes it negative. Thus $K_{21} z = 0$. Since $WW^\dagger$ is the orthogonal projector onto $(\ker W)^\perp$, it follows that $K_{21} W^\dagger W = K_{21}$ and $WW^\dagger K_{21}^T = K_{21}^T$. Multiplying the blocks gives

$$\widetilde{K} = \begin{pmatrix} W & K_{21}^T \\ K_{21} & K_{21} W^\dagger K_{21}^T \end{pmatrix}.$$

Therefore $K - \widetilde{K}$ has only its lower-right block nonzero, and

$$\|K - \widetilde{K}\|_F = \|K_{22} - K_{21} W^\dagger K_{21}^T\|_F.$$

The range argument is essential. Replacing $W^\dagger W$ by I would be incorrect for a singular landmark block.

17.5.2 Part (b): projection representation

Let $K = X^T X$ and let X_0 consist of the landmark columns. Then $W = X_0^T X_0$ and $C = X^T X_0$. Write a reduced singular value decomposition $X_0 = U\Sigma V^T$ with positive singular values. Direct multiplication yields

$$X_0(X_0^T X_0)^\dagger X_0^T = UU^T = P_{\operatorname{range} X_0}.$$

Indeed, the positive singular-value factors cancel, while the zero singular-value directions contribute zero. Substitution gives

$$\boxed{\widetilde{K} = X^T P_{\operatorname{range} X_0} X.}$$

This is the requested projection onto the span of the relevant left singular vectors.

17.5.3 Part (c): positivity

For any v, $v^T \widetilde{K} v = \|P_{\operatorname{range} X_0} X v\|^2 \geq 0$. Symmetry is immediate. Moreover,

$$K - \widetilde{K} = X^T (I - P) X \succeq 0.$$

The latter residual positivity is an additional conclusion, not a replacement for the requested answer. Alternatively, positivity follows directly because $W^\dagger \succeq 0$ and $v^T C W^\dagger C^T v \geq 0$.

17.5.4 Part (d): exact recovery

Since $\ker(X^T X) = \ker X$, $\operatorname{rank} K = \operatorname{rank} X$, and similarly $\operatorname{rank} W = \operatorname{rank} X_0$. The inclusion $\operatorname{range} X_0 \subseteq \operatorname{range} X$ plus equality of their finite dimensions gives equality of the spaces. Hence $PX = X$ and $\widetilde{K} = K$. This argument requires equality of the two ranks, not merely that both are small. Landmark selection can miss a nonzero direction even in a low-rank data set.

17.5.5 Part (e): storage accounting

For $m = 20,000,000$ and eight bytes per double, the dense matrix requires

$$8m^2 = 3,200,000,000,000,000 \text{ bytes} = 3.2 \text{ decimal PB}.$$

Storing C and W as separate dense arrays with $l = 10,000$ requires

$$8(ml + l^2) = 1,600,800,000,000 \text{ bytes} = 1.6008 \text{ decimal TB}.$$

The storage ratio is $m^2/(ml + l^2) \approx 1999.0005$. The upper block W is already present inside C; eliminating duplicate storage can save $8l^2$ bytes when the implementation actually reuses that block. The stated count follows the conventional separate-array representation. Factorization workspace, indexes, precision guards and streamed input are excluded from both counts and must be added for an executable memory plan. The reduced representation remains much too large to assume that it fits an unspecified workstation.

17.5.6 Method comparison and unresolved approximation quality

The block proof checks exact algebra and singular range conditions. The projection proof simultaneously explains positivity, residual structure and rank-dependent exactness. Both are complete for the source parts. Neither proves a small residual for arbitrary landmarks. Obtaining a deterministic or probabilistic error guarantee requires a landmark-selection rule and assumptions on the spectrum or leverage distribution.

Applying the factorized approximation to a vector can use $O(ml + l^2)$ arithmetic after preparation, compared with $O(m^2)$ for a dense matrix-vector product. A full materialized $m \times m$ output still has quadratic storage. Its preparation cost includes obtaining C, choosing landmarks and applying $W^{\dagger}$; the expression is not a promise of instantaneous compression. This matches the manuscript's general requirement that setup, evaluation and certification costs be counted separately.

Chapter 18

Solved investigations in extremal geometry and computation

18.1 Problem-family conventions

The investigations in this chapter are authored problems. They do not inherit a textbook exercise number and are not counted as original research results. Each asks for a specific quantity or conclusion under a fixed model. A conventional calculation is compared with a structural argument when that comparison adds information. Existing theorems are cited at their principal locations; the additional computations and conclusions are developed here.

18.2 VI01. Optimal allocation in a layered conductor

Problem. A series conductor has positive layer coefficients $a_1, \dots, a_m$. Choose positive resource allocations x_i with $\sum_i x_i = V$ to minimize $R(x) = \sum_i a_i/x_i$. Determine the optimum, its equality case and a certificate.

Cauchy–Schwarz gives

$$\left(\sum_i \sqrt{a_i}\right)^2 = \left(\sum_i \sqrt{a_i/x_i}\sqrt{x_i}\right)^2 \le R(x)V.$$

The lower bound is attained precisely when $\sqrt{a_i/x_i} = \lambda\sqrt{x_i}$ for all i. Thus

$$x_i^* = \frac{V\sqrt{a_i}}{\sum_j \sqrt{a_j}}, \qquad R_* = \frac{(\sum_j \sqrt{a_j})^2}{V}.$$

All coordinates are positive, so the equality point is admissible. The identity supplies a global optimum, including the boundary behavior: if any x_i tends to zero, its positive term drives the resistance to infinity.

A multiplier calculation gives $-a_i/x_i^2 + \lambda = 0$ and the same allocation, but it needs convexity or a lower-bound argument to prove global optimality. The inequality provides that argument directly. A certificate consists of the resource sum, feasibility and the matching lower bound. For rational a_i that are not rational squares, use interval square-root enclosures to certify a tolerance; do not present a rounded allocation as exactly attaining the optimum.

The model treats a_i and V as independent fixed inputs. If material properties depend on x_i or interface losses matter, the objective changes and this allocation need not remain optimal.

18.3 VI02. A balanced bridge and an unidentifiable edge

Problem. Two branches connect terminals s, t. Each branch consists of two unit conductances in series. Join their middle vertices by a conductance $b > 0$. Determine the effective conductance and whether terminal measurements identify b.

Both middle potentials equal $1/2$ by symmetry and satisfy conservation: the incoming and outgoing terminal currents each equal $1/2$, and the bridge carries zero current. Hence the total current under unit voltage is 1 and the conductance is exactly 1. The energy is four copies of $(1/2)^2$, plus $b \cdot 0^2$, also equal to 1.

Completing the energy square around this harmonic potential proves optimality for every $b > 0$. Therefore a single terminal conductance measurement cannot identify b: all positive bridge values give the same observation. Increasing a conductance has not violated Rayleigh monotonicity; the inequality is non-strict when the optimized voltage drop on that edge is zero.

A direct 2×2 linear solve reaches the same numerical answer. The structural proof identifies an entire invisible parameter family and explains the zero derivative. The result concerns this terminal configuration. An additional electrode or a changed excitation can make the bridge observable, so it is not an impossibility claim for every measurement protocol.

18.4 VI03. Exact one-edge sensitivity beyond a derivative

Problem. In a network with fixed terminal potentials, increase an interior-edge conductance by $\delta \geq 0$. Let its original optimized voltage drop be d, and let A be the free-vertex energy matrix and b its oriented incidence vector. Express the new conductance in terms of $r = b^T A^{-1} b$.

Write a perturbed potential as $u + e$, with e zero at the terminals. Original harmonicity removes the cross term. The new energy is

$$C + e^T A e + \delta(d + b^T e)^2.$$

Its stationarity equation is $(A + \delta b b^T)e = -\delta d b$. By 15.1, $e = -\delta d A^{-1} b / (1 + \delta r)$. Substitution gives

$$\boxed{C_{\text{new}} = C + \frac{\delta d^2}{1 + \delta r}.}$$

The correction is nonnegative. Its derivative at zero is d^2, and its second derivative is $-2rd^2/(1 + \delta r)^3$, agreeing with conductance concavity. If $d = 0$, no change occurs for any nonnegative δ.

One existing factorization and one solve determine r; every subsequent value of δ requires scalar arithmetic. A conventional independent refactorization for each value repeats the same rank-one structure. A practical comparison must include factorization setup and the number of queries. The formula applies to an interior edge; a nonzero Dirichlet endpoint alters the forcing convention and must be included explicitly.

18.5 VI04. Path and cut witnesses bound every network

Problem. Produce elementary lower and upper bounds on effective conductance from a single s–t path and a single cut separating s and t.

Send unit flow along the path, with signs respecting its traversal. It has the required divergence and flow energy $F = \sum_{e \in P} c_e^{-1}$. The primal–dual theorem gives $C \geq 1/F$. For a cut S containing s but not t, use potential zero on S and one outside. Its energy is $\sum_{e \in \partial S} c_e$, giving

$$\left(\sum_{e \in P} c_e^{-1} \right)^{-1} \leq C \leq \sum_{e \in \partial S} c_e.$$

Both bounds are exact on a single edge. On a long unbranched chain, the path lower bound is exact but a one-edge cut upper bound can be much larger. The gap describes the witness choice, not an error in the invariant.

Finding a minimum-resistance path and a minimum-capacity cut are different optimization tasks. Even their best values need not coincide with electrical conductance. The construction gives quick verified bounds without solving the Laplace system. It does not establish an equivalence between the maximum-flow value and electrical current under unit voltage.

18.6 VI05. Star elimination and the boundary response

Problem. An interior vertex is joined to three boundary vertices by conductances $a, b, c > 0$. Replace it by a boundary-only network preserving all Dirichlet energies.

For boundary values v_1, v_2, v_3, the center value minimizing $a(x - v_1)^2 + b(x - v_2)^2 + c(x - v_3)^2$ is $x = (av_1 + bv_2 + cv_3)/(a + b + c)$. Completing the square yields the minimum

$$\frac{ab(v_1 - v_2)^2 + ac(v_1 - v_3)^2 + bc(v_2 - v_3)^2}{a + b + c}.$$

Thus a triangle with edge conductances $ab/(a + b + c)$, $ac/(a + b + c)$ and $bc/(a + b + c)$ has the same boundary response. The identity holds for every boundary vector, so it proves equivalence of the full three-terminal Dirichlet-to-Neumann matrix, not just one experiment.

A direct nodal solve gives the same center value separately for each excitation. The eliminated form is useful for repeated boundary problems. It creates three edges from three original spokes and removes one degree of freedom; on larger graphs repeated elimination can create substantial fill. The local reduction does not imply global sparsity or globally linear complexity.

18.7 VI06. Uncertainty propagation on a three-layer chain

Problem. Let conductances satisfy $c_1 \in [1, 2]$, $c_2 \in [2, 4]$ and $c_3 \in [3, 6]$, independently. Find the sharp range of the effective conductance.

The sum of reciprocal conductances lies between $1/2+1/4+1/6 = 11/12$ and $1+1/2+1/3 = 11/6$. Consequently

$$C \in [6/11, 12/11].$$

Both endpoints are attained at permitted corners, and every intermediate value is attained by continuity along the line segment joining those corners. Hence the interval is the exact image range.

A Monte Carlo sample can illustrate this range but cannot establish that no admissible parameter vector lies outside it. Interval arithmetic combined with monotonicity supplies an exact certificate using a constant number of rational operations for this fixed three-layer instance. For a general m-layer box, the work is linear in m. Correlated measurements may narrow the range; the independent-box answer then remains valid but need not be sharp.

18.8 VI07. An attained stability bound for a cylinder density

Problem. On a cylinder of circumference ℓ and height h, use the density $\rho(y) = 1 + a \sin(2\pi y/h)$, $|a| < 1$. Calculate its extremal-length quotient for all vertical crossing curves and its normalized distance from the constant density.

Every crossing curve has weighted length at least $\int_0^h \rho(y)\,\mathrm{d}y = h$, by applying the chain rule to the increasing primitive of ρ. A straight vertical curve attains this length. The area energy is $\ell h(1 + a^2/2)$. Thus the quotient relative to the extremal value h/ℓ is

$$r = \frac{1}{1 + a^2/2}.$$

After dividing ρ by $\sqrt{\ell h(1 + a^2/2)}$, its squared distance in L^2 from $1/\sqrt{\ell h}$ is

$$2\left(1 - \frac{1}{\sqrt{1 + a^2/2}}\right) = 2(1 - \sqrt{r}).$$

This attains the near-extremizer estimate in Chapter 5. The example shows that the estimate cannot have a uniformly smaller right-hand side over this admissible class.

A numerical optimization of sampled densities might approximate the result. The analytic computation supplies the exact deficit, the exact distance and an equality family. The positivity condition on a is material: without it the primitive need not be monotone, so the argument about all crossing curves must be reconsidered.

18.9 VI08. A homological intersection inequality on a flat torus

Problem. On $\mathbb{C}/(\mathbb{Z}+\tau\mathbb{Z})$, $\tau = x+iy$, compare the extremal lengths of homology classes $\alpha = (m, n)$ and $\beta = (p, q)$ with their intersection number.

Let $v = m+n\tau$ and $w = p+q\tau$. The determinant of the two real vectors is $\operatorname{Im}(\overline{v}w) = y(mq-np)$. The elementary area inequality $|\det(v, w)| \le |v||w|$ follows from decomposing w into components parallel and perpendicular to v. The torus formula therefore gives

$$\mathsf{L}(\Gamma_\alpha)\mathsf{L}(\Gamma_\beta) = \frac{|v|^2|w|^2}{y^2} \ge (mq - np)^2.$$

For nonzero vectors, equality holds exactly when v and w are perpendicular. This includes the rectangular torus's horizontal and vertical generators.

Computing the two invariants independently obscures their common determinant. The structural inequality is exact for every modulus and every integer pair. It is a flat-torus result; extending it to arbitrary surfaces requires the appropriate intersection/extremal-length theory and is not obtained by copying the Euclidean determinant proof.

18.10 VI09. A minimax shape for two marked torus classes

Problem. Over $\tau = x + iy$, $y > 0$, minimize the larger extremal length of the fixed classes $(1, 0)$ and $(0, 1)$.

Their values are $1/y$ and $(x^2+y^2)/y$. Their product is $1+x^2/y^2 \ge 1$. Therefore their maximum is at least one. At $x = 0, y = 1$, both values equal one, attaining the lower bound. Equality in the maximum requires both values be at most one; the first gives $y \ge 1$, while the second gives $x^2 + y^2 \le y$, forcing $y = 1, x = 0$.

Thus the square torus is the unique optimum in this marked parameterization. Quotienting by changes of homology basis changes the meaning of uniqueness and must be specified separately. The minimax objective concerns two chosen classes, not all nonzero homology classes: the latter have unbounded lengths as their winding numbers grow.

A two-variable nonlinear optimizer can locate the same point numerically. The determinant lower bound and equality analysis certify a global optimum without a search, and identify the role of marking. No more general geometric optimality is inferred.

18.11 VI10. A rotated tensor invalidates a guessed linear solution

Problem. On $(0,\ell)\times(0,h)$, impose $u=0,1$ on the bottom and top and insulating sides. For a positive-definite constant tensor $A=\left(\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right)$ with $b\neq 0$, decide whether $u_0=y/h$ minimizes $\int \nabla u^T A\nabla u$.

Although u_0 solves the interior constant-coefficient equation, its side flux is $\pm b/h$, violating the insulating condition. This failure also has a direct variational witness. Let $w(x,y)=xy(h-y)$, which vanishes on the prescribed horizontal boundaries. Then

$$\int \nabla u_0^T A\nabla w = \frac{b}{h}\int_0^\ell\int_0^h y(h-y)\,\mathrm{d}y\,\mathrm{d}x = \frac{b\ell h^2}{6}\neq 0,$$

since the integral of w_y is zero. The energy of u_0+tw has a nonzero linear term and a positive quadratic term, so a sufficiently small t of the opposite sign lowers the energy. Hence u_0 is not the minimizer.

This demonstrates why an interior PDE residual alone does not certify the complete boundary-value problem. The conventional residual test must include the natural boundary condition. The diagonal-tensor formula $C=c\ell/h$ cannot be extended to a rotated tensor while retaining the same insulating boundary model.

18.12 VI11. The limiting behavior of a finite p-modulus chain

Problem. For positive fixed a_i, study $M_p=(\sum_{i=1}^m a_i^{-1/(p-1)})^{1-p}$ as $p\downarrow 1$ and the limit of $M_p^{1/p}$ as $p\to\infty$.

Put $a_*=\min_i a_i$ and $q=1/(p-1)$. Then $a_*^{-q}\le\sum a_i^{-q}\le m a_*^{-q}$. Raising to $-1/q$ reverses the inequalities, giving $a_* m^{-1/q}\le M_p\le a_*$. Letting $q\to\infty$ proves $M_p\to a_*$. The limiting linear programme assigns all density to a minimum-weight edge; multiple minimum weights cause nonuniqueness at $p=1$ even though the $p>1$ optimizer is unique.

As $p\to\infty$, each $a_i^{-1/(p-1)}\to 1$ and $(1-p)/p\to -1$, so $M_p^{1/p}\to 1/m$. This equals the smallest possible maximum coordinate under $\sum_i\rho_i\ge 1$, attained by $\rho_i=1/m$. The finite positive weights disappear from this normalized limiting objective, but not from finite-p corrections.

These limits follow from inequalities and continuity, avoiding large exponents in a numerically unstable direct formula. Log-sum-exp arithmetic can evaluate intermediate regimes, but its rounding error and underflow policy require separate implementation checks.

18.13 VI12. A certificate for canonical-weight positivity

Problem. An interval computation encloses a raw arc width by $a\in[a_-,a_+]$. Certify the buffered weight $w=\max(a-2,0)$ and decide which strict conclusions are justified.

The positive-part function is nondecreasing and one-Lipschitz. Therefore

$$w\in[\max(a_- -2,0),\max(a_+-2,0)].$$

If $a_->2$, positivity is proved. If $a_+\le 2$, the weight is exactly zero. If the interval straddles two, neither strict positivity nor exact zero follows. For $a\in[1999/1000,2001/1000]$ the valid answer is $w\in[0,1/1000]$.

Refinement can eventually decide the sign when a nonzero separation from two is known. At the threshold, an arbitrary real approximation oracle need not reveal equality after finitely

many queries. A program that classifies every straddling interval as zero is unsound. A program that keeps refining without a separation promise has no general termination guarantee for this decision.

The numerical task of enclosing the weight is well behaved even when the strict decision is not. This distinction separates Lipschitz approximation from exact comparison and is useful when constructing canonical diagrams from certified geometric inputs.

18.14 VI13. An exactly attained form-perturbation example

Problem. Give rational two-dimensional data attaining the sharp perturbation bound for $\varepsilon = 3/5$.

Use the affine constraint $x_1 = 1$, the original Euclidean energy, and

$$E = \begin{pmatrix} 9/25 & -12/25 \\ -12/25 & -9/25 \end{pmatrix}.$$

Its square is $(9/25)I$, so its eigenvalues are $\pm 3/5$. Thus $I + E$ is positive definite. The original minimizer is $u = (1, 0)$, while stationarity of the perturbed energy in x_2 gives $-12/25+(16/25)x_2 = 0$, so $v = (1, 3/4)$. The relative Euclidean error is $3/4$, exactly

$$\frac{\varepsilon}{\sqrt{1-\varepsilon^2}} = \frac{3/5}{4/5} = 3/4.$$

The original minimum is one; the new minimum is $(34/25) - (144/625)/(16/25) = 1$. Equality of the optimized energy values therefore does not imply equality of the minimizing vectors.

A generic norm-resolvent bound gives $\varepsilon/(1-\varepsilon) = 3/2$, twice the attained error. The sharp affine-geometric estimate is more informative for this output norm. All identities in this example are rational, providing a finite check of the equality case independently of floating-point plotting. It does not prove novelty of the general perturbation inequality.

18.15 VI14. Choosing precision from the sharp error factor

Problem. Suppose the relative symmetric-form perturbation is bounded by $\varepsilon = 2^{-b}$. Find the least integer b ensuring the relative minimizer error is at most 10^{-6} under the sharp theorem.

The condition is $\varepsilon/\sqrt{1-\varepsilon^2} \le 10^{-6}$. Squaring and rearranging gives

$$(10^{12}+1)2^{-2b} \le 1, \quad \text{or} \quad 2^{2b} \ge 10^{12}+1.$$

Since $2^{38} = 274877906944 < 10^{12}+1$ and $2^{40} = 1099511627776 > 10^{12}+1$, the least integer is $b = 20$. The integer comparisons provide an exact certificate; no rounded logarithm is needed.

This is precision selection from an assumed operator-form error bound. It is not a statement that twenty fraction bits in any implementation produce that form bound. Establishing the latter requires accumulated rounding, conditioning and coefficient-representation analysis. The distinction prevents an analytical precision estimate from becoming an unsupported hardware specification.

A conventional approach tests candidate bit counts numerically. The exact inequality identifies the threshold and its minimality within the model. Additional guard bits may be needed by a concrete arithmetic implementation and must be accounted for separately.

18.16 VI15. Complementary modes of the cylinder boundary operator

Problem. For one angular Fourier frequency $k > 0$ on a cylinder of height h, diagonalize the two-boundary Dirichlet-to-Neumann matrix and analyze a thin-cylinder limit.

The matrix from Chapter 12 is

$$k \begin{pmatrix} \coth(kh) & -\operatorname{csch}(kh) \\ -\operatorname{csch}(kh) & \coth(kh) \end{pmatrix}.$$

Its equal-boundary and opposite-boundary eigenvectors are $(1, 1)$ and $(1, -1)$. Using $\cosh z - 1 = 2\sinh^2(z/2)$ and $\sinh z = 2\sinh(z/2)\cosh(z/2)$ gives eigenvalues $k\tanh(kh/2)$ and $k\coth(kh/2)$.

As $h \downarrow 0$, the first is asymptotic to $k^2h/2$ and the second to $2/h$. The ratio grows like $4/(k^2h^2)$. A small geometric thickness therefore creates a strong separation between common and differential boundary modes. Scaling both modes by one typical eigenvalue can produce a poor preconditioner.

A direct eigenvalue calculation gives the values for each h. The exact modal decomposition identifies the parameter-dependent conditioning and a natural diagonal preconditioner. These conclusions concern one fixed frequency and the specified scalar cylinder problem. They do not establish uniform spectral bounds for rough or multiply connected boundaries without additional theory.

18.17 VI16. An information lower bound for analytic continuation

Problem. A forward map sends a Fourier coefficient a_k to $b_k = e^{-kh}a_k$. The observed b_k has absolute error at most δ. Establish a worst-case lower bound for recovering a_k.

Choose two admissible coefficients $a_k^{\pm} = \pm\delta e^{kh}$. Their exact forward data are $\pm\delta$, and both permit the same observation zero. Any deterministic reconstruction from that observation returns one number $\widehat{a}$. At least one of its distances from the two possible coefficients is at least half their separation, namely δe^{kh}. Thus every method has worst-case error at least that value on a class containing these two possibilities.

The direct inverse achieves the same bound for this single unconstrained coefficient: $\widehat{a} = e^{kh}\widehat{b}$, with error at most δe^{kh}. The lower and upper bounds match. No algorithm can remove the exponential instability without adding information, such as a bound on a_k, a prior regularity class, or more accurate measurements.

A small forward residual is compatible with a large inverse error. The example distinguishes conditioning from implementation quality and provides a benchmark in which an algorithmic acceleration cannot improve the information limit. For many modes, regularization trades resolution against an explicitly stated class-dependent error; that broader optimization is not settled by the scalar calculation alone.

18.18 VI17. The best uniform rank for a diagonal restriction

Problem. On ℓ^2, let $Te_j = r^{j-1}e_j$, $0 < r < 1$. Determine the smallest operator-norm error of a linear rank-at-most-N approximation.

Truncating to the first N coordinates has error r^N. For any rank-at-most-N operator A, its restriction to the $(N+1)$-dimensional span of $e_1, \ldots, e_{N+1}$ has a nonzero kernel vector. Choose

such a vector x with norm one. Then

$$\|(T-A)x\|^2 = \|Tx\|^2 = \sum_{j=1}^{N+1} r^{2j-2}|x_j|^2 \geq r^{2N}.$$

The lower bound matches truncation. This proves exact optimality among bounded linear rank-N competitors.

To guarantee error at most $\eta \in (0,1)$, choose the least N with $r^N \leq \eta$. If r and η have certified intervals, this inequality itself needs an enclosure or exact power comparison. An asymptotic formula for N does not settle an integer boundary case.

The argument does not apply to arbitrary nonlinear encodings with unrestricted decoding, nor does it charge for obtaining the input coefficients. Those are different information models. Its relevance is a mathematically exact storage-versus-operator-error contract.

18.19 VI18. Unitary scattering and perturbation of a reaction matrix

Problem. For finite Hermitian K, define $S(K) = (I+iK)(I-iK)^{-1}$. Prove a norm perturbation bound between two such matrices.

An eigenbasis gives $\|(I-iK)^{-1}\| \leq 1$. The identity $S(K) = 2(I-iK)^{-1} - I$ and the resolvent difference formula yield

$$S(K_1) - S(K_2) = 2i(I-iK_1)^{-1}(K_1-K_2)(I-iK_2)^{-1}.$$

Taking norms gives $\|S(K_1) - S(K_2)\| \leq 2\|K_1 - K_2\|$. For scalars at $K=0$, the derivative has magnitude two, so the global constant cannot be reduced uniformly. Hermiticity also ensures each $S(K)$ is unitary, as established in the physical-model chapter.

A generic inverse-matrix perturbation estimate may introduce a condition factor from each inverse; here the spectral structure bounds each by one. The result supplies robust propagation of a matrix uncertainty into a unitary scattering parameterization. It does not prove that a given empirical reaction matrix is Hermitian, that omitted channels are absent, or that a finite model captures a measured system. Those physical assumptions must be checked independently.

18.20 VI19. Conformal invariance of steady heat but not its clock

Problem. Scale a planar thermal specimen by $z \mapsto az$, $a > 0$, keeping material diffusivity D fixed. Compare a steady two-terminal Dirichlet conductance with transient relaxation times.

For a scalar isotropic steady model, change of variables cancels the gradient's a^{-2} with the area's a^2, so the dimensionless Dirichlet conductance is unchanged. In contrast, the diffusion equation $u_t = D\Delta u$ transforms to $u_t = (D/a^2)\Delta_{\text{old}} u$ in the original coordinates. Every Laplacian decay rate is divided by a^2, and corresponding relaxation times are multiplied by a^2.

Thus two specimens can have the same conformal modulus and steady conductance but different thermal time constants. In a fixed-depth physical model, geometric scaling of out-of-plane depth can introduce another factor in total conductance; it must be stated separately.

The conventional time-domain calculation solves both specimens. The scaling law determines the second transient from the first without repeating the solve, under the exact scaled boundary/initial data. An unchanged clock would produce a physically incorrect comparison even though the steady-state fields agree after coordinate transformation.

18.21 VI20. A Gaussian determinant with a missing zero mode

Problem. For an unconstrained connected graph Laplacian L, explain why $\int_{\mathbb{R}^n} \exp(-x^T Lx/2)\,\mathrm{d}x$ diverges, and obtain the finite expression after restricting to mean-zero vectors.

The constant vector lies in $\ker L$. Decompose $x = t\mathbf{1}/\sqrt{n} + y$, $y \perp \mathbf{1}$. The exponent is independent of t, so the integral includes $\int_{\mathbb{R}} \mathrm{d}t = \infty$. Connectedness implies all other eigenvalues $\lambda_2, \ldots, \lambda_n$ are positive. On the mean-zero subspace, orthogonal diagonalization and one-dimensional Gaussian integration give

$$Z_0 = (2\pi)^{(n-1)/2} \prod_{j=2}^{n} \lambda_j^{-1/2}.$$

The measure is the Euclidean measure induced on that subspace. Grounding a vertex produces a different coordinate convention; its determinant normalization cannot be substituted without the appropriate Jacobian.

A numerical log-determinant routine applied to singular L may return a nonfinite value or an arbitrary regularization-dependent result. Removing the declared gauge mode gives the correct finite model. Infinite-dimensional field determinants, ultraviolet limits and zeta regularization require additional constructions and are not justified by this finite Gaussian identity.

18.22 VI21. A phase transition between certificate regimes

Problem. A feasible potential has energy $C + e$ and a feasible unit flow has energy $C^{-1} + f$, with $e, f \geq 0$. Express the certified gap and compare improving the two witnesses.

The bracket width is

$$(C + e) - \frac{1}{C^{-1} + f} = e + \frac{C^2 f}{1 + Cf}.$$

Potential error contributes exactly e. For $Cf \ll 1$, the flow contribution is approximately $C^2 f$; for large f it saturates below C. Thus a witness-refinement strategy should account for the scale of C, not merely compare the numerical sizes of e and f.

The formula follows by rational rearrangement and has no neglected remainder. The first-order approximation can be bounded: $C^2 f - C^3 f^2 \leq C^2 f/(1 + Cf) \leq C^2 f$, because $1/(1 + x) \geq 1 - x$ for $x \geq 0$. Improving a witness is useful only after feasibility has been retained; a smaller energy from a nonconservative flow is not a stronger lower certificate.

This problem concerns verification effort. It does not assert that a particular iterative solver can achieve prescribed e, f at a fixed cost. That algorithmic relation must be measured or proved before choosing a computational stopping policy.

18.23 VI22. Resource-optimal repeated-query selection

Problem. Method A uses setup S_A and per-query cost q_A, while method B uses $S_B > S_A$ and $q_B < q_A$. Both meet the same output contract and certification costs are V_A, V_B per full batch. Determine the smallest batch size for which B is strictly cheaper.

The inequality is

$$S_B + mq_B + V_B < S_A + mq_A + V_A,$$

so $m > (S_B - S_A + V_B - V_A)/(q_A - q_B)$. The smallest nonnegative integer meeting this strict inequality is obtained by a floor plus one, unless the right-hand side is negative, in which case zero already qualifies. For a positive-query workload impose $m \geq 1$ separately.

For example, $S_A = 1, S_B = 9, q_A = 3, q_B = 1, V_A = 2, V_B = 4$ gives $m > 5$ and threshold six. Reporting only the per-query ratio three would conceal the first five batches where B has no cost advantage. If certification is per query, its cost must instead be included in each q.

The arithmetic is elementary, but the specification is consequential. Setup can include geometry construction, sparse analysis, factorization, code generation and transfer. Reusing an uncertified geometry or a changed input family without charging its rebuild violates the common contract. Real costs may vary with batch size; then the affine model itself requires revision.

18.24 VI23. A finite nonidentifiability result for model selection

Problem. Two candidate physical models induce the same probability distribution on all recorded observables. Under equal prior weights, prove that no decision rule based only on those observations has mean success probability above one half.

Let P be the common observation distribution and let $d(x) \in [0, 1]$ be the probability of choosing model one after observing x. Success under model one is $\int d \, \mathrm{d}P$; success under model two is $\int (1 - d) \, \mathrm{d}P$. Their equal-weight mean is exactly $1/2$. This covers deterministic rules as the special case $d \in \{0, 1\}$ and randomized rules without further assumptions.

The obstruction is absence of distinguishing information, not insufficient processor speed or a weakness of a particular classifier. A new observation that separates the distributions changes the problem. Priors unequal to one half permit a better default decision but do not create information in the data.

The result applies to any pair of statistical models satisfying the premise. It provides a precise caution for comparing speculative physical constructions through observables that they predict identically. It does not establish that any two particular quantum-gravity or multiverse models actually satisfy that premise; that remains a model-specific scientific question.

18.25 VI24. A robust current certificate under a changed orientation

Problem. Reverse an arbitrary subset of graph-edge orientations. Determine how to transform a potential/flow certificate and prove invariance of its bounds.

Let D be diagonal on edge space with entries ± 1, marking the reversals. Then $B_{\mathrm{new}} = DB$ and $j_{\mathrm{new}} = Dj$. The potential is unchanged. Since $D^T D = I$ and D commutes with the diagonal conductance matrix,

$$B_{\mathrm{new}}^T j_{\mathrm{new}} = B^T j, \quad (B_{\mathrm{new}} u)^T C (B_{\mathrm{new}} u) = (Bu)^T C (Bu), \quad j_{\mathrm{new}}^T C^{-1} j_{\mathrm{new}} = j^T C^{-1} j.$$

Boundary values, conservation and the primal–dual bracket are all preserved. Reversing B without reversing j generally breaks conservation, while squaring the currents can conceal the sign error in the energy. Therefore energy equality alone is an insufficient certificate test.

The transformation costs one sign operation per changed edge and does not require resolving the network. It is a metamorphic test for an implementation's orientation convention. It verifies this algebraic invariance for chosen instances; it does not replace the universal written argument or prove every parser operation correct.

Chapter 19

Quantitative network design and perturbation geometry

19.1 The mathematical design problem

The source's resistance and conductance laws concern fixed geometric or electrical data. Design introduces a second optimization: the medium itself varies under a resource constraint. Its admissible class must be stated before any optimality claim. This chapter treats a finite graph with fixed endpoints and a quadratic cost on the edge-length weights. It is a complete finite-dimensional extension of the energy model, not a theorem about arbitrary manufactured conductors.

Let $G = (V, E)$ be a finite graph with distinct vertices s, t. Give each edge a strictly positive length w_e. Assume an s–t path exists, and let d be the shortest weighted-path distance between the terminals. For a budget $B > 0$, choose nonnegative conductances c_e satisfying

$$\sum_{e \in E} w_e^2 c_e \le B.$$

Zero conductances are allowed and remove edges from the electrical network. Define the effective conductance by the minimum Dirichlet energy, with terminal potentials zero and one. If no positive-conductance terminal path remains, this minimum is zero. This convention preserves continuity on the nonnegative orthant and does not require every inactive vertex to remain connected.

19.2 A complete budget-optimality theorem

Theorem 19.1. [Z01] *Under the preceding assumptions,*

$$\max_{c \ge 0,\ \sum w_e^2 c_e \le B} C(c) = \frac{B}{d^2}.$$

A shortest path attains the maximum by assigning $c_e = B/(d w_e)$ to its edges and zero to all other edges.

Proof. For every vertex connected to s in the underlying graph, let $D(v)$ be its shortest-path distance from s. Define $u(v) = \min(D(v), d)/d$; set $u(v) = 1$ on components not connected to s. Then $u(s) = 0, u(t) = 1$. The triangle inequality for shortest-path distance gives $|D(v) - D(z)| \le w_e$ on $e = \{v, z\}$. Truncation is one-Lipschitz, so $|u(v) - u(z)| \le w_e/d$. Hence, for every feasible

design,

$$C(c) \le E_c(u) \le \frac{1}{d^2} \sum_e c_e w_e^2 \le \frac{B}{d^2}.$$

This is a universal upper bound over the entire design class.

Choose a shortest simple path P. Its length is $\sum_{e \in P} w_e = d$. Assign the stated conductances. They use exactly the budget because $\sum_{e \in P} w_e^2 B/(d w_e) = B$. The resistance of the resulting chain is $\sum_{e \in P} d w_e / B = d^2/B$, so its conductance is B/d^2. Inactive vertices can be assigned arbitrary minimizing potentials without changing the terminal conductance. Thus the upper bound is attained. □

The argument supplies an exact optimizer as well as its value. If the graph and lengths are rational, the shortest-path length and attaining conductances are rational, so an exact certificate can contain the path, a feasible distance-label potential and the resource sum. The result is recorded as a derivation here; a comprehensive literature review establishing research priority has not been completed.

19.3 A second proof through a current allocation

For a fixed unit flow j and strictly positive conductances on its support, Cauchy–Schwarz gives

$$\left(\sum_e w_e |j_e| \right)^2 \le \left(\sum_e \frac{j_e^2}{c_e} \right) \left(\sum_e w_e^2 c_e \right).$$

Therefore its dissipation is at least $(\sum w_e |j_e|)^2 / B$. Equality is attained by c_e proportional to $|j_e|/w_e$, with zero conductance wherever $j_e = 0$ if permitted. The minimum possible weighted absolute flow is d: the distance-label function from the first proof satisfies a terminal voltage difference one and edge increment bounded by w_e/d. Summation by parts yields

$$1 = \sum_e j_e (Bu)_e \le \frac{1}{d} \sum_e w_e |j_e|.$$

A unit flow on a shortest path gives equality. Minimizing flow dissipation and then taking the reciprocal again produces B/d^2.

The two proofs share finite energy duality but expose different witnesses. The first uses one potential to upper-bound every admissible design. The second identifies the optimal distribution of material for each flow and then reduces the remaining task to a shortest-path problem. They do not constitute two independent implementations or two proof-assistant verifications.

If conductances must be strictly positive on every original edge, the path design may be excluded. Assign arbitrarily small positive resources to off-path edges and allocate the remaining budget as above. The conductance approaches B/d^2 from below or equals it only in special cases. The theorem then gives a supremum, whose attainment needs separate analysis. Lower conductance bounds, resilience constraints, multiple terminal pairs and nonlinear material costs change the design problem and can destroy the single-path optimum.

19.4 Worked rational design

Consider three alternative terminal paths with lengths $(1, 2)$, $(2, 2)$ and (5), and budget $B = 12$. Their weighted lengths are $3, 4, 5$, so $d = 3$. On the shortest path assign conductances $12/(3 \cdot 1) = 4$ and $12/(3 \cdot 2) = 2$. The cost is $1^2 \cdot 4 + 2^2 \cdot 2 = 12$, the resistance is $1/4 + 1/2 = 3/4$, and the conductance is $4/3 = B/d^2$.

A feasible distance potential assigns values $0, 1/3, 1$ along the shortest path. On the second path its midpoint value is $2/3$ when using the global shortest-distance labels; the edge increments satisfy their required bounds $2/3$. The direct edge increment is one, below $5/3$. Thus the universal certificate includes every path option, even though the optimizer uses only one.

A black-box nonlinear optimization of all five edge conductances must discover a boundary solution and still needs a global upper certificate. The structural method computes a shortest path and explicit allocations. Under nonnegative weights, a conventional shortest-path algorithm suffices. Its exact operation count depends on the chosen graph representation and priority queue. This reduction has no implication that shortest-path routing optimizes latency, packet reliability or arbitrary transport networks.

19.5 First and second variations of fixed-network conductance

Return to a connected network with strictly positive conductances and fixed Dirichlet terminals. Let $u(c)$ be its harmonic potential, and write $d_e = (Bu)_e$. For a conductance perturbation $\dot{c}$, differentiate the interior stationarity equation. With B_0 restricted to free vertices and $A = B_0^T C B_0$, the free derivative satisfies

$$A\dot{u} = -B_0^T \operatorname{diag}(\dot{c}) Bu.$$

The objective derivative simplifies because the potential variation has zero boundary values and the original first variation vanishes:

$$\dot{C} = \sum_e \dot{c}_e d_e^2.$$

Differentiating once more along a straight conductance path gives

$$\ddot{C} = -2g^T A^{-1} g \le 0, \quad g = B_0^T \operatorname{diag}(\dot{c}) Bu.$$

To verify the last identity, differentiate $\sum \dot{c}_e d_e^2$ to obtain $2g^T\dot{u}$, then substitute $\dot{u} = -A^{-1}g$. All derivatives are justified locally because $A(c)$ remains nonsingular on the positive orthant and its inverse is a rational function of its entries.

The full Hessian is therefore negative semidefinite. Its entries are

$$\frac{\partial^2 C}{\partial c_e \partial c_f} = -2 d_e d_f \, b_e^T A^{-1} b_f,$$

where b_e is the free-vertex incidence column. This also bounds its rank by the number of free vertices. A direction may have zero second variation because it produces no interior forcing, even when some first derivative is nonzero. Uniform scaling is an example: $C(tc) = tC(c)$ is affine in t.

19.6 Logarithmic sensitivity is an energy distribution

Define $p_e = c_e d_e^2 / C$. Every p_e is nonnegative and $\sum_e p_e = 1$, since the optimized energy equals C. Thus

$$\frac{\partial \log C}{\partial \log c_e} = p_e.$$

The logarithmic sensitivity vector is a probability vector over the dissipated edge energies. This does not describe random edge failures; it is a deterministic derivative identity. It immediately gives a first-order bound $|\dot{\log C}| \le \max_e |\dot{\log c_e}|$.

A global finite-perturbation version does not need differentiation. If $e^{-\delta} c_e \le \widetilde{c}_e \le e^{\delta} c_e$ for every edge, monotonicity and homogeneity imply

$$e^{-\delta} C(c) \le C(\widetilde{c}) \le e^{\delta} C(c), \qquad |\log C(\widetilde{c}) - \log C(c)| \le \delta.$$

The constant one is sharp under uniform scaling. This is a complete uncertainty bound on positive networks. It does not say that the potential vector has the same stability; the sharp affine perturbation example demonstrates that optimizer sensitivity and optimized-value sensitivity differ.

19.7 A certificate for a proposed resource optimum

Suppose the objective is to maximize concave differentiable $C(c)$ under a linear resource constraint $a^T c \le B$ and lower bounds $c_e \ge \underline{c}_e > 0$. A feasible point c_* is globally optimal if there are $\lambda \ge 0$ and $\mu_e \ge 0$ such that

$$\nabla C(c_*) = \lambda a - \mu, \quad \lambda(a^T c_* - B) = 0, \quad \mu_e(c_{*,e} - \underline{c}_e) = 0.$$

Indeed, concavity gives $C(c) \le C(c_*) + \nabla C(c_*)^T (c - c_*)$. The multiplier expression is nonpositive for every feasible c by the constraints and complementary slackness. This proves sufficiency without assuming that a local numerical solver found the correct point.

For rational c_* the harmonic potential and gradient entries d_e^2 are rational, so exact feasibility and multiplier equalities form a finite global certificate. Existence of rational optimum multipliers for every rational resource problem is not asserted; the optimum itself can be irrational in more general models. An approximate multiplier witness needs a quantified residual bound rather than a boolean equality test.

19.8 Relation to continuum design

In a continuum conductivity problem the analogous derivative is an energy density, but turning the finite formula into a continuum theorem requires differentiability of the solution map, ellipticity bounds and the correct function space. Concentrating material onto a path can produce a singular coefficient outside the admissible continuum class. The finite graph's attained single-path design therefore cannot be transplanted unmodified into a positive, bounded continuum medium.

This distinction suggests a concrete subsequent research problem: fix an ellipticity floor, a material budget, a domain family and a discretization; establish convergence of the discrete optimal values and characterize whether optimal designs develop concentration. Such a theorem is not proved here. The present complete finite theorem provides a diagnostic instance and a source of upper/lower witnesses against which a more general numerical design procedure can be tested.

Chapter 20

Boundary operators and finite spectral reductions

20.1 Elimination as a variational operation

Partition a finite graph into boundary vertices B and interior vertices I, with every connected component meeting B. Its Laplacian has block form

$$L = \begin{pmatrix} L_{BB} & L_{BI} \\ L_{IB} & L_{II} \end{pmatrix}.$$

The interior block L_{II} is positive definite: extend a free interior vector by zero on B and use the graph-energy argument. For fixed boundary values v, the unique minimizing interior vector is $x = -L_{II}^{-1} L_{IB} v$. Substitution gives the boundary energy

$$E_{\min}(v) = v^T \Lambda v, \quad \Lambda = L_{BB} - L_{BI} L_{II}^{-1} L_{IB}.$$

The matrix Λ is symmetric positive semidefinite. It satisfies $\Lambda \mathbf{1} = 0$, since a constant boundary vector has the constant extension with zero energy. Differentiating the minimized energy gives twice the boundary current vector in the energy convention used here. The factor two is removed when differentiating the half-energy.

This construction is the finite Dirichlet-to-Neumann map. It retains the full boundary response while discarding interior coordinates. It does not retain every interior observable. Recovering interior potentials requires the extension operator as additional data.

20.2 Associativity of exact elimination

Proposition 20.1. [B01] *Eliminating two disjoint collections of interior variables successively yields the same final boundary quadratic form as eliminating their union at once.*

Proof. For every fixed boundary vector v, the joint interior energy is a strictly convex coercive quadratic in the two groups (x, y). It has a unique joint minimum. Taking $\min_x \min_y E(v, x, y)$ gives the same value as $\min_{x,y} E(v, x, y)$, because each inner minimum is attained and the feasible product space is unchanged. The two resulting symmetric quadratic forms agree for every v. Equality of their matrices follows by evaluation on e_i and $e_i + e_j$, which determines all diagonal and off-diagonal entries. □

A second verification can expand the block inverse formulas, but the variational proof avoids a long cancellation and exposes why the order is immaterial in exact arithmetic. Numerically,

elimination order strongly affects fill and rounding. Equality of the exact boundary operator does not imply equal implementation cost or bitwise output.

This is also a compositional verification rule. A subdomain solver can be replaced by its checked boundary form before being connected to other subdomains. The connection requires identical shared boundary variables and compatible energy units. Summing forms with mismatched normal orientations or duplicated boundary terms changes the model.

20.3 Energy-preserving assembly

Suppose subgraphs overlap only on declared interface vertices and their edge sets are disjoint. The total energy is the sum of their energies. First minimizing all strictly interior variables produces a sum of interface forms. Then minimizing the interface variables produces the same global boundary form by 20.1. This yields a finite domain-decomposition principle with an exact algebraic meaning.

Parallel subgraph elimination is possible because its independent minimizations use disjoint unknown sets once interfaces are fixed. The interface solve remains a global dependency. A claim that the entire algorithm is embarrassingly parallel would ignore this dependency. Its dimension, sparsity and condition number must appear in a cost model.

An exact interface certificate can retain each local form, its factorization or energy proof, the assembly map, and the final boundary witness. Checking only local residuals without the interface continuity and flux balance is insufficient. A set of individually correct subdomain solutions can fail to solve the global problem.

20.4 Harmonic extensions are contractions in the energy quotient

For a fixed boundary vector v, let Hv denote its harmonic extension. Any other extension is $Hv + w$, where w vanishes on the boundary. Orthogonality gives

$$E(Hv + w) = E(Hv) + E(w).$$

Thus H is the minimum-energy extension, and the boundary seminorm $\|v\|_\Lambda = (v^T \Lambda v)^{1/2}$ is exactly the energy norm of that extension. Quotienting out boundary constants makes it a genuine norm for a connected graph. A boundary vector can have small energy norm and large Euclidean amplitude along a nearly disconnected mode; the norms are not uniformly equivalent without a spectral lower bound.

If an approximate extension has interior residual r, then its energy error is $r^T L_{II}^{-1} r$. A certified lower eigenvalue bound a yields the upper bound $\|r\|_2^2/a$. Alternatively, an exact rational correction witness can be checked against $L_{II} e = r$. Either route connects the residual to an actual error. The raw residual size alone is not the certificate.

20.5 A two-terminal response determines one scalar

For two terminals, symmetry and zero row sums force

$$\Lambda = C \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Hence the complete two-terminal boundary map contains only one independent scalar. Many distinct interiors share it. The balanced bridge in VI02 gives a continuous family of examples.

An inverse problem seeking every interior edge from one such response is underdetermined before numerical noise is considered.

With more boundary vertices, there are at most $|B|(|B|-1)/2$ independent entries in a symmetric zero-row-sum matrix. This dimension count is only a necessary obstruction to parameter identification; it does not prove injectivity when the number of unknowns is smaller. Nonlinear parameter symmetries or electrical equivalences can still produce ambiguity. Establishing a full inverse-network theorem is a separate problem.

20.6 Generalized eigenvalues and a controlled mass matrix

A finite element or weighted-node model often has energy $u^T Au$ and mass $u^T Mu$, where M is positive definite. The generalized Rayleigh quotient is

$$\mathcal{R}(u) = \frac{u^T Au}{u^T Mu}.$$

Set $z = M^{1/2}u$ and $K = M^{-1/2}AM^{-1/2}$. Then $\mathcal{R}(u) = z^T Kz/\|z\|^2$. All ordinary symmetric spectral variational statements follow in these coordinates. The normalization by mass is indispensable: diagonalizing A alone computes a different spectrum when $M \neq I$.

For the least eigenvalue, the complete finite-dimensional proof is immediate from an orthogonal eigenbasis of K: the quotient is a weighted average of its eigenvalues with weights $|z_i|^2/\|z\|^2$. Its minimum is the smallest eigenvalue and is attained by its eigenvectors. Restricted trial spaces give upper bounds on this minimum, since minimization occurs over a subset. Lower bounds need an additional certificate, not another trial vector.

20.7 Residual inclusion for a symmetric eigenproblem

Proposition 20.2. [B02] *Let K be real symmetric, $\|z\| = 1$, and $r = Kz - \theta z$. At least one eigenvalue λ satisfies $|\lambda - \theta| \le \|r\|$.*

Proof. Expand $z = \sum_i a_i q_i$ in an orthonormal eigenbasis. Then $\|r\|^2 = \sum_i |a_i|^2(\lambda_i - \theta)^2$ and $\sum_i |a_i|^2 = 1$. If every spectral distance exceeded $\|r\|$, the weighted sum would exceed $\|r\|^2$, a contradiction. □

The theorem encloses some eigenvalue. It does not identify its index, multiplicity or eigenvector accuracy. To isolate the first eigenvalue or certify a gap one needs additional counting or separation information. Near repeated eigenvalues, subspace error is the appropriate object; individual eigenvector labels can change discontinuously.

For a generalized problem, first measure the residual in the transformed coordinates: $Kz - \theta z = M^{-1/2}(Au - \theta Mu)$. A small unweighted residual in the original variables may become large when M is ill-conditioned. The mass normalization is therefore part of the verification contract.

20.8 A finite heat-semigroup comparison

If $K \succeq 0$, the operator e^{-tK} has norm at most one for $t \ge 0$, because its eigenvalues are $e^{-t\lambda_i} \in [0, 1]$. For positive semidefinite symmetric K_1, K_2, differentiating $e^{-(t-s)K_1}e^{-sK_2}$ and integrating gives

$$e^{-tK_2} - e^{-tK_1} = \int_0^t e^{-(t-s)K_1}(K_1 - K_2)e^{-sK_2}\,\mathrm{d}s.$$

The contraction bound yields $\|e^{-tK_1} - e^{-tK_2}\| \leq t\|K_1 - K_2\|$. No commutation assumption is needed. This provides a complete model-level perturbation estimate for finite diffusion operators.

A conformal transformation of steady energy does not automatically preserve the mass matrix and therefore does not automatically preserve this semigroup. This is the operator counterpart of the thermal-clock calculation in VI19. A reduced-order model must state whether it preserves energy, mass, selected spectra or the full transient response; these are different requirements.

20.9 Infinite-dimensional extension requirements

The finite proofs identify the exact ingredients required for a continuum analogue: closed coercive forms, trace and extension operators, quotient constants, self-adjoint realization, compactness where discrete spectra are used, and semigroup domains. Those are substantial supporting obligations under the agreed undergraduate boundary. The finite matrix theorems do not discharge them.

A subsequent certified numerical programme can begin from exact manufactured forms and test whether a discretization preserves positivity, conservation and interface assembly. Convergence rates then require approximation and regularity results for the actual geometry. The manuscript includes no claim that every rough, nodal or infinite-type surface inherits these finite spectral conclusions without such a development.

Chapter 21

Constructive continuum comparisons and certified truncation

21.1 An exactly solvable variable medium

Let $a : [0, L] \to (0, \infty)$ be measurable with $0 < a_- \le a(x) \le a_+ < \infty$. Consider absolutely continuous potentials u with $u(0) = 0, u(L) = 1$ and finite energy $E(u) = \int_0^L a(x)u'(x)^2 \, \mathrm{d}x$. Put $R = \int_0^L a(x)^{-1} \, \mathrm{d}x$. Weighted Cauchy–Schwarz gives

$$1 = \left(\int_0^L u' \, \mathrm{d}x \right)^2 \le E(u)R.$$

The potential

$$u_*(x) = \frac{\int_0^x a(s)^{-1} \, \mathrm{d}s}{R}$$

is absolutely continuous, has the prescribed boundary values and attains $E(u_*) = 1/R$. Its flux $au_*' = 1/R$ is constant almost everywhere. Thus the continuum capacity is exactly $C = 1/R$.

For any competitor $u = u_* + e$, the cross term is $(2/R) \int e' = 0$. Hence $E(u) - C = \int ae'^2$. This proves uniqueness within the prescribed absolutely continuous class and supplies an exact energy-error identity. The assumptions that a is bounded above and below ensure the integrability and coercivity used here. Degenerate media with $a_- = 0$ require a separate admissible-space analysis.

A direct finite-difference solve approximates this problem after choosing a coefficient average. The integral formula computes the exact one-dimensional model once its reciprocal coefficient can be integrated or enclosed. A discontinuity does not require differentiating a; it is handled by the flux formulation.

21.2 Piecewise-linear restriction and its precise defect

Partition the interval into cells I_i of lengths h_i, and restrict u to functions linear on each cell. Let $\bar{a}_i = h_i^{-1} \int_{I_i} a$. The cell energy is

$$\frac{\bar{a}_i}{h_i}(u_i - u_{i-1})^2.$$

Consequently the finite-element capacity is the chain capacity

$$C_h = \left(\sum_i \frac{h_i}{\bar{a}_i} \right)^{-1}.$$

Since the trial class is smaller, $C_h \ge C$. The same inequality follows from Cauchy–Schwarz on each cell, $h_i^2 \le (\int a)(\int 1/a)$, which gives $h_i/\bar{a}_i \le \int_{I_i} 1/a$.

The discrepancy has an exact identity. Writing $\mu = \bar{a}_i$ and $\delta = a - \mu$,

$$\frac{1}{a} = \frac{1}{\mu} - \frac{\delta}{\mu^2} + \frac{\delta^2}{\mu^2 a}.$$

After integration, the linear term vanishes. Therefore

$$R - R_h = \sum_i \int_{I_i} \frac{(a - \bar{a}_i)^2}{\bar{a}_i^2 a}\,\mathrm{d}x, \quad R_h = \sum_i h_i/\bar{a}_i.$$

It follows that

$$0 \le C_h - C = \frac{R - R_h}{R R_h}, \qquad R - R_h \le \sum_i \frac{\int_{I_i}(a - \bar{a}_i)^2\,\mathrm{d}x}{\bar{a}_i^2 a_-}.$$

This is a fully specified error bound. It depends on within-cell variation, not only on the largest cell width. A discontinuity unresolved by the mesh can dominate the error even if a smooth-coefficient estimate predicts a higher convergence order.

21.3 A manufactured discontinuous example

Take $L = 1$, $a = 1$ on the first half and $a = 9$ on the second. The exact resistance is $1/2 + 1/18 = 5/9$, so $C = 9/5$. A single linear element uses $\bar{a} = 5$ and produces $C_h = 5$, a substantial overestimate. A mesh with a node at the interface represents the exact harmonic solution: its slopes are $9/5$ and $1/5$, and the potential at the interface is $9/10$.

Both exact fluxes equal $9/5$, although the gradients differ by a factor of nine. Averaging gradients across the interface without flux weighting would violate conservation. The example separates geometry resolution, coefficient averaging and solver residual. A linear system solved to machine precision can still represent the wrong discretization.

The exact continuum answer and the two finite-element results are rational. They furnish an independent diagnostic for mesh and assembly code. Their simplicity does not establish that arbitrary discontinuous two-dimensional media have equally explicit solutions.

21.4 A two-dimensional layered reduction

On a rectangle $(0, \ell) \times (0, h)$, take $A(y) = \operatorname{diag}(a_x(y), a_y(y))$ with both coefficients bounded and positive. Prescribe zero and one on the horizontal sides, and use natural insulating conditions on the vertical sides. For every admissible smooth potential, apply weighted Cauchy–Schwarz along each vertical segment:

$$\int_0^h a_y(y) u_y(x,y)^2\,\mathrm{d}y \ge \left(\int_0^h a_y(y)^{-1}\,\mathrm{d}y\right)^{-1}.$$

Integrating over x and adding the nonnegative horizontal energy gives

$$E(u) \ge \frac{\ell}{\int_0^h a_y^{-1}}.$$

The one-dimensional harmonic potential from the preceding construction, independent of x, attains this bound and satisfies the insulating sides. Thus this layered problem is exactly reducible to a one-dimensional integral.

The coefficient a_x does not influence this particular measurement. The conclusion is not that transverse conductivity is physically meaningless; the boundary experiment excites no transverse gradient. Changing electrodes or introducing off-diagonal tensor terms removes this special reduction, as VI10 shows.

21.5 An interface resistance as a singular-layer limit

Consider a thin layer of thickness ε and conductivity $a_\varepsilon = \varepsilon/r_0$, where $r_0 > 0$ is fixed. Its resistance is exactly $\varepsilon/a_\varepsilon = r_0$. Under current J, the potential jump is Jr_0 and the layer's energy is $J^2 r_0$. These quantities remain finite as the thickness tends to zero, while the gradient grows like $1/\varepsilon$.

The limiting interface law is therefore a jump condition, not continuity of potential with an ordinary bounded coefficient. Imposing continuity without the interface resistance loses a finite energy contribution. Conversely, declaring that the divergent gradient means divergent total dissipation is incorrect for this scaling.

The derivation is one-dimensional and exact. Extending it to curved interfaces or coupled transport requires a convergence theorem in a specified function space. It serves as a model for why singular or nodal geometric limits must retain the correct interface data instead of inheriting the smooth-domain equations verbatim.

21.6 Energy distortion under a controlled coordinate map

Theorem 21.1. [T01] *Let $F : U \to V$ be an orientation-preserving C^1 diffeomorphism between planar domains, respecting the two marked terminal boundary portions. Suppose its singular-value ratio satisfies $s_1/s_2 \le K$ everywhere, $s_1 \ge s_2 > 0$. Then corresponding Dirichlet energies satisfy*

$$K^{-1} E_V(v) \le E_U(v \circ F) \le K E_V(v).$$

For capacities defined by these matched admissible potential classes, $K^{-1}\operatorname{Cap}(V) \le \operatorname{Cap}(U) \le K \operatorname{Cap}(V)$.

Proof. The gradient of $v \circ F$ is $DF^T \nabla v \circ F$. Change of variables to V introduces $1/\det DF = 1/(s_1 s_2)$. In the principal singular directions, the resulting quadratic form has eigenvalues s_1/s_2 and s_2/s_1, both in $[1/K, K]$. Integrating gives the energy bounds. Pullback is a bijection between the matched admissible classes because F and its inverse preserve the designated boundary data. Taking infima proves the capacity inequalities. □

For a linear map taking a unit square to a rectangle with side lengths a, b, the vertical capacity is a/b. If $a \ge b$, $K = a/b$ and the upper factor is attained. Interchanging the marked direction attains the reciprocal factor. Thus neither uniform multiplicative constant can be improved over the admitted class.

For $K = 1$, the calculation reduces to conformal invariance. A general quasiconformal theorem for weakly differentiable maps requires additional chain-rule, area-formula and boundary-space theory. This C^1 result is complete within its stated class and does not claim that missing general development.

21.7 Complex dilatation and a computable bound

For a real-linear orientation-preserving map written $F(z) = az + b\bar{z}$, $|a| > |b|$, the maximum and minimum of $|F(z)|$ on $|z| = 1$ are $|a| + |b|$ and $|a| - |b|$. The triangle inequality gives the bounds,

and phases of z can align or oppose az and $b\bar{z}$, attaining both. Hence

$$K = \frac{|a| + |b|}{|a| - |b|} = \frac{1 + |\mu|}{1 - |\mu|}, \qquad \mu = b/a.$$

Applying this pointwise to a differentiable map with $|F_{\bar{z}}| \le k|F_z|$, $k < 1$, gives $K \le (1+k)/(1-k)$. A certified enclosure of the derivative ratio can therefore turn a proposed coordinate map into a capacity-distortion certificate under the global diffeomorphism and boundary hypotheses.

A local Jacobian bound alone does not prove global injectivity. A folded or incorrectly marked map may satisfy some local numerical conditions while failing to transfer the original problem. The certificate requires both geometric correspondence and the differential bound.

21.8 Fourier truncation with an explicit energy tail

Consider a cylinder of circumference ℓ and height h, with zero lower boundary value and upper boundary expansion in orthonormal cosine modes,

$$f(x) = \sum_{n=1}^{\infty} f_n \sqrt{2/\ell}\cos(2\pi n x/\ell), \qquad |f_n| \le Ae^{-\sigma n}, \quad \sigma > 0.$$

The harmonic solution for mode n is its boundary coefficient times $\sinh(k_n y)/\sinh(k_n h)$, where $k_n = 2\pi n/\ell$. Orthogonality and the exact mode calculation give the energy omitted after N modes as

$$E_{>N} = \sum_{n>N} k_n \coth(k_n h)|f_n|^2.$$

The exponential coefficient bound ensures uniform convergence of the series and of each fixed-order derivative, because its terms are bounded by a polynomial in n times $e^{-\sigma n}$. Thus termwise differentiation and boundary evaluation are justified by the uniform Cauchy criterion. Energy identities pass to the limit by the convergent nonnegative sum.

Since coth decreases on positive arguments, putting $q = e^{-2\sigma}$ gives

$$E_{>N} \le \frac{2\pi A^2}{\ell}\coth(2\pi h/\ell)\frac{q^{N+1}((N+1) - Nq)}{(1-q)^2}.$$

The scalar sum follows by differentiating the finite geometric tail or multiplying the series by $1 - q$ twice. This is a uniform truncation certificate over the declared coefficient class. The norm error is its square root. It does not bound unknown coefficients from a small sample of boundary data; obtaining the analytic decay parameters is part of the input contract.

21.9 A certified geometric-computational error split

Suppose a diffeomorphic approximate coordinate map has distortion bound K, and a solver in its target domain produces a capacity interval $[C_-, C_+]$. Then the original capacity lies in

$$[C_-/K,\ KC_+].$$

The interval combines geometry and numerical solution errors multiplicatively. Improving the numerical solve cannot remove the remaining factor K. Conversely, an exact conformal map gives $K = 1$ but may be expensive to construct or certify. A total-cost comparison must include that geometric work.

A truncation bound, a residual bound and a geometric distortion bound refer to different error mechanisms. They can be combined only after all are expressed for the same original observable. A small spectral tail for a function on the wrong target domain is not a certificate for the original boundary-value problem.

Chapter 22

Deficit bounds for finite extremal metrics

22.1 A convex admissible population

Let $\mathcal{A}$ be a nonempty closed convex subset of $[0,\infty)^m$, and let $a_i > 0$. For $p > 1$, define $F_p(x) = \sum_i a_i x_i^p$. Coercivity and compactness of finite-dimensional sublevel sets give a minimizer x_*. Strict convexity gives uniqueness. Path-incidence constraints $Nx \geq \mathbf{1}$ provide one example, but the results below apply to any such convex admissible set.

The derivative condition at the minimum is

$$\nabla F_p(x_*)^T(x - x_*) \geq 0 \qquad (x \in \mathcal{A}).$$

Indeed, the segment $x_* + t(x - x_*)$ is feasible for $0 \leq t \leq 1$, and the right derivative of its objective at zero cannot be negative. This condition is also sufficient by convexity. It does not require an interior minimizer or differentiable constraint functions.

A small objective deficit can now be converted into a distance estimate. The conversion is useful because a numerical solver may supply a feasible density and a lower dual bound without providing a trustworthy estimate of the density error.

22.2 A scalar remainder inequality

Lemma 22.1. [R01] *For $p \geq 2$ and $a, b \geq 0$,*

$$b^p - a^p - pa^{p-1}(b - a) \geq |b - a|^p.$$

For $1 < p < 2$ and $0 \leq a, b \leq M$, $M > 0$, the same left-hand side is at least $\frac{1}{2}p(p-1)M^{p-2}(b-a)^2$.

Proof. Taylor's integral identity gives the remainder as

$$p(p-1)(b-a)^2 \int_0^1 (1-t)(a + t(b-a))^{p-2}\,\mathrm{d}t.$$

For $p \geq 2$ and $b \geq a$, the integrand's power is at least $[t(b-a)]^{p-2}$. The resulting integral is $1/[p(p-1)]$, proving the first inequality. If $b < a$, put $d = a - b$. Then $a - td = b + (1-t)d \geq (1-t)d$. The remainder is at least $p(p-1)d^p \int_0^1 (1-t)^{p-1}\,\mathrm{d}t = (p-1)d^p \geq d^p$. For $1 < p < 2$, the negative exponent gives $(a + t(b-a))^{p-2} \geq M^{p-2}$ wherever the argument is positive. Integrating $1 - t$ gives the factor one half. Endpoint singularities are handled by replacing the argument by a positive approximation and passing to the limit in the scalar identity; the remainder itself is continuous for nonnegative a, b. □

The restriction to nonnegative scalars is material for the constant one in the $p \geq 2$ statement. Signed variants use different constants. The present domain matches nonnegative finite path densities, so no unnecessary stronger theorem is assumed.

22.3 A quantitative certificate for the optimizer

Theorem 22.2. [R02] *For $p \geq 2$ and $x \in \mathcal{A}$,*

$$\sum_i a_i |x_i - x_{*,i}|^p \leq F_p(x) - F_p(x_*).$$

For $1 < p < 2$, if both vectors have coordinates at most M, then

$$\sum_i a_i |x_i - x_{*,i}|^2 \leq \frac{2M^{2-p}}{p(p-1)} [F_p(x) - F_p(x_*)].$$

Proof. Apply the scalar inequality coordinatewise and sum with weights a_i. The sum of linear remainder terms is $\nabla F_p(x_*)^T (x - x_*)$, which is nonnegative by the variational condition. Dropping this nonnegative contribution gives each claimed bound. □

If a dual witness proves $L \leq F_p(x_*)$, then the computable quantity $F_p(x) - L$ can replace the unknown deficit on the right. A primal–dual objective bracket thus also certifies proximity to the unique optimizer. For $1 < p < 2$, a bound on M follows from an accepted upper energy U: every coordinate of both vectors is at most $(U/a_i)^{1/p}$. Taking the largest of these gives an explicit, possibly conservative M.

These inequalities quantify finite-dimensional rigidity. They do not prove existence of an extremal Borel density for an arbitrary continuum path family. The null-support counterexample in Chapter 3 is precisely why that extension needs additional hypotheses.

22.4 The quadratic case and active constraints

At $p = 2$, an exact expansion gives

$$F_2(x) - F_2(x_*) = \sum_i a_i (x_i - x_{*,i})^2 + 2 \sum_i a_i x_{*,i} (x_i - x_{*,i}).$$

The second sum is nonnegative on the feasible set. It vanishes when the feasible perturbation lies in the tangent subspace to an affine equality problem, recovering the orthogonal projection identity. Under inequality constraints it can be strictly positive. Therefore identifying the energy deficit with a squared distance without checking the linear term is generally incorrect.

For a single constraint $\sum_i x_i \geq 1$, the positive minimizer satisfies the equality constraint, and feasible perturbations with the same sum have zero cross term. Perturbations increasing the sum add an additional positive cost. A certificate format that records only $Nx \geq \mathbf{1}$ must allow both cases; it cannot silently assume every inequality is active.

22.5 A sharp quadratic example

Take $F(x, y) = x^2 + 4y^2$ and the feasible line $x + y = 1$. Differentiating in x gives $2x - 8(1 - x) = 0$, so $x_* = 4/5, y_* = 1/5$ and $F_* = 4/5$. Writing $x = 4/5 + t, y = 1/5 - t$ gives

$$F(x, y) - 4/5 = 5t^2.$$

The weighted squared distance is also $t^2 + 4t^2 = 5t^2$. If an exact lower certificate is $L = 4/5$ and a feasible output has energy at most $4/5 + 1/500$, then $|t| \le 1/50$. Both signs attain the bound with equality.

A small unweighted Euclidean distance would be $\sqrt{2}|t|$, whereas the objective's geometry uses $\sqrt{5}|t|$. Choosing the wrong norm changes numerical constants. The example supplies exact thresholds for regression tests of an optimization or certificate implementation.

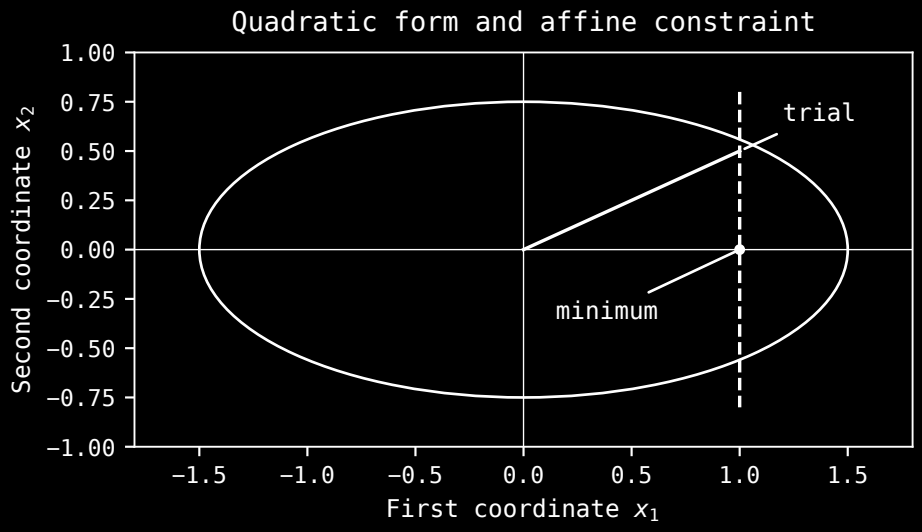


Figure 22.1: A positive quadratic form induces an ellipsoidal geometry. The dashed affine constraint has a minimum distinct from a generic trial point. The schematic depicts the role of the metric; it is not a plot of the numerical coefficients in the preceding example.

22.6 A degenerating pair of linear constraints

Let $a = (1, 0)$ and $b = (\cos\theta, \sin\theta)$, with $0 < \theta < \pi$. Minimize $\|x\|^2$ on the real plane subject to $a^T x \ge 1$ and $b^T x \ge 1$. This is a general convex constraint example; for angles with a negative component it is not a nonnegative path-incidence system.

The candidate $x_* = (a + b)/(1 + \cos\theta)$ satisfies both constraints at equality. The vector x_* is a positive linear combination of their normals, so for every feasible x, $x_*^T(x - x_*) \ge 0$. Completing the square proves optimality and uniqueness. Its squared norm is

$$\frac{2}{1 + \cos\theta} = \sec^2(\theta/2).$$

As $\theta \uparrow \pi$, the objective diverges and at $\theta = \pi$ the two constraints are incompatible. Thus a continuous family of constraint coefficients can approach infeasibility while its minimizer becomes unbounded. Stability theorems for varying path or geometric constraints require a uniform feasibility or separation hypothesis; positivity of the energy Hessian alone does not provide it.

22.7 Turning a computed deficit into an explicit stopping rule

For a finite $p \ge 2$ problem, suppose the desired weighted ℓ^p optimizer error is η. Maintain an exactly feasible primal point x and a verified dual lower bound L. Stop when $F_p(x) - L \le \eta^p$. The theorem proves the desired norm bound. The numerical optimization algorithm may be arbitrary, provided its outputs are checked for the required feasibility and energy inequalities.

For $p = 2$ and rational data, all expressions can be checked exactly. For noninteger p, evaluating x_i^p requires validated transcendental arithmetic or a separate exact algebraic representation in special rational-power cases. An ordinary floating-point power routine does not provide a certified lower or upper energy by itself. A monotone sequence of approximate gaps is not evidence that every computed gap is valid.

A practical implementation can terminate early with an "unresolved" result when the certificate generation budget or arithmetic precision is exhausted. That behavior is sound if it never labels the unresolved result accepted. The present graph checker implements the quadratic finite-witness case; the general p-optimizer stopping pipeline is a written specification and remains unimplemented.

Chapter 23

Exceptional curve families and attainment of the variational problem

23.1 The distinction resolved by the extension

The literal counterexample in Proposition 3.3 concerns constraints imposed on every curve. An area-null change of a density can alter those constraints. A useful completion of the theory therefore requires a definition of negligible *curve families*, together with a proof that the resulting convention preserves the variational value. This chapter supplies that construction for exponent two. It does not reinstate the false unrestricted smooth-density assertion.

Let $(U, \mathrm{d}A)$ be a surface with its Borel area measure. Curves in this chapter are locally rectifiable parametrized curves, and integrals are taken with arclength and multiplicity. A family $\mathcal{E}$ is *modulus-null* when its width is zero. A Borel density ρ is weakly admissible for Γ when the inequalities $\int_\gamma \rho \,\mathrm{d}s \geq 1$ fail only on a modulus-null subfamily. We work in $L^2(U)$ only after proving independence from the area-almost-everywhere representative.

The nonnegative integration facts used below are monotone convergence and countable additivity of integrals of nonnegative functions. These follow directly from the definition of the Lebesgue integral as the supremum of integrals of nonnegative simple functions: for $f_n \uparrow f$, each simple $0 \leq s < f$ is, after multiplication by any fixed $\theta < 1$, eventually dominated on an increasing sequence of measurable sets by f_n; continuity from below of measure gives $\int s \leq \lim \int f_n$. Taking the supremum over s proves monotone convergence. Apply it to partial sums to obtain the sum rule. These statements concern the usual countably additive area measure; its construction is part of the declared measure-theoretic foundation, not a new numerical test.

23.2 A zero-width criterion

Lemma 23.1. [WO1] *A curve family $\mathcal{E}$ has width zero if and only if there is a nonnegative $g \in L^2(U)$ such that $\int_\gamma g \,\mathrm{d}s = \infty$ for every $\gamma \in \mathcal{E}$.*

Proof. Given g, the density g/k is admissible for every positive integer k, and its energy tends to zero. Conversely, if the width is zero, choose admissible g_j with $\|g_j\|_2 \leq 2^{-j}$. The partial sums have L^2 norm at most one, by the triangle inequality. Monotone convergence applied to their squares gives $g = \sum_j g_j \in L^2(U)$. Nonnegative interchange of summation and the arclength integral gives $\int_\gamma g = \sum_j \int_\gamma g_j \geq \sum_j 1 = \infty$. □

The same argument shows that a countable union of modulus-null families is modulus-null. Indeed, choose witnesses g_j from the lemma and rescale each to norm at most 2^{-j}. Their sum

has finite norm and infinite integral on every curve in the union. Zero-width status is monotone under passage to a subfamily.

If N has area zero, the family of curves spending positive arclength in N is modulus-null. The function $g = \infty \mathbf{1}_N$ has zero area integral and infinite integral on every such curve. Thus altering a density on N changes its line integral only on a modulus-null family. Weak admissibility is consequently a well-defined property of the area equivalence class. A countable collection of changes of representative can be handled simultaneously by the preceding union argument.

23.3 Strong and weak constraint infima agree

Proposition 23.2. [W02] *The infimum of $\|\rho\|_2^2$ over weakly admissible densities equals the width defined using all curve constraints.*

Proof. Every strongly admissible density is weakly admissible, so the weak infimum is no larger. Let ρ be weakly admissible and let $\mathcal{E}$ contain its violations. For every $\epsilon > 0$ choose a strongly admissible density σ for $\mathcal{E}$ with $\|\sigma\|_2 < \epsilon$. Then $\rho + \sigma$ is admissible for all of Γ: on $\Gamma \setminus \mathcal{E}$ the first term suffices, and on $\mathcal{E}$ the second term suffices. Its norm is at most $\|\rho\|_2 + \epsilon$. Letting ϵ decrease to zero proves the reverse inequality. Infinite-infimum cases follow by the same comparison. □

Equality of infima does not assert that the strongly admissible population is closed or that a strong minimizer exists. The construction supplies arbitrarily small correction costs without requiring a single correction of zero cost. This quantifier distinction accounts for the role of exceptional families.

23.4 Convergence on almost every curve

Lemma 23.3 (Summable-subsequence principle)**.** [W03] *If $f_n \to f$ in $L^2(U)$, there is a subsequence f_{n_j} such that*

$$\int_\gamma |f_{n_j} - f| \, \mathrm{d}s \longrightarrow 0$$

for every curve outside one modulus-null family.

Proof. Choose n_j with $\|f_{n_j} - f\|_2 \le 2^{-j}$. Then $g = \sum_j |f_{n_j} - f|$ belongs to L^2 by the partial-sum argument used in Lemma 23.1. The family on which $\int_\gamma g = \infty$ is modulus-null, since g/k is admissible for it. On any remaining curve the nonnegative series $\sum_j \int_\gamma |f_{n_j} - f|$ is finite. Its terms therefore tend to zero. □

No uniform convergence over all curves follows. The original null-segment example illustrates why the exceptional set cannot generally be discarded. The statement is a functional-analytic convergence theorem; it does not assign an algorithmic rate to an arbitrary L^2-convergent input sequence. The relevant modulus framework and extremality criteria are treated in the primary literature [4].

23.5 Closedness, existence and uniqueness

Theorem 23.4 (Weak extremal density)**.** [W04] *If the width of Γ is finite, the nonnegative weakly admissible classes form a nonempty closed convex subset K_Γ of $L^2(U)$. There is a unique minimum-norm class ρ_*. It has energy $\mathsf{W}(\Gamma)$.*

Proof. Finiteness of the infimum implies that at least one finite-energy admissible density exists. Convexity follows by taking the union of the two exceptional families for two candidate densities; outside that union, linearity of each path integral preserves the unit lower bound.

For closedness, take $\rho_n \in K_\Gamma$ converging to ρ in L^2. The negative part satisfies $\|\rho_-\|_2 \le \|\rho - \rho_n\|_2$, so $\rho \ge 0$ almost everywhere. Use a nonnegative representative and apply Lemma 23.3. Outside the union of its exceptional family and the exceptional families of the selected ρ_n, one has

$$\int_\gamma \rho \ge \int_\gamma \rho_{n_j} - \int_\gamma |\rho_{n_j} - \rho| \ge 1 - o(1).$$

Hence ρ is weakly admissible.

Let $d = \inf_{x \in K_\Gamma} \|x\|_2$ and choose a minimizing sequence. Convexity ensures that its midpoints stay in K_Γ. The parallelogram identity, with the same calculation as in Theorem 2.3, makes the sequence Cauchy. Completeness and closedness give a minimizer. Strict convexity of the squared norm gives uniqueness. Proposition 23.2 identifies its energy with the original strong-constraint infimum. □

For completeness, the completeness of L^2 itself can be obtained by the same summable-subsequence construction: choose a subsequence of a Cauchy sequence with summable successive L^2 differences. The sum of their absolute values is finite almost everywhere and belongs to L^2. Pointwise summation defines an L^2 limit, the tail norm is bounded by the tail sum of norms, and the full Cauchy sequence has the same limit.

23.6 Quantitative attainment and an exact equality family

Theorem 23.5 (Deficit controls the density)**.** [W05] *For every $\rho \in K_\Gamma$,*

$$\|\rho - \rho_*\|_2^2 \le \|\rho\|_2^2 - \mathsf{W}(\Gamma).$$

The coefficient one is sharp, including for the full connecting family of a rectangle.

Proof. The segment $\rho_* + t(\rho - \rho_*)$ is feasible for $0 \le t \le 1$. Its squared norm has a minimum at $t = 0$, so $\langle \rho_*, \rho - \rho_* \rangle \ge 0$. Expanding $\|\rho\|_2^2$ gives the inequality.

For sharpness take $U = (0,1)^2$ and all curves joining its bottom to its top. The minimum density is $\rho_* = 1$. For $|a| < 1$ set $\rho(x,y) = 1 + a\cos(2\pi y)$. Its primitive in y increases from zero to one. On each connecting curve the integral of $\rho\,\mathrm{d}s$ is at least the total variation of this primitive and therefore at least one. Thus ρ is strongly admissible. The cosine has zero mean, so the energy deficit is exactly $\|\rho - 1\|_2^2 = a^2/2$. □

A certified upper bound U on the energy of a feasible density and a certified lower bound L on the minimum give the computable stopping condition $\|\rho - \rho_*\|_2 \le \sqrt{U - L}$. The result assumes exact feasibility or an independently bounded feasibility repair. A small objective value alone does not discharge the path constraints.

23.7 A valid smooth-density conclusion for regular condensers

Suppose a finite-area regular condenser has an energy-minimizing harmonic potential u, with continuous terminal limits zero and one, finite energy, and the modulus–capacity identity established for its complete connecting family. The density $|\nabla u|$ is admissible, although it need not be smooth at a critical point. For $\epsilon > 0$ define

$$\rho_\epsilon = \sqrt{|\nabla u|^2 + \epsilon^2}.$$

It is smooth and positive in the interior, dominates $|\nabla u|$, and has energy $\mathcal{E}(u) + \epsilon^2 \operatorname{area}(U)$. Consequently the infimum over smooth admissible densities equals the capacity. Reciprocity gives the corresponding extremal-length statement. This proof uses pointwise domination of every connecting-path integral; it does not attempt to infer trace convergence from area-norm approximation.

On an infinite-area surface a positive smooth $g \in L^2$ can replace the constant one, using $\sqrt{|\nabla u|^2 + \epsilon^2 g^2}$. Existence of such a g follows from a countable relatively compact covering and a smooth partition of unity with sufficiently small positive coefficients. The condenser hypotheses and existence of its harmonic potential remain explicit. This is a specified repair of Exercise 6.1, while Proposition 3.3 retains the source's unrestricted counterexample.

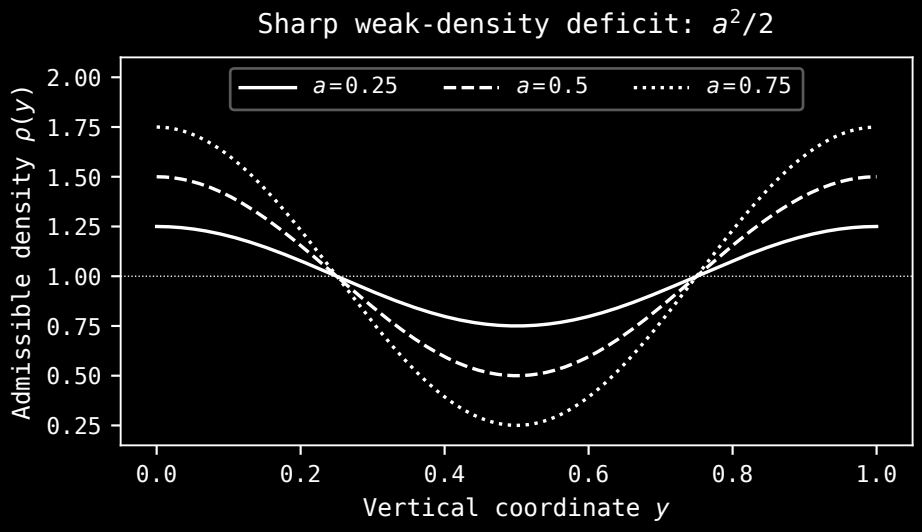


Figure 23.1: The exactly attained quadratic energy deficit of the feasible density $1 + a\cos(2\pi y)$. Every parameter in the stated range satisfies the same path constraints.

Chapter 24

Explicit geometric estimates for ring domains

24.1 Complementary continua and the model convention

Let $A \subset \widehat{\mathbb{C}}$ be a conformal annulus with complementary continua K, Q. In the planar statements K is compact and $\infty \in Q$; assume neither continuum is a singleton. Set $D = \operatorname{diam} K > 0$, $d = \operatorname{dist}(K, Q) > 0$ in the Euclidean plane, and $\delta = d/D$. The horizontal family consists of closed rectifiable curves with one nonzero winding about K and zero winding about finite points of Q. The width of its complete family is $m = \operatorname{mod} A$ by the cylinder calculation and conformal invariance. Thus the estimates below do not need smooth boundaries.

A rectifiable closed curve of nonzero winding about a point contains that point in its convex hull. To verify this, separate a point outside the convex hull from the curve by a line; the curve then lies in a half-plane on which a continuous argument about the point exists, and has winding zero. The diameter of the convex hull equals the diameter of the original compact set: the distance between two convex combinations is bounded by the largest pairwise distance. Finally, the length of a closed curve is at least twice its diameter, because the two subarcs between a maximally separated pair each have at least that distance as length. These elementary facts also hold with traversal multiplicity.

24.2 A complete tube proof of the linear upper bound

Proposition 24.1. [RG01] *With the preceding hypotheses,*

$$\operatorname{mod} A \le 2 + 2\pi + 4\delta.$$

In particular, a ring with two nondegenerate complementary continua has finite modulus.

Proof. Choose $a \in K, c \in Q$ with $|a - c| = d$. Compactness of K and closedness of the finite part of Q give such a minimizing pair. There is $b \in K$ with $s = |a - b| \ge D/2$: apply the triangle inequality to a diameter-realizing pair in K. Rescale by s and translate a to zero. Then $|b| = 1$ and $d' = |c| \le 2\delta$.

Let $E = [0, b] \cup [0, c]$ and $T = \{z : \operatorname{dist}(z, E) < 1\}$. Every horizontal curve meets $[0, c]$, since otherwise its winding number would be constant along that segment. If the curve exits T, its subarc from an intersection with E to its first exit has length at least one inside T. If it never exits, its full length is at least two because its convex hull contains $0, b$. Thus $\mathbf{1}_T$, restricted to A, is width-admissible for the horizontal family.

The unit neighborhood of a segment of length s has area $2s + \pi$. Subadditivity gives $\operatorname{area} T \le 2 + 2d' + 2\pi \le 2 + 4\delta + 2\pi$. The infimum defining width is at most this admissible energy. For spherical complementary continua, place a point of one continuum at infinity by a Möbius transformation and apply the same finite bound to the other compact continuum. Nondegeneracy is preserved. □

This completes the geometric step in source Lemma 6.10 without assuming existence of shortest horizontal curves or reducing a winding curve to a Jordan curve. It also supplies the finite-modulus conclusion of Corollary 6.11. The value $2 + 2\pi$ is an explicit sufficient constant, not an optimal universal constant.

24.3 A logarithmic upper estimate

Theorem 24.2. [RG02] *Under the same planar hypotheses,*

$$m \le 77 \log(3 + 4\delta).$$

Proof. Use the normalization from Proposition 24.1, so $a = 0$, $|b| = 1$ and $|c| = d' \le 2\delta$. Put $R = 2(1 + d')$, $\lambda = \log(4/3)$, and

$$\rho(z) = \frac{\mathbf{1}_{\{|z|<R\}}}{\lambda(1 + |z|)}.$$

Let γ be a horizontal curve. It meets $[0, c]$ at a point of radius at most d'. If it reaches radius R, the radial variation on an initial subarc gives unnormalized weighted length at least $\log((1 + R)/(1 + d')) > \log 2 > \lambda$.

Otherwise the curve stays in the disk of radius R. If its minimum radius is at least $1/2$, its winding number gives total angular variation at least 2π, and its weighted length is at least $2\pi/3 > \lambda$. If its minimum radius is below $1/2$, its maximum radius is at least one, since its convex hull contains b. An intervening subarc has weighted length at least $\int_{1/2}^{1} (1 + r)^{-1} \, \mathrm{d}r = \lambda$. All cases prove admissibility.

Polar integration gives

$$\int \rho^2 \, \mathrm{d}A = \frac{2\pi}{\lambda^2} \left(\log(1 + R) + \frac{1}{1 + R} - 1 \right) \le \frac{2\pi}{\lambda^2} \log(3 + 2d').$$

The identity $\log(4/3) = 2 \int_0^{1/7} (1 - t^2)^{-1} \, \mathrm{d}t$ gives $\lambda \ge 2/7$. Together with $\pi < 22/7$ this yields $2\pi/\lambda^2 < 77$. The latter elementary bound on π can be verified from the positive integral $\int_0^1 x^4(1 - x)^4/(1 + x^2) \, \mathrm{d}x = 22/7 - \pi$. Substituting $d' \le 2\delta$ proves the claim. □

The proof produces a witness with a computable energy. It provides no matching lower bound for the constant 77. Its purpose is a uniform geometric estimate with all scales explicit. Related ring-domain and metric-measure estimates are developed in [16]; the present argument is confined to the specified planar geometry.

24.4 An inscribed annulus with explicit thresholds

Corollary 24.3. [RG03] *If $m \ge 8192$, then A contains an essential round annulus of modulus at least $\log m$. If $A \cup K \subset D_R$, then, for all $m \ge 0$,*

$$\operatorname{diam} K \le 12R \exp(-m/77).$$

Proof. From Theorem 24.2, $\delta \ge (e^{m/77}-3)/4$. If $m \ge 77\log 6$, this is at least $e^{m/77}/8$. For any $a \in K$ the disk of radius D about a contains K, whereas its open disk of radius d avoids Q. Thus $\{D < |z-a| < d\} \subset A$ is essential when $d > D$, and its modulus is

$$\frac{1}{2\pi}\log\delta \ge \frac{m}{154\pi} - \frac{\log 8}{2\pi} > \frac{m}{484} - \frac{1}{2}.$$

At $m = 8192$ the last expression exceeds $\log m$, using $\log 2 < 7/10$; the derivative of $m/484 - 1/2 - \log m$ is positive for $m > 484$. This proves the first conclusion. Since an absolute constant in Lemma 6.10 may be enlarged, 8192 is also an explicit common sufficient threshold for the corresponding source corollary.

For the diameter statement, $d \le R$ because the complement contains $\mathbb{C} \setminus D_R$. At $m \ge 77\log 6$ the previous estimate gives $D \le 8Re^{-m/77}$. At smaller m use $D \le 2R \le 12Re^{-m/77}$. This proves the single formula in both ranges. □

The factor R is necessary. The literal Euclidean assertion with absolute diameter scale in Exercise 6.13(i) is refuted by similarity scaling in Chapter 5. The corollary is its scale-covariant replacement, not an unannounced alteration of the source. The exponential estimate is valid but deliberately nonsharp.

24.5 A quantitative two-sided pinching statement

Use chordal distance on the unit sphere embedded in $\mathbb{R}^3$, and intrinsic spherical arclength and area for densities. Write

$$s = \min(\operatorname{diam} K, \operatorname{diam} Q) > 0, \qquad d = \operatorname{dist}(K,Q) > 0, \qquad \delta = d/s.$$

Theorem 24.4. [RG04] *For the connecting-family modulus m,*

$$m \ge \frac{\delta^2}{\pi(1+\delta)^2}.$$

If $0 < \delta < 1/4$, then additionally

$$m \le \frac{2\pi}{\log(1/(4\delta))}.$$

Proof. Choose the smaller continuum, say K, and a point $a \in K$. The function $z \mapsto \operatorname{dist}(z,K)$ in chordal distance is 1-Lipschitz with respect to spherical arclength. Hence the density $1/d$ on its d-neighborhood outside K, extended by zero elsewhere, has integral at least one on each connecting curve. This neighborhood lies in the chordal cap of radius $s+d$ about a. A chordal cap of radius $r \le 2$ has area πr^2, as follows by writing $r^2 = 2(1-\cos\theta)$; for $r > 2$ the total area 4π is still at most πr^2. Thus the connecting width is at most $\pi(s+d)^2/d^2$. Reciprocity proves the lower bound.

For the upper bound, take a closest pair $a \in K, c \in Q$. There is a point of K at distance at least $\operatorname{diam} K/2 \ge s/2$ from a. There is a point of Q at distance from a at least $\operatorname{diam} Q/2 - d > s/4$. Connectedness therefore ensures that each spherical circle of chordal radius $r \in (d, s/4)$ about a meets both continua. On this circle an arc with endpoints in different continua and interior avoiding both can be chosen: start at a point of one closed set, take the first contact with the other in either orientation, and then take the last preceding contact with the first. Its interior lies in A.

Consequently every connecting-admissible density has integral at least one on the whole circle. Its circumference is $2\pi r\sqrt{1-r^2/4}$, so Cauchy–Schwarz gives a circle energy of at least the

reciprocal of that circumference. In polar area integration the radial Jacobian contributes another factor $(1 - r^2/4)^{-1/2}$. The total energy is therefore at least

$$\int_d^{s/4} \frac{\mathrm{d}r}{2\pi r(1 - r^2/4)} \geq \frac{1}{2\pi} \log \frac{s}{4d}.$$

The infimum over admissible densities has the same lower bound. Taking its reciprocal proves the result. □

These explicit functions establish the quantitative equivalence in source Proposition 6.14 for nondegenerate complementary continua. For example, if $m \geq \epsilon > 0$, then $\delta \geq \frac{1}{4} e^{-2\pi/\epsilon}$. Conversely, for $m < 1/\pi$ the first inequality yields $\delta \leq \sqrt{\pi m}/(1 - \sqrt{\pi m})$. Singleton components require a different extended-value statement and are excluded from this ratio.

24.6 Real-symmetric intervals

Corollary 24.5. [RG05] *Let $L = [a, b] \Subset (c, d) = \operatorname{int} I$ and $\delta = \min(a - c, d - b)/(b - a)$. For the slit ring $A(I, L)$,*

$$\operatorname{mod} A(I, L) \geq \frac{\delta^2}{\pi(1 + \delta)^2}.$$

For $\delta < 1/2$,

$$\operatorname{mod} A(I, L) \leq \frac{2\pi}{\log(1/(2\delta))}.$$

Proof. The planar neighborhood argument preceding (5.2) proves the first inequality with inner continuum L. For the second, take a_0 to be the endpoint of L nearest the outer slit. The inner interval reaches Euclidean distance $b - a$ from a_0; the connected outer continuum reaches infinity. In particular every circle with radius between $\delta(b - a)$ and $(b - a)/2$ meets both continua. Integrating the circle Cauchy–Schwarz bound $1/(2\pi r)$ as in Theorem 24.4 proves the asserted upper bound. □

Thus a lower bound on the real buffer gives an explicit lower modulus bound, and a lower modulus bound ϵ gives $\delta \geq \frac{1}{2} e^{-2\pi/\epsilon}$. This proves the quantitative equivalence in Lemma 6.17 without relying on a guessed outer circle that may cross a slit.

As literally printed, Corollary 6.18 concerns a conformal embedding of the whole complex plane. Such a map is affine. One proof uses the previously established Koebe growth bound: normalize f on $D(p, 2R)$; on $|z - p| = R$ it gives $|f(z) - f(p)| \leq 4R|f'(p)|$. Cauchy's derivative estimate then gives $|f''(p)| \leq 8|f'(p)|/R$ for every R, hence $f'' = 0$. A plane embedding carrying a real interval diffeomorphically to a real interval has a nonzero real affine slope. Length ratios of both complementary real intervals are preserved, including reversal of their order. Accordingly its real buffer is exactly unchanged. This also supplies a Koebe-based derivation of Exercise 6.19 under the literal global domain specification. Local-neighborhood variants require their own neighborhood and distortion hypotheses; they are not inferred from this global statement.

24.7 Divergence of nested annuli

Theorem 24.6. [RG06] *Let disjoint conformal annuli form a decreasing essential nest around nonempty compact inner continua K_n, with K_0 compact. If the sum of their moduli diverges, then $\bigcap_n K_n$ is one point. Any nonempty set surrounded by that same nest is contained in that point.*

Proof. The intersection K is nonempty by compactness and the finite intersection property. It is connected: if two disjoint open sets separated K into nonempty compact pieces, a sufficiently late K_n would be contained in their union, since otherwise the compact nested residual sets would have a point outside the union in their intersection. This would disconnect K_n.

Choose an outer Jordan curve surrounding a late annulus and its inner continua. If K were nondegenerate, the region between this curve and K would be a finite-modulus ring by Proposition 24.1. Every finite tail collection of the nested annuli lies essentially inside it, so the Grötzsch inequality bounds its modulus below by each finite tail sum. Divergence is a contradiction. Thus K is a singleton. A nonempty surrounded set is a subset of K and hence equals it. □

For the ring interpretation one may use the filled inner continua supplied by the annuli. Their decreasing intersection remains filled: its complement is the increasing union of their connected complements. The empty-set exception is necessary in Corollary 6.20. An empty set may be surrounded vacuously and is not a singleton. The displayed intersection also preserves the explicit correction of the union symbol in the supplied Proposition 6.16.

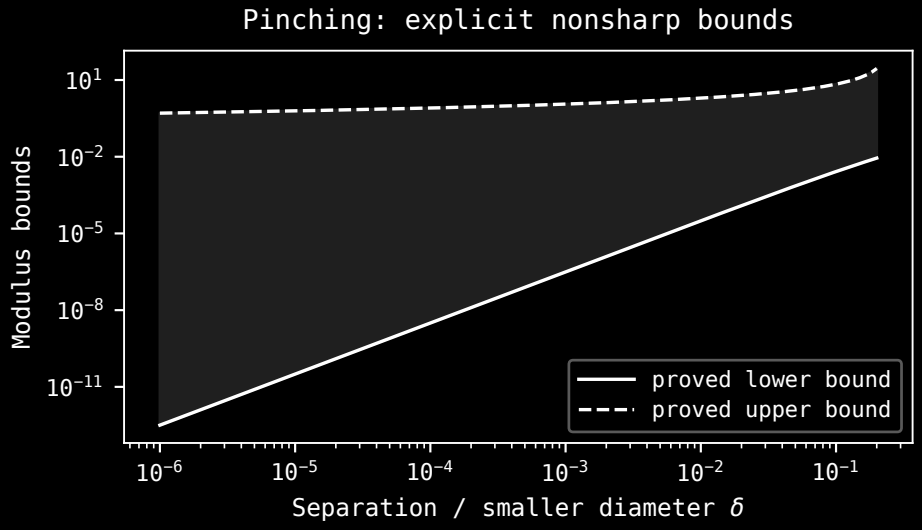


Figure 24.1: Explicit pinching bounds from Theorem 24.4. Their separation reflects sufficient nonsharp constants, not a numerical error estimate.

Chapter 25

Area distortion and a sharp annular deficit identity

25.1 A proof that avoids smooth boundary assumptions

Source Problem 6.15 asks for a lower bound on the ratio of the filled outer and inner areas. Chapter 5 gave a proof through smooth level sets. The argument here proves the stronger exponential bound directly from a conformal Laurent expansion and identifies a sharp quantitative deficit. No differentiability of the original boundary is assumed.

Let $0 < r < 1$, let f be a conformal isomorphism from $A(r, 1)$ onto a planar ring A, and choose the annular orientation so that increasing source radii surround the filled inner complement K. Let $U = A \cup K$. Suppose $0 < \operatorname{area} K \leq \operatorname{area} U < \infty$. Write

$$f(z) = \sum_{n \in \mathbb{Z}} a_n z^n.$$

Laurent expansion and termwise differentiation on compact subannuli are the complex-analysis prerequisites. They follow by Cauchy integrals on two circles enclosing the subannulus and geometric-series expansion of the two kernels, uniformly on any narrower compact subannulus. Thus the series and its derivative can be integrated termwise on each intervening circle.

25.2 The signed boundary-area identity

Lemma 25.1. [A10] *If $r < t < 1$ and $F(t)$ is the area enclosed by $f(\{|z| = t\})$ with the positive orientation around K, then*

$$F(t) = \pi \sum_{n \in \mathbb{Z}} n|a_n|^2 t^{2n}.$$

Moreover,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{F(t)}{t^2}\right) = 2\pi \sum_{n \in \mathbb{Z}} n(n-1)|a_n|^2 t^{2n-3} \geq 0.$$

Proof. The image of the circle is a smooth Jordan curve because f is injective and has nonzero derivative there. Green's area identity gives $F(t) = (2i)^{-1} \int \overline{f}\, \mathrm{d}f$. Substitute the uniformly convergent Laurent series on $z = te^{i\theta}$. Angular orthogonality $\int_0^{2\pi} e^{i(n-m)\theta}\, \mathrm{d}\theta = 2\pi \mathbf{1}_{n=m}$ leaves precisely the displayed sum. Uniform derivative convergence on a slightly wider compact subannulus justifies differentiation. For every integer n, $n(n-1) \geq 0$, with equality only at $n = 0, 1$. Hence every term of the differentiated sum is nonnegative. □

Green's identity is required only for the smooth intermediate image circles. It can be obtained directly by integrating $x\,\mathrm{d}y - y\,\mathrm{d}x$ over a smooth Jordan boundary; decomposition into finitely many graph arcs and the fundamental theorem of calculus gives twice the enclosed area. No boundary regularity of K or U enters this step.

25.3 The source area bound and its equality condition

Theorem 25.2. [A11] *Under the preceding hypotheses,*

$$\frac{\operatorname{area} U}{\operatorname{area} K} \geq \frac{1}{r^2} = e^{4\pi \bmod A} \geq 1 + 4\pi \bmod A.$$

Equality in the first inequality occurs exactly for a round annulus bounded by concentric circles.

Proof. The bounded filled interiors of the intermediate circles are nested. As $t \downarrow r$ their intersection is K, and as $t \uparrow 1$ their union is U. Indeed every point of A has an inverse image of radius strictly between r and one, so it belongs on the asserted side of a sufficiently late circle. Continuity of area under monotone limits gives $F(t) \to \operatorname{area} K$ at the inner end and $F(t) \to \operatorname{area} U$ at the outer end. The monotonicity in Lemma 25.1 yields $\operatorname{area} K/r^2 \leq \operatorname{area} U$. The source's linear bound follows from $e^x \geq 1 + x$ for $x \geq 0$.

If equality holds at the endpoints, the nondecreasing function $F(t)/t^2$ is constant. The nonnegative derivative sum forces $a_n = 0$ for all $n \notin \{0, 1\}$. Thus $f(z) = a_0 + a_1 z$, with $a_1 \neq 0$, and its boundary circles are concentric. Conversely such an affine map has precisely the stated area ratio. □

If the outer area is infinite and the inner area is positive and finite, the inequality holds in the extended sense. A zero-area inner complement does not define the ratio in the source problem. The non-divided area statement must be used instead, with its limiting conventions stated. The theorem therefore resolves the positive finite-area regime even for nonsmooth complementary continua; it does not assign a finite ratio to $0/0$.

25.4 The exact spectral deficit

Define

$$\Delta_A = \operatorname{area} U - r^{-2} \operatorname{area} K.$$

Theorem 25.3. [A12] *For the conformal parametrization above,*

$$\Delta_A = \pi \sum_{n \notin \{0,1\}} n|a_n|^2 (1 - r^{2n-2}),$$

where each summand is nonnegative. Furthermore,

$$\int_{A(r,1)} |f'(z) - a_1|^2 \,\mathrm{d}A \leq (1 + r^2)\Delta_A.$$

The constant $1 + r^2$ is attained by nonaffine univalent maps and hence is sharp for this class.

Proof. On compact radial intervals the formula follows by subtracting the two boundary-area sums. To pass to the endpoints, note first that the negative Laurent part and its derivative have square-summable circular coefficients at radius one, since it is holomorphic on $|z| > r$ and vanishes at infinity. The bounded area limit then forces $\sum_{n \geq 1} n|a_n|^2 < \infty$. At the inner boundary the positive part remains bounded, and positivity of $F(t)$ bounds the negative sum.

These observations justify both endpoint sums. Equivalently one may integrate the nonnegative derivative series and apply monotone convergence. Subtraction gives the identity.

Angular orthogonality and radial integration give

$$\int_{A(r,1)} |f' - a_1|^2 \, \mathrm{d}A = \pi \sum_{n \notin \{0,1\}} n|a_n|^2(1 - r^{2n}).$$

For $n \geq 2$, with $q = r^2 \in (0,1)$,

$$\frac{1 - q^n}{1 - q^{n-1}} = 1 + \frac{q^{n-1}(1-q)}{1 - q^{n-1}} \leq 1 + q.$$

For $n < 0$, the ratio $(q^n - 1)/(q^{n-1} - 1)$ is less than q. Compare each nonnegative term with the corresponding deficit term and sum.

Take $f(z) = z + \epsilon z^2$, where $0 < |\epsilon| < 1/2$. If $f(z) = f(w)$ in the unit disk, then $(z - w)(1 + \epsilon(z + w)) = 0$; the second factor is nonzero, so f is univalent. Only a_1 and a_2 are nonzero. Directly,

$$\Delta_A = 2\pi|\epsilon|^2(1 - r^2), \qquad \int |f' - 1|^2 = 2\pi|\epsilon|^2(1 - r^4),$$

whose ratio is $1 + r^2$. The constant is therefore attained. □

This is an exact derivative-energy conclusion in a fixed annular coordinate. It does not by itself prove a uniform boundary Hausdorff estimate: boundary traces and conformal crowding require additional hypotheses. Nor does it identify the translational coefficient a_0, which does not affect the deficit. Rotating the source annulus multiplies a_1 by a unit complex number and preserves the displayed norms.

25.5 Representation and certified evaluation

For finite Laurent data with rational real and imaginary coefficients and rational r, the normalized deficit Δ_A/π and derivative energy divided by π are exact rational numbers. A finite checker can verify the coefficient support, the ratio inequality and a separately supplied univalence certificate. For $z + \epsilon z^2$, the rational condition $4|\epsilon|^2 < 1$ is sufficient for the elementary univalence proof above. A coefficient-energy identity without such a geometric certificate is an identity for a Laurent polynomial; it is not automatically an area formula for an injective image.

For an infinite expansion, certified coefficient tails are additional input. A truncation may be checked against the exact nonnegative sums: omitted terms cannot lower the true deficit, and an upper tail bound produces a two-sided enclosure. This avoids subtracting two nearly equal floating-point areas when the annulus is close to affine. The certificate cost is linear in the number of supplied coefficients, plus the cost of the tail witness and exact rational arithmetic. It is a computation of a stated identity, not a lower bound on every possible geometric reconstruction algorithm.

25.6 Comparative problem: near-circular annular images

Consider $f(z) = z + \epsilon z^2$ on $A(r,1)$ with rational $0 < r < 1$ and $0 < \epsilon < 1/2$. The requested output is the exact normalized area ratio and the exact distance of f' from its affine coefficient in the L^2 metric. Boundary polygon quadrature is a conventional numerical route; it requires a curve approximation, a signed-area sum and a quadrature-error analysis. The coefficient calculation gives

$$\frac{\operatorname{area} U}{\operatorname{area} K} = \frac{1 + 2\epsilon^2}{r^2 + 2\epsilon^2 r^4}, \qquad \frac{1}{\pi}\int |f' - 1|^2 = 2\epsilon^2(1 - r^4).$$

Both follow from one finite coefficient sum. The gain is exactness and removal of a discretization parameter for this specified polynomial family. For a boundary supplied only as samples, recovering certified Laurent coefficients is a separate inverse task and its cost cannot be suppressed.

The formulas are independently checked by direct symbolic boundary integration in the supplement. They are described as derived identities and a sharp bound. A research-priority claim would require a more complete statement-level comparison with the literature on annular area theorems; that claim is not made in this revision.

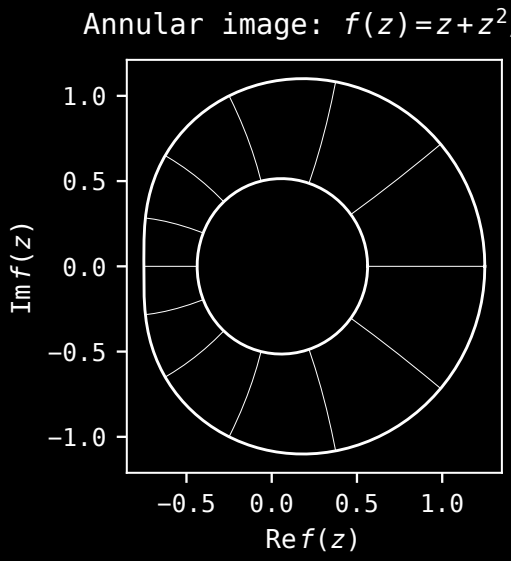


Figure 25.1: The image of $A(1/2, 1)$ under $f(z) = z + z^2/4$. This nonaffine polynomial attains the derivative-error/area-deficit ratio $1 + r^2$; the plotted radial curves are images of coordinate rays.

Chapter 26

A matrix proof of the collar estimate

26.1 The geometric setting and the remaining distinction

Let $S = \mathbb{H}/\Gamma$, where $\Gamma \subset \mathrm{PSL}_2(\mathbb{R})$ is discrete and torsion-free. The upper half-plane carries $\mathrm{d}s^2 = (\mathrm{d}x^2 + \mathrm{d}y^2)/y^2$. Let γ be a primitive simple closed geodesic of length $\ell > 0$. This chapter proves the embedded collar needed for the covering-annulus comparison. It does not require existence or uniqueness of a globally maximizing annulus.

The classification of S as a hyperbolic quotient is a checked dependency on the preceding uniformization volume, with its conformal and completeness conventions retained. All metric and matrix estimates used here are derived below. The classical collar literature provides the attribution and broader context [17, 18].

26.2 Elementary isometries and their fixed axes

A real matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ of determinant one acts by $z \mapsto (az+b)/(cz+d)$. Its derivative is $(cz+d)^{-2}$ and the imaginary part transforms as $\operatorname{Im} z/|cz+d|^2$. Substitution proves invariance of the half-plane metric. A positive dilation $z \mapsto e^{\ell} z$ translates the imaginary axis by length ℓ, since $\int_y^{e^{\ell} y} \mathrm{d}t/t = \ell$.

Geodesics are vertical lines and circles orthogonal to the real axis: the imaginary axis is a length minimizer by $\mathrm{d}s \geq |\mathrm{d}y|/y$, and its images under the isometries just computed give the stated family through every pair of interior points. The perpendicular connector of two disjoint axes is unique; in the particular normalization below it is a circle centered at zero and can be checked by Euclidean orthogonality, which the conformal metric preserves.

26.3 Distance between the imaginary axis and a semicircle

Lemma 26.1. [HO1] *Let the second axis have endpoints $0 < u < v < \infty$, and let D be its distance from the imaginary axis. Then*

$$\cosh D = \frac{u+v}{v-u}.$$

Proof. The second semicircle has center $c = (u+v)/2$ and radius $r = (v-u)/2$. Its perpendicular to the imaginary axis is the circle centered at zero with radius $s = \sqrt{uv}$, because $c^2 = r^2 + s^2$. At their intersection the angular coordinate $\theta \in (0, \pi/2)$ on this last circle satisfies $\sin\theta = r/c$. Along it the metric length element is $\mathrm{d}\theta/\sin\theta$. The distance from the imaginary axis is therefore $D = -\log\tan(\theta/2)$, and $\cosh D = 1/\sin\theta = c/r$. To verify minimality, points of the second semicircle have $x/y = (c + r\cos\varphi)/(r\sin\varphi)$. Differentiation shows that its minimum is s/r,

attained at $\cos\varphi = -r/c$. The Fermi coordinate constructed below is $\operatorname{arsinh}(x/y)$ and is exactly distance to the first axis. Its minimum is therefore $\operatorname{arsinh}(s/r) = D$. □

26.4 Trace separation of equal-length axes

Lemma 26.2. [H02] *If two distinct lifts of γ have distance D, then*

$$\sinh(D/2)\sinh(\ell/2) \geq 1.$$

Proof. Different lifts of a simple geodesic are disjoint. They cannot have one common endpoint at infinity in a discrete torsion-free group. To see the latter assertion, send that endpoint to infinity and the other fixed point of the first element to zero. The elements are $A(z) = \lambda z$ and $B(z) = \mu z + t$, with $\lambda > 1$. If their second fixed points differ, $t \neq 0$ and their commutator is a nonzero translation. Conjugating this translation by A^{-n} yields nonidentity group elements tending to the identity, contradicting discreteness.

Normalize the first axis to the imaginary axis and the second to endpoints $0 < u < v$. Put $a = \ell/2$,

$$A = \begin{pmatrix} e^{a} & 0 \\ 0 & e^{-a} \end{pmatrix}, \qquad P = \frac{1}{\sqrt{v-u}} \begin{pmatrix} v & u \\ 1 & 1 \end{pmatrix}, \qquad B = PAP^{-1}.$$

An inverse may be chosen so that B translates in the displayed direction; both choices belong to the group. Direct matrix multiplication gives

$$\operatorname{tr}(AB^{-1}) = 2\left(\cosh^2 a - \frac{u+v}{v-u}\sinh^2 a\right) = 2\left(1 - (\cosh D - 1)\sinh^2 a\right) < 2.$$

A nonidentity element of a discrete torsion-free subgroup cannot have trace strictly between -2 and 2. Such a matrix has a fixed point in $\mathbb{H}$ and is conjugate to a rotation; rational rotation angle gives torsion, while an irrational angle has nontrivial powers approaching the identity. Since the axes are different, AB^{-1} is not the identity. Its displayed trace must therefore be at most -2. Rearranging gives $(\cosh D - 1)\sinh^2 a \geq 2$, which is the claimed inequality. □

The equal-length condition is essential to this short trace argument. It holds for lifts of the same closed geodesic. The proof is not being applied without modification to arbitrary pairs of geodesics of different lengths.

26.5 Fermi coordinates and an embedded collar

Theorem 26.3. [H03] *The open collar of half-width*

$$w(\ell) = \operatorname{arsinh}\left(\frac{1}{\sinh(\ell/2)}\right)$$

about γ is embedded. It is isometric to

$$(-w, w) \times \mathbb{R}/\mathbb{Z}, \qquad \mathrm{d}s^2 = \mathrm{d}r^2 + \ell^2 \cosh^2 r \, \mathrm{d}t^2.$$

Proof. Coordinates around the imaginary axis in $\mathbb{H}$ are

$$x = e^{\tau}\tanh r, \qquad y = e^{\tau}\operatorname{sech} r.$$

They cover $\mathbb{H}$: $r = \operatorname{arsinh}(x/y)$ and $\tau = \frac{1}{2}\log(x^2 + y^2)$. Differentiation gives $\mathrm{d}s^2 = \mathrm{d}r^2 + \cosh^2 r \, \mathrm{d}\tau^2$, so $|r|$ is the distance to the axis. Translation along the axis changes τ by ℓ; put $\tau = \ell t$.

Lemma 26.2 separates the axis from every distinct translate by at least $2w$. Two open distance-w neighborhoods cannot overlap, because a point in both would give a path between the axes of length less than $2w$. The stabilizer of the axis consists of translations along it. Discreteness makes their length group cyclic; an axis-reversing orientation-preserving isometry would be an elliptic half-turn and is excluded. Primitivity makes the stabilizer exactly the cyclic group generated by the length-ℓ translation. Quotienting its strip therefore gives the claimed embedded cylinder. □

The theorem concerns the open collar. Closed collars with any strictly smaller half-width are embedded as well. At the sharp boundary width tangencies may occur, so replacing the open collar by its full closure would require a separate boundary assertion.

26.6 A stronger covering-modulus estimate

Corollary 26.4. [H04] *Let M_γ be the supremum of moduli of embedded annuli in the free homotopy class of γ. Then*

$$\frac{\pi}{\ell} - 1 < M_\gamma \le \frac{\pi}{\ell}.$$

The right equality is possible when the surface itself is the cyclic annulus.

Proof. The change of variable $y = \ell^{-1} \int_0^r \operatorname{sech} s \, \mathrm{d}s$ makes the collar conformal to a flat cylinder of circumference one and height

$$m_c = \frac{2}{\ell} \arctan\left(\frac{1}{\sinh(\ell/2)}\right).$$

The complete cyclic cover has modulus π/ℓ, obtained by integrating over all $r \in \mathbb{R}$. Their difference equals

$$\frac{\pi}{\ell} - m_c = \frac{2}{\ell} \arctan(\sinh(\ell/2)) < 1.$$

Indeed $f(x) = \arctan(\sinh(x/2))$ has $f(0) = 0$ and $f'(x) = \frac{1}{2}\operatorname{sech}(x/2) < 1/2$ for $x > 0$. Thus $M_\gamma \ge m_c > \pi/\ell - 1$.

Every embedded annulus in the class lifts injectively to the cyclic cover. Covering-space lifting follows by lifting paths from a base point: two choices differ by a loop whose image lies in the cyclic subgroup, so the endpoint is independent of the path; local lifts are holomorphic. Injectivity follows because the original annulus inclusion is injective. Essential-annulus monotonicity then bounds its modulus by π/ℓ. Taking the supremum proves the upper bound. On a cyclic surface the whole surface is an admissible annulus, giving equality. □

This supplies a complete collar-based estimate for a nonperipheral simple geodesic in a torsion-free hyperbolic quotient. If the maximal-annulus existence theorem is separately established, M_γ equals its attained modulus. No existence theorem is inferred merely from a bound on a supremum. The original Lemma 6.40's strict upper inequality requires the previously recorded annular exception. Its weaker additive lower constant is implied by the bound above in this regime.

26.7 Two counterexamples to the unrestricted crossing equality claim

Proposition 26.5. [H05] *The literal equality conclusion in source Exercise 6.31 fails both for a planar annulus paired with a slit quadrilateral and for two annuli on a flat torus.*

Proof. For the planar example take $A = \{1 < |z| < e^h\}$, $h > 0$, and let Γ be all concentric circles in A. Its width is $h/(2\pi)$. Remove the positive radial interval from A to obtain the conformal quadrilateral Q. In logarithmic coordinates $0 < x < h$, $0 < y < 2\pi$, let Δ be the horizontal straight leaves $y = \theta$, $0 < \theta < 2\pi$. They project to radial arcs joining the two circular boundary components. Their width is $2\pi/h$. Every circle meets every radial arc, so the families cross and their width product is one. Their supports are an annulus and a quadrilateral, respectively, not the same quadrilateral as asserted by the unrestricted equality conclusion.

For an independent topological example let $S = \mathbb{R}^2/\mathbb{Z}^2$. Delete the horizontal geodesic $y = 0$ and let Γ be the horizontal geodesics $y = t$, $0 < t < 1$, in the remaining annulus. Delete the vertical geodesic $x = 0$ and let Δ be the vertical geodesics $x = s$, $0 < s < 1$, in that second annulus. Each annulus is a flat cylinder of unit height and circumference. Each foliation therefore has width one. Every pair of leaves meets exactly once, and neither supporting annulus is a quadrilateral. □

The first example is already planar; an ambient-genus exclusion alone does not repair the literal source claim. The strict width-product obstruction remains valid. Both examples concern supports that differ on area-null cuts, demonstrating why uniqueness of an L^2 density cannot identify the supporting open set. The restricted rigidity statement for *two Jordan quadrilateral buffers* used in the canonical-rectangle argument needs a separate proof; neither counterexample refutes or establishes that restricted statement.

26.8 Computational verification and mathematical scope

The supplement checks the matrix trace identity symbolically, the coordinate metric by differentiation, and the gap inequality by its displayed derivative formula. These computations verify algebraic portions of the written proof. They do not replace the group-discreteness argument, the lifting theorem or the topological simplicity hypothesis.

The collar calculation provides a certified scalar comparison once a valid length enclosure is supplied. Since π/ℓ is decreasing on positive lengths, an interval $[\ell_-, \ell_+]$ with $\ell_- > 0$ gives

$$\frac{\pi}{\ell_+} - 1 < M_\gamma \leq \frac{\pi}{\ell_-}.$$

Interval rounding must be directed. If the length interval contains zero, the finite upper estimate is unavailable; a point estimate with a very small positive length cannot silently substitute for the missing lower endpoint.

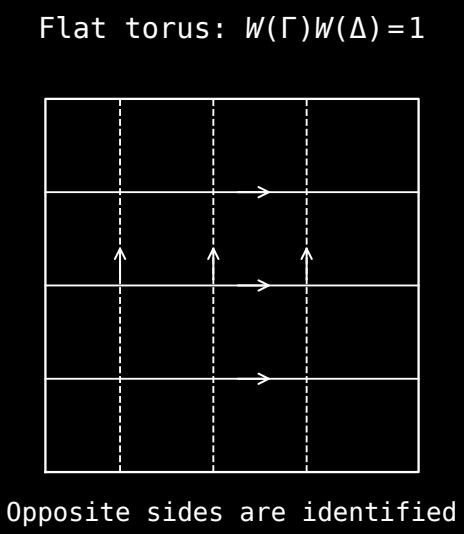


Figure 26.1: The two flat-torus foliations used in the second equality counterexample. Identification of opposite sides is part of the surface, not an unmarked planar crossing.

Chapter 27

Attainment of maximal modulus on a Riemann surface

27.1 Separating attainment from uniqueness

This chapter supplies the nondegeneracy and attainment argument omitted from the earlier review. The target is existence of an embedded annulus with greatest conformal modulus in a specified free homotopy class of essential simple closed curves. The annulus is allowed to have infinite modulus. Uniqueness is a different theorem: the puncture and torus examples already show that an unrestricted uniqueness statement is false. The finite-modulus planar uniqueness development associated with quadratic differentials remains separately identified.

Universal covering and the conformal classification of simply connected Riemann surfaces are reused from the preceding volume with their recorded written-proof status [3]. Little Montel, the omitted-value normal-family theorem, and the nonconstant limit principle for injective holomorphic maps are reused from Volume I [2]. The nondegeneracy argument, cusp construction, lifting operations and limiting embeddings needed for the present application are given explicitly.

27.2 A noncollapse criterion from four omitted points

Lemma 27.1. [MO1] *Let a, b, c, d be four distinct points of the sphere. Suppose injective holomorphic maps f_j are defined on annuli exhausting a fixed nondegenerate annulus. Assume the image of each fixed intermediate circle separates a, b from c, d. A locally spherically uniform limit of the maps is nonconstant.*

Proof. If the limit were a constant p, the image of one fixed intermediate circle would eventually lie in a spherical disk so small that it contains at most one of the four points. A Jordan curve lying in that disk has one complementary component inside the disk; its other component contains all four points outside the disk. It cannot separate the fixed two-point sets. Injectivity makes the intermediate image curve Jordan, so this contradicts the hypothesis. The same argument works when p is one of the omitted points or infinity, since it is expressed in spherical disks. □

The two points on each side are essential to this argument. For a puncture class one side has only one omitted point, and collapse to that point is possible. That case must be handled by a different construction.

27.3 Existence in a nonperipheral planar class

Proposition 27.2. [M02] *Let $S \subset \widehat{\mathbb{C}}$ and let $[\gamma]$ be a nontrivial simple closed curve class which is not peripheral about an isolated puncture on either side. Then the supremum of moduli of embedded annuli in S representing $[\gamma]$ is positive, finite and attained.*

Proof. An embedded tubular neighborhood of a smooth representative is an annulus, so the supremum m_* is positive. A simple curve may be replaced by a smooth one inside the surface: cover it by finitely many coordinate disks, approximate its successive compact subarcs by polygonal arcs in those disks while maintaining their disjointness, and smooth the finitely many vertices inside disjoint small disks. The resulting isotopy remains inside S.

The curve separates the complement of S into two fixed populations. Nontriviality makes both nonempty, and the nonperipheral assumption supplies distinct a, b on one side and c, d on the other. Place d at infinity. For every annulus in the class, its inner continuum contains a, b and its outer continuum contains c, d. The tube-density proof of Proposition 24.1, now using these fixed three finite points rather than closest points, gives the common bound

$$\operatorname{mod} A \le 2 + 2\pi + 2\frac{|a-c|}{|a-b|}.$$

Hence $m_* < \infty$.

Choose embedded annuli of moduli $m_j \to m_*$, with m_j increasing after selection, and normalized parametrizations $f_j : A(1, e^{2\pi m_j}) \to S$ preserving the chosen side orientation. The maps omit a, b, c. The normal-family theorem and a diagonal subsequence therefore give a spherically locally uniform limit on $A(1, e^{2\pi m_*})$. Every fixed intermediate circle has the stated separation, so Lemma 27.1 excludes a constant limit. The injective-limit principle gives an injective meromorphic limit.

The limit omits every point outside S. For a finite omitted point q, apply the local zero-free limit theorem to $f_j - q$, using reciprocal coordinates near infinity when required. A nonconstant limit cannot acquire a value omitted by every map. Thus the image is contained in S and the limit is a conformal embedding.

Finally its intermediate circle remains in the original free homotopy class. Its image is compactly contained in S. On a small tubular neighborhood of that image, sufficiently close image curves are homotopic by interpolation in local surface coordinates; uniform convergence provides this closeness. Equivalently subdivide the circle into finitely many arcs in contractible coordinate disks and concatenate the short homotopies, whose endpoints match. The limit annulus has modulus m_* and represents $[\gamma]$. □

Normal-family compactness alone did not establish this result: it needed both noncollapse and preservation of the class. The argument applies to arbitrary planar domains, including domains with nonlocally connected boundary, because only four complement points and a compact interior loop are used.

27.4 An elementary cusp-neighborhood construction

Lemma 27.3. [M03] *Let a discrete torsion-free subgroup of $\mathrm{PSL}_2(\mathbb{R})$ contain the primitive parabolic $T(z) = z + 1$. For every group matrix $g = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, either $c = 0$ or $|c| \ge 1$. The horoball $\{\operatorname{Im} z > 1\}$ consequently projects to an embedded punctured-disk end.*

Proof. Suppose $0 < |c| < 1$. Recursively put $g_{n+1} = g_n T g_n^{-1}$, beginning with $g_0 = g$. A direct

calculation gives

$$g_{n+1} = \begin{pmatrix} 1 - a_n c_n & a_n^2 \\ -c_n^2 & 1 + a_n c_n \end{pmatrix}.$$

Thus $|c_n| = |c_0|^{2^n} \to 0$ and no c_n vanishes. The recurrence $a_{n+1} = 1 - a_n c_n$ makes a_n bounded: once $|c_n| \le 1/2$, it obeys $|a_{n+1}| \le 1 + |a_n|/2$. Therefore $a_n \to 1$, and the displayed matrices tend to T. They are distinct from T because $c_n \neq 0$. A discrete group has no such accumulation at one of its elements, a contradiction.

If $c = 0$, the map is $z \mapsto \lambda z + t$, with $\lambda > 0$. A value $\lambda \neq 1$ conjugates T to a translation with a different scale; conjugating in the contracting direction produces nontrivial translations tending to the identity. Thus the stabilizer of infinity consists of translations. Discreteness and primitivity identify it with $\langle T \rangle$.

For $c \neq 0$ and $y > 1$,

$$\operatorname{Im} g(z) = \frac{y}{|cz + d|^2} \le \frac{1}{c^2 y} < 1.$$

Hence the horoball is disjoint from all its translates outside the stabilizer. Its quotient by T embeds, and $w = e^{2\pi i z}$ identifies that quotient with $0 < |w| < e^{-2\pi}$, an annulus of infinite modulus. □

The explicit height one is a sufficient normalized cusp neighborhood; maximal cusp height is not asserted. This construction supplies the infinite-modulus alternative without an assumed compact boundary or a finite-area hypothesis.

27.5 Attainment in a hyperbolic surface

Theorem 27.4. [MO4] *Let $S = \mathbb{H}/\Gamma$ be any hyperbolic Riemann surface and let $[\gamma]$ be a nontrivial free homotopy class represented by a simple closed curve. The supremum of embedded annular moduli in that class is attained. It is infinite in the parabolic case and finite in the hyperbolic case of the corresponding deck element.*

Proof. A simple essential loop is primitive. One verification uses its cyclic cover: if its element were the kth power of a primitive element, the simple loop would lift injectively to the corresponding annular cyclic cover with winding number k. A simple essential Jordan curve in an annulus has winding number 1 or -1, so $|k| = 1$. The cyclic stabilizer has a primitive generator because the group is discrete.

A nontrivial deck element is hyperbolic or parabolic. Elliptic elements are excluded by discreteness and torsion-freeness, as proved in the collar chapter. In the parabolic case, conjugate the generator to $T(z) = z + 1$; Lemma 27.3 provides an embedded infinite-modulus annulus in the required class.

In the hyperbolic case, write its translation length as $\ell > 0$. The cyclic cover $A_c = \mathbb{H}/\langle g \rangle$ is a finite conformal annulus of modulus π/ℓ. Every embedded annulus in $[\gamma]$ lifts injectively and essentially to A_c. This follows by the path-lifting argument in Corollary 26.4; it does not require that its geodesic representative has already been used. Thus the supremum m_* is at most π/ℓ and is positive by a tubular annulus.

Take parametrizations of a maximizing sequence and lift them to A_c, choosing the cyclic subgroup consistently. In a round coordinate for A_c, every essential intermediate circle separates its two complementary continua. Each of these has two fixed distinct points. Normality, the four-point noncollapse argument and the injective-limit principle yield an essential embedding $h : A(1, e^{2\pi m_*}) \to A_c$.

Let $q : A_c \to S$ be the cover. The maps $q \circ h_j$ were injective, and converge locally in surface charts to $q \circ h$. This limit is nonconstant, because q is locally biholomorphic. It is injective: if two distinct domain points had the same image, choose disjoint small disks around them and a common target coordinate disk. Each local nonconstant limit maps a still smaller source disk onto a target neighborhood. The zero-count stability theorem would then force the nearby approximating maps to take one common nearby target value on both source disks, contradicting their injectivity. The same local argument excludes a critical point in the limit. Therefore $q \circ h$ is a conformal embedding. Compact-loop homotopy stability, as in Proposition 27.2, preserves the class. Its image attains m_*. □

The proof concerns the actual surface, not merely the cyclic cover: injectivity after projecting the limit is checked separately. It therefore does not assume that an arbitrary annulus embedded in the cover descends injectively.

27.6 Spherical and flat cases

Corollary 27.5. [MO5] *Every nontrivial simple closed curve class on a Riemann surface admits an embedded annulus of maximal modulus, with infinite modulus allowed. Existence does not imply uniqueness.*

Proof. Uniformization reduces the nonhyperbolic cases to the sphere, plane, punctured plane and complex tori. The sphere and plane have no nontrivial simple closed curve classes. The punctured plane itself is an embedded infinite-modulus annulus and attains the supremum in its essential class.

For a torus, let ω represent the primitive lattice vector of the given class. Rescale so that $\omega = 1$ and choose a second lattice generator τ with positive imaginary part. Removing one horizontal closed geodesic leaves an open cylinder of circumference one and height $\operatorname{Im}\tau$, hence an annulus of that modulus. The torus curve-family inequality proved in the source-ordered chapter bounds every annulus in the class by the same area-to-squared-geodesic-length ratio. This annulus is maximal. Removing different parallel geodesics gives distinct maximal annuli, so uniqueness fails in this case. The hyperbolic cases were proved in Theorem 27.4. □

This establishes the existence generalization requested in Exercise 6.37 relative to the cited uniformization and normal-family dependencies. On the sphere with punctures, distinct small punctured disks in the same peripheral class all have infinite modulus. Thus the literal uniqueness statement also needs an exclusion of infinite-modulus peripheral classes. The finite-modulus planar uniqueness theorem remains a separate quadratic-differential obligation; it is not supplied by the compactness argument above.

27.7 Elementary exceptions in the canonical-arc formulation

The source construction of canonical weighted arcs invokes a Cantor limit set and quadrilaterals with disjoint ideal side closures. Those assertions require a nonelementary covering group in the stated finite-type setting. For a cyclic hyperbolic annulus the limit set has two points, not a Cantor set; the two complementary intervals share both endpoints in their closures. For a parabolic punctured disk the limit set has one point. For a simply connected disk the group is trivial. These elementary cases must be distinguished before invoking the nonempty-gap construction or the bound $-3\chi(S)$.

The finite-type negative-Euler-characteristic setting used for planar disks with at least two holes avoids those particular exceptions. The global boundary-gap correspondence, noncrossing

equality rigidity and canonical-rectangle disjointness still require their own complete development. Enumerating at most $-3\chi(S)$ arcs after those premises are granted does not establish the premises themselves.

27.8 Constructive information in the existence proof

The compactness proof selects a subsequence and proves attainment. It supplies no effective modulus of convergence of the maximizing embeddings. The explicit tube bounds and collar estimates give finite numerical constraints, but they do not convert this selection into a terminating maximizer algorithm. A constructive solution needs an input representation of S, an effective geometric separation condition, a search class of embeddings, upper and lower witnesses, and a stopping rule. Those requirements remain visible in the formalization-status record.

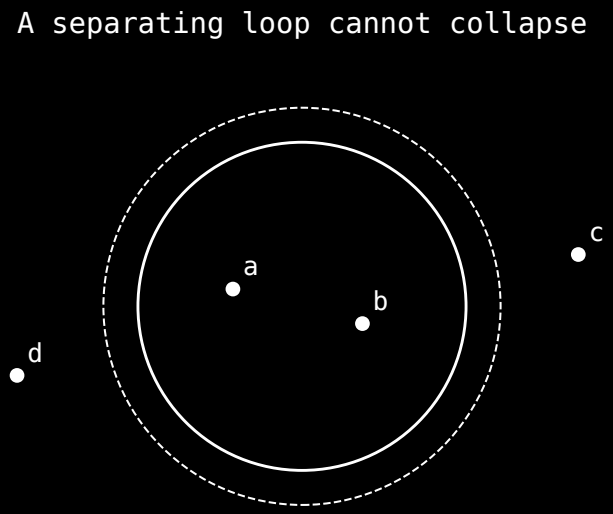


Figure 27.1: Four complementary points, two on each side of an essential loop, exclude collapse of a convergent annular embedding sequence. The picture specifies the separation used in Lemma 27.1.

Chapter 28

Regular condenser duality from variational and flow coordinates

28.1 Boundary charts and the precise regular class

Let S be a connected compact bordered Riemann surface with finitely many boundary components and analytic half-disk charts at its boundary. Divide the whole boundary into two nonempty unions E_0, E_1 of components. Fix a smooth conformal metric, and define

$$E(v) = \int_S |\nabla v|^2 \, \mathrm{d}A.$$

Terminal data are zero on E_0 and one on E_1. There are no mixed terminal changes at corners in this chapter. The square with two insulated sides is a different boundary problem, treated in Chapter 30.

This regular class gives a complete derivation of the potential, flux and curve-family identities used in the source's multiply connected discussion. Domains with arbitrary rough boundary, fractal terminal sets or irregular prime-end trace prescriptions need further definitions; they are not added by omission of the present boundary hypotheses.

28.2 Coercivity on a finite system of charts

Let $H^1_0(S)$ be the completion of smooth functions with zero boundary trace, using the sum of the L^2 norm and the Dirichlet energy. Coordinate changes on a finite compact atlas have bounded metric and area factors. On a rectangle Q, the fundamental theorem along its two coordinate directions gives

$$\int_Q |v - v_Q|^2 \le C_Q \int_Q |\nabla v|^2, \qquad v_Q = |Q|^{-1} \int_Q v.$$

Indeed write the variance as $(2|Q|)^{-1} \int_{Q\times Q} |v(x) - v(y)|^2$, split the difference along a horizontal and a vertical segment, apply Cauchy–Schwarz to each segment, and integrate. The resulting C_Q is bounded by a finite multiple of the sum of the squared side lengths. On a boundary half-rectangle where the trace vanishes on one side, the one-sided fundamental theorem directly bounds $\int_Q |v|^2$ by a finite constant times $\int_Q |\nabla v|^2$.

Choose finitely many interior or boundary coordinate rectangles covering S, with a connected overlap graph and overlaps of positive area. Such a graph is connected after discarding redundant components: otherwise the unions of the two finite subfamilies would disconnect the connected surface. On an overlap O of positive area,

$$|v_{Q_i} - v_{Q_j}| \le |O|^{-1/2} \big(\|v - v_{Q_i}\|_{L^2(O)} + \|v - v_{Q_j}\|_{L^2(O)} \big).$$

Propagate this estimate along finitely many graph paths to a boundary rectangle. The boundary estimate anchors its mean, and the local variance estimates bound all remaining means. Summing proves a finite constant C_S with

$$\|v\|^2_{L^2(S)} \le C_S E(v) \qquad (v \in H^1_0(S)).$$

The argument proves existence of a domain-dependent constant. It does not call the constant universal or claim it has been numerically optimized. A smooth function g equal to the required constants in disjoint boundary collars supplies an extension of the terminal data.

28.3 Existence and weak stationarity

Theorem 28.1. [P10] *There is a unique $u \in g + H^1_0(S)$ minimizing E. It satisfies*

$$\int_S \nabla u \cdot \nabla w \,\mathrm{d}A = 0 \qquad (w \in H^1_0(S)).$$

Its energy $C = E(u)$ is positive.

Proof. The preceding Poincaré estimate makes the Dirichlet inner product a complete norm on $H^1_0(S)$. The affine-subspace projection theorem gives existence, uniqueness and stationarity. If $E(u) = 0$, all local weak derivatives vanish; convolution in coordinate disks makes u locally constant, and connectedness makes it globally constant. Its two different terminal traces rule this out. □

Truncation to $[0, 1]$ preserves the terminal data and cannot increase the energy. For smooth approximations of the scalar truncation, the chain rule gives this inequality; pass to the H^1 limit. Uniqueness therefore implies $0 \le u \le 1$ almost everywhere.

28.4 Interior regularity and analytic reflection

Lemma 28.2. [P11] *The minimizer has a smooth harmonic representative in S and across each analytic boundary chart after subtracting its terminal value and reflecting oddly. It has nonzero normal derivative on the boundary and only finitely many critical points in the interior.*

Proof. In a conformal coordinate, two-dimensional energy stationarity is the distributional equation $\Delta u = 0$. Convolve locally with a smooth mollifier. Each u_ϵ is smooth harmonic in a smaller disk and converges to u in L^2. For a smooth harmonic v, the circle mean is constant in radius: differentiate it and apply the divergence theorem to $\Delta v = 0$. Integrating circle means against any smooth radial unit-mass kernel gives $v = v * \eta_r$. Differentiation and Cauchy–Schwarz consequently bound every derivative of v on a smaller compact disk by a constant times its L^2 norm on a larger disk.

These bounds apply uniformly to u_ϵ. Select a converging subsequence on a countable dense set, use derivative equicontinuity to promote convergence to uniform convergence on compact subsets, and repeat for each derivative by a diagonal selection. Passing through line integrals identifies the limiting derivatives, so the limit is smooth harmonic. Its L^2 limit is u, giving the desired representative.

At a boundary chart, subtract the boundary constant and reflect oddly through the real interval. The trace is zero. For a test function on the reflected disk, its odd difference restricts to a test function vanishing on that interval; weak stationarity proves that the odd extension is distributionally harmonic. No boundary delta term remains because the trace is zero. The

preceding regularity argument makes the extension smooth harmonic. This also proves that the terminal limits hold pointwise.

The mean-value property and nonconstant terminal data give $0 < u < 1$ in the interior. At a point of E_0, take a small interior tangent disk of radius a in a half-disk coordinate. The positive minimum m on its concentric circle of radius $a/2$ gives the barrier $h = m \log(a/|z - z_c|)/\log 2$ on the intermediate annulus. It is zero on the outer circle and at most u on the inner circle. The maximum principle gives $u \geq h$, so the inward normal derivative of u at the tangency point is at least $m/(a \log 2) > 0$. Apply the same argument to $1 - u$ at E_1.

The derivative of a harmonic function in a complex chart is holomorphic. Its zeros are isolated unless the function is locally constant; analytic continuation and the terminal data exclude that case. Compactness of the boundary and the nonzero boundary derivative give a collar with no critical points. The remaining critical set is a discrete subset of a compact interior set, hence finite. □

The reflection step uses the analytic boundary-chart hypothesis. A numerical mesh whose boundary only approximates such a surface requires an additional geometry-error estimate.

28.5 Flux conservation and flow coordinates

Lemma 28.3. [P12] *For every regular value $s \in (0, 1)$, the level set $\{u = s\}$ is a finite separating union of smooth closed curves and*

$$\int_{u=s} |\nabla u| \, \mathrm{d}\ell = C.$$

Except for finitely many critical separatrices, the gradient trajectories connect E_0 to E_1. Along these trajectories, $s = u$ and a local harmonic conjugate t give conformal coordinates with Jacobian $|\nabla u|^2$.

Proof. Regularity, compactness and the implicit-function theorem make a regular level a compact one-dimensional submanifold. A finite coordinate cover shows it has finitely many components. It separates the boundary partitions by continuity of u.

Green's formula applied to $u\Delta u = 0$ gives $C = \int_{E_1} \partial_n u \, \mathrm{d}\ell$; the E_0 term vanishes. The divergence theorem in $\{u > s\}$ gives the stated level flux. These formulas use classical smooth functions on a compact regular surface, so no weak normal trace is being assumed without construction.

Near a critical point choose a local conjugate and set $f = u + it$. If $f - f(p)$ vanishes to order $m \geq 2$, write it as $z^m a(z)$ with $a(0) \neq 0$. A local holomorphic mth root of a gives a coordinate ζ in which $f - f(p) = \zeta^m$. Gradient trajectories are lines of constant imaginary part of f; conformal metric factors change their speed but not their direction. Only finitely many rays can terminate at this critical point. There are finitely many critical points, so the union of their separatrices has area zero.

Along any remaining trajectory, u increases strictly. A limiting value strictly between zero and one would force accumulation in a compact region. If the gradient were bounded below there, u would increase at a fixed positive rate along normalized flow, a contradiction. An accumulated critical point instead places the trajectory on one of the excluded separatrices by the local ζ^m description. Thus the two ends lie at levels zero and one. The nonzero boundary gradient makes them reach the appropriate boundary smoothly.

Where the gradient is nonzero, the conjugate differential is the rotated differential of u. The Cauchy–Riemann equations give the Jacobian $|\nabla u|^2$ and show that t is constant on gradient trajectories. At any regular cross-section, $\mathrm{d}t = |\nabla u| \, \mathrm{d}\ell$ with the positive orientation. Cutting at finitely many critical separatrices and, where necessary, one additional regular trajectory per

component removes conjugate periods. Ordinary change of variables on the resulting coordinate pieces proves the integration formulas used below. The sum of the transverse lengths is the conserved flux C. □

The zero-area separatrix statement is used only after its local analytic description. It is not an assertion that every area-null set can be ignored by every path.

28.6 The two curve families and exact duality

Let Γ_v be the complete family of rectifiable arcs joining E_0 to E_1 in S, and let Γ_h consist of finite unions of pairwise disjoint rectifiable Jordan curves separating the two terminal partitions. The separating curves may have several components; their total line integral is the sum over those components.

Theorem 28.4. [P13] *For the regular condenser above,*

$$\mathsf{W}(\Gamma_v) = C, \qquad \mathsf{W}(\Gamma_h) = C^{-1}, \qquad \mathsf{L}(\Gamma_v) = \mathsf{W}(\Gamma_h) = C^{-1}.$$

Proof. The density $|\nabla u|$ is admissible for the vertical family: on each joining arc the absolute variation of u is at least one. Its energy is C, so $\mathsf{W}(\Gamma_v) \le C$.

Conversely, let ρ be vertical-admissible. Every nonexceptional gradient trajectory has integral at least one. Integrate these inequalities against the conjugate transverse coordinate t. Lemma 28.3 and the conformal Jacobian give

$$C \le \int_S \rho |\nabla u| \, \mathrm{d}A \le \left(\int_S \rho^2 \, \mathrm{d}A \right)^{1/2} C^{1/2}.$$

Hence $\int \rho^2 \ge C$, proving the vertical equality. The argument applies to Borel densities by nonnegative integration and the coordinate change on the finitely cut flow regions. Infinite integrals cause no difficulty.

For a separating multicurve η, orient its components as the boundary separating the region incident to E_1 from E_0. The divergence theorem gives total flux C across it, so $\int_\eta |\nabla u| \, \mathrm{d}\ell \ge C$. This remains valid for rectifiable Jordan components by approximation within a compact collar and the smooth closed flux form; alternatively use the topological invariance of its period and the inequality between period and absolute line integral. The latter invariance follows by covering a compact homotopy by finitely many coordinate disks in which the closed flux form has a primitive: integrals on adjoining edges cancel and integrals on each small boundary vanish. Redundant separating components may be discarded for the flux comparison; their nonnegative absolute line integrals only increase the admissibility bound. Thus $|\nabla u|/C$ is horizontal-admissible and has energy $1/C$.

If ρ is horizontal-admissible, every regular level multicurve has $\int_{u=s} \rho \, \mathrm{d}\ell \ge 1$. Integrating over $s \in (0,1)$, whose excluded critical values form a finite set, yields

$$1 \le \int_S \rho |\nabla u| \, \mathrm{d}A \le \left(\int_S \rho^2 \, \mathrm{d}A \right)^{1/2} C^{1/2}.$$

The same local change of variables proves the coarea equality here. Hence horizontal energy is at least $1/C$. Finally apply the already proved reciprocity of extremal length and width. □

This establishes the source Theorem 6.30 in the stated regular class, including its multiply connected and positive-genus cases. The role of separating multicurves is essential: replacing them by one chosen separating component can lose part of the terminal flux.

28.7 Uniqueness, error identities and the physical normalization

Uniqueness of the potential follows from the variational theorem. The weak extremal-density theorem identifies the vertical minimum density as $|\nabla u|$ almost everywhere, and the horizontal one as $|\nabla u|/C$. For a potential competitor v with the same trace,

$$E(v) - C = \int_S |\nabla(v-u)|^2.$$

This is an exact error identity; its scalar energy has not been identified with a pointwise error or a boundary charge density.

With physical conductivity σ multiplying the normalized scalar metric energy, conductance is σC and resistance is $(\sigma C)^{-1}$. For an electrostatic model with constant permittivity and imposed voltage, stored energy includes the conventional factor $1/2$. Surface charge is proportional to the normal derivative of the potential and need not be uniform on a boundary component. Consequently the source's verbal capacity interpretation must preserve the distinction between conductance/capacity and their reciprocal modulus. These corrections concern physical normalization, while the mathematical identities above use the declared energy without the factor $1/2$.

28.8 What this proof does not transfer automatically

A nodal surface, a metric-measure surface without analytic charts, an anisotropic medium or a nonlinear p-energy does not satisfy the hypotheses merely because its notation contains a gradient. Critical sets, conjugate coordinates and flux foliations change. The square tensor certificate uses a different weak argument that remains valid for its discontinuous matrices. The two developments have been kept separate so that each proof uses the actual available structure.

Chapter 29

Marked quadrilaterals and exact arithmetic for elliptic moduli

29.1 The standard half-plane problem

Let Q_R be the upper half-plane with four cyclically ordered ideal boundary marks $0, 1, R, \infty$, where $R > 1$. Its horizontal boundary sides are $[0, 1]$ and $[R, \infty]$; the connecting family joins those sides through the upper half-plane. The conformal rectangle modulus uses height divided by width in this marking.

A conformal disk equipped with four distinct cyclically ordered points of its ideal boundary is carried to such a configuration by its disk coordinate and a real Möbius normalization of three marks. This statement uses the ideal boundary defined through the conformal disk coordinate. It does not identify arbitrary Euclidean boundary points with distinct prime ends without further hypotheses.

29.2 An explicit rectangle coordinate

Theorem 29.1. [E10] *With the square-root branches continued from the positive values on (R, ∞), a primitive of*

$$F_R'(z) = \frac{1}{\sqrt{z}\,\sqrt{z-1}\,\sqrt{z-R}}$$

is a conformal map from Q_R onto a rectangle. Its horizontal and vertical side lengths are

$$I_0(R) = \int_0^1 \frac{\mathrm{d}t}{\sqrt{t(1-t)(R-t)}}, \quad I_1(R) = \int_1^R \frac{\mathrm{d}t}{\sqrt{t(t-1)(R-t)}}.$$

Consequently $m(R) = \operatorname{mod} Q_R = I_1(R)/I_0(R)$.

Proof. The upper half-plane is simply connected and the displayed derivative is holomorphic and nonzero there, so it has a primitive. At each finite marked point it has an integrable inverse-square-root singularity. At infinity it is $O(z^{-3/2})$. Hence the primitive has continuous finite boundary limits, including at infinity.

On the four real intervals from left to right its derivative points respectively upward, leftward, downward and rightward. This follows by taking each square root with argument in $(0, \pi/2)$ in the upper half-plane: on the boundary negative factors contribute i. The integral around a large upper semicircle tends to zero; the integrals around small half-circles at the finite marks also tend to zero. Thus the four finite side vectors close. Opposite directions and positivity of all four lengths make their image a nondegenerate rectangle, with the stated side lengths.

For any point inside this rectangle the boundary image has winding number one. Exhaust the upper half-plane by large half-disks with small indentations at the marked points. The uniform boundary-limit estimates show that the winding number remains one for sufficiently large/small truncations. The argument principle gives exactly one preimage, with multiplicity. A point outside has winding zero and no preimage. An interior point cannot map to the rectangle's boundary, since the open mapping theorem would then give nearby image points outside. The primitive is therefore a bijective holomorphic map onto the rectangle; its derivative is nonzero, so the inverse is holomorphic. □

The derivative and its boundary asymptotics prove the required rectangle construction directly. No numerical mapping routine enters the existence or injectivity proof.

29.3 Uniqueness and stronger elementary estimates

Corollary 29.2. [E11] *Every marked conformal quadrilateral is equivalent to a standard rectangle, unique up to a positive similarity preserving its marking. For $R > 1$,*

$$\frac{\log R}{\pi} \le m(R) \le \frac{\pi}{-\log(1 - 1/R)}.$$

Proof. Existence follows from the half-plane normalization and Theorem 29.1. The rectangle extremal-length formula makes the aspect ratio a conformal invariant. Rectangles with the same marked aspect ratio differ by a similarity; this gives the stated uniqueness of the standard representative.

For the first estimate, use the density $1/|z|$ in the upper half-annulus $1 < |z| < R$, extended by zero. Every connecting path from $[0, 1]$ to $[R, \infty]$ has weighted length at least $\log R$, while the energy is $\pi \log R$. Its extremal quotient is $\log R/\pi$.

Put $q = 1 - 1/R$. The map $T(z) = 1 - 1/z$ preserves the upper half-plane and carries the marks to $\infty, 0, q, 1$. Interchanging the horizontal and vertical sides inverts the rectangle modulus. Scaling the latter configuration by $1/q$ therefore yields $m(R)^{-1} = m(1/q)$. Apply the first estimate to $1/q > 1$ and invert to obtain the upper bound. □

These constants are stronger than the corresponding factors $1/(4\pi)$ and 4π printed in Exercise 6.23, with the same marking convention. Reversing that marking inverts every modulus and must not be hidden inside notation.

29.4 A complete arithmetic–geometric-mean reduction

For $a, b > 0$ set

$$J(a, b) = \int_0^\infty \frac{\mathrm{d}x}{\sqrt{(x^2 + a^2)(x^2 + b^2)}}.$$

The substitution $x = b \tan\theta$ gives $J(a, b) = \int_0^{\pi/2} (a^2 \cos^2\theta + b^2 \sin^2\theta)^{-1/2}\, \mathrm{d}\theta$.

Lemma 29.3. [E12] *If $A = (a + b)/2$ and $B = \sqrt{ab}$, then $J(a, b) = J(A, B)$. If $M(a, b)$ is the common limit of arithmetic and geometric iteration, then*

$$J(a, b) = \frac{\pi}{2M(a, b)}.$$

Proof. Set $y = (x - ab/x)/2$. This strictly increases from $-\infty$ to ∞ as x increases from zero to infinity. Direct identities give

$$(x^2 + a^2)(x^2 + b^2) = 4x^2(y^2 + A^2), \qquad \mathrm{d}y = \sqrt{y^2 + ab}\,\frac{\mathrm{d}x}{x}.$$

Thus the transformed integral is one half of the integral of $\{(y^2 + A^2)(y^2 + B^2)\}^{-1/2}$ over the entire real line. The integrand is even, proving invariance.

Assume $a \geq b$. The iterates satisfy $b \leq B \leq A \leq a$, and $A - B = (\sqrt{a} - \sqrt{b})^2/2 \leq (a - b)/2$. Their nested bounds converge to a common positive limit M. At each stage the angular integral lies between $\pi/(2a_n)$ and $\pi/(2b_n)$. Invariance and squeezing give $J(a, b) = \pi/(2M)$. □

With $k^2 = 1/R$, substitutions $t = \sin^2\theta$ in I_0 and $t = 1 + (R - 1)\sin^2\theta$ in I_1 yield

$$m(R) = \frac{K(\sqrt{1 - k^2})}{K(k)} = \frac{M(1, \sqrt{1 - 1/R})}{M(1, \sqrt{1/R})},$$

where $K(k) = \int_0^{\pi/2}(1 - k^2\sin^2\theta)^{-1/2}\,\mathrm{d}\theta$. The factor π cancels from the final ratio. The classical identity is documented in [22]. The algorithm evaluates that identity; it is not presented as a new elliptic-integral algorithm.

29.5 A rational certificate for the transcendental modulus

Theorem 29.4. [E13] *Suppose a certificate encloses each initial square root by positive rationals $b_0^- \leq \sqrt{q} \leq b_0^+$, with $a_0^- = a_0^+ = 1$. At each iteration it records rational intervals satisfying*

$$a_{j+1}^{\pm} = (a_j^{\pm} + b_j^{\pm})/2,$$

$$(b_{j+1}^-)^2 \leq a_j^- b_j^-, \qquad (b_{j+1}^+)^2 \geq a_j^+ b_j^+, \qquad 0 < b_{j+1}^- \leq b_{j+1}^+.$$

Then $M(1, \sqrt{q}) \in [b_N^-, a_N^+]$. Applying this construction to $q = 1 - 1/R$ and $q = 1/R$ gives a sound rational interval for $m(R)$ by interval division.

Proof. The initial squared inequalities and positivity enclose the true square root. Arithmetic averaging and square root of a positive product are increasing in each argument. Induction therefore encloses both true iterates. Their exact ordered values satisfy $b_N \leq M \leq a_N$, so the stated rational endpoints contain M. Both denominator endpoints in the ratio are positive. Monotonicity of division in positive numerator and denominator gives the lower ratio from numerator-lower/denominator-upper and the upper ratio from numerator-upper/denominator-lower. □

The supplied checkers also compare the radius parameter and requested tolerance with an independent model file. They reject a false bound, omitted or inconsistent iteration, nonpositive square-root endpoint, or changed problem parameter. Acceptance of the requested relative tolerance is a separate rational comparison after soundness of the enclosure is established.

The exact iteration satisfies the sharper relation

$$a_{j+1} - b_{j+1} = \frac{(a_j - b_j)^2}{2(\sqrt{a_j} + \sqrt{b_j})^2} \leq \frac{(a_j - b_j)^2}{8b_j}.$$

This proves local quadratic convergence once the positive lower iterate is controlled. Fixed arithmetic precision can create an interval floor; the implementation must then increase precision or report failure. The existence of arbitrarily accurate rational traces follows by first selecting a finite exact prefix with the desired gap and then enclosing that prefix at sufficiently high rational precision. It does not imply completion inside an arbitrary fixed precision/iteration cap.

29.6 An independent validated quadrature representation

The half-angle substitution $s = \tan(\theta/2)$ gives

$$K(\sqrt{q}) = \int_0^1 \frac{2\,\mathrm{d}s}{\sqrt{s^4 + (2-4q)s^2 + 1}}, \qquad 0 < q < 1.$$

On a rational interval $[a, b] \subset [0, 1]$, write $u = s^2$. The quadratic $u^2 + (2-4q)u + 1$ attains its maximum at one of a^2, b^2, and its minimum at one of those endpoints or at $u = 2q - 1$ when the latter belongs to the interval. All these candidate values are rational. Positive rational square-root enclosures therefore give exact lower and upper bounds on the integral over each subinterval.

Summing a finite rational partition yields a separate certificate. As mesh size tends to zero and square-root enclosure widths tend to zero, the integral bounds converge because the polynomial has a positive minimum on $[0, 1]$. More quantitatively, $|P'| \le 8$ and $P \ge p_0 > 0$ imply that the integrand is Lipschitz with constant at most $8p_0^{-3/2}$. This proves a mesh-dependent enclosure bound and supplies a computable sufficient stopping criterion after bounding $p_0^{-3/2}$ rationally.

The arithmetic–geometric iteration and the interval quadrature are different mathematical representations of the same integral. An agreement between ordinary floating-point libraries alone would not supply either certificate. The delivered quadrature is a transparent low-order validated baseline; it is not claimed to be the fastest available validated integration method.

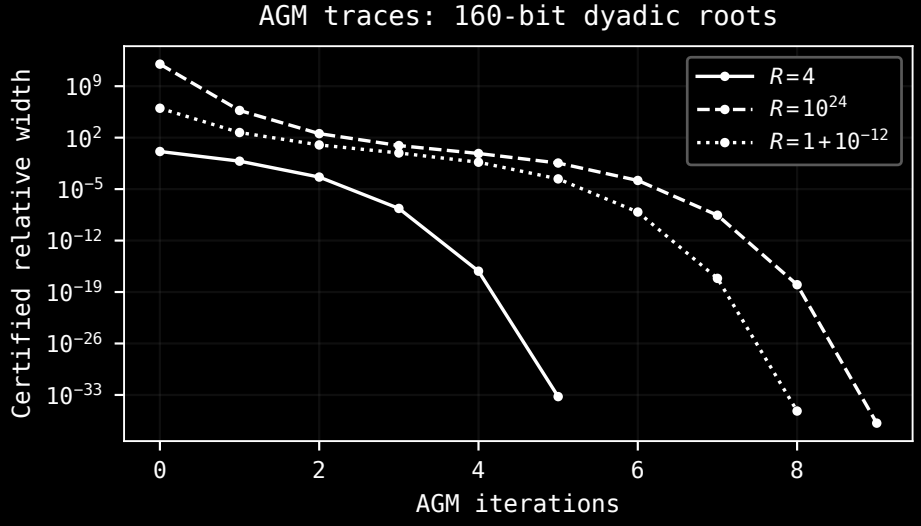


Figure 29.1: Exact rational trace gaps at three marked parameters. Points are accepted finite-stage enclosures; interpolation between the displayed iterations is only a visual guide.

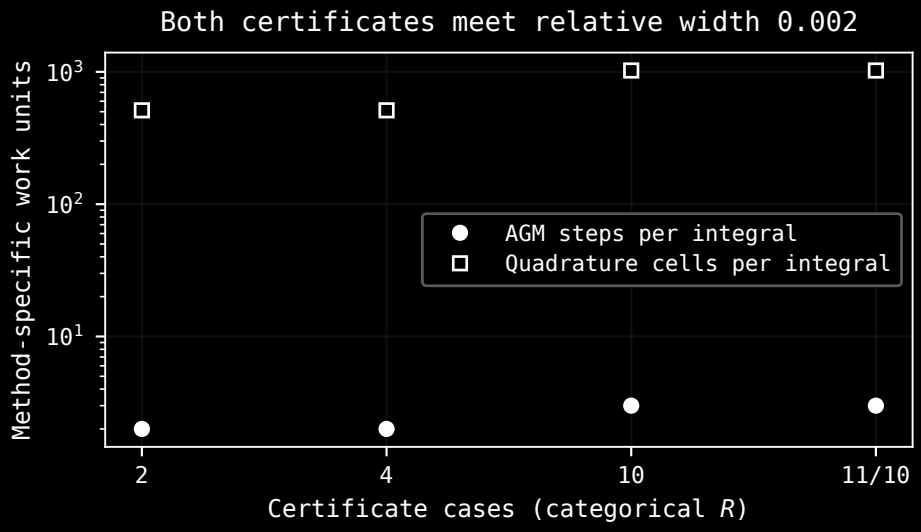


Figure 29.2: Two certified representations satisfying the same requested relative width. An AGM step and a quadrature cell are different work units; the plot is not a wall-clock speedup.

Chapter 30

A two-dimensional continuum certificate with independent implementations

30.1 The fixed mathematical problem

Let $\Omega = (0,1)^2$. The coefficient field is a measurable real symmetric matrix $A(x)$ satisfying

$$\alpha|\xi|^2 \le \xi^T A(x)\xi \le \beta|\xi|^2 \quad \text{for almost every } x \in \Omega, \quad 0 < \alpha \le \beta < \infty.$$

The input is the coefficient field, its interpretation on the square, and the terminal convention: potential zero on the left side and one on the right side. The other two sides are insulating in the weak variational formulation. The requested output is a rigorous interval containing

$$C_x[A] = \min_{v \in x + V_x} E_A(v), \qquad E_A(v) = \int_\Omega \nabla v^T A \nabla v,$$

where V_x is the closure in $H^1(\Omega)$ of smooth functions vanishing on the two vertical sides. The analogous space V_y uses the horizontal sides. These are conditions on the continuum problem, independent of the mesh subsequently chosen for its approximation.

The target is a static mathematical conductivity model. No physical measurement is used in this chapter. Tensor entries have been nondimensionalized; reinstating a common physical conductivity multiplies C_x by that conductivity. Neither dielectric dispersion nor time-dependent transport is inferred from this static identity. The classical duality literature distinguishes such regimes explicitly [19].

30.2 Sobolev spaces and the terminal traces used here

The norm of a smooth function on the closed square is

$$\|v\|_{H^1}^2 = \int_\Omega (|v|^2 + |v_x|^2 + |v_y|^2).$$

Complete this inner-product space. Limits of derivatives are distributional derivatives, since integration by parts against a compactly supported smooth test function passes through L^2 convergence. This identifies the completion with the functions and weak derivatives needed below; no pointwise differentiability of a minimizer is presumed.

For a smooth v the fundamental theorem of calculus and Cauchy–Schwarz give, at fixed y,

$$|v(0,y)|^2 \le 2\int_0^1 |v(x,y)|^2\,\mathrm{d}x + 2\int_0^1 |v_x(x,y)|^2\,\mathrm{d}x.$$

One obtains this bound by writing $v(0,y) = v(t,y) - \int_0^t v_x(s,y)\,\mathrm{d}s$, applying $(a+b)^2 \le 2a^2 + 2b^2$, and integrating in t. The same bound holds at $x = 1$ and on the horizontal sides. Thus restriction to each side extends continuously from smooth functions to an L^2 trace for the completion.

The subspace V_x consists exactly of the functions with zero vertical traces. Here is the approximation argument for the nontrivial inclusion. Extend such a function oddly through each vertical side and evenly through each horizontal side. The zero trace prevents a jump term in the distributional derivative of the odd extension: integration by parts on the two adjacent rectangles cancels the trace contributions. Reflection preserves local H^1 norms up to the finite number of copies. Convolution with a smooth even mollifier is odd about the two vertical sides and hence vanishes there. Its restriction to the square converges in H^1. Translation continuity needed for this convergence follows first for continuous functions, then for step functions on rectangles by the vanishing measure of translated symmetric differences, and finally for L^2 functions by step-function approximation. The same construction, with coordinates interchanged, identifies V_y with zero horizontal traces.

For $w \in V_x$, the one-dimensional fundamental theorem along horizontal lines, initially for smooth functions and then by completion, gives $\|w\|_2 \le \|w_x\|_2$. Consequently the gradient energy is a complete norm on V_x, equivalent to its H^1 norm. Coercivity and the projection argument already proved in the supporting chapter give a unique minimizer $u \in x + V_x$ and

$$\int_\Omega \nabla w^T A \nabla u = 0 \qquad (w \in V_x).$$

Strict positivity follows because $\int_\Omega u_x = 1$, so $E_A(u) \ge \alpha \int u_x^2 \ge \alpha$. The admissible function x gives $C_x[A] \le \beta$. All uses of weak minimizers below are therefore justified without an assumed smooth solution at material interfaces.

30.3 The conjugate-potential pairing

Put

$$R = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad B = R^T A^{-1} R.$$

For a symmetric 2×2 matrix, direct inversion gives $B = A/\det A$.

Lemma 30.1. [CT01] *For every $v \in x + V_x$ and $\psi \in y + V_y$,*

$$\int_\Omega \nabla v \cdot R \nabla \psi = 1.$$

Proof. For smooth functions, $\nabla v \cdot R\nabla\psi = v_x\psi_y - v_y\psi_x = \operatorname{div}(v\psi_y, -v\psi_x)$. The left boundary contributes zero because $v = 0$; the right boundary contributes $\int_0^1 \psi_y(1,y)\,\mathrm{d}y = 1$. On each horizontal side ψ is constant, so its tangential derivative ψ_x vanishes. The remaining boundary contributions are zero. Approximate $v - x$ and $\psi - y$ by the defining smooth sequences. The bilinear integral is continuous in their gradient L^2 norms by Cauchy–Schwarz, and the formula passes to the limit. □

The identity is a global compatibility statement. It is not inferred from a residual at mesh vertices.

Theorem 30.2. [CT02] *Let $v \in x + V_x$ and $\psi \in y + V_y$. Define*

$$U = E_A(v), \qquad F = \int_\Omega \nabla\psi^T B \nabla\psi, \qquad L = F^{-1}.$$

Then $F > 0$ and $L \le C_x[A] \le U$.

Proof. Positive definiteness and the nonconstant terminal trace imply $F > 0$. The upper bound is the variational definition. Apply Lemma 30.1 with the exact minimizer u and the supplied ψ, then apply the weighted Cauchy–Schwarz inequality:

$$1 \le \left(\int \nabla u^T A \nabla u \right) \left(\int (R\nabla\psi)^T A^{-1} (R\nabla\psi) \right) = C_x[A] F.$$

This proves the lower bound. No stationarity condition is required of either supplied witness. □

A separately computed potential and stream function therefore supply independently admissible upper and lower witnesses. An unsuccessful approximate solve can still give a valid but wide enclosure. Failure to attain a requested width is a failure of accuracy, not of the enclosure theorem.

30.4 Exact reciprocity without interface smoothness

Theorem 30.3. [CT03] *For the weak square problems just defined,*

$$C_y[B] = C_x[A]^{-1}.$$

Proof. Write $C = C_x[A]$ and $j = A\nabla u \in L^2(\Omega; \mathbb{R}^2)$. Define

$$\psi_0(x, y) = \int_0^y j_1(x, t)\, \mathrm{d}t.$$

Fubini and Cauchy–Schwarz give $\psi_0 \in L^2$ and $(\psi_0)_y = j_1$ in distributions. We verify the other derivative explicitly. For $\varphi \in C_c^\infty(\Omega)$ put $G(x, t) = \int_t^1 \varphi(x, y)\, \mathrm{d}y$. This function has zero vertical traces and belongs to V_x. Weak stationarity gives

$$0 = \int j \cdot \nabla G = \int j_1(x, t) \int_t^1 \varphi_x(x, y)\, \mathrm{d}y\, \mathrm{d}x\, \mathrm{d}t - \int j_2 \varphi = \int \psi_0 \varphi_x - \int j_2 \varphi.$$

Hence $(\psi_0)_x = -j_2$ and $\psi_0 \in H^1$.

The function $I(x) = \int_0^1 j_1(x, y)\, \mathrm{d}y$ is constant in the distributional sense: test stationarity with $g(x)$ for $g \in C_c^\infty(0, 1)$. Thus $\int I g' = 0$. A locally integrable function with zero distributional derivative is constant, as follows by convolution on compact intervals and passage to the limit. Test stationarity with $u - x \in V_x$ to obtain

$$C = \int j \cdot \nabla u = \int j \cdot \nabla x = \int_0^1 I(x)\, \mathrm{d}x = I.$$

The vertical integration defining ψ_0 has horizontal traces zero and C. To verify that these agree with the Sobolev traces, regard a smooth approximation to ψ_0 as a function of y with values in $L^2(0, 1)_x$. The one-dimensional bound $\sup_y \|w(y)\|_{L^2_x}^2 \le 2\|w\|_{L^2(\Omega)}^2 + 2\|w_y\|_{L^2(\Omega)}^2$ shows that an H^1-convergent smooth sequence has a uniformly convergent L^2_x-valued representative. Integrating its derivative and passing to the limit identifies the representative with $\int_0^y j_1(x, t)\, \mathrm{d}t$.

Its endpoint values are exactly zero and $I(x) = C$, respectively. Therefore $\psi = \psi_0/C \in y + V_y$ and $R\nabla\psi = j/C$. Its dual energy is

$$E_B(\psi) = C^{-2} \int j^T A^{-1} j = C^{-1}.$$

The pairing theorem bounds every admissible dual energy below by C^{-1}, and this particular function attains it. Hence $C_y[B] = C^{-1}$. □

This proof constructs the dual minimizer from the weak flux. It uses neither classical normal derivatives at discontinuous interfaces nor a presumed global smooth conjugate.

30.5 A continuum reference value for a checkerboard

Let a_κ equal $\kappa > 0$ in the lower-left and upper-right quarter squares, and one in the other two quarters. Set $A = a_\kappa I$. Boundaries and interfaces have area zero and do not change this weak model.

Corollary 30.4. [CT04] *The horizontal square conductance satisfies*

$$C_x[a_\kappa I] = \sqrt{\kappa}.$$

Proof. Quarter-turn rotation about the square's center interchanges the two pairs of terminal sides and interchanges the two materials. The rotated scalar field is κ/a_κ. Change variables in the energy and, if needed, replace v by $1 - v$ to preserve the terminal labels. It follows that $C_x[a_\kappa I] = C_y[(\kappa/a_\kappa)I]$. Energy scales linearly with a constant coefficient, so the right side is $\kappa C_y[a_\kappa^{-1} I]$. The reciprocity theorem makes this $\kappa/C_x[a_\kappa I]$. Positivity selects the positive square root. □

The self-dual calculation belongs to the classical two-dimensional conductivity-duality framework [19, 20]. It is used as an independently known continuum reference. The derivation here is explicitly for this finite square with mixed terminal/insulating boundaries; no appeal to an infinite periodic homogenization limit is needed.

When κ is rational, certificate containment of the reference is checked without a floating-point square root: $L^2 \le \kappa \le U^2$ together with $L, U > 0$ is equivalent to $\sqrt{\kappa} \in [L, U]$.

30.6 Conforming triangular witnesses

Subdivide the square into n^2 equal squares and divide each along its southwest–northeast diagonal. Store one rational symmetric positive-definite coefficient matrix per triangle and rational nodal values for v and ψ. Continuity of the piecewise-affine interpolants gives H^1 functions. Enforcing terminal values at every boundary vertex enforces the constant affine trace on every terminal edge. Thus the hypotheses of Theorem 30.2 are exact finite conditions.

The rotated gradient $R\nabla\psi$ is constant on each triangle. Its normal component agrees across an internal edge because that component is the tangential derivative of the continuous trace of ψ. Consequently it is divergence-free in distributions. On a horizontal boundary edge its normal component is zero, since ψ is constant there. This is a second direct verification of the flux interpretation of the dual witness.

For a lower triangle with values at southwest, southeast and northeast vertices,

$$d_x v = v_{SE} - v_{SW}, \qquad d_y v = v_{NE} - v_{SE}.$$

For the upper triangle use $d_x v = v_{NE} - v_{NW}$ and $d_y v = v_{NW} - v_{SW}$; define $d\psi$ in the same way. The actual gradient has an additional factor n, while the area is $1/(2n^2)$. For $A_T = \left(\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right)$ the exact triangle contributions are therefore

$$E_T = \tfrac{1}{2}\{a(d_x v)^2 + 2b(d_x v)(d_y v) + c(d_y v)^2\},$$

$$F_T = \frac{a(d_x \psi)^2 + 2b(d_x \psi)(d_y \psi) + c(d_y \psi)^2}{2(ac - b^2)}.$$

The exact pairing is $P_T = \{(d_x v)(d_y \psi) - (d_y v)(d_x \psi)\}/2$. All operations are rational. Summing must give $P = 1$, which is independently checked even though the written pairing theorem already implies it for correctly assembled terminal data.

30.7 What the finite checker establishes

Theorem 30.5. [CT05] *For an accepted certificate in the supplied triangular-square format, bound to a separately supplied material model, the claimed interval contains the capacity of that fixed infinite-dimensional weak condenser.*

Proof. The parser establishes the mesh dimensions, exact tensor count and vertex count. Each tensor has $a > 0$ and $ac - b^2 > 0$, hence is positive definite. Finitely many positive-definite matrices yield finite positive uniform ellipticity bounds, whether or not those bounds are numerically computed. The checker compares every tensor and the mesh against the separately specified model; the certificate cannot change the problem it claims to solve.

The checked vertex traces produce admissible continuum functions by the conforming argument above. Exact summation establishes the actual values E, F of their continuum integrals. The checker requires $E, F > 0$, $P = 1$, and exact equality of the claims with $L = 1/F$, $U = E$. Theorem 30.2 now applies. Parsing, interpolation, model binding and integration are all needed; a successful numerical solve alone is not among these proof premises. □

The C++ checker uses Boost multiprecision rational arithmetic; the independent Python checker uses integer numerators and denominators through `Fraction`. This supplies independent parsers, loops and arithmetic implementations at the application level. Their runtimes, arithmetic libraries, compilers and hardware remain trusted components of the ordinary executable checks. No proof-assistant refinement of either program has been performed.

The native approximate generator uses piecewise-linear stiffness assembly and preconditioned conjugate gradients. Its double-precision output is rounded to rational dyadic vertex values, with terminal values kept exact. Rounding changes the witnesses and their energies; the final bounds are recomputed from those rounded values. A relative residual threshold is a generation rule, not the mathematical acceptance criterion.

30.8 Completeness and termination of the specified fragment

Proposition 30.6. [CT06] *For a fixed uniformly elliptic piecewise-constant square model aligned with the nested dyadic triangulations, conforming rational potential and stream witnesses can make the certified gap arbitrarily small. Every finite certificate is checked in finitely many arithmetic operations.*

Proof. Approximate the primal and dual minimizers in H^1 by smooth functions with the correct terminal traces, using the preceding reflection argument. Interpolation of a fixed smooth function on a regular refining triangular grid converges in gradient norm: Taylor's formula on each triangle

bounds the gradient error by a constant times the mesh diameter and the supremum of its second derivatives. First choose the smooth approximation and then a fine enough grid. Continuity of the two quadratic energies proves convergence to C and $1/C$.

For a fixed finite grid, rational nodal values are dense in the real nodal values while boundary values zero and one stay exact. Polynomial continuity of each finite energy, and positivity of the limiting dual energy, show that these rational witnesses can achieve any prescribed positive gap. This is a completeness statement for the unbounded family of exact rational witnesses over all sufficiently fine dyadic meshes. It does not give a mesh-complexity bound without additional regularity.

The checker traverses $2n^2$ triangles and $(n+1)^2$ vertices. Each exact rational operation terminates by finite integer arithmetic. Under the delivered parser's explicit mesh and input-bit limits the actual accepted fragment is finite in encoding length. The unrestricted density theorem must not be confused with completeness inside those fixed implementation limits. □

The implemented meshes are bounded by $n \le 128$. The Python checker also enforces a 4096-bit bound on each parsed rational. Those are implementation limits, not a claim that every requested tolerance is reachable within them. Variable tensor denominators can increase intermediate integer lengths; an $O(n^2)$ arithmetic count alone is not an $O(n^2)$ bit-operation bound. Any common-denominator reduction and greatest-common-divisor cost must be included in bit accounting.

On nested meshes, exact Galerkin minima decrease for both energies. Approximate numerical solutions and independently rounded nodal fields need not preserve that monotonicity. Intersecting valid intervals from all completed meshes produces a monotone enclosure, provided the material model is the same. Disjoint intervals indicate a failed hypothesis, implementation defect or corrupted record; they are not averaged.

30.9 A-posteriori function error and coefficient uncertainty

For the exact minimizers u and ψ_*, orthogonality gives

$$E_A(v) - C = \|v - u\|_A^2, \qquad E_B(\psi) - C^{-1} = \|\psi - \psi_*\|_B^2.$$

Since $C \in [L, U]$, one obtains the certified estimates

$$\|v - u\|_A^2 \le U - L, \qquad \|\psi - \psi_*\|_B^2 \le F - U^{-1}.$$

These are energy-norm estimates for the actual continuum fields. Pointwise maxima, boundary flux peaks and derivatives at interface corners require separate regularity assumptions.

If coefficient uncertainty supplies Loewner bounds $A_- \preceq A \preceq A_+$, minimization shows $C_x[A_-] \le C_x[A] \le C_x[A_+]$. A lower certificate for A_- and an upper certificate for A_+ therefore enclose the uncertain-model capacity. The matrix inequalities need their own exact or validated verification; entrywise interval ordering is not generally a Loewner ordering.

Chapter 31

Work, precision and stopping rules for the continuum witnesses

31.1 The matrix problem and its relation to the continuum

Each of the two conforming witness problems has $N = n^2 - 1$ free vertex values, after imposing its two terminal sides. Let K be the free stiffness matrix and b its boundary load. Positive definiteness follows from the coercive integral and the zero-trace Poincaré inequality. Solving $Kx = b$ gives a finite-element minimizer; it does not, by itself, give the exact continuum solution.

There are therefore two separate errors. The difference between an approximate vector and the finite-element minimizer is an algebraic error. The difference between that finite-element function and the weak continuum minimizer is a discretization error. The complementary-energy certificate bounds their combined effect after the actual approximate vectors are rationalized. A small algebraic residual alone does not control the second error.

31.2 Conjugate gradients in exact arithmetic

For K positive definite and $r_0 = b - Kx_0$, put $p_0 = r_0$. Until $r_k = 0$, define

$$\alpha_k = \frac{r_k^T r_k}{p_k^T K p_k}, \quad x_{k+1} = x_k + \alpha_k p_k, \quad r_{k+1} = r_k - \alpha_k K p_k,$$

$$\beta_k = \frac{r_{k+1}^T r_{k+1}}{r_k^T r_k}, \qquad p_{k+1} = r_{k+1} + \beta_k p_k.$$

All nonzero denominators are positive. This is the classical method of Hestenes and Stiefel [21]. The following argument records the mathematical facts used for the cost comparison; finite-precision executions have a separate certificate.

Theorem 31.1. [K01] *In exact arithmetic, the nonzero directions are pairwise K-orthogonal and the residual r_k is orthogonal to the preceding directions. The iterate x_k minimizes the energy error over*

$$x_0 + \mathcal{K}_k, \qquad \mathcal{K}_k = \operatorname{span}\{r_0, Kr_0, \dots, K^{k-1}r_0\}.$$

The method terminates with the exact solution in at most N steps.

Proof. Proceed by induction. Since $p_k = r_k + \beta_{k-1}p_{k-1}$ and r_k is orthogonal to earlier directions, $r_k^T p_k = r_k^T r_k$. The stated value of α_k therefore makes r_{k+1} orthogonal to p_k. For $i < k$, the residual recurrence and the induction hypothesis give $r_{k+1}^T p_i = r_k^T p_i - \alpha_k p_k^T K p_i = 0$. Thus the

new residual is orthogonal to all directions through p_k, and hence to the earlier residuals, since each residual is a linear combination of two successive directions.

For $i < k$, use $Kp_i = (r_i - r_{i+1})/\alpha_i$ to obtain $r_{k+1}^T K p_i = 0$. For $i = k$ this same identity gives $r_{k+1}^T K p_k = -\|r_{k+1}\|^2/\alpha_k$. Also $p_k^T K p_k = \|r_k\|^2/\alpha_k$. The definition of β_k now proves $p_{k+1}^T K p_i = 0$ for every $i \le k$. This completes the orthogonality induction.

The recurrences put the first k directions in $\mathcal{K}_k$. Before termination, their K-orthogonality makes them independent, so they span this Krylov space. If $x_* = K^{-1}b$, then $K(x_* - x_k) = r_k$ is orthogonal to $\mathcal{K}_k$. Expanding the squared K-norm of $x_* - x_k - v$ for $v \in \mathcal{K}_k$ proves the minimizing property. There cannot be more than N nonzero mutually K-orthogonal vectors. At that dimension the orthogonal residual is zero, so the solution is exact. □

Finite termination in exact arithmetic is not a theorem about the number of iterations of the floating-point program. Rounding can destroy exact orthogonality; breakdown and residual drift require explicit implementation handling.

31.3 A complete polynomial error bound

Theorem 31.2. [K02] *Let $0 < \lambda_- \le \lambda_+$ bound the spectrum of K, and put $\kappa = \lambda_+/\lambda_-$. If $\kappa > 1$, the exact iterates obey*

$$\frac{\|x_* - x_k\|_K}{\|x_* - x_0\|_K} \le 2\left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}\right)^k.$$

For $\kappa = 1$ one nontrivial step suffices.

Proof. For a polynomial p of degree at most k with $p(0) = 1$, the vector $p(K)(x_* - x_0)$ differs from the original error by an element of $\mathcal{K}_k$, since $K(x_* - x_0) = r_0$. Minimality and the orthonormal spectral decomposition give an error factor bounded by $\max_{[\lambda_-,\lambda_+]} |p|$.

Define T_k by $T_0 = 1$, $T_1 = x$, and $T_{k+1} = 2xT_k - T_{k-1}$. The angle-addition formula proves $T_k(\cos\theta) = \cos(k\theta)$, hence $|T_k| \le 1$ on $[-1, 1]$. The same recurrence proves $T_k((s + s^{-1})/2) = (s^k + s^{-k})/2$ for $s > 1$. Choose

$$p(\lambda) = \frac{T_k((\lambda_+ + \lambda_- - 2\lambda)/(\lambda_+ - \lambda_-))}{T_k((\lambda_+ + \lambda_-)/(\lambda_+ - \lambda_-))}.$$

Its numerator is bounded by one on the spectral interval, and $p(0) = 1$. The denominator is at least $s^k/2$ for $s = (\sqrt{\kappa} + 1)/(\sqrt{\kappa} - 1)$, giving the result. If the spectrum is a singleton, K is a scalar multiple of the identity and the first residual direction solves the equation. □

For $0 < \epsilon < 1$, a sufficient iteration bound is $k \ge \frac{1}{2}\sqrt{\kappa}\log(2/\epsilon)$. Indeed $\log((1+t)/(1-t)) \ge 2t$ for $0 < t < 1$, by integrating $2/(1 - t^2) \ge 2$, with $t = \kappa^{-1/2}$. The independent finite-termination bound $k \le N$ can be used when smaller.

31.4 An explicit mesh-dependent condition estimate

Proposition 31.3. [K03] *For the square triangulations of Chapter 30, a stiffness matrix with $\alpha I \preceq A_T \preceq \beta I$ satisfies*

$$\lambda_{\min}(K) \ge \frac{\alpha}{12n^2}, \qquad \lambda_{\max}(K) \le 12\beta, \qquad \kappa(K) \le 144(\beta/\alpha)n^2.$$

Proof. For one triangle of area $h^2/2$, direct integration of the affine barycentric basis gives the mass matrix

$$\frac{h^2}{24}\begin{pmatrix}2&1&1\\1&2&1\\1&1&2\end{pmatrix}, \qquad h = 1/n.$$

Its smallest eigenvalue is $h^2/24$. Every free vertex belongs to at least two triangles, so the square's L^2 norm is at least $(h^2/12)$ times the sum of squares of its free nodal values. Poincaré and coercivity imply the lower stiffness bound.

For either triangle orientation, $(d_x v)^2 + (d_y v)^2 \le 4\sum_{i=1}^{3} v_i^2$. Thus its energy is at most $2\beta\sum v_i^2$. Each vertex belongs to at most six triangles, giving the upper bound 12β on the assembled quadratic form. The ratio gives the condition estimate. □

This sufficient constant is intentionally conservative. It is a proved bound on the chosen mesh family, not a measured condition number. The dual tensor has eigenvalues between $1/\beta$ and $1/\alpha$, so the same ratio estimate applies to its stiffness matrix.

The sparse matrix has $O(n^2)$ stored entries, since each vertex has a bounded mesh neighborhood. A matrix–vector product and each vector update therefore require $O(n^2)$ arithmetic operations. Combining the preceding bounds gives, at fixed ellipticity ratio, a sufficient exact-arithmetic solve cost $O(n^3 \log(1/\epsilon))$ and storage $O(n^2)$. Dense elimination of the $N \times N$ matrix requires $O(N^3) = O(n^6)$ arithmetic and $O(N^2) = O(n^4)$ storage. These are bounds for specified algorithms; optimized sparse direct methods are a distinct comparator and are included in the experiment.

31.5 Preconditioning and the measured ablation

A positive diagonal preconditioner D transforms the system to $D^{-1/2}KD^{-1/2}$. Applying Theorem 31.2 to that matrix gives its exact-arithmetic bound. The delivered Jacobi mode uses $D = \operatorname{diag} K$; the ablation uses $D = I$.

Jacobi scaling is not asserted to improve every positive-definite problem. Its benefit depends on the transformed spectrum. The paired experiment uses the same tensor model, mesh, initial affine fields, residual threshold, rounding format and exact certificate checker for the two modes. Iteration counts and measured times are retained even when they do not improve. A common code path is appropriate for an ablation, but it does not constitute an independent implementation; the separately assembled sparse-direct route and exact integral checker have different roles.

31.6 Precision limits and a useful stopping decision

For a fixed mesh, solving to a much smaller algebraic residual eventually ceases to improve the continuum bracket materially. The remaining width can be caused by unresolved interface singularities. Increasing only the dyadic precision also cannot remove that spatial error.

The correct stopping condition is the requested certificate width, for example $(U - L)/L \le \eta$. If rounding contributes materially, increase representation precision; if it does not, refine the mesh or choose a different conforming space. The supplied precision sweep distinguishes these mechanisms for two models. This is a measured decision aid with a theorem-level bound on the resulting field error. It is not an unproved a-priori mesh-complexity theorem for every discontinuous material.

31.7 Parallel and distributed interpretation

Triangle energy contributions are independent once the vertex fields are fixed. With P workers, an arithmetic work bound is $O(n^2)$ and a tree-reduction depth is $O(n^2/P + \log P)$, under a model with local access to each triangle's data. Exact rational additions have growing operand lengths; this bound counts arithmetic operations and must not be relabelled as bit complexity. A distributed implementation additionally incurs field partition, shared-vertex communication and reduction traffic.

Sparse conjugate-gradient iterations require matrix-neighbor communication and global dot products in a distributed-memory realization. Counting the number of iterations does not account for latency or overlap. No multi-node or GPU timing is reported here. The actual executed programs are CPU processes in the recorded assistant environment, not measurements of the user's workstation.

Chapter 32

New continuum experiments and exact elliptic certificates

32.1 What was executed and what a passing check means

This revision executed two new certificate families. The first verifies continuum capacity bounds from rational piecewise-affine potential and stream witnesses. The second verifies transcendental quadrilateral moduli from rational arithmetic–geometric traces or independently bounded integrals. Ordinary exact-arithmetic programs accepted the finite records. No proof assistant was used and no empirical material law was validated.

The continuum family contains 42 model–mesh instances: five scalar checkerboard contrasts $1, 2, 4, 16, 64$ and two rotated-tensor parameters $2, 16$, each on meshes $n = 2, 4, 8, 16, 32, 64$. Each initial generator instance has three recorded native runs; its delivered certificate contains the final recorded witness and is independently checked in Python rational arithmetic and C++ multiprecision rational arithmetic. The model matrices and terminal conditions are checked independently of the certificate's claimed bounds. Every scalar checkerboard interval also contains the proved positive root $C^2 = \kappa$.

A passing certificate establishes the inequalities for the stated continuum problem. It does not mean that a narrow interval was obtained, that the mesh is efficient, or that the original material has been measured. The high-contrast cases below deliberately retain broad bounds and expose incomplete numerical resolution.

32.2 Mesh convergence and high-contrast limitations

The following entries use 40-bit dyadic nodal witnesses. Relative gap means $(U - L)/L$ with the exact rational endpoints; displayed decimals summarize that ratio and are not replacement certificate endpoints. The finest mesh has 4,225 vertices and 8,192 triangles. Each free primal or dual solve has $n^2 - 1 = 4,095$ unknowns.

Model	Parameter	Gap at $n = 16$	Gap at $n = 32$	Gap at $n = 64$
checkerboard	1	0	0	0
checkerboard	2	0.00315616	0.00108311	0.000368861
checkerboard	4	0.025034	0.0109863	0.00483428
checkerboard	16	0.347561	0.213416	0.133699
checkerboard	64	2.03122	1.41691	1.02288
rotated	2	0.000933236	0.000290928	8.76641e-05
rotated	16	0.000766336	0.000239158	7.20809e-05

For contrast 64 the finest delivered scalar interval still has relative width exceeding one. The interface singularity is not resolved to a small relative tolerance by this mesh sequence. Increasing the arithmetic precision alone cannot repair that failure. The low-contrast and rotated cases have substantially smaller brackets under the same geometric subdivision, but no universal convergence exponent is inferred from six mesh values.

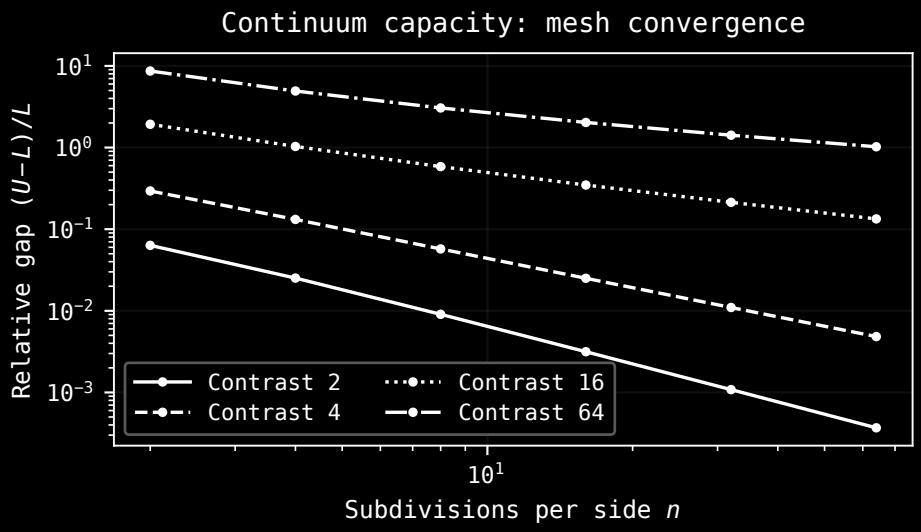


Figure 32.1: Exact capacity-certificate gaps on the supplied mesh sequence. Each plotted point comes from an accepted rational enclosure; the connecting lines do not constitute a continuum envelope in mesh size.

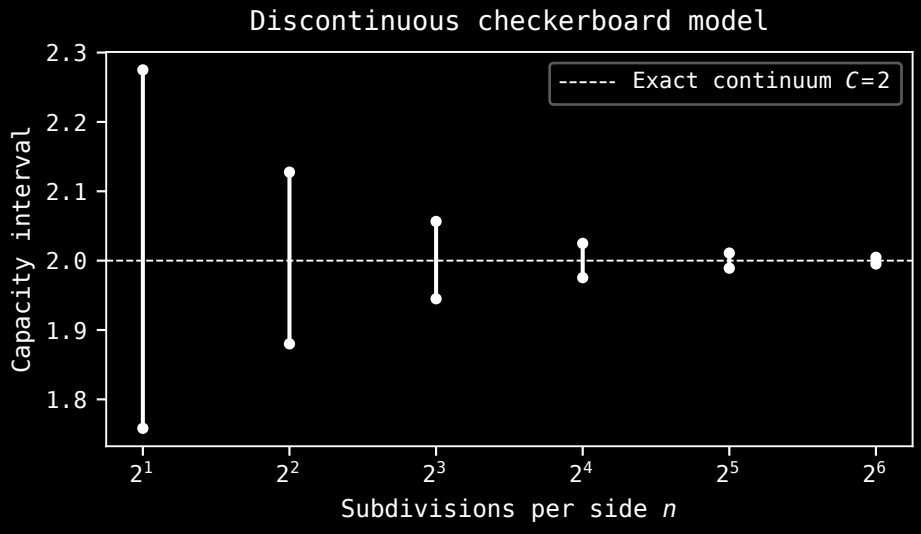


Figure 32.2: Upper and lower continuum witnesses for one scalar model. The known reference is proved by square duality, independently of the finite-element solve.

32.3 Separating rounding from spatial error

The precision sweep fixes two $n = 16$ geometries and changes the dyadic rounding depth. Ten rounded witness pairs were accepted by both rational checkers. The error plateau at high precision identifies the spatial/conforming-space limitation for these witnesses.

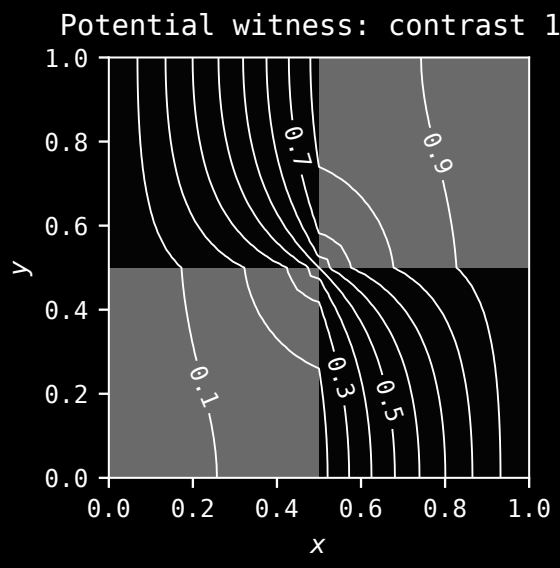


Figure 32.3: One delivered piecewise-affine potential on a discontinuous medium. Interpolation is the actual witness definition; the shading does not independently certify any pointwise PDE residual.

Model	4 bits	8 bits	16 bits	24 bits	40 bits
checkerboard_k4	0.901382	0.0280876	0.025034	0.025034	0.025034
rotated_k16	0.551558	0.00462073	0.000766385	0.000766336	0.000766336

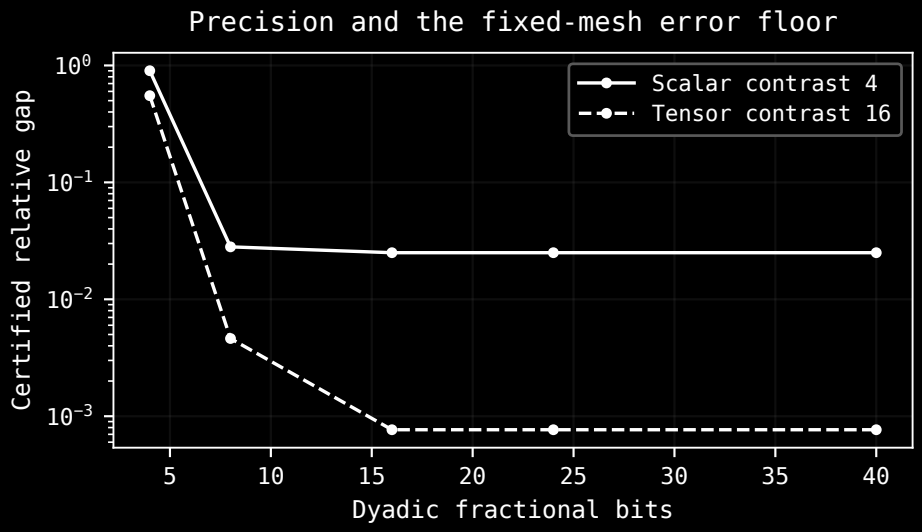


Figure 32.4: Fixed meshes under progressively finer rationalization. The certified width stops materially improving once nodal rounding is below the spatial-error scale.

The rational checkers do not test that a supplied field is a minimizer. One deliberately non-stationary interior perturbation is accepted after its true energy bounds are recomputed. This positive control distinguishes mathematical feasibility from solver convergence. Invalid model changes, terminal values and inconsistent bounds are rejected.

32.4 A matched solver comparison with an optimized control

The comparison uses nine model–mesh instances and four solution routes: native dense elimination, native unpreconditioned conjugate gradients, native Jacobi-preconditioned conjugate gradients, and independently assembled SciPy/SuperLU sparse factorization. The latter is an optimized sparse baseline, not a second name for the dense routine [23]. Each configuration has one discarded warm-up and five recorded repetitions. There are 36 configurations and 180 recorded paired primal–dual solves.

The next table gives the median time of the two linear solves, in milliseconds. Assembly, process startup, file output, rationalization and exact checking are excluded from this column and retained in separate raw fields. These exclusions are common to the displayed solve comparison; total-pipeline preference requires adding the appropriate costs. The input matrix model, free variables and output terminal conditions agree across methods. Every delivered method output is rationalized and checked by the same independently specified continuum inequality.

Model	n	Dense	CG	Jacobi CG	SuperLU
checkerboard 4	8	0.084316	0.032128	0.025959	0.24217
checkerboard 4	16	4.1967	0.37347	0.25926	0.94083
checkerboard 4	32	278.06	3.0765	2.0173	3.3064
checkerboard 16	8	0.082915	0.037156	0.025738	0.21091
checkerboard 16	16	4.3047	0.53564	0.2476	0.87331
checkerboard 16	32	277.01	5.409	1.9693	3.3841
rotated 16	8	0.082614	0.055835	0.042164	0.20941
rotated 16	16	4.3656	0.57712	0.40841	0.86688
rotated 16	32	275.74	5.0376	3.2686	4.2162

The arithmetic comparison is established by Chapter 31; this table records behavior of the particular implementations. Neither the measured ratio nor the conservative mesh-dependent theorem proves that Jacobi CG is optimal among all sparse solvers or preconditioners. Five observations describe this run's variation; they are not a universal confidence statement over machines, compiler settings or future workloads.

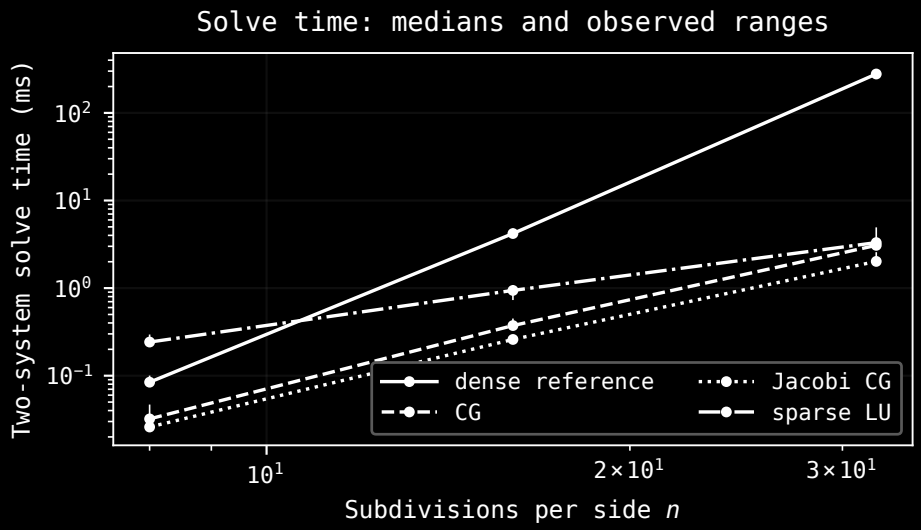


Figure 32.5: Solve-only medians for one model. The sparse-direct route is retained as a strong baseline; checker cost and assembly are separate terms.

32.5 Exact checking costs and implementation independence

The Python checker uses the standard arbitrary-precision integer and rational operations. The C++ checker uses Boost multiprecision rational arithmetic [24]. Each independently computes triangle energies and global pairing from the fixed model and actual witness. The native finite-element generator and the sparse reference assembly are different routes; identical variational functionals are intentional.

The C++ checker wall times include process startup; the native solver JSON times do not. Consequently dividing those two times without adjustment would compare different measurement boundaries. The result records retain all observed repetitions and checker timings so that a total-cost comparison can be reconstructed. The checkers establish inequalities relative to the written continuum theorem, rational implementation, parser, compiler and host arithmetic. None of those software layers has acquired formal kernel verification in this revision.

32.6 Quadrilateral certificates at large and small moduli

Eleven AGM cases use a target relative enclosure width of 10^{-20}, with 160-bit dyadic square-root bounds. The input parameters include 10^{24} and $1 + 10^{-12}$. Every accepted interval is rational; the two independent checkers use only positive rational comparisons and exact squaring. No transcendental library is needed to accept a trace.

R	Iterations	Lower endpoint	Upper endpoint
2	4	0.99999999	1.00000001
4	5	1.27926157	1.27926158
10	5	1.59887497	1.59887498
100	6	2.34681554	2.34681555
10^6	7	5.28015583	5.28015584
10^{12}	8	9.67776958	9.67776959
10^{24}	9	18.47299677	18.47299678
$1 + \frac{1}{10}$	5	0.61320251	0.61320252
$1 + \frac{1}{100}$	6	0.42553214	0.42553215
$1 + \frac{1}{1000000}$	7	0.18938833	0.18938834
$1 + \frac{1}{1000000000000}$	8	0.10332959	0.10332960

The printed endpoints are rounded outward to eight decimal places. They are intentionally wider than the exact stored intervals; the tolerance check uses the full rational records, not the table. Intervals whose printed lower and upper values differ in the last shown digit do not thereby have a 10^{-8} mathematical certificate gap.

For four moderate parameters a separate quadrature representation and an AGM trace both satisfy the same requested relative width 0.002. Their certified intervals intersect. The quadrature checker bounds the integrand on every cell of a complete rational partition and sums the interval contributions. Its rectangle-rule work grows much faster than the AGM trace length for these cases, but a cell and an iteration have different arithmetic costs.

R	AGM steps	Integral cells	AGM gap	Quadrature gap
2	2	512	5.5799e-05	0.00123492
4	2	512	0.000434393	0.001541
10	3	1024	1.19236e-06	0.00114339
11/10	3	1024	1.62173e-06	0.00118894

Across the two regimes there are 19 accepted quadrilateral certificates. Their audit additionally rejects 13 malformed or mismatched cases. These are finite checker tests; the universal soundness implication is the written theorem, and the classical AGM identity retains its attribution.

32.7 Finite audit counts and their limits

The analytical checking script passed 35 named checks, including matrix identities, metric derivatives, exact triangle integrals, rational constants, annular polynomial witnesses and a finite conjugate-gradient calculation. Twenty annular polynomial certificate cases verify an explicitly proved equality family. The continuum suites include 14 Python invalid-input rejections, 13 native invalid-input rejections, and a feasible-nonstationary modification accepted by both implementations. These populations overlap in purpose and are not added to a count of independent mathematical theorems.

The additional source arguments are individually identified in the claim register. Their universal proof status is distinct from agreement of the two rational programs. The mathematical code does not execute a new Lean, Rocq, Isabelle or HOL development, consistent with the commissioned phase.

32.8 Reproduction sequence and preserved historical evidence

Run the package-integrity checker before any replay. The reproduction runner compiles the C++ checkers, rechecks the delivered exact witnesses, runs the algebraic diagnostics and records output in a separate directory. An optional regeneration mode reconstructs the selected numerical witnesses; new timing observations are stored separately from historical records. The older network suite and its recorded results remain available and retain their original measurement scope.

The delivered native continuum benchmark has no physical or laboratory dataset, GPU run, ARM-native run, multi-node result or user-workstation measurement. It is a CPU experiment in the recorded x86-64 assistant environment. Its precise model, witness and checker remain useful independently of that hardware limitation. Those missing commissioned executions are not marked completed by this chapter.

Chapter 33

Executed calculations and reproducibility

33.1 The evidence obtained in this reconstruction

The computations in this chapter were executed for the retained revision-2 research manuscript. They are distinct from historical verification records in earlier volumes. The delivered supplement contains exact instances, sources, raw timing observations, hexadecimal native outputs and the result registers. No execution on the user's workstation, a GPU, an ARM machine or a multi-node cluster is asserted.

The measured host is a Linux x86-64 container whose processor record identifies an AMD EPYC 9V74. This identifies the reported execution environment; it does not establish exclusive access to all of the processor's physical resources. The native compiler was GCC 14.2.0 with `-std=c++17 -O2`. Python 3.13.5 was used for independent rational and symbolic checks and for document/figure tooling. The dependency versions are preserved in the supplement's environment record.

The universal mathematical arguments remain in the preceding chapters. Finite-instance arithmetic verifies those instances and helps expose implementation defects. It does not replace a proof of the general statement or establish a verified compiler or arithmetic library.

33.2 Exact potential–flow certificates

Nine rational network witnesses were accepted by the Fraction-based checker. Eight have matching lower and upper bounds: a single edge, a two-edge series chain, two parallel two-edge branches, and square grids with side lengths two through six. The ninth is a deliberately nonoptimal but feasible witness on the three-by-three grid. An interior potential coordinate is changed by 1/20 and its upper energy is recomputed, while the original feasible unit flow is retained. The resulting positive bracket tests that the checker recognizes certified approximation independently of exact optimality.

A separate native C++ checker uses `boost::multiprecision::cpp_rational` and a distinct line-based encoding. It accepted all nine corresponding witnesses. The two implementations do not share the same rational library or solver code. They do share the mathematical specification and the supplied instance data. Their agreement is useful implementation evidence, not proof of independence at every compiler, machine or conceptual layer.

The Python checker rejected twelve altered JSON witnesses. The native checker rejected eleven altered textual witnesses, including changes to signs, terminal values, conductances, denominators,

indices, claimed bounds and trailing input. The contracts of the two parsers differ, so their mutation counts are not interchangeable. Every mutation and diagnostic is recorded. A rejected malformed witness does not establish resistance to every possible malformed input.

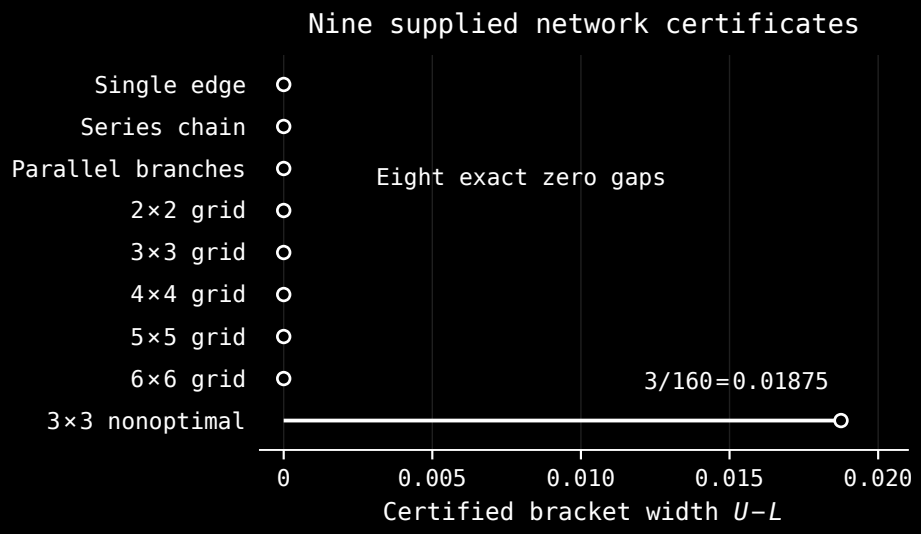


Figure 33.1: Eight exact optimizer witnesses have zero gap. A feasible nonoptimal witness has a positive gap and is still correctly accepted. The horizontal axis reports the actual bracket widths in the supplied instance set.

33.3 Exact resource-design witnesses

The budget theorem 19.1 was instantiated on five grids, from four to thirty-six vertices, with rational positive edge lengths and budget twelve. A shortest-path construction generated distance labels, conductance allocations and unit flows. A native C++ verifier checked the Lipschitz distance inequalities, budget, flow support, conservation and the matching upper/lower objective values. All five passed. Five deliberately altered witnesses were rejected.

This check does not require trusting the shortest-path algorithm's assertion that its path is shortest. Feasible distance labels give a universal upper bound, and the displayed flow gives a matching lower bound. Their equality is the certificate. The generation algorithm can therefore be changed without changing the logical acceptance rule.

The implemented design model admits zero conductances. It is not the same schema as the connected strictly positive network certificates above. This distinction appears in both the written theorem and the parser. A zero-conductance edge may carry no flow. The verifier rejects a nonzero current assigned to such an inactive edge.

33.4 Chain timing experiment

Five chain sizes, each with conductances $c_i = ((i \bmod 7) + 1)/2$, were processed by three native routines: dense Gaussian elimination, tridiagonal elimination, and the cumulative-resistance formula. Each configuration has seven repetition means. A repetition averages two dense solves or one thousand of the linear-time routines. One untimed call precedes each configuration. Allocation and matrix construction inside each function are included in the timing; compilation and input-vector construction are outside it.

The dense routine is intentionally transparent and does not exploit tridiagonal sparsity. Its

Edges	Dense (μs)	Tridiagonal (μs)	Series (μs)	Ratio
16	0.666	0.233	0.060	3.88
32	4.026	0.477	0.119	4.03
64	29.720	0.931	0.254	3.66
128	238.610	1.882	0.541	3.48
256	2032.587	4.128	1.087	3.80

Table 33.1: Median measured time per solve. The last column is tridiagonal time divided by series-formula time. These are finite implementation measurements, not universal hardware speedup factors.

much greater cost illustrates a representation error, not a fair claim that all conventional solvers are cubic on a chain. The tridiagonal method is the strong structural control; it has the same linear asymptotic arithmetic order as the series formula. The smaller observed constant for the formula is specific to the supplied code, allocator, compiler and workload.

The raw CSV contains fifteen configurations and 105 recorded repetition means. No confidence interval from repeated independent hardware sessions is reported. At microsecond and submicrosecond scales, host scheduling, clock granularity, caches and compiler optimization are important. The repeated inner loops reduce timer overhead but do not remove every source of uncertainty.

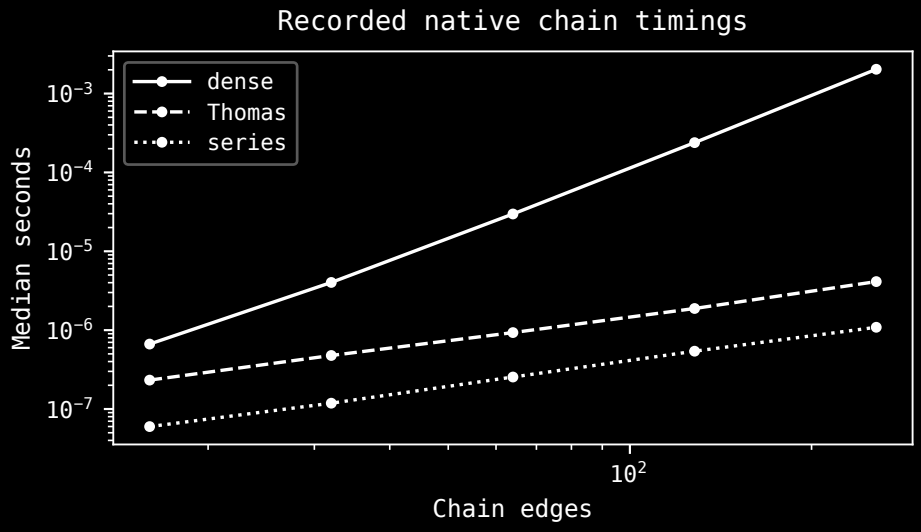


Figure 33.2: Native timings for the three matched chain calculations. Both the series formula and the tridiagonal baseline have linear arithmetic complexity; the dense implementation discards this sparsity.

33.5 Independent exact forward-error evaluation

The timed routines also exported their output vectors as hexadecimal floating-point values. A separate program reconstructs these values as exact binary rationals and generates the mathematical chain solution using Fraction arithmetic. It checks every coordinate against the exact solution. Across the fifteen configurations, 1,503 coordinates were compared, and every configuration met the declared absolute tolerance 10^{-10}.

This is stronger than agreement between two floating-point methods: a common error in

both would still be detected by the exact reference. The original timing CSV's `max_error` field records discrepancy from the floating series routine and must not be interpreted as exact error. The separate `chain_accuracy.json` file contains the exact-reference errors and residuals. In particular, a zero discrepancy of the series routine from itself is not treated as zero mathematical error.

The hexadecimal serialization preserves the reported binary value without decimal-rounding ambiguity in this representation. The analysis assumes the documented parsing and arbitrary-precision arithmetic behavior. It is not a machine-level verification of the executable or the processor.

33.6 Symbolic identities and selected source problems

The general analytical suite contains ninety-five named checks. Its populations include rational cylinder reciprocal/energy identities, perturbation equality examples, torus marking changes, Schur-complement identities and elementary physical response formulas. Ten of these are floating-point diagnostic evaluations of scalar functions; they are not validated interval enclosures. The evidence type is stored per entry.

A separate selected-problem suite performs ten exact finite checks associated with Chapter 17: a negative two-point kernel witness, feature Gram identities, a Gauss–Markov covariance difference, singular Nyström identities, rank-exact recovery and integer storage counts. These checks confirm concrete instances. The complete written source-problem arguments contain their general proofs.

The scientific checks are counted separately from file-integrity checks, source-location checks and PDF layout checks. Adding these heterogeneous counts would not produce a meaningful theorem-verification percentage.

33.7 Reproduction order

Begin by checking the archive manifest. Compile the three C++ sources with the recorded language standard and optimization flags. Run the rational witness generators and the native checkers; run the symbolic suite and the selected-problem checks. The chain benchmark may be rerun into a fresh output directory. Retain the original timing CSV and distinguish new measurements by environment and run directory.

Finally regenerate the figures from the chosen result records and compile the LaTeX manuscript. Figure regeneration is not a rerun of every scientific experiment. Conversely, changing numerical results without regenerating their tables and graphics leaves the publication inconsistent. The build instructions separate these operations, and the manifest identifies the specific payload associated with this PDF.

New timing observations can differ from the recorded medians without invalidating the mathematical formula or the archive. A failed exact identity or certificate test is different: it requires investigation before the affected claim is reused. A missing compiler or unavailable library must produce an explicit failure, never a success placeholder.

33.8 Unperformed experimental programmes

There is no GPU, ARM-native or multi-node benchmark in this review. There is no original national-laboratory measurement analysis, nuclear or high-energy experimental implementation, original-MNIST experiment, production cryptographic evaluation or quantum-hardware run. No 2560-by-1440, sixty-frame-per-second animation has been completed. The complete semantic

question-bank adjudication is also incomplete. Companion-edition obligations remain tracked separately.

These are unfinished items of the earlier commission, outside the current scholarly release scope. The successful finite network and scalar experiments cannot be used to close them. The historical records remain distinct from the fresh replay documented in the supplement.

Chapter 34

Worked synthesis: geometry, uncertainty and certificates

These sixteen investigations continue VI01–VI24. They make the output of a variational calculation usable: a conductance must have stated terminals, an interval must refer to one fixed problem, and an approximation must distinguish feasibility from accuracy. The calculations are authored expository exercises, without a research-priority claim. The square-model arguments use the spaces and pairing proved in Chapter 30; their conclusions do not require the conditional canonical geometry of Chapter 8.

34.1 VI25. An anisotropic rectangle with insulating sides

Problem. On $(0,w)\times(0,h)$, let $A=\operatorname{diag}(a,b)$ with $a,b,w,h>0$. Prescribe potential zero on the left and one on the right, leaving the horizontal sides insulating. Determine the capacity and the energy of an applied voltage V.

The candidate is $u=x/w$. For any admissible $v=u+e$, the trace of e vanishes on the vertical sides. Integration in x gives $\int e_x=0$, initially for smooth e and then by the trace-compatible approximation. Hence

$$E_A(v)=\frac{ah}{w}+\int (ae_x^2+be_y^2).$$

The second term is nonnegative and vanishes only for a constant e, whose zero traces make it zero. Thus $C=ah/w$, independently of b. At voltage difference V, the unhalved energy is CV^2; the electrostatic stored-energy model instead has $CV^2/2$ after reinstating the appropriate material and dimensional factors. The independence from b belongs to these terminals and constant diagonal tensor, not to arbitrary anisotropic boundary problems.

34.2 VI26. Rotating a two-phase laminate

Problem. Divide the unit square into two equal-width layers of scalar conductivities $a,b>0$. Compare layers perpendicular and parallel to the left–right current direction.

For perpendicular layers, a constant flux j implies slopes $j/a, j/b$. The total voltage is $j(1/(2a)+1/(2b))=1$, so

$$C_{\text{series}}=\frac{2ab}{a+b}.$$

The continuous piecewise-affine potential with these slopes has the stated traces. Its flux is constant and its cross term with every zero-terminal perturbation vanishes, proving minimality.

For parallel layers, $u = x$ has zero normal flux on the horizontal sides and no normal jump across the horizontal interface. Its energy is

$$C_{\text{parallel}} = \frac{a+b}{2}, \qquad \frac{C_{\text{parallel}}}{C_{\text{series}}} = 1 + \frac{(a-b)^2}{4ab}.$$

The conductances agree precisely for $a = b$. With $a = 1, b = 4$ the two values are $8/5$ and $5/2$, and the ratio is $25/16$. The same material fractions therefore do not determine the same directional response.

34.3 VI27. A sharp interval with uncertain geometry

Problem. In VI25 suppose independently $a \in [2,3]$, $h \in [1,2]$, $w \in [3,4]$; b remains any positive constant. Determine the exact capacity range.

Monotonicity in each factor of ah/w yields

$$\frac{1}{2} \le C \le 2.$$

The endpoints occur at $(a,h,w) = (2,1,4)$ and $(3,2,3)$. Along the straight segment between them, a and h increase while w decreases, so continuity attains every intermediate value. This proves the exact range. Correlation constraints can remove either corner and change sharpness. A floating-point sample cannot replace the corner and continuity arguments.

34.4 VI28. An off-diagonal tensor changes the boundary problem

Problem. For the unit-square terminals, take the constant tensor $A = \left(\begin{smallmatrix} a & c \\ c & b \end{smallmatrix}\right)$ with $a > 0$, $b > 0$, $ab - c^2 > 0$. Bound $C_x[A]$ without solving for the potential.

The admissible potential $v = x$ has energy a. The admissible stream function $\psi = y$ has dual energy $b/(ab - c^2)$, because $B = A/\det A$. The pairing theorem gives

$$a - \frac{c^2}{b} \le C_x[A] \le a.$$

For $a = 2, b = 3, c = 1$ this is $[5/3, 2]$. When $c = 0$, both endpoints coincide. When $c \neq 0$, the field $A\nabla x = (a, c)$ has nonzero normal flux on the insulating sides, so admissibility of x alone does not make it the minimizer. More explicitly choose $e = x(1-x)y$, which has zero vertical traces. Then $\int \nabla e^T A \nabla x = c/6$. Choosing the sign of a sufficiently small perturbation $x + te$ opposite to c strictly decreases the energy. Thus the upper bound is strict for $c \neq 0$. This is a boundary-condition test that a bulk residual alone would miss.

34.5 VI29. Removing diagonal anisotropy by changing the domain

Problem. Derive VI25 by a coordinate transformation and identify its scalar prefactor.

Set $x = \sqrt{a}\,\xi$, $y = \sqrt{b}\,\eta$. The gradient factors cancel the coefficients while the Jacobian contributes $\sqrt{ab}$:

$$E_A(v) = \sqrt{ab} \int_{(0,w/\sqrt{a})\times(0,h/\sqrt{b})} (v_\xi^2 + v_\eta^2)\,\mathrm{d}\xi\,\mathrm{d}\eta.$$

The transformed rectangle has capacity $(h/\sqrt{b})/(w/\sqrt{a})$. Multiplication by $\sqrt{ab}$ recovers ah/w. Dropping the Jacobian would give an incorrect material scale. The transformation is an anisotropic linear change of variables; it is conformal only when $a = b$.

34.6 VI30. The exact residual behind a primal–dual gap

Problem. In the square model let v, ψ be admissible, $U = E_A(v)$, $F = E_B(\psi)$ and $L = 1/F$. Give a squared-norm interpretation of $U - L$.

With the rotation R of Chapter 30, expand the square and use the exact pairing $\int \nabla v \cdot R\nabla\psi = 1$:

$$\int \left| A^{1/2}\nabla v - F^{-1}A^{-1/2}R\nabla\psi \right|^2 = U - \frac{2}{F} + \frac{F}{F^2} = U - L.$$

Equality means equality of these two fields almost everywhere. A small gap measures their disagreement in the displayed norm. The identity requires admissible traces; if the pairing is $p \neq 1$, the cross term becomes $-2p/F$ and the same claimed certificate no longer follows.

34.7 VI31. A field-error bound from an interval

Problem. Let u be the exact minimizer and let a certificate provide $0 < L \le C \le U = E_A(v)$. Bound the error of v without knowing u.

Stationarity gives $E_A(v) = C + E_A(v - u)$, so

$$E_A(v - u) \le U - L, \qquad \|\nabla(v - u)\|_2 \le \sqrt{(U - L)/\alpha}$$

when $A \succeq \alpha I$, $\alpha > 0$. On the unit square the zero-terminal Poincare inequality also bounds $\|v - u\|_2$ by this gradient bound. If $[L, U] = [99/100, 101/100]$ and $\alpha = 1/2$, the gradient bound is $1/5$. No pointwise or maximum-norm error follows without a further regularity estimate.

34.8 VI32. Which central estimate minimizes relative error?

Problem. Only $C \in [L, U]$ is known, with $0 < L \le U$. Choose a reported value z minimizing the worst-case relative error $|z - C|/C$.

Any z outside $[L, U]$ is improved by moving to the nearer endpoint. Within this interval the worst errors occur at its endpoints, so minimize $\max(z/L - 1, 1 - z/U)$. The first expression increases with z and the second decreases. Equalizing them gives

$$z_* = \frac{2LU}{L + U}, \qquad \varepsilon_* = \frac{U - L}{U + L}.$$

Thus $[2, 3]$ yields $z_* = 12/5$ and guaranteed relative error $1/5$. For absolute error, the same endpoint argument instead gives $(L+U)/2$ and radius $(U-L)/2$. The error convention determines the best reported center.

34.9 VI33. Reciprocation preserves a relative certificate radius

Problem. Convert a conductance interval to resistance and compare its minimax relative radius.

Since inversion decreases on the positive real line,

$$R \in [1/U, 1/L], \quad \frac{1/L - 1/U}{1/L + 1/U} = \frac{U - L}{U + L}.$$

The optimal relative center for R is $2/(L + U)$, whereas the reciprocal of the optimal relative center for C is $(L + U)/(2LU)$. These agree only when $L = U$, since equality is equivalent to $(U - L)^2 = 0$. Interval reciprocation is exact; reciprocating an optimally chosen point estimate solves a different minimax problem.

34.10 VI34. Combining independent certificates

Problem. Two admissible witness pairs for exactly the same model produce $[L_1, U_1]$ and $[L_2, U_2]$. What conclusion can be combined?

The common value belongs to their intersection,

$$C \in [\max(L_1, L_2), \min(U_1, U_2)].$$

For $[1, 3/2]$ and $[5/4, 7/4]$ the result is $[5/4, 3/2]$. If the intersection is empty, at least one claimed enclosure or the shared-model premise is wrong. Averaging disjoint intervals cannot repair that contradiction. Matching geometry, tensor field, terminal convention and units is therefore a prerequisite for intersection. Two different meshes may certify the same continuum model, but two different piecewise coefficient fields need an additional comparison theorem.

34.11 VI35. A coefficient uncertainty bound valid for every admissible field

Problem. Suppose $0 < \delta < 1$ and two uniformly positive tensor fields on the same square satisfy $(1-\delta)A \preceq \widetilde{A} \preceq (1+\delta)A$ almost everywhere. Propagate a capacity interval for A.

For every admissible v, the same inequalities hold between the quadratic energies. Taking the infimum over the identical feasible space yields

$$(1-\delta)C_x[A] \le C_x[\widetilde{A}] \le (1+\delta)C_x[A].$$

Consequently $C_x[A] \in [L, U]$ implies $C_x[\widetilde{A}] \in [(1-\delta)L, (1+\delta)U]$. For $[L, U] = [2, 3]$ and $\delta = 1/10$, the interval is $[9/5, 33/10]$. Scalar multiples of A show that the multiplicative constants cannot be uniformly improved. Entrywise bounds on matrix entries are not automatically Loewner-order bounds and cannot be inserted here without checking the quadratic forms.

34.12 VI36. What nested approximation spaces actually guarantee

Problem. Let admissible finite-dimensional potential spaces and stream spaces be nested under refinement. Prove monotonicity of their optimal capacity enclosures.

Write $U_n = \min E_A(v)$ in the nth potential space and $F_n = \min E_B(\psi)$ in the nth stream space. Inclusion of feasible sets gives $U_{n+1} \le U_n$ and $F_{n+1} \le F_n$. Therefore $L_n = 1/F_n$ increases and

$$[L_{n+1}, U_{n+1}] \subseteq [L_n, U_n].$$

All F_n are positive by their terminal traces. This statement concerns exact minimizers on nested spaces representing the same coefficient field. Arbitrary solver iterates need not obey the monotonicity. Without a density theorem, nesting alone also does not prove that the intervals converge to a point; a constant sequence of spaces is a counterexample to that inference.

34.13 VI37. Exact admissibility of a rational tensor

Problem. Decide when the symmetric rational tensor $A = \left(\begin{smallmatrix} a & c \\ c & b \end{smallmatrix}\right)$ is positive definite, using no floating-point eigenvalues.

For $a > 0$, complete the square:

$$as^2 + 2cst + bt^2 = a(s + ct/a)^2 + \frac{ab - c^2}{a}t^2.$$

It is positive for every nonzero (s, t) exactly when $ab - c^2 > 0$. Necessity of $a > 0$ follows by taking $(s, t) = (1, 0)$; necessity of the determinant condition follows by setting $s = -ct/a$, $t \neq 0$. For $(a, b, c) = (2, 3, 1)$ the determinant is 5 and positivity is exact. With $(1, 1, 1)$ the determinant is zero and the nonzero vector $(1, -1)$ has zero energy. Such a tensor must be rejected by the strictly elliptic certificate format, even if a numerical solver returns an output.

34.14 VI38. A complete rational network certificate by hand

Problem. A path s–m–t has conductances $2, 3$, and a direct edge s–t has conductance 1. Set $v_s = 0, v_t = 1$. Find matching potential and unit-flow certificates.

The energy is $2v_m^2 + 3(1 - v_m)^2 + 1$. Completing the square gives

$$E(v) = \frac{11}{5} + 5(v_m - 3/5)^2.$$

Thus $C = 11/5$. Orient edges from s to m to t and directly from s to t, toward increasing potential. The optimized currents are $6/5, 6/5, 1$. Dividing them by C gives the unit flow $(6/11, 6/11, 5/11)$, with conservation at m. Its resistance-weighted energy is

$$F = \frac{(6/11)^2}{2} + \frac{(6/11)^2}{3} + (5/11)^2 = \frac{5}{11}.$$

The bounds $1/F$ and $E(v)$ coincide exactly. A feasible potential $v_m = 1/2$ instead gives $U = 9/4$, still a valid upper witness. The discrepancy $9/4 - 11/5 = 1/20$ equals $5(1/2 - 3/5)^2$, displaying the exact potential-error identity.

34.15 VI39. Duality returns the original tensor

Problem. In two real dimensions define $\mathcal{D}(A) = R^T A^{-1} R$ for a symmetric positive tensor. Calculate $\mathcal{D}(\mathcal{D}(A))$ and its scaling law.

Direct inversion gives $\mathcal{D}(A) = A/\det A$. Its determinant is $1/\det A$, so

$$\mathcal{D}(\mathcal{D}(A)) = A, \qquad \mathcal{D}(tA) = t^{-1}\mathcal{D}(A) \quad (t > 0).$$

The conductance identity also swaps the terminal direction: $C_y[\mathcal{D}(A)] = 1/C_x[A]$. Applying it twice restores both the tensor and the terminals. Using A^{-1} without the rotations generally changes which directions are paired and is not this statement.

34.16 VI40. Separating material uncertainty from solve error

Problem. A certificate gives nominal $C \in [99/100, 101/100]$. Material information bounds the physical coefficient by $(1 \pm 1/20)$ times the nominal coefficient in Loewner order. Quantify the total capacity interval and compare it with perfect solution of the nominal problem.

VI35 gives

$$C_{\text{physical}} \in [1881/2000, 2121/2000].$$

Its absolute midpoint is $2001/2000$ and radius $3/50$. If the nominal capacity were known exactly to be 1, the remaining interval would still be $[19/20, 21/20]$, of radius $1/20$. Thus solving

the nominal PDE more accurately cannot remove the material uncertainty. Conversely, this bound models uncertainty in coefficients alone: an unknown domain, contact impedance or time dependence would need additional hypotheses. No measured data are introduced by this synthetic example.

34.17 Executable checks and the boundary of the examples

The supplement's `code/check_synthesis.py` checks the rational numerical examples and symbolic energy identities, including negative controls for singular tensors, broken flow conservation and inconsistent interval intersections. Its finite checks support the displayed instances. The arguments above establish the parameterized conclusions under their stated hypotheses. Neither the script nor these local models resolves the conditional global canonical geometry.

Chapter 35

Source obligations and formalization groundwork

35.1 Source locations and mathematical completion are separate

The source-ordered treatment retains all eight main subsections, the thirteen nested headings, the numbered sequence 6.1–6.45, and the four named unnumbered principles. A resolved locator is not a completed theorem. The machine-readable source register gives an interpretation, a proof location and a remaining boundary for each numbered or named entry.

The additional supporting proofs in this revision change particular statuses. They do not close the earlier, broader commission, which is outside the current release scope. The register retains the original general targets and records the precise component discharged, without converting a restricted proof or a counterexample into an unsupported general theorem.

35.2 Definitions, circuit laws and ring geometry

For §6.1 the quotient and width conventions, monotonicity, overflowing and conformal covariance remain established in the stated curve classes. The literal smooth-density claim is refuted by the null-support example. Chapter 23 now proves the modulus-null criterion, weak-admissibility equivalence, convergence outside an exceptional family, attainment in the weak class and the sharp quadratic deficit estimate. A smooth approximation result is separately proved for regular condenser connecting families.

The finite series and parallel laws and their disjoint-support equality retain their complete arguments. Countable extensions have stated finiteness and approximation conditions. These mathematical resistance/conductance rules do not imply an arbitrary time-dependent circuit model.

Chapter 24 supplies explicit rough-continuum tube and logarithmic bounds, inscribed round annuli, normalized diameter decay, a two-sided chordal pinching estimate, the real-interval bounds and a general nested-continuum divergence argument. Chapter 25 supplies a Laurent-coefficient proof of the stronger exponential area inequality for positive finite inner area and finite outer area, together with a sharp quantitative deficit. Singleton and zero-area cases remain outside ratios for which their denominators vanish.

Marked quadrilateral existence is established by an explicit holomorphic primitive in Chapter 29. Its stronger half-plane modulus inequalities hold for every $R > 1$ with the specified marking. The AGM transformation and rational enclosures are proved and implemented. The torus formula and its homological conventions remain explicit.

35.3 Dirichlet duality, crossing and transformation

Chapter 28 gives the variational, weak-regularity, reflection, critical-point, flux-coordinate and separating-multicurve proof for compact surfaces with analytic boundary and complete boundary-component terminal partitions. This is a completed regular setting for source Theorem 6.30. It does not cover arbitrary metric-measure boundaries or singular/nodal models without an additional construction.

The Non-Crossing width-product estimate now has a proof using the height images of all portions of a curve inside the support; no unsupported contiguous-crossing inference is needed. The literal equality assertion in Exercise 6.31 has a planar annulus–quadrilateral counterexample and a separate torus example. The restricted two-quadrilateral rigidity needed for the canonical rectangles remains unproved here. The product-lamination overlap argument uses its specified transverse measure and open-leaf convention.

The lower holomorphic-image inequality is developed with countably many inverse charts and an explicit critical-value argument for locally rectifiable parametrized curves. The finite-multiplicity upper inequality and power-map equality cases retain their hypotheses. The multi-terminal equality under a finite covering follows from the regular condenser proof and the area formula. The complete arbitrary irregular branched multi-component source extension is not claimed from the regular case.

35.4 Maximal annuli and the unresolved canonical geometry

Chapter 26 proves the equal-axis trace separation, embedded open collar and the resulting covering-modulus estimate. Chapter 27 proves noncollapse and attainment of maximal embedded annular modulus in the indicated Riemann-surface classes, relative to the identified earlier uniformization, covering and normal-family dependencies. The flat and parabolic alternatives are treated separately.

The nonperipheral finite-modulus Jenkins–Strebel uniqueness development is still open. The puncture and torus counterexamples settle only the corresponding unrestricted uniqueness claims. Existence is not substituted for uniqueness. The planar buffer-separation proof, boundary-gap construction and restricted equality rigidity needed to establish the full canonical arc family still require completion. The elementary-limit-set exceptions are now identified before the Cantor-gap assertions are used.

The ideal-triangulation count is proved under explicit extendability and essential/nonparallel assumptions. This combinatorial proof does not establish that the proposed canonical family satisfies those geometric premises. The exact positive-part weight operation and its interval propagation remain valid independently of those global premises.

35.5 What the executable evidence establishes

The finite-network witnesses are retained, and the new complementary-energy implementation connects finite rational triangle data to an actual two-dimensional continuum problem. Forty-two model–mesh cases, ten precision variants and 36 method-comparison outputs have accepted potential–stream certificates. Eleven high-accuracy AGM cases and eight matched AGM/quadrature cases have accepted quadrilateral certificates. The benchmark has 180 recorded paired solves; repeated observations are not separate theorems.

A certificate checker verifies its finite instance relative to the corresponding written soundness theorem and explicit arithmetic dependencies. Its test count cannot establish that theorem, the universal completeness of a certificate format, or the correctness of every compiler and

hardware layer. The finite implementation's input caps are stated separately from the unbounded certificate-family completeness arguments.

35.6 Prospective formalization and remaining source screening

The theorem register identifies every labelled mathematical proposition, source file, written argument, current proof scope, dependencies, executable links and future formalization route. New proof-assistant implementation remains deferred. No Lean, Rocq, Isabelle or HOL kernel acceptance is asserted.

The existing complete-register screen remains lexical: 5,905 source-specific identifiers, four individually developed selected questions, 108 further lexical candidates, 555 entries without imported prose, and 5,238 nonmatches whose semantic relevance is unadjudicated. This revision does not increase that screening-completion claim. The forty authored investigations are a different population: twenty-four inherited cases and sixteen worked synthesis cases. Full semantic selection of the 5,905-question register remains outside this release scope; the lexical screen is not relabelled a completed relevance review.

35.7 Publication status

This is revision 3 of the scholarly public edition, prepared under the revised Volume V completion benchmark. The public bundle contains the reading editions, scientific supplement and Markdown website assets. The internal edition adds teaching modules, worked assessments and proposed subject mappings. Its undownloaded commission brief was unavailable, so no private requirements have been invented or declared fulfilled.

Global quadratic-differential/canonical geometry and arbitrary rough-boundary extensions retain their explicit mathematical limitations. Additional specialist applications, measured-data and target-machine campaigns, complete source-question adjudication, animations and custom licensing from the previous commission are outside this release contract. Their historical ledger is retained without retroactive completion claims.

Integrity, readability and reproducibility checks concern the delivered objects. No research-priority, independent peer-review, formal-kernel or complete original-commission acceptance follows from them. Identical mathematics in another theme or companion edition remains one mathematical development.

Bibliography

[1] M. Lyubich, *Conformal Geometry and Dynamics of Quadratic Polynomials, volumes I–II*, supplied 702-page manuscript, Section 6, pp. 119–133; Sections 2–5, 7 and 49 as explicitly identified dependencies.

[2] PHI research collective, *Conformal Structures and Function Spaces*, PHI Technical Monographs, Volume I, supplied public edition. Statements and assumptions must be matched before reuse.

[3] PHI research collective, *Uniformization and Geometric Analysis*, PHI Technical Monographs, Volume V, supplied public edition. Its recorded formal-verification status is not changed by citation here.

[4] M. Badger, "Beurling's criterion and extremal metrics for Fuglede modulus," *Annales Academiae Scientiarum Fennicae Mathematica* 38 (2013), 677–689. DOI: 10.5186/aasfm.2013.3826. Primary reference for general measure-system modulus; the present null-support example is proved directly.

[5] K. Rajala, "Extremal length and duality," *Expositiones Mathematicae* 43 (2025), article 125634. DOI: 10.1016/j.exmath.2024.125634. Context for the broader duality and metric-surface programme; not a claim of complete incorporation of its proofs.

[6] N. Albin, J. Clemens, N. Fernando and P. Poggi-Corradini, "Blocking duality for p-modulus on networks and applications," *Annali di Matematica Pura ed Applicata* 198 (2019), 973–999. DOI: 10.1007/s10231-018-0806-0. Prior-work context for network modulus and duality.

[7] J. Liu, "Jenkins–Strebel differentials with poles," *Commentarii Mathematici Helvetici* 83 (2008), 211–240. DOI: 10.4171/CMH/123. A supporting-theory source; no complete proof reconstruction is asserted here.

[8] E. Arbarello and M. Cornalba, "Jenkins–Strebel differentials," *Rendiconti Lincei – Matematica e Applicazioni* 21 (2010), 115–157. DOI: 10.4171/RLM/564. Reference for the pending global quadratic-differential development.

[9] D. N. Alsheikh and T. Katsaounis, "A posteriori error estimators for discontinuous Galerkin method for diffusion problems, based on the hypercircle method," *Arabian Journal of Mathematics* 11 (2022), 407–426. DOI: 10.1007/s40065-022-00386-w. Context for complementary-energy estimates; the finite identity used here is derived explicitly.

[10] P. Christiano, J. Kelner, A. Madry, D. Spielman and S.-H. Teng, "Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs," *Proceedings of STOC* (2011), 273–282. Reference for the distinction between electrical-flow primitives and a complete maximum-flow algorithm.

[11] A. M. Polyakov, "Quantum geometry of bosonic strings," *Physics Letters B* 103 (1981), 207–210. DOI: 10.1016/0370-2693(81)90743-7. The classical Weyl-invariance calculation here does not reproduce the complete quantum construction.

[12] S. W. Hawking, “Particle creation by black holes,” *Communications in Mathematical Physics* 43 (1975), 199–220. DOI: 10.1007/BF02345020. The local Euclidean regularity model in the text has a narrower evidentiary scope.

[13] N. Seiberg and E. Witten, “Electric-magnetic duality, monopole condensation, and confinement in N=2 supersymmetric Yang–Mills theory,” *Nuclear Physics B* 426 (1994), 19–52; erratum 430 (1994), 485–486. DOI: 10.1016/0550-3213(94)90124-4. Context for physical period geometry; not a completed physical derivation in this review.

[14] M. Mohri, A. Rostamizadeh and A. Talwalkar, *Foundations of Machine Learning*, second edition, 2018. Exercises 6.7, 6.11 and 15.4; supplied source-page extracts control the inspected statements.

[15] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, second edition, corrected twelfth printing, 2017. Exercise 3.3 and its stated linear-model setting.

[16] T. Adamowicz and N. Shanmugalingam, “Non-conformal Loewner type estimates for modulus of curve families,” *Annales Academiae Scientiarum Fennicae Mathematica* 35 (2010), 609–626. DOI: 10.5186/aasfm.2010.3538. Broader modulus context; the planar estimates in this manuscript have their own displayed proofs.

[17] L. Keen, “Collars on Riemann surfaces,” in *Discontinuous Groups and Riemann Surfaces*, Annals of Mathematics Studies 79 (1974), 263–268. Reprint DOI: 10.1515/9781400881642-021.

[18] A. Basmajian, “Generalizing the hyperbolic collar lemma,” *Bulletin of the American Mathematical Society* 27 (1992), 154–158; arXiv:math/9207211. Context for the classical collar theorem and its extensions.

[19] G. P. Ortiz and W. L. Mochán, “Keller’s theorem revisited,” *New Journal of Physics* 20 (2018), 023028. DOI: 10.1088/1367-2630/aaa7e1. The static finite-square identity in this manuscript is proved with its boundary conditions; no finite-frequency conductivity identity is inferred.

[20] G. W. Milton, *The Theory of Composites*, SIAM, 2022, Chapter 3, duality transformations. DOI: 10.1137/1.9781611977486.ch3. Prior-work context for two-dimensional conductivity duality.

[21] M. R. Hestenes and E. Stiefel, “Methods of conjugate gradients for solving linear systems,” *Journal of Research of the National Bureau of Standards* 49 (1952), 409–436. DOI: 10.6028/jres.049.044. The exact-arithmetic proof and the present matrix bounds are given in the text.

[22] NIST, *Digital Library of Mathematical Functions*, Section 19.8, arithmetic–geometric mean and elliptic integrals. The integral substitution, convergence and rational-certificate proof are reproduced independently here.

[23] SciPy, documentation for `scipy.sparse.linalg.splu`. The recorded comparison used installed SciPy 1.17.0 and its SuperLU backend; software execution is evidence about that environment, not a universal timing theorem.

[24] Boost, *Multiprecision*, documentation for `cpp_rational`. The C++ certificate checkers use arbitrary-precision rational arithmetic and expose this library as an implementation dependency. They are not kernel-verified programs.